



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

Sizing Information

12/18/2025

Sizing Information

Contents

- **1 Sizing Information**
 - 1.1 Input Data for Load Estimation
 - 1.2 Basic Load Estimation
 - 1.3 Peak Load Estimation
 - 1.4 Minimum Solution Size
 - 1.5 Glossary
 - 1.6 Useful Links

Before deploying the Genesys Web Engagement (GWE) solution to your production site, you must estimate the size of solution that will be able to handle the expected user load. Genesys recommends that you download the **GWE Sizing Calculator**, an Excel workbook that you can use to help calculate the number of Web Engagement Server nodes required for your production deployment.

Click [HERE](#) to download the **GWE Sizing Calculator**. See download tips if the download was not started automatically.

The process of estimation starts from input values, usually given in the terms of business operations; for example, daily visit rates, or page view rates. Using some math, and having in mind the workflow that is applied to the input traffic, you can then produce the expected load values in terms of requests per second. Applying these values to the experimentally produced measurements, you can estimate the size of the solution required to be deployed.

Important

The exact deployment architecture and solution size will vary depending on your hardware equipment, and that the deployed system can be fine-tuned to get the best performance on given equipment and with given user load. However, the estimation can give some basic ideas for the deployment.

Input Data for Load Estimation

To estimate the load, you need to know the following data:

- Average visits rate and page views rate (per hour)
- Maximum visits rate and page views rate (per hour)

To estimate your disk space consumption, it is helpful to have the following information about the configuration of the solution:

- Number of business events per page or visit
- Portion of categorized events

You can use these numbers as inputs into the **GWE Sizing Calculator**.

Basic Load Estimation

Visits Rate Estimation

Having the maximum or average visits count and page view count per hour (or day), you can get visits rate and visit depth:

- Visits per second = Visits per hour / 3600 = Visits per Day / 24*3600
- Visit depth = Page views per hour / Visits per hour = Page views per day / Visits per day

Events Rate Estimation

Having the following input data about solution configuration:

- Average number of business events (timeout, search, and so on) per page – page custom events
- Average number of user events (SignIn/SignOut/UserInfo) per visit – custom visit events
- The fact that every visit produces one system event (visit started) and every page view produces two system events (PageEntered and PageExited)

The calculator can now estimate the average event rate produced by user visits:

- Events per second = Visits per second * (custom visit events + 1) + (Visits per Seconds * Visit Depth) * (custom page events + 2)

Requests Per Second

Assuming that each visit also contains two requests for loading client scripts and DSL, and every page view also contains one request for categorization info:

- Requests per second = Visits per second * 1 + (Visits per second * Visit depth) * 2 + Events per second

Peak Load Estimation

Having only average values for visit and page views, you can only predict some average load, and, consequently, only an *average* solution size that will be able to handle such load. However, the user load is not evenly distributed during the day, so you must estimate possible peak load during busy hours.

To cover that scenario, you can take your Visits as the Poisson distributed flow. To estimate the possible peak load, use the following calculation:

- Maximum Visits per Second = Average Visits per Second + 4 * SQRT (Average Visits per Second)

Now, having the Maximum Visits per Second value, it is possible to estimate the Maximum Events per Second and the Maximum Requests per Second, using the same approach as for Basic Load, but replacing average visits rate with maximum visits rate.

Load Estimation Example

Visits Rate Estimation

Having the maximum visits and page view rates per hour, you're calculating visits rate and visit's depth. For example, having the following values:

- maximum number of visits is 3000 per hour
- maximum number of page views is 32000 per hour

The estimated visit depth is $32000/3000 = 10.7$ pages.

- Average load: 3000 visits per hour = 0.83 visits per second
- Maximum load: $0.83 + 4*\sqrt{0.83} = 4.5$ visits per second

Events Rate Estimation

Given two business events per page, and about 10 percent of visits producing additional events:

- Average events per second = $0.83*(1 + 0.1) + (0.83*10.7)*(2 + 2) = 37$ events per second

Requests per Second

Assuming that every visit generates one request for script and another for data, and every page view generates one request for categories:

- Average requests per second = $37 \text{ events} + 0.83*2 + (0.83*10.7)*1 = 47$ requests per second

Peak Load Estimation

Using Maximum Visits per second instead of Average Visits per Second, the following values can be calculated:

- Maximum events per second = $4.5*(1 + 0.1) + (4.5*10.7)*(2 + 2) = 197$ events per second
- Maximum requests per second = $197 + 4.5*2 + (4.5*10.7)*1 = 254$ requests per second

Minimum Solution Size

The solution deployed should handle all user input, and have some N+1 redundancy. The next two sections give the minimum solution size for external and embedded versions of Cassandra, while the third section contains information on disk space requirements.

External Cassandra

- Web Engagement Server—two nodes (to support high availability and load balancing).
- Cassandra—three nodes (to provide data consistency and to allow a fault tolerance of one node). The consistency level must be `LOCAL_QUORUM`.

For more help calculating the number of Cassandra nodes you need to support data consistency in the cluster, see <http://www.ecyrd.com/cassandrascalculator/>.

Embedded Cassandra

The Web Engagement Server cluster should contain at least three nodes (to support high availability and load balancing, and to provide data consistency and allow a fault tolerance of one node). The cluster must have a consistency level of `LOCAL_QUORUM`.

Important

Starting in 8.5.0, Embedded Cassandra mode is deprecated in Web Engagement; support for this mode will be discontinued in 9.0.

Disk Space Usage

Disk space usage directly depends on the following factors:

- The event rate in the solution
- The average count of the categories applied to each event
- The replication factor specified for Cassandra (and Elasticsearch)

Glossary

Visit—series of page views from the same uniquely identified browser. Note that pages opened in different tabs or windows of the same browser will be treated as part of the same visit.

Average visits per second—service load in terms of customer site.

Request—a single request to the service. A single page view produces a series of requests to the Web Engagement Server:

- Loading of monitoring script
- Loading of categorization information
- Sending information about **PageEntered**, **PageExited**, and business events back to the Web Engagement Server

Requests per second (RPS) – actual load in terms of service performance.

Useful Links

Refer to the following links for more information about planning an external Cassandra cluster:

- For help understanding the Cassandra architecture, see <http://docs.datastax.com/en/cassandra/2.2/cassandra/architecture/archTOC.html>

- For information about hardware considerations for Cassandra nodes, see <http://docs.datastax.com/en/cassandra/2.2/cassandra/planning/planPlanningHardware.html>
- For details about Cassandra cluster configuration, refer to <http://docs.datastax.com/en/cassandra/2.2/cassandra/initialize/initSingleDS.html>
- For more about Cassandra clusters and memory, see <http://docs.datastax.com/en/cassandra/2.2/cassandra/operations/opsTuneJVM.html>
- For more help calculating the number of Cassandra nodes you need to support data consistency in the cluster, see <http://www.ecyrd.com/cassandrascalculator/>.