



This PDF is generated from authoritative online content, and is provided for convenience only. This PDF cannot be used for legal purposes. For authoritative understanding of what is and is not supported, always use the online content. To copy code samples, always use the online content.

# Genesys Interaction Recording Solution Guide

Multiple data center locations

12/18/2025

---

## Contents

- 1 Multiple data center locations
  - 1.1 Recording storage locations
  - 1.2 Recording processing
  - 1.3 User interface
  - 1.4 Encryption management
  - 1.5 Media lifecycle management

# Multiple data center locations

Genesys Interaction Recording (GIR) can be deployed in multiple data center locations.

Before you can deploy GIR in multiple data center locations, you must consider the following:

- Recording storage locations
- Recording processing
- User interface
- Encryption management
- Media lifecycle management

## Recording storage locations

Recording files can be saved in separate storage locations. The location of the recording files is directly related to the location of the agent handling the recording, so that compliance or bandwidth requirements between data centers is met. To enable each data center to record an interaction, each data center must be deployed with its own Media Control Platform (MCP) instance. Consider the following examples:

- There are data centers throughout the country where the agents reside, and Genesys server components are located in those data centers. To save network bandwidth between the data centers, all recording files must be stored in the same data center in which the recordings occur. A recording file is only transported to another data center whenever a client requests it for media file playback.
- There are data centers in different countries where the agents handle interactions. There are compliance requirements to ensure that interactions arriving in a country are recorded and stored in that country's data center only, and interactions arriving in another country are recorded and stored in that country's data center only. Your deployment model has Interaction Recording Web Services nodes installed in data centers per country.

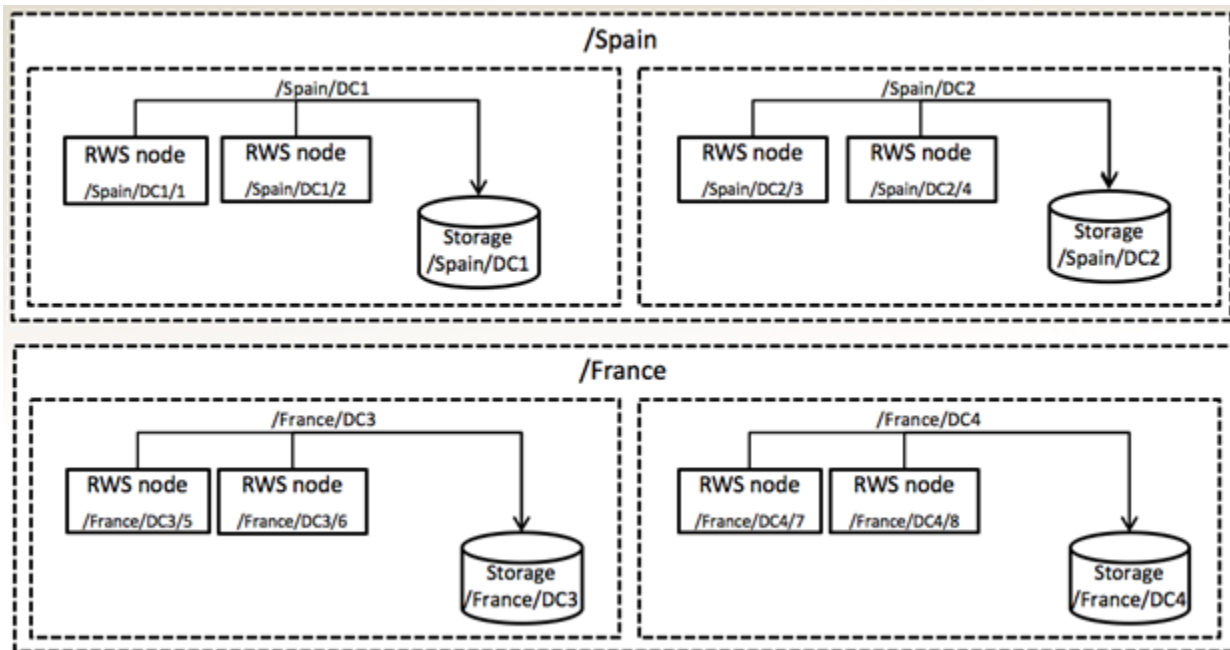
These examples require that Interaction Recording Web Services nodes be located in each data center with a local storage location. For those locations that do not have local storage configured, a centralized storage location can be used.

In the case of voice recordings, the storage location is determined based on the geo-location set in the SIP Server.

For screen recording, the mapping storage location based on the Interaction Recording Web Services nodes requires no provisioning of the location of the agent, since the connection of the Screen Recording Service to the Interaction Recording Web Services nodes represents the actual location in which the agent resides.

The following diagram is an example of how recording storage can be divided into multiple data centers. When studying this diagram consider the following:

- In a system where there is a data center in the same location as the SIP Server and GVP, recording and recording storage are located in the same data center. RWS and the Recording Crypto Server should also be deployed in all of these data centers.
- GIR recording can occur in any of these data centers. GIR operates in primary / backup across 2 such data centers. SpeechMiner and GIR components (such as Recording Processor Script or Voice Processor and Muxer), are deployed in these data centers.



## Configure voice recording storage

To configure the voice recording storage locations, refer to [Configuring the Storage Credentials for Interaction Recording Web Services > Enable Storage](#).

## Assign a recording IVR profile

A Recording IVR profile enables you to setup a separate voice recording storage location, per data center location, based on the SIP Server geo-location (see the content describing Geo-Location within the [SIP Server Deployment Guide](#)).

To setup a separate voice recording storage location, perform the steps in the [IVR Profile](#) section.

## Configure screen recording storage

To configure the screen recording storage locations, refer to [Configuration for Screen Recordings > Configuring the Storage Credentials for Interaction Recording Web Services > Step #4 Enable storage for a single or multiple locations](#).

## Recording processing

Each data center where recordings occur must be deployed with at least one Recording Processor Script or Voice Processor, to process recordings within the same data center. When the processing is complete, the recording metadata is written to the primary data center for indexing.

The Recording Processor Script or Voice Processor is deployed in the primary and backup data centers. SpeechMiner is also deployed in the primary and backup data centers, and these components are active in the primary data center while the backup data center only contains backup instances. Indexing of recording metadata also occurs in the primary data center.

For additional information, see [Deploying Recording Processor Script](#) or [Deploying Voice Processor](#).

The Recording Processor Script or Voice Processor should send recording metadata to the SpeechMiner Interaction Receiver in the primary data center according to the [IVR Profile configuration](#) above.

## Screen Recording processing

The Screen Recording Service (SRS) is always connected to Interaction Recording Web Services (RWS). This connection occurs from the agent's location to a data center in which the SIP Server and RWS are deployed. This connection manages screen recording and the upload of screen recordings to the recording storage.

The Recording Muxer Script processes screen recordings through the primary data center. Muxer operates in the primary/backup data center and all active instances operate in the primary data center.

The following instructions are applicable if you are using Workspace Web Edition (WWE), or Workspace Desktop Edition (WDE) (versions earlier than 8.5.118.10), or Screen Recording Service (SRS) (versions earlier than 8.5.340.97):

- If a data center fails and agents need to use the backup data center for continuity of operations, the SRS must be reconfigured to connect to the Interaction Recording Web Services nodes located in a backup data center. Perform one of the following procedures to reconfigure the SRS for this purpose. Once the SRS is configured to connect to that remote data center, the recording files will be transported to the remote data center and saved locally at that data center.
  - Configure the agent desktop to re-submit the data center server address to the SRS, specifying the backup data center address.
  - Update DNS mapping of RWS to specify the backup data center.
  - Use an HTTP Load Balancer to redirect HTTP requests from the SRS to the backup data center where RWS is located.

The following instructions are applicable only if you are using WDE (version 8.5.118.10 or later) and SRS (version 8.5.340.97 or later):

- If a data center fails and agents need to use a backup data center for continuity of operations, SRS must be configured to connect to the Interaction Recording Web Services nodes located in a backup data center. Perform the following procedure to configure SRS for this purpose. Once SRS is configured to connect to a remote data center, SRS will automatically failover to RWS nodes in the backup data center from the primary data center and vice versa.

- Specify a value for the [interaction-workspace] **screen-recording.htcc.peer\_uri** option in the WDE application object. For more information, refer to [Screen Recording Disaster Recovery](#).

Similar concepts apply if a custom agent desktop is used with integration through the SRS API as described in [Client Login API](#).

For eServices, RWS in both data centers connects to the Interaction Server in the primary data center. If a data center fails, or if the Interaction Server in a data center fails, the operator needs to make an HTTP POST request to RWS explicitly to switch over the Interaction Server in the primary data center to the backup data center for continuity of operations.

The following steps describe how to switch the Interaction Server being used by RWS between data centers:

1. Determine the genesys-environment ID using the following command with the ops username and password (ops:ops):

```
curl -u ops:ops http://<Interaction Recording Web Services Server>:<Interaction Recording Web Services port>/api/v2/ops/genesys-environments; echo
```

Interaction Recording Web Services returns the following output:

```
{
  "statusCode": 0,
  "paths": [
    "/genesys-environments/<genesys-environment-ID>"
  ],
  "uris": [
    "http://<Interaction Recording Web Services Server>:<Interaction Recording Web Services port>/api/v2/ops/genesys-environments/<genesys-environment-ID>"
  ]
}
```

2. Using a text editor, create a new file called **multimedia\_switch** with the following content:

```
{
  "eServicesSite": "DC2" // set to site name as per application parameter siteName
}
```

3. Execute the following command:

```
curl -u ops:ops -X POST -d @multimedia_switch http://<Interaction Recording Web Services Server>:<Interaction Recording Web Services Port>/api/v2/ops/genesys-environments/{genesys-environment-ID}/multimedia-switch --header "Content-Type: application/json"; echo
```

4. Verify that WDE is connected to same data-center:
  - a. Specify value for the option **[interaction-workspace] "disaster-recovery.eservices-site"** in the WDE application object. See: [Administrator operations in case of data center failover](#)
  - b. After the primary data center Interaction Server is brought back to working condition, make another POST to switch back to the Interaction Server in the primary data center.

### Important

RWS will monitor the **eServicesSite** value every **drMonitoringDelay** seconds. These

values are configured in the **serverSettings** section in the **application.yaml** file. If there is a change in the **eServicesSite** value during monitoring, RWS will perform a switch over to the specified **eServicesSite**. For a definition of **drMonitoringDelay** refer to the [Configuration Options](#) page.

## User interface

SpeechMiner works in active/backup mode across data centers.

To enable recordings from multiple data centers to be viewed from a single SpeechMiner user interface, assign a primary data center where SpeechMiner hosts the primary database and index.

## Data center index update

Ensure that the indexes are up-to-date in all data centers.

**SpeechMiner Index:** Configure the **Daily Index Backup** and write the backup to the backup data center. Whenever the primary data center becomes unavailable, the backup index can be configured to be the primary instance in the backup data center. For additional information, see: [SMConfig > Index](#).

**Elasticsearch Index:** In the backup data center, write a periodic task (for example, cron job) that will execute the following re-indexing command using HTTP to RWS. It is recommended that you perform this task daily.

```
POST http://<Interaction Recording Web Services Server>:<Interaction Recording Web Services Port>/api/v2/ops/contact-centers/<contact center ID (in hex format)>/recordings
```

If you are using screen recording you must also perform a re-indexing command for the screen recording index.

```
POST http://<Interaction Recording Web Services Server>:<Interaction Recording Web Services Port>/api/v2/ops/contact-centers/<contact center ID (in hex format)>/screen-recordings
```

For additional information about how to obtain the contact center ID, refer to [Configuring Features](#).

For example:

1. Using a text editor, create a new file called `update_index` with the following content:

```
{
  "operationName": "forceIndex",
  "from": "1369272257713",
  "to": "1369275857713",
```

```
    "purgeOld":false
}
```

2. Execute the following command:

```
{
  curl -u ops:ops -X POST -d @update_index
  http://<Interaction Recording Web Services Server>:<Interaction Recording Web Services
  Port>/api/v2/ops/contact-centers/<contact center ID (in hex format)>/recordings
  --header "Content-Type: application/json"; echo
}
```

3. If you are using screen recordings, you must also execute the following command:

```
{
  curl -u ops:ops -X POST -d @update_index
  http://<Interaction Recording Web Services Server>:<Interaction Recording Web Services
  Port>/api/v2/ops/contact-centers/<contact center ID (in hex format)>/screen-recordings
  --header "Content-Type: application/json"; echo
}
```

Attribute	Type	Mandatory	Description
operationName	String	yes	forceIndex
from	long	yes	Milliseconds since epoch. The starting timestamp from which recordings are indexed and rounded to the nearest hour.
to	long	no	Milliseconds since epoch. The ending timestamp until recordings are indexed and rounded to the nearest hour.
purgeOld	boolean	no	The flag that specifies if the old index should be deleted prior to re-indexing. This is necessary when the updated RWS version needs an index with an updated schema. Default value = false

Due to the rounding of time units in the API, the re-indexing command interval is a minimum of one hour.

## Encryption management

When multiple data centers are deployed, you must ensure that the Recording Crypto Server (RCS) **keystore** file is replicated to all the RCS instances in all the data centers.

To learn how to replicate the RCS keystore file, refer to the [Deploying Recording Crypto Server](#) page.

## Media lifecycle management

Media Lifecycle Management is location specific and can be configured separately for each data center. Refer to [Configuring Media Lifecycle Management](#) to configure multiple regions.