



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 Callback Private Edition Guide

Before you begin

9/3/2025

Contents

- [1 Limitations and assumptions](#)
- [2 Download the Helm charts](#)
- [3 Third-party prerequisites](#)
- [4 Storage requirements](#)
 - [4.1 Hardware requirements](#)
 - [4.2 Sizing calculator](#)
- [5 Network requirements](#)
 - [5.1 Connection topology](#)
 - [5.2 Web application firewall rules](#)
- [6 Browser requirements](#)
- [7 Genesys dependencies](#)
- [8 GDPR support](#)

Find out what to do before deploying Genesys Callback.

Related documentation:

-
-
-
-

RSS:

- [For private edition](#)

Limitations and assumptions

Not applicable

Download the Helm charts

Genesys Engagement Service (GES) is the only service that runs in the GES Docker container. The Helm charts included with the GES release provision GES and any Kubernetes infrastructure necessary for GES to run, such as load balancing, autoscaling, ingress control, and monitoring integration.

For information about how to download the Helm charts, see [Downloading your Genesys Multicloud CX containers](#).

See [Helm charts and containers for Callback](#) for the Helm chart version you must download for your release.

Third-party prerequisites

For information about setting up your Genesys Multicloud CX private edition platform, see [Software requirements](#).

Third-party services

Name	Version	Purpose	Notes
Redis	6.x	Used for caching. Only distributions of Redis that support Redis	You require the following information for GES deployment:

Name	Version	Purpose	Notes
		cluster mode are supported, however, some services may not support cluster mode.	<ul style="list-style-type: none">• Redis hostname• Redis password (only required if you use a password)• Redis port <p>You require access to the Voice Microservices deployment Redis as well (specifically the Voice Orchestration Service Redis stream). You require the following information about the Voice Microservices Redis for your GES deployment:</p> <ul style="list-style-type: none">• Redis hostname• Redis password (only required if you use a password)• Redis port
PostgreSQL	11.x	Relational database.	<p>Before you deploy Genesys Engagement Service (GES), a table and user with database owner (DBO) permissions must be deployed and ready for use by GES. You require the following information for GES deployment:</p> <ul style="list-style-type: none">• Database name• Database port• Database hostname• Database user
Consul	1.13.x	Service discovery, service mesh, and key/value store.	

Storage requirements

The primary contributor to the size of a callback record is the amount of user data that is attached to a callback. Since this is an open-ended field, and the composition will differ from customer to customer, it is difficult to state the precise storage requirements of GES for a given deployment. To

assist you, the following table lists the results of testing done in an internal Genesys development environment and shows the impact that user data has when it comes to the storage requirements for both Redis and Postgres.

Test	Redis size	Postgres size (MB)
10,000 Scheduled Callbacks with no user data	26.51 MB	41.1 MB
10,000 Scheduled Callbacks with 10 KB of user data	64.44 MB	252.91 MB
10,000 Scheduled Callbacks with 100k of user data	110.58 MB	595.79 MB

Note: This is 100k of randomized string in a single field in the user data.

Hardware requirements

Genesys strongly recommends the following hardware requirements to run GES with a single tenant. The requirements are based on running GES in a multi-tenanted environment and scaled down accordingly. Use these guidelines, coupled with the callback storage information listed above, to gauge the precise requirements needed to ensure that GES runs smoothly in your deployment.

GES

(Based on t3.medium)

- vCPUs: 1
- Memory: 2 GiB
- Network burst: 5 Gbps

Redis

(Based on cache.r5.large) Redis is essential to GES service availability. Deploy two Redis caches in a cluster; the second cache acts as a replica of the first. For more information, see Architecture.

Callback data is stored in Redis memory.

- vCPUs: 1
- Memory: 8 GiB
- Network burst: 10 Gbps

PostgreSQL

(Based on db.t3.medium)

- vCPUs: 2
- Memory: 4 GiB
- Network burst: 5 Gbps

-
- Storage: 100 GiB

Sizing calculator

The information in this section is provided to help you determine what hardware you need to run GES and third-party components. The information and formulas are based on an analysis of database disk storage and Redis memory usage requirements for callback data. The numbers provided here include only storage and memory usage for callbacks. Additional storage and memory is required for configuration data and basic operations.

Requirements per callback

Each callback record (excluding user data) requires approximately 6.5 to 7.0 kB of database disk storage, plus additional disk storage for the user data. Each kB of user data consumes approximately 3.0 kB of disk storage.

Each callback record (excluding user data) requires approximately 4.5 to 5.5 kB of Redis memory, plus an additional 1.25 kB for each kB of user data.

Use the following formulas to estimate disk storage and Redis memory requirements:

- Estimate database disk storage requirements for callback data:
$$\text{number of callbacks per day} \times (7 \text{ kB} + (3 \text{ kB} \times \text{kB of user data per callback})) \times 14 \text{ days}$$
- Estimate Redis memory requirements for callback data:
$$\text{number of callbacks per day} \times (5.5 \text{ kB} + (1.25 \text{ kB} \times \text{kB of user data per callback})) \times 14 \text{ days}$$

For example, if a tenant has an average of 100,000 callbacks per day with 1kB user data in each callback:

- The database storage requirement is approximately 14 GB.
- The Redis memory requirement is approximately 9.5 GB.

NOTE: Each callback record is stored for 14 days. If you average about 10k scheduled callbacks every day, and the scheduled callbacks are all booked as far out as possible (that is, 14 days in the future), the number of callbacks to use in storage and memory calculations is $28 \text{ days} \times 10\text{k} \text{ callbacks per day} = 280\text{k} \text{ callbacks}$.

Redis operations

The Redis operations primarily update the connectivity status to other services such as Tenant Service (specifically ORS and URS) and Genesys Web Services and Applications (GWS).

When GES is idle (zero callbacks in the past, no active callback sessions, no scheduled callbacks), GES generates about 50 Redis operations per second per GES node per tenant.

Each Immediate callback generates approximately 110 Redis operations from its creation to the end of the ORS session.

For Scheduled callbacks, assuming each callback generates 110 Redis operations when the ORS session is active (based on Immediate callback numbers), there is 1 additional Redis operation for each minute that a callback is scheduled.

For example, if a callback is scheduled for 1 hour from the time it was created, the number of Redis operations is approximately $60 + 110 = 170$.

For a callback scheduled for 1 day from the time it was created, it generates approximately $60 \times 24 + 110 = 1550$ Redis operations, using the following formula for the number of Redis operations per callback:
number of callbacks \times (110 + number of minutes until scheduled time)

Because the longevity of a callback ORS session depends on the estimated wait time (EWT), the total number of Redis operations performed by GES per minute varies, based on both the number of callbacks in the system and the EWT of the callbacks.

Use the following formula to estimate the number of Redis operations performed per minute:
Total number of Redis operations per minute = (50 base GES Redis operations per second \times 60 seconds) + number of upcoming scheduled callbacks in the system \times ((total number of active callbacks / EWT) \times 110)

Where:

- Total number of active callbacks = number of active immediate callbacks + number of active scheduled callbacks, and
- Number of active scheduled callbacks = (number of scheduled callbacks per time slot / time slot duration) \times EWT

For example, let's say we have the following scenario:

- Scheduled callbacks:
 - Time slot duration = 15 minutes
 - Maximum capacity per time slot = 100
 - Business hours = 24x7
 - Assume that all time slots are fully booked for the next 14 days
- Number of active immediate callbacks = 1,000
- Estimated wait time = 90 minutes

Using the preceding formulas, estimate the Redis operations per minute:

- Total number of scheduled callbacks = $(100 \times (60 / 15)) \times 24 \times 14 = 134,400$
- Number of active scheduled callbacks = $(100 / 15) \times 90 = 600$
- Number of upcoming scheduled callbacks = total number of scheduled callbacks - number of active scheduled callbacks = $(134,400 - 600) = 133,800$
- Total number of active callbacks = $1,000 + 600 = 1,600$
- Total number of Redis operations per minute = $(50 \times 60) + 133,800 + ((1,600 / 90) \times 110) = 138,756$

Redis keys

Each callback creates three additional Redis keys. Given the preceding calculations for Redis memory requirements for each callback, the formula for the average key size is:

$(5.5 \text{ kB} + (1.25 \text{ kB} \times \text{kB_of_user_data_per_callback})) / 3$

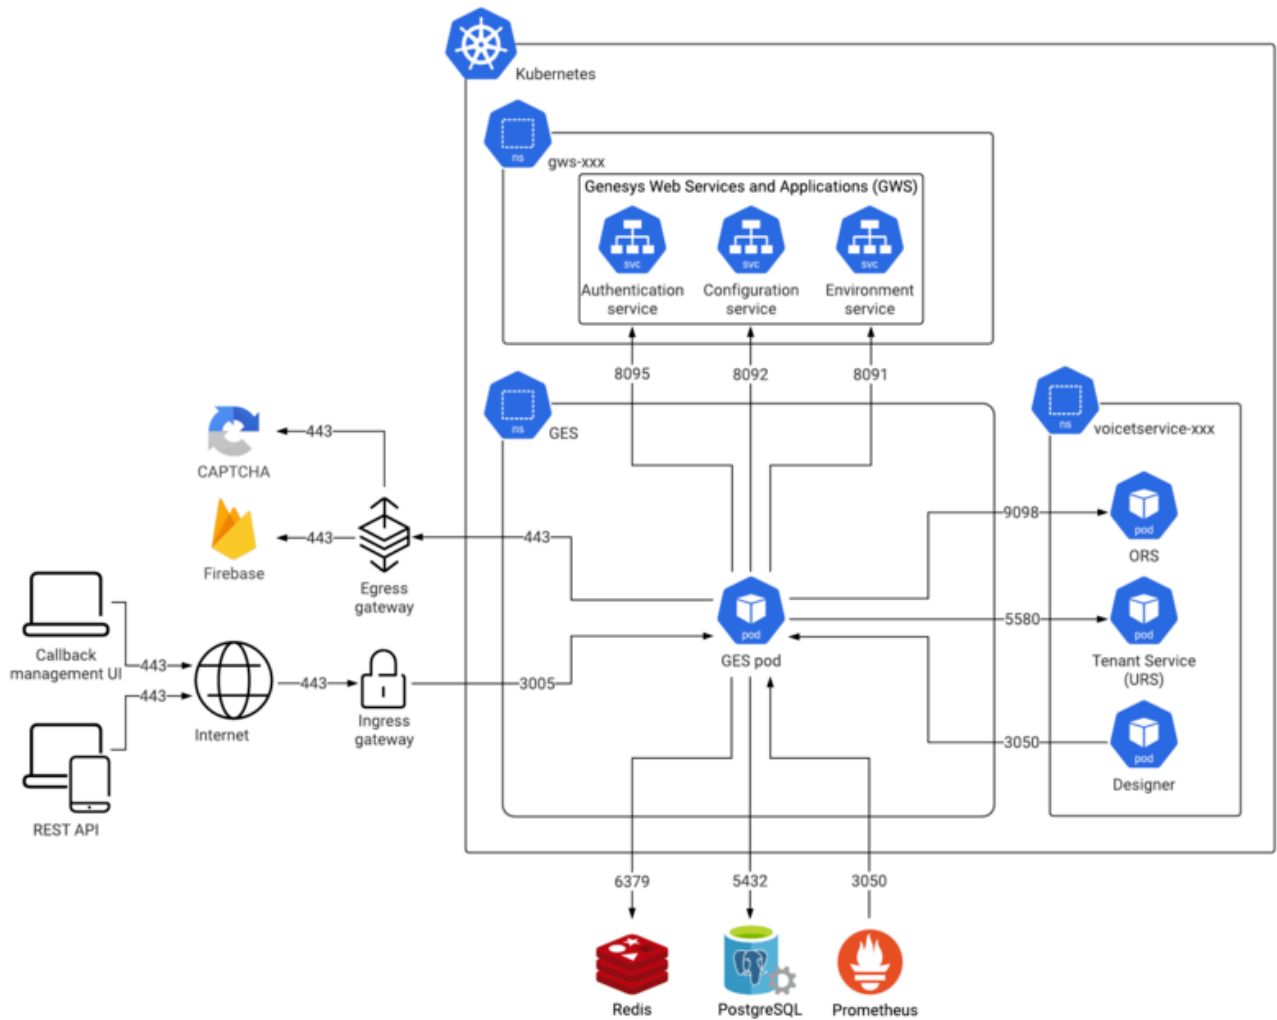
Network requirements

Incoming connections to the GES deployment are handled either through the UI or through the external API. For information about how to use the external API, see the Genesys Multicloud CX Developer Center.

Connection topology

The diagram below shows the incoming and outgoing connections amongst GES and other Genesys and third-party software such as Redis, PostgreSQL, and Prometheus. In the diagram, Prometheus is shown as being part of the broader Kubernetes deployment, although this is not a requirement. What's important is that Prometheus is able to reach the internal load balancer for GES.

The other important thing to note is that, depending on the use case, GES might communicate with Firebase and CAPTCHA over the open internet. This is not part of the default callback offering, but if you use Push Notifications with your callback service, then GES must be able to connect to Firebase over TLS. The use of Push Notifications or CAPTCHA is optional and not necessary for the basic callback scenarios.



Web application firewall rules

Information in the following sections is based on NGINX configuration used by GES in an Azure cloud environment.

Cookies and session requirements

When interacting with the UI, GES and GWS ensure that the user's browser has the appropriate session cookies. By default, UI sessions time out after 20 minutes of inactivity.

The external Engagement API does not require session management or the use of cookies, but it is important that the GES API key be provided in the request headers in the **X-API-Key** field.

For ingress to GES, allow requests to only the following paths to be forwarded to GES:

- /ges/
- /engagement/v3/callbacks/create

-
- /engagement/v3/callbacks/cancel
 - /engagement/v3/callbacks/retrieve
 - /engagement/v3/callbacks/availability/
 - /engagement/v3/callbacks/queue-status/
 - /engagement/v3/callbacks/open-for/
 - /engagement/v3/estimated-wait-time
 - /engagement/v3/call-in/requests/create
 - /engagement/v3/statistics/operations/get-statistic-ex

In addition to allowing connections to only these paths, ensure that the ccid or ContactCenterID headers on any incoming requests are empty. This enhances security of the GES deployment; it prevents the use of external APIs by an actor who has only the CCID of the contact center.

TLS/SSL certificate configuration

There are no special TLS certificate requirements for the GES/Genesys Callback web-based UI.

Subnet requirements

There are no special requirements for sizing or creating an IP subnet for GES above and beyond the demands of the broader Kubernetes cluster.

Browser requirements

The Genesys Callback user interface is supported in the following browsers.

Browsers

Name	Version	Notes
Firefox	Current release or one version previous	Genesys also supports the current ESR release. Genesys supports the transitional ESR release only during the time period in which the new ESR release is tested and certified. For more information, see Firefox ESR release cycle. Firefox updates itself automatically. Versions of Firefox are only an issue if your IT department restricts automatic updates.
Chrome	Current release or one version previous	Chrome updates itself automatically. Versions of Chrome are only an issue if your IT department restricts automatic updates.
Microsoft Edge (Legacy)	Current release	
Microsoft Edge Chromium	Current release	

Genesys dependencies

GES has dependencies on several other Genesys services. You must deploy the services on which GES depends and verify that each is working as expected *before* you provision and configure GES. If you follow this advice, then – if any issues arise during the provisioning of GES – you can be reasonably assured that the fault lies in how GES is provisioned, rather than in a downstream program.

GES/Callback requires your environment to contain supported releases of the following Genesys services, which must be deployed before you deploy Callback:

- Genesys Web Services and Applications (GWS)
- Genesys Authentication
- Voice Microservices (includes Tenant Service)
- Designer
- Agent Setup

For detailed information about the correct order of services deployment, see [Order of services deployment](#).

GDPR support

Callback records are stored for 14 days. The 14-day TTL setting starts at the Desired Callback Time. The `Callback TTL (seconds)` setting in the `CALLBACK_SETTINGS` data table has no effect on callback record storage duration; 14 days is a fixed value for all callback records.