LCG LHCC Review

Computing Fabric Overview and Status

Goal

- The goal of the Computing Fabric Area is to prepare the T0 and T1 centre at CERN. The T0 part focuses on the mass storage of the raw data, the first processing of these and the data export (e.g. raw data copies), while the T1 centre task is primarily the analysis part.

- There is currently no physical or financial distinction/separation between the T0 installation and the T1 installation at CERN. (roughly 2/3 to 1/3)

- The plan is to have a flexible, performing and efficient installation based on the current model, to be verified until 2005 taking the computing models from the Experiments as input (Phase I of the LCG project).

Strategy

- **Continue, evolve and expand the current system**
  profit from the current experience : number of total users will not change, Physics Data Challenges of LHC experiments, running Experiments (CDR of COMPASS + NA48 up 150 MB/s, they run their level 3 filter on Lxbatch)

BUT do in parallel :

- **R&D activities and Technology evaluations**
  SAN versus NAS, iSCSI, IA64 processors, ….
  PASTA, infiniband clusters, new filesystem technologies,…..

- **Computing Data Challenges to test scalabilities on larger scales**
  "bring the system to it's limit and beyond "
  we are very successful already with this approach, especially with the "beyond" part

- **Watch carefully the market trends**

View of different Fabric areas

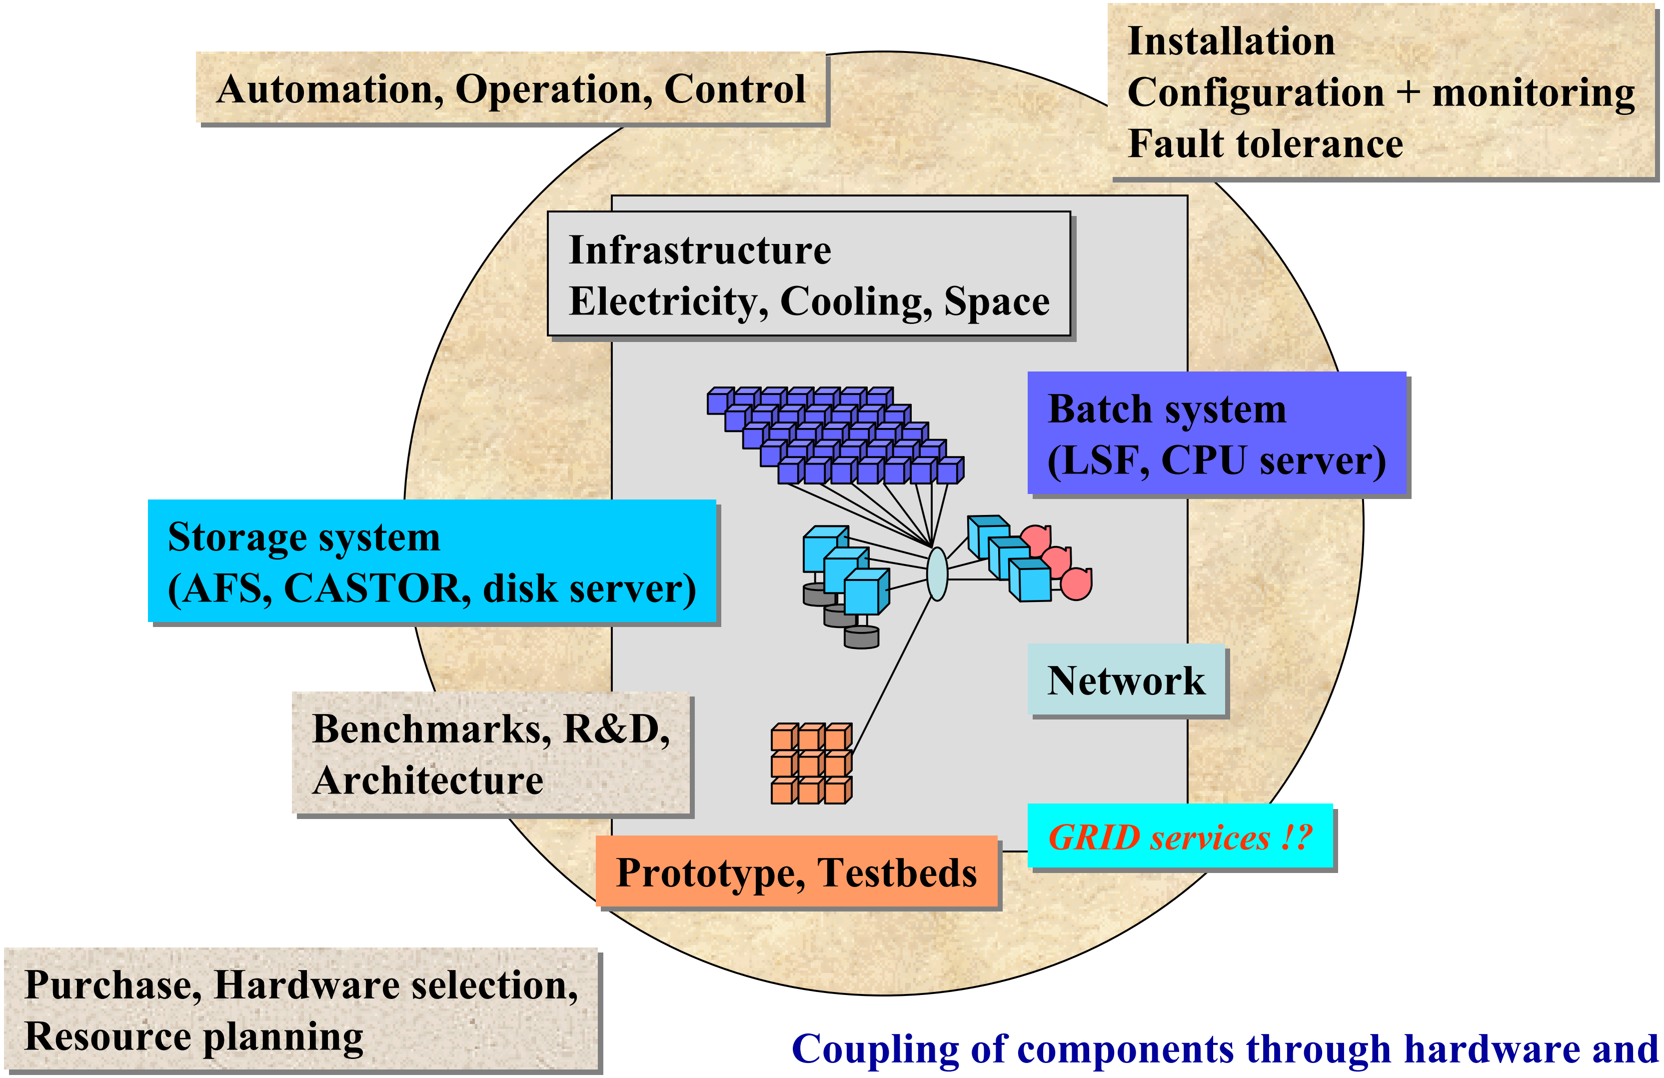

Coupling of components through hardware and software

Infrastructure

There are several components which make up the Fabric Infrastructure :

- **Material Flow**
  organization of market surveys and tenders, choice of hardware, feedback from R&D, inventories, vendor maintenance, replacement of hardware
  → major point is currently the negotiation of different purchasing procedures for the procurement of equipment in 2006

- **Electricity and cooling**
  refurbishment of the computer center to upgrade the available power from 0.8 MW today to 1.6 MW (2007) and 2.5 MW in 2008
  → development of power consumption in processors problematic

- **Automation procedures**
  Installation+Configuration+Monitoring+Fault Tolerance for all nodes
  Development based on the tools from the DataGrid project
  Already deployed on 1500 nodes, good experience, still some work to be done
  several Milestones met with little delay

Purchase

- Even with the delayed start up, large numbers of CPU & disk servers will be needed during 2006-8:
  - At least 2,600 CPU servers
    - 1,200 in peak year; c.f. purchases of 400/batch today
  - At least 1,400 disk servers
    - 550 in peak year; c.f. purchases of 70/batch today
  - Total budget: 20MCHF
- Build on our experiences to select hardware with minimal total cost of ownership.
  - Balance purchase cost against long term staff support costs, especially for
    - System management (see next section of this talk), and
    - Hardware maintenance.

➢ **Total Cost of Ownership workshop organised by the openlab, 11/12th November.**
→ we have already a very good understanding of our TCO !

Acquisition Milestones

- Agreement with SPL on acquisition strategy by December (Milestone 1.2.6.2).
  - Essential to have early involvement of SPL division given questions about purchase policy
  - likely need to select multiple vendors to ensure continuity of supply.
  - A little late, mostly due to changes in CERN structure.

- Issue Market Survey by 1st July 2004 (Milestone 1.2.6.3)
  - Based on our view of the hardware required, identify potential suppliers.
  - Input from SPL important in preparation of the Market Survey to ensure adequate qualification criteria for the suppliers.
    - Overall process will include visits to potential suppliers.

- Finance Committee Adjudication in September 2005 (Milestone 1.2.6.6)

The Power problem

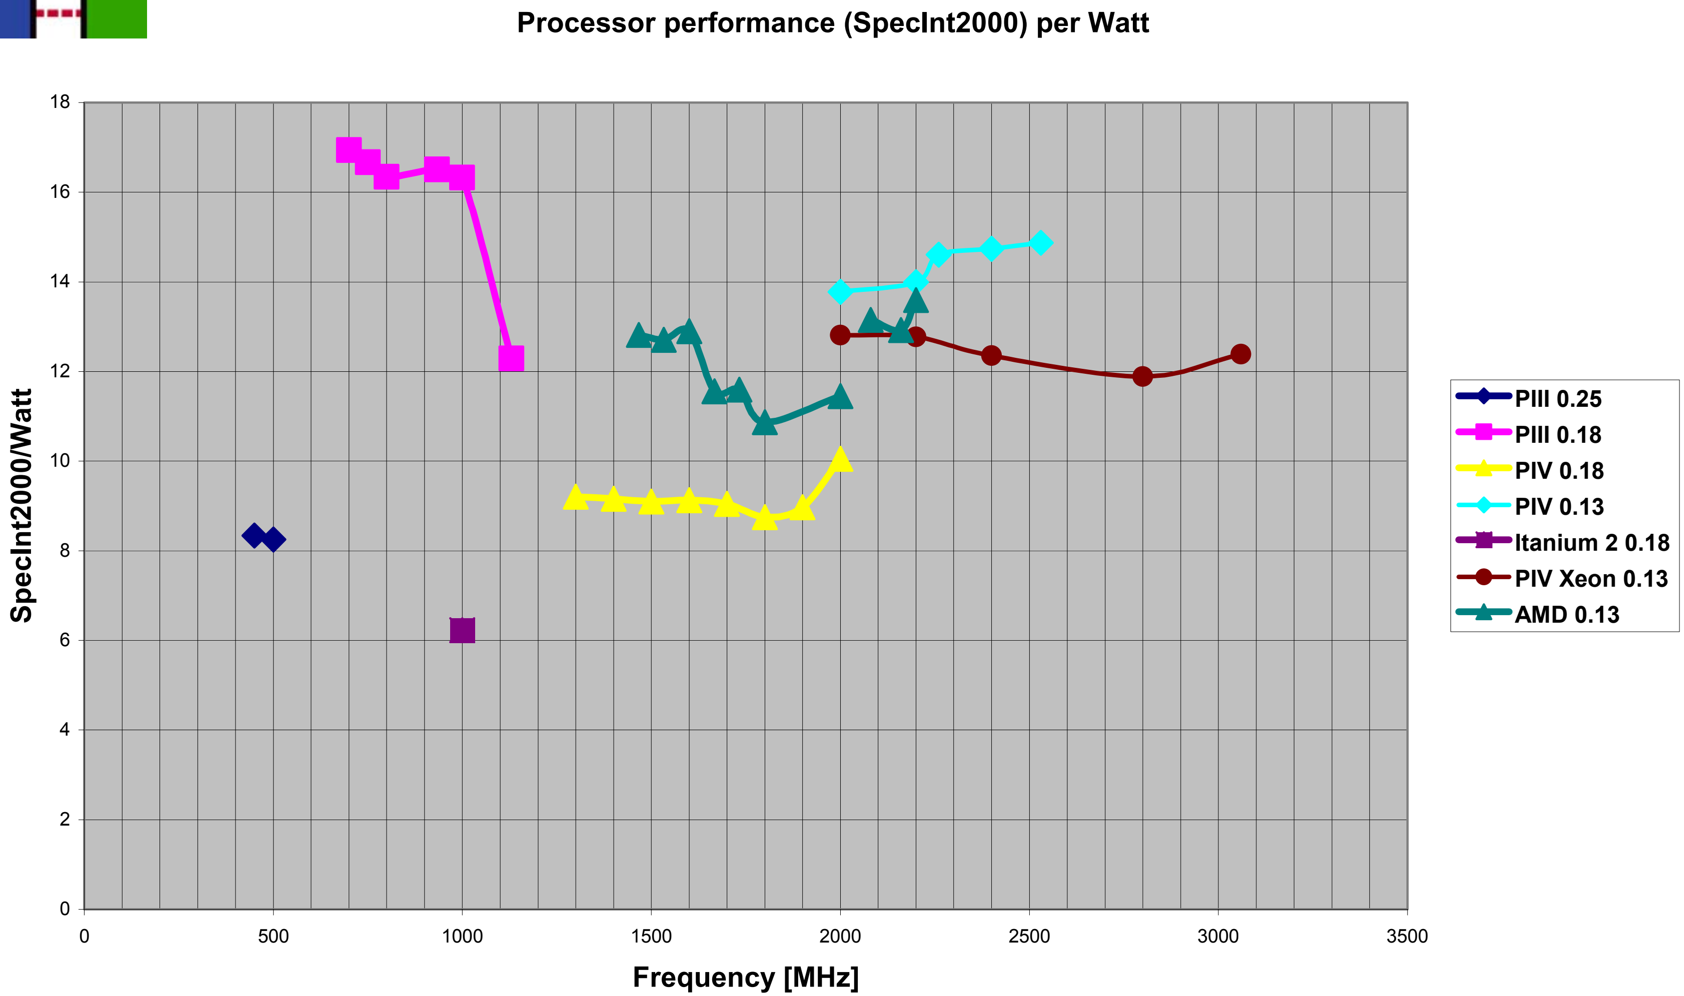

- Node power has increased from ~100W in 1999 to ~200W today, steady, linear growth
  - And, despite promises from vendors, electrical power demand seems to be directly related to Spec power:

Upgrade Timeline

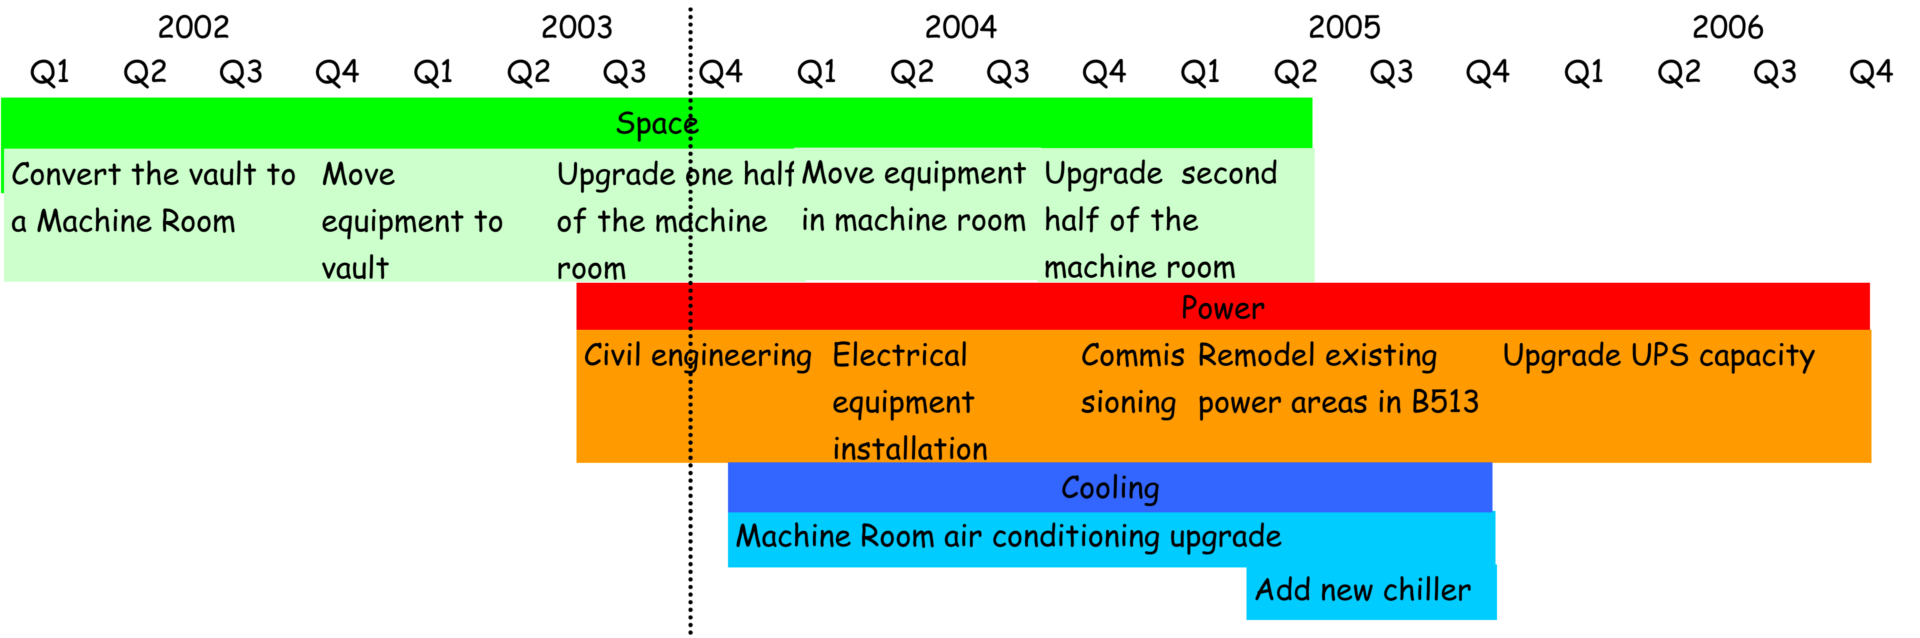

- **The power/space problem was recognised in 1999 and an upgrade plan developed after studies in 2000/1.**
- **Cost: 9.3MCHF, of which 4.3MCHF is for the new substation.**
  - **Vault upgrade was on budget. Substation civil engineering is overbudget (200KCHF), but there are potential savings in the electrical distribution.**
  - **Still some uncertainty on overall costs for air-conditioning upgrade.**

Substation Building

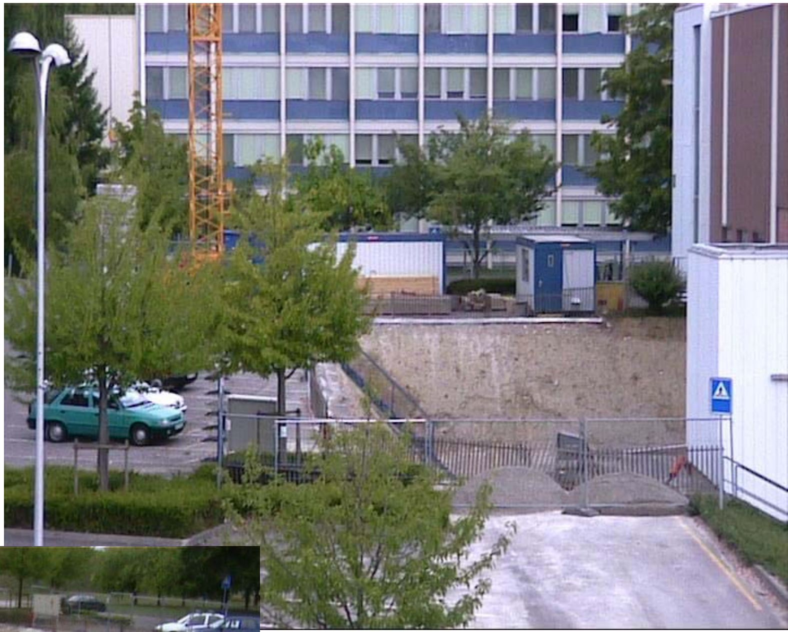

Milestone 1.2.3.3
Sub-station civil engineering starts
01-September 2003

→ started on the 18th of August

The new computer room in the vault of building 513 is now being populated

While the old room is being cleared for renovation

Upgrade Milestones

Vault converted to machine room	01/11/02	✓ 01/11/02
Right half of machine room emptied to the vault	01/08/03	✓ 15/08/03
Substation civil engineering starts	01/09/03	✓15/08/03
Substation civil engineering finishes	01/03/04	
Substation electrical installation starts	01/04/04	
Substation commissioned	07/01/05	
Elec. distrib. on RH of machine room upgraded	02/02/04	
Left half of machine room emptied	02/08/04	
Elec. distrib. on LH of machine room upgraded	01/06/05	
Machine room HVAC upgraded	01/03/05	
New 800kW UPS for physics installed	01/04/06	
Current UPS area remodelled	01/01/07	
2nd 800kW UPS added	02/04/07	
3rd 800kW UPS added	01/04/08	

On Schedule

Progress acceptable

Capacity will be installed to meet power needs.

Space and Power Summary

- Building infrastructure will be ready to support installation of production offline computing equipment from January 2006.

- The planned 2.5MW capacity will be OK for 1st year at full luminosity, but there is concern that this will not be adequate in the longer term.

- Our worst case scenario is a load of 4MW in 2010.
  - Studies show this can be met in B513, but more likely solution is to use space elsewhere on the CERN site.
  - Provision of extra power would be a 3 year project. We have time, therefore, but still need to keep a close eye on the evolution of power demand.

Fabric Management (I)

- The ELFms Large Fabric management system has been developed over the past few years to enable tight and precise control over all aspects of the local computing fabric.

- ELFms comprises
  - The EDG/WP4 quattor installation & configuration tools
  - The EDG/WP4 monitoring system, Lemon, and
  - LEAF, the LHC Era Automated Fabric system

Fabric Management (II)

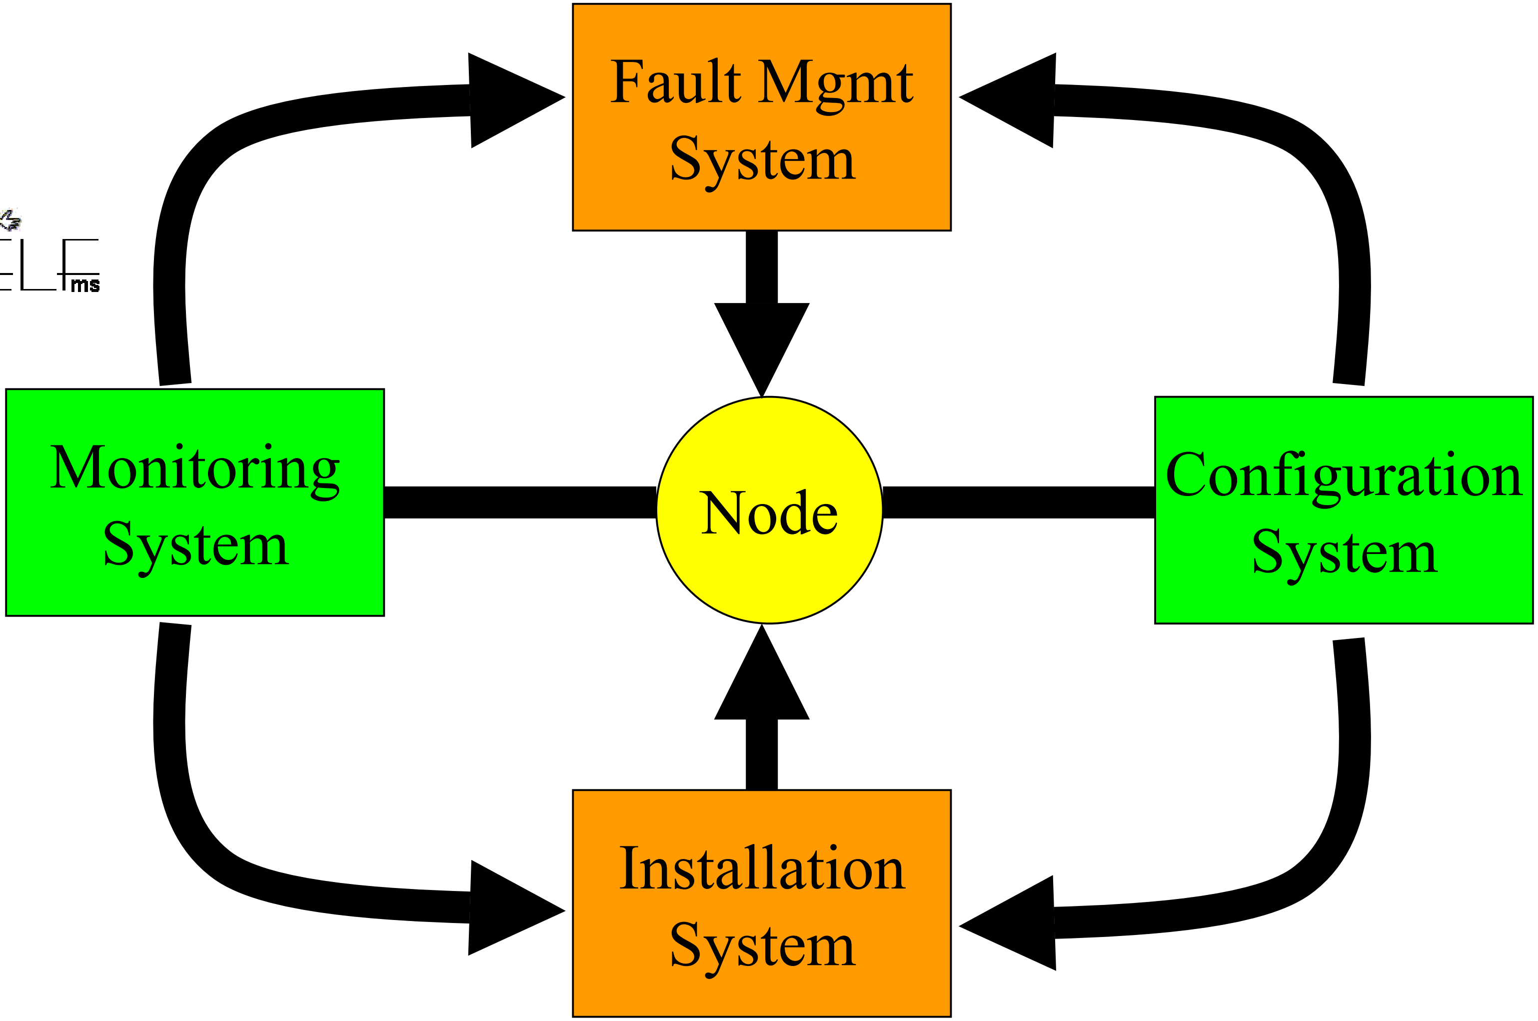

Installation+Configuration Status

- quattor is in complete control of our farms (1500 nodes).
  milestones with minimal delays on time

- We are already seeing the benefits in terms of
  - ease of installation—10 minutes for LSF upgrade,
  - speed of reaction—ssh security patch installed across all lxplus & lxbatch nodes within 1 hour of availability, and
  - homogeneous software state across the farms.

- quattor development is not complete, but future developments are desirable features, not critical issues.
  - Growing interest from elsewhere—good push to improve documentation and packaging!
  - Ported to Solaris by IT/PS

- EDG/WP4 has delivered as required.

Monitoring

Monitoring Sensor Agent
- Calls plug-in sensors to sample configured metrics
- Stores all collected data in a local disk buffer
- Sends the collected data to the global repository

Plug-in sensors
- Programs/scripts that implements a simple sensor-agent ASCII text protocol
- A C++ interface class is provided on top of the text protocol to facilitate implementation of new sensors

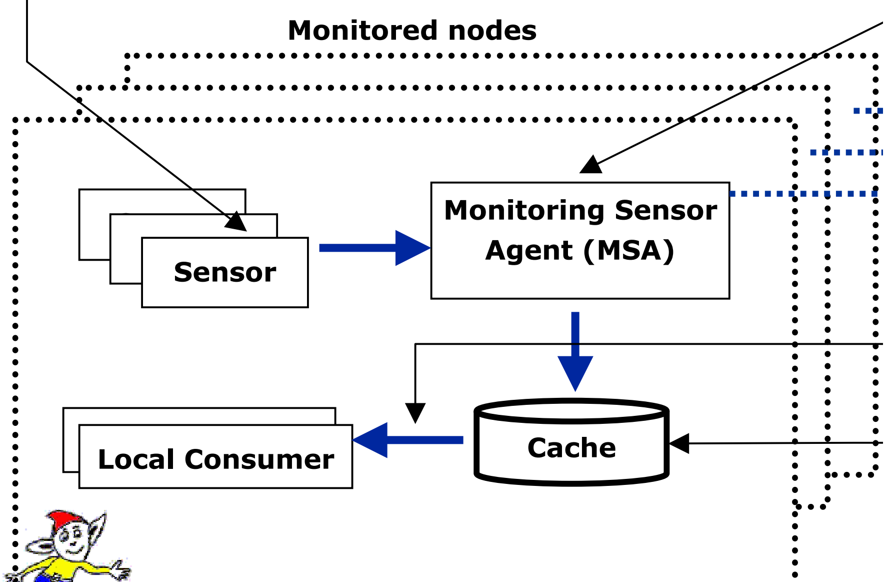

MSA in production for over 15 months, together with sensors for performance and exception metrics for basic OS and specific batch server items.

Focus now is on integrating existing monitoring for other systems, especially disk and tape servers, into the Lemon framework.

repairs

LEAF

HMS (Hardware Management System)

- **tracks systems through steps necessary for, e.g., installations & moves.**
  a Remedy workflow interfacing to ITCM, PRMS and CS group as necessary.
  - **used to manage the migration of systems to the vault.**
  now driving installation of 250+ systems.

SMS (State Management System)

- **"Give me 200 nodes, any 200. Make them like this. By then."**
  For example creation of an initial RH10 cluster
  (re)allocation of CPU nodes between lxbatch & lxshare or of disk servers.
- **Tightly coupled to Lemon to understand current state and CDB (Configuration Data Base) —which SMS must update.**

Fault Tolerance

- **We have started testing the Local Recovery Framework developed by Heidelberg within EDG/WP4.**
- **Simple recovery action code (e.g. to clean up filesystems *safely*) is available.**

Fabric Infrastructure Summary

- **The Building Fabric will be ready for start of production farm installation in January 2006.**
  - But there are concerns about a potential open ended increase of power demand.

- **CPU & disk server purchase complex**
  - Major risk is poor quality hardware and/or lack of adequate support from vendors.

- **Computing Fabric automation is well advanced.**
  - Installation and configuration tools are in place.
  - The essentials of the monitoring system, sensors and the central repository, are also in place. Displays will come. More important is to encourage users to query our repository and not each individual node.
  - LEAF is starting to show real benefits in terms of reduced human intervention for hardware moves.

Services

The focus of the computing fabric are the services and they are integral part of the IT managerial infrastructure

- Management of the farms
- Batch scheduling system
- Networking
- Linux
- Storage management

but the service is of course currently not only for the LHC Experiments
IT supports about 30 Experiments, engineers, etc.
Resource usage dominated by punctual LHC physics data challenges and running experiments (NA48, COMPASS,….)

Couplings

Physical and logical coupling

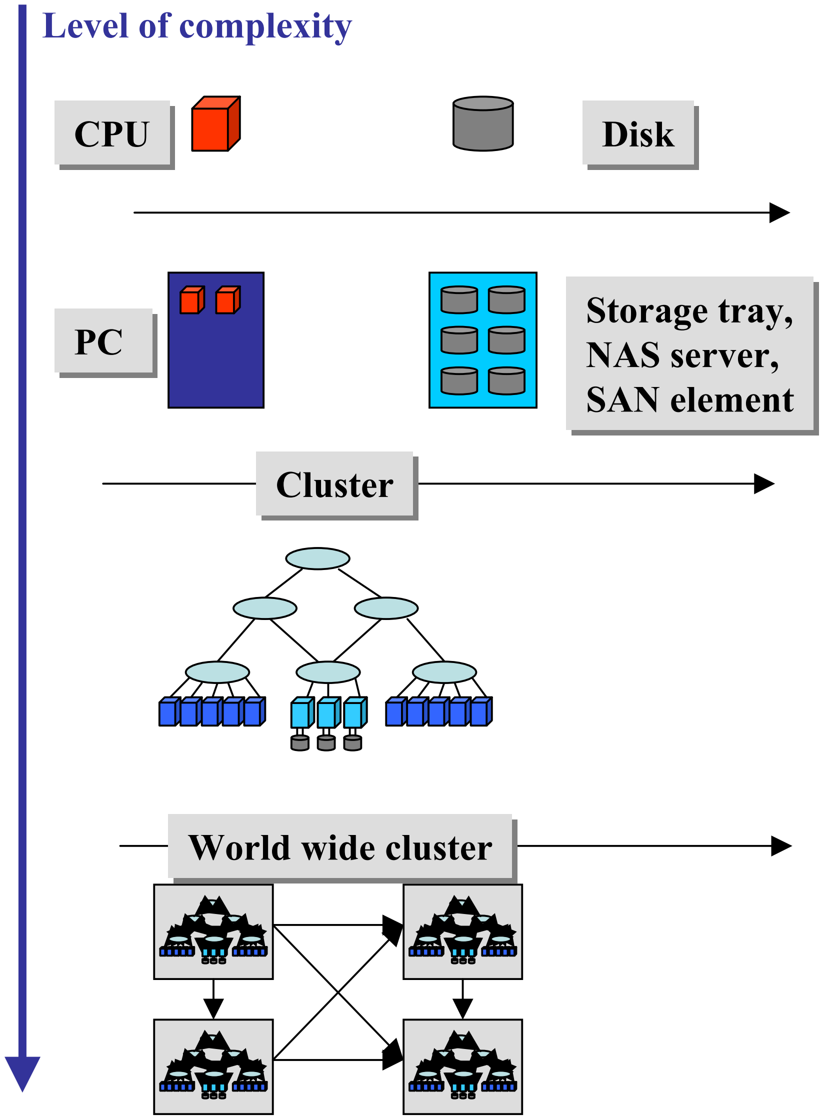

Hardware	Software
Motherboard, backplane, Bus, integrating devices (memory,Power supply, controller,..)	Operating system (Linux), driver, applications
Network (Ethernet, fibre channel, Myrinet, ….) Hubs, switches, routers	Batch system (LSF), Mass Storage (CASTOR) filesystems (AFS), Control software,
Wide area network (WAN)	Grid middleware, monitoring, firewalls

Grid-Fabric Interfaces

(Services)

Batch Scheduler

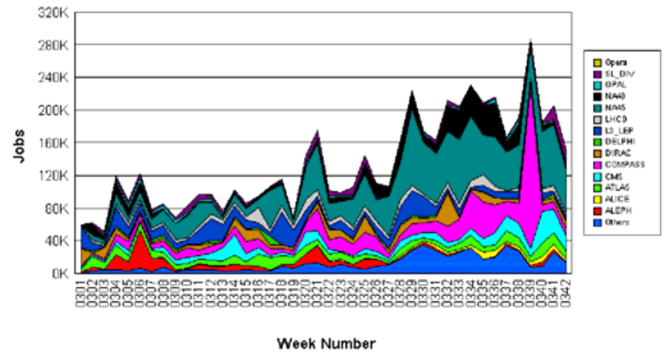

- ❑ Using LSF from Platform Ccomputing, commercial product
- ❑ deployed on 1000 nodes,
  10000 concurrent jobs in the queue on average,
  200000 jobs per week
- ❑ very good experience, fair share for optimal usage of resources
  current reliability and scalability issues are understood
  adaptation in discussion with users
  → average throughput versus peak load and real time response
- ❑ mid 2004 to start another review of available batch systems
  → choose in 2005 the batch scheduler for Phase II

Storage (I)

AFS (Andrew File System)

- A team of 2.2 FTE takes care of the shared distributed file system to provide access to the home directories (small files, programs, calibration, etc.) of about 14000 users.

- Very popular, growth rate for 2004 : 60 % (4.6 TB → 7.6 TB) expensive compared to bulk data storage (factor 5-8), automatic backup, high availability (99 %),   user perception different

- GRID job software environment distribution 'preferred' through shared file system solution per site
→ file system demands (performance, reliability, redundancy,etc.)

- Evaluation of different products have started
expect a recommendation by mid 2004, collaboration with other sites (e.g. CASPUR)

Storage (II)

CASTOR

- CERN development of a Hierarchical Storage Management system (HSM) for LHC

- Two teams are working in this area : Developer (3.8 FTE) and Support (3 FTE)
  Support to other institutes currently under negotiation (LCG, HEPCCC)

- Usage : 1.7 PB of data with 13 million files,
  250 TB disk layer and 10 PB tape storage
  Central Data Recording and data processing
  NA48 0.5 PB   COMPASS 0.6 PB   LHC Exp. 0.4 PB

- Current CASTOR implementation needs improvements → New CASTOR stager
  - A pluggable framework for intelligent and policy controlled file access scheduling
  - Evolvable storage resource sharing facility framework rather than a total solution

  detailed workplan and architecture available, presented to the user community in summer

- Carefully watching the tape technology developments (not really commodity)
  in depth knowledge and understanding is key

Linux

- 3.5 FTE team for Farms and Desktop
- Certification of new releases, bugfixes, security fixes,
  kernel expertise → improve performance and stability
- Certification group with all stakeholders : experiments, IT, accelerator, etc.
- Current distribution based on RedHat Linux
  major problem now : change in company strategy
  drop the free distributions and concentrate on the business with
  licenses and support for 'enterprise' distributions
- We are together with HEP community negotiating with RedHat
  Several alternative solutions were investigated : all need more money and/or
  more manpower
- Strategy is still to continue with Linux (2008 →)

Network

- Network infrastructure based on ethernet technology
- Need for 2008 a completely new (performance) backbone in the centre based on 10 Gbit technology. Today very few vendors offer this multiport, non-blocking, 10 Gbit router.
  We have an Enterasys product already under test (openlab, prototype)
- Timescale is tight :

Q1 2004 market survey
Q2 2004 install 2-3 different boxes, start thorough testing
→
prepare new purchasing procedures, finance committee
vendor selection, large order
→
Q3 2005 installlation of 25% of new backbone
Q3 2006 upgrade to 50%
Q3 2007 100% new backbone

Dataflow Examples

scenario for 2008

Implementation details depend on the computing models of the experiments
→ more input from the 2004 Data Challenges

→ modularity and flexibility in the architecture are important

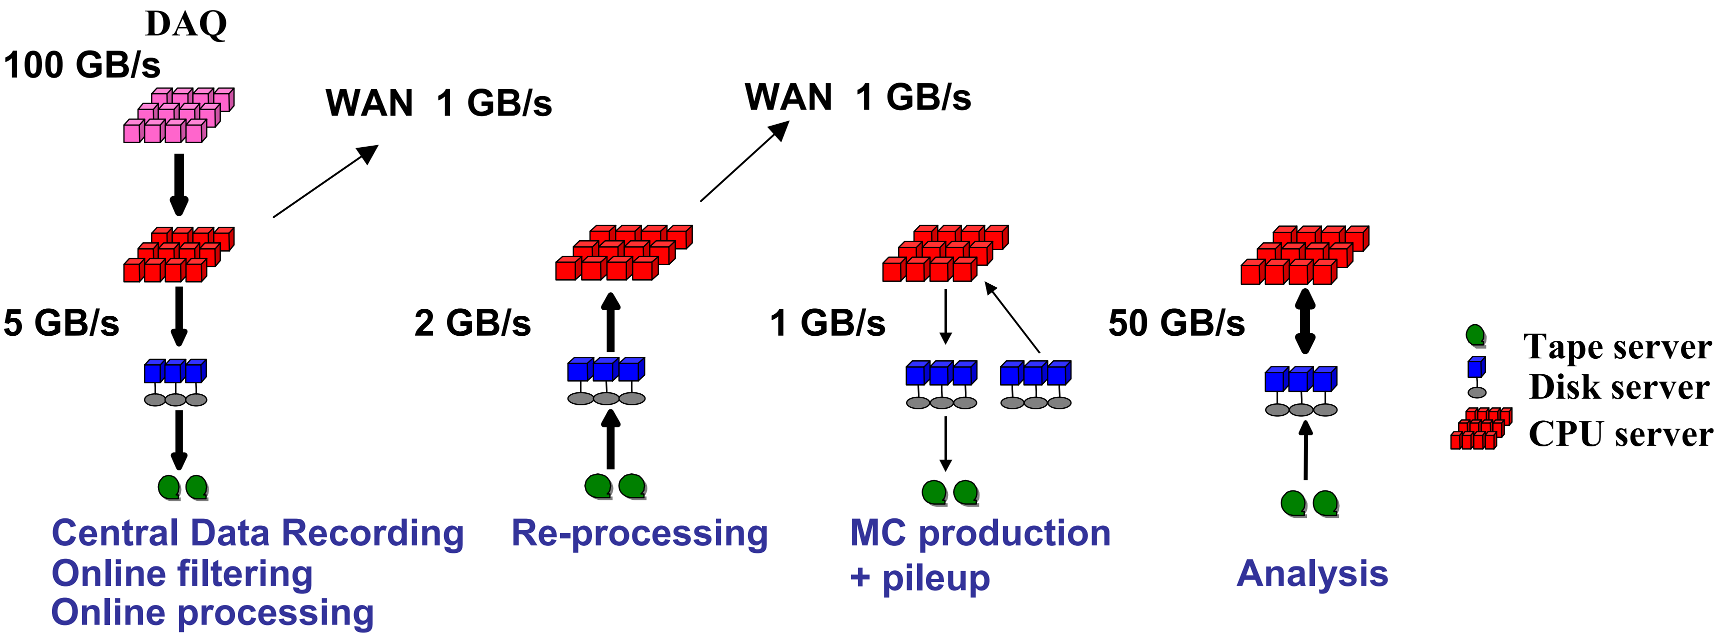

Today's schematic network topology

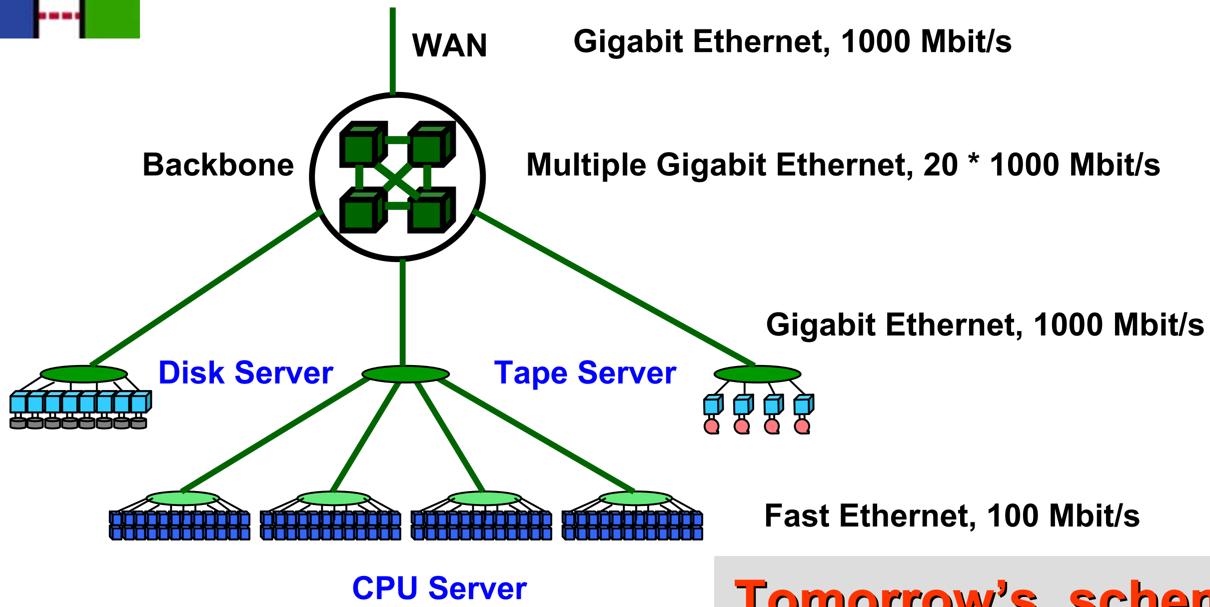

Tomorrow's schematic network topology

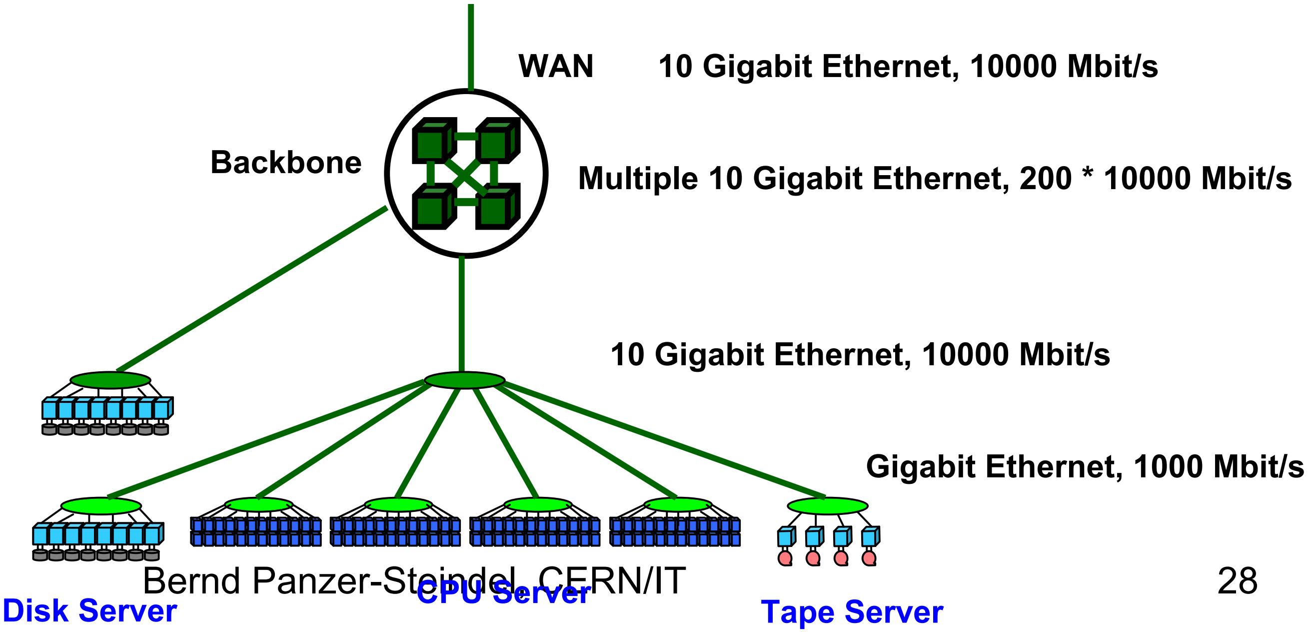

Wide Area Network

- Currently 4 lines   21 Mbit/s, 622 Mbits/s , 2.5 Gbits/s (GEANT),
  dedicated 10 Gbit/s line (starlight chicago, DATATAG),
  next year full 10 Gbit/s production line

- Needed for import and export of data, Data Challenges,
  todays data rate is 10 – 15 MB/s

- Tests of mass storage coupling starting (Fermilab and CERN)

- Next year more production like tests with the LHC experiments
  CMS-IT data streaming project inside the LCG framework
  tests on several layers : bookkeeping/production scheme, mass storage
  coupling, transfer protocols (gridftp, etc.), TCP/IP optimization

- 2008 :
  multiple 10 Gbit/s lines will be availble with the move to 40 Gbit/s connections
  CMS and LHCb will export the second copy of the raw data to the T1 center ,
  ALICE and ATLAS want to keep the second copy at CERN (still ongoing discussion)

Service Summary

- limited number of milestones
  focus on evolution of services not major changes → stability
- crucial developments in the network area
- mix of industrial and ‘home-grown’ solutions → TCO judgement
- moderate difficulties no problem so far
- separation of LHC versus non-LHC sometimes difficult

Grid – Fabric coupling

- Ideally clean interface and Grid middleware and services are one layer above the Fabric
  → reality is more complicated (intrusive)
- New research concept meets conservative production system
  → inertia and friction
- Authentication, security, storage access, repository access, job scheduler usage, etc. different implementations and concepts
  → adaptation and compromises necessary
- Regular and good collaboration between the teams established, still quite some work to be done
- Some milestones are late by several months (Lxbatch Grid integration)
  → late LCG-1 release and problem resolving in the GRID-Fabric API's more difficult than expected

Resource Planning

- Dynamic sharing of resources between the LCG prototype installation and the Lxbatch production system.
- Primarily Physics data challenges on Lxbatch and computing data challenges on the prototype
- IT Budget for the growth of the production system will be 1.7 million in 2004 and the same in 2005.
- Resource discussion and planning in the PEB

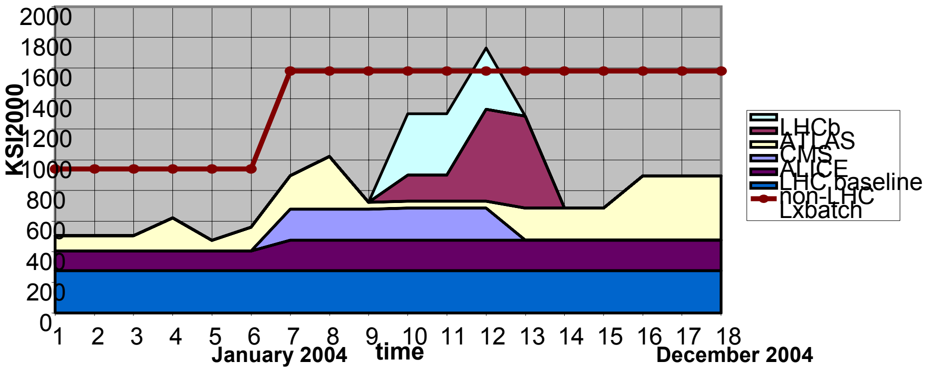

General Fabric Layout

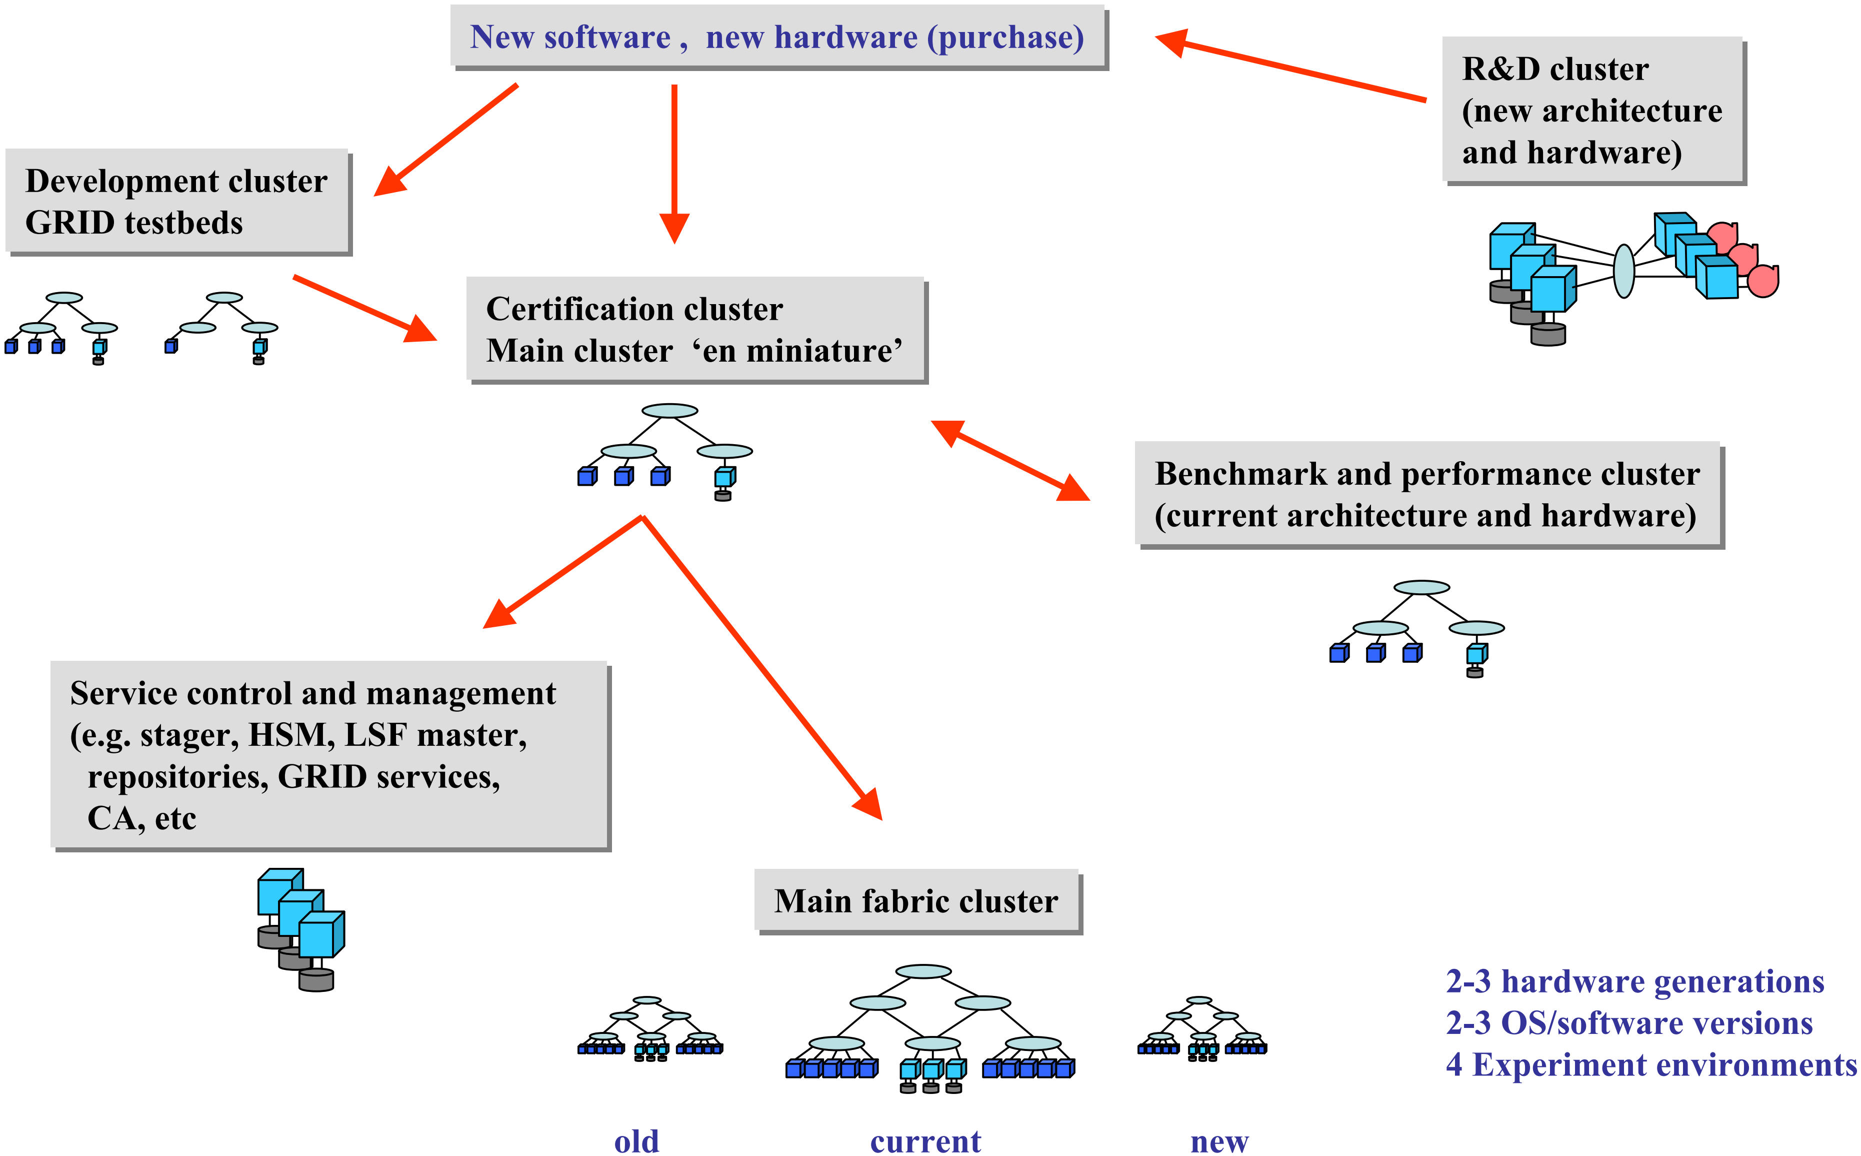

Computer center today

- **Main fabric cluster (Lxbatch/Lxplus resources)**
  - → physics production for all experiments
    Requests are made in units of Si2000
  - → 1200 CPU server, 250 disk server, ~ 1100000 Si2000, ~ 200 TB
  - → 50 tape drives (30MB/s, 200 GB cart.)
    10 silos with 6000 slots each == 12 PB capacity

- **Benchmark,performance and testbed clusters (LCG prototype resources)**
  - → computing data challenges, technology challenges,
    online tests, EDG testbeds, preparations for the LCG-1
    production system, complexity tests
  - → 600 CPU server, 60 disk server, ~500000 Si2000, ~ 60 TB
    current distribution :
    220 CPU nodes for LCG testbeds and EDG
    30 nodes for application tests (Oracle, POOL, etc.)
    200 nodes for the high performance prototype(network,ALICE DC, openlab)
    150 nodes in Lxbatch for physics DCs

Data Challenges

- Physics Data Challenges (MC event production, production schemes, middleware)

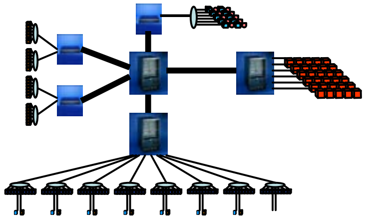

- ALICE IT Mass Storage Data Challenges
  2003 → 300 MB/s, 2004 → 450 MB/s, 2005 → 700 MB/s
  preparations for the ALICE CDR in 2008 → 1.2 GB/s

- Online DCs (ALICE event building, ATLAS DAQ)

- IT scalability and performance DCs (network, filesystems, tape storage → 1 GB/s)

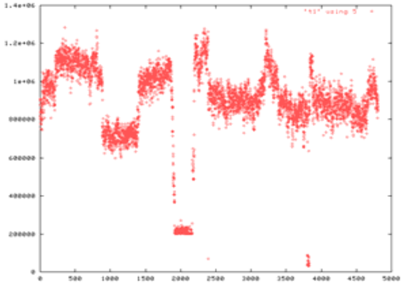

- Wide Area Network (WAN) coupling of mass storage systems, data export and import started

- Architecture testing and verification, computing models, scalability
  → needs large dedicated resources, avoid interference with production system

- Very successful Data Challenges in 2002 and 2003

LCG Materials Expenditure at CERN

	2001	2002	2003	2004	2005	TOTAL
PROTOTYPE						
Processors -		422	220	300	500	
PC R&D + admin tests		100		50		
Disk storage -		919	120	930	780	
Mass storage -		1323	100	200	200	
Systems admin contract -		380	200		400	
Physics WAN/LAN -			225	500	500	
LCG Associate s		290	310	360	360	
TOTAL PROTOTYPE	**2454**	**3434**	**1175**	**2340**	**2740**	**12143**
PREPARING CC FOR PHASE 2						
Vault		1665				
Substation			1590	805	550	
Main computer room			850	1880	500	
Air conditioning				300	450	
TOTAL PREPARING CC for PHASE 2	**150**	**1665**	**2440**	**2985**	**1500**	**8740**
TOTAL PROTOTYPE OPERATION	**2604**	**5099**	**3615**	**5325**	**4240**	**20883**

Staffing

- 25.5 FTE from IT division are allocated in the different services to LHC activities. These are fractions of people, LHC experiments not yet the dominating users of services and resources

- 12 FTE from LCG and 3 FTE from the DataGrid (EDG) project are working in the area of service developments (e.g. security, automation) and evaluation (benchmarks, data challenges, etc.)
  This number (15) will decrease to 6 by mid 2004 (EDG ends in February, end of LCG contracts (UPAS, students, etc.)
  Fellows and Staff continue until 2005

Re-costing results

Resource	Old 2006-08	New 2006-08	New- Old 2006-08	Old 2009-10	New 2009-10	New - Old 2009-10
CPU+LAN	17.7	19.5	1.8	6.3	6.8	0.5
Disk	6.3	11.9	5.6	2.2	2.9	0.7
Tape	22.5 *	27.8	5.3	19.2	17.6	-1.6
WAN	11.4	6.0	- 4.4	6.8	4.0	-2.8
Sysadmin	7.9	3.5	- 4.4	6.6	3.0	-3.6
SUM	65.8	68.7	2.9	41.1	34.3	-6.8
Budget		60.0			34.0	

* A bug in the original paper is here corrected

All units in [million CHF]

Comparison

2008 prediction		2003 status
➢ Hierarchical Ethernet network	280 GB/s	2 GB/s
➢ ~ 8000 mirrored disks	(4 PB)	2000 mirrored disk 0.25 PB
➢ ~ 3000 dual CPU nodes	(20 MSI2000)	2000 nodes 1.2 MSI2000
➢ ~ 170 tape drives	(4 GB/s)	50 drives 0.8 GB/s
➢ ~ 25 PB tape storage		10 PB

→The CMS HLT will consist of about 1000 nodes with 10 million SI2000 !!

External Fabric relations

Collaboration with Industry
openlab
HP, INTEL, IBM, Enterasys, Oracle
-- 10 Gbit networking
-- new CPU technology
-- possibly , new storage technology

LCG
-- Hardware resources
-- Manpower resources

Collaboration with India
-- filesystems
-- Quality of Service

CERN IT
- Main Fabric provider

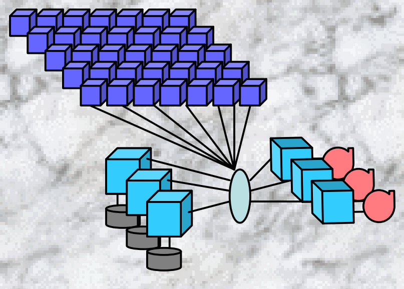

GDB working groups
-- Site coordination
-- Common fabric issues

Collaboration with CASPUR
--harware and software benchmarks
and tests, storage and network

External network
-- DataTag, Grande
-- Data Streaming project
with Fermilab

LINUX
-- Via HEPIX RedHat license
coordination
inside HEP (SLAC, Fermilab)
certification and security

CASTOR
-- SRM definition and implementation
(Berkeley, Fermi, etc.)
-- mass storage coupling tests (Fermi)
-- scheduler integration (Maui, LSF)
-- support issues (LCG, HEPCCC)

EDG, WP4
-- Installation
-- Configuration
-- Monitoring
-- Fault tolerance

GRID Technology and deployment
-- Common fabric infrastructure
-- Fabric ←→ GRID interdependencies

Online-Offline boundaries
-- workshop and discussion with
Experiments
-- Data Challenges

Timeline

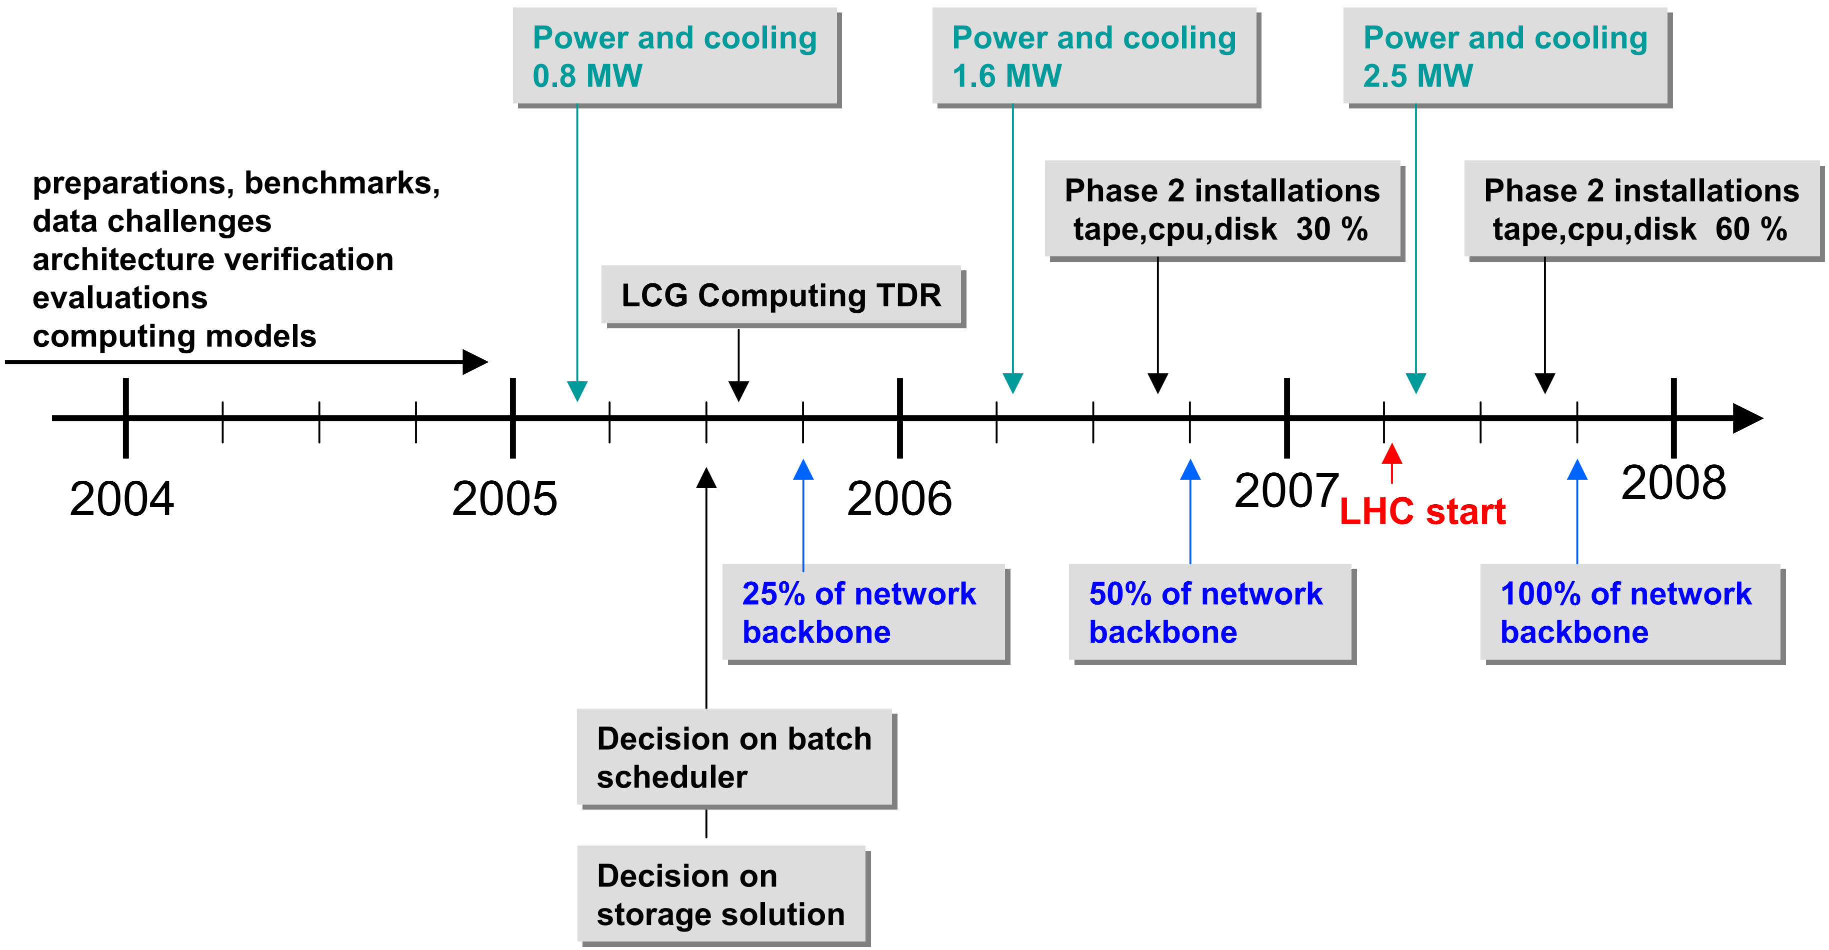
Power and cooling 0.8 MW
Power and cooling 1.6 MW
Power and cooling 2.5 MW
preparations, benchmarks, data challenges architecture verification evaluations computing models
Phase 2 installations tape,cpu,disk 30 %
Phase 2 installations tape,cpu,disk 60 %
LCG Computing TDR
2004
2005
2006
2007
LHC start
2008
25% of network backbone
50% of network backbone
100% of network backbone
Decision on batch scheduler
Decision on storage solution

Summary

- Evolution of services with focus on stability and parallel evaluation of new technologies is the strategy
- The computing models need to be defined in more detail.
- Positive collaboration with outside Institutes and Industry.
- Timescale is tight, but not problematic.
- Successful Data Challenges and most milestones on time.
- The pure technology is difficult (network backbone, storage), but the real worry is the market development.