




LCG Service Challenges: Status and Plans

Jamie.Shiers@cern.ch - <http://cern.ch/jamie/>



July 2005

Antarctica



Executive Summary



- 'Pilots' - LFC & FTS
 - Scheduled originally for mid-May
 - Multiple delays (obtaining / configuring h/w, s/w, procedures etc.)
 - LFC has been available for some weeks
 - Testing with ATLAS, ALICE, CMS, LHCb
 - FTS fully available since Monday 11th July

- SC3 Throughput Tests
 - Started - seeing 'reasonable' traffic using FTS (most T1s) + PhEDEx (FNAL + others)
 - Problems at many sites at SRM level: [monitoring page](#)
 - Holes in service over w/e (as expected)
 - Need to debug SRMs before can look at remaining FTS failures
 - We will learn a lot about running these basic services!
 - (Whilst shaking down the services significantly)
 - **Key deliverable: reliable, stress-tested core data management services**

- Site preparations: work still needed for Service Phase!
 - Valuable information through [SC Wiki](#)
 - Experiments in direct contact with some sites (e.g. Lyon)
 - This is helping to push the preparation!
 - See <http://cern.ch/LCG/> -> [Service Challenges](#)

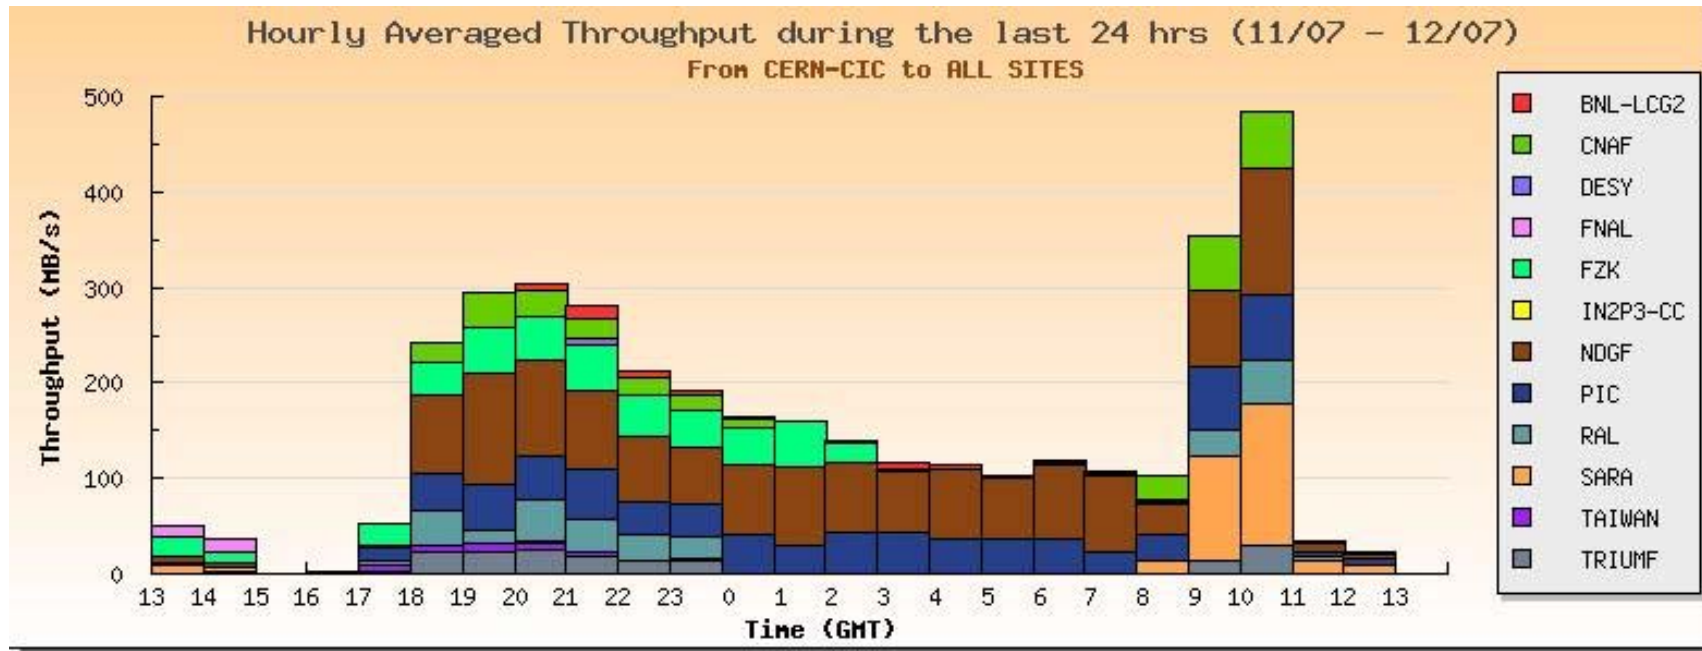
- **An awful lot has been achieved since SC2 (and SC1...) but still more ahead...**

Where are we now?



- Started Throughput Tests Monday July 11th
- Many problems seen - eventually died in evening
- All nodes at CERN end rebooted July 12th to clean state
- Daily logs:
<https://uimon.cern.ch/twiki/bin/view/LCG/ServiceChallengeThreeProgress>
<https://uimon.cern.ch/twiki/bin/view/LCG/SCThreeJCDailyLog>
- We need a clean environment to get things going
- This translates to transfers initially driven by us (see [table](#))
- Once we have a stable state - hopefully in days or less - we can discuss opening things up

Preliminary Throughput Plot



Summary of Problem Areas



- Major shortage of manpower evident at many sites
- Compensated for by enormous enthusiasm and very hard work!
- A lot of components are new and we are experiencing to-be-expected teething problems
- Service delivery spread over many groups
- Reporting lines and responsibilities not always clear...
- For a variety of reasons, we are behind schedule (that is, delivery of production-ready components end-May 2005)
- But the SCs are to show our state of unreadiness, as well as our state of readiness...

Basic Components For Setup Phase



- Each T1 to provide 10Gb network link to CERN
- Each T1 + T0 to provide SRM 1.1 interface to managed storage
 - This goes for the named T2s for the T2-T1 transfer tests too
- T0 to provide File Transfer Service; also at named T1s for T2-T1 transfer tests
- Baseline Services Working Group, Storage Management Workshop and SC3 Preparation Discussions have identified one additional data management service **for SC3**, namely the LFC
 - Not all experiments (ALICE) intend to use this
 - Nor will it be deployed for all experiments at each site
- However, as many sites support multiple experiments, and will (presumably) prefer to offer common services, this can be considered a basic component
 - Table by site below

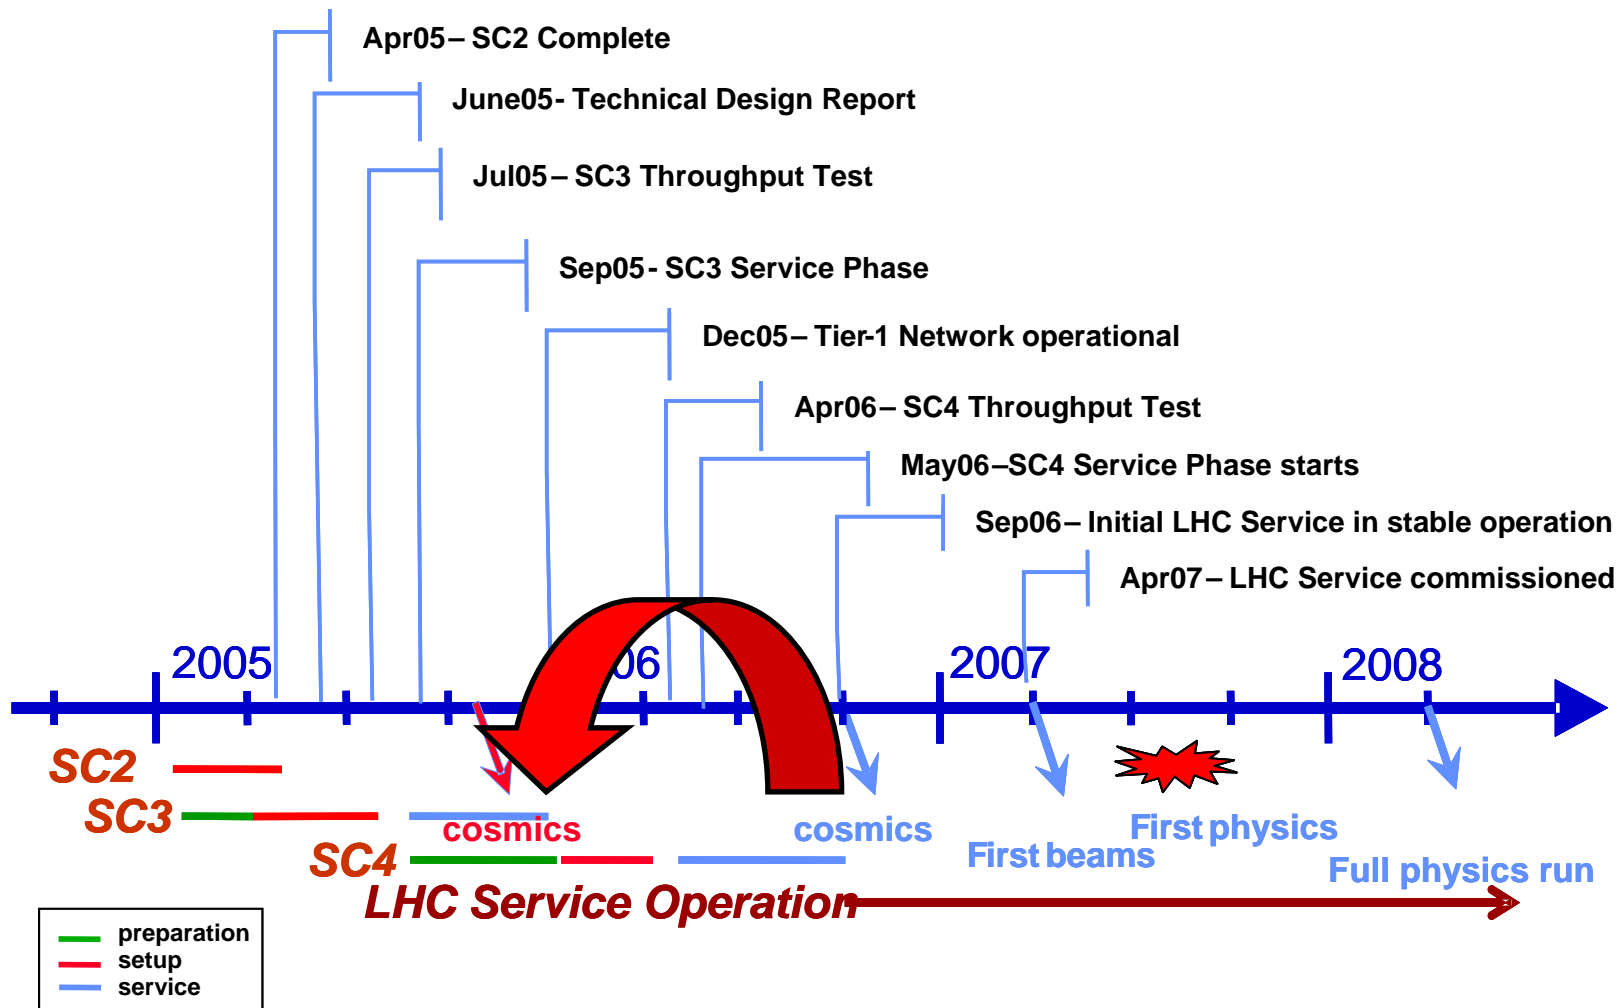
From May GDB

Site Components - Updated



- Each T1 to provide 10Gb network link to CERN
- Each site to provide SRM 1.1 interface to managed storage
 - All sites involved in SC3: T0, T1s, T2s.
- T0 to provide File Transfer Service; also at named T1s for T2-T1 transfer tests
 - Named Tier1s: BNL, CNAF, FZK, RAL; Others also setting up FTS
 - CMS T2s being supported by a number of T1s using PhEDEx
- LCG File Catalog - not involved in Throughput but **needed for Service**
 - ALICE / ATLAS: site local catalog
 - LHCb: central catalog with >1 R/O 'copies' (on ~October timescale)
 - CMS: evaluating different catalogs
 - **FNAL: Globus RLS, T0+other T1s: LFC; T2s: POOL MySQL, GRLS, ...**
- T2s - many more than foreseen
 - Running DPM or dCache, depending on T1 / local preferences / support
 - [Support load at CERN through DPM / LFC / FTS client]
- Work still needed to have these consistently available as services

LCG Deployment Schedule



SC3 - Future Milestones (TDR)



Date	Description
31 July 05	Service Challenge 3 Set-up: Set-up complete and basic service demonstrated. Performance and throughput tests complete. See Section 6.2.4 for detailed goals.
1 Sept 05	Service Challenge 3: start of stable service phase, including at least 9 Tier-1 and 10 Tier-2 centres.
31 Dec 05	Tier-0/1 high-performance network operational at CERN and 8 Tier-1s.
31 Dec 05	750 MB/s data recording demonstration at CERN: Data generator → disk → tape sustaining 750 MB/s for one week using the CASTOR mass storage system.
28 Feb 06 31 Jan 06	All required software for baseline services deployed and operational at all Tier-1s and at least 20 Tier-2 sites.



Tier2 participation by Tier1

Tier1	(Approx) Status mid-June
ASCC, Taipei	Yes; preparing for T2 support in Asia - Pacific
CNAF, Italy	Yes; workshop held end May in Bari
PIC, Spain	Yes; no Oracle service for FTS; CMS transfers with PhEDEx
IN2P3, Lyon	Yes; LAL + IN2P3
GridKA, Germany	Yes – studying with DESY
RAL, UK	Yes – plan in place for several Tier2s
BNL, USA	Yes – named ATLAS Tier2s
FNAL, USA	Yes – CMS transfers with PhEDEx; already performing transfers
TRIUMF, Canada	Yes – planning to install FTS and identify T2s for tests
NIKHEF/SARA, Netherlands	Re-evaluate on SC4 timescale (which T2s outside NL?)
Nordic Centre	Yes; preparing T1 / T2s in Nordic region
CERN	Swiss T2 plus some others not unlikely

- Virtually all Tier1s actively preparing for Tier2 support
- Much interest from Tier2 side: debugging process rapidly!
- Some Tier2s still need to identify their Tier1 centre
- This is an area where things are looking good!

T2s



		Prague	Prague, Czech Rep.	X	X		
		KFKI	Budapest, Hungary	X		X	
		SZTAKI	Budapest, Hungary	X		X	
		Eotvos University	Budapest, Hungary	X		X	
		Helsinki Institute of Physics	Helsinki, Finland			X	
		Krakow	Krakow, Poland	X	X		X
		Warszawa	Warszawa, Poland	X		X	X
		Russian Tier-2 cluster	Moscow, Russian Republic	X	X	X	X
		Technion	Haifa, Israel		X		
		Weizmann	Rehovot, Israel		X		
		Tel Aviv Univ.	Tel Aviv, Israel		X		
		PAEC-1/NCP/NUST/COMSATS	Pakistan			X	
		UERJ	Rio de Janeiro, Brazil			X	
		TIFR	Mumbai, India			X	
		VECC/SINP	Kolkata, India	X			
		Melbourne			X		
		Cape Town		X			
		Etc.					

Tier-2s



~100 identified – number still growing

Services (all 24 x 7, now - 2020)



- **Managed storage: SRM 1.1 interface, moving to 2.1 (2.2)**
 - No assumption / requirement for tertiary storage at T2s
 - Monte Carlo generation: write-through cache to T1s
 - Analysis data: read-only (?) cache from T1s; ~30 day lifetime(?)
- **Reliable network links: 1Gbit/s at T2s, 10Gbit/s at T1s, support full data rate to all T1s out of T0**
 - If network link goes down, data must be re-routed to an alternate site; pro-longed outage a major problem; need correspondingly large data buffers at T0 / T1s
- **Reliable File Transfer services**
 - Gridftp, srmcopy + higher level functionality - SERVICE
- **File catalogs, data management tools, database services**
- **Basic Services: workload management, VO management, monitoring etc.**
- **Multiple levels of experiment specific software and corresponding additional complexity**

More on Services



- 24 x 7 services do not mean that people have to be chained to the computer 24 x 7
- Services must be designed / deployed to be as reliable and recoverable as possible
 - Monitor to check that this is so - including end to end monitoring
- Cannot tolerate failure of a major component Friday evening not looked at until Monday morning... after coffee...
 - Eventually run in degraded mode?
- Need to use existing experience and technology...
 - Monitoring, alarms, operators, SMS to 2nd / 3rd level support...
- Now is the time to get these procedures in place
 - Must be able to arrange that suitable experts can have network access within reasonable time
 - Even from the beach / on the plane ...

Services at CERN



- Building on standard service model
- First level support: operations team
 - Box-level monitoring, reboot, alarms, procedures etc
- Second level support team: Grid Deployment group
 - Alerted by operators and/or alarms (and/or production managers...)
 - Follow 'smoke-tests' for applications
 - Identify appropriate 3rd level support team to call
 - Responsible for maintaining and improving procedures
 - Two people per week: complementary to System Manager on Duty
 - Provide daily report to SC meeting (09:00); interact with experiments
 - Members: IT-GD-EIS, IT-GD-SC (including me)
 - Phone numbers: 164111; 164222
- Third level support teams: by service
 - Notified through operators and / or 2nd level (by agreement)
 - Should be called (very) rarely... (Definition of a service?)

Target Date for Services



- Assumed to be ~October 2005
- FTS: goals
 - 1 experiment, 1PB in 100 days running
 - If 1GB files, 10^4 files / day per experiment
 - $n \cdot 10^4$ transfer / day (RAW files, ESD, AOD, TAG, ...)
 - >99.99% reliability if 1 manual intervention / day!
 - 99.99999% reliability if 1 manual intervention / year! (all expts)

 - (Is 'per transfer' the right way to measure this?)
- File Catalog:
 - Need to clarify remaining questions regarding deployment model(s)
 - Service expectations, as above, need to be understood
 - Deployment model on MySQL (T2s, some T1s):
 - Stop server nightly(?), copy file & restart
 - Upload file to T1 using FTS or other
 - Need willing T1 (or T2) with MySQL experience to prototype
- CASTOR2:

Where are we now?



- >Mid-point in activity (first proposal to completion)
- Demonstrated sustained disk - disk data rates of **100MB/s** to multiple Tier1 sites, **>500MB/s** out of CERN for some 10 days; **800MB/s** to a single site (FNAL)
- Now (July): demonstrate **150MB/s** to Tier1s; **1GB/s** out of CERN (disk - disk) plus **60MB/s** to tape at Tier1s
- In terms of data rate alone, have to double data rates, plus whatever is necessary in terms of 'safety factors' ,including recovering backlogs from outages etc.
- But so far, these tests have just been with dummy files, with the bare minimum software involved
- In particular, none of the experiment software has been included!
- Huge additional work: add major complexity whilst doubling rates and providing high quality services
- (BTW, neither of first two challenges fully met their goals)

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
- But... the end goal is still the same (time + functionality)
- We have a solution in principle, but it's not there yet. Potential
- And... it's not just a technical problem. ...That has got worse, as untreated...
- (much much more than we've done...)

I highlighted this as it remains an unsolved problem. ...That has got worse, as untreated...

Core Site Services



- **CERN**
 - Storage: Castor/SRM
 - File catalogue: POOL LFC Oracle
- **FNAL**
 - Storage: dCache/SRM
 - File catalogue: POOL Globus RLS
- **CNAF**
 - Storage: Castor/SRM
 - File catalogue: POOL LFC Oracle
- **RAL**
 - Storage: dCache/SRM
 - File catalogue: POOL LFC Oracle?
- **IN2P3**
 - Storage: dCache/SRM
 - File catalogue: POOL LFC Oracle
- **SARA/NIKHEF**
 - Storage: dCache/SRM
 - File catalogue: POOL LFC MySQL(?)
- **PIC**
 - Storage: Castor/SRM
 - File catalogue: POOL LFC MySQL
- **FZK**
 - Storage: dCache/SRM
 - File catalogue: POOL LFC Oracle
- **ASCC**
 - Storage: Castor/SRM
 - File catalogue: POOL LFC Oracle
- **BNL**
 - Storage: dCache/SRM
 - File catalogue: POOL LFC Oracle
- **TRIUMF**
 - Storage: dCache/SRM
 - File catalogue: POOL LFC MySQL(?)
- **NDGF**
 - Storage:
 - File catalogue:

See SC Wiki for status and progress...

Overall Schedule (Raw-ish)



Sep	Sep	Oct	Oct	Nov	Nov	Dec	Dec
ALICE	ALICE						
			ATLAS	ATLAS			
CMS	CMS			CMS	CMS		
LHCb		LHCb					

Sep	Sep	Oct	Oct	Nov	Nov	Dec	Dec
ALICE	ALICE						
				ATLAS	ATLAS		
	CMS	CMS			CMS	CMS	
		LHCb	LHCb				

SC Communication



- **Service Challenge Wiki**
 - Takes over from service-radiant wiki/web-site used in SC1 & 2
<https://uimon.cern.ch/twiki/bin/view/LCG/LCGServiceChallenges>
 - Contains Tier-0 and Tier-1 contact/configuration information and work logs for SC teams
 - Direct link from LCG home page (see under Service Challenges)
- **Weekly phonecons on-going**
 - Dial-in number: +41227676000
 - Access code: 0164222
- **Daily service meetings for CERN teams from 27th June**
- **Technical communication through service-challenge-tech@cern.ch list**
- **What else is required by Tier-1s?**
 - Daily (or frequent) meetings during SC?

SC3 Summary



- **A lot has been achieved since SC2!**
 - Particularly in the areas of monitoring, services, procedures, documentation, delivery of pilots, LCG 2.5 release, other s/w ...
 - Integration of remaining T1s, adding T2s, ...
- **Good understanding and agreement on goals of SC3**
 - What services need to run where
 - Proposed metrics to define success
 - **Outline schedule - detailed resource requirements still sketchy**
- **Concerns about readiness to run production-level services**
 - Preparations are late, but lots of pressure and effort
 - Are enough resources available to run *services*?
 - **Backups, single points of failure, vacations, ...**
- **SC3 leads to real production services by end of year**
 - Must continue to run during preparations for SC4
- **This is the build up to the LHC service - must ensure that appropriate resources are behind it**
 - **Still a number of 'pressure points' and 'single points of failure'**

Service Challenge 4

SC4



Known Knowns



- Model for 'Production' much debated and now well understood
 - All stages from data taking leading into to end-user analysis
 - Has been exercised through experiment data challenges
 - Will also be covered during Service Phase of SC3
 - **Main goal is to thoroughly stress-test the service infrastructure**

- Data types, rates, flows that correspond to above all 'clear'
 - Processing, re-processing, stripping, AOD / TAG Production etc

- Roles played by different tiers, services that they offer, services that they require etc also understood
 - Services still not fully setup; in some cases software maybe...
 - Still a large number of Tier2s with no clear Tier1
 - Expect to continue to make good progress on this prior to SC4

- Current plan is for 50 days of data taking in 2007 @ $\times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$

- **Service Challenge schedule / delivery of production system unchanged**

Known Unknowns



- End-user analysis still a mystery
 - Can easily result in significant network bandwidth / support load
 - What is the model for Analysis Facilities?
 - Dedicated PROOF farms? 100+ nodes, 50+TB disk
 - Batch mode? Single stream? Parallel?
 - More coordination with ROOT team and ARDA needed
- Startup phase of LHC unknown
 - It will certainly not be like steady-state
 - Strong pressure to exploit **(needed)** distributed resources
 - There will be a strong presence at CERN, but nevertheless fundamental need to allow detector / physics groups outside have rapid / peer access to the data
- Emphasize this in paper selection for CHEP 2006

Service Challenge 4 - SC4



- SC4 starts April 2006
- SC4 ends with the deployment of the **FULL PRODUCTION SERVICE**
- **Deadline for component (production) delivery: end January 2006**
- **Adds further complexity over SC3 - 'extra dimensions'**
 - Additional components and services, e.g. COOL and other DB-related applications
 - Analysis Use Cases
 - SRM 2.1 features required by LHC experiments **← have to monitor progress!**
 - Most Tier2s, all Tier1s at full service level
 - Anything that dropped off list for SC3...
 - **Services oriented at analysis & end-user**
 - What implications for the sites?
- **Analysis farms:**
 - Batch-like analysis at some sites (no major impact on sites)
 - Large-scale parallel interactive analysis farms and major sites
 - (100 PCs + 10TB storage) x N
- **User community:**
 - No longer small (<5) team of production users
 - 20-30 work groups of 15-25 people
 - Large (100s - 1000s) numbers of users worldwide

Analysis Use Cases (HEPCAL II)



- **Production Analysis (PA)**
 - **Goals in Context** *Create AOD/TAG data from input for physics analysis groups*
 - **Actors** *Experiment production manager*
 - **Triggers** *Need input for "individual" analysis*

- **(Sub-)Group Level Analysis (GLA)**
 - **Goals in Context** *Refine AOD/TAG data from a previous analysis step*
 - **Actors** *Analysis-group production manager*
 - **Triggers** *Need input for refined "individual" analysis*

- **End User Analysis (EA)**
 - **Goals in Context** *Find "the" physics signal*
 - **Actors** *End User*
 - **Triggers** *Publish data and get the Nobel Prize :-)*

SC4 Timeline



- Now - September: clarification of SC4 Use Cases, components, requirements, services etc.
- October 2005: SRM 2.1 testing starts; FTS/MySQL; target for post-SC3 services
- January 31st 2006: basic components delivered and in place
- February / March: integration testing
- February: SC4 planning workshop at CHEP (w/e before)
- March 31st 2006: integration testing successfully completed
- April 2006: throughput tests
- May 1st 2006: Service Phase starts (note compressed schedule!)
- September 1st 2006: Initial LHC Service in stable operation
- Summer 2007: first LHC event data

SC4 Milestones



Date	Description
31 Jan 06	All required software for baseline services deployed and operational at all Tier-1s and at least 20 Tier-2 sites.
30 Apr 06	Service Challenge 4 Set-up: Set-up complete and basic service demonstrated. Performance and throughput tests complete: Performance goal for each Tier-1 is the nominal data rate that the centre must sustain during LHC operation (see Table 7.2 below) CERN-disk → network → Tier-1-tape. Throughput test goal is to maintain for three weeks an average throughput of 1.6 GB/s from disk at CERN to tape at the Tier-1 sites. All Tier-1 sites must participate. The service must be able to support the full computing model of each experiment, including simulation and end-user batch analysis at Tier-2 centres.
31 May 06	Service Challenge 4: Start of stable service phase, including all Tier-1s and 40 Tier-2 centres.
30 Sept 06	1.6 GB/s data recording demonstration at CERN: Data generator → disk → tape sustaining 1.6 GB/s for one week using the CASTOR mass storage system.
30 Sept 06	Initial LHC Service in operation: Capable of handling the full target data rate between CERN and Tier-1s (see Table 7.2). The service will be used for extended testing of the computing systems of the four experiments, for simulation and for processing of cosmic-ray data. During the following six months each site will build up to the full throughput needed for LHC operation, which is twice the nominal data rate.
1 Apr 07	LHC Service Commissioned: A series of performance, throughput and reliability tests completed to show readiness to operate continuously at the target data rate and at twice this data rate for sustained periods.

SC4 Use Cases (?)



Not covered so far in Service Challenges:

- T0 recording to tape (and then out)
- Reprocessing at T1s
- Calibrations & distribution of calibration data
- HEPCAL II Use Cases
- Individual (mini-) productions (if / as allowed)

Additional services to be included:

- Full VOMS integration
- COOL, other AA services, experiment-specific services (e.g. ATLAS HVS)
- PROOF? xrootd? (analysis services in general...)
- Testing of next generation IBM and STK tape drives

September SC3.5 workshop (Most likely a 'virtual' one...)



- SC3 experience
 - Sites
 - Experiments
 - Outlook for remainder of service phase
 - Requirements gathering from site + experiment view points + report (by two rapporteurs from above sessions)

- SC4 preparation
 - (recent) experiment goals / plans in terms of HEPCAL use cases
 - proof / xrootd / roles / plans
 - LCG SRM status
 - targets for SC4
 - T1 plans for incorporating T2s
 - T2 plans

Remaining Challenges



- Bring core services up to robust 24 x 7 standard required
- Bring remaining Tier2 centres into the process
- Identify the additional Use Cases and functionality for SC4
- Build a cohesive service out of distributed community
- Clarity; simplicity; ease-of-use; functionality

Work Items (= Milestones?)



- Sites to deliver full production storage systems
 - As demonstrated by SC3 Throughput tests
- Experiments to provide metrics, test-jobs for SC3 (where missing...)
- SRM developers to produce plan for SRM 2.1 features
 - Delivery of first versions Oct 2005; production Jan 2006
- T1s to provide 10Gbit/s network links to CERN
 - On-going... but (v.) late in some countries
- T1s/T2s to setup File Catalog services
 - Primarily LFC (ALICE, ATLAS)
- Service model for FTS / LFC etc on MySQL
 - T1(/2?) to prototype
- Resolution of VO boxes
 - Stop-gap for SC3
- Experiments to provide test jobs to measure site readiness
- Experiments to provide detailed resource requirements
- Agreement on SC4 Use Cases
- ...

Summary and Conclusions



- Mid-way (time wise) in aggressive programme to deploy world-wide production Grid
- Services must be provided no later than **year end...**
... and then run until **career end**
- Deploying production Grids is hard...
... and requires far too much effort at the various sites
- We need to reduce this effort...
... as well dramatically increase the ease of use
- Today hard for users to 'see' their science...
... main effort is overcome the complexity of the Grid

The Service is
the Challenge

