

Amazon GameLift Servers



Amazon GameLift Servers: Hosting Guide

Table of Contents

What is Amazon GameLift Servers?	1
Key features	1
How to work with Amazon GameLift Servers	4
Game hosting options	4
Managed EC2	5
How it works	5
Core components	6
How components work together	12
Player experience	14
Provide fast gameplay to players globally	14
Help players join games quickly	14
Build creative ways to group players for game sessions	15
Deliver reliable connections throughout gameplay	15
Improve player experience based on real-world data	16
Integrate enhanced game features with other Amazon services	16
Game hosting in production	17
Monitoring hosting health and performance	17
Managing game server updates and patches	18
Optimizing performance and scaling	18
Troubleshooting and live operations	19
Tools and integration	19
Related resources	20
Key terms	20
Getting started	28
Before you start	28
Choose your path	28
I want to explore Amazon GameLift Servers quickly	28
I want to build a custom hosting solution	29
Recommended learning path	29
Set up an Amazon account	30
Sign up for an Amazon Web Services account	31
Secure IAM users	31
Set user permissions for Amazon GameLift Servers	32
Set up programmatic access for users	33

Set up programmatic access for your game	33
IAM permission examples	34
Set up an IAM service role	39
Get development tools	43
For game servers	43
For game client services	44
For resource management	45
For Amazon GameLift Servers Realtime	45
Tutorial: Quick onboarding with the Amazon GameLift Servers wrapper	46
Prerequisites	47
Overview	47
Step 1: Get and build the game server wrapper	48
Step 2: Prepare your game server build	48
Step 3: Configure the wrapper for your fleet	49
Step 4: Upload the game server build	50
Step 5: Create the managed EC2 fleet	51
Step 6: Create and connect to a game session	52
Step 7: Manage and monitor your fleet	53
Step 8: Scale your game servers	54
Troubleshoot common issues	55
Next steps	56
Amazon GameLift Servers examples	56
Custom game server example	57
Building a game hosting solution	58
Development roadmaps	58
Managed EC2 development roadmap	59
Prepare a game for hosting	64
Prepare with the game engine plugin	65
Integrate a game server	88
Package a game server build	99
Amazon GameLift Servers interactions	112
Build a backend service	117
Authenticating your players	117
Integrate game client functionality	118
Serverless backend	123
WebSocket-based backend	125

Deploy hosting fleets	127
Fleet characteristics	127
Choose a hosting option	128
Managed EC2 fleets	130
Build a hybrid hosting solution	140
Configure game session placement	142
Queue characteristics	143
Best practices	144
Create a queue	145
Set up event notification	149
Customize your game hosting solution	153
Game server builds	154
Player sessions and matchmaking	164
Game session placements	166
Hosting resources	180
Managing a game hosting solution	203
Manage game hosting resources	203
Amazon GameLift Servers console features	203
Console dashboard	204
Builds	205
Fleets	208
Aliases	229
Game session queues	232
Prepare for game launch	234
Getting ready	235
Prepare for testing	235
Prepare for launch	236
Plan for post-launch	236
Manage game hosting resources with Amazon CloudFormation	237
Best practices	237
Using Amazon CloudFormation stacks	238
Updating builds	243
Testing and troubleshooting	246
Debug fleet issues	246
Fleet creation issues	246
Server process issues	248

Fleet deletion issues	248
Remotely connect to fleet instances	249
Remote access through the console	250
Remote access with the Amazon CLI	250
Connect to an instance (server SDK 5)	253
Connect to an instance (server SDK 4.x or earlier)	254
View files on remote instances	257
Scaling	258
To manage fleet capacity in the console	258
Set hosting capacity limits	259
To set capacity limits	259
Manually set fleet capacity	261
Suspend auto scaling	261
To manually set fleet capacity	261
Auto scale fleet capacity	263
Target-based auto scaling	263
Rule-based auto scaling	265
Manage Scaling a Fleet To/From Zero	269
Scale-in behavior	270
Scale-out behavior	270
Configuring Scale To/From Zero	271
Best practices	272
Scale container fleets	272
Monitoring	275
Metrics comparison	276
Metrics availability by source	276
Topics	279
Monitor with CloudWatch	280
Metrics dimensions	280
Fleet metrics	282
Queue metrics	297
FlexMatch metrics	301
FleetIQ metrics	305
Monitor with server telemetry metrics	307
Benefits of telemetry metrics	308
Before you begin	309

Implementation	310
Available metrics	311
How it works	315
Dashboard organization and usage	316
Common monitoring scenarios	327
Troubleshooting guide	332
Logging API calls	334
Amazon GameLift Servers information in CloudTrail	335
Understanding Amazon GameLift Servers log file entries	336
Logging server messages	338
EC2 vs Container logging	339
Logging for custom servers	341
Security	345
Data protection	346
Encryption at rest	347
Encryption in transit	348
Internetwork traffic privacy	349
Identity and access management	349
Audience	349
Authenticating with identities	350
Managing access using policies	351
How Amazon GameLift Servers works with IAM	352
Identity-based policy examples	359
Troubleshooting	364
Amazon managed policies	366
Logging and monitoring with Amazon GameLift Servers	368
Compliance validation	369
Resilience	369
Infrastructure security	370
Configuration and vulnerability analysis	371
Security best practices	372
Keep fleet runtime environments up to date	372
Secure your port configurations	373
Additional security resources	374
Pricing	375
Pricing models	375

Generate pricing estimates	376
Estimate managed hosting	376
Cost optimization strategies	379
Best practices	379
Resource cost and utilization tools	381
Performance monitoring	382
Reference guides	383
Service API (Amazon SDK)	383
Manage Amazon GameLift Servers hosting resources	383
Start game sessions and join players	387
Server SDK 4 and earlier	388
Server SDK for C++: Actions	389
Server SDK for C#: Actions	412
Server SDK for Unreal: Actions	432
Game session placement events	447
Placement event syntax	447
PlacementFulfilled	448
PlacementCancelled	450
PlacementTimedOut	451
PlacementFailed	452
AMI versions	453
Service locations	455
Locations for managed hosting	455
Locations for Amazon GameLift Servers FlexMatch	456
Amazon GameLift Servers in China	457
Service endpoints and quotas	457
Amazon GameLift Servers API limits	457
API limits reference	458
Resource-level throttling	467
Amazon GameLift Servers SDK API limits	467
Release notes and SDK versions	471
SDK versions	471
Release notes	477

What is Amazon GameLift Servers?

Use Amazon GameLift Servers to deploy, operate, and scale dedicated, low-cost servers in the cloud for session-based multiplayer games. Built on Amazon global computing infrastructure, Amazon GameLift Servers helps deliver high-performance, high-reliability game servers while dynamically scaling your resource usage to meet worldwide player demand.

Key features

Global reach and low latency

- Deploy hosting in Amazon Regions and Local Zones worldwide to bring your games to new players. See a complete list of [locations supported by Amazon GameLift Servers](#).
- Deliver lowest possible latency and great player experience, wherever your players are located.
- Use multi-location placement queues to dynamically place players into game sessions with optimal player experiences.

Automatic scaling for managed fleets

- Optimize fleet capacity for cost efficiency while maintaining quality of service.
- Scale capacity based on player demand using target-based auto scaling
- Maintain capacity buffer to handle sudden player influxes.

Flexible hosting options

See more details at [Amazon GameLift Servers game hosting options](#).

- Game hosting in the cloud, managed by Amazon GameLift Servers
 - Host game servers on Amazon Elastic Compute Cloud (Amazon EC2) infrastructure. Choose from resource types including compute-optimized (C-family), memory-optimized (R-family), and cost-efficient ARM-based Graviton instances.
 - Customize how Amazon GameLift Servers manages all aspects of your game hosting solution, from deploying game servers and starting game sessions to matching players and getting them into games.
 - Use multi-location fleets to streamline global game server deployment.

- Take advantage of low-cost hosting with [Amazon EC2 Spot Instances](#).
- Deploy Windows or Linux game servers with managed EC2 hosting.
- Deploy Linux-based containers with your game server with managed container hosting.
- Game hosting managed by you, with Amazon GameLift Servers Anywhere
 - Leverage Amazon GameLift Servers game session placement and session management capabilities.
 - Self-manage game server deployments and scaling on your own hardware, on-premises infrastructure, or other cloud providers.
- Hybrid game hosting
 - Create a game hosting solution that uses a combination of Amazon GameLift Servers managed fleets and Anywhere fleets to work with existing solutions or support migration.

Game session management

- Let Amazon GameLift Servers manage game server processes for you, track game server status and availability, and place new game sessions on request.
- Use game session placement features to help players find and join game sessions that will provide the best possible player experience. Rely on Amazon GameLift Servers decision-making, or customize around criteria such as hosting cost, player latency, and geographic location.
- Create game session placement queues to efficiently process new game session requests. Use features to set a balance between placement quality and speed that fits your game and player expectations.
- Take advantage of optional player session features to reserve player slots in game sessions, validate players on connection, and monitor slot availability. Search or filter existing game sessions to find available slots.

Advanced matchmaking with FlexMatch

- Add robust matchmaking that's fully integrated with your Amazon GameLift Servers hosting solution.
- Build custom matchmaking rules based on player attributes and preferences, such as skill level or character choice, and use complex logic as needed.
- Set up team-based matches for competitive or cooperative gameplay. Balance players across team and/or match composition.

- Optimize the matchmaking algorithm to balance match quality and wait times.
- Support matches from 2 to 200 concurrent players.
- Support match backfilling to fill available player slots and keep game sessions full.

Comprehensive monitoring

- Get real-time metrics on usage and performance for game sessions and player activity, as well as hosting performance and health for managed fleets.
- Integrate with Amazon CloudWatch for alerts and dashboards.
- Access server logs and event tracking for troubleshooting.
- Analyze performance and utilization data to optimize scaling (managed fleets) or resource allocation (self-managed fleets).

Customization and extensibility

- Customize game session placement logic, including prioritizing for latency, cost, and location.
- For managed hosting, choose your game hosting resources and provide instructions for launching and running game servers on each instance. For managed containers, fine-tune resource allocation across containers.
- Define your own matchmaking rules.
- Implement custom strategies for capacity scaling.
- Add game server code to integrate with your existing Amazon services, such as databases or content storage tools.
- Create a backend service with custom features to manage player join requests (such as player parties), handle authentication, or add custom logic and persistence.

Integration with Amazon ecosystem

- Use services such as Amazon DynamoDB, Amazon Simple Storage Service, and Amazon Aurora DSQL for game state persistence.
- Implement Amazon Cognito for player authentication.
- Process game analytics with Amazon Kinesis and Amazon S3.
- Add voice chat with Amazon Chime SDK.

- Create custom game features using Amazon Lambda and Amazon API Gateway.

How to work with Amazon GameLift Servers

Use these tools to work with Amazon GameLift Servers.

Amazon CLI

Use the Amazon Command Line Interface (Amazon CLI) to make calls to the Amazon SDK, including the service API for Amazon GameLift Servers. See [Getting started with the Amazon CLI](#) in the *Amazon Command Line Interface User Guide*.

Amazon GameLift Servers console

Use the [Amazon Web Services Management Console for Amazon GameLift Servers](#) to configure resources, manage your game server deployments, and track performance and usage metrics. The Amazon GameLift Servers console is a GUI alternative to managing resources programmatically or with the Amazon CLI.

Amazon GameLift Servers SDKs

The Amazon GameLift Servers SDKs contain the libraries required to establish communication between your game clients, game servers, and game services and the Amazon GameLift Servers service. For more information, see [Get Amazon GameLift Servers development tools](#).

Amazon CloudFormation

Use Amazon CloudFormation to model, provision, and manage Amazon resources for your game hosting solution by treating infrastructure as code. Create templates that describe the resources, and Amazon CloudFormation automates the tasks of configuring and deploying resources to the locations you specify.

Amazon GameLift Servers game hosting options

Amazon GameLift Servers offers a range of options for hosting game servers for session-based multiplayer games. Whether you want to set up cloud-based hosting managed by Amazon GameLift Servers or incorporate hosting resources that you already have, you can work with Amazon GameLift Servers to build the hosting solution you need for your players. For more details about what a game hosting solution looks like, see [How hosting with Amazon GameLift Servers works](#).

Topics

- [Managed EC2](#)

Managed EC2

Offload the work of managing your production game servers onto Amazon GameLift Servers with managed EC2 hosting. Take advantage of optimizations for multiplayer game servers while still relying on the high performance and reliability of Amazon Elastic Compute Cloud (Amazon EC2) and Amazon global computing infrastructure.

[Start building an Amazon GameLift Servers managed EC2 hosting solution](#)

Characteristics

- Host multiplayer game servers that run on Amazon Linux or Windows Server operating systems.
- Use Amazon GameLift Servers to deploy and manage your custom game servers to hosting resources Amazon Web Services Cloud. Choose the hardware type and where to deploy it, and configure additional details only where you need to.
- Configure the runtime process manager to automatically maintain game server processes as needed to host game sessions.
- Set up custom health tracking to help Amazon GameLift Servers detect and resolve poor-performing game servers.
- Take advantage of Amazon GameLift Servers performance monitoring. Track metrics such as hardware performance and server process life cycles. You can also download and store game session logs.
- Use Amazon GameLift Servers to manage access between game server processes and other Amazon resources.
- For production hosting, automate game hosting deployments and management with Amazon CloudFormation templates and the Amazon Cloud Development Kit (Amazon CDK). Take advantage of continuous integration and continuous delivery (CI/CD) tools and services such as Amazon CodePipeline.

How hosting with Amazon GameLift Servers works

Amazon GameLift Servers is a service that provides dedicated, low-cost servers, infrastructure provisioning, scaling, and session management for your session-based multiplayer games.


The service offers flexible tools and features that you can customize for your games or use in collaboration with your own systems. Take advantage of Amazon GameLift Servers managed hosting so you can focus on creating gameplay experiences rather than managing server operations.

This topic describes the core components of a complete Amazon GameLift Servers hosting solution and how they work together to supply multiplayer game sessions to players.

Core components of a game hosting solution

A complete Amazon GameLift Servers hosting solution consists of several essential components that work together to deliver multiplayer game experiences. Some of these components are built and managed by you, while others are provided by Amazon GameLift Servers and managed based on your configuration choices. Understanding how these components interact and their respective roles is crucial for developing an effective game hosting solution. The components include:

- Game client runs on player devices.
- Backend service enables game clients to communicate with Amazon GameLift Servers to get game session information and join games.
- Game server software runs on game hosting resources and hosts game sessions for a group of players.
- Placement system initiates game sessions and matches players to games.
- Game hosting fleets provide game servers in one or more geographic locations.
- Game hosting management system monitors game hosting status and manages capacity.



Game client

The game client is your game software that runs on a player's device. It initiates player placement into a game session by communicating with a backend service and connects directly to a game server to participate in gameplay.

Key functions for game hosting

- Send join requests to the backend service. Include relevant player data and game session data as part of the join request.
- Collect latency data for the game client and include that information in the join request.
- Receive game session connection information from the backend service and use it to connect to a game server.

- Handle session connection scenarios such as player verification, match acceptance for Amazon GameLift Servers FlexMatch, session interruptions, or disconnects.

Who builds it

You add game hosting functionality to your game client software and set up communication with the backend service.

Backend service

The backend service is a coordination layer between game clients and the Amazon GameLift Servers service. It controls all communication with the service to make requests for game session placements and to retrieve game session and player session information. Use of a backend service is a best practice that maintains secure communication with the service on behalf of game clients and avoids having to share sensitive Amazon credentials and permissions with game clients.

Your backend service implements a player grouping strategy for your game sessions. You have a lot of flexibility in how players end up in a game session. You might set up a game session browser for players or support player-defined parties. Or you might pool game session requests as they come in, form player groups (possibly with pre-sorting), and start a game session for each group. With FlexMatch matchmaking, you can form player matches and backfill existing matches.

Key functions for game hosting

- Authenticate communications from game clients.
- Make requests to Amazon GameLift Servers to start new game sessions, get game session information, or join existing game sessions. Include relevant game session and player data in requests as needed.
- Optionally create requests for player sessions to more closely monitor game session availability. Use player sessions to reserve game session slots, validate players when they connect, and track player disconnects.
- Respond to game clients with game session connection details and other information as requested.

Who builds it

You build the backend service in an environment that's managed by you to support your game requirements and implement how players get into games. Integrate the Amazon SDK to make calls

to the Amazon GameLift Servers service API. Optionally build your backend service on Amazon using services such as Amazon Lambda, Amazon Simple Storage Service (Amazon S3), and Amazon Cognito.

Game server

The game server is your custom server software that manages game state, processes player actions, and synchronizes gameplay across multiple connected players. The game server maintains communication with the Amazon GameLift Servers service to manage game session hosting.

Key functions for game hosting

- Communicate with Amazon GameLift Servers to:
 - report status (ready to host sessions, ready to accept players, health status).
 - respond to service calls (start or end game sessions).
- Manage game session lifecycle to host one game session at a time per process.
- Coordinate with other Amazon services for added functionality.
- Optionally validate new player connections.

Who builds it

You build your game server software. You integrate the server SDK for Amazon GameLift Servers and add functionality to establish a connection with the service and support game session management.

Game hosting fleet

The hosting fleet is a collection of computing resources that run your game servers. Fleet resources can be distributed across multiple geographic locations to provide low-latency gameplay to players wherever they are. Each fleet resource runs one or more game server processes, which communicates directly with Amazon GameLift Servers. Each game server process can host one game session at a time.

Fleet characteristics and functionality vary based on the fleet's hosting type. Managed fleets deploy resources to the Amazon Web Services Cloud and are managed by Amazon GameLift Servers. Anywhere fleets are customer-provided compute resources that are managed outside of Amazon GameLift Servers.

Key functions for game hosting

- Provision game server hosting resources.
 - Managed fleets deploy cloud-based Amazon EC2 instances with a wide range of configuration settings. A fleet configuration determines the computing power of each instance in the fleet, the physical location of instances, and other details. Managed fleets add or remove EC2 instances in response to capacity scaling
 - Anywhere fleet deployments are self-managed. The fleet can consist of physical hardware or other cloud resources, and be configured as needed.
- Install runtime environment and game server software.
 - Managed Amazon EC2 fleet instances are deployed with your game server build and an Amazon Machine Image (AMI) with a compatible runtime environment.
 - Managed container fleet instances are deployed with your game server build and a container-optimized AMI with Docker tools and other components to work with Amazon ECS.
 - Anywhere fleets are deployed with your chosen game server software and operating system.
- Manage game server process life cycle. Apply pre-configured runtime instructions to start and stop processes on each fleet resource.

Who builds it

This depends on the fleet's hosting type:

- Managed Amazon EC2 fleets: You upload your game server build to Amazon GameLift Servers, which stores it for deployment to fleet instances. Amazon GameLift Servers provides the fleet's Amazon Web Services Cloud infrastructure. You configure the fleet and instruct it on how to run game servers on each fleet instance.
- Managed container fleets: You package your game server build and runtime instructions into a container image to store in Amazon ECR for deployment. Amazon GameLift Servers provides the fleet's Amazon Web Services Cloud infrastructure. You provide a container architecture and configure the fleet to host your containers.
- Anywhere fleets: You provision all infrastructure and manage server software deployment. You create Anywhere fleets to connect your active game hosting resources with Amazon GameLift Servers.

Game session placement system

The game session placement system locates available game servers to host new game sessions. The system uses real-time information about game server availability to make optimal placement decisions.

In Amazon GameLift Servers, the primary game session placement mechanism is the queue. A game session queue uses algorithms, which you can configure, to place game sessions for the best possible outcome. You can prioritize placements based on factors such as lowest hosting cost and lowest player latency, and you can configure a queue to search across multiple geographic locations. As an alternative to queues, you can designate a specific fleet to host your game sessions.

Key functions for game hosting

- Process game session placement requests received from the backend service.
- Make placements based on real-time information about hosting resource availability.
- Use player latency data and other data to prioritize placement options (queues only)
- Prompt game server processes to start new game sessions.
- Update game session connection information after the game session is ready to accept players.
- Optionally set up FlexMatch matchmaking to create player matches and request game session placement for matches.

Who builds it

Amazon GameLift Servers supplies the placement system. You configure placement behavior through how your backend service makes placement requests and optionally by setting up game session queues and FlexMatch matchmakers. Use the Amazon GameLift Servers console, the Amazon SDK, or the Amazon CLI to create and configure queues and matchmakers for your game.

Game hosting management system

The game hosting management system is the operational backbone that coordinates and monitors all aspects of your game hosting solution. This system provides the intelligence and automation that makes Amazon GameLift Servers' placement and scaling capabilities possible.

Key functions for game hosting

- Track the real-time status and availability of game server processes, game sessions, and player sessions across all fleets to support game session placement and automatic capacity scaling.
- Monitor fleet health and performance.
- Collect and analyze game hosting activity metrics.
- Provide capacity scaling tools, including automatic scaling based on player demand and fleet utilization.
- Manage updates to game server software and runtime instructions.

Who builds it

- **Managed fleets:** Amazon GameLift Servers provides availability tracking, metrics on game hosting activity and hardware performance, and capacity scaling tools. You use available Amazon tools to manage game server software updates, modify runtime instructions, and customize your use of metrics (such as setting up CloudWatch for monitoring). You configure a custom scaling policy and modify as needed.
- **Anywhere fleets:** Amazon GameLift Servers provides availability tracking and game hosting activity metrics. You manage fleet configuration changes, including updates to game server software and runtime instructions. You create systems to monitor fleet performance and manage fleet capacity scaling.

How the components work together

When a game hosting compute is deployed

- **Game server software installed:** The compute is installed with a runtime environment and your game server build
- **Game server launches:** At least one instance of the game server executable is launched, with optional launch parameters, on the compute.
- **Game server connects to Amazon GameLift Servers:** As part of its startup actions, the game server process calls the server SDK to initialize a connection to the service.
- **Game server reports ready to host a game session:** The game server process completes startup actions, then calls the server SDK to report readiness. It begins reporting health status based on its configuration.

- **Amazon GameLift Servers tracks availability:** The service records the game server process's availability to game session placement. It also updates metrics for fleets, instances, and game sessions to track usage and capacity.
- **Game server waits for game session assignment:** The game server process maintains its idle status while it waits for a prompt to start a game session.

When a player wants to join a game

- **Player initiates a game:** A player launches your game client, authenticates with your backend service, and asks to join a game.
- **Backend service requests a game session:** Your backend service calls Amazon GameLift Servers to find or create a suitable game session, based on its player grouping strategy. The request might include player or game data for use in the game session.
- **Amazon GameLift Servers places the game session:** When starting a new game session, the placement system identifies an optimal location to host the session and selects an available game server process. The service prompts the selected process to start a new game session and passes any player or game data. If FlexMatch matchmaking is in use, the matchmaker first creates a match, then requests placement for the match.
- **Game server starts the game session:** The game server initiates steps to start a game session. When complete, it reports to Amazon GameLift Servers that it's ready to accept player connections.
- **Connection information delivered to game client:** After the game server updates its status, Amazon GameLift Servers provides the game session connection information. The backend service receives this information and delivers it to the game client.
- **Player connects to the game session:** The game client uses the connection information to connect directly to the game server and begin gameplay.
- **Amazon GameLift Servers monitors game session status:** The game server process reports health status, optional player connection status, and game session status to track ongoing game session availability.
- **Game server process shuts down:** The game server process ends the game session, reports status, and then shuts itself down.

Amazon GameLift Servers and the player experience

Players expect fast, fair, and reliable multiplayer gaming experiences. Amazon GameLift Servers delivers these experiences by providing game developers with tools that directly enhance how players connect and engage with multiplayer games. This topic outlines the key player benefits Amazon GameLift Servers features provide.

Provide fast gameplay to players globally

Players want to connect to nearby servers with minimal lag. Amazon GameLift Servers ensures optimal connection experiences through intelligent placement and global infrastructure:

- **Players experience consistent performance globally** – Host games in Amazon Regions and Local Zones available globally to minimize the distance between players and game servers (see [Amazon GameLift Servers service locations](#)). You can choose from a wide range of Amazon EC2 instance types in these locations to find the best computing resources to host your game workloads.
- **Players connect to the best available servers** – Amazon GameLift Servers places players in game sessions with the lowest possible latency by gathering game client network data and searching for available servers across multiple locations. Use UDP ping beacons to collect accurate latency data. Configure latency policies to balance connection quality with acceptable wait times for players.
- **Players from different regions can play together** – Game session placements can balance game sessions with players from regional player pools while maintaining acceptable latency. Players from less-populated regions can find matches without sacrificing connection quality.

Help players join games quickly

Players want fast access to games regardless of demand. Amazon GameLift Servers provides a range of capacity scaling tools that help you ensure that players can find games during peak and quiet periods:

- **Players can easily find available servers** – With automatic scaling tools, you can ensure that hosting resources are ready when they're needed. Target tracking manages a flexible buffer of capacity to anticipate player demand patterns for your game and scale proactively.

- **Players experience smooth performance during special events** – Combine auto-scaling with manual scaling techniques to handle sudden influxes after planned events, such as game updates, marketing campaigns, or tournaments.
- **Players don't have to wait for servers to start up** – On every game hosting resource, game servers are pre-warmed so that they're always ready to host new sessions as soon as players join.

Build creative ways to group players for game sessions

Players want to compete or collaborate with others in fair, balanced ways. Amazon GameLift Servers offers a number of options for grouping players into game sessions:

- **Players are grouped to optimize player experience** – You can manage player grouping to best suit your game using a range of features. Let players search or browse game sessions or group player requests when they arrive based on skill level, latency, and other attributes. You can manage player grouping on your own or you can implement FlexMatch matchmaking with customizable match rules.
- **Players can join active matches** – Build player groups to start new game sessions, or use player sessions to track available slots in existing sessions and add new players. With FlexMatch, automatically backfill sessions with well-matched players.
- **Friends can play together as a team** – Player party support lets groups of friends play together. Game sessions can fill remaining slots with other players or restrict access. With FlexMatch, friends can form teams.
- **Players can bring custom data to game sessions** – Include game and player information in requests for game sessions and player sessions. The information gets passed on to the game server for use in a game session.
- **Players don't wait indefinitely for perfect matches** – With FlexMatch, you can relax match requirements over time to prevent player frustration from excessive queuing. Create progressive relaxation rules that prioritize the most important aspects of match quality.
- **Players can choose to accept or decline matches** – With FlexMatch, you can customize workflows for letting players accept or reject prospective matches.

Deliver reliable connections throughout gameplay

Players expect consistent gameplay without interruptions or disconnections. Amazon GameLift Servers provides comprehensive session management and protection:

- **Players enjoy high-quality performance on Amazon Cloud** – With hosting managed by Amazon GameLift Servers, players get fast, reliable game servers that run on Amazon computing infrastructure.
- **Validate players when they connect** – When you enable player sessions for your game, you can have Amazon GameLift Servers reserve player slots in game sessions and validate players on connection.
- **Players can reconnect and maintain game progress** – Amazon GameLift Servers supports reconnection for both server-side and client-side interruptions.
- **Players' games are protected from termination** – Optional game session protection feature prevents active sessions from being terminated during scaling events or other interruptions.

Improve player experience based on real-world data

Players expect games to work consistently and get better over time. Amazon GameLift Servers provides comprehensive monitoring and analytics that help optimize the player experience:

- **Players experience fewer disruptions** – Real-time monitoring of fleet performance, game sessions, and player activity can identify issues before they significantly impact gameplay. Customize graceful game session shutdowns and migrations.
- **Players benefit from proactive server health monitoring** – Amazon GameLift Servers provides continuous server health monitoring and automatically replaces unhealthy game servers to minimize hardware or software failure impact. Configure health check parameters to support different game requirements.
- **Players get data-driven improvements in gameplay** – Take advantage of game session logs and detailed analytics to reveal patterns in player behavior and server performance. Add custom logging support for game-specific events.

Integrate enhanced game features with other Amazon services

Players want integrated features like voice chat, secure authentication, and persistent progress. Amazon GameLift Servers can be integrated with other Amazon services to provide comprehensive gaming experiences:

- **Players get streamlined authentication** – Amazon Cognito integration supports various identity providers and authentication methods while maintaining secure player identities across game sessions, with streamlined login processes for returning players.

- **Players can persist their game progress across sessions** – Amazon DynamoDB integration can be used to store player progression, inventories, and persistent data with high-performance access that ensures minimal gameplay impact while supporting cross-session continuity.
- **Players benefit from analytic insights** – Amazon Kinesis and Amazon Simple Storage Service (Amazon S3) integration processes game analytics to gather insights on player behavior and preferences, enabling real-time analytics that help developers adapt to changing player patterns and keep games fresh and engaging.

Managing game hosting in production with Amazon GameLift Servers

After you've launched your game with Amazon GameLift Servers, you'll need to manage your game hosting infrastructure to ensure optimal performance, reliability, and player experience. This topic covers the key features and tools for managing game hosting lifecycles during production.

Monitoring hosting health and performance

Amazon GameLift Servers provides comprehensive monitoring tools to help you track the health and performance of your game hosting infrastructure:

- **Performance metrics** – Monitor key performance indicators for managed fleets, such as instance utilization, player sessions, and game session placements. Track these metrics in Amazon GameLift Servers or in Amazon CloudWatch, where you can also set up alarms to get notifications when metrics exceed thresholds.
- **Fleet metrics and events** – Track fleet-specific metrics including active server processes, available game sessions, and active player sessions to ensure optimal capacity.
- **Game session placement metrics and events** – Track metrics and events to monitor the health and performance of your game session placement system.
- **Game session logs** – Access and analyze game session logs to identify issues and understand player behavior patterns.
- **Realtime script logs** – For games using Real-time Servers, monitor script execution and performance through detailed logs.
- **Amazon Health Dashboard** – Stay informed about Amazon service health that might affect your Amazon GameLift Servers deployments.

You can access these monitoring tools through the Amazon GameLift Servers console, Amazon CLI, or use the Amazon SDK for Amazon GameLift Servers to create custom dashboards and monitoring solutions.

Managing game server updates and patches

Keeping your game servers updated is critical for security, performance, and adding new features. Amazon GameLift Servers provides several approaches for managing updates:

- **Build management** - Upload and manage multiple versions of your game server builds. Each build is versioned and can be deployed to different fleets.
- **Fleet replacement** - Create new fleets with AMI versions and game server build updates, and gradually shift traffic from old fleets to new ones using game session placement queues and aliases.
- **Script updates** - For Real-time Servers, update server scripts without replacing the entire fleet by uploading new script versions.
- **Container updates** - For container-based deployments, update container images and definitions. Deploy new versions of your game servers to existing fleets. Replace fleets to update AMI versions.
- **Automated deployments** - Use Amazon CodePipeline and Amazon CodeDeploy to create CI/CD pipelines for automated game server updates.

When planning updates, consider using blue/green deployment strategies to minimize disruption to active players and allow for quick rollbacks if issues are detected.

Optimizing performance and scaling

As your player base evolves, you'll need to adjust your hosting configuration to maintain optimal performance and cost-efficiency:

- **Auto scaling** - Configure fleet scaling policies based on metrics like player count or game session utilization to automatically adjust capacity.
- **Regional deployment adjustments** - Add or remove regions from your multi-region deployments based on player demographics and latency requirements.
- **Queue management** - Optimize game session placement queue configurations to balance player experience and hosting costs.

- **Instance type selection** - Analyze performance metrics to determine the most cost-effective instance types for your game server requirements.
- **Spot Instance usage** - Leverage Spot Instances for non-critical workloads to reduce costs, with appropriate fallback strategies to On-Demand instances.
- **FlexMatch tuning** - Refine matchmaking rules and algorithms based on actual player data and feedback.

Regularly review CloudWatch metrics and cost reports to identify optimization opportunities and implement changes through the Amazon GameLift Servers console or API.

Troubleshooting and live operations

Effective troubleshooting and live operations management are essential for maintaining a positive player experience:

- **Fleet event notifications** - Set up Amazon Simple Notification Service notifications for fleet events such as scaling activities, instance terminations, or game session placement failures.
- **Game session placement debugging** - Use detailed placement logs to identify and resolve issues with game session placements.
- **Server process health checks** - Monitor server process health and automatically replace processes that fail health checks.
- **Remote access** - Connect to fleet instances for direct troubleshooting using Amazon Systems Manager Session Manager.
- **Alias management** - Use aliases to quickly redirect player traffic away from problematic fleets without changing client configurations.
- **Backup and recovery** - Implement regular backups of critical game data and configuration to enable quick recovery from failures.

Establish clear incident response procedures and runbooks for common issues to minimize downtime and impact on players.

Tools and integration

Amazon GameLift Servers integrates with various Amazon services and third-party tools to enhance your production management capabilities:

- **Amazon CloudFormation** - Define and manage your GameLift resources as infrastructure as code for consistent deployments.
- **Amazon Lambda** - Create serverless functions to automate routine management tasks and respond to events.
- **Amazon EventBridge** - Build event-driven architectures that respond automatically to changes in your Amazon GameLift Servers environment.
- **Amazon SDK integration** - Use Amazon SDKs to build custom management tools and dashboards tailored to your specific needs.
- **Third-party monitoring** - Integrate with third-party monitoring and analytics platforms using CloudWatch metrics export.

Leverage these integrations to create a comprehensive management solution that aligns with your team's workflows and requirements.

Related resources

- [Monitoring Amazon GameLift with Amazon CloudWatch](#)
- [Updating Your Amazon GameLift Fleets](#)
- [Scaling Game Hosting Capacity with Amazon GameLift](#)
- [Amazon GameLift Server SDK Reference](#)

Amazon GameLift Servers key terms

This topic provides definitions for key technical terms used in Amazon GameLift Servers documentation and services.

Active game session

A game session that is currently running and accepting or hosting players. Used as a metric to monitor fleet utilization and capacity planning. See [Monitor Amazon GameLift Servers with Amazon CloudWatch](#).

Alias

A friendly identifier for a fleet that provides an abstraction layer between your game client and fleet resources. Aliases allow you to update fleet configurations without changing your game client code. See [Abstract an Amazon GameLift Servers fleet designation with an alias](#).

AMI (Amazon Machine Image)

A pre-configured virtual machine image used to create Amazon EC2 instances for managed fleets. Contains the operating system, runtime libraries, and Amazon GameLift Servers agent software needed to run game servers. See [Amazon GameLift Servers managed EC2 fleets](#).

Anywhere fleet

A type of fleet that uses customer-provided compute resources managed outside of Amazon GameLift Servers. Anywhere fleets allow you to integrate existing infrastructure or hybrid cloud solutions with Amazon GameLift Servers game session management. See [???](#).

Authentication token

A security credential generated by Amazon GameLift Servers that authenticates game server processes running on Anywhere compute resources. Required for server processes to communicate with the Amazon GameLift Servers service. See [???](#).

Backend service

A coordination layer between game clients and the Amazon GameLift Servers service that controls communication for game session placement requests and retrieves game session information. Acts as a secure intermediary to protect hosting resources from malicious attacks. See [How hosting with Amazon GameLift Servers works](#).

Game server build

A packaged version of your game server software that includes all files needed to run your game server on Amazon GameLift Servers hosting resources. Contains the game server executable, dependencies, and assets. See [Package a game server build for deployment](#).

Compute

In container fleets, represents a copy of a game server container group on a fleet instance. Each compute can host one game session at a time. See [???](#).

Container fleet

A type of managed fleet that runs containerized game servers using container images instead of traditional builds. Provides flexibility for deploying game servers with specific runtime environments and dependencies. See [???](#).

Container group

A collection of containers that work together to provide game server functionality. Similar to a container task or pod, defining how containers behave, dependencies, and resource sharing within a container fleet. See [???](#).

Destination

In game session queues, a fleet or alias that can receive game session placement requests. Queues can have multiple destinations with different priorities and latency policies. See [Customize a game session queue](#).

Fleet

A collection of computing resources that run your game servers. Fleet resources can be distributed across multiple geographic locations and each resource runs one or more game server processes. See [Deploy hosting fleets for Amazon GameLift Servers](#).

Fleet location

A specific Amazon Region or Local Zone where fleet instances are deployed. Multi-location fleets can span multiple locations to reduce latency and improve availability. See [Deploy hosting fleets for Amazon GameLift Servers](#).

FlexMatch

Amazon GameLift Servers matchmaking service that automatically forms player matches based on configurable rules and requests game session placement for completed matches. Supports skill-based matching, team balancing, and match backfilling. See [What is GameLift FlexMatch?](#)

Game client

Your game software that runs on a player's device. Initiates player placement into game sessions by communicating with a backend service and connects directly to game servers for gameplay. See [How hosting with Amazon GameLift Servers works](#).

Game property

Key-value pairs of custom data associated with a game session. Used to pass configuration information, game rules, or other metadata to game servers and clients. See [GameProperty](#) in the API Reference.

Game server

Your custom server software that manages game state, processes player actions, and synchronizes gameplay across multiple connected players. Integrates with the Amazon

GameLift Servers Server SDK to communicate with the service. See [Add Amazon GameLift Servers to your game server with the server SDK](#).

Game server activation

The process of starting and initializing a game server process on a hosting resource. Includes loading the game server executable, establishing communication with Amazon GameLift Servers, and preparing to host game sessions. See [Add Amazon GameLift Servers to your game server with the server SDK](#).

Game server container

A container within a game server container group that runs the game server application. Automatically designated as essential to the container group lifecycle. See [???](#).

Game server container group

A type of container group that manages containers running game server applications and supporting software. Required for container fleets to host game sessions and can be replicated across fleet instances. See [???](#).

Game server process

A running instance of your game server executable on a hosting resource. Each process can host one game session at a time and communicates directly with Amazon GameLift Servers to report status and receive instructions. See [Add Amazon GameLift Servers to your game server with the server SDK](#).

Game session

An instance of your multiplayer game running on a game server process. Represents a single gameplay experience that multiple players can join and participate in together. See [How hosting with Amazon GameLift Servers works](#) and [GameSession](#) in the API Reference.

Game session data

Information associated with a game session including player data, game state, session properties, and other runtime information used by game servers and clients. See [GameSession](#) in the API Reference.

Game session placement

The process of locating available game servers to host new game sessions. Uses real-time information about game server availability to make optimal placement decisions based on

factors like cost and latency. See [Integrate Amazon GameLift Servers game client functionality](#) and [StartGameSessionPlacement](#) in the API Reference.

Game session protection

A setting that prevents Amazon GameLift Servers from terminating game server processes that are hosting active game sessions during scaling events or fleet updates. Ensures ongoing games are not interrupted. See [Manually set capacity for an Amazon GameLift Servers fleet](#).

Game session queue

A placement mechanism that uses configurable algorithms to place game sessions for optimal outcomes. Can prioritize placements based on factors such as lowest hosting cost and lowest player latency across multiple geographic locations. See [Configure game session placement](#).

Idle instance

A fleet instance that is running but not currently hosting any active game sessions. Used as a metric for monitoring fleet efficiency and scaling decisions. See [Monitor Amazon GameLift Servers with Amazon CloudWatch](#).

Inbound permissions

Network access rules that define which IP addresses and port ranges can be used to connect to fleet instances. Controls access to game servers and ensures secure connections. Both managed EC2 and managed container fleets require inbound permissions configuration.

See [IpPermission](#) and [Configure network connections](#).

Latency policy

Rules in game session queues that define acceptable latency thresholds for player connections. Used to ensure players are placed in game sessions with acceptable network performance. See [Customize a game session queue](#).

Managed fleet

A type of fleet where Amazon GameLift Servers deploys and manages computing resources in the Amazon Web Services Cloud. Provides automatic scaling, health monitoring, and infrastructure management. See [Amazon GameLift Servers managed EC2 fleets](#).

Matchmaking

The process of grouping players together for multiplayer gameplay based on criteria such as skill level, latency, or other game-specific factors. In Amazon GameLift Servers, primarily handled by FlexMatch. See [What is GameLift FlexMatch?](#)

Metric group

A collection of related performance metrics used for monitoring and auto-scaling decisions. Groups metrics like active game sessions, available game servers, and player utilization. See [Auto-scale fleet capacity with Amazon GameLift Servers](#).

Per-instance container group

An optional container group type that runs additional software on each fleet instance. Useful for background services or monitoring tools, with only one copy deployed per instance. See [???](#).

Player session

A connection slot for a player in a game session. Used to reserve game session slots, validate players when they connect, and track player connection status and disconnects. See [Integrate Amazon GameLift Servers game client functionality](#) and [PlayerSession](#) in the API Reference.

Priority configuration

Settings that customize the prioritization order of a game session queue when selecting destinations for game session placement. Defines the overall queue behavior, in contrast to priority override lists which modify priority for individual placement requests. See [Prioritize game session placement](#).

Priority override list

An alternate priority list of locations that can be included in a game session placement request. Effectively replaces the queue's configured prioritization for locations for that one request only, without impacting other requests. See [Prioritize game session placement](#).

Process manager

A component that manages the lifecycle of multiple game server processes within a container or instance. Handles starting, stopping, and monitoring individual game server processes. See [Optimize game server runtime configuration on managed Amazon GameLift Servers](#).

Realtime script

JavaScript code that defines game logic for Amazon GameLift Servers Realtime Servers. Handles player connections, message routing, and game state management without requiring a custom game server. See [What is Amazon GameLift Realtime Servers?](#)

Routing strategy

Configuration that determines how aliases direct traffic to fleets. Can route to a specific fleet or distribute traffic across multiple fleets based on defined rules. See [Abstract an Amazon GameLift Servers fleet designation with an alias](#).

Runtime configuration

Settings that define how game server processes run on fleet instances, including the number of concurrent processes, launch parameters, and resource allocation. See [Optimize game server runtime configuration on managed Amazon GameLift Servers](#).

Scaling

The process of adjusting fleet capacity by adding or removing hosting resources based on player demand. Can be automatic (based on utilization metrics) or manual (administrator-controlled). See [Auto-scale fleet capacity with Amazon GameLift Servers](#).

Scaling policy

Rules that define when and how to automatically adjust fleet capacity. Can be target-based (maintain specific utilization) or rule-based (scale based on metric thresholds). See [Auto-scale fleet capacity with Amazon GameLift Servers](#).

Server SDK

Software development kit that provides the tools and libraries needed to integrate your game server with Amazon GameLift Servers. Enables communication between your game server and the service for session management, health reporting, and status updates. See [Add Amazon GameLift Servers to your game server with the server SDK](#).

Session-based multiplayer

A multiplayer game architecture where gameplay occurs within discrete, time-limited sessions. Each session has a defined start and end, with players joining and leaving as needed. See [How hosting with Amazon GameLift Servers works](#).

Spot instances

A type of Amazon EC2 instance that uses spare Amazon capacity at reduced cost. Spot instances can be interrupted when Amazon needs the capacity back, making them suitable for fault-tolerant game workloads. See [Build a queue for Spot Instances](#).

TLS certificate

A digital certificate used to enable encrypted communication between game clients and servers. Required for TLS-enabled fleets to ensure secure data transmission. See [Create an Amazon GameLift Servers managed EC2 fleet](#).

Getting started with Amazon GameLift Servers

Take advantage of these getting started resources to learn more about the Amazon GameLift Servers service and how you can start developing a custom hosting solution for your session-based multiplayer games.

Before you start

- Create an Amazon Web Services account (or designate an existing one) to use with Amazon GameLift Servers.
- Set up users with permissions for Amazon GameLift Servers and related Amazon services.
- Select an Amazon Web Services Region to work in. For development, choose a Region that's close to your location. You can change Regions at any time.

[Set up an Amazon user account](#)

Choose your path

I want to explore Amazon GameLift Servers quickly

Best for: learning, creating a proof of concept, rapid prototyping

Quick start options:

- [Game server wrapper for Amazon GameLift Servers](#) – See [Tutorial: Quick onboarding with the Amazon GameLift Servers wrapper](#). This tool is the quickest and easiest way to get your game server hosted and running game sessions with Amazon GameLift Servers, with no changes to game code required. The game server wrapper offers basic game session management functionality and streamlined game server deployment. It's ideal for doing a hands-on evaluation of Amazon GameLift Servers using your own game project or a sample project. When you're ready to build a custom game hosting solution, switch to one of the custom development options with full integration with the server SDK for Amazon GameLift Servers. If your game doesn't need a custom hosting solution, you can continue to use the game server wrapper to deploy and host your game servers in production.
- [Amazon GameLift Servers plugin for Unreal Engine or Unity](#) – The plugins give you GUI workflows and sample assets to guide you through the initial steps and deploy your game server

with a basic hosting solution. Use the plugin to set up hosting with self-managed Anywhere fleets, or deploy cloud-based, managed EC2 fleets or container fleets. When you're ready to develop a custom hosting solution, you can build on your plugin-built solutions.

- [Starter kit for Amazon GameLift Servers managed containers](#) – This kit streamlines tasks to integrate a game server, prepare a game server container image, and deploy a container fleet for hosting. For integration, the kit adds essential game session management features to your game server. The kit uses pre-configured templates to build a container fleet and an automated deployment pipeline for the game server. When you're ready to add full game session management features, follow one of the custom development roadmaps to integrate the server SDK for Amazon GameLift Servers.

I want to build a custom hosting solution

Best for: creating a production-level solution for a custom game server

Review the page [Amazon GameLift Servers game hosting options](#) and choose one of the options for your solution. Follow one of these development roadmaps to get started building a full-featured custom hosting solution for your game. The roadmaps provide detailed guidance on how to create, test, and customize each component in your hosting solution.

- [Development roadmap for hosting with Amazon GameLift Servers managed EC2](#)
- [???](#)
- [???](#)
- [???](#)

Recommended learning path

1. Start small. Try the game server wrapper or the plugin for Unreal.
2. Understand the concepts. Review the page [How hosting with Amazon GameLift Servers works](#).
3. Choose architecture. Select your hosting model based on your game's requirements.
4. Build and test. Follow the appropriate development roadmap. Build a basic version of each component and then iterate and customize.
5. Scale and optimize. Optimize your solution for production-level usage. Add advanced features such as matchmaking.

Set up an Amazon user account

Tip

Use these topics to get help with these tasks:

- Get a new Amazon Web Services account for use with Amazon GameLift Servers.
- Create a user or group with permissions to work with Amazon GameLift Servers resources.
- Set up security credentials (you need these to use the Amazon CLI tools and the Amazon GameLift Servers plugins for Unreal and Unity)
-

As with all Amazon services, you need an Amazon Web Services account to use the Amazon GameLift Servers service and tools. An Amazon Web Services account serves two primary functions: (1) it gives you a container for all the Amazon resources that you create with the account; and (2) it lets you manage security for your Amazon resources, including setting up user authentication and controlling user access permissions. There's no cost for creating an Amazon Web Services account.

Explore Amazon GameLift Servers with or without an Amazon account

You **don't** need an Amazon account to:

- Discover Amazon tools for building, running, and growing game experiences at [Amazon for Gaming](#).
- Learn more about Amazon GameLift Servers in the [product overview, FAQs, and resources](#). **Ask Amazon** to find answers to your product questions. (Try this one: "Looking for low-cost options to host my multiplayer game".)
- For a deeper dive, find out what makes Amazon GameLift Servers work in the [technical documentation](#), including developer guides for hosting and matchmaking, and the service API reference guide.
- Check out information on [Amazon GameLift Servers pricing](#) and cost optimization techniques.
- Get downloads and see code repositories for Amazon GameLift Servers SDKs, plugins, and toolkits. See [Amazon GameLift Servers Getting started](#). (You need an Amazon Web Services account to use them.)

You **do** need an Amazon account to:

- Create and manage Amazon resources using the Amazon Web Services Management Console.
- Create and manage Amazon resources using the Amazon Command Line Interface.
- Use Amazon Q with the In the Amazon GameLift Servers technical documentation to find answers, guidance, and recommendations.

Topics

- [Sign up for an Amazon Web Services account](#)
- [Secure IAM users](#)
- [Set user permissions for Amazon GameLift Servers](#)
- [Set up programmatic access for users](#)
- [Set up programmatic access for your game](#)
- [IAM permission examples for Amazon GameLift Servers](#)
- [Set up an IAM service role for Amazon GameLift Servers](#)

Sign up for an Amazon Web Services account

If you do not have an Amazon Web Services account, use the following procedure to create one.

To sign up for Amazon Web Services

1. Open <http://www.amazonaws.cn/> and choose **Sign Up**.
2. Follow the on-screen instructions.

Amazon sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to <http://www.amazonaws.cn/> and choosing **My Account**.

Secure IAM users

After you sign up for an Amazon Web Services account, safeguard your administrative user by turning on multi-factor authentication (MFA). For instructions, see [Enable a virtual MFA device for an IAM user \(console\)](#) in the *IAM User Guide*.

To give other users access to your Amazon Web Services account resources, create IAM users. To secure your IAM users, turn on MFA and only give the IAM users the permissions needed to perform their tasks.

For more information about creating and securing IAM users, see the following topics in the *IAM User Guide*:

- [Creating an IAM user in your Amazon Web Services account](#)
- [Access management for Amazon resources](#)
- [Example IAM identity-based policies](#)

Set user permissions for Amazon GameLift Servers

Create additional users or extend access permissions to existing users as needed for your Amazon GameLift Servers resources. As a best practice ([Security best practices in IAM](#)), apply least-privilege permissions for all users. For guidance on permissions syntax, see [IAM permission examples for Amazon GameLift Servers](#).

Use following instructions to set user permissions based on how you manage the users in your Amazon account.

To provide access, add permissions to your users, groups, or roles:

- Users managed in IAM through an identity provider:

Create a role for identity federation. Follow the instructions in [Create a role for a third-party identity provider \(federation\)](#) in the *IAM User Guide*.

- IAM users:
 - Create a role that your user can assume. Follow the instructions in [Create a role for an IAM user](#) in the *IAM User Guide*.
 - (Not recommended) Attach a policy directly to a user or add a user to a user group. Follow the instructions in [Adding permissions to a user \(console\)](#) in the *IAM User Guide*.

When working with IAM users, as a best practice always attach permissions to roles or user groups, not individual users.

Set up programmatic access for users

Users need programmatic access if they want to interact with Amazon outside of the Amazon Web Services Management Console. The Amazon APIs and the Amazon Command Line Interface require access keys. Whenever possible, create temporary credentials that consist of an access key ID, a secret access key, and a security token that indicates when the credentials expire.

To grant users programmatic access, choose one of the following options.

Which user needs programmatic access?	To	By
IAM	Use short-term credentials to sign programmatic requests to the Amazon CLI or Amazon APIs (directly or by using the Amazon SDKs).	Following the instructions in Using temporary credentials with Amazon resources in the <i>IAM User Guide</i> .
IAM	(Not recommended) Use long-term credentials to sign programmatic requests to the Amazon CLI or Amazon APIs (directly or by using the Amazon SDKs).	Following the instructions in Managing access keys for IAM users in the <i>IAM User Guide</i> .

If you use access keys, see [Best practices for managing Amazon access keys](#).

Set up programmatic access for your game

Most games use backend services to communicate with Amazon GameLift Servers using the Amazon SDKs. Use a backend service (acting for a game client) to request game sessions, place players into games, and other tasks. These services need programmatic access and security credentials to authenticate calls to the service API for Amazon GameLift Servers.

For Amazon GameLift Servers, you manage this access by creating a player user in Amazon Identity and Access Management (IAM). Manage player user permissions through one of the following options:

- Create an IAM role with player user permissions and allow the player user to assume the role when needed. The backend service must include code to assume this role before making requests to Amazon GameLift Servers. In accordance with security best practices, roles provide limited, temporary access. You can use roles for workloads running on Amazon resources ([IAM roles](#)) or outside of Amazon ([IAM Roles Anywhere](#)).
- Create an IAM user group with player user permissions and add your player user to the group. This option gives your player user long-term credentials, which the backend service must store and use when communicating with Amazon GameLift Servers.

For permissions policy syntax, see [Player user permission examples](#).

For more information on managing permissions for use by a workload, see [IAM Identities: Temporary credentials in IAM](#).

IAM permission examples for Amazon GameLift Servers

Use the syntax in these examples to set Amazon Identity and Access Management (IAM) permissions for users that need access to Amazon GameLift Servers resources. For more information on managing user permissions, see [Set user permissions for Amazon GameLift Servers](#). When managing permissions for users outside of the IAM Identity Center, as a best practice always attach permissions to IAM roles or user groups, not individual users.

If you're using Amazon GameLift Servers FleetIQ as a standalone solution, see [Set up your Amazon Web Services account for Amazon GameLift Servers FleetIQ](#).

Administration permission examples

These examples give a hosting administrator or developer targeted access to manage Amazon GameLift Servers game hosting resources.

Example Syntax for Amazon GameLift Servers full access resource permissions

The following example extends full access to all Amazon GameLift Servers resources.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": {
```

```
"Effect": "Allow",
"Action": "gamelift:*",
"Resource": "*"
}
}
```

Example Syntax for Amazon GameLift Servers resource permissions with support for Regions that aren't enabled by default

The following example extends access to all Amazon GameLift Servers resources and Amazon Regions that aren't enabled by default. For more information about Regions that aren't enabled by default and how to enable them, see [Managing Amazon Web Services Regions](#) in the *Amazon Web Services General Reference*.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": {
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeRegions",
      "gamelift:*"
    ],
    "Resource": "*"
  }
}
```

Example Syntax for Amazon GameLift Servers resource to access container images in Amazon ECR

The following example extends access to Amazon Elastic Container Registry (Amazon ECR) actions that Amazon GameLift Servers users need when working with managed container fleets.

JSON

```
{
  "Version": "2012-10-17",
```

```
"Statement": {
  "Effect": "Allow",
  "Action": [
    "ecr:DescribeImages",
    "ecr:BatchGetImage",
    "ecr:GetDownloadUrlForLayer"
  ],
  "Resource": "*"
}
```

Example Syntax for Amazon GameLift Servers resource and PassRole permissions

The following example extends access to all Amazon GameLift Servers resources and allows a user to pass an IAM service role to Amazon GameLift Servers. A service role gives Amazon GameLift Servers limited ability to access other resources and services on your behalf, as is described in [Set up an IAM service role for Amazon GameLift Servers](#). For example, when responding to a CreateBuild request, Amazon GameLift Servers needs access to your build files in an Amazon S3 bucket. For more information about the PassRole action, see [IAM: Pass an IAM role to a specific Amazon service](#) in the *IAM User Guide*.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "gamelift:*",
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": "iam:PassRole",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "iam:PassedToService": "gamelift.amazonaws.com"
        }
      }
    }
  ]
}
```

```
]
}
```

Player user permission examples

These examples allow a backend service or other entity to make API calls to the Amazon GameLift Servers API. They cover the common scenarios for managing game sessions, player sessions, and matchmaking. For more details, see [Set up programmatic access for your game](#).

Example Syntax for game session placement permissions

The following example extends access to the Amazon GameLift Servers APIs that use game session placement queues to create game sessions and manage player sessions.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": {
    "Sid": "PlayerPermissionsForGameSessionPlacements",
    "Effect": "Allow",
    "Action": [
      "gamelift:StartGameSessionPlacement",
      "gamelift:DescribeGameSessionPlacement",
      "gamelift:StopGameSessionPlacement",
      "gamelift:CreatePlayerSession",
      "gamelift:CreatePlayerSessions",
      "gamelift:DescribeGameSessions"
    ],
    "Resource": "*"
  }
}
```

Example Syntax for matchmaking permissions

The following example extends access to the Amazon GameLift Servers APIs that manage FlexMatch matchmaking activities. FlexMatch matches players for new or existing game sessions and initiates game session placement for games hosted on Amazon GameLift Servers. For more information about FlexMatch, see [What is Amazon GameLift Servers FlexMatch?](#)

JSON

```
{
  "Version": "2012-10-17",
  "Statement": {
    "Sid": "PlayerPermissionsForGameSessionMatchmaking",
    "Effect": "Allow",
    "Action": [
      "gamelift:StartMatchmaking",
      "gamelift:DescribeMatchmaking",
      "gamelift:StopMatchmaking",
      "gamelift:AcceptMatch",
      "gamelift:StartMatchBackfill",
      "gamelift:DescribeGameSessions"
    ],
    "Resource": "*"
  }
}
```

Example Syntax for manual game session placement permissions

The following example extends access to the Amazon GameLift Servers APIs that manually create game sessions and player sessions on specified fleets. This scenario supports games that don't use placement queues, such as games that let players join by choosing from a list of available game sessions (the "list-and-pick" method).

JSON

```
{
  "Version": "2012-10-17",
  "Statement": {
    "Sid": "PlayerPermissionsForManualGameSessions",
    "Effect": "Allow",
    "Action": [
      "gamelift:CreateGameSession",
      "gamelift:DescribeGameSessions",
      "gamelift:SearchGameSessions",
      "gamelift:CreatePlayerSession",
      "gamelift:CreatePlayerSessions",
      "gamelift:DescribePlayerSessions"
    ]
  }
}
```

```
    ],  
    "Resource": "*"    
  }  
}
```

Set up an IAM service role for Amazon GameLift Servers

Some Amazon GameLift Servers features require you to extend limited access to other Amazon resources that you own. You can do this by creating an Amazon Identity and Access Management (IAM) role. An [IAM role](#) is an IAM identity that you can create in your account that has specific permissions. An IAM role is similar to an IAM user in that it is an Amazon identity with permissions policies that determine what the identity can and cannot do in Amazon. However, instead of being uniquely associated with one person, a role is intended to be assumable by anyone who needs it. Also, a role does not have standard long-term credentials such as a password or access keys associated with it. Instead, when you assume a role, it provides you with temporary security credentials for your role session.

This topic covers how to create a role that you can use with your Amazon GameLift Servers managed fleets. If you use Amazon GameLift Servers FleetIQ to optimize game hosting on your Amazon Elastic Compute Cloud (Amazon EC2) instances, see [Set up your Amazon Web Services account for Amazon GameLift Servers FleetIQ](#).

In the following procedure, create a role with a custom permissions policy and a trust policy that allows Amazon GameLift Servers to assume the role.

Create an IAM service role for an Amazon GameLift Servers managed EC2 fleet

Step 1: Create a permissions policy.

Use the instructions and examples on this page to create a custom permissions policy for the type of Amazon GameLift Servers fleet you're working with.

To use the JSON policy editor to create a policy

1. Sign in to the Amazon Web Services Management Console and open the IAM console at <https://console.amazonaws.cn/iam/>.
2. In the navigation pane on the left, choose **Policies**.

If this is your first time choosing **Policies**, the **Welcome to Managed Policies** page appears. Choose **Get Started**.

3. At the top of the page, choose **Create policy**.
4. In the **Policy editor** section, choose the **JSON** option.
5. Enter or paste a JSON policy document. For details about the IAM policy language, see [IAM JSON policy reference](#).
6. Resolve any security warnings, errors, or general warnings generated during [policy validation](#), and then choose **Next**.

 Note

You can switch between the **Visual** and **JSON** editor options anytime. However, if you make changes or choose **Next** in the **Visual** editor, IAM might restructure your policy to optimize it for the visual editor. For more information, see [Policy restructuring](#) in the *IAM User Guide*.

7. (Optional) When you create or edit a policy in the Amazon Web Services Management Console, you can generate a JSON or YAML policy template that you can use in Amazon CloudFormation templates.

To do this, in the **Policy editor** choose **Actions**, and then choose **Generate CloudFormation template**. To learn more about Amazon CloudFormation, see [Amazon Identity and Access Management resource type reference](#) in the *Amazon CloudFormation User Guide*.

8. When you are finished adding permissions to the policy, choose **Next**.
9. On the **Review and create** page, enter a **Policy name** and a **Description** (optional) for the policy that you are creating. Review **Permissions defined in this policy** to see the permissions that are granted by your policy.
10. (Optional) Add metadata to the policy by attaching tags as key-value pairs. For more information about using tags in IAM, see [Tags for Amazon Identity and Access Management resources](#) in the *IAM User Guide*.
11. Choose **Create policy** to save your new policy.

Step 2: Create a role that Amazon GameLift Servers can assume.

To create an IAM role

1. In the navigation pane of the IAM console, choose **Roles**, and then choose **Create role**.
2. On the **Select trusted entity** page, choose the **Custom trust policy** option. This selection opens the **Custom trust policy** editor.
3. Replace the default JSON syntax with the following, and then choose **Next** to continue.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "gamelift.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

4. On the **Add permissions** page, locate and select the permissions policy that you created in Step 1. Choose **Next** to continue.
5. On the **Name, review and create** page, enter a **Role name** and a **Description** (optional) for the role that you are creating. Review the **Trust entities** and **Added permissions**.
6. Choose **Create role** to save your new role.

Create an IAM role for Amazon GameLift Servers managed containers

If you're using Amazon GameLift Servers managed containers, you need to create an IAM service role for use with a container fleet. This role grants limited permissions that Amazon GameLift Servers needs to manage your container fleet resources and take actions on your behalf.

To create an IAM role for a container fleet

1. Sign in to the Amazon Web Services Management Console and open the IAM console at <https://console.amazonaws.cn/iam/>.

2. In the navigation pane of the IAM console, choose **Roles**, and then choose **Create role**.
3. On **Select trusted entity** page, choose **Amazon service** and select the **Use case** "GameLift". Choose **Next**
4. On **Add permissions**, choose the managed policy `GameLiftContainerFleetPolicy`. Choose **Next**. See [Amazon managed policies for Amazon GameLift Servers](#) for more information about this policy.
5. On **Name, review, and create**, enter a role name and choose **Create role** to save the new role.

Permission policy syntax

- **Permissions for Amazon GameLift Servers to assume the service role**

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "gamelift.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

- **Permissions to access Amazon Regions that aren't enabled by default**

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "gamelift.amazonaws.com",
          "gamelift.ap-east-1.amazonaws.com",

```

```
        "gamelift.me-south-1.amazonaws.com",  
        "gamelift.af-south-1.amazonaws.com",  
        "gamelift.eu-south-1.amazonaws.com"  
    ],  
    },  
    "Action": "sts:AssumeRole"  
}  
]  
}
```

Get Amazon GameLift Servers development tools

Amazon GameLift Servers provides a set of SDKs and other tools to help you build game hosting solutions for your games. The SDKs add functionality to game servers, game clients, and backend services that enables them to interact with the Amazon GameLift Servers service. For the latest information about Amazon GameLift Servers SDK versions and compatibility, see [Amazon GameLift Servers release notes](#).

For game servers

Integrate and build your 64-bit game servers with the server SDK for Amazon GameLift Servers. The game server uses the server SDK to communicate with the Amazon GameLift Servers service for game session management, including starting, updating, and stopping game sessions. For help with integrating the server SDK into your game projects, see [Prepare a game for hosting with Amazon GameLift Servers](#).

Development support

- **Development OS**
 - Windows
 - Linux
- **Programming languages**

[Get the Amazon GameLift Servers SDK](#). For version-specific information and install instructions, see the included readme files in each package.

- [C++ server SDK](#)
 - [Server SDK reference](#)
 - [How to integrate](#)

- [C# server SDK](#) (Support for .NET 4, .NET 6, .NET 8 varies by version, see [SDK versions](#))
 - [Server SDK reference](#)
 - [How to integrate](#)
- [Go server SDK](#)
 - [Server SDK reference](#)
 - [How to integrate](#)

Runtime support

For a managed hosting solution, build your game server to run on one of the following Amazon machine images (AMIs). See [Amazon GameLift Servers AMI versions](#) for Amazon GameLift Servers for more details.

- [Windows Server 2016](#)
- [Amazon Linux 2](#)

Note

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

For game client services

Create a 64-bit backend service for your game and integrate it with functionality from the Amazon SDK, which includes the service API for Amazon GameLift Servers. Use the backend service to handle client-side interactions with the Amazon GameLift Servers service, including starting or finding game sessions and joining players to games.

[Get the Amazon SDK](#)

For more information about using the Amazon SDK with Amazon GameLift Servers, see the following resources:

- [Amazon GameLift Servers API Reference](#)

- Integrate game client functionality
 - [Client-side service integration for games with custom game server builds](#)
 - [Client-side service integration for games with Realtime servers](#)
- [Design a client backend service](#)

For Amazon GameLift Servers resource management

Use the following tools to create, update, and monitor your Amazon GameLift Servers managed hosting resources.

- [Amazon Web Services Management Console](#) – The Amazon Console is a web-based application that provides centralized access to all individual Amazon service consoles, including Amazon GameLift Servers. Use the Console to create or sign into an Amazon account and open the Amazon GameLift Servers console to work with your game hosting resources. You can configure and deploy hosting fleets and other resources, view usage and performance metrics, track resources in the dashboard, and many other tasks. [Go to the Amazon GameLift Servers console.](#)
- [Service API for Amazon GameLift Servers](#) – This API gives you programmatic access to all of your Amazon GameLift Servers resources. It is part of the Amazon SDK, which you can download for use with most popular programming languages. [Get the Amazon SDK.](#)
- [Amazon command line interface \(CLI\)](#) – The Amazon CLI lets you interact with Amazon services using a command-line shell. The tools provides direct access to the public APIs for Amazon services as well as customized commands that are available for a service. [Get the Amazon CLI.](#)
- [Amazon CloudFormation](#) for Amazon GameLift Servers – The Amazon CloudFormation service helps you model and set up Amazon resources to streamline infrastructure deployment and management. Create an Amazon CloudFormation template to describe the Amazon GameLift Servers resources for your hosting solution, and then use the template to build additional resources or update configurations. View the [Amazon GameLift Servers resource type reference.](#)

For Amazon GameLift Servers Realtime

Configure and deploy Realtime servers to host your multiplayer games. To allow your game clients to connect to Realtime servers, use the Amazon GameLift Servers Realtime client SDK. To get started, [download the Realtime client SDK](#). For configuration information, see [Integrating a game client for Amazon GameLift Servers Realtime](#).

SDK support

The Realtime Client SDK contains source for the following languages:

- C# (.NET)

Development environments

Build the SDK from source as needed for the following supported development operating systems and game engines:

- **Operating systems** – Windows, Linux, Android, iOS
- **Game engines** – Unity, engines that support C# libraries

Game server operating systems

You can deploy Realtime servers onto hosting resources that run on the following platforms:

- [Amazon Linux 2](#)

Note

AL2 is nearing end of support. See more details in the [Amazon Linux 2 FAQs](#).

- [Amazon Linux 2023](#)

Tutorial: Quick onboarding with the Amazon GameLift Servers wrapper

Welcome to the onboarding tutorial for Amazon GameLift Servers. In this tutorial, you rapidly deploy your game server to be hosted on a fleet of cloud-based compute resources. Use this tutorial to skip the work of integrating the server SDK for Amazon GameLift Servers into your game code, and instead deploy your game with the minimal functionality to communicate with the Amazon GameLift Servers service and run game sessions. You will set up a basic hosting solution and use it to experience the full range of features such as automatic scaling and matchmaking. It's also a great way to host a prototype of your game as part of a live demo or for testing.

Key benefits of this onboarding method:

- Quickly deploy your game server for hosting fast.

- Zero game code changes and no modifications required.
- Use this method with any game executable, regardless of game engine.
- Explore all Amazon GameLift Servers management tools, including monitoring for game session activity and hosting health.

Note

The wrapper is intended for evaluation and basic production use. Advanced features such as detailed player session management require a full server SDK integration.

Prerequisites

Before you begin, ensure you have:

- An Amazon Web Services account with appropriate permissions
- Amazon CLI installed
- Go 1.18+
- A multiplayer game server executable
- Make (Linux/Mac)
- Git installed with an active account

Overview

In this tutorial you will:

1. Get and build the wrapper
2. Prepare the game build
3. Configure the wrapper
4. Upload the game server build
5. Create the managed EC2 fleet
6. Create and connect to a game session
7. Monitor and manage your game servers
8. Scale your game servers

Step 1: Get and build the game server wrapper

Use the following commands to get the game server wrapper source and build the wrapper. These commands use SSH, but you can also go to the Github repository directly.

Windows

```
> git clone git@github.com:amazon-gamelift/amazon-gamelift-servers-game-server-wrapper.git
> cd amazon-gamelift-servers-game-server-wrapper
> powershell -file .\build.ps1
```

Mac and Linux

```
$ git clone git@github.com:amazon-gamelift/amazon-gamelift-servers-game-server-wrapper.git
$ cd amazon-gamelift-servers-game-server-wrapper
$ make
```

On a successful build, an "out" directory is added to `amazon-gamelift-servers-game-server-wrapper`. In this directory are three folders, one for each supported hosting fleet option, that contain a set of build artifacts. For this tutorial, you're deploying to managed EC2 fleets, so you'll use the folder `gamelift-servers-managed-ec2`.

Step 2: Prepare your game server build

In this step, you prepare your game server build files for uploading to the Amazon GameLift Servers.

Create the game directory

Now, prepare a game directory on your local machine. This directory must contain all the files you need to run your game server with Amazon GameLift Servers. This includes the game server wrapper, your game server build and the `config.yaml` that makes the wrapper work with your game server.

Use the following steps:

1. **Managed EC2 fleet.** In the game server wrapper output folder, find the build artifact for deploying to a Managed EC2 fleet. The build will be written to an out directory as indicated here: `out\linux\amd64\gamelift-servers-managed-ec2`.
2. Copy your game server executable and all associated files that it requires to run into the `gamelift-servers-managed-ec2` folder. You can have nested directories as needed.

An example directory structure will look like this:

```
gamelift-servers-managed-ec2
#-- config.yaml
#-- amazon-gamelift-servers-game-server-wrapper
#-- MyGame
#   #-- server-executable.exe
#   #-- my-game-settings
#   # .....
```

Step 3: Configure the wrapper for your fleet

Amazon GameLift Servers manages the lifecycle of a fleet's compute instances, spinning up new instances with your server build installed and recycling instances as needed. The service manages the game server process lifecycle that runs on each instance. A Managed EC2 fleet can have instances in multiple locations to support players wherever they are.

Edit the `config.yaml` file to configure the wrapper for logging, port setup, and server initialization.

1. **Configure the logging settings.** The game server wrapper generates log messages for each game server process. By default, the log level is set to debug for maximum verbosity. This is really useful during setup and troubleshooting and determines how detailed the log messages are - in this case, the most verbose. Options include debug, info, warn, and error (least verbose).
2. **Specify the path to the game server log directory.** The default path for game server logs is `./game-server-logs`. This directory contains all the logs generated by your game server and each instance has them. The logs are automatically uploaded to Amazon GameLift Servers where they are accessed from the **Events** tab. View the Troubleshooting section for more details.
3. **Define network port configuration.** Set the game port to whatever you prefer. For this tutorial, specify one port only, as you'll create a fleet that runs only one concurrent game server process per instance. If you decide to run multiple processes at a time, you'll need to configure enough

ports for each concurrent process. The default value is 37016 as shown in the config file, but in general, for fleets using Linux builds use ports 22 and 1026-60000. For fleets using Windows builds, use ports 1026-60000.

4. **Set up the path to the game server executable.** For `./MyGame/my-server-executable.exe` customize the path to your game server executable with the actual name and location for it. This is the entry point for launching your game server.
5. **Configure the game server arguments.** At a minimum, specify a `-port` argument and use the same game port value you defined before. The "pos" value 0 indicates this is the first argument. Add other arguments as needed. These arguments are passed to your game server when it's launched, allowing you to configure its runtime behavior.
 - a. Argument: `--port`
 - b. Value: `"{{.port number here}}"`
 - c. Position: 0 (First argument in the list)

Example configuration:

```
log-config:
  wrapper-log-level: debug
game-server-logs-dir: ./game-server-logs
ports:
  gamePort: 37016
game-server-details:
  executable-file-path: ./MyGame/my-server-executable
game-server-args:
  - arg: "--port"
    val: "{{.gameport}}"
    pos: 0
```

Step 4: Upload the game server build

You've now completed all the required elements of your game server build (game server wrapper, config.yaml, and your game server files), and you're ready to upload your game build to Amazon GameLift Servers for hosting. The quickest way to upload your game build is by using the Amazon CLI command `upload-build` as shown in the following example.

Uploading a game build with Windows:

```
aws gamelift upload-build \  
  --name gamelift-test-2025-03-11-1 \  
  --build-version gamelift-test-2025-03-11-1 \  
  --build-root out/windows/amd64/gamelift-servers-managed-ec2 \  
  --operating-system WINDOWS_2022 \  
  --server-sdk-version 5.3.0 \  
  --region us-west-2
```

 Note

For Mac and Linux builds, use `--operating-system AMAZON_LINUX_2023`

When you create the build, record the build ID from the API response to use it for fleet creation.

Step 5: Create the managed EC2 fleet

The following steps describe a minimal fleet configuration, so you can be up and running as soon as possible.

To create your fleet:

1. Log into the Amazon Web Services Management Console and navigate to Amazon GameLift Servers.
2. On the menu bar at the top of the console window, check to make sure what region your build is in. Make note of it because your fleet needs to be in the same region or you won't be able to find or choose your build.
3. In the navigation pane in the **Managed EC2** section, choose **Builds**.
4. Select the build you uploaded previously to display the **Build details page**.
5. In the Fleets section choose **Create Fleet** which displays the Define Managed EC2 fleet details page from which you are able to track status for your fleet as well as view fleet creation events on the **Events** tab.
6. Fill out the name and description and choose **Next**.
7. In the Define instance details page the build's region is shown by default. Choose any additional regions you want to add.

8. For **Fleet type** choose **On-Demand**.

9. Under **Instance types** choose **c5.large** and choose **Next**.

10 Under **Runtime configuration**, Because your uploaded game build uses the wrapper, you need to point to the wrapper executable instead. For Windows game servers this is `C:\game\amazon-gamelift-servers-game-server-wrapper.exe`. For Linux game servers, this is `/local/game/amazon-gamelift-servers-game-server-wrapper`.

For example: `LaunchPath": "/local/game/amazon-gamelift-servers-game-server-wrapper", "ConcurrentExecutions": 1, "Parameters": "-port 37016`

Also, configure the game port values with a range that allows for the port that was set in the `config.yaml` and the runtime config launch parameters. The port in `config.yaml` does not need to match the one specified in the fleet's runtime configuration, but at runtime if the configuration specifies a different port, that value overrides what is in the `config.yaml`. Launch parameters entered in the runtime config also override what is in the `config.yaml`.

11 On the **Review and create** page, double-check all configurations, then choose **Submit** to create your fleet. The fleet status will change as it spins up capacity to host the game server, and very soon shows a status of **Active**. Once activation is complete and the fleet deployed, the service launches the wrapper which is ready to receive a game session request.

Step 6: Create and connect to a game session

When your fleet status shows **Active**, this means that game servers are ready and waiting to host a game session. To start a game session, make a game session request to the Amazon GameLift Servers service. Here, you will use the Amazon CLI to make this request.

Note

Keep in mind that creating a game session using the Amazon CLI is useful for testing and becoming familiar with the process. At some point, you'll add programmatic Amazon SDK calls to to your game backend service as part of your game's matchmaking or game session placement system.

Use the following to create a game session:

```
aws gamelift create-game-session \  
--fleet-id <FLEET_ID> \  
--game-properties '[{"Key": "exampleProperty", "Value": "exampleValue"}]' \  
--maximum-player-session-count 3 \  
--region us-west-2
```

You can also pass customized game properties to your server executable. See the Game Server Arguments in the README for details. When it receives the create-game-session call, the Amazon GameLift Servers inform the wrapper to launch the game server executable and start a game session. What is in the `config.yaml` affects the game server's configuration, and the launch parameters that are set in the console determine the configuration of the game session itself.

Example format for adding game properties:

```
defaultArgs:  
  - arg: "--port"  
    val: "{{.GamePort}}"  
    pos: 0  
  - arg: "--ipAddress"  
    val: "{{.IpAddress}}"  
    pos: 1  
  - arg: "--gameSessionId"  
    val: "{{.GameSessionId}}"  
    pos: 2
```

Step 7: Manage and monitor your fleet

Now that your game server fleet is set up and a game session started, you can manage and monitor it from the Amazon GameLift Servers console. The best way to accomplish this is from the Fleet details page where you can edit fleet details or change fleet scaling and capacity on the **Scaling** tab. See the following section about scaling your game servers.

Choose the **Metrics** tab to view graphs depicting monitoring for Activity and Hardware metrics. For details about the metrics graphs, choose the **Info** link next to your fleet's ID. Additionally, while you can closely monitor your game servers from the **Metrics** tab, you can also add alarms to these metrics in the CloudWatch dashboard.

To navigate to the CloudWatch dashboard from the console:

1. Type "CloudWatch" in the search bar and choose it from the search results list to display the CloudWatch Overview.
2. Scroll down and choose **View GameLift dashboard** to see graphs based on key metrics for your fleets and game sessions.

Step 8: Scale your game servers

For the next step, you set up auto scaling. With auto scaling you dynamically scale your fleet capacity in response to game server activity. As players arrive and start game sessions, auto scaling adds more instances, and as player demand decreases, auto scaling releases unused instances. This is an effective way to minimize your hosting resources and costs while still providing a smooth, fast player experience.

As you prepare for game launch, you'll want to set up auto scaling for your fleets. Auto scaling is recommended as an effective way to minimize your hosting resources and costs while still providing a smooth, fast player experience.

To manually set up fleet capacity

To manually set up fleet capacity

1. Go to the **Scaling** tab in the details page for your fleet.
2. Select a location and choose **Edit**.
3. Change the **Desired instances** value and change the **Min** and **Max** size settings in order to scale beyond their current values, and choose **Confirm**.

Note

Use the max instances setting as a stop gap to prevent excessive scaling and costs.

To use target-based auto scaling

To use target-based auto scaling

Target-based autoscaling (target tracking) links fleet scaling to the percentage of available game sessions. As players surge to play your game and available game sessions decrease, it responds by automatically adding more instances to the fleet.

1. Under **Target-based auto scaling policy**, choose **Add policy** and set the fleet's capacity to change automatically when it reaches the threshold of the percentage of available game sessions that you set. A larger buffer can better handle surges, getting new players get into games fast, but it may also lead to higher hosting costs.
2. Choose **Confirm** to accept the changes.

Rule-based auto scaling gives you more granular control, such as the ability to link scaling to other fleet metrics and set custom thresholds and scaling responses. It offers powerful options, but it also requires using the CLI and significant testing to understand how custom rules behave in action. This tutorial focuses on first setting up the target-based approach.

Troubleshoot common issues

The following are common issues you may encounter while getting your game servers and game sessions running. When either the server or the game session isn't running properly, the first step is to check the logs which may reveal one of the issues described below for new deployments or games in production.

The following is often revealed in the logs:

- Game server process can't launch. This might be an error in the wrapper config—verify the file has the right launch path and correct launch parameters and arguments.
- Game server build isn't runnable. This is likely an error in the game code.
- Players can't connect to game sessions. This is likely a port configuration error.
- Lag or slow connections. Review scaling policies and thresholds.
- No connection. Verify the port rules and configuration for your fleet.

To view the event logs for your Amazon GameLift Servers fleet

To view the event logs for your Amazon GameLift Servers fleet

1. Open the Amazon GameLift Servers console.
2. On the Fleet Details page choose the **Events** tab and download the log. You can also monitor activity and hardware metrics for game server health and game session activations from the **Metrics** tab.

To view game session logs

To view game session logs

1. From the console, open your fleet and open the **Game sessions** tab.
2. Choose a game session ID from the list to display the Overview page for it.
3. Choose **Download logs** to download a log file locally.

To view game session logs with the CLI use the `GetGameSessionLogURL` API. Amazon GameLift Servers automatically stores the logs for 14 days.

You can also configure Amazon CloudWatch logs for your fleet. This provides additional logging capabilities and integration with other Amazon monitoring services.

For real-time log access or extended retention periods through CloudWatch:

1. At the top of your Amazon GameLift Servers console dashboard type "CloudWatch" in the search bar and choose it from the results drop-down.
2. Go to the CloudWatch Log Groups and search for specific sessions. The easiest method is to click **Search All** and filter using `gameSessionId` or `clientId`.

Next steps

- [Create a multi-location fleet to add hosting in additional locations](#)
- [Add a game session queue to make the best possible placements for game sessions across multiple locations](#)
- [Experiment with FlexMatch matchmaking by creating a matchmaker and rule set for your game](#)
- [Start working on functionality for your game client and backend service components, so players can make join requests and directly connect to game sessions](#)
- [When ready, move to a fully-integrated solution](#)

Amazon GameLift Servers examples

If you're considering using Amazon GameLift Servers to manage your custom game server, or you're interested in taking advantage of Amazon GameLift Servers Realtime, we recommend that you try the following examples before you use the service for your own game. The custom game

server example gives you experience with game hosting in the Amazon GameLift Servers console. The Amazon GameLift Servers Realtime example shows you how to prepare a game for hosting using Realtime servers.

Custom game server example

This example demonstrates the process of deploying a sample game server to Amazon GameLift Servers managed EC2 fleet for hosting. Use the sample game client to connect to a live game session. You can experience how to use Amazon GameLift Servers .tools, including the console and the Amazon CLI, to monitor the fleet's hosting performance and usage.

The example walks you through the following steps:

- Upload the sample game server build.
- Create a fleet to run the game server build.
- Get the sample game client and use it to connect to a game server and join a game session.
- Review fleet and game session metrics.

Start up multiple game clients and play the game to generate hosting data. Use the Amazon GameLift Servers console to view hosting resources, track metrics, and explore options for scaling the fleet's hosting capacity.

To get started, sign in to the [Amazon GameLift Servers console](#). In the left-side navigation, go to **Resources, Try a sample game**.

Building a game hosting solution with Amazon GameLift Servers

This section includes guidance on how to build all the components of an Amazon GameLift Servers game hosting solution, including design best practices. Whether you're starting from scratch or migrating an existing game, you'll find comprehensive information to help you architect, implement, and optimize your game hosting infrastructure using Amazon GameLift Servers.

Hosting solution components

Your hosting solution must have some version of the following components. For more detail on these components and how they work together, see [How hosting with Amazon GameLift Servers works](#).

- **Game client** runs on player devices.
- **Backend service** coordinates communication between game clients and Amazon GameLift Servers.
- **Game server** software manages multiplayer gameplay.
- **Placement system** matches players to game servers.
- **Game hosting fleets** provide computing resources in geographic locations.
- **Game hosting management system** monitors game hosting status and manages capacity.

Topics

- [Development roadmaps](#)
- [Prepare a game for hosting with Amazon GameLift Servers](#)
- [Build a backend service for Amazon GameLift Servers](#)
- [Deploy hosting fleets for Amazon GameLift Servers](#)
- [Configure game session placement](#)
- [Customize to your game hosting solution](#)

Development roadmaps

Development roadmaps provide structured, step-by-step guidance for building different types of game hosting solutions with Amazon GameLift Servers. Each roadmap is tailored to a specific

hosting approach and walks you through the essential tasks, from initial setup to deployment and testing.

Choose the roadmap that best matches your hosting requirements and technical constraints. You can also combine approaches by following multiple roadmaps to create a hybrid solution that meets your specific needs.

Topics

- [Development roadmap for hosting with Amazon GameLift Servers managed EC2](#)

Development roadmap for hosting with Amazon GameLift Servers managed EC2

This roadmap guides you through how to develop an Amazon GameLift Servers managed EC2 hosting solution for your multiplayer game. Amazon GameLift Servers offers several game hosting options; for more information on these options, see [Amazon GameLift Servers game hosting options](#).

With Amazon GameLift Servers managed hosting, your game server is hosted on Amazon Web Services Cloud-based virtual computing resources that Amazon GameLift Servers owns and operates based on your configuration. You get the security, reliability, and global availability of Amazon Elastic Compute Cloud (Amazon EC2) instances that are further optimized for use with multiplayer game hosting. Amazon GameLift Servers streamlines hosting management with tools such as automatic server deployments, life cycle handling, and capacity auto scaling.

An Amazon GameLift Servers managed solution is composed of the following components:

- A game server build, integrated with the server SDK for Amazon GameLift Servers, to deploy across all fleets.
- A game client and backend service, integrated with the Amazon SDK, to interact with the Amazon GameLift Servers service and request game sessions.
- An Amazon GameLift Servers queue or other placement mechanism to place new game sessions with available game servers across all fleets.
- (Optional) A FlexMatch matchmaker to create multi-player matches and set up game sessions for them.
- One or more Amazon GameLift Servers managed fleets, which use Amazon Elastic Compute Cloud (Amazon EC2) instances optimized for multiplayer game hosting.

This roadmap presents a streamlined path to getting your multiplayer game up and running successfully with Amazon GameLift Servers managed EC2 hosting. After you have the necessary components in place, you can continue to iterate on game development and customize your hosting solution. As you get closer to launch, see these [Prepare for launch with Amazon GameLift Servers hosting](#) for help with preparing your hosting solution for production-level usage.

** Get a jump start with the Amazon GameLift Servers plugin for Unreal Engine and Unity**

For faster deployment, try the [Amazon GameLift Servers plugin](#) for Unreal Engine and Unity. It provides guided UI workflows to quickly deploy your game server with minimal setup, so you can try out your game components in action. Then you can build on this foundation to create a custom hosting solution for your game. For more details, see [Prepare your Unreal or Unity game with the Amazon GameLift Servers plugin](#).

Step 1: Prepare your game server to work with Amazon GameLift Servers

Add functionality to your game server so that it can communicate with the Amazon GameLift Servers service when it's deployed for hosting.

- **Get the server SDK for Amazon GameLift Servers for your game project.** The server SDK is available in C++, C#, and Go. [Download an Amazon GameLift Servers server SDK](#).
- **Modify your game server code to add server SDK functionality.** For guidance, see [Prepare a game for hosting with Amazon GameLift Servers](#). At a minimum, do the following:
 - Add code to initialize the Amazon GameLift Servers SDK and establish a WebSocket connection with the Amazon GameLift Servers service. Use the server SDK action `InitSdk()`.
 - Add code to report to the Amazon GameLift Servers service when the server process is ready to host game sessions. Use the server SDK action `ProcessReady()`.
 - Implement the required callback functions `OnProcessTerminate()`, and `OnStartGameSession()`. With these functions, game server processes can maintain a connection with the Amazon GameLift Servers service, initiate a game session when prompted by Amazon GameLift Servers, and respond to a prompt to end the game server process.
 - Add code to report to the Amazon GameLift Servers service when the server process is ending a game session. Use the server SDK action `ProcessEnding()`.
- **Package your game server build.** Create an install script with your build files, dependencies and supporting software. See [Package your game build files](#). We recommend using an Amazon Simple Storage Service (Amazon S3) bucket to store versions of your game build.

- **Test your game server integration.** See [???](#).

Step 2: Prepare your game client to join hosted game sessions

Create a way for your game client to request to join a game session, get connection info, and then connect directly to a hosted game session. The most common approach is to set up backend service functionality that serves as a middleman between your game client and the Amazon GameLift Servers service. This approach protects your hosting resources and gives you greater control over how players are placed into game sessions.

- **Build backend service functionality for hosting.** The backend service communicates with the Amazon GameLift Servers service and delivers connection information to a game client. This functionality includes starting game sessions, placing players into games, and retrieving game session information. For guidance, see [Prepare a game for hosting with Amazon GameLift Servers](#). At a minimum, do the following:
 - Get the Amazon SDK for Amazon GameLift Servers and add it to your backend service project. See [Amazon GameLift Servers SDK resources for client services](#).
 - Add code to initialize an Amazon GameLift Servers client and store key settings. See [Set up the Amazon GameLift Servers API](#).
 - Add functionality to call the Amazon SDK action `CreateGameSession()` and provide game session connection information to a game client. See [Create a game session on a specific fleet](#).

Calling `CreateGameSession()` is a convenient starting point for requesting new game sessions. After you have a game session placement system in place (see Step 3), you'll replace this code with a call to `StartGameSessionPlacement()` (or `StartMatchmaking()` if you're using FlexMatch).

For guidance on designing your backend service, see [Build a backend service for Amazon GameLift Servers](#).

- **Add functionality to your game client that lets players join a hosted game session.** The game client makes requests to your backend service, not directly to Amazon GameLift Servers. After the backend service provides game session connection information, the game client connects directly with the game session to play the game.
- **Test your game client integration.** See [???](#).

Step 3: Set up game session placement

Customize how you want Amazon GameLift Servers to process requests for new game session and locate available game servers to host them. Amazon GameLift Servers automatically tracks the availability of all game servers on all fleets. When a game client sends a request to join a game session, Amazon GameLift Servers looks for the "best possible" placement based on a set of defined priorities such as minimum latency, cost, and availability.

- **Create a game session queue for placing new game session with available game servers.** Queues are the primary mechanism for game session placement. For guidance, see [Create a game session queue](#).
- **In your backend service code, convert the `CreateGameSession()` call to `StartGameSessionPlacement()`.** See [Create a game session in a multi-location queue](#).
- **Create a mechanism to notify a game client when a game session is ready to join.** While in development, you can poll for game session status using a call to `DescribeGameSessionPlacement`. Before using a queue to process high volumes, however, you'll need to enable event notifications. See [Set up event notification for game session placement](#).
- **(Optional) Add FlexMatch matchmaking components.** For guidance, see the [Amazon GameLift Servers FlexMatch developer guide](#).

Step 4: Create managed cloud-based fleets

The final piece of your solution is to set up the type of hosting resources you'll need for a production system. To start planning and configuring for production, you want to transition to working with an Amazon GameLift Servers managed fleet.

- **Package your game server build and upload to Amazon GameLift Servers.** Create an install script with your build files, dependencies and supporting software. See [Create a game server build for Amazon GameLift Servers](#). You can upload your build to Amazon GameLift Servers using either the console or the Amazon CLI.

Before uploading your build, decide in what Amazon Web Services Region you want to create a fleet. You must upload the build to the same Region. For more on choosing a fleet location, see [Geographic locations](#).

- **Create a managed EC2 fleet.** When you create a fleet, Amazon GameLift Servers immediately begins deploying your game server build for hosting. You can configure many aspects of a

managed fleet. For guidance, see [Create an Amazon GameLift Servers managed EC2 fleet](#). At minimum, do the following:

- Give the fleet a name and specify which uploaded game build to deploy.
- Choose On-Demand Instances for your fleet and select an instance type that's available in your fleet location. Spot fleets are a valuable option but require additional design and configuration.
- Create a runtime configuration for the fleet. At minimum, specify the launch path for your game server executable.
- Specify port settings to allow inbound traffic to access your game servers.
- **Add the managed fleets to your queue.** In your game session queue, add your managed fleets.
- **Test game hosting with your managed fleets.** At this point you should be able to test the entire hosting cycle, with a game client requesting a game session, getting connection info, and successfully connecting to a game session.

Step 5: Customize your managed fleets

As you prepare for game launch, you'll need to fine-tune your managed hosting resources. Some of the decisions to consider include:

- Consider adding Spot fleets for cost savings. See [Reduce game hosting costs with Spot fleets](#).
- If your game server needs to communicate other Amazon resources, set up IAM roles to manage access. See [Connect your Amazon GameLift Servers hosted game server to other Amazon resources](#).
- Determine where geographically you want to position game servers. Add remote locations to your managed fleets. See [Hosting resource customizations](#).
- Optimize fleet performance by selecting an instance type and size and configuring the runtime to run multiple server processes. See [Optimize game server runtime configuration on managed Amazon GameLift Servers](#).
- Experiment with game session placement options for managed fleets, including customizing prioritization settings. See [Customize a game session queue](#).
- Set up automatic capacity scaling to meet expected player demand. See [Scaling game hosting capacity with Amazon GameLift Servers](#).
- Set up standby fleets in other Amazon Web Services Regions and modify queues and auto scaling to handle failovers if needed.

- Set up hosting observability tools, including analytics and logging. See [Monitoring Amazon GameLift Servers](#).
- Automate your deployment using [infrastructure as code \(IaC\)](#). See [Manage Amazon GameLift Servers hosting resources using Amazon CloudFormation](#).

Amazon GameLift Servers supports the use of Amazon CloudFormation templates for any deployment-specific configurations. You can also use the Amazon Cloud Development Kit (Amazon CDK) to define your Amazon GameLift Servers resources. For more information about the Amazon CDK, see the [Amazon Cloud Development Kit \(Amazon CDK\) Developer Guide](#).

To manage the deployment of your Amazon CloudFormation stacks, we recommend using continuous integration and continuous delivery (CI/CD) tools and services such as Amazon CodePipeline. These tools help you deploy automatically or with approval whenever you build game server binary. With a CI/CD tool or service, resources deployment for a new game server version can look like this:

- Build and test your game server binary.
- Upload the binary to Amazon GameLift Servers.
- Deploy new fleets with the new build.
- Add the new fleets to your game session queue and remove the fleets with the previous build version.
- When the fleets with the previous build are no longer hosting active game sessions, delete the Amazon CloudFormation stacks of those fleets.

Prepare a game for hosting with Amazon GameLift Servers

Before you can deploy your game for hosting on Amazon GameLift Servers, you need to prepare both your game server and game client components. This involves integrating your game server with the Amazon GameLift Servers Server SDK and building a backend service that handles player authentication and game session management.

The Server SDK integration enables your game server to communicate with the Amazon GameLift Servers service, report its health status, and manage game sessions. The backend service acts as an intermediary between your game clients and the Amazon GameLift Servers service, handling player requests and coordinating game session placement.

Topics

- [Prepare your Unreal or Unity game with the Amazon GameLift Servers plugin](#)
- [Integrate a game server with Amazon GameLift Servers](#)
- [Package a game server build for deployment](#)
- [Game client/server interactions with Amazon GameLift Servers](#)

Prepare your Unreal or Unity game with the Amazon GameLift Servers plugin

The Amazon GameLift Servers plugin is a full-featured add-on to your Unreal or Unity game engine. It guides you through the basic steps to deploy your game for hosting with Amazon GameLift Servers. With the plugin's tool set and workflows, you can work in your game engine development environment to prepare your game server for hosting, set up hosting on a local machine for testing, create a simple backend service, and deploy your game server to managed cloud-based hosting.

Use the plugin to experience working with Amazon GameLift Servers and get a game hosting solution up and running fast. You can work with sample game assets or your own game project. The plugin automates a number of steps so that you can quickly build a simple working solution. When you complete the plugin's guided workflows, you'll be able to connect a game client to live hosted game sessions through Amazon GameLift Servers. After using the plugin to create a simple hosting solution, you can customize your solution to meet the needs of your game.

The plugin is available for the following game engines:

- Unreal Engine
- Unity

The plugin includes these components for each game engine:

- Plugin modules for the game engine editor. When the plugin is installed, a new main menu button gives you access to Amazon GameLift Servers functionality.
- Libraries for the Amazon GameLift Servers service API with client-side functionality.
- Libraries for the Amazon GameLift Servers server SDK (version 5).
- Sample assets for use with testing a server integration.
- Editable configurations, in the form of Amazon CloudFormation templates, that define your game server solution.

Topics

- [Plugin workflow](#)
- [Amazon GameLift Servers plugin for Unity for server SDK 4](#)

Plugin workflow

The following steps describe a typical path to preparing and deploying your game project on Amazon GameLift Servers. You complete these steps by working in the game engine editor and your game code.

1. Create a user profile that links to your Amazon account user and provides access credentials with permissions to use Amazon GameLift Servers.
2. Set up related Amazon resources that the plugin uses in the hosting solution (referred to as "bootstrapping").
3. Add server code to your project to establish communication between a running game server and the Amazon GameLift Servers service.
4. Add client code to your project that lets game clients send requests to Amazon GameLift Servers to start new game sessions and then connect to them.
5. Use the Anywhere workflow to set up your local workstation as an Anywhere compute and host your game server. Launch your game server and client locally through the plugin, connect to a game session, and test the integration.
6. Use the Managed EC2 workflow to upload your game server to Amazon GameLift Servers and deploy a simple but complete cloud hosting solution. Launch your game client locally through the plugin, request a game session and connect to it, and play your game.

When working in the plugin, you'll create and use Amazon resources. These actions might incur charges to the Amazon account in use. If you're new to Amazon, these actions might be covered under the [Amazon Free Tier](#).

Amazon GameLift Servers plugin for Unity for server SDK 4

Amazon GameLift Servers provides tools for preparing your multiplayer game servers to run on Amazon GameLift Servers. The Amazon GameLift Servers plugin for Unity makes it easier to integrate Amazon GameLift Servers into your Unity game projects and deploy Amazon GameLift Servers resources for cloud hosting. Use the plugin for Unity to access Amazon GameLift Servers APIs and deploy Amazon CloudFormation templates for common gaming scenarios.

After you've set up the plugin, you can try out the [Amazon GameLift Servers Unity sample](#) on GitHub.

Topics

- [Integrate Amazon GameLift Servers with a Unity game server project](#)
- [Integrate Amazon GameLift Servers with a Unity game client project](#)
- [Install and set up the plugin](#)
- [Test your game locally](#)
- [Deploy a scenario](#)
- [Integrate games with Amazon GameLift Servers in Unity](#)
- [Import and run a sample game](#)

Integrate Amazon GameLift Servers with a Unity game server project

This topic helps you prepare your custom game server for hosting on Amazon GameLift Servers. The game server must be able to notify Amazon GameLift Servers about its status, to start and stop game sessions when prompted, and to perform other tasks. For more information, see [Add Amazon GameLift Servers to your game server with the server SDK](#).

Prerequisites

Before integrating your game server, complete the following tasks:

- [Set up an IAM service role for Amazon GameLift Servers](#)
- [Install and set up the plugin](#)

Set up a new server process

Set up communication with Amazon GameLift Servers and report that the server process is ready to host a game session.

1. Initialize the server SDK by calling `InitSDK()`.
2. To prepare the server to accept a game session, call `ProcessReady()` with the connection port and game session location details. Include the names of callback functions that Amazon GameLift Servers service invokes, such as `OnGameSession()`, `OnGameSessionUpdate()`, `OnProcessTerminate()`, `OnHealthCheck()`. Amazon GameLift Servers might take a few minutes to provide a callback.

3. Amazon GameLift Servers updates the status of the server process to ACTIVE.
4. Amazon GameLift Servers calls `onHealthCheck` periodically.

The following code example shows how to set up a simple server process with Amazon GameLift Servers.

```
//initSDK
var initSDKOutcome = GameLiftServerAPI.InitSDK();

//processReady
// Set parameters and call ProcessReady
var processParams = new ProcessParameters(
    this.OnGameSession,
    this.OnProcessTerminate,
    this.OnHealthCheck,
    this.OnGameSessionUpdate,
    port,
    // Examples of log and error files written by the game server
    new LogParameters(new List<string>()
        {
            "C:\\game\\logs",
            "C:\\game\\error"
        })
);

var processReadyOutcome = GameLiftServerAPI.ProcessReady(processParams);

// Implement callback functions
void OnGameSession(GameSession gameSession)
{
    // game-specific tasks when starting a new game session, such as loading map
    // When ready to receive players
    var activateGameSessionOutcome = GameLiftServerAPI.ActivateGameSession();
}

void OnProcessTerminate()
{
    // game-specific tasks required to gracefully shut down a game session,
    // such as notifying players, preserving game state data, and other cleanup
    var ProcessEndingOutcome = GameLiftServerAPI.ProcessEnding();
}
```

```
bool OnHealthCheck()
{
    bool isHealthy;
    // complete health evaluation within 60 seconds and set health
    return isHealthy;
}
```

Start a game session

After game initialization is complete, you can start a game session.

1. Implement the callback function `onStartGameSession`. Amazon GameLift Servers invokes this method to start a new game session on the server process and receive player connections.
2. To activate a game session, call `ActivateGameSession()`. For more information about the SDK, see [C# server SDK for Amazon GameLift Servers 4.x -- Actions](#).

The following code example illustrates how to start a game session with Amazon GameLift Servers.

```
void OnStartGameSession(GameSession gameSession)
{
    // game-specific tasks when starting a new game session, such as loading map
    ...
    // When ready to receive players
    var activateGameSessionOutcome = GameLiftServerAPI.ActivateGameSession();
}
```

End a game session

Notify Amazon GameLift Servers when a game session is ending. As a best practice, shut down server processes after game sessions complete to recycle and refresh hosting resources.

1. Set up a function named `onProcessTerminate` to receive requests from Amazon GameLift Servers and call `ProcessEnding()`.
2. The process status changes to `TERMINATED`.

The following example describes how to end a process for a game session.

```
var processEndingOutcome = GameLiftServerAPI.ProcessEnding();
```

```
if (processReadyOutcome.Success)
    Environment.Exit(0);

// otherwise, exit with error code
Environment.Exit(errorCode);
```

Create server build and upload to Amazon GameLift Servers

After you integrate your game server with Amazon GameLift Servers, upload the build files to a fleet so that Amazon GameLift Servers can deploy it for game hosting. For more information on how to upload your server to Amazon GameLift Servers, see [Create a game server build for Amazon GameLift Servers](#).

Integrate Amazon GameLift Servers with a Unity game client project

This topic helps you set up a game client to connect to Amazon GameLift Servers hosted game sessions through a backend service. Use Amazon GameLift Servers APIs to initiate matchmaking, request game session placement, and more.

Add code to the backend service project to allow communication with the Amazon GameLift Servers service. A backend service handles all game client communication with the GameLift service. For more information about backend services, see .

A backend server handles the following game client tasks:

- Customize authentication for your players.
- Request information about active game sessions from the Amazon GameLift Servers service.
- Create a new game session.
- Add a player to an existing game session.
- Remove a player from an existing game session.

Topics

- [Prerequisites](#)
- [Initialize a game client](#)
- [Create game session on a specific fleet](#)
- [Add players to game sessions](#)
- [Remove a player from a game session](#)

Prerequisites

Before setting up game server communication with the Amazon GameLift Servers client, complete the following tasks:

- [Set up an Amazon user account](#)
- [Install and set up the plugin](#)
- [Integrate Amazon GameLift Servers with a Unity game server project](#)
- [Deploy hosting fleets for Amazon GameLift Servers](#)

Initialize a game client

Add code to initialize a game client. Run this code on launch, it's necessary for other Amazon GameLift Servers functions.

1. Initialize `AmazonGameLiftClient`. Call `AmazonGameLiftClient` with either a default client configuration or a custom configuration. For more information on how to configure a client, see [Set up the Amazon GameLift Servers API](#).
2. Generate a unique player id for each player to connect to a game session. For more information see [Generate player IDs](#).

The following examples shows how to set up an Amazon GameLift Servers client.

```
public class GameLiftClient
{
    private GameLift gl;
    //A sample way to generate random player IDs.
    bool includeBrackets = false;
    bool includeDashes = true;
    string playerId = AZ::Uuid::CreateRandom().ToString<string>(includeBrackets,
includeDashes);

    private Amazon.GameLift.Model.PlayerSession psession = null;
    public AmazonGameLiftClient aglc = null;

    public void CreateGameLiftClient()
    {
        //Access Amazon GameLift Servers service by setting up a configuration.
        //The default configuration specifies a location.
```

```
var config = new AmazonGameLiftConfig();
config.RegionEndpoint = Amazon.RegionEndpoint.USEast1;

CredentialProfile profile = null;
var nscf = new SharedCredentialsFile();
nscf.TryGetProfile(profileName, out profile);
AWSCredentials credentials = profile.GetAWSCredentials(null);
//Initialize Amazon GameLift Servers Client with default client
configuration.
    aglc = new AmazonGameLiftClient(credentials, config);

    }
}
```

Create game session on a specific fleet

Add code to start new game sessions on your deployed fleets and make them available to players. After Amazon GameLift Servers has created the new game session and returned a `GameSession`, you can add players to it.

- Place a request for a new game session.
 - If your game uses fleets, call `CreateGameSession()` with a fleet or alias ID, a session name, and maximum number of concurrent players for the game.
 - If your game uses queues, call `StartGameSessionPlacement()`.

The following example shows how to create a game session.

```
public Amazon.GameLift.Model.GameSession()
{
    var cgsreq = new Amazon.GameLift.Model.CreateGameSessionRequest();
    //A unique identifier for the alias with the fleet to create a game session in.
    cgsreq.AliasId = aliasId;
    //A unique identifier for a player or entity creating the game session
    cgsreq.CreatorId = playerId;
    //The maximum number of players that can be connected simultaneously to the game
    session.
    cgsreq.MaximumPlayerSessionCount = 4;

    //Prompt an available server process to start a game session and retrieves
    connection information for the new game session
}
```



```

    Amazon.GameLift.Model.CreateGameSessionResponse cgsres =
    aglc.CreateGameSession(cgsreq);
    string gsid = cgsres.GameSession != null ? cgsres.GameSession.GameSessionId : "N/A";
    Debug.Log((int)cgsres.HttpStatusCode + " GAME SESSION CREATED: " + gsid);
    return cgsres.GameSession;
}

```

Add players to game sessions

After Amazon GameLift Servers has created the new game session and returned a `GameSession` object, you can add players to it.

1. Reserve a player slot in a game session by creating a new player session. Use `CreatePlayerSession` or `CreatePlayerSessions` with the game session ID and a unique ID for each player.
2. Connect to the game session. Retrieve the `PlayerSession` object to get the game session's connection information. You can use this information to establish a direct connection to the server process:
 - a. Use the specified port and either the DNS name or IP address of the server process.
 - b. Use the DNS name and port of your fleets. The DNS name and port are required if your fleets have TLS certificate generation enabled.
 - c. Reference the player session ID. The player session ID is required if your game server validates incoming player connections.

The following examples demonstrates how to reserve a player spot in a game session.

```

public Amazon.GameLift.Model.PlayerSession
CreatePlayerSession(Amazon.GameLift.Model.GameSession gsession)
{
    var cpsreq = new Amazon.GameLift.Model.CreatePlayerSessionRequest();
    cpsreq.GameSessionId = gsession.GameSessionId;
    //Specify game session ID.
    cpsreq.PlayerId = playerId;
    //Specify player ID.
    Amazon.GameLift.Model.CreatePlayerSessionResponse cpsres =
    aglc.CreatePlayerSession(cpsreq);
    string psid = cpsres.PlayerSession != null ? cpsres.PlayerSession.PlayerSessionId :
    "N/A";
}

```

```
    return cpsres.PlayerSession;
}
```

The following code illustrates how to connect a player with the game session.

```
public bool ConnectPlayer(int playerId, string playerSessionId)
{
    //Call ConnectPlayer with player ID and player session ID.
    return server.ConnectPlayer(playerIdx, playerSessionId);
}
```

Remove a player from a game session

You can remove the players from the game session when they leave the game.

1. Notify the Amazon GameLift Servers service that a player has disconnected from the server process. Call `RemovePlayerSession` with the player's session ID.
2. Verify that `RemovePlayerSession` returns `Success`. Then, Amazon GameLift Servers changes the player slot to be available, which Amazon GameLift Servers can assign to a new player.

The following example illustrates how to remove a player session.

```
public void DisconnectPlayer(int playerId)
{
    //Receive the player session ID.
    string playerSessionId = playerSessions[playerIdx];
    var outcome = GameLiftServerAPI.RemovePlayerSession(playerSessionId);
    if (outcome.Success)
    {
        Debug.Log (":) PLAYER SESSION REMOVED");
    }
    else
    {
        Debug.Log(":(PLAYER SESSION REMOVE FAILED. RemovePlayerSession()
        returned " + outcome.Error.ToString());
    }
}
```

Install and set up the plugin


This section describes how to download, install, and set up the Amazon GameLift Servers plugin for Unity, version 1.0.0.

Prerequisites


- Unity for Windows 2019.4 LTS, Windows 2020.3 LTS, or Unity for MacOS
- Current version of Java
- Current version of .NET 4.x

To download and install the plugin for Unity


1. Download the Amazon GameLift Servers plugin for Unity. You can find the latest version on the [Amazon GameLift Servers plugin for Unity repository](#) page. Under the [latest release](#), choose **Assets**, and then download the `com.amazonaws.gamelift-version.tgz` file.
2. Launch Unity and choose a project.
3. In the top navigation bar, under **Window** choose **Package Manager**:




4. Under the **Package Manager** tab choose **+**, and then choose **Add package from tarball...**:



5. In the **Select packages on disk** window, navigate to the `com.amazonaws.gamelift` folder, choose the file `com.amazonaws.gamelift-version.tgz`, and then choose **Open**:



6. After Unity has loaded the plug-in, **Amazon GameLift Servers** appears as a new item in the Unity menu. It may take a few minutes to install and recompile scripts. The **Amazon GameLift Servers Plugin Settings** tab automatically opens.



7. In the **SDK** pane, choose **Use .NET 4.x**.

When configured, the status changes from **Not Configured** to **Configured**.

Test your game locally

Use Amazon GameLift Servers Local to run Amazon GameLift Servers on your local device. You can use Amazon GameLift Servers Local to verify code changes in seconds, without a network connection.

Configure local testing

1. In the plugin for Unity window, choose the **Test** tab.
2. In the **Test** pane, choose **Download Amazon GameLift Servers Local**. The plugin for Unity opens a browser window and downloads the `GameLift_06_03_2021.zip` file to your downloads folder.

The download includes the C# Server SDK, .NET source files, and a .NET component compatible with Unity.

3. Unzip the downloaded file `GameLift_06_03_2021.zip`.
4. In the **Amazon GameLift Servers Plugin Settings** window, choose **Amazon GameLift Servers Local Path**, navigate to the unzipped folder, choose the file `GameLiftLocal.jar`, and then choose **Open**.

When configured, local testing status changes from **Not Configured** to **Configured**.

5. Verify the status of the JRE. If the status is **Not Configured**, choose **Download JRE** and install the recommended Java version.

After you install and configure the Java environment, the status changes to **Configured**.

Run your local game

1. In the plugin for Unity tab, choose the **Test** tab.
2. In the **Test** pane, choose **Open Local Test UI**.
3. In the **Local Testing** window, specify a **Server executable path**. Select ... to select the path and executable name of your server application.
4. In the **Local Testing** window, specify a **GL Local port**.
5. Choose **Deploy & Run** to deploy and run the server.
6. To stop your game server, choose **Stop** or close the game server windows.

Deploy a scenario

A scenario uses an Amazon CloudFormation template to create the resources you need to deploy a cloud hosting solution for your game. This section describes the scenarios Amazon GameLift Servers provides and how to use them.

Prerequisites

To deploy the scenario, you need an IAM role for the Amazon GameLift Servers service. For information on how to create a role for Amazon GameLift Servers, see [Set up an Amazon user account](#).

Each scenario requires permissions to the following resources:

- Amazon GameLift Servers
- Amazon S3
- Amazon CloudFormation
- API Gateway
- Amazon Lambda

- Amazon WAFV2
- Amazon Cognito

Scenarios

The Amazon GameLift Servers Plug-in for Unity includes the following scenarios:

Auth only

This scenario creates a game backend service that performs player authentication without game server capability. The template creates the following resources in your account:

- An Amazon Cognito user pool to store player authentication information.
- An Amazon API Gateway REST endpoint-backed Amazon Lambda handler that starts games and views game connection information.

Single-Region fleet

This scenario creates a game backend service with a single Amazon GameLift Servers fleet. It creates the following resources:

- An Amazon Cognito user pool for a player to authenticate and start a game.
- An Amazon Lambda handler to search for an existing game session with an open player slot on the fleet. If it can't find a open slot, it creates a new game session.

Multi-Region fleet with a queue and custom matchmaker

This scenario forms matches by using Amazon GameLift Servers queues and a custom matchmaker to group together the oldest players in the waiting pool. It creates the following resources:

- An Amazon Simple Notification Service topic that Amazon GameLift Servers publishes messages to. For more information on SNS topics and notifications, see [Set up event notification for game session placement](#).
- A Lambda function that's invoked by the message that communicates placement and game connection details.
- An Amazon DynamoDB table to store placement and game connection details.
GetGameConnection calls read from this table and return the connection information to the game client.

Spot fleets with a queue and custom matchmaker

This scenario forms matches by using Amazon GameLift Servers queues and a custom matchmaker and configures three fleets. It creates the following resources:

- Two Spot fleets that contain different instance types to provide durability for Spot unavailability.
- An On-Demand fleet that acts as a backup for the other Spot fleets. For more information on designing your fleets, see [Hosting resource customizations](#).
- An Amazon GameLift Servers queue to keep server availability high and cost low. For more information and best practices about queues, see [Customize a game session queue](#).

FlexMatch

This scenario uses FlexMatch, a managed matchmaking service, to match game players together. For more information about FlexMatch, see [What is Amazon GameLift Servers FlexMatch](#). This scenario creates the following resources:

- A Lambda function to create a matchmaking ticket after it receives StartGame requests.
- A separate Lambda function to listen to FlexMatch match events.

To avoid unnecessary charges on your Amazon Web Services account, remove the resources created by each scenario after you are done using them. Delete the corresponding Amazon CloudFormation stack.

Update Amazon credentials

The Amazon GameLift Servers plugin for Unity requires security credentials to deploy a scenario. You can either create new credentials or use existing credentials.

For more information about configuring credentials, see [Understanding and getting your Amazon credentials](#).

To update Amazon credentials

1. In Unity, in the plugin for Unity tab, choose the **Deploy** tab.
2. In the **Deploy** pane, choose **Amazon Credentials**.
3. You can create new Amazon Web Services credentials or choose existing credentials.

- To create credentials, choose **Create new credentials profile**, and then specify the **New Profile Name**, **Amazon Access Key ID**, **Amazon Secret Key**, and **Amazon Web Services Region**.
 - To choose existing credentials, choose **Choose existing credentials profile** and then select a profile name and **Amazon Web Services Region**.
4. In the **Update Amazon Credentials** window, choose **Update Credentials Profile**.

Update account bootstrap

The bootstrap location is an Amazon S3 bucket used during deployment. It's used to store game server assets and other dependencies. The Amazon Web Services Region you choose for the bucket must be the same Region you will use for the scenario deployment.

For more information about Amazon S3 buckets, see [Creating, configuring, and working with Amazon Simple Storage Service buckets](#).

To update the account bootstrap location

1. In Unity, in the plugin for Unity tab, choose the **Deploy** tab.
2. In the **Deploy** pane, choose **Update Account Bootstrap**.
3. In the **Account Bootstrapping** window, you choose an existing Amazon S3 bucket or create a new Amazon S3 bucket:
 - To choose an existing bucket, choose **Choose existing Amazon S3 bucket** and **Update** to save your selection.
 - Choose **Create new Amazon S3 bucket** to create a new Amazon Simple Storage Service bucket, then choose a **Policy**. The policy specifies when the Amazon S3 bucket will be expire. Choose **Create** to create the bucket.


Deploy a game scenario

You can use a scenario to test your game with Amazon GameLift Servers. Each scenario uses a Amazon CloudFormation template to create a stack with the required resources. Most of the scenarios require a game server executable and build path. When you deploy the scenario, Amazon GameLift Servers copies game assets to the bootstrap location as part of deployment.


You must configure Amazon credentials and an Amazon account bootstrap to deploy a scenario.

To deploy a scenario

1. In Unity, in the plugin for Unity tab, choose the **Deploy** tab.
2. In the **Deploy** pane, choose **Open Deployment UI**.
3. In the **Deployment** window, choose a scenario.
4. Enter a **Game Name**. It must be unique. The game name is part of the Amazon CloudFormation stack name when you deploy the scenario.
5. Choose the **Game Server Build Folder Path**. The build folder path points to the folder containing the server executable and dependencies.
6. Choose the **Game Server Build .exe File Path**. The build executable file path points to the game server executable.
7. Choose **Start Deployment** to begin deploying a scenario. You can follow the status of the update in the **Deployment** window under **Current State**. Scenarios can take several minutes to deploy.



8. When the scenario completes deployment, the **Current State** updates to include the **Cognito Client ID** and **API Gateway Endpoint** that you can copy and paste into the game.



9. To update game settings, on the Unity menu, choose **Go To Client Connection Settings**. This displays an **Inspector** tab on the right side of the Unity screen.

10. Deselect **Local Testing Mode**.

11. Enter the **API Gateway Endpoint** and the **Cognito Client ID**. Choose the same Amazon Web Services Region you used for the scenario deployment. You can then rebuild and run the game client using the deployed scenario resources.

Deleting resources created by the scenario

To delete the resources created for the scenario, delete the corresponding Amazon CloudFormation stack.

To delete resources created by the scenario


1. In the Amazon GameLift Servers plugin for Unity **Deployment** window, select **View Amazon CloudFormation Console** to open the Amazon CloudFormation console.
2. In the Amazon CloudFormation console, choose **Stacks**, and then choose the stack that includes the game name specified during deployment.
3. Choose **Delete** to delete the stack. It may take a few minutes to delete a stack. After Amazon CloudFormation deletes the stack used by the scenario, its status changes to `ROLLBACK_COMPLETE`.

Integrate games with Amazon GameLift Servers in Unity

Integrate your Unity game with Amazon GameLift Servers by completing the following tasks:

- [Integrate Amazon GameLift Servers with a Unity game server project](#)
- [Integrate Amazon GameLift Servers with a Unity game client project](#)

The following diagram shows an example flow of integrating a game. In the diagram, a fleet with the game server is deployed to Amazon GameLift Servers. The game client communicates with the game server, which communicates with Amazon GameLift Servers.



Import and run a sample game

The Amazon GameLift Servers plugin for Unity includes a sample game you can use to explore the basics of integrating your game with Amazon GameLift Servers. In this section, you build the game client and game server and then test locally using Amazon GameLift Servers Local.

Prerequisites

- [Set up an Amazon user account](#)
- [Install and set up the plugin](#)

Build and run the sample game server

Set up the game server files of the sample game.

1. In Unity, on the menu, choose **Amazon GameLift Servers**, and then choose **Import Sample Game**.

2. In the **Import Sample Game** window, choose **Import** to import the game, its assets and dependencies.
3. Build the game server. In Unity, on the menu, choose **Amazon GameLift Servers**, and then choose **Apply Windows Sample Server Build Settings** or **Apply MacOS Sample Server Build Settings**. After you configure the game server settings, Unity recompiles the assets.
4. In Unity, on the menu, choose **File**, and then choose **Build**. Choose **Server Build**, choose **Build**, and then choose a build folder specifically for server files.

Unity builds the sample game server, placing the executable and required assets in the specified build folder.

Build and run the sample game client

Set up the game client files of the sample game.

1. In Unity, on the menu, choose **Amazon GameLift Servers**, and then choose **Apply Windows Sample Client Build Settings** or **Apply MacOS Sample Client Build Settings**. After the game client settings are configured, Unity will recompile assets.
2. In Unity, on the menu, select **Go To Client Settings**. This will display an **Inspector** tab on the right side of the Unity screen. In the **Amazon GameLift Servers Client Settings** tab, select **Local Testing Mode**.
3. Build the game client. In Unity, on the menu, choose **File**. Confirm **Server Build** is not checked, choose **Build**, and then select a build folder specifically for client files.

Unity builds the sample game client, placing the executable and required assets in the specified client build folder.

4. You've now built the game server and client. In the next steps, you run the game and see how it interacts with Amazon GameLift Servers.

Test the sample game locally

Run the sample game you imported using Amazon GameLift Servers Local.

1. Launch the game server. In Unity, in the plugin for Unity tab, choose the **Deploy** tab.
2. In the **Test** pane, choose **Open Local Test UI**.
3. In the **Local Testing** window, specify a **Game Server .exe File Path**. The path must include the executable name. For example, `C:/MyGame/GameServer/MyGameServer.exe`.

4. Choose **Deploy and Run**. The plugin for Unity launches the game server and opens an Amazon GameLift Servers Local log window. The window contains log messages including messages sent between the game server and Amazon GameLift Servers Local.
5. Launch the game client. Find the build location with the sample game client and choose the executable file .
6. In the **Amazon GameLift Servers Sample Game**, provide an email and password and then choose **Log In**. The email and password aren't validated or used.
7. In the **Amazon GameLift Servers Sample Game**, choose **Start**. The game client looks for a game session. If it can't find a session, it creates one. The game client then starts the game session. You can see game activity in the logs.

Sample game server logs

```
...
2021-09-15T19:55:3495 PID:20728 Log :) GAMELIFT AWAKE
2021-09-15T19:55:3512 PID:20728 Log :) I AM SERVER
2021-09-15T19:55:3514 PID:20728 Log :) GAMELIFT StartServer at port 33430.
2021-09-15T19:55:3514 PID:20728 Log :) SDK VERSION: 4.0.2
2021-09-15T19:55:3556 PID:20728 Log :) SERVER IS IN A GAMELIFT FLEET
2021-09-15T19:55:3577 PID:20728 Log :) PROCESSREADY SUCCESS.
2021-09-15T19:55:3577 PID:20728 Log :) GAMELIFT HEALTH CHECK REQUESTED (HEALTHY)
...
2021-09-15T19:55:3634 PID:20728 Log :) GAMELOGIC AWAKE
2021-09-15T19:55:3635 PID:20728 Log :) GAMELOGIC START
2021-09-15T19:55:3636 PID:20728 Log :) LISTENING ON PORT 33430
2021-09-15T19:55:3636 PID:20728 Log SERVER: Frame: 0 HELLO WORLD!
...
2021-09-15T19:56:2464 PID:20728 Log :) GAMELIFT SESSION REQUESTED
2021-09-15T19:56:2468 PID:20728 Log :) GAME SESSION ACTIVATED
2021-09-15T19:56:3578 PID:20728 Log :) GAMELIFT HEALTH CHECK REQUESTED (HEALTHY)
2021-09-15T19:57:3584 PID:20728 Log :) GAMELIFT HEALTH CHECK REQUESTED (HEALTHY)
2021-09-15T19:58:0334 PID:20728 Log SERVER: Frame: 8695 Connection accepted: playerId
0 joined
2021-09-15T19:58:0335 PID:20728 Log SERVER: Frame: 8696 Connection accepted: playerId
1 joined
2021-09-15T19:58:0338 PID:20728 Log SERVER: Frame: 8697 Msg rcvd from playerId 0 Msg:
CONNECT: server IP localhost
2021-09-15T19:58:0338 PID:20728 Log SERVER: Frame: 8697 Msg rcvd from player 0:CONNECT:
server IP localhost
2021-09-15T19:58:0339 PID:20728 Log SERVER: Frame: 8697 CONNECT: player index 0
```

```

2021-09-15T19:58:0339 PID:20728 Log SERVER: Frame: 8697 Msg rcvd from playerId 1 Msg:
CONNECT: server IP localhost
2021-09-15T19:58:0339 PID:20728 Log SERVER: Frame: 8697 Msg rcvd from player 1:CONNECT:
server IP localhost
2021-09-15T19:58:0339 PID:20728 Log SERVER: Frame: 8697 CONNECT: player index 1

```

Sample Amazon GameLift Servers Local logs

```

12:55:26,000 INFO || - [SocketIOServer] main - Session store / pubsub factory used:
MemoryStoreFactory (local session store only)
12:55:28,092 WARN || - [ServerBootstrap] main - Unknown channel option 'SO_LINGER' for
channel '[id: 0xe23d0a14]'
12:55:28,101 INFO || - [SocketIOServer] nioEventLoopGroup-2-1 - SocketIO server
started at port: 5757
12:55:28,101 INFO || - [SDKConnection] main - GameLift SDK server (communicates with
your game server) has started on http://localhost:5757
12:55:28,120 INFO || - [SdkWebSocketServer] WebSocketSelector-20 - WebSocket Server
started on address localhost/127.0.0.1:5759
12:55:28,166 INFO || - [StandAloneServer] main - GameLift Client server (listens for
GameLift client APIs) has started on http://localhost:8080
12:55:28,179 INFO || - [StandAloneServer] main - GameLift server sdk http listener has
started on http://localhost:5758
12:55:35,453 INFO || - [SdkWebSocketServer] WebSocketWorker-12 - onOpen
socket: /?pID=20728&sdkVersion=4.0.2&sdkLanguage=CSharp and handshake /?
pID=20728&sdkVersion=4.0.2&sdkLanguage=CSharp
12:55:35,551 INFO || - [HostProcessManager] WebSocketWorker-12 - client connected with
pID 20728
12:55:35,718 INFO || - [GameLiftSdkHttpHandler] GameLiftSdkHttpHandler-thread-0 -
GameLift API to use: ProcessReady for pId 20728
12:55:35,718 INFO || - [ProcessReadyHandler] GameLiftSdkHttpHandler-thread-0 -
Received API call for processReady from 20728
12:55:35,738 INFO || - [ProcessReadyHandler] GameLiftSdkHttpHandler-thread-0 -
onProcessReady data: port: 33430
12:55:35,739 INFO || - [HostProcessManager] GameLiftSdkHttpHandler-thread-0 -
Registered new process with pId 20728
12:55:35,789 INFO || - [GameLiftSdkHttpHandler] GameLiftSdkHttpHandler-thread-0 -
GameLift API to use: ReportHealth for pId 20728
12:55:35,790 INFO || - [ReportHealthHandler] GameLiftSdkHttpHandler-thread-0 -
Received API call for ReportHealth from 20728
12:55:35,794 INFO || - [ReportHealthHandler] GameLiftSdkHttpHandler-thread-0 -
ReportHealth data: healthStatus: true
12:56:24,098 INFO || - [GameLiftHttpHandler] Thread-12 - API to use:
GameLift.DescribeGameSessions

```

```

12:56:24,119 INFO || - [DescribeGameSessionsDispatcher] Thread-12 - Received API call
to describe game sessions with input: {"FleetId":"fleet-123"}
12:56:24,241 INFO || - [GameLiftHttpHandler] Thread-12 - API to use:
GameLift.CreateGameSession
12:56:24,242 INFO || - [CreateGameSessionDispatcher] Thread-12 - Received API call to
create game session with input: {"FleetId":"fleet-123","MaximumPlayerSessionCount":4}
12:56:24,265 INFO || - [HostProcessManager] Thread-12 - Reserved process:
20728 for gameSession: arn:aws:gamelift:local::gamesession/fleet-123/
gsess-59f6cc44-4361-42f5-95b5-fdb5825c0f3d
12:56:24,266 INFO || - [WebSocketInvoker] Thread-12 - StartGameSessionRequest:
gameSessionId=arn:aws:gamelift:local::gamesession/fleet-123/
gsess-59f6cc44-4361-42f5-95b5-fdb5825c0f3d, fleetId=fleet-123, gameSessionName=null,
maxPlayers=4, properties=[], ipAddress=127.0.0.1, port=33430, gameSessionData?=false,
matchmakerData?=false, dnsName=localhost
12:56:24,564 INFO || - [CreateGameSessionDispatcher] Thread-12 - GameSession with
id: arn:aws:gamelift:local::gamesession/fleet-123/gsess-59f6cc44-4361-42f5-95b5-
fdb5825c0f3d created
12:56:24,585 INFO || - [GameLiftHttpHandler] Thread-12 - API to use:
GameLift.DescribeGameSessions
12:56:24,585 INFO || - [DescribeGameSessionsDispatcher] Thread-12 - Received API call
to describe game sessions with input: {"FleetId":"fleet-123"}
12:56:24,660 INFO || - [GameLiftSdkHttpHandler] GameLiftSdkHttpHandler-thread-0 -
GameLift API to use: GameSessionActivate for pId 20728
12:56:24,661 INFO || - [GameSessionActivateHandler] GameLiftSdkHttpHandler-thread-0 -
Received API call for GameSessionActivate from 20728
12:56:24,678 INFO || - [GameSessionActivateHandler] GameLiftSdkHttpHandler-thread-0
- GameSessionActivate data: gameSessionId: "arn:aws:gamelift:local::gamesession/
fleet-123/gsess-59f6cc44-4361-42f5-95b5-fdb5825c0f3d"

```

Shut down server process

After you're done with your sample game, shut down the server in Unity.

1. In the game client, choose **Quit** or close the window to stop the game client.
2. In Unity, in the **Local Testing** window, choose **Stop** or close the game server windows to stop the server.

Integrate a game server with Amazon GameLift Servers

Your custom game server must be able to interact with Amazon GameLift Servers when deployed and running on hosting resources. The topics in this section describes how to integrate your game server project with the server SDK for Amazon GameLift Servers to add required and optional

functionality. This section includes general instructions that you can use with any of the language-specific server SDKs. It also includes detailed guidance on integrating your Unreal Engine and Unity game projects.

Topics

- [Add Amazon GameLift Servers to your game server with the server SDK](#)
- [Integrate the server SDK for game engines](#)

Add Amazon GameLift Servers to your game server with the server SDK

This topic provides general guidance on the server SDK functionality that you need to add to your game server code/ This functionality enables game server processes to communicate with the Amazon GameLift Servers service. Use these instructions for game servers that you plan to deploy onto Amazon GameLift Servers managed EC2 fleets, managed container fleets, or Anywhere fleets.

Game server processes communicate with the Amazon GameLift Servers service to receive instructions from the service and to report server process health and game session status. For detailed information on interactions between your game hosting solution components (game server, backend service, game client, and Amazon GameLift Servers), see [Game client/server interactions with Amazon GameLift Servers](#).

Get the server SDK

To prepare your game for hosting, add the server SDK for Amazon GameLift Servers to your game server project. The Server SDK is available in several languages. For more information about tool support for game servers, including the server SDK, see [Get Amazon GameLift Servers development tools](#).

Download the server SDK for your development language from github.com/amazon-gamelift. You can find repos for the server SDK versions, plugins, and other development tools.

Note

If you're developing using Unreal Engine or Unity, download the Amazon GameLift Servers plugin for these game engines. The server SDK is built in and ready to use. See these links for additional information:

- Unreal Engine ([Download plugin](#)) ([Integration guide](#))

- Unity ([Download plugin](#)) ([Integration guide](#))

Server SDK API references:

- [???](#)
- [???](#)
- [???](#)
-

Initialize the server process

Add code to establish communication with Amazon GameLift Servers service and report when the game server process is ready to host a game session. This code must run before any Amazon GameLift Servers code.

1. Initialize an Amazon GameLift Servers API client by calling `InitSdk()`. If you're preparing your game server to run on Amazon GameLift Servers managed EC2 fleets, use the default `InitSDK()` ([C++](#)) ([C#](#)) ([Unreal](#)) with no parameters. The API client handles connection to the Amazon GameLift Servers service for you.
2. Notify the service that the game server process is ready to host a game session. Call `ProcessReady()` ([C++](#)) ([C#](#)) ([Unreal](#)) with the following `ProcessParameters`. Each game server process must call `ProcessReady()` only once.
 - The port number for the server process. When the server process starts a game session, it provides the port to the Amazon GameLift Servers service, which updates the game session information. Your game can retrieve this information and provide it to game clients, which use it to connect to the server process and join the game session.
 - The locations of files that you want Amazon GameLift Servers to store for you. These might include game session logs and other files that the server process generates during a game session. Although Amazon GameLift Servers saves these files temporarily on the compute where the server process is running, these are only available until the instance shuts down. You can access stored files through the [Amazon GameLift Servers console](#) or by calling the Amazon GameLift Servers API operation [GetGameSessionLogUrl\(\)](#).

❗ If you're preparing your game server for use on an Amazon GameLift Servers managed containers fleet:

You don't need to specify log parameters for a container fleet. Instead, send game session and other log data to standard output. Container fleets automatically capture all container standard output as a log stream.

- The following callback functions that allow Amazon GameLift Servers to send messages or prompts to a game server process. You must implement each of these functions in your game server code. For more information, see [ProcessParameters](#) ([C++](#)) ([C#](#)) ([Unreal](#)) .
- (Optional) `onHealthCheck` – Amazon GameLift Servers calls this function regularly to request a health status report from the server.
- `onStartGameSession` – Amazon GameLift Servers calls this function in response to the client request [CreateGameSession\(\)](#).
- `onProcessTerminate` – Amazon GameLift Servers forces the server process to stop, letting it shut down gracefully.
- (Optional) `onUpdateGameSession` – Amazon GameLift Servers delivers an updated game session object to the game server or provides a status update on a match backfill request. The [FlexMatch backfill](#) feature requires this callback.

You can also set up a game server so that it can securely access other Amazon resources that you own or control. For more information, see [Connect your Amazon GameLift Servers hosted game server to other Amazon resources](#).

(Optional) Report server process health

Add code to your game server to implement the callback function `onHealthCheck()`. Amazon GameLift Servers invokes this callback method periodically to collect health metrics. To implement this callback function, do the following:

- Evaluate the health status of the server process. For example, you might report the server process as unhealthy if any external dependencies have failed.
- Complete the health evaluation and respond to the callback within 60 seconds. If Amazon GameLift Servers doesn't receive a response in that time, it automatically considers the server process to be unhealthy.

- Return a Boolean value: true for healthy, false for unhealthy.

If you don't implement a health check callback, then Amazon GameLift Servers considers the server process to be healthy unless the server doesn't respond.

The Amazon GameLift Servers service uses server process health to end unhealthy processes and clear up resources. If a server process continues to report as unhealthy or doesn't respond for three consecutive health checks, then the service might shut down the process and start a new one. The service collects metrics on a fleet's server process health.

Start a game session

Add code to implement the callback function `onStartGameSession`. Amazon GameLift Servers invokes this callback to start a game session on the server process.

The `onStartGameSession` function takes a [GameSession](#) object as an input parameter. This object includes key game session information, such as maximum players. It can also include game data and player data. The function implementation should do the following tasks:

- Initiate actions to create a new game session based on the `GameSession` properties. At minimum, the game server must associate the game session ID, which game clients reference when connecting to the server process.
- Process game data and player data as needed. This data is in the `GameSession` object.
- Notify the Amazon GameLift Servers service when a new game session is ready to accept players. Call the server API operation `ActivateGameSession()` ([C++](#)) ([C#](#)) ([Unreal](#)). In response to a successful call, the service changes the game session status to `ACTIVE`.

(Optional) Validate a new player

If you're tracking the status of player sessions, then add code to validate a new player when they connect to a game server. Amazon GameLift Servers tracks current players and available game session slots.

For validation, a game client that tries to join a game session must include a player session ID. Amazon GameLift Servers generates this ID when your game starts new game sessions by calling [StartGameSessionPlacement\(\)](#) or [StartMatchmaking\(\)](#). On these requests, an open slot in a game session is reserved for the player session.

When the game server process receives a game client connection request, it calls `AcceptPlayerSession()` ([C++](#)) ([C#](#)) ([Unreal](#)) with the player session ID. In response, Amazon GameLift Servers verifies that the player session ID corresponds to an open slot reserved in the game session. After Amazon GameLift Servers validates the player session ID, the server process accepts the connection. The player can then join the game session. If Amazon GameLift Servers doesn't validate the player session ID, then the server process denies the connection.

(Optional) Report a player session ending

If you're tracking the status of player sessions, add code to notify Amazon GameLift Servers when a player leaves the game session. This code should run whenever the server process detects a dropped connection. Amazon GameLift Servers uses this notification to track current players and available slots in the game session.

To handle dropped connections in your code, add a call to the server API operation `RemovePlayerSession()` ([C++](#)) ([C#](#)) ([Unreal](#)) with the corresponding player session ID.

End a game session

Add code to the server process shutdown sequence to notify Amazon GameLift Servers when a game session is ending. To recycle and refresh hosting resources, shut down each server process after the game session is complete.

At the start of the server process shutdown code, call the server API operation `ProcessEnding()` ([C++](#)) ([C#](#)) ([Unreal](#)). This call notifies Amazon GameLift Servers that the server process is shutting down. Amazon GameLift Servers changes the game session status and server process status to `TERMINATED`. After calling `ProcessEnding()`, it's safe for the process to shut down.

Respond to a server process shutdown notification

Add code to shut down the server process in response to a notification from the Amazon GameLift Servers service. The service sends this notification when the server process consistently reports unhealthy, or if the instance where the server process is running is being terminated. Amazon GameLift Servers can stop an instance as part of a capacity scale-down event, or in response to Spot Instance interruption. A Spot Instance interruption provides a two-minute notice, which gives the server process time to disconnect players gracefully, preserve game state data, and perform other cleanup tasks.

To handle a shutdown notification, make the following changes to your game server code:

- Implement the callback function `OnProcessTerminate()` ([C++](#)) ([C#](#)) ([Unreal](#)). This function should call code that shuts down the server process.
- Call the server API operation `GetTerminationTime()` ([C++](#)) ([C#](#)) from your game server shutdown code. If Amazon GameLift Servers has issued a call to stop the server process, then `GetTerminationTime()` returns the estimated termination time.
- At the start of your game server shutdown code, call the server API operation `ProcessEnding()` ([C++](#)) ([C#](#)) ([Unreal](#)). This call notifies the Amazon GameLift Servers service that the server process is shutting down. The service then changes the server process status to `TERMINATED`. After calling `ProcessEnding()`, it's safe for the process to shut down.

Integrate the server SDK for game engines

You can use the managed Amazon GameLift Servers service with most major game engines that support C++ or C# libraries, including O3DE, Unreal Engine, and Unity. Build the version you need for your game; see the README files with each version for build instructions and minimum requirements. For more information on available Amazon GameLift Servers SDKs, supported development platforms and operating systems, see [Get Amazon GameLift Servers development tools](#) for game servers.

In addition to the engine-specific information provided in this topic, find additional help with integrating Amazon GameLift Servers into your game servers, clients and services in the following topics:

- [Add Amazon GameLift Servers to your game server with the server SDK](#) – Detailed instructions on integrating Amazon GameLift Servers into a game server.
- [Integrate Amazon GameLift Servers game client functionality](#) – Detailed instructions on integrating into a game client or service, including creating game sessions and joining players to games.

Unreal Engine

Game servers

Prepare your game servers for hosting on Amazon GameLift Servers by adding the [server SDK for Amazon GameLift Servers for Unreal Engine](#) to your project and implementing the required server functionality. For help setting up the Unreal Engine plugin and adding Amazon GameLift Servers code, see [???](#).

Game clients and services

Enable your game clients and/or game services to interact with Amazon GameLift Servers service, such as to find available game sessions or create new ones, and add players to games. Core client functionality is provided in the [Amazon SDK for C++](#). To integrate Amazon GameLift Servers into your Unreal Engine game project, see [Integrate Amazon GameLift Servers game client functionality](#).

Unity

Game servers

Prepare your game servers for hosting on Amazon GameLift Servers by adding the [server SDK for Amazon GameLift Servers for C#](#) to your project and implementing the required server functionality. For help setting up with Unity and adding Amazon GameLift Servers code, see [Integrate Amazon GameLift Servers into a Unity project](#).

Game clients and services

Enable your game clients and/or game services to interact with Amazon GameLift Servers service, such as to find available game sessions or create new ones, and add players to games. Core client functionality is provided in the [Amazon SDK for .NET](#). To integrate Amazon GameLift Servers into your Unity game project, see [Integrate Amazon GameLift Servers game client functionality](#).

Other engines

For a full list of the Amazon GameLift Servers SDKs available for game servers and clients, see [the section called "Get development tools"](#).

Integrate Amazon GameLift Servers into a Unity project

This topic helps you set up the Amazon GameLift Servers C# Server SDK in your Unity game server projects. If you're unsure whether the managed Amazon GameLift Servers service supports the operating systems you're using, see [For game servers](#).

After you've set up your Amazon GameLift Servers Plug-in for Unity, you can try out the [Amazon GameLift Servers Unity sample](#) on GitHub.

Set up the C# server SDK for unity

Follow these steps to build the server SDK for Amazon GameLift Servers for C# and add it to your Unity game server projects.

To set up the server SDK for Amazon GameLift Servers for unity

1. **Download the [Amazon GameLift Servers server SDK](#).** To verify that your game system requirements are supported, see [Get Amazon GameLift Servers development tools](#). The Server SDK zip file includes the C# Server SDK, with source files so that you can build the SDK as needed for your project.
2. **Build the C# SDK libraries.** In an IDE, load the C# Server SDK solution file that you want to use. Use the IDEs functionality to restore NuGet files for the project. See the README.md file for the C# Server SDK for minimum requirements and additional build options. Build the solution to generate the C# SDK libraries.
3. **Check the Configuration settings.** In the Unity Editor, open your game project. Go to **File, Build Settings, Player Settings**. Under **Other Settings, Configuration**, check the following settings:
 - Scripting Runtime Version: Set to the .NET solution you're using.
4. **Add the Amazon GameLift Servers libraries to your Unity project.** In the Unity Editor, import the libraries that were produced by the solution build into the Assets/Plugins directory of your project. See the README.md file for a complete list of the libraries for the SDK version that you're using.

Add Amazon GameLift Servers server code

For more information on adding Amazon GameLift Servers functionality, see these topics:

- [Add Amazon GameLift Servers to your game server with the server SDK](#)
- [???](#)

The following code example uses a MonoBehaviour to illustrate a simple game server initialization with Amazon GameLift Servers.

```
using UnityEngine;
using Aws.GameLift.Server;
using System.Collections.Generic;

public class GameLiftServerExampleBehavior : MonoBehaviour
{
    //This is an example of a simple integration with GameLift server SDK that makes
    game server
```



```

//processes go active on Amazon GameLift
public void Start()
{
    //Set the port that your game service is listening on for incoming player
connections (hard-coded here for simplicity)
    var listeningPort = 7777;

    //InitSDK establishes a local connection with the Amazon GameLift agent to
enable
    //further communication.
    var initSDKOutcome = GameLiftServerAPI.InitSDK();
    if (initSDKOutcome.Success)
    {
        ProcessParameters processParameters = new ProcessParameters(
            (gameSession) => {
                //Respond to new game session activation request. GameLift sends
activation request
                //to the game server along with a game session object containing
game properties
                //and other settings. Once the game server is ready to receive
player connections,
                //invoke GameLiftServerAPI.ActivateGameSession()
                GameLiftServerAPI.ActivateGameSession();
            },
            () => {
                //OnProcessTerminate callback. GameLift invokes this callback
before shutting down
                //an instance hosting this game server. It gives this game server a
chance to save
                //its state, communicate with services, etc., before being shut
down.
                //In this case, we simply tell GameLift we are indeed going to shut
down.

                GameLiftServerAPI.ProcessEnding();
            },
            () => {
                //This is the HealthCheck callback.
                //GameLift invokes this callback every 60 seconds or so.
                //Here, a game server might want to check the health of
dependencies and such.
                //Simply return true if healthy, false otherwise.
                //The game server has 60 seconds to respond with its health
status.
            }
        );
    }
}

```

```

        //GameLift will default to 'false' if the game server doesn't
respond in time.
        //In this case, we're always healthy!
        return true;
    },
    //Here, the game server tells GameLift what port it is listening on for
incoming player
    //connections. In this example, the port is hardcoded for simplicity.
Active game
    //that are on the same instance must have unique ports.
    listeningPort,
    new LogParameters(new List<string>()
    {
        //Here, the game server tells GameLift what set of files to upload
when the game session ends.
        //GameLift uploads everything specified here for the developers to
fetch later.
        "/local/game/logs/myserver.log"
    }));

    //Calling ProcessReady tells GameLift this game server is ready to receive
incoming game sessions!
    var processReadyOutcome =
GameLiftServerAPI.ProcessReady(processParameters);
    if (processReadyOutcome.Success)
    {
        print("ProcessReady success.");
    }
    else
    {
        print("ProcessReady failure : " +
processReadyOutcome.Error.ToString());
    }
}
else
{
    print("InitSDK failure : " + initSDKOutcome.Error.ToString());
}
}

void OnApplicationQuit()
{
    //Make sure to call GameLiftServerAPI.ProcessEnding() when the application
quits.

```

```
//This resets the local connection with GameLift's agent.  
GameLiftServerAPI.ProcessEnding();  
}  
}
```

Package a game server build for deployment

Get your game server software ready for deployment on your hosting resources. Once deployed, the software is installed on each hosting resource, and then one or more game server process is launched and readied to host game sessions for players.

Getting a game server build ready varies depending on the type of Amazon GameLift Servers hosting options you're using. All game server builds must be integrated with the server SDK for Amazon GameLift Servers, as described in [Integrate a game server with Amazon GameLift Servers](#).

The topics in this section provide guidance on how to get your software ready for deployment to the following scenarios.

- For managed EC2 hosting, package your server software and upload it to Amazon GameLift Servers for deployment.
- For managed containers hosting, build a container image with your server software and store it in Amazon Elastic Container Registry for deployment.
- For hosting with Amazon GameLift Servers Anywhere, package your server software as needed for installation on your own hosting resources.

Note

If you're deploying an Amazon GameLift Servers Realtime configured script, see [Upload a script for Amazon GameLift Servers Realtime](#).

Topics

- [Create a game server build for Amazon GameLift Servers](#)
- [Build a container image for Amazon GameLift Servers](#)

Create a game server build for Amazon GameLift Servers

After you integrate your game server with Amazon GameLift Servers (see [Prepare a game for hosting with Amazon GameLift Servers](#)), install the game server software onto your compute resources for hosting. This process varies depending on the type of Amazon GameLift Servers hosting you're using.

Deploy for managed hosting

If you're using Amazon GameLift Servers managed EC2 hosting, you must package your game server software and upload it to Amazon GameLift Servers. When you create a managed fleet, Amazon GameLift Servers automatically deploys it to each fleet instance.

The topics in this section describe how to package your build files for uploading, create an optional build install script, and then upload the files using the [Amazon Command Line Interface \(Amazon CLI\)](#) or Amazon SDK.

Deploy for Anywhere hosting

If you're using Amazon GameLift Servers Anywhere fleets for self-managed hosting, it's your responsibility to install the game server software onto each compute in a fleet and keep it updated.

When an integrated game server process starts running, it automatically initializes and establishes communication with the Amazon GameLift Servers service. The server process starts game sessions from prompts by Amazon GameLift Servers and reports activity back to the service.

Topics

- [Package your game build files](#)
- [Add a build install script](#)
- [Create an Amazon GameLift Servers build resource for managed hosting](#)

Package your game build files

Before uploading your configured game server to Amazon GameLift Servers, package the game build files into a build directory. This process is a requirement when hosting with an EC2 managed fleet, and is a best practice when hosting with an Anywhere fleet. The build directory should include all the components that are required to run your game servers and host game sessions. This might include the following:

- **Game server binaries** – The binary files required to run the game server. A build can include binaries for multiple game servers built to run on the same platform. For a list of supported platforms, see [Get Amazon GameLift Servers development tools](#).
- **Dependencies** – Any dependent files that your game server executables require to run. Examples include assets, configuration files, and dependent libraries.

 Note

For game builds created with the server SDK for Amazon GameLift Servers for C++ (including those created with the Unreal plugin), include the OpenSSL DLL for the same version of OpenSSL that you built the server SDK with. See the server SDK README file for more details.

- **Install script** (Optional) – A script file to handle tasks that install your game build on Amazon GameLift Servers hosting servers. Place this file at the root of the build directory. Amazon GameLift Servers runs the install script as part of fleet creation.

You can set up any application in your build, including your install script, to access your resources securely on other Amazon services. For information about ways to do this, see [Connect your Amazon GameLift Servers hosted game server to other Amazon resources](#).

After you've packaged your build files, make sure that your game server can run on a clean installation of your target OS to verify that all required dependencies are included and that your install script is accurate.

Add a build install script

Create an install script for the operating system (OS) of your game build:

- Windows: Create a batch file named `install.bat`.
- Linux: Create a shell script file named `install.sh`.

When creating an install script, keep in mind the following:

- The script can't take any user input.
- Amazon GameLift Servers installs the build and recreates the file directories in your build package on a hosting server in the following locations:

- Windows fleets: C:\game
- Linux fleets: /local/game
- During the installation process for Linux fleets, the run-as user has limited access to the instance file structure. This user has full rights to the directory where your build files are installed. If your install script performs actions that require administrator permissions, then specify admin access using **sudo**. The run-as user for Windows fleets has administrator permissions by default. Permission failures related to the install script generate an event message that indicates a problem with the script.
- On Linux, Amazon GameLift Servers supports common shell interpreter languages such as bash. Add a shebang (such as `#!/bin/bash`) to the top of your install script. To verify support for your preferred shell commands, remotely access an active Linux instance and open a shell prompt. For more information, see [Remotely connect to Amazon GameLift Servers fleet instances](#).
- The install script can't rely on a VPC peering connection. A VPC peering connection isn't available until after Amazon GameLift Servers installs the build on fleet instances.

Example Windows install bash file

This example `install.bat` file installs Visual C++ runtime components required for the game server and writes the results to a log file. The script includes the component file in the build package at the root.

```
vcredist_x64.exe /install /quiet /norestart /log c:\game\vcredist_2013_x64.log
```

Example Linux install shell script

This example `install.sh` file uses bash in the install script and writes results to a log file.

```
#!/bin/bash  
echo 'Hello World' > install.log
```

This example `install.sh` file shows how you can use the Amazon CloudWatch agent to collect system-level and custom metrics, and handle log rotation. Because Amazon GameLift Servers runs in a service VPC, you must grant Amazon GameLift Servers permissions to assume an Amazon Identity and Access Management (IAM) role on your behalf. To allow Amazon GameLift Servers to assume a role, create a role that includes the Amazon managed policy `CloudWatchAgentAdminPolicy`, and use that role when you create a fleet.

```

sudo yum install -y amazon-cloudwatch-agent
sudo yum install -y https://dl.fedoraproject.org/pub/epel/epel-release-
latest-7.noarch.rpm
sudo yum install -y collectd
cat <<'EOF' > /tmp/config.json
{
    "agent": {
        "metrics_collection_interval": 60,
        "run_as_user": "root",
        "credentials": {
            "role_arn": "arn:aws:iam::account#:role/rolename"
        }
    },
    "logs": {
        "logs_collected": {
            "files": {
                "collect_list": [
                    {
                        "file_path": "/tmp/log",
                        "log_group_name": "gllog",
                        "log_stream_name": "{instance_id}"
                    }
                ]
            }
        }
    },
    "metrics": {
        "namespace": "GL_Metric",
        "append_dimensions": {
            "ImageId": "${aws:ImageId}",
            "InstanceId": "${aws:InstanceId}",
            "InstanceType": "${aws:InstanceType}"
        },
        "metrics_collected": {
            // Configure metrics you want to collect.
            // For more information, see Manually create or edit the CloudWatch agent configuration file.
        }
    }
}
EOF
sudo mv /tmp/config.json /opt/aws/amazon-cloudwatch-agent/bin/config.json

```

```
sudo /opt/aws/amazon-cloudwatch-agent/bin/amazon-cloudwatch-agent-ctl -a fetch-config -
m ec2 -s -c file:/opt/aws/amazon-cloudwatch-agent/bin/config.json
sudo systemctl enable amazon-cloudwatch-agent.service
```

Create an Amazon GameLift Servers build resource for managed hosting

When creating a build and uploading your files, you have a couple of options:

- [Create a build from a file directory](#). This is the simplest and most commonly used option.
- [Create a build with files in Amazon Simple Storage Service \(Amazon S3\)](#). With this option, you can manage your build versions in Amazon S3.

With both methods, Amazon GameLift Servers creates a new build resource with a unique build ID and other metadata. The build starts in the **Initialized** status. After Amazon GameLift Servers acquires the game server files, the build moves to **Ready** status.

When the build is ready, you can deploy it to a new Amazon GameLift Servers fleet. For more information, see [Create an Amazon GameLift Servers managed EC2 fleet](#). When Amazon GameLift Servers sets up the new fleet, it downloads the build files to each fleet instance and installs the build files.

Create a build from a file directory

To create a game build stored in any location, including a local directory, use the [upload-build](#) Amazon CLI command. This command creates a new build record in Amazon GameLift Servers and uploads files from a location that you specify.

Send an upload request. In a command line window, enter the following **upload-build** command and parameters.

```
aws gamelift upload-build \
  --name user-defined name of build \
  --operating-system supported OS \
  --server-sdk-version server SDK for Amazon GameLift Servers version \
  --build-root build path \
  --build-version user-defined build number \
  --region region name
```

- **operating-system** – The game server build's runtime environment. You must specify an OS value. You can't update this later.

- **server-sdk-version** – The version of the Amazon GameLift Servers server SDK that your game server is integrated with. If you don't provide a value, Amazon GameLift Servers uses the default value 4.0.2. If you specify an incorrect server SDK version, the game server build might fail when calling `InitSdk` to establish a connection to the Amazon GameLift Servers service.
- **build-root** – The directory path of your build files.
- **name** – A descriptive name for the new build.
- **build-version** – The version details for the build files.
- **region** – The Amazon Region where you want to create your build. Create the build in the Region where you plan to deploy fleets. If you're deploying your game in multiple Regions, create a build in each Region.

Note

View your current default Region using the [aws configure get region](#). To change your default Region, use the [aws configure set region *region* *name*](#) command.

Examples

```
aws gamelift upload-build \  
  --operating-system AMAZON_LINUX_2023 \  
  --server-sdk-version "4.0.2" \  
  --build-root "~/mygame" \  
  --name "My Game Nightly Build" \  
  --build-version "build 255" \  
  --region us-west-2
```

```
aws gamelift upload-build \  
  --operating-system WINDOWS_2016 \  
  --server-sdk-version "4.0.2" \  
  --build-root "C:\mygame" \  
  --name "My Game Nightly Build" \  
  --build-version "build 255" \  
  --region us-west-2
```

In response to your upload request, Amazon GameLift Servers provides upload progress. On a successful upload, Amazon GameLift Servers returns the new build record ID. Upload time depends on the size of your game files and the connection speed.

Create a build with files in Amazon S3

You can store your build files in Amazon S3 and upload them to Amazon GameLift Servers from there. When you create your build, you specify the S3 bucket location, and Amazon GameLift Servers retrieves the build files directly from Amazon S3.

To create a build resource

1. **Store your build files in Amazon S3.** Create a .zip file containing the packaged build files and upload it to an S3 bucket in your Amazon Web Services account. Take note of the bucket label and the file name, you'll need these when creating an Amazon GameLift Servers build.
2. **Give Amazon GameLift Servers access to your build files.** Create an IAM role by following the instructions in [Access a game build file in Amazon S3](#). After you've created the role, take note of the new role's Amazon Resource Name (ARN), you'll need this when creating a build.
3. **Create a build.** Use the Amazon GameLift Servers console or the Amazon CLI to create a new build record. You must have the `PassRole` permission, as described in [IAM permission examples for Amazon GameLift Servers](#).

Console

1. In the [Amazon GameLift Servers console](#), in the navigation pane, choose **Hosting, Builds**.
2. On the **Builds** page, choose **Create build**.
3. On the **Create build** page, under **Build settings**, do the following:
 - a. For **Name**, enter a script name.
 - b. For **Version**, enter a version. Because you can update the content of a build, version data can help you track updates.
 - c. For **Operating system (OS)**, choose the OS of your game server build. You can't update this value later.
 - d. For **Game server build**, enter the **S3 URI** of the build object that you uploaded to Amazon S3, and choose the **Object version**. If you don't remember the Amazon S3 URI and object version, choose **Browse S3** and search for the build object.
 - e. For **IAM role**, choose the role that you created that gives Amazon GameLift Servers access to your S3 bucket and build object.
4. (Optional) Under **Tags**, add tags to the build by entering **Key** and **Value** pairs.
5. Choose **Create**.

Amazon GameLift Servers assigns an ID to the new build and uploads the designated .zip file. You can view the new build, including the status, on the **Builds** page.

Amazon CLI

To define the new build and upload your server build files, use the [create-build](#) command.

1. Open a command line window and switch to a directory where you can use the Amazon CLI.
2. Enter the following **create-build** command:

```
aws gamelift create-build \  
  --name user-defined name of build \  
  --server-sdk-version server SDK for Amazon GameLift Servers version \  
  --operating-system supported OS \  
  --build-version user-defined build number \  
  --storage-location "Bucket"=S3 bucket label, "Key"=Build .zip file  
name, "RoleArn"=Access role ARN} \  
  --region region name
```

- **name** – A descriptive name for the new build.
- **server-sdk-version** – The version of the server SDK for Amazon GameLift Servers you used to integrate your game server with Amazon GameLift Servers. If you don't provide a value, Amazon GameLift Servers uses the default value 4.0.2.
- **operating-system** – The game server build's runtime environment. You must specify an OS value. You can't update this later.
- **build-version** – The version details for the build files. This information can be useful because each new version of your game server requires a new build resource.
- **storage-location**
 - **Bucket** – The name of the S3 bucket that contains your build. For example, "my_build_files".
 - **Key** – The name of the .zip file that contains your build files. For example, "my_game_build_7.0.1, 7.0.2".
 - **RoleARN** – The ARN assigned to the IAM role that you created. For example, "arn:aws-cn:iam::11112223333:role/GameLiftAccess". For an example policy, see [Access a game build file in Amazon S3](#).
- **region** – Create the build in the Amazon Region where you plan to deploy fleets. If you're deploying your game in multiple Regions, create a build in each Region.

Note

We recommend checking your current default Region using the [configure get](#) command. To change your default Region, use the [configure set](#) command.

Example

```
aws gamelift create-build \  
  --operating-system WINDOWS_2016 \  
  --storage-location  
  "Bucket"="my_game_build_files", "Key"="mygame_build_101.zip", "RoleArn"="arn:aws:iam::111  
gamelift" \  
  --name "My Game Nightly Build" \  
  --build-version "build 101" \  
  --region us-west-2
```

3. To view the new build, use the [describe-build](#) command.

Build a container image for Amazon GameLift Servers

This topic describes how to create a container image with your game server software to use with Amazon GameLift Servers. A game server container image includes the game server executable and any dependencies it needs to run. Game server container images are used with an Amazon GameLift Servers managed containers hosting solution. For details on building the complete solution, see:

- [???](#)
- [???](#)

Complete the following tasks to get your game server container image ready for deployment to an Amazon GameLift Servers container fleet. Before starting these tasks, finish integrating your game server code with the Amazon GameLift Servers server SDK.

Topics

- [Create a game server container image](#)
- [Push a container image to Amazon ECR](#)

Create a game server container image

Follow these instructions while working on a Linux-based platform or using Windows Subsystem for Linux (WSL) with Docker installed.

To create a game server container image

1. Prepare a working directory with your game server software. On a local machine, create a working directory to organize the files for your game server container. Your container image uses this file structure when deploying the container to Amazon GameLift Servers resources for hosting. For example:

```
[~/]$ mkdir -p work/glc/gamebuild && cd work && find .  
.  
./glc  
./glc/gamebuild
```

Note

If you're trying out this feature and don't have a working game server build yet, try our sample game server, [SimpleServer](#), which is available on GitHub.

2. Create a new Dockerfile using the template provided.
3. Follow the instructions in the Dockerfile template to update it for your own use.
 - Update the base image as needed.
 - Set environment variables for your game server build.
4. Build the container image. Run the `docker build`, specifying your own repository name. For example:

```
[~/work/glc]$ docker build -t <local repository name>:<optional tag> .
```

You can view your repositories and images IDs using the `docker images` command, as illustrated in this example:

Dockerfile template for a game server container image

This template contains the minimum instructions that a game server container needs to usable in an Amazon GameLift Servers fleet. Modify the content as needed for your game server.

```
# Base image
# -----
# Add the base image that you want to use,
# Make sure to use an image with the same architecture as the
# Instance type you are planning to use on your fleets.
FROM public.ecr.aws/amazonlinux/amazonlinux
#
# Game build directory
# -----
# Add your gamebuild directory to the env variable below.
# The game build provided here needs to be integrated with server sdk for Amazon
GameLift Servers.
ENV GAME_BUILD_DIRECTORY="<ADD_GAME_BUILD_DIRECTORY>" \
#
# Game executable and launch parameters
# -----
# Add the relative path to your executable in the 'GAME_EXECUTABLE' env variable
below.
# The game build provided over here needs to be integrated with server sdk for Amazon
GameLift Servers.
# This template assumes that the executable path is relative to the game build
directory.
# Add any launch parameters to pass into your executable in the 'LAUNCH_PARAMS' env
variable below.
# Add 'HOME_DIR' to identify where the game executable and logs exist.
GAME_EXECUTABLE="<ADD NAME OF EXECUTABLE WITHIN THE GAME DIRECTORY>" \
LAUNCH_PARAMS=<ADD LAUNCH PARAMETERS> \
HOME_DIR="/local/game" \

# Install dependencies as necessary
RUN yum install -y shadow-utils

RUN mkdir -p $HOME_DIR
COPY ./GAME_BUILD_DIRECTORY/ $HOME_DIR

# Change directory to home
WORKDIR $HOME_DIR
```

```
# Set up for 'gamelift' user
RUN useradd -m gamelift && \
    echo "gamelift ALL=(ALL) NOPASSWD:ALL" >> /etc/sudoers && \
    chown -R gamelift:gamelift $HOME_DIR

# Add permissions to game build
RUN chmod +x ./${GAME_EXECUTABLE}

USER gamelift

# Check directory before starting the container
RUN ls -lhrt .

# Check path before starting the container
RUN echo $PATH

# Start the game build
ENTRYPOINT ["/bin/sh", "-c", "}.${GAME_EXECUTABLE}", "$LAUNCH_PARAMS"]
```

Push a container image to Amazon ECR

After you've created a container image for deployment to Amazon GameLift Servers, store the image in a public or private repository in Amazon ECR. This repository is assigned a URI value, which Amazon GameLift Servers uses to take a snapshot of the image for deployment to a container fleet.

Note

If you don't yet have an Amazon ECR private repository, then [create one](#).

To push your container image to Amazon ECR

1. Get your Amazon ECR credentials. Before pushing a container image to Amazon ECR, first acquire your Amazon credentials in temporary form and provide them to Docker. Docker needs these credentials to log in.

```
[~/work/glc]$ aws ecr get-login-password --region us-west-2 | docker login --
username Amazon --password-stdin aws_account_id.dkr.ecr.us-west-2.amazonaws.com
WARNING! Your password will be stored unencrypted in
/home/user-name/.docker/config.json.
```

Configure a credential helper to remove this warning.
See <https://docs.docker.com/engine/reference/commandline/login/#credentials-store>

Login Succeeded

2. Copy the URI of the [Amazon ECR private repository](#) you want to use.
3. Apply an Amazon ECR tag to your container image.

Example

```
[~/work/glc]$ docker tag <IMAGE ID from above> <Amazon ECR private repository URI>:<optional tag>
```

4. Push your container image to Amazon ECR

Example

```
[~/work/glc]$ docker image push <Amazon ECR private repository URI>
```

Game client/server interactions with Amazon GameLift Servers

The components in your Amazon GameLift Servers hosting solution interact with each other in specific ways to run game sessions in response to player demand. This topic describes how components communicate with each other when your game server is hosted on Amazon GameLift Servers managed EC2 fleets, self-managed Amazon GameLift Servers Anywhere fleets, or a hybrid solution.


Hosting solution components include a game server, the Amazon GameLift Servers service, a client-side backend service, and a game client. The game server uses the [Amazon GameLift Servers server SDK](#) to interact with the Amazon GameLift Servers service. The backend service uses the [Amazon GameLift Servers service API](#) (part of the Amazon SDK) to interact with the service on behalf of the game client. When joining a game session, the game client connects directly to a game session using the game session's unique IP address and port number.

Topics

- [Interactions diagram](#)
- [Interaction behaviors](#)

Interactions diagram

The following diagram illustrates how your game hosting components interact so that the Amazon GameLift Servers service can track the status of game server availability and start game sessions in response to player demands.



Interaction behaviors

The following sections describe the sequence of events in each of the key interactions.

Initializing a game server process

On startup, a game server process establishes communication with the Amazon GameLift Servers service and reports its status as ready to host a game session.

1. A new process of the game server executable starts running on a hosting resource.
2. The game server process calls the following server SDK operations in sequence:
 - a. `InitSDK()` to initialize the server SDK, authenticate the server process, and establish communication with the Amazon GameLift Servers service.
 - b. `ProcessReady()` to communicate readiness to host a game session. This call also reports the process's connection information, which game clients use to connect to the game session, and other information.

The server process then waits for prompts from the Amazon GameLift Servers service.

3. Amazon GameLift Servers updates the status of the server process to ACTIVE and available to host a new game session.
4. Amazon GameLift Servers begins periodically calling the `onHealthCheck` callback to request a health status from the server processes. These calls continue while the server process remains in active status. The server process must respond healthy or not healthy within one minute. If the server process responds not healthy or doesn't respond, at some point the Amazon GameLift Servers service changes the server process's active status and stops sending requests to start game session.

Creating a game session

The Amazon GameLift Servers service starts a new game session in response to a request from a player to play the game.

1. A player using the game client asks to join a game session. Depending on how your game handles the player join process, the game client sends a request to the backend service.
2. If the player join process requires starting a new game session, the backend service sends a request for a new game session to the Amazon GameLift Servers service. This request calls

the service API operation `StartGameSessionPlacement()`. (As an alternative, the backend service might call `StartMatchmaking()`, or `CreateGameSession()`.)

3. The Amazon GameLift Servers service responds by creating a new `GameSessionPlacement` ticket with status `PENDING`. It returns ticket information to the backend service, so that it can track the placement ticket status and determine when the game session is ready for players. For more information, see [Set up event notification for game session placement](#).
4. The Amazon GameLift Servers service starts the game session placement process. It identifies which fleets to look at and searches those fleets for an active server process that's not hosting a game session. On locating an available server process, the Amazon GameLift Servers service does the following:
 - a. Creates a `GameSession` object with the game session settings and player data from the placement request, and sets the status to `ACTIVATING`.
 - b. Prompts the server process to start a game session. The service invokes the server process's `onStartGameSession` callback and passes the `GameSession` object.
 - c. Changes the server process's number of game sessions to 1.
5. The server process runs its `onStartGameSession` callback function. When the server process is ready to accept player connections, it calls the server SDK operation `ActivateGameSession()` and waits for player connections.
6. The Amazon GameLift Servers service updates the `GameSession` object with connection information for the server process (as reported in the call to `ProcessReady()`) and sets the game session status to `ACTIVE`. It also updates the `GameSessionPlacement` ticket status to `FULFILLED`.
7. The backend service calls `DescribeGameSessionPlacement()` to check ticket status and get game session information. When the game session is active, the backend service notifies the game client and passes the game session connection information.
8. The game client uses the connection information to connect directly to the game server process and join the game session.

Adding a player to a game

A game can optionally use player sessions to track player connections to game sessions. Player sessions can be created individually or as part of a game session placement request.

1. The backend service calls the service API operation `CreatePlayerSession()` with a game session ID.
2. The Amazon GameLift Servers service checks the game session status (must be `ACTIVE`), and looks for an open player slot in the game session. If a slot is available, then the service does the following:
 - a. Creates a new `PlayerSession` object and sets the status to `RESERVED`.
 - b. Responds to the backend service request with player session information.
3. The backend service passes player session information to the game client along with game session connection information.
4. The game client uses the connection information and player session ID to connect directly to the game server process and ask to join the game session.
5. In response to a game client join attempt, the game server process calls the service API operation `AcceptPlayerSession()` to validate the player session ID. The server process either accepts or rejects the connection.
6. The Amazon GameLift Servers service does one of the following:
 - a. If the connection is accepted, then Amazon GameLift Servers sets the `PlayerSession` status to `ACTIVE` and passes the `PlayerSession` to the game server process.
 - b. If the game server process doesn't call `AcceptPlayerSession()` for the player session ID within a certain time period after the original `CreatePlayerSession()` request, the Amazon GameLift Servers service changes the `PlayerSession` status to `TIMEDOUT` and reopens the player slot in the game session.

Removing a player

For games that use player sessions, the game server process notifies the Amazon GameLift Servers service when a player disconnects. The service uses this information to track the status of player slots in a game session and can allow new players to use open slots.

1. A player disconnects from the game session.
2. The game server process detects the lost connection and calls the server SDK operation `RemovePlayerSession()`.
3. The Amazon GameLift Servers service changes the player session status to `COMPLETED` and reopens the player slot in the game session.

Shutting down the game session

At the end of a game session, or when shutting down the game session, the server process notifies the Amazon GameLift Servers service of game session status.

1. The game server process ends the game session and initiates process shut down by calling the server SDK operation `ProcessEnding()`.
2. The game server process calls `Destroy()` to free the server SDK from memory. This step is required for [Telemetry metrics](#) to prevent it from reporting normal process exits as crashes.
3. The Amazon GameLift Servers service does the following:
 - a. Uploads game session logs to Amazon Simple Storage Service (Amazon S3).
 - b. Changes the game session status to `TERMINATED`.
 - c. Changes the server process status to `TERMINATED`.
 - d. Depending on how the hosting solution is designed, the newly available hosting resources are allocated to run a new game server process.

Build a backend service for Amazon GameLift Servers

We recommend that you implement a game client service that authenticates your players and communicates with the Amazon GameLift Servers API. By implementing a custom game client service, you can:

- Customize authentication for your players.
- Control how Amazon GameLift Servers groups players for new game sessions or to add to existing game sessions.
- Collect information from your own resources to provide game sesplayer attributes such as skill rating for matchmaking instead of trusting the client.

Using a game client service also reduces security risks introduced by game clients interacting directly with your Amazon GameLift Servers API.

Authenticating your players

You can use Amazon Cognito and player session IDs to authenticate your game clients. To manage the lifecycle and properties of your player identities, use Amazon Cognito user pools.

If you prefer, build a custom identity solution and host it on Amazon. You can also use Lambda authorizers for custom authorization logic with API Gateway.

Additional resources:

- [Using identity pools \(federated identities\)](#) (Amazon Cognito Developer Guide)
- [Getting started with user pools](#) (Amazon Cognito Developer Guide)
- [How to Set Up Player Authentication with Amazon Cognito](#) (Amazon for Games Blog)

Integrate Amazon GameLift Servers game client functionality

Integrate Amazon GameLift Servers game hosting functionality into any solution components that need to communicate with the Amazon GameLift Servers service, such as to get game session information or create new game sessions. In most cases, this functionality is built into the backend service component. Add functionality for core tasks, including:

- Request information and status of active game sessions.
- Reserve a slot for a new player in an existing game session.
- Create a new game session for a group of players.
- Request matchmaking for one or more players.
- Provide updated data for existing game sessions.

See [Game client/server interactions with Amazon GameLift Servers](#) for additional details on how game hosting components interact through the Amazon GameLift Servers SDKs.

Set up the Amazon GameLift Servers API

Add the Amazon SDK for C++ with Amazon GameLift Servers to a project.

Add code to initialize an Amazon GameLift Servers client and store key settings. This code must run before any code dependent on Amazon GameLift Servers.

1. Set up a client configuration. Use the default client configuration or create a custom client configuration object. For more information, see [AWS::Client::ClientConfiguration](#) (C++) or [AmazonGameLiftConfig](#) (C#).

A client configuration specifies a target region and endpoint to use when contacting Amazon GameLift Servers. Region identifies the set of deployed resources (fleets, queues, and

matchmakers) to use. The default client configuration sets location to the US East (N. Virginia) Region. To use any other Region, create a custom configuration.

2. Initialize an Amazon GameLift Servers client. Use [Aws::GameLift::GameLiftClient\(\)](#) (C++) or [AmazonGameLiftClient\(\)](#) (C#) with a default client configuration or a custom client configuration.
3. If you're using player sessions, add a mechanism to generate a unique identifier for each player. For more information, see [Generate player IDs](#).
4. Collect and store the following information:
 - **Target fleet** – If you're manually creating game sessions on a specific fleet, provide a fleet ID or an alias ID that points to the target fleet. If the fleet is multi-location, specify a fleet location. As a best practice, use fleet aliases so that you can switch players from one fleet to another without having to update your backend service.
 - **Target queue** – For games that use multi-fleet queues to place new game sessions, specify the name of the queue to send placement requests to.
 - **Amazon credentials** – All calls to Amazon GameLift Servers must provide credentials for the Amazon Web Services account that hosts the game. You can acquire these credentials by creating a player user, as described in [Set up programmatic access for your game](#). Depending on how you manage access for the player user, do the following:
 - If you use a role to manage player user permissions, add code to assume the role before calling an Amazon GameLift Servers API. The request to assume the role returns a set of temporary security credentials. For more information, see [Switching to an IAM role \(Amazon API\)](#) in the *IAM User Guide*.
 - If you have long-term security credentials, configure your code to locate and use stored credentials. See [Authenticate using long-term credentials](#) in the *Amazon SDKs and Tools Reference Guide*. For information on storing credentials, see the Amazon API references for [\(C++\)](#) and [\(.NET\)](#).
 - If you have temporary security credentials, add code to regularly refresh the credentials using the Amazon Security Token Service (Amazon STS), as described in [Using temporary security credentials with the Amazon SDKs](#) in the *IAM User Guide*. The code must request new credentials before the old ones expire.

Get active game sessions

Add code to discover available game sessions and manage game session settings and metadata.

Search for active game sessions

Use [SearchGameSessions](#) to get information about a specific game session, all active sessions, or sessions that meet a set of search criteria. This call returns a [GameSession](#) object for each active game session that matches your search request. This object contains the DNS name, IP address, and port, which a game client needs to connect to the game session.

Use search criteria to get a filtered list of active game sessions for players to join. For example, you can filter sessions as follows:

- Exclude game sessions that are full: `CurrentPlayerSessionCount = MaximumPlayerSessionCount`.
- Choose game sessions based on length of time that the session has been running: Evaluate `CreationTime`.
- Find game sessions based on a custom game property: `gameSessionProperties.gameMode = "brawl"`.

Manage game session data

Use any of the following operations to retrieve or update game session information.

- [DescribeGameSessionDetails\(\)](#) – Get a game session's protection status in addition to game session information.
- [UpdateGameSession\(\)](#) – Change a game session's metadata and settings as needed.
- [GetGameSessionLogUrl](#) – Access stored game session logs.

Create game sessions

Add code to start new game sessions on your deployed fleets and make them available to players. There are two options for creating game sessions, depending on how your game hosting solution manages game session placements.

Create a game session with a multi-location queue

Use [StartGameSessionPlacement](#) to place a request for a new game session in a queue. To use this operation, create a queue. This determines where Amazon GameLift Servers places the new game session. For more information about queues and how to use them, see [Configure game session placement](#).

When creating a game session placement, specify the name of the queue to use, a game session name, a maximum number of concurrent players, and an optional set of game properties. You can also optionally provide a list of players to automatically join the game session. If you include player latency data for relevant Regions, then Amazon GameLift Servers uses this information to place the new game session on a fleet that provides the ideal gameplay experience for the players.

Game session placement is an asynchronous process. After you've placed a request, you can let it succeed or time out. You can also cancel the request at any time using [StopGameSessionPlacement](#). To check the status of your placement request, call [DescribeGameSessionPlacement](#).

Create a game session on a specific fleet

Use [CreateGameSession](#) to create a new session on a specified fleet. This synchronous operation succeeds or fails depending on whether the fleet has resources available to host a new game session. After Amazon GameLift Servers creates the new game session and returns a [GameSession](#) object, you can join players to it.

When you use this operation, provide a fleet ID or alias ID, a session name, and the maximum number of concurrent players for the game. Optionally, you can include a set of game properties. Game properties are defined in an array of key-value pairs.

If you use the Amazon GameLift Servers resource protection feature to limit the number of game sessions that a player can create, then provide the game session creator's player ID.

Join a player to a game session

Add code to reserve a player slot in an active game session and connect game clients to game sessions. This action is available if your game uses player sessions with unique player IDs. For more information about player sessions, see [Amazon GameLift Servers and the player experience](#).

1. Reserve a player slot in a game session

To reserve a player slot, create a player session for the game session. There are two ways to do this:

- If you use [StartGameSessionPlacement](#) or [StartMatchmaking](#) to create a game session, you can include one or more players in the request to create player sessions for them in the new game session.

- To add players to an existing game session, call [CreatePlayerSession](#) or [CreatePlayerSessions](#) with a game session ID.

A player session request must include a unique player ID . For more information, see [Generate player IDs](#). On receiving the request, Amazon GameLift Servers verifies that the game session is accepting new players and has available player slots. If successful, Amazon GameLift Servers reserves a slot for the player, creates the new player session, and returns a [PlayerSession](#) object.

A player session can include a set of custom player data. This data is stored in the newly created player session object. Amazon GameLift Servers passes this object to the game server when the player connects directly to the game session. When requesting multiple player sessions, map each string of player data to a player ID.

2. **Connect to a game session**

Add code to the backend service to retrieve the `PlayerSession` object (such as by calling [DescribePlayerSessions\(\)](#)) and pass it back to the game client. This object contains the DNS name, IP address, and port. The game client can use this information to establish a direct connection to the server.

- You can connect using the specified port and the DNS name or IP address assigned to the server process.
- If your game server validates incoming player connections, then reference the player session ID.

After making the connection, the game client and server process communicate directly without involving Amazon GameLift Servers. The server maintains communication with Amazon GameLift Servers to report player connection status, health status, and more.

If the game server validates incoming players, then it verifies that the player session ID matches a reserved slot in the game session, and either accepts or denies the player connection. When the player disconnects, the server process reports the dropped connection.

Pass custom game data to a game session

Your game client can pass data to a game session. When you create a game session, you can include a set of game properties (key-value pairs) or game session data (string values) in the request. You can also update an existing game session with new or updated game data. This data is passed to the game server process that is hosting the game session, for use by the game server code. You cannot delete game properties.

As an example, let's say that your game defines difficulty levels: Novice, Easy, Intermediate, and Expert. A player chooses to join an Easy, game. Your game client, through the backend service, requests a new game session with the following game property: {"Key": "Difficulty", "Value": "Easy"}. In response, Amazon GameLift Servers prompts an available game server to start a new game session and passes the `GameSession` object. The game server process uses the game property provided to set the game session's difficulty level.


Learn more

- [GameProperty data type](#)
- [SearchGameSessions\(\) examples](#)
- [UpdateGameSession\(\) GameProperties parameter](#)

Standalone game session servers with a serverless backend

Using a serverless client service architecture, the backend can view the status of matchmaking tickets from a highly scalable database instead of by directly accessing the Amazon GameLift Servers API.

The following diagram shows a serverless backend built with Amazon Web Services services that matches players into games running on Amazon GameLift Servers fleets. The following list provides a description for each numbered callout in the diagram. To try out this example, see [Multiplayer Session-based Game Hosting on Amazon](#) on GitHub.



1. The game client requests an Amazon Cognito user identity from an Amazon Cognito identity pool.
2. The game client receives temporary access credentials and requests a game session through an Amazon API Gateway API.
3. API Gateway invokes an Amazon Lambda function.
4. The Lambda function requests player data from an Amazon DynamoDB NoSQL table. The function provides the Amazon Cognito identity in the request context data.
5. The Lambda function requests a match through Amazon GameLift Servers FlexMatch matchmaking.
6. FlexMatch matches a group of players with suitable latency, and then requests a game session placement through an Amazon GameLift Servers queue. The queue has fleets with one or more Amazon Web Services Region locations in it.
7. After Amazon GameLift Servers places the session on one of the fleet's locations, Amazon GameLift Servers sends an event notification to an Amazon Simple Notification Service (Amazon SNS) topic.


8. A Lambda function receives the Amazon SNS event and processes it.
9. If the matchmaking ticket is a MatchmakingSucceeded event, then the Lambda function writes the result, along with the port and IP address of the game server, to a DynamoDB table.
10. The game client makes a signed request to API Gateway to view the status of the matchmaking ticket on a specific interval.
11. API Gateway uses a Lambda function that checks the matchmaking ticket status.
12. The Lambda function checks the DynamoDB table to see if the ticket is successful. If it has succeeded, the function sends the game server's port and IP address, along with the player session ID, back to the client. If the ticket hasn't succeeded, the function sends a response verifying that the match isn't ready yet.
13. The game client connects to the game server using TCP or UDP by using the port and IP address that the backend service provides. The game client then sends the player session ID to the game server, which then validates the ID using the server SDK for Amazon GameLift Servers.

Standalone game session servers with a WebSocket-based backend

Using an Amazon API Gateway WebSocket-based architecture, you can make matchmaking requests with WebSockets and send push notifications for matchmaking completion using server-initiated messages. This architecture improves performance by having two-way communication between the client and the server.

For more information about using API Gateway WebSock APIs, see [Working with WebSocket APIs](#).

The following diagram shows a WebSocket-based backend architecture that uses API Gateway and other Amazon Web Services services to match players into games running on Amazon GameLift Servers fleets. The following list provides a description for each numbered callout in the diagram.



1. The game client requests an Amazon Cognito user identity from an Amazon Cognito identity pool.
2. The game client signs a WebSocket connection to an API Gateway API with the Amazon Cognito credentials.
3. API Gateway calls an Amazon Lambda function on the connection. The function stores the connection information in an Amazon DynamoDB table.
4. The game client sends a message to a Lambda function, through the API Gateway API over the WebSocket connection, to request a session.

5. A Lambda function receives the message and then requests a match through Amazon GameLift Servers FlexMatch matchmaking.
6. After FlexMatch matches a group of players, FlexMatch requests a game session placement through an Amazon GameLift Servers queue.
7. After Amazon GameLift Servers places the session on one of the fleet's locations, Amazon GameLift Servers sends an event notification to an Amazon Simple Notification Service (Amazon SNS) topic.
8. A Lambda function receives the Amazon SNS event and processes it.
9. If the matchmaking ticket is a MatchmakingSucceeded event, then the Lambda function requests the correct player connection from DynamoDB. The function then sends a message to the game client through the API Gateway API over the WebSocket connection. In this architecture, the game client doesn't actively poll the status of matchmaking.
10. The game client receives the port and IP address of the game server, along with the player session ID, through the WebSocket connection.
11. The game client connects to the game server using TCP or UDP using the port and IP address that the backend service provides. The game client also sends the player session ID to the game server, which then validates the ID using the server SDK for Amazon GameLift Servers.

Deploy hosting fleets for Amazon GameLift Servers

Deploying hosting resources involves creating and configuring the compute infrastructure that will run your game servers. Amazon GameLift Servers offers several fleet types to match different hosting needs, from fully managed Amazon Web Services Cloud resources to hybrid solutions that combine cloud and on-premises infrastructure.

Choose the fleet type that best fits your requirements for cost, control, scalability, and geographic distribution. You can also combine multiple fleet types in a single hosting solution to optimize for different scenarios or player populations.

Fleet characteristics

An Amazon GameLift Servers fleet is a collection of computing resources that run your game servers and host game sessions for players. Fleets can vary in the type of compute resources you use and how the fleet is managed. A fleet's size—the number of game sessions and players that it can support—depends on the number of compute resources that you give it. All Amazon GameLift Servers fleets have the following characteristics:

- The game server processes that run on all fleets are integrated with the server SDK for Amazon GameLift Servers and communicate with the Amazon GameLift Servers service in the same way. Game servers report their availability to host game sessions and players, respond to prompts to start or stop game sessions, and other interactions.
- Amazon GameLift Servers handles game session placement for all fleets in the same way. Amazon GameLift Servers keeps track of a fleet's game server status and chooses from available game servers to host a new game session. This process is used whether your game places game sessions on a single fleet or uses a [game session queue](#) to balance hosting across multiple fleets. With a queue, you can also customize placement decisions to consider factors such as resource cost and latency.
- All fleets support the use of a FlexMatch matchmaker in collaboration with a game session placement queue. The Amazon GameLift Servers service receives player match requests, forms the matches, and passes them to the game session queue to find available game servers.
- Amazon GameLift Servers collects a wide range of fleet metrics. These include status metrics for computes and server processes, as well as usage metrics for game sessions and player activity. See the complete list of available metrics at [Monitor Amazon GameLift Servers with Amazon CloudWatch](#).

Topics

- [Decision guide: Choose a hosting option](#)
- [Amazon GameLift Servers managed EC2 fleets](#)
- [Build a hybrid hosting solution](#)

Decision guide: Choose a hosting option

Amazon GameLift Servers offers several hosting options, each designed for different hosting scenarios and requirements. This guide helps you choose the type of fleet based on your game's technical needs, operational preferences, and business constraints. For high-level descriptions, see [Amazon GameLift Servers game hosting options](#).

Fleet comparison

The following table compares the key characteristics of each hosting option:

Hosting option	Best For	Management Level	Key Benefits
Managed EC2	Production hosting with full Amazon management	Fully managed	Auto-scaling, global reach, minimal operational overhead
Managed containers	Containerized applications requiring orchestration	Fully managed	Container orchestration, resource efficiency, modern deployment
Anywhere	On-premises, hybrid, or development environments	Self-managed	Hardware control, compliance requirements, cost optimization

Key decision factors

Consider these factors when choosing your hosting option:

Operational complexity

How much infrastructure management do you want to handle? Managed fleets require minimal operational overhead, while Anywhere fleets give you full control but require more management.

Scalability requirements

Do you need automatic scaling based on player demand? Managed fleets provide built-in auto-scaling, while Anywhere fleets require manual capacity management.

Geographic distribution

Where are your players located? Managed fleets can be deployed across multiple Amazon Regions globally, while Anywhere fleets depend on your own infrastructure locations.

Compliance and security

Do you have specific compliance requirements or need to keep data on-premises? Anywhere fleets allow you to maintain full control over your infrastructure and data location.

Cost considerations

What's your budget and cost optimization strategy? Consider both infrastructure costs and operational overhead when comparing options.

Hybrid solutions

You don't have to choose just one type of fleet for your game hosting solution. Many games use a hybrid approach that combines multiple fleet types to optimize for different scenarios:

- **Development and testing:** Use Anywhere fleets for development and testing, then deploy to managed fleets for production.
- **Regional optimization:** Use managed fleets in Amazon Web Services Regions with high player density and Anywhere fleets in regions where you have existing infrastructure.
- **Cost optimization:** Use managed fleets for baseline capacity and Anywhere fleets for overflow or specific workloads.

Amazon GameLift Servers managed EC2 fleets

Amazon GameLift Servers managed EC2 fleets provide cloud-based resources for production hosting. With a managed fleet, you get the flexibility, security, and reliability of Amazon Web Services Cloud resources optimized for multiplayer game hosting. Amazon GameLift Servers provides robust host management tools.

A managed EC2 fleet is a set of Amazon Elastic Compute Cloud (Amazon EC2) instances that Amazon GameLift Servers owns and operates based on your configuration. These instances are physically located in supported Amazon Web Services Regions or Local Zones. When you create a fleet, you choose an EC2 instance type that meets your game server's requirements for computing power, memory, storage, networking capabilities.

When starting each instance in the fleet, Amazon GameLift Servers deploys your game server build with the required runtime environment. The runtime environment uses the latest Amazon Machine Image (AMI) version available when the fleet is created. All instances in the fleet use the same AMI version.

Note

As a best practice, we recommend replacing your fleets every 30 days to maintain a secure and up-to-date runtime environment for your hosted game servers. This requires creating a new fleet and migrating player traffic to it. For more guidance, see [Security best practices for Amazon GameLift Servers](#).

After installing the runtime environment and your game server build on an instance, Amazon GameLift Servers starts launching game server processes. Each game server process establishes a connection to the Amazon GameLift Servers service, reports readiness to host a game session, and begins communicating health status. Amazon GameLift Servers can then prompt the server process to start a game session.

In addition to fleet deployment, Amazon GameLift Servers handles the following host management tasks so you don't have to:

- Tracks the status of all computes in the fleet and replaces stale or unhealthy computes.
- Handles authentication for communication between server processes and the Amazon GameLift Servers service.
- Automatically starts and stops game server processes on each compute, based on your runtime configuration.
- Offers auto scaling tools to adjust fleet capacity dynamically to meet player demand.
- Reports performance metrics for the fleet's EC2 instances.

See these topics about how to set up and maintain managed EC2 fleets:

- [Development roadmap for hosting with Amazon GameLift Servers managed EC2](#)
- [Create an Amazon GameLift Servers managed EC2 fleet](#)
- [Scaling game hosting capacity with Amazon GameLift Servers](#)
- [Hosting resource customizations](#)
- [Update an Amazon GameLift Servers fleet configuration](#)

Managed EC2 fleet creation workflow

For managed fleets, Amazon GameLift Servers sets up the fleet resource and also deploys a set of compute resources with your game server software installed and running. When the creation workflow is complete and successful, the fleet has one active EC2 instance in the fleet home Region and one each in the fleet's remote locations. All instances have game are ready to host game sessions.

1. Amazon GameLift Servers creates the fleet resource in the fleet's home Region and sets desired capacity in each location to one (1) instance. Fleet and location status is set to **New**.
2. Amazon GameLift Servers begins writing events to the fleet event log.
3. Amazon GameLift Servers sets fleet status to **Downloading** and begins preparing the game server software for deployment.
 - a. Gets the uploaded game server build and extracts the compressed files.
 - b. Runs install scripts, if provided.
 - c. Sets fleet status to **Validating** and begins verifying that no errors occurred when downloading and installing the build files.
4. Amazon GameLift Servers sets the fleet status to **Building**, configures fleet hardware, and allocates one EC2 instance for each fleet instance.
5. Amazon GameLift Servers sets the fleet status to **Activating**. Launches a game server process on each instance (based on the fleet's runtime instructions) and tests connectivity between the build and the Amazon GameLift Servers service.
6. When game server processes on each instance establish a connection and report readiness to host game sessions, Amazon GameLift Servers sets fleet and location statuses to **Active**. At this point, the fleet is considered ready to host game sessions.

Create an Amazon GameLift Servers managed EC2 fleet

This topic describes how to create an Amazon GameLift Servers managed EC2 fleet. Managed fleets use Amazon Elastic Compute Cloud (Amazon EC2) compute instances that are optimized for multiplayer game hosting. You can create managed fleets that deploy computes to globally to Amazon Web Services Regions and Local Zones supported by Amazon GameLift Servers.

When you create a new managed EC2 fleet resource, you immediately initiate the first phase of fleet creation. Managed fleet creation passes through several phases as Amazon GameLift

Servers deploys an EC2 instance, installs a runtime environment and your game server build on the instance, and begins launching game servers. Depending on the runtime environment your game server build requires, Amazon GameLift Servers deploys the latest version of the Amazon Machine Image (AMI) at the time of fleet creation (and all future instances in the fleet will use the same version). You can monitor a fleet's status in the console or using the Amazon Command Line Interface (Amazon CLI). When a fleet is ready to host game sessions, the status changes to ACTIVE. For more information about managed fleet creation, see these topics:

Note

As a best practice, we recommend replacing your fleets every 30 days to maintain a secure and up-to-date runtime environment for your hosted game servers. This requires creating a new fleet and migrating player traffic to it. For more guidance, see [Security best practices for Amazon GameLift Servers](#).

To create a managed EC2 fleet

Use either the Amazon GameLift Servers console or the Amazon Command Line Interface (Amazon CLI) to create a managed EC2 fleet.

Console

To create a managed EC2 fleet

1. In the [Amazon GameLift Servers console](#), in the navigation pane, choose **Fleets**.
2. On the **Fleets** page, choose **Create fleet**.
3. Choose **Managed EC2**.
4. On the **Fleet details** page do the following:
 - a. For **Name**, enter a fleet name. We recommend including the fleet type (Spot or On-demand) in your fleet names. This makes it much easier to identify fleet types when viewing a list of fleets.
 - b. For **Description**, provide a short description of the fleet.
 - c. For **Binary type**, select **Build** to identify the type game server to deploy to this fleet.
 - d. Select a **Build** from the dropdown list of previously uploaded game server builds.
5. (Optional) Under **Additional details** for the following:

- a. For **Instance role**, specify an IAM role that authorizes applications in your game build to access other Amazon resources in your account. For more information, see [Connect your Amazon GameLift Servers hosted game server to other Amazon resources](#). To create a fleet with an instance role, your account must have the IAM PassRole permission. For more information, see [IAM permission examples for Amazon GameLift Servers](#).

You can't update these settings after fleet creation.

- b. For **Metric group**, Enter the name of a new or existing fleet metric group. You can aggregate the metrics for multiple fleets by adding them to the same metric group.

You can't update the metric group after fleet creation.

6. Choose **Next**.
7. On the **Select locations** page, select one or more additional remote locations to deploy instances to. The home Region is automatically selected based on the Region you are accessing the console from. If you select additional locations, fleet instances are also deployed in these locations.

 Important

To use Regions that aren't enabled by default, enable them in your Amazon Web Services account.

- Fleets with Regions that aren't enabled that you created before February 28, 2022 are not affected by this requirement.
- To create new multi-location fleets or to update existing multi-location fleets, first enable any Regions or Local Zones that you choose to use.

For more information about Regions that aren't enabled by default and how to enable them, see [Managing Amazon Web Services Regions](#) in the *Amazon Web Services General Reference*. See [Getting started with Local Zones](#) in the *Amazon Local Zones User Guide*.

8. Choose **Next**.
9. On the **Define instance details page**, choose

- a. **On-demand** or **Spot** instances for this fleet. For more information about fleet types, see [On-Demand Instances versus Spot Instances](#).
- b. From the **Filter architecture** menu choose **x64** or **Arm**.

 Note

Graviton Arm instances require a server build for a Linux AML. Server SDK 5.1.1 or newer is required for C++ and C#. Server SDK 5.0 or newer is required for Go. These instances do not provide out-of-the-box support for Mono installation on Amazon Linux 2023 (AL2023) or Amazon Linux 2 (AL2).

For information on Amazon EC2 Arm architectures, see [Amazon Graviton Processor](#) and [Amazon EC2 instance types](#).

For information on the instance types supported by Amazon GameLift Servers, see the EC2InstanceType values under [CreateFleet\(\) request parameters](#).

10. Select an Amazon EC2 **Instance type** from the list. For more information about choosing an instance type, see [Instance types](#). After you create the fleet, you can't change the instance type.
11. Choose **Next**.
12. On the **Configure runtime** page, under **Runtime configuration** do the following:
 - a. For **Launch path**, enter the path to the game executable in your build. On Windows instances, game servers are built to the path `C:\game`. On Linux instances, game servers are built to `/local/game`. Examples: `C:\game\MyGame\server.exe` or `/local/game/MyGame/server.exe`.
 - b. (Optional) For **Launch parameters**, enter information to pass to your game executable as a set of command line parameters. Example: `+sv_port 33435 +start_lobby`.
 - c. For **Concurrent processes**, choose the number of server processes to run concurrently on each instance in the fleet. Review the Amazon GameLift Servers [limits](#) on number of concurrent server processes.

Limits on concurrent server processes per instance apply to the total of concurrent processes for all configurations. If you configure the fleet to exceed the limit, the fleet can't activate.

13. Under **Game session activation**, provide limits for activating new game sessions on the instances in this fleet:
 - a. For **Max concurrent game session activation**, enter the number of game sessions on an instance that activate at the same time. This limit is useful when launching multiple new game sessions may have an impact on the performance of other game sessions running on the instance.
 - b. For **New activation timeout**, enter how long to wait for a session to activate. If the game session doesn't move to ACTIVE status before the timeout, Amazon GameLift Servers terminates the game session activation.
14. (Optional) Under **EC2 port settings**, do the following:
 - a. Choose **Add port setting** to define access permissions for inbound traffic connecting to the server process deployed on the fleet.
 - b. For **Type**, choose **Custom TCP** or **Custom UDP**.
 - c. For **Port range**, Enter a range of port numbers that allow inbound connections. A port range must use the format `nnnnn[-nnnnn]`, with values between 1026 and 60000. Example: **1500** or **1500-20000**.
 - d. For **IP address range**, Enter a range of IP addresses. Use CIDR notation. Example: **0.0.0.0/0** (This example allows access to anyone trying to connect.)
15. (Optional) Under **Game session resource settings** do the following:
 - a. For **Game scaling protection policy**, Turn on or off scaling protection. Amazon GameLift Servers won't terminate instance with protection during a scale down event if they're hosting an active game session.
 - b. For **Resource creation limit**, enter a maximum number of game sessions a player can create during the policy period.
16. Choose **Next**.
17. (Optional) Add tags to the build by entering **Key** and **Value** pairs. Choose **Next** to continue to fleet creation review.
18. Choose **Create**. Amazon GameLift Servers assigns an ID to the new fleet and begins the fleet activation process. You can track the new fleet's status on the **Fleets** page.

You can update the fleet's metadata and configuration at any time, regardless of fleet status. For more information, see [Update an Amazon GameLift Servers fleet configuration](#). You can update fleet capacity after the fleet has reached ACTIVE status. For more information, see [Scaling game hosting capacity with Amazon GameLift Servers](#). You can also add or remove remote locations.

Amazon CLI

Use the [create-fleet](#) command to create a fleet of compute type EC2. Amazon GameLift Servers creates the fleet resource in your current default Amazon Web Services Region (or you can add a `--region` tag to specify a different Amazon Web Services Region).

Create a minimal managed fleet

The following example request creates a new fleet with the minimal settings that are required to deploy a fleet with running game servers that game clients can connect to. The new fleet has these characteristics:

- It specifies a game server build, which has been uploaded to Amazon GameLift Servers and in READY status.
- It uses c5.large On-Demand Instances with an operating system that matches the selected game build.
- It sets the fleet's home Amazon Web Services Region to us-west-2 and deploys instances to that Region only.
- Based on the runtime configuration, each compute in the fleet runs one game server process, which means that each compute can host only one game session at a time. Game session activation timeout is set to the default value of 300 seconds, and there's no limit on the number of concurrent activations.
- Players can connect to game servers using a single port setting of 33435.
- All other features are either turned off or use default settings.

```
aws gamelift create-fleet \  
  --name MinimalFleet123 \  
  --description "A basic test fleet" \  
  --region us-west-2 \  
  --ec2-instance-type c5.large \  
  --fleet-type ON_DEMAND \  
  --build-id build-1111aaaa-22bb-33cc-44dd-5555eeee66ff \  
  --
```

```
--runtime-configuration "ServerProcesses=[{LaunchPath=C:\game\Bin64.dedicated
\MultiplayerSampleProjectLauncher_Server.exe, ConcurrentExecutions=10}]" \
--ec2-inbound-permissions
"FromPort=33435,ToPort=33435,IpRange=0.0.0.0/0,Protocol=UDP"
```

Create a fully configured managed fleet

The following example request creates a production fleet with settings for all optional features. The new fleet has these characteristics:

- It specifies a game server build, which has been uploaded to Amazon GameLift Servers and in READY status.
- It uses c5.large On-Demand Instances with the operating system that matches the selected game build.
- It sets the fleet's home Amazon Web Services Region to us-west-2 and deploys instances to the home Region and one remote location sa-east-1.
- Based on the runtime configuration:
 - Each compute in the fleet runs 10 game server processes with the same launch parameters, which means that each compute can host up to 10 game sessions simultaneously.
 - On each compute, only two game sessions can be activating at the same time. Activating game sessions must be ready to host players within 300 seconds (5 minutes) or be terminated.
- Players can connect to game servers using a port in the following range 33435 to 33535.
- All game sessions in the fleet have game session protection turned on.
- Individual players are limited to creating three new game sessions within a 15-minute period.
- Metrics for this fleet are included in the metric group AMERfleets, which (for this example) aggregates metrics for a group of fleets in North, Central, and South America.

```
aws gamelift create-fleet \
--name ProdFleet123 \
--description "A fully configured prod fleet" \
--ec2-instance-type c5.large \
--region us-west-2 \
--locations "Location=sa-east-1" \
--fleet-type ON_DEMAND \
--build-id build-1111aaaa-22bb-33cc-44dd-5555eeee66ff \
```

```
--certificate-configuration "CertificateType=GENERATED" \  
--runtime-configuration "GameSessionActivationTimeoutSeconds=300,  
MaxConcurrentGameSessionActivations=2, ServerProcesses=[{LaunchPath=C:\game  
\Bin64.dedicated\MultiplayerSampleProjectLauncher_Server.exe, Parameters+=sv_port  
33435 +start_lobby, ConcurrentExecutions=10}]" \  
--new-game-session-protection-policy "FullProtection" \  
--resource-creation-limit-policy "NewGameSessionsPerCreator=3,  
PolicyPeriodInMinutes=15" \  
--ec2-inbound-permissions  
"FromPort=33435,ToPort=33535,IpRange=0.0.0.0/0,Protocol=UDP" \  
--metric-groups "AMERfleets"
```

If the create-fleet request is successful, Amazon GameLift Servers returns a set of fleet attributes that includes the configuration settings you requested and a new fleet ID. Amazon GameLift Servers then initiates the fleet activation process and sets the fleet status and the location statuses to **New**. You can track the fleet's status and view other fleet information using these CLI commands:

- [describe-fleet-events](#)
- [describe-fleet-attributes](#)
- [describe-fleet-capacity](#)
- [describe-fleet-port-settings](#)
- [describe-fleet-utilization](#)
- [describe-runtime-configuration](#)
- [describe-fleet-location-attributes](#)
- [describe-fleet-location-capacity](#)
- [describe-fleet-location-utilization](#)

You can change the fleet's capacity and other configuration settings as needed using these commands:

- [update-fleet-attributes](#)
- [update-fleet-capacity](#)
- [update-fleet-port-settings](#)
- [update-runtime-configuration](#)
- [create-fleet-locations](#)

- [delete-fleet-locations](#)

Build a hybrid hosting solution

A hybrid hosting solution combines multiple sources of game hosting resources to host your game, including Amazon GameLift Servers managed fleets that run in the Amazon Web Services Cloud and resources that you supply and manage on your own. This topic describes some common patterns for building a hybrid solution, and provides tips on how to successfully combine self-managed game hosting with cloud-based game hosting managed by Amazon GameLift Servers.

Common hybrid patterns

- **Cost optimization:** Use the most cost-effective fleet type for each scenario, such as Anywhere fleets for baseline capacity and managed fleets for peak demand.
- **Geographic flexibility:** Deploy managed fleets in high-traffic Amazon Web Services Regions and Anywhere fleets where you have existing infrastructure or specific compliance requirements.
- **Risk mitigation:** Reduce dependency on any single hosting approach by distributing load across multiple fleet types and providers.
- **Gradual migration:** Transition from on-premises or other hosting solutions to Amazon gradually while maintaining service continuity.

Development and production split

Use different fleet types for development and production environments:

- **Development:** Anywhere fleets for cost-effective development and testing
- **Production:** Managed Amazon EC2 or container fleets for scalable, reliable production hosting

Regional optimization

Optimize fleet types based on Amazon Web Services Regions characteristics:

- **High-traffic Amazon Web Services Regions:** Managed fleets with auto-scaling for variable demand
- **Specialized Amazon Web Services Regions:** Anywhere fleets for compliance, data sovereignty, or existing infrastructure

Capacity tiering

Use different fleet types for different capacity tiers:

- **Baseline capacity:** Anywhere fleets or reserved instances for predictable load
- **Burst capacity:** Managed fleets with auto-scaling for peak demand
- **Overflow capacity:** Spot instances or additional Amazon Web Services Regions for extreme peaks

Implementation considerations

When building a hybrid solution, consider these key factors:

Game session queue configuration

Configure your game session queues to include all fleet types and set appropriate priorities and latency preferences to ensure optimal placement across your hybrid infrastructure.

Monitoring and observability

Implement comprehensive monitoring across all fleet types to maintain visibility into performance, capacity, and costs across your hybrid solution.

Operational complexity

Account for the increased operational complexity of managing multiple fleet types, including different deployment processes, monitoring tools, and troubleshooting procedures.

Network connectivity

Ensure reliable network connectivity between your different hosting environments, especially for Anywhere fleets that may be on-premises or in different cloud providers.

Getting started with hybrid hosting

To implement a hybrid hosting solution:

1. **Start simple:** Begin with a single fleet type and add others gradually as your requirements become clearer.
2. **Plan your architecture:** Design your hybrid architecture based on your specific requirements for cost, performance, compliance, and operational complexity.

3. **Configure queues:** Set up game session queues that span your different fleet types with appropriate priorities and placement strategies.
4. **Test thoroughly:** Test game session placement and failover scenarios across your hybrid infrastructure before going to production.
5. **Monitor and optimize:** Continuously monitor performance and costs across all fleet types and adjust your configuration as needed.

Tips

- Use the same game client and server components with both managed and self-managed hosting resources. Provide a unified player experience across all hosting resources.
- Use the same FlexMatch matchmakers to place matches across all hosting resources.
- Centrally manage your hybrid hosting resources together while you deploy them across the globe.
- As player demand fluctuates, manage game session loads seamlessly across managed and self-managed resources.
- With the Amazon GameLift Servers Agent, you can use the same tools to manage game server life cycles on all types of hosting resources.
- Gather game and player metrics and logs across all hosting resources. Take advantage of Amazon GameLift Servers features and other Amazon services to combine data and develop cohesive observability solutions.

Configure game session placement

Game session placement is the process of finding available game servers to host new game sessions. Amazon GameLift Servers uses game session queues to intelligently place game sessions across your fleets, taking into account factors like player latency, fleet capacity, and cost optimization.

A well-configured game session queue ensures that players are connected to the best available game server, providing optimal performance while making efficient use of your hosting resources. Queues can span multiple fleets and Amazon Regions to provide global coverage and redundancy.

A game session queue is the primary mechanism that Amazon GameLift Servers uses to search for available game servers and choose them to host new game sessions. Queues offers a far more

efficient way to process large numbers of game session requests and find placements for them across multiple fleets of hosting resources. If your hosting solution uses more than one fleet, and you're processing high volumes of requests, you probably need a queue.

When your game wants to start a new game session for players, it sends a placement request to the Amazon GameLift Servers service, which funnels it to the queue. The queue's configuration determines when and how the requests are processed. When processing a placement request, Amazon GameLift Servers searches a set of fleets for a game server to host the game session. Placement succeeds when Amazon GameLift Servers finds an available game server and prompts it to start a game session.

Topics

- [Queue characteristics](#)
- [Best practices for Amazon GameLift Servers game session queues](#)
- [Create a game session queue](#)
- [Set up event notification for game session placement](#)

Queue characteristics

An Amazon GameLift Servers game session queue is an Amazon cloud resource. You can create a queue in any Amazon Web Services Region that Amazon GameLift Servers supports (see [Amazon GameLift Servers service locations](#)). Game session placement requests are sent to that location and processed there.

Automating game session placement with queues offers significant benefits for both game developers and players. These include:

- **Queues deliver the "best possible" placement.** When processing game session placement requests, a queue uses the Amazon GameLift Servers FleetIQ algorithm to prioritize placements based on a set of defined preferences, including cost, location, and player latency.
- **Queues support Spot fleets to help reduce game hosting costs.** You can configure your queues with Amazon Spot fleets, which often offer significantly lower hosting costs, as well as On-Demand fleets. Because low cost is one of the key criteria for placements, queues can always take advantage of differences in cost.
- **Queues can place new games faster during high demand.** By configuring a queue with multiple fleets, you're providing more flexible options for game session placement. But additional

fleets also provide backup capacity as needed when demand increases. For any placement request, if Amazon GameLift Servers can't place a game session in the most preferred location, it automatically moves on to evaluate other locations.

- **Queues can make game server availability more resilient.** Outages can happen. With a multi-fleet queue, a slowdown or outage doesn't have to affect player access to your game. By configuring your queue with fleets that have capacity in different Amazon Web Services Regions and availability zones, you can help make sure that players can always find a game session to join.
- **Get metrics on game session placements and queue performance.** Amazon GameLift Servers emits queue metrics, including statistics on placement successes and failures, the number of requests in the queue, and average time that requests spend in the queue. You can view these metrics in the Amazon GameLift Servers console or in CloudWatch.

To get started by creating a basic starter queue, see [Create a game session queue](#).

Topics

- [Best practices for Amazon GameLift Servers game session queues](#)
- [Create a game session queue](#)
- [Set up event notification for game session placement](#)

Best practices for Amazon GameLift Servers game session queues

A game session queue contains a list of fleets where Amazon GameLift Servers can place new game sessions. Each fleet can have hosting resources deployed in multiple geographic locations. When choosing a placement, the queue selects a fleet and a fleet location based on a set of priorities that you set for the fleet.

Consider the following guidelines and best practices:

- **Add fleets in locations that cover your players.** You can add fleets and aliases in any available location. Location is important if you're making placements based on reported player latency.
- **Use aliases for all fleets.** Assign an alias to each fleet in a queue, and use the alias names when setting destinations in your queue.
- **Use the same or a similar game build or script for all fleets.** The queue might put players into game sessions on any fleet in the queue. Players must be able to play in any game session on any fleet.

- **Create fleets in at least two locations.** By having game servers hosted in at least one other location, you mitigate the impact of Regional outages on your players. You can keep your backup fleets scaled down, and use auto scaling to increase capacity if usage increases.
- **Prioritize your game session placement.** A queue prioritizes placement choices based on several elements, including destination list order.
- **Create your queue in the same location as your client service.** By putting your queue in a location near your client service, you can minimize communication latency.
- **Use fleets with multiple locations.** Use the queue filter configuration to prevent the queue from placing game sessions in specified locations. You can use at least two multi-location fleets with different home locations to mitigate the impact of game placements during a Regional outage.

Create a game session queue

Queues are used to place new game sessions across multiple fleets and locations. Your game starts new game sessions by submitting placement requests to a queue. A queue is configured with instructions for how to process requests. Learn more about initiating game session placement requests in [Create game sessions](#).

To create a game session queue

These instructions illustrate how to create a simple working queue with minimal configuration settings and default settings. There are a number of options for customizing a queue configuration. These options help you make the best possible placements based on the needs of your game. To learn more about customizing queues for your game, see [Customize a game session queue](#). You can update most queue configuration settings at any time.

You can create a game session queue using either the Amazon GameLift Servers console or the Amazon CLI.

Console

In the [Amazon GameLift Servers console](#), select an Amazon Region to work in. Open the console's left navigation bar and choose **Queues**.

1. On the **Queues** page, choose **Create queue** to start the workflow.
2. Under **Queue settings** enter the following settings:

- a. Enter a queue name. This name must be unique to the Amazon Web Services Region that you're creating the queue in.
 - b. Keep the default **Timeout** setting, which is 600 seconds (or 10 minutes). This value controls how long Amazon GameLift Servers tries to place a new game session before stopping. Amazon GameLift Servers searches for available resources until the request times out. You can update a queue's timeout setting at any time.
 - c. Skip the **Player latency policies** section. A queue uses latency policies only when it receives placement requests that include player latency data. You can add latency policies to a queue at any time. For more on creating latency policies, see [Create a player latency policy](#).
3. Skip the **Game session placement locations** section to use the default setting of **All locations**. This setting lets you create an allow list of locations where the queue can make placements (also called a filter configuration). For more about prioritizing by location and filter configurations, see [Prioritize placements by location](#).
4. Under **Destination order**, add one or more fleets to the queue. You can identify fleets by using fleet IDs or ARNs, or by using a fleet alias. When adding multiple fleets, keep in mind that they should all be running similar game builds and be compatible with any game client that uses this queue. In addition, all fleets in a queue must have the same certificate configuration.
 - a. Select the **Region** where the fleet or alias was created. For a multi-location fleet, this is the "home" region.
 - b. For destination **Type**, select either a fleet or an alias.
 - c. Your region and type selections populate a dropdown list of existing fleets or aliases. Select one to designate as a queue destination.
 - d. To specify another fleet or alias for the queue, choose **Add destination** and repeat the previous steps.
 - e. After you've added a list of destinations, use the drag-and-drop feature to reorder the destinations. Amazon GameLift Servers uses this order when prioritizing placements by destination.
5. Skip the **Game session placement priority** section to keep the default priority order. This setting lets you customize how Amazon GameLift Servers chooses where to look for available hosting resources for new game session placements. For more information about

prioritizing placements, see [Prioritize game session placement](#). You can update a queue's placement priorities at any time.

6. Under **Location order**, keep the default values. This setting is used when prioritizing by fleet location. It provides location order to use. When using the default priority settings, location is used as a tiebreaker when the preferred destination is a fleet with multiple locations.
7. Skip the optional **Event notification settings** section. Event notifications are required for queues that process a high volume of placement requests. For queues that process low volumes, such as for development or testing purposes, you can track the status of placement requests by polling with [DescribeGameSessionPlacement](#). For more details, see [Set up event notification for game session placement](#). You can update a queue's event notification settings at any time.
8. Choose **Create** to generate a new queue with minimal customization.

Amazon CLI

Example Create a queue

The following example creates a game session queue with these configurations:

- A five minute timeout.
- Two fleet destinations.
- Filter to only allow placements in these locations: us-east-1, us-east-2, us-west-2, and ca-central-1.
- Priority order based on cost and then locations in a specified order.

```
aws gamelift create-game-session-queue \  
  --name "sample-test-queue" \  
  --timeout-in-seconds 300 \  
  --destinations DestinationArn="arn:aws:gamelift:us-east-1:111122223333:fleet/  
fleet-772266ba-8c82-4a6e-b620-a74a62a93ff8" DestinationArn="arn:aws:gamelift:us-  
east-1:111122223333:fleet/fleet-33f28fb6-aa8b-4867-85b4-ceb217bf5994" \  
  --filter-configuration "AllowedLocations=us-east-1, ca-central-1, us-east-2, us-  
west-2" \  
  --priority-configuration PriorityOrder="COST","LOCATION",LocationOrder="us-  
east-1","us-east-2","ca-central-1","us-west-2" \  
  --notification-target "arn:aws:sns:us-east-1:111122223333:gamelift-test.fifo"
```

Note

You can get fleet and alias ARN values by calling either [describe-fleet-attributes](#) or [describe-alias](#) with the fleet or alias ID.

If the `create-game-session-queue` request is successful, Amazon GameLift Servers returns a [GameSessionQueue](#) object with the new queue configuration. You can now submit requests to the queue using [StartGameSessionPlacement](#).

Example Create a queue with player latency policies

The following example creates a game session queue with these configurations:

- A ten minute timeout
- Three fleet destinations
- A set of player latency policies

```
aws gamelift create-game-session-queue \  
  --name "matchmaker-queue" \  
  --timeout-in-seconds 600 \  
  --destinations DestinationArn=arn:aws-cn:gamelift:us-east-1::alias/alias-  
a1234567-b8c9-0d1e-2fa3-b45c6d7e8910 \  
                DestinationArn=arn:aws-cn:gamelift:us-west-2::alias/alias-b0234567-  
c8d9-0e1f-2ab3-c45d6e7f8901 \  
                DestinationArn=arn:aws-cn:gamelift:us-west-2::fleet/fleet-f1234567-  
b8c9-0d1e-2fa3-b45c6d7e8912 \  
  --player-latency-policies  
    "MaximumIndividualPlayerLatencyMilliseconds=50,PolicyDurationSeconds=120" \  
  
    "MaximumIndividualPlayerLatencyMilliseconds=100,PolicyDurationSeconds=120" \  
    "MaximumIndividualPlayerLatencyMilliseconds=150" \  

```

If the `create-game-session-queue` request is successful, Amazon GameLift Servers returns a [GameSessionQueue](#) object with the new queue configuration.

Set up event notification for game session placement

You can use event notifications to monitor the status of individual placement requests. We recommend setting up event notifications for all games with high-volume placement activity.

There are two options for setting up event notifications.

- Have Amazon GameLift Servers publish event notifications to an Amazon Simple Notification Service (Amazon SNS) topic using a queue.
- Use automatically published Amazon EventBridge events and its suite of tools for managing events.

For a list of game session placement events emitted by Amazon GameLift Servers, see [Game session placement events](#).

Important

For high-volume placement systems, we recommend using standard (non-FIFO) Amazon SNS topics rather than FIFO topics. FIFO topics have lower publishing limits than standard topics, which can lead to throttling exceptions during high load. If you experience throttling with FIFO topics, you may lose queue placement notifications.

Set up an SNS topic

For Amazon GameLift Servers to publish all events generated by a game session queue to a topic, set the notification target field to a topic.

To set up an SNS topic for Amazon GameLift Servers event notification

1. Sign in to the Amazon Web Services Management Console and open the Amazon SNS console at <https://console.amazonaws.cn/sns/v3/home>.
2. From the SNS **Topics** page, choose **Create topic** and follow the instructions to create your topic.
3. Under **Access policy**, do the following:
 - a. Choose the **Advanced** method.
 - b. Add the following bolded section of the JSON object to the existing policy.

JSON

```
{
  "Version": "2012-10-17",
  "Id": "__default_policy_ID",
  "Statement": [
    {
      "Sid": "__default_statement_ID",
      "Effect": "Allow",
      "Principal": {
        "AWS": "*"
      },
      "Action": [
        "SNS:GetTopicAttributes",
        "SNS:SetTopicAttributes",
        "SNS:AddPermission",
        "SNS:RemovePermission",
        "SNS:DeleteTopic",
        "SNS:Subscribe",
        "SNS:ListSubscriptionsByTopic",
        "SNS:Publish"
      ],
      "Resource": "arn:aws-cn:sns:us-  
east-1:111122223333:your_topic_name",
      "Condition": {
        "StringEquals": {
          "AWS:SourceAccount": "your_account"
        }
      }
    },
    {
      "Sid": "__console_pub_0",
      "Effect": "Allow",
      "Principal": {
        "Service": "gamelift.amazonaws.com"
      },
      "Action": "sns:Publish",
      "Resource": "arn:aws-cn:sns:us-  
east-1:111122223333:your_topic_name",
      "Condition": {
        "ArnLike": {
```

```

        "aws:SourceArn": "arn:aws-cn:gamelift:us-
east-1:111122223333:gamesessionqueue/your_queue_name"
    }
}
]
}

```

- c. (Optional) Add additional access control to the topic by adding conditions to the resource policy.
4. Choose **Create topic**.
5. After you've created your SNS topic, add it to queues during queue creation, or edit an existing queue to add it.

Set up an SNS topic with server-side encryption

With server-side encryption (SSE), you can store sensitive data in encrypted topics. SSE protects the contents of messages in Amazon SNS topics using keys that are managed in Amazon Key Management Service (Amazon KMS). For more information about server-side encryption with Amazon SNS, see [Encryption at rest](#) in the *Amazon Simple Notification Service Developer Guide*.

To set up an SNS topic with server-side encryption, review the following topics:

- [Creating key](#) in the *Amazon Key Management Service Developer Guide*
- [Enabling SSE for a topic](#) in the *Amazon Simple Notification Service Developer Guide*

When creating your KMS key, use the following KMS key policy:

```

{
  "Effect": "Allow",
  "Principal": {
    "Service": "gamelift.amazonaws.com"
  },
  "Action": [
    "kms:Decrypt",
    "kms:GenerateDataKey"
  ],
  "Resource": "*",
  "Condition": {
    "ArnLike": {

```

```

    "aws:SourceArn":
      "arn:aws:gamelift:your_region:your_account:gamesessionqueue/your_queue_name"
    },
    "StringEquals": {
      "kms:EncryptionContext:aws:sns:topicArn":
        "arn:aws:sns:your_region:your_account:your_sns_topic_name"
    }
  }
}

```

Set up EventBridge

Amazon GameLift Servers automatically posts all game session placement events to EventBridge. With EventBridge you can set up rules to have events routed to targets for processing. For example, you can set a rule to route the event `PlacementFulfilled` to an Amazon Lambda function that handles tasks that precede connecting to a game session. For more information about EventBridge, see [What is Amazon EventBridge?](#) in the *Amazon EventBridge User Guide*.

The following are some examples of EventBridge rules to use with Amazon GameLift Servers queues:

Matches events from all Amazon GameLift Servers queues

```

{
  "source": [
    "aws.gamelift"
  ],
  "detail-type": [
    "GameLift Queue Placement Event"
  ]
}

```

Matches events from a specific queue

```

{
  "source": [
    "aws.gamelift"
  ],
  "detail-type": [
    "GameLift Queue Placement Event"
  ],
  "resources": [

```



```
    "arn:aws:gamelift:your_region:your_account:gamesessionqueue/your_queue_name"
  ]
}
```

Reference implementation with Amazon CDK

For a complete reference implementation of event-based game session placement using Amazon SNS, Amazon Lambda, and Amazon DynamoDB, see the [Event-based game session placement guidance](#) in the Amazon GameLift Toolkit. This guidance includes an Amazon CDK template for deploying a complete event-driven game session placement system with an Amazon Lambda function for processing game session placement events, an Amazon DynamoDB table for tracking game session placement state, and best practices for Amazon GameLift Servers queue configuration.

Customize to your game hosting solution

With a basic game hosting solution in place, use the following topics to customize and enhance it to improve the player experience, optimize costs, and add advanced functionality. This section covers various customization options organized by the component they primarily affect. Choose the customizations that best fit your game's requirements and player base.

Topics

- [Game server build customizations](#)
 - [Connect your Amazon GameLift Servers hosted game server to other Amazon resources](#)
 - [Let your game server access Amazon GameLift Servers fleet data](#)
 - [Set up VPC peering for Amazon GameLift Servers](#)
- [Player sessions and matchmaking customizations](#)
 - [Generate player IDs](#)
 - [Add FlexMatch matchmaking to Amazon GameLift Servers](#)
- [Game session placement customizations](#)
 - [Customize a game session queue](#)
 - [Prioritize game session placement](#)
 - [Build a queue for Spot Instances](#)
- [Hosting resource customizations](#)
 - [Choose compute resources for a managed fleet](#)

- [Customize an Amazon GameLift Servers container fleet](#)
- [Reduce game hosting costs with Spot fleets](#)
- [Optimize game server runtime configuration on managed Amazon GameLift Servers](#)
- [Work with the Amazon GameLift Servers Agent](#)
- [Abstract an Amazon GameLift Servers fleet designation with an alias](#)

Game server build customizations

Game server build customizations offer opportunities to enhance functionality to the game server, including taking advantage of other Amazon services. These customizations extend your game server capabilities beyond basic hosting to support advanced features and integrations.

Connect your Amazon GameLift Servers hosted game server to other Amazon resources

When you're creating a game server build for deployment on Amazon GameLift Servers fleets, you might want the applications in your game build to communicate directly and securely with other Amazon resources that you own. Because Amazon GameLift Servers manages your game hosting fleets, you must give Amazon GameLift Servers limited access to these resources and services.

Some example scenarios include:

- Use an Amazon CloudWatch agent to collect metrics, logs, and traces from managed EC2 fleets.
- Send instance log data to Amazon CloudWatch Logs.
- Obtain game files stored in an Amazon Simple Storage Service (Amazon S3) bucket.
- Read and write game data (such as game modes or inventory) stored in an Amazon DynamoDB database or other data storage service.
- Send signals directly to an instance using Amazon Simple Queue Service (Amazon SQS).
- Access custom resources that are deployed and running on Amazon Elastic Compute Cloud (Amazon EC2).

Amazon GameLift Servers supports these methods for establishing access:

- [Access Amazon resources with an IAM role](#)
- [Access Amazon resources with VPC peering](#)

Access Amazon resources with an IAM role

Use an IAM role to specify who can access your resources and set limits on that access. Trusted parties can "assume" a role and get temporary security credentials that authorize them to interact with the resources. When the parties make API requests related to the resource, they must include the credentials.

To set up access controlled by an IAM role, do the following tasks:

1. [Create the IAM role](#)
2. [Modify applications to acquire credentials](#)
3. [Associate a fleet with the IAM role](#)

Create the IAM role

In this step, you create an IAM role, with a set of permissions to control access to your Amazon resources and a trust policy that gives Amazon GameLift Servers rights to use the role's permissions.

For instructions on how to set up the IAM role , see [Set up an IAM service role for Amazon GameLift Servers](#). When creating the permissions policy, choose specific services, resources, and actions that your applications need to work with. As a best practice, limit the scope of the permissions as much as possible.

After you create the role, take note of the role's Amazon Resource Name (ARN). You need the role ARN during fleet creation.

Modify applications to acquire credentials

In this step, you configure your applications to acquire security credentials for the IAM role and use them when interacting with your Amazon resources . See the following table to determine how to modify your applications based on (1) the type of application, and (2) the server SDK version your game uses to communicate with Amazon GameLift Servers.

	Game server applications	Other applications
Using server SDK version 4 or earlier	Call Amazon Security Token Service (Amazon STS) AssumeRole with the role ARN.	Call Amazon Security Token Service (Amazon STS) AssumeRole with the role ARN.

Use AssumeRole() (server SDK 4)

Add code to your applications to assume the IAM role and get credentials to interact with your Amazon resources. Any application that runs on an Amazon GameLift Servers fleet instance with server SDK 4 or earlier can assume the IAM role.

In the application code, before accessing an Amazon resource, the application must call the Amazon Security Token Service (Amazon STS) [AssumeRole](#) API operation and specify the role ARN. This operation returns a set of temporary credentials that authorizes the application to access to the Amazon resource. For more information, see [Using temporary credentials with Amazon resources](#) in the *IAM User Guide*.

Associate a fleet with the IAM role

After you've created the IAM role and updated the applications in your game server build to get and use the access credentials, you can deploy a fleet. When you configure the new fleet, set the following parameters:

- [InstanceRoleArn](#) – Set this parameter to the ARN of the IAM role.
- [InstanceRoleCredentialsProvider](#) – To prompt Amazon GameLift Servers to generate a shared credentials file for each fleet instance, set this parameter to SHARED_CREDENTIAL_FILE.

You must set these values when you create the fleet. They can't be updated later.

Access Amazon resources with VPC peering

You can use Amazon Virtual Private Cloud (Amazon VPC) peering to communicate between applications running on a Amazon GameLift Servers instance and another Amazon resource. A VPC is a virtual private network that you define that includes a set of resources managed through your Amazon Web Services account. Each Amazon GameLift Servers fleet has its own VPC. With VPC peering, you can establish a direct network connection between the VPC for your fleet and for your other Amazon resources.

Amazon GameLift Servers streamlines the process of setting up VPC peering connections for your game servers. It handles peering requests, updates route tables, and configures the connections as required. For instructions about how to set up VPC peering for your game servers, see [Set up VPC peering for Amazon GameLift Servers](#).

Let your game server access Amazon GameLift Servers fleet data

There are some situations where your custom game build or Amazon GameLift Servers Realtime script may require information about the Amazon GameLift Servers fleet. For example, your game build or script might include code to:

- Monitor activity based on fleet data.
- Roll up metrics to track activity by fleet data. (Many games use this data for LiveOps activities.)
- Provide relevant data to custom game services, such as for matchmaking, additional capacity scaling, or testing.

Fleet information is available as a JSON file on each instance in the following locations:

- Windows: `C:\GameMetadata\gamelift-metadata.json`
- Linux: `/local/gamemetadata/gamelift-metadata.json`

The `gamelift-metadata.json` file includes the [attributes of an Amazon GameLift Servers fleet resource](#).

Example JSON file:

```
{
  "buildArn": "arn:aws:gamelift:us-west-2:123456789012:build/build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "buildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",
  "fleetArn": "arn:aws:gamelift:us-west-2:123456789012:fleet/fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "fleetDescription": "Test fleet for Really Fun Game v0.8",
  "fleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",
  "name": "ReallyFunGameTestFleet08",
  "fleetType": "ON_DEMAND",
  "instanceRoleArn": "arn:aws:iam::123456789012:role/S3AccessForGameLift",
  "instanceType": "c5.large",
  "serverLaunchPath": "/local/game/reallyfungame.exe"
}
```

Set up VPC peering for Amazon GameLift Servers

This topic provides guidance on how to set up a VPC peering connection between your Amazon GameLift Servers-hosted game servers and your other non-Amazon GameLift Servers resources. Use Amazon Virtual Private Cloud (VPC) peering connections to enable your game servers to communicate directly and privately with your other Amazon resources, such as a web service or a repository. You can establish VPC peering with any resources that run on Amazon and are managed by an Amazon account that you have access to.

Note

VPC peering is an advanced feature. To learn about preferred options for enabling your game servers to communicate directly and privately with your other Amazon resources, see [Connect your Amazon GameLift Servers hosted game server to other Amazon resources](#).

If you're already familiar with Amazon VPCs and VPC peering, understand that setting up peering with Amazon GameLift Servers game servers is somewhat different. You don't have access to the VPC that contains your game servers—it is controlled by the Amazon GameLift Servers service—so you can't directly request VPC peering for it. Instead, you first pre-authorize the VPC with your non-Amazon GameLift Servers resources to accept a peering request from the Amazon GameLift Servers service. Then you trigger Amazon GameLift Servers to request the VPC peering that you just authorized. Amazon GameLift Servers handles the tasks of creating the peering connection, setting up the route tables, and configuring the connection.

To set up VPC peering for an existing fleet

1. Get Amazon account ID(s) and credentials.

You need an ID and sign-in credentials for the following Amazon accounts. You can find Amazon account IDs by signing into the [Amazon Web Services Management Console](#) and viewing your account settings. To get credentials, go to the IAM console.

- Amazon account that you use to manage your Amazon GameLift Servers game servers.
- Amazon account that you use to manage your non-Amazon GameLift Servers resources.

If you're using the same account for Amazon GameLift Servers and non-Amazon GameLift Servers resources, you need ID and credentials for that account only.

2. Get identifiers for each VPC.

Get the following information for the two VPCs to be peered:

- VPC for your Amazon GameLift Servers game servers – This is your Amazon GameLift Servers fleet ID. Your game servers are deployed in Amazon GameLift Servers on a fleet of EC2 instances. A fleet is automatically placed in its own VPC, which is managed by the Amazon GameLift Servers service. You don't have direct access to the VPC, so it is identified by the fleet ID.
- VPC for your non-Amazon GameLift Servers Amazon resources – You can establish a VPC peering with any resources that run on Amazon and are managed by an Amazon account that you have access to. If you haven't already created a VPC for these resources, see [Getting started with Amazon VPC](#). Once you have created a VPC, you can find the VPC ID by signing into the [Amazon Web Services Management Console](#) for Amazon VPC and viewing your VPCs.

Note

When setting up a peering, both VPCs must exist in the same region. The VPC for your Amazon GameLift Servers fleet game servers is in the same region as the fleet.

3. Authorize a VPC peering.

In this step, you are pre-authorizing a future request from Amazon GameLift Servers to peer the VPC with your game servers with your VPC for non-Amazon GameLift Servers resources. This action updates the security group for your VPC.

To authorize the VPC peering, call the service API [CreateVpcPeeringAuthorization\(\)](#) or use the Amazon CLI command `create-vpc-peering-authorization`. Make this call using the account that manages your non-Amazon GameLift Servers resources. Identify the following information:

- Peer VPC ID – This is for the VPC with your non-Amazon GameLift Servers resources.
- Amazon GameLift Servers Amazon account ID – This is the account that you use to manage your Amazon GameLift Servers fleet.

Once you've authorized a VPC peering, the authorization remains valid for 24 hours unless revoked. You can manage your VPC peering authorizations using the following operations:

- [DescribeVpcPeeringAuthorizations\(\)](#) (Amazon CLI `describe-vpc-peering-authorizations`).
- [DeleteVpcPeeringAuthorization\(\)](#) (Amazon CLI `delete-vpc-peering-authorization`).

4. Request a peering connection.

With a valid authorization, you can request that Amazon GameLift Servers establish a peering connection.

To request a VPC peering, call the service API [CreateVpcPeeringConnection\(\)](#) or use the Amazon CLI command `create-vpc-peering-connection`. Make this call using the account that manages your Amazon GameLift Servers game servers. Use the following information to identify the two VPCs that you want to peer:

- Peer VPC ID and Amazon account ID – This is the VPC for your non-Amazon GameLift Servers resources and the account that you use to manage them. The VPC ID must match the ID on a valid peering authorization.
- Fleet ID – This identifies the VPC for your Amazon GameLift Servers game servers.

5. Track the peering connection status.

Requesting a VPC peering connection is an asynchronous operation. To track the status of a peering request and handle success or failure cases, use one of the following options:

- Continuously poll with `DescribeVpcPeeringConnections()`. This operation retrieves the VPC peering connection record, including the status of the request. If a peering connection is successfully created, the connection record also contains a CIDR block of private IP addresses that is assigned to the VPC.
- Handle fleet events associated with VPC peering connections with [DescribeFleetEvents\(\)](#), including success and failure events.

Once the peering connection is established, you can manage it using the following operations:

- [DescribeVpcPeeringConnections\(\)](#) (Amazon CLI `describe-vpc-peering-connections`).
- [DeleteVpcPeeringConnection\(\)](#) (Amazon CLI `delete-vpc-peering-connection`).

To set up VPC peering with a new fleet

You can create a new Amazon GameLift Servers fleet and request a VPC peering connection at the same time.

1. Get Amazon account ID(s) and credentials.

You need an ID and sign-in credentials for the following two Amazon accounts. You can find Amazon account IDs by signing into the [Amazon Web Services Management Console](#) and viewing your account settings. To get credentials, go to the IAM console.

- Amazon account that you use to manage your Amazon GameLift Servers game servers.
- Amazon account that you use to manage your non-Amazon GameLift Servers resources.

If you're using the same account for Amazon GameLift Servers and non-Amazon GameLift Servers resources, you need ID and credentials for that account only.

2. Get the VPC ID for your non-Amazon GameLift ServersAmazon resources.

If you haven't already created a VPC for these resources, do so now (see [Getting started with Amazon VPC](#)). Be sure that you create the new VPC in the same region where you plan to create your new fleet. If your non-Amazon GameLift Servers resources are managed under a different Amazon account or user/user group than the one you use with Amazon GameLift Servers, you'll need to use these account credentials when requesting authorization in the next step.

Once you have created a VPC, you can locate the VPC ID in Amazon VPC console by viewing your VPCs.

3. Authorize a VPC peering with non-Amazon GameLift Servers resources.

When Amazon GameLift Servers creates the new fleet and a corresponding VPC, it also sends a request to peer with the VPC for your non-Amazon GameLift Servers resources. You need to pre-authorize that request. This step updates the security group for your VPC.

Using the account credentials that manage your non-Amazon GameLift Servers resources, call the service API [CreateVpcPeeringAuthorization\(\)](#) or use the Amazon CLI command `create-vpc-peering-authorization`. Identify the following information:

- Peer VPC ID – ID of the VPC with your non-Amazon GameLift Servers resources.

- Amazon GameLift Servers Amazon account ID – ID of the account that you use to manage your Amazon GameLift Servers fleet.

Once you've authorized a VPC peering, the authorization remains valid for 24 hours unless revoked. You can manage your VPC peering authorizations using the following operations:

- [DescribeVpcPeeringAuthorizations\(\)](#) (Amazon CLI `describe-vpc-peering-authorizations`).
 - [DeleteVpcPeeringAuthorization\(\)](#) (Amazon CLI `delete-vpc-peering-authorization`).
4. Follow the instructions for [creating a new fleet using the Amazon CLI](#). Include the following additional parameters:
- *peer-vpc-aws-account-id* – ID for the account that you use to manage the VPC with your non-Amazon GameLift Servers resources.
 - *peer-vpc-id* – ID of the VPC with your non-Amazon GameLift Servers account.

A successful call to [create-fleet](#) with the VPC peering parameters generates both a new fleet and a new VPC peering request. The fleet's status is set to **New** and the fleet activation process is initiated. The peering connection request's status is set to **initiating-request**. You can track the success or failure of the peering request by calling [describe-vpc-peering-connections](#).

When requesting both a new fleet and a VPC peering connection, both actions either succeed or fail. If a fleet fails during the creation process, the VPC peering connection will not be established. Likewise, if a VPC peering connection fails for any reason, the new fleet will fail to move from status **Activating** to **Active**.

Note

The new VPC peering connection is not completed until the fleet is ready to become active. This means that the connection is not available and can't be used during the game server build installation process.

The following example creates both a new fleet and a peering connection between a pre-established VPC and the VPC for the new fleet. The pre-established VPC is uniquely identified by the combination of your non-Amazon GameLift Servers Amazon account ID and the VPC ID.

```
$ Amazon gamelift create-fleet
  --name "My_Fleet_1"
  --description "The sample test fleet"
  --ec2-instance-type "c5.large"
  --fleet-type "ON_DEMAND"
  --build-id "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
  --runtime-configuration "GameSessionActivationTimeoutSeconds=300,
                           MaxConcurrentGameSessionActivations=2,
                           ServerProcesses=[{LaunchPath=C:\game\Bin64.dedicated
\MultiplayerSampleProjectLauncher_Server.exe,
                                           Parameters=+sv_port 33435 +start_lobby,
                                           ConcurrentExecutions=10}]"
  --new-game-session-protection-policy "FullProtection"
  --resource-creation-limit-policy "NewGameSessionsPerCreator=3,
                                   PolicyPeriodInMinutes=15"
  --ec2-inbound-permissions
"FromPort=33435,ToPort=33435,IpRange=0.0.0.0/0,Protocol=UDP"

"FromPort=33235,ToPort=33235,IpRange=0.0.0.0/0,Protocol=UDP"
  --metric-groups "EMEAfleets"
  --peer-vpc-aws-account-id "111122223333"
  --peer-vpc-id "vpc-a11a11a"
```

Copyable version:

```
Amazon gamelift create-fleet --name "My_Fleet_1" --description
"The sample test fleet" --fleet-type "ON_DEMAND" --metric-groups
"EMEAfleets" --build-id "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff"
--ec2-instance-type "c5.large" --runtime-configuration
"GameSessionActivationTimeoutSeconds=300,MaxConcurrentGameSessionActivations=2,ServerProcesses
\game\Bin64.dedicated\MultiplayerSampleProjectLauncher_Server.exe,Parameters=
+sv_port 33435 +start_lobby,ConcurrentExecutions=10}]" --new-game-session-
protection-policy "FullProtection" --resource-creation-limit-policy
"NewGameSessionsPerCreator=3,PolicyPeriodInMinutes=15" --ec2-inbound-
permissions "FromPort=33435,ToPort=33435,IpRange=0.0.0.0/0,Protocol=UDP"
"FromPort=33235,ToPort=33235,IpRange=0.0.0.0/0,Protocol=UDP" --peer-vpc-aws-account-id
"111122223333" --peer-vpc-id "vpc-a11a11a"
```

Troubleshooting VPC peering issues

If you're having trouble establishing a VPC peering connection for your Amazon GameLift Servers game servers, consider these common root causes:

- An authorization for the requested connection was not found:
 - Check the status of a VPC authorization for the non-Amazon GameLift Servers VPC. It might not exist or it might have expired.
 - Check the regions of the two VPCs you're trying to peer. If they're not in the same region, they can't be peered.
- The CIDR blocks (see [Invalid VPC peering connection configurations](#)) of your two VPCs are overlapping. The IPv4 CIDR blocks that are assigned to peered VPCs cannot overlap. The CIDR block of the VPC for your Amazon GameLift Servers fleet is automatically assigned and can't be changed, so you'll need to change the CIDR block for of the VPC for your non-Amazon GameLift Servers resources. To resolve this issue:
 - Look up this CIDR block for your Amazon GameLift Servers fleet by calling `DescribeVpcPeeringConnections()`.
 - Go to the Amazon VPC console, find the VPC for your non-Amazon GameLift Servers resources, and change the CIDR block so that they don't overlap.
- The new fleet did not activate (when requesting VPC peering with a new fleet). If the new fleet failed to progress to **Active** status, there is no VPC to peer with, so the peering connection cannot succeed.

Player sessions and matchmaking customizations

Player session and matchmaking customizations give you the opportunity to develop sophisticated player management workflows, including a nuanced matchmaking system that helps you deliver balanced, engaging multiplayer experiences.

Generate player IDs

Amazon GameLift Servers uses a player session to represent a player connected to a game session. Amazon GameLift Servers creates a player session each time a player connects to a game session using a game client integrated with Amazon GameLift Servers. When a player leaves a game, the player session ends. Amazon GameLift Servers doesn't reuse player sessions.

Important

When using FlexMatch matchmaking, if you create a new matchmaking request that contains a player ID that is already included in an existing active matchmaking request, the existing request is automatically cancelled. However, a `MatchmakingCancelled`

event is not sent for the cancelled request. To monitor the status of existing matchmaking requests, use [DescribeMatchmaking](#) to poll the request status at infrequent intervals (30-60 seconds). The cancelled request will show a status of CANCELLED with the reason Cancelled due to duplicate player.

The following code example randomly generates unique player IDs:

```
bool includeBrackets = false;
bool includeDashes = true;
string playerId = AZ::Uuid::CreateRandom().ToString<string>(includeBrackets,
    includeDashes);
```

For more information about player sessions, see [Game and player sessions in the Amazon GameLift Servers console](#).

Add FlexMatch matchmaking to Amazon GameLift Servers

Use Amazon GameLift Servers FlexMatch to add player matchmaking functionality to your Amazon GameLift Servers hosted games. You can use FlexMatch with either custom game servers or Amazon GameLift Servers Realtime.

FlexMatch pairs the matchmaking service with a customizable rules engine. You design how to match players together based on player attributes and game modes that make sense for your game. FlexMatch manages the nuts and bolts of evaluating players who are looking for a game, forming matches with one or more teams, and starting game sessions to host the matches.

To use the full FlexMatch service, you must have your hosting resources set up with queues. Amazon GameLift Servers uses queues to locate the best possible hosting locations for games across multiple regions and computing types. In particular, Amazon GameLift Servers queues can use latency data, when provided by game clients, to place game sessions so that players experience the lowest possible latency when playing.

For more information on FlexMatch including detailed help with integrating matchmaking into your games, see these [Amazon GameLift Servers FlexMatch Developer Guide](#) topics:

- [How Amazon GameLift Servers FlexMatch works](#)
- [FlexMatch integration steps](#)

Game session placement customizations

Game session placement customizations give you the opportunity to fine-tune your placement system to give your player base the best possible gameplay experience. You can optimize placement for player compatibility and preferences, as well as for operational considerations such as cost efficiency, geographic distribution, and resilience against service disruptions.

Customize a game session queue

This topic describes how to customize your game session queues to make the best possible decisions about game session placement. For more information about game session queues and how they work, see [Configure game session placement](#).

These Amazon GameLift Servers features require queues:

- [Matchmaking with FlexMatch](#)
- [Build a queue for Spot Instances](#)

Topics

- [Define a queue's scope](#)
- [Build a multi-location queue](#)
- [Evaluate queue metrics](#)

Define a queue's scope

Your game's player population might have groups of players who shouldn't play together. For example, if you publish your game in two languages each language should have its own game servers.

To set up game session placement for your player population, create a separate queue for each player segment. Scope each queue to place players into the correct game servers. Some common ways to scope queues include:

- **By geographic locations.** When deploying your game servers in multiple geographic areas, you might build queues for players in each location to reduce player latency.
- **By build or script variations.** If you have more than one variation of your game server, you might be supporting player groups that can't play in the same game sessions. For example, game server builds or scripts might support different languages or device types.

- **By event types.** You might create a special queue to manage games for participants in tournaments or other special events.

Design multiple queues

Depending on your game and players, you might want to create more than one game session queue. When your game client service requests a new game session, it specifies which game session queue to use. To help you determine whether to use multiple queues, consider:

- **Variations of your game server.** You can create a separate queue for each variation of your game server. All fleets in a queue must deploy compatible game servers. This is because players who use the queue to join games must be able to play on any of the queue's game servers.
- **Different player groups.** You can customize how Amazon GameLift Servers places game sessions based on player group. For example, you might need queues customized for certain game modes that require a special instance type or runtime configuration. Or, you might want a special queue to manage placements for a tournament or other event.
- **Game session queue metrics.** You can set up queues based on how you want to collect game session placement metrics. For more information, see [Amazon GameLift Servers metrics for queues](#).

Build a multi-location queue

We recommend a multi-location design for all queues. This design can improve placement speed and hosting resiliency. A multi-location design is required to use player latency data to put players into game sessions with minimal latency. If you're building multi-location queues that use Spot Instance fleets, follow the instructions in [Reduce game hosting costs with Spot fleets](#).

One way to create a multi-location queue is to add a [multi-location fleet](#) to a queue. That way, the queue can place game sessions in any of the fleet's locations. You can also add other fleets with different configurations or home locations for redundancy. If you're using a multi-location Spot Instance fleet, follow best practices and include an On-Demand Instance fleet with the same locations.

The following example outlines the process of designing a basic multi-location queue. In this example, we use two fleets: one Spot Instance fleet and one On-Demand Instance fleet. Each fleet has the following Amazon Web Services Regions for placement locations: us-east-1, us-east-2, ca-central-1, and us-west-2.

To create a basic multi-location queue with multi-location fleets

1. Choose a location to create the queue in. You can minimize request latency by placing the queue in a location near where you deployed the client service. In this example, we create the queue in `us-east-1`.
2. Create a new queue and add your multi-location fleets as queue destinations. The destination order determines how Amazon GameLift Servers places game sessions. In this example, we list the Spot Instance fleet first and the On-Demand Instance fleet second.
3. Define the queue's game session placement priority order. This order determines where the queue searches first for an available game server. In this example, we use the default priority order.
4. Define the location order. If you don't define the location order, Amazon GameLift Servers uses the locations in alphabetical order.

Game session placement locations

Locations where the queue can place new game sessions.

Locations

Choose locations ▼

ca-central-1
Canada (Central) ✕

us-west-2
US West (Oregon) ✕

us-east-2
US East (Ohio) ✕

us-east-1
US East (N. Virginia) ✕

Destination order

An ordered list of fleets and aliases that the queue can use for game session placement.

	Region	Type	Name	
⋮	us-east-1 ▼	Fleet ▼	TestFleet-SPOT ▼	Remove
⋮	us-east-1 ▼	Fleet ▼	TestFleet-ONDEMAND ▼	Remove

Add Destination

Game session placement priority

The values that GameLift uses to prioritize game session placement. The default order is latency, cost, destination, and location.



Latency

Prioritize locations with the lowest average player latency.



Cost

Prioritize destinations with the lowest current hosting cost.



Destination

Prioritize based on the defined destination order.



Location

Prioritize based on the defined location order.

▼ Location order

An ordered list of locations that the queue can use for game session placement.

Location

ca-central-1



Remove

us-east-1



Remove

us-east-2



Remove

us-west-2



Remove

Add locations

Evaluate queue metrics

Use metrics to evaluate how well your queues are performing. You can view metrics related to queues in the [Amazon GameLift Servers console](#) or in Amazon CloudWatch. For a list and descriptions of queue metrics, see [Amazon GameLift Servers metrics for queues](#).

Queue metrics can provide insight about the following:

- **Overall queue performance** – Queue metrics indicate how successfully a queue responds to placement requests. These metrics can also help you identify when and why placements fail. For queues with manually scaled fleets, the `AverageWaitTime` and `QueueDepth` metrics can indicate when you should adjust capacity for a queue.
- **FleetIQ algorithm performance** – For placement requests using the FleetIQ algorithm, metrics show how often the algorithm finds ideal game session placement. The placement may prioritize using resources with the lowest player latency or resources with the lowest cost. There are also error metrics that identify common reasons why Amazon GameLift Servers can't find an ideal placement. For more information about metrics, see [Monitor Amazon GameLift Servers with Amazon CloudWatch](#).
- **Location specific placements** – For multi-location queues, metrics show successful placements by location. For queues that use the FleetIQ algorithm, this data provides useful insight into where player activity occurs.

When evaluating metrics for FleetIQ algorithm performance, consider the following tips:

- To track the queue's rate of finding an ideal placement, use the `PlacementsSucceeded` metric in combination with the FleetIQ metrics for lowest latency and lowest price.
- To boost a queue's rate of finding an ideal placement, review the following error metrics:
 - If the `FirstChoiceOutOfCapacity` is high, adjust capacity scaling for the queue's fleets.
 - If the `FirstChoiceNotViable` error metric is high, look at your Spot Instance fleets. Spot Instance fleets are considered not viable when the interruption rate for a particular instance type is too high. To resolve this issue, change the queue to use Spot Instance fleets with different instance types. We recommend that you include Spot Instance fleets with different instance types in each location.

Prioritize game session placement

Amazon GameLift Servers uses an algorithm to determine how to prioritize a queue's destinations and determine where to place a new game session. The algorithm is based on an ordered set of criteria. You can use the default priority order, or you can customize the order. You can edit a queue's priority order at any time.

Default priority order

1. **Latency** – If the game session placement request includes location-specific latency data for players, Amazon GameLift Servers calculates the average player latency in each location and attempts to place a game session in a fleet location with the lowest average.
2. **Cost** – If a request doesn't include latency data, or if multiple fleets have equal latency, then Amazon GameLift Servers evaluates the hosting cost of each fleet. A fleet's hosting cost varies based on fleet type (Spot or On-Demand), instance type, and location.
3. **Destination** – If multiple fleets have equal latency and costs, then Amazon GameLift Servers prioritizes fleets based on the destination order as listed in the queue configuration.
4. **Location** – For queues with multi-location fleets, if all other criteria are equal, then Amazon GameLift Servers prioritizes the fleet's locations based on alphabetical order.

Customize how a queue prioritizes game session placements

You can choose to customize how a queue prioritizes the placement criteria. The queue applies the custom prioritization to all game session placement requests that it receives.

Note

If you create a custom priority configuration and don't include all four criteria, Amazon GameLift Servers automatically appends any missing criteria in the default order.

To customize a queue's priority configuration

Use the [Amazon GameLift Servers console](#) or the Amazon Command Line Interface (Amazon CLI) to create a custom priority configuration.

Console

In the [Amazon GameLift Servers console](#), you can customize a queue's priorities when creating a new queue or updating an existing queue. Select an Amazon Region to work in.

Open the console's left navigation bar and choose **Queues**. On the Queues page, select an existing queue and choose **Edit**.

1. Go to the section **Game session placement priority**. Drag and drop each priority criteria to create the order you want.

2. Go to the section **Location order**. Add any locations that you want to prioritize. This list is useful when the queue has fleets with multiple locations. At a minimum, you must specify one location. The locations you specify here are prioritized first, followed by all other locations in the queue's destinations.
3. Choose **Save changes**.

Amazon CLI

Use the [update-game-session-queue](#) command with the `--priority-configuration` option to customize a queue's priority order. Amazon GameLift Servers updates a queue in your current default Amazon Region, or you can add a `--region` tag to specify a different Amazon Region.

The following example request adds or updates the priority configuration for a specified queue

```
aws gamelift update-game-session-queue \  
  --name "example-queue-with-priority" \  
  --priority-configuration \  
    PriorityOrder="COST','LOCATION","DESTINATION",LocationOrder="us-east-1","us-  
east-2","ca-central-1","us-west-2" \  

```

Prioritize placements by player latency

If you want to give your players the best possible player experience and ensure minimal latency, take the following steps when setting up your game session placement system:

- Set your queue to prioritize latency when choosing where to place game sessions. Latency is at the top of the priority list by default. You can also customize your queue's priority configuration and choose where to put latency in priority order.
- Set up player latency policies for your queue. Latency policies let you set hard limits on the amount of latency to allow in a game session placement. If Amazon GameLift Servers can't place a game session without exceeding the limits, the placement request will time out and fail. You can set up a single latency policy, or you can create a series of policies that gradually relax the latency limit over time. With a series of policies, you can specify very low initial latency limits, and still accommodate players with higher latencies after a short delay. For details on creating latency policies, see [Create a player latency policy](#).

- When making game session placement requests (see [StartGameSessionPlacement](#)), include latency data for each player. Player latency data includes a value for every possible location where a game session might be placed. For example, for a queue that places game sessions in Amazon Web Services Regions us-east-2 and ca-central-1, latency data might look like this:

```
"PlayerLatencies": [  
  { "LatencyInMilliseconds": 100, "PlayerId": "player1", "RegionIdentifier": "us-east-2" },  
  { "LatencyInMilliseconds": 100, "PlayerId": "player1", "RegionIdentifier": "ca-central-1" },  
  { "LatencyInMilliseconds": 150, "PlayerId": "player2", "RegionIdentifier": "us-east-2" },  
  { "LatencyInMilliseconds": 150, "PlayerId": "player2", "RegionIdentifier": "ca-central-1" }  
]
```

Prioritize placements by location

You can configure a queue to make game session placements based on a prioritized list of geographic locations. Location is one of the criteria that determines how a queue chooses where to place a new game session. By default, location is prioritized fourth, after latency, cost, and destination.

For game session placement, destination and location have somewhat different meanings:

- *Destination* refers to a specific fleet and includes all the fleet's hosting resources, wherever they're deployed. When prioritizing by destination, Amazon GameLift Servers might make a placement with any location in the fleet. Multi-location managed fleets and Anywhere fleets can have hosting resources that are deployed to one or more locations.
- *Location* refers to a specific geographic position where a fleet's hosting resources are deployed. A fleet can have multiple locations, which might include Amazon Web Services Regions, Local Zones, or custom locations (for an Anywhere fleet). A single-location managed fleet has one location and it is always an Amazon Web Services Region. A multi-location managed fleet has a home Region and can have remote locations. An Anywhere fleet has one or more custom locations.

When prioritizing placements by location, Amazon GameLift Servers looks for any queue destinations that include the priority location and searches them for an available hosting resource.

If there are multiple destinations with the priority location, Amazon GameLift Servers moves on to the next priority criteria (cost, latency, destination).

There are several ways that you can influence how a queue's locations are prioritized

- Configure how the queue handles all game session placement requests:
 - **Add a priority configuration to the queue.** A queue's priority configuration includes an ordered list of locations. You can specify one or more locations to prioritize. This list doesn't exclude any locations, it simply tells Amazon GameLift Servers where to look first for an available hosting resource. A common use for an ordered location list is when you want to funnel most traffic to one or more specific geographic locations and use additional locations as backup capacity. Add a priority configuration by calling [UpdateGameSessionQueue](#).
 - **Add a filter configuration to the queue.** A filter configuration is an allow list for the queue. It tells Amazon GameLift Servers to ignore any locations that aren't on the list when looking for an available hosting resource. There are two common uses for a filter configuration. First, for fleets with multiple locations, you might use a filter to exclude some of the fleet's locations. Second, you might want to temporarily disallow placements on a certain location; for example, a location might be experiencing transitory issues. Because you can update a queue's filter configuration at any time, you can easily add and remove locations as needed. Add a filter configuration by calling [UpdateGameSessionQueue](#).
- Use special instructions for individual placement requests:
 - **Include a priority override list in a game session placement request.** You can provide an alternate priority list of locations with any [StartGameSessionPlacement](#) request. This list effectively replaces the queue's configured prioritization for locations for that one request only. It doesn't impact any other requests. This override feature has a few requirements:
 - Use an override list only with a queue that has a priority configuration in place with `LOCATION` as the first priority.
 - Don't include player latency data in the same placement request. Including latency data sets up conflicts when prioritizing locations that Amazon GameLift Servers can't resolve.
 - Decide how you want Amazon GameLift Servers to proceed if it can't find an available resource on the priority override list. Choose between falling back to the queue's other locations, or limit placements to the override list. By default, Amazon GameLift Servers falls back to attempt placement on the queue's other locations.
 - Update the queue's filter configuration as needed, such as adding locations on the override list. The override list doesn't invalidate the filter list.

Create a player latency policy

If your placement requests include player latency data, Amazon GameLift Servers finds game sessions in locations with the lowest average latency for all players. Placing game sessions based on average player latency prevents Amazon GameLift Servers from placing most players in games with high latency. However, Amazon GameLift Servers still places players with extreme latency. To accommodate these players, create player latency policies.

A player latency policy prevents Amazon GameLift Servers from placing a requested game session anywhere that players in the request would experience latency over the maximum value. Player latency policies can also prevent Amazon GameLift Servers from matching game session requests with higher latency players.

Tip

To manage latency specific rules, such as requiring similar latency across all players in a group, you can use [Amazon GameLift Servers FlexMatch](#) to create latency-based matchmaking rules.

For example, consider this queue with a 5-minute timeout and the following player latency policies:

1. Spend 120 seconds searching for a location where all player latencies are less than 50 milliseconds.
2. Spend 120 seconds searching for a location where all player latencies are less than 100 milliseconds.
3. Spend the remaining queue time until timeout searching for a location where all player latencies are less than 200 milliseconds.

Create queue

Queue settings

Name

The name must be unique and have 1-128 characters. Valid characters: A-Z, a-z, 0-9, and - (hyphen).

Timeout

Specify how long GameLift tries to place a game session before stopping.

 seconds

Must be 10-600 seconds.

 We recommend setting player latency policies, unless you're using GameLift FlexMatch.



Player latency policies - *optional*

Add policies to help place players into games with lower latency. Use multiple policies to reduce latency requirements per policy so that each player eventually finds a match.

0 seconds left to allocate 100%

Period start

Seconds

Period end

Seconds

Max player latency

Milliseconds

Remove

Seconds

Seconds

Milliseconds

Remove

Seconds

Seconds

Milliseconds

Remove

Add policy

Build a queue for Spot Instances

You can achieve of potentially significant savings in hosting costs by using Spot fleets. For more details about Spot fleets and how to use them, see [On-Demand Instances versus Spot Instances](#).

If your game hosting solution includes Spot fleets, you must use a game session placement queue. Amazon GameLift Servers uses queues to search across multiple game hosting resources and select the best one available to host a new game session. With Spot fleets, queues are particularly important for minimizing hosting costs and avoiding possible Spot interruptions. This topic helps you set up a resilient queue that can continue to host games for players even in the event of interruptions, slowdowns and outages. You can customize how the queue prioritizes available hosting resources based on several factors including hosting cost.

Are you using FlexMatch for matchmaking? You can use a queue with Spot fleets to make game session placements for your matches.

Implementation tasks for Spot fleets

When creating or updating your game hosting solution to use Spot fleets, complete the following tasks. For more detailed guidance on how to build a queue that optimizes Spot availability and resiliency, see [Reduce game hosting costs with Spot fleets](#).

1. Choose and create a set of fleet destinations for your game session queue.

Start by deciding where you want your queue to place game sessions. A queue can search across multiple fleets to find the best possible placement. Each fleet has one instance type but can have multiple geographic locations. Queues with fleets that offer variety in both location and instance type are more likely to make successful placements. See these best practices for designing an effective and resilient Spot-optimized queue.

2. Create your Spot-optimized game session queue.

Create a queue and configure it for your Spot fleets. See [Create a game session queue](#) for help creating and configuring the new queue. You can use the Amazon GameLift Servers console or the Amazon CLI to create or edit a queue.

- Add the fleet destinations from Step 1.
- Prioritize the destination order as appropriate. By default, Amazon GameLift Servers prioritizes by cost before destination, so destination order is used only when the lowest costs between destinations are equal.
- If you want to prioritize game hosting cost before player latency, provide a custom placement priority. See [Prioritize game session placement](#).

3. Update other components in your solution to use the new queue.

When your solution uses a Spot-optimized queue to start new game sessions, the queue automatically avoids placing game sessions with fleets that have a high likelihood of interruption. It instead searches all viable fleets for resources that match your defined priorities, including player latency, hosting cost, and destination order.

- If you're not using FlexMatch – Update your backend service to specify the new Spot-optimized queue in game session requests. The backend service makes API requests to Amazon GameLift Servers on behalf of your game client (using `StartGameSessionPlacement()`), and each request must specify a queue name. For help implementing game session placements in your game client, see [Create game sessions](#).
- If you are using FlexMatch – Update your matchmaking configuration to send game session requests to the new Spot-optimized queue. When the matchmaking system forms a player match, it sends a game session placement request to the designated queue to start a new game session for the match. Only matchmaking configurations with FlexMatch mode set to "Managed" can designate a placement queue. You can update a matchmaking configuration using the Amazon CLI or the Amazon GameLift Servers console (see [Edit a matchmaking configuration](#)).

4. Review the performance of your Spot fleets and queues.

View Amazon GameLift Servers metrics in the Amazon GameLift Servers console or with Amazon CloudWatch to review performance. For more information about Amazon GameLift Servers metrics, see [Monitor Amazon GameLift Servers with Amazon CloudWatch](#). Key metrics include:

- Interruption rate – Use the `InstanceInterruptions` and `GameSessionInterruptions` metrics to track the number and frequency of Spot-related interruptions for instances and game sessions. Game sessions on reclaimed instances have a status of `TERMINATED` and a status reason of `INTERRUPTED`.
- Queue effectiveness – Track placement success rates, average wait time, and queue depth to confirm that Spot fleets don't impact your queue performance.
- Fleet usage – Monitor data on instances, game sessions and player sessions. Usage for your On-Demand fleets can be an indicator that queues are avoiding placements into your Spot fleets to avoid disruption.

Best practices for queues with Spot fleets

Use the following best practices when creating fleets and queues for Spot instances.

- **Expand your queue's geographic coverage.** Even if your players are clustered in a single Amazon Web Services Region, add adjacent locations to your Spot fleet. This approach improves the queue's ability to maintain capacity during regional slowdowns, outages, and Spot interruptions. Multi-location fleets work with both Spot and On-Demand instances.
- **Diversify your queue's instance type coverage.** Amazon GameLift Servers evaluates Spot viability based on instance type, so having Spot fleets with a variety of instance types reduces the chance that multiple Spot fleets are nonviable at the same time. Include at least two Spot fleets with different instances types each location.

Note

Pricing is based on the instances you use, not the number of fleets. Running five fleets with 10 instances each is the same as running one fleet with 50 instances of similar cost. Pricing varies by instance type, size, and location.

Tips for grouping Spot instance types:

- Use instance types in the same family, such as `m6g.medium`, `m6g.large`, and `m6g.xlarge`. Larger instance types cost more, but can also host more game sessions at a time.
- Select widely available instances types. Typically, older generation families (such as C5, M5, and R5) and common sizes (such as `.large`, `.xlarge`, and `.2xlarge`) have better availability.
- Check the 30-90 day pricing history in the Amazon GameLift Servers console. Look for instance types with consistent availability patterns.
- Use the Amazon GameLift Servers console, fleet creation tool, to explore location coverage for instance types.
- **Add On-Demand fleets for backup capacity.** Game hosting can switch to On-Demand fleets whenever Spot fleets are unavailable. Put at least one On-Demand fleet in each location to maintain low player latency. Add auto-scaling to your backup On-Demand fleets, so you can keep them scaled down until they're needed.
- **Assign aliases to all fleet destinations.** Create aliases for each of your queue's destinations. Aliases make it easier and more efficient whenever you need to replace fleets.

- **Apply a queue prioritization strategy.** You can customize how a queue prioritizes where to place game sessions (see [Prioritize game session placement](#) for more details). For Spot-optimized queues, prioritizing by cost ensures that low-cost Spot fleets are used whenever possible.

You can also prioritize certain fleets by specifying a destination order. For example, some users designate a set of primary fleets for regular use and also a set of secondary fleets as backup. In this scenario, set the queue's destination order to list the primary fleets first. Then configure the queue's priority order with destination followed by cost.

Hosting resource customizations

This section provides advanced options for configuring and managing your Amazon GameLift Servers infrastructure to meet specific performance, cost, and operational requirements. In particular, the topics in this section describe how you can customize your Amazon GameLift Servers managed hosting resources to best fit your game and your players.

Some of the decisions you want to consider:

- Where to deploy hosting resources for your players? Gameplay latency is a major factor in selecting your fleet's geographic locations, but there are other factors that vary by location, including resource type availability and cost.
- What EC2 instance types will best support your game? Choose from available instance types that use the best combination of compute architecture, memory, storage, and networking capacity.
- What size of instance type do you need? Choose an instance type size based on the resource requirements (memory and CPU) of your game server software and other factors.
- Should your fleet use On-Demand or Spot Instances? Consider whether you can take advantage of lower Spot pricing, and whether Amazon GameLift Servers sufficiently mitigates the chance of Spot interruptions to your game sessions.
- How do you want your game server software to run on each fleet instance? The runtime configuration tells Amazon GameLift Servers what server software to run and how.
- For container fleets, does the default configuration work for your game? Amazon GameLift Servers does a lot of the work for you to optimize your container fleet configurations, but you can customize most configuration settings.

Choose compute resources for a managed fleet

For Amazon GameLift Servers managed hosting, including managed EC2 and managed containers, the service deploys your game servers to fleets of computing resources in the Amazon Web Services Cloud. When you create a managed fleet, you want to configure the hosting resources to best suit your game. This topic discusses key decision points when choosing and configuring your game hosting fleets.

Topics

- [Geographic locations](#)
- [Operating systems](#)
- [Instance types](#)
- [On-Demand Instances versus Spot Instances](#)
- [Service quotas](#)

Geographic locations

Consider where you plan to deploy your game servers. In general, you want to put your game servers as close as possible to your players to deliver the best possible player experience. For Amazon GameLift Servers managed hosting, you can choose to put game servers in any of the supported Amazon Web Services Regions and Local Zones. If you're building a hybrid solution, consider how managed fleet deployment can supplement the locations of your self-managed Amazon GameLift Servers Anywhere fleets.

For most development and testing scenarios, deploying to a single location makes sense. As you prepare for launch and beyond, there are many reasons to deploy across multiple geographic locations. These include supporting a widespread group of players and improving overall game hosting resilience and reliability. Multiple locations can also boost player experience by speeding up game session placement and enabling more choices when optimizing placements for latency and cost.

For a list of locations that are supported by Amazon GameLift Servers and more information about locations for all fleet types, see [Amazon GameLift Servers service locations](#).

Multi-location fleets

A single managed fleet can deploy resources to multiple locations. You can manually set capacity for each individual location in a multi-location fleet.

Advantages of using a multi-location fleet:

- **Simplified fleet deployment and management** – You supply the game server software and fleet configuration, and Amazon GameLift Servers deploys it to fleet instances across multiple locations (build once, deploy anywhere). In a production fleet, you can view and manage all locations in a fleet instead of having to manage multiple fleets each located in a different region.
- **Local Zone availability** – If you want to use a Local Zone you must create a multi-location fleet with an Amazon Web Services Region home location and Local Zones as remote locations. Local Zones are extensions of Amazon Web Services Regions that can deliver even lower latency to areas and customers that need it. You can add a Local Zone to any multi-location fleet; you don't need to include the Local Zone's parent Amazon Web Services Region.
- **Compatibility with game session queues** – You can build game session placement queues with one or more multi-location fleets. This approach gives the queue flexibility when prioritizing and choosing locations to host a new game session.
- **Efficient resource utilization** – With auto-scaling turned on, Amazon GameLift Servers can better optimize capacity scaling across all locations in a fleet.

Tips for using multi-location fleets:

- Check for quotas on the number of locations per Amazon Web Services Region or fleet. See [Amazon GameLift Servers service quotas](#).
- Not all instance types are available in all locations. Depending on your chosen locations, you might have limited instance type options. The Amazon GameLift Servers console provides useful tools to help you find the right balance of locations and instance types.
- Consider using [???](#) to collect player latency data for all your fleet locations. Amazon GameLift Servers can use this data to optimize game sessions for low latency and prevent players from joining sessions with unacceptably high latency. These special endpoints accept UDP messages instead of traditional ICMP pings, providing accurate latency measurements to help you select optimal fleet locations.

Operating systems

All instances in a managed fleet are deployed with an Amazon machine image (AMI) that provides a complete runtime environment for your game server software. For managed EC2 fleets, you specify the game server build's operating system when you upload the build to Amazon GameLift Servers.

For managed container fleets, you specify the operating system in the container group definition. For more information about the latest AMI versions, see [Amazon GameLift Servers AMI versions](#).

AMI versions are regularly updated. When you create a new fleet, Amazon GameLift Servers assigns the latest available version of the AMI you selected for your game build. All instances that are deployed in that fleet use the same version. To keep your AMI version up to date with the latest security and software updates, you need to regularly replace your fleets. As a best practice, we recommend replacing your managed fleets every 30 days to maintain the runtime environment for your game servers. For guidance, see [Security best practices for Amazon GameLift Servers](#).

Instance types

A managed fleet's instance type determines the kind of hardware that is deployed for all fleet instances, and instance types are generally available in various sizes. All Amazon GameLift Servers managed fleet use Amazon EC2 instances, and support a wide range of instance types that offer different combinations of computing power, memory, storage, and networking capabilities. The availability of instance types varies depending on the locations you choose.

The Amazon GameLift Servers console provides useful tools to help you find the right instance types for your game build and your deployment locations. For managed container fleets, the console also offers guidance on your game's CPU power and memory requirements.

When choosing from available instance types for your game, consider:

- The compute architecture of your game server: x64 or Arm (Amazon Graviton).

Note

Graviton Arm instances require a server build for a Linux AMI. Server SDK 5.1.1 or newer is required for C++ and C#. Server SDK 5.0 or newer is required for Go. These instances do not provide out-of-the-box support for Mono installation on Amazon Linux 2023 (AL2023) or Amazon Linux 2 (AL2).

- The computing, memory, and storage requirements of your game server build.
- The size of your instance type. In addition to meeting the requirements of your game server software executables, larger instance type sizes can run multiple game server processes and/or containers on each instance. Factors to consider include cost (is it cheaper to run a few large instances or many small instances). Also consider how game session capacity might be impacted by adding or removing instances during fleet scaling events or when shutting down unhealthy

instances. If each instance runs many game server processes concurrently, adding or removing an instance can significantly affect game hosting capacity.

For more information about instance types, see [Amazon EC2 Instance Types](#).

On-Demand Instances versus Spot Instances

Amazon EC2 On-Demand Instances and Spot Instances offer the same hardware and performance, but they differ in availability and cost.

On-Demand Instances

You can acquire an On-Demand Instance when you need it and keep it for as long as you want. On-Demand Instances have a fixed cost, meaning you pay only for the amount of time that you use them. There are no long-term commitments.

Spot Instances

Spot Instances can offer a cost-efficient alternative to On-Demand Instances by utilizing unused Amazon computing capacity. Spot Instance prices fluctuate based on the supply and demand for each instance type in each location. Amazon can reclaim Spot Instances with a two-minute notification whenever it needs the capacity back, and game sessions actively running on a reclaimed instance are interrupted.

Amazon GameLift Servers offers several tools to help mitigate the likelihood of Spot interruptions to your game sessions. A spot viability algorithm tracks instance type historical data to anticipate when the risk of interruption reaches a critical point and avoids placing new game sessions on Spot Instances of that type. If an interruption does occur, your game servers can use the notification to gracefully end a game session for players.

Game hosting with Spot fleets must use a queue for game session placement. A queue is able to prioritize game session placements based on Spot fleet viability, cost, and other factors. See these topics for more information about how to take advantage of Spot for your game server hosting:

- [Reduce game hosting costs with Spot fleets](#)
- [Build a queue for Spot Instances](#)

Service quotas

You can view default service quotas for Amazon GameLift Servers and current quota status for your Amazon Web Services account using the following tools:

- For general service quota information for Amazon GameLift Servers, see [Amazon GameLift Servers endpoints and quotas](#) in the *Amazon Web Services General Reference*.
- For a list of available instance types per location for your account, open the [Service quotas](#) page of the Amazon GameLift Servers console. This page also displays your account's current usage for each instance type in each location.
- For a list of your account's current quotas for instance types per Region, run the Amazon Command Line Interface (Amazon CLI) command [describe-ec2-instance-limits](#). This command returns the number of active instances that you have in your default Region (or in another Region that you specify).

As you prepare to launch your game, fill out a launch questionnaire in the [Amazon GameLift Servers console](#). The Amazon GameLift Servers team uses the launch questionnaire to determine the correct quotas and limits for your game.

Customize an Amazon GameLift Servers container fleet

The topics in this section describe some of the optional features for Amazon GameLift Servers managed containers. You can choose to use any or all of these features.

Topics

- [Set resource limits](#)
- [Configuring NVMe Drive Access](#)
- [Designate essential containers](#)
- [Configure network connections](#)
- [Set up health checks for containers](#)
- [Set container dependencies](#)
- [Configure a container fleet](#)

Set resource limits

For each container group, you can determine how much memory and computing power the container group needs to run its software. Amazon GameLift Servers relies on this information to manage resources across the container group. It also uses this information to calculate how many game server container groups a fleet image can hold. You can also set limits for individual containers.

You can set a maximum limit on memory and computing power for a container group. By default, these resources are shared by all containers in the group. You can further customize resource management by setting limits for individual containers.

Set optional limits for individual containers

Setting container-specific resource limits allows you to exert greater control over how individual containers can use the group's resources. If you don't set container-specific limits, all containers in the group share the group resources. Sharing offers greater flexibility to use resources where they're needed. It also increases the potential for processes to compete with each other and result in container failure.

Set any of the following `ContainerDefinition` properties for any container.

- `MemoryHardLimitMebibytes` – Set a maximum memory limit for the container. If the container exceeds this limit, it results in a restart.
- `Vcpu limit` – Reserve a minimum amount of vCPU resources for the container's exclusive use. The container always has the reserved amount available to it. It can exceed this minimum at any time, if additional resources are available. (1024 CPU units is the equivalent of 1 vCPU.)

Set total resource limits for a container group

If you set limits for individual containers, you might need to modify how much memory and vCPU resources the container group needs. The goal is to allocate enough resources to optimize game server performance. Amazon GameLift Servers uses these limits to calculate how to pack game server container groups on a fleet instance. You'll also use them when choosing an instance type for a container fleet.

Calculate the total memory and vCPU needed for a container group. Consider the following:

- What are all the processes that run across all containers in the container group? Add up the resources required for these processes. Take note of any container-specific limits.
- How many concurrent game server processes do you plan to run in each container group? You determine this in your game server container image.

Based on your estimate of container group requirements, set the following `ContainerGroupDefinition` properties:

- `TotalMemoryLimitMebibytes` – Set a maximum memory limit for the container group. All containers in the group share allocated memory. If you set individual container limits, the total memory limit must be equal to or greater than the highest container-specific memory limit.
- `TotalVcpuLimit` – Set a maximum vCPU limit for the container group. All containers in the group share allocated CPU resources. If you set individual container limits, the total CPU limit must be equal to or greater than the sum of all container-specific CPU limits. As a best practice, consider setting this value to double the sum of the container CPU limits.

Example scenario

Let's say we're defining a game server container group with the following three containers:

- Container A is our game server container. We estimate the resource requirements for one game server at 512 MiB and 1024 CPU. We plan to have the container run 1 server process. Because this container runs our most critical software, we set no memory limit or vCPU reserve limit.
- Container B is a support container with resource requirements estimated at 1024 MiB and 1536 CPU. We set a memory limit of 2048 MiB, and a CPU reserve limit of 1024 CPU.
- Container C is another support container. We set a hard memory limit of 512 MiB and a CPU reserve limit of 512 CPU.

Using this information, we set the following total limits for the container group:

- Total memory limit: 7680 MiB. This value exceeds the highest memory limit (1024 MiB).
- Total CPU limit: 13312 CPU. This value exceeds the sum of the CPU limit (1024+512 CPU).

Configuring NVMe Drive Access

On d-type instances, the NVMe drive automatically mounts to the `/data` directory during host startup. To enable containers to access the SSD storage, set the following `ContainerGroupDefinition` property `MountPoints`:

- `InstancePath` – Set to `/data` to reference the auto-mounted NVMe drive on the host instance.
- `AccessLevel` – Choose the appropriate access level for your container's needs (e.g., `READ_ONLY` or `READ_WRITE`).

- **ContainerPath** – (Optional) Specify the path where the instance path will be mounted inside the container. If not specified, it defaults to the instance path.

For more information about mount points, see [ContainerMountPoint](#) in the Amazon GameLift Servers API Reference.

Designate essential containers

For a per-instance container group, designate each container as essential or non-essential. Per-instance container groups must have at least one essential support container. The essential container does the critical work of the container group. The essential container is always expected to be running. If it fails, the entire container group restarts.

Set the `ContainerDefinition` property `Essential` to either `true` or `false` for each container.

Configure network connections

You can customize network access to allow external traffic connect to any container in a container fleet. For example, you must establish network connections to the container that runs your game server processes, so that game clients can join and play your game. Game clients connect to game servers using ports and IP addresses.

In a container fleet, the connection between a client and server is not direct. Internally, a process in a container listens on a *container port*. Externally, incoming traffic connects to a fleet instance using a *connection port*. Amazon GameLift Servers maintains the mappings between internal container ports and external-facing connection ports, so that incoming traffic gets routed to the correct process on the instance.

Amazon GameLift Servers provides an extra layer of control for your network connections. Each container fleet has an *inbound permissions* setting, which allows you to control access to each external-facing connection port. For example, you could remove permissions for all connection ports to shut off all access to the fleet's containers.

You can update a fleet's inbound permissions, connection ports, and container ports.

Warning

If you provide a custom `InstanceConnectionPortRange` or `InstanceInboundPermissions`, Amazon GameLift Servers will no longer manage either value for your fleet. You must set both fields to avoid undefined behavior.

Set container port ranges

Configure container port ranges as part of each container definition. This is a required parameter for a container group definition. You need to configure enough ports to accommodate all concurrently running processes that need external access. Some containers will not need any ports.

Your game server container, which runs your game servers, needs a port for every concurrently running game server process. The game server process listens on the assigned port and reports it to Amazon GameLift Servers.

Set connection port ranges

Configure your container fleet with a set of connection ports. Connection ports provide external access to the fleet instances that are running your containers. Amazon GameLift Servers assigns connection ports and maps them to container ports as needed.

By default, Amazon GameLift Servers calculates the number of ports required for all container groups and sets a port range to accommodate them. We highly recommend you use Amazon GameLift Servers calculated values, which are updated when you deploy updates to a container group definition. If you do need to customize connection port ranges, use the following guidance.

When you create a container fleet, define a connection port range (see [ContainerFleet:InstanceConnectionPortRange](#)). Make sure that the range has enough ports to map to every container port that's defined across all containers in both container groups in the fleet. To calculate the minimum connection ports needed, use the following formula:

```
[Total number of container ports defined for containers in the game  
server container group] * [Number of game server container groups per  
instance] + [Total number of container ports defined for containers in  
the per-instance container group]
```

As a best practice, double the minimum number of connection ports.

Note

The number of connection ports can potentially limit the number of game server container groups per instance. If a fleet has only enough connection ports for one game server container group per instance, Amazon GameLift Servers will deploy only one

game server container group, even if the instances have enough compute power for multiple game server container groups.

Set inbound permissions

Inbound permissions control external access to a container fleet by specifying which connection ports to open for incoming traffic. You can use this setting to turn a fleet's network access on and off as needed.

By default, Amazon GameLift Servers calculates the number of ports required for all container groups and sets a port range to accommodate them. We highly recommend you use Amazon GameLift Servers calculated values, which are updated when you deploy updates to a container group definition. If you do need to customize connection port ranges, use the following guidance.

When you create a container fleet, define a set of inbound permissions (see [ContainerFleet:InstanceInboundPermissions](#)). Inbound permission ports should match with the fleet's connection port ranges.

Note

Since container ports are selected randomly from the `InstanceConnectionPortRange`, in order to guarantee that session connections can be made, all ports in `InstanceConnectionPortRange` should be covered by ports in `InstanceInboundPermissions`.

Example scenario

This example illustrates how to set all three network connection properties.

- Our fleet's game server container group has 1 container, which runs 1 game server process.

In the game server container group definition, we set the `PortConfiguration` parameter for this container as follows:

```
"PortConfiguration": {  
  "ContainerPortRanges": [ { "FromPort": 10, "ToPort": 20, "Protocol": "TCP" } ]  
}
```

- Our fleet also has a per-instance container group with 1 container. It has 1 process that needs network access. In the per-instance container definition, we set the `PortConfiguration` parameter for this container as follows:

```
"PortConfiguration": {  
  "ContainerPortRanges": [ { "FromPort": 25, "ToPort": 25, "Protocol": "TCP"} ] }
```

- Our fleet is configured with 20 game server container groups per fleet instance. Given this information, we can use the formula to calculate the number of connection ports we need:
 - Minimum: **21 ports** [1 game server container ports * 20 game server container groups per instance + 1 per-instance container port]
 - Best practice: **42 ports** [minimum ports * 2]

When creating the container fleet, we set the `InstanceConnectionPortRange` parameter as follows:

```
"InstanceConnectionPortRange": { "FromPort": 1010, "ToPort": 1071 }
```

- We want to allow access to all available connection ports. When creating the container fleet, we set the `InstanceInboundPermissions` parameter as follows:

```
"InstanceInboundPermissions": [  
  { "FromPort": 1010, "ToPort": 1071, "IpRange": "10.24.34.0/23", "Protocol":  
    "TCP"} ]
```

Set up health checks for containers

A container automatically restarts if it experiences a terminal failure and stops running. If a container is considered essential, it prompts the entire container group to restart.

All game server containers are automatically considered essential. Support containers can be designated essential, but they need to have a mechanism to report health. You can set health checks for non-essential support containers as well.

You can define additional custom criteria to measure container health and use a health check to test that criteria. To set up a container health check, you can define it in a Docker container image or in your container definition. If you set a health check in the container definition, it overrides any settings in the container image.

Set the following `SupportContainerDefinition` properties for a container health check:

- **Command** — Provide a command that checks some aspect of the container's health. You decide what criteria to use to measure health. The command must result in an exit value of 1 (unhealthy) or 0 (healthy).
- **StartPeriod** — Specify an initial delay before health check failures start counting. This delay gives the container time to bootstrap its processes.
- **Interval** — Decide how often to run the health check command. How quickly do you want to detect and resolve a container failure?
- **Timeout** — Decide how long to wait for success or failure before retrying the health check command. How long should the health check command take to complete?
- **Retries** — How many times should the health check command be retried before registering a failure?

Set container dependencies

Within each container group you can set dependencies between containers based on container status. A dependency impacts when the dependent container can start or shut down based the status of another container.

A key use case for dependencies is to create startup and shutdown sequences for the container group.

For example, you might want Container A to start first and complete successfully before Containers B and C start. To achieve this, first create a dependency for Container B on Container A, with the condition that Container A must complete successfully. Then create a dependency for Container C on Container A with the same condition. Startup sequences occur in reverse order for shutdown.

Configure a container fleet

When you create a container fleet, consider the following decision points. Most of these points are dependent on your container architecture and configuration.

Decide where you want to deploy your fleet

In general, you want to deploy your fleets geographically near your players to minimize latency. You can deploy your container fleet to any Amazon Web Services Region that Amazon GameLift Servers supports. If you want to deploy the same game server to additional geographic locations, you can add remote locations to the fleet including Amazon Web Services Regions

and Local Zones. For a multi-location fleet, you can adjust capacity independently in each fleet location. For more information about supported fleet locations, see [Amazon GameLift Servers service locations](#).

Choose an instance type and size for your fleet

Amazon GameLift Servers supports a wide range of Amazon EC2 instances types, all of which are available for use with a container fleet. Instance type availability and price varies by location. You can view a list of supported instance types, filtered by location, in the Amazon GameLift Servers console (under **Resources, Instance and service quotas**).

When choosing an instance type, first consider the instance family. Instance families offer various combinations of CPU, memory, storage, and networking capabilities. Get more information on [EC2 instance families](#). Within each family you have a range of instance sizes to choose from. Consider the following issues when selecting an instance size:

- What's the minimum instance size that can support your workload? Use this information to eliminate any instance types that are too small.
- What instance type sizes are a good fit for your container architecture? Ideally, you want to choose a size that can accommodate multiple copies of your game server container group with minimal wasted space.
- What scaling granularity makes sense for your game? Scale fleet capacity involves adding or removing instances, and each instance represents the ability to host a specific number of game sessions. Consider how much capacity you want to add or remove with each instance. If player demand varies by thousands from minute to minute, then it might make sense to use very large instances that can host hundreds or thousands of game sessions. By contrast, you might prefer more fine-grained scaling control with smaller instance types.
- Are there cost savings available based on size? You might find that the cost of certain instance types vary by location due to availability.

Set other optional fleet settings

You can use the following optional features when configuring a container fleet:

- Set up your game servers to access other Amazon resources. See [Connect your Amazon GameLift Servers hosted game server to other Amazon resources](#).
- Protect game sessions with active players from terminating prematurely during a scale-down event.
- Limit the number of game sessions that one individual can create on the fleet within a limited span of time.

Reduce game hosting costs with Spot fleets

When hosting multiplayer game servers using Amazon GameLift Servers managed hosting, Spot instances can provide a cost-effective alternative to On-Demand instances. The Spot pricing model offers the same hardware and performance as On-Demand but at potentially significant cost savings (up to 70-90%). However, they come with a limitation: when Amazon needs capacity back, it can reclaim these instances with a two-minute interruption notification.

Amazon GameLift Servers mitigates the risk of interruptions for game server hosting. Amazon GameLift Servers predicts the likelihood of interruptions on Spot instance types and avoids placing game sessions on any instances at risk. If a rare interruption does occur, the notification allows you to gracefully end a game session for players.

How Amazon GameLift Servers works with Spot fleets

When you set up Spot fleets for game hosting, Amazon GameLift Servers continually evaluates your Spot fleet instance types and locations for game hosting viability.

- The Spot viability algorithm analyzes recent availability patterns and historical interruption rates of Spot instance types by location.
- Based on this analysis, Amazon GameLift Servers identifies Spot instance types and locations where there's an unacceptable potential for game session interruption. It takes the following actions:
 - It marks the combination of instance type and location as temporarily non-viable.
 - It removes any non-viable Spot fleet location from consideration when placing new game sessions. As a result, game sessions are only placed in Spot fleet locations that have a high likelihood of uninterrupted game server hosting.
 - It drains the Spot fleet location of existing instances, even if Amazon doesn't reclaim them, so you're not paying for instances that can't be used for game hosting. If game session protection is turned on, instances are shut down only after active game sessions are completed.
- Amazon GameLift Servers continually re-evaluates your Spot fleet instance types and locations for game hosting viability. When a previously non-viable instance type becomes viable again based on updated historical data, you can scale up the Spot fleet again, and Amazon GameLift Servers will resume placing game sessions with it.

Design considerations

When designing your solution to use Spot fleets, consider the following issues:

- **Assess game session length** – The average length of your game sessions can impact how well Spot works for your game. With shorter game sessions, the faster turnaround keeps game sessions running on viable instance types based on the latest historical data. Longer game sessions continue to run on instance types without evaluating recent viability data, running a greater risk of interruption over time.
- **Evaluate instance type availability** – Not every fleet location offers every instance type as Spot. When choosing an instance type for your Spot fleet, use the Amazon GameLift Servers console fleet creation tool to help you find Spot instance types in the locations you need. With this tool, you can select your fleet locations and then view instance type availability across those locations.
- **Create multi-location Spot fleets** – You can create Spot fleets with multiple locations. A single multi-location Spot fleet deploys instances with the same instance type to multiple Amazon Web Services Regions or Local Zones. The Spot viability algorithm evaluates viability based on both instance type and location. If a Spot fleet location is evaluated as nonviable, it doesn't impact other locations in the fleet, which can still be used to host game sessions.
- **Create queues with Spot fleet diversity** – If you're using Spot fleets for game hosting, you need to set up a game session placement queue. For each new game session request, the queue looks for available game hosting resources and selects the best possible option. With Spot fleets, you want a queue that can search across multiple fleets that vary in both locations and instance types, and you want to include at least one On-Demand fleet as backup capacity. A well-designed multi-fleet queue that offers diverse placement options is highly resilient against interruptions, slowdowns, and outages. For additional guidance on designing a queue for Spot, see [Build a queue for Spot Instances](#).
- **Handle interruptions gracefully** – Set up your game server to minimize player impact in the event of a Spot interruption. When Amazon reclaims a Spot instance, Amazon GameLift Servers passes the termination notification to all affected server processes using the server SDK callback function `onProcessTerminate()`. Your game needs to implement this callback to end the game session gracefully. For more information, see [Respond to a server process shutdown notification](#).

 Note

Amazon makes every effort to provide the notification before it reclaims an instance, but it's possible that Amazon reclaims the Spot Instance before the warning arrives. You should also prepare your game server to handle unexpected interruptions.

- **Configure auto-scaling for your backup fleets to maintain service during a Spot interruption.** Target-tracking auto-scaling maintains a capacity buffer and scales automatically with demand. With auto-scaling, backup fleets (Spot or On-Demand) will begin increasing capacity whenever they start receiving more game session requests.

To quickly replace lost capacity when a Spot fleet becomes non-viable, a custom scaling mechanism can use available queue and fleet metrics to initiate a rapid scale-out of backup fleets. Detect when Spot fleets become non-viable with metrics such as `FirstChoiceOutOfCapacity`, `FirstChoiceNotViable`, and `PercentAvailableGameSessions`. Estimate replacement capacity needs by analyzing recent `PlacementsStarted` metric data. After scaling backup fleets to handle immediate demand, normal auto-scaling can take over.

- **Integration with FlexMatch** – If your solution uses a FlexMatch matchmaker, there are no special requirements for Spot fleets. You can configure a matchmaker to use a queue with Spot fleets. Amazon GameLift Servers automatically prioritizes match placements across Spot and On-Demand fleets, including when placing new game sessions and when backfilling empty player slots in existing game sessions.

Optimize game server runtime configuration on managed Amazon GameLift Servers

You can set up an managed EC2 fleet's runtime configuration to run multiple game server processes per instance. This uses your hosting resources more efficiently.

How a fleet manages multiple processes

Amazon GameLift Servers uses a fleet's runtime configuration to determine the type and number of processes to run on each instance. A runtime configuration contains at least one server process configuration that represents one game server executable. You can define additional server process configurations to run other types of processes related to your game. Each server process configuration contains the following information:

- The file name and path of an executable in your game build.
- (Optional) Parameters to pass to the process on launch.
- The number of processes to run concurrently.

When an instance in the fleet activates, it launches the set of server processes defined in the runtime configuration. With multiple processes, Amazon GameLift Servers staggers the launch of each process. Server processes have a limited life span. As they end, Amazon GameLift Servers launches new processes to maintain the number and type of server processes defined in the runtime configuration.

You can change the runtime configuration at any time by adding, changing, or removing server process configurations. Each instance regularly checks for updates to the fleet's runtime configuration to implement the changes. Here's how Amazon GameLift Servers adopts runtime configuration changes:

1. The instance sends a request to Amazon GameLift Servers for the latest version of the runtime configuration.
2. The instance compares its active processes to the latest runtime configuration, and then does the following:
 - If the updated runtime configuration removes a server process type, then active server processes of this type continue to run until they end. The instance doesn't replace these server processes.
 - If the updated runtime configuration decreases the number of concurrent processes for a server process type, then excess server processes of this type continue to run until they end. The instance doesn't replace these excess server processes.
 - If the updated runtime configuration adds a new server process type or increases the concurrent processes for an existing type, then the instance starts new server processes, up to the Amazon GameLift Servers maximum. In this case, the instance launches new server processes as existing processes end.

Optimize a fleet for multiple processes

To use multiple processes on a fleet, do the following:

- [Create a build](#) that contains the game server executables that you want to deploy to a fleet, and then upload the build to Amazon GameLift Servers. All game servers in a build must run on the same platform and use the server SDK for Amazon GameLift Servers.
- Create a runtime configuration with one or more server process configurations and multiple concurrent processes.
- Integrate game clients with the Amazon SDK version 2016-08-04 or later.

To optimize fleet performance, we recommend that you do the following:

- Handle server process shutdown scenarios so that Amazon GameLift Servers can recycle processes efficiently. For example:
 - Add a shutdown procedure to your game server code that calls the server API `ProcessEnding()`.
 - Implement the callback function `OnProcessTerminate()` in your game server code to handle termination requests from Amazon GameLift Servers.
- Make sure that Amazon GameLift Servers shuts down and relaunches unhealthy server processes. Report the health status back to Amazon GameLift Servers by implementing the `OnHealthCheck()` callback function in your game server code. Amazon GameLift Servers automatically shuts down server processes that are reported unhealthy for three consecutive reports. If you don't implement `OnHealthCheck()`, then Amazon GameLift Servers assumes that a server process is healthy, unless the process fails to respond to a communication.

Choose the number of processes per instance

When deciding on the number of concurrent processes to run on an instance, keep in mind the following:

- Amazon GameLift Servers limits each instance to a [maximum number of concurrent processes](#). The sum of all concurrent processes for a fleet's server process configurations can't exceed this quota.
- To maintain acceptable performance levels, the Amazon EC2 instance type might limit the number of processes that can run concurrently. Test different configurations for your game to find the right number of processes for your preferred instance type.
- Amazon GameLift Servers doesn't run more concurrent processes than the total number configured. This means that the transition from the previous runtime configuration to the new configuration might happen gradually.

Work with the Amazon GameLift Servers Agent

The Amazon GameLift Servers Agent oversees the running of game server processes on your Amazon GameLift Servers fleets. The Agent is deployed to each compute in a fleet, and it provides automated process management, hosting management, and logging for the compute. To use the

Agent, you must have your game server build integrated with the server SDK for Amazon GameLift Servers version 5.x or later.

The Amazon GameLift Servers Agent is externally available for use with Amazon GameLift Servers fleets that are not managed EC2 fleets. (Managed EC2 fleets handle the Agent's tasks automatically.) You can opt to run Amazon GameLift Servers fleets, including Anywhere fleets, with or without the Agent. Without the Agent, you must provide an alternative solution to completing the required tasks.

When deployed to a compute, the Amazon GameLift Servers Agent should be launched before any game server processes are started. On launch, the Agent completes the following tasks:

- Registers the compute with an Amazon GameLift Servers Anywhere fleet using the [RegisterCompute](#) API.
- Calls the [GetComputeAuthToken](#) API to fetch an authorization token and stores it for use by server processes that are running on the compute.
- Sets the WebSocket URL environment variable for the compute, and establishes a WebSocket connection to the Amazon GameLift Servers service.
- Requests the latest version of the fleet's runtime configuration from the Amazon GameLift Servers service.
- Starts and stops server processes according to the runtime configuration instructions.

Source code and build instructions for the Amazon GameLift Servers Agent are available in the [Amazon GameLift Servers Agent](#) GitHub.

About the Agent

The Amazon GameLift Servers Agent is designed to handle the following tasks for your fleets:

Process management

- Starts new server processes as defined in runtime instructions. The Agent might use a custom runtime configuration that is deployed with the Agent. Alternatively, you can provide a `RuntimeConfiguration` as part of your fleet definition. This approach has an advantage in that you can modify the fleet's runtime configuration at any time. The Agent periodically requests updated runtime configurations from the Amazon GameLift Servers service.
- Monitors server process activations and terminates processes when they don't activate in time.

- Sends heartbeats to Amazon GameLift Servers. If the Agent fails to send heartbeats, the compute might be marked as stale.
- Reports to Amazon GameLift Servers when a server process ends. Amazon GameLift Servers uses this information to monitor game server availability for game session placement.
- Emits fleet events for server processes, including:
 - `SERVER_PROCESS_INVALID_PATH`: The game server process launch parameters were incorrectly configured.
 - `SERVER_PROCESS_TERMINATED_UNHEALTHY`: The game server process did not report a valid health check within 3 minutes of activating and was therefore terminated.
 - `SERVER_PROCESS_FORCE_TERMINATED`: The game server process did not exit cleanly after `OnProcessTerminate()` was sent within 30 seconds.
 - `SERVER_PROCESS_CRASHED`: A game server process crashed for some reason.

Compute management

- Receives messages from the Amazon GameLift Servers service to shut down the compute.
- Prompts the compute to be terminated by Amazon GameLift Servers.

Logging

- Uploads logs to an Amazon S3 bucket in your Amazon account.

Abstract an Amazon GameLift Servers fleet designation with an alias

An Amazon GameLift Servers *alias* is used to abstract a hosting destination. Hosting destinations tell Amazon GameLift Servers where to look for available resources to host a new game session for players. Aliases are useful in the following scenarios:

- If your game doesn't use multi-fleet queues for game session placement, it requests new game sessions by specifying an Amazon GameLift Servers fleet ID. During the life of a game, you'll replace the fleet multiple times, to update a server build, update hosting hardware and operating system, or resolve performance issues. Use an alias to abstract the fleet ID so that you can seamlessly switch player traffic from an existing fleet to the new one.

- If you want to do something other than create a new game session when a game client requests one. For example, you might want to direct players who are using an out-of-date client to an upgrade website.

An alias must specify a routing strategy. There are two types. A *simple* routing strategy routes player traffic to a specified fleet ID, which you can update to redirect traffic. A *terminal* routing strategy passes a message back to the client instead of creating a new game session. You can change an alias's routing strategy at any time.

If you use a queue for game session placement, you don't need an alias to redirect traffic when replacing a fleet. With a queue, you can simply add the new fleet and remove the old fleet. This action is not visible to players, because new game session requests are automatically fulfilled using the new fleet. It doesn't impact existing game sessions. You can identify queues destinations by using either a fleet ID or alias.

Topics

- [Create an Amazon GameLift Servers alias](#)

Create an Amazon GameLift Servers alias

This topic describes how to create an Amazon GameLift Servers alias for use with game session placement.

To create an alias

Use either the Amazon GameLift Servers console or the Amazon Command Line Interface (Amazon CLI) to create an alias.

Console

In the [Amazon GameLift Servers console](#), use the navigation pane to open the **Aliases** page.

1. Choose **Create alias**.
2. Enter an alias **Name**. We recommend including meaningful characteristics in the alias name to help when viewing a list of aliases.
3. Enter an alias **Description** as needed.
4. Choose a **Routing strategy** for the alias.

- a. If you choose a **Simple** routing strategy, select a fleet ID from the list to associate with this alias. The list includes all fleets in either currently selected Amazon Web Services Region. You must create an alias in the same Region as the fleet.
 - b. If you choose a **Terminal** routing strategy, enter a string value that you want Amazon GameLift Servers to return to a game client in response to a game session request. A request with a terminal alias throws an exception with the message embedded.
5. (Optional) Add **Tags** to the alias resource. Each tag consists of a key and an optional value, both of which you define. Assign tags to Amazon resources that you want to categorize in useful ways, such as by purpose, owner, or environment. Choose **Add new tag** for each tag that you want to add.
6. When you're ready to deploy the new fleet, choose **Create**.

Amazon CLI

Use the [create-alias](#) command to create an alias. Amazon GameLift Servers creates the alias resource in your current default Amazon Web Services Region (or you can add a `--region` tag to specify a different Amazon Web Services Region).

At minimum, include an alias name and routing strategy. For a simple routing strategy, specify the ID of a fleet in the same Region as the alias. For a terminal routing strategy, provide a message string.

Managing an Amazon GameLift Servers hosting solution (launch and beyond)

This section provides guidance on how to manage your Amazon GameLift Servers resources for the long term. Successfully launching your multiplayer game is just the beginning. Managing a production hosting solution requires ongoing attention to performance, cost optimization, security updates, and player experience. This section provides guidance on the operational aspects of running Amazon GameLift Servers in production, from day-one launch activities through long-term maintenance and optimization.

Whether you're preparing for launch, managing a live game, or optimizing an established hosting solution, these topics will help you maintain reliable, cost-effective game hosting that scales with your player base.

Manage game hosting resources with Amazon GameLift Servers

Manage Amazon GameLift Servers resources, including builds, fleets, game session queues, and matchmakers. You can use the Amazon GameLift Servers console to scan resource lists, or view resource details including metrics. You can also retrieve resource information using the Amazon SDK or Amazon CLI for Amazon GameLift Servers

Amazon GameLift Servers console features

Use the Amazon GameLift Servers console to view and manage your game hosting resources and ongoing hosting activity in near real time. The console offers a graphical interface for most of the functionality of the service API for Amazon GameLift Servers. You can use the console to:

- **Use the dashboard for a high-level snapshot.** You can see the numbers and current status of all your Amazon GameLift Servers hosting resources and follow links to get details on individual resources.
- **Manage individual hosting resources.** You can create, view, and delete all Amazon GameLift Servers resources, and update their mutable properties. You can also view certain types of hosting activity, such as events and performance metrics.
- **Interact with game and player session activity.** You can track game session and player session activity by fleet, and use this information to troubleshoot game session issues. View details on

a game sessions, view player sessions for each game session, and look up player activity across multiple game sessions. You can also shut down individual game sessions as needed.

Topics

- [Hosting dashboard in the Amazon GameLift Servers console](#)
- [Game server builds](#)
- [Fleets in the Amazon GameLift Servers console](#)
- [Aliases in the Amazon GameLift Servers console](#)
- [Game session queues in the Amazon GameLift Servers console](#)

Hosting dashboard in the Amazon GameLift Servers console

Use the Amazon GameLift Servers console dashboard to get a high-level snapshot about the current status of the Amazon GameLift Servers hosting resources in your Amazon account. The **Amazon GameLift Servers dashboard** provides a view of the following:

- The number of builds in **Ready**, **Initialized**, and **Failed** statuses. Choose **View builds** for details about builds in your current Region.
- The number of fleets in all statuses. Choose **View fleets** for details about fleets in your current Region.
- Your current resources.
- New feature and service announcements.

To open the Amazon GameLift Servers dashboard

- In the [Amazon GameLift Servers console](#), in the navigation pane, choose **Dashboard**.

From the dashboard, you can:

- Prepare your game for launch by choosing **Prepare for launch** and filling out the corresponding launch questionnaire.
- Request service quota increases in preparation for launches or in response to launches by choosing **View service quotas**.

- View blog posts and detailed information about new features by choosing the link in the **Features spotlight**.

GameLift > Dashboard

Dashboard

Build status overview

Viewing data for all builds in N. Virginia region.

Ready
1

Initialized
0

Failed
0

[View builds](#)

Fleet status overview

Viewing data for all fleets in N. Virginia region.

Active
0

Deleting
0

In progress
0

New
0

Error
0

Terminated
0

[View fleets](#)

Resources (1)

[View service quotas](#)

Resource type	Count
Builds	1
Scripts	0
Fleets	0
Aliases	0
Queues	0
Matchmaking rule sets	0
Matchmaking configurations	0

Prepare for your game launch

[Learn more](#)

Fill out a launch questionnaire

Fill out our game launch questionnaire and email it to the GameLift launch team to ensure a smooth launch. The GameLift launch team will verify your GameLift setup and service limits, preparing you for launch.

[Prepare to launch](#)

Features spotlight

Updates on features available in N. Virginia region

March 22, 2022

Updates to Amazon GameLift FlexMatch for greater flexibility

October 28, 2021

New Asia Pacific (Osaka) region and Graviton2 support for Amazon GameLift

Game server builds

The build resource represents your game server software. You upload your build package to Amazon GameLift Servers for deployment to managed EC2 fleets.

View information about game server builds in the Amazon GameLift Servers console or using the or Amazon SDK for Amazon GameLift Servers.

Console

On the **Builds** page of the [Amazon GameLift Servers console](#), you can view information about and manage all the game server builds that you've uploaded to Amazon GameLift Servers for deployment on managed EC2 fleets. In the navigation pane, choose **Hosting, Managed EC2, Builds**.

The **Builds** page shows the following information for each build.

Note

The **Builds** page shows builds in your current Amazon Region only.

- **Name** – The name associated with the uploaded build.
- **Status** – The status of the build. Displays one of three status messages:
 - **Initialized** – The upload hasn't started or is still in progress.
 - **Ready** – The build is ready for fleet creation.
 - **Failed** – The build timed out before Amazon GameLift Servers received the binaries.
- **Creation time** – The date and time that you uploaded the build to Amazon GameLift Servers.
- **Build ID** – The unique ID assigned to the build on upload.
- **Version** – The version label associated with the uploaded build.
- **Operating system** – The OS that the build runs on. The build OS determines which operating system Amazon GameLift Servers installs on a fleet's instances.
- **Size** – The size, in megabytes (MB), of the build file uploaded to Amazon GameLift Servers.
- **Fleets** – The number of fleets deployed with the build.

From this page you can do any of the following:

- View build details. Choose a build's name to open its build details page.
- Create a new fleet from a build. Select a build, and then choose **Create fleet**.
- Filter and sort the build list. Use the controls at the top of the table.
- Delete a build. Select a build, and then choose **Delete**.

Build details

On the **Builds** page, choose a build's name to open its details page. The **Overview** section of the details page displays the same build summary information as the **Builds** page. The **Fleets** section shows a list of fleets that are running the build, including the same summary information as the [Fleets page](#).

Amazon SDK

Use the following Amazon CLI commands to retrieve information about this resource:

- [ListBuild](#)
- [DescribeBuild](#)

Update a game server build

When you deploy your game server build for Amazon GameLift Servers managed EC2 hosting, you upload your game server software and create an Amazon GameLift Servers build resource. After you've created an Amazon GameLift Servers build, you can update the build's metadata, but you can't update the build files themselves. To deploy updates to your game server build, you must upload a new set of files and create a new Amazon GameLift Servers build using the Amazon CLI command [upload-build](#) command. Alternatively, you can use the [create-build](#) command to upload a new build from an Amazon S3 bucket that you control. Then deploy the new build by creating a new fleet for it.

You can update a build's metadata, including the name and description. For these tasks, use the Amazon GameLift Servers console or the [update-build](#) Amazon CLI command.

Automate your game build updates

Follow these tips to help automate and streamline the process of updating game server builds for Amazon GameLift Servers managed fleets:

- **Use game session queues and swap out fleets as needed.** When sending game session requests to Amazon GameLift Servers, specify a game session queue instead of a specific fleet. With queues, you can add fleets with a new build and remove old fleets as needed. For more information, see [Configure game session placement](#).
- **Use aliases to transfer players to a new game build.** When sending game session requests to Amazon GameLift Servers, specify a fleet alias instead of a fleet ID. For more information, see [Create an Amazon GameLift Servers alias](#).

- **Set up for iterative development.** During game development, explore options for setting up a hosted test environment that supports rapid iterative development. See [???](#).

Try out these resources from the [Amazon GameLift Servers Toolkit](#) on Github:

Fast build update tool (for development only)

This tool helps you modify game server builds that are already deployed on computes in a managed EC2 fleet, saving you time during fast development iteration. The tool has several options; you can replace the entire game build or change specific files, and you can manage how to restart game server processes after the updates. You can also use it to update all computes in a fleet or target individual computes.

Visit the Amazon GameLift Servers Toolkit repo in Github to get the [fast build update tool](#) in Github and learn more about how to use it.

Production deployment sample script

This script illustrates how you can automate the process of updating game server builds that are deployed on managed EC2 fleets in production. To use this script, your Amazon GameLift Servers hosting solution must use aliases to abstract fleet IDs. The sample script handles the following basic steps: upload an updated build, create a new build and deploy to a new fleet, redirect player traffic from an existing fleet to the new fleet, and delete the old fleet. Customize the sample script to meet your specific deployment requirements.

Visit the Amazon GameLift Servers Toolkit repo in Github to get the [production deployment sample script](#) in Github and learn more about how to use it.

Fleets in the Amazon GameLift Servers console

The fleet resource represents a set of compute machines that are deployed with and run your game servers. Depending on your hosting solution, you might have managed EC2 fleets, managed container fleets, Anywhere fleets, or a combination.

View information about fleets in the Amazon GameLift Servers console or using the or Amazon SDK for Amazon GameLift Servers.

Console

You can view information on all the fleets created to host your games on Amazon GameLift Servers under your Amazon account. In the console's left-side navigation, find the hosting option you want (Anywhere, Managed EC2, Managed containers), and choose **Fleets**.

For each hosting option, the **Fleets** page lists fleets that are located in your currently selected Amazon Region. From the **Fleets** page, you can create a new fleet or view additional detail on a fleet. A fleet's [detail page](#) contains usage information, metrics, game session data, and player session data. You can also edit a fleet record or delete a fleet.

A **Fleets** page displays the following summary information by default.

- **ID** – An identifier assigned to the fleet. This ID is unique within the Amazon Web Services Region where the fleet is created.
- **Name** – A friendly name given to the fleet.
- **Status** – The status of the fleet: **New**, **Downloading**, **Building**, and **Active**.
- **Creation time** – The date and time the fleet was created.

Note

A fleet displays a warning icon for fleets that were created more than 90 days ago. As a best practice, we recommend replacing 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](#).

- **Fleet type** – The availability of the instances used to host your games, which can potentially impact hosting costs. A managed fleet can use **On-Demand** (always available) or **Spot** (availability varies) instances.
- **Active instances** – The number of EC2 instances in use for the fleet.
- **Desired instances** – The number of EC2 instances to keep active.
- **Game sessions** – The number of active game sessions running in the fleet. The data is delayed by five minutes.
- **Player sessions** – The number of active player sessions in the fleet. The data is delayed by five minutes.

Amazon SDK

Use the following Amazon CLI commands to retrieve information about this resource:

- Managed EC2 fleets
 - [ListFleets](#)
 - [ListCompute](#)
 - [DescribeFleetAttributes](#)
 - [DescribeCompute](#)
 - [DescribeFleetCapacity](#)
 - [DescribeFleetUtilization](#)
- Managed container fleets
 - [ListContainerFleets](#)
 - [DescribeContainerFleet](#)
 - [ListContainerGroupDefinitions](#)
 - [DescribeContainerGroupDefinition](#)
- [ListFleets](#)
- [ListCompute](#)
- [DescribeFleetAttributes](#)
- [DescribeCompute](#)

Update an Amazon GameLift Servers fleet configuration

Use the Amazon GameLift Servers console or the Amazon CLI to update your fleet settings, change remote locations, or delete a fleet. For managed fleets, you can't change a fleet's game server build or instance type. Instead, you must replace the fleet.

Fast Build Update Tool (for development only)

With managed EC2 fleets, to deploy a game server build update, you need to upload each new build to Amazon GameLift Servers and create a new fleet for it.

The Fast Build Update Tool lets you can bypass these steps during development, saving you time and allowing for faster development iteration. With this tool, you can quickly update your game build files across all computes in an existing fleet. The tool has several options; you can replace an entire game build or change 6 specific files, and you can manage how

to restart game server processes after the updates. You can also use it to update individual computes in a fleet.

To get the Fast Build Update Tool and learn more about how to use it, visit the Amazon GameLift Servers Toolkit repo for [The Fast Build Update Tool](#) in Github.

You can update mutable fleet attributes, port settings, and runtime configurations using the Amazon GameLift Servers console or the Amazon CLI. To change scaling limits, see [Auto-scale fleet capacity with Amazon GameLift Servers](#).

Console

1. In the [Amazon GameLift Servers console](#), in the navigation pane, choose **Fleets**.
2. Choose the fleet you want to update. A fleet must be in ACTIVE status before you can edit it.
3. On the Fleet detail page, in any of the following sections, choose **Edit**.
 - **Fleet settings**
 - Change the fleet attributes such as **Name** and **Description**.
 - Add or remove **Metric groups**, that Amazon CloudWatch uses to track aggregated Amazon GameLift Servers metrics for multiple fleets.
 - Update **Resource creation limit** settings.
 - Turn game session protection on or off.
 - **Runtime configuration** – You can change any of the following settings of your runtime configurations and add or remove runtime configurations.
 - Change the **Launch path** of your game server.
 - Add, remove, or change optional **Launch parameters**.
 - Change the number of **Concurrent processes** that your game servers run.
 - **Game session activation** – Change how you want server processes to run and host game sessions by updating **Max concurrent game session activations** and **New activation timeout**.
 - **EC2 port settings** – Update the IP addresses and port ranges that allow inbound access to the fleet.
4. Choose **Confirm** to save changes.

Amazon CLI

Use the following Amazon CLI commands to update a fleet:

- [update-fleet-attributes](#)
- [update-fleet-port-settings](#)
- [update-runtime-configuration](#)

Update fleet locations

You can add or remove a fleet's remote locations using the Amazon GameLift Servers console or the Amazon CLI. You can't change a fleet's home Region.

Console

1. In the [Amazon GameLift Servers console](#), in the navigation pane, choose **Fleets**.
2. Choose the fleet you want to update. A fleet must be in ACTIVE status before you can edit it.
3. On the Fleet detail page, choose the **Locations** tab to view the fleet's locations.
4. To add new remote locations, choose **Add** and select the locations you want to deploy instances to. This list doesn't include instances where the fleet's instance type isn't available.
5. With new locations selected, choose **Add**. Amazon GameLift Servers adds the new locations to the list, with status set to NEW. Amazon GameLift Servers then begins provisioning an instance in each added location and preparing it to host game sessions.
6. To remove existing remote locations from the fleet, use the check boxes to select one or more listed locations.
7. With one or more fleets selected, choose **Remove**. The removed locations remain in the list, with status set to DELETING. Amazon GameLift Servers then begins the process of terminating activity in the removed location. If there are active instances that are hosting game sessions, Amazon GameLift Servers uses the game server termination process to gracefully end game sessions, terminate game servers, and shut down instances.

Amazon CLI

Use the following Amazon CLI commands to update fleet locations:

- [create-fleet-locations](#)
- [delete-fleet-locations](#)

Delete a fleet

You can delete a fleet when you no longer need it. Deleting a fleet permanently removes all data associated with game sessions and player sessions, and collected metric data. As an alternative, you can retain the fleet, disable auto scaling, and manually scale the fleet to 0 instances.

Note

If the fleet has a VPC peering connection, first request authorization by calling [CreateVpcPeeringAuthorization](#). Amazon GameLift Servers deletes the VPC peering connection during fleet deletion.

You can use either the Amazon GameLift Servers console or the Amazon CLI tool to delete a fleet.

Console

1. In the [Amazon GameLift Servers console](#), in the navigation pane, choose **Fleets**.
2. Choose the fleet you want to delete. You can only delete fleets in ACTIVE or ERROR status.
3. Choose **Delete**.
4. In the **Delete fleet** dialog box, confirm the deletion by entering **delete**.
5. Choose **Delete**.

Amazon CLI

Use the following Amazon CLI command to delete a fleet:

- [delete-fleet](#)

Fleet details in the Amazon GameLift Servers console

Access a **Fleet** detail page from the dashboard or the **Fleets** page by choosing the fleet name.

On the fleet details page you can take the following actions:

- Update a fleet's attributes, port settings, and runtime configuration.
- Add or remove fleet locations.
- Change fleet capacity settings.
- Set or change target-tracking auto scaling.
- Delete a fleet.

Details

Fleet settings

- **Fleet ID** – An identifier assigned to the fleet. This ID is unique within the Amazon Web Services Region where the fleet is created.
- **Fleet name** – A friendly name given to the fleet.
- **ARN** – A unique identifier assigned to this fleet. A fleet's ARN identifies it as an Amazon GameLift Servers resource and specifies the region and Amazon account.
- **Description** – A short identifiable description of the fleet.
- **Status** – Current status of the fleet, which may be **New**, **Downloading**, **Building**, and **Active**.
- **Creation time** – The date and time when the fleet was created.

Note

A fleet displays a warning icon for fleets that were created more than 90 days ago. As a best practice, we recommend replacing 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](#).

- **Compute type** – The type of compute used to host your games. A fleet can be a **Managed EC2**, **Managed container**, or **Anywhere** fleet.
- **Operating system** – The Amazon Machine Image (AMI) that's deployed to every instance in the fleet. The AMI version is the latest version available at the time the fleet was created.
- **Certificate type** – The status of the fleet's ability to use a TLS certificate for authenticating a game server and encrypting all client/server communication.

- **Fleet type** – The availability of the instances used to host your games, which can potentially impact hosting costs. A fleet can use **On-Demand** (always available) or **Spot** (availability varies) instances.
- **EC2 instance type** – Amazon EC2 [instance type](#) selected for the fleet when it was created.
- **Instance role ARN** – An Amazon IAM role that manages access to your other Amazon resources, if one was provided during fleet creation.
- **Instance role credentials provider** – An Amazon IAM role that manages access to your other Amazon resources, if one was provided during fleet creation.
- **Metric group name** – The group used to aggregate metrics for multiple fleets.
- **Game scaling protection policy** – Indicates whether the fleet's game session protection is enabled, which prevents active game sessions from ending prematurely during a scale-down event.
- **Maximum game sessions per player** – The maximum number of sessions a player can create during the **Policy period**.
- **Policy period** – The length of time used to limit a player's number of sessions created.

Build details

The **Build details** section displays the build hosted on the fleet. Select the build name to see the full build detail page.

Runtime configuration

The **Runtime configuration** section displays the server processes to launch on each instance. It includes the path for the game server executable and optional launch parameters.

Game session activation

The **Game session activation** section displays the number of server processes that launch at the same time and how long to wait for the process to activate before terminating it.

EC2 port settings

The **Ports** section displays the fleet's connection permissions, including IP address and port setting ranges.

Metrics

The **Metrics** tab displays a graphical representation of fleet metrics over time. For more information about using metrics in Amazon GameLift Servers, see [Monitor Amazon GameLift Servers with Amazon CloudWatch](#).

Events

The **Events** tab provides a log of all events that have occurred on the fleet, including the event code, message, and time stamp. See [Event](#) descriptions in the Amazon GameLift Servers API Reference.

Scaling

The **Scaling** tab contains information about fleet capacity, including the current status and capacity changes over time. It also provides tools to update capacity limits and manage auto scaling.

Scaling capacity

View current fleet capacity settings for each fleet location. For more information about changing limits and capacity, see [Scaling game hosting capacity with Amazon GameLift Servers](#).

- **Amazon Location** – Name of a location where fleet instances are deployed.
- **Status** – Hosting status of the fleet location. Location status must be ACTIVE to be able to host games.
- **Min size** – The smallest number of instances that must be deployed in the location.
- **Desired instances** – The target number of active instances to maintain the location. When active instances and desired instances aren't the same, a scaling event is started to start or shut down instances as needed until active instances equals desired instances.
- **Max size** – The most instances that can be deployed in the location.
- **Available** – The service limit on instances minus the number of instances in use. This value tells you the maximum number of instances that you can add to the location.

Auto scaling policies

This section covers information about auto scaling policies that are applied to the fleet. You can set up or update a target-based policy. The fleet's rule-based policies, which must be defined using the

Amazon SDK or CLI, are displayed here. For more information about scaling, see [Auto-scale fleet capacity with Amazon GameLift Servers](#).

Scaling history

View graphs of capacity changes over time.

Locations

The **Locations** tab lists all locations where fleet instances are deployed. Locations include the fleet's home Region and any remote locations that have been added. You can add or remove locations directly in this tab.

- **Location** – Name of a location where fleet instances are deployed.
- **Status** – Hosting status of the fleet location. Location status tracks the process of activating the first instances in the location. Location status must be ACTIVE to be able to host games.
- **Active instances** – The number of instances with server processes running on the fleet location.
- **Active servers** – The number of game server processes able to host game sessions in the fleet location.
- **Game sessions** – The number of game sessions active on instances in the fleet location.
- **Player sessions** – The number of player sessions, which represent individual players, that are participating in game sessions that are active in the fleet location.

Game sessions

The **Game sessions** tab lists past and present game sessions hosted on the fleet, including some detail information. Choose a game session ID to access additional game session information, including player sessions. For more information about player sessions, see [Game and player sessions in the Amazon GameLift Servers console](#).

Game and player sessions in the Amazon GameLift Servers console

You can use the Amazon GameLift Servers console to work with games sessions and player sessions. For more information about game sessions and player sessions, see [Amazon GameLift Servers and the player experience](#). The Amazon GameLift Servers console provides information and tools to help you investigate issues with your game sessions.

What you can do:

- Explore game session and player session activity that is hosted on a specific fleet.
- Look up game session activity for a specific player across multiple fleets.
- Shut down a specific game session.

View game session detail

Game session and player session data is organized by the fleet that hosts the game session.

To access game session and player session information

1. In the [Amazon GameLift Servers console](#), open the left navigation pane. Select a hosting solution type and open the **Fleets** page. For example:
 - **Hosting, Anywhere, Fleets**
 - **Hosting, Managed EC2, Fleets**
 - **Hosting, Managed containers, Fleets**
2. Each **Fleets** page displays the list of fleets for your currently selected Amazon Web Services Region. Choose the fleet that you want to view game session data for.
3. In the fleet's detail page, open the **Game sessions** tab. This tab lists all game sessions that were hosted on the fleet, along with summary information.
4. Choose a game session from the list to view additional information.
5. If the game session includes player session data, choose **View player sessions** to open the player sessions lookup tool with the game session ID automatically filled in.

The **Game sessions** detail include the following information:

- **Status** – Game session status.
 - **Activating** – The instance is initiating a game session.
 - **Active** – A game session is running and available to receive players, depending on the session's [player creation policy](#).
 - **Terminated** – the game session has ended.
- **ARN** – The Amazon Resource Name of the game session.
- **Name** – Name generated for the game session.
- **Location** – The location that Amazon GameLift Servers hosted the game session in.

- **Creation time** – Date and time that Amazon GameLift Servers created the stream session.
- **Ending time** – Date and time that the game session ended.
- **DNS name** – The host name of the game session.
- **IP address** – IP address specified for the game session.
- **Port** – Port number used to connect to the game session.
- **Creator ID** – A unique identifier of the player that initiated the game session.
- **Player session creation policy** – Indicates if the game session is accepting new players.
- **Game scaling protection policy** – The type of game session protection to set on all new instances that Amazon GameLift Servers starts in the fleet.

Game data

Game property data, formatted as a string, to send to your game session on start.

Game properties

Game property data, formatted as key/value pairs, to send to your game session on start.

Matchmaking data

If the game session was created with FlexMatch, matchmaking data describes information about the matchmaking configuration and rule set. This includes each match's player attributes and team assignments. Data is in JSON format. For more information about FlexMatch matchmaking, see [Build a matchmaker](#).

Look up player session data

If your game hosting solution uses player sessions and provides unique player IDs, you can explore player-specific activity for past or present game sessions across multiple fleets. Open the Player session lookup tool using either of the following methods:

- In the Amazon GameLift Servers console, open the left navigation pane and choose **Player session lookup** and select the filter type to use.
- When viewing a fleet's game session details, choose **View player sessions**. The lookup tool opens with the game session with the game session ID filter preselected and the game session value filled in.

When using the lookup tool, you can provide any the following information:

- A player session ID to get information on a specific player session.
- A game session ID to get information on all player sessions for the requested game session. The results represent all players who reserved a slot or connected to the game session. You can optionally filter results by player session status.
- A player ID to get information on all player sessions for the requested player. The results represent all game sessions that the player participated in.

 Note

The lookup tool searches for all player session activity across the currently selected Amazon Web Services Region. If you have multiple fleets in the Region, the results include player session activity across all fleets. For multi-location fleets, the results also include player session activity across the fleets' remote locations.

The following player session data is collected for each game session:

- **Player session ID** – The identifier assigned to the player session.
- **Player ID** – A unique identifier for the player. Choose this ID to get additional player information.
- **Game session ID** – The identifier assigned to the game session.
- **Fleet ID** – The identifier assigned to the fleet that hosted the game session.
- **Status** – The status of the player session. The following are possible statuses:
 - **Reserved** – Player session has been reserved, but the players isn't connected.
 - **Active** – Player session is connected to the game server.
 - **Completed** – Player session has ended; player is no longer connected.
 - **Timed Out** – Player failed to connect.
- **Creation time** – The time the player connected to the game session.
- **Ending time** – The time the player disconnected from the game session.
- **Connection data** – IP address, DNS name and port that the player used to connect to the game session.
- **Player data** – Information about the player that was provided during player session creation.

Shut down a game session

Use the Amazon GameLift Servers console to shut down a specific game session. This feature gives you a simple and fast method for locating a game session and sending a signal to terminate it. Another termination method requires that you find the fleet instance where the game session is running, remotely access the instance, and manually shut down the game session.

You can shut down a game session for any reason. The most common reason is to resolve a game session that's failing to shut down naturally. As a result, the hosting resource for the game session can't be freed to host a new game session, and the fleet's hosting capacity is degraded.

Note

This feature relies on certain configuration settings for your hosting solution. It has the following limitations:

- The game session must be hosted on a fleet that's running a game server build with server SDK for Amazon GameLift Servers v5 or greater. If your game servers are deployed with an older version, you need to use remote access to delete the game session.
- If the game session is hosted on an Anywhere fleet, the fleet must be using the Amazon GameLift Servers Agent to manage game server processes.

To terminate a game session

1. In the [Amazon GameLift Servers console](#), open the left navigation pane. Select a hosting solution type and open the **Fleets** page. For example:
 - **Hosting, Anywhere, Fleets**
 - **Hosting, Managed EC2, Fleets**
 - **Hosting, Managed containers, Fleets**
2. Each **Fleets** page displays the list of fleets for your currently selected Amazon Web Services Region. Choose the fleet that's hosting the game session that you want to terminate.
3. In the fleet's detail page, open the **Game sessions** tab. In the list of game sessions, select the one that you want to terminate, and choose the **Terminate** button.
4. In the **Terminate game session?** window, verify that you're shutting down the right game session and choose a termination method.

- Normal game session shutdown – This option sends a signal to the server process that's hosting the game session to shut down. If your game server build was properly integrated for Amazon GameLift Servers, the server process initiates its game session shutdown sequence, notifies Amazon GameLift Servers that it's ending, and stops. Depending on your game design, the shutdown sequence might include steps to gracefully complete the game session, such as saving data and notifying active players. This method might require a small delay to complete the game session shutdown sequence.
 - Immediate game session shutdown – This option sends a signal to a process manager to shut down the server process that's hosting the game session. This option bypasses the normal game session shutdown. It's able to terminate the game session even when the server process is unable to respond.
5. Confirm game session termination. You can track the shutdown progress on the **Game sessions** console page. The game session status will change to "Terminating", and then to "Terminated" when shutdown is complete.

Related topics

- You can also shutdown game sessions using the Amazon SDK and the Amazon CLI. For more details and examples, see the Amazon GameLift Servers API Reference topic [TerminateGameSession](#).
- For more information on game server integration and how a server process responds to signals from the Amazon GameLift Servers service, see [Add Amazon GameLift Servers to your game server with the server SDK](#).

Update an Amazon GameLift Servers managed container fleet

You can update most of the properties of a managed container fleet, including the container group definitions. Depending on the settings being updated, a fleet update might initiate a new fleet deployment. In a fleet deployment, all instances in the fleet are removed and replaced with instances with the new configuration. Settings that require a deployment include:

- Container group definitions, including updates to container images
- Connection port ranges and inbound permissions
- Log configuration

You can track the status of fleet deployments in the [Amazon GameLift Servers console](#) or the Amazon Command Line Interface (Amazon CLI) to create a container fleet.

Note

You can't update a fleet's runtime environment. At fleet creation, the fleet's Amazon Machine Image (AMI) is set to the latest available version of the Linux AMI. All container images that are deployed to this fleet must be compatible with this version. To change the fleet's AMI or upgrade to a newer version, you must create a new fleet. As a best practice, we recommend replacing your fleets every 30 days to maintain a secure and up-to-date runtime environment for your hosted game servers. For more guidance, see [Security best practices for Amazon GameLift Servers](#).

Console

In the [Amazon GameLift Servers console](#), select the Amazon Web Services Region where you want to create the fleet. The container group definitions must be in the same region where you want to create the fleet.

Open the console's left navigation bar and choose **Managed containers: Fleets**. On the managed container fleets page, select a fleet from the list and choose **Edit**.

1. Update container fleet settings as needed. When you're finished, choose **Create**.
2. If your updates require a fleet deployment, you're asked to specify deployment options as follows:
 - **Game session protection.** You can choose to protect fleet instances that have active game sessions (safe deployment). With this setting, the fleet instances aren't replaced until after the game sessions end. Alternatively, you can choose to replace fleet instances regardless of game session activity (unsafe deployment). Unsafe deployments are useful during development and testing phases in order to reduce deployment time.
 - **Minimum healthy percentage.** You can manage how quickly the fleet's instances are replaced. Use this setting to maintain a minimal amount of healthy tasks during deployment. A low value prioritizes deployment speed, while a high value ensures that game server availability remains high throughout the deployment.
 - **Deployment failure strategy.** Decide what actions to take if a deployment fails. A deployment failure means that some of the updated containers have failed status checks

and are considered impaired. You can set deployments to automatically roll back all fleet instances to the previously deployed state. Alternatively you can choose to maintain some of the impaired fleet instances for use in debugging.

If your request is successful, the console displays the **Deployments** tab for the managed container fleet. Use this tab to track the status of each deployment. If you start a new deployment for the fleet, this action automatically cancels any deployment that is currently in process for the fleet.

Amazon CLI

To create a container fleet with the Amazon CLI, open a command line window and use the `update-container-fleet` command. For more information about this command, see [update-container-fleet](#) in the *Amazon CLI Command Reference*.

The following example updates an existing container fleet with the following characteristics:

- It updates the game server container group definition to use version 2.
- It specifies safe deployment options.

```
{
  "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"
}
```

Update a container group definition for an Amazon GameLift Servers container fleet

You can update most of the properties of a container group definition, including the individual container definitions. Container group definitions have a version number. When you update a container group definition, Amazon GameLift Servers saves the update and increments the

definition's version number. When configuring a container fleet, you can specify which version of a container group definition to deploy.

After updating a container group definition, you can deploy the new version to a new or existing container fleet.

Update a game server container group definition

This topic describes how to update game server container group definition using the Amazon GameLift Servers console or Amazon CLI tools. For more detailed information on optional features, see [Customize an Amazon GameLift Servers container fleet](#).

To update a container group definition:

Console

In the [Amazon GameLift Servers console](#), select the Amazon Web Services Region where you want to create the container group.

Open the console's left navigation bar and choose **Managed containers: Group definitions**. On the Container groups definition page, choose a container group definition and version to update.

After you've saved your updates, you can use the new version to create a new container fleets or you can deploy the updates to an existing container fleet.

Step 1: Define container group definition details

- You can update the total memory and vCPU limit settings.

Step 2: Add container definitions

You can make the following container definition updates:

- Update existing container definitions.
 - Add new support container definitions.
 - Remove support container definitions.
1. You can update the **ECR image URI**. Make sure to update the **Server SDK version** setting to match the new image.

2. You can update the **Internal container port range** as needed. Changes you make to these settings might impact a container fleet's connection port settings when these changes are deployed to a fleet. For more details, see [Configure network connections](#).

Step 3: Configure dependencies

- You can change dependencies as needed. For more information, see [Set container dependencies](#).

Step 3: Review and create

- Review your container group definition updates. Use **Edit** to make additional changes in any section. When you're finished, choose **Create** to generate a new version of the container group definition.

If your request is successful, the console displays the detail page for the new container group definition resource. Initially the status is **COPYING**, as Amazon GameLift Servers starts taking snapshots of all the container images for the group. When this phase is complete, the container group definition status changes to **READY**. A container group definition must be in **READY** status before you can create a container fleet with it.

Amazon CLI

When you use the Amazon CLI to create or update a container group definition, maintain your container definition configurations in a separate JSON file. You can reference the file in your CLI command. See [???](#) for schema examples.

When updating a definition, you only need to specify the values you want to update. Amazon GameLift Servers retains any values that you don't include in your update request. If you're changing a container definition. However, when changing a container definition, provide a complete set.

To update a container group definition

To update a new container group definition, use the `update-container-group-definition` CLI command. For more information about this command, see [update-container-group-definition](#) in the *Amazon CLI Command Reference*.

Example: game server container group

You can specify a container group definition version when retrieving, updating, or deleting a container group definition, or when creating or updating a container fleet. Each container group definition has a version property. In addition, the definition's ARN value specifies the version number.

This example illustrates a request for a change to a game server container group definition. It assumes that you've created a JSON file with the container definitions for this group. This example uses the ARN value for definition name, and specifies that the update is to version 1.

```
aws gamelift update-container-group-definition \
  --name arn:aws:gamelift:us-west-2:111122223333:containergroupdefinition/
  MyAdventureGameContainerGroup:1 \
  --operating-system AMAZON_LINUX_2023 \
  --container-group-type GAME_SERVER \
  --total-memory-limit-mebibytes 4096 \
  --total-vcpu-limit 1 \
  --container-definitions file://SimpleServer.json
```

Clone a container group definition

You can use the Amazon GameLift Servers console to clone an existing container group definition.

To clone a container group

1. In the [Amazon GameLift Servers console](#), go to the left navigation pane and choose **Container groups**.
2. On the **Container groups** list page, select the existing container group that you want to clone. After you select a container group, the **Clone** button is active.
3. Choose **Clone**. This action opens the container group creation wizard with pre-filled settings.
4. Enter a new name for the cloned container group. Container group in the same region must have unique names.
5. Step through the container group and container definition pages, review, and **Create** the new container group.

Delete a container group definition for an Amazon GameLift Servers container fleet

You have several options for deleting a container group definition. When you delete a container group definition, this action also deletes all the container definitions in the container group.

Container group definitions can have multiple versions. Container group versions have the same name, but have a different version number. Container group definition ARNs specify both the name and version.

You can specify a container group definition version when retrieving, updating, or deleting a container group definition, or when creating or updating a container fleet. Each container group definition has a version property. In addition, the definition's ARN value specifies the version number.

There are several ways to delete container group definitions:

- You can delete all versions of a specific definition.
- You can delete one particular version of a specific definition.
- You can keep some number of newest versions and delete the older versions of a specific definition. For example, you might delete all versions older than version 5.

You can delete a container group definition version only if it's not being used in a container fleet.

To delete a container group definition

Console

In the [Amazon GameLift Servers console](#), select the Amazon Web Services Region where you want to create the container group.

Open the console's left navigation bar and choose **Managed containers: Group definitions**. On the Container group definitions page, select the definition you want to modify and choose **Delete**.

You're prompted to select the type of deletion you want to make and specify additional settings depending on the delete type.

Amazon CLI

- To delete a container group definition, use the `delete-container-group-definition` CLI command and provide values for the type of delete you want to do. For more information about this command, see [delete-container-group-definition](#) in the *Amazon CLI Command Reference*.

This example illustrates a request to delete all versions of game server container group definition older than version 5.

Example

```
aws gamelift delete-container-group-definition \
  --name MyAdventureGameContainerGroup \
  --version-count-to-retain 5 \
```

Aliases in the Amazon GameLift Servers console

The alias resource represents a pointer to a fleet or terminal destination. They are used to direct player traffic to specific hosting destinations. Use an alias to abstract a destination so that you can more easily redirect traffic when you update your hosting fleets.

View information about aliases in the Amazon GameLift Servers console or using the or Amazon SDK for Amazon GameLift Servers.

Console

The **Alias** page displays information about Amazon GameLift Servers aliases that direct traffic specific hosting destinations. To view aliases, choose **Hosting, Aliases** in the navigation pane.

You can do the following on the aliases page:

- Create a new alias. Choose **Create alias**.
- Filter and sort the aliases table. Use the controls at the top of the table.
- View alias details. Choose an alias name to open the alias detail page.
- Delete an alias. Choose an alias and then choose **Delete**.

Alias details

The alias details page displays information about the alias.

From this page you can:

- Edit an alias. Choose **Edit**.
- View the fleets you associated with the alias.
- Delete an alias. Choose **Delete**.

Alias detail information includes:

- **ID** – The unique number used to identify the alias.
- **Description** – The description of the alias.
- **ARN** – The Amazon Resource Name of the alias.
- **Creation** – The date and time the alias was created.
- **Last updated** – The date and time that the alias was last updated.
- **Routing type** – The routing type for the alias, which can be one of these:
 - **Simple** – Routes player traffic to a specified fleet ID. You can update the fleet ID for an alias at any time.
 - **Terminal** – Passes a message back to the client. For example, you can direct players who are using an out-of-date client to a location where they can get an upgrade.
- **Tags** – Key and value pairs used to identify the alias.

Amazon SDK

Use the following Amazon CLI commands to retrieve information about this resource:

- [ListAliases](#)
- [DescribeAlias](#)

Edit an alias

You can edit an alias using the Amazon GameLift Servers console or with the Amazon CLI command [update-alias](#).

This topic describes how to edit an Amazon GameLift Servers alias for use with game session placement. You can make the following edits:

To edit an alias

Use either the Amazon GameLift Servers console or the Amazon Command Line Interface (Amazon CLI) to edit an alias.

Console

In the [Amazon GameLift Servers console](#), use the navigation pane to open the **Aliases** page.

1. Select the alias you want to edit and choose **Edit**. If you don't see the alias you want to edit, check your currently selected Amazon Web Services Region.
2. On the **Edit alias** page, you can make the following edits:
 - Change the alias name.
 - Change the alias description.
 - Change the routing strategy from simple to terminal, or from terminal to simple.
 - For an alias with a simple routing strategy, change the fleet ID the alias is associated with.
 - For an alias with a terminal routing strategy, change the message text.
3. Choose **Save changes**. When updating the fleet ID for an alias with a simple routing strategy, it may take up to 2 minutes for the transition to complete. During this time, new game session placements might take place on the old fleet.

Amazon CLI

Use the [update-alias](#) command to make changes to an alias resource. You can update an alias resource in your current default Amazon Web Services Region, or you can add a `--region` tag to specify a different Amazon Web Services Region.

You can change the following properties:

- Alias name.
- Alias description.
- Routing strategy type. Be sure to provide a fleet ID or a message string for the new routing strategy.
- Fleet ID for an existing simple routing strategy. The fleet ID must be in the same region as the alias.

- Message string for an existing terminal routing strategy.

Game session queues in the Amazon GameLift Servers console

The queue resource represents your configuration of the game session placement component. It determines how Amazon GameLift Servers searches for and chooses the best possible compute resource to host a new game session, based on your configuration choices.

View information about game session queues in the Amazon GameLift Servers console or using the or Amazon SDK for Amazon GameLift Servers.

Console

You can view information on all queues, which are used to process requests for new game sessions. The queues page shows game session queues in the currently selected Amazon Web Services Region. From the **Queues** page, you can create a new queue, delete existing queues, or open a details page for a selected queue. Each queue details page contains the queue's configuration and metrics data. For more information about queues, see [Configure game session placement](#).

The queues page displays the following summary information for each queue.

- **Queue name** – The name assigned to the queue. Requests for new game sessions specify a queue by this name.
- **Queue timeout** – Maximum length of time, in seconds, that a game session placement request remains in the queue before timing out.
- **Destinations in queue** – Number of fleets listed in the queue configuration. Amazon GameLift Servers places new game sessions on any fleet in the queue.

View queue details

You can access detailed information on any queue, including the queue configuration and metrics. To open a queue details page, go to the **Queues** page and choose a queue name.

The queue detail page displays a summary table and tabs containing additional information. On this page you can do the following:

- Update the queue's configuration, list of destinations and player latency policies. Choose **Edit**.

- Delete a queue. After you delete a queue, all requests for new game sessions that reference that queue name will fail. Choose **Delete**.

 Note

To restore a deleted queue, create a new queue with the deleted queue's name.

Details

The **Overview** section displays the queue's Amazon Resource Name (**ARN**) and the **Timeout**. You can use the ARN when referencing the queue in other actions or areas of Amazon GameLift Servers. The timeout is the maximum length of time, in seconds, that a game session placement request remains in the queue before timing out.

Event notification

The **Event notification** section lists the **SNS topic** Amazon GameLift Servers publishes event notifications to and the **Event data** that is added to all events created by this queue.

Tags

The **Tags** table displays the keys and values used to tag the resource. For more information about tagging, see [Tagging Amazon resources](#).

Metrics

The **Metrics** tab shows a graphical representation of queue metrics over time.

Queue metrics include a range of information describing placement activity across the queue, including successful placements organized by Region. You can use Region data to understand where you are hosting your games. Regional placement metrics can help to detect issues with the overall queue design.

Queue metrics are also available in Amazon CloudWatch. For descriptions of available metrics, see [Amazon GameLift Servers metrics for queues](#).

Destinations

The **Destinations** tab shows all fleets or aliases listed for the queue.

When Amazon GameLift Servers searches the destinations for available resources to host a new game session, it searches the default order listed here. As long as there is capacity on the first destination listed, Amazon GameLift Servers places new game sessions there. You can have individual game session placement requests override the default order by providing player latency data. This data tells Amazon GameLift Servers to search for an available destination with the lowest average player latency. For more information about designing your queues, see [Customize a game session queue](#).

Session placement

Configuration settings that customize the placement process for your game.

Player latency policies

The **Player latency policies** section shows all policies that the queue uses. The tables lists the policies in the order they're enforced.

Locations

The **Locations** section shows the locations that this queue can put a game session in.

Priority

The **Priority** section shows the order that the queue evaluates a game sessions details.

Location order

The **Location order** section shows the default order that the queue uses when placing game sessions. The queue uses this order if you haven't defined other types of priority.

Amazon SDK

Use the following Amazon CLI commands to retrieve information about this resource:

- [DescribeGameSessionQueues](#)

Prepare for launch with Amazon GameLift Servers hosting

Use the following checklists to validate each deployment phase of your game. Items marked **[Critical]** are critical for your production launch.

Download and complete the Amazon GameLift Servers launch questionnaire, which is available in the [Amazon GameLift Servers console](#). We want every game developer using Amazon GameLift

Servers to have a smooth launch day, and the requested information helps us help you prepare for upcoming load testing, soft launch or public launch. Plan to submit the completed questionnaire at least three (3) months before conducting your first load test.

Topics

- [Get your game ready](#)
- [Prepare for testing](#)
- [Prepare for launch](#)
- [Plan for post-launch updates](#)

Get your game ready

- **[Critical]** Verify that you've completed all the [development roadmap steps](#) for your hosting solution, and that you have all the required components in place, including an integrated game server, a backend service for game clients, hosting fleets, and a game session placement method (such as a queue).
- **[Critical]** [Create Amazon Identity and Access Management \(IAM\) roles](#) that allow your game server to access other Amazon resources while running.
- **[Critical]** Design and implement failover to other hosting resources as needed.
- [Plan the rollout of fleets to your target locations](#), considering your game's queue and fleet structure.
- [Automate your deployment](#) using infrastructure as code (IaC) with Amazon CloudFormation and the Amazon Cloud Development Kit (Amazon CDK).
- [Collect logs and analytics](#) using Amazon CloudWatch and Amazon Simple Storage Service (Amazon S3).

Prepare for testing

- **[Critical]** [Request increases for Amazon GameLift Servers service quotas](#) and other Amazon Web Services service quotas so that your live environment can scale up to production needs.
- **[Critical]** Verify that the open ports on live fleets match the range of ports that your servers could use.
- **[Critical]** Close RDP port 3389 and SSH port 22.

- Develop a plan for the DevOps management of your game. If you're using Amazon CloudWatch Logs or Amazon CloudWatch custom metrics, define alarms for severe or critical problems on the server fleet. Simulate failures and test the runbooks.
- Verify that the compute resources you're using can support the number of server processes that you want to run concurrently on each compute.
- [Tune your scaling policy](#) to be more conservative at first and provide more idle capacity than you think you need. You can optimize for cost later. Consider the use of target-based scaling policy with 20 percent idle capacity.
- For FlexMatch, use [latency rules](#) to match players who are geographically near each other. Test how this behaves under load with synthetic latency data from your load test client.
- Load test your player authentication and game session infrastructure to see if it scales effectively to meet demand.
- Verify that a server left running for several days can still accept connections.
- Raise your Amazon Web Services Support plan level to Business or Enterprise so that Amazon Web Services can respond to you during problems or outages.

Prepare for launch

- **[Critical]** [Set the fleet protection policy](#) to full protection on all live fleets so that scaling down doesn't stop active game sessions.
- **[Critical]** [Set fleet maximum sizes](#) high enough to accommodate peak anticipated demand, at minimum. We recommend that you double your maximum size for unanticipated demand.
- Encourage your entire development team to participate in the launch event and monitor your game launch in a launch room.
- Monitor player latency and player experience.

Plan for post-launch updates

- [Tune scaling policy](#) to minimize idle capacity based on player usage.
- [Modify FlexMatch rules](#) or [add hosting locations](#) based on player latency data and revised requirements.
- Optimize the runtime configuration to run as many games sessions as possible on each computing resource. Maximizing performance efficiency in this way can directly affect your fleet costs, because you might be able to run more server processes with the same compute resources.

- [Use your analytics data](#) to drive continued development, improve player experience and game longevity, and optimize monetization.

Manage Amazon GameLift Servers hosting resources using Amazon CloudFormation

You can use Amazon CloudFormation to manage your Amazon GameLift Servers resources. In Amazon CloudFormation, you create a template that models each resource and then use the template to create your resources. To update resources, you make the changes to your template and use Amazon CloudFormation to implement the updates. You can organize your resources into logical groups, called stacks and stack sets.

Using Amazon CloudFormation to maintain your Amazon GameLift Servers hosting resources offers a more efficient way to manage sets of Amazon resources. You can use version control to track template changes over time and coordinate updates made by multiple team members. You can also reuse templates. For example, when deploying a game across multiple Regions, you might use the same template to create identical resources in each Region. You can also use these templates to deploy the same sets of resources in another partition.

For more information about Amazon CloudFormation, see the [Amazon CloudFormation User Guide](#). To view template information for Amazon GameLift Servers resources, see the [Amazon GameLift Servers resource type reference](#).

Best practices

For detailed guidance on using Amazon CloudFormation, see the [Amazon CloudFormation best practices](#) in the *Amazon CloudFormation User Guide*. In addition, these best practices have special relevance with Amazon GameLift Servers.

- **Consistently manage your resources through Amazon CloudFormation.** If you change your resources outside of Amazon CloudFormation your resources will get out of sync with your resource templates.
- **Use Amazon CloudFormation stacks and stack sets to efficiently manage multiple resources.**
 - Use stacks to manage groups of connected resources. For example, a stack that contains a build, a fleet that references the build, and an alias that references the fleet. If you update your template to replace a build, Amazon CloudFormation replaces the fleets connected to the

build. Amazon CloudFormation then updates the existing aliases to point to the new fleets. For more information, see [Working with stacks](#) in the *Amazon CloudFormation User Guide*.

- Use Amazon CloudFormation stack sets if you're deploying identical stacks across multiple regions or Amazon accounts. For more information, see [Working with stack sets](#) in the *Amazon CloudFormation User Guide*.
- **If you are using Spot Instances, include an On-Demand Fleet as a back-up.** We recommend setting up your templates with two fleets in each region, one fleet with Spot Instances, and one fleet with On-Demand Instances.
- **Group your location-specific resources and global resources into separate stacks when you are managing resources in multiple locations.**
- **Place your global resources close to the services that use it.** Resources like queues and matchmaking configurations tend to receive a high volume of requests from specific sources. By placing your resources close to the source of those requests, you minimize the request travel time and can improve overall performance.
- **Place your matchmaking configuration in the same Region as the game session queue that it uses.**
- **Create a separate alias for each fleet in the stack.**

Using Amazon CloudFormation stacks

We recommend the following structures to use when setting up Amazon CloudFormation stacks for Amazon GameLift Servers resources. Your optimal stack structure varies depending on if you are deploying your game in one location or multiple locations.


Stacks for a single location

To manage Amazon GameLift Servers resources in a single location, we recommend a two-stack structure:

- **Support stack** – This stack contains resources that your Amazon GameLift Servers resources depend on. At a minimum, this stack should include the S3 bucket where you store your custom game server or Realtime script files. The stack should also include an IAM role that gives Amazon GameLift Servers permission to retrieve your files from the S3 bucket when creating an Amazon GameLift Servers build or script resource. This stack might also contain other Amazon resources that are used with your game, such as DynamoDB tables, Amazon Redshift clusters, and Lambda functions.

- **Amazon GameLift Servers stack** – This stack contains all of your Amazon GameLift Servers resources, including the build or script, a set of fleets, aliases, and game session queue. Amazon CloudFormation creates a build or script resource with files stored in the S3 bucket location and deploys the build or script to one or more fleet resources. Each fleet should have a corresponding alias. The game session queue references some or all of the fleet aliases. If you are using FlexMatch for matchmaking, this stack also contains a matchmaking configuration and rule set.

The diagram below illustrates a two-stack structure for deploying resources in a single Amazon Region.




Stacks for multiple regions


When deploying your game in more than one Region, keep in mind how resources can interact across Regions. Some resources, such as Amazon GameLift Servers fleets, can only reference other resources in the same Region. Other resources, such as an Amazon GameLift Servers queue,

are Region agnostic. To manage Amazon GameLift Servers resources in multiple Regions, we recommend the following structure.

- **Regional support stacks** – These stacks contain resources that your Amazon GameLift Servers resources depend on. This stack must include the S3 bucket where you store your custom game server or Realtime script files. It might also contain other Amazon resources for your game, such as DynamoDB tables, Amazon Redshift clusters, and Lambda functions. Many of these resources are Region specific, so you must create them in every Region. Amazon GameLift Servers also needs an IAM role that allows access to these support resources. Because an IAM role is Region agnostic, you only need one role resource, placed in any Region and referenced in all of the other support stacks.
- **Regional Amazon GameLift Servers stacks** – This stack contains the Amazon GameLift Servers resources that must exist in each region where your game is being deployed, including the build or script, a set of fleets, and aliases. Amazon CloudFormation creates a build or script resource with files in an S3 bucket location, and deploys the build or script to one or more fleet resources. Each fleet should have a corresponding alias. The game session queue references some or all of the fleet aliases. You can maintain one template to describe this type of stack and use it to create identical sets of resources in every Region.
- **Global Amazon GameLift Servers stack** – This stack contains your game session queue and matchmaking resources. These resources can be located in any Region and are usually placed in the same Region. The queue can reference fleets or aliases that are located in any Region. To place additional queues in different Regions, create additional global stacks.

The diagrams below illustrates a multistack structure for deploying resources in several Amazon Regions. The first diagram shows a structure for a single game session queue. The second diagram shows a structure with multiple queues.





Updating builds

Amazon GameLift Servers builds are immutable, as is the relationship between a build and a fleet. As a result, when you update your hosting resources to use a new set of game build files, the following need to happen:

- Create a new build using the new set of files (replacement).
- Create a new set of fleets to deploy the new game build (replacement).
- Redirect aliases to point to the new fleets (update with no interruption).

For more information, see [Update behaviors of stack resources](#) in the *Amazon CloudFormation User Guide*.

Deploy build updates automatically

When updating a stack containing related build, fleet and alias resources, the default Amazon CloudFormation behavior is to automatically perform these steps in sequence. You trigger this update by first uploading the new build files to a new S3 location. Then you modify your Amazon CloudFormation build template to point to the new S3 location. When you update your stack with the new S3 location, this triggers the following Amazon CloudFormation sequence:

1. Retrieves the new files from S3, validates the files, and creates a new Amazon GameLift Servers build.
2. Updates the build reference in the fleet template, which triggers new fleet creation.
3. After the new fleets are active, updates the fleet reference in the alias, which triggers the alias to update to target the new fleets.
4. Deletes the old fleet.
5. Deletes the old build.

If your game session queue uses fleet aliases, player traffic is automatically switched to the new fleets as soon as the aliases are updated. The old fleets are gradually drained of players as game sessions end. Auto scaling handles the task of adding and removing instances from each set of fleets as player traffic fluctuates. Alternatively, you can specify an initial desired instance count to quickly ramp up for the switch and enable auto scaling later.

You can also have Amazon CloudFormation retain resources instead of deleting them. For more information, see [RetainResources](#) in the *Amazon CloudFormation API Reference*.


Deploy build updates manually

If you want to have more control over when new fleets go live for players, you have some options. You can choose to manage aliases manually using the Amazon GameLift Servers console or the CLI. Alternatively, instead of updating your build template to replace the build and fleets, you can add a second set of build and fleet definitions to your template. When you update the template, Amazon CloudFormation creates a second build resource and corresponding fleets. Since the existing resources are not replaced, they are not deleted, and the aliases remain pointing at original fleets.

The main advantage with this approach is that it gives you the flexibility. You can create separate resources for the new version of your build, test the new resources, and control when the new

fleets go live to players. A potential drawback is that it requires twice as many resources in each Region for a brief period of time.

The following diagram illustrates this process.



How rollbacks work

When executing a resource update, if any step is not completed successfully, Amazon CloudFormation automatically initiates a rollback. This process reverses each step in sequence, deleting the newly created resources.

If you need to manually trigger a rollback, change the build template's S3 location key back to the original location and update your stack. A new Amazon GameLift Servers build and fleet are created, and the alias switches over to the new fleet after the fleet is active. If you are managing aliases separately, you need to switch them to point to the new fleets.

For more information about how to handle a rollback that fails or gets stuck, see [Continue rolling back an update](#) in the *Amazon CloudFormation User Guide*.

Testing and troubleshooting Amazon GameLift Servers

This section provides guidance for diagnosing and resolving issues with your Amazon GameLift hosting solution. Whether you're experiencing problems during development, deployment, or production operations, these resources will help you identify root causes and implement effective solutions.

- **Development and testing environments** - Set up iterative development workflows that mirror production conditions, enabling you to catch and resolve integration issues early in your development cycle.
- **Fleet deployment and management** - Diagnose common fleet creation failures, server process activation issues, scaling problems, and deployment errors that can prevent your game servers from running properly.
- **Runtime diagnostics and monitoring** - Access live fleet instances for real-time troubleshooting, performance analysis, and configuration validation to resolve issues affecting player experience.

Use these resources alongside Amazon GameLift's monitoring and logging capabilities to maintain optimal performance and availability for your multiplayer games.

Debug Amazon GameLift Servers fleet issues

This topic provides guidance on how to resolve issues with your Amazon GameLift Servers managed EC2 fleets.

Fleet creation issues

When you create a managed EC2 fleet, the Amazon GameLift Servers service initiates a workflow that creates the fleet, deploys EC2 instances with your game server build installed, and starts game server processes on each instance. For a detailed description, see . A fleet cannot host game sessions and players until it reaches **Active** status.

You can debug issues that prevent fleets from becoming active by identifying the fleet creation phase where the issue occurred and reviewing fleet creation events and logs. If the logs do not offer useful information, it's possible that the problem is due to an internal service error. In this situation, try to create the fleet again. If the problem persists, try re-uploading the game build to resolve possible file corruption). You can also contact Amazon GameLift Servers support or post a question on the forum.

Downloading and validating the build

During this phase, Amazon GameLift Servers gets your uploaded game server build, extracts the files, and runs any install scripts. If fleet creation fails during these phases, look at fleet events and logs to pinpoint the issue. Possible causes include:

- Amazon GameLift Servers can't get the compressed build file (event `FLEET_BINARY_DOWNLOAD_FAILED`). Verify that the build's storage location can be accessed, that you're creating a fleet in the same Amazon Web Services Region as the build, and that Amazon GameLift Servers has the correct permissions to access it.
- Amazon GameLift Servers can't extract the build files (event `FLEET_CREATION_EXTRACTING_BUILD`).
- An install script in the build files failed to complete successfully (event `FLEET_CREATION_FAILED_INSTALLER`).

Building fleet resources

Issues during this phase usually involve the allocation and deployment of fleet resources. Possible causes include:

- The requested instance type isn't available.
- The requested fleet type (Spot or On-Demand) isn't available.

Activating game server processes

During this phase, Amazon GameLift Servers is attempting a number of tasks and testing key elements, including the game server's viability, runtime configuration settings, and the game server's ability to connect with the Amazon GameLift Servers service using the Server SDK.

Note

In this phase, you can remotely access a fleet instance to further investigate issues. See [Remotely connect to Amazon GameLift Servers fleet instances](#).

Possible issues include:

- Server processes don't start running. This suggests an issue with the fleet's runtime configuration settings (events `FLEET_VALIDATION_LAUNCH_PATH_NOT_FOUND` or `FLEET_VALIDATION_EXECUTABLE_RUNTIME_FAILURE`. Verify that you've correctly set the launch path and optional launch parameters.

- Server processes start running, but the fleet fails to activate. If server processes start and run successfully, but the fleet does not move to **Active** status, a likely cause is that the server process is failing to communicate with the Amazon GameLift Servers service. Verify that your game server is making these correct server SDK calls (see [Initialize the server process](#)):
 - Server process fails to initialize (event `SERVER_PROCESS_SDK_INITIALIZATION_TIMEOUT`). The server process is not successfully calling `InitSdk()`.
 - Server process fails to notify Amazon GameLift Servers when it's ready to host a game session (event `SERVER_PROCESS_PROCESS_READY_TIMEOUT`). The server process initialized but didn't call `ProcessReady()` in time.
- A VPC peering connection request failed. For fleets that are created with a VPC peering connection (see [To set up VPC peering with a new fleet](#)), VPC peering is done during this **Activating** phases. If a VPC peering fails for any reason, the new fleet will fail to move to **Active** status. You can track the success or failure of the peering request by calling [describe-vpc-peering-connections](#). Be sure to check that a valid VPC peering authorization exists ([describe-vpc-peering-authorizations](#), since authorizations are only valid for 24 hours.

Server process issues

Server processes start but fail quickly or report poor health.

Other than issues with your game build, this outcome can happen when trying to run too many server processes simultaneously on the instance. The optimum number of concurrent processes depends on both the instance type and your game server's resource requirements. Try reducing the number of concurrent processes, which is set in the fleet's runtime configuration, to see if performance improves. You can change a fleet's runtime configuration using either the Amazon GameLift Servers console (edit the fleet's capacity allocation settings) or by calling the Amazon CLI command [update-runtime-configuration](#).

Fleet deletion issues

Fleet can't be terminated due to max instance count.

The error message indicates that the fleet being deleted still has active instances, which is not allowed. You must first scale a fleet down to zero active instances. This is done by manually

setting the fleet's desired instance count to "0" and then waiting for the scale-down to take effect. Be sure to turn off auto scaling, which will counteract manual settings.

VPC actions are not authorized.

This issue only applies to fleets that you have specifically created VPC peering connections for (see [Set up VPC peering for Amazon GameLift Servers](#)). This scenario occurs because the process of deleting a fleet also includes deleting the fleet's VPC and any VPC peering connections. You must first get an authorization by calling the service API for Amazon GameLift Servers [CreateVpcPeeringAuthorization\(\)](#) or use the Amazon CLI command `create-vpc-peering-authorization`. Once you have the authorization, you can delete the fleet.

Remotely connect to Amazon GameLift Servers fleet instances

You can connect to any instance in your active Amazon GameLift Servers managed EC2 or container fleets. Common reasons to remotely access an instance include:

- Troubleshoot issues with your game server integration.
- Fine-tune your runtime configuration and other fleet-specific settings.
- Get real-time game server activity, such as log tracking.
- Run benchmarking tools using actual player traffic.
- Investigate specific issues with a game session or server process.

When connecting to an instance, consider these potential issues:

- You can connect to any instance in an active fleet. Generally, you can't connect to non-active fleets, such as fleets that are in the process of activating or are in an error state. (These fleets might have limited availability for a short period of time.) For help with fleet activation issues, see [Debug Amazon GameLift Servers fleet issues](#).
- Connecting to an active instance doesn't affect the instance's hosting activity. The instance continues to start and stop server processes based on the runtime configuration. It activates and runs game sessions. The instance might shut down in response to a scale down event or other event.
- Any changes you make to files or settings on the instance might impact the instance's active game sessions and connected players.

Remote access through the console

You can connect to fleet instances directly from the Amazon GameLift Servers console using Amazon EC2 Systems Manager (SSM). This method provides secure access without requiring additional setup or credential management.

1. In the Amazon GameLift Servers console, choose **Managed EC2** or **Managed containers** from the navigation pane, and then **Fleets**.
2. Choose the fleet ID that contains the instance you want to access.
3. On the fleet details page, choose the **Instances** tab to view all compute instances for the fleet.
4. Select the instance you want to connect to, then choose **Connect**. This displays the Connect to instance dialog which informs you of the details of the connection, and allows you to view the script that will be used to connect to your instance. Confirm by choosing **Connect** again.
5. In the connection dialog, choose **Run** to create a new SSM session. The system authenticates your session through Amazon Key Management Service (Amazon KMS) and opens a terminal in your browser.

Note

Console-based remote access is available for fleets running server SDK version 5.x. For fleets running earlier SDK versions, use the Amazon CLI method described in the following section.

Remote access with the Amazon CLI

The following instructions describe how to remotely connect to an instance using the Amazon command line interface (CLI). You can also make programmatic calls using the Amazon SDK, as documented in the [service API reference for Amazon GameLift Servers](#).

Gather instance data

To connect to an Amazon GameLift Servers managed EC2 fleet instance, you need the following information:

- The ID of the instance you want to connect to. You can use either the instance ID or ARN.

- the server SDK for Amazon GameLift Servers version being used on the instance. The server SDK is integrated with the game build that is running on the instance.

The following instructions describe how complete these tasks using the Amazon CLI. You must know the fleet ID for the instance you want to connect to.

1. **Get the compute name.** Get a list of all active computes in the fleet. Call [list-compute](#) with a fleet ID or ARN. For a single-location fleet, specify the fleet identifier only. For a multi-location fleet, specify the fleet identifier and a location. With managed EC2 fleets, `list-compute` returns a list of fleet instances, and the property `ComputeName` is the instance ID. Find the compute you want to access.

Request

```
aws gamelift list-compute \  
  --fleet-id "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa" \  
  --location "sa-east-1"
```

Response

```
{  
  "ComputeList": [  
    {  
      "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",  
      "FleetArn": "arn:aws:gamelift:us-west-2::fleet/  
fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",  
      "ComputeName": "i-0abc12d3e45fa6b78",  
      "IpAddress": "00.00.000.00",  
      "DnsName":  
"b08444ki909kvqu6zpw3is24x5pyz4b6m05i3jbxvpk9craztu0lqrbbrbnbkks.uwp57060n1k6dn1nw49b78hg1  
west-2.amazongamelift.com",  
      "ComputeStatus": "Active",  
      "Location": "sa-east-1",  
      "CreationTime": "2023-07-09T22:51:45.931000-07:00",  
      "OperatingSystem": "AMAZON_LINUX_2023",  
      "Type": "c4.large"  
    }  
  ]  
}
```

2. **Find the server SDK version.** For this information you need to look up the build that is deployed to the fleet. Server SDK version is a build property.
 - a. Call [describe-fleet-attributes](#) with a fleet ID or ARN to get the fleet's build ID and ARN.
 - b. Call [describe-build](#) with the build ID or ARN to get the build's server SDK version.

For example:

Request

```
aws gamelift describe-fleet-attributes /  
  --fleet-ids "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
```

Response

```
{  
  "FleetAttributes": [  
    {  
      "FleetId": "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa",  
      "ComputeType": "EC2",  
      "BuildId": "build-3333cccc-44dd-55ee-66ff-00001111aa22",  
      . . .  
    }  
  ]  
}
```

Request

```
aws gamelift describe-build /  
  --build-id "build-3333cccc-44dd-55ee-66ff-00001111aa22"
```

Response

```
"Build": {  
  "BuildId": "build-1111aaaa-22bb-33cc-44dd-5555eeee66ff",  
  "Name": "My_Game_Server_Build_One",  
  "OperatingSystem": "AMAZON_LINUX_2023",  
  "ServerSdkVersion": "4.0.2",  
  . . .  
}
```

```
}
```

Connect to an instance (server SDK 5)

If the instance you want to connect to is running a game build with server SDK version 5.x, connect to the instance using Amazon EC2 Systems Manager (SSM). You can access remote instances that are running either Windows or Linux.

Before you start:

Complete the SSM setup steps and install the SSM plugin on your local machine. For more information, see [Setting up SSM](#) and [Install the Session Manager plugin for the Amazon CLI](#) in the *Amazon EC2 Systems Manager User Guide*.

1. **Request access credentials for the instance.** Call [get-compute-access](#) with the fleet ID and the compute name for the instance you want to connect to. Amazon GameLift Servers returns a set of temporary credentials for accessing the instance. For example:

Request

```
aws gamelift get-compute-access \  
--compute-name i-11111111a222b333c \  
--fleet-id fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa \  
--region us-west-2
```

Response

```
{  
  "ComputeName": " i-11111111a222b333c ",  
  "Credentials": {  
    "AccessKeyId": " ASIAIOSFODNN7EXAMPLE ",  
    "SecretAccessKey": " wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY ",  
    "SessionToken": " AQoDYXdzEJr...<remainder of session token>"  
  },  
  "FleetArn": " arn:aws:gamelift:us-west-2::fleet/  
fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa ",  
  "FleetId": " fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa "  
}
```

2. **Export the access credentials (optional).** You can export the credentials to environment variables and use them to configure the Amazon CLI for the default user. For more details, see [Environment variables to configure the Amazon CLI](#) in the Amazon Command Line Interface User Guide.

```
export AWS_ACCESS_KEY_ID=ASIAIOSFODNN7EXAMPLE
export AWS_SECRET_ACCESS_KEY=wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY
export AWS_SESSION_TOKEN=AQoDYXdzEJr...<remainder of session token>
```

3. **Connect to the fleet instance.** Start an SSM session with the instance you want to connect to. Include the Amazon Region or location of the instance. For more information, including how to set up SSM and the SSM plugin, see [Starting a session \(Amazon CLI\)](#) in the *Amazon EC2 Systems Manager User Guide*.

The start-session request will automatically use the credentials that you acquired in Step 1.

```
aws ssm start-session \
--target i-11111111a222b333c \
--region us-west-2 \
```

Note

If you get an access denied error, you might have an `Amazon_PROFILE` environment variable set to an Amazon profile, which causes Amazon CLI to use the wrong credentials for remote access. To resolve, temporarily unset your `Amazon_PROFILE` environment variable. Alternatively, you can create a custom Amazon profile for your remote access credentials and add the `--profile` command line parameter to your `start-session` request.

Connect to an instance (server SDK 4.x or earlier)

If the instance you want to connect to is running a game build with server SDK version 4 or earlier, use the following instructions. You can connect to instances that are running either Windows or Linux. Connect to a Windows instance using a remote desktop protocol (RDP) client. Connect to a Linux instance using an SSH client.

1. **Request access credentials for the instance.** When you have an instance ID, use the command [get-instance-access](#) to request access credentials. If successful, Amazon GameLift Servers returns the instance's operating system, IP address, and a set of credentials (user name and secret key). The credentials format depends on the instance operating system. Use the following instructions to retrieve credentials for either RDP or SSH.

- **For Windows instances** – To connect to a Windows instance, RDP requires a user name and password. The `get-instance-access` request returns these values as simple strings, so you can use the returned values as is. Example credentials:

```
"Credentials": {
  "Secret": "aA1bBB2cCCd3EEE",
  "UserName": "gl-user-remote"
}
```

- **For Linux instances** – To connect to a Linux instance, SSH requires a user name and private key. Amazon GameLift Servers issues RSA private keys and returns them as a single string, with the newline character (`\n`) indicating line breaks. To make the private key usable, take these steps: (1) convert the string to a `.pem` file, and (2) set permissions for the new file. Example credentials returned:

```
"Credentials": {
  "Secret": "-----BEGIN RSA PRIVATE KEY-----
nEXAMPLEKEYKCAQEAY7WZhaDsR1W3mR1QtvhwyORRX8gnxgDAfRt/gx42kWXsT4rXE/b5CpSgie/
\nvBoU7jLxx92pNHoFnByP+Dc21eyyz6CvjTmWA0JwfWiW5/akH7i05dSrvC7dQkW2duV5QuUdE0QW
\nZ/aNxMniGQE6XAgfwlnXVBwrerrQo+ZWQeqiUwwMkuEbLeJFLhMCvYURpUMSC1oehm449ilx9X1F
\nG50TCFe0zfl8dqqCP6GzbPaIjiU19xX/az0R9V+tpU0zEL+wmXnZt3/nHPQ5xvD20JH67km6SuPW
\noPzev/D8V+x4+bHthfSjR9Y7DvQFjfBVwHXigBdtZcU2/wei8D/HYwIDAQABaoIBAGZ1kaEvnrqu
\n/uler7vgIn5m71N5LKw4hJLAIW6tUT/fzvtcHK0SkbQCQXuriHmQ2MqyJX/0kn2NfjLV/
ufGxbL1\nmb5qwmGUnEpJaZD6QSSs3kICLwWUYUiGfc0uiSbmJoap/
GTLU0W5Mfcv36PaBUNy5p53V6G7hXb2\nbahyWyJNfjLe4M86yd2YK3V2CmK+X/
B0sShnJ36+hjrXPPWmV3N9zEmCdJjA+K15DYmhm/
tJWSD9\n81oGk9TopEp7CkIfatEATyyZiVqoRq6k64iuM9JkA30zdXzMqexXVJ1TLZVEH0E7bh1Y9d801ozR
\noQs/FiZNAx2iijCWyv01pjE73+kCgYEA9mZtyhkHkFDpwrSM1APaL8oNAbbjwEy7Z5Mqfq1
+1Ip1\nYkriL0DbLXlvRAH+yHPRit2hH0jtUNZh4Axv+cpg09qbUI3+43eEy24B7G/Uh
+GTfbjsXs0xQx/x\np9otyVwc7hsQ5TA5PZb
+mvkJ50BEKzet9XcKw0NBVELGhnEPE7cCgYEA06Vgov6YHleHui9kHuws
\nayav0elc5zkxjF9nfHFJRry21R1trw2Vdpn+9g481URrpzWV0Eihvm+xTtmaZ1Sp//1kq75XDwnU
\nwA8gkn603QE3fq2yN98BURsAKdJfJ5RL1HvGQvTe10HLYYXpJnEkHv+Unl2ajLivWUt5pbBrKbUC
\nngYBjb0+0Zk0sCcpZ29sbzjYjpIddErySiYRX5gV2uNQwAjLdp9PfN295yQ+BxMBXiIycWVQiw0bH
\noMo7yykABY70zd5wQewBQ4AdSlWSX4nGDtsiFxiI5sKuAAe0CbTosy1s8w8fxoJ5Tz1sdoxNeGs
-----"
```

```

\nArq6Wv/G16zQuAE9zK9vwwKBgF+09VI/1wJBirsDGz9whVwfFPtTkJNvJZzYt69qezx1sjgFKshy
\nWBhd4xHZtmCqpBP1AymEjr/T01bxyARmXMnIOWIANNXMGB4KGSyl1mzSVAoQ+fQr+cJ3d0dyP11j
\njjb0Ed/NY8fr1NDxAVHE8BSkdsx2f6ELEyBKJSRr9snRAoGAMrTwYneXzvTskF/S5Fyu0i0egLda
\nNWUH38v/nDCgEpIXD5Hn3qAEcju1IjmbwlvtW+nY2jVhv7UGd8MjwUTNGItbdb6nsYqM2asrnF3qS
\nVRkAKKKYeGjkuUfVTTrW0YFjXkfcR/V+QFL50ndHAKJXjW7a4ejJLncTzmZSpYzwApc=\n-----END
RSA PRIVATE KEY-----",
  "UserName": "gl-user-remote"
}

```

When using the Amazon CLI, you can automatically generate a `.pem` file by including the `--query` and `--output` parameters to your `get-instance-access` request.

To set permissions on the `.pem` file, run the following command:

```
$ chmod 400 MyPrivateKey.pem
```

2. **Open a port for the remote connection.** You can access instances in Amazon GameLift Servers fleets through any port authorized in the fleet configuration. You can view a fleet's port settings using the command [describe-fleet-port-settings](#).

As a best practice, we recommend opening ports for remote access only when you need them and closing them when you're finished. You can't update port settings after creating a fleet but before it's active. If you get stuck, re-create the fleet with the port settings open.

Use the command [update-fleet-port-settings](#) to add a port setting for the remote connection (such as 22 for SSH or 3389 for RDP). For the IP range value, specify the IP addresses for the devices you plan to use to connect (converted to CIDR format). Example:

```

$ Amazon gamelift update-fleet-port-settings
  --fleet-id "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
  --inbound-permission-authorizations
    "FromPort=22,ToPort=22,IpRange=54.186.139.221/32,Protocol=TCP"

```

The following example opens up port 3389 on a Windows fleet

```

$ Amazon gamelift update-fleet-port-settings
  --fleet-id "fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
  --inbound-permission-authorizations
    "FromPort=3389,ToPort=3389,IpRange=54.186.139.221/32,Protocol=TCP"

```


3. **Open a remote connection client.** Use Remote Desktop for Windows or SSH for Linux instances. Connect to the instance using the IP address, port setting, and access credentials.

SSH example:

```
ssh -i MyPrivateKey.pem gl-user-remote@192.0.2.0
```

View files on remote instances

Regardless of whether you have connected to an instance remotely through the console or the Amazon CLI, you have full user and administrative access to it. This means you also have the ability to cause errors or failures with game hosting. If the instance is hosting games with active players, you may run the risk of crashing game sessions and dropping players, or disrupting game shutdown processes which could cause errors in saved game data and logs.

Look for these resources on a hosting instance:

- **Game build files.** These files are the game build that you uploaded to Amazon GameLift Servers. They include one or more game server executables, assets, and dependencies. Game build files are in a root directory called `game`:
 - On Windows: `c:\game`
 - On Linux: `/local/game`
- **Game log files.** Find the log files that your game server generates in the game root directory at whatever directory path you designated.
- **Amazon GameLift Servers hosting resources.** The root directory `Whitewater` contains files used by the Amazon GameLift Servers service to manage game hosting activity. Don't modify these files for any reason.
- **Runtime configuration.** Don't access runtime configuration for individual instances. To make changes to a runtime configuration property, update the fleet's runtime configuration (see the Amazon SDK operation [UpdateRuntimeConfiguration](#) or the Amazon CLI [update-runtime-configuration](#)).
- **Fleet data.** A JSON file contains information about the fleet that the instance belongs to, for use by server processes running on the instance. The JSON file is in the following location:
 - On Windows: `C:\GameMetadata\gamelift-metadata.json`
 - On Linux: `/local/gamemetadata/gamelift-metadata.json`

Scaling game hosting capacity with Amazon GameLift Servers

Hosting capacity, measured in instances, represents the number of game sessions that Amazon GameLift Servers can host concurrently and the number of concurrent players that those game sessions can accommodate. One of the most challenging tasks with game hosting is scaling capacity to meet player demand without wasting money on resources that you don't need.

Capacity is adjusted at the fleet location level. All fleets have at least one location: the fleet's home Amazon Region. When viewing or scaling capacity, the information is listed by location, including the fleet's home Region and any additional remote locations.

You can manually set the number of instances to maintain, or you can set up auto scaling to dynamically adjust capacity as player demand changes. We recommend that you start by turning on the target-based auto scaling option. The goal of target-based auto scaling is to maintain enough hosting resources to accommodate current players plus a little extra to handle unexpected spikes in player demand. For most games, target-based auto scaling offers a highly effective scaling solution.

You can do most fleet scaling activities using the Amazon GameLift Servers console. You can also use an Amazon SDK or the Amazon Command Line Interface (Amazon CLI) with the [service API for Amazon GameLift Servers](#).

Topics

- [To manage fleet capacity in the console](#)
- [Set Amazon GameLift Servers capacity limits](#)
- [Manually set capacity for an Amazon GameLift Servers fleet](#)
- [Auto-scale fleet capacity with Amazon GameLift Servers](#)
- [Manage Scaling an Amazon GameLift Servers Fleet To/From Zero](#)
- [Scale Amazon GameLift Servers container fleets](#)

To manage fleet capacity in the console

1. Open the [Amazon GameLift Servers console](#).

2. In the navigation pane, choose **Hosting, Fleets**.
3. On the **Fleets** page, choose the name of an active fleet to open the fleet's detail page.
4. Choose the **Scaling** tab. On this tab, you can:
 - View historical scaling metrics for the entire fleet.
 - View and update capacity settings for each fleet location, including scaling limits and current capacity settings.
 - Update target-based auto scaling, view rule-based auto scaling policies applied to the entire fleet, and suspend auto scaling activity for each location.

Set Amazon GameLift Servers capacity limits

When scaling hosting capacity for an Amazon GameLift Servers fleet location, either manually or by auto scaling, consider the location's scaling limits. All fleet locations have a minimum and maximum limit that define the allowed range for the location's capacity. By default, limits on fleet locations have a minimum of 0 instances and a maximum of 1 instance. Before you can scale a fleet location, adjust the limits.

If you're using auto scaling, the maximum limit allows Amazon GameLift Servers to scale up a fleet location to meet player demand but prevents runaway hosting costs, such as during a DDOS attack. Set up an [Amazon CloudWatch alarm](#) to notify you when capacity approaches the maximum limit, so you can evaluate the situation and manually adjust as needed. (You can also [create a billing alarm](#) to monitor Amazon costs.) The minimum limit is useful to maintain hosting availability, even when player demand is low.

You can set capacity limits for a fleet's locations in the [Amazon GameLift Servers console](#) or by using the Amazon Command Line Interface (Amazon CLI).

To set capacity limits

Console

1. Open the [Amazon GameLift Servers console](#).
2. In the navigation pane, choose **Hosting, Fleets**.
3. On the **Fleets** page, choose the name of an active fleet to open the fleet's detail page.
4. On the **Scaling** tab, under **Scaling capacity**, select a fleet location, and then choose **Edit**.

5. In the **Edit scaling capacity** dialog box, set instance counts for **Min size**, **Desired instances**, and **Max size**.
6. Choose **Confirm**.

Amazon CLI

1. **Check current capacity settings.** In a command line window, use the [describe-fleet-location-capacity](#) command with the fleet ID and location that you want to change capacity for. This command returns a [FleetCapacity](#) object that includes the location's current capacity settings. Determine whether the new instance limits can accommodate the current desired instances setting.

```
aws gamelift describe-fleet-location-capacity \  
  --fleet-id <fleet identifier> \  
  --location <location name>
```

2. **Update limit settings.** In a command line window, use the [update-fleet-capacity](#) command with the following parameters. You can adjust both instance limits and desired instance count with the same command.

```
--fleet-id <fleet identifier>  
--location <location name>  
--max-size <maximum capacity for scaling>  
--min-size <minimum capacity for scaling>  
--desired-instances <fleet capacity goal>
```

Example:

```
aws gamelift update-fleet-capacity \  
  --fleet-id fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa \  
  --location us-west-2 \  
  --max-size 10 \  
  --min-size 1 \  
  --desired-instances 10
```

If your request is successful, Amazon GameLift Servers returns the fleet ID. If the new max-size or min-size value conflicts with the current desired-instances setting, Amazon GameLift Servers returns an error.

Manually set capacity for an Amazon GameLift Servers fleet

When you create a new fleet, Amazon GameLift Servers automatically sets the desired instances to one instance in each fleet location. Then, Amazon GameLift Servers deploys one new instance in each location. To change fleet capacity, you can add a target-based auto scaling policy, or you can manually set the number of instances that you want for a location.

Setting a fleet's capacity manually can be useful when you don't need auto scaling or when you need to hold capacity at a specified level. Manually setting capacity works only if you aren't using a target-based auto scaling policy. If you have a target-based auto scaling policy, it immediately resets the desired capacity based on its own scaling rules.

You can manually set capacity in the Amazon GameLift Servers console or by using the Amazon Command Line Interface (Amazon CLI). The fleet's status must be active.

Suspend auto scaling

You can suspend all auto scaling activity for each fleet location. With auto scaling suspended, the desired number of instances in the fleet location remains the same unless manually changed. When you suspend auto scaling for a location, it affects the fleet's current policies and any policies that you may define in the future.

To manually set fleet capacity

Console

1. Open the [Amazon GameLift Servers console](#).
2. In the navigation pane, choose **Hosting, Fleets**.
3. On the **Fleets** page, choose the name of an active fleet to open the fleet's detail page.
4. On the **Scaling** tab, under **Suspended auto scaling locations**, select each location that you want to suspend auto scaling for, and then choose **Suspend**.
5. Under **Scaling capacity**, select a location that you want to set manually, and then choose **Edit**.
6. In the **Edit scaling capacity** dialog box, set your preferred value for **Desired instances**, and then choose **Confirm**. This tells Amazon GameLift Servers the number of instances to maintain in an active state, ready to host game sessions.

Amazon GameLift Servers responds to the changes by deploying additional instances or shutting down unneeded ones. As Amazon GameLift Servers completes this process, the number of active instances in the location changes to match the updated desired instances value. This process may take a little time.

Amazon CLI

1. **Check current capacity settings.** In a command line window, use the [describe-fleet-location-capacity](#) command with the fleet ID and location that you want to change capacity for. This command returns a [FleetCapacity](#) object that includes the location's current capacity settings. Determine whether the instance limits can accommodate the new desired instances setting.

```
aws gamelift describe-fleet-location-capacity \  
  --fleet-id <fleet identifier> \  
  --location <location name>
```

2. **Update desired capacity.** Use the [update-fleet-capacity](#) command with the fleet ID, location, and a new value for desired instances. If this value falls outside the current limit range, you can adjust limit values in the same command.

```
--fleet-id <fleet identifier>  
--location <location name>  
--desired-instances <fleet capacity as an integer>  
--max-size <maximum capacity> [Optional]  
--min-size <minimum capacity> [Optional]
```

Example:

```
aws gamelift update-fleet-capacity \  
  --fleet-id fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa \  
  --location us-west-2 \  
  --desired-instances 5 \  
  --max-size 10 \  
  --min-size 1
```

If your request is successful, Amazon GameLift Servers returns the fleet ID. If the new desired instances setting is outside the minimum and maximum limits, Amazon GameLift Servers returns an error.

Auto-scale fleet capacity with Amazon GameLift Servers

Use auto scaling in Amazon GameLift Servers to dynamically scale your fleet capacity in response to game server activity. As players arrive and start game sessions, auto scaling can add more instances; as player demand wanes, auto scaling can terminate unneeded instances. Auto scaling is an effective way to minimize your hosting resources and costs, while still providing a smooth, fast player experience.

To use auto scaling, you create scaling policies that tell Amazon GameLift Servers when to scale up or down. There are two types of scaling policies: target-based and rule-based. The target-based approach—target tracking—is a complete solution. We recommend it as the simplest and most effective option. Rule-based scaling policies require you to define each aspect of the auto scaling decision-making process, which is useful for addressing specific issues. This solution works best as a supplement to target-based auto scaling.

You can manage target-based auto scaling using the Amazon GameLift Servers console, the Amazon Command Line Interface (Amazon CLI), or an Amazon SDK. You can manage rule-based auto scaling using the Amazon CLI or an Amazon SDK only, although you can view rule-based scaling policies in the console.

Topics

- [Target-based auto scaling](#)
- [Auto scale with rule-based policies](#)

Target-based auto scaling

Target-based auto scaling for Amazon GameLift Servers adjusts capacity levels based on the fleet metric `PercentAvailableGameSessions`. This metric represents the fleet's available buffer for sudden increases in player demand.

The primary reason for maintaining a capacity buffer is player wait time. When game session slots are ready and waiting, it takes seconds to get new players into game sessions. If no resources are available, players must wait for existing game sessions to end or for new resources to become available. It can take minutes to start up new instances and server processes.

When setting up target-based auto scaling, specify the size of the buffer that you want the fleet to maintain. Because `PercentAvailableGameSessions` measures the percentage of available

resources, the actual buffer size is a percentage of the total fleet capacity. Amazon GameLift Servers adds or removes instances to maintain the target buffer size. With a large buffer, you minimize wait time, but you also pay for extra resources that you may not use. If your players are more tolerant of wait times, you can lower costs by setting a small buffer.

To set target-based auto scaling

Console

1. Open the [Amazon GameLift Servers console](#).
2. In the navigation pane, choose **Hosting, Fleets**.
3. On the **Fleets** page, choose the name of an active fleet to open the fleet's detail page.
4. Choose the **Scaling** tab. This tab displays the fleet's historical scaling metrics and contains controls for adjusting current scaling settings.
5. Under **Scaling capacity**, check that the **Min size** and **Max size** limits are appropriate for the fleet. With auto scaling enabled, capacity adjusts between these two limits.
6. In **Target-based auto scaling policy**, choose **Edit**.
7. In the **Edit target-based auto scaling policy** dialog box, for **Percent available game sessions**, set the percentage that you want to maintain, and then choose **Confirm**. After you've confirmed the settings, Amazon GameLift Servers adds a new target-based policy under **Target-based auto scaling policy**.

Amazon CLI

1. **Set capacity limits.** Set the limit values using the [update-fleet-capacity](#) command. For more information, see [Set Amazon GameLift Servers capacity limits](#).
2. **Create a new policy.** Open a command-line window and use the [put-scaling-policy](#) command with your policy's parameter settings. To update an existing policy, specify the policy's name and provide a complete version of the updated policy.

```
--fleet-id <unique fleet identifier>
--name "<unique policy name>"
--policy-type <target- or rule-based policy>
--metric-name <name of metric>
--target-configuration <buffer size>
```

Example:


```
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"
```

Auto scale with rule-based policies

Rule-based scaling policies in Amazon GameLift Servers provide fine-grained control when auto scaling a fleet's capacity in response to player activity. For each policy, you can link scaling to one of several fleet metrics, identify a trigger point, and customize the responding scale-up or scale-down event. Rule-based policies are useful for supplementing [target-based scaling](#) to handle special circumstances.

A rule-based policy states the following: "If a fleet metric meets or crosses a threshold value for a certain length of time, then change the fleet's capacity by a specified amount." This topic describes the syntax used to construct a policy statement and provides help with creating and managing your rule-based policies.

Manage rule-based policies

Create, update, or delete rule-based policies using an Amazon SDK or the Amazon Command Line Interface (Amazon CLI) with the [Service API for Amazon GameLift Servers](#). You can view all active policies in the Amazon GameLift Servers console.

To temporarily stop all scaling policies for a fleet, use the Amazon CLI command [stop-fleet-actions](#).

To create or update a rule-based scaling policy (Amazon CLI):

1. **Set capacity limits.** Set either or both limit values using the [update-fleet-capacity](#) command. For more information, see [Set Amazon GameLift Servers capacity limits](#).
2. **Create a new policy.** Open a command line window and use the [put-scaling-policy](#) command with your policy's parameter settings. To update an existing policy, specify the policy's name and provide a complete version of the updated policy.

```
--fleet-id <unique fleet identifier>  
--name "<unique policy name>"
```

```
--policy-type <target- or rule-based policy>
--metric-name <name of metric>
--comparison-operator <comparison operator>
--threshold <threshold integer value>
--evaluation-periods <number of minutes>
--scaling-adjustment-type <adjustment type>
--scaling-adjustment <adjustment amount>
```

Example:

```
aws gamelift put-scaling-policy \
  --fleet-id fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa \
  --name "Scale up when AGS<50" \
  --policy-type RuleBased \
  --metric-name AvailableGameSessions \
  --comparison-operator LessThanThreshold \
  --threshold 50 \
  --evaluation-periods 10 \
  --scaling-adjustment-type ChangeInCapacity \
  --scaling-adjustment 1
```

To delete a rule-based scaling policy using the Amazon CLI:

- Open a command line window and use the [delete-scaling-policy](#) command with the fleet ID and policy name.

Example:

```
aws gamelift delete-scaling-policy \
  --fleet-id fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa \
  --name "Scale up when AGS<50"
```

Syntax for auto scaling rules

To construct a rule-based scaling policy statement, specify six variables:

If *<metric name>* remains *<comparison operator>* *<threshold value>* for *<evaluation period>*, then change fleet capacity using *<adjustment type>* to/by *<adjustment value>*.

For example, this policy statement starts a scale-up event whenever a fleet's extra capacity is less than what's needed to handle 50 new game sessions:

If `AvailableGameSessions` remains at less than 50 for 10 minutes, then change fleet capacity using `ChangeInCapacity` by 1 instances.

Metric name

To start a scaling event, link an auto scaling policy to one of the following fleet-specific metrics. For complete metric descriptions, see [Amazon GameLift Servers metrics for fleets](#).

- Activating game sessions
- Active game sessions
- Available game sessions
- Percent available game sessions
- Active instances
- Available player sessions
- Current player sessions
- Idle instances
- Percent idle instances

If the fleet is in a game session queue, you can use the following metrics:

- Queue depth – The number of pending game session requests this fleet is the best available hosting location for.
- Wait time – Fleet-specific wait time. The length of time that the oldest pending game session request has been waiting to be fulfilled. A fleet's wait time is equal to the oldest current request's time in queue.

Comparison operator

Tells Amazon GameLift Servers how to compare the metric data to the threshold value. Valid comparison operators include greater than (`>`), less than (`<`), greater than or equal (`>=`), and less than or equal (`<=`).

Threshold value

When the specified metric value meets or crosses the threshold value, it starts a scaling event. This value is always a positive integer.

Evaluation period

The metric must meet or cross the threshold value for the full length of the evaluation period before starting a scaling event. The evaluation period length is consecutive; if the metric retreats from the threshold, the evaluation period starts over again.

Adjustment type and value

This set of variables works together to specify how Amazon GameLift Servers should adjust the fleet's capacity when a scaling event starts. Choose from three possible adjustment types:

- **Change in capacity** – Increase or decrease the current capacity by a specified number of instances. Set the adjustment value to the number of instances to add or remove from the fleet. Positive values add instances, while negative values remove instances. For example, a value of "-10" scales down the fleet by 10 instances, regardless of the fleet's total size.
- **Percent change in capacity** – Increase or decrease the current capacity by a specified percentage. Set the adjustment value to the percentage that you want to increase or decrease the fleet capacity by. Positive values add instances, while negative values remove instances. For example, for a fleet with 50 instances, a percentage change of "20" adds 10 instances to the fleet.
- **Exact capacity** – Increase or decrease the current capacity to a specific value. Set the adjustment value to the exact number of instances that you want to maintain in the fleet.

Tips for rule-based auto scaling

The following suggestions can help you get the most out of auto scaling with rule-based policies.

Use multiple policies

You can have multiple auto scaling policies for a fleet at the same time. The most common scenario is to have a target-based policy manage most scaling needs and use rule-based policies to handle edge cases. There are no limits on using multiple policies.

With multiple policies, each policy behaves independently. There is no way to control the sequence of scaling events. For example, if you have multiple policies driving scaling up, it's possible that player activity could start multiple scaling events simultaneously. Avoid policies that start each other. For example, you could create an infinite loop if you create scale-up and scale-down policies that set capacity beyond the threshold of each other.

Set maximum and minimum capacity

Each fleet has a maximum and minimum capacity limit. This feature is important when using auto scaling. Auto scaling never sets capacity to a value outside of this range. By default, newly created fleets have a minimum of 0 and a maximum of 1. For your auto scaling policy to affect capacity as intended, increase the maximum value.

Fleet capacity is also constrained by limits on the fleet's instance type and by service quotas in your Amazon Web Services account. You can't set a minimum and maximum outside these limits and account quotas.

Track metrics after a change in capacity

After changing capacity in response to an auto scaling policy, Amazon GameLift Servers waits 10 minutes before responding to triggers from the same policy. This wait gives Amazon GameLift Servers time to add the new instances, launch the game servers, connect players, and start collecting data from the new instances. During this time, Amazon GameLift Servers evaluates the policy against the metric and tracks the policy's evaluation period, which restarts after a scaling event occurs. This means that a scaling policy could start another scaling event immediately after the wait time is over.

There is no wait time between scaling events that different auto scaling policies start.

Manage Scaling an Amazon GameLift Servers Fleet To/From Zero

Amazon GameLift Servers supports automatic scaling to and from zero instances based on game session activity. This managed capacity option allows your Fleet locations to scale-in to zero instances following a defined period of no game session activity and automatically scale-out when game sessions are requested.

Scaling to and from zero instances provides several advantages:

- **Cost optimization** – Eliminate compute costs during periods of inactivity by running zero instances when there is no game session activity.
- **Automatic reactive scale-out** – Fleets locations automatically scale-out to one instance when a game session is requested, eliminating the need for manual intervention.
- **Simplified management** – No need to manually adjust Fleet capacity to/from zero based on anticipated player demand or development needs.

When you enable Scale To/From Zero on a fleet, Amazon GameLift Servers monitors game session activity and automatically adjusts Fleet capacity:

- **Scale-in to zero** – After a configured period with no game sessions activity, Amazon GameLift Servers scales-in the Fleet location to zero instances.
- **Scale-out from zero** – When a game session creation request is received, Amazon GameLift Servers scales-out the Fleet location to one instance, allowing auto scaling to resume.
- **Continued scaling** – Following scale-out the fleet resumes using configured auto scaling policies to manage capacity.

Scale-in behavior

Amazon GameLift Servers begins the scale-in process for a Fleet location after the configured inactivity period has elapsed with no game session activity. This is defined as a period where::

- There are no active game sessions in the Fleet location.
- No requests have been made to create new game sessions in the Fleet location.

During scale-in, Amazon GameLift Servers will set minimum and desired capacity for the Fleet location to zero, rapidly scaling-in for cost savings.

Scale-out behavior

When a game session creation request is received while the Fleet location is at zero instances:

- Amazon GameLift Servers immediately initiates scale-out of one instance.
- Attempted placement of the game session may continue with other Fleets or Fleet locations, depending on configuration of Queues (if used).

Note

Scale-out from zero takes time to provision and initialize instances. Players may experience longer wait times for the first game session after a period of inactivity. For this reason this feature is best paired with multi-location Fleets and/or Queues.

Configuring Scale To/From Zero

Scaling To/From Zero is configured by updating an existing fleet.

Console

1. Open the [Amazon GameLift Servers console](#).
2. In the navigation pane, choose **Hosting, Fleets**.
3. On the **Fleets** page, choose the name of an active fleet to open the fleet's detail page.
4. Under **Scaling**, select each location for which you want to configure scale to/from zero, and then choose **Edit**.
5. In the **Edit scaling capacity** dialog box, select "Automatic" for **Minimum capacity strategy**, set your preferred value in minutes for **Set minimum capacity to 0 after** and then choose **Confirm**.

Amazon GameLift Servers will then scale-in the selected fleet locations to 0 instances once there has been no game session activity for the configured duration. Following this, when a request for a game session is made in this location, Amazon GameLift Servers will scale-out one instance as quickly as possible. This process will take a little time.

Amazon CLI

- **Configure scale to/from zero.** In a command line window, use the [update-fleet-capacity](#) command with the fleet ID, location and managed capacity configuration to configure scale to/from zero.

```
aws gamelift update-fleet-capacity \
    --fleet-id <fleet identifier> \
    --location <location name> \
    --managed-capacity-configuration
    ScaleInAfterInactivityMinutes=60,ZeroCapacityStrategy=SCALE_TO_AND_FROM_ZERO
```

Example:

```
aws gamelift update-fleet-capacity \
    --fleet-id
    fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa \
    --location us-west-2 \
```

```
--desired-instances 5 \  
--max-size 10 \  
--managed-capacity-configuration  
ScaleInAfterInactivityMinutes=60,ZeroCapacityStrategy=SCALE_TO_AND_FROM_ZERO
```

If your request is successful, Amazon GameLift Servers returns the updated capacity configuration. Amazon GameLift Servers will then scale-in the selected fleet locations to 0 instances once there has been no game session activity for the configured duration. Following this, when a request for a game session is made in this location, Amazon GameLift Servers will scale-out one instance as quickly as possible. This process will take a little time.

Best practices

Consider the following recommendations when using Scale To/From Zero:

- **Set appropriate inactivity periods** – Balance cost savings against the frequency of scale-in/scale-out cycles. Shorter wait time to scale to zero may maximize savings, but would result in more frequent cold starts.
- **Use with predictable workloads** – Scale To/From Zero works best for games with clear periods of inactivity, such as development/test environments or games with distinct off-peak hours.
- **Monitor CloudWatch metrics** – Track fleet scaling events and game session placement times to optimize your configuration.
- **Combine with scaling policies** – Use Scale To/From Zero alongside target-based or rule-based auto scaling for comprehensive capacity management.

Scale Amazon GameLift Servers container fleets

One of the most challenging tasks with game hosting is scaling capacity to meet player demand without wasting money on resources that you don't need. In a managed container fleet, you scale your fleet capacity by adding or removing fleet instances.

When you create a new fleet, Amazon GameLift Servers sets the fleet's desired capacity to one instance and deploys one instance in the fleet's home region. For a multi-location fleet, Amazon GameLift Servers deploys one instance to the home region and to each remote location. After

the fleet status reaches ACTIVE, you can raise the desired capacity to raise or lower the desired capacity to scale down.

You can use Amazon GameLift Servers scaling features to change capacity manually or set up automatic scaling based on player demand:

- Set up automatic scaling with target tracking. See [Target-based auto scaling](#).
- Manually change the capacity of your fleet. See [Manually set capacity for an Amazon GameLift Servers fleet](#).

When scaling a container fleet, consider how adding or removing instances impacts the fleet's capacity to host game sessions and players.

- Game sessions per instance
 - Each game server process running on an instance represents the capacity to host one game session.
 - Use this formula to calculate the number of game sessions that run concurrently on a container fleet instance:

```
[Game sessions per instance] = [# of game server processes per game server container] * [# of game server container groups per instance]
```

If your container architecture runs one game server process concurrently in the game server container, then game sessions per instance equal the number of game server container groups per instance.

- For game server container groups per instance, call [DescribeContainerFleet](#) to get the `GameServerContainerGroupsPerInstance` or `MaximumGameServerContainerGroupsPerInstance` value.
- Players per instance
 - You decide the number of player slots to allow in each game session. Depending on how your hosting solution handles game session placement, you might define players per game session in your matchmaking configuration or in your calls to start a game session placement.
 - Use this formula to calculate the number of players that can play your game concurrently on a container fleet instance:

$$[\text{Players per instance}] = [\text{\# of game sessions per instance}] * [\text{\# of player slots per game session}]$$

To get the current total capacity of a container fleet, call [DescribeFleetCapacity](#) or [DescribeFleetLocation Capacity](#) to get the number of game server container groups in the fleet. Active groups are those that are currently hosting game sessions. Idle groups are ready to host a new game session. Multiply these values by the number of server processes per game server container group.

Monitoring Amazon GameLift Servers

Monitoring is an important part of maintaining the reliability, availability, and performance of Amazon GameLift Servers and your other Amazon solutions. There are three primary uses for metrics with Amazon GameLift Servers: to monitor system health and set up alarms, to track game server performance and usage, and to manage capacity using manual or auto-scaling.

Amazon provides the following monitoring tools to watch Amazon GameLift Servers, report when something is wrong, and take automatic actions when appropriate:

- **Amazon GameLift Servers console** – Use the graphical interface to manage your Amazon GameLift Servers resources and track game hosting activity.
- **Server telemetry metrics** – Amazon GameLift Servers enables you to emit custom metrics directly from your game servers through SDK and plugin integration. You can define and track your own game-specific metrics alongside built-in performance, network, memory, and timing data. All metrics can be published to [Amazon Managed Service for Prometheus](#) and monitored in Amazon Grafana using fully customizable dashboards. You can also publish to Amazon CloudWatch for integration with other Amazon services. The telemetry system is fully customizable - you can create custom Prometheus queries to derive additional metrics in addition to the built in ones. For implementation guides specific to your technology stack, see [the section called “Monitor with server telemetry metrics”](#).
- **Amazon CloudWatch** – You can monitor Amazon GameLift Servers metrics in real time, as well as metrics for other Amazon resources and applications that you're running on Amazon services. Amazon CloudWatch offers a suite of monitoring features, including tools to create customized dashboards and the ability to set alarms that notify or take action when a metric reaches a specified threshold.
- **Amazon CloudTrail** – captures all API calls and related events made by or on behalf of your Amazon account for Amazon GameLift Servers and other Amazon services. Data is delivered as log files to an Amazon S3 bucket that you specify. You can identify which users and accounts called Amazon, the source IP address from which the calls were made, and when the calls occurred.
- **Game session logs** – You can output custom server messages for your game sessions to log files that are stored in Amazon S3.

Metrics comparison across monitoring sources

Amazon GameLift Servers provides metrics through three primary sources: the Amazon GameLift Servers console Fleet Activity metrics, server telemetry metrics, and Amazon CloudWatch Amazon GameLift Servers metrics. Understanding the overlap and unique capabilities of each source helps you choose the right monitoring approach for your needs.

Metrics availability by source

The following tables show which metrics are available across different monitoring sources, organized by metric category.

Instance metrics

Instance-level metrics for fleet capacity and health monitoring:

Instance metrics availability

Metric	Console	CloudWatch	Telemetry
Active instances	✓	✓	✓
Idle instances	✓	✓	✓*
Percent idle instances	✓	✓	✓*
Desired instances	✓	✓	
Max instances	✓	✓	
Min instances	✓	✓	
CPU utilization		✓	✓
Network in/out		✓	✓
Disk/Storage read/write		✓	✓
Instance interruptions	✓	✓	
Recycled instances	✓	✓	

Metric	Console	CloudWatch	Telemetry
Unhealthy instances replaced		✓	

* Available through custom Prometheus queries using telemetry data.

Game server metrics

Server process and game session metrics:

Game server metrics availability

Metric	Console	CloudWatch	Telemetry
Active game sessions	✓	✓	✓*
Activating game sessions	✓	✓	
Available game sessions	✓	✓	
Percent available game sessions	✓	✓	
Concurrent activatable game sessions	✓	✓	
Game session interruptions	✓	✓	
Active server processes			✓
Healthy game servers			✓
Crashed game sessions			✓

* Available through custom Prometheus queries using telemetry data.

Player usage metrics

Player session and concurrent user metrics:

Player usage metrics availability

Metric	Console	CloudWatch	Telemetry
Current player sessions	✓		
Player session activations	✓		
Global concurrent users (CCU)			✓
Location concurrent users (CCU)			✓
Location capacity			✓

Container fleet metrics

Container-specific metrics (available only in CloudWatch for container fleets):

Container fleet metrics availability

Metric	Console	CloudWatch	Telemetry
Active container groups		✓	✓
Idle container groups		✓	✓
Container CPU/Memory utilization		✓	✓
Container network traffic		✓	✓
Container storage operations		✓	✓

Choosing the right monitoring source

Select your monitoring approach based on your specific needs:

- **Use server telemetry metrics** to monitor comprehensive performance data from your game servers and emit custom metrics specific to your game. These metrics deliver detailed insights into server performance, network activity, memory usage, and timing data for all game engines and server SDKs. You can define custom metrics for gameplay events, business logic

performance, and application-specific data points. All dashboards are fully customizable, and you can create custom Prometheus queries to derive additional metrics from the collected data. For more information, see [the section called “Monitor with server telemetry metrics”](#).

- **Use the Amazon GameLift Servers console** for fleet management, capacity planning, and general operational oversight. The console provides an integrated view of fleet health and player activity. For more information, see [the section called “Manage game hosting resources”](#).
- **Use Amazon CloudWatch** for automated monitoring, alerting, and integration with other Amazon services. CloudWatch enables custom dashboards and alarm-based automation. For more information, see [the section called “Monitor with CloudWatch”](#).
- **Use [Amazon Managed Service for Prometheus](#)** for high-performance metrics collection and storage with PromQL querying capabilities. Prometheus provides scalable time-series data storage for server telemetry metrics.
- **Use [Amazon Managed Grafana](#)** for advanced visualization and fully customizable dashboarding. Grafana provides pre-built GameLift dashboards that you can customize and extend, plus supports creating entirely custom dashboards with multiple data sources. You can build custom queries and visualizations to track any metrics important to your game.
- **Use multiple sources** for comprehensive monitoring. Combine console oversight, server telemetry metrics, and automated alerting for complete visibility into your game hosting infrastructure.

For additional monitoring capabilities, you can also use:

- [the section called “Logging API calls”](#) – Track API calls and related events for auditing and compliance.
- [the section called “Logging server messages”](#) – Capture custom server messages and game session logs.

Topics

- [the section called “Manage game hosting resources”](#)
- [the section called “Monitor with CloudWatch”](#)
- [the section called “Monitor with server telemetry metrics”](#)
- [the section called “Logging API calls”](#)
- [the section called “Logging server messages”](#)

Monitor Amazon GameLift Servers with Amazon CloudWatch

You can monitor Amazon GameLift Servers using Amazon CloudWatch, an Amazon service that collects raw data and processes it into readable, near real-time metrics. These statistics are kept for 15 months to provide a historical perspective on how your game server hosting with Amazon GameLift Servers is performing. You can set alarms that watch for certain thresholds and send notifications or take actions when those thresholds are met. For more information, see the [Amazon CloudWatch User Guide](#).

For enhanced server-side observability with detailed telemetry metrics, Amazon GameLift Servers also supports monitoring through [Amazon Managed Service for Prometheus](#) and [Amazon Managed Grafana](#). These telemetry metrics provide detailed server performance data including:

- **Server Timings:** Server Delta Time (consistency of server tick rate), Server Tick Rate (updates per second), Server Tick Time (processing time per tick), and Server World Tick Time (game world update time)
- **Network Metrics:** Connections (total network connections), Network I/O in Bytes and Packets (data transfer statistics), and Packet Loss (transmission failure percentage)
- **Memory Metrics:** Memory Usage in Units (total memory consumption) and Physical Memory Usage as percentage (memory utilization)
- **CPU Usage:** CPU Usage percentage (resource utilization by game server process)

For implementation guides specific to your SDK or plugin, see [the section called “Monitor with server telemetry metrics”](#).

The following tables list the metrics and dimensions for Amazon GameLift Servers. All metrics that are available in CloudWatch are also available in the Amazon GameLift Servers console, which provides the data as a set of customizable graphs. To access CloudWatch metrics for your games, use the Amazon Web Services Management Console, the Amazon CLI, or the CloudWatch API.

If a metric does not have a location, it uses the home location.

Dimensions for Amazon GameLift Servers metrics

Amazon GameLift Servers supports filtering metrics by the following dimensions.

Dimension	Description
Location	Filter metrics for a fleet deployment location. If a metric does not have a location, it uses the home location.
FleetId	Filter metrics for a single fleet. This dimension can be used with all fleet metrics for instances, server processes, game sessions, and player sessions.
MetricGroup	Filter metrics for a collection of fleets. Add a fleet to a metric group by adding the metric group name to the fleet's attributes (see UpdateFleetAttributes()). This dimension can be used with all fleet metrics for instances, server processes, game sessions, and player sessions.
QueueName	Filter metrics for a single queue. This dimension is used with metrics for game session queues only.
ConfigurationName	Filter metrics for a single matchmaking configuration. This dimension is used with metrics for matchmaking configurations.
ConfigurationName-RuleName	Filter metrics for an intersect of a matchmaking configuration and matchmaking rule. This dimension is used with metrics for matchmaking rules only.
InstanceType	Filter metrics for an EC2 instance type designation, such as "c4.large". This dimension is used with metrics for spot instances.
OperatingSystem	Filter metrics for an instance's operating system. This dimension is used with metrics for spot instances.
GameServerGroup	Filter FleetIQ metrics for a game server group.

Dimension	Description
ContainerGroupType	Filter container fleet metrics to distinguish between GameServer ContainerGroup metrics on the fleet location and PerInstance ContainerGroup metrics on the fleet location. This dimension is included for certain container fleet metrics (such as <code>ContainerNetworkIn</code> , <code>ContainerStorageReadBytes</code>) allow you to distinguish between metrics for different container group types.

Amazon GameLift Servers metrics for fleets

The `Amazon GameLift` namespace includes the following metrics related to activity across a fleet or a group of fleets. Fleets are used with a managed Amazon GameLift Servers solution. The Amazon GameLift Servers service sends metrics to CloudWatch every minute.

Instances

Metric	Description
ActiveInstances	Instances with ACTIVE status, which means they are running active server processes. The count includes idle instances and those that are hosting one or more game sessions. This metric measures current total instance capacity. This metric can be used with automatic scaling. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
DesiredInstances	Target number of active instances that Amazon GameLift Servers is working to maintain in the fleet. With automatic scaling, this value is determine

Metric	Description
	d based on the scaling policies currently in force. Without automatic scaling, this value is set manually. This metric is not available when viewing data for fleet metric groups. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
IdleInstances	Active instances that are currently hosting zero (0) game sessions. This metric measures capacity that is available but unused. This metric can be used with automatic scaling. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
MaxInstances	Maximum number of instances that are allowed for the fleet. A fleet's instance maximum determines the capacity ceiling during manual or automatic scaling up. This metric is not available when viewing data for fleet metric groups. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location

Metric	Description
MinInstances	Minimum number of instances allowed for the fleet. A fleet's instance minimum determines the capacity floor during manual or automatic scaling down. This metric is not available when viewing data for fleet metric groups. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
PercentIdleInstances	Percentage of all active instances that are idle (calculated as $\text{IdleInstances} / \text{ActiveInstances}$). This metric can be used for automatic scaling. Units: Percent Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
RecycledInstances	Number of spot instances that have been recycled and replaced. Amazon GameLift Servers recycles spot instances that are not currently hosting game sessions and have a high probability of interruption. This metric is available for managed EC2 fleets only. Units: Count Relevant CloudWatch statistics: Sum, Average, Minimum, Maximum Dimensions: Location

Metric	Description
RecycledSpotInstances	Number of spot instances that have been recycled and replaced. Amazon GameLift Servers recycles spot instances that are not currently hosting game sessions and have a high probability of interruption. This metric is available for container fleets only. Units: Count Relevant CloudWatch statistics: Sum, Average, Minimum, Maximum Dimensions: Location
InstanceSpotInterruptions	Number of spot instances that have been interrupted. Units: Count Relevant CloudWatch statistics: Sum, Average, Minimum, Maximum Dimensions: Location
CPUUtilization	<i>EC2 metric. For Amazon GameLift Servers this metric represents hardware performance across all active instances in a fleet location. The percentage of physical CPU time that Amazon EC2 uses to run the instance, which includes time spent to run both the user code and Amazon EC2 code. Tools in your operating system can show a different percentage than CloudWatch due to factors such as legacy device simulation, configuration of non-legacy devices, interrupt-heavy workloads, live migration, and live update.</i> Units: Percent

Metric	Description
NetworkIn	<i>EC2 metric. For Amazon GameLift Servers this metric represents hardware performance across all active instances in a fleet location.</i> The number of bytes received on all network interfaces by the instance. This metric identifies the volume of incoming network traffic to an application on a single instance. Units: Bytes
NetworkOut	<i>EC2 metric. For Amazon GameLift Servers this metric represents hardware performance across all active instances in a fleet location.</i> The number of bytes sent out on all network interfaces by the instance. This metric identifies the volume of outgoing network traffic to an application on a single instance. Units: Bytes
DiskReadBytes	<i>EC2 metric. For Amazon GameLift Servers this metric represents hardware performance across all active instances in a fleet location.</i> Bytes read from all instance store volumes available to the instance. This metric is used to determine the volume of the data the application reads from the hard disk of the instance. You can use it to determine the speed of the application. Units: Bytes

Metric	Description
DiskWriteBytes	<i>EC2 metric. For Amazon GameLift Servers this metric represents hardware performance across all active instances in a fleet location.</i> Bytes written to all instance store volumes available to the instance. This metric is used to determine the volume of the data the application writes onto the hard disk of the instance. You can use it to determine the speed of the application. Units: Bytes
DiskReadOps	<i>EC2 metric. For Amazon GameLift Servers this metric represents hardware performance across all active instances in a fleet location.</i> Completed read operations from all instance store volumes available to the instance in a specified period of time. To calculate the average I/O operations per second (IOPS) for the period, divide the total operations in the period by the number of seconds in that period. Units: Count
DiskWriteOps	<i>EC2 metric. For Amazon GameLift Servers this metric represents hardware performance across all active instances in a fleet location.</i> Completed write operations to all instance store volumes available to the instance in a specified period of time. To calculate the average I/O operations per second (IOPS) for the period, divide the total operations in the period by the number of seconds in that period. Units: Count

Server processes

Metric	Description
ActiveServerProcesses	Server processes with ACTIVE status, which means they are running and able to host game sessions. The count includes idle server processes and those that are hosting game sessions. This metric measures current total server process capacity. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
HealthyServerProcesses	Active server processes that are reporting healthy. This metric is useful for tracking the overall health of the fleet's game servers. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
PercentHealthyServerProcesses	Percentage of all active server processes that are reporting healthy (calculated as $\text{HealthyServerProcesses} / \text{ActiveServerProcesses}$). Units: Percent Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location

Metric	Description
ServerProcessAbnormalTerminations	Server processes that were shut down due to abnormal circumstances since the last report. This metric includes terminations that were initiated by the Amazon GameLift Servers service. This occurs when a server process stops responding, consistently reports failed health checks, or does not terminate cleanly (by calling ProcessEnding()). Units: Count Relevant CloudWatch statistics: Sum, Average, Minimum, Maximum Dimensions: Location
ServerProcessActivations	Server processes that successfully transitioned from ACTIVATING to ACTIVE status since the last report. Server processes cannot host game sessions until they are active. Units: Count Relevant CloudWatch statistics: Sum, Average, Minimum, Maximum Dimensions: Location

Metric	Description
ServerProcessTerminations	Server processes that were shut down since the last report. This includes all server processes that transitioned to TERMINATED status for any reason, including normal and abnormal process terminations. Units: Count Relevant CloudWatch statistics: Sum, Average, Minimum, Maximum Dimensions: Location

Game sessions

Metric	Description
ActivatingGameSessions	Game sessions with ACTIVATING status, which means they are in the process of starting up. Game sessions cannot host players until they are active. High numbers for a sustained period of time may indicate that game sessions are not transitioning from ACTIVATING to ACTIVE status. This metric can be used with automatic scaling. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
ActiveGameSessions	Game sessions with ACTIVE status, which means they are able to host players, and are hosting zero or more players. This metric measures the total

Metric	Description
	number of game sessions currently being hosted. This metric can be used with automatic scaling. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
AvailableGameSessions	Active, healthy server processes that are not currently being used to host a game session and can start a new game session without a delay to spin up new server processes or instances. This metric can be used with automatic scaling.  Note For fleets that limit concurrent game session activations, use the metric <code>ConcurrentActivatableGameSessions</code>. That metric more accurately represents the number of new game sessions that can start without any type of delay. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location

Metric	Description
ConcurrentActivatableGameSessions	Active, healthy server processes that are not currently being used to host a game session and can immediately start a new game session. This metric differs from <code>AvailableGameSessions</code> in the following way: it does not count server processes that currently cannot activate a new game session because of limits on game session activations. (See the fleet RuntimeConfiguration optional setting <code>MaxConcurrentGameSessionActivations</code>). For fleets that don't limit game session activations, this metric is identical to <code>AvailableGameSessions</code> . Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
PercentAvailableGameSessions	Percentage of game session slots on all active server processes (healthy or unhealthy) that are not currently being used (calculated as $\text{AvailableGameSessions} / [\text{ActiveGameSessions} + \text{AvailableGameSessions} + \text{unhealthy server processes}]$). This metric can be used with automatic scaling. Units: Percent Relevant CloudWatch statistics: Average Dimensions: Location

Metric	Description
GameSessionInterruptions	Number of game sessions on spot instances that have been interrupted. Units: Count Relevant CloudWatch statistics: Sum, Average, Minimum, Maximum Dimensions: Location

Player sessions

Metric	Description
CurrentPlayerSessions	Player sessions with either ACTIVE status (player is connected to an active game session) or RESERVED status (player has been given a slot in a game session but hasn't yet connected). This metric can be used with automatic scaling. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum
PlayerSessionActivations	Player sessions that transitioned from RESERVED status to ACTIVE since the last report. This occurs when a player successfully connects to an active game session. Units: Count Relevant CloudWatch statistics: Sum, Average, Minimum, Maximum

Container fleet metrics

Container fleet metrics provide monitoring data for container-based game hosting, including container group status, resource utilization, and performance metrics.

 **Note**

Enhanced dashboards for container fleet monitoring are available through [Amazon Managed Grafana](#) when using the telemetry metrics solution. See [the section called “Monitor with server telemetry metrics”](#) for setup instructions.

Metric	Description
ActiveGameServerContainerGroups	Container groups that are currently active and ready to host game sessions. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
IdleGameServerContainerGroups	Container groups that are active but not currently hosting game sessions. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
PendingGameServerContainerGroups	Container groups that are in the process of starting up. Units: Count

Metric	Description
	Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
TerminatingGameServerContainerGroups	Container groups that are in the process of shutting down. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
UnhealthyGameServerContainerGroupsReplaced	Container groups that were replaced due to health check failures. Units: Count Relevant CloudWatch statistics: Sum, Average, Minimum, Maximum Dimensions: Location
ContainerCPUUtilizationPerInstance	Average CPU utilization across all containers on an instance. Units: Percent Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location

Metric	Description
ContainerCPUReservation	Percentage of CPU resources reserved by containers. Units: Percent Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
ContainerMemoryReservation	Percentage of memory resources reserved by containers. Units: Percent Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
ContainerMemoryUtilization	Percentage of reserved memory being used by containers. Units: Percent Relevant CloudWatch statistics: Average, Minimum, Maximum Dimensions: Location
ContainerNetworkIn	Network bytes received by containers. Units: Bytes Relevant CloudWatch statistics: Sum, Average, Minimum, Maximum Dimensions: Location

Metric	Description
ContainerNetworkOut	Network bytes sent by containers. Units: Bytes Relevant CloudWatch statistics: Sum, Average, Minimum, Maximum Dimensions: Location
ContainerStorageReadBytes	Bytes read from storage by containers. Units: Bytes Relevant CloudWatch statistics: Sum, Average, Minimum, Maximum Dimensions: Location
ContainerStorageWriteBytes	Bytes written to storage by containers. Units: Bytes Relevant CloudWatch statistics: Sum, Average, Minimum, Maximum Dimensions: Location

Amazon GameLift Servers metrics for queues

The Amazon GameLift namespace includes the following metrics related to activity across a game session placement queue. Queues are used with a managed Amazon GameLift Servers solution. The Amazon GameLift Servers service sends metrics to CloudWatch every minute.

Metric	Description
AverageWaitTime	Average amount of time that game session placement requests in the queue with status PENDING have been waiting to be fulfilled.

Metric	Description
	Units: Seconds Relevant CloudWatch statistics: Average, Minimum, Maximum, Sum Dimensions: Location
FirstChoiceNotViable	Game sessions that were successfully placed but NOT in the first-choice fleet, because that fleet was considered not viable (such as a spot fleet with a high interruption rate). This metric is based on cost, not latency. The first-choice fleet is either the first fleet listed in the queue or—when a placement request includes player latency data—it is the first fleet chosen by FleetIQ prioritization. If there are no viable spot fleets, any fleet in that region may be selected. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum, Sum
FirstChoiceOutOfCapacity	Game sessions that were successfully placed but NOT in the first-choice fleet, because that fleet had no available resources. The first-choice fleet is either the first fleet listed in the queue or—when a placement request includes player latency data — it is the first fleet chosen by your defined FleetIQ prioritization. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum, Sum

Metric	Description
LowestLatencyPlacement	Game sessions that were successfully placed in a region that offers the queue's lowest possible latency for the players. This metric is emitted only when player latency data is included in the placement request. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum, Sum
LowestPricePlacement	Game sessions that were successfully placed in a fleet with the queue's lowest possible price for the chosen region. This fleet can be either a spot fleet or an on-demand instance if the queue has no spot instances. This metric is emitted only when player latency data is included in the placement request. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum, Sum
Placement <region name>	Game sessions that are successfully placed in fleets located in the specified region. This metric breaks down the PlacementsSucceeded metric by region. Units: Count Relevant CloudWatch statistics: Sum

Metric	Description
PlacementsCanceled	Game session placement requests that were canceled before timing out since the last report. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum, Sum
PlacementsFailed	Game session placement requests that failed for any reason since the last report. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum, Sum
PlacementsStarted	New game session placement requests that were added to the queue since the last report. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum, Sum
PlacementsSucceeded	Game session placement requests that resulted in a new game session since the last report. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum, Sum

Metric	Description
PlacementsTimedOut	Game session placement requests that reached the queue's timeout limit without being fulfilled since the last report. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum, Sum
QueueDepth	Number of game session placement requests in the queue with status PENDING. Units: Count Relevant CloudWatch statistics: Average, Minimum, Maximum, Sum Dimensions: Location

Amazon GameLift Servers metrics for matchmaking

The Amazon GameLift namespace includes metrics on FlexMatch activity for matchmaking configurations and matchmaking rules. FlexMatch matchmaking is used with a managed Amazon GameLift Servers solution. The Amazon GameLift Servers service sends metrics to CloudWatch every minute.

For more information on the sequence of matchmaking activity, see [How Amazon GameLift Servers FlexMatch works](#).

Matchmaking configurations

Metric	Description
CurrentTickets	Matchmaking requests currently being processed or waiting to be processed. Units: Count

Metric	Description
	Relevant CloudWatch statistics: Average, Minimum, Maximum, Sum
MatchAcceptancesTimedOut	For matchmaking configurations that require acceptance, the potential matches that timed out during acceptance since the last report. Units: Count Relevant CloudWatch statistics: Sum
MatchesAccepted	For matchmaking configurations that require acceptance, the potential matches that were accepted since the last report. Units: Count Relevant CloudWatch statistics: Sum
MatchesCreated	Potential matches that were created since the last report. Units: Count Relevant CloudWatch statistics: Sum
MatchesPlaced	Matches that were successfully placed into a game session since the last report. Units: Count Relevant CloudWatch statistics: Sum

Metric	Description
MatchesRejected	For matchmaking configurations that require acceptance, the potential matches that were rejected by at least one player since the last report. Units: Count Relevant CloudWatch statistics: Sum
PlayersStarted	Players in matchmaking tickets that were added since the last report. Units: Count Relevant CloudWatch statistics: Sum
TicketsFailed	Matchmaking requests that resulted in failure since the last report. Units: Count Relevant CloudWatch statistics: Sum
TicketsStarted	New matchmaking requests that were created since the last report. Units: Count Relevant CloudWatch statistics: Sum
TicketsTimedOut	Matchmaking requests that reached the timeout limit since the last report. Units: Count Relevant CloudWatch statistics: Sum

Metric	Description
TimeToMatch	For matchmaking requests that were put into a potential match before the last report, the amount of time between ticket creation and potential match creation. Units: Seconds Relevant CloudWatch statistics: Data Samples, Average, Minimum, Maximum
TimeToTicketCancel	For matchmaking requests that were canceled before the last report, the amount of time between ticket creation and cancellation. Units: Seconds Relevant CloudWatch statistics: Data Samples, Average, Minimum, Maximum
TimeToTicketSuccess	For matchmaking requests that succeeded before the last report, the amount of time between ticket creation and successful match placement. Units: Seconds Relevant CloudWatch statistics: Data Samples, Average, Minimum, Maximum

Matchmaking rules

Metric	Description
RuleEvaluationsPassed	Rule evaluations during the matchmaking process that passed since the last report. This metric is limited to the top 50 rules. Units: Count

Metric	Description
	Relevant CloudWatch statistics: Sum
RuleEvaluationsFailed	Rule evaluations during matchmaking that failed since the last report. This metric is limited to the top 50 rules. Units: Count Relevant CloudWatch statistics: Sum

Amazon GameLift Servers metrics for FleetIQ

The Amazon GameLift namespace includes metrics for FleetIQ game server group and game server activity as part of a FleetIQ standalone solution for game hosting. The Amazon GameLift Servers service sends metrics to CloudWatch every minute. Also see [Monitoring your Auto Scaling groups and instances using amazon CloudWatch](#) in the *Amazon EC2 Auto Scaling User Guide*.


Metric	Description
AvailableGameServers	Game servers that are available to run a game execution and are not currently occupied with gameplay. This number includes game servers that have been claimed but are still in AVAILABLE status. Units: Count Relevant Amazon CloudWatch statistics: Sum Dimensions: GameServerGroup
UtilizedGameServers	Game servers that are currently occupied with gameplay. This number includes game servers that are in UTILIZED status. Units: Count Relevant Amazon CloudWatch statistics: Sum

Metric	Description
	Dimensions: GameServerGroup
DrainingAvailableGameServers	Game servers on instances scheduled for termination that are currently not supporting gameplay. These game servers are the lowest priority to be claimed in response to a new claim request. Units: Count Relevant Amazon CloudWatch statistics: Sum Dimensions: GameServerGroup
DrainingUtilizedGameServers	Game servers on instances scheduled for termination that are currently supporting gameplay. Units: Count Relevant Amazon CloudWatch statistics: Sum Dimensions: GameServerGroup
PercentUtilizedGameServers	Portion of game servers that are currently supporting game executions. This metric indicates the amount of game server capacity that is currently in use. It is useful for driving an Auto Scaling policy that can dynamically add and remove instances to match with player demand. Units: Percent Relevant Amazon CloudWatch statistics: Average, Minimum, Maximum Dimensions: GameServerGroup

Metric	Description
GameServerInterruptions	Game servers on Spot Instances that were interrupted due to limited Spot availability. Units: Count Relevant Amazon CloudWatch statistics: Sum Dimensions: GameServerGroup, InstanceType
InstanceInterruptions	Spot Instances that were interrupted due to limited availability. Units: Count Relevant Amazon CloudWatch statistics: Sum Dimensions: GameServerGroup, InstanceType

Monitor with server telemetry metrics

Amazon GameLift Servers can be configured to collect and publish telemetry metrics for game servers running on managed Amazon EC2 and Container fleets. These metrics become available after deploying the telemetry collector with your server build. The metrics system supports all SDKs (C++, C#, Go), all plugins (Unreal, Unity), and the Amazon GameLift Servers Game Server Wrapper. Metrics data flows to [Amazon Managed Service for Prometheus](#), [the section called "Monitor with CloudWatch"](#), and [Amazon Managed Grafana](#) dashboards (recommended for visualization).



Monitor performance across your game servers using pre-built dashboards in Amazon Managed Grafana or Amazon CloudWatch.

Benefits of telemetry metrics

The telemetry metrics system offers five key benefits:

- **Game engine-specific metrics** — Game engine plugins (Unreal, Unity) provide native integration with engine-specific performance metrics such as server tick time, frame rate, and engine-level resource utilization that are critical for game performance optimization.
- **Custom metrics support** — Define and track your own game-specific metrics using server SDK function calls to monitor custom gameplay events, business logic performance, and application-specific data points that matter to your game.
- **Automated collection** — Metrics flow automatically after telemetry collector deployment with no additional instrumentation required and direct integration with [Amazon Managed Service for Prometheus](#) and Amazon CloudWatch.

- **Multi-level monitoring** — Fleet-level metrics for capacity and scaling, instance-level metrics for resource utilization, and game session metrics for performance tracking.
- **Universal compatibility** — Works with all Amazon GameLift Servers-supported development environments, integrated with all server SDKs, and native support in game engine plugins.

Note

Telemetry metrics are available for Amazon GameLift Servers managed Amazon EC2 or container fleets running Amazon Linux 2023 or Windows.

Before you begin

Required Amazon resources

- Amazon account configured for Amazon GameLift Servers.
- Managed fleet running on:
 - Amazon EC2 with supported operating systems, OR
 - Containers with Amazon Linux 2023
- Appropriate IAM permissions

IAM requirements

The following IAM permissions are required only if you plan to use the corresponding service:

- [Amazon Managed Service for Prometheus](#) (required only if publishing metrics to Prometheus)
 - `aps:RemoteWrite` permission
 - Access to your Prometheus workspace
- **Amazon CloudWatch** (required only if publishing metrics to Amazon CloudWatch)
 - `cloudwatch:PutMetricData` permission
 - Access to metrics namespaces
- [Amazon Managed Grafana](#) (required only if using Grafana dashboards)
 - `grafana:Read` permission
 - SSO configuration for dashboard access

Implementation

Select your implementation path based on your development environment:

SDK implementation

SDK Type	SDK Setup	Custom Metrics	API Reference
Go SDK	Complete Setup Guide	Go Metrics API	Go Actions and Data Types
C# SDK	Complete Setup Guide	C# Metrics API	C# Actions and Data Types
C++ SDK	Complete Setup Guide	C++ Metrics API	C++ Actions and Data Types

Plugin implementation

Plugin	Plugin Setup	Custom Metrics	API Reference
Unreal	Complete Setup Guide	Unreal Metrics API	Unreal Actions and Data Types
Unity	Complete Setup Guide	Unity Metrics API	C# Actions and Data Types

Implementation workflow

Each implementation follows a two-step process:

1. **Complete Setup Guide (METRICS.md)** — Infrastructure deployment, Amazon infrastructure configuration, fleet setup, and Grafana dashboard configuration.
2. **API Implementation Guide (CUSTOM_METRICS.md)** — Language-specific SDK usage, metric types, custom metrics creation, and advanced configuration.

Verification

1. Validate metrics flow by checking your [Amazon Managed Service for Prometheus](#) workspace or Amazon CloudWatch console for incoming telemetry data.
2. Check dashboard visibility in [Amazon Managed Grafana](#) using the pre-built dashboards.
3. Test custom metrics by verifying they appear in your monitoring dashboards.

Note

After completing implementation, return to this page and go to the [Available metrics](#) page.

Available metrics

Metrics fall into three categories:

- Automatically collected metrics
- SDK-provided metrics
- Custom metrics

Automatic metrics collection

No code changes required for these metrics:

Instance metrics

Metric Type	Description	Use Case
CPU	Percentage utilization per instance	Resource monitoring
Memory	Physical memory usage and percentage	Capacity planning
Network I/O	Bytes and packets sent/received	Connection health
Disk I/O	Read/write operations and throughput	Storage performance

Fleet metrics

Metric Type	Description	Use Case
Active Instances	Running instances count	Fleet scaling
Game Sessions	Active and available sessions	Capacity management
Crashed game sessions	Game sessions that have crashed	Error monitoring

SDK-provided metrics

Requires SDK function calls in your code:

Server timing metrics

Metric	Description	Implementation
Server Delta Time	Difference in time between the current server tick and the previous server tick. Measures the consistency of the server's tick rate	Call <code>GetDeltaTime()</code>
Server Tick Rate	Shows the number of times per second the server is processing updates	Automatically calculated
Server Tick Time	The amount of time it takes for the server to process a single tick or update	Call <code>GetTickTime()</code>
Server World Tick Time	The amount of time it takes for the server to update the game world with each tick	Call <code>GetWorldUpdateTime()</code>

Implementation: For engine-agnostic SDKs (C++, C#, Go), you implement these metrics by calling SDK functions from your game loop with calculated timing values. For engine plugins (Unreal, Unity), these metrics are captured automatically through engine integration.

Network metrics

Metric	Description	Implementation
Connections	The total number of network connections the server has established	Automatic after InitMetrics()
Network I/O (Bytes)	The total number of bytes being sent and received by the server over the network	Automatic after InitMetrics()
Network I/O (Packets)	The total number of network packets being sent and received by the server	Automatic after InitMetrics()
Packet Loss	The percentage of network packets that are being lost during transmission	Automatic after InitMetrics()

Implementation: Integrate SDK function calls with your networking library. The SDK provides guidance for different network implementations.

Process metrics

Metric	Description	Implementation
CPU Usage (%)	The percentage of CPU resources being utilized by the game server process	Automatic after InitMetrics()
Memory Usage (Units)	The total amount of memory being consumed by the server processes	Automatic after InitMetrics()
Physical Memory Usage (%)	The percentage of the server's total physical memory that is currently being utilized	Automatic after InitMetrics()
Server Status	Game server health state	Automatic after InitMetrics()

Implementation: These metrics are automatically collected by the SDK for each game session process.

Per-process dashboard organization

Per-process metrics are available in two specialized dashboards:

- **Server Performance dashboard** — Contains server timings (delta time, tick rate, tick time, world tick time), network metrics (connections, I/O bytes/packets, packet loss), memory usage, and CPU usage for individual game sessions.
- **Instance Performance dashboard** — Features "Top N Memory Consuming Game Sessions" and "Top N CPU Consuming Game Sessions" tables that help identify which processes contribute most to instance resource consumption. Clicking on Game Session links enables deeper investigation of detailed metrics.

Per-process metrics use cases

The per-process/per-game-session metrics support the following monitoring scenarios:

- **Dive deep performance investigation** — When a host/instance has degraded performance due to specific processes or game sessions, per-process metrics help identify which process caused the issue through Top CPU and Memory consuming Game Sessions tables.
- **Game server crash investigation** — When a game session crashes, these metrics help determine if the crash was due to out of memory, CPU overload, or network bandwidth problems.
- **Investigate player reported issues** — When players report lag or interruptions during gameplay, per-process metrics help identify bottlenecks in CPU, memory, network, tick time, or world update time.
- **Identify performance changes in different builds** — Tick time, tick rate, and world update time metrics allow developers to measure how game performance changes across different server builds.
- **Detect delays and slowness in gameplay** — Tick time, tick rate, and world update time metrics reflect how fast the server updates the game, directly impacting customer experience.
- **Benchmarking** — Identify how different game scenarios affect server performance based on factors like player count, game mode, and other variables.

Dashboard organization

Metrics are organized into specialized dashboards in [Amazon Managed Grafana](#) for different monitoring scenarios. The available dashboards depend on your fleet type:

EC2 Fleet dashboards

- **EC2 Fleet Overview dashboard** — High-level fleet capacity, scaling insights, concurrent players (CCU), instances, player capacity, and crashed game sessions.
- **Instances Overview dashboard** — Aggregated host-level metrics across all instances including average CPU, memory, network, and disk utilization.
- **Instance Performance dashboard** — Detailed metrics for individual instances with "Top N Memory Consuming Game Sessions" and "Top N CPU Consuming Game Sessions" tables for identifying resource-intensive processes.
- **Server Performance dashboard (EC2)** — Game loop timing, network performance, memory, and CPU metrics for individual game sessions on EC2 instances.

Container Fleet dashboards

- **Container Fleet Overview dashboard** — High-level overview of container fleet resource utilization including CPU reservation, memory utilization, and container group status.
- **Container Performance dashboard** — Detailed metrics for individual containers within specific ECS tasks including CPU utilization, memory usage, network I/O, and storage performance.
- **Server Performance dashboard (Container)** — Game loop timing, network performance, memory, and CPU metrics for individual game sessions in containers.

For detailed dashboard information and usage instructions, see [the section called “Dashboard organization and usage”](#).

How it works

The telemetry metrics system follows a simple four-stage data flow from your game servers to visualization dashboards.

Collection: Your game server, integrated with the GameLift Server SDK or plugin, automatically emits metrics to a local telemetry collector running on the same instance. The SDK captures both automatic metrics (server lifecycle, resource usage) and custom metrics you define in your code.

Processing: The telemetry collector aggregates metrics from your game server and combines them with instance-level performance data (CPU, memory, network, disk usage). This provides a complete picture of both your game's performance and the underlying infrastructure.

Storage: Processed metrics are exported to your choice of metrics warehouse - [Amazon Managed Service for Prometheus](#) for high-performance time-series storage, Amazon CloudWatch for Amazon service integration, or both. All data transmission is authenticated and encrypted.

Visualization: [Amazon Managed Grafana](#) connects to your metrics warehouse to display pre-built GameLift dashboards. These dashboards provide fleet overviews, server performance details, and container monitoring views that help you monitor and troubleshoot your game hosting infrastructure.

 Note

All metric transmission between your game server and the telemetry collector occurs locally on the instance for security. Only the collector communicates with Amazon services using proper authentication.

Dashboard organization and usage

View your metrics on comprehensive dashboards in [Amazon Managed Grafana](#). The available dashboards depend on your fleet type:

Dashboard availability by fleet type

The following table shows which dashboards are available for each fleet type:

Dashboard	Fleet Type	Description
EC2 Fleet Overview	EC2 Fleet	Displays information on concurrent players (CCU), instances and player capacity
Instances Overview	EC2 Fleet	Displays average CPU, memory, and network utilization across all fleet instances
Instance Performance	EC2 Fleet	Displays detailed metrics (CPU, memory, disk, network) for an individual instance

Dashboard	Fleet Type	Description
Container Fleet Overview	Container Fleet	Displays average resource utilization of all containers in a managed container fleet
Container Performance	Container Fleet	Displays detailed metrics of individual containers within a specific ECS task
Server Performance	Both	Displays the network, memory and runtime performance of a specified game server process (separate versions for EC2 and Container fleets)

Managed EC2 Fleets:

- EC2 Fleet Overview provides high-level fleet capacity and scaling insights.
- Use Instances Overview and Instance Performance dashboards for host-level monitoring.
- Metrics collected via hostmetrics receiver for system-level visibility.
- Focus on EC2 instance resource utilization and performance.
- Server Performance (EC2) monitors game server application metrics independent of underlying infrastructure.

Managed Container Fleets:

- Use Container Fleet Overview and Container Performance dashboards for ECS task and container-level monitoring.
- Metrics collected via ECS Container Receiver for containerized workload visibility.
- Focus on task-level aggregation and container resource isolation.
- Server Performance (Container) monitors game server application metrics independent of underlying infrastructure.

EC2 Fleet Overview dashboard

This dashboard provides a high-level overview of your fleet's utilization and capacity globally and by location. It features graphs showing counts for game server stops, starts, and crashes, as well as the percentage of healthy game servers. You can filter by FleetID and Location.

Fleet Overview metrics

The following table shows the metrics available on the Fleet Overview dashboard:

Dashboard	Metric name	Definition
Global CCU and Capacity	CCU	Shows the number of concurrent users in all the game servers over all the instances in all locations
	Global CCU	Shows the number of concurrent users in all the game servers over all the instances globally
	Active Instances	Shows the total number of instances in the fleet that are active
	Active Processes	Shows the number of active game server processes that are ready to host a game session
	Game Server Starts	Shows the number of game sessions that started across the fleet
Global CCU and Capacity	Healthy Game Servers	Shows the average percentage of game servers that report healthy to Amazon GameLift Servers across the fleet
	Crashed Game Sessions	Shows the game session IDs of the crashed game sessions. Click the link of game session ARN to navigate to the Server Performance dashboard

Dashboard	Metric name	Definition
Location CCU	Location Players	Shows the number of concurrent players in a location, including all the game servers over all the instances in the location
	Location Capacity	Shows capacity utilization (%) in a location, and the percentage of game servers in use in the location

Note

CCU metrics require implementation in your game server code. These metrics are not automatically collected and must be implemented and reported by your application.

Instances Overview dashboard

This dashboard provides aggregated host-level metrics across all instances in your fleet. Current averages show overall health of the instances. When performance degrades, check CPU usage, memory consumption, network and disk consumption for bottlenecks. You can filter by FleetID and Location.

Instances Overview metrics

The following table shows the metrics available on the Instances Overview dashboard:

Dashboard	Metric name	Definition
Instance Summary	Average CPU Usage	Instance summary tracks average resource usage by instances. Average percentage of CPU resources being used across all fleet instances. High utilization patterns require checking individual instance metrics and resource distribution
	Peak CPU Usage	Highest CPU utilization recorded across any instance in the fleet during the time period

Dashboard	Metric name	Definition
	CPU Usage by Instance	Individual CPU utilization for each instance in the fleet, enabling identification of resource hotspots
Instance Summary	Average Memory Usage	Average percentage of RAM being used across all fleet instances. High utilization patterns require checking individual instance metrics and resource distribution like memory consumption
	Peak Memory Usage	Highest memory utilization recorded across any instance in the fleet during the time period
	Memory Usage by Instance	Individual memory utilization for each instance in the fleet, helping identify memory-intensive workloads
Network Summary	Network I/O (Bytes)	Network summary reflects average instance connectivity. Average network traffic volume (sent and received) across all fleet instances . For latency or connectivity issues, check individual instance metrics like Network I/O and Packet rates
	Network I/O (Packets)	Average network packet rate (sent and received) across all fleet instances
	Network Usage by Instance	Individual network utilization for each instance, useful for identifying network bottlenecks

Dashboard	Metric name	Definition
Disk Summary	Disk I/O Operations	Disk summary indicates average disk performance of the instances. Average disk read/write operations across all fleet instances . Slow response times suggest examining Disk operations, I/O wait times, and Pending operations at the individual instance level
	Disk I/O Throughput	Average disk read/write throughput across all fleet instances
	Disk Usage by Instance	Individual disk utilization for each instance, helping identify storage performance issues

Note

Instance-level metrics are collected via the hostmetrics receiver and provide system-level visibility into your fleet's infrastructure performance. Use this dashboard to identify overall fleet health trends and drill down to individual instances when performance issues are detected.

Instance Performance dashboard

This dashboard provides detailed performance metrics for individual instances. Current averages show overall instance health. When performance degrades, check CPU usage, memory consumption, and file system consumption for bottlenecks. It features "Top N Memory Consuming Game Sessions" and "Top N CPU Consuming Game Sessions" tables that help identify which processes contribute the most to instance resource consumption. Clicking on Game Session links enables deeper investigation of detailed metrics. You can filter by specific Instance ID.

Instance Performance metrics

The following table shows the metrics available on the Instance Performance dashboard:

Dashboard	Metric name	Definition
Load Summary	Top N Memory Consuming Game Sessions	Load summary tracks resource usage by game servers. Ranked list of game sessions consuming the most memory resources on this instance. High utilization patterns require checking individual session metrics and resource distribution. Clicking on Game Session links enables deeper investigation of detailed metrics
	Top N CPU Consuming Game Sessions	Ranked list of game sessions consuming the most CPU resources on this instance. High utilization patterns require checking individual session metrics and resource distribution like CPU and memory. Clicking on Game Session links enables deeper investigation of detailed metrics
	CPU Usage per Game Session	Individual CPU utilization breakdown showing resource consumption by each active game session
	Memory Usage per Game Session	Individual memory utilization breakdown showing resource consumption by each active game session
Current Averages	Instance CPU Usage	Overall CPU utilization for the selected instance over time
	Instance Memory Usage	Overall memory utilization for the selected instance over time
	Instance File System Usage	File system consumption for the selected instance, useful for identifying storage capacity issues

Dashboard	Metric name	Definition
Network	Instance Network I/O	Network reflects instance connectivity. Network traffic volume and packet rates for the selected instance. For latency or connectivity issues, investigate Network I/O, Packet rates, and Error counts
	Network Connections	Number of active network connections on the selected instance
	Network Error Counts	Network error statistics for identifying connectivity problems
Disk	Disk I/O Operations	Disk indicates disk performance. Disk read/write activity and utilization for the selected instance. Slow response times suggest examining Disk operations, I/O wait times, and Pending operations
	Disk I/O Wait Times	Time spent waiting for disk operations to complete
	Pending Disk Operations	Number of disk operations waiting to be processed

Note

The Top N Memory and CPU Consuming Game Sessions tables are essential for identifying performance bottlenecks and resource-intensive processes that may impact overall instance performance. These rankings enable quick identification of problematic game sessions for further investigation.

Container Fleet Overview dashboard

This dashboard provides a high-level overview of your container fleet's resource utilization and capacity. It displays average resource utilization of all containers in a managed container fleet,

including CPU reservation, memory utilization, and container group status. You can filter by FleetID and Location.

Container Fleet Overview metrics

The following table shows the metrics available on the Container Fleet Overview dashboard:

Dashboard	Metric name	Definition
Container Group Status	Active Game Server Container Groups	Container groups that are currently active and ready to host game sessions
	Idle Game Server Container Groups	Container groups that are active but not currently hosting game sessions
	Pending Game Server Container Groups	Container groups that are in the process of starting up
	Terminating Game Server Container Groups	Container groups that are in the process of shutting down
Resource Utilization	Container CPU Utilization	Average CPU utilization across all containers in the fleet
	Container Memory Utilization	Average memory utilization across all containers in the fleet
	Container CPU Reservation	Percentage of CPU resources reserved by containers across the fleet
Network Activity	Container Network In	Average network bytes received by containers across the fleet
	Container Network Out	Average network bytes sent by containers across the fleet

Note

Container fleet metrics are collected via ECS Container Receiver and provide containerized workload visibility with focus on task-level aggregation and container resource isolation.

Container Performance dashboard

This dashboard provides detailed performance metrics for individual containers within specific ECS tasks. It displays detailed metrics of individual containers including CPU utilization, memory usage, network I/O, and storage performance. You can filter by specific Container ID or ECS Task.

Container Performance metrics

The following table shows the metrics available on the Container Performance dashboard:

Dashboard	Metric name	Definition
Resource Performance	Container CPU Utilization	CPU utilization for the selected container over time
	Container Memory Utilization	Memory utilization for the selected container over time
	Container Memory Reservation	Percentage of memory resources reserved by the selected container
Network Performance	Container Network I/O	Network traffic volume (bytes in/out) for the selected container
	Network Connections	Number of active network connections for the selected container
Storage Performance	Container Storage Read/Write	Storage read and write activity for the selected container
	Container Storage I/O Operations	Number of storage I/O operations performed by the selected container

Note

Container performance metrics provide detailed visibility into individual container resource consumption and performance characteristics within ECS tasks.

Server Performance dashboard

The Server Performance dashboard shows metrics related to server timings, network activity, memory, and CPU usage for individual game sessions. You can filter by Game Session ID and export metrics directly to Amazon CloudWatch or [Amazon Managed Grafana](#).

Server Performance metrics

The following table shows the metrics available on the Server Performance dashboard:

Dashboard	Metric name	Definition
Server Timings	Server Delta Time	Difference in time between the current server tick and the previous server tick. It's a measure of the consistency of the server's tick rate
	Server Tick Rate	Shows the number of times per second the server is processing updates
	Server Tick Time	The amount of time it takes for the server to process a single tick or update
	Server World Tick Time	The amount of time it takes for the server to update the game world with each tick
Network	Connections	The total number of network connections the server has established
	Network I/O (Bytes)	The total number of bytes being sent and received by the server over the network
	Network I/O (Packets)	The total number of network packets being sent and received by the server

Dashboard	Metric name	Definition
Memory	Packet Loss	The percentage of network packets that are being lost during transmission
	Memory Usage (Units)	The total amount of memory being consumed by the server processes
	Physical Memory Usage (%)	The percentage of the server's total physical memory that is currently being utilized
CPU Usage	CPU Usage (%)	The percentage of CPU resources being utilized by the game server process

Common monitoring scenarios

Dive deep performance investigation

Scenario: A host/instance is having degraded performance due to specific processes or game sessions

Investigation steps:

- Access the Instance Performance dashboard.
- Review "Top N Memory Consuming Game Sessions" table to identify which processes contribute the most to instance memory consumption.
- Review "Top N CPU Consuming Game Sessions" table to identify which processes contribute the most to instance CPU utilization.
- Click on Game Session links to enable deeper investigation of detailed metrics.
- Analyze server timings (Server Delta Time, Server Tick Rate, Server Tick Time, Server World Tick Time) to identify performance bottlenecks.

Game server crash investigation

Scenario: A game session has crashed and you need to determine the root cause

Investigation steps:

- Access the Server Performance dashboard for the crashed game session.
- Check Memory Usage (Units) and Physical Memory Usage (%) to determine if the crash was due to out of memory.
- Review CPU Usage (%) to identify if CPU overload caused the crash.
- Analyze Network I/O (Bytes) and Network I/O (Packets) to determine if network bandwidth problems contributed to the crash.
- Examine Packet Loss percentage to identify network-related issues.

Investigate player reported issues

Scenario: Players report lag or interruption during gameplay

Investigation steps:

- Access the Server Performance dashboard for the affected game session.
- Review Server Tick Time and Server World Tick Time to identify delays in game updates.
- Check Server Tick Rate to ensure consistent server update frequency.
- Analyze CPU Usage (%) to identify processing bottlenecks.
- Review Memory Usage metrics to identify memory-related performance issues.
- Check Network I/O metrics and Packet Loss to identify network bottlenecks.

Identify performance changes in different game server builds

Scenario: You want to measure how game performance changes across different server builds

Investigation steps:

- Compare Server Tick Time metrics between different builds to measure processing efficiency changes.
- Analyze Server Tick Rate consistency across builds to identify performance regressions.
- Review Server World Tick Time to measure game world update performance changes.
- Compare Memory Usage patterns between builds to identify memory optimization improvements or regressions.
- Monitor CPU Usage trends to assess computational efficiency changes.

Detect delays and slowness in gameplay

Scenario: You need to monitor server responsiveness and game update speed

Investigation steps:

- Monitor Server Tick Time to measure how fast the server processes each update cycle.
- Track Server Tick Rate to ensure consistent game state updates per second.
- Analyze Server World Tick Time to measure game world update speed, which directly impacts customer experience.
- Set up alerts for Server Delta Time variations to detect inconsistent server performance.

Benchmarking different game scenarios

Scenario: You want to identify how different game scenarios affect server performance

Investigation steps:

- Compare server performance metrics across different player counts to understand scaling impact.
- Analyze performance differences between game modes using Server Tick Time and CPU Usage metrics.
- Monitor Memory Usage patterns across different game scenarios to identify resource-intensive features.
- Track Network I/O metrics to understand bandwidth requirements for different gameplay scenarios.
- Use the Instance Performance dashboard to identify which game scenarios produce the highest resource-consuming game sessions.

High resource utilization response

Scenario: Unusual resource spikes (CPU >85%, Memory >90%)

Investigation steps:

Identify affected resources

- Use DescribeGameSessionDetails API.

- Filter by Status if needed.
- Document affected instances.

Analyze resource usage

- Review Instance Overview dashboard.
- Compare utilization across fleet.
- Check historical patterns.

Monitor game server impact

- Check Server Performance metrics.
- Review tick times and packet loss.
- Monitor memory leaks.

Resolution steps

- Download session logs.
- Address build issues.
- Monitor improvements.

Game server crash analysis

Scenario: Multiple error-status game sessions across fleet

Investigation steps:

Initial assessment

- Access Fleet Overview dashboard.
- Review crashed session table.
- Note patterns in timing/location.

Performance analysis

- Check server timing metrics.

- Review resource utilization.
- Monitor network performance.

Infrastructure review

- Verify fleet capacity.
- Check instance health.
- Review scaling policies.

Resolution path

- Analyze server logs.
- Review code optimization.
- Implement fixes.

Fleet capacity optimization

Scenario: Game launch or benchmark study

Analysis steps:

Resource utilization

- Filter by location.
- Review P50/P95/P99 metrics.
- Analyze usage patterns.

Instance type analysis

- Compare performance by type.
- Identify scaling candidates.
- Document utilization patterns.

Optimization actions

- Adjust scaling policies.

- Modify instance types.
- Update fleet configuration.

Troubleshooting guide

Common issues and resolution steps

Missing or incomplete metrics

Symptoms

- No metrics appearing in dashboards.
- Partial metric collection.
- Delayed metric updates.

Resolution steps

A. Verify collector status

Check systemd service:

```
sudo systemctl status gamelift-telemetry-collector
```

Review collector logs:

```
sudo journalctl -u gamelift-telemetry-collector
```

- Confirm collector configuration.

B. IAM permission verification

- Check instance role permissions.
- Verify required policies:
 - `aps:RemoteWrite`
 - `cloudwatch:PutMetricData`
- Validate role trust relationships.

C. Network connectivity

- Verify endpoint access.
- Check security group rules.
- Review network ACLs.

Authentication errors

Symptoms

- SigV4 authentication failures.
- Access denied messages.
- Credential refresh issues.

Resolution steps

A. SigV4 authentication

- Verify temporary credentials.
- Check credential rotation.
- Validate instance profile.

B. AMP access

- Review workspace configuration.
- Verify remote write URL.
- Check IAM role bindings.

Dashboard issues

Symptoms

- Empty dashboards.
- Missing data points.
- Authentication failures.

Resolution steps

A. Data source configuration

- Verify Prometheus connection.
- Check Amazon CloudWatch integration.
- Test data source permissions.

B. Grafana access

- Confirm SSO configuration.
- Verify 2FA setup if required.
- Check user permissions.

Windows-specific issues

Symptoms

- Service startup failures.
- Metric collection gaps.
- Permission errors.

Resolution steps

- Verify Windows service status.
- Check Windows Event Logs.
- Review collector configuration.
- Validate Windows-specific paths.

Logging Amazon GameLift Servers API calls with Amazon CloudTrail

Amazon GameLift Servers is integrated with Amazon CloudTrail, a service that provides a record of actions taken by a user, role, or an Amazon service in Amazon GameLift Servers. CloudTrail

captures all API calls for Amazon GameLift Servers as events. The calls captured include calls from the Amazon GameLift Servers console and code calls to the Amazon GameLift Servers API operations. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Amazon GameLift Servers. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine the request that was made to Amazon GameLift Servers, the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, see the [Amazon CloudTrail User Guide](#).

Amazon GameLift Servers information in CloudTrail

CloudTrail is enabled on your Amazon Web Services account when you create the account. When activity occurs in Amazon GameLift Servers, that activity is recorded in a CloudTrail event along with other Amazon service events in **Event history**. You can view, search, and download recent events in your Amazon Web Services account. For more information, see [Viewing events with CloudTrail Event history](#).

For an ongoing record of events in your Amazon Web Services account, including events for Amazon GameLift Servers, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all Amazon Web Services Regions. The trail logs events from all Regions in the Amazon partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other Amazon services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- [Overview for creating a trail](#)
- [CloudTrail supported services and integrations](#)
- [Configuring Amazon SNS notifications for CloudTrail](#)
- [Receiving CloudTrail log files from multiple regions](#) and [Receiving CloudTrail log files from multiple accounts](#)

All Amazon GameLift Servers actions are logged by CloudTrail and are documented in the [Amazon GameLift Servers API Reference](#). For example, calls to `CreateGameSession`, `CreatePlayerSession` and `UpdateGameSession` actions generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or Amazon Identity and Access Management (IAM) user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another Amazon service.

For more information, see the [CloudTrail userIdentity element](#).

Understanding Amazon GameLift Servers log file entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the `CreateFleet` and `DescribeFleetAttributes` actions.

```
{
  "Records": [
    {
      "eventVersion": "1.04",
      "userIdentity": {
        "type": "IAMUser",
        "principalId": "AIDACKCEVSQ6C2EXAMPLE",
        "arn": "arn:aws-cn:iam::111122223333:user/myUserName",
        "accountId": "111122223333",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "userName": "myUserName"
      },
      "eventTime": "2015-12-29T23:40:15Z",
      "eventSource": "gamelift.amazonaws.com",
      "eventName": "CreateFleet",
      "awsRegion": "us-west-2",
      "sourceIPAddress": "192.0.2.0",
      "userAgent": "[]",
      "requestParameters": {
        "buildId": "build-92b6e8af-37a2-4c10-93bd-4698ea23de8d",
```



```

        "e2InboundPermissions": [
            {
                "ipRange": "10.24.34.0/23",
                "fromPort": 1935,
                "protocol": "TCP",
                "toPort": 1935
            }
        ],
        "logPaths": [
            "C:\\game\\serverErr.log",
            "C:\\game\\serverOut.log"
        ],
        "e2InstanceType": "c5.large",
        "serverLaunchPath": "C:\\game\\MyServer.exe",
        "description": "Test fleet",
        "serverLaunchParameters": "-paramX=baz",
        "name": "My_Test_Server_Fleet"
    },
    "responseElements": {
        "fleetAttributes": {
            "fleetId": "fleet-0bb84136-4f69-4bb2-bfec-a9b9a7c3d52e",
            "serverLaunchPath": "C:\\game\\MyServer.exe",
            "status": "NEW",
            "logPaths": [
                "C:\\game\\serverErr.log",
                "C:\\game\\serverOut.log"
            ],
            "description": "Test fleet",
            "serverLaunchParameters": "-paramX=baz",
            "creationTime": "Dec 29, 2015 11:40:14 PM",
            "name": "My_Test_Server_Fleet",
            "buildId": "build-92b6e8af-37a2-4c10-93bd-4698ea23de8d"
        }
    },
    "requestID": "824a2a4b-ae85-11e5-a8d6-61d5cafb25f2",
    "eventID": "c8fbea01-fbf9-4c4e-a0fe-ad7dc205ce11",
    "eventType": "AwsApiCall",
    "recipientAccountId": "111122223333"
},
{
    "eventVersion": "1.04",
    "userIdentity": {
        "type": "IAMUser",
        "principalId": "AIDACKCEVSQ6C2EXAMPLE",

```

```

        "arn": "arn:aws-cn:iam::111122223333:user/myUserName",
        "accountId": "111122223333",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "userName": "myUserName"
    },
    "eventTime": "2015-12-29T23:40:15Z",
    "eventSource": "gamelift.amazonaws.com",
    "eventName": "DescribeFleetAttributes",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "192.0.2.0",
    "userAgent": "[]",
    "requestParameters": {
        "fleetIds": [
            "fleet-0bb84136-4f69-4bb2-bfec-a9b9a7c3d52e"
        ]
    },
    "responseElements": null,
    "requestID": "82e7f0ec-ae85-11e5-a8d6-61d5cafb25f2",
    "eventID": "11daabcb-0094-49f2-8b3d-3a63c8bad86f",
    "eventType": "AwsApiCall",
    "recipientAccountId": "111122223333"
    },
]
}

```

Logging server messages in Amazon GameLift Servers

You can capture custom server messages from your Amazon GameLift Servers servers in log files. The way you configure logging depends on whether you use custom servers or Amazon GameLift Servers Realtime, and whether you're using managed EC2 fleets or container fleets (see the appropriate subsections in this chapter).

Topics

- [Logging differences: Managed EC2 vs Container fleets](#)
- [Logging server messages \(custom servers\)](#)

Logging differences: Managed EC2 vs Container fleets

Logging behavior differs significantly between managed EC2 fleets and container fleets. Understanding these differences is crucial when migrating from EC2 to containers or choosing the right fleet type for your logging requirements.

Key differences overview

Feature	Managed EC2 Fleets	Container Fleets
Log storage	Automatically uploaded to Amazon S3 after game session ends	Must configure Amazon CloudWatch Logs or custom logging solution
GetGameSessionLogUrl API	Available - returns Amazon S3 URL for log download	Not available - logs not automatically stored in Amazon S3
Log retention	14 days in Amazon S3 (automatic)	Depends on your logging configuration
Real-time monitoring	Limited - logs available only after game session ends	Available with Amazon CloudWatch Logs integration
Setup complexity	Automatic - no additional configuration required	Requires explicit logging configuration

Managed EC2 fleet logging

For managed EC2 fleets, Amazon GameLift Servers provides automatic log management:

- **Automatic upload:** Server logs are automatically uploaded to Amazon S3 when a game session ends
- **GetGameSessionLogUrl API:** Use this API to retrieve a pre-signed URL for downloading logs from Amazon S3
- **14-day retention:** Logs are retained in Amazon S3 for 14 days before automatic deletion
- **Size limits:** Log files have size limits per game session (see [Amazon GameLift Servers endpoints and quotas](#))

For more information, see [the section called “Logging for custom servers”](#) and the [GetGameSessionLogUrl API reference](#).

Container fleet logging

Container fleets require you to configure logging explicitly:

- **No automatic Amazon S3 upload:** Logs are not automatically uploaded to Amazon S3
- **GetGameSessionLogUrl not available:** This API does not work with container fleets
- **Amazon CloudWatch Logs integration:** Configure your container to send logs to Amazon CloudWatch Logs for centralized logging
- **Custom logging solutions:** Implement your own logging infrastructure using log drivers or sidecar containers
- **Real-time access:** With proper configuration, logs can be accessed in real-time during game sessions

For detailed container logging options, see [How container fleets work](#).

Migration considerations

When migrating from managed EC2 to container fleets, consider these logging changes:

- **Update log retrieval code:** Replace GetGameSessionLogUrl API calls with Amazon CloudWatch Logs queries or your custom logging solution
- **Configure log retention:** Set up appropriate retention policies in Amazon CloudWatch Logs or your logging system
- **Implement real-time monitoring:** Take advantage of real-time log access for better observability
- **Review log volume and costs:** Amazon CloudWatch Logs pricing differs from the included Amazon S3 storage in managed EC2 fleets

Recommended container logging setup

For container fleets, we recommend:

1. **Amazon CloudWatch Logs integration:** Configure your container definition to use the `awslogs` log driver

2. **Structured logging:** Use structured log formats (JSON) for better searchability and analysis
3. **Log levels:** Implement appropriate log levels to control verbosity and costs
4. **Retention policies:** Set retention periods based on your compliance and debugging needs

Logging server messages (custom servers)

You can capture custom server messages from your Amazon GameLift Servers custom servers in log files.

Important

There is a limit on the size of a log file per game session (see [Amazon GameLift Servers endpoints and quotas](#) in the *Amazon Web Services General Reference*). When a game session ends, Amazon GameLift Servers uploads the server logs to Amazon Simple Storage Service (Amazon S3). Amazon GameLift Servers will not upload logs that exceed the limit. Logs can grow very quickly and exceed the size limit. You should monitor your logs and limit the log output to necessary messages only.

Configuring logging for custom servers

With Amazon GameLift Servers custom servers, you write your own code to perform logging, which you configure as part of your server process configuration. Amazon GameLift Servers uses your logging configuration to identify the files that it must upload to S3 at the end of each game session.

The following instructions show how to configure logging using simplified code examples:

C++

To configure logging (C++)

1. Create a vector of strings that are directory paths to game server log files.

```
std::string serverLog("serverOut.log");           // Example server log file
std::vector<std::string> logPaths;
logPaths.push_back(serverLog);
```

2. Provide your vector as the [LogParameters](#) of your [ProcessParameters](#) object.

```
Aws::GameLift::Server::ProcessParameters processReadyParameter =
    Aws::GameLift::Server::ProcessParameters(
        std::bind(&Server::onStartGameSession, this, std::placeholders::_1),
        std::bind(&Server::onProcessTerminate, this),
        std::bind(&Server::OnHealthCheck, this),
        std::bind(&Server::OnUpdateGameSession, this),
        listenPort,
        Aws::GameLift::Server::LogParameters(logPaths));
```

3. Provide the [ProcessParameters](#) object when you call [ProcessReady\(\)](#).

```
Aws::GameLift::GenericOutcome outcome =
    Aws::GameLift::Server::ProcessReady(processReadyParameter);
```

For a more complete example, see [ProcessReady\(\)](#).

C#

To configure logging (C#)

1. Create a list of strings that are directory paths to game server log files.

```
List<string> logPaths = new List<string>();
logPaths.Add("C:\\game\\serverOut.txt");    // Example of a log file that the
game server writes
```

2. Provide your list as the [LogParameters](#) of your [ProcessParameters](#) object.

```
var processReadyParameter = new ProcessParameters(
    this.OnGameSession,
    this.OnProcessTerminate,
    this.OnHealthCheck,
    this.OnGameSessionUpdate,
    port,
    new LogParameters(logPaths));
```

3. Provide the [ProcessParameters](#) object when you call [ProcessReady\(\)](#).

```
var processReadyOutcome =
    GameLiftServerAPI.ProcessReady(processReadyParameter);
```

For a more complete example, see [ProcessReady\(\)](#).

Writing to logs

Your log files exist after your server process has started. You can write to the logs using any method to write to files. To capture all of your server's standard output and error output, remap the output streams to log files, as in the following examples:

C++

```
std::freopen("serverOut.log", "w+", stdout);  
std::freopen("serverErr.log", "w+", stderr);
```

C#

```
Console.SetOut(new StreamWriter("serverOut.txt"));  
Console.SetError(new StreamWriter("serverErr.txt"));
```

Accessing server logs

Log access varies by fleet type:

Managed EC2 fleets

When a game session ends, Amazon GameLift Servers automatically stores the logs in an S3 bucket and retains them for 14 days. To get the location of the logs for a game session, you can use the [GetGameSessionLogUrl](#) API operation. To download the logs, use the URL that the operation returns.

Alternatively, you can set up your own logging solution by configuring your game server to send logs directly to your preferred logging service or storage location. For more information, see [Connect your Amazon GameLift Servers hosted game server to other Amazon resources](#).

Container fleets

Container fleets capture standard output and error streams from all containers. You can configure logging using one of these options:

- Save container output as CloudWatch log streams in a specified log group

- Save container output to an S3 storage bucket
- Turn off logging (container output isn't saved)

For detailed information about configuring logging options when creating container fleets, see [LogConfiguration](#) in the Amazon GameLift Servers API Reference.

Security in Amazon GameLift Servers

If you're using Amazon GameLift Servers FleetIQ as a standalone feature with Amazon EC2, see [Security in Amazon EC2](#) in the *Amazon EC2 User Guide*.

Cloud security at Amazon is the highest priority. As an Amazon customer, you benefit from data centers and network architectures that are built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between Amazon and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – Amazon is responsible for protecting the infrastructure that runs Amazon services in the Amazon Cloud. Amazon also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the [Amazon compliance programs](#). To learn about the compliance programs that apply to Amazon GameLift Servers, see [Amazon services in scope by compliance program](#).
- **Security in the cloud** – Your responsibility is determined by the Amazon service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable Amazon and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Amazon GameLift Servers. The following topics show you how to configure Amazon GameLift Servers to meet your security and compliance objectives. You also learn how to use other Amazon services that help you to monitor and secure your Amazon GameLift Servers resources.

Topics

- [Data protection in Amazon GameLift Servers](#)
- [Identity and access management for Amazon GameLift Servers](#)
- [Logging and monitoring with Amazon GameLift Servers](#)
- [Compliance validation for Amazon GameLift Servers](#)
- [Resilience in Amazon GameLift Servers](#)
- [Infrastructure security in Amazon GameLift Servers](#)
- [Configuration and vulnerability analysis in Amazon GameLift Servers](#)
- [Security best practices for Amazon GameLift Servers](#)

Data protection in Amazon GameLift Servers

If you're using Amazon GameLift Servers FleetIQ as a standalone feature with Amazon EC2, see [Security in Amazon EC2](#) in the *Amazon EC2 User Guide*.

The Amazon [shared responsibility model](#) applies to data protection in Amazon GameLift Servers. As described in this model, Amazon is responsible for protecting the global infrastructure that runs all of the Amazon Web Services Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the Amazon Web Services services that you use. For more information about data privacy, see the [Data Privacy FAQ](#).

For data protection purposes, we recommend that you protect Amazon Web Services account credentials and set up individual users with Amazon IAM Identity Center or Amazon Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with Amazon resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with Amazon CloudTrail. For information about using CloudTrail trails to capture Amazon activities, see [Working with CloudTrail trails](#) in the *Amazon CloudTrail User Guide*.
- Use Amazon encryption solutions, along with all default security controls within Amazon Web Services services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-3 validated cryptographic modules when accessing Amazon through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-3](#).

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with Amazon GameLift Servers or other Amazon Web Services services using the console, API, Amazon CLI, or Amazon SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external

server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Amazon GameLift Servers-specific data is handled as follows:

- Game server builds and scripts that you upload to Amazon GameLift Servers are stored in Amazon S3. There is no direct customer access to this data once it is uploaded. An authorized user can get temporary access to upload files, but can't view or update the files in Amazon S3 directly. To delete scripts and builds, use the Amazon GameLift Servers console or the service API.
- Game session log data is stored in Amazon S3 for a limited period of time after the game session is completed. Authorized users can access the log data by downloading it via a link in the Amazon GameLift Servers console or by calls to the service API.
- Metric and event data is stored in Amazon GameLift Servers and can be accessed through the Amazon GameLift Servers console or by calls to the service API. Data can be retrieved on fleets, instances, game session placements, matchmaking tickets, game sessions, and player sessions. Data can also be accessed through Amazon CloudWatch and CloudWatch Events.
- Customer-supplied data is stored in Amazon GameLift Servers . Authorized users can access it by calls to the service API. Potentially sensitive data might include player data, player session and game session data (including connection info), matchmaker data, and so on.

Note

If you provide custom player IDs in your requests, it is expected that these values are anonymized UUIDs and contain no identifying player information.

For more information about data protection, see the [Amazon shared responsibility model and GDPR](#) blog post on the *Amazon Security Blog*.

Encryption at rest

At-rest encryption of Amazon GameLift Servers-specific data is handled as follows:

- Game server builds and scripts are stored in Amazon S3 buckets with server-side encryption.
- Customer-supplied data is stored in Amazon GameLift Servers in an encrypted format.

Encryption in transit

Connections to the Amazon GameLift Servers APIs are made over a secure (SSL) connection and authenticated using [Amazon Signature Version 4](#) (when connecting through the Amazon CLI or Amazon SDK, signing is handled automatically). Authentication is managed using the IAM-defined access policies for the security credentials that are used to make the connection.

Direct communication between game clients and game servers is as follows:

- For custom game servers being hosted on Amazon GameLift Servers resources, communication does not involve the Amazon GameLift Servers service. Encryption of this communication is the responsibility of the customer. You can use TLS-enabled fleets to have your game clients authenticate the game server on connection and to encrypt all communication between your game client and game server.
- Certificates for TLS-enabled fleets are created at the same time as the fleet, and their expiration dates are based on the fleet creation date.
- For Amazon GameLift Servers Realtime with TLS certificate generation enabled, traffic between game client and Realtime servers using the client SDK for Realtime is encrypted in flight. TCP traffic is encrypted using TLS 1.2, and UDP traffic is encrypted using DTLS 1.2.

Note

Amazon GameLift Servers will adjust the maximum certificate lifetimes for fleets configured to generate a certificate. Certificate lifetimes begin on fleet creation and will be changed on the following schedule:

- Until March 11, 2026, the maximum lifetime for a TLS certificate issued is 398 days.
- As of March 1, 2026, the maximum lifetime for a TLS certificate issued will be 200 days.
- As of March 1, 2027, the maximum lifetime for a TLS certificate issued will be 100 days.
- As of March 1, 2029, the maximum lifetime for a TLS certificate issued will be 47 days.

To ensure that your certificates are renewed, and to maintain up-to-date game server runtime environments, Amazon GameLift Servers recommends regularly replacing your game server fleets.

Internetwork traffic privacy

You can remotely access your Amazon GameLift Servers instances securely. For instances that use Linux, SSH provides a secure communications channel for remote access. For instances that are running Windows, use a remote desktop protocol (RDP) client. With Amazon GameLift Servers FleetIQ, remote access to your instances using Amazon Systems Manager Session Manager and Run Command is encrypted using TLS 1.2, and requests to create a connection are signed using SigV4. For help with connecting to a managed Amazon GameLift Servers instance, see [Remotely connect to Amazon GameLift Servers fleet instances](#).

Identity and access management for Amazon GameLift Servers

Amazon Identity and Access Management (IAM) is an Amazon Web Services service that helps an administrator securely control access to Amazon resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Amazon GameLift Servers resources. IAM is an Amazon Web Services service that you can use with no additional charge.

Topics

- [Audience](#)
- [Authenticating with identities](#)
- [Managing access using policies](#)
- [How Amazon GameLift Servers works with IAM](#)
- [Identity-based policy examples for Amazon GameLift Servers](#)
- [Troubleshooting Amazon GameLift Servers identity and access](#)
- [Amazon managed policies for Amazon GameLift Servers](#)

Audience

How you use Amazon Identity and Access Management (IAM) differs based on your role:

- **Service user** - request permissions from your administrator if you cannot access features (see [Troubleshooting Amazon GameLift Servers identity and access](#))
- **Service administrator** - determine user access and submit permission requests (see [How Amazon GameLift Servers works with IAM](#))

- **IAM administrator** - write policies to manage access (see [Identity-based policy examples for Amazon GameLift Servers](#))

Authenticating with identities

Authentication is how you sign in to Amazon using your identity credentials. You must be authenticated as the Amazon Web Services account root user, an IAM user, or by assuming an IAM role.

For programmatic access, Amazon provides an SDK and CLI to cryptographically sign requests. For more information, see [Amazon Signature Version 4 for API requests](#) in the *IAM User Guide*.

Amazon Web Services account root user

When you create an Amazon Web Services account, you begin with one sign-in identity called the Amazon Web Services account *root user* that has complete access to all Amazon Web Services services and resources. We strongly recommend that you don't use the root user for everyday tasks. For tasks that require root user credentials, see [Tasks that require root user credentials](#) in the *IAM User Guide*.

Federated identity

As a best practice, require human users to use federation with an identity provider to access Amazon Web Services services using temporary credentials.

A *federated identity* is a user from your enterprise directory, web identity provider, or Amazon Directory Service that accesses Amazon Web Services services using credentials from an identity source. Federated identities assume roles that provide temporary credentials.

IAM users and groups

An [IAM user](#) is an identity with specific permissions for a single person or application. We recommend using temporary credentials instead of IAM users with long-term credentials. For more information, see [Require human users to use federation with an identity provider to access Amazon using temporary credentials](#) in the *IAM User Guide*.

An [IAM group](#) specifies a collection of IAM users and makes permissions easier to manage for large sets of users. For more information, see [Use cases for IAM users](#) in the *IAM User Guide*.

IAM roles

An [IAM role](#) is an identity with specific permissions that provides temporary credentials. You can assume a role by [switching from a user to an IAM role \(console\)](#) or by calling an Amazon CLI or Amazon API operation. For more information, see [Methods to assume a role](#) in the *IAM User Guide*.

IAM roles are useful for federated user access, temporary IAM user permissions, cross-account access, cross-service access, and applications running on Amazon EC2. For more information, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Managing access using policies

You control access in Amazon by creating policies and attaching them to Amazon identities or resources. A policy defines permissions when associated with an identity or resource. Amazon evaluates these policies when a principal makes a request. Most policies are stored in Amazon as JSON documents. For more information about JSON policy documents, see [Overview of JSON policies](#) in the *IAM User Guide*.

Using policies, administrators specify who has access to what by defining which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. An IAM administrator creates IAM policies and adds them to roles, which users can then assume. IAM policies define permissions regardless of the method used to perform the operation.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you attach to an identity (user, group, or role). These policies control what actions identities can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Define custom IAM permissions with customer managed policies](#) in the *IAM User Guide*.

Identity-based policies can be *inline policies* (embedded directly into a single identity) or *managed policies* (standalone policies attached to multiple identities). To learn how to choose between managed and inline policies, see [Choose between managed policies and inline policies](#) in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples include *IAM role trust policies* and *Amazon S3 bucket policies*. In services that support resource-

based policies, service administrators can use them to control access to a specific resource. You must [specify a principal](#) in a resource-based policy.

Resource-based policies are inline policies that are located in that service. You can't use Amazon managed policies from IAM in a resource-based policy.

Other policy types

Amazon supports additional policy types that can set the maximum permissions granted by more common policy types:

- **Permissions boundaries** – Set the maximum permissions that an identity-based policy can grant to an IAM entity. For more information, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.
- **Service control policies (SCPs)** – Specify the maximum permissions for an organization or organizational unit in Amazon Organizations. For more information, see [Service control policies](#) in the *Amazon Organizations User Guide*.
- **Resource control policies (RCPs)** – Set the maximum available permissions for resources in your accounts. For more information, see [Resource control policies \(RCPs\)](#) in the *Amazon Organizations User Guide*.
- **Session policies** – Advanced policies passed as a parameter when creating a temporary session for a role or federated user. For more information, see [Session policies](#) in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how Amazon determines whether to allow a request when multiple policy types are involved, see [Policy evaluation logic](#) in the *IAM User Guide*.

How Amazon GameLift Servers works with IAM

Before you use IAM to manage access to Amazon GameLift Servers, learn what IAM features are available to use with Amazon GameLift Servers.

IAM features you can use with Amazon GameLift Servers

IAM feature	Amazon GameLift Servers support
Identity-based policies	Yes
Resource-based policies	No
Policy actions	Yes
Policy resources	Yes
Policy condition keys (service-specific)	Yes
ACLs	No
ABAC (tags in policies)	Yes
Temporary credentials	Yes
Principal permissions	Yes
Service roles	Yes
Service-linked roles	No

To get a high-level view of how Amazon GameLift Servers and other Amazon services work with most IAM features, see [Amazon services that work with IAM](#) in the *IAM User Guide*.

Identity-based policies for Amazon GameLift Servers

Supports identity-based policies: Yes

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Define custom IAM permissions with customer managed policies](#) in the *IAM User Guide*.

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. To learn about all of the elements that you can use in a JSON policy, see [IAM JSON policy elements reference](#) in the *IAM User Guide*.

Identity-based policy examples for Amazon GameLift Servers

To view examples of Amazon GameLift Servers identity-based policies, see [Identity-based policy examples for Amazon GameLift Servers](#).

Resource-based policies within Amazon GameLift Servers

Supports resource-based policies: No

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must [specify a principal](#) in a resource-based policy. Principals can include accounts, users, roles, federated users, or Amazon Web Services services.

To enable cross-account access, you can specify an entire account or IAM entities in another account as the principal in a resource-based policy. For more information, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Policy actions for Amazon GameLift Servers

Supports policy actions: Yes

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Action element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Include actions in a policy to grant permissions to perform the associated operation.

For a list of Amazon GameLift Servers actions, see [Actions defined by Amazon GameLift Servers](#) in the *Service Authorization Reference*.

Policy actions in Amazon GameLift Servers use the following prefix before the action:

```
gamelift
```

To specify multiple actions in a single statement, separate them with commas.

```
"Action": [  
    "gamelift:action1",  
    "gamelift:action2"  
]
```

You can specify multiple actions using wildcards (*). For example, to specify all actions that begin with the word `Describe`, include the following action:

```
"Action": "gamelift:Describe*"
```

To view examples of Amazon GameLift Servers identity-based policies, see [Identity-based policy examples for Amazon GameLift Servers](#).

Policy resources for Amazon GameLift Servers

Supports policy resources: Yes

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Resource JSON policy element specifies the object or objects to which the action applies. As a best practice, specify a resource using its [Amazon Resource Name \(ARN\)](#). For actions that don't support resource-level permissions, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*"
```

For a list of Amazon GameLift Servers resource types and their ARNs, see [Resources defined by Amazon GameLift Servers](#) in the *Service Authorization Reference*. To learn with which actions you can specify the ARN of each resource, see [Actions defined by Amazon GameLift Servers](#).

Some Amazon GameLift Servers resources have ARN values, which allows the resources to have their access managed using IAM policies. The Amazon GameLift Servers fleet resource has an ARN with the following syntax:

```
arn:${Partition}:gamelift:${Region}:${Account}:fleet/${FleetId}
```

For more information about the format of ARNs, see [Amazon Resource Names \(ARNs\)](#) in the *Amazon Web Services General Reference*.

For example, to specify the `fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa` fleet in your statement, use the following ARN:

```
"Resource": "arn:aws:gamelift:us-west-2:123456789012:fleet/
fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa"
```

To specify all fleets that belong to a specific account, use a wildcard (*):

```
"Resource": "arn:aws:gamelift:us-west-2:123456789012:fleet/*"
```

To view examples of Amazon GameLift Servers identity-based policies, see [Identity-based policy examples for Amazon GameLift Servers](#).

Policy condition keys for Amazon GameLift Servers

Supports service-specific policy condition keys: Yes

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Condition` element specifies when statements execute based on defined criteria. You can create conditional expressions that use [condition operators](#), such as equals or less than, to match the condition in the policy with values in the request. To see all Amazon global condition keys, see [Amazon global condition context keys](#) in the *IAM User Guide*.

For a list of Amazon GameLift Servers condition keys, see [Condition keys for Amazon GameLift Servers](#) in the *Service Authorization Reference*. To learn with which actions and resources you can use a condition key, see [Actions defined by Amazon GameLift Servers](#).

To view examples of Amazon GameLift Servers identity-based policies, see [Identity-based policy examples for Amazon GameLift Servers](#).

ACLs in Amazon GameLift Servers

Supports ACLs: No

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

ABAC with Amazon GameLift Servers

Supports ABAC (tags in policies): Yes

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes called tags. You can attach tags to IAM entities and Amazon resources, then design ABAC policies to allow operations when the principal's tag matches the tag on the resource.

To control access based on tags, you provide tag information in the [condition element](#) of a policy using the `aws:ResourceTag/key-name`, `aws:RequestTag/key-name`, or `aws:TagKeys` condition keys.

If a service supports all three condition keys for every resource type, then the value is **Yes** for the service. If a service supports all three condition keys for only some resource types, then the value is **Partial**.

For more information about ABAC, see [Define permissions with ABAC authorization](#) in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see [Use attribute-based access control \(ABAC\)](#) in the *IAM User Guide*.

For an example identity-based policy that limits access to a resource based on the tags on that resource, see [View Amazon GameLift Servers fleets based on tags](#).

Using temporary credentials with Amazon GameLift Servers

Supports temporary credentials: Yes

Temporary credentials provide short-term access to Amazon resources and are automatically created when you use federation or switch roles. Amazon recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see [Temporary security credentials in IAM](#) and [Amazon Web Services services that work with IAM](#) in the *IAM User Guide*.

Cross-service principal permissions for Amazon GameLift Servers

Supports forward access sessions (FAS): Yes

Forward access sessions (FAS) use the permissions of the principal calling an Amazon Web Services service, combined with the requesting Amazon Web Services service to make requests to downstream services. For policy details when making FAS requests, see [Forward access sessions](#).

Service roles for Amazon GameLift Servers

Supports service roles: Yes

A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see [Create a role to delegate permissions to an Amazon Web Services service](#) in the *IAM User Guide*.

Warning

Changing the permissions for a service role might break Amazon GameLift Servers functionality. Edit service roles only when Amazon GameLift Servers provides guidance to do so.

Allow your Amazon GameLift Servers-hosted game servers to access other Amazon resources, such as an Amazon Lambda function or an Amazon DynamoDB database. Because game servers are hosted on fleets that Amazon GameLift Servers manages, you need a service role that gives Amazon GameLift Servers limited access to your other Amazon resources. For more information, see [Connect your Amazon GameLift Servers hosted game server to other Amazon resources](#).

Service-linked roles for Amazon GameLift Servers

Supports service-linked roles: No

A service-linked role is a type of service role that is linked to an Amazon Web Services service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your Amazon Web Services account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

For details about creating or managing service-linked roles, see [Amazon services that work with IAM](#) in the *IAM User Guide*. Find a service in the table that includes a Yes in the **Service-linked roles** column. Choose **Yes** to view the service-linked role documentation for that service.

Identity-based policy examples for Amazon GameLift Servers

By default, users and roles don't have permission to create or modify Amazon GameLift Servers resources. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies.

To learn how to create an IAM identity-based policy by using these example JSON policy documents, see [Create IAM policies \(console\)](#) in the *IAM User Guide*.

For details about actions and resource types defined by Amazon GameLift Servers, including the format of the ARNs for each of the resource types, see [Actions, resources, and condition keys for Amazon GameLift Servers](#) in the *Service Authorization Reference*.

Topics

- [Policy best practices](#)
- [Using the Amazon GameLift Servers console](#)
- [Allow users to view their own permissions](#)
- [Allow player access for game sessions](#)
- [Allow access to one Amazon GameLift Servers queue](#)
- [View Amazon GameLift Servers fleets based on tags](#)
- [Access a game build file in Amazon S3](#)

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Amazon GameLift Servers resources in your account. These actions can incur costs for your Amazon Web Services account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with Amazon managed policies and move toward least-privilege permissions**
 - To get started granting permissions to your users and workloads, use the *Amazon managed policies* that grant permissions for many common use cases. They are available in your Amazon Web Services account. We recommend that you reduce permissions further by defining Amazon customer managed policies that are specific to your use cases. For more information, see [Amazon managed policies](#) or [Amazon managed policies for job functions](#) in the *IAM User Guide*.
- **Apply least-privilege permissions** – When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on

specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see [Policies and permissions in IAM](#) in the *IAM User Guide*.

- **Use conditions in IAM policies to further restrict access** – You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific Amazon Web Services service, such as Amazon CloudFormation. For more information, see [IAM JSON policy elements: Condition](#) in the *IAM User Guide*.
- **Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions** – IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and functional policies. For more information, see [Validate policies with IAM Access Analyzer](#) in the *IAM User Guide*.
- **Require multi-factor authentication (MFA)** – If you have a scenario that requires IAM users or a root user in your Amazon Web Services account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see [Secure API access with MFA](#) in the *IAM User Guide*.

For more information about best practices in IAM, see [Security best practices in IAM](#) in the *IAM User Guide*.

Using the Amazon GameLift Servers console

To access the Amazon GameLift Servers console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Amazon GameLift Servers resources in your Amazon Web Services account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (users or roles) with that policy.

To ensure that those entities can still use the Amazon GameLift Servers console, add permissions to users and groups with the syntax in the following examples and in [Administration permission examples](#). For more information, see [Set user permissions for Amazon GameLift Servers](#).

Users that work with Amazon GameLift Servers through Amazon CLI or Amazon API operations don't require minimum console permissions. Instead, you can limit access to only the operations

the user needs to perform. For example, a player user, acting on behalf of game clients, requires access to request game sessions, place players into games, and other tasks.

For information about the permissions required to use all Amazon GameLift Servers console features, see permissions syntax for administrators in [Administration permission examples](#).

Allow users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the Amazon CLI or Amazon API.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ViewOwnUserInfo",
      "Effect": "Allow",
      "Action": [
        "iam:GetUserPolicy",
        "iam:ListGroupsWithUser",
        "iam:ListAttachedUserPolicies",
        "iam:ListUserPolicies",
        "iam:GetUser"
      ],
      "Resource": ["arn:aws-cn:iam::*:user/${aws:username}"]
    },
    {
      "Sid": "NavigateInConsole",
      "Effect": "Allow",
      "Action": [
        "iam:GetGroupPolicy",
        "iam:GetPolicyVersion",
        "iam:GetPolicy",
        "iam:ListAttachedGroupPolicies",
        "iam:ListGroupPolicies",
        "iam:ListPolicyVersions",
        "iam:ListPolicies",
        "iam:ListUsers"
      ],
      "Resource": "*"
    }
  ]
}
```

```
}
```

Allow player access for game sessions

To place players into game sessions, game clients and backend services need permissions. For policy examples for these scenarios, see [Player user permission examples](#).

Allow access to one Amazon GameLift Servers queue

The following example provides a user with access to a specific Amazon GameLift Servers queues.

This policy grants the user permissions to add, update, and delete queue destinations with the following actions: `gamelift:UpdateGameSessionQueue`, `gamelift>DeleteGameSessionQueue`, and `gamelift:DescribeGameSessionQueues`. As shown, this policy uses the Resource element to limit access to a single queue: `gamesessionqueue/examplequeue123`.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ViewSpecificQueueInfo",
      "Effect": "Allow",
      "Action": [
        "gamelift:DescribeGameSessionQueues"
      ],
      "Resource": "arn:aws-cn:gamelift:us-east-1:555555555555:gamesessionqueue/examplequeue123"
    },
    {
      "Sid": "ManageSpecificQueue",
      "Effect": "Allow",
      "Action": [
        "gamelift:UpdateGameSessionQueue",
        "gamelift>DeleteGameSessionQueue"
      ],
      "Resource": "arn:aws-cn:gamelift:us-east-1:111122223333:gamesessionqueue/examplequeue123"
    }
  ]
}
```

```
]
}
```

View Amazon GameLift Servers fleets based on tags

You can use conditions in your identity-based policy to control access to Amazon GameLift Servers resources based on tags. This example shows how you can create a policy that allows viewing a fleet if the `Owner` tag matches the user's user name. This policy also grants the permissions necessary to complete this operation in the console.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ListFleetsInConsole",
      "Effect": "Allow",
      "Action": "gamelift:ListFleets",
      "Resource": "*"
    },
    {
      "Sid": "ViewFleetIfOwner",
      "Effect": "Allow",
      "Action": "gamelift:DescribeFleetAttributes",
      "Resource": "arn:aws-cn:gamelift:*:*:fleet/*",
      "Condition": {
        "StringEquals": {"aws:ResourceAccount": "${aws:username}"}
      }
    }
  ]
}
```

Access a game build file in Amazon S3

After you integrate your game server with Amazon GameLift Servers, upload the build files to Amazon S3. For Amazon GameLift Servers to access the build files, use the following policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:GetObjectVersion"
      ],
      "Resource": "arn:aws:s3:::bucket-name/object-name"
    }
  ]
}
```

For more information about uploading Amazon GameLift Servers game files, see [Create a game server build for Amazon GameLift Servers](#).

Troubleshooting Amazon GameLift Servers identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with Amazon GameLift Servers and Amazon Identity and Access Management (IAM).

Topics

- [I am not authorized to perform an action in Amazon GameLift Servers](#)
- [I am not authorized to perform iam:PassRole](#)
- [I want to allow people outside of my Amazon Web Services account to access my Amazon GameLift Servers resources](#)

I am not authorized to perform an action in Amazon GameLift Servers

If the Amazon Web Services Management Console tells you that you're not authorized to perform an action, contact your Amazon account administrator for assistance. Your administrator is the person that provided you with your sign-in credentials.

The following example error occurs when the `mateojackson` IAM user tries to use the console to view details about a queue but doesn't have `gamelift:DescribeGameSessionQueues` permissions:

```
User: arn:aws-cn:iam::123456789012:user/mateojackson is not authorized to perform:
gamelift:DescribeGameSessionQueues on resource: examplequeue123
```

In this case, Mateo asks his administrator to update his policies to allow him read access for the `examplequeue123` resource using the `gamelift:DescribeGameSessionQueues` action.

I am not authorized to perform `iam:PassRole`

If you receive an error that you're not authorized to perform the `iam:PassRole` action, your policies must be updated to allow you to pass a role to Amazon GameLift Servers.

Some Amazon Web Services services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named `marymajor` tries to use the console to perform an action in Amazon GameLift Servers. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws-cn:iam::123456789012:user/marymajor is not authorized to perform:
iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the `iam:PassRole` action.

If you need help, contact your Amazon administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my Amazon Web Services account to access my Amazon GameLift Servers resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Amazon GameLift Servers supports these features, see [How Amazon GameLift Servers works with IAM](#).
- To learn how to provide access to your resources across Amazon Web Services accounts that you own, see [Providing access to an IAM user in another Amazon Web Services account that you own](#) in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party Amazon Web Services accounts, see [Providing access to Amazon Web Services accounts owned by third parties](#) in the *IAM User Guide*.
- To learn how to provide access through identity federation, see [Providing access to externally authenticated users \(identity federation\)](#) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Amazon managed policies for Amazon GameLift Servers

An Amazon managed policy is a standalone policy that is created and administered by Amazon. Amazon managed policies are designed to provide permissions for many common use cases so that you can start assigning permissions to users, groups, and roles.

Keep in mind that Amazon managed policies might not grant least-privilege permissions for your specific use cases because they're available for all Amazon customers to use. We recommend that you reduce permissions further by defining [customer managed policies](#) that are specific to your use cases.

You cannot change the permissions defined in Amazon managed policies. If Amazon updates the permissions defined in an Amazon managed policy, the update affects all principal identities (users, groups, and roles) that the policy is attached to. Amazon is most likely to update an Amazon managed policy when a new Amazon Web Services service is launched or new API operations become available for existing services.

For more information, see [Amazon managed policies](#) in the *IAM User Guide*.

Amazon managed policy: GameLiftContainerFleetPolicy

You can attach GameLiftContainerFleetPolicy to your IAM roles.

The policy grants permissions for compute actions in an Amazon GameLift Servers container fleet. A container fleet is a set of hosting resources that Amazon GameLift Servers manages for you. Amazon GameLift Servers needs permissions to connect to the Amazon GameLift Servers service and other Amazon services on your behalf.

When creating a container fleet with Amazon GameLift Servers, provide an IAM service role with the `GameLiftContainerFleetPolicy` managed policy attached. For instructions on creating the service role, see [Set up an IAM service role for Amazon GameLift Servers](#).

For more information, see [GameLiftContainerFleetPolicy](#).

Permissions details

This policy includes the following permissions.

- `cloudwatch` – Allows Amazon GameLift Servers to write game session logs to an Amazon CloudWatch Events log stream in your Amazon account.
- `cloudwatch` – Allows Amazon GameLift Servers to create an CloudWatch log group to organize game session data in the log stream.
- `s3` – Allows Amazon GameLift Servers to write game session logs to an Amazon Simple Storage Service bucket in your Amazon account.
- `s3` – Allows Amazon GameLift Servers to retrieve the Amazon Web Services Region where a specified Amazon S3 bucket resides, using the API action `s3:GetBucketLocation`.
- `gamelift` – Allows Amazon GameLift Servers to retrieve an authentication token that allows a hosted game server to communicate with the Amazon GameLift Servers service through your Amazon account.

Amazon GameLift Servers updates to Amazon managed policies

View details about updates to Amazon managed policies for Amazon GameLift Servers since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the [Amazon GameLift Servers Release notes page](#).

Change	Description	Date
GameLiftContainerFleetPolicy – Change	Amazon GameLift Servers added new permissions to	February 5, 2024

Change	Description	Date
	retrieve the Amazon Web Services Region of an Amazon S3 bucket.	
GameLiftContainerFleetPolicy – New policy	Amazon GameLift Servers added new permissions to enable game server containers to run on Amazon GameLift Servers managed fleets.	November 12, 2024
Amazon GameLift Servers started tracking changes	Amazon GameLift Servers started tracking changes for its Amazon managed policies.	November 12, 2024

Logging and monitoring with Amazon GameLift Servers

Monitoring is an important part of maintaining the reliability, availability, and performance of Amazon GameLift Servers and your Amazon solutions. You should collect monitoring data from all of the parts of your Amazon solution so that you can more easily debug a multi-point failure if one occurs.

Amazon and Amazon GameLift Servers provide several tools for monitoring your game hosting resources and responding to potential incidents.

Amazon CloudWatch Alarms

Using Amazon CloudWatch alarms, you watch a single metric over a time period that you specify. If the metric exceeds a given threshold, a notification is sent to an Amazon SNS topic or Amazon Auto Scaling policy. CloudWatch alarms are triggered when their state changes and is maintained for a specified number of periods, not by being in a particular state. For more information, see [Monitor Amazon GameLift Servers with Amazon CloudWatch](#).

Amazon CloudTrail Logs

CloudTrail provides a record of actions taken by a user, role, or an Amazon service in Amazon GameLift Servers. Using the information collected by CloudTrail, you can determine the request that was made to Amazon GameLift Servers, the IP address from which the request was made, who

made the request, when it was made, and additional details. For more information, see [Logging Amazon GameLift Servers API calls with Amazon CloudTrail](#).

Compliance validation for Amazon GameLift Servers

Amazon GameLift Servers is not in scope of any Amazon compliance programs.

To learn whether an Amazon Web Services service is within the scope of specific compliance programs, see [Amazon Web Services services in Scope by Compliance Program](#) and choose the compliance program that you are interested in. For general information, see [Amazon Web Services Compliance Programs](#).

You can download third-party audit reports using Amazon Artifact. For more information, see [Downloading Reports in Amazon Artifact](#).

Your compliance responsibility when using Amazon Web Services services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. For more information about your compliance responsibility when using Amazon Web Services services, see [Amazon Security Documentation](#).

Resilience in Amazon GameLift Servers

If you're using Amazon GameLift Servers FleetIQ as a standalone feature with Amazon EC2, see [Security in Amazon EC2](#) in the *Amazon EC2 User Guide*.

The Amazon global infrastructure is built around Amazon Regions and Availability Zones. Amazon Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about Amazon Regions and Availability Zones, see [Amazon global infrastructure](#).

In addition to the Amazon global infrastructure, Amazon GameLift Servers offers the following features to help support your data resiliency needs:

- **Multi-region queues** – Amazon GameLift Servers game session queues are used to place new game sessions with available hosting resources. Queues that span multiple Regions are able to

redirect game session placements in the event of a regional outage. For more information and best practices on creating game session queues, see [Customize a game session queue](#).

- **Automatic capacity scaling** – Maintain the health and availability of your hosting resources by using Amazon GameLift Servers scaling tools. These tools provide a range of options that let you adjust fleet capacity to fit the needs of your game and players. For more information on scaling, see [Scaling game hosting capacity with Amazon GameLift Servers](#).
- **Distribution across instances** – Amazon GameLift Servers distributes incoming traffic across multiple instances, depending on fleet size. As a best practice, games in production should have multiple instances to maintain availability in case an instance becomes unhealthy or unresponsive.
- **Amazon S3 storage** – Game server builds and scripts that are uploaded to Amazon GameLift Servers are stored in Amazon S3 using the Standard storage class, which uses multiple data center replications to increase resilience. Game session logs are also stored in Amazon S3 using the Standard storage class.

Infrastructure security in Amazon GameLift Servers

If you're using Amazon GameLift Servers FleetIQ as a standalone feature with Amazon EC2, see [Security in Amazon EC2](#) in the *Amazon EC2 User Guide*.

As a managed service, Amazon GameLift Servers is protected by the Amazon global network security procedures that are described in the [Amazon Web Services: Overview of security processes](#) whitepaper.

You use Amazon published API calls to access Amazon GameLift Servers through the network. Clients must support Transport Layer Security (TLS) 1.2 or later. We recommend TLS 1.3 or later. Clients must also support cipher suites with perfect forward secrecy (PFS) such as Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Ephemeral Diffie-Hellman (ECDHE). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the [Amazon Security Token Service](#) (Amazon STS) to generate temporary security credentials to sign requests.

The Amazon GameLift Servers service places all fleets into Amazon virtual private clouds (VPCs) so that each fleet exists in a logically isolated area in the Amazon Cloud. You can use Amazon GameLift Servers policies to control access from specific VPC endpoints or specific VPCs.

Effectively, this isolates network access to a given Amazon GameLift Servers resource from only the specific VPC within the Amazon network. When you create a fleet, you specify a range of port numbers and IP addresses. These ranges limit how inbound traffic can access hosted game servers on a fleet VPC. Use standard security best practices when choosing fleet access settings.

Configuration and vulnerability analysis in Amazon GameLift Servers

If you're using Amazon GameLift Servers FleetIQ as a standalone feature with Amazon EC2, see [Security in Amazon EC2](#) in the *Amazon EC2 User Guide*.

Configuration and IT controls are a shared responsibility between Amazon and you, our customer. For more information, see the Amazon [shared responsibility model](#). Amazon handles basic security tasks like guest operating system (OS) and database patching, firewall configuration, and disaster recovery. These procedures have been reviewed and certified by the appropriate third parties. For more details, see the following resource: [Amazon Web Services: Overview of security processes](#) (whitepaper).

The following security best practices also address configuration and vulnerability analysis in Amazon GameLift Servers:

- Customers are responsible for the management of software that is deployed to Amazon GameLift Servers instances for game hosting. Specifically:
 - Customer-provided game server application software should be maintained, including updates and security patches. To update game server software, upload a new build to Amazon GameLift Servers or provide it in a new container image, create a new fleet for it, and redirect traffic to the new fleet.
 - The base Amazon Machine Image (AMI), which includes the operating system, is updated only when a new fleet is created. To patch, update, and secure the operating system and other applications that are part of the AMI, replace fleets on a regular basis, regardless of game server updates. At minimum, production fleets should be replaced at least once a year. See [Keep fleet runtime environments up to date](#) for more details.
- Customers should consider regularly updating their games with the latest SDK versions, including the Amazon SDK, the Amazon GameLift Servers Server SDK, and the Amazon GameLift Servers Client SDK for Realtime Servers.

Security best practices for Amazon GameLift Servers

If you're using Amazon GameLift Servers FleetIQ as a standalone feature with Amazon EC2, see [Security in Amazon EC2](#) in the *Amazon EC2 User Guide*.

Amazon GameLift Servers provides a number of security features to consider as you develop and implement your own security policies. The following best practices are general guidelines and don't represent a complete security solution. Because these best practices might not be appropriate or sufficient for your environment, treat them as helpful considerations rather than prescriptions.

Keep fleet runtime environments up to date

Amazon GameLift Servers highly recommends that you regularly replace managed fleets (including managed EC2 and managed container fleets) to maintain secure runtime environments for your game servers. Fleets that run for extended periods without runtime updates can contain outdated dependencies and security vulnerabilities that could compromise your game servers. For details on how responsibility is shared for software deployed to Amazon GameLift Servers fleets, see [Configuration and vulnerability analysis in Amazon GameLift Servers](#).

A managed fleet's runtime environment is determined by its Amazon Machine Image (AMI) version. 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. For details on current AMI versions, see [Amazon GameLift Servers AMI versions](#).

Recommended practices:

- **Monitor fleet age and replace fleets that are over 30 days old** – You can track a fleet's creation date in the Amazon GameLift Servers console or use the CLI to retrieve fleet attributes. Amazon GameLift Servers displays warnings in the console for fleets that are over 90 days old, and notifies account holders by email for fleets that are older than one year.

Note

Updating a fleet (such as using [UpdateFleetAttributes](#) or [UpdateContainerFleet](#)) doesn't change the AMI version. You must create a new fleet.

- **Replace fleets regularly based on security health** – Set up a regular schedule to create new fleets and retire old fleets. Consider using a service like [Amazon Q](#) to review your game code with the current AMI version, detect security issues, and suggest remediation steps.
- **Test server builds with the latest AMI versions before deployment** – You might need to modify your server build and upload it to Amazon GameLift Servers before creating a new fleet.
- **Manage fleet quotas for your Amazon account** – You can request limit increases if needed to create replacement fleets. For more information, see [Amazon GameLift Servers endpoints and quotas](#).
- **Consider automating fleet replacement** – You can automate processes to create new fleets and migrate player traffic from older fleets. For example:
 - Use Amazon CloudFormation to automate fleet creation and management. Maintain your fleet configurations as Amazon CloudFormation templates and use them to launch resource stacks.
 - Take advantage of the Amazon GameLift Servers alias feature to abstract specific fleet IDs. Fleet aliases make it easy to switch player traffic from an existing fleet to a new one with no disruption to game sessions in progress. For details, see [Abstract an Amazon GameLift Servers fleet designation with an alias](#).
 - Use blue/green deployment strategies to reduce migration risk and maintain zero downtime. With two identical production environments, you can take advantage of a full production-like testing environment, exert greater control over migration process, and ensure instant rollbacks.

Secure your port configurations

We strongly recommend against opening ports to the Internet because doing so poses a security risk. For example, the following configuration opens a remote desktop port that allows anyone on the Internet to access the instance:

```
{
  "FleetId": "<fleet identifier>",
  "InboundPermissionAuthorizations": [
    {
      "FromPort": 3389,
      "IpRange": "0.0.0.0/0",
      "Protocol": "RDP",
      "ToPort": 3389
    }
  ]
}
```

```
}
```

Instead, use [UpdateFleetPortSettings](#) to open a port with a specific IP address or range of addresses, as shown in this example:

```
{
  "FleetId": "<fleet identifier>",
  "InboundPermissionAuthorizations": [
    {
      "FromPort": 3389,
      "IpRange": "54.186.139.221/32",
      "Protocol": "TCP",
      "ToPort": 3389
    }
  ]
}
```

Additional security resources

For more information about how you can make your use of Amazon GameLift Servers more secure, see the [Amazon Well-Architected Tool Security pillar](#).

Amazon GameLift Servers pricing and cost planning

Understanding the Amazon GameLift Servers pricing structure and implementing cost-effective strategies in your game hosting solution is essential for optimizing your gaming infrastructure budget. For detailed service information, including example pricing scenarios and pricing tools, see [Amazon GameLift Servers Pricing](#).

New Amazon customers can use Amazon GameLift Servers without incurring charges under the Free Tier for up to 12 months. Stay within Free Tier usage limits to avoid charges.

Pricing models

Amazon GameLift Servers offers several pricing models designed to accommodate different usage scenarios and business needs. Each model has a different cost basis. You can control costs by adjusting the cost factors in each model.

Managed hosting

Managed EC2 and managed container fleets use Amazon EC2 instances that Amazon GameLift Servers manage. When you set up game hosting, you choose the type of Amazon EC2 instance to use.

- **Compute costs** – Pay for hourly EC2 instance usage for hosting game sessions. Cost factors for instances include:
 - Amazon Region
 - Instance family, type, and size
 - Use of Spot versus On-Demand instances
 - Operating system (managed container fleets run on Linux only)
- **Data transfer** – Pay for traffic between game clients and hosted game servers

Anywhere hosting

Anywhere hosting runs on game hosting resources that you supply and manage, so there are no compute costs. Instead costs are based on:

- **Game sessions** – Pay based on the number of game sessions placed on Anywhere computes
- **Connection minutes** – Pay for server process connection minutes

Use the **Pricing calculator for Amazon GameLift Servers** to estimate costs for different Amazon GameLift Servers configurations. Use the calculator when designing your game hosting solution or to forecast costs for configuration changes. See [Generate Amazon GameLift Servers pricing estimates](#).

Generate Amazon GameLift Servers pricing estimates

With Amazon Pricing Calculator, you can [create a pricing estimate for Amazon GameLift Servers](#). You don't need an Amazon Web Services account or in-depth knowledge of Amazon to use the calculator.

Amazon Pricing Calculator calculator guides you through the decisions that affect service costs to give you an idea of how much Amazon GameLift Servers might cost for your game project. If you're not yet sure how you plan to use Amazon GameLift Servers, then use the default values to generate an estimate. When planning for production usage, the calculator can help you test out potential scenarios and generate more accurate estimates.

Estimate Amazon GameLift Servers managed hosting

This option provides a cost estimate for hosting your games on Amazon GameLift Servers managed servers, including the costs for server instance usage and data transfer. With Amazon GameLift Servers managed hosting, there is no additional cost for FlexMatch matchmaking.

If you are hosting or plan to host game servers in more than one Amazon Region or on more than one instance type, create an estimate for each Region and instance type.

Amazon GameLift Servers instances

This section helps you estimate the type and number of compute resources that you need to host game sessions for your players. Amazon GameLift Servers uses [Amazon Elastic Compute Cloud \(Amazon EC2\) instances](#) to manage game servers. In Amazon GameLift Servers, you deploy a fleet of instances with a specific instance type and operating system. If you have or plan to have multiple fleets, create an estimate for each fleet.

To get started, open the [Configure Amazon GameLift Servers page](#) of Amazon Pricing Calculator. Add a **Description**, choose a **Region**, and then choose **Estimate Amazon GameLift Servers hosting (Instance + Data Transfer Out)**. Under **Amazon GameLift Servers instances**, complete the following fields:

- **Peak concurrent players (peak CCU)**

This is the maximum number of players who can connect to your game servers at the same time. This field indicates how much hosting capacity Amazon GameLift Servers needs to meet peak player demand. Enter the daily peak number of players that you expect to host using instances in your chosen Amazon Region.

For example, if you want to let 1,000 players connect to your game at any one time, keep the default value of **1000**.

- **Average CCU per hour as a percentage of peak daily CCU**

This is the average number of concurrent players per hour over a 24-hour period. We use this value to estimate the amount of sustained hosting capacity that Amazon GameLift Servers needs to maintain for your players. If you're not sure what percentage value to use, keep the default value of **50** percent. For games with stable player demand, we recommend entering a value of **70** percent.

For example, if your game has an average hourly CCU of 6,000 and a peak CCU of 10,000, then enter the value of **60** percent.

- **Game sessions per instance**

This is the number of game sessions that each of your game server instances can host concurrently. Factors that can affect this number include the resource requirements of your game server, the number of players to host in each game session, and player performance expectations. If you know the number of concurrent game sessions for your game, then enter that value. Alternatively, keep the default value of **20**.

- **Players per game session**

This is the average number of players who connect to a game session, as defined in your game design. If you have game modes with different number of players, estimate an average number of players per game session across your entire game. The default value is **8**.

- **Instance idle buffer %**

This is the percentage of unused hosting capacity to maintain in reserve to handle sudden spikes in player demand. Buffer size is a percentage of the total number of instances in a fleet. The default value is **10** percent.

For example, with a 20 percent idle buffer, a fleet supporting players with 100 active instances maintains 20 idle instances.

- **Spot instance %**

Amazon GameLift Servers fleets can use a combination of On-Demand Instances and Spot Instances. While On-Demand Instances offer more reliable availability, Spot Instances offer a highly cost-efficient alternative. We recommend using a combination to optimize both cost savings and availability. For information about how Amazon GameLift Servers uses Spot Instances, see [On-Demand Instances versus Spot Instances](#).

For this field, enter the percentage of Spot Instances to maintain in a fleet. We recommend a Spot Instance percentage between 50 and 85 percent. The default value is **50** percent.

For example, if you deploy a fleet with 100 instances and specify **40** percent, Amazon GameLift Servers works to maintain 60 On-Demand Instances and 40 Spot Instances.

- **Instance type**

Amazon GameLift Servers fleets can use a range of Amazon EC2 instance types that vary in computing power, memory, storage, and networking capabilities. When you configure an Amazon GameLift Servers fleet, choose an instance type that best fits your game's needs. For information about selecting an instance type with Amazon GameLift Servers, see [Choose compute resources for a managed fleet](#).

If you know the instance type that you're using or plan to use in your Amazon GameLift Servers fleet, choose that type. If you're not sure what type to choose, consider choosing **c5.large**. This is a high-availability type with average size and capabilities.

- **Operating system**

This field specifies the operating system that your game servers run on—either Linux or Windows. The default value is **Linux**.

Data transfer out (DTO)

This section helps you estimate the cost for traffic between your game clients and the game servers. Data transfer fees apply to outbound traffic only. Inbound data transfer has no cost.

On the [Configure Amazon GameLift Servers page](#) of Amazon Pricing Calculator, expand **Data transfer out (DTO)**, and then complete the following fields:

- **DTO estimate type**

You can choose to estimate DTO in either of the following two ways, depending on how you track data transfer for your game.

- **Per month (in GB)** – If you track monthly traffic for your game servers, choose this type.
- **Per player** – If you track data transfer by player, choose this type. This is the default type.

In the following field, you estimate per-player DTO based on the number of player hours that you calculated in the previous section.

- **DTO per month (in GB)**

If you chose the **Per month (in GB)** DTO estimate type, then enter your estimated monthly DTO usage in GB from each instance, per Region.

- **DTO per player**

If you chose the **Per player** DTO estimate type, then enter your game's estimated DTO usage per player in KB/sec. The default value is 4.

When you're done configuring your Amazon GameLift Servers pricing estimate, choose **Add to my estimate**. For more information about creating and managing estimates in Amazon Pricing Calculator, see [Create an estimate, configure a service, and add more services](#) in the *Amazon Pricing Calculator User Guide*.

Cost optimization strategies

Use these strategies with a managed hosting solution to help reduce your cloud hosting costs while maintaining high performance and player experience.

Best practices

Follow these tips to incorporate cost optimization practices for your game hosting. At a minimum, we recommend that you schedule regular reviews of your Amazon GameLift Servers costs and usage patterns. If you have pricing-related questions, reach out to your

Fleet configuration

- **Use the pricing calculator** – Use this tool with your game hosting data to explore options and test out potential configuration scenarios for cost savings. See [Amazon Pricing Calculator for Amazon GameLift Servers](#).

- **Save in development** – For development and testing, use smaller, less expensive instance types and run them only when you need to.
- **Match instance type to game** – Choose Amazon EC2 instance families, types, and sizes for your fleets that best suit your game's requirements. See the Instance types section in [Choose compute resources for a managed fleet](#).
- **Use Graviton instance types** – Graviton instances are powered by ARM-based processors. They offer better price-performance, higher energy efficiency, and lower costs than comparable x86-based instances with the On-Demand pricing model.
- **Use Spot instances** – Spot instances are usually lower cost than On-Demand but may not always be available. Combine Spot and On-Demand fleets to balance low cost and high availability (50-85% Spot usage recommended). Use Spot instances for non-critical game modes or during off-peak hours. To monitor Spot instance viability, track metrics such as `InstanceInterruptions` and `GameServerInterruptions`. See [On-Demand Instances versus Spot Instances](#).
- **Run servers on Linux** – Build your game server runtime for Linux. Instances that are deployed with Linux are generally more cost-efficient than those with Windows.
- **Optimize resource utilization** – Configure fleets for maximum efficiency by running as many concurrent game server processes as possible while maintaining performance. See [Optimize game server runtime configuration on managed Amazon GameLift Servers](#).
- **Create a multi-region strategy** – Deploying game hosting to multiple locations can give your players lower latency and better backup options. Balance regional coverage and cost efficiency by deploying hosting in your largest player markets first and using secondary locations for overflow capacity during peak times. See [Build a multi-location queue](#).

Fleet capacity scaling

- **Scale to zero** – When fleets are not in use, manually set fleet capacity to zero to avoid unnecessary charges. See [Manually set capacity for an Amazon GameLift Servers fleet](#).
- **Add auto scaling** – Avoid over-provisioning hosting resources by using auto scaling to adjust game hosting capacity. Match capacity to fluctuating player demand and other key metrics. See [Auto-scale fleet capacity with Amazon GameLift Servers](#).
- **Maintain a buffer** – To handle sudden spikes in player demand without making players wait, use target tracking to maintain a buffer of idle game servers. Customize the buffer based on the size and usage patterns of your player base. See [Target-based auto scaling](#).

Game session placement

- **Use placement queues** – Amazon GameLift Servers queues rely on an algorithm to determine the "best possible" hosting locations for game sessions, based on the cost of hosting resources and other factors. See [Configure game session placement](#).
- **Customize queue priorities** – You can change how a queue prioritizes hosting costs when placing game sessions. See [Prioritize game session placement](#).
- **Monitor placement metrics** – Track how queues are choosing locations for game sessions to identify optimization opportunities. See [Amazon GameLift Servers metrics for queues](#).

Data transfer

- **Track data transfers** – Monitor how data is transferred between your game clients and servers and take steps to optimize activity.
- **Use data compression** – Consider implementing data compression techniques for data transfers between game clients and servers. Compression plays a crucial role in reducing bandwidth usage as well as improving gameplay performance and latency.

Resource cost and utilization tools

Explore how to use Amazon tools to monitor and optimize your game hosting costs with Amazon GameLift Servers. For information on additional tools, see [Amazon Billing and Cost Management](#).

Amazon cost management tools

- **Billing console** – Review your Amazon bills and usage. See [Getting set up with Billing](#).
- **Free Tier usage alerts** – Set up notifications when approaching Free Tier limits. See [Tracking your Amazon Free Tier usage](#).
- **Amazon CloudWatch billing alerts** – Configure alerts when usage hits custom thresholds. See [Create a billing alarm to monitor your estimated Amazon charges](#).

Cost allocation and tracking

- **Cost allocation tags** – Tag your fleets and other resources to organize and track hosting costs. See [Organizing and tracking costs using Amazon cost allocation tags](#).

- **Cost reports** – Create reports categorized by assigned tags. See [Using the monthly cost allocation report](#).
- **Amazon Cost Explorer** – Analyze costs, trends, and forecasts with filtering and customizable views. See [Analyzing your costs and usage with Amazon Cost Explorer](#).
- **Amazon Budgets** – Track and take action on your Amazon costs and usage. See [Managing your costs with Amazon Budgets](#).

Performance monitoring with Amazon GameLift Servers and Amazon CloudWatch

Monitor these key metrics to optimize resource utilization. View metrics in the Amazon GameLift Servers console or use Amazon CloudWatch dashboards. For details on all available metrics for Amazon GameLift Servers, see [Monitor Amazon GameLift Servers with Amazon CloudWatch](#).

- **Instance metrics** – ActiveInstances, IdleInstances, PercentIdleInstances
- **Server process metrics** – ActiveServerProcesses, HealthyServerProcesses
- **Game session metrics** – ActiveGameSessions, AvailableGameSessions
- **Player session metrics** – CurrentPlayerSessions
- **Queue metrics** – AverageWaitTime, QueueDepth
- **Matchmaking metrics** – CurrentTickets, MatchesPlaced
- **Hardware performance** – CPUUtilization, NetworkIn/NetworkOut, DiskReadBytes/DiskWriteBytes
- **Resource utilization metrics**
 - Use PercentIdleInstances to determine optimal fleet size.
 - Track PercentAvailableGameSessions to ensure sufficient game hosting capacity.
 - Monitor InstanceInterruptions and GameServerInterruptions to determine Spot Instance viability.

Amazon GameLift Servers reference guides

This section contains reference documentation for using Amazon GameLift Servers.

Topics

- [Service API for Amazon GameLift Servers](#)
- [Server SDK for Amazon GameLift Servers version 4 and earlier](#)
- [Game session placement events](#)
- [Amazon GameLift Servers AMI versions](#)
- [Amazon GameLift Servers service locations](#)
- [Amazon GameLift Servers endpoints and quotas](#)
- [Amazon GameLift Servers API limits](#)
- [Amazon GameLift Servers SDK API limits](#)

Service API for Amazon GameLift Servers

Use this task-based list to find API operations when building your Amazon GameLift Servers game hosting solutions and other features. The Amazon SDK includes these operations in the `aws.gamelift` namespace. [Download the Amazon SDK](#) or [view the Amazon GameLift Servers API reference documentation](#). You can also use the API with the Amazon command line interface (Amazon CLI), as documented in the [Amazon CLI command reference](#).

The API includes two sets of operations for managed game hosting:

- [Manage Amazon GameLift Servers hosting resources](#)
- [Start game sessions and join players](#)

The Amazon GameLift Servers Service API also contains operations for use with other Amazon GameLift Servers tools and solutions. For a list of FleetIQ APIs, see [FleetIQ API operations](#). For a list of FlexMatch APIs for matchmaking, see [FlexMatch API operations](#).

Manage Amazon GameLift Servers hosting resources

Call these operations to configure hosting resources for your game servers, scale capacity to meet player demand, get performance and usage metrics, and more. Use these API operations

when hosting game servers with Amazon GameLift Servers, including Amazon GameLift Servers Realtime. You can also work in [Amazon GameLift Servers console](#) for most resource management tasks, or you can make calls with the Amazon Command Line Interface (Amazon CLI) tool.

Prepare game servers for deployment

Upload and configure your game's game server code in preparation for deployment and launching on hosting resources.

Manage custom game server builds

- [upload-build](#) – Upload build files from a local path and create a new Amazon GameLift Servers build resource. This operation, available as an Amazon CLI command, is the most common way to upload game server builds.
- [CreateBuild](#) – Create a new build using files stored in an Amazon S3 bucket.
- [ListBuilds](#) – Get a list of all builds uploaded to an Amazon GameLift Servers region.
- [DescribeBuild](#) – Retrieve information associated with a build.
- [UpdateBuild](#) – Change build metadata, including build name and version.
- [DeleteBuild](#) – Remove a build from Amazon GameLift Servers.

Manage Amazon GameLift Servers Realtime configuration scripts

- [CreateScript](#) – Upload JavaScript files and create a new Amazon GameLift Servers script resource.
- [ListScripts](#) – Get a list of all Realtime scripts uploaded to an Amazon GameLift Servers region.
- [DescribeScript](#) – Retrieve information associated with a Realtime script.
- [UpdateScript](#) – Change script metadata and upload revised script content.
- [DeleteScript](#) – Remove a Realtime script from Amazon GameLift Servers.

Set up computing resources for hosting

Configure hosting resources and deploy them with your game server build or Realtime configuration script.

Create and manage fleets

- [CreateFleet](#) – Configure and deploy a new Amazon GameLift Servers fleet of computing resources to run your game servers. Once deployed, game servers are automatically launched as configured and ready to host game sessions.
- [ListFleets](#) – Get a list of all fleets in an Amazon GameLift Servers region.
- [DeleteFleet](#) – Remove a fleet that's no longer running game servers or hosting players.
- View / update fleet locations.
 - [CreateFleetLocations](#) – Add remote locations to an existing fleet that supports multiple locations
 - [DescribeFleetLocationAttributes](#) – Get a list of all remote locations for a fleet and view the current status of each location.
 - [DeleteFleetLocations](#) – Remove remote locations from a fleet that supports multiple locations.
- View / update fleet configurations.
 - [DescribeFleetAttributes](#) / [UpdateFleetAttributes](#) – View or change a fleet's metadata and settings for game session protection and resource creation limits.
 - [DescribeFleetPortSettings](#) / [UpdateFleetPortSettings](#) – View or change the inbound permissions (IP address and port setting ranges) allowed for a fleet.
 - [DescribeRuntimeConfiguration](#) / [UpdateRuntimeConfiguration](#) – View or change what server processes (and how many) to run on each instance in a fleet.

Manage fleet capacity

- [DescribeEC2InstanceLimits](#) – Retrieve maximum number of instances allowed for the current Amazon account and the current usage level.
- [DescribeFleetCapacity](#) – Retrieve the current capacity settings for a fleet's home Region.
- [DescribeFleetLocationCapacity](#) – Retrieve the current capacity settings for each location a multi-location fleet.
- [UpdateFleetCapacity](#) – Manually adjust capacity settings for a fleet.
- Set up :
 - [PutScalingPolicy](#) – Turn on target-based auto scaling or create a custom auto scaling policy, or update an existing policy.
 - [DescribeScalingPolicies](#) – Retrieve an existing auto scaling policy.
 - [DeleteScalingPolicy](#) – Delete an auto scaling policy and stop it from affecting a fleet's capacity.
 - [StartFleetActions](#) – Restart a fleet's auto scaling policies.

- [StopFleetActions](#) – Suspend a fleet's auto scaling policies.

Monitor fleet activity.

- [DescribeFleetUtilization](#) – Retrieve statistics on the number of server processes, game sessions, and players that are currently active on a fleet.
- [DescribeFleetLocationUtilization](#) – Retrieve utilization statistics for each location in a multi-location fleet.
- [DescribeFleetEvents](#) – View logged events for a fleet during a specified time span.
- [DescribeGameSessions](#) – Retrieve game session metadata, including a game's running time and current player count.

Set up queues for game session placement

Set up multi-fleet, multi-region queues to place game sessions with the best available hosting resources for cost, latency, and resiliency.

- [CreateGameSessionQueue](#) – Create a queue for use when processing requests for game session placements.
- [DescribeGameSessionQueues](#) – Retrieve game session queues defined in an Amazon GameLift Servers region.
- [UpdateGameSessionQueue](#) – Change the configuration of a game session queue.
- [DeleteGameSessionQueue](#) – Remove a game session queue from the region.

Manage aliases

Use aliases to represent your fleets or create a terminal alternative destination. Aliases are useful when transitioning game activity from one fleet to another, such as during game server build updates.

- [CreateAlias](#) – Define a new alias and optionally assign it to a fleet.
- [ListAliases](#) – Get all fleet aliases defined in an Amazon GameLift Servers region.
- [DescribeAlias](#) – Retrieve information on an existing alias.
- [UpdateAlias](#) – Change settings for an alias, such as redirecting it from one fleet to another.
- [DeleteAlias](#) – Remove an alias from the region.

- [ResolveAlias](#) – Get the fleet ID that a specified alias points to.

Connect to managed hosting instances

View information on individual instances in a fleet, or request remote access to a specified fleet instance for troubleshooting.

- [DescribeInstances](#) – Get information on each instance in a fleet, including instance ID, IP address, location, and status.
- [GetInstanceAccess](#) – Request access credentials needed to remotely connect to a specified instance in a fleet.

Set up VPC peering

Create and manage VPC peering connections between your Amazon GameLift Servers hosting resources and other Amazon resources.

- [CreateVpcPeeringAuthorization](#) – Authorize a peering connection to one of your VPCs.
- [DescribeVpcPeeringAuthorizations](#) – Retrieve valid peering connection authorizations.
- [DeleteVpcPeeringAuthorization](#) – Delete a peering connection authorization.
- [CreateVpcPeeringConnection](#) – Establish a peering connection between the VPC for an Amazon GameLift Servers fleet and one of your VPCs.
- [DescribeVpcPeeringConnections](#) – Retrieve information on active or pending VPC peering connections with an Amazon GameLift Servers fleet.
- [DeleteVpcPeeringConnection](#) – Delete a VPC peering connection with an Amazon GameLift Servers fleet.

Start game sessions and join players

Call these operations from a backend service to start new game sessions, get information on existing game sessions, and join players to game sessions. These operations are for use with custom game servers that are hosted on Amazon GameLift Servers. If you're using Amazon GameLift Servers Realtime, manage game sessions using the [Amazon GameLift Servers Realtime client API \(C#\)](#).

- **Start new game sessions for one or more players.**

- [StartGameSessionPlacement](#) – Ask Amazon GameLift Servers to find the best available hosting resources and start a new game session. This is the preferred method for creating new game sessions. It relies on game session queues to track hosting availability across multiple regions, and uses FleetIQ algorithms to prioritize placements based on player latency, hosting cost, location, etc.
- [DescribeGameSessionPlacement](#) – Get details and status on a placement request.
- [StopGameSessionPlacement](#) – Cancel a placement request.
- [CreateGameSession](#) – Start a new, empty game session on a specific fleet location. This operation gives you greater control over where to start the game session, instead of using FleetIQ to evaluate placement options. You must add players to the new game session in a separate step.
- **Get players into existing game sessions.** Find running game sessions with available player slots and reserve them for new players.
 - [CreatePlayerSession](#) – Reserve an open slot for a player to join a game session.
 - [CreatePlayerSessions](#) – Reserve open slots for multiple players to join a game session.
- **Work with game session and player session data.** Manage information on game sessions and player sessions.
 - [SearchGameSessions](#) – Request a list of active game sessions based on a set of search criteria.
 - [DescribeGameSessions](#) – Retrieve metadata for specific game sessions, including length of time active and current player count.
 - [DescribeGameSessionDetails](#) – Retrieve metadata, including the game session protection setting, for one or more game sessions.
 - [DescribePlayerSessions](#) – Get details on player activity, including status, playing time, and player data.
 - [UpdateGameSession](#) – Change game session settings, such as maximum player count and join policy.
 - [GetGameSessionLogUrl](#) – Get the location of saved logs for a game session.

Server SDK for Amazon GameLift Servers version 4 and earlier

This reference documents the server SDK for Amazon GameLift Servers, version 4.x and earlier. The server SDK provides core functionality that your game servers use to communicate with the Amazon GameLift Servers service. For example, your game server receives prompts from the

service to start and stop game sessions and it provides regular game session status updates to the service. Integrate your game servers with the server SDK before you deploy them for hosting.

Use this server SDK reference to integrate your custom multiplayer game servers for hosting with Amazon GameLift Servers. For guidance about the integration process, see [Add Amazon GameLift Servers to your game server with the server SDK](#).

For SDK version 4.0.2, you can download it from the [official GitHub releases](#). The complete SDK package including GameLiftLocal.jar is available in the GameLift-CSharp-ServerSDK-4.0.2.zip artifact.

- Amazon GameLift Servers Anywhere
- Amazon GameLift Servers managed containers
- Amazon GameLift Servers plugin for Unreal Engine and Unity

Topics

- [C++ server SDK for Amazon GameLift Servers 4.x -- Actions](#)
- [C# server SDK for Amazon GameLift Servers 4.x -- Actions](#)
- [Server SDK \(Unreal\) for Amazon GameLift Servers -- Actions](#)

C++ server SDK for Amazon GameLift Servers 4.x -- Actions

Use the server SDK reference to integrate your multiplayer game for hosting with Amazon GameLift Servers. For guidance about the integration process, see [Add Amazon GameLift Servers to your game server with the server SDK](#).

Note

This reference is for an earlier version of the server SDK for Amazon GameLift Servers. For the latest version, see [???](#).

C++ server SDK for Amazon GameLift Servers 4.x -- Data types

Use the server SDK reference to integrate your multiplayer game for hosting with Amazon GameLift Servers. For guidance about the integration process, see [Add Amazon GameLift Servers to your game server with the server SDK](#).

Note

This reference is for an earlier version of the server SDK for Amazon GameLift Servers. For the latest version, see [???](#).

This API is defined in `GameLiftServerAPI.h`, `LogParameters.h`, and `ProcessParameters.h`.

[C++ server SDK for Amazon GameLift Servers 4.x -- Actions](#)

DescribePlayerSessionsRequest

This data type is used to specify which player session(s) to retrieve. You can use it as follows:

- Provide a `PlayerSessionId` to request a specific player session.
- Provide a `GameSessionId` to request all player sessions in the specified game session.
- Provide a `PlayerId` to request all player sessions for the specified player.

For large collections of player sessions, use the pagination parameters to retrieve results in sequential blocks.

Contents

GameSessionId

Unique game session identifier. Use this parameter to request all player sessions for the specified game session. Game session ID format is as follows: `arn:aws-cn:gamelift:<region>::gamesession/fleet-<fleet ID>/<ID string>`. The value of `<ID string>` is either a custom ID string or (if one was specified when the game session was created) a generated string.

Type: String

Required: No

Limit

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

Required: No

NextToken

Token indicating the start of the next sequential page of results. Use the token that is returned with a previous call to this action. To specify the start of the result set, do not specify a value. If a player session ID is specified, this parameter is ignored.

Type: String

Required: No

PlayerId

Unique identifier for a player. Player IDs are defined by the developer. See [Generate player IDs](#).

Type: String

Required: No

PlayerSessionId

Unique identifier for a player session.

Type: String

Required: No

PlayerSessionStatusFilter

Player session status to filter results on. 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 time-out limit (60 seconds).

Type: String

Required: No

LogParameters

This data type is used to identify which files generated during a game session that you want Amazon GameLift Servers to upload and store once the game session ends. This information is communicated to the Amazon GameLift Servers service in a [ProcessReady\(\)](#) call.

Contents

logPaths

Directory paths to game server log files that you want Amazon GameLift Servers to store for future access. These files are generated during each game session. File paths and names are defined in your game server and stored in the root game build directory. The log paths must be absolute. For example, if your game build stores game session logs in a path like MyGame\sessionlogs\, then the log path would be c:\game\MyGame\sessionLogs (on a Windows instance) or /local/game/MyGame/sessionLogs (on a Linux instance).

Type: std::vector<std::string>

Required: No

ProcessParameters

This data type contains the set of parameters sent to the Amazon GameLift Servers service in a [ProcessReady\(\)](#) call.

Contents

port

Port number the server process listens on for new player connections. The value must fall into the port range configured for any fleet deploying this game server build. This port number is included in game session and player session objects, which game sessions use when connecting to a server process.

Type: Integer

Required: Yes

logParameters

Object with a list of directory paths to game session log files.

Type: `Aws::GameLift::Server::LogParameters`

Required: No

onStartGameSession

Name of callback function that the Amazon GameLift Servers service invokes to activate a new game session. Amazon GameLift Servers calls this function in response to the client request [CreateGameSession](#). The callback function passes a [GameSession](#) object (defined in the *Amazon GameLift Servers Service API Reference*).

Type: `const std::function<void(Aws::GameLift::Model::GameSession)>
onStartGameSession`

Required: Yes

onProcessTerminate

Name of callback function that the Amazon GameLift Servers service invokes to force the server process to shut down. After calling this function, Amazon GameLift Servers waits five minutes for the server process to shut down and respond with a [ProcessEnding\(\)](#) call. If no response is received, it shuts down the server process.

Type: `std::function<void()> onProcessTerminate`

Required: No

onHealthCheck

Name of callback function that the Amazon GameLift Servers service invokes to request a health status report from the server process. Amazon GameLift Servers calls this function every 60 seconds. After calling this function Amazon GameLift Servers waits 60 seconds for a response, and if none is received, records the server process as unhealthy.

Type: `std::function<bool()> onHealthCheck`

Required: No

onUpdateGameSession

Name of callback function that the Amazon GameLift Servers service invokes to pass an updated game session object to the server process. Amazon GameLift Servers calls this function when a [match backfill](#) request has been processed in order to provide updated matchmaker

data. It passes a [GameSession](#) object, a status update (updateReason), and the match backfill ticket ID.

Type: `std::function<void(Aws::GameLift::Server::Model::UpdateGameSession)>`
`onUpdateGameSession`

Required: No

StartMatchBackfillRequest

This data type is used to send a matchmaking backfill request. The information is communicated to the Amazon GameLift Servers service in a [StartMatchBackfill\(\)](#) call.

Contents

GameSessionArn

Unique game session identifier. The API action [GetGameSessionId\(\)](#) returns the identifier in ARN format.

Type: String

Required: Yes

MatchmakingConfigurationArn

Unique identifier, in the form of an ARN, for the matchmaker to use for this request. To find the matchmaker that was used to create the original game session, look in the game session object, in the matchmaker data property. Learn more about matchmaker data in [Word with matchmaker data](#).

Type: String

Required: Yes

Players

A set of data representing all players who are currently in the game session. The matchmaker uses this information to search for new players who are good matches for the current players. See the *Amazon GameLift Servers API Reference Guide* for a description of the Player object format. To find player attributes, IDs, and team assignments, look in the game session object,

in the matchmaker data property. If latency is used by the matchmaker, gather updated latency for the current region and include it in each player's data.

Type: `std::vector`[<player>](#)

Required: Yes

TicketId

Unique identifier for a matchmaking or match backfill request ticket. If no value is provided here, Amazon GameLift Servers will generate one in the form of a UUID. Use this identifier to track the match backfill ticket status or cancel the request if needed.

Type: String

Required: No

StopMatchBackfillRequest

This data type is used to cancel a matchmaking backfill request. The information is communicated to the Amazon GameLift Servers service in a [StopMatchBackfill\(\)](#) call.

Contents

GameSessionArn

Unique game session identifier associated with the request being canceled.

Type: String

Required: Yes

MatchmakingConfigurationArn

Unique identifier of the matchmaker this request was sent to.

Type: String

Required: Yes

TicketId

Unique identifier of the backfill request ticket to be canceled.

Type: String

Required: Yes

[C++ server SDK for Amazon GameLift Servers 4.x -- Data types](#)

Topics

- [AcceptPlayerSession\(\)](#)
- [ActivateGameSession\(\)](#)
- [DescribePlayerSessions\(\)](#)
- [GetGameSessionId\(\)](#)
- [GetSdkVersion\(\)](#)
- [GetTerminationTime\(\)](#)
- [InitSDK\(\)](#)
- [ProcessEnding\(\)](#)
- [ProcessReady\(\)](#)
- [ProcessReadyAsync\(\)](#)
- [RemovePlayerSession\(\)](#)
- [StartMatchBackfill\(\)](#)
- [StopMatchBackfill\(\)](#)
- [TerminateGameSession\(\)](#)
- [UpdatePlayerSessionCreationPolicy\(\)](#)
- [Destroy\(\)](#)

AcceptPlayerSession()

Notifies the Amazon GameLift Servers service that a player with the specified player session ID has connected to the server process and needs validation. Amazon GameLift Servers verifies that the player session ID is valid—that is, that the player ID has reserved a player slot in the game session. Once validated, Amazon GameLift Servers changes the status of the player slot from RESERVED to ACTIVE.

Syntax

```
GenericOutcome AcceptPlayerSession(const std::string& playerSessionId);
```

Parameters

playerSessionId

Unique ID issued by the Amazon GameLift Servers service in response to a call to the Amazon SDK Amazon GameLift Servers API action [CreatePlayerSession](#). The game client references this ID when connecting to the server process.

Type: std::string

Required: Yes

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

This example illustrates a function for handling a connection request, including validating and rejecting invalid player session IDs.

```
void ReceiveConnectingPlayerSessionID (Connection& connection, const std::string&
playerSessionId){
    Aws::GameLift::GenericOutcome connectOutcome =
        Aws::GameLift::Server::AcceptPlayerSession(playerSessionId);
    if(connectOutcome.IsSuccess())
    {
        connectionToSessionMap.emplace(connection, playerSessionId);
        connection.Accept();
    }
    else
    {
        connection.Reject(connectOutcome.GetError().GetMessage());
    }
}
```

ActivateGameSession()

Notifies the Amazon GameLift Servers service that the server process has started a game session and is now ready to receive player connections. This action should be called as part of the `onStartGameSession()` callback function, after all game session initialization has been completed.

Syntax

```
GenericOutcome ActivateGameSession();
```

Parameters

This action has no parameters.

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

This example shows `ActivateGameSession()` being called as part of the `onStartGameSession()` callback function.

```
void onStartGameSession(Aws::GameLift::Model::GameSession myGameSession)
{
    // game-specific tasks when starting a new game session, such as loading map
    GenericOutcome outcome = Aws::GameLift::Server::ActivateGameSession();
}
```

DescribePlayerSessions()

Retrieves player session data, including settings, session metadata, and player data. Use this action to get information for a single player session, for all player sessions in a game session, or for all player sessions associated with a single player ID.

Syntax

```
DescribePlayerSessionsOutcome DescribePlayerSessions (
    const Aws::GameLift::Server::Model::DescribePlayerSessionsRequest
    &describePlayerSessionsRequest);
```

Parameters

describePlayerSessionsRequest

A [DescribePlayerSessionsRequest](#) object describing which player sessions to retrieve.

Required: Yes

Return value

If successful, returns a `DescribePlayerSessionsOutcome` object containing a set of player session objects that fit the request parameters. Player session objects have a structure identical to the Amazon SDK Amazon GameLift Servers API [PlayerSession](#) data type.

Example

This example illustrates a request for all player sessions actively connected to a specified game session. By omitting `NextToken` and setting the `Limit` value to 10, Amazon GameLift Servers returns the first 10 player sessions records matching the request.

```
// Set request parameters
Aws::GameLift::Server::Model::DescribePlayerSessionsRequest request;
request.SetPlayerSessionStatusFilter(Aws::GameLift::Server::Model::PlayerSessionStatusMapper::G
request.SetLimit(10);
request.SetGameSessionId("the game session ID");    // can use GetGameSessionId()

// Call DescribePlayerSessions
Aws::GameLift::DescribePlayerSessionsOutcome playerSessionsOutcome =
    Aws::GameLift::Server::DescribePlayerSessions(request);
```

GetGameSessionId()

Retrieves a unique identifier for the game session currently being hosted by the server process, if the server process is active. The identifier is returned in ARN format: `arn:aws-cn:gamelift:<region>::gamesession/fleet-<fleet ID>/<ID string>`.

For idle process that are not yet activated with a game session, the call returns `Success=True` and `GameSessionId=""` (an empty string).

Syntax

```
AwsStringOutcome GetGameSessionId();
```

Parameters

This action has no parameters.

Return value

If successful, returns the game session ID as an `AwsStringOutcome` object. If not successful, returns an error message.

Example

```
Aws::GameLift::AwsStringOutcome sessionIdOutcome =  
    Aws::GameLift::Server::GetGameSessionId();
```

GetSdkVersion()

Returns the current version number of the SDK in use.

Syntax

```
AwsStringOutcome GetSdkVersion();
```

Parameters

This action has no parameters.

Return value

If successful, returns the current SDK version as an `AwsStringOutcome` object. The returned string includes the version number only (ex. "3.1.5"). If not successful, returns an error message.

Example

```
Aws::GameLift::AwsStringOutcome sdkVersionOutcome =  
    Aws::GameLift::Server::GetSdkVersion();
```

GetTerminationTime()

Returns the time that a server process is scheduled to be shut down, if a termination time is available. A server process takes this action after receiving an `onProcessTerminate()` callback from the Amazon GameLift Servers service. Amazon GameLift Servers may call `onProcessTerminate()` for the following reasons: (1) when the server process has reported poor health or has not responded to Amazon GameLift Servers, (2) when terminating the instance during a scale-down event, or (3) when an instance is being terminated due to a [Spot interruption](#).

If the process has received an `onProcessTerminate()` callback, the value returned is the estimated termination time. If the process has not received an `onProcessTerminate()` callback, an error message is returned. Learn more about [shutting down a server process](#).

Syntax

```
AwsLongOutcome GetTerminationTime();
```

Parameters

This action has no parameters.

Return value

If successful, returns the termination time as an `AwsLongOutcome` object. The value is the termination time, expressed in elapsed ticks since 0001 00:00:00. For example, the date time value 2020-09-13 12:26:40 -000Z is equal to 637355968000000000 ticks. If no termination time is available, returns an error message.

Example

```
Aws::GameLift::AwsLongOutcome TermTimeOutcome =  
    Aws::GameLift::Server::GetTerminationTime();
```

InitSDK()

Initializes the Amazon GameLift Servers SDK. This method should be called on launch, before any other Amazon GameLift Servers-related initialization occurs.

Syntax

```
InitSDKOutcome InitSDK();
```

Parameters

This action has no parameters.

Return value

If successful, returns an `InitSdkOutcome` object indicating that the server process is ready to call [ProcessReady\(\)](#).

Example

```
Aws::GameLift::Server::InitSDKOutcome initOutcome =  
    Aws::GameLift::Server::InitSDK();
```

ProcessEnding()

Notifies the Amazon GameLift Servers service that the server process is shutting down. This method should be called after all other cleanup tasks, including shutting down all active game sessions. This method should exit with an exit code of 0; a non-zero exit code results in an event message that the process did not exit cleanly.

Once the method exits with a code of 0, you can terminate the process with a successful exit code. You can also exit the process with an error code. If you exit with an error code, the fleet event will indicate the process terminated abnormally (SERVER_PROCESS_TERMINATED_UNHEALTHY).

Syntax

```
GenericOutcome ProcessEnding();
```

Parameters

This action has no parameters.

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

```
Aws::GameLift::GenericOutcome outcome = Aws::GameLift::Server::ProcessEnding();  
if (outcome.Success)  
    exit(0); // exit with success  
// otherwise, exit with error code  
exit(errorCode);
```

ProcessReady()

Notifies the Amazon GameLift Servers service that the server process is ready to host game sessions. Call this method after successfully invoking [InitSDK\(\)](#) and completing setup tasks that

are required before the server process can host a game session. This method should be called only once per process.

This call is synchronous. To make an asynchronous call, use [ProcessReadyAsync\(\)](#). See [Initialize the server process](#) for more details.

Syntax

```
GenericOutcome ProcessReady(  
    const Aws::GameLift::Server::ProcessParameters &processParameters);
```

Parameters

processParameters

A [ProcessParameters](#) object communicating the following information about the server process:

- Names of callback methods, implemented in the game server code, that the Amazon GameLift Servers service invokes to communicate with the server process.
- Port number that the server process is listening on.
- Path to any game session-specific files that you want Amazon GameLift Servers to capture and store.

Required: Yes

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

This example illustrates both the [ProcessReady\(\)](#) call and callback function implementations.

```
// Set parameters and call ProcessReady  
std::string serverLog("serverOut.log");           // Example of a log file written by the  
game server  
std::vector<std::string> logPaths;  
logPaths.push_back(serverLog);  
  
int listenPort = 9339;
```

```

Aws::GameLift::Server::ProcessParameters processReadyParameter =
    Aws::GameLift::Server::ProcessParameters(
        std::bind(&Server::onStartGameSession, this, std::placeholders::_1),
        std::bind(&Server::onProcessTerminate, this),
        std::bind(&Server::OnHealthCheck, this),
        std::bind(&Server::OnUpdateGameSession, this),
        listenPort,
        Aws::GameLift::Server::LogParameters(logPaths));

Aws::GameLift::GenericOutcome outcome =
    Aws::GameLift::Server::ProcessReady(processReadyParameter);

// Implement callback functions
void Server::onStartGameSession(Aws::GameLift::Model::GameSession myGameSession)
{
    // game-specific tasks when starting a new game session, such as loading map
    GenericOutcome outcome =
        Aws::GameLift::Server::ActivateGameSession (maxPlayers);
}

void Server::onProcessTerminate()
{
    // game-specific tasks required to gracefully shut down a game session,
    // such as notifying players, preserving game state data, and other cleanup
    GenericOutcome outcome = Aws::GameLift::Server::ProcessEnding();
}

bool Server::onHealthCheck()
{
    bool health;
    // complete health evaluation within 60 seconds and set health
    return health;
}

```

ProcessReadyAsync()

Notifies the Amazon GameLift Servers service that the server process is ready to host game sessions. This method should be called once the server process is ready to host a game session. The parameters specify the names of callback functions for Amazon GameLift Servers to call in certain circumstances. Game server code must implement these functions.

This call is asynchronous. To make a synchronous call, use [ProcessReady\(\)](#). See [Initialize the server process](#) for more details.

Syntax

```
GenericOutcomeCallable ProcessReadyAsync(  
    const Aws::GameLift::Server::ProcessParameters &processParameters);
```

Parameters

processParameters

A [ProcessParameters](#) object communicating the following information about the server process:

- Names of callback methods, implemented in the game server code, that the Amazon GameLift Servers service invokes to communicate with the server process.
- Port number that the server process is listening on.
- Path to any game session-specific files that you want Amazon GameLift Servers to capture and store.

Required: Yes

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

```
// Set parameters and call ProcessReady  
std::string serverLog("serverOut.log");           // This is an example of a log file  
                                                  written by the game server  
std::vector<std::string> logPaths;  
logPaths.push_back(serverLog);  
  
int listenPort = 9339;  
  
Aws::GameLift::Server::ProcessParameters processReadyParameter =  
    Aws::GameLift::Server::ProcessParameters(  
        std::bind(&Server::onStartGameSession, this, std::placeholders::_1),  
        std::bind(&Server::onProcessTerminate, this),  
        std::bind(&Server::OnHealthCheck, this),  
        std::bind(&Server::OnUpdateGameSession, this),  
        listenPort,  
        Aws::GameLift::Server::LogParameters(logPaths));
```

```
Aws::GameLift::GenericOutcomeCallable outcome =
    Aws::GameLift::Server::ProcessReadyAsync(processReadyParameter);

// Implement callback functions
void onStartGameSession(Aws::GameLift::Model::GameSession myGameSession)
{
    // game-specific tasks when starting a new game session, such as loading map
    GenericOutcome outcome = Aws::GameLift::Server::ActivateGameSession (maxPlayers);
}

void onProcessTerminate()
{
    // game-specific tasks required to gracefully shut down a game session,
    // such as notifying players, preserving game state data, and other cleanup
    GenericOutcome outcome = Aws::GameLift::Server::ProcessEnding();
}

bool onHealthCheck()
{
    // perform health evaluation and complete within 60 seconds
    return health;
}
```

RemovePlayerSession()

Notifies the Amazon GameLift Servers service that a player with the specified player session ID has disconnected from the server process. In response, Amazon GameLift Servers changes the player slot to available, which allows it to be assigned to a new player.

Syntax

```
GenericOutcome RemovePlayerSession(
    const std::string& playerSessionId);
```

Parameters

playerSessionId

Unique ID issued by the Amazon GameLift Servers service in response to a call to the Amazon SDK Amazon GameLift Servers API action [CreatePlayerSession](#). The game client references this ID when connecting to the server process.

Type: std::string

Required: Yes

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

```
Aws::GameLift::GenericOutcome disconnectOutcome =  
    Aws::GameLift::Server::RemovePlayerSession(playerSessionId);
```

StartMatchBackfill()

Sends a request to find new players for open slots in a game session created with FlexMatch. See also the Amazon SDK action [StartMatchBackfill\(\)](#). With this action, match backfill requests can be initiated by a game server process that is hosting the game session. Learn more about the [FlexMatch backfill feature](#).

This action is asynchronous. If new players are successfully matched, the Amazon GameLift Servers service delivers updated matchmaker data by invoking the callback function `OnUpdateGameSession()`.

A server process can have only one active match backfill request at a time. To send a new request, first call [StopMatchBackfill\(\)](#) to cancel the original request.

Syntax

```
StartMatchBackfillOutcome StartMatchBackfill (  
    const Aws::GameLift::Server::Model::StartMatchBackfillRequest  
    &startBackfillRequest);
```

Parameters

StartMatchBackfillRequest

A [StartMatchBackfillRequest](#) object that communicates the following information:

- A ticket ID to assign to the backfill request. This information is optional; if no ID is provided, Amazon GameLift Servers will autogenerate one.
- The matchmaker to send the request to. The full configuration ARN is required. This value can be acquired from the game session's matchmaker data.

- The ID of the game session that is being backfilled.
- Available matchmaking data for the game session's current players.

Required: Yes

Return value

Returns a `StartMatchBackfillOutcome` object with the match backfill ticket or failure with an error message. Ticket status can be tracked using the Amazon SDK action [DescribeMatchmaking\(\)](#).

Example

```
// Build a backfill request
std::vector<Player> players;
Aws::GameLift::Server::Model::StartMatchBackfillRequest startBackfillRequest;
startBackfillRequest.SetTicketId("a ticket ID");
    //optional, autogenerated if not provided
startBackfillRequest.SetMatchmakingConfigurationArn("the matchmaker configuration
    ARN"); //from the game session matchmaker data
startBackfillRequest.SetGameSessionArn("the game session ARN");
    // can use GetGameSessionId()
startBackfillRequest.SetPlayers(players);
    //from the game session matchmaker data

// Send backfill request
Aws::GameLift::StartMatchBackfillOutcome backfillOutcome =
    Aws::GameLift::Server::StartMatchBackfill(startBackfillRequest);

// Implement callback function for backfill
void Server::OnUpdateGameSession(Aws::GameLift::Server::Model::GameSession gameSession,
    Aws::GameLift::Server::Model::UpdateReason updateReason, std::string backfillTicketId)
{
    // handle status messages
    // perform game-specific tasks to prep for newly matched players
}
```

StopMatchBackfill()

Cancels an active match backfill request that was created with [StartMatchBackfill\(\)](#). See also the Amazon SDK action [StopMatchmaking\(\)](#). Learn more about the [FlexMatch backfill feature](#).

Syntax

```
GenericOutcome StopMatchBackfill (  
    const Aws::GameLift::Server::Model::StopMatchBackfillRequest &stopBackfillRequest);
```

Parameters

StopMatchBackfillRequest

A [StopMatchBackfillRequest](#) object identifying the matchmaking ticket to cancel:

- ticket ID assigned to the backfill request being canceled
- matchmaker the backfill request was sent to
- game session associated with the backfill request

Required: Yes

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

```
// Set backfill stop request parameters  
  
Aws::GameLift::Server::Model::StopMatchBackfillRequest stopBackfillRequest;  
stopBackfillRequest.SetTicketId("the ticket ID");  
stopBackfillRequest.SetGameSessionArn("the game session ARN");  
    // can use GetGameSessionId()  
stopBackfillRequest.SetMatchmakingConfigurationArn("the matchmaker configuration ARN");  
    // from the game session matchmaker data  
  
Aws::GameLift::GenericOutcome stopBackfillOutcome =  
    Aws::GameLift::Server::StopMatchBackfillRequest(stopBackfillRequest);
```

TerminateGameSession()

This method is deprecated with version 4.0.1. Instead, the server process should call [ProcessEnding\(\)](#) after a game session has ended.

Notifies the Amazon GameLift Servers service that the server process has ended the current game session. This action is called when the server process will remain active and ready to host a new

game session. It should be called only after your game session termination procedure is complete, because it signals to Amazon GameLift Servers that the server process is immediately available to host a new game session.

This action is not called if the server process will be shut down after the game session stops. Instead, call [ProcessEnding\(\)](#) to signal that both the game session and the server process are ending.

Syntax

```
GenericOutcome TerminateGameSession();
```

Parameters

This action has no parameters.

Return value

Returns a generic outcome consisting of success or failure with an error message.

UpdatePlayerSessionCreationPolicy()

Updates the current game session's ability to accept new player sessions. A game session can be set to either accept or deny all new player sessions. See also the Amazon SDK action [UpdateGameSession\(\)](#).

Syntax

```
GenericOutcome UpdatePlayerSessionCreationPolicy(  
    Aws::GameLift::Model::PlayerSessionCreationPolicy newPlayerSessionPolicy);
```

Parameters

newPlayerSessionPolicy

String value indicating whether the game session accepts new players.

Type: `Aws::GameLift::Model::PlayerSessionCreationPolicy` enum. Valid values include:

- **ACCEPT_ALL** – Accept all new player sessions.

- **DENY_ALL** – Deny all new player sessions.

Required: Yes

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

This example sets the current game session's join policy to accept all players.

```
Aws::GameLift::GenericOutcome outcome =  
    Aws::GameLift::Server::UpdatePlayerSessionCreationPolicy(Aws::GameLift::Model::PlayerSessionCr
```

Destroy()

Cleans up memory allocated by `initSDK()` during game server initialization. Use this method after you end a game server process to avoid wasting server memory.

Syntax

```
GenericOutcome Aws::GameLift::Server::Destroy();
```

Parameters

There are no parameters.

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

This example cleans up the memory allocated by `initSDK` after a game server process has ended.

```
if (Aws::GameLift::Server::ProcessEnding().IsSuccess()) {  
    Aws::GameLift::Server::Destroy();  
    exit(0);  
}
```

C# server SDK for Amazon GameLift Servers 4.x -- Actions

Use the server SDK reference to integrate your multiplayer game for hosting with Amazon GameLift Servers. For guidance about the integration process, see [Add Amazon GameLift Servers to your game server with the server SDK](#).

Note

This reference is for an earlier version of the server SDK for Amazon GameLift Servers. For the latest version, see [???](#).

C# server SDK for Amazon GameLift Servers 4.x -- Data types

Use the server SDK reference to integrate your multiplayer game for hosting with Amazon GameLift Servers. For guidance about the integration process, see [Add Amazon GameLift Servers to your game server with the server SDK](#).

Note

This reference is for an earlier version of the server SDK for Amazon GameLift Servers. For the latest version, see [???](#).

[C# server SDK for Amazon GameLift Servers 4.x -- Actions](#)

Topics

- [LogParameters](#)
- [DescribePlayerSessionsRequest](#)
- [ProcessParameters](#)
- [StopMatchBackfillRequest](#)

LogParameters

This data type is used to identify which files generated during a game session that you want Amazon GameLift Servers to upload and store once the game session ends. This information is communicated to the Amazon GameLift Servers service in a [ProcessReady\(\)](#) call.

Contents

logPaths

List of directory paths to game server log files you want Amazon GameLift Servers to store for future access. These files are generated by a server process during each game session; file paths and names are defined in your game server and stored in the root game build directory. The log paths must be absolute. For example, if your game build stores game session logs in a path like `MyGame\sessionlogs\`, then the log path would be `c:\game\MyGame\sessionLogs` (on a Windows instance) or `/local/game/MyGame/sessionLogs` (on a Linux instance).

Type: `List<String>`

Required: No

DescribePlayerSessionsRequest

This data type is used to specify which player session(s) to retrieve. It can be used in several ways: (1) provide a `PlayerSessionId` to request a specific player session; (2) provide a `GameSessionId` to request all player sessions in the specified game session; or (3) provide a `PlayerId` to request all player sessions for the specified player. For large collections of player sessions, use the pagination parameters to retrieve results as sequential pages.

Contents

GameSessionId

Unique game session identifier. Use this parameter to request all player sessions for the specified game session. Game session ID format is as follows: `arn:aws-cn:gamelift:<region>::gamesession/fleet-<fleet ID>/<ID string>`. The value of `<ID string>` is either a custom ID string (if one was specified when the game session was created) or a generated string.

Type: `String`

Required: No

Limit

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

Required: No

NextToken

Token indicating the start of the next sequential page of results. Use the token that is returned with a previous call to this action. To specify the start of the result set, do not specify a value. If a player session ID is specified, this parameter is ignored.

Type: String

Required: No

PlayerId

Unique identifier for a player. Player IDs are defined by the developer. See [Generate player IDs](#).

Type: String

Required: No

PlayerSessionId

Unique identifier for a player session.

Type: String

Required: No

PlayerSessionStatusFilter

Player session status to filter results on. 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 time-out limit (60 seconds).

Type: String

Required: No

ProcessParameters

This data type contains the set of parameters sent to the Amazon GameLift Servers service in a [ProcessReady\(\)](#) call.

Contents

port

Port number the server process will listen on for new player connections. The value must fall into the port range configured for any fleet deploying this game server build. This port number is included in game session and player session objects, which game sessions use when connecting to a server process.

Type: Integer

Required: Yes

logParameters

Object with a list of directory paths to game session log files.

Type: `Aws::GameLift::Server::LogParameters`

Required: Yes

onStartGameSession

Name of callback function that the Amazon GameLift Servers service invokes to activate a new game session. Amazon GameLift Servers calls this function in response to the client request [CreateGameSession](#). The callback function takes a [GameSession](#) object (defined in the *Amazon GameLift Servers Service API Reference*).

Type: `void OnStartGameSessionDelegate(GameSession gameSession)`

Required: Yes

onProcessTerminate

Name of callback function that the Amazon GameLift Servers service invokes to force the server process to shut down. After calling this function, Amazon GameLift Servers waits five minutes

for the server process to shut down and respond with a [ProcessEnding\(\)](#) call before it shuts down the server process.

Type: void OnProcessTerminateDelegate()

Required: Yes

onHealthCheck

Name of callback function that the Amazon GameLift Servers service invokes to request a health status report from the server process. Amazon GameLift Servers calls this function every 60 seconds. After calling this function Amazon GameLift Servers waits 60 seconds for a response, and if none is received, records the server process as unhealthy.

Type: bool OnHealthCheckDelegate()

Required: Yes

onUpdateGameSession

Name of callback function that the Amazon GameLift Servers service invokes to pass an updated game session object to the server process. Amazon GameLift Servers calls this function when a [match backfill](#) request has been processed in order to provide updated matchmaker data. It passes a [GameSession](#) object, a status update (updateReason), and the match backfill ticket ID.

Type: void OnUpdateGameSessionDelegate (UpdateGameSession
updateGameSession)

Required: No

StartMatchBackfillRequest

This data type is used to send a matchmaking backfill request. The information is communicated to the Amazon GameLift Servers service in a [StartMatchBackfill\(\)](#) call.

Contents

GameSessionArn

Unique game session identifier. The SDK method [GetGameSessionId\(\)](#) returns the identifier in ARN format.

Type: String

Required: Yes

MatchmakingConfigurationArn

Unique identifier, in the form of an ARN, for the matchmaker to use for this request. To find the matchmaker that was used to create the original game session, look in the game session object, in the matchmaker data property. Learn more about matchmaker data in [Work with matchmaker data](#).

Type: String

Required: Yes

Players

A set of data representing all players who are currently in the game session. The matchmaker uses this information to search for new players who are good matches for the current players. See the *Amazon GameLift Servers API Reference Guide* for a description of the Player object format. To find player attributes, IDs, and team assignments, look in the game session object, in the matchmaker data property. If latency is used by the matchmaker, gather updated latency for the current region and include it in each player's data.

Type: [Player](#)[]

Required: Yes

TicketId

Unique identifier for a matchmaking or match backfill request ticket. If no value is provided here, Amazon GameLift Servers will generate one in the form of a UUID. Use this identifier to track the match backfill ticket status or cancel the request if needed.

Type: String

Required: No

StopMatchBackfillRequest

This data type is used to cancel a matchmaking backfill request. The information is communicated to the Amazon GameLift Servers service in a [StopMatchBackfill\(\)](#) call.

Contents

GameSessionArn

Unique game session identifier associated with the request being canceled.

Type: String

Required: Yes

MatchmakingConfigurationArn

Unique identifier of the matchmaker this request was sent to.

Type: String

Required: Yes

TicketId

Unique identifier of the backfill request ticket to be canceled.

Type: String

Required: Yes

[C# server SDK for Amazon GameLift Servers 4.x -- Data types](#)

Topics

- [AcceptPlayerSession\(\)](#)
- [ActivateGameSession\(\)](#)
- [DescribePlayerSessions\(\)](#)
- [GetGameSessionId\(\)](#)
- [GetSdkVersion\(\)](#)
- [GetTerminationTime\(\)](#)
- [InitSDK\(\)](#)
- [ProcessEnding\(\)](#)
- [ProcessReady\(\)](#)
- [RemovePlayerSession\(\)](#)

- [StartMatchBackfill\(\)](#)
- [StopMatchBackfill\(\)](#)
- [TerminateGameSession\(\)](#)
- [UpdatePlayerSessionCreationPolicy\(\)](#)

AcceptPlayerSession()

Notifies the Amazon GameLift Servers service that a player with the specified player session ID has connected to the server process and needs validation. Amazon GameLift Servers verifies that the player session ID is valid—that is, that the player ID has reserved a player slot in the game session. Once validated, Amazon GameLift Servers changes the status of the player slot from RESERVED to ACTIVE.

Syntax

```
GenericOutcome AcceptPlayerSession(String playerId)
```

Parameters

playerSessionId

Unique ID issued by Amazon GameLift Servers when a new player session is created. A player session ID is specified in a `PlayerSession` object, which is returned in response to a client call to the *GameLift API* actions [StartGameSessionPlacement](#), [CreateGameSession](#), [DescribeGameSessionPlacement](#), or [DescribePlayerSessions](#).

Type: String

Required: Yes

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

This example illustrates a function for handling a connection request, including validating and rejecting invalid player session IDs.

```
void ReceiveConnectingPlayerSessionID (Connection connection, String playerId){
    var acceptPlayerSessionOutcome =
    GameLiftServerAPI.AcceptPlayerSession(playerSessionId);
    if(acceptPlayerSessionOutcome.Success)
    {
        connectionToSessionMap.emplace(connection, playerId);
        connection.Accept();
    }
    else
    {
        connection.Reject(acceptPlayerSessionOutcome.Error.ErrorMessage);    }
}
```

ActivateGameSession()

Notifies the Amazon GameLift Servers service that the server process has activated a game session and is now ready to receive player connections. This action should be called as part of the `onStartGameSession()` callback function, after all game session initialization has been completed.

Syntax

```
GenericOutcome ActivateGameSession()
```

Parameters

This action has no parameters.

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

This example shows `ActivateGameSession()` being called as part of the `onStartGameSession()` delegate function.

```
void OnStartGameSession(GameSession gameSession)
{
    // game-specific tasks when starting a new game session, such as loading map
```

```
// When ready to receive players
var activateGameSessionOutcome = GameLiftServerAPI.ActivateGameSession();
}
```

DescribePlayerSessions()

Retrieves player session data, including settings, session metadata, and player data. Use this action to get information for a single player session, for all player sessions in a game session, or for all player sessions associated with a single player ID.

Syntax

```
DescribePlayerSessionsOutcome DescribePlayerSessions(DescribePlayerSessionsRequest
describePlayerSessionsRequest)
```

Parameters

describePlayerSessionsRequest

A [DescribePlayerSessionsRequest](#) object describing which player sessions to retrieve.

Required: Yes

Return value

If successful, returns a `DescribePlayerSessionsOutcome` object containing a set of player session objects that fit the request parameters. Player session objects have a structure identical to the Amazon SDK Amazon GameLift Servers API [PlayerSession](#) data type.

Example

This example illustrates a request for all player sessions actively connected to a specified game session. By omitting *NextToken* and setting the *Limit* value to 10, Amazon GameLift Servers will return the first 10 player sessions records matching the request.

```
// Set request parameters
var describePlayerSessionsRequest = new
    Aws.GameLift.Server.Model.DescribePlayerSessionsRequest()
{
    GameSessionId = GameLiftServerAPI.GetGameSessionId().Result, //gets the ID for
    the current game session
```

```
    Limit = 10,  
    PlayerSessionStatusFilter =  
    PlayerSessionStatusMapper.GetNameForPlayerSessionStatus(PlayerSessionStatus.ACTIVE)  
};  
// Call DescribePlayerSessions  
Aws::GameLift::DescribePlayerSessionsOutcome playerSessionsOutcome =  
    Aws::GameLift::Server::Model::DescribePlayerSessions(describePlayerSessionRequest);
```

GetGameSessionId()

Retrieves the ID of the game session currently being hosted by the server process, if the server process is active.

For idle process that are not yet activated with a game session, the call returns `Success=True` and `GameSessionId=""` (an empty string).

Syntax

```
AwsStringOutcome GetGameSessionId()
```

Parameters

This action has no parameters.

Return value

If successful, returns the game session ID as an `AwsStringOutcome` object. If not successful, returns an error message.

Example

```
var getGameSessionIdOutcome = GameLiftServerAPI.GetGameSessionId();
```

GetSdkVersion()

Returns the current version number of the SDK built into the server process.

Syntax

```
AwsStringOutcome GetSdkVersion()
```

Parameters

This action has no parameters.

Return value

If successful, returns the current SDK version as an `AwsStringOutcome` object. The returned string includes the version number only (ex. "3.1.5"). If not successful, returns an error message.

Example

```
var getSdkVersionOutcome = GameLiftServerAPI.GetSdkVersion();
```

GetTerminationTime()

Returns the time that a server process is scheduled to be shut down, if a termination time is available. A server process takes this action after receiving an `onProcessTerminate()` callback from the Amazon GameLift Servers service. Amazon GameLift Servers may call `onProcessTerminate()` for the following reasons: (1) for poor health (the server process has reported port health or has not responded to Amazon GameLift Servers, (2) when terminating the instance during a scale-down event, or (3) when an instance is being terminated due to a [spot-instance interruption](#).

If the process has received an `onProcessTerminate()` callback, the value returned is the estimated termination time. If the process has not received an `onProcessTerminate()` callback, an error message is returned. Learn more about [shutting down a server process](#).

Syntax

```
AwsDateTimeOutcome GetTerminationTime()
```

Parameters

This action has no parameters.

Return value

If successful, returns the termination time as an `AwsDateTimeOutcome` object. The value is the termination time, expressed in elapsed ticks since 0001 00:00:00. For example, the date time value

2020-09-13 12:26:40 -000Z is equal to 637355968000000000 ticks. If no termination time is available, returns an error message.

Example

```
var getTerminationTimeOutcome = GameLiftServerAPI.GetTerminationTime();
```

InitSDK()

Initializes the Amazon GameLift Servers SDK. This method should be called on launch, before any other Amazon GameLift Servers-related initialization occurs.

Syntax

```
InitSDKOutcome InitSDK()
```

Parameters

This action has no parameters.

Return value

If successful, returns an `InitSdkOutcome` object indicating that the server process is ready to call [ProcessReady\(\)](#).

Example

```
var initSDKOutcome = GameLiftServerAPI.InitSDK();
```

ProcessEnding()

Notifies the Amazon GameLift Servers service that the server process is shutting down. This method should be called after all other cleanup tasks, including shutting down all active game sessions. This method should exit with an exit code of 0; a non-zero exit code results in an event message that the process did not exit cleanly.

Once the method exits with a code of 0, you can terminate the process with a successful exit code. You can also exit the process with an error code. If you exit with an error code, the fleet event will indicate the process terminated abnormally (`SERVER_PROCESS_TERMINATED_UNHEALTHY`).

Syntax

```
GenericOutcome ProcessEnding()
```

Parameters

This action has no parameters.

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

```
var processEndingOutcome = GameLiftServerAPI.ProcessEnding();
if (processReadyOutcome.Success)
    Environment.Exit(0);
// otherwise, exit with error code
Environment.Exit(errorCode);
```

ProcessReady()

Notifies the Amazon GameLift Servers service that the server process is ready to host game sessions. Call this method after successfully invoking [InitSDK\(\)](#) and completing setup tasks that are required before the server process can host a game session. This method should be called only once per process.

Syntax

```
GenericOutcome ProcessReady(ProcessParameters processParameters)
```

Parameters

processParameters

A [ProcessParameters](#) object communicating the following information about the server process:

- Names of callback methods, implemented in the game server code, that the Amazon GameLift Servers service invokes to communicate with the server process.
- Port number that the server process is listening on.

- Path to any game session-specific files that you want Amazon GameLift Servers to capture and store.

Required: Yes

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

This example illustrates both the [ProcessReady\(\)](#) call and delegate function implementations.

```
// Set parameters and call ProcessReady
var processParams = new ProcessParameters(
    this.OnGameSession,
    this.OnProcessTerminate,
    this.OnHealthCheck,
    this.OnGameSessionUpdate,
    port,
    new LogParameters(new List<string>()           // Examples of log and error files
        {
            "C:\\game\\logs",
            "C:\\game\\error"
        })
    written by the game server
);

var processReadyOutcome = GameLiftServerAPI.ProcessReady(processParams);

// Implement callback functions
void OnGameSession(GameSession gameSession)
{
    // game-specific tasks when starting a new game session, such as loading map
    // When ready to receive players
    var activateGameSessionOutcome = GameLiftServerAPI.ActivateGameSession();
}

void OnProcessTerminate()
{
    // game-specific tasks required to gracefully shut down a game session,
    // such as notifying players, preserving game state data, and other cleanup
    var ProcessEndingOutcome = GameLiftServerAPI.ProcessEnding();
}
```

```
}

bool OnHealthCheck()
{
    bool isHealthy;
    // complete health evaluation within 60 seconds and set health
    return isHealthy;
}
```

RemovePlayerSession()

Notifies the Amazon GameLift Servers service that a player with the specified player session ID has disconnected from the server process. In response, Amazon GameLift Servers changes the player slot to available, which allows it to be assigned to a new player.

Syntax

```
GenericOutcome RemovePlayerSession(String playerId)
```

Parameters

playerSessionId

Unique ID issued by Amazon GameLift Servers when a new player session is created. A player session ID is specified in a `PlayerSession` object, which is returned in response to a client call to the *GameLift API* actions [StartGameSessionPlacement](#), [CreateGameSession](#), [DescribeGameSessionPlacement](#), or [DescribePlayerSessions](#).

Type: String

Required: Yes

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

```
Aws::GameLift::GenericOutcome disconnectOutcome =
    Aws::GameLift::Server::RemovePlayerSession(playerSessionId);
```

StartMatchBackfill()

Sends a request to find new players for open slots in a game session created with FlexMatch. See also the Amazon SDK action [StartMatchBackfill\(\)](#). With this action, match backfill requests can be initiated by a game server process that is hosting the game session. Learn more about the [FlexMatch backfill feature](#).

This action is asynchronous. If new players are successfully matched, the Amazon GameLift Servers service delivers updated matchmaker data using the callback function `OnUpdateGameSession()`.

A server process can have only one active match backfill request at a time. To send a new request, first call [StopMatchBackfill\(\)](#) to cancel the original request.

Syntax

```
StartMatchBackfillOutcome StartMatchBackfill (StartMatchBackfillRequest  
startBackfillRequest);
```

Parameters

StartMatchBackfillRequest

A [StartMatchBackfillRequest](#) object that communicates the following information:

- A ticket ID to assign to the backfill request. This information is optional; if no ID is provided, Amazon GameLift Servers will autogenerate one.
- The matchmaker to send the request to. The full configuration ARN is required. This value can be acquired from the game session's matchmaker data.
- The ID of the game session that is being backfilled.
- Available matchmaking data for the game session's current players.

Required: Yes

Return value

Returns a `StartMatchBackfillOutcome` object with the match backfill ticket ID or failure with an error message.

Example

```
// Build a backfill request
```

```
var startBackfillRequest = new AWS.GameLift.Server.Model.StartMatchBackfillRequest()
{
    TicketId = "a ticket ID", //optional
    MatchmakingConfigurationArn = "the matchmaker configuration ARN",
    GameSessionId = GameLiftServerAPI.GetGameSessionId().Result, // gets ID for
current game session
    //get player data for all currently connected players
    MatchmakerData matchmakerData =
        MatchmakerData.FromJson(gameSession.MatchmakerData); // gets matchmaker
data for current players
    // get matchmakerData.Players
    // remove data for players who are no longer connected
    Players = ListOfPlayersRemainingInTheGame
};

// Send backfill request
var startBackfillOutcome = GameLiftServerAPI.StartMatchBackfill(startBackfillRequest);

// Implement callback function for backfill
void OnUpdateGameSession(GameSession myGameSession)
{
    // game-specific tasks to prepare for the newly matched players and update
matchmaker data as needed
}
```

StopMatchBackfill()

Cancels an active match backfill request that was created with [StartMatchBackfill\(\)](#). See also the Amazon SDK action [StopMatchmaking\(\)](#). Learn more about the [FlexMatch backfill feature](#).

Syntax

```
GenericOutcome StopMatchBackfill (StopMatchBackfillRequest stopBackfillRequest);
```

Parameters

StopMatchBackfillRequest

A [StopMatchBackfillRequest](#) object identifying the matchmaking ticket to cancel:

- ticket ID assigned to the backfill request being canceled
- matchmaker the backfill request was sent to
- game session associated with the backfill request

Required: Yes

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

```
// Set backfill stop request parameters

var stopBackfillRequest = new AWS.GameLift.Server.Model.StopMatchBackfillRequest()
{
    TicketId = "a ticket ID", //optional, if not provided one is autogenerated
    MatchmakingConfigurationArn = "the matchmaker configuration ARN", //from the game
    session matchmaker data
    GameSessionId = GameLiftServerAPI.GetGameSessionId().Result //gets the ID for
    the current game session
};

var stopBackfillOutcome =
    GameLiftServerAPI.StopMatchBackfillRequest(stopBackfillRequest);
```

TerminateGameSession()

This method is deprecated with version 4.0.1. Instead, the server process should call [ProcessEnding\(\)](#) after a game session has ended.

Notifies the Amazon GameLift Servers service that the server process has ended the current game session. This action is called when the server process will remain active and ready to host a new game session. It should be called only after your game session termination procedure is complete, because it signals to Amazon GameLift Servers that the server process is immediately available to host a new game session.

This action is not called if the server process will be shut down after the game session stops. Instead, call [ProcessEnding\(\)](#) to signal that both the game session and the server process are ending.

Syntax

```
GenericOutcome TerminateGameSession()
```

Parameters

This action has no parameters.

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

This example illustrates a server process at the end of a game session.

```
// game-specific tasks required to gracefully shut down a game session,  
// such as notifying players, preserving game state data, and other cleanup  
  
var terminateGameSessionOutcome = GameLiftServerAPI.TerminateGameSession();  
var processReadyOutcome = GameLiftServerAPI.ProcessReady(processParams);
```

UpdatePlayerSessionCreationPolicy()

Updates the current game session's ability to accept new player sessions. A game session can be set to either accept or deny all new player sessions. (See also the [UpdateGameSession\(\)](#) action in the *Amazon GameLift Servers Service API Reference*).

Syntax

```
GenericOutcome UpdatePlayerSessionCreationPolicy(PlayerSessionCreationPolicy  
playerSessionPolicy)
```

Parameters

newPlayerSessionPolicy

String value indicating whether the game session accepts new players.

Type: [PlayerSessionCreationPolicy](#) enum. Valid values include:

- **ACCEPT_ALL** – Accept all new player sessions.
- **DENY_ALL** – Deny all new player sessions.

Required: Yes

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

This example sets the current game session's join policy to accept all players.

```
var updatePlayerSessionCreationPolicyOutcome =  
  
    GameLiftServerAPI.UpdatePlayerSessionCreationPolicy(PlayerSessionCreationPolicy.ACCEPT_ALL);
```

Server SDK (Unreal) for Amazon GameLift Servers -- Actions

Use the server SDK for Unreal to integrate your multiplayer game for hosting with Amazon GameLift Servers. For guidance about the integration process, see [Add Amazon GameLift Servers to your game server with the server SDK](#).

Note

This reference is for an earlier version of the server SDK for Amazon GameLift Servers. For the latest version, see [???](#).

This API is defined in `GameLiftServerSDK.h` and `GameLiftServerSDKModels.h`.

To set up the Unreal Engine plugin and see code examples [???](#).

Server SDK (Unreal) for Amazon GameLift Servers -- Data types

Use the Amazon GameLift Servers server SDK for Unreal reference to integrate your multiplayer game for hosting with Amazon GameLift Servers. For guidance about the integration process, see [Add Amazon GameLift Servers to your game server with the server SDK](#).

Note

This reference is for an earlier version of the server SDK for Amazon GameLift Servers. For the latest version, see [???](#).

This API is defined in `GameLiftServerSDK.h` and `GameLiftServerSDKModels.h`.

To set up the Unreal Engine plugin and see code examples [???](#).

[Server SDK \(Unreal\) for Amazon GameLift Servers -- Actions](#)

Topics

- [FDescribePlayerSessionsRequest](#)
- [FProcessParameters](#)
- [FStartMatchBackfillRequest](#)
- [FStopMatchBackfillRequest](#)

FDescribePlayerSessionsRequest

This data type is used to specify which player session(s) to retrieve. You can use it as follows:

- Provide a `PlayerSessionId` to request a specific player session.
- Provide a `GameSessionId` to request all player sessions in the specified game session.
- Provide a `PlayerId` to request all player sessions for the specified player.

For large collections of player sessions, use the pagination parameters to retrieve results in sequential blocks.

Contents

GameSessionId

Unique game session identifier. Use this parameter to request all player sessions for the specified game session. Game session ID format is as follows: `arn:aws-cn:gamelift:<region>::gamesession/fleet-<fleet ID>/<ID string>`. The value of `<ID string>` is either a custom ID string or (if one was specified when the game session was created) a generated string.

Type: String

Required: No

Limit

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

Required: No

NextToken

Token indicating the start of the next sequential page of results. Use the token that is returned with a previous call to this action. To specify the start of the result set, do not specify a value. If a player session ID is specified, this parameter is ignored.

Type: String

Required: No

PlayerId

Unique identifier for a player. Player IDs are defined by the developer. See [Generate player IDs](#).

Type: String

Required: No

PlayerSessionId

Unique identifier for a player session.

Type: String

Required: No

PlayerSessionStatusFilter

Player session status to filter results on. 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 time-out limit (60 seconds).

Type: String

Required: No

FProcessParameters

This data type contains the set of parameters sent to the Amazon GameLift Servers service in a [ProcessReady\(\)](#) call.

Contents

port

Port number the server process will listen on for new player connections. The value must fall into the port range configured for any fleet deploying this game server build. This port number is included in game session and player session objects, which game sessions use when connecting to a server process.

Type: Integer

Required: Yes

logParameters

Object with a list of directory paths to game session log files.

Type: TArray<FString>

Required: No

onStartGameSession

Name of callback function that the Amazon GameLift Servers service invokes to activate a new game session. Amazon GameLift Servers calls this function in response to the client request [CreateGameSession](#). The callback function takes a [GameSession](#) object (defined in the *Amazon GameLift Servers Service API Reference*).

Type: FOnStartGameSession

Required: Yes

onProcessTerminate

Name of callback function that the Amazon GameLift Servers service invokes to force the server process to shut down. After calling this function, Amazon GameLift Servers waits five minutes

for the server process to shut down and respond with a [ProcessEnding\(\)](#) call before it shuts down the server process.

Type: FSimpleDelegate

Required: No

onHealthCheck

Name of callback function that the Amazon GameLift Servers service invokes to request a health status report from the server process. Amazon GameLift Servers calls this function every 60 seconds. After calling this function Amazon GameLift Servers waits 60 seconds for a response, and if none is received, records the server process as unhealthy.

Type: FOnHealthCheck

Required: No

onUpdateGameSession

Name of callback function that the Amazon GameLift Servers service invokes to pass an updated game session object to the server process. Amazon GameLift Servers calls this function when a [match backfill](#) request has been processed in order to provide updated matchmaker data. It passes a [GameSession](#) object, a status update (updateReason), and the match backfill ticket ID.

Type: FOnUpdateGameSession

Required: No

FStartMatchBackfillRequest

This data type is used to send a matchmaking backfill request. The information is communicated to the Amazon GameLift Servers service in a [StartMatchBackfill\(\)](#) call.

Contents

GameSessionArn

Unique game session identifier. The API action [GetGameSessionId\(\)](#) returns the identifier in ARN format.

Type: FString

Required: Yes

MatchmakingConfigurationArn

Unique identifier, in the form of an ARN, for the matchmaker to use for this request. To find the matchmaker that was used to create the original game session, look in the game session object, in the matchmaker data property. Learn more about matchmaker data in [Work with matchmaker data](#).

Type: FString

Required: Yes

Players

A set of data representing all players who are currently in the game session. The matchmaker uses this information to search for new players who are good matches for the current players. See the *Amazon GameLift Servers API Reference Guide* for a description of the Player object format. To find player attributes, IDs, and team assignments, look in the game session object, in the matchmaker data property. If latency is used by the matchmaker, gather updated latency for the current region and include it in each player's data.

Type: TArray<[FPlayer](#)>

Required: Yes

TicketId

Unique identifier for a matchmaking or match backfill request ticket. If no value is provided here, Amazon GameLift Servers will generate one in the form of a UUID. Use this identifier to track the match backfill ticket status or cancel the request if needed.

Type: FString

Required: No

FStopMatchBackfillRequest

This data type is used to cancel a matchmaking backfill request. The information is communicated to the Amazon GameLift Servers service in a [StopMatchBackfill\(\)](#) call.

Contents

GameSessionArn

Unique game session identifier associated with the request being canceled.

Type: FString

Required: Yes

MatchmakingConfigurationArn

Unique identifier of the matchmaker this request was sent to.

Type: FString

Required: Yes

TicketId

Unique identifier of the backfill request ticket to be canceled.

Type: FString

Required: Yes

[Server SDK \(Unreal\) for Amazon GameLift Servers -- Data types](#)

Topics

- [AcceptPlayerSession\(\)](#)
- [ActivateGameSession\(\)](#)
- [DescribePlayerSessions\(\)](#)
- [GetGameSessionId\(\)](#)
- [GetSdkVersion\(\)](#)
- [InitSDK\(\)](#)
- [ProcessEnding\(\)](#)
- [ProcessReady\(\)](#)
- [RemovePlayerSession\(\)](#)
- [StartMatchBackfill\(\)](#)

- [StopMatchBackfill\(\)](#)
- [TerminateGameSession\(\)](#)
- [UpdatePlayerSessionCreationPolicy\(\)](#)

AcceptPlayerSession()

Notifies the Amazon GameLift Servers service that a player with the specified player session ID has connected to the server process and needs validation. Amazon GameLift Servers verifies that the player session ID is valid—that is, that the player ID has reserved a player slot in the game session. Once validated, Amazon GameLift Servers changes the status of the player slot from RESERVED to ACTIVE.

Syntax

```
FGameLiftGenericOutcome AcceptPlayerSession(const FString& playerId)
```

Parameters

playerSessionId

Unique ID issued by the Amazon GameLift Servers service in response to a call to the Amazon SDK Amazon GameLift Servers API action [CreatePlayerSession](#). The game client references this ID when connecting to the server process.

Type: FString

Required: Yes

Return value

Returns a generic outcome consisting of success or failure with an error message.

ActivateGameSession()

Notifies the Amazon GameLift Servers service that the server process has activated a game session and is now ready to receive player connections. This action should be called as part of the `onStartGameSession()` callback function, after all game session initialization has been completed.

Syntax

```
FGameLiftGenericOutcome ActivateGameSession()
```

Parameters

This action has no parameters.

Return value

Returns a generic outcome consisting of success or failure with an error message.

DescribePlayerSessions()

Retrieves player session data, including settings, session metadata, and player data. Use this action to get information for a single player session, for all player sessions in a game session, or for all player sessions associated with a single player ID.

Syntax

```
FGameLiftDescribePlayerSessionsOutcome DescribePlayerSessions(const  
FGameLiftDescribePlayerSessionsRequest &describePlayerSessionsRequest)
```

Parameters

describePlayerSessionsRequest

A [FDescribePlayerSessionsRequest](#) object describing which player sessions to retrieve.

Required: Yes

Return value

If successful, returns a [FDescribePlayerSessionsRequest](#) object containing a set of player session objects that fit the request parameters. Player session objects have a structure identical to the Amazon SDK Amazon GameLift Servers API [PlayerSession](#) data type.

GetGameSessionId()

Retrieves the ID of the game session currently being hosted by the server process, if the server process is active.

Syntax

```
FGameLiftStringOutcome GetGameSessionId()
```

Parameters

This action has no parameters.

Return value

If successful, returns the game session ID as an `FGameLiftStringOutcome` object. If not successful, returns an error message.

GetSdkVersion()

Returns the current version number of the SDK built into the server process.

Syntax

```
FGameLiftStringOutcome GetSdkVersion();
```

Parameters

This action has no parameters.

Return value

If successful, returns the current SDK version as an `FGameLiftStringOutcome` object. The returned string includes the version number only (ex. "3.1.5"). If not successful, returns an error message.

Example

```
Aws::GameLift::AwsStringOutcome SdkVersionOutcome =  
    Aws::GameLift::Server::GetSdkVersion();
```

InitSDK()

Initializes the Amazon GameLift Servers SDK. This method should be called on launch, before any other Amazon GameLift Servers-related initialization occurs.

Syntax

```
FGameLiftGenericOutcome InitSDK()
```

Parameters

This action has no parameters.

Return value

Returns a generic outcome consisting of success or failure with an error message.

ProcessEnding()

Notifies the Amazon GameLift Servers service that the server process is shutting down. This method should be called after all other cleanup tasks, including shutting down all active game sessions. This method should exit with an exit code of 0; a non-zero exit code results in an event message that the process did not exit cleanly.

Syntax

```
FGameLiftGenericOutcome ProcessEnding()
```

Parameters

This action has no parameters.

Return value

Returns a generic outcome consisting of success or failure with an error message.

ProcessReady()

Notifies the Amazon GameLift Servers service that the server process is ready to host game sessions. Call this method after successfully invoking [InitSDK\(\)](#) and completing setup tasks that are required before the server process can host a game session. This method should be called only once per process.

Syntax

```
FGameLiftGenericOutcome ProcessReady(FProcessParameters &processParameters)
```

Parameters

FProcessParameters

A [FProcessParameters](#) object communicating the following information about the server process:

- Names of callback methods, implemented in the game server code, that the Amazon GameLift Servers service invokes to communicate with the server process.
- Port number that the server process is listening on.
- Path to any game session-specific files that you want Amazon GameLift Servers to capture and store.

Required: Yes

Return value

Returns a generic outcome consisting of success or failure with an error message.

Example

See the sample code in [Using the Unreal Engine Plugin](#).

RemovePlayerSession()

Notifies the Amazon GameLift Servers service that a player with the specified player session ID has disconnected from the server process. In response, Amazon GameLift Servers changes the player slot to available, which allows it to be assigned to a new player.

Syntax

```
FGameLiftGenericOutcome RemovePlayerSession(const FString& playerId)
```

Parameters

playerSessionId

Unique ID issued by the Amazon GameLift Servers service in response to a call to the Amazon SDK Amazon GameLift Servers API action [CreatePlayerSession](#). The game client references this ID when connecting to the server process.

Type: FString

Required: Yes

Return value

Returns a generic outcome consisting of success or failure with an error message.

StartMatchBackfill()

Sends a request to find new players for open slots in a game session created with FlexMatch. See also the Amazon SDK action [StartMatchBackfill\(\)](#). With this action, match backfill requests can be initiated by a game server process that is hosting the game session. Learn more about the [FlexMatch backfill feature](#).

This action is asynchronous. If new players are successfully matched, the Amazon GameLift Servers service delivers updated matchmaker data using the callback function `OnUpdateGameSession()`.

A server process can have only one active match backfill request at a time. To send a new request, first call [StopMatchBackfill\(\)](#) to cancel the original request.

Syntax

```
FGameLiftStringOutcome StartMatchBackfill (FStartMatchBackfillRequest  
&startBackfillRequest);
```

Parameters

FStartMatchBackfillRequest

A [FStartMatchBackfillRequest](#) object that communicates the following information:

- A ticket ID to assign to the backfill request. This information is optional; if no ID is provided, Amazon GameLift Servers will autogenerate one.
- The matchmaker to send the request to. The full configuration ARN is required. This value can be acquired from the game session's matchmaker data.
- The ID of the game session that is being backfilled.
- Available matchmaking data for the game session's current players.

Required: Yes

Return value

If successful, returns the match backfill ticket as a `FGameLiftStringOutcome` object. If not successful, returns an error message. Ticket status can be tracked using the Amazon SDK action [DescribeMatchmaking\(\)](#).

StopMatchBackfill()

Cancels an active match backfill request that was created with [StartMatchBackfill\(\)](#). See also the Amazon SDK action [StopMatchmaking\(\)](#). Learn more about the [FlexMatch backfill feature](#).

Syntax

```
FGameLiftGenericOutcome StopMatchBackfill (FStopMatchBackfillRequest  
&stopBackfillRequest);
```

Parameters

StopMatchBackfillRequest

A [FStopMatchBackfillRequest](#) object identifying the matchmaking ticket to cancel:

- ticket ID assigned to the backfill request being canceled
- matchmaker the backfill request was sent to
- game session associated with the backfill request

Required: Yes

Return value

Returns a generic outcome consisting of success or failure with an error message.

TerminateGameSession()

This method is deprecated with version 4.0.1. Instead, the server process should call [ProcessEnding\(\)](#) after a game session has ended.

Notifies the Amazon GameLift Servers service that the server process has ended the current game session. This action is called when the server process will remain active and ready to host a new game session. It should be called only after your game session termination procedure is complete,

because it signals to Amazon GameLift Servers that the server process is immediately available to host a new game session.

This action is not called if the server process will be shut down after the game session stops. Instead, call [ProcessEnding\(\)](#) to signal that both the game session and the server process are ending.

Syntax

```
FGameLiftGenericOutcome TerminateGameSession()
```

Parameters

This action has no parameters.

Return value

Returns a generic outcome consisting of success or failure with an error message.

UpdatePlayerSessionCreationPolicy()

Updates the current game session's ability to accept new player sessions. A game session can be set to either accept or deny all new player sessions. (See also the [UpdateGameSession\(\)](#) action in the *Amazon GameLift Servers Service API Reference*).

Syntax

```
FGameLiftGenericOutcome UpdatePlayerSessionCreationPolicy(EPlayerSessionCreationPolicy policy)
```

Parameters

Policy

Value indicating whether the game session accepts new players.

Type: `EPlayerSessionCreationPolicy` enum. Valid values include:

- **ACCEPT_ALL** – Accept all new player sessions.
- **DENY_ALL** – Deny all new player sessions.

Required: Yes

Return value

Returns a generic outcome consisting of success or failure with an error message.

Game session placement events

Amazon GameLift Servers emits events for each game session placement request as it is processed. You can publish these events to an Amazon SNS topic, as described in [Set up event notification for game session placement](#). These events are also emitted to Amazon CloudWatch Events in near real time and on a best-effort basis.

This topic describes the structure of game session placement events and provides an example for each event type. For more information on the status of game session placement requests, see [GameSessionPlacement](#) in the *Amazon GameLift Servers API Reference*.

Placement event syntax

Events are represented as JSON objects. Event structure conforms to the CloudWatch Events pattern, with similar top-level fields and service-specific details.

Top-level fields include the following (see [event pattern](#) for more detail):

version

This field is always set to 0 (zero).

id

Unique tracking identifier for the event.

detail-type

Value is always `GameLift Queue Placement Event`.

source

Value is always `aws.gamelift`.

account

The Amazon account that is being used to manage Amazon GameLift Servers.

time

Event timestamp.

region

The Amazon Region where the placement request is being processed. This is the Region where the game session queue in use resides.

resources

ARN value of the game session queue that is processing the placement request.

PlacementFulfilled

The placement request has been successfully fulfilled. A new game session has been started and new player sessions have been created for each player listed in the game session placement request. Player connection information is available.

Detail syntax:

placementId

A unique identifier assigned to the game session placement request.

port

The port number for the new game session.

gameSessionArn

The ARN identifier for the new game session.

ipAddress

The IP address of the game session.

dnsName

The DNS identifier assigned to the instance that is running the new game session. The value format is different depending on whether the instance running the game session is TLS-enabled. When connecting to a game session on a TLS-enabled fleet, players must use the DNS name, not the IP address.

TLS-enabled fleets: <unique identifier>.<region identifier>.amazongamelift.com.

Non-TLS-enabled fleets: `ec2-<unique identifier>.compute.amazonaws.com`.

`startTime`

Time stamp indicating when this request was placed in the queue.

`endTime`

Time stamp indicating when this request was fulfilled.

`gameSessionRegion`

Amazon Region of the fleet that is hosting the game session. This corresponds to the region token in the `GameSessionArn`.

`placedPlayerSessions`

The collection of player sessions that have been created for each player in the game session placement request.

Example

```
{
  "version": "0",
  "id": "1111aaaa-bb22-cc33-dd44-5555eeee66ff",
  "detail-type": "GameLift Queue Placement Event",
  "source": "aws.gamelift",
  "account": "123456789012",
  "time": "2021-03-01T15:50:52Z",
  "region": "us-east-1",
  "resources": [
    "arn:aws:gamelift:us-west-2:123456789012:gamesessionqueue/MegaFrogRace-NA"
  ],
  "detail": {
    "type": "PlacementFulfilled",
    "placementId": "9999ffff-88ee-77dd-66cc-5555bb44aa",
    "port": "6262",
    "gameSessionArn": "arn:aws:gamelift:us-west-2::gamesession/fleet-2222bbbb-33cc-44dd-55ee-6666ffff77aa/4444dddd-55ee-66ff-77aa-8888bbbb99cc",
    "ipAddress": "98.987.98.987",
    "dnsName": "ec2-12-345-67-890.us-west-2.compute.amazonaws.com",
    "startTime": "2021-03-01T15:50:49.741Z",
    "endTime": "2021-03-01T15:50:52.084Z",
    "gameSessionRegion": "us-west-2",
    "placedPlayerSessions": [
```

```
{
  "playerId": "player-1"
  "playerSessionId": "psess-1232131232324124123123"
}
]
```

PlacementCancelled

The placement request was canceled with a call to the GameLift service [StopGameSessionPlacement](#).

Detail:

placementId

A unique identifier assigned to the game session placement request.

startTime

Time stamp indicating when this request was placed in the queue.

endTime

Time stamp indicating when this request was cancelled.

Example

```
{
  "version": "0",
  "id": "1111aaaa-bb22-cc33-dd44-5555eeee66ff",
  "detail-type": "GameLift Queue Placement Event",
  "source": "aws.gamelift",
  "account": "123456789012",
  "time": "2021-03-01T15:50:52Z",
  "region": "us-east-1",
  "resources": [
    "arn:aws:gamelift:us-west-2:123456789012:gamesessionqueue/MegaFrogRace-NA"
  ],
  "detail": {
    "type": "PlacementCancelled",
```

```
"placementId": "9999ffff-88ee-77dd-66cc-5555bb44aa",
"startTime": "2021-03-01T15:50:49.741Z",
"endTime": "2021-03-01T15:50:52.084Z"
}
}
```

PlacementTimedOut

Game session placement did not successfully complete before the queue's time limit expired. The placement request can be resubmitted as needed.

Detail:

placementId

A unique identifier assigned to the game session placement request.

startTime

Time stamp indicating when this request was placed in the queue.

endTime

Time stamp indicating when this request was cancelled.

Example

```
{
  "version": "0",
  "id": "1111aaaa-bb22-cc33-dd44-5555eeee66ff",
  "detail-type": "GameLift Queue Placement Event",
  "source": "aws.gamelift",
  "account": "123456789012",
  "time": "2021-03-01T15:50:52Z",
  "region": "us-east-1",
  "resources": [
    "arn:aws:gamelift:us-west-2:123456789012:gamesessionqueue/MegaFrogRace-NA"
  ],
  "detail": {
    "type": "PlacementTimedOut",
    "placementId": "9999ffff-88ee-77dd-66cc-5555bb44aa",
    "startTime": "2021-03-01T15:50:49.741Z",
```

```
    "endTime": "2021-03-01T15:50:52.084Z"  
  }  
}
```

PlacementFailed

Amazon GameLift Servers was not able to fulfill the game session request. This is generally caused by an unexpected internal error. The placement request can be resubmitted as needed.

Detail:

placementId

A unique identifier assigned to the game session placement request.

startTime

Time stamp indicating when this request was placed in the queue.

endTime

Time stamp indicating when this request failed.

Example

```
{  
  "version": "0",  
  "id": "39c978f3-ba46-3f7c-e787-55bfcca1bd31",  
  "detail-type": "GameLift Queue Placement Event",  
  "source": "aws.gamelift",  
  "account": "252386620677",  
  "time": "2021-03-01T15:50:52Z",  
  "region": "us-east-1",  
  "resources": [  
    "arn:aws:gamelift:us-west-2:252386620677:gamesessionqueue/MegaFrogRace-NA"  
  ],  
  "detail": {  
  
    "type": "PlacementFailed",  
    "placementId": "e4a1119a-39af-45cf-a990-ef150fe0d453",  
    "startTime": "2021-03-01T15:50:49.741Z",  
    "endTime": "2021-03-01T15:50:52.084Z"  
  }  
}
```

```
}
```

Amazon GameLift Servers AMI versions

The following table identifies the latest Amazon machine images (AMIs) that Amazon GameLift Servers uses for managed EC2 hosting. As described in [Configuration and vulnerability analysis in Amazon GameLift Servers](#), you must regularly create new Amazon GameLift Servers managed EC2 fleets to deploy the latest AMI version updates.

AMIs for use with server SDK for Amazon GameLift Servers 5+

Amazon GameLift Servers uses the following AMIs to host game servers that are integrated with server SDK for Amazon GameLift Servers version 5.

Amazon GameLift Servers image	Architecture	Most recent patch	Base Amazon image
Amazon Linux 2023 BASE_AMI_LINUX_2023	x86	2025-12-16	Amazon Linux 2023 AMI 2023.9.20 251208.0 x86_64 HVM kernel-6.1 (release notes)
Amazon Linux 2 BASE_AMI_LINUX_2_SDK_5	x86	2025-12-16	Amazon Linux 2 Kernel 5.10 AMI 2.0.20251208.0 x86_64 HVM gp2 (release notes)
Windows 2016 BASE_AMI_WINDOWS_2016_SDK_5	x86	2025-12-16	Windows_Server-2016-English-Full-Base-2025.12.10 (release notes)
Amazon Linux 2023 BASE_AMI_LINUX_2023_ARM	ARM64	2025-12-16	Amazon Linux 2023 AMI 2023.9.20 251208.0 arm64 HVM kernel-6.1 (release notes)
Amazon Linux 2	ARM64	2025-12-16	Amazon Linux 2 LTS Arm64 Kernel 5.10 AMI 2.0.20251208.0 arm64 HVM gp2 (release notes)

Amazon GameLift Servers image	Architecture	Most recent patch	Base Amazon image
BASE_AMI_LINUX_2_ARM			
Windows 2022 BASE_AMI_WINDOWS_2022	x86	2025-12-16	Windows_Server-2022-English-Full-Base-2025.12.10 (release notes)

AMIs for use with server SDK for Amazon GameLift Servers 4

Amazon GameLift Servers uses the following AMIs to host game servers that are integrated with server SDK for Amazon GameLift Servers version 4 or earlier.

Amazon GameLift Servers image	Architecture	Most recent patch	Base Amazon image
Amazon Linux 2 BASE_AMI_LINUX_2_SDK_4	x86	2025-12-16	Amazon Linux 2 AMI 2.0.20251208.0 x86_64 HVM gp2 (release notes)
Amazon Linux 2023 BASE_AMI_LINUX_2023_SDK_4	x86	2025-12-16	Amazon Linux 2023 AMI 2023.9.20 251208.0 x86_64 HVM kernel-6.1 (release notes)
Windows 2016 BASE_AMI_WINDOWS_2016_SDK_4	x86	2025-12-16	Windows_Server-2016-English-Full-Base-2025.12.10 (release notes)

For more information, see the following resources:

- [Amazon Linux 2023 release notes](#)

- [Amazon Linux 2 release notes](#)
- [Amazon Windows AMI version history](#)
- [Description of Software Update Services and Windows Server Update Services changes in content for 2024](#)

Amazon GameLift Servers service locations

Amazon GameLift Servers features are available across multiple Amazon Web Services Regions and Local Zones. You can design a hosting solution that puts your game servers right where your players are located.


Locations for managed hosting

Amazon GameLift Servers managed hosting deploys fleets of game server resources. Each fleet is created in an Amazon Web Services Region, which is the fleet's *home region*. A fleet's home Region is referenced in the fleet's Amazon Resource Number (ARN).

You can deploy a *single-region fleet*, with hosting resources in the home region only. Alternatively, you can deploy a *multi-location fleet*, with hosting resources in multiple geographic locations. A multi-location fleet has a home region and one or more *remote locations*. There are no limitations on how you group locations together in a fleet. When managing hosting capacity for a multi-location fleet, you can set capacity for each location individually.

Remote locations for a multi-location fleet can be other Amazon Web Services Regions or Local Zones. A *Local Zone* is an extension of an Amazon Web Services Region. It lets you place compute resources closer to users to provide lower-latency gameplay. For more information, see [Amazon Local Zones](#). The location code for a Local Zone is its parent Region code followed by a physical location identifier. For example, the code for the Los Angeles Local Zone is `us-west-2-lax-1`.

The following diagram illustrates a multi-location fleet with resources in two Amazon Web Services Regions and one Local Zone. The fleet's home Region is `us-west-2`, and it has two remote locations: `us-east-2` Region and `us-west-2-lax-1` Local Zone.



In addition to fleet resources, managed hosting with Amazon GameLift Servers can also use the following types of resources. You create these resources in a specific Amazon Web Services Region that supports the resource type.

- *Build* – This is a game server build to be hosted with a managed EC2 fleet. Create a build resource in the same region as the fleet that it will be deployed to.
- *Script* – This is a configuration script for hosting a game with Amazon GameLift Servers Realtime. Create a script resource in the same Region as the fleet that it will be deployed to.
- *Container group definition* – This is a configuration for running containers on a managed container fleet. It identifies one or more *container images* with game server software to deploy to the container fleet. Create a container group definition in the same Region as the fleet they'll be deployed to.
- *Game session queue* – This resource processes requests for game sessions and initiates new game sessions. Processing takes place in the Amazon Web Services Region where the queue is located. To reduce latency in the game session placement process, create a queue geographically near the players that will use it.

Locations for Amazon GameLift Servers FlexMatch

FlexMatch resources are used to process player requests for matchmaking. They include a matchmaking configuration resource and a rule set resource. Processing takes place in the Amazon Web Services Region where the FlexMatch resources are located. To reduce latency in the matchmaking process, create the resources geographically near the players that will use it.

A matchmaking configuration and the rule set it uses must be located in the same Amazon Web Services Region. You can create FlexMatch resources in any of the Amazon Web Services Regions that support them.

For more information about setting up FlexMatch for your hosting solution, see the [Amazon GameLift Servers FlexMatch developer guide](#).

Amazon GameLift Servers in China

When using Amazon GameLift Servers for resources in the China (Beijing) Region, operated by Sinnet, or the China (Ningxia) Region, operated by NWCD, you must have a separate Amazon (China) account. Be aware that some features are unavailable in the China Regions. For more information about using Amazon GameLift Servers in these Regions, see the following resources:

- [Amazon Web Services in China](#)
- [Amazon GameLift Servers](#) (Getting Started with Amazon Web Services in China)

Amazon GameLift Servers endpoints and quotas

- For service endpoints that you can use to connect programmatically with the service, see [Amazon GameLift Servers service endpoints](#).
- Amazon GameLift Servers has quotas on the use of service resources and actions per Amazon account. For information about service quotas, including how to view quotas for a specific service or request quota increases, see the [Service Quotas User Guide](#) in the Amazon documentation.

Amazon GameLift Servers API limits

The tables in this section list the default rate limits for Amazon GameLift Servers API operations. These limits help maintain service stability and ensure fair usage. Each API operation has specific limits based on its resource impact and usage patterns.

By default, the throttle limit for each customer account is set to 500 for burst requests and 500 for normal traffic. In addition, each API has its own default limit pair as outlined below.

API limits reference

The following table lists the default rate limits for Amazon GameLift Servers API operations. APIs marked with resource-level throttling may be subject to additional limits to prevent database hot key issues.

API limits reference

API action	Burst limit	Rate limit	Resource-level throttling	Notes
AcceptMatch	100	100	No	Limit is per account.
ClaimGameServer	40	20	No	Limit is per account.
CreateAlias	4	1	No	Limit is per account.
CreateBuild	10	1	No	Limit is per account.
CreateContainerFleet	4	1	No	Limit is per account.
CreateContainerGroupDefinition	10	1	No	Limit is per account.
CreateFleet	4	1	No	Limit is per account.
CreateFleetLocations	4	1	No	Limit is per account.
CreateGameServerGroup	5	0.1	No	Limit is per account.
CreateGameSession	50	50	No	Limit is per account.
CreateGameSessionQueue	4	1	No	Limit is per account.
CreateLocation	4	1	No	Limit is per account.

API action	Burst limit	Rate limit	Resource-level throttling	Notes
CreateMatchmakingConfiguration	4	1	No	Limit is per account.
CreateMatchmakingRuleSet	4	1	No	Limit is per account.
CreatePlayerSession	50	50	No	Limit is per account.
CreatePlayerSessions	150	150	No	Limit is per account.
CreateScript	10	1	No	Limit is per account.
CreateVpcPeeringAuthorization	1	0.1	No	Limit is per account.
CreateVpcPeeringConnection	1	0.1	No	Limit is per account.
DeleteAlias	4	1	No	Limit is per account.
DeleteBuild	10	1	No	Limit is per account.
DeleteContainerFleet	4	1	No	Limit is per account.
DeleteContainerGroupDefinition	10	1	No	Limit is per account.
DeleteFleet	4	1	No	Limit is per account.
DeleteFleetLocations	4	1	No	Limit is per account.
DeleteGameServerGroup	5	0.1	No	Limit is per account.

API action	Burst limit	Rate limit	Resource-level throttling	Notes
DeleteGameSessionQueue	4	1	No	Limit is per account.
DeleteLocation	4	1	No	Limit is per account.
DeleteMatchmakingConfiguration	4	1	No	Limit is per account.
DeleteMatchmakingRuleSet	4	1	No	Limit is per account.
DeleteScalingPolicy	4	1	No	Limit is per account.
DeleteScript	10	1	No	Limit is per account.
DeleteVpcPeeringAuthorization	1	0.1	No	Limit is per account.
DeleteVpcPeeringConnection	1	0.1	No	Limit is per account.
DeregisterGameServer	40	20	No	Limit is per account.
DeregisterCompute	10	2	No	Limit is per account.
DescribeAlias	10	2	No	Limit is per account.
DescribeBuild	20	5	No	Limit is per account.
DescribeCompute	10	2	No	Limit is per account.
DescribeContainerFleet	20	5	No	Limit is per account.

API action	Burst limit	Rate limit	Resource-level throttling	Notes
DescribeContainerGroupDefinition	20	5	No	Limit is per account.
DescribeEC2InstanceLimits	40	15	No	Limit is per account.
DescribeFleetAttributes	10	2	No	Limit is per account.
DescribeFleetCapacity	4	1	No	Limit is per account.
DescribeFleetDeployment	20	5	No	Limit is per account.
DescribeFleetEvents	10	5	No	Limit is per account.
DescribeFleetLocationAttributes	10	2	No	Limit is per account.
DescribeFleetLocationCapacity	10	5	No	Limit is per account.
DescribeFleetLocationUtilization	10	2	No	Limit is per account.
DescribeFleetPortSettings	10	2	No	Limit is per account.
DescribeFleetUtilization	10	2	No	Limit is per account.
DescribeGameServer	40	20	No	Limit is per account.

API action	Burst limit	Rate limit	Resource-level throttling	Notes
DescribeGameServerGroup	10	2	No	Limit is per account.
DescribeGameServerInstances	10	2	No	Limit is per account.
DescribeGameSessionDetails	10	5	No	Limit is per account. Consider using DescribeGameSessions instead as it has a higher default limit and supports resource level limits.
DescribeGameSessionPlacement	75	75	Yes	Limit is per account. Subject to resource-level throttling.
DescribeGameSessionQueues	10	2	No	Limit is per account.
DescribeGameSessions	20	20	Yes	Limit is per account. Subject to resource-level throttling.
DescribeInstances	10	2	No	Limit is per account.
DescribeMatchmaking	75	75	Yes	Limit is per account. Subject to resource-level throttling.
DescribeMatchmakingConfigurations	15	7	No	Limit is per account.
DescribeMatchmakingRuleSets	20	10	No	Limit is per account.
DescribePlayerSessions	40	40	Yes	Limit is per account. Subject to resource-level throttling.

API action	Burst limit	Rate limit	Resource-level throttling	Notes
DescribeRuntimeConfiguration	4	1	No	Limit is per account.
DescribeScalingPolicies	20	20	No	Limit is per account.
DescribeScript	20	5	No	Limit is per account.
DescribeVpcPeeringAuthorizations	1	0.1	No	Limit is per account.
DescribeVpcPeeringConnections	1	0.1	No	Limit is per account.
GetComputeAccess	4	1	No	Limit is per account.
GetComputeAuthToken	10	5	No	Limit is per account.
GetGameSessionLogUrl	4	1	No	Limit is per account.
GetInstanceAccess	4	1	No	Limit is per account.
ListAliases	10	2	No	Limit is per account.
ListBuilds	20	5	No	Limit is per account.
ListContainerFleets	20	5	No	Limit is per account.
ListCompute	10	2	No	Limit is per account.
ListContainerGroupDefinitions	20	5	No	Limit is per account.

API action	Burst limit	Rate limit	Resource-level throttling	Notes
ListContainerGroupDefinitionVersions	20	5	No	Limit is per account.
ListFleets	10	2	No	Limit is per account.
ListFleetDeployments	20	5	No	Limit is per account.
ListGameServerGroups	10	2	No	Limit is per account.
ListGameServers	10	2	No	Limit is per account.
ListLocations	10	2	No	Limit is per account.
ListScripts	20	5	No	Limit is per account.
ListTagsForResource	20	5	No	Limit is per account.
PutScalingPolicy	4	1	No	Limit is per account.
RegisterCompute	10	5	No	Limit is per account.
RegisterGameServer	40	20	No	Limit is per account.
RequestUploadCredentials	4	1	No	Limit is per account.
ResolveAlias	4	1	No	Limit is per account.
ResumeGameServerGroup	1	0.1	No	Limit is per account.
SearchGameSessions	10	5	No	Limit is per account.
StartFleetActions	1	0.1	No	Limit is per account.

API action	Burst limit	Rate limit	Resource-level throttling	Notes
StartGameSessionPlacement	75	75	No	Limit is per account.
StartMatchBackfill	70	70	No	Limit is per account.
StartMatchmaking	100	100	No	Limit is per account.
StopFleetActions	1	0.1	No	Limit is per account.
StopGameSessionPlacement	20	20	No	Limit is per account.
StopMatchmaking	100	100	No	Limit is per account.
SuspendGameServerGroup	1	0.1	No	Limit is per account.
TagResource	20	5	No	Limit is per account.
TerminateGameSession	10	2	No	Limit is per account.
UntagResource	20	5	No	Limit is per account.
UpdateAlias	4	1	No	Limit is per account.
UpdateBuild	10	1	No	Limit is per account.
UpdateContainerGroupDefinition	10	1	No	Limit is per account.
UpdateContainerFleet	4	1	No	Limit is per account.

API action	Burst limit	Rate limit	Resource-level throttling	Notes
UpdateFleetAttributes	4	1	No	Limit is per account.
UpdateFleetCapacity	4	1	No	Limit is per account.
UpdateFleetPortSettings	4	1	No	Limit is per account.
UpdateGameServer	40	20	No	Limit is per account.
UpdateGameServerGroup	5	0.1	No	Limit is per account.
UpdateGameSession	10	2	No	Limit is per account.
UpdateGameSessionQueue	4	1	No	Limit is per account.
UpdateMatchmakingConfiguration	4	1	No	Limit is per account.
UpdateRuntimeConfiguration	4	1	No	Limit is per account.
UpdateScript	10	1	No	Limit is per account.
ValidateMatchmakingRuleSet	4	1	No	Limit is per account.

Resource-level throttling

Some API operations are subject to resource-level throttling to prevent database hot key issues. While increasing API-level limits for these operations, they may still be throttled at the resource level.

The following API operations are subject to resource-level throttling:

- DescribeGameSessionPlacement
- DescribeGameSessions
- DescribeMatchmaking
- DescribePlayerSessions

Resource-level throttling is enforced using specific throttle keys that include resource identifiers. For example, DescribeGameSessionPlacement uses the key "Operation:GameLift/DescribeGameSessionPlacement,aws-account-placement-id:" which enforces throttling per account and placement ID combination. This prevents customers from overwhelming the system with requests for specific resources.

Amazon GameLift Servers SDK API limits

The following table lists the default rate limits for Amazon GameLift Servers SDK API operations. These limits optimize performance and prevent resource contention in game server environments. Understanding these limits is important for efficient server-side integration with Amazon GameLift Servers.

Server SDK limits are enforced using specific throttling keys that include account and process identifiers. Some limits are per-player to prevent abuse from individual players.

Note

This table is subject to change. The authoritative source for current limits is the internal SDC configuration. Contact Amazon Support for the most up-to-date information if needed.

Server SDK API limits reference

API action	Burst limit	Rate limit	Throttling key	Notes
AcceptPlayerSession	3	1	aws-account-and-process-id-and-player-id	Limit is per player and per process. Should never need a limit increase as it's per-player.
RemovePlayerSession	3	1	aws-account-and-process-id-and-player-id	Limit is per player and per process. Should never need a limit increase as it's per-player.
ActivateGameSession	3	1	aws-account-and-process-id	Limit is per server process. Should never need a limit increase as it's per-process.
ActivateHostProcessV2	3	1	aws-account-and-process-id	Limit is per server process. Should never need a limit increase as it's per-process.
StartMatchBackfill	3	1	aws-account-and-process-id	Limit is per server process. Should never need a limit increase as it's per-process.
DescribePlayerSessions	10	5	aws-account-and-process-id	Limit is per server process. Should never need a limit increase as it's per-process.
GetComputeCertificate	3	0.1	aws-account-and-process-id	Limit is per server process. Should never need a limit increase as it's per-process.
GetCustomRoleCredentials	3	0.1	aws-account-and-process-id	Limit is per server process. Should never need a limit increase as it's per-process.

API action	Burst limit	Rate limit	Throttling key	Notes
InitSDK	3	0.1	aws-account-and-process-id	Limit is per server process. Should never need a limit increase as it's per-process.
RefreshHostProcess	3	0.1	aws-account-and-process-id	Limit is per server process. Should never need a limit increase as it's per-process.
StopMatchmaking	3	1	aws-account-and-process-id	Limit is per server process. Should never need a limit increase as it's per-process.
ProcessEnding	3	1	aws-account-and-process-id	Limit is per server process. Should never need a limit increase as it's per-process.
ProcessReady	3	1	aws-account-and-process-id	Limit is per server process. Should never need a limit increase as it's per-process.
UpdateGameSession	3	1	aws-account-and-process-id	Limit is per server process. Should never need a limit increase as it's per-process.
UpdateHostProcess	3	1	aws-account-and-process-id	Limit is per server process. Should never need a limit increase as it's per-process.
GetGameSessionId	10	5	aws-account-and-process-id	Limit is per server process. Should never need a limit increase as it's per-process.
GetTerminationTime	3	1	aws-account-and-process-id	Limit is per server process. Should never need a limit increase as it's per-process.

API action	Burst limit	Rate limit	Throttling key	Notes
UpdatePlayerSessionCreationPolicy	3	1	aws-account-and-process-id	Limit is per server process. Should never need a limit increase as it's per-process.

Note

These are the default limits. If you're experiencing throttling on these APIs, it may be due to account-level throttling. To request a limit increase, contact Amazon Support.

Per-process and per-player limits are designed to accommodate normal usage patterns and should rarely need adjustment. If you're experiencing throttling with these APIs, review your implementation for potential optimization opportunities before requesting limit increases.

Account-level throttling may also apply through these additional rules:

- SdkWebSocket/AccountLevelHighUsage,aws-account
- SdkWebSocket/AccountLevelThrottle,aws-account

Amazon GameLift Servers release notes

The Amazon GameLift Servers release notes provide details about new features, updates, and fixes related to the service.

SDK versions

The following tables list all Amazon GameLift Servers releases with SDK version information. There is no requirement to use comparable SDKs for your game server and client integrations. However, earlier versions of one SDK may not fully support the latest features in another.

Amazon GameLift Servers SDKs and plugins are open source. See [Get Amazon GameLift Servers development tools](#). To get the latest versions, see the [Amazon GameLift Servers GitHub organization](#).

Current version

1. **High Time**
 2. **Fast**
 3. **Unl**
 4. **Unl**
 5. **Unl**
 6. **Unl**
 7. **Unl**
 8. **Unl**
 9. **Unl**
 10. **Unl**
 11. **Unl**
 12. **Unl**
 13. **Unl**
 14. **Unl**
 15. **Unl**
 16. **Unl**
 17. **Unl**
 18. **Unl**
 19. **Unl**
 20. **Unl**
 21. **Unl**
 22. **Unl**
 23. **Unl**
 24. **Unl**
 25. **Unl**
 26. **Unl**
 27. **Unl**
 28. **Unl**
 29. **Unl**
 30. **Unl**
 31. **Unl**
 32. **Unl**
 33. **Unl**
 34. **Unl**
 35. **Unl**
 36. **Unl**
 37. **Unl**
 38. **Unl**
 39. **Unl**
 40. **Unl**
 41. **Unl**
 42. **Unl**
 43. **Unl**
 44. **Unl**
 45. **Unl**
 46. **Unl**
 47. **Unl**
 48. **Unl**
 49. **Unl**
 50. **Unl**
 51. **Unl**
 52. **Unl**
 53. **Unl**
 54. **Unl**
 55. **Unl**
 56. **Unl**
 57. **Unl**
 58. **Unl**
 59. **Unl**
 60. **Unl**
 61. **Unl**
 62. **Unl**
 63. **Unl**
 64. **Unl**
 65. **Unl**
 66. **Unl**
 67. **Unl**
 68. **Unl**
 69. **Unl**
 70. **Unl**
 71. **Unl**
 72. **Unl**
 73. **Unl**
 74. **Unl**
 75. **Unl**
 76. **Unl**
 77. **Unl**
 78. **Unl**
 79. **Unl**
 80. **Unl**
 81. **Unl**
 82. **Unl**
 83. **Unl**
 84. **Unl**
 85. **Unl**
 86. **Unl**
 87. **Unl**
 88. **Unl**
 89. **Unl**
 90. **Unl**
 91. **Unl**
 92. **Unl**
 93. **Unl**
 94. **Unl**
 95. **Unl**
 96. **Unl**
 97. **Unl**
 98. **Unl**
 99. **Unl**
 100. **Unl**
 101. **Unl**
 102. **Unl**
 103. **Unl**
 104. **Unl**
 105. **Unl**
 106. **Unl**
 107. **Unl**
 108. **Unl**
 109. **Unl**
 110. **Unl**
 111. **Unl**
 112. **Unl**
 113. **Unl**
 114. **Unl**
 115. **Unl**
 116. **Unl**
 117. **Unl**
 118. **Unl**
 119. **Unl**
 120. **Unl**
 121. **Unl**
 122. **Unl**
 123. **Unl**
 124. **Unl**
 125. **Unl**
 126. **Unl**
 127. **Unl**
 128. **Unl**
 129. **Unl**
 130. **Unl**
 131. **Unl**
 132. **Unl**
 133. **Unl**
 134. **Unl**
 135. **Unl**
 136. **Unl**
 137. **Unl**
 138. **Unl**
 139. **Unl**
 140. **Unl**
 141. **Unl**
 142. **Unl**
 143. **Unl**
 144. **Unl**
 145. **Unl**
 146. **Unl**
 147. **Unl**
 148. **Unl**
 149. **Unl**
 150. **Unl**
 151. **Unl**
 152. **Unl**
 153. **Unl**
 154. **Unl**
 155. **Unl**
 156. **Unl**
 157. **Unl**
 158. **Unl**
 159. **Unl**
 160. **Unl**
 161. **Unl**
 162. **Unl**
 163. **Unl**
 164. **Unl**
 165. **Unl**
 166. **Unl**
 167. **Unl**
 168. **Unl**
 169. **Unl**
 170. **Unl**
 171. **Unl**
 172. **Unl**
 173. **Unl**
 174. **Unl**
 175. **Unl**
 176. **Unl**
 177. **Unl**
 178. **Unl**
 179. **Unl**
 180. **Unl**
 181. **Unl**
 182. **Unl**
 183. **Unl**
 184. **Unl**
 185. **Unl**
 186. **Unl**
 187. **Unl**
 188. **Unl**
 189. **Unl**
 190. **Unl**
 191. **Unl**
 192. **Unl**
 193. **Unl**
 194. **Unl**
 195. **Unl**
 196. **Unl**
 197. **Unl**
 198. **Unl**
 199. **Unl**
 200. **Unl**
 201. **Unl**
 202. **Unl**
 203. **Unl**
 204. **Unl**
 205. **Unl**
 206. **Unl**
 207. **Unl**
 208. **Unl**
 209. **Unl**
 210. **Unl**
 211. **Unl**
 212. **Unl**
 213. **Unl**
 214. **Unl**
 215. **Unl**
 216. **Unl**
 217. **Unl**
 218. **Unl**
 219. **Unl**
 220. **Unl**
 221. **Unl**
 222. **Unl**
 223. **Unl**
 224. **Unl**
 225. **Unl**
 226. **Unl**
 227. **Unl**
 228. **Unl**
 229. **Unl**
 230. **Unl**
 231. **Unl**
 232. **Unl**
 233. **Unl**
 234. **Unl**
 235. **Unl**
 236. **Unl**
 237. **Unl**
 238. **Unl**
 239. **Unl**
 240. **Unl**
 241. **Unl**
 242. **Unl**
 243. **Unl**
 244. **Unl**
 245. **Unl**
 246. **Unl**
 247. **Unl**
 2

Previous versions

Starting time
Offset
Unit
Go
Unity
+
Unreal
5012025-1
or
later
501204-2
or
later
5012011-3
or
later
501208-2
or
later
5012010-2
or
later
501206-0
or
later
501303-2
or
later

Setting Time
Offset
Unit

Go
Unity
+
Unreal

541303-1
or
later

541302-0
or
later

541302-2
or
later

541301-2
or
later

541301-1
or
later

541309-1
or
later

541308-2
or
later

Partying Time

Abstract

Udacity

Community

+

Unreal

~~SECRET~~ 4-1

or

ater

3.74.000

or

ater

3.74.00

er

ater

3.74.0001-1

er

ater

3.74.000-2

or

ater

3.74.00

er

ater

~~330007-0~~

or

ater

Setting Time

Asset

Unity

Go

Unity

+

Unreal

[63004-2](#)

or
later

[63003-0](#)

or
later

[63002-0](#)

or
later

[63001 2-1](#)

or
later

[62809-2](#)

or
later

[62806-1](#)

or
later

[62805-1](#)

or
later

Setting Time

Platform

Unity

Go

Unity

+

Unreal

[3.2.102-1](#)

or
later

[3.2.102-0](#)

or
later

[3.1.709-0](#)

or
later

[3.1.708-1](#)

or
later

[3.1.505-1](#)

or
later

[3.1.504-1](#)

or
later

[3.1.502-2](#)

or
later

	Fortnite
	Fortnite
	Fortnite
	Fortnite
	Fortnite
	Fortnite
	Fortnite
	Fortnite
	Fortnite
	Fortnite
	Fortnite
	Fortnite
	Fortnite
	Fortnite
	Fortnite
	Fortnite
	Fortnite
	Fortnite

Release notes

The following release notes are in chronological order, with the latest updates listed first. Amazon GameLift Servers was first released in 2016. For release notes dated earlier than those listed here, see the release date links in [SDK versions](#).

January 29, 2026: Amazon GameLift Servers now supports scaling Fleets to/from zero instances

Amazon GameLift Servers now supports automatic scaling to and from zero instances based on game session activity. Fleets scale down to zero following a defined period of no game session activity and scale up from zero when game sessions are requested, providing an option for cost optimization.

Amazon GameLift Servers API Reference links:

- [Amazon SDK action UpdateFleetCapacity](#)

Learn more:

- [Manage Scaling an Amazon GameLift Servers Fleet To/From Zero](#)

January 22, 2026: Amazon GameLift Servers Realtime now supports Node.js 24.x on Amazon Linux 2023

Amazon GameLift Servers Realtime now supports Node.js 24.x on Amazon Linux 2023. This update provides access to the latest Node.js features and security improvements while running on the modern Amazon Linux 2023 operating system.

You can now also include an optional install script when creating a Realtime fleet, allowing the installation of software during instance startup for your specific requirements. To use the new runtime version, specify Node.js 24.x when creating Realtime scripts.

Amazon GameLift Servers API Reference links:

- [Amazon SDK action CreateScript](#)

Learn more:

- [Node.js runtime versions](#)
- [Add an install script](#)

November 20, 2025: Amazon GameLift Servers now supports game builds that use the Amazon Linux 2023 operating system

Amazon GameLift Servers customers can now use the Amazon Linux 2023 operating system to host their game servers. AL2023 offers several improvements over AL2 including security.

Customers can use the newer Linux operating systems and continue to receive critical security updates when support ends for Amazon Linux 2 in June 2026.

Amazon GameLift Servers API Reference links:

- [Amazon SDK action CreateBuild](#)
- [CLI command upload-build](#)
- [CLI command create-build](#)

March 6, 2025: Amazon GameLift is now Amazon GameLift Servers and Amazon GameLift Streams

Amazon GameLift is a fully-managed service that's dedicated to helping developers build, scale, and deliver the world's most demanding games. With the general availability release of Amazon GameLift Streams, Amazon GameLift now offers both high-scaling game server and smooth-gameplay streaming capabilities.

[Amazon GameLift Servers](#) gives game developers the capability to deploy, operate, and scale dedicated game servers. You can deploy high-performance game servers in the cloud in minutes and scale up and down to meet player demand. Built on Amazon proven computing environment, Amazon GameLift Servers supports 100 million concurrent players in a single game, 100 thousand player adds per second, and 9 thousand new compute instances per minute. And with enterprise-grade security, matchmaking for the largest crowds, and pay-as-you-go flexibility, it helps you get started whether you're working on a new game idea or running a game with millions of players.

[Amazon GameLift Streams](#) helps game developers deliver game streaming experiences at up to 1080p resolution and 60 frames-per-second (fps) with no perceivable lag across devices including iOS, Android, FireOS, and PCs for gamers. Using a single Amazon offering, publishers can deploy their game content in minutes, without modifications, onto fully-managed cloud-based GPU instances and deliver them through the Amazon Network Backbone directly to any end-player device with a web browser. Players can start gaming in just a few seconds, without waiting for a download or an install and it delivers a gameplay experience that is nearly indistinguishable from playing the game locally on a PC or gaming console.

September 5, 2024: Amazon GameLift Servers improves observability of the fleet creation process

Based on customer feedback, we've clarified the Amazon GameLift Servers workflow for creating a managed EC2 fleet and getting it ready to host game sessions. Improvements include:

- We've provided more specific and accurate descriptions of each phase of the fleet creation process. This improved visibility makes it easier to pinpoint and resolve issues faster.

- The Building and Activating phases better separate instance deployment tasks (building) from tasks to start game server processes and connect to the Amazon GameLift Servers service (activating). This change makes it easier to recognize the likely cause of issues. In addition, you can now remotely connect to fleets when they're in the Activating phase.
- Two new fleet creation events communicate the success or failure of game server install scripts. If your game server build includes an install script, Amazon GameLift Servers attempts to run the script and emits one of the following new events:
 - FLEET_CREATION_COMPLETED_INSTALLER
 - FLEET_CREATION_FAILED_INSTALLER

Learn more:

- [Debug Amazon GameLift Servers fleet issues](#), *Amazon GameLift Servers Developer Guide*
- [Event data type](#), *Amazon GameLift Servers API Reference*

July 02, 2024: Amazon GameLift Servers releases new console tool to view player session data

The Amazon GameLift Servers console now offers a player session lookup tool that lets you retrieve player session information by game session ID, player session ID, or player ID. Games that use FlexMatch matchmaking automatically generate player sessions for every matched player. For all other games, player sessions are an optional feature.

You can find the player session lookup tool in the main navigation for the Amazon GameLift Servers console. View individual player sessions or compare data across multiple player sessions. You can also open player session data when viewing a game session detail page.

Learn more:

- [Game and player sessions in the Amazon GameLift Servers console](#), *Amazon GameLift Servers Developer Guide*

February 1, 2024: Amazon GameLift Servers launches support for Infrastructure as Code (IaC) tools like Terraform and Pulumi powered by Amazon Cloud Control API

You can now manage your entire Amazon GameLift Servers resource stack using Infrastructure as Code (IaC) tools. These tools include Amazon CloudFormation, and also third-party tools such as Terraform and Pulumi. With this added support, you can now focus on building your game, and leverage dev-Ops strategies to take care of resource management, continuous integration and deployment to your customers.

You can also now provision and configure all Amazon GameLift Servers resources types by using the Amazon Cloud Control API. You can continue to work with resources using the Amazon GameLift Servers APIs or the Amazon CloudFormation templates for Amazon GameLift Servers.

For details about the Amazon GameLift Servers resources available through IaC, see the https://docs.amazonaws.cn/AWSCloudFormation/latest/UserGuide/AWS_GameLift.html.

In addition, you can now automatically scale your fleets using Amazon CloudFormation templates or the Amazon Cloud Control API by using the new [Fleet](#) property: `ScalingPolicies`.

The Cloud Control API gives developers a standard set of APIs to create, read, update, delete, and list resources (CRUDL) across hundreds of Amazon services and multiple third-party tools like Terraform and Pulumi.

Learn more:

- [Amazon CloudFormation](#)
- [Amazon Cloud Control API](#)
- [Amazon CC Terraform Provider](#)
- [Pulumi](#)

December 14, 2023: Amazon GameLift Servers adds ability to update the game properties of active game sessions

You've already been able to set game properties when creating game sessions, and to search game sessions for specified properties. Now you can also add and update these properties in an active game session.

For example, your players vote on a map that they want to play on. Your game client calls `UpdateGameSession` to modify a `GameProperty` value to `{"Key": "map", "Value": "jungle"}`. Your game then implements the new map for the players in the game session.

Game administrators can also retrieve useful data from game properties by using the `SearchGameSessions` operation. For example, administrators can list game sessions that have a `Status` value of `ACTIVE` and this game property: `{"Key": "map", "Value": "desert"}`.

Learn more:

- [the section called "Integrate game client functionality"](#), *Amazon GameLift Servers Developer Guide*
- [GameProperty](#), *Amazon GameLift Servers API Reference*
- [UpdateGameSession](#), *Amazon GameLift Servers API Reference*
- [SearchGameSessions](#), *Amazon GameLift Servers API Reference*

July 20, 2023: Amazon GameLift Servers launches a new console experience

The new Amazon GameLift Servers console is now the default experience for customers in all regions, including China. It includes these improvements:

- **Improved navigation** – The new navigation pane facilitates navigation between Amazon GameLift Servers resources.
- **Amazon GameLift Servers landing page** – The new landing page provides links to helpful documentation, displays a high-level overview of Amazon GameLift Servers, and provides support through links to documentation, frequently asked questions, and Amazon Web Services re:Post.
- **Improved Amazon CloudWatch metrics** – Amazon GameLift Servers metrics are now available in both the Amazon GameLift Servers console and your CloudWatch dashboards. This update also includes new metrics for performance, utilization, hardware, and player sessions.

July 13, 2023: Amazon GameLift Servers adds fleet hardware metrics

You can now track hardware performance metrics for your Amazon GameLift Servers managed EC2 fleets. Metrics include EC2 instance metrics for CPU utilization, network traffic volume, and disk read/write activity. For Amazon GameLift Servers, these metrics describe all active instances

in a fleet location. You can view these fleet hardware metrics using an Amazon CloudWatch dashboard in the Amazon Web Services Management Console. You can also view them in the Amazon GameLift Servers console in fleet details.

Learn more:

- [Monitor Amazon GameLift Servers with Amazon CloudWatch](#) (Metrics for fleets), *Amazon GameLift Servers Developer Guide*

May 25, 2023: Amazon GameLift Servers FleetIQ adds filter to exclude game session placements on draining instances

Updated SDK versions: Amazon SDK 1.11.87

If you use Amazon GameLift Servers FleetIQ for game hosting, you can now prevent game session placements on instances that are currently draining. Draining instances are flagged for shutdown, but they can still be selected to host new game sessions if no other hosting resources are available. With this new feature, you can exclude the use of draining instances entirely.

Use this feature when calling `ClaimGameServer` to find available game servers. Add the new `FilterOption` parameter and set allowed instance statuses to `ACTIVE` only. In response, Amazon GameLift Servers FleetIQ looks only at active instances when searching for and claiming an available game server.

Learn more:

- [ClaimGameServer](#) in the *Amazon GameLift Servers API Reference*
- [How FleetIQ works](#) in the *Amazon GameLift Servers FleetIQ Developer Guide*

May 16, 2023: Amazon GameLift Servers supports cost allocation tagging for fleets

Amazon GameLift Servers customers can now use Amazon Billing cost allocation tags to organize their game hosting costs. You can assign cost allocation tags to individual Amazon GameLift Servers EC2 fleet resources to track how your fleets are contributing to the overall hosting costs.

Learn more:

- [Resource cost and utilization tools](#)

- [Using Amazon cost allocation tags](#), *Amazon Billing User Guide*

April 20, 2023: Amazon GameLift Servers launches support for Windows Server 2016

Updated SDK versions: Amazon SDK 1.11.63

Amazon GameLift Servers customers can now use the Windows Server 2016 operating system to host their game servers. This operating system is available in all Amazon Web Services Regions. Customers can use the newer Windows operating system and continue to receive critical security updates as Microsoft ends its support for Windows Server 2012 in October 2023.

Starting today, new customers who require a Windows runtime environment must specify Windows Server 2016 when creating new game server builds for hosting. Existing customers can continue to create new builds and fleets with Windows Server 2012 but must complete migration with Windows Server 2016 before the Microsoft end of support date on October 10, 2023.

This update includes the following service changes:

- When creating a game server build using Amazon GameLift Servers SDK or CLI commands, you must now explicitly set the operating system. There is no longer a default value. To deploy your game server on Windows Server 2016, use the value `WINDOWS_2016`.
- When creating a game server build using the Amazon GameLift Servers console, you must select an operating system from the available values. If you're an existing customer with active Windows Server 2012 fleets, you can choose either `WINDOWS_2012` or `WINDOWS_2016`.

Learn more:

- Amazon GameLift Servers API Reference links:
 - [CLI command `upload-build`](#)
 - [CLI command `create-build`](#)
 - [Amazon SDK action `CreateBuild`](#)
- [Amazon GameLift Servers FAQ for Windows 2012](#)

March 14, 2023: Amazon GameLift Servers launches a new console experience

The new Amazon GameLift Servers console includes these improvements:

- **Improved navigation** – The new navigation pane facilitates navigation between Amazon GameLift Servers resources.
- **Amazon GameLift Servers landing page** – The new landing page provides links to helpful documentation, displays a high-level overview of Amazon GameLift Servers, and provides support through links to documentation, frequently asked questions, and Amazon Web Services re:Post.
- **Improved Amazon CloudWatch metrics** – Amazon GameLift Servers metrics are now available in both the Amazon GameLift Servers console and your CloudWatch dashboards. This update also includes new metrics for performance, utilization, and player sessions.

Learn more:

- [Manage game hosting resources with Amazon GameLift Servers](#)
- [Building a FlexMatch matchmaker](#)

February 14, 2023: Amazon GameLift Servers now supports server side encryption for Amazon SNS topics

Server Side Encryption ((SSE)) for SNS topics encrypts your sensitive data at rest. SSE uses Amazon Key Management Service (Amazon KMS) keys to protect the contents of your SNS topics.

Learn more:

- [Set up event notification for game session placement](#)
- [FlexMatch matchmaking events](#)
- [Encryption at rest](#)

January 31, 2023: Amazon GameLift Servers server SDK supports the Go language

Updated SDK versions: Server SDK 5.0.0 for Go

Learn more:

- Download the latest version of the Amazon GameLift Servers server SDK at [Amazon GameLift Servers getting started](#)

- [???](#)

June 28, 2022: Amazon GameLift Servers launches a new opt-in console experience

The new Amazon GameLift Servers console includes these improvements:

- **Improved navigation** – The new navigation pane facilitates navigation between Amazon GameLift Servers resources.
- **Amazon GameLift Servers landing page** – The new landing page provides links to helpful documentation, displays a high-level overview of Amazon GameLift Servers, and provides support through links to documentation, frequently asked questions, and Amazon Web Services re:Post.
- **Improved Amazon CloudWatch metrics** – Amazon GameLift Servers metrics are now available in both the Amazon GameLift Servers console and your CloudWatch dashboards. This update also includes new metrics for performance, utilization, and player sessions.

Learn more:

- [Manage game hosting resources with Amazon GameLift Servers](#)
- [Building a FlexMatch matchmaker](#)

February 15, 2022: FlexMatch adds compound rule and additional improvements

FlexMatch users now have access to the following features:

- **Compound rule** – Added support for compound matchmaking rules for matches of 40 or fewer players. You can now use logical statements to create a compound rule to form a match. Without a compound rule in your rule set, to form a match, all the rules in the rule set must be true. With compound rules, you can choose which rules to apply using the following logical operators: `and`, `or`, `not`, and `xor`.
- **Flexible team selection** – Updated matchmaking property expressions to support selecting a subset of all available teams.
- **Longer string lists** – Increased the maximum number of strings from 10 to 100 in a list of strings of player attribute values.

Learn more:

- [Amazon GameLift Servers FlexMatch developer guide](#):
 - [FlexMatch rule types](#)
 - [FlexMatch property expressions](#)
- [AttributeValue: SL](#)

October 28, 2021: Amazon GameLift Servers adds support for multi-Region fleets in the Asia Pacific (Osaka) Region; Amazon GameLift ServersFleetIQ adds support for Amazon Graviton2 processors

Updated SDK versions: Amazon SDK [1.9.133](#)

Amazon GameLift Servers is now available in the Asia Pacific (Osaka) Region. Game developers can now deploy instances in Osaka using GameLift multi-Region fleet.

You can now use Graviton2-hosted game servers, based on the Arm-based processor architecture, to achieve increased performance at a lower cost when compared to the equivalent Intel-based compute options.

Highlights:

- Amazon GameLift Servers is now available in the Asia Pacific (Osaka) Region.
- Amazon GameLift Servers FleetIQ game server groups can now be configured to manage the Graviton2 instance families c6g, m6g, and r6g.

Learn more:

- [Amazon GameLift Servers multi-Region fleet](#)
- [CreateGameServerGroup](#)
- [Amazon graviton processor](#)

September 20, 2021: Amazon GameLift Servers releases plugin for Unity

The Amazon GameLift Servers plugin for Unity version 1.0.0 contains libraries and native UI that makes it easier to access Amazon GameLift Servers resources and integrate Amazon GameLift Servers into your Unity game. You can use the Amazon GameLift Servers plugin for Unity to

access Amazon GameLift Servers APIs and deploy Amazon CloudFormation templates for common gaming scenarios. The plugin also includes a sample game that works with the sample scenarios. You can use Amazon GameLift Servers Local to see messages passed between the game client and the game server to learn how a typical game interacts with Amazon GameLift Servers.

The plugin for Unity supports Unity 2019.4 LTS and 2020.3 LTS.

Highlights:

- Build, run, and modify a sample game with different scenarios, or create your own.
- Deploy sample Amazon CloudFormation scenarios for typical game scenarios including auth only, single-Region fleet, multi-Region fleets with queue and custom matchmaker, Spot Fleets with queue and custom matchmaker, and FlexMatch.

Learn more:

- [Integrating games with the Amazon GameLift Servers plugin for Unity](#)

June 30, 2021: FlexMatch adds batchDistance rule

You can use the batchDistance rule type to specify a string or numeric attribute, bringing a host of benefits to each segment.

Highlights:

- For large matches (>40 players), instead of evenly balancing players by skill only, you can now get that same balance based on skill, modes, and maps. Ensure that everyone in the match is in a skill band, band multiple numeric attributes such as league or play style, and group according to string attributes such as map or game mode. You can also create expansions over time. For example, you can create an expansion to allow a greater skill level range to enter the match the longer the player is waiting.

For matches under 40 players, you can use a new simplified rules expression.

June 3, 2021: Amazon GameLift Servers realtime client SDK and server SDK updates

Updated SDK versions: Realtime Client SDK 1.2.0, Server SDK 3.4.0 for Unreal

With this latest SDK update, you can now integrate IL2CPP into your mobile applications that use the RTS Client SDK and follow best practices with frameworks. You can also now build the Amazon GameLift Servers Server SDK for Unreal Version 4.26. This update contains components that integrate with your Windows or Linux game server, including C++ and C# versions of the Amazon GameLift Servers Server SDK, Amazon GameLift Servers Local, and an Unreal Engine plugin.

Highlights:

- Added support for IL2CPP in the RTS Client SDK and for building the native libraries as frameworks, so you can build RTS clients for the latest mobile devices.
- You can use [DescribePlayerSessions\(\)](#) to get information for a single player session, for all player sessions in a game session, or for all player sessions associated with a single player ID.
- Created Server SDK support for Unreal version 4.26.
- The existing C# SDK, version 4.0.2, has been verified compatible with Unity 2020.3. No SDK updates were required.

Learn more:

- [Amazon GameLift Servers Developer Guide:](#)
 - [DescribePlayerSessions\(\)](#)

March 23, 2021: Amazon GameLift Servers adds notifications to game session placement

Updated SDK versions: Amazon SDK [1.8.168](#)

You can now use events to monitor game session placement activity for a game session queue. Create an Amazon Simple Notification Service (Amazon SNS) topic to publish event notifications, or set up event tracking using CloudWatch Events.

Highlights:

- For each queue, you can set a custom text string to be included in all event messaging.
- When using an Amazon SNS topic, you can set additional access conditions that limit publishing to specific queues.

Learn more:

- Amazon GameLift Servers Developer Guide:
 - [Set up event notification for game session placement](#) (new)
 - [Game session placement events](#) (new)
- [API reference \(Amazon SDK\)](#)
 - New game session queue parameters `NotificationTarget` and `CustomEventData`: [GameSessionQueue](#), [CreateGameSessionQueue](#), [UpdateGameSessionQueue](#)
- [Amazon GameLift Servers forum](#)

March 16, 2021: Amazon GameLift Servers adds multi-region fleets, six new regions

Updated SDK versions: Amazon SDK [1.8.163](#)

Amazon GameLift Servers managed hosting is now available in 21 Amazon Regions. The new Regions are Cape Town (af-south-1), Bahrain (me-south-1), Hong Kong (ap-east-1), Milan (eu-south-1), Paris (eu-west-3), and Stockholm (eu-north-1).

With the new Amazon GameLift Servers multi-location fleets feature, you can now set up a single fleet to host your game servers in any or all of 20 Amazon GameLift Servers-supported Regions (Beijing Region excepted). This feature aims to significantly reduce the work required to set up and maintain Amazon GameLift Servers hosting resources globally. Multi-location fleets can be created in the following Amazon Regions: us-east-1 (N. Virginia), us-west-2 (Oregon), eu-central-1 (Frankfurt), eu-west-1 (Ireland), ap-southeast-2 (Sydney), ap-northeast-1 (Tokyo), and ap-northeast-2 (Seoul). In all other Regions, you can continue to set up single-location fleets as needed. All fleets that were created before this release are single-location fleets. Using multi-location fleets does not affect your hosting costs. Amazon GameLift Servers pricing is based on the type, location, and volume of instances that you use. (For more information, see [Amazon GameLift Servers pricing](#).) Amazon CloudFormation support for multi-location fleets will be available soon.

Note

Multi-location fleets are not available in the China Regions. Amazon GameLift Servers resources that reside in China Regions cannot interact with or be used by resources in other Amazon GameLift Servers Regions.

Highlights:

- With a multi-location fleet, explicitly add a list of remote locations. Amazon GameLift Servers deploys instances of the same type and configuration, including the build and runtime configuration, to the fleet's home Region and all added locations.
- Adjust capacity settings and scaling for each location independently. Auto scaling policies apply to an entire fleet, but you can turn them on or off by location.
- Start new game sessions at specific fleet locations. When using game session queues or matchmaking to place game sessions, you can now prioritize where new game sessions start by location, hosting cost, and player latency.
- Get hosting metrics in the Amazon GameLift Servers console, aggregated for all locations in a fleet or broken out by each fleet location.

Learn more:

- [Amazon game tech blog](#)
- [API reference \(Amazon SDK\)](#)
 - New fleet location operations: [CreateFleetLocations](#), [DescribeFleetLocationAttributes](#), [DescribeFleetLocationCapacity](#), [DescribeFleetLocationUtilization](#), [DeleteFleetLocations](#)
 - Updated fleet operations, with new multi-location support: [CreateFleet](#), [UpdateFleetCapacity](#), [DescribeEC2InstanceLimits](#), [DescribeInstances](#), [StopFleetActions](#), [StartFleetActions](#)
 - Updated game session placement operations, with new priority and filtering capability: [CreateGameSessionQueue](#), [DescribeGameSessionQueues](#), [UpdateGameSessionQueue](#)
 - Updated game session creation operations, with new location support: [CreateGameSession](#), [DescribeGameSessions](#), [DescribeGameSessionDetails](#), [SearchGameSessions](#)
- [Amazon GameLift Servers Developer Guide](#):
 - [Amazon GameLift Servers service locations](#) (updated)
 - [Hosting resource customizations](#) (new)
 - [Scaling game hosting capacity with Amazon GameLift Servers](#) (updated)
 - [Customize a game session queue](#) (new)
 - [Fleet details in the Amazon GameLift Servers console](#) (updated)
- [Amazon GameLift Servers forum](#)

February 9, 2021: Amazon GameLift Servers extends support for AMD instances, standalone FlexMatch

Updated SDK versions: Amazon SDK [1.8.139](#)

This release includes the following updates:

- Amazon GameLift Servers FleetIQ game server groups can now be configured to manage the AMD instance families C5a, M5a, and R5a. The supported Amazon EC2 instance types, as listed for the GameServerGroup [InstanceDefinition](#), now include the following:
 - c5a.large, c5a.xlarge, c5a.2xlarge, c5a.4xlarge, c5a.8xlarge, c5a.12xlarge, c5a.16xlarge, c5a.24xlarge
 - m5a.large, m5a.xlarge, m5a.2xlarge, m5a.4xlarge, m5a.8xlarge, m5a.12xlarge, m5a.16xlarge, m5a.24xlarge
 - r5a.large, r5a.xlarge, r5a.2xlarge, r5a.4xlarge, r5a.8xlarge, r5a.12xlarge, r5a.16xlarge, r5a.24xlarge

Note: AMD instances for FleetIQ are currently not available for use in the China (Beijing) Amazon Region. See [Feature availability and implementation differences](#) in China.

- Amazon GameLift Servers managed game hosting now supports AMD instances in the China (Beijing) Region, operated by Sinnet. The new AMD instance families include M5a and R5a. Supported EC2 instance types, as listed for fleet [InstanceType](#), now include the following:
 - m5a.large, m5a.xlarge, m5a.2xlarge, m5a.4xlarge, m5a.8xlarge, m5a.12xlarge, m5a.16xlarge, m5a.24xlarge
 - r5a.large, r5a.xlarge, r5a.2xlarge, r5a.4xlarge, r5a.8xlarge, r5a.12xlarge, r5a.16xlarge, r5a.24xlarge
- Amazon GameLift Servers FlexMatch can now be used as a standalone matchmaking solution in the China (Beijing) Region, operated by Sinnet. Customers can create a FlexMatch matchmaker in the Beijing Region and configure the [FlexMatchMode](#) parameter to STANDALONE. For more information about FlexMatch, either with Amazon GameLift Servers managed hosting or with a non-Amazon GameLift Servers hosting solution, in the [Amazon GameLift Servers FlexMatch Developer Guide](#).
- When setting up event notifications for Amazon GameLift Servers FlexMatch, you can now designate an Amazon SNS FIFO topic as the notification target. For more information, see:
 - [MatchmakingConfiguration NotificationTarget](#), *Amazon GameLift Servers API Reference*
 - [Set up FlexMatch event notification](#), *Amazon GameLift Servers FlexMatch Developer Guide*

- [Introducing Amazon SNS FIFO – First-in-first-out Pub/Sub messaging](#), *Amazon News Blog*

December 22, 2020: Amazon GameLift Servers server SDK supports Unreal Engine 4.25 and Unity 2020

Updated SDK versions: Amazon GameLift Servers Server SDK 4.0.2, Unreal plugin version 3.3.3

The latest version of the Amazon GameLift Servers Server SDK contains the following components:

- The updated Unreal plugin has been updated for compatibility with Unreal Engine 4.25. The API was not changed.
- The existing C# SDK, version 4.0.2, has been verified compatible with Unity 2020. No SDK updates were required.

Download the latest version of the Amazon GameLift Servers Server SDK at [Amazon GameLift Servers getting started](#).

November 24, 2020: Amazon GameLift Servers FlexMatch now available for games hosted anywhere

Updated SDK versions: Amazon SDK [1.8.95](#)

Amazon GameLift Servers FlexMatch is a customizable matchmaking service for multiplayer games. Initially designed for users of Amazon GameLift Servers managed hosting, FlexMatch can now be integrated into games that use other hosting systems, including peer-to-peer, proprietary on-premises computing, and cloud compute primitive types. Games that use Amazon GameLift Servers FleetIQ for game hosting on Amazon EC2 can now implement matchmaking with FlexMatch.

FlexMatch provides a robust matchmaking algorithm and rules language that gives you wide latitude to customize the matchmaking process so that players are matched together based on key player characteristics and reported latency. In addition, FlexMatch offers a matchmaking request workflow that supports features such as player parties, player acceptance, and match backfill. When you use FlexMatch with Amazon GameLift Servers managed hosting or Amazon GameLift Servers Realtime, the matchmaker automatically uses Amazon GameLift Servers to find hosting resources and start a new game session for newly formed matches. When using FlexMatch as a standalone service, the matchmaker delivers match results back to your game, which can then start a new game session using your hosting solution.

API operations for FlexMatch are part of the Amazon GameLift Servers service API, which is included in the Amazon SDK and the Amazon Command Line Interface (Amazon CLI). This release includes these updates to support standalone matchmaking:

- The API resource `MatchmakingConfiguration` has the following changes:
 - New property, `FlexMatchMode` indicates whether the matchmaker is being used with Amazon GameLift Servers managed hosting or as standalone matchmaking.
 - Property `GameSessionQueueArns` is not required when `FlexMatchMode` is set to `standalone`.
 - These properties are not used with standalone matchmaking: `AdditionalPlayerCount`, `BackfillMode`, `GameProperties`, `GameSessionData`.
- The automatic backfill feature is not available with standalone matchmaking.

November 24, 2020: AMD instances now available on Amazon GameLift Servers

Updated SDK versions: Amazon SDK [1.8.95](#)

The list of Amazon EC2 instance types supported by Amazon GameLift Servers now includes three new instance families: C5a, M5a, and R5a. These families consist of AMD compute-optimized instances that are powered by AMD EPYC processors running at frequencies up to 3.3. GHz. The AMD instances are x86 compatible; games that are currently running on Amazon GameLift Servers can be deployed to AMD instance types without alteration. The new instances are available in the following Amazon Regions: US East (N. Virginia and Ohio), US West (Oregon and N. California), Central Canada (Montreal), South America (Sao Paulo), EU Central (Frankfurt), EU West (London and Ireland), Asia Pacific South (Mumbai), Asia Pacific Northeast (Seoul and Tokyo), and Asia Pacific Southeast (Singapore and Sydney).

The new AMD instances include:

- c5a.large, c5a.xlarge, c5a.2xlarge, c5a.4xlarge, c5a.8xlarge, c5a.12xlarge, c5a.16xlarge, c5a.24xlarge
- m5a.large, m5a.xlarge, m5a.2xlarge, m5a.4xlarge, m5a.8xlarge, m5a.12xlarge, m5a.16xlarge, m5a.24xlarge
- r5a.large, r5a.xlarge, r5a.2xlarge, r5a.4xlarge, r5a.8xlarge, r5a.12xlarge, r5a.16xlarge, r5a.24xlarge

Learn more:

- [Amazon game tech blog](#)
- [Amazon GameLift Servers instance pricing](#)
- [Amazon EC2 instances featuring AMD EPYC processors](#)
- [Amazon GameLift Servers forum](#)

November 11, 2020: Version update to Amazon GameLift Servers server SDK**Updated SDK versions:** Amazon GameLift Servers Server SDK 4.0.2

The new Server SDK version 4.0.2 fixes a known issue with the API operation `StartMatchBackfill()`. This operation now returns a correct response to a match backfill request.

The issue did not affect the match backfill process, and there is no change to how this feature works. The issue may have impacted log messaging and error handling for match backfill requests.

Download the latest version of the Amazon GameLift Servers Server SDK at [Amazon GameLift Servers getting started](#).

November 5, 2020: New FlexMatch algorithm customizations

FlexMatch users can now adjust the following default behaviors for the matchmaking process. These customizations are set in a matchmaking rule set. There are no changes to the Amazon GameLift Servers SDKs.

- **Prioritize backfill tickets:** You can choose to raise or lower how match backfill tickets are prioritized when searching for acceptable matches. Prioritizing backfill tickets is useful when the auto-backfill feature is enabled. Use the algorithm property `backfillPriority`.
- **Pre-sort to optimize match consistency and efficiency:** Configure your matchmaker to pre-sort the ticket pool before batching tickets for evaluation. By pre-sorting tickets based on key player attributes, your resulting matches tend to have players who are more similar in those attributes. You can also boost efficiency in the evaluation process by pre-sorting on the same attributes that are used in match rules. Use the algorithm property `sortByAttributes` with the `strategy` property set to "sorted".
- **Adjust how expansion wait times are triggered:** Choose between triggering expansions based on the age of the newest (default) or oldest ticket in an incomplete match. Triggering on the oldest

ticket tends to complete matches faster, while triggering on the newest ticket leads to higher match quality. Use the algorithm property `expansionAgeSelection`.

September 17, 2020: Amazon GameLift Servers updates server SDK

Updated SDK versions: Amazon GameLift Servers Server SDK 4.0.1

The new Server SDK contains the following updates:

- C# API version 4.0.1
 - The API operation [TerminateGameSession\(\)](#) is no longer supported. Replace with a call to [ProcessEnding\(\)](#) to end both a game session and the server process.
 - The operation [GetTerminationTime\(\)](#) now returns a value of data type `AwsDateTimeOutcome`.
- C++ API version 3.4.1
 - The operation [TerminateGameSession\(\)](#) is no longer supported. Replace it with a call to [ProcessEnding\(\)](#) to end both a game session and the server process.
- Unreal Engine plugin version 3.3.2
 - The operation [TerminateGameSession\(\)](#) is no longer supported. Replace it with a call to [ProcessEnding\(\)](#) to end both a game session and the server process.
 - The callback operation `OnUpdateGameSession` is added to [FProcessParameters](#) to support match backfill.

Download the latest version of the Amazon GameLift Servers Server SDK at [Amazon GameLift Servers getting started](#).

August 27, 2020: Amazon GameLift Servers FleetIQ for game hosting with Amazon EC2 (general availability)

Updated SDK versions: Amazon SDK [1.8.36](#)

The Amazon GameLift Servers FleetIQ solution for low-cost, cloud-based game hosting on Amazon EC2 is now generally available. Amazon GameLift Servers FleetIQ gives developers the ability to host game servers directly on Amazon EC2 Spot Instances by optimizing their viability for game hosting. Game developers can use Amazon GameLift Servers FleetIQ with new games or to supplement capacity for existing games. This solution supports the use of containers or other Amazon services such as Amazon Shield and Amazon Elastic Container Service (Amazon ECS).

This general availability release includes the following updates to the Amazon GameLift Servers FleetIQ solution:

- New API operation `DescribeGameServerInstances` returns information, including status, on all active instances for an Amazon GameLift Servers FleetIQ game server group.
- New balancing strategy, `ON_DEMAND_ONLY`, configures a game server group to use On-Demand Instances only. You can update a game server group's balancing strategy at any time, making it possible to switch between using Spot Instances and On-Demand Instances as needed.
- The following preview elements have been dropped for general availability:
 - Use of custom sort keys for game server resources. Game servers can be sorted based on registration timestamp.
 - Tagging for game server resources.

April 16, 2020: Amazon GameLift Servers updates server SDK for Unity and Unreal Engine

Updated SDK versions: Amazon GameLift Servers Server SDK 4.0.0, Amazon GameLift Servers Local 1.0.5

The latest version of the Amazon GameLift Servers Server SDK contains the following updated components:

- C# SDK version 4.0.0 updated for Unity 2019.
- Unreal plugin version 3.3.1 updated for Unreal Engine versions 4.22, 4.23, and 4.24.
- Amazon GameLift Servers Local version 1.0.5 updated to test integrations that use the C# server SDK version 4.0.0.

Download the latest version of the Amazon GameLift Servers Server SDK at [Amazon GameLift Servers getting started](#).

April 2, 2020: Amazon GameLift Servers FleetIQ available for game hosting on EC2 (public preview)

Updated SDK versions: Amazon SDK [1.7.310](#)

The Amazon GameLift Servers FleetIQ feature optimizes the viability of low-cost Spot Instances for use with game hosting. This feature is now extended for customers who want to manage their

hosting resources directly rather than through the managed Amazon GameLift Servers service. This solution supports the use of containers or other Amazon services such as Amazon Shield and Amazon Elastic Container Service (Amazon ECS).

Learn more:

[GameTech blog post](#) on Amazon GameLift Servers FleetIQ

December 19, 2019: Improved Amazon resource management for Amazon GameLift Servers resources

Updated SDK versions: Amazon SDK [1.7.249](#)

You can now take advantage of Amazon resource management tools with Amazon GameLift Servers resources. In particular, all key Amazon GameLift Servers resources—builds, scripts, fleets, game session queues, matchmaking configurations, and matchmaking rule sets—are now assigned Amazon Resource Name (ARN) values. A resource ARN provides a consistent identifier that is unique across all Amazon Regions. They can be used to create resource-specific Amazon Identity and Access Management (IAM) permissions policies. Resources are now assigned an ARN and also the pre-existing resource identifier, which is not Region-specific.

In addition, Amazon GameLift Servers resources now support tagging. You can use tags to organize resources, create IAM permissions policies to manage access to groups of resources, customize Amazon cost breakdowns, etc. When managing tags for Amazon GameLift Servers resources, use the Amazon GameLift Servers API actions `TagResource()`, `UntagResource()`, and `ListTagsForResource()`.

Learn more:

- [TagResource](#) in the *Amazon GameLift Servers API Reference*
- [Tagging Amazon resources](#) in the *Amazon General Reference*
- [Amazon resource names](#) in the *Amazon General Reference*

November 14, 2019: New Amazon CloudFormation templates, updates in China (Beijing) Region

Updated SDK versions: Amazon SDK [1.7.210](#)

Amazon CloudFormation templates for Amazon GameLift Servers

Amazon GameLift Servers resources can now be created and managed through Amazon CloudFormation. The existing Amazon CloudFormation build and fleet templates have been updated to align with the current resources, and new templates are now available for scripts, queues, matchmaking configurations, and matchmaking rule sets. Amazon CloudFormation templates greatly simplify the task of managing groups of related Amazon resources, particularly when deploying games across multiple Regions.

Learn more:

- [Amazon GameLift Servers resource type reference](#) in the *Amazon CloudFormation User Guide*
- [Manage Amazon GameLift Servers hosting resources using Amazon CloudFormation](#) in the *Amazon GameLift Servers Developer Guide*