CERN IT Department

Status of LHC Computing from a Data Center Perspective

CERN School of Computing 2010 Uxbridge, U.K.

Frédéric Hemmer

IT Department Head - CERN

23rd August 2010

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

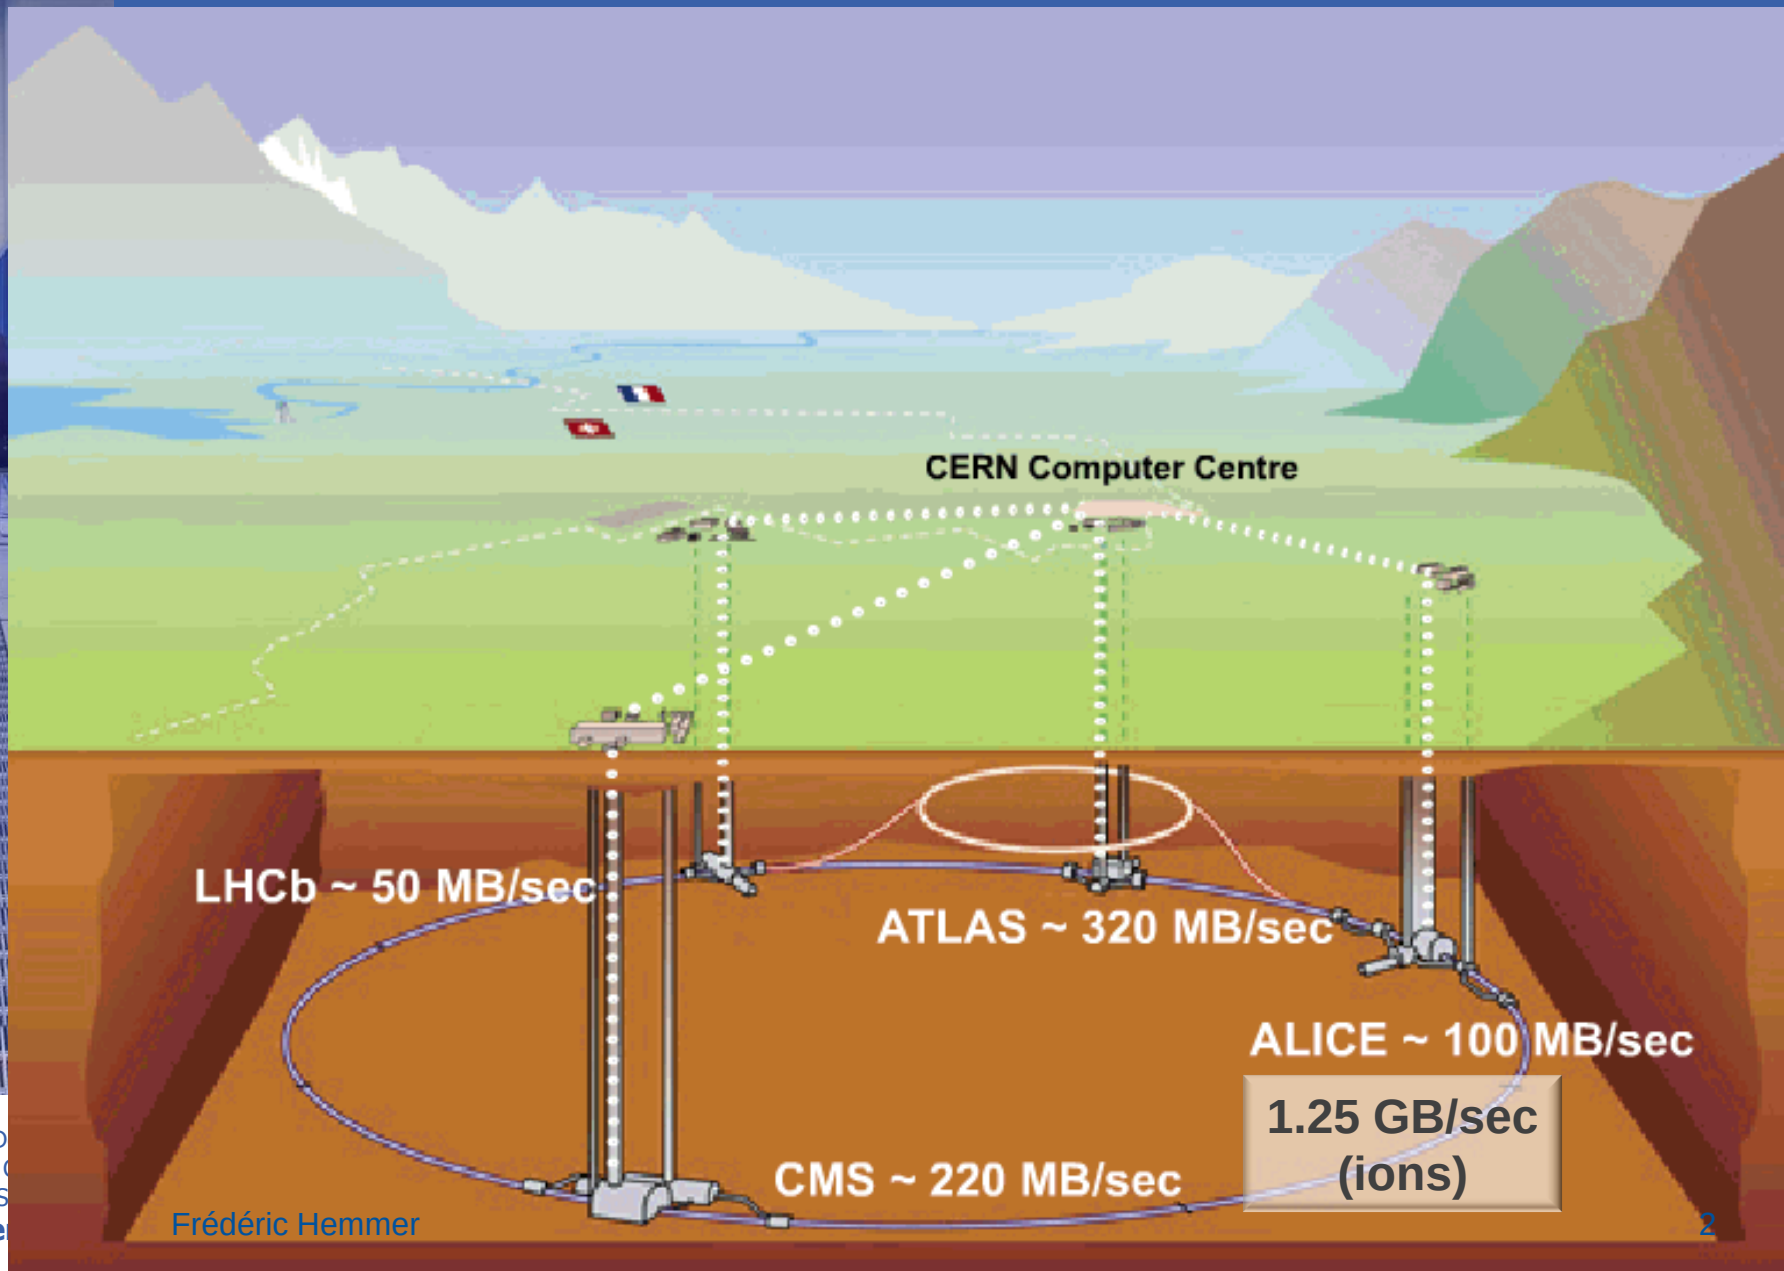
CERN Computer Centre
LHCb ~ 50 MB/sec
ATLAS ~ 320 MB/sec
ALICE ~ 100 MB/sec
1.25 GB/sec
(ions)
CMS ~ 220 MB/sec

Tier-0 Computing Tasks

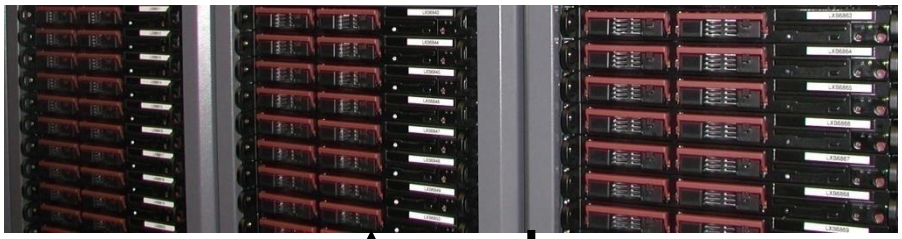

Scheduled work only!

700MB/s

420MB/s

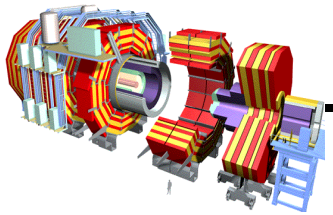

700MB/s

(1600MB/s)

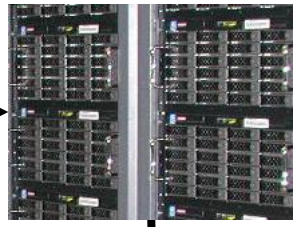

1120MB/s

(2000MB/s)

Averages! Need to be able to support 2x for recovery!

1430MB/s

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

The CERN Tier-0 in Numbers

CERN IT Department

- Data Centre Operations (Tier 0)
  - 24x7 operator support and System Administration services to support 24x7 operation of all IT services.
  - Hardware installation & retirement
    - ~7,000 hardware movements/year; ~1000 disk failures/year
  - Management and Automation framework for large scale Linux clusters

Servers	8,076
Processors	13,802
Cores	50,855
HEPSpec06	359,431

Disks	53,728
Raw disk capacity (TB)	45,331
Memory modules	48,794
RAID controllers	3,518

Tape Drives	160
Tape Cartridges	45000
Tape slots	56000
Tape Capacity (TB)	34000

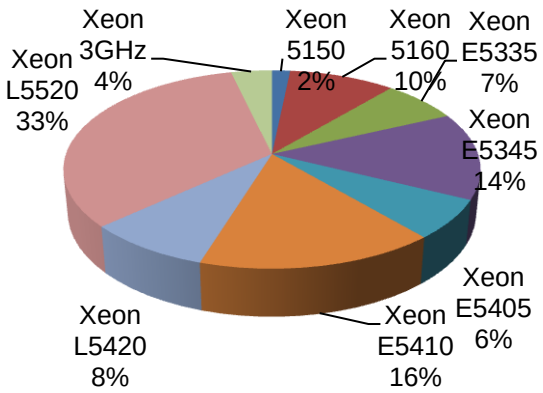

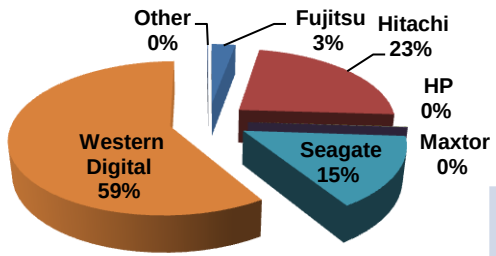

High Speed Routers (640 Mbps → 2.4 Tbps)	24
Ethernet Switches	350
10 Gbps ports	2000
Switching Capacity	4.8 Tbps

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

The WLCG Collaboration

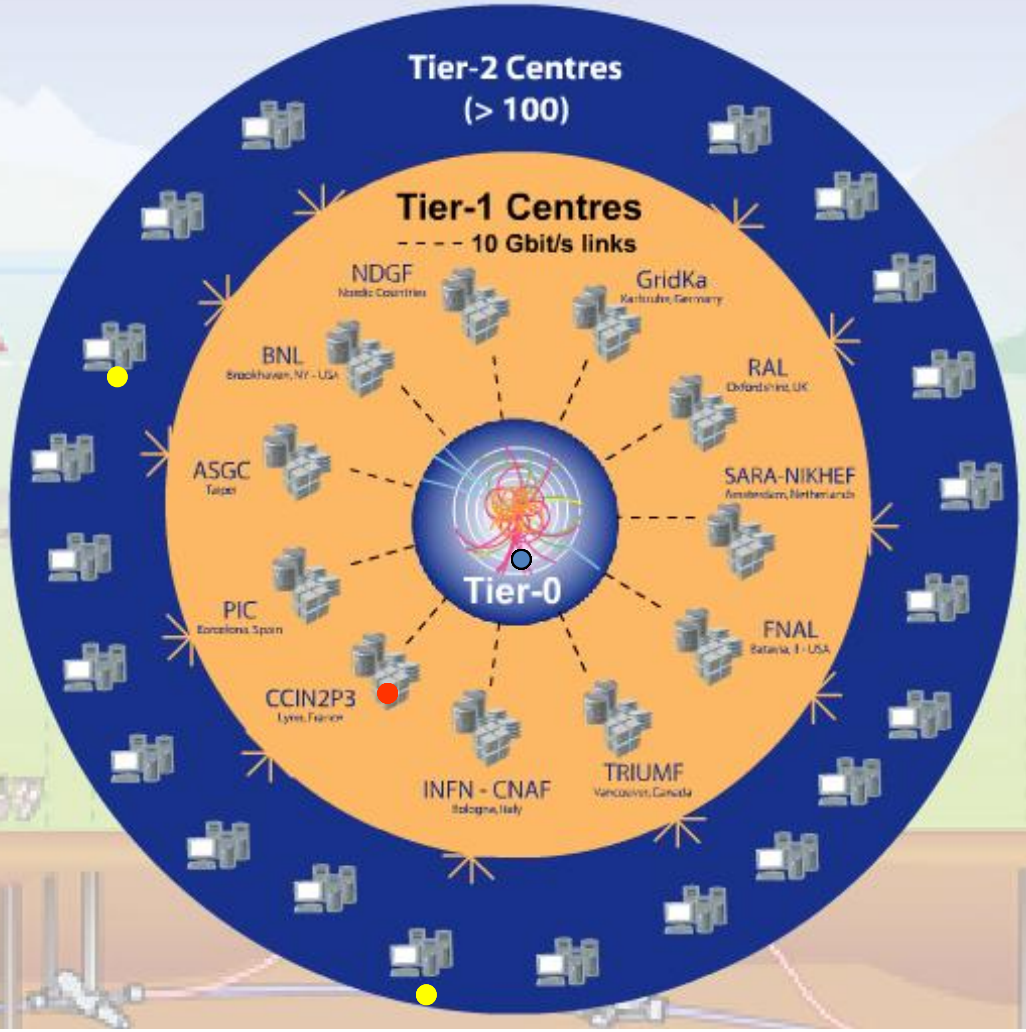

Tier-0 (CERN):
- Data recording
- Initial data reconstruction
- Data distribution

Tier-1 (11 centres):
- Permanent storage
- Re-processing
- Analysis

Tier-2 (~130 centres):
- Simulation
- End-user analysis

Today WLCG Status is:

- WLCG running increasingly high workloads:
  - ~1 million jobs/day
    - Real data processing and re-processing
    - Physics analysis
    - Simulations
  - ~100 k CPU-days/day
- Unprecedented data rates

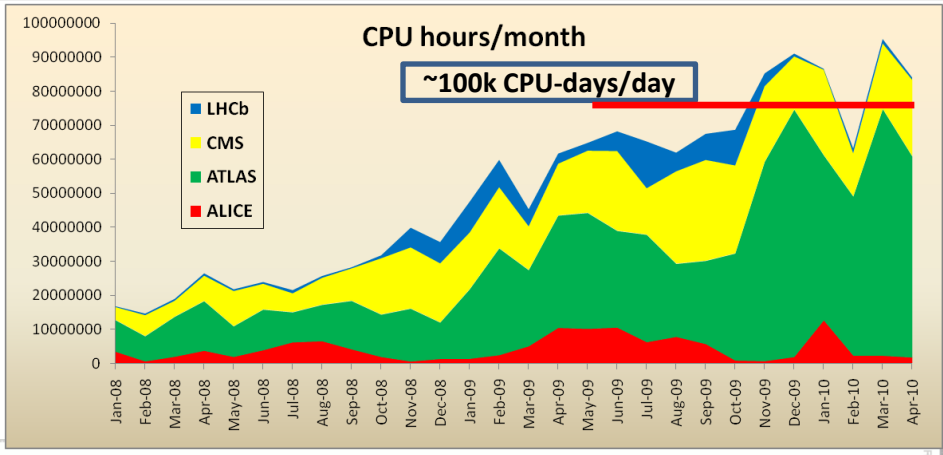

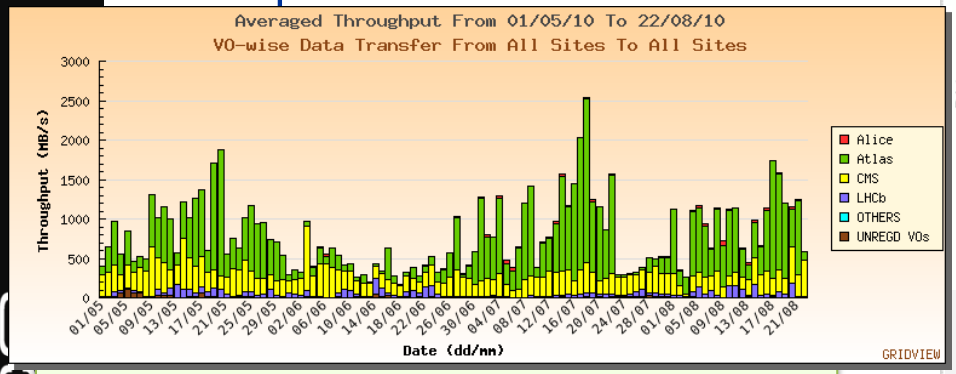

Data export during data taking:
- According to expectations on average

Castor traffic over the last 2 months:
> 4 GB/s input
> 13 GB/s served

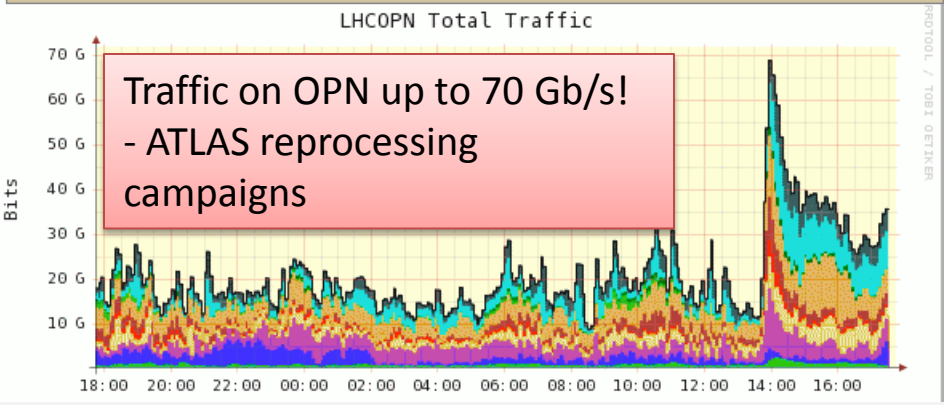

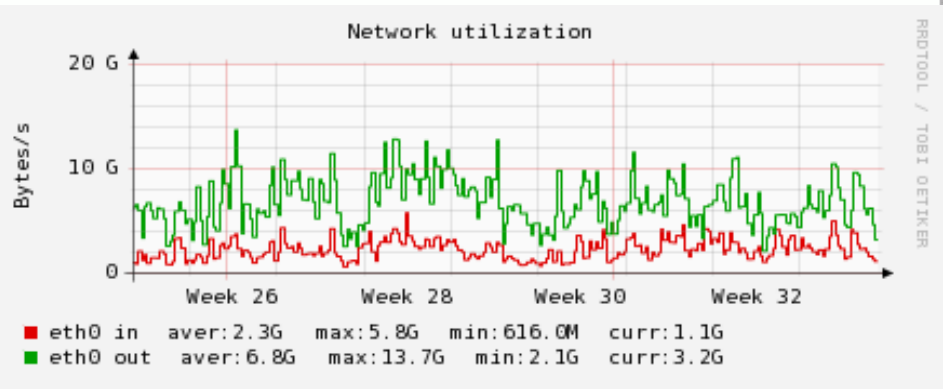

Readiness of the Computing

- Data reaches Tier 2s within hours

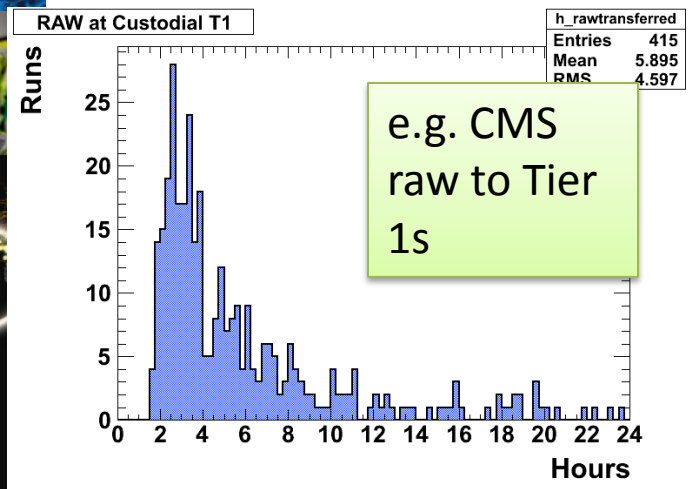

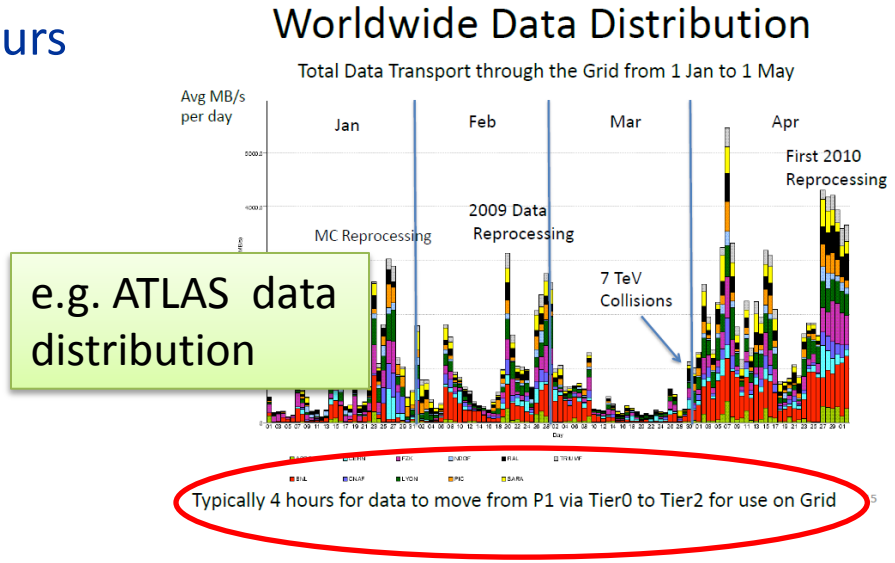

- Increasing numbers of (analysis u
  - E.g.:~500 grid users in each ATLAS/CMS; ~200 in ALICE

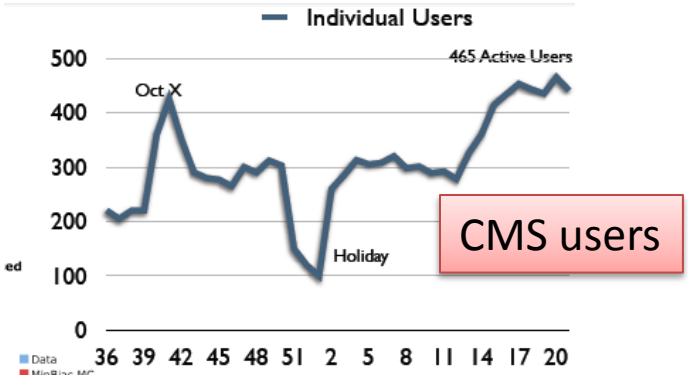

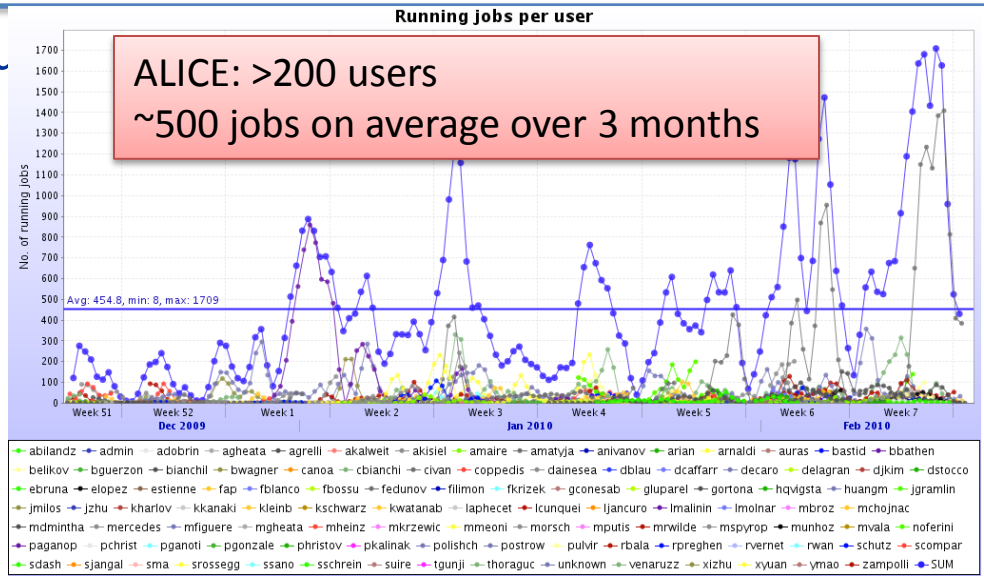

WLCG Worldwide LHC Computing Grid

So, everything fine ?

Not quite ...

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

Worldwide data distribution and analysis

From F. Gianotti
ICHEP 2010

Total throughput of ATLAS data through the Grid: from 1st January until yesterday

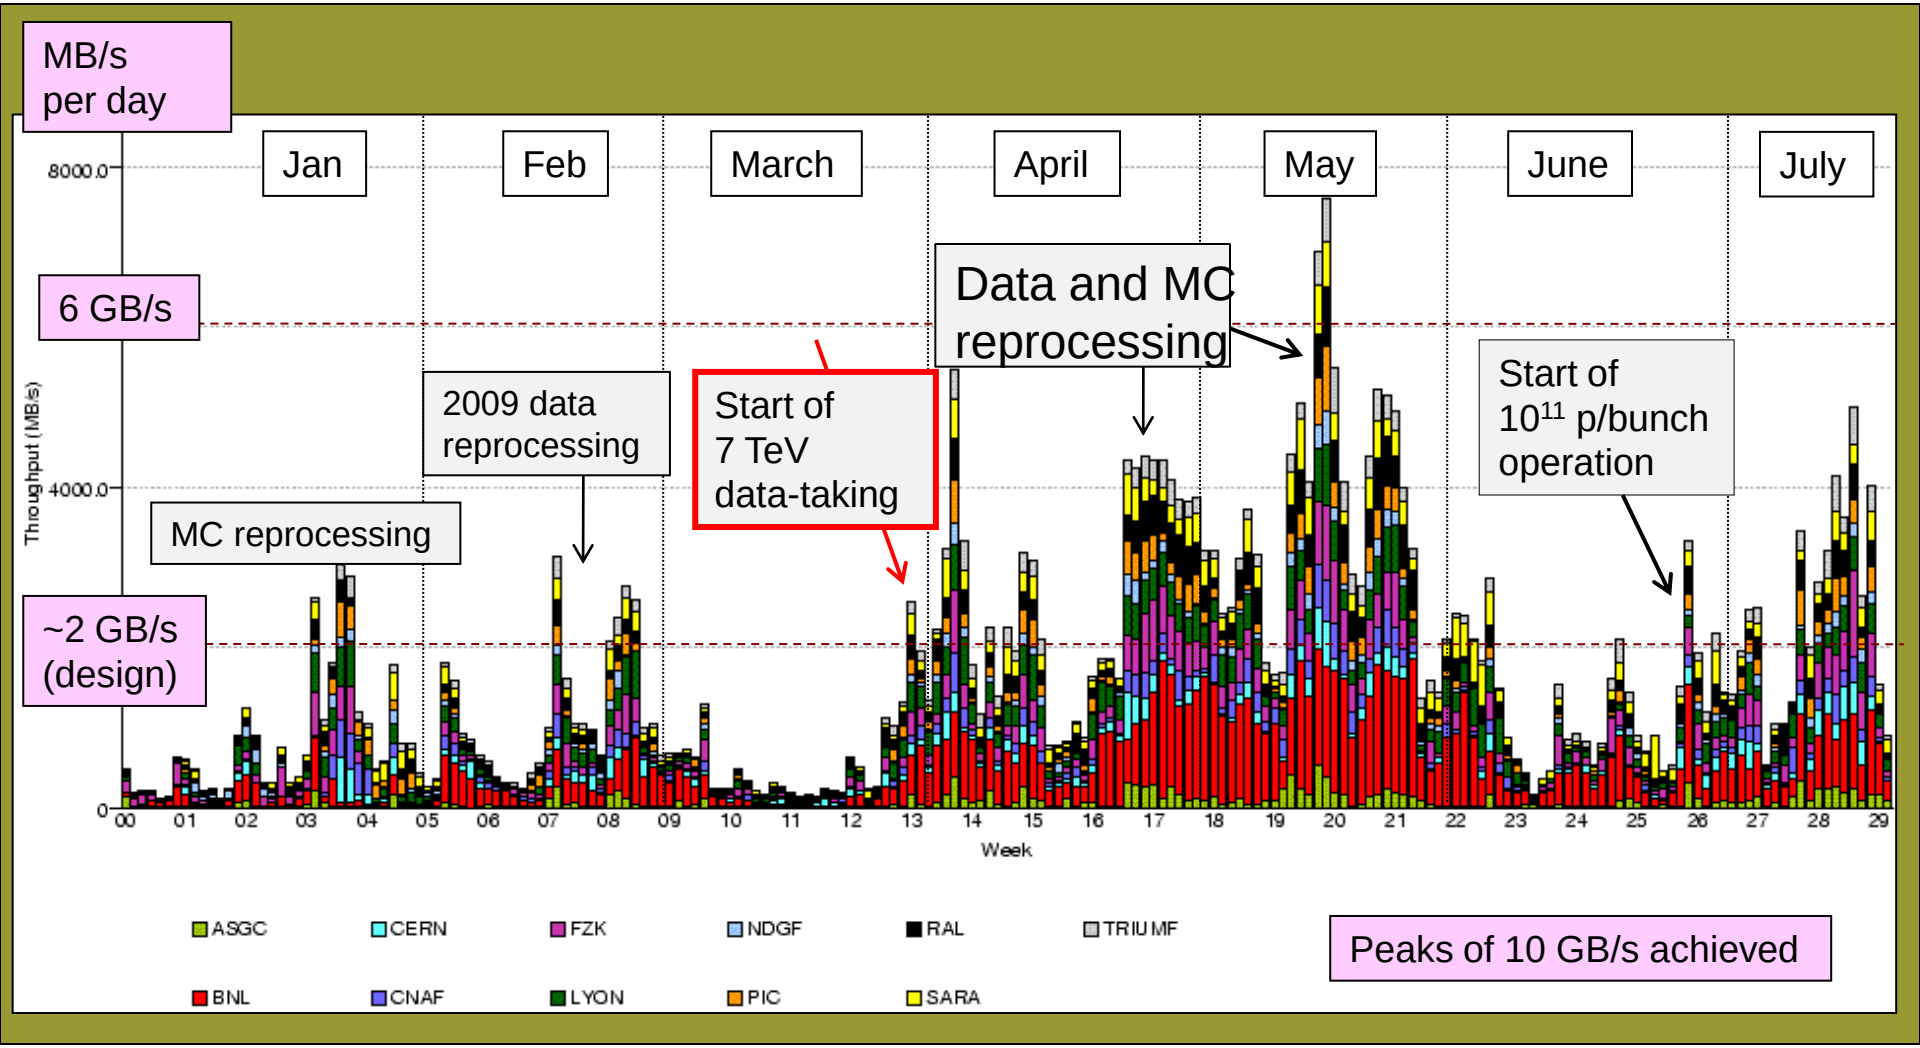

GRID-based analysis in June-July 2010:
> 1000 different users, ~ 11 million analysis jobs processed

ALICE data loss

- A “transparent” configuration change in Castor software resulted in data being directed across all available tape pools instead of to the dedicated raw data pools
  - For ALICE, ATLAS, CMS this included a pool where the tapes were re-cycled after a certain time
- The result of this was that a number of files were lost on tapes that were recycled
- For ATLAS and CMS the tapes had not been overwritten and could be fully recovered (fall back would have been to re-copy files back from Tier 1s)
- For ALICE 10k files were on tapes that were recycled, including 1973 files of 900 GeV raw data

- Actions taken:
  - Underlying problem addressed; all recycle pools removed
    - Software change procedures being reviewed now
  - Tapes sent to IBM and SUN for recovery – have been able to recover ~97% of critical (900 GeV sample) files, ~50% of all ALICE files
    - We have been very lucky, 1717 tapes were recovered by IBM & STK/Sun/Oracle
  - Work with ALICE to ensure that always 2 copies of data available
    - In HI running there is a risk for several weeks until all data is copied to Tier 1s; several options to mitigate this risk under discussion
  - As this was essentially a procedural problem: we will organise a review of Castor operations procedures (software development, deployment, operation, etc...) together with experiments and outside experts.

A Word on Computer Security

CERN IT Department

August 2010

- It is “your” responsibility to write secure code, and to fix promptly eventual flaws
- Not only for ensuring proper operation of your application
  - But also, failure to do so might bring your and other organizations in a very embarrassing position.

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

Summary

CERN IT Department

- The readiness of LHC Computing has now clearly been demonstrated
  - There are however large load fluctuations
  - Experiments Computing Models will certainly evolve
    - And WLCG Tiers will have to adapt to those changes
  - Change Management is critical for ensuring stable data taking
- Some challenges ahead
  - Large number of objects with permanent flux of change
    - Automation is critical
  - (Distributed) Software complexity
  - Computer Security remains a permanent concern

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