
Elastic Load Balancing

Getting Started Guide

API Version 2011-11-15



Elastic Load Balancing: Getting Started Guide

Copyright © 2011 Amazon Web Services LLC or its affiliates. All rights reserved.


Table of Contents

Get Started with Elastic Load Balancing	1
Define Load Balancer	2
Configure Health Check	5
Add EC2 Instances	7
Delete Your Load Balancer	10
Where Do I Go from Here?	12
Please Provide Feedback	15
About This Guide	16

Get Started with Elastic Load Balancing

Elastic Load Balancing is Amazon's web service that improves an application's availability by distributing incoming traffic between two or more Amazon Elastic Compute Cloud (Amazon EC2) instances.

You can get familiar with Elastic Load Balancing by stepping through the tasks shown in the following diagram. You'll primarily use the Amazon Web Services (AWS) Management Console, a point-and-click web-based interface.




This guide walks you through creating your first load balancer with Elastic Load Balancing.



Important

This guide assumes that you have an AWS account and have already created two or more Amazon EC2 instances with active HTTP servers. For more information, go to the [Amazon Elastic Compute Cloud User Guide](#).

Define Load Balancer



If you have an AWS account, have signed up for Amazon EC2, and have created at least two Amazon EC2 instances with active HTTP servers, you're ready to define a load balancer using the AWS Management Console.



Important

The load balancer you're about to launch will be live (and not running in a sandbox). This means that you will incur the standard Elastic Load Balancing usage fees for the load balancer until you terminate it. The total charges will be minimal (typically less than a dollar) if you complete the exercise described here in one sitting and delete your load balancer when you are finished. For more information about Elastic Load Balancing usage rates, go to the [Elastic Load Balancing product page](#).

To define a load balancer

1. Start the **Create Load Balancer** wizard:
 - a. From [the AWS Management Console](#), select **Amazon EC2** from the drop-down list box, then click **Sign in to the AWS Console**, and log in with the email address and password you used when signing up for Amazon EC2.

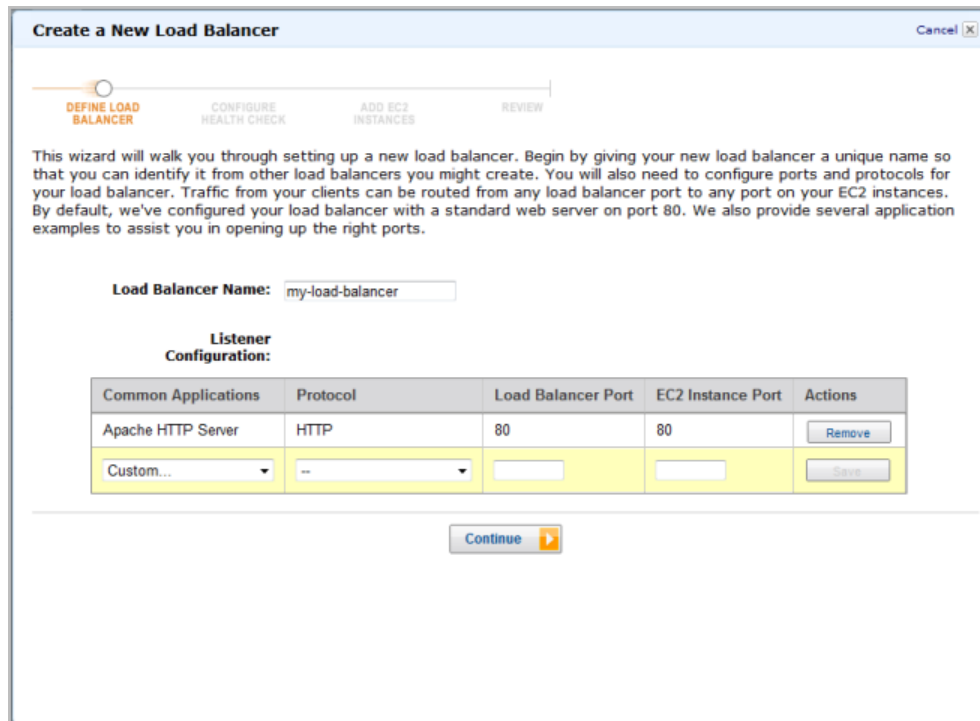


Tip

If you pause for a long period of time during this procedure, the AWS Management Console automatically logs you out. To stay logged in while you work through this tutorial, click **Settings** in the top right corner of the console window and clear the **Sign out on inactivity** check box.

- b. Click **Load Balancers** in the **Navigation** pane.
- The console displays the **Load Balancers** pane.
- c. From the **Load Balancers** pane, click **Create Load Balancers**.

The **DEFINE LOAD BALANCER** page of the **Create a New Load Balancer** wizard opens.



2. On the **DEFINE LOAD BALANCER** page, enter a name for your load balancer (e.g., myLB).

Create a New Load Balancer Cancel

DEFINE LOAD BALANCER CONFIGURE HEALTH CHECK ADD EC2 INSTANCES REVIEW

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 will also need to configure ports and protocols for your load balancer. Traffic from your clients can be routed from any load balancer port to any port on your EC2 instances. By default, we've configured your load balancer with a standard web server on port 80. We also provide several application examples to assist you in opening up the right ports.

Load Balancer Name:

Listener Configuration:

Common Applications	Protocol	Load Balancer Port	EC2 Instance Port	Actions
Apache HTTP Server	HTTP	80	80	<input type="button" value="Remove"/>
Custom.	--			

3. Leave the **Listener Configuration** set to the default value for this example.




Important

The default settings require that your Amazon EC2 HTTP servers are active and accepting requests on port 80.

4. Click **Continue**.

The **CONFIGURE HEALTH CHECK** page of the **Create a New Load Balancer** wizard opens.

Configure Health Check



Now that you have defined the load balancer's name and have used the defaults for the listener ports, you're ready to configure the health check.

Elastic Load Balancing routinely checks the health of each load-balanced Amazon EC2 instance. If Elastic Load Balancing finds an unhealthy instance, it stops sending traffic to the instance and reroutes traffic to healthy instances.

To configure the health check

1. On the **CONFIGURE HEALTH CHECK** page of the **Create a New Load Balancer** wizard, set the **Configuration Options** for your load balancer.

Create a New Load Balancer Cancel

DEFINE LOAD BALANCER **CONFIGURE HEALTH CHECK** ADD EC2 INSTANCES REVIEW

Your load balancer will automatically perform health checks on your EC2 instances and only route traffic to instances that pass the health check. If an instance fails the health check, it is automatically removed from the load balancer. Customize the health check to meet your specific needs.

Configuration Options:

Ping Protocol: HTTP

Ping Port: 80

Ping Path: /index.html

Advanced Options:

Response Timeout: 5 Seconds

Time to wait when receiving a response from the health check (2 sec - 60 sec)

- a. Leave **Ping Protocol** set to its default value of HTTP.
- b. Leave **Ping Port** set to its default value of 80.

Elastic Load Balancing uses the **Ping Port** to send health check queries to your Amazon EC2 instances.



Important

Your Amazon EC2 instances must accept incoming traffic on this port. This guide assumes that each of your instances has a working HTTP server that accepts incoming traffic on port 80.

- c. In the **Ping Path** field, replace the default value with a single forward slash ("/").


The screenshot shows the 'Create a New Load Balancer' wizard in the 'CONFIGURE HEALTH CHECK' step. The progress bar at the top indicates the current step. Below the progress bar, there is a descriptive paragraph about health checks. Under 'Configuration Options', the 'Ping Protocol' is set to 'HTTP', 'Ping Port' is '80', and 'Ping Path' is '/'. A red arrow points to the 'Ping Path' field.

Elastic Load Balancing sends health check queries to the path you specify in **Ping Path**. This example uses a single forward slash so that Elastic Load Balancing sends the query to your HTTP server's default home page, whether that default page is named `index.html`, `default.html`, or a different name.

2. Leave the **Advanced Options** set to their default values.
3. Click **Continue**.

The **ADD INSTANCES** page of the **Create a New Load Balancer** wizard opens.

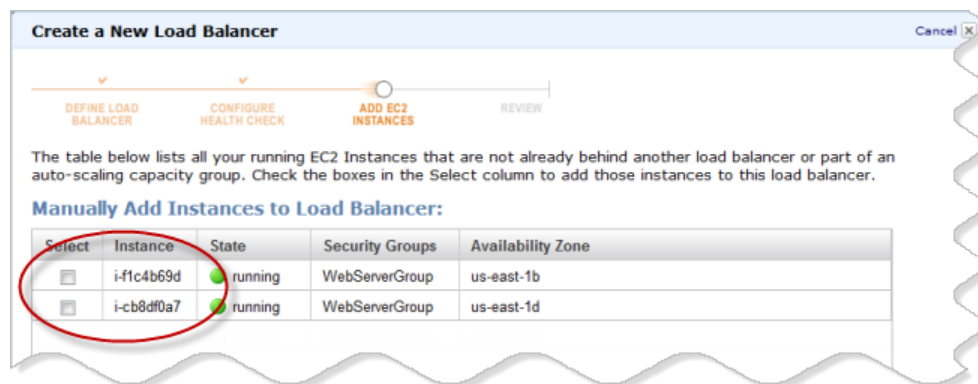
Add EC2 Instances



Now that you've configured the health check, you're ready to add Amazon EC2 instances to your load balancer.

To add Amazon EC2 instances

1. On the **ADD EC2 INSTANCES** page of the **Create a New Load Balancer** wizard, select the check box next to each instance you want to add.



Important

This guide assumes that you have already created two or more Amazon EC2 instances with active HTTP servers.

2. Click **Continue**.

The **REVIEW** page of the **Create a New Load Balancer** wizard opens.

- Review your settings. You can make changes to the settings by clicking the edit link for a specific step in the process.



Important

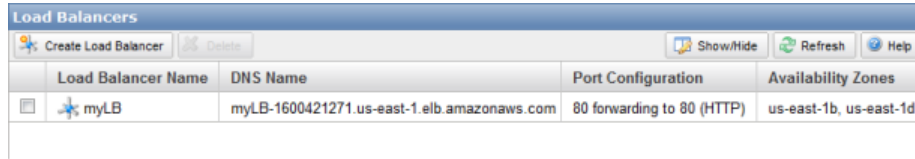
After you create a load balancer, you can modify any of the settings, except for **Load Balancer Name** and **Port Configuration**. To rename a load balancer or change its port configuration, create a replacement load balancer.

- Click **Create**.

A confirmation window opens.

- Click **Close**.

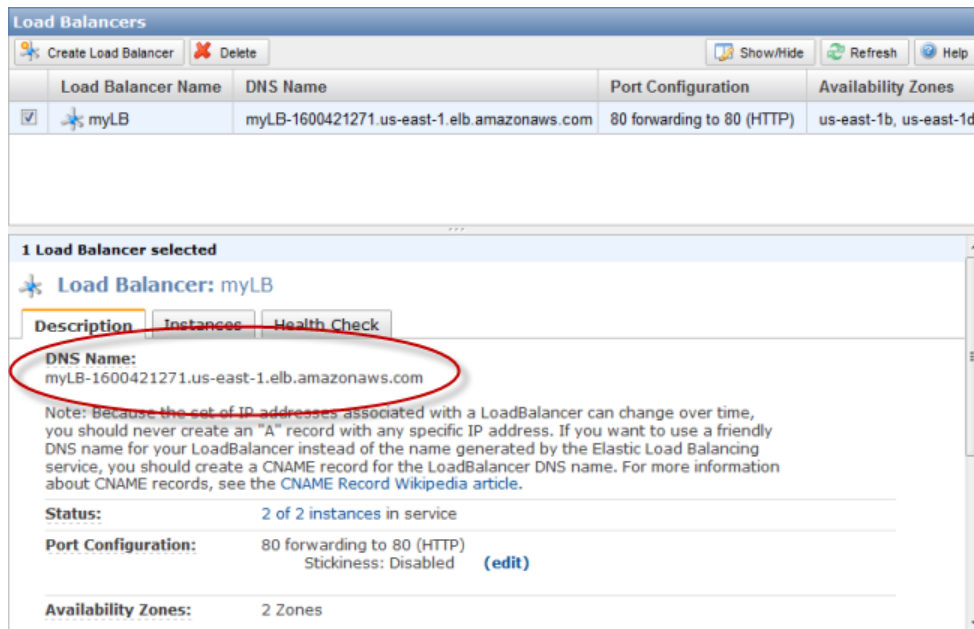
The confirmation window closes, returning you to the **Load Balancers** page. Your new load balancer now appears in the list.



Load Balancer Name	DNS Name	Port Configuration	Availability Zones
myLB	myLB-1600421271.us-east-1.elb.amazonaws.com	80 forwarding to 80 (HTTP)	us-east-1b, us-east-1d

6. Select the check box next to your load balancer.

A set of tabs opens with details about your new load balancer.



1 Load Balancer selected

Load Balancer: myLB

Description Instances Health Check

DNS Name:
myLB-1600421271.us-east-1.elb.amazonaws.com

Note: Because the set of IP addresses associated with a LoadBalancer can change over time, you should never create an "A" record with any specific IP address. If you want to use a friendly DNS name for your LoadBalancer instead of the name generated by the Elastic Load Balancing service, you should create a CNAME record for the LoadBalancer DNS name. For more information about CNAME records, see the [CNAME Record Wikipedia article](#).

Status: 2 of 2 instances in service


Port Configuration: 80 forwarding to 80 (HTTP)
Stickiness: Disabled (edit)

Availability Zones: 2 Zones

To test your load balancer, copy the **DNS Name** value that is listed in the **Description** tab and paste it into the address field of an Internet-connected web browser. If your load balancer is working, you will see the default page of your HTTP server.

Normally you'd continue using the load balancer. However, for the purposes of this tutorial, we're going to show you how to terminate the load balancer immediately.


Delete Your Load Balancer



As soon as your load balancer becomes available, you're billed for each hour or partial hour that you keep the load balancer running. After you've decided that you no longer need the load balancer, you can delete it.


To delete your load balancer

1. In the [AWS Management Console](#), on the **Amazon EC2** tab, click **Load Balancers** in the **Navigation** pane to open the **Load Balancers** page.



2. Select the check box next to the load balancer you want to delete, and then click **Delete** at the top of the **Load Balancers** page.

The **Delete Load Balancer** window appears.



3. Click **Yes, Delete**.

Elastic Load Balancing deletes the load balancer. As soon as the load balancer is deleted, you stop incurring charges for that load balancer.



Note

Even after you delete a load balancer, the Amazon EC2 instances associated with the load balancer continue to run. You will continue to incur charges on the Amazon EC2 instances while they are running.

Congratulations! You successfully created a load balancer, added Amazon EC2 instances to it, and deleted the load balancer. For more information about Elastic Load Balancing and how to continue using its rich feature set, see [Where Do I Go from Here? \(p. 12\)](#).

Your input is important to us. Help make our documentation helpful and easy to use. Please take a minute to provide feedback on your getting started experience with Elastic Load Balancing. To begin the survey, see [Please Provide Feedback \(p. 15\)](#). Thank you.

Where Do I Go from Here?

Topics

- [AWS Account and Security Credentials](#) (p. 12)
- [Other Ways to Access Elastic Load Balancing](#) (p. 12)
- [Elastic Load Balancing Resources](#) (p. 13)

Elastic Load Balancing is a rich service offering many things we haven't covered in this guide, such as session stickiness, security features, and more. This section provides links to additional resources, which will help you deepen your understanding and use of Elastic Load Balancing.

AWS Account and Security Credentials

So far you signed up for the service, got an AWS account and security credentials, and then completed a short exercise covering the essential product functions. Now that you're finished with the exercise, we recommend that you check with an administrator or coworker in your organization to determine if he or she already has an AWS account and security credentials for you to use in future interactions with AWS.

If you're an account owner or administrator and want to know more about AWS Identity and Access Management, go to the product description at <http://aws.amazon.com/iam> or to the technical documentation at [Using AWS Identity and Access Management](#).

Other Ways to Access Elastic Load Balancing

This guide has shown you how to create and delete a load balancer using the AWS Management Console. You can continue using Elastic Load Balancing through the console, or try one of the other interfaces.

Continue Using the Console

The AWS Management Console includes many other functions besides creating and deleting load balancers. To learn more about how to use Elastic Load Balancing through the console, consult the online Help on the console or go to the [Elastic Load Balancing Developer Guide](#).

Use the Command Line Interface

For information on using the Elastic Load Balancing command line interface (CLI) API, go to the [Elastic Load Balancing API Tools](#). These command line tools are a fast way to execute all of the Elastic Load Balancing functions without coding to the Query API or using a library. For a convenient listing of all Elastic Load Balancing commands, go to the [Elastic Load Balancing Quick Reference Card](#).

Use an Existing Library

If you prefer to use Elastic Load Balancing through a programmatic interface, there are libraries and resources available for the following languages:

- [Java](#)
- [PHP](#)
- [Python](#)
- [Ruby](#)
- [Windows and .NET](#)

For libraries and sample code in all languages, go to the [Amazon EC2 Sample Code & Libraries](#).

Code Directly to the Web Service API

If you want to write code directly to the Elastic Load Balancing web service API, go to the [Elastic Load Balancing Developer Guide](#). The guide describes how to create and authenticate API requests, and how to use Elastic Load Balancing through the API. For a complete description of all the API actions, go to the [Elastic Load Balancing API Reference](#).

Elastic Load Balancing Resources

The following table lists related resources that you'll find useful as you work with this service.

Resource	Description
Elastic Load Balancing Developer Guide	The Developer Guide provides conceptual information about Elastic Load Balancing and describes how to use Elastic Load Balancing features using the command line tools and Query APIs.
Elastic Load Balancing API Reference	The API Reference contains a comprehensive description of all Elastic Load Balancing Query APIs and data types.
Elastic Load Balancing Command Line Interface Quick Reference Card	The Quick Reference Card contains a comprehensive listing of all the command line tools and their options.
Elastic Load Balancing Technical FAQ	The FAQ covers the top questions developers have asked about this product.
Release Notes	The release notes give a high-level overview of the current release. They specifically note any new features, corrections, and known issues.

Elastic Load Balancing Getting Started Guide
Elastic Load Balancing Resources

Resource	Description
Discussion Forums	A community-based forum for developers to discuss technical questions related to Amazon Web Services.
AWS Support Center	The home page for AWS Technical Support, including access to our Developer Forums, Technical FAQs, Service Status page, and Premium Support.
Elastic Load Balancing product information	The primary web page for information about Elastic Load Balancing.
Contact Us	A central contact point for inquiries concerning AWS billing, account, events, abuse, etc.
Conditions of Use	Detailed information about the copyright and trademark usage at Amazon.com and other topics.

Please Provide Feedback

Your input is important to help make our documentation helpful and easy to use. Please tell us about your experience getting started with Elastic Load Balancing by completing our [Getting Started Survey](#).

Thank you.

About This Guide

This is the *Elastic Load Balancing Getting Started Guide*. It was last updated on November 21, 2011.