



Update on experiment requirements

Federico Carminati
CERN, Ginevra

At the beginning there was BS





Baseline services

- Nothing really surprising here - but a lot was clarified in terms of requirements, implementations, deployment, security, etc

- Storage management services
 - Based on SRM as the interface
- Basic transfer services
 - gridFTP, srmCopy
- Reliable file transfer service
- Grid catalogue services
- Catalogue and data management tools
- Database services
 - Required at Tier1,2
- Compute Resource Services
- Workload management

- VO management services
 - Clear need for VOMS: roles, groups, subgroups
- POSIX-like I/O service
 - local files, and include links to catalogues
- Grid monitoring tools and services
 - Focussed on job monitoring
- VO agent framework
- Applications software installation service
- Reliable messaging service
- Information system

BS WG conclusions

A: common solution mandatory;
 B: common solution required, alternative exists
 C: common solution desirable

Service	ALICE	ATLAS	CMS	LHCb
Storage Element	A	A	A	A
Basic Transfer Tools	A	A	A	A
Reliable File Transfer Service	A	A	A/B	A
Catalogue Service	B	B	B	B
Catalogue and DM tools	C	C	C	C
Compute Element	A	A	A	A
Workload Management	A/B	A	A	C
VO agents	A	A	A	A
VOMS	A	A	A	A
Database Services	A	A	A	A
Posix I/O	C	C	C	C
Application Software Installation	C	C	C	C
Job Monitoring Tools	C	C	C	C
Reliable Messaging Service	C	C	C	C
Information System	A	A	A	A



Then came SC3



SC3

- Most of the experiments planned to verify (one part of) their computing model
 - One central point was the export of data from T0 to T1's
- A large number of problems emerged
 - FTS – migration from “test-bed” to production mode exposed the functional limitations and stability of the service
 - CASTOR 2 – integration in the experiments frameworks difficult – performance, stability and functionality limited
 - SRM implementations
 - LFC performance
 - Stability and functionality of sites
- Access to data still seems problematic
- Target performance figure were not reached
- Difficult to reconcile / overload experiment and service goals



SC3

- Negative side
 - We could reach only partially our e-2-e goals
 - Tests limited to few services, not the same for all the experiments
 - Computing models still untested in full
 - The fact that we could not fully use the resources we asked for is turning into a political embarrassment
- Positive site
 - Tremendous learning exercise
 - Many problems that were there have been solved
 - Everybody has been very collaborative and helpful
 - The taskforces have demonstrated to be effective tools for cooperation
- We have now a clearer idea of our requirements



Summary of requirements



Security, authorisation, authentication

- VOMS available and stable
- Groups and roles used by all middleware
 - How to propagate to WN, DN?
- Automatic proxy & kerberos credential renewal
- Recommendations on how to develop experiment specific secure services
 - Secure service using delegated and automatically renewed user credentials
 - API or “development guide” for security delegation



Information System

- Stable access to static Information System
 - BDII or equivalent (!)
 - Split static and dynamic information?
- Access to the IS from the WN
- Publishing experiment specific info
- More complete use of Glue schema
 - Should be the same in gLite and LCG



Storage Management

- SRM 2.1 supported by all SE Services
 - Space reservation, file pinning, bulk operations
 - Disk quota management at group and user level
- Homogeneous, consistent and efficient implementation
 - Smooth transition SRM 1->2
 - SE interoperability issues must be solved
 - Operations should be “safe”
- SRM client libraries should be available to the applications
 - Full SRM 2.1 for CASTOR, d-Cache and DPM not before 3Q06?
 - Need also xrootd \Rightarrow SRM and SRM \Rightarrow xrootd (?)
- CASTOR 2
 - Pre-staging for file transfer
 - Coherent usage of multiple pools
- Do we have a consistent model for data access?



File transfer

- Applications would prefer to talk with FPS (sic!)
 - Routing, plug-ins, multiple replica / broadcast, proxy renewal, error and timeout handling
 - But retain control of FC update via plugins
- Improved FTS Clients on all sites, WNs and VOBOXes
- Ideally GUID (SE1 \Rightarrow SE2), however (SURL,SE1) \Rightarrow (SURL,SE2) acceptable
- Central entry point and transparent mechanism for all transfers Tx \leftrightarrow Ty
 - Going beyond static FTS channels (being exposed to users!)
- Full integration with SRM
- Possibility to specify type of space, lifetime of a pinned file, etc.
 - Need for tactical “hot” storage
- Dynamic priorities



File catalogue services

- LFC as global and/or local file catalogue
 - Support for replica attributes: tape, tape with cache, pinned cache, disk, archived tape, custodial flag
- POOL interface to LFC
 - Access to metadata?
- Performance
 - Read access, different kinds of queries
 - Unauthenticated read-only access, deep levels of queries should be possible for catalogue administrators even with reduced performance
 - In general, the access frequency is estimated to be $\sim 1\text{Hz}$
- Bulk operations for file and replica registration should be supported



Grid data management tools

- lcg-utils available in production
- C/C++ POSIX file access based on the LFN
 - gfal library
 - Efficient choice of “best replica” in a running job
- Possibility of multiple instances of LFC for high availability and efficiency
- Reliable registration service
- ACL support and propagation
- Bulk operations (registration, replication, staging, deletion)



Workload management

- Stable and redundant service
 - Load balance and failover
- Match making based on CPU slots / file location
- Efficient input sandbox management
- Latency for job execution and status reporting proportional to expected job duration
- Support for different priorities based on VOMS groups/roles
 - Rearrange jobs in the central and local TQ's
- Fair share across users in the same group
- Interactive access to running job sandbox (debugging)
- Computing Element accessible by services/clients other than RB
 - Monitoring and submission
- Changing of identity of a job running on the worker nodes
- Standard CPU time limits



Monitoring / accounting

- Need to monitor
 - Transfer traffic, statistics for file opening and I/O by file/dataset from SE's
 - VO specific information for global operations
 - Job status/failure/progress information
 - Publish/Subscribe to logging, bookkeeping and local batch system events for all jobs in the VO
 - Add heartbeat to SFT?
- Accounting
 - Site, user and (VOMS) group granularity
 - Aggregate by VO (user) specified tag
 - Application type (MC, Reconstruction...), executable, dataset
 - SE accounting aggregated by datasets (e.g. PFN directory)



Various

- xrootd & VOBOX at all sites
 - Should be considered a basic Grid service
 - Support model?
- Tools to for site dependent VO environment
- Secure hosting of long-lived processes
 - Standard set of secure containers
- proof?
 - Precise requirements being defined



And now the bright future



SC4

- Our last chance to test the computing model before the “reality check”
- We have to test there all the components that we intend to use for data taking
- Question is
 - Which of the requirements will be satisfied?
 - By whom?
 - Which timescale?



EGEE / LCG

- EGEE is going to “tag” gLite 1.5 by year’s end
- This will be merged with EGEE 2.7 into gLite 3.0 around March
- Experiments needs badly stability, and some need the new features of gLite
- Timing with SC4 is critical
- We have to make sure that LCG/GD, the taskforces and EGEE JRA1 have a clear idea of our requirements and priorities and work together
 - But how?



Metrics of success

- We should make a special effort to define metrics
- I am not sure we know
 - “how far” from the mark we fail
 - “how serious” is a failure
 - At least this is true for ALICE!
- This is important to provide clear targets to developers



Plans for SC4

- Experiments plans for SC4 are
 - Do what they could not do in SC3
 - Add more components of the computing model
- An attentive MW choice will be instrumental in ensuring success
 - Here I am very much afraid of riches' embarrassment



ALICE plans

- Quantitative plans still under study
 - Data storage / reconstruction / analysis on a level of 20% from a standard data taking year
- “Classical” approach
 - Large distributed MC production @ T1/2
 - Copy of the files back at CASTOR @ CERN
 - “push out” of RAW with FTS (hopefully FPS)
 - Quasi online reconstruction @ T0
 - Distributed reconstruction @ T1’s
 - Distributed analysis
 - Special emphasis on user data analysis - fast and efficient access to
 - ESDs, resource sharing production/analysis
- Usage of Raw data, alignment and calibration must be in
- Planning document in preparation



LHCb DC'06

2 data challenges envisaged in 2006

Aim to start production for 1st of these challenges end of Feb'06 - with a finalised event model

~200M events for 1st DC

Milestone	Due date	External Dependencies
2005		
Example of sub-detector alignment	September	
Analysis at all Tier-1's	November	
End of review of reconstruction software, including event model	December	
2006		
Final alignment strategy	March	
Start data processing phase of DC'06 (i) Distribution of RAW data from CERN (ii) Reconstruction/stripping at Tier-1's including CERN (iii) DST distribution to CERN & other Tier-1's	May	New version of SRM at all Tier-1's LCG services/tools from baseline service working group deployed at all Tier-1's
Alignment/calibration challenge - participation of all Tier-1's (i) Align/calibrate detector (ii) Distribute DB slice - synchronize remote DB's (iii) Reconstruct data	October	DB service to support COOL at all Tier-1's
2007		
Permanent Monte Carlo production mode ready for data taking	January	
Production system and software ready for data taking	April	

Table 6-3: High-level milestones for the computing project

LHCb LHCC computing milestones

Minimum Requirements:

LHCb

Middleware

- Production Service should NOT be testing ground for new m/w components e.g. FireMan, gLite RB, ...
- New m/w components have to be stress-tested (in a production environment) well in advance

A central FTS service

- Supporting Tx \leftrightarrow Ty transfers transparently
- FTS should handle staging

SRM (v2.1) -based SE's

- important consistent implementation across back-ends
 - Pinning/Unpinning will be essential if FTS cannot handle staging staging at time of DC04
 - Disk and MSS separate SEs at all Tier-1's
-

Minimum Requirements (cont'd):

LHCb

VO boxes to allow improved monitoring & add redundancy

- discussion at January workshop

Central LFC catalog (at CERN, read-write)

- read-only copy at Tier-1's
- read-only insecure version

IMPORTANT: data will be accessed directly from SE through protocols supported by ROOT/POOL & NOT only by GridFTP/srmcp

This is NOT a trivial statement ! It has become apparent some sites had not appreciated this



CMS input to SC4

CMS participation in SC4 will go through CMS Integration Program

→ Main goal: incremental integration of new components until Summer 2006.

For CMS SC4 is preparation for the CSA2006 milestone (Sep-Nov 2006)

→ Simulated data taking at 50Hz, all the way from T0 to T2

È CMS SC4 goals: put in place the infrastructure for CSA2006

È main goal: deliver a working service for CMS

CMS goals require a stable service from WLCG

È Most important to concentrate on improving reliability

- we need reliable components to build on so that CMS can deploy and integrate its baseline computing & software environment

È WLCG should also concentrate on improving reliability

- Establish working baseline early in 2006
 - the CMS end-to-end goals of SC3 were not all realized
- Add new components (sites?) after integration and testing at production scale, one by one if needed



CMS Priorities for 2006

Data management

- ➔ file transfer failures rates and required retries are far too high
 - 2006 must be the year of SRM being put under control
 - full interoperability of SRM servers/clients/sites
 - need to solve problems rather than hide problems under another tool layer
 - require smooth transition to SRM v2
- È Data serving capacity at WLCG sites needs to come into focus
- È Disk/Tape/Disk functionality has to be established at all T1 ☺

Job management

- È current system has scale limitations, but required base functionality
- È transition to gLite 3.0.0 needs to be non-disruptive
- È want to tackle authorization and prioritization, but scale highest prio

Distributed Data Bases

- È CMS is already testing FroNtier/Squid , expected not to be an issue



Quantitative goals

Need to maintain progress toward reliable services that scale and are on a realistic path for scaling to the experiment needs.

Service targets for the end of 2006:

→ Workload management: 100K jobs/day over LCG+OSG

- $\geq 90\%$ success rate (besides downtime)

È Data serving to jobs at sites: whichever is bigger between

- 1MB/sec/execution-slot
- 100MB/sec (Tier2) or 400MB/sec (Tier1)

È WAN transfer rates (all simultaneously, T1 rate goes to tape):

- From Tier0 to each Tier1 (RECO for 2nd pass): 5TB/day
- From each Tier1 to any other Tier1 (AOD replica): 1TB/day
- From each Tier1 to 4 Tier2 (Dataset distribution): 4TB/day (aggregate)
- From each Tier2 to one Tier1 (MC storage): 1TB/day
- $>99\%$ success rate at first attempt (besides downtime)



ATLAS SC4 Tests

- Complete Tier-0 test
 - Internal data transfer from "Event Filter" farm to Castor disk pool, Castor tape, CPU farm
 - n Calibration loop and handling of conditions data
 - n Transfer of RAW, ESD, AOD and TAG data to Tier-1s
 - n Transfer of AOD and TAG data to Tier-2s
 - n Data and dataset registration in DB (add meta-data information to meta-data DB)
- | Distributed production
 - n Full simulation chain run at Tier-2s (and Tier-1s)
 - n Data distribution to Tier-1s, other Tier-2s and CAF
- | Distributed analysis
 - n "Random" job submission accessing data at Tier-1s (some) and Tier-2s (mostly)
 - n Tests of performance of job submission, distribution and output retrieval



ATLAS SC4 Plans

- These tests run at first independently, then concurrently
- Timescale between April and October 2006
 - n Tier-0 tests in blocks of increasing scope, March, May, July
- | Every test must be preceded by a period of "service preparation" tests
 - n Checking that the underlying infrastructure is working properly



ATLAS SC4 Requirements

- Active FTS channels between all sites
 - FTS useable by anyone with a valid certificate
- | SRM "baseline WG version" deployed everywhere
- | Disk-only areas at all SEs
 - n No sudden migration of files to tape
- | Agreed (and secure) way to deploy experiment 's services
- | Full implementation of VOMS groups and policies for job submission and data management
- | Queues with different priorities for production and analysis jobs
- | Stability of middleware and service infrastructure
 - n "production quality"
- | Services run as services
 - n No sudden maintenance or upgrade interruptions

Conclusions

- We are observing steady progress in understanding and fixing problems
 - However the derivative is still not right
 - Every SC has increasingly ambitious objectives and an increasing backlog of problems to fix!
- We have to make sure that all available resources work in the same direction and with the same objectives / priorities
- (Some) experiment specific solutions are unavoidable and should be supported in a form acceptable to all parties
 - We have just no time for conflicts, we need to make work what's there, like or loath it
- Proper prioritisation of experiment requirements is now mandatory in order to satisfy them

