



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.

Setting up Genesys Multicloud CX Private Edition

Configuring logging

Contents

- 1 Logging approaches and configuration
 - 1.1 Logging architecture

Provides an overview of logging architecture in Genesys Multicloud CX private edition, different types of logging mechanisms, and related configurations.

Related documentation:

-
-

RSS:

- [For private edition](#)

Logging approaches and configuration

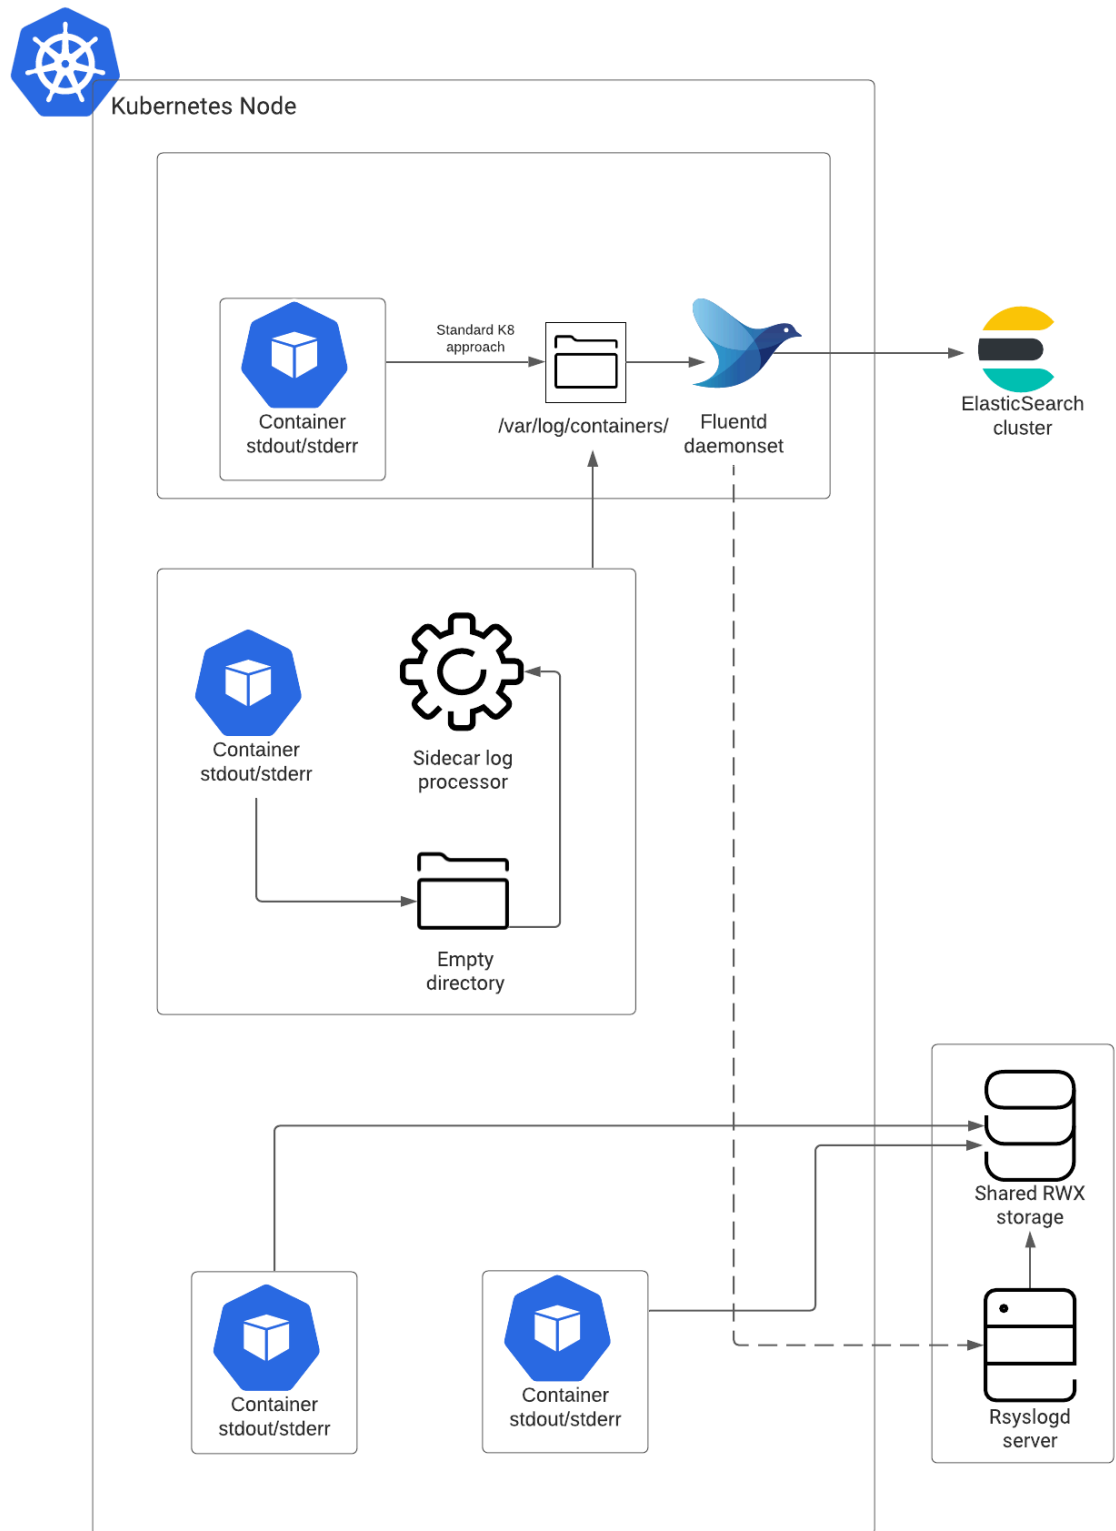
This section explains the approaches of logging used by Genesys Multicloud CX services to write log files that contain the important diagnostic information for various issues that may arise. Support of Genesys services rely on access to these application logs.

For more details, refer to [Solution-level logging approaches](#).

Logging architecture

This section explains the logging architecture of Genesys Multicloud CX private edition in detail.

Let's explore the logging architecture, the components involved, and its functionality through the following diagram.



Logging architecture components

Elasticsearch cluster

Elasticsearch cluster deployed on multiple node aggregates the structured logs from Fluentd and indexes them. This includes the logs from services that follow Secondary and Complementary logging methods. You can use a log visualizer tool like Kibana to view, search, or filter the indexed logs from Elasticsearch.

Fluentd / Fluent-bit

Fluentd is a log collector commonly used with container platforms. It collects logs from the cluster and forwards them to Elasticsearch or an externally accessible storage such as Rsyslog server or both depending on your configuration. Fluentd /Fluent-bit collects the application logs of Genesys Multicloud CX services from **/var/log/containers**. While deploying cluster wide logging each node Fluentd /Fluent-bit will be deployed to each node.

Shared RWX storage

The unstructured logs are directly written in the RWX shared storage. For services writing unstructured logs, you must mount PVC/PV. To access logs externally, use a server like NFS or S3.

Syslog server storage

Optionally, you can implement a syslog server to store the structured logs other than the Elasticsearch log store. Syslog server writes the logs in a flat file and enables you to share them externally. Genesys recommends Rsyslog server for this purpose, however you can select any syslog server of your choice. For more information, refer the deployment procedure.