



API Reference

AWS Global Accelerator



API Version 2018-08-08

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

AWS Global Accelerator: API Reference

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

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

Table of Contents

Welcome	1
Actions	3
AddCustomRoutingEndpoints	5
Request Syntax	5
Request Parameters	5
Response Syntax	6
Response Elements	6
Errors	7
See Also	8
AddEndpoints	9
Request Syntax	9
Request Parameters	9
Response Syntax	10
Response Elements	10
Errors	11
See Also	12
AdvertiseByoipCidr	13
Request Syntax	13
Request Parameters	13
Response Syntax	13
Response Elements	14
Errors	14
Examples	15
See Also	15
AllowCustomRoutingTraffic	17
Request Syntax	17
Request Parameters	17
Response Elements	19
Errors	19
Examples	19
See Also	20
CreateAccelerator	21
Request Syntax	21
Request Parameters	21

Response Syntax	23
Response Elements	24
Errors	24
Examples	25
See Also	26
CreateCrossAccountAttachment	27
Request Syntax	27
Request Parameters	27
Response Syntax	29
Response Elements	29
Errors	29
See Also	30
CreateCustomRoutingAccelerator	32
Request Syntax	32
Request Parameters	32
Response Syntax	34
Response Elements	35
Errors	35
Examples	36
See Also	37
CreateCustomRoutingEndpointGroup	38
Request Syntax	38
Request Parameters	38
Response Syntax	39
Response Elements	40
Errors	40
Examples	41
See Also	42
CreateCustomRoutingListener	43
Request Syntax	43
Request Parameters	43
Response Syntax	44
Response Elements	44
Errors	44
Examples	45
See Also	46

CreateEndpointGroup	47
Request Syntax	47
Request Parameters	47
Response Syntax	50
Response Elements	51
Errors	51
Examples	52
See Also	53
CreateListener	55
Request Syntax	55
Request Parameters	55
Response Syntax	57
Response Elements	57
Errors	57
Examples	58
See Also	59
DeleteAccelerator	60
Request Syntax	60
Request Parameters	60
Response Elements	61
Errors	61
See Also	62
DeleteCrossAccountAttachment	63
Request Syntax	63
Request Parameters	63
Response Elements	64
Errors	64
See Also	64
DeleteCustomRoutingAccelerator	66
Request Syntax	66
Request Parameters	66
Response Elements	67
Errors	67
See Also	67
DeleteCustomRoutingEndpointGroup	69
Request Syntax	69

Request Parameters	69
Response Elements	69
Errors	69
See Also	70
DeleteCustomRoutingListener	71
Request Syntax	71
Request Parameters	71
Response Elements	71
Errors	71
See Also	72
DeleteEndpointGroup	73
Request Syntax	73
Request Parameters	73
Response Elements	73
Errors	73
See Also	74
DeleteListener	75
Request Syntax	75
Request Parameters	75
Response Elements	75
Errors	75
See Also	76
DenyCustomRoutingTraffic	77
Request Syntax	77
Request Parameters	77
Response Elements	79
Errors	79
Examples	79
See Also	80
DeprovisionByoipCidr	81
Request Syntax	81
Request Parameters	81
Response Syntax	81
Response Elements	82
Errors	82
Examples	83

See Also	83
DescribeAccelerator	85
Request Syntax	85
Request Parameters	85
Response Syntax	85
Response Elements	86
Errors	86
Examples	87
See Also	88
DescribeAcceleratorAttributes	89
Request Syntax	89
Request Parameters	89
Response Syntax	89
Response Elements	89
Errors	90
Examples	90
See Also	91
DescribeCrossAccountAttachment	92
Request Syntax	92
Request Parameters	92
Response Syntax	92
Response Elements	93
Errors	93
See Also	94
DescribeCustomRoutingAccelerator	95
Request Syntax	95
Request Parameters	95
Response Syntax	95
Response Elements	96
Errors	96
Examples	97
See Also	97
DescribeCustomRoutingAcceleratorAttributes	99
Request Syntax	99
Request Parameters	99
Response Syntax	99

Response Elements	99
Errors	100
Examples	100
See Also	101
DescribeCustomRoutingEndpointGroup	102
Request Syntax	102
Request Parameters	102
Response Syntax	102
Response Elements	103
Errors	103
Examples	104
See Also	104
DescribeCustomRoutingListener	106
Request Syntax	106
Request Parameters	106
Response Syntax	106
Response Elements	107
Errors	107
See Also	107
DescribeEndpointGroup	109
Request Syntax	109
Request Parameters	109
Response Syntax	109
Response Elements	110
Errors	110
Examples	111
See Also	112
DescribeListener	113
Request Syntax	113
Request Parameters	113
Response Syntax	113
Response Elements	114
Errors	114
Examples	114
See Also	115
ListAccelerators	116

Request Syntax	116
Request Parameters	116
Response Syntax	116
Response Elements	117
Errors	118
Examples	118
See Also	119
ListByoipCidrs	121
Request Syntax	121
Request Parameters	121
Response Syntax	121
Response Elements	122
Errors	122
Examples	123
See Also	125
ListCrossAccountAttachments	126
Request Syntax	126
Request Parameters	126
Response Syntax	126
Response Elements	127
Errors	127
See Also	128
ListCrossAccountResourceAccounts	129
Response Syntax	129
Response Elements	129
Errors	129
See Also	130
ListCrossAccountResources	131
Request Syntax	131
Request Parameters	131
Response Syntax	132
Response Elements	132
Errors	133
See Also	133
ListCustomRoutingAccelerators	135
Request Syntax	135

Request Parameters	135
Response Syntax	135
Response Elements	136
Errors	136
Examples	137
See Also	138
ListCustomRoutingEndpointGroups	140
Request Syntax	140
Request Parameters	140
Response Syntax	141
Response Elements	141
Errors	142
Examples	142
See Also	143
ListCustomRoutingListeners	144
Request Syntax	144
Request Parameters	144
Response Syntax	145
Response Elements	145
Errors	145
Examples	146
See Also	147
ListCustomRoutingPortMappings	148
Request Syntax	148
Request Parameters	148
Response Syntax	149
Response Elements	150
Errors	150
See Also	151
ListCustomRoutingPortMappingsByDestination	152
Request Syntax	152
Request Parameters	152
Response Syntax	153
Response Elements	154
Errors	154
Examples	155

See Also	156
ListEndpointGroups	157
Request Syntax	157
Request Parameters	157
Response Syntax	158
Response Elements	158
Errors	159
Examples	160
See Also	161
ListListeners	162
Request Syntax	162
Request Parameters	162
Response Syntax	163
Response Elements	163
Errors	164
Examples	164
See Also	165
ListTagsForResource	166
Request Syntax	166
Request Parameters	166
Response Syntax	166
Response Elements	167
Errors	167
Examples	167
See Also	168
ProvisionByoipCidr	169
Request Syntax	169
Request Parameters	169
Response Syntax	170
Response Elements	170
Errors	170
Examples	171
See Also	172
RemoveCustomRoutingEndpoints	173
Request Syntax	173
Request Parameters	173

Response Elements	173
Errors	174
See Also	174
RemoveEndpoints	176
Request Syntax	176
Request Parameters	176
Response Elements	177
Errors	177
See Also	178
TagResource	179
Request Syntax	179
Request Parameters	179
Response Elements	180
Errors	180
Examples	180
See Also	181
UntagResource	182
Request Syntax	182
Request Parameters	182
Response Elements	183
Errors	183
Examples	183
See Also	184
UpdateAccelerator	185
Request Syntax	185
Request Parameters	185
Response Syntax	187
Response Elements	187
Errors	188
Examples	188
See Also	189
UpdateAcceleratorAttributes	190
Request Syntax	190
Request Parameters	190
Response Syntax	191
Response Elements	191

Errors	192
Examples	192
See Also	193
UpdateCrossAccountAttachment	194
Request Syntax	194
Request Parameters	195
Response Syntax	196
Response Elements	197
Errors	197
See Also	198
UpdateCustomRoutingAccelerator	199
Request Syntax	199
Request Parameters	199
Response Syntax	200
Response Elements	200
Errors	201
Examples	201
See Also	202
UpdateCustomRoutingAcceleratorAttributes	204
Request Syntax	204
Request Parameters	204
Response Syntax	205
Response Elements	205
Errors	206
Examples	206
See Also	207
UpdateCustomRoutingListener	208
Request Syntax	208
Request Parameters	208
Response Syntax	209
Response Elements	209
Errors	209
Examples	210
See Also	211
UpdateEndpointGroup	212
Request Syntax	212

Request Parameters	212
Response Syntax	215
Response Elements	216
Errors	216
Examples	217
See Also	218
UpdateListener	219
Request Syntax	219
Request Parameters	219
Response Syntax	220
Response Elements	221
Errors	221
Examples	222
See Also	222
WithdrawByoipCidr	224
Request Syntax	224
Request Parameters	224
Response Syntax	224
Response Elements	225
Errors	225
Examples	226
See Also	226
Data Types	228
Accelerator	230
Contents	230
See Also	233
AcceleratorAttributes	234
Contents	234
See Also	235
AcceleratorEvent	236
Contents	236
See Also	236
Attachment	237
Contents	237
See Also	238
ByoipCidr	239

Contents	240
See Also	240
ByoipCidrEvent	242
Contents	242
See Also	242
CidrAuthorizationContext	243
Contents	243
See Also	243
CrossAccountResource	244
Contents	244
See Also	244
CustomRoutingAccelerator	246
Contents	246
See Also	248
CustomRoutingAcceleratorAttributes	249
Contents	249
See Also	250
CustomRoutingDestinationConfiguration	251
Contents	251
See Also	252
CustomRoutingDestinationDescription	253
Contents	253
See Also	253
CustomRoutingEndpointConfiguration	255
Contents	255
See Also	255
CustomRoutingEndpointDescription	256
Contents	256
See Also	256
CustomRoutingEndpointGroup	257
Contents	257
See Also	258
CustomRoutingListener	259
Contents	259
See Also	259
DestinationPortMapping	260

Contents	260
See Also	261
EndpointConfiguration	263
Contents	263
See Also	264
EndpointDescription	265
Contents	265
See Also	266
EndpointGroup	267
Contents	267
See Also	269
EndpointIdentifier	271
Contents	271
See Also	271
IpSet	273
Contents	273
See Also	273
Listener	275
Contents	275
See Also	276
PortMapping	277
Contents	277
See Also	278
PortOverride	279
Contents	279
See Also	279
PortRange	281
Contents	281
See Also	281
Resource	282
Contents	282
See Also	282
SocketAddress	283
Contents	283
See Also	283
Tag	284

Contents	284
See Also	284
Common Parameters	285
Common Errors	288

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 the performance of your applications for local and global users. Depending on the type of accelerator you choose, you can gain additional benefits.

- By using a standard accelerator, you can improve availability of your internet applications that are used by a global audience. With a standard accelerator, Global Accelerator directs traffic to optimal endpoints over the AWS global network.
- For other scenarios, you might choose a custom routing accelerator. With a custom routing accelerator, you can use application logic to directly map one or more users to a specific endpoint among many endpoints.

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, update, or otherwise work with accelerators. That is, for example, specify `--region us-west-2` on AWS CLI commands.

By default, Global Accelerator provides you with static IP addresses that you associate with your accelerator. The static IP addresses are anycast from the AWS edge network. For IPv4, Global Accelerator provides two static IPv4 addresses. For dual-stack, Global Accelerator provides a total of four addresses: two static IPv4 addresses and two static IPv6 addresses. With a standard accelerator for IPv4, instead of using the 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 (BYOIP).

For a standard accelerator, they distribute incoming application traffic across multiple endpoint resources in multiple AWS Regions , which increases the availability of your applications. Endpoints for standard accelerators can be Network Load Balancers, Application Load Balancers, Amazon EC2 instances, or Elastic IP addresses that are located in one AWS Region or multiple AWS Regions.

For custom routing accelerators, you map traffic that arrives to the static IP addresses to specific Amazon EC2 servers in endpoints that are virtual private cloud (VPC) subnets.

 **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 like tag-based permissions with Global Accelerator to limit the users who have permissions to delete an accelerator. For more information, see [Tag-based policies](#).

For standard accelerators, 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 always directed to healthy endpoints.

For more information about understanding and using Global Accelerator, see the [AWS Global Accelerator Developer Guide](#).

This document was last published on March 20, 2024.

Actions

The following actions are supported:

- [AddCustomRoutingEndpoints](#)
- [AddEndpoints](#)
- [AdvertiseByoipCidr](#)
- [AllowCustomRoutingTraffic](#)
- [CreateAccelerator](#)
- [CreateCrossAccountAttachment](#)
- [CreateCustomRoutingAccelerator](#)
- [CreateCustomRoutingEndpointGroup](#)
- [CreateCustomRoutingListener](#)
- [CreateEndpointGroup](#)
- [CreateListener](#)
- [DeleteAccelerator](#)
- [DeleteCrossAccountAttachment](#)
- [DeleteCustomRoutingAccelerator](#)
- [DeleteCustomRoutingEndpointGroup](#)
- [DeleteCustomRoutingListener](#)
- [DeleteEndpointGroup](#)
- [DeleteListener](#)
- [DenyCustomRoutingTraffic](#)
- [DeprovisionByoipCidr](#)
- [DescribeAccelerator](#)
- [DescribeAcceleratorAttributes](#)
- [DescribeCrossAccountAttachment](#)
- [DescribeCustomRoutingAccelerator](#)
- [DescribeCustomRoutingAcceleratorAttributes](#)
- [DescribeCustomRoutingEndpointGroup](#)
- [DescribeCustomRoutingListener](#)

- [DescribeEndpointGroup](#)
- [DescribeListener](#)
- [ListAccelerators](#)
- [ListByoipCidrs](#)
- [ListCrossAccountAttachments](#)
- [ListCrossAccountResourceAccounts](#)
- [ListCrossAccountResources](#)
- [ListCustomRoutingAccelerators](#)
- [ListCustomRoutingEndpointGroups](#)
- [ListCustomRoutingListeners](#)
- [ListCustomRoutingPortMappings](#)
- [ListCustomRoutingPortMappingsByDestination](#)
- [ListEndpointGroups](#)
- [ListListeners](#)
- [ListTagsForResource](#)
- [ProvisionByoipCidr](#)
- [RemoveCustomRoutingEndpoints](#)
- [RemoveEndpoints](#)
- [TagResource](#)
- [UntagResource](#)
- [UpdateAccelerator](#)
- [UpdateAcceleratorAttributes](#)
- [UpdateCrossAccountAttachment](#)
- [UpdateCustomRoutingAccelerator](#)
- [UpdateCustomRoutingAcceleratorAttributes](#)
- [UpdateCustomRoutingListener](#)
- [UpdateEndpointGroup](#)
- [UpdateListener](#)
- [WithdrawByoipCidr](#)

AddCustomRoutingEndpoints

Associate a virtual private cloud (VPC) subnet endpoint with your custom routing accelerator.

The listener port range must be large enough to support the number of IP addresses that can be specified in your subnet. The number of ports required is: subnet size times the number of ports per destination EC2 instances. For example, a subnet defined as /24 requires a listener port range of at least 255 ports.

Note: You must have enough remaining listener ports available to map to the subnet ports, or the call will fail with a `LimitExceeded` exception.

By default, all destinations in a subnet in a custom routing accelerator cannot receive traffic. To enable all destinations to receive traffic, or to specify individual port mappings that can receive traffic, see the [AllowCustomRoutingTraffic](#) operation.

Request Syntax

```
{  
    "EndpointConfigurations": [  
        {  
            "AttachmentArn": "string",  
            "EndpointId": "string"  
        }  
    ],  
    "EndpointGroupArn": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

[EndpointConfigurations](#)

The list of endpoint objects to add to a custom routing accelerator.

Type: Array of [CustomRoutingEndpointConfiguration](#) objects

Array Members: Minimum number of 1 item. Maximum number of 20 items.

Required: Yes

[EndpointGroupArn](#)

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

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Syntax

```
{  
    "EndpointDescriptions": [  
        {  
            "EndpointId": "string"  
        }  
    ],  
    "EndpointGroupArn": "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.

[EndpointDescriptions](#)

The endpoint objects added to the custom routing accelerator.

Type: Array of [CustomRoutingEndpointDescription](#) objects

[EndpointGroupArn](#)

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

Type: String

Length Constraints: Maximum length of 255.

Errors

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

AccessDeniedException

You don't have access permission.

HTTP Status Code: 400

ConflictException

You can't use both of those options.

HTTP Status Code: 400

EndpointAlreadyExistsException

The endpoint that you specified doesn't exist.

HTTP Status Code: 400

EndpointGroupNotFoundException

The endpoint group that you specified doesn't exist.

HTTP Status Code: 400

InternalServiceErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

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

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

AddEndpoints

Add endpoints to an endpoint group. The `AddEndpoints` API operation is the recommended option for adding endpoints. The alternative options are to add endpoints when you create an endpoint group (with the [CreateEndpointGroup](#) API) or when you update an endpoint group (with the [UpdateEndpointGroup](#) API).

There are two advantages to using `AddEndpoints` to add endpoints in Global Accelerator:

- It's faster, because Global Accelerator only has to resolve the new endpoints that you're adding, rather than resolving new and existing endpoints.
- It's more convenient, because you don't need to specify the current endpoints that are already in the endpoint group, in addition to the new endpoints that you want to add.

For information about endpoint types and requirements for endpoints that you can add to Global Accelerator, see [Endpoints for standard accelerators](#) in the *AWS Global Accelerator Developer Guide*.

Request Syntax

```
{  
    "EndpointConfigurations": [  
        {  
            "AttachmentArn": "string",  
            "ClientIPPreservationEnabled": boolean,  
            "EndpointId": "string",  
            "Weight": number  
        }  
    ],  
    "EndpointGroupArn": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

EndpointConfigurations

The list of endpoint objects.

Type: Array of [EndpointConfiguration](#) objects

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

Required: Yes

EndpointGroupArn

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

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Syntax

```
{  
    "EndpointDescriptions": [  
        {  
            "ClientIPPreservationEnabled": boolean,  
            "EndpointId": "string",  
            "HealthReason": "string",  
            "HealthState": "string",  
            "Weight": number  
        }  
    ],  
    "EndpointGroupArn": "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.

EndpointDescriptions

The list of endpoint objects.

Type: Array of [EndpointDescription](#) objects

[**EndpointGroupArn**](#)

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

Type: String

Length Constraints: Maximum length of 255.

Errors

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

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

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

TransactionInProgressException

There's already a transaction in progress. Another transaction can't be processed.

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

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 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](#).

The request accepts the following data in JSON format.

[Cidr](#)

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](#)

Information about the address range.

Type: [ByoipCidr](#) object

Errors

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

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

Examples

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 V2](#)
- [AWS SDK for JavaScript V3](#)

- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

AllowCustomRoutingTraffic

Specify the Amazon EC2 instance (destination) IP addresses and ports for a VPC subnet endpoint that can receive traffic for a custom routing accelerator. You can allow traffic to all destinations in the subnet endpoint, or allow traffic to a specified list of destination IP addresses and ports in the subnet. Note that you cannot specify IP addresses or ports outside of the range that you configured for the endpoint group.

After you make changes, you can verify that the updates are complete by checking the status of your accelerator: the status changes from IN_PROGRESS to DEPLOYED.

Request Syntax

```
{  
    "AllowAllTrafficToEndpoint": boolean,  
    "DestinationAddresses": [ "string" ],  
    "DestinationPorts": [ number ],  
    "EndpointGroupArn": "string",  
    "EndpointId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

[AllowAllTrafficToEndpoint](#)

Indicates whether all destination IP addresses and ports for a specified VPC subnet endpoint can receive traffic from a custom routing accelerator. The value is TRUE or FALSE.

When set to TRUE, *all* destinations in the custom routing VPC subnet can receive traffic. Note that you cannot specify destination IP addresses and ports when the value is set to TRUE.

When set to FALSE (or not specified), you *must* specify a list of destination IP addresses that are allowed to receive traffic. A list of ports is optional. If you don't specify a list of ports, the ports that can accept traffic is the same as the ports configured for the endpoint group.

The default value is FALSE.

Type: Boolean

Required: No

DestinationAddresses

A list of specific Amazon EC2 instance IP addresses (destination addresses) in a subnet that you want to allow to receive traffic. The IP addresses must be a subset of the IP addresses that you specified for the endpoint group.

DestinationAddresses is required if AllowAllTrafficToEndpoint is FALSE or is not specified.

Type: Array of strings

Array Members: Maximum number of 100 items.

Length Constraints: Maximum length of 45.

Required: No

DestinationPorts

A list of specific Amazon EC2 instance ports (destination ports) that you want to allow to receive traffic.

Type: Array of integers

Array Members: Maximum number of 100 items.

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

Required: No

EndpointGroupArn

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

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

EndpointId

An ID for the endpoint. For custom routing accelerators, this is the virtual private cloud (VPC) subnet ID.

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](#).

EndpointGroupNotFoundException

The endpoint group that you specified doesn't exist.

HTTP Status Code: 400

InternalServiceErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

Examples

Allow destination EC2 instances to receive traffic

The following is an example for specifying all EC2 instances in a subnet endpoint to receive traffic for a custom routing accelerator.

```
aws --region us-west-2 globalaccelerator allow-custom-routing-traffic --endpoint-group-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-abcd-1234-abcd-1234abcdefgh/listener/0123vxyz/endpoint-group/ab88888example --endpoint-id subnet-abcd123example --allow-all
```

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [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.

⚠ 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, update, or otherwise work with accelerators. That is, for example, specify `--region us-west-2` on AWS CLI commands.

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](#).

The request accepts the following data in JSON format.

Enabled

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](#)

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](#)

Optionally, if you've added your own IP address pool to Global Accelerator (BYOIP), you can choose an IPv4 address from your own pool to use for the accelerator's static IPv4 address when you create an accelerator.

After you bring an address range to AWS, it appears in your account as an address pool. When you create an accelerator, you can assign one IPv4 address from your range to it. Global Accelerator assigns you a second static IPv4 address from an Amazon IP address range. If you bring two IPv4 address ranges to AWS, you can assign one IPv4 address from each range to your accelerator. This restriction is because Global Accelerator assigns each address range to a different network zone, for high availability.

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

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.

Length Constraints: Maximum length of 45.

Required: No

IpAddressType

The IP address type that an accelerator supports. For a standard accelerator, the value can be IPV4 or DUAL_STACK.

Type: String

Valid Values: IPV4 | DUAL_STACK

Required: No

Name

The name of the accelerator. The name can have a maximum of 64 characters, must contain only alphanumeric characters, periods (.), or hyphens (-), and must not begin or end with a hyphen or period.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Tags

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](#) objects

Required: No

Response Syntax

```
{  
  "Accelerator": {  
    "AcceleratorArn": "string",  
    "CreatedTime": number,  
    "DnsName": "string",  
    "DualStackDnsName": "string",  
    "Name": "string",  
    "State": "string",  
    "Type": "string"  
  }  
}
```

```
"Enabled": boolean,
"Events": [
    {
        "Message": string,
        "Timestamp": number
    }
],
"IpAddressType": string,
"IpSets": [
    {
        "IpAddresses": [ string ],
        "IpAddressFamily": 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](#)

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

Type: [Accelerator](#) object

Errors

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

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

Examples

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-1234abcdefg",
    "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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

CreateCrossAccountAttachment

Create a cross-account attachment in AWS Global Accelerator. You create a cross-account attachment to specify the *principals* who have permission to add to accelerators in their own account the resources in your account that you also list in the attachment.

A principal can be an AWS account number or the Amazon Resource Name (ARN) for an accelerator. For account numbers that are listed as principals, to add a resource listed in the attachment to an accelerator, you must sign in to an account specified as a principal. Then you can add the resources that are listed to any of your accelerators. If an accelerator ARN is listed in the cross-account attachment as a principal, anyone with permission to make updates to the accelerator can add as endpoints resources that are listed in the attachment.

For more information, see [Working with cross-account attachments and endpoints in AWS Global Accelerator](#) in the *AWS Global Accelerator Developer Guide*.

Request Syntax

```
{  
    "IdempotencyToken": "string",  
    "Name": "string",  
    "Principals": [ "string" ],  
    "Resources": [  
        {  
            "EndpointId": "string",  
            "Region": "string"  
        }  
    ],  
    "Tags": [  
        {  
            "Key": "string",  
            "Value": "string"  
        }  
    ]  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

[IdempotencyToken](#)

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

[Name](#)

The name of the cross-account attachment.

Type: String

Length Constraints: Maximum length of 64.

Pattern: [\S\s]+

Required: Yes

[Principals](#)

The principals to list in the cross-account attachment. A principal can be an AWS account number or the Amazon Resource Name (ARN) for an accelerator.

Type: Array of strings

Length Constraints: Maximum length of 256.

Pattern: (^d{12}\$|arn:.*)

Required: No

[Resources](#)

The Amazon Resource Names (ARNs) for the resources to list in the cross-account attachment. A resource can be any supported AWS resource type for Global Accelerator.

Type: Array of [Resource](#) objects

Required: No

[Tags](#)

Create tags for cross-account attachment.

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

Type: Array of [Tag](#) objects

Required: No

Response Syntax

```
{  
  "CrossAccountAttachment": [  
    {  
      "AttachmentArn": "string",  
      "CreatedTime": number,  
      "LastModifiedTime": number,  
      "Name": "string",  
      "Principals": [ "string" ],  
      "Resources": [  
        {  
          "EndpointId": "string",  
          "Region": "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.

[CrossAccountAttachment](#)

Information about the cross-account attachment.

Type: [Attachment](#) object

Errors

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

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

LimitExceededException

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

HTTP Status Code: 400

TransactionInProgressException

There's already a transaction in progress. Another transaction can't be processed.

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)

- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

CreateCustomRoutingAccelerator

Create a custom routing accelerator. A custom routing accelerator directs traffic to one of possibly thousands of Amazon EC2 instance destinations running in a single or multiple virtual private clouds (VPC) subnet endpoints.

Be aware that, by default, all destination EC2 instances in a VPC subnet endpoint cannot receive traffic. To enable all destinations to receive traffic, or to specify individual port mappings that can receive traffic, see the [AllowCustomRoutingTraffic](#) operation.

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, update, or otherwise work with accelerators. That is, for example, specify `--region us-west-2` on AWS CLI commands.

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](#).

The request accepts the following data in JSON format.

Enabled

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

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

IpAddresses

Optionally, if you've added your own IP address pool to Global Accelerator (BYOIP), you can choose an IPv4 address from your own pool to use for the accelerator's static IPv4 address when you create an accelerator.

After you bring an address range to AWS, it appears in your account as an address pool. When you create an accelerator, you can assign one IPv4 address from your range to it. Global Accelerator assigns you a second static IPv4 address from an Amazon IP address range. If you bring two IPv4 address ranges to AWS, you can assign one IPv4 address from each range to your accelerator. This restriction is because Global Accelerator assigns each address range to a different network zone, for high availability.

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

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.

Length Constraints: Maximum length of 45.

Required: No

IpAddressType

The IP address type that an accelerator supports. For a custom routing accelerator, the value must be IPV4.

Type: String

Valid Values: IPV4 | DUAL_STACK

Required: No

Name

The name of a custom routing accelerator. The name can have a maximum of 64 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

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](#) objects

Required: No

Response Syntax

```
{  
  "Accelerator": {
```

```
"AcceleratorArn": "string",
"CreatedTime": number,
"DnsName": "string",
"Enabled": boolean,
"IpAddressType": "string",
"IpSets": [
  {
    "IpAddresses": [ "string" ],
    "IpAddressFamily": "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](#)

The accelerator that is created.

Type: [CustomRoutingAccelerator](#) object

Errors

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

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

LimitExceededException

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

HTTP Status Code: 400

Examples

Create a custom routing accelerator

The following is an example of creating a custom routing accelerator with two tags, and the response (which does not include the tag information).

```
aws globalaccelerator create-custom-routing-accelerator
  --name ExampleCustomRoutingAccelerator
  --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-1234abcdefg",
    "IpAddressType": "IPV4",
    "Name": "ExampleCustomRoutingAccelerator",
    "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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

CreateCustomRoutingEndpointGroup

Create an endpoint group for the specified listener for a custom routing accelerator. An endpoint group is a collection of endpoints in one AWS Region.

Request Syntax

```
{  
    "DestinationConfigurations": [  
        {  
            "FromPort": number,  
            "Protocols": [ "string" ],  
            "ToPort": number  
        }  
    ],  
    "EndpointGroupRegion": "string",  
    "IdempotencyToken": "string",  
    "ListenerArn": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

DestinationConfigurations

Sets the port range and protocol for all endpoints (virtual private cloud subnets) in a custom routing endpoint group to accept client traffic on.

Type: Array of [CustomRoutingDestinationConfiguration](#) objects

Array Members: Minimum number of 1 item. Maximum number of 100 items.

Required: Yes

EndpointGroupRegion

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

[IdempotencyToken](#)

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](#)

The Amazon Resource Name (ARN) of the listener for a custom routing endpoint.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Syntax

```
{  
  "EndpointGroup": {  
    "DestinationDescriptions": [  
      {  
        "FromPortnumber,  
        "Protocols": \[ "string" \],  
        "ToPort": number  
      }  
    \],  
    "EndpointDescriptions": \[  
      {  
        "EndpointId": "string"  
      }  
    \],  
  },  
}
```

```
    "EndpointGroupArn": "string",
    "EndpointGroupRegion": "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.

[EndpointGroup](#)

The information about the endpoint group created for a custom routing accelerator.

Type: [CustomRoutingEndpointGroup](#) object

Errors

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

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

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

Examples

Create an endpoint group for a custom routing accelerator

The following is an example of creating an endpoint group for a custom routing accelerator, and the response.

```
aws globalaccelerator create-custom-routing-endpoint-group
    --listener-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefg/listener/0123vxyz
    --endpoint-group-region us-east-1
    --endpoint-configurations EndpointId=i-1234567890abcdef0,Weight=128
    --region us-west-2
```

```
{
    "EndpointGroup": {
```

```
        "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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

CreateCustomRoutingListener

Create a listener to process inbound connections from clients to a custom routing accelerator. Connections arrive to assigned static IP addresses on the port range that you specify.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

AcceleratorArn

The Amazon Resource Name (ARN) of the accelerator for a custom routing listener.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

IdempotencyToken

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

The port range to support for connections from clients to your accelerator.

Separately, you set port ranges for endpoints. For more information, see [About endpoints for custom routing accelerators](#).

Type: Array of [PortRange](#) objects

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

Required: Yes

Response Syntax

```
{  
  "Listener": {  
    "ListenerArn": "string",  
    "PortRanges": [  
      {  
        "FromPort": number,  
        "ToPort": 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.

Listener

The listener that you've created for a custom routing accelerator.

Type: [CustomRoutingListener](#) object

Errors

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

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

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

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

Examples

Create a listener for a custom routing accelerator

The following is an example of creating a listener for a custom routing accelerator, and the response.

```
aws --region us-west-2 globalaccelerator create-custom-routing-listener  
    --accelerator-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-  
    abcd-1234-abcd-1234abcdefgh  
    --port-ranges FromPort=80,ToPort=80 FromPort=81,ToPort=81
```

```
{  
    "Listener": {  
        "PortRange": [  
            "FromPort": 5000  
            "ToPort": 10000,  
        ],  
        "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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [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.

For more information about endpoint types and requirements for endpoints that you can add to Global Accelerator, see [Endpoints for standard accelerators](#) in the *AWS Global Accelerator Developer Guide*.

Request Syntax

```
{
  "EndpointConfigurations": [
    {
      "AttachmentArn": "string",
      "ClientIPPreservationEnabled": boolean,
      "EndpointId": "string",
      "Weight": number
    }
  ],
  "EndpointGroupRegion": "string",
  "HealthCheckIntervalSeconds": number,
  "HealthCheckPath": "string",
  "HealthCheckPort": number,
  "HealthCheckProtocol": "string",
  "IdempotencyToken": "string",
  "ListenerArn": "string",
  "PortOverrides": [
    {
      "EndpointPort": number,
      "ListenerPort": number
    }
  ],
  "ThresholdCount": number,
  "TrafficDialPercentage": number
}
```

Request Parameters

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

The request accepts the following data in JSON format.

EndpointConfigurations

The list of endpoint objects.

Type: Array of [EndpointConfiguration](#) objects

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

Required: No

EndpointGroupRegion

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

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

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.

Pattern: ^/[-a-zA-Z0-9@:_\\+.^#?&/=]*\$

Required: No

HealthCheckPort

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

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

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

The Amazon Resource Name (ARN) of the listener.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

PortOverrides

Override specific listener ports used to route traffic to endpoints that are part of this endpoint group. For example, you can create a port override in which the listener receives user traffic on

ports 80 and 443, but your accelerator routes that traffic to ports 1080 and 1443, respectively, on the endpoints.

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

Type: Array of [PortOverride](#) objects

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

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

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",  
    "PortOverrides": [  
        {  
            "EndpointPort": number,  
            "ListenerPort": number  
        }  
    ],  
    "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

The information about the endpoint group that was created.

Type: [EndpointGroup object](#)

Errors

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

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

Examples

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
    --port-overrides ListenerPort=443,EndpointPort=1443
    --endpoint-configurations EndpointId=i-1234567890abcdef0,Weight=128
    --region us-west-2
```

```
{
    "EndpointGroup": {
        "TrafficDialPercentage": 100.0,
        "EndpointDescriptions": [
            {
                "Weight": 128,
                "EndpointId": "i-1234567890abcdef0"
            }
        ],
        "PortOverrides": [
            {
                "EndpointPort": 1443,
                "ListenerPort": 443
            }
        ],
        "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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [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.

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](#).

The request accepts the following data in JSON format.

AcceleratorArn

The Amazon Resource Name (ARN) of your accelerator.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

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

[IdempotencyToken](#)

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](#)

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

Type: Array of [PortRange](#) objects

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

Required: Yes

[Protocol](#)

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

The listener that you've created.

Type: [Listener](#) object

Errors

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

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

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

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

Examples

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [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 [Identity and access management](#) 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](#).

The request accepts the following data in JSON format.

[AcceleratorArn](#)

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](#).

AcceleratorNotFoundException

The accelerator that you specified could not be disabled.

HTTP Status Code: 400

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

HTTP Status Code: 400

AssociatedListenerFoundException

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

HTTP Status Code: 400

InternalServiceErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

See Also

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

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

DeleteCrossAccountAttachment

Delete a cross-account attachment. When you delete an attachment, AWS Global Accelerator revokes the permission to use the resources in the attachment from all principals in the list of principals. AWS Global Accelerator revokes the permission for specific resources by doing the following:

- If the principal is an account ID, Global Accelerator reviews every accelerator in the account and removes cross-account endpoints from all accelerators.
- If the principal is an accelerator, Global Accelerator reviews just that accelerator and removes cross-account endpoints from it.

If there are overlapping permissions provided by multiple cross-account attachments, Global Accelerator only removes endpoints if there are no current cross-account attachments that provide access permission. For example, if you delete a cross-account attachment that lists an accelerator as a principal, but another cross-account attachment includes the account ID that owns that accelerator, endpoints will not be removed from the accelerator.

For more information, see [Working with cross-account attachments and endpoints in AWS Global Accelerator](#) in the *AWS Global Accelerator Developer Guide*.

Request Syntax

```
{  
    "AttachmentArn": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

AttachmentArn

The Amazon Resource Name (ARN) for the cross-account attachment 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](#).

AccessDeniedException

You don't have access permission.

HTTP Status Code: 400

AttachmentNotFoundException

No cross-account attachment was found.

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

TransactionInProgressException

There's already a transaction in progress. Another transaction can't be processed.

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DeleteCustomRoutingAccelerator

Delete a custom routing 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 a custom routing accelerator, by default, Global Accelerator provides you with a set of two static IP addresses.

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 [Identity and access management](#) 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](#).

The request accepts the following data in JSON format.

[AcceleratorArn](#)

The Amazon Resource Name (ARN) of the custom routing accelerator 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](#).

AcceleratorNotFoundException

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

AssociatedListenerNotFoundException

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

HTTP Status Code: 400

InternalServiceErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

See Also

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

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

DeleteCustomRoutingEndpointGroup

Delete an endpoint group from a listener for a custom routing accelerator.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

EndpointGroupArn

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](#).

EndpointGroupNotFoundException

The endpoint group that you specified doesn't exist.

HTTP Status Code: 400

InternalServiceErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DeleteCustomRoutingListener

Delete a listener for a custom routing accelerator.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

ListenerArn

The Amazon Resource Name (ARN) of the listener 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](#).

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [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](#).

The request accepts the following data in JSON format.

EndpointGroupArn

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](#).

EndpointGroupNotFoundException

The endpoint group that you specified doesn't exist.

HTTP Status Code: 400

InternalServiceErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [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](#).

The request accepts the following data in JSON format.

ListenerArn

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](#).

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DenyCustomRoutingTraffic

Specify the Amazon EC2 instance (destination) IP addresses and ports for a VPC subnet endpoint that cannot receive traffic for a custom routing accelerator. You can deny traffic to all destinations in the VPC endpoint, or deny traffic to a specified list of destination IP addresses and ports. Note that you cannot specify IP addresses or ports outside of the range that you configured for the endpoint group.

After you make changes, you can verify that the updates are complete by checking the status of your accelerator: the status changes from IN_PROGRESS to DEPLOYED.

Request Syntax

```
{  
    "DenyAllTrafficToEndpoint": boolean,  
    "DestinationAddresses": [ "string" ],  
    "DestinationPorts": [ number ],  
    "EndpointGroupArn": "string",  
    "EndpointId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

[DenyAllTrafficToEndpoint](#)

Indicates whether all destination IP addresses and ports for a specified VPC subnet endpoint *cannot* receive traffic from a custom routing accelerator. The value is TRUE or FALSE.

When set to TRUE, *no* destinations in the custom routing VPC subnet can receive traffic. Note that you cannot specify destination IP addresses and ports when the value is set to TRUE.

When set to FALSE (or not specified), you *must* specify a list of destination IP addresses that cannot receive traffic. A list of ports is optional. If you don't specify a list of ports, the ports that can accept traffic is the same as the ports configured for the endpoint group.

The default value is FALSE.

Type: Boolean

Required: No

DestinationAddresses

A list of specific Amazon EC2 instance IP addresses (destination addresses) in a subnet that you want to prevent from receiving traffic. The IP addresses must be a subset of the IP addresses allowed for the VPC subnet associated with the endpoint group.

Type: Array of strings

Array Members: Maximum number of 100 items.

Length Constraints: Maximum length of 45.

Required: No

DestinationPorts

A list of specific Amazon EC2 instance ports (destination ports) in a subnet endpoint that you want to prevent from receiving traffic.

Type: Array of integers

Array Members: Maximum number of 100 items.

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

Required: No

EndpointGroupArn

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

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

EndpointId

An ID for the endpoint. For custom routing accelerators, this is the virtual private cloud (VPC) subnet ID.

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](#).

EndpointGroupNotFoundException

The endpoint group that you specified doesn't exist.

HTTP Status Code: 400

InternalServiceErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

Examples

Deny traffic to specific destination instances

The following is an example for specifying specific instances in a subnet endpoint that cannot receive traffic for a custom routing accelerator.

```
aws --region us-west-2 globalaccelerator deny-custom-routing-traffic --endpoint-group-  
arn
```

```
arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-abcd-1234-
abcd-1234abcdefgh/listener/0123vxyz/endpoint-group/ab88888example
--endpoint-id subnet-abcd123example --destination-addresses "198.51.100.52"
"198.51.100.53" --destination-ports "80"
```

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [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.

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](#).

The request accepts the following data in JSON format.

[Cidr](#)

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": {
```

```
"Cidrstring",  
"Events    {  
        "Messagestring",  
        "Timestampnumber  
    }  
,  
"Statestring"  
}  
}
```

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](#)

Information about the address range.

Type: [ByoipCidr](#) object

Errors

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

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

Examples

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 V2](#)

- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DescribeAccelerator

Describe an accelerator.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

AcceleratorArn

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",  
        "DualStackDnsName": "string",  
        "Enabled": boolean,  
        "Events": [  
            {  
                "Message": "string",  
                "Timestamp": number  
            }  
        ]  
    }  
}
```

```
],
  "IpAddressType": "string",
  "IpSets": [
    {
      "IpAddresses": [ "string" ],
      "IpAddressFamily": "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

The description of the accelerator.

Type: [Accelerator](#) object

Errors

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

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

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

Examples

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-1234abcdefg
  --region us-west-2
```

```
{
  "Accelerator": {
    "AcceleratorArn": "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-abcd-1234-
    abcd-1234abcdefg",
    "IpAddressType": "IPV4",
    "Name": "ExampleAaccelerator",
    "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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DescribeAcceleratorAttributes

Describe the attributes of an accelerator.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

AcceleratorArn

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](#)

The attributes of the accelerator.

Type: [AcceleratorAttributes](#) object

Errors

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

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

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

Examples

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-1234abcdefg
```

{

```
"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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DescribeCrossAccountAttachment

Gets configuration information about a cross-account attachment.

Request Syntax

```
{  
    "AttachmentArn": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

AttachmentArn

The Amazon Resource Name (ARN) for the cross-account attachment to describe.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Syntax

```
{  
    "CrossAccountAttachment": {  
        "AttachmentArn": "string",  
        "CreatedTime": number,  
        "LastModifiedTime": number,  
        "Name": "string",  
        "Principals": [ "string" ],  
        "Resources": [  
            {  
                "EndpointId": "string",  
                "Region": "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.

[CrossAccountAttachment](#)

Information about the cross-account attachment.

Type: [Attachment](#) object

Errors

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

AccessDeniedException

You don't have access permission.

HTTP Status Code: 400

AttachmentNotFoundException

No cross-account attachment was found.

HTTP Status Code: 400

InternalServiceErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

See Also

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

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

DescribeCustomRoutingAccelerator

Describe a custom routing accelerator.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

AcceleratorArn

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" ],  
                "IpAddressFamily": "string",  
                "Name": "string",  
                "Type": "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](#)

The description of the custom routing accelerator.

Type: [CustomRoutingAccelerator](#) object

Errors

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

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

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

Examples

Describe a custom routing accelerator

The following is an example for describing a custom routing accelerator, and the response.

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

```
{  
    "Accelerator": {  
        "AcceleratorArn":  
            "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-abcd-1234-  
            abcd-1234abcdefg",  
        "IpAddressType": "IPV4",  
        "Name": "ExampleAaccelerator",  
        "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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DescribeCustomRoutingAcceleratorAttributes

Describe the attributes of a custom routing accelerator.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

AcceleratorArn

The Amazon Resource Name (ARN) of the custom routing accelerator to describe the attributes for.

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

The attributes of the custom routing accelerator.

Type: [CustomRoutingAcceleratorAttributes](#) object

Errors

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

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

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

Examples

Describe attributes for a custom routing accelerator

The following is an example for describing the attributes for a custom routing accelerator.

```
aws globalaccelerator describe-custom-routing-accelerator-attributes  
  --accelerator-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-  
  abcd-1234-abcd-1234abcdefg
```

{

```
"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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DescribeCustomRoutingEndpointGroup

Describe an endpoint group for a custom routing accelerator.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

EndpointGroupArn

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

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Syntax

```
{  
  "EndpointGroup": {  
    "DestinationDescriptions": [  
      {  
        "FromPort": number,  
        "Protocols": [ "string" ],  
        "ToPort": number  
      }  
    ],  
    "EndpointDescriptions": [  
      {  
        "HealthCheck": {  
          "FailureThreshold": integer,  
          "Interval": integer,  
          "Match": "string",  
          "Path": "string",  
          "Port": integer,  
          "Protocol": "string",  
          "Status": "string"  
        },  
        "Id": "string",  
        "Type": "string"  
      }  
    ]  
  }  
}
```

```
        "EndpointId": "string"
    }
],
"EndpointGroupArn": "string",
"EndpointGroupRegion": "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.

[EndpointGroup](#)

The description of an endpoint group for a custom routing accelerator.

Type: [CustomRoutingEndpointGroup](#) object

Errors

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

EndpointGroupNotFoundException

The endpoint group that you specified doesn't exist.

HTTP Status Code: 400

InternalServiceErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

Examples

Describe an endpoint group for a custom routing accelerator

The following is an example for describing an endpoint group for a custom routing accelerator, and the response.

```
aws globalaccelerator describe-custom-routing-endpoint-group
  --endpoint-group-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh/listener/6789vxyz-vxyz-6789-vxyz-6789lmnopqrs/endpoint-
group/ab88888example
```

```
{
    "EndpointGroup": {
        "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-west-2",
        "DestinationDescriptions": [
            {
                "FromPort": 80,
                "ToPort": 80,
                "Protocols": [
                    "UDP"
                ]
            }
        ],
        "EndpointDescriptions": [
            {
                "EndpointId": "subnet-1234567890abcdef0"
            }
        ]
    }
}
```

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DescribeCustomRoutingListener

The description of a listener for a custom routing accelerator.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

ListenerArn

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

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Response Syntax

```
{  
    "Listener": {  
        "ListenerArn": "string",  
        "PortRanges": [  
            {  
                "FromPort": number,  
                "ToPort": 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.

Listener

The description of a listener for a custom routing accelerator.

Type: [CustomRoutingListener](#) object

Errors

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

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DescribeEndpointGroup

Describe an endpoint group.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

EndpointGroupArn

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",  
        "HealthCheckConfig": {  
            "Interval": integer,  
            "Timeout": integer,  
            "UnhealthyThreshold": integer  
        },  
        "HealthCheckProtocol": "HTTP" | "HTTPS" | "TCP" | "HTTP2" | "TLS"| "QUIC",  
        "HealthCheckVersion": "V2",  
        "Name": "string",  
        "State": "PENDING" | "ACTIVE" | "DELETING" | "FAILED",  
        "Type": "IP" | "L4" | "L7"  
    }  
}
```

```
"EndpointGroupRegion": "string",
"HealthCheckIntervalSecondsHealthCheckPath": "string",
"HealthCheckPort": "number",
"HealthCheckProtocol": "string",
"PortOverrides": [
    {
        "EndpointPort": "number",
        "ListenerPort": "number"
    }
],
"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

The description of an endpoint group.

Type: [EndpointGroup](#) object

Errors

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

EndpointGroupNotFoundException

The endpoint group that you specified doesn't exist.

HTTP Status Code: 400

InternalServiceErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

Examples

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/ab88888example
```

```
{
  "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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DescribeListener

Describe a listener.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

ListenerArn

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](#)

The description of a listener.

Type: [Listener](#) object

Errors

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

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

Examples

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListAccelerators

List the accelerators for an AWS account.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

MaxResults

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

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": [
  {
    "AcceleratorArnCreatedTimeDnsNameDualStackDnsNameEnabledEventsMessageTimestampIpAddressTypeIpSetsIpAddressesIpAddressFamilyIpFamilyLastModifiedTimeNameStatusNextToken
```

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

The list of accelerators for a customer account.

Type: Array of [Accelerator](#) objects

NextToken

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](#).

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

Examples

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
```

```
{
"Accelerators": [
    {
        "AcceleratorArn":
"arn:aws:globalaccelerator::012345678901:accelerator/5555abcd-abcd-5555-
abcd-5555EXAMPLE1",
        "Name": "TestAccelerator",
        "IpAddressType": "IPV4",
```

```
        "Enabled": true,
        "IpSets": [
            {
                "IpFamily": "IPv4",
                "IpAddresses": [
                    "192.0.2.250",
                    "198.51.100.52"
                ]
            }
        ],
        "DnsName": "5a5a5a5a5a5a5a5a.awsglobalaccelerator.com",
        "Status": "DEPLOYED",
        "CreatedTime": 1552424416.0,
        "LastModifiedTime": 1569375641.0
    },
    {
        "AcceleratorArn":
            "arn:aws:globalaccelerator::888888888888:accelerator/8888abcd-abcd-8888-
            abcd-8888EXAMPLE2",
        "Name": "ExampleAccelerator",
        "IpAddressType": "IPV4",
        "Enabled": true,
        "IpSets": [
            {
                "IpFamily": "IPv4",
                "IpAddresses": [
                    "192.0.2.100",
                    "198.51.100.10"
                ]
            }
        ],
        "DnsName": "6a6a6a6a6a6a6a6a.awsglobalaccelerator.com",
        "Status": "DEPLOYED",
        "CreatedTime": 1575585564.0,
        "LastModifiedTime": 1579809243.0
    },
]
}
```

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [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.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

MaxResults

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

The token for the next page of results.

Type: String

Length Constraints: Maximum length of 255.

Required: No

Response Syntax

```
{
```

```
"ByoipCidrs": [
  {
    "CidrEventsMessageTimestampStateNextToken
```

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

Information about your address ranges.

Type: Array of [ByoipCidr](#) objects

NextToken

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](#).

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

InvalidNextTokenException

There isn't another item to return.

HTTP Status Code: 400

Examples

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListCrossAccountAttachments

List the cross-account attachments that have been created in AWS Global Accelerator.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

MaxResults

The number of cross-account attachment 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

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

```
{  
    "CrossAccountAttachments": [  
        ...  
    ]  
}
```

```
{  
    "AttachmentArn": "string",  
    "CreatedTime": number,  
    "LastModifiedTime": number,  
    "Name": "string",  
    "Principals": [ "string" ],  
    "Resources": [  
        {  
            "EndpointId": "string",  
            "Region": "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.

CrossAccountAttachments

Information about the cross-account attachments.

Type: Array of [Attachment](#) objects

NextToken

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](#).

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

InvalidNextTokenException

There isn't another item to return.

HTTP Status Code: 400

See Also

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

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

ListCrossAccountResourceAccounts

List the accounts that have cross-account endpoints.

For more information, see [Working with cross-account attachments and endpoints in AWS Global Accelerator](#) in the *AWS Global Accelerator Developer Guide*.

Response Syntax

```
{  
    "ResourceOwnerAwsAccountIds": [ "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.

ResourceOwnerAwsAccountIds

The account IDs of principals (resource owners) in a cross-account attachment who can add endpoints (resources) listed in the same attachment.

Type: Array of strings

Length Constraints: Fixed length of 12.

Pattern: ^\d{12}\$

Errors

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

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

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListCrossAccountResources

List the cross-account resources available to add as endpoints to an accelerator.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

[AcceleratorArn](#)

The Amazon Resource Name (ARN) of an accelerator in a cross-account attachment.

Type: String

Length Constraints: Maximum length of 255.

Required: No

[MaxResults](#)

The number of cross-account endpoints 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](#)

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

[ResourceOwnerAwsAccountId](#)

The account ID of a resource owner in a cross-account attachment.

Type: String

Length Constraints: Fixed length of 12.

Pattern: ^\d{12}\$

Required: Yes

Response Syntax

```
{  
    "CrossAccountResources": [  
        {  
            "AttachmentArn": "string",  
            "EndpointId": "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.

[CrossAccountResources](#)

The endpoints attached to an accelerator in a cross-account attachment.

Type: Array of [CrossAccountResource](#) objects

[NextToken](#)

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](#).

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

InvalidNextTokenException

There isn't another item to return.

HTTP Status Code: 400

See Also

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

- [AWS Command Line Interface](#)

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

ListCustomRoutingAccelerators

List the custom routing accelerators for an AWS account.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

MaxResults

The number of custom routing 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

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": [  
        object  
    ]  
}
```

```
{  
    "AcceleratorArn": "string",  
    "CreatedTime": number,  
    "DnsName": "string",  
    "Enabled": boolean,  
    "IpAddressType": "string",  
    "IpSets": [  
        {  
            "IpAddresses": [ "string" ],  
            "IpAddressFamily": "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](#)

The list of custom routing accelerators for a customer account.

Type: Array of [CustomRoutingAccelerator](#) objects

[NextToken](#)

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](#).

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

Examples

List custom routing accelerators

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

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

```
{  
    "Accelerators": [  
        {  
            "AcceleratorArn":  
                "arn:aws:globalaccelerator::012345678901:accelerator/5555abcd-abcd-5555-  
                abcd-5555EXAMPLE1",  
            "Name": "TestAccelerator",  
            "IpAddressType": "IPV4",  
            "Enabled": true,  
            "IpSets": [  
                {  
                    "IpFamily": "IPv4",  
                    "IpAddresses": [  
                        "192.0.2.250",  
                        "192.0.2.251",  
                        "192.0.2.252",  
                        "192.0.2.253",  
                        "192.0.2.254",  
                        "192.0.2.255  
                    ]  
                }  
            ]  
        }  
    ]  
}
```

```
        "198.51.100.52"
    ]
}
],
"DnsName": "5a5a5a5a5a5a5a5a.awsglobalaccelerator.com",
"Status": "DEPLOYED",
"CreatedTime": 1552424416.0,
"LastModifiedTime": 1569375641.0
},
{
    "AcceleratorArn":
"arn:aws:globalaccelerator::888888888888:accelerator/8888abcd-abcd-8888-
abcd-8888EXAMPLE2",
        "Name": "ExampleAccelerator",
        "IpAddressType": "IPV4",
        "Enabled": true,
        "IpSets": [
            {
                "IpFamily": "IPv4",
                "IpAddresses": [
                    "192.0.2.100",
                    "198.51.100.10"
                ]
            }
        ],
        "DnsName": "6a6a6a6a6a6a6a.awsglobalaccelerator.com",
        "Status": "DEPLOYED",
        "CreatedTime": 1575585564.0,
        "LastModifiedTime": 1579809243.0
},
]
}
```

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListCustomRoutingEndpointGroups

List the endpoint groups that are associated with a listener for a custom routing accelerator.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

ListenerArn

The Amazon Resource Name (ARN) of the listener to list endpoint groups for.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

MaxResults

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

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": [  
        {  
            "DestinationDescriptions": [  
                {  
                    "FromPort": number,  
                    "Protocols": [ "string" ],  
                    "ToPort": number  
                }  
            ],  
            "EndpointDescriptions": [  
                {  
                    "EndpointId": "string"  
                }  
            ],  
            "EndpointGroupArn": "string",  
            "EndpointGroupRegion": "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.

[EndpointGroups](#)

The list of the endpoint groups associated with a listener for a custom routing accelerator.

Type: Array of [CustomRoutingEndpointGroup](#) objects

[NextToken](#)

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](#).

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

Examples

List endpoint groups for a custom routing accelerator

The following is an example for listing the endpoint groups associated with a listener for a custom routing accelerator, and the response.

```
aws globalaccelerator list-custom-routing-endpoint-groups
    --listener-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefg/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": []  
        }  
        {  
            "EndpointGroupArn":  
                "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-abcd-1234-  
                abcd-1234abcdefgh/listener/abcdef1234/endpoint-group/ab99999example",  
            "EndpointGroupRegion": "us-east-1",  
            "EndpointDescriptions": []  
        }  
    ]  
}
```

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListCustomRoutingListeners

List the listeners for a custom routing accelerator.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

AcceleratorArn

The Amazon Resource Name (ARN) of the accelerator to list listeners for.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

MaxResults

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

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": [  
        {  
            "ListenerArn": "string",  
            "PortRanges": [  
                {  
                    "FromPort": number,  
                    "ToPort": 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.

Listeners

The list of listeners for a custom routing accelerator.

Type: Array of [CustomRoutingListener](#) objects

NextToken

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](#).

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

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

InvalidNextTokenException

There isn't another item to return.

HTTP Status Code: 400

Examples

List listeners for a custom routing accelerator

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

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

```
{
  "Listeners": [
    {
      "ListenerArn":
      "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-abcd-1234-
      abcd-1234abcdefg/listener/abcdef1234",
```

```
        "PortRanges": [
            {
                "FromPort": 5000,
                "ToPort": 10000
            }
        ],
        "Protocol": "TCP",
    }
]
}
```

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListCustomRoutingPortMappings

Provides a complete mapping from the public accelerator IP address and port to destination EC2 instance IP addresses and ports in the virtual public cloud (VPC) subnet endpoint for a custom routing accelerator. For each subnet endpoint that you add, Global Accelerator creates a new static port mapping for the accelerator. The port mappings don't change after Global Accelerator generates them, so you can retrieve and cache the full mapping on your servers.

If you remove a subnet from your accelerator, Global Accelerator removes (reclaims) the port mappings. If you add a subnet to your accelerator, Global Accelerator creates new port mappings (the existing ones don't change). If you add or remove EC2 instances in your subnet, the port mappings don't change, because the mappings are created when you add the subnet to Global Accelerator.

The mappings also include a flag for each destination denoting which destination IP addresses and ports are allowed or denied traffic.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

[AcceleratorArn](#)

The Amazon Resource Name (ARN) of the accelerator to list the custom routing port mappings for.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

[EndpointGroupArn](#)

The Amazon Resource Name (ARN) of the endpoint group to list the custom routing port mappings for.

Type: String

Length Constraints: Maximum length of 255.

Required: No

[MaxResults](#)

The number of destination port mappings that you want to return with this call. The default value is 10.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 20000.

Required: No

[NextToken](#)

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

```
{  
    "NextToken": "string",  
    "PortMappings": [  
        {  
            "AcceleratorPort": number,  
            "DestinationSocketAddress": {  
                "IpAddress": "string",  
                "Port": number  
            }  
        }  
    ]  
}
```

```
        },
        "DestinationTrafficState": "string",
        "EndpointGroupArn": "string",
        "EndpointId": "string",
        "Protocols": [ "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.

[NextToken](#)

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

Type: String

Length Constraints: Maximum length of 255.

[PortMappings](#)

The port mappings for a custom routing accelerator.

Type: Array of [PortMapping](#) objects

Errors

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

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

HTTP Status Code: 400

EndpointGroupNotFoundException

The endpoint group that you specified doesn't exist.

HTTP Status Code: 400

InternalServiceErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

InvalidNextTokenException

There isn't another item to return.

HTTP Status Code: 400

See Also

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

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

ListCustomRoutingPortMappingsByDestination

List the port mappings for a specific EC2 instance (destination) in a VPC subnet endpoint. The response is the mappings for one destination IP address. This is useful when your subnet endpoint has mappings that span multiple custom routing accelerators in your account, or for scenarios where you only want to list the port mappings for a specific destination instance.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

[DestinationAddress](#)

The endpoint IP address in a virtual private cloud (VPC) subnet for which you want to receive back port mappings.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

[EndpointId](#)

The ID for the virtual private cloud (VPC) subnet.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

MaxResults

The number of destination port mappings that you want to return with this call. The default value is 10.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 20000.

Required: No

NextToken

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

```
{  
    "DestinationPortMappings": [  
        {  
            "AcceleratorArn": "string",  
            "AcceleratorSocketAddresses": [  
                {  
                    "IpAddress": "string",  
                    "Port": number  
                }  
            ],  
            "DestinationSocketAddress": {  
                "IpAddress": "string",  
                "Port": number  
            },  
            "DestinationTrafficState": "string",  
            "EndpointGroupArn": "string",  
            "EndpointGroupRegion": "string",  
            "EndpointId": "string",  
            "IpAddressType": "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.

[DestinationPortMappings](#)

The port mappings for the endpoint IP address that you specified in the request.

Type: Array of [DestinationPortMapping](#) objects

[NextToken](#)

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](#).

EndpointNotFoundException

The endpoint that you specified doesn't exist.

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

InvalidNextTokenException

There isn't another item to return.

HTTP Status Code: 400

Examples

List the port mappings for an EC2 instance destination

The following is an example of listing the port mappings for an EC2 instance destination for a custom routing accelerator.

```
aws --region us-west-2 globalaccelerator list-custom-routing-port-mappings-by-destination  
    --endpoint-id subnet-abcd123example --destination-address "198.51.100.52"
```

```
{  
    "DestinationPortMappings": [  
        {  
            "AcceleratorArn":  
                "arn:aws:globalaccelerator::402092451327:accelerator/24ea29b8-d750-4489-8919-3095f3c4b0a7",  
            "AcceleratorSocketAddresses": [  
                {  
                    "IpAddress": "192.0.2.250",  
                    "Port": 65514  
                },  
                {  
                    "IpAddress": "198.51.100.52",  
                    "Port": 65514  
                }  
            ],  
            "EndpointGroupArn":  
                "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-abcd-1234-abcd-1234abcdefg/listener/0123vxyz/endpoint-group/ab88888example",  
            "EndpointId": "subnet-abcd123example",  
            "EndpointGroupRegion": "us-west-2",  
        }  
    ]  
}
```

```
        "DestinationSocketAddress": {  
            "IpAddress": "198.51.100.52",  
            "Port": 80  
        },  
        "IpAddressType": "IPv4",  
        "DestinationTrafficState": "DENY"  
    }  
]  
}
```

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [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.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

ListenerArn

The Amazon Resource Name (ARN) of the listener.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

MaxResults

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

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",  
            "PortOverrides": [  
                {  
                    "EndpointPort": number,  
                    "ListenerPort": number  
                }  
            ],  
            "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

The list of the endpoint groups associated with a listener.

Type: Array of [EndpointGroup](#) objects

NextToken

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](#).

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

Examples

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-1234abcd-1234abcdefgh/listener/abcdef1234
  --region us-west-2
```

```
{
  "EndpointGroups": [
    {
      "EndpointGroupArn":
        "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-abcd-1234-
abcd-1234abcdefgh/listener/abcdef1234/endpoint-group/ab8888example",
        "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/ab9999example",
        "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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListListeners

List the listeners for an accelerator.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

AcceleratorArn

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

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

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

The list of listeners for an accelerator.

Type: Array of [Listener](#) objects

NextToken

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](#).

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

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

InvalidNextTokenException

There isn't another item to return.

HTTP Status Code: 400

Examples

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-1234abcdefg
  --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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListTagsForResource

List all tags for an accelerator.

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](#).

The request accepts the following data in JSON format.

ResourceArn

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

Root level tag for the Tags parameters.

Type: Array of [Tag](#) objects

Errors

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

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

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

Examples

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-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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [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](#).

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](#).

The request accepts the following data in JSON format.

[Cidr](#)

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](#)

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

Type: [CidrAuthorizationContext](#) 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](#)

Information about the address range.

Type: [ByoipCidr](#) object

Errors

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

[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

Examples

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

RemoveCustomRoutingEndpoints

Remove endpoints from a custom routing accelerator.

Request Syntax

```
{  
    "EndpointGroupArn": "string",  
    "EndpointIds": [ "string" ]  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

EndpointGroupArn

The Amazon Resource Name (ARN) of the endpoint group to remove endpoints from.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

EndpointIds

The IDs for the endpoints. For custom routing accelerators, endpoint IDs are the virtual private cloud (VPC) subnet IDs.

Type: Array of strings

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](#).

AccessDeniedException

You don't have access permission.

HTTP Status Code: 400

ConflictException

You can't use both of those options.

HTTP Status Code: 400

EndpointGroupNotFoundException

The endpoint group that you specified doesn't exist.

HTTP Status Code: 400

EndpointNotFoundException

The endpoint that you specified doesn't exist.

HTTP Status Code: 400

InternalServiceErrorException

There was an internal error for AWS Global Accelerator.

HTTP Status Code: 400

InvalidArgumentException

An argument that you specified is invalid.

HTTP Status Code: 400

See Also

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

- [AWS Command Line Interface](#)

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

RemoveEndpoints

Remove endpoints from an endpoint group.

The RemoveEndpoints API operation is the recommended option for removing endpoints.

The alternative is to remove endpoints by updating an endpoint group by using the

[UpdateEndpointGroup](#) API operation. There are two advantages to using AddEndpoints to remove endpoints instead:

- It's more convenient, because you only need to specify the endpoints that you want to remove. With the UpdateEndpointGroup API operation, you must specify all of the endpoints in the endpoint group except the ones that you want to remove from the group.
- It's faster, because Global Accelerator doesn't need to resolve any endpoints. With the UpdateEndpointGroup API operation, Global Accelerator must resolve all of the endpoints that remain in the group.

Request Syntax

```
{  
    "EndpointGroupArn": "string",  
    "EndpointIdentifiers": [  
        {  
            "ClientIPPreservationEnabled": boolean,  
            "EndpointId": "string"  
        }  
    ]  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

[EndpointGroupArn](#)

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

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

EndpointIdentifiers

The identifiers of the endpoints that you want to remove.

Type: Array of [EndpointIdentifier](#) objects

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

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](#).

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

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

TransactionInProgressException

There's already a transaction in progress. Another transaction can't be processed.

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

TagResource

Add tags to an accelerator resource.

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](#).

The request accepts the following data in JSON format.

ResourceArn

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

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

Type: Array of [Tag](#) 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](#).

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

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

Examples

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-1234abcdefg  
--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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [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. 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](#).

The request accepts the following data in JSON format.

ResourceArn

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

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](#).

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

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

Examples

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-1234abcdefg
  --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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

UpdateAccelerator

Update an accelerator to make changes, such as the following:

- Change the name of the accelerator.
- Disable the accelerator so that it no longer accepts or routes traffic, or so that you can delete it.
- Enable the accelerator, if it is disabled.
- Change the IP address type to dual-stack if it is IPv4, or change the IP address type to IPv4 if it's dual-stack.

Be aware that 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 the accelerator, you lose the static IP addresses that are assigned to it, so you can no longer route traffic by using them.

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, update, or otherwise work with accelerators. That is, for example, specify `--region us-west-2` on AWS CLI commands.

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](#).

The request accepts the following data in JSON format.

AcceleratorArn

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

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Enabled

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

The IP address type that an accelerator supports. For a standard accelerator, the value can be IPV4 or DUAL_STACK.

Type: String

Valid Values: IPV4 | DUAL_STACK

Required: No

Name

The name of the accelerator. The name can have a maximum of 64 characters, must contain only alphanumeric characters, periods (.), or hyphens (-), and must not begin or end with a hyphen or period.

Type: String

Length Constraints: Maximum length of 255.

Required: No

Response Syntax

```
{  
  "Accelerator": {  
    "AcceleratorArn": "string",  
    "CreatedTime": number,  
    "DnsName": "string",  
    "DualStackDnsName": "string",  
    "Enabled": boolean,  
    "Events": [  
      {  
        "Message": "string",  
        "Timestamp": number  
      }  
    ],  
    "IpAddressType": "string",  
    "IpSets": [  
      {  
        "IpAddresses": [ "string" ],  
        "IpAddressFamily": "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](#)

Information about the updated accelerator.

Type: [Accelerator](#) object

Errors

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

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

Examples

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-1234abcdefg
  --name ExampleAcceleratorNew
  --region us-west-2
```

```
{
  "Accelerator":{
```

```
"AcceleratorArn": "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-abcd-1234-abcd-1234abcdefg",
    "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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

UpdateAcceleratorAttributes

Update the attributes for an accelerator.

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](#).

The request accepts the following data in JSON format.

AcceleratorArn

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

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

FlowLogsEnabled

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

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

Update the prefix for the location in the Amazon S3 bucket for the flow logs. Attribute is required if FlowLogsEnabled is true.

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](#)

Updated attributes for the accelerator.

Type: [AcceleratorAttributes](#) object

Errors

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

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

Examples

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-1234abcdefg  
--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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

UpdateCrossAccountAttachment

Update a cross-account attachment to add or remove principals or resources. When you update an attachment to remove a principal (account ID or accelerator) or a resource, AWS Global Accelerator revokes the permission for specific resources by doing the following:

- If the principal is an account ID, Global Accelerator reviews every accelerator in the account and removes cross-account endpoints from all accelerators.
- If the principal is an accelerator, Global Accelerator reviews just that accelerator and removes cross-account endpoints from it.

If there are overlapping permissions provided by multiple cross-account attachments, Global Accelerator only removes endpoints if there are no current cross-account attachments that provide access permission. For example, if you delete a cross-account attachment that lists an accelerator as a principal, but another cross-account attachment includes the account ID that owns that accelerator, endpoints will not be removed from the accelerator.

For more information, see [Working with cross-account attachments and endpoints in AWS Global Accelerator](#) in the *AWS Global Accelerator Developer Guide*.

Request Syntax

```
{  
    "AddPrincipals": [ "string" ],  
    "AddResources": [  
        {  
            "EndpointId": "string",  
            "Region": "string"  
        }  
    ],  
    "AttachmentArn": "string",  
    "Name": "string",  
    "RemovePrincipals": [ "string" ],  
    "RemoveResources": [  
        {  
            "EndpointId": "string",  
            "Region": "string"  
        }  
    ]  
}
```

```
}
```

Request Parameters

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

The request accepts the following data in JSON format.

AddPrincipals

The principals to add to the cross-account attachment. A principal is an account or the Amazon Resource Name (ARN) of an accelerator that the attachment gives permission to add the resources from another account, listed in the attachment.

To add more than one principal, separate the account numbers or accelerator ARNs, or both, with commas.

Type: Array of strings

Length Constraints: Maximum length of 256.

Pattern: (^ $\d\{12\}$ \$|arn:.*)

Required: No

AddResources

The resources to add to the cross-account attachment. A resource listed in a cross-account attachment can be added to an accelerator by the principals that are listed in the attachment.

To add more than one resource, separate the resource ARNs with commas.

Type: Array of [Resource](#) objects

Required: No

AttachmentArn

The Amazon Resource Name (ARN) of the cross-account attachment to update.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Name

The name of the cross-account attachment.

Type: String

Length Constraints: Maximum length of 64.

Pattern: [\S\s]+

Required: No

RemovePrincipals

The principals to remove from the cross-account attachment. A principal is an account or the Amazon Resource Name (ARN) of an accelerator that is given permission to add the resources from another account, listed in the cross-account attachment.

To remove more than one principal, separate the account numbers or accelerator ARNs, or both, with commas.

Type: Array of strings

Length Constraints: Maximum length of 256.

Pattern: (^d{12}\$|arn:.*)

Required: No

RemoveResources

The resources to remove from the cross-account attachment. A resource listed in a cross-account attachment can be added to an accelerator by principals that are listed in the cross-account attachment.

To remove more than one resource, separate the resource ARNs with commas.

Type: Array of [Resource](#) objects

Required: No

Response Syntax

```
{
```

```
"CrossAccountAttachment": {  
    "AttachmentArn    "CreatedTime": number,  
    "LastModifiedTime": number,  
    "Name": "string",  
    "Principals": [ "string" ],  
    "Resources": [  
        {  
            "EndpointId": "string",  
            "Region": "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.

CrossAccountAttachment

Information about the updated cross-account attachment.

Type: [Attachment](#) object

Errors

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

AccessDeniedException

You don't have access permission.

HTTP Status Code: 400

AttachmentNotFoundException

No cross-account attachment was found.

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

TransactionInProgressException

There's already a transaction in progress. Another transaction can't be processed.

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

UpdateCustomRoutingAccelerator

Update a custom routing accelerator.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

AcceleratorArn

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

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Enabled

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

The IP address type that an accelerator supports. For a custom routing accelerator, the value must be IPV4.

Type: String

Valid Values: IPV4 | DUAL_STACK

Required: No

Name

The name of the accelerator. The name can have a maximum of 64 characters, must contain only alphanumeric characters, periods (.), or hyphens (-), and must not begin or end with a hyphen or period.

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" ],  
                "IpAddressFamily": "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

Information about the updated custom routing accelerator.

Type: [CustomRoutingAccelerator](#) object

Errors

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

AcceleratorNotFoundException

The accelerator that you specified doesn't exist.

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

Examples

Update a custom routing accelerator

The following is an example for updating a custom routing accelerator.

```
aws --region us-west-2 globalaccelerator update-custom-routing-accelerator  
  --accelerator-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-  
  abcd-1234-abcd-1234abcdefg  
  --enabled
```

```
{  
    "Accelerator": {  
        "AcceleratorArn":  
            "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-abcd-1234-  
            abcd-1234abcdefg",  
        "Name": "testaccelerator",  
        "IpAddressType": "IPV4",  
        "Enabled": true,  
        "IpSets": [  
            {  
                "IpAddresses": [  
                    "192.0.2.250",  
                    "198.51.100.52"  
                ],  
                "IpFamily": "IPv4"  
            }  
        ]  
    },  
    "DnsName": "a1234567890abcdef.awsglobalaccelerator.com",  
    "Status": "IN_PROGRESS",  
    "CreatedTime": 1605833295.0,  
    "LastModifiedTime": 1605910135.0  
}  
}
```

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

UpdateCustomRoutingAcceleratorAttributes

Update the attributes for a custom routing accelerator.

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](#).

The request accepts the following data in JSON format.

AcceleratorArn

The Amazon Resource Name (ARN) of the custom routing accelerator to update attributes for.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

FlowLogsEnabled

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

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

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:

DOC-EXAMPLE-BUCKET//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

Updated custom routing accelerator.

Type: [CustomRoutingAcceleratorAttributes](#) object

Errors

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

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

Examples

Update attributes for an accelerator

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

```
aws globalaccelerator update-custom-routing-accelerator-attributes
```

```
--accelerator-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-  
abcd-1234-abcd-1234abcdefg  
--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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

UpdateCustomRoutingListener

Update a listener for a custom routing accelerator.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

ListenerArn

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

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

PortRanges

The updated port range to support for connections from clients to your accelerator. If you remove ports that are currently being used by a subnet endpoint, the call fails.

Separately, you set port ranges for endpoints. For more information, see [About endpoints for custom routing accelerators](#).

Type: Array of [PortRange](#) objects

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

Required: Yes

Response Syntax

```
{  
  "Listener": {  
    "ListenerArn": "string",  
    "PortRanges": [  
      {  
        "FromPort": number,  
        "ToPort": 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.

[Listener](#)

Information for the updated listener for a custom routing accelerator.

Type: [CustomRoutingListener](#) object

Errors

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

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

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

Examples

Update a listener for a custom routing accelerator

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

```
aws globalaccelerator update-custom-routing-listener
  --listener-arn arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
  abcd-1234-abcd-1234abcdefgh/listener/0123vxyz
  --port-ranges FromPort=5000,ToPort=10000
```

```
{
  "Listener": {
    "ListenerArn": "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
    abcd-1234-abcd-1234abcdefgh/listener/0123vxyz
    "PortRanges": [
      {
        "FromPort": 5000,
        "ToPort": 10000
      }
    ]
  }
}
```

```
    ],
    "Protocol": "TCP",
}
}
```

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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [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.

Request Syntax

```
{  
    "EndpointConfigurations": [  
        {  
            "AttachmentArn": "string",  
            "ClientIPPreservationEnabled": boolean,  
            "EndpointId": "string",  
            "Weight": number  
        }  
    ],  
    "EndpointGroupArn": "string",  
    "HealthCheckIntervalSeconds": number,  
    "HealthCheckPath": "string",  
    "HealthCheckPort": number,  
    "HealthCheckProtocol": "string",  
    "PortOverrides": [  
        {  
            "EndpointPort": number,  
            "ListenerPort": number  
        }  
    ],  
    "ThresholdCount": number,  
    "TrafficDialPercentage": number  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

[EndpointConfigurations](#)

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

Type: Array of [EndpointConfiguration](#) objects

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

Required: No

EndpointGroupArn

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

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

HealthCheckIntervalSeconds

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

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.

Pattern: ^/[-a-zA-Z0-9@:_\\\+.\~#?&/=]*\$

Required: No

HealthCheckPort

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

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

PortOverrides

Override specific listener ports used to route traffic to endpoints that are part of this endpoint group. For example, you can create a port override in which the listener receives user traffic on ports 80 and 443, but your accelerator routes that traffic to ports 1080 and 1443, respectively, on the endpoints.

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

Type: Array of [PortOverride](#) objects

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

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

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",  
        "PortOverrides": [  
            {  
                "EndpointPort": number,  
                "ListenerPort": number  
            }  
        ],  
    },  
}
```

```
    "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](#)

The information about the endpoint group that was updated.

Type: [EndpointGroup](#) object

Errors

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

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

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

Examples

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/endpoint-group/ab88888example
  --port-overrides ListenerPort=443,EndpointPort=1443
  --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"
      }
    ]
  }
}
```

```
        ],
        "PortOverrides": [
            {
                "EndpointPort": 1443,
                "ListenerPort": 443
            }
        ],
        "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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

UpdateListener

Update a listener.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

[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 to update.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

PortRanges

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

Type: Array of [PortRange](#) objects

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

Required: No

Protocol

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": [  
      "PortRange": {  
        "FromPort": 1, "ToPort": 1, "Protocol": "TCP", "PortRanges": [  
          {"FromPort": 1, "ToPort": 1, "Protocol": "TCP"}]  
        }  
      }  
    ]  
  }  
}
```

```
{  
    "FromPort": number,  
    "ToPort": number  
}  
,  
"Protocolstring"  
}  
}
```

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](#)

Information for the updated listener.

Type: [Listener](#) object

Errors

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

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

InvalidPortRangeException

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

HTTP Status Code: 400

LimitExceeded**Exception**

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

Examples

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-1234abcdefgh/listener/0123vxyz
  --port-ranges FromPort=100,ToPort=100
```

```
{
  "Listener": {
    "ListenerArn": "arn:aws:globalaccelerator::012345678901:accelerator/1234abcd-
abcd-1234-abcd-1234abcdefgh/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 V2](#)
- [AWS SDK for JavaScript V3](#)
- [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.

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](#).

The request accepts the following data in JSON format.

Cidr

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](#)

Information about the address pool.

Type: [ByoipCidr](#) object

Errors

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

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

Examples

Withdrawing an address range from advertising by

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 V2](#)

- [AWS SDK for JavaScript V3](#)
- [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](#)
- [AcceleratorAttributes](#)
- [AcceleratorEvent](#)
- [Attachment](#)
- [ByoipCidr](#)
- [ByoipCidrEvent](#)
- [CidrAuthorizationContext](#)
- [CrossAccountResource](#)
- [CustomRoutingAccelerator](#)
- [CustomRoutingAcceleratorAttributes](#)
- [CustomRoutingDestinationConfiguration](#)
- [CustomRoutingDestinationDescription](#)
- [CustomRoutingEndpointConfiguration](#)
- [CustomRoutingEndpointDescription](#)
- [CustomRoutingEndpointGroup](#)
- [CustomRoutingListener](#)
- [DestinationPortMapping](#)
- [EndpointConfiguration](#)
- [EndpointDescription](#)
- [EndpointGroup](#)

- [EndpointIdentifier](#)
- [IpSet](#)
- [Listener](#)
- [PortMapping](#)
- [PortOverride](#)
- [PortRange](#)
- [Resource](#)
- [SocketAddress](#)
- [Tag](#)

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 an accelerator's static IPv4 addresses.

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

If you have a dual-stack accelerator, you also have a second DNS name, DualStackDnsName, that points to both the A record and the AAAA record for all four static addresses for the accelerator: two IPv4 addresses and two IPv6 addresses.

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

DualStackDnsName

The Domain Name System (DNS) name that Global Accelerator creates that points to a dual-stack accelerator's four static IP addresses: two IPv4 addresses and two IPv6 addresses.

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

Note: Global Accelerator also assigns a default DNS name, DnsName, to your accelerator that points just to the static IPv4 addresses.

For more information, 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

Events

A history of changes that you make to an accelerator in Global Accelerator.

Type: Array of [AcceleratorEvent](#) objects

Required: No

IpAddressType

The IP address type that an accelerator supports. For a standard accelerator, the value can be IPV4 or DUAL_STACK.

Type: String

Valid Values: IPV4 | DUAL_STACK

Required: No

IpSets

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

Type: Array of [IpSet](#) 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 V2](#)
- [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 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 V2](#)
- [AWS SDK for Ruby V3](#)

AcceleratorEvent

A complex type that contains a `Timestamp` value and `Message` for changes that you make to an accelerator in Global Accelerator. Messages stored here provide progress or error information when you update an accelerator from IPv4 to dual-stack, or from dual-stack to IPv4. Global Accelerator stores a maximum of ten event messages.

Contents

Message

A string that contains an Event message describing changes or errors when you update an accelerator in Global Accelerator from IPv4 to dual-stack, or dual-stack to IPv4.

Type: String

Length Constraints: Maximum length of 255.

Required: No

Timestamp

A timestamp for when you update an accelerator in Global Accelerator from IPv4 to dual-stack, or dual-stack to IPv4.

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 V2](#)
- [AWS SDK for Ruby V3](#)

Attachment

A cross-account attachment in AWS Global Accelerator. A cross-account attachment specifies the *principals* who have permission to add to accelerators in their own account the resources in your account that you also list in the attachment.

Contents

AttachmentArn

The Amazon Resource Name (ARN) of the cross-account attachment.

Type: String

Length Constraints: Maximum length of 255.

Required: No

CreatedTime

The date and time that the cross-account attachment was created.

Type: Timestamp

Required: No

LastModifiedTime

The date and time that the cross-account attachment was last modified.

Type: Timestamp

Required: No

Name

The name of the cross-account attachment.

Type: String

Length Constraints: Maximum length of 64.

Pattern: [\S\s]+

Required: No

Principals

The principals included in the cross-account attachment.

Type: Array of strings

Length Constraints: Maximum length of 256.

Pattern: (^ $\d\{12\}$ \$|arn:.*)

Required: No

Resources

The resources included in the cross-account attachment.

Type: Array of [Resource](#) objects

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 V2](#)
- [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](#) 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 V2](#)

- [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 of an IP address range that you bring to 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 Global Accelerator through bring your own IP address (BYOIP).

Type: String

Length Constraints: Maximum length of 255.

Required: No

Timestamp

A timestamp for when you make a status change for an IP address range that you bring to 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 V2](#)
- [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 V2](#)
- [AWS SDK for Ruby V3](#)

CrossAccountResource

An endpoint (AWS resource) that is listed in a cross-account attachment and can be added to an accelerator by specified principals, that are also listed in the attachment.

For more information, see [Working with cross-account attachments and endpoints in AWS Global Accelerator](#) in the *AWS Global Accelerator Developer Guide*.

Contents

AttachmentArn

The Amazon Resource Name (ARN) of the cross-account attachment that specifies the endpoints (resources) that can be added to accelerators and principals that have permission to add the endpoints to accelerators.

Type: String

Length Constraints: Maximum length of 255.

Required: No

EndpointId

The endpoint ID for the endpoint that is listed in a cross-account attachment and can be added to an accelerator by specified principals.

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 V2](#)

- [AWS SDK for Ruby V3](#)

CustomRoutingAccelerator

Attributes of a custom routing accelerator.

Contents

AcceleratorArn

The Amazon Resource Name (ARN) of the custom routing 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 an accelerator's static IPv4 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.

If you have a dual-stack accelerator, you also have a second DNS name, DualStackDnsName, that points to both the A record and the AAAA record for all four static addresses for the accelerator: two IPv4 addresses and two IPv6 addresses.

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 IP address type that an accelerator supports. For a custom routing accelerator, the value must be IPV4.

Type: String

Valid Values: IPV4 | DUAL_STACK

Required: No

IpSets

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

Type: Array of [IpSet](#) 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 V2](#)
- [AWS SDK for Ruby V3](#)

CustomRoutingAcceleratorAttributes

Attributes of a custom routing 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:

DOC-EXAMPLE-BUCKET//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 V2](#)
- [AWS SDK for Ruby V3](#)

CustomRoutingDestinationConfiguration

For a custom routing accelerator, sets the port range and protocol for all endpoints (virtual private cloud subnets) in an endpoint group to accept client traffic on.

Contents

FromPort

The first port, inclusive, in the range of ports for the endpoint group that is associated with a custom routing accelerator.

Type: Integer

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

Required: Yes

Protocols

The protocol for the endpoint group that is associated with a custom routing accelerator. The protocol can be either TCP or UDP.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 2 items.

Valid Values: TCP | UDP

Required: Yes

ToPort

The last port, inclusive, in the range of ports for the endpoint group that is associated with a custom routing accelerator.

Type: Integer

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

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 V2](#)
- [AWS SDK for Ruby V3](#)

CustomRoutingDestinationDescription

For a custom routing accelerator, describes the port range and protocol for all endpoints (virtual private cloud subnets) in an endpoint group to accept client traffic on.

Contents

FromPort

The first port, inclusive, in the range of ports for the endpoint group that is associated with a custom routing accelerator.

Type: Integer

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

Required: No

Protocols

The protocol for the endpoint group that is associated with a custom routing accelerator. The protocol can be either TCP or UDP.

Type: Array of strings

Valid Values: TCP | UDP

Required: No

ToPort

The last port, inclusive, in the range of ports for the endpoint group that is associated with a custom routing accelerator.

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 V2](#)
- [AWS SDK for Ruby V3](#)

CustomRoutingEndpointConfiguration

The list of endpoint objects. For custom routing, this is a list of virtual private cloud (VPC) subnet IDs.

Contents

AttachmentArn

The Amazon Resource Name (ARN) of the cross-account attachment that specifies the endpoints (resources) that can be added to accelerators and principals that have permission to add the endpoints to accelerators.

Type: String

Length Constraints: Maximum length of 255.

Required: No

EndpointId

An ID for the endpoint. For custom routing accelerators, this is the virtual private cloud (VPC) subnet 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 V2](#)
- [AWS SDK for Ruby V3](#)

CustomRoutingEndpointDescription

A complex type for an endpoint for a custom routing accelerator. Each endpoint group can include one or more endpoints, which are virtual private cloud (VPC) subnets.

Contents

EndpointId

An ID for the endpoint. For custom routing accelerators, this is the virtual private cloud (VPC) subnet 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 V2](#)
- [AWS SDK for Ruby V3](#)

CustomRoutingEndpointGroup

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

Contents

DestinationDescriptions

For a custom routing accelerator, describes the port range and protocol for all endpoints (virtual private cloud subnets) in an endpoint group to accept client traffic on.

Type: Array of [CustomRoutingDestinationDescription](#) objects

Required: No

EndpointDescriptions

For a custom routing accelerator, describes the endpoints (virtual private cloud subnets) in an endpoint group to accept client traffic on.

Type: Array of [CustomRoutingEndpointDescription](#) 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

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 V2](#)
- [AWS SDK for Ruby V3](#)

CustomRoutingListener

A complex type for a listener for a custom routing accelerator.

Contents

ListenerArn

The Amazon Resource Name (ARN) of the listener.

Type: String

Length Constraints: Maximum length of 255.

Required: No

PortRanges

The port range to support for connections from clients to your accelerator.

Separately, you set port ranges for endpoints. For more information, see [About endpoints for custom routing accelerators](#).

Type: Array of [PortRange](#) objects

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

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 V2](#)
- [AWS SDK for Ruby V3](#)

DestinationPortMapping

The port mappings for a specified endpoint IP address (destination).

Contents

AcceleratorArn

The Amazon Resource Name (ARN) of the custom routing accelerator that you have port mappings for.

Type: String

Length Constraints: Maximum length of 255.

Required: No

AcceleratorSocketAddresses

The IP address/port combinations (sockets) that map to a given destination socket address.

Type: Array of [SocketAddress](#) objects

Required: No

DestinationSocketAddress

The endpoint IP address/port combination for traffic received on the accelerator socket address.

Type: [SocketAddress](#) object

Required: No

DestinationTrafficState

Indicates whether or not a port mapping destination can receive traffic. The value is either ALLOW, if traffic is allowed to the destination, or DENY, if traffic is not allowed to the destination.

Type: String

Valid Values: ALLOW | DENY

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 for the endpoint group.

Type: String

Length Constraints: Maximum length of 255.

Required: No

EndpointId

The ID for the virtual private cloud (VPC) subnet.

Type: String

Length Constraints: Maximum length of 255.

Required: No

IpAddressType

The IP address type that an accelerator supports. For a custom routing accelerator, the value must be IPV4.

Type: String

Valid Values: IPV4 | DUAL_STACK

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 V2](#)
- [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

AttachmentArn

The Amazon Resource Name (ARN) of the cross-account attachment that specifies the endpoints (resources) that can be added to accelerators and principals that have permission to add the endpoints to accelerators.

Type: String

Length Constraints: Maximum length of 255.

Required: No

ClientIPPreservationEnabled

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

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 endpoint fronted by the accelerator.

Client IP address preservation is supported, in specific AWS Regions, for endpoints that are Application Load Balancers, Amazon EC2 instances, and Network Load Balancers with security groups. **IMPORTANT:** You cannot use client IP address preservation with Network Load Balancers with TLS listeners.

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.

For cross-account endpoints, this must be the ARN of the resource.

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 V2](#)
- [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 endpoint. The value is true or false. The default value is true for Application Load Balancers endpoints.

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 endpoint fronted by the accelerator.

Client IP address preservation is supported, in specific AWS Regions, for endpoints that are Application Load Balancers, Amazon EC2 instances, and Network Load Balancers with security groups. **IMPORTANT:** You cannot use client IP address preservation with Network Load Balancers with TLS listeners.

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.

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

Type: String

Length Constraints: Maximum length of 255.

Required: No

HealthReason

Returns a null result.

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 V2](#)
- [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](#) 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.

Pattern: ^/[-a-zA-Z0-9@:%_\\".~#?&/=]*\$

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

PortOverrides

Allows you to override the destination ports used to route traffic to an endpoint. Using a port override lets you map a list of external destination ports (that your users send traffic to) to a list of internal destination ports that you want an application endpoint to receive traffic on.

Type: Array of [PortOverride](#) objects

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

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 V2](#)
- [AWS SDK for Ruby V3](#)

EndpointIdentifier

A complex type for an endpoint. Specifies information about the endpoint to remove from the endpoint group.

Contents

EndpointId

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

ClientIPPreservationEnabled

Indicates whether client IP address preservation is enabled for an endpoint. The value is true or false.

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 endpoint fronted by the accelerator.

Type: Boolean

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 V2](#)

- [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.

Length Constraints: Maximum length of 45.

Required: No

IpAddressFamily

The types of IP addresses included in this IP set.

Type: String

Valid Values: IPv4 | IPv6

Required: No

IpFamily

IpFamily is deprecated and has been replaced by IpAddressFamily.

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 V2](#)
- [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](#) 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 V2](#)
- [AWS SDK for Ruby V3](#)

PortMapping

Returns the ports and associated IP addresses and ports of Amazon EC2 instances in your virtual private cloud (VPC) subnets. Custom routing is a port mapping protocol in AWS Global Accelerator that statically associates port ranges with VPC subnets, which allows Global Accelerator to route to specific instances and ports within one or more subnets.

Contents

AcceleratorPort

The accelerator port.

Type: Integer

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

Required: No

DestinationSocketAddress

The EC2 instance IP address and port number in the virtual private cloud (VPC) subnet.

Type: [SocketAddress](#) object

Required: No

DestinationTrafficState

Indicates whether or not a port mapping destination can receive traffic. The value is either ALLOW, if traffic is allowed to the destination, or DENY, if traffic is not allowed to the destination.

Type: String

Valid Values: ALLOW | DENY

Required: No

EndpointGroupArn

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

Type: String

Length Constraints: Maximum length of 255.

Required: No

EndpointId

The IP address of the VPC subnet (the subnet ID).

Type: String

Length Constraints: Maximum length of 255.

Required: No

Protocols

The protocols supported by the endpoint group.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 2 items.

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 V2](#)
- [AWS SDK for Ruby V3](#)

PortOverride

Override specific listener ports used to route traffic to endpoints that are part of an endpoint group. For example, you can create a port override in which the listener receives user traffic on ports 80 and 443, but your accelerator routes that traffic to ports 1080 and 1443, respectively, on the endpoints.

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

Contents

EndpointPort

The endpoint port that you want a listener port to be mapped to. This is the port on the endpoint, such as the Application Load Balancer or Amazon EC2 instance.

Type: Integer

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

Required: No

ListenerPort

The listener port that you want to map to a specific endpoint port. This is the port that user traffic arrives to the Global Accelerator on.

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 V2](#)

- [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 V2](#)
- [AWS SDK for Ruby V3](#)

Resource

An AWS resource that is supported by AWS Global Accelerator and can be added as an endpoint for an accelerator.

Contents

EndpointId

The endpoint ID for the endpoint (AWS resource).

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Region

The AWS Region where a resource is located.

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 V2](#)
- [AWS SDK for Ruby V3](#)

SocketAddress

An IP address/port combination.

Contents

IpAddress

The IP address for the socket address.

Type: String

Length Constraints: Maximum length of 255.

Required: No

Port

The port for the socket address.

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 V2](#)
- [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 V2](#)
- [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 [Signing AWS API requests in the IAM User Guide](#).

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 [Create a signed AWS API request](#) in the *IAM User Guide*.

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'T'HHMMSS'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 [Elements of an AWS API request signature](#) in the *IAM User Guide*.

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 STS, see [AWS services that work with IAM](#) in the *IAM User Guide*.

Condition: If you're using temporary security credentials from AWS STS, 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 [Create a signed AWS API request](#) in the *IAM User Guide*.

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

NotAuthorized

You do not have permission to perform this action.

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

ValidationException

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

HTTP Status Code: 400