


Enhancement of the CMS Tabular Data Monitoring



AMERICAN UNIVERSITY OF BEIRUT

CERN Summer Student
Zeina El Kojok

Project Advisors
Nikodemus Tuckus, Federica Legger, Brij Kishor Jashal

Abstract

Monitoring of CMS distributed storage is paramount to ensure efficient use of disk and tape resources allocated to the experiment. The CMS Monitoring team currently uses **MongoDB** to store data for the **tabular data monitoring** and **DataTables** to visualize the data. However, **this pipeline can be improved and streamlined to enhance the performance.**

This project aims to implement **OpenSearch for storage** and **Grafana for access** to overcome the performance issues and to unify the tabular data monitoring pipeline with other dataset monitoring pipelines.

Current Tabular Data Monitoring Pipeline

The tabular dataset monitoring consists of tables summarizing for each dataset the content, size, creation and last access dates, and location on the CMS Rucio Storage Elements (RSEs).

Steps of the tabular dataset monitoring flow:

- Extraction and aggregation from the Rucio and DBS (Dataset Bookkeeping Service) OracleDB databases.
- Sqoop jobs create the daily database snapshots that are stored in HDFS.
- Snapshots in HDFS are processed by Spark jobs that aggregate metrics of interest for each specific use case.
- Data is injected in a stand-alone MongoDB instance that serves as a storage to a web service backed by Go and JQuery DataTables stack.

ISSUES WITH CURRENT PIPELINE

- Requires a lot of maintenance since the MongoDB is a stand alone instance
- Missing data can go unnoticed due to the absence of an alarm system
- DataTables show the monitoring data and not the relevant statistics

Current Tabular Data Monitoring Pipeline



Desired Tabular Data Monitoring Pipeline



Result 1: Database Change from MongoDB to OpenSearch

New spark jobs have been implemented such that the data is now aggregated and sent to three test OpenSearch indices directly. The spark jobs are scheduled to run once everyday to update the indices.

Index	Schema and Data Sent
test_main	<p>685,498 hits</p> <pre> { "timestamp": "2024-07-18T17:10:00", "source": "RucioReplicasDIDs", "count": 117200, "type": "RucioReplicasDIDs" } </pre>
test_detailed	<p>5,515,604 hits</p> <pre> { "timestamp": "2024-07-19T17:38:00", "source": "RucioReplicasDIDs", "count": 117200, "type": "RucioReplicasDIDs" } </pre>
test_in_tape_and_disk	<p>432,496 hits</p> <pre> { "timestamp": "2024-07-19T18:05:00", "source": "RucioReplicasDIDs", "count": 117200, "type": "RucioReplicasDIDs" } </pre>

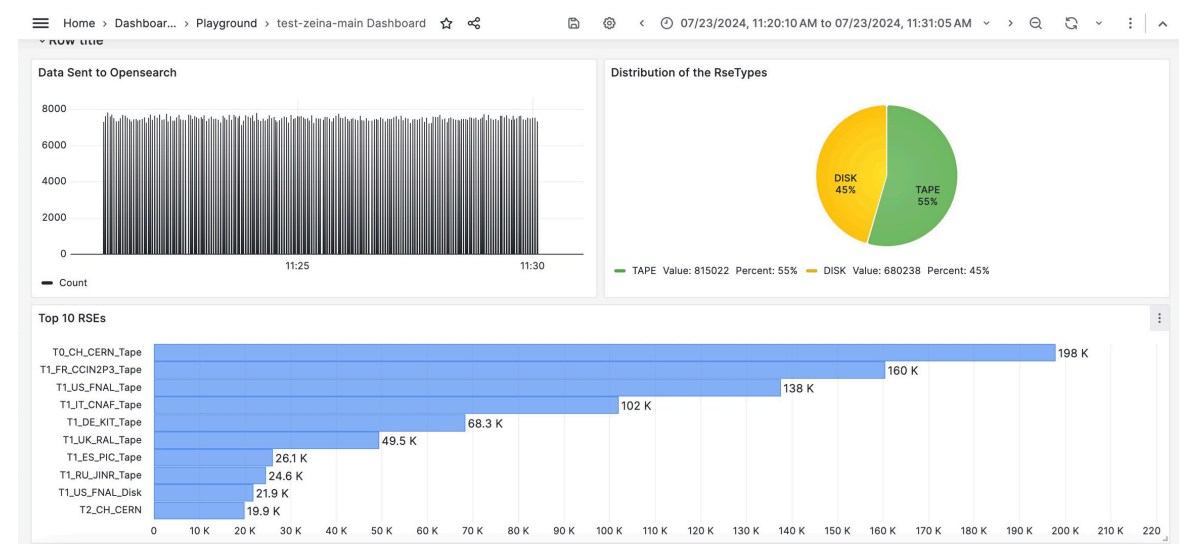
Result 2: Visualization Change from DataTables to Grafana

New Grafana dashboards have been created for each index. This change from DataTables to Grafana enhances data access and provides further insights about the data.

DataTable for the Main Table

Row Type	Dataset	Last Access	Max	Min	Avg	Sum	RepCnt	FileCnt	RSEs
100	cms-2024-07-18-17:10:00	2024-07-18 17:10:00	117200	0	117200	117200	1	1	RUCIO_CERN_Tape, TLUK_SAL_Tape

Grafana Dashboard



Advantages of Using OpenSearch and Grafana

- facilitates health checks
- reduces maintenance efforts
- allows the implementation of an alert system
- user-friendly and visually appealing
- provides a unified dashboard platform for all monitoring pipelines
- allows the creation of charts with very specific aggregations to visualize metrics commonly used by the monitoring team



Conclusion and Future Work

To improve the performance of tabular data monitoring pipeline, data storage has been moved from MongoDB to OpenSearch and the data visualization has been moved from DataTables to Grafana.

Future work includes moving the indices from the test to the production environment, adding the ActiveMQ step just before sending the data to OpenSearch and adding the relevant metrics to the dashboards.

References

- Jashal, B., Kuznetsov, V., Legger, F., & Uzunoglu, C., "The CMS monitoring applications for LHC Run 3", in 26th International Conference on Computing in High Energy 5 Nuclear Physics (CHEP2023), Norfolk, Virginia, USA, May 9 2023.
- Uzunoglu, C., Gomez, F., Jashal, B., et al., "Monitoring CMS experiment data and infrastructure for the next generation of LHC run", in ACAT 2022, Bari, Italy, Oct. 24, 2022.