



Milestones for the LCG Service Challenges next 18 months

Kors Bos, NIKHEF, Amsterdam

Les Robertson, CERN

LCG - PEB Meeting, 14 December, 2004



Service Challenge Overview



- **The LCG Service Challenges aim to prototype the Tier-0/1/2 infrastructure needed at the time of LCG startup**
 - Series of milestones progressively building up the component services, performance, capacity
 - Interspersed with ***long stable periods for experiments*** to test out their computing models
 - Fixed end-point – all components in place at full performance level 6 months before first physics data from LHC
- **Many of the components already exist**
 - But have not been proven together
 - Nor at the required data rates and reliability levels
- **Need to get the service teams who already look after infrastructure connected**
 - Initial involvement – 30 people from ~6 sites
 - Need to have active involvement from experiments at an early stage



Principles - Network



- **Not a network bandwidth challenge**
 - In 2004 10 Gbit/sec has already been proven to be possible
- **But end-to-end sustained network service at high bandwidth is still a challenge**
- **International network topology is important**
 - *Last mile* is vital
 - What will the worldwide scientific network infrastructure look like in 2007?
 - How do we integrate evolving infrastructure plans in our tight schedule?
 - Who pays for what?
- **Network performance must include grid software**
 - Not only GridFTP but also overhead of higher-level systems like file transfer service, SRM and Grid Cataloging
- **Performance must include experiment specific hard/soft/people-ware**
 - *BUT : concentrate on generic issues first*



Principles - Software



- **Service challenges test end-to-end application: from the exp. DAQ to remote tape robot**
 - Progress to be made in steps by adding more components each step
 - Start with the core and work outwards
 - file transfer → reliable file transfer service
 - disk<->disk → tape<->tape
 - Static job & data mix → dynamic job generation / Resource Broker
 -
- **Sustainability for experiment model tests is a challenge**
 - 24 hours/day for 7 months in a row
- **Redundancy and fail-over tests**
 - Data buffers for non-stop operation
 - Policy driven e.g. If one site fails other sites must take more



Tier-1 Centres (*December 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			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
				5	10	7	6	28



In parallel with the service challenge and computing model validation activity there is a permanent *baseline* service for all experiments for production, analysis, etc.

So each centre must operate this permanent service *in addition* to taking part in the service challenges

This will require hardware and people!!



2005 Q1(i)



SC2 - Robust Data Transfer Challenge

Set up infrastructure for 6 sites

- Fermi, NIKHEF/SARA, GridKa, RAL, CNAF, CCIN2P3


Test sites individually

– at least two at 500 MByte/s with CERN

Agree on sustained data rates for each participating centre

Goal – by end March sustained 500 Mbytes/s aggregate at CERN

In parallel - serve the ATLAS “Tier0 tests” (needs more discussion)





2005 Q1(ii)



In parallel with SC2 – prepare for the next service challenge (SC3)

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





2005 Q2-3(i)



SC3 - 50% service infrastructure

- Same T1s as in SC2 (Fermi, NIKHEF/SARA, GridKa, RAL, CNAF, CCIN2P3)
- Add at least two T2s
- “50%” means approximately 50% of the nominal rate of ATLAS+CMS (**may no longer be ≤ 300 MByte/s!**)


Using the 1 GByte/s *challenge* facility at CERN -

- Disk at T0 to tape at all T1 sites at 50 Mbyte/s
- Data recording at T0 from same disk buffers
- Moderate traffic disk-disk between T1s and T2s

Use ATLAS and CMS files, reconstruction, ESD skimming codes
(numbers to be worked out when the models are published)

Goal - 1 month sustained service in July

- **300 MBytes/s aggregate at CERN, 50 MBytes/s at each T1**





2005 Q2-3(ii)




In parallel with SC3 prepare additional centres using the 500 MByte/s test facility

- Test Taipei, Vancouver, Brookhaven, additional Tier-2s

Further develop framework software

- Catalogues, VO's, use experiment specific solutions





2005 – September-December (i)




50% Computing Model Validation Period

The service exercised in SC3 is made available to experiments for computing model tests

Additional sites are added as they come up to speed

End-to-end data rates –

- 300 Mbytes/s at CERN (aggregate)
- 50 Mbytes/s at Tier-1s
- Modest





2005 – September-December (ii)




**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 2 GBytes/s

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





2006 - January-August



SC4 – full computing model services


- Tier-0, **ALL** Tier-1s, all major Tier-2s operational at full target data rates (~1.2 GB/sec at Tier-0 – **needs revision?**)
- 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






2006 - September



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 6 months before first physics data





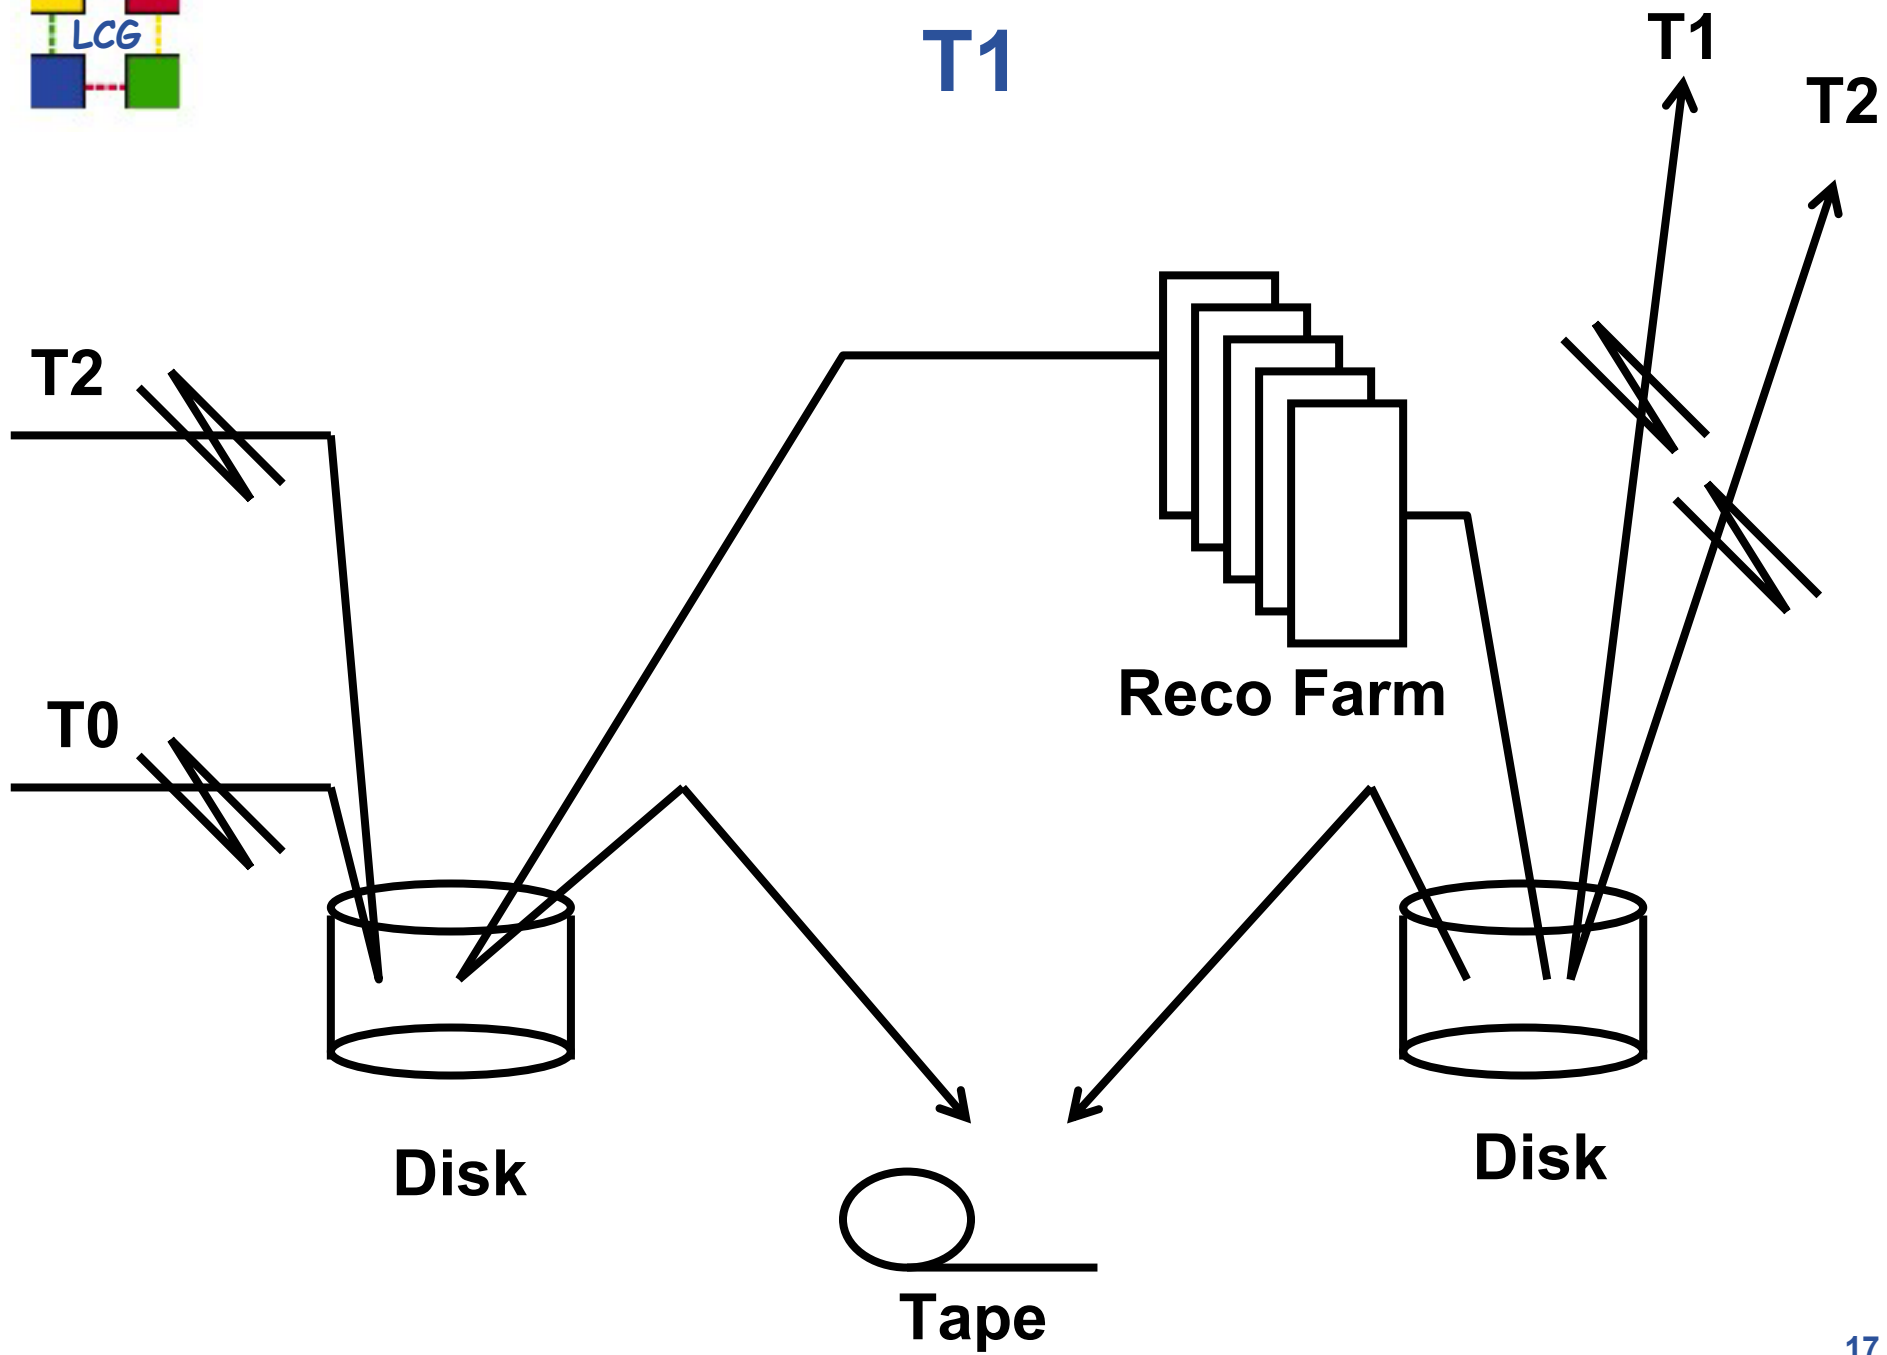
Rates

	Q1			Q2			Q3			Q4		
	network Gbit/s	transfer Mbyte/s	to tape	network Gbit/s	transfer Mbyte/s	to tape Mbyte/s	network Gbit/s	transfer Mbyte/s	to tape Mbyte/s	network Gbit/s	transfer Mbyte/s	to tape Mbyte/s
FNAL	10	500	0	10	250	50	10	500	75	10	500	100
SARA	4	200	0	4	250	50	10	300	75	10	300	100
FZK	10	500	0	10	250	50	10	300	75	10	300	100
IN2P3	1	100	0	1	100	50	10	300	75	10	300	100
RAL	2	200	0	2	100	50	2	100	50	2	100	50
CNAF	1	100	0	1	100	50	1	100	50	1	100	50
BNL				1	100	0	1	100	50	1	100	50
Taipei				1	100	0	1	100	50	1	100	50
Triumf				1	100	0	1	100	50	1	100	50
Nordic							1	100	0	1	100	50
PIC							1	100	0	1	100	50





T1





Current T1 Status



	RAL	Fermilab	Brookhaven	Karlsruhe	IN2P3	CNAF	PIC
1. Network Configuration	Contact	√	√	√	√		Contact
2. Storage Configuration		√	√	...	√		
3. Site Tuning			
4. "Service Challenge"		Jan 05		Jan 05			

	Taipei	Nordugrid	TRIUMF	NL
1. Network Configuration				√
2. Storage Configuration				√
3. Site Tuning				...
4. "Service Challenge"				Dec 04

KEY
 √ = Complete
 ... = In Progress