

# USDF Deployment Plan

## FY22

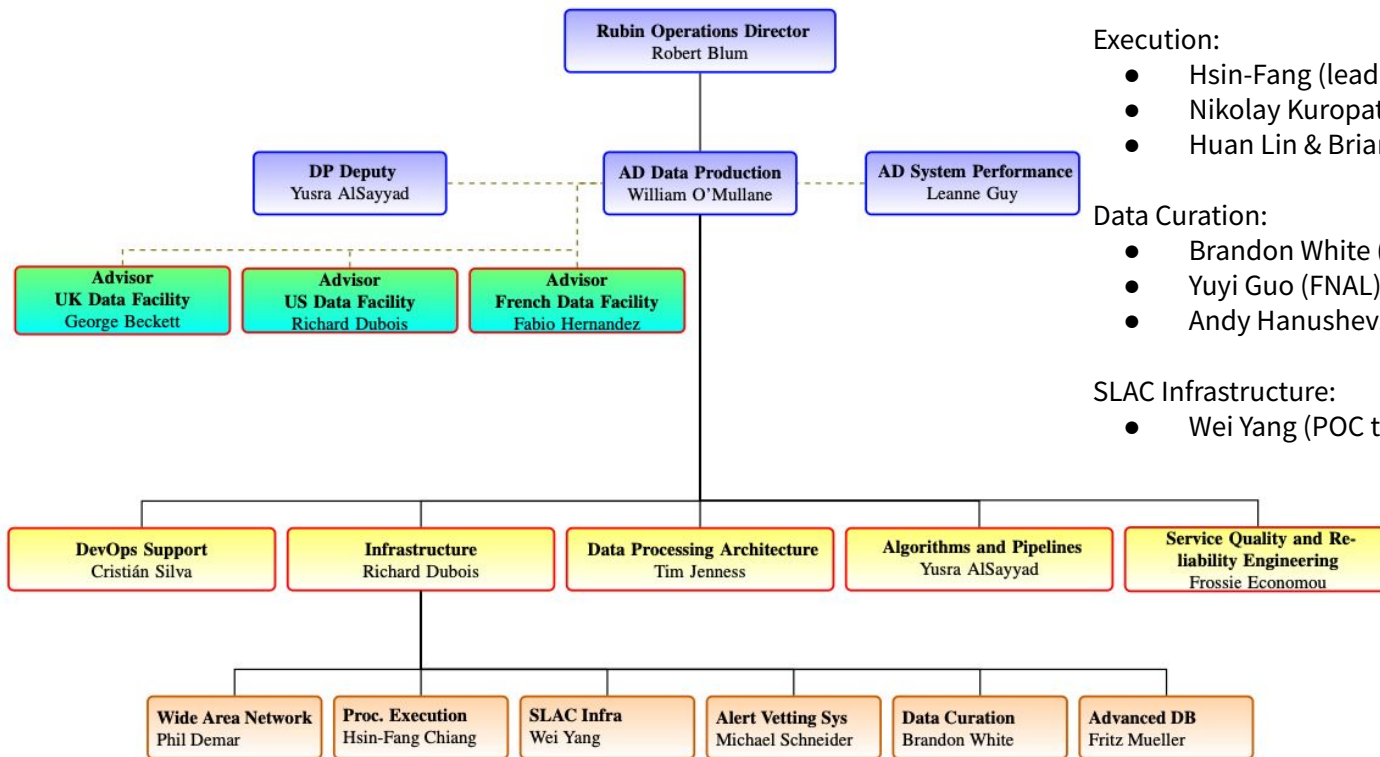


Richard Dubois  
16 Feb 2022



U.S. DEPARTMENT OF  
**ENERGY**

# USDF Who?



## Execution:

- Hsin-Fang (lead; upon return)
- Nikolay Kuropatkin, Jen Adelman-McCarthy (FNAL)
- Huan Lin & Brian Yanny filling in from V&V in DP0.2

## Data Curation:

- Brandon White (lead, FNAL)
- Yuyi Guo (FNAL)
- Andy Hanushevsky

## SLAC Infrastructure:

- Wei Yang (POC to SLAC SCS/TID)

## SCS Organization

SLAC

**Core & Networking:** Network design & implementation, 'ground zero' services (DNS, NTP, etc)

**Storage:** Posix and Object storage platforms (home, group, science data) tape archive and backup, migrations, policies, quota mgmt

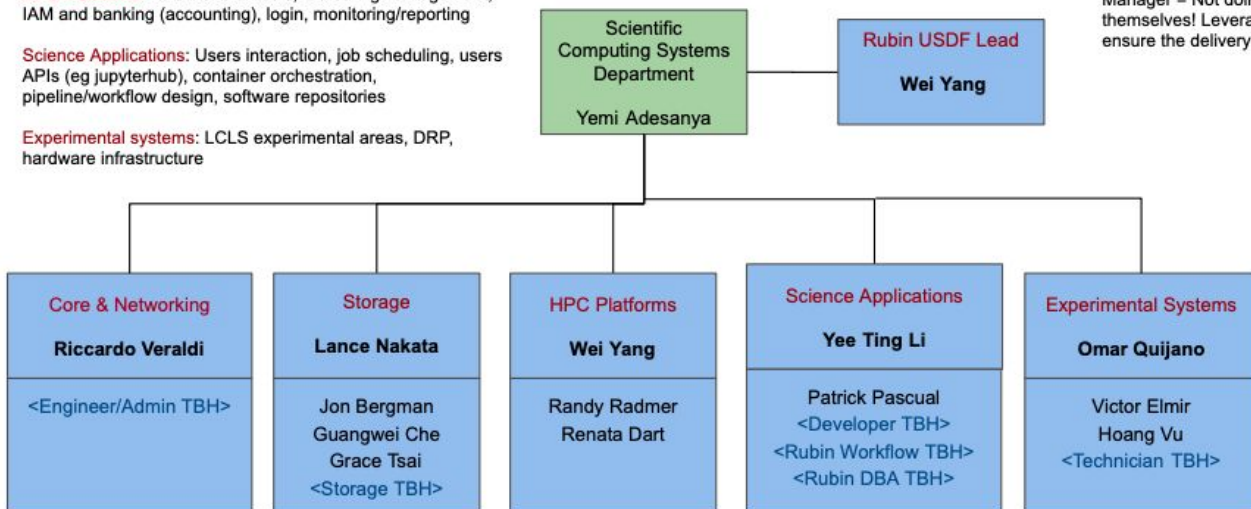
**HPC Platforms:** Data transfer tools, OS config management, IAM and banking (accounting), login, monitoring/reporting

**Science Applications:** Users interaction, job scheduling, users APIs (eg jupyterhub), container orchestration, pipeline/workflow design, software repositories

**Experimental systems:** LCLS experimental areas, DRP, hardware infrastructure

- Core team supporting all SLAC programs (biggest are LCLS and Rubin)
- Rubin has reqs out for storage engineer, software developer (to replace Brian Van Klaveren), and datacenter tech.
- FNAL to provide software dev/k8s effort starting FY23 (1.5 FTE)

**Rubin Lead:** Effectively operating as a Project Manager = Not doing all the implementation themselves! Leveraging the entire SCS team to ensure the delivery of USDF.



Plus from IT:

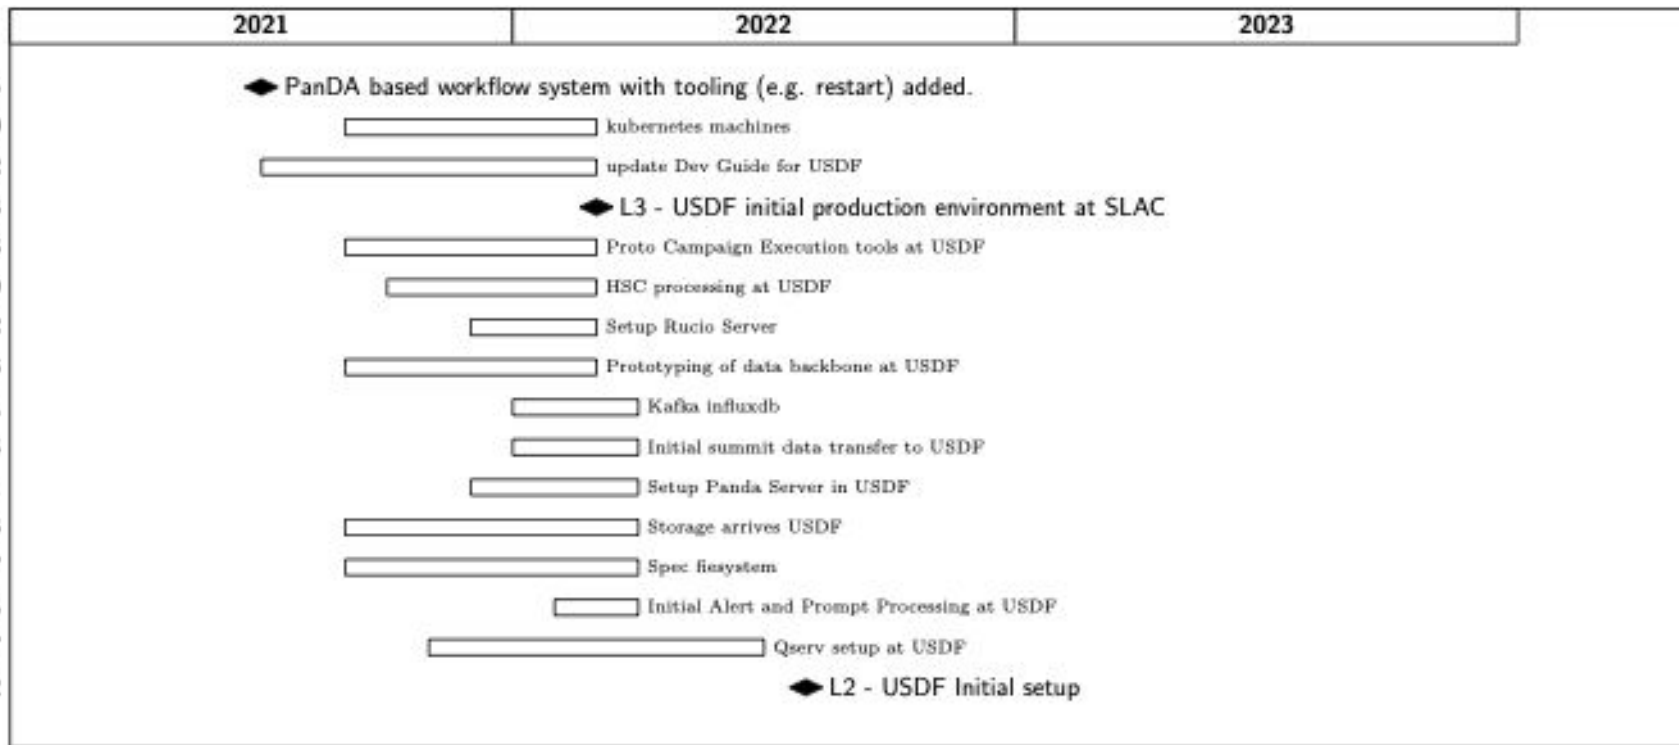
- Some sysadmin effort
- Site networking
- accounts/help desk

# Planning

---

- DMTN-189 for Data Facility scope
- RTN-021 for DF transition plan
- NCSA “shuts down” on October 1 (gulp)
  - SLAC becomes primary support for ComCam commissioning this summer
  - Driven by need to build up a secure enclave for arriving data, including ComCam
- Discussed weekly (Tues 11am PT), also #ops-usdf-arch
  - Planning focused meetings on infrastructure, k8s and security in next 3 weeks
- Two workshops so far, last one was [Jan 19-20](#); next planned for first week of April
- Upcoming: round of epic planning for the next year

# Planning: Epics & Dashboard



# Assumptions

---

- Science users in the cloud (IDF in this timeframe; hopefully forever!)
- Developers and Commissioners at SLAC
  - Fully vetted SLAC accounts (sorry!)
- ComCam data subject to security constraints
  - Develop a secure “OGA” rack to sequester data for 30 days. Must be encrypted.
- Exercise last summer to estimate (and order) hardware to support DP1 and some ComCam
  - 1k batch cores
  - 500 k8s cores
  - 4 PB storage (we have 1 PB on loan from SLAC now)
  - Ordered in *September*: k8s are here; the rest is due this month
  - Have 1 PB POSIX on loan
- Essentially all services to be deployed via k8s
- We need more of everything to really support ComCam
  - We’re working on what to order now, with \$1M in the bank and expecting \$800k more to be released this FY - uncertain due to year-long CR. This was meant to cover LSSTCam needs in FY23.


# Infrastructure Status

---

- We're building out into SRCF-I to start - 4 racks (1 or 2 for OGA)
  - We've asked for 3 more - waiting on the response
  - SDF [documentation](#)
- Construction has started on SRCF-II, expected to be usable summer '23. It's an extension of SRCF-I.
- Choices still outstanding:
  - Weka for POSIX - weka could also provide tiered access to object store
  - Object store choice: minIO vs xrootd



# SDF Load example - last week



Currently 15k cores in SDF



# Status of Activities

---

- Being tracked [here](#) for easy viewing and standups
- Loaner storage set up to mirror NCSA file tree, starting at /sdf/group/rubin/
- k8s nodes being configured
- 4 128-core Romes arrived; 4 more in transit. Installed within 10 days.
- Demonstrated RSP install; needs updating to connect to SLAC LDAP
- Working on Rucio movement of data from summit to SLAC
  - K-T has a different scheme in mind for the low-latency visit transfers
- PanDA, Rucio, cvmfs clients set up
- Starting on routine HSC reprocessing; getting the input files to SLAC
- Minimal start on updating the dev guide for SLAC

# Next round of purchases

---

- Given delivery delays, we're in a hurry to get orders out
- From the \$1M already:
  - \$125k switches etc for Rubin racks (short delivery time)
  - \$75k tape (ditto)
- k8s nodes run \$8k each; batch nodes \$100/core (128-core Romes)
- Storage is \$60/TB (uncertain what support charges will be)
- We need to balance:
  - Additional k8s nodes for RSP (12 of 16 on hand targeted at services)
  - POSIX storage for developers, commissioners vs object for data
  - Batch cycles
- We're guessing 5 PB POSIX
  - eg 12 k8s nodes (\$100k); 3k cores (\$300k); ~6 PB spinning disk might fit, depending on POSIX storage model
  - Then \$800k left (we hope!) for additional storage and cores for LSSTCam
- We need to start planning the migration from NCSA to USDF
  - Moving 5 PB won't happen overnight!

# Path to Routine Reprocessing at SLAC

- Routine reprocessing requires a full setup:
  - HSC data at SLAC and butlerized
  - PanDA server at SLAC
  - Rucio server at SLAC
  - Production hardware
  - Campaign Management tooling
  - Execution group (presumably) minding the reprocessing
- Exploratory steps (Feb) - essentially reproduce NCSA functionality
  - Ops epic: <https://jira.lsstcorp.org/browse/PREOPS-669>
  - Data should be at SLAC soon and butlerized (Brandon)
  - Fritz to try Parsl first with borrowed batch cycles
  - Try PanDA (should look the same to Fritz, given bps plugins)
  - Identify effort to drive the reprocessing
- Add Bells and Whistles
  - Production servers & storage
  - Not clear we need to involve Rucio?



# Planned Technology Choices: PanDA, Rucio, cvmfs

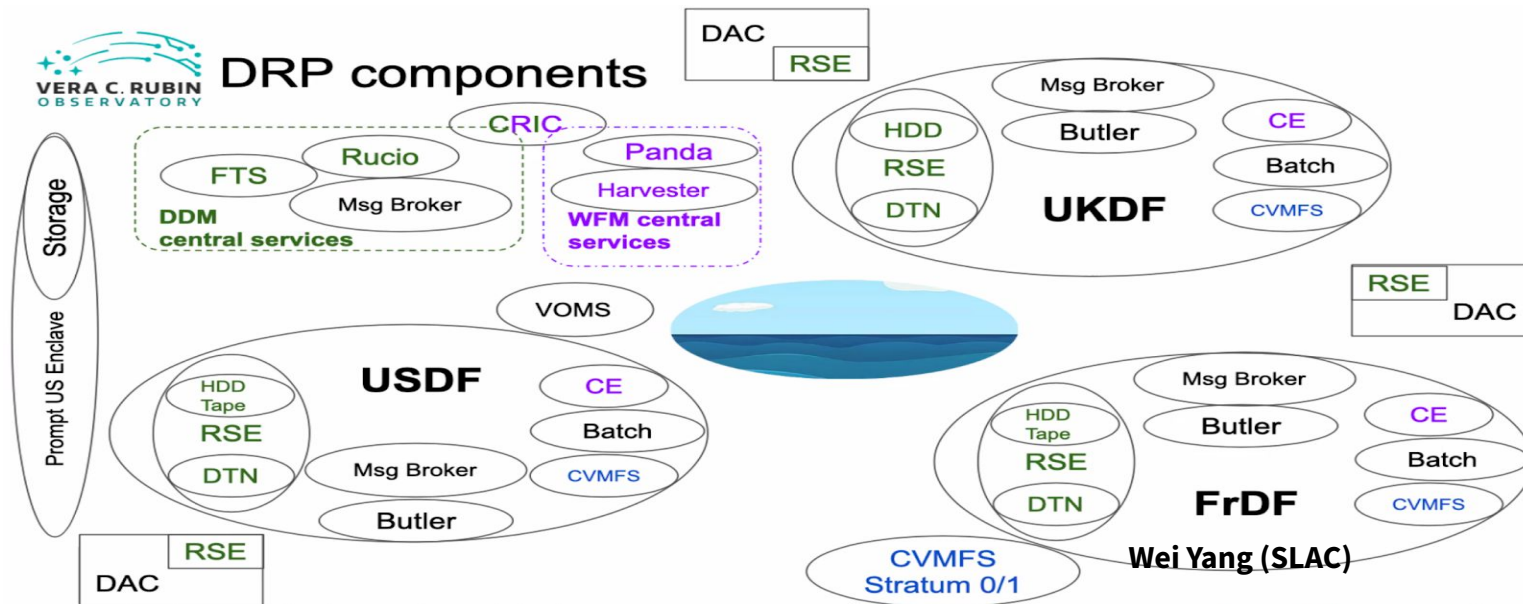
Richard Dubois  
16 Feb 2022



U.S. DEPARTMENT OF  
**ENERGY**

# Technologies Adopted by Ops

- Workflow: PanDA - rtn-013
  - Data movement: Rucio/FTS
  - Code distributions: cvmfs
- } Facing the issue of bolting them onto butler



# Technologies Adopted by Ops

---

Quick status:

- [PanDA](#)
  - In use for DP0.2
  - Multi-site testing just getting underway
  - Conditional commitment (to/from BNL) through DR1
- [Rucio](#)
  - Multi-site testing underway
  - Envisaged as the basis of the Data Backbone
  - No alternatives in mind
- [cvmfs](#)
  - 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

- History:
  - Wil liked what he heard from BNL re PanDA during the USDF search process
  - Proposed a demo by BNL to show that PanDA could work for Rubin. DP0.2 was/is the real test. RTN-013 documented the choice process.
  - We agreed to keep going with PanDA and BNL was told “PanDA through DR1, at least”
- The hype:
  - ATLAS@LHC runs millions of jobs per day at hundreds of sites and manages their production with 2 FTEs
  - BNL employs the core PanDA developers
- The reality:
  - ATLAS workflows are simpler and involve longer jobs; DAG handling was relatively new
  - Web monitoring is pretty opaque (but obviously can be learned) - deemed not good enough
  - It lumbered at the beginning (startup times not as important if long jobs)
  - BNL has 2 staff working with Tim/middleware, plus the attention of Torre Wenaus, father of PanDA
  - All the monitored information is available via a REST interface for custom web pages, but may not scale.



# Current Usage

---

- DP0.2 submits jobs to a (smallish) PanDA server at CERN, loaned for this test
  - Submissions come from RSP sessions
  - DAGs generated by bps are handled by an iDDS server (soon to be) hosted at Google (for better performance)
  - PanDA is unaware of datasets - the payloads know their inputs/outputs and assume the inputs are there in time (In ATLAS usage, PanDA is aware of datasets)
  - Connection to the middleware is via the bps plugin
- Multi-site
  - Starting with “Hello World” test among USDF, FrDF and UKDF (Lancaster, likely)
    - Wei Yang is coordinating sites configuration now
  - Will move on to running a bps submit first without Rucio, then with

# DP0.2 Experience to Date: Good, Bad, Ugly

---

- Loaned iDDS server was major performance bottleneck when large DAG was used. This triggered rewrite of job handling code in iDDS to make it more efficient and we are currently moving iDDS server to IDF so we can have a bigger machine.
- PanDA User interface still a bit opaque. PanDA team have requested feedback for suggested improvements.
- Real-time logging capture working great (payload logs captured by Pilot and then forwarded to Google Cloud Logging).
- Real time job monitoring stymied by use of Docker. Tests ongoing to use prmon in payload container.
- Execution butler got us through the initial problem of swamping CloudSQL registry but now investigating graph-backed registry.

**TimJ**

- History

- Predates me - it was the baseline when the USDF started up last year
- [DMTN-198](#) describes the data backbone plans

- The hype

- In use by all LHC experiments, DUNE - expected to handle Run 4 data volumes

- The Reality

- Rucio schedules data movement (based on rules) and relies on 3rd party transfer mechanisms. CERN invented [FTS](#) to replace Globus, which is no longer free.
- It is not meant to be triggerable, so don't expect it to leap into action
  - Need different solution for low-latency transfer from summit to USDF
  - K-T looking at object store POSTS
- Rubin may have 10x the number of files per year than LHC experiments currently
- Rucio has come along independently of butler: we've got something of a shotgun wedding going in terms of which knows/does what
- Counting on FNAL for primary Rucio support (UK also employs a core developer)

# Status & Issues

---

- Discussed at bi-weekly meetings on Mondays at 8 am PT and in #dm-rucio-testing
- Multisite testing has been underway for a couple of months among 3 DFs & NCSA.
  - Chile still being brought online.
  - Eg USDF & FrDF currently ping-ponging  $\frac{3}{4}$  million files
  - Will use it to copy HSC data from NCSA to SLAC for routine reprocessing
- Still lots to be settled:
  - Which is the source of dataset location truth? Rucio or butler?
  - The specifics of Rucio/Butler ingest are still under consideration
    - How do we reliably communicate between them?
      - Steve Pietrowicz is working on a message broker design (leaning to ActiveMQ) with K-T
    - Effective handling of metadata ingest. Will Rucio be able to handle this role?
  - Rucio scopes need careful consideration due to their role as effective Rucio namespaces
    - Database should be tested at extreme scale, using DB data volumes far greater than ATLAS

- Read-only POSIX filesystem in user space for code(/data) distribution
- [Described](#) by Fabio at the workshop
  - CC-IN2P3 hosts a stratum-0 at /cvmfs/sw.lsst.eu
  - both stable and weekly releases: typical retention period for weekly releases is 3 months, but selected releases can be pinned and kept online longer
  - distributed in two forms: set of files composing the release and (more recently) Singularity container images
  - self contained: conda, EUPS, Python interpreter, C++ compilers, ...
  - build via newinstall.sh from sources: unit tests systematically executed on both CentOS 7 and macOS
  - Singularity images are built from Docker images delivered by SquaRE via Docker hub
- Issue: is this the right mechanism for developer needs??