



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

cassandraEmbedded Section

Contents

- 1 cassandraEmbedded Section
 - 1.1 enabled
 - 1.2 clusterName
 - 1.3 seedNodes
 - 1.4 commitLogDirectory
 - 1.5 dataDirectory
 - 1.6 savedCachesDirectory
 - 1.7 listenAddress
 - 1.8 rpcAddress
 - 1.9 rpcPort
 - 1.10 nativeTransportPort
 - 1.11 storagePort
 - 1.12 sslStoragePort
 - 1.13 configFile
 - 1.14 endpointSnitch
 - 1.15 Additional options not included in the template

cassandraEmbedded Section

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.

The `cassandraEmbedded` section configures embedded Cassandra support for the Co-browse Server cluster.

`enabled`

Default Value: `true`

Valid Values: `true` or `false`

Changes Take Effect: After Co-browse server restart

Specifies whether or not Co-browse server should act as a Cassandra cluster node.

`clusterName`

Default Value: `Cluster`

Valid Values: Any string

Changes Take Effect: After Co-browse server restart

The name of the embedded Cassandra cluster node. This option is mainly used to prevent machines in one logical cluster from joining another. For more information, see http://docs.datastax.com/en/cassandra/2.1/cassandra/configuration/configCassandra_yaml_r.html?scroll=reference_ds_qfg_n1r_1k__cluster_name

`seedNodes`

Default Value: `localhost`

Valid Values: Comma-delimited list of IP addresses

Changes Take Effect: After Co-browse server restart

When a node joins a cluster, it contacts the seed node(s) listed in this option to determine the ring topology and get gossip information about the other nodes in the cluster.

Every node in the cluster should have the same list of seeds specified as a comma-delimited list of IP addresses. In multiple data center clusters, the seed list should include at least one node from each data center (replication group). For more information, see http://docs.datastax.com/en/cassandra/2.1/cassandra/configuration/configCassandra_yaml_r.html?scroll=reference_ds_qfg_n1r_1k__seed_provider.

This option is only applicable when embedded Cassandra service is activated.

commitLogDirectory

Default Value: `./storage/commitLog`

Valid Values: Valid directory path. The directory may not exist.

Changes Take Effect: After Co-browse server restart

Specifies the directory where Cassandra's commitlog directories will be located or created. If left empty, the Co-browse Server web application assumes it is running within a Jetty web container and the storage directory will be a storage sub-directory of the Jetty home directory.

This option is only applicable when embedded Cassandra service is activated.

dataDirectory

Default Value: `./storage/data`

Valid Values: Valid directory path. The directory may not exist.

Changes Take Effect: After Co-browse server restart

Specifies the directory where Cassandra's data will be located or created. If left empty, the Co-browse Server web application assumes it is running within a Jetty web container and the storage directory will be a storage sub-directory of the Jetty home directory.

This option is only applicable when embedded Cassandra service is activated.

savedCachesDirectory

Default Value: `./storage/saved_cache`

Valid Values: Valid directory path. The directory may not exist.

Changes Take Effect: After Co-browse server restart

Specifies the directory where Cassandra's saved_caches directories will be located or created. If left empty, the Co-browse Server web application assumes it is running within a Jetty web container and the storage directory will a "storage" sub-directory of Jetty home directory.

The option is applicable only when embedded Cassandra service is activated.

listenAddress

Default Value: `localhost`

Valid Values: Blank or valid address

Changes Take Effect: After Co-browse server restart

Specifies the address to bind to and to tell other Cassandra nodes to connect to. You *must* change this if you want multiple nodes to be able to communicate.

Leaving this option blank lets `InetAddress.getLocalHost()` set the address. If the node is properly configured (hostname, name resolution), the address will resolve to the address associated with the hostname.

rpcAddress

Default Value: `localhost`

Valid Values: Valid IP address or hostname.

Changes Take Effect: After Co-browse server restart

Specifies the listen address for remote procedure calls (client connections). This option is also used to configure Co-browse server as a client. See http://docs.datastax.com/en/cassandra/2.1/cassandra/configuration/configCassandra_yaml_r.html?scroll=reference_ds_qfg_n1r_1k_rpc_address. If the address is invalid, Co-browse server will not be able to connect to the embedded Cassandra service.

rpcPort

Default Value: 9160

Valid Values: Any free TCP port

Changes Take Effect: After Co-browse server restart

Specifies the port for remote procedure calls (client connections) and the Thrift service.
http://docs.datastax.com/en/cassandra/2.1/cassandra/configuration/configCassandra_yaml_r.html?scroll=reference_ds_qfg_n1r_1k_rpc_address

nativeTransportPort

Default Value: 9042

Valid Values: Any free TCP port

Changes Take Effect: After Co-browse server restart

Specifies the port for the CQL native transport to listen for clients.

storagePort

Default Value: 7000

Valid Values: Any free TCP port

Changes Take Effect: After Co-browse server restart

Specifies the TCP port for commands and data.

sslStoragePort

Default Value: 7001

Valid Values: Any free TCP port

Changes Take Effect: After Co-browse server restart

Specifies the SSL port for encrypted communication.

configFile

Default Value: none

Valid Values: Valid path to the *.yaml cassandra configuration file

Changes Take Effect: After Co-browse server restart

Specifies the Embedded Cassandra external configuration YAML file path. It overrides all Cassandra settings in the section.

endpointSnitch

Default Value: GossipingPropertyFileSnitch
Valid Values: SimpleSnitch, GossipingPropertyFileSnitch, PropertyFileSnitch, Ec2Snitch, Ec2MultiRegionSnitch, or RackInferringSnitch
Changes Take Effect: After Co-browse server restart

A snitch determines which nodes belong to which data centers and racks. They inform Cassandra about the network topology so Cassandra can route requests efficiently. They also allow Cassandra to distribute replicas by grouping machines into data centers and racks. Specifically, the replication strategy places the replicas based on the information provided by the new snitch. Also see, http://docs.datastax.com/en/cassandra/2.1/cassandra/architecture/architectureSnitchesAbout_c.html.

Additional options not included in the template

You can also configure the following options which are not included in the template:

Important

All options in this section are applied only after application restart.

[+] Click to view table

Option name	Mandatory	Default Value	Possible Values	Description
partitioner	No	org.apache.cassandra.dht.Murmur3Partitioner	org.apache.cassandra.dht.ByteOrderedPartitioner, org.apache.cassandra.dht.Murmur3Partitioner, org.apache.cassandra.dht.RandomPartitioner	<p>A partitioner determines how data is distributed across the nodes in the cluster (including replicas). Basically, a partitioner is a function for <code>ByteBuffer</code> to <code>int</code>, representing a row from the table. Partitioners typically use hashing. Each row of data is distributed across the cluster by the value of the token.</p> <p>http://docs.datastax.com/en/cassandra/2.1/cassandra/architecture/architecturePartitionerAbout_c.html</p>
commitFailurePolicy	No	stop	stop,	Policy for commit disk

Option name	Mandatory	Default Value	Possible Values	Description
			stop_commit, ignore, die	<p>failures:</p> <ul style="list-style-type: none"> • <i>die</i> - Shut down gossip and Thrift and kill the JVM, so the node can be replaced. • <i>stop</i> - Shut down gossip and Thrift, leaving the node effectively dead, but can be inspected using JMX. • <i>stop_commit</i> - Shut down the commit log, letting writes collect but continuing to service reads • <i>ignore</i> - Ignore fatal errors and let the batches fail
diskFailurePolicy	No	stop	best_effort, stop, ignore, stop_paranoid, die	<p>Sets how Cassandra responds to disk failure. Recommend settings are <i>stop</i> or <i>best_effort</i>.</p> <ul style="list-style-type: none"> • <i>die</i> - Shut down gossip and Thrift and kill the JVM for any file system errors or single SSTable errors, so the node can be replaced. • <i>stop_paranoid</i> - Shut down gossip and Thrift even for single SSTable errors.

Option name	Mandatory	Default Value	Possible Values	Description
				<ul style="list-style-type: none">• <i>stop</i> - Shut down gossip and Thrift, leaving the node effectively dead, but available for inspection using JMX.• <i>best_effort</i> - Stop using the failed disk and respond to requests based on the remaining available SSTables. This means you will see obsolete data at consistency level of ONE.• <i>ignore</i> - Ignores fatal errors and lets the requests fail; all file system errors are logged but otherwise ignored.
autoBootstrap	No	true	true, false	This setting has been removed from default configuration. It makes new (non-seed) nodes automatically migrate the right data to themselves. When initializing a fresh cluster <u>without</u> data, set this option to <i>false</i>
batchSizeWarnThreshold	No	5	Valid integer	Log WARN on any batch size exceeding this value in kilobytes. Caution should be

Option name	Mandatory	Default Value	Possible Values	Description
				taken on increasing the size of this threshold as it can lead to node instability
concurrentReads	No	32	Valid ineteger	For workloads with more data than can fit in memory, the bottleneck is ads fetching data from disk. Setting to $16 \times \text{number_of_drives}$ allows operations to queue low enough in the stack so that the OS and drives can reorder them. The default setting applies to both logical volume managed (LVM) and RAID drives
concurrentWrites	No	32	Valid ineteger	Writes in Cassandra are rarely I/O bound, so the ideal number of concurrent writes depends on the number of CPU cores in your system. The recommended value is $8 \times \text{number_of_cpu_cores}$
concurrentCounterWrites	No	32	Valid ineteger	Counter writes read the current values before incrementing and writing them back. The recommended value is $16 \times \text{number_of_drives}$
streamThroughputOutbound	No	200	Valid integer	Throttles all outbound streaming file transfers on a node to the specified throughput (Megabits/seconds).

Option name	Mandatory	Default Value	Possible Values	Description
				Cassandra does mostly sequential I/O when streaming data during bootstrap or repair, which can lead to saturating the network connection and degrading client (RPC) performance.
interDCStreamThroughputOutbound	No		Valid integer	Throttles all streaming file transfer between the data centers (Megabits/seconds).. This setting allows throttles streaming throughput between data centers in addition to throttling all network stream traffic as configured with streamThroughputOutbound
trickleFsync	No	false	true, false	When doing sequential writing, enabling this option tells fsync to force the operating system to flush the dirty buffers at a set interval trickleFsyncInterval . Enable this parameter to avoid sudden dirty buffer flushing from impacting read latencies. Recommended to use on SSDs, but not on HDDs.
trickleFsyncInterval	No	10240	Valid integer	Sets the size of the fsync in kilobytes
autoSnapshot (NODE ONLY)	No	true	true, false	Enable or disable whether a snapshot is taken

Option name	Mandatory	Default Value	Possible Values	Description
				of the data before keyspace truncation or dropping of tables. To prevent data loss, using the default setting is strongly advised. If you set to false, you will lose data on truncation or drop
incrementalBackups	No	false	true, false	Backs up data updated since the last snapshot was taken. When enabled, Cassandra creates a hard link to each SSTable flushed or streamed locally in a backups/ subdirectory of the keyspace data. Removing these links is the operator's responsibility
snapshotBeforeCompaction	No	false	true, false	Enable or disable taking a snapshot before each compaction. This option is useful to back up data when there is a data format change. Be careful using this option because Cassandra does not clean up older snapshots automatically
commitLogSync	No	periodic	periodic, batch	The method that Cassandra uses to acknowledge writes http://docs.datastax.com/en/cassandra/2.1/cassandra/dml/dml_durability_c.html
commitLogSyncPeriod	No	10000	Valid integer	The period that Cassandra uses to acknowledge writes in

Option name	Mandatory	Default Value	Possible Values	Description
				milliseconds
commitLogSegmentSize	No	32	Valid integer	Sets the size (in Mb) of the individual commitlog file segments. A commitlog segment may be archived, deleted, or recycled after all its data has been flushed to SSTables. This amount of data can potentially include commitlog segments from every table in the system. The default size is usually suitable for most commitlog archiving, but if you want a finer granularity, 8 or 16 MB is reasonable.
commitLogTotalSpace	No	8192	Valid integer	Total space used for commitlogs. If the used space goes above this value, Cassandra rounds up to the next nearest segment multiple and flushes memtables to disk for the oldest commitlog segments, removing those log segments. This reduces the amount of data to replay on start-up, and prevents infrequently-updated tables from indefinitely keeping commitlog segments. A small total commitlog space tends to cause more flush activity on less-

Option name	Mandatory	Default Value	Possible Values	Description
				active tables
concurrentCompactors	No		Valid integer	<p>Sets the number of concurrent compaction processes allowed to run simultaneously on a node, not including validation compactions for anti-entropy repair. Simultaneous compactions help preserve read performance in a mixed read-write workload by mitigating the tendency of small SSTables to accumulate during a single long-running compaction. If your data directories are backed by SSD, increase this value to the number of cores. If compaction running too slowly or too fast, adjust compactionThroughput first.</p> <p>If not set the value will be calculated: Smaller of number of disks or number of cores, with a minimum of 2 and a maximum of 8 per CPU core</p>
sstablePreemptiveOpenInterval	No	50	Valid integer	<p>When compacting, the replacement opens SSTables before they are completely written and uses in place of the prior SSTables for any range previously written (in Mb). This setting helps to smoothly transfer reads between the SSTables by reducing page cache churn and keeps hot rows hot.</p>
compactionThroughput	No	16	Valid integer	<p>Throttles compaction to the specified total throughput across the entire system (in Mb/seconds). The faster you insert data, the faster you need to compact in order to keep the SSTable count down. The</p>

Option name	Mandatory	Default Value	Possible Values	Description
				recommended value is 16 to 32 times the rate of write throughput (in MB/second). Setting the value to 0 disables compaction throttling.
compactionLargePartitionWarningThreshold	No	100	Valid integer	Logs a warning when compaction partitions larger than the set value in Mb
numTokens	No	256	Valid integer	Defines the number of tokens randomly assigned to this node on the ring when using virtual nodes (vnodes). The more tokens, relative to other nodes, the larger the proportion of data that the node stores.
memtableAllocationType	No	heap_buffers	unslabbed_heap_buffers, heap_buffers, offheap_buffers, offheap_objects	Specify the way Cassandra allocates and manages memtable memory. See Off-heap memtables in Cassandra 2.1 .
memtableCleanupThreshold	No		Valid float	Ratio of occupied non-flushing memtable size to total permitted size for triggering a flush of the largest memtable. Larger values mean larger flushes and less compaction, but also less concurrent flush activity, which can make it difficult to keep your disks saturated under heavy write load. If not set the value will be calculated as $1/(1 + \text{memtableFlushWriters})$
memtableFlushWriters	No		Valid integer	Sets the number of memtable flush writer threads. These threads

Option name	Mandatory	Default Value	Possible Values	Description
				<p>are blocked by disk I/O, and each one holds a memtable in memory while blocked. If your data directories are backed by SSD, increase this setting to the number of cores.</p> <p>If not set the value will be calculated as (Smaller of number of disks or number of cores with a minimum of 2 and a maximum of 8)</p>
memtableHeapSize	No		Valid integer	<p>Total permitted memory (in Mb) to use for memtables. Triggers a flush based on</p> <p>memtableCleanupThreshold. Cassandra stops accepting writes when the limit is exceeded until a flush completes</p> <p>If not set the value will be calculated as (1/4 heap)</p>
memtableOffheapSpace	No		Valid integer	<p>If not set the value will be calculated as (1/4 heap)</p>
fileCacheSize	No		Valid integer	<p>Total memory to use for SSTable-reading buffers.</p> <p>If not set the value will be calculated as (Smaller of 1/4 heap or 512)</p>
authenticator	No	org.apache.cassandra.auth.AllowAllAuthenticator	org.apache.cassandra.auth.AllowAllAuthenticator, org.apache.cassandra.auth.PasswordAuthenticator	<p>The authentication backend</p> <p>http://docs.datastax.com/en/cassandra/2.1/cassandra/security/secure_about_native_authenticate_c.html</p>
authorizer	No	org.apache.cassandra.auth.AllowAllAuthorizer	org.apache.cassandra.auth.AllowAllAuthorizer, org.apache.cassandra.auth.CassandraAuthorizer	<p>The authorization backend</p> <p>http://docs.datastax.com/en/cassandra/2.1/cassandra/authorizer/cassandra_authorize_c.html</p>

Option name	Mandatory	Default Value	Possible Values	Description
				cassandra/security/secure_about_native_authorize_c.html
permissionsValidity	No	2000	Valid integer	How long (in milliseconds) permissions in cache remain valid. Depending on the authorizer, such as <i>org.apache.cassandra.auth.CassandraAuthorizer</i> , permissions can be resource intensive. This setting disabled when set to 0 or when <i>org.apache.cassandra.auth.AllowAllAuthorizer</i> is set.
permissionsUpdateInterval	No		Valid integer	Refresh interval (in milliseconds) for permissions cache (if enabled). After this interval, cache entries become eligible for refresh. On next access, an async reload is scheduled and the old value is returned until it completes. If permissionsValidity , then this property must be non-zero If not set the value will be the same like permissionsValidity
writeTimeout	No	2000	Valid long	The time that the coordinator waits for write operations to complete
readTimeout	No	5000	Valid long	The time that the coordinator waits for read operations to complete
rangeTimeout	No	10000	Valid long	The time that the coordinator waits for sequential or index scans to complete
counterWriteTimeout	No	5000	Valid long	The time that the coordinator waits for counter writes to complete

Option name	Mandatory	Default Value	Possible Values	Description
casContentionTimeout	No	1000	Valid long	The time that the coordinator continues to retry a CAS (compare and set) operation that contends with other proposals for the same row.
truncateTimeout	No	60000	Valid long	The time that the coordinator waits for truncates (remove all data from a table) to complete. The long default value allows for a snapshot to be taken before removing the data. If autoSnapshot is disabled (not recommended), you can reduce this time.
requestTimeout	No	10000	Valid long	The default time for other miscellaneous operations
encryption.server.internode	No	none	none, all, dc, rack	Enable or disable inter-node encryption. You must also generate keys and provide the appropriate key and trust store locations and passwords. No custom encryption options are currently enabled http://docs.datastax.com/en/cassandra/2.1/cassandra/security/secureSSLNodeToNode_t.html
encryption.server.keystore	No	conf/.keystore	Valid path	The location of a Java keystore (JKS) suitable for use with Java Secure Socket Extension (JSSE), which is the Java version of the Secure Sockets Layer (SSL), and Transport Layer Security (TLS)

Option name	Mandatory	Default Value	Possible Values	Description
				protocols. The keystore contains the private key used to encrypt outgoing messages
encryption.server.keystorePassword	Yes	cassandra		Password for the keystore
encryption.server.truststore	Yes	conf/.truststore	Valid path	Location of the truststore containing the trusted certificate for authenticating remote servers
encryption.server.truststorePassword	Yes	cassandra		Password for the truststore
encryption.client.enabled	No	false	true, false	<p>Enable or disable client-to-node encryption. You must also generate keys and provide the appropriate key and trust store locations and passwords. No custom encryption options are currently enabled</p> <p>http://docs.datastax.com/en/cassandra/2.1/cassandra/security/secureSSLClientToNode_t.html</p>
encryption.client.keystore	Yes	conf/.keystore	Valid path	The location of a Java keystore (JKS) suitable for use with Java Secure Socket Extension (JSSE), which is the Java version of the Secure Sockets Layer (SSL), and Transport Layer Security (TLS) protocols. The keystore contains the private key used to encrypt outgoing messages
encryption.client.keystorePassword	Yes	cassandra		Password for the keystore. This must match the password used when generating the keystore and

Option name	Mandatory	Default Value	Possible Values	Description
				truststore.
encryption.client.truststore	No	conf/.truststore	Valid path	Set if encryption.client.clientAuth is true
encryption.client.truststorePassword	No	<truststore_password>		Set if encryption.client.clientAuth is true
encryption.client.clientAuth	No	false	true, false	Enables or disables certificate authentication