



Rubin Observatory Multi-Site Testing

Richard Dubois (USDF, SLAC)

2 August 2023



SLAC




U. S. DEPARTMENT OF
ENERGY

Intro to Rubin/LSST

Legacy Survey of Space and Time (LSST), at the Vera Rubin Observatory:

- Science themes: dark energy, dark matter, but also solar system and transient objects
- 8.4m telescope at Cerro Pachon (Chile)
- 3.2 Gpix camera
- All visible sky ($\sim 18000 \text{ deg}^2$) in 6 bands
- 10 years survey starting from 2025

Hernandez et al.
CHEP 2023




Key Science Drivers

Ivezic et al. 2019


Dark Matter, Dark Energy

- Weak Lensing
- Baryon acoustic oscillations
- Supernovae, Strong Lenses




Milky Way Structure & Formation

- Structure and evolutionary history
- Spatial maps of stellar characteristics
- Reach well into the halo




Cataloging the Solar System

- Potentially Hazardous Asteroids
- Near Earth Objects
- Object inventory of the Solar System



Exploring the Transient sky

- Variable stars, Supernovae
- Fill in the variability phase-space
- Discovery of new classes of transients




Rubin is a facility providing processed data to 7 independent science collaborations

Data Management System Vision

Raw Data: 20TB/night



Sequential 30s images covering the entire visible sky every few days



Prompt Data Products

Alerts: up to 10 million per night

Results of Difference Image Analysis (DIA): transient and variable sources

Solar System Objects: ~ 6 million

Data Release Data Products

Final 10yr Data Release:

- Images: 5.5 million x 3.2 Gpx
- Catalog: 15PB, 37 billion objects




via nightly alert streams



via Prompt Products Database



via Data Releases



Community Brokers
Alert Filtering Service

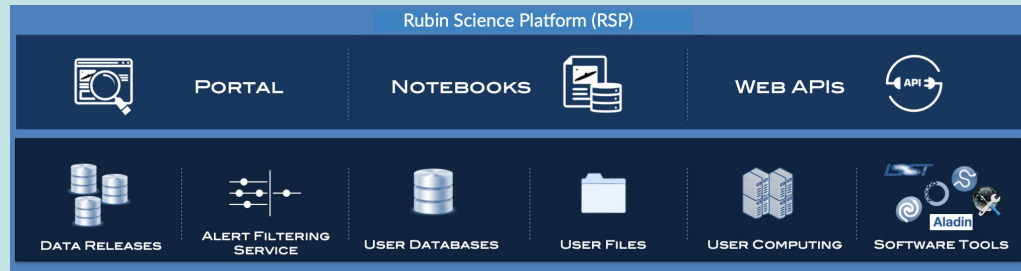


Rubin DACs (DFs & Chile)
Independent DACs (iDACs)

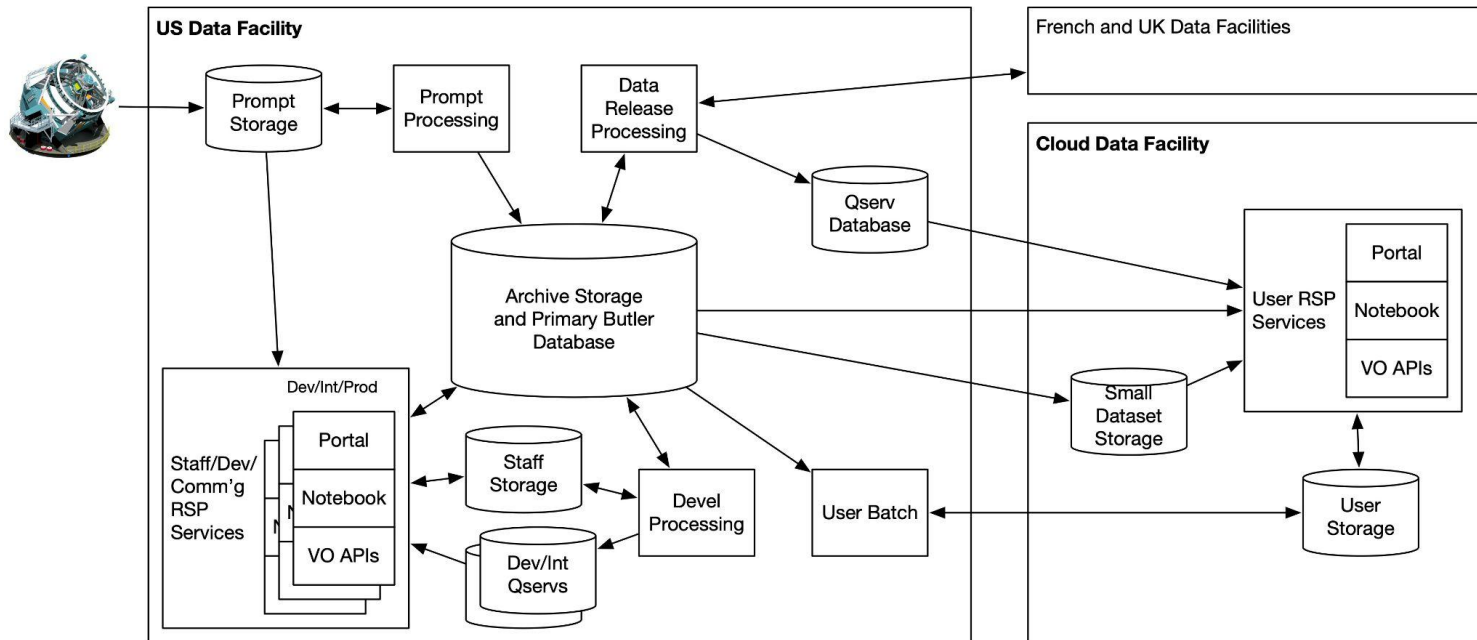
Rubin Science Platform

Provides access to Rubin Data Products and services for all science users and project staff

Access to proprietary data and the Science Platform require Rubin data rights



USDF: A Mix of On-prem and Cloud



Hybrid model: Data at SLAC but users on the Cloud.


Allows:

- Separation of security concerns
- Burst response
- Reduced risk

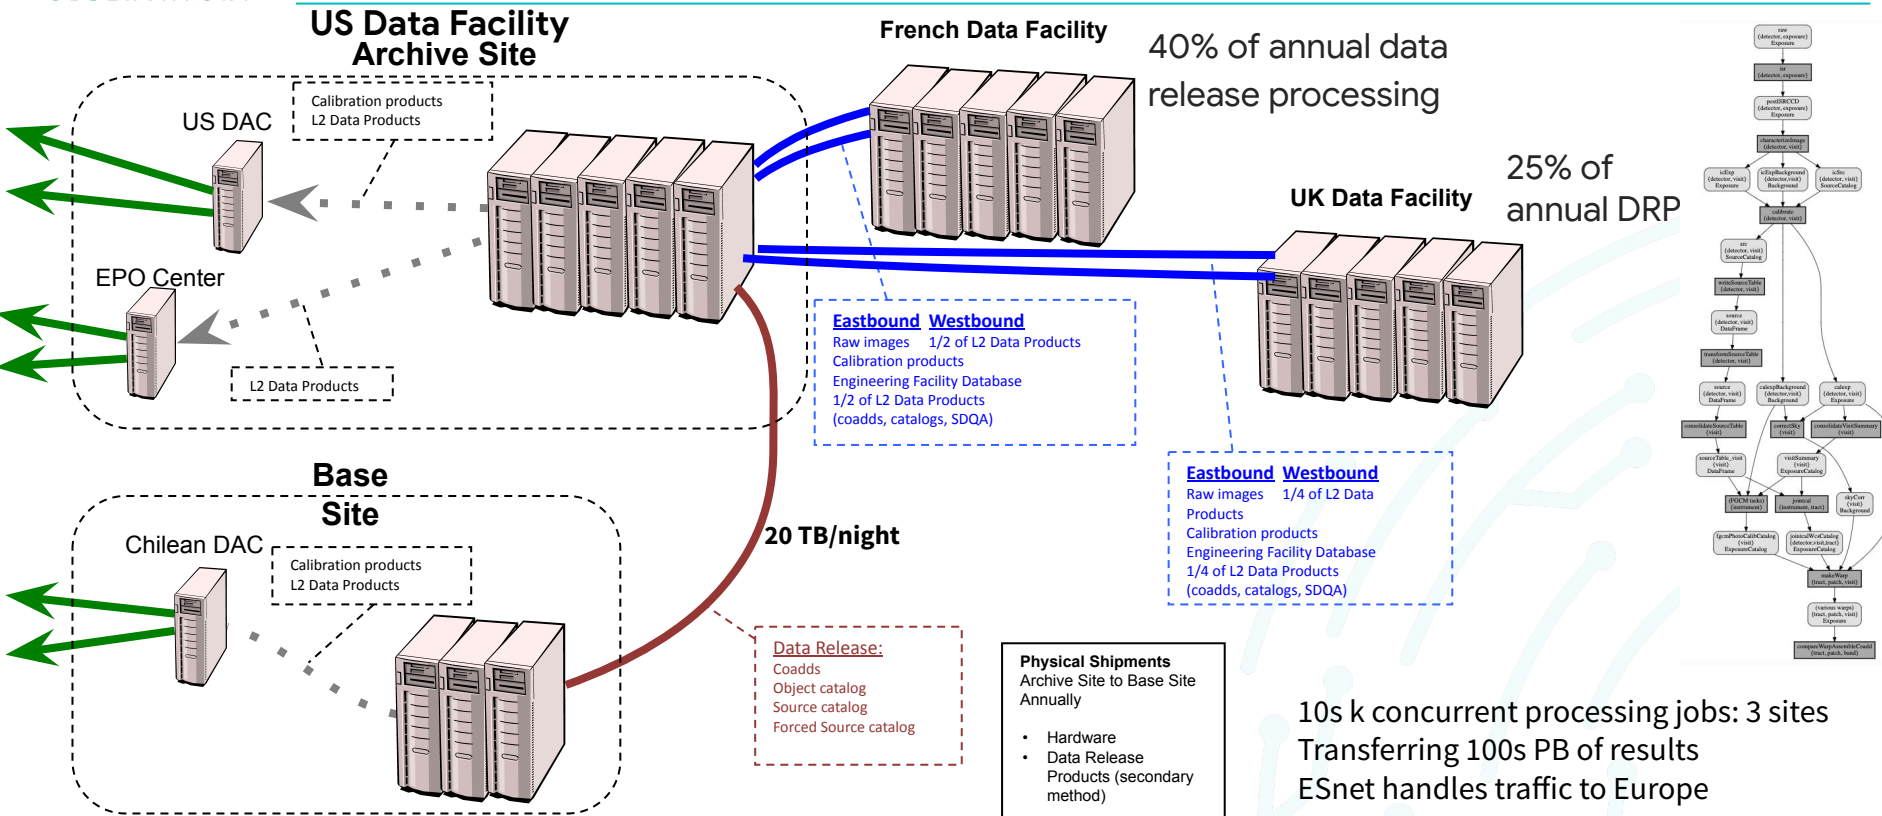
(see [DMTN-209](#))

RSP = Rubin Science Platform

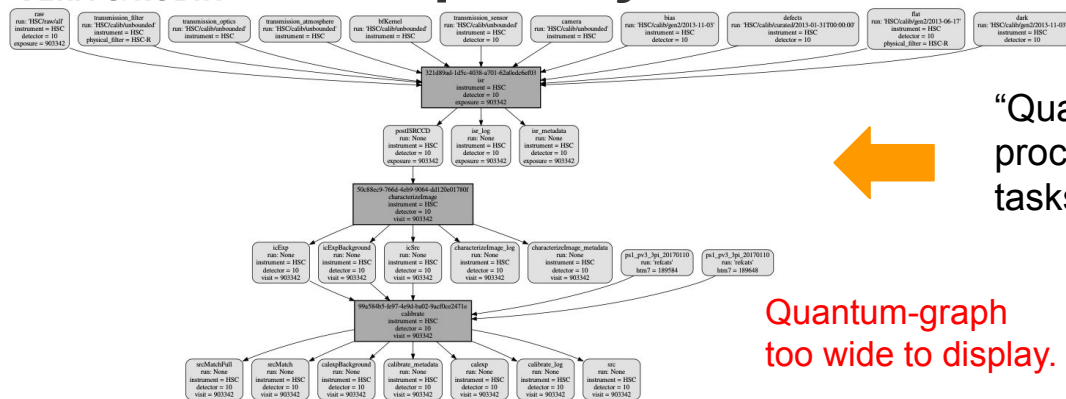
Pretty Big Data



Data Flows: Prompt & Data Release Processing



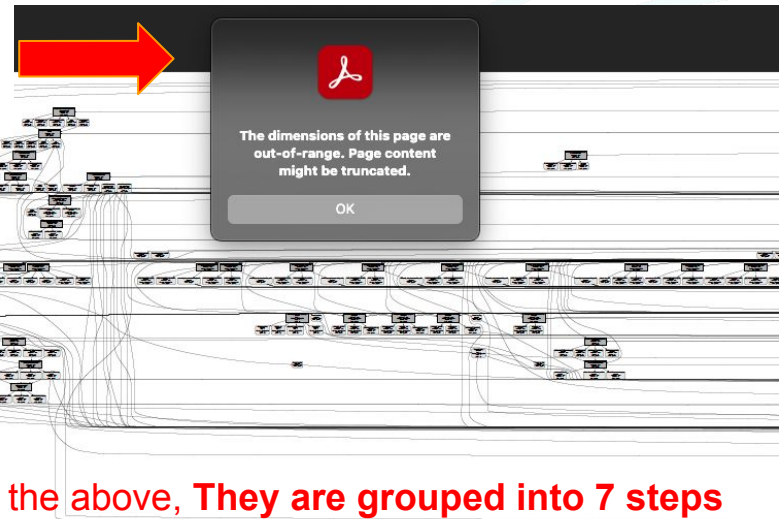
Complexity of Rubin Data Processing Pipeline



“Quantum Graph” of Rubin Science Pipeline to process a single LSST CCD image, showing 3 tasks (top to bottom) and input/outputs



Quantum-graph too wide to display.



Rubin Science Pipeline to perform single frame and coadd processing based on HSC engineering test data




Actual Rubin DRP pipeline will be a lot more complicated than the above, They are grouped into 7 steps

Projected Network Transfer Rates

Estimated Max Network Transfer Rates

SLAC outbound dominated by feeding IDACs and brokers



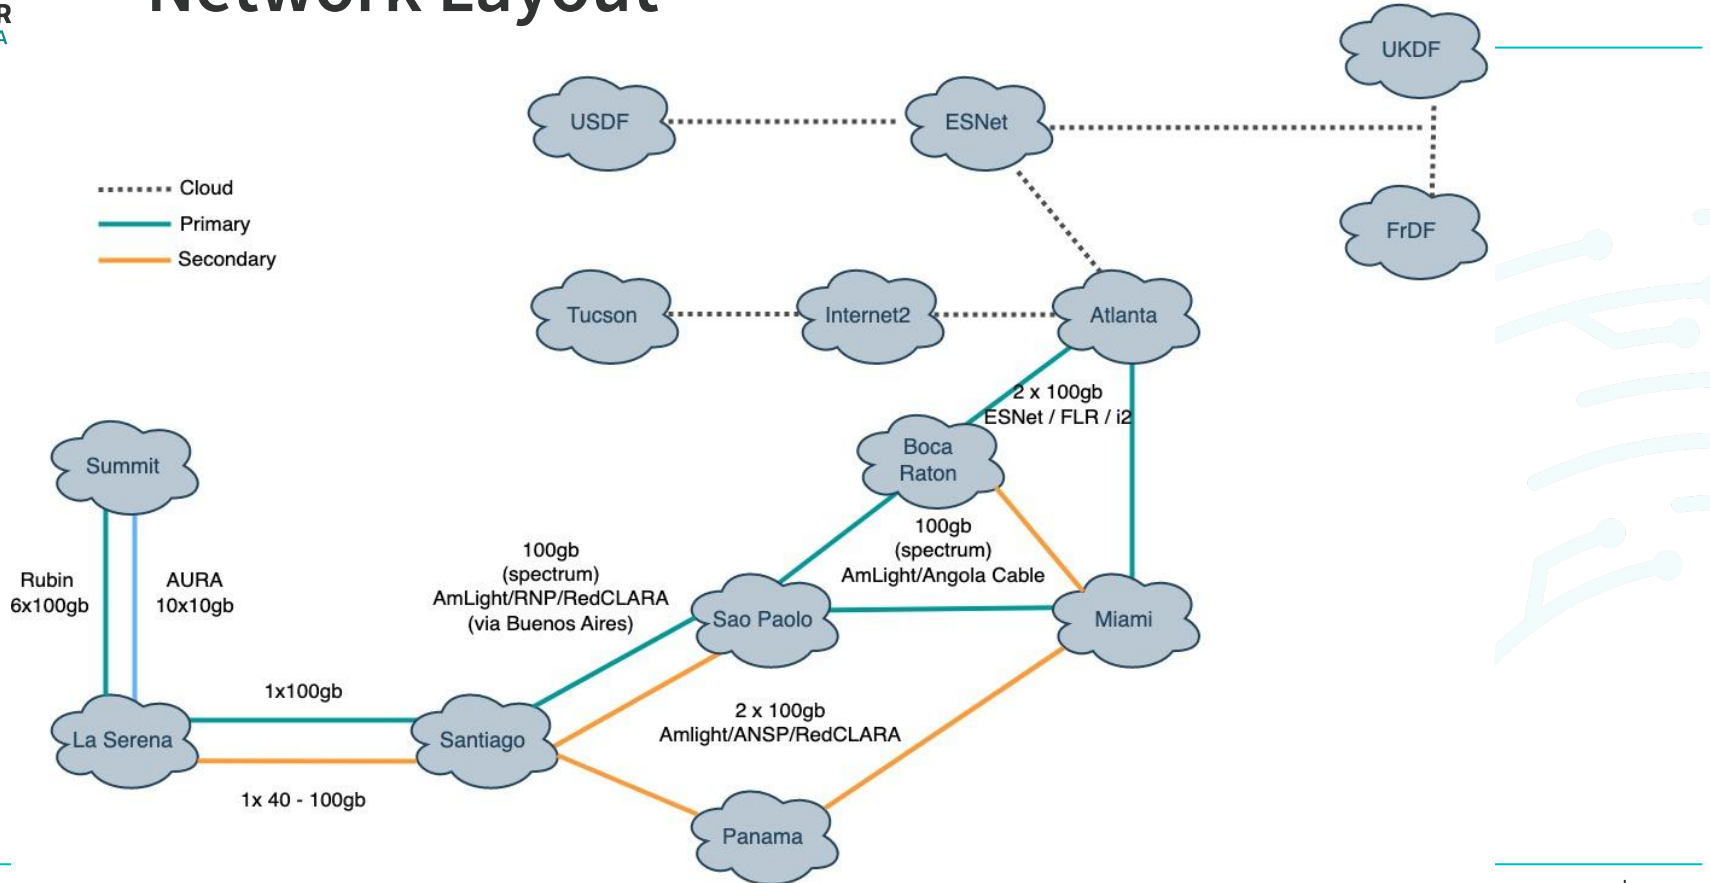
Assumes DRP transfers can proceed in parallel with processing

Technologies Adopted for Multi-Site

- [PanDA](#) - Workflow management
 - Used for DP0.2 in the Google Cloud Interim Data Facility
 - Exercised for routine CI reprocessing and now HSC-PDR2
 - Multi-site testing getting underway
- [Rucio](#) - Data movement
 - Data replication demonstrated to all Facility sites
 - Wrapping up interface to Butler
- [cvmfs](#) - code distribution
 - Stratum 0 hosted by CC-IN2P3 and in use for Rubin code in a variety of places
 - There are other options, but this appears to work

- **Rucio ecosystem:**
 - [Rucio](#): also developed by LHC ATLAS and used for 1+ decade
 - data classification, keeping track of data location, drive data movement
 - Rubin will have several times more file/object in Rucio than the current LHC ones
 - A big challenge for the backend database. Rubin will drive this forward
 - [FTS](#): also 1+ decade history
 - Think of it as a batch system dedicated to data transfer jobs.
 - Again, efficiently transfer large numbers of small files is a challenge
 - [Xrootd](#): has been around for 2+ decade
 - Mostly used as data transfer agent, to replace GridFTP
 - Rubin prefers object stores, and is driving xrootd based data transfer to/from OS/Cloud
- **Butler**
 - The original Rubin data management system
 - DB of metadata and pointers to data
 - software layer to access Rubin data
 - Must coordinate with Rucio

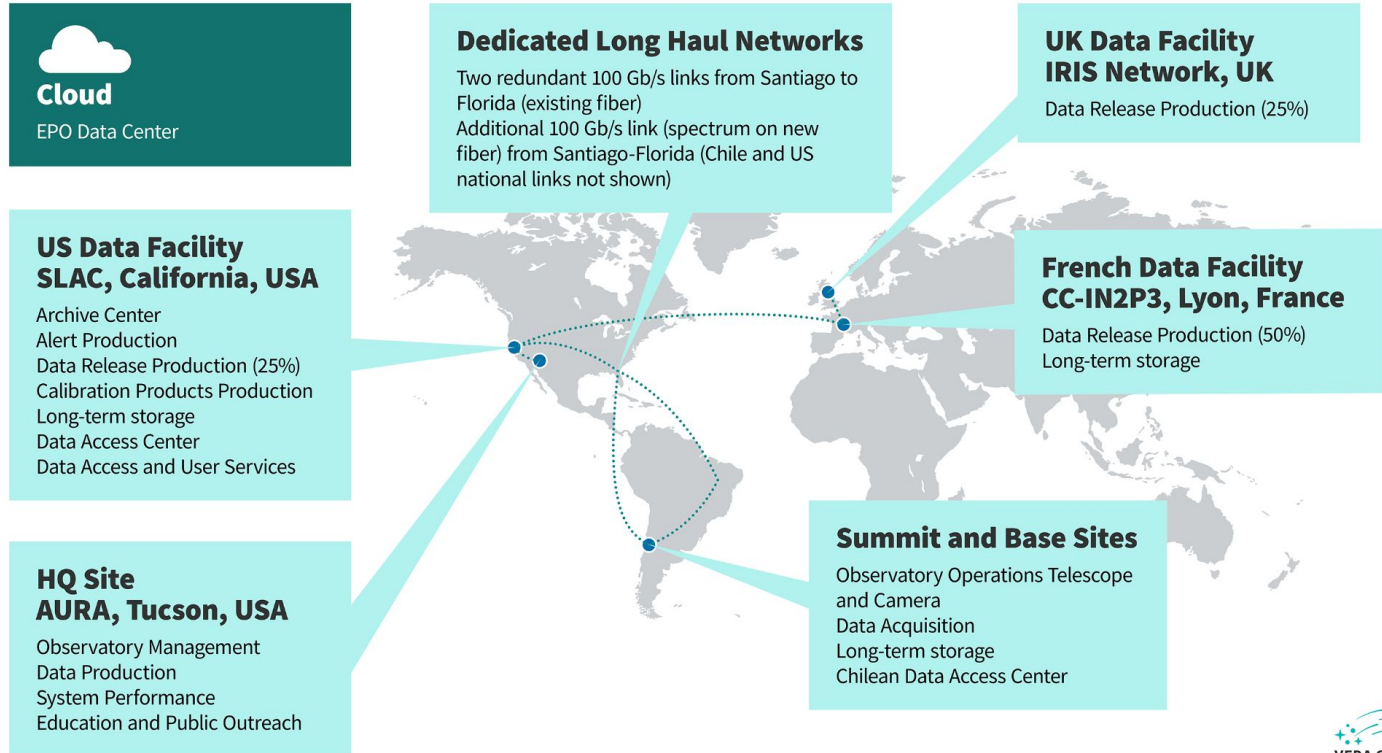
Network Layout



USDF Highlights

- Resources in place or arriving soon:
 - 10k batch cores; 30 PB disk. 50 kubernetes nodes. Various servers etc
- Ready for ComCam (actually AuxTel) Milestone completed (AuxTel is a 1 CCD camera used for atmosphere tracking)
 - Auto transfer and process AuxTel data from summit
 - Routine processing is happening; available for analysis in Rubin science platform
 - “Long Haul Network” between Summit and SLAC in use - combo of leased lines + ESNet
- Start of Hybrid model
 - Contract in place with Google - planning when the IDF transition actually needs to happen (IDF: already-running cloud instance spun up while the USDF came to be)
 - Production Qserv running with DC2 catalog; access from Google cloud demonstrated [DC2: DESC Data Challenge simulation]
- Supporting ~300 staff and commissioners
 - Successful bootcamp on USDF for commissioners
- Supporting Full Camera Testing in IR2
 - automated transfers of test data to France, using Rucio/FTS, are in place

Data Processing: An International Collaboration



- Access to 3k cores each at FrDF and UKDF
 - Demonstrated ability to submit and run jobs there to capacity (not yet at the same time)
- Rucio installed and configured:
 - Server at SLAC; Rucio Storage Elements at each site
 - Can routinely exchange data amongst sites
 - eg: transfer 7700 files, 3.5 TB - peak rate to CC-IN2P3 of 1.4 GB/s via FTS
- HSC PDR2 reprocessing at USDF
 - Not exactly multisite, but a thorough shakedown of PanDA at SLAC
 - Shakedown of “Campaign Management” tools
- Automated transfers of Full Camera data from SLAC to FrDF demonstrated
 - Transfers performed by FTS3 based on Rucio rules
- Finishing up connectors between Rucio and Butler before ramping up full multi-site capability:
 - Both Rucio and Butler act as repositories of dataset information - need to keep them in sync

Construction timeline

DP1 & DP2 - multi-site practice
DR1 in late 2025 - real thing

- **2022-09-30** : EPO Construction Finish (EPOC285) **Completed 2022-09-30**
- **2023-05-02** : TMA Handoff to Rubin (T&SC-0400-1730)
- **2023-09-22** : COMP: Camera Pre-Ship Review at SLAC (CAMM8090)
- **2024-04-23** : Camera Ready for Full System AI&T (COMC-060200-20930)
- **2024-04-26** : Dome Complete (T&SC-0400-0950)
- **2024-05-08** : 3-Mirror Optical System Ready for Testing (T&SC-1100-0400)
- **2024-07-16** : LSSTCam Ready for On Sky (First Photon) (SITCOM-122)
- **2024-10-24** : System First Light with LSSTCam (LSST-1520)
- **2025-02-19** : Test report: Final Pipelines Delivery (LDM-503-17a)
- **2025-02-19** : COMP: Science Validation Surveys Complete (COMC-0654-0430)
- **2025-02-26** : Operation Readiness Review Complete (SITCOM-130)

DMTN-232.lsst.io