

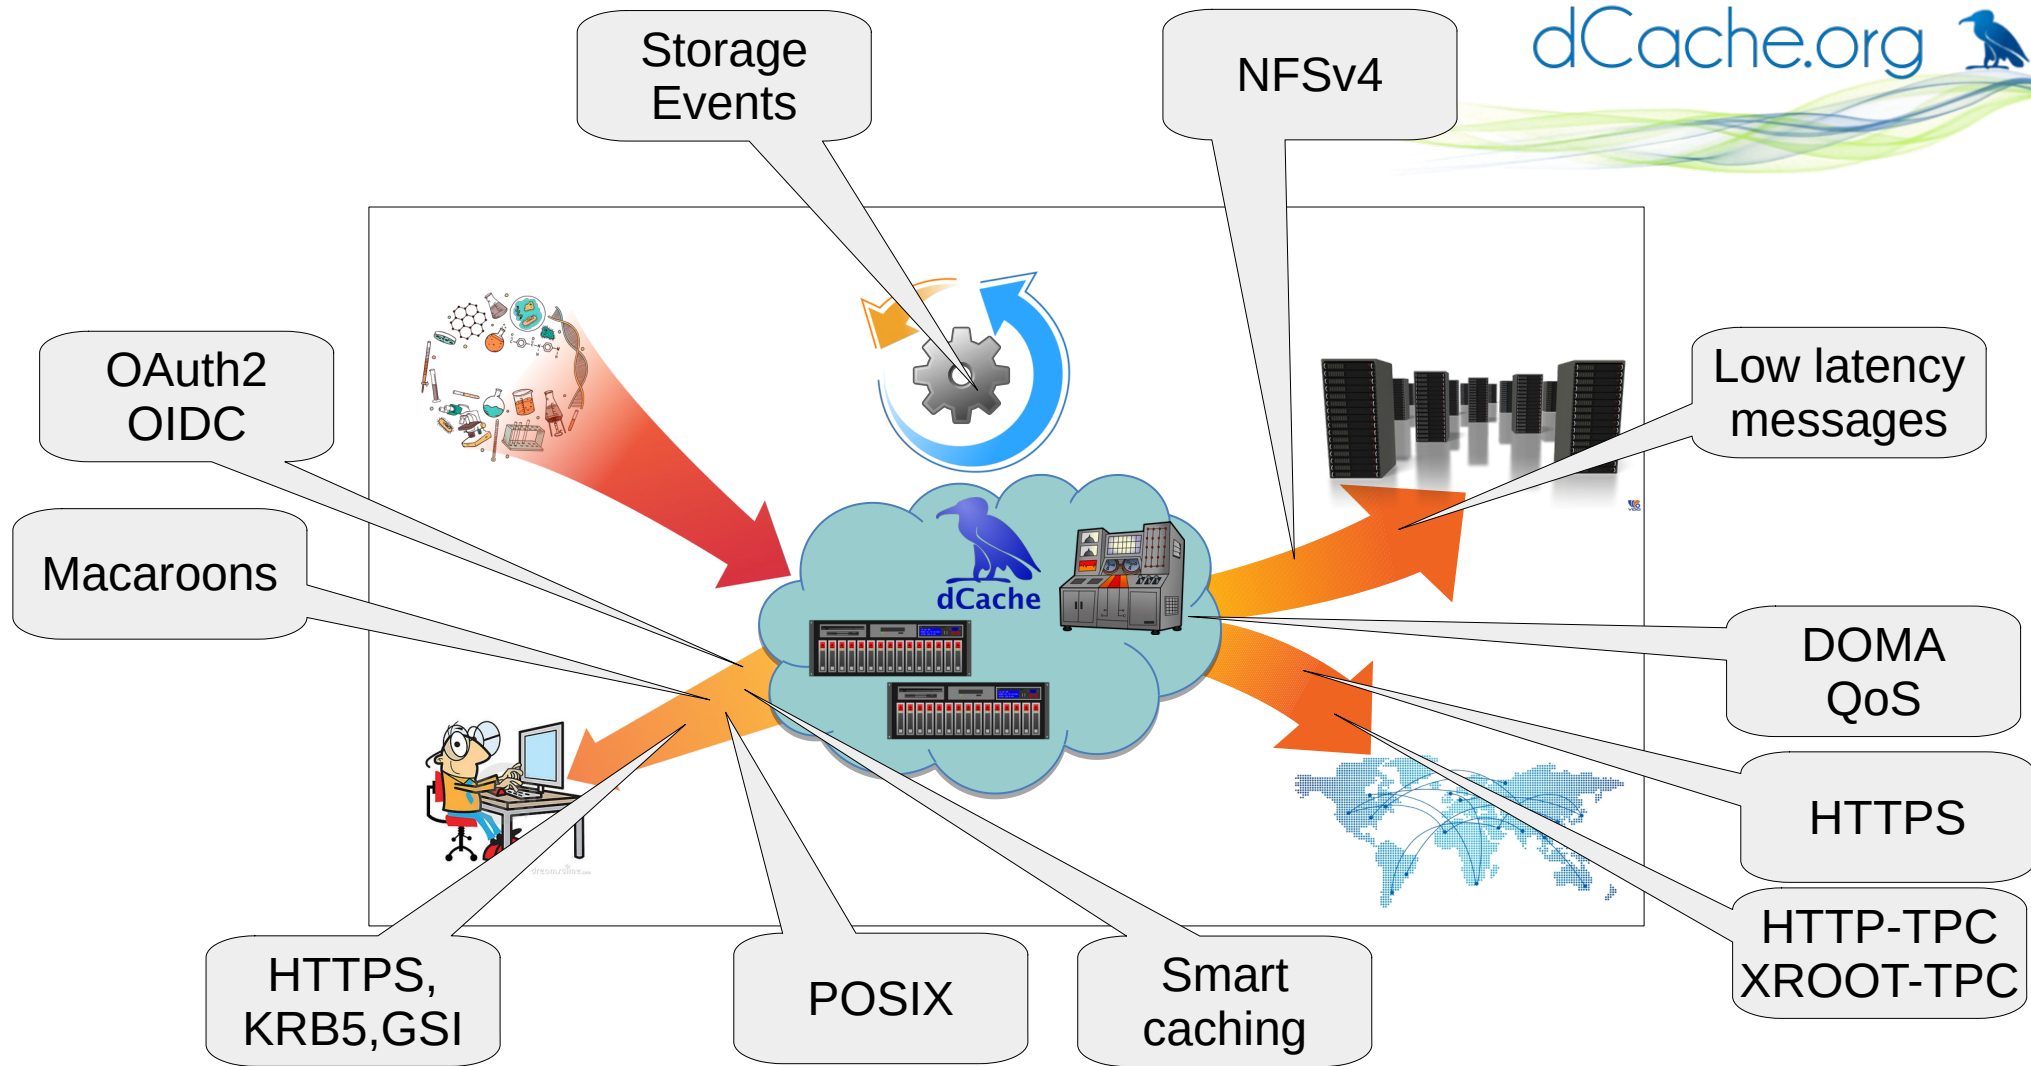
dCache User Workshop

Tigran Mkrtchyan for



Nordic e-Infrastructure
Collaboration







- WLCG / HEP
 - Almost all sites run dCache for WLCG
- DUNE / Photon Science / XFEL
 - Big local customer and stakeholder
- Astro- Bio- Life-science
 - Almost all sites have on-site communities

User Workflow Shift


- More non HEP tools and POSIX access
 - ROOT \Rightarrow Jupyter Notebook
 - Apache Spark
 - HDF5
- Growth of interactive analysis
 - Analysis Facilities
- Industry standard AuthN
 - OpenID Connect
 - OAuth2
 - Federated IDP
- Hybrid Clouds
- New 3rd-party transfers protocols
- Integration with HPC clusters



Protocol Usage (DESY-HH)



NFS access dominated by
Atlas & CMS



- x2 1.5h (2h) Zoom slots 16:00 – 18:00 CEST
 - dCache project report
 - Site session
 - Free speech session
- ~50 participants
 - Not the same people on both days
- Primary/backup speaker
 - DESY/Fermilab
- Video capture

dCache session

Breaking
changes!

Xrootd &
Co.

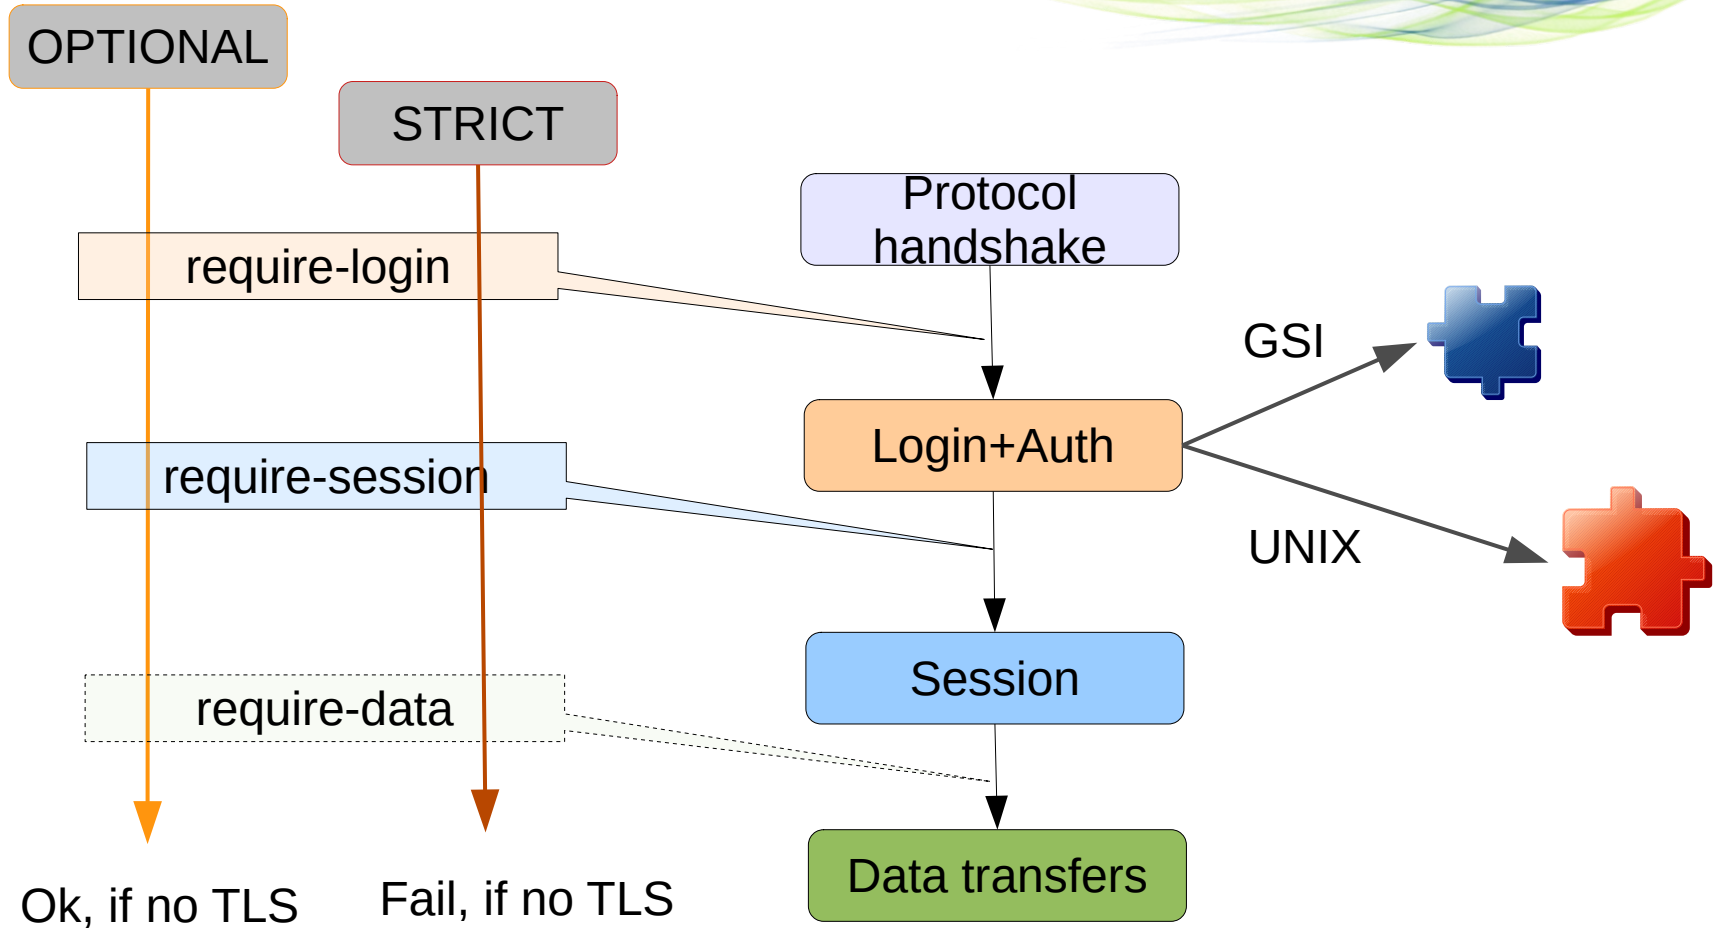
HSM & QoS

dCache.org
Watching you!

Datalakes,
fjords & deserts

Last, but
not least!

45m \Rightarrow 7m



Rule of Thumb

- GSI

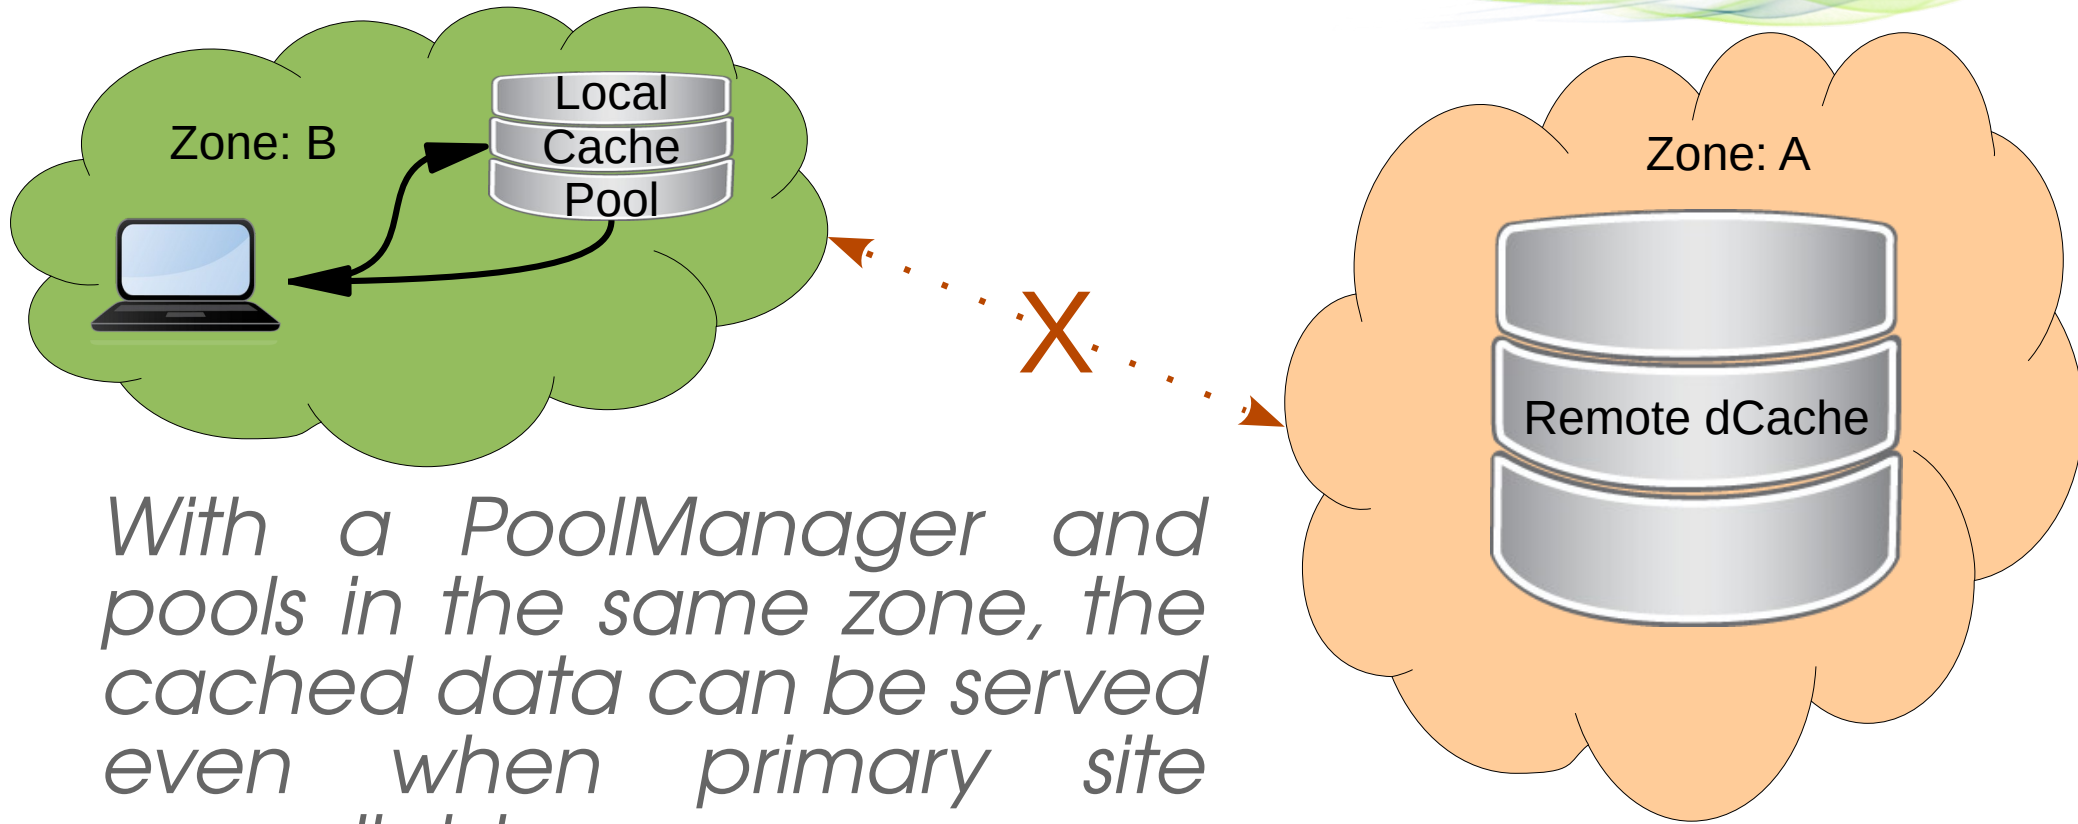
mode=**OPTIONAL**

required-**session**=true

- SciTokens

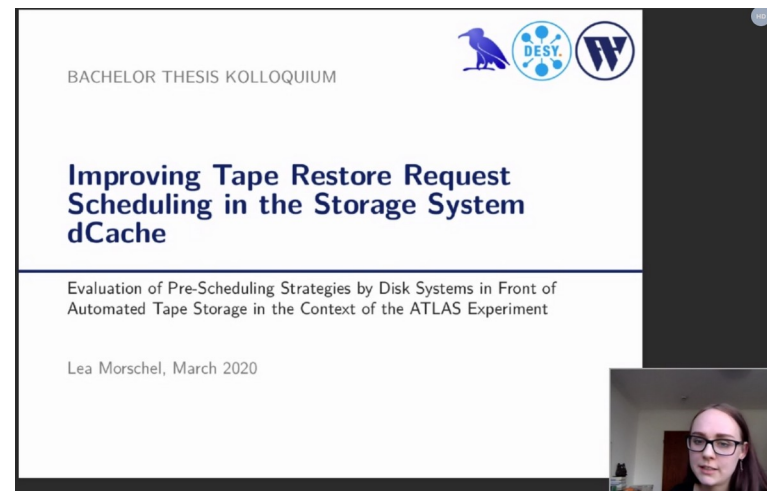
mode=**STRICT**

required-**login**=true



With a PoolManager and pools in the same zone, the cached data can be served even when primary site unavailable.

- Atlas tape carousel => WLCG Tape carousel
 - Bad habits die hard
- Tape driver interface evolution
 - Docu update
 - New drivers
 - Functionality enhancements

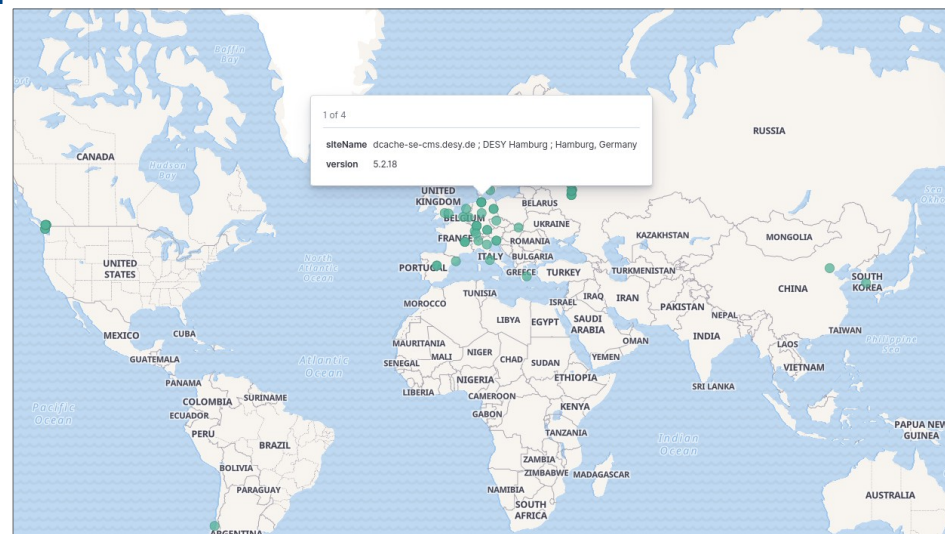


The screenshot shows a presentation slide with the following content:

- BACHELOR THESIS KOLLOQUIUM** (top left)
- Logos for **DESY** and **W** (top right)
- Improving Tape Restore Request Scheduling in the Storage System dCache** (main title)
- Evaluation of Pre-Scheduling Strategies by Disk Systems in Front of Automated Tape Storage in the Context of the ATLAS Experiment** (abstract)
- Lea Morschel, March 2020** (author and date)

A small video inset in the bottom right corner shows a woman with glasses speaking.

- See “the community”
 - Show where dCache is used
 - Good PR for the project
- See versions around
 - Version popularity
 - Planing back-ports



What We Collect, Where

- Site Name
 - `telemetry.instance.site-name=Friends Of dCache`
- Geo-location
 - `telemetry.instance.location.latitude=`
 - `telemetry.instance.location.longitude=`
- dCache version
 - auto detected
- Online capacity
 - calculated
- Stored in DB at DESY
 - `telemetry.destination.url=https://stats.dcache.org/collector`

- Sites have to enable service explicitly
- No Personal data collected
- Data used by dCache developers only
 - we might use it in our presentations!
- Site information can be removed on demand
- **Consult your SITE SECURITY OFFICER before enabling!**

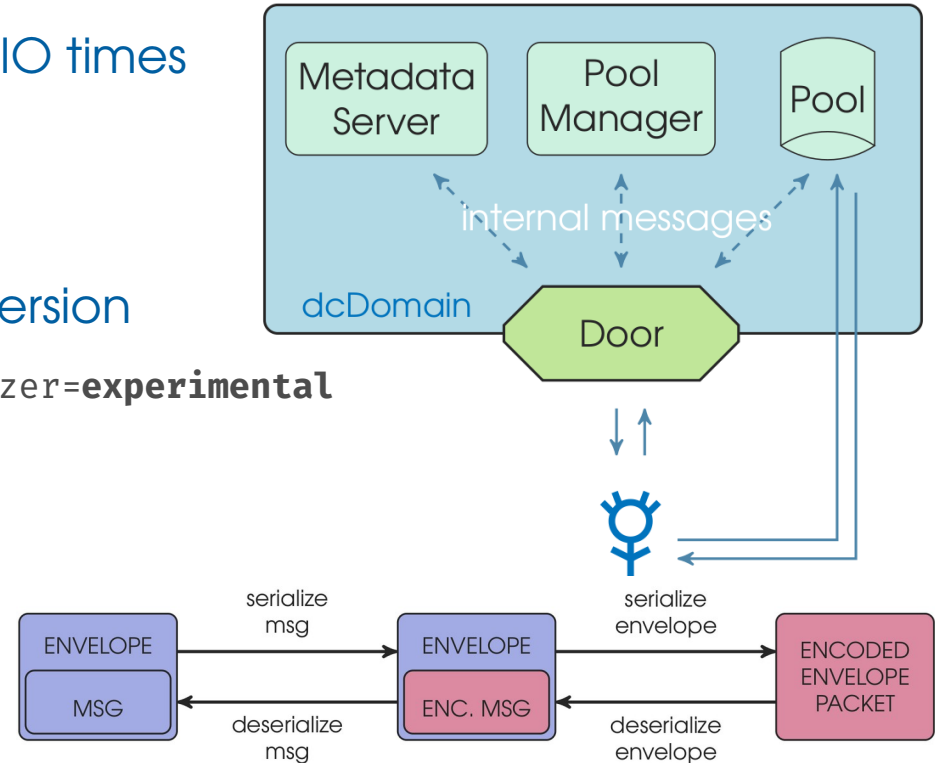
Pool Start-up

- Pool sizes increase
- File sizes decrease
- Typical pool has >1M files
- dcache-lab – 4M files, (cold) startup time 45min.

New Message Serialization

- Interactive usage / HPC via NFS/WebDAV
 - Latency for the individual user becomes noticeable!
- Inter-cell communication comparable to IO times
- New experimental message serializer:
 - (De)serializing ~ 10% faster!
 - All components run the same minor version

`dcache.broker.channel.msg-payload-serializer=experimental`

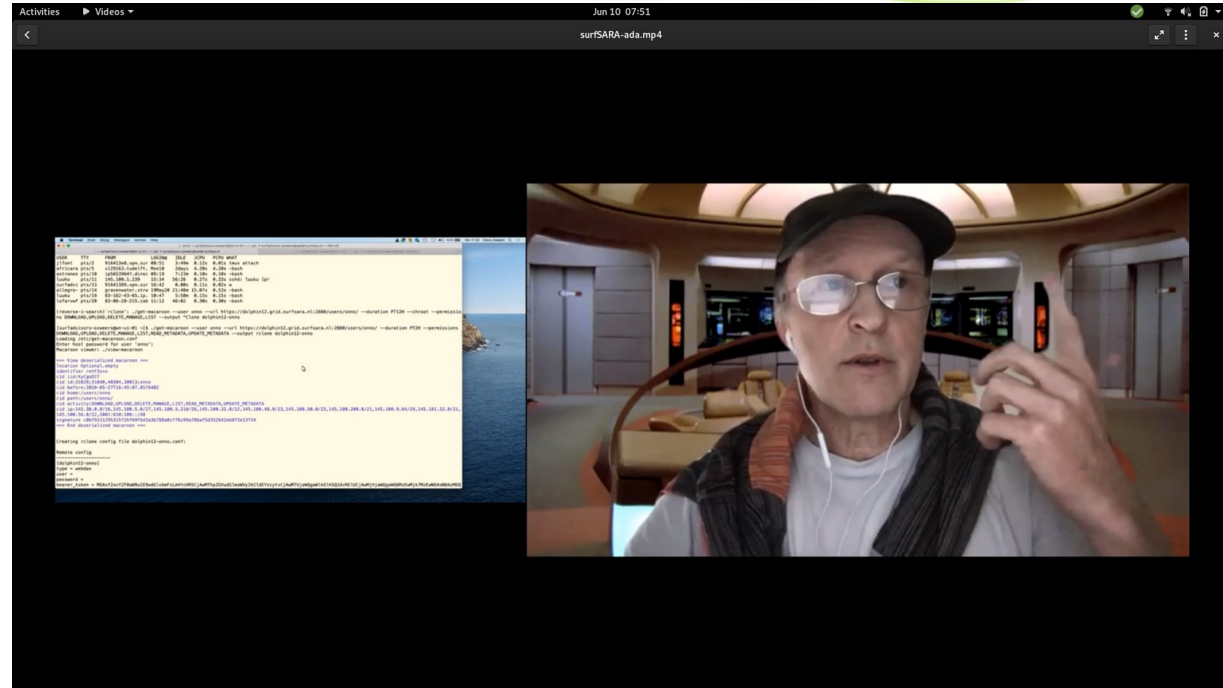


systemd or sysvinit?

- Our Zoo
 - deb – systemd (only two sites?)
 - rpm - sysvinit
- Lot of custom functionality covered by systemd
 - permission drop
 - demonizing
 - auto-restart on crash
- We need your help to make it happen

Access dCache API - sufrSARA

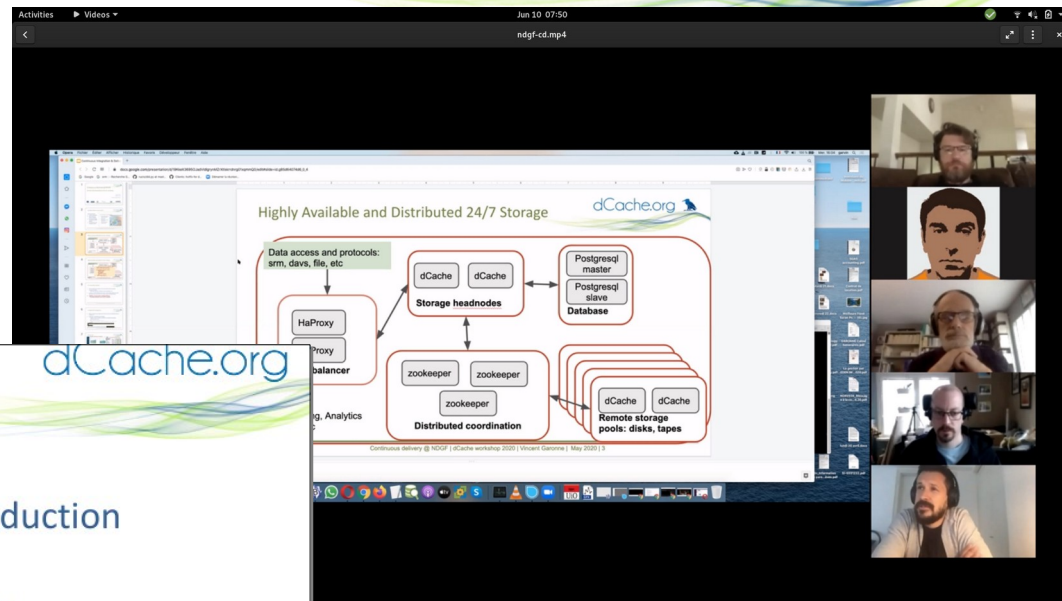
surfSara
developed tool
to access
dCache via rest-
API



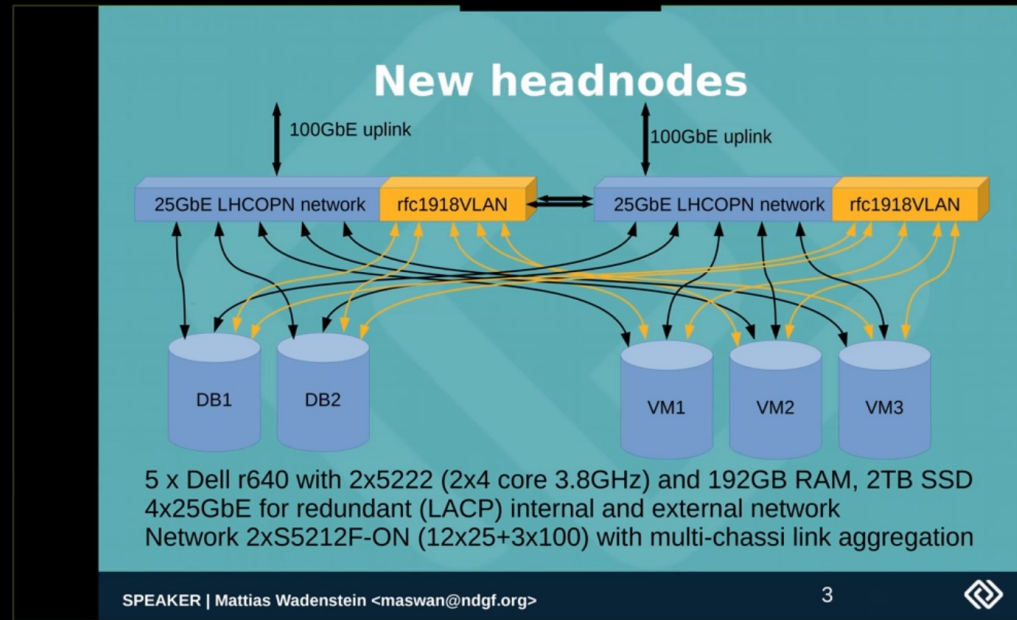
<https://github.com/sara-nl/SpiderScripts>

Pre-production Instance

- The pre-production instance is the exact clone of the production deployment model
 - Haproxy, headnodes, database, remote disk pools, etc
- Load by the ATLAS DDM functional test, hammer clouds ARC jobs, Paul's dav smoke tests and probes: data access, copy, stats and deletion
- **Validation of changes before reaching production**
 - **Fix, detect and report issues as soon as possible**







New Hardware 20



Conclusions

- Video workshop was a new experience
- Positive feedback from participants
- Lot of spontaneous discussions
- We are planning spontaneous sessions for selected topics

Support/Contact Channels

- **support  dcache.org**
 - User request tracking system
 - Place to ask for a help from developers
 - Accessible by all team member
- **security  dcache.org**
 - Request tracking system
 - To report security issues or incidents
 - Accessible by selected people
- **user-forum  dcache.org**
 - Mailing list for sysadmins self-help group
 - To ask for an advice or share experience
 - Used by (almost) all sysadmins and developers
- **dev  dcache.org**
 - Shared mailbox
 - An email to contact developers. Not for support
 - Developers can send e-mail from this address
- **srm-deployment  dcache.org**
 - Tier – 1 coordination mailing list
- **meet.desy.de/xxx**
 - Weekly Tier 1 support video call
- **workshop  dcache.org**
 - Shared mailbox
 - An e-mail used to organize workshops
- **Github issues**
 - Request tracking system
 - To report software defects and feature requests
 - Public
- **Github pull-requests**
 - Request tracking system
 - To provide code changes
 - Public

Thank You!

