

SURF site updates

LTA development meeting 2026

Natalie Danezi
Onno Zweers

Data processing services

- Spider: local compute cluster for High Throughput Computing
- Grid: federated compute infrastructure, Tier-1 site for many experiments
- MS4: custom cloud for specific machines in High Throughput setups
- dCache: massive disk and tape storage for scientific data, LTA site

Typical projects: 20TB to 20PB data volume, 100K-10M core hours per year

Compute updates

Programmable, composable, and user-centric infrastructure

- Joint acquisition HPC/HTC for next data processing platform
 - New tender every 3 years
 - Conducting interviews
- SURF evergreen storage for flexibility & performance, one platform
 - Storage becomes independent from compute (less migrations)
 - Automatic tiering (SSD / HDD / Tape)
 - Increased speed / operations per second
- Industry standard interfaces and other interoperability features
- Co-design of capabilities for scientific application workflows

Spider Ceph improvements

- Ceph version upgraded
- Multiple MDS servers deployed (5 active for /home, /project, and large project spaces)
- Mount structure of projects changed, to one mount per project

Managed K8S in production


- Shared, multi-tenant Kubernetes cluster
- No operational burden
- Built for research
- Infrastructure isolation and service deployment through declarative configurations
- Automatically restarts failed workloads and scales under load

Hybrid architecture on kubernetes HPC/HTC

- K8S to orchestrate hybrid workflows across multiple clusters
 - unified view, reproducibility, portability (no SSHing into individual clusters)
- KuHOP (Innovation Program Future Applications & Infrastructure)
 - Use interLink to bridge K8S and SLURM
 - Multi-cluster orchestration (Spinlab/HTC → ETP/HPC → Spinlab/HTC)
 - Event-driven pipeline: S3 operation to Kafka, goes into Argo events, then Argo workflows schedules jobs on Slurm clusters

Extras in development

- Integrating and expanding object store support on Spider
- SRAM access on Spider
- Token based authentication on Grid/Gina
- Application Roadmaps
 - Assessment playbook
 - Design application roadmaps for selected use cases



Next events

7th October 2026 (@SURF Utrecht)

Unlocking GPU Power for Science: 2nd Symposium on Accelerated Computing for Application Workflows (SACAW)







- One-day event by SURF and the Netherlands eScience Center to bridge complex scientific code and GPU acceleration, no GPU expertise needed
- For researchers, domain experts, and tool builders, with a Flash Talks & Pitches session to share challenges or tools

8th December 2026 (@KIT Amsterdam)

Advanced Computing User Day (ACUD)

- Annual gathering for the advanced computing community, talks, workshops, and networking
- Call for Proposals now open (deadline Sept 6, 2026) across three tracks: HPC infrastructure, applications/workflows, and user experience
- Link: <https://www.surf.nl/agenda/advanced-computing-user-day>

dCache updates



API developments

- ADA Python interface (incorporating feedback from Astron 'team rainbow') in development
 - ADA Bulk-requests for stage/unstage implemented
- ADA Bash version upload/download is supported
- dCache metadata support (labels & extended attributes)
 - bugs in current implementation but promising

dCache doors

- New domain dcache.surf.nl
- Reorganising WebDAV doors
- HAproxy in front of doors is planned
 - High available set-up of dcache
 - Performance increase
 - Redundancy

Accessibility

- Connect SRAM (SURF Research Access Management) to dCache is planned
- OIDC tokens supported

Internal setup improvements

- Set up a full DTAP for dCache
 - Currently we only have Test and Production)
- dCache on Ceph for redundancy
- Better integration with tape system
- dCache all-in-one setup script
- DDoS protection rate limiter
- User mapping - moving to multimap
- Redundancy in dCache services
 - Doubled central services, failover capability
 - Reduce downtimes

Monitoring & logging & metrics

- New monitoring interface

<https://ganglia.grid.surfsara.nl/new>

Data management

- SRM phased out (officially in dCache v12)
- SciTAGs/firefly labels for network traffic
- NFS support more reliable and stable
- QoS data placement policies
- New technique for hot file replication

Thanks!

Questions?