

# Kinesis Video Streams



# **Kinesis Video Streams: Amazon Kinesis Video Streams WebRTC Developer Guide**

# Table of Contents

<b>What Is Amazon Kinesis Video Streams with WebRTC .....</b>	<b>1</b>
Region availability .....	2
Kinesis Video Streams with WebRTC Pricing .....	3
Accessing Kinesis Video Streams with WebRTC .....	3
<b>Kinesis Video Streams with WebRTC: How It Works .....</b>	<b>4</b>
Amazon Kinesis Video Streams with WebRTC Concepts .....	4
WebRTC Technology Concepts .....	5
How STUN, TURN and ICE Work Together .....	6
Kinesis Video Streams with WebRTC Components .....	7
WebRTC Websocket APIs .....	8
ConnectAsViewer .....	8
ConnectAsMaster .....	11
SendSdpOffer .....	12
SendSdpAnswer .....	15
SendIceCandidate .....	17
Disconnect .....	19
Asynchronous Message Reception .....	20
<b>Quotas .....</b>	<b>22</b>
Control Plane API Service Quotas .....	22
Signaling API Service Quotas .....	24
TURN Service Quotas .....	25
WebRTC Ingestion Service Quotas .....	26
<b>Getting Started .....</b>	<b>27</b>
Set Up an Amazon Web Services account .....	27
Sign up for an Amazon Web Services account .....	27
Secure IAM users .....	28
Create an AmazonAccount Key .....	28
Create a Signaling Channel .....	28
Create a Signaling Channel Using the Console .....	29
Stream Live Media .....	29
WebRTC SDK in C for Embedded Devices .....	29
WebRTC SDK in JavaScript .....	32
WebRTC SDK for Android .....	37
WebRTC SDK for iOS .....	44

---

Client Metrics for the WebRTC C SDK .....	51
<b>Security .....</b>	<b>74</b>
Controlling Access to Kinesis Video Streams with WebRTC Resources Using Amazon Identity and Access Management .....	75
Policy Syntax .....	75
Actions for Kinesis Video Streams with WebRTC .....	76
Amazon Resource Names (ARNs) for Kinesis Video Streams .....	77
Granting Other IAM Accounts Access to a Kinesis Video Stream .....	77
Example Policies .....	77
Compliance validation .....	79
Resilience .....	80
Infrastructure Security in Kinesis Video Streams with WebRTC .....	80
Security Best Practices for Kinesis Video Streams with WebRTC .....	80
Implement least privilege access .....	81
Use IAM roles .....	81
Use CloudTrail to Monitor API Calls .....	81
<b>Monitoring .....</b>	<b>83</b>
Monitoring Kinesis Video Streams with WebRTC Metrics with CloudWatch .....	83
Signaling Metrics .....	84
TURN Metrics .....	85
Logging Kinesis Video Streams with WebRTC API Calls with Amazon CloudTrail .....	85
Amazon Kinesis Video Streams with WebRTC and CloudTrail .....	85
Example: Amazon Kinesis Video Streams with WebRTC Log File Entries .....	87
<b>Document History .....</b>	<b>89</b>
<b>Amazon Glossary .....</b>	<b>90</b>

# What Is Amazon Kinesis Video Streams with WebRTC

WebRTC is an open technology specification for enabling real-time communication (RTC) across browsers and mobile applications via simple APIs. It uses peering techniques for real-time data exchange between connected peers and provides low latency media streaming required for human-to-human interaction. The WebRTC specification includes a set of IETF protocols including [Interactive Connectivity Establishment](#), [Traversal Using Relay around NAT \(TURN\)](#), and [Session Traversal Utilities for NAT \(STUN\)](#) for establishing peer-to-peer connectivity, in addition to protocol specifications for reliable and secure real-time media and data streaming.

[Amazon Kinesis Video Streams](#) provides a standards-compliant WebRTC implementation as a fully managed capability. You can use Amazon Kinesis Video Streams with WebRTC to securely live stream media or perform two-way audio or video interaction between any camera IoT device and WebRTC-compliant mobile or web players. As a fully managed capability, you don't have to build, operate, or scale any WebRTC-related cloud infrastructure, such as signaling or media relay servers to securely stream media across applications and devices.

Using Kinesis Video Streams with WebRTC, you can easily build applications for live peer-to-peer media streaming, or real-time audio or video interactivity between camera IoT devices, web browsers, and mobile devices for a variety of use cases. Such applications can help parents keep an eye on their baby's room, enable homeowners to use a video doorbell to check who's at the door, enable owners of camera-enabled robot vacuums to remotely control the robot by viewing the live camera stream on a mobile phone, and so on.

If you're a first-time user of Kinesis Video Streams with WebRTC, we recommend that you read the following sections:

- [Kinesis Video Streams with WebRTC: How It Works](#)
- [WebRTC SDK in C for Embedded Devices](#)
- [Kinesis Video Streams with WebRTC SDK in JavaScript for web applications](#)
- [WebRTC SDK for Android](#)
- [WebRTC SDK for iOS](#)
- [Control plane APIs](#)
- [Data plane REST APIs](#)
- [Data plane Websocket APIs](#)

## Region availability

Amazon Kinesis Video Streams with WebRTC is available in the following regions:

Region Name	Amazon Region Code
US East (Ohio)	us-east-2
US East (N. Virginia)	us-east-1
US West (Oregon)	us-west-2
Africa (Cape Town)	af-south-1
Asia Pacific (Hong Kong)	ap-east-1
Asia Pacific (Mumbai)	ap-south-1
Asia Pacific (Seoul)	ap-northeast-2
Asia Pacific (Singapore)	ap-southeast-1
Asia Pacific (Sydney)	ap-southeast-2
Asia Pacific (Tokyo)	ap-northeast-1
Canada (Central)	ca-central-1
Europe (Frankfurt)	eu-central-1
Europe (Ireland)	eu-west-1
Europe (London)	eu-west-2
Europe (Paris)	eu-west-3
South America (Sao Paulo)	sa-east-1

# Kinesis Video Streams with WebRTC Pricing

For information about Kinesis Video Streams with WebRTC pricing, see [Amazon Kinesis Video Streams Pricing](#).

## Accessing Kinesis Video Streams with WebRTC

You can work with Kinesis Video Streams with WebRTC in any of the following ways:

### Amazon Web Services Management Console

#### [Getting Started with the Amazon Web Services Management Console](#)

The console is a browser-based interface to access and use Amazon services, including Kinesis Video Streams with WebRTC.

### Amazon SDKs

Amazon provides software development kits (SDKs) that consist of libraries and sample code for various programming languages and platforms (for example, Java, Python, Ruby, .NET, iOS, Android, and more). The SDKs provide a convenient way to create programmatic access to Kinesis Video Streams with WebRTC. For information about the Amazon SDKs, including how to download and install them, see [Tools for Amazon Web Services](#).

### Kinesis Video Streams with WebRTC HTTPS API

You can access Kinesis Video Streams with WebRTC and Amazon programmatically by using the Kinesis Video Streams with WebRTC APIs, which lets you issue API requests directly to the service. For more information, see the [Amazon Kinesis Video Streams API Reference](#).

# Kinesis Video Streams with WebRTC: How It Works

## Topics

- [Amazon Kinesis Video Streams with WebRTC Concepts](#)
- [WebRTC Technology Concepts](#)
- [How STUN, TURN and ICE Work Together](#)
- [Kinesis Video Streams with WebRTC Components](#)
- [WebRTC Websocket APIs](#)

## Amazon Kinesis Video Streams with WebRTC Concepts

The following are key terms and concepts specific to the Amazon Kinesis Video Streams with WebRTC.

### Signaling channel

A resource that enables applications to discover, set up, control, and terminate a peer-to-peer connection by exchanging signaling messages. Signaling messages are metadata that two applications exchange with each other to establish peer-to-peer connectivity. This metadata includes local media information, such as media codecs and codec parameters, and possible network candidate paths for the two applications to connect with each other for live streaming.

Streaming applications can maintain persistent connectivity with a signaling channel and wait for other applications to connect to them. Or, they can connect to a signaling channel only when they need to live stream media. A signaling channel enables applications to connect with each other in a one-to-few model, using the concept of one master connecting to multiple viewers. The application that initiates the connection assumes the responsibility of a master using the `ConnectAsMaster` API and waits for viewers. Up to 10 applications can then connect to that signaling channel by assuming the viewer responsibility by invoking the `ConnectAsViewer` API. After they're connected to a signaling channel, the master and viewer applications can send each other signaling messages to establish peer-to-peer connectivity for live media streaming.



## Peer

Any device or application (for example, a mobile or web application, webcam, home security camera, baby monitor, etc.) that is configured for real-time, two-way streaming through a Kinesis Video Streams with WebRTC.

## Master

A peer that initiates the connection and is connected to the signaling channel with the ability to discover and exchange media with any of the signaling channel's connected viewers.

### Important

Currently, a signaling channel can only have one master.

## Viewer

A peer that is connected to the signaling channel with the ability to discover and exchange media only with the signaling channel's master. A viewer cannot discover or interact with other viewers through a given signaling channel. A signaling channel can have up to 10 connected viewers.

# WebRTC Technology Concepts

As you get started with Kinesis Video Streams with WebRTC, you can also benefit from learning about several interrelated protocols and APIs of which the WebRTC technology consists.

## Session Traversal Utilities for NAT (STUN)

A protocol that is used to discover your public address and determine any restrictions in your router that would prevent a direct connection with a peer.

## Traversal Using Relays around NAT (TURN)

A server that is used to bypass the Symmetric NAT restriction by opening a connection with a TURN server and relaying all information through that server.

## Session Description Protocol (SDP)

A standard for describing the multimedia content of the connection such as resolution, formats, codecs, encryption, etc. so that both peers can understand each other once the data is transferring.

### SDP Offer

An SDP message sent by an agent which generates a session description in order to create or modify a session. It describes the aspects of desired media communication.

### SDP Answer

An SDP message sent by an answerer in response to an offer received from an offerer. The answer indicates the aspects that are accepted. For example, if all the audio and video streams in the offer are accepted.

## Interactive Connectivity Establishment (ICE)

A framework that allows your web browser to connect with peers.

### ICE Candidate

A method that the sending peer is able to use to communicate.

## How STUN, TURN and ICE Work Together

Let's take the scenario of two peers, A and B, who are both using a WebRTC peer to peer two way media streaming (for example, a video chat application). What happens when A wants to call B?

To connect to B's application, A's application must generate an SDP offer. An SDP offer contains information about the session A's application wants to establish, including what codecs to use, whether this is an audio or video session, etc. It also contains a list of ICE candidates, which are the IP and port pairs that B's application can attempt to use to connect to A.

To build the list of ICE candidates, A's application makes a series of requests to a STUN server. The server returns the public IP address and port pair that originated the request. A's application adds each pair to the list of ICE candidates, in other words, it gathers ICE candidates. Once A's application has finished gathering ICE candidates, it can return an SDP.

Next, A's application must pass the SDP to B's application through a signaling channel over which these applications communicate. The transport protocol for this exchange is not specified

in the WebRTC standard. It can be performed over HTTPS, secure WebSocket, or any other communication protocol.

Now, B's application must generate an SDP answer. B's application follows the same steps A used in the previous step: gathers ICE candidates, etc. B's application then needs to return this SDP answer to A's application.

After A and B have exchanged SDPs, they then perform a series of connectivity checks. The ICE algorithm in each application takes a candidate IP/port pair from the list it received in the other party's SDP, and sends it a STUN request. If a response comes back from the other application, the originating application considers the check successful and marks that IP/port pair as a valid ICE candidate.

After connectivity checks are finished on all of the IP/port pairs, the applications negotiate and decide to use one of the remaining, valid pairs. When a pair is selected, media begins flowing between the application.

If either of the applications can't find an IP/port pair that passes connectivity checks, they'll make STUN requests to the TURN server to obtain a media relay address. A relay address is a public IP address and port that forwards packets received to and from the application to set up the relay address. This relay address is then added to the candidate list and exchanged via the signaling channel.

## Kinesis Video Streams with WebRTC Components

Kinesis Video Streams with WebRTC includes the following components:

- **Control plane**

The control plane component is responsible for creating and maintaining the Kinesis Video Streams with WebRTC signaling channels. For more information, see the [Amazon Kinesis Video Streams API Reference](#).

- **Signaling**

The signaling component manages the WebRTC signaling endpoints that allow applications to securely connect with each other for peer-to-peer live media streaming. The signaling component includes the [Amazon Kinesis Video Signaling REST APIs](#) and a set of [Websocket APIs](#).

- **STUN**

This component manages STUN endpoints that enable applications to discover their public IP address when they are located behind a NAT or a firewall.

- **TURN**

This component manages TURN endpoints that enable media relay via the cloud when applications can't stream media peer-to-peer.

- **Kinesis Video Streams WebRTC SDKs**

These are software libraries that you can download, install, and configure on your devices and application clients to enable your camera IoT devices with WebRTC capabilities to engage in low latency peer-to-peer media streaming. These SDKs also enable Android, iOS, and web application clients to integrate Kinesis Video Streams with WebRTC signaling, TURN, and STUN capabilities with any WebRTC-compliant mobile or web players.

- [WebRTC SDK in C for Embedded Devices](#)
- [Kinesis Video Streams with WebRTC SDK in JavaScript for web applications](#)
- [WebRTC SDK for Android](#)
- [WebRTC SDK for iOS](#)

## WebRTC Websocket APIs

### Topics

- [ConnectAsViewer](#)
- [ConnectAsMaster](#)
- [SendSdpOffer](#)
- [SendSdpAnswer](#)
- [SendIceCandidate](#)
- [Disconnect](#)
- [Asynchronous Message Reception](#)

## ConnectAsViewer

Connects as a viewer to the signaling channel specified by the endpoint. Any WebSocket-compliant library can be used to connect to the endpoint obtained from the `GetSignalingEndpoint`

API call. The Amazon Resource Name (ARN) of the signaling channel and the client ID must be provided as query string parameters. There are separate endpoints for connecting as a master and as a viewer. If there is an existing connection with the same `ClientId` as specified in the request, the new connection takes precedence. The connection metadata is overwritten with the new information.

## Request

```
"X-Amz-ChannelARN": "string",  
"X-Amz-ClientId": "string"
```

- **X-Amz-ChannelARN** - ARN of the signaling channel.
  - Type: string
  - Length constraints: Minimum length of 1. Maximum length of 1024
  - Pattern: `arn:aws:kinesisvideo:[a-z0-9-]+:[0-9]+:[a-z]+/[a-zA-Z0-9_.-]+/[0-9]+`
  - Required: Yes
- **X-Amz-ClientId** - A unique identifier for the client.
  - Type: string
  - Length constraints: Minimum length of 1. Maximum length of 256.
  - Pattern: `^(?!(?i)AWS_.*)[a-zA-Z0-9_.-]`

### Note

`X-Amz-ClientId` can't start with `AWS_`.

- Required: Yes

## Response

200 OK HTTP status code with an empty body.

## Errors

- `InvalidArgumentException`

A specified parameter exceeds its restrictions, is not supported, or cannot be used. For more information, see the returned message.

HTTP Status Code: 400

- `AccessDeniedException`

The caller is not authorized to access the given channel or the token has expired.

HTTP Status Code: 403

- `ResourceNotFoundException`

The channel doesn't exist.

HTTP Status Code: 404

- `ClientLimitExceededException`

When the API is invoked at a rate that is too high or when there are more than the supported maximum number of viewers connected to the channel. For more information, see [Amazon Kinesis Video Streams with WebRTC Service Quotas](#) and [Error Retries and Exponential Backoff in Amazon](#).

HTTP Status Code: 400

## Limits/Throttling

This API is throttled at an account level if the API is invoked at too high a rate or when there are more than the supported maximum number of viewers connected to the channel. An error returned when throttled with `ClientLimitExceededException`.

## Idempotent

If a connection already exists for the specified `ClientId` and channel, the connection metadata is updated with the new information.

## Retry behavior

This is counted as a new API call.

## Concurrent calls

Concurrent calls are allowed, the connection metadata is updated for each call.

## ConnectAsMaster

Connects as a master to the signaling channel specified by the endpoint. Any WebSocket-complaint library can be used to connect to the endpoint obtained from a `GetSignalingChannelEndpoint` API call. The Amazon Resource Name (ARN) of the signaling channel must be provided as a query string parameter. There are separate endpoints for connecting as a master and as a viewer. If more than one client connects as master to a specific channel, then the most recent request takes precedence. Existing connection metadata is overwritten by the new one.

### Request

```
"X-Amz-ChannelARN": "string"
```

- **X-Amz-ChannelARN** - ARN of the signaling channel.
  - Type: string
  - Length constraints: Minimum length of 1. Maximum length of 1024.
  - Pattern: `arn:aws:kinesisvideo:[a-z0-9-]+:[0-9]+:[a-z]+/[a-zA-Z0-9_.-]+/[0-9]+`
  - Required: Yes

### Response

200 OK HTTP status code with an empty body.

### Errors

- `InvalidArgumentException`

A specified parameter exceeds its restrictions, is not supported, or cannot be used. For more information, see the returned message.

HTTP Status Code: 400

- `AccessDeniedException`

The caller is not authorized to access the given channel or the token has expired.

HTTP Status Code: 403

- `ResourceNotFoundException`

The channel doesn't exist.

HTTP Status Code: 404

- `ClientLimitExceededException`

When the API is invoked at a rate that is too high. For more information, see [Amazon Kinesis Video Streams with WebRTC Service Quotas](#) and [Error Retries and Exponential Backoff in Amazon](#).

HTTP Status Code: 400

## Limits/Throttling

This API is throttled at an account level if the API is invoked at too high a rate. An error returned when throttled with `ClientLimitExceededException`.

## Idempotent

If a connection already exists for the specified `clientId` and channel, the connection metadata is updated with the new information.

## Retry behavior

This is counted as a new API call.

## Concurrent calls

Concurrent calls are allowed, the connection metadata is updated for each call.

## SendSdpOffer

Sends the offer to the target recipient. The prerequisite is that the client must be already connected to the WebSocket endpoint obtained from the `GetSignalingChannelEndpoint` API.



If the sender type is a viewer, then it sends the offer to a master. Also, it is not necessary to specify the `RecipientClientId` and any specified value for `RecipientClientId` is ignored. If the sender type is master, the offer is sent to the target viewer specified by the `RecipientClientId`. `RecipientClientId` is a required input in this case.

A master client app is allowed to send an offer to any viewer, whereas a viewer client app is only allowed to send an offer to a master client app. If a viewer client app attempts to send an offer to another viewer client app, the request will NOT be honored. If there is an outstanding offer for the same client which is not yet delivered, it is overwritten with the new offer.

## Request

```
{
  "action": "SDP_OFFER",
  "recipientClientId": "string",
  "messagePayload": "string",
  "correlationId": "string"
}
```

- **action** - Type of the message that is being sent.
  - Type: ENUM
  - Valid values: SDP\_OFFER, SDP\_ANSWER, ICE\_CANDIDATE
  - Length constraints: Minimum length of 1. Maximum length of 256.
  - Pattern: [a-zA-Z0-9\_.-]+
  - Required: Yes
- **recipientClientId** - The unique identifier for the recipient.
  - Type: String
  - Length constraints: Minimum length of 1. Maximum length of 256.
  - Pattern: [a-zA-Z0-9\_.-]+
  - Required: Yes
- **messagePayload** - The base-64-encoded message content.
  - Type: String
  - Length constraints: Minimum length of 1. Maximum length of 10K.
  - Required: Yes
- **correlationId** - A unique identifier for the message. This is an optional parameter.

- Type: String
- Length constraints: Minimum length of 1. Maximum length of 256.
- Pattern: [a-zA-Z0-9\_.-]+
- Required: No

## Response

If the message is successfully received by the signaling backend, no response is returned. If the service encounters an error and if the `correlationId` is specified in the request, the error details are returned as a `STATUS_RESPONSE` message. For more information, see [Asynchronous Message Reception](#).

## Errors

- `InvalidArgumentException`

A specified parameter exceeds its restrictions, is not supported, or cannot be used. For more information, see the returned message.

HTTP Status Code: 400

- `ClientLimitExceededException`

When the API is invoked at a rate that is too high. For more information, see [Amazon Kinesis Video Streams with WebRTC Service Quotas](#) and [Error Retries and Exponential Backoff in Amazon](#).

HTTP Status Code: 400

## Limits/Throttling

This API is throttled at an account level if the API is invoked at too high a rate. An error returned when throttled with `ClientLimitExceededException`.

## Idempotent

This API is not idempotent.

## Retry behavior

This is counted as a new API call.

## Concurrent calls

Concurrent calls are allowed. An offer is sent once per each call.

## SendSdpAnswer

Sends the answer to the target recipient. The prerequisite is that the client must be already connected to the WebSocket endpoint obtained from the `GetSignalingChannelEndpoint` API.

If the sender type is a viewer, then it sends the answer to a master. Also, it is not necessary to specify the `RecipientClientId` and any specified value for `RecipientClientId` is ignored. If the sender type is master, the answer is sent to the target viewer specified by the `RecipientClientId`. `RecipientClientId` is a required input in this case.

A master client app is allowed to send an answer to any viewer, whereas a viewer client app is only allowed to send an answer to a master client app. If a viewer client app attempts to send an answer to another viewer client app, the request will NOT be honored. If there is an outstanding answer for the same client which is not yet delivered, it is overwritten with the new answer.

## Request

```
{
  "action": "SDP_ANSWER",
  "recipientClientId": "string",
  "messagePayload": "string",
  "correlationId": "string"
}
```

- **action** - Type of the message that is being sent.
  - Type: ENUM
  - Valid values: SDP\_OFFER, SDP\_ANSWER, ICE\_CANDIDATE
  - Length constraints: Minimum length of 1. Maximum length of 256.
  - Pattern: `[a-zA-Z0-9_.-]+`
  - Required: Yes
- **recipientClientId** - The unique identifier for the recipient.

- **Type:** String
- **Length constraints:** Minimum length of 1. Maximum length of 256.
- **Pattern:** [a-zA-Z0-9\_.-]+
- **Required:** Yes
- **messagePayload** - The base-64-encoded message content.
  - **Type:** String
  - **Length constraints:** Minimum length of 1. Maximum length of 10K.
  - **Required:** Yes
- **correlationId** - A unique identifier for the message.
  - **Type:** String
  - **Length constraints:** Minimum length of 1. Maximum length of 256.
  - **Pattern:** [a-zA-Z0-9\_.-]+
  - **Required:** No

## Response

No response is returned if the message is successfully received by the signaling backend. If the service encounters an error and if the `correlationId` is specified in the request, the error details are returned as a `STATUS_RESPONSE` message. For more information, see [Asynchronous Message Reception](#).

## Errors

- `InvalidArgumentException`

A specified parameter exceeds its restrictions, is not supported, or cannot be used. For more information, see the returned message.

HTTP Status Code: 400

- `ClientLimitExceededException`

Returned when the API is invoked at a rate that is too high. For more information, see [Amazon Kinesis Video Streams with WebRTC Service Quotas](#) and [Error Retries and Exponential Backoff in Amazon](#).

HTTP Status Code: 400

## Limits/Throttling

This API is throttled at an account level if the API is invoked at too high a rate. An error is returned when throttled with `ClientLimitExceededException`.

## Idempotent

This API is not idempotent.

## Retry behavior

This is counted as a new API call.

## Concurrent calls

Concurrent calls are allowed. An offer is sent once per each call.

## SendIceCandidate

Sends the ICE candidate to the target recipient. The prerequisite is that the client must be already connected to the WebSocket endpoint obtained from the `GetSignalingChannelEndpoint` API.

If the sender type is a viewer, then it sends the ICE candidate to a master. Also, it is not necessary to specify the `RecipientClientId` and any specified value for `RecipientClientId` is ignored. If the sender type is master, the ICE candidate is sent to the target specified by the `RecipientClientId`. `RecipientClientId` is a required input in this case.

A master client app is allowed to send an ICE candidate to any viewer, whereas a viewer client app is only allowed to send an ICE candidate to a master client app. If a viewer client app attempts to send an ICE candidate to another viewer client app, the request will NOT be honored.

## Request

```
{
  "action": "ICE_CANDIDATE",
  "recipientClientId": "string",
  "messagePayload": "string",
  "correlationId": "string"
}
```

- **action** - Type of the message that is being sent.
  - Type: ENUM
  - Valid values: SDP\_OFFER, SDP\_ANSWER, ICE\_CANDIDATE
  - Length constraints: Minimum length of 1. Maximum length of 256.
  - Pattern: [a-zA-Z0-9\_.-]+
  - Required: Yes
- **recipientClientId** - A unique identifier for the recipient.
  - Type: String
  - Length constraints: Minimum length of 1. Maximum length of 256.
  - Pattern: [a-zA-Z0-9\_.-]+
  - Required: No
- **messagePayload** - The base64-encoded message content.
  - Type: String
  - Length constraints: Minimum length of 1. Maximum length of 10K.
  - Required: Yes
- **correlationId** - A unique identifier for the message.
  - Type: String
  - Length constraints: Minimum length of 1. Maximum length of 256.
  - Pattern: [a-zA-Z0-9\_.-]+
  - Required: No

## Response

No response is returned if the message is successfully received by the signaling backend. If the service encounters an error and if the `correlationId` is specified in the request, the error details are returned as a `STATUS_RESPONSE` message. For more information, see [Asynchronous Message Reception](#).

## Errors

- `InvalidArgumentException`

A specified parameter exceeds its restrictions, is not supported, or cannot be used. For more information, see the returned message.

HTTP Status Code: 400

- `ClientLimitExceededException`

When the API is invoked at a rate that is too high. For more information, see [Amazon Kinesis Video Streams with WebRTC Service Quotas](#) and [Error Retries and Exponential Backoff in Amazon](#).

HTTP Status Code: 400

## Limits/Throttling

This API is throttled at an account level if the API is invoked at too high a rate. An error returned when throttled with `ClientLimitExceededException`.

## Idempotent

This API is not idempotent.

## Retry behavior

This is counted as a new API call.

## Concurrent calls

Concurrent calls are allowed. An offer is sent once per each call.

## Disconnect

A client can close a connection at any time. WebSocket-compliant libraries support close functionality. When the connection is closed, service marks the client as offline for the specific signaling channel and does not try to deliver any messages. The same behavior also applies in the event of idle connection timeout.

The service also sends disconnect indications to the client, for example, during deployments or server maintenance. The following are the defined indication messages:

- **GO\_AWAY**: This message is used to initiate the connection shutdown. It enables a client to gracefully process previous messages, disconnect, and reconnect to the signaling channel if needed.

- **RECONNECT\_ICE\_SERVER**: This message is used to initiate the relay connection shutdown and enables a client to gracefully disconnect, obtain a new ICE server configuration, and reconnect to the relay servers if needed.

## Asynchronous Message Reception

All response messages are asynchronously delivered to the recipient as events (for example, an SDP offer or SDP answer delivery). The following is the event message structure.

### Event

```
{
  "senderClientId": "string",
  "messageType": "string",
  "messagePayload": "string",
  "statusResponse": {
    "correlationId": "string",
    "errorType": "string",
    "statusCode": "string",
    "description": "string"
  }
}
```

- **senderClientId** - A unique identifier for the sender client.
  - Type: String
  - Length constraints: Minimum length of 1. Maximum length of 256.
  - Pattern: [a-zA-Z0-9\_.-]+
  - Required: No
- **messageType** - Type of the event.
  - Type: ENUM
  - Valid Types: SDP\_OFFER, SDP\_ANSWER, ICE\_CANDIDATE, GO\_AWAY, RECONNECT\_ICE\_SERVER, STATUS\_RESPONSE
  - Length constraints: Minimum length of 1. Maximum length of 256.
  - Pattern: [a-zA-Z0-9\_.-]+
  - Required: Yes
- **messagePayload** - The base64-encoded message content.



- **Type:** String
- **Length constraints:** Minimum length of 1. Maximum length of 10K.
- **Required:** No
- **correlationId** - An unique identifier of the message for which the status is meant. This is the same correlationId provided in the client messages (for example, SDP offer, SDP answer, or ICE candidate).
  - **Type:** String
  - **Length constraints:** Minimum length of 1. Maximum length of 256.
  - **Pattern:** [a-zA-Z0-9\_.-]+
  - **Required:** Yes
- **errorType** - A name to uniquely identify the error.
  - **Type:** String
  - **Length constraints:** Minimum length of 1. Maximum length of 256.
  - **Pattern:** [a-zA-Z0-9\_.-]+
  - **Required:** No
- **statusCode** - HTTP status code corresponding to the nature of the response.
  - **Type:** String
  - **Length constraints:** Minimum length of 1. Maximum length of 256.
  - **Pattern:** [a-zA-Z0-9\_.-]+
  - **Required:** No
- **description** - A string description explaining the status.
  - **Type:** String
  - **Length constraints:** Minimum length of 1. Maximum length of 1K.
  - **Required:** No

# Amazon Kinesis Video Streams with WebRTC Service Quotas

Kinesis Video Streams with WebRTC has the following service quotas:

## Important

The following service quotas are either soft **[s]**, which can be upgraded by submitting a support ticket, or hard **[h]**, which can't be increased. You will see [s] and [h] next to individual service quota in the tables below.

## Note

TPS stands for *transactions per second*.

## Topics

- [Control Plane API Service Quotas](#)
- [Signaling API Service Quotas](#)
- [TURN Service Quotas](#)
- [WebRTC Ingestion Service Quotas](#)

## Control Plane API Service Quotas

The following section describes service quotas for the control plane APIs.

API	Account service quota: Request	Account service quota: Channels	Channel-level service quota	Relevant Exceptions and Notes
CreateSignalingChannel	50 TPS [s]	us-east-1 and us-west-2		

API	Account service quota: Request	Account service quota: Channels	Channel-level service quota	Relevant Exceptions and Notes
		- 10,000 channels per account per region; all other supported Regions - 5,000 channels per account per region		
<b>DeleteSignalingChannel</b>	50 TPS [h]	N/A	5 TPS [h]	
<b>DescribeMediaStorageConfiguration</b>	50 TPS [h]		5 TPS [h]	
<b>DescribeSignalingChannel</b>	300 TPS [h]	N/A	5 TPS [h]	
<b>GetSignalingChannelEndpoint</b>	300 TPS [h]	N/A		
<b>ListSignalingChannels</b>	50 TPS [h]	N/A		

API	Account service quota: Request	Account service quota: Channels	Channel-level service quota	Relevant Exceptions and Notes
<b>ListTagsForResource</b>	50 TPS [h]	N/A	5 TPS [h]	
<b>TagResource</b>	50 TPS [h]	N/A	5 TPS [h]	
<b>UntagResource</b>	50 TPS [h]	N/A	5 TPS [h]	
<b>UpdateMediaStorageConfiguration</b>	10 TPS [h]		5 TPS [h]	
<b>UpdateSignalingChannel</b>	50 TPS [h]	N/A	5 TPS [h]	

## Signaling API Service Quotas

The following section describes service quotas for the signaling component in Kinesis Video Streams with WebRTC. For more information, see [Kinesis Video Streams with WebRTC: How It Works](#).

- **ConnectAsMaster**
  - API - 3 TPS per channel (h)
  - Maximum number of master connections per signaling channel - 1 (h)
  - Connection duration limit - 1 hour (h)
  - Idle connection timeout - 10 minutes (h)
  - When a client receives the GO\_AWAY message from the server, connection is terminated after a grace period of 1 minute (h)
- **ConnectAsViewer**

- API - 3 TPS per channel (h)
- Maximum number of viewer connections per channel - 10 (s)
- Connection duration limit - 1 hour (h)
- Idle connection timeout - 10 minutes (h)
- Once a client receives the GO\_AWAY message from the server, connection is terminated after a grace period of 1 minute (h)
- Disconnect
  - N/A
- GetIceServerConfig
  - API - 5 TPS per signaling channel (h)
- SendAlexaOffertoMaster
  - API - 5 TPS per signaling channel (h)
- SendICECandidate
  - API - 20 TPS per WebSocket connection (h)
  - Message payload size limit - 10k (h)
- SendSDPAnswer
  - API - 5 TPS per WebSocket connection (h)
  - Message payload size limit - 10k (h)
- SendSDPOffer
  - API - 5 TPS per WebSocket connection (h)
  - Message payload size limit - 10k (h)

## TURN Service Quotas

The following section describes service quotas for the Traversal Using Relays around NAT (TURN) component in Kinesis Video Streams with WebRTC. For more information, see [Kinesis Video Streams with WebRTC: How It Works](#).

- Bit Rate - 5Mbps (h)
- Credential Lifecycle - 5 minutes (h)
- Number of allocations - 50 per signaling channel (h)

# WebRTC Ingestion Service Quotas

The following section describes service quotas for the media recording component in Amazon Kinesis Video Streams WebRTC. For more information, see [???](#).

## JoinStorageSession

- API:
  - Per account - 50 TPS (h)
  - Per channel - 2 TPS (h)
- Streaming session quotas:
  - Bitrate - 1 Mbps (s)
  - Session duration - 1 hour (h)
  - Idle timeout - 3 minutes (h)

# Getting Started

This section describes how to perform the following tasks in Amazon Kinesis Video Streams with WebRTC:

- Set up your Amazon Web Services account and create an administrator.
- Create a Kinesis Video Streams with WebRTC signaling channel.
- Use the Kinesis Video Streams with WebRTC SDKs to configure master and viewer to perform peer-to-peer video and audio streaming over a signaling channel.

If you are new to Kinesis Video Streams with WebRTC, we recommend that you read [Kinesis Video Streams with WebRTC: How It Works](#) first.

## Topics

- [Set Up an Amazon Account and Create an Administrator](#)
- [Create a Signaling Channel](#)
- [Stream Live Media](#)

## Set Up an Amazon Account and Create an Administrator

Before you use Kinesis Video Streams with WebRTC for the first time, complete the following tasks:

### Topics

- [Sign up for an Amazon Web Services account](#)
- [Secure IAM users](#)
- [Create an AmazonAccount Key](#)

## Sign up for an Amazon Web Services account

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

### To sign up for Amazon Web Services

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

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

## Secure IAM users

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

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

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

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

## Create an AmazonAccount Key

You need an Amazon Account Key to access Kinesis Video Streams with WebRTC programmatically.

To create an Amazon Account Key, do the following:

1. Sign in to the Amazon Web Services Management Console and open the IAM console at <https://console.amazonaws.cn/iam/>.
2. Choose **Users** in the navigation bar, and choose the **Administrator** user.
3. Choose the **Security credentials** tab, and choose **Create access key**.
4. Record the **Access key ID**. Choose **Show** under **Secret access key**, and then record the **Secret access key**.

## Create a Signaling Channel

You can use the Kinesis Video Streams console, the Amazon APIs ([CreateSignalingChannel](#)), or the Amazon CLI to create your signaling channels.



## Create a Signaling Channel Using the Console

1. Sign in to the Amazon Web Services Management Console and open the [Amazon Kinesis Video Streams console](#).
2. On the **Signaling channels** page, choose **Create signaling channel**.
3. On the **Create a new signaling channel** page, type in the name for the signaling channel. Leave the default **Time-to-live (Ttl)** value of 60 seconds unchanged.
4. Choose **Create signaling channel**.
5. After the signaling channel is created, review its details on the channel's details page.

## Stream Live Media

The Kinesis Video Streams with WebRTC includes the following SDKs:

- [WebRTC SDK in C for Embedded Devices](#)
- [Kinesis Video Streams with WebRTC SDK in JavaScript for web applications](#)
- [WebRTC SDK for Android](#)
- [WebRTC SDK for iOS](#)

Each SDK includes corresponding samples and step-by-step instructions that can help you build and run those applications. You can use these samples for low latency, live, two-way audio and video streaming and data exchange between any combinations of Web/Android/iOS applications or embedded devices. In other words, you can stream live audio and video from an embedded camera device to Android or web applications or between two Android applications.

## WebRTC SDK in C for Embedded Devices

The following step-by-step instructions describe how to download, build, and run the Kinesis Video Streams with WebRTC SDK in C for embedded devices and its corresponding samples.

### Download the Kinesis Video Streams with WebRTC SDK in C

To download the Kinesis Video Streams with WebRTC SDK in C for embedded devices, run the following command:

```
$ git clone --recursive https://github.com/aws-labs/amazon-kinesis-video-streams-webrtc-sdk-c.git
```

## Build the Kinesis Video Streams with WebRTC SDK in C

### Important

Before you complete these steps on a macOS and depending on the version of the macOS you have, you must run `xcode-select --install` to download the package with the command line tools and header. Then open `/Library/Developer/CommandLineTools/Packages/macOS_SDK_headers_for_macOS_10.14.pkg` and follow the installer to install the command line tools and header. You only need to do this once and before invoking `cmake`. If you already have the command line tools and header installed, you do not need to run this command again.

Complete the following steps:

1. Install `cmake`:
  - On macOS, run `brew install cmake pkg-config srtp`
  - on Ubuntu, run `sudo apt-get install pkg-config cmake libcap2 libcap-dev`
2. Obtain the access key and the secret key of the Amazon Web Services account that you want to use for this demo.
3. Run the following command to create a build directory in your downloaded WebRTC C SDK, and execute `cmake` from it:

```
$ mkdir -p amazon-kinesis-video-streams-webrtc-sdk-c/build; cd amazon-kinesis-video-streams-webrtc-sdk-c/build; cmake ..
```

4. Now that you're in the build directory you just created with the step above, run `make` to build the WebRTC C SDK and its provided samples.

**Note**

The `kvsWebrtcClientMasterGstSample` will NOT be built if the system doesn't have `gststreamer` installed. To make sure it is built (on macOS) you must run: `brew install gststreamer gst-plugins-base gst-plugins-good`

## Run the Samples for the WebRTC SDK in C

After you complete the procedure above, you end up with the following sample applications in your build directory:

- `kvsWebrtcClientMaster` - This application sends sample H264/Opus frames (path: `/samples/h264SampleFrames` and `/samples/opusSampleFrames`) via the signaling channel. It also accepts incoming audio, if enabled in the browser. When checked in the browser, it prints the metadata of the received audio packets in your terminal.
- `kvsWebrtcClientViewer` - This application accepts sample H264/Opus frames and prints them out.
- `kvsWebrtcClientMasterGstSample` - This application sends sample H264/Opus frames from a `GStreamer` pipeline.

To run any of these samples, complete the following steps:

1. Setup your environment with your Amazon Web Services account credentials:

```
export AWS_ACCESS_KEY_ID= Your Amazon Web Services account access key
export AWS_SECRET_ACCESS_KEY= Amazon Web Services account secret key
export AWS_KVS_CACERT_PATH= Full path of your cert.pem file. It is typically
available in the certs directory inside
Kinesis-video-webrtc-native-build/certs/cert.pm
```

2. Run either of the sample applications by passing to it the name that you want to give to your signaling channel. The application creates the signaling channel using the name you provide. For example, to create a signaling channel called `myChannel` and to start sending sample H264/Opus frames via this channel, run the following command:

```
./kvsWebrtcClientMaster myChannel
```

When the command line application prints `Connection established`, you can proceed to the next step.

- Now that your signaling channel is created and the connected master is streaming media to it, you can view this stream. For example, you can view this live stream in a web application. To do so, open the WebRTC SDK Test Page using the steps in [Use the Kinesis Video Streams with WebRTC test page](#) and set the following values using the same Amazon credentials and the same signaling channel that you specified for the master above:

- Access key ID
- Secret access key
- Signaling channel name
- Client ID (optional)

Choose **Start viewer** to start live video streaming of the sample H264/Opus frames.

## Kinesis Video Streams with WebRTC SDK in JavaScript for web applications

You can find the Kinesis Video Streams with WebRTC SDK in JavaScript for web applications and its corresponding samples in [GitHub](#).

### Topics

- [Install the Kinesis Video Streams with WebRTC SDK in JavaScript](#)
- [Kinesis Video Streams with WebRTC JavaScript SDK documentation](#)
- [Use the Kinesis Video Streams with WebRTC test page](#)
- [Edit the Kinesis Video Streams with WebRTC test page](#)

## Install the Kinesis Video Streams with WebRTC SDK in JavaScript

Whether and how you install the Kinesis Video Streams with WebRTC SDK in JavaScript depends on whether the code executes in Node .js modules or browser scripts.

## NodeJS module

The preferred way to install the Kinesis Video Streams with WebRTC SDK in JavaScript for Node.js is to use [npm, the Node.js package manager](#).

The package is hosted at <https://www.npmjs.com/package/amazon-kinesis-video-streams-webrtc>.

To install this SDK in your Node . js project, use the terminal to navigate to the the same directory as your project's package . json:

Type the following:

```
npm install amazon-kinesis-video-streams-webrtc
```

You can import the SDK classes like typical Node.js modules:

```
// JavaScript
const SignalingClient = require('amazon-kinesis-video-streams-
webrtc').SignalingClient;
// TypeScript
import { SignalingClient } from 'amazon-kinesis-video-streams-webrtc';
```

## Browser

You don't have to install the SDK to use it in browser scripts. You can load the hosted SDK package directly from Amazon with a script in your HTML pages.

To use the SDK in the browser, add the following script element to your HTML pages:

```
<script src="https://unpkg.com/amazon-kinesis-video-streams-webrtc/dist/kvs-
webrtc.min.js"></script>
```

After the SDK loads in your page, the SDK is available from the global variable `KVSWebRTC` (or `window.KVSWebRTC`).

For example, `window.KVSWebRTC.SignalingClient`.

## Kinesis Video Streams with WebRTC JavaScript SDK documentation

The documentation for the SDK methods are on the GitHub readme, under [Documentation](#).

In the [Usage](#) section, there is additional information for integrating this SDK along with the Amazon SDK for JavaScript to build a web-based viewer application.

See the `examples` directory for an example of a complete application, including both a master and viewer role.

## Use the Kinesis Video Streams with WebRTC test page

Kinesis Video Streams with WebRTC also hosts a test page that you can use to either create a new signaling channel or connect to an existing channel and use it as a master or viewer.

The Kinesis Video Streams with WebRTC test page is located at <https://awslabs.github.io/amazon-kinesis-video-streams-webrtc-sdk-js/examples/index.html>.

The code for the test page is in the `examples` directory.

### Topics

- [Stream from the test page to the Amazon Web Services Management Console](#)
- [Stream from the test page to the test page](#)

## Stream from the test page to the Amazon Web Services Management Console

1. Open the [Kinesis Video Streams with WebRTC test page](#) and complete the following:
  - Amazon Web Services Region. For example, `us-west-2`.
  - The Amazon access key and the secret key for your IAM user or role. Leave the session token blank if you are using long-term Amazon credentials.
  - The name of the signaling channel to which you want to connect.

If you want to connect to a new signaling channel, choose **Create Channel** to create a signaling channel with the value provided in the box.

#### Note

Your signaling channel name must be unique for the current account and region. You can use letters, numbers, underscores (`_`), and hyphens (`-`), but not spaces.

- Whether you want to send audio, video, or both.

- ICE candidate generation. Leave STUN/TURN selected and leave Trickle ICE enabled.
2. Choose **Start Master** to connect to the signaling channel.  
  
Allow access to your camera and/or microphone, if needed.
  3. Open the [Kinesis Video Streams console](#) in the Amazon Web Services Management Console.  
  
Make sure the correct region is selected.
  4. In the left navigation, select [signaling channels](#).  
  
Select the name of the signaling channel above. Use the search bar, if needed.
  5. Expand the **Media playback viewer** section.
  6. Choose the **play** button on the video player. This joins the WebRTC session as a viewer.

The media that is being sent on the demo page should display in the Amazon Web Services Management Console.

### Stream from the test page to the test page

1. Open the [Kinesis Video Streams with WebRTC test page](#) and complete the following information:
  - Amazon Web Services Region. For example, us-west-2.
  - The Amazon access key and the secret key for your IAM user or role. Leave the session token blank if you are using long-term Amazon credentials.
  - The name of the signaling channel to which you want to connect.

If you want to connect to a new signaling channel, choose **Create Channel** to create a signaling channel with the value provided in the box.

#### Note

Your signaling channel name must be unique for the current account and region. You can use letters, numbers, underscores (\_), and hyphens (-), but not spaces.

- Whether you want to send audio, video, or both.
  - ICE candidate generation. Leave STUN/TURN selected and leave Trickle ICE enabled.
2. Choose **Start Master** to connect to the signaling channel as the master role.

Allow access to your camera and/or microphone, if needed.

3. Open another browser tab and open the [Kinesis Video Streams with WebRTC test page](#). All of the information from the previous run should load.
4. Scroll down and choose **Start Viewer** to connect to the signaling channel as the viewer role.

You should see the media being exchanged between the master and viewer.

## Edit the Kinesis Video Streams with WebRTC test page

To edit the SDK and test page for development purposes, follow the instructions below.

### Prerequisite

NodeJS version 16+

#### Note

We recommend downloading the latest long term support (LTS) version from <https://nodejs.org/en/download>.

### Edit the test page

1. Download the Kinesis Video Streams with WebRTC SDK in JavaScript.

Type the following in the terminal:

```
git clone https://github.com/aws-labs/amazon-kinesis-video-streams-webrtc-sdk-js.git
```

2. Navigate to the directory with the package.json file. The file is located in the repository's root directory.

Type the following in the terminal:

```
cd amazon-kinesis-video-streams-webrtc-sdk-js
```

3. Install dependencies.

Type the following [npm CLI](#) command in the terminal:



```
npm install
```

4. Start the web server to start serving web pages.

Type the following [npm CLI](#) command in the terminal:

```
npm run develop
```

5. In your browser, visit <http://localhost:3001/>.

You can make edits to the web page by editing the files in the `examples` directory.

## WebRTC SDK for Android

The following step-by-step instructions describe how to download, build, and run the Kinesis Video Streams with WebRTC SDK for Android and its corresponding samples.

### Note

Kinesis Video Streams does not support IPv6 addresses on Android. For more information and steps about disabling IPv6 on your Android device, see <https://support.surfshark.com/hc/en-us/articles/360011828279-How-to-disable-IPv6-on-Android->.

## Download the WebRTC SDK for Android

To download the WebRTC SDK in Android, run the following command:

```
$ git clone https://github.com/aws-labs/amazon-kinesis-video-streams-webrtc-sdk-android.git
```

## Build the WebRTC SDK in Android

To build the WebRTC SDK in Android, complete the following steps:

1. Import the Android WebRTC SDK into the Android Studio integrated development environment (IDE) by opening `amazon-kinesis-video-streams-webrtc-sdk-android/build.gradle` with **Open as Project**.

2. If you open the project for the first time, it automatically syncs. If not - initiate a sync. When you see a build error, choose to install any required SDKs by choosing **Install missing SDK package(s)**, then choose **Accept** and complete the install.
3. Configure Amazon Cognito (user pool and identity pool) settings. For details steps, see [Configure Amazon Cognito for the Android WebRTC SDK](#). This generates authentication and authorization settings required to build the Android WebRTC SDK.
4. In your Android IDE, open `awsconfiguration.json` (from `src/main/res/raw/`). The file looks like the following:

```
{
  "Version": "1.0",
  "CredentialsProvider": {
    "CognitoIdentity": {
      "Default": {
        "PoolId": "REPLACE_ME",
        "Region": "REPLACE_ME"
      }
    }
  },
  "IdentityManager": {
    "Default": {}
  },
  "CognitoUserPool": {
    "Default": {
      "AppClientSecret": "REPLACE_ME",
      "AppClientId": "REPLACE_ME",
      "PoolId": "REPLACE_ME",
      "Region": "REPLACE_ME"
    }
  }
}
```

Update `awsconfiguration.json` with the values generated by running the steps in [Configure Amazon Cognito for the Android WebRTC SDK](#).

5. Make sure your Android device is connected to the computer where you're running the Android IDE. In the Android IDE, select the connected device and then build and run the WebRTC Android SDK.

This step installs an app called `AWSKinesisVideoWebRTCDemoApp` on your Android device. Using this app, you can verify live WebRTC audio/video streaming between mobile, web and IoT device clients.

## Run the Android Sample Application

Complete the following steps:

1. On your Android device, open **AWSKinesisVideoWebRTCDemoApp** and log in using either a new (by creating it first) or an existing Amazon Cognito account.
2. In **AWSKinesisVideoWebRTCDemoApp**, navigate to the **Channel Configuration** page and either create a new signaling channel or choose an existing one.

### Note

Currently, using the sample application in this SDK, you can only run one signalling channel in **AWSKinesisVideoWebRTCDemoApp**.


3. Optional: choose a unique **Client Id** if you want to connect to this channel as a viewer. Client ID is required only if multiple viewers are connected to a channel. This helps channel's master identify respective viewers.
4. Choose the Amazon Web Services Region and whether you want to send audio or video data, or both.
5. To verify peer-to-peer streaming, do any of the following:

### Note

Ensure that you specify the same signaling channel name, Amazon region, viewer ID, and the Amazon account ID on all clients that you're using in this demo.

- Peer-to-peer streaming between two Android devices: master and viewer
  - Using procedures above, download, build, and run the Android WebRTC SDK on two Android devices.

- Open **AWSKinesisVideoWebRTCDemoApp** on one Android device in master mode (choose **START MASTER**) to start a new session (signaling channel).

 **Note**

Currently, there can only be one master for any given signaling channel.

- Open **AWSKinesisVideoWebRTCDemoApp** on your second Android device in viewer mode to connect to the signaling channel (session) started in the step above (choose **START VIEWER**).

Verify that the viewer can see master's audio/video data.

- Peer-to-peer streaming between the embedded SDK master and an Android device viewer
  - Download, build, and run the [WebRTC SDK in C for Embedded Devices](#) in master mode on a camera device.
  - Using procedures above, download, build, and run the Android WebRTC SDK on an Android device. Open **AWSKinesisVideoWebRTCDemoApp** on this Android device in viewer mode and verify that the viewer can see the embedded SDK master's audio/video data.
- Peer-to-peer streaming between Android device as master and web browser as viewer
  - Using procedures above, download, build, and run the Android WebRTC SDK on an Android device. Open **AWSKinesisVideoWebRTCDemoApp** on this Android device in master mode (choose **START MASTER**) to start a new session (signaling channel).
  - Download, build, and run the [Kinesis Video Streams with WebRTC SDK in JavaScript for web applications](#) as viewer and verify that the viewer can see the Android master's audio/video.

## Configure Amazon Cognito for the Android WebRTC SDK

### Prerequisites

- We recommend [Android Studio](#) for examining, editing, and running the application code. We recommend using the latest stable version.
- In the sample code, you provide Amazon Cognito credentials.

Follow these procedures to set up an Amazon Cognito user pool and identity pool.

## Set up a user pool

### To set up a user pool

1. Sign in to the [Amazon Cognito console](#) and verify the region is correct.
2. In the navigation on the left choose **User pools**.
3. In the **User pools** section, choose **Create user pool**.
4. Complete the following sections:
  - a. **Step 1: Configure sign-in experience** - In the **Cognito user pool sign-in options** section, select the appropriate options.  
  
Select **Next**.
  - b. **Step 2: Configure security requirements** - Select the appropriate options.  
  
Select **Next**.
  - c. **Step 3: Configure sign-up experience** - Select the appropriate options.  
  
Select **Next**.
  - d. **Step 4: Configure message delivery** - Select the appropriate options.  
  
In the **IAM role selection** field, select an existing role or create a new role.  
  
Select **Next**.
  - e. **Step 5: Integrate your app** - Select the appropriate options.  
  
In the **Initial app client** field, choose **Confidential client**.  
  
Select **Next**.
  - f. **Step 6: Review and create** - Review your selections from the previous sections, then choose **Create user pool**.
5. On the **User pools** page, select the pool that you just created.  
  
Copy the **User pool ID** and make note of this for later. In the `awsconfiguration.json` file, this is `CognitoUserPool.Default.PoolId`.
6. Select the **App integration** tab and go to the bottom of the page.
7. In the **App client list** section, choose the **App client name** you just created.

Copy the **Client ID** and make note of this for later. In the `awsconfiguration.json` file, this is `CognitoUserPool.Default.AppClientId`.

8. Show the **Client secret** and make note of this for later. In the `awsconfiguration.json` file, this is `CognitoUserPool.Default.AppClientSecret`.

## Set up an identity pool

### To set up an identity pool

1. Sign in to the [Amazon Cognito console](#) and verify the region is correct.
2. In the navigation on the left choose **Identity pools**.
3. Choose **Create identity pool**.
4. Configure the identity pool.
  - a. **Step 1: Configure identity pool trust** - Complete the following sections:
    - **User access** - Select **Authenticated access**
    - **Authenticated identity sources** - Select **Amazon Cognito user pool**

Select **Next**.
  - b. **Step 2: Configure permissions** - In the **Authenticated role** section, complete the following fields:
    - **IAM role** - Select **Create a new IAM role**
    - **IAM role name** - Enter a name and make note of it for a later step.

Select **Next**.
  - c. **Step 3: Connect identity providers** - In the **User pool details** section complete the following fields:
    - **User pool ID** - Select the user pool you created earlier.
    - **App client ID** - Select the app client ID you created earlier.

Select **Next**.

- d. **Step 4: Configure properties** - Type a name in the **Identity pool name** field.

Select **Next**.

- e. **Step 5: Review and create** - Review your selections in each of the sections, then select **Create identity pool**.

5. On the **Identity pools** page, select your new identity pool.

Copy the **Identity pool ID** and make note of this for later. In the `awsconfiguration.json` file, this is `CredentialsProvider.CognitoIdentity.Default.PoolId`.

6. Update the permissions for the IAM role.

- a. Sign in to the Amazon Web Services Management Console and open the IAM console at <https://console.amazonaws.cn/iam/>.
- b. In the navigation on the left, choose **Roles**.
- c. Find and select the role you created above.

 **Note**

Use the search bar, if needed.

- d. Select the attached permissions policy.

Select **Edit**.

- e. Select the **JSON** tab and replace the policy with the following:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "cognito-identity:*",
        "kinesisvideo:*"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

```
}
```

Select **Next**.

- f. Select the box next to **Set this new version as the default** if it isn't already selected.

Select **Save changes**.

## WebRTC SDK for iOS

The following step-by-step instructions describe how to download, build, and run the Kinesis Video Streams WebRTC SDK in iOS and its corresponding samples.

### Download the WebRTC SDK in iOS

To download the WebRTC SDK in iOS, run the following command:

```
$ git clone https://github.com/aws-labs/amazon-kinesis-video-streams-webrtc-sdk-ios.git
```

### Build the WebRTC SDK in iOS

Complete the following steps:

1. Import the iOS WebRTC SDK into the XCode integrated development environment (IDE) on an iOS computer by opening `KinesisVideoWebRTCDemoApp.xcworkspace` (path: `amazon-kinesis-video-streams-webrtc-sdk-ios/Swift/AWSKinesisVideoWebRTCDemoApp.xcworkspace`).
2. If you open the project for the first time, it automatically builds. If not, initiate a build.

You might see the following error:

```
error: The sandbox is not in sync with the Podfile.lock. Run 'pod install' or update your CocoaPods installation.
```

If you see this, do the following:

- a. Change your current working directory to `amazon-kinesis-video-streams-webrtc-sdk-ios/Swift` and run the following in the command line:

```
pod cache clean --all
pod install
```



- b. Change your current working directory to `amazon-kinesis-video-streams-webrtc-sdk-ios` and run the following at the command line:

```
$ git checkout Swift/Pods/AWSCore/AWSCore/Service/AWSService.m
```

- c. Build again.
3. Configure Amazon Cognito (user pool and identity pool) settings. For details steps, see [Configure Amazon Cognito for the iOS WebRTC SDK](#). This generates authentication and authorization settings required to build the iOS WebRTC SDK.
  4. In your IDE, open the `awsconfiguration.json` file (from `/Swift/KVSiOSApp`). The file looks like the following:

```
{
  "Version": "1.0",
  "CredentialsProvider": {
    "CognitoIdentity": {
      "Default": {
        "PoolId": "REPLACEME",
        "Region": "REPLACEME"
      }
    }
  },
  "IdentityManager": {
    "Default": {}
  },
  "CognitoUserPool": {
    "Default": {
      "AppClientSecret": "REPLACEME",
      "AppClientId": "REPLACEME",
      "PoolId": "REPLACEME",
      "Region": "REPLACEME"
    }
  }
}
```

Update `awsconfiguration.json` with the values generated by running the steps in [Configure Amazon Cognito for the Android WebRTC SDK](#).

5. In your IDE, open the `Constants.swift` file (from `/Swift/KVSiOSApp`). The file looks like the following:

```
import Foundation
import AWSCognitoIdentityProvider

let CognitoIdentityUserPoolRegion = AWSRegionType.USWest2
let CognitoIdentityUserPoolId = "REPLACEME"
let CognitoIdentityUserPoolAppClientId = "REPLACEME"
let CognitoIdentityUserPoolAppClientSecret = "REPLACEME"

let AWSCognitoUserPoolsSignInProviderKey = "UserPool"
let CognitoIdentityPoolID = "REPLACEME"

let AWSKinesisVideoEndpoint = "https://kinesisvideo.us-west-2.amazonaws.com"
let AWSKinesisVideoKey = "kinesisvideo"

let VideoProtocols = ["WSS", "HTTPS"]

let ConnectAsMaster = "connect-as-master"
let ConnectAsViewer = "connect-as-viewer"

let MasterRole = "MASTER"
let ViewerRole = "VIEWER"

let ClientID = "ConsumerViewer"
```

Update `Constants.swift` with the values generated by running the steps in [Configure Amazon Cognito for the Android WebRTC SDK](#).

6. Make sure your iOS device is connected to the Mac computer where you're running XCode. In XCode, select the connected device and then build and run the WebRTC iOS SDK.


This step installs an app called `AWSKinesisVideoWebRTCDemoApp` on your iOS device. Using this app, you can verify live WebRTC audio/video streaming between mobile, web and IoT device clients.

## Run the iOS Sample Application

Complete the following steps:

1. On your iOS device, open **`AWSKinesisVideoWebRTCDemoApp`** and log in using either a new (by creating it first) or an existing Amazon Cognito account.

2. In **AWSKinesisVideoWebRTCDemoApp**, navigate to the **Channel Configuration** page and either create a new signaling channel or choose an existing one.

 **Note**

Currently, using the sample application in this SDK, you can only run one signalling channel in **AWSKinesisVideoWebRTCDemoApp**.

3. (Optional) Choose a unique **Client Id** if you want to connect to this channel as a viewer. Client Id is required only if multiple viewers are connected to a channel. This helps channel's master identify respective viewers.
4. Choose the Amazon Web Services Region and whether you want to send audio or video data, or both.
5. To verify peer-to-peer streaming, do any of the following:

 **Note**

Ensure that you specify the same signaling channel name, Amazon region, viewer ID, and the Amazon account ID on all clients that you're using in this demo.

- Peer-to-peer streaming between two iOS devices: master and viewer
  - Using procedures above, download, build, and run the iOS WebRTC SDK on two iOS devices.
  - Open **AWSKinesisVideoWebRTCDemoApp** on one iOS device in master mode (choose **START MASTER**) to start a new session (signaling channel).

 **Note**

Currently, there can only be one master for any given signaling channel.

- Open **AWSKinesisVideoWebRTCDemoApp** on your second iOS device in viewer mode to connect to the signaling channel (session) started in the step above (choose **START VIEWER**).

Verify that the viewer can see master's audio/video data.

- Peer-to-peer streaming between the embedded SDK master and an iOS device viewer

- Download, build, and run the [WebRTC SDK in C for Embedded Devices](#) in master mode on a camera device.
- Using procedures above, download, build, and run the iOS WebRTC SDK on an iOS device. Open **AWSKinesisVideoWebRTCDemoApp** on this iOS device in viewer mode and verify that the iOS viewer can see the embedded SDK master's audio/video data.
- Peer-to-peer streaming between iOS device as master and web browser as viewer
  - Using procedures above, download, build, and run the iOS WebRTC SDK on an iOS device. Open **AWSKinesisVideoWebRTCDemoApp** on this iOS device in master mode (choose **START MASTER**) to start a new session (signaling channel).
  - Download, build, and run the [Kinesis Video Streams with WebRTC SDK in JavaScript for web applications](#) as viewer and verify that the JavaScript viewer can see the Android master's audio/video.

## Configure Amazon Cognito for the iOS WebRTC SDK

### Prerequisites

- We recommend XCode for examining, editing, and running the application code. We recommend the latest version.
- In the sample code, you provide Amazon Cognito credentials.

Follow these procedures to set up an Amazon Cognito user pool and identity pool.

### Set up a user pool

#### To set up a user pool

1. Sign in to the [Amazon Cognito console](#) and verify the region is correct.
2. In the navigation on the left choose **User pools**.
3. In the **User pools** section, choose **Create user pool**.
4. Complete the following sections:
  - a. **Step 1: Configure sign-in experience** - In the **Cognito user pool sign-in options** section, select the appropriate options.

Select **Next**.

- b. **Step 2: Configure security requirements** - Select the appropriate options.

Select **Next**.

- c. **Step 3: Configure sign-up experience** - Select the appropriate options.

Select **Next**.

- d. **Step 4: Configure message delivery** - Select the appropriate options.

In the **IAM role selection** field, select an existing role or create a new role.

Select **Next**.

- e. **Step 5: Integrate your app** - Select the appropriate options.

In the **Initial app client** field, choose **Confidential client**.

Select **Next**.

- f. **Step 6: Review and create** - Review your selections from the previous sections, then choose **Create user pool**.

5. On the **User pools** page, select the pool that you just created.

Copy the **User pool ID** and make note of this for later. In the `awsconfiguration.json` file, this is `CognitoUserPool.Default.PoolId`.

6. Select the **App integration** tab and go to the bottom of the page.

7. In the **App client list** section, choose the **App client name** you just created.

Copy the **Client ID** and make note of this for later. In the `awsconfiguration.json` file, this is `CognitoUserPool.Default.AppClientId`.

8. Show the **Client secret** and make note of this for later. In the `awsconfiguration.json` file, this is `CognitoUserPool.Default.AppClientSecret`.

## Set up an identity pool

### To set up an identity pool

1. Sign in to the [Amazon Cognito console](#) and verify the region is correct.
2. In the navigation on the left choose **Identity pools**.
3. Choose **Create identity pool**.

#### 4. Configure the identity pool.

##### a. **Step 1: Configure identity pool trust** - Complete the following sections:

- **User access** - Select **Authenticated access**
- **Authenticated identity sources** - Select **Amazon Cognito user pool**

Select **Next**.

##### b. **Step 2: Configure permissions** - In the **Authenticated role** section, complete the following fields:

- **IAM role** - Select **Create a new IAM role**
- **IAM role name** - Enter a name and make note of it for a later step.

Select **Next**.

##### c. **Step 3: Connect identity providers** - In the **User pool details** section complete the following fields:

- **User pool ID** - Select the user pool you created earlier.
- **App client ID** - Select the app client ID you created earlier.

Select **Next**.

##### d. **Step 4: Configure properties** - Type a name in the **Identity pool name** field.

Select **Next**.

##### e. **Step 5: Review and create** - Review your selections in each of the sections, then select **Create identity pool**.

#### 5. On the **Identity pools** page, select your new identity pool.

Copy the **Identity pool ID** and make note of this for later. In the `awsconfiguration.json` file, this is `CredentialsProvider.CognitoIdentity.Default.PoolId`.

#### 6. Update the permissions for the IAM role.

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

##### b. In the navigation on the left, choose **Roles**.

- c. Find and select the role you created above.

 **Note**

Use the search bar, if needed.

- d. Select the attached permissions policy.

Select **Edit**.

- e. Select the **JSON** tab and replace the policy with the following:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "cognito-identity:*",
        "kinesisvideo:*"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

Select **Next**.

- f. Select the box next to **Set this new version as the default** if it isn't already selected.

Select **Save changes**.

## Client Metrics for the WebRTC C SDK

Applications built with Amazon Kinesis Video Streams with WebRTC are comprised of various moving parts, including networking, signaling, candidates exchange, peer connection, and data exchange. Kinesis Video Streams with WebRTC in C supports various client-side metrics that enable you to monitor and track the performance and usage of these components in your applications. The supported metrics fall into two major categories: custom metrics defined specifically for the

Kinesis Video Streams' implementation of signaling and networking, and media and data-related protocol-specific metrics that are derived from the [W3C](#) standard. Note that only a subset of the W3C standard metrics is currently supported for Kinesis Video Streams with WebRTC in C.

## Topics

- [Signaling Metrics](#)
- [W3C Standard Metrics Supported for WebRTC C SDK](#)

## Signaling Metrics

Signaling metrics can be used to understand how the signaling client behaves while your application is running. You can use the `STATUS signalingClientGetMetrics (SIGNALING_CLIENT_HANDLE, PSignalingClientMetrics)` API to obtain these signaling metrics. Here's an example usage pattern:

```
SIGNALING_CLIENT_HANDLE signalingClientHandle;
SignalingClientMetrics signalingClientMetrics;
STATUS retStatus = signalingClientGetMetrics(signalingClientHandle,
&signalingClientMetrics);
printf("Signaling client connection duration: %" PRIu64 " ms",
(signalingClientMetrics.signalingClientStats.connectionDuration /
HUNDREDS_OF_NANOS_IN_A_MILLISECOND));
```

The Definition of `signalingClientStats` can be found in [Stats.h](#).

The following signaling metrics are currently supported:

Metric	Description
<b>cpApiCallLatency</b>	Calculate latency for control plane API calls. Calculation is done using Exponential Moving Average (EMA). The associated calls include: describeChannel, createChannel, getChannelEndpoint and deleteChannel.



Metric	Description	
<b>dpApiCallLatency</b>	Calculate latency for data plane API calls. Calculation is done using Exponential Moving Average (EMA). The associated calls include: <code>getIceConfig</code> .	
<b>signalingClientUptime</b>	This indicates the time for which the client object exists. Every time this metric is invoked, the most recent uptime value is emitted.	
<b>connectionDuration</b>	If connection is established, this emits the duration for which the connection is alive. Else, a value of 0 is emitted. This is different from signaling client uptime since, connections come and go, but <code>signalingClientUptime</code> is indicative of the client object itself.	
<b>numberOfMessagesSent</b>	This value is updated when the peer sends an offer, answer, or an ICE candidate.	
<b>numberOfMessagesReceived</b>	Unlike <code>numberOfMessagesSent</code> , this metric is updated for any type of signaling message. The types of signaling messages are available in <code>SIGNALING_MESSAGE_TYPE</code> .	

Metric	Description	
<b>iceRefreshCount</b>	This is incremented when <code>getIceConfig</code> is invoked. The rate at which this is invoked is based on the TTL as part of the ICE configuration received. Each time a fresh set of ICE configuration is received, a timer is set to refresh next time, given the validity of the credentials in the configuration minus some grace period.	
<b>numberOfErrors</b>	The counter is used to track the number of errors generated within the signaling client. Errors generated while getting ICE configuration, getting signaling state, tracking signaling metrics, sending signaling message, and connecting the signaling client to the web socket in order to send/receive messages are tracked.	
<b>numberOfRuntimeErrors</b>	The metric includes errors that are incurred while the core of the signaling client is running. Scenarios like reconnect failures, message receive failures, and ICE configuration refresh errors are tracked here.	

Metric	Description	
<b>numberOfReconnects</b>	The metric is incremented on every reconnect. This is a useful metric to understand the stability of the network connection in the set up.	

## W3C Standard Metrics Supported for WebRTC C SDK

A subset of the [W3C](#) standard metrics is currently supported for the applications built with the WebRTC C SDK. These fall into the following categories:

- Networking:
  - [Ice Candidate](#): these metrics provide information about the selected local and remote candidates for data exchange between the peers. This includes server source of the candidate, IP address, type of candidate selected for the communication, and candidate priority. These metrics are useful as a snapshot report.
  - [Ice Server](#): these metrics are for gathering operational information about the different ICE servers supported. This is useful when trying to understand the server that is primarily being used for communication and connectivity checks. In some instances, it is also useful to examine these metrics if the gathering of candidates fails.
  - [Ice Candidate Pair](#): these metrics are for understanding the number of bytes/packets that are being exchanged between the peers and also time-related measurements.
- Media and data:
  - [Remote Inbound RTP](#): these metrics represent the endpoint perspective of the data stream sent by the sender.
  - [Outbound RTP](#): these metrics provide information about the outgoing RTP stream. They can also be very useful when analyzing choppy streaming or streaming stops.
  - [Inbound RTP](#): these metrics provide information about the incoming media.
  - [Data channel metrics](#): these metrics can help you analyze the number of messages and bytes sent and received over a data channel. The metrics can be pulled by using the channel ID.

You can use the `STATUS rtcPeerConnectionGetMetrics` (`PRtcPeerConnection`, `PRtcRtpTransceiver`, `PRtcStats`) API to gather metrics related to ICE, RTP and the data channel. Here's a usage example:

```
RtcStats rtcStats;
rtcStats.requestedTypeOfStats = RTC_STATS_TYPE_LOCAL_CANDIDATE;
STATUS retStatus = rtcPeerConnectionGetMetrics (pRtcPeerConnection, NULL, &rtcStats);
printf("Local Candidate address: %s\n",
    rtcStats.rtcStatsObject.localIceCandidateStats.address);
```

Here's another example that shows usage pattern to get transceiver related stats:

```
RtcStats rtcStats;
PRtcRtpTransceiver pVideoRtcRtpTransceiver;
rtcStats.requestedTypeOfStats = RTC_STATS_TYPE_OUTBOUND_RTP;
STATUS retStatus = rtcPeerConnectionGetMetrics (pRtcPeerConnection,
    pVideoRtcRtpTransceiver, &rtcStats);
printf("Number of packets discarded on send: %s\n",
    rtcStats.rtcStatsObject.outboundRtpStreamStats.packetsDiscardedOnSend);
```

In the above example, if the second argument to `rtcPeerConnectionGetMetrics()` is `NULL`, data for the first transceiver in the list is returned.

Definition for `rtcStatsObject` can be found in [Stats.h](#). and definition for `RtcStats` can be found in [Include.h](#).

Sample usages of the APIs and the different metrics can be found in the [samples](#) directory in the WebRTC C SDK repository and in the [Kinesis Video Stream demos repository](#).

The following [W3C](#) standard metrics are currently supported for the applications built with the WebRTC C SDK.

## Topics

- [Networking](#)
- [Media](#)
- [Data Channel](#)

## Networking

ICE Server Metrics:

<b>Metric</b>	<b>Description</b>	
<b>URL</b>	URL of the STUN/TURN server being tracked	
<b>Port</b>	Port number used by the client	
<b>Protocol</b>	Transport protocol extracted from ICE Server URI. If the value is UDP, ICE tries TURN over UDP, else ICE tried TURN over TCP/TLS. If the URI does not contain transport, ICE tries TURN over UDP and TCP/TLS. In case of STUN server, this field is empty.	
<b>Total Requests Sent</b>	The value is updated for every srflx candidate request and while sending binding request from turn candidates.	
<b>Total Responses Received</b>	The value is updated every time a STUN binding response is received.	
<b>Total Round Trip Time</b>	The value is updated every time an equivalent response is received for a request. The request packet is tracked in a hash map with the checksum as the key.	

ICE Candidate Stats: Only the information about the selected candidate (local and remote) is included.

Metric	Description	
<b>address</b>	This indicates the IP address of the local and remote candidate.	
<b>port</b>	Port number of the candidate	
<b>protocol</b>	Protocol used to obtain the candidate. The valid values are UDP/TCP.	
<b>candidateType</b>	Type of candidate selected - host, srflx or relay.	
<b>priority</b>	Priority of the selected local and remote candidate.	
<b>url</b>	Source of the selected local candidate. This gives an indication of if the candidate selected is received from a STUN server or TURN server.	
<b>relayProtocol</b>	If TURN server is used to obtain the selected local candidate, this field indicates what protocol was used to obtain it. Valid values are TCP/UDP.	

ICE Candidate Pair Stats: Only the information about the selected candidate pairs is included.

Metric	Description	
<b>localCandidateId</b>	The ID of the selected local candidate in the pair.	

Metric	Description
<b>remoteCandidateId</b>	The ID of the selected remote candidate in the pair.
<b>state</b>	State of the candidate pair being inspected.
<b>nominated</b>	Set to TRUE since the stats are being pulled for selected candidate pair.
<b>packetsSent</b>	Number of packets sent. This is calculated in the <code>writeFrame()</code> call in the <code>writeFrame()</code> call. This information can also be extracted from outgoing RTP Stats, but since Ice candidate pair includes a <code>lastPacketSent</code> timestamp, it might be useful to calculate number of packets sent between two points in time.
<b>packetsReceived</b>	This is updated every time the <code>incomingDataHandler</code> is called.
<b>bytesSent</b>	This is calculated in the <code>iceAgentSendPacket()</code> in the <code>writeFrame()</code> call. This is useful when calculating a bit rate. Currently, this also includes the header and padding since the ICE layer is oblivious to the RTP packet format.

Metric	Description	
<b>bytesReceived</b>	This is updated every time the <code>incomingDataHandler</code> is called. Currently, this also includes the header and padding since the ICE layer is oblivious to the RTP packet format.	
<b>lastPacketSentTimestamp</b>	This is updated every time a packet is sent. This can be used in conjunction with the <code>packetsSent</code> and a recorded start time in application to current packet transfer rate.	
<b>lastPacketReceivedTimestamp</b>	This is updated on receiving data in <code>incomingDataHandler()</code> . This can be used in conjunction with <code>packetsReceived</code> to deduce the current packet receive rate. The start time has to be recorded at the application layer in the <code>transceiverOnFrame()</code> callback.	



Metric	Description	
<b>firstRequestTimestamp</b>	Recorded when the very first STUN binding request is sent out successfully in <code>iceAgentSendStunPacket()</code> . This can be used along with <code>lastRequestTimestamp</code> and <code>requestsSent</code> to find average time between STUN binding requests.	
<b>lastRequestTimestamp</b>	Recorded every time a STUN binding request is sent out successfully in <code>iceAgentSendStunPacket()</code> .	
<b>lastResponseTimestamp</b>	Recorded every time a STUN binding response is received.	
<b>totalRoundTripTime</b>	Updated when a binding response is received for a request. The request and response are mapped in a hash table based on checksum.	
<b>currentRoundTripTime</b>	Most recent round trip time updated when a binding response is received for a request on the candidate pair.	
<b>requestsReceived</b>	A counter that is updated on every STUN binding request received.	

Metric	Description	
<b>requestsSent</b>	A counter that is updated on every STUN binding request sent out in <code>iceAgentS.endStunPacket()</code> .	
<b>responsesSent</b>	A counter that is updated on every STUN binding response sent out in response to a binding request in <code>handleStunPacket()</code> .	
<b>responsesReceived</b>	A counter that is updated on every STUN binding response received in <code>handleStunPacket()</code> .	
<b>packetsDiscardedOnSend</b>	Updated when packet sending fails. In other words, this is updated when <code>iceUtilsSendData()</code> fails. This is useful to determine percentage of packets dropped in a specific duration.	
<b>bytesDiscardedOnSend</b>	Updated when packet sending fails. In other words, this is updated when <code>iceUtilsSendData()</code> fails. This is useful when determining percentage of packets dropped in a specific duration. Note that the counter also includes the header of the packets.	

## Media

### Outbound RTP Stats

Metric	Description
<b>voiceActivityFlag</b>	This is currently part of <code>RtcEncoderStats</code> defined in <code>Include.h</code> . The flag is set to <code>TRUE</code> if the last audio packet contained voice. The flag is currently not set in the samples.
<b>packetsSent</b>	This indicates the total number of RTP packets sent out for the selected SSRC. This is a part of <a href="https://www.w3.org/TR/webrtc-stats/#sentrtptime-dict">https://www.w3.org/TR/webrtc-stats/#sentrtptime-dict</a> and is included as part of outbound stats. This is incremented every time <code>writeFrame()</code> is called.
<b>bytesSent</b>	Total number of bytes excluding RTP header and padding that is sent. This is updated on every <code>writeFrame</code> call.
<b>encoderImplementation</b>	This is updated by the application layer as part of <code>RtcEncoderStats</code> object.
<b>packetsDiscardedOnSend</b>	This field is updated if the ICE agent fails to send the encrypted RTP packet for any

Metric	Description	
	reason in the <code>iceAgentSendPacket</code> call.	
<b>bytesDiscardedOnSend</b>	This field is also updated if the ICE agent fails to send the encrypted RTP packet for any reason in the <code>iceAgentSendPacket</code> call.	
<b>framesSent</b>	This is incremented only if media stream track type is <code>MEDIA_STREAM_TRACK_KIND_VIDEO</code> .	
<b>hugeFramesSent</b>	This counter is updated for frames that are 2.5 times the average size of frames. The size of the frame is obtained by calculating the fps (based on the last known frame count time and number of frames encoded in a time interval) and using the <code>targetBitrate</code> in <code>RtcEncoderStats</code> set by the application.	
<b>framesEncoded</b>	This counter is updated only for video track after successful encoding of the frame. It is updated on every <code>writeFrame</code> call.	

Metric	Description	
<b>keyFramesEncoded</b>	This counter is updated only for video track after successful encoding of the key frame. It is updated on every writeFrame call.	
<b>framesDiscardedOnSend</b>	This is updated when frame sending fails due to iceAgentSendPacket call failure. A frame comprises of a group of packets and currently, framesDiscardedOnSend fails if any packet gets discarded on while sending because of an error.	
<b>frameWidth</b>	This ideally represents the frame width of the last encoded frame. Currently , this is set to a value by the application as part of RtcEncoderStats* *and is of not much significance.	
<b>frameHeight</b>	This ideally represents the frame height of the last encoded frame. Currently , this is set to a value by the application as part of RtcEncoderStats and is of not much significance.	

Metric	Description	
<b>frameBitDepth</b>	This represents the bit depth per pixel width of the last encoded frame. Currently, this is set by the application as part of <code>RtcEncoderStats</code> and translated into outbound stats.	
<b>nackCount</b>	This value is updated every time a NACK is received on an RTP packet and a re-attempt to send the packet is made. The stack supports re-transmission of packets on receiving a NACK.	
<b>firCount</b>	The value is updated on receiving a FIR packet ( <code>onRtcpPacket-&gt;onRtcpFIRPacket</code> ). It indicates how often the stream falls behind and has to skip frames in order to catch up. FIR packet is currently not decoded to extract the fields, so, even though the count is set, no action is taken.	
<b>pliCount</b>	The value is updated on receiving a PLI packet ( <code>onRtcpPacket-&gt;onRtcpPLIPacket</code> ). It indicates that some amount of encoded video data has been lost for one or more frames.	

Metric	Description	
<b>sliCount</b>	The value is updated on receiving a SLI packet (onRtcpPacket->onRtcpSLIPacket). It indicates how often packet loss affects a single frame.	
<b>qualityLimitationResolutionChanges</b>	Currently, the stack supports this metric, however, the frame width and height are not monitored for every encoded frame.	
<b>lastPacketSentTimestamp</b>	The timestamp at which the last packet was sent. It is updated on every writeFrame call.	
<b>headerBytesSent</b>	Total number of RTP header and padding bytes sent for this SSRC excluding the actual RTP payload.	
<b>bytesDiscardedOnSend</b>	This is updated when frame sending fails due to iceAgentSendPacket call failure. A frame comprises of a group of packets, which in turn comprises of bytes and currently, bytesDiscardedOnSend fails if any packet gets discarded on while sending because of an error.	

Metric	Description	
<b>retransmittedPacketsSent</b>	The number of packets that are retransmitted on reception of PLI/SLI/NACK. Currently, the stack only counts the packet resent of NACK since PLI and SLI based retransmissions are not supported.	
<b>retransmittedBytesSent</b>	The number of bytes that are retransmitted on reception of PLI/SLI/NACK. Currently, the stack only counts the bytes resent of NACK since PLI and SLI based retransmissions are not supported.	
<b>targetBitrate</b>	This is set in the application level.	
<b>totalEncodedBytesTarget</b>	This is increased by the target frame size in bytes every time a frame is encoded. This is updated using size parameter in Frame structure.	
<b>framesPerSecond</b>	This is calculated based on the time recorded for the last known encoded frame and the number of frames sent within a second.	
<b>totalEncodeTime</b>	This is set to an arbitrary value in the application and is translated to outbound stats internally.	



Metric	Description	
<b>totalPacketSendDelay</b>	This is currently set to 0 since iceAgentSendPacket sends packet immediately.	

## Remote inbound RTP Stats:

Metric	Description	
<b>roundTripTime</b>	The value is extracted from the RTCP receiver report on receiving an RTCP packet type 201 (receiver report). The report comprises of details such as last sender report and delay since last sender report to calculate round trip time. Sender reports are generated roughly every 200 milliseconds comprising of information such as number of packets sent and bytes sent that are extracted from outbound stats.	
<b>totalRoundTripTime</b>	Sum of round trip times calculated	
<b>fractionLost</b>	Represents the fraction of RTP packets lost for the SSRC since the previous sender/receiver reportfractionLost was sent.	

Metric	Description
<b>reportsReceived</b>	Updated every time a receiver report type packet is received.
<b>roundTripTimeMeasurements</b>	Indicates the total number of reports received for the SSRC that contains valid round trip time. However, currently this value is incremented regardless so its meaning is the same as reportsReceived.

#### Inbound RTP Stats:

Metric	Description
<b>packetsReceived</b>	The counter is updated when a packet is received for a specific SSRC.
<b>jitter</b>	This metric indicates the packet Jitter measured in seconds for the specific SSRC.
<b>jitterBufferDelay</b>	This metric denotes the sum of time spent by each packet in the jitter buffer.
<b>jitterBufferEmittedCount</b>	The total number of audio samples or video frames that have come out of the jitter buffer.
<b>packetsDiscarded</b>	The counter is updated when the Jitter buffer is full and the packet cannot be pushed

Metric	Description	
	into it. This can be used to calculate percentage of packets discarded in a fixed duration.	
<b>framesDropped</b>	This value is updated when the <code>onFrameDroppedFunc()</code> is invoked.	
<b>lastPacketReceived Timestamp</b>	Represents the timestamp at which the last packet was received for this SSRC.	
<b>headerBytesReceived</b>	The counter is updated on receiving an RTP packet.	
<b>bytesReceived</b>	Number of bytes received. This does not include the header bytes. This metric can be used to calculate the incoming bit rate.	
<b>packetsFailedDecryption</b>	This is incremented when the decryption of the SRTP packet fails.	

## Data Channel

### Data Channel Metrics

Metric	Description	
<b>label</b>	Label is the name of the data channel being inspected.	

Metric	Description	
<b>protocol</b>	Since our stack uses SCTP, the protocol is set to a constant SCTP.	
<b>dataChannelIdentifier</b>	The even or odd identifier used to uniquely identify a data channel. This is updated to an odd value if the SDK is the offerer and even value if SDK is the answerer.	
<b>state</b>	State of the data channel when the stats are queried. Currently, the two states supported are RTC_DATA_CHANNEL_STATE_CONNECTING (when the channel is created) and RTC_DATA_CHANNEL_STATE_OPEN (Set in the onOpen() event).	
<b>messagesSent</b>	The counter is updated when the SDK sends messages over the data channel.	
<b>bytesSent</b>	The counter is updated with the bytes in the message that is sent out. This can be used to understand how many bytes are not sent due to failure, i.e, to understand the percentage of bytes that are sent.	
<b>messagesReceived</b>	The metric is incremented in the onMessage() callback.	

Metric	Description	
<b>bytesReceived</b>	The metric is generated in the <code>onMessage()</code> callback.	

# Security

Cloud security at Amazon is the highest priority. As an Amazon customer, you will benefit from a data center and network architecture built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between Amazon and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – Amazon is responsible for protecting the infrastructure that runs Amazon services in the Amazon Cloud. Amazon also provides you with services that you can use securely. The effectiveness of our security is regularly tested and verified by third-party auditors as part of the [Amazon compliance programs](#). To learn about the compliance programs that apply to Kinesis Video Streams, see [Amazon Services in Scope by Compliance Program](#).
- **Security in the cloud** – Your responsibility is determined by the Amazon service that you use. You are also responsible for other factors including the sensitivity of your data, your organization's requirements, and applicable laws and regulations.

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

## Topics

- [Controlling Access to Kinesis Video Streams with WebRTC Resources Using Amazon Identity and Access Management](#)
- [Compliance validation for Amazon Kinesis Video Streams with WebRTC](#)
- [Resilience in Kinesis Video Streams with WebRTC](#)
- [Infrastructure Security in Kinesis Video Streams with WebRTC](#)
- [Security Best Practices for Kinesis Video Streams with WebRTC](#)

# Controlling Access to Kinesis Video Streams with WebRTC Resources Using Amazon Identity and Access Management

By using Amazon Identity and Access Management (IAM) with Amazon Kinesis Video Streams with WebRTC, you can control whether users in your organization can perform a task using specific Kinesis Video Streams with WebRTC API operations and whether they can use specific Amazon resources.

For more information about IAM, see the following:

- [Amazon Identity and Access Management \(IAM\)](#)
- [Getting started](#)
- [IAM User Guide](#)

## Contents

- [Policy Syntax](#)
- [Actions for Kinesis Video Streams with WebRTC](#)
- [Amazon Resource Names \(ARNs\) for Kinesis Video Streams](#)
- [Granting Other IAM Accounts Access to a Kinesis Video Stream](#)
- [Example Policies for Kinesis Video Streams with WebRTC](#)

## Policy Syntax

An IAM policy is a JSON document that consists of one or more statements. Each statement is structured as follows:

```
{
  "Statement": [{
    "Effect": "effect",
    "Action": "action",
    "Resource": "arn",
    "Condition": {
      "condition": {
        "key": "value"
      }
    }
  ]
}
```

```
    }  
  ]  
}
```

There are various elements that make up a statement:

- **Effect:** The *effect* can be Allow or Deny. By default, IAM users don't have permission to use resources and API actions, so all requests are denied. An explicit allow overrides the default. An explicit deny overrides any allows.
- **Action:** The *action* is the specific API action for which you are granting or denying permission.
- **Resource:** The resource that's affected by the action. To specify a resource in the statement, you need to use its Amazon Resource Name (ARN).
- **Condition:** Conditions are optional. They can be used to control when your policy is in effect.

As you create and manage IAM policies, you might want to use the [IAM Policy Generator](#) and the [IAM Policy Simulator](#).

## Actions for Kinesis Video Streams with WebRTC

In an IAM policy statement, you can specify any API action from any service that supports IAM. For Kinesis Video Streams with WebRTC, use the following prefix with the name of the API action: `kinesisvideo:.` For example: `kinesisvideo:CreateSignalingChannel`, `kinesisvideo:ListSignalingChannels`, and `kinesisvideo:DescribeSignalingChannel`.

To specify multiple actions in a single statement, separate them with commas as follows:

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

You can also specify multiple actions using wildcards. For example, you can specify all actions whose name begins with the word "Get" as follows:

```
"Action": "kinesisvideo:Get*"
```

To specify all Kinesis Video Streams operations, use the asterisk (\*) wildcard as follows:

```
"Action": "kinesisvideo:*"
```



For the complete list of Kinesis Video Streams API actions, see the [Kinesis Video Streams API reference](#).

## Amazon Resource Names (ARNs) for Kinesis Video Streams

Each IAM policy statement applies to the resources that you specify using their ARNs.

Use the following ARN resource format for Kinesis Video Streams:

```
arn:aws:kinesisvideo:region:account-id:channel/channel-name/code
```

For example:

```
"Resource": arn:aws:kinesisvideo::*:111122223333:channel/my-channel/0123456789012
```

You can get the ARN of a channel using [DescribeSignalingChannel](#).

## Granting Other IAM Accounts Access to a Kinesis Video Stream

You might need to grant permission to other IAM accounts to perform operations on Kinesis Video Streams with WebRTC signaling channels. 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 [Creating a role to delegate permissions to an Amazon Web Service](#) in the *IAM User Guide*.

## Example Policies for Kinesis Video Streams with WebRTC

The following example policies demonstrate how you can control user access to your Kinesis Video Streams with WebRTC channels.

### Example 1: Allow users to get data from any signaling channel

This policy allows a user or group to perform the `DescribeSignalingChannel`, `GetSignalingChannelEndpoint`, `ListSignalingChannels`, and `ListTagsForResource` operations on any signaling channel.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
```

```

        "Action": [
            "kinesisvideo:Describe*",
            "kinesisvideo:Get*",
            "kinesisvideo:List*"
        ],
        "Resource": "*"
    }
]
}

```

### Example 2: Allow a user to create a signaling channel

This policy allows a user or group to perform the `CreateSignalingChannel` operation.

```

{
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kinesisvideo:CreateSignalingChannel"
      ],
      "Resource": "*"
    }
  ]
}

```

### Example 3: Allow a user full access to all Kinesis Video Streams and Kinesis Video Streams with WebRTC resources

This policy allows a user or group to perform any Kinesis Video Streams operation on any resource. This policy is appropriate for administrators.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "kinesisvideo:*",
      "Resource": "*"
    }
  ]
}

```

## Example 4: Allow a user to get data from a specific signaling channel

This policy allows a user or group to get data from a specific signaling channel.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "kinesisvideo:DescribeSignalingChannel",
      "Resource": "arn:aws:kinesisvideo:us-west-2:123456789012:channel/
channel_name/0123456789012"
    }
  ]
}
```

## Compliance validation for Amazon Kinesis Video Streams with WebRTC

To learn whether an Amazon Web Service is within the scope of specific compliance programs, see [Amazon Web Services in Scope by Compliance Program](#) and choose the compliance program that you are interested in. For general information, see [Amazon Web Services Compliance Programs](#).

You can download third-party audit reports using Amazon Artifact. For more information, see [Downloading Reports in Amazon Artifact](#).

Your compliance responsibility when using Amazon Web Services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. Amazon provides the following resources to help with compliance:

- [Security and Compliance Quick Start Guides](#) – These deployment guides discuss architectural considerations and provide steps for deploying baseline environments on Amazon that are security and compliance focused.
- [Amazon Compliance Resources](#) – This collection of workbooks and guides might apply to your industry and location.
- [Evaluating Resources with Rules](#) in the *Amazon Config Developer Guide* – The Amazon Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.

- [Amazon Security Hub](#) – This Amazon Web Service provides a comprehensive view of your security state within Amazon. Security Hub uses security controls to evaluate your Amazon resources and to check your compliance against security industry standards and best practices. For a list of supported services and controls, see [Security Hub controls reference](#).

## Resilience in Kinesis Video Streams with WebRTC

The Amazon global infrastructure is built around Amazon Web Services Regions and Availability Zones. Amazon Web Services Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. You can use Availability Zones to design and operate applications and databases that automatically fail over between Availability Zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about Amazon Web Services Regions and Availability Zones, see [Amazon Global Infrastructure](#).

## Infrastructure Security in Kinesis Video Streams with WebRTC

As a managed service, Kinesis Video Streams (including its WebRTC capability) is protected by the Amazon global network security procedures that are described in the [Amazon Web Services: Overview of Security Processes](#) whitepaper.

You use Amazon published API calls to access Kinesis Video Streams through the network. Clients must support Transport Layer Security (TLS) 1.2 or later. We recommend TLS 1.3 or later. Clients must also support cipher suites with perfect forward secrecy (PFS) such as Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Ephemeral Diffie-Hellman (ECDHE). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the [Amazon Security Token Service](#) (Amazon STS) to generate temporary security credentials to sign requests.

## Security Best Practices for Kinesis Video Streams with WebRTC

Amazon Kinesis Video Streams (including its WebRTC capability) 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.

For security best practices for your remote devices, see [Security Best Practices for Device Agents](#).

## Implement least privilege access

When granting permissions, you decide who is getting what permissions to which Kinesis Video Streams resources. You enable specific actions that you want to allow on those resources. Therefore you should grant only the permissions that are required to perform a task. Implementing least privilege access is fundamental in reducing security risk and the impact that could result from errors or malicious intent.

For example, a producer that sends data to Kinesis Video Streams requires only `PutMedia`, `GetStreamingEndpoint`, and `DescribeStream`. Do not grant producer applications permissions for all actions (\*), or for other actions such as `GetMedia`.

For more information, see [What Is Least Privilege & Why Do You Need It?](#)

## Use IAM roles

Producer and client applications must have valid credentials to access Kinesis video streams. You should not store Amazon credentials directly in a client application or in an Amazon S3 bucket. These are long-term credentials that are not automatically rotated and could have a significant business impact if they are compromised.

Instead, you should use an IAM role to manage temporary credentials for your producer and client applications to access Kinesis video streams. When you use a role, you don't have to use long-term credentials to access other resources.

For more information, see the following topics in the *IAM User Guide*:

- [IAM Roles](#)
- [Common Scenarios for Roles: Users, Applications, and Services](#)

## Use CloudTrail to Monitor API Calls

Kinesis Video Streams with WebRTC is integrated with Amazon CloudTrail, a service that provides a record of actions taken by a user, role, or an Amazon service in Kinesis Video Streams with WebRTC.

Using the information collected by CloudTrail, you can determine the request that was made to Kinesis Video Streams with WebRTC, the IP address from which the request was made, who made the request, when it was made, and additional details.

For more information, see [the section called “Logging Kinesis Video Streams with WebRTC API Calls with Amazon CloudTrail”](#).

# Monitoring

Monitoring is an important part of maintaining the reliability, availability, and performance of Amazon Kinesis Video Streams with WebRTC and your Amazon solutions. You should collect monitoring data from all of the parts of your Amazon solution so that you can more easily debug a multi-point failure if one occurs. Before you start monitoring Kinesis Video Streams with WebRTC, however, you should create a monitoring plan that includes answers to the following questions:

- What are your monitoring goals?
- What resources will you monitor?
- How often will you monitor these resources?
- What monitoring tools will you use?
- Who will perform the monitoring tasks?
- Who should be notified when something goes wrong?

After you have defined your monitoring goals and have created your monitoring plan, the next step is to establish a baseline for normal Kinesis Video Streams with WebRTC performance in your environment. You should measure Kinesis Video Streams with WebRTC performance at various times and under different load conditions. As you monitor Kinesis Video Streams with WebRTC, you should store a history of monitoring data that you've collected. You can compare current Kinesis Video Streams with WebRTC performance to this historical data to help you to identify normal performance patterns and performance anomalies, and devise methods to address issues that may arise.

## Topics

- [Monitoring Kinesis Video Streams with WebRTC Metrics with CloudWatch](#)
- [Logging Kinesis Video Streams with WebRTC API Calls with Amazon CloudTrail](#)

## Monitoring Kinesis Video Streams with WebRTC Metrics with CloudWatch

You can monitor a Kinesis Video Streams with WebRTC using Amazon CloudWatch, which collects and processes raw data from Kinesis Video Streams with WebRTC into readable, near real-time

metrics. These statistics are recorded for a period of 15 months, so that you can access historical information and gain a better perspective on how your web application or service is performing.

Kinesis Video Streams provides the following metrics:

## Topics

- [Signaling Metrics](#)
- [TURN Metrics](#)

## Signaling Metrics

Metric name	Dimensions	Unit	Description
<b>Failure</b>	Operation, Signaling ChannelName	Count	'0' is emitted if the Operation mentioned in dimension returns 200 status code response. '1' otherwise.
<b>Latency</b>	Operation, Signaling ChannelName	Milliseconds	The time measured from when the service receives the request until the service returns the response.
<b>MessagesTransferred.Count</b>	Signaling ChannelName	Count	Total number of messages transferred (sent and received) for a given channel.

The Operation dimension can apply to any of the following APIs:

- ConnectAsMaster
- ConnectAsViewer
- SendSdpOffer
- SendSdpAnswer
- SendCandidate
- SendAlexaOfferToMaster



- GetIceServerConfig
- Disconnect

## TURN Metrics

Metric name	Dimensions	Unit	Description
TURNConnectedMinutes	SignalingChannelNames	Count	'1' is emitted for each TURN allocation that is used to stream data through in a minute.

## Logging Kinesis Video Streams with WebRTC API Calls with Amazon CloudTrail

Amazon Kinesis Video Streams with WebRTC is integrated with Amazon CloudTrail, a service that provides a record of actions taken by a user, role, or an Amazon service in Amazon Kinesis Video Streams with WebRTC. CloudTrail captures all API calls for Amazon Kinesis Video Streams with WebRTC as events. The calls captured include calls from the Amazon Kinesis Video Streams console and code calls to the Amazon Kinesis Video Streams with WebRTC API operations. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Amazon Kinesis Video Streams with WebRTC. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine the request that was made to Amazon Kinesis Video Streams with WebRTC, the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, including how to configure and enable it, see the [Amazon CloudTrail User Guide](#).

## Amazon Kinesis Video Streams with WebRTC and CloudTrail

CloudTrail is enabled on your Amazon account when you create the account. When supported event activity occurs in Amazon Kinesis Video Streams with WebRTC, that activity is recorded in a CloudTrail event along with other Amazon service events in **Event history**. You can view, search,

and download recent events in your Amazon account. For more information, see [Viewing Events with CloudTrail Event History](#).

For an ongoing record of events in your Amazon account, including events for Amazon Kinesis Video Streams with WebRTC, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all Amazon Web Services Region. The trail logs events from all Regions in the Amazon partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other Amazon services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- [Overview for Creating a Trail](#)
- [CloudTrail Supported Services and Integrations](#)
- [Configuring Amazon SNS Notifications for CloudTrail](#)
- [Receiving CloudTrail Log Files from Multiple Regions](#) and [Receiving CloudTrail Log Files from Multiple Accounts](#)

Amazon Kinesis Video Streams with WebRTC supports logging the following actions as events in CloudTrail log files:

- [CreateSignalingChannel](#)
- [DeleteSignalingChannel](#)
- [DescribeSignalingChannel](#)
- [GetSignalingChannelEndpoint](#)
- [ListSignalingChannels](#)
- [ListTagsForResource](#)
- [TagResource](#)
- [UntagResource](#)
- [UpdateSignalingChannel](#)

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or Amazon Identity and Access Management (IAM) user credentials.

- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another Amazon service.

For more information, see the [CloudTrail userIdentity Element](#).

## Example: Amazon Kinesis Video Streams with WebRTC Log File Entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the [CreateSignalingChannel](#) action.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "EX_PRINCIPAL_ID",
    "arn": "arn:aws:iam::123456789012:user/Alice",
    "accountId": "123456789012",
    "accessKeyId": "EXAMPLE_KEY_ID",
    "userName": "Alice"
  },
  "eventTime": "2019-11-19T22:49:04Z",
  "eventSource": "kinesisvideo.amazonaws.com",
  "eventName": "CreateSignalingChannel",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "127.0.0.1",
  "userAgent": "aws-sdk-java/unknown-version Linux/x.xx",
  "requestParameters": {
    "channelName": "YourChannelName"
  },
  "responseElements": {
    "channelARN": "arn:aws:kinesisvideo:us-west-2:123456789012:channel/YourChannelName/1574203743620"
  },
  "requestID": "df3c99c4-1d97-49da-8569-7de6c92b4856",
  "eventID": "bb74bac2-964c-49b0-903a-3501c6bde632"
```

```
}
```

# Document History for the Amazon Kinesis Video Streams with WebRTC Developer Guide

Change	Description	Date
<a href="#">Initial publication</a>	Initial publication of the Amazon Kinesis Video Streams with WebRTC Developer Guide. <a href="#">Learn more</a>	November 4, 2019

# Amazon Glossary

For the latest Amazon terminology, see the [Amazon glossary](#) in the *Amazon Web Services Glossary Reference*.