



API Reference

# AWS App Mesh



**API Version 2019-01-25**

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

# AWS App Mesh: API Reference

Copyright © 2026 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

<b>Welcome</b> .....	<b>1</b>
<b>Actions</b> .....	<b>2</b>
CreateGatewayRoute .....	4
Request Syntax .....	4
URI Request Parameters .....	7
Request Body .....	8
Response Syntax .....	9
Response Elements .....	12
Errors .....	14
Examples .....	15
See Also .....	17
CreateMesh .....	18
Request Syntax .....	18
URI Request Parameters .....	18
Request Body .....	18
Response Syntax .....	19
Response Elements .....	20
Errors .....	21
Examples .....	22
See Also .....	23
CreateRoute .....	25
Request Syntax .....	25
URI Request Parameters .....	29
Request Body .....	30
Response Syntax .....	31
Response Elements .....	35
Errors .....	36
Examples .....	38
See Also .....	40
CreateVirtualGateway .....	41
Request Syntax .....	41
URI Request Parameters .....	42
Request Body .....	43
Response Syntax .....	44

Response Elements .....	46
Errors .....	46
Examples .....	48
See Also .....	50
CreateVirtualNode .....	51
Request Syntax .....	51
URI Request Parameters .....	53
Request Body .....	54
Response Syntax .....	55
Response Elements .....	57
Errors .....	58
Examples .....	59
See Also .....	61
CreateVirtualRouter .....	63
Request Syntax .....	63
URI Request Parameters .....	63
Request Body .....	64
Response Syntax .....	65
Response Elements .....	66
Errors .....	67
Examples .....	68
See Also .....	70
CreateVirtualService .....	71
Request Syntax .....	71
URI Request Parameters .....	71
Request Body .....	72
Response Syntax .....	73
Response Elements .....	73
Errors .....	74
Examples .....	75
See Also .....	77
DeleteGatewayRoute .....	78
Request Syntax .....	78
URI Request Parameters .....	78
Request Body .....	79
Response Syntax .....	79

---

Response Elements .....	82
Errors .....	83
Examples .....	84
See Also .....	86
DeleteMesh .....	87
Request Syntax .....	87
URI Request Parameters .....	87
Request Body .....	87
Response Syntax .....	87
Response Elements .....	88
Errors .....	89
Examples .....	90
See Also .....	91
DeleteRoute .....	92
Request Syntax .....	92
URI Request Parameters .....	92
Request Body .....	93
Response Syntax .....	93
Response Elements .....	97
Errors .....	98
Examples .....	99
See Also .....	101
DeleteVirtualGateway .....	102
Request Syntax .....	102
URI Request Parameters .....	102
Request Body .....	102
Response Syntax .....	103
Response Elements .....	104
Errors .....	105
Examples .....	106
See Also .....	108
DeleteVirtualNode .....	109
Request Syntax .....	109
URI Request Parameters .....	109
Request Body .....	109
Response Syntax .....	110

---

Response Elements .....	112
Errors .....	113
Examples .....	114
See Also .....	115
DeleteVirtualRouter .....	117
Request Syntax .....	117
URI Request Parameters .....	117
Request Body .....	117
Response Syntax .....	118
Response Elements .....	118
Errors .....	119
Examples .....	120
See Also .....	122
DeleteVirtualService .....	123
Request Syntax .....	123
URI Request Parameters .....	123
Request Body .....	123
Response Syntax .....	123
Response Elements .....	124
Errors .....	125
Examples .....	126
See Also .....	127
DescribeGatewayRoute .....	129
Request Syntax .....	129
URI Request Parameters .....	129
Request Body .....	130
Response Syntax .....	130
Response Elements .....	133
Errors .....	134
Examples .....	135
See Also .....	137
DescribeMesh .....	138
Request Syntax .....	138
URI Request Parameters .....	138
Request Body .....	138
Response Syntax .....	138

Response Elements .....	139
Errors .....	140
Examples .....	141
See Also .....	142
<b>DescribeRoute .....</b>	<b>143</b>
Request Syntax .....	143
URI Request Parameters .....	143
Request Body .....	144
Response Syntax .....	144
Response Elements .....	148
Errors .....	149
Examples .....	150
See Also .....	152
<b>DescribeVirtualGateway .....</b>	<b>153</b>
Request Syntax .....	153
URI Request Parameters .....	153
Request Body .....	153
Response Syntax .....	154
Response Elements .....	155
Errors .....	156
Examples .....	157
See Also .....	159
<b>DescribeVirtualNode .....</b>	<b>160</b>
Request Syntax .....	160
URI Request Parameters .....	160
Request Body .....	160
Response Syntax .....	161
Response Elements .....	163
Errors .....	164
Examples .....	164
See Also .....	166
<b>DescribeVirtualRouter .....</b>	<b>168</b>
Request Syntax .....	168
URI Request Parameters .....	168
Request Body .....	168
Response Syntax .....	169

Response Elements .....	169
Errors .....	170
Examples .....	171
See Also .....	172
<b>DescribeVirtualService .....</b>	<b>174</b>
Request Syntax .....	174
URI Request Parameters .....	174
Request Body .....	174
Response Syntax .....	174
Response Elements .....	175
Errors .....	176
Examples .....	177
See Also .....	178
<b>ListGatewayRoutes .....</b>	<b>180</b>
Request Syntax .....	180
URI Request Parameters .....	180
Request Body .....	181
Response Syntax .....	181
Response Elements .....	182
Errors .....	182
Examples .....	183
See Also .....	184
<b>ListMeshes .....</b>	<b>186</b>
Request Syntax .....	186
URI Request Parameters .....	186
Request Body .....	186
Response Syntax .....	187
Response Elements .....	187
Errors .....	188
Examples .....	188
See Also .....	190
<b>ListRoutes .....</b>	<b>191</b>
Request Syntax .....	191
URI Request Parameters .....	191
Request Body .....	192
Response Syntax .....	192

---

Response Elements .....	193
Errors .....	193
Examples .....	194
See Also .....	195
ListTagsForResource .....	196
Request Syntax .....	196
URI Request Parameters .....	196
Request Body .....	196
Response Syntax .....	197
Response Elements .....	197
Errors .....	197
See Also .....	198
ListVirtualGateways .....	200
Request Syntax .....	200
URI Request Parameters .....	200
Request Body .....	201
Response Syntax .....	201
Response Elements .....	201
Errors .....	202
Examples .....	203
See Also .....	204
ListVirtualNodes .....	205
Request Syntax .....	205
URI Request Parameters .....	205
Request Body .....	206
Response Syntax .....	206
Response Elements .....	206
Errors .....	207
Examples .....	208
See Also .....	210
ListVirtualRouters .....	211
Request Syntax .....	211
URI Request Parameters .....	211
Request Body .....	212
Response Syntax .....	212
Response Elements .....	212

Errors .....	213
Examples .....	214
See Also .....	215
ListVirtualServices .....	216
Request Syntax .....	216
URI Request Parameters .....	216
Request Body .....	217
Response Syntax .....	217
Response Elements .....	217
Errors .....	218
Examples .....	219
See Also .....	220
TagResource .....	222
Request Syntax .....	222
URI Request Parameters .....	222
Request Body .....	222
Response Syntax .....	223
Response Elements .....	223
Errors .....	223
See Also .....	224
UntagResource .....	225
Request Syntax .....	225
URI Request Parameters .....	225
Request Body .....	225
Response Syntax .....	226
Response Elements .....	226
Errors .....	226
See Also .....	227
UpdateGatewayRoute .....	228
Request Syntax .....	228
URI Request Parameters .....	231
Request Body .....	232
Response Syntax .....	232
Response Elements .....	236
Errors .....	237
Examples .....	238

See Also .....	240
UpdateMesh .....	241
Request Syntax .....	241
URI Request Parameters .....	241
Request Body .....	241
Response Syntax .....	242
Response Elements .....	243
Errors .....	243
See Also .....	244
UpdateRoute .....	246
Request Syntax .....	246
URI Request Parameters .....	250
Request Body .....	251
Response Syntax .....	251
Response Elements .....	256
Errors .....	257
Examples .....	258
See Also .....	260
UpdateVirtualGateway .....	262
Request Syntax .....	262
URI Request Parameters .....	263
Request Body .....	264
Response Syntax .....	264
Response Elements .....	266
Errors .....	267
Examples .....	268
See Also .....	270
UpdateVirtualNode .....	271
Request Syntax .....	271
URI Request Parameters .....	272
Request Body .....	273
Response Syntax .....	274
Response Elements .....	276
Errors .....	277
Examples .....	278
See Also .....	280

UpdateVirtualRouter .....	282
Request Syntax .....	282
URI Request Parameters .....	282
Request Body .....	283
Response Syntax .....	283
Response Elements .....	284
Errors .....	285
Examples .....	286
See Also .....	288
UpdateVirtualService .....	289
Request Syntax .....	289
URI Request Parameters .....	289
Request Body .....	290
Response Syntax .....	290
Response Elements .....	291
Errors .....	291
Examples .....	293
See Also .....	294
<b>Data Types .....</b>	<b>296</b>
AccessLog .....	302
Contents .....	302
See Also .....	302
AwsCloudMapInstanceAttribute .....	303
Contents .....	303
See Also .....	303
AwsCloudMapServiceDiscovery .....	305
Contents .....	305
See Also .....	306
Backend .....	307
Contents .....	307
See Also .....	307
BackendDefaults .....	308
Contents .....	308
See Also .....	308
ClientPolicy .....	309
Contents .....	309

See Also .....	309
ClientPolicyTls .....	310
Contents .....	310
See Also .....	310
ClientTlsCertificate .....	312
Contents .....	312
See Also .....	312
DnsServiceDiscovery .....	313
Contents .....	313
See Also .....	313
Duration .....	315
Contents .....	315
See Also .....	315
EgressFilter .....	316
Contents .....	316
See Also .....	316
FileAccessLog .....	317
Contents .....	317
See Also .....	317
GatewayRouteData .....	319
Contents .....	319
See Also .....	320
GatewayRouteHostnameMatch .....	321
Contents .....	321
See Also .....	321
GatewayRouteHostnameRewrite .....	322
Contents .....	322
See Also .....	322
GatewayRouteRef .....	323
Contents .....	323
See Also .....	325
GatewayRouteSpec .....	326
Contents .....	326
See Also .....	326
GatewayRouteStatus .....	328
Contents .....	328

---

See Also .....	328
GatewayRouteTarget .....	329
Contents .....	329
See Also .....	329
GatewayRouteVirtualService .....	330
Contents .....	330
See Also .....	330
GrpcGatewayRoute .....	331
Contents .....	331
See Also .....	331
GrpcGatewayRouteAction .....	332
Contents .....	332
See Also .....	332
GrpcGatewayRouteMatch .....	333
Contents .....	333
See Also .....	333
GrpcGatewayRouteMetadata .....	335
Contents .....	335
See Also .....	335
GrpcGatewayRouteRewrite .....	337
Contents .....	337
See Also .....	337
GrpcMetadataMatchMethod .....	338
Contents .....	338
See Also .....	339
GrpcRetryPolicy .....	340
Contents .....	340
See Also .....	341
GrpcRoute .....	342
Contents .....	342
See Also .....	342
GrpcRouteAction .....	344
Contents .....	344
See Also .....	344
GrpcRouteMatch .....	345
Contents .....	345

See Also .....	346
GrpcRouteMetadata .....	347
Contents .....	347
See Also .....	347
GrpcRouteMetadataMatchMethod .....	349
Contents .....	349
See Also .....	350
GrpcTimeout .....	351
Contents .....	351
See Also .....	351
HeaderMatchMethod .....	352
Contents .....	352
See Also .....	353
HealthCheckPolicy .....	354
Contents .....	354
See Also .....	355
HttpGatewayRoute .....	356
Contents .....	356
See Also .....	356
HttpGatewayRouteAction .....	357
Contents .....	357
See Also .....	357
HttpGatewayRouteHeader .....	358
Contents .....	358
See Also .....	358
HttpGatewayRouteMatch .....	360
Contents .....	360
See Also .....	361
HttpGatewayRoutePathRewrite .....	362
Contents .....	362
See Also .....	362
HttpGatewayRoutePrefixRewrite .....	363
Contents .....	363
See Also .....	363
HttpGatewayRouteRewrite .....	364
Contents .....	364

See Also .....	364
HttpPathMatch .....	365
Contents .....	365
See Also .....	365
HttpQueryParameter .....	366
Contents .....	366
See Also .....	366
HttpRetryPolicy .....	367
Contents .....	367
See Also .....	368
HttpRoute .....	369
Contents .....	369
See Also .....	369
HttpRouteAction .....	371
Contents .....	371
See Also .....	371
HttpRouteHeader .....	372
Contents .....	372
See Also .....	372
HttpRouteMatch .....	374
Contents .....	374
See Also .....	375
HttpTimeout .....	376
Contents .....	376
See Also .....	376
JsonFormatRef .....	377
Contents .....	377
See Also .....	377
Listener .....	378
Contents .....	378
See Also .....	379
ListenerTimeout .....	380
Contents .....	380
See Also .....	381
ListenerTls .....	382
Contents .....	382

See Also .....	382
ListenerTlsAcmeCertificate .....	384
Contents .....	384
See Also .....	384
ListenerTlsCertificate .....	385
Contents .....	385
See Also .....	385
ListenerTlsFileCertificate .....	387
Contents .....	387
See Also .....	387
ListenerTlsSdsCertificate .....	388
Contents .....	388
See Also .....	388
ListenerTlsValidationContext .....	389
Contents .....	389
See Also .....	389
ListenerTlsValidationContextTrust .....	390
Contents .....	390
See Also .....	390
Logging .....	391
Contents .....	391
See Also .....	391
LoggingFormat .....	392
Contents .....	392
See Also .....	392
MatchRange .....	393
Contents .....	393
See Also .....	393
MeshData .....	394
Contents .....	394
See Also .....	394
MeshRef .....	396
Contents .....	396
See Also .....	397
MeshServiceDiscovery .....	398
Contents .....	398

See Also .....	398
MeshSpec .....	399
Contents .....	399
See Also .....	399
MeshStatus .....	400
Contents .....	400
See Also .....	400
OutlierDetection .....	401
Contents .....	401
See Also .....	402
PortMapping .....	403
Contents .....	403
See Also .....	403
QueryParamMatch .....	404
Contents .....	404
See Also .....	404
ResourceMetadata .....	405
Contents .....	405
See Also .....	406
RouteData .....	407
Contents .....	407
See Also .....	408
RouteRef .....	409
Contents .....	409
See Also .....	411
RouteSpec .....	412
Contents .....	412
See Also .....	413
RouteStatus .....	414
Contents .....	414
See Also .....	414
ServiceDiscovery .....	415
Contents .....	415
See Also .....	415
SubjectAlternativeNameMatchers .....	416
Contents .....	416

See Also .....	416
SubjectAlternativeNames .....	417
Contents .....	417
See Also .....	417
TagRef .....	418
Contents .....	418
See Also .....	418
TcpRoute .....	419
Contents .....	419
See Also .....	419
TcpRouteAction .....	420
Contents .....	420
See Also .....	420
TcpRouteMatch .....	421
Contents .....	421
See Also .....	421
TcpTimeout .....	422
Contents .....	422
See Also .....	422
TlsValidationContext .....	423
Contents .....	423
See Also .....	423
TlsValidationContextAcmTrust .....	425
Contents .....	425
See Also .....	425
TlsValidationContextFileTrust .....	426
Contents .....	426
See Also .....	426
TlsValidationContextSdsTrust .....	427
Contents .....	427
See Also .....	427
TlsValidationContextTrust .....	428
Contents .....	428
See Also .....	428
VirtualGatewayAccessLog .....	430
Contents .....	430

See Also .....	430
VirtualGatewayBackendDefaults .....	431
Contents .....	431
See Also .....	431
VirtualGatewayClientPolicy .....	432
Contents .....	432
See Also .....	432
VirtualGatewayClientPolicyTls .....	433
Contents .....	433
See Also .....	434
VirtualGatewayClientTlsCertificate .....	435
Contents .....	435
See Also .....	435
VirtualGatewayConnectionPool .....	437
Contents .....	437
See Also .....	438
VirtualGatewayData .....	439
Contents .....	439
See Also .....	440
VirtualGatewayFileAccessLog .....	441
Contents .....	441
See Also .....	441
VirtualGatewayGrpcConnectionPool .....	442
Contents .....	442
See Also .....	442
VirtualGatewayHealthCheckPolicy .....	443
Contents .....	443
See Also .....	444
VirtualGatewayHttp2ConnectionPool .....	445
Contents .....	445
See Also .....	445
VirtualGatewayHttpConnectionPool .....	446
Contents .....	446
See Also .....	446
VirtualGatewayListener .....	447
Contents .....	447

See Also .....	447
VirtualGatewayListenerTls .....	449
Contents .....	449
See Also .....	449
VirtualGatewayListenerTlsAcmCertificate .....	451
Contents .....	451
See Also .....	451
VirtualGatewayListenerTlsCertificate .....	452
Contents .....	452
See Also .....	452
VirtualGatewayListenerTlsFileCertificate .....	454
Contents .....	454
See Also .....	454
VirtualGatewayListenerTlsSdsCertificate .....	455
Contents .....	455
See Also .....	455
VirtualGatewayListenerTlsValidationContext .....	456
Contents .....	456
See Also .....	456
VirtualGatewayListenerTlsValidationContextTrust .....	457
Contents .....	457
See Also .....	457
VirtualGatewayLogging .....	459
Contents .....	459
See Also .....	459
VirtualGatewayPortMapping .....	460
Contents .....	460
See Also .....	460
VirtualGatewayRef .....	461
Contents .....	461
See Also .....	462
VirtualGatewaySpec .....	464
Contents .....	464
See Also .....	464
VirtualGatewayStatus .....	465
Contents .....	465

See Also .....	465
VirtualGatewayTlsValidationContext .....	466
Contents .....	466
See Also .....	466
VirtualGatewayTlsValidationContextAcmTrust .....	467
Contents .....	467
See Also .....	467
VirtualGatewayTlsValidationContextFileTrust .....	468
Contents .....	468
See Also .....	468
VirtualGatewayTlsValidationContextSdsTrust .....	469
Contents .....	469
See Also .....	469
VirtualGatewayTlsValidationContextTrust .....	470
Contents .....	470
See Also .....	470
VirtualNodeConnectionPool .....	472
Contents .....	472
See Also .....	473
VirtualNodeData .....	474
Contents .....	474
See Also .....	475
VirtualNodeGrpcConnectionPool .....	476
Contents .....	476
See Also .....	476
VirtualNodeHttp2ConnectionPool .....	477
Contents .....	477
See Also .....	477
VirtualNodeHttpConnectionPool .....	478
Contents .....	478
See Also .....	478
VirtualNodeRef .....	479
Contents .....	479
See Also .....	480
VirtualNodeServiceProvider .....	482
Contents .....	482

See Also .....	482
VirtualNodeSpec .....	483
Contents .....	483
See Also .....	484
VirtualNodeStatus .....	485
Contents .....	485
See Also .....	485
VirtualNodeTcpConnectionPool .....	486
Contents .....	486
See Also .....	486
VirtualRouterData .....	487
Contents .....	487
See Also .....	488
VirtualRouterListener .....	489
Contents .....	489
See Also .....	489
VirtualRouterRef .....	490
Contents .....	490
See Also .....	491
VirtualRouterServiceProvider .....	493
Contents .....	493
See Also .....	493
VirtualRouterSpec .....	494
Contents .....	494
See Also .....	494
VirtualRouterStatus .....	495
Contents .....	495
See Also .....	495
VirtualServiceBackend .....	496
Contents .....	496
See Also .....	496
VirtualServiceData .....	497
Contents .....	497
See Also .....	498
VirtualServiceProvider .....	499
Contents .....	499

---

See Also .....	499
VirtualServiceRef .....	500
Contents .....	500
See Also .....	501
VirtualServiceSpec .....	503
Contents .....	503
See Also .....	503
VirtualServiceStatus .....	504
Contents .....	504
See Also .....	504
WeightedTarget .....	505
Contents .....	505
See Also .....	505

# Welcome

AWS App Mesh is a service mesh based on the Envoy proxy that makes it easy to monitor and control microservices. App Mesh standardizes how your microservices communicate, giving you end-to-end visibility and helping to ensure high availability for your applications.

App Mesh gives you consistent visibility and network traffic controls for every microservice in an application. You can use App Mesh with AWS Fargate, Amazon ECS, Amazon EKS, Kubernetes on AWS, and Amazon EC2.

## Note

App Mesh supports microservice applications that use service discovery naming for their components. For more information about service discovery on Amazon ECS, see [Service Discovery](#) in the *Amazon Elastic Container Service Developer Guide*. Kubernetes kube-dns and coreDNS are supported. For more information, see [DNS for Services and Pods](#) in the Kubernetes documentation.

This document was last published on May 21, 2026.

# Actions

The following actions are supported:

- [CreateGatewayRoute](#)
- [CreateMesh](#)
- [CreateRoute](#)
- [CreateVirtualGateway](#)
- [CreateVirtualNode](#)
- [CreateVirtualRouter](#)
- [CreateVirtualService](#)
- [DeleteGatewayRoute](#)
- [DeleteMesh](#)
- [DeleteRoute](#)
- [DeleteVirtualGateway](#)
- [DeleteVirtualNode](#)
- [DeleteVirtualRouter](#)
- [DeleteVirtualService](#)
- [DescribeGatewayRoute](#)
- [DescribeMesh](#)
- [DescribeRoute](#)
- [DescribeVirtualGateway](#)
- [DescribeVirtualNode](#)
- [DescribeVirtualRouter](#)
- [DescribeVirtualService](#)
- [ListGatewayRoutes](#)
- [ListMeshes](#)
- [ListRoutes](#)
- [ListTagsForResource](#)
- [ListVirtualGateways](#)
- [ListVirtualNodes](#)

- [ListVirtualRouters](#)
- [ListVirtualServices](#)
- [TagResource](#)
- [UntagResource](#)
- [UpdateGatewayRoute](#)
- [UpdateMesh](#)
- [UpdateRoute](#)
- [UpdateVirtualGateway](#)
- [UpdateVirtualNode](#)
- [UpdateVirtualRouter](#)
- [UpdateVirtualService](#)

# CreateGatewayRoute

Creates a gateway route.

A gateway route is attached to a virtual gateway and routes traffic to an existing virtual service. If a route matches a request, it can distribute traffic to a target virtual service.

For more information about gateway routes, see [Gateway routes](#).

## Request Syntax

```
PUT /v20190125/meshes/meshName/virtualGateway/virtualGatewayName/gatewayRoutes?  
meshOwner=meshOwner HTTP/1.1  
Content-type: application/json
```

```
{  
  "clientToken": "string",  
  "gatewayRouteName": "string",  
  "spec": {  
    "grpcRoute": {  
      "action": {  
        "rewrite": {  
          "hostname": {  
            "defaultTargetHostname": "string"  
          }  
        },  
      },  
      "target": {  
        "port": number,  
        "virtualService": {  
          "virtualServiceName": "string"  
        }  
      }  
    }  
  },  
  "match": {  
    "hostname": {  
      "exact": "string",  
      "suffix": "string"  
    },  
    "metadata": [  
      {  
        "invert": boolean,  
        "match": { ... },  
        "name": "string"  
      }  
    ]  
  }  
}
```

```
    }
  ],
  "port": number,
  "serviceName": "string"
}
},
"http2Route": {
  "action": {
    "rewrite": {
      "hostname": {
        "defaultTargetHostname": "string"
      },
      "path": {
        "exact": "string"
      },
      "prefix": {
        "defaultPrefix": "string",
        "value": "string"
      }
    },
    "target": {
      "port": number,
      "virtualService": {
        "virtualServiceName": "string"
      }
    }
  },
  "match": {
    "headers": [
      {
        "invert": boolean,
        "match": { ... },
        "name": "string"
      }
    ],
    "hostname": {
      "exact": "string",
      "suffix": "string"
    },
    "method": "string",
    "path": {
      "exact": "string",
      "regex": "string"
    }
  },
}
```

```

    "port": number,
    "prefix": "string",
    "queryParameters": [
      {
        "match": {
          "exact": "string"
        },
        "name": "string"
      }
    ]
  },
  "httpRoute": {
    "action": {
      "rewrite": {
        "hostname": {
          "defaultTargetHostname": "string"
        },
        "path": {
          "exact": "string"
        },
        "prefix": {
          "defaultPrefix": "string",
          "value": "string"
        }
      },
      "target": {
        "port": number,
        "virtualService": {
          "virtualServiceName": "string"
        }
      }
    },
    "match": {
      "headers": [
        {
          "invert": boolean,
          "match": { ... },
          "name": "string"
        }
      ],
      "hostname": {
        "exact": "string",
        "suffix": "string"
      }
    }
  }
}

```

```
    },
    "method": "string",
    "path": {
      "exact": "string",
      "regex": "string"
    },
    "port": number,
    "prefix": "string",
    "queryParameters": [
      {
        "match": {
          "exact": "string"
        },
        "name": "string"
      }
    ]
  },
  "priority": number
},
"tags": [
  {
    "key": "string",
    "value": "string"
  }
]
}
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh to create the gateway route in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then the account that you specify must share the mesh with your account before you can create

the resource in the service mesh. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### [virtualGatewayName](#)

The name of the virtual gateway to associate the gateway route with. If the virtual gateway is in a shared mesh, then you must be the owner of the virtual gateway resource.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request accepts the following data in JSON format.

### [clientToken](#)

Unique, case-sensitive identifier that you provide to ensure the idempotency of the request. Up to 36 letters, numbers, hyphens, and underscores are allowed.

Type: String

Required: No

### [gatewayRouteName](#)

The name to use for the gateway route.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### [spec](#)

The gateway route specification to apply.

Type: [GatewayRouteSpec](#) object

Required: Yes

## tags

Optional metadata that you can apply to the gateway route to assist with categorization and organization. Each tag consists of a key and an optional value, both of which you define. Tag keys can have a maximum character length of 128 characters, and tag values can have a maximum length of 256 characters.

Type: Array of [TagRef](#) objects

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Required: No

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "gatewayRouteName": "string",
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "grpcRoute": {
      "action": {
        "rewrite": {
          "hostname": {
            "defaultTargetHostname": "string"
          }
        }
      },
      "target": {
        "port": number,
        "virtualService": {
          "virtualServiceName": "string"
        }
      }
    }
  }
}
```

```

    }
  }
},
"match": {
  "hostname": {
    "exact": "string",
    "suffix": "string"
  },
  "metadata": [
    {
      "invert": boolean,
      "match": { ... },
      "name": "string"
    }
  ],
  "port": number,
  "serviceName": "string"
}
},
"http2Route": {
  "action": {
    "rewrite": {
      "hostname": {
        "defaultTargetHostname": "string"
      },
      "path": {
        "exact": "string"
      },
      "prefix": {
        "defaultPrefix": "string",
        "value": "string"
      }
    },
    "target": {
      "port": number,
      "virtualService": {
        "virtualServiceName": "string"
      }
    }
  },
  "match": {
    "headers": [
      {
        "invert": boolean,

```

```

        "match": { ... },
        "name": "string"
    }
],
"hostname": {
    "exact": "string",
    "suffix": "string"
},
"method": "string",
"path": {
    "exact": "string",
    "regex": "string"
},
"port": number,
"prefix": "string",
"queryParameters": [
    {
        "match": {
            "exact": "string"
        },
        "name": "string"
    }
]
},
"httpRoute": {
    "action": {
        "rewrite": {
            "hostname": {
                "defaultTargetHostname": "string"
            },
            "path": {
                "exact": "string"
            },
            "prefix": {
                "defaultPrefix": "string",
                "value": "string"
            }
        },
        "target": {
            "port": number,
            "virtualService": {
                "virtualServiceName": "string"
            }
        }
    }
}

```

```

    }
  },
  "match": {
    "headers": [
      {
        "invert": boolean,
        "match": { ... },
        "name": "string"
      }
    ],
    "hostname": {
      "exact": "string",
      "suffix": "string"
    },
    "method": "string",
    "path": {
      "exact": "string",
      "regex": "string"
    },
    "port": number,
    "prefix": "string",
    "queryParameters": [
      {
        "match": {
          "exact": "string"
        },
        "name": "string"
      }
    ]
  }
},
"priority": number
},
"status": {
  "status": "string"
},
"virtualGatewayName": "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.

### gatewayRouteName

The name of the gateway route.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### meshName

The name of the service mesh that the resource resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

An object that represents metadata for a resource.

Type: [ResourceMetadata](#) object

### spec

The specifications of the gateway route.

Type: [GatewayRouteSpec](#) object

### status

The status of the gateway route.

Type: [GatewayRouteStatus](#) object

### virtualGatewayName

The virtual gateway that the gateway route is associated with.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### BadRequestException

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### ConflictException

The request contains a client token that was used for a previous update resource call with different specifications. Try the request again with a new client token.

HTTP Status Code: 409

### ForbiddenException

You don't have permissions to perform this action.

HTTP Status Code: 403

### InternalServerErrorException

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

HTTP Status Code: 500

### LimitExceededException

You have exceeded a service limit for your account. For more information, see [Service Limits](#) in the *AWS App Mesh User Guide*.

HTTP Status Code: 400

### NotFoundException

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### ServiceUnavailableException

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### Example

The following example creates an HTTP/2 gateway route named `myGatewayRoute` that is associated to a virtual gateway named `myVirtualGateway` in the `apps` service mesh. The route routes all traffic to the virtual service named `myservicea.svc.cluster.local`.

### Sample Request

```
PUT /v20190125/meshes/apps/virtualGateways HTTP/1.1
Host: appmesh.us-west-2.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.298 Python/3.6.0 Windows/10 botocore/1.13.34
X-Amz-Date: 20200608T194401Z
Authorization: AUTHPARAMS
```

```
{
  "gatewayRouteName": "myGatewayRoute",
  "spec": {
    "http2Route": {
      "action": {
        "target": {
          "virtualService": {
```

```
    "virtualServiceName": "myservicea.svc.cluster.local"
  }
}
},
"match": {
  "prefix": "/"
}
}
},
"clientToken": "1aa1111a-1111-1111-111a-a1a1a1aa111a"
}
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 696b569b-3593-446d-832c-0a38ce1c53ac
content-type: application/json
content-length: 517
date: Mon, 08 Jun 2020 19:44:02 GMT
x-envoy-upstream-service-time: 48
server: envoy
Connection: keep-alive

{
  "gatewayRouteName": "myGatewayRoute",
  "meshName": "myApps",
  "metadata": {
    "arn": "arn:aws:appmesh:us-west-2:123456789012:mesh/apps/virtualGateway/myVirtualGateway/gatewayRoute/myGatewayRoute",
    "createdAt": 1.591642091122E9,
    "lastUpdatedAt": 1.591642091122E9,
    "meshOwner": "123456789012",
    "resourceOwner": "123456789012",
    "uid": "ff111ff2-33ff-4ff4-ffff-111f111f111f",
    "version": 1
  },
  "spec": {
    "grpcRoute": null,
    "http2Route": {
      "action": {
        "target": {
          "virtualService": {
            "virtualServiceName": "myservicea.svc.cluster.local"
          }
        }
      }
    }
  }
}
```

```
    }
  }
},
"match": {
  "prefix": "/"
}
},
"httpRoute": null
},
"status": {
  "status": "ACTIVE"
},
"virtualGatewayName": "myVirtualGateway"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# CreateMesh

Creates a service mesh.

A service mesh is a logical boundary for network traffic between services that are represented by resources within the mesh. After you create your service mesh, you can create virtual services, virtual nodes, virtual routers, and routes to distribute traffic between the applications in your mesh.

For more information about service meshes, see [Service meshes](#).

## Request Syntax

```
PUT /v20190125/meshes HTTP/1.1
Content-type: application/json

{
  "clientToken": "string",
  "meshName": "string",
  "spec": {
    "egressFilter": {
      "type": "string"
    },
    "serviceDiscovery": {
      "ipPreference": "string"
    }
  },
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ]
}
```

## URI Request Parameters

The request does not use any URI parameters.

## Request Body

The request accepts the following data in JSON format.

## clientToken

Unique, case-sensitive identifier that you provide to ensure the idempotency of the request. Up to 36 letters, numbers, hyphens, and underscores are allowed.

Type: String

Required: No

## meshName

The name to use for the service mesh.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## spec

The service mesh specification to apply.

Type: [MeshSpec](#) object

Required: No

## tags

Optional metadata that you can apply to the service mesh to assist with categorization and organization. Each tag consists of a key and an optional value, both of which you define. Tag keys can have a maximum character length of 128 characters, and tag values can have a maximum length of 256 characters.

Type: Array of [TagRef](#) objects

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Required: No

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json
```

```
{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "egressFilter": {
      "type": "string"
    },
    "serviceDiscovery": {
      "ipPreference": "string"
    }
  },
  "status": {
    "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.

### meshName

The name of the service mesh.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

The associated metadata for the service mesh.

Type: [ResourceMetadata](#) object

## spec

The associated specification for the service mesh.

Type: [MeshSpec](#) object

## status

The status of the service mesh.

Type: [MeshStatus](#) object

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ConflictException**

The request contains a client token that was used for a previous update resource call with different specifications. Try the request again with a new client token.

HTTP Status Code: 409

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **LimitExceededException**

You have exceeded a service limit for your account. For more information, see [Service Limits](#) in the *AWS App Mesh User Guide*.

HTTP Status Code: 400

## NotFoundException

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

## ServiceUnavailableException

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

## Example

The following example creates a service mesh named ecs-mesh.

## Sample Request

```
PUT /v20190125/meshes HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190227T192324Z
```

Authorization: AUTHPARAMS

```
{
  "meshName": "ecs-mesh",
  "clientToken": "34a20934-da3a-43a0-9d1b-390308a7393b"
}
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 5c2f774d-da0b-40f3-80c5-d8711eb15dce
content-type: application/json
content-length: 245
date: Wed, 27 Feb 2019 19:23:24 GMT
x-envoy-upstream-service-time: 76
server: envoy
Connection: keep-alive

{
  "meshName": "ecs-mesh",
  "metadata": {
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh",
    "createdAt": 1.551295405298E9,
    "lastUpdatedAt": 1.551295405298E9,
    "uid": "2d29a11c-f2dd-44a6-b620-33661cfdfe97",
    "version": 1
  },
  "status": {
    "status": "ACTIVE"
  }
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

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

# CreateRoute

Creates a route that is associated with a virtual router.

You can route several different protocols and define a retry policy for a route. Traffic can be routed to one or more virtual nodes.

For more information about routes, see [Routes](#).

## Request Syntax

```
PUT /v20190125/meshes/meshName/virtualRouter/virtualRouterName/routes?  
meshOwner=meshOwner HTTP/1.1  
Content-type: application/json
```

```
{  
  "clientToken": "string",  
  "routeName": "string",  
  "spec": {  
    "grpcRoute": {  
      "action": {  
        "weightedTargets": [  
          {  
            "port": number,  
            "virtualNode": "string",  
            "weight": number  
          }  
        ]  
      },  
      "match": {  
        "metadata": [  
          {  
            "invert": boolean,  
            "match": { ... },  
            "name": "string"  
          }  
        ],  
        "methodName": "string",  
        "port": number,  
        "serviceName": "string"  
      },  
      "retryPolicy": {  
        "grpcRetryEvents": [ "string" ],
```

```

    "httpRetryEvents": [ "string" ],
    "maxRetries": number,
    "perRetryTimeout": {
      "unit": "string",
      "value": number
    },
    "tcpRetryEvents": [ "string" ]
  },
  "timeout": {
    "idle": {
      "unit": "string",
      "value": number
    },
    "perRequest": {
      "unit": "string",
      "value": number
    }
  }
},
"http2Route": {
  "action": {
    "weightedTargets": [
      {
        "port": number,
        "virtualNode": "string",
        "weight": number
      }
    ]
  },
  "match": {
    "headers": [
      {
        "invert": boolean,
        "match": { ... },
        "name": "string"
      }
    ],
    "method": "string",
    "path": {
      "exact": "string",
      "regex": "string"
    },
    "port": number,
    "prefix": "string",

```

```
    "queryParameters": [
      {
        "match": {
          "exact": "string"
        },
        "name": "string"
      }
    ],
    "scheme": "string"
  },
  "retryPolicy": {
    "httpRetryEvents": [ "string" ],
    "maxRetries": number,
    "perRetryTimeout": {
      "unit": "string",
      "value": number
    },
    "tcpRetryEvents": [ "string" ]
  },
  "timeout": {
    "idle": {
      "unit": "string",
      "value": number
    },
    "perRequest": {
      "unit": "string",
      "value": number
    }
  }
},
"httpRoute": {
  "action": {
    "weightedTargets": [
      {
        "port": number,
        "virtualNode": "string",
        "weight": number
      }
    ]
  },
  "match": {
    "headers": [
      {
        "invert": boolean,
```

```
        "match": { ... },
        "name": "string"
    }
],
"method": "string",
"path": {
    "exact": "string",
    "regex": "string"
},
"port": number,
"prefix": "string",
"queryParameters": [
    {
        "match": {
            "exact": "string"
        },
        "name": "string"
    }
],
"scheme": "string"
},
"retryPolicy": {
    "httpRetryEvents": [ "string" ],
    "maxRetries": number,
    "perRetryTimeout": {
        "unit": "string",
        "value": number
    },
    "tcpRetryEvents": [ "string" ]
},
"timeout": {
    "idle": {
        "unit": "string",
        "value": number
    },
    "perRequest": {
        "unit": "string",
        "value": number
    }
}
},
"priority": number,
"tcpRoute": {
    "action": {
```

```
    "weightedTargets": [
      {
        "port": number,
        "virtualNode": "string",
        "weight": number
      }
    ],
    "match": {
      "port": number
    },
    "timeout": {
      "idle": {
        "unit": "string",
        "value": number
      }
    }
  }
},
"tags": [
  {
    "key": "string",
    "value": "string"
  }
]
}
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh to create the route in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then the account that you specify must share the mesh with your account before you can create

the resource in the service mesh. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### virtualRouterName

The name of the virtual router in which to create the route. If the virtual router is in a shared mesh, then you must be the owner of the virtual router resource.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request accepts the following data in JSON format.

### clientToken

Unique, case-sensitive identifier that you provide to ensure the idempotency of the request. Up to 36 letters, numbers, hyphens, and underscores are allowed.

Type: String

Required: No

### routeName

The name to use for the route.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### spec

The route specification to apply.

Type: [RouteSpec](#) object

Required: Yes

## tags

Optional metadata that you can apply to the route to assist with categorization and organization. Each tag consists of a key and an optional value, both of which you define. Tag keys can have a maximum character length of 128 characters, and tag values can have a maximum length of 256 characters.

Type: Array of [TagRef](#) objects

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Required: No

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "routeName": "string",
  "spec": {
    "grpcRoute": {
      "action": {
        "weightedTargets": [
          {
            "port": number,
            "virtualNode": "string",
            "weight": number
          }
        ]
      }
    },
    "match": {
```

```

    "metadata": [
      {
        "invert": boolean,
        "match": { ... },
        "name": "string"
      }
    ],
    "methodName": "string",
    "port": number,
    "serviceName": "string"
  },
  "retryPolicy": {
    "grpcRetryEvents": [ "string" ],
    "httpRetryEvents": [ "string" ],
    "maxRetries": number,
    "perRetryTimeout": {
      "unit": "string",
      "value": number
    },
    "tcpRetryEvents": [ "string" ]
  },
  "timeout": {
    "idle": {
      "unit": "string",
      "value": number
    },
    "perRequest": {
      "unit": "string",
      "value": number
    }
  }
},
"http2Route": {
  "action": {
    "weightedTargets": [
      {
        "port": number,
        "virtualNode": "string",
        "weight": number
      }
    ]
  }
},
"match": {
  "headers": [

```

```

    {
      "invert": boolean,
      "match": { ... },
      "name": "string"
    }
  ],
  "method": "string",
  "path": {
    "exact": "string",
    "regex": "string"
  },
  "port": number,
  "prefix": "string",
  "queryParameters": [
    {
      "match": {
        "exact": "string"
      },
      "name": "string"
    }
  ],
  "scheme": "string"
},
"retryPolicy": {
  "httpRetryEvents": [ "string" ],
  "maxRetries": number,
  "perRetryTimeout": {
    "unit": "string",
    "value": number
  },
  "tcpRetryEvents": [ "string" ]
},
"timeout": {
  "idle": {
    "unit": "string",
    "value": number
  },
  "perRequest": {
    "unit": "string",
    "value": number
  }
}
},
"httpRoute": {

```

```
"action": {
  "weightedTargets": [
    {
      "port": number,
      "virtualNode": "string",
      "weight": number
    }
  ]
},
"match": {
  "headers": [
    {
      "invert": boolean,
      "match": { ... },
      "name": "string"
    }
  ],
  "method": "string",
  "path": {
    "exact": "string",
    "regex": "string"
  },
  "port": number,
  "prefix": "string",
  "queryParameters": [
    {
      "match": {
        "exact": "string"
      },
      "name": "string"
    }
  ],
  "scheme": "string"
},
"retryPolicy": {
  "httpRetryEvents": [ "string" ],
  "maxRetries": number,
  "perRetryTimeout": {
    "unit": "string",
    "value": number
  },
  "tcpRetryEvents": [ "string" ]
},
"timeout": {
```

```
    "idle": {
      "unit": "string",
      "value": number
    },
    "perRequest": {
      "unit": "string",
      "value": number
    }
  },
  "priority": number,
  "tcpRoute": {
    "action": {
      "weightedTargets": [
        {
          "port": number,
          "virtualNode": "string",
          "weight": number
        }
      ]
    },
    "match": {
      "port": number
    },
    "timeout": {
      "idle": {
        "unit": "string",
        "value": number
      }
    }
  }
},
"status": {
  "status": "string"
},
"virtualRouterName": "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.

## meshName

The name of the service mesh that the route resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## metadata

The associated metadata for the route.

Type: [ResourceMetadata](#) object

## routeName

The name of the route.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## spec

The specifications of the route.

Type: [RouteSpec](#) object

## status

The status of the route.

Type: [RouteStatus](#) object

## virtualRouterName

The virtual router that the route is associated with.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### BadRequestException

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ConflictException**

The request contains a client token that was used for a previous update resource call with different specifications. Try the request again with a new client token.

HTTP Status Code: 409

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **LimitExceededException**

You have exceeded a service limit for your account. For more information, see [Service Limits](#) in the *AWS App Mesh User Guide*.

HTTP Status Code: 400

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### Example

The following example creates a route named `colorgateway-route` for the `colorgateway-vr` virtual router in the `ecs-mesh` service mesh. The route directs all traffic to the `colorgateway-vn` virtual node.

### Sample Request

```
PUT /v20190125/meshes/ecs-mesh/virtualRouter/colorgateway-vr/routes HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190227T192454Z
Authorization: AUTHPARAMS

{
  "routeName": "colorgateway-route",
  "spec": {
    "httpRoute": {
      "action": {
        "weightedTargets": [
          {
            "virtualNode": "colorgateway-vn",
            "weight": 100
          }
        ]
      },
      "match": {
        "prefix": "/"
      }
    }
  }
}
```

```
    }
  },
  "clientToken": "a785b959-201d-49d7-b81d-32e7f393968f"
}
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 159941a8-18be-423c-89fe-3beb9542ca3f
content-type: application/json
content-length: 511
date: Wed, 27 Feb 2019 19:24:54 GMT
x-envoy-upstream-service-time: 59
server: envoy
Connection: keep-alive

{
  "meshName": "ecs-mesh",
  "metadata": {
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualRouter/
colorgateway-vr/route/colorgateway-route",
    "createdAt": 1.551295495004E9,
    "lastUpdatedAt": 1.551295495004E9,
    "uid": "3251bf37-2d01-4b79-be96-d0c36c61511f",
    "version": 1
  },
  "routeName": "colorgateway-route",
  "spec": {
    "httpRoute": {
      "action": {
        "weightedTargets": [{
          "virtualNode": "colorgateway-vn",
          "weight": 100
        }]
      },
      "match": {
        "prefix": "/"
      }
    },
    "tcpRoute": null
  },
  "status": {
    "status": "ACTIVE"
  }
}
```

```
  },  
  "virtualRouterName": "colorgateway-vr"  
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# CreateVirtualGateway

Creates a virtual gateway.

A virtual gateway allows resources outside your mesh to communicate to resources that are inside your mesh. The virtual gateway represents an Envoy proxy running in an Amazon ECS task, in a Kubernetes service, or on an Amazon EC2 instance. Unlike a virtual node, which represents an Envoy running with an application, a virtual gateway represents Envoy deployed by itself.

For more information about virtual gateways, see [Virtual gateways](#).

## Request Syntax

```
PUT /v20190125/meshes/meshName/virtualGateways?meshOwner=meshOwner HTTP/1.1
```

```
Content-type: application/json
```

```
{
  "clientToken": "string",
  "spec": {
    "backendDefaults": {
      "clientPolicy": {
        "tls": {
          "certificate": { ... },
          "enforce": boolean,
          "ports": [ number ],
          "validation": {
            "subjectAlternativeNames": {
              "match": {
                "exact": [ "string" ]
              }
            },
            "trust": { ... }
          }
        }
      }
    },
    "listeners": [
      {
        "connectionPool": { ... },
        "healthCheck": {
          "healthyThreshold": number,
          "intervalMillis": number,
          "path": "string",
```

```

    "port": number,
    "protocol": "string",
    "timeoutMillis": number,
    "unhealthyThreshold": number
  },
  "portMapping": {
    "port": number,
    "protocol": "string"
  },
  "tls": {
    "certificate": { ... },
    "mode": "string",
    "validation": {
      "subjectAlternativeNames": {
        "match": {
          "exact": [ "string" ]
        }
      },
      "trust": { ... }
    }
  }
},
"logging": {
  "accessLog": { ... }
},
"tags": [
  {
    "key": "string",
    "value": "string"
  }
],
"virtualGatewayName": "string"
}

```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh to create the virtual gateway in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then the account that you specify must share the mesh with your account before you can create the resource in the service mesh. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

## Request Body

The request accepts the following data in JSON format.

### clientToken

Unique, case-sensitive identifier that you provide to ensure the idempotency of the request. Up to 36 letters, numbers, hyphens, and underscores are allowed.

Type: String

Required: No

### spec

The virtual gateway specification to apply.

Type: [VirtualGatewaySpec](#) object

Required: Yes

### tags

Optional metadata that you can apply to the virtual gateway to assist with categorization and organization. Each tag consists of a key and an optional value, both of which you define. Tag keys can have a maximum character length of 128 characters, and tag values can have a maximum length of 256 characters.

Type: Array of [TagRef](#) objects

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Required: No

### virtualGatewayName

The name to use for the virtual gateway.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "backendDefaults": {
      "clientPolicy": {
        "tls": {
          "certificate": { ... },
          "enforce": boolean,
          "ports": [ number ],
          "validation": {
            "subjectAlternativeNames": {
              "match": {
                "exact": [ "string" ]
              }
            },
            "trust": { ... }
          }
        }
      }
    }
  }
}
```

```

    }
  },
  "listeners": [
    {
      "connectionPool": { ... },
      "healthCheck": {
        "healthyThreshold": number,
        "intervalMillis": number,
        "path": "string",
        "port": number,
        "protocol": "string",
        "timeoutMillis": number,
        "unhealthyThreshold": number
      },
      "portMapping": {
        "port": number,
        "protocol": "string"
      },
      "tls": {
        "certificate": { ... },
        "mode": "string",
        "validation": {
          "subjectAlternativeNames": {
            "match": {
              "exact": [ "string" ]
            }
          },
          "trust": { ... }
        }
      }
    }
  ],
  "logging": {
    "accessLog": { ... }
  }
},
"status": {
  "status": "string"
},
"virtualGatewayName": "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.

### meshName

The name of the service mesh that the virtual gateway resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

An object that represents metadata for a resource.

Type: [ResourceMetadata](#) object

### spec

The specifications of the virtual gateway.

Type: [VirtualGatewaySpec](#) object

### status

The current status of the virtual gateway.

Type: [VirtualGatewayStatus](#) object

### virtualGatewayName

The name of the virtual gateway.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ConflictException**

The request contains a client token that was used for a previous update resource call with different specifications. Try the request again with a new client token.

HTTP Status Code: 409

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **LimitExceededException**

You have exceeded a service limit for your account. For more information, see [Service Limits](#) in the *AWS App Mesh User Guide*.

HTTP Status Code: 400

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### Example

The following example creates a virtual gateway named myVirtualGateway in the apps service mesh. The virtual gateway listens for http2 traffic on port 80.

### Sample Request

```
PUT /v20190125/meshes/apps/virtualGateways HTTP/1.1
Host: appmesh.us-west-2.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.298 Python/3.6.0 Windows/10 botocore/1.13.34
X-Amz-Date: 20200706T144617Z
Authorization: AUTHPARAMS
Content-Length: 174

{
  "spec": {
    "listeners": [{
      "portMapping": {
        "port": 80,
        "protocol": "http2"
      }
    }]
  },
  "virtualGatewayName": "myVirtualGateway",
  "clientToken": "1aa1111a-1111-1111-111a-a1a1a1aa111a"
}
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 696b569b-3593-446d-832c-0a38ce1c53ac
content-type: application/json
content-length: 513
date: Mon, 06 Jul 2020 14:46:17 GMT
x-envoy-upstream-service-time: 33
server: envoy
Connection: keep-alive

{
  "meshName": "apps",
  "metadata": {
    "arn": "arn:aws:appmesh:us-west-2:123456789012:mesh/apps/virtualGateway/myVirtualGateway",
    "createdAt": 1.594046778215E9,
    "lastUpdatedAt": 1.594046778215E9,
    "meshOwner": "123456789012",
    "resourceOwner": "123456789012",
    "uid": "ff111ff2-33ff-4ff4-ffff-111f111f111f",
    "version": 1
  },
  "spec": {
    "backendDefaults": null,
    "listeners": [{
      "healthCheck": null,
      "portMapping": {
        "port": 80,
        "protocol": "http2"
      },
      "tls": null
    }],
    "logging": null
  },
  "status": {
    "status": "ACTIVE"
  },
  "virtualGatewayName": "myVirtualGateway"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# CreateVirtualNode

Creates a virtual node within a service mesh.

A virtual node acts as a logical pointer to a particular task group, such as an Amazon ECS service or a Kubernetes deployment. When you create a virtual node, you can specify the service discovery information for your task group, and whether the proxy running in a task group will communicate with other proxies using Transport Layer Security (TLS).

You define a `listener` for any inbound traffic that your virtual node expects. Any virtual service that your virtual node expects to communicate to is specified as a backend.

The response metadata for your new virtual node contains the `arn` that is associated with the virtual node. Set this value to the full ARN; for example, `arn:aws:appmesh:us-west-2:123456789012:myMesh/default/virtualNode/myApp`) as the `APPESH_RESOURCE_ARN` environment variable for your task group's Envoy proxy container in your task definition or pod spec. This is then mapped to the `node.id` and `node.cluster` Envoy parameters.

## Note

By default, App Mesh uses the name of the resource you specified in `APPESH_RESOURCE_ARN` when Envoy is referring to itself in metrics and traces. You can override this behavior by setting the `APPESH_RESOURCE_CLUSTER` environment variable with your own name.

For more information about virtual nodes, see [Virtual nodes](#). You must be using 1.15.0 or later of the Envoy image when setting these variables. For more information about App Mesh Envoy variables, see [Envoy image](#) in the AWS App Mesh User Guide.

## Request Syntax

```
PUT /v20190125/meshes/meshName/virtualNodes?meshOwner=meshOwner HTTP/1.1
```

```
Content-type: application/json
```

```
{
  "clientToken": "string",
  "spec": {
    "backendDefaults": {
```

```
"clientPolicy": {
  "tls": {
    "certificate": { ... },
    "enforce": boolean,
    "ports": [ number ],
    "validation": {
      "subjectAlternativeNames": {
        "match": {
          "exact": [ "string" ]
        }
      },
      "trust": { ... }
    }
  }
},
"backends": [
  { ... }
],
"listeners": [
  {
    "connectionPool": { ... },
    "healthCheck": {
      "healthyThreshold": number,
      "intervalMillis": number,
      "path": "string",
      "port": number,
      "protocol": "string",
      "timeoutMillis": number,
      "unhealthyThreshold": number
    },
    "outlierDetection": {
      "baseEjectionDuration": {
        "unit": "string",
        "value": number
      },
      "interval": {
        "unit": "string",
        "value": number
      },
      "maxEjectionPercent": number,
      "maxServerErrors": number
    },
    "portMapping": {
```

```

    "port": number,
    "protocol": "string"
  },
  "timeout": { ... },
  "tls": {
    "certificate": { ... },
    "mode": "string",
    "validation": {
      "subjectAlternativeNames": {
        "match": {
          "exact": [ "string" ]
        }
      },
      "trust": { ... }
    }
  }
},
"logging": {
  "accessLog": { ... }
},
"serviceDiscovery": { ... }
},
"tags": [
  {
    "key": "string",
    "value": "string"
  }
],
"virtualNodeName": "string"
}

```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh to create the virtual node in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then the account that you specify must share the mesh with your account before you can create the resource in the service mesh. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

## Request Body

The request accepts the following data in JSON format.

### clientToken

Unique, case-sensitive identifier that you provide to ensure the idempotency of the request. Up to 36 letters, numbers, hyphens, and underscores are allowed.

Type: String

Required: No

### spec

The virtual node specification to apply.

Type: [VirtualNodeSpec](#) object

Required: Yes

### tags

Optional metadata that you can apply to the virtual node to assist with categorization and organization. Each tag consists of a key and an optional value, both of which you define. Tag keys can have a maximum character length of 128 characters, and tag values can have a maximum length of 256 characters.

Type: Array of [TagRef](#) objects

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Required: No

## virtualNodeName

The name to use for the virtual node.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "backendDefaults": {
      "clientPolicy": {
        "tls": {
          "certificate": { ... },
          "enforce": boolean,
          "ports": [ number ],
          "validation": {
            "subjectAlternativeNames": {
              "match": {
                "exact": [ "string" ]
              }
            },
            "trust": { ... }
          }
        }
      }
    }
  }
}
```

```

},
"backends": [
  { ... }
],
"listeners": [
  {
    "connectionPool": { ... },
    "healthCheck": {
      "healthyThreshold": number,
      "intervalMillis": number,
      "path": "string",
      "port": number,
      "protocol": "string",
      "timeoutMillis": number,
      "unhealthyThreshold": number
    },
    "outlierDetection": {
      "baseEjectionDuration": {
        "unit": "string",
        "value": number
      },
      "interval": {
        "unit": "string",
        "value": number
      },
      "maxEjectionPercent": number,
      "maxServerErrors": number
    },
    "portMapping": {
      "port": number,
      "protocol": "string"
    },
    "timeout": { ... },
    "tls": {
      "certificate": { ... },
      "mode": "string",
      "validation": {
        "subjectAlternativeNames": {
          "match": {
            "exact": [ "string" ]
          }
        }
      },
      "trust": { ... }
    }
  }
]

```

```
    }
  },
],
"logging": {
  "accessLog": { ... }
},
"serviceDiscovery": { ... }
},
"status": {
  "status": "string"
},
"virtualNodeName": "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.

### meshName

The name of the service mesh that the virtual node resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

The associated metadata for the virtual node.

Type: [ResourceMetadata](#) object

### spec

The specifications of the virtual node.

Type: [VirtualNodeSpec](#) object

### status

The current status for the virtual node.

Type: [VirtualNodeStatus](#) object

## virtualNodeName

The name of the virtual node.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### BadRequestException

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### ConflictException

The request contains a client token that was used for a previous update resource call with different specifications. Try the request again with a new client token.

HTTP Status Code: 409

### ForbiddenException

You don't have permissions to perform this action.

HTTP Status Code: 403

### InternalServerErrorException

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

HTTP Status Code: 500

### LimitExceededException

You have exceeded a service limit for your account. For more information, see [Service Limits](#) in the *AWS App Mesh User Guide*.

HTTP Status Code: 400

### NotFoundException

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## **Examples**

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### **Example**

The following example creates a virtual node named `colorgateway-vn` in the `ecs-mesh` service mesh.

### **Sample Request**

```
PUT /v20190125/meshes/ecs-mesh/virtualNodes HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190227T192431Z
Authorization: AUTHPARAMS

{
```

```
"spec": {
  "listeners": [
    {
      "portMapping": {
        "port": 9080,
        "protocol": "http"
      }
    }
  ],
  "serviceDiscovery": {
    "dns": {
      "hostname": "colorgateway.default.svc.cluster.local"
    }
  },
  "backends": [
    {
      "virtualService": {
        "virtualServiceName": "tcpecho.default.svc.cluster.local"
      }
    },
    {
      "virtualService": {
        "virtualServiceName": "colorteller.default.svc.cluster.local"
      }
    }
  ]
},
"virtualNodeName": "colorgateway-vn",
"clientToken": "c148ccbb-3619-49da-bb3e-4561eb5370c4"
}
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: cc29e2dd-e4c4-4d6e-9424-e9211580f70e
content-type: application/json
content-length: 687
date: Wed, 27 Feb 2019 19:24:31 GMT
x-envoy-upstream-service-time: 132
server: envoy
Connection: keep-alive
```

```
{
```

```
"meshName": "ecs-mesh",
"metadata": {
  "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualNode/
colorgateway-vn",
  "createdAt": 1.551295471546E9,
  "lastUpdatedAt": 1.551295471546E9,
  "uid": "887cfab8-a727-41b2-8cd7-2fdebfd6304e",
  "version": 1
},
"spec": {
  "backends": [{
    "virtualService": {
      "virtualServiceName": "tcpecho.default.svc.cluster.local"
    }
  }, {
    "virtualService": {
      "virtualServiceName": "colorteller.default.svc.cluster.local"
    }
  }],
  "listeners": [{
    "healthCheck": null,
    "portMapping": {
      "port": 9080,
      "protocol": "http"
    }
  }],
  "logging": null,
  "serviceDiscovery": {
    "dns": {
      "hostname": "colorgateway.default.svc.cluster.local"
    }
  }
},
"status": {
  "status": "ACTIVE"
},
"virtualNodeName": "colorgateway-vn"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# CreateVirtualRouter

Creates a virtual router within a service mesh.

Specify a `listener` for any inbound traffic that your virtual router receives. Create a virtual router for each protocol and port that you need to route. Virtual routers handle traffic for one or more virtual services within your mesh. After you create your virtual router, create and associate routes for your virtual router that direct incoming requests to different virtual nodes.

For more information about virtual routers, see [Virtual routers](#).

## Request Syntax

```
PUT /v20190125/meshes/meshName/virtualRouters?meshOwner=meshOwner HTTP/1.1
```

```
Content-type: application/json
```

```
{
  "clientToken": "string",
  "spec": {
    "listeners": [
      {
        "portMapping": {
          "port": number,
          "protocol": "string"
        }
      }
    ]
  },
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ],
  "virtualRouterName": "string"
}
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh to create the virtual router in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then the account that you specify must share the mesh with your account before you can create the resource in the service mesh. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

## Request Body

The request accepts the following data in JSON format.

### clientToken

Unique, case-sensitive identifier that you provide to ensure the idempotency of the request. Up to 36 letters, numbers, hyphens, and underscores are allowed.

Type: String

Required: No

### spec

The virtual router specification to apply.

Type: [VirtualRouterSpec](#) object

Required: Yes

### tags

Optional metadata that you can apply to the virtual router to assist with categorization and organization. Each tag consists of a key and an optional value, both of which you define.

Tag keys can have a maximum character length of 128 characters, and tag values can have a maximum length of 256 characters.

Type: Array of [TagRef](#) objects

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Required: No

### [virtualRouterName](#)

The name to use for the virtual router.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Response Syntax

```
HTTP/1.1 200
```

```
Content-type: application/json
```

```
{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "listeners": [
      {
        "portMapping": {
          "port": number,
          "protocol": "string"
        }
      }
    ]
  }
}
```

```
    }
  ]
},
"status": {
  "status": "string"
},
"virtualRouterName": "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.

### meshName

The name of the service mesh that the virtual router resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

The associated metadata for the virtual router.

Type: [ResourceMetadata](#) object

### spec

The specifications of the virtual router.

Type: [VirtualRouterSpec](#) object

### status

The current status of the virtual router.

Type: [VirtualRouterStatus](#) object

### virtualRouterName

The name of the virtual router.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ConflictException**

The request contains a client token that was used for a previous update resource call with different specifications. Try the request again with a new client token.

HTTP Status Code: 409

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **LimitExceededException**

You have exceeded a service limit for your account. For more information, see [Service Limits](#) in the *AWS App Mesh User Guide*.

HTTP Status Code: 400

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

## ServiceUnavailableException

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### Example

The following example creates a virtual router named `colorteller-vr` in the `ecs-mesh` service mesh.

### Sample Request

```
PUT /v20190125/meshes/ecs-mesh/virtualRouters HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190227T192446Z
Authorization: AUTHPARAMS

{
  "spec": {
```

```
    listeners: [
      {
        port: 80,
        protocol: http
      }
    ]},
"virtualRouterName": "colorteller-vr",
"clientToken": "e90d8ee2-c111-4431-bfa3-4725a06a84b1"
}
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 770a356d-60e5-45f7-9f47-823eb5f1b750
content-type: application/json
content-length: 337
date: Wed, 27 Feb 2019 19:24:46 GMT
x-envoy-upstream-service-time: 19
server: envoy
Connection: keep-alive

{
  "meshName": "ecs-mesh",
  "metadata": {
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualRouter/colorteller-vr",
    "createdAt": 1551295486.588,
    "lastUpdatedAt": 1551295486.588,
    "uid": "a756d1e5-ce8f-40a5-afbc-380f61f0c1e0",
    "version": 1
  },
  "spec": {
    "listeners": [
      {
        "portMapping": {
          "port": 9080,
          "protocol": "http"
        }
      }
    ]
  },
  "status": {
    "status": "ACTIVE"
  }
}
```

```
  },  
  "virtualRouterName": "colorteller-vr"  
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# CreateVirtualService

Creates a virtual service within a service mesh.

A virtual service is an abstraction of a real service that is provided by a virtual node directly or indirectly by means of a virtual router. Dependent services call your virtual service by its `virtualServiceName`, and those requests are routed to the virtual node or virtual router that is specified as the provider for the virtual service.

For more information about virtual services, see [Virtual services](#).

## Request Syntax

```
PUT /v20190125/meshes/meshName/virtualServices?meshOwner=meshOwner HTTP/1.1
Content-type: application/json
```

```
{
  "clientToken": "string",
  "spec": {
    "provider": { ... }
  },
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ],
  "virtualServiceName": "string"
}
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh to create the virtual service in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then the account that you specify must share the mesh with your account before you can create the resource in the service mesh. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

## Request Body

The request accepts the following data in JSON format.

### clientToken

Unique, case-sensitive identifier that you provide to ensure the idempotency of the request. Up to 36 letters, numbers, hyphens, and underscores are allowed.

Type: String

Required: No

### spec

The virtual service specification to apply.

Type: [VirtualServiceSpec](#) object

Required: Yes

### tags

Optional metadata that you can apply to the virtual service to assist with categorization and organization. Each tag consists of a key and an optional value, both of which you define. Tag keys can have a maximum character length of 128 characters, and tag values can have a maximum length of 256 characters.

Type: Array of [TagRef](#) objects

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Required: No

## virtualServiceName

The name to use for the virtual service.

Type: String

Required: Yes

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "provider": { ... }
  },
  "status": {
    "status": "string"
  },
  "virtualServiceName": "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.

### meshName

The name of the service mesh that the virtual service resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

An object that represents metadata for a resource.

Type: [ResourceMetadata](#) object

### spec

The specifications of the virtual service.

Type: [VirtualServiceSpec](#) object

### status

The current status of the virtual service.

Type: [VirtualServiceStatus](#) object

### virtualServiceName

The name of the virtual service.

Type: String

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ConflictException**

The request contains a client token that was used for a previous update resource call with different specifications. Try the request again with a new client token.

HTTP Status Code: 409

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **LimitExceededException**

You have exceeded a service limit for your account. For more information, see [Service Limits](#) in the *AWS App Mesh User Guide*.

HTTP Status Code: 400

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## **Examples**

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

## Example

This example creates a virtual service named `colorgateway.default.svc.cluster.local` in the `ecs-mesh` service mesh.

### Sample Request

```
PUT /v20190125/meshes/ecs-mesh/virtualServices HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190227T192448Z
Authorization: AUTHPARAMS

{
  "spec": {
    "provider": {
      "virtualNode": {
        "virtualNodeName": "colorgateway-vn"
      }
    }
  },
  "virtualServiceName": "colorgateway.default.svc.cluster.local",
  "clientToken": "0a263779-366d-400f-9c4c-0ab7ff0ef392"
}
```

### Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 6418af7b-e02b-4f0e-bb96-a3b4aab428c8
content-type: application/json
content-length: 456
date: Wed, 27 Feb 2019 19:24:48 GMT
x-envoy-upstream-service-time: 67
server: envoy
Connection: keep-alive

{
  "meshName": "ecs-mesh",
  "metadata": {
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualService/colorgateway.default.svc.cluster.local",
    "createdAt": 1.551295488701E9,
```

```
"lastUpdatedAt": 1.551295488701E9,
"uid": "f1b0ed85-7648-4be2-a0fc-c97865029c48",
"version": 1
},
"spec": {
  "provider": {
    "virtualNode": {
      "virtualNodeName": "colorgateway-vn"
    },
    "virtualRouter": null
  }
},
"status": {
  "status": "ACTIVE"
},
"virtualServiceName": "colorgateway.default.svc.cluster.local"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DeleteGatewayRoute

Deletes an existing gateway route.

## Request Syntax

```
DELETE /v20190125/meshes/meshName/virtualGateway/virtualGatewayName/
gatewayRoutes/gatewayRouteName?meshOwner=meshOwner HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### gatewayRouteName

The name of the gateway route to delete.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshName

The name of the service mesh to delete the gateway route from.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### virtualGatewayName

The name of the virtual gateway to delete the route from.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "gatewayRouteName": "string",
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "grpcRoute": {
      "action": {
        "rewrite": {
          "hostname": {
            "defaultTargetHostname": "string"
          }
        },
        "target": {
          "port": number,
          "virtualService": {
            "virtualServiceName": "string"
          }
        }
      }
    },
    "match": {
      "hostname": {
        "exact": "string",
        "suffix": "string"
      },
      "metadata": [
        {

```

```

        "invert": boolean,
        "match": { ... },
        "name": "string"
    }
],
"port": number,
"serviceName": "string"
}
},
"http2Route": {
    "action": {
        "rewrite": {
            "hostname": {
                "defaultTargetHostname": "string"
            },
            "path": {
                "exact": "string"
            },
            "prefix": {
                "defaultPrefix": "string",
                "value": "string"
            }
        },
        "target": {
            "port": number,
            "virtualService": {
                "virtualServiceName": "string"
            }
        }
    }
},
"match": {
    "headers": [
        {
            "invert": boolean,
            "match": { ... },
            "name": "string"
        }
    ],
    "hostname": {
        "exact": "string",
        "suffix": "string"
    },
    "method": "string",
    "path": {

```

```

        "exact": "string",
        "regex": "string"
    },
    "port": number,
    "prefix": "string",
    "queryParameters": [
        {
            "match": {
                "exact": "string"
            },
            "name": "string"
        }
    ]
},
"httpRoute": {
    "action": {
        "rewrite": {
            "hostname": {
                "defaultTargetHostname": "string"
            },
            "path": {
                "exact": "string"
            },
            "prefix": {
                "defaultPrefix": "string",
                "value": "string"
            }
        },
        "target": {
            "port": number,
            "virtualService": {
                "virtualServiceName": "string"
            }
        }
    },
    "match": {
        "headers": [
            {
                "invert": boolean,
                "match": { ... },
                "name": "string"
            }
        ]
    }
},

```

```
    "hostname": {
      "exact": "string",
      "suffix": "string"
    },
    "method": "string",
    "path": {
      "exact": "string",
      "regex": "string"
    },
    "port": number,
    "prefix": "string",
    "queryParameters": [
      {
        "match": {
          "exact": "string"
        },
        "name": "string"
      }
    ]
  },
  "priority": number
},
"status": {
  "status": "string"
},
"virtualGatewayName": "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.

### gatewayRouteName

The name of the gateway route.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## meshName

The name of the service mesh that the resource resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## metadata

An object that represents metadata for a resource.

Type: [ResourceMetadata](#) object

## spec

The specifications of the gateway route.

Type: [GatewayRouteSpec](#) object

## status

The status of the gateway route.

Type: [GatewayRouteStatus](#) object

## virtualGatewayName

The virtual gateway that the gateway route is associated with.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ResourceInUseException**

You can't delete the specified resource because it's in use or required by another resource.

HTTP Status Code: 409

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## **Examples**

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

## Example

The following example deletes a gateway route named `myGatewayRoute` from the `myVirtualGateway` virtual gateway that is in the apps service mesh.

### Sample Request

```
DELETE /v20190125/meshes/apps/virtualGateway/myVirtualGateway/gatewayRoutes/
myGatewayRoute HTTP/1.1
Host: appmesh.us-west-2.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.298 Python/3.6.0 Windows/10 botocore/1.13.34
X-Amz-Date: 20200608T193419Z
Authorization: AUTHPARAMS
```

### Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: a7e8bcc1-0dec-464d-9e69-5cd1f780c9eb
content-type: application/json
content-length: 618
date: Mon, 08 Jun 2020 19:34:19 GMT
x-envoy-upstream-service-time: 49
server: envoy
Connection: keep-alive

{
  "gatewayRouteName": "myGatewayRoute",
  "meshName": "myApps",
  "metadata": {
    "arn": "arn:aws:appmesh:us-west-2:123456789012:mesh/apps/virtualGateway/
myVirtualGateway/gatewayRoute/myGatewayRoute",
    "createdAt": 1.591642091122E9,
    "lastUpdatedAt": 1.591644859804E9,
    "meshOwner": "123456789012",
    "resourceOwner": "123456789012",
    "uid": "ff111ff2-33ff-4ff4-ffff-111f111f111f",
    "version": 3
  },
  "spec": {
    "grpcRoute": null,
    "http2Route": {
      "action": {
```

```
"target": {
  "virtualService": {
    "virtualServiceName": "myserviceb.svc.cluster.local"
  }
},
"match": {
  "prefix": "/"
},
"httpRoute": null
},
"status": {
  "status": "DELETED"
},
"virtualGatewayName": "myVirtualGateway"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DeleteMesh

Deletes an existing service mesh.

You must delete all resources (virtual services, routes, virtual routers, and virtual nodes) in the service mesh before you can delete the mesh itself.

## Request Syntax

```
DELETE /v20190125/meshes/meshName HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh to delete.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
```

```
    "uid": "string",
    "version": number
  },
  "spec": {
    "egressFilter": {
      "type": "string"
    },
    "serviceDiscovery": {
      "ipPreference": "string"
    }
  },
  "status": {
    "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.

### meshName

The name of the service mesh.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

The associated metadata for the service mesh.

Type: [ResourceMetadata](#) object

### spec

The associated specification for the service mesh.

Type: [MeshSpec](#) object

### status

The status of the service mesh.

Type: [MeshStatus](#) object

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ResourceInUseException**

You can't delete the specified resource because it's in use or required by another resource.

HTTP Status Code: 409

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### Example

This example command deletes a service mesh named `ecs-mesh` in your default region.

#### Sample Request

```
DELETE /v20190125/meshes/ecs-mesh HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190227T203825Z
Authorization: AUTHPARAMS
```

#### Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 5b4e49f2-dac7-4782-8be3-1e81c8599e14
content-type: application/json
content-length: 246
date: Wed, 27 Feb 2019 20:38:25 GMT
x-envoy-upstream-service-time: 35
server: envoy
Connection: keep-alive

{
  "meshName": "ecs-mesh",
  "metadata": {
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh",
    "createdAt": 1.551295405298E9,
```

```
"lastUpdatedAt": 1.551299905963E9,  
"uid": "2d29a11c-f2dd-44a6-b620-33661cfdfe97",  
"version": 1  
},  
"status": {  
  "status": "DELETED"  
}  
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DeleteRoute

Deletes an existing route.

## Request Syntax

```
DELETE /v20190125/meshes/meshName/virtualRouter/virtualRouterName/routes/routeName?  
meshOwner=meshOwner HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh to delete the route in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### routeName

The name of the route to delete.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### virtualRouterName

The name of the virtual router to delete the route in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "routeName": "string",
  "spec": {
    "grpcRoute": {
      "action": {
        "weightedTargets": [
          {
            "port": number,
            "virtualNode": "string",
            "weight": number
          }
        ]
      }
    },
    "match": {
      "metadata": [
        {
          "invert": boolean,
          "match": { ... },
          "name": "string"
        }
      ]
    },
    "methodName": "string",
    "port": number,
    "serviceName": "string"
  }
}
```

```
    },
    "retryPolicy": {
      "grpcRetryEvents": [ "string" ],
      "httpRetryEvents": [ "string" ],
      "maxRetries": number,
      "perRetryTimeout": {
        "unit": "string",
        "value": number
      },
    },
    "tcpRetryEvents": [ "string" ]
  },
  "timeout": {
    "idle": {
      "unit": "string",
      "value": number
    },
    "perRequest": {
      "unit": "string",
      "value": number
    }
  }
},
"http2Route": {
  "action": {
    "weightedTargets": [
      {
        "port": number,
        "virtualNode": "string",
        "weight": number
      }
    ]
  },
  "match": {
    "headers": [
      {
        "invert": boolean,
        "match": { ... },
        "name": "string"
      }
    ],
    "method": "string",
    "path": {
      "exact": "string",
      "regex": "string"
    }
  }
}
```

```
    },
    "port": number,
    "prefix": "string",
    "queryParameters": [
      {
        "match": {
          "exact": "string"
        },
        "name": "string"
      }
    ],
    "scheme": "string"
  },
  "retryPolicy": {
    "httpRetryEvents": [ "string" ],
    "maxRetries": number,
    "perRetryTimeout": {
      "unit": "string",
      "value": number
    },
    "tcpRetryEvents": [ "string" ]
  },
  "timeout": {
    "idle": {
      "unit": "string",
      "value": number
    },
    "perRequest": {
      "unit": "string",
      "value": number
    }
  }
},
"httpRoute": {
  "action": {
    "weightedTargets": [
      {
        "port": number,
        "virtualNode": "string",
        "weight": number
      }
    ]
  },
  "match": {
```

```

    "headers": [
      {
        "invert": boolean,
        "match": { ... },
        "name": "string"
      }
    ],
    "method": "string",
    "path": {
      "exact": "string",
      "regex": "string"
    },
    "port": number,
    "prefix": "string",
    "queryParameters": [
      {
        "match": {
          "exact": "string"
        },
        "name": "string"
      }
    ],
    "scheme": "string"
  },
  "retryPolicy": {
    "httpRetryEvents": [ "string" ],
    "maxRetries": number,
    "perRetryTimeout": {
      "unit": "string",
      "value": number
    },
    "tcpRetryEvents": [ "string" ]
  },
  "timeout": {
    "idle": {
      "unit": "string",
      "value": number
    },
    "perRequest": {
      "unit": "string",
      "value": number
    }
  }
},

```

```
"priority": number,
"tcpRoute": {
  "action": {
    "weightedTargets": [
      {
        "port": number,
        "virtualNode": "string",
        "weight": number
      }
    ]
  },
  "match": {
    "port": number
  },
  "timeout": {
    "idle": {
      "unit": "string",
      "value": number
    }
  }
}
},
"status": {
  "status": "string"
},
"virtualRouterName": "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.

### meshName

The name of the service mesh that the route resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

The associated metadata for the route.

Type: [ResourceMetadata](#) object

### routeName

The name of the route.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### spec

The specifications of the route.

Type: [RouteSpec](#) object

### status

The status of the route.

Type: [RouteStatus](#) object

### virtualRouterName

The virtual router that the route is associated with.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ResourceInUseException**

You can't delete the specified resource because it's in use or required by another resource.

HTTP Status Code: 409

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## **Examples**

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### **Example**

The following example deletes a route named `colorgateway-route` for the `colorgateway-vr` virtual router in the `ecs-mesh` service mesh.

## Sample Request

```
DELETE /v20190125/meshes/ecs-mesh/virtualRouter/colorgateway-vr/routes/colorgateway-  
route HTTP/1.1  
Host: appmesh.us-east-1.amazonaws.com  
Accept-Encoding: identity  
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46  
X-Amz-Date: 20190227T203328Z  
Authorization: AUTHPARAMS
```

## Sample Response

```
HTTP/1.1 200 OK  
x-amzn-requestid: 59fda333-ead1-4fbd-ba64-b7150e52d29a  
content-type: application/json  
content-length: 512  
date: Wed, 27 Feb 2019 20:33:28 GMT  
x-envoy-upstream-service-time: 95  
server: envoy  
Connection: keep-alive  
  
{  
  "meshName": "ecs-mesh",  
  "metadata": {  
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualRouter/  
colorgateway-vr/route/colorgateway-route",  
    "createdAt": 1.551295495004E9,  
    "lastUpdatedAt": 1.551299608556E9,  
    "uid": "3251bf37-2d01-4b79-be96-d0c36c61511f",  
    "version": 2  
  },  
  "routeName": "colorgateway-route",  
  "spec": {  
    "httpRoute": {  
      "action": {  
        "weightedTargets": [{  
          "virtualNode": "colorgateway-vn",  
          "weight": 100  
        }]  
      },  
      "match": {  
        "prefix": "/"  
      }  
    }  
  }  
}
```

```
  },
  "tcpRoute": null
},
"status": {
  "status": "DELETED"
},
"virtualRouterName": "colorgateway-vr"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DeleteVirtualGateway

Deletes an existing virtual gateway. You cannot delete a virtual gateway if any gateway routes are associated to it.

## Request Syntax

```
DELETE /v20190125/meshes/meshName/virtualGateways/virtualGatewayName?  
meshOwner=meshOwner HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh to delete the virtual gateway from.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### virtualGatewayName

The name of the virtual gateway to delete.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request does not have a request body.

## Response Syntax

HTTP/1.1 200

Content-type: application/json

```
{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "backendDefaults": {
      "clientPolicy": {
        "tls": {
          "certificate": { ... },
          "enforce": boolean,
          "ports": [ number ],
          "validation": {
            "subjectAlternativeNames": {
              "match": {
                "exact": [ "string" ]
              }
            }
          },
          "trust": { ... }
        }
      }
    }
  },
  "listeners": [
    {
      "connectionPool": { ... },
      "healthCheck": {
        "healthyThreshold": number,
        "intervalMillis": number,
        "path": "string",
        "port": number,
        "protocol": "string",
```

```
    "timeoutMillis": number,
    "unhealthyThreshold": number
  },
  "portMapping": {
    "port": number,
    "protocol": "string"
  },
  "tls": {
    "certificate": { ... },
    "mode": "string",
    "validation": {
      "subjectAlternativeNames": {
        "match": {
          "exact": [ "string" ]
        }
      },
      "trust": { ... }
    }
  }
},
"logging": {
  "accessLog": { ... }
},
"status": {
  "status": "string"
},
"virtualGatewayName": "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.

### meshName

The name of the service mesh that the virtual gateway resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## metadata

An object that represents metadata for a resource.

Type: [ResourceMetadata](#) object

## spec

The specifications of the virtual gateway.

Type: [VirtualGatewaySpec](#) object

## status

The current status of the virtual gateway.

Type: [VirtualGatewayStatus](#) object

## virtualGatewayName

The name of the virtual gateway.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### BadRequestException

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### ForbiddenException

You don't have permissions to perform this action.

HTTP Status Code: 403

### InternalServerErrorException

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

HTTP Status Code: 500

## NotFoundException

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

## ResourceInUseException

You can't delete the specified resource because it's in use or required by another resource.

HTTP Status Code: 409

## ServiceUnavailableException

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

## Example

The following example deletes a virtual gateway named myVirtualGateway from the apps service mesh.

## Sample Request

```
DELETE /v20190125/meshes/apps/virtualGateways/myVirtualGateway HTTP/1.1
Host: appmesh.us-west-2.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.298 Python/3.6.0 Windows/10 botocore/1.13.34
X-Amz-Date: 20200706T150515Z
Authorization: AUTHPARAMS
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 884f2830-6964-4d0e-9788-144631c36bb9
content-type: application/json
content-length: 514
date: Mon, 06 Jul 2020 15:05:15 GMT
x-envoy-upstream-service-time: 26
server: envoy
Connection: keep-alive

{
  "meshName": "apps",
  "metadata": {
    "arn": "arn:aws:appmesh:us-west-2:123456789012:mesh/apps/virtualGateway/
myVirtualGateway",
    "createdAt": 1.594046778215E9,
    "lastUpdatedAt": 1.594047916142E9,
    "meshOwner": "123456789012",
    "resourceOwner": "123456789012",
    "uid": "ff111ff2-33ff-4ff4-ffff-111f111f111f",
    "version": 2
  },
  "spec": {
    "backendDefaults": null,
    "listeners": [{
      "healthCheck": null,
      "portMapping": {
        "port": 80,
        "protocol": "http2"
      },
      "tls": null
    }],
    "logging": null
  }
}
```

```
  },
  "status": {
    "status": "DELETED"
  },
  "virtualGatewayName": "myVirtualGateway"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DeleteVirtualNode

Deletes an existing virtual node.

You must delete any virtual services that list a virtual node as a service provider before you can delete the virtual node itself.

## Request Syntax

```
DELETE /v20190125/meshes/meshName/virtualNodes/virtualNodeName?meshOwner=meshOwner
HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh to delete the virtual node in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### virtualNodeName

The name of the virtual node to delete.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request does not have a request body.

## Response Syntax

HTTP/1.1 200

Content-type: application/json

```
{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "backendDefaults": {
      "clientPolicy": {
        "tls": {
          "certificate": { ... },
          "enforce": boolean,
          "ports": [ number ],
          "validation": {
            "subjectAlternativeNames": {
              "match": {
                "exact": [ "string" ]
              }
            },
            "trust": { ... }
          }
        }
      }
    },
    "backends": [
      { ... }
    ],
    "listeners": [
      {
        "connectionPool": { ... },
        "healthCheck": {
          "healthyThreshold": number,
          "intervalMillis": number,

```

```

    "path": "string",
    "port": number,
    "protocol": "string",
    "timeoutMillis": number,
    "unhealthyThreshold": number
  },
  "outlierDetection": {
    "baseEjectionDuration": {
      "unit": "string",
      "value": number
    },
    "interval": {
      "unit": "string",
      "value": number
    },
    "maxEjectionPercent": number,
    "maxServerErrors": number
  },
  "portMapping": {
    "port": number,
    "protocol": "string"
  },
  "timeout": { ... },
  "tls": {
    "certificate": { ... },
    "mode": "string",
    "validation": {
      "subjectAlternativeNames": {
        "match": {
          "exact": [ "string" ]
        }
      },
      "trust": { ... }
    }
  }
},
],
"logging": {
  "accessLog": { ... }
},
"serviceDiscovery": { ... }
},
"status": {
  "status": "string"
}

```

```
  },  
  "virtualNodeName": "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.

### meshName

The name of the service mesh that the virtual node resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

The associated metadata for the virtual node.

Type: [ResourceMetadata](#) object

### spec

The specifications of the virtual node.

Type: [VirtualNodeSpec](#) object

### status

The current status for the virtual node.

Type: [VirtualNodeStatus](#) object

### virtualNodeName

The name of the virtual node.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ResourceInUseException**

You can't delete the specified resource because it's in use or required by another resource.

HTTP Status Code: 409

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### Example

This example deletes a virtual node named `colorgateway-vn` in the `ecs-mesh` service mesh.

#### Sample Request

```
DELETE /v20190125/meshes/ecs-mesh/virtualNodes/colorgateway-vn HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190227T231337Z
Authorization: AUTHPARAMS
```

#### Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 84875cd2-d36c-445a-b1a7-6a26ec798e62
content-type: application/json
content-length: 688
date: Wed, 27 Feb 2019 23:13:37 GMT
x-envoy-upstream-service-time: 25
server: envoy
Connection: keep-alive

{
  "meshName": "ecs-mesh",
  "metadata": {
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualNode/colorgateway-vn",
    "createdAt": 1.551307227168E9,
```

```
"lastUpdatedAt": 1.551309217745E9,
"uid": "4cd6303b-491e-4c2b-9108-1a4690ec9f36",
"version": 2
},
"spec": {
  "backends": [{
    "virtualService": {
      "virtualServiceName": "tcpecho.default.svc.cluster.local"
    }
  }, {
    "virtualService": {
      "virtualServiceName": "colorteller.default.svc.cluster.local"
    }
  }],
  "listeners": [{
    "healthCheck": null,
    "portMapping": {
      "port": 9080,
      "protocol": "http"
    }
  }],
  "logging": null,
  "serviceDiscovery": {
    "dns": {
      "hostname": "colorgateway.default.svc.cluster.local"
    }
  }
},
"status": {
  "status": "DELETED"
},
"virtualNodeName": "colorgateway-vn"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DeleteVirtualRouter

Deletes an existing virtual router.

You must delete any routes associated with the virtual router before you can delete the router itself.

## Request Syntax

```
DELETE /v20190125/meshes/meshName/virtualRouters/virtualRouterName?meshOwner=meshOwner
HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh to delete the virtual router in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### virtualRouterName

The name of the virtual router to delete.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "listeners": [
      {
        "portMapping": {
          "port": number,
          "protocol": "string"
        }
      }
    ]
  },
  "status": {
    "status": "string"
  },
  "virtualRouterName": "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.

### meshName

The name of the service mesh that the virtual router resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

The associated metadata for the virtual router.

Type: [ResourceMetadata](#) object

### spec

The specifications of the virtual router.

Type: [VirtualRouterSpec](#) object

### status

The current status of the virtual router.

Type: [VirtualRouterStatus](#) object

### virtualRouterName

The name of the virtual router.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ResourceInUseException**

You can't delete the specified resource because it's in use or required by another resource.

HTTP Status Code: 409

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## **Examples**

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### **Example**

This example deletes a virtual router named `co1oigateway-vr` in the `ecs-mesh` service mesh.

## Sample Request

```
DELETE /v20190125/meshes/ecs-mesh/virtualRouters/colorgateway-vr HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190227T231334Z
Authorization: AUTHPARAMS
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 77a5e726-69b4-4bd4-9016-99f76f4d1daa
content-type: application/json
content-length: 340
date: Wed, 27 Feb 2019 23:13:34 GMT
x-envoy-upstream-service-time: 23
server: envoy
Connection: keep-alive

{
  "meshName": "ecs-mesh",
  "metadata": {
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualRouter/colorgateway-vr",
    "createdAt": 1551307245.617,
    "lastUpdatedAt": 1551309214.686,
    "uid": "601f33c5-8938-4c07-a28a-3d3c91871ee4",
    "version": 2
  },
  "spec": {
    "listeners": [
      {
        "portMapping": {
          "port": 9080,
          "protocol": "http"
        }
      }
    ]
  },
  "status": {
    "status": "DELETED"
  },
}
```

```
"virtualRouterName": "colorgateway-vr"  
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DeleteVirtualService

Deletes an existing virtual service.

## Request Syntax

```
DELETE /v20190125/meshes/meshName/virtualServices/virtualServiceName?  
meshOwner=meshOwner HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh to delete the virtual service in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### virtualServiceName

The name of the virtual service to delete.

Required: Yes

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
```

```
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "provider": { ... }
  },
  "status": {
    "status": "string"
  },
  "virtualServiceName": "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.

### meshName

The name of the service mesh that the virtual service resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

An object that represents metadata for a resource.

Type: [ResourceMetadata](#) object

### spec

The specifications of the virtual service.

Type: [VirtualServiceSpec](#) object

### **status**

The current status of the virtual service.

Type: [VirtualServiceStatus](#) object

### **virtualServiceName**

The name of the virtual service.

Type: String

## **Errors**

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ResourceInUseException**

You can't delete the specified resource because it's in use or required by another resource.

HTTP Status Code: 409

### **ServiceUnavailableException**

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### Example

This example deletes a virtual service named `colorgateway.default.svc.cluster.local` in the `ecs-mesh` service mesh.

### Sample Request

```
DELETE /v20190125/meshes/ecs-mesh/virtualServices/
colorgateway.default.svc.cluster.local HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190227T231318Z
Authorization: AUTHPARAMS
```

### Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 8a5c10db-aebc-4341-8b15-4003e87150e8
content-type: application/json
```

```
content-length: 456
date: Wed, 27 Feb 2019 23:13:18 GMT
x-envoy-upstream-service-time: 78
server: envoy
Connection: keep-alive

{
  "meshName": "ecs-mesh",
  "metadata": {
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualService/
colorgateway.default.svc.cluster.local",
    "createdAt": 1.551307250696E9,
    "lastUpdatedAt": 1.55130919872E9,
    "uid": "90849766-3af0-40bc-9a83-a7b932d64fb6",
    "version": 2
  },
  "spec": {
    "provider": {
      "virtualNode": {
        "virtualNodeName": "colorgateway-vn"
      },
      "virtualRouter": null
    }
  },
  "status": {
    "status": "DELETED"
  },
  "virtualServiceName": "colorgateway.default.svc.cluster.local"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)

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

# DescribeGatewayRoute

Describes an existing gateway route.

## Request Syntax

```
GET /v20190125/meshes/meshName/virtualGateway/virtualGatewayName/
gatewayRoutes/gatewayRouteName?meshOwner=meshOwner HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### gatewayRouteName

The name of the gateway route to describe.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshName

The name of the service mesh that the gateway route resides in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### virtualGatewayName

The name of the virtual gateway that the gateway route is associated with.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "gatewayRouteName": string,
  "meshName": string,
  "metadata": {
    "arn": string,
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": string,
    "resourceOwner": string,
    "uid": string,
    "version": number
  },
  "spec": {
    "grpcRoute": {
      "action": {
        "rewrite": {
          "hostname": {
            "defaultTargetHostname": string
          }
        },
        "target": {
          "port": number,
          "virtualService": {
            "virtualServiceName": string
          }
        }
      }
    },
    "match": {
      "hostname": {
        "exact": string,
        "suffix": string
      },
      "metadata": [
        {

```

```

        "invert": boolean,
        "match": { ... },
        "name": "string"
    }
],
"port": number,
"serviceName": "string"
}
},
"http2Route": {
    "action": {
        "rewrite": {
            "hostname": {
                "defaultTargetHostname": "string"
            },
            "path": {
                "exact": "string"
            },
            "prefix": {
                "defaultPrefix": "string",
                "value": "string"
            }
        },
        "target": {
            "port": number,
            "virtualService": {
                "virtualServiceName": "string"
            }
        }
    }
},
"match": {
    "headers": [
        {
            "invert": boolean,
            "match": { ... },
            "name": "string"
        }
    ],
    "hostname": {
        "exact": "string",
        "suffix": "string"
    },
    "method": "string",
    "path": {

```

```

        "exact": "string",
        "regex": "string"
    },
    "port": number,
    "prefix": "string",
    "queryParameters": [
        {
            "match": {
                "exact": "string"
            },
            "name": "string"
        }
    ]
},
"httpRoute": {
    "action": {
        "rewrite": {
            "hostname": {
                "defaultTargetHostname": "string"
            },
            "path": {
                "exact": "string"
            },
            "prefix": {
                "defaultPrefix": "string",
                "value": "string"
            }
        },
        "target": {
            "port": number,
            "virtualService": {
                "virtualServiceName": "string"
            }
        }
    },
    "match": {
        "headers": [
            {
                "invert": boolean,
                "match": { ... },
                "name": "string"
            }
        ]
    },

```

```
    "hostname": {
      "exact": "string",
      "suffix": "string"
    },
    "method": "string",
    "path": {
      "exact": "string",
      "regex": "string"
    },
    "port": number,
    "prefix": "string",
    "queryParameters": [
      {
        "match": {
          "exact": "string"
        },
        "name": "string"
      }
    ]
  },
  "priority": number
},
"status": {
  "status": "string"
},
"virtualGatewayName": "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.

### gatewayRouteName

The name of the gateway route.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## meshName

The name of the service mesh that the resource resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## metadata

An object that represents metadata for a resource.

Type: [ResourceMetadata](#) object

## spec

The specifications of the gateway route.

Type: [GatewayRouteSpec](#) object

## status

The status of the gateway route.

Type: [GatewayRouteStatus](#) object

## virtualGatewayName

The virtual gateway that the gateway route is associated with.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## **Examples**

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### **Example**

The following example describes a gateway route named `myGatewayRoute` that is associated with a virtual gateway named `myVirtualGateway` in the apps service mesh.

## Sample Request

```
GET /v20190125/meshes/apps/virtualGateway/myVirtualGateway/gatewayRoutes/myGatewayRoute
HTTP/1.1
Host: appmesh.us-west-2.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.298 Python/3.6.0 Windows/10 botocore/1.13.34
X-Amz-Date: 20200608T194902Z
Authorization: AUTHPARAMS
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 6d60d9e4-294d-40ca-a829-37f9ebc058e8
content-type: application/json
content-length: 617
date: Mon, 08 Jun 2020 20:26:03 GMT
x-envoy-upstream-service-time: 31
server: envoy
Connection: keep-alive

{
  "gatewayRouteName": "myGatewayRoute",
  "meshName": "myApps",
  "metadata": {
    "arn": "arn:aws:appmesh:us-west-2:123456789012:mesh/apps/virtualGateway/
myVirtualGateway/gatewayRoute/myGatewayRoute",
    "createdAt": 1.591645848403E9,
    "lastUpdatedAt": 1.591645848403E9,
    "meshOwner": "123456789012",
    "resourceOwner": "123456789012",
    "uid": "ff111ff2-33ff-4ff4-ffff-111f111f111f",
    "version": 1
  },
  "spec": {
    "grpcRoute": null,
    "http2Route": {
      "action": {
        "target": {
          "virtualService": {
            "virtualServiceName": "myservicea.svc.cluster.local"
          }
        }
      }
    }
  }
}
```

```
  },
  "match": {
    "prefix": "/"
  }
},
"httpRoute": null
},
"status": {
  "status": "ACTIVE"
},
"virtualGatewayName": "myVirtualGateway"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeMesh

Describes an existing service mesh.

## Request Syntax

```
GET /v20190125/meshes/meshName?meshOwner=meshOwner HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh to describe.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
```

```
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "egressFilter": {
      "type": "string"
    },
    "serviceDiscovery": {
      "ipPreference": "string"
    }
  },
  "status": {
    "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.

### meshName

The name of the service mesh.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

The associated metadata for the service mesh.

Type: [ResourceMetadata](#) object

### spec

The associated specification for the service mesh.

Type: [MeshSpec](#) object

## [status](#)

The status of the service mesh.

Type: [MeshStatus](#) object

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### Example

This example describes a service mesh named `ecs-mesh`.

#### Sample Request

```
GET /v20190125/meshes HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190227T235715Z
Authorization: AUTHPARAMS
```

#### Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 572dcc56-18d6-4f86-b596-8e182f81afce
content-type: application/json
content-length: 114
date: Wed, 27 Feb 2019 23:57:15 GMT
x-envoy-upstream-service-time: 67
server: envoy
Connection: keep-alive

{
  "meshes": [{
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh",
    "meshName": "ecs-mesh"
  }],
}
```

```
"nextToken": null
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeRoute

Describes an existing route.

## Request Syntax

```
GET /v20190125/meshes/meshName/virtualRouter/virtualRouterName/routes/routeName?  
meshOwner=meshOwner HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh that the route resides in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### routeName

The name of the route to describe.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### virtualRouterName

The name of the virtual router that the route is associated with.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "routeName": "string",
  "spec": {
    "grpcRoute": {
      "action": {
        "weightedTargets": [
          {
            "port": number,
            "virtualNode": "string",
            "weight": number
          }
        ]
      }
    },
    "match": {
      "metadata": [
        {
          "invert": boolean,
          "match": { ... },
          "name": "string"
        }
      ]
    },
    "methodName": "string",
    "port": number,
    "serviceName": "string"
  }
}
```

```
    },
    "retryPolicy": {
      "grpcRetryEvents": [ "string" ],
      "httpRetryEvents": [ "string" ],
      "maxRetries": number,
      "perRetryTimeout": {
        "unit": "string",
        "value": number
      },
    },
    "tcpRetryEvents": [ "string" ]
  },
  "timeout": {
    "idle": {
      "unit": "string",
      "value": number
    },
    "perRequest": {
      "unit": "string",
      "value": number
    }
  }
},
"http2Route": {
  "action": {
    "weightedTargets": [
      {
        "port": number,
        "virtualNode": "string",
        "weight": number
      }
    ]
  },
  "match": {
    "headers": [
      {
        "invert": boolean,
        "match": { ... },
        "name": "string"
      }
    ],
    "method": "string",
    "path": {
      "exact": "string",
      "regex": "string"
    }
  }
}
```

```

    },
    "port": number,
    "prefix": "string",
    "queryParameters": [
      {
        "match": {
          "exact": "string"
        },
        "name": "string"
      }
    ],
    "scheme": "string"
  },
  "retryPolicy": {
    "httpRetryEvents": [ "string" ],
    "maxRetries": number,
    "perRetryTimeout": {
      "unit": "string",
      "value": number
    },
    "tcpRetryEvents": [ "string" ]
  },
  "timeout": {
    "idle": {
      "unit": "string",
      "value": number
    },
    "perRequest": {
      "unit": "string",
      "value": number
    }
  }
},
"httpRoute": {
  "action": {
    "weightedTargets": [
      {
        "port": number,
        "virtualNode": "string",
        "weight": number
      }
    ]
  },
  "match": {

```

```
    "headers": [
      {
        "invert": boolean,
        "match": { ... },
        "name": "string"
      }
    ],
    "method": "string",
    "path": {
      "exact": "string",
      "regex": "string"
    },
    "port": number,
    "prefix": "string",
    "queryParameters": [
      {
        "match": {
          "exact": "string"
        },
        "name": "string"
      }
    ],
    "scheme": "string"
  },
  "retryPolicy": {
    "httpRetryEvents": [ "string" ],
    "maxRetries": number,
    "perRetryTimeout": {
      "unit": "string",
      "value": number
    },
    "tcpRetryEvents": [ "string" ]
  },
  "timeout": {
    "idle": {
      "unit": "string",
      "value": number
    },
    "perRequest": {
      "unit": "string",
      "value": number
    }
  }
},
```

```
"priority": number,
"tcpRoute": {
  "action": {
    "weightedTargets": [
      {
        "port": number,
        "virtualNode": "string",
        "weight": number
      }
    ]
  },
  "match": {
    "port": number
  },
  "timeout": {
    "idle": {
      "unit": "string",
      "value": number
    }
  }
}
},
"status": {
  "status": "string"
},
"virtualRouterName": "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.

### meshName

The name of the service mesh that the route resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

The associated metadata for the route.

Type: [ResourceMetadata](#) object

### routeName

The name of the route.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### spec

The specifications of the route.

Type: [RouteSpec](#) object

### status

The status of the route.

Type: [RouteStatus](#) object

### virtualRouterName

The virtual router that the route is associated with.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

## InternalServerErrorException

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

HTTP Status Code: 500

## NotFoundException

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

## ServiceUnavailableException

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

## Example

This example describes a route named `colorteller-route` for the `colorteller-v1` virtual router in the `ecs-mesh` service mesh.

## Sample Request

```
GET /v20190125/meshes/ecs-mesh/virtualRouter/colorteller-vr/routes/colorteller-route
HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190228T000038Z
Authorization: AUTHPARAMS
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: f28c32fe-2d36-4fcf-a64e-01b45a4b034c
content-type: application/json
content-length: 504
date: Thu, 28 Feb 2019 00:00:38 GMT
x-envoy-upstream-service-time: 46
server: envoy
Connection: keep-alive

{
  "meshName": "ecs-mesh",
  "metadata": {
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualRouter/colorteller-vr/route/colorteller-route",
    "createdAt": 1.551311817276E9,
    "lastUpdatedAt": 1.551311817276E9,
    "uid": "1cf3109a-0d7f-438c-b17d-a3785f14ff7b",
    "version": 1
  },
  "routeName": "colorteller-route",
  "spec": {
    "httpRoute": {
      "action": {
        "weightedTargets": [{
          "virtualNode": "colorteller-vn",
          "weight": 1
        }]
      },
      "match": {
        "prefix": "/"
      }
    }
  }
}
```

```
  },
  "tcpRoute": null
},
"status": {
  "status": "ACTIVE"
},
"virtualRouterName": "colorteller-vr"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeVirtualGateway

Describes an existing virtual gateway.

## Request Syntax

```
GET /v20190125/meshes/meshName/virtualGateways/virtualGatewayName?meshOwner=meshOwner
HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh that the gateway route resides in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### virtualGatewayName

The name of the virtual gateway to describe.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request does not have a request body.

## Response Syntax

HTTP/1.1 200

Content-type: application/json

```
{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "backendDefaults": {
      "clientPolicy": {
        "tls": {
          "certificate": { ... },
          "enforce": boolean,
          "ports": [ number ],
          "validation": {
            "subjectAlternativeNames": {
              "match": {
                "exact": [ "string" ]
              }
            }
          },
          "trust": { ... }
        }
      }
    }
  },
  "listeners": [
    {
      "connectionPool": { ... },
      "healthCheck": {
        "healthyThreshold": number,
        "intervalMillis": number,
        "path": "string",
        "port": number,
        "protocol": "string",
```

```
    "timeoutMillis": number,
    "unhealthyThreshold": number
  },
  "portMapping": {
    "port": number,
    "protocol": "string"
  },
  "tls": {
    "certificate": { ... },
    "mode": "string",
    "validation": {
      "subjectAlternativeNames": {
        "match": {
          "exact": [ "string" ]
        }
      },
      "trust": { ... }
    }
  }
},
"logging": {
  "accessLog": { ... }
},
"status": {
  "status": "string"
},
"virtualGatewayName": "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.

### meshName

The name of the service mesh that the virtual gateway resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## metadata

An object that represents metadata for a resource.

Type: [ResourceMetadata](#) object

## spec

The specifications of the virtual gateway.

Type: [VirtualGatewaySpec](#) object

## status

The current status of the virtual gateway.

Type: [VirtualGatewayStatus](#) object

## virtualGatewayName

The name of the virtual gateway.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### BadRequestException

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### ForbiddenException

You don't have permissions to perform this action.

HTTP Status Code: 403

### InternalServerErrorException

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

HTTP Status Code: 500

## NotFoundException

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

## ServiceUnavailableException

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

## Example

The following example describes a virtual gateway named myVirtualGateway that is in the apps service mesh.

## Sample Request

```
GET /v20190125/meshes/apps/virtualGateways/myVirtualGateway HTTP/1.1
Host: appmesh.us-west-2.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.298 Python/3.6.0 Windows/10 botocore/1.13.34
```

```
X-Amz-Date: 20200706T145910Z
Authorization: AUTHPARAMS
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 52e5ee7c-7434-4ca0-885e-5c9d8a7bbaf8
content-type: application/json
content-length: 513
date: Mon, 06 Jul 2020 14:59:10 GMT
x-envoy-upstream-service-time: 19
server: envoy
Connection: keep-alive

{
  "meshName": "apps",
  "metadata": {
    "arn": "arn:aws:appmesh:us-west-2:123456789012:mesh/apps/virtualGateway/
myVirtualGateway",
    "createdAt": 1.594046778215E9,
    "lastUpdatedAt": 1.594046778215E9,
    "meshOwner": "123456789012",
    "resourceOwner": "123456789012",
    "uid": "ff111ff2-33ff-4ff4-ffff-111f111f111f",
    "version": 1
  },
  "spec": {
    "backendDefaults": null,
    "listeners": [{
      "healthCheck": null,
      "portMapping": {
        "port": 80,
        "protocol": "http2"
      },
      "tls": null
    }],
    "logging": null
  },
  "status": {
    "status": "ACTIVE"
  },
  "virtualGatewayName": "myVirtualGateway"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeVirtualNode

Describes an existing virtual node.

## Request Syntax

```
GET /v20190125/meshes/meshName/virtualNodes/virtualNodeName?meshOwner=meshOwner
HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh that the virtual node resides in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### virtualNodeName

The name of the virtual node to describe.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request does not have a request body.

## Response Syntax

HTTP/1.1 200

Content-type: application/json

```
{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "backendDefaults": {
      "clientPolicy": {
        "tls": {
          "certificate": { ... },
          "enforce": boolean,
          "ports": [ number ],
          "validation": {
            "subjectAlternativeNames": {
              "match": {
                "exact": [ "string" ]
              }
            },
            "trust": { ... }
          }
        }
      }
    },
    "backends": [
      { ... }
    ],
    "listeners": [
      {
        "connectionPool": { ... },
        "healthCheck": {
          "healthyThreshold": number,
          "intervalMillis": number,
```

```
    "path": "string",
    "port": number,
    "protocol": "string",
    "timeoutMillis": number,
    "unhealthyThreshold": number
  },
  "outlierDetection": {
    "baseEjectionDuration": {
      "unit": "string",
      "value": number
    },
    "interval": {
      "unit": "string",
      "value": number
    },
    "maxEjectionPercent": number,
    "maxServerErrors": number
  },
  "portMapping": {
    "port": number,
    "protocol": "string"
  },
  "timeout": { ... },
  "tls": {
    "certificate": { ... },
    "mode": "string",
    "validation": {
      "subjectAlternativeNames": {
        "match": {
          "exact": [ "string" ]
        }
      },
      "trust": { ... }
    }
  }
},
],
"logging": {
  "accessLog": { ... }
},
"serviceDiscovery": { ... }
},
"status": {
  "status": "string"
}
```

```
  },  
  "virtualNodeName": "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.

### meshName

The name of the service mesh that the virtual node resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

The associated metadata for the virtual node.

Type: [ResourceMetadata](#) object

### spec

The specifications of the virtual node.

Type: [VirtualNodeSpec](#) object

### status

The current status for the virtual node.

Type: [VirtualNodeStatus](#) object

### virtualNodeName

The name of the virtual node.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

## Example

This example describes a virtual node named `colorteller-vn` in the `ecs-mesh` service mesh.

### Sample Request

```
GET /v20190125/meshes/ecs-mesh/virtualNodes/colorteller-vn HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190228T000300Z
Authorization: AUTHPARAMS
```

### Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 083f9376-9799-4142-978b-42f5966192b8
content-type: application/json
content-length: 654
date: Thu, 28 Feb 2019 00:03:00 GMT
x-envoy-upstream-service-time: 60
server: envoy
Connection: keep-alive

{
  "meshName": "ecs-mesh",
  "metadata": {
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualNode/colorteller-
vn",
    "createdAt": 1.551311799179E9,
    "lastUpdatedAt": 1.551311799179E9,
    "uid": "0999e53a-8e0e-4c4c-8764-ae8ebecc296d",
    "version": 1
  },
  "spec": {
    "backends": [],
```

```
"listeners": [{
  "healthCheck": {
    "healthyThreshold": 2,
    "intervalMillis": 5000,
    "path": "/ping",
    "port": 9080,
    "protocol": "http",
    "timeoutMillis": 2000,
    "unhealthyThreshold": 2
  },
  "portMapping": {
    "port": 9080,
    "protocol": "http"
  }
}],
"logging": null,
"serviceDiscovery": {
  "dns": {
    "hostname": "colorteller.default.svc.cluster.local"
  }
},
"status": {
  "status": "ACTIVE"
},
"virtualNodeName": "colorteller-vn"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)

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

# DescribeVirtualRouter

Describes an existing virtual router.

## Request Syntax

```
GET /v20190125/meshes/meshName/virtualRouters/virtualRouterName?meshOwner=meshOwner
HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh that the virtual router resides in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### virtualRouterName

The name of the virtual router to describe.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "listeners": [
      {
        "portMapping": {
          "port": number,
          "protocol": "string"
        }
      }
    ]
  },
  "status": {
    "status": "string"
  },
  "virtualRouterName": "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.

### meshName

The name of the service mesh that the virtual router resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

The associated metadata for the virtual router.

Type: [ResourceMetadata](#) object

### spec

The specifications of the virtual router.

Type: [VirtualRouterSpec](#) object

### status

The current status of the virtual router.

Type: [VirtualRouterStatus](#) object

### virtualRouterName

The name of the virtual router.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

## NotFoundException

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

## ServiceUnavailableException

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### Example

This example describes a virtual router named `colorteller-vr` in the `ecs-mesh` service mesh.

### Sample Request

```
GET /v20190125/meshes/ecs-mesh/virtualRouters/colorteller-vr HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190228T000509Z
Authorization: AUTHPARAMS
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: e7e0fbe8-2015-429c-bbb5-c736c09b2500
content-type: application/json
content-length: 337
date: Thu, 28 Feb 2019 00:05:10 GMT
x-envoy-upstream-service-time: 68
server: envoy
Connection: keep-alive

{
  "meshName": "ecs-mesh",
  "metadata": {
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualRouter/
colorteller-vr",
    "createdAt": 1551311805.476,
    "lastUpdatedAt": 1551311805.476,
    "uid": "79628d34-8c17-42ba-83a5-8a42fd17ec5a",
    "version": 1
  },
  "spec": {
    "listeners": [
      {
        "portMapping": {
          "port": 9080,
          "protocol": "http"
        }
      }
    ]
  },
  "status": {
    "status": "ACTIVE"
  },
  "virtualRouterName": "colorteller-vr"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)

- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# DescribeVirtualService

Describes an existing virtual service.

## Request Syntax

```
GET /v20190125/meshes/meshName/virtualServices/virtualServiceName?meshOwner=meshOwner
HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh that the virtual service resides in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### virtualServiceName

The name of the virtual service to describe.

Required: Yes

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
```

```
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "provider": { ... }
  },
  "status": {
    "status": "string"
  },
  "virtualServiceName": "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.

### meshName

The name of the service mesh that the virtual service resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

An object that represents metadata for a resource.

Type: [ResourceMetadata](#) object

### spec

The specifications of the virtual service.

Type: [VirtualServiceSpec](#) object

### **status**

The current status of the virtual service.

Type: [VirtualServiceStatus](#) object

### **virtualServiceName**

The name of the virtual service.

Type: String

## **Errors**

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### Example

This example describes a virtual service named `colorgateway.default.svc.cluster.local` in the `ecs-mesh` service mesh.

### Sample Request

```
GET /v20190125/meshes/ecs-mesh/virtualServices/colorgateway.default.svc.cluster.local
  HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190228T000705Z
Authorization: AUTHPARAMS
```

### Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 3b822acc-bd78-4ea6-9831-8cd989dd37ec
content-type: application/json
content-length: 456
date: Thu, 28 Feb 2019 00:07:04 GMT
```

```
x-envoy-upstream-service-time: 30
server: envoy
Connection: keep-alive

{
  "meshName": "ecs-mesh",
  "metadata": {
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualService/
colorgateway.default.svc.cluster.local",
    "createdAt": 1.551311807444E9,
    "lastUpdatedAt": 1.551311807444E9,
    "uid": "dd06064b-e542-40a9-bbc7-e381a47ea0e0",
    "version": 1
  },
  "spec": {
    "provider": {
      "virtualNode": {
        "virtualNodeName": "colorgateway-vn"
      },
      "virtualRouter": null
    }
  },
  "status": {
    "status": "ACTIVE"
  },
  "virtualServiceName": "colorgateway.default.svc.cluster.local"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)

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

# ListGatewayRoutes

Returns a list of existing gateway routes that are associated to a virtual gateway.

## Request Syntax

```
GET /v20190125/meshes/meshName/virtualGateway/virtualGatewayName/gatewayRoutes?  
limit=limit&meshOwner=meshOwner&nextToken=nextToken HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### limit

The maximum number of results returned by ListGatewayRoutes in paginated output. When you use this parameter, ListGatewayRoutes returns only limit results in a single page along with a nextToken response element. You can see the remaining results of the initial request by sending another ListGatewayRoutes request with the returned nextToken value. This value can be between 1 and 100. If you don't use this parameter, ListGatewayRoutes returns up to 100 results and a nextToken value if applicable.

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

### meshName

The name of the service mesh to list gateway routes in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

## nextToken

The `nextToken` value returned from a previous paginated `ListGatewayRoutes` request where `limit` was used and the results exceeded the value of that parameter. Pagination continues from the end of the previous results that returned the `nextToken` value.

## virtualGatewayName

The name of the virtual gateway to list gateway routes in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "gatewayRoutes": [
    {
      "arn": "string",
      "createdAt": number,
      "gatewayRouteName": "string",
      "lastUpdatedAt": number,
      "meshName": "string",
      "meshOwner": "string",
      "resourceOwner": "string",
      "version": number,
      "virtualGatewayName": "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.

### gatewayRoutes

The list of existing gateway routes for the specified service mesh and virtual gateway.

Type: Array of [GatewayRouteRef](#) objects

### nextToken

The nextToken value to include in a future ListGatewayRoutes request. When the results of a ListGatewayRoutes request exceed limit, you can use this value to retrieve the next page of results. This value is null when there are no more results to return.

Type: String

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

## ServiceUnavailableException

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### Example

The following example lists gateway routes that are associated with the virtual gateway named myVirtualGateway in the apps service mesh.

### Sample Request

```
GET /v20190125/meshes/apps/virtualGateway/myVirtualGateway/gatewayRoutes HTTP/1.1
Host: appmesh.us-west-2.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.298 Python/3.6.0 Windows/10 botocore/1.13.34
X-Amz-Date: 20200608T190053Z
Authorization: AUTHPARAMS
```

### Sample Response

```
HTTP/1.1 200 OK
```

```
x-amzn-requestid: b82c3691-a7a9-4d99-acf6-0f9a81db9beb
content-type: application/json
content-length: 393
date: Mon, 08 Jun 2020 20:52:25 GMT
x-envoy-upstream-service-time: 26
server: envoy
Connection: keep-alive

{
  "gatewayRoutes": [{
    "arn": "arn:aws:appmesh:us-west-2:123456789012:mesh/apps/virtualGateway/myVirtualGateway/gatewayRoute/myGatewayRoute",
    "createdAt": 1.591642091122E9,
    "gatewayRouteName": "myGatewayRoute",
    "lastUpdatedAt": 1.591642091122E9,
    "meshName": "myApps",
    "meshOwner": "123456789012",
    "resourceOwner": "123456789012",
    "version": 1,
    "virtualGatewayName": "myVirtualGateway"
  }],
  "nextToken": null
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)



# ListMeshes

Returns a list of existing service meshes.

## Request Syntax

```
GET /v20190125/meshes?limit=limit&nextToken=nextToken HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### limit

The maximum number of results returned by ListMeshes in paginated output. When you use this parameter, ListMeshes returns only `limit` results in a single page along with a `nextToken` response element. You can see the remaining results of the initial request by sending another ListMeshes request with the returned `nextToken` value. This value can be between 1 and 100. If you don't use this parameter, ListMeshes returns up to 100 results and a `nextToken` value if applicable.

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

### nextToken

The `nextToken` value returned from a previous paginated ListMeshes request where `limit` was used and the results exceeded the value of that parameter. Pagination continues from the end of the previous results that returned the `nextToken` value.

#### Note

This token should be treated as an opaque identifier that is used only to retrieve the next items in a list and not for other programmatic purposes.

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "meshes": [
    {
      "arn": "string",
      "createdAt": number,
      "lastUpdatedAt": number,
      "meshName": "string",
      "meshOwner": "string",
      "resourceOwner": "string",
      "version": 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.

### meshes

The list of existing service meshes.

Type: Array of [MeshRef](#) objects

### nextToken

The nextToken value to include in a future ListMeshes request. When the results of a ListMeshes request exceed limit, you can use this value to retrieve the next page of results. This value is null when there are no more results to return.

Type: String

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

## Example

This example lists the service meshes for an account in the us-east-1 Region.

### Sample Request

```
GET /v20190125/meshes HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190227T235715Z
Authorization: AUTHPARAMS
```

### Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 572dcc56-18d6-4f86-b596-8e182f81afce
content-type: application/json
content-length: 114
date: Wed, 27 Feb 2019 23:57:15 GMT
x-envoy-upstream-service-time: 67
server: envoy
Connection: keep-alive

{
  "meshes": [
    {
      "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh",
      "meshName": "ecs-mesh"
    }
  ],
  "nextToken": null
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ListRoutes

Returns a list of existing routes in a service mesh.

## Request Syntax

```
GET /v20190125/meshes/meshName/virtualRouter/virtualRouterName/routes?  
limit=limit&meshOwner=meshOwner&nextToken=nextToken HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### limit

The maximum number of results returned by ListRoutes in paginated output. When you use this parameter, ListRoutes returns only limit results in a single page along with a nextToken response element. You can see the remaining results of the initial request by sending another ListRoutes request with the returned nextToken value. This value can be between 1 and 100. If you don't use this parameter, ListRoutes returns up to 100 results and a nextToken value if applicable.

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

### meshName

The name of the service mesh to list routes in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

## nextToken

The `nextToken` value returned from a previous paginated `ListRoutes` request where `limit` was used and the results exceeded the value of that parameter. Pagination continues from the end of the previous results that returned the `nextToken` value.

## virtualRouterName

The name of the virtual router to list routes in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "nextToken": "string",
  "routes": [
    {
      "arn": "string",
      "createdAt": number,
      "lastUpdatedAt": number,
      "meshName": "string",
      "meshOwner": "string",
      "resourceOwner": "string",
      "routeName": "string",
      "version": number,
      "virtualRouterName": "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 `nextToken` value to include in a future `ListRoutes` request. When the results of a `ListRoutes` request exceed `limit`, you can use this value to retrieve the next page of results. This value is `null` when there are no more results to return.

Type: String

### routes

The list of existing routes for the specified service mesh and virtual router.

Type: Array of [RouteRef](#) objects

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

## ServiceUnavailableException

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### Example

The following example lists the routes that are associated with the `colorteller-vr` virtual router in the `ecs-mesh` service mesh.

### Sample Request

```
GET /v20190125/meshes/ecs-mesh/virtualRouter/colorteller-vr/routes HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190227T235954Z
Authorization: AUTHPARAMS
```

### Sample Response

```
HTTP/1.1 200 OK
```

```
x-amzn-requestid: 5a783529-a278-44d2-b450-4e23898cb180
content-type: application/json
content-length: 236
date: Wed, 27 Feb 2019 23:59:54 GMT
x-envoy-upstream-service-time: 56
server: envoy
Connection: keep-alive

{
  "nextToken": null,
  "routes": [
    {
      "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualRouter/colorteller-vr/route/colorteller-route",
      "meshName": "ecs-mesh",
      "routeName": "colorteller-route",
      "virtualRouterName": "colorteller-vr"
    }
  ]
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ListTagsForResource

List the tags for an App Mesh resource.

## Request Syntax

```
GET /v20190125/tags?limit=limit&nextToken=nextToken&resourceArn=resourceArn HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### limit

The maximum number of tag results returned by `ListTagsForResource` in paginated output. When this parameter is used, `ListTagsForResource` returns only `limit` results in a single page along with a `nextToken` response element. You can see the remaining results of the initial request by sending another `ListTagsForResource` request with the returned `nextToken` value. This value can be between 1 and 100. If you don't use this parameter, `ListTagsForResource` returns up to 100 results and a `nextToken` value if applicable.

Valid Range: Minimum value of 1. Maximum value of 50.

### nextToken

The `nextToken` value returned from a previous paginated `ListTagsForResource` request where `limit` was used and the results exceeded the value of that parameter. Pagination continues from the end of the previous results that returned the `nextToken` value.

### resourceArn

The Amazon Resource Name (ARN) that identifies the resource to list the tags for.

Required: Yes

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "nextToken": "string",
  "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.

### [nextToken](#)

The `nextToken` value to include in a future `ListTagsForResource` request. When the results of a `ListTagsForResource` request exceed `limit`, you can use this value to retrieve the next page of results. This value is `null` when there are no more results to return.

Type: String

### [tags](#)

The tags for the resource.

Type: Array of [TagRef](#) objects

Array Members: Minimum number of 0 items. Maximum number of 50 items.

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## **See Also**

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)

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

# ListVirtualGateways

Returns a list of existing virtual gateways in a service mesh.

## Request Syntax

```
GET /v20190125/meshes/meshName/virtualGateways?  
limit=limit&meshOwner=meshOwner&nextToken=nextToken HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### limit

The maximum number of results returned by ListVirtualGateways in paginated output. When you use this parameter, ListVirtualGateways returns only limit results in a single page along with a nextToken response element. You can see the remaining results of the initial request by sending another ListVirtualGateways request with the returned nextToken value. This value can be between 1 and 100. If you don't use this parameter, ListVirtualGateways returns up to 100 results and a nextToken value if applicable.

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

### meshName

The name of the service mesh to list virtual gateways in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

## nextToken

The `nextToken` value returned from a previous paginated `ListVirtualGateways` request where `limit` was used and the results exceeded the value of that parameter. Pagination continues from the end of the previous results that returned the `nextToken` value.

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "nextToken": "string",
  "virtualGateways": [
    {
      "arn": "string",
      "createdAt": number,
      "lastUpdatedAt": number,
      "meshName": "string",
      "meshOwner": "string",
      "resourceOwner": "string",
      "version": number,
      "virtualGatewayName": "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 `nextToken` value to include in a future `ListVirtualGateways` request. When the results of a `ListVirtualGateways` request exceed `limit`, you can use this value to retrieve the next page of results. This value is `null` when there are no more results to return.

Type: String

## [virtualGateways](#)

The list of existing virtual gateways for the specified service mesh.

Type: Array of [VirtualGatewayRef](#) objects

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### Example

The following example lists the virtual gateways that are in the apps service mesh.

### Sample Request

```
GET /v20190125/meshes/apps/virtualGateways HTTP/1.1
Host: appmesh.us-west-2.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.298 Python/3.6.0 Windows/10 botocore/1.13.34
X-Amz-Date: 20200608T185502Z
Authorization: AUTHPARAMS
```

### Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 3add7f54-944b-464a-93ed-57589ed5158d
content-type: application/json
content-length: 331
date: Mon, 08 Jun 2020 20:53:50 GMT
x-envoy-upstream-service-time: 24
```

```
server: envoy
Connection: keep-alive

{
  "nextToken": null,
  "virtualGateways": [{
    "arn": "arn:aws:appmesh:us-west-2:123456789012:mesh/apps/virtualGateway/
myVirtualGateway",
    "createdAt": 1.591630673505E9,
    "lastUpdatedAt": 1.591630673505E9,
    "meshName": "apps",
    "meshOwner": "123456789012",
    "resourceOwner": "123456789012",
    "version": 1,
    "virtualGatewayName": "myVirtualGateway"
  }]
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ListVirtualNodes

Returns a list of existing virtual nodes.

## Request Syntax

```
GET /v20190125/meshes/meshName/virtualNodes?  
limit=limit&meshOwner=meshOwner&nextToken=nextToken HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### limit

The maximum number of results returned by `ListVirtualNodes` in paginated output. When you use this parameter, `ListVirtualNodes` returns only `limit` results in a single page along with a `nextToken` response element. You can see the remaining results of the initial request by sending another `ListVirtualNodes` request with the returned `nextToken` value. This value can be between 1 and 100. If you don't use this parameter, `ListVirtualNodes` returns up to 100 results and a `nextToken` value if applicable.

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

### meshName

The name of the service mesh to list virtual nodes in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

## nextToken

The `nextToken` value returned from a previous paginated `ListVirtualNodes` request where `limit` was used and the results exceeded the value of that parameter. Pagination continues from the end of the previous results that returned the `nextToken` value.

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "nextToken": "string",
  "virtualNodes": [
    {
      "arn": "string",
      "createdAt": number,
      "lastUpdatedAt": number,
      "meshName": "string",
      "meshOwner": "string",
      "resourceOwner": "string",
      "version": number,
      "virtualNodeName": "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 `nextToken` value to include in a future `ListVirtualNodes` request. When the results of a `ListVirtualNodes` request exceed `limit`, you can use this value to retrieve the next page of results. This value is `null` when there are no more results to return.

Type: String

## [virtualNodes](#)

The list of existing virtual nodes for the specified service mesh.

Type: Array of [VirtualNodeRef](#) objects

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

## Example

The following example lists the virtual nodes that reside in the `ecs-mesh` service mesh.

### Sample Request

```
GET /v20190125/meshes/ecs-mesh/virtualNodes HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190227T235730Z
Authorization: AUTHPARAMS
```

### Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 8fbb6448-5519-4f45-97d1-ed4d3238aa69
content-type: application/json
content-length: 941
date: Wed, 27 Feb 2019 23:57:31 GMT
x-envoy-upstream-service-time: 108
server: envoy
```

Connection: keep-alive

```
{
  "nextToken": null,
  "virtualNodes": [
    {
      "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualNode/colorteller-vn",
      "meshName": "ecs-mesh",
      "virtualNodeName": "colorteller-vn"
    },
    {
      "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualNode/tcpecho-vn",
      "meshName": "ecs-mesh",
      "virtualNodeName": "tcpecho-vn"
    },
    {
      "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualNode/colorteller-black-vn",
      "meshName": "ecs-mesh",
      "virtualNodeName": "colorteller-black-vn"
    },
    {
      "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualNode/colorgateway-vn",
      "meshName": "ecs-mesh",
      "virtualNodeName": "colorgateway-vn"
    },
    {
      "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualNode/colorteller-red-vn",
      "meshName": "ecs-mesh",
      "virtualNodeName": "colorteller-red-vn"
    },
    {
      "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualNode/colorteller-blue-vn",
      "meshName": "ecs-mesh",
      "virtualNodeName": "colorteller-blue-vn"
    }
  ]
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ListVirtualRouters

Returns a list of existing virtual routers in a service mesh.

## Request Syntax

```
GET /v20190125/meshes/meshName/virtualRouters?  
limit=limit&meshOwner=meshOwner&nextToken=nextToken HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### limit

The maximum number of results returned by `ListVirtualRouters` in paginated output. When you use this parameter, `ListVirtualRouters` returns only `limit` results in a single page along with a `nextToken` response element. You can see the remaining results of the initial request by sending another `ListVirtualRouters` request with the returned `nextToken` value. This value can be between 1 and 100. If you don't use this parameter, `ListVirtualRouters` returns up to 100 results and a `nextToken` value if applicable.

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

### meshName

The name of the service mesh to list virtual routers in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

## nextToken

The `nextToken` value returned from a previous paginated `ListVirtualRouters` request where `limit` was used and the results exceeded the value of that parameter. Pagination continues from the end of the previous results that returned the `nextToken` value.

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "nextToken": "string",
  "virtualRouters": [
    {
      "arn": "string",
      "createdAt": number,
      "lastUpdatedAt": number,
      "meshName": "string",
      "meshOwner": "string",
      "resourceOwner": "string",
      "version": number,
      "virtualRouterName": "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 `nextToken` value to include in a future `ListVirtualRouters` request. When the results of a `ListVirtualRouters` request exceed `limit`, you can use this value to retrieve the next page of results. This value is `null` when there are no more results to return.

Type: String

## [virtualRouters](#)

The list of existing virtual routers for the specified service mesh.

Type: Array of [VirtualRouterRef](#) objects

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### Example

The following example lists the virtual routers that reside in the `ecs-mesh` service mesh.

### Sample Request

```
GET /v20190125/meshes/ecs-mesh/virtualRouters HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190227T235738Z
Authorization: AUTHPARAMS
```

### Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: ec2895b6-ffaf-4d0a-b1dd-830901321588
content-type: application/json
content-length: 341
date: Wed, 27 Feb 2019 23:57:38 GMT
x-envoy-upstream-service-time: 43
```

```
server: envoy
Connection: keep-alive

{
  "nextToken": null,
  "virtualRouters": [
    {
      "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualRouter/
colorteller-vr",
      "meshName": "ecs-mesh",
      "virtualRouterName": "colorteller-vr"
    },
    {
      "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualRouter/
colorgateway-vr",
      "meshName": "ecs-mesh",
      "virtualRouterName": "colorgateway-vr"
    }
  ]
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ListVirtualServices

Returns a list of existing virtual services in a service mesh.

## Request Syntax

```
GET /v20190125/meshes/meshName/virtualServices?  
limit=limit&meshOwner=meshOwner&nextToken=nextToken HTTP/1.1
```

## URI Request Parameters

The request uses the following URI parameters.

### limit

The maximum number of results returned by `ListVirtualServices` in paginated output. When you use this parameter, `ListVirtualServices` returns only `limit` results in a single page along with a `nextToken` response element. You can see the remaining results of the initial request by sending another `ListVirtualServices` request with the returned `nextToken` value. This value can be between 1 and 100. If you don't use this parameter, `ListVirtualServices` returns up to 100 results and a `nextToken` value if applicable.

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

### meshName

The name of the service mesh to list virtual services in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

## nextToken

The `nextToken` value returned from a previous paginated `ListVirtualServices` request where `limit` was used and the results exceeded the value of that parameter. Pagination continues from the end of the previous results that returned the `nextToken` value.

## Request Body

The request does not have a request body.

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "nextToken": "string",
  "virtualServices": [
    {
      "arn": "string",
      "createdAt": number,
      "lastUpdatedAt": number,
      "meshName": "string",
      "meshOwner": "string",
      "resourceOwner": "string",
      "version": number,
      "virtualServiceName": "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 `nextToken` value to include in a future `ListVirtualServices` request. When the results of a `ListVirtualServices` request exceed `limit`, you can use this value to retrieve the next page of results. This value is `null` when there are no more results to return.

Type: String

## [virtualServices](#)

The list of existing virtual services for the specified service mesh.

Type: Array of [VirtualServiceRef](#) objects

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### Example

This example lists the virtual services in the ecs-mesh service mesh.

### Sample Request

```
GET /v20190125/meshes/ecs-mesh/virtualServices HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190227T235746Z
Authorization: AUTHPARAMS
```

### Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 1e2ccc92-46aa-4541-ae10-9790266e7436
content-type: application/json
content-length: 629
date: Wed, 27 Feb 2019 23:57:46 GMT
x-envoy-upstream-service-time: 9
```

```
server: envoy
Connection: keep-alive

{
  "nextToken": null,
  "virtualServices": [
    {
      "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualService/
tcpecho.default.svc.cluster.local",
      "meshName": "ecs-mesh",
      "virtualServiceName": "tcpecho.default.svc.cluster.local"
    },
    {
      "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualService/
colorteller.default.svc.cluster.local",
      "meshName": "ecs-mesh",
      "virtualServiceName": "colorteller.default.svc.cluster.local"
    },
    {
      "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualService/
colorgateway.default.svc.cluster.local",
      "meshName": "ecs-mesh",
      "virtualServiceName": "colorgateway.default.svc.cluster.local"
    }
  ]
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)

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

# TagResource

Associates the specified tags to a resource with the specified `resourceArn`. If existing tags on a resource aren't specified in the request parameters, they aren't changed. When a resource is deleted, the tags associated with that resource are also deleted.

## Request Syntax

```
PUT /v20190125/tag?resourceArn=resourceArn HTTP/1.1
Content-type: application/json
```

```
{
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ]
}
```

## URI Request Parameters

The request uses the following URI parameters.

### resourceArn

The Amazon Resource Name (ARN) of the resource to add tags to.

Required: Yes

## Request Body

The request accepts the following data in JSON format.

### tags

The tags to add to the resource. A tag is an array of key-value pairs. Tag keys can have a maximum character length of 128 characters, and tag values can have a maximum length of 256 characters.

Type: Array of [TagRef](#) objects

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Required: Yes

## Response Syntax

```
HTTP/1.1 200
```

## Response Elements

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

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## TooManyTagsException

The request exceeds the maximum allowed number of tags allowed per resource. The current limit is 50 user tags per resource. You must reduce the number of tags in the request. None of the tags in this request were applied.

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 V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# UntagResource

Deletes specified tags from a resource.

## Request Syntax

```
PUT /v20190125/untag?resourceArn=resourceArn HTTP/1.1
Content-type: application/json

{
  "tagKeys": [ "string" ]
}
```

## URI Request Parameters

The request uses the following URI parameters.

### resourceArn

The Amazon Resource Name (ARN) of the resource to delete tags from.

Required: Yes

## Request Body

The request accepts the following data in JSON format.

### tagKeys

The keys of the tags to be removed.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 50 items.

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

Required: Yes

# Response Syntax

```
HTTP/1.1 200
```

## Response Elements

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

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# UpdateGatewayRoute

Updates an existing gateway route that is associated to a specified virtual gateway in a service mesh.

## Request Syntax

```
PUT /v20190125/meshes/meshName/virtualGateway/virtualGatewayName/
gatewayRoutes/gatewayRouteName?meshOwner=meshOwner HTTP/1.1
Content-type: application/json
```

```
{
  "clientToken": "string",
  "spec": {
    "grpcRoute": {
      "action": {
        "rewrite": {
          "hostname": {
            "defaultTargetHostname": "string"
          }
        },
        "target": {
          "port": number,
          "virtualService": {
            "virtualServiceName": "string"
          }
        }
      }
    },
    "match": {
      "hostname": {
        "exact": "string",
        "suffix": "string"
      },
      "metadata": [
        {
          "invert": boolean,
          "match": { ... },
          "name": "string"
        }
      ],
      "port": number,
      "serviceName": "string"
    }
  }
}
```

```

},
"http2Route": {
  "action": {
    "rewrite": {
      "hostname": {
        "defaultTargetHostname": "string"
      },
      "path": {
        "exact": "string"
      },
      "prefix": {
        "defaultPrefix": "string",
        "value": "string"
      }
    },
    "target": {
      "port": number,
      "virtualService": {
        "virtualServiceName": "string"
      }
    }
  },
  "match": {
    "headers": [
      {
        "invert": boolean,
        "match": { ... },
        "name": "string"
      }
    ],
    "hostname": {
      "exact": "string",
      "suffix": "string"
    },
    "method": "string",
    "path": {
      "exact": "string",
      "regex": "string"
    },
    "port": number,
    "prefix": "string",
    "queryParameters": [
      {
        "match": {

```

```
        "exact": "string"
      },
      "name": "string"
    }
  ]
}
},
"httpRoute": {
  "action": {
    "rewrite": {
      "hostname": {
        "defaultTargetHostname": "string"
      },
      "path": {
        "exact": "string"
      },
      "prefix": {
        "defaultPrefix": "string",
        "value": "string"
      }
    },
    "target": {
      "port": number,
      "virtualService": {
        "virtualServiceName": "string"
      }
    }
  },
  "match": {
    "headers": [
      {
        "invert": boolean,
        "match": { ... },
        "name": "string"
      }
    ],
    "hostname": {
      "exact": "string",
      "suffix": "string"
    },
    "method": "string",
    "path": {
      "exact": "string",
      "regex": "string"
    }
  }
}
```

```
    },
    "port": number,
    "prefix": "string",
    "queryParameters": [
      {
        "match": {
          "exact": "string"
        },
        "name": "string"
      }
    ]
  },
  "priority": number
}
```

## URI Request Parameters

The request uses the following URI parameters.

### gatewayRouteName

The name of the gateway route to update.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshName

The name of the service mesh that the gateway route resides in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

## virtualGatewayName

The name of the virtual gateway that the gateway route is associated with.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request accepts the following data in JSON format.

### clientToken

Unique, case-sensitive identifier that you provide to ensure the idempotency of the request. Up to 36 letters, numbers, hyphens, and underscores are allowed.

Type: String

Required: No

### spec

The new gateway route specification to apply. This overwrites the existing data.

Type: [GatewayRouteSpec](#) object

Required: Yes

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "gatewayRouteName": "string",
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
```

```
    "uid": "string",
    "version": number
  },
  "spec": {
    "grpcRoute": {
      "action": {
        "rewrite": {
          "hostname": {
            "defaultTargetHostname": "string"
          }
        },
        "target": {
          "port": number,
          "virtualService": {
            "virtualServiceName": "string"
          }
        }
      }
    },
    "match": {
      "hostname": {
        "exact": "string",
        "suffix": "string"
      },
      "metadata": [
        {
          "invert": boolean,
          "match": { ... },
          "name": "string"
        }
      ],
      "port": number,
      "serviceName": "string"
    }
  },
  "http2Route": {
    "action": {
      "rewrite": {
        "hostname": {
          "defaultTargetHostname": "string"
        },
        "path": {
          "exact": "string"
        }
      },
      "prefix": {
```

```
        "defaultPrefix": "string",
        "value": "string"
    }
},
"target": {
    "port": number,
    "virtualService": {
        "virtualServiceName": "string"
    }
}
},
"match": {
    "headers": [
        {
            "invert": boolean,
            "match": { ... },
            "name": "string"
        }
    ],
    "hostname": {
        "exact": "string",
        "suffix": "string"
    },
    "method": "string",
    "path": {
        "exact": "string",
        "regex": "string"
    },
    "port": number,
    "prefix": "string",
    "queryParameters": [
        {
            "match": {
                "exact": "string"
            },
            "name": "string"
        }
    ]
}
},
"httpRoute": {
    "action": {
        "rewrite": {
            "hostname": {
```

```
    "defaultTargetHostname": "string"
  },
  "path": {
    "exact": "string"
  },
  "prefix": {
    "defaultPrefix": "string",
    "value": "string"
  }
},
"target": {
  "port": number,
  "virtualService": {
    "virtualServiceName": "string"
  }
}
},
"match": {
  "headers": [
    {
      "invert": boolean,
      "match": { ... },
      "name": "string"
    }
  ],
  "hostname": {
    "exact": "string",
    "suffix": "string"
  },
  "method": "string",
  "path": {
    "exact": "string",
    "regex": "string"
  },
  "port": number,
  "prefix": "string",
  "queryParameters": [
    {
      "match": {
        "exact": "string"
      },
      "name": "string"
    }
  ]
}
```

```
    }  
  },  
  "priority": number  
},  
"status": {  
  "status": "string"  
},  
"virtualGatewayName": "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.

### gatewayRouteName

The name of the gateway route.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### meshName

The name of the service mesh that the resource resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

An object that represents metadata for a resource.

Type: [ResourceMetadata](#) object

### spec

The specifications of the gateway route.

Type: [GatewayRouteSpec](#) object

### status

The status of the gateway route.

Type: [GatewayRouteStatus](#) object

### [virtualGatewayName](#)

The virtual gateway that the gateway route is associated with.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ConflictException**

The request contains a client token that was used for a previous update resource call with different specifications. Try the request again with a new client token.

HTTP Status Code: 409

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **LimitExceededException**

You have exceeded a service limit for your account. For more information, see [Service Limits](#) in the *AWS App Mesh User Guide*.

HTTP Status Code: 400

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## **Examples**

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### **Example**

The following example updates a gateway route named myGatewayRoute that is associated to the myVirtualGateway virtual gateway in the apps service mesh.

### **Sample Request**

```
PUT /v20190125/meshes/apps/virtualGateway/myVirtualGateway/gatewayRoutes/myGatewayRoute
HTTP/1.1
Host: appmesh.us-west-2.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.298 Python/3.6.0 Windows/10 botocore/1.13.34
X-Amz-Date: 20200608T192800Z
Authorization: AUTHPARAMS
```

```
{
  "spec": {
    "http2Route": {
      "action": {
        "target": {
          "virtualService": {
            "virtualServiceName": "myserviceb.svc.cluster.local"
          }
        }
      },
    },
    "match": {
      "prefix": "/"
    }
  }
},
"clientToken": "1aa1111a-1111-1111-111a-a1a1a1aa111a"
}
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 6c253f78-e011-4533-a926-e75416f06641
content-type: application/json
content-length: 617
date: Mon, 08 Jun 2020 19:28:00 GMT
x-envoy-upstream-service-time: 76
server: envoy
Connection: keep-alive

{
  "gatewayRouteName": "myGatewayRoute",
  "meshName": "myApps",
  "metadata": {
    "arn": "arn:aws:appmesh:us-west-2:123456789012:mesh/apps/virtualGateway/myVirtualGateway/gatewayRoute/myGatewayRoute",
    "createdAt": 1.591642091122E9,
    "lastUpdatedAt": 1.591644481211E9,
    "meshOwner": "123456789012",
    "resourceOwner": "123456789012",
    "uid": "ff1111ff2-33ff-4ff4-ffff-111f111f111f",
    "version": 2
  },
  "spec": {
```

```
"grpcRoute": null,
"http2Route": {
  "action": {
    "target": {
      "virtualService": {
        "virtualServiceName": "myserviceb.svc.cluster.local"
      }
    }
  },
},
"match": {
  "prefix": "/"
},
},
"httpRoute": null
},
"status": {
  "status": "ACTIVE"
},
"virtualGatewayName": "myVirtualGateway"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# UpdateMesh

Updates an existing service mesh.

## Request Syntax

```
PUT /v20190125/meshes/meshName HTTP/1.1
Content-type: application/json
```

```
{
  "clientToken": "string",
  "spec": {
    "egressFilter": {
      "type": "string"
    },
    "serviceDiscovery": {
      "ipPreference": "string"
    }
  }
}
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh to update.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request accepts the following data in JSON format.

### clientToken

Unique, case-sensitive identifier that you provide to ensure the idempotency of the request. Up to 36 letters, numbers, hyphens, and underscores are allowed.

Type: String

Required: No

### spec

The service mesh specification to apply.

Type: [MeshSpec](#) object

Required: No

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "egressFilter": {
      "type": "string"
    },
    "serviceDiscovery": {
      "ipPreference": "string"
    }
  },
  "status": {
    "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.

### meshName

The name of the service mesh.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

The associated metadata for the service mesh.

Type: [ResourceMetadata](#) object

### spec

The associated specification for the service mesh.

Type: [MeshSpec](#) object

### status

The status of the service mesh.

Type: [MeshStatus](#) object

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ConflictException**

The request contains a client token that was used for a previous update resource call with different specifications. Try the request again with a new client token.

HTTP Status Code: 409

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## **See Also**

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)

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

# UpdateRoute

Updates an existing route for a specified service mesh and virtual router.

## Request Syntax

```
PUT /v20190125/meshes/meshName/virtualRouter/virtualRouterName/routes/routeName?  
meshOwner=meshOwner HTTP/1.1  
Content-type: application/json
```

```
{  
  "clientToken": "string",  
  "spec": {  
    "grpcRoute": {  
      "action": {  
        "weightedTargets": [  
          {  
            "port": number,  
            "virtualNode": "string",  
            "weight": number  
          }  
        ]  
      },  
    },  
    "match": {  
      "metadata": [  
        {  
          "invert": boolean,  
          "match": { ... },  
          "name": "string"  
        }  
      ],  
      "methodName": "string",  
      "port": number,  
      "serviceName": "string"  
    },  
    "retryPolicy": {  
      "grpcRetryEvents": [ "string" ],  
      "httpRetryEvents": [ "string" ],  
      "maxRetries": number,  
      "perRetryTimeout": {  
        "unit": "string",  
        "value": number  
      }  
    },  
  },  
}
```

```
    "tcpRetryEvents": [ "string" ]
  },
  "timeout": {
    "idle": {
      "unit": "string",
      "value": number
    },
    "perRequest": {
      "unit": "string",
      "value": number
    }
  }
},
"http2Route": {
  "action": {
    "weightedTargets": [
      {
        "port": number,
        "virtualNode": "string",
        "weight": number
      }
    ]
  },
  "match": {
    "headers": [
      {
        "invert": boolean,
        "match": { ... },
        "name": "string"
      }
    ],
    "method": "string",
    "path": {
      "exact": "string",
      "regex": "string"
    },
    "port": number,
    "prefix": "string",
    "queryParameters": [
      {
        "match": {
          "exact": "string"
        },
        "name": "string"
      }
    ]
  }
}
```

```

    }
  ],
  "scheme": "string"
},
"retryPolicy": {
  "httpRetryEvents": [ "string" ],
  "maxRetries": number,
  "perRetryTimeout": {
    "unit": "string",
    "value": number
  },
  "tcpRetryEvents": [ "string" ]
},
"timeout": {
  "idle": {
    "unit": "string",
    "value": number
  },
  "perRequest": {
    "unit": "string",
    "value": number
  }
}
},
"httpRoute": {
  "action": {
    "weightedTargets": [
      {
        "port": number,
        "virtualNode": "string",
        "weight": number
      }
    ]
  },
  "match": {
    "headers": [
      {
        "invert": boolean,
        "match": { ... },
        "name": "string"
      }
    ],
    "method": "string",
    "path": {

```

```
    "exact": "string",
    "regex": "string"
  },
  "port": number,
  "prefix": "string",
  "queryParameters": [
    {
      "match": {
        "exact": "string"
      },
      "name": "string"
    }
  ],
  "scheme": "string"
},
"retryPolicy": {
  "httpRetryEvents": [ "string" ],
  "maxRetries": number,
  "perRetryTimeout": {
    "unit": "string",
    "value": number
  },
  "tcpRetryEvents": [ "string" ]
},
"timeout": {
  "idle": {
    "unit": "string",
    "value": number
  },
  "perRequest": {
    "unit": "string",
    "value": number
  }
}
},
"priority": number,
"tcpRoute": {
  "action": {
    "weightedTargets": [
      {
        "port": number,
        "virtualNode": "string",
        "weight": number
      }
    ]
  }
}
```

```
    ]
  },
  "match": {
    "port": number
  },
  "timeout": {
    "idle": {
      "unit": "string",
      "value": number
    }
  }
}
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh that the route resides in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### routeName

The name of the route to update.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### virtualRouterName

The name of the virtual router that the route is associated with.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request accepts the following data in JSON format.

### clientToken

Unique, case-sensitive identifier that you provide to ensure the idempotency of the request. Up to 36 letters, numbers, hyphens, and underscores are allowed.

Type: String

Required: No

### spec

The new route specification to apply. This overwrites the existing data.

Type: [RouteSpec](#) object

Required: Yes

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "routeName": "string",
  "spec": {
```

```

    "grpcRoute": {
      "action": {
        "weightedTargets": [
          {
            "port": number,
            "virtualNode": "string",
            "weight": number
          }
        ]
      },
      "match": {
        "metadata": [
          {
            "invert": boolean,
            "match": { ... },
            "name": "string"
          }
        ],
        "methodName": "string",
        "port": number,
        "serviceName": "string"
      },
      "retryPolicy": {
        "grpcRetryEvents": [ "string" ],
        "httpRetryEvents": [ "string" ],
        "maxRetries": number,
        "perRetryTimeout": {
          "unit": "string",
          "value": number
        },
        "tcpRetryEvents": [ "string" ]
      },
      "timeout": {
        "idle": {
          "unit": "string",
          "value": number
        },
        "perRequest": {
          "unit": "string",
          "value": number
        }
      }
    },
    "http2Route": {

```

```
"action": {
  "weightedTargets": [
    {
      "port": number,
      "virtualNode": "string",
      "weight": number
    }
  ]
},
"match": {
  "headers": [
    {
      "invert": boolean,
      "match": { ... },
      "name": "string"
    }
  ],
  "method": "string",
  "path": {
    "exact": "string",
    "regex": "string"
  },
  "port": number,
  "prefix": "string",
  "queryParameters": [
    {
      "match": {
        "exact": "string"
      },
      "name": "string"
    }
  ],
  "scheme": "string"
},
"retryPolicy": {
  "httpRetryEvents": [ "string" ],
  "maxRetries": number,
  "perRetryTimeout": {
    "unit": "string",
    "value": number
  },
  "tcpRetryEvents": [ "string" ]
},
"timeout": {
```

```
    "idle": {
      "unit": "string",
      "value": number
    },
    "perRequest": {
      "unit": "string",
      "value": number
    }
  }
},
"httpRoute": {
  "action": {
    "weightedTargets": [
      {
        "port": number,
        "virtualNode": "string",
        "weight": number
      }
    ]
  },
  "match": {
    "headers": [
      {
        "invert": boolean,
        "match": { ... },
        "name": "string"
      }
    ],
    "method": "string",
    "path": {
      "exact": "string",
      "regex": "string"
    },
    "port": number,
    "prefix": "string",
    "queryParameters": [
      {
        "match": {
          "exact": "string"
        },
        "name": "string"
      }
    ],
    "scheme": "string"
  }
}
```

```
    },
    "retryPolicy": {
      "httpRetryEvents": [ "string" ],
      "maxRetries": number,
      "perRetryTimeout": {
        "unit": "string",
        "value": number
      },
      "tcpRetryEvents": [ "string" ]
    },
    "timeout": {
      "idle": {
        "unit": "string",
        "value": number
      },
      "perRequest": {
        "unit": "string",
        "value": number
      }
    }
  },
  "priority": number,
  "tcpRoute": {
    "action": {
      "weightedTargets": [
        {
          "port": number,
          "virtualNode": "string",
          "weight": number
        }
      ]
    },
    "match": {
      "port": number
    },
    "timeout": {
      "idle": {
        "unit": "string",
        "value": number
      }
    }
  }
},
"status": {
```

```
    "status": "string"  
  },  
  "virtualRouterName": "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.

### meshName

The name of the service mesh that the route resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

The associated metadata for the route.

Type: [ResourceMetadata](#) object

### routeName

The name of the route.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### spec

The specifications of the route.

Type: [RouteSpec](#) object

### status

The status of the route.

Type: [RouteStatus](#) object

## virtualRouterName

The virtual router that the route is associated with.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### BadRequestException

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### ConflictException

The request contains a client token that was used for a previous update resource call with different specifications. Try the request again with a new client token.

HTTP Status Code: 409

### ForbiddenException

You don't have permissions to perform this action.

HTTP Status Code: 403

### InternalServerErrorException

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

HTTP Status Code: 500

### LimitExceededException

You have exceeded a service limit for your account. For more information, see [Service Limits](#) in the *AWS App Mesh User Guide*.

HTTP Status Code: 400

### NotFoundException

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## **Examples**

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### **Example**

The following example updates a route for the virtual router named `colorteller-vr` in the `ecs-mesh` service mesh. The route directs traffic to two weighted targets: 80% to `colorteller-blue-vn` and 20% to `colorteller-red-vn`.

### **Sample Request**

```
PUT /v20190125/meshes/ecs-mesh/virtualRouter/colorteller-vr/routes/colorteller-route
  HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190228T001532Z
Authorization: AUTHPARAMS
```

```
{
  "spec": {
    "httpRoute": {
      "action": {
        "weightedTargets": [
          {
            "virtualNode": "colorteller-blue-vn",
            "weight": 8
          },
          {
            "virtualNode": "colorteller-red-vn",
            "weight": 2
          }
        ]
      },
      "match": {
        "prefix": "/"
      }
    }
  },
  "clientToken": "e8bbfdff-5d3a-4e5c-9c32-571bad83b021"
}
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 3375b176-77ba-4f76-be44-98ec4a839b8c
content-type: application/json
content-length: 557
date: Thu, 28 Feb 2019 00:15:32 GMT
x-envoy-upstream-service-time: 56
server: envoy
Connection: keep-alive

{
  "meshName": "ecs-mesh",
  "metadata": {
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualRouter/colorteller-vr/route/colorteller-route",
    "createdAt": 1.551311817276E9,
    "lastUpdatedAt": 1.551312932979E9,
    "uid": "1cf3109a-0d7f-438c-b17d-a3785f14ff7b",
  }
}
```

```
"version": 2
},
"routeName": "colorteller-route",
"spec": {
  "httpRoute": {
    "action": {
      "weightedTargets": [{
        "virtualNode": "colorteller-blue-vn",
        "weight": 8
      }, {
        "virtualNode": "colorteller-red-vn",
        "weight": 2
      }]
    },
    "match": {
      "prefix": "/"
    }
  },
  "tcpRoute": null
},
"status": {
  "status": "ACTIVE"
},
"virtualRouterName": "colorteller-vr"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)

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

# UpdateVirtualGateway

Updates an existing virtual gateway in a specified service mesh.

## Request Syntax

```
PUT /v20190125/meshes/meshName/virtualGateways/virtualGatewayName?meshOwner=meshOwner
HTTP/1.1
```

Content-type: application/json

```
{
  "clientToken": "string",
  "spec": {
    "backendDefaults": {
      "clientPolicy": {
        "tls": {
          "certificate": { ... },
          "enforce": boolean,
          "ports": [ number ],
          "validation": {
            "subjectAlternativeNames": {
              "match": {
                "exact": [ "string" ]
              }
            },
            "trust": { ... }
          }
        }
      }
    },
    "listeners": [
      {
        "connectionPool": { ... },
        "healthCheck": {
          "healthyThreshold": number,
          "intervalMillis": number,
          "path": "string",
          "port": number,
          "protocol": "string",
          "timeoutMillis": number,
          "unhealthyThreshold": number
        },
        "portMapping": {
```

```

    "port": number,
    "protocol": "string"
  },
  "tls": {
    "certificate": { ... },
    "mode": "string",
    "validation": {
      "subjectAlternativeNames": {
        "match": {
          "exact": [ "string" ]
        }
      },
      "trust": { ... }
    }
  }
},
"logging": {
  "accessLog": { ... }
}
}
}

```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh that the virtual gateway resides in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

## virtualGatewayName

The name of the virtual gateway to update.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request accepts the following data in JSON format.

### clientToken

Unique, case-sensitive identifier that you provide to ensure the idempotency of the request. Up to 36 letters, numbers, hyphens, and underscores are allowed.

Type: String

Required: No

### spec

The new virtual gateway specification to apply. This overwrites the existing data.

Type: [VirtualGatewaySpec](#) object

Required: Yes

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
```

```

    "version": number
  },
  "spec": {
    "backendDefaults": {
      "clientPolicy": {
        "tls": {
          "certificate": { ... },
          "enforce": boolean,
          "ports": [ number ],
          "validation": {
            "subjectAlternativeNames": {
              "match": {
                "exact": [ "string" ]
              }
            },
            "trust": { ... }
          }
        }
      }
    }
  },
  "listeners": [
    {
      "connectionPool": { ... },
      "healthCheck": {
        "healthyThreshold": number,
        "intervalMillis": number,
        "path": "string",
        "port": number,
        "protocol": "string",
        "timeoutMillis": number,
        "unhealthyThreshold": number
      },
      "portMapping": {
        "port": number,
        "protocol": "string"
      },
      "tls": {
        "certificate": { ... },
        "mode": "string",
        "validation": {
          "subjectAlternativeNames": {
            "match": {
              "exact": [ "string" ]
            }
          }
        }
      }
    }
  ]
}

```

```
        },
        "trust": { ... }
    }
}
],
"logging": {
    "accessLog": { ... }
}
},
"status": {
    "status": "string"
},
"virtualGatewayName": "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.

### meshName

The name of the service mesh that the virtual gateway resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

An object that represents metadata for a resource.

Type: [ResourceMetadata](#) object

### spec

The specifications of the virtual gateway.

Type: [VirtualGatewaySpec](#) object

### status

The current status of the virtual gateway.

Type: [VirtualGatewayStatus](#) object

### [virtualGatewayName](#)

The name of the virtual gateway.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ConflictException**

The request contains a client token that was used for a previous update resource call with different specifications. Try the request again with a new client token.

HTTP Status Code: 409

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **LimitExceededException**

You have exceeded a service limit for your account. For more information, see [Service Limits](#) in the *AWS App Mesh User Guide*.

HTTP Status Code: 400

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## **Examples**

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### **Example**

The following example updates a virtual gateway named myVirtualGateway in the apps service mesh.

### **Sample Request**

```
PUT /v20190125/meshes/apps/virtualGateways/myVirtualGateway HTTP/1.1
Host: appmesh.us-west-2.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.298 Python/3.6.0 Windows/10 botocore/1.13.34
X-Amz-Date: 20200706T151140Z
Authorization: AUTHPARAMS

{
```

```
"spec": {
  "listeners": [{
    "portMapping": {
      "port": 8080,
      "protocol": "http2"
    }
  }]
},
"clientToken": "1aa1111a-1111-1111-111a-a1a1a1aa111a"
}
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 009bf6b9-c86a-40ba-b6a4-45cb41badc77
content-type: application/json
content-length: 515
date: Mon, 06 Jul 2020 15:11:40 GMT
x-envoy-upstream-service-time: 34
server: envoy
Connection: keep-alive

{
  "meshName": "apps",
  "metadata": {
    "arn": "arn:aws:appmesh:us-west-2:123456789012:mesh/apps/virtualGateway/myVirtualGateway",
    "createdAt": 1.594048239495E9,
    "lastUpdatedAt": 1.594048300694E9,
    "meshOwner": "123456789012",
    "resourceOwner": "123456789012",
    "uid": "ff111ff2-33ff-4ff4-ffff-111f111f111f",
    "version": 2
  },
  "spec": {
    "backendDefaults": null,
    "listeners": [{
      "healthCheck": null,
      "portMapping": {
        "port": 8080,
        "protocol": "http2"
      }
    }],
    "tls": null
  }
}
```

```
  ]],  
  "logging": null  
},  
"status": {  
  "status": "ACTIVE"  
},  
"virtualGatewayName": "myVirtualGateway"  
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# UpdateVirtualNode

Updates an existing virtual node in a specified service mesh.

## Request Syntax

```
PUT /v20190125/meshes/meshName/virtualNodes/virtualNodeName?meshOwner=meshOwner
```

HTTP/1.1

Content-type: application/json

```
{
  "clientToken": "string",
  "spec": {
    "backendDefaults": {
      "clientPolicy": {
        "tls": {
          "certificate": { ... },
          "enforce": boolean,
          "ports": [ number ],
          "validation": {
            "subjectAlternativeNames": {
              "match": {
                "exact": [ "string" ]
              }
            },
            "trust": { ... }
          }
        }
      }
    },
    "backends": [
      { ... }
    ],
    "listeners": [
      {
        "connectionPool": { ... },
        "healthCheck": {
          "healthyThreshold": number,
          "intervalMillis": number,
          "path": "string",
          "port": number,
          "protocol": "string",
          "timeoutMillis": number,

```

```

    "unhealthyThreshold": number
  },
  "outlierDetection": {
    "baseEjectionDuration": {
      "unit": "string",
      "value": number
    },
    "interval": {
      "unit": "string",
      "value": number
    },
    "maxEjectionPercent": number,
    "maxServerErrors": number
  },
  "portMapping": {
    "port": number,
    "protocol": "string"
  },
  "timeout": { ... },
  "tls": {
    "certificate": { ... },
    "mode": "string",
    "validation": {
      "subjectAlternativeNames": {
        "match": {
          "exact": [ "string" ]
        }
      },
      "trust": { ... }
    }
  }
}
],
"logging": {
  "accessLog": { ... }
},
"serviceDiscovery": { ... }
}
}

```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh that the virtual node resides in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### virtualNodeName

The name of the virtual node to update.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request accepts the following data in JSON format.

### clientToken

Unique, case-sensitive identifier that you provide to ensure the idempotency of the request. Up to 36 letters, numbers, hyphens, and underscores are allowed.

Type: String

Required: No

### spec

The new virtual node specification to apply. This overwrites the existing data.

Type: [VirtualNodeSpec](#) object

Required: Yes

## Response Syntax

HTTP/1.1 200

Content-type: application/json

```
{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "backendDefaults": {
      "clientPolicy": {
        "tls": {
          "certificate": { ... },
          "enforce": boolean,
          "ports": [ number ],
          "validation": {
            "subjectAlternativeNames": {
              "match": {
                "exact": [ "string" ]
              }
            },
            "trust": { ... }
          }
        }
      }
    },
    "backends": [
      { ... }
    ],
    "listeners": [
      {
        "connectionPool": { ... },
        "healthCheck": {
          "healthyThreshold": number,
          "intervalMillis": number,
```

```

    "path": "string",
    "port": number,
    "protocol": "string",
    "timeoutMillis": number,
    "unhealthyThreshold": number
  },
  "outlierDetection": {
    "baseEjectionDuration": {
      "unit": "string",
      "value": number
    },
    "interval": {
      "unit": "string",
      "value": number
    },
    "maxEjectionPercent": number,
    "maxServerErrors": number
  },
  "portMapping": {
    "port": number,
    "protocol": "string"
  },
  "timeout": { ... },
  "tls": {
    "certificate": { ... },
    "mode": "string",
    "validation": {
      "subjectAlternativeNames": {
        "match": {
          "exact": [ "string" ]
        }
      },
      "trust": { ... }
    }
  }
},
],
"logging": {
  "accessLog": { ... }
},
"serviceDiscovery": { ... }
},
"status": {
  "status": "string"
}

```

```
  },  
  "virtualNodeName": "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.

### meshName

The name of the service mesh that the virtual node resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

The associated metadata for the virtual node.

Type: [ResourceMetadata](#) object

### spec

The specifications of the virtual node.

Type: [VirtualNodeSpec](#) object

### status

The current status for the virtual node.

Type: [VirtualNodeStatus](#) object

### virtualNodeName

The name of the virtual node.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ConflictException**

The request contains a client token that was used for a previous update resource call with different specifications. Try the request again with a new client token.

HTTP Status Code: 409

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **LimitExceededException**

You have exceeded a service limit for your account. For more information, see [Service Limits](#) in the *AWS App Mesh User Guide*.

HTTP Status Code: 400

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

## Example

The following example updates a virtual node named `colorteller-vn` in the `ecs-mesh` service mesh.

## Sample Request

```
PUT /v20190125/meshes/ecs-mesh/virtualNodes/colorteller-vn HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190228T002242Z
Authorization: AUTHPARAMS

{
  "spec": {
    "backends": [],
    "listeners": [
      {
        "healthCheck": {
          "healthyThreshold": 2,
          "intervalMillis": 5000,
```

```
    "path": "/ping",
    "port": 9080,
    "protocol": "http",
    "timeoutMillis": 2000,
    "unhealthyThreshold": 2
  },
  "portMapping": {
    "port": 9080,
    "protocol": "http"
  }
},
"serviceDiscovery": {
  "dns": {
    "hostname": "colorteller-white.default.svc.cluster.local"
  }
},
"clientToken": "xxxxxxxx-xxxx-xxxx-1111-EXAMPLE11111"
}
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: d6dc2922-f4c0-4077-8107-db673f167a12
content-type: application/json
content-length: 660
date: Thu, 28 Feb 2019 00:22:42 GMT
x-envoy-upstream-service-time: 88
server: envoy
Connection: keep-alive

{
  "meshName": "ecs-mesh",
  "metadata": {
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualNode/colorteller-
vn",
    "createdAt": 1.551311799179E9,
    "lastUpdatedAt": 1.551313363236E9,
    "uid": "0999e53a-8e0e-4c4c-8764-ae8ebecc296d",
    "version": 2
  },
  "spec": {
```

```
"backends": [],
"listeners": [{
  "healthCheck": {
    "healthyThreshold": 2,
    "intervalMillis": 5000,
    "path": "/ping",
    "port": 9080,
    "protocol": "http",
    "timeoutMillis": 2000,
    "unhealthyThreshold": 2
  },
  "portMapping": {
    "port": 9080,
    "protocol": "http"
  }
}],
"logging": null,
"serviceDiscovery": {
  "dns": {
    "hostname": "colorteller-white.default.svc.cluster.local"
  }
},
"status": {
  "status": "ACTIVE"
},
"virtualNodeName": "colorteller-vn"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)

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

# UpdateVirtualRouter

Updates an existing virtual router in a specified service mesh.

## Request Syntax

```
PUT /v20190125/meshes/meshName/virtualRouters/virtualRouterName?meshOwner=meshOwner
HTTP/1.1
Content-type: application/json

{
  "clientToken": "string",
  "spec": {
    "listeners": [
      {
        "portMapping": {
          "port": number,
          "protocol": "string"
        }
      }
    ]
  }
}
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh that the virtual router resides in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

## virtualRouterName

The name of the virtual router to update.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

## Request Body

The request accepts the following data in JSON format.

### clientToken

Unique, case-sensitive identifier that you provide to ensure the idempotency of the request. Up to 36 letters, numbers, hyphens, and underscores are allowed.

Type: String

Required: No

### spec

The new virtual router specification to apply. This overwrites the existing data.

Type: [VirtualRouterSpec](#) object

Required: Yes

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
```

```
    "uid": "string",
    "version": number
  },
  "spec": {
    "listeners": [
      {
        "portMapping": {
          "port": number,
          "protocol": "string"
        }
      }
    ]
  },
  "status": {
    "status": "string"
  },
  "virtualRouterName": "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.

### meshName

The name of the service mesh that the virtual router resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### metadata

The associated metadata for the virtual router.

Type: [ResourceMetadata](#) object

### spec

The specifications of the virtual router.

Type: [VirtualRouterSpec](#) object

## status

The current status of the virtual router.

Type: [VirtualRouterStatus](#) object

## virtualRouterName

The name of the virtual router.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ConflictException**

The request contains a client token that was used for a previous update resource call with different specifications. Try the request again with a new client token.

HTTP Status Code: 409

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **LimitExceededException**

You have exceeded a service limit for your account. For more information, see [Service Limits](#) in the *AWS App Mesh User Guide*.

HTTP Status Code: 400

## NotFoundException

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

## ServiceUnavailableException

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

HTTP Status Code: 503

## TooManyRequestsException

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

## Example

The following example updates a virtual router named `colorteller-vr` in the `ecs-mesh` service mesh.

## Sample Request

```
PUT /v20190125/meshes/ecs-mesh/virtualRouters/colorteller-vr HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190228T002506Z
```

Authorization: AUTHPARAMS

```
{
  "spec": {},
  "clientToken": "6029e873-0ef2-41b9-957e-2742048c92c3"
}
```

## Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: fba0cb49-d242-46e5-aacd-c4a0fad60413
content-type: application/json
content-length: 336
date: Thu, 28 Feb 2019 00:25:06 GMT
x-envoy-upstream-service-time: 70
server: envoy
Connection: keep-alive

{
  "meshName": "ecs-mesh",
  "metadata": {
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualRouter/
colorteller-vr",
    "createdAt": 1551311805.476,
    "lastUpdatedAt": 1551313507.39,
    "uid": "79628d34-8c17-42ba-83a5-8a42fd17ec5a",
    "version": 2
  },
  "spec": {
    "listeners": [
      {
        "portMapping": {
          "port": 9080,
          "protocol": "http"
        }
      }
    ]
  },
  "status": {
    "status": "ACTIVE"
  },
  "virtualRouterName": "colorteller-vr"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# UpdateVirtualService

Updates an existing virtual service in a specified service mesh.

## Request Syntax

```
PUT /v20190125/meshes/meshName/virtualServices/virtualServiceName?meshOwner=meshOwner
HTTP/1.1
Content-type: application/json

{
  "clientToken": "string",
  "spec": {
    "provider": { ... }
  }
}
```

## URI Request Parameters

The request uses the following URI parameters.

### meshName

The name of the service mesh that the virtual service resides in.

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Length Constraints: Fixed length of 12.

### virtualServiceName

The name of the virtual service to update.

Required: Yes

## Request Body

The request accepts the following data in JSON format.

### clientToken

Unique, case-sensitive identifier that you provide to ensure the idempotency of the request. Up to 36 letters, numbers, hyphens, and underscores are allowed.

Type: String

Required: No

### spec

The new virtual service specification to apply. This overwrites the existing data.

Type: [VirtualServiceSpec](#) object

Required: Yes

## Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "meshName": "string",
  "metadata": {
    "arn": "string",
    "createdAt": number,
    "lastUpdatedAt": number,
    "meshOwner": "string",
    "resourceOwner": "string",
    "uid": "string",
    "version": number
  },
  "spec": {
    "provider": { ... }
  },
  "status": {
    "status": "string"
  },
}
```

```
"virtualServiceName": "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.

### [meshName](#)

The name of the service mesh that the virtual service resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

### [metadata](#)

An object that represents metadata for a resource.

Type: [ResourceMetadata](#) object

### [spec](#)

The specifications of the virtual service.

Type: [VirtualServiceSpec](#) object

### [status](#)

The current status of the virtual service.

Type: [VirtualServiceStatus](#) object

### [virtualServiceName](#)

The name of the virtual service.

Type: String

## Errors

### **BadRequestException**

The request syntax was malformed. Check your request syntax and try again.

HTTP Status Code: 400

### **ConflictException**

The request contains a client token that was used for a previous update resource call with different specifications. Try the request again with a new client token.

HTTP Status Code: 409

### **ForbiddenException**

You don't have permissions to perform this action.

HTTP Status Code: 403

### **InternalServerErrorException**

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

HTTP Status Code: 500

### **LimitExceededException**

You have exceeded a service limit for your account. For more information, see [Service Limits](#) in the *AWS App Mesh User Guide*.

HTTP Status Code: 400

### **NotFoundException**

The specified resource doesn't exist. Check your request syntax and try again.

HTTP Status Code: 404

### **ServiceUnavailableException**

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

HTTP Status Code: 503

### **TooManyRequestsException**

The maximum request rate permitted by the App Mesh APIs has been exceeded for your account. For best results, use an increasing or variable sleep interval between requests.

HTTP Status Code: 429

## Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You need to learn how to sign HTTP requests only if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

### Example

This example updates a virtual service named `colorgateway.default.svc.cluster.local` in the `ecs-mesh` service mesh.

### Sample Request

```
PUT /v20190125/meshes/ecs-mesh/virtualServices/colorgateway.default.svc.cluster.local
HTTP/1.1
Host: appmesh.us-east-1.amazonaws.com
Accept-Encoding: identity
User-Agent: aws-cli/1.16.56 Python/3.7.0 Darwin/17.7.0 botocore/1.12.46
X-Amz-Date: 20190228T002829Z
Authorization: AUTHPARAMS

{
  "spec": {
    "provider": {
      "virtualNode": {
        "virtualNodeName": "colorgateway-vn"
      }
    }
  },
  "clientToken": "c207a9a1-5828-4d73-9e8e-1d3b9350b2ac"
}
```

### Sample Response

```
HTTP/1.1 200 OK
x-amzn-requestid: 60d5fb65-ac74-4523-a1df-9e56da84fa84
```

```
content-type: application/json
content-length: 456
date: Thu, 28 Feb 2019 00:28:29 GMT
x-envoy-upstream-service-time: 88
server: envoy
Connection: keep-alive

{
  "meshName": "ecs-mesh",
  "metadata": {
    "arn": "arn:aws:appmesh:us-east-1:123456789012:mesh/ecs-mesh/virtualService/
colorgateway.default.svc.cluster.local",
    "createdAt": 1.551311807444E9,
    "lastUpdatedAt": 1.551313709898E9,
    "uid": "dd06064b-e542-40a9-bbc7-e381a47ea0e0",
    "version": 2
  },
  "spec": {
    "provider": {
      "virtualNode": {
        "virtualNodeName": "colorgateway-vn"
      },
      "virtualRouter": null
    }
  },
  "status": {
    "status": "ACTIVE"
  },
  "virtualServiceName": "colorgateway.default.svc.cluster.local"
}
```

## See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)

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

# Data Types

The AWS App Mesh 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:

- [AccessLog](#)
- [AwsCloudMapInstanceAttribute](#)
- [AwsCloudMapServiceDiscovery](#)
- [Backend](#)
- [BackendDefaults](#)
- [ClientPolicy](#)
- [ClientPolicyTls](#)
- [ClientTlsCertificate](#)
- [DnsServiceDiscovery](#)
- [Duration](#)
- [EgressFilter](#)
- [FileAccessLog](#)
- [GatewayRouteData](#)
- [GatewayRouteHostnameMatch](#)
- [GatewayRouteHostnameRewrite](#)
- [GatewayRouteRef](#)
- [GatewayRouteSpec](#)
- [GatewayRouteStatus](#)
- [GatewayRouteTarget](#)
- [GatewayRouteVirtualService](#)

- [GrpcGatewayRoute](#)
- [GrpcGatewayRouteAction](#)
- [GrpcGatewayRouteMatch](#)
- [GrpcGatewayRouteMetadata](#)
- [GrpcGatewayRouteRewrite](#)
- [GrpcMetadataMatchMethod](#)
- [GrpcRetryPolicy](#)
- [GrpcRoute](#)
- [GrpcRouteAction](#)
- [GrpcRouteMatch](#)
- [GrpcRouteMetadata](#)
- [GrpcRouteMetadataMatchMethod](#)
- [GrpcTimeout](#)
- [HeaderMatchMethod](#)
- [HealthCheckPolicy](#)
- [HttpGatewayRoute](#)
- [HttpGatewayRouteAction](#)
- [HttpGatewayRouteHeader](#)
- [HttpGatewayRouteMatch](#)
- [HttpGatewayRoutePathRewrite](#)
- [HttpGatewayRoutePrefixRewrite](#)
- [HttpGatewayRouteRewrite](#)
- [HttpPathMatch](#)
- [HttpQueryParameter](#)
- [HttpRetryPolicy](#)
- [HttpRoute](#)
- [HttpRouteAction](#)
- [HttpRouteHeader](#)
- [HttpRouteMatch](#)
- [HttpTimeout](#)

- [JsonFormatRef](#)
- [Listener](#)
- [ListenerTimeout](#)
- [ListenerTls](#)
- [ListenerTlsAcmCertificate](#)
- [ListenerTlsCertificate](#)
- [ListenerTlsFileCertificate](#)
- [ListenerTlsSdsCertificate](#)
- [ListenerTlsValidationContext](#)
- [ListenerTlsValidationContextTrust](#)
- [Logging](#)
- [LoggingFormat](#)
- [MatchRange](#)
- [MeshData](#)
- [MeshRef](#)
- [MeshServiceDiscovery](#)
- [MeshSpec](#)
- [MeshStatus](#)
- [OutlierDetection](#)
- [PortMapping](#)
- [QueryParameterMatch](#)
- [ResourceMetadata](#)
- [RouteData](#)
- [RouteRef](#)
- [RouteSpec](#)
- [RouteStatus](#)
- [ServiceDiscovery](#)
- [SubjectAlternativeNameMatchers](#)
- [SubjectAlternativeNames](#)
- [TagRef](#)

- [TcpRoute](#)
- [TcpRouteAction](#)
- [TcpRouteMatch](#)
- [TcpTimeout](#)
- [TlsValidationContext](#)
- [TlsValidationContextAcmTrust](#)
- [TlsValidationContextFileTrust](#)
- [TlsValidationContextSdsTrust](#)
- [TlsValidationContextTrust](#)
- [VirtualGatewayAccessLog](#)
- [VirtualGatewayBackendDefaults](#)
- [VirtualGatewayClientPolicy](#)
- [VirtualGatewayClientPolicyTls](#)
- [VirtualGatewayClientTlsCertificate](#)
- [VirtualGatewayConnectionPool](#)
- [VirtualGatewayData](#)
- [VirtualGatewayFileAccessLog](#)
- [VirtualGatewayGrpcConnectionPool](#)
- [VirtualGatewayHealthCheckPolicy](#)
- [VirtualGatewayHttp2ConnectionPool](#)
- [VirtualGatewayHttpConnectionPool](#)
- [VirtualGatewayListener](#)
- [VirtualGatewayListenerTls](#)
- [VirtualGatewayListenerTlsAcmCertificate](#)
- [VirtualGatewayListenerTlsCertificate](#)
- [VirtualGatewayListenerTlsFileCertificate](#)
- [VirtualGatewayListenerTlsSdsCertificate](#)
- [VirtualGatewayListenerTlsValidationContext](#)
- [VirtualGatewayListenerTlsValidationContextTrust](#)
- [VirtualGatewayLogging](#)

- [VirtualGatewayPortMapping](#)
- [VirtualGatewayRef](#)
- [VirtualGatewaySpec](#)
- [VirtualGatewayStatus](#)
- [VirtualGatewayTlsValidationContext](#)
- [VirtualGatewayTlsValidationContextAcmTrust](#)
- [VirtualGatewayTlsValidationContextFileTrust](#)
- [VirtualGatewayTlsValidationContextSdsTrust](#)
- [VirtualGatewayTlsValidationContextTrust](#)
- [VirtualNodeConnectionPool](#)
- [VirtualNodeData](#)
- [VirtualNodeGrpcConnectionPool](#)
- [VirtualNodeHttp2ConnectionPool](#)
- [VirtualNodeHttpConnectionPool](#)
- [VirtualNodeRef](#)
- [VirtualNodeServiceProvider](#)
- [VirtualNodeSpec](#)
- [VirtualNodeStatus](#)
- [VirtualNodeTcpConnectionPool](#)
- [VirtualRouterData](#)
- [VirtualRouterListener](#)
- [VirtualRouterRef](#)
- [VirtualRouterServiceProvider](#)
- [VirtualRouterSpec](#)
- [VirtualRouterStatus](#)
- [VirtualServiceBackend](#)
- [VirtualServiceData](#)
- [VirtualServiceProvider](#)
- [VirtualServiceRef](#)
- [VirtualServiceSpec](#)

- [VirtualServiceStatus](#)
- [WeightedTarget](#)

# AccessLog

An object that represents the access logging information for a virtual node.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### file

The file object to send virtual node access logs to.

Type: [FileAccessLog](#) object

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

# AwsCloudMapInstanceAttribute

An object that represents the AWS Cloud Map attribute information for your virtual node.

## Note

AWS Cloud Map is not available in the eu-south-1 Region.

## Contents

### key

The name of an AWS Cloud Map service instance attribute key. Any AWS Cloud Map service instance that contains the specified key and value is returned.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[a-zA-Z0-9!-~]+`

Required: Yes

### value

The value of an AWS Cloud Map service instance attribute key. Any AWS Cloud Map service instance that contains the specified key and value is returned.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `([a-zA-Z0-9!-~][ a-zA-Z0-9!-~]*){0,1}[a-zA-Z0-9!-~]{0,1}`

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

# AwsCloudMapServiceDiscovery

An object that represents the AWS Cloud Map service discovery information for your virtual node.

## Note

AWS Cloud Map is not available in the eu-south-1 Region.

## Contents

### namespaceName

The HTTP name of the AWS Cloud Map namespace to use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

### serviceName

The name of the AWS Cloud Map service to use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

### attributes

A string map that contains attributes with values that you can use to filter instances by any custom attribute that you specified when you registered the instance. Only instances that match all of the specified key/value pairs will be returned.

Type: Array of [AwsCloudMapInstanceAttribute](#) objects

Required: No

### ipPreference

The preferred IP version that this virtual node uses. Setting the IP preference on the virtual node only overrides the IP preference set for the mesh on this specific node.

Type: String

Valid Values: IPv6\_PREFERRED | IPv4\_PREFERRED | IPv4\_ONLY | IPv6\_ONLY

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

# Backend

An object that represents the backends that a virtual node is expected to send outbound traffic to.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### **virtualService**

Specifies a virtual service to use as a backend.

Type: [VirtualServiceBackend](#) object

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

# BackendDefaults

An object that represents the default properties for a backend.

## Contents

### clientPolicy

A reference to an object that represents a client policy.

Type: [ClientPolicy](#) object

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

# ClientPolicy

An object that represents a client policy.

## Contents

### tls

A reference to an object that represents a Transport Layer Security (TLS) client policy.

Type: [ClientPolicyTls](#) object

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

# ClientPolicyTls

A reference to an object that represents a Transport Layer Security (TLS) client policy.

## Contents

### validation

A reference to an object that represents a TLS validation context.

Type: [TlsValidationContext](#) object

Required: Yes

### certificate

A reference to an object that represents a client's TLS certificate.

Type: [ClientTlsCertificate](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

Required: No

### enforce

Whether the policy is enforced. The default is `True`, if a value isn't specified.

Type: Boolean

Required: No

### ports

One or more ports that the policy is enforced for.

Type: Array of integers

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

# ClientTlsCertificate

An object that represents the client's certificate.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### file

An object that represents a local file certificate. The certificate must meet specific requirements and you must have proxy authorization enabled. For more information, see [Transport Layer Security \(TLS\)](#).

Type: [ListenerTlsFileCertificate](#) object

Required: No

### sds

A reference to an object that represents a client's TLS Secret Discovery Service certificate.

Type: [ListenerTlsSdsCertificate](#) object

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

# DnsServiceDiscovery

An object that represents the DNS service discovery information for your virtual node.

## Contents

### hostname

Specifies the DNS service discovery hostname for the virtual node.

Type: String

Required: Yes

### ipPreference

The preferred IP version that this virtual node uses. Setting the IP preference on the virtual node only overrides the IP preference set for the mesh on this specific node.

Type: String

Valid Values: IPv6\_PREFERRED | IPv4\_PREFERRED | IPv4\_ONLY | IPv6\_ONLY

Required: No

### responseType

Specifies the DNS response type for the virtual node.

Type: String

Valid Values: LOADBALANCER | ENDPOINTS

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 Java V2](#)

- [AWS SDK for Ruby V3](#)

# Duration

An object that represents a duration of time.

## Contents

### unit

A unit of time.

Type: String

Valid Values: s | ms

Required: No

### value

A number of time units.

Type: Long

Valid Range: Minimum value of 0.

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

# EgressFilter

An object that represents the egress filter rules for a service mesh.

## Contents

### type

The egress filter type. By default, the type is `DROP_ALL`, which allows egress only from virtual nodes to other defined resources in the service mesh (and any traffic to `*.amazonaws.com` for AWS API calls). You can set the egress filter type to `ALLOW_ALL` to allow egress to any endpoint inside or outside of the service mesh.

#### Note

If you specify any backends on a virtual node when using `ALLOW_ALL`, you must specify all egress for that virtual node as backends. Otherwise, `ALLOW_ALL` will no longer work for that virtual node.

Type: String

Valid Values: `ALLOW_ALL` | `DROP_ALL`

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

# FileAccessLog

An object that represents an access log file.

## Contents

### path

The file path to write access logs to. You can use `/dev/stdout` to send access logs to standard out and configure your Envoy container to use a log driver, such as `awslogs`, to export the access logs to a log storage service such as Amazon CloudWatch Logs. You can also specify a path in the Envoy container's file system to write the files to disk.

#### Note

The Envoy process must have write permissions to the path that you specify here. Otherwise, Envoy fails to bootstrap properly.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### format

The specified format for the logs. The format is either `json_format` or `text_format`.

Type: [LoggingFormat](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

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

# GatewayRouteData

An object that represents a gateway route returned by a describe operation.

## Contents

### gatewayRouteName

The name of the gateway route.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshName

The name of the service mesh that the resource resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### metadata

An object that represents metadata for a resource.

Type: [ResourceMetadata](#) object

Required: Yes

### spec

The specifications of the gateway route.

Type: [GatewayRouteSpec](#) object

Required: Yes

### status

The status of the gateway route.

Type: [GatewayRouteStatus](#) object

Required: Yes

### **virtualGatewayName**

The virtual gateway that the gateway route is associated with.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# GatewayRouteHostnameMatch

An object representing the gateway route host name to match.

## Contents

### exact

The exact host name to match on.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 253.

Required: No

### suffix

The specified ending characters of the host name to match on.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 253.

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

# GatewayRouteHostnameRewrite

An object representing the gateway route host name to rewrite.

## Contents

### defaultTargetHostname

The default target host name to write to.

Type: String

Valid Values: ENABLED | DISABLED

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

# GatewayRouteRef

An object that represents a gateway route returned by a list operation.

## Contents

### **arn**

The full Amazon Resource Name (ARN) for the gateway route.

Type: String

Required: Yes

### **createdAt**

The Unix epoch timestamp in seconds for when the resource was created.

Type: Timestamp

Required: Yes

### **gatewayRouteName**

The name of the gateway route.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### **lastUpdatedAt**

The Unix epoch timestamp in seconds for when the resource was last updated.

Type: Timestamp

Required: Yes

### **meshName**

The name of the service mesh that the resource resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### **meshOwner**

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Type: String

Length Constraints: Fixed length of 12.

Required: Yes

### **resourceOwner**

The AWS IAM account ID of the resource owner. If the account ID is not your own, then it's the ID of the mesh owner or of another account that the mesh is shared with. For more information about mesh sharing, see [Working with shared meshes](#).

Type: String

Length Constraints: Fixed length of 12.

Required: Yes

### **version**

The version of the resource. Resources are created at version 1, and this version is incremented each time that they're updated.

Type: Long

Required: Yes

### **virtualGatewayName**

The virtual gateway that the gateway route is associated with.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# GatewayRouteSpec

An object that represents a gateway route specification. Specify one gateway route type.

## Contents

### grpcRoute

An object that represents the specification of a gRPC gateway route.

Type: [GrpcGatewayRoute](#) object

Required: No

### http2Route

An object that represents the specification of an HTTP/2 gateway route.

Type: [HttpGatewayRoute](#) object

Required: No

### httpRoute

An object that represents the specification of an HTTP gateway route.

Type: [HttpGatewayRoute](#) object

Required: No

### priority

The ordering of the gateway routes spec.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

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

# GatewayRouteStatus

An object that represents the current status of a gateway route.

## Contents

### status

The current status for the gateway route.

Type: String

Valid Values: ACTIVE | INACTIVE | DELETED

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

# GatewayRouteTarget

An object that represents a gateway route target.

## Contents

### virtualService

An object that represents a virtual service gateway route target.

Type: [GatewayRouteVirtualService](#) object

Required: Yes

### port

The port number of the gateway route target.

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

# GatewayRouteVirtualService

An object that represents the virtual service that traffic is routed to.

## Contents

### **virtualServiceName**

The name of the virtual service that traffic is routed to.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# GrpcGatewayRoute

An object that represents a gRPC gateway route.

## Contents

### action

An object that represents the action to take if a match is determined.

Type: [GrpcGatewayRouteAction](#) object

Required: Yes

### match

An object that represents the criteria for determining a request match.

Type: [GrpcGatewayRouteMatch](#) object

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

# GrpcGatewayRouteAction

An object that represents the action to take if a match is determined.

## Contents

### target

An object that represents the target that traffic is routed to when a request matches the gateway route.

Type: [GatewayRouteTarget](#) object

Required: Yes

### rewrite

The gateway route action to rewrite.

Type: [GrpcGatewayRouteRewrite](#) object

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

# GrpcGatewayRouteMatch

An object that represents the criteria for determining a request match.

## Contents

### hostname

The gateway route host name to be matched on.

Type: [GatewayRouteHostnameMatch](#) object

Required: No

### metadata

The gateway route metadata to be matched on.

Type: Array of [GrpcGatewayRouteMetadata](#) objects

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

Required: No

### port

The gateway route port to be matched on.

Type: Integer

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

Required: No

### serviceName

The fully qualified domain name for the service to match from the request.

Type: String

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

# GrpcGatewayRouteMetadata

An object representing the metadata of the gateway route.

## Contents

### name

A name for the gateway route metadata.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 50.

Required: Yes

### invert

Specify `True` to match anything except the match criteria. The default value is `False`.

Type: Boolean

Required: No

### match

The criteria for determining a metadata match.

Type: [GrpcMetadataMatchMethod](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

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



# GrpcGatewayRouteRewrite

An object that represents the gateway route to rewrite.

## Contents

### hostname

The host name of the gateway route to rewrite.

Type: [GatewayRouteHostnameRewrite](#) object

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

# GrpcMetadataMatchMethod

An object representing the method header to be matched.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### **exact**

The exact method header to be matched on.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

### **prefix**

The specified beginning characters of the method header to be matched on.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

### **range**

An object that represents the range of values to match on. The first character of the range is included in the range, though the last character is not. For example, if the range specified were 1-100, only values 1-99 would be matched.

Type: [MatchRange](#) object

Required: No

**regex**

The regex used to match the method header.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

**suffix**

The specified ending characters of the method header to match on.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# GrpcRetryPolicy

An object that represents a retry policy. Specify at least one value for at least one of the types of `RetryEvents`, a value for `maxRetries`, and a value for `perRetryTimeout`. Both `server-error` and `gateway-error` under `httpRetryEvents` include the Envoy `reset` policy. For more information on the `reset` policy, see the [Envoy documentation](#).

## Contents

### **maxRetries**

The maximum number of retry attempts.

Type: Long

Valid Range: Minimum value of 0.

Required: Yes

### **perRetryTimeout**

The timeout for each retry attempt.

Type: [Duration](#) object

Required: Yes

### **grpcRetryEvents**

Specify at least one of the valid values.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 5 items.

Valid Values: `cancelled` | `deadline-exceeded` | `internal` | `resource-exhausted` | `unavailable`

Required: No

### **httpRetryEvents**

Specify at least one of the following values.

- **server-error** – HTTP status codes 500, 501, 502, 503, 504, 505, 506, 507, 508, 510, and 511

- **gateway-error** – HTTP status codes 502, 503, and 504
- **client-error** – HTTP status code 409
- **stream-error** – Retry on refused stream

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 25 items.

Length Constraints: Minimum length of 1. Maximum length of 25.

Required: No

### **tcpRetryEvents**

Specify a valid value. The event occurs before any processing of a request has started and is encountered when the upstream is temporarily or permanently unavailable.

Type: Array of strings

Array Members: Fixed number of 1 item.

Valid Values: `connection-error`

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

# GrpcRoute

An object that represents a gRPC route type.

## Contents

### action

An object that represents the action to take if a match is determined.

Type: [GrpcRouteAction](#) object

Required: Yes

### match

An object that represents the criteria for determining a request match.

Type: [GrpcRouteMatch](#) object

Required: Yes

### retryPolicy

An object that represents a retry policy.

Type: [GrpcRetryPolicy](#) object

Required: No

### timeout

An object that represents types of timeouts.

Type: [GrpcTimeout](#) object

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

# GrpcRouteAction

An object that represents the action to take if a match is determined.

## Contents

### weightedTargets

An object that represents the targets that traffic is routed to when a request matches the route.

Type: Array of [WeightedTarget](#) objects

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

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

# GrpcRouteMatch

An object that represents the criteria for determining a request match.

## Contents

### metadata

An object that represents the data to match from the request.

Type: Array of [GrpcRouteMetadata](#) objects

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

Required: No

### methodName

The method name to match from the request. If you specify a name, you must also specify a `serviceName`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 50.

Required: No

### port

The port number to match on.

Type: Integer

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

Required: No

### serviceName

The fully qualified domain name for the service to match from the request.

Type: String

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

# GrpcRouteMetadata

An object that represents the match metadata for the route.

## Contents

### name

The name of the route.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 50.

Required: Yes

### invert

Specify `True` to match anything except the match criteria. The default value is `False`.

Type: Boolean

Required: No

### match

An object that represents the data to match from the request.

Type: [GrpcRouteMetadataMatchMethod](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

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



# GrpcRouteMetadataMatchMethod

An object that represents the match method. Specify one of the match values.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### **exact**

The value sent by the client must match the specified value exactly.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

### **prefix**

The value sent by the client must begin with the specified characters.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

### **range**

An object that represents the range of values to match on.

Type: [MatchRange](#) object

Required: No

### **regex**

The value sent by the client must include the specified characters.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

### **suffix**

The value sent by the client must end with the specified characters.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# GrpcTimeout

An object that represents types of timeouts.

## Contents

### idle

An object that represents an idle timeout. An idle timeout bounds the amount of time that a connection may be idle. The default value is none.

Type: [Duration](#) object

Required: No

### perRequest

An object that represents a per request timeout. The default value is 15 seconds. If you set a higher timeout, then make sure that the higher value is set for each App Mesh resource in a conversation. For example, if a virtual node backend uses a virtual router provider to route to another virtual node, then the timeout should be greater than 15 seconds for the source and destination virtual node and the route.

Type: [Duration](#) object

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

# HeaderMatchMethod

An object that represents the method and value to match with the header value sent in a request. Specify one match method.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### **exact**

The value sent by the client must match the specified value exactly.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

### **prefix**

The value sent by the client must begin with the specified characters.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

### **range**

An object that represents the range of values to match on.

Type: [MatchRange](#) object

Required: No

### **regex**

The value sent by the client must include the specified characters.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

### **suffix**

The value sent by the client must end with the specified characters.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# HealthCheckPolicy

An object that represents the health check policy for a virtual node's listener.

## Contents

### healthyThreshold

The number of consecutive successful health checks that must occur before declaring listener healthy.

Type: Integer

Valid Range: Minimum value of 2. Maximum value of 10.

Required: Yes

### intervalMillis

The time period in milliseconds between each health check execution.

Type: Long

Valid Range: Minimum value of 5000. Maximum value of 300000.

Required: Yes

### protocol

The protocol for the health check request. If you specify `grpc`, then your service must conform to the [GRPC Health Checking Protocol](#).

Type: String

Valid Values: `http` | `tcp` | `http2` | `grpc`

Required: Yes

### timeoutMillis

The amount of time to wait when receiving a response from the health check, in milliseconds.

Type: Long

Valid Range: Minimum value of 2000. Maximum value of 60000.

Required: Yes

### **unhealthyThreshold**

The number of consecutive failed health checks that must occur before declaring a virtual node unhealthy.

Type: Integer

Valid Range: Minimum value of 2. Maximum value of 10.

Required: Yes

### **path**

The destination path for the health check request. This value is only used if the specified protocol is HTTP or HTTP/2. For any other protocol, this value is ignored.

Type: String

Required: No

### **port**

The destination port for the health check request. This port must match the port defined in the [PortMapping](#) for the listener.

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

# HttpGatewayRoute

An object that represents an HTTP gateway route.

## Contents

### action

An object that represents the action to take if a match is determined.

Type: [HttpGatewayRouteAction](#) object

Required: Yes

### match

An object that represents the criteria for determining a request match.

Type: [HttpGatewayRouteMatch](#) object

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

# HttpGatewayRouteAction

An object that represents the action to take if a match is determined.

## Contents

### target

An object that represents the target that traffic is routed to when a request matches the gateway route.

Type: [GatewayRouteTarget](#) object

Required: Yes

### rewrite

The gateway route action to rewrite.

Type: [HttpGatewayRouteRewrite](#) object

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

# HttpGatewayRouteHeader

An object that represents the HTTP header in the gateway route.

## Contents

### name

A name for the HTTP header in the gateway route that will be matched on.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 50.

Required: Yes

### invert

Specify `True` to match anything except the match criteria. The default value is `False`.

Type: Boolean

Required: No

### match

An object that represents the method and value to match with the header value sent in a request. Specify one match method.

Type: [HeaderMatchMethod](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

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 Java V2](#)

- [AWS SDK for Ruby V3](#)

# HttpGatewayRouteMatch

An object that represents the criteria for determining a request match.

## Contents

### headers

The client request headers to match on.

Type: Array of [HttpGatewayRouteHeader](#) objects

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

Required: No

### hostname

The host name to match on.

Type: [GatewayRouteHostnameMatch](#) object

Required: No

### method

The method to match on.

Type: String

Valid Values: GET | HEAD | POST | PUT | DELETE | CONNECT | OPTIONS | TRACE | PATCH

Required: No

### path

The path to match on.

Type: [HttpPathMatch](#) object

Required: No

### port

The port number to match on.

Type: Integer

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

Required: No

## prefix

Specifies the path to match requests with. This parameter must always start with `/`, which by itself matches all requests to the virtual service name. You can also match for path-based routing of requests. For example, if your virtual service name is `my-service.local` and you want the route to match requests to `my-service.local/metrics`, your prefix should be `/metrics`.

Type: String

Required: No

## queryParameters

The query parameter to match on.

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

# HttpGatewayRoutePathRewrite

An object that represents the path to rewrite.

## Contents

### exact

The exact path to rewrite.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# HttpGatewayRoutePrefixRewrite

An object representing the beginning characters of the route to rewrite.

## Contents

### defaultPrefix

The default prefix used to replace the incoming route prefix when rewritten.

Type: String

Valid Values: ENABLED | DISABLED

Required: No

### value

The value used to replace the incoming route prefix when rewritten.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# HttpGatewayRouteRewrite

An object representing the gateway route to rewrite.

## Contents

### hostname

The host name to rewrite.

Type: [GatewayRouteHostnameRewrite](#) object

Required: No

### path

The path to rewrite.

Type: [HttpGatewayRoutePathRewrite](#) object

Required: No

### prefix

The specified beginning characters to rewrite.

Type: [HttpGatewayRoutePrefixRewrite](#) object

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

# HttpPathMatch

An object representing the path to match in the request.

## Contents

### exact

The exact path to match on.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

### regex

The regex used to match the path.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# HttpQueryParameter

An object that represents the query parameter in the request.

## Contents

### name

A name for the query parameter that will be matched on.

Type: String

Required: Yes

### match

The query parameter to match on.

Type: [QueryParameterMatch](#) object

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

# HttpRetryPolicy

An object that represents a retry policy. Specify at least one value for at least one of the types of `RetryEvents`, a value for `maxRetries`, and a value for `perRetryTimeout`. Both `server-error` and `gateway-error` under `httpRetryEvents` include the Envoy `reset` policy. For more information on the `reset` policy, see the [Envoy documentation](#).

## Contents

### `maxRetries`

The maximum number of retry attempts.

Type: Long

Valid Range: Minimum value of 0.

Required: Yes

### `perRetryTimeout`

The timeout for each retry attempt.

Type: [Duration](#) object

Required: Yes

### `httpRetryEvents`

Specify at least one of the following values.

- **server-error** – HTTP status codes 500, 501, 502, 503, 504, 505, 506, 507, 508, 510, and 511
- **gateway-error** – HTTP status codes 502, 503, and 504
- **client-error** – HTTP status code 409
- **stream-error** – Retry on refused stream

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 25 items.

Length Constraints: Minimum length of 1. Maximum length of 25.

Required: No

## tcpRetryEvents

Specify a valid value. The event occurs before any processing of a request has started and is encountered when the upstream is temporarily or permanently unavailable.

Type: Array of strings

Array Members: Fixed number of 1 item.

Valid Values: `connection-error`

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

# HttpRoute

An object that represents an HTTP or HTTP/2 route type.

## Contents

### action

An object that represents the action to take if a match is determined.

Type: [HttpRouteAction](#) object

Required: Yes

### match

An object that represents the criteria for determining a request match.

Type: [HttpRouteMatch](#) object

Required: Yes

### retryPolicy

An object that represents a retry policy.

Type: [HttpRetryPolicy](#) object

Required: No

### timeout

An object that represents types of timeouts.

Type: [HttpTimeout](#) object

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

# HttpRouteAction

An object that represents the action to take if a match is determined.

## Contents

### **weightedTargets**

An object that represents the targets that traffic is routed to when a request matches the route.

Type: Array of [WeightedTarget](#) objects

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

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

# HttpRequestHeader

An object that represents the HTTP header in the request.

## Contents

### name

A name for the HTTP header in the client request that will be matched on.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 50.

Required: Yes

### invert

Specify `True` to match anything except the match criteria. The default value is `False`.

Type: Boolean

Required: No

### match

The `HeaderMatchMethod` object.

Type: [HeaderMatchMethod](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

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



# HttpRouteMatch

An object that represents the requirements for a route to match HTTP requests for a virtual router.

## Contents

### headers

The client request headers to match on.

Type: Array of [HttpRouteHeader](#) objects

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

Required: No

### method

The client request method to match on. Specify only one.

Type: String

Valid Values: GET | HEAD | POST | PUT | DELETE | CONNECT | OPTIONS | TRACE | PATCH

Required: No

### path

The client request path to match on.

Type: [HttpPathMatch](#) object

Required: No

### port

The port number to match on.

Type: Integer

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

Required: No

## prefix

Specifies the path to match requests with. This parameter must always start with `/`, which by itself matches all requests to the virtual service name. You can also match for path-based routing of requests. For example, if your virtual service name is `my-service.local` and you want the route to match requests to `my-service.local/metrics`, your prefix should be `/metrics`.

Type: String

Required: No

## queryParameters

The client request query parameters to match on.

Type: Array of [HttpQueryParameter](#) objects

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

Required: No

## scheme

The client request scheme to match on. Specify only one. Applicable only for HTTP2 routes.

Type: String

Valid Values: `http` | `https`

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

# HttpTimeout

An object that represents types of timeouts.

## Contents

### idle

An object that represents an idle timeout. An idle timeout bounds the amount of time that a connection may be idle. The default value is none.

Type: [Duration](#) object

Required: No

### perRequest

An object that represents a per request timeout. The default value is 15 seconds. If you set a higher timeout, then make sure that the higher value is set for each App Mesh resource in a conversation. For example, if a virtual node backend uses a virtual router provider to route to another virtual node, then the timeout should be greater than 15 seconds for the source and destination virtual node and the route.

Type: [Duration](#) object

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

# JsonFormatRef

An object that represents the key value pairs for the JSON.

## Contents

### key

The specified key for the JSON.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Required: Yes

### value

The specified value for the JSON.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

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

# Listener

An object that represents a listener for a virtual node.

## Contents

### portMapping

The port mapping information for the listener.

Type: [PortMapping](#) object

Required: Yes

### connectionPool

The connection pool information for the listener.

Type: [VirtualNodeConnectionPool](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

Required: No

### healthCheck

The health check information for the listener.

Type: [HealthCheckPolicy](#) object

Required: No

### outlierDetection

The outlier detection information for the listener.

Type: [OutlierDetection](#) object

Required: No

### timeout

An object that represents timeouts for different protocols.

Type: [ListenerTimeout](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

Required: No

## tls

A reference to an object that represents the Transport Layer Security (TLS) properties for a listener.

Type: [ListenerTls](#) object

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

# ListenerTimeout

An object that represents timeouts for different protocols.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### grpc

An object that represents types of timeouts.

Type: [GrpcTimeout](#) object

Required: No

### http

An object that represents types of timeouts.

Type: [HttpTimeout](#) object

Required: No

### http2

An object that represents types of timeouts.

Type: [HttpTimeout](#) object

Required: No

### tcp

An object that represents types of timeouts.

Type: [TcpTimeout](#) object

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

# ListenerTls

An object that represents the Transport Layer Security (TLS) properties for a listener.

## Contents

### certificate

A reference to an object that represents a listener's Transport Layer Security (TLS) certificate.

Type: [ListenerTlsCertificate](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

Required: Yes

### mode

Specify one of the following modes.

- STRICT – Listener only accepts connections with TLS enabled.
- PERMISSIVE – Listener accepts connections with or without TLS enabled.
- DISABLED – Listener only accepts connections without TLS.

Type: String

Valid Values: STRICT | PERMISSIVE | DISABLED

Required: Yes

### validation

A reference to an object that represents a listener's Transport Layer Security (TLS) validation context.

Type: [ListenerTlsValidationContext](#) object

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

# ListenerTlsAcmeCertificate

An object that represents an AWS Certificate Manager certificate.

## Contents

### certificateArn

The Amazon Resource Name (ARN) for the certificate. The certificate must meet specific requirements and you must have proxy authorization enabled. For more information, see [Transport Layer Security \(TLS\)](#).

Type: String

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

# ListenerTlsCertificate

An object that represents a listener's Transport Layer Security (TLS) certificate.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### acm

A reference to an object that represents an AWS Certificate Manager certificate.

Type: [ListenerTlsAcmCertificate](#) object

Required: No

### file

A reference to an object that represents a local file certificate.

Type: [ListenerTlsFileCertificate](#) object

Required: No

### sds

A reference to an object that represents a listener's Secret Discovery Service certificate.

Type: [ListenerTlsSdsCertificate](#) object

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

# ListenerTlsFileCertificate

An object that represents a local file certificate. The certificate must meet specific requirements and you must have proxy authorization enabled. For more information, see [Transport Layer Security \(TLS\)](#).

## Contents

### **certificateChain**

The certificate chain for the certificate.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### **privateKey**

The private key for a certificate stored on the file system of the virtual node that the proxy is running on.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# ListenerTlsSdsCertificate

An object that represents the listener's Secret Discovery Service certificate. The proxy must be configured with a local SDS provider via a Unix Domain Socket. See App Mesh [TLS documentation](#) for more info.

## Contents

### secretName

A reference to an object that represents the name of the secret requested from the Secret Discovery Service provider representing Transport Layer Security (TLS) materials like a certificate or certificate chain.

Type: String

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

# ListenerTlsValidationContext

An object that represents a listener's Transport Layer Security (TLS) validation context.

## Contents

### trust

A reference to where to retrieve the trust chain when validating a peer's Transport Layer Security (TLS) certificate.

Type: [ListenerTlsValidationContextTrust](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

Required: Yes

### subjectAlternativeNames

A reference to an object that represents the SANs for a listener's Transport Layer Security (TLS) validation context.

Type: [SubjectAlternativeNames](#) object

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

# ListenerTlsValidationContextTrust

An object that represents a listener's Transport Layer Security (TLS) validation context trust.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### file

An object that represents a Transport Layer Security (TLS) validation context trust for a local file.

Type: [TlsValidationContextFileTrust](#) object

Required: No

### sds

A reference to an object that represents a listener's Transport Layer Security (TLS) Secret Discovery Service validation context trust.

Type: [TlsValidationContextSdsTrust](#) object

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

# Logging

An object that represents the logging information for a virtual node.

## Contents

### accessLog

The access log configuration for a virtual node.

Type: [AccessLog](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

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

# LoggingFormat

An object that represents the format for the logs.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### json

The logging format for JSON.

Type: Array of [JsonFormatRef](#) objects

Required: No

### text

The logging format for text.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1000.

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

# MatchRange

An object that represents the range of values to match on. The first character of the range is included in the range, though the last character is not. For example, if the range specified were 1-100, only values 1-99 would be matched.

## Contents

### end

The end of the range.

Type: Long

Required: Yes

### start

The start of the range.

Type: Long

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

# MeshData

An object that represents a service mesh returned by a describe operation.

## Contents

### meshName

The name of the service mesh.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### metadata

The associated metadata for the service mesh.

Type: [ResourceMetadata](#) object

Required: Yes

### spec

The associated specification for the service mesh.

Type: [MeshSpec](#) object

Required: Yes

### status

The status of the service mesh.

Type: [MeshStatus](#) object

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

# MeshRef

An object that represents a service mesh returned by a list operation.

## Contents

### **arn**

The full Amazon Resource Name (ARN) of the service mesh.

Type: String

Required: Yes

### **createdAt**

The Unix epoch timestamp in seconds for when the resource was created.

Type: Timestamp

Required: Yes

### **lastUpdatedAt**

The Unix epoch timestamp in seconds for when the resource was last updated.

Type: Timestamp

Required: Yes

### **meshName**

The name of the service mesh.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### **meshOwner**

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Type: String

Length Constraints: Fixed length of 12.

Required: Yes

### **resourceOwner**

The AWS IAM account ID of the resource owner. If the account ID is not your own, then it's the ID of the mesh owner or of another account that the mesh is shared with. For more information about mesh sharing, see [Working with shared meshes](#).

Type: String

Length Constraints: Fixed length of 12.

Required: Yes

### **version**

The version of the resource. Resources are created at version 1, and this version is incremented each time that they're updated.

Type: Long

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

# MeshServiceDiscovery

An object that represents the service discovery information for a service mesh.

## Contents

### ipPreference

The IP version to use to control traffic within the mesh.

Type: String

Valid Values: IPv6\_PREFERRED | IPv4\_PREFERRED | IPv4\_ONLY | IPv6\_ONLY

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

# MeshSpec

An object that represents the specification of a service mesh.

## Contents

### **egressFilter**

The egress filter rules for the service mesh.

Type: [EgressFilter](#) object

Required: No

### **serviceDiscovery**

An object that represents the service discovery information for a service mesh.

Type: [MeshServiceDiscovery](#) object

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

# MeshStatus

An object that represents the status of a service mesh.

## Contents

### status

The current mesh status.

Type: String

Valid Values: ACTIVE | INACTIVE | DELETED

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

# OutlierDetection

An object that represents the outlier detection for a virtual node's listener.

## Contents

### baseEjectionDuration

The base amount of time for which a host is ejected.

Type: [Duration](#) object

Required: Yes

### interval

The time interval between ejection sweep analysis.

Type: [Duration](#) object

Required: Yes

### maxEjectionPercent

Maximum percentage of hosts in load balancing pool for upstream service that can be ejected. Will eject at least one host regardless of the value.

Type: Integer

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

Required: Yes

### maxServerErrors

Number of consecutive 5xx errors required for ejection.

Type: Long

Valid Range: Minimum value of 1.

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

# PortMapping

An object that represents a port mapping.

## Contents

### port

The port used for the port mapping.

Type: Integer

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

Required: Yes

### protocol

The protocol used for the port mapping. Specify one protocol.

Type: String

Valid Values: `http` | `tcp` | `http2` | `grpc`

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

# QueryParameterMatch

An object representing the query parameter to match.

## Contents

### exact

The exact query parameter to match on.

Type: String

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

# ResourceMetadata

An object that represents metadata for a resource.

## Contents

### arn

The full Amazon Resource Name (ARN) for the resource.

Type: String

Required: Yes

### createdAt

The Unix epoch timestamp in seconds for when the resource was created.

Type: Timestamp

Required: Yes

### lastUpdatedAt

The Unix epoch timestamp in seconds for when the resource was last updated.

Type: Timestamp

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Type: String

Length Constraints: Fixed length of 12.

Required: Yes

### resourceOwner

The AWS IAM account ID of the resource owner. If the account ID is not your own, then it's the ID of the mesh owner or of another account that the mesh is shared with. For more information about mesh sharing, see [Working with shared meshes](#).

Type: String

Length Constraints: Fixed length of 12.

Required: Yes

### **uid**

The unique identifier for the resource.

Type: String

Required: Yes

### **version**

The version of the resource. Resources are created at version 1, and this version is incremented each time that they're updated.

Type: Long

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

# RouteData

An object that represents a route returned by a describe operation.

## Contents

### meshName

The name of the service mesh that the route resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### metadata

The associated metadata for the route.

Type: [ResourceMetadata](#) object

Required: Yes

### routeName

The name of the route.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### spec

The specifications of the route.

Type: [RouteSpec](#) object

Required: Yes

### status

The status of the route.

Type: [RouteStatus](#) object

Required: Yes

### **virtualRouterName**

The virtual router that the route is associated with.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# RouteRef

An object that represents a route returned by a list operation.

## Contents

### **arn**

The full Amazon Resource Name (ARN) for the route.

Type: String

Required: Yes

### **createdAt**

The Unix epoch timestamp in seconds for when the resource was created.

Type: Timestamp

Required: Yes

### **lastUpdatedAt**

The Unix epoch timestamp in seconds for when the resource was last updated.

Type: Timestamp

Required: Yes

### **meshName**

The name of the service mesh that the route resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### **meshOwner**

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Type: String

Length Constraints: Fixed length of 12.

Required: Yes

### **resourceOwner**

The AWS IAM account ID of the resource owner. If the account ID is not your own, then it's the ID of the mesh owner or of another account that the mesh is shared with. For more information about mesh sharing, see [Working with shared meshes](#).

Type: String

Length Constraints: Fixed length of 12.

Required: Yes

### **routeName**

The name of the route.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### **version**

The version of the resource. Resources are created at version 1, and this version is incremented each time that they're updated.

Type: Long

Required: Yes

### **virtualRouterName**

The virtual router that the route is associated with.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# RouteSpec

An object that represents a route specification. Specify one route type.

## Contents

### grpcRoute

An object that represents the specification of a gRPC route.

Type: [GrpcRoute](#) object

Required: No

### http2Route

An object that represents the specification of an HTTP/2 route.

Type: [HttpRoute](#) object

Required: No

### httpRoute

An object that represents the specification of an HTTP route.

Type: [HttpRoute](#) object

Required: No

### priority

The priority for the route. Routes are matched based on the specified value, where 0 is the highest priority.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

Required: No

### tcpRoute

An object that represents the specification of a TCP route.

Type: [TcpRoute](#) object

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

# RouteStatus

An object that represents the current status of a route.

## Contents

### status

The current status for the route.

Type: String

Valid Values: ACTIVE | INACTIVE | DELETED

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

# ServiceDiscovery

An object that represents the service discovery information for a virtual node.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### awsCloudMap

Specifies any AWS Cloud Map information for the virtual node.

Type: [AwsCloudMapServiceDiscovery](#) object

Required: No

### dns

Specifies the DNS information for the virtual node.

Type: [DnsServiceDiscovery](#) object

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

# SubjectAlternativeNameMatchers

An object that represents the methods by which a subject alternative name on a peer Transport Layer Security (TLS) certificate can be matched.

## Contents

### exact

The values sent must match the specified values exactly.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 254.

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

# SubjectAlternativeNames

An object that represents the subject alternative names secured by the certificate.

## Contents

### match

An object that represents the criteria for determining a SANs match.

Type: [SubjectAlternativeNameMatchers](#) object

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

# TagRef

Optional metadata that you apply to a resource to assist with categorization and organization. Each tag consists of a key and an optional value, both of which you define. Tag keys can have a maximum character length of 128 characters, and tag values can have a maximum length of 256 characters.

## Contents

### key

One part of a key-value pair that make up a tag. A key is a general label that acts like a category for more specific tag values.

Type: String

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

Required: Yes

### value

The optional part of a key-value pair that make up a tag. A value acts as a descriptor within a tag category (key).

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

# TcpRoute

An object that represents a TCP route type.

## Contents

### action

The action to take if a match is determined.

Type: [TcpRouteAction](#) object

Required: Yes

### match

An object that represents the criteria for determining a request match.

Type: [TcpRouteMatch](#) object

Required: No

### timeout

An object that represents types of timeouts.

Type: [TcpTimeout](#) object

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

# TcpRouteAction

An object that represents the action to take if a match is determined.

## Contents

### weightedTargets

An object that represents the targets that traffic is routed to when a request matches the route.

Type: Array of [WeightedTarget](#) objects

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

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

# TcpRouteMatch

An object representing the TCP route to match.

## Contents

### port

The port number to match 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# TcpTimeout

An object that represents types of timeouts.

## Contents

### idle

An object that represents an idle timeout. An idle timeout bounds the amount of time that a connection may be idle. The default value is none.

Type: [Duration](#) object

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

# TlsValidationContext

An object that represents how the proxy will validate its peer during Transport Layer Security (TLS) negotiation.

## Contents

### trust

A reference to where to retrieve the trust chain when validating a peer's Transport Layer Security (TLS) certificate.

Type: [TlsValidationContextTrust](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

Required: Yes

### subjectAlternativeNames

A reference to an object that represents the SANs for a Transport Layer Security (TLS) validation context. If you don't specify SANs on the *terminating* mesh endpoint, the Envoy proxy for that node doesn't verify the SAN on a peer client certificate. If you don't specify SANs on the *originating* mesh endpoint, the SAN on the certificate provided by the terminating endpoint must match the mesh endpoint service discovery configuration. Since SPIRE vended certificates have a SPIFFE ID as a name, you must set the SAN since the name doesn't match the service discovery name.

Type: [SubjectAlternativeNames](#) object

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



# TlsValidationContextAcmTrust

An object that represents a Transport Layer Security (TLS) validation context trust for an AWS Certificate Manager certificate.

## Contents

### **certificateAuthorityArns**

One or more ACM Amazon Resource Name (ARN)s.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 3 items.

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

# TlsValidationContextFileTrust

An object that represents a Transport Layer Security (TLS) validation context trust for a local file.

## Contents

### certificateChain

The certificate trust chain for a certificate stored on the file system of the virtual node that the proxy is running on.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# TlsValidationContextSdsTrust

An object that represents a Transport Layer Security (TLS) Secret Discovery Service validation context trust. The proxy must be configured with a local SDS provider via a Unix Domain Socket. See App Mesh [TLS documentation](#) for more info.

## Contents

### secretName

A reference to an object that represents the name of the secret for a Transport Layer Security (TLS) Secret Discovery Service validation context trust.

Type: String

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

# TlsValidationContextTrust

An object that represents a Transport Layer Security (TLS) validation context trust.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### acm

A reference to an object that represents a Transport Layer Security (TLS) validation context trust for an AWS Certificate Manager certificate.

Type: [TlsValidationContextAcmTrust](#) object

Required: No

### file

An object that represents a Transport Layer Security (TLS) validation context trust for a local file.

Type: [TlsValidationContextFileTrust](#) object

Required: No

### sds

A reference to an object that represents a Transport Layer Security (TLS) Secret Discovery Service validation context trust.

Type: [TlsValidationContextSdsTrust](#) object

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

# VirtualGatewayAccessLog

The access log configuration for a virtual gateway.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### file

The file object to send virtual gateway access logs to.

Type: [VirtualGatewayFileAccessLog](#) object

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

# VirtualGatewayBackendDefaults

An object that represents the default properties for a backend.

## Contents

### clientPolicy

A reference to an object that represents a client policy.

Type: [VirtualGatewayClientPolicy](#) object

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

# VirtualGatewayClientPolicy

An object that represents a client policy.

## Contents

### tls

A reference to an object that represents a Transport Layer Security (TLS) client policy.

Type: [VirtualGatewayClientPolicyTls](#) object

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

# VirtualGatewayClientPolicyTls

An object that represents a Transport Layer Security (TLS) client policy.

## Contents

### validation

A reference to an object that represents a Transport Layer Security (TLS) validation context.

Type: [VirtualGatewayTlsValidationContext](#) object

Required: Yes

### certificate

A reference to an object that represents a virtual gateway's client's Transport Layer Security (TLS) certificate.

Type: [VirtualGatewayClientTlsCertificate](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

Required: No

### enforce

Whether the policy is enforced. The default is `True`, if a value isn't specified.

Type: Boolean

Required: No

### ports

One or more ports that the policy is enforced for.

Type: Array of integers

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

# VirtualGatewayClientTlsCertificate

An object that represents the virtual gateway's client's Transport Layer Security (TLS) certificate.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### file

An object that represents a local file certificate. The certificate must meet specific requirements and you must have proxy authorization enabled. For more information, see [Transport Layer Security \(TLS\)](#).

Type: [VirtualGatewayListenerTlsFileCertificate](#) object

Required: No

### sds

A reference to an object that represents a virtual gateway's client's Secret Discovery Service certificate.

Type: [VirtualGatewayListenerTlsSdsCertificate](#) object

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



# VirtualGatewayConnectionPool

An object that represents the type of virtual gateway connection pool.

Only one protocol is used at a time and should be the same protocol as the one chosen under port mapping.

If not present the default value for `maxPendingRequests` is 2147483647.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### grpc

An object that represents a type of connection pool.

Type: [VirtualGatewayGrpcConnectionPool](#) object

Required: No

### http

An object that represents a type of connection pool.

Type: [VirtualGatewayHttpConnectionPool](#) object

Required: No

### http2

An object that represents a type of connection pool.

Type: [VirtualGatewayHttp2ConnectionPool](#) object

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

# VirtualGatewayData

An object that represents a virtual gateway returned by a describe operation.

## Contents

### meshName

The name of the service mesh that the virtual gateway resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### metadata

An object that represents metadata for a resource.

Type: [ResourceMetadata](#) object

Required: Yes

### spec

The specifications of the virtual gateway.

Type: [VirtualGatewaySpec](#) object

Required: Yes

### status

The current status of the virtual gateway.

Type: [VirtualGatewayStatus](#) object

Required: Yes

### virtualGatewayName

The name of the virtual gateway.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# VirtualGatewayFileAccessLog

An object that represents an access log file.

## Contents

### path

The file path to write access logs to. You can use `/dev/stdout` to send access logs to standard out and configure your Envoy container to use a log driver, such as `awslogs`, to export the access logs to a log storage service such as Amazon CloudWatch Logs. You can also specify a path in the Envoy container's file system to write the files to disk.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### format

The specified format for the virtual gateway access logs. It can be either `json_format` or `text_format`.

Type: [LoggingFormat](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

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

# VirtualGatewayGrpcConnectionPool

An object that represents a type of connection pool.

## Contents

### maxRequests

Maximum number of inflight requests Envoy can concurrently support across hosts in upstream cluster.

Type: Integer

Valid Range: Minimum value of 1.

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

# VirtualGatewayHealthCheckPolicy

An object that represents the health check policy for a virtual gateway's listener.

## Contents

### healthyThreshold

The number of consecutive successful health checks that must occur before declaring the listener healthy.

Type: Integer

Valid Range: Minimum value of 2. Maximum value of 10.

Required: Yes

### intervalMillis

The time period in milliseconds between each health check execution.

Type: Long

Valid Range: Minimum value of 5000. Maximum value of 300000.

Required: Yes

### protocol

The protocol for the health check request. If you specify `grpc`, then your service must conform to the [GRPC Health Checking Protocol](#).

Type: String

Valid Values: `http` | `http2` | `grpc`

Required: Yes

### timeoutMillis

The amount of time to wait when receiving a response from the health check, in milliseconds.

Type: Long

Valid Range: Minimum value of 2000. Maximum value of 60000.

Required: Yes

### **unhealthyThreshold**

The number of consecutive failed health checks that must occur before declaring a virtual gateway unhealthy.

Type: Integer

Valid Range: Minimum value of 2. Maximum value of 10.

Required: Yes

### **path**

The destination path for the health check request. This value is only used if the specified protocol is HTTP or HTTP/2. For any other protocol, this value is ignored.

Type: String

Required: No

### **port**

The destination port for the health check request. This port must match the port defined in the [PortMapping](#) for the listener.

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

# VirtualGatewayHttp2ConnectionPool

An object that represents a type of connection pool.

## Contents

### maxRequests

Maximum number of inflight requests Envoy can concurrently support across hosts in upstream cluster.

Type: Integer

Valid Range: Minimum value of 1.

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

# VirtualGatewayHttpConnectionPool

An object that represents a type of connection pool.

## Contents

### maxConnections

Maximum number of outbound TCP connections Envoy can establish concurrently with all hosts in upstream cluster.

Type: Integer

Valid Range: Minimum value of 1.

Required: Yes

### maxPendingRequests

Number of overflowing requests after `max_connections` Envoy will queue to upstream cluster.

Type: Integer

Valid Range: Minimum value of 1.

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

# VirtualGatewayListener

An object that represents a listener for a virtual gateway.

## Contents

### portMapping

The port mapping information for the listener.

Type: [VirtualGatewayPortMapping](#) object

Required: Yes

### connectionPool

The connection pool information for the virtual gateway listener.

Type: [VirtualGatewayConnectionPool](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

Required: No

### healthCheck

The health check information for the listener.

Type: [VirtualGatewayHealthCheckPolicy](#) object

Required: No

### tls

A reference to an object that represents the Transport Layer Security (TLS) properties for the listener.

Type: [VirtualGatewayListenerTls](#) object

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

# VirtualGatewayListenerTls

An object that represents the Transport Layer Security (TLS) properties for a listener.

## Contents

### certificate

An object that represents a Transport Layer Security (TLS) certificate.

Type: [VirtualGatewayListenerTlsCertificate](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

Required: Yes

### mode

Specify one of the following modes.

- STRICT – Listener only accepts connections with TLS enabled.
- PERMISSIVE – Listener accepts connections with or without TLS enabled.
- DISABLED – Listener only accepts connections without TLS.

Type: String

Valid Values: STRICT | PERMISSIVE | DISABLED

Required: Yes

### validation

A reference to an object that represents a virtual gateway's listener's Transport Layer Security (TLS) validation context.

Type: [VirtualGatewayListenerTlsValidationContext](#) object

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

# VirtualGatewayListenerTlsAcmCertificate

An object that represents an AWS Certificate Manager certificate.

## Contents

### certificateArn

The Amazon Resource Name (ARN) for the certificate. The certificate must meet specific requirements and you must have proxy authorization enabled. For more information, see [Transport Layer Security \(TLS\)](#).

Type: String

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

# VirtualGatewayListenerTlsCertificate

An object that represents a listener's Transport Layer Security (TLS) certificate.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### acm

A reference to an object that represents an AWS Certificate Manager certificate.

Type: [VirtualGatewayListenerTlsAcmCertificate](#) object

Required: No

### file

A reference to an object that represents a local file certificate.

Type: [VirtualGatewayListenerTlsFileCertificate](#) object

Required: No

### sds

A reference to an object that represents a virtual gateway's listener's Secret Discovery Service certificate.

Type: [VirtualGatewayListenerTlsSdsCertificate](#) object

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

# VirtualGatewayListenerTlsFileCertificate

An object that represents a local file certificate. The certificate must meet specific requirements and you must have proxy authorization enabled. For more information, see [Transport Layer Security \(TLS\)](#).

## Contents

### certificateChain

The certificate chain for the certificate.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### privateKey

The private key for a certificate stored on the file system of the mesh endpoint that the proxy is running on.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# VirtualGatewayListenerTlsSdsCertificate

An object that represents the virtual gateway's listener's Secret Discovery Service certificate. The proxy must be configured with a local SDS provider via a Unix Domain Socket. See App Mesh [TLS documentation](#) for more info.

## Contents

### secretName

A reference to an object that represents the name of the secret secret requested from the Secret Discovery Service provider representing Transport Layer Security (TLS) materials like a certificate or certificate chain.

Type: String

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

# VirtualGatewayListenerTlsValidationContext

An object that represents a virtual gateway's listener's Transport Layer Security (TLS) validation context.

## Contents

### trust

A reference to where to retrieve the trust chain when validating a peer's Transport Layer Security (TLS) certificate.

Type: [VirtualGatewayListenerTlsValidationContextTrust](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

Required: Yes

### subjectAlternativeNames

A reference to an object that represents the SANs for a virtual gateway listener's Transport Layer Security (TLS) validation context.

Type: [SubjectAlternativeNames](#) object

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

# VirtualGatewayListenerTlsValidationContextTrust

An object that represents a virtual gateway's listener's Transport Layer Security (TLS) validation context trust.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### file

An object that represents a Transport Layer Security (TLS) validation context trust for a local file.

Type: [VirtualGatewayTlsValidationContextFileTrust](#) object

Required: No

### sds

A reference to an object that represents a virtual gateway's listener's Transport Layer Security (TLS) Secret Discovery Service validation context trust.

Type: [VirtualGatewayTlsValidationContextSdsTrust](#) object

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



# VirtualGatewayLogging

An object that represents logging information.

## Contents

### accessLog

The access log configuration.

Type: [VirtualGatewayAccessLog](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

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

# VirtualGatewayPortMapping

An object that represents a port mapping.

## Contents

### port

The port used for the port mapping. Specify one protocol.

Type: Integer

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

Required: Yes

### protocol

The protocol used for the port mapping.

Type: String

Valid Values: `http` | `http2` | `grpc`

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

# VirtualGatewayRef

An object that represents a virtual gateway returned by a list operation.

## Contents

### arn

The full Amazon Resource Name (ARN) for the resource.

Type: String

Required: Yes

### createdAt

The Unix epoch timestamp in seconds for when the resource was created.

Type: Timestamp

Required: Yes

### lastUpdatedAt

The Unix epoch timestamp in seconds for when the resource was last updated.

Type: Timestamp

Required: Yes

### meshName

The name of the service mesh that the resource resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Type: String

Length Constraints: Fixed length of 12.

Required: Yes

### **resourceOwner**

The AWS IAM account ID of the resource owner. If the account ID is not your own, then it's the ID of the mesh owner or of another account that the mesh is shared with. For more information about mesh sharing, see [Working with shared meshes](#).

Type: String

Length Constraints: Fixed length of 12.

Required: Yes

### **version**

The version of the resource. Resources are created at version 1, and this version is incremented each time that they're updated.

Type: Long

Required: Yes

### **virtualGatewayName**

The name of the resource.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)

- [AWS SDK for Ruby V3](#)

# VirtualGatewaySpec

An object that represents the specification of a service mesh resource.

## Contents

### listeners

The listeners that the mesh endpoint is expected to receive inbound traffic from. You can specify one listener.

Type: Array of [VirtualGatewayListener](#) objects

Required: Yes

### backendDefaults

A reference to an object that represents the defaults for backends.

Type: [VirtualGatewayBackendDefaults](#) object

Required: No

### logging

An object that represents logging information.

Type: [VirtualGatewayLogging](#) object

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

# VirtualGatewayStatus

An object that represents the status of the mesh resource.

## Contents

### status

The current status.

Type: String

Valid Values: ACTIVE | INACTIVE | DELETED

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

# VirtualGatewayTlsValidationContext

An object that represents a Transport Layer Security (TLS) validation context.

## Contents

### trust

A reference to where to retrieve the trust chain when validating a peer's Transport Layer Security (TLS) certificate.

Type: [VirtualGatewayTlsValidationContextTrust](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

Required: Yes

### subjectAlternativeNames

A reference to an object that represents the SANs for a virtual gateway's listener's Transport Layer Security (TLS) validation context.

Type: [SubjectAlternativeNames](#) object

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

# VirtualGatewayTlsValidationContextAcmTrust

An object that represents a Transport Layer Security (TLS) validation context trust for an AWS Certificate Manager certificate.

## Contents

### certificateAuthorityArns

One or more ACM Amazon Resource Name (ARN)s.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 3 items.

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

# VirtualGatewayTlsValidationContextFileTrust

An object that represents a Transport Layer Security (TLS) validation context trust for a local file.

## Contents

### certificateChain

The certificate trust chain for a certificate stored on the file system of the virtual node that the proxy is running on.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# VirtualGatewayTlsValidationContextSdsTrust

An object that represents a virtual gateway's listener's Transport Layer Security (TLS) Secret Discovery Service validation context trust. The proxy must be configured with a local SDS provider via a Unix Domain Socket. See App Mesh [TLS documentation](#) for more info.

## Contents

### secretName

A reference to an object that represents the name of the secret for a virtual gateway's Transport Layer Security (TLS) Secret Discovery Service validation context trust.

Type: String

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

# VirtualGatewayTlsValidationContextTrust

An object that represents a Transport Layer Security (TLS) validation context trust.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### acm

A reference to an object that represents a Transport Layer Security (TLS) validation context trust for an AWS Certificate Manager certificate.

Type: [VirtualGatewayTlsValidationContextAcmTrust](#) object

Required: No

### file

An object that represents a Transport Layer Security (TLS) validation context trust for a local file.

Type: [VirtualGatewayTlsValidationContextFileTrust](#) object

Required: No

### sds

A reference to an object that represents a virtual gateway's Transport Layer Security (TLS) Secret Discovery Service validation context trust.

Type: [VirtualGatewayTlsValidationContextSdsTrust](#) object

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

# VirtualNodeConnectionPool

An object that represents the type of virtual node connection pool.

Only one protocol is used at a time and should be the same protocol as the one chosen under port mapping.

If not present the default value for `maxPendingRequests` is 2147483647.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### grpc

An object that represents a type of connection pool.

Type: [VirtualNodeGrpcConnectionPool](#) object

Required: No

### http

An object that represents a type of connection pool.

Type: [VirtualNodeHttpConnectionPool](#) object

Required: No

### http2

An object that represents a type of connection pool.

Type: [VirtualNodeHttp2ConnectionPool](#) object

Required: No

### tcp

An object that represents a type of connection pool.

Type: [VirtualNodeTcpConnectionPool](#) object

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

# VirtualNodeData

An object that represents a virtual node returned by a describe operation.

## Contents

### meshName

The name of the service mesh that the virtual node resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### metadata

The associated metadata for the virtual node.

Type: [ResourceMetadata](#) object

Required: Yes

### spec

The specifications of the virtual node.

Type: [VirtualNodeSpec](#) object

Required: Yes

### status

The current status for the virtual node.

Type: [VirtualNodeStatus](#) object

Required: Yes

### virtualNodeName

The name of the virtual node.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# VirtualNodeGrpcConnectionPool

An object that represents a type of connection pool.

## Contents

### maxRequests

Maximum number of inflight requests Envoy can concurrently support across hosts in upstream cluster.

Type: Integer

Valid Range: Minimum value of 1.

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

# VirtualNodeHttp2ConnectionPool

An object that represents a type of connection pool.

## Contents

### maxRequests

Maximum number of inflight requests Envoy can concurrently support across hosts in upstream cluster.

Type: Integer

Valid Range: Minimum value of 1.

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

# VirtualNodeHttpConnectionPool

An object that represents a type of connection pool.

## Contents

### maxConnections

Maximum number of outbound TCP connections Envoy can establish concurrently with all hosts in upstream cluster.

Type: Integer

Valid Range: Minimum value of 1.

Required: Yes

### maxPendingRequests

Number of overflowing requests after `max_connections` Envoy will queue to upstream cluster.

Type: Integer

Valid Range: Minimum value of 1.

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

# VirtualNodeRef

An object that represents a virtual node returned by a list operation.

## Contents

### arn

The full Amazon Resource Name (ARN) for the virtual node.

Type: String

Required: Yes

### createdAt

The Unix epoch timestamp in seconds for when the resource was created.

Type: Timestamp

Required: Yes

### lastUpdatedAt

The Unix epoch timestamp in seconds for when the resource was last updated.

Type: Timestamp

Required: Yes

### meshName

The name of the service mesh that the virtual node resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Type: String

Length Constraints: Fixed length of 12.

Required: Yes

### **resourceOwner**

The AWS IAM account ID of the resource owner. If the account ID is not your own, then it's the ID of the mesh owner or of another account that the mesh is shared with. For more information about mesh sharing, see [Working with shared meshes](#).

Type: String

Length Constraints: Fixed length of 12.

Required: Yes

### **version**

The version of the resource. Resources are created at version 1, and this version is incremented each time that they're updated.

Type: Long

Required: Yes

### **virtualNodeName**

The name of the virtual node.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)

- [AWS SDK for Ruby V3](#)

# VirtualNodeServiceProvider

An object that represents a virtual node service provider.

## Contents

### virtualNodeName

The name of the virtual node that is acting as a service provider.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# VirtualNodeSpec

An object that represents the specification of a virtual node.

## Contents

### backendDefaults

A reference to an object that represents the defaults for backends.

Type: [BackendDefaults](#) object

Required: No

### backends

The backends that the virtual node is expected to send outbound traffic to.

#### Important

App Mesh doesn't validate the existence of those virtual services specified in backends. This is to prevent a cyclic dependency between virtual nodes and virtual services creation. Make sure the virtual service name is correct. The virtual service can be created afterwards if it doesn't already exist.

Type: Array of [Backend](#) objects

Required: No

### listeners

The listener that the virtual node is expected to receive inbound traffic from. You can specify one listener.

Type: Array of [Listener](#) objects

Required: No

### logging

The inbound and outbound access logging information for the virtual node.

Type: [Logging](#) object

Required: No

## **serviceDiscovery**

The service discovery information for the virtual node. If your virtual node does not expect ingress traffic, you can omit this parameter. If you specify a `listener`, then you must specify service discovery information.

Type: [ServiceDiscovery](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

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

# VirtualNodeStatus

An object that represents the current status of the virtual node.

## Contents

### status

The current status of the virtual node.

Type: String

Valid Values: ACTIVE | INACTIVE | DELETED

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

# VirtualNodeTcpConnectionPool

An object that represents a type of connection pool.

## Contents

### maxConnections

Maximum number of outbound TCP connections Envoy can establish concurrently with all hosts in upstream cluster.

Type: Integer

Valid Range: Minimum value of 1.

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

# VirtualRouterData

An object that represents a virtual router returned by a describe operation.

## Contents

### meshName

The name of the service mesh that the virtual router resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### metadata

The associated metadata for the virtual router.

Type: [ResourceMetadata](#) object

Required: Yes

### spec

The specifications of the virtual router.

Type: [VirtualRouterSpec](#) object

Required: Yes

### status

The current status of the virtual router.

Type: [VirtualRouterStatus](#) object

Required: Yes

### virtualRouterName

The name of the virtual router.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# VirtualRouterListener

An object that represents a virtual router listener.

## Contents

### portMapping

An object that represents a port mapping.

Type: [PortMapping](#) object

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

# VirtualRouterRef

An object that represents a virtual router returned by a list operation.

## Contents

### **arn**

The full Amazon Resource Name (ARN) for the virtual router.

Type: String

Required: Yes

### **createdAt**

The Unix epoch timestamp in seconds for when the resource was created.

Type: Timestamp

Required: Yes

### **lastUpdatedAt**

The Unix epoch timestamp in seconds for when the resource was last updated.

Type: Timestamp

Required: Yes

### **meshName**

The name of the service mesh that the virtual router resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### **meshOwner**

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Type: String

Length Constraints: Fixed length of 12.

Required: Yes

### **resourceOwner**

The AWS IAM account ID of the resource owner. If the account ID is not your own, then it's the ID of the mesh owner or of another account that the mesh is shared with. For more information about mesh sharing, see [Working with shared meshes](#).

Type: String

Length Constraints: Fixed length of 12.

Required: Yes

### **version**

The version of the resource. Resources are created at version 1, and this version is incremented each time that they're updated.

Type: Long

Required: Yes

### **virtualRouterName**

The name of the virtual router.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)

- [AWS SDK for Ruby V3](#)

# VirtualRouterServiceProvider

An object that represents a virtual node service provider.

## Contents

### virtualRouterName

The name of the virtual router that is acting as a service provider.

Type: String

Length Constraints: Minimum length of 1. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# VirtualRouterSpec

An object that represents the specification of a virtual router.

## Contents

### listeners

The listeners that the virtual router is expected to receive inbound traffic from.

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

# VirtualRouterStatus

An object that represents the status of a virtual router.

## Contents

### status

The current status of the virtual router.

Type: String

Valid Values: ACTIVE | INACTIVE | DELETED

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

# VirtualServiceBackend

An object that represents a virtual service backend for a virtual node.

## Contents

### virtualServiceName

The name of the virtual service that is acting as a virtual node backend.

#### Important

App Mesh doesn't validate the existence of those virtual services specified in backends. This is to prevent a cyclic dependency between virtual nodes and virtual services creation. Make sure the virtual service name is correct. The virtual service can be created afterwards if it doesn't already exist.

Type: String

Required: Yes

### clientPolicy

A reference to an object that represents the client policy for a backend.

Type: [ClientPolicy](#) object

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

# VirtualServiceData

An object that represents a virtual service returned by a describe operation.

## Contents

### meshName

The name of the service mesh that the virtual service resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### metadata

An object that represents metadata for a resource.

Type: [ResourceMetadata](#) object

Required: Yes

### spec

The specifications of the virtual service.

Type: [VirtualServiceSpec](#) object

Required: Yes

### status

The current status of the virtual service.

Type: [VirtualServiceStatus](#) object

Required: Yes

### virtualServiceName

The name of the virtual service.

Type: String

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

# VirtualServiceProvider

An object that represents the provider for a virtual service.

## Contents

### Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

### virtualNode

The virtual node associated with a virtual service.

Type: [VirtualNodeServiceProvider](#) object

Required: No

### virtualRouter

The virtual router associated with a virtual service.

Type: [VirtualRouterServiceProvider](#) object

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

# VirtualServiceRef

An object that represents a virtual service returned by a list operation.

## Contents

### arn

The full Amazon Resource Name (ARN) for the virtual service.

Type: String

Required: Yes

### createdAt

The Unix epoch timestamp in seconds for when the resource was created.

Type: Timestamp

Required: Yes

### lastUpdatedAt

The Unix epoch timestamp in seconds for when the resource was last updated.

Type: Timestamp

Required: Yes

### meshName

The name of the service mesh that the virtual service resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### meshOwner

The AWS IAM account ID of the service mesh owner. If the account ID is not your own, then it's the ID of the account that shared the mesh with your account. For more information about mesh sharing, see [Working with shared meshes](#).

Type: String

Length Constraints: Fixed length of 12.

Required: Yes

### **resourceOwner**

The AWS IAM account ID of the resource owner. If the account ID is not your own, then it's the ID of the mesh owner or of another account that the mesh is shared with. For more information about mesh sharing, see [Working with shared meshes](#).

Type: String

Length Constraints: Fixed length of 12.

Required: Yes

### **version**

The version of the resource. Resources are created at version 1, and this version is incremented each time that they're updated.

Type: Long

Required: Yes

### **virtualServiceName**

The name of the virtual service.

Type: String

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



# VirtualServiceSpec

An object that represents the specification of a virtual service.

## Contents

### provider

The App Mesh object that is acting as the provider for a virtual service. You can specify a single virtual node or virtual router.

Type: [VirtualServiceProvider](#) object

**Note:** This object is a Union. Only one member of this object can be specified or returned.

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

# VirtualServiceStatus

An object that represents the status of a virtual service.

## Contents

### status

The current status of the virtual service.

Type: String

Valid Values: ACTIVE | INACTIVE | DELETED

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

# WeightedTarget

An object that represents a target and its relative weight. Traffic is distributed across targets according to their relative weight. For example, a weighted target with a relative weight of 50 receives five times as much traffic as one with a relative weight of 10. The total weight for all targets combined must be less than or equal to 100.

## Contents

### virtualNode

The virtual node to associate with the weighted target.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

### weight

The relative weight of the weighted target.

Type: Integer

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

Required: Yes

### port

The targeted port of the weighted object.

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