



# **Installation Guide LoadMaster HA for Azure**

**24 July 2024**

# Copyright

---

Visit the following page online to see Progress Software Corporation's current Product Documentation Copyright Notice/Trademark Legend: [Product Documentation Copyright Notice & Trademarks](#) | [Progress](#)

# Table of Contents

<b>Chapter 1: Introduction.</b>	<b>5</b>
Known Issues/Limitations.	8
 <b>Chapter 2: Prerequisites.</b>	 <b>9</b>
 <b>Chapter 3: Manually Configure LoadMaster HA in Azure.</b>	 <b>12</b>
Licensing Options.	12
Create the First Virtual LoadMaster in Azure.	13
Create the Second Virtual LoadMaster in Azure.	18
Enable a 10 Gb Interface (Optional).	19
 <b>Chapter 4: Create the Internal Load Balancer (ILB).</b>	 <b>23</b>
Create a Backend Pool.	27
Create Load Balancing Rules.	28
Create Inbound NAT Rules.	31
Create Outbound Rules.	33
Manage Rules to Allow Traffic on Azure Load Balancer.	35
 <b>Chapter 5: Network Security Groups.</b>	 <b>36</b>
 <b>Chapter 6: Configure the LoadMasters.</b>	 <b>37</b>

**Chapter 7: Configure GEO Clusters with HA. . . . . 40**  
    Configure GEO Clusters with HA IP Addresses. . . . . 40

**Chapter 8: LoadMaster Firmware Upgrades/Downgrades. . . . . 45**

**Chapter 9: Best Practices for Backups. . . . . 47**

**Chapter 10: Change the 'bal' User Password using the Serial Console. . . . . 48**  
    . . . . .

**Chapter 11: Troubleshooting. . . . . 50**  
    Check which LoadMaster is Active. . . . . 50  
    First/Second Unconnected. . . . . 51  
    Connection to Default Gateway Failed. . . . . 52  
    Virtual Machine Inaccessible. . . . . 52  
    Run a TCP Dump. . . . . 53  
    Sync Problems. . . . . 53  
    Misconfigured ILB. . . . . 54  
    Problems Reaching a Virtual Service. . . . . 54

**Chapter 12: References. . . . . 55**

---

# Introduction

---

## Introduction

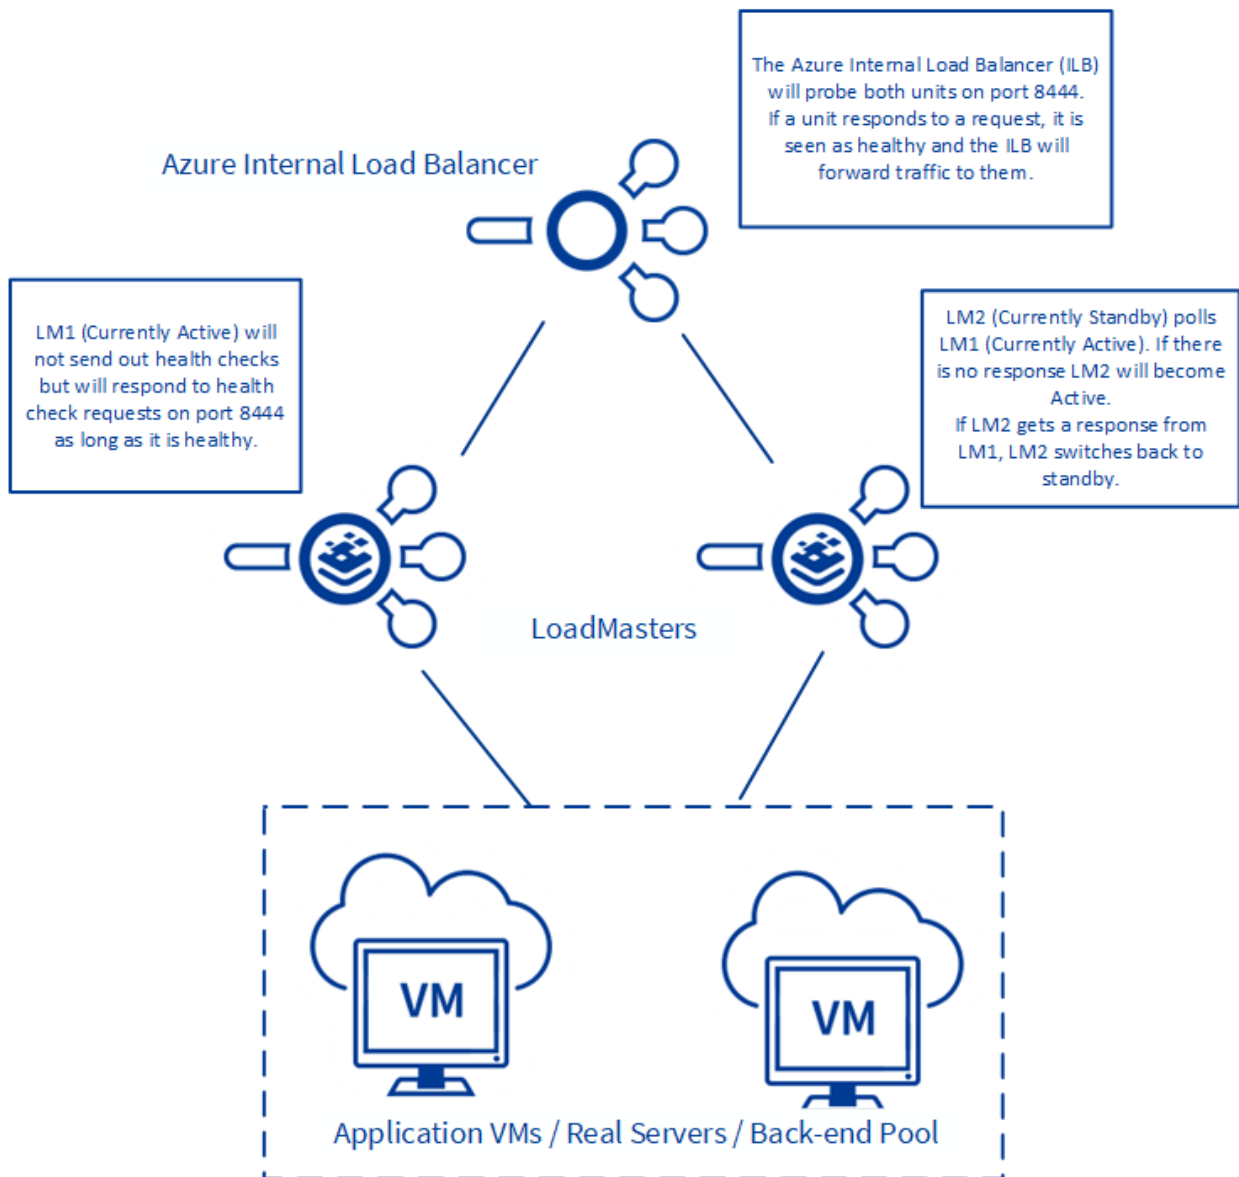
When deploying an application using the Microsoft Azure Infrastructure as a Service (IaaS) offering, you usually need to provide load balancing and other application delivery functions such as content switching, SSL Termination and IPS. Some of this functionality may also be necessary when deploying applications in Microsoft Azure Platform as a Service (PaaS). The LoadMaster for Azure enables you to address your needs of application delivery and High Availability (HA).

Deploying a single LoadMaster for Azure does not provide you with the high availability you need for your applications. When deploying a pair of LoadMasters in Azure, you can achieve high availability for your application. This document provides the details for a HA LoadMaster solution.

When using LoadMaster in High Availability on Azure, HA operates in much the same way as it does on non-cloud platforms, but with some key differences, which are listed below:

- LoadMaster HA for Azure involves two LoadMasters that synchronize settings bi-directionally. Changes made to the active unit are replicated to the standby unit and changes made to the standby unit are replicated to the active unit.
- The replication (synchronization) of settings (from active to standby) is not instant in all cases and may take a few moments to complete.
- When synchronizing the GEO settings from active to standby, any Fully Qualified Domain Name (FQDN) or cluster IP addresses that match the active unit IP address are replaced with the standby unit IP address. Likewise, when synchronizing from standby to active, the standby unit IP address is replaced with the active unit IP address.
- All user-defined settings are synchronized, with the exception of the following:
  - Default gateway (both IPv4 and IPv6)
  - IP addresses and netmasks

- Hostname
- Name server
- Domain
- Admin default gateway
- Administrative certificate settings (.cert, .pem and .setadmin files)
- Network interface settings: Link Status (Speed and Duplex), MTU and additional addresses
- Virtual LAN (VLAN) configuration
- Virtual Extensible LAN (VXLAN) configuration
- Additional routes
- The cloud HA LoadMaster does not have a "force update" option.
- By default, the active unit is always set as active and the standby unit can be standby or active if the active unit fails. The active unit is active and never becomes the standby, even if it fails. Similarly the standby unit never becomes the active unit. When the active unit comes back up it is set as active and connections are automatically directed to the active again. Either the active or standby unit can be active or standby.
- The **HA Check Port** must be set to the same port on both the active and standby units for HA to work correctly.
- Depending on the design of the Network Security Groups, you must ensure the necessary ports are open inbound to allow for the traffic.



A complete description of non-cloud LoadMaster HA can be found in the [High Availability \(HA\), Feature Description](#) document.

For details on all of the supported configurations regarding multiple Public IP addresses along with multiple interfaces, refer to the following Knowledge Base article: [Supported Network Configurations for LoadMaster in Azure](#).

#### Related Links

- [Known Issues/Limitations](#)

# Known Issues/Limitations

## Known Issues/Limitations

There are some known issues/limitations to be aware of:

- Transparency is not possible in HA setups in Azure environments. For more information and requirements, refer to the **Transparency Feature Description** document on the [Documentation page](#).
- Do not downgrade from firmware version 7.2.36 or higher to a version below 7.2.36. If you do this, the LoadMaster becomes inaccessible and you cannot recover it.
- The Virtual Service IP address must be the same IP address as the network interface.
- Alternate default gateway support is not permitted in a cloud environment.
- When setting up the LoadMaster in Azure, ensure that the **Enable OS guest diagnostics** check box is disabled because there is no support available for collection of diagnostics from the LoadMaster in Azure. If this option is enabled, the LoadMaster will not boot correctly.

The table below summarizes some of the limitations in Azure:

	Single unit, one-armed	Single unit, two-armed	HA pair one-armed	HA pair two-armed
<b>Access Control List (ACL)</b>	Requires that Server NAT is disabled if below firmware 7.2.50	Requires that Server NAT is disabled if below firmware 7.2.50	Requires that Server NAT is disabled if below firmware 7.2.50	Requires that Server NAT is disabled if below firmware 7.2.50
<b>Transparency</b>	Does not work	Works	Does not work	Does not work
<b>10Gb NICs</b>	7.2.48 and above	7.2.48 and above	7.2.48 and above	7.2.48 and above
<b>Additional addresses</b>	Works	Works	Does not work	Does not work
<b>Virtual Service not on eth0</b>	Not applicable	Works	Not applicable	Does not work



## Prerequisites

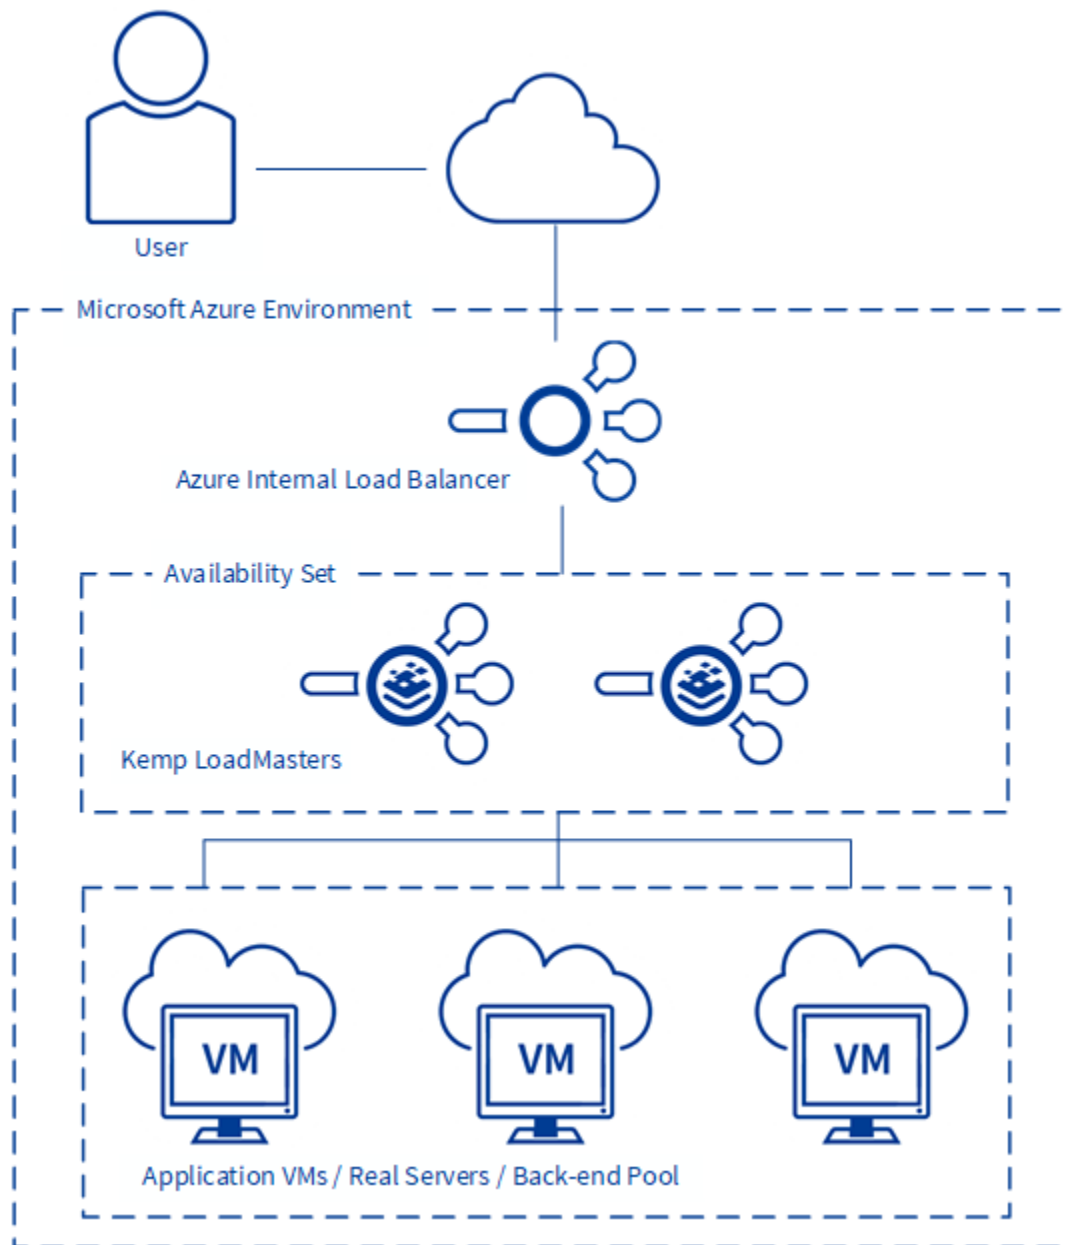
---

### Prerequisites

The following prerequisites must be met before proceeding to a high availability configuration:

- An Azure Resource Manager (ARM) (V2) Virtual Network added to Azure to place the LoadMaster VMs
- Application VMs deployed in Azure in the Virtual Network
- An Azure Internal Load Balancer deployed to create the high availability pair
- Two LoadMaster VMs deployed on the same Virtual Network as the Application VMs
- Both LoadMasters should be configured to be part of an availability set

The following diagram provides overview of the configuration described above:



To configure high availability using the LoadMaster, the following configuration must be in place:

- Application VMs are installed and configured
- LoadMaster for Azure VMs are installed and configured
- **Important:** The **HA Check Port** must be set to the same port on both the active and standby units for HA to work correctly. The same port must be configured as the probe port on the Internal Load Balancer.
- The following management Load Balanced NAT Rules may be needed to access the LoadMasters:
  - TCP Port 22 for SSH access
  - TCP Port 8443 for Management Web User Interface (WUI) access

- Additional Load Balanced Rules for any traffic that is being transmitted through the LoadMaster

---

**Note:** If using Kemp 360 Central, you must configure special NAT rules.

---

Use this table to record the necessary information required to create the LoadMaster Pair in Azure:

Fields Required for creation of LoadMaster Pair
Primary LoadMaster Name
Secondary LoadMaster Name
Pricing Tier
Password for LoadMasters
Availability Service Name
Resource Group Name
Virtual Network
Internal Load Balancer Name
Internal Load Balancer Public IP Address (PIP), if required

---

**Note:** It is not possible to bond interfaces on Azure LoadMasters.

---

---

# Manually Configure LoadMaster HA in Azure

---

## Manually Configure LoadMaster HA in Azure

---

**Note:** The steps in this section were correct at the time of writing. However, the Azure interface changes regularly so please refer to Azure documentation for up-to-date steps if needed.

---

Please complete the prerequisites documented in the earlier section.

### Related Links

- [Licensing Options](#)
- [Create the First Virtual LoadMaster in Azure](#)
- [Create the Second Virtual LoadMaster in Azure](#)
- [Enable a 10 Gb Interface \(Optional\)](#)

## Licensing Options

### Licensing Options

There are four main licensing options when deploying a LoadMaster for Azure:

- Hourly consumption
- Bring Your Own License (BYOL)

- Free version
- License Agreement - Service Provided License Agreement (SPLA)/Metered

To use the BYOL option, follow the steps below:

1. Download the **BYOL – Trial and perpetual license** version of the Virtual LoadMaster (follow the steps in the section below to do this).
2. Contact a Progress Kemp representative to get a license.
3. Update the license on your LoadMaster to apply the license change (**System Configuration > System Administration > License management**).
4. We recommend rebooting the LoadMaster after updating the license.

For more information on MELA and SPLA, refer to the relevant Feature Description on the [Documentation page](#).

## Create the First Virtual LoadMaster in Azure

### Create the First Virtual LoadMaster in Azure

The steps in this document reflect the steps in the Azure Marketplace (<http://portal.azure.com>).

To deploy a new LoadMaster, follow the steps below:

### Azure services



1. From the Azure Management Portal dashboard, click **Create a resource**.

# New

Azure Marketplace [See all](#) Popular

2. Enter **Kemp** in the search bar and press Enter on your keyboard.



### LoadMaster Load Balancer ADC

Kemp Technologies

Virtual Machine

Layer 4-7 Application Delivery  
Controller (ADC) Load Balancer,  
Content Switch and Traffic Manager

Starts at  
**\$0.20/hour**

Create  

3. Select **LoadMaster Load Balancer ADC**.

Plan



Create

4. From the drop-down menu, select the desired LoadMaster **Plan** and click **Create**.

### Project details

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

Subscription \* ⓘ

Resource group \* ⓘ  [Create new](#)

- Under **Project details**, complete the following fields:
  - Select the Azure **Subscription**.
  - Select an existing or create a new **Resource group** to deploy the LoadMaster into.

Instance details

Virtual machine name \* ⓘ  ✓

Region \* ⓘ

Availability options ⓘ

Availability zone \* ⓘ

Security type ⓘ

Image \* ⓘ  [See all images](#) | [Configure VM generation](#)

VM architecture ⓘ ☐ Arm64 ☒ x64   
 ⓘ Arm64 is not supported with the selected image.

Run with Azure Spot discount ⓘ ☐

Size \* ⓘ  [See all sizes](#)

- Under **Instance details**, complete the following fields:
  - Enter a **Virtual machine name** for the LoadMaster.
  - Select an Azure **Region**.
  - Select the relevant option from the **Availability options** drop-down list.

**Note:** The LoadMaster is compatible with Availability sets and Availability zones and it is preferred to use managed disks for production environments. Some settings are not available in all regions. Check with Microsoft if needed.

- Configure the availability settings:
  - If you are using an availability set, select an existing or new **Availability set** for the HA pair.
  - If you are using an availability zone, select an **Availability zone** to deploy the first LoadMaster into.
- Select the relevant security type for the virtual machine from the **Security type** drop-down list.
- Confirm the desired LoadMaster type is selected in the **Image** drop-down list.
- Select **x64** or **Arm64** based **VM architecture** for the virtual machine to run your applications.

8. Select the **Run with Azure Spot discount** checkbox, if required.
9. Select the desired **Size** for the virtual machine.

### Administrator account

Authentication type ⓘ

☐ SSH public key ☒ Password

Username \* ⓘ

AzureUser ✓

Password \* ⓘ

..... ✓

Confirm password \* ⓘ

..... ✓

---

**Note:** If you want to enable 10 Gb throughput for a LoadMaster virtual machine (VM) in Azure, you must select an Azure VM instance type that supports the 10 Gb Mellanox driver. For more information, refer to the [Enable a 10 Gb Interface \(Optional\)](#) section.

---

7. Under **Administrator account**, complete the following fields:
  1. Select the **Authentication type** (SSH public key or Password).

---

**Note:** We recommend using a password, but either way will work fine.

---

2. Enter a **Username**.

---

**Note:** This username is not used by the LoadMaster for Azure. The default username to access the LoadMaster is **bal**.

---

3. Enter a **Password** for the **bal** account and confirm it.

---

**Note:** The password is used to access the LoadMaster WUI.

---

4. **SSH public key source:** You can either create a new key pair, use an existing key stored in Azure, or use an existing public key.

Next : Disks >

---

**Note:** It is recommended to store SSH keys in a secure location.

---

8. Click **Next: Disks**.
9. Leave the default options for **Disk options** and **Data disks**.



Next : Networking >



10. Click **Next: Networking**.






**Network interface**


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


Virtual network \*  (new) ha\_vnet   
[Create new](#)


Subnet \*  (new) subnet (10.0.0.0/24)   
[Create new](#)


Public IP  None   
[Create new](#)

NIC network security group  ☐ None  
☐ Basic  
☒ Advanced


 This VM image has preconfigured NSG rules.

Configure network security group \* (new) HA-VNET-nsg   
[Create new](#)

Delete NIC when VM is deleted  ☒

Enable accelerated networking  ☐  
 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](#) 

Place this virtual machine behind an existing load balancing solution? ☐

**11. Under **Network interface**, complete the following fields:**

1. Select an existing or create a new **Virtual network**.
2. Select an existing or create a new **Subnet**.
3. (Optional) A **Public IP** is not required. Access is provided using the Azure Load Balancer outlined later in this guide.
4. Keep the default setting for **NIC network security group**.

---

**Note:** The security group should contain rules for port 8443 (management), 22 (SSH), and any other ports that are needed by the back-end. Do not block port 6973.

---

5. Select **Delete public IP and NIC when VM is deleted** to automatically delete the public IP and NIC when the associated virtual machine is deleted.
6. Select **Enable accelerated networking**, if the selected VM size supports.
7. (Optional) Select an existing load balancer or follow the steps outlined later in this document to create one.

**Next : Management >**

---

**12. Click **Next: Management**.**

13. You can optionally make any necessary updates to the **Identity** and **Auto-Shutdown** sections or leave them as the default settings.

---

**Note:** Ensure that the **Enable OS guest diagnostics** check box is disabled because there is no support available for collection of diagnostics from the LoadMaster in Azure. If this option is enabled, the LoadMaster will not boot correctly.

---

14. Click **Next: Monitoring**.
15. You can optionally make any necessary updates to the monitoring settings.

Next : Advanced >

16. Click **Next: Advanced**.
17. You can optionally make any necessary updates to the **Extensions** and **Custom data** sections or leave them as their defaults.

Next : Tags >

18. Click **Next: Tags**.
19. You can optionally make any necessary changes to the **Tags** section or leave the defaults.

Next : Review + create >

20. Click **Next: Review + create**.
21. You can optionally click **Download a template for automation** to download an ARM template.



22. Click **Create**.

---

**Note:** If you chose to create a new SSH key pair, you are now prompted to store the private key for the public key you created. Azure does not store the private key. After the SSH key is created, you will not be able to download the private key.

---

## Create the Second Virtual LoadMaster in Azure

### Create the Second Virtual LoadMaster in Azure

The process of setting up the second LoadMaster for Azure is similar to the first with a few exceptions, which are listed below:

- You must select the same **Resource Group** that was used during the first LoadMaster deployment.

- You must select the same **Virtual Network** that was used during the first LoadMaster deployment.
- You must select the same **Availability Set** that was created during the first LoadMaster deployment.
- If using Availability Zones, you should select a different zone to provide the necessary redundancy in the event of a data-center outage.

## Enable a 10 Gb Interface (Optional)

### Enable a 10 Gb Interface (Optional)

Follow one of the two procedures below depending on whether you are adding a single network interface or multiple network interfaces to the LoadMaster.

To enable 10 Gb throughput for a LoadMaster virtual machine (VM) in Azure, you must select an Azure VM instance type that supports the 10 Gb Mellanox driver. Accelerated Networking is supported on most general purpose and compute-optimized instance sizes with two or more vCPUs. These supported series are: D/DSv2 and F/Fs. On instances that support hyperthreading, Accelerated Networking is supported on VM instances with four or more vCPUs. Supported series are: D/Dsv3, E/Esv3, Fsv2, Lsv2, Ms/Mms and Ms/Mmsv2. Refer to the [Sizes for Linux virtual machines in Azure](#) page for further details.

### Add a Single Interface to the LoadMaster

To enable 10 Gb interfaces on the LoadMaster, perform the following steps:

1. Deploy the LoadMaster.

---

**Note:** When you instantiate a 10 Gb interface, it appears as two interfaces in the LoadMaster Web User Interface (WUI). The two interfaces are related and have the same MAC address. Only the first interface has an IP address. If you want to modify the interface, you must do this on the interface that has the IP address listed.

---



---

**Note:** For the purposes of this document, the Standard DSv2 machine size is used.

---

2. License and log into the LoadMaster.
3. Navigate to **System Configuration > System Administration > System Reboot**.
4. Click **Shutdown**.
5. In the Azure Portal, navigate to the LoadMaster Virtual Machine and click **Stop**.
6. In the Azure Portal, navigate to the Network Interface associated with the LoadMaster Virtual Machine.
7. Click **Edit accelerated networking**.
8. Click **Enabled** and select the validation box.
9. Click **Save**.
10. In the Azure Portal, navigate back to the LoadMaster Virtual Machine and click **Start**.
11. Log into the LoadMaster and verify that the Mellanox driver has instantiated correctly by performing the following steps:
  1. If the LoadMaster was deployed with a single interface, two interfaces are displayed under **System Configuration > Interfaces** on the LoadMaster WUI. If only one interface is displayed this means that the Mellanox driver has not instantiated.

```

eth0      Link encap:Ethernet  HWaddr 00:0d:3a:8d:4b:fe
          inet addr:192.168.1.4  Bcast:192.168.1.255  Mask:255.255.255.0
          inet6 addr: fe80::20d:3aff:fe8d:4bfe/64 Scope:Link
          UP BROADCAST NOTRAILERS RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:1143 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1705 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:538178 (538.1 KB)  TX bytes:1254819 (1.2 MB)

eth1      Link encap:Ethernet  HWaddr 00:0d:3a:8d:4b:fe
          UP BROADCAST NOTRAILERS RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
          RX packets:941 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1713 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:459351 (459.3 KB)  TX bytes:1262421 (1.2 MB)

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:1500  Metric:1
          RX packets:1520 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1520 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:1100982 (1.1 MB)  TX bytes:1100982 (1.1 MB)

```

2. You can also verify that two interfaces are active by checking performing an Ifconfig. To perform an Ifconfig, navigate to **System Configuration > Troubleshooting**. On the **Troubleshooting** screen, click **Ifconfig**. This displays two interfaces with the same hardware address.

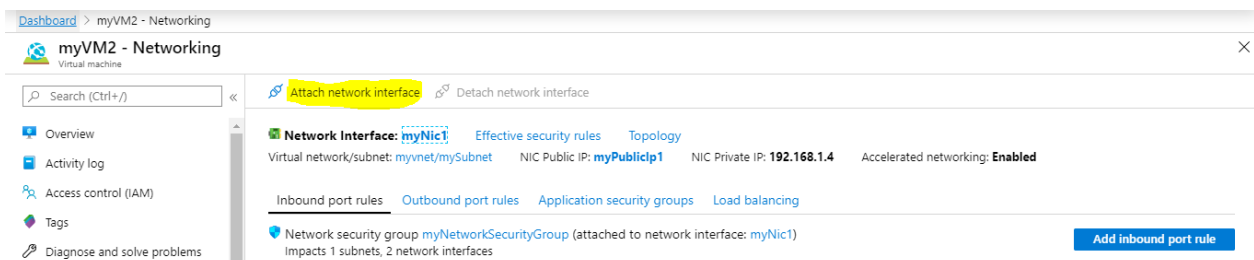
## Add Multiple Interfaces to the LoadMaster

The Azure WUI does not allow interfaces with accelerated networking to be added. You must add the interface by using the Azure command line interface (CLI) or by using PowerShell.


**Note:** You must run the command with the LoadMaster in a powered off state.

1. Create the interface using the Azure CLI similarly to the example below:  

```
PS C:\Users\test> az network nic create --resource-group testdoc --name myNic2 --vnet-name myVnet --subnet subnet2 --accelerated-networking true --public-ip-address myPublicIp2 --network-security-group myNetworkSecurityGroup --location eastus
```
2. When the interface is created, you can add this interface to the LoadMaster when it is in a powered off state. Navigate to the **Networking** tab of the LoadMaster on the Azure WUI.



### 3. Click **Attach network interface**.



4. When the attachment is complete, both interfaces appear on the Azure WUI.

5. Restart the LoadMaster.

6. Verify that the interfaces are displayed under **System Configuration > Interfaces** on the LoadMaster WUI. The LoadMaster WUI should now display four interfaces.

You can also verify that four interfaces are active by checking performing an Ifconfig. To perform an Ifconfig, navigate to **System Configuration > Troubleshooting**. On the **Troubleshooting** screen, click **Ifconfig**.

```

eth0    Link encap:Ethernet HWaddr 00:0d:3a:8d:4b:fe
        inet addr:192.168.1.4 Bcast:192.168.1.255 Mask:255.255.255.0
        inet6 addr: fe80::20d:3aff:fe8d:4bfe/64 Scope:Link
        UP BROADCAST NOTRAILERS RUNNING MULTICAST MTU:1500 Metric:1
        RX packets:1102 errors:0 dropped:0 overruns:0 frame:0
        TX packets:1641 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:407988 (407.9 KB) TX bytes:1246784 (1.2 MB)

eth1    Link encap:Ethernet HWaddr 00:0d:3a:8c:3e:81
        inet addr:192.168.2.4 Bcast:192.168.2.255 Mask:255.255.255.0
        inet6 addr: fe80::20d:3aff:fe8c:3e81/64 Scope:Link
        UP BROADCAST NOTRAILERS RUNNING MULTICAST MTU:1500 Metric:1
        RX packets:68 errors:0 dropped:0 overruns:0 frame:0
        TX packets:27 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:8147 (8.1 KB) TX bytes:4702 (4.7 KB)

eth2    Link encap:Ethernet HWaddr 00:0d:3a:8d:4b:fe
        UP BROADCAST NOTRAILERS RUNNING SLAVE MULTICAST MTU:1500 Metric:1
        RX packets:864 errors:0 dropped:0 overruns:0 frame:0
        TX packets:1648 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:344516 (344.5 KB) TX bytes:1254164 (1.2 MB)

eth3    Link encap:Ethernet HWaddr 00:0d:3a:8c:3e:81
        UP BROADCAST NOTRAILERS RUNNING SLAVE MULTICAST MTU:1500 Metric:1
        RX packets:0 errors:0 dropped:0 overruns:0 frame:0
        TX packets:15 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:0 (0.0 B) TX bytes:3256 (3.2 KB)

lo      Link encap:Local Loopback
        inet addr:127.0.0.1 Mask:255.0.0.0
        inet6 addr: ::1/128 Scope:Host
        UP LOOPBACK RUNNING MTU:1500 Metric:1
        RX packets:1575 errors:0 dropped:0 overruns:0 frame:0
        TX packets:1575 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:1084369 (1.0 MB) TX bytes:1084369 (1.0 MB)

```

Unlike the single interface case where **eth0** and **eth1** are related, for multiple interfaces, **eth0** and **eth2** and **eth1** and **eth3** are related (with the same MAC address). **eth0** and **eth1** have the IP addresses, the other interfaces without the IP addresses are related by the **HWaddr** (the MAC address).

---

## Create the Internal Load Balancer (ILB)

---

### Create the Internal Load Balancer (ILB)

An Azure Internal Load Balancer must be deployed to monitor the health of the LoadMasters and direct traffic accordingly.

There are several settings that need to be configured to provide the high availability of the LoadMasters:

- Create a back-end address pool and add the LoadMasters to the pool.
- Create Inbound NAT rules to direct traffic to the appropriate LoadMaster.
- Create a Probe to monitor the health of the LoadMasters.
- Create Load Balancing Rules to allow the necessary traffic.

Refer to the sections below for further information on each of these.

The following procedure describes how to set up an Azure Load Balancer from the Microsoft Azure portal:

---

**Note:** The steps in this document reflect the steps in the Azure Marketplace (<http://portal.azure.com>).

---

To deploy a new load balancer, follow the steps below:

## Azure services



1. From the Azure Management Portal dashboard, click **Create a resource**.

## New

Azure Marketplace [See all](#)

2. Enter **Load Balancer** in the search bar and press Enter on your keyboard.





## Load Balancer

Microsoft

### Azure Service

A load balancer that distributes incoming traffic among backend virtual machine instances.

Create 



3. Click **Create** and click **Load Balancer**.

#### Project details

Subscription \*

PLM



Resource group \*

Azure-RG1



[Create new](#)

4. Under **Project details**, complete the following fields:
1. Select the Azure **Subscription**.
  2. Select the existing **Resource Group** used to deploy the LoadMasters.

**Instance details**

Name \*

Region \*

SKU \* ⓘ  
☒ Standard  
☐ Gateway  
☐ Basic

Microsoft recommends Standard SKU load balancer for production workloads. [Learn more about pricing differences between Standard and Basic SKU](#)

Type \* ⓘ  
☒ Public  
☐ Internal

Tier \*  
☒ Regional  
☐ Global

5. Under **Instance details**, complete the following fields:
1. Enter a **Name** for the load balancer.
  2. Select the Azure **Region** used to deploy the LoadMasters.
  3. Select the relevant **SKU** for the load balancer.

**Note:** Microsoft recommends the Standard SKU for production workloads.

4. Select the **Type** of load balancer determined by **Public** access or **Internal** only.
5. Select the **Tier** of load balancer ( **Regional** or **Global**).

**Next : Frontend IP configuration >**

6. Click **Next: Frontend IP configuration**.

Basics Frontend IP configuration Backend pools Inbound rules Outbound rules Tags Review + create

A frontend IP configuration is an IP address used for inbound and/or outbound communication as defined within load balancing, inbound NAT, and outbound rules.

+ Add a frontend IP configuration

Name ↑↓

IP address ↑↓

Add a frontend IP to get started

7. Under **Frontend IP configuration**, click **Add a frontend IP configuration**.

## Add frontend IP configuration ×

Name \*  
VLM-FrontEnd ✓

IP version  
☒ IPv4 ☐ IPv6

IP type  
☒ IP address ☐ IP prefix

Public IP address \*  
 (New) LB-PIP  
[Create new](#)

Gateway Load balancer ⓘ  
 None

1. Enter a **Name** for the frontend IP address.
2. Select either **IPv4** or **IPv6** as the **IP version**.
3. Select **IP address** as the **IP type**.
4. Create a new **Public IP address** or select an existing one from the drop-down list.

---

**Note:** Select either **Dynamic** or provide a **Static IP Assignment** if creating a new **Public IP address**.

---

5. Click **OK**.
6. Click **Add**.
8. Click **Next: Backend pools**.
9. Refer to the next section for details on setting up a backend pool.

### Related Links

- [Create a Backend Pool](#)
- [Create Load Balancing Rules](#)
- [Create Inbound NAT Rules](#)
- [Create Outbound Rules](#)
- [Manage Rules to Allow Traffic on Azure Load Balancer](#)

## Create a Backend Pool

The backend pool is a collection of virtual machines (LoadMasters) which is load balanced to provide High Availability. Continuing on from the previous section, follow these steps to create a backend pool:

A backend pool is a collection of resources to which your load balancer can send traffic. A backend pool can contain virtual machines, virtual machine scale sets, and containers.

+ Add a backend pool

Name	Virtual network	Resource Name	Network interface	IP address	Availability zone
Add a backend pool to get started					

1. Click **Add a backend pool**.

Add backend pool ...

Name \*

VLM-Backend

Virtual network ⓘ

Vnet (LoadMasterQA)

Backend Pool Configuration

☐ NIC

☒ IP address

**IP addresses**  
You can only add resources IP address in the Virtual Network. The configuration is associated with the IP address and will apply to any resource which has this IP address assigned.

Backend Address Name	IP address	Resource Name	
d0c8ae19-660e-49e5-a45a-4047f...	10.0.0.4	Private Network Resource	🗑
04bc6912-65f9-4416-9120-b510...	10.0.0.5	Private Network Resource	🗑

2. Under **Add backend pool**, complete the following fields:
- 1. Enter a **Name** for the back-end pool.
  - 2. Select the **Virtual network** used for the LoadMasters.
  - 3. Select **IP address** as the **Backend Pool Configuration**.
  - 4. Select the IP address of the LoadMasters for the HA pair.
  - 5. Click **Save**.
3. Click **Next: Inbound rules**.
4. Refer to the next section for details on setting up the inbound rules.

Create Load Balancing Rules

On Azure cloud, the ILB is used to create the “Shared IP address” and to probe and route traffic to the LoadMaster instances. Load Balancing Rules must be configured for any traffic that is published through the LoadMaster:

Basics Frontend IP configuration Backend pools Inbound rules Outbound rules Tags Review + create

#### Load balancing rule

A load balancing rule distributes incoming traffic that is sent to a selected IP address and port combination across a group of backend pool instances. The load balancing rule uses a health probe to determine which backend instances are eligible to receive traffic.

+ Add a load balancing rule

Name ↑↓ Frontend IP configuration ↑↓ Backend pool ↑↓ Health probe ↑↓ Frontend Port ↑↓ Backend port ↑↓

Add a rule to get started

### 1. Click **Add a load balancing rule**.

## Add load balancing rule

Azure-ILB

A load balancing rule distributes incoming traffic that is sent to a selected IP address and port combination across a group of backend pool instances. Only backend instances that the health probe considers healthy receive new traffic.

Name *	Port443
IP Version *	<input checked="" type="radio"/> IPv4 <input type="radio"/> IPv6
Frontend IP address * ⓘ	VLMFrontEnd (To be created) ▼
Backend pool * ⓘ	VLM-Backend ▼
Protocol	<input checked="" type="radio"/> TCP <input type="radio"/> UDP
Port *	443
Backend port * ⓘ	443

### 2. Under **Add load balancing rule**, complete the following fields:


1. Enter a **Name** for the load balancing rule.
2. Select the **IP Version**.
3. Select the **Frontend IP address** created earlier.
4. Select the **Backend pool** created earlier.
5. Select the **Protocol** to use.
6. Enter the **Port** that clients will connect to.
7. Enter the **Backend port** to send traffic to.





8. For **Health Probe**, click **Create new**.

## Add load balancing rule

Azure-ILB

A load balancing rule distributes incoming traffic that is sent to a selected IP address and port combination across a group of backend pool instances. Only backend instances that the health probe considers healthy receive new traffic.

 Health probes are used to check the status of a backend pool instance. If the health probe fails to get a response from a backend instance then no new connections will be sent to that backend instance until the health probe succeeds again.

Name *	VLM-Probe
Protocol *	HTTP
Port * 	8444
Path * 	/
Interval (seconds) * 	5
Used by * 	Not used
<div>Save</div> <div>Cancel</div>	

**Note:** These two port numbers are typically the same but, if required, the traffic can be directed to a different backend port number than the one clients will connect to.

- Under **Add health probe**, complete the following fields:
  - Enter a **Name** for the health probe.
  - Select **HTTP** as the **Protocol**.
  - Enter **8444** as the **Port**.
  - Keep the **Path** as the default setting of **/**.
  - Set the **Interval** to **5**.
  - Click **OK**.
- Keep the rest of the configuration as the default settings and click **Add**.
- Add additional rules for other ports that are required in the environment.

**Note:** You can add additional load balancing rules depending on application requirements. TCP/8443 can be added to provide access to the active LoadMaster. TCP/8444 can be added to provide information on the

state of the HA pair. TCP/22 must be open if GEO partners are being used in the environment. Refer to the [Configure GEO Clusters with HA](#) section for more information on GEO clusters.

## Create Inbound NAT Rules

On Azure cloud, the ILB is used to create the "Shared IP address" and to probe and route traffic to the LoadMaster instances. To allow 'public' access to the WUI of each LoadMaster, we recommend creating ILB NAT rules:

- <SIP>:8441 maps to Node-1 port 8443
- <SIP>:8442 maps to Node-2 port 8443

**CAUTION:** If using the HA pair awareness functionality in Kemp 360 Central, you must be able to probe the shared IP address on the WUI port (for example, <SIP>:8443). This requires an ILB inbound rule for 8443 to allow access to the back-end pool. However, the ILB does not allow a port used in a NAT rule to also be used in an inbound rule. Therefore, if you want to use the HA pair awareness in Kemp 360 Central, you must create a different set of NAT rules.

Inbound NAT rules provide a translation for management access into each of the LoadMasters in the back-end pool. Each LoadMaster does not require a Public IP Address (PIP). A unique port must be configured in an Inbound NAT rule for each LoadMaster. The example rules are the following:

Target	Port	Target Port
LoadMaster1 - WUI	8441	8443
LoadMaster1 – SSH	221	22
LoadMaster2 – WUI	8442	8443
LoadMaster2 – SSH	222	22

The LoadMaster uses port 22 and 8443 by default. The remaining port numbers listed above are recommended, but you can use other port numbers if needed.

Inbound NAT rule

An inbound NAT rule forwards incoming traffic sent to a selected IP address and port combination to a specific virtual machine.

+ Add an inbound nat rule

Name ↑↓	Frontend IP configuration ↑↓	Service ↑↓	Target ↑↓	Frontend Port ↑↓
Add a rule to get started				

To create the inbound NAT rules, continuing from the previous section follow the steps below:

- 1. Click **Add an inbound NAT rule**.

### Add inbound NAT rule

New

**i** An inbound NAT rule forwards incoming traffic sent to a selected IP address and port combination to a specific virtual machine.

Name \*

Type ⓘ

Target virtual machine

Frontend IP address \* ⓘ

Frontend Port \*

Service Tag \*

Backend port \*

Protocol

Enable TCP Reset ⓘ

Idle timeout (minutes) \* ⓘ

Enable Floating IP ⓘ

VLM1-WUI-Rule ✓

☒ Azure virtual machine

☐ Backend pool

None ▾

New ▾

8441 ✓

Custom ▾

8443 ✓

☒ TCP

☐ UDP

☐

4

☐

- 2. Under **Add inbound NAT rule**, complete the following fields:
  - 1. Enter a **Name** for the rule.
  - 2. Select **Azure virtual machine** as the **Type**.
  - 3. Select the **Target virtual machine** (in this example, the first LoadMaster is selected).
  - 4. Select the **Network IP configuration**.
  - 5. Select the **Frontend IP address** created earlier.



6. Enter **8441** as the **Frontend Port**.
7. Select **Custom** as the **Service Tag**.
8. Enter **8443** as the **Backend Port**.
9. Select **TCP** as the **Protocol**.
10. Keep the remaining configuration as the default values and click **Add**.

Name ↑↓	Frontend IP configuration ↑↓	Service ↑↓	Target ↑↓	Frontend Port ↑↓
VLM1-WUI-Rule	VLM-FrontEnd	Custom	test-vm	8441
VLM2-WUI-Rule	VLM-FrontEnd	Custom	test-vm2	8442
VLM1-SSH-Rule	VLM-FrontEnd	Custom	test-vm	221
VLM2-SSH-Rule	VLM-FrontEnd	Custom	test-vm2	222

3. Create four inbound NAT rules based on the table provided earlier in this section.
4. Click **Next: Outbound Rule**.
5. Refer to the next section for details on setting up the outbound rules.

## Create Outbound Rules

Because the need for a Public IP address on the LoadMaster is not a requirement, an outbound rule is required to send traffic out of the Azure Load Balancer. This is also necessary if using GEO clusters (outlined later in this document) because the traffic must be seen as the public IP address of the Azure Load Balancer.

To create outbound rules, follow these steps:

Create load balancer ...

Basics Frontend IP configuration Backend pools Inbound rules Outbound rules Tags Review + create

### Outbound rules

An outbound rule allocates source network access translation (SNAT) ports from Frontend IP addresses to a backend pool for outbound connections to the internet.

+ Add an outbound rule

Name ↑↓	Frontend IP configuration ↑↓	Backend pool ↑↓	Protocols ↑↓	Ports Per Instance ↑↓
Add a rule to get started				

1. Click **Add an outbound rule**.

# Add outbound rule ✕

Name \*

Port443 ✓

IP Version \*

☒ IPv4

☐ IPv6

Frontend IP address \* ⓘ

1 selected ▼

Protocol

☐ All

☒ TCP

☐ UDP

Idle timeout (minutes) ⓘ

4

Max: 100

TCP Reset ⓘ

☒ Enabled

☐ Disabled

Backend pool \* ⓘ

VLM-Backend (0 instances) ▼

## Port allocation

Azure automatically assigns the number of outbound ports to use for source network address translation (SNAT) based on the number of frontend IP addresses and backend pool instances. [Learn more about outbound connectivity](#) ↗

Port allocation ⓘ

Manually choose number of outbound ports ▼

2. Under **Add load balancing rule**, complete the following fields:
  1. Enter a **Name** for the outbound rule.
  2. Select the **IP Version**.
  3. Select the **Frontend IP Address** created earlier.
  4. Select the **Backend pool** created earlier.
  5. Select **Manually choose number of outbound ports** for the **Port allocation**.
  6. Select **Maximum number of backend instances** for **Outbound ports**.
  7. Enter **2** for the **Maximum number of backend instances**.
  8. Click **Add**.
3. Click **Next: Tags**.
4. You can optionally make any necessary changes to the **Tags** section or keep the defaults.

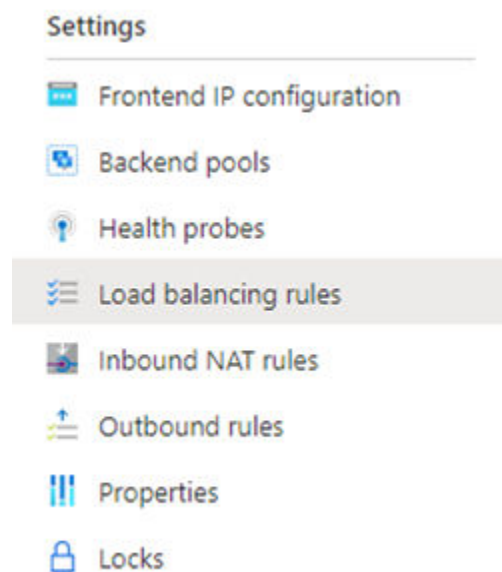
Next : Review + create >

5. Click **Next: Review + create**.
6. Click **Create**.

## Manage Rules to Allow Traffic on Azure Load Balancer

### Manage Rules to Allow Traffic on Azure Load Balancer

Once created additional rules can be added under settings of the Azure Load Balancer.



---

# Network Security Groups

---

## Network Security Groups

Network Security Groups are used in Azure to control what traffic is allowed or denied access to Virtual Machines. Depending on your configuration, you are required to update one or more Network Security Groups to allow published traffic to access the LoadMasters and backend Real Servers.

---

**Note:** The security group must contain a rule for 8443. This is the WUI port. If the LoadMaster is public-facing, other best practice, recommended ports if you are using the NAT rules mentioned in the earlier section that should be in the security group, are; 8441, 8442, 8444, 22, 221, 222, the Virtual Service ports (such as 443) and any other ports that are needed by the backend.

---

---

**Note:** For security purposes, these management ports can also be restricted to select source IP address to limit access.

---

---

**Note:** Do not block port 6973.

---

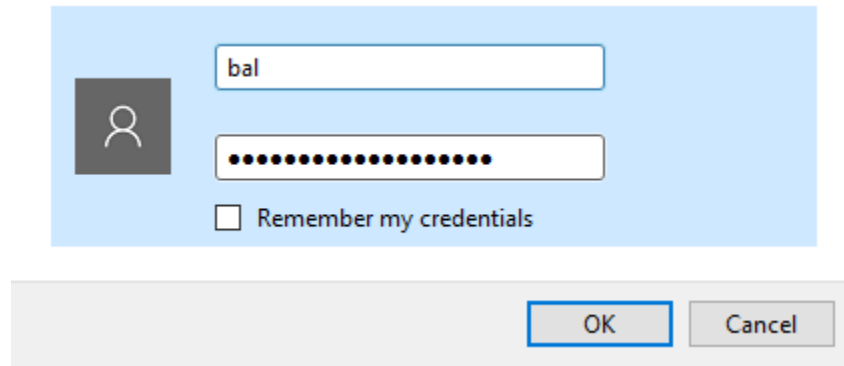
---

## Configure the LoadMasters

---

### Configure the LoadMasters

To configure LoadMaster for HA, follow the steps outlined in the sections below:



1. If the LoadMaster does not have a public address itself and you are going through the Internal Load Balancer (ILB), you can access the WUI of the LoadMaster which is the active unit:
  1. Access the WUI of active LoadMaster by going to **`https://<DNSNameURL>:8441`**.
  2. Access the WUI of the standby LoadMaster by going to **`https://<DNSNameURL>:8442`**.
  3. The default username is **bal** and the password is the password entered during the creation of the LoadMaster.
2. In the main menu, go to **System Configuration > Azure HA Parameters**.

---

Azure HA Mode	<input type="text" value="First HA Mode"/>	
Switch to Preferred Server	<input type="text" value="Prefer First"/>	
Partner Name/IP	<input type="text" value="10.0.0.4"/>	<input type="button" value="Set Partner Name/IP"/>
Health Check Port	<input type="text" value="8444"/>	<input type="button" value="Set Health Check Port"/>
Health Check on All Interfaces	<input type="checkbox"/>	

---

3. Select **First HA Mode** in the **Azure HA Mode** drop-down list.
4. Select the desired option in the **Switch to Preferred Server** drop-down list:
  - **No Preferred Host:** Each unit takes over when the other unit fails. No switchover is performed when the partner is restarted.
  - **Prefer First:** The HA1 (active) unit always takes over. This is the default option.
5. Enter the internal address of the standby LoadMaster unit in the **Partner Name/IP** text box and click **Set Partner Name/IP**.
6. Enter **8444** as the **Health Check Port** and click **Set Check Port**.

---

**Note:** The **Health Check Port** must be set to **8444** on both the active and standby units for HA to function correctly.

---

7. If using a multi-arm configuration, select the **Health Check on All Interfaces** check box.

---

**Note:** If this option is disabled, the health check listens on the primary eth0 address.

---

8. Then, access the WUI of the standby unit. Complete the following steps in the standby unit, but select **Second HA Mode** as the **Azure HA Mode** instead: In the main menu, go to **System Configuration > HA and Clustering**. to Enter the Partner Name/IP address of the slave LoadMaster unit and click **Set Partner Name/IP**.

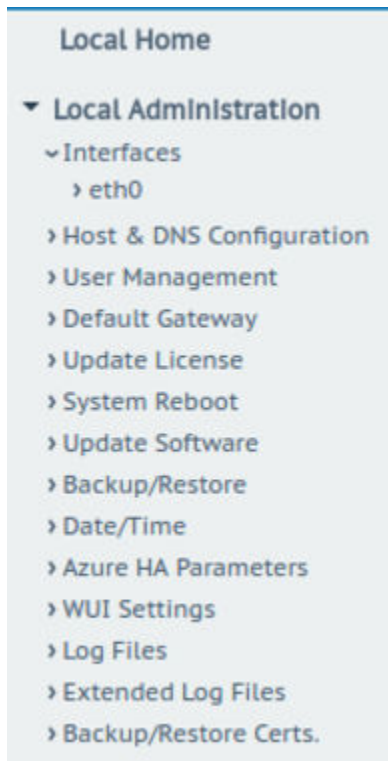
---

**Note:** HA will not work if both units have the same value selected for the **Azure HA Mode**.

---

9. After configuring both LoadMasters, reboot both units (**System Configuration > System Administration > System Reboot > Reboot**).

When HA is enabled on both devices, changes made to the Virtual Services in the active unit is replicated to the standby.



If a unit is in standby mode, WUI access is restricted to **Local Administration** only. Full WUI access is available if the unit is in an active or unchecked state.

**FIRST (ACTIVE) 06:04:06 AM**

You can tell, at a glance, which unit is active, and which is the standby, by checking the mode in the top bar of the LoadMaster.

The current status of each LoadMaster, when HA is enabled, is shown as follows:

Status	Description
<b>FIRST (ACTIVE) 06:04:06 AM</b>	This is the first LoadMaster and it is currently active.
<b>SECOND (ACTIVE) 07:46:03 AM</b>	This is the second LoadMaster and it is currently active.
<b>SECOND (STAND-BY) 05:28:41 AM</b>	This is the second unit and it is currently the standby unit.

---

## Configure GEO Clusters with HA

---

### Configure GEO Clusters with HA

If Global Server Load Balancing (GSLB)/GEO is being used in the environment, there are unique configuration requirements for setting up clusters. This section outlines the configuration and the steps can be modified based on the environment.

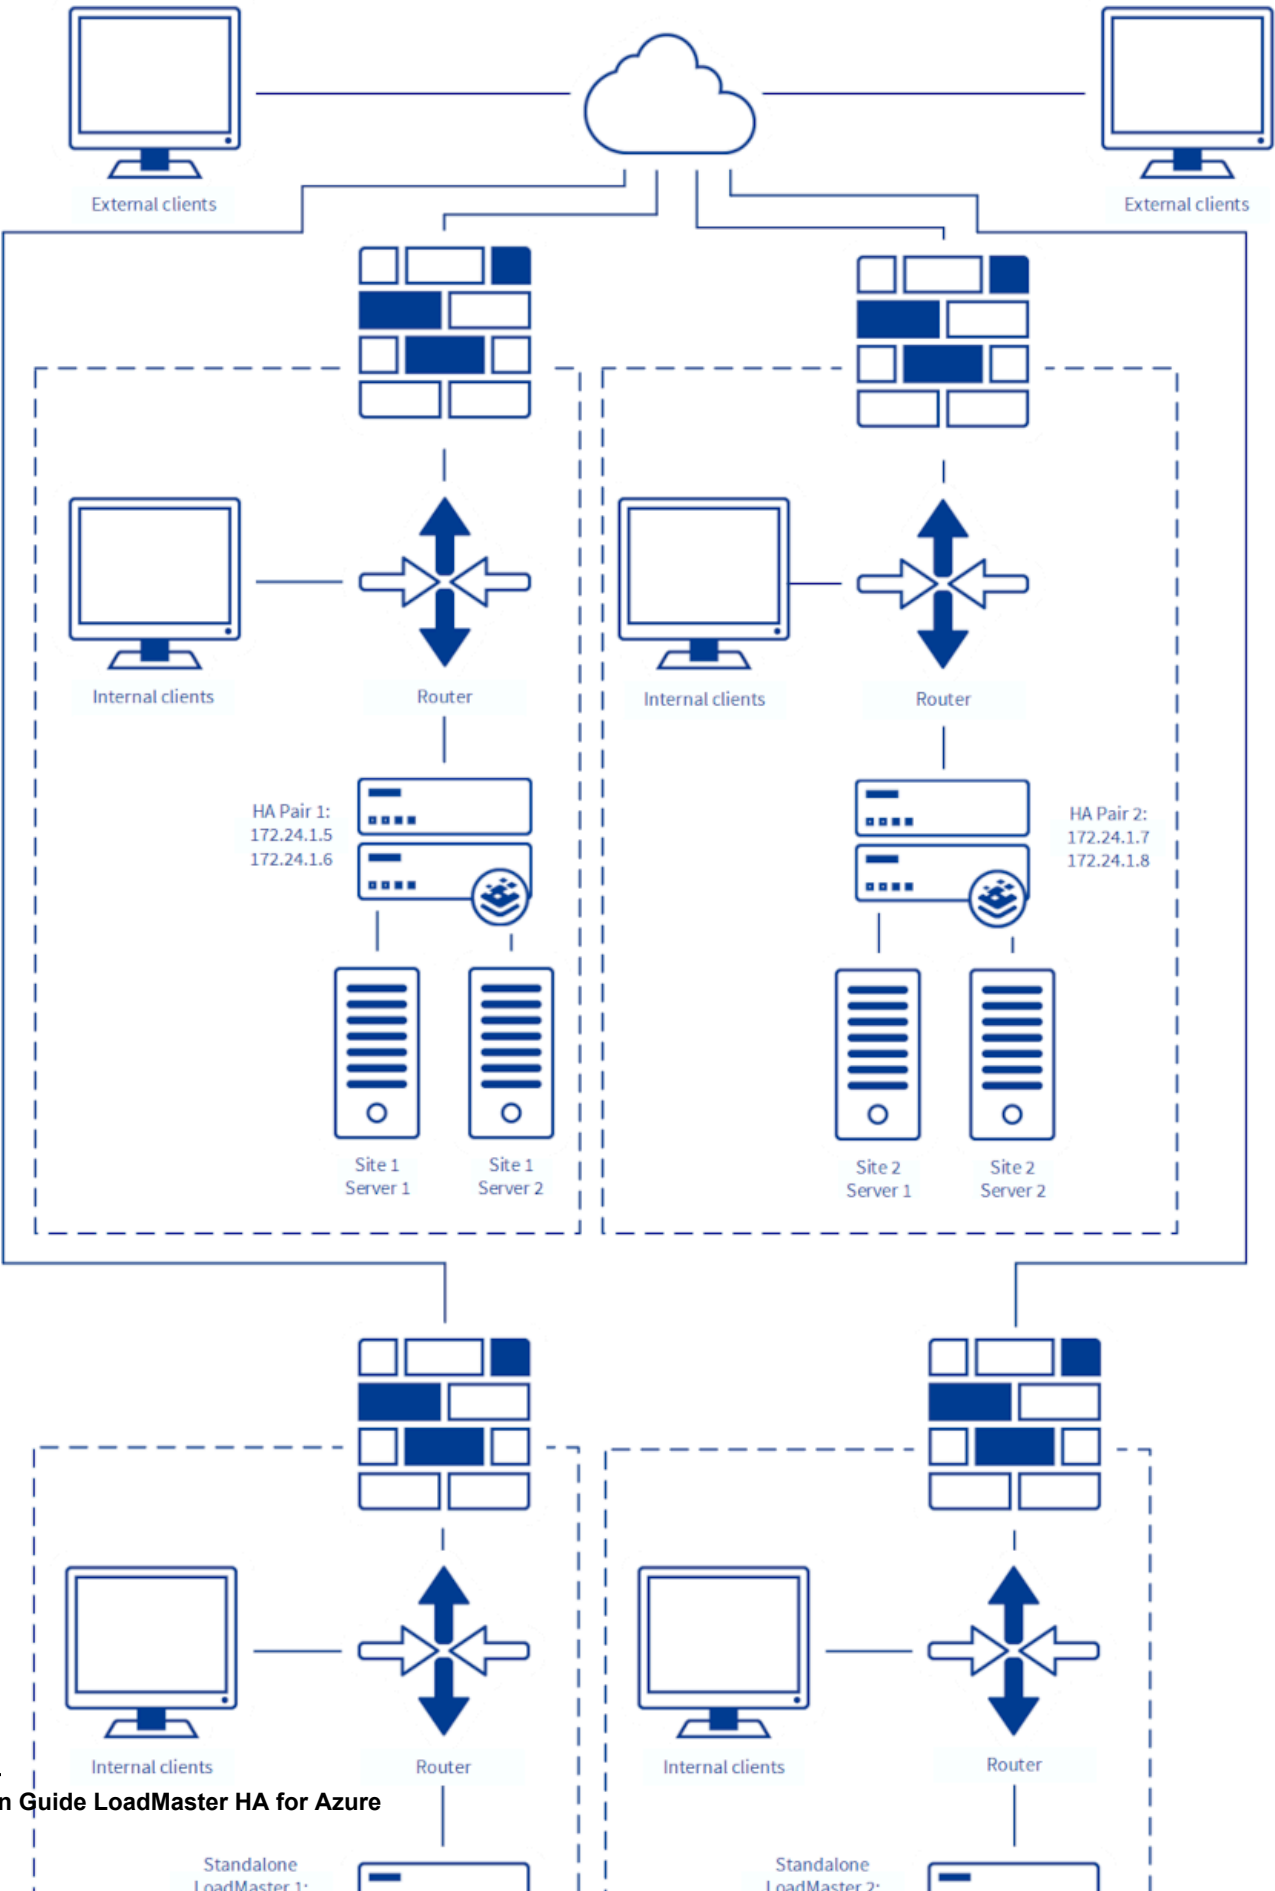
#### Related Links

- [Configure GEO Clusters with HA IP Addresses](#)

## Configure GEO Clusters with HA IP Addresses

In this example, there will be two HA Pairs of LoadMasters and two standalone LoadMasters included in the GSLB environment. The steps will be the same if more or less units are in the environment.





To configure clusters of HA IP addresses between two HA pairs and two standalone systems, follow the steps below.

On HA pair 1, complete these steps:

1. Log into the Active LoadMaster in the HA pair.
2. In the main menu, go to **Certificates & Security > Remote Access**.

The screenshot shows the 'GEO Settings' page with the following fields and buttons:

GEO Settings	
Remote GEO LoadMaster Access	<input type="text" value="52.151.122.233 52.151.1"/> <button>Set GEO LoadMaster access</button>
GEO LoadMaster Partners	<input type="text" value="52.151.122.233 52.151.1"/> <button>Set GEO LoadMaster Partners</button>
GEO LoadMaster Port	<input type="text" value="22"/> <button>Set GEO LoadMaster Port</button>
GEO Update Interface	<input type="text" value="eth0: 10.19.0.4"/>

3. Enter the IP addresses of the ILBs in front of the HA pair of LoadMasters and the IP addresses of both standalone LoadMasters in the **Remote GEO LoadMaster Access** text box and click **Set GEO LoadMaster access**.
4. Enter the IP addresses of the ILBs in front of the HA pair of LoadMasters and the IP addresses of both standalone LoadMasters in the **GEO LoadMaster Partners** text box and click **Set GEO LoadMaster Partners**.

---

**Note:** For the **GEO LoadMaster Partners** field - you do not need to specify the IP address (ILB IP Address) of the current LoadMaster you are connected to.

---

5. Reboot the Active LoadMaster.
6. Log into the new active LoadMaster.
7. In the main menu, go to **Certificates & Security > Remote Access**.
8. Re-enter the IP addresses of the ILBs in front of the HA pair of LoadMasters and the IP addresses of both standalone LoadMasters in the **Remote GEO LoadMaster Access** text box and click **Set GEO LoadMaster access**.
9. (Optional) Reboot the system again to get the original LoadMaster as active.

Repeat these steps on HA pair 2 and on the two standalone LoadMasters.

---

**Note:** Ensure the IP addresses used here are listed in the Network Security Groups for the inbound rules of TCP/22 if access is limited to specific source IP addresses.

---

Then, create the GEO clusters with the **Type Remote LM** using the IP addresses of the ILBs in front of the HA pairs and the IP addresses of both standalone LoadMasters. This only needs to be performed on one system as changes will be replicated to the other partner systems configured in the prior step.

1. In the main menu, go to **Global Balancing > Manage Clusters**.

## Add a Cluster

IP address  Name

2. In the **IP address** text box, enter the LoadMaster IP address.
3. Enter a **Name** for the cluster and click **Add Cluster**.
4. Click **Modify** on the relevant cluster.

## Modify Cluster HApair 1 Cluster

IP Address	Name	Location	Type	Checkers	Operation
52.151.122.233	<input type="text" value="HApair 1 Cluster"/> <input type="button" value="Set Name"/>	Location: 51°30'59"N 0°5'35"W <input type="button" value="Show Coordinates"/>	<input type="text" value="Remote LM"/>	Implicit	<input type="button" value="Disable"/>

5. Select **Remote LM** in the **Type** drop-down list.

Repeat these steps to add the remaining HA pair and on the two standalone LoadMasters.

Here are some example IP addresses for the scenario outlined above (involving two HA pairs and two standalone systems).

Unit	Number	IP Address
HA pair 1	LM 1	10.19.0.4
HA pair 1	LM 2	10.19.0.5
HA pair 1 ILB		52.151.122.233
HA pair 2	LM 1	10.20.0.4
HA pair 2	LM 2	10.20.0.5
HA pair 2 ILB		52.151.124.198
Standalone unit 1		20.0.54.164
Standalone unit 2		20.0.51.33

Here is an example configuration based on the scenario outlined above (involving two HA pairs and two standalone systems).

Unit	WUI Field	WUI Field Value
HA pair 1	<b>Remote GEO LoadMaster Access</b>	52.151.122.233 52.151.124.198 20.0.54.164 20.0.51.33
	<b>GEO LoadMaster Partners</b>	52.151.122.233 52.151.124.198 20.0.54.164 20.0.51.33

Unit	WUI Field	WUI Field Value
HA pair 2	<b>Remote GEO LoadMaster Access</b>	52.151.122.233 52.151.124.198 20.0.54.164 20.0.51.33
	<b>GEO LoadMaster Partners</b>	52.151.122.233 52.151.124.198 20.0.54.164 20.0.51.33
Standalone unit 1	<b>Remote GEO LoadMaster Access</b>	52.151.122.233 52.151.124.198 20.0.54.164 20.0.51.33
	<b>GEO LoadMaster Partners</b>	52.151.122.233 52.151.124.198 20.0.54.164 20.0.51.33
Standalone unit 2	<b>Remote GEO LoadMaster Access</b>	52.151.122.233 52.151.124.198 20.0.54.164 20.0.51.33
	<b>GEO LoadMaster Partners</b>	52.151.122.233 52.151.124.198 20.0.54.164 20.0.51.33

The IP addresses colored as red above correspond to the particular unit mentioned in the **Unit** column.

**Note:** For the **GEO LoadMaster Partners** field - you do not need to specify the red IP address but you can include it because it is easier for management purposes.

---

# LoadMaster Firmware Upgrades/ Downgrades

---

## LoadMaster Firmware Upgrades/Downgrades

Do not downgrade from firmware version 7.2.36 or higher to a version below 7.2.36. If you do this, the LoadMaster becomes inaccessible and you cannot recover it.

You should never leave two LoadMasters with different firmware versions paired as HA in a production environment. To avoid complications, follow the steps below in sequence and do not perform any other actions in between the steps. Please upgrade/downgrade during a maintenance window and expect service disruption because there are reboots.

The steps below are high-level, for detailed step-by-step instructions on how to upgrade the LoadMaster firmware, refer to the Updating the LoadMaster Software Feature Description on the documentation page: <https://kemptechnologies.com/loadmaster-documentation>.

## Upgrade the LoadMaster Firmware

To upgrade the LoadMaster firmware with the least disruption, follow the steps below in sequence:

1. Identify the STAND-BY unit.
2. Upgrade the LoadMaster firmware on the STAND-BY unit. Once the STAND-BY unit has rebooted, it remains in the STAND-BY state and the WUI is limited to the Local Administration options.
3. Upgrade the LoadMaster firmware on the ACTIVE unit. When the ACTIVE unit is rebooting, the STAND-BY unit becomes ACTIVE.
4. Depending on Preferred Host settings in the HA configuration, the Standby unit may failback over to the Active unit.

After these steps are completed the upgrade is finished.

## Downgrade the LoadMaster Firmware

To downgrade the LoadMaster firmware with the least disruption, follow the steps below in sequence:

1. Identify the STAND-BY unit.
2. Downgrade the LoadMaster firmware on the STAND-BY unit. Once the STANDY-BY unit has rebooted, it remains in the STAND-BY state and the WUI is limited to the Local Administration options.
3. Downgrade the LoadMaster firmware on the ACTIVE unit. When the ACTIVE unit is rebooting, the STAND-BY unit becomes ACTIVE.
4. Depending on Preferred Host settings in the HA configuration, the Standby unit may failback over to the Active unit.

After these steps are completed the downgrade is finished.

---

# Best Practices for Backups

---

## Best Practices for Backups

Hypervisor snapshots cannot be used to restore a LoadMaster to a working state. The best way to back up your LoadMaster settings is by using the native backup and restore facility in the LoadMaster WUI or API.

To back up your LoadMaster configuration, follow these steps:

1. In the main menu, go to **System Configuration > System Administration > Backup/Restore**.
2. Click **Create Backup File**.

You can create a remote host for automated backups using SCP to save backups to a remote server.

For further details on backing up and restoring the LoadMaster configuration, including certificates and cipher sets, refer to the following links:

- [Backup and Restore Technical Note](#)
- [How to Create and Restore a LoadMaster Configuration or Certificate Backup](#)

---

## Change the 'bal' User Password using the Serial Console

---

### Change the 'bal' User Password using the Serial Console

The default administrator user for the LoadMaster is called **bal**. You can change the **bal** user password using the serial console in **Local Administration > Set Password**. If you know the current password, enter it along with the new password to change the password.

---

**Note:** You can only change the password in this way if you are using the current password (you cannot change the password using an SSH key).

---

If you do not know the **bal** password, you can reset it by following these steps:

1. Log in using **pwreset** as the username and **1pwreset** as the password. This resets the password to the default value (**1fourall**).

---

**Note:** Logging in using **pwreset** only works if the LoadMaster is licensed. This does not work on newly deployed units that are unlicensed.

---

---

**Note:** This default password is only valid for the console. The old **bal** password for the WUI is still valid at this point and the password only gets updated after completing the steps below. If the LoadMaster is rebooted at this point, the default password of **1fourall** no longer works, so you would need to log in with **pwreset** and **1pwreset** again to reset the password to the default password of **1fourall**.

---

2. Log in using the **bal** username **1fourall** password.
3. Then, go to **Local Administration > Set Password**.



4. Enter the current password (**1fourall**) and the new password.

---

# Troubleshooting

---

## Troubleshooting

The sections below provide some basic troubleshooting tips. If further assistance is required, please contact Progress Kemp Support: <https://support.kemptechnologies.com>.

### Related Links

- [Check which LoadMaster is Active](#)
- [First/Second Unconnected](#)
- [Connection to Default Gateway Failed](#)
- [Virtual Machine Inaccessible](#)
- [Run a TCP Dump](#)
- [Sync Problems](#)
- [Misconfigured ILB](#)
- [Problems Reaching a Virtual Service](#)

## Check which LoadMaster is Active

### Check which LoadMaster is Active

In addition to checking the status in the top-right of the LoadMaster WUI, it is also possible to check which LoadMaster is active by accessing port 8444 through the Public IP address since the Load Balanced Rule was created for this port, that is,

**http://<PublicIPofAzureLoadBalancer>:8444**

Ensure to use HTTP, not HTTPS. On the active unit, you should see "Active/Standby is active". On the standby, you should see a 503 service unavailable error. If you see these messages, it means the LoadMasters are working correctly/

## First/Second Unconnected

### First/Second Unconnected

When initially setting up cloud HA, the active unit should have **First** in the top-right corner of the LoadMaster WUI.

The standby unit should show **Second**.

After setting up the load balancer (Internal Load Balancer (ILB) for Azure or Network Load Balancer for AWS) the units should switch from:

- First to First Unconnected
- Second to Second Unconnected

This means the LoadMasters have not been polled by the load balancer. Once the load balancer has the health check correctly set, the units should switch from:

- First Unconnected to First (Active)/First (Standby)
- Second (Unconnected) to Second (Active)/Second (Standby)

# Connection to Default Gateway Failed

## Connection to Default Gateway Failed

License Required To Continue

Please enter your KEMP ID and password below to license this LoadMaster.


If you do not have a KEMP ID, please create one by visiting:  
<https://alsi2.kemptechnologies.com/register>

KEMP ID:

Password:

Order ID# (optional):

HTTP(S) Proxy (optional):



- ✖ Attempt to retrieve Licensing Types Failed: Error occurred. Please try again later.
- ✖ Connection to Default Gateway: (10.1.1.1 - Failed)
- ⊘ Connection to DNS: Stopped
- ⊘ Resolve Licensing Server FQDN: Stopped
- ⊘ Connection to Licensing Server: Stopped

Azure blocks pings in some cases. Therefore, on older LoadMaster firmware you may see an error message like the one above when licensing. This is a red herring and can be ignored - there is likely another problem such as an incorrect Kemp ID/password. If you are running the latest version of LoadMaster firmware, this check should be skipped.

# Virtual Machine Inaccessible

## Virtual Machine Inaccessible

It takes approximately five minutes for the Virtual Machine to become accessible after booting.

# Run a TCP Dump

## Run a TCP Dump

Running a TCP dump and checking the results can also assist with troubleshooting. To do this, follow the steps below in the LoadMaster WUI:

1. In the main menu, go to **System Configuration > Troubleshooting**.
2. In the **TCP dump** section, enter the relevant IP **Address** and the Azure HA **Port**.
3. Click **Start**.
4. Let the capture run for a few minutes.
5. Click **Stop**.
6. Click **Download**.
7. Analyse the results in a packet trace analyser tool such as [Wireshark](#).

Checks from the partner LoadMaster should appear in the results. If nothing is shown there is a problem, for example Azure may be blocking the connection.

For detailed information related to TCP dump, refer to the [Packet Trace Guide Technical Note](#).

# Sync Problems

## Sync Problems

In most scenarios, the configuration settings are automatically synchronized between partners every two minutes. If a new Virtual Service is created, the settings are immediately synchronized. Because of this, creating a new Virtual Service is a good way of checking if the synchronization is working. To trace this, follow the steps below:

1. Start a TCP dump, as detailed in the [Run a TCP Dump](#) section, but use port 6973.
2. Create a Virtual Service.
3. Stop the TCP dump.
4. Download the TCP dump file.
5. Analyse the results.

After creating a Virtual Service, a lot of traffic should have been immediately triggered.

Generally, if a lot of packets are being transferred it means that the synchronization is working. If only a few packets are transferred, it may mean that the connection was unsuccessful. In this case, there may be a problem such as unmatched SSH keys.

## Misconfigured ILB

### Misconfigured ILB

It is possible that the two LoadMasters are able to communicate but the ILB might be misconfigured. Connect to both units on `http://LoadMasterAddress:8444`. On the active unit, you should see "Active/Standby is active". On the standby, you should see a 503 service unavailable error. If you see these messages, it means the LoadMasters are working correctly and the problem is elsewhere. Confirm that the health check probe on the ILB is configured correctly.

## Problems Reaching a Virtual Service

### Problems Reaching a Virtual Service

If you experience problems reaching a Virtual Service, confirm the network security group and the ILB inbound rules are configured correctly.

---

## References

---

### References

Unless otherwise specified, the following documents can be found at <https://docs.progress.com/>.

**Licensing, Feature Description**

**LoadMaster for Azure, Installation Guide**

**Azure Virtual Machines – tutorials and guides:**

<http://www.windowsazure.com/en-us/documentation/services/virtual-machines/>

**High Availability (HA), Feature Description**