Rubin Observatory

DMLT vF2F
Status of Catalog Ingest

Fritz Mueller, Leanne Guy, Hsin-Fang Chiang, Colin Slater

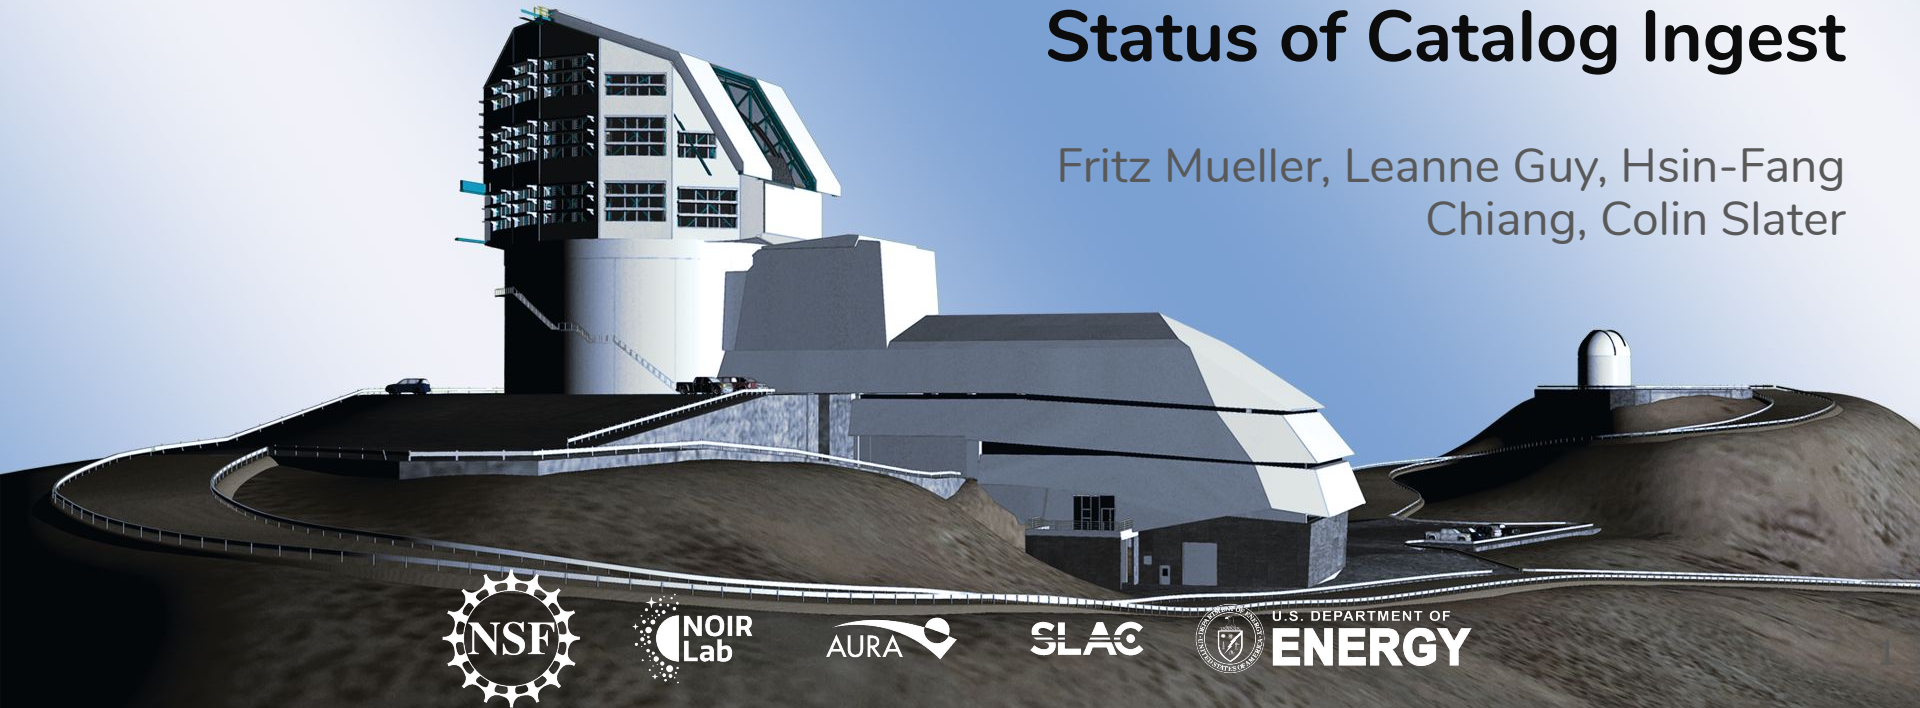

NSF NOIRLab AURA SLAC U.S. DEPARTMENT OF ENERGY

Outline

Rubin Observatory

- Recent progress automating Qserv ingest [DM-22806]
- Interaction between DPDD and SDM schema [DM-23818]
- Next steps and unanswered questions

Recent progress automating Qserv ingest

Rubin Observatory

- DM-22806: Automation of end-to-end system from science pipeline outputs to Qserv ingest
  - Triggered by RFC-243: Request for a regularly updated pipeline output dataset
  - Automation of the process to take science pipelines outputs, produce parquet files that are properly configured in the SDM format, and automatically ingest them into Qserv.
  - The initial implementation for automated ingest will be based around the reprocessed HSC data products, given that the HSC end-to-end ingest process has already been executed manually.

Recent progress automating Qserv ingest

Rubin Observatory

- **Since Feb, all HSC-RC2 Object tables have been ingested to Qserv:** w_2020_07 (Feb), w_2020_11 (Mar), w_2020_14 (Apr)
  - Often a set of data is ingested multiple times, to iteratively improve and integrate changes of the Qserv ingest system (new features or bug fixes).
  - Schema: use **cat** (aka **sdm_schema**) as the source of truth
- Just Object table so far
  - Source tables coming from pipelines soon?
  - yaml schema definition for the Source table has been made [DM-24534; just merge]
  - No problems anticipated ingesting the Source tables once they become available.
- Not in TAP_SCHEMA; but it can be done

Recent progress automating Qserv ingest

Rubin Observatory

- Original workflow plan was to use Airflow on Kubernetes, but…
  - Lots of Airflow- and k8s-specific learning to do
  - Additional delays due to HSC-PDR2
- Remediation: use Pegasus & verification cluster for first pass, accelerate prototyping
- Resulting workflow:
  - Good enough to continue ingesting HSC-RC2 tables, or to attempt larger HSC datasets on verification cluster
  - But, per Hsin-Fang: “IMO Pegasus is not the best choice here, and will want to move…” (known issues in a later slide)
  - Latest in use at https://github.com/lsst-dm/qserv-ingest-hsc-poc/

Recent progress automating Qserv ingest

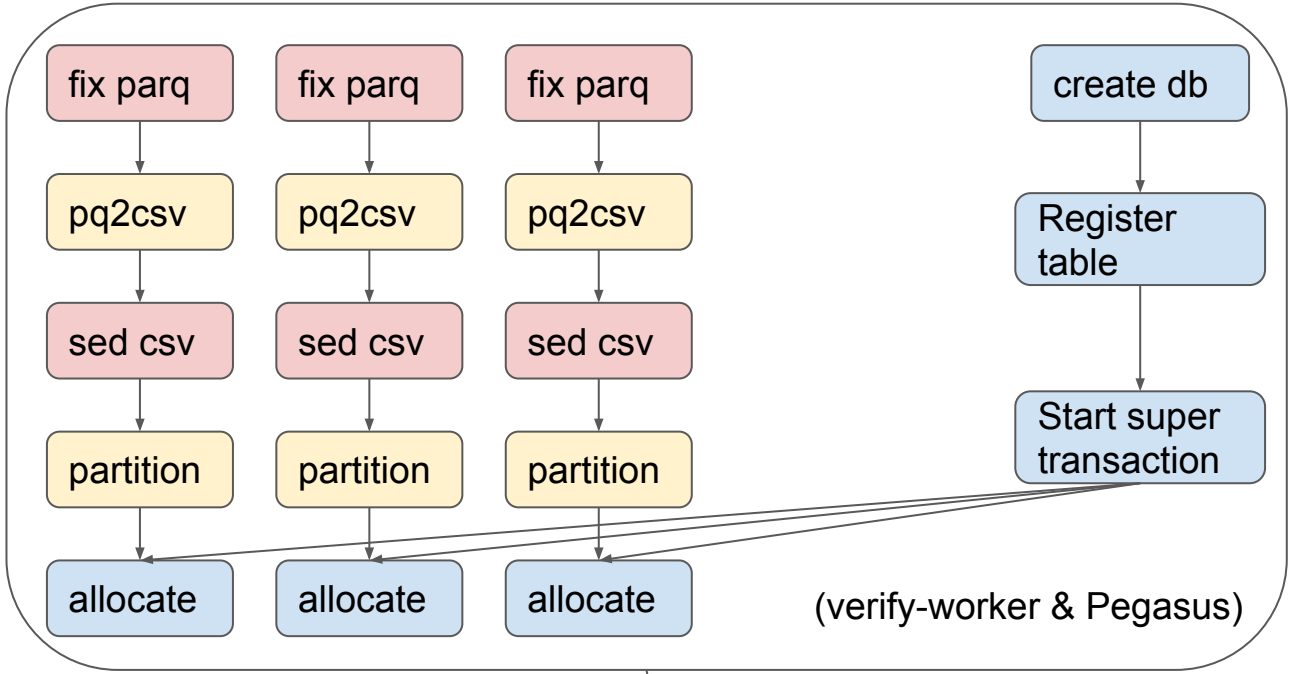

Only using one super-transaction
No restart

Inputs (not shown):
parquets from SciPi,
pre-made schema & configs

Red: will disappear
Blue: talks to qserv

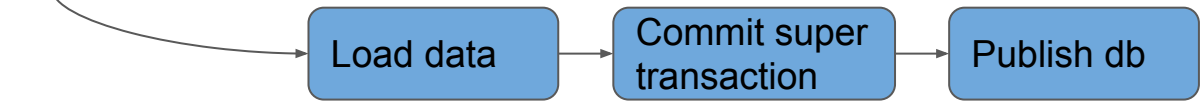

(parallel on qserv-master + workers)

Recent progress automating Qserv ingest

Rubin Observatory

- Not yet using Qserv ingest “batch mode”; avoidable overheads
- Not yet using multiple Qserv ingest super-transactions (complete or retry; no restart)
- Temporary hacks for nulls and casting data in the parquets; most hacks may be pushed into SciPi, pq2csv, partition, or ingest
- Tools don't use directly consume yaml schema definition yet (parquet_tools, partitioner, qserv table init)
- Potential issues integrating with docker and kubernetes use
- Split workflow
- For message exchange between jobs: file I/O vs. XCom
- (Ab)using shared file system. Very convenient now, but future system?
- Current implementation is only a poc made specifically for 1 RC2 table
- Want to try this in Airflow

Recent progress automating Qserv ingest

Rubin Observatory

- What's next:
  - Incremental improvement continues with known issues per previous slides
  - Continue ingesting monthly RC2
  - Include an ingest in the Qserv integration test suite
  - A small DC2 dataset was also ingested into -int; Fabrice & French working to ingest a larger DC2 dataset at CC-IN2P3
  - Forced Source table?
  - Plan workflow for IDF

Interaction between DPDD and SDM schema

Rubin Observatory

DM-23818: Produce a plan for interaction between the DPDD and the concrete SDM schema

- To date, schemas describing Rubin data products have been described in the DPDD. The DPDD (LSE-163) is a project-level change controlled document.
  - Simple changes to columns require project approval.
  - The DPDD is very large, and not very user friendly.
- "cat package" (soon to be renamed sdm_schemas) now defines SDM schema definitions for all Rubin catalog data products in YAML files" as per RFC-672
- The "cat package" serves as the **"source of truth".** Changes controlled by the DM CCB (not project CCB)

Interaction between DPDD and SDM schema

Rubin Observatory

- LDM-153: LSST Database Baseline Schema has been updated to be automatically generated from the yaml-defined schema in the CAT package.
    - Some work needed, e.g. provenance sections need updating
- Online schema browser was very useful but is currently inactive - needs some work to update it but ingest has taken priority recently.
    - Developer assigned; work to begin in June.
    - Also to be automatically generated from yaml definitions in the CAT package.
- Expect that most users will use the online browser to explore the detailed schema rather than read LDM-153 or the DPDD.

Interaction between DPDD and SDM schema

Rubin Observatory

- There is still a need for a user-facing scientific description of the data products as provided in the DPDD
- Given this, the evolution of the DPDD now needs to be addressed, then it's interaction with the SDM schema.
  - DPDD should not duplicate the information in the CAT package, LDM-153, nor the online browser.
  - If it describes SDM outputs it needs to be consistent with the CAT package.
  - How to best keep it up to date with evolving outputs of the science pipelines
  - etc

Next steps and unanswered questions

Rubin Observatory

- To address split-workflow issues, need to understand more about target system.
  - What systems will be used to support catalog ingest (DRP nodes? Dedicated? Spun up cloud per-purpose? Qserv nodes only? …?)
  - What data conduits will be available to move data between DRP, ingest drivers, and Qserv nodes (Shared file-system? Object store? Streaming only? …?)
- At what stride and cadence will parquet catalog data be received for ingest from DRP?
- QA/validation of delivered parquets (how/when/who?)
- QA/validation of partially-ingested (how/when/who?)
- What are the pathways and procedures for publishing corresponding catalog metadata in VO services in the LSP?
- What is on the books in terms of DRP->LSP “rehearsal” in order to convince ourselves there are no missing pieces?