



API Reference

Amazon GameLift Servers



API Version 2015-10-01

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

Amazon GameLift Servers: 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

Welcome	1
Actions	3
AcceptMatch	8
Request Syntax	8
Request Parameters	9
Response Elements	10
Errors	10
See Also	10
ClaimGameServer	12
Request Syntax	12
Request Parameters	13
Response Syntax	14
Response Elements	14
Errors	15
Examples	16
See Also	17
CreateAlias	18
Request Syntax	18
Request Parameters	19
Response Syntax	20
Response Elements	20
Errors	20
See Also	21
CreateBuild	23
Request Syntax	24
Request Parameters	24
Response Syntax	27
Response Elements	27
Errors	28
Examples	29
See Also	31
CreateContainerFleet	32
Request Syntax	33
Request Parameters	34

Response Syntax	42
Response Elements	43
Errors	43
Examples	45
See Also	46
CreateContainerGroupDefinition	47
Request Syntax	49
Request Parameters	51
Response Syntax	54
Response Elements	56
Errors	57
Examples	58
See Also	62
CreateFleet	63
Request Syntax	64
Request Parameters	66
Response Syntax	77
Response Elements	78
Errors	79
Examples	80
See Also	87
CreateFleetLocations	89
Request Syntax	89
Request Parameters	89
Response Syntax	90
Response Elements	91
Errors	92
Examples	93
See Also	94
CreateGameServerGroup	96
Request Syntax	97
Request Parameters	97
Response Syntax	102
Response Elements	102
Errors	103
Examples	103

See Also	107
CreateGameSession	109
Request Syntax	110
Request Parameters	110
Response Syntax	114
Response Elements	115
Errors	115
See Also	117
CreateGameSessionQueue	118
Request Syntax	119
Request Parameters	120
Response Syntax	123
Response Elements	123
Errors	124
Examples	125
See Also	127
CreateLocation	129
Request Syntax	129
Request Parameters	129
Response Syntax	130
Response Elements	130
Errors	131
See Also	132
CreateMatchmakingConfiguration	133
Request Syntax	133
Request Parameters	134
Response Syntax	139
Response Elements	139
Errors	140
Examples	141
See Also	144
CreateMatchmakingRuleSet	145
Request Syntax	145
Request Parameters	145
Response Syntax	147
Response Elements	147

Errors	147
See Also	148
CreatePlayerSession	149
Request Syntax	149
Request Parameters	149
Response Syntax	150
Response Elements	151
Errors	151
See Also	152
CreatePlayerSessions	154
Request Syntax	154
Request Parameters	154
Response Syntax	156
Response Elements	156
Errors	156
See Also	158
CreateScript	159
Request Syntax	159
Request Parameters	160
Response Syntax	162
Response Elements	163
Errors	163
Examples	164
See Also	166
CreateVpcPeeringAuthorization	168
Request Syntax	169
Request Parameters	169
Response Syntax	170
Response Elements	170
Errors	170
Examples	171
See Also	173
CreateVpcPeeringConnection	175
Request Syntax	176
Request Parameters	176
Response Elements	177

Errors	177
Examples	178
See Also	180
DeleteAlias	181
Request Syntax	181
Request Parameters	181
Response Elements	182
Errors	182
See Also	182
DeleteBuild	184
Request Syntax	184
Request Parameters	184
Response Elements	185
Errors	185
Examples	186
See Also	186
DeleteContainerFleet	187
Request Syntax	187
Request Parameters	187
Response Elements	188
Errors	188
Examples	189
See Also	189
DeleteContainerGroupDefinition	191
Request Syntax	191
Request Parameters	192
Response Elements	193
Errors	193
Examples	194
See Also	194
DeleteFleet	196
Request Syntax	196
Request Parameters	196
Response Elements	197
Errors	197
Examples	198

See Also	198
DeleteFleetLocations	200
Request Syntax	200
Request Parameters	200
Response Syntax	201
Response Elements	201
Errors	202
See Also	203
DeleteGameServerGroup	205
Request Syntax	205
Request Parameters	205
Response Syntax	206
Response Elements	207
Errors	207
Examples	208
See Also	209
DeleteGameSessionQueue	211
Request Syntax	211
Request Parameters	211
Response Elements	211
Errors	212
See Also	212
DeleteLocation	214
Request Syntax	214
Request Parameters	214
Response Elements	214
Errors	215
See Also	215
DeleteMatchmakingConfiguration	217
Request Syntax	217
Request Parameters	217
Response Elements	217
Errors	218
See Also	218
DeleteMatchmakingRuleSet	220
Request Syntax	220

Request Parameters	220
Response Elements	221
Errors	221
See Also	221
DeleteScalingPolicy	223
Request Syntax	223
Request Parameters	223
Response Elements	224
Errors	224
Examples	225
See Also	226
DeleteScript	227
Request Syntax	227
Request Parameters	227
Response Elements	228
Errors	228
Examples	229
See Also	229
DeleteVpcPeeringAuthorization	231
Request Syntax	231
Request Parameters	231
Response Elements	232
Errors	232
See Also	233
DeleteVpcPeeringConnection	234
Request Syntax	234
Request Parameters	234
Response Elements	235
Errors	235
See Also	236
DeregisterCompute	237
Request Syntax	237
Request Parameters	237
Response Elements	238
Errors	238
See Also	239

DeregisterGameServer	240
Request Syntax	240
Request Parameters	240
Response Elements	241
Errors	241
Examples	242
See Also	242
DescribeAlias	243
Request Syntax	243
Request Parameters	243
Response Syntax	244
Response Elements	244
Errors	244
See Also	245
DescribeBuild	246
Request Syntax	246
Request Parameters	246
Response Syntax	247
Response Elements	247
Errors	247
Examples	248
See Also	249
DescribeCompute	250
Request Syntax	250
Request Parameters	250
Response Syntax	251
Response Elements	252
Errors	252
See Also	253
DescribeContainerFleet	255
Request Syntax	255
Request Parameters	255
Response Syntax	256
Response Elements	257
Errors	257
Examples	258

See Also	260
DescribeContainerGroupDefinition	261
Request Syntax	261
Request Parameters	261
Response Syntax	262
Response Elements	264
Errors	265
Examples	266
See Also	267
DescribeContainerGroupPortMappings	269
Request Syntax	269
Request Parameters	270
Response Syntax	272
Response Elements	272
Errors	274
Examples	275
See Also	277
DescribeEC2InstanceLimits	279
Request Syntax	280
Request Parameters	280
Response Syntax	284
Response Elements	284
Errors	284
Examples	285
See Also	287
DescribeFleetAttributes	289
Request Syntax	289
Request Parameters	289
Response Syntax	291
Response Elements	292
Errors	292
Examples	293
See Also	295
DescribeFleetCapacity	297
Request Syntax	297
Request Parameters	298

Response Syntax	299
Response Elements	300
Errors	300
Examples	301
See Also	303
DescribeFleetDeployment	304
Request Syntax	304
Request Parameters	304
Response Syntax	305
Response Elements	306
Errors	306
Examples	307
See Also	308
DescribeFleetEvents	309
Request Syntax	309
Request Parameters	309
Response Syntax	311
Response Elements	311
Errors	312
Examples	312
See Also	314
DescribeFleetLocationAttributes	316
Request Syntax	316
Request Parameters	316
Response Syntax	318
Response Elements	318
Errors	319
Examples	320
See Also	321
DescribeFleetLocationCapacity	322
Request Syntax	322
Request Parameters	322
Response Syntax	323
Response Elements	324
Errors	324
Examples	325

See Also	326
DescribeFleetLocationUtilization	327
Request Syntax	327
Request Parameters	327
Response Syntax	328
Response Elements	329
Errors	329
Examples	330
See Also	330
DescribeFleetPortSettings	332
Request Syntax	332
Request Parameters	332
Response Syntax	333
Response Elements	334
Errors	335
Examples	336
See Also	337
DescribeFleetUtilization	339
Request Syntax	339
Request Parameters	340
Response Syntax	341
Response Elements	341
Errors	342
Examples	342
See Also	344
DescribeGameServer	346
Request Syntax	346
Request Parameters	346
Response Syntax	347
Response Elements	347
Errors	348
Examples	348
See Also	349
DescribeGameServerGroup	351
Request Syntax	351
Request Parameters	351

Response Syntax	352
Response Elements	352
Errors	352
Examples	353
See Also	354
DescribeGameServerInstances	356
Request Syntax	356
Request Parameters	356
Response Syntax	358
Response Elements	358
Errors	358
Examples	359
See Also	360
DescribeGameSessionDetails	362
Request Syntax	362
Request Parameters	363
Response Syntax	365
Response Elements	366
Errors	366
See Also	367
DescribeGameSessionPlacement	369
Request Syntax	369
Request Parameters	369
Response Syntax	370
Response Elements	371
Errors	371
See Also	372
DescribeGameSessionQueues	373
Request Syntax	373
Request Parameters	373
Response Syntax	374
Response Elements	375
Errors	375
See Also	376
DescribeGameSessions	378
Request Syntax	378

Request Parameters	379
Response Syntax	381
Response Elements	382
Errors	382
See Also	383
DescribeInstances	385
Request Syntax	385
Request Parameters	386
Response Syntax	387
Response Elements	388
Errors	388
Examples	389
See Also	391
DescribeMatchmaking	392
Request Syntax	392
Request Parameters	392
Response Syntax	393
Response Elements	394
Errors	394
See Also	395
DescribeMatchmakingConfigurations	396
Request Syntax	396
Request Parameters	396
Response Syntax	398
Response Elements	398
Errors	399
See Also	399
DescribeMatchmakingRuleSets	401
Request Syntax	401
Request Parameters	401
Response Syntax	402
Response Elements	403
Errors	403
See Also	404
DescribePlayerSessions	405
Request Syntax	405

Request Parameters	405
Response Syntax	407
Response Elements	408
Errors	408
See Also	409
DescribeRuntimeConfiguration	411
Request Syntax	411
Request Parameters	411
Response Syntax	412
Response Elements	412
Errors	412
Examples	413
See Also	414
DescribeScalingPolicies	416
Request Syntax	416
Request Parameters	416
Response Syntax	418
Response Elements	419
Errors	419
Examples	420
See Also	422
DescribeScript	423
Request Syntax	423
Request Parameters	423
Response Syntax	424
Response Elements	424
Errors	424
Examples	425
See Also	426
DescribeVpcPeeringAuthorizations	427
Response Syntax	427
Response Elements	427
Errors	427
See Also	428
DescribeVpcPeeringConnections	429
Request Syntax	429

Request Parameters	429
Response Syntax	430
Response Elements	430
Errors	430
See Also	431
GetComputeAccess	433
Request Syntax	433
Request Parameters	433
Response Syntax	434
Response Elements	435
Errors	437
See Also	437
GetComputeAuthToken	439
Request Syntax	439
Request Parameters	440
Response Syntax	440
Response Elements	441
Errors	442
See Also	443
GetGameSessionLogUrl	445
Request Syntax	445
Request Parameters	445
Response Syntax	446
Response Elements	446
Errors	446
See Also	447
GetInstanceAccess	448
Request Syntax	448
Request Parameters	449
Response Syntax	450
Response Elements	450
Errors	450
Examples	451
See Also	453
GetPlayerConnectionDetails	455
Request Syntax	455

Request Parameters	455
Response Syntax	456
Response Elements	457
Errors	457
Examples	458
See Also	459
ListAliases	461
Request Syntax	461
Request Parameters	461
Response Syntax	463
Response Elements	463
Errors	464
See Also	464
ListBuilds	466
Request Syntax	466
Request Parameters	466
Response Syntax	468
Response Elements	468
Errors	468
Examples	469
See Also	471
ListCompute	472
Request Syntax	472
Request Parameters	473
Response Syntax	475
Response Elements	475
Errors	476
See Also	477
ListContainerFleets	478
Request Syntax	478
Request Parameters	478
Response Syntax	480
Response Elements	481
Errors	481
Examples	482
See Also	484

ListContainerGroupDefinitions	485
Request Syntax	485
Request Parameters	485
Response Syntax	486
Response Elements	489
Errors	489
Examples	490
See Also	492
ListContainerGroupDefinitionVersions	493
Request Syntax	493
Request Parameters	493
Response Syntax	495
Response Elements	497
Errors	497
Examples	498
See Also	500
ListFleetDeployments	502
Request Syntax	502
Request Parameters	502
Response Syntax	503
Response Elements	504
Errors	504
Examples	505
See Also	506
ListFleets	508
Request Syntax	508
Request Parameters	508
Response Syntax	510
Response Elements	510
Errors	510
Examples	511
See Also	512
ListGameServerGroups	514
Request Syntax	514
Request Parameters	514
Response Syntax	515

Response Elements	515
Errors	516
See Also	516
ListGameServers	518
Request Syntax	518
Request Parameters	518
Response Syntax	519
Response Elements	520
Errors	520
See Also	521
ListLocations	522
Request Syntax	522
Request Parameters	523
Response Syntax	524
Response Elements	524
Errors	525
See Also	525
ListScripts	527
Request Syntax	527
Request Parameters	527
Response Syntax	528
Response Elements	528
Errors	529
Examples	530
See Also	531
ListTagsForResource	532
Request Syntax	532
Request Parameters	532
Response Syntax	533
Response Elements	533
Errors	533
See Also	534
PutScalingPolicy	536
Request Syntax	537
Request Parameters	538
Response Syntax	541

Response Elements	541
Errors	542
Examples	542
See Also	545
RegisterCompute	547
Request Syntax	547
Request Parameters	548
Response Syntax	549
Response Elements	550
Errors	550
Examples	551
See Also	552
RegisterGameServer	553
Request Syntax	553
Request Parameters	553
Response Syntax	555
Response Elements	555
Errors	556
Examples	557
See Also	558
RequestUploadCredentials	559
Request Syntax	559
Request Parameters	559
Response Syntax	560
Response Elements	560
Errors	560
Examples	561
See Also	562
ResolveAlias	563
Request Syntax	563
Request Parameters	563
Response Syntax	564
Response Elements	564
Errors	564
See Also	565
ResumeGameServerGroup	567

Request Syntax	567
Request Parameters	567
Response Syntax	568
Response Elements	569
Errors	569
Examples	570
See Also	571
SearchGameSessions	572
Request Syntax	574
Request Parameters	574
Response Syntax	577
Response Elements	578
Errors	578
Examples	580
See Also	585
StartFleetActions	586
Request Syntax	586
Request Parameters	586
Response Syntax	587
Response Elements	588
Errors	588
Examples	589
See Also	590
StartGameSessionPlacement	591
Request Syntax	592
Request Parameters	593
Response Syntax	596
Response Elements	597
Errors	598
Examples	598
See Also	601
StartMatchBackfill	602
Request Syntax	603
Request Parameters	603
Response Syntax	605
Response Elements	606

Errors	607
See Also	607
StartMatchmaking	609
Request Syntax	609
Request Parameters	610
Response Syntax	611
Response Elements	612
Errors	613
See Also	613
StopFleetActions	615
Request Syntax	615
Request Parameters	615
Response Syntax	617
Response Elements	617
Errors	617
Examples	618
See Also	619
StopGameSessionPlacement	620
Request Syntax	620
Request Parameters	620
Response Syntax	621
Response Elements	622
Errors	622
See Also	623
StopMatchmaking	624
Request Syntax	624
Request Parameters	624
Response Elements	625
Errors	625
See Also	626
SuspendGameServerGroup	627
Request Syntax	627
Request Parameters	627
Response Syntax	628
Response Elements	629
Errors	629

Examples	630
See Also	631
TagResource	632
Request Syntax	632
Request Parameters	632
Response Elements	633
Errors	633
See Also	634
TerminateGameSession	636
Request Syntax	637
Request Parameters	637
Response Syntax	638
Response Elements	639
Errors	639
Examples	641
See Also	642
UntagResource	643
Request Syntax	643
Request Parameters	643
Response Elements	644
Errors	644
See Also	645
UpdateAlias	646
Request Syntax	646
Request Parameters	646
Response Syntax	647
Response Elements	648
Errors	648
See Also	649
UpdateBuild	650
Request Syntax	650
Request Parameters	650
Response Syntax	651
Response Elements	651
Errors	652
Examples	652

See Also	653
UpdateContainerFleet	655
Request Syntax	656
Request Parameters	657
Response Syntax	661
Response Elements	662
Errors	663
Examples	664
See Also	665
UpdateContainerGroupDefinition	667
Request Syntax	668
Request Parameters	670
Response Syntax	672
Response Elements	675
Errors	675
Examples	676
See Also	678
UpdateFleetAttributes	679
Request Syntax	679
Request Parameters	680
Response Syntax	682
Response Elements	682
Errors	682
See Also	684
UpdateFleetCapacity	685
Request Syntax	686
Request Parameters	686
Response Syntax	688
Response Elements	688
Errors	689
Examples	691
See Also	693
UpdateFleetPortSettings	694
Request Syntax	694
Request Parameters	695
Response Syntax	696

Response Elements	696
Errors	696
Examples	698
See Also	699
UpdateGameServer	700
Request Syntax	700
Request Parameters	701
Response Syntax	702
Response Elements	703
Errors	703
Examples	704
See Also	705
UpdateGameServerGroup	706
Request Syntax	706
Request Parameters	707
Response Syntax	709
Response Elements	709
Errors	710
Examples	710
See Also	713
UpdateGameSession	714
Request Syntax	714
Request Parameters	714
Response Syntax	716
Response Elements	717
Errors	717
See Also	719
UpdateGameSessionQueue	720
Request Syntax	720
Request Parameters	720
Response Syntax	723
Response Elements	724
Errors	724
See Also	725
UpdateMatchmakingConfiguration	726
Request Syntax	726

Request Parameters	726
Response Syntax	731
Response Elements	732
Errors	732
See Also	733
UpdateRuntimeConfiguration	734
Request Syntax	734
Request Parameters	735
Response Syntax	735
Response Elements	736
Errors	736
See Also	737
UpdateScript	739
Request Syntax	739
Request Parameters	739
Response Syntax	741
Response Elements	742
Errors	742
Examples	743
See Also	745
ValidateMatchmakingRuleSet	746
Request Syntax	746
Request Parameters	746
Response Syntax	747
Response Elements	747
Errors	747
See Also	748
Data Types	749
Alias	753
Contents	753
See Also	754
AnywhereConfiguration	756
Contents	756
See Also	756
AttributeValue	757
Contents	757

See Also	758
AwsCredentials	759
Contents	759
See Also	760
Build	761
Contents	761
See Also	763
CertificateConfiguration	765
Contents	765
See Also	765
ClaimFilterOption	766
Contents	766
See Also	766
Compute	767
Contents	767
See Also	774
ConnectionPortRange	775
Contents	775
See Also	775
ContainerAttribute	777
Contents	777
See Also	777
ContainerDependency	779
Contents	779
See Also	780
ContainerEnvironment	781
Contents	781
See Also	781
ContainerFleet	783
Contents	783
See Also	789
ContainerFleetLocationAttributes	790
Contents	790
See Also	791
ContainerGroupDefinition	792
Contents	792

See Also	796
ContainerGroupPortMapping	797
Contents	797
See Also	798
ContainerHealthCheck	799
Contents	799
See Also	800
ContainerIdentifier	802
Contents	802
See Also	802
ContainerMountPoint	804
Contents	804
See Also	805
ContainerPortConfiguration	806
Contents	806
See Also	807
ContainerPortMapping	808
Contents	808
See Also	809
ContainerPortRange	810
Contents	810
See Also	811
DeploymentConfiguration	812
Contents	812
See Also	813
DeploymentDetails	814
Contents	814
See Also	814
DesiredPlayerSession	815
Contents	815
See Also	815
EC2InstanceCounts	817
Contents	817
See Also	818
EC2InstanceLimit	820
Contents	820

See Also	824
Event	825
Contents	825
See Also	831
FilterConfiguration	832
Contents	832
See Also	832
FleetAttributes	833
Contents	833
See Also	844
FleetCapacity	845
Contents	845
See Also	849
FleetDeployment	851
Contents	851
See Also	853
FleetUtilization	855
Contents	855
See Also	857
GameProperty	858
Contents	858
See Also	859
GameServer	860
Contents	860
See Also	863
GameServerContainerDefinition	864
Contents	864
See Also	866
GameServerContainerDefinitionInput	868
Contents	868
See Also	871
GameServerContainerGroupCounts	872
Contents	872
See Also	873
GameServerGroup	874
Contents	874

See Also	878
GameServerGroupAutoScalingPolicy	879
Contents	879
See Also	880
GameServerInstance	881
Contents	881
See Also	882
GameSession	883
Contents	883
See Also	889
GameSessionConnectionInfo	890
Contents	890
See Also	892
GameSessionCreationLimitPolicy	893
Contents	893
See Also	894
GameSessionDetail	895
Contents	895
See Also	895
GameSessionPlacement	897
Contents	897
See Also	903
GameSessionQueue	904
Contents	904
See Also	906
GameSessionQueueDestination	908
Contents	908
See Also	908
Instance	909
Contents	909
See Also	915
InstanceAccess	916
Contents	916
See Also	917
InstanceCredentials	918
Contents	918

See Also	918
InstanceDefinition	920
Contents	920
See Also	921
IpPermission	922
Contents	922
See Also	923
LaunchTemplateSpecification	924
Contents	924
See Also	925
LinuxCapabilities	926
Contents	926
See Also	927
LocationalDeployment	928
Contents	928
See Also	928
LocationAttributes	930
Contents	930
See Also	931
LocationConfiguration	932
Contents	932
See Also	932
LocationModel	933
Contents	933
See Also	934
LocationState	935
Contents	935
See Also	936
LogConfiguration	937
Contents	937
See Also	938
ManagedCapacityConfiguration	939
Contents	939
See Also	940
MatchedPlayerSession	941
Contents	941

See Also	941
MatchmakingConfiguration	943
Contents	943
See Also	948
MatchmakingRuleSet	949
Contents	949
See Also	951
MatchmakingTicket	952
Contents	952
See Also	955
PingBeacon	956
Contents	956
See Also	956
PlacedPlayerSession	957
Contents	957
See Also	957
Player	959
Contents	959
See Also	960
PlayerConnectionDetail	962
Contents	962
See Also	963
PlayerConnectionEndpoint	964
Contents	964
See Also	964
PlayerGatewayConfiguration	966
Contents	966
See Also	966
PlayerLatency	967
Contents	967
See Also	968
PlayerLatencyPolicy	969
Contents	969
See Also	969
PlayerSession	971
Contents	971

See Also	974
PriorityConfiguration	976
Contents	976
See Also	977
PriorityConfigurationOverride	978
Contents	978
See Also	979
ResourceCreationLimitPolicy	980
Contents	980
See Also	981
RoutingStrategy	982
Contents	982
See Also	983
RuntimeConfiguration	984
Contents	984
See Also	985
S3Location	986
Contents	986
See Also	987
ScalingPolicy	988
Contents	988
See Also	992
Script	994
Contents	994
See Also	996
ServerProcess	997
Contents	997
See Also	998
SupportContainerDefinition	999
Contents	999
See Also	1002
SupportContainerDefinitionInput	1003
Contents	1003
See Also	1007
Tag	1008
Contents	1008

See Also	1009
TargetConfiguration	1010
Contents	1010
See Also	1010
TargetTrackingConfiguration	1011
Contents	1011
See Also	1011
UDPEndpoint	1012
Contents	1012
See Also	1013
VpcPeeringAuthorization	1014
Contents	1014
See Also	1015
VpcPeeringConnection	1016
Contents	1016
See Also	1018
VpcPeeringConnectionStatus	1019
Contents	1019
See Also	1019
Common Parameters	1021
Common Error Types	1024

Welcome

Amazon GameLift Servers provides solutions for hosting session-based multiplayer game servers in the cloud, including tools for deploying, operating, and scaling game servers. Built on AWS global computing infrastructure, GameLift helps you deliver high-performance, high-reliability, low-cost game servers while dynamically scaling your resource usage to meet player demand.

About Amazon GameLift Servers solutions

Get more information on these Amazon GameLift Servers solutions in the [Amazon GameLift Servers Developer Guide](#).

- Amazon GameLift Servers managed hosting -- Amazon GameLift Servers offers a fully managed service to set up and maintain computing machines for hosting, manage game session and player session life cycle, and handle security, storage, and performance tracking. You can use automatic scaling tools to balance player demand and hosting costs, configure your game session management to minimize player latency, and add FlexMatch for matchmaking.
- Managed hosting with Amazon GameLift Servers Realtime -- With Amazon GameLift Servers Realtime, you can quickly configure and set up ready-to-go game servers for your game. Amazon GameLift Servers Realtime provides a game server framework with core Amazon GameLift Servers infrastructure already built in. Then use the full range of Amazon GameLift Servers managed hosting features, including FlexMatch, for your game.
- Amazon GameLift Servers FleetIQ -- Use Amazon GameLift Servers FleetIQ as a standalone service while hosting your games using EC2 instances and Auto Scaling groups. Amazon GameLift Servers FleetIQ provides optimizations for game hosting, including boosting the viability of low-cost Spot Instances gaming. For a complete solution, pair the Amazon GameLift Servers FleetIQ and FlexMatch standalone services.
- Amazon GameLift Servers FlexMatch -- Add matchmaking to your game hosting solution. FlexMatch is a customizable matchmaking service for multiplayer games. Use FlexMatch as integrated with Amazon GameLift Servers managed hosting or incorporate FlexMatch as a standalone service into your own hosting solution.

About this API Reference

This reference guide describes the low-level service API for Amazon GameLift Servers. With each topic in this guide, you can find links to language-specific SDK guides and the AWS CLI reference. Useful links:

- [Amazon GameLift Servers API operations listed by tasks](#)
- [Amazon GameLift Servers tools and resources](#)

This document was last published on August 14, 2026.

Actions

The following actions are supported:

- [AcceptMatch](#)
- [ClaimGameServer](#)
- [CreateAlias](#)
- [CreateBuild](#)
- [CreateContainerFleet](#)
- [CreateContainerGroupDefinition](#)
- [CreateFleet](#)
- [CreateFleetLocations](#)
- [CreateGameServerGroup](#)
- [CreateGameSession](#)
- [CreateGameSessionQueue](#)
- [CreateLocation](#)
- [CreateMatchmakingConfiguration](#)
- [CreateMatchmakingRuleSet](#)
- [CreatePlayerSession](#)
- [CreatePlayerSessions](#)
- [CreateScript](#)
- [CreateVpcPeeringAuthorization](#)
- [CreateVpcPeeringConnection](#)
- [DeleteAlias](#)
- [DeleteBuild](#)
- [DeleteContainerFleet](#)
- [DeleteContainerGroupDefinition](#)
- [DeleteFleet](#)
- [DeleteFleetLocations](#)
- [DeleteGameServerGroup](#)
- [DeleteGameSessionQueue](#)

- [DeleteLocation](#)
- [DeleteMatchmakingConfiguration](#)
- [DeleteMatchmakingRuleSet](#)
- [DeleteScalingPolicy](#)
- [DeleteScript](#)
- [DeleteVpcPeeringAuthorization](#)
- [DeleteVpcPeeringConnection](#)
- [DeregisterCompute](#)
- [DeregisterGameServer](#)
- [DescribeAlias](#)
- [DescribeBuild](#)
- [DescribeCompute](#)
- [DescribeContainerFleet](#)
- [DescribeContainerGroupDefinition](#)
- [DescribeContainerGroupPortMappings](#)
- [DescribeEC2InstanceLimits](#)
- [DescribeFleetAttributes](#)
- [DescribeFleetCapacity](#)
- [DescribeFleetDeployment](#)
- [DescribeFleetEvents](#)
- [DescribeFleetLocationAttributes](#)
- [DescribeFleetLocationCapacity](#)
- [DescribeFleetLocationUtilization](#)
- [DescribeFleetPortSettings](#)
- [DescribeFleetUtilization](#)
- [DescribeGameServer](#)
- [DescribeGameServerGroup](#)
- [DescribeGameServerInstances](#)
- [DescribeGameSessionDetails](#)
- [DescribeGameSessionPlacement](#)

- [DescribeGameSessionQueues](#)
- [DescribeGameSessions](#)
- [DescribeInstances](#)
- [DescribeMatchmaking](#)
- [DescribeMatchmakingConfigurations](#)
- [DescribeMatchmakingRuleSets](#)
- [DescribePlayerSessions](#)
- [DescribeRuntimeConfiguration](#)
- [DescribeScalingPolicies](#)
- [DescribeScript](#)
- [DescribeVpcPeeringAuthorizations](#)
- [DescribeVpcPeeringConnections](#)
- [GetComputeAccess](#)
- [GetComputeAuthToken](#)
- [GetGameSessionLogUrl](#)
- [GetInstanceAccess](#)
- [GetPlayerConnectionDetails](#)
- [ListAliases](#)
- [ListBuilds](#)
- [ListCompute](#)
- [ListContainerFleets](#)
- [ListContainerGroupDefinitions](#)
- [ListContainerGroupDefinitionVersions](#)
- [ListFleetDeployments](#)
- [ListFleets](#)
- [ListGameServerGroups](#)
- [ListGameServers](#)
- [ListLocations](#)
- [ListScripts](#)
- [ListTagsForResource](#)

- [PutScalingPolicy](#)
- [RegisterCompute](#)
- [RegisterGameServer](#)
- [RequestUploadCredentials](#)
- [ResolveAlias](#)
- [ResumeGameServerGroup](#)
- [SearchGameSessions](#)
- [StartFleetActions](#)
- [StartGameSessionPlacement](#)
- [StartMatchBackfill](#)
- [StartMatchmaking](#)
- [StopFleetActions](#)
- [StopGameSessionPlacement](#)
- [StopMatchmaking](#)
- [SuspendGameServerGroup](#)
- [TagResource](#)
- [TerminateGameSession](#)
- [UntagResource](#)
- [UpdateAlias](#)
- [UpdateBuild](#)
- [UpdateContainerFleet](#)
- [UpdateContainerGroupDefinition](#)
- [UpdateFleetAttributes](#)
- [UpdateFleetCapacity](#)
- [UpdateFleetPortSettings](#)
- [UpdateGameServer](#)
- [UpdateGameServerGroup](#)
- [UpdateGameSession](#)
- [UpdateGameSessionQueue](#)
- [UpdateMatchmakingConfiguration](#)

- [UpdateRuntimeConfiguration](#)
- [UpdateScript](#)
- [ValidateMatchmakingRuleSet](#)

AcceptMatch

This API works with the following fleet types: EC2, Anywhere, Container

Registers a player's acceptance or rejection of a proposed FlexMatch match. A matchmaking configuration may require player acceptance; if so, then matches built with that configuration cannot be completed unless all players accept the proposed match within a specified time limit.

When FlexMatch builds a match, all the matchmaking tickets involved in the proposed match are placed into status `REQUIRES_ACCEPTANCE`. This is a trigger for your game to get acceptance from all players in each ticket. Calls to this action are only valid for tickets that are in this status; calls for tickets not in this status result in an error.

To register acceptance, specify the ticket ID, one or more players, and an acceptance response. When all players have accepted, Amazon GameLift Servers advances the matchmaking tickets to status `PLACING`, and attempts to create a new game session for the match.

If any player rejects the match, or if acceptances are not received before a specified timeout, the proposed match is dropped. Each matchmaking ticket in the failed match is handled as follows:

- If the ticket has one or more players who rejected the match or failed to respond, the ticket status is set `CANCELLED` and processing is terminated.
- If all players in the ticket accepted the match, the ticket status is returned to `SEARCHING` to find a new match.

Learn more

[Add FlexMatch to a game client](#)

[FlexMatch events](#) (reference)

Request Syntax

```
{
  "AcceptanceType": "string",
  "PlayerIds": [ "string" ],
  "TicketId": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

AcceptanceType

Player response to the proposed match.

Type: String

Valid Values: ACCEPT | REJECT

Required: Yes

PlayerIds

A unique identifier for a player delivering the response. This parameter can include one or multiple player IDs.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

TicketId

A unique identifier for a matchmaking ticket. The ticket must be in status `REQUIRES_ACCEPTANCE`; otherwise this request will fail.

Type: String

Length Constraints: Maximum length of 128.

Pattern: `[a-zA-Z0-9-\.\ ]*`

Required: Yes

Response Elements

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

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

ClaimGameServer

This API works with the following fleet types: EC2 (FleetIQ)

Locates an available game server and temporarily reserves it to host gameplay and players. This operation is called from a game client or client service (such as a matchmaker) to request hosting resources for a new game session. In response, Amazon GameLift Servers FleetIQ locates an available game server, places it in CLAIMED status for 60 seconds, and returns connection information that players can use to connect to the game server.

To claim a game server, identify a game server group. You can also specify a game server ID, although this approach bypasses Amazon GameLift Servers FleetIQ placement optimization. Optionally, include game data to pass to the game server at the start of a game session, such as a game map or player information. Add filter options to further restrict how a game server is chosen, such as only allowing game servers on ACTIVE instances to be claimed.

When a game server is successfully claimed, connection information is returned. A claimed game server's utilization status remains AVAILABLE while the claim status is set to CLAIMED for up to 60 seconds. This time period gives the game server time to update its status to UTILIZED after players join. If the game server's status is not updated within 60 seconds, the game server reverts to unclaimed status and is available to be claimed by another request. The claim time period is a fixed value and is not configurable.

If you try to claim a specific game server, this request will fail in the following cases:

- If the game server utilization status is UTILIZED.
- If the game server claim status is CLAIMED.
- If the game server is running on an instance in DRAINING status and the provided filter option does not allow placing on DRAINING instances.

Learn more

[Amazon GameLift Servers FleetIQ Guide](#)

Request Syntax

```
{  
  "FilterOption": {
```

```
    "InstanceStatuses": [ "string" ]
  },
  "GameServerData": "string",
  "GameServerGroupName": "string",
  "GameServerId": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

GameServerGroupName

A unique identifier for the game server group where the game server is running. If you are not specifying a game server to claim, this value identifies where you want Amazon GameLift Servers FleetIQ to look for an available game server to claim.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]|^arn:.*:gameservergroup\[a-zA-Z0-9-\.]`

Required: Yes

FilterOption

Object that restricts how a claimed game server is chosen.

Type: [ClaimFilterOption](#) object

Required: No

GameServerData

A set of custom game server properties, formatted as a single string value. This data is passed to a game client or service when it requests information on game servers.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `.*\S.*`

Required: No

GameServerId

A custom string that uniquely identifies the game server to claim. If this parameter is left empty, Amazon GameLift Servers FleetIQ searches for an available game server in the specified game server group.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `[a-zA-Z0-9-\.]`⁺

Required: No

Response Syntax

```
{
  "GameServer": {
    "ClaimStatus": "string",
    "ConnectionInfo": "string",
    "GameServerData": "string",
    "GameServerGroupArn": "string",
    "GameServerGroupName": "string",
    "GameServerId": "string",
    "InstanceId": "string",
    "LastClaimTime": number,
    "LastHealthCheckTime": number,
    "RegistrationTime": number,
    "UtilizationStatus": "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.

GameServer

Object that describes the newly claimed game server.

Type: [GameServer](#) object

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

OutOfCapacityException

The specified game server group has no available game servers to fulfill a `ClaimGameServer` request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Reserve a game server for a new game session

This example reserves an available game server to host a new game session. Because the request does not specify a game server ID, Amazon GameLift Servers selects an available game server with optimal placement.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "GameServerGroupName": "MegaFrogServers_NA"
}
```

CLI command:

```
aws gamelift claim-game-server \
  --game-server-group-name MegaFrogServers_NA
```

Sample Response

```
{
  "GameServer": {
    "ClaimStatus": "CLAIMED",
    "ConnectionInfo": "192.0.2.0.80",
    "GameServerData": "",
    "GameServerGroupArn": "arn:aws:gamelift:us-west-2::GameServerGroup/MegaFrogServers_NA",
    "GameServerGroupName": "MegaFrogServers_NA",
    "GameServerId": "mega-frog-game-12345678",
    "InstanceId": "i-1234567890abcdef0",
  }
}
```

```
"LastClaimTime": 1580218197.293,  
"LastHealthCheckTime": 1580218197.293,  
"RegistrationTime": 1580218197.293,  
"UtilizationStatus": "AVAILABLE"  
}  
}
```

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](#)

CreateAlias

This API works with the following fleet types: EC2, Anywhere, Container

Creates an alias for a fleet. In most situations, you can use an alias ID in place of a fleet ID. An alias provides a level of abstraction for a fleet that is useful when redirecting player traffic from one fleet to another, such as when updating your game build.

Amazon GameLift Servers supports two types of routing strategies for aliases: simple and terminal. A simple alias points to an active fleet. A terminal alias is used to display messaging or link to a URL instead of routing players to an active fleet. For example, you might use a terminal alias when a game version is no longer supported and you want to direct players to an upgrade site.

To create a fleet alias, specify an alias name, routing strategy, and optional description. Each simple alias can point to only one fleet, but a fleet can have multiple aliases. If successful, a new alias record is returned, including an alias ID and an ARN. You can reassign an alias to another fleet by calling `UpdateAlias`.

Related actions

[All APIs by task](#)

Request Syntax

```
{
  "Description": "string",
  "Name": "string",
  "RoutingStrategy": {
    "FleetId": "string",
    "Message": "string",
    "Type": "string"
  },
  "Tags": [
    {
      "Key": "string",
      "Value": "string"
    }
  ]
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Name

A descriptive label that is associated with an alias. Alias names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `.*\S.*`

Required: Yes

RoutingStrategy

The routing configuration, including routing type and fleet target, for the alias.

Type: [RoutingStrategy](#) object

Required: Yes

Description

A human-readable description of the alias.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Tags

A list of labels to assign to the new alias resource. Tags are developer-defined key-value pairs. Tagging AWS resources are useful for resource management, access management and cost allocation. For more information, see [Tagging AWS Resources](#) in the *AWS General Reference*.

Type: Array of [Tag](#) objects

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

Required: No

Response Syntax

```
{
  "Alias": {
    "AliasArn": "string",
    "AliasId": "string",
    "CreationTime": number,
    "Description": "string",
    "LastUpdatedTime": number,
    "Name": "string",
    "RoutingStrategy": {
      "FleetId": "string",
      "Message": "string",
      "Type": "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.

[Alias](#)

The newly created alias resource.

Type: [Alias](#) object

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

CreateBuild

This API works with the following fleet types: EC2, Anywhere

Creates a new Amazon GameLift Servers build resource for your game server binary files. Combine game server binaries into a zip file for use with Amazon GameLift Servers.

Important

When setting up a new game build for Amazon GameLift Servers, we recommend using the AWS CLI command [upload-build](#). This helper command combines two tasks: (1) it uploads your build files from a file directory to an Amazon GameLift Servers Amazon S3 location, and (2) it creates a new build resource.

You can use the CreateBuild operation in the following scenarios:

- Create a new game build with build files that are in an Amazon S3 location under an AWS account that you control. To use this option, you give Amazon GameLift Servers access to the Amazon S3 bucket. With permissions in place, specify a build name, operating system, and the Amazon S3 storage location of your game build.
- Upload your build files to a Amazon GameLift Servers Amazon S3 location. To use this option, specify a build name and operating system. This operation creates a new build resource and also returns an Amazon S3 location with temporary access credentials. Use the credentials to manually upload your build files to the specified Amazon S3 location. For more information, see [Uploading Objects](#) in the *Amazon S3 Developer Guide*. After you upload build files to the Amazon GameLift Servers Amazon S3 location, you can't update them.

If successful, this operation creates a new build resource with a unique build ID and places it in INITIALIZED status. A build must be in READY status before you can create fleets with it.

Learn more

[Uploading Your Game](#)

[Create a Build with Files in Amazon S3](#)

[All APIs by task](#)

Request Syntax

```
{
  "Name": "string",
  "OperatingSystem": "string",
  "ServerSdkVersion": "string",
  "StorageLocation": {
    "Bucket": "string",
    "Key": "string",
    "ObjectVersion": "string",
    "RoleArn": "string"
  },
  "Tags": [
    {
      "Key": "string",
      "Value": "string"
    }
  ],
  "Version": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Name

A descriptive label that is associated with a build. Build names do not need to be unique. You can change this value later.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

OperatingSystem

The operating system that your game server binaries run on. This value determines the type of fleet resources that you use for this build. If your game build contains multiple executables, they all must run on the same operating system. You must specify a valid operating system in this request. There is no default value. You can't change a build's operating system later.

Note

Amazon Linux 2 (AL2) will reach end of support on 6/30/2026. See more details in the [Amazon Linux 2 FAQs](#). For game servers that are hosted on AL2 and use server SDK version 4.x for Amazon GameLift Servers, first update the game server build to server SDK 5.x, and then deploy to AL2023 instances. See [Migrate to server SDK version 5](#).

Note

Windows Server 2016 will reach end of support on 1/12/2027. For game servers that are hosted on Windows Server 2016 and use server SDK version 4.x for Amazon GameLift Servers, first update the game server build to server SDK 5.x, and then deploy to Windows Server 2022 instances. See [Migrate to server SDK version 5](#).

Type: String

Valid Values: WINDOWS_2012 | AMAZON_LINUX | AMAZON_LINUX_2 | WINDOWS_2016 | AMAZON_LINUX_2023 | WINDOWS_2022

Required: No

ServerSdkVersion

A server SDK version you used when integrating your game server build with Amazon GameLift Servers. For more information see [Integrate games with custom game servers](#). By default Amazon GameLift Servers sets this value to 4.0.2.

Type: String

Length Constraints: Maximum length of 128.

Pattern: ^\d+\.\d+\.\d+\$

Required: No

StorageLocation

Information indicating where your game build files are stored. Use this parameter only when creating a build with files stored in an Amazon S3 bucket that you own. The storage location must specify an Amazon S3 bucket name and key. The location must also specify a role ARN that you set up to allow Amazon GameLift Servers to access your Amazon S3 bucket. The S3 bucket and your new build must be in the same Region.

If a `StorageLocation` is specified, the size of your file can be found in your Amazon S3 bucket. Amazon GameLift Servers will report a `SizeOnDisk` of 0.

Type: [S3Location](#) object

Required: No

Tags

A list of labels to assign to the new build resource. Tags are developer defined key-value pairs. Tagging AWS resources are useful for resource management, access management and cost allocation. For more information, see [Tagging AWS Resources](#) in the *AWS General Reference*. Once the resource is created, you can use [TagResource](#), [UntagResource](#), and [ListTagsForResource](#) to add, remove, and view tags. The maximum tag limit may be lower than stated. See the AWS General Reference for actual tagging limits.

Type: Array of [Tag](#) objects

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

Required: No

Version

Version information that is associated with a build or script. Version strings do not need to be unique. You can change this value later.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "Build": {
    "BuildArn": "string",
    "BuildId": "string",
    "CreationTime": number,
    "Name": "string",
    "OperatingSystem": "string",
    "ServerSdkVersion": "string",
    "SizeOnDisk": number,
    "Status": "string",
    "Version": "string"
  },
  "StorageLocation": {
    "Bucket": "string",
    "Key": "string",
    "ObjectVersion": "string",
    "RoleArn": "string"
  },
  "UploadCredentials": {
    "AccessKeyId": "string",
    "SecretAccessKey": "string",
    "SessionToken": "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.

Build

The newly created build resource, including a unique build IDs and status.

Type: [Build](#) object

StorageLocation

Amazon S3 location for your game build file, including bucket name and key.

Type: [S3Location](#) object

UploadCredentials

This element is returned only when the operation is called without a storage location. It contains credentials to use when you are uploading a build file to an Amazon S3 bucket that is owned by Amazon GameLift Servers. Credentials have a limited life span. To refresh these credentials, call [RequestUploadCredentials](#).

Type: [AwsCredentials](#) object

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Create a build with files in your own S3 bucket

This example creates a custom game build resource. It uses zipped files that are stored in an S3 location in an AWS account that you control. This example assumes that you've already created an IAM role that gives Amazon GameLift Servers permission to access the S3 location. The request must specify a valid operating system value.

Sample Request

```
{
  "Name": "MegaFrogRaceServer.NA",
  "Version": "12345.678",
  "OperatingSystem": "WINDOWS_2022",
  "StorageLocation": {
    "Bucket": "MegaFrogRaceServer_NA_build_files",
    "Key": "MegaFrogRaceServer_build_123.zip",
    "RoleArn": "arn:aws:iam::111122223333:role/GameLiftAccess"
  }
}
```

Sample Response

```
{
  "Build": {
    "BuildArn": "arn:aws:gamelift:us-west-2::build/build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "BuildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "CreationTime": 1496708916.18,
    "Name": "MegaFrogRaceServer.NA",
    "OperatingSystem": "WINDOWS_2022",
    "SizeOnDisk": 0,
    "Status": "READY",
    "Version": "12345.678"
  },
  "StorageLocation": {
    "Bucket": "MegaFrogRaceServer_NA_build_files",
    "Key": "MegaFrogRaceServer_build_123.zip"
  }
}
```

```
}  
}
```

Create a game build resource for manually uploading files to Amazon GameLift Servers

This example creates a new build resource. It also gets a storage location and temporary credentials that allow you to manually upload your game build to the Amazon GameLift Servers location in Amazon S3. When you specify a storage location, Amazon GameLift Servers reports the `SizeOnDisk` as 0. You can find the actual size in Amazon S3. After you upload your build, Amazon GameLift Servers validates the build and updates the new build's status. The request must specify a valid operating system value.

Sample Request

```
{  
  "Name": "MegaFrogRaceServer.NA",  
  "Version": "12345.678",  
  "OperatingSystem": "AMAZON_LINUX_2023"  
}
```

Sample Response

```
{  
  "Build": {  
    "BuildArn": "arn:aws:gamelift:us-west-2::build/  
build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",  
    "BuildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",  
    "CreationTime": 1496708916.18,  
    "Name": "MegaFrogRaceServer.NA",  
    "OperatingSystem": "AMAZON_LINUX_2023",  
    "SizeOnDisk": 0,  
    "Status": "READY",  
    "Version": "12345.678"  
  },  
  "StorageLocation": {  
    "Bucket": "gamelift-builds-us-west-2",  
    "Key": "123456789012/build-1111aaaa-22bb-33cc-44dd-5555eeee66ff"  
  },  
  "UploadCredentials": {  
    "AccessKeyId": "AKIAIOSFODNN7EXAMPLE",  
    "SecretAccessKey": "wJalrXU3FJOQ8Hc7mB4pUbwNgt5qm4ZUJldjqAB76Zi4t6WOTD8oo3oImwujZ6",  
    "Expiration": 1496709216.18,  
    "SessionToken": "AQIDABEAAQ=="  
  }  
}
```

```
    "SecretAccessKey": "wJa1rXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY",  
    "SessionToken": "AgoGb3JpZ2luENz...EXAMPLETOKEN=="  
  }  
}
```

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](#)

CreateContainerFleet

This API works with the following fleet types: Container

Creates a managed fleet of Amazon Elastic Compute Cloud (Amazon EC2) instances to host your containerized game servers. Use this operation to define how to deploy a container architecture onto each fleet instance and configure fleet settings. You can create a container fleet in any AWS Regions that Amazon GameLift Servers supports for multi-location fleets. A container fleet can be deployed to a single location or multiple locations. Container fleets are deployed with Amazon Linux 2023 as the instance operating system.

Define the fleet's container architecture using container group definitions. Each fleet can have one of the following container group types:

- The game server container group runs your game server build and dependent software. Amazon GameLift Servers deploys one or more replicas of this container group to each fleet instance. The number of replicas depends on the computing capabilities of the fleet instance in use.
- An optional per-instance container group might be used to run other software that only needs to run once per instance, such as background services, logging, or test processes. One per-instance container group is deployed to each fleet instance.

Each container group can include the definition for one or more containers. A container definition specifies a container image that is stored in an Amazon Elastic Container Registry (Amazon ECR) public or private repository.

Request options

Use this operation to make the following types of requests. Most fleet settings have default values, so you can create a working fleet with a minimal configuration and default values, which you can customize later.

- Create a fleet with no container groups. You can configure a container fleet and then add container group definitions later. In this scenario, no fleet instances are deployed, and the fleet can't host game sessions until you add a game server container group definition. Provide the following required parameter values:
 - `FleetRoleArn`
- Create a fleet with a game server container group. Provide the following required parameter values:

- FleetRoleArn
- GameServerContainerGroupDefinitionName
- Create a fleet with a game server container group and a per-instance container group. Provide the following required parameter values:
 - FleetRoleArn
 - GameServerContainerGroupDefinitionName
 - PerInstanceContainerGroupDefinitionName

Results

If successful, this operation creates a new container fleet resource, places it in PENDING status, and initiates the [fleet creation workflow](#). For fleets with container groups, this workflow starts a fleet deployment and transitions the status to ACTIVE. Fleets without a container group are placed in CREATED status.

You can update most of the properties of a fleet, including container group definitions, and deploy the update across all fleet instances. Use [UpdateContainerFleet](#) to deploy a new game server version update across the container fleet.

Note

A managed fleet's runtime environment depends on the Amazon Machine Image (AMI) version it uses. When a new fleet is created, Amazon GameLift Servers assigns the latest available AMI version to the fleet, and all compute instances in that fleet are deployed with that version. To update the AMI version, you must create a new fleet. As a best practice, we recommend replacing your managed fleets every 30 days to maintain a secure and up-to-date runtime environment for your hosted game servers. For guidance, see [Security best practices for Amazon GameLift Servers](#).

Request Syntax

```
{  
  "BillingType": "string",  
  "Description": "string",  
  "FleetRoleArn": "string",  
  "GameServerContainerGroupDefinitionName": "string",
```



```

"GameServerContainerGroupsPerInstance": number,
"GameSessionCreationLimitPolicy": {
  "NewGameSessionsPerCreator": number,
  "PolicyPeriodInMinutes": number
},
"InstanceConnectionPortRange": {
  "FromPort": number,
  "ToPort": number
},
"InstanceInboundPermissions": [
  {
    "FromPort": number,
    "IpRange": "string",
    "Protocol": "string",
    "ToPort": number
  }
],
"InstanceType": "string",
"Locations": [
  {
    "Location": "string"
  }
],
"LogConfiguration": {
  "LogDestination": "string",
  "LogGroupArn": "string",
  "S3BucketName": "string"
},
"MetricGroups": [ "string" ],
"NewGameSessionProtectionPolicy": "string",
"PerInstanceContainerGroupDefinitionName": "string",
"PlayerGatewayMode": "string",
"Tags": [
  {
    "Key": "string",
    "Value": "string"
  }
]
}

```

Request Parameters

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

The request accepts the following data in JSON format.

 Note

In the following list, the required parameters are described first.

FleetRoleArn

The unique identifier for an AWS Identity and Access Management (IAM) role with permissions to run your containers on resources that are managed by Amazon GameLift Servers. Use an IAM service role with the `GameLiftContainerFleetPolicy` managed policy attached. For more information, see [Set up an IAM service role](#). You can't change this fleet property after the fleet is created.

IAM role ARN values use the following pattern: `arn:aws:iam::[AWS account]:role/[role name]`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^arn:.*:role\[\\w+=,.\@-\]+`

Required: Yes

BillingType

Indicates whether to use On-Demand or Spot instances for this fleet. Learn more about when to use [On-Demand versus Spot Instances](#). This fleet property can't be changed after the fleet is created.

By default, this property is set to `ON_DEMAND`.

You can't update this fleet property later.

Type: String

Valid Values: `ON_DEMAND` | `SPOT`

Required: No

Description

A meaningful description of the container fleet.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

GameServerContainerGroupDefinitionName

A container group definition resource that describes how to deploy containers with your game server build and support software onto each fleet instance. You can specify the container group definition's name to use the latest version. Alternatively, provide an ARN value with a specific version number.

Create a container group definition by calling [CreateContainerGroupDefinition](#). This operation creates a [ContainerGroupDefinition](#) resource.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-zA-Z0-9\-\_]+$|^arn:.*:containergroupdefinition\[a-zA-Z0-9\-\_]+(:[0-9]+)?$`

Required: No

GameServerContainerGroupsPerInstance

The number of times to replicate the game server container group on each fleet instance.

By default, Amazon GameLift Servers calculates the maximum number of game server container groups that can fit on each instance. This calculation is based on the CPU and memory resources of the fleet's instance type). To use the calculated maximum, don't set this parameter. If you set this number manually, Amazon GameLift Servers uses your value as long as it's less than the calculated maximum.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 5000.

Required: No

GameSessionCreationLimitPolicy

A policy that limits the number of game sessions that each individual player can create on instances in this fleet. The limit applies for a specified span of time.

Type: [GameSessionCreationLimitPolicy](#) object

Required: No

InstanceConnectionPortRange

The set of port numbers to open on each fleet instance. A fleet's connection ports map to container ports that are configured in the fleet's container group definitions.

By default, Amazon GameLift Servers calculates an optimal port range based on your fleet configuration. To use the calculated range, don't set this parameter. The values are:

- Port range: 4192 to a number calculated based on your fleet configuration. Amazon GameLift Servers uses the following formula: $4192 + [\text{\# of game server container groups per fleet instance}] * [\text{\# of container ports in the game server container group definition}] + [\text{\# of container ports in the per instance container group definition}]$

You can also choose to manually set this parameter. When manually setting this parameter, you must use port numbers that match the fleet's inbound permissions port range.

Note

If you set values manually, Amazon GameLift Servers no longer calculates a port range for you, even if you later remove the manual settings.

The port range must not overlap with the Amazon GameLift Servers reserved port range 4092–4191. This range is reserved for internal Amazon GameLift Servers services.

Type: [ConnectionPortRange](#) object

Required: No

InstanceInboundPermissions

The IP address ranges and port settings that allow inbound traffic to access game server processes and other processes on this fleet. As a best practice, when remotely accessing a fleet

instance, we recommend opening ports only when you need them and closing them when you're finished.

By default, Amazon GameLift Servers calculates an optimal port range based on your fleet configuration. To use the calculated range, don't set this parameter. The values are:

- Protocol: UDP
- Port range: 4192 to a number calculated based on your fleet configuration. Amazon GameLift Servers uses the following formula: $4192 + [\text{\# of game server container groups per fleet instance}] * [\text{\# of container ports in the game server container group definition}] + [\text{\# of container ports in the per instance container group definition}]$

You can also choose to manually set this parameter. When manually setting this parameter, you must use port numbers that match the fleet's connection port range.

 Note

If you set values manually, Amazon GameLift Servers no longer calculates a port range for you, even if you later remove the manual settings.

The port range must not overlap with the Amazon GameLift Servers reserved port range 4092-4191. This range is reserved for internal Amazon GameLift Servers services.

Type: Array of [IpPermission](#) objects

Array Members: Maximum number of 50 items.

Required: No

InstanceType

The Amazon EC2 instance type to use for all instances in the fleet. For multi-location fleets, the instance type must be available in the home region and all remote locations. Instance type determines the computing resources and processing power that's available to host your game servers. This includes including CPU, memory, storage, and networking capacity.

By default, Amazon GameLift Servers uses the `c5.large` instance type. If this instance type does not have sufficient resources for your container groups, you can choose a different

instance type that better fits your needs. See [Amazon Elastic Compute Cloud Instance Types](#) for detailed descriptions of Amazon EC2 instance types.

You can't update this fleet property later.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Locations

A set of locations to deploy container fleet instances to. You can add any AWS Region or Local Zone that's supported by Amazon GameLift Servers. Provide a list of one or more AWS Region codes, such as `us-west-2`, or Local Zone names. Also include the fleet's home Region, which is the AWS Region where the fleet is created. For a list of supported Regions and Local Zones, see [Amazon GameLift Servers service locations](#) for managed hosting.

Type: Array of [LocationConfiguration](#) objects

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

Required: No

LogConfiguration

A method for collecting container logs for the fleet. Amazon GameLift Servers saves all standard output for each container in logs, including game session logs. You can select from the following methods:

- `CLOUDWATCH` -- Send logs to an Amazon CloudWatch log group that you define. Each container emits a log stream, which is organized in the log group.
- `S3` -- Store logs in an Amazon S3 bucket that you define.
- `NONE` -- Don't collect container logs.

By default, this property is set to `CLOUDWATCH`.

Amazon GameLift Servers requires permissions to send logs other AWS services in your account. These permissions are included in the IAM fleet role for this container fleet (see `FleetRoleArn`).

Type: [LogConfiguration](#) object

Required: No

[MetricGroups](#)

The name of an AWS CloudWatch metric group to add this fleet to. You can use a metric group to aggregate metrics for multiple fleets. You can specify an existing metric group name or use a new name to create a new metric group. Each fleet can have only one metric group, but you can change this value at any time.

Type: Array of strings

Array Members: Maximum number of 1 item.

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

Required: No

[NewGameSessionProtectionPolicy](#)

Determines whether Amazon GameLift Servers can shut down game sessions on the fleet that are actively running and hosting players. Amazon GameLift Servers might prompt an instance shutdown when scaling down fleet capacity or when retiring unhealthy instances. You can also set game session protection for individual game sessions using [UpdateGameSession](#).

- **NoProtection** -- Game sessions can be shut down during active gameplay.
- **FullProtection** -- Game sessions in ACTIVE status can't be shut down.

By default, this property is set to NoProtection.

Type: String

Valid Values: NoProtection | FullProtection

Required: No

[PerInstanceContainerGroupDefinitionName](#)

The name of a container group definition resource that describes a set of axillary software. A fleet instance has one process for executables in this container group. A per-instance container group is optional. You can update the fleet to add or remove a per-instance container group at any time. You can specify the container group definition's name to use the latest version. Alternatively, provide an ARN value with a specific version number.

Create a container group definition by calling https://docs.aws.amazon.com/gamelift/latest/apireference/API_CreateContainerGroupDefinition.html. This operation creates a https://docs.aws.amazon.com/gamelift/latest/apireference/API_ContainerGroupDefinition.html resource.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-zA-Z0-9\-\_]+|^arn:.*:containergroupdefinition\[a-zA-Z0-9\-\_]+(:[0-9]+)?$`

Required: No

PlayerGatewayMode

Configures player gateway for your fleet. Player gateway provides benefits such as DDoS protection by rate limiting and validating traffic before it reaches game servers, hiding game server IP addresses from players, and providing updated endpoints when relay endpoints become unhealthy.

How it works: When enabled, game clients connect to relay endpoints instead of to your game servers. Player gateway validates player gateway tokens and routes traffic to the appropriate game server. Your game backend calls [GetPlayerConnectionDetails](#) to retrieve relay endpoints and player gateway tokens for your game clients. To learn more about this topic, see [DDoS protection with Amazon GameLift Servers player gateway](#).

Possible values include:

- **DISABLED** (default) -- Game clients connect to the game server endpoint. Use this when you do not intend to integrate your game with player gateway.
- **ENABLED** -- Player gateway is available in fleet locations where it is supported. Your game backend can call [GetPlayerConnectionDetails](#) to obtain a player gateway token and endpoints for game clients.
- **REQUIRED** -- Player gateway is available in fleet locations where it is supported, and the fleet can only use locations that support this feature. Attempting to add a remote location to your fleet which does not support player gateway will result in an `InvalidRequestException`.

Type: String

Valid Values: `DISABLED` | `ENABLED` | `REQUIRED`

Required: No

Tags

A list of labels to assign to the new fleet resource. Tags are developer-defined key-value pairs. Tagging AWS resources are useful for resource management, access management and cost allocation. For more information, see [Tagging AWS Resources](#) in the *AWS General Reference*.

Type: Array of [Tag](#) objects

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

Required: No

Response Syntax

```
{
  "ContainerFleet": {
    "BillingType": "string",
    "CreationTime": number,
    "DeploymentDetails": {
      "LatestDeploymentId": "string"
    },
    "Description": "string",
    "FleetArn": "string",
    "FleetId": "string",
    "FleetRoleArn": "string",
    "GameServerContainerGroupDefinitionArn": "string",
    "GameServerContainerGroupDefinitionName": "string",
    "GameServerContainerGroupsPerInstance": number,
    "GameSessionCreationLimitPolicy": {
      "NewGameSessionsPerCreator": number,
      "PolicyPeriodInMinutes": number
    },
    "InstanceConnectionPortRange": {
      "FromPort": number,
      "ToPort": number
    },
    "InstanceInboundPermissions": [
      {
        "FromPort": number,
        "IpRange": "string",
        "Protocol": "string",
```

```
        "ToPort": number
      },
    ],
    "InstanceType": "string",
    "LocationAttributes": [
      {
        "Location": "string",
        "PlayerGatewayStatus": "string",
        "Status": "string"
      }
    ],
    "LogConfiguration": {
      "LogDestination": "string",
      "LogGroupArn": "string",
      "S3BucketName": "string"
    },
    "MaximumGameServerContainerGroupsPerInstance": number,
    "MetricGroups": [ "string" ],
    "NewGameSessionProtectionPolicy": "string",
    "PerInstanceContainerGroupDefinitionArn": "string",
    "PerInstanceContainerGroupDefinitionName": "string",
    "PlayerGatewayMode": "string",
    "Status": "string"
  }
}
```

Response Elements

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

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

ContainerFleet

The properties for the new container fleet, including current status. All fleets are initially placed in PENDING status.

Type: [ContainerFleet](#) object

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Create a simple single-region container fleet

This example creates a container fleet with a game server container group definition only. It uses all the Amazon GameLift Servers defaults.

Sample Request

```
{
  "FleetRoleArn": "arn:aws:iam::MyAccount:role/MyRole",
  "GameServerContainerGroupDefinitionName": "arn:aws:gamelift:us-
west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:2"
}
```

Sample Response

```
{
  "ContainerFleet": {
    "BillingType": ON_DEMAND,
    "CreationTime": 1736365885.22,
    "DeploymentDetails": {
      "LatestDeploymentId": "deployment-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
    },
    "FleetArn": "arn:aws:gamelift:us-west-2:111122223333:containerfleet/
containerfleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "FleetId": "containerfleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "FleetRoleArn": "arn:aws:iam::MyAccount:role/MyRole",
    "GameServerContainerGroupDefinitionArn": "arn:aws:gamelift:us-
west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:2",
    "GameServerContainerGroupDefinitionName": "MyAdventureGameContainerGroup",
    "GameServerContainerGroupsPerInstance": number,
    "InstanceConnectionPortRange": {
      "FromPort": 4192,
      "ToPort": 4242
    },
    "InstanceInboundPermissions": [
      {
        "FromPort": 4192,
        "IpRange": "string",
        "Protocol": "UDP",

```

```
        "ToPort": 4242,
      },
    ],
    "InstanceType": "c5.large",
    "LogConfiguration": {
      "LogGroupArn": "arn:aws:logs:us-west-2:111222333444:log-group:customerLogs",
      "LogDestination": "CLOUDWATCH"
    },
    "MaximumGameServerContainerGroupsPerInstance": 10,
    "NewGameSessionProtectionPolicy": "NoProtection",
    "Status": "PENDING"
  }
}
```

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](#)

CreateContainerGroupDefinition

This API works with the following fleet types: Container

Creates a `ContainerGroupDefinition` that describes a set of containers for hosting your game server with Amazon GameLift Servers managed containers hosting. An Amazon GameLift Servers container group is similar to a container task or pod. Use container group definitions when you create a container fleet with [CreateContainerFleet](#).

A container group definition determines how Amazon GameLift Servers deploys your containers to each instance in a container fleet. You can maintain multiple versions of a container group definition.

There are two types of container groups:

- A **game server container group** has the containers that run your game server application and supporting software. A game server container group can have these container types:
 - Game server container. This container runs your game server. You can define one game server container in a game server container group.
 - Support container. This container runs software in parallel with your game server. You can define up to 8 support containers in a game server group.

When building a game server container group definition, you can choose to bundle your game server executable and all dependent software into a single game server container. Alternatively, you can separate the software into one game server container and one or more support containers.

On a container fleet instance, a game server container group can be deployed multiple times (depending on the compute resources of the instance). This means that all containers in the container group are replicated together.

- A **per-instance container group** has containers for processes that aren't replicated on a container fleet instance. This might include background services, logging, test processes, or processes that need to persist independently of the game server container group. When building a per-instance container group, you can define up to 10 support containers.

Note

This operation requires AWS Identity and Access Management (IAM) permissions to access container images in Amazon ECR repositories. See [IAM permissions for Amazon GameLift Servers](#) for help setting the appropriate permissions.

Request options

Use this operation to make the following types of requests. You can specify values for the minimum required parameters and customize optional values later.

- Create a game server container group definition. Provide the following required parameter values:
 - Name
 - ContainerGroupType (GAME_SERVER)
 - OperatingSystem
 - TotalMemoryLimitMebibytes
 - TotalVcpuLimit
 - At least one GameServerContainerDefinition
 - ContainerName
 - ImageUrl
 - PortConfiguration
 - ServerSdkVersion
- Create a per-instance container group definition. Provide the following required parameter values:
 - Name
 - ContainerGroupType (PER_INSTANCE)
 - OperatingSystem
 - TotalMemoryLimitMebibytes
 - TotalVcpuLimit
 - At least one SupportContainerDefinition

- `ImageUrl`

Results

If successful, this request creates a `ContainerGroupDefinition` resource and assigns a unique ARN value. You can update most properties of a container group definition by calling [UpdateContainerGroupDefinition](#), and optionally save the update as a new version.

Request Syntax

```
{
  "ContainerGroupType": "string",
  "GameServerContainerDefinition": {
    "ContainerName": "string",
    "DependsOn": [
      {
        "Condition": "string",
        "ContainerName": "string"
      }
    ],
    "EnvironmentOverride": [
      {
        "Name": "string",
        "Value": "string"
      }
    ],
    "ImageUri": "string",
    "LinuxCapabilities": {
      "Include": [ "string" ]
    },
    "MountPoints": [
      {
        "AccessLevel": "string",
        "ContainerPath": "string",
        "InstancePath": "string"
      }
    ],
    "PortConfiguration": {
      "ContainerPortRanges": [
        {
          "FromPort": number,
          "Protocol": "string",
```



```
        "ToPort": number
      }
    ]
  },
  "ServerSdkVersion": "string"
},
"Name": "string",
"OperatingSystem": "string",
"SupportContainerDefinitions": [
  {
    "ContainerName": "string",
    "DependsOn": [
      {
        "Condition": "string",
        "ContainerName": "string"
      }
    ],
    "EnvironmentOverride": [
      {
        "Name": "string",
        "Value": "string"
      }
    ],
    "Essential": boolean,
    "HealthCheck": {
      "Command": [ "string" ],
      "Interval": number,
      "Retries": number,
      "StartPeriod": number,
      "Timeout": number
    },
    "ImageUri": "string",
    "LinuxCapabilities": {
      "Include": [ "string" ]
    },
    "MemoryHardLimitMebibytes": number,
    "MountPoints": [
      {
        "AccessLevel": "string",
        "ContainerPath": "string",
        "InstancePath": "string"
      }
    ],
    "PortConfiguration": {
```

```
    "ContainerPortRanges": [
      {
        "FromPort": number,
        "Protocol": "string",
        "ToPort": number
      }
    ],
    "Vcpu": number
  },
  "Tags": [
    {
      "Key": "string",
      "Value": "string"
    }
  ],
  "TotalMemoryLimitMebibytes": number,
  "TotalVcpuLimit": number,
  "VersionDescription": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Name

A descriptive identifier for the container group definition. The name value must be unique in an AWS Region.

Type: String

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

Pattern: `^[a-zA-Z0-9\-.]+$`

Required: Yes

OperatingSystem

The platform that all containers in the group use. Containers in a group must run on the same operating system.

Default value: AMAZON_LINUX_2023

Note

Amazon Linux 2 (AL2) will reach end of support on 6/30/2026. See more details in the [Amazon Linux 2 FAQs](#). For game servers that are hosted on AL2 and use server SDK version 4.x for Amazon GameLift Servers, first update the game server build to server SDK 5.x, and then deploy to AL2023 instances. See [Migrate to server SDK version 5](#).

Type: String

Valid Values: AMAZON_LINUX_2023

Required: Yes

TotalMemoryLimitMebibytes

The maximum amount of memory (in MiB) to allocate to the container group. All containers in the group share this memory. If you specify memory limits for an individual container, the total value must be greater than any individual container's memory limit.

Default value: 1024

Type: Integer

Valid Range: Minimum value of 4. Maximum value of 1024000.

Required: Yes

TotalVcpuLimit

The maximum amount of vCPU units to allocate to the container group (1 vCPU is equal to 1024 CPU units). All containers in the group share this memory. If you specify vCPU limits for individual containers, the total value must be equal to or greater than the sum of the CPU limits for all containers in the group.

Default value: 1

Type: Double

Valid Range: Minimum value of 0.125. Maximum value of 10.

Required: Yes

ContainerGroupType

The type of container group being defined. Container group type determines how Amazon GameLift Servers deploys the container group on each fleet instance.

Default value: GAME_SERVER

Type: String

Valid Values: GAME_SERVER | PER_INSTANCE

Required: No

GameServerContainerDefinition

The definition for the game server container in this group. Define a game server container only when the container group type is GAME_SERVER. Game server containers specify a container image with your game server build. You can pass in your container definitions as a JSON file.

Type: [GameServerContainerDefinitionInput](#) object

Required: No

SupportContainerDefinitions

One or more definition for support containers in this group. You can define a support container in any type of container group. You can pass in your container definitions as a JSON file.

Type: Array of [SupportContainerDefinitionInput](#) objects

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

Required: No

Tags

A list of labels to assign to the container group definition resource. Tags are developer-defined key-value pairs. Tagging AWS resources are useful for resource management, access management and cost allocation. For more information, see [Tagging AWS Resources](#) in the *AWS General Reference*.

Type: Array of [Tag](#) objects

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

Required: No

VersionDescription

A description for the initial version of this container group definition.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "ContainerGroupDefinition": {
    "ContainerGroupDefinitionArn": "string",
    "ContainerGroupType": "string",
    "CreationTime": number,
    "GameServerContainerDefinition": {
      "ContainerName": "string",
      "DependsOn": [
        {
          "Condition": "string",
          "ContainerName": "string"
        }
      ],
      "EnvironmentOverride": [
        {
          "Name": "string",
          "Value": "string"
        }
      ],
      "ImageUri": "string",
      "LinuxCapabilities": {
        "Include": [ "string" ]
      },
      "MountPoints": [
        {
          "AccessLevel": "string",
```

```
        "ContainerPath": "string",
        "InstancePath": "string"
    }
],
"PortConfiguration": {
    "ContainerPortRanges": [
        {
            "FromPort": number,
            "Protocol": "string",
            "ToPort": number
        }
    ]
},
"ResolvedImageDigest": "string",
"ServerSdkVersion": "string"
},
"Name": "string",
"OperatingSystem": "string",
"Status": "string",
"StatusReason": "string",
"SupportContainerDefinitions": [
    {
        "ContainerName": "string",
        "DependsOn": [
            {
                "Condition": "string",
                "ContainerName": "string"
            }
        ],
        "EnvironmentOverride": [
            {
                "Name": "string",
                "Value": "string"
            }
        ],
        "Essential": boolean,
        "HealthCheck": {
            "Command": [ "string" ],
            "Interval": number,
            "Retries": number,
            "StartPeriod": number,
            "Timeout": number
        },
        "ImageUri": "string",
```

```

    "LinuxCapabilities": {
      "Include": [ "string" ]
    },
    "MemoryHardLimitMebibytes": number,
    "MountPoints": [
      {
        "AccessLevel": "string",
        "ContainerPath": "string",
        "InstancePath": "string"
      }
    ],
    "PortConfiguration": {
      "ContainerPortRanges": [
        {
          "FromPort": number,
          "Protocol": "string",
          "ToPort": number
        }
      ]
    },
    "ResolvedImageDigest": "string",
    "Vcpu": number
  }
],
"TotalMemoryLimitMebibytes": number,
"TotalVcpuLimit": number,
"VersionDescription": "string",
"VersionNumber": number
}
}

```

Response Elements

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

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

ContainerGroupDefinition

The properties of the new container group definition resource. You can use this resource to create a container fleet.

Type: [ContainerGroupDefinition](#) object

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Create a simple game server container group definition

This example creates a game server container group definition with a game server container definition. The request includes the minimal parameter settings and uses the default values for all other parameters. The container definition has these characteristics:

- It defines a container port range, which is required for containers that run game server processes.
- The image URI points to an ECR private repo.
- It doesn't need a health check configuration. Amazon GameLift Servers handles the health checks for a game server container.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "Name": "MyAdventureGameContainerGroup",
  "ContainerGroupType": "GAME_SERVER",
  "GameServerContainerDefinition": [
    {
      "ContainerName": "MyAdventureGameServer",
      "ImageUri": "111122223333.dkr.ecr.us-west-2.amazonaws.com/MyAdventureGameContainerImage",
      "ServerSdkVersion": 5.2.0,
      "PortConfiguration": {
        "ContainerPortRanges": [
          {
            "FromPort": 35000,
            "Protocol": "UDP",
```

```
        "ToPort": 40000
      }
    ]
  }
}
]
```

Sample Response

```
{
  "ContainerGroupDefinition": {
    "ContainerGroupDefinitionArn": "arn:aws:gamelift:us-
west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:1",
    "ContainerGroupType": "GAME_SERVER",
    "CreationTime": 1496365885.44,
    "GameServerContainerDefinition": [
      {
        "ContainerName": "MyAdventureGameServer",
        "ImageUri": "111122223333.dkr.ecr.us-west-2.amazonaws.com/
MyAdventureGameContainerImage",
        "PortConfiguration": {
          "ContainerPortRanges": [
            {
              "FromPort": 35000,
              "Protocol": "TCP",
              "ToPort": 40000
            }
          ]
        },
        "ResolvedImageDigest":
"sha256:0123456789abcdef0123456789abcdef0123456789abcdef0123456789abcdef",
        "ServerSdkVersion": "5.2.0"
      }
    ],
    "Name": "MyAdventureGameContainerGroup",
    "OperatingSystem": "AMAZON_LINUX_2023",
    "Status": "COPYING",
    "TotalMemoryLimitMebibytes": 1024,
    "TotalVcpuLimit": 1,
    "VersionNumber": "1"
  }
}
```

Create a game server container with support containers

This example creates a game server container group definition with a game server container definition and one support container. The request includes the minimal parameter settings and uses the default values for all other parameters.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "Name": "MyAdventureGameContainerGroup",
  "ContainerGroupType": "GAME_SERVER",
  "GameServerContainerDefinition": [
    {
      "ContainerName": "MyAdventureGameServer",
      "ImageUri": "111122223333.dkr.ecr.us-west-2.amazonaws.com/MyAdventureGameContainerImage",
      "ServerSdkVersion": 5.2.0,
      "PortConfiguration": {
        "ContainerPortRanges": [
          {
            "FromPort": 35000,
            "Protocol": "UDP",
            "ToPort": 40000
          }
        ]
      }
    }
  ],
  "SupportContainerDefinitions": [
    {
      "ContainerName": "MyAdventureGameDependencies",
      "ImageUri": "111122223333.dkr.ecr.us-west-2.amazonaws.com/MyAdventureGameContainerImage"
    }
  ]
}
```

Sample Response

```
{
```

```
"ContainerGroupDefinition": {
  "ContainerGroupDefinitionArn": "arn:aws:gamelift:us-
west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:1",
  "ContainerGroupType": "GAME_SERVER",
  "CreationTime": 1496365885.44,
  "GameServerContainerDefinition": [
    {
      "ContainerName": "MyAdventureGameContainer",
      "ImageUri": "111122223333.dkr.ecr.us-west-2.amazonaws.com/
MyAdventureGameContainerImage",
      "PortConfiguration": {
        "ContainerPortRanges": [
          {
            "FromPort": 35000,
            "Protocol": "TCP",
            "ToPort": 40000
          }
        ]
      },
      "ResolvedImageDigest":
"sha256:0123456789abcdef0123456789abcdef0123456789abcdef0123456789abcdef",
      "ServerSdkVersion": "5.2.0"
    }
  ],
  "Name": "MyAdventureGameContainerGroup",
  "OperatingSystem": "AMAZON_LINUX_2023",
  "Status": "READY",
  "SupportContainerDefinitions": [
    {
      "ContainerName": "MyAdventureGameDependencies",
      "ImageUri": "111122223333.dkr.ecr.us-west-2.amazonaws.com/
MyAdventureGameContainerImage"
    }
  ],
  "TotalMemoryLimitMebibytes": 1024,
  "TotalVcpuLimit": 1,
  "VersionNumber": "1"
}
```

See Also

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

- [AWS Command Line Interface 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](#)

CreateFleet

This API works with the following fleet types: EC2, Anywhere, Container

Creates a fleet of compute resources to host your game servers. Use this operation to set up a fleet for the following compute types:

Managed EC2 fleet

An EC2 fleet is a set of Amazon Elastic Compute Cloud (Amazon EC2) instances. Your game server build is deployed to each fleet instance. Amazon GameLift Servers manages the fleet's instances and controls the lifecycle of game server processes, which host game sessions for players. EC2 fleets can have instances in multiple locations. Each instance in the fleet is designated a Compute.

To create an EC2 fleet, provide these required parameters:

- Either `BuildId` or `ScriptId`
- `ComputeType` set to EC2 (the default value)
- `EC2InboundPermissions`
- `EC2InstanceType`
- `FleetType`
- `Name`
- `RuntimeConfiguration` with at least one `ServerProcesses` configuration

If successful, this operation creates a new fleet resource and places it in `NEW` status while Amazon GameLift Servers initiates the [fleet creation workflow](#). To debug your fleet, fetch logs, view performance metrics or other actions on the fleet, create a development fleet with port 22/3389 open. As a best practice, we recommend opening ports for remote access only when you need them and closing them when you're finished.

When the fleet status is `ACTIVE`, you can adjust capacity settings and turn autoscaling on/off for each location.

Note

A managed fleet's runtime environment depends on the Amazon Machine Image (AMI) version it uses. When a new fleet is created, Amazon GameLift Servers assigns the latest available AMI version to the fleet, and all compute instances in that fleet are deployed with

that version. To update the AMI version, you must create a new fleet. As a best practice, we recommend replacing your managed fleets every 30 days to maintain a secure and up-to-date runtime environment for your hosted game servers. For guidance, see [Security best practices for Amazon GameLift Servers](#).

Anywhere fleet

An Anywhere fleet represents compute resources that are not owned or managed by Amazon GameLift Servers. You might create an Anywhere fleet with your local machine for testing, or use one to host game servers with on-premises hardware or other game hosting solutions.

To create an Anywhere fleet, provide these required parameters:

- `ComputeType` set to `ANYWHERE`
- `Locations` specifying a custom location
- `Name`

If successful, this operation creates a new fleet resource and places it in `ACTIVE` status. You can register computes with a fleet in `ACTIVE` status.

Learn more

[Setting up fleets](#)

[Debug fleet creation issues](#)

[Multi-location fleets](#)

Request Syntax

```
{
  "AnywhereConfiguration": {
    "Cost": "string"
  },
  "BuildId": "string",
  "CertificateConfiguration": {
    "CertificateType": "string"
  },
  "ComputeType": "string",
  "Description": "string",
```

```

    "EC2InboundPermissions": [
      {
        "FromPort": number,
        "IpRange": "string",
        "Protocol": "string",
        "ToPort": number
      }
    ],
    "EC2InstanceType": "string",
    "FleetType": "string",
    "InstanceRoleArn": "string",
    "InstanceRoleCredentialsProvider": "string",
    "Locations": [
      {
        "Location": "string"
      }
    ],
    "LogPaths": [ "string" ],
    "MetricGroups": [ "string" ],
    "Name": "string",
    "NewGameSessionProtectionPolicy": "string",
    "PeerVpcAwsAccountId": "string",
    "PeerVpcId": "string",
    "PlayerGatewayConfiguration": {
      "GameServerIpProtocolSupported": "string"
    },
    "PlayerGatewayMode": "string",
    "ResourceCreationLimitPolicy": {
      "NewGameSessionsPerCreator": number,
      "PolicyPeriodInMinutes": number
    },
    "RuntimeConfiguration": {
      "GameSessionActivationTimeoutSeconds": number,
      "MaxConcurrentGameSessionActivations": number,
      "ServerProcesses": [
        {
          "ConcurrentExecutions": number,
          "LaunchPath": "string",
          "Parameters": "string"
        }
      ]
    },
    "ScriptId": "string",
    "ServerLaunchParameters": "string",

```



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

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Name

A descriptive label that is associated with a fleet. Fleet names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

AnywhereConfiguration

Amazon GameLift Servers Anywhere configuration options.

Type: [AnywhereConfiguration](#) object

Required: No

BuildId

The unique identifier for a custom game server build to be deployed to a fleet with compute type EC2. You can use either the build ID or ARN. The build must be uploaded to Amazon GameLift Servers and in READY status. This fleet property can't be changed after the fleet is created.

Type: String

Pattern: ^build-\\S+|^arn:.*:build\\/build-\\S+

Required: No

CertificateConfiguration

Prompts Amazon GameLift Servers to generate a TLS/SSL certificate for the fleet. Amazon GameLift Servers uses the certificates to encrypt traffic between game clients and the game servers running on Amazon GameLift Servers. By default, the `CertificateConfiguration` is `DISABLED`. You can't change this property after you create the fleet.

AWS Certificate Manager (ACM) certificates expire after 13 months. Certificate expiration can cause fleets to fail, preventing players from connecting to instances in the fleet. We recommend you replace fleets before 13 months, consider using fleet aliases for a smooth transition.

Note

ACM isn't available in all AWS regions. A fleet creation request with certificate generation enabled in an unsupported Region, fails with a 4xx error. For more information about the supported Regions, see [Supported Regions](#) in the *AWS Certificate Manager User Guide*.

Type: [CertificateConfiguration](#) object

Required: No

ComputeType

The type of compute resource used to host your game servers.

- **EC2** – The game server build is deployed to Amazon EC2 instances for cloud hosting. This is the default setting.
- **ANYWHERE** – Game servers and supporting software are deployed to compute resources that you provide and manage. With this compute type, you can also set the `AnywhereConfiguration` parameter.

Type: String

Valid Values: EC2 | ANYWHERE

Required: No

Description

A description for the fleet.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

EC2InboundPermissions

The IP address ranges and port settings that allow inbound traffic to access game server processes and other processes on this fleet. Set this parameter for managed EC2 fleets. You can leave this parameter empty when creating the fleet, but you must call https://docs.aws.amazon.com/gamelift/latest/apireference/API_UpdateFleetPortSettings to set it before players can connect to game sessions. As a best practice, we recommend opening ports for remote access only when you need them and closing them when you're finished. For Amazon GameLift Servers Realtime fleets, Amazon GameLift Servers automatically sets TCP and UDP ranges.

Type: Array of [IpPermission](#) objects

Array Members: Maximum number of 50 items.

Required: No

EC2InstanceType

The Amazon GameLift Servers-supported Amazon EC2 instance type to use with managed EC2 fleets. Instance type determines the computing resources that will be used to host your game servers, including CPU, memory, storage, and networking capacity. See [Amazon Elastic Compute Cloud Instance Types](#) for detailed descriptions of Amazon EC2 instance types.

Type: String

Valid Values: t2.micro | t2.small | t2.medium | t2.large | c3.large | c3.xlarge | c3.2xlarge | c3.4xlarge | c3.8xlarge | c4.large | c4.xlarge | c4.2xlarge | c4.4xlarge | c4.8xlarge | c5.large | c5.xlarge | c5.2xlarge | c5.4xlarge | c5.9xlarge | c5.12xlarge | c5.18xlarge | c5.24xlarge | c5a.large | c5a.xlarge | c5a.2xlarge | c5a.4xlarge | c5a.8xlarge | c5a.12xlarge | c5a.16xlarge | c5a.24xlarge | r3.large |

r3.xlarge | r3.2xlarge | r3.4xlarge | r3.8xlarge | r4.large | r4.xlarge
| r4.2xlarge | r4.4xlarge | r4.8xlarge | r4.16xlarge | r5.large |
r5.xlarge | r5.2xlarge | r5.4xlarge | r5.8xlarge | r5.12xlarge |
r5.16xlarge | r5.24xlarge | r5a.large | r5a.xlarge | r5a.2xlarge |
r5a.4xlarge | r5a.8xlarge | r5a.12xlarge | r5a.16xlarge | r5a.24xlarge
| m3.medium | m3.large | m3.xlarge | m3.2xlarge | m4.large | m4.xlarge
| m4.2xlarge | m4.4xlarge | m4.10xlarge | m5.large | m5.xlarge |
m5.2xlarge | m5.4xlarge | m5.8xlarge | m5.12xlarge | m5.16xlarge |
m5.24xlarge | m5a.large | m5a.xlarge | m5a.2xlarge | m5a.4xlarge |
m5a.8xlarge | m5a.12xlarge | m5a.16xlarge | m5a.24xlarge | c5d.large
| c5d.xlarge | c5d.2xlarge | c5d.4xlarge | c5d.9xlarge | c5d.12xlarge
| c5d.18xlarge | c5d.24xlarge | c6a.large | c6a.xlarge | c6a.2xlarge |
c6a.4xlarge | c6a.8xlarge | c6a.12xlarge | c6a.16xlarge | c6a.24xlarge
| c6i.large | c6i.xlarge | c6i.2xlarge | c6i.4xlarge | c6i.8xlarge |
c6i.12xlarge | c6i.16xlarge | c6i.24xlarge | r5d.large | r5d.xlarge |
r5d.2xlarge | r5d.4xlarge | r5d.8xlarge | r5d.12xlarge | r5d.16xlarge
| r5d.24xlarge | m6g.medium | m6g.large | m6g.xlarge | m6g.2xlarge |
m6g.4xlarge | m6g.8xlarge | m6g.12xlarge | m6g.16xlarge | c6g.medium
| c6g.large | c6g.xlarge | c6g.2xlarge | c6g.4xlarge | c6g.8xlarge |
c6g.12xlarge | c6g.16xlarge | r6g.medium | r6g.large | r6g.xlarge |
r6g.2xlarge | r6g.4xlarge | r6g.8xlarge | r6g.12xlarge | r6g.16xlarge
| c6gn.medium | c6gn.large | c6gn.xlarge | c6gn.2xlarge | c6gn.4xlarge
| c6gn.8xlarge | c6gn.12xlarge | c6gn.16xlarge | c7g.medium | c7g.large
| c7g.xlarge | c7g.2xlarge | c7g.4xlarge | c7g.8xlarge | c7g.12xlarge
| c7g.16xlarge | r7g.medium | r7g.large | r7g.xlarge | r7g.2xlarge |
r7g.4xlarge | r7g.8xlarge | r7g.12xlarge | r7g.16xlarge | m7g.medium
| m7g.large | m7g.xlarge | m7g.2xlarge | m7g.4xlarge | m7g.8xlarge |
m7g.12xlarge | m7g.16xlarge | g5g.xlarge | g5g.2xlarge | g5g.4xlarge
| g5g.8xlarge | g5g.16xlarge | r6i.large | r6i.xlarge | r6i.2xlarge |
r6i.4xlarge | r6i.8xlarge | r6i.12xlarge | r6i.16xlarge | c6gd.medium |
c6gd.large | c6gd.xlarge | c6gd.2xlarge | c6gd.4xlarge | c6gd.8xlarge |
c6gd.12xlarge | c6gd.16xlarge | c6in.large | c6in.xlarge | c6in.2xlarge
| c6in.4xlarge | c6in.8xlarge | c6in.12xlarge | c6in.16xlarge |
c7a.medium | c7a.large | c7a.xlarge | c7a.2xlarge | c7a.4xlarge |
c7a.8xlarge | c7a.12xlarge | c7a.16xlarge | c7gd.medium | c7gd.large |
c7gd.xlarge | c7gd.2xlarge | c7gd.4xlarge | c7gd.8xlarge | c7gd.12xlarge

| c7gd.16xlarge | c7gn.medium | c7gn.large | c7gn.xlarge | c7gn.2xlarge
| c7gn.4xlarge | c7gn.8xlarge | c7gn.12xlarge | c7gn.16xlarge |
c7i.large | c7i.xlarge | c7i.2xlarge | c7i.4xlarge | c7i.8xlarge |
c7i.12xlarge | c7i.16xlarge | m6a.large | m6a.xlarge | m6a.2xlarge |
m6a.4xlarge | m6a.8xlarge | m6a.12xlarge | m6a.16xlarge | m6gd.medium
| m6gd.large | m6gd.xlarge | m6gd.2xlarge | m6gd.4xlarge | m6gd.8xlarge
| m6gd.12xlarge | m6gd.16xlarge | m6i.large | m6i.xlarge | m6i.2xlarge
| m6i.4xlarge | m6i.8xlarge | m6i.12xlarge | m6i.16xlarge | m7a.medium
| m7a.large | m7a.xlarge | m7a.2xlarge | m7a.4xlarge | m7a.8xlarge |
m7a.12xlarge | m7a.16xlarge | m7gd.medium | m7gd.large | m7gd.xlarge
| m7gd.2xlarge | m7gd.4xlarge | m7gd.8xlarge | m7gd.12xlarge |
m7gd.16xlarge | m7i.large | m7i.xlarge | m7i.2xlarge | m7i.4xlarge |
m7i.8xlarge | m7i.12xlarge | m7i.16xlarge | r6gd.medium | r6gd.large |
r6gd.xlarge | r6gd.2xlarge | r6gd.4xlarge | r6gd.8xlarge | r6gd.12xlarge
| r6gd.16xlarge | r7a.medium | r7a.large | r7a.xlarge | r7a.2xlarge |
r7a.4xlarge | r7a.8xlarge | r7a.12xlarge | r7a.16xlarge | r7gd.medium |
r7gd.large | r7gd.xlarge | r7gd.2xlarge | r7gd.4xlarge | r7gd.8xlarge |
r7gd.12xlarge | r7gd.16xlarge | r7i.large | r7i.xlarge | r7i.2xlarge |
r7i.4xlarge | r7i.8xlarge | r7i.12xlarge | r7i.16xlarge | r7i.24xlarge |
r7i.48xlarge | c5ad.large | c5ad.xlarge | c5ad.2xlarge | c5ad.4xlarge |
c5ad.8xlarge | c5ad.12xlarge | c5ad.16xlarge | c5ad.24xlarge | c5n.large
| c5n.xlarge | c5n.2xlarge | c5n.4xlarge | c5n.9xlarge | c5n.18xlarge |
r5ad.large | r5ad.xlarge | r5ad.2xlarge | r5ad.4xlarge | r5ad.8xlarge |
r5ad.12xlarge | r5ad.16xlarge | r5ad.24xlarge | c6id.large | c6id.xlarge
| c6id.2xlarge | c6id.4xlarge | c6id.8xlarge | c6id.12xlarge |
c6id.16xlarge | c6id.24xlarge | c6id.32xlarge | c8g.medium | c8g.large
| c8g.xlarge | c8g.2xlarge | c8g.4xlarge | c8g.8xlarge | c8g.12xlarge
| c8g.16xlarge | c8g.24xlarge | c8g.48xlarge | m5ad.large | m5ad.xlarge
| m5ad.2xlarge | m5ad.4xlarge | m5ad.8xlarge | m5ad.12xlarge |
m5ad.16xlarge | m5ad.24xlarge | m5d.large | m5d.xlarge | m5d.2xlarge |
m5d.4xlarge | m5d.8xlarge | m5d.12xlarge | m5d.16xlarge | m5d.24xlarge
| m5dn.large | m5dn.xlarge | m5dn.2xlarge | m5dn.4xlarge | m5dn.8xlarge
| m5dn.12xlarge | m5dn.16xlarge | m5dn.24xlarge | m5n.large |
m5n.xlarge | m5n.2xlarge | m5n.4xlarge | m5n.8xlarge | m5n.12xlarge |
m5n.16xlarge | m5n.24xlarge | m6id.large | m6id.xlarge | m6id.2xlarge
| m6id.4xlarge | m6id.8xlarge | m6id.12xlarge | m6id.16xlarge

| m6id.24xlarge | m6id.32xlarge | m6idn.large | m6idn.xlarge |
m6idn.2xlarge | m6idn.4xlarge | m6idn.8xlarge | m6idn.12xlarge |
m6idn.16xlarge | m6idn.24xlarge | m6idn.32xlarge | m6in.large |
m6in.xlarge | m6in.2xlarge | m6in.4xlarge | m6in.8xlarge | m6in.12xlarge
| m6in.16xlarge | m6in.24xlarge | m6in.32xlarge | m8g.medium | m8g.large
| m8g.xlarge | m8g.2xlarge | m8g.4xlarge | m8g.8xlarge | m8g.12xlarge
| m8g.16xlarge | m8g.24xlarge | m8g.48xlarge | r5dn.large | r5dn.xlarge
| r5dn.2xlarge | r5dn.4xlarge | r5dn.8xlarge | r5dn.12xlarge |
r5dn.16xlarge | r5dn.24xlarge | r5n.large | r5n.xlarge | r5n.2xlarge |
r5n.4xlarge | r5n.8xlarge | r5n.12xlarge | r5n.16xlarge | r5n.24xlarge
| r6a.large | r6a.xlarge | r6a.2xlarge | r6a.4xlarge | r6a.8xlarge |
r6a.12xlarge | r6a.16xlarge | r6a.24xlarge | r6a.32xlarge | r6a.48xlarge
| r6id.large | r6id.xlarge | r6id.2xlarge | r6id.4xlarge | r6id.8xlarge
| r6id.12xlarge | r6id.16xlarge | r6id.24xlarge | r6id.32xlarge
| r6idn.large | r6idn.xlarge | r6idn.2xlarge | r6idn.4xlarge |
r6idn.8xlarge | r6idn.12xlarge | r6idn.16xlarge | r6idn.24xlarge |
r6idn.32xlarge | r6in.large | r6in.xlarge | r6in.2xlarge | r6in.4xlarge
| r6in.8xlarge | r6in.12xlarge | r6in.16xlarge | r6in.24xlarge |
r6in.32xlarge | r8g.medium | r8g.large | r8g.xlarge | r8g.2xlarge |
r8g.4xlarge | r8g.8xlarge | r8g.12xlarge | r8g.16xlarge | r8g.24xlarge |
r8g.48xlarge | m4.16xlarge | c6a.32xlarge | c6a.48xlarge | c6i.32xlarge
| r6i.24xlarge | r6i.32xlarge | c6in.24xlarge | c6in.32xlarge |
c7a.24xlarge | c7a.32xlarge | c7a.48xlarge | c7i.24xlarge | c7i.48xlarge
| m6a.24xlarge | m6a.32xlarge | m6a.48xlarge | m6i.24xlarge |
m6i.32xlarge | m7a.24xlarge | m7a.32xlarge | m7a.48xlarge | m7i.24xlarge
| m7i.48xlarge | r7a.24xlarge | r7a.32xlarge | r7a.48xlarge

Required: No

FleetType

Indicates whether to use On-Demand or Spot instances for this fleet. By default, this property is set to ON_DEMAND. Learn more about when to use [On-Demand versus Spot Instances](#). This fleet property can't be changed after the fleet is created.

Type: String

Valid Values: ON_DEMAND | SPOT

Required: No

InstanceRoleArn

A unique identifier for an IAM role that manages access to your AWS services. With an instance role ARN set, any application that runs on an instance in this fleet can assume the role, including install scripts, server processes, and daemons (background processes). Create a role or look up a role's ARN by using the [IAM dashboard](#) in the AWS Management Console. Learn more about using on-box credentials for your game servers at [Access external resources from a game server](#). This fleet property can't be changed after the fleet is created.

Type: String

Length Constraints: Minimum length of 1.

Required: No

InstanceRoleCredentialsProvider

Prompts Amazon GameLift Servers to generate a shared credentials file for the IAM role that's defined in InstanceRoleArn. The shared credentials file is stored on each fleet instance and refreshed as needed. Use shared credentials for applications that are deployed along with the game server executable, if the game server is integrated with server SDK version 5.x. For more information about using shared credentials, see [Communicate with other AWS resources from your fleets](#).

Type: String

Valid Values: SHARED_CREDENTIAL_FILE

Required: No

Locations

A set of remote locations to deploy additional instances to and manage as a multi-location fleet. Use this parameter when creating a fleet in AWS Regions that support multiple locations. You can add any AWS Region or Local Zone that's supported by Amazon GameLift Servers. Provide a list of one or more AWS Region codes, such as us-west-2, or Local Zone names. When using this parameter, Amazon GameLift Servers requires you to include your home location in the request. For a list of supported Regions and Local Zones, see [Amazon GameLift Servers service locations](#) for managed hosting.

Type: Array of [LocationConfiguration](#) objects

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

Required: No

LogPaths

This parameter is no longer used. To specify where Amazon GameLift Servers should store log files once a server process shuts down, use the Amazon GameLift Servers server API `ProcessReady()` and specify one or more directory paths in `logParameters`. For more information, see [Initialize the server process](#) in the *Amazon GameLift Servers Developer Guide*.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

MetricGroups

The name of an AWS CloudWatch metric group to add this fleet to. A metric group is used to aggregate the metrics for multiple fleets. You can specify an existing metric group name or set a new name to create a new metric group. A fleet can be included in only one metric group at a time.

Type: Array of strings

Array Members: Maximum number of 1 item.

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

Required: No

NewGameSessionProtectionPolicy

The status of termination protection for active game sessions on the fleet. By default, this property is set to `NoProtection`. You can also set game session protection for an individual game session by calling [UpdateGameSession](#).

- **NoProtection** - Game sessions can be terminated during active gameplay as a result of a scale-down event.
- **FullProtection** - Game sessions in `ACTIVE` status cannot be terminated during a scale-down event.

Type: String

Valid Values: NoProtection | FullProtection

Required: No

PeerVpcAwsAccountId

Used when peering your Amazon GameLift Servers fleet with a VPC, the unique identifier for the AWS account that owns the VPC. You can find your account ID in the AWS Management Console under account settings.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

PeerVpcId

A unique identifier for a VPC with resources to be accessed by your Amazon GameLift Servers fleet. The VPC must be in the same Region as your fleet. To look up a VPC ID, use the [VPC Dashboard](#) in the AWS Management Console. Learn more about VPC peering in [VPC Peering with Amazon GameLift Servers Fleets](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

PlayerGatewayConfiguration

Configuration settings for player gateway. Use this to specify advanced options for how player gateway handles connections.

Type: [PlayerGatewayConfiguration](#) object

Required: No

PlayerGatewayMode

Configures player gateway for your fleet. Player gateway provides benefits such as DDoS protection by rate limiting and validating traffic before it reaches game servers, hiding game server IP addresses from players, and providing updated endpoints when relay endpoints

become unhealthy. Note, player gateway is only available for fleets using server SDK 5.x or later game server builds.

How it works: When enabled, game clients connect to relay endpoints instead of to your game servers. Player gateway validates player gateway tokens and routes traffic to the appropriate game server. Your game backend calls [GetPlayerConnectionDetails](#) to retrieve relay endpoints and player gateway tokens for your game clients. To learn more about this topic, see [DDoS protection with Amazon GameLift Servers player gateway](#).

Possible values include:

- **DISABLED** (default) -- Game clients connect to the game server endpoint. Use this when you do not intend to integrate your game with player gateway.
- **ENABLED** -- Player gateway is available in fleet locations where it is supported. Your game backend can call [GetPlayerConnectionDetails](#) to obtain a player gateway token and endpoints for game clients.
- **REQUIRED** -- Player gateway is available in fleet locations where it is supported, and the fleet can only use locations that support this feature. Attempting to add a remote location to your fleet which does not support player gateway will result in an `InvalidRequestException`.

Type: String

Valid Values: DISABLED | ENABLED | REQUIRED

Required: No

[ResourceCreationLimitPolicy](#)

A policy that limits the number of game sessions that an individual player can create on instances in this fleet within a specified span of time.

Type: [ResourceCreationLimitPolicy](#) object

Required: No

[RuntimeConfiguration](#)

Instructions for how to launch and run server processes on the fleet. Set runtime configuration for managed EC2 fleets. For an Anywhere fleets, set this parameter only if the fleet is running the Amazon GameLift Servers Agent. The runtime configuration defines one or more server process configurations. Each server process identifies a game executable or Realtime script file and the number of processes to run concurrently.

Note

This parameter replaces the parameters `ServerLaunchPath` and `ServerLaunchParameters`, which are still supported for backward compatibility.

Type: [RuntimeConfiguration](#) object

Required: No

[ScriptId](#)

The unique identifier for a Realtime configuration script to be deployed to a fleet with compute type EC2. You can use either the script ID or ARN. Scripts must be uploaded to Amazon GameLift Servers prior to creating the fleet. This fleet property can't be changed after the fleet is created.

Type: String

Pattern: `^script-\S+|^arn:.*:script\/script-\S+`

Required: No

[ServerLaunchParameters](#)

This parameter is no longer used. Specify server launch parameters using the `RuntimeConfiguration` parameter. Requests that use this parameter instead continue to be valid.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `[A-Za-z0-9_:.+\V\\- =@;{} ,?'\[\]" ]+`

Required: No

[ServerLaunchPath](#)

This parameter is no longer used. Specify a server launch path using the `RuntimeConfiguration` parameter. Requests that use this parameter instead continue to be valid.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `[A-Za-z0-9_:.+\\/\\" data-bbox="83 136 204 154" data-label="Text">

Required: No`

Tags

A list of labels to assign to the new fleet resource. Tags are developer-defined key-value pairs. Tagging AWS resources are useful for resource management, access management and cost allocation. For more information, see [Tagging AWS Resources](#) in the *AWS General Reference*.

Type: Array of [Tag](#) objects

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

Required: No

Response Syntax

```
{
  "FleetAttributes": {
    "AnywhereConfiguration": {
      "Cost": "string"
    },
    "BuildArn": "string",
    "BuildId": "string",
    "CertificateConfiguration": {
      "CertificateType": "string"
    },
    "ComputeType": "string",
    "CreationTime": number,
    "Description": "string",
    "FleetArn": "string",
    "FleetId": "string",
    "FleetType": "string",
    "InstanceRoleArn": "string",
    "InstanceRoleCredentialsProvider": "string",
    "InstanceType": "string",
    "LogPaths": [ "string" ],
    "MetricGroups": [ "string" ],
    "Name": "string",
    "NewGameSessionProtectionPolicy": "string",
```

```
"OperatingSystem": "string",
"PlayerGatewayConfiguration": {
  "GameServerIpProtocolSupported": "string"
},
"PlayerGatewayMode": "string",
"ResourceCreationLimitPolicy": {
  "NewGameSessionsPerCreator": number,
  "PolicyPeriodInMinutes": number
},
"ScriptArn": "string",
"ScriptId": "string",
"ServerLaunchParameters": "string",
"ServerLaunchPath": "string",
"Status": "string",
"StoppedActions": [ "string " ],
"TerminationTime": number
},
"LocationStates": [
  {
    "Location": "string",
    "PlayerGatewayStatus": "string",
    "Status": "string"
  }
]
}
```

Response Elements

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

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

[FleetAttributes](#)

The properties for the new fleet, including the current status. All fleets are placed in NEW status on creation.

Type: [FleetAttributes](#) object

[LocationStates](#)

The fleet's locations and life-cycle status of each location. For new fleets, the status of all locations is set to NEW. During fleet creation, Amazon GameLift Servers updates each location

status as instances are deployed there and prepared for game hosting. This list includes an entry for the fleet's home Region. For fleets with no remote locations, only one entry, representing the home Region, is returned.

Type: Array of [LocationState](#) objects

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

NotReadyException

The operation failed because Amazon GameLift Servers has not yet finished validating this compute. We recommend attempting 8 to 10 retries over 3 to 5 minutes with [exponential backoffs and jitter](#).

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Create fleet with minimal configuration

This example creates an EC2 fleet for a game server build with a minimal configuration and a placeholder launch path. You can use this fleet to create queues and matchmakers, test Amazon GameLift Servers Server API calls, etc. When you're ready to start hosting game sessions, complete the configuration settings using the `UpdateFleet` operations. If the fleet is in a Region that supports multiple locations, you can add remote locations with `CreateFleetLocations` later. Note: this example generates a TLS certificate for the fleet, which can only be enabled during fleet creation.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "Name": "My_Fleet_1",
  "Description": "A minimal sample fleet",
  "BuildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "CertificateConfiguration": {
    "CertificateType": "GENERATED"
  },
  "ComputeType": "EC2",
  "EC2InstanceType": "c4.large",
  "FleetType": "ON_DEMAND",
  "RuntimeConfiguration": {
    "ServerProcesses": [
      {
        "LaunchPath": "/local/game/mygame.exe",
        "ConcurrentExecutions": 1
      }
    ]
  }
}
```

Sample Response

```
{
  "FleetAttributes": {
    "BuildArn": "arn:aws:gamelift:us-west-2::build/build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "BuildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "CertificateConfiguration": {
      "CertificateType": "GENERATED"
    },
    "ComputeType": "EC2",
    "CreationTime": 1496365885.44,
    "Description": "A minimal sample fleet",
    "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "FleetArn": "arn:aws:gamelift:us-west-2::fleet/fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "FleetType": "ON_DEMAND",
    "InstanceType": "c4.large",
    "MetricGroups": [
      "default"
    ],
    "Name": "My_Fleet_1",
    "NewGameSessionProtectionPolicy": "NoProtection",
  }
}
```



```
        "OperatingSystem": "AMAZON_LINUX_2023",
        "Status": "NEW"
    }
}
```

Create fleet with minimal configuration (Windows)

This example creates a fleet for a game server build with a minimal configuration and a placeholder launch path. You can use this fleet to create queues and matchmakers, test Amazon GameLift Servers Server API calls, etc. Once you're ready to start hosting game sessions, complete the configuration settings with the `UpdateFleet` operations. If the fleet is created in a Region that supports multiple locations, you can add remote locations with `CreateFleetLocations` later.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "Name": "My_Fleet_1",
  "Description": "A minimal sample fleet",
  "BuildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "ComputeType": "EC2",
  "EC2InstanceType": "c4.large",
  "FleetType": "ON_DEMAND",
  "RuntimeConfiguration": {
    "ServerProcesses": [
      {
        "LaunchPath": "c:\\game\\mygame.exe",
        "ConcurrentExecutions": 1
      }
    ]
  }
}
```

Sample Response

```
{
  "FleetAttributes": {
    "BuildArn": "arn:aws:gamelift:us-west-2::build/build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "BuildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "CertificateConfiguration": {
      "CertificateType": "DISABLED"
    }
  }
}
```

```
    },
    "ComputeType": "EC2",
    "CreationTime": 1496365885.44,
    "Description": "A minimal sample fleet",
    "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "FleetArn": "arn:aws:gamelift:us-west-2::fleet/
fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "FleetType": "ON_DEMAND",
    "InstanceType": "c4.large",
    "MetricGroups": [
        "default"
    ],
    "Name": "My_Fleet_1",
    "NewGameSessionProtectionPolicy": "NoProtection",
    "OperatingSystem": "WINDOWS_2022",
    "Status": "ACTIVE"
}
}
```

Create fleet with full configuration

This example creates an EC2 fleet with complete configuration details. The new fleet, which is created in the AWS Region `us-west-2`, also deploys instances to two remote locations. In this example, the runtime configuration defines two server process configurations. The first configuration calls for three concurrent processes of the game server to run in standard mode. The second configuration calls for one process of the game server to run in a test mode. As a result, all fleet instances will maintain a total of four processes concurrently. This example also references an instance role ARN, which allows the processes running on each instance to access other AWS resources. The example opens port 22 for debugging the fleet. The game server build in this example runs on Windows.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "Name": "My_Fleet_1",
  "Description": "A fully configured sample fleet with",
  "BuildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "CertificateConfiguration": {
    "CertificateType": "GENERATED"
  }
}
```

```

    },
    "ComputeType": "EC2",
    "EC2InstanceType": "c5.large",
    "EC2InboundPermissions": [
        {
            "FromPort": 33435,
            "ToPort": 33435,
            "IpRange": "10.24.34.0/23",
            "Protocol": "UDP"
        },
        {
            "FromPort": 22,
            "ToPort": 22,
            "IpRange": "10.24.34.0/23",
            "Protocol": "TCP"
        }
    ],
    "FleetType": "ON_DEMAND",
    "Locations": [
        {
            "Location": "us-east-2"
        },
        {
            "Location": "ca-central-1"
        }
    ],
    "NewGameSessionProtectionPolicy": "FullProtection",
    "RuntimeConfiguration": {
        "ServerProcesses": [
            {
                "LaunchPath": "c:\\game\\mygame.exe",
                "Parameters": "+map Winter444",
                "ConcurrentExecutions": 3
            },
            {
                "LaunchPath": "c:\\game\\mygame.exe",
                "Parameters": "-dev -console +map Winter444",
                "ConcurrentExecutions": 1
            }
        ],
        "MaxConcurrentGameSessionActivations": 2,
        "GameSessionActivationTimeoutSeconds": 300
    },
    "ResourceCreationLimitPolicy": {
        "NewGameSessionsPerCreator": 3,
        "PolicyPeriodInMinutes": 15
    },
    "MetricGroups": ["EMEAfleets"],
    "InstanceRoleArn": "arn:aws:iam::444455556666:role/S3AccessForGameLift",
}

```

Sample Response

```

{
    "FleetAttributes": {

```

```

    "BuildArn": "arn:aws:gamelift:us-west-2::build/
build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "BuildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "CertificateConfiguration": {
        "CertificateType": "GENERATED"
    },
    "ComputeType": "EC2",
    "CreationTime": 1496375088.502,
    "Description": "A fully configured sample fleet",
    "FleetArn": "arn:aws:gamelift:us-west-2::fleet/
fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "FleetType": "ON_DEMAND",
    "InstanceRoleArn": "arn:aws:iam::444455556666:role/S3AccessForGameLift",
    "MetricGroups": [
        "EMEAfleets"
    ],
    "Name": "My_Fleet_One",
    "NewGameSessionProtectionPolicy": "FullProtection",
    "OperatingSystem": "WINDOWS_2022",
    "ResourceCreationLimitPolicy": {
        "NewGameSessionsPerCreator": 3,
        "PolicyPeriodInMinutes": 15
    }
},
"LocationStates": [
    {
        "Location": "us-east-2",
        "Status": "NEW"
    },
    {
        "Location": "ca-central-1",
        "Status": "NEW"
    }
]
    "Status": "NEW",
}

```

Create a Realtime Servers fleet

This example creates a fleet using a Realtime script that has been uploaded to Amazon GameLift Servers. All Realtime servers are deployed onto Linux machines. This example represents a simple yet complete fleet configuration. You can change configuration settings with the

[UpdateRuntimeConfiguration](#) operation. The new fleet, which is created in the AWS Region us-west-2, also deploys instances to two remote locations.

In this example, the uploaded Realtime script includes multiple script files, with the `Init()` function located in the script file called "myMainScript.js". This file is identified as the launch script in the runtime configuration.

Note

We recommend using a minimal version of the Realtime script when creating your fleets (see this [working code example](#)). This will make it much easier to troubleshoot fleet creation issues. After the fleet has reached ACTIVE status, you can update your Realtime script as needed.

Sample Request

```
{
  "Name": "My_Realtime_Fleet_1",
  "Description": "A complete Realtime sample fleet",
  "CertificateConfiguration": {
    "CertificateType": "GENERATED"
  },
  "EC2InstanceType": "c4.large",
  "FleetType": "SPOT",
  "ComputeType": "EC2",
  "Locations": [
    {"Location": "us-east-2"},
    {"Location": "ca-central-1"}
  ],
  "RuntimeConfiguration": {
    "ServerProcesses": [
      {"LaunchPath": "/local/game/myMainScript.js",
        "Parameters": "+map Winter444",
        "ConcurrentExecutions": 5}
    ]
  },
  "ScriptId": "script-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
}
```

Sample Response

```
{
  "FleetAttributes": {
    "CertificateConfiguration": {
      "CertificateType": "GENERATED"
    },
    "ComputeType": "EC2",
    "CreationTime": 1496375088.502,
    "Description": "A complete Realtime sample fleet",
    "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "FleetArn": "arn:aws:gamelift:us-west-2::fleet/fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "FleetType": "SPOT",
    "MetricGroups": [
      "default"
    ],
    "Name": "My_Realtime_Fleet_1",
    "NewGameSessionProtectionPolicy": "NoProtection",
    "OperatingSystem": "AMAZON_LINUX_2023",
    "ScriptArn": "arn:aws:gamelift:us-west-2::script/script-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "ScriptId": "script-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "Status": "NEW"
  },
  "LocationStates": [
    {
      "Location": "us-east-2",
      "Status": "NEW"
    },
    {
      "Location": "ca-central-1",
      "Status": "NEW"
    }
  ]
}
```

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](#)

CreateFleetLocations

This API works with the following fleet types: EC2, Anywhere, Container

Adds remote locations to an EC2 and begins populating the new locations with instances. The new instances conform to the fleet's instance type, auto-scaling, and other configuration settings.

Note

You can't add remote locations to a fleet that resides in an AWS Region that doesn't support multiple locations. Fleets created prior to March 2021 can't support multiple locations.

To add fleet locations, specify the fleet to be updated and provide a list of one or more locations.

If successful, this operation returns the list of added locations with their status set to NEW. Amazon GameLift Servers initiates the process of starting an instance in each added location. You can track the status of each new location by monitoring location creation events using [DescribeFleetEvents](#).

Learn more

[Setting up fleets](#)

[Update fleet locations](#)

[Amazon GameLift Servers service locations](#) for managed hosting.

Request Syntax

```
{
  "FleetId": "string",
  "Locations": [
    {
      "Location": "string"
    }
  ]
}
```

Request Parameters

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

The request accepts the following data in JSON format.

 Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet to add locations to. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

Locations

A list of locations to deploy additional instances to and manage as part of the fleet. You can add any Amazon GameLift Servers-supported AWS Region as a remote location, in the form of an AWS Region code such as `us-west-2`.

Type: Array of [LocationConfiguration](#) objects

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

Required: Yes

Response Syntax

```
{
  "FleetArn": "string",
  "FleetId": "string",
  "LocationStates": [
    {
      "Location": "string",
      "PlayerGatewayStatus": "string",
      "Status": "string"
    }
  ]
}
```

```
}  
]  
}
```

Response Elements

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

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

FleetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet[/[a-z]*fleet-[a-zA-Z0-9\ -]]+$`

FleetId

A unique identifier for the fleet that was updated with new locations.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\ -]]+$|^arn:.*:[a-z]*fleet[/[a-z]*fleet-[a-zA-Z0-9\ -]]+$`

LocationStates

The remote locations that are being added to the fleet, and the life-cycle status of each location. For new locations, the status is set to NEW. During location creation, Amazon GameLift Servers updates each location's status as instances are deployed there and prepared for game hosting. This list does not include the fleet home Region or any remote locations that were already added to the fleet.

Type: Array of [LocationState](#) objects

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidFleetStatusException

The requested operation would cause a conflict with the current state of a resource associated with the request and/or the fleet. Resolve the conflict before retrying.

HTTP Status Code: 400

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

NotReadyException

The operation failed because Amazon GameLift Servers has not yet finished validating this compute. We recommend attempting 8 to 10 retries over 3 to 5 minutes with [exponential backoffs and jitter](#).

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Add remote locations to an existing fleet

This example adds three five remote locations to a fleet that resides in the AWS Region eu-central-1. As shown, the status of each new location is initially set to NEW.

Sample Request

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "Locations": [
    { "Location": "me-south-1"},
    { "Location": "eu-west-2"},
    { "Location": "eu-south-1"},
    { "Location": "eu-west-3"},
    { "Location": "eu-north-1"}
  ]
}
```

Sample Response

```
{
```

```
"FleetArn": "arn:aws:gamelift:us-west-2::fleet/
fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
"FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
"LocationStates": [
  {
    "Location": "eu-north-1",
    "Status": "NEW"
  },
  {
    "Location": "eu-west-3",
    "Status": "NEW"
  },
  {
    "Location": "eu-south-1",
    "Status": "NEW"
  },
  {
    "Location": "eu-west-2",
    "Status": "NEW"
  },
  {
    "Location": "me-south-1",
    "Status": "NEW"
  }
]
```

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](#)

CreateGameServerGroup

This API works with the following fleet types: EC2 (FleetIQ)

Creates a Amazon GameLift Servers FleetIQ game server group for managing game hosting on a collection of Amazon Elastic Compute Cloud instances for game hosting. This operation creates the game server group, creates an Auto Scaling group in your AWS account, and establishes a link between the two groups. You can view the status of your game server groups in the Amazon GameLift Servers console. Game server group metrics and events are emitted to Amazon CloudWatch.

Before creating a new game server group, you must have the following:

- An Amazon Elastic Compute Cloud launch template that specifies how to launch Amazon Elastic Compute Cloud instances with your game server build. For more information, see [Launching an Instance from a Launch Template](#) in the *Amazon Elastic Compute Cloud User Guide*.
- An IAM role that extends limited access to your AWS account to allow Amazon GameLift Servers FleetIQ to create and interact with the Auto Scaling group. For more information, see [Create IAM roles for cross-service interaction](#) in the *Amazon GameLift Servers FleetIQ Developer Guide*.

To create a new game server group, specify a unique group name, IAM role and Amazon Elastic Compute Cloud launch template, and provide a list of instance types that can be used in the group. You must also set initial maximum and minimum limits on the group's instance count. You can optionally set an Auto Scaling policy with target tracking based on a Amazon GameLift Servers FleetIQ metric.

Once the game server group and corresponding Auto Scaling group are created, you have full access to change the Auto Scaling group's configuration as needed. Several properties that are set when creating a game server group, including maximum/minimum size and auto-scaling policy settings, must be updated directly in the Auto Scaling group. Keep in mind that some Auto Scaling group properties are periodically updated by Amazon GameLift Servers FleetIQ as part of its balancing activities to optimize for availability and cost.

Learn more

[Amazon GameLift Servers FleetIQ Guide](#)

Request Syntax

```
{
  "AutoScalingPolicy": {
    "EstimatedInstanceWarmup": number,
    "TargetTrackingConfiguration": {
      "TargetValue": number
    }
  },
  "BalancingStrategy": "string",
  "GameServerGroupName": "string",
  "GameServerProtectionPolicy": "string",
  "InstanceDefinitions": [
    {
      "InstanceType": "string",
      "WeightedCapacity": "string"
    }
  ],
  "LaunchTemplate": {
    "LaunchTemplateId": "string",
    "LaunchTemplateName": "string",
    "Version": "string"
  },
  "MaxSize": number,
  "MinSize": number,
  "RoleArn": "string",
  "Tags": [
    {
      "Key": "string",
      "Value": "string"
    }
  ],
  "VpcSubnets": [ "string" ]
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

GameServerGroupName

An identifier for the new game server group. This value is used to generate unique ARN identifiers for the Amazon EC2 Auto Scaling group and the Amazon GameLift Servers FleetIQ game server group. The name must be unique per Region per AWS account.

Type: String

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

Pattern: [a-zA-Z0-9-\.]+

Required: Yes

InstanceDefinitions

The Amazon EC2 instance types and sizes to use in the Auto Scaling group. The instance definitions must specify at least two different instance types that are supported by Amazon GameLift Servers FleetIQ. For more information on instance types, see [EC2 Instance Types](#) in the *Amazon Elastic Compute Cloud User Guide*. You can optionally specify capacity weighting for each instance type. If no weight value is specified for an instance type, it is set to the default value "1". For more information about capacity weighting, see [Instance Weighting for Amazon EC2 Auto Scaling](#) in the Amazon EC2 Auto Scaling User Guide.

Type: Array of [InstanceDefinition](#) objects

Array Members: Minimum number of 2 items. Maximum number of 20 items.

Required: Yes

LaunchTemplate

The Amazon EC2 launch template that contains configuration settings and game server code to be deployed to all instances in the game server group. You can specify the template using either the template name or ID. For help with creating a launch template, see [Creating a Launch Template for an Auto Scaling Group](#) in the *Amazon Elastic Compute Cloud Auto Scaling User Guide*. After the Auto Scaling group is created, update this value directly in the Auto Scaling group using the AWS console or APIs.

Note

If you specify network interfaces in your launch template, you must explicitly set the property `AssociatePublicIpAddress` to "true". If no network interface is specified in the launch template, Amazon GameLift Servers FleetIQ uses your account's default VPC.

Type: [LaunchTemplateSpecification](#) object

Required: Yes

[MaxSize](#)

The maximum number of instances allowed in the Amazon EC2 Auto Scaling group. During automatic scaling events, Amazon GameLift Servers FleetIQ and EC2 do not scale up the group above this maximum. After the Auto Scaling group is created, update this value directly in the Auto Scaling group using the AWS console or APIs.

Type: Integer

Valid Range: Minimum value of 1.

Required: Yes

[MinSize](#)

The minimum number of instances allowed in the Amazon EC2 Auto Scaling group. During automatic scaling events, Amazon GameLift Servers FleetIQ and Amazon EC2 do not scale down the group below this minimum. In production, this value should be set to at least 1. After the Auto Scaling group is created, update this value directly in the Auto Scaling group using the AWS console or APIs.

Type: Integer

Valid Range: Minimum value of 0.

Required: Yes

[RoleArn](#)

The Amazon Resource Name ([ARN](#)) for an IAM role that allows Amazon GameLift Servers to access your Amazon EC2 Auto Scaling groups.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^arn:.*:role\[\\w+=,.\@-]\+`

Required: Yes

AutoScalingPolicy

Configuration settings to define a scaling policy for the Auto Scaling group that is optimized for game hosting. The scaling policy uses the metric "PercentUtilizedGameServers" to maintain a buffer of idle game servers that can immediately accommodate new games and players. After the Auto Scaling group is created, update this value directly in the Auto Scaling group using the AWS console or APIs.

Type: [GameServerGroupAutoScalingPolicy](#) object

Required: No

BalancingStrategy

Indicates how Amazon GameLift Servers FleetIQ balances the use of Spot Instances and On-Demand Instances in the game server group. Method options include the following:

- SPOT_ONLY - Only Spot Instances are used in the game server group. If Spot Instances are unavailable or not viable for game hosting, the game server group provides no hosting capacity until Spot Instances can again be used. Until then, no new instances are started, and the existing nonviable Spot Instances are terminated (after current gameplay ends) and are not replaced.
- SPOT_PREFERRED - (default value) Spot Instances are used whenever available in the game server group. If Spot Instances are unavailable, the game server group continues to provide hosting capacity by falling back to On-Demand Instances. Existing nonviable Spot Instances are terminated (after current gameplay ends) and are replaced with new On-Demand Instances.
- ON_DEMAND_ONLY - Only On-Demand Instances are used in the game server group. No Spot Instances are used, even when available, while this balancing strategy is in force.

Type: String

Valid Values: SPOT_ONLY | SPOT_PREFERRED | ON_DEMAND_ONLY

Required: No

GameServerProtectionPolicy

A flag that indicates whether instances in the game server group are protected from early termination. Unprotected instances that have active game servers running might be terminated during a scale-down event, causing players to be dropped from the game. Protected instances cannot be terminated while there are active game servers running except in the event of a forced game server group deletion (see). An exception to this is with Spot Instances, which can be terminated by AWS regardless of protection status. This property is set to NO_PROTECTION by default.

Type: String

Valid Values: NO_PROTECTION | FULL_PROTECTION

Required: No

Tags

A list of labels to assign to the new game server group resource. Tags are developer-defined key-value pairs. Tagging AWS resources is useful for resource management, access management, and cost allocation. For more information, see [Tagging AWS Resources](#) in the *AWS General Reference*.

Type: Array of [Tag](#) objects

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

Required: No

VpcSubnets

A list of virtual private cloud (VPC) subnets to use with instances in the game server group. By default, all Amazon GameLift Servers FleetIQ-supported Availability Zones are used. You can use this parameter to specify VPCs that you've set up. This property cannot be updated after the game server group is created, and the corresponding Auto Scaling group will always use the property value that is set with this request, even if the Auto Scaling group is updated directly.

Type: Array of strings

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

Length Constraints: Minimum length of 15. Maximum length of 24.

Pattern: ^subnet-[0-9a-z]+\$

Required: No

Response Syntax

```
{
  "GameServerGroup": {
    "AutoScalingGroupArn": "string",
    "BalancingStrategy": "string",
    "CreationTime": number,
    "GameServerGroupArn": "string",
    "GameServerGroupName": "string",
    "GameServerProtectionPolicy": "string",
    "InstanceDefinitions": [
      {
        "InstanceType": "string",
        "WeightedCapacity": "string"
      }
    ],
    "LastUpdatedTime": number,
    "RoleArn": "string",
    "Status": "string",
    "StatusReason": "string",
    "SuspendedActions": [ "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.

GameServerGroup

The newly created game server group object, including the new ARN value for the Amazon GameLift Servers FleetIQ game server group and the object's status. The Amazon EC2 Auto Scaling group ARN is initially null, since the group has not yet been created. This value is added once the game server group status reaches ACTIVE.

Type: [GameServerGroup](#) object

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Create a new game server group

This example creates a game server group and Auto Scaling group that deploys your game server software to Spot instances of two instance types. The game server group is set up to use Spot

Instances only, so if no Spot Instances of these types are available, the game server group will effectively stop starting new game servers until Spot resources are available again. This request also triggers creation of a target-tracking Auto Scaling policy. With this policy, the Auto Scaling group will maintain approximately 75% game server utilization. The 25% of idle capacity enables the game server group to immediately accommodate sudden increases in player demand.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "GameServerGroupName": "MegaFrogServers_NA",
  "RoleArn": "arn:aws:iam:123456789012::role/GameLiftGsgRole",
  "MinSize": 1,
  "MaxSize": 1000,
  "GameServerProtectionPolicy": "NO_PROTECTION",
  "BalancingStrategy": "SPOT_ONLY",
  "LaunchTemplate": {
    "LaunchTemplateId": "lt-1234567890abcdef",
  },
  "InstanceDefinitions": [
    {"InstanceType": "c5.xlarge"},
    {"InstanceType": "m5.xlarge"}
  ],
  "AutoScalingPolicy": {
    "TargetTrackingConfiguration": {
      "TargetValue": 75
    }
  }
}
```

CLI command:

```
aws gamelift create-game-server-group \
  --game-server-group-name MegaFrogServers_NA \
  --role-arn arn:aws:iam::123456789012:role/GameLiftGSGRole \
  --min-size 1 \
  --max-size 1000 \
  --game-server-protection-policy NO_PROTECTION \
  --balancing-strategy SPOT_ONLY \
  --launch-template LaunchTemplateId=lt-012ab345cde6789ff \
```

```
--instance-definitions '[{"InstanceType": "c5.xlarge"}, {"InstanceType":
"m5.xlarge"}]' \
    --auto-scaling-policy '{"TargetTrackingConfiguration": {"TargetValue": 75}}'
```

Sample Response

```
{
  "GameServerGroup": {
    "AutoScalingGroupArn": "arn:aws:autoscaling:us-
west-2:123456789012:autoScalingGroup:1111aaaa-22bb-33cc-44dd-5555eeee66ff:autoScalingGroupName/
MegaFrogServers_NA",
    "BalancingStrategy": "SPOT_ONLY",
    "CreationTime": 1496365885.44,
    "GameServerGroupArn": "arn:aws:gamelift:us-west-2::GameServerGroup/
MegaFrogServers_NA",
    "GameServerGroupName": " MegaFrogServers_NA",
    "GameServerProtectionPolicy": "NO_PROTECTION",
    "InstanceDefinitions": [
      {
        "InstanceType": "c5.xlarge",
        "WeightedCapacity": "1"
      }
      {
        "InstanceType": "m5.xlarge",
        "WeightedCapacity": "1"
      }
    ],
    "LastUpdatedTime": 1496365885.44,
    "RoleArn": "arn:aws:iam:123456789012::role/GameLiftGsgRole",
    "Status": "NEW",
    "StatusReason": "",
    "SuspendedActions": []
  }
}
```

Create a new game server group with weighted instances

This example creates a game server group with three weighted instance types. Because a balancing strategy is not defined in this example, the default "SPOT_PREFERRED" is used.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "GameServerGroupName": "MegaFrogServers_NA",
  "RoleArn": "arn:aws:iam:123456789012::role/GameLiftGsgRole",
  "MinSize": 1,
  "MaxSize": 200,
  "InstanceDefinitions": [
    {
      "InstanceType": "c5.2xlarge",
      "WeightedCapacity": "1"
    },
    {
      "InstanceType": "c5.4xlarge",
      "WeightedCapacity": "2"
    },
    {
      "InstanceType": "c5.24xlarge",
      "WeightedCapacity": "12"
    }
  ],
  "LaunchTemplate": {
    "LaunchTemplateName": "MegaFrogServers"
  }
}
```

CLI command:

```
aws gamelift create-game-server-group \
  --game-server-group-name MegaFrogServers_NA \
  --role-arn arn:aws:iam:123456789012::role/GameLiftGsgRole \
  --min-size 1 \
  --max-size 200 \
  --launch-template {"LaunchTemplateName": MegaFrogServers} \
  --instance-definitions '[{"IntanceType":"m5.2xlarge","WeightedCapacity":"1"},
{"InstanceType":"m5.4xlarge","WeightedCapacity":"2"},
{"InstanceType":"m5.24xlarge","WeightedCapacity":"12"}]'
```

Sample Response

```
{
  "GameServerGroup": {
```

```
    "AutoScalingGroupArn": "arn:aws:autoscaling:us-west-2:123456789012:autoScalingGroup:1111aaaa-22bb-33cc-44dd-5555eeee66ff:autoScalingGroupName/MegaFrogServers_NA",
    "BalancingStrategy": "SPOT_PREFERRED",
    "CreationTime": 1496365885.44,
    "GameServerGroupArn": "arn:aws:gamelift:us-west-2::GameServerGroup/MegaFrogServers_NA",
    "GameServerGroupName": " MegaFrogServers_NA",
    "GameServerProtectionPolicy": "NO_PROTECTION",
    "InstanceDefinitions": [
      {
        "InstanceType": "c5.2xlarge",
        "WeightedCapacity": "1"
      },
      {
        "InstanceType": "c5.4xlarge",
        "WeightedCapacity": "2"
      },
      {
        "InstanceType": "c5.24xlarge",
        "WeightedCapacity": "12"
      }
    ],
    "LastUpdatedTime": 1496365885.44,
    "RoleArn": "arn:aws:iam:123456789012::role/GameLiftGsgRole",
    "Status": "NEW",
    "StatusReason": "",
    "SuspendedActions": []
  }
```

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](#)

CreateGameSession

This API works with the following fleet types: EC2, Anywhere, Container

Creates a multiplayer game session for players in a specific fleet location. This operation prompts an available server process to start a game session and retrieves connection information for the new game session. As an alternative, consider using the Amazon GameLift Servers game session placement feature with [StartGameSessionPlacement](#), which uses the FleetIQ algorithm and queues to optimize the placement process.

When creating a game session, you specify exactly where you want to place it and provide a set of game session configuration settings. The target fleet must be in ACTIVE status.

You can use this operation in the following ways:

- To create a game session on an instance in a fleet's home Region, provide a fleet or alias ID along with your game session configuration.
- To create a game session on an instance in a fleet's remote location, provide a fleet or alias ID and a location name, along with your game session configuration.
- To create a game session on an instance in an Anywhere fleet, specify the fleet's custom location.

If successful, Amazon GameLift Servers initiates a workflow to start a new game session and returns a `GameSession` object containing the game session configuration and status. When the game session status is ACTIVE, it is updated with connection information and you can create player sessions for the game session. By default, newly created game sessions are open to new players. You can restrict new player access by using [UpdateGameSession](#) to change the game session's player session creation policy.

Amazon GameLift Servers retains logs for active for 14 days. To access the logs, call [GetGameSessionLogUrl](#) to download the log files.

Available in Amazon GameLift Servers Local.

Learn more

[Start a game session](#)

[All APIs by task](#)

Request Syntax

```
{
  "AliasId": "string",
  "CreatorId": "string",
  "FleetId": "string",
  "GameProperties": [
    {
      "Key": "string",
      "Value": "string"
    }
  ],
  "GameSessionData": "string",
  "GameSessionId": "string",
  "IdempotencyToken": "string",
  "Location": "string",
  "MaximumPlayerSessionCount": number,
  "Name": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

MaximumPlayerSessionCount

The maximum number of players that can be connected simultaneously to the game session.

Type: Integer

Valid Range: Minimum value of 0.

Required: Yes

AliasId

A unique identifier for the alias associated with the fleet to create a game session in. You can use either the alias ID or ARN value. Each request must reference either a fleet ID or alias ID, but not both.

Type: String

Pattern: `^alias-\S+|^arn:.*:alias\/alias-\S+`

Required: No

CreatorId

A unique identifier for a player or entity creating the game session.

If you add a resource creation limit policy to a fleet, the `CreateGameSession` operation requires a `CreatorId`. Amazon GameLift Servers limits the number of game session creation requests with the same `CreatorId` in a specified time period.

If you your fleet doesn't have a resource creation limit policy and you provide a `CreatorId` in your `CreateGameSession` requests, Amazon GameLift Servers limits requests to one request per `CreatorId` per second.

To not limit `CreateGameSession` requests with the same `CreatorId`, don't provide a `CreatorId` in your `CreateGameSession` request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

FleetId

A unique identifier for the fleet to create a game session in. You can use either the fleet ID or ARN value. Each request must reference either a fleet ID or alias ID, but not both.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\-]+|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\-]+`

Required: No

GameProperties

A set of key-value pairs that can store custom data in a game session. For example: {"Key": "difficulty", "Value": "novice"}. For an example, see [Create a game session with custom properties](#).

Note

- Avoid using periods (".") in property keys if you plan to search for game sessions by properties. Property keys containing periods cannot be searched and will be filtered out from search results due to search index limitations.
- If you use SearchGameSessions API, there is a limit of 500 game property keys across all game sessions and all fleets per region. If the limit is exceeded, there will potentially be game session entries missing from SearchGameSessions API results.

Type: Array of [GameProperty](#) objects

Array Members: Maximum number of 16 items.

Required: No

GameSessionData

A set of custom game session properties, formatted as a single string value. This data is passed to a game server process with a request to start a new game session. For more information, see [Start a game session](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 262144.

Required: No

GameSessionId

This parameter is deprecated. Use IdempotencyToken instead.

Custom string that uniquely identifies a request for a new game session. Maximum token length is 48 characters. If provided, this string is included in the new game session's ID.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 48.

Pattern: [a-zA-Z0-9-]+

Required: No

IdempotencyToken

Custom string that uniquely identifies the new game session request. This is useful for ensuring that game session requests with the same idempotency token are processed only once. Subsequent requests with the same string return the original `GameSession` object, with an updated status. Maximum token length is 48 characters. If provided, this string is included in the new game session's ID. The value is always a full ARN in the following format: For Home Region game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>`. For Remote Location game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<location>/<ID string>`. Idempotency tokens remain in use for 30 days after a game session has ended; game session objects are retained for this time period and then deleted.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 48.

Pattern: [a-zA-Z0-9-]+

Required: No

Location

A fleet's remote location to place the new game session in. If this parameter is not set, the new game session is placed in the fleet's home Region. Specify a remote location with an AWS Region code such as `us-west-2`. When using an Anywhere fleet, this parameter is required and must be set to the Anywhere fleet's custom location.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: ^[A-Za-z0-9\-]+

Required: No

Name

A descriptive label that is associated with a game session. Session names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "GameSession": {
    "ComputeName": "string",
    "CreationTime": number,
    "CreatorId": "string",
    "CurrentPlayerSessionCount": number,
    "DnsName": "string",
    "FleetArn": "string",
    "FleetId": "string",
    "GameProperties": [
      {
        "Key": "string",
        "Value": "string"
      }
    ],
    "GameSessionData": "string",
    "GameSessionId": "string",
    "IpAddress": "string",
    "Location": "string",
    "MatchmakerData": "string",
    "MaximumPlayerSessionCount": number,
    "Name": "string",
    "PlayerGatewayStatus": "string",
    "PlayerSessionCreationPolicy": "string",
    "Port": number,
    "Status": "string",
    "StatusReason": "string",
    "TerminationTime": number
  }
}
```

```
}  
}
```

Response Elements

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

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

GameSession

Object that describes the newly created game session record.

Type: [GameSession](#) object

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

FleetCapacityExceededException

The specified fleet has no available instances to fulfill a `CreateGameSession` request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 400

IdempotentParameterMismatchException

A game session with this custom ID string already exists in this fleet. Resolve this conflict before retrying this request.

HTTP Status Code: 400

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidFleetStatusException

The requested operation would cause a conflict with the current state of a resource associated with the request and/or the fleet. Resolve the conflict before retrying.

HTTP Status Code: 400

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TerminalRoutingStrategyException

The service is unable to resolve the routing for a particular alias because it has a terminal `RoutingStrategy` associated with it. The message returned in this exception is the message defined in the routing strategy itself. Such requests should only be retried if the routing strategy for the specified alias is modified.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

CreateGameSessionQueue

This API works with the following fleet types: EC2, Anywhere, Container

Creates a placement queue that processes requests for new game sessions. A queue uses FleetIQ algorithms to locate the best available placement locations for a new game session, and then prompts the game server process to start a new game session.

A game session queue is configured with a set of destinations (Amazon GameLift Servers fleets or aliases) that determine where the queue can place new game sessions. These destinations can span multiple AWS Regions, can use different instance types, and can include both Spot and On-Demand fleets. If the queue includes multi-location fleets, the queue can place game sessions in any of a fleet's remote locations.

You can configure a queue to determine how it selects the best available placement for a new game session. Queues can prioritize placement decisions based on a combination of location, hosting cost, and player latency. You can set up the queue to use the default prioritization or provide alternate instructions using `PriorityConfiguration`.

Request options

Use this operation to make these common types of requests.

- Create a queue with the minimum required parameters.
 - Name
 - Destinations (This parameter isn't required, but a queue can't make placements without at least one destination.)
- Create a queue with placement notification. Queues that have high placement activity must use a notification system, such as with Amazon Simple Notification Service (Amazon SNS) or Amazon CloudWatch.
 - Required parameters Name and Destinations
 - NotificationTarget
- Create a queue with custom prioritization settings. These custom settings replace the default prioritization configuration for a queue.
 - Required parameters Name and Destinations
 - PriorityConfiguration
- Create a queue with special rules for processing player latency data.

- Required parameters Name and Destinations
- PlayerLatencyPolicies

Results

If successful, this operation returns a new `GameSessionQueue` object with an assigned queue ARN. Use the queue's name or ARN when submitting new game session requests with [StartGameSessionPlacement](#) or [StartMatchmaking](#).

Learn more

[Design a game session queue](#)

[Create a game session queue](#)

Related actions

[CreateGameSessionQueue](#) | [DescribeGameSessionQueues](#) | [UpdateGameSessionQueue](#) | [DeleteGameSessionQueue](#) | [All APIs by task](#)

Request Syntax

```
{
  "CustomEventData": "string",
  "Destinations": [
    {
      "DestinationArn": "string"
    }
  ],
  "FilterConfiguration": {
    "AllowedLocations": [ "string" ]
  },
  "Name": "string",
  "NotificationTarget": "string",
  "PlayerLatencyPolicies": [
    {
      "MaximumIndividualPlayerLatencyMilliseconds": number,
      "PolicyDurationSeconds": number
    }
  ],
  "PriorityConfiguration": {
```

```
    "LocationOrder": [ "string" ],
    "PriorityOrder": [ "string" ]
  },
  "Tags": [
    {
      "Key": "string",
      "Value": "string"
    }
  ],
  "TimeoutInSeconds": number
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Name

A descriptive label that is associated with game session queue. Queue names must be unique within each Region.

Type: String

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

Pattern: [a-zA-Z0-9-]+

Required: Yes

CustomEventData

Information to be added to all events that are related to this game session queue.

Type: String

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

Pattern: `[\s\S]*`

Required: No

Destinations

A list of fleets and/or fleet aliases that can be used to fulfill game session placement requests in the queue. Destinations are identified by either a fleet ARN or a fleet alias ARN, and are listed in order of placement preference.

Type: Array of [GameSessionQueueDestination](#) objects

Required: No

FilterConfiguration

A list of locations where a queue is allowed to place new game sessions. Locations are specified in the form of AWS Region codes, such as `us-west-2`. If this parameter is not set, game sessions can be placed in any queue location.

Type: [FilterConfiguration](#) object

Required: No

NotificationTarget

An SNS topic ARN that is set up to receive game session placement notifications. See [Setting up notifications for game session placement](#).

Type: String

Length Constraints: Minimum length of 0. Maximum length of 300.

Pattern: `[a-zA-Z0-9:_-]*(\.\fifo)?`

Required: No

PlayerLatencyPolicies

A set of policies that enforce a sliding cap on player latency when processing game sessions placement requests. Use multiple policies to gradually relax the cap over time if Amazon GameLift Servers can't make a placement. Policies are evaluated in order starting with the lowest maximum latency value.

Type: Array of [PlayerLatencyPolicy](#) objects

Required: No

PriorityConfiguration

Custom settings to use when prioritizing destinations and locations for game session placements. This configuration replaces the FleetIQ default prioritization process. Priority types that are not explicitly named will be automatically applied at the end of the prioritization process.

Type: [PriorityConfiguration](#) object

Required: No

Tags

A list of labels to assign to the new game session queue resource. Tags are developer-defined key-value pairs. Tagging AWS resources are useful for resource management, access management and cost allocation. For more information, see [Tagging AWS Resources](#) in the *AWS General Reference*.

Type: Array of [Tag](#) objects

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

Required: No

TimeoutInSeconds

The maximum time, in seconds, that a new game session placement request remains in the queue. When a request exceeds this time, the game session placement changes to a TIMED_OUT status. If you don't specify a request timeout, the queue uses a default value.

Note

The minimum value is 10 and the maximum value is 600.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

Response Syntax

```
{
  "GameSessionQueue": {
    "CustomEventData": "string",
    "Destinations": [
      {
        "DestinationArn": "string"
      }
    ],
    "FilterConfiguration": {
      "AllowedLocations": [ "string" ]
    },
    "GameSessionQueueArn": "string",
    "Name": "string",
    "NotificationTarget": "string",
    "PlayerLatencyPolicies": [
      {
        "MaximumIndividualPlayerLatencyMilliseconds": number,
        "PolicyDurationSeconds": number
      }
    ],
    "PriorityConfiguration": {
      "LocationOrder": [ "string" ],
      "PriorityOrder": [ "string" ]
    },
    "TimeoutInSeconds": number
  }
}
```

Response Elements

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

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

GameSessionQueue

An object that describes the newly created game session queue.

Type: GameSessionQueue object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Create and configure a game session queue

In this example, we want to create a game session queue with two single-location destinations, each residing in different Regions. We configure the queue so that requests for new game sessions expire after 10 minutes. The queue will process game session requests with player latency data, so we provide a set of latency cap polices that initially start at 100mm and then relax to 200ms after one minute. We opt to use the default sort approach for FleetIQ.

Sample Request

```
{
  "Name": "matchmaker-queue",
  "Destinations": [
    { "DestinationArn": "arn:aws:gamelift:us-west-2::fleet/
fleet-1a2b3c4d-5e6f-7a8b-9c0d-1e2f3a4b5c6d" },
    { "DestinationArn": "arn:aws:gamelift:sa-east-1::fleet/
fleet-5c6d3c4d-5e6f-7a8b-9a0b-1e2f3a4b5a2b" }
  ],
  "NotificationTarget": "arn:aws:sns:us-west-2:111122223333:My_Placement_SNS_Topic",
  "PlayerLatencyPolicies": [
    { "MaximumIndividualPlayerLatencyMilliseconds": 200 },
    { "MaximumIndividualPlayerLatencyMilliseconds": 100, "PolicyDurationSeconds":
60 }
  ],
  "TimeoutInSeconds": 600
}
```

Sample Response

```
{
  "GameSessionQueue": {
    "Name": "matchmaker-queue",
    "GameSessionQueueArn": "arn:aws:gamelift:us-
west-2:111122223333:gamesessionqueue/matchmaker-queue",
    "TimeoutInSeconds": 600,
    "NotificationTarget": "arn:aws:sns:us-
west-2:111122223333:My_Placement_SNS_Topic",
    "PlayerLatencyPolicies": [
      {
        "MaximumIndividualPlayerLatencyMilliseconds": 100,
```

```

        "PolicyDurationSeconds": 60
    },
    {
        "MaximumIndividualPlayerLatencyMilliseconds": 200
    }
],
"Destinations": [
    { "DestinationArn": "arn:aws:gamelift:us-west-2::fleet/
fleet-1a2b3c4d-5e6f-7a8b-9c0d-1e2f3a4b5c6d" },
    { "DestinationArn": "arn:aws:gamelift:us-east-1::fleet/
fleet-5c6d3c4d-5e6f-7a8b-9c0d-1e2f3a4b5a2b" }
]
}
}

```

Create a game session queue with multi-location fleets

In this example, we want to create a game session queue to place game sessions with two multi-location Spot fleets and one single-location On-Demand fleet. The multi-location fleets have remote locations in (us-west-1, us-east-2, sa-east-1).

We also want to change how FleetIQ prioritizes destinations and locations for placement. We opt to have FleetIQ prioritize by location first (with a custom location order provided), and then by destination list order. Based on our priority configuration, the queue will always try to place game sessions in the us-west-1 remote location of the first listed destination fleet. If no game servers are available there, the queue will try to place in the us-west-1 home Region of the second listed destination fleet, and so on.

Sample Request

```

{
    "Name": "matchmaker-queue",
    "Destinations": [
        { "DestinationArn": "arn:aws:gamelift:us-west-2::fleet/
fleet-1a2b3c4d-5e6f-7a8b-9c0d-1e2f3a4b5c6d" },
        { "DestinationArn": "arn:aws:gamelift:us-west-1::fleet/
fleet-2a3b4c5d-6e7f-8a9b-0c1d-2e3f4a5b6c7d" },
        { "DestinationArn": "arn:aws:gamelift:us-west-2::fleet/
fleet-5c6d3c4d-5e6f-7a8b-9a0b-1e2f3a4b5a2b" }
    ],
    "NotificationTarget": "arn:aws:sns:us-west-2:111122223333:My_Placement_SNS_Topic",
    "PriorityConfiguration": {

```

```
    "PriorityOrder": "LOCATION,DESTINATION",
    "LocationOrder": "us-west-1,us-west-2,us-east-2, sa-east-1"
  },
  "TimeoutInSeconds": 600
}
```

Sample Response

```
{
  "GameSessionQueue": {
    "Name": "matchmaker-queue",
    "GameSessionQueueArn": "arn:aws:gamelift:us-west-2:111122223333:gamesessionqueue/matchmaker-queue",
    "TimeoutInSeconds": 600,
    "Destinations": [
      { "DestinationArn": "arn:aws:gamelift:us-west-2::fleet/fleet-1a2b3c4d-5e6f-7a8b-9c0d-1e2f3a4b5c6d" },
      { "DestinationArn": "arn:aws:gamelift:us-west-1::fleet/fleet-2a3b4c5d-6e7f-8a9b-0c1d-2e3f4a5b6c7d" },
      { "DestinationArn": "arn:aws:gamelift:us-west-2::fleet/fleet-5c6d3c4d-5e6f-7a8b-9a0b-1e2f3a4b5a2b" }
    ],
    "PriorityConfiguration": {
      "PriorityOrder": "LOCATION,DESTINATION",
      "LocationOrder": "us-west-1,us-west-2,us-east-2, sa-east-1"
    },
    "NotificationTarget": "arn:aws:sns:us-west-2:111122223333:My_Placement_SNS_Topic",
  }
}
```

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](#)

CreateLocation

This API works with the following fleet types: Anywhere

Creates a custom location for use in an Anywhere fleet.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

LocationName

A descriptive name for the custom location.

Type: String

Length Constraints: Minimum length of 8. Maximum length of 64.

Pattern: ^custom-[A-Za-z0-9\ -]+

Required: Yes

Tags

A list of labels to assign to the new resource. Tags are developer-defined key-value pairs. Tagging AWS resources are useful for resource management, access management, and cost allocation. For more information, see [Tagging AWS Resources](#) in the *AWS General Reference*.

Type: Array of [Tag](#) objects

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

Required: No

Response Syntax

```
{
  "Location": {
    "LocationArn": "string",
    "LocationName": "string",
    "PingBeacon": {
      "UDPEndpoint": {
        "Domain": "string",
        "Port": number
      }
    }
  }
}
```

Response Elements

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

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

Location

The details of the custom location you created.

Type: [LocationModel](#) object

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

CreateMatchmakingConfiguration

This API works with the following fleet types: EC2, Anywhere, Container

Defines a new matchmaking configuration for use with FlexMatch. Whether you are using FlexMatch with Amazon GameLift Servers hosting or as a standalone matchmaking service, the matchmaking configuration sets out rules for matching players and forming teams. If you're also using Amazon GameLift Servers hosting, it defines how to start game sessions for each match. Your matchmaking system can use multiple configurations to handle different game scenarios. All matchmaking requests identify the matchmaking configuration to use and provide player attributes consistent with that configuration.

To create a matchmaking configuration, you must provide the following: configuration name and FlexMatch mode (with or without Amazon GameLift Servers hosting); a rule set that specifies how to evaluate players and find acceptable matches; whether player acceptance is required; and the maximum time allowed for a matchmaking attempt. When using FlexMatch with Amazon GameLift Servers hosting, you also need to identify the game session queue to use when starting a game session for the match.

In addition, you must set up an Amazon Simple Notification Service topic to receive matchmaking notifications. Provide the topic ARN in the matchmaking configuration.

Learn more

[Design a FlexMatch matchmaker](#)

[Set up FlexMatch event notification](#)

Request Syntax

```
{
  "AcceptanceRequired": boolean,
  "AcceptanceTimeoutSeconds": number,
  "AdditionalPlayerCount": number,
  "BackfillMode": "string",
  "CustomEventData": "string",
  "Description": "string",
  "FlexMatchMode": "string",
  "GameProperties": [
    {
      "Key": "string",
```

```
    "Value": "string"
  }
],
"GameSessionData": "string",
"GameSessionQueueArns": [ "string" ],
"Name": "string",
"NotificationTarget": "string",
"RequestTimeoutSeconds": number,
"RuleSetName": "string",
"Tags": [
  {
    "Key": "string",
    "Value": "string"
  }
]
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

AcceptanceRequired

A flag that determines whether a match that was created with this configuration must be accepted by the matched players. To require acceptance, set to TRUE. With this option enabled, matchmaking tickets use the status REQUIRES_ACCEPTANCE to indicate when a completed potential match is waiting for player acceptance.

Type: Boolean

Required: Yes

Name

A unique identifier for the matchmaking configuration. This name is used to identify the configuration associated with a matchmaking request or ticket.

Type: String

Length Constraints: Maximum length of 128.

Pattern: [a-zA-Z0-9-\.]*

Required: Yes

RequestTimeoutSeconds

The maximum duration, in seconds, that a matchmaking ticket can remain in process before timing out. Requests that fail due to timing out can be resubmitted as needed.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 43200.

Required: Yes

RuleSetName

A unique identifier for the matchmaking rule set to use with this configuration. You can use either the rule set name or ARN value. A matchmaking configuration can only use rule sets that are defined in the same Region.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: [a-zA-Z0-9-\.]*|^arn:.*:matchmakingruleset\[a-zA-Z0-9-\.]*

Required: Yes

AcceptanceTimeoutSeconds

The length of time (in seconds) to wait for players to accept a proposed match, if acceptance is required.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 600.

Required: No

AdditionalPlayerCount

The number of player slots in a match to keep open for future players. For example, if the configuration's rule set specifies a match for a single 12-person team, and the additional player

count is set to 2, only 10 players are selected for the match. This parameter is not used if `FlexMatchMode` is set to `STANDALONE`.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

BackfillMode

The method used to backfill game sessions that are created with this matchmaking configuration. Specify `MANUAL` when your game manages backfill requests manually or does not use the match backfill feature. Specify `AUTOMATIC` to have Amazon GameLift Servers create a backfill request whenever a game session has one or more open slots. Learn more about manual and automatic backfill in [Backfill Existing Games with FlexMatch](#). Automatic backfill is not available when `FlexMatchMode` is set to `STANDALONE`.

Type: String

Valid Values: `AUTOMATIC` | `MANUAL`

Required: No

CustomEventData

Information to be added to all events related to this matchmaking configuration.

Type: String

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

Required: No

Description

A human-readable description of the matchmaking configuration.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

FlexMatchMode

Indicates whether this matchmaking configuration is being used with Amazon GameLift Servers hosting or as a standalone matchmaking solution.

- **STANDALONE** - FlexMatch forms matches and returns match information, including players and team assignments, in a [MatchmakingSucceeded](#) event.
- **WITH_QUEUE** - FlexMatch forms matches and uses the specified Amazon GameLift Servers queue to start a game session for the match.

Type: String

Valid Values: STANDALONE | WITH_QUEUE

Required: No

GameProperties

A set of key-value pairs that can store custom data in a game session. For example: {"Key": "difficulty", "Value": "novice"}. This information is added to the new `GameSession` object that is created for a successful match. This parameter is not used if `FlexMatchMode` is set to `STANDALONE`.

Note

- Avoid using periods (".") in property keys if you plan to search for game sessions by properties. Property keys containing periods cannot be searched and will be filtered out from search results due to search index limitations.
- If you use `SearchGameSessions` API, there is a limit of 500 game property keys across all game sessions and all fleets per region. If the limit is exceeded, there will potentially be game session entries missing from `SearchGameSessions` API results.

Type: Array of [GameProperty](#) objects

Array Members: Maximum number of 16 items.

Required: No

GameSessionData

A set of custom game session properties, formatted as a single string value. This data is passed to a game server process with a request to start a new game session. For more information, see

[Start a game session](#). This information is added to the new `GameSession` object that is created for a successful match. This parameter is not used if `FlexMatchMode` is set to `STANDALONE`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 4096.

Required: No

[**GameSessionQueueArns**](#)

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers game session queue resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::gamesessionqueue/<queue name>`. Queues can be located in any Region. Queues are used to start new Amazon GameLift Servers-hosted game sessions for matches that are created with this matchmaking configuration. If `FlexMatchMode` is set to `STANDALONE`, do not set this parameter.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `[a-zA-Z0-9:/- ]+`

Required: No

[**NotificationTarget**](#)

An SNS topic ARN that is set up to receive matchmaking notifications. See [Setting up notifications for matchmaking](#) for more information.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 300.

Pattern: `[a-zA-Z0-9:_/- ]*(.fifo)?`

Required: No

[**Tags**](#)

A list of labels to assign to the new matchmaking configuration resource. Tags are developer-defined key-value pairs. Tagging AWS resources are useful for resource management, access

management and cost allocation. For more information, see [Tagging AWS Resources](#) in the *AWS General Reference*.

Type: Array of [Tag](#) objects

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

Required: No

Response Syntax

```
{
  "Configuration": {
    "AcceptanceRequired": boolean,
    "AcceptanceTimeoutSeconds": number,
    "AdditionalPlayerCount": number,
    "BackfillMode": "string",
    "ConfigurationArn": "string",
    "CreationTime": number,
    "CustomEventData": "string",
    "Description": "string",
    "FlexMatchMode": "string",
    "GameProperties": [
      {
        "Key": "string",
        "Value": "string"
      }
    ],
    "GameSessionData": "string",
    "GameSessionQueueArns": [ "string" ],
    "Name": "string",
    "NotificationTarget": "string",
    "RequestTimeoutSeconds": number,
    "RuleSetArn": "string",
    "RuleSetName": "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.

Configuration

Object that describes the newly created matchmaking configuration.

Type: [MatchmakingConfiguration](#) object

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Create a matchmaking configuration for use with GameLift managed hosting

The following example creates a matchmaking configuration for a game that is being hosted on Amazon GameLift Servers servers, identifying a game session queue and providing a set of game properties to be passed on to new game sessions. Player acceptance is required, with a 60-second timeout, and auto-backfill is enabled.

Sample Request

```
{
  "Name": "MM-1v1",
  "Description": "Head-to-head sprint format",
  "RequestTimeoutSeconds": 120,
  "RuleSetName": "mm_one_v_one",
  "NotificationTarget": "arn:aws:sns:us-west-2:111122223333:My_Matchmaking_SNS_Topic",
  "FlexMatchMode": "WITH_QUEUE",
  "GameSessionQueueArns": "arn:aws:gamelift:us-west-2:111122223333:gamesessionqueue/My_Game_Session_Queue_One",
  "AcceptanceRequired": true,
  "AcceptanceTimeoutSeconds": 60,
  "BackfillMode": "AUTOMATIC",
  "AdditionalPlayerCount": 8,
  "GameProperties": [
    {
      "Key": "level",
      "Value": "10"
    },
    {
      "Key": "difficulty",
      "Value": "hard"
    }
  ]
}
```

```
}
```

Sample Response

```
{
  "Configuration": {
    "AcceptanceRequired": true,
    "AcceptanceTimeoutSeconds": 60,
    "AdditionalPlayerCount": 8,
    "BackfillMode": "AUTOMATIC",
    "ConfigurationArn": "arn:aws:gamelift:us-west-2:111122223333:matchmakingconfiguration/MM-1v1",
    "CreationTime": 1496365885.44,
    "Description": "Head-to-head sprint format",
    "FlexMatchMode": "WITH_QUEUE",
    "GameProperties": [
      {
        "Key": "level",
        "Value": "10"
      },
      {
        "Key": "difficulty",
        "Value": "hard"
      }
    ],
    "GameSessionQueueArns": "arn:aws:gamelift:us-west-2:111122223333:gamesessionqueue/My_Game_Session_Queue_One",
    "Name": "MM-1v1",
    "NotificationTarget": "arn:aws:sns:us-west-2:111122223333:My_Matchmaking_SNS_Topic",
    "RequestTimeoutSeconds": 120,
    "RuleSetArn": "arn:aws:gamelift:us-west-2:111122223333:matchmakingruleset/mm_one_v_one",
    "RuleSetName": "mm_one_v_one"
  }
}
```

Create a matchmaking configuration for a standalone FlexMatch system

The following example creates a matchmaking configuration for a game that is hosted on resources other than Amazon GameLift Servers game servers. This includes games that are hosted on Amazon EC2 with Amazon GameLift Servers FleetIQ. This configuration omits the game session

queue, game properties, and additional player count. Player acceptance is required, with a 60-second timeout.

Sample Request

```
{
  "Name": "MM-1v1",
  "Description": "Head-to-head sprint format",
  "RequestTimeoutSeconds": 120,
  "RuleSetName": "mm_one_v_one",
  "NotificationTarget": "arn:aws:sns:us-
west-2:111122223333:My_Matchmaking_SNS_Topic",
  "FlexMatchMode": "STANDALONE",
  "AcceptanceRequired": true,
  "AcceptanceTimeoutSeconds": 60,
  "BackfillMode": "MANUAL"
}
```

Sample Response

```
{
  "Configuration": {
    "AcceptanceRequired": true,
    "AcceptanceTimeoutSeconds": 60,
    "AdditionalPlayerCount": 0,
    "BackfillMode": "MANUAL",
    "ConfigurationArn": "arn:aws:gamelift:us-
west-2:111122223333:matchmakingconfiguration/MM-1v1",
    "CreationTime": 1496365885.44,
    "Description": "Head-to-head sprint format",
    "FlexMatchMode": "STANDALONE",
    "GameSessionQueueArns": "",
    "Name": "MM-1v1",
    "NotificationTarget": "arn:aws:sns:us-
west-2:111122223333:My_Matchmaking_SNS_Topic",
    "RequestTimeoutSeconds": 120,
    "RuleSetArn": "arn:aws:gamelift:us-west-2:111122223333:matchmakingruleset/
mm_one_v_one",
    "RuleSetName": "mm_one_v_one"
  }
}
```

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](#)

CreateMatchmakingRuleSet

This API works with the following fleet types: EC2, Anywhere, Container

Creates a new rule set for FlexMatch matchmaking. A rule set describes the type of match to create, such as the number and size of teams. It also sets the parameters for acceptable player matches, such as minimum skill level or character type.

To create a matchmaking rule set, provide unique rule set name and the rule set body in JSON format. Rule sets must be defined in the same Region as the matchmaking configuration they are used with.

Since matchmaking rule sets cannot be edited, it is a good idea to check the rule set syntax using [ValidateMatchmakingRuleSet](#) before creating a new rule set.

Learn more

- [Build a rule set](#)
- [Design a matchmaker](#)
- [Matchmaking with FlexMatch](#)

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Name

A unique identifier for the matchmaking rule set. A matchmaking configuration identifies the rule set it uses by this name value. Note that the rule set name is different from the optional name field in the rule set body.

Type: String

Length Constraints: Maximum length of 128.

Pattern: `[a-zA-Z0-9-\.]*`

Required: Yes

RuleSetBody

A collection of matchmaking rules, formatted as a JSON string. Comments are not allowed in JSON, but most elements support a description field.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 65535.

Required: Yes

Tags

A list of labels to assign to the new matchmaking rule set resource. Tags are developer-defined key-value pairs. Tagging AWS resources are useful for resource management, access management and cost allocation. For more information, see [Tagging AWS Resources](#) in the *AWS General Reference*.

Type: Array of [Tag](#) objects

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

Required: No

Response Syntax

```
{
  "RuleSet": {
    "CreationTime": number,
    "RuleSetArn": "string",
    "RuleSetBody": "string",
    "RuleSetName": "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.

RuleSet

The newly created matchmaking rule set.

Type: [MatchmakingRuleSet](#) object

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

CreatePlayerSession

This API works with the following fleet types: EC2, Anywhere, Container

Reserves an open player slot in a game session for a player. New player sessions can be created in any game session with an open slot that is in ACTIVE status and has a player creation policy of ACCEPT_ALL. You can add a group of players to a game session with [CreatePlayerSessions](#).

To create a player session, specify a game session ID, player ID, and optionally a set of player data.

If successful, a slot is reserved in the game session for the player and a new `PlayerSessions` object is returned with a player session ID. The player references the player session ID when sending a connection request to the game session, and the game server can use it to validate the player reservation with the Amazon GameLift Servers service. Player sessions cannot be updated.

The maximum number of players per game session is 200. It is not adjustable.

Related actions

[All APIs by task](#)

Request Syntax

```
{
  "GameSessionId": "string",
  "PlayerData": "string",
  "PlayerId": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

GameSessionId

An identifier for the game session that is unique across all regions to add a player to. The value is always a full ARN in the following format: For Home Region game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>`. For Remote Location game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<location>/<ID string>`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `[a-zA-Z0-9:/-]+`

Required: Yes

PlayerId

A unique identifier for a player. Player IDs are developer-defined.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

PlayerData

Developer-defined information related to a player. Amazon GameLift Servers does not use this data, so it can be formatted as needed for use in the game.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

Response Syntax

```
{
  "PlayerSession": {
    "CreationTime": number,
    "DnsName": "string",
```

```
"FleetArn": "string",  
"FleetId": "string",  
"GameSessionId": "string",  
"IpAddress": "string",  
"PlayerData": "string",  
"PlayerId": "string",  
"PlayerSessionId": "string",  
"Port": number,  
"Status": "string",  
"TerminationTime": number  
}  
}
```

Response Elements

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

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

PlayerSession

Object that describes the newly created player session record.

Type: [PlayerSession](#) object

Errors

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

GameSessionFullException

The game instance is currently full and cannot allow the requested player(s) to join. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 400

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidGameSessionStatusException

The requested operation would cause a conflict with the current state of a resource associated with the request and/or the game instance. Resolve the conflict before retrying.

HTTP Status Code: 400

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TerminalRoutingStrategyException

The service is unable to resolve the routing for a particular alias because it has a terminal RoutingStrategy associated with it. The message returned in this exception is the message defined in the routing strategy itself. Such requests should only be retried if the routing strategy for the specified alias is modified.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

CreatePlayerSessions

This API works with the following fleet types: EC2, Anywhere, Container

Reserves open slots in a game session for a group of players. New player sessions can be created in any game session with an open slot that is in ACTIVE status and has a player creation policy of ACCEPT_ALL. To add a single player to a game session, use [CreatePlayerSession](#)

To create player sessions, specify a game session ID and a list of player IDs. Optionally, provide a set of player data for each player ID.

If successful, a slot is reserved in the game session for each player, and new `PlayerSession` objects are returned with player session IDs. Each player references their player session ID when sending a connection request to the game session, and the game server can use it to validate the player reservation with the Amazon GameLift Servers service. Player sessions cannot be updated.

The maximum number of players per game session is 200. It is not adjustable.

Related actions

[All APIs by task](#)

Request Syntax

```
{
  "GameSessionId": "string",
  "PlayerDataMap": {
    "string" : "string"
  },
  "PlayerIds": [ "string" ]
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

GameSessionId

An identifier for the game session that is unique across all regions to add players to.

The value is always a full ARN in the following format: For Home Region game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>`. For Remote Location game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<location>/<ID string>`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `[a-zA-Z0-9:/- ]+`

Required: Yes

PlayerIds

List of unique identifiers for the players to be added.

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 1024.

Required: Yes

PlayerDataMap

Map of string pairs, each specifying a player ID and a set of developer-defined information related to the player. Amazon GameLift Servers does not use this data, so it can be formatted as needed for use in the game. Any player data strings for player IDs that are not included in the `PlayerIds` parameter are ignored.

Type: String to string map

Key Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

Response Syntax

```
{
  "PlayerSessions": [
    {
      "CreationTime": number,
      "DnsName": "string",
      "FleetArn": "string",
      "FleetId": "string",
      "GameSessionId": "string",
      "IpAddress": "string",
      "PlayerData": "string",
      "PlayerId": "string",
      "PlayerSessionId": "string",
      "Port": number,
      "Status": "string",
      "TerminationTime": number
    }
  ]
}
```

Response Elements

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

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

PlayerSessions

A collection of player session objects created for the added players.

Type: Array of [PlayerSession](#) objects

Errors

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

GameSessionFullException

The game instance is currently full and cannot allow the requested player(s) to join. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 400

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidGameSessionStatusException

The requested operation would cause a conflict with the current state of a resource associated with the request and/or the game instance. Resolve the conflict before retrying.

HTTP Status Code: 400

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TerminalRoutingStrategyException

The service is unable to resolve the routing for a particular alias because it has a terminal RoutingStrategy associated with it. The message returned in this exception is the message defined in the routing strategy itself. Such requests should only be retried if the routing strategy for the specified alias is modified.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

CreateScript

This API works with the following fleet types: EC2, Anywhere

Creates a new script record for your Amazon GameLift Servers Realtime script. Realtime scripts are JavaScript that provide configuration settings and optional custom game logic for your game. The script is deployed when you create a Amazon GameLift Servers Realtime fleet to host your game sessions. Script logic is executed during an active game session.

To create a new script record, specify a script name and provide the script file(s). The script files and all dependencies must be zipped into a single file. You can pull the zip file from either of these locations:

- A locally available directory. Use the *ZipFile* parameter for this option.
- An Amazon Simple Storage Service (Amazon S3) bucket under your AWS account. Use the *StorageLocation* parameter for this option. You'll need to have an Identity Access Management (IAM) role that allows the Amazon GameLift Servers service to access your S3 bucket.

If the call is successful, a new script record is created with a unique script ID. If the script file is provided as a local file, the file is uploaded to an Amazon GameLift Servers-owned S3 bucket and the script record's storage location reflects this location. If the script file is provided as an S3 bucket, Amazon GameLift Servers accesses the file at this storage location as needed for deployment.

Learn more

[Amazon GameLift Servers Amazon GameLift Servers Realtime](#)

[Set Up a Role for Amazon GameLift Servers Access](#)

Related actions

[All APIs by task](#)

Request Syntax

```
{  
  "Name": "string",  
  "NodeJsVersion": "string",
```

```
"StorageLocation": {
  "Bucket": "string",
  "Key": "string",
  "ObjectVersion": "string",
  "RoleArn": "string"
},
"Tags": [
  {
    "Key": "string",
    "Value": "string"
  }
],
"Version": "string",
"ZipFile": blob
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Name

A descriptive label that is associated with a script. Script names do not need to be unique. You can use [UpdateScript](#) to change this value later.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

NodeJsVersion

The Node.js version used for execution of your Realtime script. The valid values are 10.x | 24.x. By default, NodeJsVersion is 10.x. This value cannot be updated later.

Note

Node.js 10 will reach end of support on September 30, 2026. See more details in the [Node.js 10 FAQs](#). For migration guidance, see [Migrating from Node.js 10 to 24](#).

Type: String

Length Constraints: Maximum length of 16.

Pattern: `^[0-9]+\.(?:x|[0-9]+)$`

Required: No

StorageLocation

The location of the Amazon S3 bucket where a zipped file containing your Realtime scripts is stored. The storage location must specify the Amazon S3 bucket name, the zip file name (the "key"), and a role ARN that allows Amazon GameLift Servers to access the Amazon S3 storage location. The S3 bucket must be in the same Region where you want to create a new script. By default, Amazon GameLift Servers uploads the latest version of the zip file; if you have S3 object versioning turned on, you can use the `ObjectVersion` parameter to specify an earlier version.

Type: [S3Location](#) object

Required: No

Tags

A list of labels to assign to the new script resource. Tags are developer-defined key-value pairs. Tagging AWS resources are useful for resource management, access management and cost allocation. For more information, see [Tagging AWS Resources](#) in the *AWS General Reference*. Once the resource is created, you can use [TagResource](#), [UntagResource](#), and [ListTagsForResource](#) to add, remove, and view tags. The maximum tag limit may be lower than stated. See the AWS General Reference for actual tagging limits.

Type: Array of [Tag](#) objects

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

Required: No

Version

Version information that is associated with a build or script. Version strings do not need to be unique. You can use [UpdateScript](#) to change this value later.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

ZipFile

A data object containing your Realtime scripts and dependencies as a zip file. The zip file can have one or multiple files. Maximum size of a zip file is 5 MB.

When using the AWS CLI tool to create a script, this parameter is set to the zip file name. It must be prepended with the string "fileb://" to indicate that the file data is a binary object. For example: `--zip-file fileb://myRealtimeScript.zip`.

Type: Base64-encoded binary data object

Length Constraints: Maximum length of 5000000.

Required: No

Response Syntax

```
{
  "Script": {
    "CreationTime": number,
    "Name": "string",
    "NodeJsVersion": "string",
    "ScriptArn": "string",
    "ScriptId": "string",
    "SizeOnDisk": number,
    "StorageLocation": {
      "Bucket": "string",
      "Key": "string",
      "ObjectVersion": "string",
      "RoleArn": "string"
    },
  },
  "Version": "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.

[Script](#)

The newly created script record with a unique script ID and ARN. The new script's storage location reflects an Amazon S3 location: (1) If the script was uploaded from an S3 bucket under your account, the storage location reflects the information that was provided in the *CreateScript* request; (2) If the script file was uploaded from a local zip file, the storage location reflects an S3 location controls by the Amazon GameLift Servers service.

Type: [Script](#) object

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Create a script from a local zip file

This example creates a script record and uploads a Realtime script from a zip file that is stored locally.

Sample Request

```
{
  "Name": "My_Realtime_Script_1",
  "Version": "1.0.1",
  "NodeJsVersion": "24.x",
  "ZipFile": <zip file data>
}
```

CLI syntax:

```
aws gamelift create-script
  --name "My_Realtime_Script_1"
  --node-js-version "24.x"
  --script-version "1.0.1"
  --zip-file fileb://myrealtime_script.zip
```

Sample Response

```
{
```

```
"Script": {
  "CreationTime": 1496708916.18,
  "Name": "My_Realtime_Script_1",
  "NodeJsVersion": "24.x",
  "ScriptArn": "arn:aws:gamelift:us-west-2::script/
script-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "ScriptId": "script-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "SizeOnDisk": 9000,
  "StorageLocation": {
    "Bucket": "prod-gamescale-scripts-us-west-2",
    "Key": "123456789012/script-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
  },
  "Version": "1.0.1"
}
```

Create a script with a file in Amazon S3

This example creates a script record and uploads a Realtime server script from a zip file that is stored in an Amazon S3 account.

Sample Request

```
{
  "Name": "My_Realtime_Script_2",
  "Version": "12345.678",
  "NodeJsVersion": "24.x",
  "StorageLocation": {
    "Bucket": "my_realtime_script_files",
    "Key": "myRealtimeScript.zip",
    "RoleArn": "arn:aws:iam::111122223333:role/GameLiftAccess"
  }
}
```

CLI syntax:

```
aws gamelift create-script
--name "My_Realtime_Script_2"
--node-js-version "24.x"
--script-version "12345.678"
--storage-location
"Bucket=my_realtime_script_files,
Key=myRealtimeScript.zip,
```

```
RoleArn=arn:aws:iam::123456789012:role/GameLiftAccess"
```

Sample Response

```
{
  "Script": {
    "CreationTime": 1496708916.18,
    "Name": "My_Realtime_Script_2",
    "NodeJsVersion": "24.x",
    "ScriptId": "script-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "ScriptArn": "arn:aws:gamelift:us-west-2::script/
script-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "SizeOnDisk": 0,
    "StorageLocation": {
      "Bucket": "my_realtime_script_files",
      "Key": "myRealtimeScript.zip"
      "RoleArn": "arn:aws:iam::111122223333:role/GameLiftAccess"
      "ObjectVersion": null
    },
    "Version": "12345.678"
  }
}
```

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](#)

CreateVpcPeeringAuthorization

This API works with the following fleet types: EC2

Requests authorization to create or delete a peer connection between the VPC for your Amazon GameLift Servers fleet and a virtual private cloud (VPC) in your AWS account. VPC peering enables the game servers on your fleet to communicate directly with other AWS resources. After you've received authorization, use [CreateVpcPeeringConnection](#) to establish the peering connection. For more information, see [VPC Peering with Amazon GameLift Servers Fleets](#).

You can peer with VPCs that are owned by any AWS account you have access to, including the account that you use to manage your Amazon GameLift Servers fleets. You cannot peer with VPCs that are in different Regions.

To request authorization to create a connection, call this operation from the AWS account with the VPC that you want to peer to your Amazon GameLift Servers fleet. For example, to enable your game servers to retrieve data from a DynamoDB table, use the account that manages that DynamoDB resource. Identify the following values: (1) The ID of the VPC that you want to peer with, and (2) the ID of the AWS account that you use to manage Amazon GameLift Servers. If successful, VPC peering is authorized for the specified VPC.

To request authorization to delete a connection, call this operation from the AWS account with the VPC that is peered with your Amazon GameLift Servers fleet. Identify the following values: (1) VPC ID that you want to delete the peering connection for, and (2) ID of the AWS account that you use to manage Amazon GameLift Servers.

The authorization remains valid for 24 hours unless it is canceled. You must create or delete the peering connection while the authorization is valid.

Note

Amazon GameLift Servers uses the caller's credentials to update peer-VPC resources. The IAM user that calls this operation must have the following Amazon EC2 permissions enabled:

- `ec2:AcceptVpcPeeringConnection`
- `ec2:AuthorizeSecurityGroupEgress`
- `ec2:AuthorizeSecurityGroupIngress`
- `ec2:CreateRoute`

- `ec2:DescribeRouteTables`
- `ec2:DescribeSecurityGroups`

Related actions

[All APIs by task](#)

Request Syntax

```
{  
  "GameLiftAwsAccountId": "string",  
  "PeerVpcId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[GameLiftAwsAccountId](#)

A unique identifier for the AWS account that you use to manage your Amazon GameLift Servers fleet. You can find your Account ID in the AWS Management Console under account settings.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

[PeerVpcId](#)

A unique identifier for a VPC with resources to be accessed by your Amazon GameLift Servers fleet. The VPC must be in the same Region as your fleet. To look up a VPC ID, use the [VPC](#)

[Dashboard](#) in the AWS Management Console. Learn more about VPC peering in [VPC Peering with Amazon GameLift Servers Fleets](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

Response Syntax

```
{
  "VpcPeeringAuthorization": {
    "CreationTime": number,
    "ExpirationTime": number,
    "GameLiftAwsAccountId": "string",
    "PeerVpcAwsAccountId": "string",
    "PeerVpcId": "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.

[VpcPeeringAuthorization](#)

Details on the requested VPC peering authorization, including expiration.

Type: [VpcPeeringAuthorization](#) object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Authorize VPC peering between your Amazon GameLift Servers fleet and resources on your Amazon GameLift Servers account

In this example, you want your Amazon GameLift Servers hosted game servers to access a web service. You manage the Amazon GameLift Servers fleet and the web service through the same AWS account (account ID *111122223333*). The web service already has a VPC set up, with ID *vpc-a12bc345*.

When making this request, use credentials for AWS account *111122223333*.

Sample Request

```
POST / HTTP/1.1
Host: gamelift.us-west-2.amazonaws.com;
Accept-Encoding: identity
Content-Length: 77
User-Agent: aws-cli/1.11.36 Python/2.7.9 Windows/7 botocore/1.4.93
Content-Type: application/x-amz-json-1.0
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20170406/us-west-2/
gamelift/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target,
Signature=wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
```

```
X-Amz-Date: 20170406T004805Z
X-Amz-Target: GameLift.CreateVpcPeeringAuthorization

{
  "GameLiftAwsAccountId": "111122223333",
  "PeerVpcId": "vpc-a12bc345"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: b34f8665-EXAMPLE
Content-Type: application/x-amz-json-1.1
Content-Length: 225
Date: Thu, 06 Apr 2017 00:48:07 GMT

{"VpcPeeringAuthorization":
  {"CreationTime": 1503608847.489,
    "ExpirationTime": 1503695247,
    "GameLiftAwsAccountId": "111122223333",
    "PeerVpcAwsAccountId": "111122223333",
    "PeerVpcId": "vpc-a12bc345"}
}
```

Authorize VPC peering between your Amazon GameLift Servers fleet and resources on a different account

As in the previous example, you want your game servers to access a web service. But in this example, the Amazon GameLift Servers fleet and the web service are managed through different AWS accounts. Your Amazon GameLift Servers account ID is `111122223333`, while the web service account ID is `444455556666`. A VPC on account `444455556666` with the web service is set up with the ID `vpc-c67ef890`.

When making this request, use credentials for AWS account `444455556666`. If you don't have rights to this account, work with the account owner to make the request. You'll need to provide your Amazon GameLift Servers account ID.

Sample Request

```
POST / HTTP/1.1
Host: gamelift.us-west-2.amazonaws.com;
Accept-Encoding: identity
```

```
Content-Length: 82
User-Agent: aws-cli/1.11.36 Python/2.7.9 Windows/7 botocore/1.4.93
Content-Type: application/x-amz-json-1.0
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20170406/us-west-2/
gamelift/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target,
Signature=wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
X-Amz-Date: 20170406T004805Z
X-Amz-Target: GameLift.CreateVpcPeeringAuthorization

{
  "GameLiftAwsAccountId": "111122223333",
  "PeerVpcId": "vpc-c67ef890"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: b34f8665-EXAMPLE
Content-Type: application/x-amz-json-1.1
Content-Length: 225
Date: Thu, 06 Apr 2017 00:48:07 GMT

{"VpcPeeringAuthorization":
  {"CreationTime": 1503608847.489,
    "ExpirationTime": 1503695247,
    "GameLiftAwsAccountId": "111122223333",
    "PeerVpcAwsAccountId": "444455556666",
    "PeerVpcId": "vpc-c67ef890"}
}
```

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](#)

CreateVpcPeeringConnection

This API works with the following fleet types: EC2

Establishes a VPC peering connection between a virtual private cloud (VPC) in an AWS account with the VPC for your Amazon GameLift Servers fleet. VPC peering enables the game servers on your fleet to communicate directly with other AWS resources. You can peer with VPCs in any AWS account that you have access to, including the account that you use to manage your Amazon GameLift Servers fleets. You cannot peer with VPCs that are in different Regions. For more information, see [VPC Peering with Amazon GameLift Servers Fleets](#).

Before calling this operation to establish the peering connection, you first need to use [CreateVpcPeeringAuthorization](#) and identify the VPC you want to peer with. Once the authorization for the specified VPC is issued, you have 24 hours to establish the connection. These two operations handle all tasks necessary to peer the two VPCs, including acceptance, updating routing tables, etc.

To establish the connection, call this operation from the AWS account that is used to manage the Amazon GameLift Servers fleets. Identify the following values: (1) The ID of the fleet you want to be enable a VPC peering connection for; (2) The AWS account with the VPC that you want to peer with; and (3) The ID of the VPC you want to peer with. This operation is asynchronous. If successful, a connection request is created. You can use continuous polling to track the request's status using [DescribeVpcPeeringConnections](#) , or by monitoring fleet events for success or failure using [DescribeFleetEvents](#) .

Note

Amazon GameLift Servers uses the caller's credentials to update peer-VPC resources. The IAM user that calls this operation must have the following Amazon EC2 permissions enabled:

- `ec2:AcceptVpcPeeringConnection`
- `ec2:AuthorizeSecurityGroupEgress`
- `ec2:AuthorizeSecurityGroupIngress`
- `ec2:CreateRoute`
- `ec2:DescribeRouteTables`
- `ec2:DescribeSecurityGroups`

Related actions

[All APIs by task](#)

Request Syntax

```
{  
  "FleetId": "string",  
  "PeerVpcAwsAccountId": "string",  
  "PeerVpcId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet. You can use either the fleet ID or ARN value. This tells Amazon GameLift Servers which GameLift VPC to peer with.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\-]+`

Required: Yes

PeerVpcAwsAccountId

A unique identifier for the AWS account with the VPC that you want to peer your Amazon GameLift Servers fleet with. You can find your Account ID in the AWS Management Console under account settings.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

PeerVpcId

A unique identifier for a VPC with resources to be accessed by your Amazon GameLift Servers fleet. The VPC must be in the same Region as your fleet. To look up a VPC ID, use the [VPC Dashboard](#) in the AWS Management Console. Learn more about VPC peering in [VPC Peering with Amazon GameLift Servers Fleets](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

Response Elements

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

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Peer the VPC for your Amazon GameLift Servers fleet with a VPC on your Amazon GameLift Servers account

This example builds on Example 1 in [CreateVpcPeeringAuthorization](#). If authorization succeeded, the next step is to tell Amazon GameLift Servers to request the peering connection. In this example, you want your game servers that are running on an Amazon GameLift Servers fleet to be able to access a web service. The web service is managed through the same AWS account that you use to manage your Amazon GameLift Servers fleet. To request the peering, provide the following details for the two VPCs to peer: (1) the Amazon GameLift Servers fleet ID, and (2) the account and VPC for the web service. The account ID and VPC for the web service must be the same one you used in the authorization.

The fleet's ID is *fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa*. The AWS account (the same as the one you use to manage your Amazon GameLift Servers fleet) is *111122223333*. The VPC ID for the web service is *vpc-a12bc345*.

To make this request, sign in using your credentials for AWS account 111122223333.

Sample Request

```
POST / HTTP/1.1
Host: gamelift.us-west-2.amazonaws.com;
Accept-Encoding: identity
Content-Length: 141
User-Agent: aws-cli/1.11.36 Python/2.7.9 Windows/7 botocore/1.4.93
Content-Type: application/x-amz-json-1.0
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20170406/us-west-2/
gamelift/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target,
Signature=wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY
X-Amz-Date: 20170406T004805Z
X-Amz-Target: GameLift.CreateVpcPeeringConnection
```

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "PeerVpcAwsAccountId": "111122223333",
  "PeerVpcId": "vpc-a12bc345"}
```

Peer the VPC for your Amazon GameLift Servers fleet with a VPC on a different account

This example builds on Example 2 in [CreateVpcPeeringAuthorization](#). If authorization succeeded, the next step is to tell Amazon GameLift Servers to request the peering connection. As in the previous example, you want your game servers that are running on an Amazon GameLift Servers fleet to be able to access a web service. But in this example, the web service is managed by a different account from the one that you use to manage your Amazon GameLift Servers fleet. To request the peering, provide the following details for the two VPCs to peer: (1) the Amazon GameLift Servers fleet ID, and (2) the account and VPC for the web service. The account ID and VPC for the web service must be the same one you used in the authorization.

The fleet's ID is *fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa*. Your Amazon GameLift Servers account ID is *111122223333*. The AWS account with the web service is *444455556666*. The VPC ID for the web service is *vpc-c67ef890*.

To make this request, sign in using your credentials for AWS account 111122223333 (your Amazon GameLift Servers account).

Sample Request

```
POST / HTTP/1.1
Host: gamelift.us-west-2.amazonaws.com;
Accept-Encoding: identity
Content-Length: 141
User-Agent: aws-cli/1.11.36 Python/2.7.9 Windows/7 botocore/1.4.93
Content-Type: application/x-amz-json-1.0
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20170406/us-west-2/gamelift/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
X-Amz-Date: 20170406T004805Z
X-Amz-Target: GameLift.CreateVpcPeeringConnection

{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
```

```
"PeerVpcAwsAccountId": "444455556666",  
"PeerVpcId": "vpc-c67ef890"}
```

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](#)

DeleteAlias

This API works with the following fleet types: EC2, Anywhere, Container

Deletes an alias. This operation removes all record of the alias. Game clients attempting to access a server process using the deleted alias receive an error. To delete an alias, specify the alias ID to be deleted.

Related actions

[All APIs by task](#)

Request Syntax

```
{
  "AliasId": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[AliasId](#)

A unique identifier of the alias that you want to delete. You can use either the alias ID or ARN value.

Type: String

Pattern: `^alias-\S+|^arn:.*:alias\/alias-\S+`

Required: Yes

Response Elements

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

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

DeleteBuild

This API works with the following fleet types: EC2

Deletes a build. This operation permanently deletes the build resource and any uploaded build files. Deleting a build does not affect the status of any active fleets using the build, but you can no longer create new fleets with the deleted build.

To delete a build, specify the build ID.

Learn more

[Upload a Custom Server Build](#)

[All APIs by task](#)

Request Syntax

```
{  
  "BuildId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[BuildId](#)

A unique identifier for the build to delete. You can use either the build ID or ARN value.

Type: String

Pattern: ^build-\S+|^arn:.*:build\/build-\S+

Required: Yes

Response Elements

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

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Remove a build

This example removes a build from your Amazon GameLift Servers account. After the build is deleted, you can no longer use it to create new fleets. This operation can't be undone.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "BuildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
}
```

Sample Response

```
HTTP/1.1 200 OK undefined
```

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](#)

DeleteContainerFleet

This API works with the following fleet types: Container

Deletes all resources and information related to a container fleet and shuts down currently running fleet instances, including those in remote locations. The container fleet must be in ACTIVE status to be deleted.

To delete a fleet, specify the fleet ID to be terminated. During the deletion process, the fleet status is changed to DELETING.

Learn more

[Setting up Amazon GameLift Servers Fleets](#)

Request Syntax

```
{
  "FleetId": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[FleetId](#)

A unique identifier for the container fleet to delete. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\`

Required: Yes

Response Elements

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

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Delete a fleet that is no longer in use

This example deletes an existing fleet.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "FleetId": "containerfleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: b34f8665-EXAMPLE
Date: Thu, 06 Apr 2017 00:48:07 GMT
```

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](#)

DeleteContainerGroupDefinition

This API works with the following fleet types: Container

Request options:

Deletes a container group definition.

- Delete an entire container group definition, including all versions. Specify the container group definition name, or use an ARN value without the version number.
- Delete a particular version. Specify the container group definition name and a version number, or use an ARN value that includes the version number.
- Keep the newest versions and delete all older versions. Specify the container group definition name and the number of versions to retain. For example, set `VersionCountToRetain` to 5 to delete all but the five most recent versions.

Result

If successful, Amazon GameLift Servers removes the container group definition versions that you request deletion for. This request will fail for any requested versions if the following is true:

- If the version is being used in an active fleet
- If the version is being deployed to a fleet in a deployment that's currently in progress.
- If the version is designated as a rollback definition in a fleet deployment that's currently in progress.

Learn more

- [Manage a container group definition](#)

Request Syntax

```
{
  "Name": "string",
  "VersionCountToRetain": number,
  "VersionNumber": number
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Name

The unique identifier for the container group definition to delete. You can use either the Name or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-zA-Z0-9\-\_]+\$|^arn:.*:containergroupdefinition\[a-zA-Z0-9\-\_]+(:[0-9]+)?\$`

Required: Yes

VersionCountToRetain

The number of most recent versions to keep while deleting all older versions.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

VersionNumber

The specific version to delete.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

Response Elements

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

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Delete a container group definition version

This example deletes the 8th version of a container group definition. It assumes we're working in the same AWS Region as the container group we want to delete. To delete a container group definition in a different region, we could make the request using the definition's ARN value or we could specify the other region.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "Name": "MyAdventureGameContainerGroup",
  "Version": 8
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: b34f8665-EXAMPLE
Date: Thu, 14 Apr 2024 00:48:07 GMT
```

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](#)

DeleteFleet

This API works with the following fleet types: EC2, Anywhere

Deletes all resources and information related to a fleet and shuts down any currently running fleet instances, including those in remote locations.

Note

If the fleet being deleted has a VPC peering connection, you first need to get a valid authorization (good for 24 hours) by calling [CreateVpcPeeringAuthorization](#). You don't need to explicitly delete the VPC peering connection.

To delete a fleet, specify the fleet ID to be terminated. During the deletion process, the fleet status is changed to DELETING. When completed, the status switches to TERMINATED and the fleet event FLEET_DELETED is emitted.

Learn more

[Setting up Amazon GameLift Servers Fleets](#)

Request Syntax

```
{  
  "FleetId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet to be deleted. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

Response Elements

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

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidFleetStatusException

The requested operation would cause a conflict with the current state of a resource associated with the request and/or the fleet. Resolve the conflict before retrying.

HTTP Status Code: 400

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Delete a fleet that is no longer in use

This example attempts to delete an existing fleet.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: b34f8665-EXAMPLE
Date: Thu, 06 Apr 2017 00:48:07 GMT
```

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](#)

DeleteFleetLocations

This API works with the following fleet types: EC2, Anywhere, Container

Removes locations from a multi-location fleet. When deleting a location, all game server process and all instances that are still active in the location are shut down.

To delete fleet locations, identify the fleet ID and provide a list of the locations to be deleted.

If successful, GameLift sets the location status to DELETING, and begins to shut down existing server processes and terminate instances in each location being deleted. When completed, the location status changes to TERMINATED.

Learn more

[Setting up Amazon GameLift Servers fleets](#)

Request Syntax

```
{
  "FleetId": "string",
  "Locations": [ "string" ]
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet to delete locations for. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\`

Required: Yes

Locations

The list of fleet locations to delete. Specify locations in the form of an AWS Region code, such as `us-west-2`.

Type: Array of strings

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

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-\_]+\`

Required: Yes

Response Syntax

```
{
  "FleetArn": "string",
  "FleetId": "string",
  "LocationStates": [
    {
      "Location": "string",
      "PlayerGatewayStatus": "string",
      "Status": "string"
    }
  ]
}
```

Response Elements

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

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

FleetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\]\+$`

FleetId

A unique identifier for the fleet that location attributes are being deleted for.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\]\+$|^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\]\+$`

LocationStates

The remote locations that are being deleted, with each location status set to DELETING.

Type: Array of [LocationState](#) objects

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

DeleteGameServerGroup

This API works with the following fleet types: EC2 (FleetIQ)

Terminates a game server group and permanently deletes the game server group record. You have several options for how these resources are impacted when deleting the game server group. Depending on the type of delete operation selected, this operation might affect these resources:

- The game server group
- The corresponding Auto Scaling group
- All game servers that are currently running in the group

To delete a game server group, identify the game server group to delete and specify the type of delete operation to initiate. Game server groups can only be deleted if they are in ACTIVE or ERROR status.

If the delete request is successful, a series of operations are kicked off. The game server group status is changed to DELETE_SCHEDULED, which prevents new game servers from being registered and stops automatic scaling activity. Once all game servers in the game server group are deregistered, Amazon GameLift Servers FleetIQ can begin deleting resources. If any of the delete operations fail, the game server group is placed in ERROR status.

Amazon GameLift Servers FleetIQ emits delete events to Amazon CloudWatch.

Learn more

[Amazon GameLift Servers FleetIQ Guide](#)

Request Syntax

```
{  
  "DeleteOption": "string",  
  "GameServerGroupName": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

 Note

In the following list, the required parameters are described first.

GameServerGroupName

A unique identifier for the game server group. Use either the name or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]|^arn:.*:gameservergroup\[a-zA-Z0-9-\.]`

Required: Yes

DeleteOption

The type of delete to perform. Options include the following:

- **SAFE_DELETE** – (default) Terminates the game server group and Amazon EC2 Auto Scaling group only when it has no game servers that are in UTILIZED status.
- **FORCE_DELETE** – Terminates the game server group, including all active game servers regardless of their utilization status, and the Amazon EC2 Auto Scaling group.
- **RETAIN** – Does a safe delete of the game server group but retains the Amazon EC2 Auto Scaling group as is.

Type: String

Valid Values: **SAFE_DELETE** | **FORCE_DELETE** | **RETAIN**

Required: No

Response Syntax

```
{  
  "GameServerGroup": {
```

```
"AutoScalingGroupArn": "string",
"BalancingStrategy": "string",
"CreationTime": number,
"GameServerGroupArn": "string",
"GameServerGroupName": "string",
"GameServerProtectionPolicy": "string",
"InstanceDefinitions": [
  {
    "InstanceType": "string",
    "WeightedCapacity": "string"
  }
],
"LastUpdatedTime": number,
"RoleArn": "string",
"Status": "string",
"StatusReason": "string",
"SuspendedActions": [ "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.

GameServerGroup

An object that describes the deleted game server group resource, with status updated to DELETE_SCHEDULED.

Type: [GameServerGroup](#) object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Delete a game server group

This example deletes a game server group only if there is no hosting activity taking place on instances in the group.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "GameServerGroupName": "MegaFrogServers_NA",
  "DeleteOption": [ "SAFE_DELETE" ]
}
```

CLI command:

```
aws gamelift delete-game-server-group \
  --game-server-group MegaFrogServers_NA \
```

```
--delete-option SAFE_DELETE
```

Sample Response

```
{
  "GameServerGroup": {
    "AutoScalingGroupArn": "arn:aws:autoscaling:us-west-2:123456789012:autoScalingGroup:1111aaaa-22bb-33cc-44dd-5555eeee66ff:autoScalingGroupName/MegaFrogServers_NA",
    "BalancingStrategy": "SPOT_ONLY",
    "CreationTime": 1496365885.44,
    "GameServerGroupArn": "arn:aws:gamelift:us-west-2::GameServerGroup/MegaFrogServers_NA",
    "GameServerGroupName": "MegaFrogServers_NA",
    "GameServerProtectionPolicy": "NO_PROTECTION",
    "InstanceDefinitions": [
      {
        "InstanceType": "c5.2xlarge",
        "WeightedCapacity": "1"
      },
      {
        "InstanceType": "m5.2xlarge",
        "WeightedCapacity": "1"
      }
    ],
    "LastUpdatedTime": 1496365885.44,
    "RoleArn": "arn:aws:iam:123456789012::role/GameLiftGsgRole",
    "Status": "DELETE_SCHEDULED",
    "StatusReason": "",
    "SuspendedActions": []
  }
}
```

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](#)

DeleteGameSessionQueue

This API works with the following fleet types: EC2, Anywhere, Container

Deletes a game session queue. Once a queue is successfully deleted, unfulfilled [StartGameSessionPlacement](#) requests that reference the queue will fail. To delete a queue, specify the queue name.

Request Syntax

```
{  
  "Name": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Name

A descriptive label that is associated with game session queue. Queue names must be unique within each Region. You can use either the queue ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-]+|^arn:.*:gamesessionqueue\/[a-zA-Z0-9-]+`

Required: Yes

Response Elements

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

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

DeleteLocation

This API works with the following fleet types: Anywhere

Deletes a custom location.

Before deleting a custom location, review any fleets currently using the custom location and deregister the location if it is in use. For more information, see [DeregisterCompute](#).

Request Syntax

```
{  
  "LocationName": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

LocationName

The location name of the custom location to be deleted.

Type: String

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

Pattern: `^custom-[A-Za-z0-9\-\_]+|^arn:.*:location\/custom-\S+`

Required: Yes

Response Elements

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

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

DeleteMatchmakingConfiguration

This API works with the following fleet types: EC2, Anywhere, Container

Permanently removes a FlexMatch matchmaking configuration. To delete, specify the configuration name. A matchmaking configuration cannot be deleted if it is being used in any active matchmaking tickets.

Request Syntax

```
{  
  "Name": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Name

A unique identifier for the matchmaking configuration. You can use either the configuration name or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]*|^arn:.*:matchmakingconfiguration\[a-zA-Z0-9-\.]*`

Required: Yes

Response Elements

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

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

DeleteMatchmakingRuleSet

This API works with the following fleet types: EC2, Anywhere, Container

Deletes an existing matchmaking rule set. To delete the rule set, provide the rule set name. Rule sets cannot be deleted if they are currently being used by a matchmaking configuration.

Learn more

- [Build a rule set](#)

Request Syntax

```
{  
  "Name": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Name

A unique identifier for the matchmaking rule set to be deleted. (Note: The rule set name is different from the optional "name" field in the rule set body.) You can use either the rule set name or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]*|^arn:.*:matchmakingruleset\[a-zA-Z0-9-\.]*`

Required: Yes

Response Elements

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

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

DeleteScalingPolicy

This API works with the following fleet types: EC2, Container

Deletes a fleet scaling policy. Once deleted, the policy is no longer in force and Amazon GameLift Servers removes all record of it. To delete a scaling policy, specify both the scaling policy name and the fleet ID it is associated with.

To temporarily suspend scaling policies, use [StopFleetActions](#). This operation suspends all policies for the fleet.

Request Syntax

```
{
  "FleetId": "string",
  "Name": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet to be deleted. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

Name

A descriptive label that is associated with a fleet's scaling policy. Policy names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

Response Elements

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

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Delete a scaling policy

To delete a policy, we must specify both the fleet ID and name. The combination of these two values is what uniquely identifies the policy. Once a delete request is received, the policy is put into status DELETING and no longer affects the fleet's scaling activity.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
POST / HTTP/1.1
Host: gamelift.us-west-2.amazonaws.com;
Accept-Encoding: identity
Content-Length: 336
User-Agent: aws-cli/1.11.36 Python/2.7.9 Windows/7 botocore/1.4.93
Content-Type: application/x-amz-json-1.0
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20170406/us-west-2/
gamelift/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target,
Signature=wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY
X-Amz-Date: 20170406T004805Z
X-Amz-Target: GameLift.DeleteScalingPolicy
```

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "Name": "My_Target_Policy_1"
}
```

CLI syntax:

```
$aws gamelift delete-scaling-policy
--fleet-id "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
--name "My_Target_Policy_1"
```


Sample Response

```
HTTP/1.1 200 OK undefined
```

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](#)

DeleteScript

This API works with the following fleet types: EC2

Deletes a Realtime script. This operation permanently deletes the script record. If script files were uploaded, they are also deleted (files stored in an S3 bucket are not deleted).

To delete a script, specify the script ID. Before deleting a script, be sure to terminate all fleets that are deployed with the script being deleted. Fleet instances periodically check for script updates, and if the script record no longer exists, the instance will go into an error state and be unable to host game sessions.

Learn more

[Amazon GameLift Servers Amazon GameLift Servers Realtime](#)

Related actions

[All APIs by task](#)

Request Syntax

```
{  
  "ScriptId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[ScriptId](#)

A unique identifier for the Realtime script to delete. You can use either the script ID or ARN value.

Type: String

Pattern: `^script-\S+|^arn:.*:script\/script-\S+`

Required: Yes

Response Elements

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

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Delete an existing script

This example illustrates one usage of DeleteScript.

Sample Request

```
{
    "ScriptId": "script-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
}
```

CLI syntax:

```
aws gamelift delete-script --script-id "script-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
```

Sample Response

```
HTTP/1.1 200 OK undefined
```

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](#)

DeleteVpcPeeringAuthorization

This API works with the following fleet types: EC2

Cancels a pending VPC peering authorization for the specified VPC. If you need to delete an existing VPC peering connection, use [DeleteVpcPeeringConnection](#).

Related actions

[All APIs by task](#)

Request Syntax

```
{
  "GameLiftAwsAccountId": "string",
  "PeerVpcId": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[GameLiftAwsAccountId](#)

A unique identifier for the AWS account that you use to manage your Amazon GameLift Servers fleet. You can find your Account ID in the AWS Management Console under account settings.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

PeerVpcId

A unique identifier for a VPC with resources to be accessed by your Amazon GameLift Servers fleet. The VPC must be in the same Region as your fleet. To look up a VPC ID, use the [VPC Dashboard](#) in the AWS Management Console. Learn more about VPC peering in [VPC Peering with Amazon GameLift Servers Fleets](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

Response Elements

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

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

DeleteVpcPeeringConnection

This API works with the following fleet types: EC2

Removes a VPC peering connection. To delete the connection, you must have a valid authorization for the VPC peering connection that you want to delete..

Once a valid authorization exists, call this operation from the AWS account that is used to manage the Amazon GameLift Servers fleets. Identify the connection to delete by the connection ID and fleet ID. If successful, the connection is removed.

Related actions

[All APIs by task](#)

Request Syntax

```
{  
  "FleetId": "string",  
  "VpcPeeringConnectionId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[FleetId](#)

A unique identifier for the fleet. This fleet specified must match the fleet referenced in the VPC peering connection record. You can use either the fleet ID or ARN value.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\ -]+`

Required: Yes

VpcPeeringConnectionId

A unique identifier for a VPC peering connection.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

Response Elements

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

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

DeregisterCompute

This API works with the following fleet types: Anywhere

Removes a compute resource from an Anywhere fleet. Deregistered computes can no longer host game sessions through Amazon GameLift Servers. Use this operation with an Anywhere fleet that doesn't use the Amazon GameLift Servers Agent For Anywhere fleets with the Agent, the Agent handles all compute registry tasks for you.

To deregister a compute, call this operation from the compute that's being deregistered and specify the compute name and the fleet ID.

Request Syntax

```
{
  "ComputeName": "string",
  "FleetId": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

ComputeName

The unique identifier of the compute resource to deregister. For an Anywhere fleet compute, use the registered compute name.

Type: String

Length Constraints: Maximum length of 1024.

Pattern: `^[a-zA-Z0-9\-\_](\/[a-zA-Z0-9\-\_])?|^arn:.*:compute\[a-zA-Z0-9\-\_](\/[a-zA-Z0-9\-\_])?`

Required: Yes

FleetId

A unique identifier for the fleet the compute resource is currently registered to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

Response Elements

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

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

DeregisterGameServer

This API works with the following fleet types: EC2 (FleetIQ)

Removes the game server from a game server group. As a result of this operation, the deregistered game server can no longer be claimed and will not be returned in a list of active game servers.

To deregister a game server, specify the game server group and game server ID. If successful, this operation emits a CloudWatch event with termination timestamp and reason.

Learn more

[Amazon GameLift Servers FleetIQ Guide](#)

Request Syntax

```
{
  "GameServerGroupName": "string",
  "GameServerId": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

GameServerGroupName

A unique identifier for the game server group where the game server is running.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]+\|^arn:.*:gameservergroup\[a-zA-Z0-9-\.]+\`

Required: Yes

GameServerId

A custom string that uniquely identifies the game server to deregister.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: [a-zA-Z0-9-\.]+

Required: Yes

Response Elements

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

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Remove a game server from active status

This example deregisters a game server so that it can no longer be claimed for gameplay.

Sample Request

```
{
  "GameServerGroupName": "MegaFrogServers_NA",
  "GameServerId": "mega-frog-game-12345678"
}
```

CLI command:

```
aws gamelift deregister-game-server \
  --game-server-group-name MegaFrogServers_NA \
  --GameServerId mega-frog-game-12345678
```

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](#)

DescribeAlias

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves properties for an alias. This operation returns all alias metadata and settings. To get an alias's target fleet ID only, use `ResolveAlias`.

To get alias properties, specify the alias ID. If successful, the requested alias record is returned.

Related actions

[All APIs by task](#)

Request Syntax

```
{  
  "AliasId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[AliasId](#)

The unique identifier for the fleet alias that you want to retrieve. You can use either the alias ID or ARN value.

Type: String

Pattern: `^alias-\S+|^arn:.*:alias\/alias-\S+`

Required: Yes

Response Syntax

```
{
  "Alias": {
    "AliasArn": "string",
    "AliasId": "string",
    "CreationTime": number,
    "Description": "string",
    "LastUpdatedTime": number,
    "Name": "string",
    "RoutingStrategy": {
      "FleetId": "string",
      "Message": "string",
      "Type": "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.

Alias

The requested alias resource.

Type: [Alias](#) object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

DescribeBuild

This API works with the following fleet types: EC2

Retrieves properties for a custom game build. To request a build resource, specify a build ID. If successful, an object containing the build properties is returned.

Learn more

[Upload a Custom Server Build](#)

[All APIs by task](#)

Request Syntax

```
{  
  "BuildId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[BuildId](#)

A unique identifier for the build to retrieve properties for. You can use either the build ID or ARN value.

Type: String

Pattern: `^build-\S+|^arn:.*:build\/build-\S+`

Required: Yes

Response Syntax

```
{
  "Build": {
    "BuildArn": "string",
    "BuildId": "string",
    "CreationTime": number,
    "Name": "string",
    "OperatingSystem": "string",
    "ServerSdkVersion": "string",
    "SizeOnDisk": number,
    "Status": "string",
    "Version": "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.

Build

Set of properties describing the requested build.

Type: [Build](#) object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

View a build resource

This example retrieves the information, including current status, about a specific build.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "BuildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
}
```

Sample Response

```
{
  "Build": {
    "BuildArn": "arn:aws:gamelift:us-west-2::build/build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "BuildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "CreationTime": 1496708916.18,
    "Name": "My_Game_Server_Build_One",
    "OperatingSystem": "AMAZON_LINUX_2023",
    "SizeOnDisk": 1304924,
    "Status": "READY",
    "Version": "12345.678"
  }
}
```

```
}
```

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](#)

DescribeCompute

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves properties for a specific compute resource in an Amazon GameLift Servers fleet. You can list all computes in a fleet by calling [ListCompute](#).

Request options

Provide the fleet ID and compute name. The compute name varies depending on the type of fleet.

- For a compute in a managed EC2 fleet, provide an instance ID. Each instance in the fleet is a compute.
- For a compute in a managed container fleet, provide a compute name. In a container fleet, each game server container group on a fleet instance is assigned a compute name.
- For a compute in an Anywhere fleet, provide a registered compute name. Anywhere fleet computes are created when you register a hosting resource with the fleet.

Results

If successful, this operation returns details for the requested compute resource. Depending on the fleet's compute type, the result includes the following information:

- For a managed EC2 fleet, this operation returns information about the EC2 instance.
- For a managed container fleet, this operation returns information about the container group and the instance it's running on.
- For an Anywhere fleet, this operation returns information about the registered compute.

Request Syntax

```
{  
  "ComputeName": "string",  
  "FleetId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

 Note

In the following list, the required parameters are described first.

ComputeName

The unique identifier of the compute resource to retrieve properties for. For a managed container fleet or Anywhere fleet, use a compute name. For an EC2 fleet, use an instance ID. To retrieve a fleet's compute identifiers, call [ListCompute](#).

Type: String

Length Constraints: Maximum length of 1024.

Pattern: `^[a-zA-Z0-9\-\_](\/[a-zA-Z0-9\-\_])?|^arn:.*:compute\[a-zA-Z0-9\-\_](\/[a-zA-Z0-9\-\_])?`

Required: Yes

FleetId

A unique identifier for the fleet that the compute belongs to. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+|^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\_]+`

Required: Yes

Response Syntax

```
{
  "Compute": {
    "ComputeArn": "string",
    "ComputeName": "string",
```

```
"ComputeStatus": "string",
"ContainerAttributes": [
  {
    "ContainerName": "string",
    "ContainerRuntimeId": "string"
  }
],
"CreationTime": number,
"DnsName": "string",
"FleetArn": "string",
"FleetId": "string",
"GameLiftAgentEndpoint": "string",
"GameLiftServiceSdkEndpoint": "string",
"GameServerContainerGroupDefinitionArn": "string",
"InstanceId": "string",
"IpAddress": "string",
"Location": "string",
"OperatingSystem": "string",
"Type": "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.

Compute

The set of properties for the requested compute resource.

Type: [Compute](#) object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

DescribeContainerFleet

This API works with the following fleet types: Container

Retrieves the properties for a container fleet. When requesting attributes for multiple fleets, use the pagination parameters to retrieve results as a set of sequential pages.

Request options

- Get container fleet properties for a single fleet. Provide either the fleet ID or ARN value.

Results

If successful, a `ContainerFleet` object is returned. This object includes the fleet properties, including information about the most recent deployment.

Note

Some API operations limit the number of fleet IDs that allowed in one request. If a request exceeds this limit, the request fails and the error message contains the maximum allowed number.

Request Syntax

```
{  
  "FleetId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the container fleet to retrieve. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\`

Required: Yes

Response Syntax

```
{
  "ContainerFleet": {
    "BillingType": "string",
    "CreationTime": number,
    "DeploymentDetails": {
      "LatestDeploymentId": "string"
    },
    "Description": "string",
    "FleetArn": "string",
    "FleetId": "string",
    "FleetRoleArn": "string",
    "GameServerContainerGroupDefinitionArn": "string",
    "GameServerContainerGroupDefinitionName": "string",
    "GameServerContainerGroupsPerInstance": number,
    "GameSessionCreationLimitPolicy": {
      "NewGameSessionsPerCreator": number,
      "PolicyPeriodInMinutes": number
    },
    "InstanceConnectionPortRange": {
      "FromPort": number,
      "ToPort": number
    },
    "InstanceInboundPermissions": [
      {
        "FromPort": number,
        "IpRange": "string",
```

```
        "Protocol": "string",
        "ToPort": number
    },
    ],
    "InstanceType": "string",
    "LocationAttributes": [
        {
            "Location": "string",
            "PlayerGatewayStatus": "string",
            "Status": "string"
        }
    ],
    "LogConfiguration": {
        "LogDestination": "string",
        "LogGroupArn": "string",
        "S3BucketName": "string"
    },
    "MaximumGameServerContainerGroupsPerInstance": number,
    "MetricGroups": [ "string" ],
    "NewGameSessionProtectionPolicy": "string",
    "PerInstanceContainerGroupDefinitionArn": "string",
    "PerInstanceContainerGroupDefinitionName": "string",
    "PlayerGatewayMode": "string",
    "Status": "string"
}
}
```

Response Elements

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

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

ContainerFleet

The properties for the requested container fleet, including current status.

Type: [ContainerFleet](#) object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Request attributes for a container fleet

This example retrieves attributes for an managed container fleet.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
```

```
"FleetId": "containerfleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
}
```

Sample Response

```
{
  "ContainerFleet": {
    "BillingType": ON_DEMAND,
    "CreationTime": 1736365885.22,
    "DeploymentDetails": {
      "LatestDeploymentId": "deployment-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
    },
    "FleetArn": "arn:aws:gamelift:us-west-2:111122223333:containerfleet/
containerfleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "FleetId": "containerfleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "FleetRoleArn": "arn:aws:iam::MyAccount:role/MyRole",
    "GameServerContainerGroupDefinitionArn": "arn:aws:gamelift:us-
west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:2",
    "GameServerContainerGroupDefinitionName": "MyAdventureGameContainerGroup",
    "GameServerContainerGroupsPerInstance": number,
    "InstanceConnectionPortRange": {
      "FromPort": 4192,
      "ToPort": 4242
    },
    "InstanceInboundPermissions": [
      {
        "FromPort": 4192,
        "IpRange": "string",
        "Protocol": "UDP",
        "ToPort": 4242,
      }
    ],
    "InstanceType": "c5.large",
    "LogConfiguration": {
      "LogGroupArn": "arn:aws:logs:us-west-2:111222333444:log-group:customerLogs",
      "LogDestination": "CLOUDWATCH"
    },
    "MaximumGameServerContainerGroupsPerInstance": 10,
    "NewGameSessionProtectionPolicy": "NoProtection",
    "Status": "PENDING"
  }
}
```

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](#)

DescribeContainerGroupDefinition

This API works with the following fleet types: Container

Retrieves the properties of a container group definition, including all container definitions in the group.

Request options:

- Retrieve the latest version of a container group definition. Specify the container group definition name only, or use an ARN value without a version number.
- Retrieve a particular version. Specify the container group definition name and a version number, or use an ARN value that includes the version number.

Results:

If successful, this operation returns the complete properties of a container group definition version.

Learn more

- [Manage a container group definition](#)

Request Syntax

```
{  
  "Name": "string",  
  "VersionNumber": number  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Name

The unique identifier for the container group definition to retrieve properties for. You can use either the Name or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-zA-Z0-9\-\-]+\$|^arn:.*:containergroupdefinition\[a-zA-Z0-9\-\-]+(:[0-9]+)?\$`

Required: Yes

VersionNumber

The specific version to retrieve.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

Response Syntax

```
{
  "ContainerGroupDefinition": {
    "ContainerGroupDefinitionArn": "string",
    "ContainerGroupType": "string",
    "CreationTime": number,
    "GameServerContainerDefinition": {
      "ContainerName": "string",
      "DependsOn": [
        {
          "Condition": "string",
          "ContainerName": "string"
        }
      ],
      "EnvironmentOverride": [
        {
          "Name": "string",
          "Value": "string"
        }
      ]
    }
  }
}
```

```

    ],
    "ImageUri": "string",
    "LinuxCapabilities": {
        "Include": [ "string" ]
    },
    "MountPoints": [
        {
            "AccessLevel": "string",
            "ContainerPath": "string",
            "InstancePath": "string"
        }
    ],
    "PortConfiguration": {
        "ContainerPortRanges": [
            {
                "FromPort": number,
                "Protocol": "string",
                "ToPort": number
            }
        ]
    },
    "ResolvedImageDigest": "string",
    "ServerSdkVersion": "string"
},
"Name": "string",
"OperatingSystem": "string",
"Status": "string",
"StatusReason": "string",
"SupportContainerDefinitions": [
    {
        "ContainerName": "string",
        "DependsOn": [
            {
                "Condition": "string",
                "ContainerName": "string"
            }
        ]
    },
    {
        "Name": "string",
        "Value": "string"
    }
],
"Essential": boolean,

```

```

    "HealthCheck": {
      "Command": [ "string" ],
      "Interval": number,
      "Retries": number,
      "StartPeriod": number,
      "Timeout": number
    },
    "ImageUri": "string",
    "LinuxCapabilities": {
      "Include": [ "string" ]
    },
    "MemoryHardLimitMebibytes": number,
    "MountPoints": [
      {
        "AccessLevel": "string",
        "ContainerPath": "string",
        "InstancePath": "string"
      }
    ],
    "PortConfiguration": {
      "ContainerPortRanges": [
        {
          "FromPort": number,
          "Protocol": "string",
          "ToPort": number
        }
      ]
    },
    "ResolvedImageDigest": "string",
    "Vcpu": number
  }
],
"TotalMemoryLimitMebibytes": number,
"TotalVcpuLimit": number,
"VersionDescription": "string",
"VersionNumber": number
}
}

```

Response Elements

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

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

ContainerGroupDefinition

The properties of the requested container group definition resource.

Type: [ContainerGroupDefinition](#) object

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Retrieve a container group definition version

This example retrieves the 3rd version of a container group definition. It assumes we're working in the same AWS Region as the container group we want to retrieve. To retrieve a container group definition in a different region, we could make the request using the definition's ARN value or we could specify the other region.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "Name": "MyAdventureGameContainerGroup",
  "Version": 3
}
```

Sample Response

```
{
  "ContainerGroupDefinition": {
    "ContainerGroupDefinitionArn": "arn:aws:gamelift:us-west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:3",
    "ContainerGroupType": "GAME_SERVER",
    "CreationTime": 1496365885.44,
    "GameServerContainerDefinition": [
      {
        "ContainerName": "MyAdventureGameContainer",
        "ImageUri": "111122223333.dkr.ecr.us-west-2.amazonaws.com/MyAdventureGameContainerImage",
        "PortConfiguration": {
          "ContainerPortRanges": [
            {
              "FromPort": 35000,
              "Protocol": "TCP",
              "ToPort": 40000
            }
          ]
        }
      }
    ]
  },
}
```

```
    "LinuxCapabilities": {
      "Include": ["NET_BIND_SERVICE", "SYS_PTRACE"]
    },
    "ResolvedImageDigest":
"sha256:0123456789abcdef0123456789abcdef0123456789abcdef0123456789abcdef",
    "ServerSdkVersion": "5.2.0"
  }
],
"Name": "MyAdventureGameContainerGroup",
"OperatingSystem": "AMAZON_LINUX_2023",
"Status": "READY",
"SupportContainerDefinitions": [
  {
    "ContainerName": "MyAdventureGameDependencies",
    "ImageUri": "111122223333.dkr.ecr.us-west-2.amazonaws.com/
MyAdventureGameContainerImage",
    "LinuxCapabilities": {
      "Include": ["NET_ADMIN"]
    }
  }
],
"TotalMemoryLimitMebibytes": 1024,
"TotalVcpuLimit": 1,
"VersionNumber": "3"
}
}
```

See Also

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

- [AWS Command Line Interface 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](#)

DescribeContainerGroupPortMappings

This API works with the following fleet types: Container

Retrieves the port mappings for a container group running on a container fleet. Port mappings show how container ports are mapped to connection ports on the fleet instance. Use this operation to find the connection port for a specific container on a fleet instance.

Request options

- Get port mappings for a game server container group. Provide the fleet ID, set `ContainerGroupType` to `GAME_SERVER`, and specify the `ComputeName` for the game server container group.
- Get port mappings for a per-instance container group. Provide the fleet ID, set `ContainerGroupType` to `PER_INSTANCE`, and specify the `InstanceId` for the instance.
- Optionally filter results to a single container by providing a `ContainerName`.

Results

This operation returns the fleet ID, fleet ARN, location, container group definition ARN, container group type, compute name (for game server container groups), instance ID, and a list of `ContainerGroupPortMapping` objects. Each object contains the container name, runtime ID, and a list of port mappings that show how container ports map to connection ports on the instance.

Learn more

[Connect to containers](#)

[Create a container group definition](#)

Request Syntax

```
{
  "ComputeName": "string",
  "ContainerGroupType": "string",
  "ContainerName": "string",
  "FleetId": "string",
  "InstanceId": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

ContainerGroupType

The type of container group to retrieve port mappings for.

- `GAME_SERVER` -- Get port mappings for a game server container group.
- `PER_INSTANCE` -- Get port mappings for a per-instance container group.

Type: String

Valid Values: `GAME_SERVER` | `PER_INSTANCE`

Required: Yes

FleetId

A unique identifier for the container fleet. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

ComputeName

A unique identifier for the compute resource for which to retrieve port mappings. For a container fleet, a compute represents a game server container group running on a fleet instance. You can use either the compute name or ARN value.

When `ContainerGroupType` is `GAME_SERVER`, this parameter is required.

When `ContainerGroupType` is `PER_INSTANCE`, do not provide this parameter. If you provide a compute name with `PER_INSTANCE`, the request fails with an `InvalidRequestException`.

Type: String

Length Constraints: Maximum length of 1024.

Pattern: `^[a-zA-Z0-9\-\-](\/[a-zA-Z0-9\-\-])?|^arn:.*:compute\/[a-zA-Z0-9\-\-](\/[a-zA-Z0-9\-\-])?`

Required: No

ContainerName

A container name to filter the results. When provided, the operation returns port mappings for the specified container only. If no container with the specified name exists in the container group, the request fails with a `NotFoundException`.

If not provided, the operation returns port mappings for all containers in the container group.

Type: String

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

Pattern: `^[a-zA-Z0-9\-\-]+$`

Required: No

InstanceId

A unique identifier for the fleet instance to retrieve port mappings for.

When `ContainerGroupType` is `PER_INSTANCE`, this parameter is required.

When `ContainerGroupType` is `GAME_SERVER`, this parameter is optional. If you provide an instance ID, it must match the instance that's running the specified compute. If the instance ID doesn't match, the request fails with an `InvalidRequestException`.

Type: String

Pattern: `[a-zA-Z0-9\-\-\.-]+`

Required: No

Response Syntax

```
{
  "ComputeName": "string",
  "ContainerGroupDefinitionArn": "string",
  "ContainerGroupPortMappings": [
    {
      "ContainerName": "string",
      "ContainerPortMappings": [
        {
          "ConnectionPort": number,
          "ContainerPort": number,
          "Protocol": "string"
        }
      ],
      "ContainerRuntimeId": "string"
    }
  ],
  "ContainerGroupType": "string",
  "FleetArn": "string",
  "FleetId": "string",
  "InstanceId": "string",
  "Location": "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.

ComputeName

A unique identifier for the compute resource running the game server container group. Returned when ContainerGroupType is GAME_SERVER.

Type: String

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

Pattern: `^[a-zA-Z0-9\-\_]+(\/[a-zA-Z0-9\-\_]+)?`

ContainerGroupDefinitionArn

The Amazon Resource Name ([ARN](#)) that is assigned to the container group definition. The ARN value also identifies the specific container group definition version in use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:containergroupdefinition\[a-zA-Z0-9\-\_]+(:[0-9]+)?$`

ContainerGroupPortMappings

A list of `ContainerGroupPortMapping` objects that describe the port mappings for each container in the container group.

Type: Array of [ContainerGroupPortMapping](#) objects

ContainerGroupType

The type of container group that was specified in the request. Valid values are `GAME_SERVER` or `PER_INSTANCE`.

Type: String

Valid Values: `GAME_SERVER` | `PER_INSTANCE`

FleetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`. In a GameLift fleet ARN, the resource ID matches the `FleetId` value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\_]+$`

FleetId

A unique identifier for the container fleet.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-.]+`

InstanceId

A unique identifier for the fleet instance. For GAME_SERVER requests, this is the instance running the specified compute. For PER_INSTANCE requests, this is the instance specified in the request.

Type: String

Pattern: `[a-zA-Z0-9\-.]+`

Location

The location of the fleet instance, expressed as an AWS Region code, such as `us-west-2`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-.]+`

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Get port mappings for a game server container group

This example retrieves port mappings for a game server container group running on a specific compute in a container fleet.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "FleetId": "containerfleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
  "ContainerGroupType": "GAME_SERVER",
  "ComputeName": "62c5ff7f7a9a445d84877074c80aeafc"
}
```

Sample Response

```
{
  "FleetId": "containerfleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
  "FleetArn": "arn:aws:gamelift:us-west-2:111122223333:containerfleet/
containerfleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
  "Location": "us-west-2",
}
```

```

    "ContainerGroupDefinitionArn": "arn:aws:gamelift:us-
west-2:111122223333:containergroupdefinition/MyGameServerContainerGroup:1",
    "ContainerGroupType": "GAME_SERVER",
    "ComputeName": "62c5ff7f7a9a445d84877074c80aeafc",
    "InstanceId": "i-1234567890abcdef0",
    "ContainerGroupPortMappings": [
        {
            "ContainerName": "MyGameServer",
            "ContainerRuntimeId": "a1b2c3d4e5f6",
            "ContainerPortMappings": [
                {
                    "ContainerPort": 7777,
                    "ConnectionPort": 49200,
                    "Protocol": "UDP"
                },
                {
                    "ContainerPort": 8080,
                    "ConnectionPort": 49201,
                    "Protocol": "TCP"
                }
            ]
        },
        {
            "ContainerName": "MySidecar",
            "ContainerRuntimeId": "f6e5d4c3b2a1",
            "ContainerPortMappings": [
                {
                    "ContainerPort": 9090,
                    "ConnectionPort": 49202,
                    "Protocol": "TCP"
                }
            ]
        }
    ]
}

```

Get port mappings for a per-instance container group

This example retrieves port mappings for a per-instance container group (such as daemon or logging containers) running on a specific instance in a container fleet.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "FleetId": "containerfleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
  "ContainerGroupType": "PER_INSTANCE",
  "InstanceId": "i-1234567890abcdef0"
}
```

Sample Response

```
{
  "FleetId": "containerfleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
  "FleetArn": "arn:aws:gamelift:us-west-2:111122223333:containerfleet/containerfleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
  "Location": "us-west-2",
  "ContainerGroupDefinitionArn": "arn:aws:gamelift:us-west-2:111122223333:containergroupdefinition/MyDaemonContainerGroup:1",
  "ContainerGroupType": "PER_INSTANCE",
  "InstanceId": "i-1234567890abcdef0",
  "ContainerGroupPortMappings": [
    {
      "ContainerName": "MyLoggingAgent",
      "ContainerRuntimeId": "b2c3d4e5f6a1",
      "ContainerPortMappings": [
        {
          "ContainerPort": 8125,
          "ConnectionPort": 49300,
          "Protocol": "UDP"
        }
      ]
    }
  ]
}
```

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](#)

DescribeEC2InstanceLimits

This API works with the following fleet types: EC2

Retrieves the instance limits and current utilization for an AWS Region or location. Instance limits control the number of instances, per instance type, per location, that your AWS account can use. Learn more at [Amazon EC2 Instance Types](#). The information returned includes the maximum number of instances allowed and your account's current usage across all fleets. This information can affect your ability to scale your Amazon GameLift Servers fleets. You can request a limit increase for your account by using the **Service limits** page in the Amazon GameLift Servers console.

Instance limits differ based on whether the instances are deployed in a fleet's home Region or in a remote location. For remote locations, limits also differ based on the combination of home Region and remote location. All requests must specify an AWS Region (either explicitly or as your default settings). To get the limit for a remote location, you must also specify the location. To learn more about how Amazon GameLift Servers handles locations, see [Amazon GameLift Servers service locations](#). For example, the following requests all return different results:

- Request specifies the Region `ap-northeast-1` with no location. The result is limits and usage data on all of the fleets that reside in `ap-northeast-1`, for all instance types that are deployed in `ap-northeast-1`.
- Request specifies the Region `ap-northeast-1` with location `us-west-2`. The result is limits and usage data on all of the fleets that reside in `ap-northeast-1`, for all instance types that are deployed in `us-west-2`.
- Request specifies the Region `us-east-1` with location `ap-northeast-1`. The result is limits and usage data on all of the fleets that reside in `us-east-1`, for all instance types that are deployed in `ap-northeast-1`. These limits do not affect fleets in any other Regions that deploy instances to `ap-northeast-1`.

This operation can be used in the following ways:

- To get limit and usage data for all instance types that are deployed in an AWS Region by fleets that reside in the same Region: Specify the Region only. Optionally, specify a single instance type to retrieve information for.

- To get limit and usage data for all instance types that are deployed to a remote location by fleets that reside in different AWS Region: Provide both the AWS Region and the remote location. Optionally, specify a single instance type to retrieve information for.

If successful, an `EC2InstanceLimits` object is returned with limits and usage data for each requested instance type.

Learn more

[Setting up Amazon GameLift Servers fleets](#)

Request Syntax

```
{  
  "EC2InstanceType": "string",  
  "Location": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

EC2InstanceType

Name of an Amazon EC2 instance type that is supported in Amazon GameLift Servers. A fleet instance type determines the computing resources of each instance in the fleet, including CPU, memory, storage, and networking capacity. Do not specify a value for this parameter to retrieve limits for all instance types.

Type: String

Valid Values: `t2.micro` | `t2.small` | `t2.medium` | `t2.large` | `c3.large` | `c3.xlarge` | `c3.2xlarge` | `c3.4xlarge` | `c3.8xlarge` | `c4.large` | `c4.xlarge` | `c4.2xlarge` | `c4.4xlarge` | `c4.8xlarge` | `c5.large` | `c5.xlarge` |

c5.2xlarge | c5.4xlarge | c5.9xlarge | c5.12xlarge | c5.18xlarge |
c5.24xlarge | c5a.large | c5a.xlarge | c5a.2xlarge | c5a.4xlarge |
c5a.8xlarge | c5a.12xlarge | c5a.16xlarge | c5a.24xlarge | r3.large |
r3.xlarge | r3.2xlarge | r3.4xlarge | r3.8xlarge | r4.large | r4.xlarge
| r4.2xlarge | r4.4xlarge | r4.8xlarge | r4.16xlarge | r5.large |
r5.xlarge | r5.2xlarge | r5.4xlarge | r5.8xlarge | r5.12xlarge |
r5.16xlarge | r5.24xlarge | r5a.large | r5a.xlarge | r5a.2xlarge |
r5a.4xlarge | r5a.8xlarge | r5a.12xlarge | r5a.16xlarge | r5a.24xlarge
| m3.medium | m3.large | m3.xlarge | m3.2xlarge | m4.large | m4.xlarge
| m4.2xlarge | m4.4xlarge | m4.10xlarge | m5.large | m5.xlarge |
m5.2xlarge | m5.4xlarge | m5.8xlarge | m5.12xlarge | m5.16xlarge |
m5.24xlarge | m5a.large | m5a.xlarge | m5a.2xlarge | m5a.4xlarge |
m5a.8xlarge | m5a.12xlarge | m5a.16xlarge | m5a.24xlarge | c5d.large
| c5d.xlarge | c5d.2xlarge | c5d.4xlarge | c5d.9xlarge | c5d.12xlarge
| c5d.18xlarge | c5d.24xlarge | c6a.large | c6a.xlarge | c6a.2xlarge |
c6a.4xlarge | c6a.8xlarge | c6a.12xlarge | c6a.16xlarge | c6a.24xlarge
| c6i.large | c6i.xlarge | c6i.2xlarge | c6i.4xlarge | c6i.8xlarge |
c6i.12xlarge | c6i.16xlarge | c6i.24xlarge | r5d.large | r5d.xlarge |
r5d.2xlarge | r5d.4xlarge | r5d.8xlarge | r5d.12xlarge | r5d.16xlarge
| r5d.24xlarge | m6g.medium | m6g.large | m6g.xlarge | m6g.2xlarge |
m6g.4xlarge | m6g.8xlarge | m6g.12xlarge | m6g.16xlarge | c6g.medium
| c6g.large | c6g.xlarge | c6g.2xlarge | c6g.4xlarge | c6g.8xlarge |
c6g.12xlarge | c6g.16xlarge | r6g.medium | r6g.large | r6g.xlarge |
r6g.2xlarge | r6g.4xlarge | r6g.8xlarge | r6g.12xlarge | r6g.16xlarge
| c6gn.medium | c6gn.large | c6gn.xlarge | c6gn.2xlarge | c6gn.4xlarge
| c6gn.8xlarge | c6gn.12xlarge | c6gn.16xlarge | c7g.medium | c7g.large
| c7g.xlarge | c7g.2xlarge | c7g.4xlarge | c7g.8xlarge | c7g.12xlarge
| c7g.16xlarge | r7g.medium | r7g.large | r7g.xlarge | r7g.2xlarge |
r7g.4xlarge | r7g.8xlarge | r7g.12xlarge | r7g.16xlarge | m7g.medium
| m7g.large | m7g.xlarge | m7g.2xlarge | m7g.4xlarge | m7g.8xlarge |
m7g.12xlarge | m7g.16xlarge | g5g.xlarge | g5g.2xlarge | g5g.4xlarge
| g5g.8xlarge | g5g.16xlarge | r6i.large | r6i.xlarge | r6i.2xlarge |
r6i.4xlarge | r6i.8xlarge | r6i.12xlarge | r6i.16xlarge | c6gd.medium |
c6gd.large | c6gd.xlarge | c6gd.2xlarge | c6gd.4xlarge | c6gd.8xlarge |
c6gd.12xlarge | c6gd.16xlarge | c6in.large | c6in.xlarge | c6in.2xlarge
| c6in.4xlarge | c6in.8xlarge | c6in.12xlarge | c6in.16xlarge |

c7a.medium | c7a.large | c7a.xlarge | c7a.2xlarge | c7a.4xlarge |
c7a.8xlarge | c7a.12xlarge | c7a.16xlarge | c7gd.medium | c7gd.large |
c7gd.xlarge | c7gd.2xlarge | c7gd.4xlarge | c7gd.8xlarge | c7gd.12xlarge
| c7gd.16xlarge | c7gn.medium | c7gn.large | c7gn.xlarge | c7gn.2xlarge
| c7gn.4xlarge | c7gn.8xlarge | c7gn.12xlarge | c7gn.16xlarge |
c7i.large | c7i.xlarge | c7i.2xlarge | c7i.4xlarge | c7i.8xlarge |
c7i.12xlarge | c7i.16xlarge | m6a.large | m6a.xlarge | m6a.2xlarge |
m6a.4xlarge | m6a.8xlarge | m6a.12xlarge | m6a.16xlarge | m6gd.medium
| m6gd.large | m6gd.xlarge | m6gd.2xlarge | m6gd.4xlarge | m6gd.8xlarge
| m6gd.12xlarge | m6gd.16xlarge | m6i.large | m6i.xlarge | m6i.2xlarge
| m6i.4xlarge | m6i.8xlarge | m6i.12xlarge | m6i.16xlarge | m7a.medium
| m7a.large | m7a.xlarge | m7a.2xlarge | m7a.4xlarge | m7a.8xlarge |
m7a.12xlarge | m7a.16xlarge | m7gd.medium | m7gd.large | m7gd.xlarge
| m7gd.2xlarge | m7gd.4xlarge | m7gd.8xlarge | m7gd.12xlarge |
m7gd.16xlarge | m7i.large | m7i.xlarge | m7i.2xlarge | m7i.4xlarge |
m7i.8xlarge | m7i.12xlarge | m7i.16xlarge | r6gd.medium | r6gd.large |
r6gd.xlarge | r6gd.2xlarge | r6gd.4xlarge | r6gd.8xlarge | r6gd.12xlarge
| r6gd.16xlarge | r7a.medium | r7a.large | r7a.xlarge | r7a.2xlarge |
r7a.4xlarge | r7a.8xlarge | r7a.12xlarge | r7a.16xlarge | r7gd.medium |
r7gd.large | r7gd.xlarge | r7gd.2xlarge | r7gd.4xlarge | r7gd.8xlarge |
r7gd.12xlarge | r7gd.16xlarge | r7i.large | r7i.xlarge | r7i.2xlarge |
r7i.4xlarge | r7i.8xlarge | r7i.12xlarge | r7i.16xlarge | r7i.24xlarge |
r7i.48xlarge | c5ad.large | c5ad.xlarge | c5ad.2xlarge | c5ad.4xlarge |
c5ad.8xlarge | c5ad.12xlarge | c5ad.16xlarge | c5ad.24xlarge | c5n.large
| c5n.xlarge | c5n.2xlarge | c5n.4xlarge | c5n.9xlarge | c5n.18xlarge |
r5ad.large | r5ad.xlarge | r5ad.2xlarge | r5ad.4xlarge | r5ad.8xlarge |
r5ad.12xlarge | r5ad.16xlarge | r5ad.24xlarge | c6id.large | c6id.xlarge
| c6id.2xlarge | c6id.4xlarge | c6id.8xlarge | c6id.12xlarge |
c6id.16xlarge | c6id.24xlarge | c6id.32xlarge | c8g.medium | c8g.large
| c8g.xlarge | c8g.2xlarge | c8g.4xlarge | c8g.8xlarge | c8g.12xlarge
| c8g.16xlarge | c8g.24xlarge | c8g.48xlarge | m5ad.large | m5ad.xlarge
| m5ad.2xlarge | m5ad.4xlarge | m5ad.8xlarge | m5ad.12xlarge |
m5ad.16xlarge | m5ad.24xlarge | m5d.large | m5d.xlarge | m5d.2xlarge |
m5d.4xlarge | m5d.8xlarge | m5d.12xlarge | m5d.16xlarge | m5d.24xlarge
| m5dn.large | m5dn.xlarge | m5dn.2xlarge | m5dn.4xlarge | m5dn.8xlarge
| m5dn.12xlarge | m5dn.16xlarge | m5dn.24xlarge | m5n.large |

m5n.xlarge | m5n.2xlarge | m5n.4xlarge | m5n.8xlarge | m5n.12xlarge | m5n.16xlarge | m5n.24xlarge | m6id.large | m6id.xlarge | m6id.2xlarge | m6id.4xlarge | m6id.8xlarge | m6id.12xlarge | m6id.16xlarge | m6id.24xlarge | m6id.32xlarge | m6idn.large | m6idn.xlarge | m6idn.2xlarge | m6idn.4xlarge | m6idn.8xlarge | m6idn.12xlarge | m6idn.16xlarge | m6idn.24xlarge | m6idn.32xlarge | m6in.large | m6in.xlarge | m6in.2xlarge | m6in.4xlarge | m6in.8xlarge | m6in.12xlarge | m6in.16xlarge | m6in.24xlarge | m6in.32xlarge | m8g.medium | m8g.large | m8g.xlarge | m8g.2xlarge | m8g.4xlarge | m8g.8xlarge | m8g.12xlarge | m8g.16xlarge | m8g.24xlarge | m8g.48xlarge | r5dn.large | r5dn.xlarge | r5dn.2xlarge | r5dn.4xlarge | r5dn.8xlarge | r5dn.12xlarge | r5dn.16xlarge | r5dn.24xlarge | r5n.large | r5n.xlarge | r5n.2xlarge | r5n.4xlarge | r5n.8xlarge | r5n.12xlarge | r5n.16xlarge | r5n.24xlarge | r6a.large | r6a.xlarge | r6a.2xlarge | r6a.4xlarge | r6a.8xlarge | r6a.12xlarge | r6a.16xlarge | r6a.24xlarge | r6a.32xlarge | r6a.48xlarge | r6id.large | r6id.xlarge | r6id.2xlarge | r6id.4xlarge | r6id.8xlarge | r6id.12xlarge | r6id.16xlarge | r6id.24xlarge | r6id.32xlarge | r6idn.large | r6idn.xlarge | r6idn.2xlarge | r6idn.4xlarge | r6idn.8xlarge | r6idn.12xlarge | r6idn.16xlarge | r6idn.24xlarge | r6idn.32xlarge | r6in.large | r6in.xlarge | r6in.2xlarge | r6in.4xlarge | r6in.8xlarge | r6in.12xlarge | r6in.16xlarge | r6in.24xlarge | r6in.32xlarge | r8g.medium | r8g.large | r8g.xlarge | r8g.2xlarge | r8g.4xlarge | r8g.8xlarge | r8g.12xlarge | r8g.16xlarge | r8g.24xlarge | r8g.48xlarge | m4.16xlarge | c6a.32xlarge | c6a.48xlarge | c6i.32xlarge | r6i.24xlarge | r6i.32xlarge | c6in.24xlarge | c6in.32xlarge | c7a.24xlarge | c7a.32xlarge | c7a.48xlarge | c7i.24xlarge | c7i.48xlarge | m6a.24xlarge | m6a.32xlarge | m6a.48xlarge | m6i.24xlarge | m6i.32xlarge | m7a.24xlarge | m7a.32xlarge | m7a.48xlarge | m7i.24xlarge | m7i.48xlarge | r7a.24xlarge | r7a.32xlarge | r7a.48xlarge

Required: No

Location

The name of a remote location to request instance limits for, in the form of an AWS Region code such as us-west-2.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-\_]+`

Required: No

Response Syntax

```
{
  "EC2InstanceLimits": [
    {
      "CurrentInstances": number,
      "EC2InstanceType": "string",
      "InstanceLimit": number,
      "Location": "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.

EC2InstanceLimits

The maximum number of instances for the specified instance type.

Type: Array of [EC2InstanceLimit](#) objects

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Request EC2 instance limits for home region (ap-northeast-1) with no location

In this example, we request the instance limits and current usage for the `m5.large` instance type. The API call is made to the AWS Region `ap-northeast-1` with no `Location` specified, so the results show limits and usage for fleets that reside in `ap-northeast-1`, for instances deployed in `ap-northeast-1`. The response indicates that our account allows up to 200 `m5.large` instances and currently has 25 deployed.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "EC2InstanceType": "m5.large"
}
```

Sample Response

```
{
```

```
"EC2InstanceLimits": [  
  {  
    "EC2InstanceType": "m5.large",  
    "CurrentInstances": 25,  
    "InstanceLimit": 200,  
    "Location": "ap-northeast-1"  
  }  
]  
}
```

Request EC2 instance limits for home region (ap-northeast-1) with remote location (us-west-2)

In this example, we request the instance limits and current usage for the `m5.large` instance type at a remote location. The API call is made to the AWS Region `ap-northeast-1` with `Location` set to `us-west-2`. The results show limits and usage for fleets that reside in `ap-northeast-1`, for instances deployed in `us-west-2`. The response indicates that our account allows up to 100 `m5.large` instances and currently has 10 deployed.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{  
  "EC2InstanceType": "m5.large",  
  "Location": "us-west-2"  
}
```

Sample Response

```
{  
  "EC2InstanceLimits": [  
    {  
      "EC2InstanceType": "m5.large",  
      "CurrentInstances": 10,  
      "InstanceLimit": 100,  
      "Location": "us-west-2"  
    }  
  ]  
}
```

Request EC2 instance limits for home region (us-east-1) with remote location (ap-northeast-1)

In this example, we request the instance limits and current usage for the `m5.large` instance type at a remote location. The API call is made to the AWS Region `us-east-1` with `Location` set to `ap-northeast-1`. The results show limits and usage for fleets that reside in `us-east-1`, for instances deployed in `ap-northeast-1`. These limits do not affect fleets in other Regions that also deploy instances to `ap-northeast-1`. The response indicates that our account allows up to 150 `m5.large` instances and currently has none deployed.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "EC2InstanceType": "m5.large",
  "Location": "ap-northeast-1"
}
```

Sample Response

```
{
  "EC2InstanceLimits": [
    {
      "EC2InstanceType": "m5.large",
      "CurrentInstances": 0,
      "InstanceLimit": 150,
      "Location": "ap-northeast-1"
    }
  ]
}
```

See Also

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

- [AWS Command Line Interface 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](#)

DescribeFleetAttributes

This API works with the following fleet types: EC2, Anywhere

Retrieves core fleet-wide properties for fleets in an AWS Region. Properties include the computing hardware and deployment configuration for instances in the fleet.

You can use this operation in the following ways:

- To get attributes for specific fleets, provide a list of fleet IDs or fleet ARNs.
- To get attributes for all fleets, do not provide a fleet identifier.

When requesting attributes for multiple fleets, use the pagination parameters to retrieve results as a set of sequential pages.

If successful, a `FleetAttributes` object is returned for each fleet requested, unless the fleet identifier is not found.

Note

Some API operations limit the number of fleet IDs that allowed in one request. If a request exceeds this limit, the request fails and the error message contains the maximum allowed number.

Learn more

[Setting up Amazon GameLift Servers fleets](#)

Request Syntax

```
{  
  "FleetIds": [ "string" ],  
  "Limit": number,  
  "NextToken": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

 Note

In the following list, the required parameters are described first.

FleetIds

A list of unique fleet identifiers to retrieve attributes for. You can use either the fleet ID or ARN value. To retrieve attributes for all current fleets, do not include this parameter.

Type: Array of strings

Array Members: Minimum number of 1 item.

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: No

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages. This parameter is ignored when the request specifies one or a list of fleet IDs.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value. This parameter is ignored when the request specifies one or a list of fleet IDs.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "FleetAttributes": [
    {
      "AnywhereConfiguration": {
        "Cost": "string"
      },
      "BuildArn": "string",
      "BuildId": "string",
      "CertificateConfiguration": {
        "CertificateType": "string"
      },
      "ComputeType": "string",
      "CreationTime": number,
      "Description": "string",
      "FleetArn": "string",
      "FleetId": "string",
      "FleetType": "string",
      "InstanceRoleArn": "string",
      "InstanceRoleCredentialsProvider": "string",
      "InstanceType": "string",
      "LogPaths": [ "string" ],
      "MetricGroups": [ "string" ],
      "Name": "string",
      "NewGameSessionProtectionPolicy": "string",
      "OperatingSystem": "string",
      "PlayerGatewayConfiguration": {
        "GameServerIpProtocolSupported": "string"
      },
      "PlayerGatewayMode": "string",
      "ResourceCreationLimitPolicy": {
        "NewGameSessionsPerCreator": number,
        "PolicyPeriodInMinutes": number
      },
      "ScriptArn": "string",
      "ScriptId": "string",
      "ServerLaunchParameters": "string",
      "ServerLaunchPath": "string",
```

```
    "Status": "string",
    "StoppedActions": [ "string" ],
    "TerminationTime": 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.

FleetAttributes

A collection of objects containing attribute metadata for each requested fleet ID. Attribute objects are returned only for fleets that currently exist.

Type: Array of [FleetAttributes](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Request attributes for a list of fleets

This example retrieves attributes for an EC2 fleet.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "FleetIds": [
    "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  ]
}
```

Sample Response

```
{
  "FleetAttributes": [
    {
      "BuildArn": "arn:aws:gamelift:us-west-2::build/build-3333cccc-44dd-55ee-66ff-00001111aa22",
    }
  ]
}
```

```
    "BuildId": "build-3333cccc-44dd-55ee-66ff-00001111aa22",
    "CertificateConfiguration": {
      "CertificateType": "DISABLED"
    },
    "ComputeType": "EC2",
    "CreationTime": 1568836191.995,
    "Description": "On-demand hosts for v2 North America",
    "FleetArn": "arn:aws:gamelift:us-west-2::fleet/fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "FleetType": "SPOT",
    "InstanceType": "c5.large",
    "MetricGroups": [
      "default"
    ],
    "Name": "MegaFrogRaceServer.NA.v2-spot",
    "NewGameSessionProtectionPolicy": "NoProtection",
    "OperatingSystem": "WINDOWS_2022",
    "Status": "ACTIVE"
  }
]
```

Request attributes for all fleets

This example returns fleet attributes for all fleets with any status. This example uses the pagination parameters to return one fleet at a time.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "Limit": 1
  "NextToken":
    "eyJhd3NBZDZlNvdW50SWQlOmsicyI6IjMwMjc3NjAxNjM5OCJ9LCJldWlsZEIkdjIjP7InMiOiJldWlsZC01NWYxZTZmMS1jY"
}
```

Sample Response

```
{
  "FleetAttributes": [
```

```

    {
      "FleetId": "fleet-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
      "FleetArn": "arn:aws:gamelift:us-west-2::fleet/
fleet-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
      "FleetType": "SPOT",
      "InstanceType": "c4.large",
      "Description": "On-demand hosts for v2 North America",
      "Name": "MegaFrogRaceServer.NA.v2-spot",
      "CreationTime": 1568838275.379,
      "Status": "ACTIVATING",
      "BuildId": "build-3333cccc-44dd-55ee-66ff-00001111aa22",
      "BuildArn": "arn:aws:gamelift:us-west-2::build/
build-3333cccc-44dd-55ee-66ff-00001111aa22",
      "ServerLaunchPath": "C:\\\\game\\\\MegaFrogRace_Server.exe",
      "NewGameSessionProtectionPolicy": "NoProtection",
      "OperatingSystem": "WINDOWS_2022",
      "MetricGroups": [
        "default"
      ],
      "CertificateConfiguration": {
        "CertificateType": "GENERATED"
      }
    }
  ],
  "NextToken":
"eyJhd3NBZjY2NvdW50SWQ10nsicyI6IjQwMTY4MDEwMjY5NCJ9LCJmbGVldElkIjpw7InMiOiJmbGVldC00ZjcyY2E4ZS1iM
}

```

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](#)

DescribeFleetCapacity

This API works with the following fleet types: EC2, Container

Retrieves the resource capacity settings for one or more fleets. For a container fleet, this operation also returns counts for game server container groups.

With multi-location fleets, this operation retrieves data for the fleet's home Region only. To retrieve capacity for remote locations, see https://docs.aws.amazon.com/gamelift/latest/apireference/API_DescribeFleetLocationCapacity.html.

This operation can be used in the following ways:

- To get capacity data for one or more specific fleets, provide a list of fleet IDs or fleet ARNs.
- To get capacity data for all fleets, do not provide a fleet identifier.

When requesting multiple fleets, use the pagination parameters to retrieve results as a set of sequential pages.

If successful, a `FleetCapacity` object is returned for each requested fleet ID. Each `FleetCapacity` object includes a `Location` property, which is set to the fleet's home Region. Capacity values are returned only for fleets that currently exist.

Note

Some API operations may limit the number of fleet IDs that are allowed in one request. If a request exceeds this limit, the request fails and the error message includes the maximum allowed.

Learn more

[Setting up Amazon GameLift Servers fleets](#)

[GameLift metrics for fleets](#)

Request Syntax

```
{
```



```
"FleetIds": [ "string" ],  
"Limit": number,  
"NextToken": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetIds

A unique identifier for the fleet to retrieve capacity information for. You can use either the fleet ID or ARN value. Leave this parameter empty to retrieve capacity information for all fleets.

Type: Array of strings

Array Members: Minimum number of 1 item.

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\-]+|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\-]+`

Required: No

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages. This parameter is ignored when the request specifies one or a list of fleet IDs.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value. This parameter is ignored when the request specifies one or a list of fleet IDs.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "FleetCapacity": [
    {
      "FleetArn": "string",
      "FleetId": "string",
      "GameServerContainerGroupCounts": {
        "ACTIVE": number,
        "IDLE": number,
        "PENDING": number,
        "TERMINATING": number
      },
      "InstanceCounts": {
        "ACTIVE": number,
        "DESIRED": number,
        "IDLE": number,
        "MAXIMUM": number,
        "MINIMUM": number,
        "PENDING": number,
        "TERMINATING": number
      },
      "InstanceType": "string",
      "Location": "string",
      "ManagedCapacityConfiguration": {
        "ScaleInAfterInactivityMinutes": number,
        "ZeroCapacityStrategy": "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.

FleetCapacity

A collection of objects that contains capacity information for each requested fleet ID. Capacity objects are returned only for fleets that currently exist. Changes in desired instance value can take up to 1 minute to be reflected.

Type: Array of [FleetCapacity](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Request capacity status for a list of fleets

This example retrieves fleet capacity information for a list of two fleets. The first result shows capacity for a container fleet that's configured to hold five game server container groups per instance. The second result shows a fleet in the middle of a scale down event: instances are being terminated so that the active instances count matches the desired instances count.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "FleetIds": [
    "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "arn:aws:gamelift:us-west-2::fleet/fleet-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
  ]
}
```

Sample Response

```
{
  "FleetCapacity": [
    {
```

```
    "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "FleetArn": "arn:aws:gamelift:us-east-2::fleet/
fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "InstanceCounts": {
        "DESIRED": 10,
        "MINIMUM": 1,
        "MAXIMUM": 20,
        "PENDING": 0,
        "ACTIVE": 10,
        "IDLE": 3,
        "TERMINATING": 0
    },
    "InstanceType": "c5.large",
    "Location": "us-west-2",
    "GameServerContainerGroupCounts": {
        "ACTIVE": 50,
        "IDLE": 15,
        "PENDING": 0,
        "TERMINATING": 0
    },
},
{
    "FleetId": "fleet-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "FleetArn": "arn:aws:gamelift:us-east-2::fleet/
fleet-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "InstanceCounts": {
        "DESIRED": 13,
        "MINIMUM": 1,
        "MAXIMUM": 20,
        "PENDING": 0,
        "ACTIVE": 15,
        "IDLE": 2,
        "TERMINATING": 2
    }
    "InstanceType": "c5.large",
    "Location": "us-west-2"
}
]
```

See Also

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

- [AWS Command Line Interface 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](#)

DescribeFleetDeployment

This API works with the following fleet types: Container

Retrieves information about a managed container fleet deployment.

Request options

- Get information about the latest deployment for a specific fleet. Provide the fleet ID or ARN.
- Get information about a specific deployment. Provide the fleet ID or ARN and the deployment ID.

Results

If successful, a `FleetDeployment` object is returned.

Request Syntax

```
{  
  "DeploymentId": "string",  
  "FleetId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the container fleet. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\-]+|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\-]+`

Required: Yes

DeploymentId

A unique identifier for the deployment to return information for.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `^[a-zA-Z0-9\-\-]+`

Required: No

Response Syntax

```
{
  "FleetDeployment": {
    "CreationTime": number,
    "DeploymentConfiguration": {
      "ImpairmentStrategy": "string",
      "MinimumHealthyPercentage": number,
      "ProtectionStrategy": "string"
    },
    "DeploymentId": "string",
    "DeploymentStatus": "string",
    "FleetId": "string",
    "GameServerBinaryArn": "string",
    "PerInstanceBinaryArn": "string",
    "RollbackGameServerBinaryArn": "string",
    "RollbackPerInstanceBinaryArn": "string"
  },
  "LocationalDeployments": {
    "string" : {
      "DeploymentStatus": "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.

FleetDeployment

The requested deployment information.

Type: [FleetDeployment](#) object

LocationalDeployments

If the deployment is for a multi-location fleet, the requests returns the deployment status in each fleet location.

Type: String to [LocationalDeployment](#) object map

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

Key Pattern: `^[a-zA-Z0-9\-\-]+$`

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Retrieve information on a fleet deployment

This example gets information on a specific deployment for a container fleet.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "FleetId": "containerfleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "DeploymentId": "deployment-3333aaaa-44bb-55cc-66dd-7777eeee88ff"
}
```

Sample Response

```
{
  "FleetDeployment": {
    "CreationTime": 1736365885.22,
    "DeploymentConfiguration": {
      "ImpairmentStrategy": "ROLLBACK",
      "MinimumHealthyPercentage": 30,
      "ProtectionStrategy": "WITH_PROTECTION"
    },
    "DeploymentId": "deployment-3333aaaa-44bb-55cc-66dd-7777eeee88ff",
    "DeploymentStatus": "COMPLETE",
    "FleetId": "containerfleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  }
}
```

```
"GameServerBinaryArn": "arn:aws:gamelift:us-west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:2",
  "RollbackGameServerBinaryArn": "arn:aws:gamelift:us-west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:1",
}
```

See Also

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

- [AWS Command Line Interface 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](#)

DescribeFleetEvents

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves entries from a fleet's event log. Fleet events are initiated by changes in status, such as during fleet creation and termination, changes in capacity, etc. If a fleet has multiple locations, events are also initiated by changes to status and capacity in remote locations.

You can specify a time range to limit the result set. Use the pagination parameters to retrieve results as a set of sequential pages.

If successful, a collection of event log entries matching the request are returned.

Learn more

[Setting up Amazon GameLift Servers fleets](#)

Request Syntax

```
{  
  "EndTime": number,  
  "FleetId": "string",  
  "Limit": number,  
  "NextToken": "string",  
  "StartTime": number  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet to get event logs for. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

EndTime

The most recent date to retrieve event logs for. If no end time is specified, this call returns entries from the specified start time up to the present. Format is a number expressed in Unix time as milliseconds (ex: "1469498468.057").

Type: Timestamp

Required: No

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

StartTime

The earliest date to retrieve event logs for. If no start time is specified, this call returns entries starting from when the fleet was created to the specified end time. Format is a number expressed in Unix time as milliseconds (ex: "1469498468.057").

Type: Timestamp

Required: No

Response Syntax

```
{
  "Events": [
    {
      "Count": number,
      "EventCode": "string",
      "EventId": "string",
      "EventTime": number,
      "Message": "string",
      "PreSignedLogUrl": "string",
      "ResourceId": "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.

Events

A collection of objects containing event log entries for the specified fleet.

Type: Array of [Event](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Request events for a specified time span

This example returns all events that occurred for the fleet during the specified time span. In this example, the time span is on January 21, 2020, from 15:45:00 to 16:15:00 (local time).

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "StartTime": 1579647600,
  "EndTime": 1579649400,
  "Limit": 5
}
```

Sample Response

```
{
  "Events": [
    {
      "EventId": "a37b6892-5d07-4d3b-8b47-80244ecf66b9",
      "ResourceId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
      "EventCode": "FLEET_STATE_ACTIVE",
      "Message": "Fleet fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa changed state to ACTIVE",
      "EventTime": 1579649342.191
    },
    {
      "EventId": "67da4ec9-92a3-4d95-886a-5d6772c24063",
      "ResourceId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
      "EventCode": "FLEET_STATE_ACTIVATING",
      "Message": "Fleet fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa changed state to ACTIVATING",
      "EventTime": 1579649321.427
    },
    {
      "EventId": "23813a46-a9e6-4a53-8847-f12e6a8381ac",
      "ResourceId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
      "EventCode": "FLEET_STATE_BUILDING",
      "Message": "Fleet fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa changed state to BUILDING",
      "EventTime": 1579649321.243
    },
    {
      "EventId": "3bf217d0-1d44-42f9-9202-433ed475d2e8",
      "ResourceId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
      "EventCode": "FLEET_STATE_VALIDATING",
      "Message": "Fleet fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa changed state to VALIDATING",
    }
  ]
}
```



```

      "EventTime": 1579649197.449
    },
    {
      "EventId": "2ecd0130-5986-44eb-99a7-62df27741084",
      "ResourceId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
      "EventCode": "FLEET_VALIDATION_LAUNCH_PATH_NOT_FOUND",
      "Message": "Failed to find a valid path",
      "EventTime": 1569319075.839,
      "PreSignedLogUrl": "https://gamelift-event-logs-prod-us-west-2.s3.us-
west-2.amazonaws.com/logs/fleet-83422059-8329-42a2-a4d6-c4444386a6f8/
events/2ecd0130-5986-44eb-99a7-62df27741084/FLEET_VALIDATION_LAUNCH_PATH_NOT_FOUND.txt?
X-Amz-Security-Token=IQoJb3JpZ2luX2VjEB8aCXVzLXdlc3QtMiJHMEUCIHV5K
%2FLPx8h310D%2FAvx0%2FZxsDy5XA3cJ0wPdu3T0eBa%2FAiEA1yovokcZYy
%2FV4CWW6l26aFyiSH0%2Bxz%2FBMAhEHYHMQNcqkQMImP%2F%2F%2F%2F%2F
%2F%2F%2F%2FARAAGgw3NDEwNjE1OTIxNzEiDI8rsZtzLzlwEDQhXSrlAtl5Ae
%2Fgo6FCIzqXPbXfb0nSvFYqeDlriZarEpKqKrUt8mXQv9iqHResqCph9AKo49lwgSYTT2QoSxnrD7%2FUgv
%2BZm2pVuczvuKtUA0fcx6s0GxpjIAzdIE%2F5P%2FB7B9M%2BVZ
%2F9KF82hbJi0HTE6Y7BjKsEgFCvk4UXILhfjtan9iQl8%2F21ZTurAcJbm7Y5tuLF9SWSK3%2BEa7VX0cCK4D401sMjmdR
%2FIaXoHkNvg0RVTa0hIqdvpaDQlsSBNdqTXbjHTu6fETE9Y9Ky%2BiJK5KiUG
%2F59GjCpDcvS1FqKeLUEmKT7wysGmvjMc2n%2Fr
%2F9VxQfte7w9srXw1LAQuwhiXAAyI5ICMZ5JvzjzQwTqD4CHTVKUUDwL%2BRZzbuuqkJ0bZml02CkRGp
%2B74RTAzLbWptVqZTIzfzctiCTmWxb%2FmKyELRYsVLrwNJ%2BGJ7%2BCrN0RC%2FjlgfLYIZyeAqjPgAu5HjgX
%2BM7jCo9M7wBTInAXK0FQuf9dvA84SuwX0JFp17LYGjrHMKv0qC3GfbTMrZ6kzeNV9awKCpXB2Gnx9z2KvI1JdqirWvpvH
%2F9C6%2B4jIZPME3jXmZcEHqqw5uvAVF7aeIavtUZU8pxpDIWT0YE4p3Kriy2AA7ziCRKtVfjV839InyLk8LUjsioWK2qJl
%2BYUq8%2FDTL1Lxqj1S%2Fi04TI0Wo7ilAo%2FKKWWF4guuNDexj8E00ynSp1yImB
%2BZf2Fua3044W4eEXAMPLE33333&X-Amz-Algorithm=AWS4-HMAC-SHA256&X-Amz-
Date=20170621T231808Z&X-Amz-SignedHeaders=host&X-Amz-Expires=900&X-Amz-
Credential=AKIAIOSFODNN7EXAMPLE%2F20170621%2Fus-west-2%2Fs3%2Faws4_request&X-Amz-
Signature=wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY"
    }
  ],
  "NextToken":
    "eyJhd3NBZj2NvdW50SWQ0i0nsicyI6IjMwMjc3NjAxNjM50CJ9LCJidWlsZE1kIjpw7InMi0iJidWlsZC01NWYxZTZmMS1jY
}

```

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](#)

DescribeFleetLocationAttributes

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves information on a fleet's remote locations, including life-cycle status and any suspended fleet activity.

This operation can be used in the following ways:

- To get data for specific locations, provide a fleet identifier and a list of locations. Location data is returned in the order that it is requested.
- To get data for all locations, provide a fleet identifier only. Location data is returned in no particular order.

When requesting attributes for multiple locations, use the pagination parameters to retrieve results as a set of sequential pages.

If successful, a `LocationAttributes` object is returned for each requested location. If the fleet does not have a requested location, no information is returned.

Learn more

[Setting up Amazon GameLift Servers fleets](#)

[Amazon GameLift Servers service locations](#) for managed hosting

Request Syntax

```
{  
  "FleetId": "string",  
  "Limit": number,  
  "Locations": [ "string" ],  
  "NextToken": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet to retrieve remote locations for. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages. This limit is not currently enforced.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

Locations

A list of fleet locations to retrieve information for. Specify locations in the form of an AWS Region code, such as `us-west-2`.

Type: Array of strings

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

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-\_]+\`

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "FleetArn": "string",
  "FleetId": "string",
  "LocationAttributes": [
    {
      "LocationState": {
        "Location": "string",
        "PlayerGatewayStatus": "string",
        "Status": "string"
      },
      "StoppedActions": [ "string" ],
      "UpdateStatus": "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.

FleetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is

arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: ^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\-]+\

FleetId

A unique identifier for the fleet that location attributes were requested for.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: ^[a-z]*fleet-[a-zA-Z0-9\-\-]+\|arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\-]+\

LocationAttributes

Location-specific information on the requested fleet's remote locations.

Type: Array of [LocationAttributes](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Retrieve remote fleet locations

This example retrieves information on all remote locations for a fleet. The requested fleet's home Region is `us-west-2`. It can deploy instances in the following AWS Regions: `us-west-2`, `us-west-1`, and `ca-central-1`. In this example, auto-scaling has been suspended in `ca-central-1`, and there is a fleet update that has not yet been completed for that location.

Sample Request

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
}
```

Sample Response

```
{
```

```
"FleetArn": "arn:aws:gamelift:us-west-2::fleet/
fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
"FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
"LocationAttributes": [
  {
    "LocationState": {
      "Location": "us-west-1",
      "Status": "ACTIVE"
    }
  },
  {
    "LocationState": {
      "Location": "ca-central-1",
      "Status": "ACTIVE"
    },
    "StoppedActions": [ "AUTO_SCALING" ],
    "UpdateStatus": "PENDING_UPDATE"
  }
],
"NextToken": "string"
}
```

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](#)

DescribeFleetLocationCapacity

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves the resource capacity settings for a fleet location. The data returned includes the current capacity (number of EC2 instances) and some scaling settings for the requested fleet location. For a managed container fleet, this operation also returns counts for game server container groups.

Use this operation to retrieve capacity information for a fleet's remote location or home Region (you can also retrieve home Region capacity by calling `DescribeFleetCapacity`).

To retrieve capacity data, identify a fleet and location.

If successful, a `FleetCapacity` object is returned for the requested fleet location.

Learn more

[Setting up Amazon GameLift Servers fleets](#)

[Amazon GameLift Servers service locations](#) for managed hosting

[GameLift metrics for fleets](#)

Request Syntax

```
{
  "FleetId": "string",
  "Location": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet to request location capacity for. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\-]+|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\-]+`

Required: Yes

Location

The fleet location to retrieve capacity information for. Specify a location in the form of an AWS Region code, such as `us-west-2`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-\-]+`

Required: Yes

Response Syntax

```
{
  "FleetCapacity": {
    "FleetArn": "string",
    "FleetId": "string",
    "GameServerContainerGroupCounts": {
      "ACTIVE": number,
      "IDLE": number,
      "PENDING": number,
      "TERMINATING": number
    },
    "InstanceCounts": {
      "ACTIVE": number,
      "DESIRED": number,
      "IDLE": number,
```

```
    "MAXIMUM": number,
    "MINIMUM": number,
    "PENDING": number,
    "TERMINATING": number
  },
  "InstanceType": "string",
  "Location": "string",
  "ManagedCapacityConfiguration": {
    "ScaleInAfterInactivityMinutes": number,
    "ZeroCapacityStrategy": "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.

[FleetCapacity](#)

Resource capacity information for the requested fleet location. Capacity objects are returned only for fleets and locations that currently exist. Changes in desired instance value can take up to 1 minute to be reflected.

Type: [FleetCapacity](#) object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Request capacity information for a fleet location

This example retrieves fleet capacity information for a remote fleet location. The returned `FleetCapacity` object describes the capacity at the specified location only. To get a fleet's current total capacity, first call `DescribeFleetLocationAttributes` to get a list of the fleet's remote locations, and then call `DescribeFleetLocationCapacity` for each remote location plus the home Region, and sum the results for active instances.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "Location": "sa-east-1"
}
```

Sample Response

```
{
```

```
    "FleetCapacity": {
      "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
      "FleetArn": "arn:aws:gamelift:us-west-2::fleet/
fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
      "InstanceCounts": {
        "DESIRED": 10,
        "MINIMUM": 1,
        "MAXIMUM": 20,
        "PENDING": 0,
        "ACTIVE": 10,
        "IDLE": 3,
        "TERMINATING": 0
      },
      "InstanceType": "c5.large",
      "Location": "sa-east-1"
    }
  ]
}
```

See Also

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

- [AWS Command Line Interface 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](#)

DescribeFleetLocationUtilization

This API works with the following fleet types: EC2, Anywhere

Retrieves current usage data for a fleet location. Utilization data provides a snapshot of current game hosting activity at the requested location. Use this operation to retrieve utilization information for a fleet's remote location or home Region (you can also retrieve home Region utilization by calling `DescribeFleetUtilization`).

To retrieve utilization data, identify a fleet and location.

If successful, a `FleetUtilization` object is returned for the requested fleet location.

Learn more

[Setting up Amazon GameLift Servers fleets](#)

[Amazon GameLift Servers service locations](#) for managed hosting

[GameLift metrics for fleets](#)

Request Syntax

```
{
  "FleetId": "string",
  "Location": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet to request location utilization for. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

Location

The fleet location to retrieve utilization information for. Specify a location in the form of an AWS Region code, such as `us-west-2`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-\_]+\$`

Required: Yes

Response Syntax

```
{
  "FleetUtilization": {
    "ActiveGameSessionCount": number,
    "ActiveServerProcessCount": number,
    "CurrentPlayerSessionCount": number,
    "FleetArn": "string",
    "FleetId": "string",
    "Location": "string",
    "MaximumPlayerSessionCount": number
  }
}
```

Response Elements

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

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

FleetUtilization

Utilization information for the requested fleet location. Utilization objects are returned only for fleets and locations that currently exist.

Type: [FleetUtilization](#) object

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Request utilization for a fleet location

This example retrieves usage data for the remote fleet location sa-east-1. The fleet's home Region is us-west-2.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "FleetId": "arn:aws:gamelift:us-west-2::fleet/
fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "Location": "sa-east-1"
}
```

Sample Response

```
{
  "FleetUtilization": {
    "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "FleetArn": "arn:aws:gamelift:us-west-2::fleet/
fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "ActiveServerProcessCount": 100,
    "ActiveGameSessionCount": 62,
    "CurrentPlayerSessionCount": 329,
    "MaximumPlayerSessionCount": 1000,
    "Location": "sa-east-1"
  }
}
```

See Also

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

- [AWS Command Line Interface 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](#)

DescribeFleetPortSettings

This API works with the following fleet types: EC2

Retrieves a fleet's inbound connection permissions. Connection permissions specify IP addresses and port settings that incoming traffic can use to access server processes in the fleet. Game server processes that are running in the fleet must use a port that falls within this range.

Use this operation in the following ways:

- To retrieve the port settings for a fleet, identify the fleet's unique identifier.
- To check the status of recent updates to a fleet remote location, specify the fleet ID and a location. Port setting updates can take time to propagate across all locations.

If successful, a set of `IpPermission` objects is returned for the requested fleet ID. When specifying a location, this operation returns a pending status. If the requested fleet has been deleted, the result set is empty.

Learn more

[Setting up Amazon GameLift Servers fleets](#)

Request Syntax

```
{  
  "FleetId": "string",  
  "Location": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet to retrieve port settings for. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

Location

A remote location to check for status of port setting updates. Use the AWS Region code format, such as `us-west-2`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-\_]+\$`

Required: No

Response Syntax

```
{
  "FleetArn": "string",
  "FleetId": "string",
  "InboundPermissions": [
    {
      "FromPort": number,
      "IpRange": "string",
      "Protocol": "string",
      "ToPort": number
    }
  ],
  "Location": "string",
  "UpdateStatus": "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.

FleetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet[/[a-z]*fleet-[a-zA-Z0-9\ -]]+$`

FleetId

A unique identifier for the fleet that was requested.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\ -]+$`

InboundPermissions

The port settings for the requested fleet ID.

Type: Array of [IpPermission](#) objects

Array Members: Maximum number of 50 items.

Location

The requested fleet location, expressed as an AWS Region code, such as `us-west-2`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\ -]+$`

UpdateStatus

The current status of updates to the fleet's port settings in the requested fleet location. A status of `PENDING_UPDATE` indicates that an update was requested for the fleet but has not yet been completed for the location.

Type: String

Valid Values: `PENDING_UPDATE`

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Request inbound connection permissions for a fleet

This example retrieves connection settings for a specified fleet.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
}
```

Sample Response

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "FleetArn": "arn:aws:gamelift:us-east-2::fleet/fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "InboundPermissions": [
    {
      "FromPort": 33400,
      "ToPort": 33500,
      "IpRange": "0.0.0.0/0",
      "Protocol": "UDP"
    },
    {
      "FromPort": 1900,
      "ToPort": 2000,
      "IpRange": "0.0.0.0/0",
      "Protocol": "TCP"
    }
  ]
}
```

Check port setting updates in a remote location

This example retrieves the current status of recent port setting updates for a specified remote location.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "Location": "us-west-2"
}
```

Sample Response

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "FleetArn": "arn:aws:gamelift:us-east-2::fleet/fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "InboundPermissions": [
    {
      "FromPort": 33400,
      "ToPort": 33500,
      "IpRange": "0.0.0.0/0",
      "Protocol": "UDP"
    },
    {
      "FromPort": 1900,
      "ToPort": 2000,
      "IpRange": "0.0.0.0/0",
      "Protocol": "TCP"
    }
  ],
  "Location": "us-west-2",
  "Status": "PENDING_UPDATE"
}
```

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](#)

DescribeFleetUtilization

This API works with the following fleet types: EC2

Retrieves utilization statistics for one or more fleets. Utilization data provides a snapshot of how the fleet's hosting resources are currently being used. For fleets with remote locations, this operation retrieves data for the fleet's home Region only. See [DescribeFleetLocationUtilization](#) to get utilization statistics for a fleet's remote locations.

This operation can be used in the following ways:

- To get utilization data for one or more specific fleets, provide a list of fleet IDs or fleet ARNs.
- To get utilization data for all fleets, do not provide a fleet identifier.

When requesting multiple fleets, use the pagination parameters to retrieve results as a set of sequential pages.

If successful, a [FleetUtilization](#) object is returned for each requested fleet ID, unless the fleet identifier is not found. Each fleet utilization object includes a `Location` property, which is set to the fleet's home Region.

Note

Some API operations may limit the number of fleet IDs allowed in one request. If a request exceeds this limit, the request fails and the error message includes the maximum allowed.

Learn more

[Setting up Amazon GameLift Servers Fleets](#)

[GameLift Metrics for Fleets](#)

Request Syntax

```
{
  "FleetIds": [ "string" ],
  "Limit": number,
  "NextToken": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetIds

A unique identifier for the fleet to retrieve utilization data for. You can use either the fleet ID or ARN value. To retrieve attributes for all current fleets, do not include this parameter.

Type: Array of strings

Array Members: Minimum number of 1 item.

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: No

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages. This parameter is ignored when the request specifies one or a list of fleet IDs.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value. This parameter is ignored when the request specifies one or a list of fleet IDs.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "FleetUtilization": [
    {
      "ActiveGameSessionCount": number,
      "ActiveServerProcessCount": number,
      "CurrentPlayerSessionCount": number,
      "FleetArn": "string",
      "FleetId": "string",
      "Location": "string",
      "MaximumPlayerSessionCount": 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.

FleetUtilization

A collection of objects containing utilization information for each requested fleet ID. Utilization objects are returned only for fleets that currently exist.

Type: Array of [FleetUtilization](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Request utilization data for a list of fleets

This example retrieves current usage information for a specific fleet. The fleet's home Region is us-west-2.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "FleetIds": ["fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"]
}
```

Sample Response

```
{
  "FleetUtilization": [
    {
      "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
      "FleetArn": "arn:aws:gamelift:us-west-2::fleet/fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
      "ActiveServerProcessCount": 100,
      "ActiveGameSessionCount": 62,
      "CurrentPlayerSessionCount": 329,
      "MaximumPlayerSessionCount": 1000,
      "Location": "us-west-2"
    }
  ]
}
```

Request utilization data for all fleets

This example returns fleet usage data for all fleets with any status. This example uses the pagination parameters to return results for two fleets at a time. The home Region for both fleets is `us-west-2`.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "Limit": 2
  "NextToken":
    "eyJhd3NBZjY2NvdW50SWQi0nsicyI6IjMwMjc3NjAxNjM5MCJ9LCJidWlsZElkIjpw7InMi0iJidWlsZC01NWYxZTZmMS1jY
}
```

Sample Response

```
{
  "FleetUtilization": [
    {
      "FleetId": "fleet-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
      "FleetArn": "arn:aws:gamelift:us-west-2::fleet/fleet-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
      "ActiveServerProcessCount": 100,
      "ActiveGameSessionCount": 13,
      "CurrentPlayerSessionCount": 98,
      "MaximumPlayerSessionCount": 1000,
      "Location": "us-west-2"
    },
    {
      "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
      "FleetArn": "arn:aws:gamelift:us-west-2::fleet/fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
      "ActiveServerProcessCount": 100,
      "ActiveGameSessionCount": 62,
      "CurrentPlayerSessionCount": 329,
      "MaximumPlayerSessionCount": 1000,
      "Location": "us-west-2"
    }
  ],
  "NextToken": "eyJhd3NBWY2NvdW50SWQi0nsicyI6IjMwMjc3NjAxNjM5OCJ9LCJidWlsZEIkJjY7InMiOiJidWlsZC01NWYxZTZmMS1jY"
}
```

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](#)

DescribeGameServer

This API works with the following fleet types: EC2 (FleetIQ)

Retrieves information for a registered game server. Information includes game server status, health check info, and the instance that the game server is running on.

To retrieve game server information, specify the game server ID. If successful, the requested game server object is returned.

Learn more

[Amazon GameLift Servers FleetIQ Guide](#)

Request Syntax

```
{
  "GameServerGroupName": "string",
  "GameServerId": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

GameServerGroupName

A unique identifier for the game server group where the game server is running.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]+\|^arn:.*:gameservergroup\[a-zA-Z0-9-\.]+\`

Required: Yes

GameServerId

A custom string that uniquely identifies the game server information to be retrieved.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: [a-zA-Z0-9-\.\.]+

Required: Yes

Response Syntax

```
{
  "GameServer": {
    "ClaimStatus": "string",
    "ConnectionInfo": "string",
    "GameServerData": "string",
    "GameServerGroupArn": "string",
    "GameServerGroupName": "string",
    "GameServerId": "string",
    "InstanceId": "string",
    "LastClaimTime": number,
    "LastHealthCheckTime": number,
    "RegistrationTime": number,
    "UtilizationStatus": "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.

GameServer

Object that describes the requested game server.

Type: GameServer object

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Retrieve details for a game server

This example retrieves attributes for a specified game server.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "GameServerGroupName": "MegaFrogServers_NA",
```

```
"GameServerId": "mega-frog-game-12345678"
}
```

CLI command:

```
aws gamelift describe-game-server \
  --game-server-group-name MegaFrogServers_NA \
  --game-server-id mega-frog-game-12345678
```

Sample Response

```
{
  "GameServer": {
    "ClaimStatus": "",
    "ConnectionInfo": "192.0.2.0.80",
    "GameServerData": "",
    "GameServerGroupArn": "arn:aws:gamelift:us-west-2::GameServerGroup/MegaFrogServers_NA",
    "GameServerGroupName": "MegaFrogServers_NA",
    "GameServerId": "mega-frog-game-12345678",
    "InstanceId": "i-1234567890abcdef0",
    "LastClaimTime": 1580218197.293,
    "LastHealthCheckTime": 1580218197.293,
    "RegistrationTime": 1580218197.293,
    "UtilizationStatus": "AVAILABLE"
  }
}
```

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](#)

DescribeGameServerGroup

This API works with the following fleet types: EC2 (FleetIQ)

Retrieves information on a game server group. This operation returns only properties related to Amazon GameLift Servers FleetIQ. To view or update properties for the corresponding Auto Scaling group, such as launch template, auto scaling policies, and maximum/minimum group size, access the Auto Scaling group directly.

To get attributes for a game server group, provide a group name or ARN value. If successful, a `GameServerGroup` object is returned.

Learn more

[Amazon GameLift Servers FleetIQ Guide](#)

Request Syntax

```
{  
  "GameServerGroupName": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[GameServerGroupName](#)

A unique identifier for the game server group. Use either the name or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]+\|^arn:.*:gameservergroup\[a-zA-Z0-9-\.]+\`

Required: Yes

Response Syntax

```
{
  "GameServerGroup": {
    "AutoScalingGroupArn": "string",
    "BalancingStrategy": "string",
    "CreationTime": number,
    "GameServerGroupArn": "string",
    "GameServerGroupName": "string",
    "GameServerProtectionPolicy": "string",
    "InstanceDefinitions": [
      {
        "InstanceType": "string",
        "WeightedCapacity": "string"
      }
    ],
    "LastUpdatedTime": number,
    "RoleArn": "string",
    "Status": "string",
    "StatusReason": "string",
    "SuspendedActions": [ "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.

GameServerGroup

An object with the property settings for the requested game server group resource.

Type: [GameServerGroup](#) object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Retrieve a game server group

This example retrieves information about a game server group by providing the group name.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "GameServerGroupName": "MegaFrogServers_NA"
}
```

CLI command:

```
aws gamelift describe-game-server-group \
```



```
--game-server-group-name MegaFrogServers_NA
```

Sample Response

```
{
  "GameServerGroup": {
    "AutoScalingGroupArn": "arn:aws:autoscaling:us-west-2:123456789012:autoScalingGroup:1111aaaa-22bb-33cc-44dd-5555eeee66ff:autoScalingGroupName/MegaFrogServers_NA",
    "BalancingStrategy": "SPOT_PREFERRED",
    "CreationTime": 1496365885.44,
    "GameServerGroupArn": "arn:aws:gamelift:us-west-2::GameServerGroup/MegaFrogServers_NA",
    "GameServerGroupName": " MegaFrogServers_NA",
    "GameServerProtectionPolicy": "NO_PROTECTION",
    "InstanceDefinitions": [
      {
        "InstanceType": "c5.2xlarge",
        "WeightedCapacity": "1"
      },
      {
        "InstanceType": "c5.4xlarge",
        "WeightedCapacity": "2"
      }
    ],
    "LastUpdatedTime": 1496365885.44,
    "RoleArn": "arn:aws:iam:123456789012::role/GameLiftGsgRole",
    "Status": "ACTIVE",
    "StatusReason": "",
    "SuspendedActions": []
  }
}
```

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](#)

DescribeGameServerInstances

This API works with the following fleet types: EC2 (FleetIQ)

Retrieves status information about the Amazon EC2 instances associated with a Amazon GameLift Servers FleetIQ game server group. Use this operation to detect when instances are active or not available to host new game servers.

To request status for all instances in the game server group, provide a game server group ID only. To request status for specific instances, provide the game server group ID and one or more instance IDs. Use the pagination parameters to retrieve results in sequential segments. If successful, a collection of `GameServerInstance` objects is returned.

This operation is not designed to be called with every game server claim request; this practice can cause you to exceed your API limit, which results in errors. Instead, as a best practice, cache the results and refresh your cache no more than once every 10 seconds.

Learn more

[Amazon GameLift Servers FleetIQ Guide](#)

Request Syntax

```
{  
  "GameServerGroupName": "string",  
  "InstanceIds": [ "string" ],  
  "Limit": number,  
  "NextToken": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

GameServerGroupName

A unique identifier for the game server group. Use either the name or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]|^arn:.*:gameservergroup\[a-zA-Z0-9-\.]`

Required: Yes

InstanceIds

The Amazon EC2 instance IDs that you want to retrieve status on. Amazon EC2 instance IDs use a 17-character format, for example: `i-1234567890abcdef0`. To retrieve all instances in the game server group, leave this parameter empty.

Type: Array of strings

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

Length Constraints: Fixed length of 19.

Pattern: `^i-[0-9a-zA-Z]{17}`

Required: No

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "GameServerInstances": [
    {
      "GameServerGroupArn": "string",
      "GameServerGroupName": "string",
      "InstanceId": "string",
      "InstanceStatus": "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.

GameServerInstances

The collection of requested game server instances.

Type: Array of [GameServerInstance](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Retrieve game server instance status

This example retrieves the current status of the first twenty instances in a game server group. Since no instance IDs are specified in this request, there is no error information to be returned.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "GameServerGroupName": "MegaFrogServers_NA",
  "Limit": "20"
}
```

CLI command:

```
aws gamelift describe-game-server-instances \  
  --game-server-group-name MegaFrogServers_NA \  
  --limit 20
```

Sample Response

```
{  
  "GameServerInstances": [  
    {  
      "GameServerGroupArn": "arn:aws:gamelift:us-west-2::GameServerGroup/  
MegaFrogServers_NA",  
      "GameServerGroupName": " MegaFrogServers_NA",  
      "InstanceId": "i-1234567890abcdef0",  
      "InstanceStatus": "ACTIVE"  
    },  
    {  
      "GameServerGroupArn": "arn:aws:gamelift:us-west-2::GameServerGroup/  
MegaFrogServers_NA",  
      "GameServerGroupName": " MegaFrogServers_NA",  
      "InstanceId": "i-9876543210abcdef1",  
      "InstanceStatus": "DRAINING"  
    }  
  ],  
  "NextToken":  
    "eyJhd3NBZjY2NvdW50SWQ0I0nsicyI6IjMwMjc3NjAxNjM5OCJ9LCJidWlsZE1kIjpw7InMiOiJidWlsZC01NWYxZTZmMS1jY"  
}
```

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](#)

DescribeGameSessionDetails

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves additional game session properties, including the game session protection policy in force, a set of one or more game sessions in a specific fleet location. You can optionally filter the results by current game session status.

This operation can be used in the following ways:

- To retrieve details for all game sessions that are currently running on all locations in a fleet, provide a fleet or alias ID, with an optional status filter. This approach returns details from the fleet's home Region and all remote locations.
- To retrieve details for all game sessions that are currently running on a specific fleet location, provide a fleet or alias ID and a location name, with optional status filter. The location can be the fleet's home Region or any remote location.
- To retrieve details for a specific game session, provide the game session ID. This approach looks for the game session ID in all fleets that reside in the AWS Region defined in the request.

Use the pagination parameters to retrieve results as a set of sequential pages.

If successful, a `GameSessionDetail` object is returned for each game session that matches the request.

Learn more

[Find a game session](#)

[All APIs by task](#)

Request Syntax

```
{  
  "AliasId": "string",  
  "FleetId": "string",  
  "GameSessionId": "string",  
  "Limit": number,  
  "Location": "string",  
  "NextToken": "string",  
  "StatusFilter": "string"
```

```
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

AliasId

A unique identifier for the alias associated with the fleet to retrieve all game sessions for. You can use either the alias ID or ARN value.

Type: String

Pattern: `^alias-\S+|^arn:.*:alias\/alias-\S+`

Required: No

FleetId

A unique identifier for the fleet to retrieve all game sessions active on the fleet. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\`

Required: No

GameSessionId

An identifier for the game session that is unique across all regions to retrieve. The value is always a full ARN in the following format: For Home Region game session -

arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>. For Remote Location game session - arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<location>/<ID string>.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: [a-zA-Z0-9:/-]+

Required: No

Limit

The maximum number of results to return. Use this parameter with NextToken to get results as a set of sequential pages.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

Location

A fleet location to get game session details for. You can specify a fleet's home Region or a remote location. Use the AWS Region code format, such as us-west-2.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: ^[A-Za-z0-9\-]+

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

StatusFilter

Game session status to filter results on. Possible game session statuses include ACTIVE, TERMINATED, ACTIVATING and TERMINATING (the last two are transitory).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "GameSessionDetails": [
    {
      "GameSession": {
        "ComputeName": "string",
        "CreationTime": number,
        "CreatorId": "string",
        "CurrentPlayerSessionCount": number,
        "DnsName": "string",
        "FleetArn": "string",
        "FleetId": "string",
        "GameProperties": [
          {
            "Key": "string",
            "Value": "string"
          }
        ],
        "GameSessionData": "string",
        "GameSessionId": "string",
        "IpAddress": "string",
        "Location": "string",
        "MatchmakerData": "string",
        "MaximumPlayerSessionCount": number,
        "Name": "string",
        "PlayerGatewayStatus": "string",
        "PlayerSessionCreationPolicy": "string",
```

```
        "Port": number,
        "Status": "string",
        "StatusReason": "string",
        "TerminationTime": number
    },
    "ProtectionPolicy": "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.

GameSessionDetails

A collection of properties for each game session that matches the request.

Type: Array of [GameSessionDetail](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TerminalRoutingStrategyException

The service is unable to resolve the routing for a particular alias because it has a terminal RoutingStrategy associated with it. The message returned in this exception is the message defined in the routing strategy itself. Such requests should only be retried if the routing strategy for the specified alias is modified.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

DescribeGameSessionPlacement

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves information, including current status, about a game session placement request.

To get game session placement details, specify the placement ID.

This operation is not designed to be continually called to track game session status. This practice can cause you to exceed your API limit, which results in errors. Instead, you must configure an Amazon Simple Notification Service (SNS) topic to receive notifications from FlexMatch or queues. Continuously polling with `DescribeGameSessionPlacement` should only be used for games in development with low game session usage. For a reference implementation of event-based game session placement tracking, see [Event-based game session placement guidance](#) in the Amazon GameLift Toolkit.

Request Syntax

```
{  
  "PlacementId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

PlacementId

A unique identifier for a game session placement to retrieve.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 48.

Pattern: `[a-zA-Z0-9-]+`

Required: Yes

Response Syntax

```
{
  "GameSessionPlacement": {
    "DnsName": "string",
    "EndTime": number,
    "GameProperties": [
      {
        "Key": "string",
        "Value": "string"
      }
    ],
    "GameSessionArn": "string",
    "GameSessionData": "string",
    "GameSessionId": "string",
    "GameSessionName": "string",
    "GameSessionQueueName": "string",
    "GameSessionRegion": "string",
    "IpAddress": "string",
    "MatchmakerData": "string",
    "MaximumPlayerSessionCount": number,
    "PlacedPlayerSessions": [
      {
        "PlayerId": "string",
        "PlayerSessionId": "string"
      }
    ],
    "PlacementId": "string",
    "PlayerGatewayStatus": "string",
    "PlayerLatencies": [
      {
        "LatencyInMilliseconds": number,
        "PlayerId": "string",
        "RegionIdentifier": "string"
      }
    ],
    "Port": number,
```

```
    "PriorityConfigurationOverride": {  
        "LocationOrder": [ "string" ],  
        "PlacementFallbackStrategy": "string"  
    },  
    "StartTime": number,  
    "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.

GameSessionPlacement

Object that describes the requested game session placement.

Type: [GameSessionPlacement](#) object

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

DescribeGameSessionQueues

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves the properties for one or more game session queues. When requesting multiple queues, use the pagination parameters to retrieve results as a set of sequential pages. When specifying a list of queues, objects are returned only for queues that currently exist in the Region.

Learn more

[View Your Queues](#)

Request Syntax

```
{
  "Limit": number,
  "Names": [ "string" ],
  "NextToken": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages. You can request up to 50 results.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

Names

A list of queue names to retrieve information for. You can use either the queue ID or ARN value. To request settings for all queues, leave this parameter empty.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-]+|^arn:.*:gamesessionqueue\/[a-zA-Z0-9-]+`

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "GameSessionQueues": [
    {
      "CustomEventData": "string",
      "Destinations": [
        {
          "DestinationArn": "string"
        }
      ],
      "FilterConfiguration": {
        "AllowedLocations": [ "string" ]
      },
      "GameSessionQueueArn": "string",
```

```
"Name": "string",
"NotificationTarget": "string",
"PlayerLatencyPolicies": [
  {
    "MaximumIndividualPlayerLatencyMilliseconds": number,
    "PolicyDurationSeconds": number
  }
],
"PriorityConfiguration": {
  "LocationOrder": [ "string" ],
  "PriorityOrder": [ "string" ]
},
"TimeoutInSeconds": 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.

GameSessionQueues

A collection of objects that describe the requested game session queues.

Type: Array of [GameSessionQueue](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

DescribeGameSessions

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves a set of one or more game sessions in a specific fleet location. You can optionally filter the results by current game session status.

This operation can be used in the following ways:

- To retrieve all game sessions that are currently running on all locations in a fleet, provide a fleet or alias ID, with an optional status filter. This approach returns all game sessions in the fleet's home Region and all remote locations.
- To retrieve all game sessions that are currently running on a specific fleet location, provide a fleet or alias ID and a location name, with optional status filter. The location can be the fleet's home Region or any remote location.
- To retrieve a specific game session, provide the game session ID. This approach looks for the game session ID in all fleets that reside in the AWS Region defined in the request.

Use the pagination parameters to retrieve results as a set of sequential pages.

If successful, a `GameSession` object is returned for each game session that matches the request.

This operation is not designed to be continually called to track game session status. This practice can cause you to exceed your API limit, which results in errors. Instead, you must configure an Amazon Simple Notification Service (SNS) topic to receive notifications from FlexMatch or queues. Continuously polling with `DescribeGameSessions` should only be used for games in development with low game session usage.

Available in Amazon GameLift Servers Local.

Learn more

[Find a game session](#)

[All APIs by task](#)

Request Syntax

```
{  
  "AliasId": "string",
```

```
"FleetId": "string",  
"GameSessionId": "string",  
"Limit": number,  
"Location": "string",  
"NextToken": "string",  
"StatusFilter": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

AliasId

A unique identifier for the alias associated with the fleet to retrieve game sessions for. You can use either the alias ID or ARN value.

Type: String

Pattern: `^alias-\S+|^arn:.*:alias\/alias-\S+`

Required: No

FleetId

A unique identifier for the fleet to retrieve game sessions for. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\`

Required: No

GameSessionId

An identifier for the game session that is unique across all regions to retrieve. The value is always a full ARN in the following format: For Home Region game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>`. For Remote Location game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<location>/<ID string>`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `[a-zA-Z0-9:/-]+`

Required: No

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

Location

A fleet location to get game sessions for. You can specify a fleet's home Region or a remote location. Use the AWS Region code format, such as `us-west-2`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-]+`

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

StatusFilter

Game session status to filter results on. You can filter on the following states: ACTIVE, TERMINATED, ACTIVATING, and TERMINATING. The last two are transitory and used for only very brief periods of time.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "GameSessions": [
    {
      "ComputeName": "string",
      "CreationTime": number,
      "CreatorId": "string",
      "CurrentPlayerSessionCount": number,
      "DnsName": "string",
      "FleetArn": "string",
      "FleetId": "string",
      "GameProperties": [
        {
          "Key": "string",
          "Value": "string"
        }
      ],
      "GameSessionData": "string",
      "GameSessionId": "string",
      "IpAddress": "string",
      "Location": "string",
      "MatchmakerData": "string",
      "MaximumPlayerSessionCount": number,
      "Name": "string",
```

```
    "PlayerGatewayStatus": "string",
    "PlayerSessionCreationPolicy": "string",
    "Port": number,
    "Status": "string",
    "StatusReason": "string",
    "TerminationTime": 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.

GameSessions

A collection of properties for each game session that matches the request.

Type: Array of [GameSession](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TerminalRoutingStrategyException

The service is unable to resolve the routing for a particular alias because it has a terminal RoutingStrategy associated with it. The message returned in this exception is the message defined in the routing strategy itself. Such requests should only be retried if the routing strategy for the specified alias is modified.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

DescribeInstances

This API works with the following fleet types: EC2, Container

Retrieves information about the EC2 instances in an Amazon GameLift Servers managed fleet, including instance ID, connection data, and status. You can use this operation with a multi-location fleet to get location-specific instance information. As an alternative, use the operations https://docs.aws.amazon.com/gamelift/latest/apireference/API_ListCompute and https://docs.aws.amazon.com/gamelift/latest/apireference/API_DescribeCompute to retrieve information for compute resources, including EC2 and Anywhere fleets.

You can call this operation in the following ways:

- To get information on all instances in a fleet's home Region, specify the fleet ID.
- To get information on all instances in a fleet's remote location, specify the fleet ID and location name.
- To get information on a specific instance in a fleet, specify the fleet ID and instance ID.

Use the pagination parameters to retrieve results as a set of sequential pages.

If successful, this operation returns Instance objects for each requested instance, listed in no particular order. If you call this operation for an Anywhere fleet, you receive an `InvalidRequestException`.

Learn more

[Remotely connect to fleet instances](#)

[Debug fleet issues](#)

Related actions

[All APIs by task](#)

Request Syntax

```
{  
  "FleetId": "string",  
  "InstanceId": "string",  
  "Limit": number,  
}
```



```
"Location": "string",  
"NextToken": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet to retrieve instance information for. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\-]+|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\-]+`

Required: Yes

InstanceId

A unique identifier for an instance to retrieve. Specify an instance ID or leave blank to retrieve all instances in the fleet.

Type: String

Pattern: `[a-zA-Z0-9\.\-]+`

Required: No

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

Location

The name of a location to retrieve instance information for, in the form of an AWS Region code such as `us-west-2`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\ -]+`

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "Instances": [
    {
      "CreationTime": number,
      "DnsName": "string",
      "FleetArn": "string",
      "FleetId": "string",
      "InstanceId": "string",
      "IpAddress": "string",
      "Location": "string",
      "OperatingSystem": "string",
```

```
    "Status": "string",  
    "Type": "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.

Instances

A collection of objects containing properties for each instance returned.

Type: Array of [Instance](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Get instance data on a fleet's home Region

This example requests information on instances in a fleet's home Region. If the specified fleet also has instances deployed to remote locations, they aren't included in the response.

Sample Request

```
{
  "FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
  "Limit": "1"
}
```

Sample Response

```
{
  "Instances": [
    {
      "FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
      "FleetArn": "arn:aws:gamelift:us-west-2::fleet/fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
      "InstanceId": "i-11111111a222b333c,
```

```
        "IpAddress": "192.0.2.0",
        "DnsName": "ec2-192-0-2-0.us-west-2.compute.amazonaws.com",
        "OperatingSystem": "WIN_2012",
        "Type": "c4.large",
        "Status": "Active",
        "CreationTime": 1515032347.867,
        "Location": "us-west-2
    }
]
}
```

Get instance data on a fleet's remote location

This example requests information on fleet instances deployed to a remote location in the AWS Region sa-east-1. This example retrieves information on fleet instances in the sa-east-1 only.

Sample Request

```
{
  "FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
  "Location": "sa-east-1",
  "Limit": "2"
}
```

Sample Response

```
{
  "Instances": [
    {
      "FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
      "FleetArn": "arn:aws:gamelift:us-west-2::fleet/fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
      "InstanceId": "i-11111111a222b333c",
      "IpAddress": "192.0.2.0",
      "DnsName": "ec2-192-0-2-0.us-west-2.compute.amazonaws.com",
      "OperatingSystem": "WIN_2012",
      "Type": "c4.large",
      "Status": "Active",
      "CreationTime": 1515032347.867,
      "Location": "sa-east-1"
    },
    {
```

```
    "FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
    "FleetArn": "arn:aws:gamelift:us-east-2::fleet/
fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
    "InstanceId": "i-12312311a222b444d",
    "IpAddress": "192.0.2.0",
    "DnsName": "ec2-192-0-2-0.us-west-2.compute.amazonaws.com",
    "OperatingSystem": "WIN_2012",
    "Type": "c4.large",
    "Status": "Active",
    "CreationTime": 1515032999.867,
    "Location": "sa-east-1"
  }
]
```

See Also

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

- [AWS Command Line Interface 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](#)

DescribeMatchmaking

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves one or more matchmaking tickets. Use this operation to retrieve ticket information, including--after a successful match is made--connection information for the resulting new game session.

To request matchmaking tickets, provide a list of up to 10 ticket IDs. If the request is successful, a ticket object is returned for each requested ID that currently exists.

This operation is not designed to be continually called to track matchmaking ticket status. This practice can cause you to exceed your API limit, which results in errors. Instead, as a best practice, set up an Amazon Simple Notification Service to receive notifications, and provide the topic ARN in the matchmaking configuration.

Learn more

[Add FlexMatch to a game client](#)

[Set Up FlexMatch event notification](#)

Request Syntax

```
{  
  "TicketIds": [ "string" ]  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

TicketIds

A unique identifier for a matchmaking ticket. You can include up to 10 ID values.

Type: Array of strings

Length Constraints: Maximum length of 128.

Pattern: `[a-zA-Z0-9-\.\.]*`

Required: Yes

Response Syntax

```
{
  "TicketList": [
    {
      "ConfigurationArn": "string",
      "ConfigurationName": "string",
      "EndTime": number,
      "EstimatedWaitTime": number,
      "GameSessionConnectionInfo": {
        "DnsName": "string",
        "GameSessionArn": "string",
        "IpAddress": "string",
        "MatchedPlayerSessions": [
          {
            "PlayerId": "string",
            "PlayerSessionId": "string"
          }
        ],
        "PlayerGatewayStatus": "string",
        "Port": number
      },
      "Players": [
        {
          "LatencyInMs": {
            "string" : number
          },
          "PlayerAttributes": {
            "string" : {
              "N": number,
```



```
        "S": "string",
        "SDM": {
            "string" : number
        },
        "SL": [ "string" ]
    },
    "PlayerId": "string",
    "Team": "string"
}
],
"StartTime": number,
"Status": "string",
"StatusMessage": "string",
"StatusReason": "string",
"TicketId": "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.

TicketList

A collection of existing matchmaking ticket objects matching the request.

Type: Array of [MatchmakingTicket](#) objects

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

DescribeMatchmakingConfigurations

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves the details of FlexMatch matchmaking configurations.

This operation offers the following options: (1) retrieve all matchmaking configurations, (2) retrieve configurations for a specified list, or (3) retrieve all configurations that use a specified rule set name. When requesting multiple items, use the pagination parameters to retrieve results as a set of sequential pages.

If successful, a configuration is returned for each requested name. When specifying a list of names, only configurations that currently exist are returned.

Learn more

[Setting up FlexMatch matchmakers](#)

Request Syntax

```
{  
  "Limit": number,  
  "Names": [ "string" ],  
  "NextToken": "string",  
  "RuleSetName": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Limit

The maximum number of results to return. Use this parameter with NextToken to get results as a set of sequential pages. This parameter is limited to 10.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

Names

A unique identifier for the matchmaking configuration(s) to retrieve. You can use either the configuration name or ARN value. To request all existing configurations, leave this parameter empty.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]*|^arn:.*:matchmakingconfiguration\[a-zA-Z0-9-\.]*`

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

RuleSetName

A unique identifier for the matchmaking rule set. You can use either the rule set name or ARN value. Use this parameter to retrieve all matchmaking configurations that use this rule set.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]*|^arn:.*:matchmakingruleset\[a-zA-Z0-9-\.]*`

Required: No

Response Syntax

```
{
  "Configurations": [
    {
      "AcceptanceRequired": boolean,
      "AcceptanceTimeoutSeconds": number,
      "AdditionalPlayerCount": number,
      "BackfillMode": "string",
      "ConfigurationArn": "string",
      "CreationTime": number,
      "CustomEventData": "string",
      "Description": "string",
      "FlexMatchMode": "string",
      "GameProperties": [
        {
          "Key": "string",
          "Value": "string"
        }
      ],
      "GameSessionData": "string",
      "GameSessionQueueArns": [ "string" ],
      "Name": "string",
      "NotificationTarget": "string",
      "RequestTimeoutSeconds": number,
      "RuleSetArn": "string",
      "RuleSetName": "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.

Configurations

A collection of requested matchmaking configurations.

Type: Array of [MatchmakingConfiguration](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

DescribeMatchmakingRuleSets

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves the details for FlexMatch matchmaking rule sets. You can request all existing rule sets for the Region, or provide a list of one or more rule set names. When requesting multiple items, use the pagination parameters to retrieve results as a set of sequential pages. If successful, a rule set is returned for each requested name.

Learn more

- [Build a rule set](#)

Request Syntax

```
{
  "Limit": number,
  "Names": [ "string" ],
  "NextToken": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Limit

The maximum number of results to return. Use this parameter with NextToken to get results as a set of sequential pages.

Type: Integer

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

Required: No

Names

A list of one or more matchmaking rule set names to retrieve details for. (Note: The rule set name is different from the optional "name" field in the rule set body.) You can use either the rule set name or ARN value.

Type: Array of strings

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

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]*|^arn:.*:matchmakingruleset\/[a-zA-Z0-9-\.]*`

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "NextToken": "string",
  "RuleSets": [
    {
      "CreationTime": number,
      "RuleSetArn": "string",
      "RuleSetBody": "string",
      "RuleSetName": "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.

RuleSets

A collection of requested matchmaking rule set objects.

Type: Array of [MatchmakingRuleSet](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

DescribePlayerSessions

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves properties for one or more player sessions.

This action can be used in the following ways:

- To retrieve a specific player session, provide the player session ID only.
- To retrieve all player sessions in a game session, provide the game session ID only.
- To retrieve all player sessions for a specific player, provide a player ID only.

To request player sessions, specify either a player session ID, game session ID, or player ID. You can filter this request by player session status. If you provide a specific `PlayerSessionId` or `PlayerId`, Amazon GameLift Servers ignores the filter criteria. Use the pagination parameters to retrieve results as a set of sequential pages.

If successful, a `PlayerSession` object is returned for each session that matches the request.

Related actions

[All APIs by task](#)

Request Syntax

```
{
  "GameSessionId": "string",
  "Limit": number,
  "NextToken": "string",
  "PlayerId": "string",
  "PlayerSessionId": "string",
  "PlayerSessionStatusFilter": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

GameSessionId

An identifier for the game session that is unique across all regions to retrieve player sessions for. The value is always a full ARN in the following format: For Home Region game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>`. For Remote Location game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<location>/<ID string>`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `[a-zA-Z0-9:/-]+`

Required: No

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages. If a player session ID is specified, this parameter is ignored.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value. If a player session ID is specified, this parameter is ignored.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

PlayerId

A unique identifier for a player to retrieve player sessions for.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

PlayerSessionId

A unique identifier for a player session to retrieve.

Type: String

Pattern: ^psess-\S+

Required: No

PlayerSessionStatusFilter

Player session status to filter results on. Note that when a PlayerSessionId or PlayerId is provided in a DescribePlayerSessions request, then the PlayerSessionStatusFilter has no effect on the response.

Possible player session statuses include the following:

- **RESERVED** -- The player session request has been received, but the player has not yet connected to the server process and/or been validated.
- **ACTIVE** -- The player has been validated by the server process and is currently connected.
- **COMPLETED** -- The player connection has been dropped.
- **TIMEDOUT** -- A player session request was received, but the player did not connect and/or was not validated within the timeout limit (60 seconds).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
```

```
"NextToken": "string",
"PlayerSessions": [
  {
    "CreationTime": number,
    "DnsName": "string",
    "FleetArn": "string",
    "FleetId": "string",
    "GameSessionId": "string",
    "IpAddress": "string",
    "PlayerData": "string",
    "PlayerId": "string",
    "PlayerSessionId": "string",
    "Port": number,
    "Status": "string",
    "TerminationTime": number
  }
]
```

Response Elements

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

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

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

PlayerSessions

A collection of objects containing properties for each player session that matches the request.

Type: Array of [PlayerSession](#) objects

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

DescribeRuntimeConfiguration

This API works with the following fleet types: EC2

Retrieves a fleet's runtime configuration settings. The runtime configuration determines which server processes run, and how, on computes in the fleet. For managed EC2 fleets, the runtime configuration describes server processes that run on each fleet instance. You can update a fleet's runtime configuration at any time using [UpdateRuntimeConfiguration](#).

To get the current runtime configuration for a fleet, provide the fleet ID.

If successful, a RuntimeConfiguration object is returned for the requested fleet. If the requested fleet has been deleted, the result set is empty.

Learn more

[Setting up Amazon GameLift Servers fleets](#)

[Running multiple processes on a fleet](#)

Request Syntax

```
{  
  "FleetId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[FleetId](#)

A unique identifier for the fleet to get the runtime configuration for. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\ -]+$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\ -]+$`

Required: Yes

Response Syntax

```
{
  "RuntimeConfiguration": {
    "GameSessionActivationTimeoutSeconds": number,
    "MaxConcurrentGameSessionActivations": number,
    "ServerProcesses": [
      {
        "ConcurrentExecutions": number,
        "LaunchPath": "string",
        "Parameters": "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.

RuntimeConfiguration

Instructions that describe how server processes are launched and maintained on computes in the fleet.

Type: [RuntimeConfiguration](#) object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Request the runtime configuration for a fleet

This example retrieves the current runtime configuration for a specified fleet. As shown, the requested fleet is configured to run four concurrent processes of the game server executable, one with debug mode turned on. The property `MaxConcurrentGameSessionActivations` is set to the default value, which places no limit on concurrent activations.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
```

```
}
```

Sample Response

```
{
  "RuntimeConfiguration": {
    "ServerProcesses": [
      {
        "LaunchPath": "C:\\game\\Bin64.Release.Dedicated\\MegaFrogRace_Server.exe",
        "Parameters": "+gamelift_start_server",
        "ConcurrentExecutions": 3
      },
      {
        "LaunchPath": "C:\\game\\Bin64.Release.Dedicated\\MegaFrogRace_Server.exe",
        "Parameters": "+gamelift_start_server +debug",
        "ConcurrentExecutions": 1
      }
    ],
    "MaxConcurrentGameSessionActivations": 2147483647,
    "GameSessionActivationTimeoutSeconds": 300
  }
}
```

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](#)

DescribeScalingPolicies

This API works with the following fleet types: EC2, Container

Retrieves all scaling policies applied to a fleet.

To get a fleet's scaling policies, specify the fleet ID. You can filter this request by policy status, such as to retrieve only active scaling policies. Use the pagination parameters to retrieve results as a set of sequential pages. If successful, set of `ScalingPolicy` objects is returned for the fleet.

A fleet may have all of its scaling policies suspended. This operation does not affect the status of the scaling policies, which remains ACTIVE.

Request Syntax

```
{
  "FleetId": "string",
  "Limit": number,
  "Location": "string",
  "NextToken": "string",
  "StatusFilter": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet for which to retrieve scaling policies. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\-]+|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\-]+`

Required: Yes

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

Location

The fleet location. If you don't specify this value, the response contains the scaling policies of every location in the fleet.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-\-]+`

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

StatusFilter

Scaling policy status to filter results on. A scaling policy is only in force when in an `ACTIVE` status.

- **ACTIVE** -- The scaling policy is currently in force.
- **UPDATEREQUESTED** -- A request to update the scaling policy has been received.
- **UPDATING** -- A change is being made to the scaling policy.
- **DELETEREQUESTED** -- A request to delete the scaling policy has been received.
- **DELETING** -- The scaling policy is being deleted.
- **DELETED** -- The scaling policy has been deleted.
- **ERROR** -- An error occurred in creating the policy. It should be removed and recreated.

Type: String

Valid Values: ACTIVE | UPDATE_REQUESTED | UPDATING | DELETE_REQUESTED | DELETING | DELETED | ERROR

Required: No

Response Syntax

```
{
  "NextToken": "string",
  "ScalingPolicies": [
    {
      "ComparisonOperator": "string",
      "EvaluationPeriods": number,
      "FleetArn": "string",
      "FleetId": "string",
      "Location": "string",
      "MetricName": "string",
      "Name": "string",
      "PolicyType": "string",
      "ScalingAdjustment": number,
      "ScalingAdjustmentType": "string",
      "Status": "string",
      "TargetConfiguration": {
        "TargetValue": number
      },
      "Threshold": number,
      "UpdateStatus": "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](#)

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

[ScalingPolicies](#)

A collection of objects containing the scaling policies matching the request.

Type: Array of [ScalingPolicy](#) objects

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

View active scaling policies

In this example, we want to get a list of active scaling policies for a specific fleet. We can do this by retrieving the list of all policies for the fleet filtered by status. The sample request illustrates the use of pagination parameters, `Limit` and `NextToken`, to retrieve multiple results in sequential sets.

Sample Request

```
POST / HTTP/1.1
Host: gamelift.us-west-2.amazonaws.com;
Accept-Encoding: identity
Content-Length: 336
User-Agent: aws-cli/1.11.36 Python/2.7.9 Windows/7 botocore/1.4.93
Content-Type: application/x-amz-json-1.0
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20170406/us-west-2/
gamelift/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target,
Signature=wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
X-Amz-Date: 20170406T004805Z
X-Amz-Target: GameLift.DescribeScalingPolicies

{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "Limit": "2",
  "StatusFilter": "ACTIVE"
}
```

CLI syntax:

```
$aws gamelift describe-scaling-policies
--fleet-id "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
--limit "2"
--status-filter "ACTIVE"
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: b34f8665-EXAMPLE
Content-Type: application/x-amz-json-1.1
Content-Length: 600
Date: Thu, 06 Apr 2017 00:48:07 GMT

{
  "ScalingPolicies": [
    {
      "Status": "ACTIVE",
      "Name": "My_Target_Policy_1",
      "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
      "PolicyType": "TargetBased",
      "MetricName": "PercentAvailableGameSessions",
      "TargetConfiguration": {"TargetValue": 15}
    },
    {
      "Status": "ACTIVE",
      "EvaluationPeriods": "1",
      "Name": "My_Rule_Policy_1",
      "ComparisonOperator": "GreaterThanThreshold",
      "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
      "PolicyType": "RuleBased",
      "Threshold": "1.0",
      "ScalingAdjustment": "10",
      "MetricName": "QueueDepth",
      "ScalingAdjustmentType": "PercentChangeInCapacity"
    }
  ],
  "NextToken":
    "eyJhd3NBZjY2NvdW50SWQ0nsicyI6IjMwMjc3NjAxNjM5OCJ9LCJidWlsZElkIjpw7InMiOiJidWlsZC01NWYxZTZmMS1jY"
}
```

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](#)

DescribeScript

This API works with the following fleet types: EC2

Retrieves properties for a Realtime script.

To request a script record, specify the script ID. If successful, an object containing the script properties is returned.

Learn more

[Amazon GameLift Servers Amazon GameLift Servers Realtime](#)

Related actions

[All APIs by task](#)

Request Syntax

```
{  
  "ScriptId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

ScriptId

A unique identifier for the Realtime script to retrieve properties for. You can use either the script ID or ARN value.

Type: String

Pattern: `^script-\S+|^arn:.*:script\/script-\S+`

Required: Yes

Response Syntax

```
{
  "Script": {
    "CreationTime": number,
    "Name": "string",
    "NodeJsVersion": "string",
    "ScriptArn": "string",
    "ScriptId": "string",
    "SizeOnDisk": number,
    "StorageLocation": {
      "Bucket": "string",
      "Key": "string",
      "ObjectVersion": "string",
      "RoleArn": "string"
    },
    "Version": "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.

Script

A set of properties describing the requested script.

Type: [Script](#) object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

View a script record

This example illustrates one usage of DescribeScript.

Sample Request

```
{
    "ScriptId": "script-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
}
```

CLI syntax:

```
aws gamelift describe-script --script-id "script-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
```

Sample Response

```
{
```



```
"Script": {
  "CreationTime": 1496708916.18,
  "Name": "My_Realtime_Script_2",
  "NodeJsVersion": "24.x",
  "ScriptArn": "arn:aws:gamelift:us-west-2::script/
script-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "ScriptId": "script-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "SizeOnDisk": 0,
  "StorageLocation": {
    "Bucket": "my_realtime_script_files",
    "Key": "myRealtimeScript.zip"
    "RoleArn": "arn:aws:iam::111122223333:role/GameLiftAccess"
    "ObjectVersion": null
  },
  "Version": "12345.678"
}
```

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](#)

DescribeVpcPeeringAuthorizations

This API works with the following fleet types: EC2

Retrieves valid VPC peering authorizations that are pending for the AWS account. This operation returns all VPC peering authorizations and requests for peering. This includes those initiated and received by this account.

Related actions

[All APIs by task](#)

Response Syntax

```
{
  "VpcPeeringAuthorizations": [
    {
      "CreationTime": number,
      "ExpirationTime": number,
      "GameLiftAwsAccountId": "string",
      "PeerVpcAwsAccountId": "string",
      "PeerVpcId": "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.

[VpcPeeringAuthorizations](#)

A collection of objects that describe all valid VPC peering operations for the current AWS account.

Type: Array of [VpcPeeringAuthorization](#) objects

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

DescribeVpcPeeringConnections

This API works with the following fleet types: EC2

Retrieves information on VPC peering connections. Use this operation to get peering information for all fleets or for one specific fleet ID.

To retrieve connection information, call this operation from the AWS account that is used to manage the Amazon GameLift Servers fleets. Specify a fleet ID or leave the parameter empty to retrieve all connection records. If successful, the retrieved information includes both active and pending connections. Active connections identify the IPv4 CIDR block that the VPC uses to connect.

Related actions

[All APIs by task](#)

Request Syntax

```
{  
  "FleetId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[FleetId](#)

A unique identifier for the fleet. You can use either the fleet ID or ARN value.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\-]+`

Required: No

Response Syntax

```
{
  "VpcPeeringConnections": [
    {
      "FleetArn": "string",
      "FleetId": "string",
      "GameLiftVpcId": "string",
      "IPv4CidrBlock": "string",
      "PeerVpcId": "string",
      "Status": {
        "Code": "string",
        "Message": "string"
      },
      "VpcPeeringConnectionId": "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.

VpcPeeringConnections

A collection of VPC peering connection records that match the request.

Type: Array of [VpcPeeringConnection](#) objects

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

GetComputeAccess

This API works with the following fleet types: EC2, Container

Requests authorization to remotely connect to a hosting resource in a Amazon GameLift Servers managed fleet. This operation is not used with Amazon GameLift Servers Anywhere fleets.

Request options

Provide the fleet ID and compute name. The compute name varies depending on the type of fleet.

- For a compute in a managed EC2 fleet, provide an instance ID. Each instance in the fleet is a compute.
- For a compute in a managed container fleet, provide a compute name. In a container fleet, each game server container group on a fleet instance is assigned a compute name.

Results

If successful, this operation returns a set of temporary AWS credentials, including a two-part access key and a session token.

- With a managed EC2 fleet (where compute type is EC2), use these credentials with Amazon EC2 Systems Manager (SSM) to start a session with the compute. For more details, see [Starting a session \(AWS CLI\)](#) in the *Amazon EC2 Systems Manager User Guide*.
- With a managed container fleet, use these credentials and the target value with SSM to connect to the fleet instance with the container group. When connected to the instance, use Docker commands to interact with the container using the container name and runtime ID.

Request Syntax

```
{  
  "ComputeName": "string",  
  "FleetId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

 Note

In the following list, the required parameters are described first.

ComputeName

A unique identifier for the compute resource that you want to connect to. For an EC2 fleet, use an instance ID. For a managed container fleet, use a compute name. You can retrieve a fleet's compute names by calling [ListCompute](#).

Type: String

Length Constraints: Maximum length of 1024.

Pattern: `^[a-zA-Z0-9\-\_]+(\/[a-zA-Z0-9\-\_]+)?$|^arn:.*:compute\[a-zA-Z0-9\-\_]+(\/[a-zA-Z0-9\-\_]+)?`

Required: Yes

FleetId

A unique identifier for the fleet that holds the compute resource that you want to connect to. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+$|^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\_]+$`

Required: Yes

Response Syntax

```
{  
  "ComputeArn": "string",  
  "ComputeName": "string",  
}
```

```
"ContainerIdentifiers": [  
  {  
    "ContainerName": "string",  
    "ContainerRuntimeId": "string"  
  }  
],  
"Credentials": {  
  "AccessKeyId": "string",  
  "SecretAccessKey": "string",  
  "SessionToken": "string"  
},  
"FleetArn": "string",  
"FleetId": "string",  
"Target": "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.

ComputeArn

The Amazon Resource Name ([ARN](#)) that is assigned to an Amazon GameLift Servers compute resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::compute/compute-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`.

Type: String

Length Constraints: Maximum length of 1024.

Pattern: `^arn:.*:compute\/[a-zA-Z0-9\-\_]+(\/[a-zA-Z0-9\-\_]+)?`

ComputeName

The identifier of the compute resource to be accessed. This value might be either a compute name or an instance ID.

Type: String

Length Constraints: Maximum length of 1024.

Pattern: `^[a-zA-Z0-9\-\-]+(\/[a-zA-Z0-9\-\-]+)?$|^arn:.*:compute\[a-zA-Z0-9\-\-]+(\/[a-zA-Z0-9\-\-]+)?`

ContainerIdentifiers

For a managed container fleet, a list of containers on the compute. Use the container runtime ID with Docker commands to connect to a specific container.

Type: Array of [ContainerIdentifier](#) objects

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

Credentials

A set of temporary AWS credentials for use when connecting to the compute resource with Amazon EC2 Systems Manager (SSM).

Type: [AwsCredentials](#) object

FleetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\-]+`

FleetId

The ID of the fleet that holds the compute resource to be accessed.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\-]+$|^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\-]+`

Target

The instance ID where the compute resource is running.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 400.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

GetComputeAuthToken

This API works with the following fleet types: EC2, Anywhere, Container

Requests an authentication token from Amazon GameLift Servers for a compute resource in an Amazon GameLift Servers fleet. Game servers that are running on the compute use this token to communicate with the Amazon GameLift Servers service, such as when calling the Amazon GameLift Servers server SDK action `InitSDK()`. Authentication tokens are valid for a limited time span, so you need to request a fresh token before the current token expires.

Request options

- For managed EC2 fleets (compute type EC2), auth token retrieval and refresh is handled automatically. All game servers that are running on all fleet instances have access to a valid auth token.
- For Anywhere fleets (compute type ANYWHERE), if you're using the Amazon GameLift Servers Agent, auth token retrieval and refresh is handled automatically for any compute where the Agent is running. If you're not using the Agent, create a mechanism to retrieve and refresh auth tokens for computes that are running game server processes.
- For managed container fleets, if you're using the Amazon GameLift Servers Agent, auth token retrieval and refresh is handled automatically for any container or Anywhere compute where the Agent is running. If you're not using the Agent, use the alternative authentication method instead of using this operation.

Learn more

- [Create an Anywhere fleet](#)
- [Test your integration](#)
- [Server SDK reference guides](#) (for version 5.x)

Request Syntax

```
{  
  "ComputeName": "string",  
  "FleetId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

ComputeName

The name of the compute resource you are requesting the authentication token for. For an Anywhere fleet compute, use the registered compute name. For an EC2 fleet instance, use the instance ID.

Type: String

Length Constraints: Maximum length of 1024.

Pattern: `^[a-zA-Z0-9\-\_]+(\/[a-zA-Z0-9\-\_]+)?$|^arn:.*:compute\[a-zA-Z0-9\-\_]+(\/[a-zA-Z0-9\-\_]+)?`

Required: Yes

FleetId

A unique identifier for the fleet that the compute is registered to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+$|^arn:.*:[a-z]*fleet\[a-zA-Z0-9\-\_]+$`

Required: Yes

Response Syntax

```
{
```

```
"AuthToken": "string",  
"ComputeArn": "string",  
"ComputeName": "string",  
"ExpirationTimestamp": number,  
"FleetArn": "string",  
"FleetId": "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.

AuthToken

A valid temporary authentication token.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[a-zA-Z0-9\-]+`

ComputeArn

The Amazon Resource Name ([ARN](#)) that is assigned to an Amazon GameLift Servers compute resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::compute/compute-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`.

Type: String

Length Constraints: Maximum length of 1024.

Pattern: `^arn:.*:compute[/[a-zA-Z0-9\-]+]([/[a-zA-Z0-9\-]+])?`

ComputeName

The name of the compute resource that the authentication token is issued to.

Type: String

Length Constraints: Maximum length of 1024.

Pattern: `^[a-zA-Z0-9\-\_]+(\/[a-zA-Z0-9\-\_]+)?$|^arn:.*:compute\[a-zA-Z0-9\-\_]+(\/[a-zA-Z0-9\-\_]+)?`

[ExpirationTimestamp](#)

The amount of time until the authentication token is no longer valid.

Type: Timestamp

[FleetArn](#)

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\_]+$`

[FleetId](#)

A unique identifier for the fleet that the compute is registered to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+$|^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\_]+$`

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

GetGameSessionLogUrl

This API works with the following fleet types: EC2

Retrieves the location of stored game session logs for a specified game session on Amazon GameLift Servers managed fleets. When a game session is terminated, Amazon GameLift Servers automatically stores the logs in Amazon S3 and retains them for 14 days. Use this URL to download the logs.

Note

See the [AWS Service Limits](#) page for maximum log file sizes. Log files that exceed this limit are not saved.

[All APIs by task](#)

Request Syntax

```
{  
  "GameSessionId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[GameSessionId](#)

An identifier for the game session that is unique across all regions to get logs for. The value is always a full ARN in the following format: For Home Region game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>`. For

Remote Location game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<location>/<ID string>`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `[a-zA-Z0-9:/-]+`

Required: Yes

Response Syntax

```
{
  "PreSignedUrl": "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.

[PreSignedUrl](#)

Location of the requested game session logs, available for download. This URL is valid for 15 minutes, after which S3 will reject any download request using this URL. You can request a new URL any time within the 14-day period that the logs are retained.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

GetInstanceAccess

This API works with the following fleet types: EC2

Requests authorization to remotely connect to an instance in an Amazon GameLift Servers managed fleet. Use this operation to connect to instances with game servers that use Amazon GameLift Servers server SDK 4.x or earlier. To connect to instances with game servers that use server SDK 5.x or later, call https://docs.aws.amazon.com/gamelift/latest/apireference/API_GetComputeAccess.

To request access to an instance, specify IDs for the instance and the fleet it belongs to. You can retrieve instance IDs for a fleet by calling [DescribeInstances](#) with the fleet ID.

If successful, this operation returns an IP address and credentials. The returned credentials match the operating system of the instance, as follows:

- For a Windows instance: returns a user name and secret (password) for use with a Windows Remote Desktop client.
- For a Linux instance: returns a user name and secret (RSA private key) for use with an SSH client. You must save the secret to a `.pem` file. If you're using the AWS CLI, see the example [Get credentials for a Linux instance](#) for tips on automatically saving the secret to a `.pem` file.

Learn more

[Remotely connect to fleet instances](#)

[Debug fleet issues](#)

Related actions

[All APIs by task](#)

Request Syntax

```
{  
  "FleetId": "string",  
  "InstanceId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet that contains the instance you want to access. You can request access to instances in EC2 fleets with the following statuses: ACTIVATING, ACTIVE, or ERROR. Use either a fleet ID or an ARN value.

Note

You can access fleets in ERROR status for a short period of time before Amazon GameLift Servers deletes them.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-.]+|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-.]+`

Required: Yes

InstanceId

A unique identifier for the instance you want to access. You can access an instance in any status.

Type: String

Pattern: `[a-zA-Z0-9\-.]+`

Required: Yes

Response Syntax

```
{
  "InstanceAccess": {
    "Credentials": {
      "Secret": "string",
      "UserName": "string"
    },
    "FleetId": "string",
    "InstanceId": "string",
    "IpAddress": "string",
    "OperatingSystem": "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.

InstanceAccess

The connection information for a fleet instance, including IP address and access credentials.

Type: [InstanceAccess](#) object

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Get credentials for a Linux instance

This example requests a set of credentials to remotely connect to a fleet instance deployed with a Linux Amazon machine image (AMI). The instance is running a game server build with server SDK version 4.02.

When calling `GetInstanceAccess` programmatically, save the returned value of `Secret` (an RSA private key) as a `.pem` file in the proper format. The returned value uses a newline (`\n`) to indicate a line break.

If you're calling `GetInstanceAccess` with the AWS CLI, you can automatically save the RSA private key as a `.pem` file. The CLI syntax in this example saves the secret to a file called `MyPrivateKey.pem`. After saving the private key, you can update the file permissions with the following command:

```
$ chmod 400 MyPrivateKey.pem
```

Sample Request

```
{"FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",  
  "InstanceId": "i-11111111a222b333c"  
}
```

CLI syntax:

```
aws gamelift get-instance-access \
```

```
--fleet-id "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa" \
--instance-id "i-11111111a222b333c" \
--query 'InstanceAccess.Credentials.Secret' \
--output text > MyPrivateKey.pem
```

Sample Response

```
{"InstanceAccess": {
  "Credentials": {
    "Secret": "-----BEGIN RSA PRIVATE KEY-----
nEXAMPLEKEYKCAQEAY7WZhaDsrA1W3mRlQtvhwyORRX8gnxgDAfRt/gx42kWXsT4rXE/b5CpSgie/
\nvBoU7jLxx92pNHOfnByP+Dc21eyyz6CvjTmWA0JwfWiW5/akH7i05dSrvC7dQkW2duV5QuUdEQW
\nZ/aNxMniGQE6XAgfwlnXVBwrerrQo+ZWQeqiUwwMkuEbLeJFLhMCvYURpUMSC1oehm449ilx9X1F
\nG50TCFe0zfl8dqqCP6GzbPaIjiU19xX/az0R9V+tpU0zEL+wmXnZt3/nHPQ5xvD20JH67km6SuPW
\noPzev/D8V+x4+bHthfSjR9Y7DvQFjFBVwHXigBdtZcU2/wei8D/HYwIDAQABaoIBAGZ1kaEvnrrqu
\n/uler7vgIn5m7lN5LKw4hJLAIW6tUT/fzvtcHK0SkbQCQXuriHmQ2MQyJX/0kn2NfjLV/
ufGxbL1\nmb5qwMGUnEpJaZD6QSSs3kICLwWUYUiGfc0uiSbmJoap/
GTLU0W5Mfcv36PaBUNy5p53V6G7hXb2\nbahyWyJNfjLe4M86yd2YK3V2CmK+X/
B0sShnJ36+hjrXPPWmV3N9zEmCdJjA+K15DYmhm/
tJWSD9\n81oGk9TopEp7CkIfatEATyyZiVqoRq6k64iuM9JkA30zdXzMqexXVJ1TLZVEH0E7bh1Y9d801ozR
\noQs/FiZNAx2iijCWyv0lpjE73+kCgYEA9mZtyhkHkFDpwrSM1APaL8oNAbbjwEy7Z5Mqfq1
+1Ip1\nYkriL0DbLXlvrAH+yHPRit2hH0jtUNZh4Axv+cpq09qbUI3+43eEy24B7G/Uh+GTfbjsXs0xQx/
x\np9otyVwc7hsQ5TA5PZb+mvkJ50BEKzet9XcKw0NBVELGhnEpe7cCgYEA06Vgov6YHleHui9kHuws
\nayav0elc5zkxjF9nfHFJRry21R1trw2Vdpn+9g481URrpzWV0Eihvm+xTtmaZ1Sp//1kq75XDwnU
\nWA8gkn603QE3fq2yN98BURsAKdJfJ5RL1HvGQvTe10HLYYXpJnEkHv+Unl2ajLivWUt5pbBrKbUC
\nngYBjb0+0Zk0sCcpZ29sbzjYjpIddErySIyRX5gV2uNQwAjlDp9PfN295yQ+BxMBXiIycWVQiw0bH
\nnoMo7yykABY70zd5wQewBQ4AdSlWSX4nGDtsiFxiWiI5sKuAAe0CbTosy1s8w8fxoJ5Tz1sd0xNeGs
\nArq6Wv/G16zQuAE9zK9vwwKBgF+09VI/1wJBirsDGz9whVwFFPrTkJNvJZzYt69qezx1sjgFKshy
\nWBhd4xHZtmCqpBP1AymEjr/T0lbxyARmXMnIOWIANXMGb4KGSyl1mzSVAoQ+fqR+cJ3d0dyP11j
\nnjjb0Ed/NY8frlNDxAVHE8BSkdsx2f6ELEyBKJSRr9snRAoGAMrTwYneXzvTskF/S5Fyu0i0egLda
\nNWUH38v/nDCgEpIXD5Hn3qAEcju1IjmbwlvTW+nY2jVhv7UGd8MjwUTNGItDb6nsYqM2asrnF3qS
\nVRkAKKKYeGjkpUfVTrW0YFjXkfcR/V+QFL50ndHAKJXjW7a4ejJLncTzmZSpYzwApc=\n-----END RSA
PRIVATE KEY-----",
    "UserName": "gl-user-remote"
  },
  "FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
  "InstanceId": "i-11111111a222b333c",
  "IpAddress": "192.0.2.0",
  "OperatingSystem": "AMAZON_LINUX_2"
}
```

Get credentials for a Windows instance

This example requests a set of credentials to remotely connect to a fleet instance deployed with a Windows Amazon machine image (AMI). The instance is running a game server build with server SDK version 4.02.

Sample Request

```
{"FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",  
  "InstanceId": "i-11111111a222b333c"  
}
```

CLI syntax:

```
aws gamelift get-instance-access \  
  --fleet-id "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa" \  
  --instance-id "i-11111111a222b333c"
```

Sample Response

```
{"InstanceAccess": {  
  "Credentials": {  
    "Secret": "aA1bBB2cCCd3EEE",  
    "UserName": "gl-user-remote"  
  },  
  "FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",  
  "InstanceId": "i-11111111a222b333c",  
  "IpAddress": "192.0.2.0",  
  "OperatingSystem": "WIN_2012"  
}
```

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](#)

GetPlayerConnectionDetails

This API works with the following fleet types: EC2 (server SDK 5.x or later), Container

Retrieves connection details for game clients to connect to game sessions.

Player gateway benefits: DDoS protection with negligible impact to latency.

To enable player gateway on your fleet, set `PlayerGatewayMode` to `ENABLED` or `REQUIRED` when calling [CreateFleet](#) or [CreateContainerFleet](#).

How to use: After creating a game session and adding players, call this operation with the game session ID and player IDs. When player gateway is enabled, the response includes connection endpoints and player gateway tokens that your game clients can use to connect to the game session through player gateway. To learn more about player gateway integration, see [DDoS protection with Amazon GameLift Servers player gateway](#).

When player gateway is disabled or in locations where player gateway is not supported, this operation returns game server connection information without player gateway tokens, so that your game clients directly connect to the game server endpoint.

Request Syntax

```
{
  "GameSessionId": "string",
  "PlayerIds": [ "string" ]
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

GameSessionId

An identifier for the game session that is unique across all regions for which to retrieve player connection details. The value is always a full ARN in the following format: For Home Region game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>`. For Remote Location game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<location>/<ID string>`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `[a-zA-Z0-9:/-]+`

Required: Yes

PlayerIds

List of unique identifiers for players. Connection details are returned for each player in this list.

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 1024.

Required: Yes

Response Syntax

```
{
  "GameSessionId": "string",
  "PlayerConnectionDetails": [
    {
      "Endpoints": [
        {
          "IpAddress": "string",
          "Port": number
        }
      ],
      "Expiration": number,
      "PlayerGatewayToken": "string",
```

```
    "PlayerId": "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.

GameSessionId

An identifier for the game session that is unique across all regions for which the player connection details were retrieved. The value is always a full ARN in the following format: For Home Region game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>`. For Remote Location game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<location>/<ID string>`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `[a-zA-Z0-9:/- ]+`

PlayerConnectionDetails

A collection of player connection detail objects, one for each requested player.

Type: Array of [PlayerConnectionDetail](#) objects

Array Members: Minimum number of 1 item.

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidGameSessionStatusException

The requested operation would cause a conflict with the current state of a resource associated with the request and/or the game instance. Resolve the conflict before retrying.

HTTP Status Code: 400

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Get player connection details with player gateway enabled

This example retrieves connection details for a game client to connect to a game session through player gateway.

Sample Request

```
{
  "GameSessionId": "arn:aws:gamelift:ap-northeast-2::gamesession/
fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa/gsess-4444dddd-55ee-66ff-77aa-8888bbbb99cc",
  "PlayerIds": ["player1"]
}
```

Sample Response

```
{
  "GameSessionId": "arn:aws:gamelift:ap-northeast-2::gamesession/
fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa/gsess-4444dddd-55ee-66ff-77aa-8888bbbb99cc",
  "PlayerConnectionDetails": [
    {
      "PlayerId": "player1",
      "Endpoints": [
        {
          "IpAddress": "192.0.2.0",
          "Port": 7777
        },
        {
          "IpAddress": "198.51.100.0",
          "Port": 7777
        },
        {
          "IpAddress": "203.0.113.0",
          "Port": 7777
        }
      ],
      "PlayerGatewayToken": "wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY",
      "Expiration": "1736365885.22"
    }
  ]
}
```

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](#)

ListAliases

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves all aliases for this AWS account. You can filter the result set by alias name and/or routing strategy type. Use the pagination parameters to retrieve results in sequential pages.

Note

Returned aliases are not listed in any particular order.

Related actions

[All APIs by task](#)

Request Syntax

```
{
  "Limit": number,
  "Name": "string",
  "NextToken": "string",
  "RoutingStrategyType": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Limit

The maximum number of results to return. Use this parameter with NextToken to get results as a set of sequential pages.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

Name

A descriptive label that is associated with an alias. Alias names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1.

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1.

Required: No

RoutingStrategyType

The routing type to filter results on. Use this parameter to retrieve only aliases with a certain routing type. To retrieve all aliases, leave this parameter empty.

Possible routing types include the following:

- **SIMPLE** -- The alias resolves to one specific fleet. Use this type when routing to active fleets.
- **TERMINAL** -- The alias does not resolve to a fleet but instead can be used to display a message to the user. A terminal alias throws a `TerminalRoutingStrategyException` with the [RoutingStrategy](#) message embedded.

Type: String

Valid Values: SIMPLE | TERMINAL

Required: No

Response Syntax

```
{
  "Aliases": [
    {
      "AliasArn": "string",
      "AliasId": "string",
      "CreationTime": number,
      "Description": "string",
      "LastUpdatedTime": number,
      "Name": "string",
      "RoutingStrategy": {
        "FleetId": "string",
        "Message": "string",
        "Type": "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.

Aliases

A collection of alias resources that match the request parameters.

Type: Array of [Alias](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

ListBuilds

This API works with the following fleet types: EC2

Retrieves build resources for all builds associated with the AWS account in use. You can limit results to builds that are in a specific status by using the `Status` parameter. Use the pagination parameters to retrieve results in

Note

Build resources are not listed in any particular order.

Learn more

[Upload a Custom Server Build](#)

[All APIs by task](#)

Request Syntax

```
{
  "Limit": number,
  "NextToken": "string",
  "Status": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1.

Required: No

Status

Build status to filter results by. To retrieve all builds, leave this parameter empty.

Possible build statuses include the following:

- **INITIALIZED** -- A new build has been defined, but no files have been uploaded. You cannot create fleets for builds that are in this status. When a build is successfully created, the build status is set to this value.
- **READY** -- The game build has been successfully uploaded. You can now create new fleets for this build.
- **FAILED** -- The game build upload failed. You cannot create new fleets for this build.

Type: String

Valid Values: INITIALIZED | READY | FAILED

Required: No

Response Syntax

```
{
  "Builds": [
    {
      "BuildArn": "string",
      "BuildId": "string",
      "CreationTime": number,
      "Name": "string",
      "OperatingSystem": "string",
      "ServerSdkVersion": "string",
      "SizeOnDisk": number,
      "Status": "string",
      "Version": "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.

[Builds](#)

A collection of build resources that match the request.

Type: Array of [Build](#) objects

[NextToken](#)

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Retrieve all builds

This example retrieves game server builds in the current Region. The sample request illustrates how to use the pagination parameters, `Limit` and `NextToken`, to retrieve the results in sequential sets.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "Limit": 2,
  "NextToken":
    "eyJhd3NBWY2NvdW50SWQIOnsicyIGIjMwMjc3NjAxNjM5OCJ9LCJidWlsZE1kIjpw7InMiOiJidWlsZC00NDRlZjQxZS1hM"
}
```

Sample Response

```
{
```

```
"Builds": [
  {
    "BuildArn": "arn:aws:gamelift:us-west-2::build/build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "BuildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "CreationTime": 1495664528.723,
    "Name": "My_Game_Server_Build_One",
    "OperatingSystem": "WINDOWS_2022",
    "SizeOnDisk": 8567781,
    "Status": "READY",
    "Version": "12345.678"
  },
  {
    "BuildArn": "arn:aws:gamelift:us-west-2::build/build-3333cccc-44dd-55ee-66ff-7777aaaa88bb",
    "BuildId": "build-3333cccc-44dd-55ee-66ff-7777aaaa88bb",
    "CreationTime": 1495528748.555,
    "Name": "My_Game_Server_Build_Two",
    "OperatingSystem": "AMAZON_LINUX_2023",
    "SizeOnDisk": 8567781,
    "Status": "INITIALIZED",
    "Version": "23456.789"
  }
],
"NextToken":
"eyJhd3NBW52NvdW50SWQiOnsicyI6IjMwMjc3NjAxNjM5OCJ9LCJidWlsZEIkJp7InMiOiJidWlsZC01NWYxZTZmMS1jY"
}
```

Retrieve failed builds

This example retrieves all game server builds in the current Region that failed to upload. It uses the pagination parameters to retrieve two builds at a time. With no `NextToken` provided, this request tries to retrieve the first two results. There is only one result, so only one build resource and no `NextToken` is returned.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "Limit": 2,
  "Status": "FAILED"
}
```

Sample Response

```
{
  "Builds":
  [
    {"BuildArn": "arn:aws:gamelift:us-west-2::build/build-3333cccc-44dd-55ee-66ff-7777aaaa88bb",
      "BuildId": "build-3333cccc-44dd-55ee-66ff-7777aaaa88bb",
      "CreationTime": 1495528748.555,
      "Name": "My_Game_Server_Build_Two",
      "OperatingSystem": "AMAZON_LINUX_2023",
      "SizeOnDisk": 8567781,
      "Status": "FAILED",
      "Version": "23456.789"}
  ]
}
```

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](#)

ListCompute

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves information on the compute resources in an Amazon GameLift Servers fleet. Use the pagination parameters to retrieve results in a set of sequential pages.

Request options

- Retrieve a list of all computes in a fleet. Specify a fleet ID.
- Retrieve a list of all computes in a specific fleet location. Specify a fleet ID and location.
- Retrieve a list of managed container fleets based on compute status. Specify a fleet ID, status, and (optionally) location.

Results

If successful, this operation returns information on a set of computes. Depending on the type of fleet, the result includes the following information:

- For a managed EC2 fleet (compute type EC2), this operation returns information about the EC2 instance. Compute names are EC2 instance IDs.
- For an Anywhere fleet (compute type ANYWHERE), this operation returns compute names and details from when the compute was registered with `RegisterCompute`. This includes `GameLiftServiceSdkEndpoint` or `GameLiftAgentEndpoint`.
- For a managed container fleet, this operation returns information about the fleet's container groups and the instances they're running on. Compute names refer to game server container groups, which can include multiple containers.

Request Syntax

```
{  
  "ComputeStatus": "string",  
  "ContainerGroupDefinitionName": "string",  
  "FleetId": "string",  
  "Limit": number,  
  "Location": "string",  
  "NextToken": "string"
```

```
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet to retrieve compute resources for.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

ComputeStatus

The status of computes in a managed container fleet, based on the success of the latest update deployment.

- **ACTIVE** -- The compute is deployed with the correct container definitions. It is ready to process game servers and host game sessions.
- **IMPAIRED** -- An update deployment to the compute failed, and the compute is deployed with incorrect container definitions.

Type: String

Valid Values: ACTIVE | IMPAIRED

Required: No

ContainerGroupDefinitionName

For computes in a managed container fleet, the name of the deployed container group definition.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-zA-Z0-9\-\_]+\$|^arn:.*:containergroupdefinition\[a-zA-Z0-9\-\_]+(:[0-9]+)?\$`

Required: No

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

Location

The name of a location to retrieve compute resources for. For an Amazon GameLift Servers Anywhere fleet, use a custom location. For a managed fleet, provide a AWS Region or Local Zone code (for example: `us-west-2` or `us-west-2-lax-1`).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-\_]+\$`

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "ComputeList": [
    {
      "ComputeArn": "string",
      "ComputeName": "string",
      "ComputeStatus": "string",
      "ContainerAttributes": [
        {
          "ContainerName": "string",
          "ContainerRuntimeId": "string"
        }
      ],
      "CreationTime": number,
      "DnsName": "string",
      "FleetArn": "string",
      "FleetId": "string",
      "GameLiftAgentEndpoint": "string",
      "GameLiftServiceSdkEndpoint": "string",
      "GameServerContainerGroupDefinitionArn": "string",
      "InstanceId": "string",
      "IpAddress": "string",
      "Location": "string",
      "OperatingSystem": "string",
      "Type": "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.

ComputeList

A list of compute resources in the specified fleet.

Type: Array of [Compute](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

ListContainerFleets

This API works with the following fleet types: Container

Retrieves a collection of container fleet resources in an AWS Region. For fleets that have multiple locations, this operation retrieves fleets based on their home Region only.

Request options

- Get a list of all fleets. Call this operation without specifying a container group definition.
- Get a list of fleets filtered by container group definition. Provide the container group definition name or ARN value.
- To get a list of all Amazon GameLift Servers Realtime fleets with a specific configuration script, provide the script ID.

Use the pagination parameters to retrieve results as a set of sequential pages.

If successful, this operation returns a collection of container fleets that match the request parameters. A NextToken value is also returned if there are more result pages to retrieve.

Note

Fleet IDs are returned in no particular order.

Request Syntax

```
{  
  "ContainerGroupDefinitionName": "string",  
  "Limit": number,  
  "NextToken": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

ContainerGroupDefinitionName

The container group definition to filter the list on. Use this parameter to retrieve only those fleets that use the specified container group definition. You can specify the container group definition's name to get fleets with the latest versions. Alternatively, provide an ARN value to get fleets with a specific version number.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-zA-Z0-9\-.]+|^arn:.*:containergroupdefinition\[a-zA-Z0-9\-.]+(:[0-9]+)?$`

Required: No

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "ContainerFleets": [
    {
      "BillingType": "string",
      "CreationTime": number,
      "DeploymentDetails": {
        "LatestDeploymentId": "string"
      },
      "Description": "string",
      "FleetArn": "string",
      "FleetId": "string",
      "FleetRoleArn": "string",
      "GameServerContainerGroupDefinitionArn": "string",
      "GameServerContainerGroupDefinitionName": "string",
      "GameServerContainerGroupsPerInstance": number,
      "GameSessionCreationLimitPolicy": {
        "NewGameSessionsPerCreator": number,
        "PolicyPeriodInMinutes": number
      },
      "InstanceConnectionPortRange": {
        "FromPort": number,
        "ToPort": number
      },
      "InstanceInboundPermissions": [
        {
          "FromPort": number,
          "IpRange": "string",
          "Protocol": "string",
          "ToPort": number
        }
      ],
      "InstanceType": "string",
      "LocationAttributes": [
        {
          "Location": "string",
          "PlayerGatewayStatus": "string",
          "Status": "string"
        }
      ],
      "LogConfiguration": {
        "LogDestination": "string",
```

```
    "LogGroupArn": "string",
    "S3BucketName": "string"
  },
  "MaximumGameServerContainerGroupsPerInstance": number,
  "MetricGroups": [ "string" ],
  "NewGameSessionProtectionPolicy": "string",
  "PerInstanceContainerGroupDefinitionArn": "string",
  "PerInstanceContainerGroupDefinitionName": "string",
  "PlayerGatewayMode": "string",
  "Status": "string"
}
],
"NextToken": "string"
}
```

Response Elements

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

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

ContainerFleets

A collection of container fleet objects for all fleets that match the request criteria.

Type: Array of [ContainerFleet](#) objects

Array Members: Minimum number of 1 item.

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

List container fleets with a container group definition

This example retrieves the properties of all fleets with a specified container group definition. Because the request uses the definition name, it returns fleets that are deployed with the latest version of the container group definition. The request uses the pagination parameters to retrieve one fleet at a time. The example response includes a `NextToken`, which indicates that there are still more results to retrieve.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "ContainerGroupDefinitionName": "MyAdventureGameContainerGroup",
```

```

    "Limit": 1,
    "NextToken":
    "eyJhd3NBW52NvdW50SWQl0nsicyI6IjMwMjc3NjAxNjM50CJ9LCJidWlsZE1kIj7InMi0iJidWlsZC01NWYxZTZmMS1jY
  }

```

Sample Response

```

{
  "ContainerFleets": [
    {
      "BillingType": ON_DEMAND,
      "CreationTime": 1736365885.22,
      "DeploymentDetails": {
        "LatestDeploymentId": "deployment-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
      },
      "FleetArn": "arn:aws:gamelift:us-west-2:111122223333:containerfleet/
containerfleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
      "FleetId": "containerfleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
      "FleetRoleArn": "arn:aws:iam::MyAccount:role/MyRole",
      "GameServerContainerGroupDefinitionArn": "arn:aws:gamelift:us-
west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:2",
      "GameServerContainerGroupDefinitionName": "MyAdventureGameContainerGroup",
      "GameServerContainerGroupsPerInstance": number,
      "InstanceConnectionPortRange": {
        "FromPort": 4192,
        "ToPort": 4242
      },
      "InstanceInboundPermissions": [
        {
          "FromPort": 4192,
          "IpRange": "string",
          "Protocol": "UDP",
          "ToPort": 4242,
        }
      ],
      "InstanceType": "c5.large",
      "LogConfiguration": {
        "LogGroupArn": "arn:aws:logs:us-west-2:111222333444:log-group:customerLogs",
        "LogDestination": "CLOUDWATCH"
      },
      "MaximumGameServerContainerGroupsPerInstance": 10,
      "NewGameSessionProtectionPolicy": "NoProtection",
      "Status": "PENDING"
    }
  ]
}

```

```
    }  
  ],  
  "NextToken":  
    "eyJhd3NBdW50SWQi0nsicyI6IjMwMjc3NjAxNjM5OCJ9LCJidWlsZEIkJp7InMi0iJidWlsZC01NWYxZTZmMS1jY  
}
```

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](#)

ListContainerGroupDefinitions

This API works with the following fleet types: Container

Retrieves container group definitions for the AWS account and AWS Region. Use the pagination parameters to retrieve results in a set of sequential pages.

This operation returns only the latest version of each definition. To retrieve all versions of a container group definition, use [ListContainerGroupDefinitionVersions](#).

Request options:

- Retrieve the most recent versions of all container group definitions.
- Retrieve the most recent versions of all container group definitions, filtered by type. Specify the container group type to filter on.

Results:

If successful, this operation returns the complete properties of a set of container group definition versions that match the request.

Note

This operation returns the list of container group definitions in no particular order.

Request Syntax

```
{  
  "ContainerGroupType": "string",  
  "Limit": number,  
  "NextToken": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

ContainerGroupType

The type of container group to retrieve. Container group type determines how Amazon GameLift Servers deploys the container group on each fleet instance.

Type: String

Valid Values: GAME_SERVER | PER_INSTANCE

Required: No

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages.

Type: Integer

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

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "ContainerGroupDefinitions": [
    {
```

```
"ContainerGroupDefinitionArn": "string",
"ContainerGroupType": "string",
"CreationTime": number,
"GameServerContainerDefinition": {
  "ContainerName": "string",
  "DependsOn": [
    {
      "Condition": "string",
      "ContainerName": "string"
    }
  ],
  "EnvironmentOverride": [
    {
      "Name": "string",
      "Value": "string"
    }
  ],
  "ImageUri": "string",
  "LinuxCapabilities": {
    "Include": [ "string" ]
  },
  "MountPoints": [
    {
      "AccessLevel": "string",
      "ContainerPath": "string",
      "InstancePath": "string"
    }
  ],
  "PortConfiguration": {
    "ContainerPortRanges": [
      {
        "FromPort": number,
        "Protocol": "string",
        "ToPort": number
      }
    ]
  },
  "ResolvedImageDigest": "string",
  "ServerSdkVersion": "string"
},
"Name": "string",
"OperatingSystem": "string",
"Status": "string",
"StatusReason": "string",
```

```
"SupportContainerDefinitions": [  
  {  
    "ContainerName": "string",  
    "DependsOn": [  
      {  
        "Condition": "string",  
        "ContainerName": "string"  
      }  
    ],  
    "EnvironmentOverride": [  
      {  
        "Name": "string",  
        "Value": "string"  
      }  
    ],  
    "Essential": boolean,  
    "HealthCheck": {  
      "Command": [ "string" ],  
      "Interval": number,  
      "Retries": number,  
      "StartPeriod": number,  
      "Timeout": number  
    },  
    "ImageUri": "string",  
    "LinuxCapabilities": {  
      "Include": [ "string" ]  
    },  
    "MemoryHardLimitMebibytes": number,  
    "MountPoints": [  
      {  
        "AccessLevel": "string",  
        "ContainerPath": "string",  
        "InstancePath": "string"  
      }  
    ],  
    "PortConfiguration": {  
      "ContainerPortRanges": [  
        {  
          "FromPort": number,  
          "Protocol": "string",  
          "ToPort": number  
        }  
      ]  
    },  
  },  
]
```

```
        "ResolvedImageDigest": "string",
        "Vcpu": number
    }
],
"TotalMemoryLimitMebibytes": number,
"TotalVcpuLimit": number,
"VersionDescription": "string",
"VersionNumber": 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.

ContainerGroupDefinitions

A result set of container group definitions that match the request.

Type: Array of [ContainerGroupDefinition](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Retrieve all game server container group definitions

This example retrieves all container group definitions of type `GAME_SERVER`. It uses the pagination parameters to retrieve two definitions at a time. The response includes a `NextToken`, which indicates that there are still more results to retrieve.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "ContainerGroupType": "GAME_SERVER",
  "Limit": 2
}
```

Sample Response

```
{
  "ContainerGroupDefinitions": [
    {
      "ContainerGroupDefinitionArn": "arn:aws:gamelift:us-
west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:3",
```

```

    "ContainerGroupType": "GAME_SERVER",
    "CreationTime": 1496365885.44,
    "GameServerContainerDefinition": {
        "ContainerName": "MyAdventureGameContainer",
        "ImageUri": "111122223333.dkr.ecr.us-west-2.amazonaws.com/
MyAdventureGameContainerImage",
        "PortConfiguration": {
            "ContainerPortRanges": [
                {
                    "FromPort": 35000,
                    "Protocol": "TCP",
                    "ToPort": 40000
                }
            ]
        },
        "LinuxCapabilities": {
            "Include": ["NET_BIND_SERVICE", "SYS_PTRACE"]
        },
        "ResolvedImageDigest":
"sha256:0123456789abcdef0123456789abcdef0123456789abcdef0123456789abcdef",
        "ServerSdkVersion": "5.2.0"
    },
    "Name": "MyAdventureGameContainerGroup",
    "OperatingSystem": "AMAZON_LINUX_2023",
    "Status": "READY",
    "TotalMemoryLimitMebibytes": 1024,
    "TotalVcpuLimit": 1,
    "VersionNumber": "3"
},
{
    "ContainerGroupDefinitionArn": "arn:aws:gamelift:us-
west-2:111122223333:containergroupdefinition/MyRPGGameContainerGroup:1",
    "ContainerGroupType": "GAME_SERVER",
    "CreationTime": 1496365892.11,
    "GameServerContainerDefinition": {
        "ContainerName": "MyRPGGameContainer",
        "ImageUri": "111122223333.dkr.ecr.us-west-2.amazonaws.com/
MyRPGGameContainerImage",
        "PortConfiguration": {
            "ContainerPortRanges": [
                {
                    "FromPort": 35000,
                    "Protocol": "TCP",
                    "ToPort": 40000
                }
            ]
        }
    }
}

```

```
    }
  ]
},
"LinuxCapabilities": {
  "Include": ["NET_BIND_SERVICE"]
},
"ResolvedImageDigest":
"sha256:0123456789abcdef0123456789abcdef0123456789abcdef0123456789abcdef",
"ServerSdkVersion": "5.2.0"
}
"Name": "MyRPGGameContainerGroup",
"OperatingSystem": "AMAZON_LINUX_2023",
"Status": "COPYING",
"TotalMemoryLimitMebibytes": 1024,
"TotalVcpuLimit": 1,
"VersionNumber": "1"
},
],
"NextToken":
"eyJhd3NBdW50SWQiOnsicyI6IjMwMjc3NjAxNjM5OCJ9LCJidWlsZEIkJp7InMiOiJidWlsZC01NWYxZTZmMS1jY
}
```

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](#)

ListContainerGroupDefinitionVersions

This API works with the following fleet types: Container

Retrieves all versions of a container group definition. Use the pagination parameters to retrieve results in a set of sequential pages.

Request options:

- Get all versions of a specified container group definition. Specify the container group definition name or ARN value. (If the ARN value has a version number, it's ignored.)

Results:

If successful, this operation returns the complete properties of a set of container group definition versions that match the request.

Note

This operation returns the list of container group definitions in descending version order (latest first).

Learn more

- [Manage a container group definition](#)

Request Syntax

```
{  
  "Limit": number,  
  "Name": "string",  
  "NextToken": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

 Note

In the following list, the required parameters are described first.

Name

The unique identifier for the container group definition to retrieve properties for. You can use either the Name or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-zA-Z0-9\-\_]+\$|^arn:.*:containergroupdefinition\[a-zA-Z0-9\-\_]+(:[0-9]+)?\$`

Required: Yes

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages.

Type: Integer

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

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "ContainerGroupDefinitions": [
    {
      "ContainerGroupDefinitionArn": "string",
      "ContainerGroupType": "string",
      "CreationTime": number,
      "GameServerContainerDefinition": {
        "ContainerName": "string",
        "DependsOn": [
          {
            "Condition": "string",
            "ContainerName": "string"
          }
        ],
        "EnvironmentOverride": [
          {
            "Name": "string",
            "Value": "string"
          }
        ],
        "ImageUri": "string",
        "LinuxCapabilities": {
          "Include": [ "string" ]
        },
        "MountPoints": [
          {
            "AccessLevel": "string",
            "ContainerPath": "string",
            "InstancePath": "string"
          }
        ],
        "PortConfiguration": {
          "ContainerPortRanges": [
            {
              "FromPort": number,
              "Protocol": "string",
              "ToPort": number
            }
          ]
        },
        "ResolvedImageDigest": "string",
```

```
    "ServerSdkVersion": "string"
  },
  "Name": "string",
  "OperatingSystem": "string",
  "Status": "string",
  "StatusReason": "string",
  "SupportContainerDefinitions": [
    {
      "ContainerName": "string",
      "DependsOn": [
        {
          "Condition": "string",
          "ContainerName": "string"
        }
      ],
      "EnvironmentOverride": [
        {
          "Name": "string",
          "Value": "string"
        }
      ],
      "Essential": boolean,
      "HealthCheck": {
        "Command": [ "string " ],
        "Interval": number,
        "Retries": number,
        "StartPeriod": number,
        "Timeout": number
      },
      "ImageUri": "string",
      "LinuxCapabilities": {
        "Include": [ "string " ]
      },
      "MemoryHardLimitMebibytes": number,
      "MountPoints": [
        {
          "AccessLevel": "string",
          "ContainerPath": "string",
          "InstancePath": "string"
        }
      ],
      "PortConfiguration": {
        "ContainerPortRanges": [
          {
```

```
        "FromPort": number,
        "Protocol": "string",
        "ToPort": number
    }
]
},
"ResolvedImageDigest": "string",
"Vcpu": number
}
],
"TotalMemoryLimitMebibytes": number,
"TotalVcpuLimit": number,
"VersionDescription": "string",
"VersionNumber": 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.

ContainerGroupDefinitions

A result set of container group definitions that match the request.

Type: Array of [ContainerGroupDefinition](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Retrieve versions of a container group definition

This example retrieves all versions of a container group definition. It uses the pagination parameters to retrieve two definitions at a time. The response includes a `NextToken`, which indicates that there are still more results to retrieve.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "Name": "MyAdventureGameContainerGroup",
  "Limit": 2
}
```

Sample Response

```
{
  "ContainerGroupDefinitions": [
    {
      "ContainerGroupDefinitionArn": "arn:aws:gamelift:us-west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:3",
      "ContainerGroupType": "GAME_SERVER",
      "CreationTime": 1496365885.44,
      "GameServerContainerDefinition": {
        "ContainerName": "MyAdventureGameContainer",
        "ImageUri": "111122223333.dkr.ecr.us-west-2.amazonaws.com/MyAdventureGameContainerImage",
        "PortConfiguration": {
          "ContainerPortRanges": [
            {
              "FromPort": 35000,
              "Protocol": "TCP",
              "ToPort": 40000
            }
          ]
        },
        "LinuxCapabilities": {
          "Include": ["NET_BIND_SERVICE", "SYS_PTRACE"]
        },
        "ResolvedImageDigest": "sha256:0123456789abcdef0123456789abcdef0123456789abcdef0123456789abcdef",
        "ServerSdkVersion": "5.2.0"
      },
      "Name": "MyAdventureGameContainerGroup",
      "OperatingSystem": "AMAZON_LINUX_2023",
      "Status": "READY",
      "TotalMemoryLimitMebibytes": 1024,
      "TotalVcpuLimit": 1,
      "VersionNumber": "3"
    },
  ],
}
```

```

{
  "ContainerGroupDefinitionArn": "arn:aws:gamelift:us-
west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:2",
  "ContainerGroupType": "GAME_SERVER",
  "CreationTime": 1496365892.11,
  "GameServerContainerDefinition": {
    "ContainerName": "MyAdventureGameContainer",
    "ImageUri": "111122223333.dkr.ecr.us-west-2.amazonaws.com/
MyAdventureGameContainerImage",
    "PortConfiguration": {
      "ContainerPortRanges": [
        {
          "FromPort": 35000,
          "Protocol": "TCP",
          "ToPort": 40000
        }
      ]
    },
    "LinuxCapabilities": {
      "Include": ["NET_BIND_SERVICE"]
    },
    "ResolvedImageDigest":
"sha256:0123456789abcdef0123456789abcdef0123456789abcdef0123456789abcdef",
    "ServerSdkVersion": "5.2.0"
  }
  "Name": "MyAdventureGameContainerGroup",
  "OperatingSystem": "AMAZON_LINUX_2023",
  "Status": "COPYING",
  "TotalMemoryLimitMebibytes": 1024,
  "TotalVcpuLimit": 1,
  "VersionNumber": "2"
},
],
"NextToken":
"eyJhd3NBWY2NvdW50SWQlOmsicyI6IjMwMjc3NjAxNjM5OCJ9LCJidWlsZElkIjpw7InMiOiJidWlsZC01NWYxZTZmMS1jY
}

```

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](#)

ListFleetDeployments

This API works with the following fleet types: Container

Retrieves a collection of container fleet deployments in an AWS Region. Use the pagination parameters to retrieve results as a set of sequential pages.

Request options

- Get a list of all deployments. Call this operation without specifying a fleet ID.
- Get a list of all deployments for a fleet. Specify the container fleet ID or ARN value.

Results

If successful, this operation returns a list of deployments that match the request parameters. A NextToken value is also returned if there are more result pages to retrieve.

Note

Deployments are returned starting with the latest.

Request Syntax

```
{  
  "FleetId": "string",  
  "Limit": number,  
  "NextToken": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the container fleet. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: No

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "FleetDeployments": [
    {
```

```
    "CreationTime": number,
    "DeploymentConfiguration": {
      "ImpairmentStrategy": "string",
      "MinimumHealthyPercentage": number,
      "ProtectionStrategy": "string"
    },
    "DeploymentId": "string",
    "DeploymentStatus": "string",
    "FleetId": "string",
    "GameServerBinaryArn": "string",
    "PerInstanceBinaryArn": "string",
    "RollbackGameServerBinaryArn": "string",
    "RollbackPerInstanceBinaryArn": "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.

[FleetDeployments](#)

The requested deployment information.

Type: Array of [FleetDeployment](#) objects

[NextToken](#)

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

List deployments for a container fleet

This example retrieves the properties of all deployments for a specified container fleets. The request uses the pagination parameters to retrieve one fleet at a time. The example response includes a `NextToken`, which indicates that there are still more results to retrieve.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "FleetId": "containerfleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "Limit": 1
}
```

Sample Response

```
{
  "FleetDeployments": [
    {
      "CreationTime": 1736365885.22,
      "DeploymentConfiguration": {
        "ImpairmentStrategy": "ROLLBACK",
        "MinimumHealthyPercentage": 30,
        "ProtectionStrategy": "WITH_PROTECTION"
      },
      "DeploymentId": "deployment-3333aaaa-44bb-55cc-66dd-7777eeee88ff",
      "DeploymentStatus": "COMPLETE",
      "FleetId": "containerfleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
      "GameServerBinaryArn": "arn:aws:gamelift:us-west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:2",
      "RollbackGameServerBinaryArn": "arn:aws:gamelift:us-west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:1",
    }
  ],
  "NextToken": "eyJhd3NBdW50SWQi0nsicyI6IjMwMjc3NjAxNjM5OCJ9LCJidWlsZEIkJp7InMi0iJidWlsZC01NWYxZTZmMS1jY"
}
```

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](#)

ListFleets

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves a collection of fleet resources in an AWS Region. You can filter the result set to find only those fleets that are deployed with a specific build or script. For fleets that have multiple locations, this operation retrieves fleets based on their home Region only.

You can use operation in the following ways:

- To get a list of all fleets in a Region, don't provide a build or script identifier.
- To get a list of all fleets where a specific game build is deployed, provide the build ID.
- To get a list of all Amazon GameLift Servers Realtime fleets with a specific configuration script, provide the script ID.

Use the pagination parameters to retrieve results as a set of sequential pages.

If successful, this operation returns a list of fleet IDs that match the request parameters. A NextToken value is also returned if there are more result pages to retrieve.

Note

Fleet IDs are returned in no particular order.

Request Syntax

```
{  
  "BuildId": "string",  
  "Limit": number,  
  "NextToken": "string",  
  "ScriptId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

BuildId

A unique identifier for the build to request fleets for. Use this parameter to return only fleets using a specified build. Use either the build ID or ARN value.

Type: String

Pattern: `^build-\S+|^arn:.*:build\/build-\S+`

Required: No

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

ScriptId

A unique identifier for the Realtime script to request fleets for. Use this parameter to return only fleets using a specified script. Use either the script ID or ARN value.

Type: String

Pattern: `^script-\S+|^arn:.*:script\/script-\S+`

Required: No

Response Syntax

```
{
  "FleetIds": [ "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.

FleetIds

A set of fleet IDs that match the list request.

Type: Array of strings

Array Members: Minimum number of 1 item.

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+`

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

List fleets in a Region

This example retrieves the fleet IDs of all fleets with their home Region in the currently selected Region. It uses the pagination parameters to retrieve two fleet IDs at a time. The example response includes a `NextToken`, which indicates that there are still more results to retrieve.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "Limit": 2,
  "NextToken":
    "eyJhd3NBZjY2NvdW50SWQi0nsicyI6IjMwMjc3NjAxNjM5MCJ9LCJidWlsZEIkJp7InMi0iJidWlsZC01NWYxZTZmMS1jY
```

```
}
```

Sample Response

```
{
  "FleetIds": [
    "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
    "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa"
  ],
  "NextToken":
  "eyJhd3NBWY2NvdW50SWQi0nsicyI6IjMwMjc3NjAxNjM5OCJ9LCJidWlsZEIkJp7InMi0iJidWlsZC01NWYxZTZmMS1jY"
}
```

List all fleets in a Region with a specific build or script

This example retrieves the IDs of fleets in the currently selected Region that are deployed with a specified game build. If you're working with Realtime Servers, you can opt to provide a script ID in place of a build ID. This example does not specify the limit parameter, so results can include up to 16 fleet IDs.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "Build": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
}
```

Sample Response

```
{
  "FleetIds": ["fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"]
}
```

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](#)

ListGameServerGroups

This API works with the following fleet types: EC2 (FleetIQ)

Lists a game server groups.

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Limit

The game server groups' limit.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

NextToken

Specify the pagination token from a previous request to retrieve the next page of results.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "GameServerGroups": [
    {
      "AutoScalingGroupArn": "string",
      "BalancingStrategy": "string",
      "CreationTime": number,
      "GameServerGroupArn": "string",
      "GameServerGroupName": "string",
      "GameServerProtectionPolicy": "string",
      "InstanceDefinitions": [
        {
          "InstanceType": "string",
          "WeightedCapacity": "string"
        }
      ],
      "LastUpdatedTime": number,
      "RoleArn": "string",
      "Status": "string",
      "StatusReason": "string",
      "SuspendedActions": [ "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.

GameServerGroups

The game server groups' game server groups.

Type: Array of [GameServerGroup](#) objects

NextToken

Specify the pagination token from a previous request to retrieve the next page of results.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

ListGameServers

This API works with the following fleet types: EC2 (FleetIQ)

Retrieves information on all game servers that are currently active in a specified game server group. You can opt to sort the list by game server age. Use the pagination parameters to retrieve results in a set of sequential segments.

Learn more

[Amazon GameLift Servers FleetIQ Guide](#)

Request Syntax

```
{  
  "GameServerGroupName": "string",  
  "Limit": number,  
  "NextToken": "string",  
  "SortOrder": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

GameServerGroupName

An identifier for the game server group to retrieve a list of game servers from. Use either the name or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.\+|^arn:.*:gameservergroup\/[a-zA-Z0-9-\.\+]`

Required: Yes

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

SortOrder

Indicates how to sort the returned data based on game server registration timestamp. Use `ASCENDING` to retrieve oldest game servers first, or use `DESCENDING` to retrieve newest game servers first. If this parameter is left empty, game servers are returned in no particular order.

Type: String

Valid Values: `ASCENDING` | `DESCENDING`

Required: No

Response Syntax

```
{
```

```
"GameServers": [  
  {  
    "ClaimStatus": "string",  
    "ConnectionInfo": "string",  
    "GameServerData": "string",  
    "GameServerGroupArn": "string",  
    "GameServerGroupName": "string",  
    "GameServerId": "string",  
    "InstanceId": "string",  
    "LastClaimTime": number,  
    "LastHealthCheckTime": number,  
    "RegistrationTime": number,  
    "UtilizationStatus": "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.

GameServers

A collection of game server objects that match the request.

Type: Array of [GameServer](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

ListLocations

This API works with the following fleet types: EC2, Anywhere, Container

Lists all custom and AWS locations where Amazon GameLift Servers can host game servers. This operation also returns UDP ping beacon information for locations, which you can use to measure network latency between player devices and potential hosting locations.

The set of locations returned by this API depends on the AWS Region you specify when calling this API (or your default Region if you don't specify one).

When you call from:

- **A home Region of a fleet that supports multi-locations:** API returns information for *all* hosting locations
- **A home Region of a fleet that supports a single location:** API returns information for that location

Note that if you call this API using a location that doesn't have a service endpoint, such as one that can only be a remote location in a multi-location fleet, the API returns an error.

Consult the table of supported locations in [Amazon GameLift Servers service locations](#) to identify home Regions that support single and multi-location fleets.

Avoid calling ListLocations before each latency measurement, as it's not designed for high-volume requests. The beacon endpoint information returned by this API are static and won't change as long as Amazon GameLift Servers supports game hosting in that location.

Learn more

[Service locations](#)

[UDP ping beacons](#)

Request Syntax

```
{  
  "Filters": [ "string" ],  
  "Limit": number,  
  "NextToken": "string"
```

```
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Filters

Filters the list for AWS or CUSTOM locations. Use this parameter to narrow down results to only AWS-managed locations (Amazon EC2 or container) or only your custom locations (such as an Amazon GameLift Servers Anywhere fleet).

Type: Array of strings

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

Valid Values: AWS | CUSTOM

Required: No

Limit

The maximum number of results to return. Use this parameter with NextToken to get results as a set of sequential pages.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 200.

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "Locations": [
    {
      "LocationArn": "string",
      "LocationName": "string",
      "PingBeacon": {
        "UDPEndpoint": {
          "Domain": "string",
          "Port": 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.

Locations

A collection of locations, including both AWS and custom locations. Each location includes a name and ping beacon information that can be used to measure network latency between player devices and the location.

Type: Array of [LocationModel](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

ListScripts

This API works with the following fleet types: EC2

Retrieves script records for all Realtime scripts that are associated with the AWS account in use.

Learn more

[Amazon GameLift Servers Amazon GameLift Servers Realtime](#)

Related actions

[All APIs by task](#)

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Limit

The maximum number of results to return. Use this parameter with NextToken to get results as a set of sequential pages.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1.

Required: No

Response Syntax

```
{
  "NextToken": "string",
  "Scripts": [
    {
      "CreationTime": number,
      "Name": "string",
      "NodeJsVersion": "string",
      "ScriptArn": "string",
      "ScriptId": "string",
      "SizeOnDisk": number,
      "StorageLocation": {
        "Bucket": "string",
        "Key": "string",
        "ObjectVersion": "string",
        "RoleArn": "string"
      },
      "Version": "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

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1.

Scripts

A set of properties describing the requested script.

Type: Array of [Script](#) objects

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Retrieve all scripts

This example retrieves the Realtime scripts in the current Region. The example illustrates using the pagination parameters to retrieve the results in sequential sets. This sample request uses a `NextToken` value that was returned in a previous `ListScripts` request. The response shows two script records; the first script was uploaded from an Amazon S3 bucket, and the second script was uploaded from a local zip file.

Sample Request

```
{
  "Limit": 2,
  "NextToken":
    "eyJhd3NBW50SWQ0nsicyI6IjMwMjc3NjAxNjM5OCJ9LCJidWlsZElkIjp7InMiOiJidWlsZC00NDRLZjQxZS1hM
}
```

CLI syntax:

```
aws gamelift list-scripts
  -limit 2
  -next-token
    "eyJhd3NBW50SWQ0nsicyI6IjMwMjc3NjAxNjM5OCJ9LCJidWlsZElkIjp7InMiOiJidWlsZC00NDRLZjQxZS1hM
```

Sample Response

```
{
  "Scripts": [
    {
      "CreationTime": 1496708916.18,
      "Name": "My_Realtime_Script_2",
      "NodeJsVersion": "24.x",
      "ScriptArn": "arn:aws:gamelift:us-west-2::script/
script-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
      "ScriptId": "script-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
      "SizeOnDisk": 0,
      "StorageLocation": {
        "Bucket": "my_realtime_script_files",
        "Key": "myRealtimeScript.zip"
        "RoleArn": "arn:aws:iam::111122223333:role/GameLiftAccess"
        "ObjectVersion": null
      },
    },
  ],
}
```

```
    "Version": "12345.678"
  },
  {
    "CreationTime": 1495528748.555,
    "Name": "My_Realtime_Script_1",
    "NodeJsVersion": "24.x",
    "ScriptArn": "arn:aws:gamelift:us-west-2::script/
script-3333cccc-44dd-55ee-66ff-7777aaaa88bb",
    "ScriptId": "script-3333cccc-44dd-55ee-66ff-7777aaaa88bb",
    "SizeOnDisk": 9000,
    "StorageLocation": {
      "Bucket": "prod-gamescale-scripts-us-west-2",
      "Key": "123456789012/script-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
    },
    "Version": "1.0.1"
  }
  "NextToken":
  "kyJhd3NBY2NvdW50SWQi0nsicyI6IjMwMjc3NjAxNjM5OCJ9LCJidWlsZEIkJp7InMi0iJidWlsZC01NWYxZTZmMS1jY
}
```

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

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves all tags assigned to a Amazon GameLift Servers resource. Use resource tags to organize AWS resources for a range of purposes. This operation handles the permissions necessary to manage tags for Amazon GameLift Servers resources that support tagging.

To list tags for a resource, specify the unique ARN value for the resource.

Learn more

[Tagging AWS Resources](#) in the *AWS General Reference*

[AWS Tagging Strategies](#)

Related actions

[All APIs by task](#)

Request Syntax

```
{  
  "ResourceARN": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[ResourceARN](#)

The Amazon Resource Name ([ARN](#)) that uniquely identifies the Amazon GameLift Servers resource that you want to retrieve tags for. Amazon GameLift Servers includes resource ARNs

in the data object for the resource. You can retrieve the ARN by calling a `List` or `Describe` operation for the resource type.

Type: String

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

Required: Yes

Response Syntax

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

Response Elements

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

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

Tags

The collection of tags assigned to the resource.

Type: Array of [Tag](#) objects

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

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

PutScalingPolicy

This API works with the following fleet types: EC2, Container

Creates or updates a scaling policy for a fleet. Scaling policies are used to automatically scale a fleet's hosting capacity to meet player demand. An active scaling policy instructs Amazon GameLift Servers to track a fleet metric and automatically change the fleet's capacity when a certain threshold is reached. There are two types of scaling policies: target-based and rule-based. Use a target-based policy to quickly and efficiently manage fleet scaling; this option is the most commonly used. Use rule-based policies when you need to exert fine-grained control over auto-scaling.

Fleets can have multiple scaling policies of each type in force at the same time; you can have one target-based policy, one or multiple rule-based scaling policies, or both. We recommend caution, however, because multiple auto-scaling policies can have unintended consequences.

Learn more about how to work with auto-scaling in [Set Up Fleet Automatic Scaling](#).

Target-based policy

A target-based policy tracks a single metric: `PercentAvailableGameSessions`. This metric tells us how much of a fleet's hosting capacity is ready to host game sessions but is not currently in use. This is the fleet's buffer; it measures the additional player demand that the fleet could handle at current capacity. With a target-based policy, you set your ideal buffer size and leave it to Amazon GameLift Servers to take whatever action is needed to maintain that target.

For example, you might choose to maintain a 10% buffer for a fleet that has the capacity to host 100 simultaneous game sessions. This policy tells Amazon GameLift Servers to take action whenever the fleet's available capacity falls below or rises above 10 game sessions. Amazon GameLift Servers will start new instances or stop unused instances in order to return to the 10% buffer.

To create or update a target-based policy, specify a fleet ID and name, and set the policy type to "TargetBased". Specify the metric to track (`PercentAvailableGameSessions`) and reference a `TargetConfiguration` object with your desired buffer value. Exclude all other parameters. On a successful request, the policy name is returned. The scaling policy is automatically in force as soon as it's successfully created. If the fleet's auto-scaling actions are temporarily suspended, the new policy will be in force once the fleet actions are restarted.

Rule-based policy

A rule-based policy tracks specified fleet metric, sets a threshold value, and specifies the type of action to initiate when triggered. With a rule-based policy, you can select from several available fleet metrics. Each policy specifies whether to scale up or scale down (and by how much), so you need one policy for each type of action.

For example, a policy may make the following statement: "If the percentage of idle instances is greater than 20% for more than 15 minutes, then reduce the fleet capacity by 10%."

A policy's rule statement has the following structure:

If [MetricName] is [ComparisonOperator] [Threshold] for [EvaluationPeriods] minutes, then [ScalingAdjustmentType] to/by [ScalingAdjustment].

To implement the example, the rule statement would look like this:

If [PercentIdleInstances] is [GreaterThanThreshold] [20] for [15] minutes, then [PercentChangeInCapacity] to/by [10].

To create or update a scaling policy, specify a unique combination of name and fleet ID, and set the policy type to "RuleBased". Specify the parameter values for a policy rule statement. On a successful request, the policy name is returned. Scaling policies are automatically in force as soon as they're successfully created. If the fleet's auto-scaling actions are temporarily suspended, the new policy will be in force once the fleet actions are restarted.

Request Syntax

```
{
  "ComparisonOperator": "string",
  "EvaluationPeriods": number,
  "FleetId": "string",
  "MetricName": "string",
  "Name": "string",
  "PolicyType": "string",
  "ScalingAdjustment": number,
  "ScalingAdjustmentType": "string",
  "TargetConfiguration": {
    "TargetValue": number
  },
  "Threshold": number
}
```


Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet to apply this policy to. You can use either the fleet ID or ARN value. The fleet cannot be in any of the following statuses: ERROR or DELETING.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

MetricName

Name of the Amazon GameLift Servers-defined metric that is used to trigger a scaling adjustment. For detailed descriptions of fleet metrics, see [Monitor Amazon GameLift Servers with Amazon CloudWatch](#).

- **ActivatingGameSessions** -- Game sessions in the process of being created.
- **ActiveGameSessions** -- Game sessions that are currently running.
- **ActiveInstances** -- Fleet instances that are currently running at least one game session.
- **AvailableGameSessions** -- Additional game sessions that fleet could host simultaneously, given current capacity.
- **AvailablePlayerSessions** -- Empty player slots in currently active game sessions. This includes game sessions that are not currently accepting players. Reserved player slots are not included.
- **CurrentPlayerSessions** -- Player slots in active game sessions that are being used by a player or are reserved for a player.

- **IdleInstances** -- Active instances that are currently hosting zero game sessions.
- **PercentAvailableGameSessions** -- Unused percentage of the total number of game sessions that a fleet could host simultaneously, given current capacity. Use this metric for a target-based scaling policy.
- **PercentIdleInstances** -- Percentage of the total number of active instances that are hosting zero game sessions.
- **QueueDepth** -- Pending game session placement requests, in any queue, where the current fleet is the top-priority destination.
- **WaitTime** -- Current wait time for pending game session placement requests, in any queue, where the current fleet is the top-priority destination.

Type: String

Valid Values: `ActivatingGameSessions` | `ActiveGameSessions` | `ActiveInstances` | `AvailableGameSessions` | `AvailablePlayerSessions` | `CurrentPlayerSessions` | `IdleInstances` | `PercentAvailableGameSessions` | `PercentIdleInstances` | `QueueDepth` | `WaitTime` | `ConcurrentActivatableGameSessions`

Required: Yes

Name

A descriptive label that is associated with a fleet's scaling policy. Policy names do not need to be unique. A fleet can have only one scaling policy with the same name.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

ComparisonOperator

Comparison operator to use when measuring the metric against the threshold value.

Type: String

Valid Values: `GreaterThanOrEqualToThreshold` | `GreaterThanThreshold` | `LessThanThreshold` | `LessThanOrEqualToThreshold`

Required: No

EvaluationPeriods

Length of time (in minutes) the metric must be at or beyond the threshold before a scaling event is triggered.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

PolicyType

The type of scaling policy to create. For a target-based policy, set the parameter *MetricName* to 'PercentAvailableGameSessions' and specify a *TargetConfiguration*. For a rule-based policy set the following parameters: *MetricName*, *ComparisonOperator*, *Threshold*, *EvaluationPeriods*, *ScalingAdjustmentType*, and *ScalingAdjustment*.

Type: String

Valid Values: RuleBased | TargetBased

Required: No

ScalingAdjustment

Amount of adjustment to make, based on the scaling adjustment type.

Type: Integer

Required: No

ScalingAdjustmentType

The type of adjustment to make to a fleet's instance count:

- **ChangeInCapacity** -- add (or subtract) the scaling adjustment value from the current instance count. Positive values scale up while negative values scale down.
- **ExactCapacity** -- set the instance count to the scaling adjustment value.
- **PercentChangeInCapacity** -- increase or reduce the current instance count by the scaling adjustment, read as a percentage. Positive values scale up while negative values scale down; for example, a value of "-10" scales the fleet down by 10%.

Type: String

Valid Values: `ChangeInCapacity` | `ExactCapacity` | `PercentChangeInCapacity`

Required: No

TargetConfiguration

An object that contains settings for a target-based scaling policy.

Type: [TargetConfiguration](#) object

Required: No

Threshold

Metric value used to trigger a scaling event.

Type: Double

Required: No

Response Syntax

```
{  
  "Name": "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.

Name

A descriptive label that is associated with a fleet's scaling policy. Policy names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Create a target-based scaling policy

This example sets up auto-scaling using a target-based scaling policy. For this fleet, we want to maintain a 15% capacity buffer for our game, so that our fleet will always be able to immediately accommodate some additional game sessions. For a target-based policy, we need to specify a fleet

ID, policy name and type, metric name (set this parameter to "PercentAvailableGameSessions"), and target configuration (buffer size).

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
POST / HTTP/1.1
Host: gamelift.us-west-2.amazonaws.com;
Accept-Encoding: identity
Content-Length: 338
User-Agent: aws-cli/1.11.36 Python/2.7.9 Windows/7 botocore/1.4.93
Content-Type: application/x-amz-json-1.0
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20170406/us-west-2/
gamelift/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target,
Signature=wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
X-Amz-Date: 20170406T004805Z
X-Amz-Target: GameLift.PutScalingPolicy

{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "Name": "My_Target_Policy_1",
  "PolicyType": "TargetBased",
  "MetricName": "PercentAvailableGameSessions",
  "TargetConfiguration": {"TargetValue": 15}
}
```

CLI syntax:

```
$aws gamelift put-scaling-policy
--fleet-id "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
--name "My_Target_Policy_1"
--policy-type "TargetBased"
--metric-name "PercentAvailableGameSessions"
--target-configuration "TargetValue=5"
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: b34f8665-EXAMPLE
Content-Type: application/x-amz-json-1.1
Content-Length: 607
```

```
Date: Thu, 06 Apr 2017 00:48:07 GMT
```

```
{
  "Name": "My_Target_Policy_1"
}
```

Create a rule-based scaling policy

This example illustrates using a rule-based policy to supplement a target-based policy. While the target policy does most of the work of ensuring that capacity tracks with player demand, a well-formed rule-based policy can handle special circumstances and edge cases. For example, the target-based approach becomes less efficient when fleets have just few instances. We can mitigate this issue by creating a rule that maintains at least one idle instance ready to host new game sessions. At low capacity, the two policies do not conflict; at higher capacity, the rule-based policy loses relevance.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
POST / HTTP/1.1
Host: gamelift.us-west-2.amazonaws.com;
Accept-Encoding: identity
Content-Length: 336
User-Agent: aws-cli/1.11.36 Python/2.7.9 Windows/7 botocore/1.4.93
Content-Type: application/x-amz-json-1.0
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20170406/us-west-2/
gamelift/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target,
Signature=wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
X-Amz-Date: 20170406T004805Z
X-Amz-Target: GameLift.PutScalingPolicy

{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "Name": "My_Rule_Policy_1",
  "PolicyType": "RuleBased",
  "MetricName": "IdleInstances",
  "ComparisonOperator": "LessThanThreshold",
  "Threshold": "2"
  "EvaluationPeriods": "5"
  "ScalingAdjustmentType": "ChangeInCapacity"
```

```
    "ScalingAdjustment": "1"
  }
}
```

CLI syntax:

```
$aws gamelift put-scaling-policy
--fleet-id "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
--name "My_Rule_Policy_1"
--policy-type "RuleBased"
--metric-name "IdleInstances"
--comparison-operator "LessThanThreshold"
--threshold "2"
--evaluation-periods "5"
--scaling-adjustment-type "ChangeInCapacity"
--scaling-adjustment "1"
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: b34f8665-EXAMPLE
Content-Type: application/x-amz-json-1.1
Content-Length: 600
Date: Thu, 06 Apr 2017 00:48:07 GMT

{
  "Name": "My_Rule_Policy_1"
}
```

See Also

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

- [AWS Command Line Interface 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](#)

RegisterCompute

This API works with the following fleet types: Anywhere

Registers a compute resource in an Amazon GameLift Servers Anywhere fleet.

For an Anywhere fleet that's running the Amazon GameLift Servers Agent, the Agent handles all compute registry tasks for you. For an Anywhere fleet that doesn't use the Agent, call this operation to register fleet computes.

To register a compute, give the compute a name (must be unique within the fleet) and specify the compute resource's DNS name or IP address. Provide a fleet ID and a fleet location to associate with the compute being registered. You can optionally include the path to a TLS certificate on the compute resource.

If successful, this operation returns compute details, including an Amazon GameLift Servers SDK endpoint or Agent endpoint. Game server processes running on the compute can use this endpoint to communicate with the Amazon GameLift Servers service. Each server process includes the SDK endpoint in its call to the Amazon GameLift Servers server SDK action `InitSDK()`.

To view compute details, call [DescribeCompute](#) with the compute name.

Learn more

- [Create an Anywhere fleet](#)
- [Test your integration](#)
- [Server SDK reference guides](#) (for version 5.x)

Request Syntax

```
{  
  "CertificatePath": "string",  
  "ComputeName": "string",  
  "DnsName": "string",  
  "FleetId": "string",  
  "IpAddress": "string",  
  "Location": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

ComputeName

A descriptive label for the compute resource.

Type: String

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

Pattern: `^[a-zA-Z0-9\-\_]+(\/[a-zA-Z0-9\-\_]+)?`

Required: Yes

FleetId

A unique identifier for the fleet to register the compute to. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+$|^arn:.*:[a-z]*fleet\[a-zA-Z0-9\-\_]+$`

Required: Yes

CertificatePath

The path to a TLS certificate on your compute resource. Amazon GameLift Servers doesn't validate the path and certificate.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

DnsName

The DNS name of the compute resource. Amazon GameLift Servers requires either a DNS name or IP address.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^[0-9a-zA-Z_\-\.]+`

Required: No

IpAddress

The IP address of the compute resource. Amazon GameLift Servers requires either a DNS name or IP address. When registering an Anywhere fleet, an IP address is required.

Type: String

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

Pattern: `^[0-9A-Fa-f\:\.]+`

Required: No

Location

The name of a custom location to associate with the compute resource being registered. This parameter is required when registering a compute for an Anywhere fleet.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-]+`

Required: No

Response Syntax

```
{
  "Compute": {
    "ComputeArn": "string",
```

```
"ComputeName": "string",
"ComputeStatus": "string",
"ContainerAttributes": [
  {
    "ContainerName": "string",
    "ContainerRuntimeId": "string"
  }
],
"CreationTime": number,
"DnsName": "string",
"FleetArn": "string",
"FleetId": "string",
"GameLiftAgentEndpoint": "string",
"GameLiftServiceSdkEndpoint": "string",
"GameServerContainerGroupDefinitionArn": "string",
"InstanceId": "string",
"IpAddress": "string",
"Location": "string",
"OperatingSystem": "string",
"Type": "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.

[Compute](#)

The details of the compute resource you registered.

Type: [Compute](#) object

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

NotReadyException

The operation failed because Amazon GameLift Servers has not yet finished validating this compute. We recommend attempting 8 to 10 retries over 3 to 5 minutes with [exponential backoffs and jitter](#).

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Register your laptop as a testing compute resource

The following example registers a development laptop with the IP address of 10.1.2.3.4 to an Anywhere fleet with the fleet ID `fleet-1234`.

Sample Request

```
{
  "FleetId": "fleet-1234",
  "ComputeName": "DevLaptop",
  "IpAddress": "10.1.2.3",
  "Location": "custom-location1"
}
```

Sample Response

```
{
  FleetId = fleet-1234,
  ComputeName = DevLaptop,
  Status = ACTIVE,
  IpAddress = 10.1.2.3,
  GameLiftServiceSdkEndpoint = wss://12345678.execute-api.amazonaws.com/
}
```

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](#)

RegisterGameServer

This API works with the following fleet types: EC2 (FleetIQ)

Creates a new game server resource and notifies Amazon GameLift Servers FleetIQ that the game server is ready to host gameplay and players. This operation is called by a game server process that is running on an instance in a game server group. Registering game servers enables Amazon GameLift Servers FleetIQ to track available game servers and enables game clients and services to claim a game server for a new game session.

To register a game server, identify the game server group and instance where the game server is running, and provide a unique identifier for the game server. You can also include connection and game server data.

Once a game server is successfully registered, it is put in status AVAILABLE. A request to register a game server may fail if the instance it is running on is in the process of shutting down as part of instance balancing or scale-down activity.

Learn more

[Amazon GameLift Servers FleetIQ Guide](#)

Request Syntax

```
{  
  "ConnectionInfo": "string",  
  "GameServerData": "string",  
  "GameServerGroupName": "string",  
  "GameServerId": "string",  
  "InstanceId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

GameServerGroupName

A unique identifier for the game server group where the game server is running.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.\.]+\|^arn:.*:gameservergroup\[a-zA-Z0-9-\.\.\.]+`

Required: Yes

GameServerId

A custom string that uniquely identifies the game server to register. Game server IDs are developer-defined and must be unique across all game server groups in your AWS account.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `[a-zA-Z0-9-\.\.]+`

Required: Yes

InstanceId

The unique identifier for the instance where the game server is running. This ID is available in the instance metadata. EC2 instance IDs use a 17-character format, for example: `i-1234567890abcdef0`.

Type: String

Length Constraints: Fixed length of 19.

Pattern: `^i-[0-9a-zA-Z]{17}$`

Required: Yes

ConnectionInfo

Information that is needed to make inbound client connections to the game server. This might include the IP address and port, DNS name, and other information.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `.*\S.*`

Required: No

GameServerData

A set of custom game server properties, formatted as a single string value. This data is passed to a game client or service when it requests information on game servers.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `.*\S.*`

Required: No

Response Syntax

```
{
  "GameServer": {
    "ClaimStatus": "string",
    "ConnectionInfo": "string",
    "GameServerData": "string",
    "GameServerGroupArn": "string",
    "GameServerGroupName": "string",
    "GameServerId": "string",
    "InstanceId": "string",
    "LastClaimTime": number,
    "LastHealthCheckTime": number,
    "RegistrationTime": number,
    "UtilizationStatus": "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.

GameServer

Object that describes the newly registered game server.

Type: [GameServer](#) object

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Register a game server

This example illustrates how an game server process notifies GameLift that it is ready to host a game.

Sample Request

```
{
  "GameServerGroupName": "MegaFrogServers_NA",
  "ConnectionInfo": "192.0.2.0.80",
  "GameServerId": "mega-frog-game-12345678",
  "InstanceId": "i-1234567890abcdef0"
}
```

CLI command:

```
aws gamelift register-game-server \
  --game-server-group-name MegaFrogServers_NA \
  --connection-info "192.0.2.0.80" \
  --game-server-id mega-frog-game-12345678 \
  --instance-id i-1234567890abcdef0
```

Sample Response

```
{
  "GameServer": {
    "ClaimStatus": "",
    "ConnectionInfo": "192.0.2.0.80",
    "GameServerData": "",
    "GameServerGroupArn": "arn:aws:gamelift:us-west-2::GameServerGroup/MegaFrogServers_NA",
    "GameServerGroupName": "MegaFrogServers_NA",
    "GameServerId": "mega-frog-game-12345678",
    "InstanceId": "i-1234567890abcdef0",
    "LastClaimTime": ,
    "LastHealthCheckTime": 1580218197.293,
    "RegistrationTime": 1580218197.293,
    "UtilizationStatus": "AVAILABLE"
  }
}
```

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](#)

RequestUploadCredentials

This API works with the following fleet types: EC2

Retrieves a fresh set of credentials for use when uploading a new set of game build files to Amazon GameLift Servers's Amazon S3. This is done as part of the build creation process; see [CreateBuild](#).

To request new credentials, specify the build ID as returned with an initial `CreateBuild` request. If successful, a new set of credentials are returned, along with the S3 storage location associated with the build ID.

Learn more

[Create a Build with Files in S3](#)

[All APIs by task](#)

Request Syntax

```
{  
  "BuildId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

BuildId

A unique identifier for the build to get credentials for. You can use either the build ID or ARN value.

Type: String

Pattern: ^build-\\S+|^arn:.*:build\\/build-\\S+

Required: Yes

Response Syntax

```
{
  "StorageLocation": {
    "Bucket": "string",
    "Key": "string",
    "ObjectVersion": "string",
    "RoleArn": "string"
  },
  "UploadCredentials": {
    "AccessKeyId": "string",
    "SecretAccessKey": "string",
    "SessionToken": "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.

StorageLocation

Amazon S3 path and key, identifying where the game build files are stored.

Type: [S3Location](#) object

UploadCredentials

AWS credentials required when uploading a game build to the storage location. These credentials have a limited lifespan and are valid only for the build they were issued for.

Type: [AwsCredentials](#) object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Refresh access credentials for uploading a build

This example obtains new, valid access credentials for uploading a build file to a Amazon GameLift Servers Amazon S3 location. Credentials have a limited lift span. The build ID required for this operation is returned in response to the original CreateBuild request.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "BuildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
}
```


Sample Response

```
{
  "StorageLocation": {
    "Bucket": "gamelift-builds-us-west-2",
    "Key": "123456789012/build-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
  },
  "UploadCredentials": {
    "AccessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "SecretAccessKey": "wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY",
    "SessionToken": "AgoGb3JpZ2luENz...EXAMPLETOKEN=="
  }
}
```

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](#)

ResolveAlias

This API works with the following fleet types: EC2, Anywhere, Container

Attempts to retrieve a fleet ID that is associated with an alias. Specify a unique alias identifier.

If the alias has a SIMPLE routing strategy, Amazon GameLift Servers returns a fleet ID. If the alias has a TERMINAL routing strategy, the result is a `TerminalRoutingStrategyException`.

Related actions

[All APIs by task](#)

Request Syntax

```
{  
  "AliasId": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[AliasId](#)

The unique identifier of the alias that you want to retrieve a fleet ID for. You can use either the alias ID or ARN value.

Type: String

Pattern: `^alias-\S+|^arn:.*:alias\/alias-\S+`

Required: Yes

Response Syntax

```
{  
  "FleetArn": "string",  
  "FleetId": "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.

FleetArn

The Amazon Resource Name ([ARN](#)) associated with the GameLift fleet resource that this alias points to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet[/[a-z]*fleet-[a-zA-Z0-9\ -]]+$`

FleetId

The fleet identifier that the alias is pointing to.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\ -]]+`

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TerminalRoutingStrategyException

The service is unable to resolve the routing for a particular alias because it has a terminal RoutingStrategy associated with it. The message returned in this exception is the message defined in the routing strategy itself. Such requests should only be retried if the routing strategy for the specified alias is modified.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

ResumeGameServerGroup

This API works with the following fleet types: EC2 (FleetIQ)

Reinstates activity on a game server group after it has been suspended. A game server group might be suspended by the [SuspendGameServerGroup](#) operation, or it might be suspended involuntarily due to a configuration problem. In the second case, you can manually resume activity on the group once the configuration problem has been resolved. Refer to the game server group status and status reason for more information on why group activity is suspended.

To resume activity, specify a game server group ARN and the type of activity to be resumed. If successful, a GameServerGroup object is returned showing that the resumed activity is no longer listed in SuspendedActions.

Learn more

[Amazon GameLift Servers FleetIQ Guide](#)

Request Syntax

```
{
  "GameServerGroupName": "string",
  "ResumeActions": [ "string" ]
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[GameServerGroupName](#)

A unique identifier for the game server group. Use either the name or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.\+|^arn:.*:gameservergroup\[a-zA-Z0-9-\.\+]`

Required: Yes

ResumeActions

The activity to resume for this game server group.

Type: Array of strings

Array Members: Fixed number of 1 item.

Valid Values: REPLACE_INSTANCE_TYPES

Required: Yes

Response Syntax

```
{
  "GameServerGroup": {
    "AutoScalingGroupArn": "string",
    "BalancingStrategy": "string",
    "CreationTime": number,
    "GameServerGroupArn": "string",
    "GameServerGroupName": "string",
    "GameServerProtectionPolicy": "string",
    "InstanceDefinitions": [
      {
        "InstanceType": "string",
        "WeightedCapacity": "string"
      }
    ],
    "LastUpdatedTime": number,
    "RoleArn": "string",
    "Status": "string",
    "StatusReason": "string",
    "SuspendedActions": [ "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.

GameServerGroup

An object that describes the game server group resource, with the `SuspendedActions` property updated to reflect the resumed activity.

Type: [GameServerGroup](#) object

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Restart a game server group activity

This example restores instance balancing activities for the game server group.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "GameServerGroupName": "MegaFrogServers_NA",
  "ResumeActions": [ "REPLACE_INSTANCE_TYPES" ]
}
```

CLI command:

```
aws gamelift resume-game-server-group \
  --game-server-group MegaFrogServers_NA \
  --resume-actions REPLACE_INSTANCE_TYPES
```

Sample Response

```
{
  "GameServerGroup": {
    "AutoScalingGroupArn": "arn:aws:autoscaling:us-west-2:123456789012:autoScalingGroup:1111aaaa-22bb-33cc-44dd-5555eeee66ff:autoScalingGroupName/MegaFrogServers_NA",
    "BalancingStrategy": "SPOT_PREFERRED",
    "CreationTime": 1496365885.44,
    "GameServerGroupArn": "arn:aws:gamelift:us-west-2::GameServerGroup/MegaFrogServers_NA",
    "GameServerGroupName": " MegaFrogServers_NA",
    "GameServerProtectionPolicy": "NO_PROTECTION",
    "InstanceDefinitions": [
      {
        "InstanceType": "c5.2xlarge",
        "WeightedCapacity": "1"
      },
      {
        "InstanceType": "c5.4xlarge",
```

```
        "WeightedCapacity": "2"
    },
    ],
    "LastUpdatedTime": 1496365885.44,
    "RoleArn": "arn:aws:iam:123456789012::role/GameLiftGsgRole",
    "Status": "ACTIVE",
    "StatusReason": "",
    "SuspendedActions": []
}
```

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](#)

SearchGameSessions

This API works with the following fleet types: EC2, Anywhere, Container

Retrieves all active game sessions that match a set of search criteria and sorts them into a specified order.

This operation is not designed to continually track game session status because that practice can cause you to exceed your API limit and generate errors. Instead, configure an Amazon Simple Notification Service (Amazon SNS) topic to receive notifications from a matchmaker or a game session placement queue.

When searching for game sessions, you specify exactly where you want to search and provide a search filter expression, a sort expression, or both. A search request can search only one fleet, but it can search all of a fleet's locations.

This operation can be used in the following ways:

- To search all game sessions that are currently running on all locations in a fleet, provide a fleet or alias ID. This approach returns game sessions in the fleet's home Region and all remote locations that fit the search criteria.
- To search all game sessions that are currently running on a specific fleet location, provide a fleet or alias ID and a location name. For location, you can specify a fleet's home Region or any remote location.

Use the pagination parameters to retrieve results as a set of sequential pages.

If successful, a `GameSession` object is returned for each game session that matches the request. Search finds game sessions that are in `ACTIVE` status only. To retrieve information on game sessions in other statuses, use [DescribeGameSessions](#).

To set search and sort criteria, create a filter expression using the following game session attributes. For game session search examples, see the Examples section of this topic.

- **gameSessionId** -- An identifier for the game session that is unique across all regions. You must use the full ARN value.
- **gameSessionName** -- Name assigned to a game session. Game session names do not need to be unique to a game session.

- **gameSessionProperties** -- A set of key-value pairs that can store custom data in a game session. For example: {"Key": "difficulty", "Value": "novice"}. The filter expression must specify the https://docs.aws.amazon.com/gamelift/latest/apireference/API_GameProperty -- a Key and a string Value to search for the game sessions.

For example, to search for the above key-value pair, specify the following search filter:
`gameSessionProperties.difficulty = "novice"`. All game property values are searched as strings.

For examples of searching game sessions, see the ones below, and also see [Search game sessions by game property](#).

 Note

- Avoid using periods (".") in property keys if you plan to search for game sessions by properties. Property keys containing periods cannot be searched and will be filtered out from search results due to search index limitations.
- If you use SearchGameSessions API, there is a limit of 500 game property keys across all game sessions and all fleets per region. If the limit is exceeded, there will potentially be game session entries missing from SearchGameSessions API results.

- **maximumSessions** -- Maximum number of player sessions allowed for a game session.
- **creationTimeMillis** -- Value indicating when a game session was created. It is expressed in Unix time as milliseconds.
- **playerSessionCount** -- Number of players currently connected to a game session. This value changes rapidly as players join the session or drop out.
- **hasAvailablePlayerSessions** -- Boolean value indicating whether a game session has reached its maximum number of players. It is highly recommended that all search requests include this filter attribute to optimize search performance and return only sessions that players can join.

 Note

Returned values for `playerSessionCount` and `hasAvailablePlayerSessions` change quickly as players join sessions and others drop out. Results should be considered a snapshot in time. Be sure to refresh search results often, and handle sessions that fill up before a player can join.

[All APIs by task](#)

Request Syntax

```
{  
  "AliasId": "string",  
  "FilterExpression": "string",  
  "FleetId": "string",  
  "Limit": number,  
  "Location": "string",  
  "NextToken": "string",  
  "SortExpression": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[AliasId](#)

A unique identifier for the alias associated with the fleet to search for active game sessions. You can use either the alias ID or ARN value. Each request must reference either a fleet ID or alias ID, but not both.

Type: String

Pattern: `^alias-\S+|^arn:.*:alias\/alias-\S+`

Required: No

[FilterExpression](#)

String containing the search criteria for the session search. If no filter expression is included, the request returns results for all game sessions in the fleet that are in ACTIVE status.

A filter expression can contain one or multiple conditions. Each condition consists of the following:

- **Operand** -- Name of a game session attribute. Valid values are `gameSessionName`, `gameSessionId`, `gameSessionProperties`, `maximumSessions`, `creationTimeMillis`, `playerSessionCount`, `hasAvailablePlayerSessions`.
- **Comparator** -- Valid comparators are: `=`, `<>`, `<`, `>`, `<=`, `>=`.
- **Value** -- Value to be searched for. Values may be numbers, boolean values (`true/false`) or strings depending on the operand. String values are case sensitive and must be enclosed in single quotes. Special characters must be escaped. Boolean and string values can only be used with the comparators `=` and `<>`. For example, the following filter expression searches on `gameSessionName`: `"FilterExpression": "gameSessionName = 'Matt\\'s Awesome Game 1'"`.

To chain multiple conditions in a single expression, use the logical keywords `AND`, `OR`, and `NOT` and parentheses as needed. For example: `x AND y AND NOT z, NOT (x OR y)`.

Session search evaluates conditions from left to right using the following precedence rules:

1. `=`, `<>`, `<`, `>`, `<=`, `>=`
2. Parentheses
3. `NOT`
4. `AND`
5. `OR`

For example, this filter expression retrieves game sessions hosting at least ten players that have an open player slot: `"maximumSessions>=10 AND hasAvailablePlayerSessions=true"`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

FleetId

A unique identifier for the fleet to search for active game sessions. You can use either the fleet ID or ARN value. Each request must reference either a fleet ID or alias ID, but not both.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\-]+|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\-]+`

Required: No

Limit

The maximum number of results to return. Use this parameter with `NextToken` to get results as a set of sequential pages. The maximum number of results returned is 20, even if this value is not set or is set higher than 20.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

Location

A fleet location to search for game sessions. You can specify a fleet's home Region or a remote location. Use the AWS Region code format, such as `us-west-2`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-\-]+`

Required: No

NextToken

A token that indicates the start of the next sequential page of results. Use the token that is returned with a previous call to this operation. To start at the beginning of the result set, do not specify a value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

SortExpression

Instructions on how to sort the search results. If no sort expression is included, the request returns results in random order. A sort expression consists of the following elements:

- **Operand** -- Name of a game session attribute. Valid values are `gameSessionName`, `gameSessionId`, `gameSessionProperties`, `maximumSessions`, `creationTimeMillis`, `playerSessionCount`, `hasAvailablePlayerSessions`.
- **Order** -- Valid sort orders are `ASC` (ascending) and `DESC` (descending).

For example, this sort expression returns the oldest active sessions first: `"SortExpression": "creationTimeMillis ASC"`. Results with a null value for the sort operand are returned at the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "GameSessions": [
    {
      "ComputeName": "string",
      "CreationTime": number,
      "CreatorId": "string",
      "CurrentPlayerSessionCount": number,
      "DnsName": "string",
      "FleetArn": "string",
      "FleetId": "string",
      "GameProperties": [
        {
          "Key": "string",
          "Value": "string"
        }
      ],
      "GameSessionData": "string",
      "GameSessionId": "string",
```



```
    "IpAddress": "string",
    "Location": "string",
    "MatchmakerData": "string",
    "MaximumPlayerSessionCount": number,
    "Name": "string",
    "PlayerGatewayStatus": "string",
    "PlayerSessionCreationPolicy": "string",
    "Port": number,
    "Status": "string",
    "StatusReason": "string",
    "TerminationTime": 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.

GameSessions

A collection of objects containing game session properties for each session that matches the request.

Type: Array of [GameSession](#) objects

NextToken

A token that indicates where to resume retrieving results on the next call to this operation. If no token is returned, these results represent the end of the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TerminalRoutingStrategyException

The service is unable to resolve the routing for a particular alias because it has a terminal `RoutingStrategy` associated with it. The message returned in this exception is the message defined in the routing strategy itself. Such requests should only be retried if the routing strategy for the specified alias is modified.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Search game sessions

In this example, we want to find all game sessions that have at least two players already connected. We also want to filter out active game sessions that are not accepting new players.

This example illustrates a search that includes all of the fleet's locations. The results include a matching game session in the fleet's home Region (us-west-2) and another in a remote location (ca-central-1).

Sample Request

```
{
  "AliasId": "MOG-base",
  "FilterExpression": "playerSessionCount>=2 AND hasAvailablePlayerSessions=true",
  "Limit": 2
}
```

CLI syntax:

```
aws gamelift search-game-sessions --alias-id "MOG-base" --filter-expression
"playerSessionCount>=2 AND hasAvailablePlayerSessions=true" --limit 2
```

Sample Response

```
{
  "GameSessions": [
    {
      "CreationTime": 1469498468.057,
      "CurrentPlayerSessionCount": 5,
      "FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
      "GameProperties": [
        {"Key": "difficulty", "Value": "easy"},
        {"Key": "gameMap", "Value": "Snowfall"},
        {"Key": "gameMode", "Value": "Explore"}
      ],
      "GameSessionId": "arn:aws:gamelift:us-west-2::gamesession/
fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa/
gsess-4444dddd-55ee-66ff-77aa-8888bbbb99cc",
      "IpAddress": "192.0.2.0",
      "MaximumPlayerSessionCount": 10,
      "Name": "Matt's Awesome Game win123",
      "Port": "8080",
    }
  ]
}
```

```

        "Status": "ACTIVE",
        "Location": "us-west-2",
        "ComputeName": "i-1234567890abcdef0"
    },
    {
        "CreationTime": 1469498497.792,
        "CurrentPlayerSessionCount": 3,
        "FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
        "GameProperties": [
            {"Key": "difficulty", "Value": "insane"},
            {"Key": "gameMap", "Value": "Dystopia"},
            {"Key": "gameMode", "Value": "FFA"}
        ],
        "GameSessionId": "arn:aws:gamelift:ca-central-1::gamesession/
fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa/
gsess-7777dddd-55ee-66ff-44aa-8888bbbb99cc",
        "IpAddress": "192.0.2.0",
        "MaximumPlayerSessionCount": 10,
        "Name": "Matt's Awesome Game win456",
        "Port": "8080",
        "Status": "ACTIVE",
        "Location": "ca-central-1",
        "ComputeName": "i-0987654321fedcba0"
    }
]
}

```

Search and sort game sessions

In this example, we want to find all game sessions that allow 20 or more players and are currently accepting new players. We want the results to be sorted so that the newest game sessions are returned first.

This example illustrates a search of a single fleet location. The requested fleet, which resides in us-west-2, also has game sessions in remote locations, including ap-southeast-2. As shown, the results are limited to the requested fleet location.

Sample Request

```

{
    "FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
    "Location": "ap-southeast-2",
    "FilterExpression": "maximumSessions>=20 AND hasAvailablePlayerSessions=true",

```

```
"SortExpression": "creationTimeMillis DESC"
"Limit": 2
}
```

CLI syntax:

```
aws gamelift search-game-sessions --fleet-id
"fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa" --location "ap-southeast-2" --filter-
expression "maximumSessions=20 AND hasAvailablePlayerSessions=true" --sort-expression
"creationTimeMillis DESC"
```

Sample Response

```
{
  "GameSessions": [
    {
      "CreationTime": 1469498497.792,
      "CurrentPlayerSessionCount": 3,
      "FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
      "GameProperties": [
        {"Key": "difficulty", "Value": "hard"},
        {"Key": "gameMap", "Value": "Dystopia"},
        {"Key": "gameMode", "Value": "Brawl"}
      ],
      "GameSessionId": "arn:aws:gamelift:ap-southeast-2::gamesession/
fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa/
gssess-7777dddd-55ee-66ff-44aa-8888bbbb99cc",
      "IpAddress": "192.0.2.0",
      "MaximumPlayerSessionCount": 20,
      "Name": "Matt's Awesome Game win456",
      "Port": "8080",
      "Status": "ACTIVE",
      "Location": "ap-southeast-2",
      "ComputeName": "i-abcdef1234567890a"
    },
    {
      "CreationTime": 1469498468.057,
      "CurrentPlayerSessionCount": ,
      "FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
      "GameProperties": [
        {"Key": "difficulty", "Value": "easy"},
        {"Key": "gameMap", "Value": "Snowfall"},
        {"Key": "gameMode", "Value": "Explore"}
      ]
    }
  ]
}
```

```
    ],
    "GameSessionId": "arn:aws:gamelift:ap-southeast-2::gamesession/
fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa/
gsess-4444dddd-55ee-66ff-77aa-8888bbbb99cc",
    "IpAddress": "192.0.2.0",
    "MaximumPlayerSessionCount": 50,
    "Name": "Matt's Awesome Game win123",
    "Port": "8080",
    "Status": "ACTIVE",
    "Location": "ap-southeast-2",
    "ComputeName": "i-fedcba0987654321f"
  }
]
}
```

Search game sessions by custom game properties

This example searches for game sessions based on game map and game mode information, which is stored as key-value pairs in the `GameProperties` of a `GameSession`. In this example, we want to find all game sessions where `gameMode` is `Ffa` (free-for-all), and `gameMap` is either `"Suzuka"` or `"Silverstone"`. We are sorting our results by `gameSessionProperties.difficulty` (with possible values of `"novice"`, `"easy"`, `"normal"`, `"hard"`, or `"insane"`). Note: `Value` is evaluated as a string, so the sorted results will be listed by the alphabetic order of the difficulty values.

Sample Request

```
{
  "FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
  "Location": "us-west-2",
  "FilterExpression": "gameSessionProperties.gameMode = 'Ffa' AND
gameSessionProperties.gameMap = 'Suzuka' OR gameSessionProperties.gameMap =
'Silverstone'",
  "SortExpression": "gameSessionProperties.difficulty ASC"
  "Limit": 2
}
```

CLI syntax:

```
aws gamelift search-game-sessions --fleet-id "9999ffff-88ee-77dd-66cc-5555bbbb44aa"
--filter-expression "gameSessionProperties.gameMode = 'Ffa' AND
gameSessionProperties.gameMap = 'Suzuka' OR gameSessionProperties.gameMap =
'Silverstone'" --sort-expression "gameSessionProperties.difficulty DESC"
```

Sample Response

```
{
  "GameSessions": [
    {
      "CreationTime": 1469498468.057,
      "CurrentPlayerSessionCount": 5,
      "FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
      "GameProperties": [
        {"Key": "difficulty", "Value": "easy"},
        {"Key": "gameMap", "Value": "Suzuka"},
        {"Key": "gameMode", "Value": "Ffa"}
      ],
      "GameSessionId": "arn:aws:gamelift:us-west-2::gamesession/
fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa/
gsess-4444dddd-55ee-66ff-77aa-8888bbbb99cc",
      "IpAddress": "192.0.2.0",
      "MaximumPlayerSessionCount": 10,
      "Name": "Matt's Awesome Game win123",
      "Port": "8080",
      "Status": "ACTIVE",
      "Location": "us-west-2",
      "ComputeName": "i-1111222233334444a"
    },
    {
      "CreationTime": 1469498497.792,
      "CurrentPlayerSessionCount": 3,
      "FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
      "GameProperties": [
        {"Key": "difficulty", "Value": "normal"},
        {"Key": "gameMap", "Value": "Silverstone"},
        {"Key": "gameMode", "Value": "Ffa"}
      ],
      "GameSessionId": "arn:aws:gamelift:us-west-2::gamesession/
fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa/
gsess-7777dddd-55ee-66ff-44aa-8888bbbb99cc",
      "IpAddress": "192.0.2.0",
      "MaximumPlayerSessionCount": 10,
      "Name": "Matt's Awesome Game win456",
      "Port": "8080",
      "Status": "ACTIVE",
      "Location": "us-west-2",
      "ComputeName": "i-5555666677778888b"
    }
  ]
}
```

```
]
}
```

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](#)

StartFleetActions

This API works with the following fleet types: EC2, Container

Resumes certain types of activity on fleet instances that were suspended with [StopFleetActions](#). For multi-location fleets, fleet actions are managed separately for each location. Currently, this operation is used to restart a fleet's auto-scaling activity.

This operation can be used in the following ways:

- To restart actions on instances in the fleet's home Region, provide a fleet ID and the type of actions to resume.
- To restart actions on instances in one of the fleet's remote locations, provide a fleet ID, a location name, and the type of actions to resume.

If successful, Amazon GameLift Servers once again initiates scaling events as triggered by the fleet's scaling policies. If actions on the fleet location were never stopped, this operation will have no effect.

Learn more

[Setting up Amazon GameLift Servers fleets](#)

Request Syntax

```
{  
  "Actions": [ "string" ],  
  "FleetId": "string",  
  "Location": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Actions

List of actions to restart on the fleet.

Type: Array of strings

Array Members: Fixed number of 1 item.

Valid Values: AUTO_SCALING

Required: Yes

FleetId

A unique identifier for the fleet to restart actions on. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

Location

The fleet location to restart fleet actions for. Specify a location in the form of an AWS Region code, such as `us-west-2`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-\_]+\$`

Required: No

Response Syntax

```
{
```

```
"FleetArn": "string",  
"FleetId": "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.

FleetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-]+`

FleetId

A unique identifier for the fleet to restart actions on.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-]+`

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Restart automatic scaling activity for a fleet

In this example, we want to resume the use a fleet's scaling policies that had be previously suspended. Once started, the scaling policies immediately begin tracking their respective metrics.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "Actions": ["AUTO_SCALING"]
}
```

Sample Response

```
{
  "FleetArn": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "FleetId": "arn:aws:gamelift:us-west-2::fleet/
fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
}
```

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](#)

StartGameSessionPlacement

This API works with the following fleet types: EC2, Anywhere, Container

Makes a request to start a new game session using a game session queue. When processing a placement request, Amazon GameLift Servers looks for the best possible available resource to host the game session, based on how the queue is configured to prioritize factors such as resource cost, latency, and location. After selecting an available resource, Amazon GameLift Servers prompts the resource to start a game session. A placement request can include a list of players to create a set of player sessions. The request can also include information to pass to the new game session, such as to specify a game map or other options.

Request options

Use this operation to make the following types of requests.

- Request a placement using the queue's default prioritization process (see the default prioritization described in [PriorityConfiguration](#)). Include these required parameters:
 - `GameSessionQueueName`
 - `MaximumPlayerSessionCount`
 - `PlacementID`
- Request a placement and prioritize based on latency. Include these parameters:
 - Required parameters `GameSessionQueueName`, `MaximumPlayerSessionCount`, `PlacementID`.
 - `PlayerLatencies`. Include a set of latency values for destinations in the queue. When a request includes latency data, Amazon GameLift Servers automatically reorder the queue's locations priority list based on lowest available latency values. If a request includes latency data for multiple players, Amazon GameLift Servers calculates each location's average latency for all players and reorders to find the lowest latency across all players. To obtain accurate latency measurements, use Amazon GameLift Servers's UDP ping beacons. These endpoints allow you to measure actual UDP network latency between player devices and potential hosting locations, resulting in more accurate placement decisions than using ICMP pings. For more information on using UDP ping beacons to measure latency, refer to [UDP ping beacons](#) in the Amazon GameLift Servers Developer Guide.
 - Don't include `PriorityConfigurationOverride`.

- Prioritize based on a custom list of locations. If you're using a queue that's configured to prioritize location first (see [PriorityConfiguration](#) for game session queues), you can optionally use the *PriorityConfigurationOverride* parameter to substitute a different location priority list for this placement request. Amazon GameLift Servers searches each location on the priority override list to find an available hosting resource for the new game session. Specify a fallback strategy to use in the event that Amazon GameLift Servers fails to place the game session in any of the locations on the override list.
- Request a placement and prioritized based on a custom list of locations.
- You can request new player sessions for a group of players. Include the *DesiredPlayerSessions* parameter and include at minimum a unique player ID for each. You can also include player-specific data to pass to the new game session.

Result

If successful, this operation generates a new game session placement request and adds it to the game session queue for processing. You can track the status of individual placement requests by calling [DescribeGameSessionPlacement](#) or by monitoring queue notifications. When the request status is FULFILLED, a new game session has started and the placement request is updated with connection information for the game session (IP address and port). If the request included player session data, Amazon GameLift Servers creates a player session for each player ID in the request.

The request results in a `InvalidRequestException` in the following situations:

- If the request includes both *PlayerLatencies* and *PriorityConfigurationOverride* parameters.
- If the request includes the *PriorityConfigurationOverride* parameter and specifies a queue that doesn't prioritize locations.

Amazon GameLift Servers continues to retry each placement request until it reaches the queue's timeout setting. If a request times out, you can resubmit the request to the same queue or try a different queue.

Request Syntax

```
{
  "DesiredPlayerSessions": [
    {
      "PlayerData": "string",
```

```
    "PlayerId": "string"
  },
],
"GameProperties": [
  {
    "Key": "string",
    "Value": "string"
  }
],
"GameSessionData": "string",
"GameSessionName": "string",
"GameSessionQueueName": "string",
"MaximumPlayerSessionCount": number,
"PlacementId": "string",
"PlayerLatencies": [
  {
    "LatencyInMilliseconds": number,
    "PlayerId": "string",
    "RegionIdentifier": "string"
  }
],
"PriorityConfigurationOverride": {
  "LocationOrder": [ "string" ],
  "PlacementFallbackStrategy": "string"
}
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

GameSessionQueueName

Name of the queue to use to place the new game session. You can use either the queue name or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-]+\^arn:.*:gamesessionqueue\[a-zA-Z0-9-]+\`

Required: Yes

MaximumPlayerSessionCount

The maximum number of players that can be connected simultaneously to the game session.

Type: Integer

Valid Range: Minimum value of 0.

Required: Yes

PlacementId

A unique identifier to assign to the new game session placement. This value is developer-defined. The value must be unique across all Regions and cannot be reused.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 48.

Pattern: `[a-zA-Z0-9-]+\`

Required: Yes

DesiredPlayerSessions

Set of information on each player to create a player session for.

Type: Array of [DesiredPlayerSession](#) objects

Required: No

GameProperties

A set of key-value pairs that can store custom data in a game session. For example: `{"Key": "difficulty", "Value": "novice"}`.

Note

- Avoid using periods (".") in property keys if you plan to search for game sessions by properties. Property keys containing periods cannot be searched and will be filtered out from search results due to search index limitations.
- If you use SearchGameSessions API, there is a limit of 500 game property keys across all game sessions and all fleets per region. If the limit is exceeded, there will potentially be game session entries missing from SearchGameSessions API results.

Type: Array of [GameProperty](#) objects

Array Members: Maximum number of 16 items.

Required: No

GameSessionData

A set of custom game session properties, formatted as a single string value. This data is passed to a game server process with a request to start a new game session. For more information, see [Start a game session](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 262144.

Required: No

GameSessionName

A descriptive label that is associated with a game session. Session names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

PlayerLatencies

A set of values, expressed in milliseconds, that indicates the amount of latency that a player experiences when connected to a fleet location (AWS Regions or custom locations for Amazon

GameLift Servers Anywhere fleets). This information is used to try to place the new game session where it can offer the best possible gameplay experience for the players. To obtain accurate latency measurements, use Amazon GameLift Servers's UDP ping beacons. These endpoints allow you to measure actual UDP network latency between player devices and potential hosting locations, resulting in more accurate placement decisions than using ICMP pings. For more information on using UDP ping beacons to measure latency, refer to [UDP ping beacons](#) in the Amazon GameLift Servers Developer Guide.

Type: Array of [PlayerLatency](#) objects

Required: No

[PriorityConfigurationOverride](#)

A prioritized list of locations to use for the game session placement and instructions on how to use it. This list overrides a queue's prioritized location list for this game session placement request only. You can include AWS Regions, local zones, and custom locations (for Anywhere fleets). You can choose to limit placements to locations on the override list only, or you can prioritize locations on the override list first and then fall back to the queue's other locations if needed. Choose a fallback strategy to use in the event that Amazon GameLift Servers fails to place a game session in any of the locations on the priority override list.

Type: [PriorityConfigurationOverride](#) object

Required: No

Response Syntax

```
{
  "GameSessionPlacement": {
    "DnsName": "string",
    "EndTime": number,
    "GameProperties": [
      {
        "Key": "string",
        "Value": "string"
      }
    ],
    "GameSessionArn": "string",
    "GameSessionData": "string",
    "GameSessionId": "string",
```

```
"GameSessionName": "string",
"GameSessionQueueName": "string",
"GameSessionRegion": "string",
"IpAddress": "string",
"MatchmakerData": "string",
"MaximumPlayerSessionCount": number,
"PlacedPlayerSessions": [
  {
    "PlayerId": "string",
    "PlayerSessionId": "string"
  }
],
"PlacementId": "string",
"PlayerGatewayStatus": "string",
"PlayerLatencies": [
  {
    "LatencyInMilliseconds": number,
    "PlayerId": "string",
    "RegionIdentifier": "string"
  }
],
"Port": number,
"PriorityConfigurationOverride": {
  "LocationOrder": [ "string" ],
  "PlacementFallbackStrategy": "string"
},
"StartTime": number,
"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.

GameSessionPlacement

Object that describes the newly created game session placement. This object includes all the information provided in the request, as well as start/end time stamps and placement status.

Type: [GameSessionPlacement](#) object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Request a new game session placement with player latency data

This example starts a new game session placement. The request calls for player sessions for two players, and provides each player's latency data for two Regions.

Amazon GameLift Servers uses the latency data provided to determine what order to use when looking for a fleet to host the new game session. It does this by calculating the average player

latency for each Region and ordering the queue's destinations starting with the lowest average latency. If the queue "matchmaker-queue" has a latency policy, however, things may change. For example, let's say matchmaker-queue has a policy that caps latency at 130 milliseconds for 60 seconds, followed by no cap. In this scenario, using the sample request below, the following sequence plays out:

1. Amazon GameLift Servers calculates average latency for each Region: us-east-1 = 110 and us-west-2 = 100.
2. Amazon GameLift Servers reorders the queue's destinations based on lowest average latency, and prioritizes destinations in Region us-west-2.
3. The queue has a latency cap of 130 ms in force for the first 60 seconds of a placement. Amazon GameLift Servers looks for any individual latency values that are greater than 130 ms. There is one: Player 2 reports a 150 ms latency when connected to Region us-west-2. As a result, Amazon GameLift Servers temporarily drops all us-west-2 fleets as valid destinations.
4. Amazon GameLift Servers tries to place the new game session on fleets in Region us-east-1, followed by fleets in Regions with no latency information (if any). If available resources are found, the game session is placed and the request fulfilled.
5. If no available resources are found, Amazon GameLift Servers starts a new round of placement attempts, restarting at step 3. If 60 seconds have passed and the latency policy is no longer in force, then fleets in Region us-west-2 are once more valid destinations -- and are preferred based on their low average latency.
6. Amazon GameLift Servers continues to attempt to place the new game session until it is successful or until the queue's timeout limit is reached.

Sample Request

```
{
  "DesiredPlayerSessions": [
    { "PlayerData": "level:10", "PlayerId": "player1" },
    { "PlayerData": "level:11", "PlayerId": "player2" }
  ],
  "GameProperties": [
    { "Key": "map", "Value": "winter" }
  ],
  "GameSessionName": "matchmaker-1234567890",
  "GameSessionQueueName": "matchmaker-queue",
  "MaximumPlayerSessionCount": 4,
```

```
"PlacementId": "Place-12345",
"PlayerLatencies": [
  { "LatencyInMilliseconds": 100, "PlayerId": "player1", "RegionIdentifier": "us-east-1" },
  { "LatencyInMilliseconds": 50, "PlayerId": "player1", "RegionIdentifier": "us-west-2" },
  { "LatencyInMilliseconds": 120, "PlayerId": "player2", "RegionIdentifier": "us-east-1" },
  { "LatencyInMilliseconds": 150, "PlayerId": "player2", "RegionIdentifier": "us-west-2" }
]
```

Request a game session placement with priority configuration override

This example requests a game session placement using an alternate list of location priorities.

Amazon GameLift Servers uses the priority configuration override to change how it searches a queue's destinations when looking for a resource to host the new game session. The queue "custom-locations" has a `PriorityConfiguration` setting that prioritizes locations first and includes a prioritized location list. For this one placement request only, that list is overridden with the list in *PriorityConfigurationOverride*. The following sequence plays out:

1. Amazon GameLift Servers looks for the first location on the override list (us-west-2) in each destination in the queue, using the queue's destination priority order. In each fleet location, it searches for an available hosting resource. When one is found, the game session is placed and the request fulfilled.
2. If no available resources are found in the first location, Amazon GameLift Servers looks at destinations for the next location on the override list. and so on. As soon as an available resource is found, the game session is placed and the search stops.
3. If Amazon GameLift Servers searches all locations on the override list and finds no available resources, the first placement pass has failed. Because this placement request specified no fallback strategy, Amazon GameLift Servers doesn't look at any additional queue locations. The placement request remains in the queue and is processed again in turn until a placement succeeds or until the queue's timeout limit is reached.

Sample Request

```
{
  "GameSessionName": "customlocations-1234567890",
  "GameSessionQueueName": "custom-locations",
  "MaximumPlayerSessionCount": 4,
  "PlacementId": "Place-12345",
  "PriorityConfigurationOverride": {
    "LocationOrder": ["us-west-2", "us-east-1"],
    "PlacementFallbackStrategy": "NONE"
  }
}
```

See Also

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

- [AWS Command Line Interface 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](#)

StartMatchBackfill

This API works with the following fleet types: EC2, Anywhere, Container

Finds new players to fill open slots in currently running game sessions. The backfill match process is essentially identical to the process of forming new matches. Backfill requests use the same matchmaker that was used to make the original match, and they provide matchmaking data for all players currently in the game session. FlexMatch uses this information to select new players so that backfilled match continues to meet the original match requirements.

When including player latency data in backfill requests, you can use Amazon GameLift Servers's UDP ping beacons to obtain accurate measurements of network latency between player devices and the game session's hosting location. For more information on using UDP ping beacons to measure latency, refer to [UDP ping beacons](#) in the Amazon GameLift Servers Developer Guide.

When using FlexMatch with Amazon GameLift Servers managed hosting, you can request a backfill match from a client service by calling this operation with a `GameSessions` ID. You also have the option of making backfill requests directly from your game server. In response to a request, FlexMatch creates player sessions for the new players, updates the `GameSession` resource, and sends updated matchmaking data to the game server. You can request a backfill match at any point after a game session is started. Each game session can have only one active backfill request at a time; a subsequent request automatically replaces the earlier request.

When using FlexMatch as a standalone component, request a backfill match by calling this operation without a game session identifier. As with newly formed matches, matchmaking results are returned in a matchmaking event so that your game can update the game session that is being backfilled.

To request a backfill match, specify a unique ticket ID, the original matchmaking configuration, and matchmaking data for all current players in the game session being backfilled. Optionally, specify the `GameSession` ARN. If successful, a match backfill ticket is created and returned with status set to `QUEUED`. Track the status of backfill tickets using the same method for tracking tickets for new matches.

Only game sessions created by FlexMatch are supported for match backfill.

Learn more

[Backfill existing games with FlexMatch](#)

[Matchmaking events](#) (reference)

[How Amazon GameLift Servers FlexMatch works](#)

Request Syntax

```
{
  "ConfigurationName": "string",
  "GameSessionArn": "string",
  "Players": [
    {
      "LatencyInMs": {
        "string" : number
      },
      "PlayerAttributes": {
        "string" : {
          "N": number,
          "S": "string",
          "SDM": {
            "string" : number
          },
          "SL": [ "string" ]
        }
      },
      "PlayerId": "string",
      "Team": "string"
    }
  ],
  "TicketId": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

ConfigurationName

Name of the matchmaker to use for this request. You can use either the configuration name or ARN value. The ARN of the matchmaker that was used with the original game session is listed in the `GameSession` object, `MatchmakerData` property.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]*|^arn:.*:matchmakingconfiguration\[a-zA-Z0-9-\.]*`

Required: Yes

Players

Match information on all players that are currently assigned to the game session. This information is used by the matchmaker to find new players and add them to the existing game.

You can include up to 199 `Players` in a `StartMatchBackfill` request.

- `PlayerID`, `PlayerAttributes`, `Team` -- This information is maintained in the `GameSession` object, `MatchmakerData` property, for all players who are currently assigned to the game session. The matchmaker data is in JSON syntax, formatted as a string. For more details, see [Match Data](#).

The backfill request must specify the team membership for every player. Do not specify team if you are not using backfill.

- `LatencyInMs` -- If the matchmaker uses player latency, include a latency value, in milliseconds, for the Region that the game session is currently in. Do not include latency values for any other Region.

Type: Array of [Player](#) objects

Required: Yes

GameSessionArn

An identifier for the game session that is unique across all regions. The value is always a full ARN in the following format: For Home Region game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>`. For Remote Location game session - `arn:aws:gamelift:<home_region>::gamesession/`

<fleet ID>/<location>/<ID string>. When using FlexMatch as a standalone matchmaking solution, this parameter is not needed.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: [a-zA-Z0-9:/-]+

Required: No

TicketId

A unique identifier for a matchmaking ticket. If no ticket ID is specified here, Amazon GameLift Servers will generate one in the form of a UUID. Use this identifier to track the match backfill ticket status and retrieve match results.

Type: String

Length Constraints: Maximum length of 128.

Pattern: [a-zA-Z0-9-\.]*

Required: No

Response Syntax

```
{
  "MatchmakingTicket": {
    "ConfigurationArn": "string",
    "ConfigurationName": "string",
    "EndTime": number,
    "EstimatedWaitTime": number,
    "GameSessionConnectionInfo": {
      "DnsName": "string",
      "GameSessionArn": "string",
      "IpAddress": "string",
      "MatchedPlayerSessions": [
        {
          "PlayerId": "string",
          "PlayerSessionId": "string"
        }
      ]
    }
  }
}
```

```
    ],  
    "PlayerGatewayStatus": "string",  
    "Port": number  
  },  
  "Players": [  
    {  
      "LatencyInMs": {  
        "string" : number  
      },  
      "PlayerAttributes": {  
        "string" : {  
          "N": number,  
          "S": "string",  
          "SDM": {  
            "string" : number  
          },  
          "SL": [ "string" ]  
        }  
      },  
      "PlayerId": "string",  
      "Team": "string"  
    }  
  ],  
  "StartTime": number,  
  "Status": "string",  
  "StatusMessage": "string",  
  "StatusReason": "string",  
  "TicketId": "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.

MatchmakingTicket

Ticket representing the backfill matchmaking request. This object includes the information in the request, ticket status, and match results as generated during the matchmaking process.

Type: [MatchmakingTicket](#) object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

StartMatchmaking

This API works with the following fleet types: EC2, Anywhere, Container

Uses FlexMatch to create a game match for a group of players based on custom matchmaking rules. With games that use Amazon GameLift Servers managed hosting, this operation also triggers Amazon GameLift Servers to find hosting resources and start a new game session for the new match. Each matchmaking request includes information on one or more players and specifies the FlexMatch matchmaker to use. When a request is for multiple players, FlexMatch attempts to build a match that includes all players in the request, placing them in the same team and finding additional players as needed to fill the match.

To start matchmaking, provide a unique ticket ID, specify a matchmaking configuration, and include the players to be matched. You must also include any player attributes that are required by the matchmaking configuration's rule set. If successful, a matchmaking ticket is returned with status set to QUEUED.

Track matchmaking events to respond as needed and acquire game session connection information for successfully completed matches. Ticket status updates are tracked using event notification through Amazon Simple Notification Service, which is defined in the matchmaking configuration.

Learn more

[Add FlexMatch to a game client](#)

[Set Up FlexMatch event notification](#)

[How Amazon GameLift Servers FlexMatch works](#)

Request Syntax

```
{
  "ConfigurationName": "string",
  "Players": [
    {
      "LatencyInMs": {
        "string" : number
      },
      "PlayerAttributes": {
        "string" : {
          "N": number,
          "S": "string",
```



```
    "SDM": {
      "string" : number
    },
    "SL": [ "string" ]
  },
  "PlayerId": "string",
  "Team": "string"
},
"TicketId": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

ConfigurationName

Name of the matchmaking configuration to use for this request. Matchmaking configurations must exist in the same Region as this request. You can use either the configuration name or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]*|^arn:.*:matchmakingconfiguration\[a-zA-Z0-9-\.]*`

Required: Yes

Players

Information on each player to be matched. This information must include a player ID, and may contain player attributes and latency data to be used in the matchmaking process. After a successful match, `Player` objects contain the name of the team the player is assigned to.

You can include up to 10 Players in a StartMatchmaking request.

For accurate player latency data, use Amazon GameLift Servers's UDP ping beacons to measure network latency between player devices and potential hosting locations. UDP ping beacons provide more accurate measurements than ICMP pings because they use the same protocol (UDP) that most game servers use. For more information on using UDP ping beacons to measure latency, refer to [UDP ping beacons](#) in the Amazon GameLift Servers Developer Guide.

Type: Array of [Player](#) objects

Required: Yes

TicketId

A unique identifier for a matchmaking ticket. If no ticket ID is specified here, Amazon GameLift Servers will generate one in the form of a UUID. Use this identifier to track the matchmaking ticket status and retrieve match results.

Type: String

Length Constraints: Maximum length of 128.

Pattern: [a-zA-Z0-9-\.]*

Required: No

Response Syntax

```
{
  "MatchmakingTicket": {
    "ConfigurationArn": "string",
    "ConfigurationName": "string",
    "EndTime": number,
    "EstimatedWaitTime": number,
    "GameSessionConnectionInfo": {
      "DnsName": "string",
      "GameSessionArn": "string",
      "IpAddress": "string",
      "MatchedPlayerSessions": [
        {
          "PlayerId": "string",
          "PlayerSessionId": "string"
        }
      ]
    }
  }
}
```

```
    ],
    "PlayerGatewayStatus": "string",
    "Port": number
  },
  "Players": [
    {
      "LatencyInMs": {
        "string": number
      },
      "PlayerAttributes": {
        "string": {
          "N": number,
          "S": "string",
          "SDM": {
            "string": number
          },
          "SL": [ "string" ]
        }
      },
      "PlayerId": "string",
      "Team": "string"
    }
  ],
  "StartTime": number,
  "Status": "string",
  "StatusMessage": "string",
  "StatusReason": "string",
  "TicketId": "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.

MatchmakingTicket

Ticket representing the matchmaking request. This object include the information included in the request, ticket status, and match results as generated during the matchmaking process.

Type: MatchmakingTicket object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

StopFleetActions

This API works with the following fleet types: EC2, Container

Suspends certain types of activity in a fleet location. Currently, this operation is used to stop auto-scaling activity. For multi-location fleets, fleet actions are managed separately for each location.

Stopping fleet actions has several potential purposes. It allows you to temporarily stop auto-scaling activity but retain your scaling policies for use in the future. For multi-location fleets, you can set up fleet-wide auto-scaling, and then opt out of it for certain locations.

This operation can be used in the following ways:

- To stop actions on instances in the fleet's home Region, provide a fleet ID and the type of actions to suspend.
- To stop actions on instances in one of the fleet's remote locations, provide a fleet ID, a location name, and the type of actions to suspend.

If successful, Amazon GameLift Servers no longer initiates scaling events except in response to manual changes using [UpdateFleetCapacity](#). To restart fleet actions again, call [StartFleetActions](#).

Learn more

[Setting up Amazon GameLift Servers Fleets](#)

Request Syntax

```
{  
  "Actions": [ "string" ],  
  "FleetId": "string",  
  "Location": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Actions

List of actions to suspend on the fleet.

Type: Array of strings

Array Members: Fixed number of 1 item.

Valid Values: AUTO_SCALING

Required: Yes

FleetId

A unique identifier for the fleet to stop actions on. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

Location

The fleet location to stop fleet actions for. Specify a location in the form of an AWS Region code, such as `us-west-2`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-\_]+\$`

Required: No

Response Syntax

```
{  
  "FleetArn": "string",  
  "FleetId": "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.

[FleetArn](#)

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\]\+$`

[FleetId](#)

A unique identifier for the fleet to stop actions on.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\]\+`

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Stop auto-scaling activity in a fleet's remote location

In this example, we want to suspend all auto-scaling activity for a fleet's capacity in the remote location of `ca-central-1`. The fleet in question resides in the AWS Region `us-west-2` (which is implicitly defined as our default Region). Once the policies are suspended, fleet capacity remains at the last "desired instances" setting unless fleet capacity is manually adjusted.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "Location": "ca-central-1",
  "Actions": ["AUTO_SCALING"]
}
```

Sample Response

```
{
  "FleetArn": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "FleetId": "arn:aws:gamelift:us-west-2::fleet/fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
}
```

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](#)

StopGameSessionPlacement

This API works with the following fleet types: EC2, Anywhere, Container

Cancels a game session placement that's in PENDING status. To stop a placement, provide the placement ID value.

Results

If successful, this operation removes the placement request from the queue and moves the GameSessionPlacement to CANCELLED status.

This operation results in an `InvalidRequestException` (400) error if a game session has already been created for this placement. You can clean up an unneeded game session by calling [TerminateGameSession](#).

Request Syntax

```
{
  "PlacementId": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[PlacementId](#)

A unique identifier for a game session placement to stop.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 48.

Pattern: [a-zA-Z0-9-]+

Required: Yes

Response Syntax

```
{
  "GameSessionPlacement": {
    "DnsName": "string",
    "EndTime": number,
    "GameProperties": [
      {
        "Key": "string",
        "Value": "string"
      }
    ],
    "GameSessionArn": "string",
    "GameSessionData": "string",
    "GameSessionId": "string",
    "GameSessionName": "string",
    "GameSessionQueueName": "string",
    "GameSessionRegion": "string",
    "IpAddress": "string",
    "MatchmakerData": "string",
    "MaximumPlayerSessionCount": number,
    "PlacedPlayerSessions": [
      {
        "PlayerId": "string",
        "PlayerSessionId": "string"
      }
    ],
    "PlacementId": "string",
    "PlayerGatewayStatus": "string",
    "PlayerLatencies": [
      {
        "LatencyInMilliseconds": number,
        "PlayerId": "string",
        "RegionIdentifier": "string"
      }
    ],
    "Port": number,
    "PriorityConfigurationOverride": {
      "LocationOrder": [ "string" ],

```

```
    "PlacementFallbackStrategy": "string"  
  },  
  "StartTime": number,  
  "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.

GameSessionPlacement

Object that describes the canceled game session placement, with CANCELLED status and an end time stamp.

Type: [GameSessionPlacement](#) object

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

StopMatchmaking

This API works with the following fleet types: EC2, Anywhere, Container

Cancels a matchmaking ticket or match backfill ticket that is currently being processed. To stop the matchmaking operation, specify the ticket ID. If successful, work on the ticket is stopped, and the ticket status is changed to CANCELLED.

This call is also used to turn off automatic backfill for an individual game session. This is for game sessions that are created with a matchmaking configuration that has automatic backfill enabled. The ticket ID is included in the MatchmakerData of an updated game session object, which is provided to the game server.

Note

If the operation is successful, the service sends back an empty JSON struct with the HTTP 200 response (not an empty HTTP body).

Learn more

[Add FlexMatch to a game client](#)

Request Syntax

```
{
  "TicketId": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

TicketId

A unique identifier for a matchmaking ticket.

Type: String

Length Constraints: Maximum length of 128.

Pattern: [a-zA-Z0-9-\.]*

Required: Yes

Response Elements

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

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

SuspendGameServerGroup

This API works with the following fleet types: EC2 (FleetIQ)

Temporarily stops activity on a game server group without terminating instances or the game server group. You can restart activity by calling [ResumeGameServerGroup](#). You can suspend the following activity:

- **Instance type replacement** - This activity evaluates the current game hosting viability of all Spot instance types that are defined for the game server group. It updates the Auto Scaling group to remove nonviable Spot Instance types, which have a higher chance of game server interruptions. It then balances capacity across the remaining viable Spot Instance types. When this activity is suspended, the Auto Scaling group continues with its current balance, regardless of viability. Instance protection, utilization metrics, and capacity scaling activities continue to be active.

To suspend activity, specify a game server group ARN and the type of activity to be suspended. If successful, a GameServerGroup object is returned showing that the activity is listed in `SuspendedActions`.

Learn more

[Amazon GameLift Servers FleetIQ Guide](#)

Request Syntax

```
{  
  "GameServerGroupName": "string",  
  "SuspendActions": [ "string" ]  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

GameServerGroupName

A unique identifier for the game server group. Use either the name or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]+\|^arn:.*:gameservergroup\/[a-zA-Z0-9-\.]+`

Required: Yes

SuspendActions

The activity to suspend for this game server group.

Type: Array of strings

Array Members: Fixed number of 1 item.

Valid Values: REPLACE_INSTANCE_TYPES

Required: Yes

Response Syntax

```
{
  "GameServerGroup": {
    "AutoScalingGroupArn": "string",
    "BalancingStrategy": "string",
    "CreationTime": number,
    "GameServerGroupArn": "string",
    "GameServerGroupName": "string",
    "GameServerProtectionPolicy": "string",
    "InstanceDefinitions": [
      {
        "InstanceType": "string",
        "WeightedCapacity": "string"
      }
    ],
    "LastUpdatedTime": number,
    "RoleArn": "string",
    "Status": "string",
    "StatusReason": "string",
```

```
    "SuspendedActions": [ "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.

GameServerGroup

An object that describes the game server group resource, with the `SuspendedActions` property updated to reflect the suspended activity.

Type: GameServerGroup object

Errors

For information about the errors that are common to all actions, see Common Error Types.

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Suspend a game server group activity

This example suspends instance balancing activities for the game server group.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "GameServerGroupName": "MegaFrogServers_NA",
  "SuspendActions": [ "REPLACE_INSTANCE_TYPES" ]
}
```

CLI command:

```
aws gamelift suspend-game-server-group \
  --game-server-group MegaFrogServers_NA \
  --suspend-actions REPLACE_INSTANCE_TYPES
```

Sample Response

```
{
  "GameServerGroup": {
    "AutoScalingGroupArn": "arn:aws:autoscaling:us-west-2:123456789012:autoScalingGroup:1111aaaa-22bb-33cc-44dd-5555eeee66ff:autoScalingGroupName/MegaFrogServers_NA",
    "BalancingStrategy": "SPOT_PREFERRED",
    "CreationTime": 1496365885.44,
    "GameServerGroupArn": "arn:aws:gamelift:us-west-2::GameServerGroup/MegaFrogServers_NA",
    "GameServerGroupName": " MegaFrogServers_NA",
    "GameServerProtectionPolicy": "NO_PROTECTION",
    "InstanceDefinitions": [
      {
        "InstanceType": "c5.2xlarge",
        "WeightedCapacity": "1"
      },
    ],
  },
}
```

```
{
  {
    "InstanceType": "c5.4xlarge",
    "WeightedCapacity": "2"
  },
  "LastUpdatedTime": 1496365885.44,
  "RoleArn": "arn:aws:iam:123456789012::role/GameLiftGsgRole",
  "Status": "ACTIVE",
  "StatusReason": "",
  "SuspendedActions": [REPLACE_INSTANCE_TYPES]
}
```

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

This API works with the following fleet types: EC2, Anywhere, Container

Assigns a tag to an Amazon GameLift Servers resource. You can use tags to organize resources, create IAM permissions policies to manage access to groups of resources, customize AWS cost breakdowns, and more. This operation handles the permissions necessary to manage tags for Amazon GameLift Servers resources that support tagging.

To add a tag to a resource, specify the unique ARN value for the resource and provide a tag list containing one or more tags. The operation succeeds even if the list includes tags that are already assigned to the resource.

Learn more

[Tagging AWS Resources](#) in the *AWS General Reference*

[AWS Tagging Strategies](#)

Related actions

[All APIs by task](#)

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

ResourceARN

The Amazon Resource Name ([ARN](#)) that uniquely identifies the Amazon GameLift Servers resource that you want to assign tags to. Amazon GameLift Servers includes resource ARNs in the data object for the resource. You can retrieve the ARN by calling a `List` or `Describe` operation for the resource type.

Type: String

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

Required: Yes

Tags

A list of one or more tags to assign to the specified Amazon GameLift Servers resource. Tags are developer-defined and structured as key-value pairs. The maximum tag limit may be lower than stated. See [Tagging AWS Resources](#) for tagging limits.

Type: Array of [Tag](#) objects

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

Required: Yes

Response Elements

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

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

TerminateGameSession

This API works with the following fleet types: EC2, Anywhere, Container

Ends a game session that's currently in progress. Use this action to terminate any game session that isn't in ERROR status. Terminating a game session is the most efficient way to free up a server process when it's hosting a game session that's in a bad state or not ending properly. You can use this action to terminate a game session that's being hosted on any type of Amazon GameLift Servers fleet compute, including computes for managed EC2, managed container, and Anywhere fleets. The game server must be integrated with Amazon GameLift Servers server SDK 5.x or greater.

Request options

Request termination for a single game session. Provide the game session ID and the termination mode. There are two potential methods for terminating a game session:

- Initiate a graceful termination using the normal game session shutdown sequence. With this mode, the Amazon GameLift Servers service prompts the server process that's hosting the game session by calling the server SDK callback method `OnProcessTerminate()`. The callback implementation is part of the custom game server code. It might involve a variety of actions to gracefully end a game session, such as notifying players, before stopping the server process.
- Force an immediate game session termination. With this mode, the Amazon GameLift Servers service takes action to stop the server process, which ends the game session without the normal game session shutdown sequence.

Results

If successful, game session termination is initiated. During this activity, the game session status is changed to `TERMINATING`. When completed, the server process that was hosting the game session has been stopped and replaced with a new server process that's ready to host a new game session. The old game session's status is changed to `TERMINATED` with a status reason that indicates the termination method used.

Learn more

[Add Amazon GameLift Servers to your game server](#)

Amazon GameLift Servers server SDK 5 reference guide for `OnProcessTerminate()` ([C++](#)) ([C#](#)) ([Unreal](#)) ([Go](#))

Request Syntax

```
{  
  "GameSessionId": "string",  
  "TerminationMode": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[GameSessionId](#)

An identifier for the game session that is unique across all regions to be terminated.

The value is always a full ARN in the following format: For Home Region game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>`. For Remote Location game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<location>/<ID string>`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `[a-zA-Z0-9:/-]+`

Required: Yes

[TerminationMode](#)

The method to use to terminate the game session. Available methods include:

- **TRIGGER_ON_PROCESS_TERMINATE** – Prompts the Amazon GameLift Servers service to send an `OnProcessTerminate()` callback to the server process and initiate the normal game session shutdown sequence. The `OnProcessTerminate` method, which is implemented in the game server code, must include a call to the server SDK action `ProcessEnding()`, which is how the server process signals to Amazon GameLift Servers that a game session is ending. If the server process doesn't call `ProcessEnding()`, the game session termination won't conclude successfully.
- **FORCE_TERMINATE** – Prompts the Amazon GameLift Servers service to stop the server process immediately. Amazon GameLift Servers takes action (depending on the type of fleet) to shut down the server process without the normal game session shutdown sequence.

 Note

This method is not available for game sessions that are running on Anywhere fleets unless the fleet is deployed with the Amazon GameLift Servers Agent. In this scenario, a force terminate request results in an invalid or bad request exception.

Type: String

Valid Values: `TRIGGER_ON_PROCESS_TERMINATE` | `FORCE_TERMINATE`

Required: Yes

Response Syntax

```
{
  "GameSession": {
    "ComputeName": "string",
    "CreationTime": number,
    "CreatorId": "string",
    "CurrentPlayerSessionCount": number,
    "DnsName": "string",
    "FleetArn": "string",
    "FleetId": "string",
    "GameProperties": [
      {
        "Key": "string",
        "Value": "string"
      }
    ]
  }
}
```

```
    ],  
    "GameSessionData": "string",  
    "GameSessionId": "string",  
    "IpAddress": "string",  
    "Location": "string",  
    "MatchmakerData": "string",  
    "MaximumPlayerSessionCount": number,  
    "Name": "string",  
    "PlayerGatewayStatus": "string",  
    "PlayerSessionCreationPolicy": "string",  
    "Port": number,  
    "Status": "string",  
    "StatusReason": "string",  
    "TerminationTime": number  
  }  
}
```

Response Elements

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

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

GameSession

Properties describing a game session.

A game session in ACTIVE status can host players. When a game session ends, its status is set to TERMINATED.

Amazon GameLift Servers retains a game session resource for 30 days after the game session ends. You can reuse idempotency token values after this time. Game session logs are retained for 14 days.

All APIs by task

Type: [GameSession](#) object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidGameSessionStatusException

The requested operation would cause a conflict with the current state of a resource associated with the request and/or the game instance. Resolve the conflict before retrying.

HTTP Status Code: 400

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

NotReadyException

The operation failed because Amazon GameLift Servers has not yet finished validating this compute. We recommend attempting 8 to 10 retries over 3 to 5 minutes with [exponential backoffs and jitter](#).

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Terminate a game session

In this example, we want to gracefully end a game session that's entered a bad state. The game session is being hosted on a managed EC2 fleet.

Sample Request

```
{
  "GameSessionId": "arn:aws:gamelift:us-west-2::gamesession/
fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa/gsess-4444dddd-55ee-66ff-77aa-8888bbbb99cc",
  "TerminationMode": "TRIGGER_ON_PROCESS_TERMINATE"
}
```

CLI syntax:

```
aws gamelift terminate-game-session --game-session-id "arn:aws:gamelift:us-
west-2::gamesession/fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa/
gsess-4444dddd-55ee-66ff-77aa-8888bbbb99cc" --termination-mode
"TRIGGER_ON_PROCESS_TERMINATE"
```

Sample Response

```
{
  "GameSession": {
    "CreationTime": 1469498468.057,
    "CurrentPlayerSessionCount": 5,
    "FleetArn": "arn:aws:gamelift:us-west-2::fleet/
fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
    "FleetId": "fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa",
    "GameProperties": [
      {"Key": "difficulty", "Value": "easy"},
      {"Key": "gameMap", "Value": "Snowfall"},
      {"Key": "gameMode", "Value": "Explore"}
    ],
    "GameSessionId": "arn:aws:gamelift:us-west-2::gamesession/
fleet-9999ffff-88ee-77dd-66cc-5555bbbb44aa/
gsess-4444dddd-55ee-66ff-77aa-8888bbbb99cc",
    "IpAddress": "192.0.2.0",
    "Location": "us-west-2",
    "MaximumPlayerSessionCount": 10,
    "Name": "Matt's Awesome Game win123",
```



```
    "PlayerSessionCreationPolicy": "ACCEPT_ALL",  
    "Port": "8080",  
    "Status": "ACTIVE",  
    "ComputeName": "i-9999aaaabbbbccccc0"  
  }  
}
```

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

This API works with the following fleet types: EC2, Anywhere, Container

Removes a tag assigned to a Amazon GameLift Servers resource. You can use resource tags to organize AWS resources for a range of purposes. This operation handles the permissions necessary to manage tags for Amazon GameLift Servers resources that support tagging.

To remove a tag from a resource, specify the unique ARN value for the resource and provide a string list containing one or more tags to remove. This operation succeeds even if the list includes tags that aren't assigned to the resource.

Learn more

[Tagging AWS Resources](#) in the *AWS General Reference*

[AWS Tagging Strategies](#)

Related actions

[All APIs by task](#)

Request Syntax

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

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

ResourceARN

The Amazon Resource Name ([ARN](#)) that uniquely identifies the Amazon GameLift Servers resource that you want to remove tags from. Amazon GameLift Servers includes resource ARNs in the data object for the resource. You can retrieve the ARN by calling a `List` or `Describe` operation for the resource type.

Type: String

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

Required: Yes

TagKeys

A list of one or more tag keys to remove from the specified Amazon GameLift Servers resource.

Type: Array of strings

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

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

Required: Yes

Response Elements

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

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

TaggingFailedException

The requested tagging operation did not succeed. This may be due to invalid tag format or the maximum tag limit may have been exceeded. Resolve the issue before retrying.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

UpdateAlias

This API works with the following fleet types: EC2, Anywhere, Container

Updates properties for an alias. Specify the unique identifier of the alias to be updated and the new property values.

When reassigning an alias to a new fleet, provide an updated routing strategy. If successful, the updated alias record is returned.

Related actions

[All APIs by task](#)

Request Syntax

```
{
  "AliasId": "string",
  "Description": "string",
  "Name": "string",
  "RoutingStrategy": {
    "FleetId": "string",
    "Message": "string",
    "Type": "string"
  }
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

AliasId

A unique identifier for the alias that you want to update. You can use either the alias ID or ARN value.

Type: String

Pattern: `^alias-\S+|^arn:.*:alias\/alias-\S+`

Required: Yes

Description

A human-readable description of the alias.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Name

A descriptive label that is associated with an alias. Alias names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `.*\S.*`

Required: No

RoutingStrategy

The routing configuration, including routing type and fleet target, for the alias.

Type: [RoutingStrategy](#) object

Required: No

Response Syntax

```
{
  "Alias": {
    "AliasArn": "string",
    "AliasId": "string",
    "CreationTime": number,
    "Description": "string",
```

```
"LastUpdatedTime": number,
"Name": "string",
"RoutingStrategy": {
  "FleetId": "string",
  "Message": "string",
  "Type": "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.

Alias

The updated alias resource.

Type: [Alias](#) object

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

UpdateBuild

This API works with the following fleet types: EC2

Updates metadata in a build resource, including the build name and version. To update the metadata, specify the build ID to update and provide the new values. If successful, a build object containing the updated metadata is returned.

Learn more

[Upload a Custom Server Build](#)

[All APIs by task](#)

Request Syntax

```
{
  "BuildId": "string",
  "Name": "string",
  "Version": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

BuildId

A unique identifier for the build to update. You can use either the build ID or ARN value.

Type: String

Pattern: `^build-\S+|^arn:.*:build\/build-\S+`

Required: Yes

Name

A descriptive label that is associated with a build. Build names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Version

Version information that is associated with a build or script. Version strings do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "Build": {
    "BuildArn": "string",
    "BuildId": "string",
    "CreationTime": number,
    "Name": "string",
    "OperatingSystem": "string",
    "ServerSdkVersion": "string",
    "SizeOnDisk": number,
    "Status": "string",
    "Version": "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.

Build

The updated build resource.

Type: [Build](#) object

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Change a build resource

This example updates a build resource with a new name and version number, which are the only elements that can be changed. The returned build object verifies that the changes were made successfully.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "BuildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "Name": "My_Game_Server_Build_Foo",
  "Version": "12345.f00"
}
```

Sample Response

```
{
  "Build": {
    "BuildArn": "arn:aws:gamelift:us-west-2::build/build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "BuildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "CreationTime": 1496708916.18,
    "Name": "My_Game_Server_Build_Foo",
    "OperatingSystem": "AMAZON_LINUX_2023",
    "SizeOnDisk": 1304924,
    "Status": "READY",
    "Version": "12345.f00"
  }
}
```

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](#)

UpdateContainerFleet

This API works with the following fleet types: Container

Updates the properties of a managed container fleet. Depending on the properties being updated, this operation might initiate a fleet deployment. You can track deployments for a fleet using https://docs.aws.amazon.com/gamelift/latest/apireference/API_DescribeFleetDeployment.html.

Note

A managed fleet's runtime environment, which depends on the fleet's Amazon Machine Image {AMI} version, can't be updated. You must create a new fleet. As a best practice, we recommend replacing your managed fleets every 30 days to maintain a secure and up-to-date runtime environment for your hosted game servers. For guidance, see [Security best practices for Amazon GameLift Servers](#).

Request options

As with CreateContainerFleet, many fleet properties use common defaults or are calculated based on the fleet's container group definitions.

- Update fleet properties that result in a fleet deployment. Include only those properties that you want to change. Specify deployment configuration settings.
- Update fleet properties that don't result in a fleet deployment. Include only those properties that you want to change.

Changes to the following properties initiate a fleet deployment:

- GameServerContainerGroupDefinition
- PerInstanceContainerGroupDefinition
- GameServerContainerGroupsPerInstance
- InstanceInboundPermissions
- InstanceConnectionPortRange
- LogConfiguration

Results

If successful, this operation updates the container fleet resource, and might initiate a new deployment of fleet resources using the deployment configuration provided. A deployment replaces existing fleet instances with new instances that are deployed with the updated fleet properties. The fleet is placed in `UPDATING` status until the deployment is complete, then return to `ACTIVE`.

You can have only one update deployment active at a time for a fleet. If a second update request initiates a deployment while another deployment is in progress, the first deployment is cancelled.

Request Syntax

```
{
  "DeploymentConfiguration": {
    "ImpairmentStrategy": "string",
    "MinimumHealthyPercentage": number,
    "ProtectionStrategy": "string"
  },
  "Description": "string",
  "FleetId": "string",
  "GameServerContainerGroupDefinitionName": "string",
  "GameServerContainerGroupsPerInstance": number,
  "GameSessionCreationLimitPolicy": {
    "NewGameSessionsPerCreator": number,
    "PolicyPeriodInMinutes": number
  },
  "InstanceConnectionPortRange": {
    "FromPort": number,
    "ToPort": number
  },
  "InstanceInboundPermissionAuthorizations": [
    {
      "FromPort": number,
      "IpRange": "string",
      "Protocol": "string",
      "ToPort": number
    }
  ],
  "InstanceInboundPermissionRevocations": [
    {
      "FromPort": number,
      "IpRange": "string",
      "Protocol": "string",

```

```
    "ToPort": number
  },
],
"LogConfiguration": {
  "LogDestination": "string",
  "LogGroupArn": "string",
  "S3BucketName": "string"
},
"MetricGroups": [ "string" ],
"NewGameSessionProtectionPolicy": "string",
"PerInstanceContainerGroupDefinitionName": "string",
"RemoveAttributes": [ "string" ]
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the container fleet to update. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\-]+|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\-]+`

Required: Yes

DeploymentConfiguration

Instructions for how to deploy updates to a container fleet, if the fleet update initiates a deployment. The deployment configuration lets you determine how to replace fleet instances and what actions to take if the deployment fails.

Type: [DeploymentConfiguration](#) object

Required: No

Description

A meaningful description of the container fleet.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

GameServerContainerGroupDefinitionName

The name or ARN value of a new game server container group definition to deploy on the fleet. If you're updating the fleet to a specific version of a container group definition, use the ARN value and include the version number. If you're updating the fleet to the latest version of a container group definition, you can use the name value. You can't remove a fleet's game server container group definition, you can only update or replace it with another definition.

Update a container group definition by calling [UpdateContainerGroupDefinition](#). This operation creates a [ContainerGroupDefinition](#) resource with an incremented version.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-zA-Z0-9\-\_]+\$|^arn:.*:containergroupdefinition\[a-zA-Z0-9\-\_]+(:[0-9]+)?\$`

Required: No

GameServerContainerGroupsPerInstance

The number of times to replicate the game server container group on each fleet instance. By default, Amazon GameLift Servers calculates the maximum number of game server container groups that can fit on each instance. You can remove this property value to use the calculated

value, or set it manually. If you set this number manually, Amazon GameLift Servers uses your value as long as it's less than the calculated maximum.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 5000.

Required: No

GameSessionCreationLimitPolicy

A policy that limits the number of game sessions that each individual player can create on instances in this fleet. The limit applies for a specified span of time.

Type: [GameSessionCreationLimitPolicy](#) object

Required: No

InstanceConnectionPortRange

A revised set of port numbers to open on each fleet instance. By default, Amazon GameLift Servers calculates an optimal port range based on your fleet configuration. If you previously set this parameter manually, you can't reset this to use the calculated settings.

The port range must not overlap with the Amazon GameLift Servers reserved port range 4092-4191. This range is reserved for internal Amazon GameLift Servers services.

Type: [ConnectionPortRange](#) object

Required: No

InstanceInboundPermissionAuthorizations

A set of ports to add to the container fleet's inbound permissions.

The port range must not overlap with the Amazon GameLift Servers reserved port range 4092-4191. This range is reserved for internal Amazon GameLift Servers services.

Type: Array of [IpPermission](#) objects

Array Members: Maximum number of 50 items.

Required: No

InstanceInboundPermissionRevocations

A set of ports to remove from the container fleet's inbound permissions.

Type: Array of [IpPermission](#) objects

Array Members: Maximum number of 50 items.

Required: No

LogConfiguration

The method for collecting container logs for the fleet.

Type: [LogConfiguration](#) object

Required: No

MetricGroups

The name of an AWS CloudWatch metric group to add this fleet to.

Type: Array of strings

Array Members: Maximum number of 1 item.

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

Required: No

NewGameSessionProtectionPolicy

The game session protection policy to apply to all new game sessions that are started in this fleet. Game sessions that already exist are not affected.

Type: String

Valid Values: NoProtection | FullProtection

Required: No

PerInstanceContainerGroupDefinitionName

The name or ARN value of a new per-instance container group definition to deploy on the fleet. If you're updating the fleet to a specific version of a container group definition, use the ARN value and include the version number. If you're updating the fleet to the latest version of a container group definition, you can use the name value.

Update a container group definition by calling [UpdateContainerGroupDefinition](#). This operation creates a [ContainerGroupDefinition](#) resource with an incremented version.

To remove a fleet's per-instance container group definition, leave this parameter empty and use the parameter `RemoveAttributes`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-zA-Z0-9\-\_]+\$|^arn:.*:containergroupdefinition\[a-zA-Z0-9\-\_]+(:[0-9]+)?\$`

Required: No

[RemoveAttributes](#)

If set, this update removes a fleet's per-instance container group definition. You can't remove a fleet's game server container group definition.

Type: Array of strings

Array Members: Fixed number of 1 item.

Valid Values: `PER_INSTANCE_CONTAINER_GROUP_DEFINITION`

Required: No

Response Syntax

```
{
  "ContainerFleet": {
    "BillingType": "string",
    "CreationTime": number,
    "DeploymentDetails": {
      "LatestDeploymentId": "string"
    },
    "Description": "string",
    "FleetArn": "string",
    "FleetId": "string",
    "FleetRoleArn": "string",
    "GameServerContainerGroupDefinitionArn": "string",
    "GameServerContainerGroupDefinitionName": "string",
    "GameServerContainerGroupsPerInstance": number,
    "GameSessionCreationLimitPolicy": {
```

```

    "NewGameSessionsPerCreator": number,
    "PolicyPeriodInMinutes": number
  },
  "InstanceConnectionPortRange": {
    "FromPort": number,
    "ToPort": number
  },
  "InstanceInboundPermissions": [
    {
      "FromPort": number,
      "IpRange": "string",
      "Protocol": "string",
      "ToPort": number
    }
  ],
  "InstanceType": "string",
  "LocationAttributes": [
    {
      "Location": "string",
      "PlayerGatewayStatus": "string",
      "Status": "string"
    }
  ],
  "LogConfiguration": {
    "LogDestination": "string",
    "LogGroupArn": "string",
    "S3BucketName": "string"
  },
  "MaximumGameServerContainerGroupsPerInstance": number,
  "MetricGroups": [ "string" ],
  "NewGameSessionProtectionPolicy": "string",
  "PerInstanceContainerGroupDefinitionArn": "string",
  "PerInstanceContainerGroupDefinitionName": "string",
  "PlayerGatewayMode": "string",
  "Status": "string"
}
}

```

Response Elements

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

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

ContainerFleet

A collection of container fleet objects for all fleets that match the request criteria.

Type: [ContainerFleet](#) object

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

NotReadyException

The operation failed because Amazon GameLift Servers has not yet finished validating this compute. We recommend attempting 8 to 10 retries over 3 to 5 minutes with [exponential backoffs and jitter](#).

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Update a single-region container fleet with a new game server container group

This example updates the fleet with a new version of the fleet's game server container group. No other fleet properties are changed. The update initiates a fleet deployment.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "DeploymentConfiguration": {
    "ImpairmentStrategy": "ROLLBACK",
    "MinimumHealthyPercentage": 75,
    "ProtectionStrategy": "WITH_PROTECTION"
  },
  "FleetId": "containerfleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "GameServerContainerGroupDefinitionName": "arn:aws:gamelift:us-west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:2"
}
```

Sample Response

```
{
  "ContainerFleet": {
    "BillingType": ON_DEMAND,
    "CreationTime": 1736365885.22,
    "DeploymentDetails": {
```

```
    "LatestDeploymentId": "deployment-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
  },
  "FleetArn": "arn:aws:gamelift:us-west-2:111122223333:containerfleet/
containerfleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "FleetId": "containerfleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "FleetRoleArn": "arn:aws:iam::MyAccount:role/MyRole",
  "GameServerContainerGroupDefinitionArn": "arn:aws:gamelift:us-
west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:2",
  "GameServerContainerGroupDefinitionName": "MyAdventureGameContainerGroup",
  "GameServerContainerGroupsPerInstance": number,
  "InstanceConnectionPortRange": {
    "FromPort": 4192,
    "ToPort": 4242
  },
  "InstanceInboundPermissions": [
    {
      "FromPort": 4192,
      "IpRange": "string",
      "Protocol": "UDP",
      "ToPort": 4242,
    }
  ],
  "InstanceType": "c5.large",
  "LogConfiguration": {
    "LogGroupArn": "arn:aws:logs:us-west-2:111222333444:log-group:customerLogs",
    "LogDestination": "CLOUDWATCH"
  },
  "MaximumGameServerContainerGroupsPerInstance": 10,
  "NewGameSessionProtectionPolicy": "NoProtection",
  "Status": "UPDATING"
}
}
```

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](#)

UpdateContainerGroupDefinition

This API works with the following fleet types: Container

Updates properties in an existing container group definition. This operation doesn't replace the definition. Instead, it creates a new version of the definition and saves it separately. You can access all versions that you choose to retain.

The only property you can't update is the container group type.

Request options:

- Update based on the latest version of the container group definition. Specify the container group definition name only, or use an ARN value without a version number. Provide updated values for the properties that you want to change only. All other values remain the same as the latest version.
- Update based on a specific version of the container group definition. Specify the container group definition name and a source version number, or use an ARN value with a version number. Provide updated values for the properties that you want to change only. All other values remain the same as the source version.
- Change a game server container definition. Provide a complete set of container definitions, including the updated definition.
- Add or change a support container definition. Provide a complete set of container definitions, including the updated definition.
- Remove a support container definition. Provide a complete set of container definitions, excluding the definition to remove. If the container group has only one support container definition, provide an empty set.

Results:

If successful, this operation returns the complete properties of the new container group definition version.

If the container group definition version is used in an active fleets, the update automatically initiates a new fleet deployment of the new version. You can track a fleet's deployments using [ListFleetDeployments](#).

Request Syntax

```
{
  "GameServerContainerDefinition": {
    "ContainerName": "string",
    "DependsOn": [
      {
        "Condition": "string",
        "ContainerName": "string"
      }
    ],
    "EnvironmentOverride": [
      {
        "Name": "string",
        "Value": "string"
      }
    ],
    "ImageUri": "string",
    "LinuxCapabilities": {
      "Include": [ "string" ]
    },
    "MountPoints": [
      {
        "AccessLevel": "string",
        "ContainerPath": "string",
        "InstancePath": "string"
      }
    ],
    "PortConfiguration": {
      "ContainerPortRanges": [
        {
          "FromPort": number,
          "Protocol": "string",
          "ToPort": number
        }
      ]
    },
    "ServerSdkVersion": "string"
  },
  "Name": "string",
  "OperatingSystem": "string",
  "SourceVersionNumber": number,
  "SupportContainerDefinitions": [
```

```
{
  "ContainerName": "string",
  "DependsOn": [
    {
      "Condition": "string",
      "ContainerName": "string"
    }
  ],
  "EnvironmentOverride": [
    {
      "Name": "string",
      "Value": "string"
    }
  ],
  "Essential": boolean,
  "HealthCheck": {
    "Command": [ "string" ],
    "Interval": number,
    "Retries": number,
    "StartPeriod": number,
    "Timeout": number
  },
  "ImageUri": "string",
  "LinuxCapabilities": {
    "Include": [ "string" ]
  },
  "MemoryHardLimitMebibytes": number,
  "MountPoints": [
    {
      "AccessLevel": "string",
      "ContainerPath": "string",
      "InstancePath": "string"
    }
  ],
  "PortConfiguration": {
    "ContainerPortRanges": [
      {
        "FromPort": number,
        "Protocol": "string",
        "ToPort": number
      }
    ]
  },
  "Vcpu": number
}
```

```
    }  
  ],  
  "TotalMemoryLimitMebibytes": number,  
  "TotalVcpuLimit": number,  
  "VersionDescription": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Name

A descriptive identifier for the container group definition. The name value must be unique in an AWS Region.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-zA-Z0-9\-\_]+|^arn:.*:containergroupdefinition\[a-zA-Z0-9\-\_]+(:[0-9]+)?$`

Required: Yes

GameServerContainerDefinition

An updated definition for the game server container in this group. Define a game server container only when the container group type is `GAME_SERVER`. You can pass in your container definitions as a JSON file.

Type: [GameServerContainerDefinitionInput](#) object

Required: No

OperatingSystem

The platform that all containers in the group use. Containers in a group must run on the same operating system.

Note

Amazon Linux 2 (AL2) will reach end of support on 6/30/2026. See more details in the [Amazon Linux 2 FAQs](#). For game servers that are hosted on AL2 and use server SDK version 4.x for Amazon GameLift Servers, first update the game server build to server SDK 5.x, and then deploy to AL2023 instances. See [Migrate to server SDK version 5](#).

Type: String

Valid Values: AMAZON_LINUX_2023

Required: No

SourceVersionNumber

The container group definition version to update. The new version starts with values from the source version, and then updates values included in this request.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

SupportContainerDefinitions

One or more definitions for support containers in this group. You can define a support container in any type of container group. You can pass in your container definitions as a JSON file.

Type: Array of [SupportContainerDefinitionInput](#) objects

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

Required: No

TotalMemoryLimitMebibytes

The maximum amount of memory (in MiB) to allocate to the container group. All containers in the group share this memory. If you specify memory limits for an individual container, the total value must be greater than any individual container's memory limit.

Type: Integer

Valid Range: Minimum value of 4. Maximum value of 1024000.

Required: No

TotalVcpuLimit

The maximum amount of vCPU units to allocate to the container group (1 vCPU is equal to 1024 CPU units). All containers in the group share this memory. If you specify vCPU limits for individual containers, the total value must be equal to or greater than the sum of the CPU limits for all containers in the group.

Type: Double

Valid Range: Minimum value of 0.125. Maximum value of 10.

Required: No

VersionDescription

A description for this update to the container group definition.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Response Syntax

```
{
  "ContainerGroupDefinition": {
    "ContainerGroupDefinitionArn": "string",
    "ContainerGroupType": "string",
    "CreationTime": number,
    "GameServerContainerDefinition": {
      "ContainerName": "string",
```

```
"DependsOn": [  
  {  
    "Condition": "string",  
    "ContainerName": "string"  
  }  
],  
"EnvironmentOverride": [  
  {  
    "Name": "string",  
    "Value": "string"  
  }  
],  
"ImageUri": "string",  
"LinuxCapabilities": {  
  "Include": [ "string " ]  
},  
"MountPoints": [  
  {  
    "AccessLevel": "string",  
    "ContainerPath": "string",  
    "InstancePath": "string"  
  }  
],  
"PortConfiguration": {  
  "ContainerPortRanges": [  
    {  
      "FromPort": number,  
      "Protocol": "string",  
      "ToPort": number  
    }  
  ]  
},  
"ResolvedImageDigest": "string",  
"ServerSdkVersion": "string"  
},  
"Name": "string",  
"OperatingSystem": "string",  
"Status": "string",  
"StatusReason": "string",  
"SupportContainerDefinitions": [  
  {  
    "ContainerName": "string",  
    "DependsOn": [  
      {
```



```
        "Condition": "string",
        "ContainerName": "string"
    },
],
"EnvironmentOverride": [
    {
        "Name": "string",
        "Value": "string"
    }
],
"Essential": boolean,
"HealthCheck": {
    "Command": [ "string" ],
    "Interval": number,
    "Retries": number,
    "StartPeriod": number,
    "Timeout": number
},
"ImageUri": "string",
"LinuxCapabilities": {
    "Include": [ "string" ]
},
"MemoryHardLimitMebibytes": number,
"MountPoints": [
    {
        "AccessLevel": "string",
        "ContainerPath": "string",
        "InstancePath": "string"
    }
],
"PortConfiguration": {
    "ContainerPortRanges": [
        {
            "FromPort": number,
            "Protocol": "string",
            "ToPort": number
        }
    ]
},
"ResolvedImageDigest": "string",
"Vcpu": number
}
],
"TotalMemoryLimitMebibytes": number,
```

```
"TotalVcpuLimit": number,  
"VersionDescription": "string",  
"VersionNumber": number  
}  
}
```

Response Elements

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

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

ContainerGroupDefinition

The properties of the updated container group definition version.

Type: [ContainerGroupDefinition](#) object

Errors

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

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Update a game server container group definition for a new support container

This example adds a new support container definition to a game server container group that previously had one game server container defined. The request has these characteristics:

- A source version isn't provided, and the container group Name parameter is an ARN without a version number. As a result, this update will use the latest version as the source.
- The request changes the total memory limit for the container group.
- The new support container definition supplies the minimum required parameters.

As shown, an update request includes only the properties that you want to change or add.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "Name": "arn:aws:gamelift:us-west-2:111122223333:containergroupdefinition/
MyAdventureGameContainerGroup",
  "TotalMemoryLimitMebibytes": 2048,
```

```
"SupportContainerDefinitions": [  
  {  
    "ContainerName": "MyAdventureGameDependencies",  
    "ImageUri": "111122223333.dkr.ecr.us-west-2.amazonaws.com/  
MyAdventureGameDependenciesContainerImage"  
  }  
],  
"VersionDescription": "New support container for game server dependencies and groups  
total memory limit increased to 2048"  
}
```

Sample Response

```
{  
  "ContainerGroupDefinition": {  
    "ContainerGroupDefinitionArn": "arn:aws:gamelift:us-  
west-2:111122223333:containergroupdefinition/MyAdventureGameContainerGroup:4",  
    "ContainerGroupType": "GAME_SERVER",  
    "CreationTime": 1496365885.44,  
    "GameServerContainerDefinition": [  
      {  
        "ContainerName": "MyAdventureGameServer",  
        "ImageUri": "111122223333.dkr.ecr.us-west-2.amazonaws.com/  
MyAdventureGameContainerImage",  
        "PortConfiguration": {  
          "ContainerPortRanges": [  
            {  
              "FromPort": 35000,  
              "Protocol": "TCP",  
              "ToPort": 40000  
            }  
          ]  
        },  
        "LinuxCapabilities": {  
          "Include": ["NET_BIND_SERVICE", "SYS_PTRACE"]  
        },  
        "ResolvedImageDigest":  
"sha256:0123456789abcdef0123456789abcdef0123456789abcdef0123456789abcdef",  
        "ServerSdkVersion": "5.2.0"  
      }  
    ],  
    "Name": "MyAdventureGameContainerGroup",  
    "OperatingSystem": "AMAZON_LINUX_2023",  
  }  
}
```

```
"Status": "COPYING",
"SupportContainerDefinitions": [
  {
    "ContainerName": "MyAdventureGameDependencies",
    "ImageUri": "111122223333.dkr.ecr.us-west-2.amazonaws.com/
MyAdventureGameDependenciesContainerImage"
    "ResolvedImageDigest":
      "sha256:0123456789abcdef0123456789abcdef0123456789fedcba",
  }
],
"TotalMemoryLimitMebibytes": 2048,
"TotalVcpuLimit": 1,
"VersionDescription": "New support container for game server dependencies and
groups total memory limit increased to 2048",
"VersionNumber": 4
}
```

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](#)

UpdateFleetAttributes

This API works with the following fleet types: EC2, Anywhere

Updates a fleet's mutable attributes, such as game session protection and resource creation limits.

To update fleet attributes, specify the fleet ID and the property values that you want to change. If successful, Amazon GameLift Servers returns the identifiers for the updated fleet.

Note

A managed fleet's runtime environment, which depends on the fleet's Amazon Machine Image {AMI} version, can't be updated. You must create a new fleet. As a best practice, we recommend replacing your managed fleets every 30 days to maintain a secure and up-to-date runtime environment for your hosted game servers. For guidance, see [Security best practices for Amazon GameLift Servers](#).

Learn more

[Setting up Amazon GameLift Servers fleets](#)

Request Syntax

```
{
  "AnywhereConfiguration": {
    "Cost": "string"
  },
  "Description": "string",
  "FleetId": "string",
  "MetricGroups": [ "string" ],
  "Name": "string",
  "NewGameSessionProtectionPolicy": "string",
  "ResourceCreationLimitPolicy": {
    "NewGameSessionsPerCreator": number,
    "PolicyPeriodInMinutes": number
  }
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet to update attribute metadata for. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

AnywhereConfiguration

Amazon GameLift Servers Anywhere configuration options.

Type: [AnywhereConfiguration](#) object

Required: No

Description

A human-readable description of a fleet.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

MetricGroups

The name of a metric group to add this fleet to. Use a metric group in Amazon CloudWatch to aggregate the metrics from multiple fleets. Provide an existing metric group name, or create a new metric group by providing a new name. A fleet can only be in one metric group at a time.

Type: Array of strings

Array Members: Maximum number of 1 item.

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

Required: No

Name

A descriptive label that is associated with a fleet. Fleet names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

NewGameSessionProtectionPolicy

The game session protection policy to apply to all new game sessions created in this fleet. Game sessions that already exist are not affected. You can set protection for individual game sessions using [UpdateGameSession](#) .

- **NoProtection** -- The game session can be terminated during a scale-down event.
- **FullProtection** -- If the game session is in an ACTIVE status, it cannot be terminated during a scale-down event.

Type: String

Valid Values: NoProtection | FullProtection

Required: No

ResourceCreationLimitPolicy

Policy settings that limit the number of game sessions an individual player can create over a span of time.

Type: [ResourceCreationLimitPolicy](#) object

Required: No

Response Syntax

```
{
  "FleetArn": "string",
  "FleetId": "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.

FleetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\]\+$`

FleetId

A unique identifier for the fleet that was updated.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\]\+$`

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidFleetStatusException

The requested operation would cause a conflict with the current state of a resource associated with the request and/or the fleet. Resolve the conflict before retrying.

HTTP Status Code: 400

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

UpdateFleetCapacity

This API works with the following fleet types: EC2, Container

Updates capacity settings for a managed EC2 fleet or managed container fleet. For these fleets, you adjust capacity by changing the number of instances in the fleet. Fleet capacity determines the number of game sessions and players that the fleet can host based on its configuration. For fleets with multiple locations, use this operation to manage capacity settings in each location individually.

- **Minimum/maximum size:** Set hard limits on the number of Amazon EC2 instances allowed. If Amazon GameLift Servers receives a request--either through manual update or automatic scaling--it won't change the capacity to a value outside of this range.
- **Desired capacity:** As an alternative to automatic scaling, manually set the number of Amazon EC2 instances to be maintained. Before changing a fleet's desired capacity, check the maximum capacity of the fleet's Amazon EC2 instance type by calling [DescribeEC2InstanceLimits](#).

To update capacity for a fleet's home Region, or if the fleet has no remote locations, omit the `Location` parameter. The fleet must be in `ACTIVE` status.

To update capacity for a fleet's remote location, set the `Location` parameter to the location to update. The location must be in `ACTIVE` status.

If successful, Amazon GameLift Servers updates the capacity settings and returns the identifiers for the updated fleet and/or location. If a requested change to desired capacity exceeds the instance type's limit, the `LimitExceeded` exception occurs.

Updates often prompt an immediate change in fleet capacity, such as when current capacity is different than the new desired capacity or outside the new limits. In this scenario, Amazon GameLift Servers automatically initiates steps to add or remove instances in the fleet location. You can track a fleet's current capacity by calling [DescribeFleetCapacity](#) or [DescribeFleetLocationCapacity](#).

Use `ManagedCapacityConfiguration` with the `"SCALE_TO_AND_FROM_ZERO"` `ZeroCapacityStrategy` to enable Amazon GameLift Servers to fully manage the `MinSize` value, switching between 0 and 1 based on game session activity. This is ideal for eliminating compute costs during periods of no game activity. It is particularly beneficial during development when you're away from your desk, iterating on builds for extended periods, in production environments serving low-traffic locations, or for games with long, predictable downtime windows. By automatically managing capacity

between 0 and 1 instances, you avoid paying for idle instances while maintaining the ability to serve game sessions when demand arrives. Note that while scale-out is triggered immediately upon receiving a game session request, actual game session availability depends on your server process startup time, so this approach works best with multi-location Fleets where cold-start latency is tolerable. With a "MANUAL" ZeroCapacityStrategy Amazon GameLift Servers will not modify Fleet MinSize values automatically and will not scale out from zero instances in response to game sessions. This is configurable per-location.

Learn more

[Scaling fleet capacity](#)

Request Syntax

```
{
  "DesiredInstances": number,
  "FleetId": "string",
  "Location": "string",
  "ManagedCapacityConfiguration": {
    "ScaleInAfterInactivityMinutes": number,
    "ZeroCapacityStrategy": "string"
  },
  "MaxSize": number,
  "MinSize": number
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[FleetId](#)

A unique identifier for the fleet to update capacity settings for. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

DesiredInstances

The number of Amazon EC2 instances you want to maintain in the specified fleet location. This value must fall between the minimum and maximum size limits. Changes in desired instance value can take up to 1 minute to be reflected when viewing the fleet's capacity settings.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

Location

The name of a remote location to update fleet capacity settings for, in the form of an AWS Region code such as `us-west-2`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-\_]+\$`

Required: No

ManagedCapacityConfiguration

Configuration for Amazon GameLift Servers-managed capacity scaling options.

Type: [ManagedCapacityConfiguration](#) object

Required: No

MaxSize

The maximum number of instances that are allowed in the specified fleet location. If this parameter is not set, the default is 1.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

MinSize

The minimum number of instances that are allowed in the specified fleet location. If this parameter is not set, the default is 0. This parameter cannot be set when using a `ManagedCapacityConfiguration` where `ZeroCapacityStrategy` has a value of `SCALE_TO_AND_FROM_ZERO`.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

Response Syntax

```
{
  "FleetArn": "string",
  "FleetId": "string",
  "Location": "string",
  "ManagedCapacityConfiguration": {
    "ScaleInAfterInactivityMinutes": number,
    "ZeroCapacityStrategy": "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.

FleetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is

arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: ^arn:.*:[a-z]*fleet\[a-z\]*fleet-[a-zA-Z0-9\-\-]+\$

FleetId

A unique identifier for the fleet that was updated.

Type: String

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

Pattern: ^[a-z]*fleet-[a-zA-Z0-9\-\-]+\$

Location

The remote location being updated, expressed as an AWS Region code, such as us-west-2.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: ^[A-Za-z0-9\-\-]+\$

ManagedCapacityConfiguration

Configuration for Amazon GameLift Servers-managed capacity scaling options.

Type: [ManagedCapacityConfiguration](#) object

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidFleetStatusException

The requested operation would cause a conflict with the current state of a resource associated with the request and/or the fleet. Resolve the conflict before retrying.

HTTP Status Code: 400

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

HTTP Status Code: 400

Examples

Increase fleet capacity

This example requests an increase in fleet capacity and raises the maximum allowed instances to accommodate the request. Because no location is specified, the update is applied to the fleet's home Region (where the fleet was created), which is `us-west-2`. Note that a fleet's home Region is always identified in the fleet ARN.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "FleetId": "arn:aws:gamelift:us-west-2::fleet/
fleet-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "DesiredInstances": 10,
  "MaxSize": 10
}
```

Sample Response

```
{
  "FleetArn": "arn:aws:gamelift:us-west-2::fleet/
fleet-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "FleetId": "fleet-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "Location": "us-west-2",
  "ManagedCapacityConfiguration": {
    "ZeroCapacityStrategy": "MANUAL"
  }
}
```

Update capacity settings for a fleet's remote location

This example adjusts the minimum/maximum allowed instances for one remote location in a multi-location fleet. For fleets that use auto-scaling, allowed instance settings define the scaling range

for a fleet location. In this example, the updates are being applied to the remote location `sa-east-1`.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "FleetId": "arn:aws:gamelift:us-west-2::fleet/
fleet-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "Location": "sa-east-1",
  "MinSize": 10,
  "MaxSize": 100
}
```

Sample Response

```
{
  "FleetArn": "arn:aws:gamelift:us-west-2::fleet/
fleet-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "FleetId": "fleet-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "Location": "sa-east-1",
  "ManagedCapacityConfiguration": {
    "ZeroCapacityStrategy": "MANUAL"
  }
}
```

Configure ManagedCapacityConfiguration to Scale to/from Zero

This example configures `ManagedCapacityConfiguration` such that fleet capacity will automatically scale in to zero following a period (60 minutes in this example) where no game sessions are placed on the fleet and will scale out from zero in response to game session activity.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the `Authorization` header field.

Sample Request

```
{
  "FleetId": "arn:aws:gamelift:us-west-2::fleet/
fleet-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
```

```
"ManagedCapacityConfiguration": {  
    "ScaleInAfterInactivityMinutes": 60,  
    "ZeroCapacityStrategy": "SCALE_TO_AND_FROM_ZERO"  
}  
}
```

Sample Response

```
{  
    "FleetArn": "arn:aws:gamelift:us-west-2::fleet/  
fleet-1111aaaa-22bb-33cc-44dd-5555eeee66ff",  
    "FleetId": "fleet-1111aaaa-22bb-33cc-44dd-5555eeee66ff",  
    "Location": "us-west-2",  
    "ManagedCapacityConfiguration": {  
        "ScaleInAfterInactivityMinutes": 60,  
        "ZeroCapacityStrategy": "SCALE_TO_AND_FROM_ZERO"  
    }  
}
```

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](#)

UpdateFleetPortSettings

This API works with the following fleet types: EC2

Updates permissions that allow inbound traffic to connect to game sessions in the fleet.

To update settings, specify the fleet ID to be updated and specify the changes to be made. List the permissions you want to add in `InboundPermissionAuthorizations`, and permissions you want to remove in `InboundPermissionRevocations`. Permissions to be removed must match existing fleet permissions.

If successful, the fleet identifiers for the updated fleet are returned. For fleets with remote locations, port setting updates can take time to propagate across all locations. You can check the status of updates in each location by calling `DescribeFleetPortSettings` with a location name.

Learn more

[Setting up Amazon GameLift Servers fleets](#)

Request Syntax

```
{
  "FleetId": "string",
  "InboundPermissionAuthorizations": [
    {
      "FromPort": number,
      "IpRange": "string",
      "Protocol": "string",
      "ToPort": number
    }
  ],
  "InboundPermissionRevocations": [
    {
      "FromPort": number,
      "IpRange": "string",
      "Protocol": "string",
      "ToPort": number
    }
  ]
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet to update port settings for. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

InboundPermissionAuthorizations

A collection of port settings to be added to the fleet resource.

Type: Array of [IpPermission](#) objects

Array Members: Maximum number of 50 items.

Required: No

InboundPermissionRevocations

A collection of port settings to be removed from the fleet resource.

Type: Array of [IpPermission](#) objects

Array Members: Maximum number of 50 items.

Required: No

Response Syntax

```
{
  "FleetArn": "string",
  "FleetId": "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.

FleetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet[/][a-z]*fleet-[a-zA-Z0-9\-.]+`

FleetId

A unique identifier for the fleet that was updated.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-.]+`

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidFleetStatusException

The requested operation would cause a conflict with the current state of a resource associated with the request and/or the fleet. Resolve the conflict before retrying.

HTTP Status Code: 400

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Update fleet to open to enable SSH access

The following example opens port 22 for SSH access:

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "InboundPermissionAuthorizations": [
    {
      "FromPort": 22,
      "IpRange": "54.186.139.221/32",
      "Protocol": "TCP",
      "ToPort": 22
    }
  ]
}
```

Sample Response

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
}
```

Update fleet to enable Windows Remote Desktop access

The following example opens port 3389 for Windows Remote Desktop access:

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "InboundPermissionAuthorizations": [
```

```
{
  "FromPort": 3389,
  "IpRange": "54.186.139.221/32",
  "Protocol": "TCP",
  "ToPort": 3389
}
```

Sample Response

```
{
  "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
}
```

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](#)

UpdateGameServer

This API works with the following fleet types: EC2 (FleetIQ)

Updates information about a registered game server to help Amazon GameLift Servers FleetIQ track game server availability. This operation is called by a game server process that is running on an instance in a game server group.

Use this operation to update the following types of game server information. You can make all three types of updates in the same request:

- To update the game server's utilization status from AVAILABLE (when the game server is available to be claimed) to UTILIZED (when the game server is currently hosting games). Identify the game server and game server group and specify the new utilization status. You can't change the status from UTILIZED to AVAILABLE .
- To report health status, identify the game server and game server group and set health check to HEALTHY. If a game server does not report health status for a certain length of time, the game server is no longer considered healthy. As a result, it will be eventually deregistered from the game server group to avoid affecting utilization metrics. The best practice is to report health every 60 seconds.
- To change game server metadata, provide updated game server data.

Once a game server is successfully updated, the relevant statuses and timestamps are updated.

Learn more

[Amazon GameLift Servers FleetIQ Guide](#)

Request Syntax

```
{  
  "GameServerData": "string",  
  "GameServerGroupName": "string",  
  "GameServerId": "string",  
  "HealthCheck": "string",  
  "UtilizationStatus": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

GameServerGroupName

A unique identifier for the game server group where the game server is running.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]|^arn:.*:gameservergroup\[a-zA-Z0-9-\.]`

Required: Yes

GameServerId

A custom string that uniquely identifies the game server to update.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `[a-zA-Z0-9-\.]`

Required: Yes

GameServerData

A set of custom game server properties, formatted as a single string value. This data is passed to a game client or service when it requests information on game servers.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `.*\S.*`

Required: No

HealthCheck

Indicates health status of the game server. A request that includes this parameter updates the game server's *LastHealthCheckTime* timestamp.

Type: String

Valid Values: HEALTHY

Required: No

UtilizationStatus

Indicates if the game server is available or is currently hosting gameplay. You can update a game server status from AVAILABLE to UTILIZED, but you can't change a the status from UTILIZED to AVAILABLE.

Type: String

Valid Values: AVAILABLE | UTILIZED

Required: No

Response Syntax

```
{
  "GameServer": {
    "ClaimStatus": "string",
    "ConnectionInfo": "string",
    "GameServerData": "string",
    "GameServerGroupArn": "string",
    "GameServerGroupName": "string",
    "GameServerId": "string",
    "InstanceId": "string",
    "LastClaimTime": number,
    "LastHealthCheckTime": number,
    "RegistrationTime": number,
    "UtilizationStatus": "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.

GameServer

Object that describes the newly updated game server.

Type: [GameServer](#) object

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Report that a game server is now hosting gameplay

This example changes the utilization status of a game server and reports the results of a health check. Prior to this call, the game server's utilization status is AVAILABLE and the claim status is CLAIMED.

Sample Request

```
{
  "GameServerGroupName": "MegaFrogServers_NA",
  "HealthCheck": "HEALTHY",
  "UtilizationStatus": "UTILIZED"
}
```

CLI command:

```
aws gamelift update-game-server \
  --game-server-group-name MegaFrogServers_NA \
  --health-check HEALTHY \
  --utilization-status UTILIZED
```

Sample Response

```
{
  "GameServer": {
    "ClaimStatus": "",
    "ConnectionInfo": "192.0.2.0.80",
    "GameServerData": "",
    "GameServerGroupArn": "arn:aws:gamelift:us-west-2::GameServerGroup/MegaFrogServers_NA",
    "GameServerGroupName": "MegaFrogServers_NA",
    "GameServerId": "mega-frog-game-12345678",
    "InstanceId": "i-1234567890abcdef0",
    "LastClaimTime": 1580218290.293,
    "LastHealthCheckTime": 1580218294.293,
    "RegistrationTime": 1580218197.293,
    "UtilizationStatus": "UTILIZED"
  }
}
```

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](#)

UpdateGameServerGroup

This API works with the following fleet types: EC2 (FleetIQ)

Updates Amazon GameLift Servers FleetIQ-specific properties for a game server group. Many Auto Scaling group properties are updated on the Auto Scaling group directly, including the launch template, Auto Scaling policies, and maximum/minimum/desired instance counts.

To update the game server group, specify the game server group ID and provide the updated values. Before applying the updates, the new values are validated to ensure that Amazon GameLift Servers FleetIQ can continue to perform instance balancing activity. If successful, a `GameServerGroup` object is returned.

Note

Target tracking Auto Scaling policies on the Auto Scaling group cannot be updated through the AWS Management Console. Instead, use the Amazon Elastic Compute Cloud Auto Scaling [PutScalingPolicy](#) API action to update these policies.

Learn more

[Amazon GameLift Servers FleetIQ Guide](#)

Request Syntax

```
{
  "BalancingStrategy": "string",
  "GameServerGroupName": "string",
  "GameServerProtectionPolicy": "string",
  "InstanceDefinitions": [
    {
      "InstanceType": "string",
      "WeightedCapacity": "string"
    }
  ],
  "RoleArn": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

GameServerGroupName

A unique identifier for the game server group. Use either the name or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]|^arn:.*:gameservergroup\[a-zA-Z0-9-\.]`

Required: Yes

BalancingStrategy

Indicates how Amazon GameLift Servers FleetIQ balances the use of Spot Instances and On-Demand Instances in the game server group. Method options include the following:

- **SPOT_ONLY** - Only Spot Instances are used in the game server group. If Spot Instances are unavailable or not viable for game hosting, the game server group provides no hosting capacity until Spot Instances can again be used. Until then, no new instances are started, and the existing nonviable Spot Instances are terminated (after current gameplay ends) and are not replaced.
- **SPOT_PREFERRED** - (default value) Spot Instances are used whenever available in the game server group. If Spot Instances are unavailable, the game server group continues to provide hosting capacity by falling back to On-Demand Instances. Existing nonviable Spot Instances are terminated (after current gameplay ends) and are replaced with new On-Demand Instances.
- **ON_DEMAND_ONLY** - Only On-Demand Instances are used in the game server group. No Spot Instances are used, even when available, while this balancing strategy is in force.

Type: String

Valid Values: SPOT_ONLY | SPOT_PREFERRED | ON_DEMAND_ONLY

Required: No

GameServerProtectionPolicy

A flag that indicates whether instances in the game server group are protected from early termination. Unprotected instances that have active game servers running might be terminated during a scale-down event, causing players to be dropped from the game. Protected instances cannot be terminated while there are active game servers running except in the event of a forced game server group deletion (see). An exception to this is with Spot Instances, which can be terminated by AWS regardless of protection status. This property is set to NO_PROTECTION by default.

Type: String

Valid Values: NO_PROTECTION | FULL_PROTECTION

Required: No

InstanceDefinitions

An updated list of Amazon EC2 instance types to use in the Auto Scaling group. The instance definitions must specify at least two different instance types that are supported by Amazon GameLift Servers FleetIQ. This updated list replaces the entire current list of instance definitions for the game server group. For more information on instance types, see [EC2 Instance Types](#) in the *Amazon EC2 User Guide*. You can optionally specify capacity weighting for each instance type. If no weight value is specified for an instance type, it is set to the default value "1". For more information about capacity weighting, see [Instance Weighting for Amazon EC2 Auto Scaling](#) in the Amazon EC2 Auto Scaling User Guide.

Type: Array of [InstanceDefinition](#) objects

Array Members: Minimum number of 2 items. Maximum number of 20 items.

Required: No

RoleArn

The Amazon Resource Name ([ARN](#)) for an IAM role that allows Amazon GameLift Servers to access your Amazon EC2 Auto Scaling groups.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: ^arn:.*:role\[\\w+=, .@-\\]+

Required: No

Response Syntax

```
{
  "GameServerGroup": {
    "AutoScalingGroupArn": "string",
    "BalancingStrategy": "string",
    "CreationTime": number,
    "GameServerGroupArn": "string",
    "GameServerGroupName": "string",
    "GameServerProtectionPolicy": "string",
    "InstanceDefinitions": [
      {
        "InstanceType": "string",
        "WeightedCapacity": "string"
      }
    ],
    "LastUpdatedTime": number,
    "RoleArn": "string",
    "Status": "string",
    "StatusReason": "string",
    "SuspendedActions": [ "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.

GameServerGroup

An object that describes the game server group resource with updated properties.

Type: [GameServerGroup](#) object

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Update instance definitions for the game server group

This example specifies a new set of instance definitions for the game server group. Previously, the game server group used a Spot-only balancing strategy with c3.xlarge and c4.xlarge instance types. This request removes the old instance definitions and replaces them. As shown in the response, because the new definitions do not specify `WeightedCapacity`, these are set to the default value of 1.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{
  "GameServerGroupName": "MegaFrogServers_NA",
  "InstanceDefinitions": [
    {"InstanceType": "c4.xlarge"},
    {"InstanceType": "c5.xlarge"}
  ]
}
```

CLI command:

```
aws gamelift update-game-server-group \
  --game-server-group MegaFrogServers_NA \
  --instance-definitions '[{"InstanceType": "c4.xlarge"}, {"InstanceType": "c5.xlarge"}]'
```

Sample Response

```
{
  "GameServerGroup": {
    "AutoScalingGroupArn": "arn:aws:autoscaling:us-west-2:123456789012:autoScalingGroup:1111aaaa-22bb-33cc-44dd-5555eeee66ff:autoScalingGroupName/MegaFrogServers_NA",
    "BalancingStrategy": "SPOT_ONLY",
    "CreationTime": 1496365885.44,
    "GameServerGroupArn": "arn:aws:gamelift:us-west-2::GameServerGroup/MegaFrogServers_NA",
    "GameServerGroupName": " MegaFrogServers_NA",
    "GameServerProtectionPolicy": "NO_PROTECTION",
    "InstanceDefinitions": [
      {
        "InstanceType": "c4.xlarge",
        "WeightedCapacity": "1"
      },
      {
        "InstanceType": "c5.xlarge",
        "WeightedCapacity": "1"
      }
    ],
  },
}
```

```
"LastUpdatedTime": 1496365885.44,  
"RoleArn": "arn:aws:iam:123456789012::role/GameLiftGsgRole",  
"Status": "ACTIVE",  
"StatusReason": "",  
"SuspendedActions": []  
}
```

Change the balancing strategy for a game server group

This example changes the game server group's balancing strategy from preferring Spot instances to using On-Demand instances only. In this example scenario, the game developer wants to eliminate the possibility of game session interruptions during an upcoming highly-anticipated content launch. Later, once player demand settles, the balancing strategy can be updated again to use Spot instances for cost savings.

HTTP requests are authenticated using an [AWS Signature Version 4](#) signature in the Authorization header field.

Sample Request

```
{  
  "GameServerGroupName": "MegaFrogServers_NA",  
  "BalancingStrategy": "ON_DEMAND_ONLY"  
}
```

CLI command:

```
aws gamelift update-game-server-group \  
  --game-server-group MegaFrogServers_NA \  
  --balancing-strategy ON_DEMAND_ONLY
```

Sample Response

```
{  
  "GameServerGroup": {  
    "AutoScalingGroupArn": "arn:aws:autoscaling:us-west-2:123456789012:autoScalingGroup:1111aaaa-22bb-33cc-44dd-5555eeee66ff:autoScalingGroupName/MegaFrogServers_NA",  
    "BalancingStrategy": "ON_DEMAND_ONLY",  
    "CreationTime": 1496365885.44,  
    "GameServerGroupArn": "arn:aws:gamelift:us-west-2::GameServerGroup/MegaFrogServers_NA",  
    "GameServerGroupName": "MegaFrogServers_NA",  
    "RoleArn": "arn:aws:iam:123456789012::role/GameLiftGsgRole",  
    "Status": "ACTIVE",  
    "StatusReason": "",  
    "SuspendedActions": []  
  }  
}
```

```
"GameServerGroupName": " MegaFrogServers_NA",
"GameServerProtectionPolicy": "FULL_PROTECTION",
"InstanceDefinitions": [
  {
    "InstanceType": "c4.xlarge",
    "WeightedCapacity": "1"
  },
  {
    "InstanceType": "c5.xlarge",
    "WeightedCapacity": "1"
  }
],
"LastUpdatedTime": 1496365885.44,
"RoleArn": "arn:aws:iam:123456789012::role/GameLiftGsgRole",
"Status": "ACTIVE",
"StatusReason": "",
"SuspendedActions": []
}
```

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](#)

UpdateGameSession

This API works with the following fleet types: EC2, Anywhere, Container

Updates the mutable properties of a game session.

To update a game session, specify the game session ID and the values you want to change.

If successful, the updated GameSession object is returned.

[All APIs by task](#)

Request Syntax

```
{
  "GameProperties": [
    {
      "Key": "string",
      "Value": "string"
    }
  ],
  "GameSessionId": "string",
  "MaximumPlayerSessionCount": number,
  "Name": "string",
  "PlayerSessionCreationPolicy": "string",
  "ProtectionPolicy": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

GameSessionId

An identifier for the game session that is unique across all regions to update. The value is always a full ARN in the following format: For Home Region game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>`. For Remote Location game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<location>/<ID string>`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `[a-zA-Z0-9:/-]+`

Required: Yes

GameProperties

A set of key-value pairs that can store custom data in a game session. For example: `{"Key": "difficulty", "Value": "novice"}`. You can use this parameter to modify game properties in an active game session. This action adds new properties and modifies existing properties. There is no way to delete properties. For an example, see [Update the value of a game property](#).

Note

- Avoid using periods (".") in property keys if you plan to search for game sessions by properties. Property keys containing periods cannot be searched and will be filtered out from search results due to search index limitations.
- If you use `SearchGameSessions` API, there is a limit of 500 game property keys across all game sessions and all fleets per region. If the limit is exceeded, there will potentially be game session entries missing from `SearchGameSessions` API results.

Type: Array of [GameProperty](#) objects

Array Members: Maximum number of 16 items.

Required: No

MaximumPlayerSessionCount

The maximum number of players that can be connected simultaneously to the game session.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

Name

A descriptive label that is associated with a game session. Session names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

PlayerSessionCreationPolicy

A policy that determines whether the game session is accepting new players.

Type: String

Valid Values: ACCEPT_ALL | DENY_ALL

Required: No

ProtectionPolicy

Game session protection policy to apply to this game session only.

- NoProtection -- The game session can be terminated during a scale-down event.
- FullProtection -- If the game session is in an ACTIVE status, it cannot be terminated during a scale-down event.

Type: String

Valid Values: NoProtection | FullProtection

Required: No

Response Syntax

```
{  
  "GameSession": {  
    "ComputeName": "string",
```

```
"CreationTime": number,
"CreatorId": "string",
"CurrentPlayerSessionCount": number,
"DnsName": "string",
"FleetArn": "string",
"FleetId": "string",
"GameProperties": [
  {
    "Key": "string",
    "Value": "string"
  }
],
"GameSessionData": "string",
"GameSessionId": "string",
"IpAddress": "string",
"Location": "string",
"MatchmakerData": "string",
"MaximumPlayerSessionCount": number,
"Name": "string",
"PlayerGatewayStatus": "string",
"PlayerSessionCreationPolicy": "string",
"Port": number,
"Status": "string",
"StatusReason": "string",
"TerminationTime": number
}
```

Response Elements

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

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

GameSession

The updated game session properties.

Type: [GameSession](#) object

Errors

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

ConflictException

The requested operation would cause a conflict with the current state of a service resource associated with the request. Resolve the conflict before retrying this request.

HTTP Status Code: 400

InternalServiceException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidGameSessionStatusException

The requested operation would cause a conflict with the current state of a resource associated with the request and/or the game instance. Resolve the conflict before retrying.

HTTP Status Code: 400

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

NotReadyException

The operation failed because Amazon GameLift Servers has not yet finished validating this compute. We recommend attempting 8 to 10 retries over 3 to 5 minutes with [exponential backoffs and jitter](#).

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

UpdateGameSessionQueue

This API works with the following fleet types: EC2, Anywhere, Container

Updates the configuration of a game session queue, which determines how the queue processes new game session requests. To update settings, specify the queue name to be updated and provide the new settings. When updating destinations, provide a complete list of destinations.

Learn more

[Using Multi-Region Queues](#)

Request Syntax

```
{
  "CustomEventData": "string",
  "Destinations": [
    {
      "DestinationArn": "string"
    }
  ],
  "FilterConfiguration": {
    "AllowedLocations": [ "string" ]
  },
  "Name": "string",
  "NotificationTarget": "string",
  "PlayerLatencyPolicies": [
    {
      "MaximumIndividualPlayerLatencyMilliseconds": number,
      "PolicyDurationSeconds": number
    }
  ],
  "PriorityConfiguration": {
    "LocationOrder": [ "string" ],
    "PriorityOrder": [ "string" ]
  },
  "TimeoutInSeconds": number
}
```

Request Parameters

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

The request accepts the following data in JSON format.

 Note

In the following list, the required parameters are described first.

Name

A descriptive label that is associated with game session queue. Queue names must be unique within each Region. You can use either the queue ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-]+|^arn:.*:gamesessionqueue\[a-zA-Z0-9-]+\`

Required: Yes

CustomEventData

Information to be added to all events that are related to this game session queue.

Type: String

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

Pattern: `[\s\S]*`

Required: No

Destinations

A list of fleets and/or fleet aliases that can be used to fulfill game session placement requests in the queue. Destinations are identified by either a fleet ARN or a fleet alias ARN, and are listed in order of placement preference. When updating this list, provide a complete list of destinations.

Type: Array of [GameSessionQueueDestination](#) objects

Required: No

FilterConfiguration

A list of locations where a queue is allowed to place new game sessions. Locations are specified in the form of AWS Region codes, such as `us-west-2`. If this parameter is not set, game

sessions can be placed in any queue location. To remove an existing filter configuration, pass in an empty set.

Type: [FilterConfiguration](#) object

Required: No

NotificationTarget

An SNS topic ARN that is set up to receive game session placement notifications. See [Setting up notifications for game session placement](#).

Type: String

Length Constraints: Minimum length of 0. Maximum length of 300.

Pattern: `[a-zA-Z0-9:_-]*(\.\fifo)?`

Required: No

PlayerLatencyPolicies

A set of policies that enforce a sliding cap on player latency when processing game sessions placement requests. Use multiple policies to gradually relax the cap over time if Amazon GameLift Servers can't make a placement. Policies are evaluated in order starting with the lowest maximum latency value. When updating policies, provide a complete collection of policies.

Type: Array of [PlayerLatencyPolicy](#) objects

Required: No

PriorityConfiguration

Custom settings to use when prioritizing destinations and locations for game session placements. This configuration replaces the FleetIQ default prioritization process. Priority types that are not explicitly named will be automatically applied at the end of the prioritization process. To remove an existing priority configuration, pass in an empty set.

Type: [PriorityConfiguration](#) object

Required: No

TimeoutInSeconds

The maximum time, in seconds, that a new game session placement request remains in the queue. When a request exceeds this time, the game session placement changes to a TIMED_OUT status.

Note

The minimum value is 10 and the maximum value is 600.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

Response Syntax

```
{
  "GameSessionQueue": {
    "CustomEventData": "string",
    "Destinations": [
      {
        "DestinationArn": "string"
      }
    ],
    "FilterConfiguration": {
      "AllowedLocations": [ "string" ]
    },
    "GameSessionQueueArn": "string",
    "Name": "string",
    "NotificationTarget": "string",
    "PlayerLatencyPolicies": [
      {
        "MaximumIndividualPlayerLatencyMilliseconds": number,
        "PolicyDurationSeconds": number
      }
    ],
    "PriorityConfiguration": {
      "LocationOrder": [ "string" ],
      "PriorityOrder": [ "string" ]
    }
  }
}
```

```
    },  
    "TimeoutInSeconds": number  
  }  
}
```

Response Elements

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

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

[GameSessionQueue](#)

An object that describes the newly updated game session queue.

Type: [GameSessionQueue](#) object

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

UpdateMatchmakingConfiguration

This API works with the following fleet types: EC2, Anywhere, Container

Updates settings for a FlexMatch matchmaking configuration. These changes affect all matches and game sessions that are created after the update. To update settings, specify the configuration name to be updated and provide the new settings.

Learn more

[Design a FlexMatch matchmaker](#)

Request Syntax

```
{
  "AcceptanceRequired": boolean,
  "AcceptanceTimeoutSeconds": number,
  "AdditionalPlayerCount": number,
  "BackfillMode": "string",
  "CustomEventData": "string",
  "Description": "string",
  "FlexMatchMode": "string",
  "GameProperties": [
    {
      "Key": "string",
      "Value": "string"
    }
  ],
  "GameSessionData": "string",
  "GameSessionQueueArns": [ "string" ],
  "Name": "string",
  "NotificationTarget": "string",
  "RequestTimeoutSeconds": number,
  "RuleSetName": "string"
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

Name

A unique identifier for the matchmaking configuration to update. You can use either the configuration name or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]*|^arn:.*:matchmakingconfiguration\[a-zA-Z0-9-\.]*`

Required: Yes

AcceptanceRequired

A flag that indicates whether a match that was created with this configuration must be accepted by the matched players. To require acceptance, set to TRUE. With this option enabled, matchmaking tickets use the status `REQUIRES_ACCEPTANCE` to indicate when a completed potential match is waiting for player acceptance.

Type: Boolean

Required: No

AcceptanceTimeoutSeconds

The length of time (in seconds) to wait for players to accept a proposed match, if acceptance is required.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 600.

Required: No

AdditionalPlayerCount

The number of player slots in a match to keep open for future players. For example, if the configuration's rule set specifies a match for a single 12-person team, and the additional player

count is set to 2, only 10 players are selected for the match. This parameter is not used if `FlexMatchMode` is set to `STANDALONE`.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

BackfillMode

The method that is used to backfill game sessions created with this matchmaking configuration. Specify `MANUAL` when your game manages backfill requests manually or does not use the match backfill feature. Specify `AUTOMATIC` to have GameLift create a match backfill request whenever a game session has one or more open slots. Learn more about manual and automatic backfill in [Backfill Existing Games with FlexMatch](#). Automatic backfill is not available when `FlexMatchMode` is set to `STANDALONE`.

Type: String

Valid Values: `AUTOMATIC` | `MANUAL`

Required: No

CustomEventData

Information to add to all events related to the matchmaking configuration.

Type: String

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

Required: No

Description

A description for the matchmaking configuration.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

FlexMatchMode

Indicates whether this matchmaking configuration is being used with Amazon GameLift Servers hosting or as a standalone matchmaking solution.

- **STANDALONE** - FlexMatch forms matches and returns match information, including players and team assignments, in a [MatchmakingSucceeded](#) event.
- **WITH_QUEUE** - FlexMatch forms matches and uses the specified Amazon GameLift Servers queue to start a game session for the match.

Type: String

Valid Values: STANDALONE | WITH_QUEUE

Required: No

GameProperties

A set of key-value pairs that can store custom data in a game session. For example: {"Key": "difficulty", "Value": "novice"}. This information is added to the new `GameSession` object that is created for a successful match. This parameter is not used if `FlexMatchMode` is set to `STANDALONE`.

Note

- Avoid using periods (".") in property keys if you plan to search for game sessions by properties. Property keys containing periods cannot be searched and will be filtered out from search results due to search index limitations.
- If you use `SearchGameSessions` API, there is a limit of 500 game property keys across all game sessions and all fleets per region. If the limit is exceeded, there will potentially be game session entries missing from `SearchGameSessions` API results.

Type: Array of [GameProperty](#) objects

Array Members: Maximum number of 16 items.

Required: No

GameSessionData

A set of custom game session properties, formatted as a single string value. This data is passed to a game server process with a request to start a new game session. For more information,

see [Start a game session](#). This information is added to the game session that is created for a successful match. This parameter is not used if `FlexMatchMode` is set to `STANDALONE`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 4096.

Required: No

[**GameSessionQueueArns**](#)

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers game session queue resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::gamesessionqueue/<queue name>`. Queues can be located in any Region. Queues are used to start new Amazon GameLift Servers-hosted game sessions for matches that are created with this matchmaking configuration. If `FlexMatchMode` is set to `STANDALONE`, do not set this parameter.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `[a-zA-Z0-9:/-]+`

Required: No

[**NotificationTarget**](#)

An SNS topic ARN that is set up to receive matchmaking notifications. See [Setting up notifications for matchmaking](#) for more information.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 300.

Pattern: `[a-zA-Z0-9:_/-]*(.fifo)?`

Required: No

[**RequestTimeoutSeconds**](#)

The maximum duration, in seconds, that a matchmaking ticket can remain in process before timing out. Requests that fail due to timing out can be resubmitted as needed.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 43200.

Required: No

RuleSetName

A unique identifier for the matchmaking rule set to use with this configuration. You can use either the rule set name or ARN value. A matchmaking configuration can only use rule sets that are defined in the same Region.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9-\.]*|^arn:.*:matchmakingruleset\/[a-zA-Z0-9-\.]*`

Required: No

Response Syntax

```
{
  "Configuration": {
    "AcceptanceRequired": boolean,
    "AcceptanceTimeoutSeconds": number,
    "AdditionalPlayerCount": number,
    "BackfillMode": "string",
    "ConfigurationArn": "string",
    "CreationTime": number,
    "CustomEventData": "string",
    "Description": "string",
    "FlexMatchMode": "string",
    "GameProperties": [
      {
        "Key": "string",
        "Value": "string"
      }
    ],
    "GameSessionData": "string",
    "GameSessionQueueArns": [ "string" ],
    "Name": "string",
    "NotificationTarget": "string",
    "RequestTimeoutSeconds": number,
    "RuleSetArn": "string",
```

```
    "RuleSetName": "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.

Configuration

The updated matchmaking configuration.

Type: [MatchmakingConfiguration](#) object

Errors

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

InternalServerError

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

UpdateRuntimeConfiguration

This API works with the following fleet types: EC2

Updates the runtime configuration for the specified fleet. The runtime configuration tells Amazon GameLift Servers how to launch server processes on computes in managed EC2 and Anywhere fleets. You can update a fleet's runtime configuration at any time after the fleet is created; it does not need to be in ACTIVE status.

To update runtime configuration, specify the fleet ID and provide a RuntimeConfiguration with an updated set of server process configurations.

If successful, the fleet's runtime configuration settings are updated. Fleet computes that run game server processes regularly check for and receive updated runtime configurations. The computes immediately take action to comply with the new configuration by launching new server processes or by not replacing existing processes when they shut down. Updating a fleet's runtime configuration never affects existing server processes.

Learn more

[Setting up Amazon GameLift Servers fleets](#)

Request Syntax

```
{
  "FleetId": "string",
  "RuntimeConfiguration": {
    "GameSessionActivationTimeoutSeconds": number,
    "MaxConcurrentGameSessionActivations": number,
    "ServerProcesses": [
      {
        "ConcurrentExecutions": number,
        "LaunchPath": "string",
        "Parameters": "string"
      }
    ]
  }
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet to update runtime configuration for. You can use either the fleet ID or ARN value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\_]+\$|^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\-\_]+\$`

Required: Yes

RuntimeConfiguration

Instructions for launching server processes on fleet computes. Server processes run either a custom game build executable or a Amazon GameLift Servers Realtime script. The runtime configuration lists the types of server processes to run, how to launch them, and the number of processes to run concurrently.

Type: [RuntimeConfiguration](#) object

Required: Yes

Response Syntax

```
{  
  "RuntimeConfiguration": {
```

```
"GameSessionActivationTimeoutSeconds": number,
"MaxConcurrentGameSessionActivations": number,
"ServerProcesses": [
  {
    "ConcurrentExecutions": number,
    "LaunchPath": "string",
    "Parameters": "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.

RuntimeConfiguration

The runtime configuration currently in use by computes in the fleet. If the update is successful, all property changes are shown.

Type: [RuntimeConfiguration](#) object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidFleetStatusException

The requested operation would cause a conflict with the current state of a resource associated with the request and/or the fleet. Resolve the conflict before retrying.

HTTP Status Code: 400

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

LimitExceededException

The requested operation would cause the resource to exceed the allowed service limit. Resolve the issue before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

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](#)

UpdateScript

This API works with the following fleet types: EC2

Updates Realtime script metadata and content.

To update script metadata, specify the script ID and provide updated name and/or version values.

To update script content, provide an updated zip file by pointing to either a local file or an Amazon S3 bucket location. You can use either method regardless of how the original script was uploaded. Use the *Version* parameter to track updates to the script.

If the call is successful, the updated metadata is stored in the script record and a revised script is uploaded to the Amazon GameLift Servers service. Once the script is updated and acquired by a fleet instance, the new version is used for all new game sessions.

Learn more

[Amazon GameLift Servers Amazon GameLift Servers Realtime](#)

Related actions

[All APIs by task](#)

Request Syntax

```
{
  "Name": "string",
  "ScriptId": "string",
  "StorageLocation": {
    "Bucket": "string",
    "Key": "string",
    "ObjectVersion": "string",
    "RoleArn": "string"
  },
  "Version": "string",
  "ZipFile": blob
}
```

Request Parameters

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

The request accepts the following data in JSON format.

 Note

In the following list, the required parameters are described first.

ScriptId

A unique identifier for the Realtime script to update. You can use either the script ID or ARN value.

Type: String

Pattern: `^script-\S+|^arn:.*:script\/script-\S+`

Required: Yes

Name

A descriptive label that is associated with a script. Script names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

StorageLocation

The location of the Amazon S3 bucket where a zipped file containing your Realtime scripts is stored. The storage location must specify the Amazon S3 bucket name, the zip file name (the "key"), and a role ARN that allows Amazon GameLift Servers to access the Amazon S3 storage location. The S3 bucket must be in the same Region where you want to create a new script. By default, Amazon GameLift Servers uploads the latest version of the zip file; if you have S3 object versioning turned on, you can use the `ObjectVersion` parameter to specify an earlier version.

Type: [S3Location](#) object

Required: No

Version

Version information that is associated with a build or script. Version strings do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

ZipFile

A data object containing your Realtime scripts and dependencies as a zip file. The zip file can have one or multiple files. Maximum size of a zip file is 5 MB.

When using the AWS CLI tool to create a script, this parameter is set to the zip file name. It must be prepended with the string "fileb://" to indicate that the file data is a binary object. For example: `--zip-file fileb://myRealtimeScript.zip`.

Type: Base64-encoded binary data object

Length Constraints: Maximum length of 5000000.

Required: No

Response Syntax

```
{
  "Script": {
    "CreationTime": number,
    "Name": "string",
    "NodeJsVersion": "string",
    "ScriptArn": "string",
    "ScriptId": "string",
    "SizeOnDisk": number,
    "StorageLocation": {
      "Bucket": "string",
      "Key": "string",
      "ObjectVersion": "string",
      "RoleArn": "string"
    },
  },
  "Version": "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.

[Script](#)

The newly created script record with a unique script ID. The new script's storage location reflects an Amazon S3 location: (1) If the script was uploaded from an S3 bucket under your account, the storage location reflects the information that was provided in the *CreateScript* request; (2) If the script file was uploaded from a local zip file, the storage location reflects an S3 location controls by the Amazon GameLift Servers service.

Type: [Script](#) object

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

NotFoundException

The requested resource was not found. The resource was either not created yet or deleted.

HTTP Status Code: 400

UnauthorizedException

The client failed authentication. Clients should not retry such requests.

HTTP Status Code: 400

Examples

Update a script from a local zip file

This example updates the Realtime script with a zip file that is stored locally.

Sample Request

```
{
  "ScriptId": "script-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "Version": "1.0.2",
  "ZipFile": <zip file data>
}
```

CLI syntax:

```
aws gamelift create-script
--script-id "script-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
--script-version "1.0.2"
--zip-file fileb://myrealtime_script.zip
```

Sample Response

```
{
  "Script": {
    "CreationTime": 1496708916.18,
    "Name": "My_Realtime_Script_1",
    "NodeJsVersion": "24.x",
    "ScriptArn": "arn:aws:gamelift:us-west-2::script/script-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "ScriptId": "script-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "SizeOnDisk": 10000,
    "StorageLocation": {
      "Bucket": "prod-gamescale-scripts-us-west-2",
      "Key": "123456789012/script-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
    }
  },
}
```

```
    "Version": "1.0.2"
  }
}
```

Update a script with a file in Amazon S3

This example updates the Realtime server script with a zip file that is stored in an S3 account.

Sample Request

```
{
  "ScriptId": "script-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "Version": "1.0.2",
  "StorageLocation": {
    "Bucket": "my_realtime_script_files",
    "Key": "myRealtimeScript.zip",
    "RoleArn": "arn:aws:iam::111122223333:role/GameLiftAccess"
  }
}
```

CLI syntax:

```
aws gamelift update-script
--script-id "script-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
--script-version "1.0.2"
--storage-location
"Bucket=my_realtime_script_files,
Key=myRealtimeScript.zip,
RoleArn=arn:aws:iam::123456789012:role/GameLiftAccess"
```

Sample Response

```
{
  "Script": {
    "CreationTime": 1496708916.18,
    "Name": "My_Realtime_Script_1",
    "NodeJsVersion": "24.x",
    "ScriptArn": "arn:aws:gamelift:us-west-2::script/
script-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "ScriptId": "script-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
    "SizeOnDisk": 0,
    "StorageLocation": {
      "Bucket": "my_realtime_script_files",
```

```
        "Key": "myRealtimeScript.zip"
        "RoleArn": "arn:aws:iam::111122223333:role/GameLiftAccess"
        "ObjectVersion": null
    },
    "Version": "1.0.2"
}
```

See Also

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

- [AWS Command Line Interface 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](#)

ValidateMatchmakingRuleSet

This API works with the following fleet types: EC2, Anywhere, Container

Validates the syntax of a matchmaking rule or rule set. This operation checks that the rule set is using syntactically correct JSON and that it conforms to allowed property expressions. To validate syntax, provide a rule set JSON string.

Learn more

- [Build a rule set](#)

Request Syntax

```
{  
  "RuleSetBody": "string"  
}
```

Request Parameters

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

The request accepts the following data in JSON format.

Note

In the following list, the required parameters are described first.

[RuleSetBody](#)

A collection of matchmaking rules to validate, formatted as a JSON string.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 65535.

Required: Yes

Response Syntax

```
{  
  "Valid": boolean  
}
```

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.

Valid

A response indicating whether the rule set is valid.

Type: Boolean

Errors

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

InternalServerErrorException

The service encountered an unrecoverable internal failure while processing the request. Clients can retry such requests immediately or after a waiting period.

HTTP Status Code: 500

InvalidRequestException

One or more parameter values in the request are invalid. Correct the invalid parameter values before retrying.

HTTP Status Code: 400

UnsupportedRegionException

The requested operation is not supported in the Region specified.

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](#)

Data Types

The Amazon GameLift 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:

- [Alias](#)
- [AnywhereConfiguration](#)
- [AttributeValue](#)
- [AwsCredentials](#)
- [Build](#)
- [CertificateConfiguration](#)
- [ClaimFilterOption](#)
- [Compute](#)
- [ConnectionPortRange](#)
- [ContainerAttribute](#)
- [ContainerDependency](#)
- [ContainerEnvironment](#)
- [ContainerFleet](#)
- [ContainerFleetLocationAttributes](#)
- [ContainerGroupDefinition](#)
- [ContainerGroupPortMapping](#)
- [ContainerHealthCheck](#)
- [ContainerIdentifier](#)
- [ContainerMountPoint](#)
- [ContainerPortConfiguration](#)

- [ContainerPortMapping](#)
- [ContainerPortRange](#)
- [DeploymentConfiguration](#)
- [DeploymentDetails](#)
- [DesiredPlayerSession](#)
- [EC2InstanceCounts](#)
- [EC2InstanceLimit](#)
- [Event](#)
- [FilterConfiguration](#)
- [FleetAttributes](#)
- [FleetCapacity](#)
- [FleetDeployment](#)
- [FleetUtilization](#)
- [GameProperty](#)
- [GameServer](#)
- [GameServerContainerDefinition](#)
- [GameServerContainerDefinitionInput](#)
- [GameServerContainerGroupCounts](#)
- [GameServerGroup](#)
- [GameServerGroupAutoScalingPolicy](#)
- [GameServerInstance](#)
- [GameSession](#)
- [GameSessionConnectionInfo](#)
- [GameSessionCreationLimitPolicy](#)
- [GameSessionDetail](#)
- [GameSessionPlacement](#)
- [GameSessionQueue](#)
- [GameSessionQueueDestination](#)
- [Instance](#)
- [InstanceAccess](#)

- [InstanceCredentials](#)
- [InstanceDefinition](#)
- [IpPermission](#)
- [LaunchTemplateSpecification](#)
- [LinuxCapabilities](#)
- [LocationalDeployment](#)
- [LocationAttributes](#)
- [LocationConfiguration](#)
- [LocationModel](#)
- [LocationState](#)
- [LogConfiguration](#)
- [ManagedCapacityConfiguration](#)
- [MatchedPlayerSession](#)
- [MatchmakingConfiguration](#)
- [MatchmakingRuleSet](#)
- [MatchmakingTicket](#)
- [PingBeacon](#)
- [PlacedPlayerSession](#)
- [Player](#)
- [PlayerConnectionDetail](#)
- [PlayerConnectionEndpoint](#)
- [PlayerGatewayConfiguration](#)
- [PlayerLatency](#)
- [PlayerLatencyPolicy](#)
- [PlayerSession](#)
- [PriorityConfiguration](#)
- [PriorityConfigurationOverride](#)
- [ResourceCreationLimitPolicy](#)
- [RoutingStrategy](#)
- [RuntimeConfiguration](#)

- [S3Location](#)
- [ScalingPolicy](#)
- [Script](#)
- [ServerProcess](#)
- [SupportContainerDefinition](#)
- [SupportContainerDefinitionInput](#)
- [Tag](#)
- [TargetConfiguration](#)
- [TargetTrackingConfiguration](#)
- [UDPEndpoint](#)
- [VpcPeeringAuthorization](#)
- [VpcPeeringConnection](#)
- [VpcPeeringConnectionStatus](#)

Alias

Properties that describe an alias resource.

Related actions

[All APIs by task](#)

Contents

Note

In the following list, the required parameters are described first.

AliasArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers alias resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::alias/alias-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`. In a GameLift alias ARN, the resource ID matches the alias ID value.

Type: String

Pattern: `^arn:.*:alias\/alias-\S+`

Required: No

AliasId

A unique identifier for the alias. Alias IDs are unique within a Region.

Type: String

Pattern: `^alias-\S+`

Required: No

CreationTime

A time stamp indicating when this data object was created. Format is a number expressed in Unix time as milliseconds (for example `"1469498468.057"`).

Type: Timestamp

Required: No

Description

A human-readable description of an alias.

Type: String

Required: No

LastUpdatedTime

The time that this data object was last modified. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

Name

A descriptive label that is associated with an alias. Alias names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `.*\S.*`

Required: No

RoutingStrategy

The routing configuration, including routing type and fleet target, for the alias.

Type: [RoutingStrategy](#) 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](#)

AnywhereConfiguration

Amazon GameLift Servers configuration options for your Anywhere fleets.

Contents

Note

In the following list, the required parameters are described first.

Cost

The cost to run your fleet per hour. Amazon GameLift Servers uses the provided cost of your fleet to balance usage in queues. For more information about queues, see [Setting up queues](#) in the *Amazon GameLift Servers Developer Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 11.

Pattern: `^\d{1,5}(:\.\d{1,5})?$`

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](#)

AttributeValue

Values for use in player attribute key-value pairs. This object lets you specify an attribute value using any of the valid data types: string, number, string array, or data map. Each `AttributeValue` object can use only one of the available properties.

Contents

Note

In the following list, the required parameters are described first.

N

For number values, expressed as double.

Type: Double

Required: No

S

For single string values. Maximum string length is 100 characters.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Required: No

SDM

For a map of up to 10 data type:value pairs. Maximum length for each string value is 100 characters.

Type: String to double map

Key Length Constraints: Minimum length of 1. Maximum length of 100.

Required: No

SL

For a list of up to 100 strings. Maximum length for each string is 100 characters. Duplicate values are not recognized; all occurrences of the repeated value after the first of a repeated value are ignored.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 100.

Required: No

See Also

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

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

AwsCredentials

AWS account security credentials that allow interactions with Amazon GameLift Servers resources. The credentials are temporary and valid for a limited time span. You can request fresh credentials at any time.

AWS security credentials consist of three parts: an access key ID, a secret access key, and a session token. You must use all three parts together to authenticate your access requests.

You need AWS credentials for the following tasks:

- To upload a game server build directly to Amazon GameLift Servers S3 storage using `CreateBuild`. To get access for this task, call https://docs.aws.amazon.com/gamelift/latest/apireference/API_RequestUploadCredentials.html.
- To remotely connect to an active Amazon GameLift Servers fleet instances. To get remote access, call https://docs.aws.amazon.com/gamelift/latest/apireference/API_GetComputeAccess.html.

Contents

Note

In the following list, the required parameters are described first.

AccessKeyId

The access key ID that identifies the temporary security credentials.

Type: String

Length Constraints: Minimum length of 1.

Required: No

SecretAccessKey

The secret access key that can be used to sign requests.

Type: String

Length Constraints: Minimum length of 1.

Required: No

SessionToken

The token that users must pass to the service API to use the temporary credentials.

Type: String

Length Constraints: Minimum length 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](#)

Build

Properties describing a custom game build.

[All APIs by task](#)

Contents

Note

In the following list, the required parameters are described first.

BuildArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers build resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::build/build-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`. In a GameLift build ARN, the resource ID matches the *BuildId* value.

Type: String

Pattern: `^arn:.*:build\/build-\\S+`

Required: No

BuildId

A unique identifier for the build.

Type: String

Pattern: `^build-\\S+`

Required: No

CreationTime

A time stamp indicating when this data object was created. Format is a number expressed in Unix time as milliseconds (for example `"1469498468.057"`).

Type: Timestamp

Required: No

Name

A descriptive label that is associated with a build. Build names do not need to be unique. It can be set using [CreateBuild](#) or [UpdateBuild](#).

Type: String

Required: No

OperatingSystem

Operating system that the game server binaries are built to run on. This value determines the type of fleet resources that you can use for this build.

Note

Amazon Linux 2 (AL2) will reach end of support on 6/30/2026. See more details in the [Amazon Linux 2 FAQs](#). For game servers that are hosted on AL2 and use server SDK version 4.x for Amazon GameLift Servers, first update the game server build to server SDK 5.x, and then deploy to AL2023 instances. See [Migrate to server SDK version 5](#).

Type: String

Valid Values: WINDOWS_2012 | AMAZON_LINUX | AMAZON_LINUX_2 | WINDOWS_2016 | AMAZON_LINUX_2023 | WINDOWS_2022

Required: No

ServerSdkVersion

The Amazon GameLift Servers Server SDK version used to develop your game server.

Type: String

Length Constraints: Maximum length of 128.

Pattern: ^\d+\.\d+\.\d+\$

Required: No

SizeOnDisk

File size of the uploaded game build, expressed in bytes. When the build status is `INITIALIZED` or when using a custom Amazon S3 storage location, this value is 0.

Type: Long

Valid Range: Minimum value of 0.

Required: No

Status

Current status of the build.

Possible build statuses include the following:

- **INITIALIZED** -- A new build has been defined, but no files have been uploaded. You cannot create fleets for builds that are in this status. When a build is successfully created, the build status is set to this value.
- **READY** -- The game build has been successfully uploaded. You can now create new fleets for this build.
- **FAILED** -- The game build upload failed. You cannot create new fleets for this build.

Type: String

Valid Values: `INITIALIZED` | `READY` | `FAILED`

Required: No

Version

Version information that is associated with a build or script. Version strings do not need to be unique.

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](#)

CertificateConfiguration

Determines whether a TLS/SSL certificate is generated for a fleet. This feature must be enabled when creating the fleet. All instances in a fleet share the same certificate. The certificate can be retrieved by calling the [Amazon GameLift Servers Server SDK](#) operation `GetInstanceCertificate`.

Contents

Note

In the following list, the required parameters are described first.

CertificateType

Indicates whether a TLS/SSL certificate is generated for a fleet.

Valid values include:

- **GENERATED** - Generate a TLS/SSL certificate for this fleet.
- **DISABLED** - (default) Do not generate a TLS/SSL certificate for this fleet.

Type: String

Valid Values: DISABLED | GENERATED

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](#)

ClaimFilterOption

Filters which game servers may be claimed when calling `ClaimGameServer`.

Contents

Note

In the following list, the required parameters are described first.

InstanceStatuses

List of instance statuses that game servers may be claimed on. If provided, the list must contain the `ACTIVE` status.

Type: Array of strings

Valid Values: `ACTIVE` | `DRAINING`

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](#)

Compute

An Amazon GameLift Servers compute resource for hosting your game servers. Computes in an Amazon GameLift Servers fleet differs depending on the fleet's compute type property as follows:

- For managed EC2 fleets, a compute is an EC2 instance.
- For Anywhere fleets, a compute is a computing resource that you provide and is registered to the fleet.
- For managed container fleets, a compute is a container group in the fleet.

Contents

Note

In the following list, the required parameters are described first.

ComputeArn

The ARN that is assigned to a compute resource and uniquely identifies it. ARNs are unique across locations. Instances in managed EC2 fleets are not assigned a Compute ARN.

Type: String

Length Constraints: Maximum length of 1024.

Pattern: `^arn:.*:compute[/[a-zA-Z0-9\-\]]+(\[/[a-zA-Z0-9\-\]]+)?`

Required: No

ComputeName

A descriptive label for the compute resource. For instances in a managed EC2 fleet, the compute name is the same value as the InstanceId ID.

Type: String

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

Pattern: `^[a-zA-Z0-9\-\]]+(\[/[a-zA-Z0-9\-\]]+)?`

Required: No

ComputeStatus

Current status of the compute. A compute must have an ACTIVE status to host game sessions. Valid values include PENDING, ACTIVE, TERMINATING, and IMPAIRED.

Note

While the ComputeStatus enum type is valid for Container based servers, the result may also include other non-enumerated string values such as "Active" for fleets which are not Container-based.

Type: String

Valid Values: PENDING | ACTIVE | TERMINATING | IMPAIRED

Required: No

ContainerAttributes

A set of attributes for each container in the compute.

Type: Array of [ContainerAttribute](#) objects

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

Required: No

CreationTime

A time stamp indicating when this data object was created. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

DnsName

The DNS name of a compute resource. Amazon GameLift Servers requires a DNS name or IP address for a compute.

Type: String

Required: No

FleetArn

The Amazon Resource Name (ARN) of the fleet that the compute belongs to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\-]+\`

Required: No

FleetId

A unique identifier for the fleet that the compute belongs to.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\-]+\`

Required: No

GameLiftAgentEndpoint

The endpoint of the Amazon GameLift Servers Agent.

Type: String

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

Required: No

GameLiftServiceSdkEndpoint

The Amazon GameLift Servers SDK endpoint connection for a registered compute resource in an Anywhere fleet. The game servers on the compute use this endpoint to connect to the Amazon GameLift Servers service.

Type: String

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

Required: No

GameServerContainerGroupDefinitionArn

The game server container group definition for the compute.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^[a-zA-Z0-9\-$]|^arn:.*:containergroupdefinition\[a-zA-Z0-9\-$]+(:[0-9]+)?$`

Required: No

InstanceID

The InstanceID of the EC2 instance that is hosting the compute.

Type: String

Pattern: `[a-zA-Z0-9\.-]+`

Required: No

IpAddress

The IP address of a compute resource. Amazon GameLift Servers requires a DNS name or IP address for a compute.

Type: String

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

Pattern: `^[0-9A-Fa-f\:\.]+`

Required: No

Location

The name of the custom location you added to the fleet that this compute resource resides in.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-$]+`

Required: No

OperatingSystem

The type of operating system on the compute resource.

Note

Amazon Linux 2 (AL2) will reach end of support on 6/30/2026. See more details in the [Amazon Linux 2 FAQs](#). For game servers that are hosted on AL2 and use server SDK version 4.x for Amazon GameLift Servers, first update the game server build to server SDK 5.x, and then deploy to AL2023 instances. See [Migrate to server SDK version 5](#).

Type: String

Valid Values: WINDOWS_2012 | AMAZON_LINUX | AMAZON_LINUX_2 | WINDOWS_2016 | AMAZON_LINUX_2023 | WINDOWS_2022

Required: No

Type

The Amazon EC2 instance type that the fleet uses. For registered computes in an Amazon GameLift Servers Anywhere fleet, this property is empty.

Type: String

Valid Values: t2.micro | t2.small | t2.medium | t2.large | c3.large | c3.xlarge | c3.2xlarge | c3.4xlarge | c3.8xlarge | c4.large | c4.xlarge | c4.2xlarge | c4.4xlarge | c4.8xlarge | c5.large | c5.xlarge | c5.2xlarge | c5.4xlarge | c5.9xlarge | c5.12xlarge | c5.18xlarge | c5.24xlarge | c5a.large | c5a.xlarge | c5a.2xlarge | c5a.4xlarge | c5a.8xlarge | c5a.12xlarge | c5a.16xlarge | c5a.24xlarge | r3.large | r3.xlarge | r3.2xlarge | r3.4xlarge | r3.8xlarge | r4.large | r4.xlarge | r4.2xlarge | r4.4xlarge | r4.8xlarge | r4.16xlarge | r5.large | r5.xlarge | r5.2xlarge | r5.4xlarge | r5.8xlarge | r5.12xlarge | r5.16xlarge | r5.24xlarge | r5a.large | r5a.xlarge | r5a.2xlarge | r5a.4xlarge | r5a.8xlarge | r5a.12xlarge | r5a.16xlarge | r5a.24xlarge | m3.medium | m3.large | m3.xlarge | m3.2xlarge | m4.large | m4.xlarge | m4.2xlarge | m4.4xlarge | m4.10xlarge | m5.large | m5.xlarge | m5.2xlarge | m5.4xlarge | m5.8xlarge | m5.12xlarge | m5.16xlarge | m5.24xlarge | m5a.large | m5a.xlarge | m5a.2xlarge | m5a.4xlarge | m5a.8xlarge | m5a.12xlarge | m5a.16xlarge | m5a.24xlarge | c5d.large | c5d.xlarge | c5d.2xlarge | c5d.4xlarge | c5d.9xlarge | c5d.12xlarge

| c5d.18xlarge | c5d.24xlarge | c6a.large | c6a.xlarge | c6a.2xlarge | c6a.4xlarge | c6a.8xlarge | c6a.12xlarge | c6a.16xlarge | c6a.24xlarge | c6i.large | c6i.xlarge | c6i.2xlarge | c6i.4xlarge | c6i.8xlarge | c6i.12xlarge | c6i.16xlarge | c6i.24xlarge | r5d.large | r5d.xlarge | r5d.2xlarge | r5d.4xlarge | r5d.8xlarge | r5d.12xlarge | r5d.16xlarge | r5d.24xlarge | m6g.medium | m6g.large | m6g.xlarge | m6g.2xlarge | m6g.4xlarge | m6g.8xlarge | m6g.12xlarge | m6g.16xlarge | c6g.medium | c6g.large | c6g.xlarge | c6g.2xlarge | c6g.4xlarge | c6g.8xlarge | c6g.12xlarge | c6g.16xlarge | r6g.medium | r6g.large | r6g.xlarge | r6g.2xlarge | r6g.4xlarge | r6g.8xlarge | r6g.12xlarge | r6g.16xlarge | c6gn.medium | c6gn.large | c6gn.xlarge | c6gn.2xlarge | c6gn.4xlarge | c6gn.8xlarge | c6gn.12xlarge | c6gn.16xlarge | c7g.medium | c7g.large | c7g.xlarge | c7g.2xlarge | c7g.4xlarge | c7g.8xlarge | c7g.12xlarge | c7g.16xlarge | r7g.medium | r7g.large | r7g.xlarge | r7g.2xlarge | r7g.4xlarge | r7g.8xlarge | r7g.12xlarge | r7g.16xlarge | m7g.medium | m7g.large | m7g.xlarge | m7g.2xlarge | m7g.4xlarge | m7g.8xlarge | m7g.12xlarge | m7g.16xlarge | g5g.xlarge | g5g.2xlarge | g5g.4xlarge | g5g.8xlarge | g5g.16xlarge | r6i.large | r6i.xlarge | r6i.2xlarge | r6i.4xlarge | r6i.8xlarge | r6i.12xlarge | r6i.16xlarge | c6gd.medium | c6gd.large | c6gd.xlarge | c6gd.2xlarge | c6gd.4xlarge | c6gd.8xlarge | c6gd.12xlarge | c6gd.16xlarge | c6in.large | c6in.xlarge | c6in.2xlarge | c6in.4xlarge | c6in.8xlarge | c6in.12xlarge | c6in.16xlarge | c7a.medium | c7a.large | c7a.xlarge | c7a.2xlarge | c7a.4xlarge | c7a.8xlarge | c7a.12xlarge | c7a.16xlarge | c7gd.medium | c7gd.large | c7gd.xlarge | c7gd.2xlarge | c7gd.4xlarge | c7gd.8xlarge | c7gd.12xlarge | c7gd.16xlarge | c7gn.medium | c7gn.large | c7gn.xlarge | c7gn.2xlarge | c7gn.4xlarge | c7gn.8xlarge | c7gn.12xlarge | c7gn.16xlarge | c7i.large | c7i.xlarge | c7i.2xlarge | c7i.4xlarge | c7i.8xlarge | c7i.12xlarge | c7i.16xlarge | m6a.large | m6a.xlarge | m6a.2xlarge | m6a.4xlarge | m6a.8xlarge | m6a.12xlarge | m6a.16xlarge | m6gd.medium | m6gd.large | m6gd.xlarge | m6gd.2xlarge | m6gd.4xlarge | m6gd.8xlarge | m6gd.12xlarge | m6gd.16xlarge | m6i.large | m6i.xlarge | m6i.2xlarge | m6i.4xlarge | m6i.8xlarge | m6i.12xlarge | m6i.16xlarge | m7a.medium | m7a.large | m7a.xlarge | m7a.2xlarge | m7a.4xlarge | m7a.8xlarge | m7a.12xlarge | m7a.16xlarge | m7gd.medium | m7gd.large | m7gd.xlarge | m7gd.2xlarge | m7gd.4xlarge | m7gd.8xlarge | m7gd.12xlarge |

m7gd.16xlarge | m7i.large | m7i.xlarge | m7i.2xlarge | m7i.4xlarge |
m7i.8xlarge | m7i.12xlarge | m7i.16xlarge | r6gd.medium | r6gd.large |
r6gd.xlarge | r6gd.2xlarge | r6gd.4xlarge | r6gd.8xlarge | r6gd.12xlarge
| r6gd.16xlarge | r7a.medium | r7a.large | r7a.xlarge | r7a.2xlarge |
r7a.4xlarge | r7a.8xlarge | r7a.12xlarge | r7a.16xlarge | r7gd.medium |
r7gd.large | r7gd.xlarge | r7gd.2xlarge | r7gd.4xlarge | r7gd.8xlarge |
r7gd.12xlarge | r7gd.16xlarge | r7i.large | r7i.xlarge | r7i.2xlarge |
r7i.4xlarge | r7i.8xlarge | r7i.12xlarge | r7i.16xlarge | r7i.24xlarge |
r7i.48xlarge | c5ad.large | c5ad.xlarge | c5ad.2xlarge | c5ad.4xlarge |
c5ad.8xlarge | c5ad.12xlarge | c5ad.16xlarge | c5ad.24xlarge | c5n.large
| c5n.xlarge | c5n.2xlarge | c5n.4xlarge | c5n.9xlarge | c5n.18xlarge |
r5ad.large | r5ad.xlarge | r5ad.2xlarge | r5ad.4xlarge | r5ad.8xlarge |
r5ad.12xlarge | r5ad.16xlarge | r5ad.24xlarge | c6id.large | c6id.xlarge
| c6id.2xlarge | c6id.4xlarge | c6id.8xlarge | c6id.12xlarge |
c6id.16xlarge | c6id.24xlarge | c6id.32xlarge | c8g.medium | c8g.large
| c8g.xlarge | c8g.2xlarge | c8g.4xlarge | c8g.8xlarge | c8g.12xlarge
| c8g.16xlarge | c8g.24xlarge | c8g.48xlarge | m5ad.large | m5ad.xlarge
| m5ad.2xlarge | m5ad.4xlarge | m5ad.8xlarge | m5ad.12xlarge |
m5ad.16xlarge | m5ad.24xlarge | m5d.large | m5d.xlarge | m5d.2xlarge |
m5d.4xlarge | m5d.8xlarge | m5d.12xlarge | m5d.16xlarge | m5d.24xlarge
| m5dn.large | m5dn.xlarge | m5dn.2xlarge | m5dn.4xlarge | m5dn.8xlarge
| m5dn.12xlarge | m5dn.16xlarge | m5dn.24xlarge | m5n.large |
m5n.xlarge | m5n.2xlarge | m5n.4xlarge | m5n.8xlarge | m5n.12xlarge |
m5n.16xlarge | m5n.24xlarge | m6id.large | m6id.xlarge | m6id.2xlarge
| m6id.4xlarge | m6id.8xlarge | m6id.12xlarge | m6id.16xlarge
| m6id.24xlarge | m6id.32xlarge | m6idn.large | m6idn.xlarge |
m6idn.2xlarge | m6idn.4xlarge | m6idn.8xlarge | m6idn.12xlarge |
m6idn.16xlarge | m6idn.24xlarge | m6idn.32xlarge | m6in.large |
m6in.xlarge | m6in.2xlarge | m6in.4xlarge | m6in.8xlarge | m6in.12xlarge
| m6in.16xlarge | m6in.24xlarge | m6in.32xlarge | m8g.medium | m8g.large
| m8g.xlarge | m8g.2xlarge | m8g.4xlarge | m8g.8xlarge | m8g.12xlarge
| m8g.16xlarge | m8g.24xlarge | m8g.48xlarge | r5dn.large | r5dn.xlarge
| r5dn.2xlarge | r5dn.4xlarge | r5dn.8xlarge | r5dn.12xlarge |
r5dn.16xlarge | r5dn.24xlarge | r5n.large | r5n.xlarge | r5n.2xlarge |
r5n.4xlarge | r5n.8xlarge | r5n.12xlarge | r5n.16xlarge | r5n.24xlarge
| r6a.large | r6a.xlarge | r6a.2xlarge | r6a.4xlarge | r6a.8xlarge |

r6a.12xlarge | r6a.16xlarge | r6a.24xlarge | r6a.32xlarge | r6a.48xlarge
| r6id.large | r6id.xlarge | r6id.2xlarge | r6id.4xlarge | r6id.8xlarge
| r6id.12xlarge | r6id.16xlarge | r6id.24xlarge | r6id.32xlarge
| r6idn.large | r6idn.xlarge | r6idn.2xlarge | r6idn.4xlarge |
r6idn.8xlarge | r6idn.12xlarge | r6idn.16xlarge | r6idn.24xlarge |
r6idn.32xlarge | r6in.large | r6in.xlarge | r6in.2xlarge | r6in.4xlarge
| r6in.8xlarge | r6in.12xlarge | r6in.16xlarge | r6in.24xlarge |
r6in.32xlarge | r8g.medium | r8g.large | r8g.xlarge | r8g.2xlarge |
r8g.4xlarge | r8g.8xlarge | r8g.12xlarge | r8g.16xlarge | r8g.24xlarge |
r8g.48xlarge | m4.16xlarge | c6a.32xlarge | c6a.48xlarge | c6i.32xlarge
| r6i.24xlarge | r6i.32xlarge | c6in.24xlarge | c6in.32xlarge |
c7a.24xlarge | c7a.32xlarge | c7a.48xlarge | c7i.24xlarge | c7i.48xlarge
| m6a.24xlarge | m6a.32xlarge | m6a.48xlarge | m6i.24xlarge |
m6i.32xlarge | m7a.24xlarge | m7a.32xlarge | m7a.48xlarge | m7i.24xlarge
| m7i.48xlarge | r7a.24xlarge | r7a.32xlarge | r7a.48xlarge

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](#)

ConnectionPortRange

The set of port numbers to open on each instance in a container fleet. Connection ports are used by inbound traffic to connect with processes that are running in containers on the fleet.

The port range must not overlap with the Amazon GameLift Servers reserved port range 4092-4191. This range is reserved for internal Amazon GameLift Servers services.

Contents

Note

In the following list, the required parameters are described first.

FromPort

Starting value for the port range.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 60000.

Required: Yes

ToPort

Ending value for the port. Port numbers are end-inclusive. This value must be equal to or greater than `FromPort`.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 60000.

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](#)

ContainerAttribute

A unique identifier for a container in a container fleet compute.

Returned by: [DescribeCompute](#)

Contents

Note

In the following list, the required parameters are described first.

ContainerName

The identifier for a container that's running in a compute.

Type: String

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

Pattern: `^[a-zA-Z0-9\ -]+$`

Required: No

ContainerRuntimeId

The runtime ID for the container that's running in a compute. This value is unique within the compute.

Type: String

Length Constraints: Minimum length 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](#)

ContainerDependency

A container's dependency on another container in the same container group. The dependency impacts how the dependent container is able to start or shut down based the status of the other container.

For example, *ContainerA* is configured with the following dependency: a START dependency on *ContainerB*. This means that *ContainerA* can't start until *ContainerB* has started. It also means that *ContainerA* must shut down before *ContainerB*.

Part of: [GameServerContainerDefinition](#), [GameServerContainerDefinitionInput](#), [SupportContainerDefinition](#), [SupportContainerDefinitionInput](#)

Contents

Note

In the following list, the required parameters are described first.

Condition

The condition that the dependency container must reach before the dependent container can start. Valid conditions include:

- **START** - The dependency container must have started.
- **COMPLETE** - The dependency container has run to completion (exits). Use this condition with nonessential containers, such as those that run a script and then exit. The dependency container can't be an essential container.
- **SUCCESS** - The dependency container has run to completion and exited with a zero status. The dependency container can't be an essential container.
- **HEALTHY** - The dependency container has passed its Docker health check. Use this condition with dependency containers that have health checks configured. This condition is confirmed at container group startup only.

Type: String

Valid Values: START | COMPLETE | SUCCESS | HEALTHY

Required: Yes

ContainerName

A descriptive label for the container definition that this container depends on.

Type: String

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

Pattern: `^[a-zA-Z0-9\ -]+$`

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](#)

ContainerEnvironment

An environment variable to set inside a container, in the form of a key-value pair.

Part of: [GameServerContainerDefinition](#), [GameServerContainerDefinitionInput](#), [SupportContainerDefinition](#), [SupportContainerDefinitionInput](#)

Contents

Note

In the following list, the required parameters are described first.

Name

The environment variable name.

Type: String

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

Required: Yes

Value

The environment variable value.

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](#)

ContainerFleet

Describes an Amazon GameLift Servers managed container fleet.

Contents

Note

In the following list, the required parameters are described first.

BillingType

Indicates whether the fleet uses On-Demand or Spot instances for this fleet. Learn more about when to use [On-Demand versus Spot Instances](#). You can't update this fleet property.

By default, this property is set to ON_DEMAND.

Type: String

Valid Values: ON_DEMAND | SPOT

Required: No

CreationTime

A time stamp indicating when this data object was created. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

DeploymentDetails

Information about the most recent deployment for the container fleet.

Type: [DeploymentDetails](#) object

Required: No

Description

A meaningful description of the container fleet.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

FleetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`. In a GameLift fleet ARN, the resource ID matches the `FleetId` value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\]\+$`

Required: No

FleetId

A unique identifier for the container fleet to retrieve.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\]\+`

Required: No

FleetRoleArn

The unique identifier for an AWS Identity and Access Management (IAM) role with permissions to run your containers on resources that are managed by Amazon GameLift Servers. See [Set up an IAM service role](#). This fleet property can't be changed.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^arn:.*:role\[\\w+=,.\@-\]\+`

Required: No

GameServerContainerGroupDefinitionArn

The Amazon Resource Name ([ARN](#)) that is assigned to the fleet's game server container group. The ARN value also identifies the specific container group definition version in use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:containergroupdefinition\[a-zA-Z0-9\-\]+(:[0-9]+)?$`

Required: No

GameServerContainerGroupDefinitionName

The name of the fleet's game server container group definition, which describes how to deploy containers with your game server build and support software onto each fleet instance.

Type: String

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

Pattern: `^[a-zA-Z0-9\-\]+$`

Required: No

GameServerContainerGroupsPerInstance

The number of times to replicate the game server container group on each fleet instance.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 5000.

Required: No

GameSessionCreationLimitPolicy

A policy that limits the number of game sessions that each individual player can create on instances in this fleet. The limit applies for a specified span of time.

Type: [GameSessionCreationLimitPolicy](#) object

Required: No

InstanceConnectionPortRange

The set of port numbers to open on each instance in a container fleet. Connection ports are used by inbound traffic to connect with processes that are running in containers on the fleet.

The port range must not overlap with the Amazon GameLift Servers reserved port range 4092-4191. This range is reserved for internal Amazon GameLift Servers services.

Type: [ConnectionPortRange](#) object

Required: No

InstanceInboundPermissions

The IP address ranges and port settings that allow inbound traffic to access game server processes and other processes on this fleet.

Type: Array of [IpPermission](#) objects

Array Members: Maximum number of 50 items.

Required: No

InstanceType

The Amazon EC2 instance type to use for all instances in the fleet. Instance type determines the computing resources and processing power that's available to host your game servers. This includes including CPU, memory, storage, and networking capacity. You can't update this fleet property.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

LocationAttributes

Information about the container fleet's remote locations where fleet instances are deployed.

Type: Array of [ContainerFleetLocationAttributes](#) objects

Required: No

LogConfiguration

The method that is used to collect container logs for the fleet. Amazon GameLift Servers saves all standard output for each container in logs, including game session logs.

- CLOUDWATCH -- Send logs to an Amazon CloudWatch log group that you define. Each container emits a log stream, which is organized in the log group.
- S3 -- Store logs in an Amazon S3 bucket that you define.
- NONE -- Don't collect container logs.

Type: [LogConfiguration](#) object

Required: No

MaximumGameServerContainerGroupsPerInstance

The calculated maximum number of game server container group that can be deployed on each fleet instance. The calculation depends on the resource needs of the container group and the CPU and memory resources of the fleet's instance type.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 5000.

Required: No

MetricGroups

The name of an AWS CloudWatch metric group to add this fleet to. Metric groups aggregate metrics for multiple fleets.

Type: Array of strings

Array Members: Maximum number of 1 item.

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

Required: No

NewGameSessionProtectionPolicy

Determines whether Amazon GameLift Servers can shut down game sessions on the fleet that are actively running and hosting players. Amazon GameLift Servers might prompt an instance

shutdown when scaling down fleet capacity or when retiring unhealthy instances. You can also set game session protection for individual game sessions using [UpdateGameSession](#).

- **NoProtection** -- Game sessions can be shut down during active gameplay.
- **FullProtection** -- Game sessions in ACTIVE status can't be shut down.

Type: String

Valid Values: NoProtection | FullProtection

Required: No

PerInstanceContainerGroupDefinitionArn

The Amazon Resource Name ([ARN](#)) that is assigned to the fleet's per-instance container group. The ARN value also identifies the specific container group definition version in use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:containergroupdefinition\[a-zA-Z0-9\]+(:[0-9]+)?$`

Required: No

PerInstanceContainerGroupDefinitionName

The name of the fleet's per-instance container group definition.

Type: String

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

Pattern: `^[a-zA-Z0-9\]+`

Required: No

PlayerGatewayMode

Indicates whether player gateway is enabled for this container fleet. Player gateway provides benefits such as DDoS protection with negligible impact to latency.

If **ENABLED** or **REQUIRED**, game clients can use player gateway to connect with the game server. If **DISABLED**, game clients cannot use player gateway. Instead, they have to directly connect to the game server.

Type: String

Valid Values: DISABLED | ENABLED | REQUIRED

Required: No

Status

The current status of the container fleet.

- PENDING -- A new container fleet has been requested.
- CREATING -- A new container fleet resource is being created.
- CREATED -- A new container fleet resource has been created. No fleet instances have been deployed.
- ACTIVATING -- New container fleet instances are being deployed.
- ACTIVE -- The container fleet has been deployed and is ready to host game sessions.
- UPDATING -- Updates to the container fleet is being updated. A deployment is in progress.
- EXPIRED -- The container fleet has been expired. The fleet is scaled down to zero instances and cannot host new game sessions.

Type: String

Valid Values: PENDING | CREATING | CREATED | ACTIVATING | ACTIVE | UPDATING
| DELETING

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](#)

ContainerFleetLocationAttributes

Details about a location in a multi-location container fleet.

Contents

Note

In the following list, the required parameters are described first.

Location

A location identifier.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-]+`

Required: No

PlayerGatewayStatus

The current status of player gateway in this location for this container fleet. Note, even if a container fleet has `PlayerGatewayMode` configured as `ENABLED`, player gateway might not be available in a specific location. For more information about locations where player gateway is supported, see [Amazon GameLift Servers service locations](#).

Possible values include:

- `ENABLED` -- Player gateway is available for this container fleet location.
- `DISABLED` -- Player gateway is not available for this container fleet location.

Type: String

Valid Values: `DISABLED` | `ENABLED`

Required: No

Status

The status of fleet activity in the location.

- **PENDING** -- A new container fleet has been requested.
- **CREATING** -- A new container fleet resource is being created.
- **CREATED** -- A new container fleet resource has been created. No fleet instances have been deployed.
- **ACTIVATING** -- New container fleet instances are being deployed.
- **ACTIVE** -- The container fleet has been deployed and is ready to host game sessions.
- **UPDATING** -- Updates to the container fleet is being updated. A deployment is in progress.
- **EXPIRED** -- The container fleet has been expired. The fleet is scaled down to zero instances and cannot host new game sessions.

Type: String

Valid Values: **PENDING** | **CREATING** | **CREATED** | **ACTIVATING** | **ACTIVE** | **UPDATING** | **DELETING**

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](#)

ContainerGroupDefinition

The properties that describe a container group resource. You can update all properties of a container group definition properties. Updates to a container group definition are saved as new versions.

Used with: [CreateContainerGroupDefinition](#)

Returned by: [DescribeContainerGroupDefinition](#), [ListContainerGroupDefinitions](#), [UpdateContainerGroupDefinition](#)

Contents

Note

In the following list, the required parameters are described first.

Name

A descriptive identifier for the container group definition. The name value is unique in an AWS Region.

Type: String

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

Pattern: `^[a-zA-Z0-9\-\_]+$`

Required: Yes

ContainerGroupDefinitionArn

The Amazon Resource Name ([ARN](#)) that is assigned to an Amazon GameLift Servers ContainerGroupDefinition resource. It uniquely identifies the resource across all AWS Regions. Format is `arn:aws:gamelift:[region]::containergroupdefinition/[container group definition name]:[version]`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:containergroupdefinition\[a-zA-Z0-9\-\]+(:[0-9]+)?$`

Required: No

ContainerGroupType

The type of container group. Container group type determines how Amazon GameLift Servers deploys the container group on each fleet instance.

Type: String

Valid Values: `GAME_SERVER` | `PER_INSTANCE`

Required: No

CreationTime

A time stamp indicating when this data object was created. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

GameServerContainerDefinition

The definition for the game server container in this group. This property is used only when the container group type is `GAME_SERVER`. This container definition specifies a container image with the game server build.

Type: [GameServerContainerDefinition](#) object

Required: No

OperatingSystem

The platform that all containers in the container group definition run on.

Note

Amazon Linux 2 (AL2) will reach end of support on 6/30/2026. See more details in the [Amazon Linux 2 FAQs](#). For game servers that are hosted on AL2 and use server SDK version 4.x for Amazon GameLift Servers, first update the game server build to server SDK 5.x, and then deploy to AL2023 instances. See [Migrate to server SDK version 5](#).

Type: String

Valid Values: AMAZON_LINUX_2023

Required: No

Status

Current status of the container group definition resource. Values include:

- **COPYING** -- Amazon GameLift Servers is in the process of making copies of all container images that are defined in the group. While in this state, the resource can't be used to create a container fleet.
- **READY** -- Amazon GameLift Servers has copied the registry images for all containers that are defined in the group. You can use a container group definition in this status to create a container fleet.
- **FAILED** -- Amazon GameLift Servers failed to create a valid container group definition resource. For more details on the cause of the failure, see `StatusReason`. A container group definition resource in failed status will be deleted within a few minutes.

Type: String

Valid Values: READY | COPYING | FAILED

Required: No

StatusReason

Additional information about a container group definition that's in **FAILED** status. Possible reasons include:

- An internal issue prevented Amazon GameLift Servers from creating the container group definition resource. Delete the failed resource and call [CreateContainerGroupDefinition](#) again.
- An access-denied message means that you don't have permissions to access the container image on ECR. See [IAM permission examples](#) for help setting up required IAM permissions for Amazon GameLift Servers.
- The `ImageUri` value for at least one of the containers in the container group definition was invalid or not found in the current AWS account.
- At least one of the container images referenced in the container group definition exceeds the allowed size. For size limits, see [Amazon GameLift Servers endpoints and quotas](#).

- At least one of the container images referenced in the container group definition uses a different operating system than the one defined for the container group.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

SupportContainerDefinitions

The set of definitions for support containers in this group. A container group definition might have zero support container definitions. Support container can be used in any type of container group.

Type: Array of [SupportContainerDefinition](#) objects

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

Required: No

TotalMemoryLimitMebibytes

The amount of memory (in MiB) on a fleet instance to allocate for the container group. All containers in the group share these resources.

You can set a limit for each container definition in the group. If individual containers have limits, this total value must be greater than any individual container's memory limit.

Type: Integer

Valid Range: Minimum value of 4. Maximum value of 1024000.

Required: No

TotalVcpuLimit

The amount of vCPU units on a fleet instance to allocate for the container group (1 vCPU is equal to 1024 CPU units). All containers in the group share these resources. You can set a limit for each container definition in the group. If individual containers have limits, this total value must be equal to or greater than the sum of the limits for each container in the group.

Type: Double

Valid Range: Minimum value of 0.125. Maximum value of 10.

Required: No

VersionDescription

An optional description that was provided for a container group definition update. Each version can have a unique description.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

VersionNumber

Indicates the version of a particular container group definition. This number is incremented automatically when you update a container group definition. You can view, update, or delete individual versions or the entire container group definition.

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](#)

ContainerGroupPortMapping

Describes the port mappings for a single container in a container group. Each mapping shows how a container port maps to a connection port on the fleet instance.

Returned by: [DescribeContainerGroupPortMappings](#)

Contents

Note

In the following list, the required parameters are described first.

ContainerName

The name of the container, as defined in the container group definition.

Type: String

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

Pattern: `^[a-zA-Z0-9\-\_]+$`

Required: No

ContainerPortMappings

A list of ContainerPortMapping objects that describe the port mappings for this container.

Type: Array of [ContainerPortMapping](#) objects

Required: No

ContainerRuntimeId

The runtime ID for the container that's running in a compute. This value is unique within the compute.

Type: String

Length Constraints: Minimum length 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](#)

ContainerHealthCheck

Instructions on when and how to check the health of a support container in a container fleet. These properties override any Docker health checks that are set in the container image. For more information on container health checks, see [HealthCheck command](#) in the *Amazon Elastic Container Service API*. Game server containers don't have a health check parameter; Amazon GameLift Servers automatically handles health checks for these containers.

The following example instructs the container to initiate a health check command every 60 seconds and wait 10 seconds for it to succeed. If it fails, retry the command 3 times before flagging the container as unhealthy. It also tells the container to wait 100 seconds after launch before counting failed health checks.

```
{"Command": [ "CMD-SHELL", "ps cax | grep \"processmanager\" || exit 1" ],  
"Interval": 60, "Timeout": 10, "Retries": 3, "StartPeriod": 100 }
```

Part of: [SupportContainerDefinition](#), [SupportContainerDefinitionInput](#)

Contents

Note

In the following list, the required parameters are described first.

Command

A string array that specifies the command that the container runs to determine if it's healthy.

Type: Array of strings

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

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

Required: Yes

Interval

The time period (in seconds) between each health check.

Type: Integer

Valid Range: Minimum value of 60. Maximum value of 300.

Required: No

Retries

The number of times to retry a failed health check before flagging the container unhealthy. The first run of the command does not count as a retry.

Type: Integer

Valid Range: Minimum value of 5. Maximum value of 10.

Required: No

StartPeriod

The optional grace period (in seconds) to give a container time to bootstrap before the first failed health check counts toward the number of retries.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 300.

Required: No

Timeout

The time period (in seconds) to wait for a health check to succeed before counting a failed health check.

Type: Integer

Valid Range: Minimum value of 30. Maximum value of 60.

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](#)

ContainerIdentifier

A unique identifier for a container in a compute on a managed container fleet instance. This information makes it possible to remotely connect to a specific container on a fleet instance.

Related to: [ContainerAttribute](#)

Use with: [GetComputeAccess](#)

Contents

Note

In the following list, the required parameters are described first.

ContainerName

The identifier for a container that's running in a compute.

Type: String

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

Pattern: `^[a-zA-Z0-9\-\_]+$`

Required: No

ContainerRuntimeId

The runtime ID for the container that's running in a compute. This value is unique within the compute. It is returned as a `ContainerAttribute` value in a `Compute` object.

Type: String

Length Constraints: Minimum length 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](#)

ContainerMountPoint

A mount point that binds a container to a file or directory on the host system.

Part of: [GameServerContainerDefinition](#), https://docs.aws.amazon.com/gamelift/latest/apireference/API_GameServerContainerDefinitionInput.html, [SupportContainerDefinition](#), https://docs.aws.amazon.com/gamelift/latest/apireference/API_SupportContainerDefinitionInput.html

Contents

Note

In the following list, the required parameters are described first.

InstancePath

The path to the source file or directory.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `^\\[\\s\\S]*$`

Required: Yes

AccessLevel

The type of access for the container.

Type: String

Valid Values: `READ_ONLY` | `READ_AND_WRITE`

Required: No

ContainerPath

The mount path on the container. If this property isn't set, the instance path is used.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `^(\\+[^\\]+\\* )+$`

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](#)

ContainerPortConfiguration

A set of port ranges that can be opened on the container. A process that's running in the container can bind to a port number, making it accessible to inbound traffic when it's mapped to a container fleet's connection port.

Each container port range specifies a network protocol. When the configuration supports more than one protocol, we recommend that you use a different range for each protocol. If your ranges have overlapping port numbers, Amazon GameLift Servers maps a duplicated container port number to different connection ports. For example, if you include 1935 in port ranges for both TCP and UDP, it might result in the following mappings:

- container port 1935 (tcp) => connection port 2001
- container port 1935 (udp) => connection port 2002

Part of: [GameServerContainerDefinition](#), [GameServerContainerDefinitionInput](#), [SupportContainerDefinition](#), [SupportContainerDefinitionInput](#)

Contents

Note

In the following list, the required parameters are described first.

ContainerPortRanges

A set of one or more container port number ranges. The ranges can't overlap if the ranges' network protocols are the same. Overlapping ranges with different protocols is allowed but not recommended.

Type: Array of [ContainerPortRange](#) objects

Array Members: Minimum number of 1 item. Maximum number of 100 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](#)

ContainerPortMapping

Describes a mapping between a container port and a connection port on a fleet instance. You define container ports in a container group definition. Amazon GameLift Servers assigns connection ports when it deploys the container group to an instance.

Part of: [ContainerGroupPortMapping](#)

Contents

Note

In the following list, the required parameters are described first.

ConnectionPort

The port number on the fleet instance that maps to the container port. Connection ports are assigned by Amazon GameLift Servers when the container group is deployed to an instance.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 60000.

Required: No

ContainerPort

The port number on the container. This port is defined in the container group definition. Container port numbers must be unique within a container group definition.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 60000.

Required: No

Protocol

The network protocol for the port mapping. Valid values are TCP or UDP.

Type: String

Valid Values: TCP | UDP

Required: No

See Also

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

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

ContainerPortRange

A set of one or more port numbers that can be opened on the container, and the supported network protocol.

Part of: [ContainerPortConfiguration](#)

Contents

Note

In the following list, the required parameters are described first.

FromPort

A starting value for the range of allowed port numbers.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 60000.

Required: Yes

Protocol

The network protocol that these ports support.

Type: String

Valid Values: TCP | UDP

Required: Yes

ToPort

An ending value for the range of allowed port numbers. Port numbers are end-inclusive. This value must be equal to or greater than FromPort.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 60000.

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](#)

DeploymentConfiguration

Set of rules for processing a deployment for a container fleet update.

Contents

Note

In the following list, the required parameters are described first.

ImpairmentStrategy

Determines what actions to take if a deployment fails. If the fleet is multi-location, this strategy applies across all fleet locations. With a rollback strategy, updated fleet instances are rolled back to the last successful deployment. Alternatively, you can maintain a few impaired containers for the purpose of debugging, while all other tasks return to the last successful deployment.

Type: String

Valid Values: MAINTAIN | ROLLBACK

Required: No

MinimumHealthyPercentage

Sets a minimum level of healthy tasks to maintain during deployment activity.

Type: Integer

Valid Range: Minimum value of 30. Maximum value of 75.

Required: No

ProtectionStrategy

Determines how fleet deployment activity affects active game sessions on the fleet. With protection, a deployment honors game session protection, and delays actions that would interrupt a protected active game session until the game session ends. Without protection, deployment activity can shut down all running tasks, including active game sessions, regardless of game session protection.

Type: String

Valid Values: WITH_PROTECTION | IGNORE_PROTECTION

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](#)

DeploymentDetails

Information about the most recent deployment for the container fleet.

Contents

Note

In the following list, the required parameters are described first.

LatestDeploymentId

A unique identifier for a fleet deployment.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `^[a-zA-Z0-9\-\_]+$`

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](#)

DesiredPlayerSession

Player information for use when creating player sessions using a game session placement request.

Contents

Note

In the following list, the required parameters are described first.

PlayerData

Developer-defined information related to a player. Amazon GameLift Servers does not use this data, so it can be formatted as needed for use in the game.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

PlayerId

A unique identifier for a player to associate with the player session.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

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](#)

EC2InstanceCounts

Resource capacity settings. Fleet capacity is measured in Amazon EC2 instances. Pending and terminating counts are non-zero when the fleet capacity is adjusting to a scaling event or if access to resources is temporarily affected.

Contents

Note

In the following list, the required parameters are described first.

ACTIVE

Actual number of instances that are ready to host game sessions.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

DESIRED

Requested number of active instances. Amazon GameLift Servers takes action as needed to maintain the desired number of instances. Capacity is scaled up or down by changing the desired instances. A change in the desired instances value can take up to 1 minute to be reflected when viewing a fleet's capacity settings.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

IDLE

Number of active instances that are not currently hosting a game session.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

MAXIMUM

The maximum instance count value allowed.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

MINIMUM

The minimum instance count value allowed.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

PENDING

Number of instances that are starting but not yet active.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

TERMINATING

Number of instances that are no longer active but haven't yet been terminated.

Type: Integer

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](#)

EC2InstanceLimit

The Amazon GameLift Servers service limits for an Amazon EC2 instance type and current utilization. Amazon GameLift Servers allows AWS accounts a maximum number of instances, per instance type, per AWS Region or location, for use with Amazon GameLift Servers. You can request an limit increase for your account by using the **Service limits** page in the Amazon GameLift Servers console.

Contents

Note

In the following list, the required parameters are described first.

CurrentInstances

The number of instances for the specified type and location that are currently being used by the AWS account.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

EC2InstanceType

The name of an Amazon EC2 instance type. See [Amazon Elastic Compute Cloud Instance Types](#) for detailed descriptions.

Type: String

Valid Values: t2.micro | t2.small | t2.medium | t2.large | c3.large | c3.xlarge | c3.2xlarge | c3.4xlarge | c3.8xlarge | c4.large | c4.xlarge | c4.2xlarge | c4.4xlarge | c4.8xlarge | c5.large | c5.xlarge | c5.2xlarge | c5.4xlarge | c5.9xlarge | c5.12xlarge | c5.18xlarge | c5.24xlarge | c5a.large | c5a.xlarge | c5a.2xlarge | c5a.4xlarge | c5a.8xlarge | c5a.12xlarge | c5a.16xlarge | c5a.24xlarge | r3.large | r3.xlarge | r3.2xlarge | r3.4xlarge | r3.8xlarge | r4.large | r4.xlarge | r4.2xlarge | r4.4xlarge | r4.8xlarge | r4.16xlarge | r5.large |

r5.xlarge | r5.2xlarge | r5.4xlarge | r5.8xlarge | r5.12xlarge |
r5.16xlarge | r5.24xlarge | r5a.large | r5a.xlarge | r5a.2xlarge |
r5a.4xlarge | r5a.8xlarge | r5a.12xlarge | r5a.16xlarge | r5a.24xlarge
| m3.medium | m3.large | m3.xlarge | m3.2xlarge | m4.large | m4.xlarge
| m4.2xlarge | m4.4xlarge | m4.10xlarge | m5.large | m5.xlarge |
m5.2xlarge | m5.4xlarge | m5.8xlarge | m5.12xlarge | m5.16xlarge |
m5.24xlarge | m5a.large | m5a.xlarge | m5a.2xlarge | m5a.4xlarge |
m5a.8xlarge | m5a.12xlarge | m5a.16xlarge | m5a.24xlarge | c5d.large
| c5d.xlarge | c5d.2xlarge | c5d.4xlarge | c5d.9xlarge | c5d.12xlarge
| c5d.18xlarge | c5d.24xlarge | c6a.large | c6a.xlarge | c6a.2xlarge |
c6a.4xlarge | c6a.8xlarge | c6a.12xlarge | c6a.16xlarge | c6a.24xlarge
| c6i.large | c6i.xlarge | c6i.2xlarge | c6i.4xlarge | c6i.8xlarge |
c6i.12xlarge | c6i.16xlarge | c6i.24xlarge | r5d.large | r5d.xlarge |
r5d.2xlarge | r5d.4xlarge | r5d.8xlarge | r5d.12xlarge | r5d.16xlarge
| r5d.24xlarge | m6g.medium | m6g.large | m6g.xlarge | m6g.2xlarge |
m6g.4xlarge | m6g.8xlarge | m6g.12xlarge | m6g.16xlarge | c6g.medium
| c6g.large | c6g.xlarge | c6g.2xlarge | c6g.4xlarge | c6g.8xlarge |
c6g.12xlarge | c6g.16xlarge | r6g.medium | r6g.large | r6g.xlarge |
r6g.2xlarge | r6g.4xlarge | r6g.8xlarge | r6g.12xlarge | r6g.16xlarge
| c6gn.medium | c6gn.large | c6gn.xlarge | c6gn.2xlarge | c6gn.4xlarge
| c6gn.8xlarge | c6gn.12xlarge | c6gn.16xlarge | c7g.medium | c7g.large
| c7g.xlarge | c7g.2xlarge | c7g.4xlarge | c7g.8xlarge | c7g.12xlarge
| c7g.16xlarge | r7g.medium | r7g.large | r7g.xlarge | r7g.2xlarge |
r7g.4xlarge | r7g.8xlarge | r7g.12xlarge | r7g.16xlarge | m7g.medium
| m7g.large | m7g.xlarge | m7g.2xlarge | m7g.4xlarge | m7g.8xlarge |
m7g.12xlarge | m7g.16xlarge | g5g.xlarge | g5g.2xlarge | g5g.4xlarge
| g5g.8xlarge | g5g.16xlarge | r6i.large | r6i.xlarge | r6i.2xlarge |
r6i.4xlarge | r6i.8xlarge | r6i.12xlarge | r6i.16xlarge | c6gd.medium |
c6gd.large | c6gd.xlarge | c6gd.2xlarge | c6gd.4xlarge | c6gd.8xlarge |
c6gd.12xlarge | c6gd.16xlarge | c6in.large | c6in.xlarge | c6in.2xlarge
| c6in.4xlarge | c6in.8xlarge | c6in.12xlarge | c6in.16xlarge |
c7a.medium | c7a.large | c7a.xlarge | c7a.2xlarge | c7a.4xlarge |
c7a.8xlarge | c7a.12xlarge | c7a.16xlarge | c7gd.medium | c7gd.large |
c7gd.xlarge | c7gd.2xlarge | c7gd.4xlarge | c7gd.8xlarge | c7gd.12xlarge
| c7gd.16xlarge | c7gn.medium | c7gn.large | c7gn.xlarge | c7gn.2xlarge
| c7gn.4xlarge | c7gn.8xlarge | c7gn.12xlarge | c7gn.16xlarge |

c7i.large | c7i.xlarge | c7i.2xlarge | c7i.4xlarge | c7i.8xlarge |
c7i.12xlarge | c7i.16xlarge | m6a.large | m6a.xlarge | m6a.2xlarge |
m6a.4xlarge | m6a.8xlarge | m6a.12xlarge | m6a.16xlarge | m6gd.medium
| m6gd.large | m6gd.xlarge | m6gd.2xlarge | m6gd.4xlarge | m6gd.8xlarge
| m6gd.12xlarge | m6gd.16xlarge | m6i.large | m6i.xlarge | m6i.2xlarge
| m6i.4xlarge | m6i.8xlarge | m6i.12xlarge | m6i.16xlarge | m7a.medium
| m7a.large | m7a.xlarge | m7a.2xlarge | m7a.4xlarge | m7a.8xlarge |
m7a.12xlarge | m7a.16xlarge | m7gd.medium | m7gd.large | m7gd.xlarge
| m7gd.2xlarge | m7gd.4xlarge | m7gd.8xlarge | m7gd.12xlarge |
m7gd.16xlarge | m7i.large | m7i.xlarge | m7i.2xlarge | m7i.4xlarge |
m7i.8xlarge | m7i.12xlarge | m7i.16xlarge | r6gd.medium | r6gd.large |
r6gd.xlarge | r6gd.2xlarge | r6gd.4xlarge | r6gd.8xlarge | r6gd.12xlarge
| r6gd.16xlarge | r7a.medium | r7a.large | r7a.xlarge | r7a.2xlarge |
r7a.4xlarge | r7a.8xlarge | r7a.12xlarge | r7a.16xlarge | r7gd.medium |
r7gd.large | r7gd.xlarge | r7gd.2xlarge | r7gd.4xlarge | r7gd.8xlarge |
r7gd.12xlarge | r7gd.16xlarge | r7i.large | r7i.xlarge | r7i.2xlarge |
r7i.4xlarge | r7i.8xlarge | r7i.12xlarge | r7i.16xlarge | r7i.24xlarge |
r7i.48xlarge | c5ad.large | c5ad.xlarge | c5ad.2xlarge | c5ad.4xlarge |
c5ad.8xlarge | c5ad.12xlarge | c5ad.16xlarge | c5ad.24xlarge | c5n.large
| c5n.xlarge | c5n.2xlarge | c5n.4xlarge | c5n.9xlarge | c5n.18xlarge |
r5ad.large | r5ad.xlarge | r5ad.2xlarge | r5ad.4xlarge | r5ad.8xlarge |
r5ad.12xlarge | r5ad.16xlarge | r5ad.24xlarge | c6id.large | c6id.xlarge
| c6id.2xlarge | c6id.4xlarge | c6id.8xlarge | c6id.12xlarge |
c6id.16xlarge | c6id.24xlarge | c6id.32xlarge | c8g.medium | c8g.large
| c8g.xlarge | c8g.2xlarge | c8g.4xlarge | c8g.8xlarge | c8g.12xlarge
| c8g.16xlarge | c8g.24xlarge | c8g.48xlarge | m5ad.large | m5ad.xlarge
| m5ad.2xlarge | m5ad.4xlarge | m5ad.8xlarge | m5ad.12xlarge |
m5ad.16xlarge | m5ad.24xlarge | m5d.large | m5d.xlarge | m5d.2xlarge |
m5d.4xlarge | m5d.8xlarge | m5d.12xlarge | m5d.16xlarge | m5d.24xlarge
| m5dn.large | m5dn.xlarge | m5dn.2xlarge | m5dn.4xlarge | m5dn.8xlarge
| m5dn.12xlarge | m5dn.16xlarge | m5dn.24xlarge | m5n.large |
m5n.xlarge | m5n.2xlarge | m5n.4xlarge | m5n.8xlarge | m5n.12xlarge |
m5n.16xlarge | m5n.24xlarge | m6id.large | m6id.xlarge | m6id.2xlarge
| m6id.4xlarge | m6id.8xlarge | m6id.12xlarge | m6id.16xlarge
| m6id.24xlarge | m6id.32xlarge | m6idn.large | m6idn.xlarge |
m6idn.2xlarge | m6idn.4xlarge | m6idn.8xlarge | m6idn.12xlarge |

m6idn.16xlarge | m6idn.24xlarge | m6idn.32xlarge | m6in.large | m6in.xlarge | m6in.2xlarge | m6in.4xlarge | m6in.8xlarge | m6in.12xlarge | m6in.16xlarge | m6in.24xlarge | m6in.32xlarge | m8g.medium | m8g.large | m8g.xlarge | m8g.2xlarge | m8g.4xlarge | m8g.8xlarge | m8g.12xlarge | m8g.16xlarge | m8g.24xlarge | m8g.48xlarge | r5dn.large | r5dn.xlarge | r5dn.2xlarge | r5dn.4xlarge | r5dn.8xlarge | r5dn.12xlarge | r5dn.16xlarge | r5dn.24xlarge | r5n.large | r5n.xlarge | r5n.2xlarge | r5n.4xlarge | r5n.8xlarge | r5n.12xlarge | r5n.16xlarge | r5n.24xlarge | r6a.large | r6a.xlarge | r6a.2xlarge | r6a.4xlarge | r6a.8xlarge | r6a.12xlarge | r6a.16xlarge | r6a.24xlarge | r6a.32xlarge | r6a.48xlarge | r6id.large | r6id.xlarge | r6id.2xlarge | r6id.4xlarge | r6id.8xlarge | r6id.12xlarge | r6id.16xlarge | r6id.24xlarge | r6id.32xlarge | r6idn.large | r6idn.xlarge | r6idn.2xlarge | r6idn.4xlarge | r6idn.8xlarge | r6idn.12xlarge | r6idn.16xlarge | r6idn.24xlarge | r6idn.32xlarge | r6in.large | r6in.xlarge | r6in.2xlarge | r6in.4xlarge | r6in.8xlarge | r6in.12xlarge | r6in.16xlarge | r6in.24xlarge | r6in.32xlarge | r8g.medium | r8g.large | r8g.xlarge | r8g.2xlarge | r8g.4xlarge | r8g.8xlarge | r8g.12xlarge | r8g.16xlarge | r8g.24xlarge | r8g.48xlarge | m4.16xlarge | c6a.32xlarge | c6a.48xlarge | c6i.32xlarge | r6i.24xlarge | r6i.32xlarge | c6in.24xlarge | c6in.32xlarge | c7a.24xlarge | c7a.32xlarge | c7a.48xlarge | c7i.24xlarge | c7i.48xlarge | m6a.24xlarge | m6a.32xlarge | m6a.48xlarge | m6i.24xlarge | m6i.32xlarge | m7a.24xlarge | m7a.32xlarge | m7a.48xlarge | m7i.24xlarge | m7i.48xlarge | r7a.24xlarge | r7a.32xlarge | r7a.48xlarge

Required: No

InstanceLimit

The number of instances that is allowed for the specified instance type and location.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

Location

An AWS Region code, such as us-west-2.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\ -]+`

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](#)

Event

Log entry describing an event that involves Amazon GameLift Servers resources (such as a fleet). In addition to tracking activity, event codes and messages can provide additional information for troubleshooting and debugging problems.

Contents

Note

In the following list, the required parameters are described first.

Count

The number of times that this event occurred.

Type: Long

Valid Range: Minimum value of 1. Maximum value of 101.

Required: No

EventCode

The type of event being logged.

Fleet state transition events:

- **FLEET_CREATED** -- A fleet resource was successfully created with a status of NEW. Event messaging includes the fleet ID.
- **FLEET_STATE_DOWNLOADING** -- Fleet status changed from NEW to DOWNLOADING. Amazon GameLift Servers is downloading the compressed build and running install scripts.
- **FLEET_STATE_VALIDATING** -- Fleet status changed from DOWNLOADING to VALIDATING. Amazon GameLift Servers has successfully installed build and is now validating the build files.
- **FLEET_STATE_BUILDING** -- Fleet status changed from VALIDATING to BUILDING. Amazon GameLift Servers has successfully verified the build files and is now launching a fleet instance.
- **FLEET_STATE_ACTIVATING** -- Fleet status changed from BUILDING to ACTIVATING. Amazon GameLift Servers is launching a game server process on the fleet instance and is testing its connectivity with the Amazon GameLift Servers service.

- **FLEET_STATE_ACTIVE** -- The fleet's status changed from **ACTIVATING** to **ACTIVE**. The fleet is now ready to host game sessions.
- **FLEET_STATE_ERROR** -- The Fleet's status changed to **ERROR**. Describe the fleet event message for more details.

Fleet creation events (ordered by fleet creation activity):

- **FLEET_BINARY_DOWNLOAD_FAILED** -- The build failed to download to the fleet instance.
- **FLEET_CREATION_EXTRACTING_BUILD** -- The game server build was successfully downloaded to an instance, and Amazon GameLift Servers is now extracting the build files from the uploaded build. Failure at this stage prevents a fleet from moving to **ACTIVE** status. Logs for this stage display a list of the files that are extracted and saved on the instance. Access the logs by using the URL in *PreSignedLogUrl*.
- **FLEET_CREATION_RUNNING_INSTALLER** -- The game server build files were successfully extracted, and Amazon GameLift Servers is now running the build's install script (if one is included). Failure in this stage prevents a fleet from moving to **ACTIVE** status. Logs for this stage list the installation steps and whether or not the install completed successfully. Access the logs by using the URL in *PreSignedLogUrl*.
- **FLEET_CREATION_COMPLETED_INSTALLER** -- The game server build files were successfully installed and validation of the installation will begin soon.
- **FLEET_CREATION_FAILED_INSTALLER** -- The installed failed while attempting to install the build files. This event indicates that the failure occurred before Amazon GameLift Servers could start validation.
- **FLEET_CREATION_VALIDATING_RUNTIME_CONFIG** -- The build process was successful, and the GameLift is now verifying that the game server launch paths, which are specified in the fleet's runtime configuration, exist. If any listed launch path exists, Amazon GameLift Servers tries to launch a game server process and waits for the process to report ready. Failures in this stage prevent a fleet from moving to **ACTIVE** status. Logs for this stage list the launch paths in the runtime configuration and indicate whether each is found. Access the logs by using the URL in *PreSignedLogUrl*.
- **FLEET_VALIDATION_LAUNCH_PATH_NOT_FOUND** -- Validation of the runtime configuration failed because the executable specified in a launch path does not exist on the instance.
- **FLEET_VALIDATION_EXECUTABLE_RUNTIME_FAILURE** -- Validation of the runtime configuration failed because the executable specified in a launch path failed to run on the fleet instance.

- **FLEET_VALIDATION_TIMED_OUT** -- Validation of the fleet at the end of creation timed out. Try fleet creation again.
- **FLEET_ACTIVATION_FAILED** -- The fleet failed to successfully complete one of the steps in the fleet activation process. This event code indicates that the game build was successfully downloaded to a fleet instance, built, and validated, but was not able to start a server process. For more information, see [Debug Fleet Creation Issues](#).
- **FLEET_ACTIVATION_FAILED_NO_INSTANCES** -- Fleet creation was not able to obtain any instances based on the input fleet attributes. Try again at a different time or choose a different combination of fleet attributes such as fleet type, instance type, etc.
- **FLEET_INITIALIZATION_FAILED** -- A generic exception occurred during fleet creation. Describe the fleet event message for more details.

VPC peering events:

- **FLEET_VPC_PEERING_SUCCEEDED** -- A VPC peering connection has been established between the VPC for an Amazon GameLift Servers fleet and a VPC in your AWS account.
- **FLEET_VPC_PEERING_FAILED** -- A requested VPC peering connection has failed. Event details and status information provide additional detail. A common reason for peering failure is that the two VPCs have overlapping CIDR blocks of IPv4 addresses. To resolve this, change the CIDR block for the VPC in your AWS account. For more information on VPC peering failures, see <https://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/invalid-peering-configurations.html>
- **FLEET_VPC_PEERING_DELETED** -- A VPC peering connection has been successfully deleted.

Spot instance events:

- **INSTANCE_INTERRUPTED** -- A spot instance was interrupted by EC2 with a two-minute notification.
- **INSTANCE_RECYCLED** -- A spot instance was determined to have a high risk of interruption and is scheduled to be recycled once it has no active game sessions.

Server process events:

- **SERVER_PROCESS_INVALID_PATH** -- The game server executable or script could not be found based on the Fleet runtime configuration. Check that the launch path is correct based on the operating system of the Fleet.
- **SERVER_PROCESS_SDK_INITIALIZATION_TIMEOUT** -- The server process did not call `InitSDK()` within the time expected (5 minutes). Check your game session log to see why

`InitSDK()` was not called in time. This event is not emitted for managed container fleets and Anywhere fleets unless they're deployed with the Amazon GameLift Servers Agent.

- `SERVER_PROCESS_PROCESS_READY_TIMEOUT` -- The server process did not call `ProcessReady()` within the time expected (5 minutes) after calling `InitSDK()`. Check your game session log to see why `ProcessReady()` was not called in time.
- `SERVER_PROCESS_CRASHED` -- The server process exited without calling `ProcessEnding()`. Check your game session log to see why `ProcessEnding()` was not called.
- `SERVER_PROCESS_TERMINATED_UNHEALTHY` -- The server process did not report a valid health check for too long and was therefore terminated by GameLift. Check your game session log to see if the thread became stuck processing a synchronous task for too long.
- `SERVER_PROCESS_FORCE_TERMINATED` -- The server process did not exit cleanly within the time expected after `OnProcessTerminate()` was sent. Check your game session log to see why termination took longer than expected.
- `SERVER_PROCESS_PROCESS_EXIT_TIMEOUT` -- The server process did not exit cleanly within the time expected (30 seconds) after calling `ProcessEnding()`. Check your game session log to see why termination took longer than expected.

Game session events:

- `GAME_SESSION_ACTIVATION_TIMEOUT` -- `GameSession` failed to activate within the expected time. Check your game session log to see why `ActivateGameSession()` took longer to complete than expected.

Other fleet events:

- `FLEET_SCALING_EVENT` -- A change was made to the fleet's capacity settings (desired instances, minimum/maximum scaling limits). Event messaging includes the new capacity settings.
- `FLEET_NEW_GAME_SESSION_PROTECTION_POLICY_UPDATED` -- A change was made to the fleet's game session protection policy setting. Event messaging includes both the old and new policy setting.
- `FLEET_DELETED` -- A request to delete a fleet was initiated.
- `FLEET_EXPIRED` -- The fleet has been expired. The fleet is scaled down to zero instances and can no longer host game sessions.
- `GENERIC_EVENT` -- An unspecified event has occurred.

Type: String

Valid Values: GENERIC_EVENT | FLEET_CREATED | FLEET_DELETED
| FLEET_SCALING_EVENT | FLEET_STATE_DOWNLOADING
| FLEET_STATE_VALIDATING | FLEET_STATE_BUILDING |
FLEET_STATE_ACTIVATING | FLEET_STATE_ACTIVE | FLEET_STATE_ERROR |
FLEET_STATE_PENDING | FLEET_STATE_CREATING | FLEET_STATE_CREATED
| FLEET_STATE_UPDATING | FLEET_INITIALIZATION_FAILED |
FLEET_BINARY_DOWNLOAD_FAILED | FLEET_VALIDATION_LAUNCH_PATH_NOT_FOUND |
FLEET_VALIDATION_EXECUTABLE_RUNTIME_FAILURE | FLEET_VALIDATION_TIMED_OUT
| FLEET_ACTIVATION_FAILED | FLEET_ACTIVATION_FAILED_NO_INSTANCES
| FLEET_NEW_GAME_SESSION_PROTECTION_POLICY_UPDATED |
SERVER_PROCESS_INVALID_PATH | SERVER_PROCESS_SDK_INITIALIZATION_TIMEOUT
| SERVER_PROCESS_PROCESS_READY_TIMEOUT | SERVER_PROCESS_CRASHED |
SERVER_PROCESS_TERMINATED_UNHEALTHY | SERVER_PROCESS_FORCE_TERMINATED |
SERVER_PROCESS_PROCESS_EXIT_TIMEOUT |
SERVER_PROCESS_SDK_INITIALIZATION_FAILED |
SERVER_PROCESS_MISCONFIGURED_CONTAINER_PORT |
GAME_SESSION_ACTIVATION_TIMEOUT | FLEET_CREATION_EXTRACTING_BUILD |
FLEET_CREATION_RUNNING_INSTALLER |
FLEET_CREATION_VALIDATING_RUNTIME_CONFIG | FLEET_VPC_PEERING_SUCCEEDED
| FLEET_VPC_PEERING_FAILED | FLEET_VPC_PEERING_DELETED |
INSTANCE_INTERRUPTED | INSTANCE_RECYCLED | INSTANCE_REPLACED_UNHEALTHY
| FLEET_CREATION_COMPLETED_INSTALLER | FLEET_CREATION_FAILED_INSTALLER
| COMPUTE_LOG_UPLOAD_FAILED | GAME_SERVER_CONTAINER_GROUP_CRASHED |
PER_INSTANCE_CONTAINER_GROUP_CRASHED |
GAME_SERVER_CONTAINER_GROUP_REPLACED_UNHEALTHY | LOCATION_STATE_PENDING
| LOCATION_STATE_CREATING | LOCATION_STATE_CREATED |
LOCATION_STATE_ACTIVATING | LOCATION_STATE_ACTIVE |
LOCATION_STATE_UPDATING | LOCATION_STATE_ERROR | LOCATION_STATE_DELETING
| LOCATION_STATE_DELETED

Required: No

EventId

A unique identifier for a fleet event.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

EventTime

Time stamp indicating when this event occurred. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

Message

Additional information related to the event.

Type: String

Length Constraints: Minimum length of 1.

Required: No

PreSignedLogUrl

Location of stored logs with additional detail that is related to the event. This is useful for debugging issues. The URL is valid for 15 minutes. You can also access fleet creation logs through the Amazon GameLift Servers console.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

ResourceId

A unique identifier for an event resource, such as a fleet ID.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

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](#)

FilterConfiguration

A list of fleet locations where a game session queue can place new game sessions. You can use a filter to temporarily exclude specific locations from receiving placements. For queues that have multi-location fleets, you can use a filter configuration allow placement with some, but not all, of a fleet's locations.

Contents

Note

In the following list, the required parameters are described first.

AllowedLocations

A list of locations to allow game session placement in, in the form of AWS Region codes such as `us-west-2`.

Type: Array of strings

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

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-]+`

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](#)

FleetAttributes

Describes an Amazon GameLift Servers fleet of game hosting resources. Attributes differ based on the fleet's compute type, as follows:

- EC2 fleet attributes identify a Build resource (for fleets with customer game server builds) or a Script resource (for Amazon GameLift Servers Realtime fleets).
- Amazon GameLift Servers Anywhere fleets have an abbreviated set of attributes, because most fleet configurations are set directly on the fleet's computes. Attributes include fleet identifiers and descriptive properties, creation/termination time, and fleet status.

Returned by: https://docs.aws.amazon.com/gamelift/latest/apireference/API_DescribeFleetAttributes

Contents

Note

In the following list, the required parameters are described first.

AnywhereConfiguration

A set of attributes that are specific to an Anywhere fleet.

Type: [AnywhereConfiguration](#) object

Required: No

BuildArn

The Amazon Resource Name ([ARN](#)) associated with the Amazon GameLift Servers build resource that is deployed on instances in this fleet. In a GameLift build ARN, the resource ID matches the BuildId value. This attribute is used with fleets where ComputeType is "EC2".

Type: String

Pattern: `^arn:.*:build\/build-\/S+`

Required: No

BuildId

A unique identifier for the build resource that is deployed on instances in this fleet. This attribute is used with fleets where `ComputeType` is "EC2".

Type: String

Pattern: `^build-\S+`

Required: No

CertificateConfiguration

Determines whether a TLS/SSL certificate is generated for a fleet. This feature must be enabled when creating the fleet. All instances in a fleet share the same certificate.

Type: [CertificateConfiguration](#) object

Required: No

ComputeType

The type of compute resource used to host your game servers. You can use your own compute resources with Amazon GameLift Servers Anywhere or use Amazon EC2 instances with managed Amazon GameLift Servers.

Type: String

Valid Values: EC2 | ANYWHERE

Required: No

CreationTime

A time stamp indicating when this data object was created. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

Description

A human-readable description of the fleet.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

FleetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`. In a GameLift fleet ARN, the resource ID matches the `FleetId` value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\]\+$`

Required: No

FleetId

A unique identifier for the fleet.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\]\+`

Required: No

FleetType

Indicates whether the fleet uses On-Demand or Spot instances. For more information, see [On-Demand versus Spot Instances](#). This fleet property can't be changed after the fleet is created.

Type: String

Valid Values: `ON_DEMAND` | `SPOT`

Required: No

InstanceRoleArn

A unique identifier for an IAM role that manages access to your AWS services. With an instance role ARN set, any application that runs on an instance in this fleet can assume the role, including install scripts, server processes, and daemons (background processes). Create a role or look up a role's ARN by using the [IAM dashboard](#) in the AWS Management Console. Learn more

about using on-box credentials for your game servers at [Access external resources from a game server](#). This attribute is used with fleets where `ComputeType` is EC2.

Type: String

Length Constraints: Minimum length of 1.

Required: No

InstanceRoleCredentialsProvider

Indicates that fleet instances maintain a shared credentials file for the IAM role defined in `InstanceRoleArn`. Shared credentials allow applications that are deployed with the game server executable to communicate with other AWS resources. This property is used only when the game server is integrated with the server SDK version 5.x. For more information about using shared credentials, see [Communicate with other AWS resources from your fleets](#). This attribute is used with fleets where `ComputeType` is EC2.

Type: String

Valid Values: `SHARED_CREDENTIAL_FILE`

Required: No

InstanceType

The Amazon EC2 instance type that the fleet uses. Instance type determines the computing resources of each instance in the fleet, including CPU, memory, storage, and networking capacity. See [Amazon Elastic Compute Cloud Instance Types](#) for detailed descriptions. This attribute is used with fleets where `ComputeType` is EC2.

Type: String

Valid Values: `t2.micro` | `t2.small` | `t2.medium` | `t2.large` | `c3.large` | `c3.xlarge` | `c3.2xlarge` | `c3.4xlarge` | `c3.8xlarge` | `c4.large` | `c4.xlarge` | `c4.2xlarge` | `c4.4xlarge` | `c4.8xlarge` | `c5.large` | `c5.xlarge` | `c5.2xlarge` | `c5.4xlarge` | `c5.9xlarge` | `c5.12xlarge` | `c5.18xlarge` | `c5.24xlarge` | `c5a.large` | `c5a.xlarge` | `c5a.2xlarge` | `c5a.4xlarge` | `c5a.8xlarge` | `c5a.12xlarge` | `c5a.16xlarge` | `c5a.24xlarge` | `r3.large` | `r3.xlarge` | `r3.2xlarge` | `r3.4xlarge` | `r3.8xlarge` | `r4.large` | `r4.xlarge` | `r4.2xlarge` | `r4.4xlarge` | `r4.8xlarge` | `r4.16xlarge` | `r5.large` | `r5.xlarge` | `r5.2xlarge` | `r5.4xlarge` | `r5.8xlarge` | `r5.12xlarge` | `r5.16xlarge` | `r5.24xlarge` | `r5a.large` | `r5a.xlarge` | `r5a.2xlarge` |

r5a.4xlarge | r5a.8xlarge | r5a.12xlarge | r5a.16xlarge | r5a.24xlarge
| m3.medium | m3.large | m3.xlarge | m3.2xlarge | m4.large | m4.xlarge
| m4.2xlarge | m4.4xlarge | m4.10xlarge | m5.large | m5.xlarge |
m5.2xlarge | m5.4xlarge | m5.8xlarge | m5.12xlarge | m5.16xlarge |
m5.24xlarge | m5a.large | m5a.xlarge | m5a.2xlarge | m5a.4xlarge |
m5a.8xlarge | m5a.12xlarge | m5a.16xlarge | m5a.24xlarge | c5d.large
| c5d.xlarge | c5d.2xlarge | c5d.4xlarge | c5d.9xlarge | c5d.12xlarge
| c5d.18xlarge | c5d.24xlarge | c6a.large | c6a.xlarge | c6a.2xlarge |
c6a.4xlarge | c6a.8xlarge | c6a.12xlarge | c6a.16xlarge | c6a.24xlarge
| c6i.large | c6i.xlarge | c6i.2xlarge | c6i.4xlarge | c6i.8xlarge |
c6i.12xlarge | c6i.16xlarge | c6i.24xlarge | r5d.large | r5d.xlarge |
r5d.2xlarge | r5d.4xlarge | r5d.8xlarge | r5d.12xlarge | r5d.16xlarge
| r5d.24xlarge | m6g.medium | m6g.large | m6g.xlarge | m6g.2xlarge |
m6g.4xlarge | m6g.8xlarge | m6g.12xlarge | m6g.16xlarge | c6g.medium
| c6g.large | c6g.xlarge | c6g.2xlarge | c6g.4xlarge | c6g.8xlarge |
c6g.12xlarge | c6g.16xlarge | r6g.medium | r6g.large | r6g.xlarge |
r6g.2xlarge | r6g.4xlarge | r6g.8xlarge | r6g.12xlarge | r6g.16xlarge
| c6gn.medium | c6gn.large | c6gn.xlarge | c6gn.2xlarge | c6gn.4xlarge
| c6gn.8xlarge | c6gn.12xlarge | c6gn.16xlarge | c7g.medium | c7g.large
| c7g.xlarge | c7g.2xlarge | c7g.4xlarge | c7g.8xlarge | c7g.12xlarge
| c7g.16xlarge | r7g.medium | r7g.large | r7g.xlarge | r7g.2xlarge |
r7g.4xlarge | r7g.8xlarge | r7g.12xlarge | r7g.16xlarge | m7g.medium
| m7g.large | m7g.xlarge | m7g.2xlarge | m7g.4xlarge | m7g.8xlarge |
m7g.12xlarge | m7g.16xlarge | g5g.xlarge | g5g.2xlarge | g5g.4xlarge
| g5g.8xlarge | g5g.16xlarge | r6i.large | r6i.xlarge | r6i.2xlarge |
r6i.4xlarge | r6i.8xlarge | r6i.12xlarge | r6i.16xlarge | c6gd.medium |
c6gd.large | c6gd.xlarge | c6gd.2xlarge | c6gd.4xlarge | c6gd.8xlarge |
c6gd.12xlarge | c6gd.16xlarge | c6in.large | c6in.xlarge | c6in.2xlarge
| c6in.4xlarge | c6in.8xlarge | c6in.12xlarge | c6in.16xlarge |
c7a.medium | c7a.large | c7a.xlarge | c7a.2xlarge | c7a.4xlarge |
c7a.8xlarge | c7a.12xlarge | c7a.16xlarge | c7gd.medium | c7gd.large |
c7gd.xlarge | c7gd.2xlarge | c7gd.4xlarge | c7gd.8xlarge | c7gd.12xlarge
| c7gd.16xlarge | c7gn.medium | c7gn.large | c7gn.xlarge | c7gn.2xlarge
| c7gn.4xlarge | c7gn.8xlarge | c7gn.12xlarge | c7gn.16xlarge |
c7i.large | c7i.xlarge | c7i.2xlarge | c7i.4xlarge | c7i.8xlarge |
c7i.12xlarge | c7i.16xlarge | m6a.large | m6a.xlarge | m6a.2xlarge |

m6a.4xlarge | m6a.8xlarge | m6a.12xlarge | m6a.16xlarge | m6gd.medium
| m6gd.large | m6gd.xlarge | m6gd.2xlarge | m6gd.4xlarge | m6gd.8xlarge
| m6gd.12xlarge | m6gd.16xlarge | m6i.large | m6i.xlarge | m6i.2xlarge
| m6i.4xlarge | m6i.8xlarge | m6i.12xlarge | m6i.16xlarge | m7a.medium
| m7a.large | m7a.xlarge | m7a.2xlarge | m7a.4xlarge | m7a.8xlarge |
m7a.12xlarge | m7a.16xlarge | m7gd.medium | m7gd.large | m7gd.xlarge
| m7gd.2xlarge | m7gd.4xlarge | m7gd.8xlarge | m7gd.12xlarge |
m7gd.16xlarge | m7i.large | m7i.xlarge | m7i.2xlarge | m7i.4xlarge |
m7i.8xlarge | m7i.12xlarge | m7i.16xlarge | r6gd.medium | r6gd.large |
r6gd.xlarge | r6gd.2xlarge | r6gd.4xlarge | r6gd.8xlarge | r6gd.12xlarge
| r6gd.16xlarge | r7a.medium | r7a.large | r7a.xlarge | r7a.2xlarge |
r7a.4xlarge | r7a.8xlarge | r7a.12xlarge | r7a.16xlarge | r7gd.medium |
r7gd.large | r7gd.xlarge | r7gd.2xlarge | r7gd.4xlarge | r7gd.8xlarge |
r7gd.12xlarge | r7gd.16xlarge | r7i.large | r7i.xlarge | r7i.2xlarge |
r7i.4xlarge | r7i.8xlarge | r7i.12xlarge | r7i.16xlarge | r7i.24xlarge |
r7i.48xlarge | c5ad.large | c5ad.xlarge | c5ad.2xlarge | c5ad.4xlarge |
c5ad.8xlarge | c5ad.12xlarge | c5ad.16xlarge | c5ad.24xlarge | c5n.large
| c5n.xlarge | c5n.2xlarge | c5n.4xlarge | c5n.9xlarge | c5n.18xlarge |
r5ad.large | r5ad.xlarge | r5ad.2xlarge | r5ad.4xlarge | r5ad.8xlarge |
r5ad.12xlarge | r5ad.16xlarge | r5ad.24xlarge | c6id.large | c6id.xlarge
| c6id.2xlarge | c6id.4xlarge | c6id.8xlarge | c6id.12xlarge |
c6id.16xlarge | c6id.24xlarge | c6id.32xlarge | c8g.medium | c8g.large
| c8g.xlarge | c8g.2xlarge | c8g.4xlarge | c8g.8xlarge | c8g.12xlarge
| c8g.16xlarge | c8g.24xlarge | c8g.48xlarge | m5ad.large | m5ad.xlarge
| m5ad.2xlarge | m5ad.4xlarge | m5ad.8xlarge | m5ad.12xlarge |
m5ad.16xlarge | m5ad.24xlarge | m5d.large | m5d.xlarge | m5d.2xlarge |
m5d.4xlarge | m5d.8xlarge | m5d.12xlarge | m5d.16xlarge | m5d.24xlarge
| m5dn.large | m5dn.xlarge | m5dn.2xlarge | m5dn.4xlarge | m5dn.8xlarge
| m5dn.12xlarge | m5dn.16xlarge | m5dn.24xlarge | m5n.large |
m5n.xlarge | m5n.2xlarge | m5n.4xlarge | m5n.8xlarge | m5n.12xlarge |
m5n.16xlarge | m5n.24xlarge | m6id.large | m6id.xlarge | m6id.2xlarge
| m6id.4xlarge | m6id.8xlarge | m6id.12xlarge | m6id.16xlarge
| m6id.24xlarge | m6id.32xlarge | m6idn.large | m6idn.xlarge |
m6idn.2xlarge | m6idn.4xlarge | m6idn.8xlarge | m6idn.12xlarge |
m6idn.16xlarge | m6idn.24xlarge | m6idn.32xlarge | m6in.large |
m6in.xlarge | m6in.2xlarge | m6in.4xlarge | m6in.8xlarge | m6in.12xlarge

| m6in.16xlarge | m6in.24xlarge | m6in.32xlarge | m8g.medium | m8g.large
| m8g.xlarge | m8g.2xlarge | m8g.4xlarge | m8g.8xlarge | m8g.12xlarge
| m8g.16xlarge | m8g.24xlarge | m8g.48xlarge | r5dn.large | r5dn.xlarge
| r5dn.2xlarge | r5dn.4xlarge | r5dn.8xlarge | r5dn.12xlarge |
r5dn.16xlarge | r5dn.24xlarge | r5n.large | r5n.xlarge | r5n.2xlarge |
r5n.4xlarge | r5n.8xlarge | r5n.12xlarge | r5n.16xlarge | r5n.24xlarge
| r6a.large | r6a.xlarge | r6a.2xlarge | r6a.4xlarge | r6a.8xlarge |
r6a.12xlarge | r6a.16xlarge | r6a.24xlarge | r6a.32xlarge | r6a.48xlarge
| r6id.large | r6id.xlarge | r6id.2xlarge | r6id.4xlarge | r6id.8xlarge
| r6id.12xlarge | r6id.16xlarge | r6id.24xlarge | r6id.32xlarge
| r6idn.large | r6idn.xlarge | r6idn.2xlarge | r6idn.4xlarge |
r6idn.8xlarge | r6idn.12xlarge | r6idn.16xlarge | r6idn.24xlarge |
r6idn.32xlarge | r6in.large | r6in.xlarge | r6in.2xlarge | r6in.4xlarge
| r6in.8xlarge | r6in.12xlarge | r6in.16xlarge | r6in.24xlarge |
r6in.32xlarge | r8g.medium | r8g.large | r8g.xlarge | r8g.2xlarge |
r8g.4xlarge | r8g.8xlarge | r8g.12xlarge | r8g.16xlarge | r8g.24xlarge |
r8g.48xlarge | m4.16xlarge | c6a.32xlarge | c6a.48xlarge | c6i.32xlarge
| r6i.24xlarge | r6i.32xlarge | c6in.24xlarge | c6in.32xlarge |
c7a.24xlarge | c7a.32xlarge | c7a.48xlarge | c7i.24xlarge | c7i.48xlarge
| m6a.24xlarge | m6a.32xlarge | m6a.48xlarge | m6i.24xlarge |
m6i.32xlarge | m7a.24xlarge | m7a.32xlarge | m7a.48xlarge | m7i.24xlarge
| m7i.48xlarge | r7a.24xlarge | r7a.32xlarge | r7a.48xlarge

Required: No

LogPaths

This parameter is no longer used. Game session log paths are now defined using the Amazon GameLift Servers server API `ProcessReady()` `logParameters`. See more information in the [Server API Reference](#).

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

MetricGroups

Name of a metric group that metrics for this fleet are added to. In Amazon CloudWatch, you can view aggregated metrics for fleets that are in a metric group. A fleet can be included in only one metric group at a time. This attribute is used with fleets where `ComputeType` is `EC2`.

Type: Array of strings

Array Members: Maximum number of 1 item.

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

Required: No

Name

A descriptive label that is associated with a fleet. Fleet names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

NewGameSessionProtectionPolicy

The type of game session protection to set on all new instances that are started in the fleet. This attribute is used with fleets where `ComputeType` is `EC2`.

- **NoProtection** -- The game session can be terminated during a scale-down event.
- **FullProtection** -- If the game session is in an `ACTIVE` status, it cannot be terminated during a scale-down event.

Type: String

Valid Values: `NoProtection` | `FullProtection`

Required: No

OperatingSystem

The operating system of the fleet's computing resources. A fleet's operating system is determined by the OS of the build or script that is deployed on this fleet. This attribute is used with fleets where `ComputeType` is `EC2`.

Note

Amazon Linux 2 (AL2) will reach end of support on 6/30/2026. See more details in the [Amazon Linux 2 FAQs](#). For game servers that are hosted on AL2 and use server SDK version 4.x for Amazon GameLift Servers, first update the game server build to server SDK 5.x, and then deploy to AL2023 instances. See [Migrate to server SDK version 5](#).

Type: String

Valid Values: WINDOWS_2012 | AMAZON_LINUX | AMAZON_LINUX_2 | WINDOWS_2016 | AMAZON_LINUX_2023 | WINDOWS_2022

Required: No

PlayerGatewayConfiguration

Configuration settings for player gateway on this fleet.

Type: [PlayerGatewayConfiguration](#) object

Required: No

PlayerGatewayMode

Indicates whether player gateway is enabled for this fleet. Player gateway provides benefits such as DDoS protection with negligible impact to latency.

If ENABLED or REQUIRED, game clients can use player gateway to connect with the game server. If DISABLED, game clients cannot use player gateway. Instead, they have to directly connect to the game server.

Type: String

Valid Values: DISABLED | ENABLED | REQUIRED

Required: No

ResourceCreationLimitPolicy

A policy that puts limits on the number of game sessions that a player can create within a specified span of time. With this policy, you can control players' ability to consume available resources.

The policy is evaluated when a player tries to create a new game session. On receiving a `CreateGameSession` request, Amazon GameLift Servers checks that the player (identified by `CreatorId`) has created fewer than game session limit in the specified time period.

The purpose of this policy is to prevent a single player from consuming a large share of available hosting resources. For example, setting `NewGameSessionsPerCreator` to 4 and `PolicyPeriodInMinutes` to 10 limits each player to creating 4 game sessions every 10 minutes. Setting these values too high (for example, 200 game sessions every 1000 minutes) still allows a single player to rapidly consume resources. We recommend keeping these values small.

Type: [ResourceCreationLimitPolicy](#) object

Required: No

ScriptArn

The Amazon Resource Name ([ARN](#)) associated with the GameLift script resource that is deployed on instances in this fleet. In a GameLift script ARN, the resource ID matches the `ScriptId` value.

Type: String

Pattern: `^arn:.*:script\/script-\S+`

Required: No

ScriptId

A unique identifier for the Realtime script resource that is deployed on instances in this fleet. This attribute is used with fleets where `ComputeType` is "EC2".

Type: String

Pattern: `^script-\S+`

Required: No

ServerLaunchParameters

This parameter is no longer used. Server launch parameters are now defined using the fleet's runtime configuration. Requests that use this parameter continue to be valid.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `[A-Za-z0-9_:.+\\/\\"'`- =@;{} ,?'\[\]""]+`

Required: No

ServerLaunchPath

This parameter is no longer used. Server launch paths are now defined using the fleet's [RuntimeConfiguration](#). Requests that use this parameter continue to be valid.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `[A-Za-z0-9_:.+\\/\\"'`-]+`

Required: No

Status

Current status of the fleet. Possible fleet statuses include the following:

- **NEW** -- A new fleet resource has been defined and Amazon GameLift Servers has started creating the fleet. Desired instances is set to 1.
- **DOWNLOADING/VALIDATING/BUILDING** -- Amazon GameLift Servers is download the game server build, running install scripts, and then validating the build files. When complete, Amazon GameLift Servers launches a fleet instance.
- **ACTIVATING** -- Amazon GameLift Servers is launching a game server process and testing its connectivity with the Amazon GameLift Servers service.
- **ACTIVE** -- The fleet is now ready to host game sessions.
- **ERROR** -- An error occurred when downloading, validating, building, or activating the fleet.
- **EXPIRED** -- The fleet has been expired. The fleet is scaled down to zero instances and cannot host new game sessions.
- **DELETING** -- Hosts are responding to a delete fleet request.
- **TERMINATED** -- The fleet no longer exists.

Type: String

Valid Values: NEW | DOWNLOADING | VALIDATING | BUILDING | ACTIVATING | ACTIVE | DELETING | ERROR | TERMINATED | NOT_FOUND

Required: No

StoppedActions

A list of fleet activity that has been suspended using [StopFleetActions](#). This includes fleet auto-scaling. This attribute is used with fleets where ComputeType is EC2.

Type: Array of strings

Array Members: Fixed number of 1 item.

Valid Values: AUTO_SCALING

Required: No

TerminationTime

A time stamp indicating when this data object was terminated. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

See Also

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

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

FleetCapacity

Current resource capacity settings for managed EC2 fleets and managed container fleets. For multi-location fleets, location values might refer to a fleet's remote location or its home Region.

Returned by: [DescribeFleetCapacity](#), [DescribeFleetLocationCapacity](#), [UpdateFleetCapacity](#)

Contents

Note

In the following list, the required parameters are described first.

FleetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\]\+$`

Required: No

FleetId

A unique identifier for the fleet associated with the location.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\]\+$`

Required: No

GameServerContainerGroupCounts

The number and status of game server container groups deployed in a container fleet.

Type: [GameServerContainerGroupCounts](#) object

Required: No

InstanceCounts

The current number of instances in the fleet, listed by instance status. Counts for pending and terminating instances might be non-zero if the fleet is adjusting to a scaling event or if access to resources is temporarily affected.

Type: [EC2InstanceCounts](#) object

Required: No

InstanceType

The Amazon EC2 instance type that is used for instances in a fleet. Instance type determines the computing resources in use, including CPU, memory, storage, and networking capacity. See [Amazon Elastic Compute Cloud Instance Types](#) for detailed descriptions.

Type: String

Valid Values: t2.micro | t2.small | t2.medium | t2.large | c3.large | c3.xlarge | c3.2xlarge | c3.4xlarge | c3.8xlarge | c4.large | c4.xlarge | c4.2xlarge | c4.4xlarge | c4.8xlarge | c5.large | c5.xlarge | c5.2xlarge | c5.4xlarge | c5.9xlarge | c5.12xlarge | c5.18xlarge | c5.24xlarge | c5a.large | c5a.xlarge | c5a.2xlarge | c5a.4xlarge | c5a.8xlarge | c5a.12xlarge | c5a.16xlarge | c5a.24xlarge | r3.large | r3.xlarge | r3.2xlarge | r3.4xlarge | r3.8xlarge | r4.large | r4.xlarge | r4.2xlarge | r4.4xlarge | r4.8xlarge | r4.16xlarge | r5.large | r5.xlarge | r5.2xlarge | r5.4xlarge | r5.8xlarge | r5.12xlarge | r5.16xlarge | r5.24xlarge | r5a.large | r5a.xlarge | r5a.2xlarge | r5a.4xlarge | r5a.8xlarge | r5a.12xlarge | r5a.16xlarge | r5a.24xlarge | m3.medium | m3.large | m3.xlarge | m3.2xlarge | m4.large | m4.xlarge | m4.2xlarge | m4.4xlarge | m4.10xlarge | m5.large | m5.xlarge | m5.2xlarge | m5.4xlarge | m5.8xlarge | m5.12xlarge | m5.16xlarge | m5.24xlarge | m5a.large | m5a.xlarge | m5a.2xlarge | m5a.4xlarge | m5a.8xlarge | m5a.12xlarge | m5a.16xlarge | m5a.24xlarge | c5d.large | c5d.xlarge | c5d.2xlarge | c5d.4xlarge | c5d.9xlarge | c5d.12xlarge | c5d.18xlarge | c5d.24xlarge | c6a.large | c6a.xlarge | c6a.2xlarge | c6a.4xlarge | c6a.8xlarge | c6a.12xlarge | c6a.16xlarge | c6a.24xlarge

| c6i.large | c6i.xlarge | c6i.2xlarge | c6i.4xlarge | c6i.8xlarge |
c6i.12xlarge | c6i.16xlarge | c6i.24xlarge | r5d.large | r5d.xlarge |
r5d.2xlarge | r5d.4xlarge | r5d.8xlarge | r5d.12xlarge | r5d.16xlarge
| r5d.24xlarge | m6g.medium | m6g.large | m6g.xlarge | m6g.2xlarge |
m6g.4xlarge | m6g.8xlarge | m6g.12xlarge | m6g.16xlarge | c6g.medium
| c6g.large | c6g.xlarge | c6g.2xlarge | c6g.4xlarge | c6g.8xlarge |
c6g.12xlarge | c6g.16xlarge | r6g.medium | r6g.large | r6g.xlarge |
r6g.2xlarge | r6g.4xlarge | r6g.8xlarge | r6g.12xlarge | r6g.16xlarge
| c6gn.medium | c6gn.large | c6gn.xlarge | c6gn.2xlarge | c6gn.4xlarge
| c6gn.8xlarge | c6gn.12xlarge | c6gn.16xlarge | c7g.medium | c7g.large
| c7g.xlarge | c7g.2xlarge | c7g.4xlarge | c7g.8xlarge | c7g.12xlarge
| c7g.16xlarge | r7g.medium | r7g.large | r7g.xlarge | r7g.2xlarge |
r7g.4xlarge | r7g.8xlarge | r7g.12xlarge | r7g.16xlarge | m7g.medium
| m7g.large | m7g.xlarge | m7g.2xlarge | m7g.4xlarge | m7g.8xlarge |
m7g.12xlarge | m7g.16xlarge | g5g.xlarge | g5g.2xlarge | g5g.4xlarge
| g5g.8xlarge | g5g.16xlarge | r6i.large | r6i.xlarge | r6i.2xlarge |
r6i.4xlarge | r6i.8xlarge | r6i.12xlarge | r6i.16xlarge | c6gd.medium |
c6gd.large | c6gd.xlarge | c6gd.2xlarge | c6gd.4xlarge | c6gd.8xlarge |
c6gd.12xlarge | c6gd.16xlarge | c6in.large | c6in.xlarge | c6in.2xlarge
| c6in.4xlarge | c6in.8xlarge | c6in.12xlarge | c6in.16xlarge |
c7a.medium | c7a.large | c7a.xlarge | c7a.2xlarge | c7a.4xlarge |
c7a.8xlarge | c7a.12xlarge | c7a.16xlarge | c7gd.medium | c7gd.large |
c7gd.xlarge | c7gd.2xlarge | c7gd.4xlarge | c7gd.8xlarge | c7gd.12xlarge
| c7gd.16xlarge | c7gn.medium | c7gn.large | c7gn.xlarge | c7gn.2xlarge
| c7gn.4xlarge | c7gn.8xlarge | c7gn.12xlarge | c7gn.16xlarge |
c7i.large | c7i.xlarge | c7i.2xlarge | c7i.4xlarge | c7i.8xlarge |
c7i.12xlarge | c7i.16xlarge | m6a.large | m6a.xlarge | m6a.2xlarge |
m6a.4xlarge | m6a.8xlarge | m6a.12xlarge | m6a.16xlarge | m6gd.medium
| m6gd.large | m6gd.xlarge | m6gd.2xlarge | m6gd.4xlarge | m6gd.8xlarge
| m6gd.12xlarge | m6gd.16xlarge | m6i.large | m6i.xlarge | m6i.2xlarge
| m6i.4xlarge | m6i.8xlarge | m6i.12xlarge | m6i.16xlarge | m7a.medium
| m7a.large | m7a.xlarge | m7a.2xlarge | m7a.4xlarge | m7a.8xlarge |
m7a.12xlarge | m7a.16xlarge | m7gd.medium | m7gd.large | m7gd.xlarge
| m7gd.2xlarge | m7gd.4xlarge | m7gd.8xlarge | m7gd.12xlarge |
m7gd.16xlarge | m7i.large | m7i.xlarge | m7i.2xlarge | m7i.4xlarge |
m7i.8xlarge | m7i.12xlarge | m7i.16xlarge | r6gd.medium | r6gd.large |

r6gd.xlarge | r6gd.2xlarge | r6gd.4xlarge | r6gd.8xlarge | r6gd.12xlarge
| r6gd.16xlarge | r7a.medium | r7a.large | r7a.xlarge | r7a.2xlarge |
r7a.4xlarge | r7a.8xlarge | r7a.12xlarge | r7a.16xlarge | r7gd.medium |
r7gd.large | r7gd.xlarge | r7gd.2xlarge | r7gd.4xlarge | r7gd.8xlarge |
r7gd.12xlarge | r7gd.16xlarge | r7i.large | r7i.xlarge | r7i.2xlarge |
r7i.4xlarge | r7i.8xlarge | r7i.12xlarge | r7i.16xlarge | r7i.24xlarge |
r7i.48xlarge | c5ad.large | c5ad.xlarge | c5ad.2xlarge | c5ad.4xlarge |
c5ad.8xlarge | c5ad.12xlarge | c5ad.16xlarge | c5ad.24xlarge | c5n.large
| c5n.xlarge | c5n.2xlarge | c5n.4xlarge | c5n.9xlarge | c5n.18xlarge |
r5ad.large | r5ad.xlarge | r5ad.2xlarge | r5ad.4xlarge | r5ad.8xlarge |
r5ad.12xlarge | r5ad.16xlarge | r5ad.24xlarge | c6id.large | c6id.xlarge
| c6id.2xlarge | c6id.4xlarge | c6id.8xlarge | c6id.12xlarge |
c6id.16xlarge | c6id.24xlarge | c6id.32xlarge | c8g.medium | c8g.large
| c8g.xlarge | c8g.2xlarge | c8g.4xlarge | c8g.8xlarge | c8g.12xlarge
| c8g.16xlarge | c8g.24xlarge | c8g.48xlarge | m5ad.large | m5ad.xlarge
| m5ad.2xlarge | m5ad.4xlarge | m5ad.8xlarge | m5ad.12xlarge |
m5ad.16xlarge | m5ad.24xlarge | m5d.large | m5d.xlarge | m5d.2xlarge |
m5d.4xlarge | m5d.8xlarge | m5d.12xlarge | m5d.16xlarge | m5d.24xlarge
| m5dn.large | m5dn.xlarge | m5dn.2xlarge | m5dn.4xlarge | m5dn.8xlarge
| m5dn.12xlarge | m5dn.16xlarge | m5dn.24xlarge | m5n.large |
m5n.xlarge | m5n.2xlarge | m5n.4xlarge | m5n.8xlarge | m5n.12xlarge |
m5n.16xlarge | m5n.24xlarge | m6id.large | m6id.xlarge | m6id.2xlarge
| m6id.4xlarge | m6id.8xlarge | m6id.12xlarge | m6id.16xlarge
| m6id.24xlarge | m6id.32xlarge | m6idn.large | m6idn.xlarge |
m6idn.2xlarge | m6idn.4xlarge | m6idn.8xlarge | m6idn.12xlarge |
m6idn.16xlarge | m6idn.24xlarge | m6idn.32xlarge | m6in.large |
m6in.xlarge | m6in.2xlarge | m6in.4xlarge | m6in.8xlarge | m6in.12xlarge
| m6in.16xlarge | m6in.24xlarge | m6in.32xlarge | m8g.medium | m8g.large
| m8g.xlarge | m8g.2xlarge | m8g.4xlarge | m8g.8xlarge | m8g.12xlarge
| m8g.16xlarge | m8g.24xlarge | m8g.48xlarge | r5dn.large | r5dn.xlarge
| r5dn.2xlarge | r5dn.4xlarge | r5dn.8xlarge | r5dn.12xlarge |
r5dn.16xlarge | r5dn.24xlarge | r5n.large | r5n.xlarge | r5n.2xlarge |
r5n.4xlarge | r5n.8xlarge | r5n.12xlarge | r5n.16xlarge | r5n.24xlarge
| r6a.large | r6a.xlarge | r6a.2xlarge | r6a.4xlarge | r6a.8xlarge |
r6a.12xlarge | r6a.16xlarge | r6a.24xlarge | r6a.32xlarge | r6a.48xlarge
| r6id.large | r6id.xlarge | r6id.2xlarge | r6id.4xlarge | r6id.8xlarge

| r6id.12xlarge | r6id.16xlarge | r6id.24xlarge | r6id.32xlarge
| r6idn.large | r6idn.xlarge | r6idn.2xlarge | r6idn.4xlarge |
r6idn.8xlarge | r6idn.12xlarge | r6idn.16xlarge | r6idn.24xlarge |
r6idn.32xlarge | r6in.large | r6in.xlarge | r6in.2xlarge | r6in.4xlarge
| r6in.8xlarge | r6in.12xlarge | r6in.16xlarge | r6in.24xlarge |
r6in.32xlarge | r8g.medium | r8g.large | r8g.xlarge | r8g.2xlarge |
r8g.4xlarge | r8g.8xlarge | r8g.12xlarge | r8g.16xlarge | r8g.24xlarge |
r8g.48xlarge | m4.16xlarge | c6a.32xlarge | c6a.48xlarge | c6i.32xlarge
| r6i.24xlarge | r6i.32xlarge | c6in.24xlarge | c6in.32xlarge |
c7a.24xlarge | c7a.32xlarge | c7a.48xlarge | c7i.24xlarge | c7i.48xlarge
| m6a.24xlarge | m6a.32xlarge | m6a.48xlarge | m6i.24xlarge |
m6i.32xlarge | m7a.24xlarge | m7a.32xlarge | m7a.48xlarge | m7i.24xlarge
| m7i.48xlarge | r7a.24xlarge | r7a.32xlarge | r7a.48xlarge

Required: No

Location

The fleet location for the instance count information, expressed as an AWS Region code, such as `us-west-2`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-]+`

Required: No

ManagedCapacityConfiguration

Configuration settings for managed capacity scaling.

Type: [ManagedCapacityConfiguration](#) 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](#)

FleetDeployment

Describes a container fleet deployment with updates to the fleet.

Contents

Note

In the following list, the required parameters are described first.

CreationTime

A time stamp indicating when this data object was created. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

DeploymentConfiguration

Instructions for how to deploy updates to a container fleet and what actions to take if the deployment fails.

Type: [DeploymentConfiguration](#) object

Required: No

DeploymentId

A unique identifier for the deployment.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `^[a-zA-Z0-9\-\_]+$`

Required: No

DeploymentStatus

The status of fleet deployment activity in the location.

- `IN_PROGRESS` -- The deployment is in progress.
- `IMPAIRED` -- The deployment failed and the fleet has some impaired containers.
- `COMPLETE` -- The deployment has completed successfully.
- `ROLLBACK_IN_PROGRESS` -- The deployment failed and rollback has been initiated.
- `ROLLBACK_IN_COMPLETE` -- The deployment failed and rollback has been completed.
- `CANCELLED` -- The deployment was cancelled.

Type: String

Valid Values: `IN_PROGRESS` | `IMPAIRED` | `COMPLETE` | `ROLLBACK_IN_PROGRESS` | `ROLLBACK_IN_COMPLETE` | `CANCELLED` | `PENDING`

Required: No

FleetId

A unique identifier for the container fleet.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\-]+`

Required: No

GameServerBinaryArn

The unique identifier for the version of the game server container group definition that is being deployed.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `[a-zA-Z0-9:/-]+`

Required: No

PerInstanceBinaryArn

The unique identifier for the version of the per-instance container group definition that is being deployed.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: [a-zA-Z0-9:/-]+

Required: No

RollbackGameServerBinaryArn

The unique identifier for the version of the game server container group definition to roll back to if deployment fails. Amazon GameLift Servers sets this property to the container group definition version that the fleet used when it was last active.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: [a-zA-Z0-9:/-]+

Required: No

RollbackPerInstanceBinaryArn

The unique identifier for the version of the per-instance container group definition to roll back to if deployment fails. Amazon GameLift Servers sets this property to the container group definition version that the fleet used when it was last active.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: [a-zA-Z0-9:/-]+

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](#)

FleetUtilization

Current resource utilization statistics in a specified fleet or location. The location value might refer to a fleet's remote location or its home region.

Contents

Note

In the following list, the required parameters are described first.

ActiveGameSessionCount

The number of active game sessions that are currently being hosted across all instances in the fleet location.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

ActiveServerProcessCount

The number of server processes in ACTIVE status that are currently running across all instances in the fleet location.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

CurrentPlayerSessionCount

The number of active player sessions that are currently being hosted across all instances in the fleet location.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

FleetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\]\+$`

Required: No

FleetId

A unique identifier for the fleet associated with the location.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\]\+`

Required: No

Location

The fleet location for the fleet utilization information, expressed as an AWS Region code, such as `us-west-2`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-\]\+`

Required: No

MaximumPlayerSessionCount

The maximum number of players allowed across all game sessions that are currently being hosted across all instances in the fleet location.

Type: Integer

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](#)

GameProperty

This key-value pair can store custom data about a game session. For example, you might use a GameProperty to track a game session's map, level of difficulty, or remaining time. The difficulty level could be specified like this: {"Key": "difficulty", "Value": "Novice"}.

You can set game properties when creating a game session. You can also modify game properties of an active game session. When searching for game sessions, you can filter on game property keys and values. You can't delete game properties from a game session.

For examples of working with game properties, see [Create a game session with properties](#).

Contents

Note

In the following list, the required parameters are described first.

Key

The game property identifier.

Note

- Avoid using periods (".") in property keys if you plan to search for game sessions by properties. Property keys containing periods cannot be searched and will be filtered out from search results due to search index limitations.
- If you use SearchGameSessions API, there is a limit of 500 game property keys across all game sessions and all fleets per region. If the limit is exceeded, there will potentially be game session entries missing from SearchGameSessions API results.

Type: String

Length Constraints: Maximum length of 32.

Required: Yes

Value

The game property value.

Type: String

Length Constraints: Maximum length of 96.

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](#)

GameServer

Properties describing a game server that is running on an instance in a game server group.

A game server is created by a successful call to `RegisterGameServer` and deleted by calling `DeregisterGameServer`. A game server is claimed to host a game session by calling `ClaimGameServer`.

Contents

Note

In the following list, the required parameters are described first.

ClaimStatus

Indicates when an available game server has been reserved for gameplay but has not yet started hosting a game. Once it is claimed, the game server remains in CLAIMED status for a maximum of one minute. During this time, game clients connect to the game server to start the game and trigger the game server to update its utilization status. After one minute, the game server claim status reverts to null.

Type: String

Valid Values: CLAIMED

Required: No

ConnectionInfo

The port and IP address that must be used to establish a client connection to the game server.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `.*\S.*`

Required: No

GameServerData

A set of custom game server properties, formatted as a single string value. This data is passed to a game client or service when it requests information on game servers.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `.*\S.*`

Required: No

GameServerGroupArn

The ARN identifier for the game server group where the game server is located.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^arn:.*:gameservergroup\[a-zA-Z0-9-\.\]*`

Required: No

GameServerGroupName

A unique identifier for the game server group where the game server is running.

Type: String

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

Pattern: `[a-zA-Z0-9-\.\]+`

Required: No

GameServerId

A custom string that uniquely identifies the game server. Game server IDs are developer-defined and are unique across all game server groups in an AWS account.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `[a-zA-Z0-9-\.]`⁺

Required: No

InstanceId

The unique identifier for the instance where the game server is running. This ID is available in the instance metadata. EC2 instance IDs use a 17-character format, for example: `i-1234567890abcdef0`.

Type: String

Length Constraints: Fixed length of 19.

Pattern: `^i-[0-9a-zA-Z]{17}$`

Required: No

LastClaimTime

Timestamp that indicates the last time the game server was claimed. The format is a number expressed in Unix time as milliseconds (for example `"1469498468.057"`). This value is used to calculate when a claimed game server's status should revert to null.

Type: Timestamp

Required: No

LastHealthCheckTime

Timestamp that indicates the last time the game server was updated with health status. The format is a number expressed in Unix time as milliseconds (for example `"1469498468.057"`). After game server registration, this property is only changed when a game server update specifies a health check value.

Type: Timestamp

Required: No

RegistrationTime

Timestamp that indicates when the game server registered. The format is a number expressed in Unix time as milliseconds (for example `"1469498468.057"`).

Type: Timestamp

Required: No

UtilizationStatus

Indicates whether the game server is currently available for new games or is busy. Possible statuses include:

- **AVAILABLE** - The game server is available to be claimed. A game server that has been claimed remains in this status until it reports game hosting activity.
- **UTILIZED** - The game server is currently hosting a game session with players.

Type: String

Valid Values: AVAILABLE | UTILIZED

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](#)

GameServerContainerDefinition

Describes the game server container in an existing game server container group. A game server container identifies a container image with your game server build. A game server container is automatically considered essential; if an essential container fails, the entire container group restarts.

You can update a container definition and deploy the updates to an existing fleet. When creating or updating a game server container group definition, use the property https://docs.aws.amazon.com/gamelift/latest/apireference/API_GameServerContainerDefinitionInput.

Part of: [ContainerGroupDefinition](#)

Returned by: [CreateContainerGroupDefinition](#), [DescribeContainerGroupDefinition](#), [ListContainerGroupDefinitions](#), [ListContainerGroupDefinitionVersions](#), [UpdateContainerGroupDefinition](#)

Contents

Note

In the following list, the required parameters are described first.

ContainerName

The container definition identifier. Container names are unique within a container group definition.

Type: String

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

Pattern: `^[a-zA-Z0-9\-\_]+$`

Required: No

DependsOn

Indicates that the container relies on the status of other containers in the same container group during startup and shutdown sequences. A container might have dependencies on multiple containers.

Type: Array of [ContainerDependency](#) objects

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

Required: No

EnvironmentOverride

A set of environment variables that's passed to the container on startup. See the [ContainerDefinition::environment](#) parameter in the *Amazon Elastic Container Service API Reference*.

Type: Array of [ContainerEnvironment](#) objects

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

Required: No

ImageUri

The URI to the image that Amazon GameLift Servers uses when deploying this container to a container fleet. For a more specific identifier, see `ResolvedImageDigest`.

Type: String

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

Pattern: `^[a-zA-Z0-9-_\.\@\/:]+$`

Required: No

LinuxCapabilities

Linux-specific modifications that are applied to the default Docker container configuration, such as Linux capabilities. For more information see [LinuxCapabilities](#).

Type: [LinuxCapabilities](#) object

Required: No

MountPoints

A mount point that binds a path inside the container to a file or directory on the host system and lets it access the file or directory.

Type: Array of [ContainerMountPoint](#) objects

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

Required: No

PortConfiguration

The set of ports that are available to bind to processes in the container. For example, a game server process requires a container port to allow game clients to connect to it. Container ports aren't directly accessed by inbound traffic. Amazon GameLift Servers maps these container ports to externally accessible connection ports, which are assigned as needed from the container fleet's `ConnectionPortRange`.

Type: [ContainerPortConfiguration](#) object

Required: No

ResolvedImageDigest

A unique and immutable identifier for the container image. The digest is a SHA 256 hash of the container image manifest.

Type: String

Pattern: `^sha256:[a-fA-F0-9]{64}$`

Required: No

ServerSdkVersion

The Amazon GameLift Servers server SDK version that the game server is integrated with. Only game servers using 5.2.0 or higher are compatible with container fleets.

Type: String

Length Constraints: Maximum length of 128.

Pattern: `^\d+\.\d+\.\d+$`

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](#)

GameServerContainerDefinitionInput

Describes the configuration for a container that runs your game server executable. This definition includes container configuration, resources, and start instructions. Use this data type when creating or updating a game server container group definition. For properties of a deployed container, see [GameServerContainerDefinition](#). A game server container is automatically considered essential; if an essential container fails, the entire container group restarts.

Use with: [CreateContainerGroupDefinition](#), [UpdateContainerGroupDefinition](#)

Contents

Note

In the following list, the required parameters are described first.

ContainerName

A string that uniquely identifies the container definition within a container group.

Type: String

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

Pattern: `^[a-zA-Z0-9\-\_]+$`

Required: Yes

ImageUri

The location of the container image to deploy to a container fleet. Provide an image in an Amazon Elastic Container Registry public or private repository. The repository must be in the same AWS account and AWS Region where you're creating the container group definition. For limits on image size, see [Amazon GameLift Servers endpoints and quotas](#). You can use any of the following image URI formats:

- Image ID only: `[AWS account].dkr.ecr.[AWS region].amazonaws.com/[repository ID]`
- Image ID and digest: `[AWS account].dkr.ecr.[AWS region].amazonaws.com/[repository ID]@[digest]`

- Image ID and tag: `[AWS account].dkr.ecr.[AWS region].amazonaws.com/[repository ID]:[tag]`

Type: String

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

Pattern: `^[a-zA-Z0-9-_\.\@\/:]+$`

Required: Yes

PortConfiguration

A set of ports that Amazon GameLift Servers can assign to processes in a container. The container port configuration must have enough ports for each container process that accepts inbound traffic connections. For example, a game server process requires a container port to allow game clients to connect to it. A container port configuration can have one or more container port ranges. Each range specifies starting and ending values as well as the supported network protocol.

Container ports aren't directly accessed by inbound traffic. Amazon GameLift Servers maps each container port to an externally accessible connection port (see the container fleet property `ConnectionPortRange`).

Type: [ContainerPortConfiguration](#) object

Required: Yes

ServerSdkVersion

The Amazon GameLift Servers server SDK version that the game server is integrated with. Only game servers using 5.2.0 or higher are compatible with container fleets.

Type: String

Length Constraints: Maximum length of 128.

Pattern: `^\d+\.\d+\.\d+$`

Required: Yes

DependsOn

Establishes dependencies between this container and the status of other containers in the same container group. A container can have dependencies on multiple different containers.

You can use dependencies to establish a startup/shutdown sequence across the container group. For example, you might specify that *ContainerB* has a START dependency on *ContainerA*. This dependency means that *ContainerB* can't start until after *ContainerA* has started. This dependency is reversed on shutdown, which means that *ContainerB* must shut down before *ContainerA* can shut down.

Type: Array of [ContainerDependency](#) objects

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

Required: No

EnvironmentOverride

A set of environment variables to pass to the container on startup. See the [ContainerDefinition::environment](#) parameter in the *Amazon Elastic Container Service API Reference*.

Type: Array of [ContainerEnvironment](#) objects

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

Required: No

LinuxCapabilities

Linux-specific modifications that are applied to the default Docker container configuration, such as Linux capabilities. For more information see [LinuxCapabilities](#).

Type: [LinuxCapabilities](#) object

Required: No

MountPoints

A mount point that binds a path inside the container to a file or directory on the host system and lets it access the file or directory.

Type: Array of [ContainerMountPoint](#) 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](#)

GameServerContainerGroupCounts

The number and status of game server container groups that are deployed across a container fleet. Combine this count with the number of server processes that each game server container group runs to learn how many game sessions the fleet is capable of hosting concurrently. For example, if a fleet has 50 game server container groups, and the game server container in each group runs 1 game server process, then the fleet has the capacity to run host 50 game sessions at a time.

Returned by: https://docs.aws.amazon.com/gamelift/latest/apireference/API_DescribeFleetCapacity.html, https://docs.aws.amazon.com/gamelift/latest/apireference/API_DescribeFleetLocationCapacity.html

Contents

Note

In the following list, the required parameters are described first.

ACTIVE

The number of container groups that have active game sessions.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

IDLE

The number of container groups that have no active game sessions.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

PENDING

The number of container groups that are starting up but haven't yet registered.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

TERMINATING

The number of container groups that are in the process of shutting down.

Type: Integer

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](#)

GameServerGroup

Properties that describe a game server group resource. A game server group manages certain properties related to a corresponding Amazon EC2 Auto Scaling group.

A game server group is created by a successful call to `CreateGameServerGroup` and deleted by calling `DeleteGameServerGroup`. Game server group activity can be temporarily suspended and resumed by calling `SuspendGameServerGroup` and `ResumeGameServerGroup`, respectively.

Contents

Note

In the following list, the required parameters are described first.

AutoScalingGroupArn

A generated unique ID for the Amazon EC2 Auto Scaling group that is associated with this game server group.

Type: String

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

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*`

Required: No

BalancingStrategy

Indicates how Amazon GameLift Servers FleetIQ balances the use of Spot Instances and On-Demand Instances in the game server group. Method options include the following:

- `SPOT_ONLY` - Only Spot Instances are used in the game server group. If Spot Instances are unavailable or not viable for game hosting, the game server group provides no hosting capacity until Spot Instances can again be used. Until then, no new instances are started, and the existing nonviable Spot Instances are terminated (after current gameplay ends) and are not replaced.
- `SPOT_PREFERRED` - (default value) Spot Instances are used whenever available in the game server group. If Spot Instances are unavailable, the game server group continues to provide

hosting capacity by falling back to On-Demand Instances. Existing nonviable Spot Instances are terminated (after current gameplay ends) and are replaced with new On-Demand Instances.

- `ON_DEMAND_ONLY` - Only On-Demand Instances are used in the game server group. No Spot Instances are used, even when available, while this balancing strategy is in force.

Type: String

Valid Values: `SPOT_ONLY` | `SPOT_PREFERRED` | `ON_DEMAND_ONLY`

Required: No

CreationTime

A time stamp indicating when this data object was created. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

GameServerGroupArn

A generated unique ID for the game server group.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^arn:.*:gameservergroup\/[a-zA-Z0-9-\.]*`

Required: No

GameServerGroupName

A developer-defined identifier for the game server group. The name is unique for each Region in each AWS account.

Type: String

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

Pattern: `[a-zA-Z0-9-\.]+`

Required: No

GameServerProtectionPolicy

A flag that indicates whether instances in the game server group are protected from early termination. Unprotected instances that have active game servers running might be terminated during a scale-down event, causing players to be dropped from the game. Protected instances cannot be terminated while there are active game servers running except in the event of a forced game server group deletion (see). An exception to this is with Spot Instances, which can be terminated by AWS regardless of protection status.

Type: String

Valid Values: NO_PROTECTION | FULL_PROTECTION

Required: No

InstanceDefinitions

The set of Amazon EC2 instance types that Amazon GameLift Servers FleetIQ can use when balancing and automatically scaling instances in the corresponding Auto Scaling group.

Type: Array of [InstanceDefinition](#) objects

Array Members: Minimum number of 2 items. Maximum number of 20 items.

Required: No

LastUpdatedTime

A timestamp that indicates when this game server group was last updated.

Type: Timestamp

Required: No

RoleArn

The Amazon Resource Name ([ARN](#)) for an IAM role that allows Amazon GameLift Servers to access your Amazon EC2 Auto Scaling groups.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: ^arn:.*:role\/[\w+=,.\@-]+

Required: No

Status

The current status of the game server group. Possible statuses include:

- **NEW** - Amazon GameLift Servers FleetIQ has validated the `CreateGameServerGroup()` request.
- **ACTIVATING** - Amazon GameLift Servers FleetIQ is setting up a game server group, which includes creating an Auto Scaling group in your AWS account.
- **ACTIVE** - The game server group has been successfully created.
- **DELETE_SCHEDULED** - A request to delete the game server group has been received.
- **DELETING** - Amazon GameLift Servers FleetIQ has received a valid `DeleteGameServerGroup()` request and is processing it. Amazon GameLift Servers FleetIQ must first complete and release hosts before it deletes the Auto Scaling group and the game server group.
- **DELETED** - The game server group has been successfully deleted.
- **ERROR** - The asynchronous processes of activating or deleting a game server group has failed, resulting in an error state.

Type: String

Valid Values: NEW | ACTIVATING | ACTIVE | DELETE_SCHEDULED | DELETING | DELETED | ERROR

Required: No

StatusReason

Additional information about the current game server group status. This information might provide additional insight on groups that are in ERROR status.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

SuspendedActions

A list of activities that are currently suspended for this game server group. If this property is empty, all activities are occurring.

Type: Array of strings

Array Members: Fixed number of 1 item.

Valid Values: REPLACE_INSTANCE_TYPES

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](#)

GameServerGroupAutoScalingPolicy

Configuration settings for intelligent automatic scaling that uses target tracking. These settings are used to add an Auto Scaling policy when creating the corresponding Auto Scaling group. After the Auto Scaling group is created, all updates to Auto Scaling policies, including changing this policy and adding or removing other policies, is done directly on the Auto Scaling group.

Contents

Note

In the following list, the required parameters are described first.

TargetTrackingConfiguration

Settings for a target-based scaling policy applied to Auto Scaling group. These settings are used to create a target-based policy that tracks the Amazon GameLift Servers FleetIQ metric "PercentUtilizedGameServers" and specifies a target value for the metric. As player usage changes, the policy triggers to adjust the game server group capacity so that the metric returns to the target value.

Type: [TargetTrackingConfiguration](#) object

Required: Yes

EstimatedInstanceWarmup

Length of time, in seconds, it takes for a new instance to start new game server processes and register with Amazon GameLift Servers FleetIQ. Specifying a warm-up time can be useful, particularly with game servers that take a long time to start up, because it avoids prematurely starting new instances.

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](#)

GameServerInstance

Additional properties, including status, that describe an EC2 instance in a game server group. Instance configurations are set with game server group properties (see `DescribeGameServerGroup` and with the EC2 launch template that was used when creating the game server group.

Retrieve game server instances for a game server group by calling `DescribeGameServerInstances`.

Contents

Note

In the following list, the required parameters are described first.

GameServerGroupArn

A generated unique identifier for the game server group that includes the game server instance.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^arn:.*:gameservergroup\[a-zA-Z0-9-\.\]*`

Required: No

GameServerGroupName

A developer-defined identifier for the game server group that includes the game server instance. The name is unique for each Region in each AWS account.

Type: String

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

Pattern: `[a-zA-Z0-9-\.\]+`

Required: No

InstanceId

The unique identifier for the instance where the game server is running. This ID is available in the instance metadata. EC2 instance IDs use a 17-character format, for example: `i-1234567890abcdef0`.

Type: String

Length Constraints: Fixed length of 19.

Pattern: `^i-[0-9a-zA-Z]{17}$`

Required: No

InstanceStatus

Current status of the game server instance

Type: String

Valid Values: ACTIVE | DRAINING | SPOT_TERMINATING

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](#)

GameSession

Properties describing a game session.

A game session in **ACTIVE** status can host players. When a game session ends, its status is set to **TERMINATED**.

Amazon GameLift Servers retains a game session resource for 30 days after the game session ends. You can reuse idempotency token values after this time. Game session logs are retained for 14 days.

[All APIs by task](#)

Contents

Note

In the following list, the required parameters are described first.

ComputeName

A descriptive label for the compute resource. The compute resource that is hosting the game session. For EC2 fleets, this is the EC2 instance ID. For Container fleets, each game server container group on a fleet instance is assigned a compute name. For Anywhere fleets, this is the custom compute name.

Type: String

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

Pattern: `^[a-zA-Z0-9\-\_]+(\.[a-zA-Z0-9\-\_]+)?`

Required: No

CreationTime

A time stamp indicating when this data object was created. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

CreatorId

A unique identifier for a player. This ID is used to enforce a resource protection policy (if one exists), that limits the number of game sessions a player can create.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

CurrentPlayerSessionCount

Number of players currently in the game session.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

DnsName

The DNS identifier assigned to the instance that is running the game session. Values have the following format:

- TLS-enabled fleets: <unique identifier>.<region identifier>.amazongamelift.com.
- Non-TLS-enabled fleets: ec2-<unique identifier>.compute.amazonaws.com. (See [Amazon EC2 Instance IP Addressing](#).)

When connecting to a game session that is running on a TLS-enabled fleet, you must use the DNS name, not the IP address.

Type: String

Required: No

FleetArn

The Amazon Resource Name ([ARN](#)) associated with the GameLift fleet that this game session is running on.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet\|[a-z]*fleet-[a-zA-Z0-9\-\-]+$`

Required: No

FleetId

A unique identifier for the fleet that the game session is running on.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\-]+$`

Required: No

GameProperties

A set of key-value pairs that can store custom data in a game session. For example: `{"Key": "difficulty", "Value": "novice"}`.

Note

- Avoid using periods (".") in property keys if you plan to search for game sessions by properties. Property keys containing periods cannot be searched and will be filtered out from search results due to search index limitations.
- If you use `SearchGameSessions` API, there is a limit of 500 game property keys across all game sessions and all fleets per region. If the limit is exceeded, there will potentially be game session entries missing from `SearchGameSessions` API results.

Type: Array of [GameProperty](#) objects

Array Members: Maximum number of 16 items.

Required: No

GameSessionData

A set of custom game session properties, formatted as a single string value. This data is passed to a game server process with a request to start a new game session. For more information, see [Start a game session](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 262144.

Required: No

GameSessionId

An identifier for the game session that is unique across all regions. The value is always a full ARN in the following format: For Home Region game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>`. For Remote Location game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<location>/<ID string>`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

IpAddress

The IP address of the game session. To connect to a Amazon GameLift Servers game server, an app needs both the IP address and port number.

Type: String

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

Pattern: `^[0-9A-Fa-f\:\\.]+`

Required: No

Location

The fleet location where the game session is running. This value might specify the fleet's home Region or a remote location. Location is expressed as an AWS Region code such as `us-west-2`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-]+`

Required: No

MatchmakerData

Information about the matchmaking process that resulted in the game session, if matchmaking was used. Data is in JSON syntax, formatted as a string. Information includes the matchmaker ID as well as player attributes and team assignments. For more details on matchmaker data, see [Match Data](#). Matchmaker data is updated whenever new players are added during a successful backfill (see [StartMatchBackfill](#)).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 390000.

Required: No

MaximumPlayerSessionCount

The maximum number of players that can be connected simultaneously to the game session.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

Name

A descriptive label that is associated with a game session. Session names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

PlayerGatewayStatus

Indicates whether player gateway is available for use for this game session. Note, even if a fleet has PlayerGatewayMode configured as ENABLED, player gateway might not be available in a specific location. For more information about locations where player gateway is supported, see [Amazon GameLift Servers service locations](#).

Possible values include:

- ENABLED -- Player gateway is available for routing player connections for this game session.

- **DISABLED** -- Player gateway is not available for this game session.

Type: String

Valid Values: DISABLED | ENABLED

Required: No

PlayerSessionCreationPolicy

Indicates whether the game session is accepting new players.

Type: String

Valid Values: ACCEPT_ALL | DENY_ALL

Required: No

Port

The port number for the game session. To connect to a Amazon GameLift Servers game server, an app needs both the IP address and port number.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 60000.

Required: No

Status

Current status of the game session. A game session must have an **ACTIVE** status to have player sessions.

Type: String

Valid Values: ACTIVE | ACTIVATING | TERMINATED | TERMINATING | ERROR

Required: No

StatusReason

Provides additional information about game session status.

- **INTERRUPTED** -- The game session was hosted on an EC2 Spot instance that was reclaimed, causing the active game session to be stopped.

- `TRIGGERED_ON_PROCESS_TERMINATE` – The game session was stopped by calling `TerminateGameSession` with the termination mode `TRIGGER_ON_PROCESS_TERMINATE`.
- `FORCE_TERMINATED` – The game session was stopped by calling `TerminateGameSession` with the termination mode `FORCE_TERMINATE`.

Type: String

Valid Values: `INTERRUPTED` | `TRIGGERED_ON_PROCESS_TERMINATE` | `FORCE_TERMINATED`

Required: No

TerminationTime

A time stamp indicating when this data object was terminated. Format is a number expressed in Unix time as milliseconds (for example `"1469498468.057"`).

Type: Timestamp

Required: No

See Also

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

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

GameSessionConnectionInfo

Connection information for a new game session that is created in response to a start matchmaking request. Once a match is made, the FlexMatch engine creates a new game session for it. This information, including the game session endpoint and player sessions for each player in the original matchmaking request, is added to the matchmaking ticket.

Contents

Note

In the following list, the required parameters are described first.

DnsName

The DNS identifier assigned to the instance that is running the game session. Values have the following format:

- TLS-enabled fleets: `<unique identifier>.<region identifier>.amazongamelift.com`.
- Non-TLS-enabled fleets: `ec2-<unique identifier>.compute.amazonaws.com`. (See [Amazon EC2 Instance IP Addressing](#).)

When connecting to a game session that is running on a TLS-enabled fleet, you must use the DNS name, not the IP address.

Type: String

Required: No

GameSessionArn

An identifier for the game session that is unique across all regions. The value is always a full ARN in the following format: For Home Region game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>`. For Remote Location game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<location>/<ID string>`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: [a-zA-Z0-9:/-]+

Required: No

IpAddress

The IP address of the game session. To connect to a Amazon GameLift Servers game server, an app needs both the IP address and port number.

Type: String

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

Pattern: ^[0-9A-Fa-f\:\.]+

Required: No

MatchedPlayerSessions

A collection of player session IDs, one for each player ID that was included in the original matchmaking request.

Type: Array of [MatchedPlayerSession](#) objects

Required: No

PlayerGatewayStatus

The current status of player gateway for the game session. Note, even if a fleet has PlayerGatewayMode configured as ENABLED, player gateway might not be available in a specific location. For more information about locations where player gateway is supported, see [supported locations](#).

Possible values include:

- ENABLED -- Player gateway is available for this game session.
- DISABLED -- Player gateway is not available for this game session.

Type: String

Valid Values: DISABLED | ENABLED

Required: No

Port

The port number for the game session. To connect to a Amazon GameLift Servers game server, an app needs both the IP address and port number.

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](#)

GameSessionCreationLimitPolicy

A policy that puts limits on the number of game sessions that a player can create within a specified span of time. With this policy, you can control players' ability to consume available resources.

The policy is evaluated when a player tries to create a new game session. On receiving a `CreateGameSession` request, Amazon GameLift Servers checks that the player (identified by `CreatorId`) has created fewer than game session limit in the specified time period.

Contents

Note

In the following list, the required parameters are described first.

NewGameSessionsPerCreator

A policy that puts limits on the number of game sessions that a player can create within a specified span of time. With this policy, you can control players' ability to consume available resources.

The policy evaluates when a player tries to create a new game session. On receiving a `CreateGameSession` request, Amazon GameLift Servers checks that the player (identified by `CreatorId`) has created fewer than game session limit in the specified time period.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

PolicyPeriodInMinutes

The time span used in evaluating the resource creation limit policy.

Type: Integer

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](#)

GameSessionDetail

A game session's properties plus the protection policy currently in force.

Contents

Note

In the following list, the required parameters are described first.

GameSession

Object that describes a game session.

Type: [GameSession](#) object

Required: No

ProtectionPolicy

Current status of protection for the game session.

- **NoProtection** -- The game session can be terminated during a scale-down event.
- **FullProtection** -- If the game session is in an ACTIVE status, it cannot be terminated during a scale-down event.

Type: String

Valid Values: NoProtection | FullProtection

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](#)

GameSessionPlacement

Represents a potential game session placement, including the full details of the original placement request and the current status.

Note

If the game session placement status is PENDING, the properties for game session ID/ARN, region, IP address/DNS, and port aren't final. A game session is not active and ready to accept players until placement status reaches FULFILLED. When the placement is in PENDING status, Amazon GameLift Servers may attempt to place a game session multiple times before succeeding. With each attempt it creates a https://docs.aws.amazon.com/gamelift/latest/apireference/API_GameSession object and updates this placement object with the new game session properties.

Contents

Note

In the following list, the required parameters are described first.

DnsName

The DNS identifier assigned to the instance that is running the game session. Values have the following format:

- TLS-enabled fleets: <unique identifier>.<region identifier>.amazongamelift.com.
- Non-TLS-enabled fleets: ec2-<unique identifier>.compute.amazonaws.com. (See [Amazon EC2 Instance IP Addressing](#).)

When connecting to a game session that is running on a TLS-enabled fleet, you must use the DNS name, not the IP address.

Type: String

Required: No

EndTime

Time stamp indicating when this request was completed, canceled, or timed out.

Type: Timestamp

Required: No

GameProperties

A set of key-value pairs that can store custom data in a game session. For example: {"Key": "difficulty", "Value": "novice"}.

Note

- Avoid using periods (".") in property keys if you plan to search for game sessions by properties. Property keys containing periods cannot be searched and will be filtered out from search results due to search index limitations.
- If you use SearchGameSessions API, there is a limit of 500 game property keys across all game sessions and all fleets per region. If the limit is exceeded, there will potentially be game session entries missing from SearchGameSessions API results.

Type: Array of [GameProperty](#) objects

Array Members: Maximum number of 16 items.

Required: No

GameSessionArn

An identifier for the game session that is unique across all regions. The value is always a full ARN in the following format: For Home Region game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>`. For Remote Location game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<location>/<ID string>`. This value is the same as GameSessionId. This value isn't final until placement status is FULFILLED.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

GameSessionData

A set of custom game session properties, formatted as a single string value. This data is passed to a game server process with a request to start a new game session. For more information, see [Start a game session](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 262144.

Required: No

GameSessionId

An identifier for the game session that is unique across all regions. The value is always a full ARN in the following format: For Home Region game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>`. For Remote Location game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<location>/<ID string>`. This value is the same as `GameSessionArn`. This value isn't final until placement status is `FULFILLED`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

GameSessionName

A descriptive label that is associated with a game session. Session names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

GameSessionQueueName

A descriptive label that is associated with game session queue. Queue names must be unique within each Region.

Type: String

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

Pattern: [a-zA-Z0-9-]+

Required: No

GameSessionRegion

Name of the Region where the game session created by this placement request is running. This value isn't final until placement status is FULFILLED.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

IpAddress

The IP address of the game session. To connect to a Amazon GameLift Servers game server, an app needs both the IP address and port number. This value isn't final until placement status is FULFILLED.

Type: String

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

Pattern: ^[0-9A-Fa-f\:\.]+

Required: No

MatchmakerData

Information on the matchmaking process for this game. Data is in JSON syntax, formatted as a string. It identifies the matchmaking configuration used to create the match, and contains data on all players assigned to the match, including player attributes and team assignments. For more details on matchmaker data, see [Match Data](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 390000.

Required: No

MaximumPlayerSessionCount

The maximum number of players that can be connected simultaneously to the game session.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

PlacedPlayerSessions

A collection of information on player sessions created in response to the game session placement request. These player sessions are created only after a new game session is successfully placed (placement status is FULFILLED). This information includes the player ID, provided in the placement request, and a corresponding player session ID.

Type: Array of [PlacedPlayerSession](#) objects

Required: No

PlacementId

A unique identifier for a game session placement.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 48.

Pattern: [a-zA-Z0-9-]+

Required: No

PlayerGatewayStatus

The current status of player gateway for the game session placement. Note, even if a fleet has PlayerGatewayMode configured as ENABLED, player gateway might not be available in a specific location. For more information about locations where player gateway is supported, see [Amazon GameLift Servers service locations](#).

Possible values include:

- ENABLED -- Player gateway is available for this game session placement.
- DISABLED -- Player gateway is not available for this game session placement.

Type: String

Valid Values: DISABLED | ENABLED

Required: No

PlayerLatencies

A set of values, expressed in milliseconds, that indicates the amount of latency that a player experiences when connected to a fleet location (AWS Regions or custom locations for Amazon GameLift Servers Anywhere fleets).

Type: Array of [PlayerLatency](#) objects

Required: No

Port

The port number for the game session. To connect to a Amazon GameLift Servers game server, an app needs both the IP address and port number. This value isn't final until placement status is FULFILLED.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 60000.

Required: No

PriorityConfigurationOverride

An alternative priority list of locations that's included with a game session placement request. When provided, the list overrides a queue's location order list for this game session placement request only. The list might include AWS Regions, local zones, and custom locations (for Anywhere fleets). The fallback strategy tells Amazon GameLift Servers what action to take (if any) in the event that it failed to place a new game session.

Type: [PriorityConfigurationOverride](#) object

Required: No

StartTime

Time stamp indicating when this request was placed in the queue. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

Status

Current status of the game session placement request.

- **PENDING** -- The placement request is in the queue waiting to be processed. Game session properties are not yet final.
- **FULFILLED** -- A new game session has been successfully placed. Game session properties are now final.
- **CANCELLED** -- The placement request was canceled.
- **TIMED_OUT** -- A new game session was not successfully created before the time limit expired. You can resubmit the placement request as needed.
- **FAILED** -- Amazon GameLift Servers is not able to complete the process of placing the game session. Common reasons are the game session terminated before the placement process was completed, or an unexpected internal error.

Type: String

Valid Values: PENDING | FULFILLED | CANCELLED | TIMED_OUT | FAILED

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](#)

GameSessionQueue

Configuration for a game session placement mechanism that processes requests for new game sessions. A queue can be used on its own or as part of a matchmaking solution.

Contents

Note

In the following list, the required parameters are described first.

CustomEventData

Information that is added to all events that are related to this game session queue.

Type: String

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

Pattern: `[\s\S]*`

Required: No

Destinations

A list of fleets and/or fleet aliases that can be used to fulfill game session placement requests in the queue. Destinations are identified by either a fleet ARN or a fleet alias ARN, and are listed in order of placement preference.

Type: Array of [GameSessionQueueDestination](#) objects

Required: No

FilterConfiguration

A list of locations where a queue is allowed to place new game sessions. Locations are specified in the form of AWS Region codes, such as `us-west-2`. If this parameter is not set, game sessions can be placed in any queue location.

Type: [FilterConfiguration](#) object

Required: No

GameSessionQueueArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers game session queue resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::gamesessionqueue/<queue name>`. In a Amazon GameLift Servers game session queue ARN, the resource ID matches the *Name* value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^arn:.*:gamesessionqueue\/[a-zA-Z0-9-]+`

Required: No

Name

A descriptive label that is associated with game session queue. Queue names must be unique within each Region.

Type: String

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

Pattern: `[a-zA-Z0-9-]+`

Required: No

NotificationTarget

An SNS topic ARN that is set up to receive game session placement notifications. See [Setting up notifications for game session placement](#).

Type: String

Length Constraints: Minimum length of 0. Maximum length of 300.

Pattern: `[a-zA-Z0-9:_-]*(\\.fifo)?`

Required: No

PlayerLatencyPolicies

A set of policies that enforce a sliding cap on player latency when processing game sessions placement requests. Use multiple policies to gradually relax the cap over time if Amazon

GameLift Servers can't make a placement. Policies are evaluated in order starting with the lowest maximum latency value.

Type: Array of [PlayerLatencyPolicy](#) objects

Required: No

PriorityConfiguration

Custom settings to use when prioritizing destinations and locations for game session placements. This configuration replaces the FleetIQ default prioritization process. Priority types that are not explicitly named will be automatically applied at the end of the prioritization process.

Type: [PriorityConfiguration](#) object

Required: No

TimeoutInSeconds

The maximum time, in seconds, that a new game session placement request remains in the queue. When a request exceeds this time, the game session placement changes to a TIMED_OUT status.

Note

The minimum value is 10 and the maximum value is 600.

Type: Integer

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](#)

GameSessionQueueDestination

A fleet or alias designated in a game session queue. Queues fulfill requests for new game sessions by placing a new game session on any of the queue's destinations.

Contents

Note

In the following list, the required parameters are described first.

DestinationArn

The Amazon Resource Name (ARN) that is assigned to fleet or fleet alias. ARNs, which include a fleet ID or alias ID and a Region name, provide a unique identifier across all Regions.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: [a-zA-Z0-9:/-]+

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](#)

Instance

Represents a virtual computing instance that runs game server processes and hosts game sessions. In Amazon GameLift Servers, one or more instances make up a managed EC2 fleet.

Contents

Note

In the following list, the required parameters are described first.

CreationTime

A time stamp indicating when this data object was created. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

DnsName

The DNS identifier assigned to the instance that is running the game session. Values have the following format:

- TLS-enabled fleets: <unique identifier>.<region identifier>.amazongamelift.com.
- Non-TLS-enabled fleets: ec2-<unique identifier>.compute.amazonaws.com. (See [Amazon EC2 Instance IP Addressing](#).)

When connecting to a game session that is running on a TLS-enabled fleet, you must use the DNS name, not the IP address.

Type: String

Required: No

FleetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is

arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: ^arn:.*:[a-z]*fleet\[a-z\]*fleet-[a-zA-Z0-9\-\-]+\$

Required: No

FleetId

A unique identifier for the fleet that the instance belongs to.

Type: String

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

Pattern: ^[a-z]*fleet-[a-zA-Z0-9\-\-]+\$

Required: No

InstanceId

A unique identifier for the instance.

Type: String

Pattern: [a-zA-Z0-9\-\-]+

Required: No

IpAddress

IP address that is assigned to the instance.

Type: String

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

Pattern: ^[0-9A-Fa-f\:\.\-]+\$

Required: No

Location

The fleet location of the instance, expressed as an AWS Region code, such as us-west-2.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-\_]+`

Required: No

OperatingSystem

Operating system that is running on this EC2 instance.

Note

Amazon Linux 2 (AL2) will reach end of support on 6/30/2026. See more details in the [Amazon Linux 2 FAQs](#). For game servers that are hosted on AL2 and use server SDK version 4.x for Amazon GameLift Servers, first update the game server build to server SDK 5.x, and then deploy to AL2023 instances. See [Migrate to server SDK version 5](#).

Type: String

Valid Values: `WINDOWS_2012` | `AMAZON_LINUX` | `AMAZON_LINUX_2` | `WINDOWS_2016` | `AMAZON_LINUX_2023` | `WINDOWS_2022`

Required: No

Status

Current status of the instance. Possible statuses include the following:

- **PENDING** -- The instance is in the process of being created and launching server processes as defined in the fleet's run-time configuration.
- **ACTIVE** -- The instance has been successfully created and at least one server process has successfully launched and reported back to Amazon GameLift Servers that it is ready to host a game session. The instance is now considered ready to host game sessions.
- **TERMINATING** -- The instance is in the process of shutting down. This may happen to reduce capacity during a scaling down event or to recycle resources in the event of a problem.

Type: String

Valid Values: `PENDING` | `ACTIVE` | `TERMINATING`

Required: No

Type

EC2 instance type that defines the computing resources of this instance.

Type: String

Valid Values: t2.micro | t2.small | t2.medium | t2.large | c3.large | c3.xlarge | c3.2xlarge | c3.4xlarge | c3.8xlarge | c4.large | c4.xlarge | c4.2xlarge | c4.4xlarge | c4.8xlarge | c5.large | c5.xlarge | c5.2xlarge | c5.4xlarge | c5.9xlarge | c5.12xlarge | c5.18xlarge | c5.24xlarge | c5a.large | c5a.xlarge | c5a.2xlarge | c5a.4xlarge | c5a.8xlarge | c5a.12xlarge | c5a.16xlarge | c5a.24xlarge | r3.large | r3.xlarge | r3.2xlarge | r3.4xlarge | r3.8xlarge | r4.large | r4.xlarge | r4.2xlarge | r4.4xlarge | r4.8xlarge | r4.16xlarge | r5.large | r5.xlarge | r5.2xlarge | r5.4xlarge | r5.8xlarge | r5.12xlarge | r5.16xlarge | r5.24xlarge | r5a.large | r5a.xlarge | r5a.2xlarge | r5a.4xlarge | r5a.8xlarge | r5a.12xlarge | r5a.16xlarge | r5a.24xlarge | m3.medium | m3.large | m3.xlarge | m3.2xlarge | m4.large | m4.xlarge | m4.2xlarge | m4.4xlarge | m4.10xlarge | m5.large | m5.xlarge | m5.2xlarge | m5.4xlarge | m5.8xlarge | m5.12xlarge | m5.16xlarge | m5.24xlarge | m5a.large | m5a.xlarge | m5a.2xlarge | m5a.4xlarge | m5a.8xlarge | m5a.12xlarge | m5a.16xlarge | m5a.24xlarge | c5d.large | c5d.xlarge | c5d.2xlarge | c5d.4xlarge | c5d.9xlarge | c5d.12xlarge | c5d.18xlarge | c5d.24xlarge | c6a.large | c6a.xlarge | c6a.2xlarge | c6a.4xlarge | c6a.8xlarge | c6a.12xlarge | c6a.16xlarge | c6a.24xlarge | c6i.large | c6i.xlarge | c6i.2xlarge | c6i.4xlarge | c6i.8xlarge | c6i.12xlarge | c6i.16xlarge | c6i.24xlarge | r5d.large | r5d.xlarge | r5d.2xlarge | r5d.4xlarge | r5d.8xlarge | r5d.12xlarge | r5d.16xlarge | r5d.24xlarge | m6g.medium | m6g.large | m6g.xlarge | m6g.2xlarge | m6g.4xlarge | m6g.8xlarge | m6g.12xlarge | m6g.16xlarge | c6g.medium | c6g.large | c6g.xlarge | c6g.2xlarge | c6g.4xlarge | c6g.8xlarge | c6g.12xlarge | c6g.16xlarge | r6g.medium | r6g.large | r6g.xlarge | r6g.2xlarge | r6g.4xlarge | r6g.8xlarge | r6g.12xlarge | r6g.16xlarge | c6gn.medium | c6gn.large | c6gn.xlarge | c6gn.2xlarge | c6gn.4xlarge | c6gn.8xlarge | c6gn.12xlarge | c6gn.16xlarge | c7g.medium | c7g.large | c7g.xlarge | c7g.2xlarge | c7g.4xlarge | c7g.8xlarge | c7g.12xlarge

| c7g.16xlarge | r7g.medium | r7g.large | r7g.xlarge | r7g.2xlarge | r7g.4xlarge | r7g.8xlarge | r7g.12xlarge | r7g.16xlarge | m7g.medium | m7g.large | m7g.xlarge | m7g.2xlarge | m7g.4xlarge | m7g.8xlarge | m7g.12xlarge | m7g.16xlarge | g5g.xlarge | g5g.2xlarge | g5g.4xlarge | g5g.8xlarge | g5g.16xlarge | r6i.large | r6i.xlarge | r6i.2xlarge | r6i.4xlarge | r6i.8xlarge | r6i.12xlarge | r6i.16xlarge | c6gd.medium | c6gd.large | c6gd.xlarge | c6gd.2xlarge | c6gd.4xlarge | c6gd.8xlarge | c6gd.12xlarge | c6gd.16xlarge | c6in.large | c6in.xlarge | c6in.2xlarge | c6in.4xlarge | c6in.8xlarge | c6in.12xlarge | c6in.16xlarge | c7a.medium | c7a.large | c7a.xlarge | c7a.2xlarge | c7a.4xlarge | c7a.8xlarge | c7a.12xlarge | c7a.16xlarge | c7gd.medium | c7gd.large | c7gd.xlarge | c7gd.2xlarge | c7gd.4xlarge | c7gd.8xlarge | c7gd.12xlarge | c7gd.16xlarge | c7gn.medium | c7gn.large | c7gn.xlarge | c7gn.2xlarge | c7gn.4xlarge | c7gn.8xlarge | c7gn.12xlarge | c7gn.16xlarge | c7i.large | c7i.xlarge | c7i.2xlarge | c7i.4xlarge | c7i.8xlarge | c7i.12xlarge | c7i.16xlarge | m6a.large | m6a.xlarge | m6a.2xlarge | m6a.4xlarge | m6a.8xlarge | m6a.12xlarge | m6a.16xlarge | m6gd.medium | m6gd.large | m6gd.xlarge | m6gd.2xlarge | m6gd.4xlarge | m6gd.8xlarge | m6gd.12xlarge | m6gd.16xlarge | m6i.large | m6i.xlarge | m6i.2xlarge | m6i.4xlarge | m6i.8xlarge | m6i.12xlarge | m6i.16xlarge | m7a.medium | m7a.large | m7a.xlarge | m7a.2xlarge | m7a.4xlarge | m7a.8xlarge | m7a.12xlarge | m7a.16xlarge | m7gd.medium | m7gd.large | m7gd.xlarge | m7gd.2xlarge | m7gd.4xlarge | m7gd.8xlarge | m7gd.12xlarge | m7gd.16xlarge | m7i.large | m7i.xlarge | m7i.2xlarge | m7i.4xlarge | m7i.8xlarge | m7i.12xlarge | m7i.16xlarge | r6gd.medium | r6gd.large | r6gd.xlarge | r6gd.2xlarge | r6gd.4xlarge | r6gd.8xlarge | r6gd.12xlarge | r6gd.16xlarge | r7a.medium | r7a.large | r7a.xlarge | r7a.2xlarge | r7a.4xlarge | r7a.8xlarge | r7a.12xlarge | r7a.16xlarge | r7gd.medium | r7gd.large | r7gd.xlarge | r7gd.2xlarge | r7gd.4xlarge | r7gd.8xlarge | r7gd.12xlarge | r7gd.16xlarge | r7i.large | r7i.xlarge | r7i.2xlarge | r7i.4xlarge | r7i.8xlarge | r7i.12xlarge | r7i.16xlarge | r7i.24xlarge | r7i.48xlarge | c5ad.large | c5ad.xlarge | c5ad.2xlarge | c5ad.4xlarge | c5ad.8xlarge | c5ad.12xlarge | c5ad.16xlarge | c5ad.24xlarge | c5n.large | c5n.xlarge | c5n.2xlarge | c5n.4xlarge | c5n.9xlarge | c5n.18xlarge | r5ad.large | r5ad.xlarge | r5ad.2xlarge | r5ad.4xlarge | r5ad.8xlarge | r5ad.12xlarge | r5ad.16xlarge | r5ad.24xlarge | c6id.large | c6id.xlarge

| c6id.2xlarge | c6id.4xlarge | c6id.8xlarge | c6id.12xlarge |
c6id.16xlarge | c6id.24xlarge | c6id.32xlarge | c8g.medium | c8g.large
| c8g.xlarge | c8g.2xlarge | c8g.4xlarge | c8g.8xlarge | c8g.12xlarge
| c8g.16xlarge | c8g.24xlarge | c8g.48xlarge | m5ad.large | m5ad.xlarge
| m5ad.2xlarge | m5ad.4xlarge | m5ad.8xlarge | m5ad.12xlarge |
m5ad.16xlarge | m5ad.24xlarge | m5d.large | m5d.xlarge | m5d.2xlarge |
m5d.4xlarge | m5d.8xlarge | m5d.12xlarge | m5d.16xlarge | m5d.24xlarge
| m5dn.large | m5dn.xlarge | m5dn.2xlarge | m5dn.4xlarge | m5dn.8xlarge
| m5dn.12xlarge | m5dn.16xlarge | m5dn.24xlarge | m5n.large |
m5n.xlarge | m5n.2xlarge | m5n.4xlarge | m5n.8xlarge | m5n.12xlarge |
m5n.16xlarge | m5n.24xlarge | m6id.large | m6id.xlarge | m6id.2xlarge
| m6id.4xlarge | m6id.8xlarge | m6id.12xlarge | m6id.16xlarge
| m6id.24xlarge | m6id.32xlarge | m6idn.large | m6idn.xlarge |
m6idn.2xlarge | m6idn.4xlarge | m6idn.8xlarge | m6idn.12xlarge |
m6idn.16xlarge | m6idn.24xlarge | m6idn.32xlarge | m6in.large |
m6in.xlarge | m6in.2xlarge | m6in.4xlarge | m6in.8xlarge | m6in.12xlarge
| m6in.16xlarge | m6in.24xlarge | m6in.32xlarge | m8g.medium | m8g.large
| m8g.xlarge | m8g.2xlarge | m8g.4xlarge | m8g.8xlarge | m8g.12xlarge
| m8g.16xlarge | m8g.24xlarge | m8g.48xlarge | r5dn.large | r5dn.xlarge
| r5dn.2xlarge | r5dn.4xlarge | r5dn.8xlarge | r5dn.12xlarge |
r5dn.16xlarge | r5dn.24xlarge | r5n.large | r5n.xlarge | r5n.2xlarge |
r5n.4xlarge | r5n.8xlarge | r5n.12xlarge | r5n.16xlarge | r5n.24xlarge
| r6a.large | r6a.xlarge | r6a.2xlarge | r6a.4xlarge | r6a.8xlarge |
r6a.12xlarge | r6a.16xlarge | r6a.24xlarge | r6a.32xlarge | r6a.48xlarge
| r6id.large | r6id.xlarge | r6id.2xlarge | r6id.4xlarge | r6id.8xlarge
| r6id.12xlarge | r6id.16xlarge | r6id.24xlarge | r6id.32xlarge
| r6idn.large | r6idn.xlarge | r6idn.2xlarge | r6idn.4xlarge |
r6idn.8xlarge | r6idn.12xlarge | r6idn.16xlarge | r6idn.24xlarge |
r6idn.32xlarge | r6in.large | r6in.xlarge | r6in.2xlarge | r6in.4xlarge
| r6in.8xlarge | r6in.12xlarge | r6in.16xlarge | r6in.24xlarge |
r6in.32xlarge | r8g.medium | r8g.large | r8g.xlarge | r8g.2xlarge |
r8g.4xlarge | r8g.8xlarge | r8g.12xlarge | r8g.16xlarge | r8g.24xlarge |
r8g.48xlarge | m4.16xlarge | c6a.32xlarge | c6a.48xlarge | c6i.32xlarge
| r6i.24xlarge | r6i.32xlarge | c6in.24xlarge | c6in.32xlarge |
c7a.24xlarge | c7a.32xlarge | c7a.48xlarge | c7i.24xlarge | c7i.48xlarge
| m6a.24xlarge | m6a.32xlarge | m6a.48xlarge | m6i.24xlarge |

m6i.32xlarge | m7a.24xlarge | m7a.32xlarge | m7a.48xlarge | m7i.24xlarge
| m7i.48xlarge | r7a.24xlarge | r7a.32xlarge | r7a.48xlarge

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](#)

InstanceAccess

Information and credentials that you can use to remotely connect to an instance in an EC2 managed fleet. This data type is returned in response to a call to https://docs.aws.amazon.com/gamelift/latest/apireference/API_GetInstanceAccess.

Contents

Note

In the following list, the required parameters are described first.

Credentials

Security credentials that are required to access the instance.

Type: [InstanceCredentials](#) object

Required: No

FleetId

A unique identifier for the fleet containing the instance to be accessed.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-.]+`

Required: No

InstanceId

A unique identifier for the instance to be accessed.

Type: String

Pattern: `[a-zA-Z0-9\-.]+`

Required: No

IpAddress

IP address assigned to the instance.

Type: String

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

Pattern: `^[0-9A-Fa-f\:\\.]+`

Required: No

OperatingSystem

Operating system that is running on the instance.

Type: String

Valid Values: `WINDOWS_2012` | `AMAZON_LINUX` | `AMAZON_LINUX_2` | `WINDOWS_2016` | `AMAZON_LINUX_2023` | `WINDOWS_2022`

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](#)

InstanceCredentials

A set of credentials that allow remote access to an instance in an EC2 managed fleet. These credentials are returned in response to a call to https://docs.aws.amazon.com/gamelift/latest/apireference/API_GetInstanceAccess, which requests access for instances that are running game servers with the Amazon GameLift Servers server SDK version 4.x or earlier.

Contents

Note

In the following list, the required parameters are described first.

Secret

Secret string. For Windows instances, the secret is a password for use with Windows Remote Desktop. For Linux instances, it's a private key for use with SSH.

Type: String

Length Constraints: Minimum length of 1.

Required: No

UserName

A user name for logging in.

Type: String

Length Constraints: Minimum length 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](#)

InstanceDefinition

An allowed instance type for a game server group. All game server groups must have at least two instance types defined for it. Amazon GameLift Servers FleetIQ periodically evaluates each defined instance type for viability. It then updates the Auto Scaling group with the list of viable instance types.

Contents

Note

In the following list, the required parameters are described first.

InstanceType

An Amazon EC2 instance type designation.

Type: String

Valid Values: c4.large | c4.xlarge | c4.2xlarge | c4.4xlarge | c4.8xlarge | c5.large | c5.xlarge | c5.2xlarge | c5.4xlarge | c5.9xlarge | c5.12xlarge | c5.18xlarge | c5.24xlarge | c5a.large | c5a.xlarge | c5a.2xlarge | c5a.4xlarge | c5a.8xlarge | c5a.12xlarge | c5a.16xlarge | c5a.24xlarge | c6g.medium | c6g.large | c6g.xlarge | c6g.2xlarge | c6g.4xlarge | c6g.8xlarge | c6g.12xlarge | c6g.16xlarge | r4.large | r4.xlarge | r4.2xlarge | r4.4xlarge | r4.8xlarge | r4.16xlarge | r5.large | r5.xlarge | r5.2xlarge | r5.4xlarge | r5.8xlarge | r5.12xlarge | r5.16xlarge | r5.24xlarge | r5a.large | r5a.xlarge | r5a.2xlarge | r5a.4xlarge | r5a.8xlarge | r5a.12xlarge | r5a.16xlarge | r5a.24xlarge | r6g.medium | r6g.large | r6g.xlarge | r6g.2xlarge | r6g.4xlarge | r6g.8xlarge | r6g.12xlarge | r6g.16xlarge | m4.large | m4.xlarge | m4.2xlarge | m4.4xlarge | m4.10xlarge | m5.large | m5.xlarge | m5.2xlarge | m5.4xlarge | m5.8xlarge | m5.12xlarge | m5.16xlarge | m5.24xlarge | m5a.large | m5a.xlarge | m5a.2xlarge | m5a.4xlarge | m5a.8xlarge | m5a.12xlarge | m5a.16xlarge | m5a.24xlarge | m6g.medium | m6g.large | m6g.xlarge | m6g.2xlarge | m6g.4xlarge | m6g.8xlarge | m6g.12xlarge | m6g.16xlarge

Required: Yes

WeightedCapacity

Instance weighting that indicates how much this instance type contributes to the total capacity of a game server group. Instance weights are used by Amazon GameLift Servers FleetIQ to calculate the instance type's cost per unit hour and better identify the most cost-effective options. For detailed information on weighting instance capacity, see [Instance Weighting](#) in the *Amazon Elastic Compute Cloud Auto Scaling User Guide*. Default value is "1".

Type: String

Length Constraints: Minimum length of 1. Maximum length of 3.

Pattern: `^[\u0031-\u0039][\u0030-\u0039]{0,2}$`

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](#)

IpPermission

A range of IP addresses and port settings that allow inbound traffic to connect to processes on an instance in a fleet. Processes are assigned an IP address/port number combination, which must fall into the fleet's allowed ranges.

For Amazon GameLift Servers Realtime fleets, Amazon GameLift Servers automatically opens two port ranges, one for TCP messaging and one for UDP.

Contents

Note

In the following list, the required parameters are described first.

FromPort

A starting value for a range of allowed port numbers.

For fleets using Linux builds, only ports 22 and 1026-60000 are valid.

For fleets using Windows builds, only ports 1026-60000 are valid.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 60000.

Required: Yes

IpRange

A range of allowed IP addresses. This value must be expressed in CIDR notation. Example: "000.000.000.000/[subnet mask]" or optionally the shortened version "0.0.0.0/[subnet mask]".

Type: String

Pattern: `[^\s]+`

Required: Yes

Protocol

The network communication protocol used by the fleet.

Type: String

Valid Values: TCP | UDP

Required: Yes

ToPort

An ending value for a range of allowed port numbers. Port numbers are end-inclusive. This value must be equal to or greater than FromPort.

For fleets using Linux builds, only ports 22 and 1026-60000 are valid.

For fleets using Windows builds, only ports 1026-60000 are valid.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 60000.

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](#)

LaunchTemplateSpecification

An Amazon Elastic Compute Cloud launch template that contains configuration settings and game server code to be deployed to all instances in a game server group. The launch template is specified when creating a new game server group.

Contents

Note

In the following list, the required parameters are described first.

LaunchTemplateId

A unique identifier for an existing Amazon EC2 launch template.

Type: String

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

Pattern: `[\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDFFF\r\n\t]+`

Required: No

LaunchTemplateName

A readable identifier for an existing Amazon EC2 launch template.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `[a-zA-Z0-9\(\)\.\-/_]+`

Required: No

Version

The version of the Amazon EC2 launch template to use. If no version is specified, the default version will be used. With Amazon EC2, you can specify a default version for a launch template. If none is set, the default is the first version created.

Type: String

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

Pattern: `[\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDFFF\r\n\t]+`

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](#)

LinuxCapabilities

A set of Linux capabilities that are added to a container's default Docker configuration for a container defined in the [ContainerGroupDefinition](#). For more detailed information about these Linux capabilities, see the [capabilities\(7\)](#) Linux manual page.

Modifying capabilities on an existing container: To remove a capability, update the Include list with only the needed capabilities. To revert back to default capabilities, omit LinuxCapabilities within the ContainerDefinition.

Part of: [GameServerContainerDefinition](#), [GameServerContainerDefinitionInput](#), [SupportContainerDefinition](#), [SupportContainerDefinitionInput](#)

Returned by: [CreateContainerGroupDefinition](#), [DescribeContainerGroupDefinition](#), [ListContainerGroupDefinitions](#), [ListContainerGroupDefinitionVersions](#), [UpdateContainerGroupDefinition](#)

Contents

Note

In the following list, the required parameters are described first.

Include

The list of Linux capabilities to add to the container's default configuration. Specify each capability as a string from the set of supported capability names (for example, NET_BIND_SERVICE or SYS_PTRACE).

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 37 items.

Valid Values: AUDIT_CONTROL | AUDIT_WRITE | BLOCK_SUSPEND | CHOWN | DAC_OVERRIDE | DAC_READ_SEARCH | FOWNER | FSETID | IPC_LOCK | IPC_OWNER | KILL | LEASE | LINUX_IMMUTABLE | MAC_ADMIN | MAC_OVERRIDE | MKNOD | NET_ADMIN | NET_BIND_SERVICE | NET_BROADCAST | NET_RAW | SETFCAP | SETGID | SETPCAP | SETUID | SYS_ADMIN | SYS_BOOT | SYS_CHROOT

| SYS_MODULE | SYS_NICE | SYS_PACCT | SYS_PTRACE | SYS_RAWIO |
SYS_RESOURCE | SYS_TIME | SYS_TTY_CONFIG | SYSLOG | WAKE_ALARM

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](#)

LocationalDeployment

For a multi-location container fleet, describes the progress of a deployment across all fleet locations.

Contents

Note

In the following list, the required parameters are described first.

DeploymentStatus

The status of fleet deployment activity in the location.

- IN_PROGRESS -- The deployment is in progress.
- IMPAIRED -- The deployment failed and the fleet has some impaired containers.
- COMPLETE -- The deployment has completed successfully.
- ROLLBACK_IN_PROGRESS -- The deployment failed and rollback has been initiated.
- ROLLBACK_IN_COMPLETE -- The deployment failed and rollback has been completed.
- CANCELLED -- The deployment was cancelled.

Type: String

Valid Values: IN_PROGRESS | IMPAIRED | COMPLETE | ROLLBACK_IN_PROGRESS | ROLLBACK_COMPLETE | CANCELLED | PENDING

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](#)

LocationAttributes

Details about a location in a multi-location fleet.

Contents

Note

In the following list, the required parameters are described first.

LocationState

A fleet location and its current life-cycle state.

Type: [LocationState](#) object

Required: No

StoppedActions

A list of fleet actions that have been suspended in the fleet location.

Type: Array of strings

Array Members: Fixed number of 1 item.

Valid Values: AUTO_SCALING

Required: No

UpdateStatus

The status of fleet activity updates to the location. The status PENDING_UPDATE indicates that StopFleetActions or StartFleetActions has been requested but the update has not yet been completed for the location.

Type: String

Valid Values: PENDING_UPDATE

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](#)

LocationConfiguration

A remote location where a multi-location fleet can deploy game servers for game hosting.

Contents

Note

In the following list, the required parameters are described first.

Location

An AWS Region code, such as `us-west-2`. For a list of supported Regions and Local Zones, see [Amazon GameLift Servers service locations](#) for managed hosting.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-]+\`

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](#)

LocationModel

Properties of a location, which can include its name, ARN (for custom locations), and ping beacon information.

Contents

Note

In the following list, the required parameters are described first.

LocationArn

The Amazon Resource Name ([ARN](#)) that is assigned to a custom location resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::location/location-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`.

Type: String

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

Pattern: `^arn:.*:location\/custom-\S+`

Required: No

LocationName

The location's name.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-]+`

Required: No

PingBeacon

Information about the UDP ping beacon for this location. Ping beacons are fixed endpoints that you can use to measure network latency between a player device and an Amazon GameLift Servers hosting location.

Type: [PingBeacon](#) 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](#)

LocationState

A fleet location and its life-cycle state. A location state object might be used to describe a fleet's remote location or home Region. Life-cycle state tracks the progress of launching the first instance in a new location and preparing it for game hosting, and then removing all instances and deleting the location from the fleet.

- **NEW** -- A new fleet location has been defined and desired instances is set to 1.
- **DOWNLOADING/VALIDATING/BUILDING/ACTIVATING** -- Amazon GameLift Servers is setting up the new fleet location, creating new instances with the game build or Realtime script and starting server processes.
- **ACTIVE** -- Hosts can now accept game sessions.
- **ERROR** -- An error occurred when downloading, validating, building, or activating the fleet location.
- **DELETING** -- Hosts are responding to a delete fleet location request.
- **TERMINATED** -- The fleet location no longer exists.
- **NOT_FOUND** -- The fleet location was not found. This could be because the custom location was removed or not created.

Contents

Note

In the following list, the required parameters are described first.

Location

The fleet location, expressed as an AWS Region code such as `us-west-2`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-]+`

Required: No

PlayerGatewayStatus

The current status of player gateway in this location for this fleet. Note, even if a fleet has PlayerGatewayMode configured as ENABLED, player gateway might not be available in a specific location. For more information about locations where player gateway is supported, see [Amazon GameLift Servers service locations](#).

Possible values include:

- ENABLED -- Player gateway is available for this fleet location.
- DISABLED -- Player gateway is not available for this fleet location.

Type: String

Valid Values: DISABLED | ENABLED

Required: No

Status

The life-cycle status of a fleet location.

Type: String

Valid Values: NEW | DOWNLOADING | VALIDATING | BUILDING | ACTIVATING | ACTIVE | DELETING | ERROR | TERMINATED | NOT_FOUND

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](#)

LogConfiguration

A method for collecting container logs for the fleet. Amazon GameLift Servers saves all standard output for each container in logs, including game session logs. You can select from the following methods:

Contents

Note

In the following list, the required parameters are described first.

LogDestination

The type of log collection to use for a fleet.

- **CLOUDWATCH** -- (default value) Send logs to an Amazon CloudWatch log group that you define. Each container emits a log stream, which is organized in the log group.
- **S3** -- Store logs in an Amazon S3 bucket that you define. This bucket must reside in the fleet's home AWS Region.
- **NONE** -- Don't collect container logs.

Type: String

Valid Values: NONE | CLOUDWATCH | S3

Required: No

LogGroupArn

If log destination is CLOUDWATCH, logs are sent to the specified log group in Amazon CloudWatch.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: [a-zA-Z0-9:/\-*]+

Required: No

S3BucketName

If log destination is S3, logs are sent to the specified Amazon S3 bucket name.

Type: String

Length Constraints: Minimum length 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](#)

ManagedCapacityConfiguration

Use ManagedCapacityConfiguration with the "SCALE_TO_AND_FROM_ZERO" ZeroCapacityStrategy to enable Amazon GameLift Servers to fully manage the MinSize value, switching between 0 and 1 based on game session activity. This is ideal for eliminating compute costs during periods of no game activity. It is particularly beneficial during development when you're away from your desk, iterating on builds for extended periods, in production environments serving low-traffic locations, or for games with long, predictable downtime windows. By automatically managing capacity between 0 and 1 instances, you avoid paying for idle instances while maintaining the ability to serve game sessions when demand arrives. Note that while scale-out is triggered immediately upon receiving a game session request, actual game session availability depends on your server process startup time, so this approach works best with multi-location Fleets where cold-start latency is tolerable. With a "MANUAL" ZeroCapacityStrategy Amazon GameLift Servers will not modify Fleet MinSize values automatically and will not scale out from zero instances in response to game sessions.

Contents

Note

In the following list, the required parameters are described first.

ScaleInAfterInactivityMinutes

Length of time, in minutes, that Amazon GameLift Servers will wait before scaling in your MinSize and DesiredInstances to 0 after a period with no game session activity. Default: 30 minutes.

Type: Integer

Valid Range: Minimum value of 5. Maximum value of 1440.

Required: No

ZeroCapacityStrategy

The strategy Amazon GameLift Servers will use to automatically scale your capacity to and from zero instances in response to game session activity. Game session activity refers to any active running sessions or game session requests.

Possible ZeroCapacityStrategy types include:

- **MANUAL** -- (default value) Amazon GameLift Servers will not update capacity to and from zero on your behalf.
- **SCALE_TO_AND_FROM_ZERO** -- Amazon GameLift Servers will automatically scale out MinSize and DesiredInstances from 0 to 1 in response to a game session request, and will scale in MinSize and DesiredInstances to 0 after a period with no game session activity. The duration of this scale in period can be configured using ScaleInAfterInactivityMinutes.

Type: String

Valid Values: MANUAL | SCALE_TO_AND_FROM_ZERO

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](#)

MatchedPlayerSession

Represents a new player session that is created as a result of a successful FlexMatch match. A successful match automatically creates new player sessions for every player ID in the original matchmaking request.

When players connect to the match's game session, they must include both player ID and player session ID in order to claim their assigned player slot.

Contents

Note

In the following list, the required parameters are described first.

PlayerId

A unique identifier for a player

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

PlayerSessionId

A unique identifier for a player session. PlayerSessionId will only be populated for player sessions that are in ACTIVE or RESERVED status when the ticket is completed.

Type: String

Pattern: ^psess-\S+

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](#)

MatchmakingConfiguration

Guidelines for use with FlexMatch to match players into games. All matchmaking requests must specify a matchmaking configuration.

Contents

Note

In the following list, the required parameters are described first.

AcceptanceRequired

A flag that indicates whether a match that was created with this configuration must be accepted by the matched players. To require acceptance, set to TRUE. When this option is enabled, matchmaking tickets use the status `REQUIRES_ACCEPTANCE` to indicate when a completed potential match is waiting for player acceptance.

Type: Boolean

Required: No

AcceptanceTimeoutSeconds

The length of time (in seconds) to wait for players to accept a proposed match, if acceptance is required. If any player rejects the match or fails to accept before the timeout, the ticket continues to look for an acceptable match.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 600.

Required: No

AdditionalPlayerCount

The number of player slots in a match to keep open for future players. For example, if the configuration's rule set specifies a match for a single 12-person team, and the additional player count is set to 2, only 10 players are selected for the match. This parameter is not used when `FlexMatchMode` is set to `STANDALONE`.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

BackfillMode

The method used to backfill game sessions created with this matchmaking configuration. MANUAL indicates that the game makes backfill requests or does not use the match backfill feature. AUTOMATIC indicates that GameLift creates backfill requests whenever a game session has one or more open slots. Learn more about manual and automatic backfill in [Backfill existing games with FlexMatch](#). Automatic backfill is not available when FlexMatchMode is set to STANDALONE.

Type: String

Valid Values: AUTOMATIC | MANUAL

Required: No

ConfigurationArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers matchmaking configuration resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::matchmakingconfiguration/<matchmaking configuration name>`. In a Amazon GameLift Servers configuration ARN, the resource ID matches the *Name* value.

Type: String

Pattern: `^arn:.*:matchmakingconfiguration\[a-zA-Z0-9-\.\]*`

Required: No

CreationTime

A time stamp indicating when this data object was created. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

CustomEventData

Information to attach to all events related to the matchmaking configuration.

Type: String

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

Required: No

Description

A descriptive label that is associated with matchmaking configuration.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

FlexMatchMode

Indicates whether this matchmaking configuration is being used with Amazon GameLift Servers hosting or as a standalone matchmaking solution.

- **STANDALONE** - FlexMatch forms matches and returns match information, including players and team assignments, in a [MatchmakingSucceeded](#) event.
- **WITH_QUEUE** - FlexMatch forms matches and uses the specified Amazon GameLift Servers queue to start a game session for the match.

Type: String

Valid Values: STANDALONE | WITH_QUEUE

Required: No

GameProperties

A set of key-value pairs that can store custom data in a game session. For example: {"Key": "difficulty", "Value": "novice"}. This information is added to the new GameSession object that is created for a successful match. This parameter is not used when FlexMatchMode is set to STANDALONE.

Note

- Avoid using periods (".") in property keys if you plan to search for game sessions by properties. Property keys containing periods cannot be searched and will be filtered out from search results due to search index limitations.

- If you use SearchGameSessions API, there is a limit of 500 game property keys across all game sessions and all fleets per region. If the limit is exceeded, there will potentially be game session entries missing from SearchGameSessions API results.

Type: Array of [GameProperty](#) objects

Array Members: Maximum number of 16 items.

Required: No

GameSessionData

A set of custom game session properties, formatted as a single string value. This data is passed to a game server process with a request to start a new game session. For more information, see [Start a game session](#). This information is added to the new GameSession object that is created for a successful match. This parameter is not used when FlexMatchMode is set to STANDALONE.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 4096.

Required: No

GameSessionQueueArns

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers game session queue resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::gamesessionqueue/<queue name>`. Queues can be located in any Region. Queues are used to start new Amazon GameLift Servers-hosted game sessions for matches that are created with this matchmaking configuration. This property is not set when FlexMatchMode is set to STANDALONE.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `[a-zA-Z0-9:/- ]+`

Required: No

Name

A unique identifier for the matchmaking configuration. This name is used to identify the configuration associated with a matchmaking request or ticket.

Type: String

Length Constraints: Maximum length of 128.

Pattern: `[a-zA-Z0-9-\.]*`

Required: No

NotificationTarget

An SNS topic ARN that is set up to receive matchmaking notifications.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 300.

Pattern: `[a-zA-Z0-9:_/-]*(.fifo)?`

Required: No

RequestTimeoutSeconds

The maximum duration, in seconds, that a matchmaking ticket can remain in process before timing out. Requests that fail due to timing out can be resubmitted as needed.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 43200.

Required: No

RuleSetArn

The Amazon Resource Name ([ARN](#)) associated with the GameLift matchmaking rule set resource that this configuration uses.

Type: String

Pattern: `^arn:.*:matchmakingruleset\[a-zA-Z0-9-\.]*`

Required: No

RuleSetName

A unique identifier for the matchmaking rule set to use with this configuration. A matchmaking configuration can only use rule sets that are defined in the same Region.

Type: String

Length Constraints: Maximum length of 128.

Pattern: [a-zA-Z0-9-\.\]*

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](#)

MatchmakingRuleSet

Set of rule statements, used with FlexMatch, that determine how to build your player matches. Each rule set describes a type of group to be created and defines the parameters for acceptable player matches.

A rule set may define the following elements for a match. For detailed information and examples showing how to construct a rule set, see [Build a FlexMatch rule set](#).

- **Teams -- Required.** A rule set must define one or multiple teams for the match and set minimum and maximum team sizes. For example, a rule set might describe a 4x4 match that requires all eight slots to be filled.
- **Player attributes -- Optional.** These attributes specify a set of player characteristics to evaluate when looking for a match. Matchmaking requests that use a rule set with player attributes must provide the corresponding attribute values. For example, an attribute might specify a player's skill or level.
- **Rules -- Optional.** Rules define how to evaluate potential players for a match based on player attributes. A rule might specify minimum requirements for individual players, teams, or entire matches. For example, a rule might require each player to meet a certain skill level, each team to have at least one player in a certain role, or the match to have a minimum average skill level. or may describe an entire group--such as all teams must be evenly matched or have at least one player in a certain role.
- **Expansions -- Optional.** Expansions allow you to relax the rules after a period of time when no acceptable matches are found. This feature lets you balance getting players into games in a reasonable amount of time instead of making them wait indefinitely for the best possible match. For example, you might use an expansion to increase the maximum skill variance between players after 30 seconds.

Contents

Note

In the following list, the required parameters are described first.

RuleSetBody

A collection of matchmaking rules, formatted as a JSON string. Comments are not allowed in JSON, but most elements support a description field.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 65535.

Required: Yes

CreationTime

A time stamp indicating when this data object was created. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

RuleSetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers matchmaking rule set resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::matchmakingruleset/<ruleset name>`. In a GameLift rule set ARN, the resource ID matches the *RuleSetName* value.

Type: String

Pattern: `^arn:.*:matchmakingruleset\/[a-zA-Z0-9-\.]*`

Required: No

RuleSetName

A unique identifier for the matchmaking rule set

Type: String

Length Constraints: Maximum length of 128.

Pattern: `[a-zA-Z0-9-\.]*`

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](#)

MatchmakingTicket

Ticket generated to track the progress of a matchmaking request. Each ticket is uniquely identified by a ticket ID, supplied by the requester, when creating a matchmaking request.

Contents

Note

In the following list, the required parameters are described first.

ConfigurationArn

The Amazon Resource Name ([ARN](#)) associated with the GameLift matchmaking configuration resource that is used with this ticket.

Type: String

Pattern: `^arn:.*:matchmakingconfiguration\[a-zA-Z0-9-\.\]*`

Required: No

ConfigurationName

Name of the matchmaking configuration that is used with this ticket. Matchmaking configurations determine how players are grouped into a match and how a new game session is created for the match.

Type: String

Length Constraints: Maximum length of 128.

Pattern: `\[a-zA-Z0-9-\.\]*`

Required: No

EndTime

Time stamp indicating when the matchmaking request stopped being processed due to successful completion, timeout, or cancellation. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

EstimatedWaitTime

Average amount of time (in seconds) that players are currently waiting for a match. If there is not enough recent data, this property may be empty.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

GameSessionConnectionInfo

Connection information for a new game session. Once a match is made, the FlexMatch engine creates a new game session for it. This information is added to the matchmaking ticket, which you can retrieve by calling [DescribeMatchmaking](#) .

Type: [GameSessionConnectionInfo](#) object

Required: No

Players

A set of `Player` objects, each representing a player to find matches for. Players are identified by a unique player ID and may include latency data for use during matchmaking. If the ticket is in status `COMPLETED`, the `Player` objects include the team the players were assigned to in the resulting match.

Type: Array of [Player](#) objects

Required: No

StartTime

Time stamp indicating when this matchmaking request was received. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

Status

Current status of the matchmaking request.

- **QUEUED** -- The matchmaking request has been received and is currently waiting to be processed.
- **SEARCHING** -- The matchmaking request is currently being processed.
- **REQUIRES_ACCEPTANCE** -- A match has been proposed and the players must accept the match. This status is used only with requests that use a matchmaking configuration with a player acceptance requirement.
- **PLACING** -- The FlexMatch engine has matched players and is in the process of placing a new game session for the match.
- **COMPLETED** -- Players have been matched and a game session is ready to host the players. A ticket in this state contains the necessary connection information for players.
- **FAILED** -- The matchmaking request was not completed.
- **CANCELLED** -- The matchmaking request was canceled. This may be the result of a `StopMatchmaking` operation or a proposed match that one or more players failed to accept.
- **TIMED_OUT** -- The matchmaking request was not successful within the duration specified in the matchmaking configuration.

Note

Matchmaking requests that fail to successfully complete (statuses `FAILED`, `CANCELLED`, `TIMED_OUT`) can be resubmitted as new requests with new ticket IDs.

Type: String

Valid Values: `CANCELLED` | `COMPLETED` | `FAILED` | `PLACING` | `QUEUED` | `REQUIRES_ACCEPTANCE` | `SEARCHING` | `TIMED_OUT`

Required: No

StatusMessage

Additional information about the current status.

Type: String

Required: No

StatusReason

Code to explain the current status. For example, a status reason may indicate when a ticket has returned to SEARCHING status after a proposed match fails to receive player acceptances.

Type: String

Required: No

TicketId

A unique identifier for a matchmaking ticket.

Type: String

Length Constraints: Maximum length of 128.

Pattern: [a-zA-Z0-9-\.]*

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](#)

PingBeacon

Information about a UDP ping beacon that can be used to measure network latency between a player device and an Amazon GameLift Servers hosting location.

Ping beacons provide more accurate latency measurements than ICMP pings because they use UDP, which is the same protocol used by most game servers. For more information on UDP ping beacons, including guidance on how to implement latency measurements and best practices for caching beacon endpoints, refer to [UDP ping beacons](#) in the Amazon GameLift Servers Developer Guide.

Contents

Note

In the following list, the required parameters are described first.

UDPEndpoint

The domain name and port of the UDP ping beacon. Your game client can send UDP messages to this endpoint and receive responses to measure network latency.

Type: [UDPEndpoint](#) 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](#)

PlacedPlayerSession

Information about a player session. This object contains only the player ID and player session ID. To retrieve full details on a player session, call [DescribePlayerSessions](#) with the player session ID.

Contents

Note

In the following list, the required parameters are described first.

PlayerId

A unique identifier for a player that is associated with this player session.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

PlayerSessionId

A unique identifier for a player session.

Type: String

Pattern: ^psess-\S+

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](#)

Player

Represents a player in matchmaking. When starting a matchmaking request, a player has a player ID, attributes, and may have latency data. Team information is added after a match has been successfully completed.

Contents

Note

In the following list, the required parameters are described first.

LatencyInMs

A set of values, expressed in milliseconds, that indicates the amount of latency that a player experiences when connected to a fleet location (AWS Regions or custom locations for Amazon GameLift Servers Anywhere fleets). If this property is present, FlexMatch considers placing the match only in Regions for which latency is reported.

If a matchmaker has a rule that evaluates player latency, players must report latency in order to be matched. If no latency is reported in this scenario, FlexMatch assumes that no Regions are available to the player and the ticket is not matchable.

To collect accurate latency data, use Amazon GameLift Servers's UDP ping beacons, which provide fixed endpoints in each Amazon GameLift Servers hosting location. These endpoints allow you to measure actual UDP network latency, which provides more accurate results than ICMP pings because UDP is the same protocol used by most game servers. For more information on using UDP ping beacons to measure latency, refer to [UDP ping beacons](#) in the Amazon GameLift Servers Developer Guide.

Type: String to integer map

Key Length Constraints: Minimum length of 1.

Valid Range: Minimum value of 1.

Required: No

PlayerAttributes

A collection of key:value pairs containing player information for use in matchmaking. Player attribute keys must match the *playerAttributes* used in a matchmaking rule set.

Example: "PlayerAttributes": {"skill": {"N": "23"}, "gameMode": {"S": "deathmatch"}}.

You can provide up to 10 PlayerAttributes.

Type: String to [AttributeValue](#) object map

Key Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

PlayerId

A unique identifier for a player

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Team

Name of the team that the player is assigned to in a match. Team names are defined in a matchmaking rule set.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

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](#)

PlayerConnectionDetail

Connection information for a game client to connect to a game session. This object contains the IP address(es), port(s), and authentication details your game client needs to establish a connection.

With player gateway enabled: Contains relay endpoints and a player gateway token. Your game client must prepend player gateway token to each payload for validation and connection through relay endpoints.

With player gateway disabled: Contains game server endpoint. Player gateway token and expiration fields are empty.

Contents

Note

In the following list, the required parameters are described first.

Endpoints

List of connection endpoints for the game client. Your game client uses these IP address(es) and port(s) to connect to the game session.

When player gateway is enabled, these are relay endpoints with benefits such as DDoS protection. When disabled, this is the game server endpoint.

Type: Array of [PlayerConnectionEndpoint](#) objects

Required: No

Expiration

When player gateway is enabled, this is the timestamp indicating when player gateway token expires. Your game backend should call [GetPlayerConnectionDetails](#) to retrieve fresh connection information for your game clients before this time. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

This value is empty when player gateway is disabled.

Type: Timestamp

Required: No

PlayerGatewayToken

Access token that your game client must prepend to all traffic sent through player gateway. Player gateway verifies identity and authorizes connection based on this token.

This value is empty when player gateway is disabled.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1024.

Required: No

PlayerId

A unique identifier for a player associated with this connection.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

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](#)

PlayerConnectionEndpoint

Network address(es) and port(s) for connecting to a game session.

Contents

Note

In the following list, the required parameters are described first.

IpAddress

IP address for connecting to the game session. When player gateway is enabled, this is a player gateway IP address. When player gateway is disabled, this is the game server IP address.

Type: String

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

Pattern: `^[0-9A-Fa-f\:\.]+`

Required: No

Port

Port number for connecting to the game session. When player gateway is enabled, this is a player gateway port. When player gateway is disabled, this is the game server port.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 60000.

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](#)

PlayerGatewayConfiguration

Configuration settings for player gateway. Use these settings to specify advanced options for how player gateway handles connections.

Contents

Note

In the following list, the required parameters are described first.

GameServerIpProtocolSupported

The IP protocol that your game servers support for player connections through player gateway. If the value is set to IPv4, GameLift will install and execute a lightweight IP translation software on fleet instances to receive and transform incoming IPv6 traffic to IPv4. If the value is set to DUAL_STACK, the lightweight IP translation software will not be installed on fleet instances. DUAL_STACK provides slightly better performance than IPv4.

Type: String

Valid Values: IPv4 | DUAL_STACK

Required: No

See Also

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

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

PlayerLatency

Regional latency information for a player, used when requesting a new game session. This value indicates the amount of time lag that exists when the player is connected to a fleet in the specified location (an AWS Region or a custom location for Amazon GameLift Servers Anywhere fleets). The relative difference between a player's latency values for multiple locations are used to determine which fleets are best suited to place a new game session for the player.

Contents

Note

In the following list, the required parameters are described first.

LatencyInMilliseconds

Amount of time that represents the time lag experienced by the player when connected to the specified Region.

Type: Float

Required: No

PlayerId

A unique identifier for a player associated with the latency data.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

RegionIdentifier

Name of the Region or custom location that is associated with the latency value. For Amazon GameLift Servers Anywhere fleets, use the custom location name.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

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](#)

PlayerLatencyPolicy

Sets a latency cap for individual players when placing a game session. With a latency policy in force, a game session cannot be placed in a fleet location where a player reports latency higher than the cap. Latency policies are used only with placement request that provide player latency information. Player latency policies can be stacked to gradually relax latency requirements over time.

Contents

Note

In the following list, the required parameters are described first.

MaximumIndividualPlayerLatencyMilliseconds

The maximum latency value that is allowed for any player, in milliseconds. All policies must have a value set for this property.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

PolicyDurationSeconds

The length of time, in seconds, that the policy is enforced while placing a new game session. A null value for this property means that the policy is enforced until the queue times out.

Type: Integer

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](#)

PlayerSession

Represents a player session. Player sessions are created either for a specific game session, or as part of a game session placement or matchmaking request. A player session can represent a reserved player slot in a game session (when status is `RESERVED`) or actual player activity in a game session (when status is `ACTIVE`). A player session object, including player data, is automatically passed to a game session when the player connects to the game session and is validated. After the game session ends, player sessions information is retained for 30 days and then removed.

Related actions

[All APIs by task](#)

Contents

Note

In the following list, the required parameters are described first.

CreationTime

A time stamp indicating when this data object was created. Format is a number expressed in Unix time as milliseconds (for example `"1469498468.057"`).

Type: Timestamp

Required: No

DnsName

The DNS identifier assigned to the instance that is running the game session. Values have the following format:

- TLS-enabled fleets: `<unique identifier>.<region identifier>.amazongamelift.com`.
- Non-TLS-enabled fleets: `ec2-<unique identifier>.compute.amazonaws.com`. (See [Amazon EC2 Instance IP Addressing](#).)

When connecting to a game session that is running on a TLS-enabled fleet, you must use the DNS name, not the IP address.

Type: String

Required: No

FleetArn

The Amazon Resource Name ([ARN](#)) associated with the GameLift fleet that the player's game session is running on.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet\[/a-z]*fleet-[a-zA-Z0-9\-\-]+$`

Required: No

FleetId

A unique identifier for the fleet that the player's game session is running on.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\-]+$`

Required: No

GameSessionId

An identifier for the game session that is unique across all regions that the player session is connected to. The value is always a full ARN in the following format: For Home Region game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<ID string>`. For Remote Location game session - `arn:aws:gamelift:<home_region>::gamesession/<fleet ID>/<location>/<ID string>`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

IpAddress

The IP address of the game session. To connect to a Amazon GameLift Servers game server, an app needs both the IP address and port number.

Type: String

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

Pattern: `^[0-9A-Fa-f\:\\.]+`

Required: No

PlayerData

Developer-defined information related to a player. Amazon GameLift Servers does not use this data, so it can be formatted as needed for use in the game.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

PlayerId

A unique identifier for a player that is associated with this player session.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

PlayerSessionId

A unique identifier for a player session.

Type: String

Pattern: `^psess-\S+`

Required: No

Port

Port number for the game session. To connect to a Amazon GameLift Servers server process, an app needs both the IP address and port number.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 60000.

Required: No

Status

Current status of the player session.

Possible player session statuses include the following:

- **RESERVED** -- The player session request has been received, but the player has not yet connected to the server process and/or been validated.
- **ACTIVE** -- The player has been validated by the server process and is currently connected.
- **COMPLETED** -- The player connection has been dropped.
- **TIMEDOUT** -- A player session request was received, but the player did not connect and/or was not validated within the timeout limit (60 seconds).

Type: String

Valid Values: RESERVED | ACTIVE | COMPLETED | TIMEDOUT

Required: No

TerminationTime

A time stamp indicating when this data object was terminated. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

See Also

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

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

PriorityConfiguration

Custom prioritization settings to use with a game session queue. Prioritization settings determine how the queue selects a game hosting resource to start a new game session. This configuration replaces the default prioritization process for queues.

By default, a queue makes game session placements based on the following criteria:

- When a game session request does not include player latency data, Amazon GameLift Servers places game sessions based on the following priorities: (1) the queue's default destination order, and (2) for multi-location fleets, an alphabetic list of locations.
- When a game session request includes player latency data, Amazon GameLift Servers re-orders the queue's destinations to make placements where the average player latency is lowest. It reorders based the following priorities: (1) the lowest average latency across all players, (2) the lowest hosting cost, (3) the queue's default destination order, and (4) for multi-location fleets, an alphabetic list of locations.

Contents

Note

In the following list, the required parameters are described first.

LocationOrder

The prioritization order to use for fleet locations, when the `PriorityOrder` property includes `LOCATION`. Locations can include AWS Region codes (such as `us-west-2`), local zones, and custom locations (for Anywhere fleets). Each location must be listed only once. For details, see [Amazon GameLift Servers service locations](#).

Type: Array of strings

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

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-\_]+`

Required: No

PriorityOrder

A custom sequence to use when prioritizing where to place new game sessions. Each priority type is listed once.

- **LATENCY** -- Amazon GameLift Servers prioritizes locations where the average player latency is lowest. Player latency data is provided in each game session placement request.
- **COST** -- Amazon GameLift Servers prioritizes queue destinations with the lowest current hosting costs. Cost is evaluated based on the destination's location, instance type, and fleet type (Spot or On-Demand).
- **DESTINATION** -- Amazon GameLift Servers prioritizes based on the list order of destinations in the queue configuration.
- **LOCATION** -- Amazon GameLift Servers prioritizes based on the provided order of locations, as defined in `LocationOrder`.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 4 items.

Valid Values: LATENCY | COST | DESTINATION | LOCATION

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](#)

PriorityConfigurationOverride

An alternate list of prioritized locations for use with a game session queue. When this property is included in a [StartGameSessionPlacement](#) request, the alternate list overrides the queue's default location priorities, as defined in the queue's [PriorityConfiguration](#) setting (*LocationOrder*). The override is valid for an individual placement request only. Use this property only with queues that have a `PriorityConfiguration` setting that prioritizes `LOCATION` first.

Note

A priority configuration override list does not override a queue's `FilterConfiguration` setting, if the queue has one. Filter configurations are used to limit placements to a subset of the locations in a queue's destinations. If the override list includes a location that's not on in the `FilterConfiguration` allowed list, Amazon GameLift Servers won't attempt to place a game session there.

Contents

Note

In the following list, the required parameters are described first.

LocationOrder

A prioritized list of hosting locations. The list can include AWS Regions (such as `us-west-2`), local zones, and custom locations (for Anywhere fleets). Each location must be listed only once. For details, see [Amazon GameLift Servers service locations](#).

Type: Array of strings

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

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\-\-]+`

Required: Yes

PlacementFallbackStrategy

Instructions for how to proceed if placement fails in every location on the priority override list. Valid strategies include:

- `DEFAULT_AFTER_SINGLE_PASS` -- After attempting to place a new game session in every location on the priority override list, try to place a game session in queue's other locations. This is the default behavior.
- `NONE` -- Limit placements to locations on the priority override list only.

Type: String

Valid Values: `DEFAULT_AFTER_SINGLE_PASS` | `NONE`

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](#)

ResourceCreationLimitPolicy

A policy that puts limits on the number of game sessions that a player can create within a specified span of time. With this policy, you can control players' ability to consume available resources.

The policy is evaluated when a player tries to create a new game session. On receiving a `CreateGameSession` request, Amazon GameLift Servers checks that the player (identified by `CreatorId`) has created fewer than game session limit in the specified time period.

The purpose of this policy is to prevent a single player from consuming a large share of available hosting resources. For example, setting `NewGameSessionsPerCreator` to 4 and `PolicyPeriodInMinutes` to 10 limits each player to creating 4 game sessions every 10 minutes. Setting these values too high (for example, 200 game sessions every 1000 minutes) still allows a single player to rapidly consume resources. We recommend keeping these values small.

Contents

Note

In the following list, the required parameters are described first.

NewGameSessionsPerCreator

A policy that puts limits on the number of game sessions that a player can create within a specified span of time. With this policy, you can control players' ability to consume available resources.

The policy is evaluated when a player tries to create a new game session. On receiving a `CreateGameSession` request, Amazon GameLift Servers checks that the player (identified by `CreatorId`) has created fewer than game session limit in the specified time period.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

PolicyPeriodInMinutes

The time span used in evaluating the resource creation limit policy.

Type: Integer

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](#)

RoutingStrategy

The routing configuration for a fleet alias.

Related actions

[All APIs by task](#)

Contents

Note

In the following list, the required parameters are described first.

FleetId

A unique identifier for the fleet that the alias points to. This value is the fleet ID, not the fleet ARN.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\ -]+`

Required: No

Message

The message text to be used with a terminal routing strategy.

Type: String

Required: No

Type

The type of routing strategy for the alias.

Possible routing types include the following:

- **SIMPLE** - The alias resolves to one specific fleet. Use this type when routing to active fleets.

- **TERMINAL** - The alias does not resolve to a fleet but instead can be used to display a message to the user. A terminal alias throws a `TerminalRoutingStrategyException` with the message embedded.

Type: String

Valid Values: SIMPLE | TERMINAL

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](#)

RuntimeConfiguration

A set of instructions that define the set of server processes to run on computes in a fleet. Server processes run either an executable in a custom game build or a Amazon GameLift Servers Realtime script. Amazon GameLift Servers launches the processes, manages their life cycle, and replaces them as needed. Computes check regularly for an updated runtime configuration.

An Amazon GameLift Servers instance is limited to 50 processes running concurrently. To calculate the total number of processes defined in a runtime configuration, add the values of the `ConcurrentExecutions` parameter for each server process. Learn more about [Running Multiple Processes on a Fleet](#).

Contents

Note

In the following list, the required parameters are described first.

GameSessionActivationTimeoutSeconds

The maximum amount of time (in seconds) allowed to launch a new game session and have it report ready to host players. During this time, the game session is in status `ACTIVATING`. If the game session does not become active before the timeout, it is ended and the game session status is changed to `TERMINATED`.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 600.

Required: No

MaxConcurrentGameSessionActivations

The number of game sessions in status `ACTIVATING` to allow on an instance or compute. This setting limits the instance resources that can be used for new game activations at any one time.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 2147483647.

Required: No

ServerProcesses

A collection of server process configurations that identify what server processes to run on fleet computes.

Type: Array of [ServerProcess](#) objects

Array Members: Minimum number of 1 item. Maximum number of 50 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](#)

S3Location

The location in Amazon S3 where build or script files are stored for access by Amazon GameLift Servers.

Contents

Note

In the following list, the required parameters are described first.

Bucket

An Amazon S3 bucket identifier. The name of the S3 bucket.

Note

Amazon GameLift Servers doesn't support uploading from Amazon S3 buckets with names that contain a dot (.).

Type: String

Length Constraints: Minimum length of 1.

Required: No

Key

The name of the zip file that contains the build files or script files.

Type: String

Length Constraints: Minimum length of 1.

Required: No

ObjectVersion

The version of the file, if object versioning is turned on for the bucket. Amazon GameLift Servers uses this information when retrieving files from an S3 bucket that you own. Use this

parameter to specify a specific version of the file. If not set, the latest version of the file is retrieved.

Type: String

Length Constraints: Minimum length of 1.

Required: No

RoleArn

The Amazon Resource Name ([ARN](#)) for an IAM role that allows Amazon GameLift Servers to access the S3 bucket.

Type: String

Length Constraints: Minimum length 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](#)

ScalingPolicy

Rule that controls how a fleet is scaled. Scaling policies are uniquely identified by the combination of name and fleet ID.

Contents

Note

In the following list, the required parameters are described first.

ComparisonOperator

Comparison operator to use when measuring a metric against the threshold value.

Type: String

Valid Values: `GreaterThanOrEqualToThreshold` | `GreaterThanThreshold` | `LessThanThreshold` | `LessThanOrEqualToThreshold`

Required: No

EvaluationPeriods

Length of time (in minutes) the metric must be at or beyond the threshold before a scaling event is triggered.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

FleetArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers fleet resource and uniquely identifies it. ARNs are unique across all Regions. Format is `arn:aws:gamelift:<region>::fleet/fleet-a1234567-b8c9-0d1e-2fa3-b45c6d7e8912`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet\/[a-z]*fleet-[a-zA-Z0-9\ -]+$`

Required: No

FleetId

A unique identifier for the fleet that is associated with this scaling policy.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\ -]+$`

Required: No

Location

The fleet location.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `^[A-Za-z0-9\ -]+$`

Required: No

MetricName

Name of the Amazon GameLift Servers-defined metric that is used to trigger a scaling adjustment. For detailed descriptions of fleet metrics, see [Monitor Amazon GameLift Servers with Amazon CloudWatch](#).

- **ActivatingGameSessions** -- Game sessions in the process of being created.
- **ActiveGameSessions** -- Game sessions that are currently running.
- **ActiveInstances** -- Fleet instances that are currently running at least one game session.
- **AvailableGameSessions** -- Additional game sessions that fleet could host simultaneously, given current capacity.
- **AvailablePlayerSessions** -- Empty player slots in currently active game sessions. This includes game sessions that are not currently accepting players. Reserved player slots are not included.

- **CurrentPlayerSessions** -- Player slots in active game sessions that are being used by a player or are reserved for a player.
- **IdleInstances** -- Active instances that are currently hosting zero game sessions.
- **PercentAvailableGameSessions** -- Unused percentage of the total number of game sessions that a fleet could host simultaneously, given current capacity. Use this metric for a target-based scaling policy.
- **PercentIdleInstances** -- Percentage of the total number of active instances that are hosting zero game sessions.
- **QueueDepth** -- Pending game session placement requests, in any queue, where the current fleet is the top-priority destination.
- **WaitTime** -- Current wait time for pending game session placement requests, in any queue, where the current fleet is the top-priority destination.

Type: String

Valid Values: `ActivatingGameSessions` | `ActiveGameSessions` | `ActiveInstances` | `AvailableGameSessions` | `AvailablePlayerSessions` | `CurrentPlayerSessions` | `IdleInstances` | `PercentAvailableGameSessions` | `PercentIdleInstances` | `QueueDepth` | `WaitTime` | `ConcurrentActivatableGameSessions`

Required: No

Name

A descriptive label that is associated with a fleet's scaling policy. Policy names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

PolicyType

The type of scaling policy to create. For a target-based policy, set the parameter *MetricName* to 'PercentAvailableGameSessions' and specify a *TargetConfiguration*. For a rule-based policy set the following parameters: *MetricName*, *ComparisonOperator*, *Threshold*, *EvaluationPeriods*, *ScalingAdjustmentType*, and *ScalingAdjustment*.

Type: String

Valid Values: RuleBased | TargetBased

Required: No

ScalingAdjustment

Amount of adjustment to make, based on the scaling adjustment type.

Type: Integer

Required: No

ScalingAdjustmentType

The type of adjustment to make to a fleet's instance count.

- **ChangeInCapacity** -- add (or subtract) the scaling adjustment value from the current instance count. Positive values scale up while negative values scale down.
- **ExactCapacity** -- set the instance count to the scaling adjustment value.
- **PercentChangeInCapacity** -- increase or reduce the current instance count by the scaling adjustment, read as a percentage. Positive values scale up while negative values scale down.

Type: String

Valid Values: ChangeInCapacity | ExactCapacity | PercentChangeInCapacity

Required: No

Status

Current status of the scaling policy. The scaling policy can be in force only when in an ACTIVE status. Scaling policies can be suspended for individual fleets. If the policy is suspended for a fleet, the policy status does not change.

- **ACTIVE** -- The scaling policy can be used for auto-scaling a fleet.
- **UPDATE_REQUESTED** -- A request to update the scaling policy has been received.
- **UPDATING** -- A change is being made to the scaling policy.
- **DELETE_REQUESTED** -- A request to delete the scaling policy has been received.
- **DELETING** -- The scaling policy is being deleted.
- **DELETED** -- The scaling policy has been deleted.
- **ERROR** -- An error occurred in creating the policy. It should be removed and recreated.

Type: String

Valid Values: ACTIVE | UPDATE_REQUESTED | UPDATING | DELETE_REQUESTED | DELETING | DELETED | ERROR

Required: No

TargetConfiguration

An object that contains settings for a target-based scaling policy.

Type: [TargetConfiguration](#) object

Required: No

Threshold

Metric value used to trigger a scaling event.

Type: Double

Required: No

UpdateStatus

The current status of the fleet's scaling policies in a requested fleet location. The status PENDING_UPDATE indicates that an update was requested for the fleet but has not yet been completed for the location.

Type: String

Valid Values: PENDING_UPDATE

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](#)

Script

Properties describing a Realtime script.

Related actions

[All APIs by task](#)

Contents

Note

In the following list, the required parameters are described first.

CreationTime

A time stamp indicating when this data object was created. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

Name

A descriptive label that is associated with a script. Script names do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

NodeJsVersion

The Node.js version used for execution of your Realtime script. The valid values are 10.x | 24.x. By default, NodeJsVersion is 10.x. This value cannot be updated later.

Note

Node.js 10 will reach end of support on September 30, 2026. See more details in the [Node.js 10 FAQs](#). For migration guidance, see [Migrating from Node.js 10 to 24](#).

Type: String

Length Constraints: Maximum length of 16.

Pattern: `^[0-9]+\.(?:x|[0-9]+)$`

Required: No

ScriptArn

The Amazon Resource Name ([ARN](#)) that is assigned to a Amazon GameLift Servers script resource and uniquely identifies it. ARNs are unique across all Regions. In a GameLift script ARN, the resource ID matches the *ScriptId* value.

Type: String

Pattern: `^arn:.*:script/script-\S+`

Required: No

ScriptId

A unique identifier for the Realtime script

Type: String

Pattern: `^script-\S+`

Required: No

SizeOnDisk

The file size of the uploaded Realtime script, expressed in bytes. When files are uploaded from an S3 location, this value remains at "0".

Type: Long

Valid Range: Minimum value of 0.

Required: No

StorageLocation

The location of the Amazon S3 bucket where a zipped file containing your Realtime scripts is stored. The storage location must specify the Amazon S3 bucket name, the zip file name (the

"key"), and a role ARN that allows Amazon GameLift Servers to access the Amazon S3 storage location. The S3 bucket must be in the same Region where you want to create a new script. By default, Amazon GameLift Servers uploads the latest version of the zip file; if you have S3 object versioning turned on, you can use the `ObjectVersion` parameter to specify an earlier version.

Type: [S3Location](#) object

Required: No

Version

Version information that is associated with a build or script. Version strings do not need to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

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](#)

ServerProcess

A set of instructions for launching server processes on fleet computes. Server processes run either an executable in a custom game build or a Amazon GameLift Servers Realtime script. Server process configurations are part of a fleet's runtime configuration.

Contents

Note

In the following list, the required parameters are described first.

ConcurrentExecutions

The number of server processes using this configuration that run concurrently on each instance or compute.

Type: Integer

Valid Range: Minimum value of 1.

Required: Yes

LaunchPath

The location of a game build executable or Realtime script. Game builds and Realtime scripts are installed on instances at the root:

- Windows (custom game builds only): C:\game. Example: "C:\game\MyGame\server.exe"
- Linux: /local/game. Examples: "/local/game/MyGame/server.exe" or "/local/game/MyRealtimeScript.js"

Note

Amazon GameLift Servers doesn't support the use of setup scripts that launch the game executable. For custom game builds, this parameter must indicate the executable that calls the server SDK operations `initSDK()` and `ProcessReady()`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `[A-Za-z0-9_:.+\\/\- ]+`

Required: Yes

Parameters

An optional list of parameters to pass to the server executable or Realtime script on launch.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `[A-Za-z0-9_:.+\\/\- =@{ },?'\" ]+`

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](#)

SupportContainerDefinition

Describes a support container in a container group. A support container might be in a game server container group or a per-instance container group. Support containers don't run game server processes.

You can update a support container definition and deploy the updates to an existing fleet. When creating or updating a game server container group definition, use the property [GameServerContainerDefinitionInput](#).

Part of: [ContainerGroupDefinition](#)

Returned by: [CreateContainerGroupDefinition](#), [DescribeContainerGroupDefinition](#), [ListContainerGroupDefinitions](#), [ListContainerGroupDefinitionVersions](#), [UpdateContainerGroupDefinition](#)

Contents

Note

In the following list, the required parameters are described first.

ContainerName

The container definition identifier. Container names are unique within a container group definition.

Type: String

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

Pattern: `^[a-zA-Z0-9\ -]+$`

Required: No

DependsOn

Indicates that the container relies on the status of other containers in the same container group during its startup and shutdown sequences. A container might have dependencies on multiple containers.

Type: Array of [ContainerDependency](#) objects

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

Required: No

EnvironmentOverride

A set of environment variables that's passed to the container on startup. See the [ContainerDefinition::environment](#) parameter in the *Amazon Elastic Container Service API Reference*.

Type: Array of [ContainerEnvironment](#) objects

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

Required: No

Essential

Indicates whether the container is vital to the container group. If an essential container fails, the entire container group restarts.

Type: Boolean

Required: No

HealthCheck

A configuration for a non-terminal health check. A support container automatically restarts if it stops functioning or if it fails this health check.

Type: [ContainerHealthCheck](#) object

Required: No

ImageUri

The URI to the image that Amazon GameLift Servers deploys to a container fleet. For a more specific identifier, see `ResolvedImageDigest`.

Type: String

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

Pattern: `^[a-zA-Z0-9-_\.\@\/:]+$`

Required: No

LinuxCapabilities

Linux-specific modifications that are applied to the default Docker container configuration, such as Linux capabilities. For more information see [LinuxCapabilities](#).

Type: [LinuxCapabilities](#) object

Required: No

MemoryHardLimitMebibytes

The amount of memory that Amazon GameLift Servers makes available to the container. If memory limits aren't set for an individual container, the container shares the container group's total memory allocation.

Related data type: [ContainerGroupDefinition TotalMemoryLimitMebibytes](#)

Type: Integer

Valid Range: Minimum value of 4. Maximum value of 1024000.

Required: No

MountPoints

A mount point that binds a path inside the container to a file or directory on the host system and lets it access the file or directory.

Type: Array of [ContainerMountPoint](#) objects

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

Required: No

PortConfiguration

A set of ports that allow access to the container from external users. Processes running in the container can bind to a one of these ports. Container ports aren't directly accessed by inbound traffic. Amazon GameLift Servers maps these container ports to externally accessible connection ports, which are assigned as needed from the container fleet's `ConnectionPortRange`.

Type: [ContainerPortConfiguration](#) object

Required: No

ResolvedImageDigest

A unique and immutable identifier for the container image. The digest is a SHA 256 hash of the container image manifest.

Type: String

Pattern: `^sha256:[a-fA-F0-9]{64}$`

Required: No

Vcpu

The number of vCPU units that are reserved for the container. If no resources are reserved, the container shares the total vCPU limit for the container group.

Related data type: [ContainerGroupDefinition TotalVcpuLimit](#)

Type: Double

Valid Range: Minimum value of 0.125. Maximum value of 10.

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](#)

SupportContainerDefinitionInput

Describes a support container in a container group. You can define a support container in either a game server container group or a per-instance container group. Support containers don't run game server processes.

This definition includes container configuration, resources, and start instructions. Use this data type when creating or updating a container group definition. For properties of a deployed support container, see [SupportContainerDefinition](#).

Use with: [CreateContainerGroupDefinition](#), [UpdateContainerGroupDefinition](#)

Contents

Note

In the following list, the required parameters are described first.

ContainerName

A string that uniquely identifies the container definition within a container group.

Type: String

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

Pattern: `^[a-zA-Z0-9\-\_]+$`

Required: Yes

ImageUri

The location of the container image to deploy to a container fleet. Provide an image in an Amazon Elastic Container Registry public or private repository. The repository must be in the same AWS account and AWS Region where you're creating the container group definition. For limits on image size, see [Amazon GameLift Servers endpoints and quotas](#). You can use any of the following image URI formats:

- Image ID only: `[AWS account].dkr.ecr.[AWS region].amazonaws.com/[repository ID]`

- Image ID and digest: [AWS account].dkr.ecr.[AWS region].amazonaws.com/[repository ID]@[digest]
- Image ID and tag: [AWS account].dkr.ecr.[AWS region].amazonaws.com/[repository ID]:[tag]

Type: String

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

Pattern: `^[a-zA-Z0-9-_\.\@\/:]+$`

Required: Yes

DependsOn

Establishes dependencies between this container and the status of other containers in the same container group. A container can have dependencies on multiple different containers.

.

You can use dependencies to establish a startup/shutdown sequence across the container group. For example, you might specify that *ContainerB* has a START dependency on *ContainerA*. This dependency means that *ContainerB* can't start until after *ContainerA* has started. This dependency is reversed on shutdown, which means that *ContainerB* must shut down before *ContainerA* can shut down.

Type: Array of [ContainerDependency](#) objects

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

Required: No

EnvironmentOverride

A set of environment variables to pass to the container on startup. See the [ContainerDefinition::environment](#) parameter in the *Amazon Elastic Container Service API Reference*.

Type: Array of [ContainerEnvironment](#) objects

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

Required: No

Essential

Flags the container as vital for the container group to function properly. If an essential container fails, the entire container group restarts. At least one support container in a per-instance container group must be essential. When flagging a container as essential, also configure a health check so that the container can signal that it's healthy.

Type: Boolean

Required: No

HealthCheck

Configuration for a non-terminal health check. A container automatically restarts if it stops functioning. With a health check, you can define additional reasons to flag a container as unhealthy and restart it. If an essential container fails a health check, the entire container group restarts.

Type: [ContainerHealthCheck](#) object

Required: No

LinuxCapabilities

Linux-specific modifications that are applied to the default Docker container configuration, such as Linux capabilities. For more information see [LinuxCapabilities](#).

Type: [LinuxCapabilities](#) object

Required: No

MemoryHardLimitMebibytes

A specified amount of memory (in MiB) to reserve for this container. If you don't specify a container-specific memory limit, the container shares the container group's total memory allocation.

Related data type: [ContainerGroupDefinition](#)TotalMemoryLimitMebibytes

Type: Integer

Valid Range: Minimum value of 4. Maximum value of 1024000.

Required: No

MountPoints

A mount point that binds a path inside the container to a file or directory on the host system and lets it access the file or directory.

Type: Array of [ContainerMountPoint](#) objects

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

Required: No

PortConfiguration

A set of ports that Amazon GameLift Servers can assign to processes in a container. The container port configuration must have enough ports for each container process that accepts inbound traffic connections. A container port configuration can have one or more container port ranges. Each range specifies starting and ending values as well as the supported network protocol.

Container ports aren't directly accessed by inbound traffic. Amazon GameLift Servers maps each container port to an externally accessible connection port (see the container fleet property `ConnectionPortRange`).

Type: [ContainerPortConfiguration](#) object

Required: No

Vcpu

The number of vCPU units to reserve for this container. The container can use more resources when needed, if available. If you don't reserve CPU units for this container, it shares the container group's total vCPU limit.

Related data type: [ContainerGroupDefinition](#) `TotalCpuLimit`

Type: Double

Valid Range: Minimum value of 0.125. Maximum value of 10.

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](#)

Tag

A label that you can assign to a Amazon GameLift Servers resource.

Learn more

[Tagging AWS Resources](#) in the *AWS General Reference*

[AWS Tagging Strategies](#)

Related actions

[All APIs by task](#)

Contents

Note

In the following list, the required parameters are described first.

Key

The key for a developer-defined key value pair for tagging an AWS resource.

Type: String

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

Required: Yes

Value

The value for a developer-defined key value pair for tagging an AWS resource.

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](#)

TargetConfiguration

Settings for a target-based scaling policy. A target-based policy tracks a particular fleet metric specifies a target value for the metric. As player usage changes, the policy triggers Amazon GameLift Servers to adjust capacity so that the metric returns to the target value. The target configuration specifies settings as needed for the target based policy, including the target value.

Contents

Note

In the following list, the required parameters are described first.

TargetValue

Desired value to use with a target-based scaling policy. The value must be relevant for whatever metric the scaling policy is using. For example, in a policy using the metric `PercentAvailableGameSessions`, the target value should be the preferred size of the fleet's buffer (the percent of capacity that should be idle and ready for new game sessions).

Type: Double

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](#)

TargetTrackingConfiguration

Settings for a target-based scaling policy as part of a [GameServerGroupAutoScalingPolicy](#) . These settings are used to create a target-based policy that tracks the Amazon GameLift Servers FleetIQ metric "PercentUtilizedGameServers" and specifies a target value for the metric. As player usage changes, the policy triggers to adjust the game server group capacity so that the metric returns to the target value.

Contents

Note

In the following list, the required parameters are described first.

TargetValue

Desired value to use with a game server group target-based scaling policy.

Type: Double

Valid Range: Minimum value of 0.

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](#)

UDPEndpoint

The domain name and port information for a UDP endpoint.

For Amazon GameLift Servers ping beacons, this endpoint is a fixed address that remains available as long as Amazon GameLift Servers supports game hosting in that location. For more information on UDP ping beacons, including guidance on how to implement latency measurements and best practices for caching beacon endpoints, refer to [UDP ping beacons](#) in the Amazon GameLift Servers Developer Guide.

Contents

Note

In the following list, the required parameters are described first.

Domain

The domain name of the UDP endpoint.

For Amazon GameLift Servers ping beacons, this is typically in the format `gamelift-ping.<region>.api.aws` for AWS Regions, and `gamelift-ping-<local-zone-code>.<region>.api.aws` for AWS Local Zones.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Port

The port number of the UDP endpoint. For Amazon GameLift Servers ping beacons, this is typically port 7770.

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](#)

VpcPeeringAuthorization

Represents an authorization for a VPC peering connection between the VPC for an Amazon GameLift Servers fleet and another VPC on an account you have access to. This authorization must exist and be valid for the peering connection to be established. Authorizations are valid for 24 hours after they are issued.

Related actions

[All APIs by task](#)

Contents

Note

In the following list, the required parameters are described first.

CreationTime

Time stamp indicating when this authorization was issued. Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

ExpirationTime

Time stamp indicating when this authorization expires (24 hours after issuance). Format is a number expressed in Unix time as milliseconds (for example "1469498468.057").

Type: Timestamp

Required: No

GameLiftAwsAccountId

A unique identifier for the AWS account that you use to manage your Amazon GameLift Servers fleet. You can find your Account ID in the AWS Management Console under account settings.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

PeerVpcAwsAccountId

The authorization's peer VPC AWS account ID.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

PeerVpcId

A unique identifier for a VPC with resources to be accessed by your Amazon GameLift Servers fleet. The VPC must be in the same Region as your fleet. To look up a VPC ID, use the [VPC Dashboard](#) in the AWS Management Console. Learn more about VPC peering in [VPC Peering with Amazon GameLift Servers Fleets](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

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](#)

VpcPeeringConnection

Represents a peering connection between a VPC on one of your AWS accounts and the VPC for your Amazon GameLift Servers fleets. This record may be for an active peering connection or a pending connection that has not yet been established.

Related actions

[All APIs by task](#)

Contents

Note

In the following list, the required parameters are described first.

FleetArn

The Amazon Resource Name ([ARN](#)) associated with the GameLift fleet resource for this connection.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `^arn:.*:[a-z]*fleet\[a-z]*fleet-[a-zA-Z0-9\-\-]+\`

Required: No

FleetId

A unique identifier for the fleet. This ID determines the ID of the Amazon GameLift Servers VPC for your fleet.

Type: String

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

Pattern: `^[a-z]*fleet-[a-zA-Z0-9\-\-]+\`

Required: No

GameLiftVpcId

A unique identifier for the VPC that contains the Amazon GameLift Servers fleet for this connection. This VPC is managed by Amazon GameLift Servers and does not appear in your AWS account.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

IpV4CidrBlock

CIDR block of IPv4 addresses assigned to the VPC peering connection for the GameLift VPC. The peered VPC also has an IPv4 CIDR block associated with it; these blocks cannot overlap or the peering connection cannot be created.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

PeerVpcId

A unique identifier for a VPC with resources to be accessed by your Amazon GameLift Servers fleet. The VPC must be in the same Region as your fleet. To look up a VPC ID, use the [VPC Dashboard](#) in the AWS Management Console. Learn more about VPC peering in [VPC Peering with Amazon GameLift Servers Fleets](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Status

The status information about the connection. Status indicates if a connection is pending, successful, or failed.

Type: [VpcPeeringConnectionStatus](#) object

Required: No

VpcPeeringConnectionId

A unique identifier that is automatically assigned to the connection record. This ID is referenced in VPC peering connection events, and is used when deleting a connection.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

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](#)

VpcPeeringConnectionStatus

Represents status information for a VPC peering connection. Status codes and messages are provided from EC2 (see [VpcPeeringConnectionStateReason](#)). Connection status information is also communicated as a fleet event.

Contents

Note

In the following list, the required parameters are described first.

Code

Code indicating the status of a VPC peering connection.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Message

Additional messaging associated with the connection status.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

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](#)

Common Parameters

The following list contains the parameters that all actions use for signing Signature Version 4 requests with a query string. Any action-specific parameters are listed in the topic for that action. For more information about Signature Version 4, see [Signing AWS API requests](#) in the *IAM User Guide*.

X-Amz-Algorithm

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

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

Type: string

Valid Values: AWS4-HMAC-SHA256

Required: Conditional

X-Amz-Credential

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

For more information, see [Create a signed AWS API request](#) in the *IAM User Guide*.

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

Type: string

Required: Conditional

X-Amz-Date

The date that is used to create the signature. The format must be ISO 8601 basic format (YYYYMMDD'T'HHMMSS'Z'). For example, the following date time is a valid X-Amz-Date value: 20120325T120000Z.

Condition: X-Amz-Date is optional for all requests; it can be used to override the date used for signing requests. If the Date header is specified in the ISO 8601 basic format, X-Amz-Date is not required. When X-Amz-Date is used, it always overrides the value of the Date header. For more information, see [Elements of an AWS API request signature](#) in the *IAM User Guide*.

Type: string

Required: Conditional

X-Amz-Security-Token

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

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

Type: string

Required: Conditional

X-Amz-Signature

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

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

Type: string

Required: Conditional

X-Amz-SignedHeaders

Specifies all the HTTP headers that were included as part of the canonical request. For more information about specifying signed headers, see [Create a signed AWS API request](#) in the *IAM User Guide*.

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

Type: string

Required: Conditional

Common Error Types

This section lists common error types that this AWS service may return. Not all services return all error types listed here. For errors specific to an API action for this service, see the topic for that API action.

AccessDeniedException

You don't have permission to perform this action. Verify that your IAM policy includes the required permissions.

HTTP Status Code: 403

ExpiredTokenException

The security token included in the request has expired. Request a new security token and try again.

HTTP Status Code: 403

IncompleteSignature

The request signature doesn't conform to AWS standards. Verify that you're using valid AWS credentials and that your request is properly formatted. If you're using an SDK, ensure it's up to date.

HTTP Status Code: 403

InternalFailure

The request can't be processed right now because of an internal server issue. Try again later. If the problem persists, contact AWS Support.

HTTP Status Code: 500

MalformedHttpRequestException

The request body can't be processed. This typically happens when the request body can't be decompressed using the specified content encoding algorithm. Verify that the content encoding header matches the compression format used.

HTTP Status Code: 400

NotAuthorized

You don't have permissions to perform this action. Verify that your IAM policy includes the required permissions.

HTTP Status Code: 401

OptInRequired

Your AWS account needs a subscription for this service. Verify that you've enabled the service in your account.

HTTP Status Code: 403

RequestAbortedException

The request was aborted before a response could be returned. This typically happens when the client closes the connection.

HTTP Status Code: 400

RequestEntityTooLargeException

The request entity is too large. Reduce the size of the request body and try again.

HTTP Status Code: 413

RequestTimeoutException

The request timed out. The server didn't receive the complete request within the expected time frame. Try again.

HTTP Status Code: 408

ServiceUnavailable

The service is temporarily unavailable. Try again later.

HTTP Status Code: 503

ThrottlingException

Your request rate is too high. The AWS SDKs automatically retry requests that receive this exception. Reduce the frequency of requests.

HTTP Status Code: 400

UnknownOperationException

The action or operation isn't recognized. Verify that the action name is spelled correctly and that it's supported by the API version you're using.

HTTP Status Code: 404

UnrecognizedClientException

The X.509 certificate or AWS access key ID you provided doesn't exist in our records. Verify that you're using valid credentials and that they haven't expired.

HTTP Status Code: 403

ValidationError

The input doesn't meet the required format or constraints. Check that all required parameters are included and that values are valid.

HTTP Status Code: 400