



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

Elasticsearch

12/15/2025

Elasticsearch

Contents

- [1 Elasticsearch](#)
 - [1.1 Configuring Web Services to use a standalone Elasticsearch cluster](#)
 - [1.2 Configuring Web Services to use an embedded Elasticsearch cluster](#)
 - [1.3 Next Step](#)

Web Services uses **Elasticsearch** — an open-source, full-text search engine with a RESTful web interface — to store both real-time and historical statistics. Real-time statistics reflect the current state of the object (User, Queue, Skill), while historical statistics are stored as time-based events.

For new deployments, Genesys recommends that you set up a cluster of Elasticsearch nodes that is separate from your Web Services nodes. See [Configuring Web Services to use a standalone Elasticsearch cluster](#) for details. It's possible to set up an embedded Elasticsearch cluster, which means that Elasticsearch is included in your Web Services nodes. **Note: Genesys does not recommend using this approach in a production environment.** See [Configuring Web Services to use an embedded Elasticsearch cluster](#) for details.

Configuring Web Services to use a standalone Elasticsearch cluster

You can configure Web Services to work with a standalone Elasticsearch cluster by completing the steps below.

Important

Genesys recommends that you use a standalone Elasticsearch cluster for new deployments of Web Services. Contact your Genesys representative for information about how to migrate from embedded to standalone.

Prerequisites

- You have deployed and configured a cluster of Elasticsearch nodes. Refer to the [Elasticsearch documentation](#) for details. **Note:** Genesys recommends that you use the **latest stable 1.x version of Elasticsearch**.

Start

Complete the following steps for each Web Services node:

1. Copy the **installation_path/config-templates/elasticsearch.yml.sample** file to the configuration folder on each node in your Elasticsearch cluster, rename the file to **elasticsearch.yml**, and then edit the file to set correct option values.
2. Copy all files from the **installation_path/elasticsearch/templates/** folder to the templates directory of the configuration folder, for example, **/etc/elasticsearch/templates**, on each node in your Elasticsearch cluster.
3. Web Services keeps Elasticsearch in sync with Cassandra and removes statistical data for deleted objects by performing index verification. By default, Web Services runs index verification when its node is started, but you can also configure scheduled index verification by setting **enableScheduledIndexVerification** to true in the **elasticSearchSettings** option. By default, the index verification takes place every 720 minutes (12 hours), but you can change this timing by setting **indexVerificationInterval** in **elasticSearchSettings**. **Note:** Genesys recommends that you only configure one Web Services node for index verification to avoid excessive requests to Elasticsearch and Cassandra.

4. If your deployment uses statistics, make sure you complete the [reporting configuration steps](#), including setting the `nodeId`.
5. Set the `crClusterName` option to the name of the cluster. All Web Services nodes with the same cluster name will form the cluster.
6. Set the `enableElasticSearchIndexing` option to true. **Note:** In this case, Web Services writes statistics values to both Elasticsearch and Cassandra, but only reads them from Elasticsearch.
7. Set the `elasticSearchSettings` option to appropriate values for your environment.

End

Configuring Web Services to use an embedded Elasticsearch cluster

You can configure Web Services to work with an embedded Elasticsearch cluster by completing the steps below.

Important

Genesys does not recommend using an embedded Elasticsearch cluster in a production environment. Contact your Genesys representative for information about how to migrate from embedded to standalone.

Complete the following steps for each Web Services node that you want to host Elasticsearch:

Start

1. Copy `installation_path/config-templates/elasticsearch.yml.sample` to the `installation_path/config/elasticsearch.yml` file, and then edit the file to set correct option values.
2. Web Services keeps Elasticsearch in sync with Cassandra and removes statistical data for deleted objects by performing index verification. By default, Web Services runs index verification when its node is started, but you can also configure scheduled index verification by setting `enableScheduledIndexVerification` to true in the `elasticSearchSettings` option. By default, the index verification takes place every 720 minutes (12 hours), but you can change this timing by setting `indexVerificationInterval` in `elasticSearchSettings`. **Note:** Genesys recommends that you only configure one Web Services node for index verification to avoid excessive requests to Elasticsearch and Cassandra.
3. If your deployment uses statistics, make sure you complete the [reporting configuration steps](#), including setting the `nodeId`.
4. Set the `crClusterName` option to the name of the cluster. All Web Services nodes with the same cluster name will form the cluster.
5. Set the `enableElasticSearchIndexing` option to true. **Note:** In this case, Web Services writes statistics values to both Elasticsearch and Cassandra, but only reads them from Elasticsearch.
6. Set the `elasticSearchSettings` option to appropriate values for your environment.

End

Next Step

- [Back to Configuring features](#)