
AWS Global Accelerator

API Reference

API Version 2018-08-08



AWS Global Accelerator: API Reference

Copyright © 2019 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

Welcome	1
Actions	3
CreateAccelerator	4
Request Syntax	4
Request Parameters	4
Response Syntax	5
Response Elements	5
Errors	5
Example	6
See Also	6
CreateEndpointGroup	7
Request Syntax	7
Request Parameters	7
Response Syntax	9
Response Elements	9
Errors	9
Example	10
See Also	11
CreateListener	12
Request Syntax	12
Request Parameters	12
Response Syntax	13
Response Elements	13
Errors	14
Example	14
See Also	15
DeleteAccelerator	16
Request Syntax	16
Request Parameters	16
Response Elements	16
Errors	16
See Also	17
DeleteEndpointGroup	18
Request Syntax	18
Request Parameters	18
Response Elements	18
Errors	18
See Also	18
DeleteListener	20
Request Syntax	20
Request Parameters	20
Response Elements	20
Errors	20
See Also	20
DescribeAccelerator	22
Request Syntax	22
Request Parameters	22
Response Syntax	22
Response Elements	22
Errors	23
Example	23
See Also	24
DescribeAcceleratorAttributes	25
Request Syntax	25

Request Parameters	25
Response Syntax	25
Response Elements	25
Errors	25
See Also	26
DescribeEndpointGroup	27
Request Syntax	27
Request Parameters	27
Response Syntax	27
Response Elements	27
Errors	28
See Also	28
DescribeListener	29
Request Syntax	29
Request Parameters	29
Response Syntax	29
Response Elements	29
Errors	30
See Also	30
ListAccelerators	31
Request Syntax	31
Request Parameters	31
Response Syntax	31
Response Elements	32
Errors	32
See Also	32
ListEndpointGroups	33
Request Syntax	33
Request Parameters	33
Response Syntax	33
Response Elements	34
Errors	34
See Also	35
ListListeners	36
Request Syntax	36
Request Parameters	36
Response Syntax	36
Response Elements	37
Errors	37
See Also	37
UpdateAccelerator	39
Request Syntax	39
Request Parameters	39
Response Syntax	40
Response Elements	40
Errors	40
See Also	40
UpdateAcceleratorAttributes	42
Request Syntax	42
Request Parameters	42
Response Syntax	43
Response Elements	43
Errors	43
Example	43
See Also	44
UpdateEndpointGroup	45
Request Syntax	45

Request Parameters	45
Response Syntax	47
Response Elements	47
Errors	47
Example	48
See Also	48
UpdateListener	49
Request Syntax	49
Request Parameters	49
Response Syntax	50
Response Elements	50
Errors	50
See Also	51
Data Types	52
Accelerator	53
Contents	53
See Also	54
AcceleratorAttributes	55
Contents	55
See Also	55
EndpointConfiguration	56
Contents	56
See Also	56
EndpointDescription	57
Contents	57
See Also	58
EndpointGroup	59
Contents	59
See Also	60
IpSet	61
Contents	61
See Also	61
Listener	62
Contents	62
See Also	63
PortRange	64
Contents	64
See Also	64
Common Parameters	65
Common Errors	67

Welcome

This is the *AWS Global Accelerator API Reference*. This guide is for developers who need detailed information about AWS Global Accelerator API actions, data types, and errors. For more information about Global Accelerator features, see the [AWS Global Accelerator Developer Guide](#).

AWS Global Accelerator is a network layer service in which you create accelerators to improve availability and performance for internet applications used by a global audience.

Global Accelerator provides you with static IP addresses that you associate with your accelerator. These IP addresses are anycast from the AWS edge network and distribute incoming application traffic across multiple endpoint resources in multiple AWS Regions, which increases the availability of your applications. Endpoints can be Elastic IP addresses, Network Load Balancers, and Application Load Balancers that are located in one AWS Region or multiple Regions.

Global Accelerator uses the AWS global network to route traffic to the optimal regional endpoint based on health, client location, and policies that you configure. The service reacts instantly to changes in health or configuration to ensure that internet traffic from clients is directed to only healthy endpoints.

Global Accelerator includes components that work together to help you improve performance and availability for your applications:

Static IP address

AWS Global Accelerator provides you with a set of static IP addresses which are anycast from the AWS edge network and serve as the single fixed entry points for your clients. If you already have Elastic Load Balancing or Elastic IP address resources set up for your applications, you can easily add those to Global Accelerator to allow the resources to be accessed by a Global Accelerator static IP address.

Accelerator

An accelerator directs traffic to optimal endpoints over the AWS global network to improve availability and performance for your internet applications that have a global audience. Each accelerator includes one or more listeners.

Network zone

A network zone services the static IP addresses for your accelerator from a unique IP subnet. Similar to an AWS Availability Zone, a network zone is an isolated unit with its own set of physical infrastructure. When you configure an accelerator, Global Accelerator allocates two IPv4 addresses for it. If one IP address from a network zone becomes unavailable due to IP address blocking by certain client networks, or network disruptions, then client applications can retry on the healthy static IP address from the other isolated network zone.

Listener

A listener processes inbound connections from clients to Global Accelerator, based on the protocol and port that you configure. Each listener has one or more endpoint groups associated with it, and traffic is forwarded to endpoints in one of the groups. You associate endpoint groups with listeners by specifying the Regions that you want to distribute traffic to. Traffic is distributed to optimal endpoints within the endpoint groups associated with a listener.

Endpoint group

Each endpoint group is associated with a specific AWS Region. Endpoint groups include one or more endpoints in the Region. You can increase or reduce the percentage of traffic that would be otherwise directed to an endpoint group by adjusting a setting called a *traffic dial*. The traffic dial

lets you easily do performance testing or blue/green deployment testing for new releases across different AWS Regions, for example.

Endpoint

An endpoint is an Elastic IP address, Network Load Balancer, or Application Load Balancer. Traffic is routed to endpoints based on several factors, including the geo-proximity to the user, the health of the endpoint, and the configuration options that you choose, such as endpoint weights. For each endpoint, you can configure weights, which are numbers that you can use to specify the proportion of traffic to route to each one. This can be useful, for example, to do performance testing within a Region.

This document was last published on January 16, 2019.

Actions

The following actions are supported:

- [CreateAccelerator](#) (p. 4)
- [CreateEndpointGroup](#) (p. 7)
- [CreateListener](#) (p. 12)
- [DeleteAccelerator](#) (p. 16)
- [DeleteEndpointGroup](#) (p. 18)
- [DeleteListener](#) (p. 20)
- [DescribeAccelerator](#) (p. 22)
- [DescribeAcceleratorAttributes](#) (p. 25)
- [DescribeEndpointGroup](#) (p. 27)
- [DescribeListener](#) (p. 29)
- [ListAccelerators](#) (p. 31)
- [ListEndpointGroups](#) (p. 33)
- [ListListeners](#) (p. 36)
- [UpdateAccelerator](#) (p. 39)
- [UpdateAcceleratorAttributes](#) (p. 42)
- [UpdateEndpointGroup](#) (p. 45)
- [UpdateListener](#) (p. 49)

CreateAccelerator

Create an accelerator. An accelerator includes one or more listeners that process inbound connections and direct traffic to one or more endpoint groups, each of which includes endpoints, such as Network Load Balancers. To see an AWS CLI example of creating an accelerator, scroll down to **Example**.

Request Syntax

```
{  
  "Enabled": boolean,  
  "IdempotencyToken": "string",  
  "IpAddressType": "string",  
  "Name": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 65\)](#).

The request accepts the following data in JSON format.

Enabled (p. 4)

Indicates whether an accelerator is enabled. The value is true or false. The default value is true.

If the value is set to true, an accelerator cannot be deleted. If set to false, the accelerator can be deleted.

Type: Boolean

Required: No

IdempotencyToken (p. 4)

A unique, case-sensitive identifier that you provide to ensure the idempotency—that is, the uniqueness—of an accelerator.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

IpAddressType (p. 4)

The value for the address type must be IPv4.

Type: String

Valid Values: `IPv4`

Required: No

Name (p. 4)

The name of an accelerator. The name can have a maximum of 32 characters, must contain only alphanumeric characters or hyphens (-), and must not begin or end with a hyphen.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Syntax

```
{
  "Accelerator": {
    "AcceleratorArn": "string",
    "CreatedTime": number,
    "Enabled": boolean,
    "IpAddressType": "string",
    "IpSets": [
      {
        "IpAddresses": [ "string" ],
        "IpFamily": "string"
      }
    ],
    "LastModifiedTime": number,
    "Name": "string",
    "Status": "string"
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Accelerator (p. 5)

The accelerator that is created by specifying a listener and the supported IP address types.

Type: [Accelerator \(p. 53\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 67\)](#).

InternalServerErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

LimitExceededException

Processing your request would cause you to exceed an AWS Global Accelerator limit.

HTTP Status Code: 400

Example

Create an accelerator

The following is an example of creating an accelerator and the response.

```
aws globalaccelerator create-accelerator
  --name ExampleAccelerator
  --region us-west-2
  --idempotencytoken dcba4321-dcba-4321-dcba-dcba4321
```

```
{
  "Accelerator": {
    "AcceleratorArn": "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh",
    "IpAddressType": "IPv4",
    "Name": "ExampleAccelerator",
    "Enabled": true,
    "Status": "IN_PROGRESS",
    "IpSets": [
      {
        "IpAddresses": [
          "192.0.2.250",
          "192.0.2.52"
        ],
        "IpFamily": "IPv4"
      }
    ],
    "CreatedTime": 1542394847.0,
    "LastModifiedTime": 1542394847.0
  }
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

CreateEndpointGroup

Create an endpoint group for the specified listener. An endpoint group is a collection of endpoints in one AWS Region. To see an AWS CLI example of creating an endpoint group, scroll down to **Example**.

Request Syntax

```
{
  "EndpointConfigurations": [
    {
      "EndpointId": "string",
      "Weight": number
    }
  ],
  "EndpointGroupRegion": "string",
  "HealthCheckIntervalSeconds": number,
  "HealthCheckPath": "string",
  "HealthCheckPort": number,
  "HealthCheckProtocol": "string",
  "IdempotencyToken": "string",
  "ListenerArn": "string",
  "ThresholdCount": number,
  "TrafficDialPercentage": number
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 65\)](#).

The request accepts the following data in JSON format.

EndpointConfigurations (p. 7)

The list of endpoint objects.

Type: Array of [EndpointConfiguration \(p. 56\)](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

Required: No

EndpointGroupRegion (p. 7)

The name of the AWS Region where the endpoint group is located. A listener can have only one endpoint group in a specific Region.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

HealthCheckIntervalSeconds (p. 7)

The time—10 seconds or 30 seconds—between each health check for an endpoint. The default value is 30.

Type: Integer

Valid Range: Minimum value of 10. Maximum value of 30.

Required: No

HealthCheckPath (p. 7)

If the protocol is HTTP/S, then this specifies the path that is the destination for health check targets. The default value is slash (/).

Type: String

Length Constraints: Maximum length of 255.

Required: No

HealthCheckPort (p. 7)

The port that AWS Global Accelerator uses to check the health of endpoints that are part of this endpoint group. The default port is the listener port that this endpoint group is associated with. If listener port is a list of ports, Global Accelerator uses the first port in the list.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 65535.

Required: No

HealthCheckProtocol (p. 7)

The protocol that AWS Global Accelerator uses to check the health of endpoints that are part of this endpoint group. The default value is TCP.

Type: String

Valid Values: `TCP` | `HTTP` | `HTTPS`

Required: No

IdempotencyToken (p. 7)

A unique, case-sensitive identifier that you provide to ensure the idempotency—that is, the uniqueness—of the request.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

ListenerArn (p. 7)

The Amazon Resource Name (ARN) of the listener.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

ThresholdCount (p. 7)

The number of consecutive health checks required to set the state of a healthy endpoint to unhealthy, or to set an unhealthy endpoint to healthy. The default value is 3.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 10.

Required: No

TrafficDialPercentage (p. 7)

The percentage of traffic to send to an AWS Region. Additional traffic is distributed to other endpoint groups for this listener.

Use this action to increase (dial up) or decrease (dial down) traffic to a specific Region. The percentage is applied to the traffic that would otherwise have been routed to the Region based on optimal routing.

The default value is 100.

Type: Float

Valid Range: Minimum value of 0. Maximum value of 100.

Required: No

Response Syntax

```
{
  "EndpointGroup": {
    "EndpointDescriptions": [
      {
        "EndpointId": "string",
        "HealthReason": "string",
        "HealthState": "string",
        "Weight": number
      }
    ],
    "EndpointGroupArn": "string",
    "EndpointGroupRegion": "string",
    "HealthCheckIntervalSeconds": number,
    "HealthCheckPath": "string",
    "HealthCheckPort": number,
    "HealthCheckProtocol": "string",
    "ThresholdCount": number,
    "TrafficDialPercentage": number
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

EndpointGroup (p. 9)

The information about the endpoint group that was created.

Type: [EndpointGroup \(p. 59\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 67\)](#).

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

HTTP Status Code: 400

EndpointGroupAlreadyExistsException

The endpoint group that you specified already exists.

HTTP Status Code: 400

InternalServiceErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

LimitExceededException

Processing your request would cause you to exceed an AWS Global Accelerator limit.

HTTP Status Code: 400

ListenerNotFoundException

The listener that you specified doesn't exist.

HTTP Status Code: 400

Example

Create an endpoint group

The following is an example of creating an endpoint group, and the response.

```
aws globalaccelerator create-endpoint-group
    --listener-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh/listener/0123vxyz
    --endpoint-group-region us-east-1
    --endpoint-configurations EndpointId=eipalloc-eip01234567890abc,Weight=128
    --region us-west-2
    --idempotencytoken dcba4321-dcba-4321-dcba-dcba4321
```

```
{
  "EndpointGroup": {
    "TrafficDialPercentage": 100.0,
    "EndpointDescriptions": [
      {
        "Weight": 128,
        "EndpointId": "eipalloc-eip01234567890abc"
      }
    ],
    "EndpointGroupArn": "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh/listener/0123vxyz/endpoint-group/098765zyxwvu",
```

```
    "EndpointGroupRegion": "us-east-1"  
  }  
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

CreateListener

Create a listener to process inbound connections from clients to an accelerator. Connections arrive to assigned static IP addresses on a port, port range, or list of port ranges that you specify. To see an AWS CLI example of creating a listener, scroll down to **Example**.

Request Syntax

```
{
  "AcceleratorArn": "string",
  "ClientAffinity": "string",
  "IdempotencyToken": "string",
  "PortRanges": [
    {
      "FromPort": number,
      "ToPort": number
    }
  ],
  "Protocol": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 65\)](#).

The request accepts the following data in JSON format.

AcceleratorArn (p. 12)

The Amazon Resource Name (ARN) of your accelerator.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

ClientAffinity (p. 12)

Client affinity lets you direct all requests from a user to the same endpoint, if you have stateful applications, regardless of the port and protocol of the client request. Client affinity gives you control over whether to always route each client to the same specific endpoint.

AWS Global Accelerator uses a consistent-flow hashing algorithm to choose the optimal endpoint for a connection. If client affinity is `NONE`, Global Accelerator uses the "five-tuple" (5-tuple) properties—source IP address, source port, destination IP address, destination port, and protocol—to select the hash value, and then chooses the best endpoint. However, with this setting, if someone uses different ports to connect to Global Accelerator, their connections might not be always routed to the same endpoint because the hash value changes.

If you want a given client to always be routed to the same endpoint, set client affinity to `SOURCE_IP` instead. When you use the `SOURCE_IP` setting, Global Accelerator uses the "two-tuple" (2-tuple) properties—source (client) IP address and destination IP address—to select the hash value.

The default value is `NONE`.

Type: String

Valid Values: NONE | SOURCE_IP

Required: No

IdempotencyToken (p. 12)

A unique, case-sensitive identifier that you provide to ensure the idempotency—that is, the uniqueness—of the request.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

PortRanges (p. 12)

The list of port ranges to support for connections from clients to your accelerator.

Type: Array of [PortRange \(p. 64\)](#) objects

Array Members: Minimum number of 1 item. Maximum number of 10 items.

Required: Yes

Protocol (p. 12)

The protocol for connections from clients to your accelerator.

Type: String

Valid Values: TCP | UDP

Required: Yes

Response Syntax

```
{
  "Listener": {
    "ClientAffinity": "string",
    "ListenerArn": "string",
    "PortRanges": [
      {
        "FromPort": number,
        "ToPort": number
      }
    ],
    "Protocol": "string"
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Listener (p. 13)

The listener that you've created.

Type: [Listener \(p. 62\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 67\)](#).

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

HTTP Status Code: 400

InternalServerErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

InvalidPortRangeException

The port numbers that you specified are not valid numbers or are not unique for this accelerator.

HTTP Status Code: 400

LimitExceededException

Processing your request would cause you to exceed an AWS Global Accelerator limit.

HTTP Status Code: 400

Example

Create a listener

The following is an example of creating a listener and the response.

```
aws globalaccelerator create-listener
  --accelerator-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefg
  --port-ranges FromPort=80,ToPort=80 FromPort=81,ToPort=81
  --protocol TCP
  --region us-west-2
  --idempotencytoken dcba4321-dcba-4321-dcba-dcba4321
```

```
{
  "Listener": {
    "PortRanges": [
      {
        "ToPort": 80,
        "FromPort": 80
      },
      {

```

```
        "ToPort": 81,  
        "FromPort": 81  
    }  
  ],  
  "ClientAffinity": "NONE",  
  "Protocol": "TCP",  
  "ListenerArn": "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-  
abcd-1234-abcd-1234abcdefgh/listener/0123vxyz"  
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DeleteAccelerator

Delete an accelerator. Note: before you can delete an accelerator, you must disable it and remove all dependent resources (listeners and endpoint groups).

Request Syntax

```
{  
  "AcceleratorArn": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 65).

The request accepts the following data in JSON format.

AcceleratorArn (p. 16)

The Amazon Resource Name (ARN) of an accelerator.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

For information about the errors that are common to all actions, see [Common Errors](#) (p. 67).

AcceleratorNotDisabledException

The accelerator that you specified could not be disabled.

HTTP Status Code: 400

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

HTTP Status Code: 400

AssociatedListenerFoundException

The accelerator that you specified has a listener associated with it. You must remove all dependent resources from an accelerator before you can delete it.

HTTP Status Code: 400

InternalServiceErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DeleteEndpointGroup

Delete an endpoint group from a listener.

Request Syntax

```
{  
  "EndpointGroupArn": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 65\)](#).

The request accepts the following data in JSON format.

EndpointGroupArn (p. 18)

The Amazon Resource Name (ARN) of the endpoint group to delete.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 67\)](#).

EndpointGroupNotFoundException

The endpoint group that you specified doesn't exist.

HTTP Status Code: 400

InternalServerErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DeleteListener

Delete a listener from an accelerator.

Request Syntax

```
{  
  "ListenerArn": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 65\)](#).

The request accepts the following data in JSON format.

ListenerArn (p. 20)

The Amazon Resource Name (ARN) of the listener.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 67\)](#).

AssociatedEndpointGroupFoundException

The listener that you specified has an endpoint group associated with it. You must remove all dependent resources from a listener before you can delete it.

HTTP Status Code: 400

InternalServiceErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

ListenerNotFoundException

The listener that you specified doesn't exist.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DescribeAccelerator

Describe an accelerator. To see an AWS CLI example of describing an accelerator, scroll down to **Example**.

Request Syntax

```
{  
  "AcceleratorArn": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 65\)](#).

The request accepts the following data in JSON format.

AcceleratorArn (p. 22)

The Amazon Resource Name (ARN) of the accelerator to describe.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Syntax

```
{  
  "Accelerator": {  
    "AcceleratorArn": "string",  
    "CreatedTime": number,  
    "Enabled": boolean,  
    "IpAddressType": "string",  
    "IpSets": [  
      {  
        "IpAddresses": [ "string" ],  
        "IpFamily": "string"  
      }  
    ],  
    "LastModifiedTime": number,  
    "Name": "string",  
    "Status": "string"  
  }  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Accelerator (p. 22)

The description of the accelerator.

Type: [Accelerator \(p. 53\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 67\)](#).

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

HTTP Status Code: 400

InternalServerErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

Example

Describe an accelerator

The following is an example for describing an accelerator and the response.

```
aws globalaccelerator describe-accelerator
  --accelerator-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh
  --region us-west-2
```

```
{
  "Accelerator": {
    "AcceleratorArn": "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh",
    "IpAddressType": "IPv4",
    "Name": "ExampleAccelerator",
    "Enabled": true,
    "Status": "IN_PROGRESS",
    "IpSets": [
      {
        "IpAddresses": [
          "192.0.2.250",
          "192.0.2.52"
        ],
        "IpFamily": "IPv4"
      }
    ],
    "CreatedTime": 1542394847.0,
    "LastModifiedTime": 1542395013.0
  }
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DescribeAcceleratorAttributes

Describe the attributes of an accelerator.

Request Syntax

```
{  
  "AcceleratorArn": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 65\)](#).

The request accepts the following data in JSON format.

AcceleratorArn (p. 25)

The Amazon Resource Name (ARN) of the accelerator with the attributes that you want to describe.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Syntax

```
{  
  "AcceleratorAttributes": {  
    "FlowLogsEnabled": boolean,  
    "FlowLogsS3Bucket": "string",  
    "FlowLogsS3Prefix": "string"  
  }  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

AcceleratorAttributes (p. 25)

The attributes of the accelerator.

Type: [AcceleratorAttributes \(p. 55\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 67\)](#).

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

HTTP Status Code: 400

InternalServerErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DescribeEndpointGroup

Describe an endpoint group.

Request Syntax

```
{  
  "EndpointGroupArn": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 65\)](#).

The request accepts the following data in JSON format.

EndpointGroupArn (p. 27)

The Amazon Resource Name (ARN) of the endpoint group to describe.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Syntax

```
{  
  "EndpointGroup": {  
    "EndpointDescriptions": [  
      {  
        "EndpointId": "string",  
        "HealthReason": "string",  
        "HealthState": "string",  
        "Weight": number  
      }  
    ],  
    "EndpointGroupArn": "string",  
    "EndpointGroupRegion": "string",  
    "HealthCheckIntervalSeconds": number,  
    "HealthCheckPath": "string",  
    "HealthCheckPort": number,  
    "HealthCheckProtocol": "string",  
    "ThresholdCount": number,  
    "TrafficDialPercentage": number  
  }  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[EndpointGroup \(p. 27\)](#)

The description of an endpoint group.

Type: [EndpointGroup \(p. 59\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 67\)](#).

EndpointGroupNotFoundException

The endpoint group that you specified doesn't exist.

HTTP Status Code: 400

InternalServiceErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DescribeListener

Describe a listener.

Request Syntax

```
{  
  "ListenerArn": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 65\)](#).

The request accepts the following data in JSON format.

ListenerArn (p. 29)

The Amazon Resource Name (ARN) of the listener to describe.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Syntax

```
{  
  "Listener": {  
    "ClientAffinity": "string",  
    "ListenerArn": "string",  
    "PortRanges": [  
      {  
        "FromPort": number,  
        "ToPort": number  
      }  
    ],  
    "Protocol": "string"  
  }  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Listener (p. 29)

The description of a listener.

Type: [Listener \(p. 62\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 67\)](#).

InternalServerErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

ListenerNotFoundException

The listener that you specified doesn't exist.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

ListAccelerators

List the accelerators for an AWS account.

Request Syntax

```
{  
  "MaxResults": number,  
  "NextToken": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 65).

The request accepts the following data in JSON format.

MaxResults (p. 31)

The number of Global Accelerator objects that you want to return with this call. The default value is 10.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 100.

Required: No

NextToken (p. 31)

The token for the next set of results. You receive this token from a previous call.

Type: String

Length Constraints: Maximum length of 255.

Required: No

Response Syntax

```
{  
  "Accelerators": [  
    {  
      "AcceleratorArn": "string",  
      "CreatedTime": number,  
      "Enabled": boolean,  
      "IpAddressType": "string",  
      "IpSets": [  
        {  
          "IpAddresses": [ "string ],  
          "IpFamily": "string"  
        }  
      ],  
      "LastModifiedTime": number,  
      "Name": "string",  
      "Status": "string"  
    }  
  ]  
}
```

```
    }  
  ],  
  "NextToken": "string"  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Accelerators (p. 31)

The list of accelerators for a customer account.

Type: Array of [Accelerator](#) (p. 53) objects

NextToken (p. 31)

The token for the next set of results. You receive this token from a previous call.

Type: String

Length Constraints: Maximum length of 255.

Errors

For information about the errors that are common to all actions, see [Common Errors](#) (p. 67).

InternalServerErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidNextTokenException

There isn't another item to return.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

ListEndpointGroups

List the endpoint groups that are associated with a listener.

Request Syntax

```
{
  "ListenerArn": "string",
  "MaxResults": number,
  "NextToken": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 65).

The request accepts the following data in JSON format.

ListenerArn (p. 33)

The Amazon Resource Name (ARN) of the listener.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

MaxResults (p. 33)

The number of endpoint group objects that you want to return with this call. The default value is 10.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 100.

Required: No

NextToken (p. 33)

The token for the next set of results. You receive this token from a previous call.

Type: String

Length Constraints: Maximum length of 255.

Required: No

Response Syntax

```
{
  "EndpointGroups": [
    {
      "EndpointDescriptions": [
        {
          "EndpointId": "string",
          "HealthReason": "string",

```

```
        "HealthState": "string",  
        "Weight": number  
    }  
  ],  
  "EndpointGroupArn": "string",  
  "EndpointGroupRegion": "string",  
  "HealthCheckIntervalSeconds": number,  
  "HealthCheckPath": "string",  
  "HealthCheckPort": number,  
  "HealthCheckProtocol": "string",  
  "ThresholdCount": number,  
  "TrafficDialPercentage": number  
}  
],  
"NextToken": "string"  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

EndpointGroups (p. 33)

The list of the endpoint groups associated with a listener.

Type: Array of [EndpointGroup \(p. 59\)](#) objects

NextToken (p. 33)

The token for the next set of results. You receive this token from a previous call.

Type: String

Length Constraints: Maximum length of 255.

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 67\)](#).

InternalServerErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

InvalidNextTokenException

There isn't another item to return.

HTTP Status Code: 400

ListenerNotFoundException

The listener that you specified doesn't exist.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

ListListeners

List the listeners for an accelerator.

Request Syntax

```
{
  "AcceleratorArn": "string",
  "MaxResults": number,
  "NextToken": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 65).

The request accepts the following data in JSON format.

AcceleratorArn (p. 36)

The Amazon Resource Name (ARN) of the accelerator for which you want to list listener objects.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

MaxResults (p. 36)

The number of listener objects that you want to return with this call. The default value is 10.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 100.

Required: No

NextToken (p. 36)

The token for the next set of results. You receive this token from a previous call.

Type: String

Length Constraints: Maximum length of 255.

Required: No

Response Syntax

```
{
  "Listeners": [
    {
      "ClientAffinity": "string",
      "ListenerArn": "string",
      "PortRanges": [
```

```
    {
      "FromPort": number,
      "ToPort": number
    }
  ],
  "Protocol": "string"
}
],
"NextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Listeners (p. 36)

The list of listeners for an accelerator.

Type: Array of [Listener \(p. 62\)](#) objects

NextToken (p. 36)

The token for the next set of results. You receive this token from a previous call.

Type: String

Length Constraints: Maximum length of 255.

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 67\)](#).

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

HTTP Status Code: 400

InternalServerErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidNextTokenException

There isn't another item to return.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

UpdateAccelerator

Update an accelerator.

Request Syntax

```
{  
  "AcceleratorArn": "string",  
  "Enabled": boolean,  
  "IpAddressType": "string",  
  "Name": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 65\)](#).

The request accepts the following data in JSON format.

AcceleratorArn (p. 39)

The Amazon Resource Name (ARN) of the accelerator to update.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Enabled (p. 39)

Indicates whether an accelerator is enabled. The value is true or false. The default value is true.

If the value is set to true, the accelerator cannot be deleted. If set to false, the accelerator can be deleted.

Type: Boolean

Required: No

IpAddressType (p. 39)

The value for the address type must be IPv4.

Type: String

Valid Values: `IPv4`

Required: No

Name (p. 39)

The name of the accelerator. The name can have a maximum of 32 characters, must contain only alphanumeric characters or hyphens (-), and must not begin or end with a hyphen.

Type: String

Length Constraints: Maximum length of 255.

Required: No

Response Syntax

```
{
  "Accelerator": {
    "AcceleratorArn": "string",
    "CreatedTime": number,
    "Enabled": boolean,
    "IpAddressType": "string",
    "IpSets": [
      {
        "IpAddresses": [ "string" ],
        "IpFamily": "string"
      }
    ],
    "LastModifiedTime": number,
    "Name": "string",
    "Status": "string"
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Accelerator (p. 40)

Information about the updated accelerator.

Type: [Accelerator \(p. 53\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 67\)](#).

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

HTTP Status Code: 400

InternalServerErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

UpdateAcceleratorAttributes

Update the attributes for an accelerator. To see an AWS CLI example of updating an accelerator to enable flow logs, scroll down to **Example**.

Request Syntax

```
{  
  "AcceleratorArn": "string",  
  "FlowLogsEnabled": boolean,  
  "FlowLogsS3Bucket": "string",  
  "FlowLogsS3Prefix": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 65\)](#).

The request accepts the following data in JSON format.

[AcceleratorArn \(p. 42\)](#)

The Amazon Resource Name (ARN) of the accelerator that you want to update.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

[FlowLogsEnabled \(p. 42\)](#)

Update whether flow logs are enabled. The default value is false. If the value is true, `FlowLogsS3Bucket` and `FlowLogsS3Prefix` must be specified.

For more information, see [Flow Logs](#) in the *AWS Global Accelerator Developer Guide*.

Type: Boolean

Required: No

[FlowLogsS3Bucket \(p. 42\)](#)

The name of the Amazon S3 bucket for the flow logs. Attribute is required if `FlowLogsEnabled` is true. The bucket must exist and have a bucket policy that grants AWS Global Accelerator permission to write to the bucket.

Type: String

Length Constraints: Maximum length of 255.

Required: No

[FlowLogsS3Prefix \(p. 42\)](#)

Update the prefix for the location in the Amazon S3 bucket for the flow logs. Attribute is required if `FlowLogsEnabled` is true. If you don't specify a prefix, the flow logs are stored in the root of the bucket.

Type: String

Length Constraints: Maximum length of 255.

Required: No

Response Syntax

```
{
  "AcceleratorAttributes": {
    "FlowLogsEnabled": boolean,
    "FlowLogsS3Bucket": "string",
    "FlowLogsS3Prefix": "string"
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

AcceleratorAttributes (p. 43)

Updated attributes for the accelerator.

Type: [AcceleratorAttributes](#) (p. 55) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#) (p. 67).

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

HTTP Status Code: 400

InternalServerErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

Example

Update attributes for an accelerator

The following is an example for updating an accelerator to enable flow logs.


```
aws globalaccelerator update-accelerator-attributes
  --accelerator-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh
  --flow-logs-enabled
  --flow-logs-s3-bucket your-s3-bucket-name
  --flow-logs-s3-prefix your-bucket-prefix
  --region us-west-2
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

UpdateEndpointGroup

Update an endpoint group. To see an AWS CLI example of updating an endpoint group, scroll down to **Example**.

Request Syntax

```
{
  "EndpointConfigurations": [
    {
      "EndpointId": "string",
      "Weight": number
    }
  ],
  "EndpointGroupArn": "string",
  "HealthCheckIntervalSeconds": number,
  "HealthCheckPath": "string",
  "HealthCheckPort": number,
  "HealthCheckProtocol": "string",
  "ThresholdCount": number,
  "TrafficDialPercentage": number
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 65).

The request accepts the following data in JSON format.

EndpointConfigurations (p. 45)

The list of endpoint objects.

Type: Array of [EndpointConfiguration](#) (p. 56) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

Required: No

EndpointGroupArn (p. 45)

The Amazon Resource Name (ARN) of the endpoint group.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

HealthCheckIntervalSeconds (p. 45)

The time—10 seconds or 30 seconds—between each health check for an endpoint. The default value is 30.

Type: Integer

Valid Range: Minimum value of 10. Maximum value of 30.

Required: No

HealthCheckPath (p. 45)

If the protocol is HTTP/S, then this specifies the path that is the destination for health check targets. The default value is slash (/).

Type: String

Length Constraints: Maximum length of 255.

Required: No

HealthCheckPort (p. 45)

The port that AWS Global Accelerator uses to check the health of endpoints that are part of this endpoint group. The default port is the listener port that this endpoint group is associated with. If the listener port is a list of ports, Global Accelerator uses the first port in the list.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 65535.

Required: No

HealthCheckProtocol (p. 45)

The protocol that AWS Global Accelerator uses to check the health of endpoints that are part of this endpoint group. The default value is TCP.

Type: String

Valid Values: `TCP` | `HTTP` | `HTTPS`

Required: No

ThresholdCount (p. 45)

The number of consecutive health checks required to set the state of a healthy endpoint to unhealthy, or to set an unhealthy endpoint to healthy. The default value is 3.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 10.

Required: No

TrafficDialPercentage (p. 45)

The percentage of traffic to send to an AWS Region. Additional traffic is distributed to other endpoint groups for this listener.

Use this action to increase (dial up) or decrease (dial down) traffic to a specific Region. The percentage is applied to the traffic that would otherwise have been routed to the Region based on optimal routing.

The default value is 100.

Type: Float

Valid Range: Minimum value of 0. Maximum value of 100.

Required: No

Response Syntax

```
{
  "EndpointGroup": {
    "EndpointDescriptions": [
      {
        "EndpointId": "string",
        "HealthReason": "string",
        "HealthState": "string",
        "Weight": number
      }
    ],
    "EndpointGroupArn": "string",
    "EndpointGroupRegion": "string",
    "HealthCheckIntervalSeconds": number,
    "HealthCheckPath": "string",
    "HealthCheckPort": number,
    "HealthCheckProtocol": "string",
    "ThresholdCount": number,
    "TrafficDialPercentage": number
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

EndpointGroup (p. 47)

The information about the endpoint group that was updated.

Type: [EndpointGroup \(p. 59\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 67\)](#).

EndpointGroupNotFoundException

The endpoint group that you specified doesn't exist.

HTTP Status Code: 400

InternalServerErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

LimitExceededException

Processing your request would cause you to exceed an AWS Global Accelerator limit.

HTTP Status Code: 400

Example

Update an endpoint group

The following is an example of adding an endpoint to an endpoint group and the response.

```
aws globalaccelerator update-accelerator
  --accelerator-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh
  --name ExampleAcceleratorNew
  --region us-west-2
```

```
{
  "Accelerator": {
    "AcceleratorArn": "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh",
    "IpAddressType": "IPv4",
    "Name": "ExampleAcceleratorNew",
    "Enabled": true,
    "Status": "IN_PROGRESS",
    "IpSets": [
      {
        "IpAddresses": [
          "192.0.2.250",
          "192.0.2.52"
        ],
        "IpFamily": "IPv4"
      }
    ],
    "CreatedTime": 1542394847.0,
    "LastModifiedTime": 1542395654.0
  }
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

UpdateListener

Update a listener.

Request Syntax

```
{
  "ClientAffinity": "string",
  "ListenerArn": "string",
  "PortRanges": [
    {
      "FromPort": number,
      "ToPort": number
    }
  ],
  "Protocol": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 65).

The request accepts the following data in JSON format.

ClientAffinity (p. 49)

Client affinity lets you direct all requests from a user to the same endpoint, if you have stateful applications, regardless of the port and protocol of the client request. Client affinity gives you control over whether to always route each client to the same specific endpoint.

AWS Global Accelerator uses a consistent-flow hashing algorithm to choose the optimal endpoint for a connection. If client affinity is `NONE`, Global Accelerator uses the "five-tuple" (5-tuple) properties—source IP address, source port, destination IP address, destination port, and protocol—to select the hash value, and then chooses the best endpoint. However, with this setting, if someone uses different ports to connect to Global Accelerator, their connections might not be always routed to the same endpoint because the hash value changes.

If you want a given client to always be routed to the same endpoint, set client affinity to `SOURCE_IP` instead. When you use the `SOURCE_IP` setting, Global Accelerator uses the "two-tuple" (2-tuple) properties—source (client) IP address and destination IP address—to select the hash value.

The default value is `NONE`.

Type: String

Valid Values: `NONE` | `SOURCE_IP`

Required: No

ListenerArn (p. 49)

The Amazon Resource Name (ARN) of the listener to update.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

PortRanges (p. 49)

The updated list of port ranges for the connections from clients to the accelerator.

Type: Array of [PortRange \(p. 64\)](#) objects

Array Members: Minimum number of 1 item. Maximum number of 10 items.

Required: No

Protocol (p. 49)

The updated protocol for the connections from clients to the accelerator.

Type: String

Valid Values: TCP | UDP

Required: No

Response Syntax

```
{
  "Listener": {
    "ClientAffinity": "string",
    "ListenerArn": "string",
    "PortRanges": [
      {
        "FromPort": number,
        "ToPort": number
      }
    ],
    "Protocol": "string"
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Listener (p. 50)

Information for the updated listener.

Type: [Listener \(p. 62\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 67\)](#).

InternalServerErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

InvalidPortRangeException

The port numbers that you specified are not valid numbers or are not unique for this accelerator.

HTTP Status Code: 400

LimitExceededException

Processing your request would cause you to exceed an AWS Global Accelerator limit.

HTTP Status Code: 400

ListenerNotFoundException

The listener that you specified doesn't exist.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

Data Types

The AWS Global Accelerator API contains several data types that various actions use. This section describes each data type in detail.

Note

The order of each element in a data type structure is not guaranteed. Applications should not assume a particular order.

The following data types are supported:

- [Accelerator](#) (p. 53)
- [AcceleratorAttributes](#) (p. 55)
- [EndpointConfiguration](#) (p. 56)
- [EndpointDescription](#) (p. 57)
- [EndpointGroup](#) (p. 59)
- [IpSet](#) (p. 61)
- [Listener](#) (p. 62)
- [PortRange](#) (p. 64)

Accelerator

An accelerator is a complex type that includes one or more listeners that process inbound connections and then direct traffic to one or more endpoint groups, each of which includes endpoints, such as load balancers.

Contents

AcceleratorArn

The Amazon Resource Name (ARN) of the accelerator.

Type: String

Length Constraints: Maximum length of 255.

Required: No

CreatedTime

The date and time that the accelerator was created.

Type: Timestamp

Required: No

Enabled

Indicates whether the accelerator is enabled. The value is true or false. The default value is true.

If the value is set to true, the accelerator cannot be deleted. If set to false, accelerator can be deleted.

Type: Boolean

Required: No

IpAddressType

The value for the address type must be IPv4.

Type: String

Valid Values: `IPV4`

Required: No

IpSets

IP address set associated with the accelerator.

Type: Array of [IpSet \(p. 61\)](#) objects

Required: No

LastModifiedTime

The date and time that the accelerator was last modified.

Type: Timestamp

Required: No

Name

The name of the accelerator. The name can have a maximum of 32 characters, must contain only alphanumeric characters or hyphens (-), and must not begin or end with a hyphen.

Type: String

Length Constraints: Maximum length of 255.

Required: No

Status

Describes the deployment status of the accelerator.

Type: String

Valid Values: `DEPLOYED` | `IN_PROGRESS`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

AcceleratorAttributes

Attributes of an accelerator.

Contents

FlowLogsEnabled

Indicates whether flow logs are enabled. The default value is `false`. If the value is `true`, `FlowLogsS3Bucket` and `FlowLogsS3Prefix` must be specified.

For more information, see [Flow Logs](#) in the *AWS Global Accelerator Developer Guide*.

Type: Boolean

Required: No

FlowLogsS3Bucket

The name of the Amazon S3 bucket for the flow logs. Attribute is required if `FlowLogsEnabled` is `true`. The bucket must exist and have a bucket policy that grants AWS Global Accelerator permission to write to the bucket.

Type: String

Length Constraints: Maximum length of 255.

Required: No

FlowLogsS3Prefix

The prefix for the location in the Amazon S3 bucket for the flow logs. Attribute is required if `FlowLogsEnabled` is `true`. If you don't specify a prefix, the flow logs are stored in the root of the bucket.

Type: String

Length Constraints: Maximum length of 255.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

EndpointConfiguration

A complex type for endpoints.

Contents

EndpointId

An ID for the endpoint. If the endpoint is a Network Load Balancer or Application Load Balancer, this is the Amazon Resource Name (ARN) of the resource. If the endpoint is an Elastic IP address, this is the Elastic IP address allocation ID.

Type: String

Length Constraints: Maximum length of 255.

Required: No

Weight

The weight associated with the endpoint. When you add weights to endpoints, you configure AWS Global Accelerator to route traffic based on proportions that you specify. For example, you might specify endpoint weights of 4, 5, 5, and 6 (sum=20). The result is that 4/20 of your traffic, on average, is routed to the first endpoint, 5/20 is routed both to the second and third endpoints, and 6/20 is routed to the last endpoint. For more information, see [Endpoint Weights](#) in the *AWS Global Accelerator Developer Guide*.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 255.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

EndpointDescription

A complex type for an endpoint. Each endpoint group can include one or more endpoints, such as load balancers.

Contents

EndpointId

An ID for the endpoint. If the endpoint is a Network Load Balancer or Application Load Balancer, this is the Amazon Resource Name (ARN) of the resource. If the endpoint is an Elastic IP address, this is the Elastic IP address allocation ID.

Type: String

Length Constraints: Maximum length of 255.

Required: No

HealthReason

The reason code associated with why the endpoint is not healthy. If the endpoint state is healthy, a reason code is not provided.

If the endpoint state is **unhealthy**, the reason code can be one of the following values:

- **Timeout:** The health check requests to the endpoint are timing out before returning a status.
- **Failed:** The health check failed, for example because the endpoint response was invalid (malformed).

If the endpoint state is **initial**, the reason code can be one of the following values:

- **ProvisioningInProgress:** The endpoint is in the process of being provisioned.
- **InitialHealthChecking:** Global Accelerator is still setting up the minimum number of health checks for the endpoint that are required to determine its health status.

Type: String

Length Constraints: Maximum length of 255.

Required: No

HealthState

The health status of the endpoint.

Type: String

Valid Values: `INITIAL` | `HEALTHY` | `UNHEALTHY`

Required: No

Weight

The weight associated with the endpoint. When you add weights to endpoints, you configure AWS Global Accelerator to route traffic based on proportions that you specify. For example, you might specify endpoint weights of 4, 5, 5, and 6 (sum=20). The result is that 4/20 of your traffic, on average, is routed to the first endpoint, 5/20 is routed both to the second and third endpoints, and 6/20 is routed to the last endpoint. For more information, see [Endpoint Weights](#) in the *AWS Global Accelerator Developer Guide*.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 255.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

EndpointGroup

A complex type for the endpoint group. An AWS Region can have only one endpoint group for a specific listener.

Contents

EndpointDescriptions

The list of endpoint objects.

Type: Array of [EndpointDescription \(p. 57\)](#) objects

Required: No

EndpointGroupArn

The Amazon Resource Name (ARN) of the endpoint group.

Type: String

Length Constraints: Maximum length of 255.

Required: No

EndpointGroupRegion

The AWS Region that this endpoint group belongs.

Type: String

Length Constraints: Maximum length of 255.

Required: No

HealthCheckIntervalSeconds

The time—10 seconds or 30 seconds—between health checks for each endpoint. The default value is 30.

Type: Integer

Valid Range: Minimum value of 10. Maximum value of 30.

Required: No

HealthCheckPath

If the protocol is HTTP/S, then this value provides the ping path that Global Accelerator uses for the destination on the endpoints for health checks. The default is slash (/).

Type: String

Length Constraints: Maximum length of 255.

Required: No

HealthCheckPort

The port that Global Accelerator uses to perform health checks on endpoints that are part of this endpoint group.

The default port is the port for the listener that this endpoint group is associated with. If the listener port is a list, Global Accelerator uses the first specified port in the list of ports.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 65535.

Required: No

HealthCheckProtocol

The protocol that Global Accelerator uses to perform health checks on endpoints that are part of this endpoint group. The default value is TCP.

Type: String

Valid Values: `TCP` | `HTTP` | `HTTPS`

Required: No

ThresholdCount

The number of consecutive health checks required to set the state of a healthy endpoint to unhealthy, or to set an unhealthy endpoint to healthy. The default value is 3.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 10.

Required: No

TrafficDialPercentage

The percentage of traffic to send to an AWS Region. Additional traffic is distributed to other endpoint groups for this listener.

Use this action to increase (dial up) or decrease (dial down) traffic to a specific Region. The percentage is applied to the traffic that would otherwise have been routed to the Region based on optimal routing.

The default value is 100.

Type: Float

Valid Range: Minimum value of 0. Maximum value of 100.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

IpSet

A complex type for the set of IP addresses for an accelerator.

Contents

IpAddresses

The array of IP addresses in the IP address set. An IP address set can have a maximum of two IP addresses.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 2 items.

Required: No

IpFamily

The types of IP addresses included in this IP set.

Type: String

Length Constraints: Maximum length of 255.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

Listener

A complex type for a listener.

Contents

ClientAffinity

Client affinity lets you direct all requests from a user to the same endpoint, if you have stateful applications, regardless of the port and protocol of the client request. Client affinity gives you control over whether to always route each client to the same specific endpoint.

AWS Global Accelerator uses a consistent-flow hashing algorithm to choose the optimal endpoint for a connection. If client affinity is `NONE`, Global Accelerator uses the "five-tuple" (5-tuple) properties—source IP address, source port, destination IP address, destination port, and protocol—to select the hash value, and then chooses the best endpoint. However, with this setting, if someone uses different ports to connect to Global Accelerator, their connections might not be always routed to the same endpoint because the hash value changes.

If you want a given client to always be routed to the same endpoint, set client affinity to `SOURCE_IP` instead. When you use the `SOURCE_IP` setting, Global Accelerator uses the "two-tuple" (2-tuple) properties—source (client) IP address and destination IP address—to select the hash value.

The default value is `NONE`.

Type: String

Valid Values: `NONE` | `SOURCE_IP`

Required: No

ListenerArn

The Amazon Resource Name (ARN) of the listener.

Type: String

Length Constraints: Maximum length of 255.

Required: No

PortRanges

The list of port ranges for the connections from clients to the accelerator.

Type: Array of [PortRange \(p. 64\)](#) objects

Array Members: Minimum number of 1 item. Maximum number of 10 items.

Required: No

Protocol

The protocol for the connections from clients to the accelerator.

Type: String

Valid Values: `TCP` | `UDP`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

PortRange

A complex type for a range of ports for a listener.

Contents

FromPort

The first port in the range of ports, inclusive.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 65535.

Required: No

ToPort

The last port in the range of ports, inclusive.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 65535.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

Common Parameters

The following list contains the parameters that all actions use for signing Signature Version 4 requests with a query string. Any action-specific parameters are listed in the topic for that action. For more information about Signature Version 4, see [Signature Version 4 Signing Process](#) in the *Amazon Web Services General Reference*.

Action

The action to be performed.

Type: string

Required: Yes

Version

The API version that the request is written for, expressed in the format YYYY-MM-DD.

Type: string

Required: Yes

X-Amz-Algorithm

The hash algorithm that you used to create the request signature.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Valid Values: `AWS4-HMAC-SHA256`

Required: Conditional

X-Amz-Credential

The credential scope value, which is a string that includes your access key, the date, the region you are targeting, the service you are requesting, and a termination string ("aws4_request"). The value is expressed in the following format: `access_key/YYYYMMDD/region/service/aws4_request`.

For more information, see [Task 2: Create a String to Sign for Signature Version 4](#) in the *Amazon Web Services General Reference*.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Required: Conditional

X-Amz-Date

The date that is used to create the signature. The format must be ISO 8601 basic format (YYYYMMDD'THHMMSS'Z'). For example, the following date time is a valid X-Amz-Date value: `20120325T120000Z`.

Condition: X-Amz-Date is optional for all requests; it can be used to override the date used for signing requests. If the Date header is specified in the ISO 8601 basic format, X-Amz-Date is

not required. When X-Amz-Date is used, it always overrides the value of the Date header. For more information, see [Handling Dates in Signature Version 4](#) in the *Amazon Web Services General Reference*.

Type: string

Required: Conditional

X-Amz-Security-Token

The temporary security token that was obtained through a call to AWS Security Token Service (AWS STS). For a list of services that support temporary security credentials from AWS Security Token Service, go to [AWS Services That Work with IAM](#) in the *IAM User Guide*.

Condition: If you're using temporary security credentials from the AWS Security Token Service, you must include the security token.

Type: string

Required: Conditional

X-Amz-Signature

Specifies the hex-encoded signature that was calculated from the string to sign and the derived signing key.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Required: Conditional

X-Amz-SignedHeaders

Specifies all the HTTP headers that were included as part of the canonical request. For more information about specifying signed headers, see [Task 1: Create a Canonical Request For Signature Version 4](#) in the *Amazon Web Services General Reference*.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Required: Conditional

Common Errors

This section lists the errors common to the API actions of all AWS services. For errors specific to an API action for this service, see the topic for that API action.

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

IncompleteSignature

The request signature does not conform to AWS standards.

HTTP Status Code: 400

InternalFailure

The request processing has failed because of an unknown error, exception or failure.

HTTP Status Code: 500

InvalidAction

The action or operation requested is invalid. Verify that the action is typed correctly.

HTTP Status Code: 400

InvalidClientTokenId

The X.509 certificate or AWS access key ID provided does not exist in our records.

HTTP Status Code: 403

InvalidParameterCombination

Parameters that must not be used together were used together.

HTTP Status Code: 400

InvalidParameterValue

An invalid or out-of-range value was supplied for the input parameter.

HTTP Status Code: 400

InvalidQueryParameter

The AWS query string is malformed or does not adhere to AWS standards.

HTTP Status Code: 400

MalformedQueryString

The query string contains a syntax error.

HTTP Status Code: 404

MissingAction

The request is missing an action or a required parameter.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 403

MissingParameter

A required parameter for the specified action is not supplied.

HTTP Status Code: 400

OptInRequired

The AWS access key ID needs a subscription for the service.

HTTP Status Code: 403

RequestExpired

The request reached the service more than 15 minutes after the date stamp on the request or more than 15 minutes after the request expiration date (such as for pre-signed URLs), or the date stamp on the request is more than 15 minutes in the future.

HTTP Status Code: 400

ServiceUnavailable

The request has failed due to a temporary failure of the server.

HTTP Status Code: 503

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

ValidationError

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400