



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.

Voice Microservices Private Edition Guide

Before you begin

5/25/2026

Contents

- 1 Limitations and assumptions
- 2 Download the Helm charts
- 3 Third-party prerequisites
- 4 Storage requirements
 - 4.1 Choosing Voicemail storage
- 5 Network requirements
- 6 Browser requirements
- 7 Genesys dependencies
- 8 GDPR support
 - 8.1 Multi-Tenant Inbound Voice: Voicemail Service
 - 8.2 GDPR multi-region support
 - 8.3 Standalone Scripts
 - 8.4 Limitations

Find out what to do before deploying the Voicemail Service.

Related documentation:

-
-
-

RSS:

- [For private edition](#)

Limitations and assumptions

Voice Voicemail Service is integrated with Workspace Web Edition and Web Services and Applications. The integration brings an appearance of actionable voice mailbox information in the Workspace Web Edition UI, presenting users with Message Waiting Indicator(s) for each voice mailbox assigned to them either directly as a personal mailbox or as a group mailbox via membership in a group(s) having a mailbox provisioned. Users still have access to voicemail from Workspace Web Edition by dialing directly to a voicemail access number, which is 5555.

Download the Helm charts

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

The following table identifies the Helm chart version associated with the Voicemail service.

Service name	Helm chart version
voice-voicemail	voice-voicemail-100.0.xx.tgz

Third-party prerequisites

See the [Third-party prerequisites for the Voice Services](#).

Storage requirements

Choosing Voicemail storage

To store mailbox metadata and messages, consider the following supported options for storage in the Private Edition deployment:

1. Persistent Volumes & Persistent Volume Claims
2. Azure Blob Storage
3. AWS S3 Bucket

See the following sections to learn how to use these storage options and to find information about their limitations.

Persistent Volume & Claim

- Persistent Volume (PV) is a piece of storage that can be mounted to a Voicemail Service deployment inside the Kubernetes cluster.
- Voicemail Service requires a separate storage class and PV to be created for a Voicemail storage.
- If the customer wants to extend the deployment to more than one Kubernetes cluster, Voicemail Service requires to mount the same PV for all the Kubernetes cluster for that customer.
- Create the Persistent Volume Claim (PVC) from the Voicemail PV.
- The access mode for the PVC must be **ReadWriteMany**, since the Voicemail Service will edit the existing data while updating the mailbox settings or the message state.
- Use the sizing doc, which you can find on the Genesys SIP Feature Server landing page, to calculate the required storage space.

Here is the sample Kubernetes YAML file for creating PVCs for a Voicemail Service. The PVC creation is controlled by the Voicemail Service Helm chart by overriding the **values.yaml**.

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: voice-voicemail-pvc
  namespace: voice
  labels:
    servicename: voice-voicemail
spec:
  storageClassName: voice-voicemail
  volumeMode: Filesystem
  accessModes:
    - ReadWriteMany
  resources:
    requests:
      storage: 20Gi
```

Limitations

1. Replication strategies are not available for the data.
2. Retention limit: Admins can't configure the auto-expiration for a Voicemail message.

-
3. When a customer has more than one Kubernetes cluster deployed, the PV for all the Kubernetes clusters must be created from a single storage drive, so that the data from one Kubernetes cluster is shared among other Kubernetes clusters.

Azure blob

- Unlike PV, the Azure Blob Storage provides options to replicate and configure Time to live for the files and can be accessed from any Kubernetes cluster by using the storage access keys.
- Create the Azure Storage with the blob storage.
- The access keys for the blob storage must be securely mounted to the Voicemail pod. You can do one of the following:
 - Store access keys in Azure Key Vault and mount it via a Container Storage Interface (CSI) driver.
 - Create access keys as a Kubernetes secret and volume mount the Kubernetes secret. (This option is considered less secured than the CSI driver approach.)
- The **values.yaml** file can be overridden for configuring either a Kubernetes secret or CSI driver, which is explained in Override Helm chart values.

AWS S3 Bucket

- Like Azure Blob Storage, S3 bucket provides options to replicate and configure Time to live for the files and can be accessed from any Kubernetes cluster by using the access and secret keys.
- Create a new S3 bucket or a folder inside the existing bucket.
- The access and secret keys for blob storage needs to be securely mounted to the Voicemail pod as below.
 - Create access/secret keys as Kubernetes secret and volume mount the Kubernetes secret.
- The values.yaml file can be overridden for configuring Kubernetes secret, which is explained in Override Helm chart values.

Network requirements

For more information, see Network requirements in the *Configure and deploy* section of this guide.

Browser requirements

Not applicable.

Genesys dependencies

For information about dependencies for Voicemail Service, see additional prerequisites on the Deploy Voicemail page. For detailed information about the correct order of services deployment, see Order of

services deployment.

GDPR support

Customer data that is likely to identify an individual, or a combination of other held data to identify an individual is considered as Personally Identifiable Information (PII). Customer name, phone number, email address, bank details, and IP address are some examples of PII.

Multi-Tenant Inbound Voice: Voicemail Service

According to EU GDPR:

- When a customer requests to access personal data that is available with the contact center, the PII associated with the client is exported from the database in client-understandable format. You use the **Export Me** request to do this.
- When a customer requests to delete personal data, the PII associated with that client is deleted from the database within 30 days. However, the Voicemail service is designed in a way that the Customer PII data is deleted in one day using the **Forget Me** request.

Both **Export Me** and **Forget Me** requests depend only on Caller ID/ANI input from the customer. The following PII data is deleted or exported during the **Forget Me** or **Export Me** request process, respectively:

- Voicemail Message
- Caller ID/ANI

GDPR feature is supported only when **StorageInterface** is configured as **BlobStorage**, and **Voicemail service** is configured with Azure storage account data store.

Adding caller_id tag during voicemail deposit

Index tag **caller_id** is included in voicemail messages and metadata blob files during voicemail deposit. Using the index tags, you can easily filter the **Forget Me** messages instead of searching every mailbox.

GDPR multi-region support

In voicemail service, all voicemail metadata files are stored in master region and voicemail messages are deposited/stored in the respective region. Therefore, it is required to connect all the regions of a tenant to perform Forget Me, Undo Forget Me, or Export Me processes for GDPR inputs.

To provide multi-region support for GDPR, follow these steps while performing GDPR operation:

1. Get the list of regions of a tenant.
2. Ensure all regions storage accounts are up. If any one of storage accounts is down, you cannot perform the GDPR operation.
3. GDPR operates in the master region files, first.

4. Then, GDPR operates in all the non-master region files.

Standalone Scripts

You can invoke the **Forget Me** and **Export Me** APIs from a standalone Node.js script. This script can be executed by a user or an automated scheduler. When a user executes the script:

- The script authenticates with the user auth token.
- The user must have the bearer or the basic token.

In the case of an automated scheduler, the script uses the client credential (also known as system account) and processes the request. In this scenario, the user has to configure the GWS URL as an environment variable. The script would generate the auth token for the client and access the GDPR APIs. The script can be integrated into the GitHub Actions pipeline and invoked from the GitHub pipeline.

Script Inputs

Parameter	Value	Is it mandatory	Description
-i or --input	file-path	Yes	<p>Input.json is the same as the JSON input passed to the REST API.</p> <p>The client credentials ccid (contact center id) must be included as key-value pair in the input.json file because ccid cannot be fetched from auth token of the client credentials.</p> <p>Sample value "ccid" : "2c5ea4c0-4067-11e9-8bad-9b1deb4d3b7d"</p>
-o or --output	output-location	Yes	<p>output.json</p> <p>Forget Me Operation:</p> <p>The output.json is the same as the response from the Forget Me API.</p> <p>Export Me Operation:</p> <pre>{ "caseid": "123456789", "consumers": { //The message media is exported in the output location and the filename is the same as the message IDs. "555551212": ["filename of</pre>

			<pre>message1", "filename of message2"], "555556161": [], "555556162": ["filename of message1"] } }</pre> <p>Undo Operation:</p> <p>The output.json is the same as the response from Undo API.</p> <p>execution.log</p> <p>Execution logs are available in the execution.log file</p>
-u or --user	User token	Either user token or client credentials	User token is fetched from GWS
-c or --client	Client credentials	Either user token or client credentials	Client credential is required when scheduling the script. Client credentials can be obtained by requesting the GWS team.
-p or --operation	forgetme exportme undoforgot	Yes	Type of operation to be done when the script is executed.

User token example:

```
node gdpr.js -i "t2026" -o "t2026" -u "dDIwMjZcXGRlZmF1bHQ6cGFzc3dvvW=" -p "forgetme" --basic
node gdpr.js -i "t2026" -o "t2026" -u "dDIwMjZcXGRlZmF1bHQ6cGFzc3dvvW=" -p "undoforget" --basic
node gdpr.js -i "t2026" -o "t2026" -u "dDIwMjZcXGRlZmF1bHQ6cGFzc3dvvW=" -p "exportme" --basic
```

Client credentials example:

```
node gdpr.js -i "t2026" -o "t2026" -c "iPdZIMR5qHAohE4wsC0La0eAopUyJDZalmwN6FPH9rjUcztZ" -p "forgetme"
node gdpr.js -i "t2026" -o "t2026" -c "iPdZIMR5qHAohE4wsC0La0eAopUyJDZalmwN6FPH9rjUcztZ" -p "undoforget"
node gdpr.js -i "t2026" -o "t2026" -c "iPdZIMR5qHAohE4wsC0La0eAopUyJDZalmwN6FPH9rjUcztZ" -p "exportme"
```

Limitations

MWI count is not updated automatically on deleting the files during **Forgetme** operation. It is updated during the next voicemail message deposit or voicemail message delete of a mailbox.

-
- If the **Forgetme** rule is first executed at 10:00 UTC in the day, then the file **X** marked for **Forgetme** at 10.01 UTC same day, the **Forgetme** rule does not delete the file 'X' on the second day at 10:00 UTC since it does not meet the **file has not been modified in one day** condition. However, it gets deleted in next day.
 - If the message is deposited and not read by any agent, the **Forget Me** API is executed and marked for deletion. Before deleting the file, if the agent reads the message/forward, the message metadata and the last modified time are updated. In such cases, the file may not be deleted in one day because the last modified date condition is not met.