



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.

Web Services and Applications Deployment Guide

Multiple Data Center Deployment

Contents

- 1 Multiple Data Center Deployment
 - 1.1 Overview
 - 1.2 Architecture
 - 1.3 Incoming traffic distribution
 - 1.4 Business continuity (smart failover)
 - 1.5 Configuration
 - 1.6 Limitations

Multiple Data Center Deployment

Genesys Web Services 8.6 supports deployment with multiple (two or more) data centers. This section describes this type of deployment.

Overview

Genesys Web Services 8.6 supports deployments in two and three data centers.


- Typically one data center is designated as the "primary". This is the data center where the primary read/write instances of the Configuration Database and GWS Database (if used/deployed).
- The other data center(s) will each have a deployment that mirrors the primary data center in terms of the GWS Nodes, Redis cluster, and Elasticsearch. Note that Redis and Elasticsearch operate independently in each data center.
- GWS in the primary data center should always be active (except for maintenance or disaster scenarios). GWS in the secondary data center(s) can be active or standby.

Architecture

A typical GWS data center in a multiple data center deployment consists of the following components:

- 2+ GWS API Service nodes
- 2+ GWS Platform Service nodes
- 2 GWS Stat Index nodes running in primary/secondary mode
- 2 GWS Configuration Index nodes running in primary/secondary mode
- 3+ Redis Cluster primary nodes
- 3+ Elasticsearch nodes

The following diagram illustrates the architecture of this sample multiple data center deployment.



Medium deployments may also be deployed across multiple data centers.

Note the following restrictions of this architecture:

- Each data center must have a dedicated list of Genesys servers, such as Configuration Servers, Stat Servers, and T-Servers.
- Each GWS data center must have its own standalone and dedicated Elasticsearch Cluster.
- Each GWS data center must have its own standalone and dedicated Redis Cluster.
- The GWS node identity must be unique across the entire Cluster.

Incoming traffic distribution

GWS relies on a third-party reverse proxy to enable traffic distribution both within and across data centers. The reverse proxy provides session stickiness based on association with sessions on the backend server; this rule is commonly referred to as *Application-Controlled Session Stickiness*. In other words, when a GWS node creates a new session and returns Set-Cookie in response, the load-balancer issues its own stickiness cookie. GWS uses a JSESSIONID cookie by default, but this can be reconfigured by using the following option in the **application.yaml** file:

```
jetty:
  cookies:
    name: <HTTP Session Cookie Name>
```

Business continuity (smart failover)

Depending on the components that are functional/non-functional in a given data center (primary or secondary), the behavior will be different when the agent logs back into the other data center. All these behaviors are dependent on the agent via their desktop application (WWE or custom application) re-logging into the other data center. This re-login can be handled in two different ways:

1. The agent receives the appropriate error from the application and terminates that instance of the application, start a new instance and re-logs in the application but connected to the GWS cluster in the other data center.
2. The desktop application detects the error situation and then automatically navigates the agent through the necessary step to re-log them into the other data center. In WWE, this process is called smart failover.

The following are the common conditions that could cause an agent to re-log into another data center:

Data Center condition	What capabilities are available when logged into the other data center
GWS/WWE is down in either one of the data centers.	Agent stays in current state. The status of active interactions in a failing data center is as follows: • Voice - All existing interactions are active. • Digital - All existing interactions are active. The Contact History view resets to the Default view.
Primary data center is down (for example, due to a data center network issue).	Agent will go into ready state and is ready for new interactions (all channels). Following is the status of the interactions that were active in the primary data center: • Voice - All existing interactions are dropped. • Digital - All existing interactions are dropped. The Contact History view resets to the Default view.
Secondary data center is down (for example, due to a data center network issue).	Agent will go into ready state and is ready for new voice interactions. Following is the status of the interactions that were active in the primary data center: • Voice - All existing interactions are dropped. • Digital - All existing interactions are active. The Contact History view resets to the Default view.
Voice Service is down in primary or secondary data center.	Agent will go into ready state and is ready for new voice interactions. Following is the status of the interactions that were active in the primary data center: • Voice - All existing interactions are dropped. • Digital - All existing interactions are active. The Contact History view resets to the Default view.
Digital Service is down in the primary region	Agent changes state based on the loss of the digital/social interactions. Following is the status of the interactions that were active in the primary data center: • Voice - All existing interactions are active.


- **Digital** - Existing digital /social interactions are dropped.

The **Contact History** view resets to the Default view.

Configuration

This section describes the additional configuration required to set up a multiple data center deployment.

The topology of a GWS Cluster can be considered as a standard directory tree where a leaf node is a GWS data center. The following diagram shows a GWS Cluster with two geographical regions (US and EU), and three GWS data centers (East and West in the US region, and EU as its own data center).



Important

In a Disaster Recovery (DR) environment (for instance, having two separate SIP Servers connected to a WS Cluster Application with different locations specified), it is possible that the **Place** can contain two DNs of the type **Extension**. However, the two DNs must belong to different switches and the API server will be mapped only to one of them during the **StartContactCenterSession** operation. In such cases, the DN visibility depends on the **nodePath** parameter in the **application.yaml** file in the API server and on the **location** attribute of the SIP Server connection in the WS Cluster Application.

Genesys Web Services and Applications

The position of each node inside the GWS Cluster is specified by the mandatory property **nodePath** provided in **application.yaml**. The value of this property is in the standard file path format, and uses the forward slash (/) symbol as a delimiter. This property has the following syntax:

```
nodePath: <path-to-node-in-cluster>/<node-identity>
```

Where:

- `<path-to-node-in-cluster>` is the path inside the cluster with all logical sub-groups.
- `<node-identity>` is the unique identity of the node. Genesys recommends that you use the name of the host on which this data center is running for this parameter.

For example:

```
nodePath: /US/West/api-node-1
```

Each server will be aware of servers or nodes in the same path of above. In addition to the configuration options set in the standard deployment procedure, set the following configuration options in **application.yaml** for all GWS nodes to enable the multiple data center functionality:

```
serverSettings:  
  nodePath: <path-to-node-in-cluster>/<node-identity>
```

```
statistics:  
  locationAwareMonitoringDistribution: true  
  enableMultipleDataCenterMonitoring: true
```

- Set both **locationAwareMonitoringDistribution** and **enableMultipleDataCenterMonitoring** on all nodes. Without them, the statistics information in a Multiple Data Center environment will not be displayed properly.}}

In addition, set the following options on all Stat nodes:

```
serverSettings:  
  elasticSearchSettings:  
    enableScheduledIndexVerification: true  
    enableIndexVerificationAtStartup: true
```

GWS API Service node in the secondary Data Center might need to connect to the GWS API Service node in the primary Data Center if users are authenticated against the password stored in the configuration database. This connectivity enables changing the password by a user connected to the secondary Data Center.

The following should be configured in CME or GA:

```
Section "master-configuration-data-center" in Cloud Cluster application options should  
contain following options:  
location: <location of the primary region>  
redirectUri: <URI of primary region LB>
```

If the SQL DB is used for persistent data storage, it must be configured to replicate from the primary Data Center to the secondary Data Center database using the native replication capabilities of the database.

Additionally, the GWS API Service nodes in the secondary Data Center must be configured with a connection to the database in the primary Data Center, see <https://docs.genesys.com/Documentation/GWS/8.6.0DRAFT/Dep/ConfigurationPremise>.

Limitations

The cross-site monitoring is only supported via Team Communicator, using the WWE option `teamlead.monitoring-cross-site-based-on-activity-enabled`.

The My Agents view is not supported for cross-site operations.