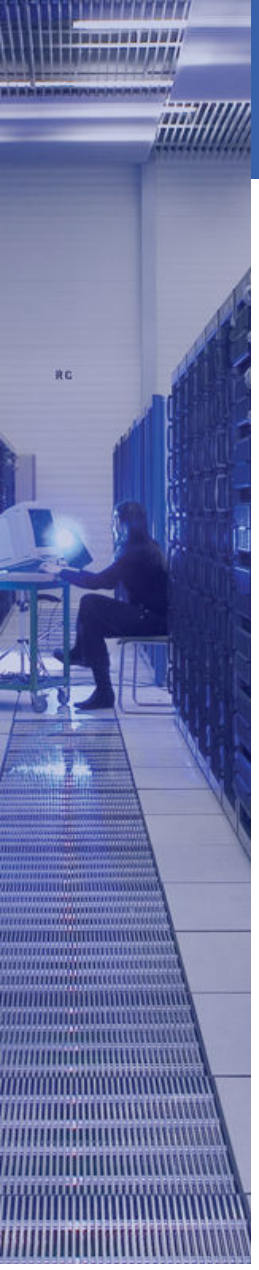

CERN IT Department

Introduction to CERN Computing Services

Bernd Panzer-Steindel, CERN/IT

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

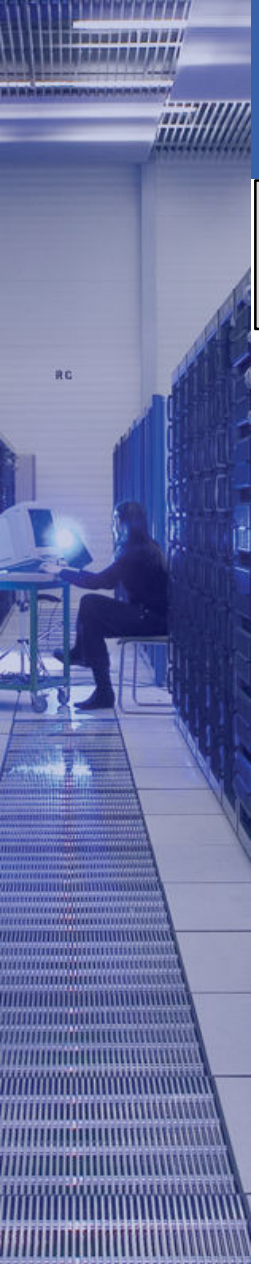

IT Department

CERN IT Department

220 people working in the IT Department

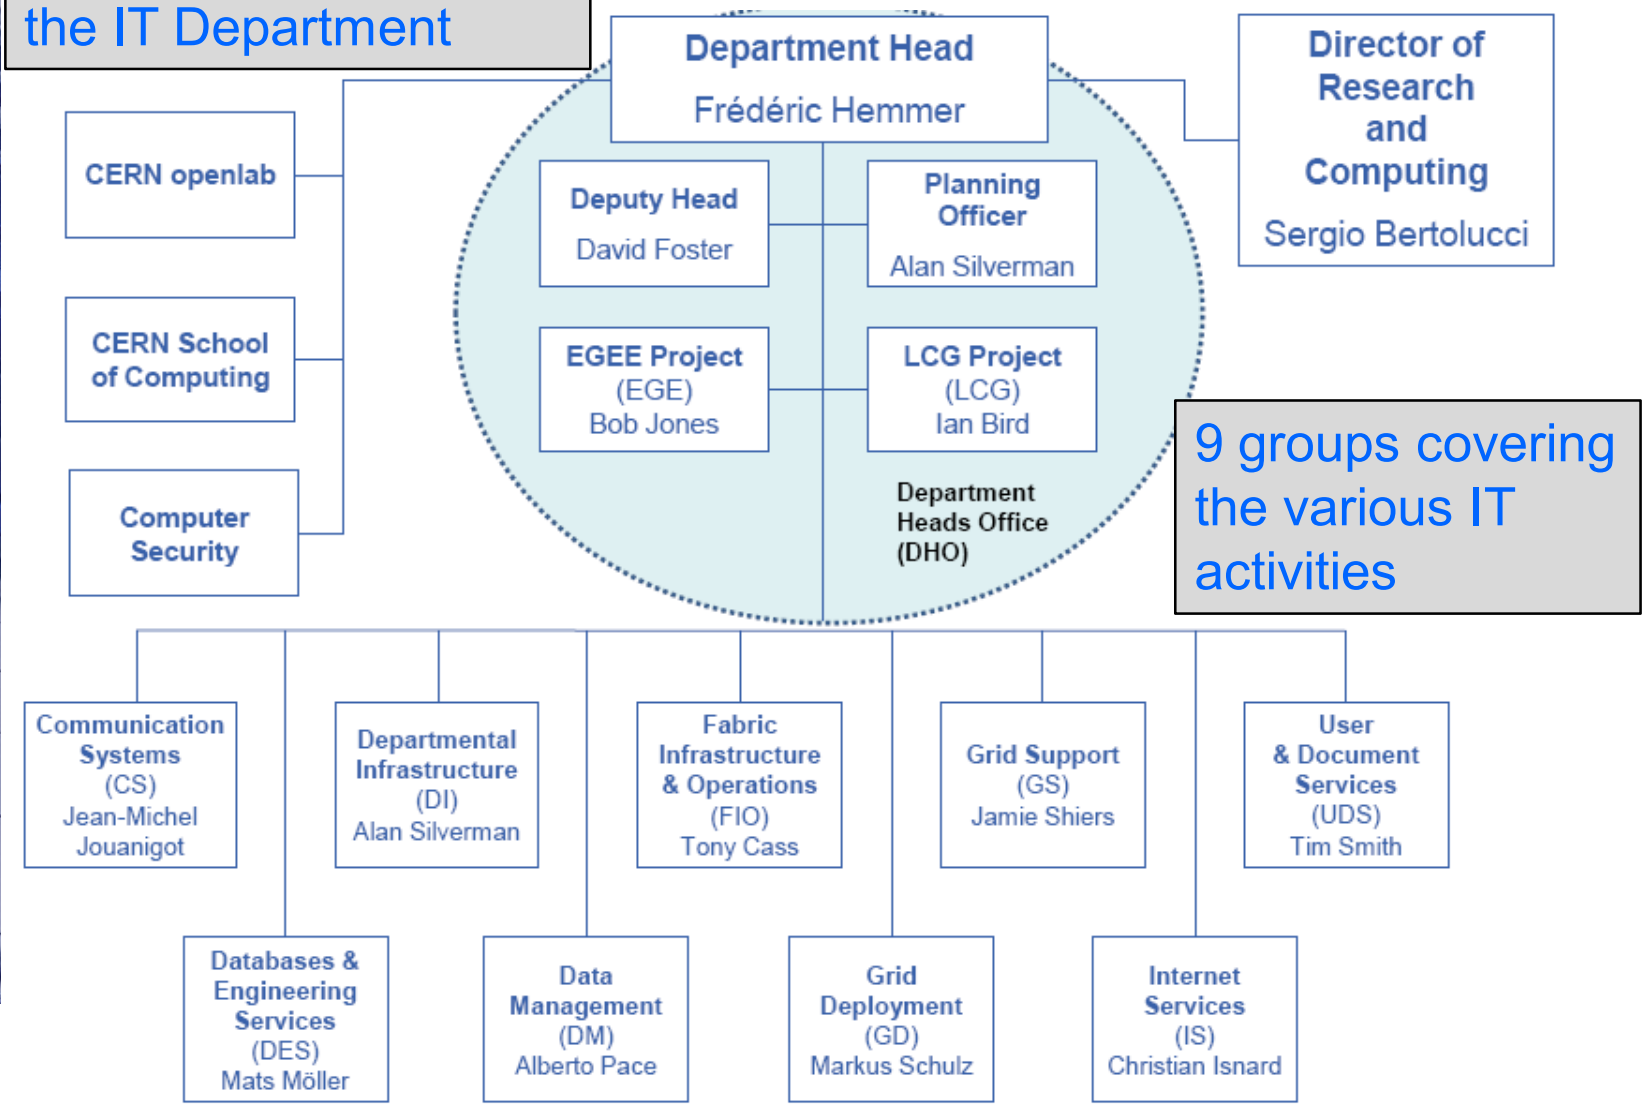

9 groups covering the various IT activities

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Location

CERN IT Department

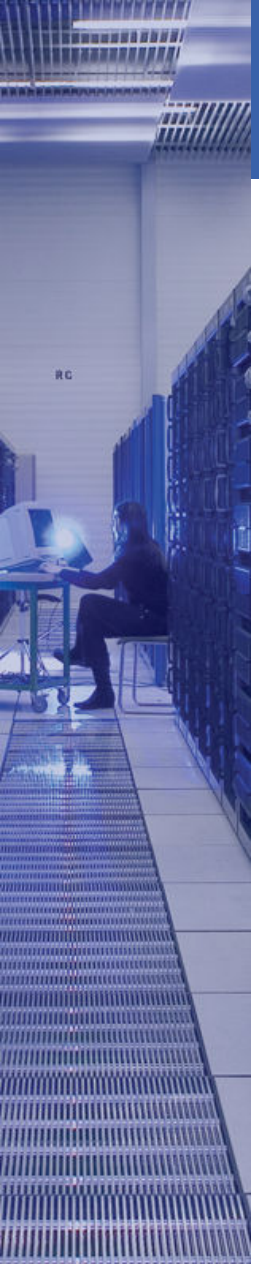

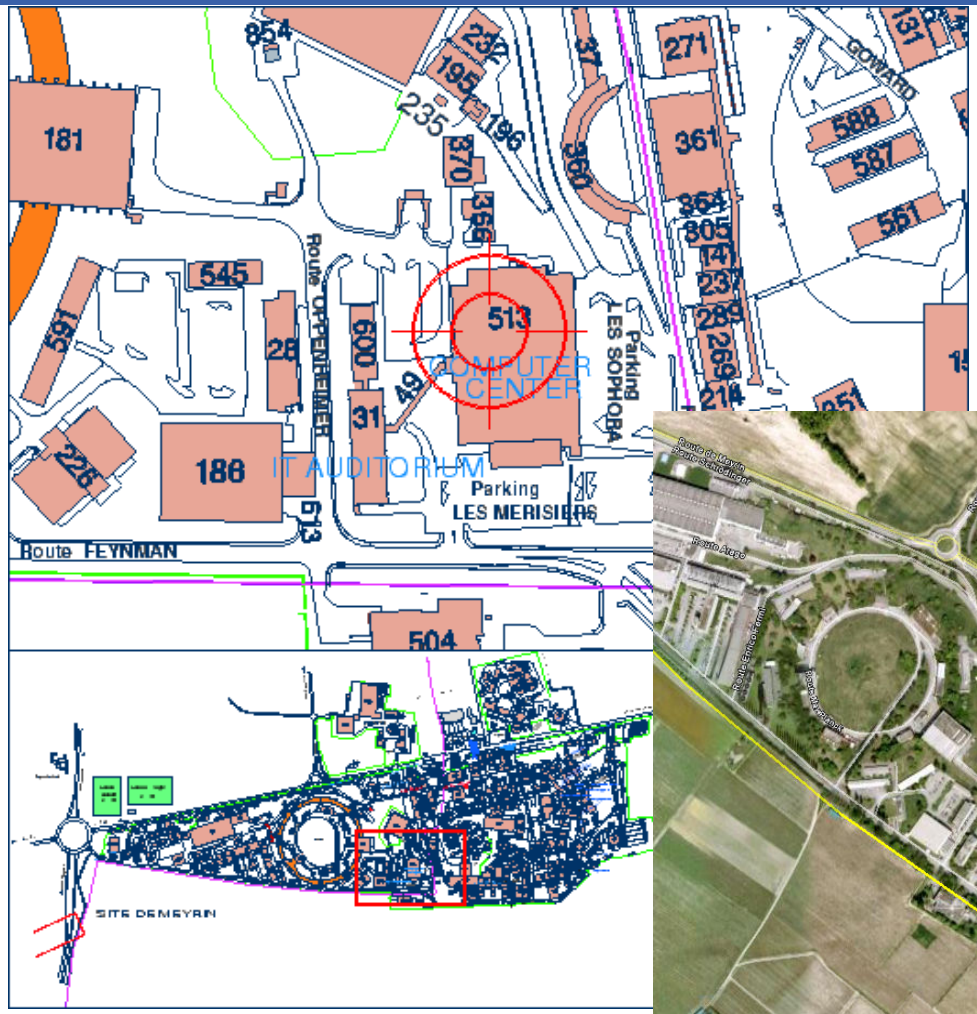

Building 513
Opposite of
Restaurant Nr. 2

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Building

CERN IT Department

Large building with 2700 m2 surface for computing equipment, capacity for 2.9 MW electricity and 2.9 MW air and water cooling

Chillers

Transformers

CERN IT Dep
CH-1211 Geneve 23
Switzerland
www.cern.ch/it

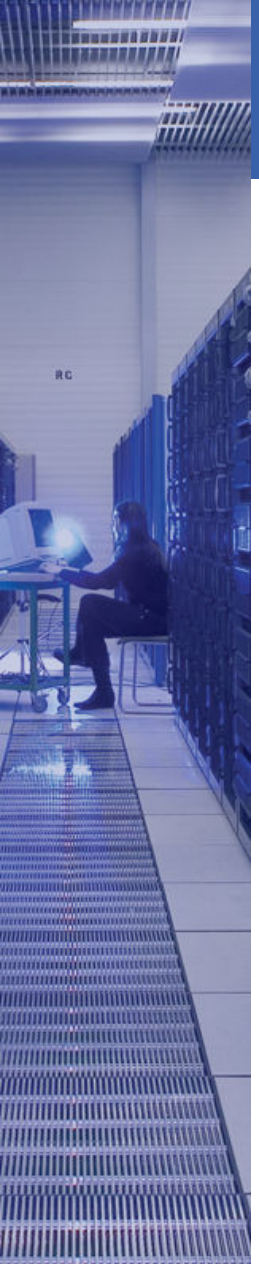

IT Mandate

CERN IT Department

To provide the information technology required for the fulfillment of the laboratory's mission in an efficient and effective manner through building world-class competencies in the technical analysis, design, procurement, implementation, operation and support of computing infrastructure and services.

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Tasks overview

Users need access to

Communication tools : mail, web, twiki, GSM, …

Productivity tools : office software, software development, compiler, visualization tools, engineering software, …

Computing capacity : CPU processing, data repositories, personal storage, software repositories, metadata repositories, …

Needs underlying infrastructure

- Network and telecom equipment,
- Processing, storage and database computing equipment,
- Management and monitoring software
- Maintenance and operations
- Authentication and security

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

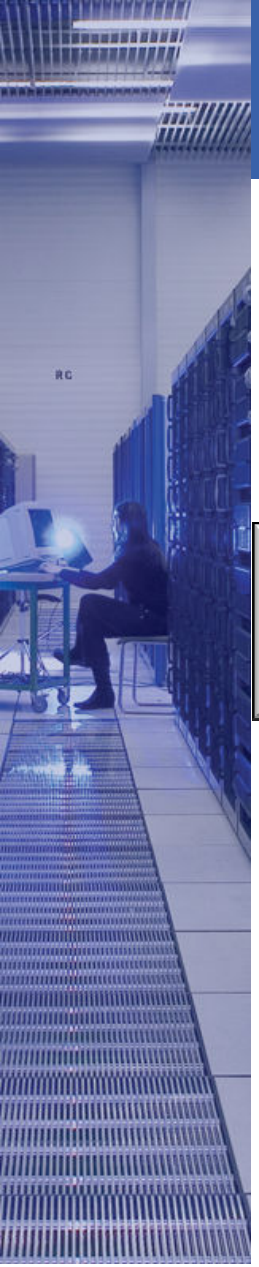

Infrastructure and Services

CERN IT Department

Software environment and productivity tools

User registration and authentication
→ 22000 registered users

Mail
→2M emails/day, 99% spam
18000 mail boxes

Web services
→ >8000 web sites

Tool accessibility
Windows, Office, CadCam

Home directories (DFS, AFS)
→60 TB, backup service,
→1 Billion files

PC management
Software and patch installations

Infrastructure needed :
> 300 PC server and 100 TB disk space

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Network overview

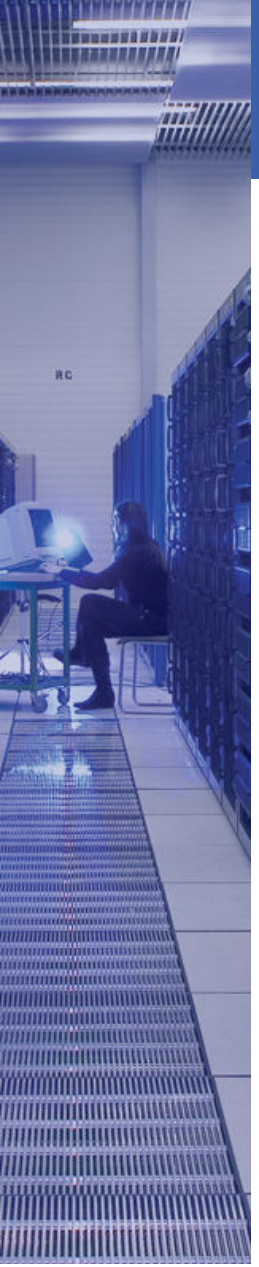

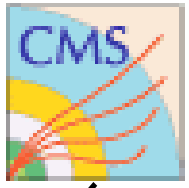

Experiments

All CERN buildings
12000 users

Central, high speed network backbone

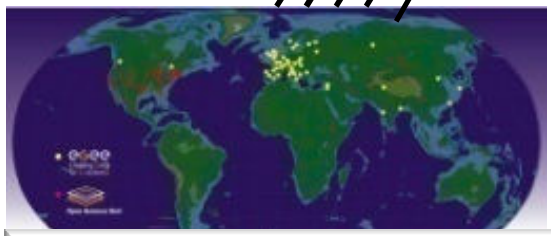

World Wide Grid centers

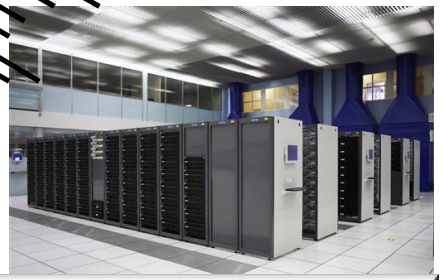

Computer center
Processing clusters

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

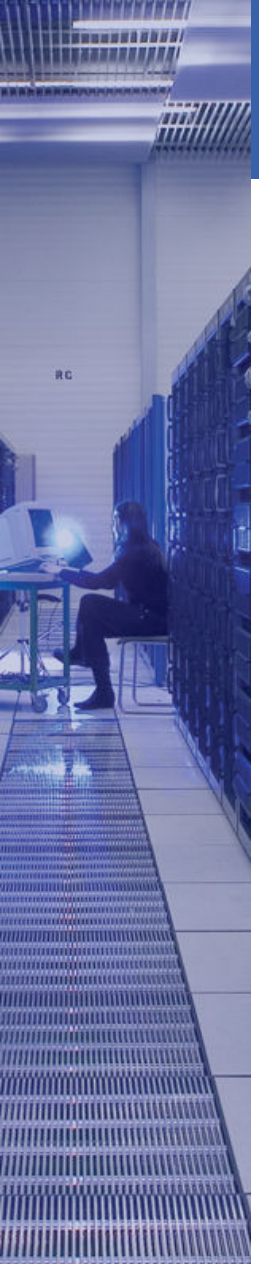

Network

- Hierarchical network topology based on Ethernet
- 150+ very high performance routers
- 3’700+ subnets
- 2200+ switches (increasing)
- 50’000 active user devices (exploding)
- 80’000 sockets – 5’000 km of UTP cable
- 5’000 km of fibers (CERN owned)
- 140 Gbps of WAN connectivity

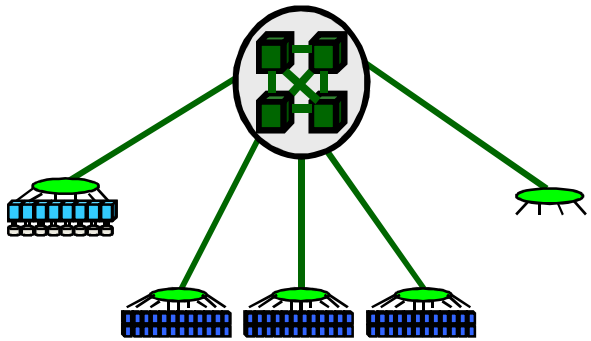

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Network topology

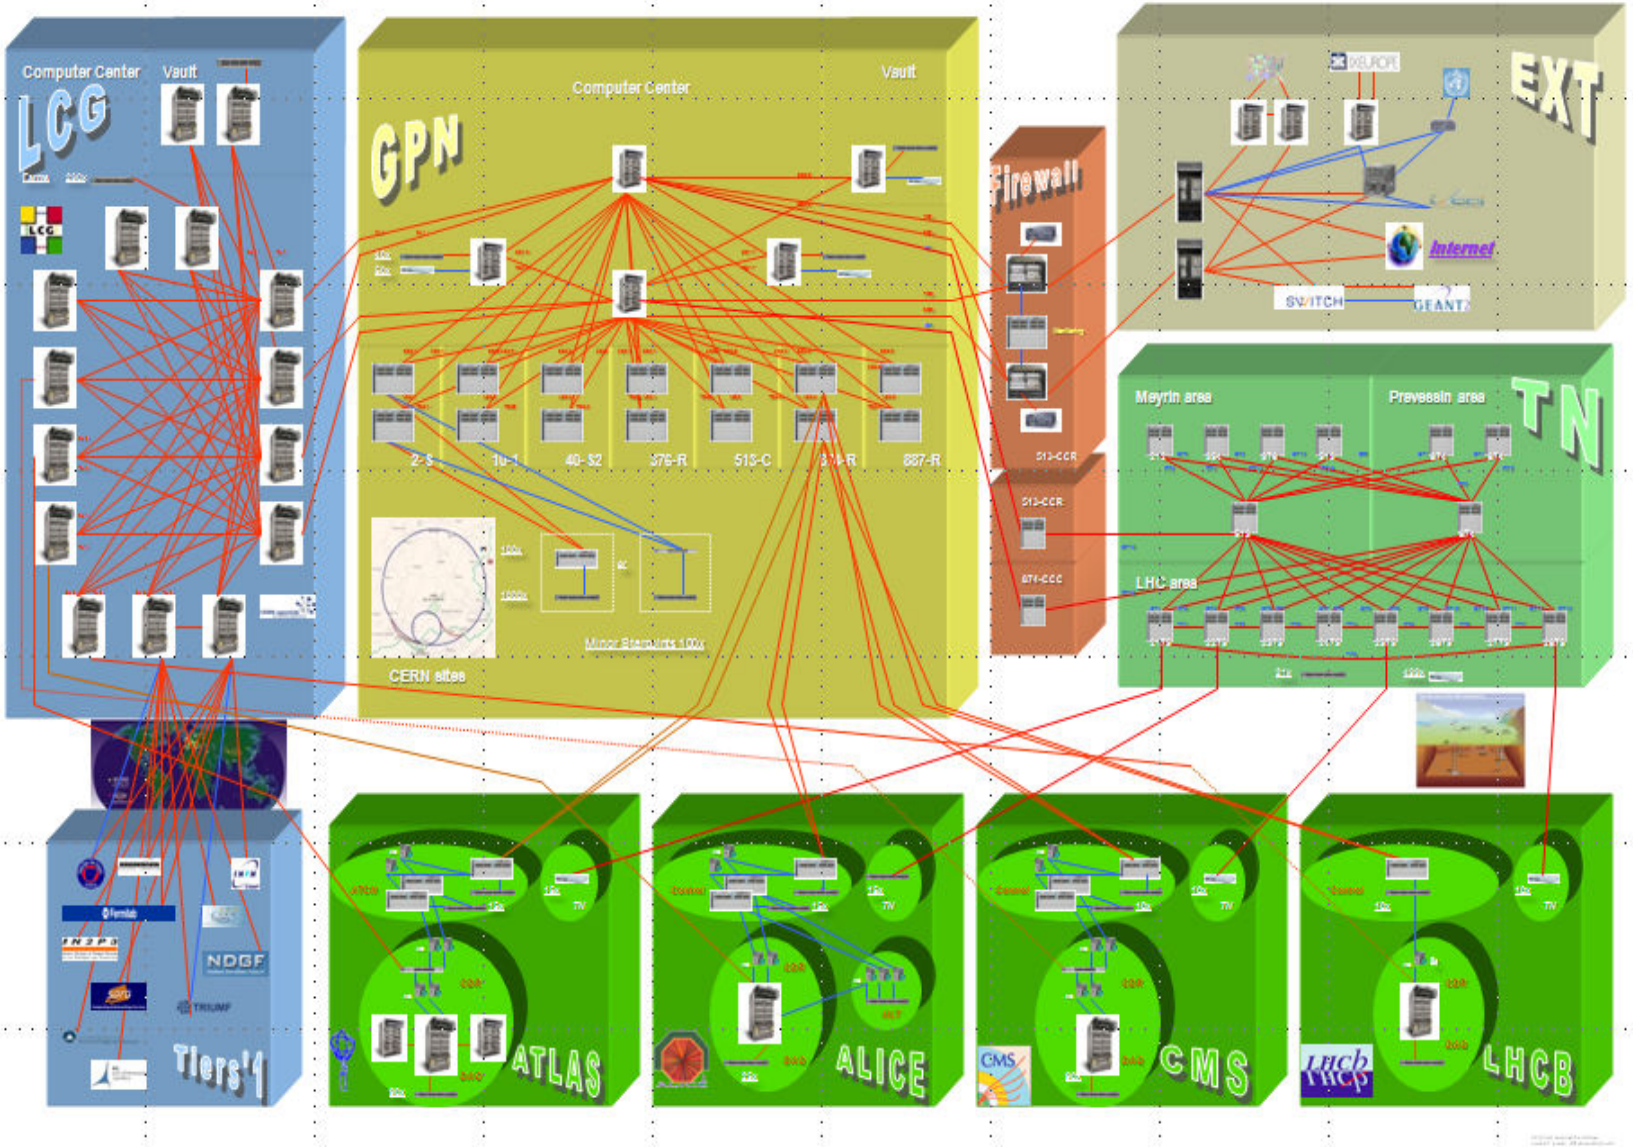

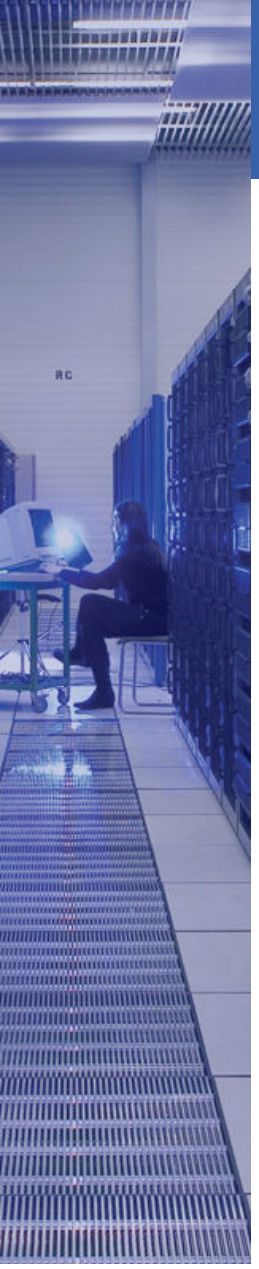

Data base services

More than 200 ORACLE data base instances on > 300 service nodes

-- bookkeeping of physics events for the experiments
-- meta data for the physics events (e.g. detector conditions)
-- management of data processing
-- highly compressed and filtered event data
-- …….

-- LHC magnet parameters

-- Human resource information
-- Financial bookkeeping
-- material bookkeeping and material flow control
-- LHC and detector construction details
-- ……..

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Monitoring

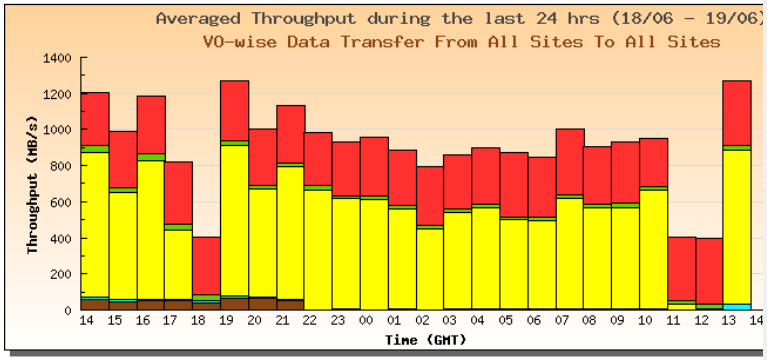

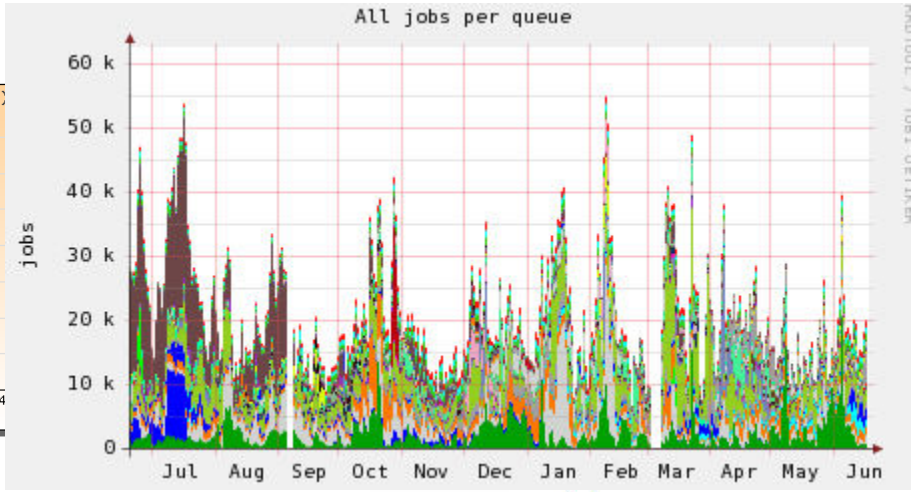

Large scale montitoring

Surveillance of all nodes in the Computer center

Hundreds of parameters in various time intervalls, from minutes to hours per node and service

Data base storage and interactive visalization

Batch cluster
type: host
cluster: 1
name: lxbatch

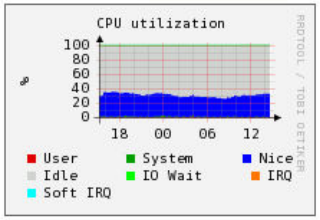

Fileservers
type: host
cluster: 1
name: castor2

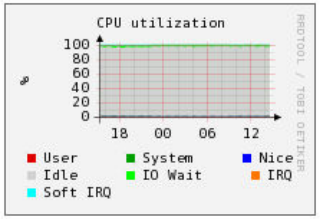

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Security I

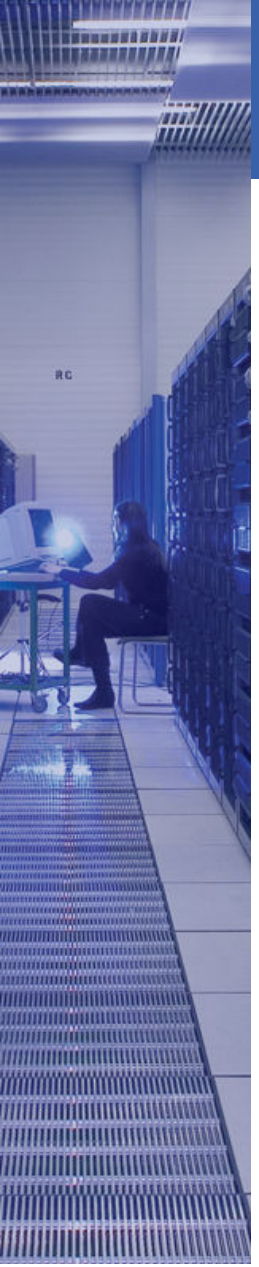

CERN IT Department

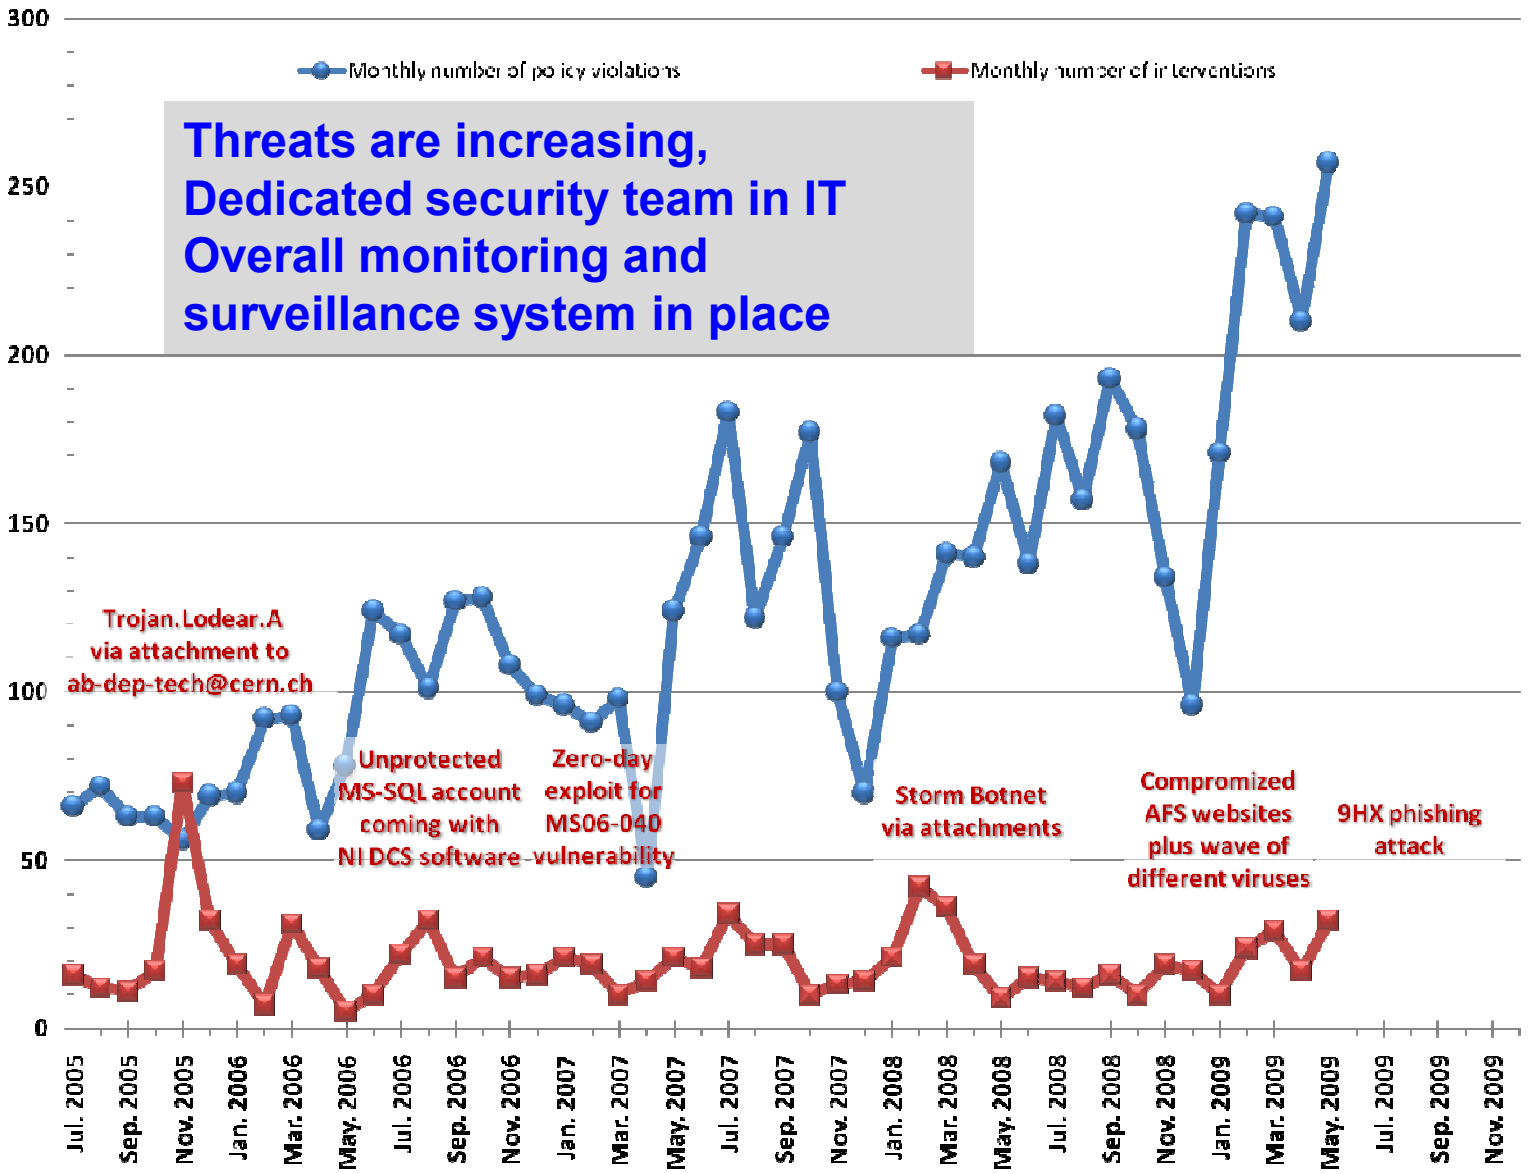

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

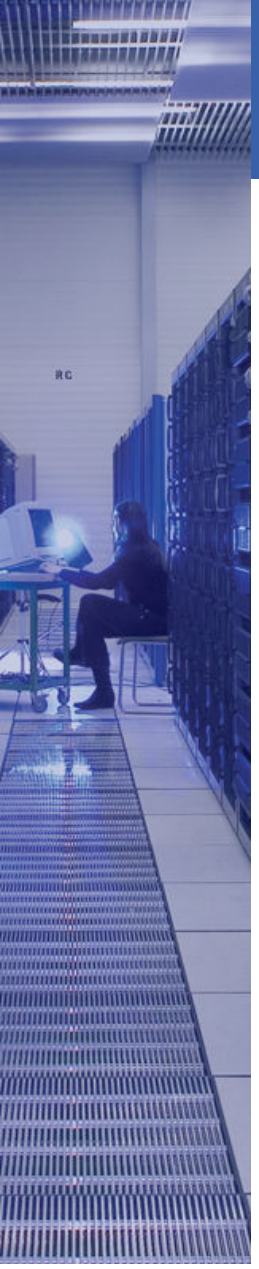

Security II

CERN IT Department

- Software which is not required for a user's professional duties introduces an unnecessary risk and should not be installed or used on computers connected to CERN's networks
- The use of P2P file sharing software such as KaZaA, eDonkey, BitTorrent, etc. is not permitted. IRC has been used in many break-ins and is also not permitted and Skype is tolerated provided it is configured according to http://cern.ch/security/skype
- Download of illegal or pirated data (software, music, video, etc) is *not* permitted
- Passwords must not be divulged or easily guessed and protect access to unattended equipment

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

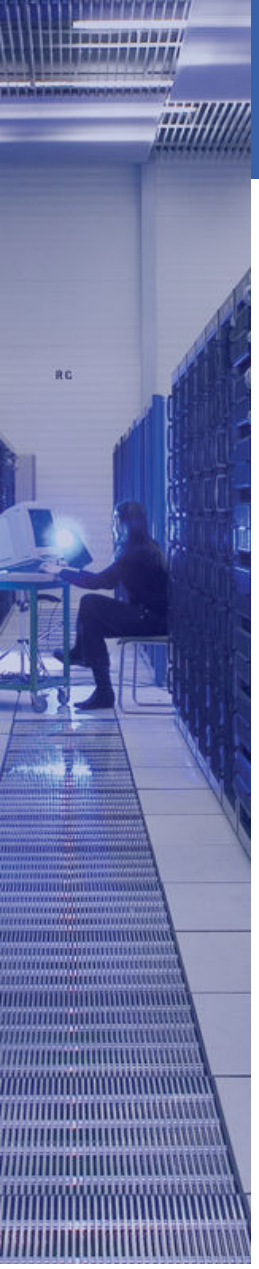

Security III

Please read the following documents
and follow their rules

http://cern.ch/ComputingRules

http://cern.ch/security

http://cern.ch/security/software-restrictions

In case of incidents or questions please contact :

Computer.Security@cern.ch

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Physics dataflow

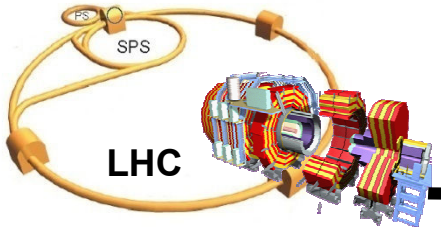

4 Experiments

1000 million Events/s

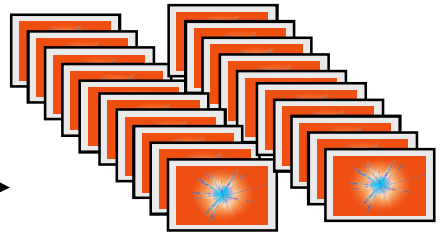

Filter and first selection

1 GB/s

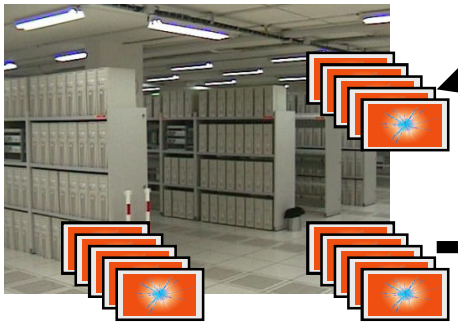

Create sub-samples

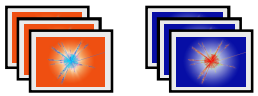

World-Wide Analysis

Physics
Explanation of nature

4 GB/s

3 GB/s

Store on disk and tape

Export copies

$$\sigma_{f\bar{f}} \approx \sigma^0_{f\bar{f}} \times \frac{s\Gamma_Z^2}{(s-m_Z^2)^2+s^2\Gamma}$$

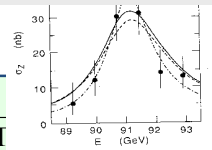

$$\sigma^0_{f\bar{f}} = \frac{12\pi}{m_Z^2}\frac{\Gamma_{ee}\Gamma_{f\bar{f}}}{\Gamma_Z^2} \quad \text{and} \quad \Gamma_{f\bar{f}} = \frac{G_F m_Z^3}{6\pi\sqrt{2}} \times (v_f^2 + a_f^2) \times N_{col}$$

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CERN Computing Fabric

CERN IT Department

large scale !

Purchasing

Physical installation

Logistic !

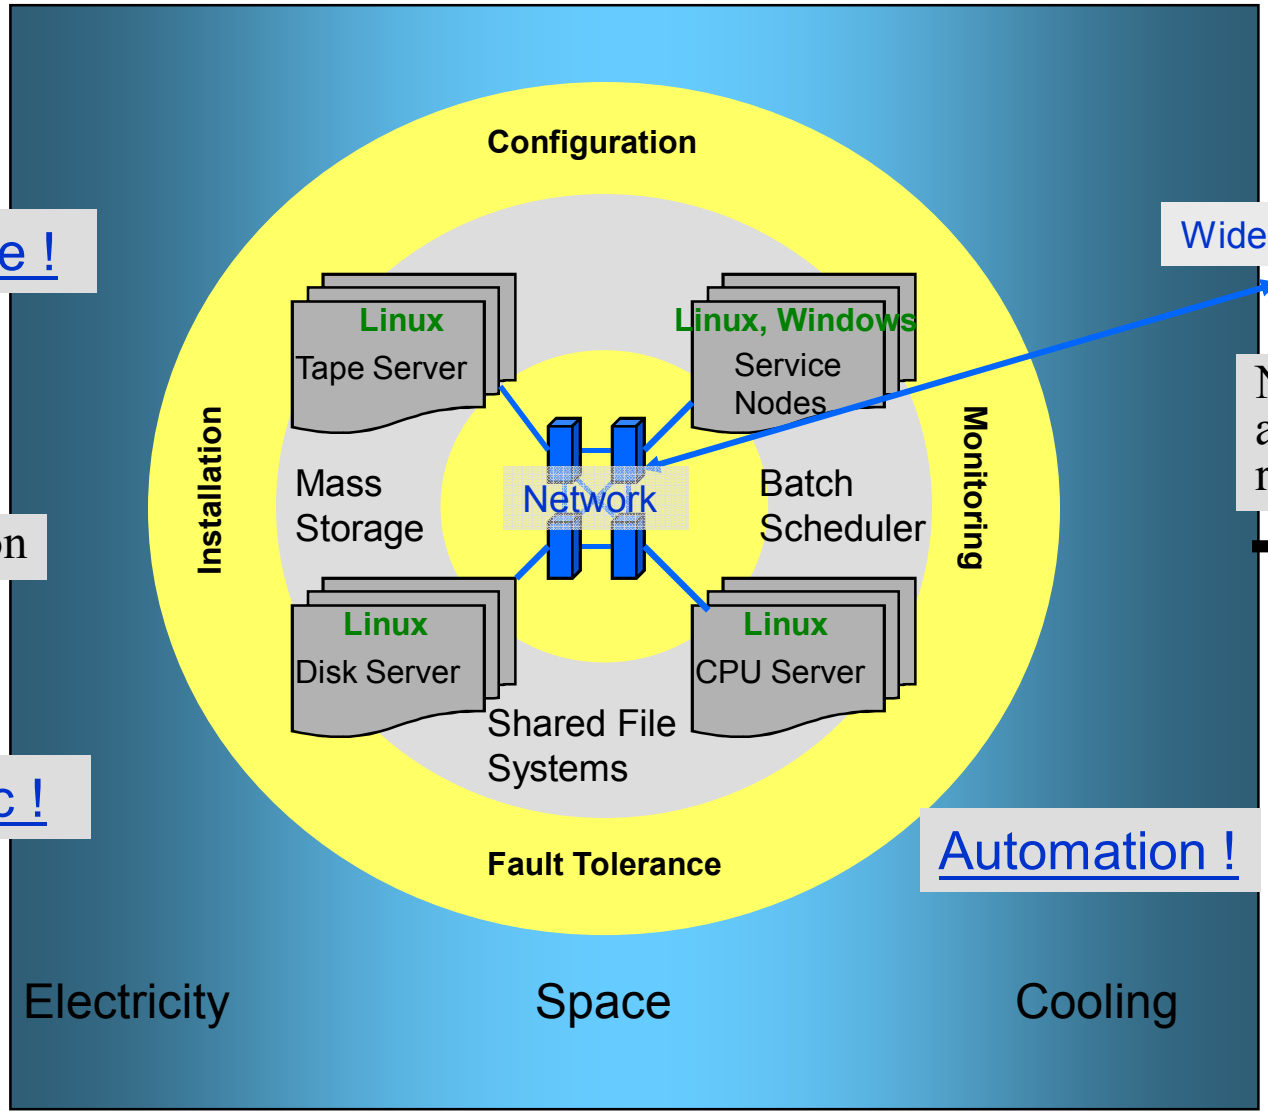

Wide Area Network

Node repair and replacement

Automation !

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

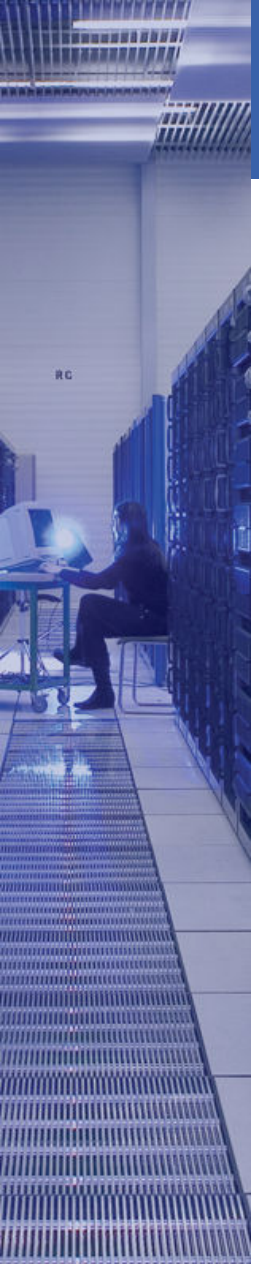

Logistics

- Today we have about 6000 PC‘s installed in the center
- Assume 3 years lifetime for the equipment
  → Key factors = power consumption, performance, reliability
- Experiment requests require investments of ~ 10 MCHF/year for new PC hardware

Infrastructur and operation setup needed for :

- ~1000 nodes installed per year and ~1000 nodes removed per year
- Installation in racks, cabling, automatic installation, Linux software environment
- Equipment replacement rate, e.g. 2 disk errors per day
  several nodes in repair per week
  50 node crashes per day

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Physical and logical connectivity

Complexity

Components

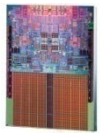



Hardware

Software

CPU,disk,memory, motherbord

PC, disk server


Operating system
Device drivers

Network, Interconnects

Cluster, Local fabric

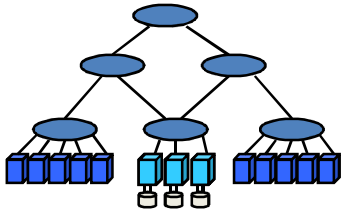


Resource Management software

Wide area network

World Wide Cluster

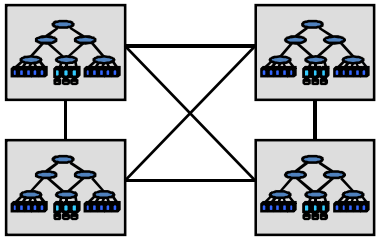
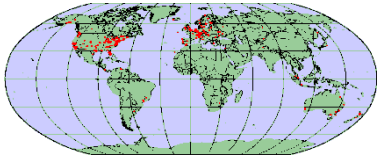

Grid and Cloud Management software

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Computing building blocks

commodity market components
not cheap but cost effective !
simple components, but many of them

CPU server or Service node

dual CPU, quad core,
16 GB memory

Tape server

=

CPU server
+ fibre channel connection
+ tape drive

TCO Total Cost of Ownership

Disk server

=

CPU server
+RAID controler
+24 SATA disks

market trends more important than technology trends

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

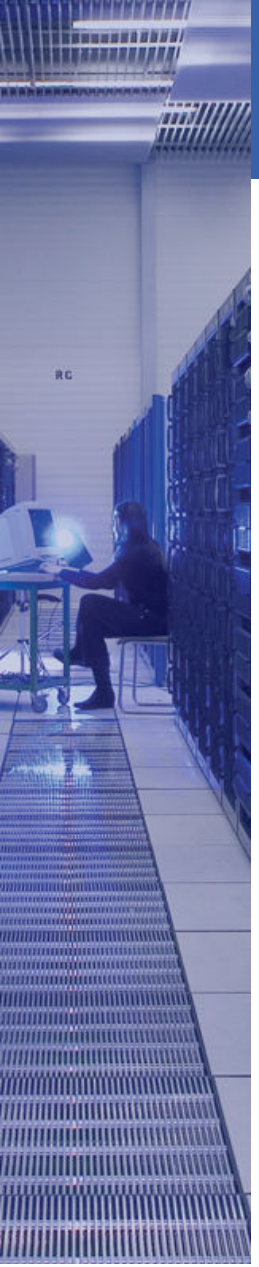

Software 'Glue'

- **Management of the basic hardware and software : installation, configuration and monitoring system (quattor, lemon, elfms)**
  Which version of Linux ? How to upgrade the software ?
  What is going on in the farm ? Load ? Failures ?

- **Management of the processor computing resources : Batch system (LSF from Platform Computing)**
  Where are free processors ? How to set priorities between different users ? sharing of the resources ? How are the results coming back ?

- **Management of the storage (disk and tape) : CASTOR (CERN developed Hierarchical Storage Management system)**
  Where are the files ? How can one access them ?
  How much space is available ? what is on disk, what is on tape ?

TCO Total Cost of Ownership buy or develop ?!

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Data and control flow

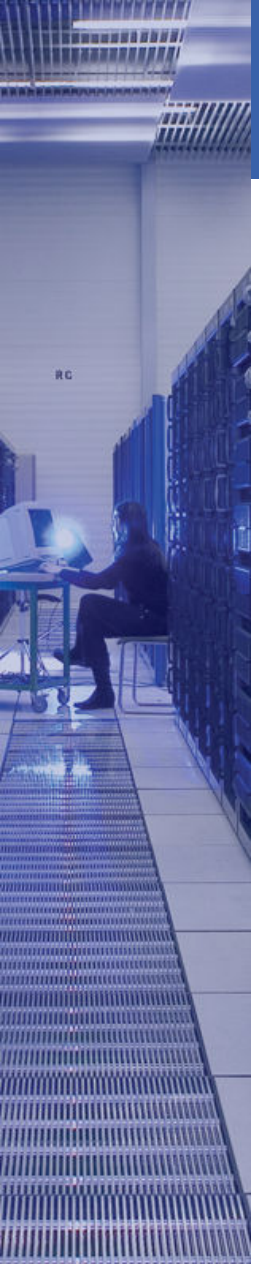

Users

Repositories
- code
- metadata
- .....

Interactive
Lxplus

Data processing
Lxbatch

Bookkeeping
Data base

Disk storage

Tape storage

CASTOR

Hierarchical Mass Storage Management System (HSM)

Interactive Service: Lxplus

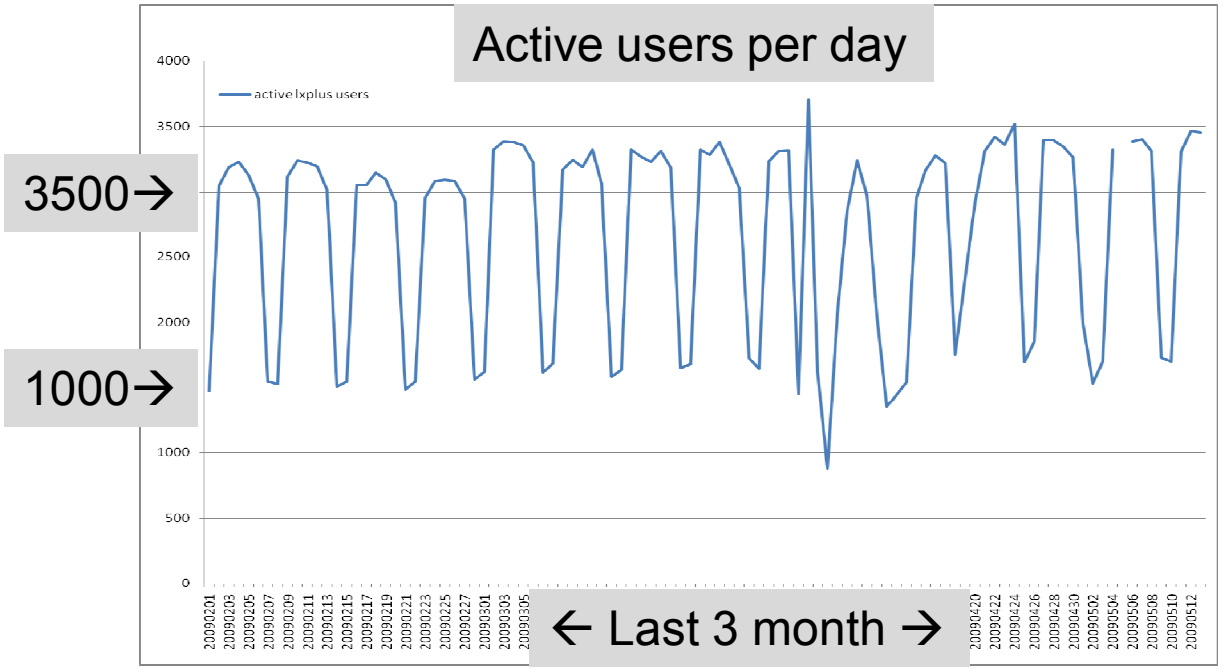

- Interactive computing facility
- 70 CPU nodes running Linux (RedHat)
- Access via ssh, xterminal from Desktop/Notebooks ,Windows,Linux,Mac
- Used for compilation of programs, short program execution tests, some Interactive analysis of data, submission of longer tasks (jobs) into the Lxbatch facility, internet access, program development, etc.

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

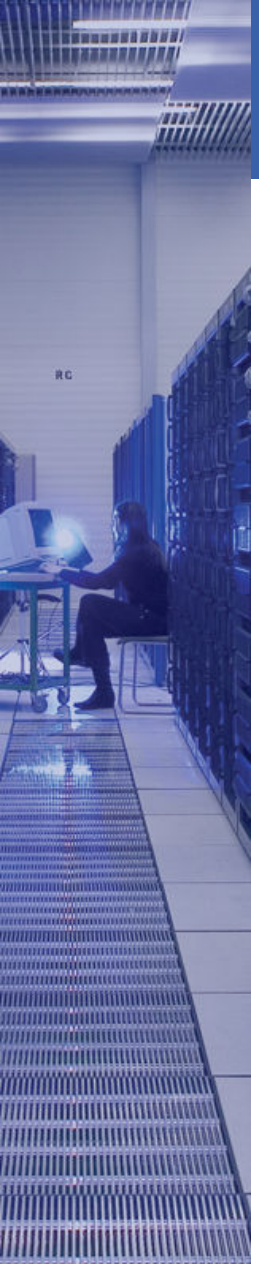

Processing Facility: Lxbatch

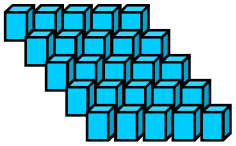

- Jobs are submitted from Lxplus or channeled through GRID interfaces world-wide
- Today about 2000 nodes with 14000 processors (cores)
- About 100000 user jobs are run per day
- Reading and writing 260 TB per day = 8 PB per month (comparison : we expect ~10 PB of raw data from LHC per year)
- Uses LSF as a management tool to schedule the various jobs from a large number of users.
- Expect a resource growth rate of ~30% per year

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

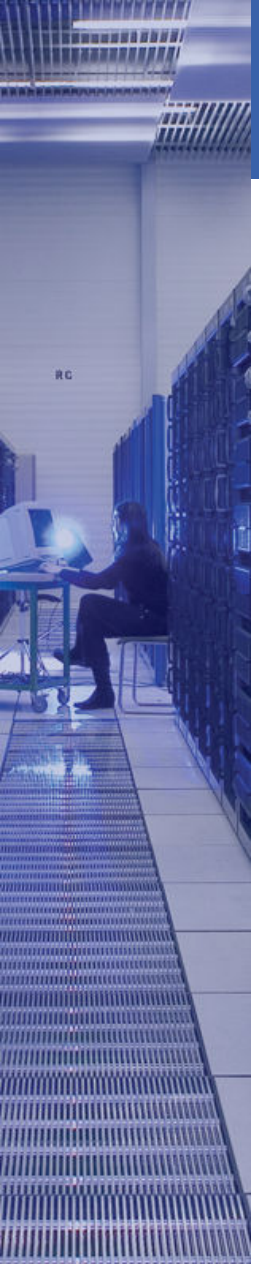

Mass Storage

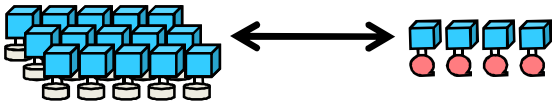

- Large disk cache in front of a long term storage tape system
- 1100 disk servers with 9 PB usable capacity
- Redundant disk configuration, 2-3 disk failures per day needs to be part of the operational procedures
- Logistics again : need to store all data forever on tape
  20 PB storage added per year, plus a complete copy every 4 years (change of technology)
- CASTOR data management system, developed at CERN, manages the user IO requests
- Expect a resource growth rate of ~30% per year

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

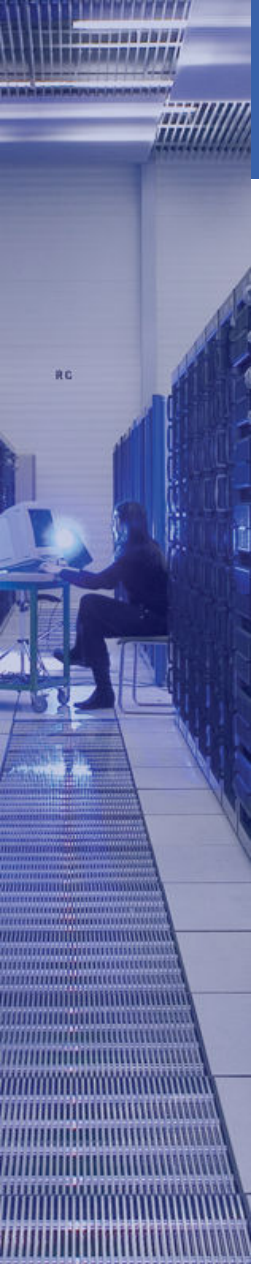

Mass Storage

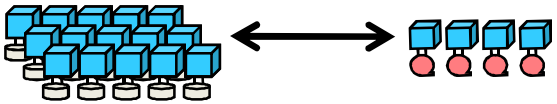

115 million files

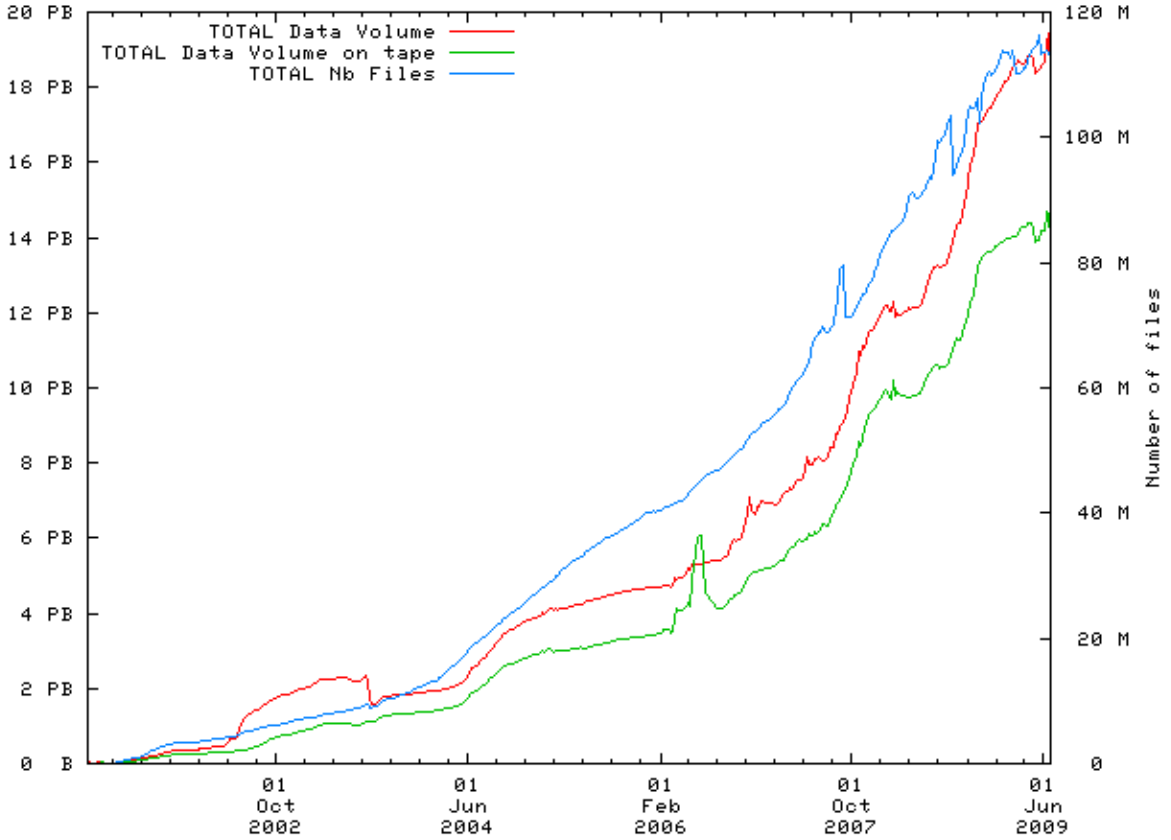

19 PB of data on tape already today

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Storage

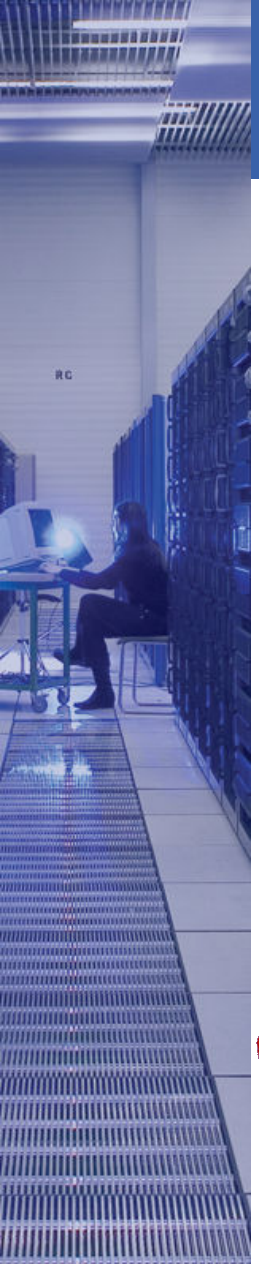

backup
per hour and per day

Administrative Data
1 million electronic documents
55000 electronic signatures per month
60000 emails per day
250000 orders per year

Users

> 1000 million user files

accessibility → 24*7*52 = always
tape storage → forever

continuous storage

Physics Data
20000 TB and
50 million files per year

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CERN Computer Center

ORACLE Data Base Server

240 CPU and disk server
150 TB

2.9 MW Electricity
and Cooling
2700 m2

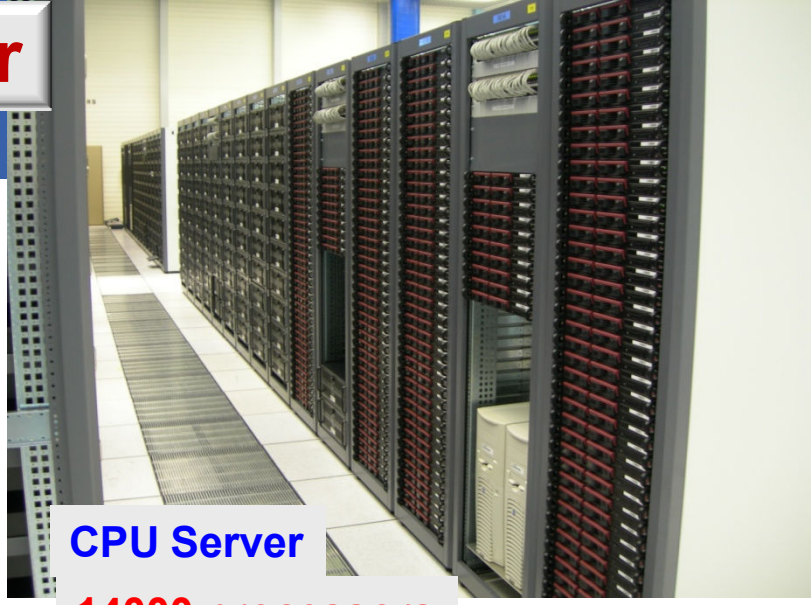

CPU Server

14000 processors

160 Gbits/s

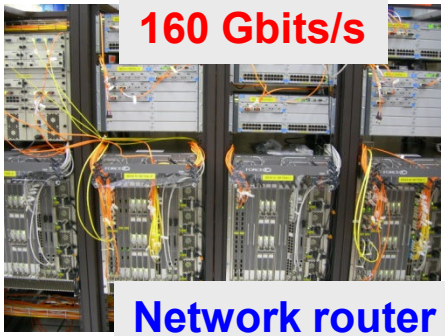

Network router

Tape Server
and
Tape Library

160 tape drives , 50000 tapes
40000 TB capacity

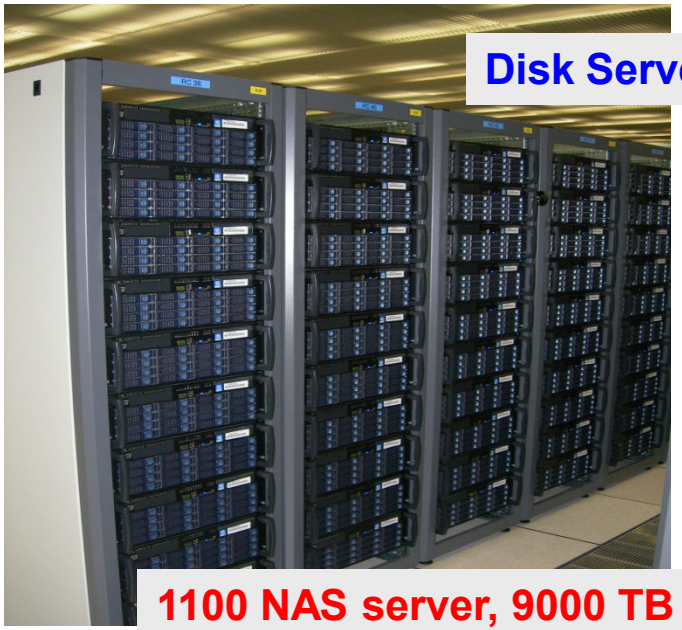

Disk Server

1100 NAS server, 9000 TB
22000 disks

Computing Resources

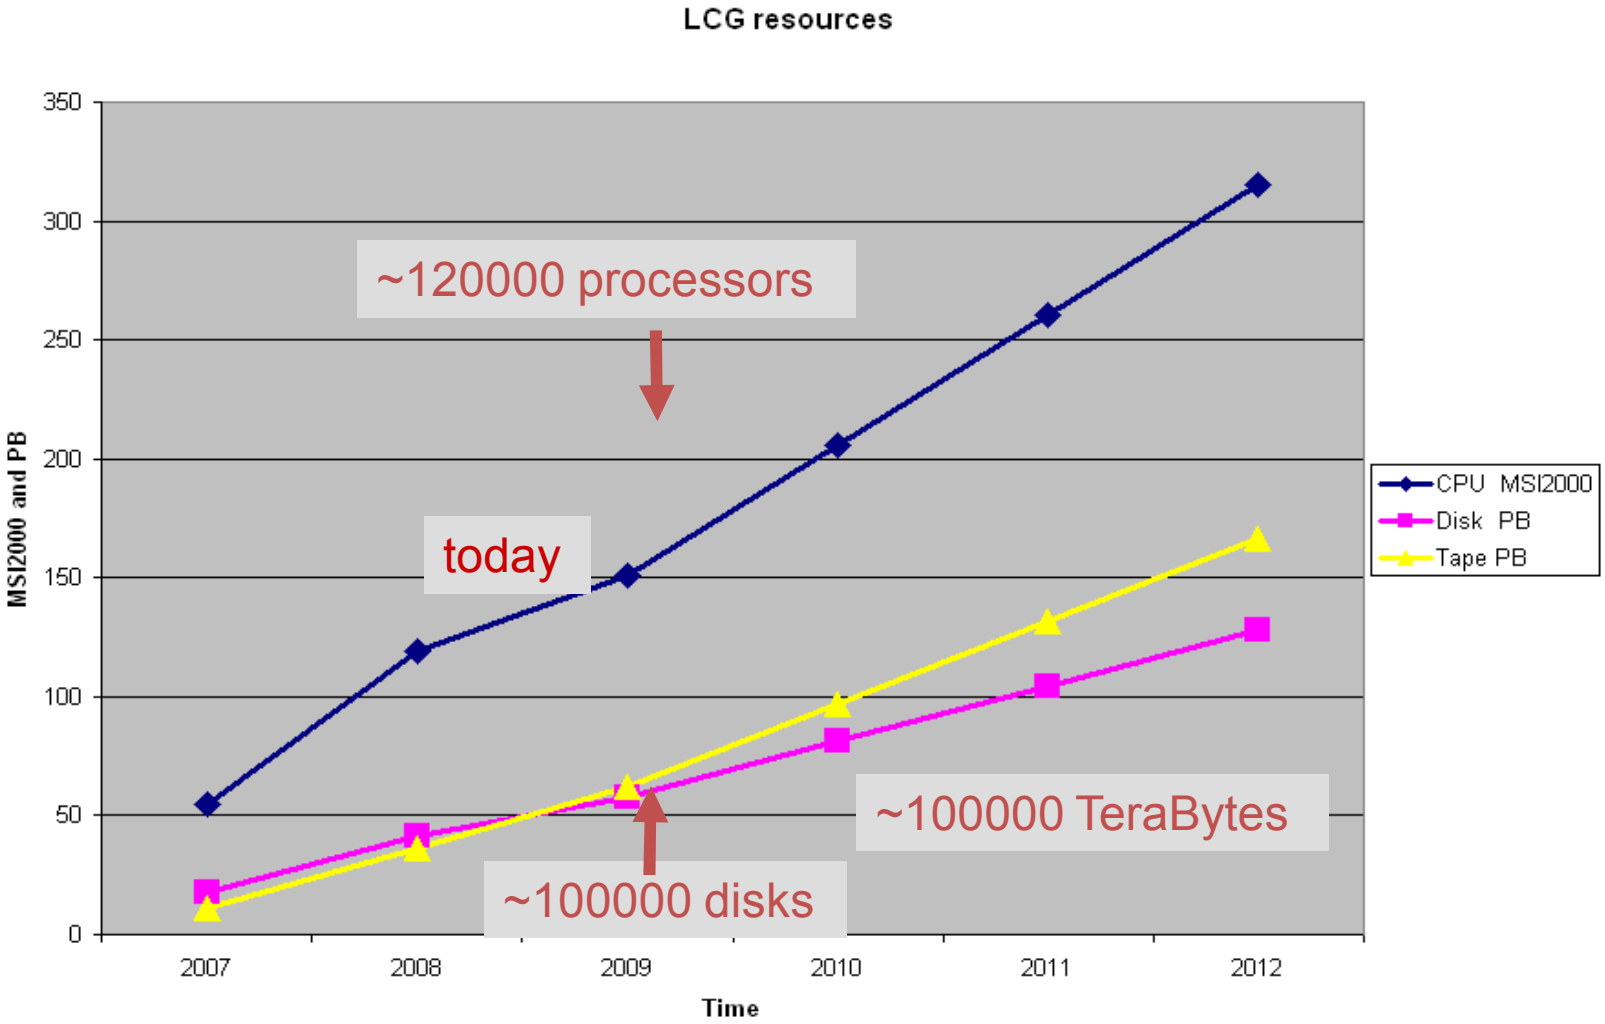

CERN can only contribute ~15% of these resource
→ need a world-wide collaboration

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

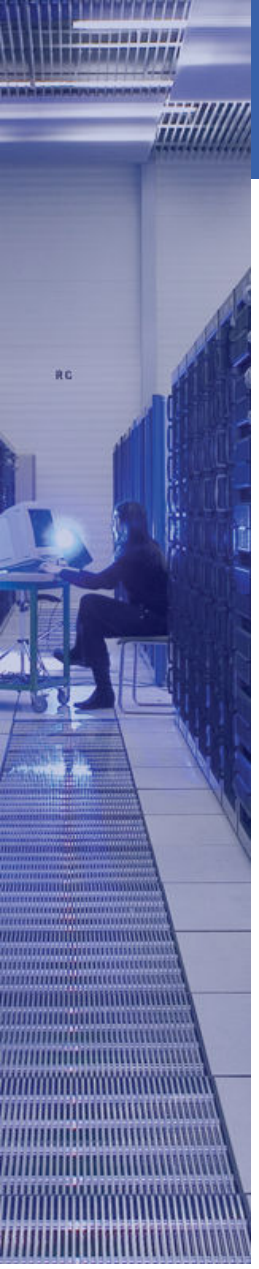

WEB and GRID

The World Wide Web provides seamless access to information that is stored in many millions of different geographical locations

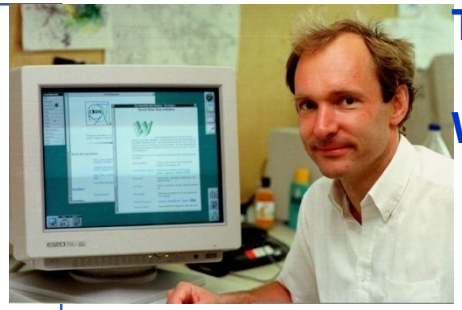

Tim Berners-Lee invented the World Wide Web at CERN in 1989

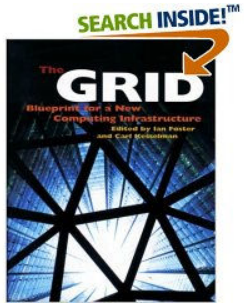

The Grid is an infrastructure that provides seamless access to computing power and data storage capacity distributed over the globe

Foster and Kesselman 1997

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

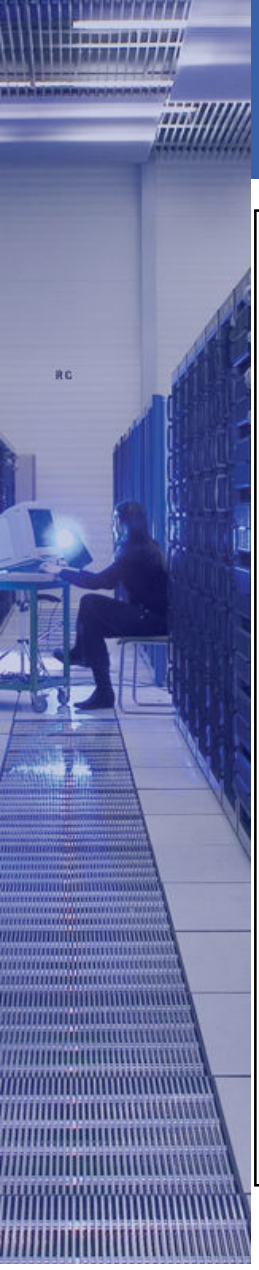

How does the GRID work ?

CERN IT Department

- It relies on advanced software, called **middleware**.
- Middleware automatically finds the **data** the scientist needs, and the **computing power** to analyse it.
- Middleware balances the load on different resources. It also handles **security, accounting monitoring** and much more.

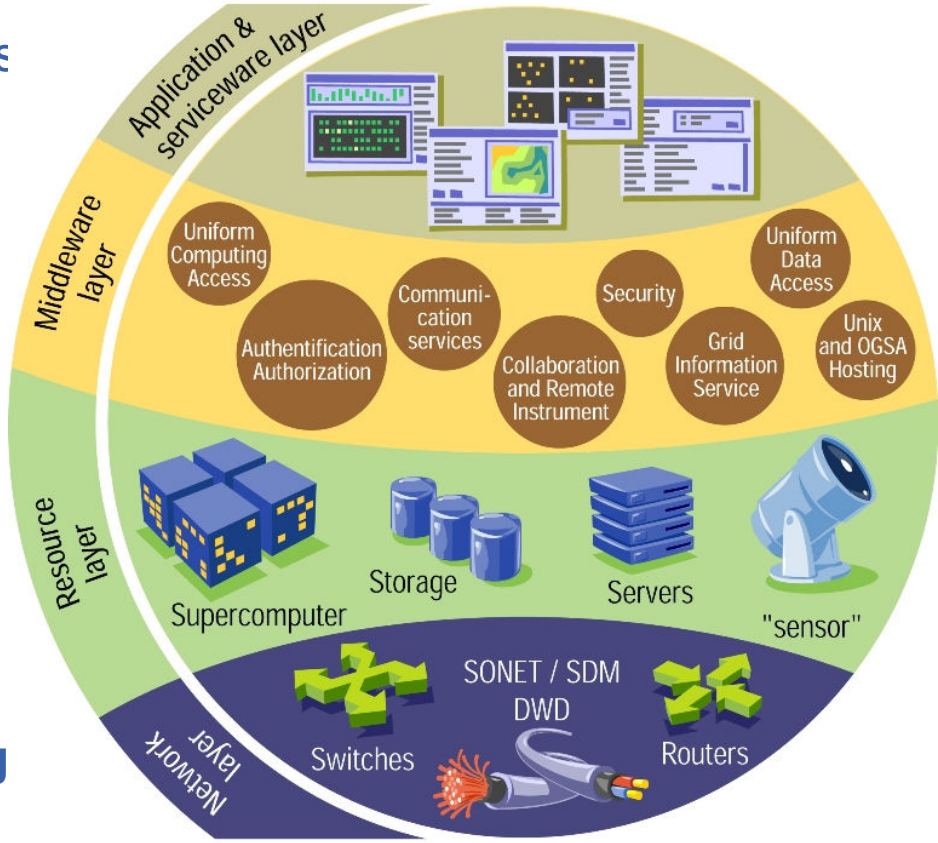

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

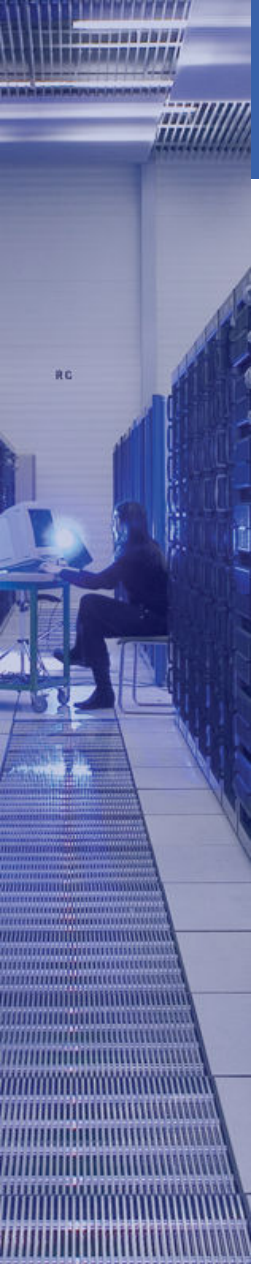

GRID infrastructure at CERN

GRID services

Work load management :
submission and scheduling of jobs in the GRID,
transfer into the local fabric batch scheduling systems

Storage elements and file transfer services :
Distribution of physics data sets world-wide
Interface to the local fabric storage systems

Experiment specific infrastructure :
Definition of physics data sets
Bookkeeping of logical file names and their physical location world-wide
Definition of data processing campaigns
→ which data need processing with what programs

Various types of PC's needed > 400

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

LHC Computing Grid : LCG

CERN IT Department

Grid Projects Collaborating in LHC Computing Grid

Open Science Grid

Today

Active sites	: >	300
Countries involved	: ~	70
Available processors	: ~	100000
Available space	: ~	100 PB

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/i

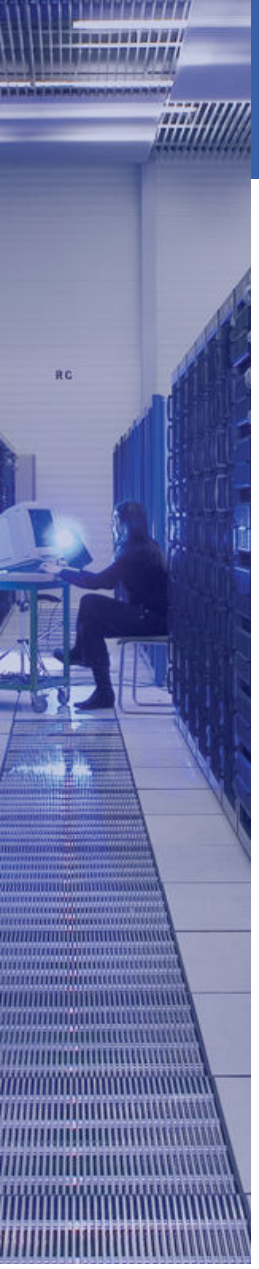

Links I

CERN IT Department

IT department

http://it-div.web.cern.ch/it-div/

http://it-div.web.cern.ch/it-div/need-help/

Monitoring

http://sls.cern.ch/sls/index.php

http://lemonweb.cern.ch/lemon-status/

http://gridview.cern.ch/GRIDVIEW/dt_index.php

http://gridportal.hep.ph.ic.ac.uk/rtm/

Lxplus

http://plus.web.cern.ch/plus/

Lxbatch

http://batch.web.cern.ch/batch/

CASTOR

http://castor.web.cern.ch/castor/

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

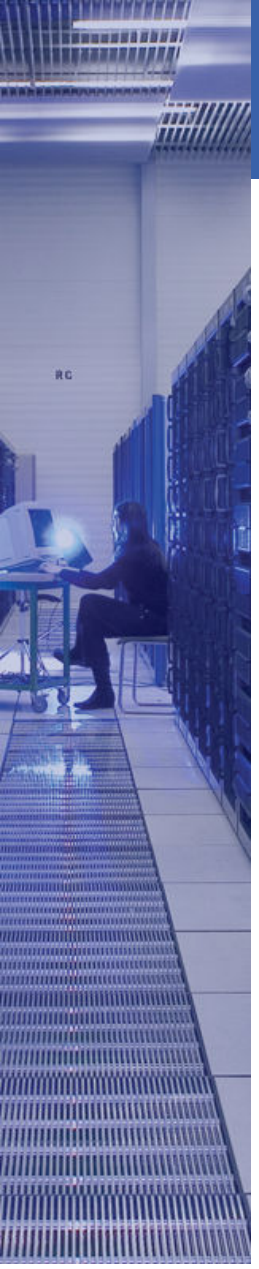

Links II

Windows, Web, Mail

https://winservices.web.cern.ch/winservices/

Grid

http://lcg.web.cern.ch/LCG/public/default.htm

http://www.eu-egee.org/

Computing and Physics

http://www.particle.cz/conferences/chep2009/programme.aspx

In case of further questions don't hesitate to contact me:

Bernd.Panzer-Steindel@cern.ch

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it