



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.

GDL Deployment Guide

[Overview](#)

12/18/2025

Contents

- [1 GDL Overview](#)
- [2 Docker Compose Files](#)
- [3 Management Station API](#)
- [4 Volume and Log Files](#)

GDL Overview

Genesys Data Layer (GDL) is built on the Apache Kafka distributed data pipeline and streaming platform. Genesys products write data to GDL and Genesys InfoMart (GIM) reads from it to populate the GIM database, which is used for historical reporting.

Docker Compose Files

GDL consists of three components: Kafka, Zookeeper and Management station. Based on the number of components, the different cluster types are:

- Zookeeper - 1, Kafka - 1, and Management station.
- Zookeeper - 1, Kafka - 3, and Management station.
- Zookeeper - 3, Kafka - 3, and Management station.

These cluster types can be defined through docker-compose files available in the IP (Installation Package) files. The different docker-compose files are:

- docker-compose-single.yml - Creates 1 Zookeeper, 1 Kafka, and the Management Station.
- docker-compose-half.yml - Creates a cluster with 1 Zookeeper, 3 Kafka, and the Management Station.
- docker-compose-full.yml - Creates a cluster with 3 Zookeepers, 3 Kafka, and the Management Station.

The definitions and structure of Docker Compose files can be referred [here](#).

All the mandatory configuration parameters required for GDL lab versions are pre-configured in the docker-compose files.

Management Station API

The Management station API allows the Genesys Data Layer users to view and modify the GDL cluster properties and configuration through RESTful API. It allows to view the number of brokers, run test cases, create topics and volumes, etc.

The main API of GDL Management Station API are:

-
1. Cluster API (Operations pertaining to metadata in Kafka Cluster) to verify the number of brokers and access points of Kafka.
 2. Command API (Execute deployment test) to execute the deployment test cases of Kafka layers.
 3. Topic API (Operations pertaining to Kafka topics) to create, retrieve, update, and delete Kafka topics.

Volume and Log Files

Named volumes and log files are used by each Zookeeper and Kafka to persist the application log under `/var/lib/zookeeper`, and `/var/lib/kafka` respectively.

The server log for debugging purposes are under `/mnt/log/kafka`, `/mnt/log/zookeeper` and `/mnt/log/tenant-api` for the 3 components (Kafka, Zookeeper, and tenant-api) respectively.

The naming convention in the docker-compose file reflects the instance that uses the volume. For example, in `docker-compose-single.yml`, the volume names used in the compose files are

1. **zookeeper:**
2. **kafka:**
3. **tenant-api**