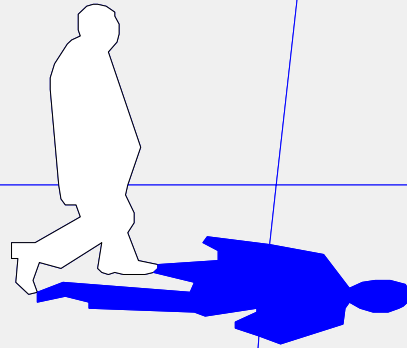


# CERN, June 2006

## The Pilot WLCG Service: Last steps before full production



- i) Issues Related to Running Production Services
- ii) Operational Concerns Seen from Five (5) Service Challenges
- iii) Roadmap for rest of 2006, early 2007

Jamie Shiers, CERN

# Abstract

- 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...**

- 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

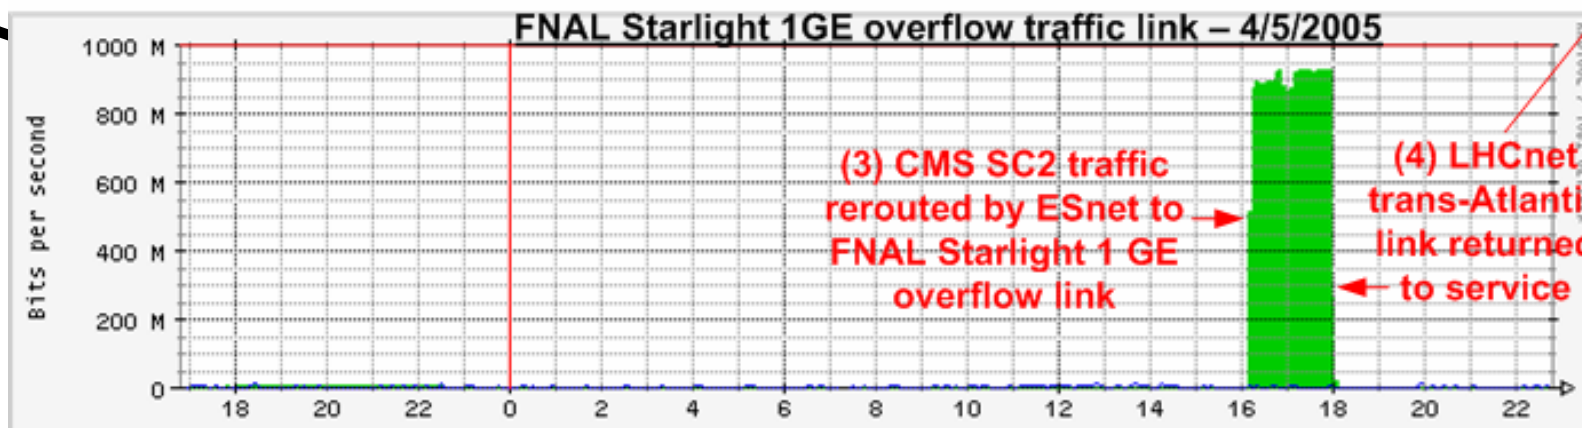
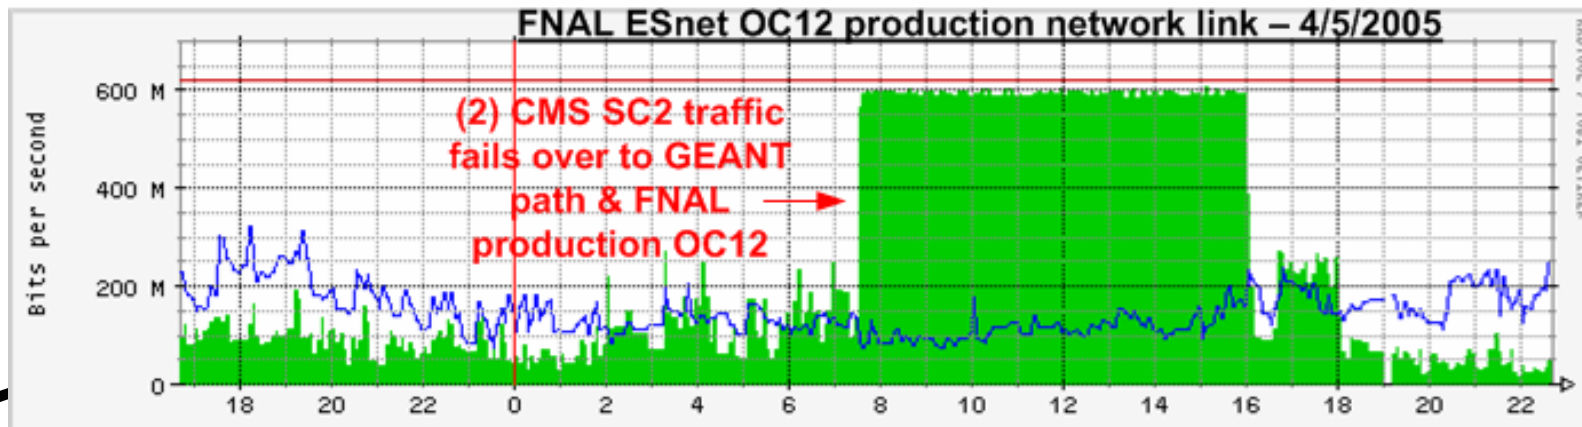
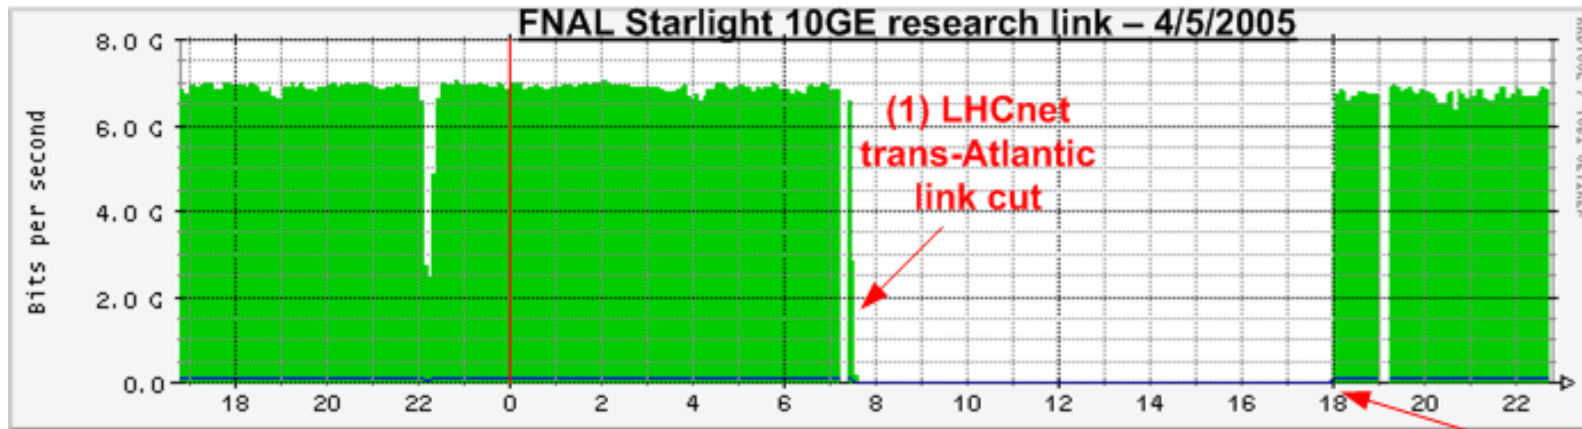
# Grid Computing

- 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

# Distributed Systems...

- "A distributed system is one in which the failure of a computer you **didn't even know existed** can render your own computer unusable."

Leslie Lamport



# Production WLCG Services

(b) So What Happens *When*<sup>1</sup> it Doesn't Work?

<sup>1</sup>*Something* doesn't work all of the time



# The 1<sup>st</sup> 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>&gt; 50%</i>	<i>&gt; 20%</i>	
<b>Acceptance of data from the Tier-0 Centre during accelerator operation</b>	<b>12</b>	<b>12</b>	<b>24</b>	<b>99%</b>
<b>Other essential services – prime service hours</b>	<b>2</b>	<b>2</b>	<b>4</b>	<b>98%</b>
<b>Other essential services – outside prime service hours</b>	<b>24</b>	<b>48</b>	<b>48</b>	<b>97%</b>

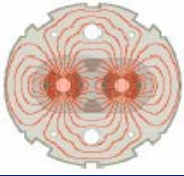


## Problem Response Time and Availability targets Tier-2 Centres

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

# CERN (Tier0) MoU Commitments

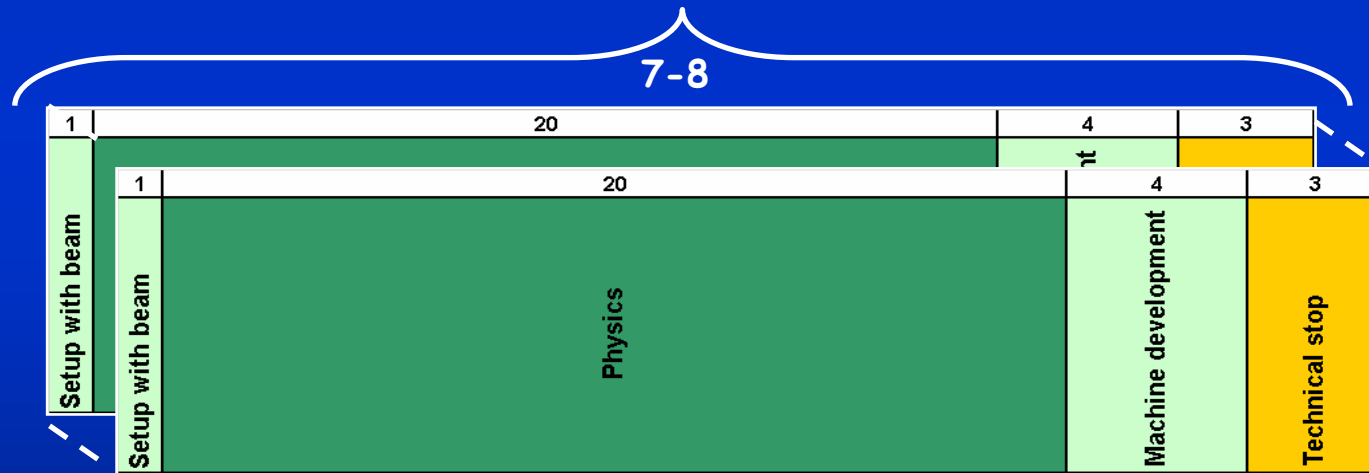

Service	Maximum delay in responding to operational problems			Average availability <sup>[1]</sup> on an annual basis	
	DOWN	Degradation > 50%	Degradation > 20%	BEAM ON	BEAM OFF
Raw data recording	4 hours	6 hours	6 hours	99%	n/a
Event reconstruction / data distribution (beam ON)	6 hours	6 hours	12 hours	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 <sup>[2]</sup> - prime service hours <sup>[3]</sup>	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 Operations

Beyond EGEE / OSG

# Introduction

- Whilst *WLCG* is built upon existing *Grid* infrastructures - and must use procedures / tools etc at the underlying level as much as possible, there are aspects of the *WLCG* service that require additional procedures / agreements etc.
- Two real-life examples follow
- These could eventually be built into procedures of the underlying *Grids*...
- ... But we need it now...

# Scheduled Interventions

- Need procedures for announcing and handling scheduled interventions
- The WLCG Management Board has agreed the following:
  - Interruptions of up to **4 hours** must be announced at least **one day** in advance;
  - Interruptions greater than **4 hours** but less than **12** must be announced at the weekly operations meeting **prior** to the event;
  - Interruptions greater than **12 hours** must be announced at the operations meeting of the preceding **week**.
- This is particularly important for services which affect outside users (e.g. CASTOR at CERN!)
  - LHCb are also keen that batch queues are appropriately closed / drained
    - (A revised version is attached to the agenda pending MB approval)



# Site Offline Procedure (or Emergency Contact)

- So what happens when a site goes offline?
- **Follow operations procedures**
- But these are on the Web...
- So the person who lives closest drives home and uses his/her private Internet connection
- Or we have a procedure...
- **And don't tell me it'll never happen (again...)**

# Pragmatic Solution

- I have compiled a table of contacts (e-mail, phone, mobiles) from replies from site contacts / GOCDDB
- I have printed it, stuck it on my door and in the corridor in B28
- I have loaded all numbers into my mobile phone **but I haven't called them**
- This goes beyond GOCDDB in any case
  - CERN MOD, SMOD, GMOD, central computer operator (5011), ...
  - Control room number at some sites ...
- OK - its not "nice", but the next time Tony Cass calls to tell me he's about to shutdown the Computer Centre, at least I'll have a better answer than
  - "Romain thinks he might have Steve Traylen's number at home"

TRIUMF	<a href="mailto:sc@triumf.ca">sc@triumf.ca</a>	+1 604 222 7333 (main TRIUMF control room)
--------	--	--

## Operations and Service Contacts

See the list of [Grid Operations Contacts](#) as well as the [Grid Operations Meetings and shift rota](#).

See also the [Resource Centers](#) page on the [CIC portal](#).

The CERN IT Manager-on-Duty can be contacted via [mod@cern.ch](mailto:mod@cern.ch) (phone 163013 / +41764873013). See also the CERN [IT Service Status](#) page.

The IT Service Manager on Duty (SMOD) [Web](#) and [e-mail address](#).

IT service managers can be contacted through [it-dep-service-managers@cern.ch](mailto:it-dep-service-managers@cern.ch)

## CIC-On-Duty

[project-eu-egee-sa1-cic-on-duty@cern.ch](mailto:project-eu-egee-sa1-cic-on-duty@cern.ch)

## VO Managers

Alice	<a href="mailto:project-lcq-vo-alice-admin@cern.ch">project-lcq-vo-alice-admin@cern.ch</a>
Atlas	<a href="mailto:project-lcq-vo-atlas-admin@cern.ch">project-lcq-vo-atlas-admin@cern.ch</a>
CMS	<a href="mailto:project-lcq-vo-cms-admin@cern.ch">project-lcq-vo-cms-admin@cern.ch</a>
LHCb	<a href="mailto:project-lcq-vo-lhcb-admin@cern.ch">project-lcq-vo-lhcb-admin@cern.ch</a>

## Site Offline Procedure

1. If a site goes completely off-line (e.g. major power failure) and they want to let the world know, then they should contact their Regional Operations Centre by phone and ask them to make the broadcast.
2. If the site *is also* the ROC, then the ROC should phone one of the other ROCs and ask them to make the broadcast.
3. We already have a backup grid operator-on-duty team each week, so if the primary one goes off-line, then they call the backu

# 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
  - 2<sup>nd</sup> 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

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
FNAL, USA	-	-	28%	-	200
BNL, USA	-	22%	-	-	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
<b>TOTAL</b>					<b>1600MB/s</b>

# 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)

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

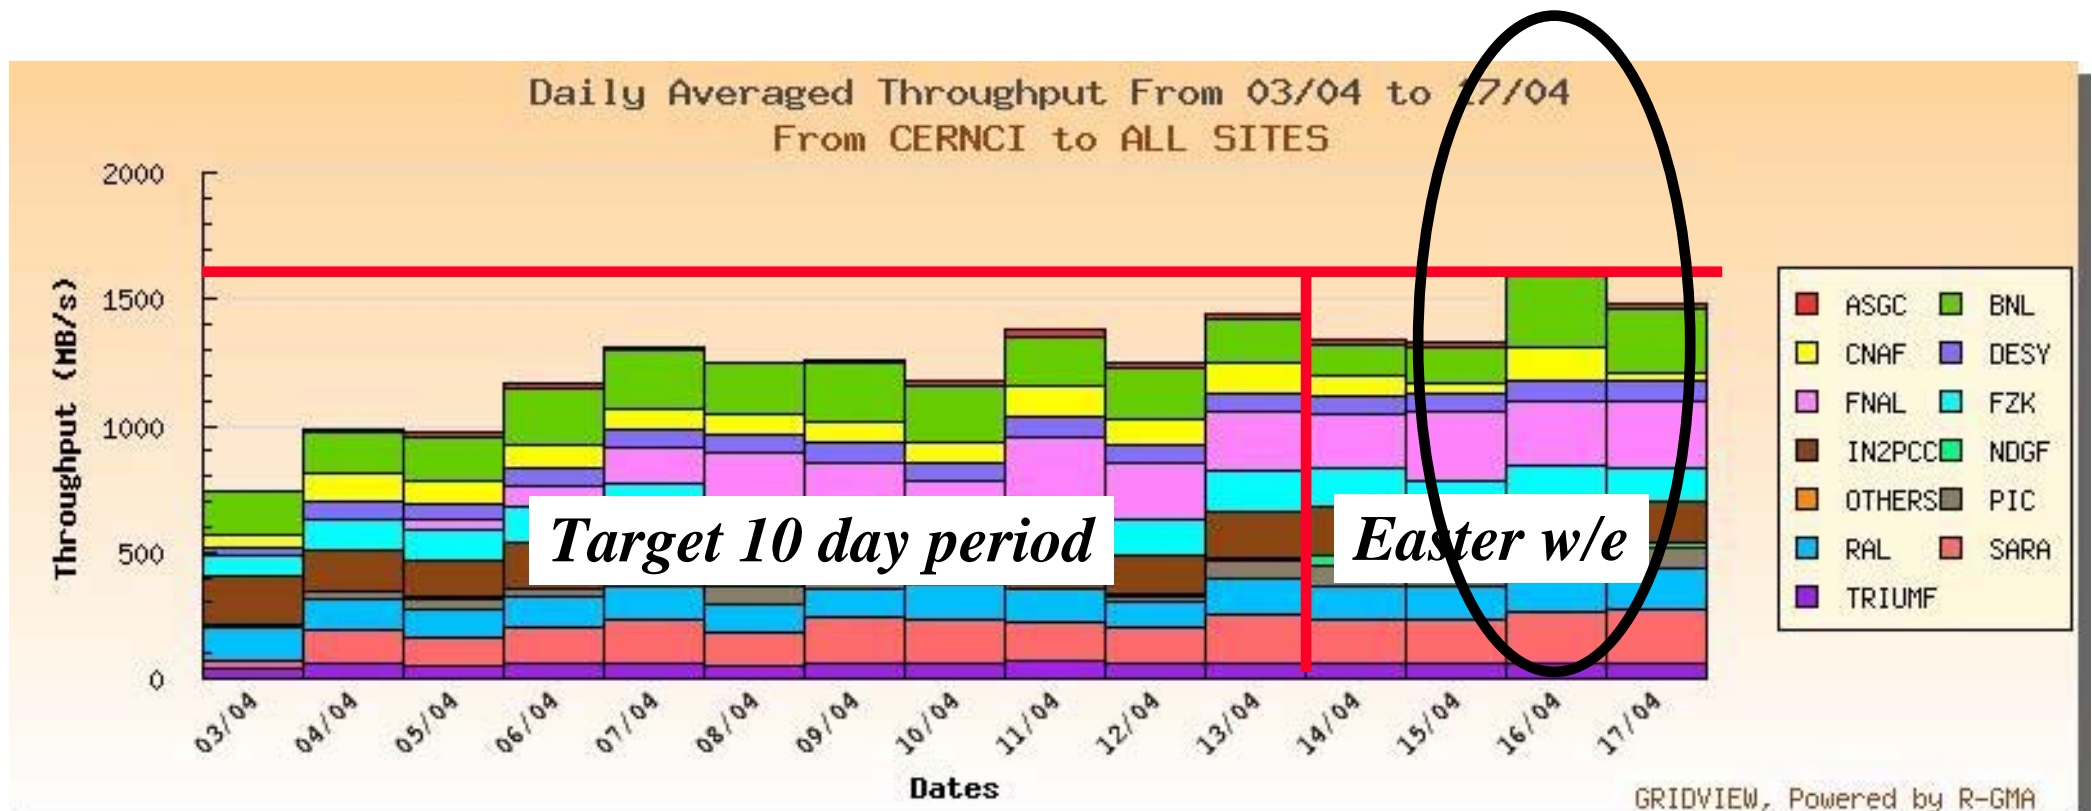
**Successful demonstration of both disk and tape targets.**

- Dropped based on experience of first week of disk - disk tests

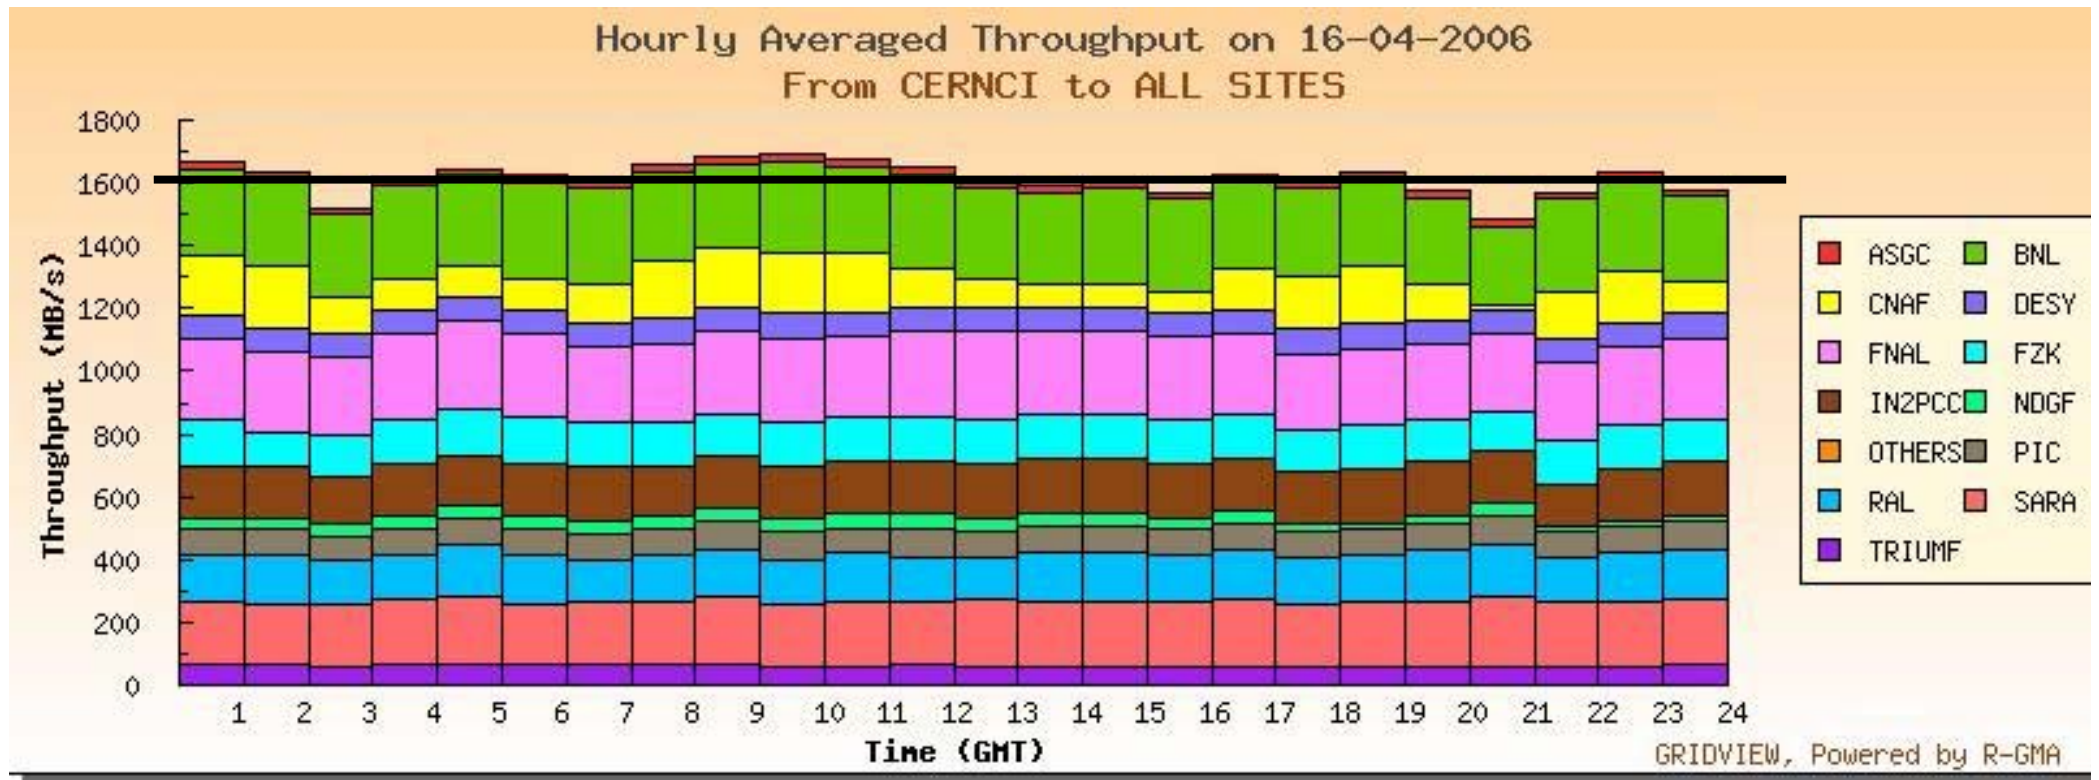


# SC4 T0-T1: Results

- **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 for 16/4/2006*

# 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!

# Disk - Tape Targets

- Realisation during SC4 that we were simply “turning up all the knobs” in an attempt to meet site & global targets
  - Not necessarily under conditions representative of LHC data taking
- Could continue in this way for future disk - tape tests *but*
- Recommend moving to realistic conditions as soon as possible
  - At least some components of distributed storage system not necessarily optimised for this use case (focus was on local use cases...)
  - ☛ If we **do** need another round of upgrades, know that this can take 6+ months!
- Proposal: benefit from ATLAS (and other?) Tier0+Tier1 export tests in June + Service Challenge Technical meeting (also June)
  - Work on operational issues can (must) continue in parallel
  - As must deployment / commissioning of new tape sub-systems at the sites
  - e.g. milestone on sites to perform disk - tape transfers at > (>>) nominal rates?
- This will provide some feedback by late June / early July
  - Input to further tests performed over the summer

# Combined Tier0 + Tier1 Export Rates

<i>Centre</i>	<i>ATLAS</i>	<i>CMS*</i>	<i>LHCb<sup>+</sup></i>	<i>ALICE</i>	<i>Combined (ex-ALICE)</i>	<i>Nominal</i>
ASGC	60.0	10	-	-	70	100
CNAF	59.0	25	23	? (20%)	108	200
PIC	48.6	30	23	-	103	100
IN2P3	90.2	15	23	? (20%)	138	200
GridKA	74.6	15	23	? (20%)	95	200
RAL	59.0	10	23	? (10%)	118	150
BNL	196.8	-	-	-	200	200
TRIUMF	47.6	-	-	-	50	50
SARA	87.6	-	23	-	113	150
NDGF	48.6	-	-	-	50	50
FNAL	-	50	-	-	50	200
US site	-	-	-	? 20%)		
Totals				300	~1150	1600

- \* CMS target rates double by end of year
- + Mumbai rates - scheduled delayed by ~1 month (start July)
- ? ALICE rates - 300MB/s aggregate (Heavy Ion running)

# SC4 - Successes & Remaining Work

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

# SC4 Conclusions



- We have demonstrated - through the SC3 re-run and more convincingly through SC4 - that we can send data to the Tier1 sites at the required rates for extended periods
  - Disk - tape rates are reasonably encouraging but still require full deployment of production tape solutions across all sites to meet targets
- **Demonstrations of the needed data rates corresponding to experiment transfer patterns must now be proven**
- As well as an acceptable - and affordable - service level
- Moving from dTeam to experiment transfers will hopefully also help drive the migration to full production service
  - Rather than the current 'best' (where 'best' is clearly +ve!) effort



# SC4 - Meeting with LHCC Referees

- Following presentation of SC4 status to LHCC referees, I was asked to write a report (originally confidential to Management Board) summarising issues & concerns
- **I did not want to do this!**
- This report started with some (uncontested) observations
- Made some recommendations
  - Somewhat luke-warm reception to some of these at MB
  - ... but I still believe that they make sense! (So I'll show them anyway...)
- Rated site-readiness according to a few simple metrics...
- **We are not ready yet!**

# Disclaimer

- Please find a report reviewing Site Monitoring and Operation in SC4 attached to the following page:
- <https://twiki.cern.ch/twiki/bin/view/LCG/ForManagementBoard>
- (It is not attached to the MB agenda and/or Wiki as it should be considered confidential to MB members).
- **Two seconds later it was attached to the agenda, so no longer confidential...**
- *In the table below tentative service levels are given, based on the experience in April 2006. It is proposed that each site checks these assessments and provides corrections as appropriate and that these are then reviewed on a site-by-site basis.*
- **(By definition, T0-T1 transfers involve source&sink)**

# Observations

1. Several sites took a **long time** to ramp up to the performance levels required, despite having taken part in a similar test during January. This appears to indicate that the data transfer service is **not yet integrated** in the normal site operation;
2. **Monitoring** of data rates to **tape** at the Tier1 sites is not provided at many of the sites, neither 'real-time' nor after-the-event reporting. This is considered to be a **major hole** in offering services at the required level for LHC data taking;
3. Sites **regularly fail to detect problems** with transfers terminating at that site - these are often picked up by manual monitoring of the transfers at the CERN end. This manual monitoring has been provided on an exceptional basis 16 x 7 during much of SC4 - **this is not sustainable in the medium to long term**;
4. Service interventions of **some hours** up to **two days** during the service challenges have occurred **regularly** and are expected to be a part of life, i.e. it must be assumed that these will occur during LHC data taking and thus sufficient capacity to recover rapidly from backlogs from corresponding scheduled downtimes needs to be demonstrated;
5. **Reporting** of operational problems - both on a daily and weekly basis - is weak and inconsistent. **In order to run an effective distributed service these aspects must be improved considerably in the immediate future.**

# Recommendations

- All sites should provide a schedule for implementing monitoring of data rates to input disk buffer and to tape. This monitoring information should be published so that it can be viewed by the COD, the service support teams and the corresponding VO support teams. (See June internal review of LCG Services.)
- Sites should provide a schedule for implementing monitoring of the basic services involved in acceptance of data from the Tier0. This includes the local hardware infrastructure as well as the data management and relevant grid services, and should provide alarms as necessary to initiate corrective action. (See June internal review of LCG Services.)
- ✓ A procedure for announcing scheduled interventions has been approved by the Management Board (main points next)
- All sites should maintain a daily operational log - visible to the partners listed above - and submit a weekly report covering all main operational issues to the weekly operations hand-over meeting. It is essential that these logs report issues in a complete and open way - including reporting of human errors - and are not 'sanitised'. Representation at the weekly meeting on a regular basis is also required.
- ✓ Recovery from scheduled downtimes of individual Tier1 sites for both short (~4 hour) and long (~48 hour) interventions at full nominal data rates needs to be demonstrated. Recovery from scheduled downtimes of the Tier0 - and thus affecting transfers to all Tier1s - up to a minimum of 8 hours must also be demonstrated. A plan for demonstrating this capability should be developed in the Service Coordination meeting before the end of May.
- ✓ Continuous low-priority transfers between the Tier0 and Tier1s must take place to exercise the service permanently and to iron out the remaining service issues. These transfers need to be run as part of the service, with production-level monitoring, alarms and procedures, and not as a "special effort" by individuals.

# Site Readiness - Metrics

- Ability to ramp-up to nominal data rates - see results of SC4 disk - disk transfers [2];
- Stability of transfer services - see table 1 below;
- Submission of weekly operations report (with appropriate reporting level);
- Attendance at weekly operations meeting;
- Implementation of site monitoring and daily operations log;
- Handling of scheduled and unscheduled interventions with respect to procedure proposed to LCG Management Board.

# Site Readiness

Site	Ramp-up	Stability	Weekly Report	Weekly Meeting	Monitoring Operations /	Interventions	Average
CERN	2-3	2	3	1	2	1	2
ASGC	4	4	2	3	4	3	3
TRIUMF	1	1	4	2	1-2	1	2
FNAL	2	3	4	1	2	3	2.5
BNL	2	1-2	4	1	2	2	2
NDGF	4	4	4	4	4	2	3.5
PIC	2	3	3	1	4	3	3
RAL	2	2	1-2	1	2	2	2
SARA	2	2	3	2	3	3	2.5
CNAF	3	3	1	2	3	3	2.5
IN2P3	2	2	4	2	2	2	2.5
FZK	3	3	2	2	3	3	3

- 1 - always meets targets
- 2 - usually meets targets
- 3 - sometimes meets targets
- 4 - rarely meets targets

# Site Readiness

Site	Ramp-up	Stability	Weekly Report	Weekly Meeting	Monitoring Operations /	Interventions	Average
CERN	2-3	2	3	1	2	1	2
ASGC	4	4	2	3	4	3	3
TRIUMF	1	1	4	2	1-2	1	2
FNAL	2	3	4	1	2	3	2.5
BNL	2	1-2	4	1	2	2	2
NDGF	4	4	4	4	4	2	3.5
PIC	2	3	3	1	4	3	3
RAL	2	2	1-2	1	2	2	2
SARA	2	2	3	2	3	3	2.5
CNAF	3	3	1	2	3	3	2.5
IN2P3	2	2	4	2	2	2	2.5
FZK	3	3	2	2	3	3	3

- 1 - always meets targets
- 2 - usually meets targets
- 3 - sometimes meets targets
- 4 - rarely meets targets

# SC4 Disk - Disk Average Daily Rates

Site/Date	3	4	5	6	7	8	9	10	11	12	13	14	15	16	Av. (Nom.)
<b>ASGC</b>	0	7	23	23	0	0	12	22	33	25	26	21	19	22	<b>17(100)</b>
<b>TRIUMF</b>	44	42	55	62	56	55	61	62	69	63	63	60	60	62	<b>58(50)</b>
<b>FNAL</b>	0	0	38	80	145	247	198	168	289	224	159	218	269	258	<b>164(200)</b>
<b>BNL</b>	170	103	173	218	227	205	239	220	199	204	168	122	139	284	<b>191(200)</b>
<b>NDGF</b>	0	0	0	0	0	14	0	0	0	0	14	38	32	35	<b>10(50)</b>
<b>PIC</b>	0	18	41	22	58	75	80	49	0	24	72	76	75	84	<b>48(100<sup>[1]</sup>)</b>
<b>RAL</b>	129	86	117	128	137	109	117	137	124	106	142	139	131	151	<b>125(150)</b>
<b>SARA</b>	30	78	106	140	176	130	179	173	158	135	190	170	175	206	<b>146(150)</b>
<b>CNAF</b>	55	71	92	95	83	80	81	82	121	96	123	77	44	132	<b>88(200)</b>
<b>IN2P3</b>	200	114	148	179	193	137	182	86	133	157	183	193	167	166	<b>160(200)</b>
<b>FZK</b>	81	80	118	142	140	127	38	97	174	141	159	152	144	139	<b>124(200)</b>

<sup>[1]</sup> The agreed target for PIC is 60MB/s, pending the availability of their 10Gb/s link to CERN.



### All Wikis

[ACPP](#)  
[ADCgroup](#)  
[AISgroup](#)  
[ALICE](#)  
[ALPHA](#)  
[AliceSPD](#)  
[AthenaFCalTBAna](#)  
[Atlas](#)  
[CERNSearch](#)  
[CMS](#)  
[CS](#)  
[Controls](#)  
[DESgroup](#)  
[Dbaservices](#)  
[DefaultWeb](#)  
[EGEE](#)  
[ELFms](#)  
[ETICS](#)  
[EgeePtf](#)  
[FIOgroup](#)  
[HCC](#)  
[HRonDemand](#)  
[Know](#)  
[Lar](#)  
[LCG](#)  
[LCGAAWorkbook](#)  
[LHCatHome](#)  
[LHCOPN](#)  
[LHCb](#)  
[LHCgas](#)  
[LcgProcurementInfo](#)  
[LinuxSupport](#)  
[Main](#)  
[PHESS](#)  
[PSSGroup](#)  
[Plugins](#)  
[SPI](#)  
[SRMDev](#)  
[Sandbox](#)  
[SupComp05](#)  
[TWiki](#)

[TWikiGuest](#)

## Week two (April 10 on)

- Week 2 average to sum of Tier1 sites is 1262 MB/s - 79% of the target.

Site	Disk-Disk	Week1 Average	Week2 Average	Apr10	Apr11	April12	April13	April14	April15	April16	April17
TRIUMF	50	54	63	62	69	63	63	60	60	62	63
BNL	200	191	199	220	199	204	168	122	139	284	257
FNAL	200	101	231	168	289	224	159	218	269	258	261
PIC	60	49	78 (5 days)	49	-	24	72	76	75	84	82
<a href="#">RAL</a>	150	118	136	137	124	106	142	139	131	151	160
SARA	150	120	178	173	158	135	190	170	175	206	213
IN2P3	200	165	157	86	133	157	183	193	167	166	167
FZK	200	104	142	97	174	141	159	152	144	139	130
CNAF	200	80	88	82	121	96	123	77	44	132	32
ASGC	100		24	22	33	25	26	21	19	22	24
NDGF	50		28 (5 days)	-	-	-	14	38	32	35	20
DESY	60	70	74	71	77	69	72	76	73	76	76
<b>TOTAL (T1s)</b>	<b>1600</b>			1096	1300	1175	1046	1266	1255	1539	1409

## Week one (April 3 on)

Site	Disk-Disk	Apr3	Apr4	Apr5	Apr6	Apr7	Apr8	Apr9	Weekly average	Average from startup	Target
TRIUMF	50	44	42	55	62	56	55	61	54	54 (>100%)	50
BNL	200	170	103	173	218	227	205	239	191	191 (>95%)	200
FNAL	200	-	-	38	80	145	247	198	101	141 (>70%)	200
PIC	60	-	18	41	22	58	75	80	49	42 (70%)	60
<a href="#">RAL</a>	150	129	86	117	128	137	109	117	118	118 (~80%)	150
SARA	150	30	78	106	140	176	130	179	120	120 (80%)	150
IN2P3	200	200	114	148	179	193	137	182	165	165 (>80%)	200
FZK	200	81	80	118	142	140	127	38	104	104	200
CNAF	200	55	71	92	95	83	80	81	80	80	200
ASGC	100	-	7	23	23	-	-	12			100
NDGF	50	-	-	-	-	-	14	-			50
DESY	60	-	68	63	75	74	68	74		70	60
<b>TOTAL (T1s)</b>	<b>1600</b>	<b>709</b>	<b>599</b>	<b>911</b>	<b>1089</b>	<b>1215</b>	<b>1179</b>	<b>1187</b>		<b>984 (61.5% of target)</b>	

# Site Readiness - Summary

- I believe that these subjective metrics paint a fairly realistic picture
- The ATLAS and other Challenges will provide more data points
- I know the support of multiple VOs, standard Tier1 responsibilities, plus others taken up by individual sites / projects represent significant effort
- But at some stage we have to adapt the plan to reality
- If a small site is late things can probably be accommodated
- 💣 If a major site is late we have a major problem

# Site Readiness - Next Steps

- Discussion at MB was to repeat review but with rotating reviewers
- Clear candidate for next phase would be ATLAS T0-T1 transfers
- As this involves all Tier1s except FNAL, suggestion is that FNAL nominate a co-reviewer
  - e.g. Ian Fisk + Harry Renshall
- Metrics to be established in advance and agreed by MB and Tier1s
- (This test also involves a strong Tier0 component which may have to be factored out)
- Possible metrics next:

# June Readiness Review

- Readiness for start date
  - Date at which required information was communicated
- T0-T1 transfer rates as daily average 100% of target
  - List the daily rate, the total average, histogram the distribution
  - Separate disk and tape contributions
  - Ramp-up efficiency (# hours, # days)
- MoU targets for pseudo accelerator operation
  - Service availability, time to intervene
- Problems and their resolution (using standard channels)
  - # tickets, details
- Site report / analysis
  - Sites own report of the 'run', similar to that produced by IN2P3

# **WLCG Service**

Experiment Production Activities During WLCG Pilot

Aka SC4 Service Phase June - September Inclusive

# Overview

- All 4 LHC experiments will run **major production exercises** during WLCG pilot / SC4 Service Phase
- These will test all aspects of the respective Computing Models plus stress Site Readiness to run (collectively) full production services
- These plans have been assembled from the material presented at the Mumbai workshop, with follow-up by Harry Renshall with each experiment, together with input from Bernd Panzer (TO) and the Pre-production team, and summarised on the SC4 planning page.
- We have also held a number of meetings with representatives from all experiments to confirm that we have all the necessary input (all activities: PPS, SC, Tier0, ...) and to spot possible clashes in schedules and / or resource requirements. (See "LCG Resource Scheduling Meetings" under LCG Service Coordination Meetings).
  - fyi; the LCG Service Coordination Meetings (LCGSCM) focus on the CERN component of the service; we also held a WLCGSCM at CERN last December.
- The conclusions of these meetings has been presented to the weekly operations meetings and the WLCG Management Board in written form (documents, presentations)
  - See for example these points on the MB agenda page for May 24 2006
- **The Service Challenge Technical meeting (21 June IT amphi) will list the exact requirements by VO and site with timetable, contact details etc.**

# DTEAM Activities

- Background disk-disk transfers from the Tier0 to all Tier1s will start from June 1st.
  - These transfers will continue - but with low priority - until further notice (it is assumed until the end of SC4) to debug site monitoring, operational procedures and the ability to ramp-up to full nominal rates rapidly (a matter of hours, not days).
  - These transfers will use the disk end-points established for the April SC4 tests.
  - Once these transfers have satisfied the above requirements, a schedule for ramping to full nominal disk - tape rates will be established.
  - The current resources available at CERN for DTEAM only permit transfers up to 800MB/s and thus can be used to test ramp-up and stability, but not to drive all sites at their full nominal rates for pp running.
  - All sites (Tier0 + Tier1s) are expected to operate the required services (as already established for SC4 throughput transfers) in full production mode.
- **(Transfer) SERVICE COORDINATOR**

# ATLAS

- ATLAS will start a major exercise on June 19th. This exercise is described in more detail in <https://uimon.cern.ch/twiki/bin/view/Atlas/DDMSc4>, and is scheduled to run for 3 weeks.
- However, preparation for this challenge has already started and will ramp-up in the coming weeks.
- **That is, the basic requisites must be met prior to that time, to allow for preparation and testing before the official starting date of the challenge.**
- The sites in question will be ramped up in phases - the exact schedule is still to be defined.
- The target data rates that should be supported from CERN to each Tier1 supporting ATLAS are given in the table below.
- **40% of these data rates must be written to tape, the remainder to disk.**
- It is a requirement that the tapes in question are at least unloaded having been written.
- Both disk and tape data maybe recycled after 24 hours.
- **Possible targets: 4 / 8 / all Tier1s meet (75-100%) of nominal rates for 7 days**



# ATLAS Rates by Site

Centre	ATLAS SC4	Nominal (pp) MB/s (all experiments)
ASGC	60.0	100
CNAF	59.0	200
PIC	48.6	100
IN2P3	90.2	200
GridKA	74.6	200
RAL	59.0	150
BNL	196.8	200
TRIUMF	47.6	50
SARA	87.6	150
NDGF	48.6	50
FNAL	-	200

~25MB/s to tape, remainder to disk

# ATLAS Preparations

Site	TB disk (24 hr lifetime)	TB disk+tape (24 hr lifetime)	Data Rate (MB/s)
ASGC	3	2	60
BNL	10	7	196
CNAF	2	3	69
FZK	4	3	74
IN2P3	6	4	90
NDGF	3	2	49
NIKHEF/SARA	6	4	88
PIC	3	2	48
RAL	3	2	59
TRIUMF	3	2	48

# ATLAS ramp-up - request

- Overall goals: raw data to the Atlas T1 sites at an aggregate of 320 MB/sec, ESD data at 250 MB/sec and AOD data at 200 MB/sec.
  - The distribution over sites is close to the agreed MoU shares.
  - The raw data should be written to tape and the tapes ejected at some point. The ESD and AOD data should be written to disk only.
- Both the tapes and disk can be recycled after some hours (we suggest 24) as the objective is to simulate the permanent storage of these data.)
- It is intended to ramp up these transfers starting now at about 25% of the total, increasing to 50% during the week of 5 to 11 June and 75% during the week of 12 to 18 June.
- For each Atlas T1 site we would like to know SRM end points for the disk only data and for the disk backed up to tape (or that will become backed up to tape).
  - These should be for Atlas data only, at least for the period of the tests.
- During the 3 weeks from 19 June the target is to have a period of at least 7 contiguous days of stable running at the full rates.
- Sites can organise recycling of disk and tape as they wish but it would be good to have buffers of at least 3 days to allow for any unattended weekend operation.

# ATLAS T2 Requirements

- (ATLAS) expects that some Tier-2s will participate on a voluntary basis.
- There are no particular requirements on the Tier-2s, besides having a SRM-based Storage Element.
- An FTS channel to and from the associated Tier-1 should be set up on the Tier-1 FTS server and tested (under an ATLAS account).
- The nominal rate to a Tier-2 is 20 MB/s. We ask that they keep the data for 24 hours so, this means that the SE should have a minimum capacity of 2 TB.
- For support, we ask that there is someone knowledgeable of the SE installation that is available during office hours to help to debug problems with data transfer.
- Don't need to install any part of DDM/DQ2 at the Tier-2. The control on "which data goes to which site" will be of the responsibility of the Tier-0 operation team so, the people at the Tier-2 sites will not have to use or deal with DQ2.
- See <https://twiki.cern.ch/twiki/bin/view/Atlas/ATLASServiceChallenges>

# CMS

- The CMS plans for June include 20 MB/sec aggregate Phedex (FTS) traffic to/from temporary disk at each Tier 1 (SC3 functionality re-run) and the ability to run 25000 jobs/day at end of June.
- This activity will continue through-out the remainder of WLCG pilot / SC4 service phase (see [Wiki](#) for more information)
- It will be followed by a MAJOR activity in the - similar (AFAIK) in scope / size to the June ATLAS tests - [CSA06](#)
- The lessons learnt from the ATLAS tests should feedback - inter alia - into the services and perhaps also CSA06 itself (the model - not scope or goals)

# CMS CSA06

A 50-100 million event exercise to test the workflow and dataflow associated with the data handling and data access model of CMS

- Receive from HLT (previously simulated) events with online tag
- Prompt reconstruction at Tier-0, including determination and application of calibration constants
- Streaming into physics datasets (5-7)
- Local creation of AOD
- Distribution of AOD to all participating Tier-1s
- Distribution of some FEVT to participating Tier-1s
- Calibration jobs on FEVT at some Tier-1s
- Physics jobs on AOD at some Tier-1s
- Skim jobs at some Tier-1s with data propagated to Tier-2s
- Physics jobs on skimmed data at some Tier-2s

# ALICE

- In conjunction with on-going transfers driven by the other experiments, ALICE will begin to transfer data at 300MB/s out of CERN - corresponding to heavy-ion data taking conditions (1.25GB/s during data taking but spread over the four months shutdown, i.e.  $1.25/4=300\text{MB/s}$ ).
  - The Tier1 sites involved are CNAF (20%), CCIN2P3 (20%), GridKA (20%), SARA (10%), RAL (10%), US (one centre) (20%).
  - Time of the exercise - July 2006, duration of exercise - 3 weeks (including set-up and debugging), the transfer type is disk-tape.
  - Goal of exercise: test of service stability and integration with ALICE FTD (File Transfer Daemon).
- **Primary objective: 7 days of sustained transfer to all T1s.**
- As a follow-up of this exercise, ALICE will test a synchronous transfer of data from CERN (after first pass reconstruction at T0), coupled with a second pass reconstruction at T1. The data rates, necessary production and storage capacity to be specified later.
  - More details are given in the ALICE documents attached to the MB agenda of 30th May 2006.

Last updated 12 June to add scheduled dates of **24 July - 6 August** for T0 to T1 data export tests.

# LHCb

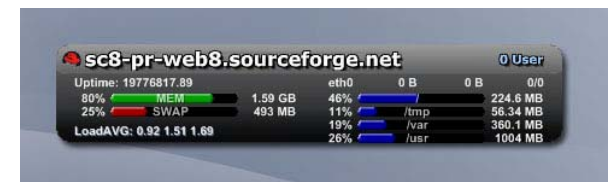
- Starting from July LHCb will **distribute** "raw" data from CERN and store data on **tape** at each Tier1.
- CPU resources are required for the **reconstruction** and **stripping** of these data, as well as at Tier1s for **MC event generation**.
- The exact resource requirements by site and time profile are provided in the updated LHCb spreadsheet that can be found on <https://twiki.cern.ch/twiki/bin/view/LCG/SC4ExperimentPlans> under
  - "LHCb plans".
- (Detailed breakdown of resource requirements in Spreadsheet)



# The Dashboard



- Sounds like a conventional problem for a 'dashboard'
- But there is not one single viewpoint...
  - Funding agency - how well are the resources provided being used?
  - VO manager - how well is my production proceeding?
  - Site administrator - are my services up and running? MoU targets?
  - Operations team - are there any alarms?
  - LHCC referee - how is the overall preparation progressing? Areas of concern?
  - ...
- Nevertheless, much of the information that would need to be collected is common...
- So separate the collection from presentation (views...)
- As well as the discussion on metrics...



# Summary of Key Issues

- There are clearly many areas where a great deal still remains to be done, **including:**
- Getting stable, reliable, data transfers up to **full** rates
- Identifying and testing all **other** data transfer needs
- Understanding experiments' data **placement** policy
- Bringing services up to **required** level - functionality, availability, (operations, support, upgrade schedule, ...)
- Delivery and commissioning of needed **resources**
- Enabling remaining sites to rapidly and effectively **participate**
- **Accurate and concise monitoring, reporting and accounting**
- Documentation, training, information dissemination...

# Monitoring of Data Management

- GridView is far from sufficient in terms of data management monitoring
- We cannot really tell what is going on:
  - Globally;
  - At individual sites.
- This is an area where we urgently need to improve things
- Service Challenge Throughput tests are one thing...
- **But providing a reliable service for data distribution during accelerator operation is yet another...**
- **Cannot just 'go away' for the weekend; staffing; coverage etc.**

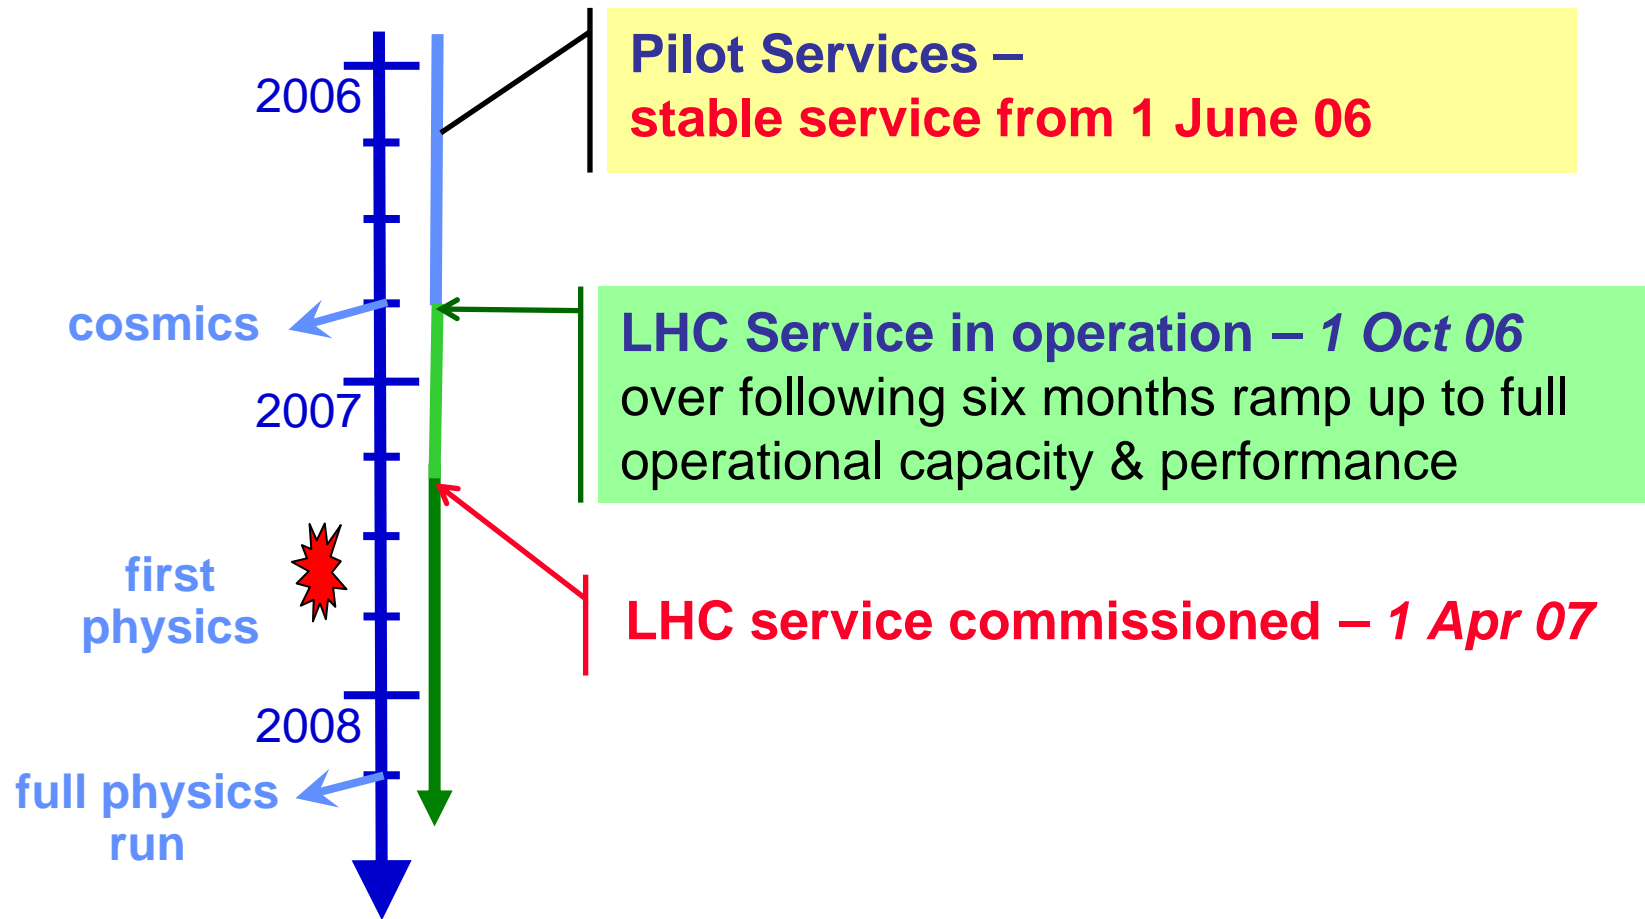
# The Carminati Maxim

- What is not there for SC4 (aka WLCG pilot) will not be there for WLCG production (and vice-versa)
- This means:
- We have to be using - consistantly, systematically, daily, ALWAYS - all of the agreed tools and procedures that have been put in place by Grid projects such as EGEE, OSG, ...
- **BY USING THEM WE WILL FIND - AND FIX - THE HOLES**
- If we continue to use - or invent more - stop-gap solutions, then these will continue well into production, resulting in confusion, duplication of effort, waste of time, ...
- (None of which can we afford)

# Issues & Concerns

- **Operations:** we have to be much more formal and systematic about logging and reporting. Much of the activity e.g. on the Service Challenge throughput phases - including major service interventions - has not been systematically reported by all sites. Nor do sites regularly and systematically participate. Network operations needs to be included (site; global)
- **Support:** move to GGUS as primary (sole?) entry point advancing well. Need to continue efforts in this direction and ensure that support teams behind are correctly staffed and trained.
- **Monitoring and Accounting:** we are well behind what is desirable here. Many activities - need better coordination and direction. The recently available **SAM** monitoring shows how valuable this is! (LFC, FTS etc.)
- **Services:** all of the above need to be in place by June 1<sup>st</sup>(!) and fully debugged through WLCG pilot phase. In conjunction with the specific services, based on Grid Middleware, Data Management products (CASTOR, dCache, ... ) etc.

# WLCG Service Deadlines



# SC4 - the Pilot LHC Service from June 2006

A stable service on which experiments can make a full demonstration of experiment offline chain

- DAQ → Tier-0 → Tier-1  
data recording, calibration, reconstruction
- Offline analysis - Tier-1 ↔ Tier-2 data exchange  
simulation, batch and end-user analysis

And sites can test their operational readiness

- Service metrics → MoU service levels
- Grid services
- Mass storage services, including magnetic tape

Extension to most Tier-2 sites

Evolution of SC3 rather than lots of new functionality

In parallel -

- Development and deployment of distributed database services (3D project)
- Testing and deployment of new mass storage services (SRM 2.x)

# Future Workshops

- Suggest 'regional' workshops to analyse results of experiment activities in SC4 during Q3/Q4 this year
- A 'global' workshop early 2007 focussing on experiment plans for 2007
- Another just prior to CHEP



# SC Tech Meeting

## Morning (09:00 - 12:30)

- Understanding Disk - Disk and Disk - Tape Results (Maarten)
- Why is it so hard to setup basic services? (Gavin)
- What features are missing in core services that are required for operations? (James)
- Moving from here to full production services and data rates (based on experiment and DTEAM challenges/tests) (Harry)
- Each Tier1 should prepare a few slides addressing specific issues regarding:
  - Problems seen during the disk-disk and disk-tape transfers and steps taken/planned to address them
  - Problems seen in implementing the agreed services, including a timeline
  - Problems encountered in the gLite 3.0 upgrade (maybe this has been covered to death elsewhere...)
  - Features seen as missing in core services / middleware required for operations

## Afternoon (14:00 - )

- Production Activities and Requirements by Experiment
  - ATLAS - Dario Barberis(?)
  - CMS - Ian Fisk
  - ALICE - Patricia Mendez, Latchezar Betev
  - LHCb - Umberto Marconi
- Specifically, each experiment should address:
  - What they want to achieve over the next few months with details of the specific tests and production runs.
  - Specific actions, timeline, sites involved.
  - If they have had bad experiences with specific sites then this should be discussed and resolved.

# 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.
- **Dates:** from 23 January 2007 09:00 to 25 January 2007 18:00
- **Location:** *CERN*  
Room: Main auditorium

## Sep 1-2, 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

# Conclusions

- 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***



