



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.

Deployment Guide

Configure a Cluster of Co-browse Servers

12/21/2025

Configure a Cluster of Co-browse Servers

Genesys Co-browse supports load balancing using [Stickiness](#).

Load balancing is enabled by configuring a cluster of Co-browse Servers. Cassandra is embedded in Genesys Co-browse Server, so when you set up a cluster of Genesys Co-browse servers, each server may also act as a Cassandra node. You configure the Cassandra nodes by setting configuration options in the `cassandraEmbedded` section of the Co-browse Server application.

Complete the following steps to implement load balancing:

1. Set up 2 or more Co-browse Server nodes

Tip

To determine the how many nodes your Co-browse cluster needs, use the [Genesys Co-browse Sizing Calculator](#).

8.5.000

For Co-browse 8.5.000, you must set up a cluster of Co-browse Servers to enable load balancing. For each Co-browse Server in your planned cluster, complete the procedures on the [Install Genesys Co-browse Server](#) page.

If the Co-browse servers reside on the same machine, you must configure the applications to use different ports for Jetty and the embedded Cassandra. This will prevent port conflicts among your Co-browse Servers.

Important

Every Co-browse Server in the cluster generally plays the same role as the others, except some embedded Cassandra nodes act as seed nodes. This means that to see consistent behavior on the cluster, regardless of which server serves requests, all Co-browse Servers should have the same options set in their application objects in Configuration Server. The rule of thumb is to configure the cluster servers the same, unless it is absolutely necessary to do otherwise (for example, a port is busy on a machine). This simplifies maintenance of production deployments.

8.5.001+

For Co-browse 8.5.001+, you must set up a cluster of Co-browse Nodes to enable load balancing. To do this, complete the procedures to create Application objects for a Co-browse Cluster and Co-browse Nodes. Follow the installation steps outlined in the **8.5.001+** tab in the [Creating the Co-browse Server Application Object in Genesys Administrator](#) section.

If the Co-browse Nodes reside on the same machine, you must configure the applications to use different ports for Jetty and the embedded Cassandra. This will prevent port conflicts among your Co-browse Nodes.

2. Configure the Cassandra cluster

Configure the Cassandra cluster

Prerequisite: You have completed [Set up 2 or more Co-browse Servers](#).

External Cassandra Cluster Setup

Important

For supported versions of Cassandra, see [Genesys Co-browse in the Supported Operating Environment Reference Guide](#).

External Cassandra cluster deployment is thoroughly described in the official Cassandra documentation. See the following:

- [Planning a cluster deployment](#)
- [Initializing a cluster](#)
- [External Cassandra Access Configuration](#)
- [Replication Strategy Configuration](#)
- [Replication Factor Configuration](#)
- [Cassandra Native Transport Configuration](#)

Replication Strategy

By default, Co-browse server activates *NetworkTopologyStrategy* as a [replication strategy](#). *NetworkTopologyStrategy* is [recommended](#) for production Cassandra deployments and is supported by [GossipingPropertyFileSnitch](#). *GossipingPropertyFileSnitch* relies on the `cassandra-rackdc.properties` file. Make sure the data center names defined in this file (one for each Cassandra node) correspond to data center names defined in the [replicationStrategyParams](#)

option.

The `cassandra-rackdc.properties` file location depends on the type of Cassandra cluster:

- External Cassandra—see the `conf` subdirectory of the Cassandra installation directory. For more information, see http://docs.datastax.com/en/cassandra/2.0/cassandra/architecture/architectureSnitchGossipPF_c.html.
- Embedded Cassandra—see the `resources` subdirectory of the server directory. For example, `<Co-browse Server Install Directory>/server/resources`.

Replication Factor

Replication factor configures the number of copies of data to keep in the cluster. Typically, three copies is enough for most scenarios (provided you have at least three nodes in your cluster). The replication factor can be increased to achieve higher redundancy levels. Set the replication factor in the `replicationStrategyParams` option to a number less than or equal to the number of nodes.

Cassandra Native Transport

The Co-browse Server acts as the Cassandra CQL client, so, native transport should be fully configured for all nodes of your external Cassandra cluster. To verify, open the `cassandra.yaml` file for each Cassandra node and check whether:

1. `native_transport_port` is set to any valid port.
2. `start_native_transport` is set to `true`.

Embedded Cassandra Cluster Setup

Important

Starting in 8.5.0, Embedded Cassandra mode is deprecated in Genesys Co-browse; support for this mode will be discontinued in 9.0.

An Embedded Cassandra cluster is setup similarly to an external Cassandra except that embedded Cassandra node settings are provided either through **Configuration Server options** or an **external `cassandra.yaml` file**.

Start of procedure

Complete the steps below for each Co-browse application you created in **Set up 2 or more Co-browse Servers**:

1. Open Genesys Administrator and navigate to **PROVISIONING > Environment > Applications**.
2. Select the Co-browse application and click **Edit**.

3. In the **Options** tab, locate the **cassandraEmbedded** section and update the following options:
 1. **listenAddress** — Set this value to the IP of the node (**listenAddress** in **cassandra.yml**).
 2. **rpcAddress** — Set this value to the IP of the node (**rpcAddress** in **cassandra.ymal**).
 3. **seedNodes** — Set this value to the IP of the first node (**seedNodes** in **cassandra.yaml**).
 4. **clusterName** (optional) — This name should be the same for each node (**name** in **cassandra.yaml**).
4. Click **Save & Close**.

End of procedure

3. Verify the status of the Cassandra cluster

Prerequisite: You have a separate installation of Cassandra 2.x (2.1.3+, the same version used in Co-browse).

1. Start the first node and wait until it starts listening.
2. Start all other nodes in your cluster.
3. Open a command line and run `<cassandra home>\bin\cassandra-cli.bat -h <ip of first node> -p <cassandra rpcPort>`, where `<ip of first node>` is the IP of the first node in your cluster and `<cassandra rpcPort>` is the value you configured for **rpcPort**.
4. Enter the following command: `describe cluster`. The output should look similar to the following:

```
[default@unknown] describe cluster;  
Cluster Information:  
Snitch: org.apache.cassandra.locator.SimpleSnitch  
Partitioner: org.apache.cassandra.dht.RandomPartitioner  
Schema versions:  
6c960880-1719-11e3-0000-242d50cf1fbf: [192.168.00.1, 192.168.00.2, 192.168.00.3]
```

The list of IP address in square brackets (`[192.168.00.1, 192.168.00.2 ...]`) should match all the nodes in your cluster.

End of procedure

4. Configure the load balancer

See [Configuring a Load Balancer for Co-browse Cluster](#) for details about configuring the load balancer and sample configurations for **Nginx** and **Apache**.

Tip

In Co-browse 8.5.002+, the agent side uses the **cluster URL** while the

end user (customer) side uses the URL in the **Website Instrumentation**. You can have two load balancers, an internal load balancer for agents which you specify in the **cluster URL** option and a public load balancer for end users to use in the JS instrumentation. Depending on your infrastructure's setup, two load balancers may benefit traffic.

5. Modify the website instrumentation

You must modify the URLs in your Co-browse instrumentation scripts to point to your configured load balancer. See **Website Instrumentation** for details about modifying the script.

Important

Starting with Co-browse 8.5.002+, the customer side always uses the URL in the JS instrumentation for css-proxy and url-proxy.

If you are using the Co-Browse proxy to instrument your site, you will need to modify the URLs in the in proxy's `map.xml` file. See **Test with the Co-browse Proxy** for details about modifying the xml file.

Warning

The Co-browse proxy should only be used in a lab environment, not in production.

6. Modify the agent side and controller configuration

Configure the Co-browse Server applications

- 8.5.000—Modify the `url` option in the **cluster section** of all your Co-browse Server applications.
- 8.5.001—Modify the `url` option in the **cluster section** of your Co-browse Cluster application.

See the **cluster section** for details.

Tip

In Co-browse 8.5.002+, the agent side uses the **cluster URL** while the end user (customer) side uses the URL in the **Website Instrumentation**. You can have two load balancers, an internal load balancer for agents which you specify in the **cluster URL** option and a public load balancer for end users to use in the JS instrumentation. Depending on your infrastructure's setup, two load balancers may benefit traffic.

You must also set up a similar configuration for the Genesys Co-browse Plug-in for Workspace Desktop Edition. To support this, you might consider setting up two load balancers:

- public — This load balancer should have a limited set of Co-browse resources. For example, it should not include session history resources.
- private — This load balancer should have all Co-browse resources and it should be placed in the network so that it is accessible only from the corporate intranet. It should only be used for internal applications, such as Workspace Desktop Edition.

Complete the procedure below to configure the plug-in to support the Co-browse cluster:

Configure the Co-browse Plug-in for Workspace Desktop Edition

See [Configuring Workspace Desktop Edition to allow the Plug-in to work with co-browsing](#).

7. Configure Genesys Workspace Web Edition

If you use Workspace Web Edition on the agent side, you must configure it to work with Co-browse. For instructions, see [Configure Genesys Workspace Web Edition to Work with Co-browse](#).

8. Launch and test

To test your set-up, create a Co-browse session, join it as an agent and do some co-browsing. If you can do this, your configuration was successful.

End of procedure