

EFD Transport and Transformation

Kian-Tat Lim

The logo for the Large Synoptic Survey Telescope (LSST). The letters 'LSST' are rendered in a bold, black, sans-serif font. The letter 'S' is filled with a blue-to-white gradient, giving it a three-dimensional, glowing appearance. The letters are outlined in white.

Large Synoptic Survey Telescope

Requirements

LSE-61 DM System Requirements

DMS-REQ-0102: Provide Engineering & Facility Database Archive

- Engineering and Facility data from the Observatory Control System and associated metadata shall be permanently archived by the DMS and available for public access within [24] hours of their generation by the OCS.

DMS-REQ-0358: DM EFD Query Performance

- The DM copy of the EFD shall support at least [5] simultaneous queries, assuming each query lasts no more than [10 sec].

(Very) High-Level Design



LDM-148 DM System Design

[The Archiver service] also includes an EFD Transformation service that extracts all information (including telemetry, events, configurations, and commands) from the EFD and its large file annex, transforms it into a form more suitable for querying by image timestamp, and loads it into the permanently archived "Transformed EFD" database in the Data Backbone.



DMTN-050 EFD Handling within DM

2017-12-06 (original) to 2018-04-25 (schema considerations)

- Transformed EFD is contained within the Consolidated Database (and therefore is relational).
- Transformation extracts from the Base, inserts into NCSA, which is then replicated back to the Base.
- Expect latencies of < 5 min in normal operation.
- Initial schema is a copy of the EFD plus columns/tables to match telemetry to exposures.
- LFA ingested into the Data Backbone as files; "pointer" URLs point to DBB.

DMTN-082 On accessing EFD data in the Science Platform



2018-08-13 to 2018-09-11 (informed by discussions at PCW)

- Concerns about latency for Commissioning uses, potential difficulty updating ETL, and strict linkage of EFD and Transformed EFD schemas (although the latter is a bit of a red herring).
- Desire to reuse SQuaSH tooling and experience.
- Need for progressive refinement/aggregation of items (effectively replaces part of Calibration Database).
- Proposal to use Kafka to transport EFD data directly from SAL to so-called "DM-EFD".
 - Telemetry (potentially aggregated), events, and commands all stored.
 - LFA and operator logbook not stored in "DM-EFD" but instead ingested into DBB.
- TAP interface and custom EFD access Python library proposed.

SQR-029 DM-EFD prototype implementation



2019-01-17 to 2019-02-25

- Since John commanded all of you to read this, I do not need to summarize.
- Question about SAL topic evolution: answer is that all clients of that topic must redeploy synchronously. Currently, this is done by announcing the tag of `ts_xml` to be used and then rebuilding and redeploying all relevant CSCs. SAL messages are not known to contain version information.

Questions

Kafka + Confluent Schema Registry

- Can we select Kafka + Confluent Schema Registry for transport?
- Should SAL produce "plain text" or Avro messages?



Questions

InfluxDB + Chronograf + Kapacitor

- Can we choose InfluxData/TICK?
 - Alternatives have been stalled, but `mysql dump + copy + ingest` with addition of exposure join tables (as proposed in DMTN-050) could possibly be built relatively easily.
 - Is it cheaper to develop aggregations in Kapacitor (or Kafka) vs. other ETL options?
 - Is InfluxDB storage capacity sufficient?
 - Is InfluxDB+Chronograf better than RDBMS+Grafana (demonstrated by AndresV)?
 - Should we add a relational database consumer for a permanent archive or is InfluxDB backup sufficient?
 - If delayed data comes in beyond the InfluxDB retention policy, does it get aggregated?
 - How many InfluxDB instances will there be?