

The Worldwide LHC Computing Grid

Introduction & Workshop Goals

Jamie Shiers, CERN, Geneva, Switzerland



Worldwide LHC Computing Grid
Distributed Production Environment for Physics data Processing



Overview

- Reminder of Service Challenge Goals and Status
- Outline of LHC Commissioning and Initial Operation
- Reminder of Service Availability Targets
- Experiment Activities Prior to First Collisions
- Brief Summary of T2 Workshop Questionnaire Responses

Abstract (SC4 Service Start)

- The production phase of the Service Challenge 4 - aka the **Pilot WLCG Service** - started at the beginning of **June 2006**. This leads to the full production WLCG service from **October 2006**.

- Thus the WLCG pilot is the **final** opportunity to shakedown not only the services provided as part of the WLCG computing environment - including their functionality - but also the **operational** and **support** procedures that are required to offer a full production service.

- This talk will **focus on operational** aspects of the service, together with the currently planned production / test activities of the LHC experiments to validate their computing models and the service itself.

- 💣 Despite the huge achievements over the last 18 months or so, we still have a very long way to go. Some sites / regions may not make it - at least not in time. Have to focus on a few key regions...

Service Challenges - Reminder

- Purpose
 - Understand what it takes to operate a real grid service – run for weeks/months at a time (not just limited to experiment Data Challenges)
 - Trigger and verify Tier-1 & large Tier-2 planning and deployment –
- tested with realistic usage patterns
 - Get the essential grid services ramped up to target levels of reliability, availability, scalability, end-to-end performance
 - Four progressive steps from October 2004 thru September 2006
 - End 2004 - SC1 – data transfer to subset of Tier-1s
 - Spring 2005 – SC2 – include mass storage, all Tier-1s, some Tier-2s
 - 2nd half 2005 – SC3 – Tier-1s, >20 Tier-2s – first set of baseline services
 - **Jun-Sep 2006 – SC4 – pilot service**
- **Autumn 2006 – LHC service in continuous operation
– ready for data taking in 2007**

SC4 - Executive Summary

We have shown that we can drive transfers at full nominal rates to:

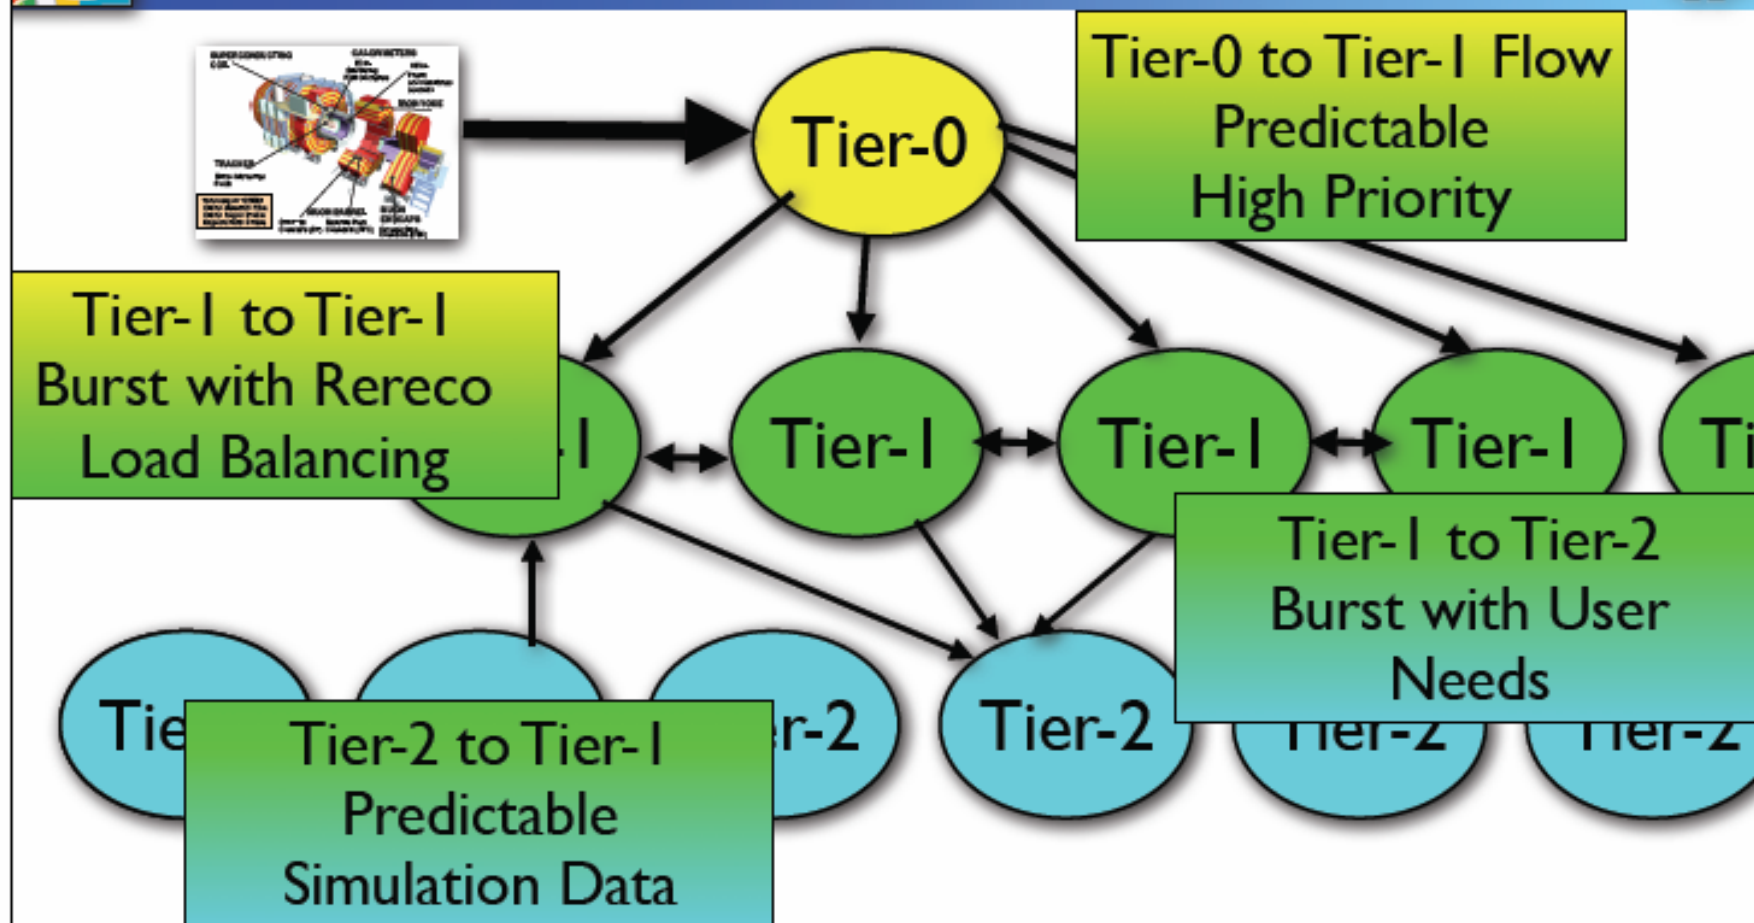
- Most sites simultaneously;
- All sites in groups (modulo network constraints - PIC);
- At the target nominal rate of 1.6GB/s expected in pp running

In addition, several sites exceeded the disk - tape transfer targets

- There is no reason to believe that we cannot drive all sites at or above nominal rates for sustained periods.

But

- There are still major operational issues to resolve - and most importantly - a full end-to-end demo under realistic conditions



Tier-2 centers may have relationships with Tier-1 centers for management, support, and operations

➔ Data access may come from a variety of Tier-1 centers

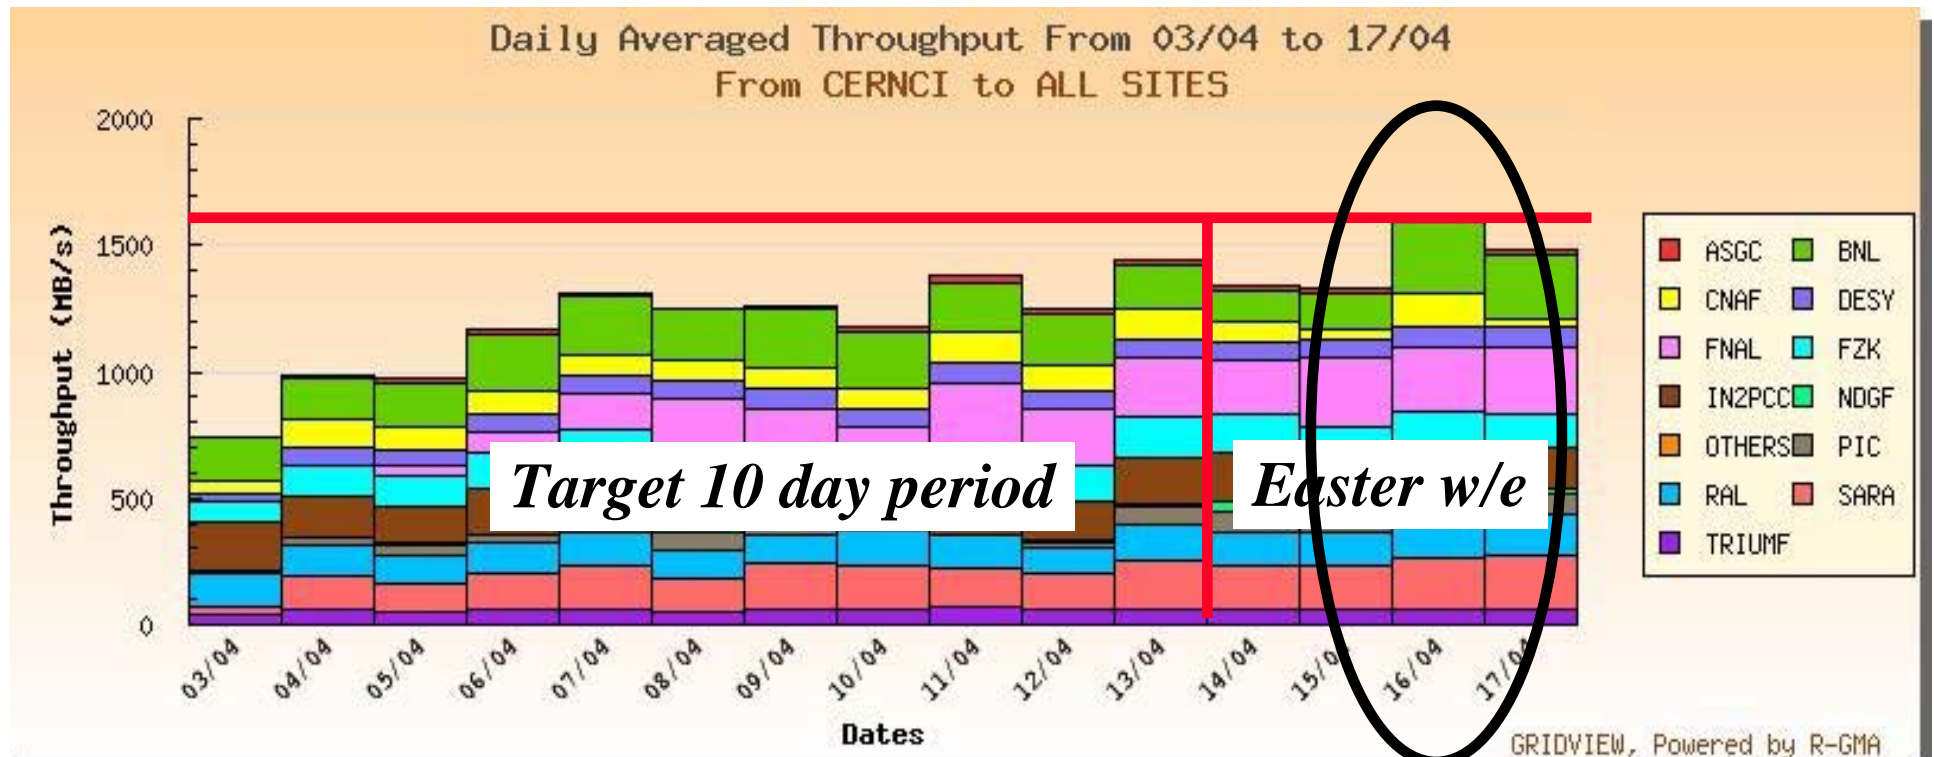
Heat

Nominal Tier0 - Tier1 Data Rates (pp)

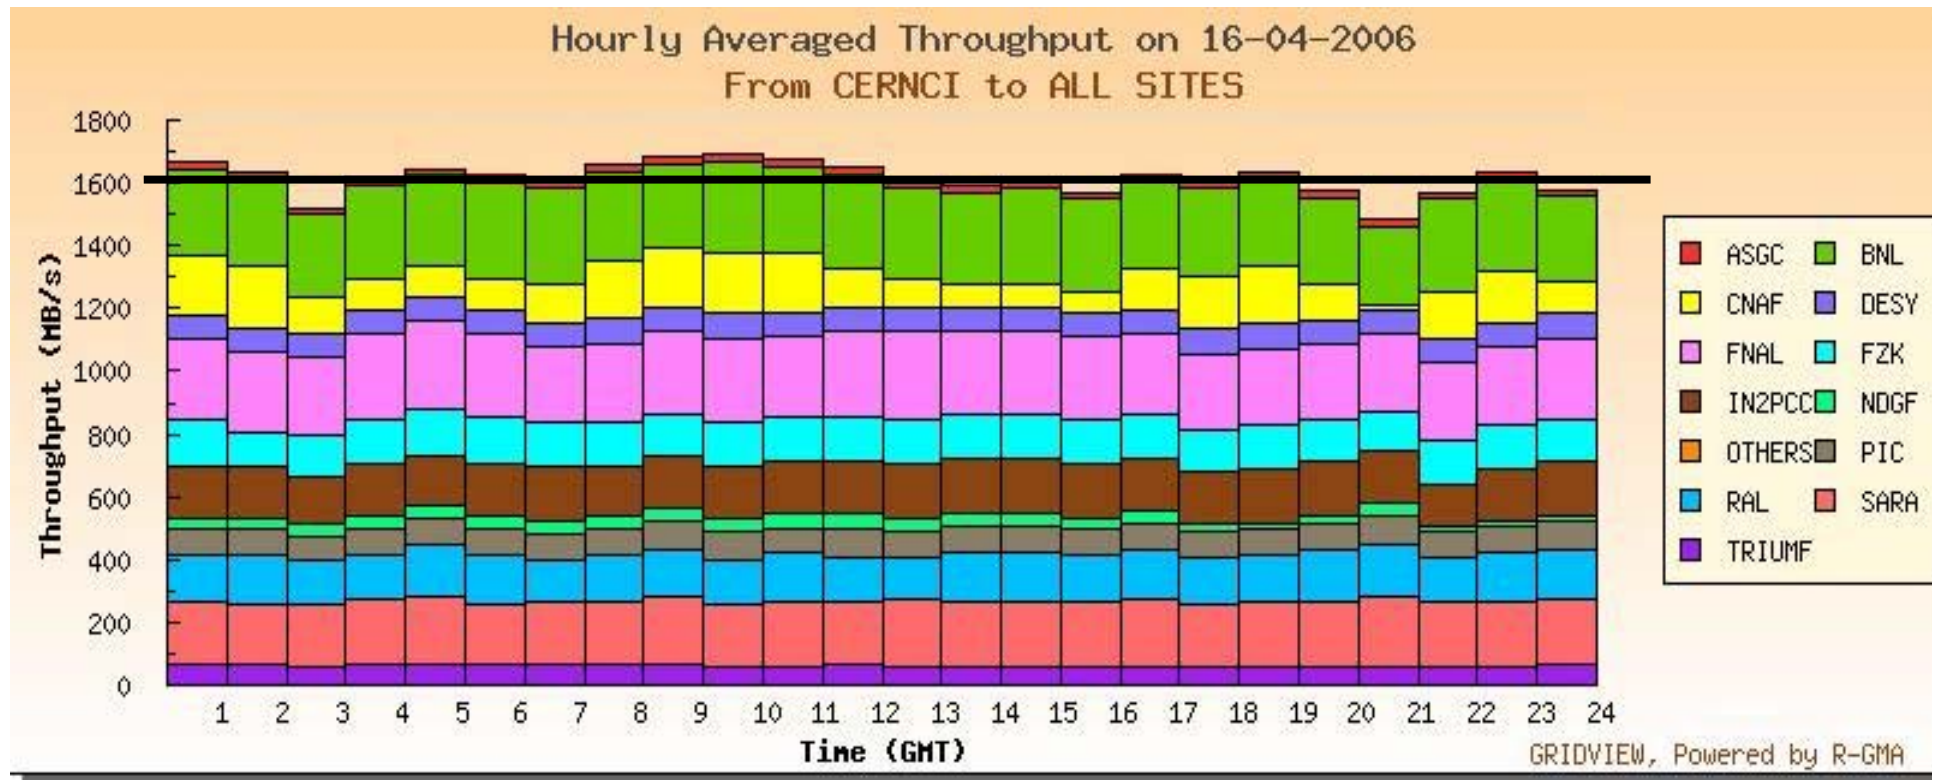
<i>Tier1 Centre</i>	<i>ALICE</i>	<i>ATLAS</i>	<i>CMS</i>	<i>LHCb</i>	<i>Target</i>
IN2P3, Lyon	9%	13%	10%	27%	200
GridKA, Germany	20%	10%	8%	10%	200
CNAF, Italy	7%	7%	13%	11%	200
BNL, USA	-	22%	-	-	200
FNAL, USA	-	-	28%	-	200
RAL, UK	-	7%	3%	15%	150
NIKHEF, NL	(3%)	13%	-	23%	150
ASGC, Taipei	-	8%	10%	-	100
PIC, Spain	-	4% (5)	6% (5)	6.5%	100
Nordic Data Grid Facility	-	6%	-	-	50
TRIUMF, Canada	-	4%	-	-	50
TOTAL					1.6GB/s

SC4 T0-T1: Results

- **Target:** sustained disk - disk transfers at **1.6GB/s** out of CERN at full nominal rates for **~10 days**
- **Result:** just managed this rate on **Good Sunday (1/10)**



Easter Sunday: > 1.6GB/s including DESY



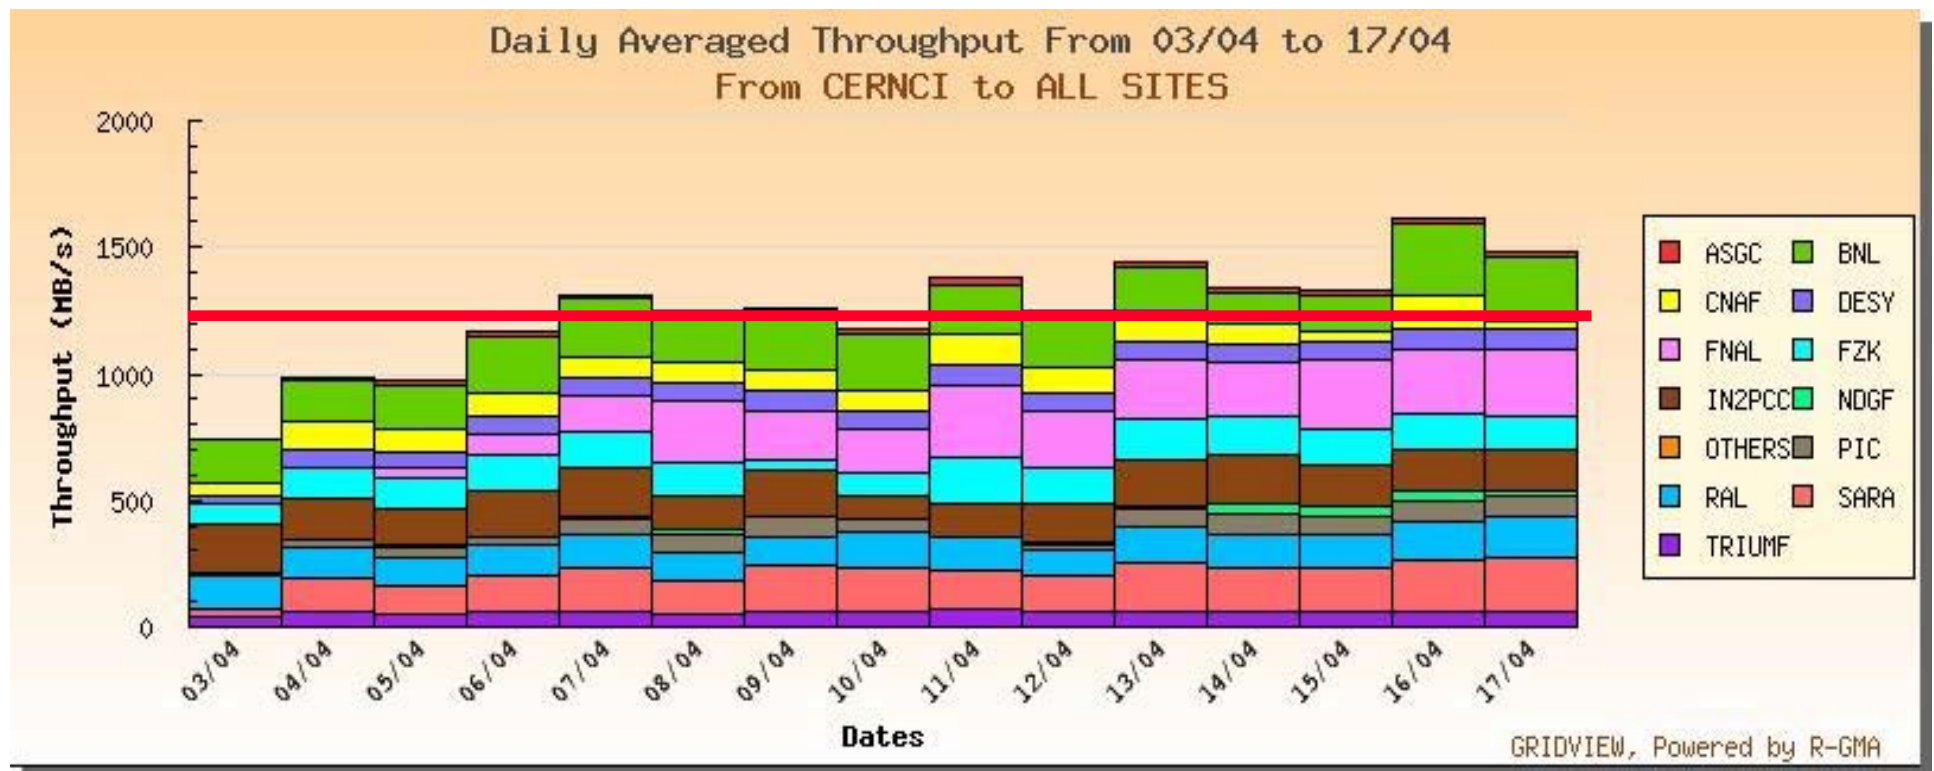
GridView reports 1614.5MB/s as daily average

Heat

Tier0 - Tier1 Rates (Megatable)

<i>Tier1 Centre</i>	<i>ALICE</i>	<i>ATLAS</i>	<i>CMS</i>	<i>LHCb</i>	<i>Total</i>
IN2P3, Lyon	27.9	75.4	43	22.6	168.9
GridKA, Germany	60.0	63.8	37	18.5	179.3
CNAF, Italy	34.6	107.0	55	18.1	214.7
BNL, USA	-	186.5	-	-	186.5
FNAL, USA	-	-	110	-	110
RAL, UK	8.8	76.8	7	18.5	111.1
NIKHEF, NL	13.8	72.0	-	21.2	107.0
ASGC, Taipei	-	35.7	37	-	72.7
PIC, Spain	-	20.6	19	15.7	55.3
Nordic Data Grid Facility	21.2	20.6	-	-	41.8
TRIUMF, Canada	-	19.2	-	-	19.2
US ALICE	46.4	-	-	-	46.4
TOTAL					1312.9


SC4 Revisited



Global Inter-Site Rates

<i>Centre</i>	<i>T0->T1</i>	<i>T1->T2</i>	<i>T2->T1</i>	<i>T1<->T1</i>
	Predictable – Data Taking	<i>Bursty – User Needs</i>	Predictable – Simulation	Scheduled Reprocessing
IN2P3, Lyon	168.9	286.2	85.5	498.0
GridKA, Germany	179.3	384.9	84.1	395.6
CNAF, Italy	214.7	321.3	58.4	583.8
FNAL, USA	110	415.0	52.6	417.0
BNL, USA	186.5	137.7	24.8	358.0
RAL, UK	111.1	108.3	36.0	479.4
NIKHEF, NL	107.0	34.1	6.1	310.4
ASGC, Taipei	72.7	126.5	19.3	241.2
PIC, Spain	55.3	167.1	23.3	294.5
Nordic Data Grid Facility	41.8	-	-	62.4
TRIUMF, Canada	19.2	-	-	59.0

Site by Site Debugging



ATLAS T1 - T1 Rates

- Take ATLAS as the example - highest inter-T1 rates due to multiple ESD copies
- Given spread of resources offered by T1s to ATLAS, requires "pairing of sites" to store ESD mirrors
- Reprocessing performed ~1 month after data taking with better calibrations & at end of year with better calibrations & algorithms
- Continuous or continual? (i.e. is network load constant or peaks+troughs?)

FZK (10%) + CCIN2P3 (13%)	BNL (22%)
CNAF (7%)	RAL (7%)
NIKHEF/SARA (13%)	TRIUMF (4%) + ASGC (8%)
PIC (4-6%)	NDGF (6%)

Meeting the LCG challenge

Example: Tier-2 individual transfer tests

Initial focus was on getting SRMs understood and deployed.....

	Receiving											
	RAL Tier-1	Lancaster	Manchester	Edinburgh	Glasgow	Birmingham	Oxford	Cam	Durham	QMUL	IC-HEP	RAL-PPD
RAL Tier-1		-800Mb/s	350Mb/s	156Mb/s	166 Mb/s	289 Mb/s	252 Mb/s			118 Mb/s	84Mb/s	397 Mb/s
Lancaster												
Manchester	150 Mb/s											
Edinburgh	440Mb/s											
Glasgow	331Mb/s											
Birmingham	461 Mb/s											
IC-HEP												
Oxford	456 Mb/s											
Cambridge	74 Mb/s											
Durham	193 Mb/s											
QMUL	172 Mb/s											
IC-HEP												
RAL-PPD	388 Mb/s											

- Big variation in what sites could achieve
 - Internal networking configuration issues
 - Site connectivity (and contention)
 - SRM setup and level of optimisation
- Rates to RAL were generally better than from RAL
 - Availability and setup of gridFTP servers at Tier-2s
 - SRM setup and level of optimisation
- Scheduling tests was not straightforward
 - Availability of local site staff
 - Status of hardware deployment
 - Availability of Tier-1
 - Need to avoid first tests during certain periods (local impacts)

Example rates from throughput tests









The Scoville Scale


- The **Scoville scale** is a measure of the *hotness* of a chilli pepper. These fruits of the Capsicum genus contain capsaicin, a chemical compound which stimulates thermoreceptor nerve endings in the tongue, and the number of Scoville heat units (SHU) indicates the amount of capsaicin present. Many hot sauces use their Scoville rating in advertising as a selling point.
- It is named after Wilbur Scoville, who developed the Scoville Organoleptic Test in 1912[1]. As originally devised, a solution of the pepper extract is diluted in sugar water until the 'heat' is no longer detectable to a panel of (usually five) tasters; the degree of dilution gives its measure on the Scoville scale. Thus a sweet pepper, containing no capsaicin at all, has a Scoville rating of zero, meaning no heat detectable even undiluted. Conversely, the hottest chiles, such as habaneros, have a rating of 300,000 or more, indicating that their extract has to be diluted 300,000-fold before the capsaicin present is undetectable. The greatest weakness of the Scoville Organoleptic Test is its imprecision, because it relies on human subjectivity.



Scoville Scale - cont.

Scoville rating			
No heat	TABASCO® <i>SWEET & Spicy</i> Pepper Sauce		100 - 600
600 - 800	TABASCO® Green Pepper Sauce		600 - 800
30,000 - 50,000	TABASCO® Garlic Pepper Sauce		1200 - 1800
100,000 - 325,000	TABASCO® Chipotle Pepper Sauce		2000 - 2500
15,000,000 - 16,000,000	TABASCO® Pepper Sauce		2500 - 5000
	TABASCO® Habanero Sauce		7000 - 8000

A Brief History...



- SC1 - December 2004: did **not** meet its goals of:
 - Stable running for ~2 weeks with 3 named Tier1 sites...
 - But more sites took part than foreseen...
- SC2 - April 2005: met **throughput** goals, but still
 - No reliable file transfer service (or real services in general...)
 - Very limited functionality / complexity
- SC3 "classic" - July 2005: added several components and **raised** bar
 - SRM interface to storage at all sites;
 - Reliable file transfer service using gLite FTS;
 - Disk - disk targets of 100MB/s per site; 60MB/s to tape
 - **Numerous issues seen - investigated and debugged over many months**
- SC3 "Casablanca edition" - Jan / Feb re-run
 - Showed that we had resolved many of the issues seen in July 2005
 - Network bottleneck at CERN, but most sites at or above targets
 - **Good step towards SC4(?)**

SC4 Schedule

- Disk - disk Tier0-Tier1 tests at the full nominal rate are scheduled for April. (from weekly con-call minutes...)
- The proposed schedule is as follows:
 - April 3rd (Monday) - April 13th (Thursday before Easter) - sustain an average daily rate to each Tier1 at or above the full nominal rate. (This is the week of the GDB + HEPiX + LHC OPN meeting in Rome...)
 - Any loss of average rate $\geq 10\%$ needs to be:
 - *accounted for (e.g. explanation / resolution in the operations log)*
 - *compensated for by a corresponding increase in rate in the following days*
 - We should continue to run at the same rates unattended over Easter weekend (14 - 16 April).
 - From Tuesday April 18th - Monday April 24th we should perform the tape tests at the rates in the table below.

Excellent report produced by IN2P3, covering disk and tape transfers, together with analysis of issues.

Successful demonstration of both disk and tape targets.

Concerns - April 25 MB

- Site maintenance and support coverage during throughput tests
 - After **5 attempts**, have to assume that this will not change in immediate future - better design and build the system to handle this
 - (This applies also to CERN)
- **Unplanned schedule changes, e.g. FZK missed disk - tape tests**
 - Some (successful) tests since ...
- Monitoring, showing the data rate to tape at remote sites and also of overall status of transfers
- Debugging of rates to specific sites [which has been done...]
- **Future throughput tests using more realistic scenarios**

SC4 - Remaining Challenges

- Full nominal rates *to tape* at all Tier1 sites - sustained!
- Proven ability to **ramp-up rapidly** to nominal rates at LHC start-of-run
- Proven ability to **recover** from backlogs
 - T1 unscheduled interruptions of 4 - 8 hours
 - T1 scheduled interruptions of 24 - 48 hours(!)
 - **💣 T0 unscheduled interruptions of 4 - 8 hours**
- Production scale & quality operations and monitoring
- **Monitoring and reporting is still a grey area**
 - I particularly like TRIUMF's and RAL's pages with lots of useful info!

- The Service Challenge programme this year must show that we can run **reliable services**
- Grid reliability is the **product of many components**
 - middleware, grid operations, computer centres,
- Target for September
 - 90% site availability
 - 90% user job success
- Requires a major effort by everyone to monitor, measure, debug

Too modest?
Too ambitious?

First data will arrive next year

NOT an option to get things going later



Production *WLCG* Services

(a) The building blocks

Grid Computing

- Today there are many definitions of *Grid computing*:
- The definitive definition of a Grid is provided by [\[1\]](#) Ian Foster in his article "[What is the Grid? A Three Point Checklist](#)" [\[2\]](#).
- The three points of this checklist are:
 - Computing resources are not administered centrally.
 - Open standards are used.
 - Non trivial quality of service is achieved.
- ... **Some sort of Distributed System at least...**
 - that crosses Management / Enterprise domains


The Worldwide LHC Computing Grid

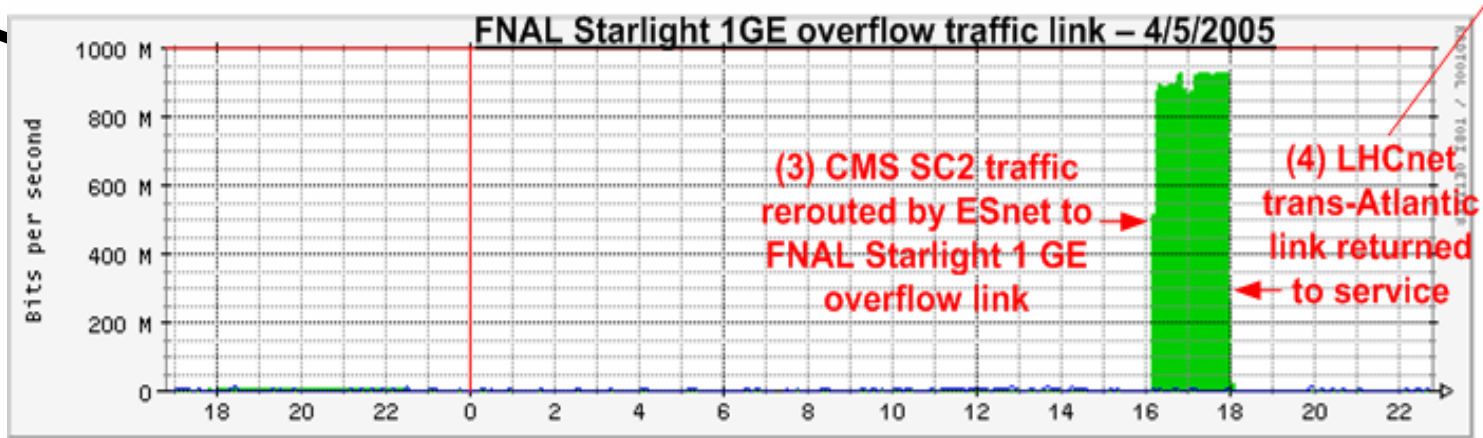
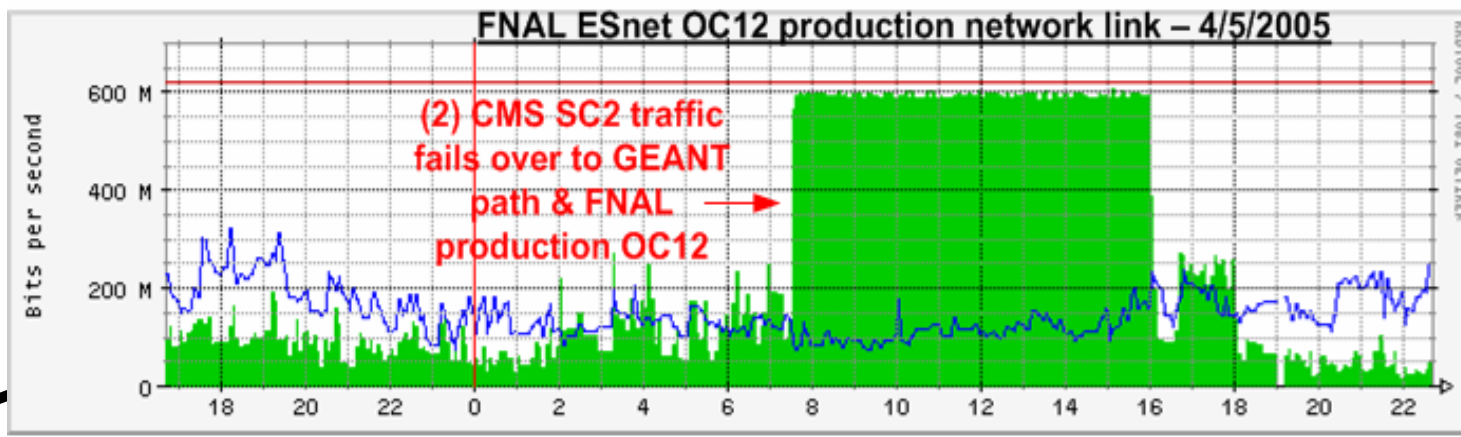
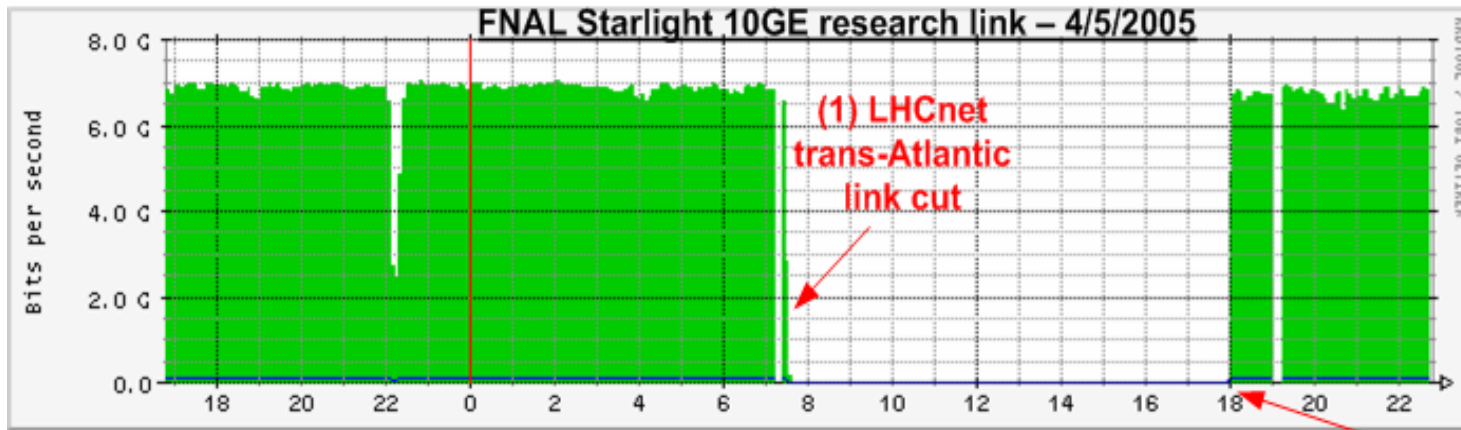
Introduction & Workshop Goals

Jamie Shiers, CERN, Geneva, Switzerland



Worldwide LHC Computing Grid
Distributed Production Environment for Physics data Processing





CHNET (France & Switzerland).

Production WLCG Services

(b) So What Happens *When*¹ it Doesn't Work?

¹ *Something* doesn't work all of the time

The 1st Law Of (Grid) Computing

- **Murphy's law** (also known as **Finagle's law** or **Sod's law**) is a popular **adage** in **Western culture**, which broadly states that things will go wrong in any given situation. "If there's more than one way to do a job, and one of those ways will result in disaster, then somebody will do it that way." It is most commonly formulated as "Anything that can go wrong will go wrong." In **American culture** the law was named after **Major Edward A. Murphy, Jr.**, a development **engineer** working for a brief time on **rocket sled** experiments done by the **United States Air Force** in 1949.
- ... first received public attention during a press conference ... it was that nobody had been severely injured during the rocket sled [of testing the **human** tolerance for **g-forces** during rapid deceleration.]. Stapp replied that it was because they took Murphy's Law under consideration.



Problem Response Time and Availability targets Tier-1 Centres

<i>Service</i>	<i>Maximum delay in responding to operational problems (hours)</i>			<i>Availability</i>
	<i>Service interruption</i>	<i>Degradation of the service</i>		
		<i>> 50%</i>	<i>> 20%</i>	
Acceptance of data from the Tier-0 Centre during accelerator operation	12	12	24	99%
Other essential services – prime service hours	2	2	4	98%
Other essential services – outside prime service hours	24	48	48	97%




Problem Response Time and Availability targets Tier-2 Centres

Service	<i>Maximum delay in responding to operational problems</i>		<i>availability</i>
	<i>Prime time</i>	<i>Other periods</i>	
End-user analysis facility	2 hours	72 hours	95%
Other services	12 hours	72 hours	95%

CERN (Tier0) MoU Commitments


Service	Maximum delay in responding to operational problems			Average availability ^[1] on an annual basis	
	DOWN	Degradation > 50%	Degradation > 20%	BEAM ON	BEAM OFF
Raw data recording	4 hours	6 hours	6 hours	99%	n/a
Event reconstruction / data distribution (beam ON)	6 hours	6 hours	12 hours	99%	n/a
Networking service to Tier-1 Centres (beam ON)	6 hours	6 hours	12 hours	99%	n/a
All other Tier-0 services	12 hours	24 hours	48 hours	98%	98%
All other services ^[2] - prime service hours ^[3]	1 hour	1 hour	4 hours	98%	98%
All other services - outside prime service hours	12 hours	24 hours	48 hours	97%	97%



Breakdown of a normal year

- From Chamonix XIV -

Service upgrade slots?




~ 140-160 days for physics per year
 Not forgetting ion and TOTEM operation
 Leaves ~ 100-120 days for proton luminosity running
 ? Efficiency for physics 50% ?
 ~ 50 days ~ 1200 h ~ $4 \cdot 10^6$ s of proton luminosity running / year




WLCG


- WLCG:

- A federation of fractal Grids...





EGEE - Close-up





So why this workshop?

- As we have seen, the WLCG “FZK federation” is one of the largest and most complex
 - Highest data rate; all experiments; many countries; no clearly established management / collaboration infrastructure covering entire federation
 - In EGEE-speak, covers DECH; Central Europe & Russia!
 - **Lonely Planet’s Central Europe guide covers Germany, Liechtenstein, Switzerland, Czech Republic, Slovakia, Poland, Hungary, Austria, and Slovenia**
 - (Asia-Pacific is also a large & complex region - primarily ATLAS & CMS - but fewer sites and lower data rates / volumes...)
- Need to foster existing relationships with the goal that the federation is self-monitoring & self-managing (cf GridPP) prior to first data
 - Next workshop prior to April GDB in Prague?
 - (IMHO, there are good reasons for these workshops to be at Tier2 sites and to move around...)
- Something to discuss in more detail tonight?





The LHC Machine

- Some clear indications regarding LHC startup schedule and operation are now available
 - Press release issued last Friday
 - Comparing *our* actual status with 'the plan', we (globally) are arguably **one year late!**
 - One site possibly two years late...
 - We still have an awful lot of work to do
- **Not the time to relax!**



Press Release - Extract

- CERN confirms LHC start-up for 2007
 - Geneva, 23 June 2006. First collisions in the ... LHC ... in **November 2007** said ... Lyn Evans at the 137th meeting of the CERN Council ...
 - A **two month** run in 2007, with beams colliding at an energy of **0.9 TeV** will allow the LHC accelerator and detector teams to run-in their equipment ready for a full **14 TeV** energy run to start in **Spring 2008**
 - *Service Challenge '07?*
 - The schedule announced today ensures the fastest route to a high-energy physics run with substantial quantities of data in 2008, while optimising the commissioning schedules for both the accelerator and the detectors that will study its particle collisions. It foresees closing the LHC's 27 km ring in August 2007 for equipment commissioning. Two months of running, starting in November 2007, will allow the accelerator and detector teams to test their equipment with low-energy beams. After a winter shutdown in which commissioning will continue without beam, the high-energy run will begin. Data collection will continue until a pre-determined amount of data has been accumulated, allowing the experimental collaborations to announce their first results.



Important Milestones


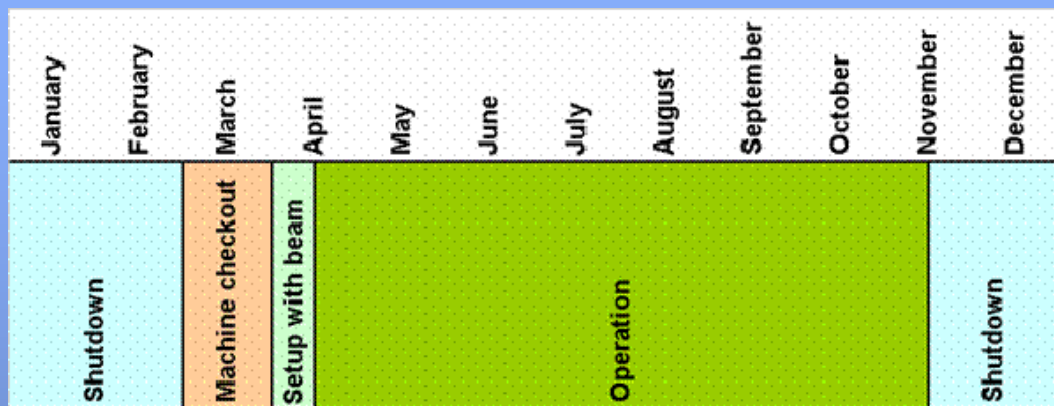
Last magnet delivered	October 2006
Last magnet tested	December 2006
Last magnet installed	March 2007
Machine closed	August 2007
First collisions	November 2007



Machine Commissioning

- Sectors 7-8 and 8-1 will be fully commissioned up to 7 TeV in 2006-2007
- The other sectors will be commissioned up to the field needed for de-Gaussing (1.2 TeV)
- Initial operation will be at 900 GeV (CM) with a static machine (no ramp, no squeeze) to debug machine and detectors and to give a significant sample of W and Z
- Full commissioning up to 7 TeV will be done in the winter 2008 shutdown


Breakdown of a normal year



~ 140-160 days for physics per year
 Not forgetting ion and TOTEM operation
 Leaves ~ 100-120 days for proton luminosity running
 ? Efficiency for physics 50% ?
 ~ 1200 h or ~ 4 10⁶ s of proton luminosity running / year

Conclusions

- All key objectives have been reached for the end of 2005 and installation is now proceeding smoothly.
 - Three quarters of the machine has been liberated for magnet installation and interconnect work is proceeding in 2 octants in parallel. Magnet installation is now steady at 25/wk . Installation will finish end March 2007. The machine will be closed in August 2007.
 - Every effort is being made to establish colliding beams before the end of 2007 at reduced energy. The full commissioning up to 7 TeV will be done during the 2008 winter shutdown ready for a Physics run at full energy in 2008.



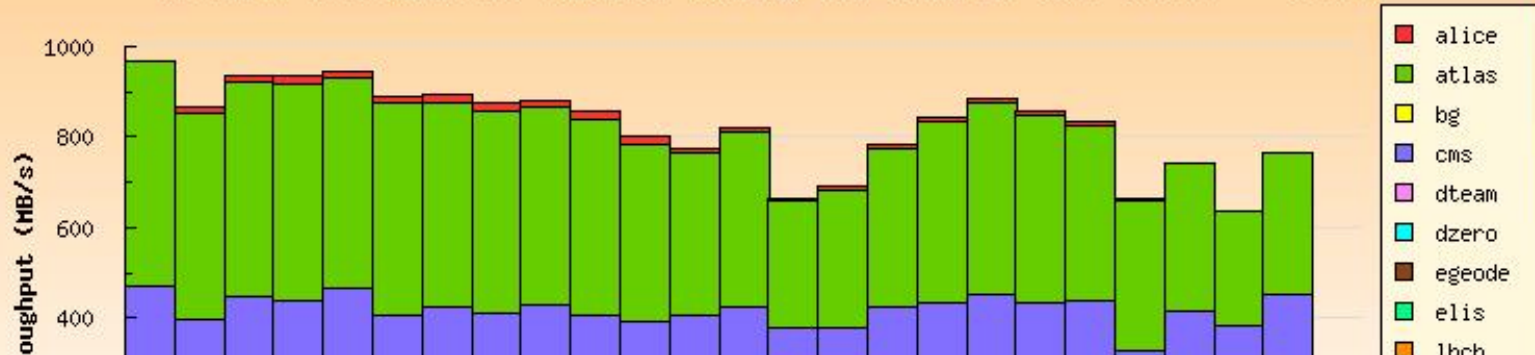
Machine checkout and later

	56	67	78	81	12	23	34	45
Oct						Minimum HWC		
	Operations testing				Minimum HWC	Operations testing		
Nov	Full Machine Checkout (TI8 & TI2, Access, Vacuum, Equipment Tests, Cycle and Set, BIC and INB)							
	Beam Commissioning to 450GeV 16days estimated, 60%efficiency assumed							
Dec	Engineering run (Collisions at 450GeV + Ramp Commissioning)							



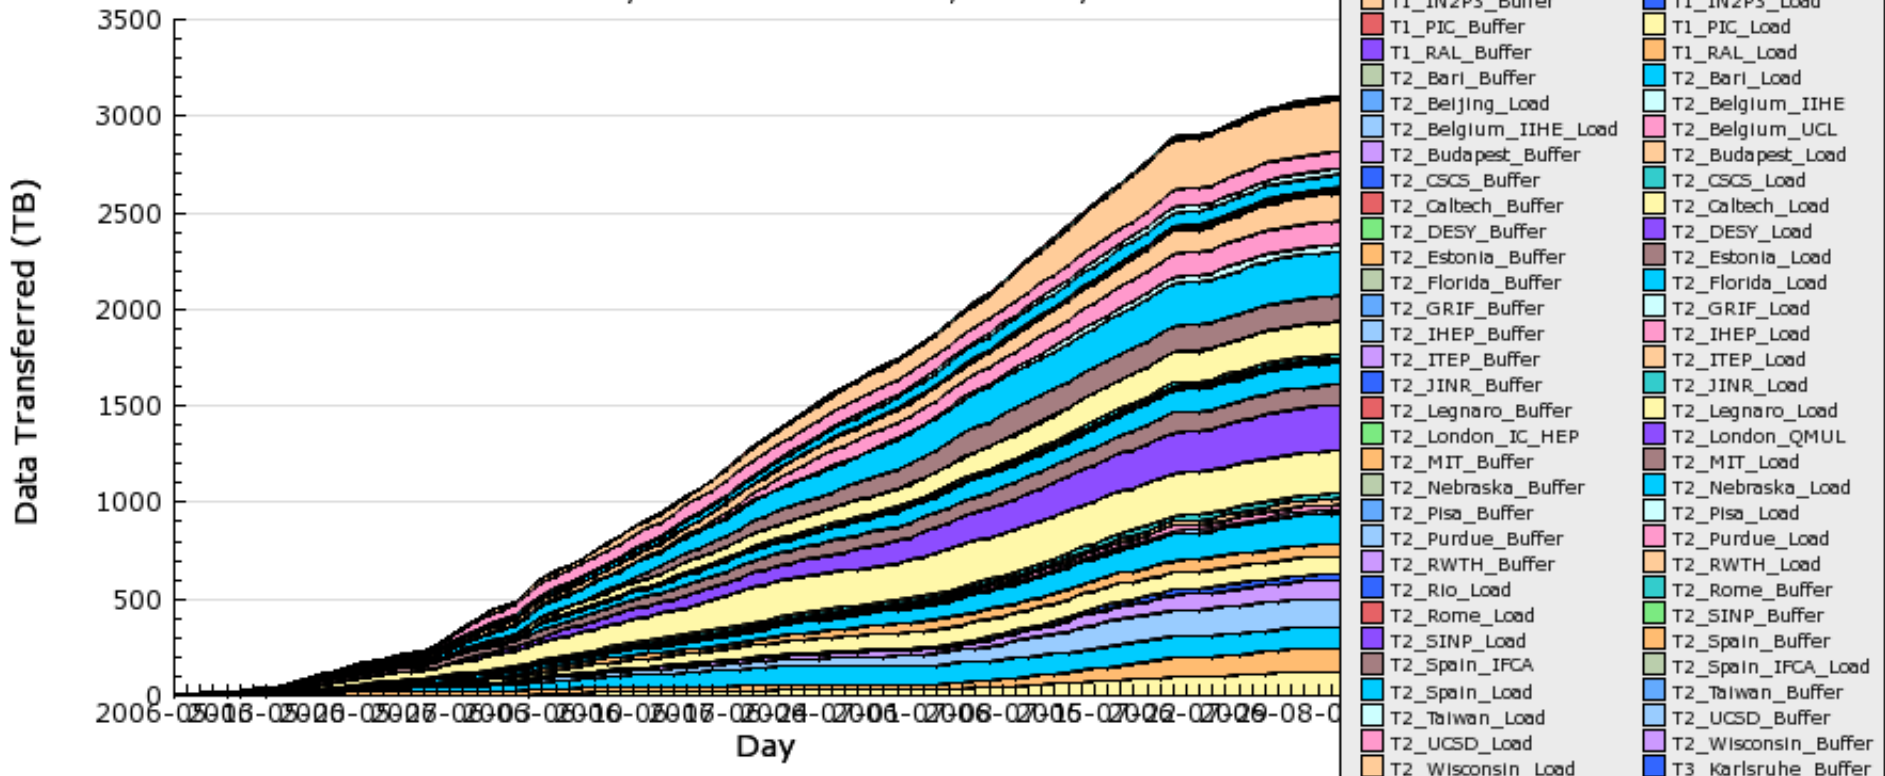
• Exp
cha

V0-wise Averaged Throughput during the last 24 hrs (06/08 - 07/08)



PhEDEx SC4 Data Transfers By Destinations matching

Last 91 Days at 2006-08-11 09:08, last entry 2006-08-11 GMT





Plans Prior to First Collisions

- Between now and first collisions these activities will continue, progressively ramping up in scope and scale
- Still significant work to involve ~100 Tier2s in a distributed, reliable service
- Still much work to do to attain data rates for prolonged periods (weeks) including recovery from site failure
 - power, cooling, service issues



And Beyond...

- First collisions LHC expected November 2007
 - These will be at 'low' energy - 450 GeV per beam
 - Main target will be understanding detectors, trigger and offline software
 - 'Re-discover' existing physics - excellent for calibration!
 - ☛ Data rates will be full nominal values! (Machine efficiency?)
- First full energy run in 2008: 7 + 7 TeV
 - Physics discovery run!
 - Heavy Ions in 2009? Data export schedule?
- Typically takes ~years to fully understand detector and software chain
 - Much of the initial 'analysis' will be done starting from RAW/ESD datasets
 - Big impact on network load - larger datasets, transferred more frequently
 - ☛ Potential mismatch with 'steady-state' planning?
 - Much larger initial bandwidth requirement (but do you really believe it will go down?)
 - Those sites that have it will be more 'competitive' (and vice-versa...)
- Rate calculations have overhead for recovering backlogs due to down-time
 - But not for recovery from human and / or software error!
 - ☛ e.g. bug in alignment / calibration / selection / classification code -> junk data!



WLCG Tier2 Workshop

- This was the 2nd SC4 Workshop with primary focus on “new Tier2s”
 - i.e. those not (fully) involved in SC activities so far
 - 1-2 people obviously didn't know this from responses
- Complementary to Mumbai “Tier1” workshop
- Attempted to get Tier2s heavily involved in:
 - Planning the workshop (content)
 - The event itself
 - Chairing sessions, giving presentations and tutorials, ...
 - Less successful in this than hoped - room for improvement!



Workshop Feedback

- >160 people registered and participated!
 - This is very large for a workshop - about same as Mumbai
 - Some comments related directly to this
- Requests for more tutorials, particularly “hands-on”
- Requests for more direct Tier2 involvement
 - Feedback sessions, planning concrete actions etc.
- **Your active help in preparing / defining future events will be much appreciated**
 - Please not just the usual suspects...




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?

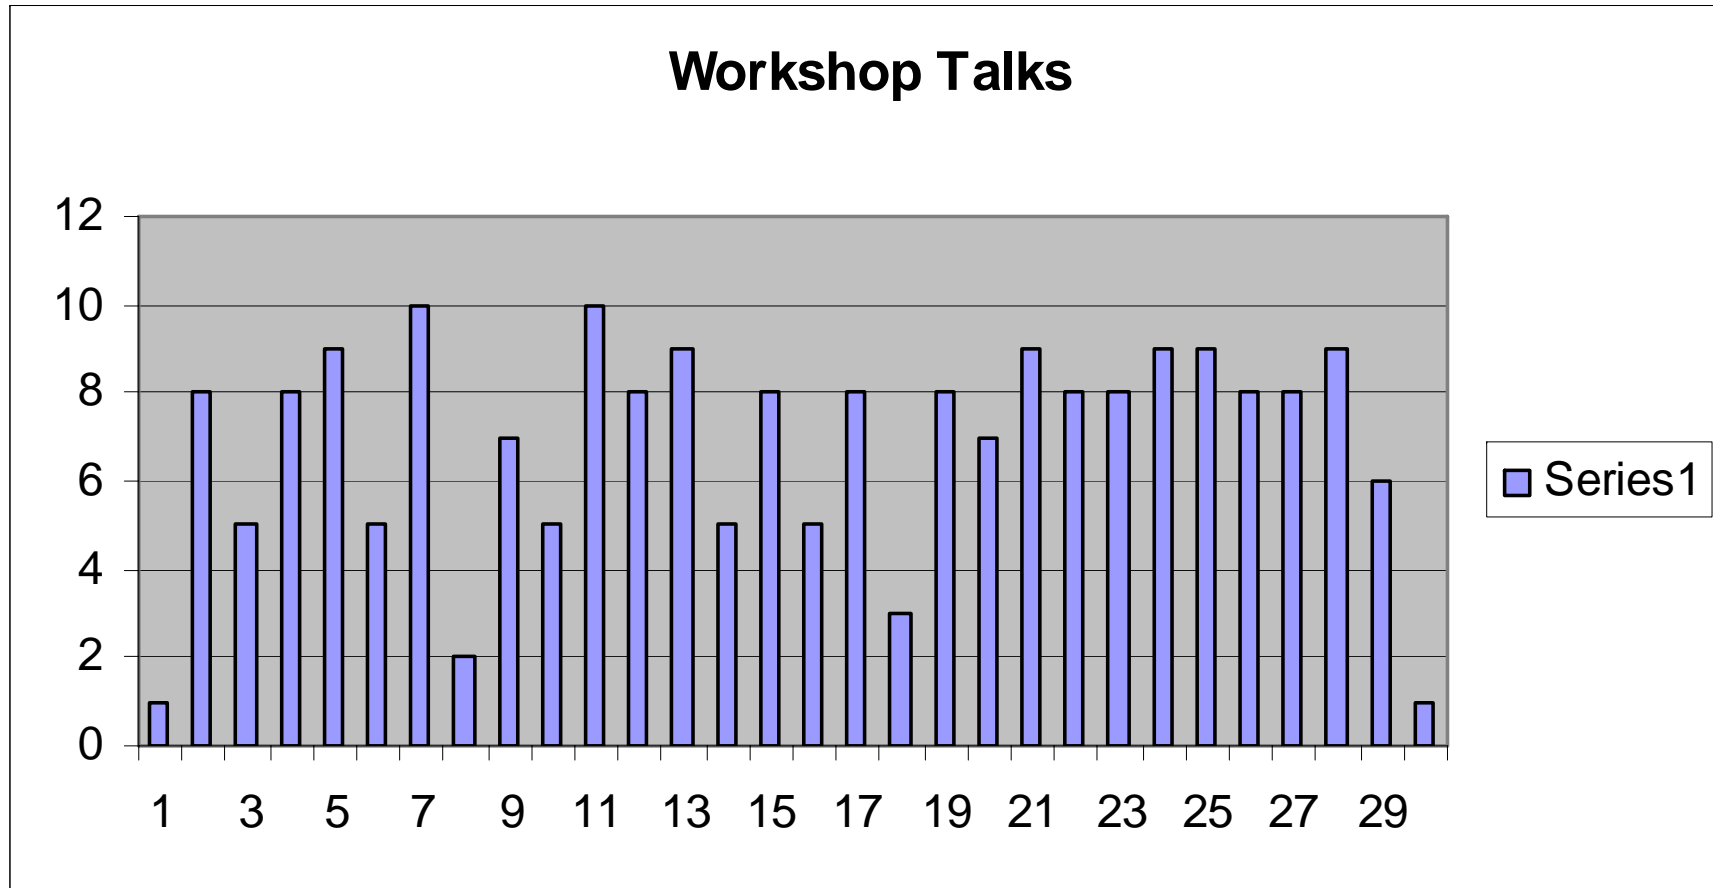


Tutorial Rating - 10=best





Workshop Rating





Future Workshops

- Suggest 'regional' workshops to analyse results of experiment activities in SC4 during Q3/Q4 this year - important to drill down to details / problems / solutions
- A 'global' workshop early 2007 focussing on experiment plans for 2007
- Another just prior to CHEP
- Given the size of the WLCG collaboration, these events are likely to be **BIG!**
- Few suitable meeting rooms at CERN - need to plan well in advance
- Something like 2 per year? Co-locate with CHEP / other events where possible?
- Quite a few comments suggesting HEPiX-like issues. Co-locate with HEPiX?
- A one-size-fits-all event is probably not going to succeed...

Communication



- ❑ U.S. Resource Allocation Committee (RAC)
- ❑ Tier1-Tier2 meetings every 2 months (1-2 days)
 - ❑ Rotate among sites
 - ❑ Attended by management, site coordinators, and most importantly technical staff from all sites
 - ❑ Site requirements are discussed and decided jointly
 - At UC - <http://indico.cern.ch/conferenceDisplay.py?confId=a062200>
 - At HU- <http://indico.cern.ch/conferenceDisplay.py?confId=4897>
 - ❑ Dec 8 – next meeting at UTA
- ❑ Site trouble ticket system
 - ❑ <http://www.usatlas.bnl.gov/twiki/bin/view/Support/RTUsersGuide>
- ❑ Panda mailing list, various Savannah, various hypernews
- ❑ Regular tutorials (examples: production, DDM, user...)



Jan 23-25 2007, CERN

- This workshop will cover: For each LHC experiment, detailed plans / requirements / timescales for 2007 activities.
- Exactly what (technical detail) is required where (sites by name), by which date, coordination & follow-up, responsables, contacts, etc etc. There will also be an initial session covering the status of the various software / middleware and outlook. **Do we also cover operations / support?**
- From feedback received so far, looks like an explicit interactive planning session would be a good idea
 - Dates: 23 January 2007 09:00 to 25 January 2007 18:00 (whole week now booked)
 - Location: *CERN*, Room: [Main auditorium](#)
- **Do we need tutorials? If so, what topics? Who can help?**
- **Other ideas? Expert panel Q&A? International advisory panel?**



Agenda Ideas - 1/2

- 9:00 Introduction
 - (Neil Geddes, WLCG Collaboration Board chair) 15'
- 9:15 Physics at the LHC
 - (A CERN star, e.g. John Ellis, or a Tier2 person who could e.g. summarise the Krakow workshop?) 45'
- 10:00 Machine status 30 (45')
- 10:30 coffee
- 11:00 Status of the experiments 60'

- Break

- 13:00 What's new in WLCG 45'
- 13:45 'Database-related' services (including those using caching techniques at T>0) 45'
- 14:30 SRM 2.2 services 45'
- 15:15 gLite 3.x 45'

N.B. we have main auditorium until 16:00 except Wednesday (18:00)

Additional rooms for BOFs, break-out sessions etc.



Agenda Ideas - 2/2

- 09:00 - 13:00 - 2007 Experiment Dress Rehearsals
 - 4 x 1 hour (4 x 45' + coffee??)
- 14:00 - 15:45 Service Planning for 2007
 - (start with a proposal then try to discuss details)
- Federation reports + site feedback 11-12 x 30'
 - (adjust to federation size?)
- Tours, BOFs, Tutorials, "ask the expert" sessions etc.
- Some explicit sessions about running stable, reliable services?
 - e.g. from Tier1s, but also 'mega-Tier2s', e.g. SLAC, DESY, ...
- Some 'HEPiX'-style sessions? (Which could include the above...)
- Tier2 (and Tier1!) involvement as much as possible!
- A WLCG Collaboration Board meeting will take place, probably Thursday afternoon

Day 2

Day 3



Sep 1-2 2007, Victoria, BC

- Workshop focussing on service needs for initial data taking: commissioning, calibration and alignment, early physics. Target audience: all active sites plus experiments
- *We start with a detailed update on the schedule and operation of the accelerator for 2007/2008, followed by similar sessions from each experiment.*
- *We wrap-up with a session on operations and support, leaving a slot for parallel sessions (e.g. 'regional' meetings, such as GridPP etc.) before the foreseen social event on Sunday evening.*
- Dates: 1-2 September 2007
- Location: *Victoria, BC, Canada*, co-located with CHEP 2007



Topics to be addressed (sites)

- Grid services offered by this site
 - BDII, CE, SRM-enabled SE, ...
 - LFC, VO-boxes (both local (e.g. ALICE) and remote (e.g. ATLAS))
 - 3D services (SQUID, local MySQL or other DB services)
 - O/S; middleware; hardware (CPU, disk, tape) status and outlook
 - Support and operations staff + expertise and outlook
 - Issues & Concerns
- Participation to date in SC4
 - Activities; results; issues
- Participation in remainder of 2006
 - and beyond??



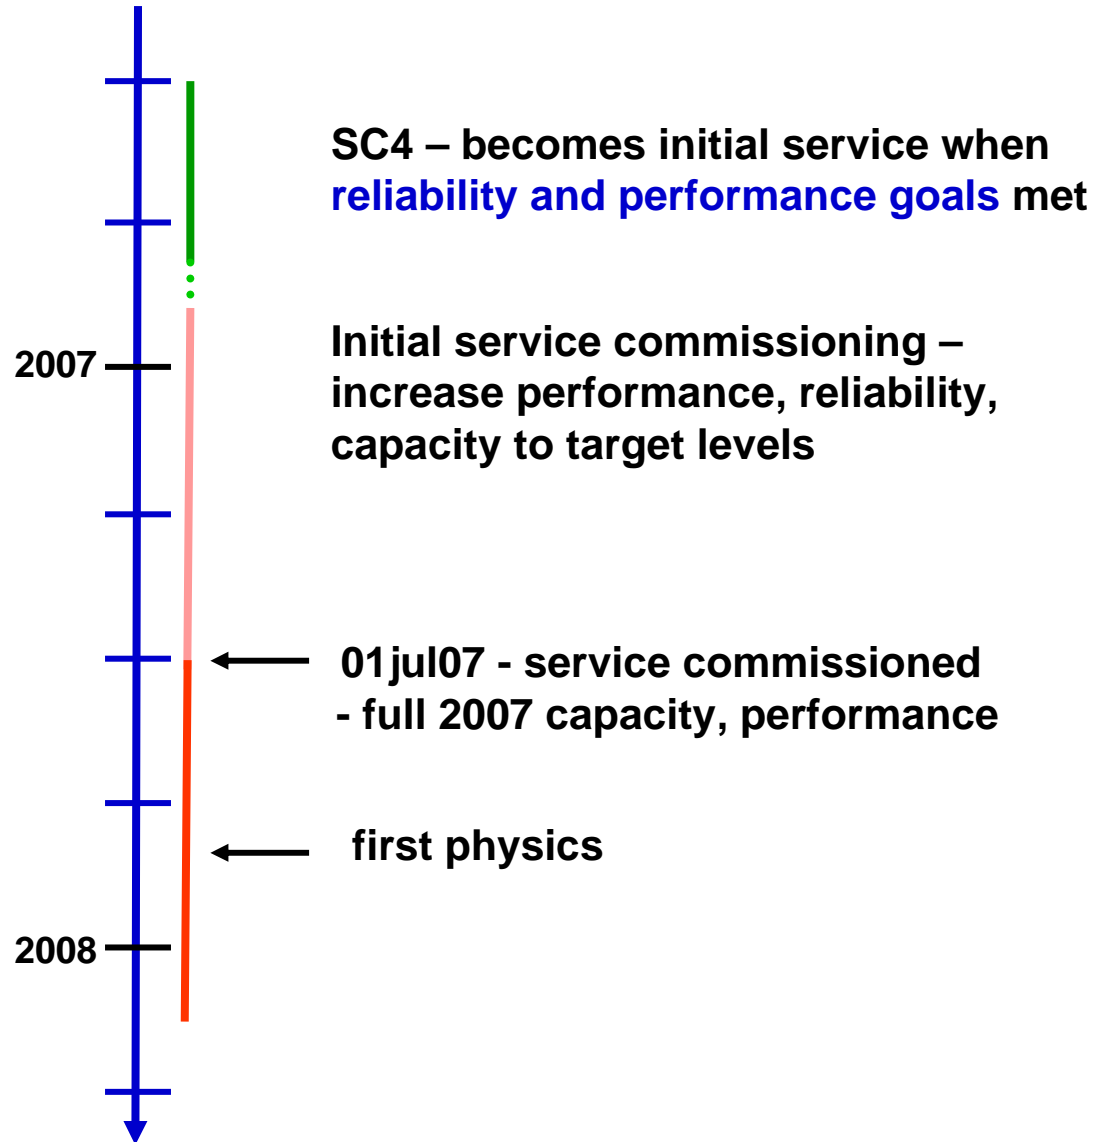
Topics to be addressed (experiments)

- *Goals of the SC4 activities (up-front metrics)*
- *Sites involved; requirements by site (e.g. services to be deployed, disk / tape storage areas, clean-up policy)*
- *Results obtained wrt metrics and foreseen timetable*
- *Problems encountered; how were they solved; how could things be improved for the future*
- ...



Commissioning Schedule 2006

Introduce residual services
Full FTS T1/T2 T1/T1 services
3D SRM v2.2 VOMS roles





Summary & Conclusions

- Deploying a Worldwide Production Grid is not without its challenges
- Much has been accomplished; much still outstanding
- My two top issues?
 - Collaboration & communication at such a scale requires significant and constant effort
 - We are not yet at the level that this is just basic infrastructure
 - “Design for failure” - i.e. assume that things don’t work, rather than hope that they always do!
 - A lesson from our “founding fathers” - the creators of the Internet?

