Computing evolution

Donatella Lucchesi
Università di Padova e Sez. INFN Padova

Vincenzo Vagnoni
INFN Bologna

The LTS1 2014-Workshop on the Long Term Strategy of INFN-CSN1
The next 10 years of accelerator based experiments

Status

Resource usage and future requests dominated by LHC experiments → focus only on LHC

- Computing Model evolves continuously driven by technology evolution
- One example: network

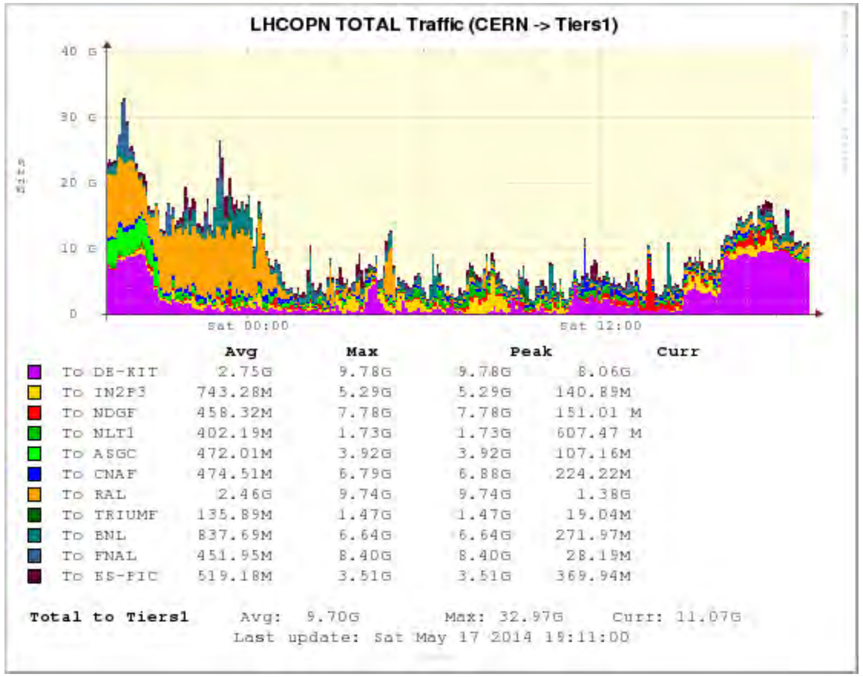

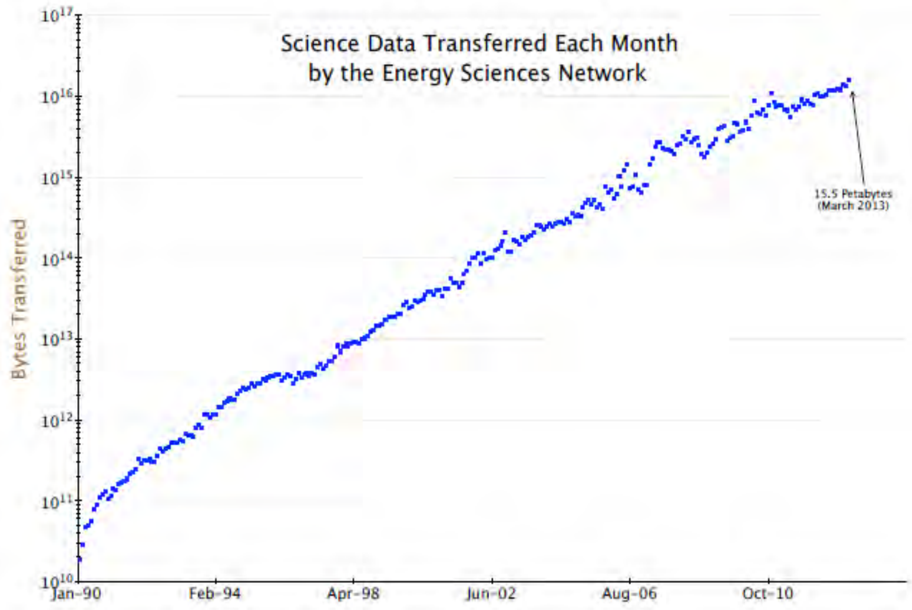

Current Computing Model

The old hierarchical model has evolved towards a full mesh with little distinction between Tier-N roles thanks to networks improvements

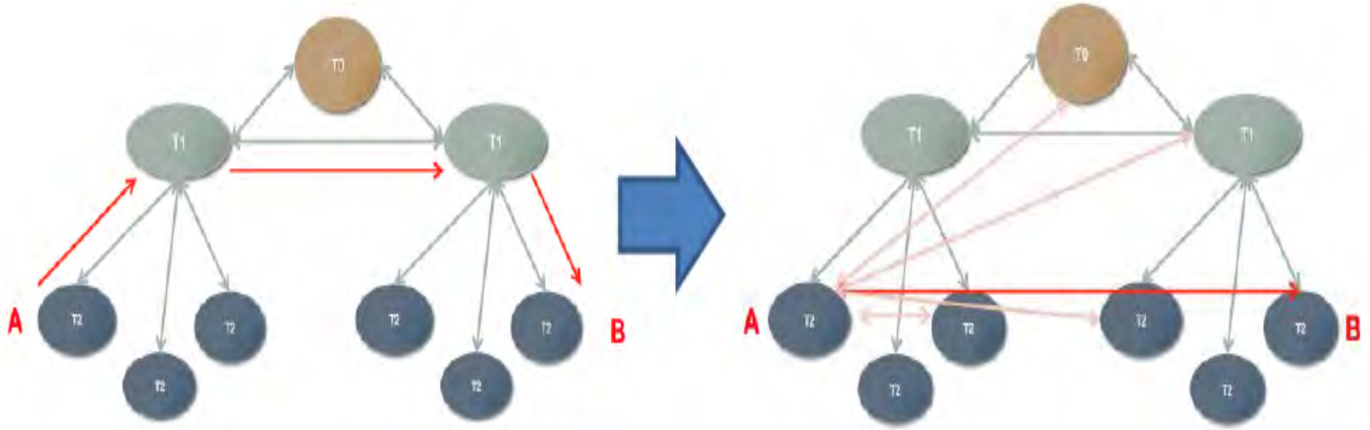

Tier-N differences:
- tape system → long term data custodial for the moment
- services available

LHC pledged resources for 2014

	Tier-0	Tier-1	Tier-2
CPU (kHS06)	401	750	1036
Disk (PB)	30.3	80.8	89.9
Tape (PB)	73.5	116	-

Example: Distributed Tier-0

- Tier-0 located at CERN and Wigner Center (Budapest, Hungary)
- two dedicated 100 Gb/s links
- in production
  - >1000 worker nodes installed
  - disk cache scaled with CPU
- No significant difference in efficiency between Geneva and Wigner
- Detailed performance study will set the scale for remote data access

The Future: definition

- Run2: 2015-2018
  - Up to 1.5e34, 25 ns, 13-14 TeV
  - Up to 50 fb-1/y
  - **<PU> = 40**

we know what we need

- Run3: 2020-2022
  - Up to 2.5e34, 25 ns, 13-14 TeV
  - Up to 100 fb-1/y
  - **<PU> = 60**

we know what we expect, but..

- Run4: 2025-2028
  - Up to ~5e34, 25 ns 13-14 TeV
  - Up to 300 fb-1/y
  - **<PU> = 140+**

into the darkness!

What drives the needs for 2015+

ATLAS and CMS: keep low the inclusive muon and electron p_T thresholds

LHCb: increase in the signal rate by factor 2.5 due to energy, to profit from this it must reach at least 10 kHz

ALICE: increase the TPC and TRD readout by factor 2

Hz	ALICE	ATLAS	CMS	LHCb
2012	400 Hz 330 MB/s (p-p) 540 MB/s (p-Pb)	550 Hz 440 MB/s	460+360 Hz 328 MB/S	5000 Hz 300 MB/s
2015	500 Hz 525 MB/s (p-p) 810 MB/s (p-Pb) 3750 MB/s (Pb-Pb)	1000 Hz 800-1000 MB/s	1000 Hz 600 MB/S	10000 Hz 750 MB/s

Resource Needs for 2015+

LHC pledged resources for 2015

	Tier-0	Tier-1	Tier-2
CPU (kHS06)	687	988	1286
Disk (PB)	49	91.5	106.6
Tape (PB)	95.8	172.9	-

The expected needs already include:
- several code improvements
- use of HLT farms

Projections using the current CM, RSG

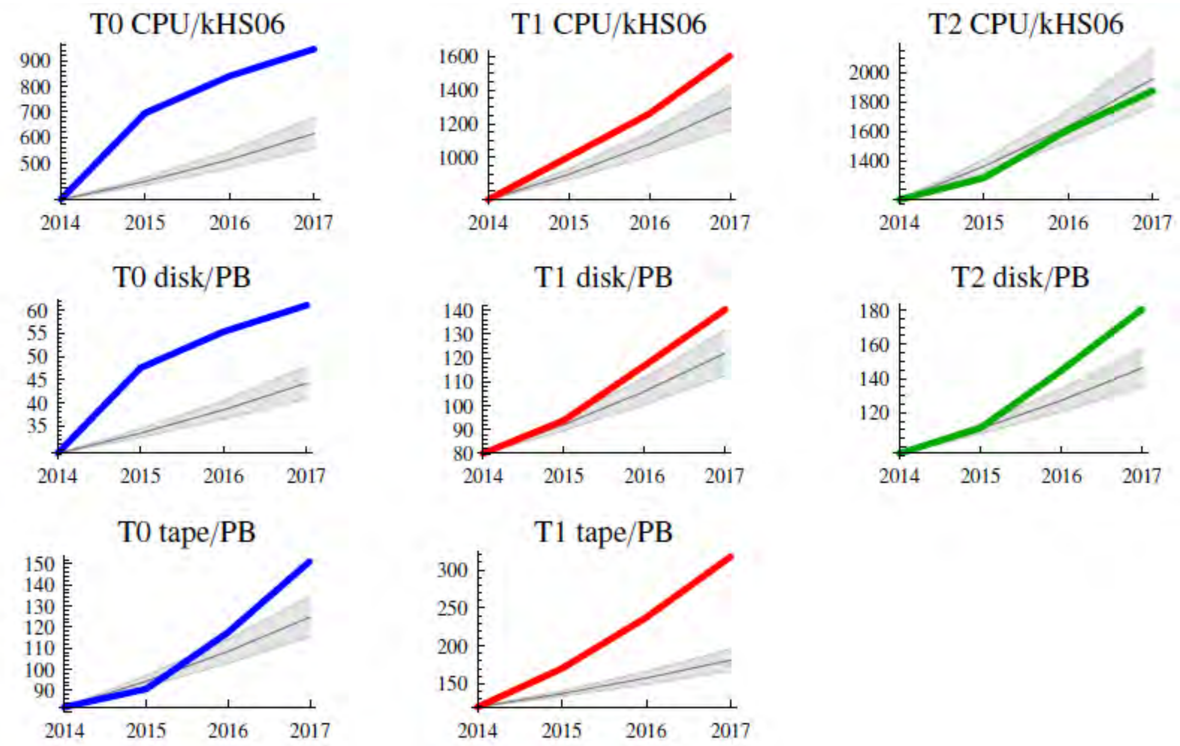

Starting from 2014 pledges

Beyond 2020: back on the envelope calculation

Expected rates: ATLAS & CMS 5-10 kHz LHCb ~20 kHz

Assumptions:

CPU:

- scales linearly with trigger rate
- more than linearly with event complexity
- linearly with data integrated

Putting together all of the above, the expected CPU needs are roughly:

Run2/Run1 ~6

Run3/Run1 ~10

Run4/Run1 >100

T. Boccali, A. De Salvo

Beyond 2020: back on the envelope calculation

Expected rates: ATLAS & CMS 5-10 kHz   LHCb ~20 kHz

Assumptions:

Disk:
- scales linearly with trigger rate
- less than linearly with event complexity
- linearly with data integrated

T. Boccali, A. De Salvo

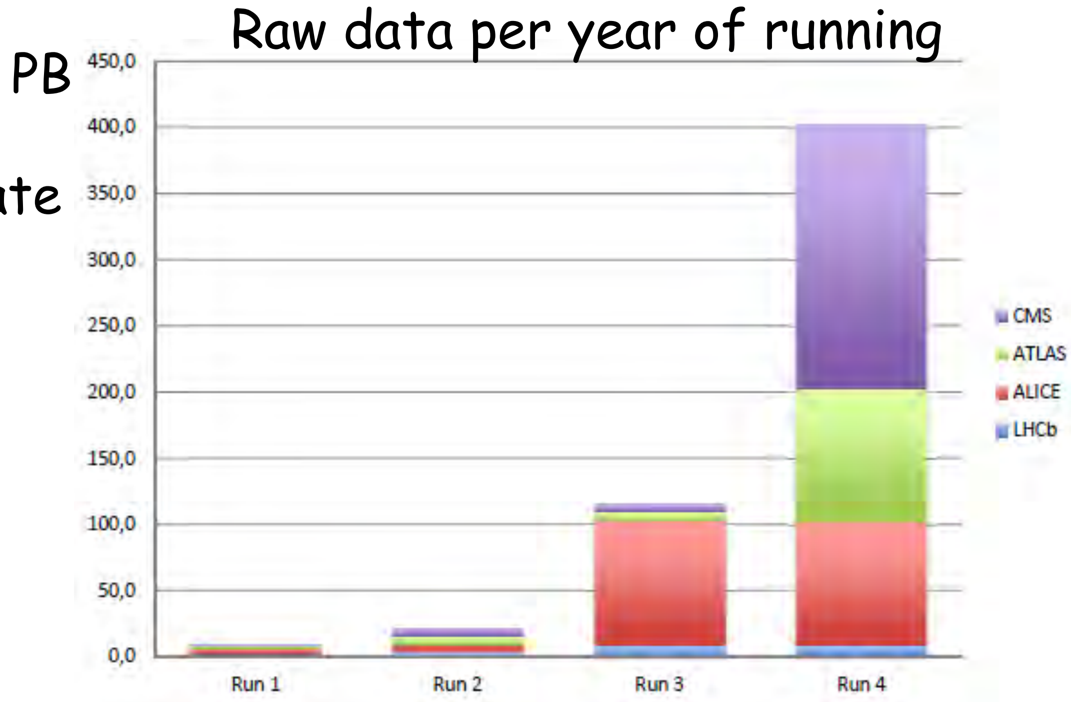

Very rough estimate using a simple extrapolation of current data volume scaled by the output rates. To be added: derived data (ESD, AOD), simulation, user data

The Question

Where to find resources
(with less and less money)?

Technology and Market up to now

Moore's law
The number of transistors on integrated circuits doubles approximately every two years

Kryder's law
The capacity of Hard Drives doubles approximately every two years

↓

Every two years for the same price we have ~2x resources

Nielsen's law
The bandwidth available to users increases by 50% annually

Technology and Market in the future

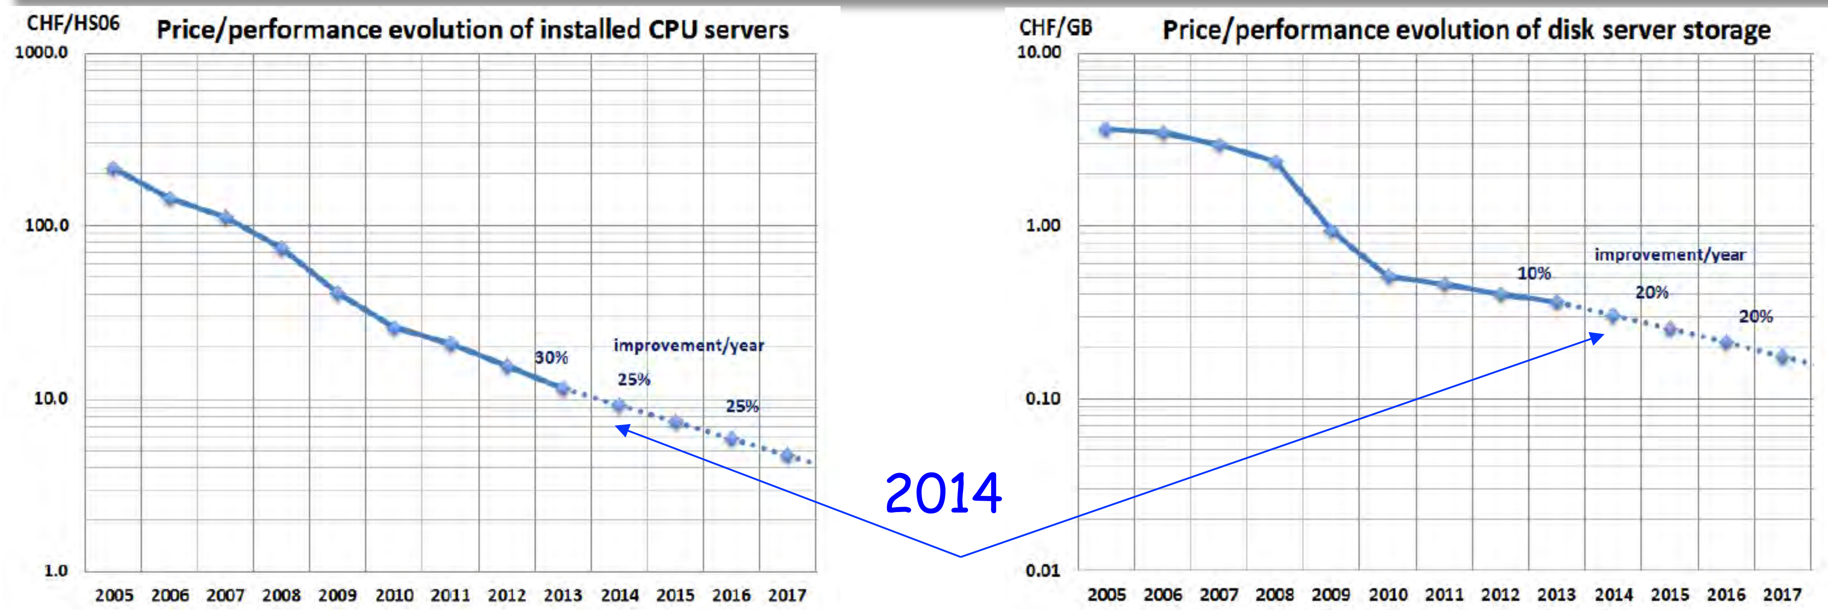

Study of B. Panzer for the CERN Computing Centre

What we can gain: +25%/y for CPU and +20%/y for disk
i.e. a factor 2 over 3 years
Expected factor ~10 in 2025

Not enough!
Even worse, moving to multi-core devices

The (missing) Strategy

From the executive summary of the CM (CERN-LHCC-2014-014)
It is unlikely that simple extrapolations of the present computing models will be affordable, and work must be invested to explore how the models should evolve.

In the mean time use some handles:

1. improve experiment's software and efficiency of the current infrastructure
2. be able to exploit new technologies
3. access new (~free) resources

1. Experiment's Software and Infrastructure

LHC experiments already improved their code:

- ✓ faster algorithms, faster libraries → reduced CPU use per event
- ✓ reduced memory consumption → allow to use any available machine (e.g. High Level Trigger farms)
- ✓ adaptation to changing architectures, discussed later
- ✓ increase the percentage of coordinated activities, less chaotic user analysis
- ✓ distributed Tiers:
    - ✓ exploit resources in different centers (ex. distributed Tier-0)
    - ✓ usage of Tier-1 and Tier-2 for any kind of jobs, e.g. exploit Tier-2s during period of partial inactivity for Tier-1-like jobs

2. Usage of New Technologies

Market driven by commercial products

Videogames: based on GPU

~exponential computational increase with GPU

Tablets and Smartphones:
low consumption and low cost CPU
possible use of low consumption + accelerators (GPU, INTEL, ...)

INTEL: many cores CPU architectures

2. Usage of New Technologies

To exploit efficiently these resources an important investment in new code development is necessary:

- It does not mean that all of the experiment software codes should be re-written, only the part that is convenient to port to a parallel paradigm
- However, this is a major change in the infrastructural code (frameworks)
- and a huge effort in the tools: Geant, root, …

3. Access new (~free) resources

US Example

When Grids Collide

Open Science Grid

XSEDE

Extreme Science and Engineering
Discovery Environment

Overview

- 2012 LHC data collection rates higher than first planned (1000Hz vs. 150Hz)
- Additional data was "parked" to be reduced during 2 year shutdown
- Delays the science from data at the end

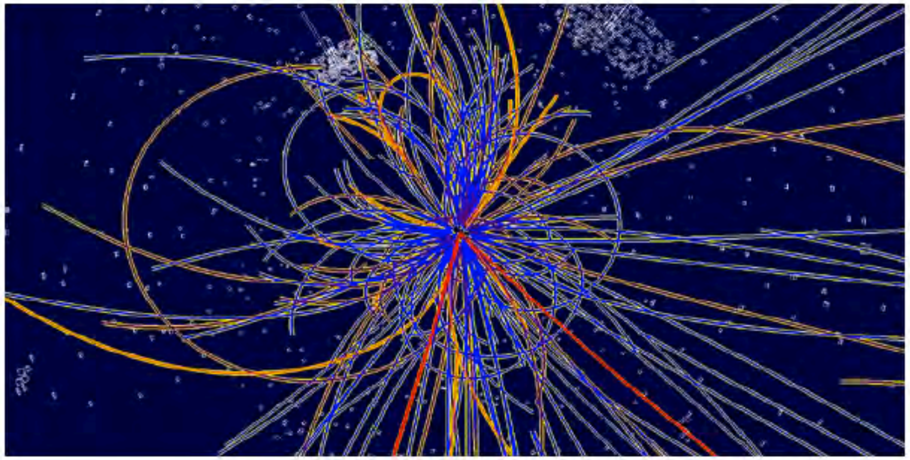

Results

- Work completed in February to March 2013

- 400 million collision events
- 125TB in, ~150 TB out
- ~2 million SUs

- Good experience regarding OSG-XSEDE compatibility

Collaboration wins!

European Initiative: EU-T0

Association of major European research institutes and funding agencies – CERN, CIEMAT-ES, DESY-GE, IFAE-ES, IN2P3-FR, INFN-IT, KIT-GE and STFC-UK – which between them support research projects in disciplines ranging from particle, nuclear, astroparticle physics to cosmology, astrophysics and photon science

EU-T0 will federate several major European e-Infrastructures for research funded by the partners, into a virtual Tier-0 centre, enabling and coordinating development of research computing services for many disciplines of physics, and potentially beyond

European Initiative: EU-T0

The EU-T0 data research and innovation hub will provide:

- Federated computing infrastructure and interoperable services to support research workflows of multidisciplinary projects
- Software services and tools to the research communities
- Data management and data preservation services and solutions
- Governance and coordination functions across the federated centres

European Initiative: EU-T0

Two main calls in preparation for September:

- The EU-T0 data backbone: heterogeneous storage managed in a federated way; interoperable (EUDAT-compatible); fulfilling the real-time ingestion and the archive access requirements

- The EU-T0 cloud: hybrid distributed computing architecture configurable for all users, and serving also the so-called long-tail science

To conclude

- LHC data processing and analysis will require a lot of resources in the next 15 years
- Technologies and e-infrastructures are changing fast
- We need to keep up with evolutions of the large-scale market
- It is very important to get substantial support from EC funds via the EU-T0 initiative