



LoadMaster HA for AWS

Installation Guide

UPDATED: 28 July 2023

© 2022 Progress Software Corporation and/or one of its subsidiaries or affiliates. All rights reserved.

These materials and all Progress® software products are copyrighted and all rights are reserved by Progress Software Corporation. The information in these materials is subject to change without notice, and Progress Software Corporation assumes no responsibility for any errors that may appear therein. The references in these materials to specific platforms supported are subject to change.

#1 Load Balancer in Price/Performance, 360 Central, 360 Vision, Chef, Chef (and design), Chef Habitat, Chef Infra, Code Can (and design), Compliance at Velocity, Corticon, Corticon.js, DataDirect (and design), DataDirect Cloud, DataDirect Connect, DataDirect Connect64, DataDirect XML Converters, DataDirect XQuery, DataRPM, Defrag This, Deliver More Than Expected, DevReach (and design), Driving Network Visibility, Flowmon, Inspec, Ipswitch, iMacros, K (stylized), Kemp, Kemp (and design), Kendo UI, Kinvey, LoadMaster, MessageWay, MOVEit, NativeChat, OpenEdge, Powered by Chef, Powered by Progress, Progress, Progress Software Developers Network, SequeLink, Sitefinity (and Design), Sitefinity, Sitefinity (and design), Sitefinity Insight, SpeedScript, Stylized Design (Arrow/3D Box logo), Stylized Design (C Chef logo), Stylized Design of Samurai, TeamPulse, Telerik, Telerik (and design), Test Studio, WebSpeed, WhatsConfigured, WhatsConnected, WhatsUp, and WS_FTP are registered trademarks of Progress Software Corporation or one of its affiliates or subsidiaries in the U.S. and/or other countries.

Analytics360, AppServer, BusinessEdge, Chef Automate, Chef Compliance, Chef Desktop, Chef Workstation, Corticon Rules, Data Access, DataDirect Autonomous REST Connector, DataDirect Spy, DevCraft, Fiddler, Fiddler Classic, Fiddler Everywhere, Fiddler Jam, FiddlerCap, FiddlerCore, FiddlerScript, Hybrid Data Pipeline, iMail, InstaRelinker, JustAssembly, JustDecompile, JustMock, KendoReact, OpenAccess, PASOE, Pro2, ProDataSet, Progress Results, Progress Software, ProVision, PSE Pro, Push Jobs, SafeSpaceVR, Sitefinity Cloud, Sitefinity CMS, Sitefinity Digital Experience Cloud, Sitefinity Feather, Sitefinity Thunder, SmartBrowser, SmartComponent, SmartDataBrowser, SmartDataObjects, SmartDataView, SmartDialog, SmartFolder, SmartFrame, SmartObjects, SmartPanel, SmartQuery, SmartViewer, SmartWindow, Supermarket, SupportLink, Unite UX, and WebClient are trademarks or service marks of Progress Software Corporation and/or its subsidiaries or affiliates in the U.S. and other countries. Java is a registered trademark of Oracle and/or its affiliates. Any other marks contained herein may be trademarks of their respective owners.

Please refer to the NOTICE.txt or Release Notes – Third-Party Acknowledgements file applicable to a particular Progress product/hosted service offering release for any related required third-party acknowledgements.

Table of Contents

1 Introduction	5
1.1 Document Purpose	5
1.2 Intended Audience	5
1.3 Prerequisites	5
2 AWS Network Load Balancing Service Architecture	7
3 Terminology Differences	11
4 Using LoadMaster HA for AWS	12
5 Configure GEO Clusters with HA	14
5.1 Configure the Master IP Address for one HA Pair	14
5.2 Configure Clusters of HA IP Addresses Between Two HA Pairs	14
5.3 Configuring Clusters of HA IP Addresses Between a HA Pair and a Standalone System	16
5.4 Configure Clusters of HA IP Addresses Between Two HA Pairs and Two Standalone LoadMasters	17
6 Creating AWS HA Pairs	20
6.1 Create the Network Load Balancer in AWS	20
6.2 Configure the LoadMaster	24
6.2.1 Virtual Service Restrictions	26
7 LoadMaster Firmware Upgrades/Downgrades	27
7.1 Upgrade the LoadMaster Firmware	27
7.2 Downgrade the LoadMaster Firmware	27
8 Master/Slave Unconnected	29

References	30
Last Updated Date	31

1 Introduction

The Kemp LoadMaster system can be deployed as a single unit or in an active/standby dual-unit configuration (High Availability (HA)). HA allows two physical or virtual machines to become one logical device. Only one of these units is active and handling traffic at any one time. One unit is active and the other is a hot standby (passive). This provides redundancy and resiliency, meaning if one LoadMaster goes down for any reason, the hot standby can become active, therefore minimizing any downtime.

The AWS Network Load Balancer is used to achieve HA in AWS when using Kemp LoadMasters. The Network Load Balancer does not leverage multiplexing therefore the LoadMaster persistence options can be enabled.

Network Load Balancing automatically distributes incoming application traffic across multiple Amazon EC2 instances in the cloud. Network Load Balancing ensures that only healthy Amazon EC2 instances receive traffic by detecting unhealthy instances and re-routing traffic across the remaining healthy instances.

Placing the Kemp LoadMasters behind the Network Load Balancer enables advanced application delivery functionality.

1.1 Document Purpose

The purpose of this document is to provide information and step-by-step instructions on how to configure HA when using the Kemp LoadMaster in AWS.

1.2 Intended Audience

This document is intended to be read by anyone who is interested in finding out how to configure HA when using the Kemp LoadMaster in an AWS environment.

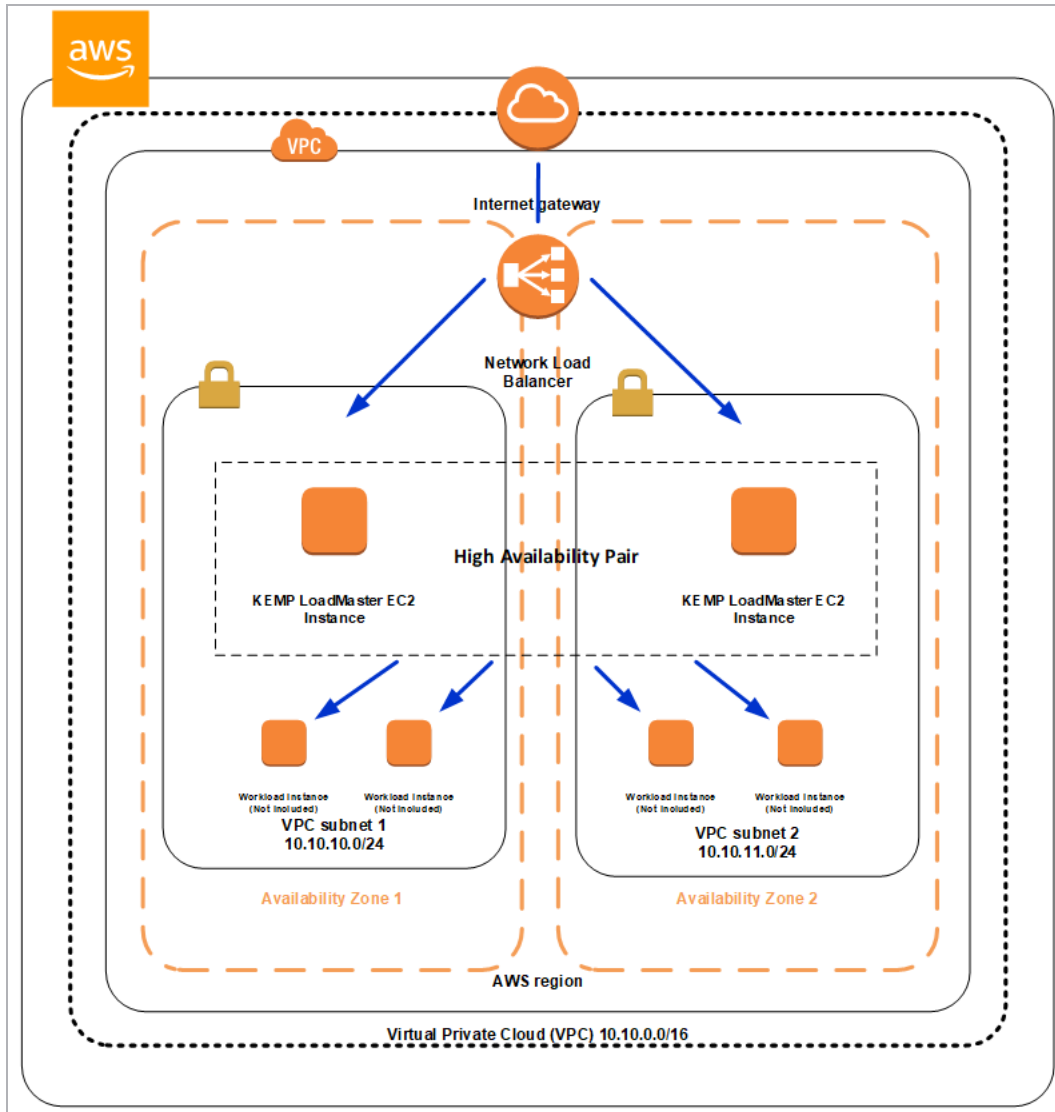
1.3 Prerequisites

This document assumes that you already have two LoadMaster HA instances which are configured and accessible using the User Interface (UI). For instructions on how to do this, refer to the **LoadMaster for AWS Feature Description** on the [Kemp Documentation Page](#). One should be designated as a master and the other as a slave.

Due to AWS limitations, it is not possible to bond interfaces on AWS LoadMasters.

For step-by-step instructions on how to deploy a LoadMaster in AWS, refer to the [LoadMaster for AWS Installation Guide](#).

2 AWS Network Load Balancing Service Architecture



There are two logical components in the Network Load Balancing service architecture:

- Load balancers
- A controller service

The load balancers are resources that monitor traffic and handle requests that come in through the Internet, that is, the Kemp LoadMaster.

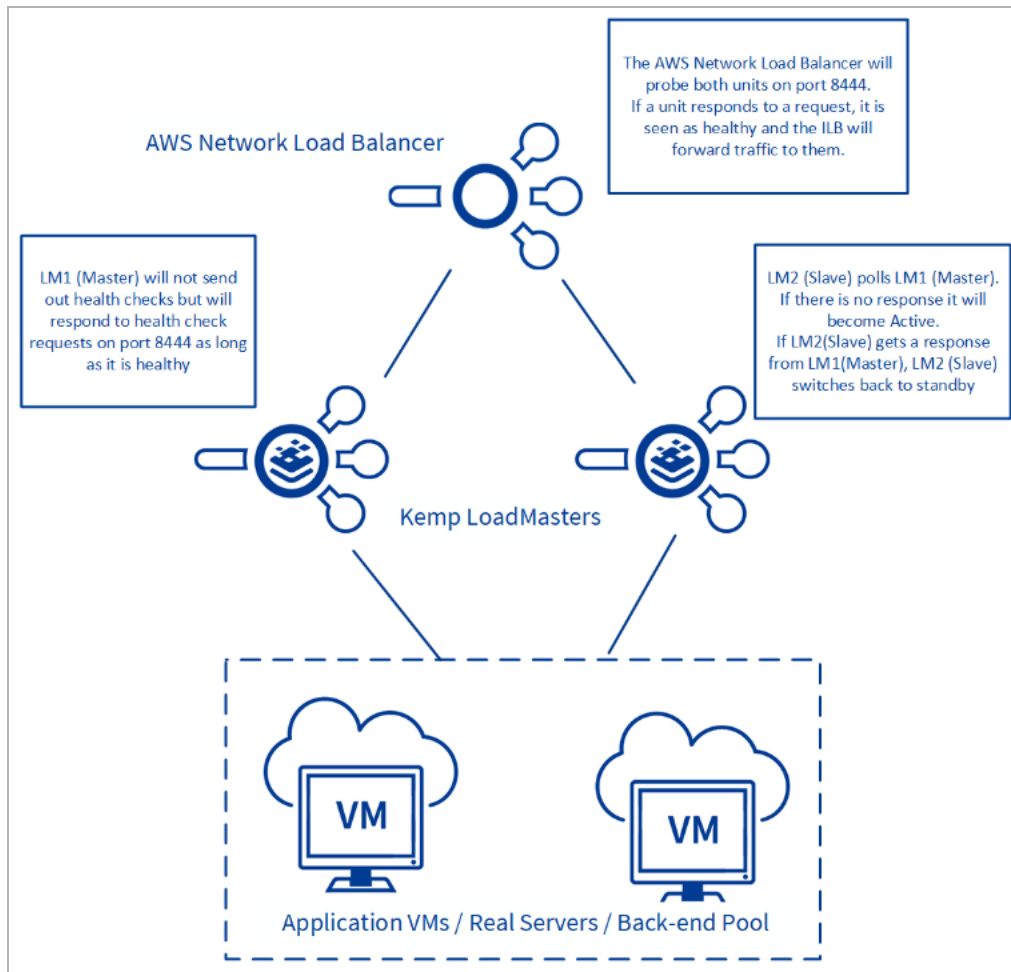
The controller service monitors the load balancers and verifies that load balancers are behaving properly.

Once you create a network load balancer, you must configure it to accept incoming traffic and route requests to your EC2 instances. These configuration parameters are stored by the controller, and the controller ensures that all of the load balancers are operating with the correct configuration.

Network Load Balancing will perform health checks on back-end instances, using the configuration that you supply.

To discover the availability of your EC2 instances, the load balancer periodically sends pings, attempts connections, or sends requests to test the EC2 instances. These tests are called health checks. Instances that are healthy at the time of the health check are marked as **InService** and the instances that are unhealthy at the time of the health check are marked as **OutOfService**. The load balancer performs health checks on all registered instances, whether the instance is in a healthy state or an unhealthy state. When using AWS VLMs in HA mode – one unit is active and in service, the other is stand-by and out-of-service.

2 AWS Network Load Balancing Service Architecture



The load balancer routes traffic only to the healthy instances. When the load balancer determines that an instance is unhealthy, it stops routing traffic to that instance. The load balancer resumes routing traffic to the instance when it has been restored to a healthy state.

The load balancer checks the health of the registered instances using either the default health check configuration provided by Network Load Balancing or a health check configuration that you configure.

The health checks must reach the defined target set in the Network Load Balancing configuration for the number of successful checks before the instance is considered to be “in service” and healthy. For example, for any instance registered with Network Load Balancing - if you set the interval for health checks to 20 seconds, and you set the number of successful health checks to 10, then it will take at least 200 seconds before Network Load Balancing will route traffic to the instance.

The health check also defines a failure threshold. For example, if you set the interval to 20 seconds and you set the failure threshold at 4, then when an instance no longer responds to requests - at least 80 seconds will elapse before it is taken out of service. However, if an instance is terminated, traffic will no longer be sent to the terminated instance, but there can be a delay before the load balancer is aware that the instance was terminated. For this reason, it is important to de-register your instances before terminating them; instances are removed from service in a much shorter amount of time if they are de-registered.

3 Terminology Differences

There are some terminology differences between "normal" (non-cloud) High Availability (HA) LoadMaster units and cloud HA units. The table below outlines these differences:

Platform	Unit number	Terminology	Statuses
Non-cloud	Unit 1	HA first	Master/Standby
Non-cloud	Unit 2	HA second	Master/Standby
Cloud	Unit 1	Master	Active/Standby
Cloud	Unit 2	Slave	Active/Standby

4 Using LoadMaster HA for AWS

When using LoadMaster in High Availability on AWS, HA operates in much the same way as it does on non-cloud platforms, but with some key differences due to how HA interacts with the AWS Elastic IP feature:

- LoadMaster HA for AWS involves two LoadMasters that synchronize. Changes made to the active LoadMaster are replicated to the standby LoadMaster.
- When synchronizing the GEO settings from active to standby, any Fully Qualified Domain Name (FQDN) or cluster IP addresses that match the active's IP address are replaced with the standby's IP address. Likewise, when synchronizing from standby to active, the standby's IP address is replaced with the active's 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
 - 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.
- Both devices are capable of responding to Network Load Balancer health check requests.
 - The LoadMaster that is currently handling client traffic will respond with the status code **200 OK** to the AWS health check - meaning that it is healthy. Meanwhile, the standby LoadMaster will respond with the status code **503** -- meaning that it is unhealthy. In this way, all client requests are redirected by the Network Load Balancer to the healthy LoadMaster.
 - The “standby” LoadMaster (the LoadMaster which is not handling traffic) polls the “active” LoadMaster to check the availability of the service. If the probe is successful, it remains in “standby” mode, otherwise it takes over as the “active” - answering 200 OK to the AWS health check becoming capable to handle traffic.

If the master unit fails, connections are directed to the slave unit. The master unit never assumes the slave role and the slave never becomes the master. When the master unit becomes available again after a failure, connections are automatically directed to the master unit again. The master can be active or standby. The slave can be active or standby.

You can set the **Switch to Preferred Server** option to **No Preferred Host** which allows the slave unit to maintain the active state when the master unit comes back online.

For HA to work, the two LoadMasters must have different values set for the **AWS HA Mode**.

Kemp recommends always using NIC0 for HA checks on AWS.

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

5 Configure GEO Clusters with HA

There is no shared IP addresses for cloud HA LoadMasters, so individual LoadMaster IP addresses must be used when configuring GEO clusters. When creating a GEO cluster, use the IP address of the master LoadMaster. If there is a failover and the standby machine takes the active role, the cluster IP address changes automatically.

To configure a working configuration of a HA pair with GEO clusters, follow the steps below, depending on your type of configuration.

5.1 Configure the Master IP Address for one HA Pair

To configure the master IP address as a cluster on a HA pair, follow these steps in the LoadMaster WUI:

1. In the main menu, go to **Certificates & Security > Remote Access**.
2. Enter the IP address of the master LoadMaster in the **Remote GEO LoadMaster Access** text box and click **Set GEO LoadMaster access**.
3. In the main menu, go to **Global Balancing > Manage Clusters**.
4. In the **IP address** text box, enter the master LoadMaster IP address.
5. Enter a **Name** for the cluster and click **Add Cluster**.
6. Select **Remote LM** in the **Type** drop-down list.

5.2 Configure Clusters of HA IP Addresses Between Two HA Pairs

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

On HA pair 1, complete these steps:

1. In the main menu, go to **Certificates & Security > Remote Access**.

2. Enter the IP address of the master LoadMaster and the IP addresses of both partners in the second HA pair in the **Remote GEO LoadMaster Access** text box and click **Set GEO LoadMaster access**.
3. Enter the IP addresses of both partners in the second HA pair In the **GEO LoadMaster Partners** text box and click **Set GEO LoadMaster Partners**.

On HA pair 2, complete these steps:

1. In the main menu, go to **Certificates & Security > Remote Access**.
2. Enter the IP address of the master LoadMaster and the IP addresses of both partners in the first HA pair in the **Remote GEO LoadMaster Access** text box and click **Set GEO LoadMaster access**.
3. Enter the IP addresses of both partners in the first HA pair In the **GEO LoadMaster Partners** text box and click **Set GEO LoadMaster Partners**.

Then, complete these steps:

1. Reboot the master LoadMaster of the first HA pair.
2. Submit the **GEO LoadMaster Partners** configuration again on the new active LoadMaster.
3. Reboot the master LoadMaster of the second HA pair.
4. Submit the **GEO LoadMaster Partners** configuration again on the new active LoadMaster.

Then, create the GEO clusters with the **TypeRemote LM** using the master IP address of both HA pairs:

1. In the main menu, go to **Global Balancing > Manage Clusters**.
2. In the **IP address** text box, enter the master LoadMaster IP address.
3. Enter a **Name** for the cluster and click **Add Cluster**.
4. Select **Remote LM** in the **Type** drop-down list.

5.3 Configuring Clusters of HA IP Addresses Between a HA Pair and a Standalone System

To configure clusters of HA IP addresses between a HA pair and a standalone system, follow the steps below.

On the HA pair LoadMaster, complete these steps:

1. In the main menu, go to **Certificates & Security > Remote Access**.
2. Enter the IP address of the master LoadMaster and the IP address of the standalone LoadMaster in the **Remote GEO LoadMaster Access** text box and click **Set GEO LoadMaster access**.
3. Enter the IP address of the standalone system in the **GEO LoadMaster Partners** text box and click **Set GEO LoadMaster Partners**.

On the standalone system, complete these steps:

1. In the main menu, go to **Certificates & Security > Remote Access**.
2. Enter the IP address of the standalone LoadMaster and the IP addresses of both partners in the HA pair in the **Remote GEO LoadMaster Access** text box and click **Set GEO LoadMaster access**.

Then, complete these steps:

1. Reboot the master LoadMaster of the HA pair.
2. Submit the **GEO LoadMaster Partners** configuration again on the new active LoadMaster.
3. On the standalone LoadMaster, submit the **GEO LoadMaster Partners** configuration again.

Then, create the GEO clusters with the **TypeRemote LM** using the master IP address of the HA pair and the IP address of the standalone LoadMaster:

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

2. In the **IP address** text box, enter the LoadMaster IP address.
3. Enter a **Name** for the cluster and click **Add Cluster**.
4. Select **Remote LM** in the **Type** drop-down list.

5.4 Configure Clusters of HA IP Addresses Between Two HA Pairs and Two Standalone LoadMasters

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. In the main menu, go to **Certificates & Security > Remote Access**.
2. Enter the IP address of the master LoadMaster, the IP addresses of both partners of the second HA pair, and the IP addresses of both standalone LoadMasters in the **Remote GEO LoadMaster Access** text box and click **Set GEO LoadMaster access**.
3. Enter the IP addresses of both partners of the second HA pair and the IP addresses of both standalone LoadMasters in the **GEO LoadMaster Partners** text box and click **Set GEO LoadMaster Partners**.

On HA pair 2, complete these steps:

1. In the main menu, go to **Certificates & Security > Remote Access**.
2. Enter the IP address of the master LoadMaster, the IP addresses of both partners of the first HA pair, and the IP addresses of both standalone LoadMasters in the **Remote GEO LoadMaster Access** text box and click **Set GEO LoadMaster access**.
3. Enter the IP addresses of both partners of the first HA pair and the IP addresses of both standalone LoadMasters in the **GEO LoadMaster Partners** text box and click **Set GEO LoadMaster Partners**.

On standalone LoadMaster 1, complete these steps:

1. In the main menu, go to **Certificates & Security > Remote Access**.

2. Enter the IP address of standalone LoadMaster 1, the IP addresses of both partners of both HA pairs, and the IP addresses of the second standalone LoadMaster in the **Remote GEO LoadMaster Access** text box and click **Set GEO LoadMaster access**.
3. Enter the IP addresses of both partners of HA pairs and the IP address of the second standalone LoadMasters in the **GEO LoadMaster Partners** text box and click **Set GEO LoadMaster Partners**.

On standalone LoadMaster 2, complete these steps:

1. In the main menu, go to **Certificates & Security > Remote Access**.
2. Enter the IP address of standalone LoadMaster 2, the IP addresses of both partners of both HA pairs, and the IP addresses of the first standalone LoadMaster in the **Remote GEO LoadMaster Access** text box and click **Set GEO LoadMaster access**.
3. Enter the IP addresses of both partners of HA pairs and the IP address of the first standalone LoadMasters in the **GEO LoadMaster Partners** text box and click **Set GEO LoadMaster Partners**.

Then, complete these steps:

1. Reboot the master LoadMaster of the first HA pair.
2. Submit the **GEO LoadMaster Partners** configuration again on the new active LoadMaster.
3. Reboot the master LoadMaster of the second HA pair.
4. Submit the **GEO LoadMaster Partners** configuration again on the new active LoadMaster.
5. On both standalone LoadMasters, submit the **GEO LoadMaster Partners** configuration again.

Then, create the GEO clusters with the **TypeRemote LM** using the master IP address of both HA pairs and the IP addresses of both standalone LoadMasters:

1. In the main menu, go to **Global Balancing > Manage Clusters**.
2. In the **IP address** text box, enter the LoadMaster IP address.
3. Enter a **Name** for the cluster and click **Add Cluster**.

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

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

HA Pair 1

172.24.1.5

172.24.1.6

Remote GEO LoadMaster Access: 172.24.1.7 172.24.1.8 172.24.1.5 172.24.1.9 172.24.1.10

GEO LoadMaster Partners: 172.24.1.7 172.24.1.8 172.24.1.9 172.24.1.10

HA Pair 2

172.24.1.7

172.24.1.8

Remote GEO LoadMaster Access: 172.24.1.8 172.24.1.5 172.24.1.6 172.24.1.9 172.24.1.10

GEO LoadMaster Partners: 172.24.1.5 172.24.1.6 172.24.1.9 172.24.1.10

Standalone LoadMaster 1

172.24.1.9

Remote GEO LoadMaster Access: 172.24.1.7 172.24.1.8 172.24.1.5 172.24.1.6 172.24.1.9 172.24.1.10

GEO LoadMaster Partners: 172.24.1.7 172.24.1.8 172.24.1.5 172.24.1.6 172.24.1.10

Standalone LoadMaster 2

172.24.1.10

Remote GEO LoadMaster Access: 172.24.1.7 172.24.1.8 172.24.1.5 172.24.1.6 172.24.1.9 172.24.1.10

GEO LoadMaster Partners: 172.24.1.7 172.24.1.8 172.24.1.5 172.24.1.6 172.24.1.9

6 Creating AWS HA Pairs

This document assumes that you already have two LoadMaster HA instances which are configured and accessible using the User Interface (UI). One should be designated as a master and the other as a slave.

TCP port 6973 must be allowed in the inbound rules for each LoadMaster in the HA pair to allow for synchronization traffic.

For further information and steps on how to deploy an individual LoadMaster instance, refer to the [LoadMaster for AWS Installation Guide](#) document.


There are options when choosing a Load Balancer in AWS to support the LoadMaster HA Pair. This document covers the deployment of a Network Load Balancer to allow for different persistence options on the Kemp LoadMaster. You can choose the Classic Load Balancer if persistence is not required on the LoadMaster.

6.1 Create the Network Load Balancer in AWS

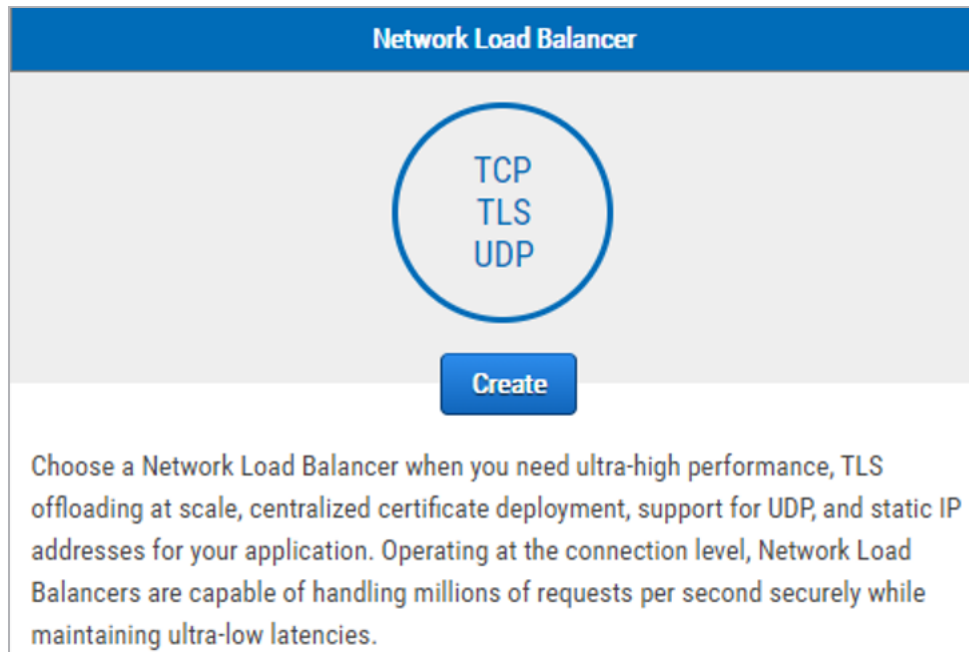
Before creating a Network Load Balancer in AWS, it is best practice to have an existing Elastic IP (EIP) allocated. To find out how to allocate an Elastic IP address, refer to the following AWS link: [Elastic IP addresses](#).

To create AWS HA pairs, carry out the following steps:

1. Open the Amazon EC2 console.
2. Navigate to **Load Balancing > Load Balancers**.



3. Click **Create Load Balancer**.



4. Click **Create** for **Network Load Balancer**.

Step 1: Configure Load Balancer

Basic Configuration

To configure your load balancer, provide a name, select a scheme, specify one or more listeners, and select a network. The default configuration is an Internet-facing load balancer in the selected network with a listener that receives TCP traffic on port 80.

Name ⓘ

Scheme ⓘ ☒ internet-facing
☐ internal

Listeners

A listener is a process that checks for connection requests, using the protocol and port that you configured.

Load Balancer Protocol	Load Balancer Port
TCP	80

[Add listener](#)

Availability Zones

Specify the Availability Zones to enable for your load balancer. The load balancer routes traffic to the targets in these Availability Zones only. You can specify only one subnet per Availability Zone. You may also add one Elastic IP per Availability Zone if you wish.

[Create and manage Elastic IPs in the VPC console](#)

VPC ⓘ

Availability Zones

☒ us-east-1a

IPv4 address ⓘ

☐ us-east-1f

Temporary limitation
Choose your Availability Zones and subnets with care. After you create the load balancer, you cannot disable the enabled subnets, but you can enable additional ones.

5. Set the following options:

6 Creating AWS HA Pairs

- a) Enter a **Name** for the Load Balancer.
- b) Select whether to provide access to **internal** only or allow for an **internet-facing** load balancer.
- c) Add one or more **Listeners** based on the application being load balanced.
- d) Select the **VPC** and **Availability Zones** to deploy the Network Load Balancer in.
- e) Select **Choose an Elastic IP** and select the EIP that was allocated earlier.
- f) Add optional **Tags**.

6. Click **Next: Configuration Security Settings**.

Step 2: Configure Security Settings

Select default certificate

AWS Certificate Manager (ACM) is the preferred tool to provision and store server certificates. If you previously stored a server certificate using IAM, you can deploy it to your load balancer. [Learn more](#) about TLS listeners and certificate management.

Certificate type ⓘ

- ☒ Choose a certificate from ACM (recommended)
- ☐ Upload a certificate to ACM (recommended)
- ☐ Choose a certificate from IAM
- ☐ Upload a certificate to IAM

[Request a new certificate from ACM](#)
AWS Certificate Manager makes it easy to provision, manage, deploy, and renew SSL Certificates on the AWS platform. ACM manages certificate renewals for you. [Learn more](#)

Certificate name ⓘ ⓘ

Select Security Policy

Security policy ⓘ

ALPN Policy

ALPN, or Application-Layer Protocol Negotiation, is a TLS extension that includes the protocol negotiation within the exchange of hello messages. Selecting a policy (anything other than None) enables this listener attribute across all the TLS listeners within this load balancer. Once you have selected a policy, you can manage this listener attribute at the individual listener level.

ALPN Policy ⓘ

- ☒ **None**
Do not accept ALPN.
- ☐ **HTTP1Only**
Allow only HTTP/1 connections.
- ☐ **HTTP2Only**
Allow only HTTP/2 connections.
- ☐ **HTTP2Optional**
Prefer HTTP/1 connections; accept HTTP/2 connections.
- ☐ **HTTP2Preferred**
Prefer HTTP/2 connections; allow falling back to HTTP/1.

7. If you selected a TLS Listener, you can create or upload a certificate to encrypt the traffic.
8. Click **Next: Configure Routing**.

Step 3: Configure Routing
Your load balancer routes requests to the targets in this target group using the protocol and port that you specify.

Target group

Target group ⓘ

Name ⓘ

Target type
☒ Instance
☐ IP

Protocol ⓘ

Port ⓘ

Health checks

Protocol ⓘ

Path ⓘ

▼ **Advanced health check settings**

Port ⓘ ☐ traffic port
☒ override

Healthy threshold ⓘ

Unhealthy threshold ⓘ

Timeout ⓘ seconds

Interval ⓘ ☒ 10 seconds
☐ 30 seconds

Success codes ⓘ

9. Set the following options:

- Select a new **Target group**.
- Enter a **Name** for the target group.
- Select **Instance** for **Target type**.
- Select the **Protocol** and **Port** for routing traffic to the Kemp LoadMasters.
- Select **HTTP** as the health check **Protocol**.
- Enter **/** for the health check **Path**.
- Under **Advanced health check settings**, select **override** for **Port**.
- Enter **8444** for the **override** port.
- Enter **2** as the **Healthy threshold**.
- Select **10 seconds** as the **Interval**.

10. Click **Next: Register Targets**.

6 Creating AWS HA Pairs

Add to registered on port 80							
Search Instances							
<input type="checkbox"/>	Instance	Name	State	Security groups	Zone	Subnet ID	Subnet CIDR
<input checked="" type="checkbox"/>	i-09201a5d90e2f164b	Kemp-HA-1	running	Kemp Load Balancer ADC - B...	us-east-2a	subnet-f3b5679a	172.31.0.0/20
<input checked="" type="checkbox"/>	i-0a240cae9baff7535	Kemp-HA-2	running	Kemp Load Balancer ADC - B...	us-east-2a	subnet-f3b5679a	172.31.0.0/20
<input type="checkbox"/>	i-00305d89d9dbf4e40	MELA-VLM-AWS-FNT	running	License Agreement Based L...	us-east-2a	subnet-f3b5679a	172.31.0.0/20

11. Select the two LoadMasters in the HA pair and click **Add to registered**.
12. Click **Next: Review**.
13. Click **Create**.

6.2 Configure the LoadMaster

Complete the following steps to configure the LoadMaster settings:

1. Log in to the UI of the master LoadMaster.
2. In the main menu, go to **System Configuration > AWS HA Parameters**.

AWS HA Mode	Master HA Mode	
Switch to Preferred Server	Prefer Master	
Partner Name/IP	172.31.85.198	Set Partner Name/IP
Health Check Port	8444	Set Health Check Port
Health Check on All Interfaces	<input type="checkbox"/>	

3. Select **Master HA Mode** from the **AWS 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 Master:** The HA1 (master) unit always takes over. This is the default option.
5. Enter the IP address of the slave LoadMaster in the **Partner Name/IP** text box and click **Set Partner Name/IP**.

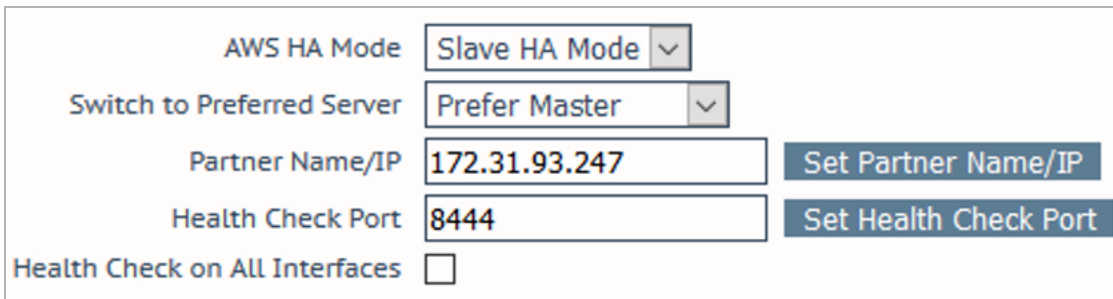
6. Enter the health check port selected earlier in the **Health Check Port** text box and click **Set Health Check Port**.

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

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

8. Log in to the UI of the slave LoadMaster.

9. In the main menu, go to **AWS HA Parameters**.



10. Select **Slave HA Mode** from the **AWS HA Mode** drop-down list.

11. Enter the IP address of the master LoadMaster in the **Partner Name/IP** text box and click **Set Partner Name/IP**.

12. Enter the health check port selected earlier in the **Health Check Port** text box and click **Set Health Check Port**.


The **Health Check Port** must be the same on both the master and slave units in order for HA to function correctly.

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

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

In the Amazon EC2 console, go to the ELB and select the **Instances** tab. The master instance should be marked as **InService**. The slave instance should be marked as **OutOfService**.

In the LoadMaster, set up a HTTP and HTTPS Virtual Service with Real Servers. These should then be available using the ELB Domain and they should properly fail over.



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.

6.2.1 Virtual Service Restrictions

There are some situations where Virtual Service settings may prevent HA from functioning correctly. Please follow the guidelines below to avoid any issues:

- Do not set up a Virtual Service on the same port as the health check port
- Do not set up a TCP Virtual Service on port 6973 on the interface where HA sync is configured
- Do not set up a TCP Virtual Service on port 22 on a LoadMaster interface port

7 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 Kemp documentation page: <https://kemptechnologies.com/loadmaster-documentation>.

7.1 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 Slave unit may failback over to the Master unit.

After these steps are completed the upgrade is finished.

7.2 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 Slave unit may failback over to the Master unit.

After these steps are completed the downgrade is finished.

8 Master/Slave Unconnected

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

The slave unit should show **SLAVE**.

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

- Master to Master Unconnected
- Slave to Slave 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:

- Master Unconnected to Master (Active)/Master (Standby)
- Slave (Unconnected) to Slave (Active)/Slave (Standby)

References

Unless otherwise specified, the following documents can be found at <http://kemptechnologies.com/documentation>.

LoadMaster for AWS, Installation Guide

High Availability (HA), Feature Description

Last Updated Date

This document was last updated on 28 July 2023.