



WLCG – Worldwide LHC Computing Grid

Tier 1 Feb CCRC'08 Site Reviews

WLCG Workshop
22 April 2008

Harry Renshall
WLCG Grid Support





Agenda

- Sites were asked to briefly summarise their experiences of the February run of CCRC'08 highlighting what went well, what went badly, any surprises, any suggestions to the Tier 0 and WLCG and what is planned to be done differently for the May full scale dress rehearsal (given that February was more for functional tests on as full a scale as possible) . With thanks we have short reports from 4 Tier 1:
 - ASGC (9% of ATLAS T1 resources, 14% of CMS T1 resources) from Min Hong Tsai
 - FZK (25% of ALICE T1 resources, 10% of ATLAS T1 resources, 12% of CMS T1 resources, 9% of LHCb T1 resources) from Andreas Heiss
 - PIC (5% of ATLAS T1 resources, 6% of CMS T1 resources, 6% of LHCb T1 resources) from Gonzalo Merino
 - TRIUMF (5% of ATLAS T1 resources) from Reda Tafirout
- And two longer reports to be presented by themselves
 - RAL (2% of ALICE T1 resources, 16% of ATLAS T1 resources, 8% of CMS T1 resources, 35% of LHCb T1 resources) from Derek Ross
 - US-CMS (FNAL providing 40% of CMS T1 resources) from Ian Fisk



ASGC Report (1/4)

- Preparation

Refactor LAN network:

- * Upgrade core switch with 10G interfaces
- * Installed additional edge switches and replace multiple 1G trunked links with 10G links. We observed that trunked 1G links did not provide good load balancing characteristics.
- * Blade disk servers network switches with 1G uplinks for multiple servers caused bottlenecks. Replaced aggregated uplinks with dedicated 1G links for each disk server.

Resources: Installation of CPU and storage resources procured in late 2007

Grid Services

- * Additional CE added to distribute load
- * LFC mysql migrated to oracle
- * Upgrade grid services to meet baseline versions

Castor preparation

- * Install SRM2 and configured required storage classes
- * Castor repack deployment

WAN upgrade: April 4th - ASGC AMS-CERN upgraded from 5 to 10G

Facilities: March 21th - 2MW generator installed to provide backup power



ASGC Report (2/4)

Went Well:

ATLAS:

- * ATLAS provided clear milestone information early on and status updates which was helpful for sites in planning for CCRC.

CMS:

- * T0-ASGC transfer for 9 days above 13MB/s target with peak daily transfer at 87MB/s

- * T1-T1 transfer for 11 days above 24 MB/s target

- * CMS T1 prestaging: 13.2 TB, 5632 files, 360 tapes, 18 min stage request time

- * total stage time: 22 hours, tape to buffer rate ~150MB/s

Went badly:

CMS

- * T2 transfers succeeded with ASGC to TW, however failed with KNU and TIFR in February

- * KNU achieved better results 35MB/s (ASGC to KNU Mar 25) and 20MB/s (KNU to ASGC Apr 1)

- * ASGC-TIFR obtained steady but low transfer rates of 2.5MB/s on Mar 26

- * TIFR are currently limited by a 34Mbps link connection



ASGC Report (3/4)

Surprises:

Encountered DB table lock issues for Castor SRM

- * experienced high DB loading and table locks affecting data transfers
- * separated SRM and Castor tables into separate DBs so they don't affect each other.
- * added index to SRM tables as recommended by CERN to improve performance
- * loading has dropped but some long lasting table locks still exists which are regularly cleaned.

Castor upgrade to from 2.1.4 to 2.1.6 - March 19 (after the Feb run)

- * less out of memory disk server crashes
- * heavier CPU loading on disk servers observed
- * IO error on import disk pools
 - * maybe related to incompatibility of FC cards between server and arrays
 - * working with vendor to resolve this issue
- * GridFTP2 only supports one service class per disk server
 - * ASGC castor was configured with multiple service classes sharing a single disk pool. This configuration is currently not possible.



ASGC Report (4/4)

Suggestions

Automated SRM endpoint updates for FTS channels

- * This is currently a manual process and some sites do not update the endpoint information promptly causing transfers to fail (a script to refresh the FTS endpoint cache is now distributed).

ATLAS: Explicit and uniform information about cleanup of dataset at sites:

- * what part are done centrally and what are on site side.
- * Twiki pages for sites to follow in CCRC in experimental issues.

Plans for May

- * Lack of disk server with FC cards: need to migrate services and free servers to act as disk servers
- * Address disk array controller thermal issues: working with facility engineers to cool disk controllers located at the back of array in hot aisle
- * Discuss with Experiment support and Castor team on how to best configure Castor to meet CCRC objectives
- * Add additional RAC nodes for castor DB
- * investigate DNS HA solutions grid services
- * Improve number and quality of recovery procedures for on-call engineers for 24 by 7 operations
- * upgrade HK and JP network connectivity to 2.5 Gbps in June (this also improves connectivity to China, Korea)

CCRC'08 @ GridKa

Steinbuch Centre for Computing



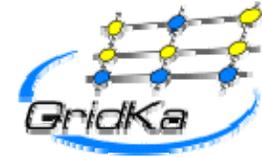
Forschungszentrum Karlsruhe
in der Helmholtz-Gemeinschaft



Universität Karlsruhe (TH)
Research University · founded 1825



Problems and issues *before* CCRC08



- **Requirements of the experiments,
e.g. required space tokens and sizes**
 - late (it's no problem to setup a space token ...
... if you have the disk space available!)
 - different information sources (wiki, slides, meetings, ...)
 - different levels of detail

Problems and issues *during* CCRC08

- **FTS**
 - base line release very late.
 - proxy renewal problem reappearing.
- **GridKa storage system**
 - initially no increase of size of existing spaces possible.
(later fixed by dCache patch)
 - disk controller failure affected CMS write pools and limited data import rate.
 - many (~ 5) dCache updates/patches necessary.
- **Would like more detailed feedback from experiments**
 - what are they doing, what kind of jobs are running etc.
- **Very useful information on Atlas activities from Kors Bos during Feb. run (“about CCRC” emails)**

GridKa improvements for May CCRC08

- **Storage system**
 - dCache upgraded to version 1.8.0-14
 - added more pools
 - improved TSS interface to TSM tape system
- **LFC**
 - moved backend from MySQL to Oracle RAC (3 nodes)
 - new, more powerful frontend
- **Additional CEs and BDIs** to cope with higher load caused by GridKa (and Tier-2's) higher # of CPUs

Issues from PIC



- We did not see a constant and “predictable” load of jobs on the CPUs that we could use to learn which are the needs for the farm in terms of bandwidth per job slot, etc.
 - The bandwidth per job slot to the storage (actually to each Storage Class) has important deployment (hardware!) implications. We do need realistic tests or estimations of these.
- The SRMv2.2 classes and tokens did work quite well. Most of the experiments set clear specifications for space under each token/class for the Feb tests. OK.
- Now we have learnt the impact of configuring static reservations in the storage, I see more than ever the need of as clear as possible requirements from the experiments (last chance ☺)
 - [Space and bandwidth per token](#) (*follow link in .ppt*)
 - [Megatable!](#) (key parameters as the input rate CERN→T1 tape/disk are still not clear to us)

Issues from PIC



- Publication of the SRMv2.2 in the Information System:
 - Available and Used space: We are currently publishing information of TxD1 space. Is there a consistent and commonly agreed way for publishing T1D0? Disk buffer size and space on tape?
- dCache configuration for SRMv2.2:
 - In our prod configuration (still 1.8.0-12p6) we are effectively forced to associate a Path to a Storage Class. Isn't this contradictory with the original idea to keep them orthogonal?
- Deleted files on tape:
 - CMS has deleted a lot of the data it had stored on tape since Oct-07. If the files that go to tape are not grouped in “tape families” a priori, this forces us to “repack” which is a very (human and drive) resources demanding operation. Massive deletions of files in the “default” tape family should be discouraged.



TRIUMF Report (1/2)

- SRM v2:
The SRM v2 configuration and space tokens definition was implemented "just in time" for Feb. 1st. The SRM v2 was deployed and tested few days earlier on a virtual testbed then ported to our production system without serious problems. We had a few issues in publishing the space tokens properly in BDII, but we managed to fix it.

DISK:

ATLAS is managing centrally datasets deletions. This was exercised during CCRC'08 phase 1 and seems to work. The only issue is that there is no call back to the HSM system when files are deleted from the tape end point.

ATLAS is now more organized with respect to datasets naming convention and directory structures. This is important for dCache sites where directories are tagged for specific things (to distinguish between permanent disk and tape buffer space being one example). At TRIUMF we separate disk pools and HSM pools traffics.



TRIUMF Report (2/2)

- **TAPE:**

Our tape writing infrastructure is able to write specific datasets (based on directory paths in PNFS/dCache) to specific range of tapes. This way tape recycling is easier. We effectively separated CCRC data from normal production data. At one point we had a glitch in our tape system, we fixed it quickly and data was flushed to tapes without problems. Our tape disk buffer space was set to 8 TB to handle this kind of scenario and does not affect ATLAS sending us data.
- **Overall:**

From TRIUMF point of view, things went well I think and ATLAS tests went pretty smoothly when writing to both our disk and tape end points.