

# An Introduction to S3DF

Jay Srinivasan  
Division Director, SCS



## S3DF:

- What is S3DF? What is it NOT?
- Benefits of S3DF
- SRCF and S3DF architecture
- S3DF components: Compute, Storage and Services
- Support model:
  - Hardware contributions, installation and maintenance
  - Operations

# What is the SLAC Shared Science Data Facility?



- **Common shared scientific computing infrastructure** to tackle the challenges of massive throughput data analytics at SLAC
- Will support several key mission areas:
  - **Critical, data-heavy, scientific computing workflows** in SLAC users facilities (LCLS, Rubin, UED, CryoEM, SSRL)
  - Baseline scientific computing resources for ALL SLAC users
- This infrastructure will offer:
  - “Ever-green” computing and storage platforms (disk and archival storage)
  - Broad set of capabilities and services to support scientific computing
    - Interactive access
    - Batch computing
    - Service Compute
      - User-defined Kubernetes-based services for databases, web-based portals, data management, etc.
      - Jupyter notebook portals

# What the SLAC Shared Science Data Facility is NOT



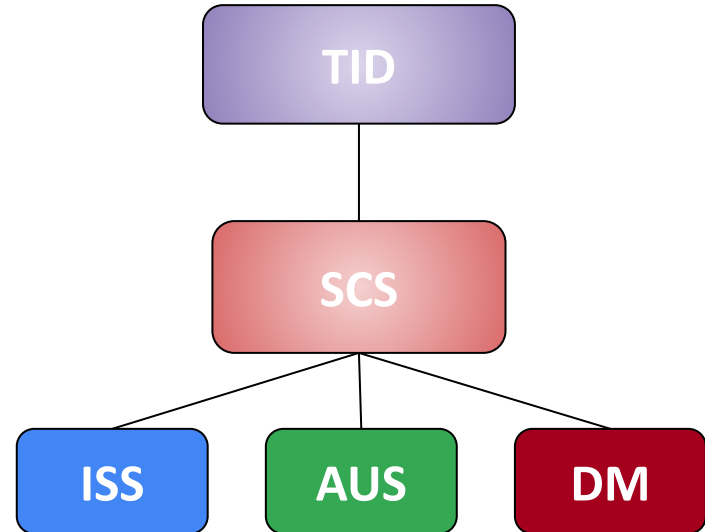
- Unix Desktop support
  - This is a SLAC IT function
  - We will provide interactive access to our systems, but the systems are optimized for scientific computing, and not desktop use
    - No printing, web-browsing, etc.
- Not a drop-in replacement for AFS
  - AFS has been used as a shared filesystem to support science groups for a long time. The S3DF storage is designed as a parallel filesystem to support computing on S3DF.
  - We will make some specific storage areas available via gateways, but this is not a “catch-all” storage area for *all* science output.
- Not a “personal” website.
  - While web-based services are offered through S3DF, the aim is to support delivery and distribution of science results from S3DF or portals to S3DF services.
- Not a place to run general (centralized) databases
  - The S3DF will support some databases that are specifically tied to scientific workflows that use the computing resources of our systems

# S3DF: Benefits

- Expect S3DF to be:
  - **Efficient** in terms of scalability, utilization and support
  - **Sustainable** deployments of compute, network and storage resources
  - Provide **baseline services** to provide meaningful resources for all users
  - Allow for **long-term planning**
- Emphasise data-centric scientific computing to differentiate this system from earlier SLAC Unix support




# S3DF: Managed by SCS

- SCS: Newly formed Division within TID
  - Composed of parts of old CDS Division with staff from ID as well
- Three Departments
  - Infrastructure & Systems Support (ISS)
    - Deployment and Management of S3DF; Storage; Networking; Experimental Systems
  - Application & User Services (AUS)
    - User engagement; User environment on S3DF; Application support for facilities; Batch system; Kubernetes services; Coact Platform
  - Data Management (DM)
    - Data workflow support for facilities; data management services



# SRCF-II Building Plans

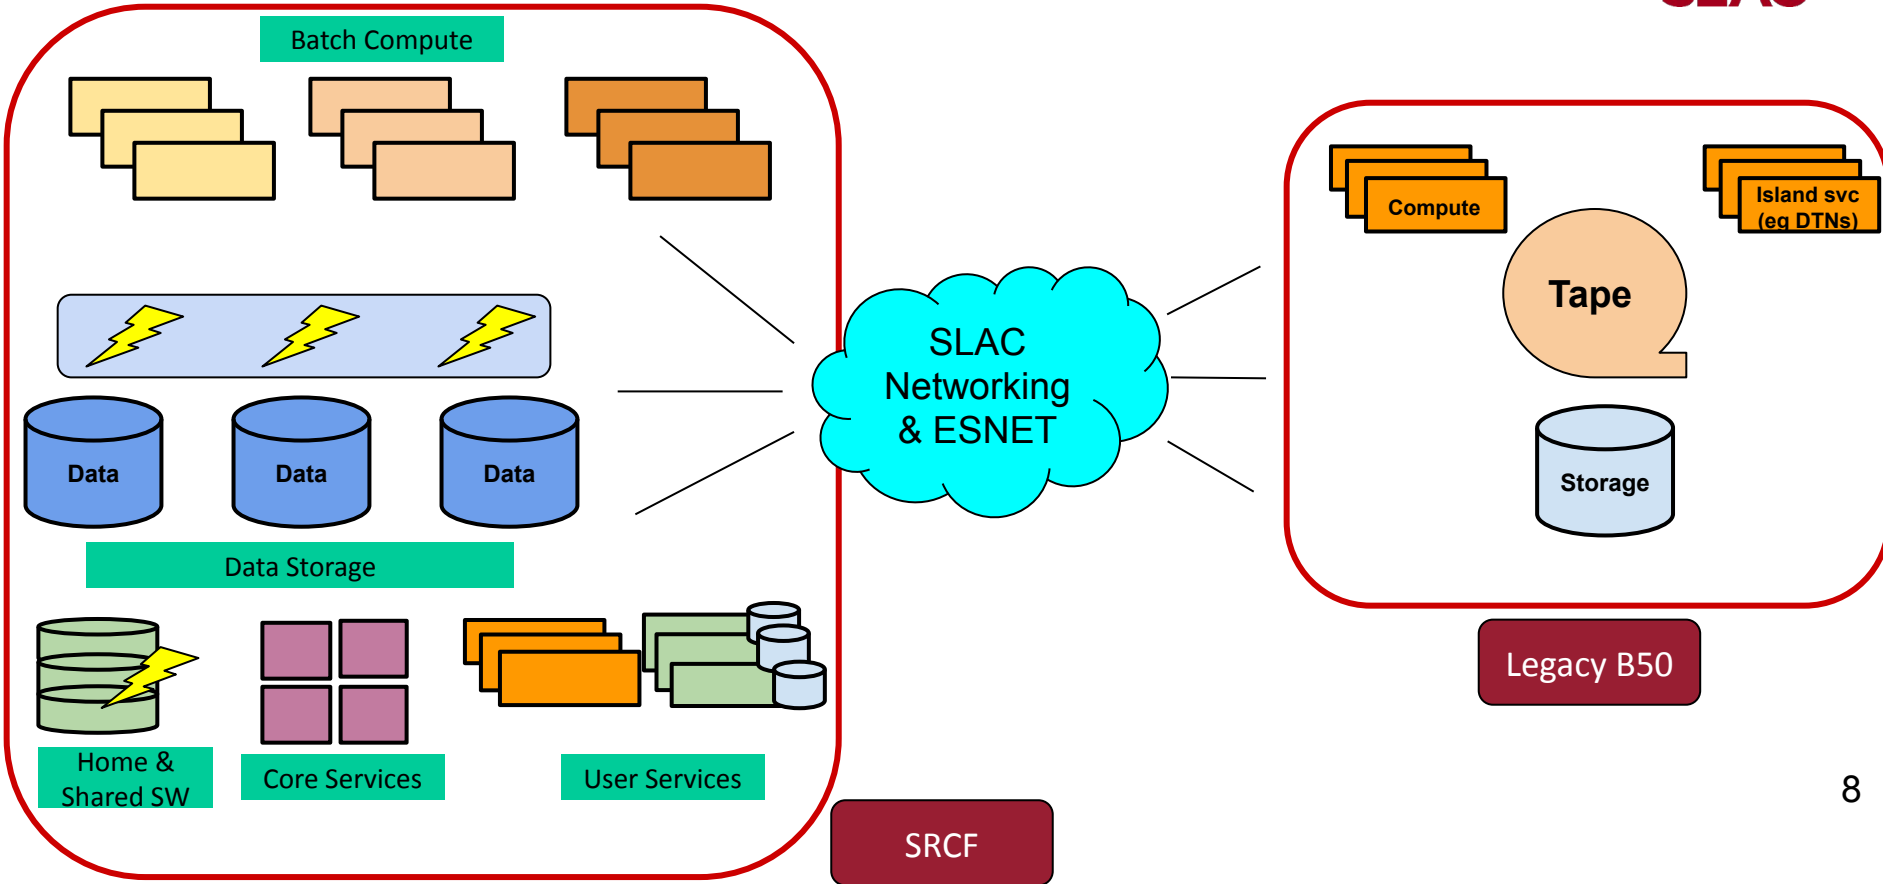
## SRCF Phase II



- 6MW Facility with air cooling (SLAC can use up to 2.5 MW)
- Flywheel + Generator allows for resilient power
- SLAC has over 100 racks
- 400 Gbps Networking to SLAC backbone

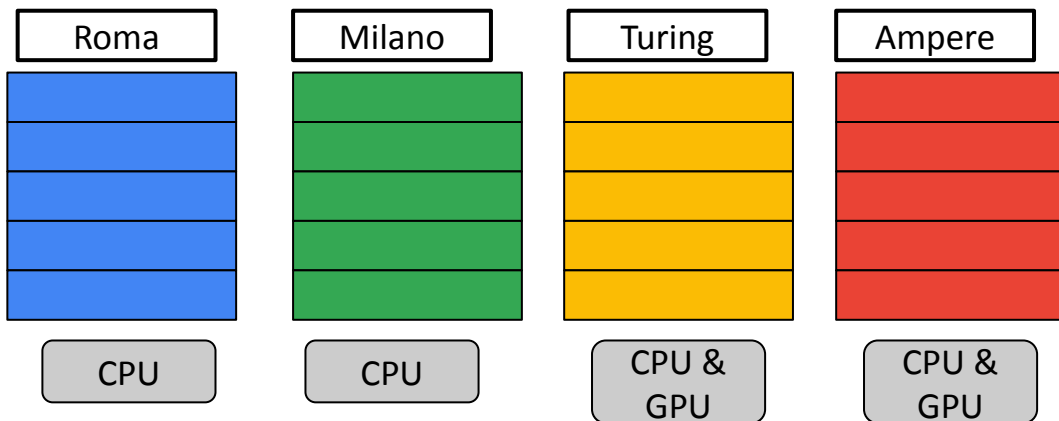
# S3DF Architecture

SLAC





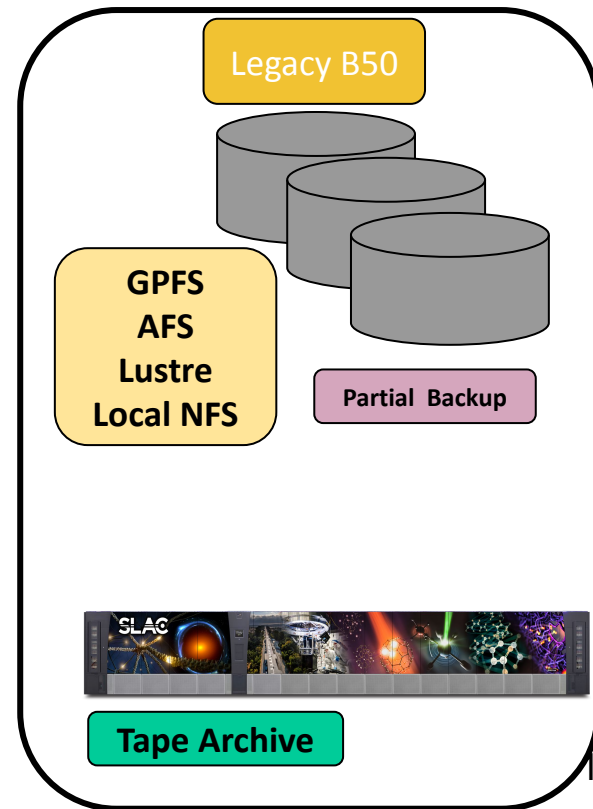
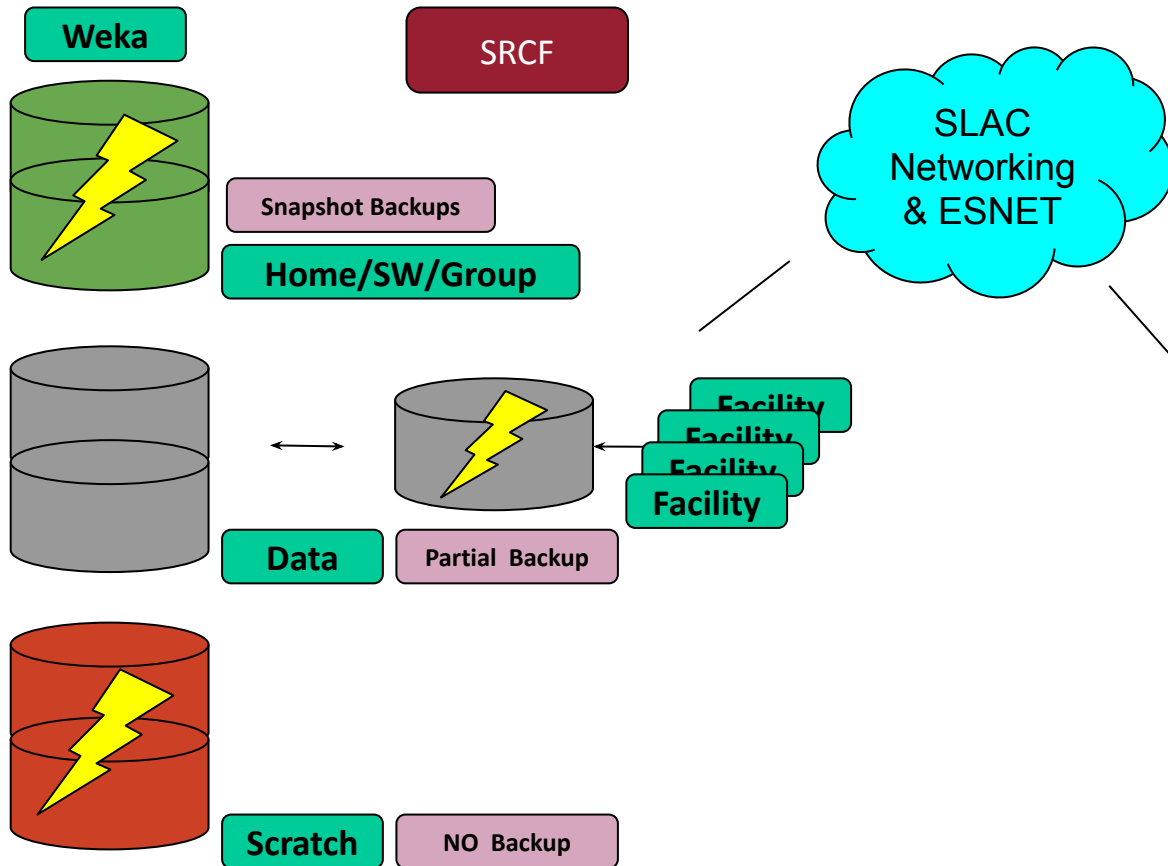
# S3DF Batch Compute Systems (Clusters)



- Compute Clusters (Partitions) are available to Facilities
- Facilities are entities that contribute to the S3DF - for example Rubin, AD, etc.
- Users within a Facility are assigned to Repos - which are a group of resources for a common project/experiment

- Access to compute resources is gated by Slurm Batch System
- Users specify Partition + Repo to gain Slurm access to compute resources
- Coact (<https://coact.slac.stanford.edu>) is the portal to look at and control all user allocations and setup, resource management, etc.

# S3DF Storage Systems



S3DF provides a Kubernetes-based (K8s) “Service Compute” offering to users.

- Supports resilient website and web applications for several groups
- Runs on a K8s cluster, which has access to S3DF storage
- Users may set up portals which are front-ends to web applications (for user access to resources, etc.)
- Workflow control services may also be run on the K8s cluster - with limited compute usage

S3DF provides VMs as well to users

- Limited offering based on very specific requirements
- Has some constraints in terms of resources available to the VM
- Preferable to use K8s offering instead for overall resilience and manageability

S3DF manages some Databases for selected, resourced, use cases (with support)

- Will need considerable user engagement & support

# S3DF Support model: Hardware contributions, installation and maintenance

## Rules of Engagement:

- Two calls per year to add resources to the system; hardware types to be determined by the **Steering Committee**
- Each contributing group is granted the **computing hours** corresponding to the hardware they buy for up to 5 years
  - Controlled via the banking/accounting mechanisms (Coact, but still being enhanced)
  - All users get low-priority access to the compute cycles of the entire system (low priority jobs may be preempted by higher priority ones)
- Each contributing group can **purchase storage hardware** for use up to 5 years
  - In addition each user gets access to a high-speed shared home and sw repos plus some scratch space (scratch policy will depend on actual usage, time and available resources)

Note: each contribution pays for the actual hardware plus 15% overhead for spares, installation and infrastructure costs

# Support model: operations



- **24/7 escalations** - priority for critical issues impacting production/uptime
- On-call rotations
  - Dedicated SLAC extension x7333 (S3DF) forwarded to the cellphone of the active team member in the rotation
  - Integrated with x4357 (HELP) for triage, escalation and redirection
- ServiceNow **ticket queues for users** (with email injection via sdf-help) - plus several internal ones
  - Integrated with it-help for triage, escalation and redirection
- **SLACK channels** - for rapid communication and notification

# S3DF Steering Committee

- User community will engage with SCS leadership to keep S3DF aligned with the lab mission
- The S3DF Steering Committee will lead this engagement
  - Representatives across the lab from all involved Directorates
- Effective, inclusive decision-making by engaging a small group (~10) that can represent all user groups
- The Steering Committee will be consulted on:
  - S3DF operational policies and priorities
  - Discretionary use of lab-funded (“free”) resources
- The Steering Committee is NOT
  - A technical design review team
  - A ‘Management By Committee’ team
- Membership will rotate periodically



Current 2024 S3DF Steering Committee Members:

- Jay Srinivasan
- Yemi Adesanya
- Richard Dubois
- Johannes Voss
- Greg White
- Jana Thayer
- Vivek Thampy
- Auralee Linscott Edelen
- Tina Cartaro
- Frederico Fiuza

**Questions?**