CHEP – Mumbai, February 2006

The LCG Service Challenges

Focus on SC3 Re-run; Outlook for 2006

Jamie Shiers, LCG Service Manager

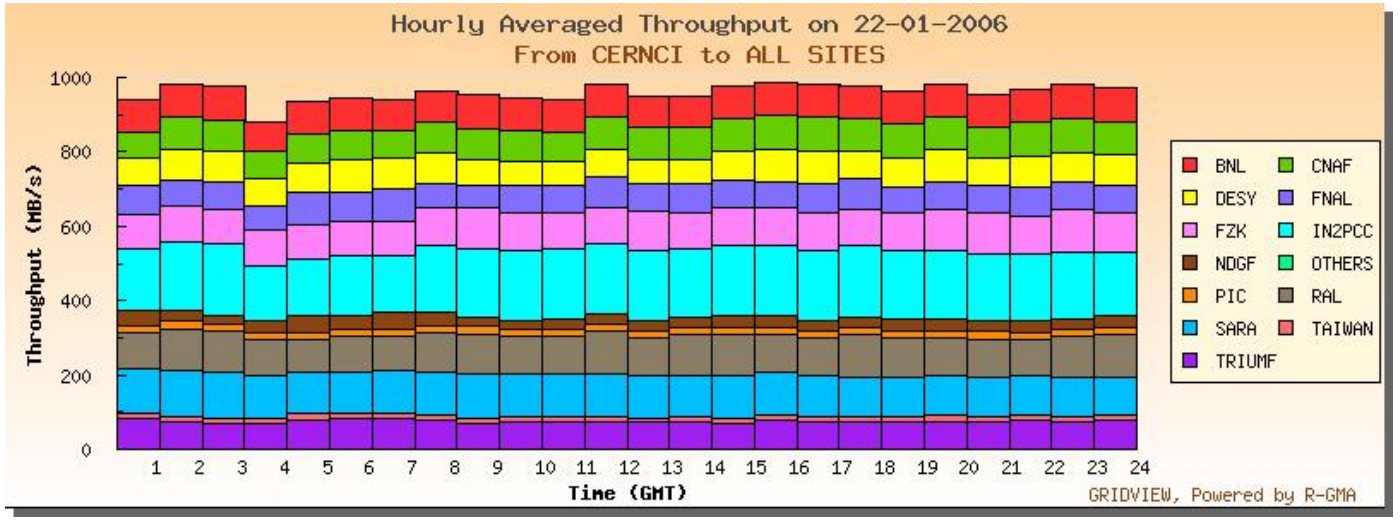

Abstract

- The LCG Service Challenges are aimed at achieving the goal of a **production quality** world-wide Grid that meets the requirements of the LHC experiments in terms of functionality and scale.
- This talk highlights the **main goals** of the Service Challenge programme, **significant milestones** as well as the key services that have been validated in production by the LHC experiments.
- The LCG Service Challenge programme currently involves both the 4 LHC experiments as well as many sites, including the Tier0, **all Tier1s** as well as a number of **key Tier2s**, allowing all primary data flows to be demonstrated.
- The functionality so far addresses all primary offline Use Cases of the experiments except for analysis, the latter being addressed in the final challenge - scheduled to run from April until September 2006 - prior to delivery of the **full production** Worldwide LHC Computing Service.

Agenda

- Overview of the Service Challenge Programme
- What we have – and have not – achieved so far
- **The "SC3 re-run" (focusing on Tier0-Tier1 aspects)**
- **Other aspects are also very important, but this is a key area where we did not meet our 2005 targets**
- Timetable and high level tasks for the future
- Summary and Conclusions

Introduction

- The (W)LCG Service Challenges are about **preparing**, **hardening** and **delivering** the production Worldwide LHC Computing Environment (WLCG)
- The date for delivery of the production LHC Computing Environment is **30 September 2006**
- Production Services are required as from **1 September 2005** and
  - service phase of Service Challenge 3

  1 May 2006 ← now 1 June after validation of next m/w release
  - service phase of Service Challenge 4

➢ **This is not a drill.**

SC3 Goals

➢ **Much more than just a throughput test!**

Activities kicked off to address these issues were successful

☺ **More data management services:**
- SRM required at all sites
- Reliable File Transfer Service based on gLite FTS
- LFC file catalog deployed as required
  - Global catalog for LHCb, site local for ALICE + ATLAS
- [Other services as per BSWG]

☺ **More sites:**
- *All* Tier1s took part – this was better than foreseen!
- Many Tier2s – now above **20** sites, covering most regions. This too is working well!
- Workshops held in many countries / regions (T1 + T2s + experiments) – **this should continue!**
  - UK, Italy, France, Germany, Asia-Pacific, North America, Nordic region, …
  - **(A-P w/s early May 2006; North American w/s around September GDB???)**

☺ **All experiments:**
- Clear goals established together with metrics for measuring success
- List of issues / requirements has been produced – **plan for addressing remaining issues**

☹ **Throughput targets:**
- 50% higher than SC2 but using SRM & FTS as above (150MB/s to disk at T1s)
- 60MB/s to tape at Tier1s (following disk – disk tests)
- Modest T2->T1 targets, representing MC upload (3 x 1GB file / hour)

SC3(4) Service Summary

- Services identified through combination of Baseline Services Working Group, Storage Management Workshop and 1-1 discussions with experiments
  - Timeline of BSWG & service setup lead time did not allow to wait for 'final report' before starting to implement services
- For new services (LFC, FTS), two flavours established at CERN
  - 'Pilot' – to allow experiments to gain experience with functionality, adapt their s/w to interfaces etc.
  - 'SC3' – full production services
  - ➢ **This separation proved useful and needs to continue with Pre-Production System!**
- New services for sites: LFC (most sites), FTS (T1s), SRM (DPM, dCache at T2s)
- Support lists established for these services, plus global 'catch-call'
  - Concrete plan for moving VOs to GGUS is being discussed with VOs
- 'SC3' services being re-deployed for full production
  - Some of this work was done during end-Oct / early Nov intervention
- ➢ **Services by site; VO variations; Detailed Functionality and Timeline exists**

SC3 Throughput Tests

- **Unfortunately, July Throughput Tests did not meet targets**
  - Compounded by service instability
  - Continued through to the end, i.e. disk – disk, disk – tape and T2 – T1 components
  - Spent most of August debugging and fixing
  - dCache workshop held in DESY identified concrete actions / configurations / dCache improvements
  - Improvements also in CASTOR SRM & gLite FTS
  - All software upgrades now released & deployed
- **Disk – disk rates obtained in July around 1/2 target, without stability!**

SC3 Throughput: Disk & Tape

Disk target:
150MB/s/site
1GB/s (CERN)

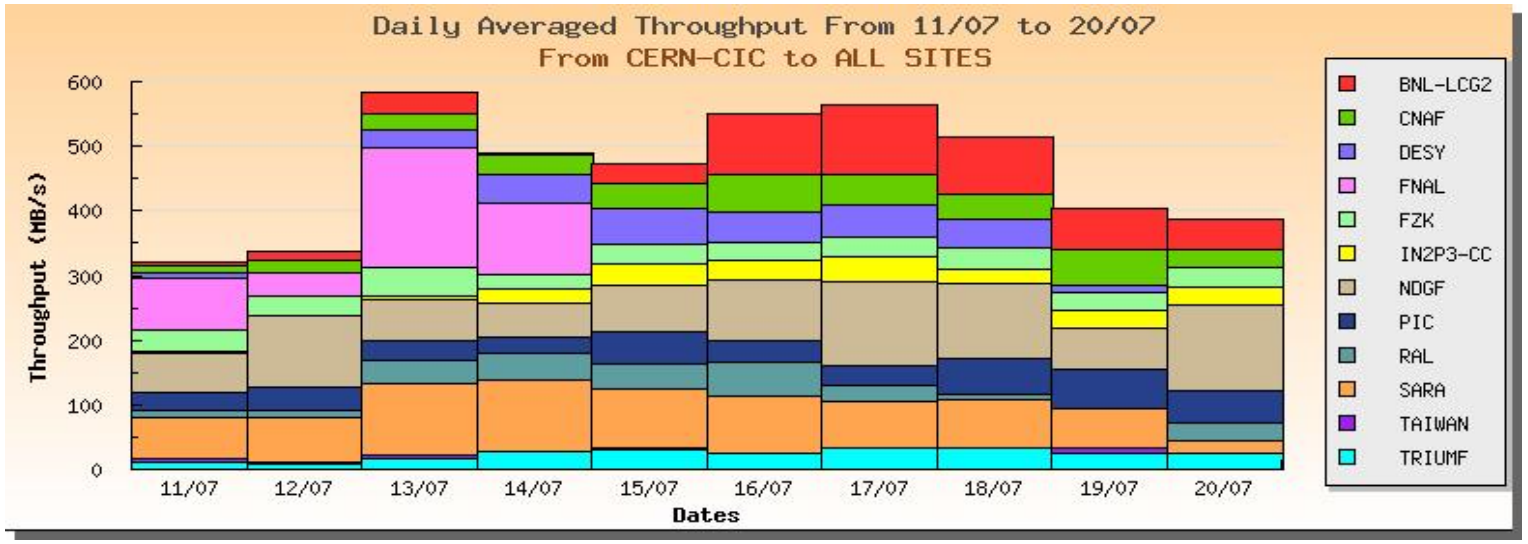

Tape target:
60MB/s/site
(Not all sites)

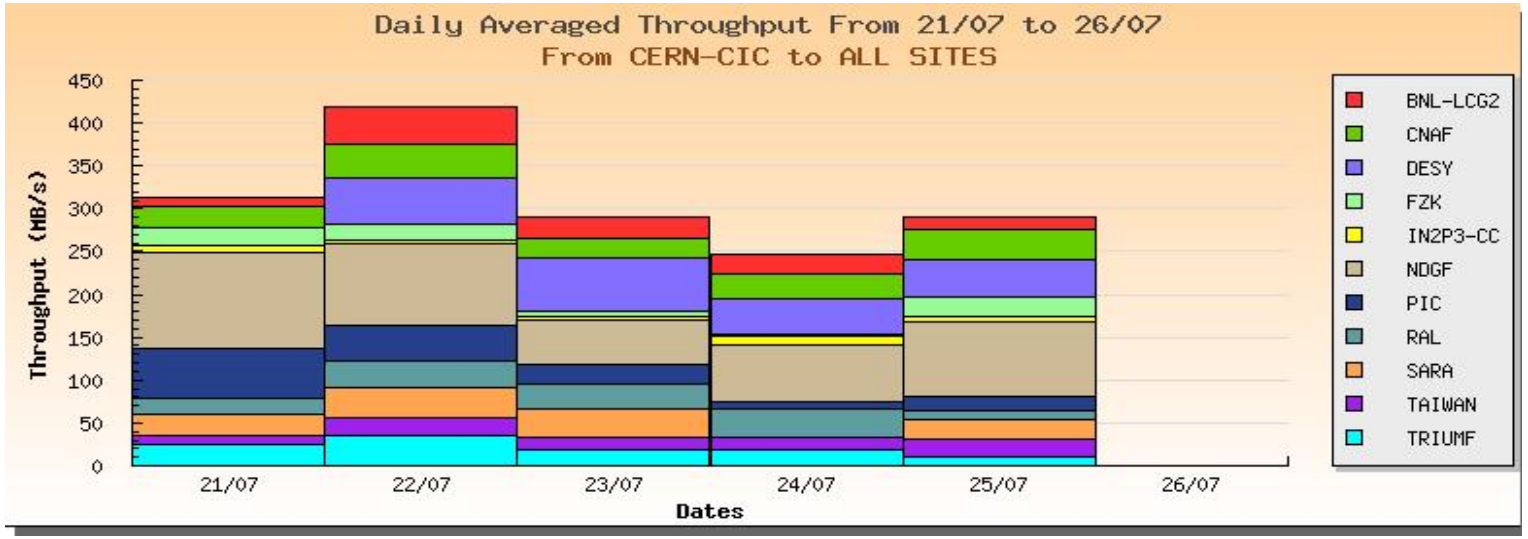

James Casey

Results of SC3 in terms of Transfers

- Target data rates 50% higher than during SC2
- All T1s (most supporting T2s) participated in this challenge
- Transfers between SRMs (not the case in SC1/2)
- Important step to gain experience with the services before SC4

Site	MoU Target (Tape)	Daily average MB/s (Disk)
ASGC	100	10
BNL	200	107
FNAL	200	185
GridKa	200	42
CC-IN2P3	200	40
CNAF	200	50
NDGF	50	129
PIC	100	54
RAL	150	52
SARA/NIKHEF	150	111
TRIUMF	50	34

Rates during July throughput tests. Better single-site rates since, but need to rerun tests…

For this we need dCache 1.6.6(+) to be released/deployed, latest FTS (now), network upgrades etc.

January '06 (<CHEP)

Post-debugging

- Now can achieve same rate as before with fewer sites
  - Still need to add in other sites, and see what the new upper limit is

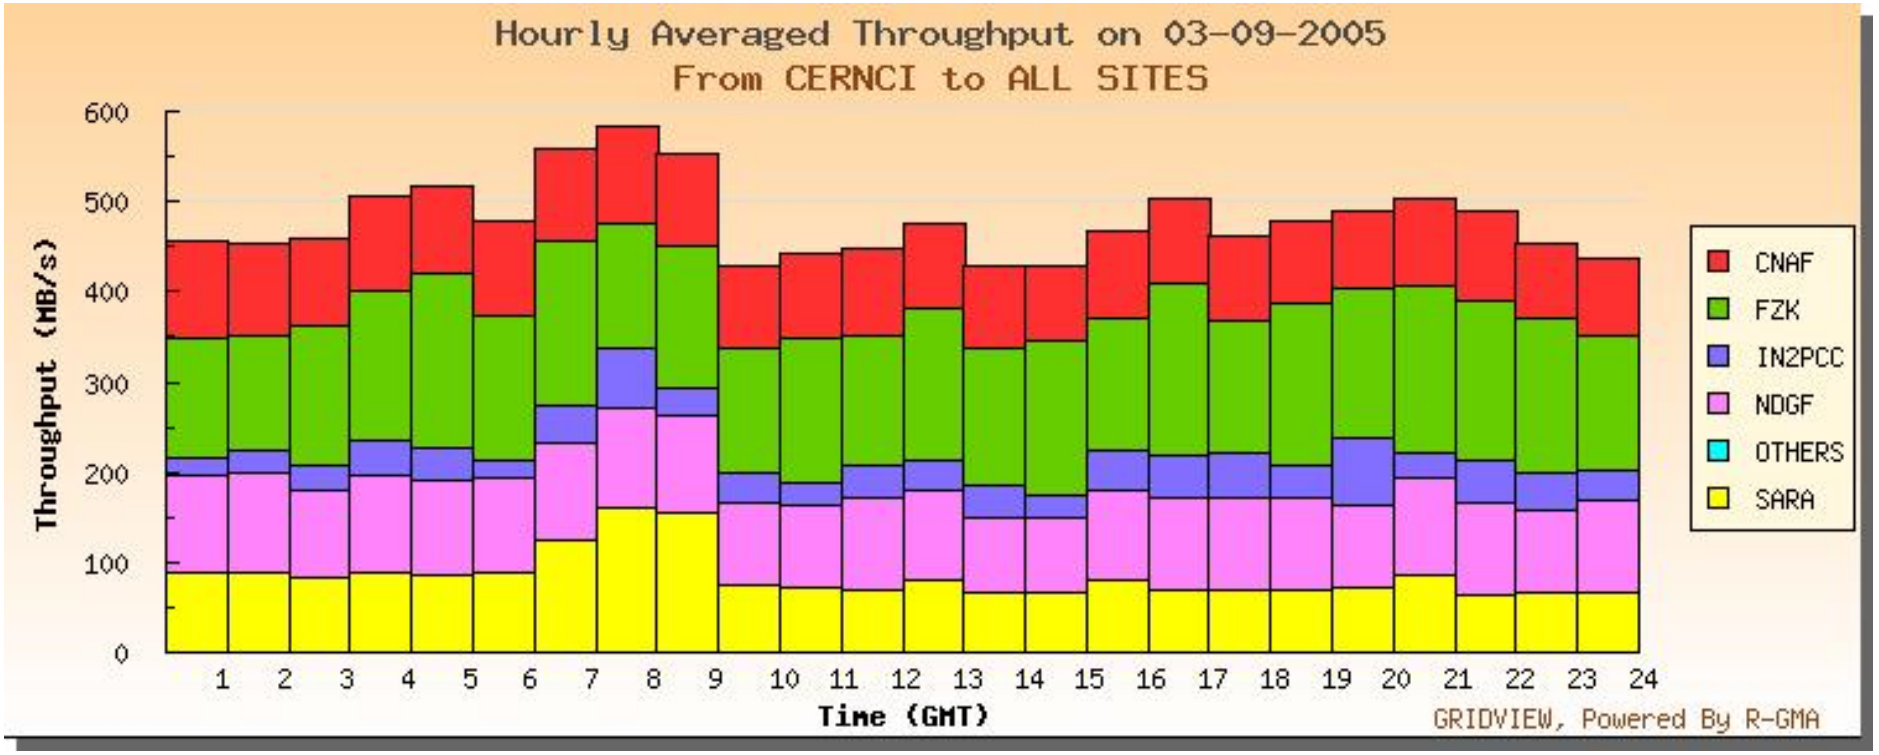

Pre-Requisites for Re-Run of Throughput Tests

- Deployment of gLite FTS 1.4 (srmcp support)
  - ✓ Done at CERN in recent intervention
- dCache 1.6.6 (or later) release and deployed at all dCache sites.
  - ✓ Released – some sites already planning upgrade
- CASTOR2 clients and CASTORSRM version 2.2.8 (or later) at all CASTOR sites (ASGC, CNAF, PIC).
- Upgrade to CERN internal network infrastructure.
  - Partly done – remainder during Christmas shutdown
  - ➢ N.B. intend to keep Grid running over Xmas! (Close to last chance…)
- 10Gbit/s network connections at operational at the following sites:
  - IN2P3, GridKA, CNAF, NIKHEF/SARA, BNL, FNAL

dCache – the Upgrade (CHEP 2006)

- For the last two years, the dCache/SRM Storage Element has been successfully integrated into the LCG framework and is in heavy production at several dozens of sites, spanning a range from single host installations up to those with some hundreds of TB of disk space, delivering more than 50 TB per day to clients. Based on the permanent feedback from our users and the detailed reports given by representatives of large dCache sites during our workshop at DESY end of August 2005, the dCache team has been identified important areas of improvement.
- This includes a more sophisticated handling of the various supported tape back-ends, the introduction of multiple I/O queues per pool with different properties to account for the diverse behaviours of the different I/O protocols and the possibility to have one dCache instance spread over more than one physical site.
- … changes in the name-space management as short and long term perspective to keep up with future requirements.
- … initiative to make dCache a widely scalable storage element by introducing dCache, the Book, plans for improved packaging and more convenient source code license terms.
- Finally I would like to cover the dCache part of the German e-science project, d-Grid, which will allow for improved scheduling of tape to disk restore operations as well as advanced job scheduling by providing extended information exchange between storage elements and Job Scheduler.

Disk – Disk Rates (SC3 Repeat)

Centre (Nominal rate)	ALICE	ATLAS	CMS	LHCb	Target Data Rate MBytes/sec (Site target)
Canada, TRIUMF		X			50 (++)
France, CC-IN2P3 (200MB/s)	X	X	X	X	150 (100?)
Germany, GridKA (200MB/s)	X	X	X	X	150
Italy, CNAF (200MB/s)	X	X	X	X	150
Netherlands, NIKHEF/SARA	X	X		X	150
Nordic Data Grid Facility	X	X	X		50
Spain, PIC Barcelona		X	X	X	100 (30)
Taipei, ASGC		X	X		100 (75)
UK, RAL	X	X	X	X	150
USA, BNL (200MB/s)		X			150
USA, FNAL (200MB/s)			X		150

- These are the nominal data rates capped at 150MB/s

Disk – Tape Rates (SC3 Re-run)

Centre	*ALICE*	*ATLAS*	*CMS*	*LHCb*	*Target Data Rate MB/s*
Canada, TRIUMF		X			50
France, CC-IN2P3	X	X	X	X	50
Germany, GridKA	X	X	X	X	50
Italy, CNAF	X	X	X	X	50
Netherlands, NIKHEF/SARA	X	X		X	50
Nordic Data Grid Facility	X	X	X		50
Spain, PIC Barcelona		X	X	X	50
Taipei, ASGC		X	X		50
UK, RAL	X	X	X	X	50
USA, BNL		X			50
USA, FNAL			X		50

- Target ~5 drives of current technology
- Rate per site is 50MB/s (60MB/s was July target)

February 2006

SC3 Summary

- Underlined the complexity of reliable, high rate sustained transfers to be met during LHC data taking
- Many issues understood and resolved – need to confirm by re-run of Throughput exercise
- We are now well into the Service Phase (Sep – Dec)
- Collaboration with sites & experiments has been excellent
- We are continuing to address problems as they are raised
- Together with preparing for SC4 and WLCG pilot / production
- The experiment view have been summarised by Nick Brook
- See also GDB presentations of this week

SC3 Tier0 – Tier1 Disk – Disk Rerun

- Involved all Tier1 sites + DESY
  - BNL, CNAF, DESY, FNAL, FZK, IN2P3, NDGF, PIC, RAL, SARA, TAIWAN, TRIUMF
- Preparation phase significantly smoother than previously
  - Although a number of the problems seen had occurred before…
  - As usual, better documentation…
- Sites clearly have a much (much) better handle on the systems now…
  - What to tune, preferred settings etc.
- We still do not have the stability required / desired…
  - The daily **average** needs to meet / exceed targets
  - We need to handle this without **"heroic efforts"** at all times of day / night
  - We need to **sustain** this over many (100) days
  - We need to test **recovery** from problems (individual **sites** – also **Tier0**)
- **But a big improvement, also in rates achieved**
  - Limited by h/w configuration at CERN to ~1GB/s – the target (average)

SC3 Re-run Results

- ➢ We managed to obtain close to 1GB/s for extended periods
- ☺ Several sites exceeded the targets agreed
- ☺ Several sites reached or even exceeded their 'nominal rates'
- ☹ Still see quite a few operational problems
- ☹ And the software is not completely debugged yet…
- ➢ But its an encouraging step on reaching the final goals…

Hourly Averaged Throughput on 23-01-2006
From CERNCI to FZK

250

Hourly Averaged Throughput on 23-01-2006
From CERNCI to IN2PCC

250

Averaged Throughput during the last 24 hrs (23/01 - 24/01)
From CERNCI to RAL

Hourly Averaged Throughput on 19-01-2006
From CERNCI to SARA

250

Hourly Averaged Throughput on 22-01-2006
From CERNCI to TRIUMF

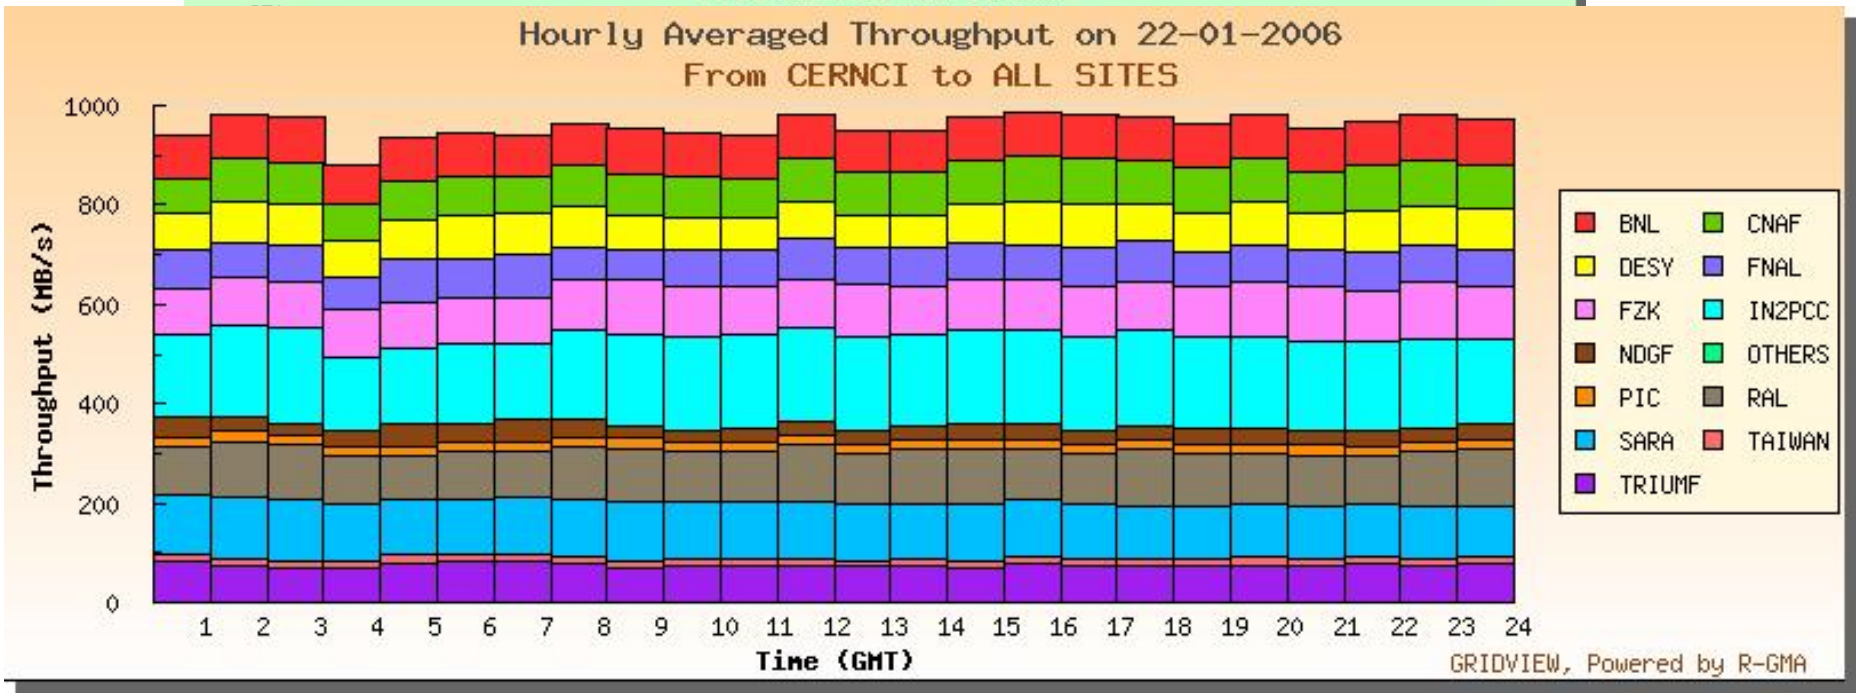

Currently Scheduled Throughput Tests

- January 2006 – rerun of SC3 disk – disk transfers (max 150MB/s)
  - All Tier1s and DESY participated. Achieved ~1GB/s out of CERN; good rates to sites
- **February 2006 – rerun of SC3 disk – tape transfers (50MB/s – was 60MB/s in July)**
  - Sites should allocate 5 current generation drives and understand issues involved
- March 2006 – T0-T1 "loop-back" tests at 2 x nominal rate
  - CERN, using new tape technology and corresponding infrastructure
- April 2006 – T0-T1 disk-disk (nominal rates) disk-tape (50-75MB/s)
  - All Tier1s – disk rates at BNL, FNAL, CNAF, FZK, IN2P3 go up to 200MB/s
- July 2006 – T0-T1 disk-tape (nominal rates)
  - All Tier1s – rates 50 – 200MB/s depending on VOs supported & resources provided
- T1-T1; T1-T2; T2-T1 and other rates TBD according to CTDRs
  - All Tier1s; 20 – 40 Tier2s; all VOs; all offline Use Cases
- **Still significant work ahead for experiments, T0, T1s and T2s!**

Further Throughput Tests

- Need to define throughput tests for TierX ←→ TierY transfers
  - Mumbai (February) & CERN (June) workshops
- Need to demonstrate sustained average throughput at or **above** MoU targets from Tier0 to tape at all Tier1s
- Need to demonstrate recovery of backlog(s)
  - ~4 hour downtimes at individual Tier1s : ~1 day downtime? : more?
  - ~4 hour downtime at the Tier0 (!) : ~1 day downtime?
    - Use operations log to establish realistic Use Cases
- Full end-to-end throughput demonstration

➢ **All the above confirmed using experiment-driven transfers**

SC3 Services – Lessons (re-)Learnt

- It takes a **L O N G** time to put services into (full) production
- A lot of experience gained in ***running*** these services Grid-wide
- Merge of 'SC' and 'CERN' daily operations meeting has been good
- **Still need to improve 'Grid operations' and 'Grid support'**
- A CERN 'Grid Operations Room' needs to be established
- Need to be more rigorous about:
  - Announcing scheduled downtimes;
  - Reporting unscheduled ones;
  - Announcing experiment plans;
  - Reporting experiment results;
  - Attendance at 'V-meetings';
  - …

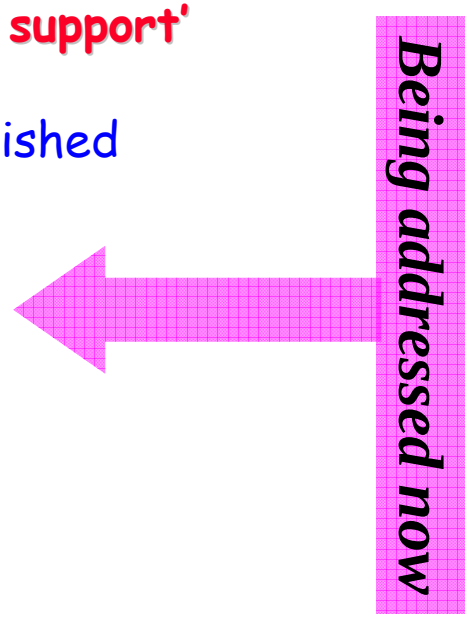

- **A daily OPS 'meeting' is foreseen for LHC preparation / commissioning**

WLCG - Major Challenges Ahead

1. Get data rates at all Tier1s up to MoU Values
   - This is currently our biggest challenge **but good progress recently!**
   - Plan is to work with a few key sites and gradually expand
   - (Focus on highest-data rate sites initially…)

2. (Re-)deploy Required Services at Sites so that they meet MoU Targets
   - Tier0 will have all services re-deployed prior to SC4 Service Phase (WLCG Pilot)
   - Plans are being shared with Tier1s and Tier2s, as will be experience
   - LCG Service Coordination team will be proactive in driving this forward
   - ➢ **A lot of work, but no major show-stopper foreseen**

3. **Understand other key Use Cases for verification / validation**
   - This includes also other inter-Tier transfers (see workshops)
   - Many will be tested by experiment production
   - Which should be explicitly tested as dedicated "Service Tests"?

Timeline - 2006

January	SC3 disk repeat – nominal rates capped at 150MB/s SRM 2.1 delivered (?)	July	Tape Throughput tests at full nominal rates! (150MB/s for RAL)
February	CHEP w/s – T1-T1 Use Cases, SC3 disk – tape repeat (50MB/s, 5 drives)	August	T2 Milestones – debugging of tape results if needed
March	Detailed plan for SC4 service agreed (M/W + DM service enhancements)	September	LHCC review – rerun of tape tests if required?
April	SC4 disk – disk (nominal) and disk – tape (reduced) throughput tests	October	WLCG Service Officially opened. Capacity continues to build up.
May	Start of SC4 production Tests by experiments of 'T1 Use Cases'	November	1st WLCG 'conference' All sites have network / tape h/w in production(?)
June	'Tier2 workshop' – identification of key Use Cases and Milestones for T2s	December	'Final' service / middleware review leading to early 2007 upgrades for LHC data taking??

O/S Upgrade? Sometime before April 2007!

Conclusions

- A great deal of progress in less than one year…
- Which is all we have left until FULL PRODUCTION
- Focus now is on SERVICE and STABILITY
- Service levels & functionality (including data transfers) defined in WLCG MoU
- ☺ A huge amount of work by many people… Thanks to all!
- ☺ (From CSO & LHCC Referees too!)

Acknowledgements

- The above represents the (hard) work of a large number of people across many sites and within all LHC collaborations

- Thanks to everybody for their participation, dedication and enthusiasm!

The LHC Computing Grid – (The Worldwide LCG)