



USDF Update

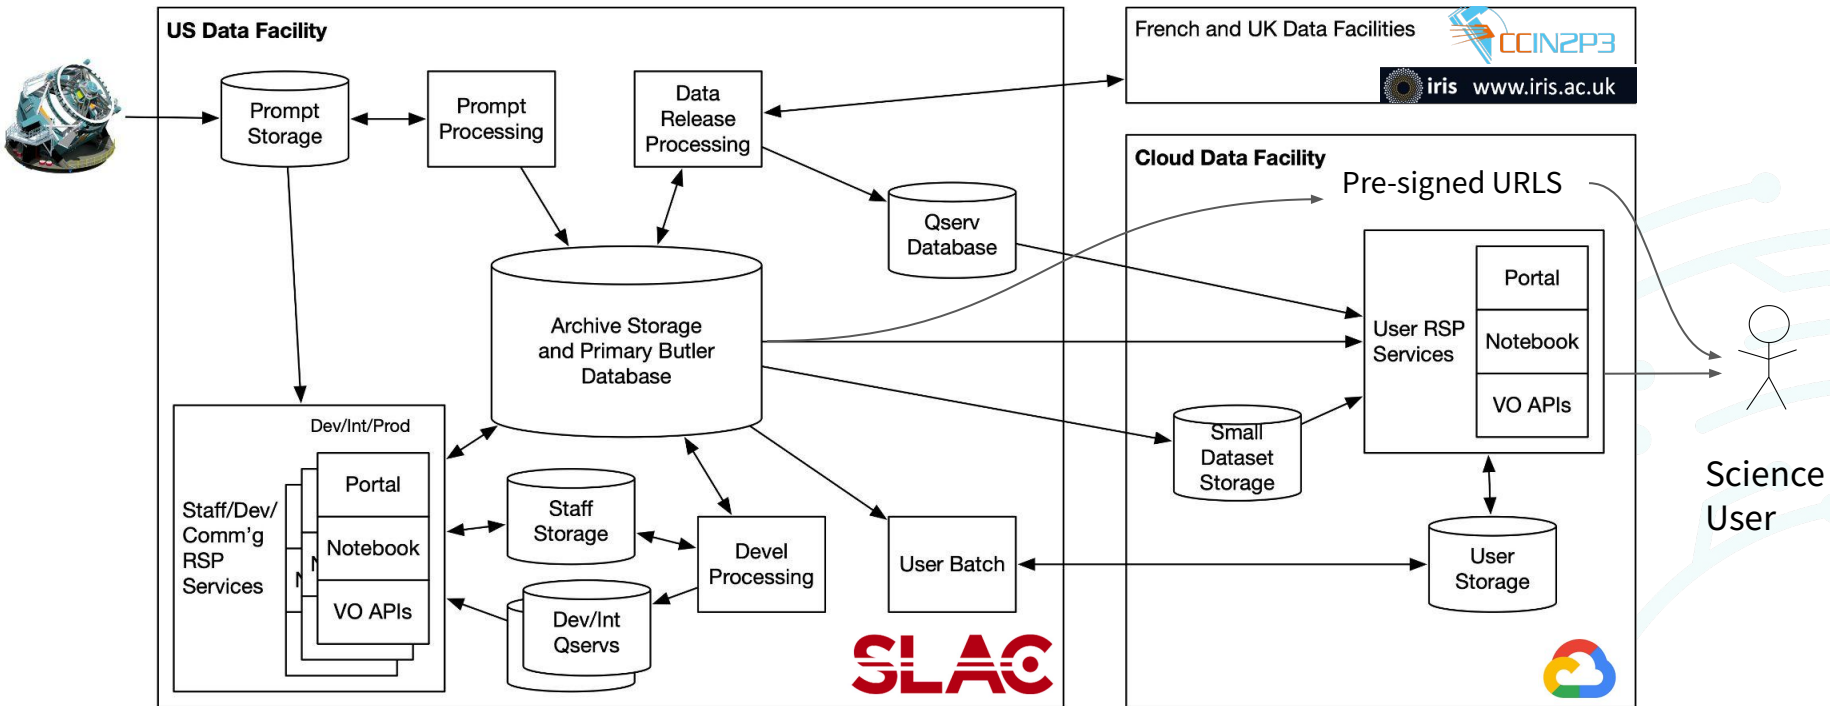
Richard Dubois
Jim Chiang
Steve Pietrowicz

24 Oct 2023



U.S. DEPARTMENT OF
ENERGY

USDF and Data Flow



RSP = Rubin Science Platform

Cloud-SLAC Division of Scope

Cloud

- Science users, with access provided by the Rubin Science Platform
- Provide personal storage/CPU + cloud access to coadds
 - Effective 500 cores added each year + 2 PB → 10 PB storage

SLAC

- Prompt and DRP processing
- Qserv catalogue server
- Storage archive for all data: cutouts, image access
- Serving alerts to the community
- Home for developers and staff (and commissioners)
- What about user batch? Work submitted from the cloud and run at SLAC... ?

USDF Status

- 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
 - Early throughput testing underway
- 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
 - DP0.3 worked
- Prompt Processing; Alerts distribution in place - scales to full rate; MPC mirroring in place.
- Supporting ~300 staff and commissioners
- Supporting Full Camera Testing in IR2
 - automated transfers of test data to France, using Rucio/FTS, are in place

Multisite & Reprocessing Status

- HSC PDR2 used to shake out USDF use of PanDA and s3df
- PanDA and Rucio servers in production at SLAC (supplanting CERN DOMA for PanDA)
- Gearing up for Multisite HSC campaign
 - Raw data (via Rucio) and calibs in place in all 3 DFs
 - Testing Rucio/butler connections now
 - Steps 1 and 2 of DRP will be manually submitted

USDF Issues 1

- Still not fully recovered from loss of EFD in May
 - Open source influxdb has made it impossible to restore old data to a live EFD version (also true at the base)
 - We are gearing up to get the Enterprise version which brings proper backups and HA operation. May need to retreat to local disk for the database for performance reasons. Might repurpose 2 Qserv nodes
 - Targeted for November.
- Later than expected access to SRCF-II
 - Making do trying to expedite embargo rack provision. Only flash server + Arista switches go there. Likely will split across 2 physical racks.
 - Encourages spending hardware budget early to allow for installation delays - we have \$2M of \$4.3M in hand for FY24 already (unusual).

USDF Issues 2 - USDF/S3DF reliability

- Not there yet...
 - Ideally must be (more) reliable for ComCam commissioning data taking
- USDF
 - Embarking on program to review and document k8s services, starting with PanDA and Rucio
 - Not all our upgrades work either...
 - Limping along with no real DBA - maybe being addressed with IT contract company
 - Bulking up monitoring to get early warning of problems
 - Mobu is running, but errors not being tracked yet
- S3DF
 - Lots of moving parts, and more errors than I'd like to see
 - I keep Amedeo apprised of the issues
 - Many barn doors *are* being closed - hopefully becoming one-time problems
 - ingress/egress may still be issues, especially when we start exercising cloud access to USDF services and multi-site processing hits servers at SLAC

Staffing Issues

- DBA
 - No luck at Fermilab for operations DBA effort
 - Investigating local-to-SLAC IT company (erginc.com) : first resume is in
 - Will need SOW and contract <\$250k (to avoid having to compete)
- Full stack dev-ops
 - K-T to turn into pumpkin after ORR...
 - Developing job ad now
- Replace Richard
 - May 31 deadline
 - Job ad in prep for quick release

Looking ahead

- Current list of USDF epics [here](#) for FY24Q1-Q2, roughly organized as:
 - Multi-site processing and scaling; WAN throughput testing
 - US DAC prep for transition from IDF and access to SLAC services
 - Embargo rack implementation; LHN throughput testing
 - Consolidated db, LFA exchanged with summit: LSE-72-required publication of events and metrics to the Summit
 - embargo_transfer (post-embargo data movement out of embargo rack)
 - Backup and disaster recovery strategies
 - Monitoring
 - Acquire hardware for next round
- Issues
 - Need client-server butler to test out DAC services; user batch?
 - What additional services need to be set up? (see Life of Photon)
 - Will Cassandra be up to the task?
 - Waiting on secret satellite catalogue
 - Still need to iterate on who looks after what in day-to-day operations

Past Activity

- Over the Last few Months
 - Rucio data transfers
 - Rucio Hermes Daemon deployed; Kafka connectivity between USDF/UKDF/FrDF
 - HTCondor at the Summit in progress
 - MTHederService created
 - MTOODS created
 - OODS container deployment handed off to T&S

Upcoming Activities

- Continued development
 - BPS & BPS plugin additions/changes/fixes (Michelle, Nick)
 - OODS updates for S3 (Steve)
 - Rucio automatic butler ingestion (ingestd, Rucio registration merge job code) (Steve)
 - MTHederService additions (Felipe)
- Support and updates
 - HeaderServices (Felipe)
 - Rucio transfers (Greg)
 - OODS for Auxtel, ComCam, LSSTCam (Steve)
 - BPS (Michelle, Nick)
 - HTCondor at SLAC and Summit (Greg, consulting with Michelle and Steve)
 - New NCSA employee talking to A2 about commissioning work