
AWS Global Accelerator

API Reference

API Version 2018-08-08



AWS Global Accelerator: API Reference

Copyright © 2020 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
AdvertiseByoipCidr	4
Request Syntax	4
Request Parameters	4
Response Syntax	4
Response Elements	4
Errors	5
Example	5
See Also	6
CreateAccelerator	7
Request Syntax	7
Request Parameters	7
Response Syntax	8
Response Elements	9
Errors	9
Example	9
See Also	10
CreateEndpointGroup	11
Request Syntax	11
Request Parameters	11
Response Syntax	13
Response Elements	13
Errors	14
Example	14
See Also	15
CreateListener	16
Request Syntax	16
Request Parameters	16
Response Syntax	17
Response Elements	17
Errors	18
Example	18
See Also	19
DeleteAccelerator	20
Request Syntax	20
Request Parameters	20
Response Elements	20
Errors	20
See Also	21
DeleteEndpointGroup	22
Request Syntax	22
Request Parameters	22
Response Elements	22
Errors	22
See Also	22
DeleteListener	24
Request Syntax	24
Request Parameters	24
Response Elements	24
Errors	24
See Also	25
DeprovisionByoipCidr	26
Request Syntax	26

Request Parameters	26
Response Syntax	26
Response Elements	26
Errors	27
Example	27
See Also	28
DescribeAccelerator	29
Request Syntax	29
Request Parameters	29
Response Syntax	29
Response Elements	29
Errors	30
Example	30
See Also	31
DescribeAcceleratorAttributes	32
Request Syntax	32
Request Parameters	32
Response Syntax	32
Response Elements	32
Errors	32
Example	33
See Also	33
DescribeEndpointGroup	34
Request Syntax	34
Request Parameters	34
Response Syntax	34
Response Elements	34
Errors	35
Example	35
See Also	36
DescribeListener	37
Request Syntax	37
Request Parameters	37
Response Syntax	37
Response Elements	37
Errors	38
Example	38
See Also	38
ListAccelerators	40
Request Syntax	40
Request Parameters	40
Response Syntax	40
Response Elements	41
Errors	41
Example	41
See Also	42
ListByoipCidrs	44
Request Syntax	44
Request Parameters	44
Response Syntax	44
Response Elements	45
Errors	45
Example	45
See Also	46
ListEndpointGroups	48
Request Syntax	48
Request Parameters	48

Response Syntax	48
Response Elements	49
Errors	49
Example	50
See Also	50
ListListeners	52
Request Syntax	52
Request Parameters	52
Response Syntax	52
Response Elements	53
Errors	53
Example	54
See Also	54
ListTagsForResource	55
Request Syntax	55
Request Parameters	55
Response Syntax	55
Response Elements	55
Errors	56
Example	56
See Also	56
ProvisionByoipCidr	58
Request Syntax	58
Request Parameters	58
Response Syntax	58
Response Elements	59
Errors	59
Example	59
See Also	60
TagResource	61
Request Syntax	61
Request Parameters	61
Response Elements	61
Errors	61
Example	62
See Also	62
UntagResource	63
Request Syntax	63
Request Parameters	63
Response Elements	63
Errors	63
Example	64
See Also	64
UpdateAccelerator	65
Request Syntax	65
Request Parameters	65
Response Syntax	66
Response Elements	66
Errors	66
Example	67
See Also	67
UpdateAcceleratorAttributes	68
Request Syntax	68
Request Parameters	68
Response Syntax	69
Response Elements	69
Errors	69

Example	70
See Also	70
UpdateEndpointGroup	71
Request Syntax	71
Request Parameters	71
Response Syntax	73
Response Elements	73
Errors	73
Example	74
See Also	74
UpdateListener	76
Request Syntax	76
Request Parameters	76
Response Syntax	77
Response Elements	77
Errors	77
Example	78
See Also	78
WithdrawByoipCidr	80
Request Syntax	80
Request Parameters	80
Response Syntax	80
Response Elements	80
Errors	81
Example	81
See Also	82
Data Types	83
Accelerator	84
Contents	84
See Also	85
AcceleratorAttributes	86
Contents	86
See Also	86
ByoipCidr	87
Contents	87
See Also	88
ByoipCidrEvent	89
Contents	89
See Also	89
CidrAuthorizationContext	90
Contents	90
See Also	90
EndpointConfiguration	91
Contents	91
See Also	91
EndpointDescription	93
Contents	93
See Also	94
EndpointGroup	95
Contents	95
See Also	96
IpSet	97
Contents	97
See Also	97
Listener	98
Contents	98
See Also	99

PortRange	100
Contents	100
See Also	100
Tag	101
Contents	101
See Also	101
Common Parameters	102
Common Errors	104

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 service in which you create *accelerators* to improve availability and performance of your applications for local and global users. Global Accelerator directs traffic to optimal endpoints over the AWS global network. This improves the availability and performance of your internet applications that are used by a global audience. Global Accelerator is a global service that supports endpoints in multiple AWS Regions, which are listed in the [AWS Region Table](#).

Important

Global Accelerator is a global service that supports endpoints in multiple AWS Regions but you must specify the US West (Oregon) Region to create or update accelerators.

By default, Global Accelerator provides you with static IP addresses that you associate with your accelerator. (Instead of using the IP addresses that Global Accelerator provides, you can configure these entry points to be IPv4 addresses from your own IP address ranges that you bring to Global Accelerator.) The static 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 Network Load Balancers, Application Load Balancers, EC2 instances, or Elastic IP addresses 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

By default, AWS Global Accelerator provides you with a set of static IP addresses that are anycast from the AWS edge network and serve as the single fixed entry points for your clients. Or you can configure these entry points to be IPv4 addresses from your own IP address ranges that you bring to Global Accelerator (BYOIP). For more information, see [Bring Your Own IP Addresses \(BYOIP\)](#) in the *AWS Global Accelerator Developer Guide*. If you already have load balancers, EC2 instances, or Elastic IP addresses set up for your applications, you can easily add those to Global Accelerator to allow the resources to be accessed by the static IP addresses.

Important

The static IP addresses remain assigned to your accelerator for as long as it exists, even if you disable the accelerator and it no longer accepts or routes traffic. However, when you *delete* an accelerator, you lose the static IP addresses that are assigned to it, so you can no longer route traffic by using them. You can use IAM policies with Global Accelerator to limit the users who have permissions to delete an accelerator. For more information, see [Authentication and Access Control](#) in the *AWS Global Accelerator Developer Guide*.

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.

DNS name

Global Accelerator assigns each accelerator a default Domain Name System (DNS) name, similar to `a1234567890abcdef.awsglobalaccelerator.com`, that points to your Global Accelerator

static IP addresses. Depending on the use case, you can use your accelerator's static IP addresses or DNS name to route traffic to your accelerator, or set up DNS records to route traffic using your own custom domain name.

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, by default, 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 a Network Load Balancer, Application Load Balancer, EC2 instance, or Elastic IP address. 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 September 24, 2020.

Actions

The following actions are supported:

- [AdvertiseByoipCidr](#) (p. 4)
- [CreateAccelerator](#) (p. 7)
- [CreateEndpointGroup](#) (p. 11)
- [CreateListener](#) (p. 16)
- [DeleteAccelerator](#) (p. 20)
- [DeleteEndpointGroup](#) (p. 22)
- [DeleteListener](#) (p. 24)
- [DeprovisionByoipCidr](#) (p. 26)
- [DescribeAccelerator](#) (p. 29)
- [DescribeAcceleratorAttributes](#) (p. 32)
- [DescribeEndpointGroup](#) (p. 34)
- [DescribeListener](#) (p. 37)
- [ListAccelerators](#) (p. 40)
- [ListByoipCidrs](#) (p. 44)
- [ListEndpointGroups](#) (p. 48)
- [ListListeners](#) (p. 52)
- [ListTagsForResource](#) (p. 55)
- [ProvisionByoipCidr](#) (p. 58)
- [TagResource](#) (p. 61)
- [UntagResource](#) (p. 63)
- [UpdateAccelerator](#) (p. 65)
- [UpdateAcceleratorAttributes](#) (p. 68)
- [UpdateEndpointGroup](#) (p. 71)
- [UpdateListener](#) (p. 76)
- [WithdrawByoipCidr](#) (p. 80)

AdvertiseByoipCidr

Advertises an IPv4 address range that is provisioned for use with your AWS resources through bring your own IP addresses (BYOIP). It can take a few minutes before traffic to the specified addresses starts routing to AWS because of propagation delays. To see an AWS CLI example of advertising an address range, scroll down to **Example**.

To stop advertising the BYOIP address range, use [WithdrawByoipCidr](#).

For more information, see [Bring Your Own IP Addresses \(BYOIP\)](#) in the *AWS Global Accelerator Developer Guide*.

Request Syntax

```
{
  "Cidr": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 102\)](#).

The request accepts the following data in JSON format.

Cidr (p. 4)

The address range, in CIDR notation. This must be the exact range that you provisioned. You can't advertise only a portion of the provisioned range.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Syntax

```
{
  "ByoipCidr": {
    "Cidr": "string",
    "Events": [
      {
        "Message": "string",
        "Timestamp": number
      }
    ],
    "State": "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.

ByoipCidr (p. 4)

Information about the address range.

Type: [ByoipCidr \(p. 87\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 104\)](#).

AccessDeniedException

You don't have access permission.

HTTP Status Code: 400

ByoipCidrNotFoundException

The CIDR that you specified was not found or is incorrect.

HTTP Status Code: 400

IncorrectCidrStateException

The CIDR that you specified is not valid for this action. For example, the state of the CIDR might be incorrect for this action.

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

Example

Advertise address range

The following is an example of advertising an address range, and the response.

```
aws globalaccelerator advertise-byoip-cidr --cidr "198.51.100.0/24"
```

```
{
  "ByoipCidr": {
    "Cidr": "198.51.100.0/24",
    "State": "PENDING_ADVERTISING"
  }
}
```

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 V3](#)

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**.

Important

Global Accelerator is a global service that supports endpoints in multiple AWS Regions but you must specify the US West (Oregon) Region to create or update accelerators.

Request Syntax

```
{
  "Enabled": boolean,
  "IdempotencyToken": "string",
  "IpAddresses": [ "string" ],
  "IpAddressType": "string",
  "Name": "string",
  "Tags": [
    {
      "Key": "string",
      "Value": "string"
    }
  ]
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 102\)](#).

The request accepts the following data in JSON format.

Enabled (p. 7)

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. 7)

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

IpAddresses (p. 7)

Optionally, if you've added your own IP address pool to Global Accelerator (BYOIP), you can choose IP addresses from your own pool to use for the accelerator's static IP addresses when you create an

accelerator. You can specify one or two addresses, separated by a comma. Do not include the /32 suffix.

Only one IP address from each of your IP address ranges can be used for each accelerator. If you specify only one IP address from your IP address range, Global Accelerator assigns a second static IP address for the accelerator from the AWS IP address pool.

Note that you can't update IP addresses for an existing accelerator. To change them, you must create a new accelerator with the new addresses.

For more information, see [Bring Your Own IP Addresses \(BYOIP\)](#) in the *AWS Global Accelerator Developer Guide*.

Type: Array of strings

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

Required: No

IpAddressType (p. 7)

The value for the address type must be IPv4.

Type: String

Valid Values: IPV4

Required: No

Name (p. 7)

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

Tags (p. 7)

Create tags for an accelerator.

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

Type: Array of [Tag \(p. 101\)](#) objects

Required: No

Response Syntax

```
{
  "Accelerator": {
    "AcceleratorArn": "string",
    "CreatedTime": number,
    "DnsName": "string",
    "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. 8)

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

Type: [Accelerator \(p. 84\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 104\)](#).

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 with two tags, and the response (which does not include the tag information).

```
aws globalaccelerator create-accelerator
  --name ExampleAccelerator
  --tags Key="Name",Value="Example Name" Key="Project",Value="Example Project"
  --region us-west-2
  --ip-addresses 192.0.2.250 198.51.100.52
```



```
{
  "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",
          "198.51.100.52"
        ],
        "IpFamily": "IPv4"
      }
    ],
    "DnsName": "a1234567890abcdef.awsglobalaccelerator.com",
    "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 V3](#)

CreateEndpointGroup

Create an endpoint group for the specified listener. An endpoint group is a collection of endpoints in one AWS Region. A resource must be valid and active when you add it as an endpoint.

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

Request Syntax

```
{
  "EndpointConfigurations": [
    {
      "ClientIPPreservationEnabled": boolean,
      "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. 102).

The request accepts the following data in JSON format.

EndpointConfigurations (p. 11)

The list of endpoint objects.

Type: Array of [EndpointConfiguration](#) (p. 91) objects

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

Required: No

EndpointGroupRegion (p. 11)

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. 11)

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. 11)

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. 11)

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. 11)

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. 11)

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. 11)

The Amazon Resource Name (ARN) of the listener.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

ThresholdCount (p. 11)

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. 11)

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": [
      {
        "ClientIPPreservationEnabled": boolean,
        "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. 13)

The information about the endpoint group that was created.

Type: [EndpointGroup \(p. 95\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 104\)](#).

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

HTTP Status Code: 400

AccessDeniedException

You don't have access permission.

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=i-1234567890abcdef0,Weight=128
    --region us-west-2
```

```
{
```

```
"EndpointGroup": {
  "TrafficDialPercentage": 100.0,
  "EndpointDescriptions": [
    {
      "Weight": 128,
      "EndpointId": "i-1234567890abcdef0"
    }
  ],
  "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 V3](#)

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 \(p. 102\)](#).

The request accepts the following data in JSON format.

AcceleratorArn (p. 16)

The Amazon Resource Name (ARN) of your accelerator.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

ClientAffinity (p. 16)

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. 16)

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. 16)

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

Type: Array of [PortRange \(p. 100\)](#) objects

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

Required: Yes

Protocol (p. 16)

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. 17)

The listener that you've created.

Type: [Listener \(p. 98\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 104\)](#).

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
```

```
{
  "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 V3](#)

DeleteAccelerator

Delete an accelerator. Before you can delete an accelerator, you must disable it and remove all dependent resources (listeners and endpoint groups). To disable the accelerator, update the accelerator to set `Enabled` to `false`.

Important

When you create an accelerator, by default, Global Accelerator provides you with a set of two static IP addresses. Alternatively, you can bring your own IP address ranges to Global Accelerator and assign IP addresses from those ranges.

The IP addresses are assigned to your accelerator for as long as it exists, even if you disable the accelerator and it no longer accepts or routes traffic. However, when you *delete* an accelerator, you lose the static IP addresses that are assigned to the accelerator, so you can no longer route traffic by using them. As a best practice, ensure that you have permissions in place to avoid inadvertently deleting accelerators. You can use IAM policies with Global Accelerator to limit the users who have permissions to delete an accelerator. For more information, see [Authentication and Access Control](#) in the *AWS Global Accelerator Developer Guide*.

Request Syntax

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

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 102).

The request accepts the following data in JSON format.

AcceleratorArn (p. 20)

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. 104).

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

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 V3](#)

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 \(p. 102\)](#).

The request accepts the following data in JSON format.

EndpointGroupArn (p. 22)

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 \(p. 104\)](#).

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

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 V3](#)

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. 102).

The request accepts the following data in JSON format.

ListenerArn (p. 24)

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. 104).

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

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 V3](#)

DeprovisionByoipCidr

Releases the specified address range that you provisioned to use with your AWS resources through bring your own IP addresses (BYOIP) and deletes the corresponding address pool. To see an AWS CLI example of deprovisioning an address range, scroll down to **Example**.

Before you can release an address range, you must stop advertising it by using [WithdrawByoipCidr](#) and you must not have any accelerators that are using static IP addresses allocated from its address range.

For more information, see [Bring Your Own IP Addresses \(BYOIP\)](#) in the *AWS Global Accelerator Developer Guide*.

Request Syntax

```
{
  "Cidr": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 102).

The request accepts the following data in JSON format.

Cidr (p. 26)

The address range, in CIDR notation. The prefix must be the same prefix that you specified when you provisioned the address range.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Syntax

```
{
  "ByoipCidr": {
    "Cidr": "string",
    "Events": [
      {
        "Message": "string",
        "Timestamp": number
      }
    ],
    "State": "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.

ByoipCidr (p. 26)

Information about the address range.

Type: [ByoipCidr](#) (p. 87) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#) (p. 104).

AccessDeniedException

You don't have access permission.

HTTP Status Code: 400

ByoipCidrNotFoundException

The CIDR that you specified was not found or is incorrect.

HTTP Status Code: 400

IncorrectCidrStateException

The CIDR that you specified is not valid for this action. For example, the state of the CIDR might be incorrect for this action.

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

Example

Deprovision address range

The following is an example of deprovisioning an address range, and the response.

```
aws globalaccelerator deprovision-byoip-cidr --cidr "198.51.100.0/24"
```

```
{
  "ByoipCidr": {
    "Cidr": "198.51.100.0/24",
    "State": "PENDING_DEPROVISIONING"
  }
}
```

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 V3](#)

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. 102).

The request accepts the following data in JSON format.

AcceleratorArn (p. 29)

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,  
    "DnsName": "string",  
    "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. 29)

The description of the accelerator.

Type: [Accelerator \(p. 84\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 104\)](#).

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",
          "198.51.100.52"
        ],
        "IpFamily": "IPv4"
      }
    ],
    "DnsName": "a1234567890abcdef.awsglobalaccelerator.com",
  }
}
```

```
"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 V3](#)

DescribeAcceleratorAttributes

Describe the attributes of an accelerator. To see an AWS CLI example of describing the attributes of 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. 102).

The request accepts the following data in JSON format.

AcceleratorArn (p. 32)

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. 32)

The attributes of the accelerator.

Type: [AcceleratorAttributes](#) (p. 86) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#) (p. 104).

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 attributes for an accelerator

The following is an example for describing the attributes for an accelerator.

```
aws globalaccelerator describe-accelerator-attributes
  --accelerator-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh
```

```
{
  "AcceleratorAttributes": {
    "FlowLogsEnabled": true
    "FlowLogsS3Bucket": flowlogs-abc
    "FlowLogsS3Prefix": bucketprefix-abc
  }
}
```

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 V3](#)

DescribeEndpointGroup

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

Request Syntax

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

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 102).

The request accepts the following data in JSON format.

EndpointGroupArn (p. 34)

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": [  
      {  
        "ClientIPPreservationEnabled": boolean,  
        "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. 34)

The description of an endpoint group.

Type: [EndpointGroup \(p. 95\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 104\)](#).

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

Example

Describe an endpoint group

The following is an example for describing an endpoint group, and the response.

```
aws globalaccelerator describe-endpoint-group
  --endpoint-group-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh/listener/6789vxyz-vxyz-6789-vxyz-6789lmnopqrs/endpoint-group/
ab8888example
```

```
{
  "EndpointGroup": {
    "TrafficDialPercentage": 100.0,
    "EndpointDescriptions": [
      {
        "Weight": 128,
        "EndpointId": "i-1234567890abcdef0"
      },
      {
        "Weight": 128,
        "EndpointId": "arn:aws:elasticloadbalancing:us-
east-1:000123456789:loadbalancer/app/ALBTesting/alb01234567890xyz"
      },
      {
        "Weight": 128,
        "EndpointId": "arn:aws:elasticloadbalancing:us-
east-1:000123456789:loadbalancer/net/NLBTesting/alb01234567890qrs"
      }
    ]
  }
}
```

```
    }  
  ],  
  "EndpointGroupArn": "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-  
abcd-1234-abcd-1234abcdefgh/listener/6789vxyz-vxyz-6789-vxyz-6789lmnopqrs/endpoint-  
group/4321abcd-abcd-4321-abcd-4321abcdefg",  
  "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 V3](#)

DescribeListener

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

Request Syntax

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

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 102\)](#).

The request accepts the following data in JSON format.

ListenerArn (p. 37)

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. 37)

The description of a listener.

Type: [Listener \(p. 98\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 104\)](#).

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

Example

Describe a listener

The following is an example for describing a listener, and the response.

```
aws globalaccelerator describe-listener
  --listener-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh/listener/abcdef1234
  --region us-west-2
```

```
{
  "Listener": {
    "ListenerArn": "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh/listener/abcdef1234",
    "PortRanges": [
      {
        "FromPort": 80,
        "ToPort": 80
      }
    ],
    "Protocol": "TCP",
    "ClientAffinity": "NONE"
  }
}
```

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 V3](#)

ListAccelerators

List the accelerators for an AWS account. To see an AWS CLI example of listing the accelerators for an AWS account, scroll down to **Example**.

Request Syntax

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

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 102\)](#).

The request accepts the following data in JSON format.

MaxResults (p. 40)

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. 40)

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,
      "DnsName": "string",
      "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. 40)

The list of accelerators for a customer account.

Type: Array of [Accelerator \(p. 84\)](#) objects

NextToken (p. 40)

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 \(p. 104\)](#).

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

Example

List accelerators

The following is an example for listing the accelerators for an AWS account, and the response.

```
aws globalaccelerator list-accelerators --region us-west-2
```


- [AWS SDK for Ruby V3](#)

ListByoipCidrs

Lists the IP address ranges that were specified in calls to [ProvisionByoipCidr](#), including the current state and a history of state changes.

To see an AWS CLI example of listing BYOIP CIDR addresses, scroll down to **Example**.

Request Syntax

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

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 102).

The request accepts the following data in JSON format.

MaxResults (p. 44)

The maximum number of results to return with a single call. To retrieve the remaining results, make another call with the returned `nextToken` value.

Type: Integer

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

Required: No

NextToken (p. 44)

The token for the next page of results.

Type: String

Length Constraints: Maximum length of 255.

Required: No

Response Syntax

```
{
  "ByoipCidrs": [
    {
      "Cidr": "string",
      "Events": [
        {
          "Message": "string",
          "Timestamp": number
        }
      ],
      "State": "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.

ByoipCidrs (p. 44)

Information about your address ranges.

Type: Array of [ByoipCidr \(p. 87\)](#) objects

NextToken (p. 44)

The token for the next page of results.

Type: String

Length Constraints: Maximum length of 255.

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 104\)](#).

AccessDeniedException

You don't have access permission.

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

InvalidNextTokenException

There isn't another item to return.

HTTP Status Code: 400

Example

List BYOIP CIDR addresses

The following is an example of listing BYOIP CIDR addresses and the response.

```
aws globalaccelerator list-byoip-cidrs
```

```

{
  "ByoipCidrs": [
    {
      "Cidr": "198.51.100.0/24",
      "State": "DEPROVISIONED",
      "Events": [
        {
          "Message": "CIDR is deprovisioned",
          "Timestamp": 1584977476.0
        },
        {
          "Message": "Initiated CIDR deprovisioning",
          "Timestamp": 1584977440.0
        },
        {
          "Message": "CIDR is no longer advertising",
          "Timestamp": 1584977429.0
        },
        {
          "Message": "Initiated CIDR withdrawal",
          "Timestamp": 1584977348.0
        },
        {
          "Message": "CIDR is advertising",
          "Timestamp": 1584977334.0
        },
        {
          "Message": "Initiated CIDR advertisement",
          "Timestamp": 1584977298.0
        },
        {
          "Message": "CIDR is provisioned",
          "Timestamp": 1584977280.0
        },
        {
          "Message": "Initiated CIDR provisioning",
          "Timestamp": 1584977239.0
        }
      ]
    }
  ]
}
{
  "Cidr": "203.0.113.25/24",
  "State": "READY",
  "Events": [
    {
      "Message": "CIDR is provisioned",
      "Timestamp": 1584977802.0
    },
    {
      "Message": "Initiated CIDR provisioning",
      "Timestamp": 1584977367.0
    }
  ]
}
]
}

```

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 V3](#)

ListEndpointGroups

List the endpoint groups that are associated with a listener. To see an AWS CLI example of listing the endpoint groups for listener, scroll down to **Example**.

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. 102).

The request accepts the following data in JSON format.

ListenerArn (p. 48)

The Amazon Resource Name (ARN) of the listener.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

MaxResults (p. 48)

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. 48)

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": [
        {
```

```
        "ClientIPPreservationEnabled": boolean,
        "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. 48)

The list of the endpoint groups associated with a listener.

Type: Array of [EndpointGroup \(p. 95\)](#) objects

NextToken (p. 48)

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 \(p. 104\)](#).

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

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

Example

List endpoint groups

The following is an example for listing the endpoint groups for listener, and the response.

```
aws globalaccelerator list-endpoint-groups
  --listener-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh/listener/abcdef1234
  --region us-west-2
```

```
{
  "EndpointGroups": [
    {
      "EndpointGroupArn":
"arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-abcd-1234-abcd-1234abcdefgh/
listener/abcdef1234/endpoint-group/ab88888example",
      "EndpointGroupRegion": "eu-central-1",
      "EndpointDescriptions": [],
      "TrafficDialPercentage": 100.0,
      "HealthCheckPort": 80,
      "HealthCheckProtocol": "TCP",
      "HealthCheckIntervalSeconds": 30,
      "ThresholdCount": 3
    }
    {
      "EndpointGroupArn":
"arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-abcd-1234-abcd-1234abcdefgh/
listener/abcdef1234/endpoint-group/ab99999example",
      "EndpointGroupRegion": "us-east-1",
      "EndpointDescriptions": [],
      "TrafficDialPercentage": 50.0,
      "HealthCheckPort": 80,
      "HealthCheckProtocol": "TCP",
      "HealthCheckIntervalSeconds": 30,
      "ThresholdCount": 3
    }
  ]
}
```

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 V3](#)

ListListeners

List the listeners for an accelerator. To see an AWS CLI example of listing the listeners for an accelerator, scroll down to **Example**.

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. 102).

The request accepts the following data in JSON format.

AcceleratorArn (p. 52)

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. 52)

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. 52)

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. 52)

The list of listeners for an accelerator.

Type: Array of [Listener \(p. 98\)](#) objects

NextToken (p. 52)

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 \(p. 104\)](#).

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

InvalidNextTokenException

There isn't another item to return.

HTTP Status Code: 400

Example

List listeners

The following is an example for listing the listeners for an accelerator, and the response.

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

```
{
  "Listeners": [
    {
      "ListenerArn": "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh/listener/abcdef1234",
      "PortRanges": [
        {
          "FromPort": 80,
          "ToPort": 80
        }
      ],
      "Protocol": "TCP",
      "ClientAffinity": "NONE"
    }
  ]
}
```

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 V3](#)

ListTagsForResource

List all tags for an accelerator. To see an AWS CLI example of listing tags for an accelerator, scroll down to **Example**.

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

Request Syntax

```
{  
  "ResourceArn": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 102).

The request accepts the following data in JSON format.

ResourceArn (p. 55)

The Amazon Resource Name (ARN) of the accelerator to list tags for. An ARN uniquely identifies an accelerator.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1011.

Required: Yes

Response Syntax

```
{  
  "Tags": [  
    {  
      "Key": "string",  
      "Value": "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.

Tags (p. 55)

Root level tag for the Tags parameters.

Type: Array of [Tag](#) (p. 101) objects

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 104\)](#).

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

List tags for an accelerator

The following is an example for listing tags for an accelerator.

```
aws globalaccelerator list-tags-for-resource
  --accelerator-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefg
```

```
{
  "Tags": [
    {
      "Key": "Project",
      "Value": "A123456"
    }
  ]
}
```

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 V3](#)

ProvisionByoipCidr

Provisions an IP address range to use with your AWS resources through bring your own IP addresses (BYOIP) and creates a corresponding address pool. After the address range is provisioned, it is ready to be advertised using [AdvertiseByoipCidr](#).

To see an AWS CLI example of provisioning an address range for BYOIP, scroll down to **Example**.

For more information, see [Bring Your Own IP Addresses \(BYOIP\)](#) in the *AWS Global Accelerator Developer Guide*.

Request Syntax

```
{
  "Cidr": "string",
  "CidrAuthorizationContext": {
    "Message": "string",
    "Signature": "string"
  }
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 102).

The request accepts the following data in JSON format.

[Cidr](#) (p. 58)

The public IPv4 address range, in CIDR notation. The most specific IP prefix that you can specify is /24. The address range cannot overlap with another address range that you've brought to this or another Region.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

[CidrAuthorizationContext](#) (p. 58)

A signed document that proves that you are authorized to bring the specified IP address range to Amazon using BYOIP.

Type: [CidrAuthorizationContext](#) (p. 90) object

Required: Yes

Response Syntax

```
{
  "ByoipCidr": {
    "Cidr": "string",
    "Events": [
      {

```

```
    "Message": "string",  
    "Timestamp": number  
  },  
  ],  
  "State": "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.

ByoipCidr (p. 58)

Information about the address range.

Type: [ByoipCidr \(p. 87\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 104\)](#).

AccessDeniedException

You don't have access permission.

HTTP Status Code: 400

IncorrectCidrStateException

The CIDR that you specified is not valid for this action. For example, the state of the CIDR might be incorrect for this action.

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

Example

Provisioning an address range for BYOIP

The following is an example of provisioning an address range for BYOIP and the response.

For more information about creating the values for `text_message` and `signature`, see [Bring Your Own IP Addresses \(BYOIP\)](#) in the *AWS Global Accelerator Developer Guide*.

```
aws globalaccelerator provision-byoip-cidr
--cidr 203.0.113.25/24
--cidr-authorization-context Message="$text_message",Signature="$signed_message"
```

```
{
  "ByoipCidr": {
    "Cidr": "203.0.113.25/24",
    "State": "PENDING_PROVISIONING"
  }
}
```

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 V3](#)

TagResource

Add tags to an accelerator resource. To see an AWS CLI example of adding tags to an accelerator, scroll down to **Example**.

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

Request Syntax

```
{
  "ResourceArn": "string",
  "Tags": [
    {
      "Key": "string",
      "Value": "string"
    }
  ]
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 102\)](#).

The request accepts the following data in JSON format.

ResourceArn (p. 61)

The Amazon Resource Name (ARN) of the Global Accelerator resource to add tags to. An ARN uniquely identifies a resource.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1011.

Required: Yes

Tags (p. 61)

The tags to add to a resource. A tag consists of a key and a value that you define.

Type: Array of [Tag \(p. 101\)](#) objects

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 \(p. 104\)](#).

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

Add tags to an accelerator

The following is an example for adding tags to an accelerator. When successful, this command has no output.

Note

Instead of using command line arguments, you can provide the tags in a JSON file. Then, for example, with a file called `tags.json` in the current folder, specify `file://tags.json` with the `--tags` parameter.

```
aws globalaccelerator tag-resource
  --resource-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh
  --tags Key="Name",Value="Example Name" Key="Project",Value="Example Project"
```

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 V3](#)

UntagResource

Remove tags from a Global Accelerator resource. When you specify a tag key, the action removes both that key and its associated value. To see an AWS CLI example of removing tags from an accelerator, scroll down to **Example**. The operation succeeds even if you attempt to remove tags from an accelerator that was already removed.

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

Request Syntax

```
{  
  "ResourceArn": "string",  
  "TagKeys": [ "string" ]  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 102).

The request accepts the following data in JSON format.

ResourceArn (p. 63)

The Amazon Resource Name (ARN) of the Global Accelerator resource to remove tags from. An ARN uniquely identifies a resource.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1011.

Required: Yes

TagKeys (p. 63)

The tag key pairs that you want to remove from the specified resources.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Length Constraints: Minimum length of 1. Maximum length of 128.

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. 104).

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

Remove tags from an accelerator

The following is an example for removing tags from an accelerator. When successful, this command has no output.

```
aws globalaccelerator untag-resource
  --resource-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh
  --tag-keys Key="Name" Key="Project"
```

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 V3](#)

UpdateAccelerator

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

Important

Global Accelerator is a global service that supports endpoints in multiple AWS Regions but you must specify the US West (Oregon) Region to create or update accelerators.

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 \(p. 102\)](#).

The request accepts the following data in JSON format.

AcceleratorArn (p. 65)

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

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Enabled (p. 65)

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. 65)

The value for the address type must be IPv4.

Type: String

Valid Values: IPV4

Required: No

Name (p. 65)

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,
    "DnsName": "string",
    "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. 66)

Information about the updated accelerator.

Type: [Accelerator \(p. 84\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 104\)](#).

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 an accelerator

The following is an example for updating an accelerator to change the name.

```
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",
          "198.51.100.52"
        ],
        "IpFamily":"IPv4"
      }
    ],
    "DnsName":"a1234567890abcdef.awsglobalaccelerator.com",
    "CreatedTime":1232394847.0,
    "LastModifiedTime":1232395654.0
  }
}
```

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 V3](#)

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. 102).

The request accepts the following data in JSON format.

AcceleratorArn (p. 68)

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. 68)

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. 68)

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. 68)

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. If you specify slash (/) for the S3 bucket prefix, the log file bucket folder structure will include a double slash (/), like the following:

s3-bucket_name//AWSLogs/aws_account_id

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. 69)

Updated attributes for the accelerator.

Type: [AcceleratorAttributes](#) (p. 86) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#) (p. 104).

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

HTTP Status Code: 400

AccessDeniedException

You don't have access permission.

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

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 flowlogs-abc
  --flow-logs-s3-prefix bucketprefix-abc
  --region us-west-2
```

```
{
  "AcceleratorAttributes": {
    "FlowLogsEnabled": true
    "FlowLogsS3Bucket": flowlogs-abc
    "FlowLogsS3Prefix": bucketprefix-abc
  }
}
```

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 V3](#)

UpdateEndpointGroup

Update an endpoint group. A resource must be valid and active when you add it as an endpoint.

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

Request Syntax

```
{
  "EndpointConfigurations": [
    {
      "ClientIPPreservationEnabled": boolean,
      "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 \(p. 102\)](#).

The request accepts the following data in JSON format.

EndpointConfigurations (p. 71)

The list of endpoint objects. A resource must be valid and active when you add it as an endpoint.

Type: Array of [EndpointConfiguration \(p. 91\)](#) objects

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

Required: No

EndpointGroupArn (p. 71)

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

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

HealthCheckIntervalSeconds (p. 71)

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. 71)

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. 71)

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. 71)

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. 71)

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. 71)

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": [
      {
        "ClientIPPreservationEnabled": boolean,
        "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. 73)

The information about the endpoint group that was updated.

Type: [EndpointGroup](#) (p. 95) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#) (p. 104).

AccessDeniedException

You don't have access permission.

HTTP Status Code: 400

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 endpoints to an endpoint group, and the response.

```
aws globalaccelerator update-endpoint-group
  --endpoint-group-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh/listener/6789vxyz-vxyz-6789-vxyz-6789lmnopqrs/endpoint-group/
ab8888example
  --endpoint-configurations
    EndpointId=i-1234567890abcdef0,Weight=128
    EndpointId=arn:aws:elasticloadbalancing:us-east-1:000123456789:loadbalancer/app/
ALBTesting/alb01234567890xyz,Weight=128
    EndpointId=arn:aws:elasticloadbalancing:us-east-1:000123456789:loadbalancer/net/
NLBTesting/alb01234567890qrs,Weight=128
```

```
{
  "EndpointGroup": {
    "TrafficDialPercentage": 100.0,
    "EndpointDescriptions": [
      {
        "Weight": 128,
        "EndpointId": "i-1234567890abcdef0"
      },
      {
        "Weight": 128,
        "EndpointId": "arn:aws:elasticloadbalancing:us-
east-1:000123456789:loadbalancer/app/ALBTesting/alb01234567890xyz"
      },
      {
        "Weight": 128,
        "EndpointId": "arn:aws:elasticloadbalancing:us-
east-1:000123456789:loadbalancer/net/NLBTesting/alb01234567890qrs"
      }
    ],
    "EndpointGroupArn": "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh/listener/6789vxyz-vxyz-6789-vxyz-6789lmnopqrs/endpoint-
group/4321abcd-abcd-4321-abcd-4321abcdefg",
    "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 V3](#)

UpdateListener

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

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 \(p. 102\)](#).

The request accepts the following data in JSON format.

ClientAffinity (p. 76)

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. 76)

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

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

PortRanges (p. 76)

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

Type: Array of [PortRange \(p. 100\)](#) objects

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

Required: No

Protocol (p. 76)

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. 77)

Information for the updated listener.

Type: [Listener \(p. 98\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 104\)](#).

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

Example

Update a listener

The following is an example of updating the port range for a listener, and the response.

```
aws globalaccelerator update-listener
  --listener-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-abcd-1234-
abcd-1234abcdefg/1234abcd-1234abcdefg/listener/0123vxyz
  --port-ranges FromPort=100,ToPort=100
```

```
{
  "Listener": {
    "ListenerArn": "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefg/1234abcd-1234abcdefg/listener/0123vxyz",
    "PortRanges": [
      {
        "FromPort": 100,
        "ToPort": 100
      }
    ],
    "Protocol": "TCP",
    "ClientAffinity": "NONE"
  }
}
```

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 V3](#)

WithdrawByoipCidr

Stops advertising an address range that is provisioned as an address pool. You can perform this operation at most once every 10 seconds, even if you specify different address ranges each time. To see an AWS CLI example of withdrawing an address range for BYOIP so it will no longer be advertised by AWS, scroll down to **Example**.

It can take a few minutes before traffic to the specified addresses stops routing to AWS because of propagation delays.

For more information, see [Bring Your Own IP Addresses \(BYOIP\)](#) in the *AWS Global Accelerator Developer Guide*.

Request Syntax

```
{  
  "Cidr": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 102\)](#).

The request accepts the following data in JSON format.

Cidr (p. 80)

The address range, in CIDR notation.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Syntax

```
{  
  "ByoipCidr": {  
    "Cidr": "string",  
    "Events": [  
      {  
        "Message": "string",  
        "Timestamp": number  
      }  
    ],  
    "State": "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.

ByoipCidr (p. 80)

Information about the address pool.

Type: [ByoipCidr \(p. 87\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 104\)](#).

AccessDeniedException

You don't have access permission.

HTTP Status Code: 400

ByoipCidrNotFoundException

The CIDR that you specified was not found or is incorrect.

HTTP Status Code: 400

IncorrectCidrStateException

The CIDR that you specified is not valid for this action. For example, the state of the CIDR might be incorrect for this action.

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

Example

Withdrawing an address range from advertising by AWS

The following is an example of withdrawing an address range from advertising by AWS and the response.

```
aws globalaccelerator withdraw-byoip-cidr
--cidr 203.0.113.25/24
```

```
{
  "ByoipCidr": {
    "Cidr": "203.0.113.25/24",
    "State": "PENDING_WITHDRAWING"
  }
}
```


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 V3](#)

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. 84)
- [AcceleratorAttributes](#) (p. 86)
- [ByoipCidr](#) (p. 87)
- [ByoipCidrEvent](#) (p. 89)
- [CidrAuthorizationContext](#) (p. 90)
- [EndpointConfiguration](#) (p. 91)
- [EndpointDescription](#) (p. 93)
- [EndpointGroup](#) (p. 95)
- [IpSet](#) (p. 97)
- [Listener](#) (p. 98)
- [PortRange](#) (p. 100)
- [Tag](#) (p. 101)

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

DnsName

The Domain Name System (DNS) name that Global Accelerator creates that points to your accelerator's static IP addresses.

The naming convention for the DNS name is the following: A lowercase letter a, followed by a 16-bit random hex string, followed by .awsglobalaccelerator.com. For example: a1234567890abcdef.awsglobalaccelerator.com.

For more information about the default DNS name, see [Support for DNS Addressing in Global Accelerator](#) in the *AWS Global Accelerator Developer Guide*.

Type: String

Length Constraints: Maximum length of 255.

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

The static IP addresses that Global Accelerator associates with the accelerator.

Type: Array of [IpSet \(p. 97\)](#) 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 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 V3](#)

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. If you specify slash (/) for the S3 bucket prefix, the log file bucket folder structure will include a double slash (//), like the following:

```
s3-bucket_name//AWSLogs/aws_account_id
```

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 V3](#)

ByoipCidr

Information about an IP address range that is provisioned for use with your AWS resources through bring your own IP address (BYOIP).

The following describes each BYOIP state that your IP address range can be in.

- **PENDING_PROVISIONING** — You've submitted a request to provision an IP address range but it is not yet provisioned with AWS Global Accelerator.
- **READY** — The address range is provisioned with AWS Global Accelerator and can be advertised.
- **PENDING_ADVERTISING** — You've submitted a request for AWS Global Accelerator to advertise an address range but it is not yet being advertised.
- **ADVERTISING** — The address range is being advertised by AWS Global Accelerator.
- **PENDING_WITHDRAWING** — You've submitted a request to withdraw an address range from being advertised but it is still being advertised by AWS Global Accelerator.
- **PENDING_DEPROVISIONING** — You've submitted a request to deprovision an address range from AWS Global Accelerator but it is still provisioned.
- **DEPROVISIONED** — The address range is deprovisioned from AWS Global Accelerator.
- **FAILED_PROVISION** — The request to provision the address range from AWS Global Accelerator was not successful. Please make sure that you provide all of the correct information, and try again. If the request fails a second time, contact AWS support.
- **FAILED_ADVERTISING** — The request for AWS Global Accelerator to advertise the address range was not successful. Please make sure that you provide all of the correct information, and try again. If the request fails a second time, contact AWS support.
- **FAILED_WITHDRAW** — The request to withdraw the address range from advertising by AWS Global Accelerator was not successful. Please make sure that you provide all of the correct information, and try again. If the request fails a second time, contact AWS support.
- **FAILED_DEPROVISION** — The request to deprovision the address range from AWS Global Accelerator was not successful. Please make sure that you provide all of the correct information, and try again. If the request fails a second time, contact AWS support.

Contents

Cidr

The address range, in CIDR notation.

Type: String

Length Constraints: Maximum length of 255.

Required: No

Events

A history of status changes for an IP address range that you bring to AWS Global Accelerator through bring your own IP address (BYOIP).

Type: Array of [ByoipCidrEvent](#) (p. 89) objects

Required: No

State

The state of the address pool.

Type: String

Valid Values: PENDING_PROVISIONING | READY | PENDING_ADVERTISING | ADVERTISING
| PENDING_WITHDRAWING | PENDING_DEPROVISIONING | DEPROVISIONED |
FAILED_PROVISION | FAILED_ADVERTISING | FAILED_WITHDRAW | FAILED_DEPROVISION

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 V3](#)

ByoipCidrEvent

A complex type that contains a `Message` and a `Timestamp` value for changes that you make in the status an IP address range that you bring to AWS Global Accelerator through bring your own IP address (BYOIP).

Contents

Message

A string that contains an `Event` message describing changes that you make in the status of an IP address range that you bring to AWS Global Accelerator through bring your own IP address (BYOIP).

Type: String

Length Constraints: Maximum length of 255.

Required: No

Timestamp

A timestamp when you make a status change for an IP address range that you bring to AWS Global Accelerator through bring your own IP address (BYOIP).

Type: Timestamp

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 V3](#)

CidrAuthorizationContext

Provides authorization for Amazon to bring a specific IP address range to a specific AWS account using bring your own IP addresses (BYOIP).

For more information, see [Bring Your Own IP Addresses \(BYOIP\)](#) in the *AWS Global Accelerator Developer Guide*.

Contents

Message

The plain-text authorization message for the prefix and account.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Signature

The signed authorization message for the prefix and account.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

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 V3](#)

EndpointConfiguration

A complex type for endpoints. A resource must be valid and active when you add it as an endpoint.

Contents

ClientIPPreservationEnabled

Indicates whether client IP address preservation is enabled for an Application Load Balancer endpoint. The value is true or false. The default value is true for new accelerators.

If the value is set to true, the client's IP address is preserved in the `X-Forwarded-For` request header as traffic travels to applications on the Application Load Balancer endpoint fronted by the accelerator.

For more information, see [Preserve Client IP Addresses in AWS Global Accelerator](#) in the *AWS Global Accelerator Developer Guide*.

Type: Boolean

Required: No

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. For Amazon EC2 instances, this is the EC2 instance ID. A resource must be valid and active when you add it as an endpoint.

An Application Load Balancer can be either internal or internet-facing.

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 V3](#)

EndpointDescription

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

Contents

ClientIPPreservationEnabled

Indicates whether client IP address preservation is enabled for an Application Load Balancer endpoint. The value is true or false. The default value is true for new accelerators.

If the value is set to true, the client's IP address is preserved in the `X-Forwarded-For` request header as traffic travels to applications on the Application Load Balancer endpoint fronted by the accelerator.

For more information, see [Viewing Client IP Addresses in AWS Global Accelerator](#) in the *AWS Global Accelerator Developer Guide*.

Type: Boolean

Required: No

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. For EC2 instances, this is the EC2 instance ID.

An Application Load Balancer can be either internal or internet-facing.

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 V3](#)

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 \(p. 93\)](#) 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 where the endpoint group is located.

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 V3](#)

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 V3](#)

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 \(p. 100\)](#) 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 V3](#)

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 V3](#)

Tag

A complex type that contains a `Tag` key and `Tag` value.

Contents

Key

A string that contains a `Tag` key.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 128.

Required: Yes

Value

A string that contains a `Tag` value.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Required: Yes

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 V3](#)

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