Federated dCache

Pools at Different Pool-Sites, Single Federated Management

Christian Voß, for DESY-HH dCache Operations Team and dCache Development Team

Mass-Storage for LHC and Others

dCache as Central Mass Storage for Many Communities

dCache.org

- Central element in overall storage strategy
- Collaborative development under open source licence by
  - DESY
  - Fermilab
  - Nordic E-Infrastructure Collaboration (a.k.a. NDGF)
- Particle Physics in general
  - In production at 9 of 13 WLCG Tier-1 centres
  - Seamless integration of tertiary storage (e.g. tape)
  - In use at over 60 Tier-2 sites world wide
  - 75% of all remote LHC data stored on dCache
  - In addition: Tevatron and HERA data
- DESY
  - Raw-Data for smaller Particle Physics experiments
  - Raw-Data for Photon Science Archival
  - Mass-storage for user data during analysis
  - Long-term archival

Features

- Highly horizontally scalable storage system
- Expose a single unified namespace
- Supports many protocols
- Supports many authorisation schemes
- Micro-service architecture

Basic Setup

Standard Single Site Setup

- Use dCache: Access to `/pnfs/desy.de/cms`

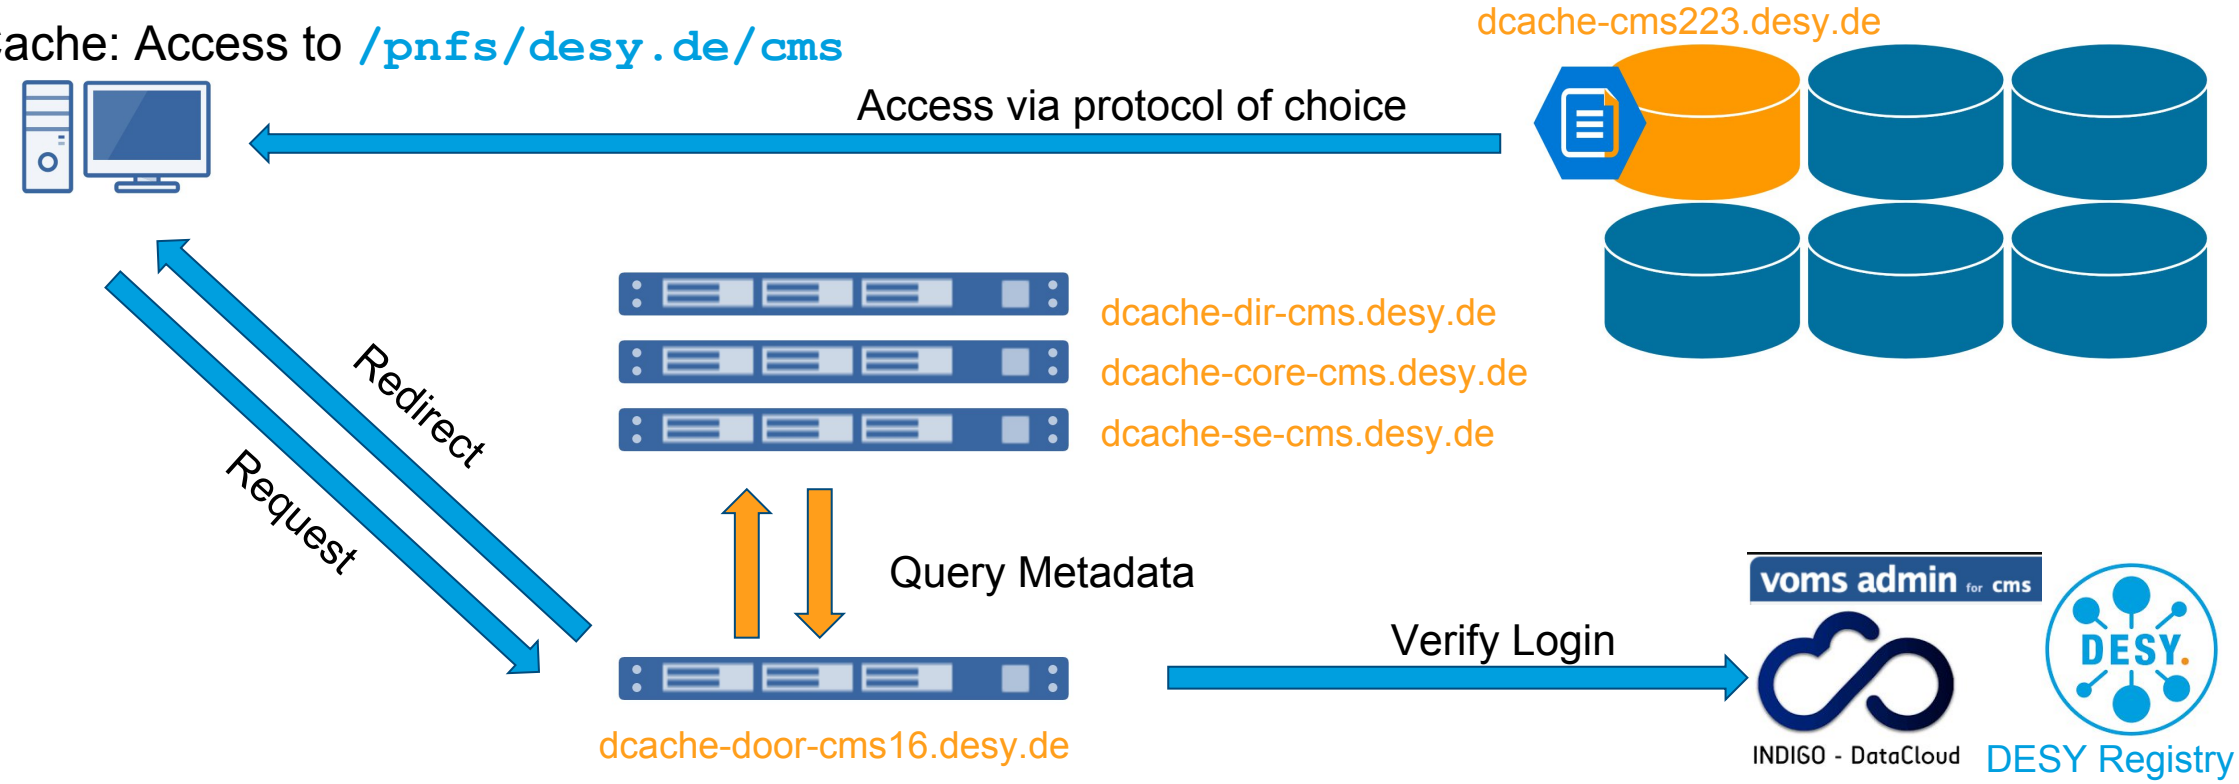

Placing dCache Pools Off-Site

Forming a Federated dCache

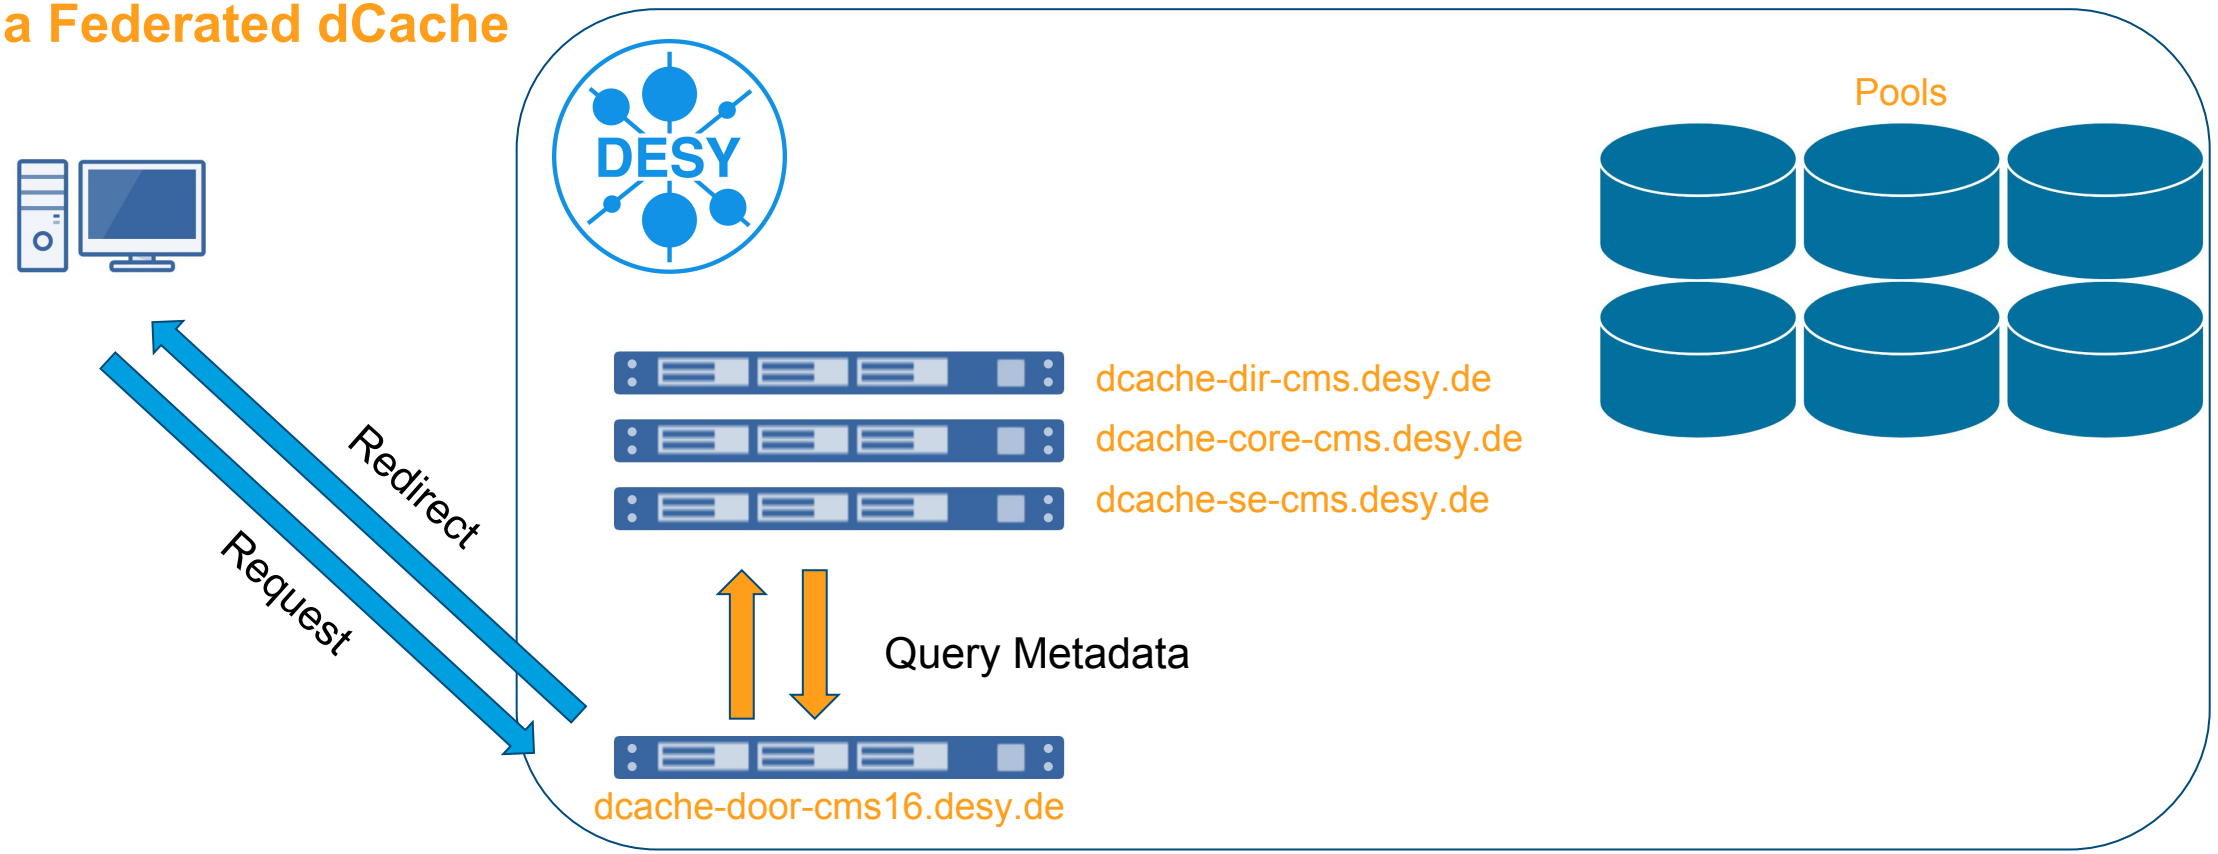

- Use dCache: Access to `/pnfs/remote-site.de/`

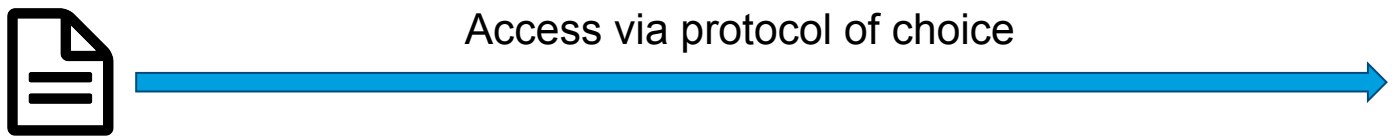

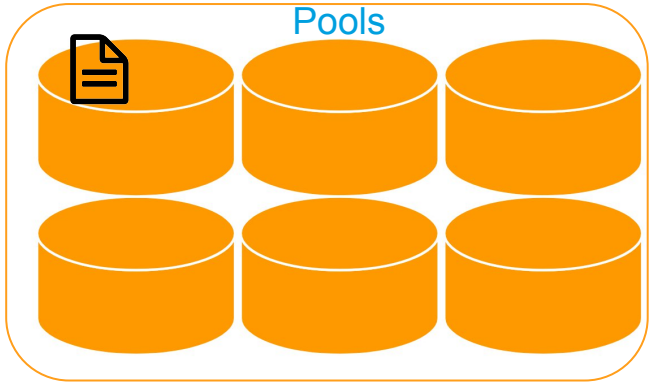

Placing dCache Pools Off-Site

Forming a Federated dCache

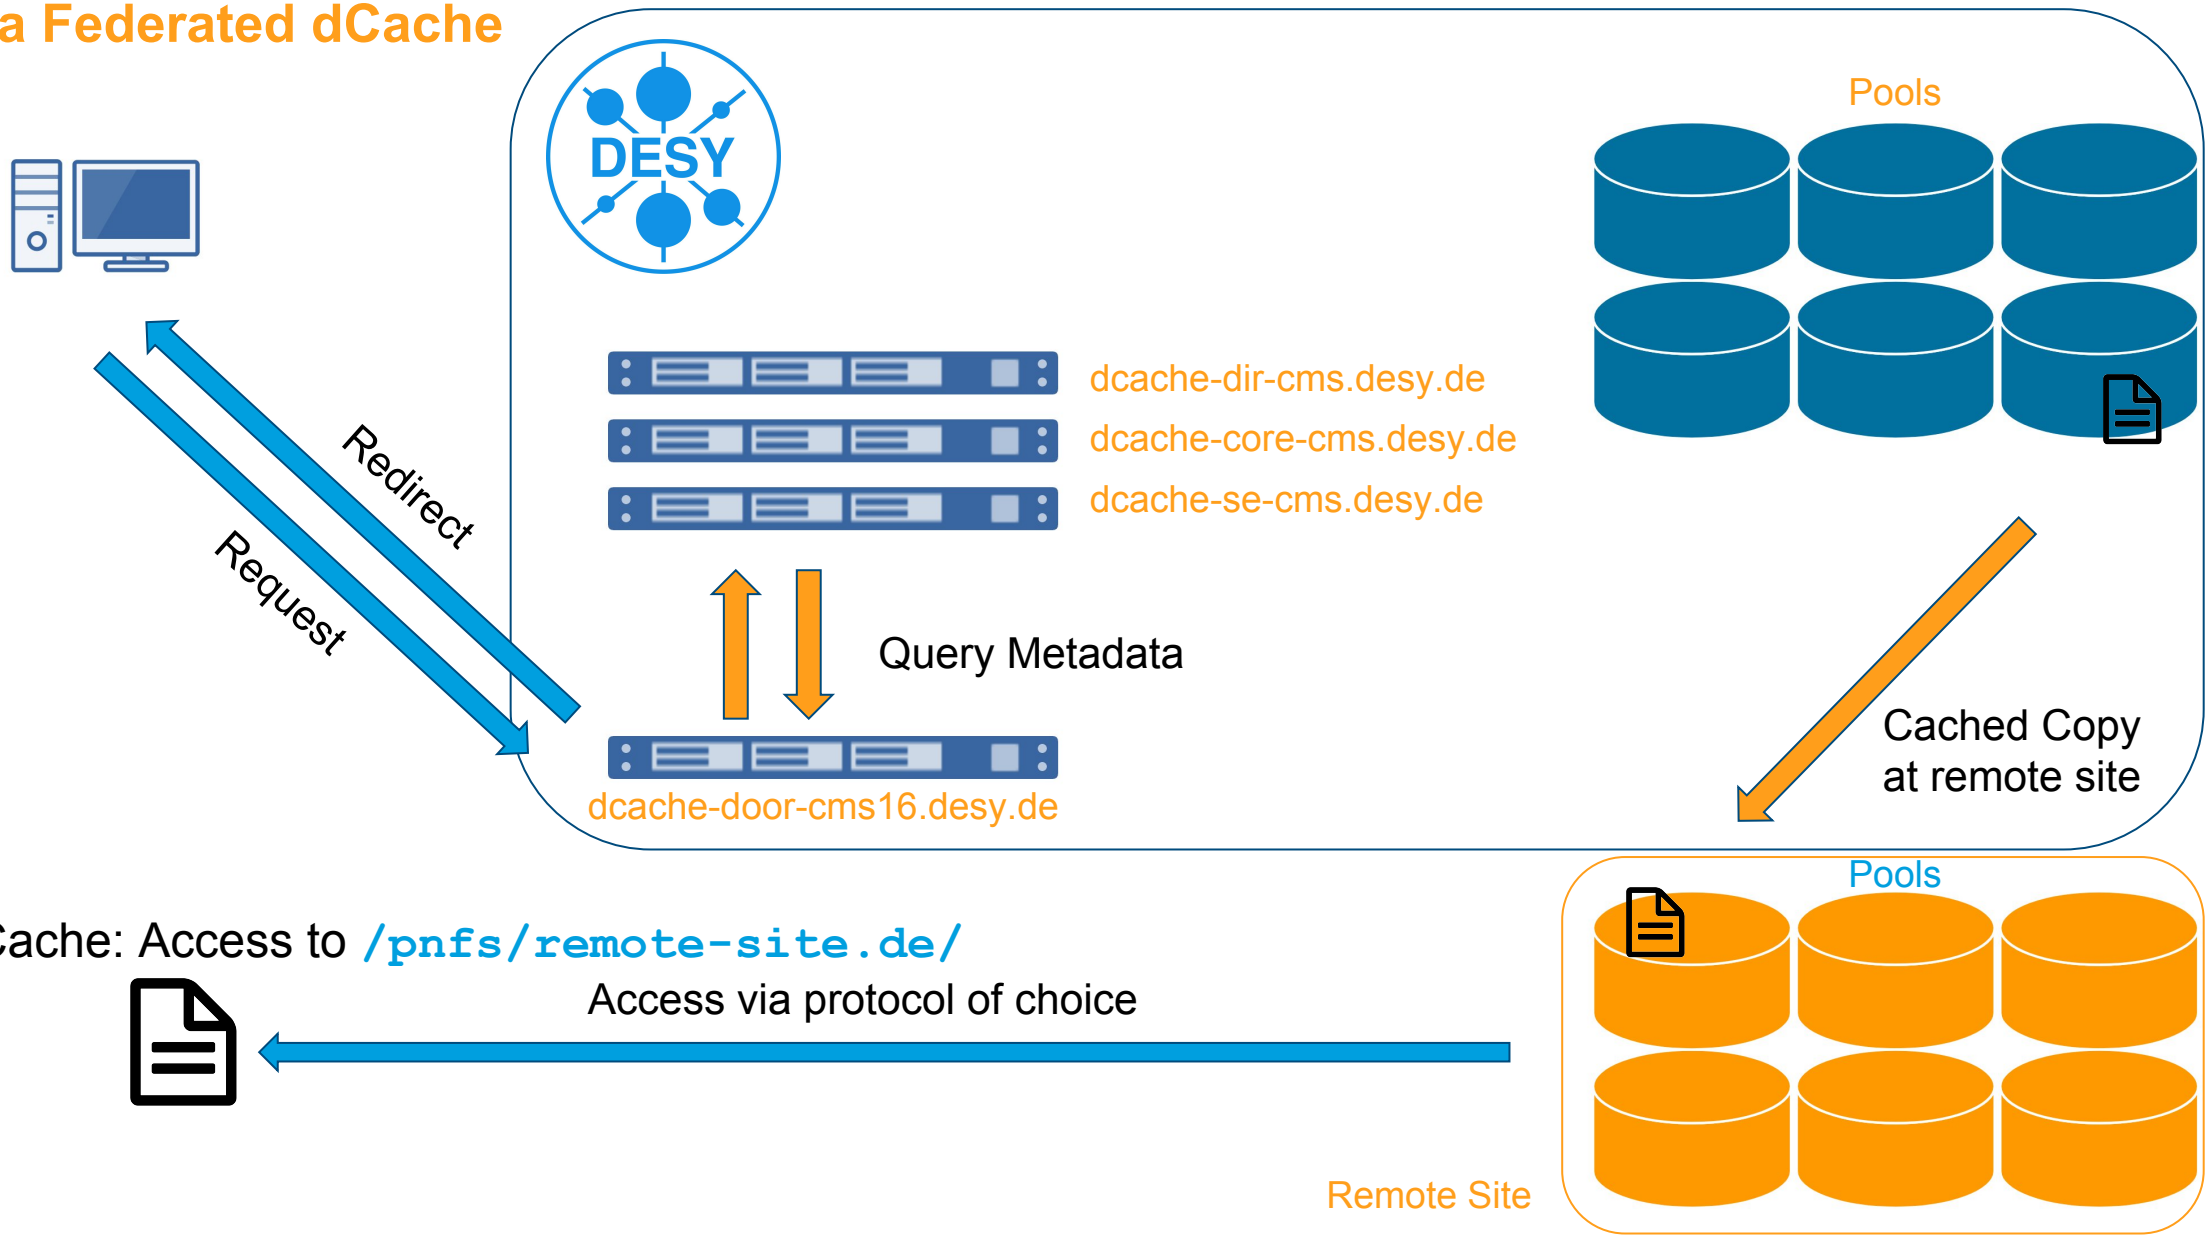

- Use dCache: Access to `/pnfs/remote-site.de/`

Advantages of Using dCache

Easy Example: Fully dependent Remote Site

- Small requirements: require space on a file system at a site and require container support or Java support
- Offer Tarball/Container with full deployment and configuration, need only a mount for the disk path
- All other services remain at central host site➜ limiting the need for storage admins at remote sites

Smaller Communities

- Could used to create a uniform namespace (Could eliminate need for dedicated catalogs/Data Management):
  - `/fsp-condensed-matter/university-of-rostock/DFG-projects/`
  - `/fsp-condensed-matter/university-of-lubeck/PETRA-III-results/`
- Local files will be read locally, remote files will be copied to local site on demand automatically
- Access through many protocols and often more importantly: POSIX-like access

Larger Communities

- Host site would provide an independent storage endpoint and namespace
- All files written to these paths and endpoint would be stored at the NHR ➜ for experiments: a regular site
- Full range of supported protocols: NFS, WebDAV, XrootD, Staging capabilities at remote site

Simple Setup on a Remote Site

Simplest Setup Possible – Expected Loops in Existing Network Security

Cache-Only-Site

dcache-remote-site.desy.de

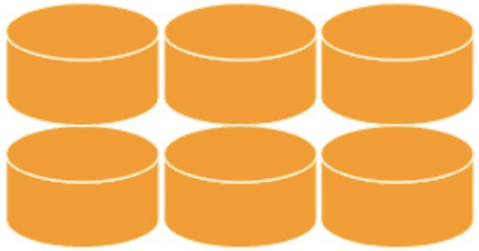

Hidden

- Option for local Compute at Unis
- A connection to DESY would allow access from DESY

Cache-Only-Site

Bring-Online

dcache-random-nhr.desy.de

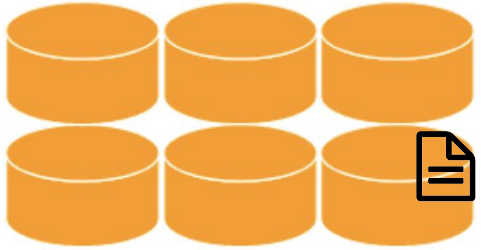

Internal Transfer

Hidden

Summary

Advantages and Disadvantages

Features

- Centralised namespace (think of it like Rucio)
- Centralised interface/connection to the AAI
- Resilience configurable
- Configured as permanent storage or cache

Advantages

- Provide a variety of protocols for the users
- Expose only one endpoint to experiments
- Single, centralised configuration and administration
- Little load on admin at remote sites (pools only)
- Reduce invest costs at remote site (pools only)
- Similar setups in production: NDGF-Tier1, Great Lakes Tier2

Disadvantages

- Reduced independence of remote site
- Reliance on stability and responsiveness of local/remote sites
- Reliance on WAN for all metadata operations (can be mitigated)
- Setup requires expert knowledge ➜ **future funding important to make this more steam lined**