



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 Info Mart Private Edition Guide

GSP metrics and alerts

12/26/2025

---

## Contents

- [1 Metrics](#)
- [2 Alerts](#)

---

Find the metrics GSP exposes and the alerts defined for GSP.

### Related documentation:

- 

Service	CRD or annotations?	Port	Endpoint/Selector	Metrics update interval
GSP	PodMonitor	9249	Endpoint: / Selector: <pre>matchLabels:   app: {{ template "gsp.fullname" . }}</pre> where the value of <code>gsp.fullname</code> depends on deployment parameters such as Helm release name, <code>.Values.fullnameOverride</code> , and <code>.Values.nameOverride</code> .	30 seconds

See details about:

- GSP metrics
- GSP alerts

## Metrics

GSP exposes some standard Apache Flink and Kafka metrics as well as Genesys-defined metrics, which are exposed via the Flink API. Therefore, all GSP metrics start with the prefix **flink\_** but in some cases the values are calculated by GSP.

You can query Prometheus directly to see all the metrics Flink and the Flink Kafka connector expose through GSP.

- For full information about the standard Flink metrics, see the Apache Flink documentation.
- For full information about the Kafka metrics, see the Apache Kafka or Confluent Kafka documentation.

The following metrics are likely to be particularly useful. The naming convention is `_`. Genesys does not commit to maintain other currently available GSP metrics not documented on this page.

Metric and description	Metric details	Indicator of
<b>flink_taskmanager_job_task_operations_errors_numInvalidRecords</b>		Error

Metric and description	Metric details	Indicator of
Number of invalid input records.	<b>Type:</b> Gauge <b>Label:</b> <b>Sample value:</b> 0	
<b>flink_jobmanager_numRunningJobs</b> Number of running Flink jobs. If less than 1, there is a problem.	<b>Unit:</b> <b>Type:</b> Gauge <b>Label:</b> <b>Sample value:</b> 1	Error
<b>flink_taskmanager_job_task_operator_user_errors_numOversizedMessages</b> Number of messages exceeding the <b>max.request.size</b> Kafka option.	<b>Unit:</b> <b>Type:</b> Gauge <b>Label:</b> • operator_name <b>Sample value:</b> 0	Error
<b>flink_taskmanager_job_task_operator_tenant_error_total</b> Number of issues encountered, such as errors or warnings.	<b>Unit:</b> <b>Type:</b> Gauge <b>Label:</b> • operator_name • tenant • error <b>Sample value:</b>	Error
<b>flink_taskmanager_job_task_operator_currentInputWatermark</b> The last watermark received by this operator/task, in milliseconds since the Unix Epoch (00:00:00 UTC on 1 January 1970). <b>Note:</b> For operators/tasks with two inputs, this is the earlier of the last received watermarks.	<b>Unit:</b> milliseconds <b>Type:</b> Gauge <b>Label:</b> • operator_name <b>Sample value:</b>	Latency
<b>flink_taskmanager_job_task_operator_currentOutputWatermark</b> The last watermark this operator has emitted, in milliseconds since the Unix Epoch.	<b>Unit:</b> milliseconds <b>Type:</b> Gauge <b>Label:</b> • operator_name: • Sink:_Agent_State_Facts • Sink:_Interaction_Facts <b>Sample value:</b>	Latency
<b>flink_taskmanager_job_task_operator_records_lag_max</b> The maximum lag in terms of the number of records for any partition in this window. An increasing value over time is your best indication that the consumer group is not	<b>Unit:</b> <b>Type:</b> Gauge <b>Label:</b> <b>Sample value:</b>	Latency

Metric and description	Metric details	Indicator of
keeping up with the producers.		
<b>flink_taskmanager_job_task_operator_records_consumed_rate</b> The average number of records consumed per second.	<b>Unit:</b> <b>Type:</b> Gauge <b>Label:</b> <b>Sample value:</b>	Traffic
<b>flink_taskmanager_job_task_operator_numCallsCreated</b> Total number of EventCallCreated events GSP received since it started processing.	<b>Unit:</b> <b>Type:</b> Gauge <b>Label:</b> <b>Sample value:</b>	Traffic
<b>flink_taskmanager_job_task_operator_numCallsCreatedPerSecond</b> Number of EventCallCreated events per second (CPS).	<b>Unit:</b> <b>Type:</b> Gauge <b>Label:</b> <b>Sample value:</b>	Traffic
<b>flink_taskmanager_job_task_operator_numThreadsCreated</b> Total number of CallThreads GSP received since it started processing.	<b>Unit:</b> <b>Type:</b> Gauge <b>Label:</b> <b>Sample value:</b>	Traffic
<b>flink_taskmanager_job_task_operator_numCallThreadsCreatedPerSecond</b> Number of CallThreads per second (CTHPS).	<b>Unit:</b> <b>Type:</b> Gauge <b>Label:</b> <b>Sample value:</b>	Traffic
<b>flink_taskmanager_job_task_operator_numChainsProcessed</b> Total number of EventOCSCChainStartProcessing events GSP received since it started processing.	<b>Unit:</b> <b>Type:</b> Gauge <b>Label:</b> <b>Sample value:</b>	Traffic
<b>flink_taskmanager_job_task_operator_numChainsProcessedPerSecond</b> Number of EventOCSCChainStartProcessing events per second (CPS).	<b>Unit:</b> <b>Type:</b> Gauge <b>Label:</b> <b>Sample value:</b>	Traffic
<b>flink_(job task)manager_Status_JVM_CPU_Load</b> The recent CPU usage for the JVM process. The value is a double in the [0.0,1.0] interval, where a value of 0.0 means that none of the CPUs were running threads from the JVM process, while a value of 1.0 means that all CPUs were actively running threads from the JVM 100% of the time during the recent period being observed. A negative value means usage data is not available. For more information, see <a href="https://docs.oracle.com/javase/7/docs/jre/api/management/extension/com/sun/management/OperatingSystemMXBean.html#getProcessCpuLoad()">https://docs.oracle.com/javase/7/docs/jre/api/management/extension/com/sun/management/OperatingSystemMXBean.html#getProcessCpuLoad()</a> .	<b>Unit:</b> <b>Type:</b> Gauge <b>Label:</b> • pod <b>Sample value:</b>	Saturation
<b>flink_(job task)manager_Status_JVM_Memory_Direct_TotalCapacity</b>		Saturation

Metric and description	Metric details	Indicator of
The total capacity of all buffers in the direct buffer pool.	<b>Type:</b> Gauge <b>Label:</b> <ul style="list-style-type: none"> <li>pod</li> </ul> <b>Sample value:</b>	
<b>flink_(job task)manager_Status_JVM:Memory_Direct_MemoryUsed</b> The amount of memory used by the JVM for the direct buffer pool.	<b>Unit:</b> bytes <b>Type:</b> Gauge <b>Label:</b> <ul style="list-style-type: none"> <li>pod</li> </ul> <b>Sample value:</b>	Saturation
<b>flink_(job task)manager_Status_JVM:Memory_NonHeap_Max</b> The maximum amount of non-heap memory that can be used for memory management.	<b>Unit:</b> bytes <b>Type:</b> Gauge <b>Label:</b> <ul style="list-style-type: none"> <li>pod</li> </ul> <b>Sample value:</b>	Saturation
<b>flink_(job task)manager_Status_JVM:Memory_NonHeap_Used</b> The amount of non-heap memory currently used.	<b>Unit:</b> bytes <b>Type:</b> Gauge <b>Label:</b> <ul style="list-style-type: none"> <li>pod</li> </ul> <b>Sample value:</b>	Saturation
<b>flink_(job task)manager_Status_JVM:Memory_Heap_Max</b> The maximum amount of heap memory that can be used for memory management.	<b>Unit:</b> bytes <b>Type:</b> Gauge <b>Label:</b> <ul style="list-style-type: none"> <li>pod</li> </ul> <b>Sample value:</b>	Saturation
<b>flink_(job task)manager_Status_JVM:Memory_Heap_Used</b> The amount of heap memory currently used.	<b>Unit:</b> bytes <b>Type:</b> Gauge <b>Label:</b> <ul style="list-style-type: none"> <li>pod</li> </ul> <b>Sample value:</b>	Saturation

## Alerts

The alerts are based on Flink and Kubernetes cluster metrics.

---

The following alerts are defined for GSP.

Alert	Severity	Description	Based on	Threshold
GspFlinkJobDown	Critical	Triggered when the GSP Flink job is not running (number of running jobs equals to 0 or metric is not available)	flink_jobmanager_numRunningJobs	For 5 minutes
GspOOMKilled	Critical	Triggered when a GSP pod is restarted because of OOMKilled	kube_pod_container_status_restarts_total	0
GspNoTmRegistered	Critical	Triggered when there are no registered TaskManagers (or metric not available)	flink_jobmanager_numRegisteredTaskManagers	For 5 minutes
GspUnknownPerson	High	Triggered when GSP encounters unknown person(s)	flink_taskmanager_job_taskmanager_error_total	For 5 minutes