Data Platform Components Monitoring

Estimated time to read: 3 minutes

Overview

In this tutorial we describe ways to monitor **Data Platforms** components that handle data ingestion, processing, aggregation and replication.

System Monitoring

System Monitoring, in the **Administration** section, contains 6 probes that can be used to monitor **Data Platforms** components.

In the **Status** tab there are two probes:

1. **Online Replication Queue** tells us how many messages are waiting to be published to Kafka, with information regarding what needs to be done for replicating data and building the Canonical Data Model (CDM) events; these messages are produced whenever a transaction is closed in CM MES.
2. **ODS Replication Queue** provides the same information, only for messages created in ODS by the Initial Sync mechanism (which is used to run replication/CDM building for historical data).

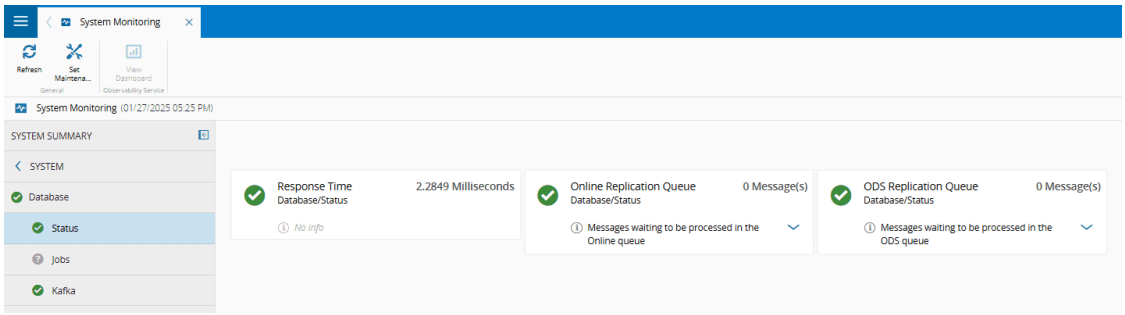

In the **Database -> Kafka** tab there are four more probes:

1. **Kafka Lag** tells us the lag building the replication documents.
2. **CDM Lag** tells us the lag building the CDM documents.
3. **IoT Lag** tells us the lag forwarding IoT events to the destination topics.
4. **Replication Lag** tells us the lag materializing replication documents in the destination database.

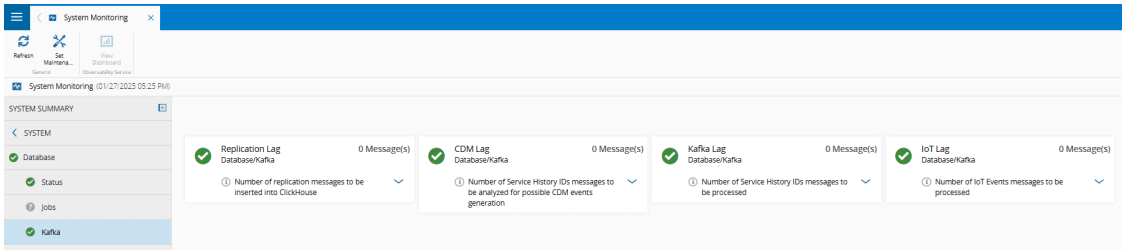

The location of the probes is represented in this simplified flow chart, which can be used to help diagnose potential issues with **Data Platform** components:

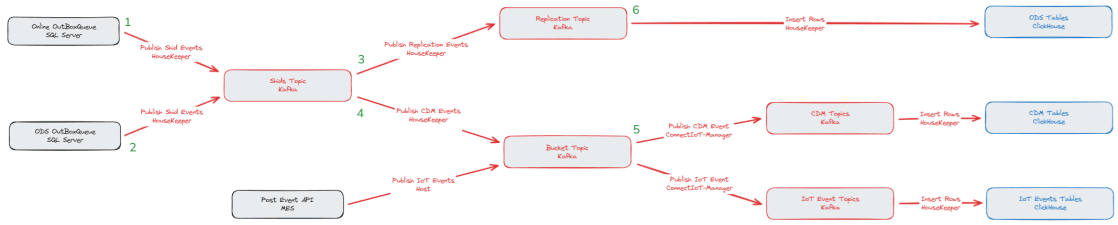

A quick way to look at it:

1. If **Online Replication Queue** shows a lot of messages waiting to be published, then it is possible that there is an issue with either **Kafka** or Data Platform component **HouseKeeper**.
2. If **ODS Replication Queue** shows a lot of messages waiting to be published, then the diagnosis is the same, i.e., a potential issue with either **Kafka** or **HouseKeeper**.
3. If **Kafka Lag** is too high, then there might be an issue with **HouseKeeper**, but it is more likely that either the MES **SQL Server** database is too slow, or more **HouseKeeper** instances are needed to keep up with the workload being generated by CM MES.
4. If **CDM Lag** is too high, then the diagnosis is similar, i.e., an issue with **HouseKeeper**, a slow **SQL Server**, or the need for more **HouseKeeper** instances.
5. If **IoT Lag** is too high, then there might be an issue with the component **ConnectIoT-Manager**.
6. If **Replication Lag** is too high, then the mostly likely diagnosis is an issue with **ClickHouse** or **HouseKeeper**.

By looking at the status of these probes it should be easy to determine if the CDM building and the ODS replication are working as expected and without issues.

Data Orchestrator

Another **Data Platform** process that can be monitored in MES is the aggregation of the Data Warehouse data sets. This can be done by navigating to the **Data Orchestrator**, in the **Administration** section:

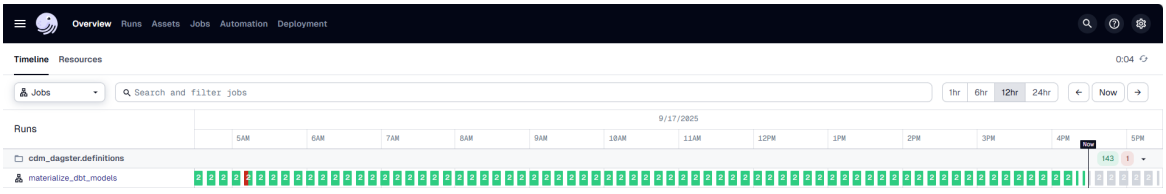

The **Data Orchestrator** is **Data Platforms** job runner. If you navigate to the **Runs** tab you can see the status of the job that aggregates the data in the Data Warehouse data sets, which goes by the name `materialize_dbt_models` and runs every 5 minutes:

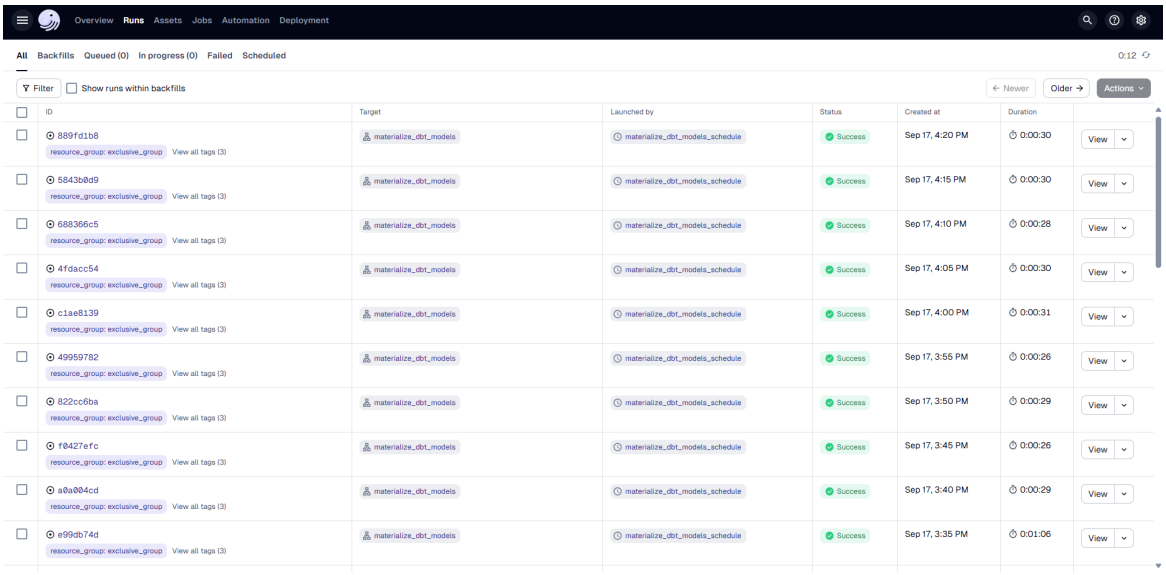

If you select **View** for a specific run, you can get details on the aggregations that were performed: