

The Worldwide LHC Computing Grid

Report from the Jan 22 - 26
WLCG Collaboration week

GDB, February 7th 2007



Worldwide LHC Computing Grid
Distributed Production Environment for Physics data Processing



Overview

- Dates were initially chosen at the June 2006 Tier2 workshop - avoiding clashes isn't easy...
 - Initially 3 days, but by July GDB had already expanded to a full week including tutorials!
 - *Then more and more things attached - dCache workshop (later moved to DESY in previous week, (preceded by virtualisation workshop)), joint operations workshop, VOM(R)S...*
- Close to 2 full weeks all in all!



Previous Workshop Gripes

- Why no visit to e.g. ATLAS?
- Why no introduction to particle physics?
- These things could clearly have been arranged
- *Why no suggestion in the meeting Wiki?*
- Why no name badges? (We had CERN access cards, but not for 'locals'...)
- Start time (11:00 on Mon&Tue) (dictated by room availability)
- Better coffee, air conditioning, ...
- More involvement of Tier1s and Tier2s





Agenda

- Every single suggestion received by no matter what means since last June was pursued
- Only one didn't make it!
 - ✘ Report from Virtualisation workshop
 - ☹ No response from the suggestor when contacted...
- Clearly very much driven by WLCG Commissioning schedule as established at September 2006 LHCC Referees' Review...

WLCG Commissioning Schedule

experiments

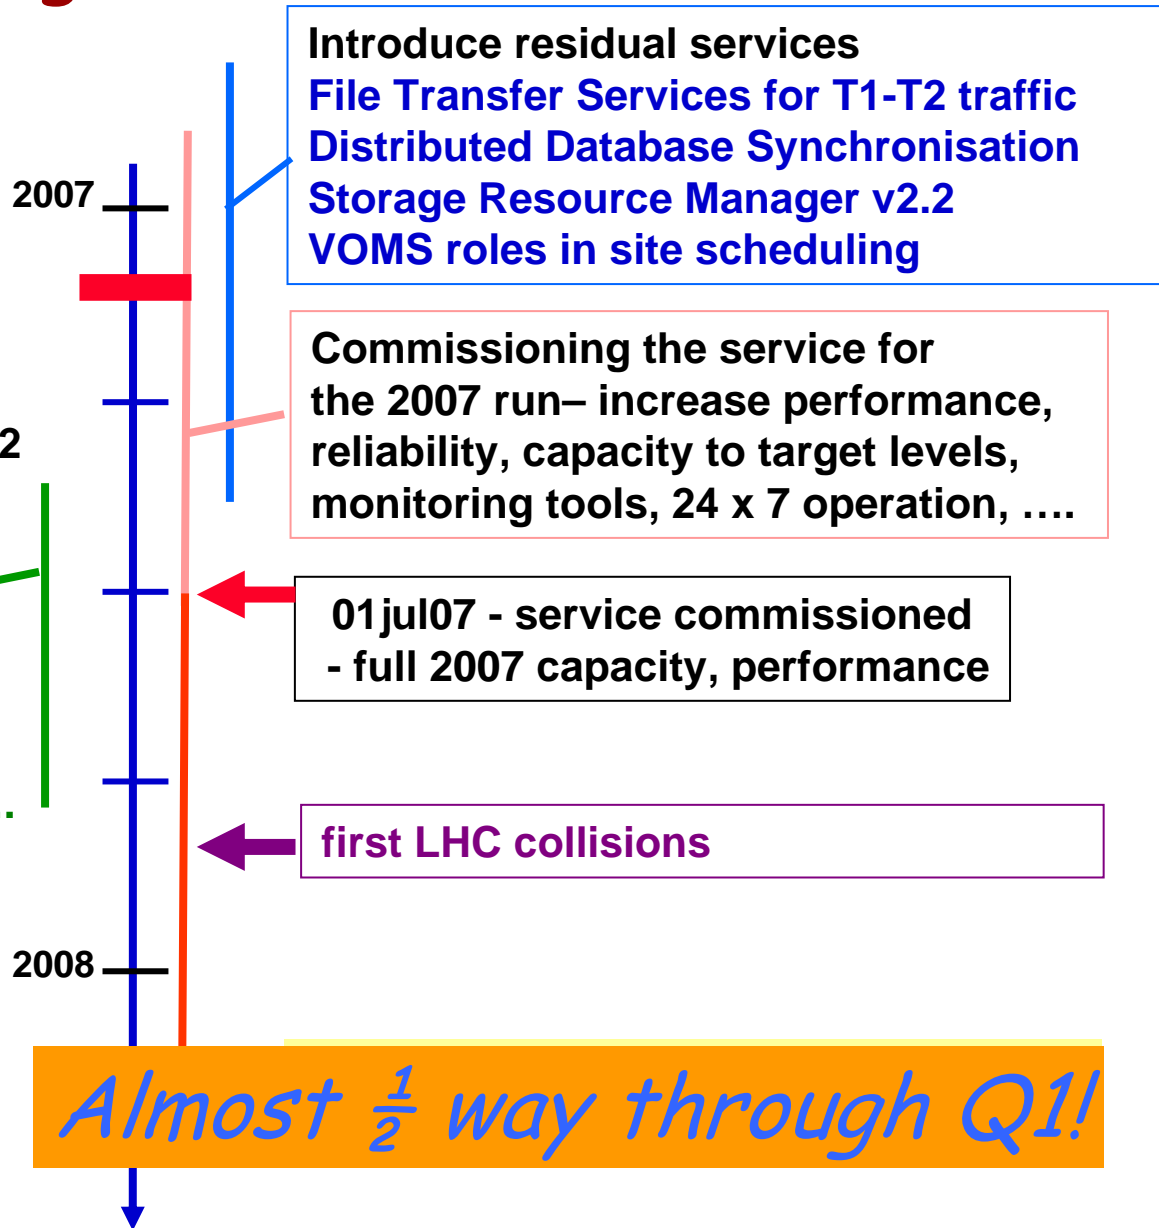
- Continued testing of computing models, basic services
- Testing the full data flow
DAQ → Tier-0 → Tier-1 ↔ Tier-2
- Building up end-user analysis support

Dress Rehearsals

- Exercising the computing systems, ramping up job rates, data management performance,



services





Workshop Participation

- 270 people registered (almost) by deadline
- 15 others missed it... (where have they been living?) (I know...)
- **Walk-ins actually represent a significant overhead**
- How many people actually came?
 - IP addresses: 120 too few, 240 OK (but unstable...)
 - Coffee: ordered for 270 - ran out
 - BOF sessions: >200 in parallel BOFs Monday evening
- Certainly **more** than foreseen...
 - My guess was 180: Sophie & Gavin both bet on 197...
- We had a single registration for all events during the week, so 270 was the **OR** of all sessions...
- Anything else (single registration) unmanageable (IMHO)...



Workshop Results

- All talks available on the agenda page:
 - <http://indico.cern.ch/conferenceDisplay.py?confId=3738>
- As are videos of (most) sessions
 - Some short sessions glued together
- Notes from the panel discussions
 - Some still to come!
- Also the results of Michel Jouvin's survey...
- ☺ **Most people seemed quite happy(?)**
- Some conclusions from the various BOF & experiment sessions follow...



Site / Region Reports

- Suggestions from UKI ROC mailed to all speakers, as well as proposal that the GridPP deployment talk (GridPPnn) is used as an example:
- *"Regarding the T1 reports with information on associated T2s - the emphasis should be on problems encountered and solutions implemented.*
- *The talks should be limited to a single slide to describe the regional makeup - i.e. how many sites, experiments supported, CPUs and storage deployed etc. Descriptions of hardware purchased and reasons for those purchases might be of interest but again a single slide. You mentioned metrics ... For example amount of scheduled downtime/time unavailability with explanations of major gaps and not just a lot of numbers. Speakers should be encouraged to raise questions which have been put forward by sites in the region and stop to discuss possible answers rather than just skipping over the question(s). These might be left to towards the end of the talks. In putting the emphasis on problems encountered some care will be needed to stop the session becoming a moaning forum! There are lots of interesting approaches to organising these sessions. You could imagine for example that each speaker has to give the 3 biggest successes in the T1 region and the 3 biggest problems."*

➤ Question for the future - should we (GDB/MB/someone else) propose a template? Or are 'suggestions' good enough?

Brief Summary of BOF sessions

Longer Summaries can also be found on the
W/S agenda page





dCache Workshop

- Not part of the WLCG week in the end but...
- **Significant +ve change since first dCache workshop held at DESY September 2005**
- ~70 participants (x 2), Tier1 and Tier2 sites, extensive experience in running dCache services, plus developers
- Focus on SRM v2.2 & dCache 1.8 release

<https://indico.desy.de/conferenceDisplay.py?confId=138>

Version 1.8 is the first Version featuring the SRM 2.2 protocol.

A 'beta' version will be available early April.





Data Management Roadmap

- Very well attended session - IT amphitheatre was completely stowed out
- 🌶️ Confirms Data Management as 'the **HOT** topic'
- Some discussion (again) about priorities (TCG)
- **"Dark Storage"** - you heard it here first!
- 😊 **LFC bulk methods & FTS 2.0 shortly in PRO**
- ☹️ **SRM v2.2 DELAYED - ASK AGENT** - see BOF
- 💣 **Stability of storage services at sites #1 issue**



DB BOF & Mini-workshop

- Will be covered in next talk...
- Agenda pages:

<http://indico.cern.ch/sessionDisplay.py?sessionId=23&slotId=0&confId=3738#2007-01-22>

<http://indico.cern.ch/conferenceDisplay.py?confId=10132>



The WLCG VOM(R)S BOF

<http://indico.cern.ch/sessionDisplay.py?sessionId=19&slotId=0&confId=3738#2007-01-22>

☺ A success in terms of:

- Participation
 - Several VOs: LHC Experiments, AUGER, VOCE and other authorities: TCG, JSPG, NDGF
- Communication between voms/vomrs developers in view of any new feature proposals or changes in api.

• Outcome:

- Generic Attributes (GAs) are needed for ATLAS and LHCb.
- VO managers must provide written requirements in <http://littleblue.cnaf.infn.it/twiki/bin/view/VOMS/VomsVomrsRequirements>
- **Deadline for submitting requirements is MID FEBRUARY 2007.**
- GAs should be handled at the level of VO-specific software and will NOT be interpreted by the grid middleware.



DPM BOF

Input from:

- site administrators (survey + BOF session)
- experiments

6 major items identified (in order of priority):

- RFIO
- support for secondary groups
- administrative and monitoring tools
- ACLs on pools
- quotas
- better scheduler

☺ RFIO: wrong dependency found - fix in progress!

- About 15 people, from ROCs, VOs, OMC, GGUS.
- Presentations by: GGUS, ALICE, ATLAS, CMS.
- People in general seem happy about the new release process and the new features implemented
- **We need to cater for and focus on the VO needs:** for instance CMS is thinking about the integration of the CRAB support mailing list into GGUS.
- The most important and recurrent complaint at the moment is the long time it often takes to solve problems
 - this is mainly a resource issue -> more dedicated supporters are needed
 - it makes people look for (and find) ways to bypass the official support channels
 - **We need to address the lack of manpower behind the system so that the time to solve a problem becomes shorter.**



Monitoring BOF

Aims

- Communicate the goals and structure of the new groups
- Describe the present state and approach
- Seek feedback and further community contribution

BOF


- Well attended with 80+ CBLFs
- ☺ There were useful questions and discussions
- Feedback confirmed the proposed workplans
- ☺ Some new volunteers came forward
- Some new leads for the co-ordinators to follow-up

- The status and functionalities of **WLCG SRM 2.2** implementations have been presented
 1. **Attendance: about 40 people, SRM 2.2 developers, MSS developers, experiment representatives, site representatives.**
 2. **Good participation. Many questions about the SRM implementation by the experiments. The answers were considered satisfactory.**
- Question on space reservation in dCache for Tape1Disk0
- ✓ SRM 2.2 open issues discussed and agreed
- ✓ Aggressive testing plan proposed and agreed
 - **Basic functionality tests passed by Jan 2007**
 - **Use cases and interoperability tests by Feb 2007**
 - **Stress tests to be performed till end of March 2007**
- New GSSD group introduced to the SRM 2.2 developers



SRM v2.2 - Summary


- Major effort in recent weeks in improving test results
- But we know from experience that moving from 'test' under light load to 'production' under full user load takes time & effort
- Will be followed closely on weekly (phone) and monthly (face-to-face) basis
- Production deployment issues must be carefully addressed
 - Migration strategy, existing data, catalogs etc etc



Summary of experiment sessions

Only ALICE held tutorials *per se* - probably a name we should drop. **All** experiments held sessions which were well attended and judged useful.





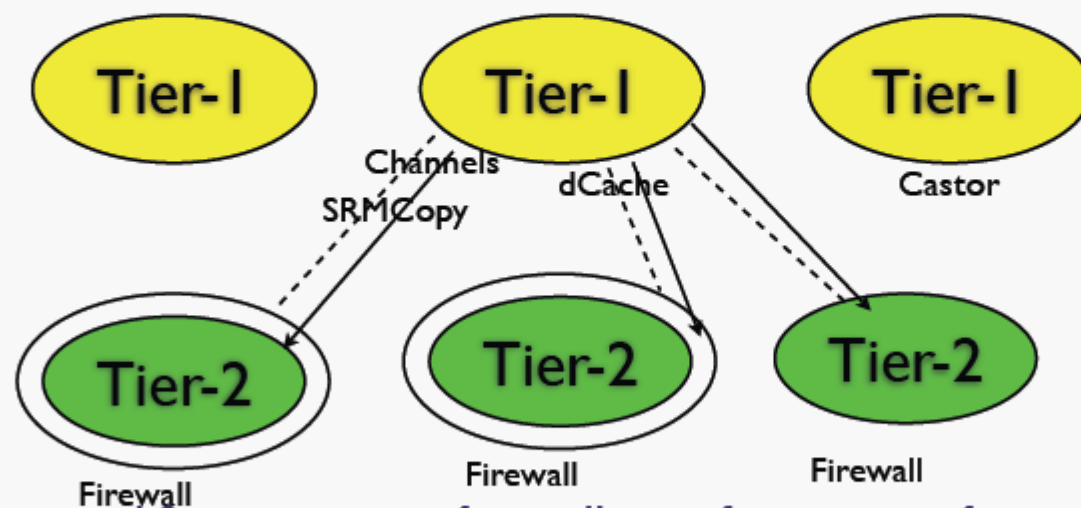
- DDM migration to version 0.3
 - Oracle backend + better handling of OPEN datasets
 - T0 exercise with 0.3
 - No need of extra VOBOXes
- SRM
 - Need to have SRM2.2 in production for May
 - At the T1s and T0
 - Important especially for reprocessing **But no reprocessing this year???**
- LFC
 - Need deployment of new server ASAP (**CERN 1st half Feb, T1s 2nd half...**)
 - Efficient bulk GUID lookup
- Security Model
 - Need to start running DDM services with service certificates
 - And implement delegation
 - Technically already possible, tricky from the policy point of view.
- ARDA monitoring now in place
- Main problems:
 - ☹ Lack of coordination in LCG Application software (versions of POOLvsROOT, RFIO libraries etc ...)
 - ☹ Instability of SRM storage elements
- See <http://indico.cern.ch/sessionDisplay.py?sessionId=5&slotId=3&confId=3738#2007-01-26> for a description of ATLAS DDM activities expected in 2007



Proposed Changes


Proposal is to move the channels to the Tier-1 of the source

- Example CNAF to PISA, CNAF to Nebraska, CNAF to KNU
- As a first iteration we should be able to use a Tier-1 to STAR for some of them



The simplification comes from all transfers coming from one Tier-1 and going to any Tier-2, instead of coming from any Tier-1 and going to 7-8 Tier-1s

- Fewer configurations



- The *ALICE* session concentrated on different tutorials regarding specific aspects of *ALICE* software
- During the whole morning and part of the afternoon about 25 persons participated to the talks
- The global AliEn WMS Data Management Systems were explained to site managers (mostly dedicated to *ALICE*)
- The rest of the tutorial covered Monitoring, AliRoot and AliEn software
- Interesting discussions allowed direct contact between site managers and *ALICE* experts
- Considered as an interesting session for all present



LHCb Session

- Agenda page <http://indico.cern.ch/conferenceDisplay.py?confId=11345>
- Philippe went through the presentation he gave at the plenary session and that was the starting point for further more detailed discussions with T1s.
- Main topics:
 - Testing of glexec ("Glex Eck") at some of the sites;
 - Data owned by a single "abstract" DN (a service certificate).
 - DIRAC has to manage the disk quota at the site level. Tricky under the security point of view and traceability problems.
 - Instability at dCache centres (RAL and IN2P3) over the last weeks preventing successfully transfers and reconstruction activity.
 - Both RAL & IN2P3 site admins confirmed that there is a problem on their SRM endpoints and the only solution would be to pass to a more stable version (1.7) of dCache. ??? IN2P3 already on 1.7 ???
 - (Or presumably CASTOR in the case or RAL...)

Operations Workshop Conclusions

Ian Bird
CERN

wLCG Collaboration workshop
22-26 January 2007

<http://indico.cern.ch/materialDisplay.py?contribId=42&sessionId=22&materialId=slides&confId=3738>

(I'm not going to summarise the summary)



ROYAL INSTITUTE
OF TECHNOLOGY

**Operations Workshop,
ARM and COD meetings
Stockholm 11-15 June**

Interoperation workshop
Indiana – May or June?



Issues Raised

- Follow-up on 24x7 services and 'piquet'
 - Stand-by duty in CERN staff rules & regulations
- Follow-up on intervention scheduling and machine operation cycle / schedule
 - Interventions at least inconvenient time
 - There is no 'most convenient' time...
 - Piquet required for full period of LHC run - not just when beam is on!
- Follow-up on schedule for deploying residual services - in particular SRM v2.2 services
- Service stability issues - Thursday morning session
- *Integration of Grid and Experiment service operations, both globally and regionally*
 - Ideally should be in place for testing in Dress Rehearsals
 - (If we know what it means by then, of course...)




Wireless Access

- Concerned before the workshop about the implications of signing ~100 temporary network access requests
- ✓ Agreed with security team that provided a means of contacting these people was available (e.g. registration details) : OK
- 💣 **Unauthorised Network/P2P activity has been detected from VISITOR-xxxxxx. All UNAUTHORISED SOFTWARE MUST BE REMOVED and a confirmation with the NAMES OF THE APPLICATIONS removed sent back to us. Further details are given below.**
- Caused quite a stir...



Future

- *One* WLCG collaboration week per year is probably all that anyone can stand - *and anything less probably useless*
- Tentatively foresee this ~midway between previous year's data taking and current
 - e.g. April 2008 (also avoids Easter, Geneva car show, "the-week-formerly-referred-to-as-ski-week", etc.)
- Need sufficient time to analyse results of pilot run - uncertain schedule - January 2008 would be too early...
- Will need also smaller, focussed events
 - e.g. SRM v2.2 Service Deployment workshop July in DESY
 - WLCG "weekend" prior to CHEP (review progress on 'residual services', service hardening, dress rehearsals & last minute update on schedule and scope of Pilot Run)
- *i.e. maximum one big event at CERN per year*
- *Need 'Volunteers' to host smaller interim events...*



What to im

- Wireless availability & stability - seems unavoidable requirement these days
 - No-one (well, just one person) complained about power sockets (which are Swiss anyway... See picture...)
- Webcasting - I had understood that it would be *both* Webcast live & recorded
 - Apparently two different services
- Ran out of 'supplies' at WLCG Collaboration drink
 - Isn't this sufficient motivation to register?
- Experiment visit - very hard to get confirmed names!





PSS 2007 Commissioning - Main Points CERN IT Department
(as of Dec. 2006)

Power tests on magnet circuits

- Failed to start with full protection system
- 10 to 15 days on standby before start after visit
- 10 days of normal operating temperature after visit
- Failed to start with limited protection system
- 10 to 15 days on standby before start after visit
- 10 days of normal operating temperature after visit

All special function equipment to be tested to 400GeV and more
Transfer fees, injection systems, extraction systems
of 1000000 particles (more)


Vacuum closed end August 2007
Global test of Access Control System October 2007
Engineering run in 2007

LHC Schedule from a WLF Service Perspective - 8



Summary (From Neil Geddes)

- Good attendance at plenary/summary sessions (100+)
 - Little discussion
- More discussion in “BOF”s
- Lots of corridor discussion and side ad-hoc meetings
- Felt to be well worthwhile, but a long week.
- Worth doing - but not too often





Thanks...

- To all those who came...
- To all of the speakers - particularly "Guest Stars" ...
- To all those who 'agreed' to chair a session / BOF
- To all those who 'volunteered' to be on a panel / take notes etc.
 - Definition of volunteering was given at July GDB...
- UDS for Indico and audio-visual support
- Fabienne for arranging coffee breaks, collaboration drink, transport to ATLAS and many other things...
- Ioannis, our cheerful guide
- Anyone I've forgotten... (sorry...)



Workshops

Events in this category:

♥ 2008

April 2008

21 - 25 [WLCG Collaboration Workshop \(Tier0/Tier1/Tier2\)](#)

♥ 2007

September 2007

01 - 02 [WLCG Collaboration Workshop \(Tier0/Tier1/Tier2\)](#)

January 2007

22 - 26 [WLCG Collaboration Workshop \(Tier0/Tier1/Tier2\)](#)

♥ 2006

March 2006

27 - 28 [xRootd Workshop](#)

Tools

◀ [Display](#)

▶ [Overview](#)

▶ [Calendar](#)

▶ [Map](#)

▶ [Modify](#)

▶ [Statistics](#)

▶ [Search](#)

▶ [Help](#)

Add Event

[Lecture](#)

[Meeting](#)

[Conference](#)




WLCG@CHEP'07

- 1-2 September, Victoria, Vancouver Island
 - I know that this clashes with GridPP19...
 - ☛ Getting to Island on Friday night requires preparation!
- 1 room (capacity 400) + 3 rooms (capacity 150 each)
 - BOFs for Data Management, Monitoring, Databases, Operations ???
- **Need to confirm by May 2007!**
- Sandwich lunch & coffee breaks will be provided on site in registration fee
- Registration will be joint with CHEP
- Wireless is being investigated
 - <http://www.chep2007.com/wlwg.html>
 - <http://indico.cern.ch/conferenceDisplay.py?confId=3578>



My Summary

- All in all, I believe it went rather well...
- Short of importing coffee and cakes from Bologna (**hint, hint**) I'm not sure what really could be improved... But let's try...
- **The main issues were exposed and discussed - widely agreed that this is useful/necessary**
- Hopefully, it will improve further still the excellent collaborative spirit that we've built up over the years...



END



The Worldwide LHC Computing Grid

What's New in WLCG Production Services?

Collaboration Workshop, Jan 2007



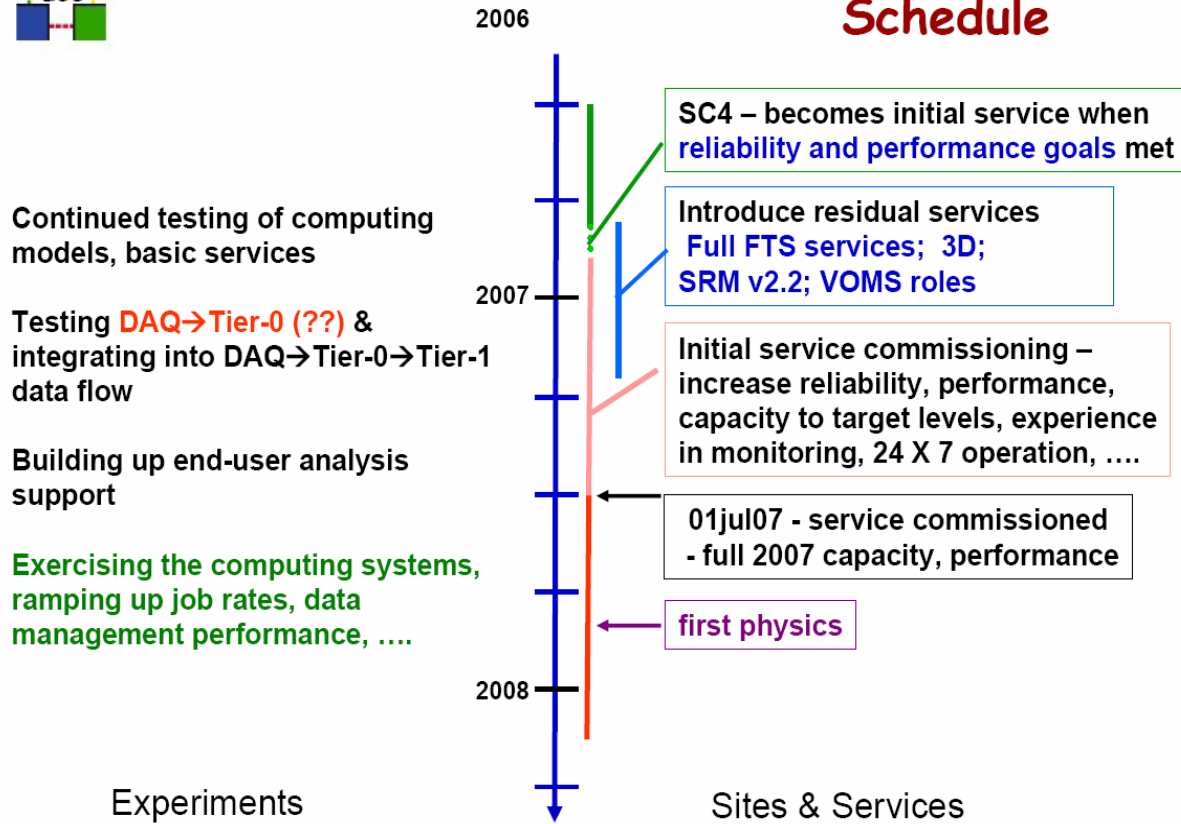
Worldwide LHC Computing Grid
Distributed Production Environment for Physics data Processing



What's New in WLCG?

1. Stability
 - Thursday
 2. Stability
 - 25th January 2007
 3. Stability
 - The Day After Tomorrow
- And a bit of functionality too...
- Today...

WLCG Commissioning Schedule



- **Still an ambitious programme ahead**
- Timely testing of full data chain from DAQ to T-2 chain was major item from last CR
 - DAQ→ T-0 still largely untested

Inter-Site Rates - Revised Megatable

Centre	T0→T1	T1→T2	T2→T1	T1↔T1
	Predictable – Data Taking	Bursty – User Needs	Predictable – Simulation	Scheduled Reprocessing
IN2P3, Lyon	220	286.2	85.5	498.0
GridKA, Germany	220	384.9	84.1	395.6
CNAF, Italy	190	321.3	58.4	583.8
FNAL, USA	110	415.0	52.6	417.0
BNL, USA	300	137.7	24.8	358.0
RAL, UK	120	108.3	36.0	479.4
NIKHEF, NL	160	34.1	6.1	318.4
ASGC, Taipei	120	126.5	19.3	241.2
PIC, Spain	100	167.1	23.3	294.5
Nordic Data Grid Facility	60	-	-	62.4
TRIUMF, Canada	60	-	-	58.0

Continued testing of computing models, basic services

Testing **DAQ→Tier-0 (??)** & integrating into DAQ→Tier-0→Tier-1 data flow

Building up end-user analysis support

Exercising the computing systems, ramping up job rates, data management performance,

WLCG Commissioning Schedule

2006

2007

2008

SC4 – becomes initial service when reliability and performance goals met

Introduce residual services
Full FTS services; 3D; **gLite 3.x**;
SRM v2.2; VOMS roles; **SL(C)4**

Initial service commissioning – increase performance, reliability, capacity to target levels, experience in monitoring, 24 x 7 operation,

01 jul07 - service commissioned - full 2007 capacity, performance

first collisions in the LHC. Full FTS services demonstrated at 2008 data rates for all required Tx-Ty channels, over extended periods, including recovery (T0-T1).



FDR = First Dress Rehearsal

A complete exercise of the full chain from trigger to (distributed) analysis, to be performed in 2007, a few months before data taking starts

Some details

- Generate
- Filter
- Pass
- Mix
- Run L

Services must be in place well before to allow adequate time for testing - 'SM' is $T_0 - 2\text{months}$

- Produce byte streams → emulate the raw data
- Send raw data to Point 1, pass through HLT nodes (flag mode) and SFO, write out events by streams, closing files at boundary of luminosity blocks.
- Send events from Point 1 to Tier0
- Perform calibration & alignment at Tier0 (also outside ?)
- Run reconstruction at Tier0 (and maybe Tier1s ?) → produce ESD, AOD, TAGs
- Distribute ESD, AOD, TAGs to Tier1s and Tier2s
- Perform distributed analysis (possibly at Tier2s) using TAGs
- MCTruth propagated down to ESD only (no truth in AOD or TAGs)

Ambitious goals... need to plan it carefully (both in terms of effort needed and of technical issues and implications)



Operations Issues

- WLCG Collaboration workshop 'pauses' Thursday lunchtime & restarts Friday with BOF sessions and Experiment Tutorials
- Thursday afternoon:
 - Joint Operations Workshop (continues Friday am)
 - [WLCG Collaboration Meeting (closed)]
- Final Plenary session:
 - **Integration of Grid and Experiment service operations, both globally and regionally**
 - Experiment Operations Coordinators, Tier1 Representatives, EGEE / OSG / NDGF / WLCG Operations Coordinators