



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.

Genesys Pulse Sizing

Genesys Pulse Hardware Sizing and Performance Information

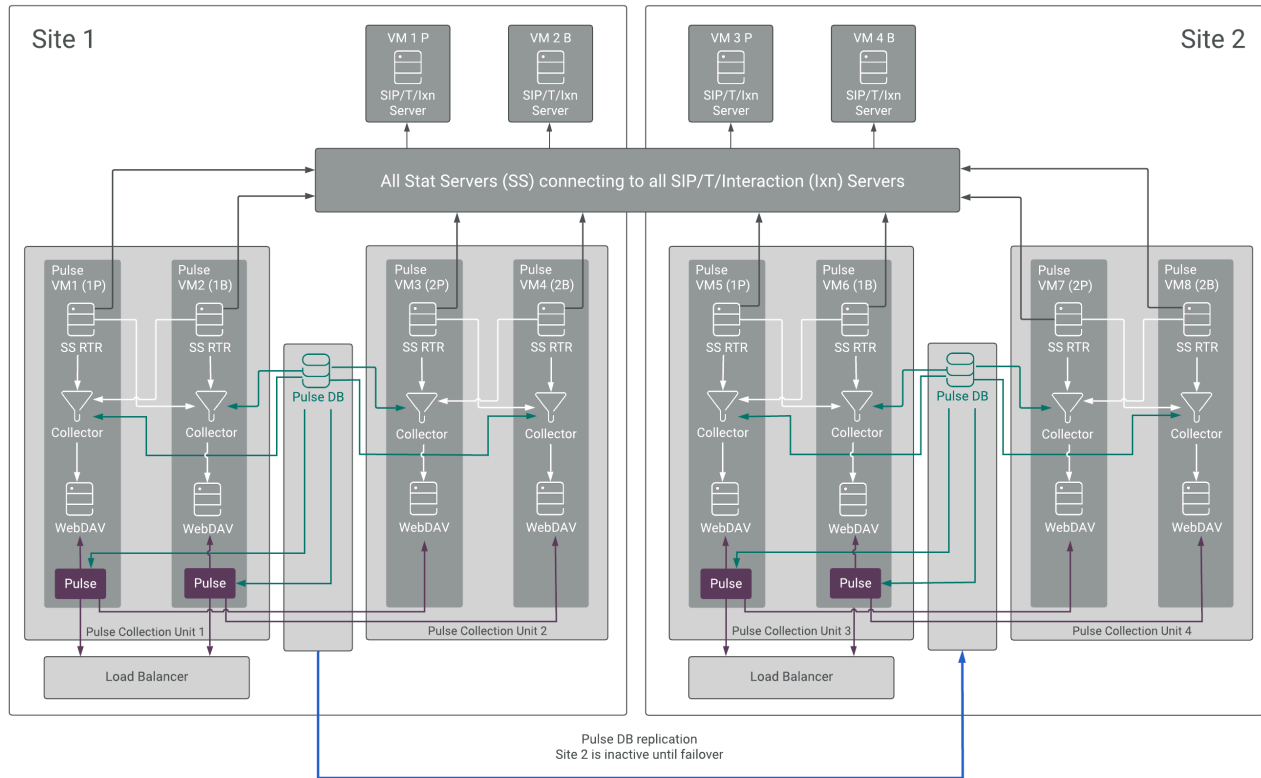
Genesys Pulse Hardware Sizing and Performance Information

This article describes the result of load testing performed using Genesys Pulse release 9.0.0 and serves as a guideline for Genesys Pulse capacity and resource usage. For a production deployment, you must test Genesys Pulse in your own environment under a production load to ensure its performance meets your expectations and is sized properly.

Contents

- [1 Genesys Pulse Hardware Sizing and Performance Information](#)
 - [1.1 Architecture Example](#)
 - [1.2 Environment Configuration](#)
 - [1.3 Dashboard Configuration](#)

Architecture Example



Genesys Pulse Architecture Example



Important

Genesys recommends to have two Configuration Server Proxies (Primary for this VM and backup for another VM) on each Genesys Pulse VM. At least one dedicated Configuration Server Proxy is required for two Genesys Pulse instances to provide appropriate performance.

1. Each Genesys Pulse Collector can handle about 300K statistics with a 10-second refresh rate.
2. Genesys Pulse Collector adjusts the load based on the refresh rate.

Important

If you use change-based notifications for anything other than Agent Current State, you might have a significant number of notifications for each of these statistics per second. You must configure sensitivity appropriately and account for these additional notifications in sizing.

3. If you use Genesys Pulse Collector Cluster configuration, the statistics are dynamically distributed across the Collectors in the cluster.
4. Genesys recommends you have a dedicated Stat Server for each Genesys Pulse Collector. If you to have multiple Collectors for each Stat Server, contact Customer Care to evaluate your solution.
5. Each Genesys Pulse Collector supports only a single connection to Stat Server (HA pair). All Stat Servers used for all Collectors in a cluster must be connected to the same sources (for example, T-Servers, SIP Servers, Interaction Servers).
6. Each Genesys Pulse can render the statistics and views across all Genesys Pulse Collectors in a cluster.
7. You can have multiple Genesys Pulse Collectors for each Genesys Pulse.
8. Determine the number of Genesys Pulse instances depending on the number of concurrent users and numbers of Genesys Pulse Collectors depending on the number of Stat Request notifications per second.
9. For High Availability environments, you need a pair of Genesys Pulse Collectors to present a node.
10. Genesys Pulse components (Genesys Pulse, Genesys Pulse Collector, Stat Server, DB Server) are operational in Site 1 (primary region) and Site 2 (failover region).
11. The Genesys Pulse is normally accessed only in Site 1.
12. Genesys Pulse in Site 2 can provide reports, but users only connect to the Site 2 if the failover from Site 1 is initiated.
13. Genesys Pulse database is replicated from Site 1 to Site 2 through a periodic backup and restore process.
14. WebDav is used when Genesys Pulse needs to pull snapshot data from Genesys Pulse Collector that is not installed on the local (same) host. Genesys supports lighttpd http server with lighttpd-mod-webdav.

Which Genesys Pulse Collector handles which requests? Genesys Pulse Collectors are grouped into active-active HA pairs so Genesys Pulse Collectors from Node 1 always process the same requests (for example, each subscribes to same statistics). Similarly, Genesys Pulse Collectors from Node 2 process the same stats.

Which widget uses which node? They function as a round robin of widget distribution - basically, half the widgets go to Node 1 and half go to Node 2.

All Stat Servers are configured in HA pairs and they work in active-active mode. For example, the backup Stat Server calculates statistics the same way as Primary.

Both Genesys Pulse Collectors from the same node are connected to both Stat Servers from this node in terms in TCP connections. However, configure them in Configuration Server with VM1 SS as primary and VM2 SS as backup. Then both Genesys Pulse Collectors on VM1 and on VM2 are connected to primary Stat Server (which is VM1 SS), so there is only one connection. When Genesys

Pulse Collector starts it sees that it is connected to Primary VM1 SS that has a backup configured and then it connects to backup automatically.

Environment Configuration

2 nodes Amazon AWS:

	vCPU	ECU	Memory (GiB)	Instance Storage (GB)
c3.8xlarge	32	108	60	2 x 320 SSD

Each node has cluster of 3 Genesys Pulse Collectors, 3 Stat Servers and 1 Genesys Pulse instance.

Genesys Pulse Collectors and Stat Servers are connected to each other in HA mode (1 Genesys Pulse Collector on each node is connected to Stat Server on the same node and Stat Server on another node).

Key performance impacting dimensions of the environment:

Widget and Statistic notification rate is 10 seconds.

Important

Change based statistics (and quick updates feature of Genesys Pulse) are not used. If you use change based statistics in your configuration (and they are not Agent State statistics) make sure you set sensitivity such that Stat Server sends notifications no often than once a second. Amount of Stat Server notifications with change based statistics can be unpredictable and be an order of magnitude greater than with time based so they must be carefully tuned (using sensitivity) and tested.

Number of layouts (widgets) = 2400.

Total Number of Statistics Requests to Stat Servers = Sum of (Number of Objects * Number of Stat) across all widgets = 1,740,000 (1.7M)

As a result we get 1.7M notifications from Stat Servers every 10 seconds so 170K updates per second.

Genesys Pulse DB size is 31260672 bytes for ~ 2500 layouts so 12.5Kbytes per layout (widget)

This load was split between 3 Genesys Pulse Collectors/Stat Server pairs (in cluster).

Each Genesys Pulse Collector in cluster processed 800 layouts and consumed the following:

- Memory (about 2K per statistic for 64bit version): up to 5GB
- CPU up to 800% (this means 8 cores)

- Bandwidth: ~1Gbit/s between Genesys Pulse Collector and Stat Server (on the other node)

Important

Because Genesys Pulse Collector queues the messages that come from Stat Server the memory consumption above is achieved when there is no backlog - i.e Collector is able to process all layouts and save them as it gets the data.
If Genesys Pulse Collector cannot process all Stat Server messages it receives, it expands its memory size (use the **statistic-request-handling/max-stat-data-queue-size** option to control it).

Genesys Pulse resource consumption:

- Memory usage 9MB per logged in user session
- CPU usage peaks at 800% (8 cores)
- User with 5 widgets with refresh rate 10 sec (total sum of number of statistic * number of object in these widgets is 5781) receives about 600KB per minute (so average traffic is about 80Kbit/s).
- Genesys Pulse instance handled 200-300 concurrent users

Dashboard Configuration

For the best experience Genesys recommends not to exceed the following number of Grid widgets per dashboard:

Browser	Number of grid widgets
Google Chrome	6
Mozilla Firefox	6
Microsoft Internet Explorer 11 (support ends in release 9.0.008.02)	3
Microsoft Edge	6

with the total number of displayed grid columns <= 30 per dashboard.

Measurements were performed using the following client configuration:

- (VM) Client 1:
 - Windows 7 Professional 64-bit RAM 6.00 GB, intel Xeon (R) CPU E7-4850 @ 2.00 GHz
 - Screen resolution: 1920x1080
 - Total video memory: 32 MB
 - Browsers: Internet Explorer 11, Chrome 60.0.8112.90, Firefox 55.0.2

- (VM) Client 2:
 - Windows 10 Professional 32-bit, 5.95 GB, intel Xeon (R) CPU E7-4850 @ 2.00 GHz
 - Screen resolution: 1920x1080
 - Total video memory: 32 MB
 - Browsers: Microsoft Edge 40.15063.0.0