dCache - Inter-Disciplinary Storage

vCHEP 2021

Tigran Mkrtchyan for the dCache collaboration

HELMHOLTZ RESEARCH FOR GRAND CHALLENGES

The MISSION

“... to provide a system for storing and retrieving huge amounts of data, distributed among a large number of heterogeneous server nodes, under a single virtual filesystem tree with a variety of standard access methods.”

https://dcache.org/about/

Scientific Data Challenges

Ingest

- High data ingest rate
- Multiple parallel streams
- High durability
- Effective handling of large number of files

Analysis

- High CPU efficiency
- Chaotic access
- Standard access protocols
- Access control
- Local user management

Sharing & Exchange

- 3rd party copy
- Effective WAN Access
- In-flight data protection
- Identity federation
- Access control

Long Term Preservation

- High Reliability
- Self-healing
- Automatic technology migration
- Persistent identifier

Strategic Communities

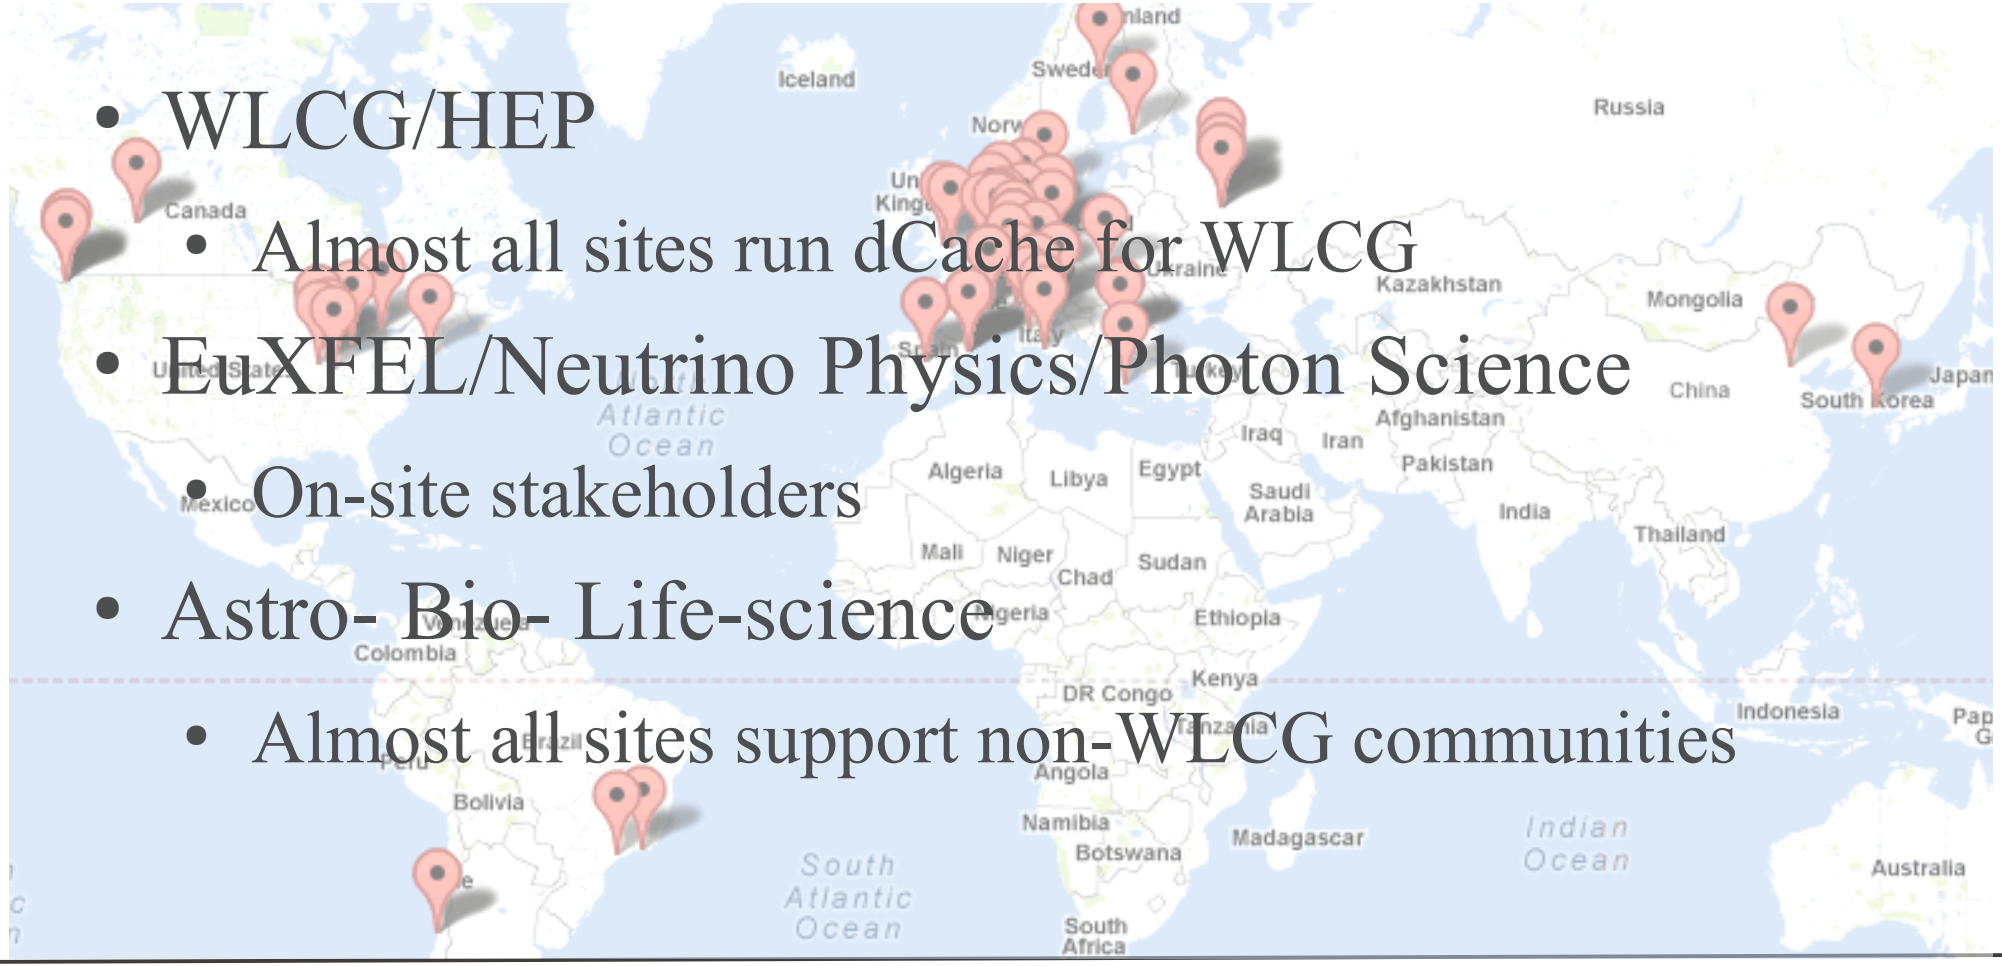

- WLCG/HEP
  - Almost all sites run dCache for WLCG
- EuXFEL/Neutrino Physics/Photon Science
  - On-site stakeholders
- Astro- Bio- Life-science
  - Almost all sites support non-WLCG communities

Data Access Variety

- ROOT-IO
- Non-HEP tool chain
  - Active use of Jupyter Notebooks
  - Non-ROOT data formats
- Industry standard AuthN
  - Tokens based authentication
  - Federated IdP
- Use of private clouds
  - Data access from a container
- Use of HPC resources

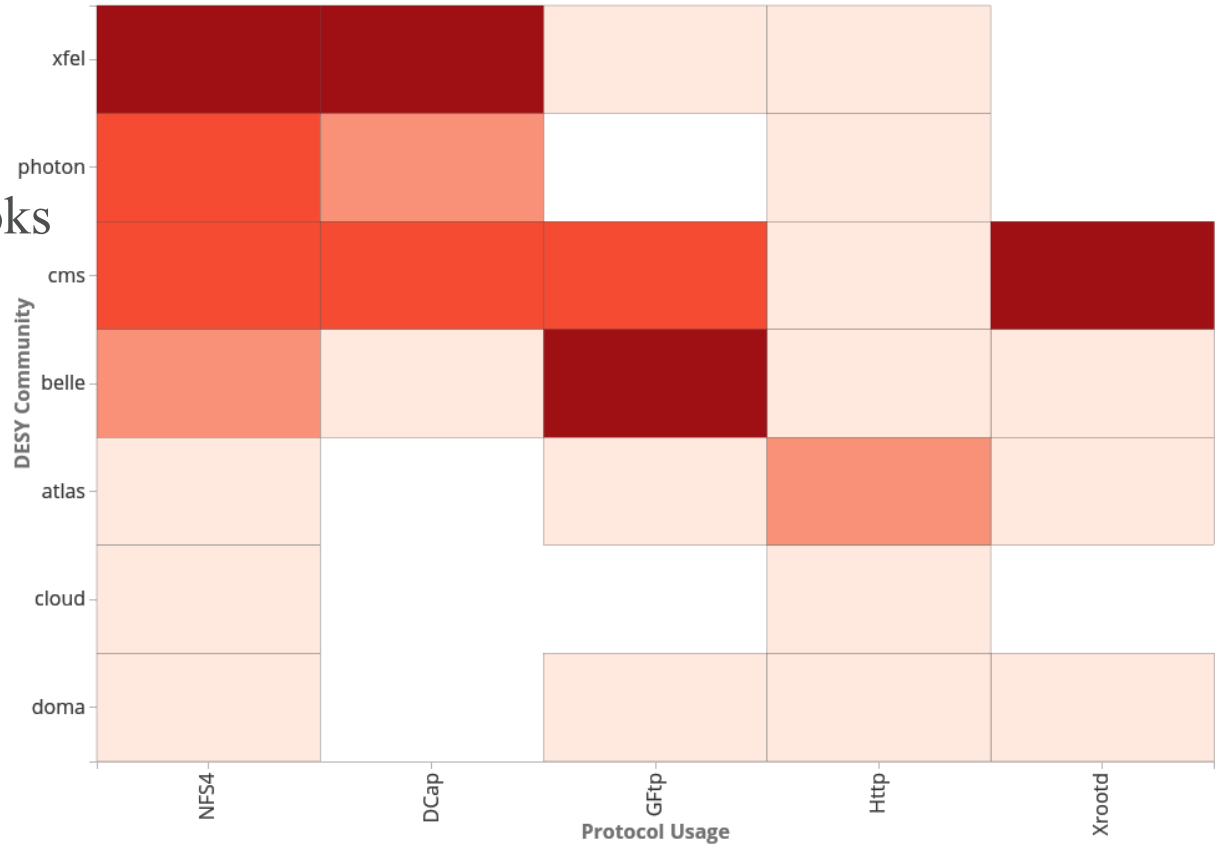

HSM, Tape, QoS

- ATLAS “Tape carousel” ⇒ WLCG “Data carousel”
  - Stay tuned for Lea’s presentation: Improving Performance of Tape Restore
- High number of small files by Photon Science
  - ~4MB, 10^6 files per directory
  - See Svenja’s presentation at HEPiX https://indico.cern.ch/event/995485/contributions/4256474/
- Multi-media copy guarantees
  - Stay tuned for Al’s presentation: From Resilience to Quality of Service
- Integration with CERN Tape Archive (CTA)
  - Close work with CERN team to bring CTA support to dCache

v1 Bulk REST-API (like SRM, but different)

STAGE

- Request to stage many files at once

CANCEL

- Cancel bulk request

DELETE

- Cancel bulk request + clear history/status

EVICT

- unpin cached copy

PIN

- Pin cached copies with a lifetime

FILEINFO

- Request status many files at once (locality, checksum)

StoRM

Automatic Metadata Population

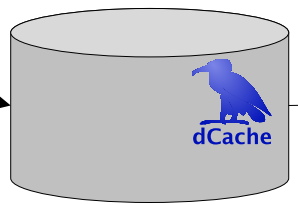

Event: new data

extract metadata

Automatic catalog update

metadata catalog

#Anna Bertha Ludwig

#hand #1895 #x-ray

Metadata Population

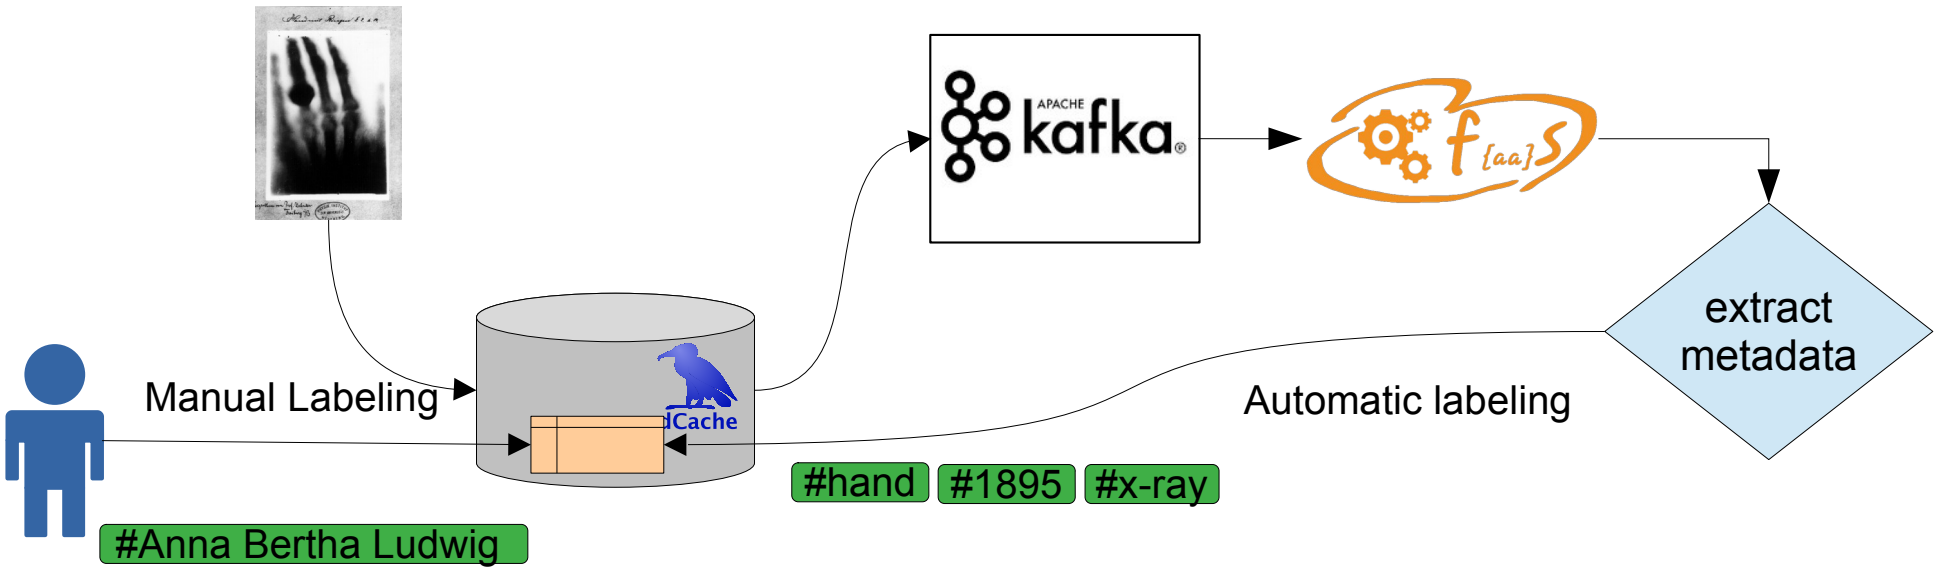

User Metadata/Labeling in dCache

- Extended attributes
  - Exposed via NFS, WebDAV, REST
- Label-based virtual **read-only** directories (WIP)
  - List all files with a given label
- dCache rules applies
  - Visible through all protocols
  - Respect file/dir permissions

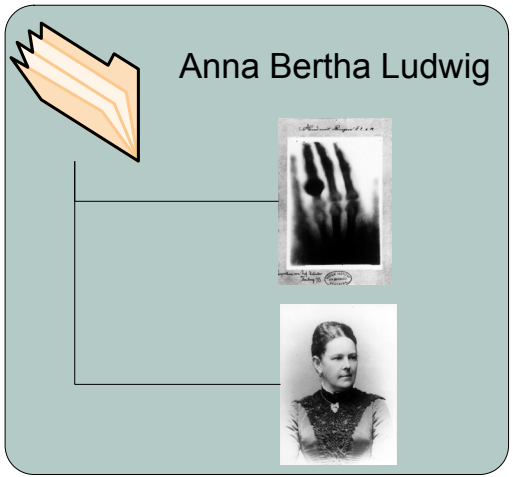

Keycloak Integration

- Open source SSO & IAM solution
  - Generates *identity* and *access* tokens.
  - Popular at many-sites including DESY
  - Delegated group membership management
- Supports standards-based OpenID-connect, OAuth2 and LDAP
  - ***preferred_username*** to map to LDAP accounts
  - ***eduPersonEntitlement*** mapped to group membership
- dCache's gPlazma config can be combined with X509, VOMS and others
  - Some code changes are needed to improve integration

Keycloak Integration

PaNOSC

Laser-Driven Proton Acceleration from Cryogenic Hydrogen Target

Description
2D particle-in-cell simulation of the interaction of high-intensity laser pulse (parameters are relevant to L4 laser) with a cryogenic hydrogen target. Only protons with energy above 300 MeV at the end of the simulation are tracked and their position and energy are visualized. Two different groups of protons accelerated by different mechanisms can be distinguished from each other in space: Protons originated from the target interior and from the target rear side.

Citation	Dana Scully; (2020), Re-polarization of the aft quantum plasma collector, DOI:10.9563/if.2015.87.012
Keywords	X-ray excited optical luminescence,
Type	Proposal
Author	Laima Reinhold
Other	Stuff

Stolen from Michael Schuh

Datasets

PaNOSC Test Dataset 11
HEIMDAL @ ESS

Name
my-environment
Descritption
re-produce research
Flavour
jupyter_small
Spawn

Preview Visualization

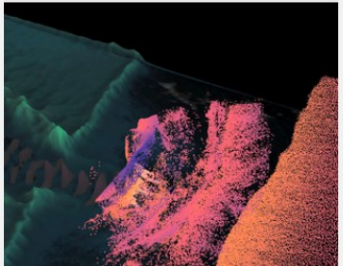

Summary & Conclusions

- The dCache team has been providing a reliable software to manage scientific data for over 20 years.
- Seamless integration into the site's infrastructure makes dCache a natural part of any data center.
- Multi-protocol and authentication scheme capabilities allow to support multiple communities even on a single instance.
- In close cooperation with experiments we address today's and future data management challenges.

Thank You!

More info:

https://dcache.org

To steal and contribute:

https://github.com/dCache/dcache

Help and support:

support@dcache.org, user-forum@dcache.org

Developers:

dev@dcache.org