



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.

Conversation Manager Product Guide

[Sizing Calculator](#)

12/14/2025

Sizing Calculator

Before deploying Context Services (CS) for your Conversation Manager solution, you must estimate the size of solution that will be able to handle the expected user load. Genesys recommends that you download the **CS Storage Sizing Calculator**, an Excel spreadsheet that you can use to help calculate the number of Context Services nodes required for your production deployment.

CS Storage Sizing Calculator

[CS-StorageSizing.xlsx](#)

Using the Sizing Calculator Spreadsheet

To use the sizing calculator, The following information is required:

- Approximate event size in bytes for event start/end/update of service/state/task (including extension data). You can extract this information from the JSON of the events.
- Number of update events per services/states/tasks
- Number of states/tasks per service
- Number of Cassandra nodes and replication factor
- Retention policy in days and number of conversation per day

The sizing calculator takes into account Cassandra storage specifics such as Replication Factor (storage size is multiplied by 3 as default replication), compression (ratio is estimated to 0.40 by default), and compaction overhead which requires more storage during cycles of data re-organization by Cassandra.

Input Example

Genesys Context Services - Sizing Input Parameters			
Events Parameter	Event Size (bytes)	Number of Events	Comments
service start event	246	1	
service end event	227		
service update event(s)	291	0	
state start event	210	5	
state end event	225		
state update event(s)	192	0	
task start event	208	10	
task end event	228		
task update event(s)	165	0	
Volume Parameter	Value	Comments	
Number of Conversations per Day	160000		
Conversations life (days)	90		
Number of storage nodes	4		
Replication Factor	3		
Constants	Value	Comments	
Compression ratio	0,50		
CommitLog Size (MB)	8192,00		

Output Example

Calculations		
Description	Value	Comments
Number of events per conversation	32	
Size of conversation in storage (bytes)	7008	
Number of Activities	2560000	
Description	Total	Comments
Estimated disk size (MB)	80117,80	

Genesys Context Services - Sizing Output	
Data Size (MB)	80 118
Cluster Data Size (MB)	240 353
Data Size per node MB)	60 088
Disk usage per node (MB)	128 369
Grand Total (MB)	513 475

Example of Hardware Sizing

This example shows an initial GMS and CS cluster deployment and ensures that performance is correct in scale.

In this example, GMS and CS 8.5.003.00 and later instance is used with external Cassandra 2.0.8.

Host	Specifications	Software Components
4 similar hosts	Windows Server 2008 R2 Standard SP 1 Intel Xeon X3440 @ 2.53 GHz, quad core 8 GB RAM Low Price disks	JMeter 2.11 GMS 8.5.003.00 and later Xmx is 2G for java Cassandra 2.0.8 external
Genesys Framework	Windows Server 2008 R2 Standard SP1	Runs Genesys Framework/Suite 8: <ul style="list-style-type: none">• Config Server: 8.1.100.14• DB Server: 8.1.100.04• Message Server: 8.1.100.02

GMS and CS are configured to handle a constant average CPU of 50%. This allows for burst and Cassandra background activities to happen correctly, high load prevent cleaning to occur correctly, and increasing latencies over time.

Configuration

Start from a database with 2.5GB data on each node. Complex statistics (requiring fetch of "start" event during processing of "end" event) are enabled for States/Tasks. All statistics are available for Services.

Special settings and enhancements:

- Cassandra 2.0.8
- Replication Factor = 2
- Consistency Level = 1/1 (read/write)
- Correctly Balanced Cassandra node tokens for each node to own 25% data

Options:

- `cvview/data-validation=false`

JMeter:

- Modified to add constant time between events (200ms), allowing constant load of GMS nodes

Results

Captured measure	Value	Comments
Database size	68 GB / node (after 120h+ run by set of 48h) = 272 GB	272 GB corresponds to 16.5 M conversations or 264 M activities that make a ratio of 0.54KB for each activity. Multiply by replication factor to reflect correct disk size.
Throughput	Approximately 1000 tx/s (approximately 30.5 scenario/s)	Constant through 24h testing
Size of conversation	Approximately 10kB in JSON	This matches a conversation of 16 activities, with extension data as returned by Query Service, State, and Tasks by CustomerId
CPU	Approximately 20% to 60%	Each node was around an average of 30% with peaks to 60% from time to time
Memory	2.4GB	For each GMS process, less than 2 GB for Cassandra processes.
I/O Disc	Approximately 3 MB/s (Cassandra measure)	Disk I/Os fluctuate a lot (between 0 MB/s and 20 MB/s)
I/O Network	Approximately 3 MB/s	Between 20Mbps and 60Mbps.

Conclusion

This test showed successful long running tests for more than 48h with an average of 1000 tx/s on a 4 node cluster GMS with external Cassandra 2.0.8. To scale on higher throughput, you should allow more nodes into the cluster.

In the testing environment, it was possible to have up to 2x better throughput but Cassandra could not sustain very long before getting late in compaction state, resulting in increase of latencies after a few hours load (4-8h).

The throughput tested in this scenario is about 2x maximum speed required (approximately 25 CAPS, approximately 400 tx/s expected).