Developer Guide

Amazon GameLift Streams

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

Amazon GameLift Streams: Developer Guide

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

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

Table of Contents

What is Amazon GameLift Streams? 1
Features 1
How to get started with Amazon GameLift Streams 1
Accessing Amazon GameLift Streams 2
Terms of use 3
Setting up 4
Sign up for an AWS account 4
Create a user with administrative access 5
Get programmatic access 6
Download the Web SDK 7
Download the AWS CLI 7
Set up billing alerts 7
Getting started 8
Choosing a configuration 8
Starting point 9
Cost optimizations 10
Deciding on a configuration 11
How your configuration choices impact next steps 12
Next steps 13
Configuration options 13
Runtime environments 13
Stream classes 14
Runtime environment and stream class compatibility 19
Your first stream 20
Prerequisites 21
Step 1: Upload your application to an Amazon S3 bucket 22
Step 2: Configure your application for Amazon GameLift Streams 24
Step 3: Manage how Amazon GameLift Streams streams your application 28
Step 4: Test your stream in Amazon GameLift Streams 34
Step 5: Clean up (don't skip) 35
Managing your streams 37
Key concepts 37
Applications 39
Before you upload 39

Upload your application to an Amazon S3 bucket ... 40
Create an application ... 41
Edit an application ... 47
Delete an application ... 50
Application log bucket permission policy ... 52
Linked stream groups ... 53
Stream groups ... 53
About stream capacity ... 54
Capacity and service quotas ... 56
About locations ... 58
Create a stream group ... 58
Edit general settings ... 69
Edit capacity ... 70
Capacity scale-down behavior ... 72
Add locations in a stream group ... 72
Remove locations in a stream group ... 74
Delete a stream group ... 75
Linked applications ... 76
Stream group lifecycle ... 77
Stream group maintenance ... 77
Multi-application stream groups ... 78
Limitations and requirements ... 78
About default applications ... 79
Change default application ... 80
Link an application ... 81
Unlink an application ... 83
Stream sessions ... 85
About stream sessions ... 85
Testing a stream in the console ... 85
Stream session lifecycle ... 86
Timeout values affecting stream sessions ... 89
Terminating a stream session ... 90
Reconnecting to a stream session ... 90
Export stream session files ... 91
How it works ... 91
Cost impact ... 92

Export files (Console) ... 92
Export files (CLI) ... 93
Amazon GameLift Streams backend service and web client ... 96
Supported browsers and input ... 96
Known issues ... 98
Limitations ... 98
IPv6 support ... 99
Required ports ... 99
Setting up a web server and client ... 99
Prerequisites ... 100
Download the Web SDK ... 100
Set up your streaming resources ... 100
Set up a backend server ... 101
Launch a web client ... 101
Clean up streaming resources ... 102
Customize stream appearance ... 104
Loading screen ... 104
Locale preference ... 104
Mouse movement handling ... 105
Mouse input modes ... 105
Pointer lock ... 106
Best practices ... 107
Data channel communication ... 108
Features ... 108
Using data channels ... 109
On the client-side ... 109
On the application-side ... 109
Launch checklist ... 113
Notify the Amazon GameLift Streams team ... 113
Compatibility and performance testing ... 113
Capacity reservation ... 113
Performance testing at scale ... 114
Pre-launch setup ... 114
Additional tips ... 114
Need Further Assistance? ... 114
Security ... 115

Data protection ... 116
Encryption at rest ... 117
Encryption in transit ... 118
Protection of end-user streams ... 118
Session isolation in Linux stream classes ... 118
Session isolation in Windows stream classes ... 119
Encryption key management ... 119
Inter-network traffic privacy ... 119
Identity and Access Management ... 120
Audience ... 120
Authenticating with identities ... 121
Managing access using policies ... 122
How Amazon GameLift Streams works with IAM ... 123
Identity-based policy examples ... 129
Troubleshooting ... 132
Compliance validation ... 133
Resilience ... 133
Infrastructure Security ... 133
Reuse and multi-tenancy ... 134
Interface VPC endpoints ... 135
Configuration and vulnerability analysis ... 137
Security best practices ... 138
Monitoring Amazon GameLift Streams ... 139
Monitor with CloudWatch ... 139
Stream group capacity and usage ... 140
Stream group performance and resource utilization ... 141
Stream status ... 142
Customer engagement ... 143
Data channels ... 143
Logging API calls with CloudTrail ... 144
Amazon GameLift Streams data events in CloudTrail ... 146
Amazon GameLift Streams management events in CloudTrail ... 147
Amazon GameLift Streams event examples ... 148
Real-time performance stats ... 152
Receive performance stats ... 152
Performance stats reference ... 154

Troubleshoot 156
Access denied 156
Application issues 157
Preliminary checks 157
Proton issues 157
Application issues due to screen resolution 158
Application terminates at start of stream session 158
Unreal Engine application crashes or requires additional dependencies 158
Performance issues 159
Game performance is reduced when streaming on Amazon GameLift Streams 159
Windows applications experience slow load times or stuttering issues 160
Stream connectivity issues 165
Stream input issues 166
General input troubleshooting 166
Gamepad and microphone inputs don't work on native Linux applications 166
Key input appears stuck on MacOS client 167
Mouse movement behaves differently on Amazon GameLift Streams 167
Stream session issues 168
Stream session does not start 168
Stream session terminated 169
Compatibility with Proton 169
High-level steps to test and troubleshoot 169
Known issues with Proton 170
Set up a local machine 170
Set up a remote machine 171
Troubleshoot on Proton 176
Profiling Unreal Engine performance 180
Regions, quotas, and limitations 183
AWS Regions and streaming locations 183
Service endpoints 184
Streaming locations 184
Supported locations by stream class 185
Service quotas 186
Service quotas 187
API rate limits 201
Other limitations 202

Usage and bills 204
Review your Amazon GameLift Streams bills and usage 204
Best practices to manage Amazon GameLift Streams costs 205
Create billing alerts to monitor usage 205
Scale stream groups to zero capacity 205
Delete original application files 206

What is Amazon GameLift Streams?

With Amazon GameLift Streams, game publishers and others can provide on-demand, low-latency streaming experiences to players and viewers globally. Amazon GameLift Streams uses its own streaming technology combined with the AWS global infrastructure to operate and maintain application streaming at scale. Publishers have the flexibility to provision both on-demand and reserved streaming resources to effectively manage capacity and costs.

Topics

- Features
- How to get started with Amazon GameLift Streams
- Accessing Amazon GameLift Streams

Features

Amazon GameLift Streams offers these key features:

- Streaming technology that delivers real-time gameplay experiences with minimal player-to-cloud latency to any device with a browser using the AWS global footprint.
- Seamless play of games at high definition (1080p) resolution and 60 fps without requiring downloads, turning any browser-based device into a powerful gaming machine.
- Scaling tools to adjust your streaming capacity to meet customer demand. For example, with these tools you can keep game streaming costs in line while maintaining enough capacity to accommodate new players into stream sessions quickly.
- Stream performance analysis using the Amazon GameLift Streams console to track metrics, view stream logs, and review data on stream resource usage.
- Direct streaming of Windows and Linux-based games with little to no modification.
- Amazon GameLift Streams SDK to help you integrate your existing identity services, storefront, and client applications.

How to get started with Amazon GameLift Streams

If you're a first-time Amazon GameLift Streams user, we recommend that you begin with the following topics:

- Setting up Amazon GameLift Streams as a developer covers one-time setup tasks, including getting an AWS account with user access and setting up the software you need for development with Amazon GameLift Streams.
- Starting your first stream in Amazon GameLift Streams guides you through the critical steps in the content streaming workflow. Starting with your content, such as a game build, you will provision Amazon GameLift Streams streaming cloud resources and initiate a streaming session.

Accessing Amazon GameLift Streams

You can create, access, and manage your application content and streaming resources with the following tools:

- AWS Management Console – Provides a web interface that you can use to create and manage your Amazon GameLift Streams applications and stream groups.
- AWS Command Line Interface (AWS CLI) – Provides commands for a broad set of AWS services and is supported on Windows, Mac, and Linux. For more information on this tool, see the AWS Command Line Interface page.
- AWS SDK – Provides language-specific APIs and takes care of connection details, such as calculating signatures, handling request retries, and error handling. For documentation on the Amazon GameLift Streams service API, see the Amazon GameLift Streams API Reference. For more general information on the AWS SDK, see Tools to Build on AWS.

For additional information on supported AWS Regions, see Regions, quotas, and limitations.

Terms of use for Amazon GameLift Streams

Before you use Amazon GameLift Streams, be sure you can comply with all applicable legal requirements, including license terms applicable to the applications you intend to stream, and the locations where you intend to be streaming.

- For more information about the AWS requirements, see Section 43 of the AWS Service Terms.
- For more information about the service level agreements, see the Amazon GameLift Streams Service Level Agreement.

Setting up Amazon GameLift Streams as a developer

To start using the Amazon GameLift Streams service with your projects, complete these basic setup tasks. If you already have an AWS account and a user under that account that you want to use with Amazon GameLift Streams, you can skip to Download the Web SDK.

For more information on what you can do with an AWS account, see Getting started with AWS.

After you've completed these setup tasks, we recommend that you go to Starting your first stream in Amazon GameLift Streams and step through the tutorial, which covers the entire workflow for getting your content streaming in a web client.

Topics

- Sign up for an AWS account
- Create a user with administrative access
- Get programmatic access
- Download the Amazon GameLift Streams Web SDK
- Download the AWS CLI
- Set up billing alerts

Sign up for an AWS account

If you do not have an AWS account, complete the following steps to create one.

To sign up for an AWS account

1. Open https://portal.aws.amazon.com/billing/signup.
2. Follow the online instructions.

   Part of the sign-up procedure involves receiving a phone call or text message and entering a verification code on the phone keypad.

   When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to a user, and use only the root user to perform tasks that require root user access.

AWS 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 https://aws.amazon.com/ and choosing **My Account**.

Create a user with administrative access

After you sign up for an AWS account, secure your AWS account root user, enable AWS IAM Identity Center, and create an administrative user so that you don't use the root user for everyday tasks.

Secure your AWS account root user

1. Sign in to the AWS Management Console as the account owner by choosing **Root user** and entering your AWS account email address. On the next page, enter your password.

   For help signing in by using root user, see Signing in as the root user in the *AWS Sign-In User Guide*.

2. Turn on multi-factor authentication (MFA) for your root user.

   For instructions, see Enable a virtual MFA device for your AWS account root user (console) in the *IAM User Guide*.

Create a user with administrative access

1. Enable IAM Identity Center.

   For instructions, see Enabling AWS IAM Identity Center in the *AWS IAM Identity Center User Guide*.

2. In IAM Identity Center, grant administrative access to a user.

   For a tutorial about using the IAM Identity Center directory as your identity source, see Configure user access with the default IAM Identity Center directory in the *AWS IAM Identity Center User Guide*.

Sign in as the user with administrative access

- To sign in with your IAM Identity Center user, use the sign-in URL that was sent to your email address when you created the IAM Identity Center user.

For help signing in using an IAM Identity Center user, see Signing in to the AWS access portal in the *AWS Sign-In User Guide*.

Assign access to additional users

1. In IAM Identity Center, create a permission set that follows the best practice of applying least-privilege permissions.

   For instructions, see Create a permission set in the *AWS IAM Identity Center User Guide*.

2. Assign users to a group, and then assign single sign-on access to the group.

   For instructions, see Add groups in the *AWS IAM Identity Center User Guide*.

Get programmatic access

In addition to your user sign-in credentials for the AWS Management Console, you need credentials for programmatic access, such as when working with the AWS Command Line Interface (AWS CLI). Programmatic credentials consist of a two-part set of access keys. Use one of the following methods to generate your access keys:

- Method 1 – If you're using an administrative user created with the IAM Identity Center, see Getting IAM role credentials for AWS CLI access to generate temporary security credentials for short-term access to AWS resources. When following these instructions, make sure you're signed in through your account's AWS access portal URL with your administrative user name and password (not your root user).
- Method 2 – If you're using an existing IAM user and you haven't yet transitioned to using the IAM Identity Center, see Managing access keys for IAM users (console) to generate long-term credentials for your user.

> **Note**
>
> As a best practice, use temporary credentials instead of long-term access keys. Temporary credentials include an access key ID, a secret access key, and a security token that indicates when the credentials expire. For more information, see Best practices for managing AWS access keys in the *AWS General Reference*.

Download the Amazon GameLift Streams Web SDK

You can get started without any additional materials by using the in-console streaming experience. We recommend this as a starting point because it allows you to evaluate how your application performs on the Amazon GameLift Streams without setting up any additional infrastructure. For more information, proceed to Getting started with Amazon GameLift Streams.

When you're ready to build your own Amazon GameLift Streams integration, download the Amazon GameLift Streams Web SDK, available in the Resources section of the Getting Started product page. Amazon GameLift Streams is built to be integrated into your web applications. You will need to integrate our JavaScript-based Web SDK to setup streaming from your website or browser-based applications. The download also contains a sample web server that uses the Amazon GameLift Streams service, and a sample web client for connecting to streams.

For more information about setting up your own Amazon GameLift Streams solution, refer to Amazon GameLift Streams backend service and web client.

Download the AWS CLI

To use Amazon GameLift Streams with your content, we recommend that you get the AWS Command Line Interface (AWS CLI). The AWS CLI is an open source tool that gives you equivalent AWS SDK functionality by running commands from a terminal program.

1. Download and install the latest version of the AWS CLI for your operating system. See these install instructions in the *AWS Command Line Interface User Guide*.
2. Configure the tool with your user access credentials and other preferences, as described in Setting up the AWS CLI. With this configuration, you won't have to explicitly specify your credentials and other settings with every command.
3. Use the following command to verify your installation and get a list of available Amazon GameLift Streams commands:

   ```
   aws gameliftstreams help
   ```

Set up billing alerts

A stream group incurs cost per active stream capacity per second. To make sure your cost and usage stays within your budget, see Create billing alerts to monitor usage.

Getting started with Amazon GameLift Streams

This section can help you successfully get started streaming your applications and games through Amazon GameLift Streams. The topics in this section cover the end-to-end process—from uploading your application to Amazon GameLift Streams, to testing how your content performs in a stream. It also covers important steps to help you prepare for streaming, such as choosing the right runtime and stream class configuration to optimize performance and cost.

Topics

- Choosing a configuration in Amazon GameLift Streams
- Configuration options
- Starting your first stream in Amazon GameLift Streams

Choosing a configuration in Amazon GameLift Streams

This guide can help you choose the optimal runtime environment and configuration settings for streaming your applications and games through Amazon GameLift Streams. The configuration settings directly impact your content's performance and the costs associated with running it on Amazon GameLift Streams. There are several options to support a wide variety of applications and graphical fidelity.

You can find the complete list of configuration options in Configuration options.

The following key terms can help you understand how these configuration options work together:

- *Runtimes* refer to the underlying operating system and software environment that will execute your application on Amazon GameLift Streams. The main runtime environment options are Windows, Linux, and Proton.
- *Stream classes* represent the different resource configurations available within Amazon GameLift Streams, varying in operating system, CPU, GPU, RAM, and other specifications. The stream class is a configuration option of a stream group that defines both the hardware resources allocated to a stream session and the tenancy model (how many concurrent streams can run on a single virtual machine).
- *Multi-tenancy* enables multiple users to share the same underlying hardware resources, which can be a cost-effective option for applications that don't require maximum hardware capabilities.

A stream class with multi-tenancy can host multiple streams for the cost of one resource. "High" stream classes have 1:2 tenancy, while "Ultra" stream classes have 1 tenancy.

When setting up your Amazon GameLift Streams configuration, the runtime environment you choose determines the specific stream class options that are compatible and available to you. Matching your application's requirements with the right runtime environment and stream class is key to optimizing performance and cost-efficiency in Amazon GameLift Streams.

The cost to stream depends on the stream class. For a detailed list of cost, refer to the Amazon GameLift Streams Pricing page.

Starting point

Depending on your application, these are good starting points to get started streaming. Later, you can explore other configuration options to optimize the cost.

For Windows applications

We recommend using the Microsoft Windows Server 2022 Base runtime environment for Microsoft Windows applications. There are four hardware configurations available for this runtime, the NVIDIA-based `gen6n_pro_win2022` , `gen6n_ultra_win2022` , `gen5n_win2022`, and `gen4n_win2022` stream classes. In this environment, Amazon GameLift Streams supports games and other 3D applications using DirectX 11 or DirectX 12, and game engines including Unity 2022.3, Unreal Engine 4.27, and Unreal Engine 5 up through . Streaming is supported over both IPv4 and IPv6.

This combination of runtime environment and stream classes provides a predictable, well-supported configuration with the highest compatibility and best performance for your Windows-based content.

For Linux applications

Use the Ubuntu 22.04 LTS runtime environment for applications built to run natively on Linux. To optimize for performance, choose one of the Pro or Ultra stream classes (`gen6n_pro` , `gen6n_ultra` , `gen5n_ultra` , or `gen4n_ultra`). To optimize for cost, choose one of the Small, Medium, or High stream classes (`gen6n_high` , `gen6n_medium` , `gen6n_small` , `gen5n_high` , or `gen4n_high`) that support *multi-tenancy*— a cost-effective option where multiple concurrent stream sessions share the same compute resources.

Important

The Linux runtime in Amazon GameLift Streams does not support streaming over IPv6. Clients must stream applications over IPv4.

Cost optimizations

While the starting point recommendations are a great place to begin, you might want to consider other configuration options to optimize for cost while maintaining good performance.

Use Proton runtime environment

Many Windows applications can run in the Proton runtime environment. Proton is a game-optimized compatibility layer that runs on Linux. The stream class options for this runtime include powerful GPU resources running on NVIDIA hardware, with support for DirectX 11 and, beginning with Proton 8.0-5, DirectX 12. Visit the Proton wiki for more details about this option. If you choose to explore running your application on Proton, we recommend that you start your testing using Proton 9.0-2.

Important

Proton runtimes in Amazon GameLift Streams do not support streaming over IPv6. Clients must stream applications over IPv4.

Important

The compatibility of your Windows application in a Proton runtime environment depends on your specific application requirements. For example, Proton 9.0-2 has better support than Proton 8.0-2c for Unreal Engine 5. In general, the newer your game, the newer version of Proton you will need. We strongly recommend thoroughly testing this runtime in your local environment to ensure optimal performance. Use our Proton troubleshooting guide to help you in this effort.

Compile your application to Linux

Another cost-saving option is to target your application to run natively on Linux. Test the application on your end first to ensure that the Linux version of your application performs as needed. If your application successfully runs on Linux, then you can follow the Amazon GameLift Streams configuration options for Linux applications.

For information about cross-compiling Unreal Engine applications to Linux, refer to the Cross-Compile Toolchain section in Unreal Engine's developer guide.

Deciding on a configuration

To determine the best runtime and stream class configuration, consider the following key questions.

1. **What platform is your application or game built for?** If you have a Windows application, the Windows runtime environment is the simplest to set up. If your application is built for Linux, the Linux runtime environment is the most straightforward. To save costs for streaming a Windows application, you can explore the Proton runtime environment or compile the application to Linux.
2. **How important is performance versus cost for your use case?** The Windows runtime environment may offer the best performance, but it can be more expensive to run. Comparitively, the Proton runtime environment is more cost-effective, though you might experience slightly lower performance or potential compatibility issues. This is because Windows-based applications might require certain functionality that is not yet fully supported in the available Proton runtimes. As a result, you could experience functional or graphical differences when running your application on the Proton environment. We recommend that you test your application on the different runtime environments and stream classes to evaluate the performance and cost trade-offs. For a full list of runtime environment options, see Runtime environments.
3. **What are the graphical requirements of your application?** The graphical requirements of your application can help determine which stream class configuration is most appropriate. If your application demands high-performance GPUs, you should consider using stream classes with greater amounts of video memory (VRAM) and system memory (RAM). For example, the gen5n and gen6n stream classes deliver up to 3x better performance for graphics-intensive applications compared to the gen4n stream classes. If your application requires maximum GPU and CPU resources, you should consider the "pro" stream classes. Conversely, if your application

can operate effectively at a lower graphical fidelity, you might save on costs by using stream classes that support multi-tenancy (any of the "small", "medium", or "high" stream classes). This allows multiple users to share the same underlying hardware resources. For a full description of stream class options, see Stream classes.

4. **How much effort are you willing to invest in the setup?** The simplest way to set up your application is to run it natively using the Windows or Linux runtimes, since they are more likely to be compatible with your application out-of-the-box. In contrast, the Proton runtime environment will require more hands-on testing to identify the optimal Proton configuration for your needs. Consider the time and resources you can allocate to the setup and testing process when deciding between the runtime environment options.
5. **Have you tested your application on the various runtime environments and stream classes?** We recommend testing your content on different runtime environments and stream classes to see how it performs. This helps you determine the best fit based on factors like stability, graphics quality, feature functionality, and input responsiveness.

How your configuration choices impact next steps

The configuration you select directly impacts the next phases of setting up your streaming environment. Specifically:

- **Creating an Amazon GameLift Streams application**: When you upload your game or application to Amazon GameLift Streams, you'll need to specify the runtime environment you want to use. This choice will determine the type of stream group you can use.
- **Linking to a stream group**: If you already have an existing stream group, your runtime environment choice will need to match the configuration of that group. For example, if you select the Windows runtime, you can only link your application to a stream group that's set up for Windows applications.
- **Creating a stream group**: When you create a new stream group, you must choose a stream class that's compatible with your chosen runtime. The stream class you choose should match the graphics requirements and compute power that your application requires.

By understanding how the configuration settings you choose influences these subsequent steps, you can better plan your overall streaming implementation and ensure a smooth integration process.

Next steps

Depending on the configuration you've chosen, there are a few different approaches you can take to set up your application for streaming.

If you've selected the Windows or Linux runtime

For Windows or Linux runtimes, the next steps are to set up streaming in Amazon GameLift Streams and then test the stream. For more information, proceed to Starting your first stream in Amazon GameLift Streams.

If you're considering using Proton

An application's compatibility with Proton depends on the application's specific requirements. Therefore, we recommend that you test your application on different Proton versions before bringing it to Amazon GameLift Streams. This helps you identify the Proton setup that provides the best performance and compatibility for your needs. By testing outside of Amazon GameLift Streams, you can validate the application's performance and functionality, and debug issues that are specific to the runtime. For information, see Testing and troubleshooting compatibility with Proton for Amazon GameLift Streams.

When you've selected a specific Proton configuration, you're ready to set up streaming in Amazon GameLift Streams. For more information, proceed to Starting your first stream in Amazon GameLift Streams.

Configuration options

Runtime environments

Runtimes refer to the underlying operating system and software environment that executes your application on Amazon GameLift Streams. The main runtime options are Windows, Linux, and Proton. You specify the runtime environment in Step 2: Configure your application for Amazon GameLift Streams of the getting started worflow.

Proton is a compatibility layer that enables many Windows applications to run in a Linux-based environment. If you plan to use Proton, we recommend that you test how your application runs on a local machine. For more information, refer to Testing and troubleshooting compatibility with Proton for Amazon GameLift Streams.

Runtime	Description
Microsoft Windows Server 2022 Base	Compatible with Windows applications. Supports using IPv4 and IPv6 in stream sessions.
Ubuntu 22.04 LTS	Compatible with Linux applications. Does not support using IPv6 in stream sessions.
Proton 9.0-2	Compatible with Windows applications. Based on the Proton experimental_9.0 branch. Recommended version to start testing compatiblity with Proton. Does not support using IPv6 in stream sessions.
Proton 8.0-5	Compatible with Windows applications. Based on the Proton experimental_8.0 branch. Does not support using IPv6 in stream sessions.
Proton 8.0-2c	Compatible with Windows applications. Based on the Proton experimental_8.0 branch. Does not support using IPv6 in stream sessions.

Limitations

Gamepad support is not available on Ubuntu 22.04 LTS. Other runtime environments support gamepads, depending on the end user's operating system and browser. For more information, see Supported browsers and input.

Stream classes

Stream classes represent the different resource configurations available within Amazon GameLift Streams, varying in CPU, GPU, RAM, and other specifications. The stream class is a configuration option of a stream group that defines both the hardware resources allocated to a stream session and the tenancy model (how many concurrent streams can run on a single virtual machine). You specify the stream class in Step 3: Manage how Amazon GameLift Streams streams your application of the getting started workflow.

Stream class	Amazon EC2 configuration	Description
`gen6n_pro_win2022`	Windows runtime on a g6.4xlarge Amazon EC2 instance	(NVIDIA, pro) Supports applications with extremely high 3D scene complexity which require maximum resources. Runs applicati

Stream class	Amazon EC2 configuration	Description
		ons on Microsoft Windows Server 2022 Base and supports DirectX 12. Compatibl e with Unreal Engine versions up through 5.6, 32 and 64-bit applications, and anti-cheat technology. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 16. RAM: 64 GB. VRAM: 24 GB. Tenancy: Supports up to one concurrent stream session.
`gen6n_pro`	Linux runtime on a g6.4xlarge Amazon EC2 instance	(NVIDIA, pro) Supports applications with extremely high 3D scene complexity which require maximum resources. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 16. RAM: 64 GB. VRAM: 24 GB. Tenancy: Supports up to one concurrent stream session.

Stream class	Amazon EC2 configuration	Description
`gen6n_ultra_win2022`	Windows runtime on a g6.2xlarge Amazon EC2 instance	(NVIDIA, ultra) Supports applications with high 3D scene complexity. Runs applications on Microsoft Windows Server 2022 Base and supports DirectX 12. Compatible with Unreal Engine versions up through 5.6, 32 and 64-bit applications, and anti-cheat technology. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 24 GB. Tenancy: Supports up to one concurrent stream session.
`gen6n_ultra`	Linux runtime on a g6.2xlarge Amazon EC2 instance	(NVIDIA, ultra) Supports applications with high 3D scene complexity. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 24 GB. Tenancy: Supports up to one concurrent stream session.
`gen6n_high`	Linux runtime on a g6.2xlarge Amazon EC2 instance with 2:1 tenancy	(NVIDIA, high) Supports applications with moderate-to-high 3D scene complexity. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 4. RAM: 16 GB. VRAM: 12 GB. Tenancy: Supports up to two concurrent stream sessions.

Stream class	Amazon EC2 configuration	Description
`gen6n_medium`	Linux runtime on a g6.2xlarge Amazon EC2 instance with 4:1 tenancy	(NVIDIA, medium) Supports applications with moderate 3D scene complexity. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 2. RAM: 8 GB. VRAM: 6 GB. Tenancy: Supports up to four concurrent stream sessions.
`gen6n_small`	Linux runtime on a g6.4xlarge Amazon EC2 instance with 12:1 tenancy	(NVIDIA, small) Supports applications with lightweight 3D scene complexity and low CPU usage. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 1. RAM: 4 GB. VRAM: 2 GB. Tenancy: Supports up to twelve concurrent stream sessions.
`gen5n_win2022`	Windows runtime on a g5.2xlarge Amazon EC2 instance	(NVIDIA, ultra) Supports applications with extremely high 3D scene complexity. Runs applications on Microsoft Windows Server 2022 Base and supports DirectX 12 and DirectX 11. Supports Unreal Engine up through version 5.6, 32 and 64-bit applications, and anti-cheat technology. Uses NVIDIA A10G Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 24 GB. Tenancy: Supports one concurrent stream session.

Stream class	Amazon EC2 configuration	Description
`gen5n_ultra`	Linux runtime on a g5.2xlarge Amazon EC2 instance	(NVIDIA, ultra) Supports applications with extremely high 3D scene complexity. Uses NVIDIA A10G Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 24 GB. Tenancy: Supports one concurrent stream session.
`gen5n_high`	Linux runtime on a g5.2xlarge Amazon EC2 instance with 2:1 tenancy	(NVIDIA, high) Supports applications with moderate-to-high 3D scene complexity. Uses NVIDIA A10G Tensor Core GPU. Resources per application: vCPUs: 4. RAM: 16 GB. VRAM: 12 GB. Tenancy: Supports up to two concurrent stream sessions.
`gen4n_win 2022`	Windows runtime on a g4dn.2xlarge Amazon EC2 instance	(NVIDIA, ultra) Supports applications with high 3D scene complexity. Runs applicati ons on Microsoft Windows Server 2022 Base and supports DirectX 12 and DirectX 11. Supports Unreal Engine up through version 5.6, 32 and 64-bit applications, and anti-cheat technology. Uses NVIDIA T4 Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 16 GB. Tenancy: Supports one concurrent stream session.

Stream class	Amazon EC2 configuration	Description
gen4n_ultra	Linux runtime on a g4dn.2xlarge Amazon EC2 instance	(NVIDIA, ultra) Supports applications with high 3D scene complexity. Uses NVIDIA T4 Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 16 GB. Tenancy: Supports one concurrent stream session.
gen4n_high	Linux runtime on a g4dn.2xlarge Amazon EC2 instance with 2:1 tenancy	(NVIDIA, high) Supports applications with moderate-to-high 3D scene complexity. Uses NVIDIA T4 Tensor Core GPU. Resources per application: vCPUs: 4. RAM: 16 GB. VRAM: 8 GB. Tenancy: Supports up to two concurrent stream sessions.

Runtime environment and stream class compatibility

Runtime environment	Compatible stream classes
Windows	gen6n_pro_win2022 gen6n_ultra_win2022 gen5n_win2022 gen4n_win2022
Linux (Ubuntu 22.04 LTS)	gen6n_pro gen6n_ultra

Runtime environment	Compatible stream classes
	gen6n_high gen6n_medium gen6n_small gen5n_ultra gen5n_high gen4n_ultra gen4n_high
Proton	gen6n_pro gen6n_ultra gen6n_high gen6n_medium gen6n_small gen5n_ultra gen5n_high gen4n_ultra gen4n_high

Starting your first stream in Amazon GameLift Streams

This tutorial walks you through the steps to get started with Amazon GameLift Streams to stream your application or game. Amazon GameLift Streams runs your application and streams them directly to your end-users' web browser. You'll learn how to upload and configure the application you want to stream, and how to manage the way Amazon GameLift Streams streams. By the end,

you will test how your application streams on Amazon GameLift Streams by interacting with it directly in the Amazon GameLift Streams console.

⚠ **Before you start, understand Amazon GameLift Streams pricing.**

You can find the cost of Amazon GameLift Streams in the Pricing page. To learn more, refer to Managing usage and bills for Amazon GameLift Streams.
You incur cost for using Amazon GameLift Streams, specifically when you:

- Create an Amazon GameLift Streams application in Step 2: Configure your application for Amazon GameLift Streams
- Create a stream group in Step 3: Manage how Amazon GameLift Streams streams your application

Do not skip Step 5: Clean up (don't skip). To avoid unnecessary charges after you're done trying Amazon GameLift Streams, you must clean up all your resources.

Topics

- Prerequisites
- Step 1: Upload your application to an Amazon S3 bucket
- Step 2: Configure your application for Amazon GameLift Streams
- Step 3: Manage how Amazon GameLift Streams streams your application
- Step 4: Test your stream in Amazon GameLift Streams
- Step 5: Clean up (don't skip)

Prerequisites

Complete the following tasks before you start the tutorial.

- Sign up for an AWS account and create a user with administrative access, if you don't already have one. Refer to the Setting up topic in this guide for help with this task. You do not need to download the Amazon GameLift Streams Web SDK or set up the AWS CLI at this time — you will complete the following steps using the AWS Management Console.

- Get a version of your application content files with no digital rights management (DRM). Collect the files required to run the application, including executables and assets, into a folder, but *do not* compress the folder.

Step 1: Upload your application to an Amazon S3 bucket

Amazon GameLift Streams uses Amazon Simple Storage Service (Amazon S3) to store your application or game files in the cloud and access it for streaming. In this step, you upload your application files to an Amazon S3 bucket. Complete this step in the Amazon S3 Console.

> **Note**
>
> The Amazon S3 storage class that Amazon GameLift Streams requires is the default **S3 Standard**. Other storage classes like **S3 Glacier** or objects being moved to **Infrequent Access** or **Archive Access** by **S3 Intelligent-Tiering** are not supported by Amazon GameLift Streams.
> To optimize storage cost, you can delete the application from your S3 bucket after you complete Step 2: Configure your application for Amazon GameLift Streams and the application is in **Ready** status.

Application limitations

Name	Default	Adjuse	Description
Files per application	30,000 files	Yes*	The maximum number of files that you can have in an application, in this account.
Single file size	80 GiB	No	The maximum size of a single file in an application. Note that a gibibyte (GiB) equals 1024*1024*1024 bytes.

Name	Default	Adjuse	Description
Application size	100 GiB	Yes*	The maximum total size of an Amazon GameLift Streams application, in this account. Note that a gibibyte (GiB) equals 1024*1024*1024 bytes.

*To request an increase, sign in to the AWS Management Console and open the Service Quotas console to Amazon GameLift Streams, where you can review your current quotas in the **Applied account-level quota value** column and submit a request to increase a value.

To upload your application to Amazon S3

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. Create an Amazon S3 bucket. Enter a bucket name and select an AWS Region. This region must be the same as the application and stream group that you will create later. See AWS Regions and streaming locations supported by Amazon GameLift Streams for a list of AWS Regions where Amazon GameLift Streams is available. For the remaining fields, keep the default settings.

   For more instructions, refer to Creating a bucket in the *Amazon Simple Storage Service User Guide*.
3. Open the new bucket and upload the folder with your application files.

> **Warning**
>
> You must upload your application files as an uncompressed folder. Don't upload a `.zip` folder.

> **Warning**
>
> Ensure that the application files you uploaded are the correct ones, and are within the application file size limitations. If you want to update your files later, you must repeat Step 2: Configure your application for Amazon GameLift Streams, which can take a few minutes.

Step 2: Configure your application for Amazon GameLift Streams

What is an application in Amazon GameLift Streams?

An Amazon GameLift Streams application is a resource that contains a game or interactive application that runs on Amazon GameLift Streams infrastructure and delivers gameplay experiences to players through cloud streaming. The application executes on AWS compute instances and renders game content that is streamed directly to players' devices over the internet, eliminating the need for players to download, install, or run the game locally.

In this step, you configure the application you want to stream with Amazon GameLift Streams by creating an Amazon GameLift Streams application. When you create an Amazon GameLift Streams application, you provide the Amazon S3 URI to the application folder you uploaded to your Amazon S3 bucket, and the relative path to a valid executable or script file. Complete this step in the Amazon GameLift Streams console.

To create an Amazon GameLift Streams application using the Amazon GameLift Streams console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console. Choose the same AWS Region as the Amazon S3 bucket where you uploaded your set of files. For more information, refer to Choosing a Region in the *AWS Management Console Getting Started Guide*.
2. In the navigation bar, choose **Applications** and then choose **Create application**.
3. In **Runtime settings**, enter the following:

   - **Runtime environment**

     This is the runtime environment to run your application on. Amazon GameLift Streams can run on either Windows, Ubuntu 22.04 LTS, or Proton.

You cannot edit this field after the creation workflow.

Choose from one of the following runtime environments.

- For Linux applications:
  - Ubuntu 22.04 LTS (`UBUNTU, 22_04_LTS`)
- For Windows applications:
  - Microsoft Windows Server 2022 Base (`WINDOWS, 2022`)
  - Proton 9.0-2 (`PROTON, 20250516`)
  - Proton 8.0-5 (`PROTON, 20241007`)
  - Proton 8.0-2c (`PROTON, 20230704`)

Review the descriptions and use the comparison checklist to help you select the optimal runtime environment for your application.

4. In **General settings**, enter the following:

   a. **Description**

   This is a human-readable label for your application. This value does not have to be unique. For best practice, use a meaningful description, name, or label for the application. You can edit this field at any time.

   b. **Base path**

   This is the Amazon S3 URI to your application's root folder in the Amazon S3 bucket. The folder and any subfolders should contain your build executable and any supporting files.

   A valid URI is the bucket prefix that contains all the files needed to run and stream the application. Example: A bucket called `mygamebuild` contains three complete versions of the game build files, each in a separate folder. You want to stream the build in the folder `mygamebuild-EN101`. In this example, the URI is `s3://amzn-s3-demo-bucket/mygamebuild-EN101`.

   You cannot edit this field after the creation workflow.

   c. **Executable launch path**

This is the Amazon S3 URI of the executable file that Amazon GameLift Streams will stream. The file must be contained within the application's root folder. For Windows applications, the file must be a valid Windows executable or batch file with a filename ending in .exe, .cmd, or .bat. For Linux applications, the file must be a valid Linux binary executable or a script that contains an initial interpreter line starting with a shebang ('#!').

You cannot edit this field after the creation workflow.

5. (Optional) In **Application log path**, enter the following:

   a. **Application log path**

   This is the path (or paths) to the application folder or file which contains logs that you want to save. Specify each log path relative to your application base path. If you use this feature, then at the end of every stream session, Amazon GameLift Streams will copy the file(s) that you specify to the Amazon S3 bucket that you name. The copy operation is not performed recursively in an application folder's subfolders.

   To disable logging, remove all application log paths and clear the application log output destination.

   You can edit this field at any time.

   b. **Application log output**

   This is the URI to the Amazon S3 bucket where Amazon GameLift Streams will copy application log files. This field is required if you specify an application log path.

   To disable logging, remove all application log paths and clear the application log output destination.

   You can edit this field at any time.

   To save log files on your behalf, Amazon GameLift Streams must be given permission to your S3 bucket for saving. If you let Amazon GameLift Streams create the bucket for logging, the permission policy will be applied automatically upon creation. If you provide your own bucket, you will need to apply the permission policy, yourself.

   Bucket permission policy template

Copy the following policy code and apply it to the bucket that you want to use for application logs. Be sure to replace **amzn-s3-demo-bucket** with the name of your existing S3 bucket.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PutPolicy",
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "gameliftstreams.amazonaws.com"
        ]
      },
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::amzn-s3-demo-bucket/*",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "your 12-digit account id"
        }
      }
    }
  ]
}
```

6. (Optional) In **Tags**, assign tags to this application.

   Tags are labels that can help you organize your AWS resources. For more information, refer to Tagging your AWS resources.

   For example to track application versions, use a tag such as `application-version : my-game-1121`.

7. Choose **Create application**.

Amazon GameLift Streams takes a few minutes to prepare your application. In the **Applications** page, the new application is in **Processing** status. When your application is in **Ready** status, you can go to the next step, Step 3: Manage how Amazon GameLift Streams streams your application.

If the request returns an error, or if the application is created but in **Error** status, make sure that you're working with user credentials that include access to both Amazon S3 and Amazon GameLift Streams.

> **Note**
>
> When an application is in **Ready** status, you can safely delete the application files in your Amazon S3 bucket, without affecting your new application. This also helps optimize storage cost. For more information, see Delete an application.

For more information, refer to Prepare an application in Amazon GameLift Streams.

Step 3: Manage how Amazon GameLift Streams streams your application

What is a stream group?

Manage how Amazon GameLift Streams streams your applications by using a stream group. A stream group is a collection of compute resources that Amazon GameLift Streams uses to stream your application to end users. When you create a stream group, you specify the hardware configuration (CPU, GPU, RAM) that will run your game (known as the *stream class*), the geographical locations where your game can run, and the number of streams that can run simultaneously in each location (known as *stream capacity*). You can link an application when you create the stream group, or wait until later, but you must link at least one application before you can stream from a stream group. After a stream group has been created, Amazon GameLift Streams allocates compute resources in the locations where you have allocated stream capacity. At this point, you can also associate additional applications to the stream group so that you have a choice of which one to stream.

With your application ready, the next thing you need is compute resources for Amazon GameLift Streams to stream it. In this step, you manage how Amazon GameLift Streams streams your application by creating a stream group. Complete this step in the Amazon GameLift Streams console.

To create a stream group in the Amazon GameLift Streams console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console. Choose the AWS Region where you want to create your stream group. This Region must be the same as that of the application that you want to stream with the stream group. For more information, refer to Choosing a Region in the *AWS Management Console Getting Started Guide*.
2. To open the creation workflow, in the navigation pane, choose **Stream groups**, and then choose **Create stream group**.
3. In **Define stream group**, enter the following:

   a. **Description**

   A human-readable label for your stream group. This value doesn't have to be unique. As a best practice, use a meaningful description, name, or label for the stream group. You can edit this field at any time.

   b. **Tags**

   Tags are labels that can help you organize your AWS resources. For more information, refer to Tagging your AWS resources.
4. In **Select stream class**, choose a stream class for the stream group.

   - **Stream class options**

   The type of compute resources to run and stream applications with. This choice impacts the quality of the streaming experience and the cost. You can specify only one stream class per stream group. Choose the class that best fits your application.

Stream class	Description
`gen6n_pro_win2022`	(NVIDIA, pro) Supports applications with extremely high 3D scene complexity which require maximum resources. Runs applications on Microsoft Windows Server 2022 Base and supports DirectX 12. Compatible with Unreal Engine versions up through 5.6, 32 and 64-bit applications, and anti-cheat technology. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 16. RAM: 64 GB. VRAM: 24 GB.

Stream class	Description
	Tenancy: Supports up to one concurrent stream session.
gen6n_pro	(NVIDIA, pro) Supports applications with extremely high 3D scene complexity which require maximum resources. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 16. RAM: 64 GB. VRAM: 24 GB. Tenancy: Supports up to one concurrent stream session.
gen6n_ultra_win2022	(NVIDIA, ultra) Supports applications with high 3D scene complexity. Runs applications on Microsoft Windows Server 2022 Base and supports DirectX 12. Compatible with Unreal Engine versions up through 5.6, 32 and 64-bit applications, and anti-cheat technology. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 24 GB. Tenancy: Supports up to one concurrent stream session.
gen6n_ultra	(NVIDIA, ultra) Supports applications with high 3D scene complexity. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 24 GB. Tenancy: Supports up to one concurrent stream session.
gen6n_high	(NVIDIA, high) Supports applications with moderate-to-high 3D scene complexity. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 4. RAM: 16 GB. VRAM: 12 GB. Tenancy: Supports up to two concurrent stream sessions.

Stream class	Description
gen6n_medium	(NVIDIA, medium) Supports applications with moderate 3D scene complexity. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 2. RAM: 8 GB. VRAM: 6 GB. Tenancy: Supports up to four concurrent stream sessions.
gen6n_small	(NVIDIA, small) Supports applications with lightweight 3D scene complexity and low CPU usage. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 1. RAM: 4 GB. VRAM: 2 GB. Tenancy: Supports up to twelve concurrent stream sessions.
gen5n_win2022	(NVIDIA, ultra) Supports applications with extremely high 3D scene complexity. Runs applications on Microsoft Windows Server 2022 Base and supports DirectX 12 and DirectX 11. Supports Unreal Engine up through version 5.6, 32 and 64-bit applications, and anti-cheat technology. Uses NVIDIA A10G Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 24 GB. Tenancy: Supports one concurrent stream session.
gen5n_high	(NVIDIA, high) Supports applications with moderate-to-high 3D scene complexity. Uses NVIDIA A10G Tensor Core GPU. Resources per application: vCPUs: 4. RAM: 16 GB. VRAM: 12 GB. Tenancy: Supports up to two concurrent stream sessions.
gen5n_ultra	(NVIDIA, ultra) Supports applications with extremely high 3D scene complexity. Uses NVIDIA A10G Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 24 GB. Tenancy: Supports one concurrent stream session.

Stream class	Description
`gen4n_win2022`	(NVIDIA, ultra) Supports applications with high 3D scene complexity. Runs applications on Microsoft Windows Server 2022 Base and supports DirectX 12 and DirectX 11. Supports Unreal Engine up through version 5.6, 32 and 64-bit applicati ons, and anti-cheat technology. Uses NVIDIA T4 Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 16 GB. Tenancy: Supports one concurrent stream session.
`gen4n_high`	(NVIDIA, high) Supports applications with moderate-to-high 3D scene complexity. Uses NVIDIA T4 Tensor Core GPU. Resources per application: vCPUs: 4. RAM: 16 GB. VRAM: 8 GB. Tenancy: Supports up to two concurrent stream sessions.
`gen4n_ultra`	(NVIDIA, ultra) Supports applications with high 3D scene complexity. Uses NVIDIA T4 Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 16 GB. Tenancy: Supports one concurrent stream session.

To continue, choose **Next**.

5. In **Link application**, choose an application that you want to stream, or select "**No application**" to choose one at a later time. You can edit the stream group after it has been created to add or remove applications. You can only link an application that's in `Ready` status and has a runtime that's compatible with the stream class you've chosen. By default, these are the only applications that are shown in the table. To see all applications in `Ready` status, choose `All runtimes` in the drop down list.

> **Note**
>
> If you don't see your application listed, then check the current AWS Region setting. You can only link an application to a stream group that's in the same Region.

To continue, choose **Next**.

6. In **Configure stream settings**, under **Locations and capacity**, choose one or more locations where your stream group will have capacity to stream your application. By default, the region where you create the stream group, known as the *primary location*, has already been added to your stream group and cannot be removed. You can add additional locations by checking the box next to each location that you want to add. For lower latency and better quality streaming, you should choose locations closer to your users.

   For each location, you can specify its *streaming capacity*. Stream capacity represents the number of concurrent streams that can be active at a time. You set stream capacity per location in each stream group.

   - **Always-on capacity:** This setting, if non-zero, indicates minimum streaming capacity which is allocated to you and is never released back to the service. You pay for this base level of capacity at all times, whether used or idle.
   - **Maximum capacity:** This indicates the maximum capacity that the service can allocate for you. Newly created streams may take a few minutes to start. Capacity is released back to the service when idle. You pay for capacity that is allocated to you until it is released.
   - **Target-idle capacity:** This indicates idle capacity which the service pre-allocates and hold for you in anticipation of future activity. This helps to insulate your users from capacity-allocation delays. You pay for capacity which is held in this intentional idle state.

   You can increase or decrease your total stream capacity at any time to meet changes in user demand for a location by adjusting either capacity. Amazon GameLift Streams fulfills streaming requests using the idle, pre-allocated resources in the always-on capacity pool if any are available. If all always-on capacity is in use, Amazon GameLift Streams will provision additional compute resources up to the maximum number specified in on-demand capacity. As allocated capacity scales, the change is reflected in your total cost for the stream group.

Linked applications will be automatically replicated to each enabled location. An application must finish replicating in a remote location before the remote location can host a stream. To check on the replication status, open the stream group after it has been created and refer to the **Replication status** column in the table of linked applications. Click on the current status to see the replication status for each added location.

> **Note**
>
> Application data will be stored in all enabled locations including the primary location for this stream group. Stream session data will be stored in both the primary location and the location where the streaming occurred.

7. In **Review and create stream group**, verify your stream group configuration and make changes as needed. When everything is correct, choose **Create stream group**.

For more information, refer to Manage streaming with an Amazon GameLift Streams stream group.

Step 4: Test your stream in Amazon GameLift Streams

What is a stream session?

Refers to the stream itself. This is an instance of a stream that Amazon GameLift Streams transmits from the server to the end-user. A stream session runs on a compute resource, or stream capacity, that a stream group has allocated. Also referred to as *stream* for short.

You can see how your application streams by running it directly in the Amazon GameLift Streams console. When you start a stream, Amazon GameLift Streams uses one of the compute resources that your stream group allocates. So, you must have available capacity in your stream group.

To test your stream in the Amazon GameLift Streams console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console.
2. You can test a stream in several ways. Start from the **Stream groups** page or **Test stream** page and follow these steps:

   a. Select a stream group that you want to use to stream.

b. If you're starting from the **Stream groups** page, choose **Test stream**. If you're starting from the **Test stream** page, select **Choose**. This opens the **Test stream** configuration page for the selected stream group.

c. In **Linked applications**, select an application.

d. In **Location**, choose a location with available capacity.

e. (Optional) In **Program configurations**, enter command-line arguments or environment variables to pass to the application as it launches.

f. Confirm your selection, and choose **Test stream**.

3. After your stream loads, you can do the following actions in your stream:

a. To connect input, such as your mouse, keyboard, and gamepad (except microphones, which are not supported in **Test stream**), choose **Attach input**. You automatically attach your mouse when you move the cursor into the stream window.

b. To have files that were created during the streaming session exported to an Amazon S3 bucket at the end of the session, choose **Export files** and specify the bucket details. Exported files can be found on the **Sessions** page.

c. To view the stream in fullscreen, choose **Fullscreen**. Press **Escape** to reverse this action.

4. To end the stream, choose **Terminate session**. When the stream disconnects, the stream capacity becomes available to start another stream.

> **Note**
>
> The **Test stream** feature in the Amazon GameLift Streams console does not support microphones.

Step 5: Clean up (don't skip)

> **Avoid unnecessary cost**
>
> A stream groups incurs costs when it has allocated capacity, even if that capacity is unused. To avoid unnecessary costs, scale your stream group capacities to your required size. We suggest during development you scale your always-on capacity to zero when not in use. For more information, refer to Best practices to manage Amazon GameLift Streams costs.

After you complete the tutorial and no longer need to stream your application, follow these steps to clean up your Amazon GameLift Streams resources.

To delete a stream group using the Amazon GameLift Streams console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console.
2. To view a list of your existing stream groups, in the navigation pane, choose **Stream groups**.
3. Choose the name of the stream group that you want to delete.
4. On the stream group detail page, choose **Delete**.
5. In the **Delete** dialog box, confirm the delete action.

Amazon GameLift Streams begins releasing compute resources and deleting the stream group. During this time, the stream group is in **Deleting** status. After Amazon GameLift Streams deletes the stream group, you can no longer retrieve it.

To delete an application using the Amazon GameLift Streams console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console.
2. In the navigation bar, choose **Applications** to view a list of your existing applications. Choose the application you want to delete.
3. In the application detail page, choose **Delete**.
4. In the **Delete** dialog box, confirm the delete action.

Amazon GameLift Streams begins deleting the application. During this time, the application is in `Deleting` status. After Amazon GameLift Streams deletes the application, you can no longer retrieve it.

For more information, refer to Delete a stream group and Delete an application.

Managing your streams with Amazon GameLift Streams

This section provides detailed information about how to stream with Amazon GameLift Streams. Learn about the streaming resources (an *application* and *stream group*), the properties to scale your streaming (*stream capacity* and *locations*), and the stream itself (a *stream session*). You can handle all of the tasks required to set up streaming with Amazon GameLift Streams by using the Amazon GameLift Streams console or Amazon GameLift Streams CLI commands.

If it's your first time using Amazon GameLift Streams, then refer to Starting your first stream in Amazon GameLift Streams, which walks you through the whole workflow.

Topics

- Key concepts
- Prepare an application in Amazon GameLift Streams
- Manage streaming with an Amazon GameLift Streams stream group
- Overview of multi-application stream groups
- Start stream sessions with Amazon GameLift Streams
- Export stream session files

Key concepts

Application

An Amazon GameLift Streams application is a resource that contains a game or interactive application that runs on Amazon GameLift Streams infrastructure and delivers gameplay experiences to players through cloud streaming. The application executes on AWS compute instances and renders game content that is streamed directly to players' devices over the internet, eliminating the need for players to download, install, or run the game locally.

Multi-application stream groups

A stream group that's linked to multiple applications. This many-to-one relationship allows you to stream multiple applications by using the same configuration that you've set up in a single stream group . When you start a stream session, you specify any linked applications. Then, Amazon GameLift Streams streams that application by using available stream capacity in this stream group.

Multi-location stream groups

A stream group that's configured to host applications and stream sessions from multiple locations, in addition to the primary location (the AWS Region where you created the stream group). You manage capacity for each location.

Multi-tenancy

Tenancy refers to how many concurrent streams can be supported by a single compute resource in Amazon GameLift Streams. *Multi-tenancy* is a feature that enables multiple users to share the same underlying hardware resources, which can be a cost-effective option for applications that don't require maximum hardware capabilities. A stream class with multi-tenancy can host multiple streams for the cost of one resource. "High" stream classes support multi-tenancy, allowing two applications to run concurrently on a single compute resource, while "Ultra" stream classes do not support multi-tenancy.

Stream group

Manage how Amazon GameLift Streams streams your applications by using a stream group. A stream group is a collection of compute resources that Amazon GameLift Streams uses to stream your application to end users. When you create a stream group, you specify the hardware configuration (CPU, GPU, RAM) that will run your game (known as the *stream class*), the geographical locations where your game can run, and the number of streams that can run simultaneously in each location (known as *stream capacity*). You can link an application when you create the stream group, or wait until later, but you must link at least one application before you can stream from a stream group. After a stream group has been created, Amazon GameLift Streams allocates compute resources in the locations where you have allocated stream capacity. At this point, you can also associate additional applications to the stream group so that you have a choice of which one to stream.

Stream capacity

Represents the number of concurrent streams that can be active at a time. You set stream capacity per location in each stream group. You configure always-on capacity and maximum capacity. Maximum capacity represents the total streams possible, combining always-on capacity with additional capacity provisioned as needed.

Stream session

Refers to the stream itself. This is an instance of a stream that Amazon GameLift Streams transmits from the server to the end-user. A stream session runs on a compute resource, or stream capacity, that a stream group has allocated. Also referred to as *stream* for short.

Prepare an application in Amazon GameLift Streams

To set up streaming with Amazon GameLift Streams, first you upload the game or other application that you want to stream, then you configure an application resource within Amazon GameLift Streams to define metadata about your game. An Amazon GameLift Streams application consists of the files you uploaded (executable and any supporting files) and a configuration that instructs Amazon GameLift Streams what executable to run when streaming.

Each Amazon GameLift Streams application represents a single version of your content. If you have multiple versions, you must create a separate application for each version. After you create an application, you cannot update the files. If you need to update the executable or any supporting files, you must create a new Amazon GameLift Streams application.

Before you upload

Before you create an Amazon GameLift Streams application, verify that your game adheres to the following limitations.

Name	Default	Adjuse	Description
Files per application	30,000 files	Yes*	The maximum number of files that you can have in an application, in this account.
Single file size	80 GiB	No	The maximum size of a single file in an application. Note that a gibibyte (GiB) equals 1024*1024*1024 bytes.
Application size	100 GiB	Yes*	The maximum total size of an Amazon GameLift Streams application, in this account. Note that a gibibyte (GiB) equals 1024*1024*1024 bytes.

*To request an increase, sign in to the AWS Management Console and open the Service Quotas console to Amazon GameLift Streams, where you can review your current quotas in the **Applied account-level quota value** column and submit a request to increase a value.

> **Note**
>
> To save yourself time and effort, verify that the files you are ready to upload are the correct version of your application. While you can upload new versions later, you will need to repeat the Create an application step for each version.

Upload your application to an Amazon S3 bucket

Now that you have prepared your game for Amazon GameLift Streams, it's time to upload it to an Amazon Simple Storage Service (Amazon S3) bucket in your AWS account.

> **Note**
>
> The Amazon S3 storage class that Amazon GameLift Streams requires is the default **S3 Standard**. Other storage classes like **S3 Glacier** or objects being moved to **Infrequent Access** or **Archive Access** by **S3 Intelligent-Tiering** are not supported by Amazon GameLift Streams.
> To optimize storage cost, you can delete the application from your S3 bucket after you complete Create an application and the application is in **Ready** status.

To upload your application to Amazon S3

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. Create an Amazon S3 bucket. Enter a bucket name and select an AWS Region. This region must be the same as the application and stream group that you will create later. See AWS Regions and streaming locations supported by Amazon GameLift Streams for a list of AWS Regions where Amazon GameLift Streams is available. For the remaining fields, keep the default settings.

   For more instructions, refer to Creating a bucket in the *Amazon Simple Storage Service User Guide*.

3. Open the new bucket and upload the folder with your application files.

> **Warning**
>
> You must upload your application files as an uncompressed folder. Don't upload a `.zip` folder.

Create an application

An Amazon GameLift Streams application is a resource that contains a game or interactive application that runs on Amazon GameLift Streams infrastructure and delivers gameplay experiences to players through cloud streaming. The application executes on AWS compute instances and renders game content that is streamed directly to players' devices over the internet, eliminating the need for players to download, install, or run the game locally.

When you create an Amazon GameLift Streams application, you provide the Amazon S3 URI to the application folder you uploaded to your Amazon S3 bucket, and the relative path to a valid executable or script file.

Amazon GameLift Streams does not keep your application files in sync with the files in the Amazon S3 bucket. If you want to update the files in your Amazon GameLift Streams application, you must create a new Amazon GameLift Streams application.

Console

To create an Amazon GameLift Streams application using the Amazon GameLift Streams console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console. Choose the same AWS Region as the Amazon S3 bucket where you uploaded your set of files. For more information, refer to Choosing a Region in the *AWS Management Console Getting Started Guide*.
2. In the navigation bar, choose **Applications** and then choose **Create application**.
3. In **Runtime settings**, enter the following:
    - **Runtime environment**

This is the runtime environment to run your application on. Amazon GameLift Streams can run on either Windows, Ubuntu 22.04 LTS, or Proton.

You cannot edit this field after the creation workflow.

Choose from one of the following runtime environments.

- For Linux applications:
  - Ubuntu 22.04 LTS (`UBUNTU, 22_04_LTS`)
- For Windows applications:
  - Microsoft Windows Server 2022 Base (`WINDOWS, 2022`)
  - Proton 9.0-2 (`PROTON, 20250516`)
  - Proton 8.0-5 (`PROTON, 20241007`)
  - Proton 8.0-2c (`PROTON, 20230704`)

Review the descriptions and use the comparison checklist to help you select the optimal runtime environment for your application.

4. In **General settings**, enter the following:

   a. **Description**

   This is a human-readable label for your application. This value does not have to be unique. For best practice, use a meaningful description, name, or label for the application. You can edit this field at any time.

   b. **Base path**

   This is the Amazon S3 URI to your application's root folder in the Amazon S3 bucket. The folder and any subfolders should contain your build executable and any supporting files.

   A valid URI is the bucket prefix that contains all the files needed to run and stream the application. Example: A bucket called `mygamebuild` contains three complete versions of the game build files, each in a separate folder. You want to stream the build in the folder `mygamebuild-EN101`. In this example, the URI is `s3://amzn-s3-demo-bucket/mygamebuild-EN101`.

You cannot edit this field after the creation workflow.

c. **Executable launch path**

This is the Amazon S3 URI of the executable file that Amazon GameLift Streams will stream. The file must be contained within the application's root folder. For Windows applications, the file must be a valid Windows executable or batch file with a filename ending in .exe, .cmd, or .bat. For Linux applications, the file must be a valid Linux binary executable or a script that contains an initial interpreter line starting with a shebang ('#!').

You cannot edit this field after the creation workflow.

5. (Optional) In **Application log path**, enter the following:

a. **Application log path**

This is the path (or paths) to the application folder or file which contains logs that you want to save. Specify each log path relative to your application base path. If you use this feature, then at the end of every stream session, Amazon GameLift Streams will copy the file(s) that you specify to the Amazon S3 bucket that you name. The copy operation is not performed recursively in an application folder's subfolders.

To disable logging, remove all application log paths and clear the application log output destination.

You can edit this field at any time.

b. **Application log output**

This is the URI to the Amazon S3 bucket where Amazon GameLift Streams will copy application log files. This field is required if you specify an application log path.

To disable logging, remove all application log paths and clear the application log output destination.

You can edit this field at any time.

To save log files on your behalf, Amazon GameLift Streams must be given permission to your S3 bucket for saving. If you let Amazon GameLift Streams create the bucket for logging, the permission policy will be applied automatically upon creation. If you provide your own bucket, you will need to apply the permission policy, yourself.

Bucket permission policy template

Copy the following policy code and apply it to the bucket that you want to use for application logs. Be sure to replace **amzn-s3-demo-bucket** with the name of your existing S3 bucket.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PutPolicy",
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "gameliftstreams.amazonaws.com"
        ]
      },
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::amzn-s3-demo-bucket/*",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "your 12-digit account id"
        }
      }
    }
  ]
}
```

6. (Optional) In **Tags**, assign tags to this application.

   Tags are labels that can help you organize your AWS resources. For more information, refer to Tagging your AWS resources.

   For example to track application versions, use a tag such as `application-version : my-game-1121`.

7. Choose **Create application**.

CLI

Prerequisite

You must configure the AWS CLI with your user credentials and your chosen AWS Region. For setup instructions, refer to Download the AWS CLI.

To create an application using the AWS CLI

In your AWS CLI use the CreateApplication command, customized for your content.

```
aws gameliftstreams create-application \
    --description "MyGame v1" \
    --runtime-environment '{"Type":"PROTON", "Version":"20241007"}' \
    --executable-path "launcher.exe" \
    --application-source-uri "s3://amzn-s3-demo-bucket/example"
```

where

- `description`:

  This is a human-readable label for your application. This value does not have to be unique. For best practice, use a meaningful description, name, or label for the application. You can edit this field at any time.

- `runtime-environment`:

  This is the runtime environment to run your application on. Amazon GameLift Streams can run on either Windows, Ubuntu 22.04 LTS, or Proton.

  You cannot edit this field after the creation workflow.

  Choose from one of the following runtime environments.

  - For Linux applications:
    - Ubuntu 22.04 LTS (`Type=UBUNTU, Version=22_04_LTS`)
  - For Windows applications:
    - Microsoft Windows Server 2022 Base (`Type=WINDOWS, Version=2022`)
    - Proton 9.0-2 (`Type=PROTON, Version=20250516`)
    - Proton 8.0-5 (`Type=PROTON, Version=20241007`)
    - Proton 8.0-2c (`Type=PROTON, Version=20230704`)

- `application-source-uri`:

This is the Amazon S3 URI to your application's root folder in the Amazon S3 bucket. The folder and any subfolders should contain your build executable and any supporting files.

A valid URI is the bucket prefix that contains all the files needed to run and stream the application. Example: A bucket called `mygamebuild` contains three complete versions of the game build files, each in a separate folder. You want to stream the build in the folder `mygamebuild-EN101`. In this example, the URI is `s3://amzn-s3-demo-bucket/mygamebuild-EN101`.

You cannot edit this field after the creation workflow.

- `executable-path`:

  This is the relative path and file name of the executable file that Amazon GameLift Streams will stream. Specify a path relative to the `application-source-uri`. The file must be contained within the application's root folder. For Windows applications, the file must be a valid Windows executable or batch file with a filename ending in .exe, .cmd, or .bat. For Linux applications, the file must be a valid Linux binary executable or a script that contains an initial interpreter line starting with a shebang ('#!').

  You cannot edit this field after the creation workflow.

If the request is successful, Amazon GameLift Streams returns a response similar to the following:

```
{
    "Arn": "arn:aws:gameliftstreams:us-west-2:111122223333:application/a-9ZY8X7Wv6",
    "Description": "MyGame v1",
    "RuntimeEnvironment": {
        "Type": "PROTON",
        "Version": "20241007"
    },
    "ExecutablePath": "launcher.exe",
    "ApplicationSourceUri": "s3://amzn-s3-demo-bucket/example",
    "Id": "a-9ZY8X7Wv6",
    "Status": "PROCESSING",
    "CreatedAt": "2022-11-18T15:47:11.924000-08:00",
    "LastUpdatedAt": "2022-11-18T15:47:11.924000-08:00"
}
```

To check the status of your application, call the GetApplication command, as shown in the following example.

```
aws gameliftstreams get-application /
    --identifier a-9ZY8X7Wv6
```

Amazon GameLift Streams takes a few minutes to prepare your application. During this time, the new application is in **Processing** status. When your application is in **Ready** status, you can go to the next step, Create a stream group.

If the request returns an error, or if the application is created but placed in an **Error** status, make sure that you're working with user credentials that include access to both Amazon S3 and Amazon GameLift Streams.

> **Note**
>
> When an application is in **Ready** status, Amazon GameLift Streams has successfully copied your application files to its private Amazon S3 bucket. You can delete your original application files without affecting your new application. This also helps you optimize on storage cost. For more information, see Delete an application.

Edit an application

You can update the settings for any application in **Ready** status. If you make changes to an existing application, these changes impact the streaming behavior for both new and existing stream groups.

Console

To edit an application in the Amazon GameLift Streams console

1. In the navigation bar, choose **Applications** to view a list of your existing applications. Choose the application you want to edit.
2. In the application details page, locate the section that contains the settings you want to change and choose **Edit** or **Manage tags** accordingly.
3. You can change the following settings:

Short description

This is a human-readable label for your application. This value does not have to be unique. For best practice, use a meaningful description, name, or label for the application. You can edit this field at any time.

Application log path

This is the path (or paths) to the application folder or file which contains logs that you want to save. Specify each log path relative to your application base path. If you use this feature, then at the end of every stream session, Amazon GameLift Streams will copy the file(s) that you specify to the Amazon S3 bucket that you name. The copy operation is not performed recursively in an application folder's subfolders.

To disable logging, remove all application log paths and clear the application log output destination.

You can edit this field at any time.

Application log output

This is the URI to the Amazon S3 bucket where Amazon GameLift Streams will copy application log files. This field is required if you specify an application log path.

To disable logging, remove all application log paths and clear the application log output destination.

You can edit this field at any time.

To save log files on your behalf, Amazon GameLift Streams must be given permission to your S3 bucket for saving. If you let Amazon GameLift Streams create the bucket for logging, the permission policy will be applied automatically upon creation. If you provide your own bucket, you will need to apply the permission policy, yourself.

For more information, see Application log bucket permission policy.

Tags

Tags are labels that can help you organize your AWS resources. For more information, refer to Tagging your AWS resources.

For example to track application versions, use a tag such as `application-version : my-game-1121`.

4. Choose **Save changes**. The Amazon GameLift Streams console returns to the application details page, displaying the updated settings.

CLI

Prerequisite

You must configure the AWS CLI with your user credentials and your chosen AWS Region. For setup instructions, refer to Download the AWS CLI.

To edit an application using the AWS CLI

In your AWS CLI use the UpdateApplication command, customized for your content.

```
aws gameliftstreams update-application \
    --identifier a-9ZY8X7Wv6 \
    --description "MyGame v2" \
    --application-log-paths '[".\\logs"]' \
    --application-log-output-uri "s3://amzn-s3-demo-bucket/mygame"
```

where

- `identifier`: The application to edit.

  This value is an Amazon Resource Name (ARN) or ID that uniquely identifies the application resource.

  ARN example: `arn:aws:gameliftstreams:us-west-2:111122223333:application/a-9ZY8X7Wv6`

  ID example: `a-9ZY8X7Wv6`
- `description`:

  This is a human-readable label for your application. This value does not have to be unique. For best practice, use a meaningful description, name, or label for the application. You can edit this field at any time.
- `application-log-paths`:

This is the path (or paths) to the application folder or file which contains logs that you want to save. Specify each log path relative to your application base path. If you use this feature, then at the end of every stream session, Amazon GameLift Streams will copy the file(s) that you specify to the Amazon S3 bucket that you name. The copy operation is not performed recursively in an application folder's subfolders.

To disable logging, remove all application log paths and clear the application log output destination.

You can edit this field at any time.

- `application-log-output-uri`:

  This is the URI to the Amazon S3 bucket where Amazon GameLift Streams will copy application log files. This field is required if you specify an application log path.

  To disable logging, remove all application log paths and clear the application log output destination.

  You can edit this field at any time.

  To save log files on your behalf, Amazon GameLift Streams must be given permission to your S3 bucket for saving. If you let Amazon GameLift Streams create the bucket for logging, the permission policy will be applied automatically upon creation. If you provide your own bucket, you will need to apply the permission policy, yourself.

  For more information, see Application log bucket permission policy.

Delete an application

Delete an application if you no longer need it. This action permanently deletes the application, including the application content files stored with Amazon GameLift Streams. However, this does not delete the original files that you uploaded to your Amazon S3 bucket; you can delete these any time after Amazon GameLift Streams creates an application, which is the only time Amazon GameLift Streams accesses your Amazon S3 bucket.

You can only delete an application that meets the following conditions:

- The application is in the **Ready** or **Error** state.

- An application is not streaming in any ongoing stream session. You must wait until the client ends the stream session or call TerminateStreamSession in the Amazon GameLift Streams API to end the stream.

If the application is linked to any stream groups, you must unlink it from all associated stream groups before you can delete it. In the console, a dialog box will walk you through this process.

Console

To delete an application using the Amazon GameLift Streams console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console.
2. In the navigation bar, choose **Applications** to view a list of your existing applications. Choose the application you want to delete.
3. In the application detail page, choose **Delete**.
4. In the **Delete** dialog box, confirm the delete action.

CLI

Prerequisite

You must configure the AWS CLI with your user credentials and your chosen AWS Region. For setup instructions, refer to Download the AWS CLI.

To delete an application using the AWS CLI

In your AWS CLI use the DeleteApplication command, customized for your content.

```
aws gameliftstreams delete-application \
    --identifier arn:aws:gameliftstreams:us-west-2:111122223333:application/
a-9ZY8X7Wv6
```

where

- `identifier`: The application to delete.

  This value is an Amazon Resource Name (ARN) or ID that uniquely identifies the application resource.

ARN example: `arn:aws:gameliftstreams:us-west-2:111122223333:application/a-9ZY8X7Wv6`

ID example: `a-9ZY8X7Wv6`

Amazon GameLift Streams begins deleting the application. During this time, the application is in `Deleting` status. After Amazon GameLift Streams deletes the application, you can no longer retrieve it.

Application log bucket permission policy

If you provide your own application log Amazon S3 bucket, you will need to apply a permission policy to the bucket so that Amazon GameLift Streams can save log files to the bucket. Use the following template to update the permissions in Amazon S3.

Bucket permission policy template

Copy the following policy code and apply it to the bucket that you want to use for application logs. Be sure to replace **amzn-s3-demo-bucket** with the name of your existing S3 bucket.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PutPolicy",
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "gameliftstreams.amazonaws.com"
        ]
      },
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::amzn-s3-demo-bucket/*",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "your 12-digit account id"
        }
      }
    }
  ]
```

```
}
```

> **Note**
>
> Amazon GameLift Streams does not permit cross-account resource access. The Amazon S3 bucket must be owned by the same AWS account as the application resource. Although this is strongly enforced by the service, it is a best practice to always include `aws:SourceAccount` or `aws:SourceArn` conditions to prevent the confused deputy problem when granting permission to any AWS service.

Linked stream groups

If you want to stream multiple applications by using the same pool of compute resources, you can link multiple applications to the same stream group. Similarly, if you want to stream an application by using different sets of compute resources, you can link an application to multiple stream groups.

For more information about linking applications to stream groups, refer to Overview of multi-application stream groups.

Manage streaming with an Amazon GameLift Streams stream group

After you set up an Amazon GameLift Streams application, you're ready to manage and deploy compute resources to run and stream your application. An Amazon GameLift Streams *stream group* represents a collection of these compute resources. You specify the maximum number of concurrent streams to support by scaling the stream capacity.

Amazon GameLift Streams allocates compute resources in the AWS Region where you create a stream group. You can also add remote locations to a stream group and manage capacity per location. It's a best practice to host stream sessions in locations that are geographically near your end users. This helps minimize latency and improve stream quality. For more information, refer to AWS Regions and streaming locations supported by Amazon GameLift Streams.

In a stream group, you can specify one or more Amazon GameLift Streams applications that the stream group can stream. A single application can be in multiple stream groups, so you can set up different configurations or types of compute resources to stream the same application. For

example, to provide two graphics-quality options for streaming an application, you can set up two stream groups with different stream class configurations and link them to the same application.

Conversely, a single stream group can have multiple applications: the *default application*, which you can set when you create the stream group, and additional *linked applications*. For more information, refer to Overview of multi-application stream groups.

How you relate your stream groups and applications together depends on your use case, but the relationship can be many-to-many.

Stream groups should be recreated every 3-4 weeks to pick up important service updates and fixes. For more information, refer to Stream group lifecycle.

Topics

- About stream capacity
- Capacity and service quotas
- About locations
- Create a stream group
- Edit general settings
- Edit capacity
- Capacity scale-down behavior
- Add locations in a stream group
- Remove locations in a stream group
- Delete a stream group
- Linked applications
- Stream group lifecycle
- Stream group maintenance

About stream capacity

You manage the number of streams you can deliver concurrently to end-users by setting the stream group's capacity, or *stream capacity*. Stream capacity represents the number of concurrent stream sessions a stream group can support. It is configured at each location.

- **Always-on capacity:** This setting, if non-zero, indicates minimum streaming capacity which is allocated to you and is never released back to the service. You pay for this base level of capacity at all times, whether used or idle.
- **Maximum capacity:** This indicates the maximum capacity that the service can allocate for you. Newly created streams may take a few minutes to start. Capacity is released back to the service when idle. You pay for capacity that is allocated to you until it is released.
- **Target-idle capacity:** This indicates idle capacity which the service pre-allocates and hold for you in anticipation of future activity. This helps to insulate your users from capacity-allocation delays. You pay for capacity which is held in this intentional idle state.

If you have a stream group with an maximum capacity set to 100 at a location, this means the stream group has enough resources to stream to 100 end-users concurrently at that location. You can increase or decrease the stream capacity at any time, at each location (up to your current quota amount) to meet changes in user demand.

Amazon GameLift Streams first tries to fulfill new session requests using idle capacity which is already allocated to you. If this causes the amount of idle capacity to drop below your target idle capacity, then new capacity is allocated asynchronously. If no idle capacity is available, the request is paused while new capacity is allocated on demand, up to the maximum capacity for the stream group. If the maximum is reached and there is still no idle capacity available, the session request will wait for an existing session to terminate and free capacity.

When sessions terminate, the corresponding capacity is marked as idle. If there is more idle capacity than the target idle value, the excess capacity will be deallocated and returned to the service after a brief delay. The service will not deallocate idle capacity if that would drop your capacity level below the configured minimum (which could be zero).

When specifying the stream capacity in stream groups with multi-tenant stream classes (which can stream more than 1 session per compute resource), the capacity must be a multiple of the tenancy. For example, the `gen6n_high` stream class has a multi-tenancy of 2. That means each compute resource that gets allocated in your stream group can stream to 2 clients. Therefore, the capacity you request must be in multiples of 2.

Scaling the capacity reflects in your total cost for the stream group. Ensure that you set up billing alerts to manage your Amazon GameLift Streams costs. Refer to Create billing alerts to monitor usage.

To change stream group capacity, edit your stream group settings and enter new values for the capacity settings. When you change always-on capacity, Amazon GameLift Streams adjusts allocated resources to match the new value by provisioning new resources or shutting down existing ones. Increasing always-on capacity can take more than a few minutes if resources aren't immediately available. Decreasing always-on capacity takes a few minutes to deprovision allocated resources.

Example: Stream capacity configurations

The following examples demonstrate common stream capacity configurations for different use cases:

1. **Cost-conscious development phase:** You are a developer who wants to save costs. You set `Minimum (always-on) capacity` = 0, `Maximum capacity` = 10, and `Target Idle (pre-warmed) capacity` = 1. This keeps at least one session available for fast start up.
2. **Planned event with fixed demand:** You want fast session starts for a planned event with known demand. You set `Minimum (always-on) capacity` = 200, `Maximum capacity` = 200, and `Target Idle (pre-warmed) capacity` = 0. You pay only for 200 capacity. No scaling delays happen because demand is known.
3. **Large-scale event with burst capacity:** You are planning for 1,000 users with 100 new sessions per minute at peak times. You set `Minimum` = 0, `Maximum` = 1,000, and `Target Idle` = 100. This saves money when idle. This keeps at least 100 sessions available for fast start up.

> **Note**
>
> The `OnDemandCapacity` input parameter is deprecated. Use `MaximumCapacity` instead when configuring capacity through the API.

Capacity and service quotas

Usage of Amazon GameLift Streams is subject to service quotas that limit the total number of GPUs (compute resources) you can configure for streaming in your account. The default quotas and the utilization of the quotas can be viewed in the Service Quota Console for GameLift Streams. Understanding how these quotas interact with stream capacity helps you plan your streaming infrastructure and avoid capacity limitations.

More specifically, the GPU service quotas specify the maximum number of GPUs of a particular stream class family you can request per location across all stream groups in your account. For example, if your account has a limit of 5 gen6n GPUs in us-west-2, the sum of gen6n GPUs needed to provide the total stream capacity in us-west-2 for all of your stream groups must be less than or equal to 5. This includes GPUs for both always-on and on-demand capacity.

When calculating the total stream capacity provided by these GPUs, it's important to remember that multi-tenant stream classes support streaming more than one session per GPU. Therefore, if you're using multi-tenant stream classes in your stream groups, such as gen6n_high, you will need to take this into account when determining how the capacity will count against your quota. Single-tenant stream classes, such as gen6n_ultra and gen6n_ultra_win2022, dedicate one GPU per stream session.

Example: How quotas affect capacity

The following example demonstrates how service quotas interact with stream capacity across multiple stream groups and locations. In this example, assume your account has a quota of 10 gen6n GPUs per location.

1. **Create a single-tenant stream group:** You create a stream group using the gen6n_ultra stream class with 5 total capacity (always-on plus on-demand) in us-east-2. Because this stream class has 1:1 tenancy (1 stream per GPU), you need 5 GPUs for 5 total capacity. This leaves you with 5 remaining GPUs in us-east-2.
2. **Create a multi-tenant stream group:** You create another stream group using the gen6n_high stream class with 6 total capacity in us-east-2. Because this stream class has 1:2 tenancy (2 streams per GPU), you only need 3 GPUs for 6 total capacity. This leaves you with 2 remaining GPUs in us-east-2.
3. **Add capacity in other locations:** After creating these stream groups, you have 2 remaining GPUs in us-east-2, but you still have 10 GPUs available in other locations such as us-west-2 or eu-west-1. You can add these locations to either of the stream groups you created earlier or create new stream groups that have these locations.

This example shows that quotas are enforced per location and across all of your stream groups, allowing you to distribute your streaming capacity across multiple geographic regions while staying within service limits.

> **Note**
>
> You can view your Applied account level or default quota, including the utilization of those quotas, in the Service Quotas console by selecting the GameLift Streams as the AWS service. For more information, see Amazon GameLift Streams service quotas.

About locations

The location is where Amazon GameLift Streams allocates compute resources to host your application and stream to users. For lower latency and better quality, you should choose locations closer to your users. By default, you can stream from the AWS Region where you created your stream group, known as the *primary location*. Additionally, a stream group can extend its coverage to stream from other supported locations, known as *remote locations*.

For a complete list of supported locations, refer to AWS Regions and streaming locations.

Multi-location stream group

A stream group that's configured to host applications and stream sessions from multiple locations, in addition to the primary location (the AWS Region where you created the stream group). You manage capacity for each location.

Create a stream group

Console

To create a stream group in the Amazon GameLift Streams console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console. Choose the AWS Region where you want to create your stream group. This Region must be the same as that of the application that you want to stream with the stream group. For more information, refer to Choosing a Region in the *AWS Management Console Getting Started Guide*.
2. To open the creation workflow, in the navigation pane, choose **Stream groups**, and then choose **Create stream group**.
3. In **Define stream group**, enter the following:

a. **Description**

A human-readable label for your stream group. This value doesn't have to be unique. As a best practice, use a meaningful description, name, or label for the stream group. You can edit this field at any time.

b. **Tags**

Tags are labels that can help you organize your AWS resources. For more information, refer to Tagging your AWS resources.

4. In **Select stream class**, choose a stream class for the stream group.

- **Stream class options**

The type of compute resources to run and stream applications with. This choice impacts the quality of the streaming experience and the cost. You can specify only one stream class per stream group. Choose the class that best fits your application.

Stream class	Description
`gen6n_pro_win2022`	(NVIDIA, pro) Supports applications with extremely high 3D scene complexity which require maximum resources. Runs applications on Microsoft Windows Server 2022 Base and supports DirectX 12. Compatible with Unreal Engine versions up through 5.6, 32 and 64-bit applications, and anti-cheat technology. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 16. RAM: 64 GB. VRAM: 24 GB. Tenancy: Supports up to one concurrent stream session.
`gen6n_pro`	(NVIDIA, pro) Supports applications with extremely high 3D scene complexity which require maximum resources. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 16. RAM: 64 GB. VRAM: 24 GB.

Stream class	Description
	Tenancy: Supports up to one concurrent stream session.
`gen6n_ultra_win2022`	(NVIDIA, ultra) Supports applications with high 3D scene complexity. Runs applications on Microsoft Windows Server 2022 Base and supports DirectX 12. Compatible with Unreal Engine versions up through 5.6, 32 and 64-bit applications, and anti-cheat technology. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 24 GB. Tenancy: Supports up to one concurrent stream session.
`gen6n_ultra`	(NVIDIA, ultra) Supports applications with high 3D scene complexity. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 24 GB. Tenancy: Supports up to one concurrent stream session.
`gen6n_high`	(NVIDIA, high) Supports applications with moderate-to-high 3D scene complexity. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 4. RAM: 16 GB. VRAM: 12 GB. Tenancy: Supports up to two concurrent stream sessions.
`gen6n_medium`	(NVIDIA, medium) Supports applications with moderate 3D scene complexity. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 2. RAM: 8 GB. VRAM: 6 GB. Tenancy: Supports up to four concurrent stream sessions.

Stream class	Description
gen6n_small	(NVIDIA, small) Supports applications with lightweight 3D scene complexity and low CPU usage. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 1. RAM: 4 GB. VRAM: 2 GB. Tenancy: Supports up to twelve concurrent stream sessions.
gen5n_win2022	(NVIDIA, ultra) Supports applications with extremely high 3D scene complexity. Runs applications on Microsoft Windows Server 2022 Base and supports DirectX 12 and DirectX 11. Supports Unreal Engine up through version 5.6, 32 and 64-bit applications, and anti-cheat technology. Uses NVIDIA A10G Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 24 GB. Tenancy: Supports one concurrent stream session.
gen5n_high	(NVIDIA, high) Supports applications with moderate-to-high 3D scene complexity. Uses NVIDIA A10G Tensor Core GPU. Resources per application: vCPUs: 4. RAM: 16 GB. VRAM: 12 GB. Tenancy: Supports up to two concurrent stream sessions.
gen5n_ultra	(NVIDIA, ultra) Supports applications with extremely high 3D scene complexity. Uses NVIDIA A10G Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 24 GB. Tenancy: Supports one concurrent stream session.

Stream class	Description
gen4n_win2022	(NVIDIA, ultra) Supports applications with high 3D scene complexity. Runs applications on Microsoft Windows Server 2022 Base and supports DirectX 12 and DirectX 11. Supports Unreal Engine up through version 5.6, 32 and 64-bit applicati ons, and anti-cheat technology. Uses NVIDIA T4 Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 16 GB. Tenancy: Supports one concurrent stream session.
gen4n_high	(NVIDIA, high) Supports applications with moderate-to-high 3D scene complexity. Uses NVIDIA T4 Tensor Core GPU. Resources per application: vCPUs: 4. RAM: 16 GB. VRAM: 8 GB. Tenancy: Supports up to two concurrent stream sessions.
gen4n_ultra	(NVIDIA, ultra) Supports applications with high 3D scene complexity. Uses NVIDIA T4 Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 16 GB. Tenancy: Supports one concurrent stream session.

To continue, choose **Next**.

5. In **Link application**, choose an application that you want to stream, or select "**No application**" to choose one at a later time. You can edit the stream group after it has been created to add or remove applications. You can only link an application that's in Ready status and has a runtime that's compatible with the stream class you've chosen. By default, these are the only applications that are shown in the table. To see all applications in Ready status, choose `All runtimes` in the drop down list.

> **Note**
>
> If you don't see your application listed, then check the current AWS Region setting. You can only link an application to a stream group that's in the same Region.

To continue, choose **Next**.

6. In **Configure stream settings**, under **Locations and capacity**, choose one or more locations where your stream group will have capacity to stream your application. By default, the region where you create the stream group, known as the *primary location*, has already been added to your stream group and cannot be removed. You can add additional locations by checking the box next to each location that you want to add. For lower latency and better quality streaming, you should choose locations closer to your users.

   For each location, you can specify its *streaming capacity*. Stream capacity represents the number of concurrent streams that can be active at a time. You set stream capacity per location in each stream group.

   - **Always-on capacity:** This setting, if non-zero, indicates minimum streaming capacity which is allocated to you and is never released back to the service. You pay for this base level of capacity at all times, whether used or idle.
   - **Maximum capacity:** This indicates the maximum capacity that the service can allocate for you. Newly created streams may take a few minutes to start. Capacity is released back to the service when idle. You pay for capacity that is allocated to you until it is released.
   - **Target-idle capacity:** This indicates idle capacity which the service pre-allocates and hold for you in anticipation of future activity. This helps to insulate your users from capacity-allocation delays. You pay for capacity which is held in this intentional idle state.

   You can increase or decrease your total stream capacity at any time to meet changes in user demand for a location by adjusting either capacity. Amazon GameLift Streams fulfills streaming requests using the idle, pre-allocated resources in the always-on capacity pool if any are available. If all always-on capacity is in use, Amazon GameLift Streams will provision additional compute resources up to the maximum number specified in on-demand capacity. As allocated capacity scales, the change is reflected in your total cost for the stream group.

Linked applications will be automatically replicated to each enabled location. An application must finish replicating in a remote location before the remote location can host a stream. To check on the replication status, open the stream group after it has been created and refer to the **Replication status** column in the table of linked applications. Click on the current status to see the replication status for each added location.

> **Note**
>
> Application data will be stored in all enabled locations including the primary location for this stream group. Stream session data will be stored in both the primary location and the location where the streaming occurred.

7. In **Review and create stream group**, verify your stream group configuration and make changes as needed. When everything is correct, choose **Create stream group**.

CLI

Prerequisite

You must configure the AWS CLI with your user credentials and your chosen AWS Region. For setup instructions, refer to Download the AWS CLI.

To create a stream group using the AWS CLI

In your AWS CLI use the CreateStreamGroup command, customized for your content.

```
aws gameliftstreams create-stream-group \
    --description "Test_gen4_high" \
    --default-application-identifier arn:aws:gameliftstreams:us-
west-2:111122223333:application/a-9ZY8X7Wv6 \
    --stream-class gen4n_high \
    --location-configurations '[{"LocationName": "us-east-1", "AlwaysOnCapacity": 2,
 "MaximumCapacity": 6, "TargetIdleCapacity": 1}]'
```

where

`description`:

A human-readable label for your stream group. This value doesn't have to be unique. As a best practice, use a meaningful description, name, or label for the stream group. You can edit this field at any time.

`default-application-identifier`

The Amazon Resource Name (ARN) value or ID assigned to an Amazon GameLift Streams application resource. The application must be in `READY` status.

ARN example: `arn:aws:gameliftstreams:us-west-2:111122223333:application/a-9ZY8X7Wv6`

ID example: `a-9ZY8X7Wv6`

`stream-class`

Stream class options

The type of compute resources to run and stream applications with. This choice impacts the quality of the streaming experience and the cost. You can specify only one stream class per stream group. Choose the class that best fits your application.

Stream class	Description
`gen6n_pro_win2022`	(NVIDIA, pro) Supports applications with extremely high 3D scene complexity which require maximum resources. Runs applications on Microsoft Windows Server 2022 Base and supports DirectX 12. Compatible with Unreal Engine versions up through 5.6, 32 and 64-bit applications, and anti-cheat technology. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 16. RAM: 64 GB. VRAM: 24 GB. Tenancy: Supports up to one concurrent stream session.
`gen6n_pro`	(NVIDIA, pro) Supports applications with extremely high 3D scene complexity which require maximum resources. Uses NVIDIA L4 Tensor Core GPU.

Stream class	Description
	Resources per application: vCPUs: 16. RAM: 64 GB. VRAM: 24 GB. Tenancy: Supports up to one concurrent stream session.
`gen6n_ultra_win2022`	(NVIDIA, ultra) Supports applications with high 3D scene complexity. Runs applications on Microsoft Windows Server 2022 Base and supports DirectX 12. Compatible with Unreal Engine versions up through 5.6, 32 and 64-bit applications, and anti-cheat technology. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 24 GB. Tenancy: Supports up to one concurrent stream session.
`gen6n_ultra`	(NVIDIA, ultra) Supports applications with high 3D scene complexity. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 24 GB. Tenancy: Supports up to one concurrent stream session.
`gen6n_high`	(NVIDIA, high) Supports applications with moderate-to-high 3D scene complexity. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 4. RAM: 16 GB. VRAM: 12 GB. Tenancy: Supports up to two concurrent stream sessions.
`gen6n_medium`	(NVIDIA, medium) Supports applications with moderate 3D scene complexity. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 2. RAM: 8 GB. VRAM: 6 GB. Tenancy: Supports up to four concurrent stream sessions.

Stream class	Description
gen6n_small	(NVIDIA, small) Supports applications with lightweight 3D scene complexity and low CPU usage. Uses NVIDIA L4 Tensor Core GPU. Resources per application: vCPUs: 1. RAM: 4 GB. VRAM: 2 GB. Tenancy: Supports up to twelve concurrent stream sessions.
gen5n_win2022	(NVIDIA, ultra) Supports applications with extremely high 3D scene complexity. Runs applications on Microsoft Windows Server 2022 Base and supports DirectX 12 and DirectX 11. Supports Unreal Engine up through version 5.6, 32 and 64-bit applications, and anti-cheat technology. Uses NVIDIA A10G Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 24 GB. Tenancy: Supports one concurrent stream session.
gen5n_ultra	(NVIDIA, ultra) Supports applications with extremely high 3D scene complexity. Uses NVIDIA A10G Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 24 GB. Tenancy: Supports one concurrent stream session.
gen5n_high	(NVIDIA, high) Supports applications with moderate-to-high 3D scene complexity. Uses NVIDIA A10G Tensor Core GPU. Resources per application: vCPUs: 4. RAM: 16 GB. VRAM: 12 GB. Tenancy: Supports up to two concurrent stream sessions.

Stream class	Description
gen4n_win2022	(NVIDIA, ultra) Supports applications with high 3D scene complexity. Runs applications on Microsoft Windows Server 2022 Base and supports DirectX 12 and DirectX 11. Supports Unreal Engine up through version 5.6, 32 and 64-bit applicati ons, and anti-cheat technology. Uses NVIDIA T4 Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 16 GB. Tenancy: Supports one concurrent stream session.
gen4n_ultra	(NVIDIA, ultra) Supports applications with high 3D scene complexity. Uses NVIDIA T4 Tensor Core GPU. Resources per application: vCPUs: 8. RAM: 32 GB. VRAM: 16 GB. Tenancy: Supports one concurrent stream session.
gen4n_high	(NVIDIA, high) Supports applications with moderate-to-high 3D scene complexity. Uses NVIDIA T4 Tensor Core GPU. Resources per application: vCPUs: 4. RAM: 16 GB. VRAM: 8 GB. Tenancy: Supports up to two concurrent stream sessions.

`location-configurations`

A set of locations to add to this stream group, and their capacities. By default, if no capacities are specified, Amazon GameLift Streams will only allocate enough always-on stream capacity to start one stream in the location where the stream group is created. For a complete list of locations that Amazon GameLift Streams supports, refer to AWS Regions and streaming locations supported by Amazon GameLift Streams.

Values for capacity must be whole number multiples of the tenancy value of the stream group's stream class.

If the request is successful, then Amazon GameLift Streams returns a response similar to the following:

```
{
    "Arn": "arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/
sg-1AB2C3De4",
    "Description": "Test_gen4_high",
    "DefaultApplication": {
        "Id": "a-9ZY8X7Wv6"
    },
    "StreamClass": "gen4n_high",
    "Id": "sg-1AB2C3De4",
    "Status": "ACTIVATING",
    "LastUpdatedAt": "2024-11-18T15:49:01.482000-08:00",
    "CreatedAt": "2024-11-18T15:49:01.482000-08:00"
}
```

Amazon GameLift Streams begins searching for unallocated computing resources and provisioning them for the new stream group, which can take several minutes. During this time, the new stream group is in **Activating** status.

You can adjust the stream group's capacity when its status is **Active**. For more information, refer to Edit capacity.

When the stream group is in **Active** status, it's ready to deploy resources for streaming. To start streaming, refer to Start stream sessions with Amazon GameLift Streams.

Edit general settings

Amazon GameLift Streams groups the following settings together in the console under **Stream group settings**: **Status**, **Stream group ID**, **Description**, **Stream group ARN**, and **Stream class**. Of these, the only one that you can update without creating a new stream group is **Description**.

Console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console.
2. In the navigation bar, choose **Stream groups** to view a list of your existing stream groups. Choose the stream group you want to edit.
3. In the stream group detail page, choose **Edit settings**.

4. To update the description, enter a new value.

CLI

Prerequisite

You must configure the AWS CLI with your user credentials and your chosen AWS Region. For setup instructions, refer to Download the AWS CLI.

To edit a stream group's description using the AWS CLI

In your AWS CLI use the UpdateStreamGroup command, customized for your content.

```
aws gameliftstreams update-stream-group \
    --identifier arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/
sg-1AB2C3De4 \
    --description "MyGame - Ultra"
```

where

`identifier`

An Amazon Resource Name (ARN) or ID that uniquely identifies the stream group resource.

ARN example: `arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/sg-1AB2C3De4`

ID example: `sg-1AB2C3De4`

`description`

A human-readable label for your stream group. This value doesn't have to be unique. As a best practice, use a meaningful description, name, or label for the stream group. You can edit this field at any time.

Edit capacity

Scale your stream groups by adjusting the capacity for each location.

Refer to Amazon GameLift Streams service quotas to learn more about stream group capacity quotas per AWS account, per location, and how to increase these quotas.

Console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console.
2. In the navigation bar, choose **Stream groups** to view a list of your existing stream groups. Choose the stream group you want to edit.
3. In the stream group detail page, choose **Edit configuration**.
4. For each location, enter new always-on capacity, maximum capacity, and target-idle capacity values in the relevant cells in the table. Values for capacity must be whole number multiples of the tenancy value of the stream group's stream class.

   If you set the always-on capacity value to zero, the stream group won't allocate any hosts to stream.

CLI

Prerequisite

You must configure the AWS CLI with your user credentials and your chosen AWS Region. For setup instructions, refer to Download the AWS CLI.

To edit stream capacity using the AWS CLI

In your AWS CLI use the UpdateStreamGroup command, customized for your content.

```
aws gameliftstreams update-stream-group \
    --identifier arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/
sg-1AB2C3De4 \
    --location-configurations '[{"LocationName": "us-east-1", "AlwaysOnCapacity": 4,
 "MaximumCapacity": 8}, \
        {"LocationName": "ap-northeast-1", "AlwaysOnCapacity": 0,
 "MaximumCapacity": 2, "TargetIdleCapacity": 1}]'
```

where

`identifier`

An Amazon Resource Name (ARN) or ID that uniquely identifies the stream group resource.

ARN example: `arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/sg-1AB2C3De4`

ID example: `sg-1AB2C3De4`

`location-configurations`

A set of locations to update in this stream group with their new capacities. Values for capacity must be whole number multiples of the tenancy value of the stream group's stream class.

When you update a stream group location's capacity, Amazon GameLift Streams will begin processing your request, which can take a some time. During this time, Amazon GameLift Streams works to allocate or release resources in the stream group as needed to meet the desired always-on stream capacity you set. You can view the provisioning status of your stream capacity by viewing the **Stream group details** page in the Amazon GameLift Streams console, or by calling the GetStreamGroup API.

When your stream group is in **Active** status, has available stream capacity, and the application has finished replicating to the location where you want to stream, you can start streaming. For more information, refer to Start stream sessions with Amazon GameLift Streams.

Capacity scale-down behavior

When you scale down capacity, Amazon GameLift Streams waits until the host is idle before releasing it. Since a host can support 1 or 2 sessions, the host is idle only when all active sessions on the host end. A stream session ends when the user ends their session or the session times out. Therefore, in extreme situations when existing sessions are allowed to reach the maximum possible duration, it may take up to 24 hours to reach the desired capacity. If you want to force all active stream sessions in a stream group to end, you can delete the stream group in the console or by using the DeleteStreamGroup API, or you can use the TerminateStreamSession API to end active sessions one at a time.

Add locations in a stream group

Console

To add locations to a stream group using the Amazon GameLift Streams console

1. In the navigation bar, choose **Stream groups** to view a list of your existing stream groups. Choose the stream group you want to add new locations to.
2. In the **Stream group details** page, choose **Edit configuration**.

3. Select the checkbox next to the location(s) you want to add to this stream group, and then set their capacities.
4. Review the summary of your selected locations, including the cost for stream capacity. Choose **Save** to confirm your selection.

CLI

Prerequisite

You must configure the AWS CLI with your user credentials and your chosen AWS Region. For setup instructions, refer to Download the AWS CLI.

To add locations to a stream group using the AWS CLI

In your AWS CLI use the AddStreamGroupLocations command, customized for your content.

```
aws gameliftstreams add-stream-group-locations \
    --identifier arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/
sg-1AB2C3De4
    --location-configurations '[{"LocationName": "us-east-1", "AlwaysOnCapacity": 2,
 "MaximumCapacity": 4, "TargetIdleCapacity": 1}]'
```

where

`identifier`

An Amazon Resource Name (ARN) or ID that uniquely identifies the stream group resource.

ARN example: `arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/sg-1AB2C3De4`

ID example: `sg-1AB2C3De4`

`location-configurations`

A set of locations to add to this stream group, and their capacities. For a complete list of locations that Amazon GameLift Streams supports, refer to AWS Regions and streaming locations supported by Amazon GameLift Streams.

Values for capacity must be whole number multiples of the tenancy value of the stream group's stream class.

When your application has completed replicating to the new location(s) and your stream group has available stream capacity, you can start streaming from the new location(s). For more information on streaming, refer to Start stream sessions with Amazon GameLift Streams. Amazon GameLift Streams will begin processing your request. During this time, Amazon GameLift Streams works to replicate your application and allocate compute resources in the new locations. You can view the status of the replication from the **Linked applications** section of the **Stream group details** page by clicking on the status in the **Replication status** column.

Remove locations in a stream group

To stop using compute resources from specific locations, you can remove the locations from your stream group. You cannot remove the primary location of a stream group. However, if you don't want compute resources in that location, then you can set the stream capacities to zero.

> **Warning**
>
> When you remove a location in a stream group, Amazon GameLift Streams disconnects active streams in that location, which stops the stream of any connected end users.

Console

To remove locations from a stream group using the Amazon GameLift Streams console

1. In the navigation pane, choose **Stream groups** to view a list of your existing stream groups.
2. Choose the name of the stream group that you want to remove locations from.
3. In the **Stream group details** page, choose **Edit configuration**.
4. Uncheck the checkbox next to the name of the location that you want to remove.
5. Choose **Save**.

CLI

Prerequisite

You must configure the AWS CLI with your user credentials and your chosen AWS Region. For setup instructions, refer to Download the AWS CLI.

To remove locations from a stream group using the AWS CLI

In your AWS CLI use the RemoveStreamGroupLocations command, customized for your content.

```
aws gameliftstreams remove-stream-group-locations \
    --identifier arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/
sg-1AB2C3De4
    --locations us-east-1 eu-central-1
```

where

`identifier`

An Amazon Resource Name (ARN) or ID that uniquely identifies the stream group resource.

ARN example: `arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/sg-1AB2C3De4`

ID example: `sg-1AB2C3De4`

`locations`

A set of locations to remove from this stream group. For a complete list of locations that Amazon GameLift Streams supports, refer to AWS Regions and streaming locations supported by Amazon GameLift Streams.

Delete a stream group

You can delete a stream group that's in any status. This action permanently deletes the stream group and releases its compute resources. If there are streams in process, then this action stops them and your end users can no longer view the stream.

As a best practice, before you delete a stream group, check for streams in process and take steps to stop them.

Console

To delete a stream group using the Amazon GameLift Streams console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console.
2. To view a list of your existing stream groups, in the navigation pane, choose **Stream groups**.
3. Choose the name of the stream group that you want to delete.

4. On the stream group detail page, choose **Delete**.
5. In the **Delete** dialog box, confirm the delete action.

CLI

Prerequisite

You must configure the AWS CLI with your user credentials and your chosen AWS Region. For setup instructions, refer to Download the AWS CLI.

To delete your stream group using the AWS CLI

In your AWS CLI use the DeleteStreamGroup command, customized for your content.

```
aws gameliftstreams delete-stream-group \
    --identifier arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/
sg-1AB2C3De4
```

where

`identifier`

An Amazon Resource Name (ARN) or ID that uniquely identifies the stream group resource.

ARN example: `arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/sg-1AB2C3De4`

ID example: `sg-1AB2C3De4`

Amazon GameLift Streams begins releasing compute resources and deleting the stream group. During this time, the stream group is in **Deleting** status. After Amazon GameLift Streams deletes the stream group, you can no longer retrieve it.

Linked applications

If you want to stream multiple application using the same pool of compute resources, then you can link multiple applications to the same stream group. Similarly, if you want to stream an application using different sets of compute resources, then you can link an application to multiple stream groups.

For more information about linking applications to stream groups, refer to Overview of multi-application stream groups.

Stream group lifecycle

Stream groups have a maximum lifespan of 365 days. As a best practice, we recommend that you recreate stream groups every 3-4 weeks to receive important service updates and fixes and ensure optimal performance. Recreating a stream group does not affect your uploaded applications.

As stream groups age, the following restrictions apply:

- **At 180 days**: You can no longer update the stream group with new application associations
- **At 365 days**: The stream group expires and can no longer stream sessions

The account associated with the stream group will receive two reminder notifications from AWS Health: one on the 45-day and a second reminder on the 150-day. These notifications will remind you that application association functionality will be lost on the 180-day. There will also be one final notification on 335-day reminding you that stream groups will expire on 365-day. Maintenance warnings also appear on the AWS Health dashboard and on stream group pages in the Amazon GameLift Streams console.

To find the expiration date of a stream group, view the **Stream group details** page on the console, or use the `ExpiresAt` field in the GetStreamGroup API response.

An expired stream group has a status of `EXPIRED` and becomes read-only. You cannot update it or start new stream sessions. To regain functionality, recreate the stream group.

Stream group maintenance

Whenever a feature is released that requires a new stream group to use it, you will see a "Maintenance required" message at the top of the stream group's detail page to inform you that it is outdated. Recreating a stream group is a manual process, but to help you do it, use the **Create Stream Group** button in the message to start the process. Some of the fields will be filled in for you.

Stream group maintenance is also required when the stream group is over 180 days old. You will no longer be able to link new applications to these older stream groups until they are recreated. At 365 days, streaming from the stream group will not be possible, and no changes to the stream group will be permitted.

Overview of multi-application stream groups

A *multi-application stream group* is a stream group that's linked to multiple applications. This enables you to stream multiple applications by using the same set of compute resources in a single stream group.

A common use case for multi-application stream groups is to release different versions of your game. For example, suppose that you created a stream group and set the default application to the original version of your game. Then, suppose you create additional applications that contains other versions of your game and link them to the stream group. Since these applications are associated with the same stream group, you only have to manage a single set of compute resources, or stream capacity, to stream all of these games. This means, regardless of which application an end-user streams, the application runs on a compute resource from the same set that this stream group has allocated.

Here are other possible real-life examples:

- A game streaming platform that offers different streaming tiers to customers.
- A quality assurance team that's testing multiple versions of a game.
- To simplify stream capacity management by using a single stream group for multiple applications.
- To enable a set of applications to stream from the same pool of stream capacity.

Limitations and requirements

You can only associate applications to stream groups that have compatible runtime environments and stream classes. For more information, refer to Runtime environment and stream class compatibility.

The following association limits apply to applications and stream groups. These limits are fixed within the service for all customers.

Name	Default	Adjustable	Description
Applications in a stream group	250	No	The maximum number of Amazon GameLift Streams applications that can be

Name	Default	Adjustable	Description
			associated to a stream group.
Stream group associations per application	100	No	The maximum number of stream groups that an Amazon GameLift Streams application can be associated to.

About default applications

Each stream group has one *default application*, which is initially the first application that you add to the stream group. The default application is automatically pre-cached on all always-on compute resources, which can help reduce application load time during stream startup. The Amazon GameLift Streams service can also cache other linked applications during its optimization processes.

Characteristics of default applications and other linked applications:

- The default application is pre-cached (on pre-allocated compute resources such as your always-on capacity) to help reduce application load time during stream startup.
- The default application can be changed. Note that when you switch default applications in a stream group, it can take up to a few hours for the new default application to be pre-cached in all locations.
- At least one linked application is required before you can start streaming from the stream group. The first application linked is automatically made the default application.
- If you unlink the default application of a stream group, Amazon GameLift Streams will automatically choose a new default application from the remaining associated applications, if there are any.
- The same application can be the default application for multiple stream groups.
- The set of linked applications is mutable until the stream group is 180 days old. In practical terms, this means that you can link and unlink applications until the stream group is 180 days old. After that, you will only be able to unlink applications from a stream group throughout the remainder of the stream group's lifecycle.

Change default application

When you link the first application to a stream group, it automatically becomes the default application and receives pre- caching benefits. You can change the default application at any time to give these benefits to a different application.

> **Note**
>
> When you switch default applications in a stream group, it can take up to a few hours for the new default application to be pre-cached in all locations.

Console

To change the default application using the Amazon GameLift Streams console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console.
2. In the navigation bar, choose **Stream groups** to view a list of your existing stream groups.
3. Select a stream group to view its details.
4. In **Linked applications**, select the application that you want to make the default.
5. Choose **Make default**.

CLI

Prerequisite

You must configure the AWS CLI with your user credentials and your chosen AWS Region. For setup instructions, refer to Download the AWS CLI.

To change the default application using the AWS CLI

In your AWS CLI use the UpdateStreamGroup command, customized for your content. The application that you want to make the default must already be associated to the stream group.

```
aws gameliftstreams update-stream-group \
    --identifier arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/
sg-1AB2C3De4 \
    --default-application-identifier a-9ZY8X7Wv6
```

where

- `identifier`:

  A stream group that has an application that you want to make the default.

  This value can be an Amazon Resource Name (ARN) or ID that uniquely identifies the stream group resource.

  ARN example: `arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/sg-1AB2C3De4`

  ID example: `sg-1AB2C3De4`

- `default-application-identifier`:

  The application that you want to make the default in this stream group.

  This value is an Amazon Resource Name (ARN) or ID that uniquely identifies the application resource.

  ARN example: `arn:aws:gameliftstreams:us-west-2:111122223333:application/a-9ZY8X7Wv6`

  ID example: `a-9ZY8X7Wv6`

Link an application

When you link, or associate, an application to a stream group, the stream group will be able to stream the application. If it is the first application in the stream group, it will automatically become the *default application*. You can link and unlink additional applications to a stream group until it reaches 180 days old. After that, you will only be able to unlink applications from a stream group throughout the remainder of the group's lifecycle.

> **Important**
>
> You cannot link an application to a stream group that is over 180 days old. To associate different applications to the stream group, you will first need to recreate it. For instructions on how to recreate a stream group, refer to Stream group maintenance.

Before you link an application, ensure that the stream group is in **Active** status.

Console

To link using the Amazon GameLift Streams console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console.
2. In the navigation bar, choose **Stream groups** to view a list of your existing stream groups.
3. Select a stream group to view its details.
4. In **Linked applications**, choose **Link application**.
5. Select an application that you want to link. Confirm your selection and choose **Link application**.

CLI

Prerequisite

You must configure the AWS CLI with your user credentials and your chosen AWS Region. For setup instructions, refer to Download the AWS CLI.

To link an application(s) using the AWS CLI

In your AWS CLI use the AssociateApplications command, customized for your content.

```
aws gameliftstreams associate-applications \
    --identifier arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/
sg-1AB2C3De4 \
    --application-identifiers a-9ZY8X7Wv6 a-1Z78C7Wv6
```

where

- `identifier`:

  A stream group to link these applications with.

  This value can be an Amazon Resource Name (ARN) or ID that uniquely identifies the stream group resource.

  ARN example: `arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/sg-1AB2C3De4`

ID example: `sg-1AB2C3De4`

- `application-identifiers`:

  A set of applications that you want to link with this stream group.

  This value is an Amazon Resource Name (ARN) or ID that uniquely identifies the application resource.

  ARN example: `arn:aws:gameliftstreams:us-west-2:111122223333:application/a-9ZY8X7Wv6`

  ID example: `a-9ZY8X7Wv6`

Unlink an application

When you unlink, or disassociate, an application from a stream group, you can no longer stream this application by using that stream group's allocated compute resources. Any streams in process will continue until they terminate, which helps avoid interrupting an end-user's stream. Amazon GameLift Streams will not initiate new streams using this stream group. The unlink action does not affect the stream capacity of a stream group.

If you unlink the default application of a stream group, Amazon GameLift Streams will automatically choose a new default application from the remaining associated applications, if there are any.

Console

To unlink using the Amazon GameLift Streams console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console.
2. In the navigation bar, choose **Stream groups** to view a list of your existing stream groups.
3. Select a stream group to view its details.
4. In **Linked applications**, select the application(s) that you want to unlink. Choose **Unlink applications**.
5. In the **Unlink applications** dialog, confirm the unlink action and choose **Unlink**.

CLI

Prerequisite

You must configure the AWS CLI with your user credentials and your chosen AWS Region. For setup instructions, refer to Download the AWS CLI.

To unlink an application(s) using the AWS CLI

In your AWS CLI use the DisassociateApplications command, customized for your content.

```
aws gameliftstreams disassociate-applications \
    --identifier arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/
sg-1AB2C3De4 \
    --application-identifiers a-9ZY8X7Wv6 a-1Z78C7Wv6
```

where

- `identifier`:

  A stream group to unlink these applications from.

  This value can be an Amazon Resource Name (ARN) or ID that uniquely identifies the stream group resource.

  ARN example: `arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/sg-1AB2C3De4`

  ID example: `sg-1AB2C3De4`
- `application-identifiers`:

  A set of applications that you want to unlink from this stream group.

  This value is an Amazon Resource Name (ARN) or ID that uniquely identifies the application resource.

  ARN example: `arn:aws:gameliftstreams:us-west-2:111122223333:application/a-9ZY8X7Wv6`

  ID example: `a-9ZY8X7Wv6`

Start stream sessions with Amazon GameLift Streams

This section covers stream sessions, the actual instance of a stream where an end user or player can interact with your application or play your game. You'll learn about how to test your own stream session and understand the stream session lifecycle.

For launching stream sessions to end users, you must integrate Amazon GameLift Streams into your own service. For more information, refer to Amazon GameLift Streams backend service and web client.

About stream sessions

The prerequisites to start a stream session are an application in **Ready** status, a stream group that has available capacity in the location where you want to stream, and the application replicated to the location where you want to stream. A stream session runs on one of the compute resources that a stream group has allocated. When you start a stream, you must specify a stream group and an application to stream using their ARN or ID values.

When you successfully start a stream session, you receive a unique identifier for that stream session. Then, you use that ID to connect the stream session to an end user. For more information, refer to StartStreamSession in the *Amazon GameLift Streams API Reference*.

Testing a stream in the console

The most direct way for you to test how your application streams is through the Amazon GameLift Streams console. When you start a stream, Amazon GameLift Streams uses one of the compute resources that your stream group allocates. So, you must have available capacity in your stream group.

To test your stream in the Amazon GameLift Streams console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console.
2. You can test a stream in several ways. Start from the **Stream groups** page or **Test stream** page and follow these steps:

    a. Select a stream group that you want to use to stream.

    b. If you're starting from the **Stream groups** page, choose **Test stream**. If you're starting from the **Test stream** page, select **Choose**. This opens the **Test stream** configuration page for the selected stream group.

c. In **Linked applications**, select an application.

d. In **Location**, choose a location with available capacity.

e. (Optional) In **Program configurations**, enter command-line arguments or environment variables to pass to the application as it launches.

f. Confirm your selection, and choose **Test stream**.

3. After your stream loads, you can do the following actions in your stream:

a. To connect input, such as your mouse, keyboard, and gamepad (except microphones, which are not supported in **Test stream**), choose **Attach input**. You automatically attach your mouse when you move the cursor into the stream window.

b. To have files that were created during the streaming session exported to an Amazon S3 bucket at the end of the session, choose **Export files** and specify the bucket details. Exported files can be found on the **Sessions** page.

c. To view the stream in fullscreen, choose **Fullscreen**. Press **Escape** to reverse this action.

4. To end the stream, choose **Terminate session**. When the stream disconnects, the stream capacity becomes available to start another stream.

> **Note**
>
> The **Test stream** feature in the Amazon GameLift Streams console does not support microphones.

Stream session lifecycle

When working with stream sessions in Amazon GameLift Streams, this diagram can help you understand the different states that a stream session transitions to throughout its lifecycle.

- StartStreamSession creates a new stream session, which begins in `ACTIVATING` state. When Amazon GameLift Streams finds available resources to host the stream, the stream session transitions to `ACTIVE`. When a client connects to the active stream, the stream session transitions to `CONNECTED`.
- When a client disconnects from a stream, the stream session transitions to `PENDING_CLIENT_RECONNECTION` state. CreateStreamSessionConnection transitions the stream session to `RECONNECTING`, and will either initiate the client to reconnect to the stream

or create a new stream session. When a stream session is ready for the client to reconnect, it transitions to `ACTIVE`. When the client reconnects, it transitions back to `CONNECTED`. If a client is disconnected for longer than `ConnectionTimeoutSeconds`, the stream session ends.

- When a client doesn't connect to a stream session in `ACTIVE` or `PENDING_CLIENT_RECONNECTION` state within the period of time specified by `ConnectionTimeoutSeconds`, then it transitions to `TERMINATED`.
- TerminateStreamSession initiates termination of the stream, and the stream session transitions to `TERMINATING` state. When the stream session terminates successfully, it transitions to `TERMINATED`.
- A stream session in any state, except `TERMINATED`, can transition to `ERROR`. When an API call returns `ERROR` as a Status value, check the value of StatusReason for a short description of the cause of the error. You can also call GetStreamSession to check these values.

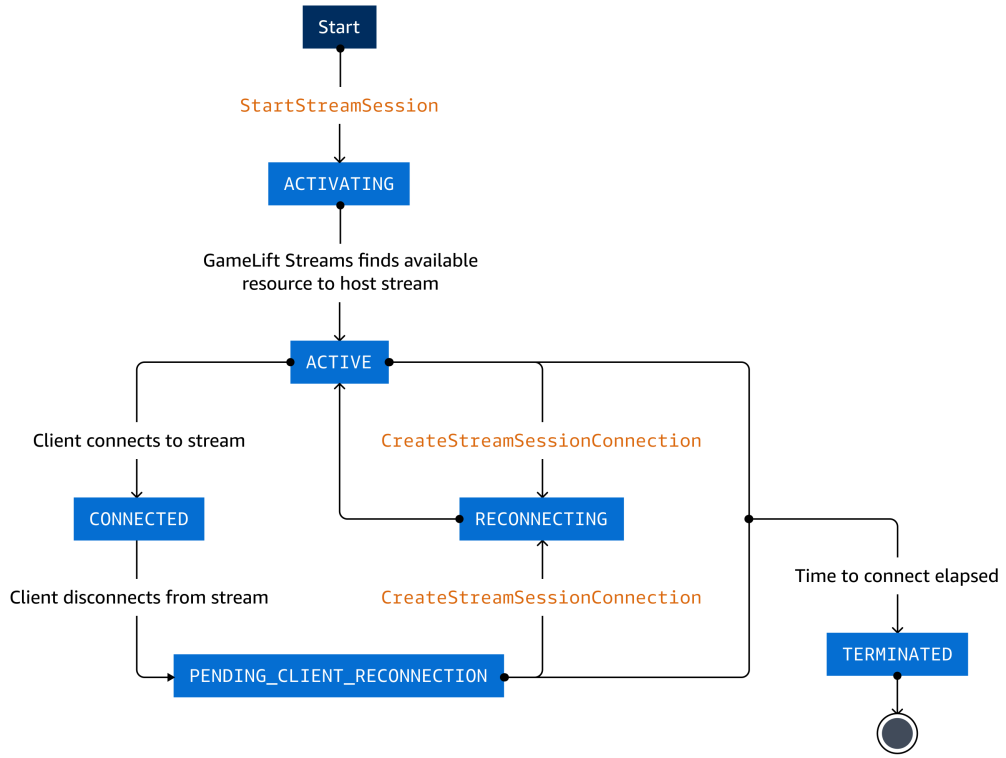
Start
StartStreamSession
ACTIVATING
GameLift Streams finds available resource to host stream
ACTIVE
Client connects to stream
CreateStreamSessionConnection
CONNECTED
RECONNECTING
Client disconnects from stream
CreateStreamSessionConnection
Time to connect elapsed
PENDING_CLIENT_RECONNECTION
TERMINATED

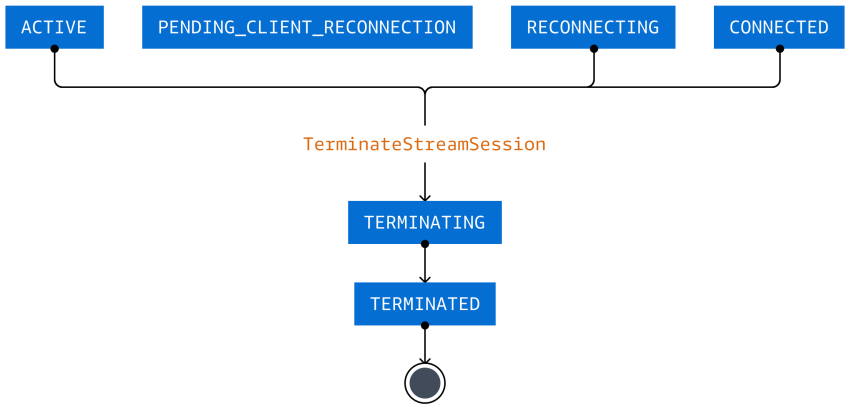
ACTIVE
PENDING_CLIENT_RECONNECTION
RECONNECTING
CONNECTED
TerminateStreamSession
TERMINATING
TERMINATED

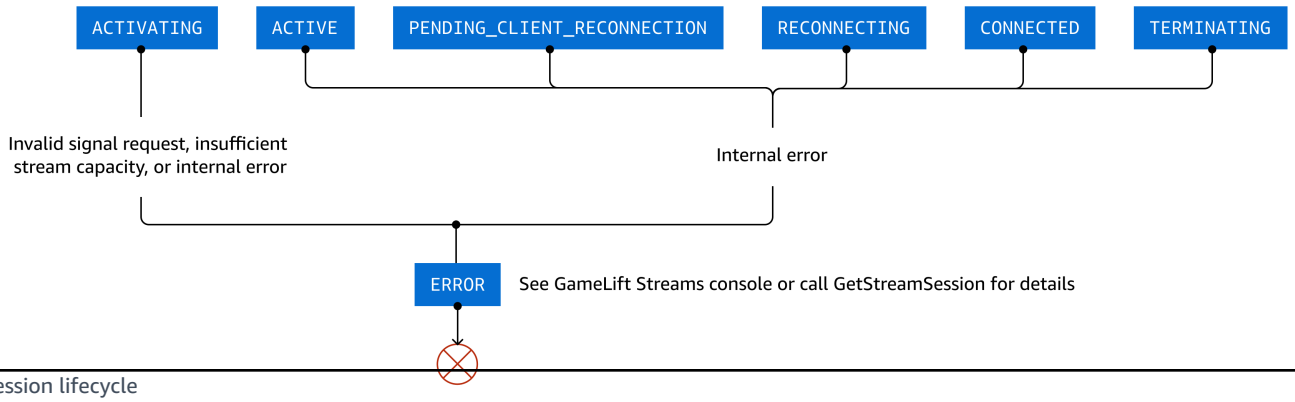
ACTIVATING
ACTIVE
PENDING_CLIENT_RECONNECTION
RECONNECTING
CONNECTED
TERMINATING
Invalid signal request, insufficient stream capacity, or internal error
Internal error
ERROR
See GameLift Streams console or call GetStreamSession for details

Timeout values affecting stream sessions

Stream sessions are governed by several timeout values that control different aspects of the session lifecycle. In roughly chronological order of when you might typically encounter them during the stream session lifecycle, they include the following:

Placement timeout

Time limit for Amazon GameLift Streams to find compute resources to host a stream session using available capacity. Placement timeout varies based on the capacity type used to fulfill your stream request:

- Always-on capacity: 75 seconds
- On-demand capacity:
  - Linux/Proton runtimes: 90 seconds
  - Windows runtime: 10 minutes
- Behavior: If Amazon GameLift Streams cannot identify available resources within this time, the stream session `Status` changes to `ERROR` with a `StatusReason` of `placementTimeout`.

Connection timeout

Length of time Amazon GameLift Streams waits for a client to connect or reconnect to a stream session.

- Parameter: `ConnectionTimeoutSeconds` in StartStreamSession
- Range: 1 - 3600 seconds (1 hour)
- Default: 120 seconds (2 minutes)
- Behavior: Timer starts when the stream session reaches `ACTIVE` or `PENDING_CLIENT_RECONNECTION` status. If no client connects before the timeout, the session `Status` transitions to `TERMINATED`.

Idle timeout

Length of time that a stream session can be idle with no user input.

- Parameter: (internal)
- Value: 60 minutes
- Behavior: If no user input is received during this time, Amazon GameLift Streams terminates the stream session.

Session length timeout

Maximum duration Amazon GameLift Streams keeps a stream session open.

- Parameter: `SessionLengthSeconds` in StartStreamSession
- Range: 1 - 86400 seconds (24 hours)
- Default: 43200 seconds (12 hours)
- Behavior: Terminates the stream session regardless of any existing client connection when the time limit is reached.

Terminating a stream session

If you need to force a stream session to terminate, you have the following options:

- **Use the TerminateStreamSession API:** To use TerminateStreamSession, you will need the stream group ID and the stream session ID. You can use ListStreamSessions or ListStreamSessionsByAccount with the `--status CONNECTED` parameter to get a list of stream sessions that have a client connected.
- **Remove the session's location from its stream group:** Removing the location from the stream group where the session is streaming will terminate all active stream sessions in that location. You can remove a location in a stream group from the console or by using the RemoveStreamGroupLocations API.
- **Delete the session's stream group:** Deleting a stream group will terminate all active stream sessions in all locations of the stream group. You can delete a stream group from the console or by using the DeleteStreamGroup API. Use with caution since you will be abruptly ending client connections.

Reconnecting to a stream session

If a client gets disconnected from a stream session without ending the session, it can reconnect to the session within the time specified by `ConnectionTimeoutSeconds` when the stream session was started. To reconnect to a session, you need the stream session's ID. For details, see CreateStreamSessionConnection in the *Amazon GameLift Streams API Reference*. You can see an example of reconnecting to a stream session in the React Starter Sample.

Export stream session files

During a stream session, your application can generate output files that help you debug or verify your application, such as logs, diagnostic information, crash dumps, save files, user data, and screenshots. The export stream session files feature collects files that are created or modified during a session and exports them as a compressed ZIP file to a provided Amazon S3 location. The feature also collects performance stats for the session every second, which are included in the export ZIP file.

Warning

Before you export files, be aware of the following things:

- Files may contain sensitive information written by your application, including credentials information.
- File sizes may be large depending on your application size, which impacts your Amazon S3 storage cost.
- If you select an Amazon S3 bucket in an AWS Region that differs from the Region of the stream group, then the exported stream session files will move across regions.

How it works

You must manually invoke this operation on an active stream session to export the files generated during that session. The stream session must be active, specifically in one of the following statuses `ACTIVE`, `CONNECTED`, `PENDING_CLIENT_RECONNECTION`, and `RECONNECTING`. At the end of the session, Amazon GameLift Streams exports the files to your bucket in Amazon Simple Storage Service (Amazon S3). Thus, all exported data is within your ownership and is subject to the Amazon S3 bucket's permissions policy.

Here's a walkthrough of the stream session lifecycle with export files activated:

1. Amazon GameLift Streams begins a session by connecting the user to your application that's running on the compute resource.
2. While your application streams, it creates or modifies files in the filesystem of the runtime environment.

3. When the session ends, Amazon GameLift Streams gets a copy of all the new or modified files in the filesystem and exports the files to your Amazon S3 bucket.

Amazon GameLift Streams collects the following generated and modified files. Find them in the corresponding folders in the `.zip` archive.

- `application/`: The folder where your application or game is stored.
- `profile/`: The user's profile folder contains the user's personal settings, configurations, and data.
- `temp/`: The system's temp folder contains temporary files and data that your application and the system creates. This can include cache files, log files, or intermediate processing data.
- `stats/`: Contains `perf_stats_v1.csv`, which holds performance stats for the session collected per-second. This includes application-level stats (CPU and memory utilization) and system-level stats (CPU, memory, GPU, and VRAM utilization). For a detailed description of each stat included in the CSV file, see the section called “Performance stats reference”

To delete the files, delete the object in the Amazon S3 bucket.

Cost impact

You incur a cost for having the files stored in Amazon S3. A stream session might generate a large amount of data depending on your application. Be aware that with many stream sessions that have this feature enabled, the cost can add up.

For more information, refer to Amazon S3 pricing.

Export files (Console)

To enable export stream session files in the Amazon GameLift Streams console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console.
2. In the navigation bar, choose **Sessions** to view a list of active and recent stream sessions within the last 90 days.
3. In the **Active sessions** tab, select an active stream session.
4. Choose **Export files** to enable the export files feature for that stream session.

5. In the **Export stream sessions file** dialog box, choose either **Create a new S3 bucket** or **Select an existing S3 bucket**. Follow the steps in the console to create or select an S3 object to store the exported data into.

   > **Warning**
   >
   > If the ZIP file name matches an existing one in the directory, the previous one will be overwritten.

6. Choose **Confirm**. You can now find the session listed in the **Exported files** tab.
7. Wait for the session to end and for the files to export.

Amazon GameLift Streams will export the files when the session is in **Terminated** state. When a session has terminated, it will move from the **Active sessions** tab to the **Recent sessions** tab.

You can check the status of the export process in the **Session exports** tab. If the status is **Pending**, the stream session is still active, so Amazon GameLift Streams hasn't exported the files yet. If the status is **Succeeded**, you can download the files from Amazon S3 using the link provided. If the status is **Failed**, hover over the status to see the reason for the failure.

Export files (CLI)

Prerequisite

You must configure the AWS CLI with your user credentials and your chosen AWS Region. For setup instructions, refer to Download the AWS CLI.

To export stream session files in the AWS CLI

In your AWS CLI use the ExportStreamSessionFiles command, customized for your content.

```
aws gameliftstreams export-stream-session-files \
    --identifier arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/
sg-1AB2C3De4 \
    --stream-session-identifier arn:aws:gameliftstreams:us-
west-2:111122223333:streamsession/sg-1AB2C3De4/ABC123def4567
    --ouput-uri s3://amzn-s3-demo-bucket/prefix
```

Where

`identifier`

An Amazon Resource Name (ARN) or ID that uniquely identifies the stream group resource.

ARN example: `arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/sg-1AB2C3De4`

ID example: `sg-1AB2C3De4`

`stream-session-identifier`

An Amazon Resource Name (ARN) or ID that uniquely identifies the stream session resource.

ARN example: `arn:aws:gameliftstreams:us-west-2:111122223333:streamsession/sg-1AB2C3De4/ABC123def4567`

ID example: `ABC123def4567`

`output-uri`

The Amazon S3 bucket URI where Amazon GameLift Streams uploads the set of compressed exported files for this stream session.

There are two valid formats that you can provide. If the URI has a `.zip` or `.ZIP` file extension, then Amazon GameLift Streams stores the exported files at the provided URI. Otherwise, Amazon GameLift Streams generates the name for a compressed folder and stores it at the URI. The generated name follows the pattern: `date-time-applicationId-streamGroupId-streamSessionId`. For example:

- If you provide a URI called `s3://amzn-s3-demo-bucket/MyGame_Session1.zip`, then Amazon GameLift Streams saves the files in that exact ZIP folder.
- If you provide a URI called `s3://amzn-s3-demo-bucket/MyGame_Session1/`, then Amazon GameLift Streams will save the files at `s3://amzn-s3-demo-bucket/MyGame_Session1/YYYYMMDD-HHMMSS-applicationId-streamGroupId-sessionId.zip`.

Be sure that your ZIP file name complies with the Object key naming guidelines in the *Amazon Simple Storage Service User Guide*.

> **Warning**
>
> If the ZIP file name matches an existing one in the directory, the previous one will be overwritten.

You can check on the status of the active session by calling the GetStreamSession API. From the stream session summary, you can get details about the status of exported files. If the status is **Pending**, then the stream session is still active, so Amazon GameLift Streams hasn't exported the files yet. If the status is **Succeeded**, navigate to the output URI to see the files in Amazon S3. If the status is **Failed**, check the `StatusReason` in the `ExportFilesMetaData`.

Amazon GameLift Streams backend service and web client

Amazon GameLift Streams enables you to stream applications through a web browser. With the Amazon GameLift Streams Web SDK, you can set up a backend streaming service. Then, end users connect to a stream through a web client. They can play your game or interact with your application all through the cloud.

The Amazon GameLift Streams Web SDK includes a sample backend server and a sample web client, which you can use to get started on creating a backend service. You can also use these samples to test how Amazon GameLift Streams streams, without additional development. To get started, refer to Setting up a web server and client with Amazon GameLift Streams.

Topics

- Supported browsers and input
- Required ports
- Setting up a web server and client with Amazon GameLift Streams
- Customize stream appearance
- Locale preference
- Mouse movement handling
- Data channel communication between an application and web client

Supported browsers and input

The following lists the supported platforms and browsers for viewing Amazon GameLift Streams streams and their compatible input peripherals. Browsers must also be compatible with advanced video coding (AVC), also known as H.264.

Overall, we recommend Google Chrome, Microsoft Edge, or a custom Chromium-based desktop application for the best end-user experience and maximum compatibility, particularly with game controllers.

To learn more about which controllers are compatible with which browsers, see the Web Gamepad API. Although some guidance may not apply to Amazon GameLift Streams, we expect most game controllers to connect successfully via Bluetooth.

Operating system	Browser	Input
Windows	Chrome, Edge	Keyboard, mouse, microphone, game controller (including haptic feedback)
	Firefox	Keyboard, mouse, microphone, game controller
Mac	Chrome, Edge, Safari	Keyboard, mouse, microphone, game controller (in Bluetooth mode) (including haptic feedback)
	Firefox	Keyboard, mouse, microphone
Linux	Chrome, Edge, Firefox	Keyboard, mouse
Android	Chrome, Edge	Simple touch-to-mouse emulation , microphone, external physical mouse, keyboard and game controller (in Bluetooth mode)
iOS	Chrome, Edge, Firefox, Safari	Simple touch-to-mouse emulation , microphone, external physical mouse, keyboard and game controller (in Bluetooth mode)

Known issues

Following are known issues with browsers and input:

- Safari will immediately exit fullscreen whenever `Esc` is pressed. This cannot be overriden.
- “Embedded” or “in-app” browser views like those inside mobile apps such as LinkedIn, Yelp, Instagram, and others are not supported on iOS. These tend to disable the browser WebRTC support necessary for realtime interactive streaming. We recommend detecting non-standard browser strings and prompting the user to open in Safari.
- If the screen resolution in your application is not set to 1080p, mouse tracking might be impacted. We recommend disabling the selection of any other resolution, if possible. We also recommend disabling windowed mode, and only run in full screen.
- To support plug and play of game controllers on Proton, despite the lack of support for them in native Linux applications, games running in Proton runtime environments will *always* show a game controller connected, even if none are plugged in on the client. This could be an issue for games that prompt for controller input even when the controller is idle and unused. We recommend that games show input UI based on the last input method.

Limitations

- Most runtime environments support game controllers, except for Ubuntu 22.04 LTS. If you need game controller support, consider creating the game using another runtime environment. For a list of other runtime environments, refer to Runtime environments.
- The PlayStation 5 and Luna game controllers are not supported in Firefox.
- Haptic feedback support:
  - Haptic feedback on the PlayStation 4 and Xbox Series S/X controllers are supported in Chrome, Edge, and Safari.
  - Haptics on the PlayStation 5 DualSense controller is only supported in the Safari browser.
  - Firefox does not support haptic feedback on any controller.
  - Android and iOS devices do not support haptic feedback on any controller.
- The **Test stream** feature in the Amazon GameLift Streams console does not support microphones.

IPv6 support

Streaming to IPv6-only clients is supported only with Windows runtime applications.

Runtime	Streaming over IPv4	Streaming over IPv6
Microsoft Windows Server 2022 Base	Yes	Yes
Ubuntu 22.04 LTS	Yes	No
Proton runtimes	Yes	No

Required ports

To integrate Amazon GameLift Streams, ensure that your network infrastructure has the necessary ports open and accessible. The following is a list of the ports you should plan to have open on your network to communicate with Amazon GameLift Streams.

Port	Protocol	Purpose
443	(HTTPS) TCP	AWS APIs, including Amazon GameLift Streams
33435-33465	UDP	Web RTC

Setting up a web server and client with Amazon GameLift Streams

In this tutorial, you will set up a web client application that integrates Amazon GameLift Streams' streaming service. Then, you will use the Amazon GameLift Streams Web SDK, a JavaScript library, and sample code that you can start with. The sample code includes a simple Amazon GameLift Streams backend web server and a simple web client. By the end of this tutorial, you can start a stream by using the sample code.

If it's your first time using Amazon GameLift Streams, we highly recommend starting with the Starting your first stream in Amazon GameLift Streams tutorial, which walks you through uploading a game to Amazon S3 and testing streaming it from within the Amazon GameLift Streams console in your browser.

Prerequisites

- An AWS account with proper credentials for programmatic access. For more information, see Setting up Amazon GameLift Streams as a developer.
- The AWS SDK.
- An Amazon GameLift Streams-supported web browser — see Supported browsers and input.
- Node.js — see Node.js downloads page.

Download the Web SDK

For this tutorial, you will need to download the following materials from the Resources section of the Getting Started product page:

- **Amazon GameLift Streams Web SDK bundle**: This includes sample code for a simple backend service and web client.
- **Amazon GameLift Streams Web SDK API Reference**: This API reference documents Amazon GameLift Streams API wrappers for JavaScript.

Set up your streaming resources

You must have stream resources—an application and a stream group—to start a stream. Specifically, you must have:

- An application in **Ready** status.
- A stream group in **Active** status with available stream capacity.
- For streaming in locations other than the primary location, the application must have finished replicating to that location.

To set up an application and a stream group using either the Amazon GameLift Streams console or Amazon GameLift Streams CLI, refer to Prepare an application in Amazon GameLift Streams and

Manage streaming with an Amazon GameLift Streams stream group, respectively. Alternatively, for an end-to-end walkthrough in the Amazon GameLift Streams console, refer to Starting your first stream in Amazon GameLift Streams.

Set up a backend server

The backend server is responsible for handling tasks such as authenticating users, configuring stream parameters, and performing Amazon GameLift Streams service API calls on behalf of end-users. Review the sample code and the Amazon GameLift Streams Web SDK API Reference to learn more about setting this up. Specifically, see the server.js file in the Amazon GameLift Streams Web SDK package.

> **Important**
>
> This code is example code for testing and evaluation purposes only and should not be used in a production capacity.

To run the sample backend service

1. Open a terminal or command prompt and navigate to the folder `AmazonGameLiftStreamsWebSDK\GameLiftStreamsSampleGamePublisherService\`.
2. Run the following commands:

```
npm install
node server.js
```

With the sample backend service running, end-users can connect to a stream through the web client. Test the web client in the next step.

Launch a web client

The web client application is responsible for receiving and decoding Amazon GameLift Streams streams, streaming to end-users, and providing the web browser UI for end-users to engage with the application. Review the sample code and the Amazon GameLift Streams Web SDK API Reference to learn more about how to integrate the JavaScript Amazon GameLift Streams Web

SDK into your own web client application. Specifically, see `public/index.html` in the Amazon GameLift Streams Web SDK package. You can also look at the web page source when you launch a web client in your browser.

> **Note**
>
> The Windows runtime in Amazon GameLift Streams supports stream sessions over IPv4 or IPv6. However, Linux and Proton runtime environments only support streaming over IPv4.

To launch a web client application

1. Open a web browser and navigate to `http://localhost:port/`. The port number is set by the backend server; by default, this is HTTP port 8000.
2. Play the game or use the software.
   a. To attach input, such as your mouse, choose **Attach input**.
   b. To exit the game, choose the **Esc** key.
   c. To stop the server process, choose **Ctrl+C** key.

Clean up streaming resources

> **Warning**
>
> A stream group incurs costs when it has allocated streaming capacity, even if that capacity is unused. To avoid unnecessary costs, scale your stream groups to your required size. We suggest during development that you scale always-on capacity and target-idle capacity in your stream groups to zero when not in use. For more information, refer to Scale stream groups to zero capacity.

After you complete the tutorial and no longer need to stream your application, follow these steps to clean up your Amazon GameLift Streams resources.

Deleting a stream group

When you delete a stream group, Amazon GameLift Streams works to release all stream capacity.

To delete a stream group using the Amazon GameLift Streams console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console.
2. To view a list of your existing stream groups, in the navigation pane, choose **Stream groups**.
3. Choose the name of the stream group that you want to delete.
4. On the stream group detail page, choose **Delete**.
5. In the **Delete** dialog box, confirm the delete action.

Amazon GameLift Streams begins releasing compute resources and deleting the stream group. During this time, the stream group is in **Deleting** status. After Amazon GameLift Streams deletes the stream group, you can no longer retrieve it.

Deleting an application

You can only delete an application that meets the following conditions:

- The application is in the **Ready** or **Error** state.
- An application is not streaming in any ongoing stream session. You must wait until the client ends the stream session or call TerminateStreamSession in the Amazon GameLift Streams API to end the stream.

If the application is linked to any stream groups, you must unlink it from all associated stream groups before you can delete it. In the console, a dialog box will walk you through this process.

To delete an application using the Amazon GameLift Streams console

1. Sign in to the AWS Management Console and open the Amazon GameLift Streams console.
2. In the navigation bar, choose **Applications** to view a list of your existing applications. Choose the application you want to delete.
3. In the application detail page, choose **Delete**.
4. In the **Delete** dialog box, confirm the delete action.

Amazon GameLift Streams begins deleting the application. During this time, the application is in `Deleting` status. After Amazon GameLift Streams deletes the application, you can no longer retrieve it.

Customize stream appearance

Loading screen

When a customer opens a web browser to view a stream, the web client starts establishing a connection to the Amazon GameLift Streams stream session. While the stream session loads, you can display a custom background and logo to the customer's screen.

The Amazon GameLift Streams Web SDK sample client, in the `GameLiftStreamsSampleGamePublisherService/public/LoadingScreen/loadingscreen.js` file, demonstrates how you can implement an animated logo in your front-end web client. The default loading screen consists of 2 images: background and foreground. The foreground image is positioned in the middle and has a pulse animation. The animation plays only while the stream session is connecting.

To enable a loading screen

1. In the Amazon GameLift Streams Web SDK sample client, navigate to the `GameLiftStreamsSampleGamePublisherService/public/LoadingScreen/` folder.
2. Add your background and foreground images using the default names, `Background.png` and `LoadingLogo.png`. If you want to rename them or use a different image format, you must update the code in `GameLiftStreamsSampleGamePublisherService/public/loadingscreen.js`.
3. (Optional) In `GameLiftStreamsSampleGamePublisherService/public/loadingscreen.js`, update the JavaScript code to implement different animations.

Locale preference

In Amazon GameLift Streams, you can set the locale preference per stream. This is useful if your application retrieves location-specific information from the end user's operating system, such as time or currency.

Amazon GameLift Streams supports the following languages:

Value	Description
en_US	U.S. English (default)

Value	Description
ja_jp.UTF-8	Japanese

To change the locale setting

When you call StartStreamSession using the Amazon GameLift Streams API, add LANG=*<language>* to your AdditionalEnvironmentVariables. Since locale preference is unique per user, you set this at the stream-session level. If you don't set this, the stream uses U.S. English by default.

Example Example

```
aws gameliftstreams start-stream-session \
   --identifier arn:aws:gameliftstreams:us-west-2:123456789012:streamgroup/1AB2C3De4 \
   --protocol WebRTC \
   --signal-request "[webrtc-ice-offer json string]" \
   --user-id xnshijwh \
   --additional-environment-variables '{"LANG": "ja_JP.UTF-8"}'
```

Mouse movement handling

Mouse movement handling is critical for delivering responsive and intuitive user experiences in streamed applications. Amazon GameLift Streams automatically optimizes mouse input transmission based on your application's cursor behavior, ensuring that mouse movements feel natural whether the cursor is hidden or visible. Understanding how Amazon GameLift Streams processes mouse events helps you design applications that work seamlessly with the streaming service and provide the best possible user experience.

Mouse input modes

Amazon GameLift Streams uses two distinct modes for transmitting mouse events to your application, automatically selecting the appropriate mode based on cursor visibility:

Relative mode

In relative mode, mouse updates are transmitted as small, incremental differences from the previous position. This mode is ideal for applications that require precise, continuous mouse movement tracking, such as first-person shooter (FPS) games or interfaces that use 3D

orientation. Amazon GameLift Streams uses relative mode when the operating system cursor is hidden or fully transparent.

Absolute mode

In absolute mode, the mouse cursor position is transmitted as an exact screen coordinate. This mode works well for applications that rely on precise cursor positioning, such as point-and-click games or any UI with clickable elements. Amazon GameLift Streams uses absolute mode when the operating system cursor is visible, even if your application displays a custom cursor image.

This automatic selection ensures optimal performance for different application types without requiring manual configuration.

Pointer lock

Pointer lock is a web API feature that captures the mouse cursor within a specific element, hiding the cursor and preventing it from leaving the designated area. This feature is particularly valuable for games that require unrestricted mouse movement for camera control or aiming, without the distraction of a visible cursor or the limitation of reaching window edges.

Amazon GameLift Streams provides automatic pointer lock functionality through the `autoPointerLock` property in the Web SDK's `InputConfiguration` interface. This feature integrates with the requestPointerLock API to provide intuitive and context-aware mouse capture.

Automatic pointer lock behavior

Amazon GameLift Streams automatically enables pointer lock when the application is fullscreen and the remote cursor is invisible on the stream host. This behavior aligns well with common game development patterns:

- **FPS/TPS games and 3D orientation control** - The pointer is automatically locked and the cursor is hidden, providing unrestricted camera control essential for FPS gameplay.
- **Point-and-click games and UI control** - When games make the cursor visible for menu interactions or strategy gameplay, the pointer remains visible and unlocked, preserving the intended user experience.

Configuration options

The `autoPointerLock` property accepts the following values:

`true`

The mouse is always captured when the remote cursor is invisible.

`false`

The mouse is never captured, regardless of cursor visibility.

`'fullscreen'` (default)

The mouse is only captured when the video element is in fullscreen mode and the remote cursor is invisible.

> **Important**
>
> `autoPointerLock` has no effect in the Safari browser or on iOS platforms due to platform limitations.

Best practices

To ensure optimal mouse handling in your streamed applications:

- **Always stream fullscreen** - Your application should already be running in fullscreen mode to work properly on our service. In addition, we recommend using browser support to make the stream a fullscreen element for the best end-user experience. This will help avoid problems such as alignment issues between the system cursor and software cursor.
- **Hide the cursor for relative motion** - If your application expects relative mouse motion (such as FPS-style camera controls or drag-based interactions), hide the operating system cursor during those interactions. In some scenarios, you might need to hide the cursor on mouse-down and show it again on mouse-up.
- **Show the cursor for absolute positioning** - When your application needs precise cursor positioning for UI interactions, ensure the operating system cursor remains visible to enable absolute coordinate mode.
- **Test different input scenarios** - Verify that your application handles both relative and absolute mouse modes correctly, as Amazon GameLift Streams may switch between modes based on your cursor visibility changes.

- **Test different window modes** - Test your application's mouse handling in both windowed and fullscreen modes, if applicable. Determine which `autoPointerLock` setting is best for your input configuration.

Data channel communication between an application and web client

Data channels allow you to securely communicate arbitrary messages between your Amazon GameLift Streams application and the web client (the JavaScript code running in the end-user's web browser). This allows end-users to interact with the application that Amazon GameLift Streams is streaming, via the web browser where they're viewing the stream.

Here are some example use cases of data channels in Amazon GameLift Streams:

- Users can open URLs in the application in their local browser.
- Users can pass content in the clipboard back and forth to the application.
- Users can upload content from their local machine to the application.
- Developers can implement UI in the browser that sends commands to the application.
- Users can pass schemas to control display of visualization layers.

Features

Message size limits

Amazon GameLift Streams Web SDK imposes a maximum size limit of 64 KB (65536 bytes) per message. This ensures that the message size limits are compatible with most browsers, and that the communication has low impact on the total bandwidth of the stream.

Metrics

Metrics on your data channel usage are sent to your AWS account when a stream session ends. For more information, refer to Data channels in the *Monitoring Amazon GameLift Streams* section.

Using data channels

The Amazon GameLift Streams Web SDK provides the `sendApplicationMessage` function that sends a message as a byte array to the application. The message is processed by a callback function, `clientConnection.applicationMessage` that you define.

If the client sends messages before the application connects to the data channel port, the messages are queued. Then, when the application connects, it receives the messages. However, if the application sends messages before the client connects to the data channel port, the messages are lost. The application must check the connection state of the clients before sending a message.

On the client-side

Write the following code in your web client application.

1. Define the callback function to receive incoming messages from the application.

```
function streamApplicationMessageCallback(message) {
    console.log('Received ' + message.length + ' bytes of message from
    Application');
}
```

2. Set `clientConnection.applicationMessage` to your callback function.

```
clientConnection: {
    connectionState: streamConnectionStateCallback,
    channelError: streamChannelErrorCallback,
    serverDisconnect: streamServerDisconnectCallback,
    applicationMessage: streamApplicationMessageCallback,
}
```

3. Call the `GameLiftStreams.sendApplicationMessage` function to send messages to your application. You can call this any time, as long as the stream session is active and the input is attached.

As an example, refer to the Amazon GameLift Streams Web SDK sample client, which demonstrates how to set up a simple data channel on the client-side.

On the application-side

Write the following logic in your application.

Step 1. Connect to the data channel port

When your application starts, connect to port `40712` on `localhost`. Your application should maintain this connection for the entire duration of execution. If the application closes the connection, it cannot be reopened.

Step 2. Listen for events

An event starts with a fixed-size header, followed by variable-length associated data. When your application receives an event, parse the event to retrieve the information.

Event format

- **Header**: A 4-byte header in the form `abcc`
  - `a` : Client id byte. This identifies a specific client connection, in the case of multiple connections (due to disconnection and reconnection).
  - `b` : Event type byte. `0` - the client connected, `1` - the client disconnected, `2` - a message is sent from the client. Other event types may be received with future Amazon GameLift Streams service updates, and should be ignored.
  - `cc` : Length of the associated event data. This is represented as 2 bytes with big-endian ordering (first byte is the most significant). If the event type is `2`, the event data represents the message contents from the client.
- **Data**: The remaining bytes contain the event data, such as a client message. The length of the data is indicated by `cc` in the header.

To listen for events

1. Read the four header bytes to retrieve the client id, event type, and length of the event data.
2. Read the variable-length event data, regardless of the client id and event type, according to the length described in the header. It's important to read the data unconditionally so that event data is never left in the buffer, where it could be confused with the next event header. Do not make assumptions about the length of the data based on event types.
3. Take appropriate action based on the event type, if recognized by your application. This action might include logging an incoming connection or disconnection, or parsing the client message and triggering application logic.

Step 3. Transmit messages to the client

The application should transmit messages with the same four-byte header format used by incoming events.

To transmit a message to the client

1. Write the header with the following properties:

    a. a : Client id byte. If your message is in response to a client message, it should reuse the same client id as the incoming client message, to avoid race conditions such as delivering a response from an old client connection to a newly reconnected client. If your application is sending an unsolicited message to the client, it should set the client id to match the most recent "client connection" event (event type 0).

    b. b : The event type of outgoing messages must always be 2. The client ignores messages with other event types.

    c. cc : Length of the message, in bytes.

2. Write the message bytes.

The message delivers to the specified client, unless the client disconnects. When a disconnected client reconnects, a new client ID is assigned via a client connected event. Any undelivered messages for the old client ID are discarded.

Example

The following pseudo-code demonstrates the logic to communicate messages in the application-side. For a complete example using Winsock, refer to Complete Winsock Client Code in the Windows Sockets 2 documentation.

```
connection = connect_to_tcp_socket("localhost:40712")
loop:
    while has_pending_bytes(connection):
        client_id = read_unsigned_byte(connection)
        event_type = read_unsigned_byte(connection)
        event_length = 256 * read_unsigned_byte(connection)
        event_length = event_length + read_unsigned_byte(connection)
        event_data = read_raw_bytes(connection, event_length)
        if message_type == 0:
            app_process_client_connected(client_id)
        else if message_type == 1:
            app_process_client_disconnected(client_id)
        else if message_type == 2:
            app_process_client_message(client_id, event_data)
        else:
            log("ignoring unrecognized event type")
    while app_has_outgoing_messages():
        target_client_id, message_bytes = app_next_outgoing_message()
        message_length = length(message_bytes)
        write_unsigned_byte(connection, target_client_id)
        write_unsigned_byte(connection, 2)
        write_unsigned_byte(connection, message_length / 256)
        write_unsigned_byte(connection, message_length mod 256)
        write_raw_bytes(connection, message_bytes)
```

Amazon GameLift Streams launch checklist

Preparing for a successful launch on Amazon GameLift Streams involves planning and coordination. Follow this detailed checklist to ensure a smooth experience in the weeks leading up to your event.

Notify the Amazon GameLift Streams team

Action: At least 8 weeks in advance, inform your technical account manager, your account team, or your account solution architect about your launch timeline and expected peak concurrent streams.

Reason: Understanding the scale of your production workload helps us ensure that your service limits are adequate, and adjust them if necessary. We also provide guidance on capacity availability and recommendations for the launch.

Compatibility and performance testing

Action: Test your application at scale, and in all of the locations where you have capacity, to confirm a positive customer experience. Amazon GameLift Streams offers NVIDIA based stream classes with different levels of performance and runtimes supported.

Reason: Thorough testing helps identify and resolve any potential compatibility and performance issues before the launch. Keep in mind the following about stream classes:

- The "high" stream classes support multi-tenancy, allowing two applications to run concurrently on a single instance. If you're using the "high" stream class, test with at least 2 concurrent streams to see how your application performs with shared resources, such as the CPU, GPU, and memory.

Capacity reservation

Action: At least 8 weeks before launch, reach out to your account team to reserve capacity, especially if you anticipate a critical, large-scale need. Decide on the stream classes and streaming locations based on your compatibility testing, performance requirements, and budget. Provide the start/end times and the required capacity. AWS requires that all capacity reservations be finalized 6-8 weeks before the reservation need-by date.

Reason: Amazon GameLift Streams operates on a first-come, first-serve basis using on-demand capacity. Reservations are essential to guarantee the necessary capacity.

Performance testing at scale

Action: Conduct thorough load testing of your APIs and your Amazon GameLift Streams configurations to observe its performance under load (latency, resolution, and frame rate). Be sure to check the Amazon GameLift Streams API rate limits to ensure that you have sufficient headroom for your launch and beyond. If you believe you will need a limit increase, you should reach out to your account manager or submit a support ticket.

Reason: Load-testing reveals how your application and Amazon GameLift Streams configurations will perform under stress before the launch. This is crucial to ensure smooth performance at scale.

Pre-launch setup

Action: At least 2-3 days before launch, create your final application resources and stream groups. Validate streaming performance and scale up capacity as needed.

Reason: This ensures that all components are working as expected, minimizing the risk of unexpected issues and allowing for easier diagnosis and recovery during the event.

Additional tips

- **Consistency is key**: Using the same existing stream groups throughout a launch event maintains consistency in the Amazon GameLift Streams backend, simplifying troubleshooting.
- **Monitor closely**: Closely monitor performance and user feedback to quickly address any issues. Build an operational dashboard. Monitor stream capacity, usage, and performance using Amazon CloudWatch (see Monitor with CloudWatch for details). Refer to the Well-Architected Framework for additional guidance.

Need Further Assistance?

If you have any questions or require further support, don't hesitate to reach out to us at Amazon GameLift Streams support. We're here to help ensure your launch is successful and seamless.

Security in Amazon GameLift Streams

Cloud security at AWS is the highest priority. As an AWS 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 AWS and you. The shared responsibility model describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS 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 AWS Compliance Programs. To learn about the compliance programs that apply to Amazon GameLift Streams, see AWS Services in Scope by Compliance Program.
- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. Amazon GameLift Streams is designed to run programs that you provide, and that you are solely responsible for the content and security of those programs. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Amazon GameLift Streams. The following topics show you how to configure Amazon GameLift Streams to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your Amazon GameLift Streams resources.

Topics

- Data protection in Amazon GameLift Streams
- Identity and Access Management for Amazon GameLift Streams
- Compliance validation for Amazon GameLift Streams
- Resilience in Amazon GameLift Streams
- Infrastructure Security in Amazon GameLift Streams
- Configuration and vulnerability analysis in Amazon GameLift Streams
- Security best practices for Amazon GameLift Streams

Data protection in Amazon GameLift Streams

The AWS shared responsibility model applies to data protection in Amazon GameLift Streams. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS 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 AWS services that you use. For more information about data privacy, see the Data Privacy FAQ. For information about data protection in Europe, see the AWS Shared Responsibility Model and GDPR blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS 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 AWS resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail. For information about using CloudTrail trails to capture AWS activities, see Working with CloudTrail trails in the *AWS CloudTrail User Guide*.
- Use AWS encryption solutions, along with all default security controls within AWS 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 AWS 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 Streams or other AWS services using the console, API, AWS CLI, or AWS 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 Streams handles service-specific data as follows:

- **Customer-supplied applications** – Amazon GameLift Streams stores customer data, if provided, in internal service-managed Amazon S3 buckets and on NVME storage drives attached to Amazon EC2 instances. All data is stored with service-managed encryption at rest. There is no direct customer access to this copy of the data. To delete an application, use the Amazon GameLift Streams console or the service API.
- **Customer-supplied metadata** – Customers may provide metadata to Amazon GameLift Streams APIs including descriptions, connection information, and opaque identifiers such as customer IDs. This metadata is always associated with specific customer resources.
- **Customer-generated data** – If an application writes new data as part of its normal operation, this customer-generated data is retained until the end of the user session. At the end of the session, generated data can optionally be exported to an Amazon S3 bucket destination of the customer's choice. Customer-generated data otherwise does not leave the Amazon EC2 instance where it was generated. For more information about data handling, refer to the topics on Session isolation.
- **Metrics and event data** – Amazon GameLift Streams metric and event data, which can be accessed through the Amazon GameLift Streams console or by calls to the service API. Data is available on applications, stream groups, and stream sessions. Authorized users can also access this data through Amazon CloudWatch and CloudWatch Events.

> **Important**
>
> If you provide customer IDs or other identifiers to Amazon GameLift Streams, it is expected that these values are anonymized references and do not contain any sensitive or personal information. Amazon GameLift Streams does not redact any metadata fields.

For more information about data protection, see the AWS shared responsibility model and GDPR blog post on the *AWS Security Blog*.

Encryption at rest

At-rest encryption of Amazon GameLift Streams-specific data is handled as follows:

- Application content is stored in service-managed encrypted Amazon S3 buckets and additionally on hardware-encrypted NVME drives attached to service-managed Amazon EC2 instances.

Encryption in transit

Calls to the Amazon GameLift Streams APIs are made over a secure (SSL) connection and authenticated using AWS Signature Version 4 (when connecting through the AWS CLI or AWS SDK, signing is handled automatically). Calling entities use security credentials, which are authenticated by applying the IAM access policies that are defined for Amazon GameLift Streams resources.

In the context of multi-location stream groups, in order to stream an application from any location in the stream group that has been allocated streaming capacity, Amazon GameLift Streams securely replicates applications to those locations.

Similarly, Amazon GameLift Streams will save log data and session files, when requested, to customer-named Amazon S3 buckets at the end of a session. If the bucket is not in the same location as the session, Amazon GameLift Streams will transfer the files securely to the AWS Region where the bucket is located.

Protection of end-user streams

Individual end-user streams are direct connections between the end-user's web browser and the Amazon GameLift Streams backend hosts. Those streams are protected with industry-standard WebRTC encryption, and both endpoints of the stream are positively identified by cryptographic identifiers which are part of the `SignalRequest` and `SignalResponse` values negotiated through the stream session APIs.

Data channel messages are also covered by the WebRTC encryption used for streams. These messages are decrypted by Amazon GameLift Streams and passed locally on-the-host to the customer's application through an unencrypted API. If end-to-end encryption which even the Amazon GameLift Streams service cannot decrypt is required, this additional layer of encryption is the responsibility of the application developer.

Session isolation in Linux stream classes

On Linux stream classes (Ubuntu and Proton runtimes), Amazon GameLift Streams uses *container isolation*. Every session runs in a new Linux container which is discarded after use. This means each new session runs in a fresh environment, isolated from other users sharing the compute resource (if running in a shared-resource stream class). No data from prior sessions exists when a new session starts up.

Session isolation in Windows stream classes

On Windows stream classes (Microsoft Windows Server runtimes), Amazon GameLift Streams uses *software isolation*. The service relies on a software agent to reset critical system state between sessions. Some folders are preserved across multiple sessions to allow for performance optimizations, such as on-host disk caching. The software agent automatically removes any files that were generated in the user's profile directory during the prior stream session. However, the agent does not remove any files that existed prior to the application running and were modified while the application was running. Nor does it remove any Windows registry keys that the application had added. Customers should be aware that it is their responsibility to avoid damaging the integrity of the overall operating system. Applications are executed as the Administrator user, which may permit modification to critical system-level files, including changes that persist across multiple sessions. It is the responsibility of the customer to secure their applications and guard against creating unsafe or unstable operating system modifications.

Customers are responsible for cleaning up those modified files and added registry keys from previous sessions when the application launches. This is an important step to protect confidential or sensitive information that the application writes to the user's profile directory. To do this, customers can write their own custom script that performs the following actions:

- Restore any files outside of the `%USERPROFILE%` directory that were modified by the application.
- Clean up any sensitive or user-specific registry keys that the application added.

Encryption key management

The service uses AWS-managed encryption keys. Each region uses a separate KMS key. Customer-managed keys (CMKs) are not supported.

Application files provided to Amazon GameLift Streams cannot be republished or exported from the service. The customer can use the service console or APIs to delete applications. Drives which previously held these application files can be completely purged by deleting the associated stream groups.

Inter-network traffic privacy

Amazon GameLift Streams uses public-facing networks to host stream sessions. Each stream group consists of one or more service-managed VPC networks which are isolated from other

stream groups and from other customers. Inbound network connections are denied except for authenticated, service-brokered WebRTC stream connections. Customer applications may connect out from these VPCs to other public addresses without restriction.

Additionally, there is no way for a customer to make a stream or their application data publicly-accessible using service API calls or settings alone. All service interactions are gated by AWS-authenticated API calls. If the customer wishes to make a stream accessible to the public they must create their own client web application which makes the authenticated calls to start and display a stream.

Identity and Access Management for Amazon GameLift Streams

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Amazon GameLift Streams resources. IAM is an AWS service that you can use with no additional charge.

Topics

- Audience
- Authenticating with identities
- Managing access using policies
- How Amazon GameLift Streams works with IAM
- Identity-based policy examples for Amazon GameLift Streams
- Troubleshooting Amazon GameLift Streams identity and access

Audience

How you use AWS 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 Streams identity and access)
- **Service administrator** - determine user access and submit permission requests (see How Amazon GameLift Streams works with IAM)
- **IAM administrator** - write policies to manage access (see Identity-based policy examples for Amazon GameLift Streams)

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be authenticated as the AWS account root user, an IAM user, or by assuming an IAM role.

You can sign in as a federated identity using credentials from an identity source like AWS IAM Identity Center (IAM Identity Center), single sign-on authentication, or Google/Facebook credentials. For more information about signing in, see How to sign in to your AWS account in the *AWS Sign-In User Guide*.

For programmatic access, AWS provides an SDK and CLI to cryptographically sign requests. For more information, see AWS Signature Version 4 for API requests in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity called the AWS account *root user* that has complete access to all AWS 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 AWS services using temporary credentials.

A *federated identity* is a user from your enterprise directory, web identity provider, or Directory Service that accesses AWS services using credentials from an identity source. Federated identities assume roles that provide temporary credentials.

For centralized access management, we recommend AWS IAM Identity Center. For more information, see What is IAM Identity Center? in the *AWS IAM Identity Center User Guide*.

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 AWS 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 AWS CLI or AWS 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 AWS by creating policies and attaching them to AWS identities or resources. A policy defines permissions when associated with an identity or resource. AWS evaluates these policies when a principal makes a request. Most policies are stored in AWS 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 AWS managed policies from IAM in a resource-based policy.

Other policy types

AWS 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 AWS Organizations. For more information, see Service control policies in the *AWS 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 *AWS 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 AWS determines whether to allow a request when multiple policy types are involved, see Policy evaluation logic in the *IAM User Guide*.

How Amazon GameLift Streams works with IAM

Before you use IAM to manage access to Amazon GameLift Streams, learn what IAM features are available to use with Amazon GameLift Streams.

IAM features you can use with Amazon GameLift Streams

IAM feature	Amazon GameLift Streams 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)	Partial. ABAC is only supported for applications and stream groups.
Temporary credentials	Yes
Principal permissions	Yes
Service roles	No
Service-linked roles	No

To get a high-level view of how Amazon GameLift Streams and other AWS services work with most IAM features, see AWS services that work with IAM in the *IAM User Guide*.

Identity-based policies for Amazon GameLift Streams

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 Streams

To view examples of Amazon GameLift Streams identity-based policies, see Identity-based policy examples for Amazon GameLift Streams.

Resource-based policies within Amazon GameLift Streams

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 AWS 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 Streams

Supports policy actions: Yes

Administrators can use AWS 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.

Policy actions in Amazon GameLift Streams use the following prefix before the action:

```
gameliftstreams
```

To specify multiple actions in a single statement, separate them with commas.

Example

```
"Action": [
            "gameliftstreams:action1",
            "gameliftstreams:action2"
               ]
```

To view examples of Amazon GameLift Streams identity-based policies, see Identity-based policy examples for Amazon GameLift Streams.

Policy resources for Amazon GameLift Streams

Supports policy resources: Yes

Administrators can use AWS 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": "*"
```

To view examples of Amazon GameLift Streams identity-based policies, see Identity-based policy examples for Amazon GameLift Streams.

Policy condition keys for Amazon GameLift Streams

Supports service-specific policy condition keys: Yes

Administrators can use AWS 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 AWS global condition keys, see AWS global condition context keys in the *IAM User Guide*.

To view examples of Amazon GameLift Streams identity-based policies, see Identity-based policy examples for Amazon GameLift Streams.

ACLs in Amazon GameLift Streams

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 Streams

Supports ABAC (tags in policies): Partial

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 AWS 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*.

Using temporary credentials with Amazon GameLift Streams

Supports temporary credentials: Yes

Temporary credentials provide short-term access to AWS resources and are automatically created when you use federation or switch roles. AWS recommends that you dynamically generate

temporary credentials instead of using long-term access keys. For more information, see Temporary security credentials in IAM and AWS services that work with IAM in the *IAM User Guide*.

Cross-service principal permissions for Amazon GameLift Streams

Supports forward access sessions (FAS): Yes

Forward access sessions (FAS) use the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. For policy details when making FAS requests, see Forward access sessions.

When creating new application resources, Amazon GameLift Streams uses the permissions of the calling principal to access the Amazon S3 bucket that contains the customer's application files. Amazon GameLift Streams also examines the calling principal to verify opt-in eligibility for certain cross-region functionality, such as multi-location stream groups.

Service roles for Amazon GameLift Streams

Supports service roles: No

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 AWS service in the *IAM User Guide*.

> **Warning**
>
> Changing the permissions for a service role might break Amazon GameLift Streams functionality. Edit service roles only when Amazon GameLift Streams provides guidance to do so.

Service-linked roles for Amazon GameLift Streams

Supports service-linked roles: No

A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS 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 AWS services that work with IAM. Find a service in the table that includes a Yes in the **Service-linked role** column. Choose the **Yes** link to view the service-linked role documentation for that service.

Identity-based policy examples for Amazon GameLift Streams

By default, users and roles don't have permission to create or modify Amazon GameLift Streams 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 Streams, including the format of the ARNs for each of the resource types, see Actions, resources, and condition keys for Amazon GameLift Streams in the *Service Authorization Reference*.

Topics

- Policy best practices
- Using the Amazon GameLift Streams console
- Allow users to view their own permissions

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Amazon GameLift Streams resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with AWS managed policies and move toward least-privilege permissions** – To get started granting permissions to your users and workloads, use the *AWS managed policies* that grant permissions for many common use cases. They are available in your AWS account. We recommend that you reduce permissions further by defining AWS customer managed policies that are specific to your use cases. For more information, see AWS managed policies or AWS 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 AWS service, such as 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 AWS 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 Streams console

To access the Amazon GameLift Streams console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Amazon GameLift Streams resources in your AWS 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.

You don't need to allow minimum console permissions for users that are making calls only to the AWS CLI or the AWS API. Instead, allow access to only the actions that match the API operation that they're trying to perform.

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 AWS CLI or AWS API.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ViewOwnUserInfo",
            "Effect": "Allow",
            "Action": [
                "iam:GetUserPolicy",
                "iam:ListGroupsForUser",
                "iam:ListAttachedUserPolicies",
                "iam:ListUserPolicies",
                "iam:GetUser"
            ],
            "Resource": ["arn:aws: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": "*"
        }
    ]
}
```

Troubleshooting Amazon GameLift Streams identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with Amazon GameLift Streams and IAM.

Topics

- I am not authorized to perform an action in Amazon GameLift Streams
- I want to allow people outside of my AWS account to access my Amazon GameLift Streams resources

I am not authorized to perform an action in Amazon GameLift Streams

If you receive an error that you're not authorized to perform an action, your policies must be updated to allow you to perform the action.

The following example error occurs when the `mateojackson` IAM user tries to use the console to view details about a fictional *`my-example-widget`* resource but doesn't have the fictional `gameliftstreams:`*`GetWidget`* permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform:
 gameliftstreams:GetWidget on resource: my-example-widget
```

In this case, the policy for the `mateojackson` user must be updated to allow access to the *`my-example-widget`* resource by using the `gameliftstreams:`*`GetWidget`* action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my AWS account to access my Amazon GameLift Streams resources

This is not possible with Amazon GameLift Streams. All API access is restricted to the account which owns the resources. Instead, customers who wish to share content externally are responsible for using their account to initiate new stream sessions on behalf of other users using Amazon GameLift Streams APIs, and forwarding the appropriate connection information to those external users' web browsers.

Compliance validation for Amazon GameLift Streams

To learn whether an AWS service is within the scope of specific compliance programs, see AWS services in Scope by Compliance Program and choose the compliance program that you are interested in. For general information, see AWS Compliance Programs.

You can download third-party audit reports using AWS Artifact. For more information, see Downloading Reports in AWS Artifact.

Your compliance responsibility when using AWS 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 AWS services, see AWS Security Documentation.

Resilience in Amazon GameLift Streams

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS 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 AWS Regions and Availability Zones, see AWS Global Infrastructure.

In addition to the data redundancy provided by the AWS global infrastructure, Amazon GameLift Streams is built with a resilient multi-Availability Zone infrastructure. In the case of an Availability Zone outage, individual existing sessions might be affected, but the service will continue to load-balance new sessions across healthy Availability Zones.

Infrastructure Security in Amazon GameLift Streams

As a managed service, Amazon GameLift Streams is protected by AWS global network security. For information about AWS security services and how AWS protects infrastructure, see AWS Cloud Security. To design your AWS environment using the best practices for infrastructure security, see Infrastructure Protection in *Security Pillar AWS Well-Architected Framework*.

You use AWS published API calls to access Amazon GameLift Streams through the network. Clients must support the following:

- Transport Layer Security (TLS). We require TLS 1.2 and recommend TLS 1.3.
- Cipher suites with perfect forward secrecy (PFS) such as DHE (Ephemeral Diffie-Hellman) or ECDHE (Elliptic Curve Ephemeral Diffie-Hellman). Most modern systems such as Java 7 and later support these modes.

Reuse and multi-tenancy in Amazon GameLift Streams

Amazon GameLift Streams doesn't share any compute resources across stream groups or with other AWS customers. Some Amazon GameLift Streams stream groups rely on internal resource sharing.

Reuse of compute resources

Within a stream group, resources are reused over time to serve multiple sessions with minimal downtime. The specific details of reuse are different between Windows and non-Windows stream groups.

Non-Windows stream groups with stream classes such as `gen4n_high`, `gen5n_ultra`, `gen6n_ultra` or `gen6n_pro` execute your applications inside of dedicated per-session containers. Each stream session begins with a copy of the application files and an empty user profile folder. When a session terminates, all file system modifications are discarded and all processes launched by your application are terminated as part of container cleanup.

Windows-based stream groups with stream classes such as `gen4n_win2022`, `gen5n_win2022`, `gen6n_ultra_win2022`, or `gen6n_pro_win2022` execute your applications directly on the host operating system. Each stream session begins with a copy of the application files and an empty user profile folder. When a session terminates, the user profile folder and application folder are fully reset. Sub-processes launched by your application are terminated. If your application modifies files outside of the user profile folder and the application folder, or modifies the system registry, then those changes might persist across multiple sessions.

For any stream group configuration, the underlying compute resources and operating system environment will be reused over time to launch new stream sessions. Under the Shared Responsibility Model, it is your responsibility to maintain the security of your applications and avoid executing untrusted code or modifying critical operating system files.

Multi-tenant stream groups

Stream groups are either single-tenant or multi-tenant, depending on your selection of stream class. Multi-tenant stream classes such as `gen4n_high` or `gen5n_high` share one GPU across multiple simultaneous sessions. In this context, multi-tenancy refers to running more than one session at a time on the underlying hardware. The hardware is still dedicated to your stream group and is not shared across stream groups or with other AWS customers.

This multi-tenant stream group model is unique to Amazon GameLift Streams and comes with important security and performance implications. The security posture of a multi-tenant stream group is equivalent to hosting multiple application containers on a single physical server. This posture isn't inherently insecure, but it might amplify the impact of existing security vulnerabilities in your applications. Under the Shared Responsibility Model, it is your responsibility to maintain the security of your applications.

Amazon GameLift Streams makes efforts to ensure that multi-tenant sessions do not interfere with each other. However, if an application consumes CPU or GPU resources without regard for the defined limits of the stream class, this can have an impact on other streams that are trying to use the same shared resources. For example, in a "high" stream group with two tenants per GPU, a greedy application can negatively impact up to one other stream. Your application should regulate its own resource consumption. If your application cannot self-regulate and your use case has no tolerance for potential "noisy neighbor" performance variations, a single-tenant stream class, such as `gen5n_win2022`, `gen6n_pro_win2022`, `gen5n_ultra`, or `gen6n_ultra`, is recommended.

Interface VPC endpoints in Amazon GameLift Streams

You can improve the security posture of your VPC by configuring Amazon GameLift Streams to use an interface VPC endpoint. Interface endpoints are powered by AWS PrivateLink, a technology that allows you to privately access Amazon GameLift Streams APIs by using private IP addresses. AWS PrivateLink restricts all network traffic between your VPC and Amazon GameLift Streams to the Amazon network. You don't need an internet gateway, a NAT device, or a virtual private gateway.

For more information about AWS PrivateLink and VPC endpoints, see VPC endpoints in the *Amazon VPC User Guide*.

> **Note**
>
> AWS PrivateLink is only applicable to API endpoints. Amazon GameLift Streams managed stream sessions always use public network addresses.

Creating the VPC endpoints for Amazon GameLift Streams

To create the VPC endpoint for the Amazon GameLift Streams service, use the Access an AWS service using an interface VPC endpoint procedure in the *Amazon VPC User Guide* to create the following endpoint:

- `com.amazonaws.region.gameliftstreams`

> **Note**
>
> *region* represents the Region identifier for an AWS Region supported by Amazon GameLift Streams, such as `us-east-2` for the US East (Ohio) Region.

Creating a VPC endpoint policy for Amazon GameLift Streams

You can attach an endpoint policy to your VPC endpoint that controls access to Amazon GameLift Streams. The policy specifies the following information:

- The principal that can perform actions.
- The actions that can be performed.
- The resources on which actions can be performed.

For more information, see Control access to VPC endpoints using endpoint policies in the *Amazon VPC User Guide*.

Example Example: VPC endpoint policy for Amazon GameLift Streams

The following is an example of an endpoint policy for Amazon GameLift Streams. When attached to an endpoint, this policy grants permission to create and list stream groups.

```
{
  "Statement":[
    {
      "Effect": "Allow",
      "Principal": "*",
      "Action": [
        "gameliftstreams:CreateStreamGroup",
```

```
        "gameliftstreams:ListStreamGroups"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

Configuration and vulnerability analysis in Amazon GameLift Streams

Configuration and IT controls are a shared responsibility between AWS and you, our customer. For more information, see the AWS shared responsibility model. AWS 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 Streams:

- Customers are responsible for the management of software deployed to Amazon GameLift Streams stream groups for stream hosting. Specifically:
  - Customer-provided application content and software should be maintained, including updates and security patches. To update, create a new Amazon GameLift Streams application and deploy it to new stream groups.
  - At this time, the operating system and runtime environment for a stream group is updated only when you create a new stream group. To patch, update, and secure the operating system and other applications that are part of the runtime environment, we recommend that you recycle stream groups every two to four weeks, regardless of application updates.
- Customers should consider regularly updating their games with the latest SDK versions, including the AWS SDK and the Amazon GameLift Streams Web Client SDK.

Security best practices for Amazon GameLift Streams

Amazon GameLift Streams 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.

- At this time, the operating system and runtime environment for a stream group is updated only when you create a new stream group. To patch, update, and secure the operating system and other applications that are part of the runtime environment, we recommend that you recycle stream groups every two to four weeks, regardless of application updates.
- Best practices for security, identity, and compliance

Monitoring Amazon GameLift Streams

Monitoring is an important part of maintaining the reliability, availability, and performance of Amazon GameLift Streams and your other AWS solutions. AWS provides the following monitoring tools to watch Amazon GameLift Streams, report when something is wrong, and take automatic actions when appropriate:

- *Amazon CloudWatch* monitors your AWS resources and the applications you run on AWS in real time. You can collect and track metrics, create customized dashboards, and set alarms that notify you or take actions when a specified metric reaches a threshold that you specify. For more information, see the Amazon CloudWatch User Guide.
- With *Amazon CloudWatch Logs* you can monitor, store, and access your log files from services like Amazon Elastic Compute Cloud, AWS CloudTrail, and other sources. CloudWatch Logs can monitor information in the log files and notify you when your services meet certain thresholds. You can also archive your log data in highly durable storage. For more information, see the Amazon CloudWatch Logs User Guide.
- *AWS CloudTrail* captures API calls and related events made by or on behalf of your AWS account and delivers the log files to an Amazon Simple Storage Service bucket that you specify. You can identify which users and accounts called AWS, the source IP address from which the calls were made, and when the calls occurred. For more information, see the AWS CloudTrail User Guide.
- *Real-time performance stats* collect application-level and shared system-level performance stats during stream sessions. You can receive these stats in real-time on the client or post-session as a CSV file in exported session files. Using this feature, you can monitor the CPU, memory, GPU, and VRAM utilization of your stream. For more information, see the section called “Real-time performance stats”.

Monitor Amazon GameLift Streams with Amazon CloudWatch

You can monitor Amazon GameLift Streams using CloudWatch, which collects raw data and processes it into readable, near real-time metrics. These statistics are kept for 15 months, so that you can access historical information and gain a better perspective on how your web application or service is performing. You can also 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.

Amazon GameLift Streams provides metrics to help customers monitor the following:

- Stream group capacity and usage.
- Stream performance and resource usage.
- Stream status to resolve issues and support users.
- Customer engagement across content offerings.
- Data channel usage.

The following tables list the dimensions and metrics for Amazon GameLift Streams.

Stream group capacity and usage

Use these metrics to help scale resources to meet demand. These metrics are published every minute.

> **Important**
>
> **For stream groups created before September 5, 2025**
> Due to an issue with CloudWatch's data retention policy, accurate capacity metrics are only available for the last 15 days. For capacity metrics that are older than 15 days, no data will be visible when the period is 1 minute, and the data shown will be inaccurate when the period is 5 minutes or more.
> As a workaround, you can add `SUM(METRICS())/5` math (for example, when using a 5-minute period) to a sum-type statistic in your CloudWatch graph as a workaround to see accurate capacity counts beyond the 15 day, 1-minute metrics retention limitation.
> To pick up the fix for this issue, recreate your stream groups.

Metric	Description	Dimension	Unit
ActiveCapacity	The number of compute resources that are provisioned and ready to stream. It includes resources that are currently streaming and resources that are idle and ready to respond to new stream requests.	(StreamGroupId, Location)	Count

Metric	Description	Dimension	Unit
IdleCapacity	The numerical portion of active capacity that is not currently streaming. It represents the availability of compute resources to respond to new stream requests.	(StreamGroupId, Location)	Count

Stream group performance and resource utilization

These metrics are published every minute.

Metric	Description	Dimension	Unit
MemoryUtilization	% of available memory used by the stream.	(StreamGroupId, Location), (ApplicationId, StreamClass)	Percentage
CPUUtilization	% of available CPU used by the stream.	(StreamGroupId, Location), (ApplicationId, StreamClass)	Percentage
FrameCaptureRate	Rate at which frames are captured from the application.	(StreamGroupId, Location), (ApplicationId, StreamClass)	None

Metric	Description	Dimension	Unit
AudioCaptureRate	Rate at which audio samples are captured from the application.	(StreamGroupId, Location), (ApplicationId, StreamClass)	None
RoundTripTime	Round trip time between client and server.	(StreamGroupId, Location), (ApplicationId, StreamClass)	ms

Stream status

These metrics are published at the end of a stream session.

Metric	Description	Dimension	Unit
TerminatedStreamSessions	Number of sessions ended in state `TERMINATED`	(StreamGroupId, Location), (ApplicationId, StreamClass)	Count
ErroredStreamSessions	Number of sessions ended in state `ERROR`	(StreamGroupId, Location), (ApplicationId, StreamClass)	Count

Customer engagement

These metrics are published at the end of a stream session..

Metric	Description	Dimension	Unit
Session Length	Stream session duration	(StreamGroupId, Location), (ApplicationId, StreamClass)	Seconds

Data channels

These metrics are published at the end of a stream session.

Metric	Description	Dimension	Unit
DataChannel-ApplicationConnected	Number of times your application connects to the data channel port. This number is at most 1 per stream session.	(StreamGroupId, Location), (ApplicationId, StreamClass)	Count
DataChannel-ApplicationMessage	Number of messages your application has sent to your client.	(StreamGroupId, Location), (ApplicationId, StreamClass)	Count
DataChannel-ApplicationMessageBytes	Total bytes of messages your application has sent to your client.	(StreamGroupId, Location), (Applicat	Bytes

Metric	Description	Dimension	Unit
		ionId, StreamClass)	
DataChannel-ClientMessage	Number of messages your client has sent to your application.	(StreamGroupId, Location), (ApplicationId, StreamClass)	Count
DataChannel-ClientMessageBytes	Total bytes of messages your client has sent to your application.	(StreamGroupId, Location), (ApplicationId, StreamClass)	Bytes

Logging Amazon GameLift Streams API calls using AWS CloudTrail

Amazon GameLift Streams is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service. CloudTrail captures all API calls for Amazon GameLift Streams as events. The calls captured include calls from the Amazon GameLift Streams console and code calls to the Amazon GameLift Streams API operations. Using the information collected by CloudTrail, you can determine the request that was made to Amazon GameLift Streams, the IP address from which the request was made, when it was made, and additional details.

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 user or user credentials.
- Whether the request was made on behalf of an IAM Identity Center user.
- Whether the request was made with temporary security credentials for a role or federated user.

- Whether the request was made by another AWS service.

CloudTrail is active in your AWS account when you create the account and you automatically have access to the CloudTrail **Event history**. The CloudTrail **Event history** provides a viewable, searchable, downloadable, and immutable record of the past 90 days of recorded management events in an AWS Region. For more information, see Working with CloudTrail Event history in the *AWS CloudTrail User Guide*. There are no CloudTrail charges for viewing the **Event history**.

For an ongoing record of events in your AWS account past 90 days, create a trail or a CloudTrail Lake event data store.

CloudTrail trails

A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. All trails created using the AWS Management Console are multi-Region. You can create a single-Region or a multi-Region trail by using the AWS CLI. Creating a multi-Region trail is recommended because you capture activity in all AWS Regions in your account. If you create a single-Region trail, you can view only the events logged in the trail's AWS Region. For more information about trails, see Creating a trail for your AWS account and Creating a trail for an organization in the *AWS CloudTrail User Guide*.

You can deliver one copy of your ongoing management events to your Amazon S3 bucket at no charge from CloudTrail by creating a trail, however, there are Amazon S3 storage charges. For more information about CloudTrail pricing, see AWS CloudTrail Pricing. For information about Amazon S3 pricing, see Amazon S3 Pricing.

CloudTrail Lake event data stores

CloudTrail Lake lets you run SQL-based queries on your events. CloudTrail Lake converts existing events in row-based JSON format to Apache ORC format. ORC is a columnar storage format that is optimized for fast retrieval of data. Events are aggregated into *event data stores*, which are immutable collections of events based on criteria that you select by applying advanced event selectors. The selectors that you apply to an event data store control which events persist and are available for you to query. For more information about CloudTrail Lake, see Working with AWS CloudTrail Lake in the *AWS CloudTrail User Guide*.

CloudTrail Lake event data stores and queries incur costs. When you create an event data store, you choose the pricing option you want to use for the event data store. The pricing option determines the cost for ingesting and storing events, and the default and maximum

retention period for the event data store. For more information about CloudTrail pricing, see AWS CloudTrail Pricing.

Amazon GameLift Streams data events in CloudTrail

Data events provide information about the resource operations performed on or in a resource (for example, starting a stream session in a stream group). These are also known as data plane operations. Data events are often high-volume activities. By default, CloudTrail doesn't log data events. The CloudTrail **Event history** doesn't record data events.

Additional charges apply for data events. For more information about CloudTrail pricing, see AWS CloudTrail Pricing.

You can log data events for the Amazon GameLift Streams resource types by using the CloudTrail console, AWS CLI, or CloudTrail API operations. For more information about how to log data events, see Logging data events with the AWS Management Console and Logging data events with the AWS Command Line Interface in the *AWS CloudTrail User Guide*.

The following table lists the Amazon GameLift Streams resource types for which you can log data events. The **Resource type (console)** column shows the value to choose from the **Resource type** list on the CloudTrail console. The **resources.type value** column shows the `resources.type` value, which you would specify when configuring advanced event selectors using the AWS CLI or CloudTrail APIs. The **Data APIs logged to CloudTrail** column shows the API calls logged to CloudTrail for the resource type.

Resource type (console)	resources.type value	Data APIs logged to CloudTrail
GameLift Streams application	`AWS::GameLiftStreams::Application`	• StartStreamSession
GameLift Streams stream group	`AWS::GameLiftStreams::StreamGroup`	• CreateStreamSessionConnection • ExportStreamSessionFiles • GetStreamSession • ListStreamSessions

Resource type (console)	resources.type value	Data APIs logged to CloudTrail
		• ListStreamSessions ByAccount • StartStreamSession • TerminateStreamSession

You can configure advanced event selectors to filter on the `eventName`, `readOnly`, and `resources.`ARN fields to log only those events that are important to you. For more information about these fields, see AdvancedFieldSelector in the *AWS CloudTrail API Reference*.

Amazon GameLift Streams management events in CloudTrail

Management events provide information about management operations that are performed on resources in your AWS account. These are also known as control plane operations. By default, CloudTrail logs management events.

Amazon GameLift Streams logs the following Amazon GameLift Streams control plane operations to CloudTrail as *management events*.

- AddStreamGroupLocations
- AssociateApplications
- CreateApplication
- CreateStreamGroup
- DeleteApplication
- DeleteStreamGroup
- DisassociateApplications
- GetApplication
- GetStreamGroup
- ListApplications
- ListStreamGroups
- ListTagsForResource
- RemoveStreamGroupLocations

- TagResource
- UntagResource
- UpdateApplication
- UpdateStreamGroup

Amazon GameLift Streams event examples

An event represents a single request from any source and includes information about the requested API operation, the date and time of the operation, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so events don't appear in any specific order.

The following example shows a CloudTrail management event that demonstrates the CreateApplication operation.


```
{
      "eventVersion": "1.09",
      "userIdentity": {
         "type": "AssumedRole",
         "principalId": "AROA123456789EXAMPLE:assume-temporary-gameliftstreams-access-
role",
         "arn": "arn:aws:sts::111122223333:assumed-role/GameLiftStreamsTestRole/assume-
temporary-gameliftstreams-access-role",
         "accountId": "111122223333",
         "accessKeyId": "ASIAIOSFODNN7EXAMPLE",
         "sessionContext": {
               "sessionIssuer": {
                  "type": "Role",
                  "principalId": "AROA123456789EXAMPLE",
                  "arn": "arn:aws:iam::111122223333:role/GameLiftStreamsTestRole",
                  "accountId": "111122223333",
                  "userName": "GameLiftStreamsTestRole"
               },
               "webIdFederationData": {},
               "attributes": {
                  "creationDate": "2025-07-23T21:18:19Z",
                  "mfaAuthenticated": "false"
               }
         }
      },
      "eventTime": "2025-07-23T21:58:54Z",
```

```
      "eventSource": "gameliftstreams.amazonaws.com",
      "eventName": "CreateApplication",
      "awsRegion": "us-west-2",
      "sourceIPAddress": "203.0.113.0",
      "userAgent": "aws-sdk-javascript/2.0.0 Linux/4.14.291-218.527.amzn2.x86_64
 OpenJDK_64-Bit_Server_VM/11.0.17+9-LTS Java/11.0.17 vendor/Amazon.com_Inc. exec-env/
AWS_ECS_FARGATE io/sync http/Apache cfg/retry-mode/legacy",
      "requestParameters": {
         "ApplicationSourceUri": "s3://amzn-s3-demo-bucket/MyGame",
         "Description": "MyGame canary - Proton 8",
         "RuntimeEnvironment": {
               "Type": "PROTON",
               "Version": "20230704"
         },
         "ClientToken": "a1b2c3d4-5678-90ab-cdef-EXAMPLE33333",
         "ExecutablePath": "MyGame100.exe"
      },
      "responseElements": {
         "Status": "INITIALIZED",
         "ApplicationSourceUri": "s3://amzn-s3-demo-bucket/MyGame",
         "Description": "MyGame canary - Proton 8",
         "RuntimeEnvironment": {
               "Type": "PROTON",
               "Version": "20230704"
         },
         "LastUpdatedAt": 1753307934.293,
         "CreatedAt": 1753307934.293,
         "Id": "a-9ZY8X7Wv6",
         "Arn": "arn:aws:gameliftstreams:us-west-2:111122223333:application/
a-9ZY8X7Wv6",
         "ExecutablePath": "MyGame100.exe"
      },
      "requestID": "a1b2c3d4-5678-90ab-cdef-EXAMPLE11111",
      "eventID": "a1b2c3d4-5678-90ab-cdef-EXAMPLEbbbbb",
      "readOnly": false,
      "eventType": "AwsApiCall",
      "managementEvent": true,
      "recipientAccountId": "111122223333",
      "eventCategory": "Management"
}
```


The following example shows a CloudTrail data event from a trail log that demonstrates the StartStreamSession operation.

```
{
    "Records": [
        {
            "eventVersion": "1.09",
            "userIdentity": {
                "type": "AssumedRole",
                "principalId": "AROA123456789EXAMPLE:assume-temporary-gameliftstreams-
access-role",
                "arn": "arn:aws:sts::111122223333:assumed-role/GameLiftStreamsTestRole/
assume-temporary-gameliftstreams-access-role",
                "accountId": "111122223333",
                "accessKeyId": "ASIAIOSFODNN7EXAMPLE",
                "sessionContext": {
                    "sessionIssuer": {
                        "type": "Role",
                        "principalId": "AROA123456789EXAMPLE",
                        "arn": "arn:aws:iam::111122223333:role/
GameLiftStreamsTestRole",
                        "accountId": "111122223333",
                        "userName": "GameLiftStreamsTestRole"
                    },
                    "attributes": {
                        "creationDate": "2025-07-23T21:18:19Z",
                        "mfaAuthenticated": "false"
                    }
                }
            },
            "eventTime": "2025-07-23T23:43:46Z",
            "eventSource": "gameliftstreams.amazonaws.com",
            "eventName": "StartStreamSession",
            "awsRegion": "us-east-2",
            "sourceIPAddress": "203.0.113.0",
            "userAgent": "Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36
 (KHTML, like Gecko) Chrome/138.0.0.0 Safari/537.36",
            "requestParameters": {
                "Identifier": "sg-1AB2C3De4",
                "Description": "StreamGroup sg-1AB2C3De4 Application a-9ZY8X7Wv6
 Console stream",
                "AdditionalLaunchArgs": [],
                "UserId": "a1b2c3d4-5678-90ab-cdef-EXAMPLE11111",
                "Locations": [
                    "us-east-2"
                ],
```

```
                "SignalRequest": "***",
                "Protocol": "WebRTC",
                "ApplicationIdentifier": "a-9ZY8X7Wv6",
                "ClientToken": "a1b2c3d4-5678-90ab-cdef-EXAMPLE22222",
                "ConnectionTimeoutSeconds": 100,
                "AdditionalEnvironmentVariables": {}
            },
            "responseElements": {
                "Status": "ACTIVATING",
                "ApplicationArn": "arn:aws:gameliftstreams:us-
west-2:111122223333:application/a-9ZY8X7Wv6",
                "Description": "StreamGroup sg-1AB2C3De4 Application a-9ZY8X7Wv6
 Console stream",
                "LastUpdatedAt": 1.753314225925E9,
                "CreatedAt": 1.753314225925E9,
                "AdditionalEnvironmentVariables": {},
                "ConnectionTimeoutSeconds": 100,
                "AdditionalLaunchArgs": [],
                "StreamGroupId": "sg-1AB2C3De4",
                "UserId": "a1b2c3d4-5678-90ab-cdef-EXAMPLE11111",
                "SessionLengthSeconds": 43200,
                "SignalRequest": "***",
                "Arn": "arn:aws:gameliftstreams:us-west-2:111122223333:streamsession/
sg-1AB2C3De4/ABC123def4567",
                "Protocol": "WebRTC",
                "WebSdkProtocolUrl": "https://123456789012.cloudfront.net/
e3b0c44298fc1c149afbf4c8996fb92427ae41e4649b934ca495991b7852b855.js"
            },
            "requestID": "a1b2c3d4-5678-90ab-cdef-EXAMPLEaaaaa",
            "eventID": "a1b2c3d4-5678-90ab-cdef-EXAMPLEbbbbb",
            "readOnly": false,
            "resources": [
                {
                    "accountId": "111122223333",
                    "type": "AWS::GameLiftStreams::StreamGroup",
                    "ARN": "arn:aws:gameliftstreams:us-west-2:111122223333:streamgroup/
sg-1AB2C3De4"
                },
                {
                    "accountId": "111122223333",
                    "type": "AWS::GameLiftStreams::Application",
                    "ARN": "arn:aws:gameliftstreams:us-west-2:111122223333:application/
a-9ZY8X7Wv6"
                }
```

```
            ],
            "eventType": "AwsApiCall",
            "managementEvent": false,
            "recipientAccountId": "111122223333",
            "eventCategory": "Data"
        }
    ]
}
```


For information about CloudTrail record contents, see CloudTrail record contents in the *AWS CloudTrail User Guide*.

Real-time performance stats

Amazon GameLift Streams collects performance stats during active stream sessions, measuring resource utilization every second. Use these stats to monitor your application's performance, identify resource bottlenecks, and optimize your streaming experience.

Performance stats include both application-level stats (CPU and memory utilization for your specific application) and system-level stats (CPU, memory, GPU, and VRAM utilization for the shared compute infrastructure).

You can receive performance stats in two ways:

- **In real-time during the session:** Use the Amazon GameLift Streams Web SDK to receive stats as they're collected. This enables you to build performance overlays and monitor resource utilization as you interact with the application.
- **Post-session as a CSV file:** When you export session files, the stats are included as `stats/perf_stats_v1.csv`. This provides a complete record for post-session analysis and debugging.

Receive performance stats

Receive stats in real-time

To receive performance stats in your client application during an active session, set the `SharedWithClient` parameter to **`true`** when calling the `StartStreamSession` API. The Amazon GameLift Streams Web SDK provides a `performanceStats` callback that triggers whenever new stats arrive from the streaming session.

> **Warning**
>
> Do not enable `SharedWithClient` for production sessions with end users. Enable it only when the client is trusted, such as for internal debugging and testing.

When initializing the Amazon GameLift Streams Web SDK, set `clientConnection.performanceStats` to a callback function that will receive performance stats.

```
const gls = new gameliftstreams.GameLiftStreams({
    videoElement: document.getElementById('streamVideoElement'),
    audioElement: document.getElementById('streamAudioElement'),
    inputConfiguration: {
        ...
    },
    clientConnection: {
        ...
        performanceStats: (perfStats) => {
            // Your callback logic here
            console.log('CPU: ' + perfStats.application.cpuNormalized);
            console.log('Memory: ' + perfStats.application.memoryMB + ' MB');
            console.log('GPU: ' + perfStats.system.gpuPercent + '%');
        },
    }
});
```

The callback receives a `PerformanceStats` object containing both application-level and system-level stats. For details on the interface structure, see the Amazon GameLift Streams Web SDK documentation on the Getting Started product page.

The Amazon GameLift Streams console also includes a built-in performance overlay when using the test stream feature, allowing you to monitor stats in real-time without any implementation work.

You can combine performance stats with WebRTC stats provided by the `getVideoRTCStats()` and `getAudioRTCStats()` functions in the Amazon GameLift Streams Web SDK. This combination provides a complete picture of streaming performance, including network stats, client frame rate, and resource utilization.

Receive stats post-session

Amazon GameLift Streams automatically collects performance stats during every stream session. When you export session files, the stats are included as `stats/perf_stats_v1.csv` in the exported ZIP file. This provides a complete record of all stats collected during the session for post-session analysis and debugging.

For more information about exporting session files, see the section called “Export stream session files”.

Performance stats reference

The following table lists all performance stats collected by Amazon GameLift Streams. Application stats are specific to the current session, while shared system stats reflect the total utilization of the shared compute by sessions on multi-tenant stream classes.

> **Normalized stats on multi-tenant stream classes**
>
> Amazon GameLift Streams supports multi-tenant stream classes where multiple sessions may share the same compute instance. Normalized stats (application CPU and memory utilization) measure your application's resource usage relative to its allocated fair share. The fair share is calculated by dividing the total available CPU and memory on the compute instance evenly based on the stream class tenancy.
> A value of 1.0 means your application is using exactly its fair share allocation. Values below 1.0 indicate you're using less than your allocation. Values exceeding 1.0 indicate overutilization, which may lead to performance degradation for your session. On multi-tenant stream classes (tenancy greater than 1), overutilization may also impact other sessions sharing the same compute instance.

The stat names listed in the following table are used as CSV column headers in the exported file. When receiving stats in real-time via the Amazon GameLift Streams Web SDK, these stats are available through the `PerformanceStats` interface with property names in camel case. For the exact interface structure and property names, see the Amazon GameLift Streams Web SDK API reference guide on the Getting Started product page.

Stat name (CSV column)	Description	Scope
`timestamp`	Time when the measurement was taken, in ISO 8601 format.	All
`app_cpu_normalized`	Application's CPU usage normalized against the fair share allocation, where 1.0 represents the target fair share limit. Usage over 1.0 indicates overutilization, which may lead to performance issues	Application
`app_mem_mb`	Total memory (RAM) used by the application (measured in MiB)	Application
`app_mem_normalized`	Application's memory usage normalized against the fair share allocation, where 1.0 represents the target fair share limit. Usage over 1.0 indicates overutilization, which may lead to performance issues	Application
`shared_sys_cpu_pct`	Percentage of total CPU usage across the shared compute.	Shared System
`shared_sys_mem_mb`	Total memory used on the instance (measured in MiB).	Shared System
`shared_sys_mem_pct`	Percentage of total memory in use across the shared compute.	Shared System
`shared_sys_gpu_pct`	Percentage of total GPU utilization across the shared compute.	Shared System
`shared_sys_vram_mb`	Total VRAM (GPU memory) used on the shared compute (measured in MiB).	Shared System
`shared_sys_vram_pct`	Percentage of total VRAM (GPU memory) in use across the shared compute.	Shared System

Troubleshooting Amazon GameLift Streams

Topics

- Access denied when making a request to Amazon GameLift Streams service
- Application issues
- Performance issues
- Stream connectivity and network performance issues
- Stream input issues
- Stream session issues
- Testing and troubleshooting compatibility with Proton for Amazon GameLift Streams
- Profiling Unreal Engine performance

Access denied when making a request to Amazon GameLift Streams service

If you encounter AccessDenied exceptions when making calls to Amazon GameLift Streams APIs or working with resources in the console, your AWS Identity and Access Management (IAM) role might have insufficient permissions for Amazon GameLift Streams. Check the following:

- If the IAM role has an explicit "deny-all" policy, you must explicitly list Amazon GameLift Streams as an exception to that policy by adding `"gameliftstreams:*"` to the NotAction element. For example:

```
{
    "Sid": "DenyAllExceptListedIfNoMFA",
    "Effect": "Deny",
    "NotAction": [
        "iam:CreateVirtualMFADevice",
        "iam:EnableMFADevice",
        "iam:GetUser",
        "iam:ListMFADevices",
        "iam:ListVirtualMFADevices",
        "iam:ResyncMFADevice",
        "sts:GetSessionToken",
        "gameliftstreams:*"                    // Add this
    ],
    "Resource": "*",
    "Condition": {
        "BoolIfExists": {"aws:MultiFactorAuthPresent": "false"}
    }
}
```

- For more information, see Identity and Access Management for Amazon GameLift Streams in the Security chapter, and review Troubleshooting access denied error messages in the *IAM User Guide*.

Application issues

This section identifies potential causes for issues that prevent applications from running or cause them to appear differently on Amazon GameLift Streams.

Preliminary checks

- Run your application on a different machine to verify that it's correctly packaged. This confirms that your application content doesn't contain any hardcoded paths, missing assets, libraries, or binaries that might not work on other devices.

Proton issues

- **Verify that your application is compatible with Proton.** Test your application on a local environment without the Amazon GameLift Streams server to verify that it's compatible with

Proton. For instructions, see Testing and troubleshooting compatibility with Proton for Amazon GameLift Streams.

Application issues due to screen resolution

Applications might freeze, crash, or render incorrectly if you attempt to use a full-screen resolution that is not 1920x1080. We recommend that you use a borderless fullscreen window to run your application and do not attempt to change the resolution.

Application terminates at start of stream session

If your application terminates immediately when a stream session starts, review the following for potential causes and solutions:

- **Verify runtime.** In the Amazon GameLift Streams application configuration, confirm that the file you specified in the **Executable launch path** is an executable file or script and is correct for the runtime environment that you selected. Windows applications should have a file type of ".exe", ".bat", or ".cmd" and target either of the Windows or Proton runtimes. Native Linux applications should be executable files that target the Ubuntu 22.04 LTS runtime.
- **Verify required DLLs.** Your Windows application might be missing required DLLs. For example, if your application is a debug build, then it requires the debug version of the Microsoft C and C ++ (MSVC) runtime libraries. To resolve this, we recommend that you package your build and DLLs side-by-side. For instructions, refer to Prepare a test machine to run a debug executable by Microsoft.

In general, we recommend that you test your build on a clean machine first, before trying on Amazon GameLift Streams. For instructions about testing on an Amazon EC2 instance, refer to Set up a remote machine.

Unreal Engine application crashes or requires additional dependencies

If your Unreal Engine application fails to start, crashes, or requires you to install additional dependencies, such as the Microsoft C and C++ (MSVC) runtime, try the following:

- **Use the correct executable.** For your application to work correctly with Amazon GameLift Streams, set the application path to the full executable that's located in the `Binaries/Win64/` (or similar) subfolder. Unreal Engine produces two executables: a small bootstrap executable at

the root of the folder, and a platform target executable in the `Binaries/Win64/` subfolder. The bootstrap executable at the root attempts to validate pre-conditions are correct and can create false positives on Amazon GameLift Streams that prevent application launch. If the platform target executable is missing, the application might not have been built correctly. For example, see the following folder structure of a sample Unreal application:

```
BuildApp
|-> MyUnrealApp.exe
|-> MyUnrealApp
            |-> Binaries
                    |-> Win64
                            |-> MyUnrealApp.exe
```

- **Turn off Unreal Engine Asserts.** Disable the `check`, `verify`, and `ensure` macros. They can prevent the application from creating crash dumps. For more information, see Asserts in Unreal Engine documentation.
  - Define `USE_CHECKS_IN_SHIPPING=0` in your build to disable `check` and `verify` macros.
  - Use `-handleensurepercent=0` command-line argument to disable `ensure` macros.

Performance issues

This section identifies potential causes for game performance issues when running on Amazon GameLift Streams, and offers suggestions for optimizing your streams on the service.

Game performance is reduced when streaming on Amazon GameLift Streams

If your game runs well on your own machine but experiences performance issues when you stream it on Amazon GameLift Streams, consider the following:

- Your machine might have more powerful hardware than Amazon GameLift Streams. Make sure to test the application on a machine with similar performance to the hardware that Amazon GameLift Streams uses:
  - gen4n: comparable to NVIDIA RTX 2060 GPU
  - gen5n: comparable to NVIDIA RTX 3080 GPU
  - gen6n: comparable to NVIDIA RTX 4060 GPU

This verifies that your application's rendering settings are compatible with the GPU and that the performance meets your expectations.

- The problem might be due to your network connection or Amazon GameLift Streams's settings. Try the troubleshooting tips in the Stream connectivity issues section.

If your game is slow even when running locally, you'll need to optimize its performance. The best optimization methods will depend on the specific engine or framework you're using.

- For Unreal Engine games, refer to Profiling Unreal Engine performance.

Windows applications experience slow load times or stuttering issues

If your game is experiencing long load times or stuttering behavior, we recommend the following course of action:

1. Ensure your application is packaged and optimized for loading performance using your engine vendor's guidance around optimizing content and shader performance.
2. Ensure your application is set to be the default application in a stream group.
3. Optimize application first launch on the service by caching shaders as part of your application packaging.

There are two approaches to enabling shader caching:

- **Driver-based caching** – This approach is specific to the runtime environment GPU and driver version. This option can be applied to all applications and is therefore the default recommended approach. The steps for this approach will need to be replicated for every GPU/driver combination.
- **Engine-based caching** – This approach enables shader caching through the game engine, if available. It places the burden of creating a pre-baked pipeline state object (PSO) cache on the developer. It also assumes that the engine is capable of handling cache support for different drivers on the same GPU hardware.

As a best practice, we recommend implementing driver-based caching first, because it does not require deep understanding of how PSO caching is implemented for the given engine.

With these implementations, shader files can be exported and packaged with your application so that they don't have to be generated with every new stream start.

To implement a driver-based caching fix for a Windows runtime application

1. Start streaming your default application and play it extensively to generate shaders for the application.

   > **Important**
   >
   > Be sure to visit all the areas or levels of the environment to generate as many shaders as possible.

2. *Before* closing the stream, enable the export feature in your active stream session. For details, see Export stream session files.
3. Download the stream session export .zip file from the Amazon S3 bucket you specified in the previous step. You can find a download link on the Amazon GameLift Streams console on the **Sessions** page.
4. Locate the shaders folder within the stream session export. It's usually saved to this location: `AppData\Local\NVIDIA\DXCache`. Upload the generated shader files (`*.nvph`) to your application's Amazon S3 bucket.
5. Create a `.bat` file that will copy the shader files into the NVIDIA caching folder at runtime. This folder is usually located at: `C:\Users\Administrator\AppData\Local\NVIDIA\DXCache`. Upload the `.bat` file to the Amazon S3 application bucket.
6. Create a new Amazon GameLift Streams application with the `.bat` file as the executable path.

When your application starts streaming, your `.bat` file will copy the pre-generated shaders to the shader cache before launching the application, improving the stream loading performance.

> **Note**
>
> You might need to repeat these steps whenever you update your application or link the Amazon GameLift Streams application to a new stream group. Newer stream groups can contain updated GPU drivers from the service.

The following example `.bat` file assumes that the shader files are stored under the Amazon S3 bucket prefix `Shaders\`. You can use a different folder structure.

```
@echo off
set CURRENT_PATH=%cd%
set DXCACHE_DIR=%CURRENT_PATH%\Shaders
set NVIDIA_DXCACHE_DIR=C:\Users\Administrator\AppData\Local\NVIDIA\DXCache

if not exist "%NVIDIA_DXCACHE_DIR%" (
    mkdir "%NVIDIA_DXCACHE_DIR%"
)

xcopy /s /f "%DXCACHE_DIR%" "%NVIDIA_DXCACHE_DIR%"

start %CURRENT_PATH%\app.exe
```

To implement a driver-based caching fix for a Proton runtime application

1. Start streaming your default application with the following environment variable override:

   ```
   "__GL_SHADER_DISK_CACHE_PATH" : "/home/unpriv/games"
   ```

2. Play application extensively to generate shaders.

   > **Important**
   >
   > Be sure to visit all the areas or levels of the environment to generate as many shaders as possible.

3. *Before* closing the stream, enable the export feature in your active stream session. For details, see Export stream session files.

4. Download the stream session export .zip file from the Amazon S3 bucket you specified in the previous step. You can find a download link on the Amazon GameLift Streams console on the **Sessions** page.

5. Locate the shaders folders and files within the stream session export:

   a. `application\GLCache` folder

   b. if application uses DX11: `application\path-to-exe\exe-name.dxvk-cache` file

c. if application uses DX12: `application\path-to-exe\vkd3d-proton.cache.write` file

6. Upload the generated shader files to your application's Amazon S3 bucket:

   a. Copy the `GLCache` folder into the root directory of your application.

   b. If available, copy the `.dxvk-cache` or `vkd3d-proton.cache.write` cache file to the folder containing the application executable.

7. Create a new Amazon GameLift Streams application with the same Proton configuration.

8. Run the application with the same environment variable override:

```
"__GL_SHADER_DISK_CACHE_PATH" : "/home/unpriv/games"
```

When your application starts streaming, it will use the pre-generated shaders, improving the stream loading performance.

> **Note**
>
> You might need to repeat these steps whenever you update your application or link the Amazon GameLift Streams application to a new stream group. Newer stream groups can contain updated GPU drivers from the service.

To implement an engine-based caching fix for an application using Unreal Engine

For this approach, you can use Unreal Engine features to create a pipeline state object (PSO) cache for your Amazon GameLift Streams application. A PSO cache lets you deliver pre-compiled graphics pipeline states with decreased runtime compilation times, which can reduce hitches during loading and rendering. This requires advanced knowledge of Unreal Engine, and therefore we will not cover all the engine-specific details here. For additional instructions, refer to the guidance from Unreal Engine in Creating a Bundled PSO Cache, "Collection Flow" section.

1. Generate shaders for your application that has PSO logging enabled.

   a. Create a new Amazon GameLift Streams application using the packaged build with the PSO-enabled application.

b. Start a stream with `-logPSO` command in your PSO logging app. You can use the command-line arguments option on the **Test stream** configuration page in the Amazon GameLift Streams console.

> **Important**
>
> Be sure to visit all the areas or levels of the environment to generate as many shaders as possible.

c. *Before* closing the stream, enable the export feature in your active stream session. For details, see Export stream session files.

d. Quit the application from the menu or by using Unreal shutdown commands. If you close the stream directly, the Unreal shaders pipeline file won't be generated.

e. Download the stream session export .zip file from the Amazon S3 bucket you specified in the export step. You can find a download link on the Amazon GameLift Streams console on the **Sessions** page.

2. Package the Unreal shaders pipeline file into your Amazon GameLift Streams application.

a. Locate the recorded PSO files (`rec.pipelinecache`) in the stream session export under `Saved/CollectedPSOs`. Unpack the PSO files using Unreal commands.

b. Package a new Unreal build with the generated output from the unpacking. Follow the Unreal guidance, sections Converting PSO caches and Including PSO caches in your Application.

> **Important**
>
> When running the Unreal command in the "Converting PSO Caches" section, make sure that you use the same driver's version input files. For example: for DX12, use only the SM6 files as inputs. Otherwise you'll get an error when packaging the new application.

c. Create a new Amazon GameLift Streams application for the new packaged build with the PSO files.

d. When starting and testing streams, confirm that PSO cache is being loaded. Check the game logs for the following line:

```
Opened FPipelineCacheFile: ../../...
```

> **Note**
>
> You might need to repeat these steps whenever you update your application or link the Amazon GameLift Streams application to a new stream group. Newer stream groups can contain updated GPU drivers from the service.

Stream connectivity and network performance issues

When you set up your Amazon GameLift Streams backend service, check the following:

- Choose the closest AWS Region possible to the end user. High latency from your clients to the Region hosting your stream can impact stream quality. Refer to AWS Regions and streaming locations supported by Amazon GameLift Streams for a list of locations where you can stream from. You can ping AWS console endpoints in the Region to get an approximate latency measurement.
- Verify your stream group has capacity for new streams.
- Verify that `ConnectionTimeoutSeconds` is reasonably set to allow end users plenty of time to connect before their web client times out.

Advise your end users to check the following:

- Ensure firewalls allow access to UDP port range 33435-33465 to allow streaming from Amazon GameLift Streams. If Amazon GameLift Streams can't reach these ports, it can lead to streaming issues, such as a black or grey screen.
- Verify that your internet connection can sustain a connection speed of at least 10 Mbps for a 1080p stream. If you detect network issues while playing on Amazon GameLift Streams, your internet speed might be fluctuating and you might not be getting at least 10 Mbps consistently. Run an internet speed test and continue through the troubleshooting steps.
- Use a wired network if possible. When using Wi-Fi, move your device close to your router for stronger signal strength.

- If you're using a Wi-Fi router with both 2.4 GHz and 5 GHz bands, try connecting to a different band. If you're unsure how to switch your router to a different band, visit the support pages of the manufacturer or provider of your Wi-Fi router. You can also contact their customer service.
- Identify if others on the same network (especially when on home Wi-Fi) are running high-bandwidth applications like video streaming, downloading, online gaming, or backups.
- Close other applications on your device that take up bandwidth.
- Don't use a VPN or proxy while streaming. They can cause higher latencies and impact gameplay.
- Verify you're using Wi-Fi instead of cellular networks when playing on an iPad or iPhone. Using a cellular network can result in connectivity issues.
- MacOS users should disable Location Services as it will cause the Wi-Fi to pause from time to time, which will lead to a poor streaming experience.

Stream input issues

This section identifies potential causes and solutions for issues related to user input in a stream session.

General input troubleshooting

- Test to see if the issue is browser-specifc. Overall, we recommend Google Chrome, Microsoft Edge, or a custom Chromium-based desktop application for the best end-user experience and maximum compatibility, particularly with game controllers.
- Log input events sent from client and received by the application to identify where there is an input mismatch in your front end code.
- Be sure to check Supported browsers and input for additional information on supported browsers and input devices, including known issues and limitations.

Gamepad and microphone inputs don't work on native Linux applications

Gamepad and microphone inputs are not supported in native Linux applications. See Supported browsers and input for additional information on supported input devices, including known issues and limitations.

Key input appears stuck on MacOS client

On MacOS clients, keys might suddenly appear to be stuck when the **Command** modifier key and another key are simultaneously pressed, repeating the key event. For example, the arrow key might get stuck when the **Command** key is also pressed. In a game, if arrow keys are used to spin the camera, this would make the camera rotate endlessly.

- Issue: The **Command** key on MacOS maps to the **Meta** key event, which maps to the **Windows** key on Microsoft Windows. The issue is a bug affecting MacOS browsers when **Command** and another key are pressed simultaneously, where the **Meta** key is reset when released but the arrow key is not reset because the browser didn't capture a keyup event for the arrow key, so the Web SDK client won't send a keyup event to the server and the streaming application would still think the key is being pressed.
- Solution: If you are not using the **Command** key, you can filter it out using the Web SDK keyboard filter mechanism (`keyboardFilter`) found in the Web SDK's `InputConfiguration` interface.

Mouse movement behaves differently on Amazon GameLift Streams

If mouse movement behaves differently when streaming with Amazon GameLift Streams, such as moving more quickly than expected, you might need to adjust the mouse-handling and cursor management logic in your application.

- Issue: Amazon GameLift Streams uses a heuristic to pick whether to transmit mouse events in “relative” or “absolute” mode. In relative mode, new mouse updates are provided as small, incremental differences from the previous update. In absolute mode, the mouse cursor is continually forced to a screen position that is synchronized with the client. When the operating system cursor is visible over the streamed content, the heuristic always picks absolute coordinates. This can cause unexpectedly large movement deltas if your application is expecting small, relative updates.
- Solution: If your application expects relative mouse motion (for example, FPS-style camera controls or drag-based interactions), hide the operating system cursor during mouse interactions. For instance, hide the cursor on mouse-down and show it again on mouse-up. This ensures dragging motions use relative coordinates, with absolute position synchronized only when the button is released.

For more information on mouse movement in Amazon GameLift Streams, see Mouse movement handling.

Stream session issues

This section identifies potential causes and solutions for issues related to a stream session starting or terminating unexpectedly.

Stream session does not start

Potential causes:

- Application is hung or crashed. Refer to the Application issues section for troubleshooting instructions.
- Stream group status is not `Active`. Verify the status of the stream group.
- On-demand capacity is taking longer to spin up than the timeout specified by `ConnectionTimeoutSeconds` in the StartStreamSession API. On the Windows runtime, the on-demand spin up time can take 5 minutes or more.
- No available capacity in the streaming location. Verify that your allocated capacity is greater than your in-use capacity, or that you have on-demand capacity that's not in use (allocated capacity is less than always-on capacity plus on-demand capacity). In the console, you can find these values in the list of stream groups or on the stream group detail page. Using the service API, you can find these values using GetStreamGroup. A few scenarios where available capacity is temporarily at zero include the following:
  - If you just increased always-on capacity in the streaming location, wait a few minutes for the capacity to be allocated.
  - If you only have 1 available capacity in the streaming location and your client unexpectedly disconnected, the previous session might still be in a disconnected state. Wait a few minutes for the session to timeout and try again.
  - If you recently added a location to your stream group and the application did not exist at the location, the application might not have finished replicating there. Check the replication status on the stream group details page in the console. Alternatively, you can use the GetApplication API and check the `ReplicationStatuses` value to verify that the `Status` of the desired streaming location is `COMPLETED`.

- Network conditions are so poor that frames, especially the first frame, are not being sent. Check network conditions between the client and the streaming location and adjust or try a different location.

Stream session terminated

Stream sessions automatically terminate when an application crashes or quits, or when the client connection is lost. Sessions can also terminate due to the following timeout values:

- **Placement timeout**: Timeout value for Amazon GameLift Streams to find compute resources to host a stream session.
- **Connection timeout**: Timeout value for a client to connect or reconnect to a stream session.
- **Idle timeout**: Maximum time that a stream session can be idle with no user input.
- **Session length timeout**: Maximum time for a stream session.

For a detailed explanation of each timeout and its possible values, refer to Timeout values affecting stream sessions.

Testing and troubleshooting compatibility with Proton for Amazon GameLift Streams

If your Amazon GameLift Streams application runs on a Proton runtime environment, this section can help you troubleshoot compatibility issues between your application and the Proton layer. These instructions include a set of scripts that installs Proton to your own machine, simulating the environment that Amazon GameLift Streams would use. By troubleshooting without the Amazon GameLift Streams service, you can focus on troubleshooting issues specific to your application and the runtime environment.

High-level steps to test and troubleshoot

1. Acquire an Ubuntu 22.04 LTS machine. You can use either a local machine or an Amazon EC2 cloud-based desktop. Choose from the following topics for instructions:
   - Set up a local machine
   - Set up a remote machine

2. Install the Proton runtime environment to test and debug your application. Refer to Troubleshoot on Proton for guidance.

Known issues with Proton

Refer to the Proton GitHub wiki for the latest compatibility and troubleshooting resources. You can also search for issues in the Proton GitHub issue tracker. Following are some specific issues to be aware of that our customers have encountered when running Windows applications on Proton:

Godot applications on Proton

- Godot-based applications running on Proton might encounter a black screen if the Amazon Vulkan capture layer is enabled. To mitigate this issue, disable shared textures when streaming by setting the environment variable `VK_LAYER_AMZN_BLITSURFACE_SHARED_TEXTURES=0`.

Unreal Engine applications on Proton

- If you run into problems on Proton 8.x with Electra Media Player, (an Unreal Engine plugin) we recommend using the fixes found in https://github.com/ValveSoftware/wine/pull/257.

Set up a local machine to troubleshoot Proton

Proton is a compatibility layer that enables Windows applications to run on Linux. As such, you must have an Ubuntu machine to test and troubleshoot with. If you don't have a local Ubuntu machine, you can set up a remote machine using Amazon EC2. To do so, follow the steps in Set up a remote machine instead.

Prerequisites

- Ubuntu 22.04 LTS. For installation instructions, you can use Ubuntu's Install Ubuntu Desktop tutorial.
- NVIDIA GPU

Install GPU drivers

Installing the latest GPU drivers can prevent your application from poor performance and crashes.

To check what GPU driver your system uses

1. Run the following command in a terminal:

```
lshw -C display | grep driver
```

2. If the correct drivers are installed, you should see the following output, or similar, where *<gpu>* is `nvidia` for NVIDIA: `configuration: driver=<gpu> latency=0`

To install the latest NVIDIA GPU drivers

Follow the instructions in NVIDIA drivers installation.

Verify GPU drivers

Verify that GPU drivers are installed and working correctly. One way to verify this is by running the vkcube application in a terminal.

1. Install the `vulkan-tools` apt package using the following command.

```
sudo apt install -y vulkan-tools
```

2. Run `vkcube`.
3. Review the output.
   - If your system is properly using the correct GPU, you will see output similar to the following, with the name of your GPU: `Selected GPU 0: AMD Radeon Pro V520 (RADV NAVI12), type: 2`
   - If your application isn&t able to use the GPU correctly, you might see different output similar to the following: `Selected GPU 0: llvmpipe (LLVM 15.0.7, 256 bits), type: 4`

     In this case, check the GPU drivers and re-install if needed.

Next step

With your local Ubuntu machine ready, the next step is to set up Proton. For instructions, refer to Troubleshoot on Proton.

Set up a remote Amazon EC2 machine to troubleshoot Proton

If you don't have a local Ubuntu machine, follow these instructions to set up a remote machine instead.

In this step, you will set up your remote Ubuntu machine using Amazon Elastic Compute Cloud (Amazon EC2), which you will use to troubleshoot your application's compatibility with Proton for Amazon GameLift Streams. This topic describes how to set up an Amazon EC2 instance with Ubuntu 22.04 LTS, necessary GPU drivers, and the Amazon DCV Server for a visual remote desktop.

Launch an Amazon EC2 Instance with Ubuntu 22.04 LTS AMI

1. Navigate to Amazon EC2 in the AWS Management Console.
2. Select **Launch Instances**.
3. Enter "Amazon GameLift Streams Testing" for **Name**.
4. Select **Ubuntu Server 22.04 LTS (HVM)** for **Application and OS Images (Amazon Machine Image)**.
5. Select **g4dn.2xlarge** for **Instance Type**.
6. For **Key pair (login)**, choose a key pair if you want to use SSH to access the instance. We recommend using an instance profile with the `AmazonSSMManagedInstanceCore` policy to connect to your instances using AWS Systems Manager Session Manager. For more details, follow Adding Session Manager permissions to an existing IAM role.
7. For **Network settings**, create a new security group:
8. For **Security Group Name**, enter **DCV.**
9. Add **Inbound Security Group Rules** with **Type** `Custom TCP`, **Port Range** `8443`, and **Source Type** `Anywhere` to allow access using Amazon DCV.
10. Increase storage to at least **256GB** and choose **gp3** as the storage type.
11. Choose **Launch Instance**.

Your instance should now be launched.

Follow the instructions in Connect to your Linux instance to connect to the instance using SSH or AWS Systems Manager Session Manager.

Install GPU drivers

G4dn - NVIDIA GPU

Install additional modules and Linux firmware by running the following commands:

```
sudo apt install linux-modules-extra-aws linux-firmware
```

```
# Install the AWS CLI required for NVIDIA driver installation
curl "https://awscli.amazonaws.com/awscli-exe-linux-x86_64.zip" -o "awscliv2.zip"
sudo apt install unzip
unzip awscliv2.zip
sudo ./aws/install
```

Follow the instructions on the NVIDIA GRID drivers for Ubuntu and Debian in Install NVIDIA drivers on Linux.

Set up user environment

Set up your user environment so it can use the GPU by running the following commands. This does the following things:

- Add you to the `video` groups to give you access to a video device, and the `render` group to give you access to a rendering device.
- Install the AWS CLI, which is required for NVIDIA drivers and for downloading your applications or games from Amazon S3.

```
sudo adduser user

# Add the current user to the video and render group
sudo usermod -a -G video user
sudo usermod -a -G render user
sudo adduser user sudo

# Install the AWS CLI
curl "https://awscli.amazonaws.com/awscli-exe-linux-x86_64.zip" -o "awscliv2.zip"
sudo apt install unzip
unzip awscliv2.zip
sudo ./aws/install

sudo reboot
```

Installation and configuration of Amazon DCV

Reconnect to the instance using SSH or AWS Systems Manager Session Manager and follow the instructions from Installing the Amazon DCV Server on Linux for Ubuntu.

- Verify that the server is correctly configured as described in the documentation.

- Follow the steps in Install and configure NVIDIA drivers for NVIDIA GPU.
- Add the Amazon DCV user to video group, as explained in step 7 of the Installing the Server guide (navigate to the Ubuntu tab).

There is no need to install any optional parts of the Amazon DCV Server.

When you're done, run the following command to start the Amazon DCV Server:

```
sudo systemctl start dcvserver
sudo systemctl enable dcvserver
```

Connecting to the Ubuntu Server using the Amazon DCV client

Reconnect to your Ubuntu instance and create a session for a user by running:

```
sudo dcv create-session --owner user --user user my-session --type console
```

You can now use the Amazon DCV Client to access your Ubuntu instance using its public IP address. When you launch a Amazon DCV client, a window appears, allowing you to access your Ubuntu instance through a visual display.

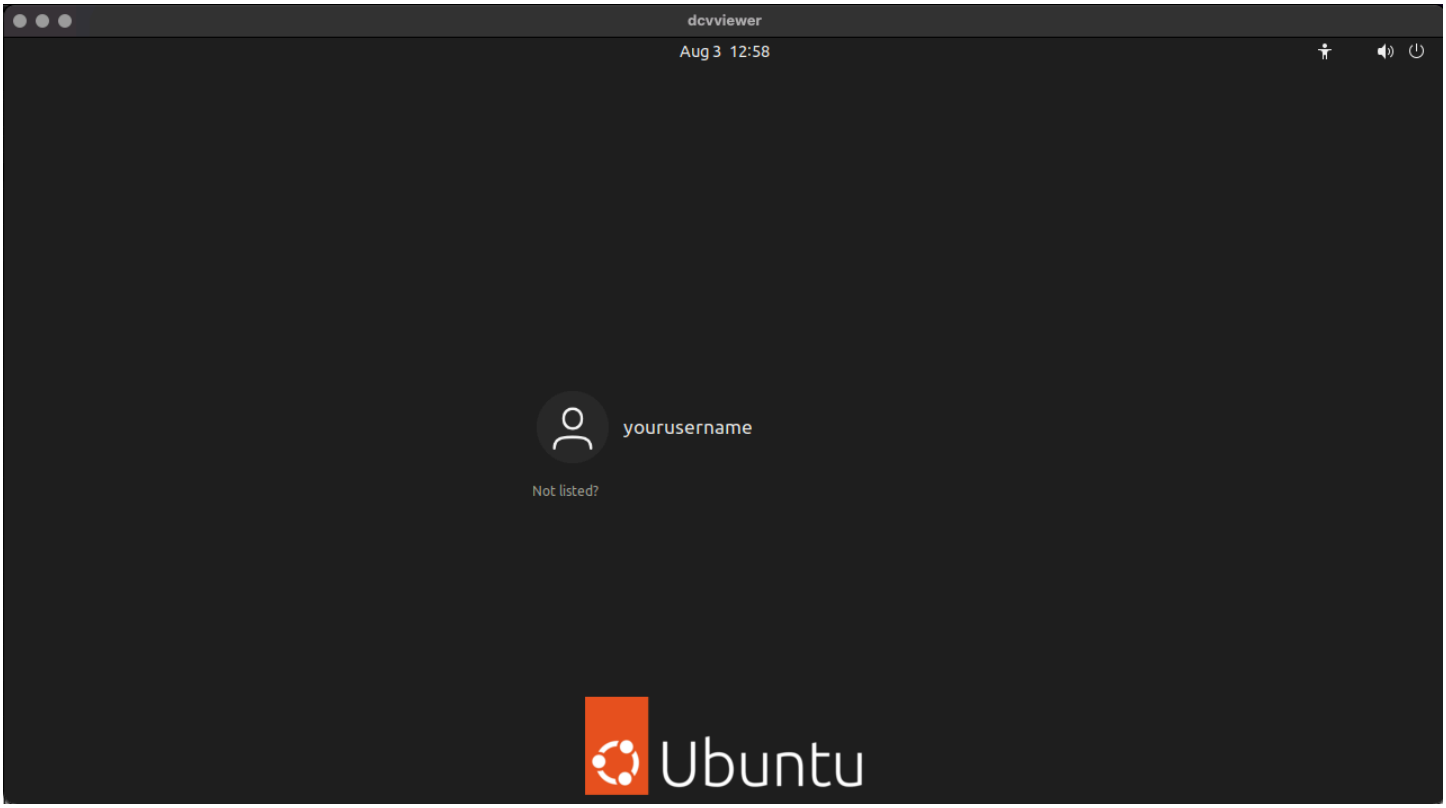

Verify GPU drivers

Verify that GPU drivers are installed and working correctly. One way to verify this is by running the vkcube application in a terminal.

1. Install the `vulkan-tools` apt package using the following command.

```
sudo apt install -y vulkan-tools
```

2. Run `vkcube`.
3. Review the output.
   - If your system is properly using the correct GPU, you will see output similar to the following, with the name of your GPU: `Selected GPU 0: AMD Radeon Pro V520 (RADV NAVI12), type: 2`
   - If your application isn&t able to use the GPU correctly, you might see different output similar to the following: `Selected GPU 0: llvmpipe (LLVM 15.0.7, 256 bits), type: 4`

     In this case, check the GPU drivers and re-install if needed.

Set up Podman (Proton only)

If you're using a Proton runtime, you must install Podman, a container that's used by Proton's build process. Complete the following steps in a terminal.

1. Install Podman, a container that Proton's build process uses.

```
sudo apt install podman
```

2. In the files `/etc/subgid` and `/etc/subgid`
   a. Verify that the files list your Linux machine user name and ID. You can either open the files or use the `cat` command to see what's in the files. Format example: `test:165536:65536`, where `test` corresponds to your user name.
   b. If they're not listed, add them in. Format example: `test:165536:65536`, where `test` corresponds to your user name.

```
$ cat /etc/subuid
              ceadmin:100000:65536
              test:165536:65536

              $ cat /etc/subgid
              ceadmin:100000:65536
              test:165536:65536
```

For more information, refer to Basic Setup and Use of Podman in a Rootless environment in Podman's documentation.

Next step

You now have an Amazon EC2 instance and environment setup to troubleshoot compatibility issues with Amazon GameLift Streams. The next step is to set up Proton. For instructions, refer to Troubleshoot on Proton.

Troubleshoot compatibility on Proton

In this step, you will set up Proton on your own machine, so you can troubleshoot compatibility issues between your Amazon GameLift Streams application and Proton. Running your application in a simulated environment without the Amazon GameLift Streams server can help you identify issues specific to your application and runtime environment.

Prerequisites

- Ubuntu 22.04 LTS with GPU drivers installed. For instructions, refer to Set up a local machine or Set up a remote machine.

Install Proton

To install Proton on your Ubuntu 22.04 LTS machine, use the following script to clone, build, and configure the version of Proton you want to test from the Proton GitHub repository.

1. Copy and paste the following code into a file called `proton-setup.sh` on your Ubuntu 22.04 LTS machine.

```
#!/bin/bash
# This is a script to build Proton. The default build is a tag from the
# experimental_9.0 branch of Proton, but can be changed as a parameter to this
 script.
#
# Usage: ./proton-setup.sh [optional proton_branch_name {default:
 experimental-9.0-20241121b}]
set -e

sudo apt install -y podman make git

# clone proton from github, recurse submodules
# if no proton git link is supplied, use a default tag from the experimental_8.0
 branch
PROTON_BRANCH=${1:-"experimental-9.0-20241121b"}
PROTON_BUILD_DIR=protonBuild
PROTON_DIR=$(pwd)/proton
if git clone https://github.com/ValveSoftware/Proton.git --recurse-submodules --
branch $PROTON_BRANCH proton;
then
 echo "Successfully cloned Proton and its submodules."
else
 echo "Warning: a proton directory/repository already exists. It is recommended to
 delete this folder and re-run this script unless it is a valid repository with
 initialized submodules."
fi

if [ -d $PROTON_BUILD_DIR ];
then
 echo "Error: protonBuild directory already exists. Delete this folder first to
 create a fresh build of Proton before re-running this script."
 exit 1
fi
mkdir $PROTON_BUILD_DIR
cd $PROTON_BUILD_DIR
$PROTON_DIR/configure.sh --enable-ccache --container-engine=podman

# build proton
echo "Building Proton"
make
echo "Done building Proton!"

# prepare proton for execution
```

```
cd dist
mkdir compatdata
if [ -e ./dist ]; then
  PROTON_FILES=dist
elif [ -e ./files ]; then
  PROTON_FILES=files
fi
cp version $PROTON_FILES/
echo "Finished installing proton. Proton binary location: $(pwd)/proton"
echo "STEAM_COMPAT_DATA_PATH: $(pwd)/compatdata"
echo "STEAM_COMPAT_CLIENT_INSTALL_PATH: anything"
```

2. In this step you will run the Proton setup script to clone and install Proton and additional dependencies. The script accepts as an argument the tag or branch name for the Proton version you want to install. To simulate one of the custom builds of Proton that Amazon GameLift Streams provides, use the instructions for that version, below.

   > **Note**
   >
   > Expect the cloning from GitHub to take some time. There are many submodules to download, totalling several gigabytes.

   In your terminal, run the `proton-setup.sh` script and specify the Proton version branch:

   - **Built-in Proton versions**
     - For Proton 9.0-2 (`PROTON-20250516`), use `experimental-9.0-20241121b`.

       ```
       proton-setup.sh experimental-9.0-20241121b
       ```

     - For Proton 8.0-5 (`PROTON-20241007`), use `experimental-8.0-20240205`.

       ```
       proton-setup.sh experimental-8.0-20240205
       ```

       Typically, no additional source code is needed. However, if you run into problems with Electra Media Player, (an Unreal Engine plugin) we recommend using the fixes found in https://github.com/ValveSoftware/wine/pull/257.

> **Note**
>
> For Proton 8.0-2c (`PROTON-20230704`), Amazon GameLift Streams uses a proprietary build, which is not available to buld locally.

- **Recommended custom Proton version**

  For a custom Proton version, we recommend using the Proton experimental_8.0 branch.

  ```
  proton-setup.sh experimental_8.0
  ```

- **Other custom Proton versions**

  For other Proton versions, use an exact branch or tag name listed in Proton releases.

  ```
  proton-setup.sh branch-or-tag-name
  ```

If your installation is successful, the output in your terminal should be similar to the following:

```
...
Done building Proton!
Finished preparing proton. Proton binary location: /home/test/protonBuild/dist/
proton
STEAM_COMPAT_DATA_PATH: /home/test/protonBuild/dist/compatdata
STEAM_COMPAT_CLIENT_INSTALL_PATH: anything
```

Take note of the following variables from the output because you will need them to run Proton in the next step:

- Proton binary location
- `STEAM_COMPAT_DATA_PATH`
- `STEAM_COMPAT_CLIENT_INSTALL_PATH`

Run your application on Proton

The following steps assume that the application executable is located in `path/myapplication/bin/application.exe`. Replace it with the path and file name for your application.

- In a terminal, navigate to the folder where your application executable is located.

```
cd path/myapplication/bin/application.exe
```

- Run your application on Proton. Use the Proton binary location and the environment variables that you got in the previous step.

```
STEAM_COMPAT_DATA_PATH=/home/test/protonBuild/dist/compatdata
 STEAM_COMPAT_CLIENT_INSTALL_PATH=anything /home/test/protonBuild/dist/proton run
 application.exe
```

The application should now attempt to start. If the application starts locally, but not on Amazon GameLift Streams, it may be due to a configuration issue when calling Amazon GameLift Streams APIs. Verify that the API call parameters are correct. Otherwise, continue to the next step for debugging.

Debug the application through log files

If your application has issues running on the local Proton environment, check the output log. The log contains output from your application and runtime environment. Trace where your application is failing to discover issues on the application side.

To dump the log output into a text file, such as `proton.log`, use the following command:

```
STEAM_COMPAT_DATA_PATH=/home/test/protonBuild/dist/compatdata
 STEAM_COMPAT_CLIENT_INSTALL_PATH=anything /home/test/protonBuild/dist/proton run
 application.exe &>proton.log
```

Proton also indicates if the issue is due to a Wine plugin, unimplemented function, missing dlls, and so on. For more information, see Wine HQ's Debugging Wine guide. If you find a Proton or Wine error in the logs that you can't fix on the application side, reach out to your AWS Account Manager or post a question in AWS re:Post for help with further debugging.

Profiling Unreal Engine performance

In this section, learn how to analyze your Unreal Engine game or application performance. This can help you identify area's to optimize, leading to smoother streaming in Amazon GameLift Streams.

You can use Unreal Engine's console and its built-in stat commands to get a detailed look at your game's performance. You can access the console in a non-shippable build or the **Editor**. A non-shippable build refers to a project that was built using a debug or development configuration.

To access the console

In non-shippable builds and the Play In Editor mode, press the tilde (**~**) key to open the console. Double-press the tilde key to expand the console.

Here are some tips for using the console:

- Type a keyword to list all possible commands containing that keyword. Scroll through the list using the arrow keys.
- Scroll through the history by using the arrow keys or Page up and Page down keys.
- Logs are saved in a `.txt` file in your project's `Saved/Logs` directory

To profile your game's performance

1. Start by running the `stat fps` and `stat unit` commands. This will give you an overview of where your game struggles with performance.
   - `stat fps`: Shows the current frames per second.
   - `stat unit`: Breaks down the frame into several subsections.
     - **Frame**: Total wall-clock time starting from when the simulation of the frame starts to when the presentation of the frame is on the screen.
     - **Game**: Total CPU time taken by the game simulation thread per frame.
     - **Draw**: Total CPU time for the rendering threads to translate the scene to commands for the GPU and submit them to the GPU.
     - **GPU**: Total time for the GPU to process all commands.
     - **Draws**: Total number of draws submitted for the frame.
     - **Prims**: Total number of triangles drawn.
2. Play through the game and identify areas with low performance, indicated by decreased FPS and increased time in **Game**, **Draw**, or **GPU**.
3. Run `stat game` to see how time is spent for the various gameplay groups.
4. Refine the stats for specific gameplay factors like AI, animation, physics, gameplay, scripting, and so on. Here are a few examples:

- `stat ai`: Time to compute AI behavior.
- `stat anim`: Time to compute skinned meshes.
- `stat physics`: Time to compute physics simulations.

5. Run `stat drawcount` to see which render areas generate the most draws. The list shows the render passes that emit draws, and the number of draws emitted each frame. You can get more information by analyzing the GPU stats in the next step.
6. Run `stat gpu` to see which render types consume the most GPU time.
7. Refine the rendering types into broad groups, such as lights, shadows, lumen (lighting), hair, post processing, and so on. Here are a few common examples:

   - `stat lightrendering`: GPU time to render lights and shadows.
   - `stat shadowrendering`: GPU time to update the various shadows.
   - `stat scenerendering`: GPU time to render the scene.

This section covers only a subset of available commands. Depending on your game's features, look into stats for areas such as asset streaming, virtual texturing, CPU task workload distribution, threading, sound, particles, and so on. For more information, refer to Stat commands.

Regions, quotas, and limitations

Amazon GameLift Streams is available across multiple AWS Regions, offering dual-stack service endpoints that support both IPv4 and IPv6 connectivity. The service operates from primary locations including US East (Ohio), US West (Oregon), Asia Pacific (Tokyo), and Europe (Frankfurt), with the ability to manage additional AWS Regions and locations, collectively referred to as *remote locations*, for optimized latency and stream quality.

The service infrastructure is governed by three main categories of constraints:

- Service quotas
- API rate limits
- Fixed service limitations

These include restrictions on application sizes, number of applications per region, file management capacities, and GPU allocations across different stream classes and regions. The service implements specific API rate limits for various operations, ranging from 1 to 20 requests per second, ensuring stable service performance. Additionally, there are fixed service limitations concerning stream group configurations, GPU deployments, and application associations that apply uniformly across all customers.

AWS Regions and streaming locations supported by Amazon GameLift Streams

An AWS Region is a collection of AWS resources in a geographic area. Each AWS Region is isolated and independent of the other Regions. For general information about AWS Regions, see Managing AWS Regions in the *AWS General Reference*.

The following table lists the AWS Regions where the Amazon GameLift Streams service is available and the endpoints for each Region. You create all Amazon GameLift Streams application and stream group resources in a specified Region, whether you work in the Amazon GameLift Streams console, use the AWS Command Line Interface (AWS CLI), or make programmatic calls. The Region where you create these resources is known as the *primary location*. Use your primary location's endpoint to connect to the Amazon GameLift Streams service programmatically.

Service endpoints

Amazon GameLift Streams supports dual-stack service endpoints, allowing clients and resources to interact with the service using IPv6 or IPv4.

Region Name	Region	Endpoint	Protocol	
US East (Ohio)	us-east-2	gameliftstreams.us-east-2.api.aws	HTTPS	
US West (Oregon)	us-west-2	gameliftstreams.us-west-2.api.aws	HTTPS	
Asia Pacific (Tokyo)	ap-northeast-1	gameliftstreams.ap-northeast-1.api.aws	HTTPS	
Europe (Frankfurt)	eu-central-1	gameliftstreams.eu-central-1.api.aws	HTTPS	

Streaming locations

Amazon GameLift Streams supports streaming from all the following locations from any of the service endpoints. We recommend that you choose streaming locations that are geographically close to your users to optimize latency and stream quality.

Region name	AWS Region		
US East (N. Virginia)	us-east-1		
US East (Ohio)	us-east-2		
US West (Oregon)	us-west-2		
Asia Pacific (Mumbai)	ap-south-1		

Region name	AWS Region		
Asia Pacific (Seoul)	ap-northeast-2		
Asia Pacific (Sydney)	ap-southeast-2		
Asia Pacific (Tokyo)	ap-northeast-1		
Europe (Frankfurt)	eu-central-1		
Europe (Ireland)	eu-west-1		
Europe (London)	eu-west-2		
Europe (Stockholm)	eu-north-1		
South America (São Paulo)	sa-east-1		

Supported locations by stream class in Amazon GameLift Streams

The following table shows the availability of each stream class family across all supported AWS Regions and streaming locations.

Region name	Region	gen6*	gen5*	gen4*
Asia Pacific (Mumbai)	ap-south-1	✓ Yes	✓ Yes	✓ Yes
Asia Pacific (Seoul)	ap-northeast-2	✓ Yes	✓ Yes	✓ Yes
Asia Pacific (Sydney)	ap-southeast-2	✓ Yes	✓ Yes	✓ Yes
Asia Pacific (Tokyo)	ap-northeast-1	✓ Yes	✓ Yes	✓ Yes
Europe (Frankfurt)	eu-central-1	✓ Yes	✓ Yes	✓ Yes
Europe (Ireland)	eu-west-1	✗ No	✓ Yes	✓ Yes

Region name	Region	gen6*	gen5*	gen4*
Europe (London)	eu-west-2	✓ Yes	✓ Yes	✓ Yes
Europe (Stockholm)	eu-north-1	✓ Yes	✓ Yes	✓ Yes
South America (São Paulo)	sa-east-1	✓ Yes	✓ Yes	✓ Yes
US East (N. Virginia)	us-east-1	✓ Yes	✓ Yes	✓ Yes
US East (Ohio)	us-east-2	✓ Yes	✓ Yes	✓ Yes
US West (Oregon)	us-west-2	✓ Yes	✓ Yes	✓ Yes

Amazon GameLift Streams service quotas

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account.

Many of the service quotas in Amazon GameLift Streams limit the total number of GPUs (compute resources) you can configure for streaming in your account. More specifically, these GPU service quotas specify the maximum number of GPUs of a particular stream class family you can request per location across all stream groups in your account. For example, if your account has a limit of 5 `gen5n` GPUs in `us-west-2`, the sum of `gen5n` GPUs needed to provide the total stream capacity in `us-west-2` for all of your stream groups must be less than or equal to 5. This includes GPUs for both always-on and on-demand capacity.

For more information about how quotas interact with stream capacity, see Capacity and service quotas. Also be sure to check API rate limits and Other limitations for additional limitations to be aware of in Amazon GameLift Streams.

View your default or applied account level quota and the utilization in the Service Quotas console by selecting the GameLift Streams as the AWS service.

For general information about service quotas, see AWS service quotas in the *AWS General Reference*.

Service quotas

In the following table, the GPU quotas are all 0 by default. However, your account's applied quotas might be different. To check, sign in to the AWS Management Console and open the Service Quotas console to Amazon GameLift Streams, where you can review your current quotas in the **Applied account-level quota value** column and utilization of these quotas in the Utilization column and submit a request to increase these values.

Name	Default	Adjustable	Description
Application size (GiB)	Each supported Region: 100	Yes	The maximum total size (in GiB) of an application, in this account. Note that a gibibyte (GiB) equals 1024*1024*1024 bytes.
Applications	Each supported Region: 20	Yes	The maximum number of applications that you can create in this account, per AWS Region.
Files per application	Each supported Region: 30,000	Yes	The maximum number of files that you can have in an application, in this account.
Gen4n GPUs, ap-northeast-1	Each supported Region: 0	Yes	The maximum number of Gen4n GPUs you can configure for streaming in the ap-northeast-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support

Name	Default	Adjustable	Description
			streaming more than one session per GPU.
Gen4n GPUs, ap-northeast-2	Each supported Region: 0	Yes	The maximum number of Gen4n GPUs you can configure for streaming in the ap-northeast-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU.
Gen4n GPUs, ap-south-1	Each supported Region: 0	Yes	The maximum number of Gen4n GPUs you can configure for streaming in the ap-south-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU.

Name	Default	Adjustable	Description
Gen4n GPUs, ap-southeast-2	Each supported Region: 0	Yes	The maximum number of Gen4n GPUs you can configure for streaming in the ap-southeast-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU.
Gen4n GPUs, eu-central-1	Each supported Region: 0	Yes	The maximum number of Gen4n GPUs you can configure for streaming in the eu-central-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU.

Name	Default	Adjustable	Description
Gen4n GPUs, eu-north-1	Each supported Region: 0	Yes	The maximum number of Gen4n GPUs you can configure for streaming in the eu-north-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU.
Gen4n GPUs, eu-west-1	Each supported Region: 0	Yes	The maximum number of Gen4n GPUs you can configure for streaming in the eu-west-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU.
Gen4n GPUs, eu-west-2	Each supported Region: 0	Yes	The maximum number of Gen4n GPUs you can configure for streaming in the eu-west-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU.

Name	Default	Adjustable	Description
Gen4n GPUs, sa-east-1	Each supported Region: 0	Yes	The maximum number of Gen4n GPUs you can configure for streaming in the sa-east-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU.
Gen4n GPUs, us-east-1	Each supported Region: 0	Yes	The maximum number of Gen4n GPUs you can configure for streaming in the us-east-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU.
Gen4n GPUs, us-east-2	Each supported Region: 0	Yes	The maximum number of Gen4n GPUs you can configure for streaming in the us-east-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU.

Name	Default	Adjustable	Description
Gen4n GPUs, us-west-2	Each supported Region: 0	Yes	The maximum number of Gen4n GPUs you can configure for streaming in the us-west-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU.
Gen5n GPUs, ap-northeast-1	Each supported Region: 0	Yes	The maximum number of Gen5n GPUs you can configure for streaming in the ap-northeast-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU.
Gen5n GPUs, ap-northeast-2	Each supported Region: 0	Yes	The maximum number of Gen5n GPUs you can configure for streaming in the ap-northeast-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU.

Name	Default	Adjustable	Description
Gen5n GPUs, ap-south-1	Each supported Region: 0	Yes	The maximum number of Gen5n GPUs you can configure for streaming in the ap-south-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU.
Gen5n GPUs, ap-southeast-2	Each supported Region: 0	Yes	The maximum number of Gen5n GPUs you can configure for streaming in the ap-southeast-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU.

Name	Default	Adjustable	Description
Gen5n GPUs, eu-central-1	Each supported Region: 0	Yes	The maximum number of Gen5n GPUs you can configure for streaming in the eu-central-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU.
Gen5n GPUs, eu-north-1	Each supported Region: 0	Yes	The maximum number of Gen5n GPUs you can configure for streaming in the eu-north-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU.
Gen5n GPUs, eu-west-1	Each supported Region: 0	Yes	The maximum number of Gen5n GPUs you can configure for streaming in the eu-west-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU.

Name	Default	Adjustable	Description
Gen5n GPUs, eu-west-2	Each supported Region: 0	Yes	The maximum number of Gen5n GPUs you can configure for streaming in the eu-west-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU.
Gen5n GPUs, sa-east-1	Each supported Region: 0	Yes	The maximum number of Gen5n GPUs you can configure for streaming in the sa-east-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU.
Gen5n GPUs, us-east-1	Each supported Region: 0	Yes	The maximum number of Gen5n GPUs you can configure for streaming in the us-east-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU.

Name	Default	Adjustable	Description
Gen5n GPUs, us-east-2	Each supported Region: 0	Yes	The maximum number of Gen5n GPUs you can configure for streaming in the us-east-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU.
Gen5n GPUs, us-west-2	Each supported Region: 0	Yes	The maximum number of Gen5n GPUs you can configure for streaming in the us-west-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU.
Gen6n GPUs, ap-northeast-1	Each supported Region: 0	Yes	The maximum number of Gen6n GPUs you can configure for streaming in the ap-northeast-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU.

Name	Default	Adjustable	Description
Gen6n GPUs, ap-northeast-2	Each supported Region: 0	Yes	The maximum number of Gen6n GPUs you can configure for streaming in the ap-northeast-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU.
Gen6n GPUs, ap-south-1	Each supported Region: 0	Yes	The maximum number of Gen6n GPUs you can configure for streaming in the ap-south-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU.

Name	Default	Adjustable	Description
Gen6n GPUs, ap-southeast-2	Each supported Region: 0	Yes	The maximum number of Gen6n GPUs you can configure for streaming in the ap-southeast-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU.
Gen6n GPUs, eu-central-1	Each supported Region: 0	Yes	The maximum number of Gen6n GPUs you can configure for streaming in the eu-central-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU.

Name	Default	Adjustable	Description
Gen6n GPUs, eu-north-1	Each supported Region: 0	Yes	The maximum number of Gen6n GPUs you can configure for streaming in the eu-north-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU.
Gen6n GPUs, eu-west-2	Each supported Region: 0	Yes	The maximum number of Gen6n GPUs you can configure for streaming in the eu-west-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU.
Gen6n GPUs, sa-east-1	Each supported Region: 0	Yes	The maximum number of Gen6n GPUs you can configure for streaming in the sa-east-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU.

Name	Default	Adjustable	Description
Gen6n GPUs, us-east-1	Each supported Region: 0	Yes	The maximum number of Gen6n GPUs you can configure for streaming in the us-east-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU.
Gen6n GPUs, us-east-2	Each supported Region: 0	Yes	The maximum number of Gen6n GPUs you can configure for streaming in the us-east-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU.
Gen6n GPUs, us-west-2	Each supported Region: 0	Yes	The maximum number of Gen6n GPUs you can configure for streaming in the us-west-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU.

Name	Default	Adjustable	Description
Stream groups	Each supported Region: 5	Yes	The maximum number of stream groups that you can create in this account, per AWS Region. A stream group is a collection of compute resources that stream your application to end users.

Amazon GameLift Streams API rate limits

These limits reflect the maximum rate of requests per second from your AWS account to the Amazon GameLift Streams service in an AWS Region.

API operation	Requests per second
AddStreamGroupLocations	5
AssociateApplications	5
CreateApplication	5
CreateStreamGroup	1
CreateStreamSessionConnection	20
DeleteApplication	5
DeleteStreamGroup	5
DisassociateApplications	5
ExportStreamSessionFiles	20

API operation	Requests per second
GetApplication	10
GetStreamGroup	10
GetStreamSession	20
ListApplications	10
ListStreamGroups	10
ListStreamSessions	20
ListStreamSessionsByAccount	20
ListTagsForResource	10
RemoveStreamGroupLocations	5
StartStreamSession	20
TagResource	10
TerminateStreamSession	20
UntagResource	10
UpdateApplication	5
UpdateStreamGroup	5

Other Amazon GameLift Streams limitations

This page lists other limitations to be aware of as you create your streaming solution. These limits are fixed within the service for all customers.

Name	Limitation	Description
Applications in a stream group	100	The maximum number of Amazon GameLift Streams applications that can be associated to a stream group.
GPUs in a stream group	2500	The maximum number of GPUs in a stream group across all Regions and remote locations.
Single file size (GiB)	80 GiB	The maximum size (in GiB) of a single file in an application. Note that a gibibyte (GiB) equals 1024*1024*1024 bytes.
Stream group associations per application	100	The maximum number of stream groups that an Amazon GameLift Streams application can be associated to.

Managing usage and bills for Amazon GameLift Streams

This topic covers how to monitor and manage your Amazon GameLift Streams usage, costs, and billing to optimize your streaming expenses.

Also see the Amazon GameLift Streams Pricing page for the following information:

- **Cost breakdown**: Understand what AWS charges you for when you use Amazon GameLift Streams.
- **Amazon GameLift Streams rates**: See how much Amazon GameLift Streams costs and compare different options.
- **Stream capacity reservation**: Plan ahead and ensure that you have enough stream capacity to meet your customer demands.

Review your Amazon GameLift Streams bills and usage

You can review your Amazon GameLift Streams bills and usage by using the AWS Billing and Cost Management tools in the AWS Console or AWS CLI.

To view your bill through the AWS Console, refer to Viewing your bill in the AWS Billing User Guide.

To view your bill through the AWS CLI, call `GetCostAndUsage` using the Billing and Cost Management API. For example, use the following command to retrieve a monthly bill for Amazon GameLift Streams, and replace the dates with ones relevant to you.

Example : Use `GetCostAndUsage` API to view bill

```
aws ce get-cost-and-usage /
    --time-period Start=2023-01-01,End=2023-01-31 /
    --granularity MONTHLY /
    --metrics BlendedCost /
    --filter Amazon GameLift Streams-bill-filter.json
```

where the filter, such as `Amazon GameLift Streams-bill-filter.json`, specifies the Amazon GameLift Streams service as follows:

```
{
    "Dimensions": {
        "Key": "SERVICE",
        "Values": ["Amazon Amazon GameLift Streams"]
    }
}
```

Best practices to manage Amazon GameLift Streams costs

We strongly recommend that you use the following tools and techniques to manage your Amazon GameLift Streams costs to avoid unexpected costs.

Create billing alerts to monitor usage

Set up billing alerts using AWS Budgets, which enables you to track your costs and usage, and respond quickly to alerts to avoid unexpected costs. You can also configure the billing alert to trigger actions that help you stay within budget. By default, budgets include all of your AWS services. To specify a budget for Amazon GameLift Streams only, add a budget filter.

For more information, see the following topics:

- Creating a budget
- Best practices for AWS Budgets

Scale stream groups to zero capacity

Allocated stream capacity continues to incur costs even when they're not currently hosting stream sessions. Scale stream groups to zero capacity when not in use to avoid unnecessary cost. This prevents your stream group from allocating resources. When you set always-on and on-demand stream capacity to zero, all connected streams end. When you're ready, you can reuse your stream group by scaling capacity back up.

For instructions, refer to Edit capacity.

> **Warning**
>
> Avoid deleting a stream group, unless you don't plan to use the stream group again. If you delete a stream group, you cannot restore the original stream group and must create a new one.

Delete original application files

To optimize storage cost, you can delete the original application files that you uploaded to an Amazon S3 bucket. It's safe to delete the files if the application is in **Ready** status. At that point, Amazon GameLift Streams has a snapshot of the application files and no longer accesses your original files.