

LCG Services: Status of SC3

Jamie.Shiers@cern.ch

September 6 2005

SC3 Status - From July PEB

- **Going well:**
 - **All Tier1s** now involved and participating in transfers
 - Integration of **Tier2s** (progressed further still since GDB)
 - Understanding of **services** required by site; setup, monitoring, use(!)
 - Preparations for Service Phase (expts + sites in direct contact)
 - At least one site (BNL) has reached 150MB/s (disk) + 60MB/s (tape)
- **Main concern:**
 - SRM instabilities need to be resolved, then understand transfers
 - Need to run at 50MB/s per site **stably** (initially)
 - Once this can be achieved, increase rate to 100 then 150MB/s
- **Procedure:**
 - Take 1 - 2 sites (and install DPM?) and test (more later...)
 - dCache expert workshop at DESY end August - also needed for CASTOR?
 - Medium term goal: 150MB/s sustainable, low effort / site
 - Long term goal: 225 - 250MB/s sustainable
 - Establish regular T0+T1 technical / service meetings (4 x year?)
- **Storage is the primary area on which we need to focus**

pp data rates - 'weighted'

Centre	ALICE	ATLAS	CMS	LHCb	Rate into T1 (pp)
ASCC, Taipei	-	8%	10%	-	100
CNAF, Italy	7%	7%	13%	11%	200
PIC, Spain	-	5%	5%	6.5%	100
IN2P3, Lyon	9%	13%	10%	27%	200
GridKA, Germany	?	?	?	?	200
RAL, UK	-	7%	3%	15%	150
BNL, USA	-	22%	-	-	200
FNAL, USA	-	-	28%	-	200
TRIUMF, Canada	-	4%	-	-	50
NIKHEF/SARA, NL	3%	13%	-	23%	150
Nordic Data Grid Facility	6%	6%	-	-	50
Totals	-	-	-	-	1,600


***Full AOD & TAG to all T1s (probably not in early days)
...But ESD might well be larger than target...***

Nominal Data Rates


- 2 years before data taking can transfer from SRM at CERN to DPM at T1 at ~target data rate
- Stably, reliably, days on end
- Need to do this to all T1s at target data rates to tape
- Plus factor 2 for backlogs / peaks
- Need to have fully debugged recovery procedures
- Data flows from re-processing need to be discussed
 - New ESD copied back to CERN (and to another T1 for ATLAS)
 - AOD and TAG copied to other T1s, T0, T2s (subset for AOD?)

SC3 Throughput: Disk & Tape

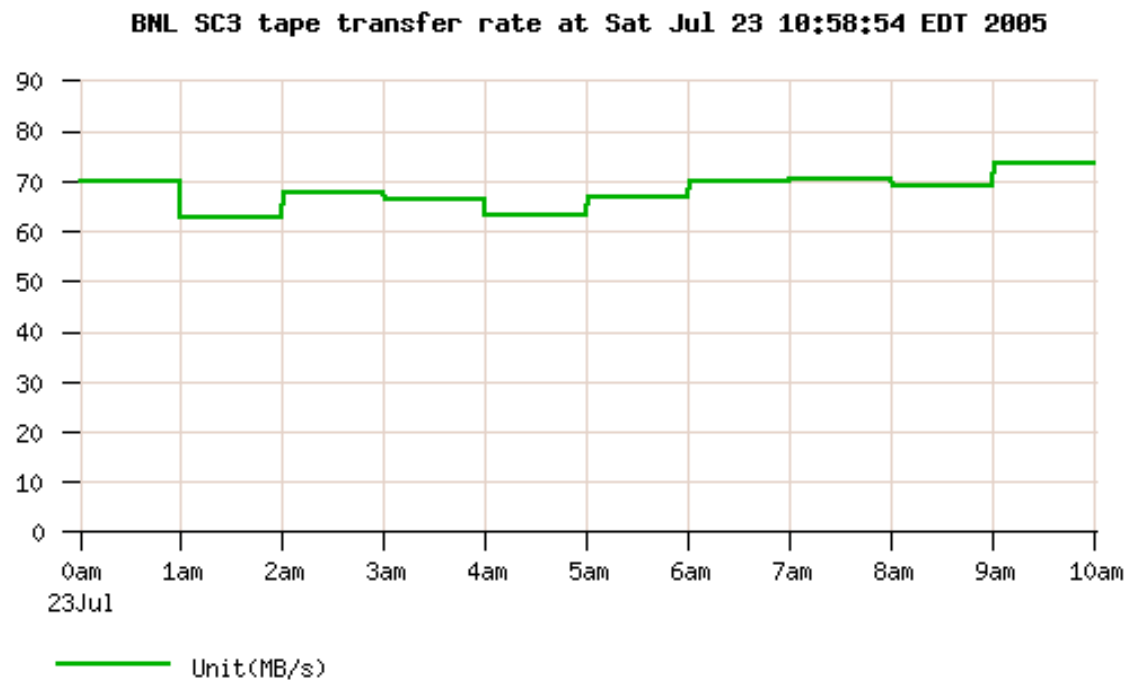

**Disk target:
150MB/s/site
1GB/s (CERN)**



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



Transfers to BNL



What Remained to be done? [In July]

- Baseline services setup at all participating sites
- Validation through sample jobs provided by experiments
- Agreement on resource requirements and schedule
- Agreement of metrics
- Resolution of outstanding issues (VO-boxes, experiment-specific services, clear definition of support lines, software components, releases and dependencies etc.)
- ...


This is monitored through weekly calls + site Wikis

Plan for August [In July]

- Need to understand transfers to various sites
- Propose first half August site debugging
 - ASCC, BNL, DESY, FNAL, GridKA, IN2P3, PIC, TRIUMF so far...
 - (dCache sites except ASCC + PIC)
- Need to eliminate possible sources of problem
 - Source, destination, tuning parameters etc
 - Try e.g. DPM@CERN->DPM@NDGF, DPM@CERN->dCache@DESY
 - Establish recommended dCache setup, try at FZK, SARA, (RAL)
- Eventually run multi-site throughput 2nd half August
- Service preparations need to continue in parallel
- (Beyond August - Service Phase)

Post-debugging (from JC)

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



Proposal # 1 - Reference Nodes

- Proposal is to place a "reference" node at each Tier1 site
 - 'standard' OS installed - i.e. SL3/4
 - The DPM software would be installed on it
- This can be used for system debugging, network tuning and regression tests
 - This would be used to run background iperf tests and file replication for regression analysis
 - Can remove a link from the transfer chain to eliminate source or destination SRM as cause of problem
- Does not have to be most up-to-date hardware
 - But good network connectivity and NIC essential
 - Modest storage requirements - data can be cleaned up
 - [1 TB is the smallest unit we talk about, right?]

Monitoring on-going transfers

- FTS used gridftp performance markers
 - Has 120 seconds marker-to-marker timeout
 - Has global transfer time set much higher (~1hr)
 - dCache does not send the performance markers
 - This initially caused all long-hop transfers to time out
 - Have to disable this feature
 - Had the effect of if any problem occurs, it takes 1hr to fail !
 - Bad for channel utilization
- dCache developers have promised to implement this feature

New FTS Version

- Provides:
 - Multi-VO support (scheduling);
 - VO agent framework (e.g. for experiment customisation, such as catalog manipulation)
 - Better monitoring information
- Goal:
 - Deploy production version ready for SC3 Service Phase

- This has been done

- Support for srmcp developed - deployment most likely end Sep / early Oct

Where can we do better?

- Despite significant effort devoted to documentation, communication, support lists etc - and not aided by moving targets - need to improve clarity of requirements for each site
- Propose to target this for SC4, adding incremental levels of detail as they become clear:
 - Step 1: Timeline, Target Data Rates, Site Responsibilities etc.
 - Step 2: Detailed List of Services by Site
 - Step 3: Detailed List of Use Cases to be Tested
 - ...
- Will also produce a single document describing the agreed setup by site for SC3
 - Attempts to do this in May this year generated quite some heat!

What Needs to Change?

- Current support is clearly an interim step towards standard users support / operations
 - See Wiki page for current details...
- Discuss in detail at operations workshop
- Target: full migration prior to Sep 2006 Service
- Implement gradually from now on as services become well understood?

Up-coming Schedule

- Scheduled discussed at July / August PEBs is below
- Main concern is support load at startup of Service Phase
- Understanding that multiple experiments will run in parallel from ~October on
 - But with support priority to 'scheduled' experiment

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

Re-run of Throughput Tests (to disk)

- Need to schedule re-run of disk-disk throughput tests

- Requirements:

- dCache 1.6.6 at all sites with recommended config
- castor2 clients and srm server upgraded at relevant sites

- Schedule ~1 week when experiment activity permits
- End-October / Early November??
- Prior to SC 2005???

September Summary

- Continuing progress in setup of services at different sites
- Intensive debugging during August has resulted in some clear actions - and already improved stability & throughput
- The Service Phase has started - ALICE is already active!
- Need to repeat disk - disk throughput tests once sites have performed necessary upgrades
- Preparations for SC4 and the Full Production Service are already underway!
- First planning document distributed - review site plans at October GDB (need to be submitted prior to end September)


CHEP Workshop 10 - 12 Feb, Mumbai

- Current thinking is to spend roughly the same amount of time (1 day) on the following three topics:
 1. SRM / storage related issues (dCache, CASTOR, DPM, ...)
 2. Services required at sites for LCG Computing
 3. Experiment / ARDA Use Cases for SC4

▪ Please remember to tick the Workshop box when you register for CHEP!

- Participation to Workshop on 'Service Challenges for LHC experiment': **Y**

<http://www.tifr.res.in/~chep06/>



Backup Slides

Slides from GDB / previous PEB

Introduction



- Neither SC1 nor SC2 *fully* met their goals
 - ☺ SC2 met / exceeded its throughput goals
 - ☹ **But not its service goals...**
- Multiple threads started early 2005 to address:
 - Bringing experiments into loop (SC3+)
 - Bringing T2s into loop (ditto)
 - Preparing for full production services
 - Addressing problems beyond 'throughput goals'
 - e.g. site / experiment goals, additional services etc
- ☺ **All Tier1s are now involved! Many Tier2s! New s/w successfully deployed!**
 - Will not comment on individual successes / issues - site slots for that!
- ☺ **Successful workshops, tutorials (April, May, June) and site visits!**
- Throughput tests gradually approaching target (more later)
- ☝ **Need to understand the problems areas and address them**
- ☺ **Acknowledge progress / successes / hard-work of many!**

Executive Summary (updated since PEB)

- '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
 - Using "Quick Fix" release from previous Friday...

- SC3 Throughput Tests have started!
 - Seeing '**SC2-level**' 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 we 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...**

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)
 - IN2P3 to host one copy; CNAF? Taiwan? RAL?
 - 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

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(other) area where things are looking good!

T2s

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

Services at CERN

- Building on 'standard service model'
 1. First level support: operations team
 - Box-level monitoring, reboot, alarms, procedures etc
 2. 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 Service 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
 3. Third level support teams: by service
 - Notified by 2nd level and / or through operators (by agreement)
 - Should be called (very) rarely... **(Definition of a service?)**

Services elsewhere

- Several services require DB behind them
 - CASTOR/dCache/DPM etc
 - FTS
 - LFC
- LFC (today) and FTS (October?) will support MySQL as well as Oracle database backend
 - CASTOR also does this today (PIC)
- Knowledge of community being leveraged to provide guidance - through Wiki - on how to do these
 - e.g. proposal for DB backup at T2s archiving recovery set at T1
 - (stop server; copy file & restart; archive at T1 or hot backup as sample options)

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 ...

SC3 - Deadlines and Deliverables

- May 31st 2005: basic components delivered and in place
- June 2005: integration testing
- June 13 - 15: planning workshop - experiment issues
- June 30th 2005: integration testing successfully completed
- July 1 - 10: start disk - disk throughput tests
 - Assume a number of false starts / difficulties
- July 11 - 20: disk tests
- July 21 - 27: tape tests
- July 28 - 31: T2 tests


Service 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 - cern.ch/LCG](http://cern.ch/LCG) -> [Service Challenges](#)
 - Contains Tier-0 and Tier-1 contact/configuration information and work logs for SC teams
- [Weekly phone-cons on-going](#)
 - Dial-in number: +41227676000
 - Access code: 0164222
- [Daily service meetings for CERN teams from 27th June](#)
 - B28 R-015: [standing agenda and minutes](#) via Wiki
- [Technical communication through service-challenge-tech@cern.ch list](#)
- [What else is required by Tier-1s?](#)
 - Daily (or frequent) meetings during SC?



SC Meetings / Workshops

- Not enough support for September workshop
 - Despite +ve feedback from April & June workshops
- Propose to continue with CHEP workshop nevertheless
- I believe weekly con-calls are useful
 - Judging on length / number of people joining etc
- There are still many issues we need to discuss / resolve
- Please bring up issues that worry you!
- GDBs in September / October?



SC3 Summary



- **There has been a great deal of progress 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'**