



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 Knowledge Center Deployment Guide

Software Configuration

5/10/2026

Contents

- 1 Software Configuration
 - 1.1 JVM Settings
 - 1.2 Co-Locating On Same Host
 - 1.3 Genesys Knowledge Center Server/Cluster Configuration

Software Configuration

Important

Disclaimer: This page contains information related to the 8.5.x release of the product. It will be updated soon to reflect the 9.0 release.

JVM Settings

We recommended you run the most recent version of Java Virtual Machine. The minimum version supported is Java 1.8 64 bit. The recommended (default) memory settings for the Knowledge Center Server java process (Xms/Xmx) is 4 Gb. Extending it to 16 Gb might be a good idea when using large amounts of data. Adding an extra 4 Gb of heap for the process recommended on the nodes that are used to execute history archiving.

Important

It is strongly advised to not set the Xmx larger than 30Gb.

Co-Locating On Same Host

We do not recommend co-locating two or more nodes of the Genesys Knowledge Center Server on the same host. Losing the host due to hardware or network failure will result in two or more nodes being lost by the cluster. This will lead to massive network operations, potential loss of data (in the case where the number of replications is configured to 1) and even a complete outage of the cluster (for example, for clusters with 3 nodes).

Genesys Knowledge Center Server can be co-located on the same host with other Genesys or 3rd party solutions.

When co-locating please ensure that:

- The host is equipped with enough resources, especially RAM
- the solutions you are co-locating with are not causing any CPU spike when any other application on the host is blocked

Do not co-locate with other disks' read-write intensive solutions as it could easily overrun capabilities of your disk system.

Important

In most cases it is recommended to plan your deployment with 50% of your memory allocated to the running application. The remaining 50% is used for the OS filesystem cache allowing for different software (including Genesys Knowledge Center) to work faster and minimizing disk operations.

Genesys Knowledge Center Server/Cluster Configuration

Genesys Knowledge Center Server/Cluster is configured by default to enable smooth and high performance operations. Detailed information on every option exposed, valid values, and considerations behind the solutions are available on the [Configurations Options](#) page of the product documentation.

The section below provides additional information and recommendations.

Cluster Setup

The starting point of any cluster configuration is to create one application of the Application Cluster type and as many as needed applications of the Genesys Knowledge Center Server type in Genesys Administrator.

While doing this please ensure that you are creating the application only for the actual nodes of the Server that you will use. Having a spare application added to the cluster will mislead the cluster members while calculating the required number of nodes to be online and to enable cluster functionality.

Generally a cluster expects the $N/2+1$ node to be online to start servicing the clients (where N is the number of Genesys Knowledge Center Server applications connected to the cluster and is enabled).

For example, if you have configured 3 nodes in the Genesys Administrator and started just one, the started server will refuse any client's requests as it will treat the overall cluster as not started yet. The cluster will become functional only when the second node of The Genesys Knowledge Center Server joins the cluster.

In rare cases when you still need to run such a configuration there are several ways to enforce cluster functionality:

- Disable applications that you are not going to run (by unchecking the **Enable** checkbox in every such application in Genesys Administrator)
- Manually setting the **minimumMasterNodes** value option in the **index** section into a desired value (for example, 1 in the example below).

Important

We do not recommended using manual settings, especially manually defining the **minimumMasterNodes** option as any incorrect values could result in data corruption.

Tip

Running nodes needs to be restarted if you applying one of the listed below recommendations.

Internode Communication

When running a production cluster it is required to disable the multicast node discovery by setting it to *false* **enabled** option in **multicast** section of the cluster application object in the configuration. It is also a good idea to disable multicast in lab deployments.

Multicast may lead to incidental joining of the undesired node to the cluster that will trigger data relocations to rebalance the data between nodes.

Host On-Boarding

Having the majority of the configuration in common between all nodes in the cluster there are few parameters that must be configured to onboard the node on a particular host. An example of these parameters is the folder for log files that is configured in the **log** section of the Genesys Knowledge Center Server application.

The other configuration that you need to set properly is the location of the indexed knowledge that is configured in option **path.data** of **gms.yml** file.

Important

It is recommended to store your indexed knowledge on the fastest disk you have on the host. By default the data will be placed into the folder where you have installed Genesys Knowledge Center Server.

If you have two disks attached to the host (for example, one spinning and another SSD) you can reconfigure the application to store the data in the fastest disc (SSD) while using the spinning disc for the application binaries.