Updating brochures

CERN Grid & ATLAS computing 4 pages brochures: pages 2 to 4

ATLAS 2011 Fact Sheet: pages 5 to 7

ATLAS visitor centre smart board slide: pages 8 to 9

“The grid”- by CERN IT

2006 : https://cds.cern.ch/record/976156

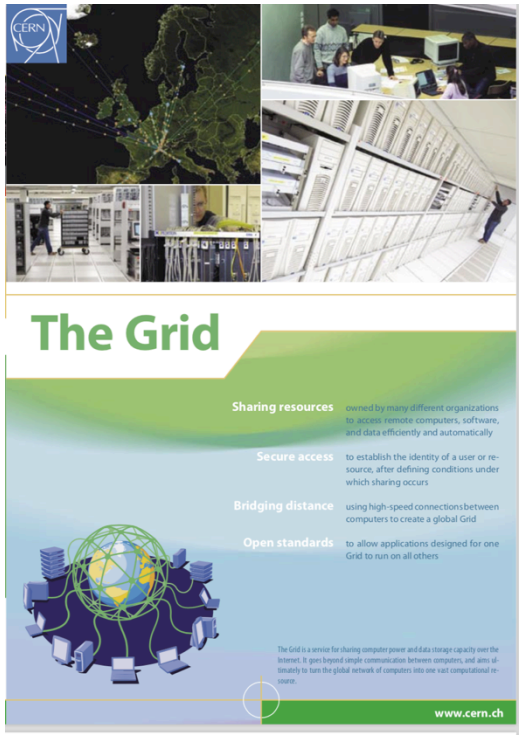

- CERN data centre in 2006:
  - 5000 PC’s
  - 5 Petabytes storage in disk & tape
  - Can write 1 GB/s on tape

- Resource used by ATLAS in 2017:
  - 421 kHS06 = 20 000 laptops (double core)
  - 25 Petabytes of data on disk; 58 on tape
  - Detector sends 2 GB/s ; 1 to 1.5 kHz
  - Total 2017 raw data volume = 10 Petabytes

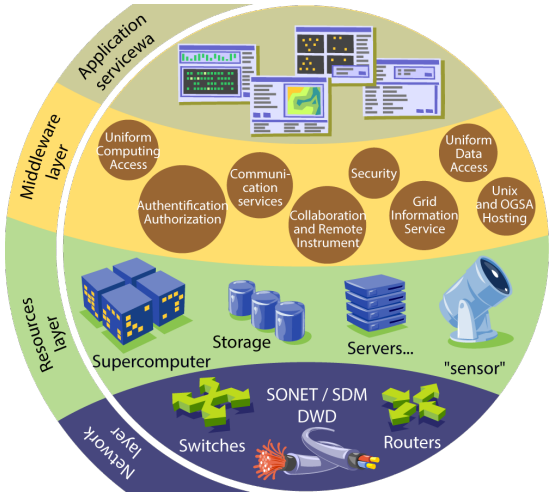

“The grid”- by CERN IT

2012 update: https://cds.cern.ch/record/1421296

- CERN data centre in 2012
  - LHC experiments produce 25 Petabytes per year, the equivalent of 5.3 million DVD movies
  - T0 = 65 000 processor cores, corresponding to ~ 20% of the LHC computing capacity used for LHC
  - 30 Petabytes storage in disk, 45 on tape
  - Optical fibre links to experiments and T1: multiples of 10 gigabits/s

- Resource used by ATLAS only in 2017:
  - Total 2017 raw data volume = 10 Petabytes
  - 421 kHS06 = 40 000 cores
  - 25 Petabytes of data on disk; 58 on tape
  - Detector sends 2 GB/s ; 1 to 1.5 kHz

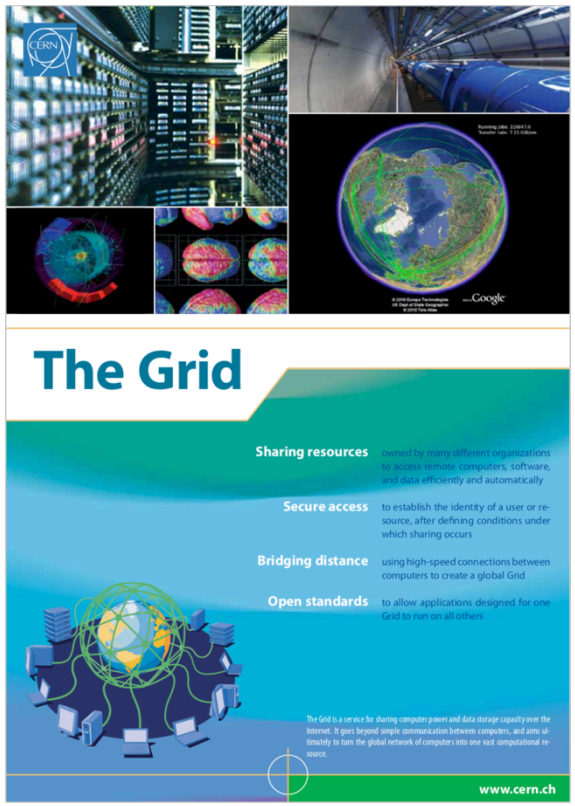

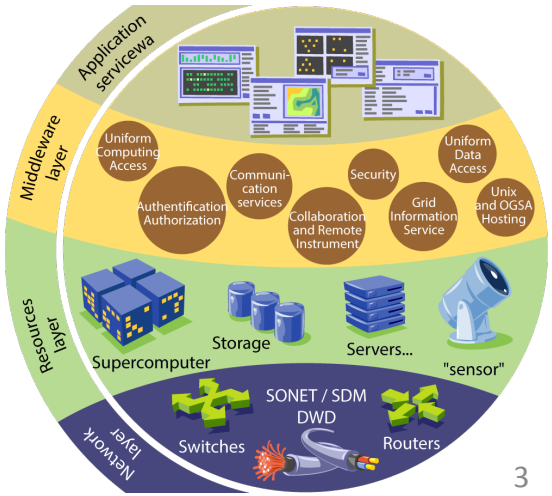

“Computing in the ATLAS Experiment” – 2013

CDS link ????

Page 1: the ATLAS detectors produces more than 5 Petabytes per year	2017 raw data = 10 Petabytes.
Page 3: computing infrastructure	ATLAS moved for run 2 to a “network oriented” model of ~ 30 “nucleus + satellites”
Page 3: Total capacity CERN: 111 kHS06, disk 9 PB, tape 18 PB T1: 260 kHS06, disk 27 PB, tape 30 PB T2: 266 kHS06, disk 47 PB, no tapes Data transfer rate to T0: 320 MB/s Data export rate to T1: 1.2 GB/s	**2017 ATLAS only:** T0: 421 kHS06, disk 25 PB, tape 58 PB T1: 826 kHS06, disk 67 PB, tape 102 PB T2: 1505 kHS06, disk 78 PB, no tapes Data transfer rate to T0: 2 GB/s Every week 1 PB is moved to/in T0, 5 PB moved on the grid for prod & analysis
Page 3: network links - Links from CERN to T1: > 10 Gb/s - Links between T1 and T2: > 1 Gb/s	2016 ECFA R&D session talk by S.Campana https://indico.cern.ch/event/524795/ - 2016: 10 to 100 Gb links - 2026: 1 to 10 Tb links

ATLAS fact sheet / 2010
Page 9: computing

https://cds.cern.ch/record/1457044

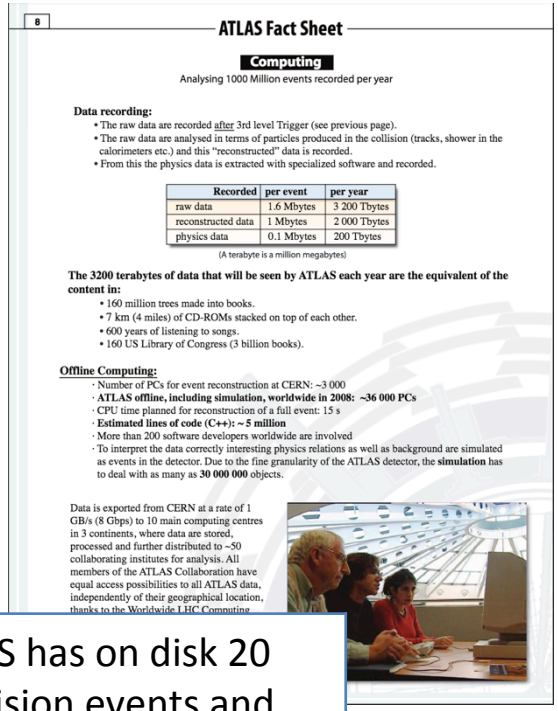

8

ATLAS Fact Sheet

Computing

Analysing 1000 Million events recorded per year

Data recording:

- The raw data are recorded after 3rd level Trigger (see previous page).
- The raw data are analysed in terms of particles produced in the collision (tracks, shower in the calorimeters etc.) and this "reconstructed" data is recorded.
- From this the physics data is extracted with specialized software and recorded.

Recorded	per event	per year
raw data	1.6 Mbytes	3 200 Tbytes
reconstructed data	1 Mbytes	2 000 Tbytes
physics data	0.1 Mbytes	200 Tbytes

(A terabyte is a million megabytes)

The 3200 terabytes of data that will be seen by ATLAS each year are the equivalent of the content in:

- 160 million trees made into books.
- 7 km (4 miles) of CD-ROMs stacked on top of each other.
- 600 years of listening to songs.
- 160 US Library of Congress (3 billion books).

Offline Computing:

- Number of PCs for event reconstruction at CERN: ~3 000
- **ATLAS offline, including simulation, worldwide in 2008: ~36 000 PCs**
- CPU time planned for reconstruction of a full event: 15 s
- **Estimated lines of code (C++): ~ 5 million**
- More than 200 software developers worldwide are involved
- To interpret the data correctly interesting physics relations as well as background are simulated as events in the detector. Due to the fine granularity of the ATLAS detector, the **simulation** has to deal with as many as **30 000 000** objects.

Data is exported from CERN at a rate of 1 GB/s (8 Gbps) to 10 main computing centres in 3 continents, where data are stored, processed and further distributed to ~50 collaborating institutes for analysis. All members of the ATLAS Collaboration have equal access possibilities to all ATLAS data, independently of their geographical location, thanks to the Worldwide LHC Computing

Computing **Analysing 1000 Million events recorded per year**	2017: ATLAS has on disk 20 billions collision events and 16 billions simulated events
Data recording: - The raw data are recorded after **the 3rd level Trigger** - the raw data are analysed in terms of particles produced in the collision (tracks, shower in the calorimeters, etc.) and this "reconstructed" data is recorded. - From this the physics data is extracted with specialized software and recorded	**HLT**

Recorded	per event	per year
raw data	1.6 Mbytes	3 200 Tbytes
reconstructed data	1 Mbytes	2 000 Tbytes
physics data	0.1 Mbytes	200 Tbytes

(A terabyte is a million megabytes)

Recorded	Per event	Per year
Raw data	1.0 MB	10 PB
Reconstructed data	0.3 MB	Scales-> 3 PB
Physics data: 20 trains producing each 5 formats	Target: total to be the same as reconstructed data	idem

Numbers taken from the oct 2017 report to C-RSG

Brochure:	2017:
The 3200 terabytes of data that will be seen by ATLAS each year are the equivalent of the content in:	Raw data = 10 Petabytes Access/move 6 Petabytes per week TDAQ output: 2 GB/s
• 160 million trees made into books • 7 km of CD-ROMs stacked on top of each other • 600 years of listening to songs • 160 US library of Congress (3 billions books)	• Blu-ray standard = 50 GB -> ATLAS fills 200 10^3 disks per year. • Netflix = 1 GB/hour (3 for HD) • Whole of LHC data = Facebook uploads

Offline computing [2008]:	**Update 10 years later [2018]:**
Number of PCs for event reconstruction at CERN: ~ 3000	For ATLAS only: 400 kHS06 @ T0 = 40 000 cores (PC's have 2 these days) [1]
ATLAS offline, including simulation, worldwide: ~36 000 PCs	For ATLAS only: 2 M HS06 = 200 000 cores [1]
CPU time planed for reconstruction of a full event: 15 s	Still about right, although depends on pileup [1]
Estimated lines of code (C++): ~ 5 millions	~ 4 millions in C++ , 1.4 in Python [2]
More than 200 software developers worldwide are involved	Still about right, even optimistic [3]
The simulation has to deal with ~ 30 000 000 objects	No reason why it should have changed
Data is exported from CERN at a rate of 1 GB/s to 10 main computing centres in 3 continents.	Reports now mention the flow of data moved per week within cern (1 PB) and on the grid for prod & analysis (5 PB) [1]

References:

[1] April 2018 report to C-RSG https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/ComputingModel

[2] June 2018 Sw&Comp week plenary talk by C.Adam, slide 10 by Ed Moise
https://indico.cern.ch/event/646942/contributions/2974818/attachments/1674621/2688132/OutreachAndComm.pdf

[3] Feb 2018 ATLAS week plenary talk by Ed Moise, slide 11
https://indico.cern.ch/event/648945/contributions/2855642/attachments/1605619/2547705/Plenary23022018.pdf

AVC Smart board slide

Moved to a network and resources driven “nucleus + satellites” model for run 2

25% of the 2017 simulation done on HPC / cloud

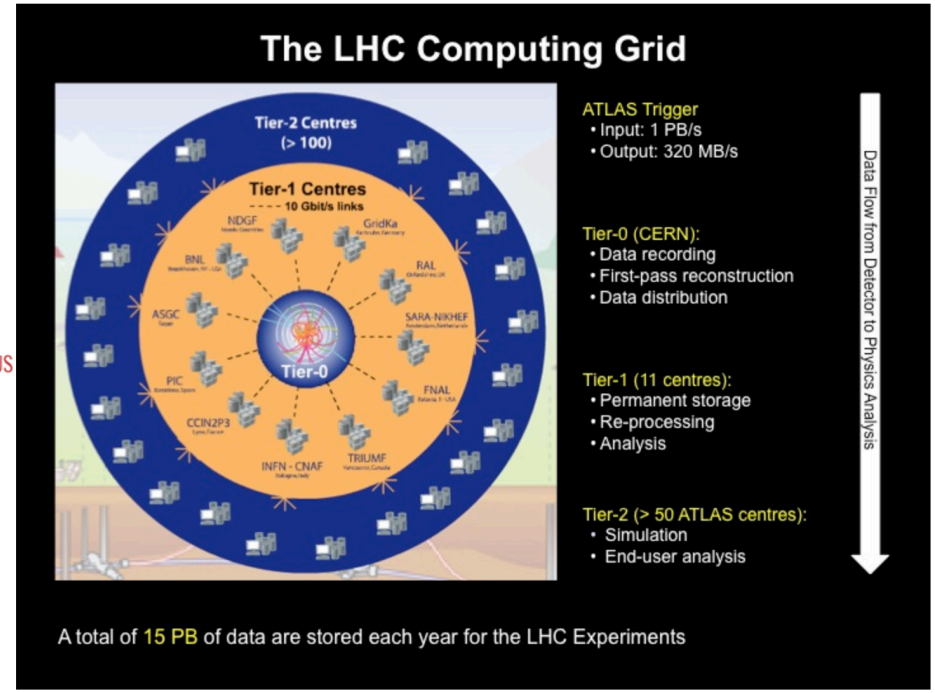

◀ PREVIOUS

NEXT ▶

ATLAS Trigger output: 2 GB/s

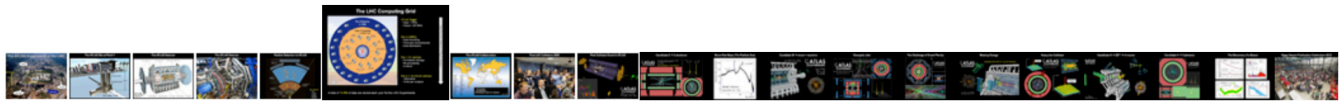

ATLAS only: produced 10 PB raw data in 2017
ATLAS only: 25 PB on disk, 58 on tape in 2017

As an example, 2018 summer student lecture slides [courtesy A. Sfyrla]

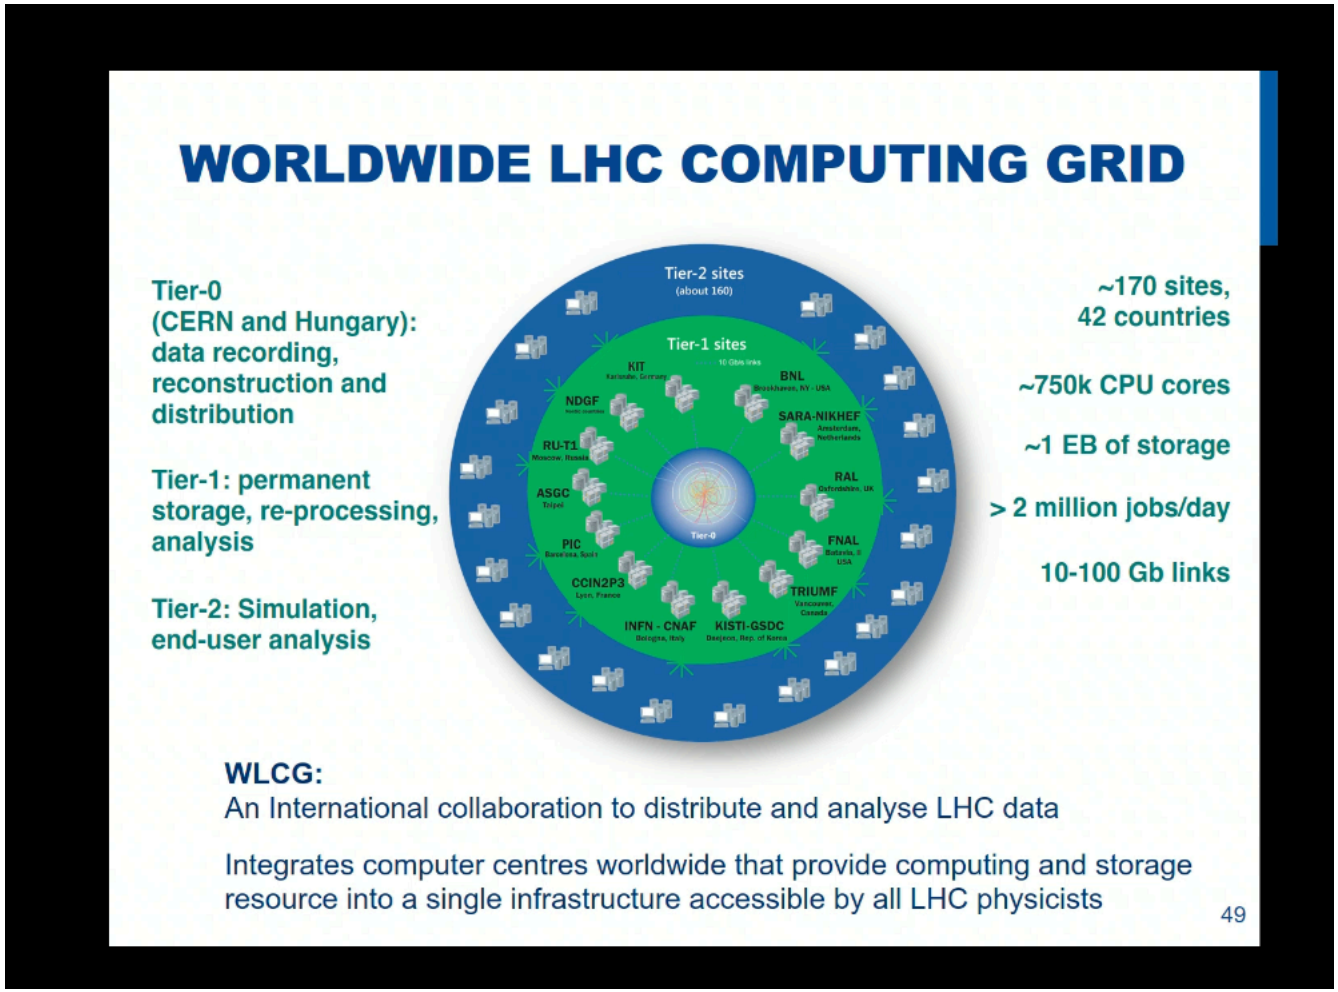

Very similar slide in the WLCG management talks, e.g.
https://www.slideshare.net/HelixNebulaTheScienc/wlcg-hn-scicloud030417-1 [slide 3]

What about future ?

ECFA 2016 HL-LHC workshop, R&D session on sw & computing
https://indico.cern.ch/event/524795/