



EUROPEAN MIDDLEWARE INITIATIVE

EMI and dCache: the vanguard of standards adoption

Paul Millar

What am I talking about?

- dCache in EMI
- Adoption of existing standards
- Aggressive implementation of emerging standards
- Involvement with standards creation.

dCache in EMI

dCache in EMI

- Reworking our dependencies
- Easier to deploy
 - Install is trivial:
 - No script to run, just install the RPM (and setting up a database)
 - Upgrading is trivial:
 - No script to run
 - Head-nodes may be updated independently from pool nodes.



dCache in EMI

- Easier to configure
 - More flexible system
 - Easier to understand
 - Defaults are separated
 - Empty configuration file is valid
 - Very positive feedback from early adopters:

“After spending a couple of hours with the new config files I wouldn't go back.”

Gerard Bernabeu; PIC, Tier-1 facility, at Göttingen dCache workshop 2011

dCache in EMI

- Core knowledge and code contribution is **wide spread**
FermiLab, NDGF, SNIC, DESY
- Multiple **inter-locking support networks**:
 - EGI/EMI support,
 - support@dcache.org,
 - Dedicated Tier-1 support meetings,
 - dCache support, funded by German government

Adoption of existing standards

Existing standards: GridFTP v2

- GFD.47 describes several extensions to **GridFTP**.
 - Most interesting is GET/PUT that allows **direct communication** between pools.
- dCache was **first** storage server to implement GridFTP v2.
- We also contributed code to add support in the Globus client libraries (in production and in use)

Existing standards: WebDAV

- WebDAV is **filesystem** access over **HTTP**
 - Widely supported already, in all client platforms (Windows, Mac, Linux, ...)
- dCache supports reading and writing with WebDAV.
- Supports X.509 communities.
- Also some support for non-X509
 - Read-only anonymous access now
 - username+password coming soon

Aggressive implementation of emerging standards

Emerging standards: NFSv4.1/pNFS

- NFS v4.1 / pNFS
 - supports client talking to pool
- dCache is **first publicly available** implementation
- Client is the Linux kernel, or Windows device driver, ..
- Mount dCache as a regular filesystem:
 - Support off-the-shelf software.
 - Security is via Kerberos.

Emerging standards: **GLUE 2.0**

- GLUE is an **info-model** for describing grid resources.
- Current version is **GLUE v1.3**
 - **Backwards compatible** with prev. versions
 - Somewhat **ad-hoc design**
 - Hard to publish what's deployed
 - SRM must be published twice,
 - Naming conventions used to tie objects.
 - Problems in representing cluster partitioning.
- Solution: **GLUE v2.0**

GLUE 2.0

- dCache was involved with creating GLUE 2.0:
 - abstract model,
 - LDAP rendering and XML rendering
- dCache supports GLUE 2.0 with EMI-1
- For more details, see **Elisabetta's talk**



Establishing new standards

Setting the scene

- **SRM**
 - for managing storage
- **GSI**
 - V similar to **SSLv3** but incompatible
 - Clients can delegate
- **Delegation**
 - Allows **SRM** server to act “on behalf of” end-users
 - Creates key-pair on server



Delegation Services

Name	Tech.	C / Native		Java / JVM		Supported
		Client	Server	Client	Server	
Globus Credential Delegation Service	SOAP	Y	Y	N	N	N
GridSite Delegation Service	SOAP	Y	Y	Y	Y	Y
Globus New Delegation Service	REST	N	N	N	N	Y
IVOA Delegation Service	REST	N	N	Y	Y	M

Taking it further

- GDS is a “standard” only within EMI.
 - Not endorsed by any standards body.
- No obviously applicable standard
- Looking at starting up an **OGF WG**
 - Suggest writing up GDS as an experience report
 - **BoF session** at OGF-31 meeting was encouraging and supportive.
 - Next step is to **establish the WG.**



Establishing standards: StAR

- Currently there's no standard for storage **accounting information**
 - There is for CPU usage.
- EMI has established a group to build this, dCache is involved.
- Working within the OGF group on accounting.
- EMI's Jon Nilsen is *nearly* co-chair.
- For more details, see **Jon's talk**

Conclusions

- dCache is committed to standards
 - Want to be **competitive to industry**
 - Attract new, **non-HEP communities**
- Implemented many standards:
FTP, GridFTP, GridFTPv2, SRM v1.1, SRM v2.2, WebDAV, NFS v4.1/pNFS, GLUE 1.3, GLUE 2.0, ...
- Through membership of EMI, we can push for more standardisation.



Thank you!

EMI is partially funded by the European Commission under Grant Agreement RI-261611

GLUE 2.0: the standard

- Breaks backwards compatibility
 - More structured, hierarchical approach.
- Multiple documents
 - Technology agnostic, general document
 - Multiple documents that map generic document to specific technologies
- General document finished 2009
- At least two renditions:
 - LDAP (completed), XML (WiP)

GLUE 2.0: now (+/- a bit)

- EMI-1 includes support for GLUE 2.0
 - LDAP schema support
 - Information published as v1.3 and v2.0
 - Publish different versions side-by-side.
 - Available v2.0 information similar to v1.3
 - Clients use v1.3 initially.
 - With EMI-2, clients will use GLUE 2.0.
- QA tools for validating GLUE v2.0
(LDAP Schemata are quite primitive)

GLUE 2.0: in the future

- Clients move from using v1.3 to use v2.0
 - Coming with EMI-2
- Enrich publishing of GLUE 2.0 information
- Drop support for GLUE 1.3

Why delegate?

- 3rd party copy

srmCopy

- Reserving bandwidth

srmPrepareTo(Get|Put),
srm(Get|Put)Done, srmCopy

- X.509 backed tape

SrmBringOnline, srmPrepareTo(Get|Put),
srm(Get|Put)Done, srmCopy

- Federated SRM

Everything except srmPing.



Why switch from GSI to SSL?

- GSI is **not a standard** (SSL is)
- Coupling delegation with transport negotiation is **inflexible**.
- It's **not widely used** outside of Grid
- Only libraries are coming from a **single vendor**: Globus
- Hard to add advanced features; e.g., **no hardware acceleration**.

What to do about delegation?

- If delegation isn't needed then pure SSL works fine.
- Delegation requires extra functionality
- Solution is a “service” that allows delegation.



(NB. we're **not** talking about a single, per-site shared service; rather, each service has a common extra API)

Introducing the winner

- GDS is a **de facto standard**.
 - Developed ~2005 by Andrew McNab for web management software.
 - Adopted by gLite, after going through a review process
- Current version is v2.0.0
 - In production (FTS, GridSite, ...)
- **Two** independently developed libraries (Java and C), both provide **client** and **server**.

GDS: an EMI standard

- Other services in EMI also need to delegate
 - EMI ES (execution service), FTS, ...
- Agreement to use GDS within EMI.
- Current API docs need tidying up:
 - Conflates documenting software with documenting the standard,
 - Leaves some things too vague.
- Some work underway in this area