



HA for AWS

Installation Guide

UPDATED: 27 July 2023

Copyright Notices

© 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	4
1.1 Document Purpose	4
1.2 Intended Audience	4
1.3 Prerequisites	4
2 AWS Elastic Load Balancing Service Architecture	6
3 Terminology Differences	10
4 Using LoadMaster HA for AWS	11
5 Creating AWS HA Pairs	13
5.1 Create the Classic Load Balancer in AWS	13
5.2 Modify the Inbound Security Rule	19
5.3 Configure the LoadMaster	19
5.3.1 Virtual Service Restrictions	21
6 LoadMaster Firmware Upgrades/Downgrades	22
6.1 Upgrade the LoadMaster Firmware	22
6.2 Downgrade the LoadMaster Firmware	22
7 Master/Slave Unconnected	24
References	25
Last Updated Date	26

1 Introduction

The Kemp LoadMaster system can be deployed as a single unit or in an active/standby dual-unit configuration (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.

AWS Elastic Load Balancing is used to achieve HA in AWS when using Kemp LoadMasters.

Elastic Load Balancing automatically distributes incoming application traffic across multiple Amazon EC2 instances in the cloud. Elastic 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 Elastic 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.

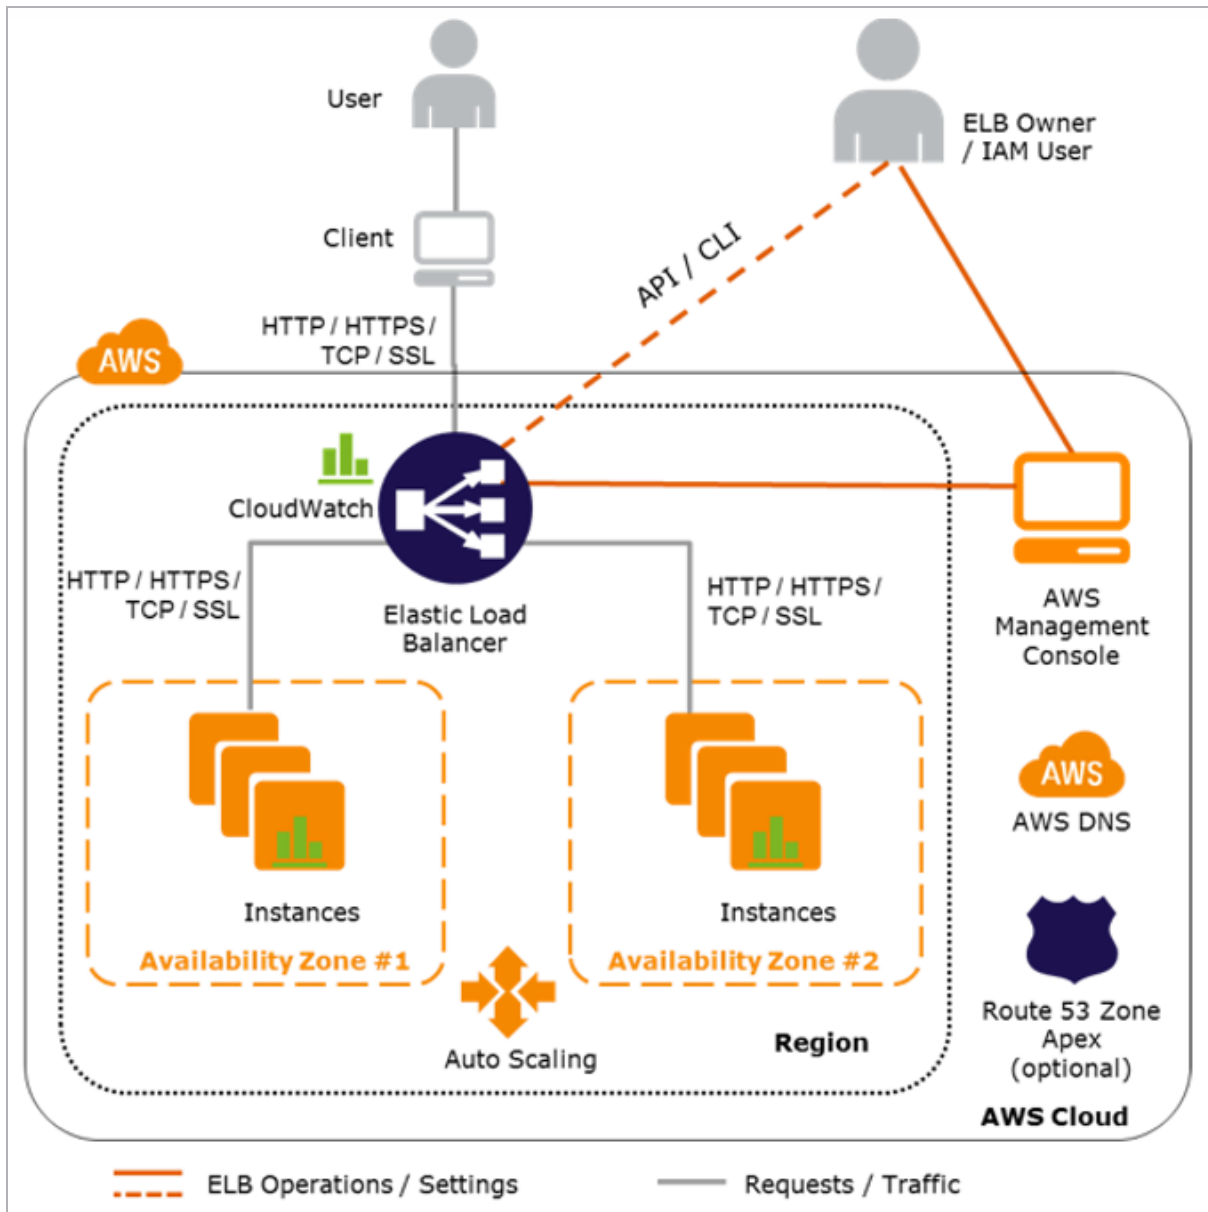
1.3 Prerequisites

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, please refer to the [LoadMaster for AWS, Feature Description](#) document.

2 AWS Elastic Load Balancing Service Architecture



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

2 AWS Elastic 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.

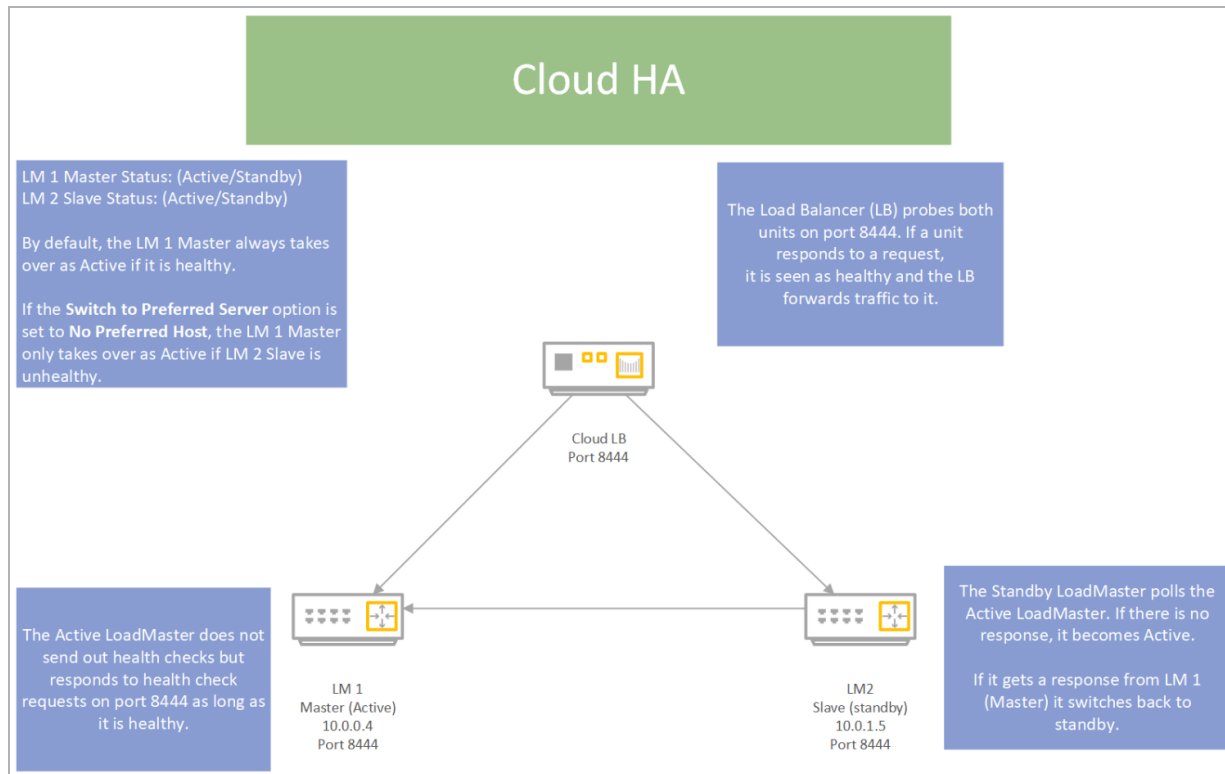
The controller service monitors the load balancers and verifies that load balancers are behaving properly.

Once you create a classic 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.

Elastic 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 Elastic 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 Elastic Load Balancing or a health check configuration that you configure.

The health checks must reach the defined target set in the Elastic 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 Elastic 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 Elastic 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 classic 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 Classic 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**.

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

5 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.

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

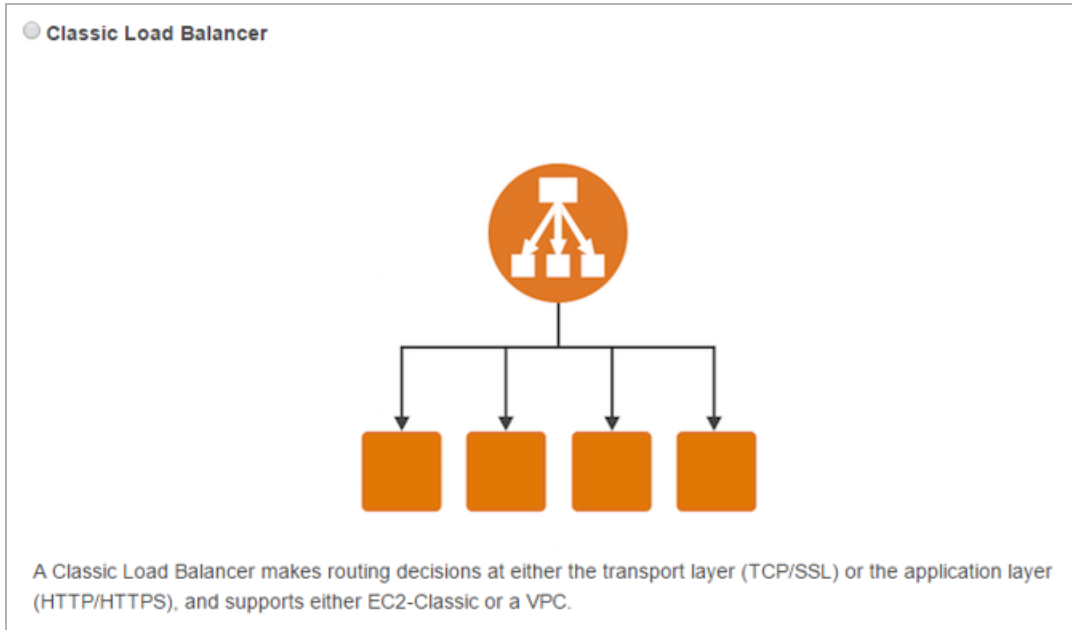
5.1 Create the Classic Load Balancer in AWS

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. Select **Classic Load Balancer**.

Step 1: Define Load Balancer

Basic Configuration

This wizard will walk you through setting up a new load balancer. Begin by giving your new load balancer a unique name so that you can identify it from other load balancers you might create. You can also specify any port on your EC2 instances. By default, we've configured your load balancer with a standard web server on port 80.

Load Balancer name:

Create LB inside: ⌵

Create an internal load balancer: ☐ [\(what's this?\)](#)

Listener Configuration:

Load Balancer Protocol	Load Balancer Port	Instance Protocol
<input type="text" value="HTTP"/> ⌵	<input type="text" value="80"/>	<input type="text" value="HTTP"/> ⌵
<input type="text" value="HTTPS (Secure HTTP)"/> ⌵	<input type="text" value="443"/>	<input type="text" value="HTTPS (Secure HTTP)"/> ⌵

Select Subnets

You will need to select a Subnet for each Availability Zone where you wish traffic to be routed by your load balancer. If you have instances in only one Availability Zone, please select at least two Subnets.

VPC vpc-6b475509 (172.31.0.0/16) | Default

Available Subnets

Actions	Availability Zone	Subnet ID
<input type="checkbox"/>	us-east-1b	subnet-00ec9f28
<input type="checkbox"/>	us-east-1c	subnet-e6764792
<input type="checkbox"/>	us-east-1d	subnet-46124100

Selected Subnets

Actions	Availability Zone	Subnet ID
<input checked="" type="checkbox"/>	us-east-1b	subnet-00ec9f28
<input checked="" type="checkbox"/>	us-east-1c	subnet-e6764792
<input checked="" type="checkbox"/>	us-east-1d	subnet-46124100

5. Set the name and Virtual Private Cloud (VPC).

6. Add listeners as necessary. Typically HTTP and HTTPS are used.

Other listeners can be added as required.

7. Add selected subnets, as needed.

Step 2: Assign Security Groups
 You have selected the option of having your Elastic Load Balancer inside of a VPC, which allows you to assign security groups to your load balancer. Please select the security groups to assign to

Assign a security group: ☐ Create a new security group
☒ Select an existing security group

Security Group ID	Name	Description
<input type="checkbox"/> sg-74e70c11	default	default VPC security group
<input type="checkbox"/> sg-11289975	launch-wizard-10	launch-wizard-10 created 2015-02-16T15:22:47.984+00:00
<input type="checkbox"/> sg-1ae95a7e	launch-wizard-11	launch-wizard-11 created 2015-02-17T15:26:09.911+00:00
<input type="checkbox"/> sg-7fd1621b	launch-wizard-12	launch-wizard-12 created 2015-02-17T15:50:21.996+00:00
<input type="checkbox"/> sg-d5145cb1	launch-wizard-13	launch-wizard-13 created 2015-04-24T10:36:54.143+01:00
<input type="checkbox"/> sg-e928608d	launch-wizard-14	launch-wizard-14 created 2015-04-24T10:53:13.969+01:00
<input type="checkbox"/> sg-66236b02	launch-wizard-15	launch-wizard-15 created 2015-04-24T11:35:14.956+01:00
<input type="checkbox"/> sg-0345b567	launch-wizard-2	launch-wizard-2 created 2014-12-15T12:44:30.927+00:00
<input type="checkbox"/> sg-798a581d	launch-wizard-3	launch-wizard-3 created 2015-01-21T12:26:10.113+00:00
<input type="checkbox"/> sg-639b3b07	launch-wizard-4	launch-wizard-4 created 2015-02-03T17:47:12.426+00:00
<input type="checkbox"/> sg-5895363c	launch-wizard-5	launch-wizard-5 created 2015-02-04T11:22:26.240+00:00
<input type="checkbox"/> sg-6be74b0f	launch-wizard-6	launch-wizard-6 created 2015-02-06T14:26:13.807+00:00
<input type="checkbox"/> sg-39c16d5d	launch-wizard-7	launch-wizard-7 created 2015-02-06T16:40:56.512+00:00
<input type="checkbox"/> sg-c79c32a3	launch-wizard-8	launch-wizard-8 created 2015-02-08T22:31:48.526+00:00
<input type="checkbox"/> sg-7a3f8e1e	launch-wizard-9	launch-wizard-9 created 2015-02-16T14:28:50.297+00:00

8. Assign a security group which allows load-balanced traffic.

Step 3: Configure Security Settings

Select Certificate

An SSL Certificate allows you to configure the HTTPS/SSL listeners of your load balancer. You may select a previously uploaded certificate below.

Certificate Type: ☒ Choose an **existing** SSL Certificate
☐ Upload a **new** SSL Certificate

Certificate Name:

Select a Cipher

Configure SSL negotiation settings for the HTTPS/SSL listeners of your load balancer. You may select one of the Security Policies listed below.

☒ Predefined Security Policy

☐ Custom Security Policy

SSL Protocols

- ☐ Protocol-SSLv2
- ☒ Protocol-TLSv1
- ☐ Protocol-SSLv3
- ☒ Protocol-TLSv1.1
- ☒ Protocol-TLSv1.2

SSL Options

- ☒ Server Order Preference

SSL Ciphers

- ☒ ECDHE-ECDSA-AES128-GCM-SHA256
- ☒ ECDHE-RSA-AES128-GCM-SHA256
- ☒ ECDHE-ECDSA-AES128-SHA256

9. Select an SSL certificate, if needed. This is only needed if load balancing HTTPS on the Internal Load Balancer (ILB).

10. Select the desired ciphers and protocols.

Step 4: Configure Health Check

Your load balancer will automatically perform health checks on your EC2 instances

Ping Protocol

Ping Port

Ping Path

Advanced Details

Response Timeout seconds

Health Check Interval seconds

Unhealthy Threshold

Healthy Threshold

11. Select **HTTP** as the **Ping Protocol** for the health check.
12. Enter **/** as the **Ping Path**.
13. Set the **Ping Port** to **8444**.

The port can be changed as needed.

Kemp strongly recommend not lowering the timeout and interval values. If the settings are lower, it may cause false health check failures.

Step 5: Add EC2 Instances

The table below lists all your running EC2 instances. Check the boxes in the Select column to add those instances to this load balancer.

VPC vpc-6b475509 (172.31.0.0/16) | Default

Select	Instance ID	Instance Name	Status	Availability Zone
<input type="checkbox"/>	i-0b2815a8	lm-hvm-1-4	running	LMINST
<input checked="" type="checkbox"/>	i-d8ba6d65	QA-AC-7.1-29-VLM_AWS_5-MASTER	running	LMINST
<input checked="" type="checkbox"/>	i-d8ba6d66	QA-AC-7.1-29-VLM_AWS_5-SLAVE	running	LMINST
<input type="checkbox"/>	i-5d5d31b1		stopped	LMINST
<input type="checkbox"/>	i-d38c1a7a	7.1-27-1002	running	LMINST
<input type="checkbox"/>	i-c5d03425		running	default
<input type="checkbox"/>	i-3c6bb9eb		running	LMINST
<input type="checkbox"/>	i-1c422bf0	testtest	stopped	LMINST
<input type="checkbox"/>	i-c17e997f	Fab-one	running	LMINST

Availability Zone Distribution
2 instances in us-east-1b

☒ Enable Cross-Zone Load Balancing

☒ Enable Connection Draining seconds

14. Select the LoadMaster instances.

Step 6: Add Tags

Apply tags to your resources to help organize and identify them.

A tag consists of a case-sensitive key-value pair. For example, you could define a tag with key = Name and value = Webserver. [Learn more](#) about tagging your Amazon EC2 resources.

Key	Value
<input type="text"/>	<input type="text"/>

Create Tag

15. Specify any tags, as desired.

Step 7: Review

Please review the load balancer details before continuing

▼ Define Load Balancer

Load Balancer name: AWS-HA-TEST
Scheme: internet-facing
Port Configuration: 80 (HTTP) forwarding to 80 (HTTP)
443 (HTTPS) forwarding to 443 (HTTPS)

▼ Configure Health Check

Ping Target: HTTP:8444/
Timeout: 2 seconds
Interval: 5 seconds
Unhealthy Threshold: 2
Healthy Threshold: 2

▼ Add EC2 Instances

Cross-Zone Load Balancing: Enabled
Connection Draining: Disabled
Instances: i-dbb6d65 (QA-AC-7.1-29-29-VLM_AWS_5-MASTER), i-d8ba6d66 (QA-AC-7.1-29-29-VLM_AWS_5-SLAVE)

▼ VPC Information

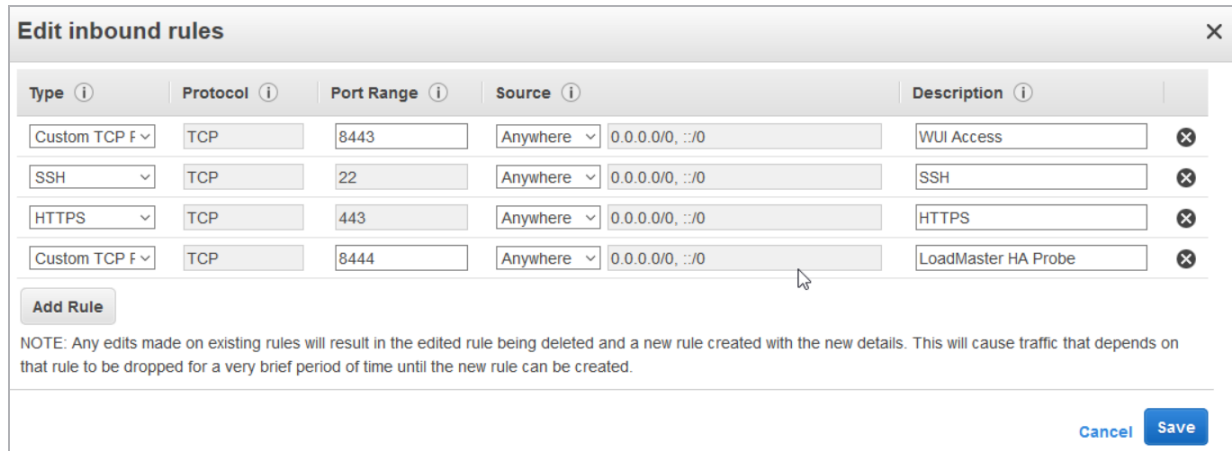
VPC: vpc-6b475509 (Default)
Subnets: subnet-00ec9f28, subnet-e6764792, subnet-46124100

▼ Security Groups

Security Groups: sg-a7db16c2

16. Review the settings and click **Create** to create the Classic Load Balancer (ELB).

5.2 Modify the Inbound Security Rule



Type	Protocol	Port Range	Source	Description
Custom TCP F	TCP	8443	Anywhere	WUI Access
SSH	TCP	22	Anywhere	SSH
HTTPS	TCP	443	Anywhere	HTTPS
Custom TCP F	TCP	8444	Anywhere	LoadMaster HA Probe

NOTE: Any edits made on existing rules will result in the edited rule being deleted and a new rule created with the new details. This will cause traffic that depends on that rule to be dropped for a very brief period of time until the new rule can be created.

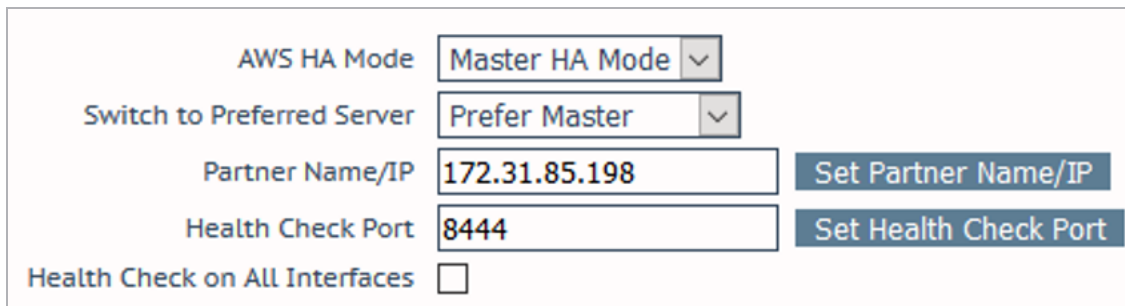
Kemp recommends limiting the source address whenever possible to further secure access to management interfaces, and so on.

Modify the inbound security rule for the individual instance to allow port 8444 traffic. If this is not done, the health check in the load balancer group will not pass and both instances are shown as OutOfService.

5.3 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

Health Check Port: 8444

Health Check on All Interfaces: ☐

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**.

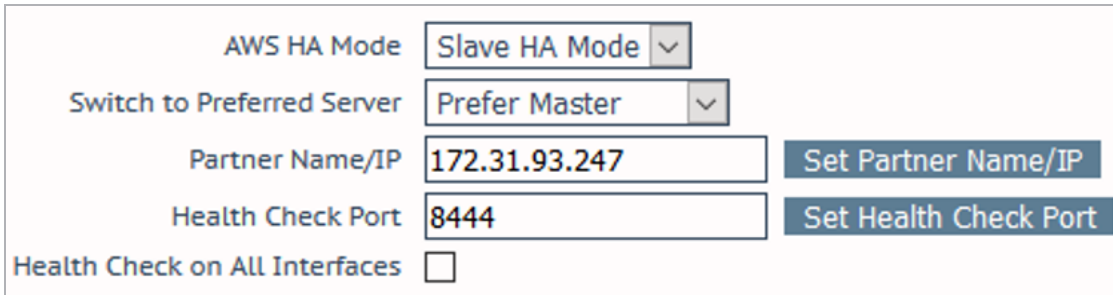
6. Enter the health check port selected earlier in the **Health Check Port** text box and click **Set Health Check Port**.

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**.



The screenshot shows a configuration form for AWS HA Parameters. It includes the following fields and controls:

- AWS HA Mode:** A dropdown menu currently showing "Slave HA Mode".
- Switch to Preferred Server:** A dropdown menu currently showing "Prefer Master".
- Partner Name/IP:** A text input field containing "172.31.93.247". To its right is a blue button labeled "Set Partner Name/IP".
- Health Check Port:** A text input field containing "8444". To its right is a blue button labeled "Set Health Check Port".
- Health Check on All Interfaces:** A checkbox that is currently unchecked.

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**.

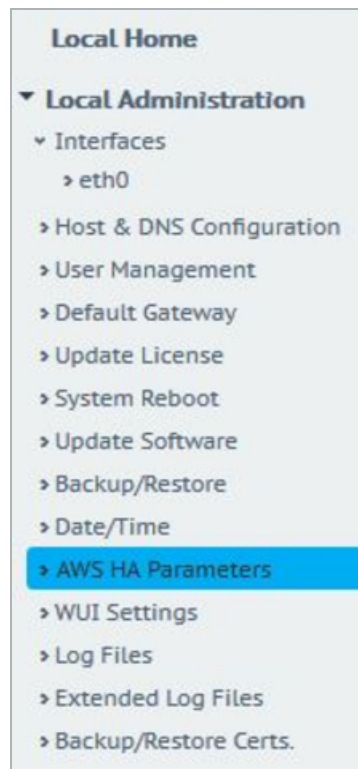
The **Health Check Port** must be the same on both the master and slave units in order for HA to function correctly.

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.

5.3.1 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

6 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>.

6.1 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.

6.2 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.

6 LoadMaster Firmware Upgrades/Downgrades

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.

7 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 Classic Load Balancer (ELB) 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, Feature Description

High Availability (HA), Feature Description

Last Updated Date

This document was last updated on 27 July 2023.