WLCG – Worldwide LHC Computing Grid

SRM v2.2 Status

LHCC Referees meeting,
September 24th 2007

Jamie Shiers
WLCG Service Coordination

Agenda

- WLCG Management Follow-up on Production Deployment of SRM v2.2 services
- Very brief reminder of the problem
- Status of the Deployment Schedule

WLCG Management Follow-up

- Following many discussions and a number of attempts, there are now weekly conference calls to follow the Production Deployment of SRM v2.2 services for the WLCG.
- And only this.
- We will continue as long as necessary – presumably until end 2007 at least
- At that stage the standard mechanisms for handling any residual bugs / service problems should suffice

Members

- Jean-Philippe Baud - IT/GD <Jean-Philippe.Baud@cern.ch>
  Tony Cass - IT/FIO <Tony.Cass@cern.ch>
  Flavia Donno - IT/GD <Flavia.Donno@cern.ch>
  Patrick Fuhrmann - IT/LCG <Patrick.Fuhrmann@cern.ch>
  John Gordon - IT/EGE <j.c.gordon@rl.ac.uk>
  Volker Guelzow - IT/LCG <volker.guelzow@desy.de>
  Harry Renshall - IT/GD <Harry.Renshall@cern.ch>
  Jamie Shiers - IT/GD <Jamie.Shiers@cern.ch>
  Gene Oleynik <oleynik@fnal.gov>

- Web archive: here

Study of SRM 2.2 specification

- In September 2006 *very different interpretations* of the spec
- *6 Releases of the SRM v2.2* specification document: July, September, December 2006 and January(2x), April 2007
- Study of the spec (*state/activity diagrams*): many unspecified behaviours.
- A list of about *50 open issues* has been compiled in *September 2006*.
- Last *30 points discussed* and agreed during the *WLCG Workshop* in *January 2007*. Other major points delayed to SRM 3.0.

- *The study of the specifications*, the discussions and testing of the open issues have helped *insure coherence* in the protocol definition and *consistency between SRM implementations*.

https://twiki.cern.ch/twiki/bin/view/SRMDev/IssuesInTheSpecifications

SRM v2.2 implementations

- *CASTOR2* : developed by CERN and RAL. **Used at ASGC, CERN, CNAF, RAL**
  SRM v2.2 support in v.2.1.4.
- *dCache* : developed by DESY and FNAL. **Used at other WLCG Tier1s & ~50 Tier2s (+)**
  SRM 2.2 support in v1.8.

- *DPM* : developed by CERN. **Used at ~100 sites**
  SRM v2.2 support in v1.6.5 in production.
- *StoRM* : developed by INFN and ICTP.
  SRM v2.2 interface for many filesystems: GPFS, Lustre, XFS and POSIX generic filesystem.
  SRM v2.2 support in v1.3.15.
- *BeStMan* : developed by LBNL.
  SRM v2.2 support in v2.2.0.0.

Top 5 Issues (MB – May/June)

Experiment	ALICE	ATLAS	CMS	LHCb
Issue #1	xrootd-CASTOR2	CASTOR@CERN	CASTOR: functionality & performance	Data Access from T1 MSS
Issue #2	xrootd-DPM	Integration of DDM/FTS/SRM/ GridFTP/LFC etc	SRM I/F with functionality & performance	glexec usage
Issue #3	FTS Service	(Lack of) SRM 2.2	FTS Service	File management
Issue #4	gLite WMS	Data Storage Management Tools	Workload management	Deployment procedure
Issue #5	VOMS	Stability of the Information System	Information system	Information system

The Deployment Schedule

- Named sites to be ready in 2 phases for experiment testing (mid / end July);
  - This phase went well – only very minor delays of 1-2 days in a few cases;
- Verify sites correctly configured;
  - Numerous problems encountered, partly due to lack of adequate documentation on new features
- Experiment testing (LHCb, ATLAS, CMS);
  - Testing foreseen in August delayed.
  - Now being actively (and positively) pursued.
- Upgrade named sites in agreed order.
  - Delays of a few weeks since plan agreed in July

➢ **We need to be more realistic about manpower that can work on these issues not only during vacation / conference periods but also at 'normal' times – everyone has multiple responsibilities**

Component Service Readiness

Service	Comments
SL4	WN & UI delivered. Other components will be ported Q4 2007+
LFC	Stable & in production. Additional bulk methods (ATLAS) delivered and being tested.
FTS	**Will only run at CERN and (most) T1s. Most have deployed it already, the rest soon (< end September). Still some issues related to experiments' CMs to be worked out.**
VOMS	Server and management interfaces work Still issues over how proxies, roles, groups, attributes will be used – being analysed
WLM	Working solution(?), but changing…
3D	**Tier 0 and Tier 1 database infrastructure in place and streams replication used by ATLAS, CMS (T0 online-offline) and LHCb; 3D monitoring integrated with experiment dashboards**
SRM v2.2	**Extensive testing since ~1 year. Production schedule being developed. (CERN+major T1s < end 2007; rest < end Feb 2008.)**

SRM v2.2 – GDB Review

- All implementations delivered for testing
- Bugs, inconsistencies, configuration issues
- Critical issues identified
- Experiment testing one month late
- Good progress but not there yet.

- *This was the status < 1 month ago (end August) – much of the detail has since changed (for the better…)* **(no more critical issues!)**
- *This is clearly a result of the extensive testing and focussed effort on resolving the issues*
- *The WLCG Collaboration Workshop in Victoria, together with discussions during the CHEP conference helped resolve further issues*

- *We are now agreeing concrete production deployment dates by named site*
- *Experiment testing is still required – as many problems as possible should be fixed before production – some will only be found and fixed later*

j.c.gordon@rl.ac.uk

SRM v2.2 Production Deployment

- Details of SRM v2.2 production at CERN now being finalised. Plan is for one 'endpoint' per LHC experiment, plus a public one for the rest (as for CASTOR2).
  - Target < end October 2007
  - Tier1s running CASTOR – wait at least one month after CERN
- SRM v2.2 will be deployed at FZK during the week of November 5 with experts on site. Other major Tier1s (and some Tier2s – e.g. DESY) will follow up until end 2007.
- Remaining sites – including those that source dCache through OSG – will be upgraded by end-Feb 2008.
- DPM is already available for Tier2s; STORM also for INFN(+) sites
- CCRC'08 (Feb) is foreseen to run on SRM v2.2

Draft dCache Upgrade Schedule

- Oct 29 : NDGF (for sure)
- Nov 5 : gridKa (for sure), SARA (likely)
- Nov 12 : Edinburgh workshop (for sure)
- Nov 19 : (nothing yet)
- Nov 26 : IN2P3 (likely)
- Dec 3 : (nothing yet)
- Dec 10 : BNL (rather sure)
- Dec 17 : PIC, RAL (they actually said : end of the year)
- Fermi : December to January (Jon needs a special feature in dCache)
- TRIUMF : no reply yet.

SRM clients: FTS, lcg_util/GFAL

- lcg_util / GFAL testing done
  - Various bugs / issues reported, now fixed in latest release
- Tests on pilot FTS service
  - Both dteam tests and experiment tests: wide-scale testing to discover issues early
  - Last release of FTS 2.0 fixes known integration issues
- Many SRM issues identified and resolved
  - Specification conformity issues reported to SRM providers and solved
- Major work is now done w.r.t. client integration
  - We do anticipate there will be more integration / conformity issues as we ramp up production, but these will be solved as they occur

DPM

- SRM 2.2 been deployed in production since January 2007
  - Not all T2 sites have yet upgraded
- SRM 2.2. conformance tests and stress-tests done
  - All issues resolved
  - DPM-managed SRM copy functionality has still to be provided

Number of calls per week

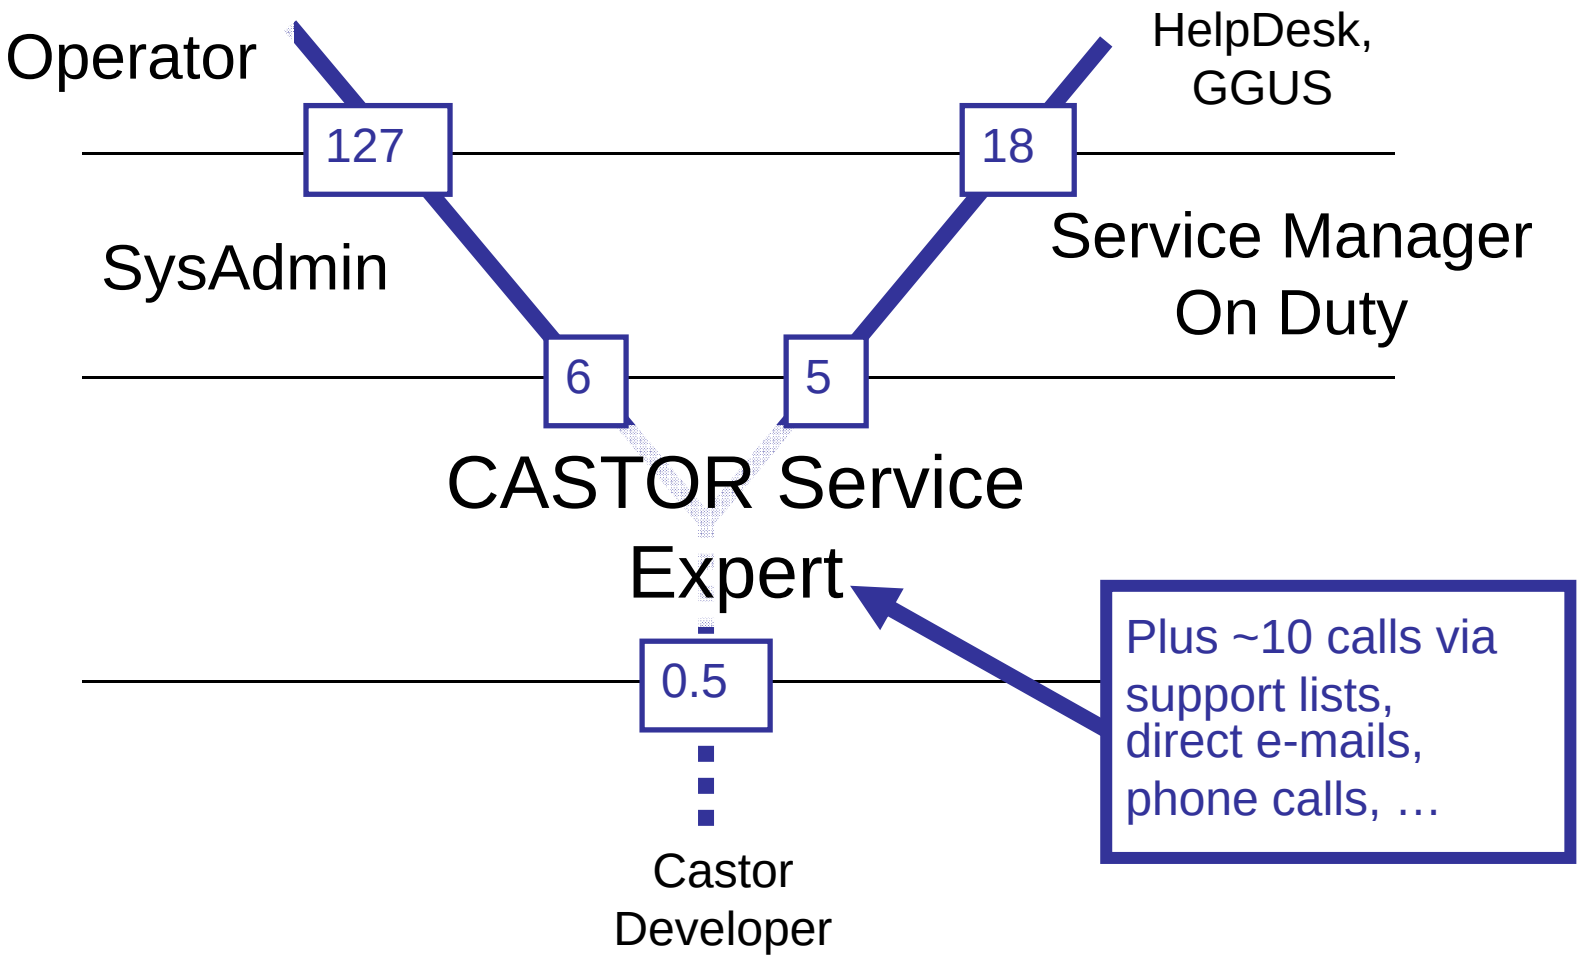

SRM V2.2/dCache 1.8

- All US LHC Tier-1 and Tier-2s use dCache implementation for grid accessible managed storage.
- SRM V1/dCache 1.7 distributed through VDT (as RPMs) since 03/07 (currently dCache 1.7.0-42).
- Test stand for functionality and performance testing in place for SRM V2.2/dCache 1.8.
- OSG UCSD site participating in Grid Storage Services Deployment (GSSD) group and testing with CMS Phedex.
- OSG BNL Tier-1 participating in and contributing to GSSD production tests.
- OSG will deploy separate instance of GSSD testing for local use and contribution to production deployment.
- SRM V2.2/dCache 1.8 scheduled for deployment as part of OSG 1.0 release in Feb '08. We recognize this is late for the WLCG; Discussions are underway to address this.

Funding for storage and data movement support for WLCG

- dCache contributions at Fermilab funded by Fermilab Computing Division, US-CMS and OSG.
- SciDAC CEDPS project includes extensions to data movement (GridFTP), enhancements to integration of GridFTP and dCache.
- DOE funded LambdaStation provides enhanced features for integration of dCache with managed wide area networks.

Organization

- Gene Oleynik - Storage Section Head, Overall leader of Fermilab dCache Contributions.
- Timur Perelmutov - dCache/SRM at Fermilab technical project lead.
  - Timur Perelmutov - SRM
  - Dmitry Litvintsev (50%) - SRM
  - Alex Kulyavtsev - Resilient Manager
  - Vladimir Podstavkov - Resilient Manager
- Expect to hire another developer this coming year.
- Ted Hesselroth - OSG Storage Middleware

Support: WLCG Support expected to follow this model

- We operate within a 3 tiered support model:
  - Level 1 - Local site expertise. Install hardware, configure and maintain systems, first level of troubleshooting
  - Level 2 - More expertise; advise, assist and troubleshoot installations, configuration issues and transfer problems
  - Level 3 support - Developer expertise, identify and fix bugs
- Fermilab developers provide all three levels of support (with help from operational groups) for the Fermilab dCache systems at CDF, CMS, and our public dCache system (MINOS MiniBooNe, etc.)
- Our expectations are
  - L1 expertise is available at and provided by the deployment sites
  - L2 support is funded the by external stakeholder organizations (e.g. OSG storage activities staff, WLCG/EGEE funded support staff)
  - Developers provide L3 support in a steady state situation (ie effort limited).

FTEs in Support

- Fermilab core developers (total 3.5 FTEs) spend 1.5 FTE in local and global support.
- US CMS Tier-1 Facility provides Level 1 and Level 2 support of 2 FTEs.
- US CMS Tier-2 Sites provide local and community support at 1 FTE per site.
- OSG Level 2 and Extensions staff is currently 2.75 FTE (reevaluated annually).

Current Issues and Action items in Fermilab dCache Contributions & Delivery for the WLCG

- We recognize the software schedule has slipped. For Fermilab the reasons are various and mainly related to the current large install dCache base for our running experiments.
- A major cause has been continuing instability in the CDF dCache 1.7 system since its upgrade in March '07. This is critical for CDF data analysis. Good news this is quickly winding down.
- We are adding to our dCache contributions for WLCG delivery and schedule:
  - Dmitry will be 100% on SRM starting next week (When he is back from vacation)
  - Vladimir has joined the resilient dCache effort
  - US CMS are increasing their contributions through the hire another developer in the next few months.
- We are working with OSG to re-align the priorities and increase contributions through the GSSD.

SRM Collaboration Support

- GridPP / Edinburgh already has a large knowledge base for supporting GridPP sites
- It has agreed that this can be extended to provide 'WLCG Collaboration Support'
- An SRM workshop in 2nd week November will contribute to this
- Deliverable: provide sufficient documentation for other sites to install / configure dCache & DPM
- Open to GridPP sites & other priority sites as indicated by experiments (cannot be all…)

SRM v2.2 - Summary

- The ‘SRM saga’ has clearly been much longer than desirable – and at times fraught
- Special thanks are due to all those involved, for their hard-work over an extended period and (in anticipation) for the delivery of successful production services
- When the time is right, it would make sense to learn from this exercise…
- Large-scale collaboration is part of our world:
  - What did we do well? What could we improve?