



The LCG Service Challenges: Ramping up the LCG Service

Jamie Shiers, CERN-IT-GD-SC

<http://cern.ch/jamie/SC.ppt>

April 2005

Antarctica

SC3 Planning Slides

N.B. these slides are a placeholder

The final slides will be posted as soon as possible


Agenda

- Goals and Timelines of the LCG Service Challenges
- Review of SC1 and SC2
- Summary of LHC Experiments' Computing Models
- Outline of SC3 and SC4
- After that it's the **FULL PRODUCTION SERVICE!**
- Plans for involving Tier2 sites in the Service Challenges
- Detailed SC3 planning

LCG Service Challenges - Overview

- LHC will enter production (physics) in April 2007
 - Will generate an enormous volume of data
 - Will require huge amount of processing power
- LCG 'solution' is a world-wide Grid
 - Many components understood, deployed, tested..
- But...
 - Unprecedented scale
 - Humungous challenge of getting large numbers of institutes and individuals, all with existing, sometimes conflicting commitments, to work together
- LCG must be ready at full production capacity, functionality and reliability in less than 2 years from now
 - Issues include h/w acquisition, personnel hiring and training, vendor rollout schedules etc.
- Should not limit ability of physicist to exploit performance of detectors nor LHC's physics potential
 - Whilst being stable, reliable and easy to use

LCG Deployment Schedule



Service Challenges

- **Purpose**
 - Understand what it takes to operate a real grid service - run for days/weeks at a time (outside of experiment Data Challenges)
 - Trigger/encourage the Tier1 & large Tier-2 planning - move towards real resource planning - based on realistic usage patterns
 - Get the essential grid services ramped up to target levels of reliability, availability, scalability, end-to-end performance
 - Set out milestones needed to achieve goals during the service challenges
- **NB: This is focussed on Tier 0 - Tier 1/large Tier 2**
 - Data management, batch production and analysis
- **Short term goal - by end 2004 -**
have in place a robust and reliable data management service and support infrastructure and robust batch job submission

From early proposal, May 2004

Why Service Challenges?

To test Tier-0 ↔ Tier-1 ↔ Tier-2 services

- **Network service**
 - Sufficient bandwidth: ~10 Gbit/sec
 - Backup path
 - Quality of service: security, help desk, error reporting, bug fixing, ..
 - **Robust file transfer service**
 - File servers
 - File Transfer Software (GridFTP)
 - Data Management software (SRM, dCache)
 - Archiving service: tapeservers, taperobots, tapes, tapedrives, ..
 - **Sustainability**
 - Weeks in a row un-interrupted 24/7 operation
 - Manpower implications: ~7 fte/site
 - Quality of service: helpdesk, error reporting, bug fixing, ..
- **Towards a stable production environment for experiments**

Whither Service Challenges?

- First discussions: GDB May - June 2004
 - May 18 - Lessons from Data Challenges and planning for the next steps (+ Discussion) (1h10') ([transparencies](#))
 - June 15 - Progress with the service plan team (10') ([document](#))
- Other discussions: PEB June 2004
 - June 8 - Service challenges - proposal (40') ([transparencies](#))
 - June 29 - Service challenges - status and further reactions (30') ([transparencies](#))
- May 2004 HEPiX
 - [LCG Service Challenges](#) Slides from Ian Bird (CERN)
- My involvement: from January 2005
 - Current Milestones:
<http://lcg.web.cern.ch/LCG/PEB/Planning/deployment/Grid%20Deployment%20Schedule.htm>

Key Principles

- Service challenges result in a series of services that exist in parallel with baseline production service
- Rapidly and successively approach production needs of LHC
- Initial focus: core (data management) services
- Swiftly expand out to cover full spectrum of production and analysis chain

- Must be as realistic as possible, including end-end testing of key experiment use-cases over extended periods with recovery from glitches and longer-term outages

- Necessary resources and commitment pre-requisite to success!
- Effort should not be under-estimated!

Informal Overview

- In addition to planned GDB meetings, Service Challenge Meetings, Network Meetings etc:
- Visits to all Tier1 sites (initially)
 - Goal is to meet as many of the players as possible
 - Not just GDB representatives! Equivalents of Vlado etc.
- Current Schedule:
 - Aim to complete many of European sites by Easter
 - "Round world" trip to BNL / FNAL / Triumpf / ASCC in April
- Need to address also Tier2s
 - Cannot be done in the same way!
 - Work through existing structures, e.g.
 - HEPiX, national and regional bodies etc.
 - e.g. GridPP (12)
- Talking of SC Update at May HEPiX (FZK) with more extensive programme at Fall HEPiX (SLAC)
 - Maybe some sort of North American T2-fest around this?

SC1 Review

- ☹ SC1 **did not** complete its goals successfully
 - Dec04 - Service Challenge I complete
 - mass store (disk) - mass store (disk)
 - 3 T1s (Lyon, Amsterdam, Chicago) (others also participated...)
 - 500 MB/sec (individually and aggregate)
 - 2 weeks sustained
 - Software; GridFTP plus some scripts

- **We did not meet the milestone of 500MB/s for 2 weeks**
 - We need to do these challenges to see what actually goes wrong
 - A lot of things do, and did, go wrong
 - We need better test plans for validating the infrastructure before the challenges (network throughput, disk speeds, etc...)


- OK, so we're off to a great start with the Service Challenges...

SC2 - Overview

- "Service Challenge 2"
 - Throughput test from Tier-0 to Tier-1 sites
 - Started 14th March
- Set up Infrastructure to 7 Sites
 - NIKHEF/SARA, IN2P3, FNAL, BNL, FZK, INFN, RAL
- 100MB/s to each site
 - 500MB/s combined to all sites at same time
 - 500MB/s to a few sites individually
- Goal : by end March, sustained 500 MB/s at CERN

SC2 met its throughput targets

- >600MB/s daily average for 10 days was achieved: Midday 23rd March to Midday 2nd April
 - Not without outages, but system showed it could recover rate again from outages
 - Load reasonable evenly divided over sites (give network bandwidth constraints of Tier-1 sites)




Division of Data between sites




Site	Average throughput (MB/s)	Data Moved (TB)
BNL	61	51
FNAL	61	51
GridKA	133	109
IN2P3	91	75
INFN	81	67
RAL	72	58
SARA	106	88
TOTAL	600	500

Individual site tests

- Overlapped with LCG Storage Management Workshop
 - Sites can pick days in next two weeks when they have the capacity
 - 500MB/s to disk
 - 60MB/s to tape
- FNAL was running 500MB/s disk tests at the time...



LCG Service Challenges – Deploying the Service



SC2 Summary

- SC2 met its throughput goals - and with more sites than originally planned!
 - A big improvement from SC1
- ☹ But we still don't have something we can call a service
 - Monitoring is better
 - We see outages when they happen, and we understood why they happen
 - First step towards operations guides
- Some advances in infrastructure and software will happen before SC3
 - gLite transfer software
 - SRM service more widely deployed
- We have to understand how to incorporate these elements

SC1/2 - Conclusions

- Setting up the infrastructure and achieving reliable transfers, even at much lower data rates than needed for LHC, is complex and requires a lot of technical work + coordination
- Even within one site - people are working very hard & are stressed. Stressed people do not work at their best. Far from clear how this scales to SC3/SC4, let alone to LHC production phase
- Compound this with the multi-site / multi-partner issue, together with time zones etc and you have a large "non-technical" component to an already tough problem (example of technical problem follows...)
- But... the end point is fixed (time + functionality)
- We should be careful not to over-complicate the problem or potential solutions
- And not forget there is still a humungous amount to do...
- (much much more than we've done...)

Computing Model Summary - Goals

- Present key features of LHC experiments' Computing Models in a consistent manner
- High-light the commonality
- Emphasize the key differences
- Define these 'parameters' in a central place (LCG web)
 - Update with change-log as required
- Use these parameters as input to requirements for Service Challenges
- To enable partners (T0/T1 sites, experiments, network providers) to have a clear understanding of what is required of them
- Define precise terms and 'factors'

Where do these numbers come from?

- Obtained from LHC Computing Models as reviewed in January
- Part of plan is to understand how sensitive overall model is to variations in key parameters
- Iteration with experiments is on-going
 - i.e. I have tried to clarify any questions that I have had
- Any mis-representation or mis-interpretation is entirely my responsibility
- Sanity check: compare with numbers from MoU Task Force
- (Actually the following LCG document now uses these numbers!)

http://cern.ch/LCG/documents/LHC_Computing_Resources_report.pdf

Nominal	These are the raw figures produced by multiplying e.g. event size x trigger rate.
Headroom	A factor of 1.5 that is applied to cater for peak rates.
Efficiency	A factor of 2 to ensure networks run at less than 50% load.
Recovery	A factor of 2 to ensure that backlogs can be cleared within 24 - 48 hours and to allow the load from a failed Tier1 to be switched over to others.
Total Requirement	<p>A factor of 6 must be applied to the nominal values to obtain the bandwidth that must be provisioned.</p> <p>Arguably this is an over-estimate, as "Recovery" and "Peak load" conditions are presumably relatively infrequent, and can also be smoothed out using appropriately sized transfer buffers.</p> <p>But as there may be under-estimates elsewhere...</p>

LHC Parameters (Computing Models)

Year	pp operations		Heavy Ion operations	
	Beam time (seconds/year)	Luminosity ($\text{cm}^{-2}\text{s}^{-1}$)	Beam time (seconds/year)	Luminosity ($\text{cm}^{-2}\text{s}^{-1}$)
2007	5×10^6	5×10^{32}	-	-
2008	(1.8 x) 10^7	2×10^{33}	(2.6 x) 10^6	5×10^{26}
2009	10^7	2×10^{33}	10^6	5×10^{26}
2010	10^7	10^{34}	10^6	5×10^{26}

(Real time given in brackets above)

LHC Schedule - "Chamonix" workshop




- First collisions: two months after first turn on in August 2007
- 32 weeks of operation, 16 weeks of shutdown, 4 weeks commissioning = 140 days physics / year (5 lunar months)



Overview of pp running

Experiment	SIM	SIMESD	RAW	Trigger	RECO	AOD	TAG
ALICE	400KB	40KB	1MB	100Hz	200KB	50KB	10KB
ATLAS	2MB	500KB	1.6MB	200Hz	500KB	100KB	1KB
CMS	2MB	400KB	1.5MB	150Hz	250KB	50KB	10KB
LHCb		400KB	25KB	2KHz	75KB	25KB	1KB



Overview of Heavy Ion running

Experiment	SIM	SIMESD	RAW	Trigger	RECO	AOD	TAG
ALICE	300MB	2.1MB	12.5MB	100Hz	2.5MB	250KB	10KB
ATLAS			5MB	50Hz			
CMS			7MB	50Hz	1MB	200KB	TBD
LHCb	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Tier-1 Centres (January 2004)

				ALICE	ATLAS	CMS	LHCb	
1	GridKa	Karlsruhe	Germany	X	X	X	X	4
2	CCIN2P3	Lyon	France	X	X	X	X	4
3	CNAF	Bologna	Italy	X	X	X	X	4
4	NIKHEF/SARA	Amsterdam	Netherlands	X	X		X	3
5	Nordic	Distributed	Dk, No, Fi, Se	X	X			1
6	PIC	Barcelona	Spain		X	X	X	3
7	RAL	Didcot	UK	X	X	X	X	4
8	Triumf	Vancouver	Canada		X			1
9	BNL	Brookhaven	US		X			1
10	FNAL	Batavia, Ill.	US			X		1
11	ASCC	Taipei	Taiwan		X	X		2
				6	10	7	6	

x – announced at January GDB

pp / AA data rates (equal split)

<i>Centre</i>	<i>ALICE</i>	<i>ATLAS</i>	<i>CMS</i>	<i>LHCb</i>	<i>Rate into T1</i>	<i>Rate into T1 (AA)</i>
ASCC, Taipei	0	1	1	0	118.7	28.2
CNAF, Italy	1	1	1	1	205.0	97.2
PIC, Spain	0	1	1	1	179.0	28.2
IN2P3, Lyon	1	1	1	1	205.0	97.2
GridKA, Germany	1	1	1	1	205.0	97.2
RAL, UK	1	1	1	1	205.0	97.2
BNL, USA	0	1	0	0	72.2	11.3
FNAL, USA	0	0	1	0	46.5	16.9
TRIUMF, Canada	0	1	0	0	72.2	11.3
NIKHEF/SARA, Netherlands	1	1	0	1	158.5	80.3
Nordic Centre	1	1	0	0	98.2	80.3
Totals	6	10	7	6		

N.B. these calculations assume equal split as in Computing Model documents. It is clear that this is not the 'final' answer...

Streaming

- All experiments foresee RAW data streaming, but with different approaches:
 - CMS: $O(50)$ streams based on trigger path
 - Classification is immutable, defined by L1+HLT
 - Atlas: 4 streams based on event types
 - Primary physics, Express line, Calibration, Debugging and diagnostic
 - LHCb: >4 streams based on trigger category
 - B-exclusive, Di-muon, D^* Sample, B-inclusive
 - Streams are not created in the first pass, but during the “stripping” process
- Not clear what is the best/right solution. Probably bound to evolve in time.

Reprocessing

- Data need to be reprocessed several times because of:
 - Improved software
 - More accurate calibration and alignment
- **Reprocessing mainly at T1 centers**
 - LHCb is planning on using the T0 during the shutdown - not obvious it is available
- **Number of passes per year**

Alice	Atlas	CMS	LHCb
3	2	2	4

- But experience shows the reprocessing requires huge effort!
- Use these numbers in the calculation but 2 / year will be good going!

pp / AA data rates - comments

<i>Centre</i>	<i>ALICE</i>	<i>ATLAS</i>	<i>CMS</i>	<i>LHCb</i>	<i>Rate into T1</i>	<i>Rate into T1 (AA)</i>
ASCC, Taipei	0	1	1	0	118.7	28.2
CNAF, Italy	1	1	1	1	205.0	97.2
PIC, Spain	0	1	1	1	179.0	28.2
IN2P3, Lyon	1	1	1	1	205.0	97.2
GridKA, Germany	1	1	1	1	205.0	97.2
RAL, UK	1	1	1	1	205.0	97.2
BNL, USA	0	1	0	0	72.2	11.3
FNAL, USA	0	0	1	0	46.5	16.9
TRIUMF, Canada	0	1	0	0	72.2	11.3
NIKHEF/SARA, Netherlands	1	1	0	1	158.5	80.3
Nordic Centre	1	1	0	0	98.2	80.3
Totals	6	10	7	6		

Sites expected to take fraction proportional to the communities that they serve / resources they have (same thing?)

Will likely narrow spread to ~150MB/s (200MB/s peaks?)

Base Requirements for T1s

- Provisioned bandwidth comes in units of 10Gbits/sec although this is an evolving parameter
 - *From Reply to Questions from Computing MoU Task Force...*
 - Since then, some parameters of the Computing Models have changed
 - Given the above quantisation, relatively insensitive to small-ish changes
 - Important to understand implications of multiple-10Gbit links, particularly for sites with Heavy Ion programme
 - Spread of AA distribution during shutdown probably means 1 link sufficient...
- For now, planning for 10Gbit links to all Tier1s

T1/T2 Roles

Tier1

- Keep certain portions of RAW, ESD, sim ESD
- Full copies of AOD + TAG, calibration data
- Official physics group large scale data analysis
- ALICE + LHCb:
 - also contribute to simulation

Tier2

- Keep certain portions of AOD and full copies of TAG for real + simulated data
 - LHCb: sim only at T2s
- Selected ESD samples
- Produce simulated data
- General end-user analysis

Based on "T1 Services for T2 Centres" document

(Just type this into Google)

MC Data

	Units	ALICE		ATLAS	CMS	LHCb
		p-p	Pb-Pb	p-p	p-p	
Time to reconstruct 1 event	kSI2k sec	5.4	675	15	25	2.4
Time to simulate 1 event	kSI2k sec	35	15000	100	45	50

Tier2 sites offer 10 – 1000 kSI2K years
ATLAS: 16MSI2K years over ~30 sites in 2008
CMS: 20MSI2K years over ~20 sites in 2008

Parameter	Unit	ALICE		ATLAS	CMS	LHCb
		p-p	Pb-Pb			
Events/year	Giga	1	0.1	2	1.5	20
Events SIM/year	Giga	1	0.01	0.4	1.5	4
Ratio SIM/data	%	100%	10%	20%	100%	20%

GridPP Estimates of T2 Networking

	Number of T1s	Number of T2s	Total T2 CPU	Total T2 Disk	Average T2 CPU	Average T2 Disk	Network In	Network Out
			KSI2K	TB	KSI2K	TB	Gb/s	Gb/s
ALICE	6	21	13700	2600	652	124	0.010	0.600
ATLAS	10	30	16200	6900	540	230	0.140	0.034
CMS	6 to 10	25	20725	5450	829	218	1.000	0.100
LHCb	6	14	7600	23	543	2	0.008	0.008

The CMS figure of 1Gb/s into a T2 comes from the following:

- Each T2 has ~10% of current RECO data and 1/2 AOD (real+MC sample)
- These data are refreshed every 3 weeks
 - compatible with frequency of (possible) major selection pass at T1s
- See CMS Computing Model S-30 for more details

Service Challenge 3

Goals and Timeline for
Service Challenge 3

Service Challenge 3 - Phases

High level view:

- Setup phase (includes Throughput Test)
 - 2 weeks sustained in July 2005
 - "Obvious target" - GDB of July 20th
 - Primary goals:
 - 150MB/s disk - disk to Tier1s;
 - 60MB/s disk (T0) - tape (T1s)
 - Secondary goals:
 - Include a few named T2 sites (T2 -> T1 transfers)
 - Encourage remaining T1s to start disk - disk transfers
- Service phase
 - September - end 2005
 - Start with ALICE & CMS, add ATLAS and LHCb October/November
 - All offline use cases except for analysis
 - More components: WMS, VOMS, catalogs, experiment-specific solutions
 - Implies production setup (CE, SE, ...)

SC3 - Production Services

- SC3 is a relatively small step wrt SC2 *in terms of throughput!*
- We know we can do it technology-wise, but do we have a solution that will scale?
- **Let's make it a priority for the coming months to streamline our operations**
- And not just throw resources at the problem...
 - which we don't have...
- Whilst not forgetting 'real' goals of SC3... *i.e. services!*

SC3 - Service Phase

- It sounds easy:
“all offline Use Cases except for analysis”
- And in some senses it is:
these are well understood and tested
- So it's clear what we have to do:
 - Work with the experiments to understand and agree on the experiment-specific solutions that need to be deployed
 - Agree on a realistic and achievable work-plan that is consistent with overall goals / constraints
- Either that or send a 'droid looking for Obi-Wan Kenobi...

Service Phase - Priorities

- Experiments have repeatedly told us to focus on reliability and functionality
- This we need to demonstrate as a first step...
- But cannot lose sight of need to pump up data rates - whilst maintaining production service - to pretty impressive "DC" figures

SC3 Preparation Workshop

- This (proposed) workshop will focus on very detailed technical planning for the whole SC3 exercise.
- It is intended to be as interactive as possible, i.e. not presentations to an audience largely in a different (wireless) world.
- There will be sessions devoted to specific experiment issues, Tier1 issues, Tier2 issues as well as the general service infrastructure.
- Planning for SC3 has already started and will continue prior to the workshop.
- This is an opportunity to get together to iron out concerns and issues that cannot easily be solved by e-mail, phone conferences and/or other meetings prior to the workshop.

SC3 Preparation W/S Agenda

- 4 x 1/2 days devoted to experiments
 - in B160 1-009, phone conferencing possible
- 1 day focussing on T1/T2 issues together with output of above
 - In 513 1-024, VRVS available
- Dates are 13 - 15 June (Monday - Wednesday)
- Even though conference room booked tentatively in February, little flexibility in dates even then!

SC3 on

- SC3 is significantly more complex than previous challenges
- It includes experiments s/w, additional m/w, Tier2s etc
 - Proving we can transfer dummy files from A-B proves nothing
 - Obviously need to show that basic infrastructure works...
- Preparation for SC3 includes:
 - Understanding experiments' Computing Models
 - Agreeing involvement of experiments' production teams
 - Visiting all (involved) Tier1s (multiple times)
 - Preparing for the involvement of 50-100 Tier2s
- Short of resources at all levels:
 - "Managerial" - discussing with experiments and Tier1s (visiting)
 - "Organizational" - milestones, meetings, workshops, ...
 - "Technical" - preparing challenges and running CERN end - 24 x 7 ???

2005 Q1 - SC3 preparation

Prepare for the next service challenge (SC3)
 -- in parallel with SC2 (reliable file transfer) -

Build up 1 GByte/s *challenge* facility at CERN

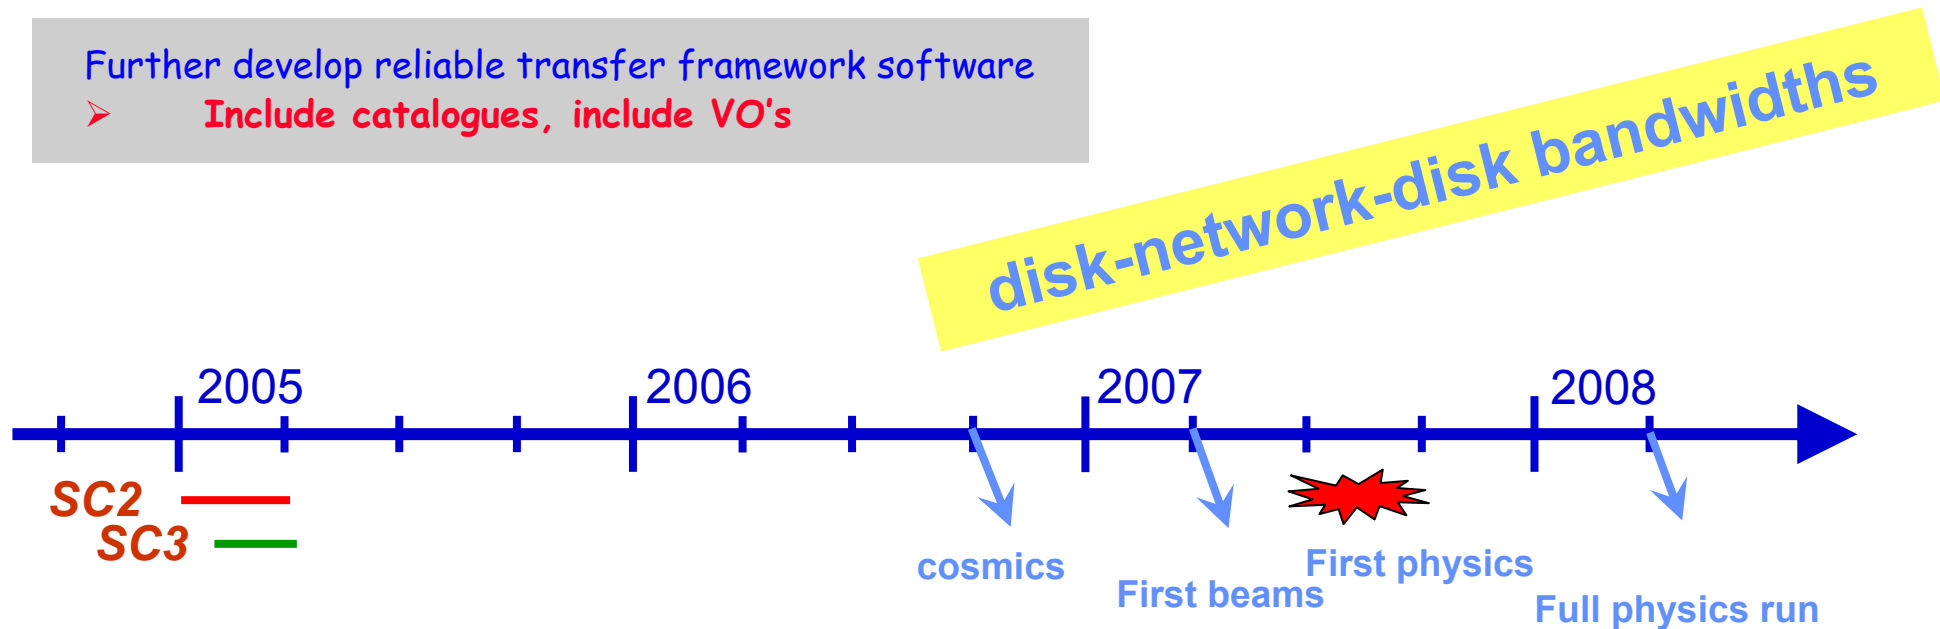
- The current 500 MByte/s facility used for SC2 will become the *testbed* from April onwards (10 ftp servers, 10 disk servers, network equipment)


Build up infrastructure at each external centre

- Average *capability* ~150 MB/sec at a Tier-1 (to be agreed with each T-1)

Further develop reliable transfer framework software

- Include catalogues, include VO's





SC3 - Milestone Decomposition

- **File transfer goals:**
 - Build up disk - disk transfer speeds to 150MB/s with 1GB/s out of CERN
 - SC2 was 100MB/s - agreed by site
 - Include tape - transfer speeds of 60MB/s with 300MB/s out of CERN
- **Tier1 goals:**
 - Bring in additional Tier1 sites wrt SC2 (at least wrt the original plan...)
 - PIC and Nordic most likely added later: SC4?
- **Tier2 goals:**
 - Start to bring Tier2 sites into challenge
 - Agree services T2s offer / require
 - On-going plan (more later) to address this via GridPP, INFN etc.
- **Experiment goals:**
 - Address main offline use cases *except* those related to analysis
 - i.e. real data flow out of T0-T1-T2; simulation in from T2-T1
- **Service goals:**
 - Include CPU (to generate files) and storage
 - Start to add additional components
 - Catalogs, VOs, experiment-specific solutions etc, 3D involvement, ...
 - Choice of software components, validation, fallback, ...

SC3 - Experiment Goals

- Meetings on-going to discuss goals of SC3 and experiment involvement
- Focus on:
 - First demonstrate robust infrastructure;
 - Add 'simulated' experiment-specific usage patterns;
 - Add experiment-specific components;
 - Run experiments offline frameworks but don't preserve data;
 - Exercise primary Use Cases *except* analysis (SC4)
 - Service phase: data is preserved...
- Has significant implications on resources beyond file transfer services
 - Storage; CPU; Network... Both at CERN and participating sites (T1/T2)
 - May have different partners for experiment-specific tests (e.g. not all T1s)
- In effect, experiments' usage of SC during service phase = data challenge
- Must be exceedingly clear on goals / responsibilities during each phase!

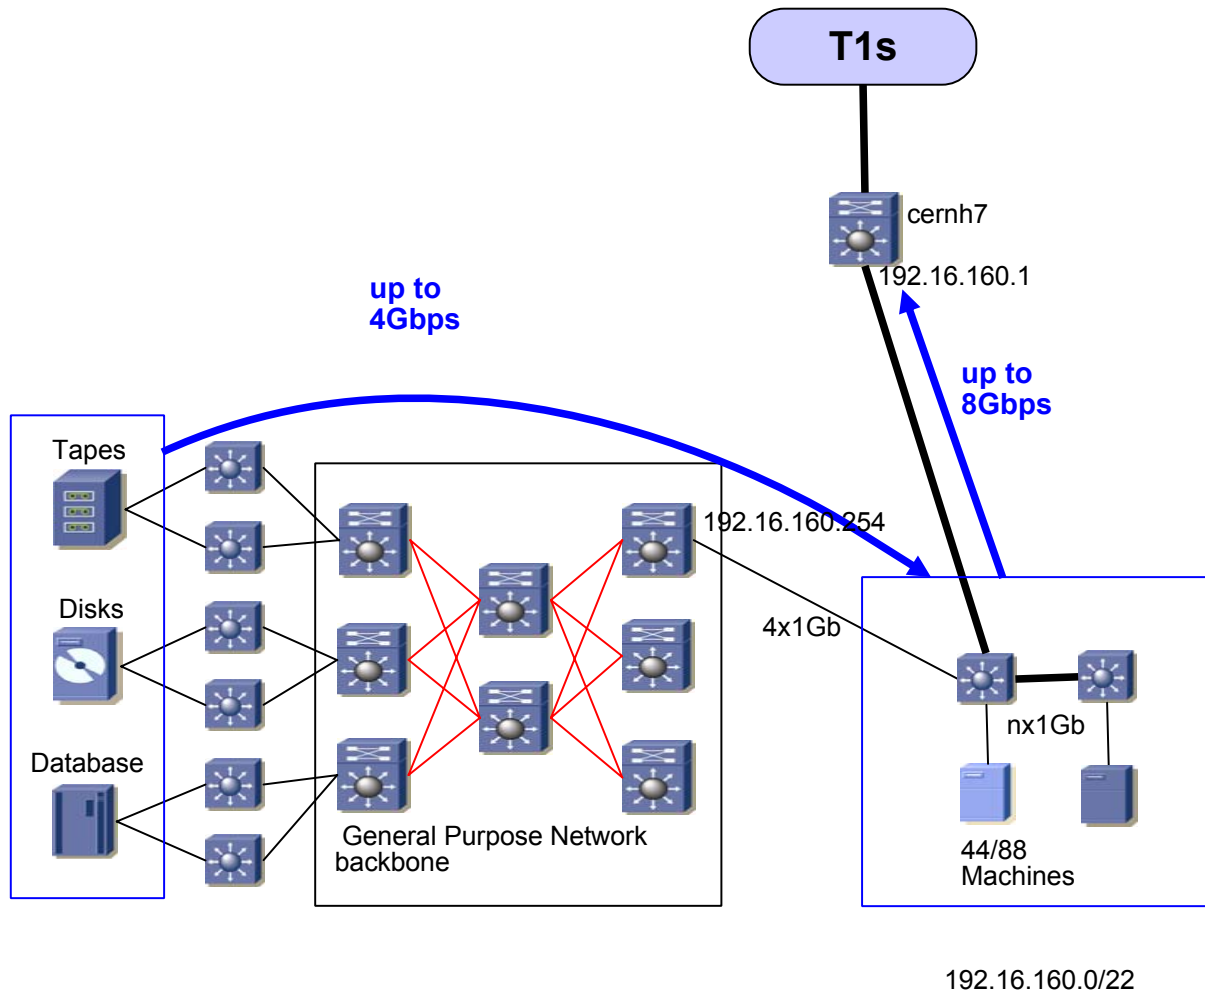
SC3 - Experiment Involvement Cont.

- Regular discussions with experiments have started
 - ATLAS: at DM meetings
 - ALICE+CMS: every ~2 weeks
 - LHCb: no regular slot yet, but discussions started...
- Anticipate to start first with ALICE and CMS (exactly when TDB)
ATLAS and LHCb around October
 - T2 sites being identified in common with these experiments
 - More later...
 - List of experiment-specific components and the sites where they need to be deployed being drawn up
 - Need this on April timeframe for adequate preparation & testing

'SC3' pilot services


- Not really restricted to SC3, except timewise...
- gLite FTS pilot
 - For experiments to start to get experience with it, integrating with their frameworks etc
- File catalogs
 - Both LFC and FiReMan... and RLS continues for the time being...
- Expect these to start in May, based on s/w delivered this week (code free April 15) for release end May
 - Max one more cycle for SC3
 - end June is too late for throughput phase!
 - LFC is already part of LCG releases

Service Challenge 3 Network – CERN Side



Routing
 cernh7, one GPN backbone router and the 44-88 machines will be connected at layer 2. The 160.16.160.0/22 prefix will be shared among their interfaces. The correct routing must be configured on each one of the 44-88 machines: default towards the GPN, tier1s prefixes towards cernh7.

Security
 Very strict access-list must be configured on cernh7 and the GPN backbone router.



2005 Sep-Dec - SC4 preparation


In parallel with the SC3 model validation period,
in preparation for the first 2006 service challenge (SC4) -

Using 500 MByte/s test facility

- test PIC and Nordic T1s
- and T2's that are ready (Prague, LAL, UK, INFN, ..)

Build up the production facility at CERN to 3.6 GBytes/s

Expand the capability at all Tier-1s to full nominal data rate



Historical slides from Les / Ian

2006 Jan-Aug - SC4

SC4 - full computing model services


- Tier-0, ALL Tier-1s, all major Tier-2s operational at full target data rates (~2 GB/sec at Tier-0)
- acquisition - reconstruction - recording - distribution, *PLUS* ESD skimming, servicing Tier-2s

Goal - stable test service for one month - April 2006

100% Computing Model Validation Period (May-August 2006)

Tier-0/1/2 full model test - All experiments

- 100% nominal data rate, with processing load scaled to 2006 cpus



SC4 Planning


- Discussing a joint workshop with ARDA focussing on SC4 after the summer
 - Tentative dates: during week of October 10 - 14
 - Clash with GDB and HEPiX?
- After we have concrete results from SC3?
- When a more concrete model for analysis has appeared?
- In all events, early enough for SC4 planning...

2006 Sep - LHC service available

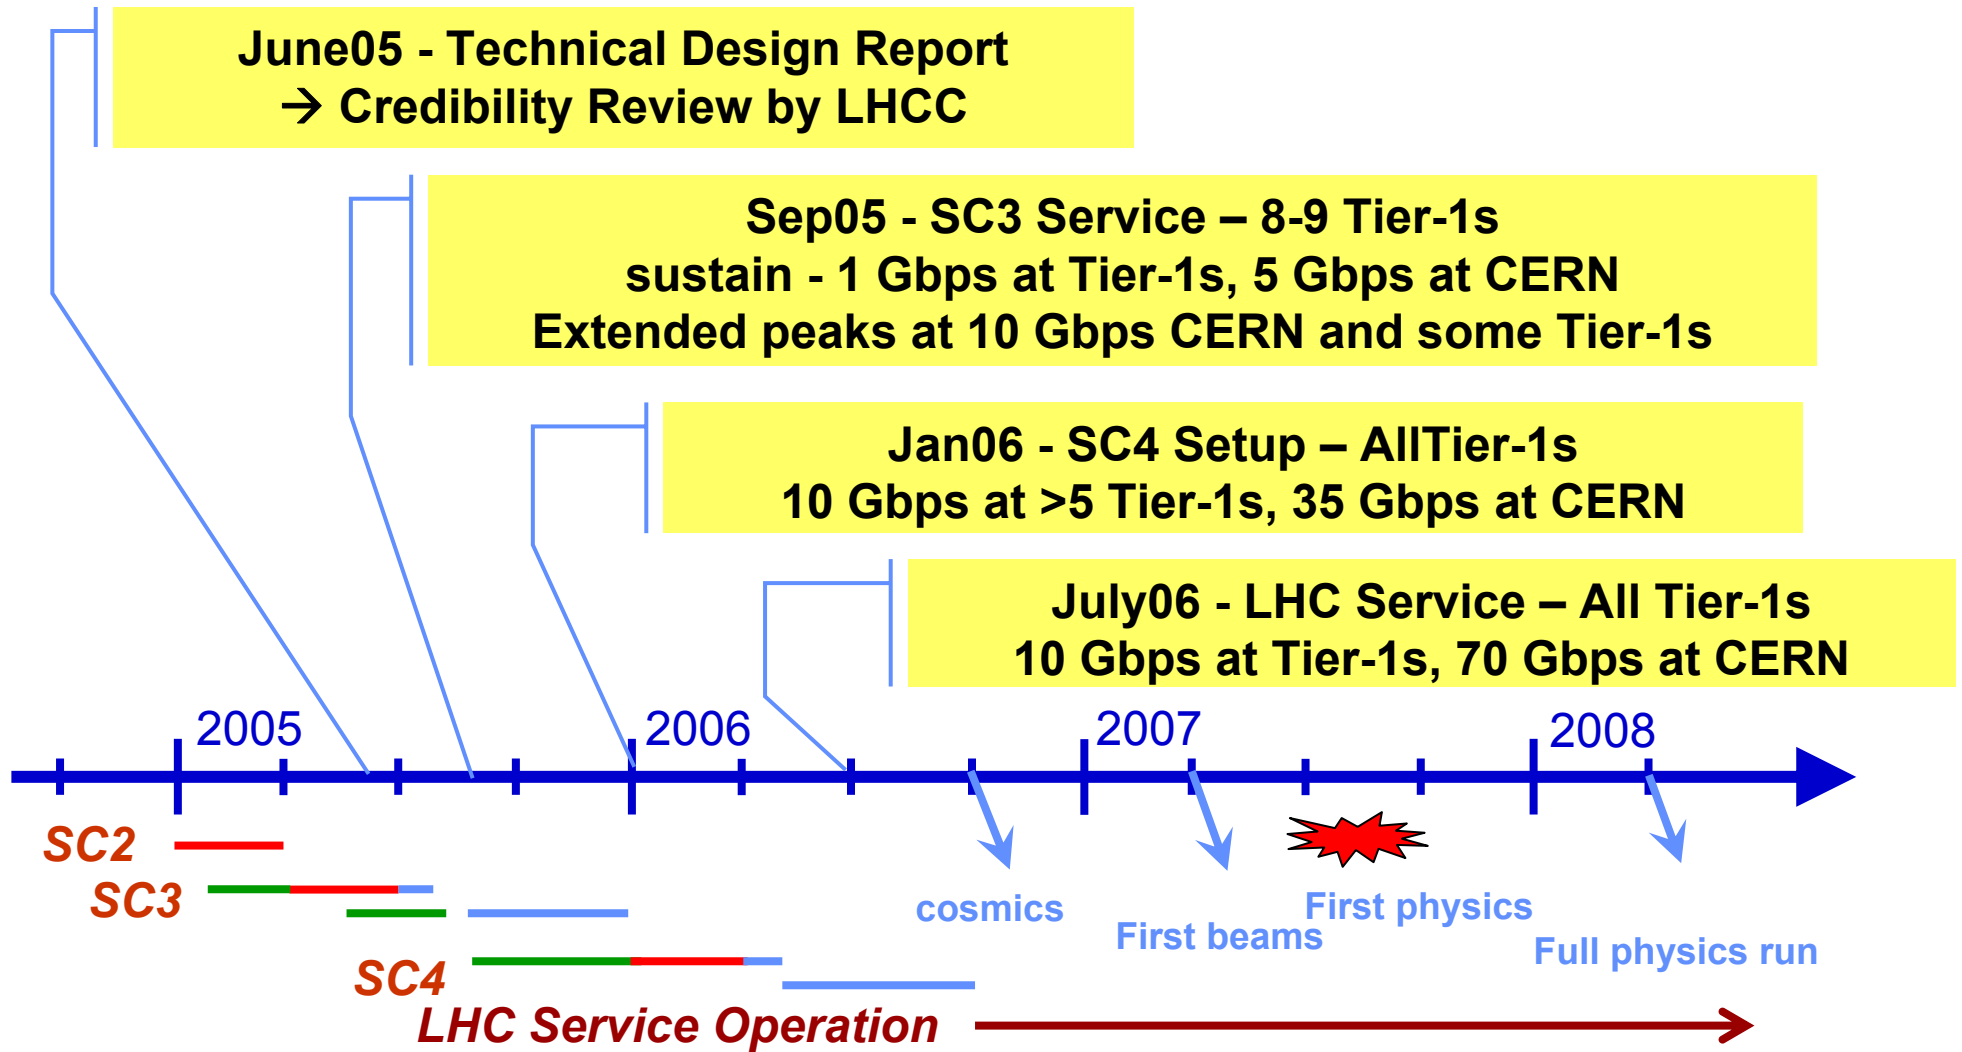
The SC4 service becomes the permanent LHC service - available for experiments' testing, commissioning, processing of cosmic data, etc.

All centres ramp-up to capacity needed at LHC startup

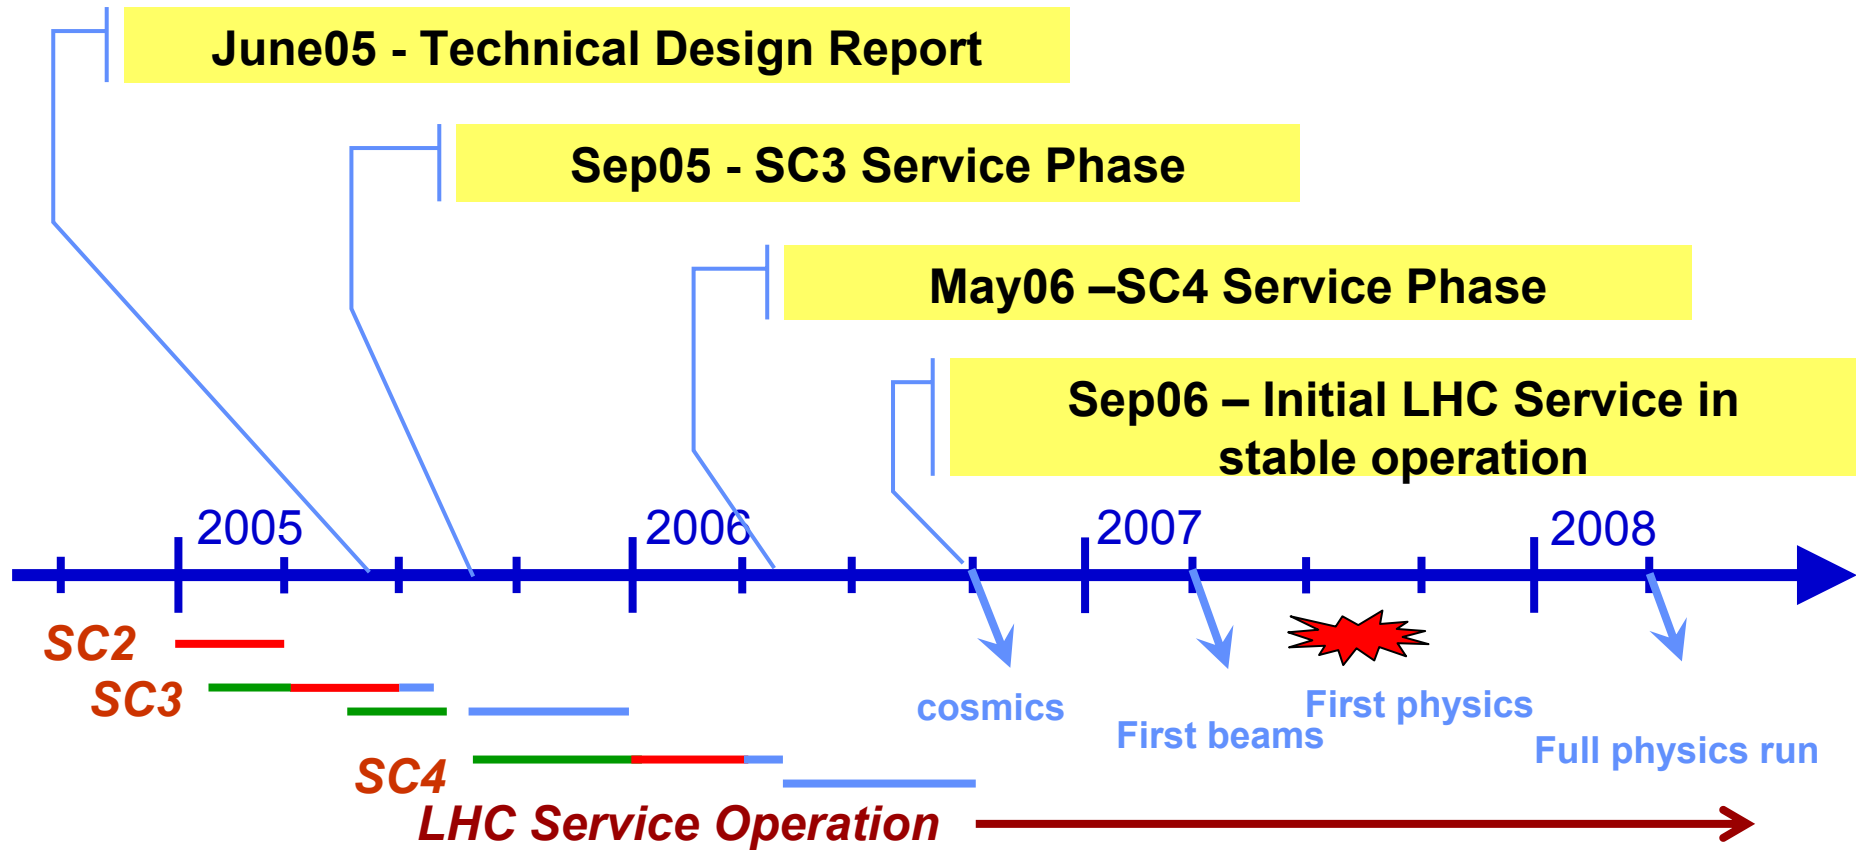
- TWICE nominal performance
- Milestone to demonstrate this 3 months before first physics data
→ April 2007



Key dates for Connectivity



Key dates for Services





- Additional threads started to address:
 - Experiment involvement;
 - Bringing T2s in SC3;
 - Longer-term goals of bringing all T2s into the LCG Service (Challenges)
- The enthusiasm and support provided to these new activities is much appreciated
- We have a lot of work ahead...
- ...but the problem is beginning to become tractable(?)

Conclusions

- To be ready to fully exploit LHC, significant resources need to be allocated to a series of Service Challenges by all concerned parties
- These challenges should be seen as an essential on-going and long-term commitment to achieving production LCG
- The countdown has started - we are already in (pre-)production mode
- Next stop: 2020

Postscript

**The Service is
the Challenge**