

# **GENESYS**

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.

## Voice Microservices Private Edition Guide

Voice SIP Proxy Service metrics and alerts

### Contents

- 1 Metrics
- 2 Alerts

Find the metrics Voice SIP Proxy Service exposes and the alerts defined for Voice SIP Proxy Service.

Service	CRD or annotations?	Port	Endpoint/Selector	Metrics update interval
Voice SIP Proxy Service	Supports both CRD and annotations	11400	http://:11400/metrics	30 seconds

#### See details about:

- Voice SIP Proxy Service metrics
- Voice SIP Proxy Service alerts

#### Metrics

Voice SIP Proxy Service exposes Genesys-defined, SIP Proxy Service-specific metrics as well as some standard Kafka metrics. You can query Prometheus directly to see all the metrics that the SIP Proxy Service exposes. The following metrics are likely to be particularly useful. Genesys does not commit to maintain other currently available SIP Proxy Service metrics not documented on this page.

Metric and description	Metric details	Indicator of
sipproxy_requests_total  Total number of received requests.	Unit: N/A  Type: counter Label: method Sample value:	Traffic
sipproxy_rejected_requests_tot The total number of rejected requests.	Unit: N/A cal Type: counter Label: Sample value:	Errors
sipproxy_requests_processed_s The total number of received requests that were processed by SIP Proxy itself.	thrit: NA  Type: counter Label: method Sample value:	Traffic
sipproxy_requests_forwarded_t The total number of forwarded requests.	Unit: N/A  otal type: counter Label: method, request_direction, sip_node_id Sample value:	Traffic
sipproxy_requests_sip_node_re	s <b>eletted</b> /total	Errors

Metric and description	Metric details	Indicator of
Total count of sip-node reselection.	Type: counter Label: Sample value:	
sipproxy_responses_forwarded_ Total count of forwarded responses.	Unit: N/A  total - lype: counter  Label: method, sip_node_id, request_direction Sample value:	Traffic
sipproxy_response_latency SIP response latency.	Unit:  Type: histogram Label: le, sip_node_id, request_direction, target, node_in_cache Sample value:	Latency
sipproxy_register_processed_to  Total number of REGISTER requests that SIP Proxy received for processing.	tanit: N/A  Type: counter Label: Sample value:	Traffic
sipproxy_register_rejected_total  Total number of REGISTER requests for processing that were rejected.	Unit: N/A  Type: counter Label: Sample value:	Errors
sipproxy_calls_per_second_cour Current calculated calls per second.	Unit: N/A nt Type: gauge Label: Sample value:	Saturation
sipproxy_active_sip_nodes_cour Current number of active SIP nodes.	Unit: N/A nt Type: gauge Label: Sample value:	
sipproxy_sip_nodes_count  Current number of discovered SIP nodes.	Unit: N/A  Type: gauge Label: Sample value:	
sipproxy_tenants_count  Current count of discovered tenants.	Unit: N/A  Type: gauge Label: Sample value:	
sipproxy_consul_record_process Current number of errors while processing records got from Consul.	SUnit: N/A sing_errors_count Type: counter Label: Sample value:	
sipproxy_consul_errors_count	Unit: N/A	

Metric and description	Metric details	Indicator of
Current number of Consul errors.	Type: counter Label: Sample value:	
sipproxy_sip_node_is_capacity_a	avalitible	
Indicates whether SIP node has available capacity or not.	Type: gauge Label: sip_node_id Sample value:	
service_version_info	Unit: N/A	
Displays the version of Voice SIP Proxy Service that is currently running. In the case of this metric, the labels provide the important information. The metric value is always 1 and does not provide any information.	Type: gauge Label: version Sample value: service_version_info{version="100.0.100001	06"}
sipproxy_health_level		
Health level of the SIP Proxy node:	Unit: N/A	
-1 - fail 0 - starting 1 - degraded 2 - pass	Type: gauge Label: Sample value:	
sipproxy_envoy_proxy_status	Unit: N/A	
Status of the Envoy proxy:	Type: gauge	
-1 – error 0 – disconnected 1 – connected	Label: Sample value: 1	
sipproxy_config_node_status	Unit: N/A	
Status of the Config node connection:	Type: gauge	
0 – disconnected 1 – connected	Label: Sample value: 1	
sip_server_transactions_created	a_ <b>LBit</b> ai N/A	
Total number of created server transactions.	Type: counter Label: Sample value:	Traffic
sip_client_transactions_created	_ <del>Utiti</del> N/A	
Total number of created client transactions.	Type: counter Label: Sample value:	Traffic
sip_server_transactions_deleted	a_ <b>UBita</b> i N/A	
Total number of deleted server transactions.	Type: counter Label: Sample value:	Traffic
sip_client_transactions_deleted	_totat N/A	Traffic

Metric and description	Metric details	Indicator of
Total number of deleted client transactions.	Type: counter Label: Sample value:	
sip_client_transactions_count  Current number of client transactions.	Unit: N/A  Type: gauge Label: Sample value:	Saturation
sip_server_transactions_count  Current number of server transactions.	Unit: N/A  Type: gauge Label: Sample value:	Saturation
sip_server_transactions_rejected  Total number of server transactions rejected for internal reasons.	Type: counter Label: Sample value:	Errors
sip_proxy_contexts_count  Current number of active SIP Proxy forwarding contexts.	Unit: N/A  Type: gauge Label: Sample value:	Saturation
sip_received_bytes_total  Total traffic received, measured in bytes.	Unit: bytes  Type: counter Label: transport Sample value:	Traffic
sip_sent_bytes_total  Total traffic sent, measured in bytes.	Unit: bytes  Type: counter Label: transport Sample value:	Traffic
sip_transport_errors_total  Total number of transport errors.	Unit: N/A  Type: counter Label: transport, address Sample value:	Errors
sip_stream_transport_wait_drai	n_total <sup>N/A</sup>	
Total number of requests to wait for drain events on stream transports.	Type: counter Label: Sample value:	
sip_stream_transport_flood_total  Total number of flood events on the stream transports.	aunit: N/A  Type: counter Label: Sample value:	
http_client_request_duration_se	edbnidsseconds	Latency

Metric and description	Metric details	Indicator of
The time duration between the HTTP client request and the response, measured in seconds.	Type: histogram Label: le, target_service_name Sample value:	
http_client_response_count The number of HTTP client responses received.	Unit: N/A  Type: counter Label: target_service_name, status Sample value:	Traffic
log_output_bytes_total  The total amount of log output, measured in bytes.	<pre>log_output_bytes_total{level="info",format=" manager"} 4184 log_output_bytes_total{level="info",format=" node"} 135</pre>	txt",module="sipproxy_node@sipproxy- txt",module="sipproxy_node@sipproxy@sip"} 181 json",module="sipproxy_node@config-
kafka_consumer_recv_message  Number of messages received from Kafka.	s_total  Type: counter Label: Sample value:	Traffic
kafka_consumer_error_total  Number of Kafka consumer errors.	Unit:  Type: counter Label: Sample value:	Errors
kafka_consumer_latency  Consumer latency is the time difference between when the message is produced and when the message is consumed. That is, the time when the consumer received the message minus the time when the producer produced the message.	Unit: Type: histogram Label: Sample value:	Latency
kafka_consumer_rebalance_total Number of Kafka consumer rebalance events.	al Unit: Type: counter Label: Sample value:	
kafka_consumer_state  Current state of the Kafka consumer.	Unit:  Type: gauge Label: Sample value:	
kafka_producermessages_tot	aUnit: Type: counter	Traffic

Metric and description	Metric details	Indicator of
Kafka.	Label: Sample value:	
kafka_producer_queue_depth	Unit:	
Number of Kafka producer pending events.	Type: gauge Label: kafka_location Sample value:	Saturation
kafka_producer_queue_age_sec	oldi: seconds	
Age of the oldest producer pending event in seconds.	Type: gauge Label: kafka_location Sample value:	
kafka_producer_error_total	Unit:	
Number of Kafka producer errors.	Type: counter Label: kafka_location Sample value:	Errors
kafka_producer_state	Unit:	
Current state of the Kafka producer.	Type: gauge Label: kafka_location Sample value:	
kafka_producer_biggest_event_	Unit: size	
Biggest event size so far.	Type: gauge Label: kafka_location, topic Sample value: 231	
kafka_max_request_size	Unit:	
Exposed config to compare with biggest event size.	Type: gauge Label: kafka_location Sample value: 1000000	
kafka_producer_dropped_event	Unit: number	
Number of dropped events.	Type: gauge Label: Sample value:	

## Alerts

The following alerts are defined for Voice SIP Proxy Service.

Alert	Severity	Description	Based on	Threshold
Too many Kafka pending events	Critical	Too many Kafka producer pending events for pod {{ \$labels.pod }}.	kafka_producer_que	Too many Kafka producer pending le depth events for service {{

Alert	Severity	Description	Based on	Threshold
		This alert means there are issues with SIP REGISTER processing on this voice-sipproxy.  Actions:  • Make sure there are no issues with Kafka or with the {{ \$labels.pod }} pod's CPU and network.		\$labels.container }} (more than 100 in 5 minutes).
		Actions:		
SIP server response time too high	Warning	<ul> <li>If the alarm is triggered for multiple sipproxynodes, make sure there are no issues on {{ \$labels.sip_node} }}.</li> <li>If the alarm is triggered only for sipproxynode {{ \$labels.pod }}, check to see if there is an issue with the service related to the topic (CPU, memory, or network overload).</li> </ul>	_id sipproxy_response_la	SIP response latency for more than 95% of messages forwarded to {{ \$labels.sip_node_idate} ffy ffy ff ff for sipproxy-node {{ \$labels.pod }}.
Pod status failed	Warning	Actions:  • Restart the pod and check to see if there are any issues with the pod after restart.	kube_pod_status_ph	Pod {{ \$labels.pod }} is in Failed astate.
Pod status Unknown	Warning	Pod {{ \$labels.pod }} is in Unknown state.	kube_pod_status_ph	Pod {{ \$labels.pod }} is in Unknown state for 5 minutes.

Alert	Severity	Description	Based on	Threshold
		Restart the pod and check to see if there are any issues with the pod after restart.		
Pod status Pending	Warning	Pod {{ \$labels.pod }} is in Pending state.  Actions:  Restart the pod and check to see if there are any issues with the pod after restart.	kube_pod_status_ph	Pod {{ \$labels.pod }} is in Pending state for 5 <sup>a</sup> ffiinutes.
Pod status NotReady	Critical	Pod {{ \$labels.pod }} is in NotReady state.  Actions:  Restart the pod and check to see if there are any issues with the pod after restart.	kube_pod_status_rea	Pod {{ \$labels.pod }} is in NotReady state for 5 adhinutes.
Container restarted repeatedly	Critical	Container { { \$labels.container } } was repeatedly restarted.  Actions:  • Check to see if a new version of the image was deployed. Also check for issues with the Kubernetes cluster.	kube_pod_container	Container { {
No sip-nodes available for 2 minutes	Critical	No sip-nodes are available for the pod {{ \$labels.pod }}.	sipproxy_active_sip_	No sip-nodes are available for the node to slabels pod }} for 2 minutes.

Alert	Severity	Description	Based on	Threshold
		<ul> <li>If the alarm is triggered for multiple services, make sure there are no issues with sip-nodes.</li> <li>If the alarm is triggered only for pod {{ \$labels.pod }}, check to see if there is any issues with the pod.</li> </ul>		
sip-node capacity limit reached	Warning	The sip-node {{     \$labels.sip_node_id     }} hit capacity limit on {{     \$labels.pod }}.  Actions:  If alarm is     triggered for     multiple     services make     sure there is no     issues with sip-     node {{         \$labels.sip_node         }}.  If alarm is     triggered only     for pod {{         \$labels.pod }}     check if there     is any issue     with the pod	sipproxy_sip_node_is _id	The sip-node {{  \$labels.sip_node_id }} hit capacity limit on {{  \$ \$ \$ \$ \$ \$ \$ \$ \$ \$ \$ \$ \$ \$ \$ \$ \$ \$
Pod CPU greater than 80%	Critical	Critical CPU load for pod {{ \$labels.pod }}.  Actions:  • Check whether the horizontal pod autoscaler has triggered	container_cpu_usage container_spec_cpu_	Container { {

Alert	Severity	Description	Based on	Threshold
		and the maximum number of pods has been reached.  • Check Grafana for abnormal load.  • Collect the service logs for pod {{ \$labels.pod }} and raise an investigation ticket.		
Pod CPU greater than 65%	Warning	High CPU load for pod {{ \$labels.pod }}.  Actions:  Check whether the horizontal pod autoscaler has triggered and the maximum number of pods has been reached.  Check Grafana for abnormal load.  Collect the service logs for pod {{ \$labels.pod }} and raise an investigation ticket.	container_cpu_usage container_spec_cpu_	Container { {  \$labels.container }} CPU usage
Pod memory greater than 80%	Critical	Critical memory usage for pod { { \$labels.pod } }.  Actions:  • Check whether the horizontal pod autoscaler	container_memory_v kube_pod_container_	Container { {     \$labels.container     }} memory usage     Verkiegaset-80%etor - smillufesequests_m

Alert	Severity	Description	Based on	Threshold
		has triggered and the maximum number of pods has been reached.  • Check Grafana for abnormal load.  • Restart the service for pod {{ \$labels.pod }}.		
Pod memory greater than 65%	Warning	Pod {{ \$labels.pod }} has high memory usage.  Actions:  Check whether the horizontal pod autoscaler has triggered and the maximum number of pods has been reached.  Check Grafana for abnormal load.  Collect the service logs for pod {{ \$labels.pod }} and raise an investigation ticket	container_memory_v kube_pod_container	Container { {  \$labels.container }} memory usage  VOX:18866666666666666666666666666666666666
Config node fail	Warning	The request to the config node failed.  Action:  • Check if there is any problem with pod {{ \$labels.pod }} and config node.	http_client_response	Requests to the config node fail for 5 consecutive minutes.