Data Management Update

Brian Bockelman
19 March 2012

WLCG TEG OMG

- WLCG TEG = Worldwide LHC Computing Grid Technology Evolution Group.
- These 5 working groups have been mandated to produce a document describing the *status quo* and come with a *recommendation* on how the WLCG should evolve in the next 2-3 years.
- The working groups will be giving short reports during the AHM. I'm the co-chair of the Data Management TEG.

Status Quo

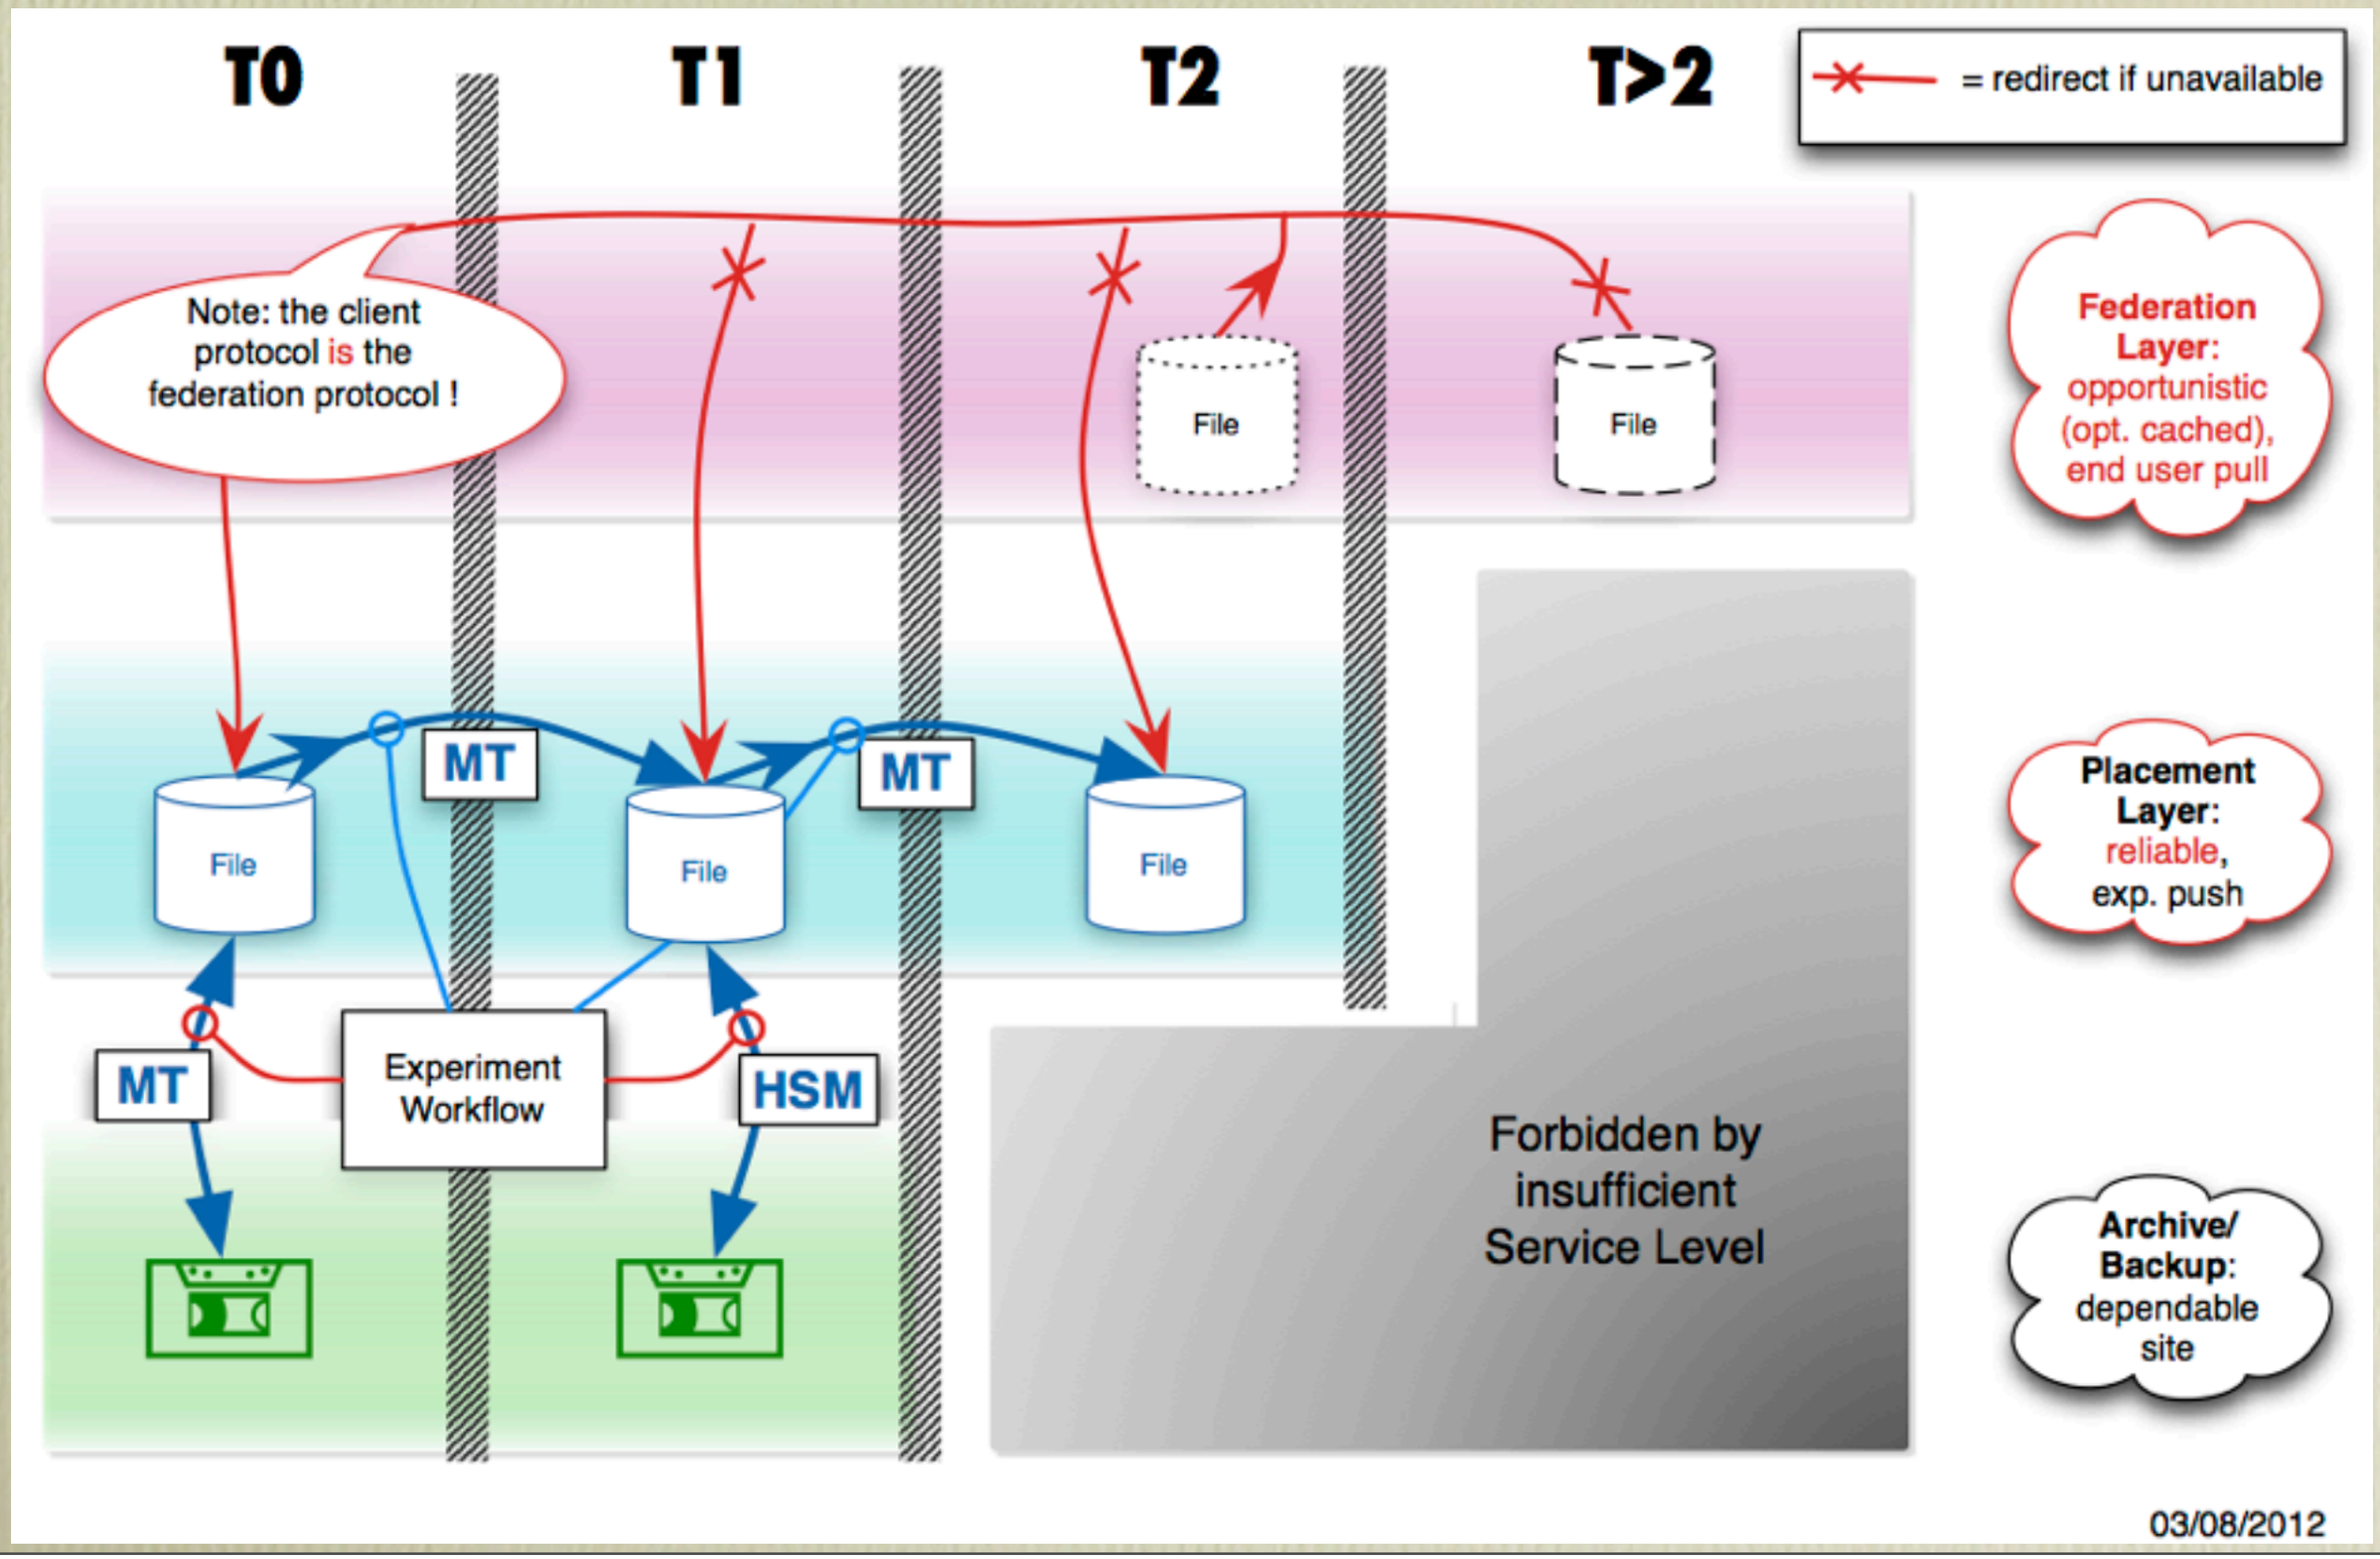
T0
T1
T2
T>2
= redirect if unavailable
Note: the client protocol is the federation protocol !
File
File
Federation Layer: opportunistic (opt. cached), end user pull
File
File
File
MT
MT
MT
Experiment Workflow
HSM
Forbidden by insufficient Service Level
Placement Layer: reliable, exp. push
Archive/ Backup: dependable site
03/08/2012

Storage Management

- Strongly endorse the principle that all storage is either an archive or a disk cache:
    - An archive may have high-latency for access to data and file loss is considered exceptional.
    - File loss in a cache is not-exceptional: files have a finite lifetime and a cache eviction policy.
        - Cache eviction policy may be something like LRU or when the disk is struck by lightning.

Storage Management

- Considering an architecture where file locality policy is a function of the endpoint.
  - The archive and disk would be represented as different endpoints instead of different SRM storage tokens.
- Allow the tape and disk systems to evolve at different speeds: take CERN EOS and CERN Castor as an example.

Storage Management

- Storage Management: SRM has several sets of functionality groups, not necessarily uniformly used by VOs or sites.
  - If groups want to replace parts of SRM, they should try to do complete functionality groups at a time. Compatibility level between sites is defined to be FTS.
  - I.e., WebDAV or GridFTP likely replace the LHC T2 use case for SRM. Nebraska is actively investigating this!

Transfer Protocols

- Preference is toward consolidating LAN access methods and direct ROOT access.
  - Likely, Xrootd and POSIX-based will be the most common in a few years. However, we can really do anything ROOT supports well.
- For WAN, dominate protocols are GridFTP and (for ALICE) Xrootd.
  - Support is being added for HTTP.
  - Idea is to encourage experiments to support what FTS supports.

Interest in HTTP

- The European Middleware Initiative (EMI) is making a sustained investment in providing HTTP support for its storage.
  - Adding 3rd-party-copy and scalable authz.
- Likely to encourage leveraging this work, as HTTP is ubiquitous, is less tied to any one project, and popular with users.

Storage Accounting

- WLCG has been looking for the “silver bullet” for storage accounting for quite some time.
  - Installed capacity currently is taken BDII (EGI) or OIM (OSG).
- In 2013, proposal is to provide this information via the upcoming “Storage Accounting Record” (StAR).

Security

- There's little interest in new security paradigms for data.
- Two pushes:
  - Properly prioritize security based on data value; make sure high-value data is secure before tackling low-value data.
  - Have a greater emphasis on auditing: think of the disgruntled admin or disgruntled user use case.

Stuff for Sites

- Storage as a Service: WLCG should open a better dialog about site requirements.
  - Currently, LHC experiments state requirements in terms of TB per site.
    - Extend requirements to include IOPS per core, MB/s per core, latency.
    - Have an honest expectation for file loss rate: don't pretend to require that files aren't lost.
- And, accordingly, make sure sites can measure these for supported storage systems.

Want to read more?

- Check out our final report, due in a few weeks.