LHC Computing Models and Their Evolution

Borut Paul Kersevan
Jozef Stefan Inst., Slovenia

for the LHC experiments

XLDB Workshop, CERN 2013

Introduction

- LHC delivered **billions** of recorded collision events to the LHC experiments from proton-proton and proton-lead collisions in the Run 1 period (2009-2013).
  - This translates to ~ 100 PB of data recorded at CERN.
  - several 100 PB more storage needed across the **Worldwide LHC Computing Grid** to provide space for archival, replication, simulation and analysis.
- The challenge how to process and analyze the data and produce timely physics results was substantial but in the end resulted in a great success.

Data transferred from CERN across the world for access and processing

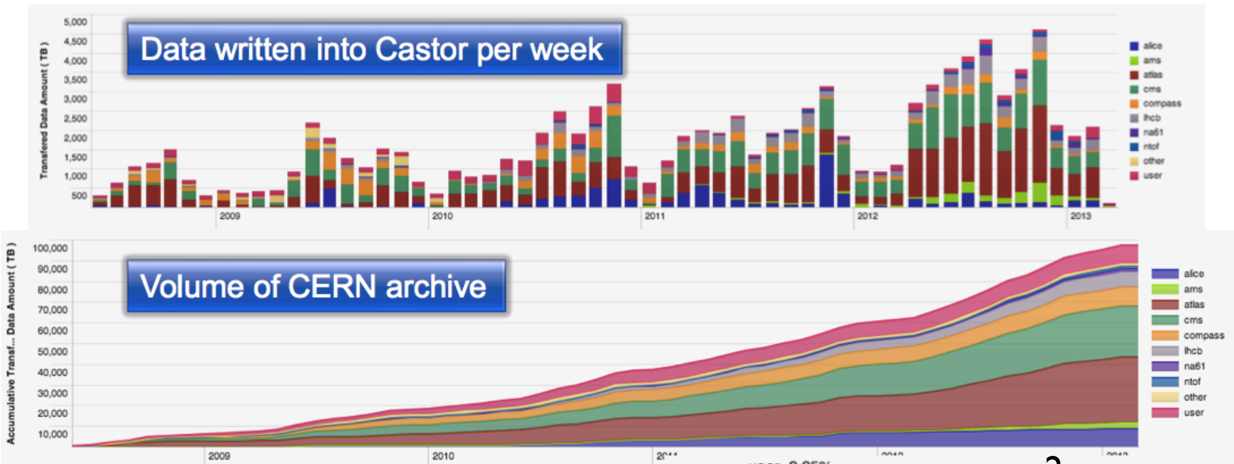

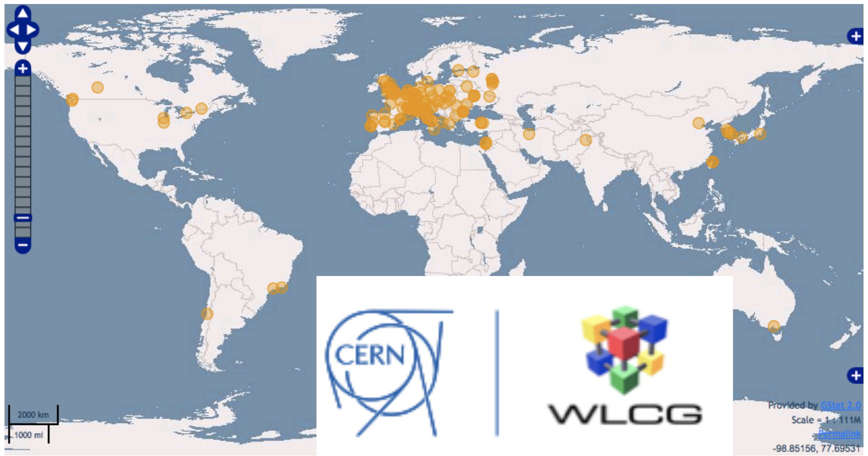

WLCG tiered structure

- The LHC experiments rely on **distributed computing resources:**
  - **WLCG - a global solution, based on the Grid technologies/middleware.**
    - distributing the data for processing, user access, local analysis facilities etc.
    - at time of inception envisaged as the seed for global adoption of the technologies.
  - **Tiered structure:**
    - Tier-0 at CERN: the central facility for data processing and archival,
    - 11 Tier-1s: big computing centers with high quality of service used for most complex/intensive processing operations and archival,
    - ~140 Tier-2s: computing centers across the world used primarily for data analysis and simulation.
  - **WLCG and LHC computing a big success in Run 1!**
    - **Computing was not a limiting factor for the Physics program of the LHC experiments.**
    - **Many thanks to our Grid sites for their excellent performance and contributions!**

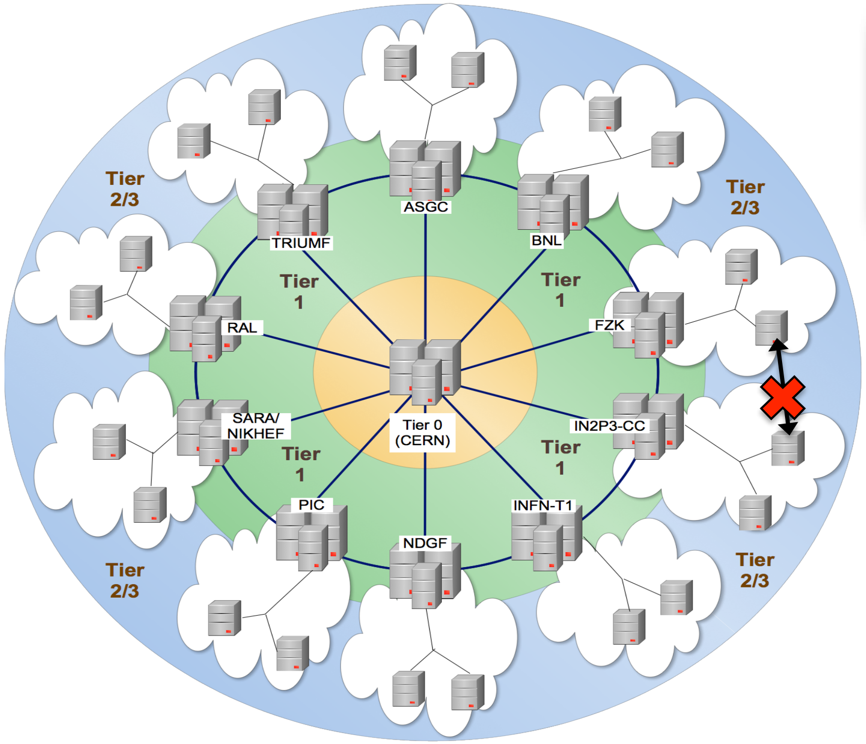

Capacity:

- ~350,000 CPU cores
- ~200 PB of disk space
- ~200 PB of tape space

- **Hierarchical tier organization** based on **MONARC** (**MO**DELS OF **N**ETWORKED **A**NALYSIS AT **R**EGIONAL **C**ENTERS) network topology
- In **ATLAS** sites are grouped into **clouds** for organizational reasons
- Possible communications:
  - Optical Private Network
    - T0-T1
    - T1-T1
  - National networks
    - Intra-cloud T1-T2
- Restricted communications: General public network
  - Inter-cloud T1-T2
  - Inter-cloud T2-T2

Operating the Distributed Computing Environment

- **Distributed computing introduces a new scale w.r.t. a local computing facility in terms of data and job management.**
  - No ‘industrial’ standard or simple rules on what are optimal solutions for data placement and job brokerage to ensure the optimal usage/minimal job latency/..
  - The LHC experiments composed their Computing Models based on best knowledge of the new and evolving system, including their experiment specifics.

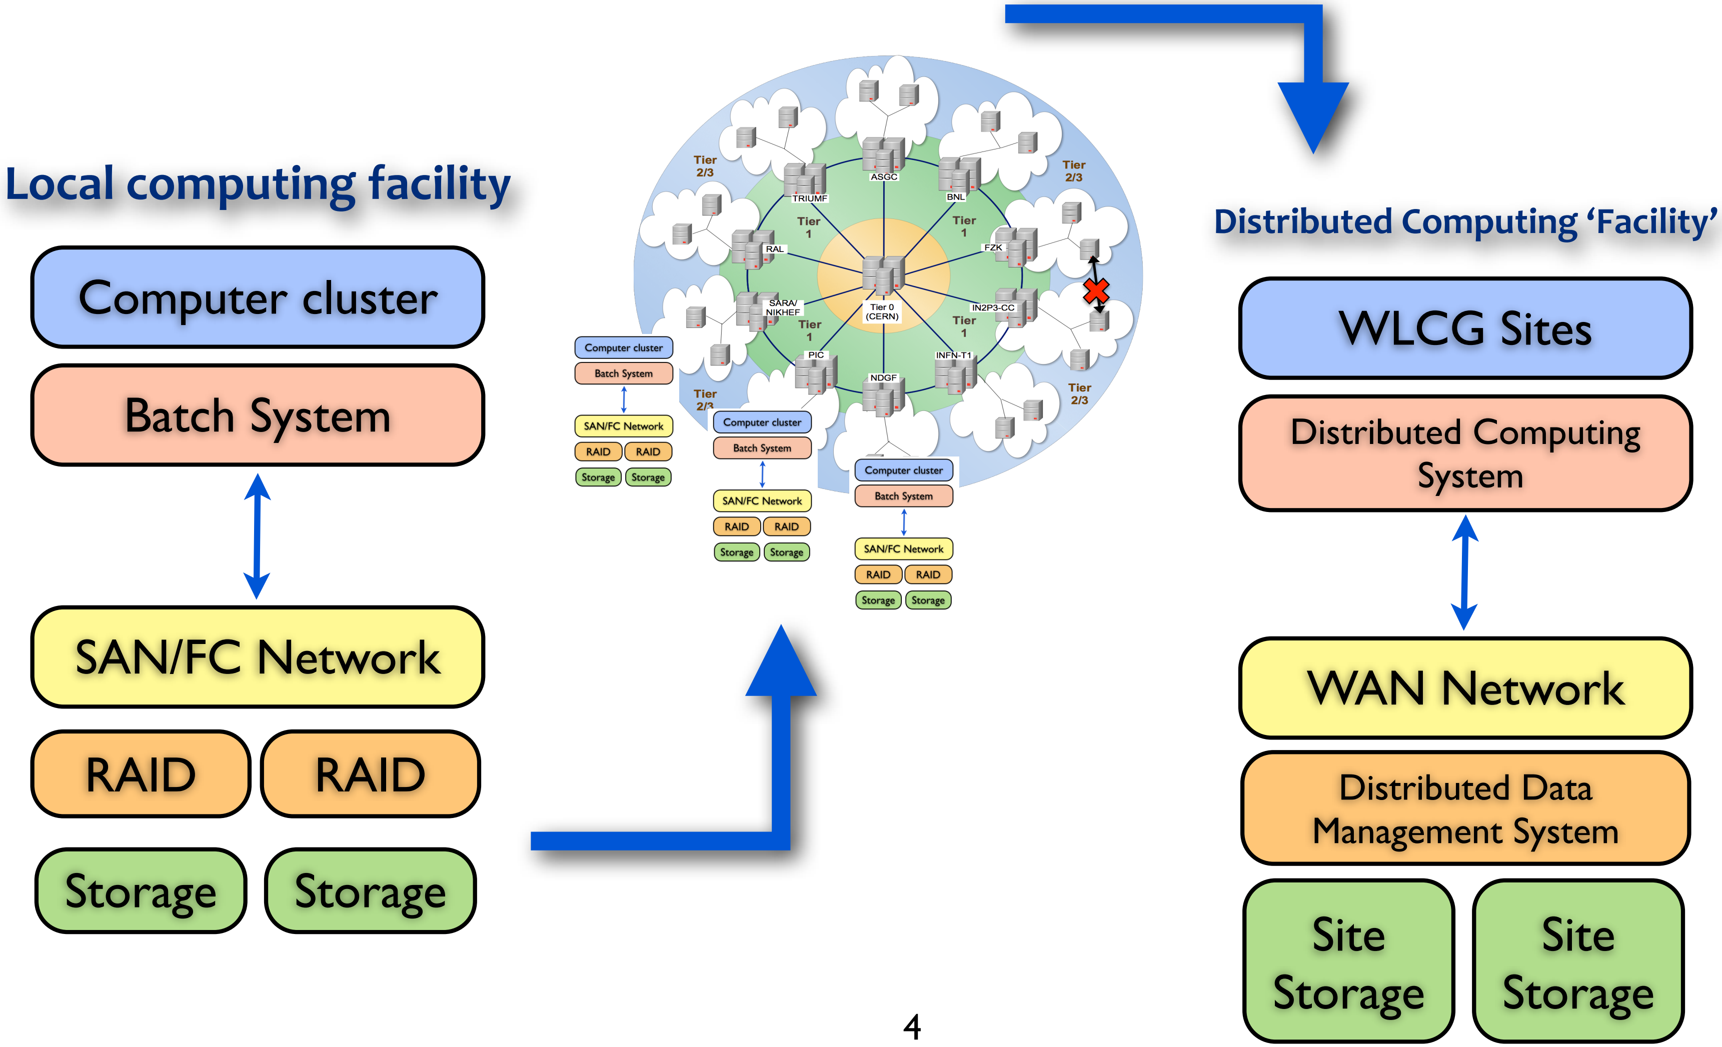

Experiment Computing Models

- The LHC experiment Computing Technical Design Reports were produced in ~2005, with the best knowledge available at the time.
  - **No plan survives the reality, in this case the arrival of the data:**
    - Operational experience introduced significant modifications and improvements in Run 1.
      - E.g. moving away from the Monarc model, all Tiers perform similar activities and pass the data between them.
    - Significant technological evolution until today also impacted (and continues to impact) the optimal operational models:
      - For example network bandwidths increased more than anticipated, one can make better use of storage resources with more dynamic data movement.
  - **'In-house' technical solutions custom-made by experiments to optimize the operation proliferated:**
    - Several solutions to the same problems.
    - Expensive to maintain and develop in terms of manpower.
    - Awareness that searching for common use cases between experiments and global community could boost the activities and economize manpower is becoming crucial in view of the current financial climate.

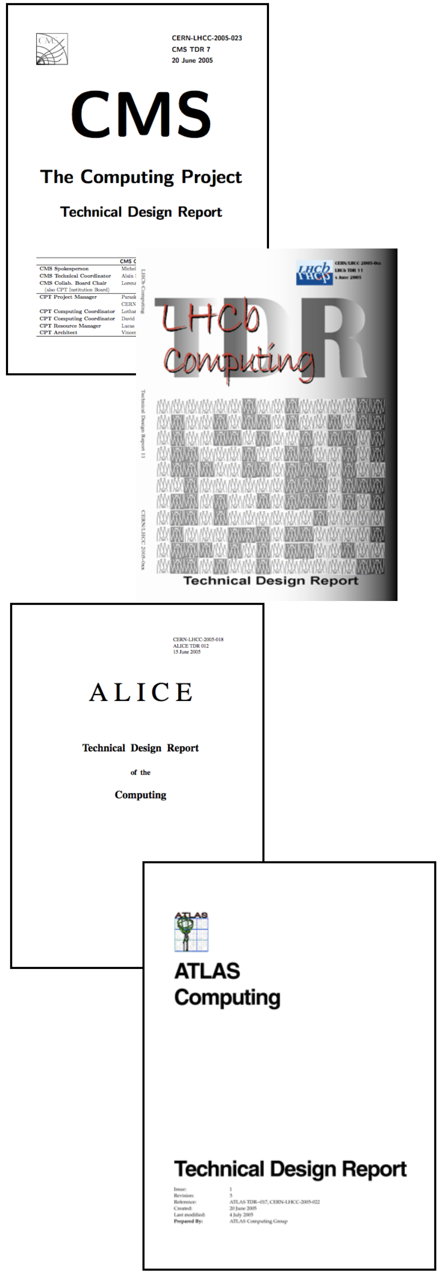

Typical Data Processing Flows

- Generic steps: real data reconstruction, data simulation, data analysis:
  - **Simulation of data the most computationally intensive task!**
    - Simulated data an essential component in final physics analysis! **And a lot of it is needed!**

Real Data Flow (ATLAS)

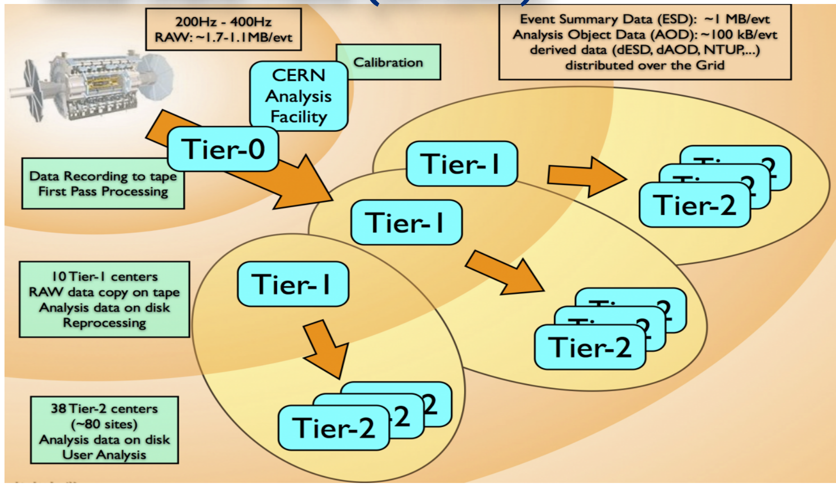

Analysis (ALICE)

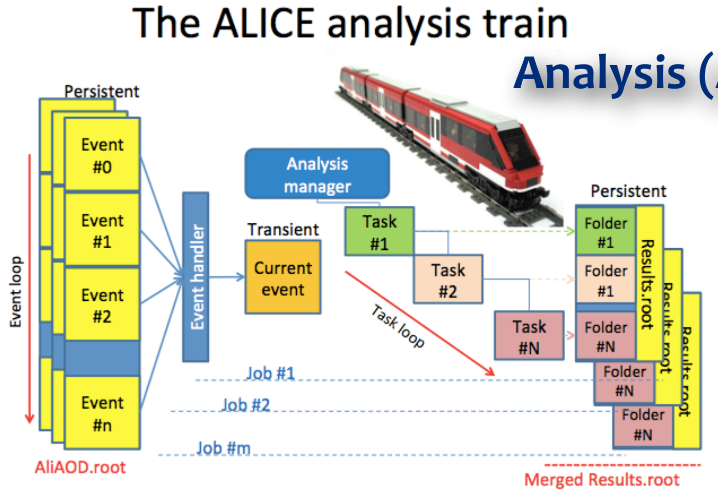

Real data processing (LHCB)

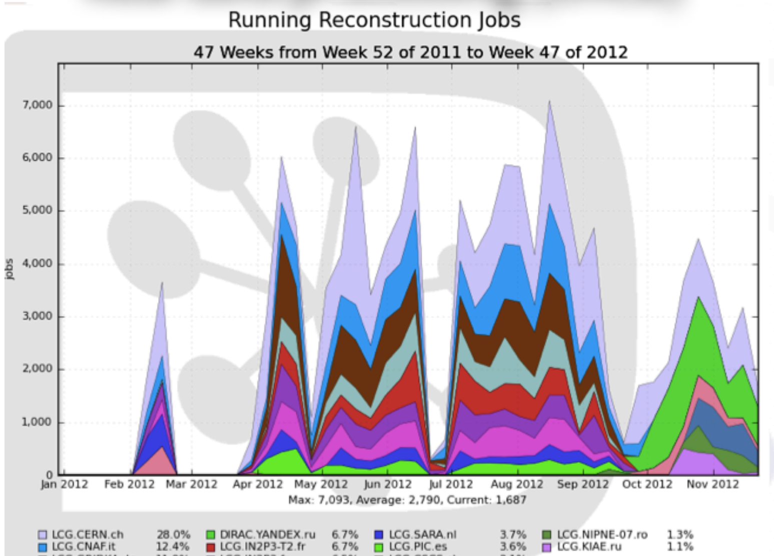

Simulation production (CMS)

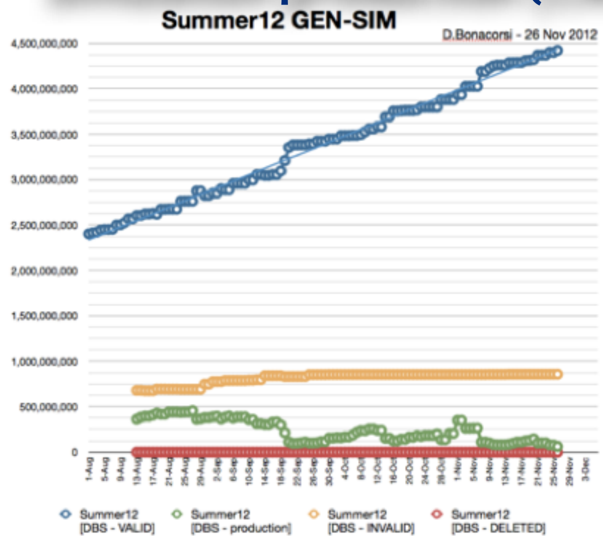

Data Processing and Management Tools

- **At time of inception no global/commercial solution for the distributed computing (Grid middleware developed) needed for our 'Big Data' handling:**
  - Even today, 'Cloud facilities' or HPC centers **are not distributed resources,** neither in terms of CPU nor storage.
    - We could build a Grid of Cloud facilities but not a Cloud of Grid facilities.
- **'In-house' experiment specific topmost job and data management layers, also tying in different Grid and local (batch/storage..) setup flavors.**
  - Distributed Computing Systems (job handling): AliEn(Alice),PanDA(ATLAS),Crab(CMS), Dirac(LHCb)
  - Distributed Data Management Systems (file placement, replication and access handling): AliEn(Alice),DQ2/Rucio(ATLAS), PhedEX(CMS), Dirac(LHCb)
- **Lower layers generally common (WLCG deliverables/Grid middleware):**
  - Storage (Castor, EOS, dCache,DPM..),
  - File transfer services (CERN FTS2 and FTS3),
  - File/access catalogues (LFC),
  - Virtual machines and remote filesystems for software access (CERNVM, CVMFS),
  - Database caching (Frontier/Squid for ORACLE DB access),
  - Monitoring tools (SAM, PerfSonar, DashBoard),
  - Information infrastructure (BDII).

Processing millions of jobs and PB of data weekly

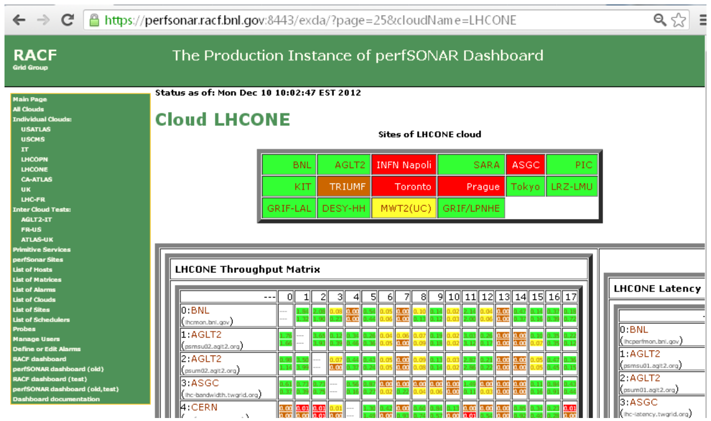

... And we are successful!

- Successfully processing millions of jobs and hundreds of PB of data!
- .. there is of course always room for improvement and we are working on it.

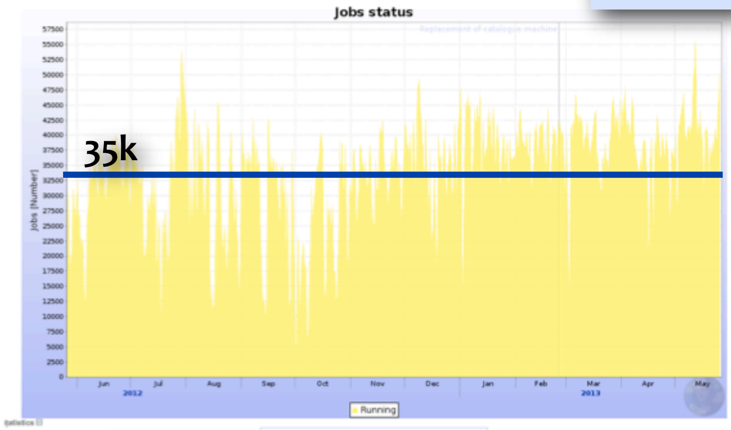

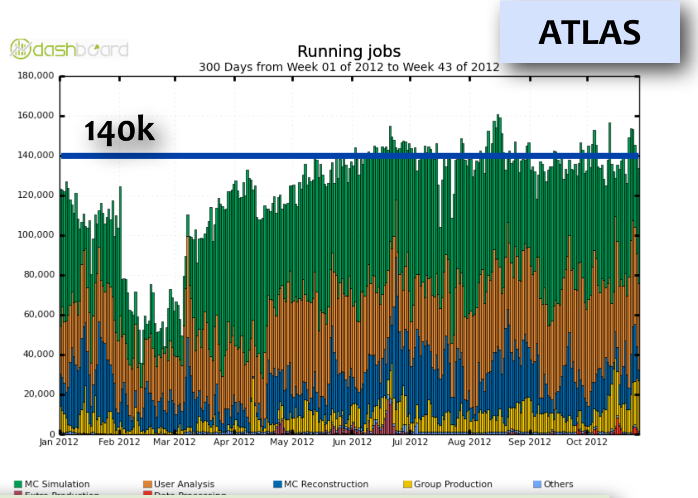

Ratio in WLCG Tier1+Tier2 CPU in 2012: ATLAS(40%):CMS(30%):ALICE(20%):LHCb(10%)

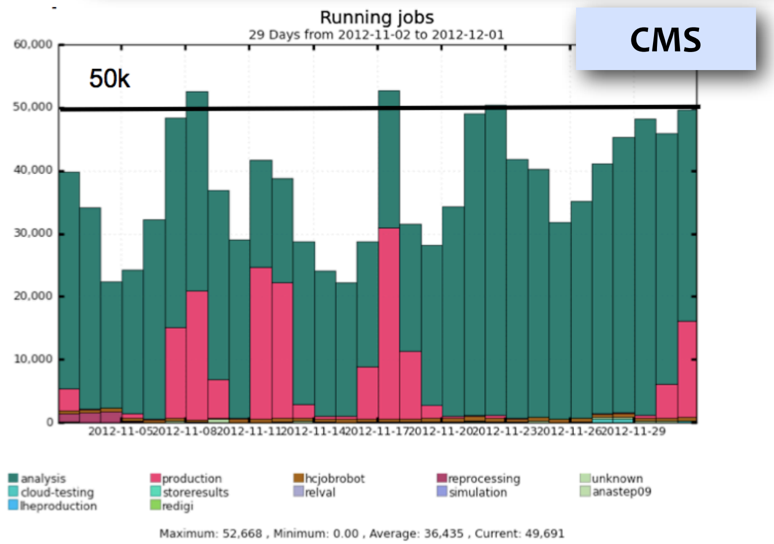

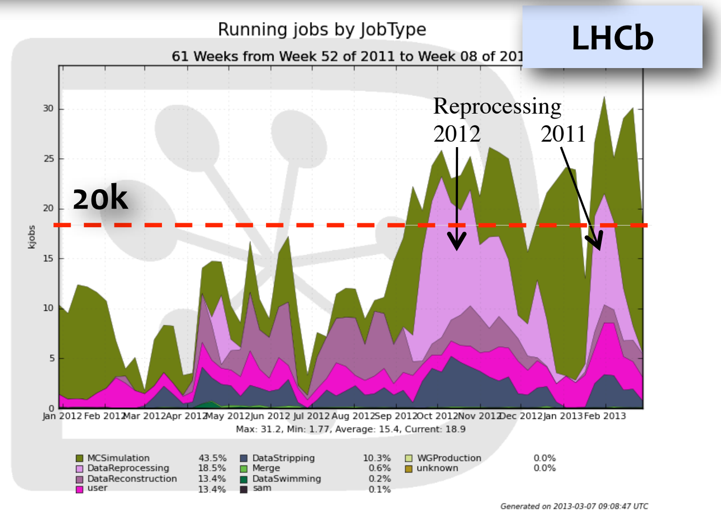

Processing the Experiment Big Data

- The simplest solution in processing our data is **using the data affinity** for the jobs:
  - data is staged to the site where the compute resources are located and data access by analysis code is from local, site-resident, storage.
- **However:**
  - In our distributed computing environment we do not have enough disk space to host all our data on every WLCG site:
    - Thus we distribute (pre-place) our data across our sites.
  - The popularity of certain data sets is very hard to judge in advance:
    - Thus the computing capacity at a site might not match the demand for certain data sets.
- **Different approaches are being implemented:**
  - **Dynamic or on-demand data replication:**
    - **Dynamic:** If certain data is popular on one site (i.e. processed often), make additional copies on sites with spare CPU (and disk) capacity. (e.g. ATLAS PD2P service).
    - **On-demand:** The popular data on one site can be copied locally by jobs in another site (XRootD, HTTP federations).
  - **Remote access:**
    - The popular data on one site can be remotely accessed from jobs in another site (XRootD, HTTP federations).
  - **Both approaches have the underlying scenario that puts the WAN between the data and the executing analysis code.**
    - Inserting the WAN is a change that potentially requires special measures to ensure the smooth flow of data between disk and computing system, and therefore the "smooth" job execution needed to make effective use of the compute resources.

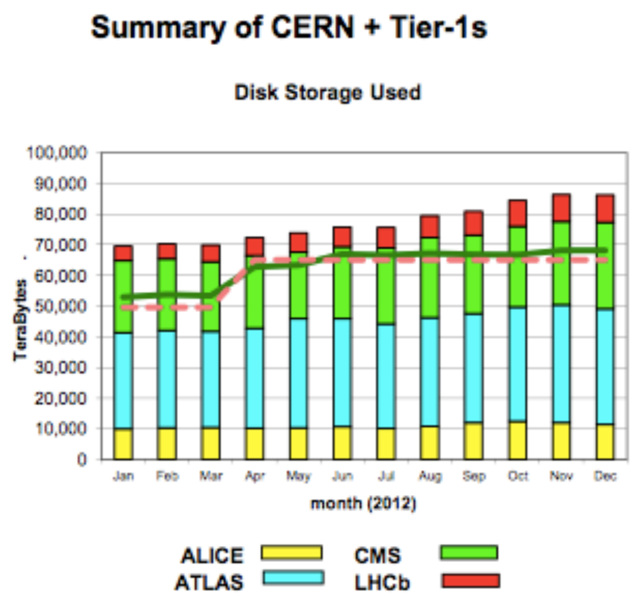

Data transfers in Run 1 (ATLAS)

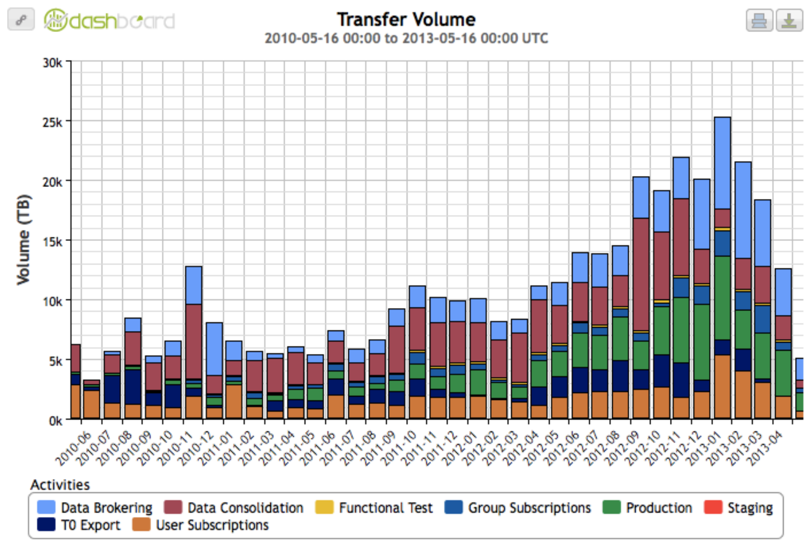

Note on Networking

- Networking is obviously an essential ingredient in our distributed computing models.
- Is this something to worry about?
  - Maybe not, millions of Netflix users generate much more traffic than HEP users do
    - If it works for them it must work for us too!
  - Maybe yes, because Netflix users don't compare well with us
  - Commercial Internet providers optimize their infrastructure for mainstream clients and not for the specific needs of the HEP community, e.g.
    - Traffic pattern characterized by small flows (tiny "transactions").
    - A few lost packets at 10 Gbps cause 80-fold throughput drop.
    - Connectivity issues between NRENs and "Commercials".
    - Availability/Reliability issues.
- Networking usage increases with time due to:
  - Bigger real/simulated data volume to transfer.
  - Processing activities becoming possible on all Tiers (e.g. data reprocessing on Tier-2s).
  - Trading off less pre-placed data for network transfers to optimize/reduce the disk usage.

• ATLAS: Transfer rate reached 5-10 GB/s permanently
→ ATLAS adapts to available network infrastructure
• Many different activities competing for network resources

Content Delivery Networks: Compare to Netflix (I.Fisk)

- HEP problem harder than Netflix?
  - ★ Netflix delivers streaming video content to about 20M subscribers
  - ★ Routinely quoted as the single largest user of bandwidth in the US
    - More than 30% of the traffic
- HEP: <# clients, < distribution, > bandwidth per client
- However, **much** larger data set
  - ★ HEP can't make many multiple static copies
  - ★ —need different strategies instead:
  - ★ make **dynamic** replicas and clean up when no longer useful
  - ★ access data directly over the wide area networks

	NETFLIX	LHC Computing
Bandwidth per client	1.5Mbit/sec	1MByte/sec
Clients	1M*	100k cores
Serving	1.5Tbits	0.8Tbits
Total Data Distributed	12TB	20PB
Annual Budget	>$4B	< $.04B

Similar Problems: Not all files are equally accessed

e.g. Forward Physics ;-)

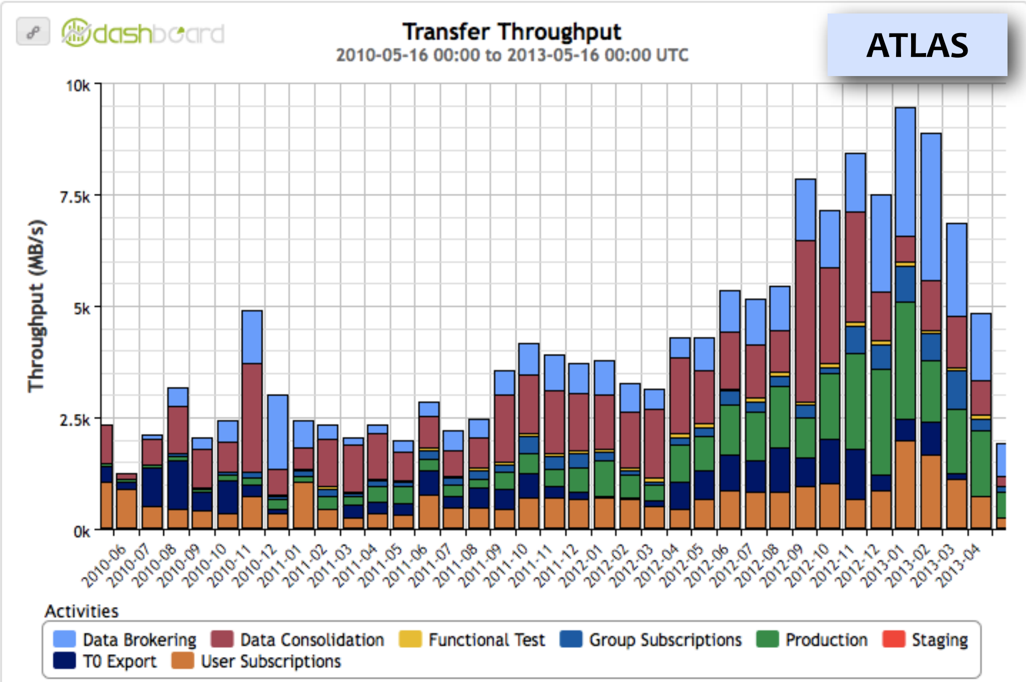

Looking for Opportunistic Resources

- All our resource planning is done using **average** CPU (and disk) consumption rates.
  - And we are **using all available resources all the time** for diverse experiment activities.
  - The experiment analysis activities **peak** before big conference periods, can lead to congestions and backlogs in all the Grid demands.
- **Several venues to explore during Long Shutdown 1 (2013-2014):**
  - **Optimizing/changing our workflows, both in analysis and on the grid.**
    - It will necessarily involve also a change in the ways people analyze the data!
  - **Finding opportunistic resources:**
    - High Performance Computing centres have a lot of CPU available, we could use the available idle cycles for (a subset of) our activities,
      - e.g. MC event generation, possibly simulation.
    - Cloud resources: Again, for a subset of our activities, similar to HPC
      - If we are really hard pressed, even use commercial resources (?)
        - Exploring setups with Amazon EC2, Google Cloud ...
    - Opportunistic offers of big computing centers:
      - The experiments need to be able to simply and quickly integrate such resources into their distributed computing environment.
    - Volunteer computing resources: exploiting virtualization (CernVM), BOINC..
  - **Looking for solutions to speed up our code and accommodate our needs.**
- **A lot of activity foreseen in the experiment Software & Computing during LS1 to tackle this.**

Resource Projections

- **The LHC experiments have even more ambitions Physics plans for Run 2!**
  - with a more demanding LHC environment: higher energy, more complex collisions...
  - We plan to record more data: Physics motivated, the measurements will become even more involved..
- **The demands in computing resources to accommodate the Run 2 physics needs increase:**
  - The experiments need to show that they exercise due diligence in evolving their Computing models, software and operational models to optimally use the required resources.
  - **We are working on further improving both our software and distributed computing.**

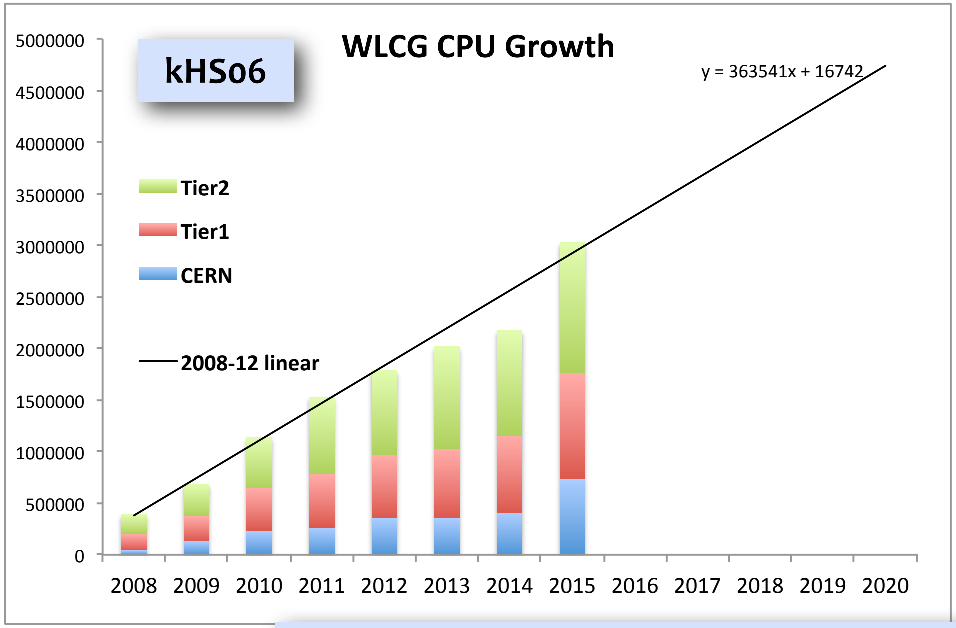

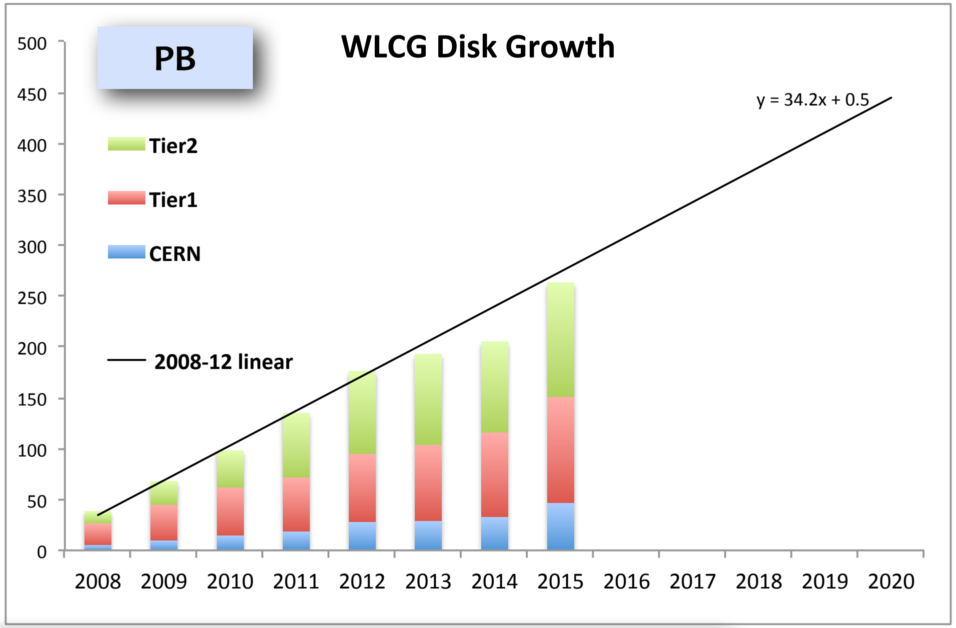

Resource growth (mostly) due to technological progress with (almost) constant budgets

LHC Upgrade Timeline - the Challenge to Computing Repeats periodically!

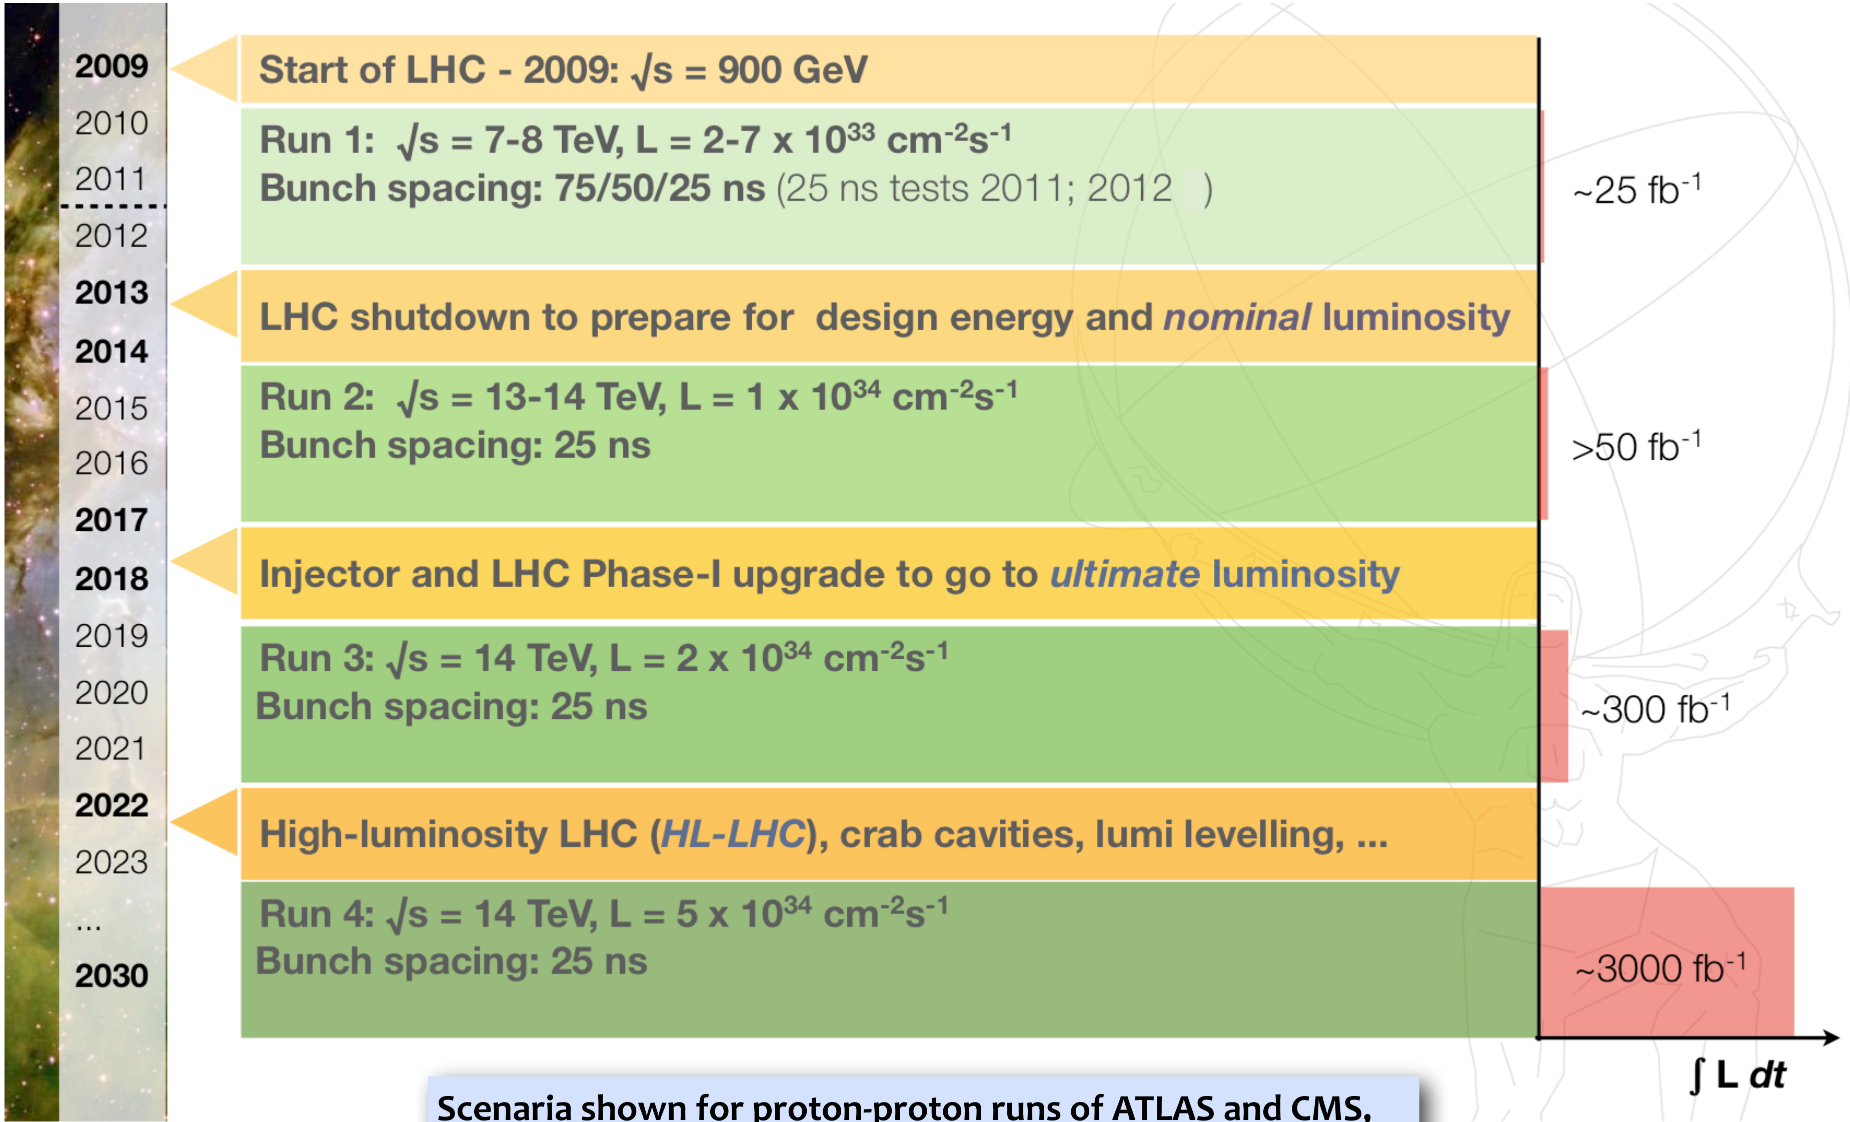

Scenaria shown for proton-proton runs of ATLAS and CMS, LHCb and Alice follow different strategies.

Summary

- The LHC experiment Distributed Computing and software performance was a great success in LHC Run 1.
- **Growing investment in software development required during LS1(from reconstruction to distributed computing) is needed, if the computing is not to be the limiting factor in Run 2 physics results throughput:**
  - An ambitious plan is taking shape for the S&C LS1 activities.
    - We **need** to do it to meet the Run 2 challenges.
    - try exploit additional resources, e.g. in high-performance computing centres.
    - also seek manpower savings in Tiers operation.
  - Evolution: key points:
    - Need to demonstrate that we are doing as much as possible to make best use of resources available!
    - **Software:**
      - Parallelism, new architectures, etc – significant challenges and need of expertise.
      - Requires some investment.
    - **Commonality:**
      - Between experiments,
      - With other sciences!
    - **Simplicity:**
      - Reduce complexity where possible:
        - Grid services,
        - Deployments (e.g. a central service is simpler).
    - **Focus our efforts:**
      - Where we must:
        - e.g. data management tools..
      - Cannot afford (nor should we) do everything ourselves!
    - Collaborate : To bring in expertise and to share ours!