



Application Load Balancers

Elastic Load Balancing



Elastic Load Balancing: Application Load Balancers

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

What is an Application Load Balancer?	1
Application Load Balancer components	1
Application Load Balancer overview	2
Benefits of migrating from a Classic Load Balancer	3
Related services	4
Pricing	5
Getting started	6
Before you begin	6
Step 1: Configure your target group	6
Step 2: Choose a load balancer type	7
Step 3: Configure your load balancer and listener	7
Step 4: Test your load balancer	8
Step 5: (Optional) Delete your load balancer	9
Tutorial: Create an Application Load Balancer using the AWS CLI	10
Before you begin	10
Create your load balancer	10
Add an HTTPS listener	12
Add path-based routing	13
Delete your load balancer	13
Load balancers	14
Subnets for your load balancer	15
Availability Zone subnets	15
Local Zone subnets	16
Outpost subnets	16
Load balancer security groups	17
Load balancer state	18
Load balancer attributes	18
IP address type	21
Load balancer resource map	22
Resource map components	22
Load balancer connections	23
Connection idle timeout	23
HTTP client keepalive duration	24
Cross-zone load balancing	25

Deletion protection	26
Desync mitigation mode	27
Host header preservation	28
AWS WAF	31
Create a load balancer	32
Step 1: Configure a target group	6
Step 2: Register targets	34
Step 3: Configure a load balancer and a listener	34
Step 4: Test the load balancer	8
Update Availability Zones	39
Update security groups	39
Recommended rules	40
Update the associated security groups	42
Update the address type	43
Update tags	44
Delete a load balancer	45
Zonal shift	45
Start a zonal shift	47
Update a zonal shift	47
Cancel a zonal shift	48
Listeners and rules	49
Listener configuration	49
Listener rules	50
Default rules	50
Rule priority	51
Rule actions	51
Rule conditions	51
Rule action types	51
Fixed-response actions	52
Forward actions	53
Redirect actions	55
Rule condition types	59
HTTP header conditions	60
HTTP request method conditions	61
Host conditions	61
Path conditions	62

Query string conditions	64
Source IP address conditions	64
Create an HTTP listener	65
Prerequisites	65
Add an HTTP listener	65
Create an HTTPS listener	66
SSL certificates	67
Security policies	70
Add an HTTPS listener	92
Update listener rules	94
Requirements	95
Add a rule	95
Edit a rule	97
Reorder rules	98
Delete a rule	99
Update an HTTPS listener	100
Replace the default certificate	100
Add certificates to the certificate list	101
Remove certificates from the certificate list	101
Update the security policy	102
Use mutual TLS authentication	103
Before you begin	104
HTTP headers	107
Configuring mutual TLS	108
Connection logs	114
Authenticate users	114
Prepare to use an OIDC-compliant IdP	114
Prepare to use Amazon Cognito	115
Prepare to use Amazon CloudFront	117
Configure user authentication	117
Authentication flow	120
User claims encoding and signature verification	122
Timeout	125
Authentication logout	126
X-forwarded headers	127
X-Forwarded-For	127

X-Forwarded-Proto	131
X-Forwarded-Port	131
Update tags	132
Update listener tags	132
Update rule tags	133
Delete a listener	134
Target groups	135
Routing configuration	136
Target type	136
IP address type	138
Protocol version	138
Registered targets	140
Target group attributes	141
Routing algorithms	143
Modify the routing algorithm of a target group	144
Automatic Target Weights (ATW)	144
Anomaly detection	145
Anomaly mitigation	146
Deregistration delay	148
Slow start mode	149
Create a target group	150
Configure health checks	152
Health check settings	152
Target health status	154
Health check reason codes	156
Check the health of your targets	157
Modify the health check settings of a target group	158
Cross-zone load balancing	158
Turn off cross-zone load balancing	159
Turn on cross-zone load balancing	160
Target group health	161
Unhealthy state actions	161
Requirements and considerations	162
Monitoring	162
Example	162
Modify target group health settings	164

Using Route 53 DNS failover for your load balancer	165
Register targets	166
Target security groups	166
Shared subnets	167
Register or deregister targets	167
Sticky sessions	170
Duration-based stickiness	171
Application-based stickiness	173
Lambda functions as targets	176
Prepare the Lambda function	177
Create a target group for the Lambda function	169
Receive events from the load balancer	179
Respond to the load balancer	180
Multi-value headers	181
Enable health checks	183
Deregister the Lambda function	185
Update tags	185
Delete a target group	186
Monitor your load balancers	188
CloudWatch metrics	189
Application Load Balancer metrics	189
Metric dimensions for Application Load Balancers	209
Statistics for Application Load Balancer metrics	209
View CloudWatch metrics for your load balancer	211
Access logs	213
Access log files	214
Access log entries	215
Example log entries	230
Processing access log files	232
Enable access logs	232
Disable access logs	243
Connection logs	243
Connection log files	244
Connection log entries	246
Example log entries	249
Processing connection log files	250

Enable connection logs	250
Disable connection logs	256
Request tracing	257
Syntax	257
Limitations	258
CloudTrail logs	258
Elastic Load Balancing information in CloudTrail	259
Understanding Elastic Load Balancing log file entries	260
Troubleshoot your load balancers	263
A registered target is not in service	263
Clients cannot connect to an internet-facing load balancer	265
Requests sent to a custom domain aren't received by the load balancer	265
HTTPS requests sent to the load balancer return "NET::ERR_CERT_COMMON_NAME_INVALID"	266
Load balancer shows elevated processing times	266
The load balancer sends a response code of 000	266
The load balancer generates an HTTP error	266
HTTP 400: Bad request	267
HTTP 401: Unauthorized	267
HTTP 403: Forbidden	268
HTTP 405: Method not allowed	268
HTTP 408: Request timeout	268
HTTP 413: Payload too large	268
HTTP 414: URI too long	268
HTTP 460	269
HTTP 463	269
HTTP 464	269
HTTP 500: Internal server error	269
HTTP 501: Not implemented	270
HTTP 502: Bad gateway	270
HTTP 503: Service unavailable	271
HTTP 504: Gateway timeout	271
HTTP 505: Version not supported	271
HTTP 507: Insufficient Storage	271
HTTP 561: Unauthorized	271
A target generates an HTTP error	272

An AWS Certificate Manager certificate is not available for use	272
Multi-Line headers are not supported	272
Troubleshoot unhealthy targets using the resource map	272
Quotas	275
Document history	279

What is an Application Load Balancer?

Elastic Load Balancing automatically distributes your incoming traffic across multiple targets, such as EC2 instances, containers, and IP addresses, in one or more Availability Zones. It monitors the health of its registered targets, and routes traffic only to the healthy targets. Elastic Load Balancing scales your load balancer as your incoming traffic changes over time. It can automatically scale to the vast majority of workloads.

Elastic Load Balancing supports the following load balancers: Application Load Balancers, Network Load Balancers, Gateway Load Balancers, and Classic Load Balancers. You can select the type of load balancer that best suits your needs. This guide discusses Application Load Balancers. For more information about the other load balancers, see the [User Guide for Network Load Balancers](#), the [User Guide for Gateway Load Balancers](#), and the [User Guide for Classic Load Balancers](#).


Application Load Balancer components

A *load balancer* serves as the single point of contact for clients. The load balancer distributes incoming application traffic across multiple targets, such as EC2 instances, in multiple Availability Zones. This increases the availability of your application. You add one or more listeners to your load balancer.

A *listener* checks for connection requests from clients, using the protocol and port that you configure. The rules that you define for a listener determine how the load balancer routes requests to its registered targets. Each rule consists of a priority, one or more actions, and one or more conditions. When the conditions for a rule are met, then its actions are performed. You must define a default rule for each listener, and you can optionally define additional rules.

Each *target group* routes requests to one or more registered targets, such as EC2 instances, using the protocol and port number that you specify. You can register a target with multiple target groups. You can configure health checks on a per target group basis. Health checks are performed on all targets registered to a target group that is specified in a listener rule for your load balancer.

The following diagram illustrates the basic components. Notice that each listener contains a default rule, and one listener contains another rule that routes requests to a different target group. One target is registered with two target groups.



For more information, see the following documentation:

- [Load balancers](#)
- [Listeners](#)
- [Target groups](#)

Application Load Balancer overview

An Application Load Balancer functions at the application layer, the seventh layer of the Open Systems Interconnection (OSI) model. After the load balancer receives a request, it evaluates the listener rules in priority order to determine which rule to apply, and then selects a target from the target group for the rule action. You can configure listener rules to route requests to different target groups based on the content of the application traffic. Routing is performed independently for each target group, even when a target is registered with multiple target groups. You can configure the routing algorithm used at the target group level. The default routing algorithm is round robin; alternatively, you can specify the least outstanding requests routing algorithm.

You can add and remove targets from your load balancer as your needs change, without disrupting the overall flow of requests to your application. Elastic Load Balancing scales your load balancer as traffic to your application changes over time. Elastic Load Balancing can scale to the vast majority of workloads automatically.

You can configure health checks, which are used to monitor the health of the registered targets so that the load balancer can send requests only to the healthy targets.

For more information, see [How Elastic Load Balancing works](#) in the *Elastic Load Balancing User Guide*.

Benefits of migrating from a Classic Load Balancer

Using an Application Load Balancer instead of a Classic Load Balancer has the following benefits:

- Support for [Path conditions](#). You can configure rules for your listener that forward requests based on the URL in the request. This enables you to structure your application as smaller services, and route requests to the correct service based on the content of the URL.
- Support for [Host conditions](#). You can configure rules for your listener that forward requests based on the host field in the HTTP header. This enables you to route requests to multiple domains using a single load balancer.
- Support for routing based on fields in the request, such as [HTTP header conditions](#) and methods, query parameters, and source IP addresses.
- Support for routing requests to multiple applications on a single EC2 instance. You can register an instance or IP address with multiple target groups, each on a different port.
- Support for redirecting requests from one URL to another.
- Support for returning a custom HTTP response.
- Support for registering targets by IP address, including targets outside the VPC for the load balancer.
- Support for registering Lambda functions as targets.
- Support for the load balancer to authenticate users of your applications through their corporate or social identities before routing requests.
- Support for containerized applications. Amazon Elastic Container Service (Amazon ECS) can select an unused port when scheduling a task and register the task with a target group using this port. This enables you to make efficient use of your clusters.
- Support for monitoring the health of each service independently, as health checks are defined at the target group level and many CloudWatch metrics are reported at the target group level. Attaching a target group to an Auto Scaling group enables you to scale each service dynamically based on demand.
- Access logs contain additional information and are stored in compressed format.
- Improved load balancer performance.

For more information about the features supported by each load balancer type, see [Product comparisons](#) for Elastic Load Balancing.

Related services

Elastic Load Balancing works with the following services to improve the availability and scalability of your applications.

- **Amazon EC2** — Virtual servers that run your applications in the cloud. You can configure your load balancer to route traffic to your EC2 instances.
- **Amazon EC2 Auto Scaling** — Ensures that you are running your desired number of instances, even if an instance fails, and enables you to automatically increase or decrease the number of instances as the demand on your instances changes. If you enable Auto Scaling with Elastic Load Balancing, instances that are launched by Auto Scaling are automatically registered with the target group, and instances that are terminated by Auto Scaling are automatically de-registered from the target group.
- **AWS Certificate Manager** — When you create an HTTPS listener, you can specify certificates provided by ACM. The load balancer uses certificates to terminate connections and decrypt requests from clients. For more information, see [SSL certificates](#).
- **Amazon CloudWatch** — Enables you to monitor your load balancer and take action as needed. For more information, see [CloudWatch metrics for your Application Load Balancer](#).
- **Amazon ECS** — Enables you to run, stop, and manage Docker containers on a cluster of EC2 instances. You can configure your load balancer to route traffic to your containers. For more information, see [Service load balancing](#) in the *Amazon Elastic Container Service Developer Guide*.
- **AWS Global Accelerator** — Improves the availability and performance of your application. Use an accelerator to distribute traffic across multiple load balancers in one or more AWS Regions. For more information, see the [AWS Global Accelerator Developer Guide](#).
- **Route 53** — Provides a reliable and cost-effective way to route visitors to websites by translating domain names (such as `www.example.com`) into the numeric IP addresses (such as `192.0.2.1`) that computers use to connect to each other. AWS assigns URLs to your resources, such as load balancers. However, you might want a URL that is easy for users to remember. For example, you can map your domain name to a load balancer. For more information, see [Routing traffic to an ELB load balancer](#) in the *Amazon Route 53 Developer Guide*.
- **AWS WAF** — You can use AWS WAF with your Application Load Balancer to allow or block requests based on the rules in a web access control list (web ACL). For more information, see [Application Load Balancers and AWS WAF](#).

To view information about services that are integrated with your load balancer, select your load balancer in the AWS Management Console and choose the **Integrated services** tab.

Pricing

With your load balancer, you pay only for what you use. For more information, see [Elastic Load Balancing pricing](#).

Getting started with Application Load Balancers

This tutorial provides a hands-on introduction to Application Load Balancers through the AWS Management Console, a web-based interface. To create your first Application Load Balancer, complete the following steps.

Tasks

- [Before you begin](#)
- [Step 1: Configure your target group](#)
- [Step 2: Choose a load balancer type](#)
- [Step 3: Configure your load balancer and listener](#)
- [Step 4: Test your load balancer](#)
- [Step 5: \(Optional\) Delete your load balancer](#)

For demos of common load balancer configurations, see [Elastic Load Balancing demos](#).

Before you begin

- Decide which two Availability Zones you will use for your EC2 instances. Configure your virtual private cloud (VPC) with at least one public subnet in each of these Availability Zones. These public subnets are used to configure the load balancer. You can launch your EC2 instances in other subnets of these Availability Zones instead.
- Launch at least one EC2 instance in each Availability Zone. Be sure to install a web server, such as Apache or Internet Information Services (IIS), on each EC2 instance. Ensure that the security groups for these instances allow HTTP access on port 80.

Step 1: Configure your target group

Create a target group, which is used in request routing. The default rule for your listener routes requests to the registered targets in this target group. The load balancer checks the health of targets in this target group using the health check settings defined for the target group.

To configure your target group using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.

2. In the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose **Create target group**.
4. Under **Basic configuration**, keep the **Target type** as instance.
5. For **Target group name**, enter a name for the new target group.
6. Keep the default protocol (**HTTP**) and port (**80**).
7. Select the **VPC** containing your instances. Keep the protocol version as **HTTP1**.
8. For **Health checks**, keep the default settings.
9. Choose **Next**.
10. On the **Register targets** page, complete the following steps. This is an optional step for creating the load balancer. However, you must register this target if you want to test your load balancer and ensure that it is routing traffic to this target.
 - a. For **Available instances**, select one or more instances.
 - b. Keep the default port 80, and choose **Include as pending below**.
11. Choose **Create target group**.

Step 2: Choose a load balancer type

Elastic Load Balancing supports different types of load balancers. For this tutorial, you create an Application Load Balancer.

To create an Application Load Balancer using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation bar, choose a Region for your load balancer. Be sure to choose the same Region that you used for your EC2 instances.
3. In the navigation pane, under **Load Balancing**, choose **Load Balancers**.
4. Choose **Create Load Balancer**.
5. For **Application Load Balancer**, choose **Create**.

Step 3: Configure your load balancer and listener

To create an Application Load Balancer, you must first provide basic configuration information for your load balancer, such as a name, scheme, and IP address type. Then, you provide information

about your network, and one or more listeners. A listener is a process that checks for connection requests. It is configured with a protocol and a port for connections from clients to the load balancer. For more information about supported protocols and ports, see [Listener configuration](#).

To configure your load balancer and listener

1. For **Load balancer name**, enter a name for your load balancer. For example, my-a1b.
2. For **Scheme** and **IP address type**, keep the default values.
3. For **Network mapping**, select the VPC that you used for your EC2 instances. Select at least two Availability Zones and one subnet per zone. For each Availability Zone that you used to launch your EC2 instances, select the Availability Zone and then select one public subnet for that Availability Zone.
4. For **Security groups**, we select the default security group for the VPC that you selected in the previous step. You can choose a different security group instead. The security group must include rules that allow the load balancer to communicate with registered targets on both the listener port and the health check port. For more information, see [Security group rules](#).
5. For **Listeners and routing**, keep the default protocol and port, and select your target group from the list. This configures a listener that accepts HTTP traffic on port 80 and forwards traffic to the selected target group by default. For this tutorial, you are not creating an HTTPS listener.
6. For **Default action**, select the target group that you created and registered in Step 1: Configure your target group.
7. (Optional) Add a tag to categorize your load balancer. Tag keys must be unique for each load balancer. Allowed characters are letters, spaces, numbers (in UTF-8), and the following special characters: + - = . _ : / @. Do not use leading or trailing spaces. Tag values are case-sensitive.
8. Review your configuration, and choose **Create load balancer**. A few default attributes are applied to your load balancer during creation. You can view and edit them after creating the load balancer. For more information, see [Load balancer attributes](#).

Step 4: Test your load balancer

After creating the load balancer, verify that it's sending traffic to your EC2 instances.

To test your load balancer

1. After you are notified that your load balancer was created successfully, choose **Close**.

2. In the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Select the newly created target group.
4. Choose **Targets** and verify that your instances are ready. If the status of an instance is `initial`, it's probably because the instance is still in the process of being registered, or it has not passed the minimum number of health checks to be considered healthy. After the status of at least one instance is `healthy`, you can test your load balancer.
5. In the navigation pane, under **Load Balancing**, choose **Load Balancers**.
6. Select the newly created load balancer.
7. Choose **Description** and copy the DNS name of the load balancer (for example, `my-load-balancer-1234567890abcdef.elb.us-east-2.amazonaws.com`). Paste the DNS name into the address field of an internet-connected web browser. If everything is working, the browser displays the default page of your server.
8. (Optional) To define additional listener rules, see [Add a rule](#).

Step 5: (Optional) Delete your load balancer

As soon as your load balancer becomes available, you are billed for each hour or partial hour that you keep it running. When you no longer need a load balancer, you can delete it. As soon as the load balancer is deleted, you stop incurring charges for it. Note that deleting a load balancer does not affect the targets registered with the load balancer. For example, your EC2 instances continue to run after deleting the load balancer created in this guide.

To delete your load balancer using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, under **Load Balancing**, choose **Load Balancers**.
3. Select the checkbox for the load balancer, choose **Actions**, then choose **Delete**.
4. When prompted for confirmation, choose **Yes, Delete**.

Tutorial: Create an Application Load Balancer using the AWS CLI

This tutorial provides a hands-on introduction to Application Load Balancers through the AWS CLI.

Before you begin

- Use the following command to verify that you are running a version of the AWS CLI that supports Application Load Balancers.

```
aws elbv2 help
```

If you get an error message that `elbv2` is not a valid choice, update your AWS CLI. For more information, see [Installing the AWS Command Line Interface](#) in the *AWS Command Line Interface User Guide*.

- Launch your EC2 instances in a virtual private cloud (VPC). Ensure that the security groups for these instances allow access on the listener port and the health check port. For more information, see [Target security groups](#).
- Decide if you will create an IPv4 or dualstack load balancer. Use IPv4 if you want clients to communicate with the load balancer using IPv4 addresses only. Use dualstack if you want clients to communicate with the load balancer using IPv4 and IPv6 addresses. You can also use dualstack to communicate with backend targets, such as IPv6 applications or dualstack subnets, using IPv6.
- Be sure to install a web server, such as Apache or Internet Information Services (IIS), on each EC2 instance. Ensure that the security groups for these instances allow HTTP access on port 80.

Create your load balancer

To create your first load balancer, complete the following steps.

To create a load balancer

1. Use the [create-load-balancer](#) command to create a load balancer. You must specify two subnets that are not from the same Availability Zone.

```
aws elbv2 create-load-balancer --name my-load-balancer \  
--subnets subnet-0e3f5cac72EXAMPLE subnet-081ec835f3EXAMPLE --security-groups  
sg-07e8ffd50fEXAMPLE
```

Use the [create-load-balancer](#) command to create a **dualstack** load balancer.

```
aws elbv2 create-load-balancer --name my-load-balancer \  
--subnets subnet-0e3f5cac72EXAMPLE subnet-081ec835f3EXAMPLE --security-groups  
sg-07e8ffd50fEXAMPLE --ip-address-type dualstack
```

The output includes the Amazon Resource Name (ARN) of the load balancer, with the following format:

```
arn:aws:elasticloadbalancing:us-east-2:123456789012:loadbalancer/app/my-load-  
balancer/1234567890123456
```

2. Use the [create-target-group](#) command to create a target group, specifying the same VPC that you used for your EC2 instances.

You can create IPv4 and IPv6 target groups to associate with dualstack load balancers. The target group's IP address type determines the IP version that the load balancer will use to both communicate with, and check the health of, your backend targets.

```
aws elbv2 create-target-group --name my-targets --protocol HTTP --port 80 \  
--vpc-id vpc-0598c7d356EXAMPLE --ip-address-type [ipv4 or ipv6]
```

The output includes the ARN of the target group, with this format:

```
arn:aws:elasticloadbalancing:us-east-2:123456789012:targetgroup/my-  
targets/1234567890123456
```

3. Use the [register-targets](#) command to register your instances with your target group:

```
aws elbv2 register-targets --target-group-arn targetgroup-arn \  
--targets Id=i-0abcdef1234567890 Id=i-1234567890abcdef0
```

4. Use the [create-listener](#) command to create a listener for your load balancer with a default rule that forwards requests to your target group:

```
aws elbv2 create-listener --load-balancer-arn loadbalancer-arn \  
--protocol HTTP --port 80 \  
--default-actions Type=forward,TargetGroupArn=targetgroup-arn
```

The output contains the ARN of the listener, with the following format:

```
arn:aws:elasticloadbalancing:us-east-2:123456789012:listener/app/my-load-balancer/1234567890123456/1234567890123456
```

5. (Optional) You can verify the health of the registered targets for your target group using this [describe-target-health](#) command:

```
aws elbv2 describe-target-health --target-group-arn targetgroup-arn
```

Add an HTTPS listener

If you have a load balancer with an HTTP listener, you can add an HTTPS listener as follows.

To add an HTTPS listener to your load balancer

1. Create an SSL certificate for use with your load balancer using one of the following methods:
 - Create or import the certificate using AWS Certificate Manager (ACM). For more information, see [Request a certificate](#) or [Importing certificates](#) in the *AWS Certificate Manager User Guide*.
 - Upload the certificate using AWS Identity and Access Management (IAM). For more information, see [Working with server certificates](#) in the *IAM User Guide*.
2. Use the [create-listener](#) command to create the listener with a default rule that forwards requests to your target group. You must specify an SSL certificate when you create an HTTPS listener. Note that you can specify an SSL policy other than the default using the `--ssl-policy` option.

```
aws elbv2 create-listener --load-balancer-arn loadbalancer-arn \  
--protocol HTTPS --port 443 \  
--certificates CertificateArn=certificate-arn \  
--default-actions Type=forward,TargetGroupArn=targetgroup-arn
```

Add path-based routing

If you have a listener with a default rule that forwards requests to one target group, you can add a rule that forwards requests to another target group based on URL. For example, you can route general requests to one target group and requests to display images to another target group.

To add a rule to a listener with a path pattern

1. Use the [create-target-group](#) command to create a target group:

```
aws elbv2 create-target-group --name my-targets --protocol HTTP --port 80 \  
--vpc-id vpc-0598c7d356EXAMPLE
```

2. Use the [register-targets](#) command to register your instances with your target group:

```
aws elbv2 register-targets --target-group-arn targetgroup-arn \  
--targets Id=i-0abcdef1234567890 Id=i-1234567890abcdef0
```

3. Use the [create-rule](#) command to add a rule to your listener that forwards requests to the target group if the URL contains the specified pattern:

```
aws elbv2 create-rule --listener-arn listener-arn --priority 10 \  
--conditions Field=path-pattern,Values='/img/*' \  
--actions Type=forward,TargetGroupArn=targetgroup-arn
```

Delete your load balancer

When you no longer need your load balancer and target group, you can delete them as follows:

```
aws elbv2 delete-load-balancer --load-balancer-arn loadbalancer-arn  
aws elbv2 delete-target-group --target-group-arn targetgroup-arn
```

Application Load Balancers

A *load balancer* serves as the single point of contact for clients. Clients send requests to the load balancer, and the load balancer sends them to targets, such as EC2 instances. To configure your load balancer, you create [target groups](#), and then register targets with your target groups. You also create [listeners](#) to check for connection requests from clients, and listener rules to route requests from clients to the targets in one or more target groups.

For more information, see [How Elastic Load Balancing works](#) in the *Elastic Load Balancing User Guide*.

Contents

- [Subnets for your load balancer](#)
- [Load balancer security groups](#)
- [Load balancer state](#)
- [Load balancer attributes](#)
- [IP address type](#)
- [Application Load Balancer resource map](#)
- [Load balancer connections](#)
- [Cross-zone load balancing](#)
- [Deletion protection](#)
- [Desync mitigation mode](#)
- [Host header preservation](#)
- [Application Load Balancers and AWS WAF](#)
- [Create an Application Load Balancer](#)
- [Availability Zones for your Application Load Balancer](#)
- [Security groups for your Application Load Balancer](#)
- [IP address types for your Application Load Balancer](#)
- [Tags for your Application Load Balancer](#)
- [Delete an Application Load Balancer](#)
- [Zonal shift](#)

Subnets for your load balancer

When you create an Application Load Balancer, you must enable the zones that contain your targets. To enable a zone, specify a subnet in the zone. Elastic Load Balancing creates a load balancer node in each zone that you specify.

Considerations

- Your load balancer is most effective when you ensure that each enabled zone has at least one registered target.
- If you register targets in a zone but do not enable the zone, these registered targets do not receive traffic from the load balancer.
- If you enable multiple zones for your load balancer, the zones must be of the same type. For example, you can't enable both an Availability Zone and a Local Zone.
- You can specify a subnet that was shared with you.

Application Load Balancers support the following types of subnets.

Subnet types

- [Availability Zone subnets](#)
- [Local Zone subnets](#)
- [Outpost subnets](#)

Availability Zone subnets

You must select at least two Availability Zone subnets. The following restrictions apply:

- Each subnet must be from a different Availability Zone.
- To ensure that your load balancer can scale properly, verify that each Availability Zone subnet for your load balancer has a CIDR block with at least a /27 bitmask (for example, 10.0.0.0/27) and at least eight free IP addresses per subnet. These eight IP addresses are required to allow the load balancer to scale out if needed. Your load balancer uses these IP addresses to establish connections with the targets. Without them your Application Load Balancer could experience difficulties with node replacement attempts, causing it to enter a failed state.

Note: If an Application Load Balancers subnet runs out of usable IP addresses while attempting to scale, the Application Load Balancer will run with insufficient capacity. During this time old nodes will continue to serve traffic, but the stalled scaling attempt may cause 5xx errors or timeouts when attempting to establish a connection.

Local Zone subnets

You can specify one or more Local Zone subnets. The following restrictions apply:

- You cannot use AWS WAF with the load balancer.
- You cannot use a Lambda function as a target.
- You cannot use sticky sessions or application stickiness.

Outpost subnets

You can specify a single Outpost subnet. The following restrictions apply:

- You must have installed and configured an Outpost in your on-premises data center. You must have a reliable network connection between your Outpost and its AWS Region. For more information, see the [AWS Outposts User Guide](#).
- The load balancer requires two `large` instances on the Outpost for the load balancer nodes. The supported instance types are shown in the following table. The load balancer scales as needed, resizing the nodes one size at a time (from `large` to `xlarge`, then `xlarge` to `2xlarge`, and then `2xlarge` to `4xlarge`). After scaling the nodes to the largest instance size, if you need additional capacity, the load balancer adds `4xlarge` instances as load balancer nodes. If you do not have sufficient instance capacity or available IP addresses to scale the load balancer, the load balancer reports an event to the [AWS Health Dashboard](#) and the load balancer state is `active_impaired`.
- You can register targets by instance ID or IP address. If you register targets in the AWS Region for the Outpost, they are not used.
- The following features are not available: Lambda functions as targets, AWS WAF integration, sticky sessions, authentication support, and integration with AWS Global Accelerator.

An Application Load Balancer can be deployed on c5/c5d, m5/m5d, or r5/r5d instances on an Outpost. The following table shows the size and EBS volume per instance type that the load balancer can use on an Outpost:

Instance type and size	EBS volume (GB)
c5/c5d	
large	50
xlarge	50
2xlarge	50
4xlarge	100
m5/m5d	
large	50
xlarge	50
2xlarge	100
4xlarge	100
r5/r5d	
large	50
xlarge	100
2xlarge	100
4xlarge	100

Load balancer security groups

A *security group* acts as a firewall that controls the traffic allowed to and from your load balancer. You can choose the ports and protocols to allow for both inbound and outbound traffic.

The rules for the security groups that are associated with your load balancer must allow traffic in both directions on both the listener and the health check ports. Whenever you add a listener to a load balancer or update the health check port for a target group, you must review your security group rules to ensure that they allow traffic on the new port in both directions. For more information, see [Recommended rules](#).

Load balancer state

A load balancer can be in one of the following states:

provisioning

The load balancer is being set up.

active

The load balancer is fully set up and ready to route traffic.

active_impaired

The load balancer is routing traffic but does not have the resources it needs to scale.

failed

The load balancer could not be set up.

Load balancer attributes

The following are the load balancer attributes:

`access_logs.s3.enabled`

Indicates whether access logs stored in Amazon S3 are enabled. The default is `false`.

`access_logs.s3.bucket`

The name of the Amazon S3 bucket for the access logs. This attribute is required if access logs are enabled. For more information, see [Enable access logs](#).

`access_logs.s3.prefix`

The prefix for the location in the Amazon S3 bucket.

`client_keep_alive.seconds`

The client keepalive value, in seconds. The default is 3600 seconds.

`deletion_protection.enabled`

Indicates whether deletion protection is enabled. The default is `false`.

`idle_timeout.timeout_seconds`

The idle timeout value, in seconds. The default is 60 seconds.

`ipv6.deny_all_igw_traffic`

Blocks internet gateway (IGW) access to the load balancer, preventing unintended access to your internal load balancer through an internet gateway. It is set to `false` for internet-facing load balancers and `true` for internal load balancers. This attribute does not prevent non-IGW internet access (such as, through peering, Transit Gateway, AWS Direct Connect, or AWS VPN).

`routing.http.desync_mitigation_mode`

Determines how the load balancer handles requests that might pose a security risk to your application. The possible values are `monitor`, `defensive`, and `strictest`. The default is `defensive`.

`routing.http.drop_invalid_header_fields.enabled`

Indicates whether HTTP headers with header fields that are not valid are removed by the load balancer (`true`), or routed to targets (`false`). The default is `false`. Elastic Load Balancing requires that valid HTTP header names conform to the regular expression `[-A-Za-z0-9]+`, as described in the HTTP Field Name Registry. Each name consists of alphanumeric characters or hyphens. Select `true` if you want HTTP headers that do not conform to this pattern, to be removed from requests.

`routing.http.preserve_host_header.enabled`

Indicates whether the Application Load Balancer should preserve the Host header in the HTTP request and send it to targets without any change. The possible values are `true` and `false`. The default is `false`.

`routing.http.x_amzn_tls_version_and_cipher_suite.enabled`

Indicates whether the two headers (`x-amzn-tls-version` and `x-amzn-tls-cipher-suite`), which contain information about the negotiated TLS version and cipher suite, are added to the client request before sending it to the target. The `x-amzn-tls-version` header

has information about the TLS protocol version negotiated with the client, and the `x-amzn-tls-cipher-suite` header has information about the cipher suite negotiated with the client. Both headers are in OpenSSL format. The possible values for the attribute are `true` and `false`. The default is `false`.

`routing.http.xff_client_port.enabled`

Indicates whether the X-Forwarded-For header should preserve the source port that the client used to connect to the load balancer. The possible values are `true` and `false`. The default is `false`.

`routing.http.xff_header_processing.mode`

Enables you to modify, preserve, or remove the X-Forward-For header in the HTTP request before the Application Load Balancer sends the request to the target. The possible values are `append`, `preserve`, and `remove`. The default is `append`.

- If the value is `append`, the Application Load Balancer adds the client IP address (of the last hop) to the X-Forward-For header in the HTTP request before it sends it to targets.
- If the value is `preserve`, the Application Load Balancer preserves the X-Forward-For header in the HTTP request, and sends it to targets without any change.
- If the value is `remove`, the Application Load Balancer removes the X-Forward-For header in the HTTP request before it sends it to targets.

`routing.http2.enabled`

Indicates whether HTTP/2 is enabled. The default is `true`.

`waf.fail_open.enabled`

Indicates whether to allow a AWS WAF-enabled load balancer to route requests to targets if it is unable to forward the request to AWS WAF. The possible values are `true` and `false`. The default is `false`.

Note

The `routing.http.drop_invalid_header_fields.enabled` attribute was introduced to offer HTTP desync protection. The `routing.http.desync_mitigation_mode` attribute was added to provide more comprehensive protection from HTTP desync for your applications. You aren't required to use both attributes and may choose either, depending on your application's requirements.

IP address type

You can set the types of IP addresses that clients can use to access your internet-facing and internal load balancers.

Application Load Balancers support the following IP address types:

ipv4

Clients must connect to the load balancer using IPv4 addresses (for example, 192.0.2.1)

dualstack

Clients can connect to the load balancer using both IPv4 addresses (for example, 192.0.2.1) and IPv6 addresses (for example, 2001:0db8:85a3:0:0:8a2e:0370:7334).

Considerations

- The load balancer communicates with targets based on the IP address type of the target group.
- When you enable dualstack mode for the load balancer, Elastic Load Balancing provides an AAAA DNS record for the load balancer. Clients that communicate with the load balancer using IPv4 addresses resolve the A DNS record. Clients that communicate with the load balancer using IPv6 addresses resolve the AAAA DNS record.
- Access to your internal dualstack load balancers through the internet gateway is blocked to prevent unintended internet access. However, this does not prevent non-IGW internet access (such as, through peering, Transit Gateway, AWS Direct Connect, or AWS VPN).

dualstack-without-public-ipv4

Clients must connect to the load balancer using IPv6 addresses (for example, 2001:0db8:85a3:0:0:8a2e:0370:7334).

Considerations

- Application Load Balancer authentication only supports IPv4 when connecting to an Identity Provider (IdP) or Amazon Cognito endpoint. Without a public IPv4 address the load balancer cannot complete the authentication process, resulting in HTTP 500 errors.

For more information about IP address types, see [IP address types for your Application Load Balancer](#).

Application Load Balancer resource map

The Application Load Balancer resource map provides an interactive display of your load balancer's architecture, including its associated listeners, rules, target groups, and targets. The resource map also highlights the relationships and routing paths between all resources, producing a visual representation of your load balancer's configuration.

To view your Application Load Balancer's resource map using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. Choose the **Resource map** tab to display the load balancer's resource map.

Resource map components

Map views

There are two views available in the Application Load Balancer resource map: **Overview**, and **Unhealthy Target Map**. **Overview** is selected by default and displays all of your load balancer's resources. Selecting the **Unhealthy Target Map** view will only display the unhealthy targets and the resources associated to them.

The **Unhealthy Target Map** view can be used to troubleshoot targets that are failing health checks. For more information, see [Troubleshoot unhealthy targets using the resource map](#).

Resource groups

The Application Load Balancer resource map contains four resource groups, one for each resource type. The resource groups are **Listeners**, **Rules**, **Target groups**, and **Targets**.

Resource tiles

Each resource within a group has its own tile, which displays details about that specific resource.

- Hovering over a resource tile highlights the relationships between it and other resources.
- Selecting a resource tile highlights the relationships between it and other resources, and displays additional details about that resource.

- **rule conditions:** The conditions for each rule.
- **target group health summary:** The number of registered targets for each health status.
- **target health status** The targets current health status and description.

Note

You can turn off **Show resource details** to hide additional details within the resource map.

- Each resource tile contains a link that, when selected, navigates to that resource's details page.
 - **Listeners** - Select the listeners protocol:port. For example, HTTP:80
 - **Rules** - Select the rules action. For example, Forward to target group
 - **Target groups** - Select the target group name. For example, my-target-group
 - **Targets** - Select the targets ID. For example, i-1234567890abcdef0

Export the resource map

Selecting **Export** gives you the option of exporting the current view of your Application Load Balancer's resource map as a PDF.

Load balancer connections

When processing a request, the load balancer maintains two connections: one connection with the client and one connection with a target. The connection between the load balancer and the client is also referred to as the front-end connection. The connection between the load balancer and the target is also referred to as the back-end connection.

Connection idle timeout

The connection idle timeout is the period of time an existing client or target connection can remain inactive, with no data being sent or received, before the load balancer closes the connection.

To ensure that lengthy operations such as file uploads have time to complete, send at least 1 byte of data before each idle timeout period elapses and increase the length of the idle timeout period as needed. We also recommend that you configure the idle timeout of your application to be larger than the idle timeout configured for the load balancer. Otherwise, if the application closes the

TCP connection to the load balancer ungracefully, the load balancer might send a request to the application before it receives the packet indicating that the connection is closed. If this is the case, then the load balancer sends an HTTP 502 Bad Gateway error to the client.

By default, Elastic Load Balancing sets the idle timeout value for your load balancer to 60 seconds, or 1 minute. Use the following procedure to set a different idle timeout value.

To update the connection idle timeout value using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Attributes** tab, choose **Edit**.
5. Under **Traffic configuration**, enter a value for **Connection idle timeout**. The valid range is 1 through 4000 seconds.
6. Choose **Save changes**.

To update the idle timeout value using the AWS CLI

Use the [modify-load-balancer-attributes](#) command with the `idle_timeout.timeout_seconds` attribute.

HTTP client keepalive duration

The HTTP client keepalive duration is the maximum length of time an Application Load Balancer will maintain a persistent HTTP connection to a client. After the configured HTTP client keepalive duration has elapsed, the Application Load Balancer accepts one request and returns a response that gracefully closes the connection.

The type of response sent by the load balancer depends on the HTTP version used by the client connection. For clients connected using HTTP 1.x, the load balancer sends an HTTP header containing the field `Connection: close`. For clients connected using HTTP/2, the load balancer sends a GOAWAY frame.

By default, Application Load Balancers set the HTTP client keepalive duration value to 3600 seconds, or 1 hour. The HTTP client keepalive duration cannot be turned off or set below the minimum of 60 seconds, but you can increase the HTTP client keepalive duration to a maximum of 604800 seconds, or 7 days. The Application Load Balancer begins the HTTP client keepalive

duration period when an HTTP connection to a client is initially established. The duration period continues to run when there's no traffic, and does not reset until a new connection is established.

Note

When switching the IP address type of your Application Load Balancer to `dualstack-without-public-ipv4` the load balancer waits for all active connections to complete. To decrease the amount of time it takes to switch your Application Load Balancers IP address type, consider lowering the HTTP client keepalive duration.

The Application Load Balancer assigns the HTTP client keepalive duration once during the initial connection. When updating the HTTP client keepalive duration, this can result in simultaneous connections with different HTTP client keepalive duration values. Existing connections will retain the HTTP client keepalive duration value applied during its initial connection, while any new connections will receive the updated HTTP client keepalive duration value.

To update the client keepalive duration value using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Attributes** tab, choose **Edit**.
5. Under **Traffic configuration**, enter a value for **HTTP client keep alive duration**. The valid range is 60 through 604800 seconds.
6. Choose **Save changes**.

To update the client keepalive duration value using the AWS CLI

Use the [modify-load-balancer-attributes](#) command with the `client_keep_alive.seconds` attribute.

Cross-zone load balancing

With Application Load Balancers, cross-zone load balancing is on by default and cannot be changed at the load balancer level. For more information, see the [Cross-zone load balancing](#) section in the *Elastic Load Balancing User Guide*.

Turning off cross-zone load balancing is possible at the target group level. For more information, see [the section called “Turn off cross-zone load balancing”](#).

Deletion protection

To prevent your load balancer from being deleted accidentally, you can enable deletion protection. By default, deletion protection is disabled for your load balancer.

If you enable deletion protection for your load balancer, you must disable it before you can delete the load balancer.

To enable deletion protection using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Attributes** tab, choose **Edit**.
5. Under **Configuration**, turn on **Deletion protection**.
6. Choose **Save changes**.

To disable deletion protection using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Attributes** tab, choose **Edit**.
5. Under **Configuration** page, turn off **Deletion protection**.
6. Choose **Save changes**.

To enable or disable deletion protection using the AWS CLI

Use the [modify-load-balancer-attributes](#) command with the `deletion_protection.enabled` attribute.

Desync mitigation mode

Desync mitigation mode protects your application from issues due to HTTP desync. The load balancer classifies each request based on its threat level, allows safe requests, and then mitigates risk as specified by the mitigation mode that you specify. The desync mitigation modes are monitor, defensive, and strictest. The default is the defensive mode, which provides durable mitigation against HTTP desync while maintaining the availability of your application. You can switch to strictest mode to ensure that your application receives only requests that comply with [RFC 7230](#).

The `http_desync_guardian` library analyzes HTTP requests to prevent HTTP desync attacks. For more information, see [HTTP Desync Guardian](#) on GitHub.

Classifications

The classifications are as follows:

- **Compliant** — Request complies with RFC 7230 and poses no known security threats.
- **Acceptable** — Request does not comply with RFC 7230 but poses no known security threats.
- **Ambiguous** — Request does not comply with RFC 7230 but poses a risk, as various web servers and proxies could handle it differently.
- **Severe** — Request poses a high security risk. The load balancer blocks the request, serves a 400 response to the client, and closes the client connection.

If a request does not comply with RFC 7230, the load balancer increments the `DesyncMitigationMode_NonCompliant_Request_Count` metric. For more information, see [Application Load Balancer metrics](#).

The classification for each request is included in the load balancer access logs. If the request does not comply, the access logs include a classification reason code. For more information, see [Classification reasons](#).

Modes

The following table describes how Application Load Balancers treat requests based on mode and classification.

Classification	Monitor mode	Defensive mode	Strictest mode
Compliant	Allowed	Allowed	Allowed
Acceptable	Allowed	Allowed	Blocked
Ambiguous	Allowed	Allowed ¹	Blocked
Severe	Allowed	Blocked	Blocked

¹ Routes the requests but closes the client and target connections. You might incur additional charges if your load balancer receives a large number of Ambiguous requests in Defensive mode. This is because the increased number of new connections per second contributes to the Load Balancer Capacity Units (LCU) used per hour. You can use the `NewConnectionCount` metric to compare how your load balancer establishes new connections in Monitor mode and Defensive mode.

To update desync mitigation mode using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Attributes** tab, choose **Edit**.
5. Under **Packet handling**, for **Desync mitigation mode**, choose **Defensive**, **Strictest**, or **Monitor**.
6. Choose **Save changes**.

To update desync mitigation mode using the AWS CLI

Use the [modify-load-balancer-attributes](#) command with the `routing.http.desync_mitigation_mode` attribute set to `monitor`, `defensive`, or `strictest`.

Host header preservation

When you enable the **Preserve host header** attribute, the Application Load Balancer preserves the Host header in the HTTP request, and sends the header to targets without any modification. If the

Application Load Balancer receives multiple Host headers, it preserves all of them. Listener rules are applied only to the first Host header received.

By default, when the **Preserve host header** attribute is not enabled, the Application Load Balancer modifies the Host header in the following manner:

When host header preservation is not enabled, and listener port is a non-default port: When not using the default ports (ports 80 or 443) we append the port number to the host header if it isn't already appended by the client. For example, the Host header in the HTTP request with Host: `www.example.com` would be modified to Host: `www.example.com:8080`, if the listener port is a non-default port such as 8080.

When host header preservation is not enabled, and the listener port is a default port (port 80 or 443): For default listener ports (either port 80 or 443), we do not append the port number to the outgoing host header. Any port number that was already in the incoming host header, is removed.

The following table shows more examples of how Application Load Balancers treat host headers in the HTTP request based on listener port.

Listener port	Example request	Host header in the request	Host header preservation is disabled (default behavior)	Host header preservation is enabled
Request is sent on default HTTP/HTTPS listener.	GET / index.ht ml HTTP/1.1 Host: example.com	example.com	example.com	example.com
Request is sent on default HTTP listener and host header has a port (for	GET / index.ht ml HTTP/1.1 Host: example.c om:80	example.com:80	example.com	example.com:80

Listener port	Example request	Host header in the request	Host header preservation is disabled (default behavior)	Host header preservation is enabled
example, 80 or 443).				
Request has an absolute path.	GET https:// dns_name/ index.html HTTP/1.1 Host: example.com	example.com	dns_name	example.com
Request is sent on a non-default listener port (for example, 8080)	GET / index.ht ml HTTP/1.1 Host: example.com	example.com	example.c om:8080	example.com
Request is sent on a non-default listener port and host header has port (for example, 8080).	GET / index.ht ml HTTP/1.1 Host: example.c om:8080	example.c om:8080	example.c om:8080	example.c om:8080

To enable host header preservation using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Attributes** tab, choose **Edit**.
5. Under **Packet handling**, turn on **Preserve host header**.

6. Choose **Save changes**.

To enable host header preservation using the AWS CLI

Use the [modify-load-balancer-attributes](#) command with the `routing.http.preserve_host_header.enabled` attribute set to `true`.

Application Load Balancers and AWS WAF

You can use AWS WAF with your Application Load Balancer to allow or block requests based on the rules in a web access control list (web ACL). For more information, see [Working with web ACLs](#) in the *AWS WAF Developer Guide*.

By default, if the load balancer cannot get a response from AWS WAF, it returns an HTTP 500 error and does not forward the request. If you need your load balancer to forward requests to targets even if it is unable to contact AWS WAF, you can enable AWS WAF fail open. To check whether your load balancer integrates with AWS WAF, select your load balancer in the AWS Management Console and choose the **Integrated services** tab.

Pre-defined web ACLs

When enabling AWS WAF integration you can choose to automatically create a new web ACL with pre-defined rules. The pre-defined web ACL includes three AWS managed rules which offer protections against the most common security threats.

- `AWSManagedRulesAmazonIpReputationList` - The Amazon IP reputation list rule group blocks IP addresses typically associated with bots or other threats. For more information, see [Amazon IP reputation list managed rule group](#) in the *AWS WAF Developer Guide*.
- `AWSManagedRulesCommonRuleSet` - The core rule set (CRS) rule group provides protection against exploitation of a wide range of vulnerabilities, including some of the high risk and commonly occurring vulnerabilities described in OWASP publications such as [OWASP Top 10](#). For more information, see [Core rule set \(CRS\) managed rule group](#) in the *AWS WAF Developer Guide*.
- `AWSManagedRulesKnownBadInputsRuleSet` - The Known bad inputs rule group blocks request patterns that are known to be invalid and are associated with exploitation or discovery of vulnerabilities. For more information, see [Known bad inputs managed rule group](#) in the *AWS WAF Developer Guide*.

To enable AWS WAF using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Integrations** tab, expand **AWS Web Application Firewall (WAF)**, and choose **Associate a WAF web ACL**.
5. Under **Web ACL**, choose **Auto-create pre-defined web ACL**, or select an existing web ACL.
6. Under **Rule action**, choose **Block**, or **Count**.
7. Choose **Confirm**.

To enable AWS WAF fail open using the AWS CLI

Use the [modify-load-balancer-attributes](#) command with the `waf.fail_open.enabled` attribute set to `true`.

Create an Application Load Balancer

A load balancer takes requests from clients and distributes them across targets in a target group.

Before you begin, ensure that you have a virtual private cloud (VPC) with at least one public subnet in each of the zones used by your targets. For more information, see [the section called "Subnets for your load balancer"](#).

To create a load balancer using the AWS CLI, see [Tutorial: Create an Application Load Balancer using the AWS CLI](#).

To create a load balancer using the AWS Management Console, complete the following tasks.

Tasks

- [Step 1: Configure a target group](#)
- [Step 2: Register targets](#)
- [Step 3: Configure a load balancer and a listener](#)
- [Step 4: Test the load balancer](#)

Step 1: Configure a target group

Configuring a target group allows you to register targets such as EC2 instances. The target group that you configure in this step is used as the target group in the listener rule when you configure your load balancer. For more information, see [Target groups for your Application Load Balancers](#).

To configure your target group using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose **Target Groups**.
3. Choose **Create target group**.
4. In the **Basic configuration** section, set the following parameters:
 - a. For **Choose a target type**, select **Instances** to specify targets by instance ID or **IP addresses** to specify targets by only IP address. If the target type is a **Lambda function**, you can enable health checks by selecting **Enable** in the **Health checks** section.
 - b. For **Target group name**, enter a name for the target group.
 - c. Modify the **Port** and **Protocol** as needed.
 - d. If the target type is **Instances** or **IP addresses**, choose **IPv4** or **IPv6** as the **IP address type**, otherwise skip to the next step.

Note that only targets that have the selected IP address type can be included in this target group. The IP address type cannot be changed after the target group is created.
 - e. For **VPC**, select a virtual private cloud (VPC) with the targets that you want to include in your target group.
 - f. For **Protocol version**, select **HTTP1** when the request protocol is HTTP/1.1 or HTTP/2; select **HTTP2**, when the request protocol is HTTP/2 or gRPC; and select **gRPC**, when the request protocol is gRPC.
5. In the **Health checks** section, modify the default settings as needed. For **Advanced health check settings**, choose the health check port, count, timeout, interval, and specify success codes. If health checks consecutively exceed the **Unhealthy threshold** count, the load balancer takes the target out of service. If health checks consecutively exceed the **Healthy threshold** count, the load balancer puts the target back in service. For more information, see [Health checks for your target groups](#).
6. (Optional) Add one or more tags as follows:

- a. Expand the **Tags** section.
 - b. Choose **Add tag**.
 - c. Enter the tag **Key** and tag **Value**. Allowed characters are letters, spaces, numbers (in UTF-8), and the following special characters: + - = . _ : / @. Do not use leading or trailing spaces. Tag values are case-sensitive.
7. Choose **Next**.

Step 2: Register targets

You can register EC2 instances, IP addresses, or Lambda functions as targets in a target group. This is an optional step to create a load balancer. However, you must register your targets to ensure that your load balancer routes traffic to them.

1. In the **Register targets** page, add one or more targets as follows:
 - If the target type is **Instances**, select one or more instances, enter one or more ports, and then choose **Include as pending below**.
 - If the target type is **IP addresses**, do the following:
 - a. Select a network **VPC** from the list, or choose **Other private IP addresses**.
 - b. Enter the IP address manually, or find the IP address using instance details. You can enter up to five IP addresses at a time.
 - c. Enter the ports for routing traffic to the specified IP addresses.
 - d. Choose **Include as pending below**.
 - If the target type is **Lambda**, select a Lambda function, or enter a Lambda function ARN, and then choose **Include as pending below**.
2. Choose **Create target group**.

Step 3: Configure a load balancer and a listener

To create an Application Load Balancer, you must first provide basic configuration information for your load balancer, such as a name, scheme, and IP address type. Then, you provide information about your network, and one or more listeners. A listener is a process that checks for connection requests. It is configured with a protocol and a port for connections from clients to the load balancer. For more information about supported protocols and ports, see [Listener configuration](#).

To configure your load balancer and listener using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose **Load Balancers**.
3. Choose **Create Load Balancer**.
4. Under **Application Load Balancer**, choose **Create**.
5. **Basic configuration**
 - a. For **Load balancer name**, enter a name for your load balancer. For example, **my-a1b**. The name of your Application Load Balancer must be unique within your set of Application Load Balancers and Network Load Balancers for the Region. Names can have a maximum of 32 characters, and can contain only alphanumeric characters and hyphens. They can not begin or end with a hyphen, or with `internal-`. The name of your Application Load Balancer cannot be changed after it's created.
 - b. For **Scheme**, choose **Internet-facing** or **Internal**. An internet-facing load balancer routes requests from clients to targets over the internet. An internal load balancer routes requests to targets using private IP addresses.
 - c. For **IP address type**, choose **IPv4**, **Dualstack**, or **Dualstack without public IPv4**. Choose **IPv4** if your clients use IPv4 addresses to communicate with the load balancer. Choose **Dualstack** if your clients use both IPv4 and IPv6 addresses to communicate with the load balancer. Choose **Dualstack without public IPv4** if your clients use only IPv6 addresses to communicate with the load balancer.
6. **Network mapping**
 - a. For **VPC**, select the VPC that you used for your EC2 instances. If you selected **Internet-facing** for **Scheme**, only VPCs with an internet gateway are available for selection.
 - b. For **Mappings**, enable zones for your load balancer by selecting subnets as follows:
 - Subnets from two or more Availability Zones
 - Subnets from one or more Local Zones
 - One Outpost subnet

For more information, see [the section called "Subnets for your load balancer"](#).

For internal load balancers, the IPv4 and IPv6 addresses are assigned from the subnet CIDR.

If you enabled **Dualstack** mode for the load balancer, select subnets with both IPv4 and IPv6 CIDR blocks.

7. For **Security groups**, select an existing security group, or create a new one.

The security group for your load balancer must allow it to communicate with registered targets on both the listener port and the health check port. The console can create a security group for your load balancer on your behalf with rules that allow this communication. You can also create a security group and select it instead. For more information, see [Recommended rules](#).

(Optional) To create a new security group for your load balancer, choose **Create a new security group**.

8. For **Listeners and routing**, the default listener accepts HTTP traffic on port 80. You can keep the default protocol and port, or choose different ones. For **Default action**, choose the target group that you created. You can optionally choose **Add listener** to add another listener (for example, an HTTPS listener).
9. (Optional) If using an HTTPS listener

For **Security policy**, we recommend that you always use the latest predefined security policy.

- a. For **Default SSL/TLS certificate**, the following options are available:

- If you created or imported a certificate using AWS Certificate Manager, select **From ACM**, then select the certificate from **Select a certificate**.
- If you imported a certificate using IAM, select **From IAM**, and then select your certificate from **Select a certificate**.
- If you have a certificate to import but ACM is not available in your Region, select **Import**, then select **To IAM**. Type the name of the certificate in the **Certificate name** field. In **Certificate private key**, copy and paste the contents of the private key file (PEM-encoded). In **Certificate body**, copy and paste the contents of the public key certificate file (PEM-encoded). In **Certificate Chain**, copy and paste the contents of the certificate chain file (PEM-encoded), unless you are using a self-signed certificate and it's not important that browsers implicitly accept the certificate.

- b. (Optional) To enable mutual authentication, under **Client certificate handling** enable **Mutual authentication (mTLS)**.

When enabled, the default mutual TLS mode is **passthrough**.

If you select **Verify with Trust Store**:

- By default, connections with expired client certificates are rejected. To change this behavior expand **Advanced mTLS settings**, then under **Client certificate expiration** select **Allow expired client certificates**.
- Under **Trust Store** choose an existing trust store, or choose **New trust store**.
 - If you chose **New trust store**, provide a **Trust store name**, the **S3 URI Certificate Authority location**, and optionally an **S3 URI Certificate revocation list location**.

10. (Optional) You can integrate other services with your load balancer during creation, under **Optimize with service integrations**.

- You can choose to include **AWS WAF** security protections for your load balancer, with an existing or automatically created web ACL. After creation, web ACLs can be managed in the [AWS WAF console](#). For more information, see [Associating or disassociating a web ACL with an AWS resource](#) in the *AWS WAF Developer Guide*.
- You can choose to have **AWS Global Accelerator** create an accelerator for you and associate your load balancer with the accelerator. The accelerator name can have the following characters (up to 64 characters): a-z, A-Z, 0-9, . (period), and - (hyphen). After the accelerator is created, you can manage it in the [AWS Global Accelerator console](#). For more information, see [Add an accelerator when you create a load balancer](#) in the *AWS Global Accelerator Developer Guide*.

11. **Tag and create**

- (Optional) Add a tag to categorize your load balancer. Tag keys must be unique for each load balancer. Allowed characters are letters, spaces, numbers (in UTF-8), and the following special characters: + - = . _ : / @. Do not use leading or trailing spaces. Tag values are case-sensitive.
- Review your configuration, and choose **Create load balancer**. A few default attributes are applied to your load balancer during creation. You can view and edit them after creating the load balancer. For more information, see [Load balancer attributes](#).

Step 4: Test the load balancer

After creating your load balancer, you can verify that your EC2 instances pass the initial health check. You can then check that the load balancer is sending traffic to your EC2 instance. To delete the load balancer, see [Delete an Application Load Balancer](#).

To test the load balancer

1. After the load balancer is created, choose **Close**.
2. In the navigation pane, choose **Target Groups**.
3. Select the newly created target group.
4. Choose **Targets** and verify that your instances are ready. If the status of an instance is `initial`, it's typically because the instance is still in the process of being registered. This status can also indicate that the instance has not passed the minimum number of health checks to be considered healthy. After the status of at least one instance is healthy, you can test your load balancer. For more information, see [Target health status](#).
5. In the navigation pane, choose **Load Balancers**.
6. Select the newly created load balancer.
7. Choose **Description** and copy the DNS name of the internet facing or internal load balancer (for example, `my-load-balancer-1234567890abcdef.elb.us-east-2.amazonaws.com`).
 - For internet facing load balancers, paste the DNS name into the address field of an internet connected web browser.
 - For internal load balancers, paste the DNS name into the address field of a web browser which has private connectivity to the VPC.

If everything is configured correctly, the browser displays the default page of your server.

8. If the web page does not display, refer to the following documents for additional configuration help and troubleshooting steps.
 - For DNS related issues, see [Routing traffic to an ELB load balancer](#) in the *Amazon Route 53 Developer Guide*.
 - For Load Balancer related issues, see [Troubleshoot your Application Load Balancers](#).

Availability Zones for your Application Load Balancer

You can enable or disable the Availability Zones for your load balancer at any time. After you enable an Availability Zone, the load balancer starts routing requests to the registered targets in that Availability Zone. Your load balancer is most effective if you ensure that each enabled Availability Zone has at least one registered target.

After you disable an Availability Zone, the targets in that Availability Zone remain registered with the load balancer, but the load balancer will not route requests to them.

To update Availability Zones using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Network mapping** tab, choose **Edit subnets**.
5. To enable an Availability Zone, select its check box and select one subnet. If there is only one available subnet, it is selected for you.
6. To change the subnet for an enabled Availability Zone, choose one of the other subnets from the list.
7. To disable an Availability Zone, clear its check box.
8. Choose **Save changes**.

To update Availability Zones using the AWS CLI

Use the [set-subnets](#) command.

Security groups for your Application Load Balancer

The security group for your Application Load Balancer controls the traffic that is allowed to reach and leave the load balancer. You must ensure that your load balancer can communicate with registered targets on both the listener port and the health check port. Whenever you add a listener to your load balancer or update the health check port for a target group used by the load balancer to route requests, you must verify that the security groups associated with the load balancer allow traffic on the new port in both directions. If they don't, you can edit the rules for the currently

associated security groups or associate different security groups with the load balancer. You can choose the ports and protocols to allow. For example, you can open Internet Control Message Protocol (ICMP) connections for the load balancer to respond to ping requests (however, ping requests are not forwarded to any instances).

Recommended rules

The following rules are recommended for an internet-facing load balancer.

Inbound		
Source	Port Range	Comment
0.0.0.0/0	<i>listener</i>	Allow all inbound traffic on the load balancer listener port
Outbound		
Destination	Port Range	Comment
<i>instance security group</i>	<i>instance listener</i>	Allow outbound traffic to instances on the instance listener port
<i>instance security group</i>	<i>health check</i>	Allow outbound traffic to instances on the health check port

The following rules are recommended for an internal load balancer.

Inbound		
Source	Port Range	Comment
<i>VPC CIDR</i>	<i>listener</i>	Allow inbound traffic from the VPC CIDR on the load balancer listener port
Outbound		

Destination	Port Range	Comment
<i>instance security group</i>	<i>instance listener</i>	Allow outbound traffic to instances on the instance listener port
<i>instance security group</i>	<i>health check</i>	Allow outbound traffic to instances on the health check port

The following rules are recommended for an Application Load Balancer used as a target of a Network Load Balancer.

Inbound

Source	Port Range	Comment
<i>client IP addresses/ CIDR</i>	<i>alb listener</i>	Allow inbound client traffic on the load balancer listener port
<i>VPC CIDR</i>	<i>alb listener</i>	Allow inbound client traffic via AWS PrivateLink on the load balancer listener port
<i>VPC CIDR</i>	<i>alb listener</i>	Allow inbound health traffic from the Network Load Balancer

Outbound

Destination	Port Range	Comment
<i>instance security group</i>	<i>instance listener</i>	Allow outbound traffic to instances on the instance listener port

*instance security
group*

health check

Allow outbound traffic to instances on the health check port

Note that the security groups for your Application Load Balancer use connection tracking to track information about traffic coming from the Network Load Balancer. This happens regardless of the security group rules set for your Application Load Balancer. To learn more about Amazon EC2 connection tracking, see [Security group connection tracking](#) in the *Amazon EC2 User Guide*.

To ensure your targets receive traffic exclusively from the load balancer, restrict the security groups associated with your targets to accept traffic solely from the load balancer. This can be achieved by setting the load balancer's security group as the source in the ingress rule of the target's security group.

We also recommend that you allow inbound ICMP traffic to support Path MTU Discovery. For more information, see [Path MTU Discovery](#) in the *Amazon EC2 User Guide*.

Update the associated security groups

You can update the security groups associated with your load balancer at any time.

To update security groups using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Security** tab, choose **Edit**.
5. To associate a security group with your load balancer, select it. To remove a security group association, choose the **X** icon for the security group.
6. Choose **Save changes**.

To update security groups using the AWS CLI

Use the [set-security-groups](#) command.

IP address types for your Application Load Balancer

You can configure your Application Load Balancer so that clients can communicate with the load balancer using IPv4 addresses only, or using both IPv4 and IPv6 addresses (dualstack). The load balancer communicates with targets based on the IP address type of the target group. For more information, see [IP address type](#).

Dualstack requirements

- You can set the IP address type when you create the load balancer and update it at any time.
- The virtual private cloud (VPC) and subnets that you specify for the load balancer must have associated IPv6 CIDR blocks. For more information, see [IPv6 addresses](#) in the *Amazon EC2 User Guide*.
- The route tables for the load balancer subnets must route IPv6 traffic.
- The security groups for the load balancer must allow IPv6 traffic.
- The network ACLs for the load balancer subnets must allow IPv6 traffic.

To set the IP address type at creation

Configure settings as described in [???](#).

To update the IP address type using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Network mapping** tab, choose **Edit IP address type**.
5. For **IP address type**, choose **IPv4** to support IPv4 addresses only, **Dualstack** to support both IPv4 and IPv6 addresses, or **Dualstack without public IPv4** to support IPv6 addresses only.
6. Choose **Save changes**.

To update the IP address type using the AWS CLI

Use the [set-ip-address-type](#) command.

Tags for your Application Load Balancer

Tags help you to categorize your load balancers in different ways, for example, by purpose, owner, or environment.

You can add multiple tags to each load balancer. If you add a tag with a key that is already associated with the load balancer, it updates the value of that tag.

When you are finished with a tag, you can remove it from your load balancer.

Restrictions

- Maximum number of tags per resource—50
- Maximum key length—127 Unicode characters
- Maximum value length—255 Unicode characters
- Tag keys and values are case sensitive. Allowed characters are letters, spaces, and numbers representable in UTF-8, plus the following special characters: + - = . _ : / @. Do not use leading or trailing spaces.
- Do not use the `aws :` prefix in your tag names or values because it is reserved for AWS use. You can't edit or delete tag names or values with this prefix. Tags with this prefix do not count against your tags per resource limit.

To update the tags for a load balancer using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Tags** tab, choose **Manage tags**, and then do one or more of the following:
 - a. To update a tag, edit the values of **Key** and **Value**.
 - b. To add a new tag, choose **Add tag** and then enter values for **Key** and **Value**.
 - c. To delete a tag, choose the **Remove** button next to the tag.
5. When you have finished updating tags, choose **Save changes**.

To update the tags for a load balancer using the AWS CLI

Use the [add-tags](#) and [remove-tags](#) commands.

Delete an Application Load Balancer

As soon as your load balancer becomes available, you are billed for each hour or partial hour that you keep it running. When you no longer need the load balancer, you can delete it. As soon as the load balancer is deleted, you stop incurring charges for it.

You can't delete a load balancer if deletion protection is enabled. For more information, see [Deletion protection](#).

Note that deleting a load balancer does not affect its registered targets. For example, your EC2 instances continue to run and are still registered to their target groups. To delete your target groups, see [Delete a target group](#).

To delete a load balancer using the console

1. If you have a DNS record for your domain that points to your load balancer, point it to a new location and wait for the DNS change to take effect before deleting your load balancer.

Example:

- If the record is a CNAME record with a Time To Live (TTL) of 300 seconds, wait at least 300 seconds before continuing to the next step.
 - If the record is a Route 53 Alias(A) record, wait at least 60 seconds.
 - If using Route 53, the record change takes 60 seconds to propagate to all global Route 53 name servers. Add this time to the TTL value of the record that is being updated.
2. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
 3. On the navigation pane, choose **Load Balancers**.
 4. Select the load balancer, and then choose **Actions, Delete load balancer**.
 5. When prompted for confirmation, enter **confirm** and then choose **Delete**.

To delete a load balancer using the AWS CLI

Use the [delete-load-balancer](#) command.

Zonal shift

Zonal shift is a capability in Amazon Route 53 Application Recovery Controller (Route 53 ARC). With zonal shift, you can shift a load balancer resource away from an impaired Availability Zone

with a single action. This way, you can continue operating from other healthy Availability Zones in an AWS Region.

When you start a zonal shift, your load balancer stops sending traffic for the resource to the affected Availability Zone. Route 53 ARC creates the zonal shift immediately. However, it can take a short time, typically up to a few minutes, to complete existing, in-progress connections in the affected Availability Zone. For more information, see [How a zonal shift works: health checks and zonal IP addresses](#) in the *Amazon Route 53 Application Recovery Controller Developer Guide*.

Zonal shifts are only supported on Application Load Balancers and Network Load Balancers with cross-zone load balancing turned off. If you turn on cross-zone load balancing, you can't start a zonal shift. For more information, see [Resources supported for zonal shifts](#) in the *Amazon Route 53 Application Recovery Controller Developer Guide*.

Before you use a zonal shift, review the following:

- Cross-zone load balancing isn't supported with zonal shifts. You must turn off cross-zone load balancing to use this capability.
- Zonal shift isn't supported when you use an Application Load Balancer as an accelerator endpoint in AWS Global Accelerator.
- You can start a zonal shift for a specific load balancer only for a single Availability Zone. You can't start a zonal shift for multiple Availability Zones.
- AWS proactively removes zonal load balancer IP addresses from DNS when multiple infrastructure issues impact services. Always check current Availability Zone capacity before you start a zonal shift. If your load balancers have cross-zone load balancing turned off and you use a zonal shift to remove a zonal load balancer IP address, the Availability Zone affected by the zonal shift also loses target capacity.
- When an Application Load Balancer is a target of a Network Load Balancer, always start the zonal shift from the Network Load Balancer. If you start a zonal shift from the Application Load Balancer, the Network Load Balancer doesn't recognize the shift and continues to send traffic to the Application Load Balancer.

For more guidance and information, see [Best practices with Route 53 ARC zonal shifts](#) in the *Amazon Route 53 Application Recovery Controller Developer Guide*.

Start a zonal shift

The steps in this procedure explain how to start a zonal shift using the Amazon EC2 console. For steps to start a zonal shift using the Route 53 ARC console, see [Starting a zonal shift](#) in the *Amazon Route 53 Application Recovery Controller Developer Guide*.

To start a zonal shift using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Load Balancers**.
3. Select the load balancer name.
4. On the **Integrations** tab, under **Route 53 Application Recovery Controller**, choose **Start zonal shift**.
5. Select the Availability Zone that you want to move traffic away from.
6. Choose or enter an expiration for the zonal shift. A zonal shift can initially be set from 1 minute up to three days (72 hours).

All zonal shifts are temporary. You must set an expiration, but you can update active shifts later to set a new expiration.

7. Enter a comment. You can update the zonal shift later to edit the comment, if you like.
8. Select the check box to acknowledge that starting a zonal shift will reduce capacity for your application by shifting traffic away from the Availability Zone.
9. Choose **Start**.

To start a zonal shift using the AWS CLI

To work with zonal shift programmatically, see the [Zonal Shift API Reference Guide](#).

Update a zonal shift

The steps in this procedure explain how to update a zonal shift using the Amazon EC2 console. For steps to update a zonal shift using the Amazon Route 53 Application Recovery Controller console, see [Updating a zonal shift](#) in the *Amazon Route 53 Application Recovery Controller Developer Guide*.

To update a zonal shift using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.

2. On the navigation pane, under **Load Balancing**, choose **Load Balancers**.
3. Select a load balancer name that has an active zonal shift.
4. On the **Integrations** tab, under **Route 53 Application Recovery Controller**, choose **Update zonal shift**.

This opens the Route 53 ARC console to continue the update.

5. For **Set zonal shift expiration**, optionally select or enter an expiration.
6. For **Comment**, optionally edit the existing comment or enter a new comment.
7. Choose **Update**.

To update a zonal shift using the AWS CLI

To work with zonal shift programmatically, see the [Zonal Shift API Reference Guide](#).

Cancel a zonal shift

The steps in this procedure explain how to cancel a zonal shift using the Amazon EC2 console. For steps to cancel a zonal shift using the Amazon Route 53 Application Recovery Controller console, see [Canceling a zonal shift](#) in the *Amazon Route 53 Application Recovery Controller Developer Guide*.

To cancel a zonal shift using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Load Balancers**.
3. Select a load balancer name that has an active zonal shift.
4. On the **Integrations** tab, under **Route 53 Application Recovery Controller**, choose **Cancel zonal shift**.

This opens the Route 53 ARC console to continue the cancelation.

5. Choose **Cancel zonal shift**.
6. On the confirmation dialog, choose **Confirm**.

To cancel a zonal shift using the AWS CLI

To work with zonal shift programmatically, see the [Zonal Shift API Reference Guide](#).

Listeners for your Application Load Balancers

A *listener* is a process that checks for connection requests, using the protocol and port that you configure. Before you start using your Application Load Balancer, you must add at least one listener. If your load balancer has no listeners, it can't receive traffic from clients. The rules that you define for your listeners determine how the load balancer routes requests to the targets that you register, such as EC2 instances.

Contents

- [Listener configuration](#)
- [Listener rules](#)
- [Rule action types](#)
- [Rule condition types](#)
- [Create an HTTP listener for your Application Load Balancer](#)
- [Create an HTTPS listener for your Application Load Balancer](#)
- [Listener rules for your Application Load Balancer](#)
- [Update an HTTPS listener for your Application Load Balancer](#)
- [Mutual authentication with TLS in Application Load Balancer](#)
- [Authenticate users using an Application Load Balancer](#)
- [HTTP headers and Application Load Balancers](#)
- [Tags for your listeners and rules](#)
- [Delete a listener for your Application Load Balancer](#)

Listener configuration

Listeners support the following protocols and ports:

- **Protocols:** HTTP, HTTPS
- **Ports:** 1-65535

You can use an HTTPS listener to offload the work of encryption and decryption to your load balancer so that your applications can focus on their business logic. If the listener protocol is

HTTPS, you must deploy at least one SSL server certificate on the listener. For more information, see [Create an HTTPS listener for your Application Load Balancer](#).

If you must ensure that the targets decrypt HTTPS traffic instead of the load balancer, you can create a Network Load Balancer with a TCP listener on port 443. With a TCP listener, the load balancer passes encrypted traffic through to the targets without decrypting it. For more information, see the [User Guide for Network Load Balancers](#).

Application Load Balancers provide native support for WebSockets. You can upgrade an existing HTTP/1.1 connection into a WebSocket (ws or wss) connection by using an HTTP connection upgrade. When you upgrade, the TCP connection used for requests (to the load balancer as well as to the target) becomes a persistent WebSocket connection between the client and the target through the load balancer. You can use WebSockets with both HTTP and HTTPS listeners. The options that you choose for your listener apply to WebSocket connections as well as to HTTP traffic. For more information, see [How the WebSocket Protocol Works](#) in the *Amazon CloudFront Developer Guide*.

Application Load Balancers provide native support for HTTP/2 with HTTPS listeners. You can send up to 128 requests in parallel using one HTTP/2 connection. You can use the protocol version to send the request to the targets using HTTP/2. For more information, see [Protocol version](#). Because HTTP/2 uses front-end connections more efficiently, you might notice fewer connections between clients and the load balancer. You can't use the server-push feature of HTTP/2.

For more information, see [Request routing](#) in the *Elastic Load Balancing User Guide*.


Listener rules

Every listener has a default action, also known as the default rule. The default rule cannot be deleted and is always performed last. Additional rules can be created and consist of a priority, one or more actions, and one or more conditions. You can add or edit rules at any time. For more information, see [Edit a rule](#).

Default rules

When you create a listener, you define actions for the default rule. Default rules can't have conditions. If the conditions for none of a listener's rules are met, then the action for the default rule is performed.

The following is an example of a default rule as shown in the console:

Priority	Conditions (If)	Actions (Then) 
Last (default)	<i>If no other rule applies</i>	Forward to target group <ul style="list-style-type: none"> • my-targets: 1 (100%) • Group-level stickiness: Off

Rule priority

Each rule has a priority. Rules are evaluated in priority order, from the lowest value to the highest value. The default rule is evaluated last. You can change the priority of a nondefault rule at any time. You cannot change the priority of the default rule. For more information, see [Update rule priority](#).

Rule actions

Each rule action has a type, a priority, and the information required to perform the action. For more information, see [Rule action types](#).

Rule conditions

Each rule condition has a type and configuration information. When the conditions for a rule are met, then its actions are performed. For more information, see [Rule condition types](#).

Rule action types

The following are the supported action types for a listener rule:

`authenticate-cognito`

[HTTPS listeners] Use Amazon Cognito to authenticate users. For more information, see [Authenticate users using an Application Load Balancer](#).

`authenticate-oidc`

[HTTPS listeners] Use an identity provider that is compliant with OpenID Connect (OIDC) to authenticate users.

`fixed-response`

Return a custom HTTP response. For more information, see [Fixed-response actions](#).

forward

Forward requests to the specified target groups. For more information, see [Forward actions](#).

redirect

Redirect requests from one URL to another. For more information, see [Redirect actions](#).

The action with the lowest priority is performed first. Each rule must include exactly one of the following actions: `forward`, `redirect`, or `fixed-response`, and it must be the last action to be performed.

If the protocol version is gRPC or HTTP/2, the only supported actions are `forward` actions.

Fixed-response actions

You can use `fixed-response` actions to drop client requests and return a custom HTTP response. You can use this action to return a 2XX, 4XX, or 5XX response code and an optional message.

When a `fixed-response` action is taken, the action and the URL of the redirect target are recorded in the access logs. For more information, see [Access log entries](#). The count of successful `fixed-response` actions is reported in the `HTTP_Fixed_Response_Count` metric. For more information, see [Application Load Balancer metrics](#).

Example Example fixed response action for the AWS CLI

You can specify an action when you create or modify a rule. For more information, see the [create-rule](#) and [modify-rule](#) commands. The following action sends a fixed response with the specified status code and message body.

```
[
  {
    "Type": "fixed-response",
    "FixedResponseConfig": {
      "StatusCode": "200",
      "ContentType": "text/plain",
      "MessageBody": "Hello world"
    }
  }
]
```

Forward actions

You can use forward actions to route requests to one or more target groups. If you specify multiple target groups for a forward action, you must specify a weight for each target group. Each target group weight is a value from 0 to 999. Requests that match a listener rule with weighted target groups are distributed to these target groups based on their weights. For example, if you specify two target groups, each with a weight of 10, each target group receives half the requests. If you specify two target groups, one with a weight of 10 and the other with a weight of 20, the target group with a weight of 20 receives twice as many requests as the other target group.

By default, configuring a rule to distribute traffic between weighted target groups does not guarantee that sticky sessions are honored. To ensure that sticky sessions are honored, enable target group stickiness for the rule. When the load balancer first routes a request to a weighted target group, it generates a cookie named `AWSALBTG` that encodes information about the selected target group, encrypts the cookie, and includes the cookie in the response to the client. The client should include the cookie that it receives in subsequent requests to the load balancer. When the load balancer receives a request that matches a rule with target group stickiness enabled and contains the cookie, the request is routed to the target group specified in the cookie.

Application Load Balancers do not support cookie values that are URL encoded.

With CORS (cross-origin resource sharing) requests, some browsers require `SameSite=None; Secure` to enable stickiness. In this case, Elastic Load Balancing generates a second cookie, `AWSALBTGCORS`, which includes the same information as the original stickiness cookie plus this `SameSite` attribute. Clients receive both cookies.

Example Example forward action with one target group

You can specify an action when you create or modify a rule. For more information, see the [create-rule](#) and [modify-rule](#) commands. The following action forwards requests to the specified target group.

```
[
  {
    "Type": "forward",
    "ForwardConfig": {
      "TargetGroups": [
        {
          "TargetGroupArn": "arn:aws:elasticloadbalancing:us-
west-2:123456789012:targetgroup/my-targets/73e2d6bc24d8a067"
```

```

    }
  ]
}
]

```

Example Example forward action with two weighted target groups

The following action forwards requests to the two specified target groups, based on the weight of each target group.

```

[
  {
    "Type": "forward",
    "ForwardConfig": {
      "TargetGroups": [
        {
          "TargetGroupArn": "arn:aws:elasticloadbalancing:us-
west-2:123456789012:targetgroup/blue-targets/73e2d6bc24d8a067",
          "Weight": 10
        },
        {
          "TargetGroupArn": "arn:aws:elasticloadbalancing:us-
west-2:123456789012:targetgroup/green-targets/09966783158cda59",
          "Weight": 20
        }
      ]
    }
  }
]

```

Example Example forward action with stickiness enabled

If you have a forward action with multiple target groups and one or more of the target groups has [sticky sessions](#) enabled, you must enable target group stickiness.

The following action forwards requests to the two specified target groups, with target group stickiness enabled. Requests that do not contain the stickiness cookies are routed based on the weight of each target group.

```

[
  {

```

```
    "Type": "forward",
    "ForwardConfig": {
      "TargetGroups": [
        {
          "TargetGroupArn": "arn:aws:elasticloadbalancing:us-
west-2:123456789012:targetgroup/blue-targets/73e2d6bc24d8a067",
          "Weight": 10
        },
        {
          "TargetGroupArn": "arn:aws:elasticloadbalancing:us-
west-2:123456789012:targetgroup/green-targets/09966783158cda59",
          "Weight": 20
        }
      ],
      "TargetGroupStickinessConfig": {
        "Enabled": true,
        "DurationSeconds": 1000
      }
    }
  }
]
```

Redirect actions

You can use `redirect` actions to redirect client requests from one URL to another. You can configure redirects as either temporary (HTTP 302) or permanent (HTTP 301) based on your needs.

A URI consists of the following components:

```
protocol://hostname:port/path?query
```

You must modify at least one of the following components to avoid a redirect loop: protocol, hostname, port, or path. Any components that you do not modify retain their original values.

protocol

The protocol (HTTP or HTTPS). You can redirect HTTP to HTTP, HTTP to HTTPS, and HTTPS to HTTPS. You cannot redirect HTTPS to HTTP.

hostname

The hostname. A hostname is not case-sensitive, can be up to 128 characters in length, and consists of alpha-numeric characters, wildcards (* and ?), and hyphens (-).

port

The port (1 to 65535).

path

The absolute path, starting with the leading "/". A path is case-sensitive, can be up to 128 characters in length, and consists of alpha-numeric characters, wildcards (* and ?), & (using &), and the following special characters: _-.\$/~"@:+

query

The query parameters. The maximum length is 128 characters.

You can reuse URI components of the original URL in the target URL using the following reserved keywords:

- `{protocol}` - Retains the protocol. Use in the protocol and query components.
- `{host}` - Retains the domain. Use in the hostname, path, and query components.
- `{port}` - Retains the port. Use in the port, path, and query components.
- `{path}` - Retains the path. Use in the path and query components.
- `{query}` - Retains the query parameters. Use in the query component.

When a `redirect` action is taken, the action is recorded in the access logs. For more information, see [Access log entries](#). The count of successful `redirect` actions is reported in the `HTTP_Redirect_Count` metric. For more information, see [Application Load Balancer metrics](#).

Example Example redirect actions using the console

The following rule sets up a permanent redirect to a URL that uses the HTTPS protocol and the specified port (40443), but retains the original hostname, path, and query parameters. This screen is equivalent to "https://{host}:40443/{path}?{query}".

Action types

Forward to target groups

Redirect to URL

Return fixed response

Redirect to URL | [Info](#)

Redirect client requests from one URL to another. You cannot redirect HTTPS to HTTP. To avoid a redirect loop, you must modify at least one of the following components: protocol, port, hostname or path. Components that you do not modify retain their original values.

URI parts

Full URL

Protocol : Port
To retain the original port enter #{port}.

1-65535

Custom host, path, query
Select to modify host, path and query. If no changes are made, settings from the request URL are retained.

Status code

The following rule sets up a permanent redirect to a URL that retains the original protocol, port, hostname, and query parameters, and uses the `#{path}` keyword to create a modified path. This screen is equivalent to `"#{protocol}://#{host}:#{port}/new/#{path}?#{query}"`.

Action types Forward to target groups Redirect to URL Return fixed response**Redirect to URL** | [Info](#)

Redirect client requests from one URL to another. You cannot redirect HTTPS to HTTP. To avoid a redirect loop, you must modify at least one of the following components: protocol, port, hostname or path. Components that you do not modify retain their original values.

URI parts**Full URL****Protocol : Port**

To retain the original port enter #{port}.

#{protocol} ▼

#{port}

1-65535

 Custom host, path, query

Select to modify host, path and query. If no changes are made, settings from the request URL are retained.

Host

Specify a host or retain the original host by using #{host}. Not case sensitive.

#{host}

Maximum 128 characters. Allowed characters are a-z, A-Z, 0-9; the following special characters: -,; and wildcards (* and ?). At least one "." is required. Only alphabetical characters are allowed after the final "." character.

Path

Specify a path or retain the original path by using #{path}. Case sensitive.

/new/#{path}

Maximum 128 characters. Allowed characters are a-z, A-Z, 0-9; the following special characters: _-.\$/~'"@:~; & (using &); and wildcards (* and ?).

Query - optional

Specify a query or retain the original query by using #{query}. Not case sensitive.

#{query}

Maximum 128 characters.

Status code

301 - Permanently moved ▼

Example Example redirect action for the AWS CLI

You can specify an action when you create or modify a rule. For more information, see the [create-rule](#) and [modify-rule](#) commands. The following action redirects an HTTP request to an HTTPS request on port 443, with the same host name, path, and query string as the HTTP request.

```
[
  {
    "Type": "redirect",
    "RedirectConfig": {
      "Protocol": "HTTPS",
      "Port": "443",
      "Host": "#{host}",
      "Path": "/#{path}",
      "Query": "#{query}",
      "StatusCode": "HTTP_301"
    }
  }
]
```

Rule condition types

The following are the supported condition types for a rule:

host-header

Route based on the host name of each request. For more information, see [Host conditions](#).

http-header

Route based on the HTTP headers for each request. For more information, see [HTTP header conditions](#).

http-request-method

Route based on the HTTP request method of each request. For more information, see [HTTP request method conditions](#).

path-pattern

Route based on path patterns in the request URLs. For more information, see [Path conditions](#).

query-string

Route based on key/value pairs or values in the query strings. For more information, see [Query string conditions](#).

source-ip

Route based on the source IP address of each request. For more information, see [Source IP address conditions](#).

Each rule can optionally include up to one of each of the following conditions: `host-header`, `http-request-method`, `path-pattern`, and `source-ip`. Each rule can also optionally include one or more of each of the following conditions: `http-header` and `query-string`.

You can specify up to three match evaluations per condition. For example, for each `http-header` condition, you can specify up to three strings to be compared to the value of the HTTP header in the request. The condition is satisfied if one of the strings matches the value of the HTTP header. To require that all of the strings are a match, create one condition per match evaluation.

You can specify up to five match evaluations per rule. For example, you can create a rule with five conditions where each condition has one match evaluation.

You can include wildcard characters in the match evaluations for the `http-header`, `host-header`, `path-pattern`, and `query-string` conditions. There is a limit of five wildcard characters per rule.

Rules are applied only to visible ASCII characters; control characters (0x00 to 0x1f and 0x7f) are excluded.

For demos, see [Advanced request routing](#).

HTTP header conditions

You can use HTTP header conditions to configure rules that route requests based on the HTTP headers for the request. You can specify the names of standard or custom HTTP header fields. The header name and the match evaluation are not case-sensitive. The following wildcard characters are supported in the comparison strings: `*` (matches 0 or more characters) and `?` (matches exactly 1 character). Wildcard characters are not supported in the header name.

Example Example HTTP header condition for the AWS CLI

You can specify conditions when you create or modify a rule. For more information, see the [create-rule](#) and [modify-rule](#) commands. The following condition is satisfied by requests with a User-Agent header that matches one of the specified strings.

```
[
  {
    "Field": "http-header",
    "HTTPHeaderConfig": {
      "HTTPHeaderName": "User-Agent",
      "Values": ["*Chrome*", "*Safari*"]
    }
  }
]
```

```
}  
]
```

HTTP request method conditions

You can use HTTP request method conditions to configure rules that route requests based on the HTTP request method of the request. You can specify standard or custom HTTP methods. The match evaluation is case-sensitive. Wildcard characters are not supported; therefore, the method name must be an exact match.

We recommend that you route GET and HEAD requests in the same way, because the response to a HEAD request may be cached.

Example Example HTTP method condition for the AWS CLI

You can specify conditions when you create or modify a rule. For more information, see the [create-rule](#) and [modify-rule](#) commands. The following condition is satisfied by requests that use the specified method.

```
[  
  {  
    "Field": "http-request-method",  
    "HttpRequestMethodConfig": {  
      "Values": ["CUSTOM-METHOD"]  
    }  
  }  
]
```

Host conditions

You can use host conditions to define rules that route requests based on the host name in the host header (also known as *host-based routing*). This enables you to support multiple subdomains and different top-level domains using a single load balancer.

A hostname is not case-sensitive, can be up to 128 characters in length, and can contain any of the following characters:

- A–Z, a–z, 0–9
- - .
- * (matches 0 or more characters)

- ? (matches exactly 1 character)

You must include at least one "." character. You can include only alphabetical characters after the final "." character.

Example hostnames

- **example.com**
- **test.example.com**
- ***.example.com**

The rule ***.example.com** matches **test.example.com** but doesn't match **example.com**.

Example host header condition for the AWS CLI

You can specify conditions when you create or modify a rule. For more information, see the [create-rule](#) and [modify-rule](#) commands. The following condition is satisfied by requests with a host header that matches the specified string.

```
[
  {
    "Field": "host-header",
    "HostHeaderConfig": {
      "Values": ["*.example.com"]
    }
  }
]
```

Path conditions

You can use path conditions to define rules that route requests based on the URL in the request (also known as *path-based routing*).

The path pattern is applied only to the path of the URL, not to its query parameters. It is applied only to visible ASCII characters; control characters (0x00 to 0x1f and 0x7f) are excluded.

The rule evaluation is performed only after URI normalization occurs.

A path pattern is case-sensitive, can be up to 128 characters in length, and can contain any of the following characters.

- A–Z, a–z, 0–9
- _ - . \$ / ~ " ' @ : +
- & (using &#38;)
- * (matches 0 or more characters)
- ? (matches exactly 1 character)

If the protocol version is gRPC, conditions can be specific to a package, service, or method.

Example HTTP path patterns

- /img/*
- /img/*/pics

Example gRPC path patterns

- /package
- /package.service
- /package.service/method

The path pattern is used to route requests but does not alter them. For example, if a rule has a path pattern of /img/*, the rule forwards a request for /img/picture.jpg to the specified target group as a request for /img/picture.jpg.

Example Example path pattern condition for the AWS CLI

You can specify conditions when you create or modify a rule. For more information, see the [create-rule](#) and [modify-rule](#) commands. The following condition is satisfied by requests with a URL that contains the specified string.

```
[
  {
    "Field": "path-pattern",
    "PathPatternConfig": {
      "Values": ["/img/*"]
    }
  }
]
```


Query string conditions

You can use query string conditions to configure rules that route requests based on key/value pairs or values in the query string. The match evaluation is not case-sensitive. The following wildcard characters are supported: * (matches 0 or more characters) and ? (matches exactly 1 character).

Example Example query string condition for the AWS CLI

You can specify conditions when you create or modify a rule. For more information, see the [create-rule](#) and [modify-rule](#) commands. The following condition is satisfied by requests with a query string that includes either a key/value pair of "version=v1" or any key set to "example".

```
[
  {
    "Field": "query-string",
    "QueryStringConfig": {
      "Values": [
        {
          "Key": "version",
          "Value": "v1"
        },
        {
          "Value": "*example*"
        }
      ]
    }
  }
]
```

Source IP address conditions

You can use source IP address conditions to configure rules that route requests based on the source IP address of the request. The IP address must be specified in CIDR format. You can use both IPv4 and IPv6 addresses. Wildcard characters are not supported. You cannot specify the 255.255.255.255/32 CIDR for the source IP rule condition.

If a client is behind a proxy, this is the IP address of the proxy, not the IP address of the client.

This condition is not satisfied by the addresses in the X-Forwarded-For header. To search for addresses in the X-Forwarded-For header, use an `http-header` condition.

Example Example source IP condition for the AWS CLI

You can specify conditions when you create or modify a rule. For more information, see the [create-rule](#) and [modify-rule](#) commands. The following condition is satisfied by requests with a source IP address in one of the specified CIDR blocks.

```
[
  {
    "Field": "source-ip",
    "SourceIpConfig": {
      "Values": ["192.0.2.0/24", "198.51.100.10/32"]
    }
  }
]
```

Create an HTTP listener for your Application Load Balancer

A listener checks for connection requests. You define a listener when you create your load balancer, and you can add listeners to your load balancer at any time.

The information on this page helps you create an HTTP listener for your load balancer. To add an HTTPS listener to your load balancer, see [Create an HTTPS listener for your Application Load Balancer](#).

Prerequisites

- To add a forward action to the default listener rule, you must specify an available target group. For more information, see [Create a target group](#).
- You can specify the same target group in multiple listeners, but these listeners must belong to the same load balancer. To use a target group with a load balancer, you must verify that it is not used by a listener for any other load balancer.

Add an HTTP listener

You configure a listener with a protocol and a port for connections from clients to the load balancer, and a target group for the default listener rule. For more information, see [Listener configuration](#).

To add an HTTP listener using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Listeners and rules** tab, choose **Add listener**.
5. For **Protocol : Port**, choose **HTTP** and keep the default port or enter a different port.
6. For **Default actions**, choose one of the following:
 - **Forward to target groups** – Choose one or more target groups to forward traffic to. To add target groups choose **Add target group**. If using more than one target group, select a weight for each target group and review the associated percentage. You must enable group-level stickiness on a rule, if you've enabled stickiness on one or more of the target groups.
 - **Redirect to URL** – Specify the URL that client requests will be redirected to. This can be done by entering each part separately on the **URI parts** tab, or by entering the full address on the **Full URL** tab. For **Status code** you can configure redirects as either temporary (HTTP 302) or permanent (HTTP 301) based on your needs.
 - **Return fixed response** – Specify the **Response code** that will be returned to dropped client requests. Additionally, you can specify the **Content type** and **Response body**, but they're not required.
7. Choose **Add**.

To add an HTTP listener using the AWS CLI

Use the [create-listener](#) command to create the listener and default rule, and the [create-rule](#) command to define additional listener rules.

Create an HTTPS listener for your Application Load Balancer

A listener checks for connection requests. You define a listener when you create your load balancer, and you can add listeners to your load balancer at any time.

To create an HTTPS listener, you must deploy at least one SSL server certificate on your load balancer. The load balancer uses a server certificate to terminate the front-end connection and then decrypt requests from clients before sending them to the targets. You must also specify

a security policy, which is used to negotiate secure connections between clients and the load balancer.

If you need to pass encrypted traffic to targets without the load balancer decrypting it, you can create a Network Load Balancer or Classic Load Balancer with a TCP listener on port 443. With a TCP listener, the load balancer passes encrypted traffic through to the targets without decrypting it.

Application Load Balancers do not support ED25519 keys.

The information on this page helps you create an HTTPS listener for your load balancer. To add an HTTP listener to your load balancer, see [Create an HTTP listener for your Application Load Balancer](#).

Contents

- [SSL certificates](#)
 - [Default certificate](#)
 - [Certificate list](#)
 - [Certificate renewal](#)
- [Security policies](#)
 - [TLS 1.3 security policies](#)
 - [FIPS security policies](#)
 - [FS supported policies](#)
 - [TLS 1.0 - 1.2 security policies](#)
 - [TLS protocols and ciphers](#)
- [Add an HTTPS listener](#)

SSL certificates

The load balancer requires X.509 certificates (SSL/TLS server certificates). Certificates are a digital form of identification issued by a certificate authority (CA). A certificate contains identification information, a validity period, a public key, a serial number, and the digital signature of the issuer.

When you create a certificate for use with your load balancer, you must specify a domain name. The domain name on the certificate must match the custom domain name record so that we can verify the TLS connection. If they do not match, the traffic is not encrypted.

You must specify a fully qualified domain name (FQDN) for your certificate, such as `www.example.com` or an apex domain name such as `example.com`. You can also use an asterisk (*) as a wild card to protect several site names in the same domain. When you request a wild-card certificate, the asterisk (*) must be in the leftmost position of the domain name and can protect only one subdomain level. For instance, `*.example.com` protects `corp.example.com`, and `images.example.com`, but it cannot protect `test.login.example.com`. Also note that `*.example.com` protects only the subdomains of `example.com`, it does not protect the bare or apex domain (`example.com`). The wild-card name appears in the **Subject** field and in the **Subject Alternative Name** extension of the certificate. For more information about public certificates, see [Requesting a public certificate](#) in the *AWS Certificate Manager User Guide*.

We recommend that you create certificates for your load balancer using [AWS Certificate Manager \(ACM\)](#). ACM supports RSA certificates with 2048, 3072, and 4096-bit key lengths, and all ECDSA certificates. ACM integrates with Elastic Load Balancing so that you can deploy the certificate on your load balancer. For more information, see the [AWS Certificate Manager User Guide](#).

Alternatively, you can use SSL/TLS tools to create a certificate signing request (CSR), then get the CSR signed by a CA to produce a certificate, then import the certificate into ACM or upload the certificate to AWS Identity and Access Management (IAM). For more information about importing certificates into ACM, see [Importing certificates](#) in the *AWS Certificate Manager User Guide*. For more information about uploading certificates to IAM, see [Working with server certificates](#) in the *IAM User Guide*.

Default certificate

When you create an HTTPS listener, you must specify exactly one certificate. This certificate is known as the *default certificate*. You can replace the default certificate after you create the HTTPS listener. For more information, see [Replace the default certificate](#).

If you specify additional certificates in a [certificate list](#), the default certificate is used only if a client connects without using the Server Name Indication (SNI) protocol to specify a hostname or if there are no matching certificates in the certificate list.

If you do not specify additional certificates but need to host multiple secure applications through a single load balancer, you can use a wildcard certificate or add a Subject Alternative Name (SAN) for each additional domain to your certificate.

Certificate list

After you create an HTTPS listener, it has a default certificate and an empty certificate list. You can optionally add certificates to the certificate list for the listener. Using a certificate list enables the load balancer to support multiple domains on the same port and provide a different certificate for each domain. For more information, see [Add certificates to the certificate list](#).

The load balancer uses a smart certificate selection algorithm with support for SNI. If the hostname provided by a client matches a single certificate in the certificate list, the load balancer selects this certificate. If a hostname provided by a client matches multiple certificates in the certificate list, the load balancer selects the best certificate that the client can support. Certificate selection is based on the following criteria in the following order:

- Public key algorithm (prefer ECDSA over RSA)
- Hashing algorithm (prefer SHA over MD5)
- Key length (prefer the largest)
- Validity period

The load balancer access log entries indicate the hostname specified by the client and the certificate presented to the client. For more information, see [Access log entries](#).

Certificate renewal

Each certificate comes with a validity period. You must ensure that you renew or replace each certificate for your load balancer before its validity period ends. This includes the default certificate and certificates in a certificate list. Renewing or replacing a certificate does not affect in-flight requests that were received by the load balancer node and are pending routing to a healthy target. After a certificate is renewed, new requests use the renewed certificate. After a certificate is replaced, new requests use the new certificate.

You can manage certificate renewal and replacement as follows:

- Certificates provided by AWS Certificate Manager and deployed on your load balancer can be renewed automatically. ACM attempts to renew certificates before they expire. For more information, see [Managed renewal](#) in the *AWS Certificate Manager User Guide*.
- If you imported a certificate into ACM, you must monitor the expiration date of the certificate and renew it before it expires. For more information, see [Importing certificates](#) in the *AWS Certificate Manager User Guide*.

- If you imported a certificate into IAM, you must create a new certificate, import the new certificate to ACM or IAM, add the new certificate to your load balancer, and remove the expired certificate from your load balancer.

Security policies

Elastic Load Balancing uses a Secure Socket Layer (SSL) negotiation configuration, known as a security policy, to negotiate SSL connections between a client and the load balancer. A security policy is a combination of protocols and ciphers. The protocol establishes a secure connection between a client and a server and ensures that all data passed between the client and your load balancer is private. A cipher is an encryption algorithm that uses encryption keys to create a coded message. Protocols use several ciphers to encrypt data over the internet. During the connection negotiation process, the client and the load balancer present a list of ciphers and protocols that they each support, in order of preference. By default, the first cipher on the server's list that matches any one of the client's ciphers is selected for the secure connection.

Considerations:

- Application Load Balancers support SSL renegotiation for target connections only.
- Application Load Balancers do not support custom security policies.
- The `ELBSecurityPolicy-TLS13-1-2-2021-06` policy is the default security policy for HTTPS listeners created using the AWS Management Console.
- The `ELBSecurityPolicy-2016-08` policy is the default security policy for HTTPS listeners created using the AWS CLI.
- When you create an HTTPS listener, selecting a security policy is required.
 - We recommend the `ELBSecurityPolicy-TLS13-1-2-2021-06` security policy, which includes TLS 1.3, and is backwards compatible with TLS 1.2.
- You can choose the security policy that is used for front-end connections, but not backend connections.
 - For backend connections, if any of your HTTPS listeners are using a TLS 1.3 security policy, the `ELBSecurityPolicy-TLS13-1-0-2021-06` security policy is used. Otherwise, the `ELBSecurityPolicy-2016-08` security policy is used for backend connections.
- To meet compliance and security standards that require disabling certain TLS protocol versions, or to support legacy clients requiring deprecated ciphers, you can use one of the `ELBSecurityPolicy-TLS-` security policies. To view the TLS protocol version for requests to

your Application Load Balancer, enable access logging for your load balancer and examine the corresponding access log entries. For more information, see [Access logs for your Application Load Balancer](#).

- You can restrict which security policies are available to users across your AWS accounts and AWS Organizations by using the [Elastic Load Balancing condition keys](#) in your IAM and service control policies (SCPs), respectively. For more information, see [Service control policies \(SCPs\)](#) in the *AWS Organizations User Guide*
- Application Load Balancers support TLS resumption using PSK (TLS 1.3) and session IDs/session Tickets (TLS 1.2 and older). Resumptions are only supported in connections to the same Application Load Balancer IP address. The 0-RTT Data feature and early_data extension are not implemented.

TLS 1.3 security policies

Elastic Load Balancing provides the following TLS 1.3 security policies for Application Load Balancers:

- ELBSecurityPolicy-TLS13-1-2-2021-06 **(Recommended)**
- ELBSecurityPolicy-TLS13-1-2-Res-2021-06
- ELBSecurityPolicy-TLS13-1-2-Ext1-2021-06
- ELBSecurityPolicy-TLS13-1-2-Ext2-2021-06
- ELBSecurityPolicy-TLS13-1-1-2021-06
- ELBSecurityPolicy-TLS13-1-0-2021-06
- ELBSecurityPolicy-TLS13-1-3-2021-06

FIPS security policies

Important

All secure listeners attached to an Application Load Balancer must use either FIPS security policies or non-FIPS security policies; they cannot be mixed. If an existing Application Load Balancer has two or more listeners using non-FIPS policies and you want the listeners to use FIPS security policies instead, remove all listeners until there is only one. Change the security policy of the listener to FIPS and then create additional listeners using FIPS

security policies. Alternatively, you can create a new Application Load Balancer with new listeners using only FIPS security policies.

The Federal Information Processing Standard (FIPS) is a US and Canadian government standard that specifies the security requirements for cryptographic modules that protect sensitive information. To learn more, see [Federal Information Processing Standard \(FIPS\) 140](#) on the *AWS Cloud Security Compliance* page.

All FIPS policies leverage the AWS-LC FIPS validated cryptographic module. To learn more, see the [AWS-LC Cryptographic Module](#) page on the *NIST Cryptographic Module Validation Program* site.

Elastic Load Balancing provides the following FIPS security policies for Application Load Balancers:

- ELBSecurityPolicy-TLS13-1-3-FIPS-2023-04
- ELBSecurityPolicy-TLS13-1-2-Res-FIPS-2023-04
- ELBSecurityPolicy-TLS13-1-2-FIPS-2023-04 **(Recommended)**
- ELBSecurityPolicy-TLS13-1-2-Ext0-FIPS-2023-04
- ELBSecurityPolicy-TLS13-1-2-Ext1-FIPS-2023-04
- ELBSecurityPolicy-TLS13-1-2-Ext2-FIPS-2023-04
- ELBSecurityPolicy-TLS13-1-1-FIPS-2023-04
- ELBSecurityPolicy-TLS13-1-0-FIPS-2023-04

FS supported policies

Elastic Load Balancing provides the following FS (Forward Secrecy) supported security policies for Application Load Balancers:

- ELBSecurityPolicy-FS-1-2-Res-2020-10
- ELBSecurityPolicy-FS-1-2-Res-2019-08
- ELBSecurityPolicy-FS-1-2-2019-08
- ELBSecurityPolicy-FS-1-1-2019-08
- ELBSecurityPolicy-FS-2018-06

TLS 1.0 - 1.2 security policies

Elastic Load Balancing provides the following TLS 1.0 - 1.2 security policies for Application Load Balancers:

- ELBSecurityPolicy-TLS-1-2-Ext-2018-06
- ELBSecurityPolicy-TLS-1-2-2017-01
- ELBSecurityPolicy-TLS-1-1-2017-01
- ELBSecurityPolicy-2016-08
- ELBSecurityPolicy-TLS-1-0-2015-04
- ELBSecurityPolicy-2015-05 (**identical to ELBSecurityPolicy-2016-08**)

TLS protocols and ciphers

TLS 1.3

The following table describes the supported TLS protocols and ciphers for the available TLS 1.3 security policies.

Note: The ELBSecurityPolicy- prefix has been removed from the policy names in the security policies row.

Example: Security policy ELBSecurityPolicy-TLS13-1-2-2021-06 is displayed as TLS13-1-2-2021-06.

Security policies	TLS13-1-2-2021-06	TLS13-1-3-2021-06	TLS13-1-2-Res-2021-06	TLS13-1-2-Ext2-2021-06	TLS13-1-2-Ext1-2021-06	TLS13-1-1-2021-06	TLS13-1-0-2021-06
TLS Protocols							
Protocol- TLSv1							✓

Security policies	TLS13-1-2-2021-06	TLS13-1-3-2021-06	TLS13-1-2-Res-2021-06	TLS13-1-2-Ext2-2021-06	TLS13-1-2-Ext1-2021-06	TLS13-1-1-2021-06	TLS13-1-0-2021-06
Protocol-TLSv1.1						✓	✓
Protocol-TLSv1.2	✓		✓	✓	✓	✓	✓
Protocol-TLSv1.3	✓	✓	✓	✓	✓	✓	✓
TLS Ciphers							
TLS_AES_128_GCM_SHA256	✓	✓	✓	✓	✓	✓	✓
TLS_AES_256_GCM_SHA384	✓	✓	✓	✓	✓	✓	✓
TLS_CHACHA20_POLY1305_SHA256	✓	✓	✓	✓	✓	✓	✓
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	✓		✓	✓	✓	✓	✓

Security policies	TLS13-1-2-2021-06	TLS13-1-3-2021-06	TLS13-1-2-Res-2021-06	TLS13-1-2-Ext2-2021-06	TLS13-1-2-Ext1-2021-06	TLS13-1-1-1-2021-06	TLS13-1-0-2021-06
ECDHE- RSA- AES128- GCM- SHA256	✓		✓	✓	✓	✓	✓
ECDHE- ECDSA- AES128- SHA256	✓			✓	✓	✓	✓
ECDHE- RSA- AES128- SHA256	✓			✓	✓	✓	✓
ECDHE- ECDSA- AES128- SHA				✓		✓	✓
ECDHE- RSA- AES128- SHA				✓		✓	✓
ECDHE- ECDSA- AES256 -GCM- SHA384	✓		✓	✓	✓	✓	✓

Security policies	TLS13-1-2-2021-06	TLS13-1-3-2021-06	TLS13-1-2-Res-2021-06	TLS13-1-2-Ext2-2021-06	TLS13-1-2-Ext1-2021-06	TLS13-1-1-1-2021-06	TLS13-1-0-2021-06
ECDHE- RSA- AES256- GCM- SHA384	✓		✓	✓	✓	✓	✓
ECDHE- ECDSA- AES256- SHA384	✓			✓	✓	✓	✓
ECDHE- RSA- AES256- SHA384	✓			✓	✓	✓	✓
ECDHE- RSA- AES256- SHA				✓		✓	✓
ECDHE- ECDSA- AES256- SHA				✓		✓	✓
AES128- GCM- SHA256				✓	✓	✓	✓

Security policies	TLS13-1-2-2021-06	TLS13-1-3-2021-06	TLS13-1-2-Res-2021-06	TLS13-1-2-Ext2-2021-06	TLS13-1-2-Ext1-2021-06	TLS13-1-1-2021-06	TLS13-1-0-2021-06
AES128-SHA256				✓	✓	✓	✓
AES128-SHA				✓		✓	✓
AES256-GCM-SHA384				✓	✓	✓	✓
AES256-SHA256				✓	✓	✓	✓
AES256-SHA				✓		✓	✓

To create an HTTPS listener that uses a TLS 1.3 policy using the CLI

Use the [create-listener](#) command with any [TLS 1.3 security policy](#).

The example uses the `ELBSecurityPolicy-TLS13-1-2-2021-06` security policy.

```
aws elbv2 create-listener --name my-listener \
--protocol HTTPS --port 443 \
--ssl-policy ELBSecurityPolicy-TLS13-1-2-2021-06
```

To modify an HTTPS listener to use a TLS 1.3 policy using the CLI

Use the [modify-listener](#) command with any [TLS 1.3 security policy](#).

The example uses the `ELBSecurityPolicy-TLS13-1-2-2021-06` security policy.

```
aws elbv2 modify-listener \
```

```
--listener-arn arn:aws:elasticloadbalancing:us-east-1:123456789012:listener/app/my-load-balancer/abcdef01234567890/1234567890abcdef0 \  
--ssl-policy ELBSecurityPolicy-TLS13-1-2-2021-06
```

To view the security policies used by a listener using the CLI

Use the [describe-listeners](#) command with the arn of your listener.

```
aws elbv2 describe-listeners \  
--listener-arn arn:aws:elasticloadbalancing:us-east-1:123456789012:listener/app/my-load-balancer/abcdef01234567890/1234567890abcdef0
```

To view the configuration of a TLS 1.3 security policy using the CLI

Use the [describe-ssl-policies](#) command with any [TLS 1.3 security policy](#).

The example uses the `ELBSecurityPolicy-TLS13-1-2-2021-06` security policy.

```
aws elbv2 describe-ssl-policies \  
--names ELBSecurityPolicy-TLS13-1-2-2021-06
```

FIPS

Important

Policies `ELBSecurityPolicy-TLS13-1-1-FIPS-2023-04` and `ELBSecurityPolicy-TLS13-1-0-FIPS-2023-04` are provided for legacy compatibility only. While they utilize FIPS cryptography using the FIPS140 module, they may not conform to the latest NIST guidance for TLS configuration.

The following table describes the supported TLS protocols and ciphers for the available FIPS security policies.

Note: The `ELBSecurityPolicy-` prefix has been removed from the policy names in the security policies row.

Example: Security policy `ELBSecurityPolicy-TLS13-1-2-FIPS-2023-04` is displayed as `TLS13-1-2-FIPS-2023-04`.

Security policies

TLS13-1-3-FIPS-2023-04

TLS13-1-2-Res-FIPS-2023-04

TLS13-1-2-FIPS-2023-04

TLS13-1-2-Ext0-FIPS-2023-04

TLS13-1-2-Ext1-FIPS-2023-04

TLS13-1-2-Ext2-FIPS-2023-04

TLS13-1-1-FIPS-2023-04

TLS13-1-0-FIPS-2023-04

TLS ProtocolsProtocol-
TLSv1

✓

Protocol-
TLSv1.1

✓

✓

Protocol-
TLSv1.2

✓

✓

✓

✓

✓

✓

✓

Protocol- ✓
TLSv1.3

✓

✓

✓

✓

✓

✓

✓

TLS Ciphers

TLS_AES_128_GCM_SHA256

✓

✓

✓

✓

✓

✓

✓

TLS_AES_256_GCM_SHA384

✓

✓

✓

✓

✓

✓

✓

ECDHE-ECDSA-AES128-GCM-SHA256

✓

✓

✓

✓

✓

✓

✓

Security policies

	TLS13-1-3-FIPS-2023-04	TLS13-1-2-Res-FIPS-2023-04	TLS13-1-2-FIPS-2023-04	TLS13-1-2-Ext0-FIPS-2023-04	TLS13-1-2-Ext1-FIPS-2023-04	TLS13-1-2-Ext2-FIPS-2023-04	TLS13-1-1-FIPS-2023-04	TLS13-1-0-FIPS-2023-04
ECDHE-RSA-AES128-GCM-SHA256		✓	✓	✓	✓	✓	✓	✓
ECDHE-ECDSA-AES128-SHA256			✓	✓	✓	✓	✓	✓
ECDHE-RSA-AES128-SHA256			✓	✓	✓	✓	✓	✓
ECDHE-ECDSA-AES128-SHA				✓		✓	✓	✓
ECDHE-RSA-AES128-SHA				✓		✓	✓	✓

Security policies

	TLS13-1-3-FIPS-2023-04	TLS13-1-2-Res-FIPS-2023-04	TLS13-1-2-FIPS-2023-04	TLS13-1-2-Ext0-FIPS-2023-04	TLS13-1-2-Ext1-FIPS-2023-04	TLS13-1-2-Ext2-FIPS-2023-04	TLS13-1-1-FIPS-2023-04	TLS13-1-0-FIPS-2023-04
ECDHE- ECDSA- AES256 -GCM- SHA3 84		✓	✓	✓	✓	✓	✓	✓
ECDHE- RSA- AES256- GCM- SHA384	✓	✓	✓	✓	✓	✓	✓	✓
ECDHE- ECDSA- AES256 - SHA384			✓	✓	✓	✓	✓	✓
ECDHE- RSA- AES256- S HA384			✓	✓	✓	✓	✓	✓
ECDHE- RSA- AES256- SHA				✓		✓	✓	✓

Security policies	TLS13-1-3-FIPS-2023-04	TLS13-1-2-Res-FIPS-2023-04	TLS13-1-2-FIPS-2023-04	TLS13-1-2-Ext0-FIPS-2023-04	TLS13-1-2-Ext1-FIPS-2023-04	TLS13-1-2-Ext2-FIPS-2023-04	TLS13-1-1-FIPS-2023-04	TLS13-1-0-FIPS-2023-04
ECDHE- ECDSA- AES256 -SHA				✓		✓	✓	✓
AES128- GCM- SHA256					✓	✓	✓	✓
AES128- SH A256					✓	✓	✓	✓
AES128- SHA						✓	✓	✓
AES256- GCM- SHA384					✓	✓	✓	✓
AES256- SH A256					✓	✓	✓	✓
AES256- SHA						✓	✓	✓

To create an HTTPS listener that uses a FIPS policy using the CLI

Use the [create-listener](#) command with any [FIPS security policy](#).

The example uses the `ELBSecurityPolicy-TLS13-1-2-FIPS-2023-04` security policy.

```
aws elbv2 create-listener --name my-listener \
--protocol HTTPS --port 443 \
--ssl-policy ELBSecurityPolicy-TLS13-1-2-FIPS-2023-04
```

To modify an HTTPS listener to use a FIPS policy using the CLI

Use the [modify-listener](#) command with any [FIPS security policy](#).

The example uses the `ELBSecurityPolicy-TLS13-1-2-FIPS-2023-04` security policy.

```
aws elbv2 modify-listener \
--listener-arn arn:aws:elasticloadbalancing:us-east-1:123456789012:listener/app/my-load-balancer/abcdef01234567890/1234567890abcdef0 \
--ssl-policy ELBSecurityPolicy-TLS13-1-2-FIPS-2023-04
```

To view the security policies used by a listener using the CLI

Use the [describe-listeners](#) command with the arn of your listener.

```
aws elbv2 describe-listeners \
--listener-arn arn:aws:elasticloadbalancing:us-east-1:123456789012:listener/app/my-load-balancer/abcdef01234567890/1234567890abcdef0
```

To view the configuration of a FIPS security policy using the CLI

Use the [describe-ssl-policies](#) command with any [FIPS security policy](#).

The example uses the `ELBSecurityPolicy-TLS13-1-2-FIPS-2023-04` security policy.

```
aws elbv2 describe-ssl-policies \
--names ELBSecurityPolicy-TLS13-1-2-FIPS-2023-04
```

FS

The following table describes the supported TLS protocols and ciphers for the available FS supported security policies.

Note: The `ELBSecurityPolicy-` prefix has been removed from the policy names in the security policies row.

Example: Security policy `ELBSecurityPolicy-FS-2018-06` is displayed as `FS-2018-06`.

Security policies	Default	FS-1-2-Res-2020-10	FS-1-2-Res-2019-08	FS-1-2-2019-08	FS-1-1-2019-08	FS-2018-06
TLS Protocols						
Protocol-TLSv1	✓					✓
Protocol-TLSv1.1	✓				✓	✓
Protocol-TLSv1.2	✓	✓	✓	✓	✓	✓
TLS Ciphers						
ECDHE-ECDSA-AES128-GCM-SHA256	✓	✓	✓	✓	✓	✓
ECDHE-RSA-AES128-GCM-SHA256	✓	✓	✓	✓	✓	✓
ECDHE-ECDSA-AES128-SHA256	✓		✓	✓	✓	✓

Security policies	Default	FS-1-2-Res-2020-10	FS-1-2-Res-2019-08	FS-1-2-2019-08	FS-1-1-2019-08	FS-2018-06
ECDHE-RSA-AES128-SHA256	✓		✓	✓	✓	✓
ECDHE-ECDSA-AES128-SHA	✓			✓	✓	✓
ECDHE-RSA-AES128-SHA	✓			✓	✓	✓
ECDHE-ECDSA-AES256-GCM-SHA384	✓	✓	✓	✓	✓	✓
ECDHE-RSA-AES256-GCM-SHA384	✓	✓	✓	✓	✓	✓

Security policies	Default	FS-1-2-Res-2020-10	FS-1-2-Res-2019-08	FS-1-2-2019-08	FS-1-1-2019-08	FS-2018-06
ECDHE- ECDSA- AES256- SHA384	✓		✓	✓	✓	✓
ECDHE- RSA- AES256-S HA384	✓		✓	✓	✓	✓
ECDHE- RSA- AES256-S HA	✓			✓	✓	✓
ECDHE- ECDSA- AES256- SHA	✓			✓	✓	✓
AES128- GCM- SHA256	✓					
AES128- SHA256	✓					
AES128- SHA	✓					

Security policies	Default	FS-1-2-Res-2020-10	FS-1-2-Res-2019-08	FS-1-2-2019-08	FS-1-1-2019-08	FS-2018-06
AES256-GCM-SHA384	✓					
AES256-SHA256	✓					
AES256-SHA	✓					

To create an HTTPS listener that uses a FS supported policy using the CLI

Use the [create-listener](#) command with any [FS supported security policy](#).

The example uses the `ELBSecurityPolicy-FS-2018-06` security policy.

```
aws elbv2 create-listener --name my-listener \
--protocol HTTPS --port 443 \
--ssl-policy ELBSecurityPolicy-FS-2018-06
```

To modify an HTTPS listener to use a FS supported policy using the CLI

Use the [modify-listener](#) command with any [FS supported security policy](#).

The example uses the `ELBSecurityPolicy-FS-2018-06` security policy.

```
aws elbv2 modify-listener \
--listener-arn arn:aws:elasticloadbalancing:us-east-1:123456789012:listener/app/my-load-balancer/abcdef01234567890/1234567890abcdef0 \
--ssl-policy ELBSecurityPolicy-FS-2018-06
```


To view the security policies used by a listener using the CLI

Use the [describe-listeners](#) command with the arn of your listener.

```
aws elbv2 describe-listeners \
--listener-arn arn:aws:elasticloadbalancing:us-east-1:123456789012:listener/app/my-load-balancer/abcdef01234567890/1234567890abcdef0
```

To view the configuration of a FS supported security policy using the CLI

Use the [describe-ssl-policies](#) command with any [FS supported security policy](#).

The example uses the ELBSecurityPolicy-FS-2018-06 security policy.

```
aws elbv2 describe-ssl-policies \
--names ELBSecurityPolicy-FS-2018-06
```

TLS 1.0 - 1.2

The following table describes the supported TLS protocols and ciphers for the available TLS 1.0-1.2 security policies.

Note: The ELBSecurityPolicy- prefix has been removed from the policy names in the security policies row.

Example: Security policy ELBSecurityPolicy-TLS-1-2-Ext-2018-06 is displayed as TLS-1-2-Ext-2018-06.

Security policies	Default	TLS-1-2-Ext-2018-06	TLS-1-2-2017-01	TLS-1-1-2017-01	TLS-1-0-2015-04*
TLS Protocols					
Protocol-TLSv1	✓				✓

Security policies	Default	TLS-1-2-Ext-2018-06	TLS-1-2-2017-01	TLS-1-1-2017-01	TLS-1-0-2015-04*
Protocol-TLSv1.1	✓			✓	✓
Protocol-TLSv1.2	✓	✓	✓	✓	✓
TLS Ciphers					
ECDHE-ECD SA-AES128 -GCM-SHA2 56	✓	✓	✓	✓	✓
ECDHE-RSA -AES128-G CM-SHA256	✓	✓	✓	✓	✓
ECDHE-ECD SA-AES128- SHA256	✓	✓	✓	✓	✓
ECDHE-RSA -AES128-S HA256	✓	✓	✓	✓	✓
ECDHE-ECD SA-AES128- SHA	✓	✓		✓	✓

Security policies	Default	TLS-1-2-Ext-2018-06	TLS-1-2-2017-01	TLS-1-1-2017-01	TLS-1-0-2015-04*
ECDHE-RSA-AES128-SHA	✓	✓		✓	✓
ECDHE-ECD SA-AES256 -GCM-SHA3 84	✓	✓	✓	✓	✓
ECDHE-RSA -AES256-G CM-SHA384	✓	✓	✓	✓	✓
ECDHE-ECD SA-AES256- SHA384	✓	✓	✓	✓	✓
ECDHE-RSA -AES256-S HA384	✓	✓	✓	✓	✓
ECDHE-RSA- AES256-SHA	✓	✓		✓	✓
ECDHE-ECD SA-AES256- SHA	✓	✓		✓	✓
AES128-GC M-SHA256	✓	✓	✓	✓	✓

Security policies	Default	TLS-1-2-Ext-2018-06	TLS-1-2-2017-01	TLS-1-1-2017-01	TLS-1-0-2015-04*
AES128-SH A256	✓	✓	✓	✓	✓
AES128-SHA	✓	✓		✓	✓
AES256-GC M-SHA384	✓	✓	✓	✓	✓
AES256-SH A256	✓	✓	✓	✓	✓
AES256-SHA	✓	✓		✓	✓
DES-CBC3- SHA					✓

* Do not use this policy unless you must support a legacy client that requires the DES-CBC3-SHA cipher, which is a weak cipher.

To create an HTTPS listener that uses a TLS 1.0-1.2 policy using the CLI

Use the [create-listener](#) command with any [TLS 1.0-1.2 supported security policy](#).

The example uses the `ELBSecurityPolicy-2016-08` security policy.

```
aws elbv2 create-listener --name my-listener \  
--protocol HTTPS --port 443 \  
--ssl-policy ELBSecurityPolicy-2016-08
```

To modify an HTTPS listener to use a TLS 1.0-1.2 policy using the CLI

Use the [modify-listener](#) command with any [TLS 1.0-1.2 supported security policy](#).

The example uses the `ELBSecurityPolicy-2016-08` security policy.

```
aws elbv2 modify-listener \  
--listener-arn arn:aws:elasticloadbalancing:us-east-1:123456789012:listener/app/my-load-balancer/abcdef01234567890/1234567890abcdef0 \  
--ssl-policy ELBSecurityPolicy-2016-08
```

To view the security policies used by a listener using the CLI

Use the [describe-listeners](#) command with the arn of your listener.

```
aws elbv2 describe-listeners \  
--listener-arn arn:aws:elasticloadbalancing:us-east-1:123456789012:listener/app/my-load-balancer/abcdef01234567890/1234567890abcdef0
```

To view the configuration of a TLS 1.0-1.2 security policy using the CLI

Use the [describe-ssl-policies](#) command with any [TLS 1.0-1.2 supported security policy](#).

The example uses the `ELBSecurityPolicy-2016-08` security policy.

```
aws elbv2 describe-ssl-policies \  
--names ELBSecurityPolicy-2016-08
```

Add an HTTPS listener

You configure a listener with a protocol and a port for connections from clients to the load balancer, and a target group for the default listener rule. For more information, see [Listener configuration](#).

Prerequisites

- To create an HTTPS listener, you must specify a certificate and a security policy. The load balancer uses the certificate to terminate the connection and decrypt requests from clients before routing them to targets. The load balancer uses the security policy when negotiating SSL connections with the clients.

- To add a forward action to the default listener rule, you must specify an available target group. For more information, see [Create a target group](#).
- You can specify the same target group in multiple listeners, but these listeners must belong to the same load balancer. To use a target group with a load balancer, you must verify that it is not used by a listener for any other load balancer.

To add an HTTPS listener using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Listeners and rules** tab, choose **Add listener**.
5. For **Protocol : Port**, choose **HTTPS** and keep the default port or enter a different port.
6. (Optional) To enable authentication, under **Authentication** select **Use OpenID or Amazon Cognito**, and provide the requested information. For more information, see [Authenticate users using an Application Load Balancer](#).
7. For **Default actions**, do one of the following:
 - **Forward to target groups** – Choose one or more target groups to forward traffic to. To add target groups choose **Add target group**. If using more than one target group, select a weight for each target group and review the associated percentage. You must enable group-level stickiness on a rule, if you've enabled stickiness on one or more of the target groups.
 - **Redirect to URL** – Specify the URL that client requests will be redirected to. This can be done by entering each part separately on the **URI parts** tab, or by entering the full address on the **Full URL** tab. For **Status code** you can configure redirects as either temporary (HTTP 302) or permanent (HTTP 301) based on your needs.
 - **Return fixed response** – Specify the **Response code** that will be returned to dropped client requests. Additionally, you can specify the **Content type** and **Response body**, but they're not required.
8. For **Security policy**, we recommend that you always use the latest predefined security policy.
9. For **Default SSL/TLS certificate**, the following options are available:
 - If you created or imported a certificate using AWS Certificate Manager, select **From ACM**, then select the certificate from **Select a certificate**.

- If you imported a certificate using IAM, select **From IAM**, and then select your certificate from **Select a certificate**.
 - If you have a certificate to import but ACM is not available in your Region, select **Import**, then select **To IAM**. Type the name of the certificate in the **Certificate name** field. In **Certificate private key**, copy and paste the contents of the private key file (PEM-encoded). In **Certificate body**, copy and paste the contents of the public key certificate file (PEM-encoded). In **Certificate Chain**, copy and paste the contents of the certificate chain file (PEM-encoded), unless you are using a self-signed certificate and it's not important that browsers implicitly accept the certificate.
10. (Optional) To enable mutual authentication, under **Client certificate handling** enable **Mutual authentication (mTLS)**.

When enabled, the default mutual TLS mode is **passthrough**.

If you select **Verify with Trust Store**:

- By default, connections with expired client certificates are rejected. To change this behavior expand **Advanced mTLS settings**, then under **Client certificate expiration** select **Allow expired client certificates**.
 - Under **Trust Store** choose an existing trust store, or choose **New trust store**.
 - If you chose **New trust store**, provide a **Trust store name**, the **S3 URI Certificate Authority location**, and optionally an **S3 URI Certificate revocation list location**.
11. Choose **Save**.

To add an HTTPS listener using the AWS CLI

Use the [create-listener](#) command to create the listener and default rule, and the [create-rule](#) command to define additional listener rules.

Listener rules for your Application Load Balancer

The rules that you define for your listener determine how the load balancer routes requests to the targets in one or more target groups.

Each rule consists of a priority, one or more actions, and one or more conditions. For more information, see [Listener rules](#).

Requirements

- Each rule must include exactly one of the following actions: `forward`, `redirect`, or `fixed-response`, and it must be the last action to be performed.
- Each rule can include zero or one of the following conditions: `host-header`, `http-request-method`, `path-pattern`, and `source-ip`, and zero or more of the following conditions: `http-header` and `query-string`.
- You can specify up to three comparison strings per condition and up to five per rule.
- A `forward` action routes requests to its target group. Before you add a `forward` action, create the target group and add targets to it. For more information, see [Create a target group](#).

Add a rule

You define a default rule when you create a listener, and you can define additional nondefault rules at any time.

To add a rule using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer to view its details.
4. On the **Listeners and rules** tab, do one of the following:
 - a. Select the text in the **Protocol:Port** column to open the detail page for the listener.

On the **Rules** tab, choose **Add rule**.

- b. Select the listener you want to add a rule to.

Choose **Manage rules**, then **Add rule**.

5. You can specify a name for your rule under **Name and tags**, although it's not required.

To add additional tags select the **Add additional tags** text.

6. Choose **Next**.
7. Choose **Add condition**.
8. Add one or more of the following conditions:

- **Host header** – Define the host header. For example: *.example.com. To save the condition, choose **Confirm**.

Maximum 128 characters. Not case sensitive. Allowed characters are a-z, A-Z, 0-9; the following special characters: -_.; and wildcards (* and ?).

- **Path** – Define the path. For example: /item/* . To save the condition, choose **Confirm**.

Maximum 128 characters. Case sensitive. Allowed characters are a-z, A-Z, 0-9; the following special characters: _.\$/~"@:~+; &; and wildcards (* and ?).

- **HTTP request method** – Define the HTTP request method. To save the condition, choose **Confirm**.

Maximum 40 characters. Case sensitive. Allowed characters are A-Z, and the following special characters: -_. Wildcards are not supported.

- **Source IP** – Define the source IP address in CIDR format. To save the condition, choose **Confirm**.

Both IPv4 and IPv6 CIDRs are allowed. Wildcards are not supported.

- **HTTP header** – Enter the name of the header and add one or more comparison strings. To save the condition, choose **Confirm**.

- **HTTP header name** – Rule will assess requests containing this header to confirm matching values.

Maximum 40 characters. Not case sensitive. Allowed characters are a-z, A-Z, 0-9, and the following special characters: *?~!#\$%&'!+.^_`|~. Wildcards are not supported.

- **HTTP header value** – Enter strings to compare against the HTTP header value.

Maximum 128 characters. Not case sensitive. Allowed characters are a-z, A-Z, 0-9; spaces; the following special characters: !"#%&'()+,./:;#=>@[]^_`{|}~-; and wildcards (* and ?).

- **Query string** – Route requests based on key:value pairs or values in the query string. To save the condition, choose **Confirm**.

Maximum 128 characters. Not case sensitive. Allowed characters are a-z, A-Z, 0-9; the following special characters: _.\$/~"@:~+&!;,=; and wildcards (* and ?).

9. Choose **Next**.

10. Define one of the following actions for your rule:

- **Forward to target groups** – Choose one or more target groups to forward traffic to. To add target groups choose **Add target group**. If using more than one target group, select a weight for each target group and review the associated percentage. You must enable group-level stickiness on a rule, if you've enabled stickiness on one or more of the target groups.
 - **Redirect to URL** – Specify the URL that client requests will be redirected to. This can be done by entering each part separately on the **URI parts** tab, or by entering the full address on the **Full URL** tab. For **Status code** you can configure redirects as either temporary (HTTP 302) or permanent (HTTP 301) based on your needs.
 - **Return fixed response** – Specify the **Response code** that will be returned to dropped client requests. Additionally, you can specify the **Content type** and **Response body**, but they're not required.
11. Choose **Next**.
 12. Specify the **Priority** of your rule by entering a value from 1-50000.
 13. Choose **Next**.
 14. Review all the details and settings currently configured for your new rule. Once you're satisfied with your selections, choose **Create**.

To add a rule using the AWS CLI

Use the [create-rule](#) command to create the rule. Use the [describe-rules](#) command to view information about the rule.

Edit a rule

You can edit the action and conditions for a rule at any time. Rule updates do not take effect immediately, so requests could be routed using the previous rule configuration for a short time after you update a rule. Any in-flight requests are completed.

To edit a rule using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Listeners and rules** tab, do one of the following:
 - Select the text in the **Protocol:Port** column to open the detail page for the listener.

- i. On the **Rules** tab, in the **Listener rules** section, select the text in the **Name tag** column for the rule you want to edit.

Choose **Actions**, then **Edit rule**.

- ii. On the **Rules** tab, in the **Listener rules** section, select the rule you want to edit.

Choose **Actions**, then **Edit rule**.

5. Modify the name and tags as needed. To add additional tags select the **Add additional tags** text.
6. Choose **Next**
7. Modify the conditions as needed. You can add, edit an existing, or delete conditions.
8. Choose **Next**
9. Modify the actions as needed.
10. Choose **Next**
11. Modify the rule priority as needed. You can enter a value from 1-50000.
12. Choose **Next**
13. Review all the details and updated settings configured for your rule. Once you're satisfied with your selections, choose **Save changes**.

To edit a rule using the AWS CLI

Use the [modify-rule](#) command.

Update rule priority

Rules are evaluated in priority order, from the lowest value to the highest value. The default rule is evaluated last. You can change the priority of a nondefault rule at any time. You cannot change the priority of the default rule.

To update rule priority using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.

4. On the **Listeners and rules** tab, do one of the following:
 - a. Select the text in the **Protocol:Port** or **Rules** columns to open the detail page for the listener.
 - i. Choose **Actions**, then **Reprioritize rules**.
 - ii. On the **Rules** tab, in the **Listener rules** section, choose **Actions** then **Reprioritize rules**.
 - b. Select the listener.
 - Choose **Manage rules**, then **Reprioritize rules**
5. In the **Listener rules** section the **Priority** column displays current rules priority. You can update a rules priority by entering a value from 1-50000.
6. Once you're satisfied with your changes, choose **Save changes**.

To update rule priorities using the AWS CLI

Use the [set-rule-priorities](#) command.

Delete a rule

You can delete the nondefault rules for a listener at any time. You cannot delete the default rule for a listener. When you delete a listener, all its rules are deleted.

To delete a rule using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Listeners and rules** tab, do one of the following:
 - a. Select the text in the **Protocol:Port** or **Rules** columns to open the detail page for the listener.
 - i. Select the rule you want to delete.
 - ii. Choose **Actions**, then **Delete rule**
 - iii. Type `confirm` in the text field, then choose **Delete**.

- b. Select the text in the **Name tag** column to open the detail page for the rule.
 - i. Choose **Actions**, then **Delete rule**.
 - ii. Type `confirm` in the text field, then choose **Delete**.

To delete a rule using the AWS CLI

Use the [delete-rule](#) command.

Update an HTTPS listener for your Application Load Balancer

After you create an HTTPS listener, you can replace the default certificate, update the certificate list, or replace the security policy.

Tasks

- [Replace the default certificate](#)
- [Add certificates to the certificate list](#)
- [Remove certificates from the certificate list](#)
- [Update the security policy](#)

Replace the default certificate

You can replace the default certificate for your listener using the following procedure. For more information, see [SSL certificates](#).

To change the default certificate using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Listeners and rules** tab, choose the text in the **Protocol:Port** column to open the detail page for the listener.
5. On the **Certificates** tab, choose **Change default**.
6. Within the **ACM and IAM certificates** table, select a new default certificate.

7. Choose **Save as default**.

To change the default certificate using the AWS CLI

Use the [modify-listener](#) command.

Add certificates to the certificate list

You can add certificates to the certificate list for your listener using the following procedure. When you first create an HTTPS listener, the certificate list is empty. You can add one or more certificates. You can optionally add the default certificate to ensure that this certificate is used with the SNI protocol even if it is replaced as the default certificate. For more information, see [SSL certificates](#).

To change the default certificate using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Listeners and rules** tab, choose the text in the **Protocol:Port** column to open the detail page for the listener.
5. On the **Certificates** tab, choose **Add certificate**.
6. Within the **ACM and IAM certificates** table, select the certificates to add and choose **Include as pending below**.
7. If you have a certificate that isn't managed by ACM or IAM, choose **Import certificate**, complete the form, and choose **Import**.
8. Choose **Add pending certificates**.

To add a certificate to the certificate list using the AWS CLI

Use the [add-listener-certificates](#) command.

Remove certificates from the certificate list

You can remove certificates from the certificate list for an HTTPS listener using the following procedure. To remove the default certificate for an HTTPS listener, see [Replace the default certificate](#).

To remove certificates from the certificate list using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Listeners and rules** tab, select the text in the **Protocol:Port** column to open the detail page for the listener.
5. On the **Certificates** tab, select the check boxes for the certificates and choose **Remove**.
6. When prompted for confirmation, enter **confirm** and choose **Remove**.

To remove a certificate from the certificate list using the AWS CLI

Use the [remove-listener-certificates](#) command.

Update the security policy

When you create an HTTPS listener, you can select the security policy that meets your needs. When a new security policy is added, you can update your HTTPS listener to use the new security policy. Application Load Balancers do not support custom security policies. For more information, see [Security policies](#).

Using FIPS policies on your Application Load Balancer:

All secure listeners attached to an Application Load Balancer must use either FIPS security policies or non-FIPS security policies; they cannot be mixed. If an existing Application Load Balancer has two or more listeners using non-FIPS policies and you want the listeners to use FIPS security policies instead, remove all listeners until there is only one. Change the security policy of the listener to FIPS and then create additional listeners using FIPS security policies. Alternatively, you can create a new Application Load Balancer with new listeners using only FIPS security policies.

To update the security policy using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Listeners and rules** tab, select the text in the **Protocol:Port** column to open the detail page for the listener.

5. On the **Details** page, choose **Actions**, then **Edit listener**.
6. In the **Secure listener settings** section, under **Security policy**, choose a new security policy.
7. Choose **Save changes**.

To update the security policy using the AWS CLI

Use the [modify-listener](#) command.

Mutual authentication with TLS in Application Load Balancer

Mutual TLS authentication is a variation of transport layer security (TLS). Traditional TLS establishes secure communications between a server and client, where the server needs to provide its identity to its clients. With mutual TLS, a load balancer negotiates mutual authentication between the client and the server while negotiating TLS. When you use mutual TLS with Application Load Balancer, you simplify authentication management and reduce the load on your applications.

By using mutual TLS with Application Load Balancer, your load balancer can manage client authentication to help ensure that only trusted clients communicate with your backend applications. When you use this feature, Application Load Balancer authenticates clients with certificates from third-party certificate authority (CA) or by using the AWS Private Certificate Authority (PCA), optionally, with revocation checks. Application Load Balancer passes on client certificate information to the backend, which your applications can use for authorization. By using mutual TLS in Application Load Balancer, you can get built-in, scalable, managed authentication for certificate-based entities, that uses established libraries.

Mutual TLS for Application Load Balancers provides the following two options for validating your X.509v3 client certificates:

Note: X.509v1 client certificates are not supported.

- **Mutual TLS passthrough:** When you use mutual TLS passthrough mode, Application Load Balancer sends the whole client certificate chain to the target using HTTP headers. Then, by using the client certificate chain, you can implement corresponding authentication and authorization logic in your application.
- **Mutual TLS verify:** When you use mutual TLS verify mode, Application Load Balancer performs X.509 client certificate authentication for clients when a load balancer negotiates TLS connections.

To get started with mutual TLS in Application Load Balancer using passthrough, you only need to configure the listener to accept any certificates from clients. To use mutual TLS with verification, you must do the following:

- Create a new trust store resource.
- Upload your certificate authority (CA) bundle and, optionally, revocation lists.
- Attach the trust store to the listener that is configured to verify client certificates.

For step-by-step procedures to configure mutual TLS verify mode with your Application Load Balancer, see [Configuring mutual TLS on an Application Load Balancer](#).

Before you begin configuring mutual TLS on your Application Load Balancer

Before you begin configuring mutual TLS on your Application Load Balancer, be aware of the following:

Quotas

Application Load Balancers include certain limits related to the amount of trust stores, CA certificates, and certificate revocation lists in use within your AWS account.

For more information, see [Quotas for your Application Load Balancers](#).

Requirements for certificates

Application Load Balancers support the following for certificates used with mutual TLS authentication:

- Supported certificate: X.509v3
- Supported public keys: RSA 2K – 8K or ECDSA secp256r1, secp384r1, secp521r1
- Supported signature algorithms: SHA256, 384, 512 with RSA/SHA256, 384, 512 with EC/SHA256,384,512 hash with RSASSA-PSS with MGF1

CA certificate bundles

The following applies to certificate authority (CA) bundles:

- Application Load Balancers upload each certificate authority (CA) certificate bundle as a batch. Application Load Balancers don't support uploading individual certificates. If you need to add new certificates, you must upload the certificates bundle file.

- To replace a CA certificate bundle, use the [ModifyTrustStore](#) API.

Certificate order for passthrough

When you use mutual TLS passthrough, the Application Load Balancer inserts headers to present the clients certificate chain to the backend targets. The order of presentation starts with the leaf certificates and finishes with the root certificate.

Session resumption

Session resumption is not supported while using mutual TLS passthrough or verify modes with an Application Load Balancer.

HTTP headers

Application Load Balancers use X-Amzn-Mtls headers to send certificate information when it negotiates client connections using mutual TLS. For more information and example headers, see [HTTP headers and mutual TLS](#).

CA certificate files

CA certificate files must satisfy the following requirements:

- Certificate file must use PEM (Privacy Enhanced Mail) format.
- Certificate contents must be enclosed within the -----BEGIN CERTIFICATE----- and -----END CERTIFICATE----- boundaries.
- Comments must be preceded by a # character and must not contain any - characters.
- There cannot be any blank lines.

Example certificate that is not accepted (invalid):

```
# comments

Certificate:
  Data:
    Version: 3 (0x2)
    Serial Number: 01
  Signature Algorithm: ecdsa-with-SHA384
  Issuer: C=US, O=EXAMPLE, OU=EXAMPLE, CN=EXAMPLE
  Validity
    Not Before: Jan 11 23:57:57 2024 GMT
    Not After : Jan 10 00:57:57 2029 GMT
  Subject: C=US, O=EXAMPLE, OU=EXAMPLE, CN=EXAMPLE
  Subject Public Key Info:
```

```

    Public Key Algorithm: id-ecPublicKey
      Public-Key: (384 bit)
        pub:
          00:01:02:03:04:05:06:07:08
      ASN1 OID: secp384r1
      NIST CURVE: P-384
  X509v3 extensions:
    X509v3 Key Usage: critical
      Digital Signature, Key Encipherment, Certificate Sign, CRL Sign
    X509v3 Basic Constraints: critical
      CA:TRUE
    X509v3 Subject Key Identifier:
      00:01:02:03:04:05:06:07:08
    X509v3 Subject Alternative Name:
      URI:EXAMPLE.COM
  Signature Algorithm: ecdsa-with-SHA384
    00:01:02:03:04:05:06:07:08
-----BEGIN CERTIFICATE-----
Base64-encoded certificate
-----END CERTIFICATE-----

```

Example certificates that are accepted (valid):

1. Single certificate (PEM-encoded):

```

# comments
-----BEGIN CERTIFICATE-----
Base64-encoded certificate
-----END CERTIFICATE-----

```

2. Multiple certificates (PEM-encoded):

```

# comments
-----BEGIN CERTIFICATE-----
Base64-encoded certificate
-----END CERTIFICATE-----
# comments
-----BEGIN CERTIFICATE-----
Base64-encoded certificate
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
Base64-encoded certificate
-----END CERTIFICATE-----

```

HTTP headers and mutual TLS

This section describes the HTTP headers that Application Load Balancers use to send certificate information when negotiating connections with clients using mutual TLS. The specific `X-Amzn-Mtls` headers that the Application Load Balancer uses depends on the mutual TLS mode that you've specified: passthrough mode or verify mode.

For information about other HTTP headers supported by Application Load Balancers, see [HTTP headers and Application Load Balancers](#).

HTTP header for passthrough mode

For mutual TLS in passthrough mode, Application Load Balancers use the following header.

`X-Amzn-Mtls-Clientcert`

This header contains the URL-encoded PEM format of the entire client certificate chain presented in the connection, with `+=/` as safe characters.

Example header contents:

```
X-Amzn-Mtls-Clientcert: -----BEGIN%20CERTIFICATE-----%0AMIID<...reduced...>do0g%3D%3D%0A-----END%20CERTIFICATE-----%0A-----BEGIN%20CERTIFICATE-----%0AMIID1<...reduced...>3eZlyKA%3D%3D%0A-----END%20CERTIFICATE-----%0A
```

HTTP headers for verify mode

For mutual TLS in verify mode, Application Load Balancers use the following headers.

`X-Amzn-Mtls-Clientcert-Serial-Number`

This header contains a hexadecimal representation of the leaf certificate serial number.

Example header contents:

```
X-Amzn-Mtls-Clientcert-Serial-Number: 03A5B1
```

`X-Amzn-Mtls-Clientcert-Issuer`

This header contains an RFC2253 string representation of the issuer's distinguished name (DN).

Example header contents:

```
X-Amzn-Mtls-Clientcert-Issuer:  
CN=rootcamtls.com,OU=rootCA,O=mTLS,L=Seattle,ST=Washington,C=US
```

X-Amzn-Mtls-Clientcert-Subject

This header contains an RFC2253 string representation of the subject's distinguished name (DN).

Example header contents:

```
X-Amzn-Mtls-Clientcert-Subject: CN=client_.com,OU=client-3,O=mTLS,ST=Washington,C=US
```

X-Amzn-Mtls-Clientcert-Validity

This header contains an ISO8601 format of the notBefore and notAfter date.

Example header contents:

```
X-Amzn-Mtls-Clientcert-Validity:  
NotBefore=2023-09-21T01:50:17Z;NotAfter=2024-09-20T01:50:17Z
```

X-Amzn-Mtls-Clientcert-Leaf

This header contains a URL-encoded PEM format of the leaf certificate, with +=/ as safe characters.

Example header contents:

```
X-Amzn-Mtls-Clientcert-Leaf: -----BEGIN%20CERTIFICATE-----%0AMIIG<...reduced...>NmUlw  
%0A-----END%20CERTIFICATE-----%0A
```

Configuring mutual TLS on an Application Load Balancer

This section includes the procedures for configuring mutual TLS verify mode for authentication on Application Load Balancers.

To use mutual TLS passthrough mode, you only need to configure the listener to accept any certificates from clients. When you use mutual TLS passthrough, the Application Load Balancer sends the whole client certificate chain to the target using HTTP headers, which enables you to

implement corresponding authentication and authorization logic in your application. For more information, see [Create an HTTPS listener for your Application Load Balancer](#).

When you use mutual TLS in verify mode, the Application Load Balancer performs X.509 client certificate authentication for clients when a load balancer negotiates TLS connections.

To utilize mutual TLS verify mode, perform the following:

- Create a new trust store resource.
- Upload your certificate authority (CA) bundle and, optionally, revocation lists.
- Attach the trust store to the listener that is configured to verify client certificates.

Follow the procedures in this section to configure mutual TLS verify mode on your Application Load Balancer in the AWS Management Console. To configure mutual TLS by using API operations instead of the console, see the [Application Load Balancer API Reference Guide](#).

Tasks

- [Create a trust store](#)
- [Associate a trust store](#)
- [View trust store details](#)
- [Modify a trust store](#)
- [Delete a trust store](#)

Create a trust store

There are three ways that you can create a trust store: when you create an Application Load Balancer, when you create a secure listener, and by using the Trust Store console. When you add a trust store when you create a load balancer or listener, the trust store is automatically associated with the new listener. When you create a trust store by using the Trust Store console, you must associate it with a listener yourself.

This section covers creating a trust store using the Trust Store console, but the steps used while creating an Application Load Balancer or listener are the same. For more info, see [Configure a load balancer and a listener](#) and [Add an HTTPS listener](#).

Prerequisites:

- To create a trust store, you must have a certificate bundle from your Certificate Authority (CA).

To create a trust store using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Trust Stores**.
3. Select **Create trust store**.
4. **Trust store configuration**
 - a. For **Trust store name** enter a name for your trust store.
 - b. For **Certificate authority bundle** enter the Amazon S3 path to the ca certificate bundle you want your trust store to use.

Optional: Use **Object version** to select a previous version of the ca certificate bundle. Otherwise the current version is used.

5. For **Revocations** you can optionally add a certificate revocation list to your trust store.
 - Under **Certificate revocation list** enter the Amazon S3 path to the certificate revocation list you want your trust store to use.

Optional: Use **Object version** to select a previous version of the certificate revocation list. Otherwise the current version is used.

6. For **Trust store tags** you can optionally enter up to 50 tags to apply to your trust store.
7. Select **Create trust store**.

Associate a trust store

After you create a trust store, you must associate it with a listener before your Application Load Balancer can begin using the trust store. You can have only one trust store associated to each of your secure listeners, but one trust store can be associated to multiple listeners.

This section covers associating a trust store to an existing listener. Alternatively, you can associate a trust store while creating an Application Load Balancer or listener. For more info, see [Configure a load balancer and a listener](#) and [Add an HTTPS listener](#).

To associate a trust store using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.

3. Select the load balancer to view its details page.
4. On the **Listeners and rules** tab, choose the link in the **Protocol:Port** column to open the details page for the secure listener.
5. On the **Security** tab, choose **Edit secure listener settings**.
6. (Optional) If mutual TLS is not enabled, select **Mutual authentication (mTLS)** under **Client certificate handling** and then choose **Verify with trust store**.
7. Under **Trust store**, choose the trust store that you created.
8. Choose **Save changes**.

View trust store details

CA certificate bundles

The CA certificate bundle is a required component of the trust store. It's a collection of trusted root and intermediate certificates that have been validated by a certificate authority. These validated certificates ensure the client can trust the certificate being presented is owned by the load balancer.

You can view the contents of the current CA certificate bundle in your trust store at any time.

View a CA certificate bundle

To view a CA certificate bundle using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Trust Stores**.
3. Select the trust store to view the details page.
4. Choose **Actions**, then **Get CA bundle**.
5. Choose **Share link**, or **Download**.

Certificate revocation lists

Optionally, you can create a certificate revocation list for a trust store. Revocation lists are released by certificate authorities and contain data for certificates that have been revoked. Application Load Balancers only support certificate revocation lists in the PEM format.

When a certificate revocation list is added to a trust store, it's given a revocation ID. The revocation IDs increase for every revocation list added to the trust store, and they cannot be changed. If a certificate revocation list is deleted from a trust store, its revocation ID is also deleted and is not reused for the life of the trust store.

Note

Application Load Balancers cannot revoke certificates that have a negative serial number, within a certificate revocation list.

View a certificate revocation list

To view a revocation list using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Trust Stores**.
3. Select the trust store to view the details page.
4. On the **Certificate revocation lists** tab, select **Actions**, then **Get revocation list**.
5. Choose **Share link**, or **Download**.

Modify a trust store

A trust store can only contain one CA certificate bundle at a time, but you can replace the CA certificate bundle at any time after the trust store is created.

Replace a CA certificate bundle

To replace a CA certificate bundle using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Trust Stores**.
3. Select the trust store to view the details page.
4. Choose **Actions**, then **Replace CA bundle**.
5. On the **Replace CA bundle** page, under **Certificate authority bundle** enter the Amazon S3 location of the desired CA bundle.

6. (Optional) Use **Object version** to select a previous version of the certificate revocation list. Otherwise the current version is used.
7. Select **Replace CA bundle**.

Add a certificate revocation list

To add a revocation list using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Trust Stores**.
3. Select the trust store to view its details page.
4. On the **Certificate revocation lists** tab, select **Actions**, then **Add revocation list**.
5. On the **Add revocation list** page, under **Certificate revocation list** enter the Amazon S3 location of the desired certificate revocation list
6. (Optional) Use **Object version** to select a previous version of the certificate revocation list. Otherwise the current version is used.
7. Select **Add revocation list**

Delete a certificate revocation list

To delete a revocation list using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Trust Stores**.
3. Select the trust store to view the details page.
4. On the **Certificate revocation lists** tab, select **Actions**, then **Delete revocation list**.
5. Confirm the deletion by typing `confirm`.
6. Select **Delete**.

Delete a trust store

When you no longer have use for a trust store, you can delete it.

Note: You cannot delete a trust store that is currently associated with a listener.

To delete a trust store using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Trust Stores**.
3. Select the trust store to view its details page.
4. Choose **Actions**, then **Delete trust store**.
5. Confirm the deletion by typing `confirm`.
6. Select **Delete**

Connection logs for Application Load Balancers

Elastic Load Balancing provides connection logs that capture attributes about the requests sent to your Application Load Balancers. Connection logs contain information such as the client IP address and port, client certificate information, connection results, and TLS ciphers being used. These connection logs can then be used to review request patterns, and other trends.

To learn more about connection logs, see [Connection logs for your Application Load Balancer](#)

Authenticate users using an Application Load Balancer

You can configure an Application Load Balancer to securely authenticate users as they access your applications. This enables you to offload the work of authenticating users to your load balancer so that your applications can focus on their business logic.

The following use cases are supported:

- Authenticate users through an identity provider (IdP) that is OpenID Connect (OIDC) compliant.
- Authenticate users through social IdPs, such as Amazon, Facebook, or Google, through the user pools supported by Amazon Cognito.
- Authenticate users through corporate identities, using SAML, OpenID Connect (OIDC), or OAuth, through the user pools supported by Amazon Cognito.

Prepare to use an OIDC-compliant IdP

Do the following if you are using an OIDC-compliant IdP with your Application Load Balancer:

- Create a new OIDC app in your IdP. The IdP's DNS must be publicly resolvable.

- You must configure a client ID and a client secret.
- Get the following endpoints published by the IdP: authorization, token, and user info. You can locate this information in the config.
- The IdP endpoints certificates should be issued by a trusted public certificate authority.
- The DNS entries for the endpoints must be publicly resolvable, even if they resolve to private IP addresses.
- Allow one of the following redirect URLs in your IdP app, whichever your users will use, where DNS is the domain name of your load balancer and CNAME is the DNS alias for your application:
 - <https://DNS/oauth2/idpresponse>
 - <https://CNAME/oauth2/idpresponse>

Prepare to use Amazon Cognito

Regions Available

Amazon Cognito integration for Application Load Balancers is available in the following regions:

- US East (N. Virginia)
- US East (Ohio)
- US West (N. California)
- US West (Oregon)
- Canada (Central)
- Europe (Stockholm)
- Europe (Milan)
- Europe (Frankfurt)
- Europe (Zurich)
- Europe (Ireland)
- Europe (London)
- Europe (Paris)
- South America (São Paulo)
- Asia Pacific (Tokyo)
- Asia Pacific (Seoul)
- Asia Pacific (Osaka)

- Asia Pacific (Mumbai)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Jakarta)
- Middle East (UAE)
- Middle East (Bahrain)
- Africa (Cape Town)
- Israel (Tel Aviv)

Do the following if you are using Amazon Cognito user pools with your Application Load Balancer:

- Create a user pool. For more information, see [Amazon Cognito user pools](#) in the *Amazon Cognito Developer Guide*.
- Create a user pool client. You must configure the client to generate a client secret, use code grant flow, and support the same OAuth scopes that the load balancer uses. For more information, see [Configuring a user pool app client](#) in the *Amazon Cognito Developer Guide*.
- Create a user pool domain. For more information, see [Adding a Domain name for your user pool](#) in the *Amazon Cognito Developer Guide*.
- Verify that the requested scope returns an ID token. For example, the default scope, `openid` returns an ID token but the `aws:cognito.signin.user.admin` scope does not.

Note: Application Load Balancers do not support customized access tokens issued by Amazon Cognito. For more information, see [Pre token generation](#) in the *Amazon Cognito Developer Guide*.

- To federate with a social or corporate IdP, enable the IdP in the federation section. For more information, see [Add social sign-in to a user pool](#) or [Add sign-in with a SAML IdP to a user pool](#) in the *Amazon Cognito Developer Guide*.
- Allow the following redirect URLs in the callback URL field for Amazon Cognito, where DNS is the domain name of your load balancer, and CNAME is the DNS alias for your application (if you are using one):
 - `https://DNS/oauth2/idpresponse`
 - `https://CNAME/oauth2/idpresponse`
- Allow your user pool domain on your IdP app's callback URL. Use the format for your IdP. For example:
 - `https://domain-prefix.auth.region.amazoncognito.com/saml2/idpresponse`

- <https://user-pool-domain/oauth2/idpresponse>

The callback URL in the app client settings must use all lowercase letters.

To enable a user to configure a load balancer to use Amazon Cognito to authenticate users, you must grant the user permission to call the `cognito-idp:DescribeUserPoolClient` action.

Prepare to use Amazon CloudFront

Enable the following settings if you are using a CloudFront distribution in front of your Application Load Balancer:

- Forward request headers (all) — Ensures that CloudFront does not cache responses for authenticated requests. This prevents them from being served from the cache after the authentication session expires. Alternatively, to reduce this risk while caching is enabled, owners of a CloudFront distribution can set the time-to-live (TTL) value to expire before the authentication cookie expires.
- Query string forwarding and caching (all) — Ensures that the load balancer has access to the query string parameters required to authenticate the user with the IdP.
- Cookie forwarding (all) — Ensures that CloudFront forwards all authentication cookies to the load balancer.

Configure user authentication

You configure user authentication by creating an authenticate action for one or more listener rules. The `authenticate-cognito` and `authenticate-oidc` action types are supported only with HTTPS listeners. For descriptions of the corresponding fields, see [AuthenticateCognitoActionConfig](#) and [AuthenticateOidcActionConfig](#) in the *Elastic Load Balancing API Reference version 2015-12-01*.

The load balancer sends a session cookie to the client to maintain authentication status. This cookie always contains the `secure` attribute, because user authentication requires an HTTPS listener. This cookie contains the `SameSite=None` attribute with CORS (cross-origin resource sharing) requests.

For a load balancer supporting multiple applications that require independent client authentication, each listener rule with an authenticate action should have a unique cookie name. This ensures that clients are always authenticated with the IdP before being routed to the target group specified in the rule.

Application Load Balancers do not support cookie values that are URL encoded.

By default, the `SessionTimeout` field is set to 7 days. If you want shorter sessions, you can configure a session timeout as short as 1 second. For more information, see [Session timeout](#).

Set the `OnUnauthenticatedRequest` field as appropriate for your application. For example:

- **Applications that require the user to log in using a social or corporate identity**—This is supported by the default option, `authenticate`. If the user is not logged in, the load balancer redirects the request to the IdP authorization endpoint and the IdP prompts the user to log in using its user interface.
- **Applications that provide a personalized view to a user that is logged in or a general view to a user that is not logged in**—To support this type of application, use the `allow` option. If the user is logged in, the load balancer provides the user claims and the application can provide a personalized view. If the user is not logged in, the load balancer forwards the request without the user claims and the application can provide the general view.
- **Single-page applications with JavaScript that loads every few seconds**—If you use the `deny` option, the load balancer returns an HTTP 401 Unauthorized error to AJAX calls that have no authentication information. But if the user has expired authentication information, it redirects the client to the IdP authorization endpoint.

The load balancer must be able to communicate with the IdP token endpoint (`TokenEndpoint`) and the IdP user info endpoint (`UserInfoEndpoint`). Application Load Balancers only support IPv4 when communicating with these endpoints. If your IdP uses public addresses, ensure the security groups for your load balancer and the network ACLs for your VPC allow access to the endpoints. When using an internal load balancer or the **IP address type** `dualstack-without-public-ipv4`, a NAT gateway can enable the load balancer to communicate with the endpoints. For more information, see [NAT gateway basics](#) in the *Amazon VPC User Guide*.

Use the following [create-rule](#) command to configure user authentication.

```
aws elbv2 create-rule --listener-arn listener-arn --priority 10 \  
--conditions Field=path-pattern,Values="/login" --actions file://actions.json
```

The following is an example of the `actions.json` file that specifies an `authenticate-oidc` action and a `forward` action. `AuthenticationRequestExtraParams` allows you to pass extra parameters to an IdP during authentication. Please follow documentation provided by your identity provider to determine the fields that are supported

```
[{
  "Type": "authenticate-oidc",
  "AuthenticateOidcConfig": {
    "Issuer": "https://idp-issuer.com",
    "AuthorizationEndpoint": "https://authorization-endpoint.com",
    "TokenEndpoint": "https://token-endpoint.com",
    "UserInfoEndpoint": "https://user-info-endpoint.com",
    "ClientId": "abcdefghijklmnopqrstuvwxy123456789",
    "ClientSecret": "123456789012345678901234567890",
    "SessionCookieName": "my-cookie",
    "SessionTimeout": 3600,
    "Scope": "email",
    "AuthenticationRequestExtraParams": {
      "display": "page",
      "prompt": "login"
    },
    "OnUnauthenticatedRequest": "deny"
  },
  "Order": 1
},
{
  "Type": "forward",
  "TargetGroupArn": "arn:aws:elasticloadbalancing:region-code:account-id:targetgroup/target-group-name/target-group-id",
  "Order": 2
}]
```

The following is an example of the `actions.json` file that specifies an `authenticate-cognito` action and a `forward` action.

```
[{
  "Type": "authenticate-cognito",
  "AuthenticateCognitoConfig": {
    "UserPoolArn": "arn:aws:cognito-idp:region-code:account-id:userpool/user-pool-id",
    "UserPoolClientId": "abcdefghijklmnopqrstuvwxy123456789",
    "UserPoolDomain": "userPoolDomain1",
    "SessionCookieName": "my-cookie",
    "SessionTimeout": 3600,
    "Scope": "email",
    "AuthenticationRequestExtraParams": {
      "display": "page",
      "prompt": "login"
    }
  }
},
{
  "Type": "forward",
  "TargetGroupArn": "arn:aws:elasticloadbalancing:region-code:account-id:targetgroup/target-group-name/target-group-id",
  "Order": 2
}]
```



```


    },
    "OnUnauthenticatedRequest": "deny"
  },
  "Order": 1
},
{
  "Type": "forward",
  "TargetGroupArn": "arn:aws:elasticloadbalancing:region-code:account-
id:targetgroup/target-group-name/target-group-id",
  "Order": 2
}]

```

For more information, see [Listener rules](#).

Authentication flow

The following network diagram is a visual representation of how an Application Load Balancer uses OIDC to authenticate users.



The numbered items below, highlight and explain elements shown in the preceding network diagram.

1. User sends an HTTPS request to a website hosted behind an Application Load Balancer. When the conditions for a rule with an authenticate action are met, the load balancer checks for an authentication session cookie in the request headers.
2. If the cookie is not present, the load balancer redirects the user to the IdP authorization endpoint so that the IdP can authenticate the user.
3. After the user is authenticated, the IdP sends the user back to the load balancer with an authorization grant code.
4. The load balancer presents the authorization grant code to the IdP token endpoint.
5. Upon receiving a valid authorization grant code, the IdP provides the ID token and access token to the Application Load Balancer.
6. The Application Load Balancer then sends the access token to the user info endpoint.
7. The user info endpoint exchanges the access token for user claims.
8. The Application Load Balancer redirects the user with the AWSELB authentication session cookie to the original URI. Because most browsers limit the cookie size to 4K, the load balancer shards a cookie that is greater than 4K in size into multiple cookies. If the total size of the user claims and access token received from the IdP is greater than 11K bytes in size, the load balancer returns an HTTP 500 error to the client and increments the `ELBAuthUserClaimsSizeExceeded` metric.
9. The Application Load Balancer validates the cookie and forwards the user info to targets in the `X-AMZN-OIDC-*` HTTP headers set. For more information, see [User claims encoding and signature verification](#).
10. The target sends a response back to the Application Load Balancer.
11. The Application Load Balancer sends the final response to the user.

Every new request goes through steps 1 through 11, while subsequent requests go through steps 9 through 11. That is, every subsequent request starts at step 9 as long as the cookie has not expired.

The `AWSALBAuthNonce` cookie is added to the request header after the user authenticates at the IdP. This does not change how the Application Load Balancer processes redirect requests from the IdP.

If the IdP provides a valid refresh token in the ID token, the load balancer saves the refresh token and uses it to refresh the user claims each time the access token expires, until the session times out or the IdP refresh fails. If the user logs out, the refresh fails and the load balancer redirects the user

to the IdP authorization endpoint. This enables the load balancer to drop sessions after the user logs out. For more information, see [Session timeout](#).

Note

The cookie expiry is different from the authentication session expiry. The cookie expiry is an attribute of the cookie, which is set to 7 days. The actual length of the authentication session is determined by the session timeout configured on the Application Load Balancer for the authentication feature. This session timeout is included in the Auth cookie value, which is also encrypted.

User claims encoding and signature verification

After your load balancer authenticates a user successfully, it sends the user claims received from the IdP to the target. The load balancer signs the user claim so that applications can verify the signature and verify that the claims were sent by the load balancer.

The load balancer adds the following HTTP headers:

`x-amzn-oidc-accesstoken`

The access token from the token endpoint, in plain text.

`x-amzn-oidc-identity`

The subject field (sub) from the user info endpoint, in plain text.

Note: The sub claim is the best way to identify a given user.

`x-amzn-oidc-data`

The user claims, in JSON web tokens (JWT) format.

Access tokens and user claims are different from ID tokens. Access tokens and user claims only allow access to server resources, while ID tokens carry additional information to authenticate a user. The Application Load Balancer creates a new access token when authenticating a user and only passes the access tokens and claims to the backend, however it does not pass the ID token information.

These tokens follow the JWT format but are not ID tokens. The JWT format includes a header, payload, and signature that are base64 URL encoded, and includes padding characters at the end.

An Application Load Balancer uses ES256 (ECDSA using P-256 and SHA256) to generate the JWT signature.

The JWT header is a JSON object with the following fields:

```
{
  "alg": "algorithm",
  "kid": "12345678-1234-1234-1234-123456789012",
  "signer": "arn:aws:elasticloadbalancing:region-code:account-id:loadbalancer/
app/load-balancer-name/load-balancer-id",
  "iss": "url",
  "client": "client-id",
  "exp": "expiration"
}
```

The JWT payload is a JSON object that contains the user claims received from the IdP user info endpoint.

```
{
  "sub": "1234567890",
  "name": "name",
  "email": "alias@example.com",
  ...
}
```

Because the load balancer does not encrypt the user claims, we recommend that you configure the target group to use HTTPS. If you configure your target group to use HTTP, be sure to restrict the traffic to your load balancer using security groups.

To ensure security, you must verify the signature before doing any authorization based on the claims and validate that the `signer` field in the JWT header contains the expected Application Load Balancer ARN.

To get the public key, get the key ID from the JWT header and use it to look up the public key from the endpoint. The endpoint for each AWS Region is as follows:

```
https://public-keys.auth.elb.region.amazonaws.com/key-id
```

For AWS GovCloud (US), the endpoints are as follows:

```
https://s3-us-gov-west-1.amazonaws.com/aws-elb-public-keys-prod-us-gov-west-1/key-id
```

```
https://s3-us-gov-east-1.amazonaws.com/aws-elb-public-keys-prod-us-gov-east-1/key-id
```

The following example shows how to get the key ID, public key, and payload in Python 3.x:

```
import jwt
import requests
import base64
import json

# Step 1: Validate the signer
expected_alb_arn = 'arn:aws:elasticloadbalancing:region-code:account-id:loadbalancer/
app/load-balancer-name/load-balancer-id'

encoded_jwt = headers.dict['x-amzn-oidc-data']
jwt_headers = encoded_jwt.split('.')[0]
decoded_jwt_headers = base64.b64decode(jwt_headers)
decoded_jwt_headers = decoded_jwt_headers.decode("utf-8")
decoded_json = json.loads(decoded_jwt_headers)
received_alb_arn = decoded_json['signer']

assert expected_alb_arn == received_alb_arn, "Invalid Signer"

# Step 2: Get the key id from JWT headers (the kid field)
kid = decoded_json['kid']

# Step 3: Get the public key from regional endpoint
url = 'https://public-keys.auth.elb.' + region + '.amazonaws.com/' + kid
req = requests.get(url)
pub_key = req.text

# Step 4: Get the payload
payload = jwt.decode(encoded_jwt, pub_key, algorithms=['ES256'])
```

The following example shows how to get the key ID, public key, and payload in Python 2.7:

```
import jwt
import requests
import base64
import json

# Step 1: Validate the signer
expected_alb_arn = 'arn:aws:elasticloadbalancing:region-code:account-id:loadbalancer/
app/load-balancer-name/load-balancer-id'
```

```
encoded_jwt = headers.dict['x-amzn-oidc-data']
jwt_headers = encoded_jwt.split('.')[0]
decoded_jwt_headers = base64.b64decode(jwt_headers)
decoded_json = json.loads(decoded_jwt_headers)
received_alb_arn = decoded_json['signer']

assert expected_alb_arn == received_alb_arn, "Invalid Signer"

# Step 2: Get the key id from JWT headers (the kid field)
kid = decoded_json['kid']

# Step 3: Get the public key from regional endpoint
url = 'https://public-keys.auth.elb.' + region + '.amazonaws.com/' + kid
req = requests.get(url)
pub_key = req.text

# Step 4: Get the payload
payload = jwt.decode(encoded_jwt, pub_key, algorithms=['ES256'])
```

Considerations

- These examples do not cover how to validate the signature of the issuer with the signature in the token.
- Standard libraries are not compatible with the padding that is included in the Application Load Balancer authentication token in JWT format.

Timeout

Session timeout

The refresh token and the session timeout work together as follows:

- If the session timeout is shorter than the access token expiration, the load balancer honors the session timeout. If the user has an active session with the IdP, the user might not be prompted to log in again. Otherwise, the user is redirected to log in.
 - If the IdP session timeout is longer than the Application Load Balancer session timeout, the user does not have to supply credentials to log in again. Instead, the IdP redirects back to the Application Load Balancer with a new authorization grant code. Authorization codes are single use, even if there is no re-login.

- If the IdP session timeout is equal to or shorter than the Application Load Balancer session timeout, the user is asked to supply credentials to log in again. After the user logs in, IdP redirects back to the Application Load Balancer with a new authorization grant code, and the rest of the authentication flow continues until the request reaches the backend.
- If the session timeout is longer than the access token expiration and the IdP does not support refresh tokens, the load balancer keeps the authentication session until it times out. Then, it has the user log in again.
- If the session timeout is longer than the access token expiration and the IdP supports refresh tokens, the load balancer refreshes the user session each time the access token expires. The load balancer has the user log in again only after the authentication session times out or the refresh flow fails.

Client login timeout

A client must initiate and complete the authentication process within 15 minutes. If a client fails to complete authentication within the 15-minute limit, it receives an HTTP 401 error from the load balancer. This timeout can't be changed or removed.

For example, if a user loads the login page through the Application Load Balancer, they must complete the login process within 15 minutes. If the user waits and then attempts to log in after the 15-minute timeout has expired, the load balancer returns an HTTP 401 error. The user will have to refresh the page and attempt logging in again.

Authentication logout

When an application needs to log out an authenticated user, it should set the expiration time of the authentication session cookie to -1 and redirect the client to the IdP logout endpoint (if the IdP supports one). To prevent users from reusing a deleted cookie, we recommend that you configure as short an expiration time for the access token as is reasonable. If a client provides a load balancer with a session cookie that has an expired access token with a non-NULL refresh token, the load balancer contacts the IdP to determine whether the user is still logged in.

The client logout landing page is an unauthenticated page. This means that it cannot be behind an Application Load Balancer rule that requires authentication.

- When a request is sent to the target, the application must set the expiry to -1 for all authentication cookies. Application Load Balancers support cookies up to 16K in size and can therefore create up to 4 shards to send to the client.

- If the IdP has a logout endpoint, it should issue a redirect to the IdP logout endpoint, for example, the [LOGOUT Endpoint](#) documented in the *Amazon Cognito Developer Guide*.
- If the IdP does not have a logout endpoint, the request goes back to the client logout landing page, and the login process is restarted.
- Assuming that the IdP has a logout endpoint, the IdP must expire access tokens and refresh tokens, and redirect the user back to the client logout landing page.
- Subsequent requests follow the original authentication flow.

HTTP headers and Application Load Balancers

HTTP requests and HTTP responses use header fields to send information about the HTTP messages. HTTP headers are added automatically. Header fields are colon-separated name-value pairs that are separated by a carriage return (CR) and a line feed (LF). A standard set of HTTP header fields is defined in RFC 2616, [Message Headers](#). There are also non-standard HTTP headers available that are automatically added and widely used by the applications. Some of the non-standard HTTP headers have an X-Forwarded prefix. Application Load Balancers support the following X-Forwarded headers.

For more information about HTTP connections, see [Request routing](#) in the *Elastic Load Balancing User Guide*.

X-Forwarded headers

- [X-Forwarded-For](#)
- [X-Forwarded-Proto](#)
- [X-Forwarded-Port](#)

X-Forwarded-For

The X-Forwarded-For request header helps you identify the IP address of a client when you use an HTTP or HTTPS load balancer. Because load balancers intercept traffic between clients and servers, your server access logs only contain the IP address of the load balancer. To see the IP address of the client, use the `routing.http.xff_header_processing.mode` attribute. This attribute enables you to modify, preserve, or remove the X-Forwarded-For header in the HTTP request before the Application Load Balancer sends the request to the target. The possible

values for this attribute are `append`, `preserve`, and `remove`. The default value for this attribute is `append`.

Important

The `X-Forwarded-For` header should be used with caution due to the potential for security risks. The entries can only be considered trustworthy if added by systems that are properly secured within the network.

Append

By default, the Application Load Balancer stores the IP address of the client in the `X-Forwarded-For` request header and passes the header to your server. If the `X-Forwarded-For` request header is not included in the original request, the load balancer creates one with the client IP address as the request value. Otherwise, the load balancer appends the client IP address to the existing header and then passes the header to your server. The `X-Forwarded-For` request header may contain multiple IP addresses that are comma separated.

The `X-Forwarded-For` request header takes the following form:

```
X-Forwarded-For: client-ip-address
```

The following is an example `X-Forwarded-For` request header for a client with an IP address of `203.0.113.7`.

```
X-Forwarded-For: 203.0.113.7
```

The following is an example `X-Forwarded-For` request header for a client with an IPv6 address of `2001:DB8::21f:5bff:febf:ce22:8a2e`.

```
X-Forwarded-For: 2001:DB8::21f:5bff:febf:ce22:8a2e
```

When the client port preservation attribute (`routing.http.xff_client_port.enabled`) is enabled on the load balancer, the `X-Forwarded-For` request header includes the `client-port-number` appended to the `client-ip-address`, separated by a colon. The header then takes the following form:

```
IPv4 -- X-Forwarded-For: client-ip-address:client-port-number
```

```
IPv6 -- X-Forwarded-For: [client-ip-address]:client-port-number
```

For IPv6, note that when the load balancer appends the `client-ip-address` to the existing header, it encloses the address in square brackets.

The following is an example `X-Forwarded-For` request header for a client with an IPv4 address of `12.34.56.78` and a port number of `8080`.

```
X-Forwarded-For: 12.34.56.78:8080
```

The following is an example `X-Forwarded-For` request header for a client with an IPv6 address of `2001:db8:85a3:8d3:1319:8a2e:370:7348` and a port number of `8080`.

```
X-Forwarded-For: [2001:db8:85a3:8d3:1319:8a2e:370:7348]:8080
```

Preserve

The `preserve` mode in the attribute ensures that the `X-Forwarded-For` header in the HTTP request is not modified in any way before it is sent to targets.

Remove

The `remove` mode in the attribute removes the `X-Forwarded-For` header in the HTTP request before it is sent to targets.

Note

If you enable the client port preservation attribute (`routing.http.xff_client_port.enabled`), and also select `preserve` or `remove` for the `routing.http.xff_header_processing.mode` attribute, the Application Load Balancer overrides the client port preservation attribute. It keeps the `X-Forwarded-For` header unchanged, or removes it depending on the mode you select, before it sends it to the targets.

The following table shows examples of the X-Forwarded-For header that the target receives when you select either the append, preserve or the remove mode. In this example, the IP address of the last hop is 127.0.0.1.

Request description	Example request	XFF in append mode	XFF in preserve mode	XFF in remove mode
Request is sent with no XFF header	GET / index.ht ml HTTP/1.1 Host: example.com	X-Forward ed-For: 127.0.0.1	Not present	Not present
Request is sent with an XFF header and a client IP address.	GET / index.ht ml HTTP/1.1 Host: example.com X-Forward ed-For: 127.0.0.4	X-Forward ed-For: 127.0.0.4, 127.0.0.1	X-Forward ed-For: 127.0.0.4	Not present
Request is sent with an XFF header with multiple client IP addresses.	GET / index.ht ml HTTP/1.1 Host: example.com X-Forward ed-For: 127.0.0.4, 127.0.0.8	X-Forward ed-For: 127.0.0.4, 127.0.0.8, 127.0.0.1	X-Forward ed-For: 127.0.0.4, 127.0.0.8	Not present

To modify, preserve, or remove the X-Forwarded-For header using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.

2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Attributes** tab, choose **Edit**.
5. In the **Traffic configuration** section, under **Packet handling**, for **X-Forwarded-For header** choose **Append** (default), **Preserve**, or **Remove**.
6. Choose **Save changes**.

To modify, preserve, or remove the X-Forwarded-For header using the AWS CLI

Use the [modify-load-balancer-attributes](#) command with the `routing.http.xff_header_processing.mode` attribute.

X-Forwarded-Proto

The X-Forwarded-Proto request header helps you identify the protocol (HTTP or HTTPS) that a client used to connect to your load balancer. Your server access logs contain only the protocol used between the server and the load balancer; they contain no information about the protocol used between the client and the load balancer. To determine the protocol used between the client and the load balancer, use the X-Forwarded-Proto request header. Elastic Load Balancing stores the protocol used between the client and the load balancer in the X-Forwarded-Proto request header and passes the header along to your server.

Your application or website can use the protocol stored in the X-Forwarded-Proto request header to render a response that redirects to the appropriate URL.

The X-Forwarded-Proto request header takes the following form:

```
X-Forwarded-Proto: originatingProtocol
```

The following example contains an X-Forwarded-Proto request header for a request that originated from the client as an HTTPS request:

```
X-Forwarded-Proto: https
```

X-Forwarded-Port

The X-Forwarded-Port request header helps you identify the destination port that the client used to connect to the load balancer.

Tags for your listeners and rules

Tags help you to categorize your listeners and rules in different ways. For example, you can tag a resource by purpose, owner, or environment.

You can add multiple tags to each listener and rule. Tag keys must be unique for each listener and rule. If you add a tag with a key that is already associated with the listener and rule, it updates the value of that tag.

When you are finished with a tag, you can remove it.

Restrictions

- Maximum number of tags per resource—50
- Maximum key length—127 Unicode characters
- Maximum value length—255 Unicode characters
- Tag keys and values are case-sensitive. Allowed characters are letters, spaces, and numbers representable in UTF-8, plus the following special characters: + - = . _ : / @. Do not use leading or trailing spaces.
- Do not use the `aws :` prefix in your tag names or values because it is reserved for AWS use. You can't edit or delete tag names or values with this prefix. Tags with this prefix do not count against your tags per resource limit.

Update listener tags

To update the tags for a listener using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Load Balancers**.
3. Choose the name of the load balancer that contains the listener you want to update, to open its details page.
4. On the **Listeners and rules** tab, do one of the following:
 - a. Select the text in the **Protocol:Port** column to open the detail page for the listener.

On the **Tags** tab, choose **Manage tags**.
 - b. Select the listener you want to update tags on.

Choose **Manage listener**, then **Manage tags**.

- c. Select the text in the **Tags** column to open the listener details page, on the tags tab.

Choose **Manage tags**.

5. On the **Manage tags** page, do one or more of the following:
 - a. To update a tag, enter new values for **Key** and **Value**.
 - b. To add a tag, choose **Add new tag** and enter values for **Key** and **Value**.
 - c. To delete a tag, choose **Remove** next to the tag.
6. When you have finished updating tags, choose **Save changes**.

To update the tags for a listener using the AWS CLI

Use the [add-tags](#) and [remove-tags](#) commands.

Update rule tags

To update the tags for a rule using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Load Balancers**.
3. Choose the name of the load balancer that contains the rule you want to update, to open its details page.
4. On the **Listeners and rules** tab, select the text in the **Protocol:Port** column of the listener containing the rule you want to update, to open the detail page of the listener
5. On the listener details page, do one of the following:
 - a. Select the text in the **Name tag** column to open the detail page for the rule.

On the rule details page, choose **Manage tags**.

- b. Select the text in the **Tags** column for the rule you want to update.

In the tags summary pop up choose **Manage tags**.

6. On the **Manage tags** page, do one or more of the following:
 - a. To update a tag, enter new values for **Key** and **Value**.

- b. To add a tag, choose **Add new tag** and enter values for **Key** and **Value**.
 - c. To delete a tag, choose **Remove** next to the tag.
7. When you have finished updating tags, choose **Save changes**.

To update the tags for a rule using the AWS CLI

Use the [add-tags](#) and [remove-tags](#) commands.

Delete a listener for your Application Load Balancer

You can delete a listener at any time. When you delete a load balancer, all its listeners are deleted.

To delete a listener using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, choose **Load Balancers**.
3. Select the load balancer.
4. On the **Listeners and rules** tab, select the check box for the listener and choose **Manage listener, Delete listener**.
5. When prompted for confirmation, enter **confirm** and choose **Delete**.

To delete a listener using the AWS CLI

Use the [delete-listener](#) command.

Target groups for your Application Load Balancers

Target groups route requests to individual registered targets, such as EC2 instances, using the protocol and port number that you specify. You can register a target with multiple target groups. You can configure health checks on a per target group basis. Health checks are performed on all targets registered to a target group that is specified in a listener rule for your load balancer.

Each target group is used to route requests to one or more registered targets. When you create each listener rule, you specify a target group and conditions. When a rule condition is met, traffic is forwarded to the corresponding target group. You can create different target groups for different types of requests. For example, create one target group for general requests and other target groups for requests to the microservices for your application. You can use each target group with only one load balancer. For more information, see [Application Load Balancer components](#).

You define health check settings for your load balancer on a per target group basis. Each target group uses the default health check settings, unless you override them when you create the target group or modify them later on. After you specify a target group in a rule for a listener, the load balancer continually monitors the health of all targets registered with the target group that are in an Availability Zone enabled for the load balancer. The load balancer routes requests to the registered targets that are healthy.

Contents

- [Routing configuration](#)
- [Target type](#)
- [IP address type](#)
- [Protocol version](#)
- [Registered targets](#)
- [Target group attributes](#)
- [Routing algorithms](#)
- [Automatic Target Weights \(ATW\)](#)
- [Deregistration delay](#)
- [Slow start mode](#)
- [Create a target group](#)
- [Health checks for your target groups](#)

- [Cross-zone load balancing for target groups](#)
- [Target group health](#)
- [Register targets with your target group](#)
- [Sticky sessions for your Application Load Balancer](#)
- [Lambda functions as targets](#)
- [Tags for your target group](#)
- [Delete a target group](#)

Routing configuration

By default, a load balancer routes requests to its targets using the protocol and port number that you specified when you created the target group. Alternatively, you can override the port used for routing traffic to a target when you register it with the target group.

Target groups support the following protocols and ports:

- **Protocols:** HTTP, HTTPS
- **Ports:** 1-65535

When a target group is configured with the HTTPS protocol or uses HTTPS health checks, if any HTTPS listener is using a TLS 1.3 security policy, the `ELBSecurityPolicy-TLS13-1-0-2021-06` security policy will be used for target connections. Otherwise, the `ELBSecurityPolicy-2016-08` security policy is used. The load balancer establishes TLS connections with the targets using certificates that you install on the targets. The load balancer does not validate these certificates. Therefore, you can use self-signed certificates or certificates that have expired. Because the load balancer, and its targets are in a virtual private cloud (VPC), traffic between the load balancer and the targets is authenticated at the packet level, so it is not at risk of man-in-the-middle attacks or spoofing even if the certificates on the targets are not valid. Traffic that leaves AWS will not have these same protections, and additional steps may be needed to secure traffic further.

Target type

When you create a target group, you specify its target type, which determines the type of target you specify when registering targets with this target group. After you create a target group, you cannot change its target type.

The following are the possible target types:

`instance`

The targets are specified by instance ID.

`ip`

The targets are IP addresses.

`lambda`

The target is a Lambda function.

When the target type is `ip`, you can specify IP addresses from one of the following CIDR blocks:


- The subnets of the VPC for the target group
- 10.0.0.0/8 ([RFC 1918](#))
- 100.64.0.0/10 ([RFC 6598](#))
- 172.16.0.0/12 (RFC 1918)
- 192.168.0.0/16 (RFC 1918)

 **Important**

You can't specify publicly routable IP addresses.

All of the supported CIDR blocks enable you to register the following targets with a target group:

- Instances in a VPC that is peered to the load balancer VPC (same Region or different Region).
- AWS resources that are addressable by IP address and port (for example, databases).
- On-premises resources linked to AWS through AWS Direct Connect or a Site-to-Site VPN connection.

 **Note**

For Application Load Balancers deployed within a Local Zone, the `ip` targets must be in the same Local Zone to receive traffic.

For more information, see [What is AWS Local Zones?](#)

If you specify targets using an instance ID, traffic is routed to instances using the primary private IP address specified in the primary network interface for the instance. If you specify targets using IP addresses, you can route traffic to an instance using any private IP address from one or more network interfaces. This enables multiple applications on an instance to use the same port. Each network interface can have its own security group.

If the target type of your target group is `Lambda`, you can register a single Lambda function. When the load balancer receives a request for the Lambda function, it invokes the Lambda function. For more information, see [Lambda functions as targets](#).

You can configure Amazon Elastic Container Service (Amazon ECS) as a target of your Application Load Balancer. For more information, see [Creating an Application Load Balancer](#) in the *Amazon Elastic Container Service User Guide for AWS Fargate*.

IP address type

When creating a new target group, you can select the IP address type of your target group. This controls the IP version used to communicate with targets and check their health status.

Application Load Balancers support both IPv4 and IPv6 target groups. The default selection is IPv4.

Considerations

- All IP addresses within a target group must have the same IP address type. For example, you can't register an IPv4 target with an IPv6 target group.
- IPv6 target groups can only be used with `dualstack` load balancers.
- IPv6 target groups support IP and Instance type targets.

Protocol version

By default, Application Load Balancers send requests to targets using HTTP/1.1. You can use the protocol version to send requests to targets using HTTP/2 or gRPC.

The following table summarizes the result for the combinations of request protocol and target group protocol version.

Request protocol	Protocol version	Result
HTTP/1.1	HTTP/1.1	Success
HTTP/2	HTTP/1.1	Success
gRPC	HTTP/1.1	Error
HTTP/1.1	HTTP/2	Error
HTTP/2	HTTP/2	Success
gRPC	HTTP/2	Success if targets support gRPC
HTTP/1.1	gRPC	Error
HTTP/2	gRPC	Success if a POST request
gRPC	gRPC	Success

Considerations for the gRPC protocol version

- The only supported listener protocol is HTTPS.
- The only supported action type for listener rules is forward.
- The only supported target types are instance and ip.
- The load balancer parses gRPC requests and routes the gRPC calls to the appropriate target groups based on the package, service, and method.
- The load balancer supports unary, client-side streaming, server-side streaming, and bi-directional streaming.
- You must provide a custom health check method with the format `/package.service/method`.
- You must specify the gRPC status codes to use when checking for a successful response from a target.
- You cannot use Lambda functions as targets.

Considerations for the HTTP/2 protocol version

- The only supported listener protocol is HTTPS.
- The only supported action type for listener rules is `forward`.
- The only supported target types are `instance` and `ip`.
- The load balancer supports streaming from clients. The load balancer does not support streaming to the targets.

Registered targets

Your load balancer serves as a single point of contact for clients and distributes incoming traffic across its healthy registered targets. You can register each target with one or more target groups.

If demand on your application increases, you can register additional targets with one or more target groups in order to handle the demand. The load balancer starts routing traffic to a newly registered target as soon as the registration process completes and the target passes the first initial health check, irrespective of the configured threshold.

If demand on your application decreases, or you need to service your targets, you can deregister targets from your target groups. Deregistering a target removes it from your target group, but does not affect the target otherwise. The load balancer stops routing requests to a target as soon as it is deregistered. The target enters the `draining` state until in-flight requests have completed. You can register the target with the target group again when you are ready for it to resume receiving requests.

If you are registering targets by instance ID, you can use your load balancer with an Auto Scaling group. After you attach a target group to an Auto Scaling group, Auto Scaling registers your targets with the target group for you when it launches them. For more information, see [Attaching a load balancer to your Auto Scaling group](#) in the *Amazon EC2 Auto Scaling User Guide*.

Limits

- You cannot register the IP addresses of another Application Load Balancer in the same VPC. If the other Application Load Balancer is in a VPC that is peered to the load balancer VPC, you can register its IP addresses.
- You cannot register instances by instance ID if they are in a VPC that is peered to the load balancer VPC (same Region or different Region). You can register these instances by IP address.

Target group attributes

The following target group attributes are supported if the target group type is `instance` or `ip`:

`deregistration_delay.timeout_seconds`

The amount of time for Elastic Load Balancing to wait before deregistering a target. The range is 0–3600 seconds. The default value is 300 seconds.

`load_balancing.algorithm.type`

The load balancing algorithm determines how the load balancer selects targets when routing requests. The value is `round_robin`, `least_outstanding_requests`, or `weighted_random`. The default is `round_robin`.

`load_balancing.algorithm.anomaly_mitigation`

Only available when `load_balancing.algorithm.type` is `weighted_random`. Indicates whether anomaly mitigation is enabled. The value is `on` or `off`. The default is `off`.

`load_balancing.cross_zone.enabled`

Indicates whether cross zone load balancing is enabled. The value is `true`, `false` or `use_load_balancer_configuration`. The default is `use_load_balancer_configuration`.

`slow_start.duration_seconds`

The time period, in seconds, during which the load balancer sends a newly registered target a linearly increasing share of the traffic to the target group. The range is 30–900 seconds (15 minutes). The default is 0 seconds (disabled).

`stickiness.enabled`

Indicates whether sticky sessions are enabled. The value is `true` or `false`. The default is `false`.

`stickiness.app_cookie.cookie_name`

The name of the application cookie. The application cookie name cannot have the following prefixes: `AWSALB`, `AWSALBAPP`, or `AWSALBTG`; they're reserved for use by the load balancer.

`stickiness.app_cookie.duration_seconds`

The application-based cookie expiration period, in seconds. After this period, the cookie is considered stale. The minimum value is 1 second and the maximum value is 7 days (604800 seconds). The default value is 1 day (86400 seconds).

`stickiness.lb_cookie.duration_seconds`

The duration-based cookie expiration period, in seconds. After this period, the cookie is considered stale. The minimum value is 1 second and the maximum value is 7 days (604800 seconds). The default value is 1 day (86400 seconds).

`stickiness.type`

The type of stickiness. The possible values are `lb_cookie` and `app_cookie`.

`target_group_health.dns_failover.minimum_healthy_targets.count`

The minimum number of targets that must be healthy. If the number of healthy targets is below this value, mark the zone as unhealthy in DNS, so that traffic is routed only to healthy zones. The possible values are `off`, or an integer from 1 to the maximum number of targets. When `off`, DNS fail away is disabled, meaning each target group independently contributes to DNS fail over. The default is 1.

`target_group_health.dns_failover.minimum_healthy_targets.percentage`

The minimum percentage of targets that must be healthy. If the percentage of healthy targets is below this value, mark the zone as unhealthy in DNS, so that traffic is routed only to healthy zones. The possible values are `off`, or an integer from 1 to the maximum number of targets. When `off`, DNS fail away is disabled, meaning each target group independently contributes to DNS fail over. The default is 1.

`target_group_health.unhealthy_state_routing.minimum_healthy_targets.count`

The minimum number of targets that must be healthy. If the number of healthy targets is below this value, send traffic to all targets, including unhealthy targets. The range is 1 to the maximum number of targets. The default is 1.

`target_group_health.unhealthy_state_routing.minimum_healthy_targets.percentage`

The minimum percentage of targets that must be healthy. If the percentage of healthy targets is below this value, send traffic to all targets, including unhealthy targets. The possible values are `off` or an integer from 1 to 100. The default is `off`.

The following target group attribute is supported if the target group type is `lambda`:

`lambda.multi_value_headers.enabled`

Indicates whether the request and response headers exchanged between the load balancer and the Lambda function include arrays of values or strings. The possible values are `true` or `false`. The default value is `false`. For more information, see [Multi-value headers](#).

Routing algorithms

A routing algorithm is the method used by the load balancer when determining which targets will receive requests. The **round robin** routing algorithm is used by default to route requests at the target group level. The **least outstanding requests** and **weighted random** routing algorithms are also available based on the needs of your application. A target group can only have one active routing algorithm at a time, however the routing algorithm can be updated whenever needed.

If you enable sticky sessions, the selected routing algorithm is used for the initial target selection. Future requests from the same client will be forwarded to the same target, bypassing the selected routing algorithm.

Round robin

- The round robin routing algorithm routes requests evenly across healthy targets in the target group, in a sequential order.
- This algorithm is commonly used when the requests being received are similar in complexity, the registered targets are similar in processing capability, or if you need to distribute requests equally among targets.

Least outstanding requests

- The least outstanding requests routing algorithm routes requests to the targets with the lowest number of in progress requests.
- This algorithm is commonly used when the requests being received vary in complexity, the registered targets vary in processing capability.
- When a load balancer that supports HTTP/2 is using targets that only support HTTP/1.1, it converts the request to multiple HTTP/1.1 requests. In this configuration the least outstanding requests algorithm will treat each HTTP/2 request as multiple requests.

- When using WebSockets, the target is selected using the least outstanding requests algorithm. Once selected, the load balancer creates a connection to the target and sends all messages over this connection.
- The least outstanding requests routing algorithm can not be used with slow start mode.

Weighted random

- The weighted random routing algorithm routes requests evenly across healthy targets in the target group, in a random order.
- This algorithm supports Automatic Target Weights (ATW) anomaly mitigation.
- The weighted random routing algorithm can not be used with slow start mode.

Modify the routing algorithm of a target group

You can modify the routing algorithm for your target group at any time.

To modify the routing algorithm using the new console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. On the target groups detail page, on the **Attributes** tab, choose **Edit**.
5. On the **Edit target group attributes** page, in the **Traffic configuration** section, under **Load balancing algorithm**, choose **Round robin**, **Least outstanding requests**, or **Weighted random**.
6. Choose **Save changes**.

To modify the routing algorithm using the AWS CLI

Use the [modify-target-group-attributes](#) command with the `load_balancing.algorithm.type` attribute.

Automatic Target Weights (ATW)

Automatic Target Weights (ATW) constantly monitors the targets running your applications, detecting significant performance deviations, known as anomalies. ATW provides the ability

to dynamically adjust the amount of traffic routed to targets, through real time data anomaly detection.

Automatic Target Weights (ATW) performs anomaly detection on every Application Load Balancer in your account automatically. When anomalous targets are identified, ATW can automatically attempt to stabilize them by reducing the amount of traffic they're routed, known as anomaly mitigation. ATW continuously optimizes traffic distribution to maximize per-target success rates while minimizing target group failure rates.

Considerations:

- Anomaly detection currently monitors HTTP 5xx response codes coming from, and connection failures to, your targets. Anomaly detection is always on and cannot be turned off.
- ATW is not supported when using Lambda as a target.

Anomaly detection

ATW anomaly detection monitors for any targets that are displaying a significant deviation in behavior from other targets in their target group. These deviations, called anomalies, are determined by comparing the percent errors of one target with the percent errors of other targets in the target group. These errors can be both connection errors and HTTP error codes. Targets reporting significantly higher than their peers are then considered anomalous.

Anomaly detection requires a minimum of three healthy targets in the target group. When a target is registered to a target group it has to first pass the health checks to start receiving traffic. Once the target is receiving target, ATW begins monitoring the target and continuously publishes the anomaly result. For targets without anomalies, the anomaly result is `normal`. For targets with anomalies, the anomaly result is `anomalous`.

ATW anomaly detection works independently from target group health checks. A target can be passing all target group health checks, but still be marked anomalous due to an elevated error rate. Targets becoming anomalous does not affect their target group health check status.

Anomaly detection status

ATW continuously publishes the status of the anomaly detections it performs on targets. You can view the current status at any time using the AWS Management Console or AWS CLI.

To view anomaly detection status using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. On the target groups detail page, choose the **Targets** tab.
5. Within the **Registered targets** table, you can view each targets anomaly status in the **Anomaly detection result** column.

If no anomalies were detected, the result is normal.

If anomalies were detected, the result is anomalous.

To view anomaly detection results using the AWS CLI

Use the [describe-target-health](#) command with the `Include.member.N` attribute value set to `AnomalyDetection`.

Anomaly mitigation

Important

The anomaly mitigation function of ATW is only available when using the **Weighted random** routing algorithm.

ATW anomaly mitigation routes traffic away from anomalous targets automatically, giving them an opportunity to recover.

During mitigation:

- ATW periodically adjusts the amount of traffic routed to anomalous targets. Currently, the period is every five seconds.
- ATW reduces the amount of traffic routed to anomalous targets to the minimum amount required to perform anomaly mitigation.
- Targets which are no longer detected as anomalous will gradually have more traffic routed to them until they reach parity with other normal targets in the target group.

Turn on ATW anomaly mitigation

You can turn on anomaly mitigation at any time.

To turn on anomaly mitigation using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. On the target groups detail page, on the **Attributes** tab, choose **Edit**.
5. On the **Edit target group attributes** page, in the **Traffic configuration** section, under **Load balancing algorithm**, ensure **Weighted random** is selected.

Note: When the weighted random algorithm is initially selected, anomaly detection is on by default.

6. Under **Anomaly mitigation**, ensure **Turn on anomaly mitigation** is selected.
7. Choose **Save changes**.

To turn on anomaly mitigation using the AWS CLI

Use the [modify-target-group-attributes](#) command with the `load_balancing.algorithm.anomaly_mitigation` attribute.

Anomaly mitigation status

Whenever ATW is performing mitigation on a target, you can view the current status at any time using the AWS Management Console or AWS CLI.

To view anomaly mitigation status using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. On the target groups detail page, choose the **Targets** tab.
5. Within the **Registered targets** table, you can view each targets anomaly mitigation status in the **Mitigation in effect** column.

If mitigation is not in progress, the status is `yes`.

If mitigation is in progress, the status is `no`.

To view anomaly mitigation status using the AWS CLI

Use the [describe-target-health](#) command with the `Include.member.N` attribute value set to `AnomalyDetection`.

Deregistration delay

Elastic Load Balancing stops sending requests to targets that are deregistering. By default, Elastic Load Balancing waits 300 seconds before completing the deregistration process, which can help in-flight requests to the target to complete. To change the amount of time that Elastic Load Balancing waits, update the deregistration delay value.

The initial state of a deregistering target is `draining`. After the deregistration delay elapses, the deregistration process completes and the state of the target is `unused`. If the target is part of an Auto Scaling group, it can be terminated and replaced.

If a deregistering target has no in-flight requests and no active connections, Elastic Load Balancing immediately completes the deregistration process, without waiting for the deregistration delay to elapse. However, even though target deregistration is complete, the status of the target is displayed as `draining` until the deregistration delay timeout expires. After the timeout expires, the target transitions to an `unused` state.

If a deregistering target terminates the connection before the deregistration delay elapses, the client receives a 500-level error response.

To update the deregistration delay value using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. On the **Group details** tab, in the **Attributes** section, choose **Edit**.
5. On the **Edit attributes** page, change the value of **Deregistration delay** as needed.
6. Choose **Save changes**.

To update the deregistration delay value using the AWS CLI

Use the [modify-target-group-attributes](#) command with the `deregistration_delay.timeout_seconds` attribute.

Slow start mode

By default, a target starts to receive its full share of requests as soon as it is registered with a target group and passes an initial health check. Using slow start mode gives targets time to warm up before the load balancer sends them a full share of requests.

After you enable slow start for a target group, its targets enter slow start mode when they are considered healthy by the target group. A target in slow start mode exits slow start mode when the configured slow start duration period elapses or the target becomes unhealthy. The load balancer linearly increases the number of requests that it can send to a target in slow start mode. After a healthy target exits slow start mode, the load balancer can send it a full share of requests.

Considerations

- When you enable slow start for a target group, the healthy targets registered with the target group do not enter slow start mode.
- When you enable slow start for an empty target group and then register targets using a single registration operation, these targets do not enter slow start mode. Newly registered targets enter slow start mode only when there is at least one healthy target that is not in slow start mode.
- If you deregister a target in slow start mode, the target exits slow start mode. If you register the same target again, it enters slow start mode when it is considered healthy by the target group.
- If a target in slow start mode becomes unhealthy, the target exits slow start mode. When the target becomes healthy, it enters slow start mode again.
- You cannot enable slow start mode when using the **least outstanding requests** or **weighted random** routing algorithms.

To update the slow start duration value using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.

3. Choose the name of the target group to open its details page.
4. On the **Group details** tab, in the **Attributes** section, choose **Edit**.
5. On the **Edit attributes** page, change the value of **Slow start duration** as needed. To disable slow start mode, set the duration to 0.
6. Choose **Save changes**.

To update the slow start duration value using the AWS CLI

Use the [modify-target-group-attributes](#) command with the `slow_start.duration_seconds` attribute.

Create a target group

You register your targets with a target group. By default, the load balancer sends requests to registered targets using the port and protocol that you specified for the target group. You can override this port when you register each target with the target group.

After you create a target group, you can add tags.

To route traffic to the targets in a target group, specify the target group in an action when you create a listener or create a rule for your listener. For more information, see [Listener rules](#). You can specify the same target group in multiple listeners, but these listeners must belong to the same Application Load Balancer. To use a target group with a load balancer, you must verify that the target group is not in use by a listener for any other load balancer.

You can add or remove targets from your target group at any time. For more information, see [Register targets with your target group](#). You can also modify the health check settings for your target group. For more information, see [Modify the health check settings of a target group](#).

To create a target group using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose **Create target group**.
4. For **Choose a target type**, select **Instances** to register targets by instance ID, **IP addresses** to register targets by IP address, or **Lambda function** to register a Lambda function as a target.

5. For **Target group name**, type a name for the target group. This name must be unique per region per account, can have a maximum of 32 characters, must contain only alphanumeric characters or hyphens, and must not begin or end with a hyphen.
6. (Optional) For **Protocol** and **Port**, modify the default values as needed.
7. If the target type is **Instances** or **IP addresses**, choose **IPv4** or **IPv6** as the **IP address type**, otherwise skip to the next step.

Note that only targets that have the selected IP address type can be included in this target group. The IP address type cannot be changed after the target group is created.

8. For **VPC**, select a virtual private cloud (VPC). Note that for **IP addresses** target types, the VPCs available for selection are those that support the **IP address type** that you chose in the previous step.
9. (Optional) For **Protocol version**, modify the default value as needed.
10. (Optional) In the **Health checks** section, modify the default settings as needed.
11. If the target type is **Lambda function**, you can enable health checks by selecting **Enable** in the **Health checks** section.
12. (Optional) Add one or more tags as follows:
 - a. Expand the **Tags** section.
 - b. Choose **Add tag**.
 - c. Enter the tag key and the tag value.

13. Choose **Next**.

14. (Optional) Add one or more targets as follows:

- If the target type is **Instances**, select one or more instances, enter one or more ports, and then choose **Include as pending below**.

Note: The instances must have an assigned primary IPv6 address to be registered with a IPv6 target group.

- If the target type is **IP addresses**, do the following:
 - a. Select a network **VPC** from the list, or choose **Other private IP addresses**.
 - b. Enter the IP address manually, or find the IP address using instance details. You can enter up to five IP addresses at a time.
 - c. Enter the ports for routing traffic to the specified IP addresses.
 - d. Choose **Include as pending below**.

- If the target type is a **Lambda function**, specify a single Lambda function or omit this step and specify a Lambda function later.
15. Choose **Create target group**.
 16. (Optional) You can specify the target group in a listener rule. For more information, see [Listener Rules](#).

To create a target group using the AWS CLI

Use the [create-target-group](#) command to create the target group, the [add-tags](#) command to tag your target group, and the [register-targets](#) command to add targets.

Health checks for your target groups

Your Application Load Balancer periodically sends requests to its registered targets to test their status. These tests are called *health checks*.

Each load balancer node routes requests only to the healthy targets in the enabled Availability Zones for the load balancer. Each load balancer node checks the health of each target, using the health check settings for the target groups with which the target is registered. After your target is registered, it must pass one health check to be considered healthy. After each health check is completed, the load balancer node closes the connection that was established for the health check.

If a target group contains only unhealthy registered targets, the load balancer routes requests to all those targets, regardless of their health status. This means that if all targets fail health checks at the same time in all enabled Availability Zones, the load balancer fails open. The effect of the fail-open is to allow traffic to all targets in all enabled Availability Zones, regardless of their health status, based on the load balancing algorithm.

Health checks do not support WebSockets.

Health check settings

You configure health checks for the targets in a target group as described in the following table. The setting names used in the table are the names used in the API. The load balancer sends a health check request to each registered target every **HealthCheckIntervalSeconds** seconds, using the specified port, protocol, and health check path. Each health check request is independent and the result lasts for the entire interval. The time that it takes for the target to respond does not affect the interval for the next health check request. If the health checks exceed

UnhealthyThresholdCount consecutive failures, the load balancer takes the target out of service. When the health checks exceed **HealthyThresholdCount** consecutive successes, the load balancer puts the target back in service.

Setting	Description
HealthCheckProtocol	<p>The protocol the load balancer uses when performing health checks on targets. For Application Load Balancers the possible protocols are HTTP and HTTPS. The default is the HTTP protocol.</p> <p>These protocols use the HTTP GET method to send health check requests.</p>
HealthCheckPort	<p>The port the load balancer uses when performing health checks on targets. The default is to use the port on which each target receives traffic from the load balancer.</p>
HealthCheckPath	<p>The destination for health checks on the targets.</p> <p>If the protocol version is HTTP/1.1 or HTTP/2, specify a valid URI (<i>/path?query</i>). The default is <code>/</code>.</p> <p>If the protocol version is gRPC, specify the path of a custom health check method with the format <code>/package.service/method</code>. The default is <code>/AWS.ALB/healthcheck</code>.</p>
HealthCheckTimeoutSeconds	<p>The amount of time, in seconds, during which no response from a target means a failed health check. The range is 2–120 seconds. The default is 5 seconds if the target type is <code>instance</code> or <code>ip</code> and 30 seconds if the target type is <code>lambda</code>.</p>

Setting	Description
HealthCheckIntervalSeconds	The approximate amount of time, in seconds, between health checks of an individual target. The range is 5–300 seconds. The default is 30 seconds if the target type is <code>instance</code> or <code>ip</code> and 35 seconds if the target type is <code>lambda</code> .
HealthyThresholdCount	The number of consecutive successful health checks required before considering an unhealthy target healthy. The range is 2–10. The default is 5.
UnhealthyThresholdCount	The number of consecutive failed health checks required before considering a target unhealthy. The range is 2–10. The default is 2.
Matcher	<p>The codes to use when checking for a successful response from a target. These are called Success codes in the console.</p> <p>If the protocol version is HTTP/1.1 or HTTP/2, the possible values are from 200 to 499. You can specify multiple values (for example, "200,202") or a range of values (for example, "200-299"). The default value is 200.</p> <p>If the protocol version is gRPC, the possible values are from 0 to 99. You can specify multiple values (for example, "0,1") or a range of values (for example, "0-5"). The default value is 12.</p>

Target health status

Before the load balancer sends a health check request to a target, you must register it with a target group, specify its target group in a listener rule, and ensure that the Availability Zone of the target

is enabled for the load balancer. Before a target can receive requests from the load balancer, it must pass the initial health checks. After a target passes the initial health checks, its status is Healthy.

The following table describes the possible values for the health status of a registered target.

Value	Description
<code>initial</code>	<p>The load balancer is in the process of registering the target or performing the initial health checks on the target.</p> <p>Related reason codes: <code>Elb.RegistrationInProgress</code> <code>Elb.InitialHealthChecking</code></p>
<code>healthy</code>	<p>The target is healthy.</p> <p>Related reason codes: None</p>
<code>unhealthy</code>	<p>The target did not respond to a health check or failed the health check.</p> <p>Related reason codes: <code>Target.ResponseCodeMismatch</code> <code>Target.Timeout</code> <code>Target.FailedHealthChecks</code> <code>Elb.InternalError</code></p>
<code>unused</code>	<p>The target is not registered with a target group, the target group is not used in a listener rule, the target is in an Availability Zone that is not enabled, or the target is in the stopped or terminated state.</p> <p>Related reason codes: <code>Target.NotRegistered</code> <code>Target.NotInUse</code> <code>Target.InvalidState</code> <code>Target.IpUnusable</code></p>
<code>draining</code>	<p>The target is deregistering and connection draining is in process.</p> <p>Related reason code: <code>Target.DeregistrationInProgress</code></p>

Value	Description
unavailable	Health checks are disabled for the target group. Related reason code: Target.HealthCheck Disabled

Health check reason codes

If the status of a target is any value other than `Healthy`, the API returns a reason code and a description of the issue, and the console displays the same description. Reason codes that begin with `Elb` originate on the load balancer side and reason codes that begin with `Target` originate on the target side. For more information about possible causes for health check failures, see [Troubleshooting](#).

Reason code	Description
<code>Elb.InitialHealthChecking</code>	Initial health checks in progress
<code>Elb.InternalError</code>	Health checks failed due to an internal error
<code>Elb.RegistrationInProgress</code>	Target registration is in progress
<code>Target.DeregistrationInProgress</code>	Target deregistration is in progress
<code>Target.FailedHealthChecks</code>	Health checks failed
<code>Target.HealthCheckDisabled</code>	Health checks are disabled
<code>Target.InvalidState</code>	Target is in the stopped state Target is in the terminated state Target is in the terminated or stopped state Target is in an invalid state

Reason code	Description
Target.IpUnusable	The IP address cannot be used as a target, as it is in use by a load balancer
Target.NotInUse	Target group is not configured to receive traffic from the load balancer Target is in an Availability Zone that is not enabled for the load balancer
Target.NotRegistered	Target is not registered to the target group
Target.ResponseCodeMismatch	Health checks failed with these codes: [code]
Target.Timeout	Request timed out

Check the health of your targets

You can check the health status of the targets registered with your target groups.

To check the health of your targets using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. On the **Targets** tab, the **Status** column indicates the status of each target.
5. If the status is any value other than **Healthy**, the **Status details** column contains more information. For help with health check failures, see [Troubleshooting](#).

To check the health of your targets using the AWS CLI

Use the [describe-target-health](#) command. The output of this command contains the target health state. If the status is any value other than **Healthy**, the output also includes a reason code.

To receive email notifications about unhealthy targets

Use CloudWatch alarms to trigger a Lambda function to send details about unhealthy targets. For step-by-step instructions, see the following blog post: [Identifying unhealthy targets of your load balancer](#).

Modify the health check settings of a target group

You can modify the health check settings for your target group at any time.

To modify the health check settings of a target group using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. On the **Group details** tab, in the **Health check settings** section, choose **Edit**.
5. On the **Edit health check settings** page, modify the settings as needed, and then choose **Save changes**.

To modify the health check settings of a target group using the AWS CLI

Use the [modify-target-group](#) command.

Cross-zone load balancing for target groups

The nodes for your load balancer distribute requests from clients to registered targets. When cross-zone load balancing is on, each load balancer node distributes traffic across the registered targets in all registered Availability Zones. When cross-zone load balancing is off, each load balancer node distributes traffic only across the registered targets in its Availability Zone. This could be if zonal failure domains are preferred over regional, ensuring that a healthy zone isn't impacted by an unhealthy zone, or for overall latency improvements.

With Application Load Balancers, cross-zone load balancing is always turned on at the load balancer level, and cannot be turned off. For target groups, the default is to use the load balancer setting, but you can override the default by explicitly turning cross-zone load balancing off at the target group level.

Considerations

- Target stickiness is not supported when cross-zone load balancing is off.

- Lambda functions as targets are not supported when cross-zone load balancing is off.
- Attempting to turn off cross-zone load balancing through the `ModifyTargetGroupAttributes` API if any targets have parameter `AvailabilityZone` set to `all` results in an error.
- When registering targets, the `AvailabilityZone` parameter is required. Specific Availability Zone values are only allowed when cross-zone load balancing is off. Otherwise, the parameter is ignored and treated as `all`.

Best practices

- Plan for enough target capacity across all Availability Zones that you expect to utilize, per target group. If you can't plan for enough capacity across all participating Availability Zones, we recommend that you keep cross-zone load balancing on.
- When configuring your Application Load Balancer with multiple target groups, ensure all target groups are participating in the same Availability Zones, within the configured Region. This is to avoid an Availability Zone being empty while cross-zone load balancing is off, as this triggers a 503 error for all HTTP requests that enter the empty Availability Zone.
- Avoid creating empty subnets. Application Load Balancers expose zonal IP addresses through DNS for the empty subnets, which triggers 503 errors for HTTP requests.
- There can be occurrences where a target group with cross-zone load balancing turned off has enough planned target capacity per Availability Zone, but all targets in an Availability Zone become unhealthy. When there is at least one target group with all unhealthy targets, the IP addresses of the load balancer nodes are removed from DNS. After the target group has at least one healthy target, the IP addresses are restored to DNS.

Turn off cross-zone load balancing

You can turn off cross-zone load balancing for your Application Load Balancer target groups at any time.

To turn off cross-zone load balancing using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, select **Target Groups**.
3. Select the name of the target group to open its details page.

4. On the **Attributes** tab, select **Edit**.
5. On the **Edit target group attributes** page, select **Off** for **Cross-zone load balancing**.
6. Choose **Save changes**.

To turn off cross-zone load balancing using the AWS CLI

Use the [modify-target-group-attributes](#) command and set the `load_balancing.cross_zone.enabled` attribute to `false`.

```
aws elbv2 modify-target-group-attributes --target-group-arn my-targetgroup-arn --attributes Key=load_balancing.cross_zone.enabled,Value=false
```

The following is an example response:

```
{
  "Attributes": [
    {
      "Key": "load_balancing.cross_zone.enabled",
      "Value": "false"
    },
  ]
}
```

Turn on cross-zone load balancing

You can turn on cross-zone load balancing for your Application Load Balancer target groups at any time. The cross-zone load balancing setting at the target group level overrides the setting at the load balancer level.

To turn on cross-zone load balancing using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, select **Target Groups**.
3. Select the name of the target group to open its details page.
4. On the **Attributes** tab, select **Edit**.
5. On the **Edit target group attributes** page, select **On** for **Cross-zone load balancing**.
6. Choose **Save changes**.

To turn on cross-zone load balancing using the AWS CLI

Use the [modify-target-group-attributes](#) command and set the `load_balancing.cross_zone.enabled` attribute to `true`.

```
aws elbv2 modify-target-group-attributes --target-group-arn my-targetgroup-arn --attributes Key=load_balancing.cross_zone.enabled,Value=true
```

The following is an example response:

```
{
  "Attributes": [
    {
      "Key": "load_balancing.cross_zone.enabled",
      "Value": "true"
    },
  ]
}
```

Target group health

By default, a target group is considered healthy as long as it has at least one healthy target. If you have a large fleet, having only one healthy target serving traffic is not sufficient. Instead, you can specify a minimum count or percentage of targets that must be healthy, and what actions the load balancer takes when the healthy targets fall below the specified threshold. This improves availability.

Unhealthy state actions

You can configure healthy thresholds for the following actions:

- **DNS failover** – When the healthy targets in a zone fall below the threshold, we mark the IP addresses of the load balancer node for the zone as unhealthy in DNS. Therefore, when clients resolve the load balancer DNS name, the traffic is routed only to healthy zones.
- **Routing failover** – When the healthy targets in a zone fall below the threshold, the load balancer sends traffic to all targets that are available to the load balancer node, including unhealthy targets. This increases the chances that a client connection succeeds, especially when targets temporarily fail to pass health checks, and reduces the risk of overloading the healthy targets.

Requirements and considerations

- You can't use this feature with target groups where the target is a Lambda function. If the Application Load Balancer is the target of a Network Load Balancer or Global Accelerator, do not configure a threshold for DNS failover.
- If you specify both types of thresholds for an action (count and percentage), the load balancer takes the action when either threshold is breached.
- If you specify thresholds for both actions, the threshold for DNS failover must be greater than or equal to the threshold for routing failover, so that DNS failover occurs either with or before routing failover.
- If you specify the threshold as a percentage, we calculate the value dynamically, based on the total number of targets that are registered with the target groups.
- The total number of targets is based on whether cross-zone load balancing is off or on. If cross-zone load balancing is off, each node sends traffic only to the targets in its own zone, which means that the thresholds apply to the number of targets in each enabled zone separately. If cross-zone load balancing is on, each node sends traffic to all targets in all enabled zones, which means that the specified thresholds apply to the total number targets in all enabled zones.
- With DNS failover, we remove the IP addresses for the unhealthy zones from the DNS hostname for the load balancer. However, the local client DNS cache might contain these IP addresses until the time-to-live (TTL) in the DNS record expires (60 seconds).
- When DNS failover occurs, this impacts all target groups associated with the load balancer. Ensure that you have enough capacity in your remaining zones to handle this additional traffic, especially if cross-zone load balancing is off.
- With DNS failover, if all load balancer zones are considered unhealthy, the load balancer sends traffic to all zones, including the unhealthy zones.
- There are factors other than whether there are enough healthy targets that might lead to DNS failover, such as the health of the zone.

Monitoring

To monitor the health of your target groups, see [CloudWatch metrics for target group health](#).

Example

The following example demonstrates how target group health settings are applied.

Scenario

- A load balancer that supports two Availability Zones, A and B
- Each Availability Zone contains 10 registered targets
- The target group has the following target group health settings:
 - DNS failover - 50%
 - Routing failover - 50%
- Six targets fail in Availability Zone B

If cross-zone load balancing is off

- The load balancer node in each Availability Zone can send traffic only to the 10 targets in its Availability Zone.
- There are 10 healthy targets in Availability Zone A, which meets the required percentage of healthy targets. The load balancer continues to distribute traffic between the 10 healthy targets.
- There are only 4 healthy targets in Availability Zone B, which is 40% of the targets for the load balancer node in Availability Zone B. Because this is less than the required percentage of healthy targets, the load balancer takes the following actions:
 - DNS failover - Availability Zone B is marked as unhealthy in DNS. Because clients can't resolve the load balancer name to the load balancer node in Availability Zone B, and Availability Zone A is healthy, clients send new connections to Availability Zone A.
 - Routing failover - When new connections are sent explicitly to Availability Zone B, the load balancer distributes traffic to all targets in Availability Zone B, including the unhealthy targets. This prevents outages among the remaining healthy targets.

If cross-zone load balancing is on

- Each load balancer node can send traffic to all 20 registered targets across both Availability Zones.
- There are 10 healthy targets in Availability Zone A and 4 healthy targets in Availability Zone B, for a total of 14 healthy targets. This is 70% of the targets for the load balancer nodes in both Availability Zones, which meets the required percentage of healthy targets.
- The load balancer distributes traffic between the 14 healthy targets in both Availability Zones.

Modify target group health settings

You can modify the target group health settings for your target group as follows.

To modify target group health settings using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. On the **Attributes** tab, choose **Edit**.
5. Check whether cross-zone load balancing is turned on or turned off. Update this setting as needed to ensure that you have enough capacity to handle the additional traffic if a zone fails.
6. Expand **Target group health requirements**.
7. For **Configuration type**, we recommend that you choose **Unified configuration**, which sets the same threshold for both actions.
8. For **Healthy state requirements**, do one of the following:
 - Choose **Minimum healthy target count**, and then enter a number from 1 to the maximum number of targets for your target group.
 - Choose **Minimum healthy target percentage**, and then enter a number from 1 to 100.
9. Choose **Save changes**.

To modify target group health settings using the AWS CLI

Use the [modify-target-group-attributes](#) command. The following example sets the healthy threshold for both unhealthy state actions to 50%.

```
aws elbv2 modify-target-group-attributes \  
--target-group-arn arn:aws:elasticloadbalancing:region:123456789012:targetgroup/my-  
targets/73e2d6bc24d8a067 \  
--attributes  
Key=target_group_health.dns_failover.minimum_healthy_targets.percentage,Value=50 \  
  
Key=target_group_health.unhealthy_state_routing.minimum_healthy_targets.percentage,Value=50
```

Using Route 53 DNS failover for your load balancer

If you use Route 53 to route DNS queries to your load balancer, you can also configure DNS failover for your load balancer using Route 53. In a failover configuration, Route 53 checks the health of the target group targets for the load balancer to determine whether they are available. If there are no healthy targets registered with the load balancer, or if the load balancer itself is unhealthy, Route 53 routes traffic to another available resource, such as a healthy load balancer or a static website in Amazon S3.

For example, suppose that you have a web application for `www.example.com`, and you want redundant instances running behind two load balancers residing in different Regions. You want the traffic to be primarily routed to the load balancer in one Region, and you want to use the load balancer in the other Region as a backup during failures. If you configure DNS failover, you can specify your primary and secondary (backup) load balancers. Route 53 directs traffic to the primary load balancer if it is available, or to the secondary load balancer otherwise.

Using evaluate target health

- When evaluate target health is set to Yes on an alias record for an Application Load Balancer, Route 53 evaluates the health of the resource specified by the `alias target` value. For an Application Load Balancer, Route 53 uses the target group health checks associated with the load balancer.
- When all the target groups in an Application Load Balancer are healthy, Route 53 marks the alias record as healthy. If a target group contains at least one healthy target, the target group health check passes. Route 53 then returns records according to your routing policy. If the failover routing policy is used, Route 53 returns the primary record.
- If any of the target groups in an Application Load Balancer are unhealthy, the alias record fails the Route 53 health check (fail-open). If using evaluate target health, this would fail the failover routing policy.
- If all of the target groups in an Application Load Balancer are empty (no targets), then Route 53 considers the record unhealthy (fail-open). If using evaluate target health, this would fail the failover routing policy.

For more information, see [Configuring DNS failover](#) in the *Amazon Route 53 Developer Guide*.

Register targets with your target group

You register your targets with a target group. When you create a target group, you specify its target type, which determines how you register its targets. For example, you can register instance IDs, IP addresses, or Lambda functions. For more information, see [Target groups for your Application Load Balancers](#).

If demand on your currently registered targets increases, you can register additional targets in order to handle the demand. When your target is ready to handle requests, register it with your target group. The load balancer starts routing requests to the target as soon as the registration process completes and the target passes the initial health checks.

If demand on your registered targets decreases, or you need to service a target, you can deregister it from your target group. The load balancer stops routing requests to a target as soon as you deregister it. When the target is ready to receive requests, you can register it with the target group again.

When you deregister a target, the load balancer waits until in-flight requests have completed. This is known as *connection draining*. The status of a target is `draining` while connection draining is in progress.

When you deregister a target that was registered by IP address, you must wait for the deregistration delay to complete before you can register the same IP address again.

If you are registering targets by instance ID, you can use your load balancer with an Auto Scaling group. After you attach a target group to an Auto Scaling group and the group scales out, the instances launched by the Auto Scaling group are automatically registered with the target group. If you detach the target group from the Auto Scaling group, the instances are automatically deregistered from the target group. For more information, see [Attaching a load balancer to your Auto Scaling group](#) in the *Amazon EC2 Auto Scaling User Guide*.

Target security groups

When you register EC2 instances as targets, you must ensure that the security groups for your instances allow the load balancer to communicate with your instances on both the listener port and the health check port.

Recommended rules

Inbound

Source	Port Range	Comment
<i>load balancer security group</i>	<i>instance listener</i>	Allow traffic from the load balancer on the instance listener port
<i>load balancer security group</i>	<i>health check</i>	Allow traffic from the load balancer on the health check port

We also recommend that you allow inbound ICMP traffic to support Path MTU Discovery. For more information, see [Path MTU Discovery](#) in the *Amazon EC2 User Guide*.

Shared subnets

Participants can create an Application Load Balancer in a shared VPC. Participants can't register a target that runs in a subnet that is not shared with them.

Register or deregister targets

The target type of your target group determines how you register targets with that target group. For more information, see [Target type](#).

Contents

- [Register or deregister targets by instance ID](#)
- [Register or deregister targets by IP address](#)
- [Register or deregister a Lambda function](#)
- [Register or deregister targets using the AWS CLI](#)

Register or deregister targets by instance ID

Note

When registering targets by instance ID for a IPv6 target group, the targets must have an assigned primary IPv6 address. To learn more, see [IPv6 addresses](#) in the *Amazon EC2 User Guide*

The instance must be in the virtual private cloud (VPC) that you specified for the target group. The instance must also be in the running state when you register it.

To register or deregister targets by instance ID using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. Choose the **Targets** tab.
5. To register instances, choose **Register targets**. Select one or more instances, enter the default instance port as needed, and then choose **Include as pending below**. When you are finished adding instances, choose **Register pending targets**.

Note:

- The instances must have an assigned primary IPv6 address to be registered with a IPv6 target group.
 - AWS GovCloud (US) Regions don't support assigning a primary IPv6 address using the console. You must use the API to assign primary IPv6 addresses in AWS GovCloud (US) Regions.
6. To deregister instances, select the instances and then choose **Deregister**.

Register or deregister targets by IP address

IPv4 targets

The IP addresses that you register must be from one of the following CIDR blocks:

- The subnets of the VPC for the target group
- 10.0.0.0/8 (RFC 1918)
- 100.64.0.0/10 (RFC 6598)
- 172.16.0.0/12 (RFC 1918)
- 192.168.0.0/16 (RFC 1918)

You cannot register the IP addresses of another Application Load Balancer in the same VPC. If the other Application Load Balancer is in a VPC that is peered to the load balancer VPC, you can register its IP addresses.

IPv6 targets

- The IP addresses that you register must be within the VPC CIDR block or within a peered VPC CIDR block.

To register or deregister targets by IP address using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. Choose the **Targets** tab.
5. To register IP addresses, choose **Register targets**. For each IP address, select the network, enter the IP address and port, and choose **Include as pending below**. When you are finished specifying addresses, choose **Register pending targets**.
6. To deregister IP addresses, select the IP addresses and then choose **Deregister**. If you have many registered IP addresses, you might find it helpful to add a filter or change the sort order.

Register or deregister a Lambda function

You can register a single Lambda function with each target group. Elastic Load Balancing must have permissions to invoke the Lambda function. If you no longer need to send traffic to your Lambda function, you can deregister it. After you deregister a Lambda function, in-flight requests fail with HTTP 5XX errors. To replace a Lambda function, it is better to create a new target group instead. For more information, see [Lambda functions as targets](#).

To register or deregister a Lambda function using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. Choose the **Targets** tab.

5. If there is no Lambda function registered, choose **Register**. Select the Lambda function and choose **Register**.
6. To deregister a Lambda function, choose **Deregister**. When prompted for confirmation, choose **Deregister**.

Register or deregister targets using the AWS CLI

Use the [register-targets](#) command to add targets and the [deregister-targets](#) command to remove targets.

Sticky sessions for your Application Load Balancer

By default, an Application Load Balancer routes each request independently to a registered target based on the chosen load-balancing algorithm. However, you can use the sticky session feature (also known as session affinity) to enable the load balancer to bind a user's session to a specific target. This ensures that all requests from the user during the session are sent to the same target. This feature is useful for servers that maintain state information in order to provide a continuous experience to clients. To use sticky sessions, the client must support cookies.

Application Load Balancers support both duration-based cookies and application-based cookies. Sticky sessions are enabled at the target group level. You can use a combination of duration-based stickiness, application-based stickiness, and no stickiness across your target groups.

The key to managing sticky sessions is determining how long your load balancer should consistently route the user's request to the same target. If your application has its own session cookie, then you can use application-based stickiness and the load balancer session cookie follows the duration specified by the application's session cookie. If your application does not have its own session cookie, then you can use duration-based stickiness to generate a load balancer session cookie with a duration that you specify.

The content of load balancer generated cookies are encrypted using a rotating key. You cannot decrypt or modify load balancer generated cookies.

For both stickiness types, the Application Load Balancer resets the expiry of the cookies it generates after every request. If a cookie expires, the session is no longer sticky and the client should remove the cookie from its cookie store.

Requirements

- An HTTP/HTTPS load balancer.
- At least one healthy instance in each Availability Zone.

Considerations

- Sticky sessions are not supported if [cross-zone load balancing is disabled](#). Attempting to enable sticky sessions while cross-zone load balancing is disabled will fail.
- For application-based cookies, cookie names have to be specified individually for each target group. However, for duration-based cookies, AWSALB is the only name used across all target groups.
- If you are using multiple layers of Application Load Balancers, you can enable sticky sessions across all layers with application-based cookies. However, with duration-based cookies, you can enable sticky sessions only on one layer, because AWSALB is the only name available.
- If the Application Load Balancer receives both an AWSALBCORS and an AWSALB duration-based stickiness cookie, the value in AWSALBCORS will take precedence.
- Application-based stickiness does not work with weighted target groups.
- If you have a [forward action](#) with multiple target groups, and sticky sessions are enabled for one or more of the target groups, you must enable stickiness at the target group level.
- WebSocket connections are inherently sticky. If the client requests a connection upgrade to WebSockets, the target that returns an HTTP 101 status code to accept the connection upgrade is the target used in the WebSockets connection. After the WebSockets upgrade is complete, cookie-based stickiness is not used.
- Application Load Balancers use the `Expires` attribute in the cookie header instead of the `Max-Age` attribute.
- Application Load Balancers do not support cookie values that are URL encoded.

Duration-based stickiness

Duration-based stickiness routes requests to the same target in a target group using a load balancer generated cookie (AWSALB). The cookie is used to map the session to the target. If your application does not have its own session cookie, you can specify your own stickiness duration and manage how long your load balancer should consistently route the user's request to the same target.

When a load balancer first receives a request from a client, it routes the request to a target (based on the chosen algorithm), and generates a cookie named AWSALB. It encodes information about the selected target, encrypts the cookie, and includes the cookie in the response to the client. The load balancer generated cookie has its own expiry of 7 days which is non-configurable.

In subsequent requests, the client should include the AWSALB cookie. When the load balancer receives a request from a client that contains the cookie, it detects it and routes the request to the same target. If the cookie is present but cannot be decoded, or if it refers to a target that was deregistered or is unhealthy, the load balancer selects a new target and updates the cookie with information about the new target.

For cross-origin resource sharing (CORS) requests, some browsers require `SameSite=None; Secure` to enable stickiness. To support these browsers the load balancer always generates a second stickiness cookie, AWSALBCORS, which includes the same information as the original stickiness cookie, as well as the `SameSite` attribute. Clients receive both cookies, including non CORS requests.

To enable duration-based stickiness using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. On the **Group details** tab, in the **Attributes** section, choose **Edit**.
5. On the **Edit attributes** page, do the following:
 - a. Select **Stickiness**.
 - b. For **Stickiness type**, select **Load balancer generated cookie**.
 - c. For **Stickiness duration**, specify a value between 1 second and 7 days.
 - d. Choose **Save changes**.

To enable duration-based stickiness using the AWS CLI

Use the [modify-target-group-attributes](#) command with the `stickiness.enabled` and `stickiness.lb_cookie.duration_seconds` attributes.

Use the following command to enable duration-based stickiness.

```
aws elbv2 modify-target-group-attributes --target-group-arn ARN --attributes
Key=stickiness.enabled,Value=true
Key=stickiness.lb_cookie.duration_seconds,Value=time-in-seconds
```

Your output should be similar to the following example.

```
{
  "Attributes": [
    ...
    {
      "Key": "stickiness.enabled",
      "Value": "true"
    },
    {
      "Key": "stickiness.lb_cookie.duration_seconds",
      "Value": "86500"
    },
    ...
  ]
}
```

Application-based stickiness

Application-based stickiness gives you the flexibility to set your own criteria for client-target stickiness. When you enable application-based stickiness, the load balancer routes the first request to a target within the target group based on the chosen algorithm. The target is expected to set a custom application cookie that matches the cookie configured on the load balancer to enable stickiness. This custom cookie can include any of the cookie attributes required by the application.

When the Application Load Balancer receives the custom application cookie from the target, it automatically generates a new encrypted application cookie to capture stickiness information. This load balancer generated application cookie captures stickiness information for each target group that has application-based stickiness enabled.

The load balancer generated application cookie does not copy the attributes of the custom cookie set by the target. It has its own expiry of 7 days which is non-configurable. In the response to the client, the Application Load Balancer only validates the name with which the custom cookie was configured at the target group level and not the value or the expiry attribute of the custom cookie.

As long as the name matches, the load balancer sends both cookies, the custom cookie set by the target, and the application cookie generated by the load balancer, in the response to the client.

In subsequent requests, clients have to send back both cookies to maintain stickiness. The load balancer decrypts the application cookie, and checks whether the configured duration of stickiness is still valid. It then uses the information in the cookie to send the request to the same target within the target group to maintain stickiness. The load balancer also proxies the custom application cookie to the target without inspecting or modifying it. In subsequent responses, the expiry of the load balancer generated application cookie and the duration of stickiness configured on the load balancer are reset. To maintain stickiness between client and target, the expiry of the cookie, and the duration of stickiness should not elapse.

If a target fails or becomes unhealthy, the load balancer stops routing requests to that target, and chooses a new healthy target based on the chosen load balancing algorithm. The load balancer treats the session as now being "stuck" to the new healthy target, and continues routing requests to the new healthy target even if the failed target comes back.

With cross-origin resource sharing (CORS) requests, to enable stickiness, the load balancer adds the `SameSite=None; Secure` attributes to the load balancer generated application cookie only if the user-agent version is Chromium80 or above.

Since most browsers limit cookies to 4K in size, the load balancer shards application cookies greater than 4K into multiple cookies. Application Load Balancers support cookies up to 16K in size and can therefore create up to 4 shards that it sends to the client. The application cookie name that the client sees begins with "AWSALBAPP-" and includes a fragment number. For example, if the cookie size is 0-4K, the client sees AWSALBAPP-0. If the cookie size is 4-8k, the client sees AWSALBAPP-0 and AWSALBAPP-1, and so on.

To enable application-based stickiness using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. On the **Group details** tab, in the **Attributes** section, choose **Edit**.
5. On the **Edit attributes** page, do the following:
 - a. Select **Stickiness**.
 - b. For **Stickiness type**, select **Application-based cookie**.

- c. For **Stickiness duration**, specify a value between 1 second and 7 days.
- d. For **App cookie name**, enter a name for your application-based cookie.

Do not use AWSALB, AWSALBAPP, or AWSALBTG for the cookie name; they're reserved for use by the load balancer.

- e. Choose **Save changes**.

To enable application-based stickiness using the AWS CLI

Use the [modify-target-group-attributes](#) command with the following attributes:

- `stickiness.enabled`
- `stickiness.type`
- `stickiness.app_cookie.cookie_name`
- `stickiness.app_cookie.duration_seconds`

Use the following command to enable application-based stickiness.

```
aws elbv2 modify-target-group-attributes --target-group-arn ARN --attributes
Key=stickiness.enabled,Value=true Key=stickiness.type,Value=app_cookie
Key=stickiness.app_cookie.cookie_name,Value=my-cookie-name
Key=stickiness.app_cookie.duration_seconds,Value=time-in-seconds
```

Your output should be similar to the following example.

```
{
  "Attributes": [
    ...
    {
      "Key": "stickiness.enabled",
      "Value": "true"
    },
    {
      "Key": "stickiness.app_cookie.cookie_name",
      "Value": "MyCookie"
    },
    {
```



```
        "Key": "stickiness.type",
        "Value": "app_cookie"
    },
    {
        "Key": "stickiness.app_cookie.duration_seconds",
        "Value": "86500"
    },
    ...
]
}
```

Manual rebalancing

When scaling up, if the number of targets increase considerably, there is potential for unequal distribution of load due to stickiness. In this scenario, you can rebalance the load on your targets using the following two options:

- Set an expiry on the cookie generated by the application that is prior to the current date and time. This will prevent clients from sending the cookie to the Application Load Balancer, which will restart the process of establishing stickiness.
- Set a very short duration on the load balancer's application-based stickiness configuration, for example, 1 second. This forces the Application Load Balancer to reestablish stickiness even if the cookie set by the target has not expired.

Lambda functions as targets

You can register your Lambda functions as targets and configure a listener rule to forward requests to the target group for your Lambda function. When the load balancer forwards the request to a target group with a Lambda function as a target, it invokes your Lambda function and passes the content of the request to the Lambda function, in JSON format.

Limits

- The Lambda function and target group must be in the same account and in the same Region.
- The maximum size of the request body that you can send to a Lambda function is 1 MB. For related size limits, see [HTTP header limits](#).
- The maximum size of the response JSON that the Lambda function can send is 1 MB.

- WebSockets are not supported. Upgrade requests are rejected with an HTTP 400 code.
- Local Zones are not supported.
- Automatic Target Weights (ATW) is not supported.

Contents

- [Prepare the Lambda function](#)
- [Create a target group for the Lambda function](#)
- [Receive events from the load balancer](#)
- [Respond to the load balancer](#)
- [Multi-value headers](#)
- [Enable health checks](#)
- [Deregister the Lambda function](#)

For a demo, see [Lambda target on Application Load Balancer](#).

Prepare the Lambda function

The following recommendations apply if you are using your Lambda function with an Application Load Balancer.

Permissions to invoke the Lambda function

If you create the target group and register the Lambda function using the AWS Management Console, the console adds the required permissions to your Lambda function policy on your behalf. Otherwise, after you create the target group and register the function using the AWS CLI, you must use the [add-permission](#) command to grant Elastic Load Balancing permission to invoke your Lambda function. We recommend that you use the `aws:SourceAccount` and `aws:SourceArn` condition keys to restrict function invocation to the specified target group. For more information, see [The confused deputy problem](#) in the *IAM User Guide*,

```
aws lambda add-permission \  
--function-name lambda-function-arn-with-alias-name \  
--statement-id elb1 \  
--principal elasticloadbalancing.amazonaws.com \  
--action lambda:InvokeFunction \  
--source-arn target-group-arn \  

```

```
--source-account target-group-account-id
```

Lambda function versioning

You can register one Lambda function per target group. To ensure that you can change your Lambda function and that the load balancer always invokes the current version of the Lambda function, create a function alias and include the alias in the function ARN when you register the Lambda function with the load balancer. For more information, see [AWS Lambda function versioning and aliases](#) and [Traffic shifting using aliases](#) in the *AWS Lambda Developer Guide*.

Function Timeout

The load balancer waits until your Lambda function responds or times out. We recommend that you configure the timeout of the Lambda function based on your expected run time. For information about the default timeout value and how to change it, see [Basic AWS Lambda function configuration](#). For information about the maximum timeout value you can configure, see [AWS Lambda limits](#).

Create a target group for the Lambda function

Create a target group, which is used in request routing. If the request content matches a listener rule with an action to forward it to this target group, the load balancer invokes the registered Lambda function.

To create a target group and register the Lambda function using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose **Create target group**.
4. For **Choose a target type**, select **Lambda function**.
5. For **Target group name**, type a name for the target group.
6. (Optional) To enable health checks, choose **Enable** in the **Health checks** section.
7. (Optional) Add one or more tags as follows:
 - a. Expand the **Tags** section.
 - b. Choose **Add tag**.
 - c. Enter the tag key and the tag value.

8. Choose **Next**.
9. Specify a single Lambda function or omit this step and specify a Lambda function later.
10. Choose **Create target group**.

To create a target group and register the Lambda function using the AWS CLI

Use the [create-target-group](#) and [register-targets](#) commands.

Receive events from the load balancer

The load balancer supports Lambda invocation for requests over both HTTP and HTTPS. The load balancer sends an event in JSON format. The load balancer adds the following headers to every request: `X-Amzn-Trace-Id`, `X-Forwarded-For`, `X-Forwarded-Port`, and `X-Forwarded-Proto`.

If the `content-encoding` header is present, the load balancer Base64 encodes the body and sets `isBase64Encoded` to `true`.

If the `content-encoding` header is not present, Base64 encoding depends on the content type. For the following types, the load balancer sends the body as is and sets `isBase64Encoded` to `false`: `text/*`, `application/json`, `application/javascript`, and `application/xml`. Otherwise, the load balancer Base64 encodes the body and sets `isBase64Encoded` to `true`.

The following is an example event.

```
{
  "requestContext": {
    "elb": {
      "targetGroupArn":
"arn:aws:elasticloadbalancing:region:123456789012:targetgroup/my-target-
group/6d0ecf831eec9f09"
    }
  },
  "httpMethod": "GET",
  "path": "/",
  "queryStringParameters": {parameters},
  "headers": {
    "accept": "text/html,application/xhtml+xml",
    "accept-language": "en-US,en;q=0.8",
    "content-type": "text/plain",
    "cookie": "cookies",
```

```
    "host": "lambda-846800462-us-east-2.elb.amazonaws.com",
    "user-agent": "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11_6)",
    "x-amzn-trace-id": "Root=1-5bdb40ca-556d8b0c50dc66f0511bf520",
    "x-forwarded-for": "72.21.198.66",
    "x-forwarded-port": "443",
    "x-forwarded-proto": "https"
  },
  "isBase64Encoded": false,
  "body": "request_body"
}
```

Respond to the load balancer

The response from your Lambda function must include the Base64 encoding status, status code, and headers. You can omit the body.

To include a binary content in the body of the response, you must Base64 encode the content and set `isBase64Encoded` to `true`. The load balancer decodes the content to retrieve the binary content and sends it to the client in the body of the HTTP response.

The load balancer does not honor hop-by-hop headers, such as `Connection` or `Transfer-Encoding`. You can omit the `Content-Length` header because the load balancer computes it before sending responses to clients.

The following is an example response from a **nodejs** based Lambda function.

```
{
  "isBase64Encoded": false,
  "statusCode": 200,
  "statusDescription": "200 OK",
  "headers": {
    "Set-cookie": "cookies",
    "Content-Type": "application/json"
  },
  "body": "Hello from Lambda (optional)"
}
```

For Lambda function templates that work with Application Load Balancers, see [application-load-balancer-serverless-app](#) on github. Alternatively, open the [Lambda console](#), choose **Applications**, **Create a application**, and select one of the following from the AWS Serverless Application Repository:

- ALB-Lambda-Target-UploadFiletoS3
- ALB-Lambda-Target-BinaryResponse
- ALB-Lambda-Target-WhatisMyIP

Multi-value headers

If requests from a client or responses from a Lambda function contain headers with multiple values or contains the same header multiple times, or query parameters with multiple values for the same key, you can enable support for multi-value header syntax. After you enable multi-value headers, the headers and query parameters exchanged between the load balancer and the Lambda function use arrays instead of strings. If you do not enable multi-value header syntax and a header or query parameter has multiple values, the load balancer uses the last value that it receives.

Contents

- [Requests with multi-value headers](#)
- [Responses with multi-value headers](#)
- [Enable multi-value headers](#)

Requests with multi-value headers

The names of the fields used for headers and query string parameters differ depending on whether you enable multi-value headers for the target group.

The following example request has two query parameters with the same key:

```
http://www.example.com?&myKey=val1&myKey=val2
```

With the default format, the load balancer uses the last value sent by the client and sends you an event that includes query string parameters using `queryStringParameters`. For example:

```
"queryStringParameters": { "myKey": "val2"},
```

If you enable multi-value headers, the load balancer uses both key values sent by the client and sends you an event that includes query string parameters using `multiValueQueryStringParameters`. For example:

```
"multiValueQueryStringParameters": { "myKey": ["val1", "val2"] },
```

Similarly, suppose that the client sends a request with two cookies in the header:

```
"cookie": "name1=value1",  
"cookie": "name2=value2",
```

With the default format, the load balancer uses the last cookie sent by the client and sends you an event that includes headers using `headers`. For example:

```
"headers": {  
  "cookie": "name2=value2",  
  ...  
},
```

If you enable multi-value headers, the load balancer uses both cookies sent by the client and sends you an event that includes headers using `multiValueHeaders`. For example:

```
"multiValueHeaders": {  
  "cookie": ["name1=value1", "name2=value2"],  
  ...  
},
```

If the query parameters are URL-encoded, the load balancer does not decode them. You must decode them in your Lambda function.

Responses with multi-value headers

The names of the fields used for headers differ depending on whether you enable multi-value headers for the target group. You must use `multiValueHeaders` if you have enabled multi-value headers and `headers` otherwise.

With the default format, you can specify a single cookie:

```
{  
  "headers": {  
    "Set-cookie": "cookie-name=cookie-value;Domain=myweb.com;Secure;HttpOnly",  
    "Content-Type": "application/json"  
  },  
}
```

If you enable multi-value headers, you must specify multiple cookies as follows:

```
{
  "multiValueHeaders": {
    "Set-cookie": ["cookie-name=cookie-
value;Domain=myweb.com;Secure;HttpOnly", "cookie-name=cookie-value;Expires=May 8,
2019"],
    "Content-Type": ["application/json"]
  },
}
```

The load balancer might send the headers to the client in a different order than the order specified in the Lambda response payload. Therefore, do not count on headers being returned in a specific order.

Enable multi-value headers

You can enable or disable multi-value headers for a target group with the target type `lambda`.

To enable multi-value headers using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. On the **Group details** tab, in the **Attributes** section, choose **Edit**.
5. Select or clear **Multi value headers**.
6. Choose **Save changes**.

To enable multi-value headers using the AWS CLI

Use the [modify-target-group-attributes](#) command with the `lambda.multi_value_headers.enabled` attribute.

Enable health checks

By default, health checks are disabled for target groups of type `lambda`. You can enable health checks in order to implement DNS failover with Amazon Route 53. The Lambda function can check the health of a downstream service before responding to the health check request. If the response

from the Lambda function indicates a health check failure, the health check failure is passed to Route 53. You can configure Route 53 to fail over to a backup application stack.

You are charged for health checks as you are for any Lambda function invocation.

The following is the format of the health check event sent to your Lambda function. To check whether an event is a health check event, check the value of the user-agent field. The user agent for health checks is `ELB-HealthChecker/2.0`.

```
{
  "requestContext": {
    "elb": {
      "targetGroupArn":
"arn:aws:elasticloadbalancing:region:123456789012:targetgroup/my-target-
group/6d0ecf831eec9f09"
    }
  },
  "httpMethod": "GET",
  "path": "/",
  "queryStringParameters": {},
  "headers": {
    "user-agent": "ELB-HealthChecker/2.0"
  },
  "body": "",
  "isBase64Encoded": false
}
```

To enable health checks for a target group using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. On the **Group details** tab, in the **Health check settings** section, choose **Edit**.
5. For **Health checks**, select **Enable**.
6. Choose **Save changes**.

To enable health checks for a target group using the AWS CLI

Use the [modify-target-group](#) command with the `--health-check-enabled` option.

Deregister the Lambda function

If you no longer need to send traffic to your Lambda function, you can deregister it. After you deregister a Lambda function, in-flight requests fail with HTTP 5XX errors.

To replace a Lambda function, we recommend that you create a new target group, register the new function with the new target group, and update the listener rules to use the new target group instead of the existing one.

To deregister the Lambda function using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. On the **Targets** tab, choose **Deregister**.
5. When prompted for confirmation, choose **Deregister**.

To deregister the Lambda function using the AWS CLI

Use the [deregister-targets](#) command.

Tags for your target group

Tags help you to categorize your target groups in different ways, for example, by purpose, owner, or environment.

You can add multiple tags to each target group. Tag keys must be unique for each target group. If you add a tag with a key that is already associated with the target group, it updates the value of that tag.

When you are finished with a tag, you can remove it.

Restrictions

- Maximum number of tags per resource—50
- Maximum key length—127 Unicode characters
- Maximum value length—255 Unicode characters

- Tag keys and values are case-sensitive. Allowed characters are letters, spaces, and numbers representable in UTF-8, plus the following special characters: + - = . _ : / @. Do not use leading or trailing spaces.
- Do not use the `aws :` prefix in your tag names or values because it is reserved for AWS use. You can't edit or delete tag names or values with this prefix. Tags with this prefix do not count against your tags per resource limit.

To update the tags for a target group using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Choose the name of the target group to open its details page.
4. On the **Tags** tab, choose **Manage tags** and do one or more of the following:
 - a. To update a tag, enter new values for **Key** and **Value**.
 - b. To add a tag, choose **Add tag** and enter values for **Key** and **Value**.
 - c. To delete a tag, choose **Remove** next to the tag.
5. When you have finished updating tags, choose **Save changes**.

To update the tags for a target group using the AWS CLI

Use the [add-tags](#) and [remove-tags](#) commands.

Delete a target group

You can delete a target group if it is not referenced by the forward actions of any listener rules. Deleting a target group does not affect the targets registered with the target group. If you no longer need a registered EC2 instance, you can stop or terminate it.

To delete a target group using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the navigation pane, under **Load Balancing**, choose **Target Groups**.
3. Select the target group and choose **Actions, Delete**.
4. When prompted for confirmation, choose **Yes, delete**.

To delete a target group using the AWS CLI

Use the [delete-target-group](#) command.

Monitor your Application Load Balancers

You can use the following features to monitor your load balancers, analyze traffic patterns, and troubleshoot issues with your load balancers and targets.

CloudWatch metrics

You can use Amazon CloudWatch to retrieve statistics about data points for your load balancers and targets as an ordered set of time-series data, known as *metrics*. You can use these metrics to verify that your system is performing as expected. For more information, see [CloudWatch metrics for your Application Load Balancer](#).

Access logs

You can use access logs to capture detailed information about the requests made to your load balancer and store them as log files in Amazon S3. You can use these access logs to analyze traffic patterns and to troubleshoot issues with your targets. For more information, see [Access logs for your Application Load Balancer](#).

Connection logs

You can use connection logs to capture attributes about the requests sent to your load balancer, and store them as log files in Amazon S3. You can use these connection logs to determine the client IP address and port, client certificate information, connection results, and TLS ciphers being used. These connection logs can then be used to review request patterns, and other trends. For more information, see [Connection logs for your Application Load Balancer](#).

Request tracing

You can use request tracing to track HTTP requests. The load balancer adds a header with a trace identifier to each request it receives. For more information, see [Request tracing for your Application Load Balancer](#).

CloudTrail logs

You can use AWS CloudTrail to capture detailed information about the calls made to the Elastic Load Balancing API and store them as log files in Amazon S3. You can use these CloudTrail logs to determine which calls were made, the source IP address where the call came from, who made the call, when the call was made, and so on. For more information, see [Logging API calls for your Application Load Balancer using AWS CloudTrail](#).

CloudWatch metrics for your Application Load Balancer

Elastic Load Balancing publishes data points to Amazon CloudWatch for your load balancers and your targets. CloudWatch enables you to retrieve statistics about those data points as an ordered set of time-series data, known as *metrics*. Think of a metric as a variable to monitor, and the data points as the values of that variable over time. For example, you can monitor the total number of healthy targets for a load balancer over a specified time period. Each data point has an associated time stamp and an optional unit of measurement.

You can use metrics to verify that your system is performing as expected. For example, you can create a CloudWatch alarm to monitor a specified metric and initiate an action (such as sending a notification to an email address) if the metric goes outside what you consider an acceptable range.

Elastic Load Balancing reports metrics to CloudWatch only when requests are flowing through the load balancer. If there are requests flowing through the load balancer, Elastic Load Balancing measures and sends its metrics in 60-second intervals. If there are no requests flowing through the load balancer or no data for a metric, the metric is not reported.

For more information, see the [Amazon CloudWatch User Guide](#).

Contents

- [Application Load Balancer metrics](#)
- [Metric dimensions for Application Load Balancers](#)
- [Statistics for Application Load Balancer metrics](#)
- [View CloudWatch metrics for your load balancer](#)

Application Load Balancer metrics

- [Load balancers](#)
- [Targets](#)
- [Target group health](#)
- [Lambda functions](#)
- [User authentication](#)

The AWS/ApplicationELB namespace includes the following metrics for load balancers.

Metric	Description
ActiveConnectionCount	<p>The total number of concurrent TCP connections active from clients to the load balancer and from the load balancer to targets.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer • AvailabilityZone , LoadBalancer
AnomalousHostCount	<p>The number of hosts detected with anomalies.</p> <p>Reporting criteria: Always reported</p> <p>Statistics: The most useful statistics are Average, Minimum, and Maximum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • TargetGroup , LoadBalancer • TargetGroup , AvailabilityZone , LoadBalancer
ClientTLSNegotiationErrorCount	<p>The number of TLS connections initiated by the client that did not establish a session with the load balancer due to a TLS error. Possible causes include a mismatch of ciphers or protocols or the client failing to verify the server certificate and closing the connection.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer • AvailabilityZone , LoadBalancer

Metric	Description
ConsumedLCUs	<p>The number of load balancer capacity units (LCU) used by your load balancer. You pay for the number of LCUs that you use per hour. For more information, see Elastic Load Balancing pricing.</p> <p>Reporting criteria: Always reported</p> <p>Statistics: All</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer
DesyncMitigationMode_NonCompliant_Request_Count	<p>The number of requests that do not comply with RFC 7230.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer • AvailabilityZone , LoadBalancer
DroppedInvalidHeaderRequestCount	<p>The number of requests where the load balancer removed HTTP headers with header fields that are not valid before routing the request. The load balancer removes these headers only if the <code>routing.http.drop_invalid_header_fields.enabled</code> attribute is set to true.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: All</p> <p>Dimensions</p> <ul style="list-style-type: none"> • AvailabilityZone , LoadBalancer

Metric	Description
MitigatedHostCount	<p>The number of targets under mitigation.</p> <p>Reporting criteria: Always reported</p> <p>Statistics: The most useful statistics are Average, Minimum, and Maximum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • TargetGroup , LoadBalancer • TargetGroup , AvailabilityZone , LoadBalancer
ForwardedInvalidHeaderRequestCount	<p>The number of requests routed by the load balancer that had HTTP headers with header fields that are not valid. The load balancer forwards requests with these headers only if the <code>routing.http.drop_invalid_header_fields.enabled</code> attribute is set to false.</p> <p>Reporting criteria: Always reported</p> <p>Statistics: All</p> <p>Dimensions</p> <ul style="list-style-type: none"> • AvailabilityZone , LoadBalancer

Metric	Description
GrpcRequestCount	<p>The number of gRPC requests processed over IPv4 and IPv6.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is Sum. Minimum, Maximum, and Average all return 1.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer , TargetGroup • AvailabilityZone , LoadBalancer , TargetGroup • TargetGroup • AvailabilityZone , TargetGroup
HTTP_Fixed_Response_Count	<p>The number of fixed-response actions that were successful.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The only meaningful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer • AvailabilityZone , LoadBalancer
HTTP_Redirect_Count	<p>The number of redirect actions that were successful.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The only meaningful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer • AvailabilityZone , LoadBalancer

Metric	Description
HTTP_Redirect_Url_Limit_Exceeded_Count	<p>The number of redirect actions that couldn't be completed because the URL in the response location header is larger than 8K.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The only meaningful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none">• LoadBalancer• AvailabilityZone , LoadBalancer
HTTPCode_ELB_3XX_Count	<p>The number of HTTP 3XX redirection codes that originate from the load balancer. This count does not include response codes generated by targets.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The only meaningful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none">• LoadBalancer• AvailabilityZone , LoadBalancer

Metric	Description
HTTPCode_ELB_4XX_Count	<p>The number of HTTP 4XX client error codes that originate from the load balancer. This count does not include response codes generated by targets.</p> <p>Client errors are generated when requests are malformed or incomplete. These requests were not received by the target, other than in the case where the load balancer returns an HTTP 460 error code. This count does not include any response codes generated by the targets.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is Sum. Minimum, Maximum, and Average all return 1.</p> <p>Dimensions</p> <ul style="list-style-type: none">• LoadBalancer• AvailabilityZone , LoadBalancer
HTTPCode_ELB_5XX_Count	<p>The number of HTTP 5XX server error codes that originate from the load balancer. This count does not include any response codes generated by the targets.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is Sum. Minimum, Maximum, and Average all return 1.</p> <p>Dimensions</p> <ul style="list-style-type: none">• LoadBalancer• AvailabilityZone , LoadBalancer

Metric	Description
HTTPCode_ELB_500_Count	<p>The number of HTTP 500 error codes that originate from the load balancer.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The only meaningful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none">• LoadBalancer• AvailabilityZone , LoadBalancer
HTTPCode_ELB_502_Count	<p>The number of HTTP 502 error codes that originate from the load balancer.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The only meaningful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none">• LoadBalancer• AvailabilityZone , LoadBalancer
HTTPCode_ELB_503_Count	<p>The number of HTTP 503 error codes that originate from the load balancer.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The only meaningful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none">• LoadBalancer• AvailabilityZone , LoadBalancer

Metric	Description
HTTPCode_ELB_504_Count	<p>The number of HTTP 504 error codes that originate from the load balancer.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The only meaningful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none">• LoadBalancer• AvailabilityZone , LoadBalancer
IPv6ProcessedBytes	<p>The total number of bytes processed by the load balancer over IPv6. This count is included in ProcessedBytes .</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none">• LoadBalancer• AvailabilityZone , LoadBalancer
IPv6RequestCount	<p>The number of IPv6 requests received by the load balancer.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is Sum. Minimum, Maximum, and Average all return 1.</p> <p>Dimensions</p> <ul style="list-style-type: none">• LoadBalancer• AvailabilityZone , LoadBalancer

Metric	Description
NewConnectionCount	<p>The total number of new TCP connections established from clients to the load balancer and from the load balancer to targets.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none">• LoadBalancer• AvailabilityZone , LoadBalancer
NonStickyRequestCount	<p>The number of requests where the load balancer chose a new target because it couldn't use an existing sticky session. For example, the request was the first request from a new client and no stickiness cookie was presented, a stickiness cookie was presented but it did not specify a target that was registered with this target group, the stickiness cookie was malformed or expired, or an internal error prevented the load balancer from reading the stickiness cookie.</p> <p>Reporting criteria: Stickiness is enabled on the target group.</p> <p>Statistics: The only meaningful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none">• LoadBalancer• AvailabilityZone , LoadBalancer

Metric	Description
ProcessedBytes	<p>The total number of bytes processed by the load balancer over IPv4 and IPv6 (HTTP header and HTTP payload). This count includes traffic to and from clients and Lambda functions, and traffic from an Identity Provider (IdP) if user authentication is enabled.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none">• LoadBalancer• AvailabilityZone , LoadBalancer
RejectedConnection Count	<p>The number of connections that were rejected because the load balancer had reached its maximum number of connections.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none">• LoadBalancer• AvailabilityZone , LoadBalancer

Metric	Description
RequestCount	<p>The number of requests processed over IPv4 and IPv6. This metric is only incremented for requests where the load balancer node was able to choose a target. Requests that are rejected before a target is chosen are not reflected in this metric.</p> <p>Reporting criteria: Always reported</p> <p>Statistics: The most useful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer • LoadBalancer , AvailabilityZone • LoadBalancer , TargetGroup • LoadBalancer , AvailabilityZone , TargetGroup
RuleEvaluations	<p>The number of rules processed by the load balancer given a request rate averaged over an hour.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer

The AWS/ApplicationELB namespace includes the following metrics for targets.

Metric	Description
HealthyHostCount	<p>The number of targets that are considered healthy.</p> <p>Reporting criteria: Reported if there are registered targets.</p>

Metric	Description
	<p>Statistics: The most useful statistics are Average, Minimum, and Maximum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer , TargetGroup • LoadBalancer , AvailabilityZone , TargetGroup
<p>HTTPCode_Target_2XX_Count , HTTPCode_Target_3XX_Count , HTTPCode_Target_4XX_Count , HTTPCode_Target_5XX_Count</p>	<p>The number of HTTP response codes generated by the targets. This does not include any response codes generated by the load balancer.</p> <p>Reporting criteria: Reported if there are registered targets.</p> <p>Statistics: The most useful statistic is Sum. Minimum, Maximum, and Average all return 1.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer • AvailabilityZone , LoadBalancer • TargetGroup , LoadBalancer • TargetGroup , AvailabilityZone , LoadBalancer

Metric	Description
RequestCountPerTarget	<p>The average request count per target, in a target group. You must specify the target group using the <code>TargetGroup</code> dimension. This metric does not apply if the target is a Lambda function.</p> <p>This count uses the total number of requests received by the target group, divided by the number of healthy targets in the target group. If there are no healthy targets in the target group, it is divided by the total number of registered targets.</p> <p>Reporting criteria: Always reported</p> <p>Statistics: The only valid statistic is <code>Sum</code>. This represents the average not the sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • <code>TargetGroup</code> • <code>TargetGroup</code> , <code>AvailabilityZone</code> • <code>LoadBalancer</code> , <code>TargetGroup</code> • <code>LoadBalancer</code> , <code>AvailabilityZone</code> , <code>TargetGroup</code>
TargetConnectionErrorCount	<p>The number of connections that were not successfully established between the load balancer and target. This metric does not apply if the target is a Lambda function.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is <code>Sum</code>.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • <code>LoadBalancer</code> • <code>AvailabilityZone</code> , <code>LoadBalancer</code> • <code>TargetGroup</code> , <code>LoadBalancer</code> • <code>TargetGroup</code> , <code>AvailabilityZone</code> , <code>LoadBalancer</code>

Metric	Description
TargetResponseTime	<p>The time elapsed, in seconds, after the request leaves the load balancer until the target starts to send the response headers. This is equivalent to the <code>target_processing_time</code> field in the access logs.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistics are Average and pNN.NN (percentiles).</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer • AvailabilityZone , LoadBalancer • TargetGroup , LoadBalancer • TargetGroup , AvailabilityZone , LoadBalancer
TargetTLSTransactionErrorCount	<p>The number of TLS connections initiated by the load balancer that did not establish a session with the target. Possible causes include a mismatch of ciphers or protocols. This metric does not apply if the target is a Lambda function.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer • AvailabilityZone , LoadBalancer • TargetGroup , LoadBalancer • TargetGroup , AvailabilityZone , LoadBalancer

Metric	Description
UnHealthyHostCount	<p>The number of targets that are considered unhealthy.</p> <p>Reporting criteria: Reported if there are registered targets.</p> <p>Statistics: The most useful statistics are Average, Minimum, and Maximum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer , TargetGroup • LoadBalancer , AvailabilityZone , TargetGroup

The AWS/ApplicationELB namespace includes the following metrics for target group health. For more information, see [the section called “Target group health”](#).

Metric	Description
HealthyStateDNS	<p>The number of zones that meet the DNS healthy state requirements.</p> <p>Statistics: The most useful statistic is Min.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer , TargetGroup • AvailabilityZone , LoadBalancer , TargetGroup
HealthyStateRouting	<p>The number of zones that meet the routing healthy state requirements.</p> <p>Statistics: The most useful statistic is Min.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer , TargetGroup • AvailabilityZone , LoadBalancer , TargetGroup

Metric	Description
UnhealthyRoutingRequestCount	<p>The number of requests that are routed using the routing failover action (fail open).</p> <p>Statistics: The most useful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer , TargetGroup • AvailabilityZone , LoadBalancer , TargetGroup
UnhealthyStateDNS	<p>The number of zones that do not meet the DNS healthy state requirements and therefore were marked unhealthy in DNS.</p> <p>Statistics: The most useful statistic is Min.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer , TargetGroup • AvailabilityZone , LoadBalancer , TargetGroup
UnhealthyStateRouting	<p>The number of zones that do not meet the routing healthy state requirements, and therefore the load balancer distributes traffic to all targets in the zone, including the unhealthy targets.</p> <p>Statistics: The most useful statistic is Min.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer , TargetGroup • AvailabilityZone , LoadBalancer , TargetGroup

The AWS/ApplicationELB namespace includes the following metrics for Lambda functions that are registered as targets.

Metric	Description
LambdaInternalError	<p>The number of requests to a Lambda function that failed because of an issue internal to the load balancer or AWS Lambda. To get the error reason codes, check the error_reason field of the access log.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The only meaningful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • TargetGroup • TargetGroup , LoadBalancer
LambdaTargetProcessedBytes	<p>The total number of bytes processed by the load balancer for requests to and responses from a Lambda function.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The only meaningful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer
LambdaUserError	<p>The number of requests to a Lambda function that failed because of an issue with the Lambda function. For example, the load balancer did not have permission to invoke the function, the load balancer received JSON from the function that is malformed or missing required fields, or the size of the request body or response exceeded the maximum size of 1 MB. To get the error reason codes, check the error_reason field of the access log.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The only meaningful statistic is Sum.</p>

Metric	Description
	<p>Dimensions</p> <ul style="list-style-type: none"> TargetGroup TargetGroup , LoadBalancer

The AWS/ApplicationELB namespace includes the following metrics for user authentication.

Metric	Description
ELBAuthError	<p>The number of user authentications that could not be completed because an authenticate action was misconfigured, the load balancer couldn't establish a connection with the IdP, or the load balancer couldn't complete the authentication flow due to an internal error. To get the error reason codes, check the error_reason field of the access log.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The only meaningful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> LoadBalancer AvailabilityZone , LoadBalancer
ELBAuthFailure	<p>The number of user authentications that could not be completed because the IdP denied access to the user or an authorization code was used more than once. To get the error reason codes, check the error_reason field of the access log.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The only meaningful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> LoadBalancer

Metric	Description
ELBAuthLatency	<ul style="list-style-type: none"> • AvailabilityZone , LoadBalancer <p>The time elapsed, in milliseconds, to query the IdP for the ID token and user info. If one or more of these operations fail, this is the time to failure.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: All statistics are meaningful.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer • AvailabilityZone , LoadBalancer
ELBAuthRefreshTokenSuccess	<p>The number of times the load balancer successfully refreshed user claims using a refresh token provided by the IdP.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The only meaningful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer • AvailabilityZone , LoadBalancer
ELBAuthSuccess	<p>The number of authenticate actions that were successful. This metric is incremented at the end of the authentication workflow, after the load balancer has retrieved the user claims from the IdP.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> • LoadBalancer • AvailabilityZone , LoadBalancer

Metric	Description
ELBAuthUserClaimsSizeExceeded	<p>The number of times that a configured IdP returned user claims that exceeded 11K bytes in size.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The only meaningful statistic is Sum.</p> <p>Dimensions</p> <ul style="list-style-type: none"> LoadBalancer AvailabilityZone , LoadBalancer

Metric dimensions for Application Load Balancers

To filter the metrics for your Application Load Balancer, use the following dimensions.

Dimension	Description
AvailabilityZone	Filters the metric data by Availability Zone.
LoadBalancer	Filters the metric data by load balancer. Specify the load balancer as follows: <code>app/load-balancer-name/1234567890123456</code> (the final portion of the load balancer ARN).
TargetGroup	Filters the metric data by target group. Specify the target group as follows: <code>targetgroup/target-group-name/1234567890123456</code> (the final portion of the target group ARN).

Statistics for Application Load Balancer metrics

CloudWatch provides statistics based on the metric data points published by Elastic Load Balancing. Statistics are metric data aggregations over specified period of time. When you request statistics, the returned data stream is identified by the metric name and dimension. A dimension is

a name-value pair that uniquely identifies a metric. For example, you can request statistics for all the healthy EC2 instances behind a load balancer launched in a specific Availability Zone.

The `Minimum` and `Maximum` statistics reflect the minimum and maximum values of the data points reported by the individual load balancer nodes in each sampling window. For example, suppose there are 2 load balancer nodes that make up the Application Load Balancer. One node has `HealthyHostCount` with a `Minimum` of 2, a `Maximum` of 10, and an `Average` of 6, while the other node has `HealthyHostCount` with a `Minimum` of 1, a `Maximum` of 5, and an `Average` of 3. Therefore, the load balancer has a `Minimum` of 1, a `Maximum` of 10, and an `Average` of about 4.

We recommend you monitor for non-zero `UnHealthyHostCount` in the `Minimum` statistic, and alarm on non-zero value for more than one data point. Using the `Minimum` will detect when targets are considered unhealthy by every node and Availability Zone of your load balancer. Alarming on `Average` or `Maximum` is useful if you want to be alerted to potential problems, and we recommend customers review this metric and investigate non-zero occurrences. Mitigating failures automatically can be done following best practices of using load balancer health check in Amazon EC2 Auto Scaling, or Amazon Elastic Container Service (Amazon ECS).

The `Sum` statistic is the aggregate value across all load balancer nodes. Because metrics include multiple reports per period, `Sum` is only applicable to metrics that are aggregated across all load balancer nodes.

The `SampleCount` statistic is the number of samples measured. Because metrics are gathered based on sampling intervals and events, this statistic is typically not useful. For example, with `HealthyHostCount`, `SampleCount` is based on the number of samples that each load balancer node reports, not the number of healthy hosts.

A percentile indicates the relative standing of a value in a data set. You can specify any percentile, using up to two decimal places (for example, `p95.45`). For example, the 95th percentile means that 95 percent of the data is below this value and 5 percent is above. Percentiles are often used to isolate anomalies. For example, suppose that an application serves the majority of requests from a cache in 1-2 ms, but in 100-200 ms if the cache is empty. The maximum reflects the slowest case, around 200 ms. The average doesn't indicate the distribution of the data. Percentiles provide a more meaningful view of the application's performance. By using the 99th percentile as an Auto Scaling trigger or a CloudWatch alarm, you can target that no more than 1 percent of requests take longer than 2 ms to process.

View CloudWatch metrics for your load balancer

You can view the CloudWatch metrics for your load balancers using the Amazon EC2 console. These metrics are displayed as monitoring graphs. The monitoring graphs show data points if the load balancer is active and receiving requests.

Alternatively, you can view metrics for your load balancer using the CloudWatch console.

To view metrics using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. To view metrics filtered by target group, do the following:
 - a. In the navigation pane, choose **Target Groups**.
 - b. Select your target group, and then choose the **Monitoring** tab.
 - c. (Optional) To filter the results by time, select a time range from **Showing data for**.
 - d. To get a larger view of a single metric, select its graph.
3. To view metrics filtered by load balancer, do the following:
 - a. In the navigation pane, choose **Load Balancers**.
 - b. Select your load balancer, and then choose the **Monitoring** tab.
 - c. (Optional) To filter the results by time, select a time range from **Showing data for**.
 - d. To get a larger view of a single metric, select its graph.

To view metrics using the CloudWatch console

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Metrics**.
3. Select the **ApplicationELB** namespace.
4. (Optional) To view a metric across all dimensions, enter its name in the search field.
5. (Optional) To filter by dimension, select one of the following:
 - To display only the metrics reported for your load balancers, choose **Per AppELB Metrics**. To view the metrics for a single load balancer, enter its name in the search field.
 - To display only the metrics reported for your target groups, choose **Per AppELB, per TG Metrics**. To view the metrics for a single target group, enter its name in the search field.

- To display only the metrics reported for your load balancers by Availability Zone, choose **Per AppELB, per AZ Metrics**. To view the metrics for a single load balancer, enter its name in the search field. To view the metrics for a single Availability Zone, enter its name in the search field.
- To display only the metrics reported for your load balancers by Availability Zone and target group, choose **Per AppELB, per AZ, per TG Metrics**. To view the metrics for a single load balancer, enter its name in the search field. To view the metrics for a single target group, enter its name in the search field. To view the metrics for a single Availability Zone, enter its name in the search field.

To view metrics using the AWS CLI

Use the following [list-metrics](#) command to list the available metrics:

```
aws cloudwatch list-metrics --namespace AWS/ApplicationELB
```

To get the statistics for a metric using the AWS CLI

Use the following [get-metric-statistics](#) command get statistics for the specified metric and dimension. CloudWatch treats each unique combination of dimensions as a separate metric. You can't retrieve statistics using combinations of dimensions that were not specially published. You must specify the same dimensions that were used when the metrics were created.

```
aws cloudwatch get-metric-statistics --namespace AWS/ApplicationELB \  
--metric-name UnHealthyHostCount --statistics Average --period 3600 \  
--dimensions Name=LoadBalancer,Value=app/my-load-balancer/50dc6c495c0c9188 \  
Name=TargetGroup,Value=targetgroup/my-targets/73e2d6bc24d8a067 \  
--start-time 2016-04-18T00:00:00Z --end-time 2016-04-21T00:00:00Z
```

The following is example output:

```
{  
  "Datapoints": [  
    {  
      "Timestamp": "2016-04-18T22:00:00Z",  
      "Average": 0.0,  
      "Unit": "Count"  
    },  
  ],  
}
```

```
{
  "Timestamp": "2016-04-18T04:00:00Z",
  "Average": 0.0,
  "Unit": "Count"
},
...
],
"Label": "UnHealthyHostCount"
}
```

Access logs for your Application Load Balancer

Elastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and troubleshoot issues.

Access logs is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logs for your load balancer, Elastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify as compressed files. You can disable access logs at any time.

You are charged storage costs for Amazon S3, but not charged for the bandwidth used by Elastic Load Balancing to send log files to Amazon S3. For more information about storage costs, see [Amazon S3 pricing](#).

Contents

- [Access log files](#)
- [Access log entries](#)
- [Example log entries](#)
- [Processing access log files](#)
- [Enable access logs for your Application Load Balancer](#)
- [Disable access logs for your Application Load Balancer](#)

Access log files

Elastic Load Balancing publishes a log file for each load balancer node every 5 minutes. Log delivery is eventually consistent. The load balancer can deliver multiple logs for the same period. This usually happens if the site has high traffic.

The file names of the access logs use the following format:

```
bucket[/prefix]/AWSLogs/aws-account-id/elasticloadbalancing/region/yyyy/mm/dd/aws-account-id_elasticloadbalancing_region_app.load-balancer-id_end-time_ip-address_random-string.log.gz
```

bucket

The name of the S3 bucket.

prefix

(Optional) The prefix (logical hierarchy) for the bucket. The prefix that you specify must not include the string AWSLogs. For more information, see [Organizing objects using prefixes](#).

AWSLogs

We add the portion of the file name starting with AWSLogs after the bucket name and optional prefix that you specify.

aws-account-id

The AWS account ID of the owner.

region

The Region for your load balancer and S3 bucket.

yyyy/mm/dd

The date that the log was delivered.

load-balancer-id

The resource ID of the load balancer. If the resource ID contains any forward slashes (/), they are replaced with periods (.).

end-time

The date and time that the logging interval ended. For example, an end time of 20140215T2340Z contains entries for requests made between 23:35 and 23:40 in UTC or Zulu time.

ip-address

The IP address of the load balancer node that handled the request. For an internal load balancer, this is a private IP address.

random-string

A system-generated random string.

The following is an example log file name with a prefix:

```
s3://my-bucket/my-prefix/AWSLogs/123456789012/elasticloadbalancing/us-east-2/2022/05/01/123456789012_elasticloadbalancing_us-east-2_app.my-loadbalancer.1234567890abcdef_20220215T2340Z_172.160.001.192_20sg8hgm.log.gz
```

The following is an example log file name without a prefix:

```
s3://my-bucket/AWSLogs/123456789012/elasticloadbalancing/us-east-2/2022/05/01/123456789012_elasticloadbalancing_us-east-2_app.my-loadbalancer.1234567890abcdef_20220215T2340Z_172.160.001.192_20sg8hgm.log.gz
```

You can store your log files in your bucket for as long as you want, but you can also define Amazon S3 lifecycle rules to archive or delete log files automatically. For more information, see [Object lifecycle management](#) in the *Amazon Simple Storage Service User Guide*.

Access log entries

Elastic Load Balancing logs requests sent to the load balancer, including requests that never made it to the targets. For example, if a client sends a malformed request, or there are no healthy targets to respond to the request, the request is still logged. Elastic Load Balancing does not log health check requests.

Each log entry contains the details of a single request (or connection in the case of WebSockets) made to the load balancer. For WebSockets, an entry is written only after the connection is closed.

If the upgraded connection can't be established, the entry is the same as for an HTTP or HTTPS request.

Important

Elastic Load Balancing logs requests on a best-effort basis. We recommend that you use access logs to understand the nature of the requests, not as a complete accounting of all requests.

Contents

- [Syntax](#)
- [Actions taken](#)
- [Classification reasons](#)
- [Error reason codes](#)

Syntax

The following table describes the fields of an access log entry, in order. All fields are delimited by spaces. When new fields are introduced, they are added to the end of the log entry. You should ignore any fields at the end of the log entry that you were not expecting.

Field	Description
type	The type of request or connection. The possible values are as follows (ignore any other values): <ul style="list-style-type: none">• <code>http</code> — HTTP• <code>https</code> — HTTP over TLS• <code>h2</code> — HTTP/2 over TLS• <code>grpcs</code> — gRPC over TLS• <code>ws</code> — WebSockets• <code>wss</code> — WebSockets over TLS

Field	Description
time	The time when the load balancer generated a response to the client, in ISO 8601 format. For WebSockets, this is the time when the connection is closed.
elb	The resource ID of the load balancer. If you are parsing access log entries, note that resources IDs can contain forward slashes (/).
client:port	The IP address and port of the requesting client. If there is a proxy in front of the load balancer, this field contains the IP address of the proxy.
target:port	<p>The IP address and port of the target that processed this request.</p> <p>If the client didn't send a full request, the load balancer can't dispatch the request to a target, and this value is set to -.</p> <p>If the target is a Lambda function, this value is set to -.</p> <p>If the request is blocked by AWS WAF, this value is set to - and the value of <code>elb_status_code</code> is set to 403.</p>
request_processing_time	<p>The total time elapsed (in seconds, with millisecond precision) from the time the load balancer received the request until the time it sent the request to a target.</p> <p>This value is set to -1 if the load balancer can't dispatch the request to a target. This can happen if the target closes the connection before the idle timeout or if the client sends a malformed request.</p> <p>This value can also be set to -1 if the registered target does not respond before the idle timeout.</p> <p>If AWS WAF is enabled for your Application Load Balancer or the target type is a Lambda function, the time it takes for the client to send the required data for POST requests is counted towards <code>request_processing_time</code>.</p>

Field	Description
target_processing_time	<p>The total time elapsed (in seconds, with millisecond precision) from the time the load balancer sent the request to a target until the target started to send the response headers.</p> <p>This value is set to -1 if the load balancer can't dispatch the request to a target. This can happen if the target closes the connection before the idle timeout or if the client sends a malformed request.</p> <p>This value can also be set to -1 if the registered target does not respond before the idle timeout.</p> <p>If AWS WAF is not enabled for your Application Load Balancer, the time it takes for the client to send the required data for POST requests is counted towards target_processing_time .</p>
response_processing_time	<p>The total time elapsed (in seconds, with millisecond precision) from the time the load balancer received the response header from the target until it started to send the response to the client. This includes both the queuing time at the load balancer and the connection acquisition time from the load balancer to the client.</p> <p>This value is set to -1 if the load balancer doesn't receive a response from a target. This can happen if the target closes the connection before the idle timeout or if the client sends a malformed request.</p>
elb_status_code	<p>The status code of the response generated by the load balancer, fixed response rule, or AWS WAF custom response code for Block actions.</p>
target_status_code	<p>The status code of the response from the target. This value is recorded only if a connection was established to the target and the target sent a response. Otherwise, it is set to -.</p>
received_bytes	<p>The size of the request, in bytes, received from the client (requester). For HTTP requests, this includes the headers. For WebSockets, this is the total number of bytes received from the client on the connection.</p>

Field	Description
sent_bytes	The size of the response, in bytes, sent to the client (requester). For HTTP requests, this includes the headers. For WebSockets, this is the total number of bytes sent to the client on the connection.
"request"	The request line from the client, enclosed in double quotes and logged using the following format: HTTP method + protocol://host:port/uri + HTTP version. The load balancer preserves the URL sent by the client, as is, when recording the request URI. It does not set the content type for the access log file. When you process this field, consider how the client sent the URL.
"user_agent"	A User-Agent string that identifies the client that originated the request, enclosed in double quotes. The string consists of one or more product identifiers, product[/version]. If the string is longer than 8 KB, it is truncated.
ssl_cipher	[HTTPS listener] The SSL cipher. This value is set to - if the listener is not an HTTPS listener.
ssl_protocol	[HTTPS listener] The SSL protocol. This value is set to - if the listener is not an HTTPS listener.
target_group_arn	The Amazon Resource Name (ARN) of the target group.
"trace_id"	The contents of the X-Amzn-Trace-Id header, enclosed in double quotes.
"domain_name"	[HTTPS listener] The SNI domain provided by the client during the TLS handshake, enclosed in double quotes. This value is set to - if the client doesn't support SNI or the domain doesn't match a certificate and the default certificate is presented to the client.
"chosen_cert_arn"	[HTTPS listener] The ARN of the certificate presented to the client, enclosed in double quotes. This value is set to <code>session-reused</code> if the session is reused. This value is set to - if the listener is not an HTTPS listener.

Field	Description
matched_rule_priority	The priority value of the rule that matched the request. If a rule matched, this is a value from 1 to 50,000. If no rule matched and the default action was taken, this value is set to 0. If an error occurs during rules evaluation, it is set to -1. For any other error, it is set to -.
request_creation_time	The time when the load balancer received the request from the client, in ISO 8601 format.
"actions_executed"	The actions taken when processing the request, enclosed in double quotes. This value is a comma-separated list that can include the values described in Actions taken . If no action was taken, such as for a malformed request, this value is set to -.
"redirect_url"	The URL of the redirect target for the location header of the HTTP response, enclosed in double quotes. If no redirect actions were taken, this value is set to -.
"error_reason"	The error reason code, enclosed in double quotes. If the request failed, this is one of the error codes described in Error reason codes . If the actions taken do not include an authenticate action or the target is not a Lambda function, this value is set to -.
"target:port_list"	<p>A space-delimited list of IP addresses and ports for the targets that processed this request, enclosed in double quotes. Currently, this list can contain one item and it matches the target:port field.</p> <p>If the client didn't send a full request, the load balancer can't dispatch the request to a target, and this value is set to -.</p> <p>If the target is a Lambda function, this value is set to -.</p> <p>If the request is blocked by AWS WAF, this value is set to - and the value of elb_status_code is set to 403.</p>

Field	Description
"target_status_code_list"	<p>A space-delimited list of status codes from the responses of the targets, enclosed in double quotes. Currently, this list can contain one item and it matches the target_status_code field.</p> <p>This value is recorded only if a connection was established to the target and the target sent a response. Otherwise, it is set to -.</p>
"classification"	<p>The classification for desync mitigation, enclosed in double quotes. If the request does not comply with RFC 7230, the possible values are Acceptable, Ambiguous, and Severe.</p> <p>If the request complies with RFC 7230, this value is set to -.</p>
"classification_reason"	<p>The classification reason code, enclosed in double quotes. If the request does not comply with RFC 7230, this is one of the classification codes described in Classification reasons. If the request complies with RFC 7230, this value is set to -.</p>
conn_trace_id	<p>The connection traceability ID is a unique opaque ID used to identify each connection. After a connection is established with a client, subsequent requests from this client will contain this ID in their respective access log entries. This ID acts as a foreign key to create a link between the connection and access logs.</p>

Actions taken

The load balancer stores the actions that it takes in the actions_executed field of the access log.

- `authenticate` — The load balancer validated the session, authenticated the user, and added the user information to the request headers, as specified by the rule configuration.
- `fixed-response` — The load balancer issued a fixed response, as specified by the rule configuration.
- `forward` — The load balancer forwarded the request to a target, as specified by the rule configuration.

- `redirect` — The load balancer redirected the request to another URL, as specified by the rule configuration.
- `waf` — The load balancer forwarded the request to AWS WAF to determine whether the request should be forwarded to the target. If this is the final action, AWS WAF determined that the request should be rejected. By default, requests rejected by AWS WAF will be logged as "403" in the `elb_status_code` field. When AWS WAF is configured to reject requests with a Custom Response Code, the `elb_status_code` field will reflect the configured response code.
- `waf-failed` — The load balancer attempted to forward the request to AWS WAF, but this process failed.

Classification reasons

If a request does not comply with RFC 7230, the load balancer stores one of the following codes in the `classification_reason` field of the access log. For more information, see [Desync mitigation mode](#).

Code	Description	Classification
<code>AmbiguousUri</code>	The request URI contains control characters.	Ambiguous
<code>BadContentLength</code>	The Content-Length header contains a value that cannot be parsed or is not a valid number.	Severe
<code>BadHeader</code>	A header contains a null character or carriage return.	Severe
<code>BadTransferEncoding</code>	The Transfer-Encoding header contains a bad value.	Severe
<code>BadUri</code>	The request URI contains a null character or carriage return.	Severe
<code>BadMethod</code>	The request method is malformed.	Severe
<code>BadVersion</code>	The request version is malformed.	Severe
<code>BothTeClpPresent</code>	The request contains both a Transfer-Encoding header and a Content-Length header.	Ambiguous

Code	Description	Classification
DuplicateContentLength	There are multiple Content-Length headers with the same value.	Ambiguous
EmptyHeader	A header is empty or there is a line with only spaces.	Ambiguous
GetHeadZeroContentLength	There is a Content-Length header with a value of 0 for a GET or HEAD request.	Acceptable
MultipleContentLength	There are multiple Content-Length headers with different values.	Severe
MultipleTransferEncodingChunked	There are multiple Transfer-Encoding: chunked headers.	Severe
NonCompliantHeader	A header contains a non-ASCII or control character.	Acceptable
NonCompliantVersion	The request version contains a bad value.	Acceptable
SpaceInUri	The request URI contains a space that is not URL encoded.	Acceptable
SuspiciousHeader	There is a header that can be normalized to Transfer-Encoding or Content-Length using common text normalization techniques.	Ambiguous
UndefinedContentLengthSemantics	There is a Content-Length header defined for a GET or HEAD request.	Ambiguous

Code	Description	Classification
UndefinedTransferEncodingSemantics	There is a Transfer-Encoding header defined for a GET or HEAD request.	Ambiguous

Error reason codes

If the load balancer cannot complete an authenticate action, the load balancer stores one of the following reason codes in the `error_reason` field of the access log. The load balancer also increments the corresponding CloudWatch metric. For more information, see [Authenticate users using an Application Load Balancer](#).

Code	Description	Metric
AuthInvalidCookie	The authentication cookie is not valid.	ELBAuthFailure
AuthInvalidGrantError	The authorization grant code from the token endpoint is not valid.	ELBAuthFailure
AuthInvalidIdToken	The ID token is not valid.	ELBAuthFailure
AuthInvalidStateParam	The state parameter is not valid.	ELBAuthFailure
AuthInvalidTokenResponse	The response from the token endpoint is not valid.	ELBAuthFailure
AuthInvalidUserInfoResponse	The response from the user info endpoint is not valid.	ELBAuthFailure

Code	Description	Metric
AuthMissingCodeParam	The authentication response from the authorization endpoint is missing a query parameter named 'code'.	ELBAuthFailure
AuthMissingHostHeader	The authentication response from the authorization endpoint is missing a host header field.	ELBAuthError
AuthMissingStateParam	The authentication response from the authorization endpoint is missing a query parameter named 'state'.	ELBAuthFailure
AuthTokenEpRequestFailed	There is an error response (non-2XX) from the token endpoint.	ELBAuthError
AuthTokenEpRequestTimeout	The load balancer is unable to communicate with the token endpoint.	ELBAuthError
AuthUnhandledException	The load balancer encountered an unhandled exception.	ELBAuthError
AuthUserInfoEpRequestFailed	There is an error response (non-2XX) from the IdP user info endpoint.	ELBAuthError
AuthUserInfoEpRequestTimeout	The load balancer is unable to communicate with the IdP user info endpoint.	ELBAuthError
AuthUserInfoResponseSizeExceeded	The size of the claims returned by the IdP exceeded 11K bytes.	ELBAuthUserClaimsSizeExceeded

If a request to a weighted target group fails, the load balancer stores one of the following error codes in the `error_reason` field of the access log.

Code	Description
<code>AWSALBTGCookieInvalid</code>	The AWSALBTG cookie, which is used with weighted target groups, is not valid. For example, the load balancer returns this error when cookie values are URL encoded.
<code>WeightedTargetGroupsUnhandledException</code>	The load balancer encountered an unhandled exception.

If a request to a Lambda function fails, the load balancer stores one of the following reason codes in the `error_reason` field of the access log. The load balancer also increments the corresponding CloudWatch metric. For more information, see the Lambda [Invoke](#) action.

Code	Description	Metric
<code>LambdaAccessDenied</code>	The load balancer did not have permission to invoke the Lambda function.	<code>LambdaUserError</code>
<code>LambdaBadRequest</code>	Lambda invocation failed because the client request headers or body did not contain only UTF-8 characters.	<code>LambdaUserError</code>
<code>LambdaConnectionError</code>	The load balancer cannot connect to Lambda.	<code>LambdaInternalError</code>
<code>LambdaConnectionTimeout</code>	An attempt to connect to Lambda timed out.	<code>LambdaInternalError</code>
<code>LambdaEC2AccessDeniedException</code>	Amazon EC2 denied access to Lambda during function initialization.	<code>LambdaUserError</code>

Code	Description	Metric
LambdaEC2ThrottledException	Amazon EC2 throttled Lambda during function initialization.	LambdaUserError
LambdaEC2UnexpectedException	Amazon EC2 encountered an unexpected exception during function initialization.	LambdaUserError
LambdaENILimitReachedException	Lambda couldn't create a network interface in the VPC specified in the configuration of the Lambda function because the limit for network interfaces was exceeded.	LambdaUserError
LambdaInvalidResponse	The response from the Lambda function is malformed or is missing required fields.	LambdaUserError
LambdaInvalidRuntimeException	The specified version of the Lambda runtime is not supported.	LambdaUserError
LambdaInvalidSecurityGroupIDException	The security group ID specified in the configuration of the Lambda function is not valid.	LambdaUserError
LambdaInvalidSubnetIDException	The subnet ID specified in the configuration of the Lambda function is not valid.	LambdaUserError
LambdaInvalidZipFileException	Lambda could not unzip the specified function zip file.	LambdaUserError

Code	Description	Metric
LambdaKMS AccessDeniedException	Lambda could not decrypt environment variables because access to the KMS key was denied. Check the KMS permissions of the Lambda function.	LambdaUserError
LambdaKMS DisabledException	Lambda could not decrypt environment variables because the specified KMS key is disabled. Check the KMS key settings of the Lambda function.	LambdaUserError
LambdaKMS InvalidStateException	Lambda could not decrypt environment variables because the state of the KMS key is not valid. Check the KMS key settings of the Lambda function.	LambdaUserError
LambdaKMS NotFoundException	Lambda could not decrypt environment variables because the KMS key was not found. Check the KMS key settings of the Lambda function.	LambdaUserError
LambdaRequestTooLarge	The size of the request body exceeded 1 MB.	LambdaUserError
LambdaResourceNotFound	The Lambda function could not be found.	LambdaUserError
LambdaResponseTooLarge	The size of the response exceeded 1 MB.	LambdaUserError
LambdaServiceException	Lambda encountered an internal error.	LambdaInternalError

Code	Description	Metric
LambdaSubnetIPAddressLimitReachedException	Lambda could not set up VPC access for the Lambda function because one or more subnets have no available IP addresses.	LambdaUserError
LambdaThrottling	The Lambda function was throttled because there were too many requests.	LambdaUserError
LambdaUnhandled	The Lambda function encountered an unhandled exception.	LambdaUserError
LambdaUnhandledException	The load balancer encountered an unhandled exception.	LambdaInternalError
LambdaWebSocketNotSupported	WebSockets are not supported with Lambda.	LambdaUserError

If the load balancer encounters an error when forwarding requests to AWS WAF, it stores one of the following error codes in the `error_reason` field of the access log.

Code	Description
WAFConnectionError	The load balancer cannot connect to AWS WAF.
WAFConnectionTimeout	The connection to AWS WAF timed out.
WAFResponseReadTimeout	A request to AWS WAF timed out.
WAFServiceError	AWS WAF returned a 5XX error.
WAFUnhandledException	The load balancer encountered an unhandled exception.

Example log entries

The following are example log entries. Note that the text appears on multiple lines only to make them easier to read.

Example HTTP Entry

The following is an example log entry for an HTTP listener (port 80 to port 80):

```
http 2018-07-02T22:23:00.186641Z app/my-loadbalancer/50dc6c495c0c9188
192.168.131.39:2817 10.0.0.1:80 0.000 0.001 0.000 200 200 34 366
"GET http://www.example.com:80/ HTTP/1.1" "curl/7.46.0" - -
arn:aws:elasticloadbalancing:us-east-2:123456789012:targetgroup/my-
targets/73e2d6bc24d8a067
"Root=1-58337262-36d228ad5d99923122bbe354" "-" "-"
0 2018-07-02T22:22:48.364000Z "forward" "-" "-" "10.0.0.1:80" "200" "-" "-"
```

Example HTTPS Entry

The following is an example log entry for an HTTPS listener (port 443 to port 80):

```
https 2018-07-02T22:23:00.186641Z app/my-loadbalancer/50dc6c495c0c9188
192.168.131.39:2817 10.0.0.1:80 0.086 0.048 0.037 200 200 0 57
"GET https://www.example.com:443/ HTTP/1.1" "curl/7.46.0" ECDHE-RSA-AES128-GCM-SHA256
TLSv1.2
arn:aws:elasticloadbalancing:us-east-2:123456789012:targetgroup/my-
targets/73e2d6bc24d8a067
"Root=1-58337281-1d84f3d73c47ec4e58577259" "www.example.com" "arn:aws:acm:us-
east-2:123456789012:certificate/12345678-1234-1234-1234-123456789012"
1 2018-07-02T22:22:48.364000Z "authenticate,forward" "-" "-" "10.0.0.1:80" "200" "-"
"-" TID_123456
```

Example HTTP/2 Entry

The following is an example log entry for an HTTP/2 stream.

```
h2 2018-07-02T22:23:00.186641Z app/my-loadbalancer/50dc6c495c0c9188
10.0.1.252:48160 10.0.0.66:9000 0.000 0.002 0.000 200 200 5 257
"GET https://10.0.2.105:773/ HTTP/2.0" "curl/7.46.0" ECDHE-RSA-AES128-GCM-SHA256
TLSv1.2
arn:aws:elasticloadbalancing:us-east-2:123456789012:targetgroup/my-
targets/73e2d6bc24d8a067
```

```
"Root=1-58337327-72bd00b0343d75b906739c42" "-" "-"
1 2018-07-02T22:22:48.364000Z "redirect" "https://example.com:80/" "-" "10.0.0.66:9000"
"200" "-" "-"
```

Example WebSockets Entry

The following is an example log entry for a WebSockets connection.

```
ws 2018-07-02T22:23:00.186641Z app/my-loadbalancer/50dc6c495c0c9188
10.0.0.140:40914 10.0.1.192:8010 0.001 0.003 0.000 101 101 218 587
"GET http://10.0.0.30:80/ HTTP/1.1" "-" - -
arn:aws:elasticloadbalancing:us-east-2:123456789012:targetgroup/my-
targets/73e2d6bc24d8a067
"Root=1-58337364-23a8c76965a2ef7629b185e3" "-" "-"
1 2018-07-02T22:22:48.364000Z "forward" "-" "-" "10.0.1.192:8010" "101" "-" "-"
```

Example Secured WebSockets Entry

The following is an example log entry for a secured WebSockets connection.

```
wss 2018-07-02T22:23:00.186641Z app/my-loadbalancer/50dc6c495c0c9188
10.0.0.140:44244 10.0.0.171:8010 0.000 0.001 0.000 101 101 218 786
"GET https://10.0.0.30:443/ HTTP/1.1" "-" ECDHE-RSA-AES128-GCM-SHA256 TLSv1.2
arn:aws:elasticloadbalancing:us-west-2:123456789012:targetgroup/my-
targets/73e2d6bc24d8a067
"Root=1-58337364-23a8c76965a2ef7629b185e3" "-" "-"
1 2018-07-02T22:22:48.364000Z "forward" "-" "-" "10.0.0.171:8010" "101" "-" "-"
```

Example Entries for Lambda Functions

The following is an example log entry for a request to a Lambda function that succeeded:

```
http 2018-11-30T22:23:00.186641Z app/my-loadbalancer/50dc6c495c0c9188
192.168.131.39:2817 - 0.000 0.001 0.000 200 200 34 366
"GET http://www.example.com:80/ HTTP/1.1" "curl/7.46.0" - -
arn:aws:elasticloadbalancing:us-east-2:123456789012:targetgroup/my-
targets/73e2d6bc24d8a067
"Root=1-58337364-23a8c76965a2ef7629b185e3" "-" "-"
0 2018-11-30T22:22:48.364000Z "forward" "-" "-" "-" "-" "-" "-"
```

The following is an example log entry for a request to a Lambda function that failed:


```
http 2018-11-30T22:23:00.186641Z app/my-loadbalancer/50dc6c495c0c9188
192.168.131.39:2817 - 0.000 0.001 0.000 502 - 34 366
"GET http://www.example.com:80/ HTTP/1.1" "curl/7.46.0" - -
arn:aws:elasticloadbalancing:us-east-2:123456789012:targetgroup/my-
targets/73e2d6bc24d8a067
"Root=1-58337364-23a8c76965a2ef7629b185e3" "-" "-"
0 2018-11-30T22:22:48.364000Z "forward" "-" "LambdaInvalidResponse" "-" "-" "-" "
```

Processing access log files

The access log files are compressed. If you open the files using the Amazon S3 console, they are uncompressed and the information is displayed. If you download the files, you must uncompress them to view the information.

If there is a lot of demand on your website, your load balancer can generate log files with gigabytes of data. You might not be able to process such a large amount of data using line-by-line processing. Therefore, you might have to use analytical tools that provide parallel processing solutions. For example, you can use the following analytical tools to analyze and process access logs:

- Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. For more information, see [Querying Application Load Balancer logs](#) in the *Amazon Athena User Guide*.
- [Loggly](#)
- [Splunk](#)
- [Sumo logic](#)

Enable access logs for your Application Load Balancer

When you enable access logs for your load balancer, you must specify the name of the S3 bucket where the load balancer will store the logs. The bucket must have a bucket policy that grants Elastic Load Balancing permission to write to the bucket.

Tasks

- [Step 1: Create an S3 bucket](#)
- [Step 2: Attach a policy to your S3 bucket](#)
- [Step 3: Configure access logs](#)

- [Step 4: Verify bucket permissions](#)
- [Troubleshooting](#)

Step 1: Create an S3 bucket

When you enable access logs, you must specify an S3 bucket for the access logs. You can use an existing bucket, or create a bucket specifically for access logs. The bucket must meet the following requirements.

Requirements

- The bucket must be located in the same Region as the load balancer. The bucket and the load balancer can be owned by different accounts.
- The only server-side encryption option that's supported is Amazon S3-managed keys (SSE-S3). For more information, see [Amazon S3-managed encryption keys \(SSE-S3\)](#).

To create an S3 bucket using the Amazon S3 console

1. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Choose **Create bucket**.
3. On the **Create bucket** page, do the following:
 - a. For **Bucket name**, enter a name for your bucket. This name must be unique across all existing bucket names in Amazon S3. In some Regions, there might be additional restrictions on bucket names. For more information, see [Bucket restrictions and limitations](#) in the *Amazon Simple Storage Service User Guide*.
 - b. For **AWS Region**, select the Region where you created your load balancer.
 - c. For **Default encryption**, choose **Amazon S3-managed keys (SSE-S3)**.
 - d. Choose **Create bucket**.

Step 2: Attach a policy to your S3 bucket

Your S3 bucket must have a bucket policy that grants Elastic Load Balancing permission to write the access logs to the bucket. Bucket policies are a collection of JSON statements written in the access policy language to define access permissions for your bucket. Each statement includes information about a single permission and contains a series of elements.

If you're using an existing bucket that already has an attached policy, you can add the statement for Elastic Load Balancing access logs to the policy. If you do so, we recommend that you evaluate the resulting set of permissions to ensure that they are appropriate for the users that need access to the bucket for access logs.

Available bucket policies

The bucket policy that you'll use depends on the AWS Region and the type of zone.

Regions available as of August 2022 or later

This policy grants permissions to the specified log delivery service. Use this policy for load balancers in Availability Zones and Local Zones in the following Regions:

- Asia Pacific (Hyderabad)
- Asia Pacific (Melbourne)
- Canada West (Calgary)
- Europe (Spain)
- Europe (Zurich)
- Israel (Tel Aviv)
- Middle East (UAE)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "logdelivery.elasticloadbalancing.amazonaws.com"
      },
      "Action": "s3:PutObject",
      "Resource": "my-s3-arn"
    }
  ]
}
```

Replace *my-s3-arn* with the ARN of the location for your access logs. The ARN that you specify depends on whether you plan to specify a prefix when you enable access logs in [step 3](#).

- ARN format with a prefix

```
arn:aws:s3::bucket-name/prefix/AWSLogs/aws-account-id/*
```

Example ARN, where *bucket-name* is my-s3-bucket, *prefix* is my-prefix, and *aws-account-id* is 1111-2222-3333

```
arn:aws:s3::my-s3-bucket/my-prefix/AWSLogs/1111-2222-3333/*
```

- ARN format without a prefix

```
arn:aws:s3::bucket-name/AWSLogs/aws-account-id/*
```

Example ARN, where *bucket-name* is my-s3-bucket, and *aws-account-id* is 1111-2222-3333

```
arn:aws:s3::my-s3-bucket/AWSLogs/1111-2222-3333/*
```

Using NotPrincipal when Effect is Deny.

If the Amazon S3 bucket policy uses Effect with the value Deny and includes NotPrincipal as shown in the example below, ensure that logdelivery.elasticloadbalancing.amazonaws.com is included in the Service list.

```
{
  "Effect": "Deny",
  "NotPrincipal": {
    "Service": [
      "logdelivery.elasticloadbalancing.amazonaws.com",
      "example.com"
    ]
  }
},
```

Regions available before August 2022

This policy grants permissions to the specified Elastic Load Balancing account ID. Use this policy for load balancers in Availability Zones or Local Zones in the Regions in the list below.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::elb-account-id:root"
      },
      "Action": "s3:PutObject",
      "Resource": "my-s3-arn"
    }
  ]
}
```

Replace *elb-account-id* with the ID of the AWS account for Elastic Load Balancing for your Region:

- US East (N. Virginia) – 127311923021
- US East (Ohio) – 033677994240
- US West (N. California) – 027434742980
- US West (Oregon) – 797873946194
- Africa (Cape Town) – 098369216593
- Asia Pacific (Hong Kong) – 754344448648
- Asia Pacific (Jakarta) – 589379963580
- Asia Pacific (Mumbai) – 718504428378
- Asia Pacific (Osaka) – 383597477331
- Asia Pacific (Seoul) – 600734575887
- Asia Pacific (Singapore) – 114774131450
- Asia Pacific (Sydney) – 783225319266
- Asia Pacific (Tokyo) – 582318560864
- Canada (Central) – 985666609251
- Europe (Frankfurt) – 054676820928
- Europe (Ireland) – 156460612806
- Europe (London) – 652711504416
- Europe (Milan) – 635631232127

- Europe (Paris) – 009996457667
- Europe (Stockholm) – 897822967062
- Middle East (Bahrain) – 076674570225
- South America (São Paulo) – 507241528517

Replace *my-s3-arn* with the ARN of the location for your access logs. The ARN that you specify depends on whether you plan to specify a prefix when you enable access logs in [step 3](#).

- ARN format with a prefix

```
arn:aws:s3::bucket-name/prefix/AWSLogs/aws-account-id/*
```

Example ARN, where *bucket-name* is my-s3-bucket, *prefix* is my-prefix, and *aws-account-id* is 1111-2222-3333

```
arn:aws:s3::my-s3-bucket/my-prefix/AWSLogs/1111-2222-3333/*
```

- ARN format without a prefix

```
arn:aws:s3::bucket-name/AWSLogs/aws-account-id/*
```

Example ARN, where *bucket-name* is my-s3-bucket, and *aws-account-id* is 1111-2222-3333

```
arn:aws:s3::my-s3-bucket/AWSLogs/1111-2222-3333/*
```

Using NotPrincipal when Effect is Deny.

If the Amazon S3 bucket policy uses Effect with the value Deny and includes NotPrincipal as shown in the example below, ensure that logdelivery.elasticloadbalancing.amazonaws.com is included in the Service list.

```
{
  "Effect": "Deny",
  "NotPrincipal": {
    "Service": [
      "logdelivery.elasticloadbalancing.amazonaws.com",
```

```

    "example.com"
  ]
}
},

```

AWS GovCloud (US) Regions

This policy grants permissions to the specified Elastic Load Balancing account ID. Use this policy for load balancers in Availability Zones or Local Zones in the AWS GovCloud (US) Regions in the list below.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws-us-gov:iam::elb-account-id:root"
      },
      "Action": "s3:PutObject",
      "Resource": "my-s3-arn"
    }
  ]
}

```

Replace *elb-account-id* with the ID of the AWS account for Elastic Load Balancing for your AWS GovCloud (US) Region:

- AWS GovCloud (US-West) – 048591011584
- AWS GovCloud (US-East) – 190560391635

Replace *my-s3-arn* with the ARN of the location for your access logs. The ARN that you specify depends on whether you plan to specify a prefix when you enable access logs in [step 3](#).

- ARN format with a prefix

```
arn:aws-us-gov:s3::bucket-name/prefix/AWSLogs/aws-account-id/*
```

Example ARN, where *bucket-name* is *my-s3-bucket*, *prefix* is *my-prefix*, and *aws-account-id* is *1111-2222-3333*

```
arn:aws-us-gov:s3::my-s3-bucket/my-prefix/AWSLogs/1111-2222-3333/*
```

- ARN format without a prefix

```
arn:aws-us-gov:s3::bucket-name/AWSLogs/aws-account-id/*
```

Example ARN, where *bucket-name* is my-s3-bucket, and *aws-account-id* is 1111-2222-3333

```
arn:aws-us-gov:s3::my-s3-bucket/AWSLogs/1111-2222-3333/*
```

Using NotPrincipal when Effect is Deny.

If the Amazon S3 bucket policy uses Effect with the value Deny and includes NotPrincipal as shown in the example below, ensure that logdelivery.elasticloadbalancing.amazonaws.com is included in the Service list.

```
{
  "Effect": "Deny",
  "NotPrincipal": {
    "Service": [
      "logdelivery.elasticloadbalancing.amazonaws.com",
      "example.com"
    ]
  }
},
```

Outposts Zones

The following policy grants permissions to the specified log delivery service. Use this policy for load balancers in Outposts Zones.

```
{
  "Effect": "Allow",
  "Principal": {
    "Service": "logdelivery.elb.amazonaws.com"
  },
  "Action": "s3:PutObject",
  "Resource": "my-s3-arn,"
```



```

    "Condition": {
      "StringEquals": {
        "s3:x-amz-acl": "bucket-owner-full-control"
      }
    }
  }
}

```

Replace *my-s3-arn* with the ARN of the location for your access logs. The ARN that you specify depends on whether you plan to specify a prefix when you enable access logs in [step 3](#).

- ARN format with a prefix

```
arn:aws:s3::bucket-name/prefix/AWSLogs/aws-account-id/*
```

Example ARN, where *bucket-name* is my-s3-bucket, *prefix* is my-prefix, and *aws-account-id* is 1111-2222-3333

```
arn:aws:s3::my-s3-bucket/my-prefix/AWSLogs/1111-2222-3333/*
```

- ARN format without a prefix

```
arn:aws:s3::bucket-name/AWSLogs/aws-account-id/*
```

Example ARN, where *bucket-name* is my-s3-bucket, and *aws-account-id* is 1111-2222-3333

```
arn:aws:s3::my-s3-bucket/AWSLogs/1111-2222-3333/*
```

Using NotPrincipal when Effect is Deny.

If the Amazon S3 bucket policy uses Effect with the value Deny and includes NotPrincipal as shown in the example below, ensure that logdelivery.elasticloadbalancing.amazonaws.com is included in the Service list.

```

{
  "Effect": "Deny",
  "NotPrincipal": {
    "Service": [
      "logdelivery.elasticloadbalancing.amazonaws.com",

```

```
        "example.com"  
    ]  
}  
,
```

To attach a bucket policy for access logs to your bucket using the Amazon S3 console

1. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Select the name of the bucket to open its details page.
3. Choose **Permissions** and then choose **Bucket policy, Edit**.
4. Update the bucket policy to grant the required permissions.
5. Choose **Save changes**.

Step 3: Configure access logs

Use the following procedure to configure access logs to capture and deliver log files to your S3 bucket.

Requirements

The bucket must meet the requirements described in [step 1](#), and you must attach a bucket policy as described in [step 2](#). If you specify a prefix, it must not include the string "AWSLogs".

To enable access logs for your load balancer using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose **Load Balancers**.
3. Select the name of your load balancer to open its details page.
4. On the **Attributes** tab, choose **Edit**.
5. For **Monitoring**, turn on **Access logs**.
6. For **S3 URI**, enter the S3 URI for your log files. The URI that you specify depends on whether you're using a prefix.
 - URI with a prefix: `s3://bucket-name/prefix`
 - URI without a prefix: `s3://bucket-name`
7. Choose **Save changes**.

To enable access logs using the AWS CLI

Use the [modify-load-balancer-attributes](#) command.

To manage the S3 bucket for your access logs

Be sure to disable access logs before you delete the bucket that you configured for access logs. Otherwise, if there is a new bucket with the same name and the required bucket policy but created in an AWS account that you don't own, Elastic Load Balancing could write the access logs for your load balancer to this new bucket.

Step 4: Verify bucket permissions

After access logs are enabled for your load balancer, Elastic Load Balancing validates the S3 bucket and creates a test file to ensure that the bucket policy specifies the required permissions. You can use the Amazon S3 console to verify that the test file was created. The test file is not an actual access log file; it doesn't contain example records.

To verify a test file was created in your bucket using the Amazon S3 console

1. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Select the name of the bucket that you specified for access logs.
3. Navigate to the test file, `ELBAccessLogTestFile`. The location depends on whether you're using a prefix.
 - Location with a prefix: `my-bucket/prefix/AWSLogs/123456789012/ELBAccessLogTestFile`
 - Location without a prefix: `my-bucket/AWSLogs/123456789012/ELBAccessLogTestFile`

Troubleshooting

If you receive an access denied error, the following are possible causes:

- The bucket policy does not grant Elastic Load Balancing permission to write access logs to the bucket. Verify that you are using the correct bucket policy for the Region. Verify that the resource ARN uses the same bucket name that you specified when you enabled access logs. Verify that the resource ARN does not include a prefix if you did not specify a prefix when you enabled access logs.

- The bucket uses an unsupported server-side encryption option. The bucket must use Amazon S3-managed keys (SSE-S3).

Disable access logs for your Application Load Balancer

You can disable access logs for your load balancer at any time. After you disable access logs, your access logs remain in your S3 bucket until you delete them. For more information, see [Working with buckets](#) in the *Amazon Simple Storage Service User Guide*.

To disable access logs using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose **Load Balancers**.
3. Select the name of your load balancer to open its details page.
4. On the **Attributes** tab, choose **Edit**.
5. For **Monitoring**, turn off **Access logs**.
6. Choose **Save changes**.

To disable access logs using the AWS CLI

Use the [modify-load-balancer-attributes](#) command.

Connection logs for your Application Load Balancer

Elastic Load Balancing provides connection logs that capture detailed information about requests sent to your load balancer. Each log contains information such as the client's IP address and port, listener port, the TLS cipher and protocol used, TLS handshake latency, connection status, and client certificate details. You can use these connection logs to analyze request patterns and troubleshoot issues.

Connection logs is an optional feature of Elastic Load Balancing that is disabled by default. After you enable connection logs for your load balancer, Elastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify, as compressed files. You can disable connection logs at any time.

You are charged storage costs for Amazon S3, but not charged for the bandwidth used by Elastic Load Balancing to send log files to Amazon S3. For more information about storage costs, see [Amazon S3 pricing](#).

Contents

- [Connection log files](#)
- [Connection log entries](#)
- [Example log entries](#)
- [Processing connection log files](#)
- [Enable connection logs for your Application Load Balancer](#)
- [Disable connection logs for your Application Load Balancer](#)

Connection log files

Elastic Load Balancing publishes a log file for each load balancer node every 5 minutes. Log delivery is eventually consistent. The load balancer can deliver multiple logs for the same period. This usually happens if the site has high traffic.

The file names of the connection logs use the following format:

```
bucket[/prefix]/AWSLogs/aws-account-id/elasticloadbalancing/region/yyyy/mm/dd/  
conn_log.aws-account-id_elasticloadbalancing_region_app.load-balancer-id_end-time_ip-  
address_random-string.log.gz
```

bucket

The name of the S3 bucket.

prefix

(Optional) The prefix (logical hierarchy) for the bucket. The prefix that you specify must not include the string AWSLogs. For more information, see [Organizing objects using prefixes](#).

AWSLogs

We add the portion of the file name starting with AWSLogs after the bucket name and optional prefix that you specify.

aws-account-id

The AWS account ID of the owner.

region

The Region for your load balancer and S3 bucket.

yyyy/mm/dd

The date that the log was delivered.

load-balancer-id

The resource ID of the load balancer. If the resource ID contains any forward slashes (/), they are replaced with periods (.).

end-time

The date and time that the logging interval ended. For example, an end time of 20140215T2340Z contains entries for requests made between 23:35 and 23:40 in UTC or Zulu time.

ip-address

The IP address of the load balancer node that handled the request. For an internal load balancer, this is a private IP address.

random-string

A system-generated random string.

The following is an example log file name with a prefix:

```
s3://my-bucket/my-prefix/AWSLogs/123456789012/elasticloadbalancing/us-east-2/2022/05/01/conn_log.123456789012_elasticloadbalancing_us-east-2_app.my-loadbalancer.1234567890abcdef_20220215T2340Z_172.160.001.192_20sg8hgm.log.gz
```

The following is an example log file name without a prefix:

```
s3://my-bucket/AWSLogs/123456789012/elasticloadbalancing/us-east-2/2022/05/01/conn_log.123456789012_elasticloadbalancing_us-east-2_app.my-loadbalancer.1234567890abcdef_20220215T2340Z_172.160.001.192_20sg8hgm.log.gz
```

You can store your log files in your bucket for as long as you want, but you can also define Amazon S3 lifecycle rules to archive or delete log files automatically. For more information, see [Object lifecycle management](#) in the *Amazon Simple Storage Service User Guide*.

Connection log entries

Each connection attempt has an entry in a connection log file. How client requests are sent is determined by the connection being persistent, or nonpersistent. Nonpersistent connections have a single request, which creates a single entry in the access log and connection log. Persistent connections have multiple requests, which creates multiple entries in the access log and a single entry in the connection log.

Contents

- [Syntax](#)
- [Error reason codes](#)

Syntax

Connection log entries use the following format:

```
[timestamp] [client_ip] [client_port] [listener_port] [tls_protocol] [tls_cipher]
[tls_handshake_latency] [leaf_client_cert_subject] [leaf_client_cert_validity]
[leaf_client_cert_serial_number] [tls_verify_status]
```

The following table describes the fields of a connection log entry, in order. All fields are delimited by spaces. When new fields are introduced, they are added to the end of the log entry. You should ignore any fields at the end of the log entry that you were not expecting.

Field	Description
timestamp	The time, in ISO 8601 format, when the load balancer successfully established or failed to establish a connection.
client_ip	The IP address of the requesting client.
client_port	The port of the requesting client.
listener_port	The port of the load balancer listener receiving the client request.
tls_protocol	[HTTPS listener] The SSL/TLS protocol used during handshakes. This field is set to - for non SSL/TLS requests.

Field	Description
tls_cipher	[HTTPS listener] The SSL/TLS protocol used during handshakes. This field is set to - for non SSL/TLS requests.
tls_handshake_latency	[HTTPS listener] The total time in seconds, with a millisecond precision, elapsed while establishing a successful handshake. This field is set to - when: <ul style="list-style-type: none"> The incoming request is not a SSL/TLS request. The handshake is not established successfully.
leaf_client_cert_subject	[HTTPS listener] The subject name of the leaf client certificate. This field is set to - when: <ul style="list-style-type: none"> The incoming request is not a SSL/TLS request. The load balancer listener is not configured with mTLS enabled. The server is not able to load/parse the leaf client certificate.
leaf_client_cert_validity	[HTTPS listener] The validity, with <code>not-before</code> and <code>not-after</code> in ISO 8601 format, of the leaf client certificate. This field is set to - when: <ul style="list-style-type: none"> The incoming request is not a SSL/TLS request. The load balancer listener is not configured with mTLS enabled. The server is not able to load/parse the leaf client certificate.
leaf_client_cert_serial_number	[HTTPS listener] The serial number of the leaf client certificate. This field is set to - when: <ul style="list-style-type: none"> The incoming request is not a SSL/TLS request. The load balancer listener is not configured with mTLS enabled. The server is not able to load/parse the leaf client certificate.
tls_verify_status	[HTTPS listener] The status of the connection request. This value is <code>Success</code> if the connection is established successfully. On an unsuccessful connection the value is <code>Failed:\$error_code</code> .

Field	Description
conn_trace_id	The connection traceability ID is a unique opaque ID used to identify each connection. After a connection is established with a client, subsequent requests from this client will contain this ID in their respective access log entries. This ID acts as a foreign key to create a link between the connection and access logs.

Error reason codes

If the load balancer is unable to establish a connection, the load balancer stores one of the following reason codes in the connection log.

Code	Description
ClientCertificateMaxChainDepthExceeded	The maximum client certificate chain depth has been exceeded
ClientCertificateMaxSizeExceeded	The maximum client certificate size has been exceeded
ClientCertificateCrlHit	Client certificate has been revoked by the CA
ClientCertificateCrlProcessingError	CRL processing error
ClientCertificateUntrusted	Client certificate is untrusted
ClientCertificateNotYetValid	Client certificate is yet not valid

Code	Description
ClientCertificateExpired	Client certificate is expired
ClientCertificateTypeUnsupported	Client certificate type is unsupported
ClientCertificateInvalid	Client certificate is invalid
ClientCertificateRejected	Client certificate is rejected by custom server validation
UnmappedConnectionError	Unmapped runtime connection error

Example log entries

The following are example connection log entries.

The following is an example log entry for a successful connection with a HTTPS listener with mutual TLS verify mode enabled on port 443:

```
2023-10-04T17:05:15.514108Z 203.0.113.1 36280 443 TLSv1.2 ECDHE-RSA-AES128-GCM-SHA256 4.036 "CN=amazondomains.com,O=endEntity,L=Seattle,ST=Washington,C=US"
NotBefore=2023-09-21T22:43:21Z;NotAfter=2026-06-17T22:43:21Z FEF257372D5C14D4 Success
```

The following is an example log entry for a failed connection with a HTTPS listener with mutual TLS verify mode enabled on port 443.:

```
2023-10-04T17:05:15.514108Z 203.0.113.1 36280 443 TLSv1.2 ECDHE-RSA-AES128-GCM-SHA256 - "CN=amazondomains.com,O=endEntity,L=Seattle,ST=Washington,C=US"
NotBefore=2023-09-21T22:43:21Z;NotAfter=2026-06-17T22:43:21Z FEF257372D5C14D4
Failed:ClientCertUntrusted
```

Processing connection log files

The connection log files are compressed. If you open the files using the Amazon S3 console, they are uncompressed and the information is displayed. If you download the files, you must uncompress them to view the information.

If there is a lot of demand on your website, your load balancer can generate log files with gigabytes of data. You might not be able to process such a large amount of data using line-by-line processing. Therefore, you might have to use analytical tools that provide parallel processing solutions. For example, you can use the following analytical tools to analyze and process connection logs:

- Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL.
- [Loggly](#)
- [Splunk](#)
- [Sumo logic](#)

Enable connection logs for your Application Load Balancer

When you enable connection logs for your load balancer, you must specify the name of the S3 bucket where the load balancer will store the logs. The bucket must have a bucket policy that grants Elastic Load Balancing permission to write to the bucket.

Tasks

- [Step 1: Create an S3 bucket](#)
- [Step 2: Attach a policy to your S3 bucket](#)
- [Step 3: Configure connection logs](#)
- [Step 4: Verify bucket permissions](#)
- [Troubleshooting](#)

Step 1: Create an S3 bucket

When you enable connection logs, you must specify an S3 bucket for the connection logs. You can use an existing bucket, or create a bucket specifically for connection logs. The bucket must meet the following requirements.

Requirements

- The bucket must be located in the same Region as the load balancer. The bucket and the load balancer can be owned by different accounts.
- The only server-side encryption option that's supported is Amazon S3-managed keys (SSE-S3). For more information, see [Amazon S3-managed encryption keys \(SSE-S3\)](#).

To create an S3 bucket using the Amazon S3 console

1. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Choose **Create bucket**.
3. On the **Create bucket** page, do the following:
 - a. For **Bucket name**, enter a name for your bucket. This name must be unique across all existing bucket names in Amazon S3. In some Regions, there might be additional restrictions on bucket names. For more information, see [Bucket restrictions and limitations](#) in the *Amazon Simple Storage Service User Guide*.
 - b. For **AWS Region**, select the Region where you created your load balancer.
 - c. For **Default encryption**, choose **Amazon S3-managed keys (SSE-S3)**.
 - d. Choose **Create bucket**.

Step 2: Attach a policy to your S3 bucket

Your S3 bucket must have a bucket policy that grants Elastic Load Balancing permission to write the connection logs to the bucket. Bucket policies are a collection of JSON statements written in the access policy language to define access permissions for your bucket. Each statement includes information about a single permission and contains a series of elements.

If you're using an existing bucket that already has an attached policy, you can add the statement for Elastic Load Balancing connection logs to the policy. If you do so, we recommend that you evaluate the resulting set of permissions to ensure that they are appropriate for the users that need access to the bucket for connection logs.

Available bucket policies

The bucket policy that you'll use depends on the AWS Region and the type of zone.

Regions available as of August 2022 or later

This policy grants permissions to the specified log delivery service. Use this policy for load balancers in Availability Zones and Local Zones in the following Regions:

- Asia Pacific (Hyderabad)
- Asia Pacific (Melbourne)
- Europe (Spain)
- Europe (Zurich)
- Israel (Tel Aviv)
- Middle East (UAE)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "logdelivery.elasticloadbalancing.amazonaws.com"
      },
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::bucket-name/prefix/AWSLogs/aws-account-id/*"
    }
  ]
}
```

Regions available before August 2022

This policy grants permissions to the specified Elastic Load Balancing account ID. Use this policy for load balancers in Availability Zones or Local Zones in the Regions in the list below.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::elb-account-id:root"
      },
    }
  ]
}
```

```
    "Action": "s3:PutObject",
    "Resource": "arn:aws:s3:::bucket-name/prefix/AWSLogs/aws-account-id/*"
  }
]
```

Replace *elb-account-id* with the ID of the AWS account for Elastic Load Balancing for your Region:

- US East (N. Virginia) – 127311923021
- US East (Ohio) – 033677994240
- US West (N. California) – 027434742980
- US West (Oregon) – 797873946194
- Africa (Cape Town) – 098369216593
- Asia Pacific (Hong Kong) – 754344448648
- Asia Pacific (Jakarta) – 589379963580
- Asia Pacific (Mumbai) – 718504428378
- Asia Pacific (Osaka) – 383597477331
- Asia Pacific (Seoul) – 600734575887
- Asia Pacific (Singapore) – 114774131450
- Asia Pacific (Sydney) – 783225319266
- Asia Pacific (Tokyo) – 582318560864
- Canada (Central) – 985666609251
- Europe (Frankfurt) – 054676820928
- Europe (Ireland) – 156460612806
- Europe (London) – 652711504416
- Europe (Milan) – 635631232127
- Europe (Paris) – 009996457667
- Europe (Stockholm) – 897822967062
- Middle East (Bahrain) – 076674570225
- South America (São Paulo) – 507241528517
- AWS GovCloud (US-West) – 048591011584
- AWS GovCloud (US-East) – 190560391635

Replace *my-s3-arn* with the ARN of the location for your connection logs. The ARN that you specify depends on whether you plan to specify a prefix when you enable connection logs in [step 3](#).

- ARN example with a prefix

```
arn:aws:s3:::bucket-name/prefix/AWSLogs/aws-account-id/*
```

- ARN example without a prefix

```
arn:aws:s3:::bucket-name/AWSLogs/aws-account-id/*
```

Using `NotPrincipal` when `Effect` is `Deny`.

If the Amazon S3 bucket policy uses `Effect` with the value `Deny` and includes `NotPrincipal` as shown in the example below, ensure that `logdelivery.elasticloadbalancing.amazonaws.com` is included within the `Service` list.

```
{
  "Effect": "Deny",
  "NotPrincipal": {
    "Service": [
      "logdelivery.elasticloadbalancing.amazonaws.com",
      "example.com"
    ]
  },
}
```

To attach a bucket policy for connection logs to your bucket using the Amazon S3 console

1. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Select the name of the bucket to open its details page.
3. Choose **Permissions** and then choose **Bucket policy, Edit**.
4. Update the bucket policy to grant the required permissions.
5. Choose **Save changes**.

Step 3: Configure connection logs

Use the following procedure to configure connection logs to capture and deliver log files to your S3 bucket.

Requirements

The bucket must meet the requirements described in [step 1](#), and you must attach a bucket policy as described in [step 2](#). If you specify a prefix, it must not include the string "AWSLogs".

To enable connection logs for your load balancer using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose **Load Balancers**.
3. Select the name of your load balancer to open its details page.
4. On the **Attributes** tab, choose **Edit**.
5. For **Monitoring**, turn on **Connection logs**.
6. For **S3 URI**, enter the S3 URI for your log files. The URI that you specify depends on whether you're using a prefix.
 - URI with a prefix: `s3://bucket-name/prefix`
 - URI without a prefix: `s3://bucket-name`
7. Choose **Save changes**.

To enable connection logs using the AWS CLI

Use the [modify-load-balancer-attributes](#) command.

To manage the S3 bucket for your connection logs

Be sure to disable connection logs before you delete the bucket that you configured for connection logs. Otherwise, if there is a new bucket with the same name and the required bucket policy but created in an AWS account that you don't own, Elastic Load Balancing could write the connection logs for your load balancer to this new bucket.

Step 4: Verify bucket permissions

After connection logs are enabled for your load balancer, Elastic Load Balancing validates the S3 bucket and creates a test file to ensure that the bucket policy specifies the required permissions. You can use the Amazon S3 console to verify that the test file was created. The test file is not an actual connection log file; it doesn't contain example records.

To verify that Elastic Load Balancing created a test file in your S3 bucket

1. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Select the name of the bucket that you specified for connection logs.
3. Navigate to the test file, `ELBConnectionLogTestFile`. The location depends on whether you're using a prefix.
 - Location with a prefix: `my-bucket/prefix/AWSLogs/123456789012/ELBConnectionLogTestFile`
 - Location without a prefix: `my-bucket/AWSLogs/123456789012/ELBConnectionLogTestFile`

Troubleshooting

If you receive an access denied error, the following are possible causes:

- The bucket policy does not grant Elastic Load Balancing permission to write connection logs to the bucket. Verify that you are using the correct bucket policy for the Region. Verify that the resource ARN uses the same bucket name that you specified when you enabled connection logs. Verify that the resource ARN does not include a prefix if you did not specify a prefix when you enabled connection logs.
- The bucket uses an unsupported server-side encryption option. The bucket must use Amazon S3-managed keys (SSE-S3).

Disable connection logs for your Application Load Balancer

You can disable connection logs for your load balancer at any time. After you disable connection logs, your connection logs remain in your S3 bucket until you delete them. For more information, see [Working with buckets](#) in the *Amazon Simple Storage Service User Guide*.

To disable connection logs using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose **Load Balancers**.
3. Select the name of your load balancer to open its details page.
4. On the **Attributes** tab, choose **Edit**.

5. For **Monitoring**, turn off **Connection logs**.
6. Choose **Save changes**.

To disable connection logs using the AWS CLI

Use the [modify-load-balancer-attributes](#) command.

Request tracing for your Application Load Balancer

When the load balancer receives a request from a client, it adds or updates the **X-Amzn-Trace-Id** header before sending the request to the target. Any services or applications between the load balancer and the target can also add or update this header.

You can use request tracing to track HTTP requests from clients to targets or other services. If you enable access logs, the contents of the **X-Amzn-Trace-Id** header are logged. For more information, see [Access logs for your Application Load Balancer](#).

Syntax

The **X-Amzn-Trace-Id** header contains fields with the following format:

```
Field=version-time-id
```

Field

The name of the field. The supported values are `Root` and `Self`.

An application can add arbitrary fields for its own purposes. The load balancer preserves these fields but does not use them.

version

The version number.

time

The epoch time, in seconds.

id

The trace identifier.

Examples

If the **X-Amzn-Trace-Id** header is not present on an incoming request, the load balancer generates a header with a `Root` field and forwards the request. For example:

```
X-Amzn-Trace-Id: Root=1-67891233-abcdef012345678912345678
```

If the **X-Amzn-Trace-Id** header is present and has a `Root` field, the load balancer inserts a `Self` field and forwards the request. For example:

```
X-Amzn-Trace-Id: Self=1-67891233-12456789abcdef012345678;Root=1-67891233-abcdef012345678912345678
```

If an application adds a header with a `Root` field and a custom field, the load balancer preserves both fields, inserts a `Self` field, and forwards the request:

```
X-Amzn-Trace-Id: Self=1-67891233-12456789abcdef012345678;Root=1-67891233-abcdef012345678912345678;CalledFrom=app
```

If the **X-Amzn-Trace-Id** header is present and has a `Self` field, the load balancer updates the value of the `Self` field.

Limitations

- The load balancer updates the header when it receives an incoming request, not when it receives a response.
- If the HTTP headers are greater than 7 KB, the load balancer rewrites the **X-Amzn-Trace-Id** header with a `Root` field.
- With WebSockets, you can trace only until the upgrade request is successful.

Logging API calls for your Application Load Balancer using AWS CloudTrail

Elastic Load Balancing is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in Elastic Load Balancing. CloudTrail captures all API calls for Elastic Load Balancing as events. The calls captured include calls from the AWS Management Console and code calls to the Elastic Load Balancing API operations. If you create a trail, you can

enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Elastic Load Balancing. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine the request that was made to Elastic Load Balancing, the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, see the [AWS CloudTrail User Guide](#).

To monitor other actions for your load balancer, such as when a client makes a request to your load balancer, use access logs. For more information, see [Access logs for your Application Load Balancer](#).

Elastic Load Balancing information in CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When activity occurs in Elastic Load Balancing, that activity is recorded in a CloudTrail event along with other AWS service events in **Event history**. You can view, search, and download recent events in your AWS account. For more information, see [Viewing events with CloudTrail event history](#).

For an ongoing record of events in your AWS account, including events for Elastic Load Balancing, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all AWS Regions. The trail logs events from all Regions in the AWS partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- [Overview for creating a trail](#)
- [CloudTrail supported services and integrations](#)
- [Configuring Amazon SNS notifications for CloudTrail](#)
- [Receiving CloudTrail log files from multiple regions](#) and [Receiving CloudTrail log files from multiple accounts](#)

All Elastic Load Balancing actions for Application Load Balancers are logged by CloudTrail and are documented in the [Elastic Load Balancing API Reference version 2015-12-01](#). For example, calls to the `CreateLoadBalancer` and `DeleteLoadBalancer` actions generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see the [CloudTrail userIdentity element](#).

Understanding Elastic Load Balancing log file entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The log files include events for all AWS API calls for your AWS account, not just Elastic Load Balancing API calls. You can locate calls to the Elastic Load Balancing API by checking for eventSource elements with the value `elasticloadbalancing.amazonaws.com`. To view a record for a specific action, such as `CreateLoadBalancer`, check for eventName elements with the action name.

The following are example CloudTrail log records for Elastic Load Balancing for a user who created an Application Load Balancer and then deleted it using the AWS CLI. You can identify the CLI using the userAgent elements. You can identify the requested API calls using the eventName elements. Information about the user (Alice) can be found in the userIdentity element.

Example Example: CreateLoadBalancer

```
{
  "eventVersion": "1.03",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "123456789012",
    "arn": "arn:aws:iam::123456789012:user/Alice",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "userName": "Alice"
  },
  "eventTime": "2016-04-01T15:31:48Z",
  "eventSource": "elasticloadbalancing.amazonaws.com",
  "eventName": "CreateLoadBalancer",
```

```

"awsRegion": "us-west-2",
"sourceIPAddress": "198.51.100.1",
"userAgent": "aws-cli/1.10.10 Python/2.7.9 Windows/7 boto-core/1.4.1",
"requestParameters": {
  "subnets": ["subnet-8360a9e7", "subnet-b7d581c0"],
  "securityGroups": ["sg-5943793c"],
  "name": "my-load-balancer",
  "scheme": "internet-facing"
},
"responseElements": {
  "loadBalancers": [{
    "type": "application",
    "loadBalancerName": "my-load-balancer",
    "vpcId": "vpc-3ac0fb5f",
    "securityGroups": ["sg-5943793c"],
    "state": {"code": "provisioning"},
    "availabilityZones": [
      {"subnetId": "subnet-8360a9e7", "zoneName": "us-west-2a"},
      {"subnetId": "subnet-b7d581c0", "zoneName": "us-west-2b"}
    ],
    "dnsName": "my-load-balancer-1836718677.us-west-2.elb.amazonaws.com",
    "canonicalHostedZoneId": "Z2P70J7HTTTPLU",
    "createdTime": "Apr 11, 2016 5:23:50 PM",
    "loadBalancerArn": "arn:aws:elasticloadbalancing:us-
west-2:123456789012:loadbalancer/app/my-load-balancer/ffcddace1759e1d0",
    "scheme": "internet-facing"
  ]
},
"requestID": "b9960276-b9b2-11e3-8a13-f1ef1EXAMPLE",
"eventID": "6f4ab5bd-2daa-4d00-be14-d92efEXAMPLE",
"eventType": "AwsApiCall",
"apiVersion": "2015-12-01",
"recipientAccountId": "123456789012"
}

```

Example Example: DeleteLoadBalancer

```

{
  "eventVersion": "1.03",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "123456789012",
    "arn": "arn:aws:iam::123456789012:user/Alice",

```

```
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "userName": "Alice"
  },
  "eventTime": "2016-04-01T15:31:48Z",
  "eventSource": "elasticloadbalancing.amazonaws.com",
  "eventName": "DeleteLoadBalancer",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "198.51.100.1",
  "userAgent": "aws-cli/1.10.10 Python/2.7.9 Windows/7 boto-core/1.4.1",
  "requestParameters": {
    "loadBalancerArn": "arn:aws:elasticloadbalancing:us-
west-2:123456789012:loadbalancer/app/my-load-balancer/ffcddace1759e1d0"
  },
  "responseElements": null,
  "requestID": "349598b3-000e-11e6-a82b-298133eEXAMPLE",
  "eventID": "75e81c95-4012-421f-a0cf-babdaEXAMPLE",
  "eventType": "AwsApiCall",
  "apiVersion": "2015-12-01",
  "recipientAccountId": "123456789012"
}
```

Troubleshoot your Application Load Balancers

The following information can help you troubleshoot issues with your Application Load Balancer.

Issues

- [A registered target is not in service](#)
- [Clients cannot connect to an internet-facing load balancer](#)
- [Requests sent to a custom domain aren't received by the load balancer](#)
- [HTTPS requests sent to the load balancer return "NET::ERR_CERT_COMMON_NAME_INVALID"](#)
- [Load balancer shows elevated processing times](#)
- [The load balancer sends a response code of 000](#)
- [The load balancer generates an HTTP error](#)
- [A target generates an HTTP error](#)
- [An AWS Certificate Manager certificate is not available for use](#)
- [Multi-Line headers are not supported](#)
- [Troubleshoot unhealthy targets using the resource map](#)

A registered target is not in service

If a target is taking longer than expected to enter the InService state, it might be failing health checks. Your target is not in service until it passes one health check. For more information, see [Health checks for your target groups](#).

Verify that your instance is failing health checks and then check for the following issues:

A security group does not allow traffic

The security group associated with an instance must allow traffic from the load balancer using the health check port and health check protocol. You can add a rule to the instance security group to allow all traffic from the load balancer security group. Also, the security group for your load balancer must allow traffic to the instances.

A network access control list (ACL) does not allow traffic

The network ACL associated with the subnets for your instances must allow inbound traffic on the health check port and outbound traffic on the ephemeral ports (1024-65535). The network

ACL associated with the subnets for your load balancer nodes must allow inbound traffic on the ephemeral ports and outbound traffic on the health check and ephemeral ports.

The ping path does not exist

Create a target page for the health check and specify its path as the ping path.

The connection times out

First, verify that you can connect to the target directly from within the network using the private IP address of the target and the health check protocol. If you can't connect, check whether the instance is over-utilized, and add more targets to your target group if it is too busy to respond. If you can connect, it is possible that the target page is not responding before the health check timeout period. Choose a simpler target page for the health check or adjust the health check settings.

The target did not return a successful response code

By default, the success code is 200, but you can optionally specify additional success codes when you configure health checks. Confirm the success codes that the load balancer is expecting and that your application is configured to return these codes on success.

The target response code was malformed or there was an error connecting to the target

Verify that your application responds to the load balancer's health check requests. Some applications require additional configuration to respond to health checks, such as a virtual host configuration to respond to the HTTP host header sent by the load balancer. The host header value contains the private IP address of the target, followed by the health check port when not using a default port. If the target uses a default health check port, the host header value contains only the private IP address of the target. For example, if your target's private IP address is `10.0.0.10` and its health check port is `8080`, the HTTP Host header sent by the load balancer in health checks is `Host: 10.0.0.10:8080`. If your target's private IP address is `10.0.0.10` and its health check port is `80` then the HTTP Host header sent by the load balancer in health checks is `Host: 10.0.0.10`. A virtual host configuration to respond to that host, or a default configuration, may be required to successfully health check your application. Health check requests have the following attributes: the `User-Agent` is set to `ELB-HealthChecker/2.0`, the line terminator for message-header fields is the sequence CRLF, and the header terminates at the first empty line followed by a CRLF.

Clients cannot connect to an internet-facing load balancer

If the load balancer is not responding to requests, check for the following issues:

Your internet-facing load balancer is attached to a private subnet

You must specify public subnets for your load balancer. A public subnet has a route to the Internet Gateway for your virtual private cloud (VPC).

A security group or network ACL does not allow traffic

The security group for the load balancer and any network ACLs for the load balancer subnets must allow inbound traffic from the clients and outbound traffic to the clients on the listener ports.

Requests sent to a custom domain aren't received by the load balancer

If the load balancer is not receiving requests sent to a custom domain, check for the following issues:

The custom domain name does not resolve to the load balancer IP address

- Confirm what IP address the custom domain name resolves to using a command line interface.
 - **Linux, macOS, or Unix** – You can use the `dig` command within Terminal. Ex.`dig example.com`
 - **Windows** – You can use the `nslookup` command within Command Prompt. Ex.`nslookup example.com`
- Confirm what IP address the load balancers DNS name resolves to using a command line interface.
- Compare the results of the two outputs. The IP addresses must match.

If using Route 53 to host your custom domain, see [My domain is unavailable on the internet](#) in the *Amazon Route 53 Developer Guide*.

HTTPS requests sent to the load balancer return "NET::ERR_CERT_COMMON_NAME_INVALID"

If HTTPS requests are receiving `NET::ERR_CERT_COMMON_NAME_INVALID` from the load balancer, check the following possible causes:

- The domain name used in the HTTPS request does not match the alternate name specified in the listeners associated ACM certificate.
- The load balancers default DNS name is being used. The default DNS name cannot be used to make HTTPS requests as a public certificate cannot be requested for the `*.amazonaws.com` domain.

Load balancer shows elevated processing times

The load balancer counts processing times differently based on configuration.

- If AWS WAF is associated with your Application Load Balancer and a client sends an HTTP POST request, the time to send the data for POST requests is reflected in the `request_processing_time` field in the load balancer access logs. This behavior is expected for HTTP POST requests.
- If AWS WAF is not associated with your Application Load Balancer and a client sends an HTTP POST request, the time to send the data for POST requests is reflected in the `target_processing_time` field in the load balancer access logs. This behavior is expected for HTTP POST requests.

The load balancer sends a response code of 000

With HTTP/2 connections, if the compressed length of any of the headers exceeds 8 K bytes or if the number of requests served through one connection exceeds 10,000, the load balancer sends a GOAWAY frame and closes the connection with a TCP FIN.

The load balancer generates an HTTP error

The following HTTP errors are generated by the load balancer. The load balancer sends the HTTP code to the client, saves the request to the access log, and increments the `HTTPCode_ELB_4XX_Count` or `HTTPCode_ELB_5XX_Count` metric.

Errors

- [HTTP 400: Bad request](#)
- [HTTP 401: Unauthorized](#)
- [HTTP 403: Forbidden](#)
- [HTTP 405: Method not allowed](#)
- [HTTP 408: Request timeout](#)
- [HTTP 413: Payload too large](#)
- [HTTP 414: URI too long](#)
- [HTTP 460](#)
- [HTTP 463](#)
- [HTTP 464](#)
- [HTTP 500: Internal server error](#)
- [HTTP 501: Not implemented](#)
- [HTTP 502: Bad gateway](#)
- [HTTP 503: Service unavailable](#)
- [HTTP 504: Gateway timeout](#)
- [HTTP 505: Version not supported](#)
- [HTTP 507: Insufficient Storage](#)
- [HTTP 561: Unauthorized](#)

HTTP 400: Bad request

Possible causes:

- The client sent a malformed request that does not meet the HTTP specification.
- The request header exceeded 16 K per request line, 16 K per single header, or 64 K for the entire request header.
- The client closed the connection before sending the full request body.

HTTP 401: Unauthorized

You configured a listener rule to authenticate users, but one of the following is true:

- You configured `OnUnauthenticatedRequest` to deny unauthenticated users or the IdP denied access.
- The size of the claims returned by the IdP exceeded the maximum size supported by the load balancer.
- A client submitted an HTTP/1.0 request without a host header, and the load balancer was unable to generate a redirect URL.
- The requested scope doesn't return an ID token.
- You don't complete the login process before the client login timeout expires. For more information see, [Client login timeout](#).

HTTP 403: Forbidden

You configured an AWS WAF web access control list (web ACL) to monitor requests to your Application Load Balancer and it blocked a request.

HTTP 405: Method not allowed

The client used the TRACE method, which is not supported by Application Load Balancers.

HTTP 408: Request timeout

The client did not send data before the idle timeout period expired. Sending a TCP keep-alive does not prevent this timeout. Send at least 1 byte of data before each idle timeout period elapses. Increase the length of the idle timeout period as needed.

HTTP 413: Payload too large

Possible causes:

- The target is a Lambda function and the request body exceeds 1 MB.
- The request header exceeded 16 K per request line, 16 K per single header, or 64 K for the entire request header.

HTTP 414: URI too long

The request URL or query string parameters are too large.

HTTP 460

The load balancer received a request from a client, but the client closed the connection with the load balancer before the idle timeout period elapsed.

Check whether the client timeout period is greater than the idle timeout period for the load balancer. Ensure that your target provides a response to the client before the client timeout period elapses, or increase the client timeout period to match the load balancer idle timeout, if the client supports this.

HTTP 463

The load balancer received an **X-Forwarded-For** request header with too many IP addresses. The upper limit for IP addresses is 30.

HTTP 464

The load balancer received an incoming request protocol that is incompatible with the version config of the target group protocol.

Possible causes:

- The request protocol is an HTTP/1.1, while the target group protocol version is a gRPC or HTTP/2.
- The request protocol is a gRPC, while the target group protocol version is an HTTP/1.1.
- The request protocol is an HTTP/2 and the request is not POST, while target group protocol version is a gRPC.

HTTP 500: Internal server error

Possible causes:

- You configured an AWS WAF web access control list (web ACL) and there was an error executing the web ACL rules.
- The load balancer is unable to communicate with the IdP token endpoint or the IdP user info endpoint.
 - Verify that the IdP's DNS is publicly resolvable.

- Verify that the security groups for your load balancer and the network ACLs for your VPC allow outbound access to these endpoints.
- Verify that your VPC has internet access. If you have an internal-facing load balancer, use a NAT gateway to enable internet access.
- The user claim received from the IdP is greater than 11KB in size.

HTTP 501: Not implemented

The load balancer received a **Transfer-Encoding** header with an unsupported value. The supported values for **Transfer-Encoding** are `chunked` and `identity`. As an alternative, you can use the **Content-Encoding** header.

HTTP 502: Bad gateway

Possible causes:

- The load balancer received a TCP RST from the target when attempting to establish a connection.
- The load balancer received an unexpected response from the target, such as "ICMP Destination unreachable (Host unreachable)", when attempting to establish a connection. Check whether traffic is allowed from the load balancer subnets to the targets on the target port.
- The target closed the connection with a TCP RST or a TCP FIN while the load balancer had an outstanding request to the target. Check whether the keep-alive duration of the target is shorter than the idle timeout value of the load balancer.
- The target response is malformed or contains HTTP headers that are not valid.
- The target response header exceeded 32 K for the entire response header.
- The deregistration delay period elapsed for a request being handled by a target that was deregistered. Increase the delay period so that lengthy operations can complete.
- The target is a Lambda function and the response body exceeds 1 MB.
- The target is a Lambda function that did not respond before its configured timeout was reached.
- The target is a Lambda function that returned an error or the function was throttled by the Lambda service.
- The load balancer encountered an SSL handshake error when connecting to a target.

For more information see [How do I troubleshoot Application Load Balancer HTTP 502 errors](#) in the AWS Support Knowledge Center.

HTTP 503: Service unavailable

The target groups for the load balancer have no registered targets.

HTTP 504: Gateway timeout

Possible causes:

- The load balancer failed to establish a connection to the target before the connection timeout expired (10 seconds).
- The load balancer established a connection to the target but the target did not respond before the idle timeout period elapsed.
- The networks ACL or SecurityGroup Policies did not allow traffic from the targets to the load balancer nodes on the ephemeral ports (1024-65535).
- The target returns a content-length header that is larger than the entity body. The load balancer timed out waiting for the missing bytes.
- The target is a Lambda function and the Lambda service did not respond before the connection timeout expired.
- The load balancer encountered an SSL handshake timeout (10 seconds) when connecting to a target.

HTTP 505: Version not supported

The load balancer received an unexpected HTTP version request. For example, the load balancer established an HTTP/1 connection but received an HTTP/2 request.

HTTP 507: Insufficient Storage

The redirect URL is too long.

HTTP 561: Unauthorized

You configured a listener rule to authenticate users, but the IdP returned an error code when authenticating the user. Check your access logs for the related [error reason code](#).

A target generates an HTTP error

The load balancer forwards valid HTTP responses from targets to the client, including HTTP errors. The HTTP errors generated by a target are recorded in the `HTTPCode_Target_4XX_Count` and `HTTPCode_Target_5XX_Count` metrics.

An AWS Certificate Manager certificate is not available for use

When deciding to use an HTTPS listener with your Application Load Balancer, AWS Certificate Manager requires you to validate domain ownership before issuing a certificate. If this step is missed during setup, the certificate remains in the `Pending Validation` state, and not available for use until validated.

- If using email validation, see [Email validation](#) in the *AWS Certificate Manager User Guide*.
- If using DNS validation, see [DNS validation](#) in the *AWS Certificate Manager User Guide*.

Multi-Line headers are not supported

Application Load Balancers do not support multi-line headers, including the `message/http media type` header. When a multi-line header is provided the Application Load Balancer appends a colon character, `:`, before passing it to the target.

Troubleshoot unhealthy targets using the resource map

If your Application Load Balancer targets are failing health checks, you can use the resource map to find unhealthy targets and take actions based on the failure reason code. For more information, see [Application Load Balancer resource map](#).

Resource map provides two views: **Overview**, and **Unhealthy Target Map**. **Overview** is selected by default and displays all of your load balancer's resources. Selecting the **Unhealthy Target Map** view will display only the unhealthy targets in each target group associated to the Application Load Balancer.

Note

You must enable **Show resource details** to view the health check summary and error messages for all applicable resources within the resource map. When not enabled, you must select each resource to view its details.

The **Target groups** column displays a summary of the healthy and unhealthy targets for each target group. This can help determine if all the targets are failing health checks, or only specific targets are failing. If all targets in a target group are failing health checks, check the configuration of the target group. Select a target groups name to open its detail page in a new tab.

The **Targets** column displays the TargetID and the current health check status for each target. When a target is unhealthy, the health check failure reason code is displayed. When a single target is failing a health check, verify the target has sufficient resources and confirm that applications running on the target are available. Select a targets ID to open its detail page in a new tab.

Selecting **Export** gives you the option of exporting the current view of your Application Load Balancer's resource map as a PDF.

Verify that your instance is failing health checks and then based on the failure reason code check for the following issues:

- **Unhealthy: HTTP Response Mismatch**

- Verify the application running on the target is sending the correct HTTP response to the Application Load Balancer's health check requests.
- Alternatively, you can update the Application Load Balancer's health check request to match the response from the application running on the target.

- **Unhealthy: Request timed out**

- Verify the security groups and network access control lists (ACL) associated with your targets and Application Load Balancer are not blocking connectivity.
- Verify the target has sufficient resources available to accept connections from the Application Load Balancer.
- Verify the status of any applications running on the target.
- The Application Load Balancer's health check responses can be viewed in each target's application logs. For more information, see [Health check reason codes](#).

- **Unhealthy: FailedHealthChecks**

- Verify the status of any applications running on the target.
- Verify the target is listening for traffic on the health check port.

When using an HTTPS listener

You choose which security policy is used for front-end connections. The security policy used for back-end connections is automatically selected based on the front-end security policy in use.

- If your HTTPS listener is using a TLS 1.3 security policy for front-end connections, the `ELBSecurityPolicy-TLS13-1-0-2021-06` security policy is used for back-end connections.
- If your HTTPS listener is not using a TLS 1.3 security policy for front-end connections, the `ELBSecurityPolicy-2016-08` security policy is used for back-end connections.

For more information, see [Security policies](#).

- Verify the target is providing a server certificate and key in the correct format specified by the security policy.
- Verify the target supports one or more matching ciphers, and a protocol provided by the Application Load Balancer to establish TLS handshakes.

Quotas for your Application Load Balancers

Your AWS account has default quotas, formerly referred to as limits, for each AWS service. Unless otherwise noted, each quota is Region-specific. You can request increases for some quotas, and other quotas cannot be increased.

To view the quotas for your Application Load Balancers, open the [Service Quotas console](#). In the navigation pane, choose **AWS services** and select **Elastic Load Balancing**. You can also use the [describe-account-limits](#) (AWS CLI) command for Elastic Load Balancing.

To request a quota increase, see [Requesting a quota increase](#) in the *Service Quotas User Guide*. If the quota is not yet available in Service Quotas, use the [Elastic Load Balancing limit increase form](#).

Load balancers

Your AWS account has the following quotas related to Application Load Balancers.

Name	Default	Adjustable
Application Load Balancers per Region	50	Yes
Certificates per Application Load Balancer (excluding default certificates)	25	Yes
Listeners per Application Load Balancer	50	Yes
Target Groups per Action per Application Load Balancer	5	No
Target Groups per Application Load Balancer	100	No
Targets per Application Load Balancer	1,000	Yes

Target groups

The following quotas are for target groups.

Name	Default	Adjustable
Target Groups per Region	3,000 *	Yes
Targets per Target Group per Region (instances or IP addresses)	1,000	Yes
Targets per Target Group per Region (Lambda functions)	1	No
Load balancers per target group	1	No

* This quota is shared by Application Load Balancers and Network Load Balancers.

Rules

The following quotas are for rules.

Name	Default	Adjustable
Rules per Application Load Balancer (excluding default rules)	100	Yes
Condition Values per Rule	5	No
Condition Wildcards per Rule	5	No
Match evaluations per rule	5	No

Trust stores

The following quotas are for trust stores.

Name	Default	Adjustable
Trust stores per account	20	Yes

Name	Default	Adjustable
Number of listeners using mTLS in verify mode, per load balancer.	2	No

Certificate authority certificates

The following quotas are for CA certificates.

Name	Default	Adjustable
CA certificates per trust store	25	Yes
CA certificate size	16KB	No
Maximum certificate chain depth	4	No

Certificate revocation lists

The following quotas are for certificate revocation lists.

Name	Default	Adjustable
Revocation lists per trust store	30	Yes
Revocation entries per trust store	500,000	Yes
Revocation list file size	50MB	No

HTTP headers

The following are the size limits for HTTP headers.

Name	Default	Adjustable
Request line	16 K	No

Name	Default	Adjustable
Single header	16 K	No
Entire response header	32 K	No
Entire request header	64 K	No

Document history for Application Load Balancers

The following table describes the releases for Application Load Balancers.

Change	Description	Date
Resource map	This release adds support to view your load balancer resources and relationships in a visual format.	March 8, 2024
One click WAF	This release adds support for configuring the behavior of your load balancer if it integrates with one click AWS WAF.	February 6, 2024
Mutual TLS	This release adds support for mutual TLS authentication.	November 26, 2023
Automatic Target Weights	This release adds support for the automatic target weights algorithm.	November 26, 2023
FIPS 140-3 TLS termination	This release adds security policies that use FIPS 140-3 cryptographic modules when terminating TLS connections.	November 20, 2023
Register targets using IPv6	This release adds support to register instances as targets when addressed by IPv6.	October 2, 2023
Security policies supporting TLS 1.3	This release adds support for TLS 1.3 predefined security policies.	March 22, 2023

Zonal shift	This release adds support to route traffic away from a single impaired Availability Zone through integration with the Amazon Route 53 Application Recovery Controller.	November 28, 2022
Turn off cross-zone load balancing	This release adds support to turn off cross-zone load balancing.	November 28, 2022
Target group health	This release adds support to configure the minimum count or percentage of targets that must be healthy, and what actions the load balancer takes when the threshold is not met.	November 28, 2022
Cross-zone load balancing	This release adds support to configure cross-zone load balancing at the target group level.	November 17, 2022
IPv6 target groups	This release adds support to configure IPv6 target groups for Application Load Balancers.	November 23, 2021
IPv6 internal load balancers	This release adds support to configure IPv6 target groups for Application Load Balancers.	November 23, 2021

AWS PrivateLink and static IP addresses	This release adds support to use AWS PrivateLink and expose static IP addresses by forwarding traffic directly from Network Load Balancers to Application Load Balancers.	September 27, 2021
Client port preservation	This release adds an attribute to preserve the source port that the client used to connect to the load balancer.	July 29, 2021
TLS headers	This release adds an attribute to indicate that the TLS headers, which contain information about the negotiated TLS version and cipher suite, are added to the client request before sending it to the target.	July 21, 2021
Additional ACM certificates	This release supports RSA certificates with 2048, 3072, and 4096-bit key lengths, and all ECDSA certificates.	July 14, 2021
Application-based stickiness	This release adds an application-based cookie to support sticky sessions for your load balancer.	February 8, 2021
Security policy for FS supporting TLS version 1.2	This release adds a security policy for Forward Secrecy (FS) supporting TLS version 1.2.	November 24, 2020

WAF fail open support	This release adds support for configuring the behavior of your load balancer if it integrates with AWS WAF.	November 13, 2020
gRPC and HTTP/2 support	This release adds support for gRPC workloads and end-to-end HTTP/2.	October 29, 2020
Outpost support	You can provision an Application Load Balancer on your AWS Outposts.	September 8, 2020
Desync mitigation mode	This release adds support for desync mitigation mode.	August 17, 2020
Least outstanding requests	This release adds support for the least outstanding requests algorithm.	November 25, 2019
Weighted target groups	This release adds support for forward actions with multiple target groups. Requests are distributed to these target groups based on the weight you specify for each target group.	November 19, 2019
New attribute	This release adds support for the routing.http.drop_invalid_header_fields.enabled attribute.	November 15, 2019
Security policies for FS	This release adds support for three additional predefined forward secrecy security policies.	October 8, 2019

Advanced request routing	This release adds support for additional condition types for your listener rules.	March 27, 2019
Lambda functions as a target	This release adds support for registering Lambda functions as a target.	November 29, 2018
Redirect actions	This release adds support for the load balancer to redirect requests to a different URL.	July 25, 2018
Fixed-response actions	This release adds support for the load balancer to return a custom HTTP response.	July 25, 2018
Security policies for FS and TLS 1.2	This release adds support for two additional predefined security policies.	June 6, 2018
User authentication	This release adds support for the load balancer to authenticate users of your applications using their corporate or social identities before routing requests.	May 30, 2018
Resource-level permissions	This release adds support for resource-level permissions and tagging condition keys.	May 10, 2018
Slow start mode	This release adds support for slow start mode, which gradually increases the share of requests the load balancer sends to a newly registered target while it warms up.	March 24, 2018

SNI support	This release adds support for Server Name Indication (SNI).	October 10, 2017
IP addresses as targets	This release adds support for registering IP addresses as targets.	August 31, 2017
Host-based routing	This release adds support for routing requests based on the host names in the host header.	April 5, 2017
Security policies for TLS 1.1 and TLS 1.2	This release adds security policies for TLS 1.1 and TLS 1.2.	February 6, 2017
IPv6 support	This release adds support for IPv6 addresses.	January 25, 2017
Request tracing	This release adds support for request tracing.	November 22, 2016
Percentiles support for the TargetResponseTime metric	This release adds support for the new percentile statistics supported by Amazon CloudWatch.	November 17, 2016
New load balancer type	This release of Elastic Load Balancing introduces Application Load Balancers.	August 11, 2016