# News from LCG-France and WLCG

Catherine Biscarat, LPSC/IN2P3/CNRS
Technical coordinator of LCG-France

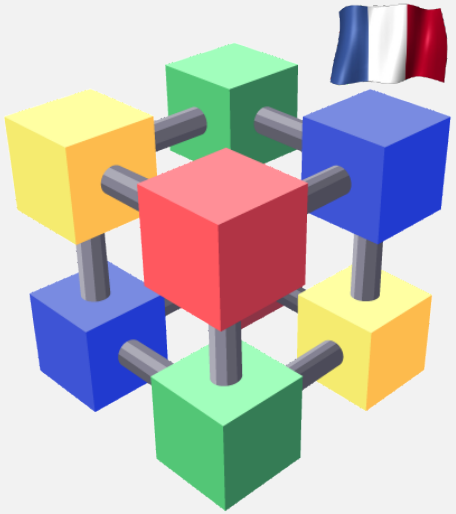

LCG France

"FR-cloud, Regional Centers Meeting",
11-12 April 2018, Hong Kong

# My very first time in Hong Kong

My fourth participation in this workshop.
Fruitful meetings and discussions

Thanks for the invitation !

# In this talk

- LCG-France
  - Organisation
  - Sites
  - Pledges
  - Operations in brief
- WLCG
  - The challenges ahead of us
  - How WLCG get organised

# LCG-France

# LCG-France

- LCG-France is a project to organise the French contribution to WLCG
  - Project launch in 2004
  - Funding agencies: CNRS/IN2P3 and CEA/DSM
  - Parties: sites and experiments
- Goal: to contribute at the 10% level to the computing of the LHC experiments (MoU, T1+T2)

- Organisation of the project
  - Scientific leader:L. Duflot
  - Technical leader:  CB
  - Close collaboration with CC-IN2P3 (E. Fede) and experiments (L. Poggioli for ATLAS)
  - The Management Board
  - The Executive Board
  - Technical meetings
  - Workshops

http://www.lhc-france.fr/

**Our partners and friends**

National Grid Initiative | LHC computing | Sites connectivity

CC-telecom team

# People

The key of the project

Last LCG-France workshop, November 2017, Clermont-Ferrand (LPC)

# WLCG sites in France

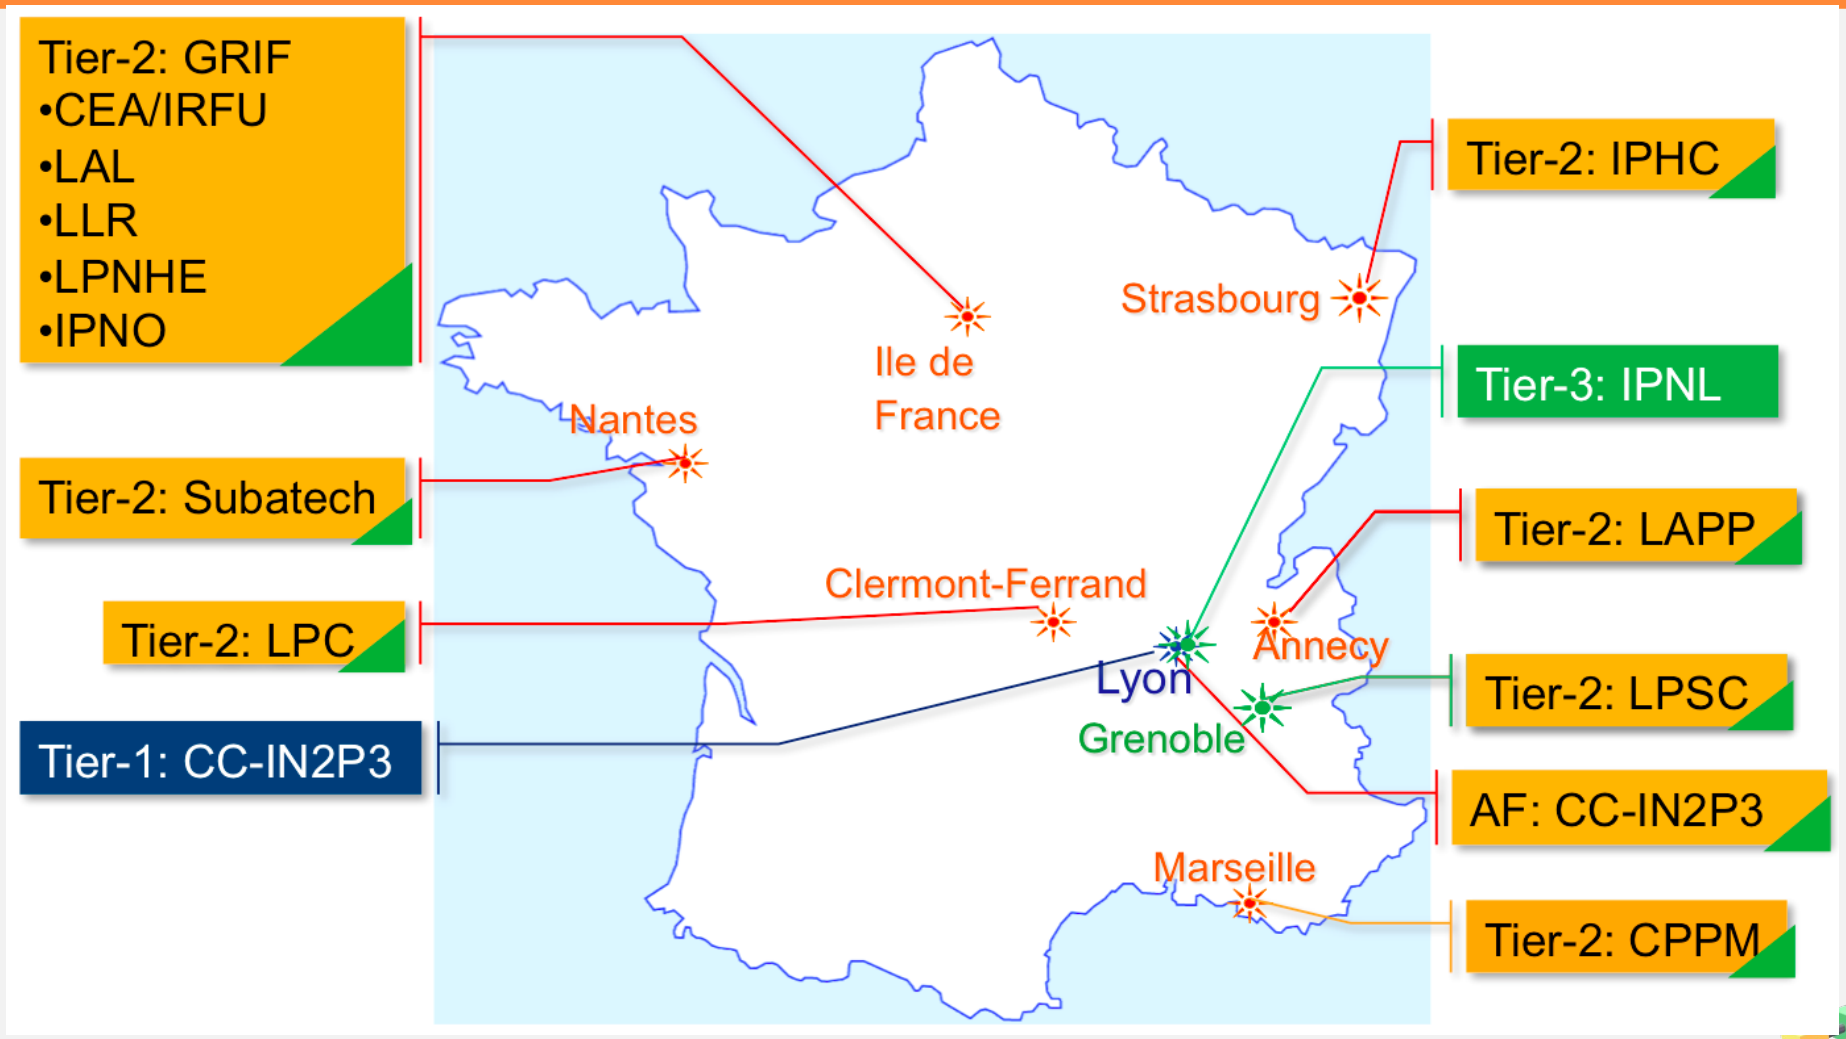

# WLCG sites in France

**All sites are**

- Giving support to the experiments at their labs;
- CC-IN2P3 serves the IN2P3 community (all 4 exp)
- Open to non-LHC VO (except one)
- EGI sites (certified) attached to the french NGI

**The French NGI: France Grilles**

- Five sites are not be IN2P3
- France Grilles operates basic MW deployment (EGI specifics), monitoring and security
- Large move of non-LCG sites to cloud
  - To serve tail of sciences

| Role | Site | ALICE | ATLAS | CMS | LHCb |
|---|---|---|---|---|---|
| Tier-1 | IN2P3-CC | ✓ | ✓ | ✓ | ✓ |
| Tier-2 | IN2P3-CC-T2 (AF) | | | ✓ | |
| | IN2P3-CPPM | | ✓ | | ✓ |
| | GRIF | ✓ | ✓ | ✓ | ✓ |
| | IN2P3-LPC | ✓ | ✓ | | ✓ |
| | IN2P3-IPHC | ✓ | | ✓ | |
| | IN2P3-LAPP | | ✓ | | ✓ |
| | IN2P3-LPSC | ✓ | ✓ | | |
| | IN2P3-SUBATECH | ✓ | | | |
| Tier-3 | IN2P3-IPNL | ✓ | | ✓ | |

# French landscape

These changes did not impact (yet) neither our organisation nor our fundings.

A new research infrastructure “IR T2” is being thought. It would put together HPC, cloud and grid sites. The goal is to get the best of each.

This project could be seen as a similar to the EU project “EOSC” (interconnect).

Shown in Beijing in 2016

## Computing landscape in France

- Most LCG sites are affiliated to a University & the CNRS
- New directives from CNRS and the French Ministry
- They are on the same line
  - Large numbers of small data centres
  - Waste in resources (room, fluid and manpower cost)
  - Promote a few data centre per region
- CNRS
  - The CC-IN2P3 has been classified “national centre”,
    - To promote a synergy between the CNRS HPC centre (IDRIS) and the CC-IN2P3
  - A new body “compute and data” is in charge of this reorganisation (created Jan. 2016)
- May have an impact on our funding scheme and site organisation

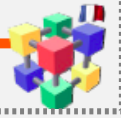

LCG-France, biscarat@lpsc.in2p3.fr FR-Cloud Regional Centers meeting, 7-8 Apr. 2016, IHEP Beijing 15

# LCG-France budget

- LCG-France is a national project with fundings coming from CEA and IN2P3

- The IN2P3 budget is invested
  - mostly at the T1 (renewal and growth) – goal is to be at the 8-10% level of WLCG
  - to help T2 to maintain their resources (renewal of 70% of the capacities)
  - to promote meetings and travelling (LCG-France and WLCG workshops, …)

- This year: the IN2P3 directorate, the labs and the french representatives of the LHC experiments have signed a MoU for 2018-2022
  - Same level of contribution to the experiments as previous years
    - According to physics/detector support
    - At the T1: 45% ATLAS, 25% CMS, 15% ALICE and 15% LHCb
    - Also globally true at T2
  - A flat budget is expected

# LCG-France – pledged resources

- The Tier-2s pledge (CPU and disk) about the same amount of capacities as the Tier-1.
  - Bring additional fundings (regional).
- Additional capacities are offered on top on pledges (mostly at Tier-2s).

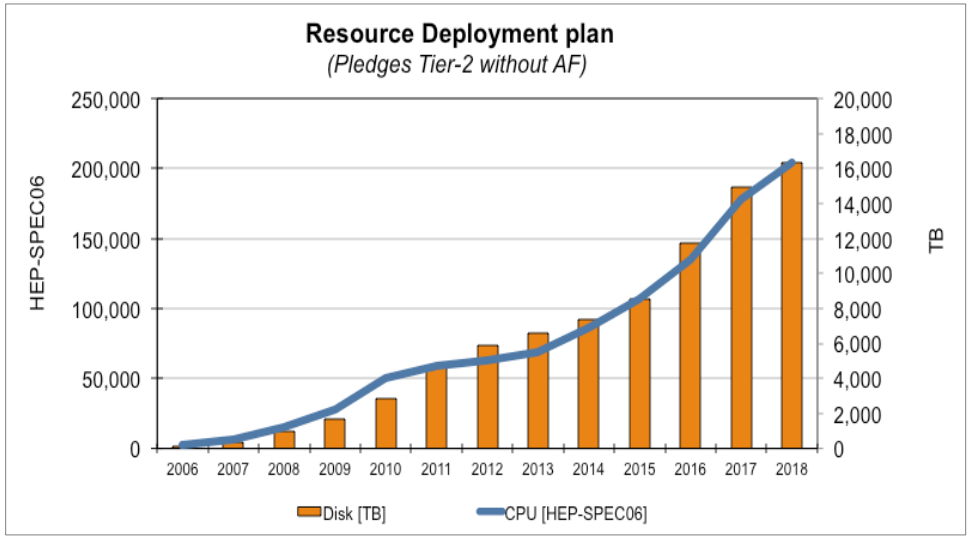


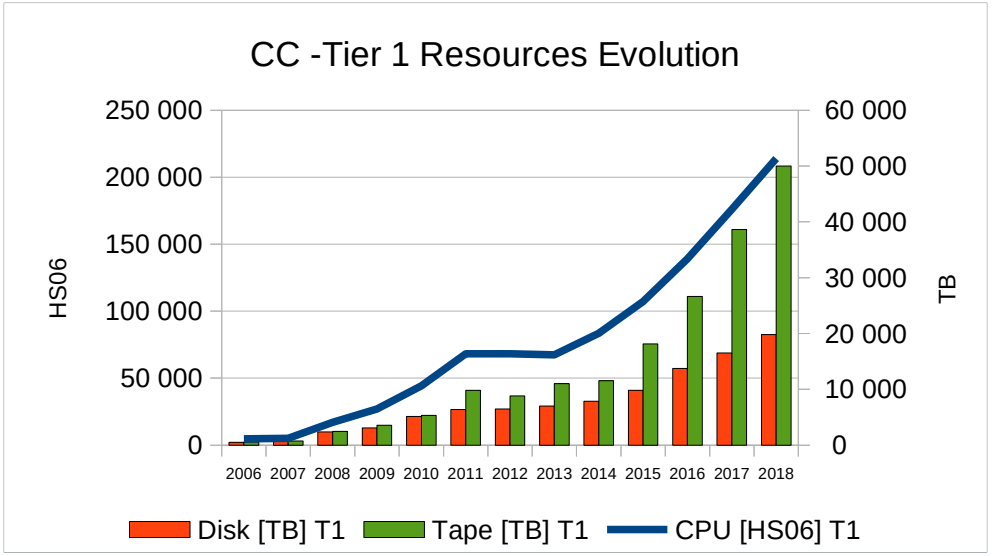

# LCG-France – share of LCG-France

- The amount of resources offered by LCG-France fulfil 7-8% of the required resources at Tier-2s and 9-10% at Tier-1s (depending on the media).
- Broken by experiments part of ALICE, ATLAS, CMS and LHCb requirements:
  - Tier-2s: 9%, 8%, 5%, 12%  (2 sites are T2D LHCb)
  - Tier-1: 11%, 11%, 8%, 12% (given by the distribution of the budget)

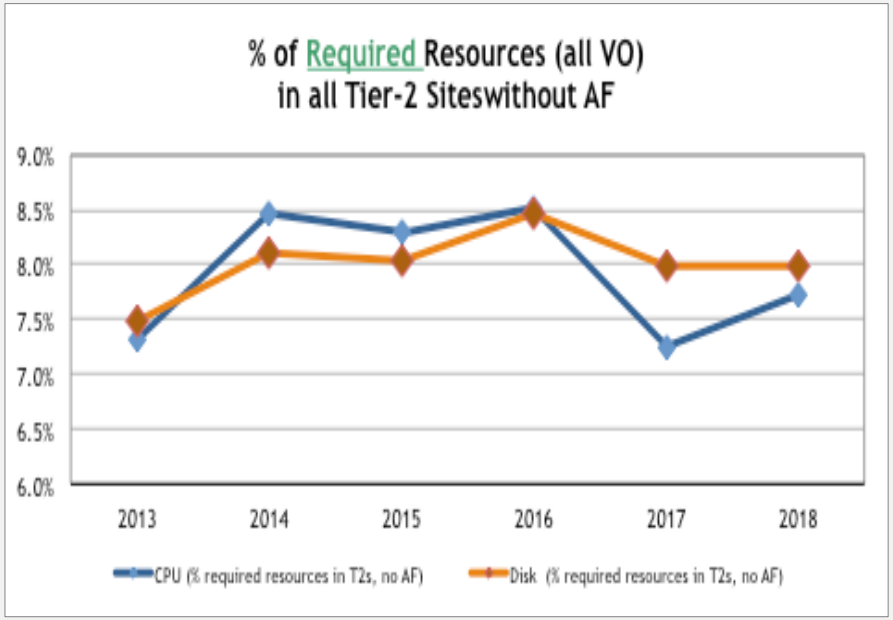


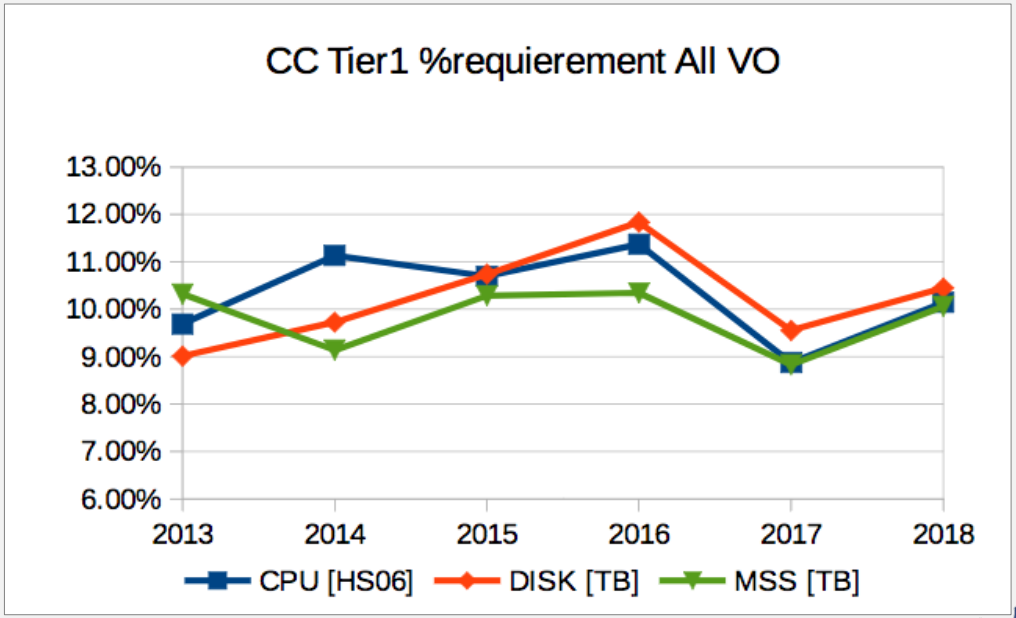

# Evolution of experiment requests

- 2016-2017: exceptional requests du to the good performance of the LHC
- 2017-2018: relatively small requests compared to previous year
- 2018-2019: for Tier-1s and Tier-2s requests are below 20%, except for ATLAS disk
- 2019 has to be finalised at the coming C-RRB on April 24th 2018 (https://indico.cern.ch/event/707294/)

- Flat budget: CERN's model 20%/15%/15% growth per year on CPU/disk/tape

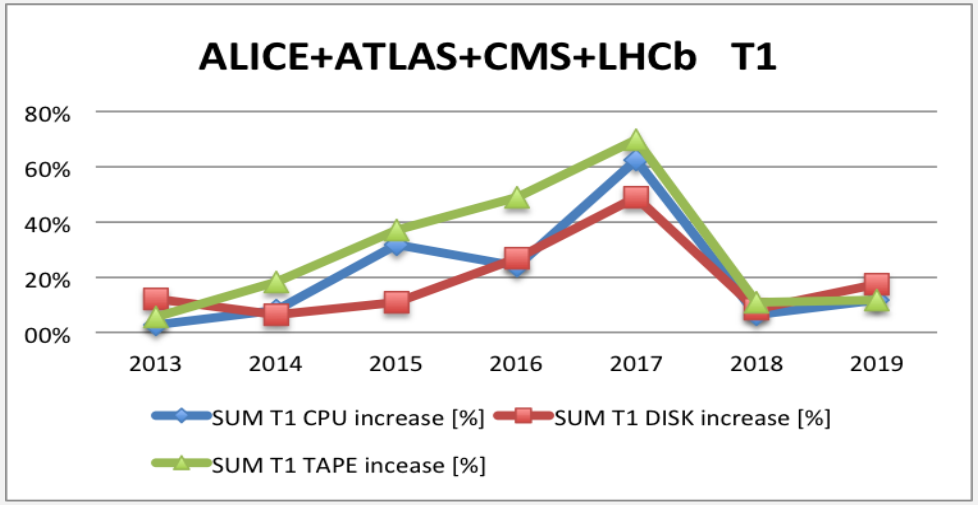


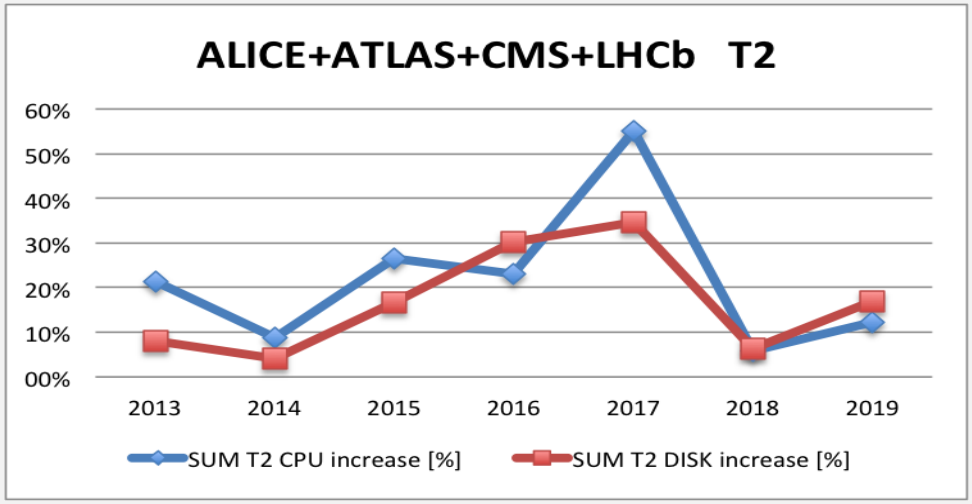


*Details per experiment in back-up slides*

As of November 20017 (autumn C-RRB)

# Operations in France



# Operations

- FR sites are doing well (A/R, resources, exp. needs)
- Useful links and meetings (with minutes):
  - WLCG Monthly Operation coordination meetings:
    https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOpsCoordination
  - WLCG Operation portal
    https://wlcg-ops.web.cern.ch/
  - Grid deployment board (with pre-GDB):
    https://twiki.cern.ch/twiki/bin/view/LCG/WLCGGDBDocs
- Hot topics:
  - EL7 migration/containers
  - IPv6 (J. Bernier yesterday)

We have discussed this yesterday
(site report session)
Five slides in back-up

Ops coordination

- Highlights
- Agenda
- Attendance
- Operations News
- Middleware News
  - Discussion
- Tier 0 News
- Tier 1 Feedback
- Tier 2 Feedback
- Experiments Reports
  - ALICE
  - ATLAS
  - CMS
    - Discussion
  - LHCb
- Ongoing Task Forces and Working Groups
  - Accounting TF
  - Archival Storage WG
  - Information System Evolution TF
  - IPv6 Validation and Deployment TF
  - Machine/Job Features TF
  - Monitoring
  - MW Readiness WG
  - Network Throughput WG
  - Squid Monitoring and HTTP Proxy Discovery TFs
  - Traceability WG
  - Container WG
- Special topics
  - CNAF recovery update
  - SAM recalculation policy draft
  - Tape metrics in WLCG Storage Space Accounting prototype
- Action list
  - Specific actions for experiments
  - Specific actions for sites
- AOB

# Ahead of us



# LHC schedule

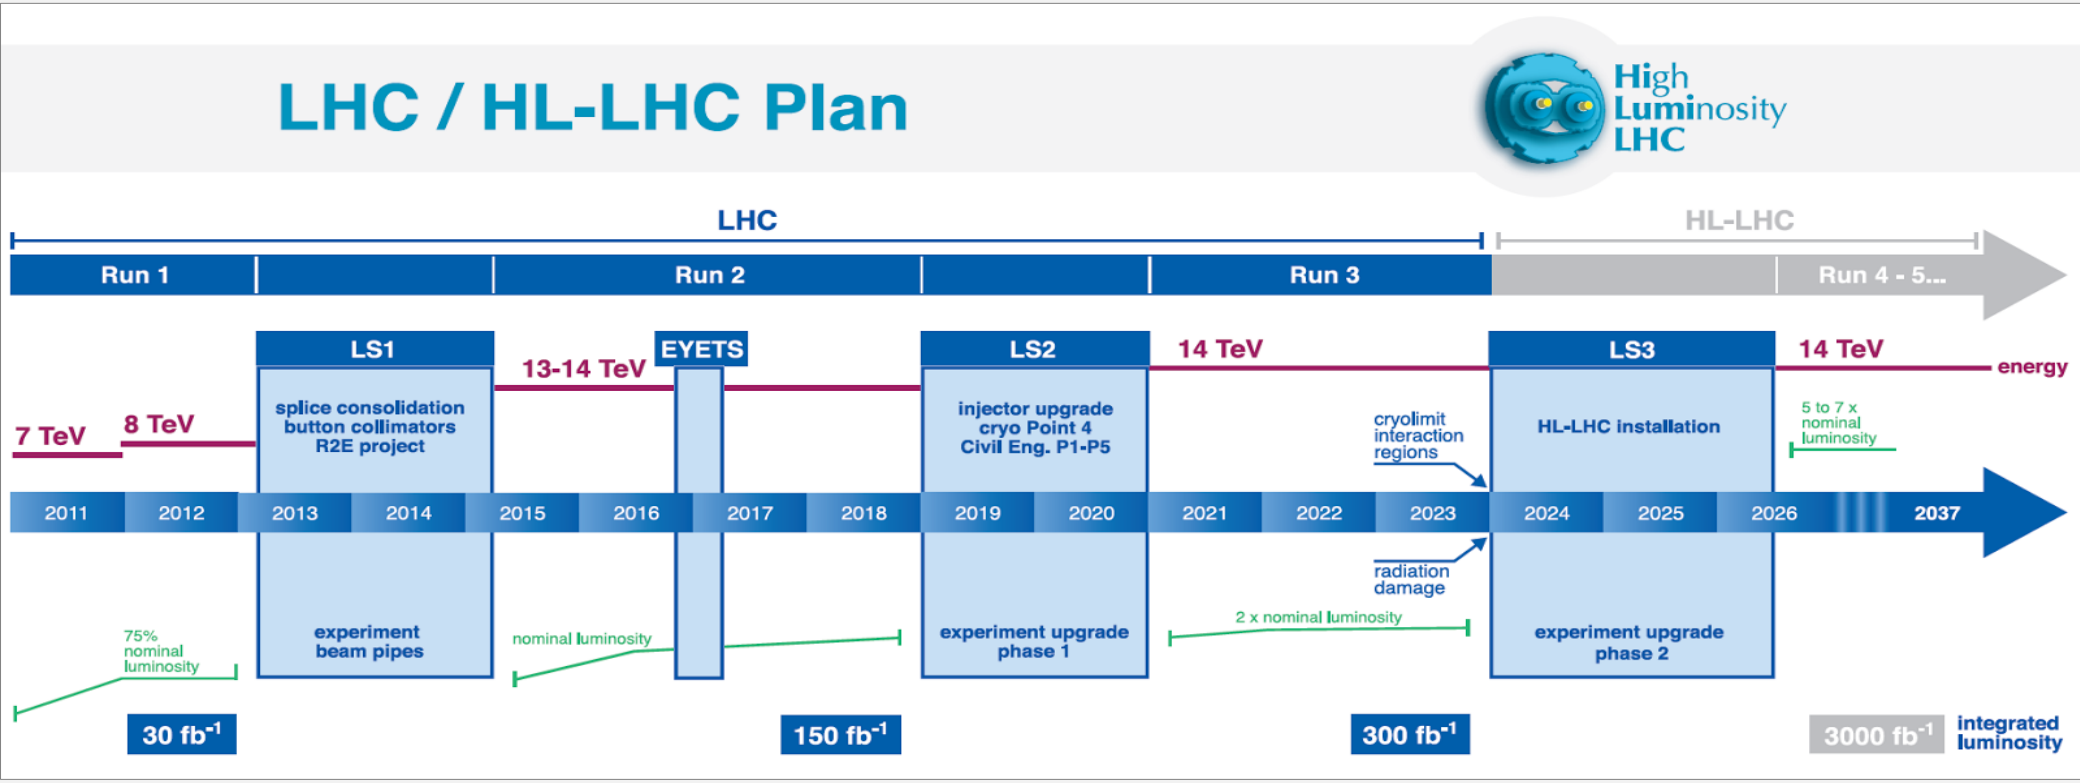


Pile-up : complexity of events
Integrated lumi.: nb of events

# What it means

**Run 3:** upgrade of ALICE and LHCb
**Run 4:** upgrade of ATLAS and CMS

**Run 2:** almost done, many improvements on the SW, the data placement ...

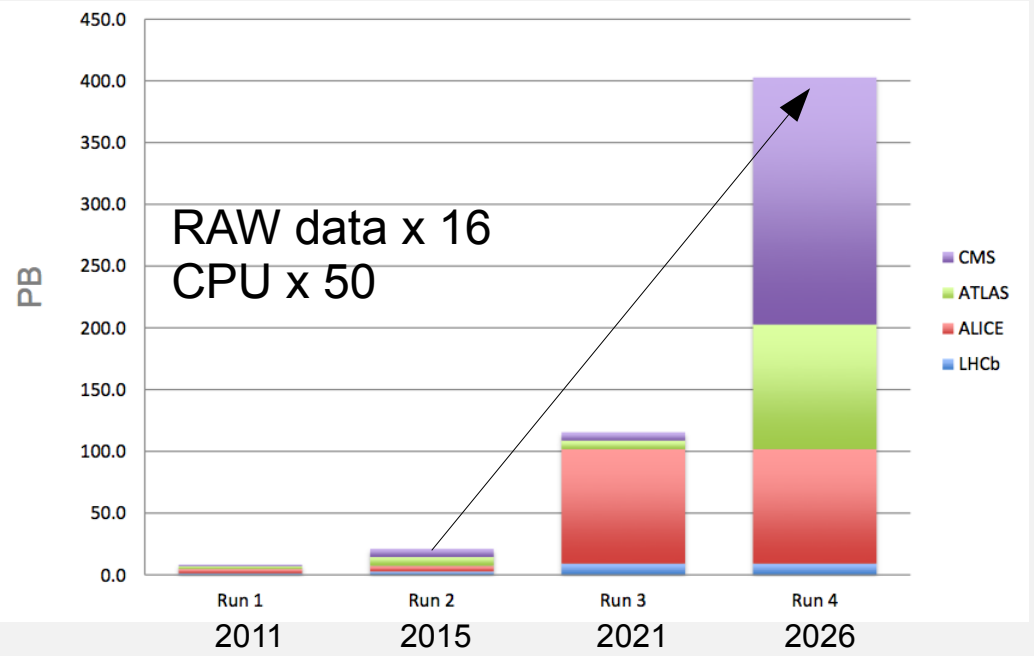


RAW data size per LHC Run
Ref : ECFA High luminosity LHC experiments workshop (2013)

# Run 3 (2021-2023)

**ALICE**

- New Online-Offline farm at CERN (reco. and data compression) → 60 PB to store / year
- Strategy: minimal replication, suppression, additional Analysis Facilities
- Pledges with a constant budget: for T2s +20% CPU and + 20% disk in sites until 2020 ; afterwards constant disk capacity & focus on CPU growth
  - Focus on CPU, diminution of storage end-points (regional « end-points »)

**LHCb**

- Computing Technical Design Report to be published this month
  - Online calibration & reconstruction (turbo stream),
  - Direct access to single events after event classification (event index)
- Major data volume growth, expected to be manageable in a « flat budget »
- Most of Tiers-2 are disk-less (run2)

**ATLAS and CMS**

- Changes more adiabatic as far as grid computing is concerned.
- ATLAS wish gain on storage overhead (operation) → storage consolidation

**Budget**

- Should stay in the flat budget « envelop »

# Run 4 (HL-LHC)

## Resource needs

- Large data volume, large processing needs
  - More registered events (x 5-10)
  - More collisions per event (x 60-200)
- Does not scale with a flat budget extrapolation
- The storage is the most problematic (site and exp.).

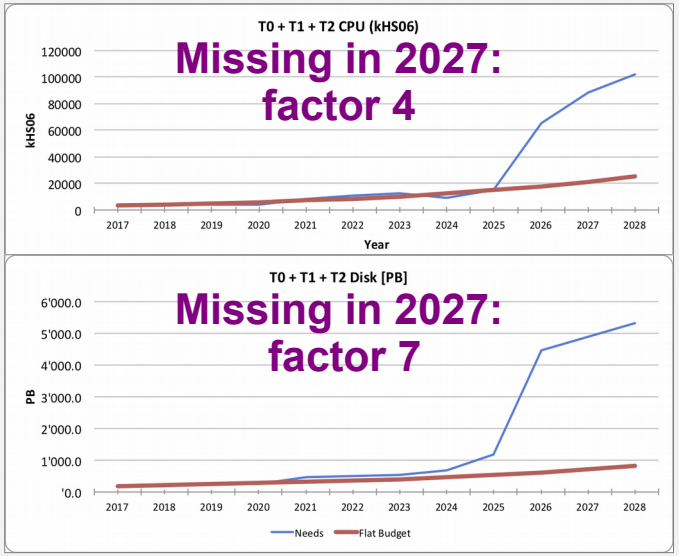


ATLAS example: resource needs with a Run2 extrapolation (including decreasing cost) vs a flat budget growth of resources

## Usage of resources

- Make the most of them
- This includes SW

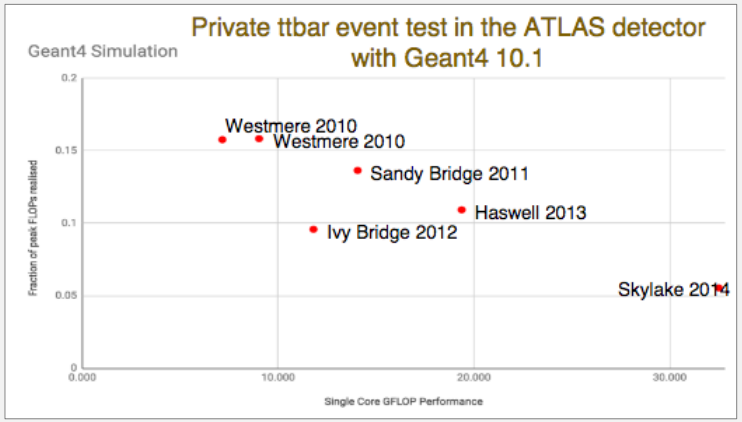


## Challenges

- Storage
- Technical: new architecture (HPC, GPU)
- Human: code rewrite, convergence in-between experiments
- Other big players are coming: DUNE (FNAL), Belle-2 et SKA, FAIR (GSI Darmstardt), CTA...

# The first step

## HEP Software Fundation

- Established in 2015
- To facilitate cooperation and common efforts in High Energy Physics software and computing internationally.
- Phylosophy og HSF: do-ocracy

*HSF web site*

## HSF Community white paper

- Mandate given by WLCG (2016) to HSF
  - The HEP community (LHC+)
  - 10 year horizon (SW and computing)
  - Plan on the directions and R&D
- The CWP is published
  - About 300 colleagues signed
  - Individual WG papers
    http://hepsoftwarefoundation.org/activities/cwp.html

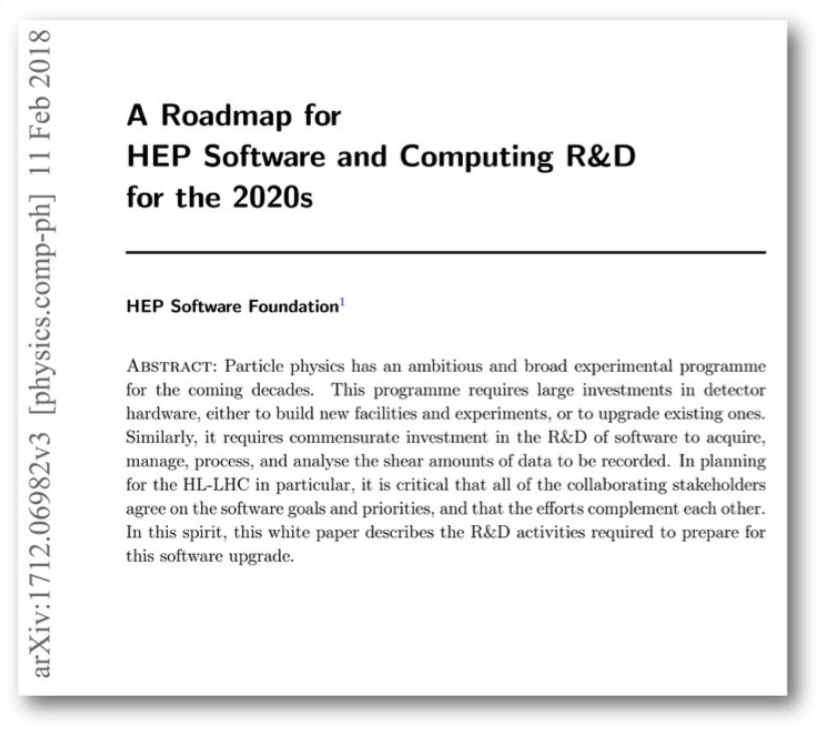


*The Community White Paper*
*https://arxiv.org/abs/1712.06982*
*Summary given in Naples.*

# CWP – working groups

13 working groups - focus on R&D to prepare the Run 4



**Facilities and distributed computing**
**- ideas -**

- Shared infrastructures
- Computing resources integration (HPC, cloud, commercial, arch.)
- Processing of small chunks
- Optimise the scheduling
- Storage consolidation (site disk caches, data lakes, vs fundings)
- Use of the network (“cheaper”), integrated into applications
- (interactive) Analysis facilities
- Convergence of experiments
- Cost optimisation (HW, RH)

# Data lakes – the original idea

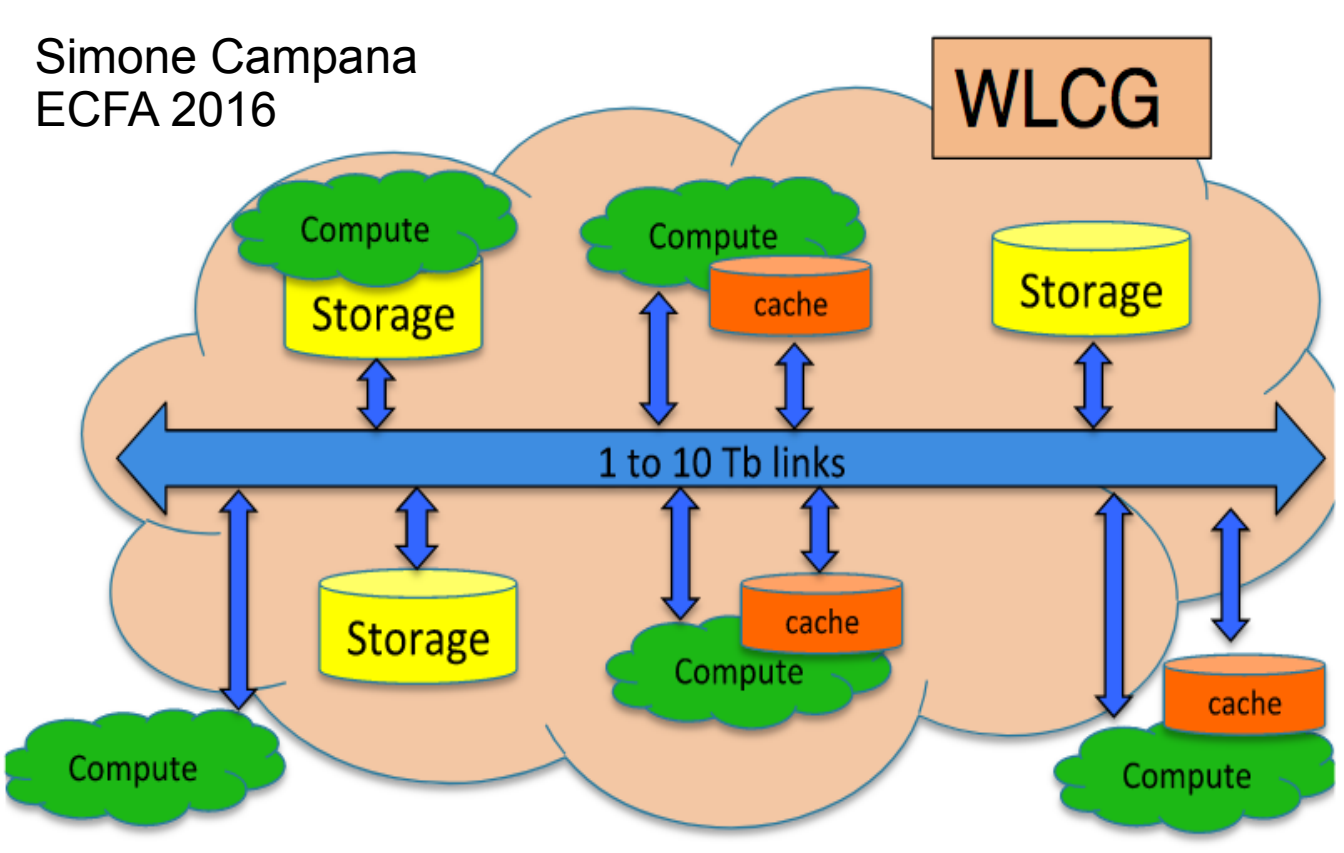


We want to keep control of the data.
We need to deploy data to a diverse set of resources (HPC, clouds, grid sites).

# Towards the HL-LHC TDR (2020)

- The LHCC requires a document from WLCG about its strategy for Run 4
  - R&D and milestones towards the preparation of the HL-LHC computing TDR (2020)

- The document will be based on the CWP
  - With a WLCG view
  - Key points are :
    - SW performance
    - Algorithmics improvements
    - Reducing data volumes
    - Managing operation cost
    - Optimising HW cost
    - Commonalties

- Document du this month (review in the WLCG MB)

1. Introduction
2. Computing Models
3. Experiment Software
4. System Performance & Efficiency
   - Cost Model
   - Software performance
   - I/O performance
5. Data & Compute Infrastructure
   - Storage consolidation
   - Caching
   - Storage, access, transfer protocols
   - Data Lakes
   - Network
   - Processing resources
   - Cloud analysis
6. Sustainability
   - Common solutions
   - Security infrastructure
7. Workplan
8. Appendix: technology evolution
9. Appendix: likely benefits

Ref. Ian Bird

# Joint WLCG and HSF workshop

- Following step to go further: Joint WLCG and HSF workshop in Naples two weeks ago

*Live Notes*
*Agenda*

- France is concerned by the HL-LHC
  - We envision the CC-IN2P3 to play a major rôle in the data workflow (data lakes)
  - FR Tier2s will contribute to the overall effort (data is our expertise)
  - We may take this occasion to re-think our regional model
- LCG-France had ~10 people attending and we will do a debrief on April 27th
  - Our to get involved ? How to make sure we will be ready for run 4 ?

- WLCG has done a debrief yesterday at the GDB
  - I did not have a chance to look though the slides yet

- My personal thoughts on this joint workshop:
  - Very dense with talks; money and cost optimisation was everywhere
  - Of particular interest for the LCG-France community: Technology watch; Common data management and data lakes; Workload management; Performance and cost modelling

# Common data management and data lakes

- One of the priority of WLCG
- Follow-up on the CWP with concrete examples
  - Distributed storage (EOS, dcache, XDC)
  - Data provisioning (Ruccio, Dynamo)
  - Caching technologies (Xcache, http cache)

- We need to understand the benefit of “data lakes”
  - Integrated distributed storage (and compute)
  - Level of service (tape/disk/SSD)/cache
  - Workload management knows the data in use
  - Gain in quality of service, in volume to be stored

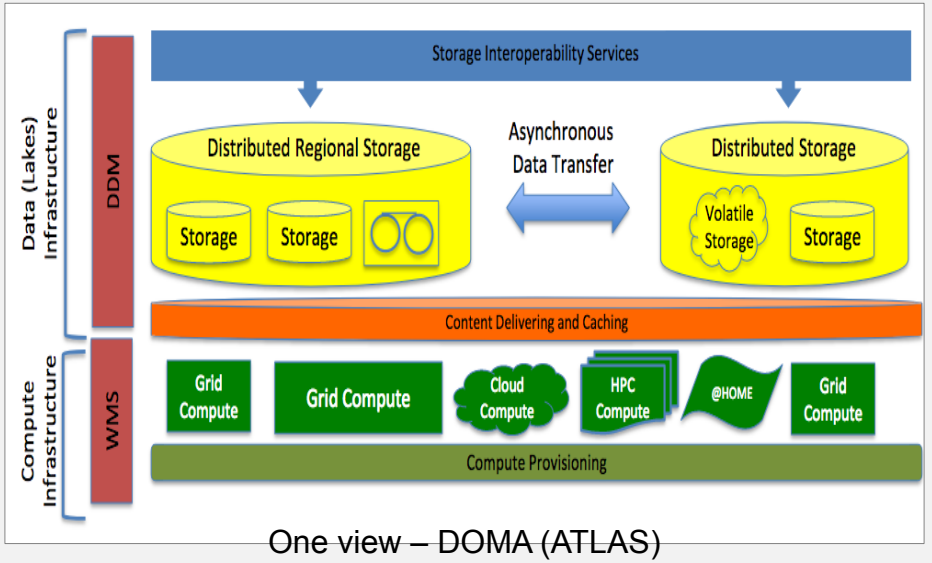


One view – DOMA (ATLAS)

- Proposal to create a DOMA (Data Organisation Management Access) evolution project
  - Experiments, MW developers, storage providers, facilities
  - Key points : commonalties, standards, sharing with other communities

# Workload management

**MW: ARC, PanDA, DIRAC, OSG**

- No merging perspectives, neither standardisation
- Common approaches, potential for common R&D ? provisioning heterogeneous resources, caching, dynamic workflows, fine grained processing → we need to follow-up on this

**Experiments: ATLAS, CMS**

- Few distinctions
- Many agreements
  - Biggest problem is disk
  - Rely more on tape
  - Use of carousel mode

- *Cost ?*

## ATLAS & CMS - commonality & distinctions

- Agree: computing can't be allowed to blow the budget. A constrained resource. Technical and sociological challenge to expose to analysts the costs and tradeoffs of their processing.
- Agree: expectation management is vital. Understanding of what we can do, by when.
- Agree: addressing CPU shortfall demands much work, but we see paths.
- Agree: the bigger problem and challenge is disk storage.
- Agree: there will be managed production stages in the processing chain when we can expect the inputs for a campaign will not all fit on disk.
- Differ: CMS has one agreed compact analysis data product, ATLAS has ~100; if this persists, ATLAS likely will be hit by this earlier in the chain than CMS.
- Agree: The best bet for reducing storage budget is to rely more on tape and less on disk.
- Agree: tape usage will need to be orchestrated by WM & DM (eg train processing in 'carousel' mode) to reach the needed efficiency and turnaround time to meet analyst expectations. Generally, a full (re)architecting of the workflow + storage stack with WM and DM closely interacting.
- Agree: Interested to make greater tape usage part of near term R&D.
- Agree: we need to calculate, model, analyse just what such workflows imply.
- Differ: ATLAS couples this interest with granular processing/streaming R&D in the data lake context.
- Agree: common interest in data/analysis facility to explore 'declarative programming' for analysis.

3

# System performances and cost modelling

- New WLCG-HSF working group (launched in November 2017)
  - SystemPerformanceModeling@cern.ch, 35 participants (5 from LCG-France)
  - Workload, workflow, framework developers, people who plan, engineer and operate IT sites
- Goal: quantify the impact (in term of cost) of the changes in the HL-LHC computing model
  - WLCG oriented, could be of use for CTA, SKA...

- We have to understand/guess:
  - Payloads → requirements for infrastructure → cost
    - Create a feed-back loop to optimise the amount of the physics we can do within budget
  - The resources needed by experiments given the LHC properties+exp. model
  - Network, data lakes … manpower, complexity
  - Future HL-LHC payloads (chain of microkernels)

*More information : Indico wiki*

# Technology watch

Ref. Naples Workshop

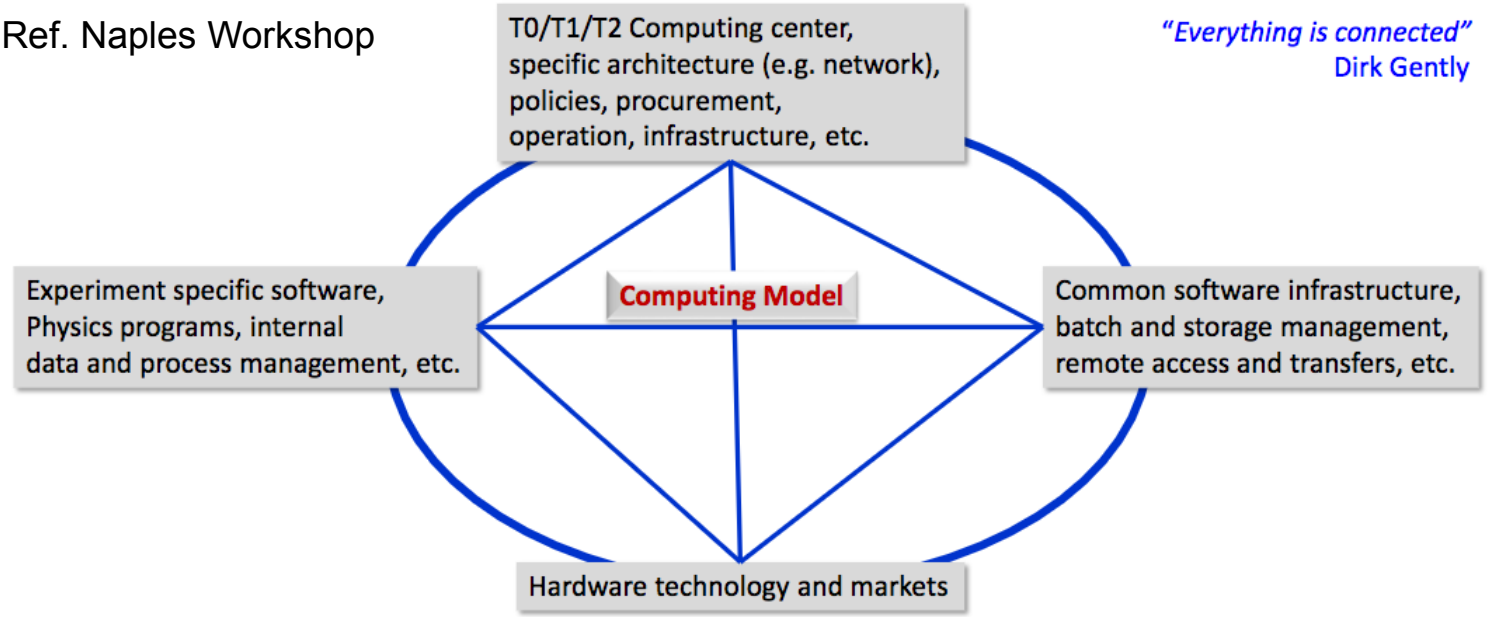


- Proposal to launch a new WG "technology watch" (long term WG)
- Coordination of sites seeking for technology improvement
- Important inputs for the "cost modelling" WG

# Other communities

- Other communities are emerging
  - Similar datasets as the LHC
  - Similar timescales
- We will share the same infrastructure
  - Standard protocols, convergence

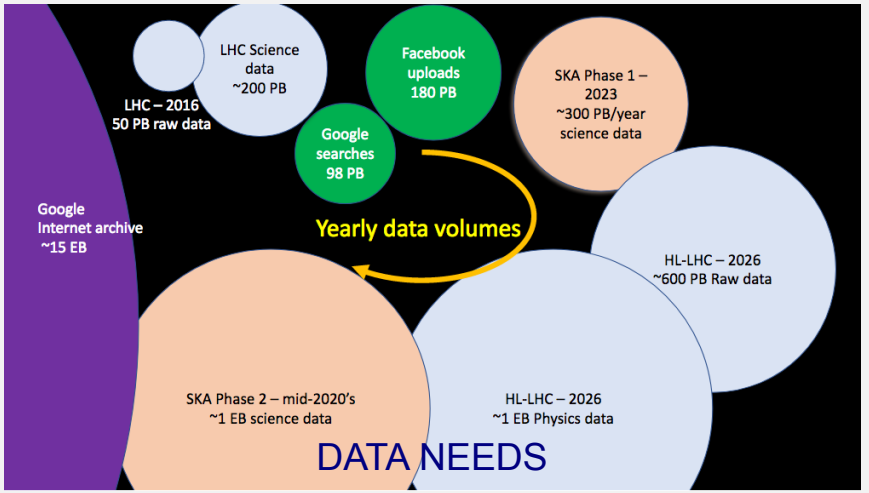


## WLCG/HEP are concerned

- Understand the landscape in participating country → Computing Scientific Form
  - https://indico.cern.ch/event/663273/
  - To identify longterm solutions for publicly funded scientific computing. Develop strategies for data-intensive applications with large computing requirements for the next decade.
- Bring other big players in our circle, e. g.
  - Belle-2: a close partner of WLCG
  - SKA: collaboration agreement signed with CERN (July 2017)

# Conclusion



# Conclusion

## LCG-France is doing well

- Each site is working well, providing its share of resources, active community
- Planning to play a major role in the HL-LHC computing

## Sites have to be prepared

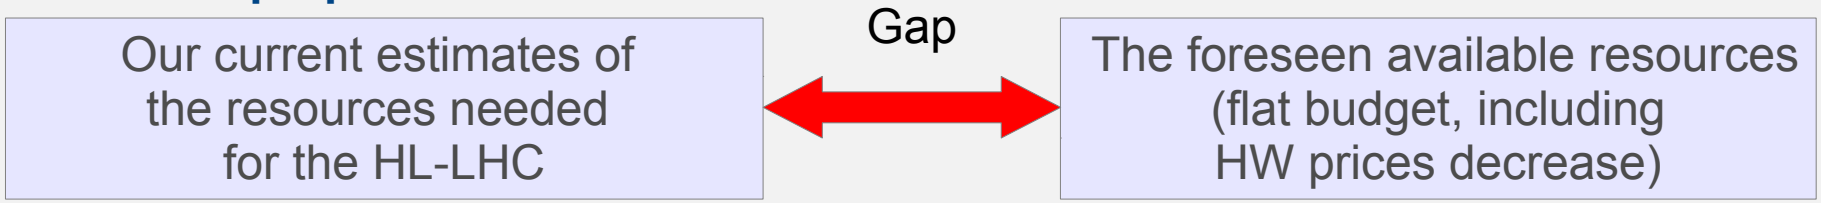


## Solutions

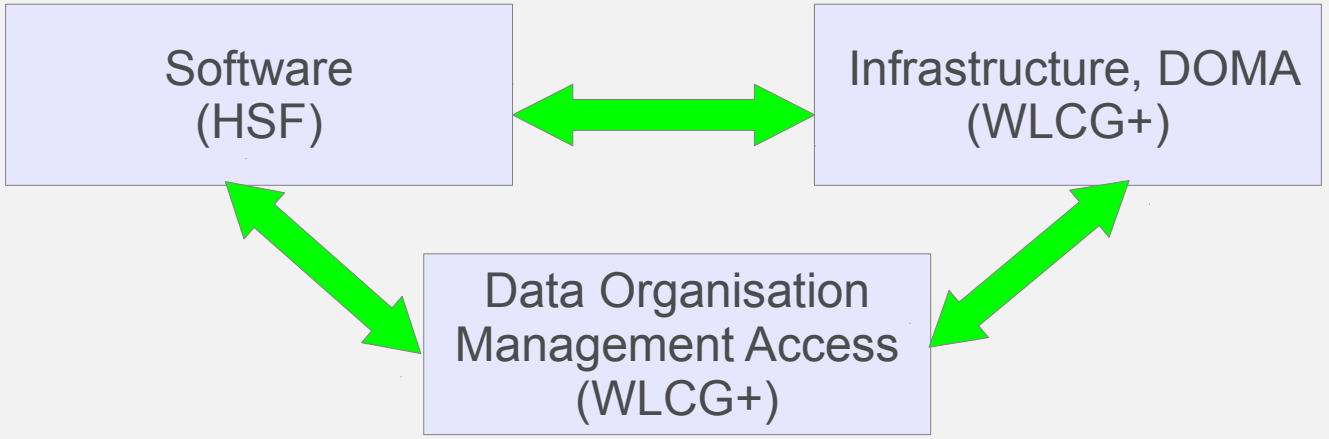


A long process starting now !

# Back-up slides

# Evolution of experiment requests

- 2016-2017: exceptional requests du to the good performance of the LHC
- 2017-2018: relatively small requests compared to previous year
- 2018-2019: for Tier-1s and Tier-2s requests are below 20%, except for ATLAS disk
- 2019 has to be finalised at the coming C-RRB on April 24th 2018 (https://indico.cern.ch/event/707294/)

Tier-1s

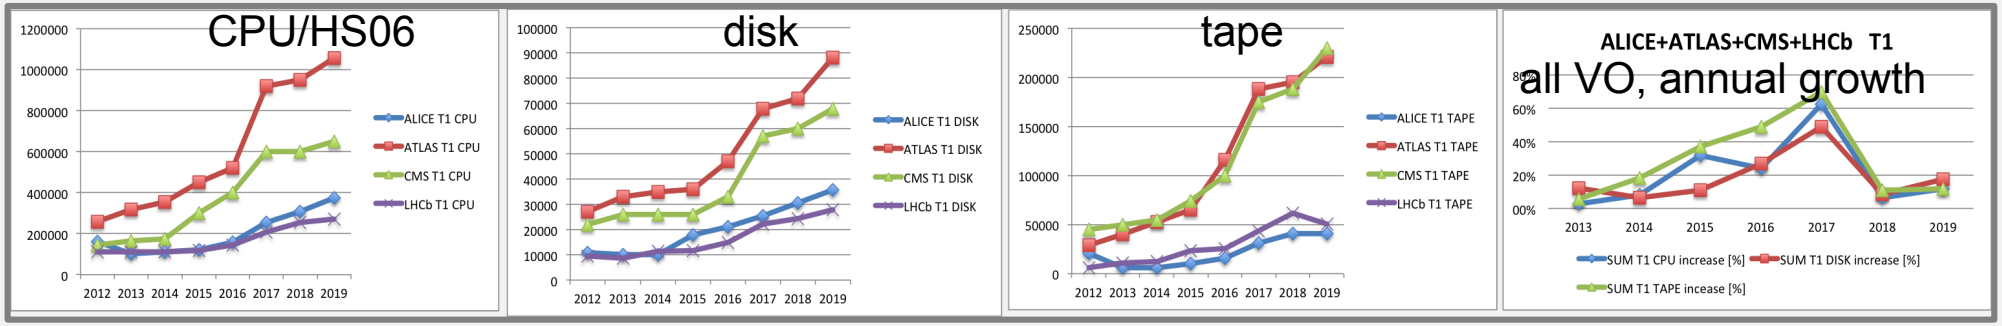


Tier-2s

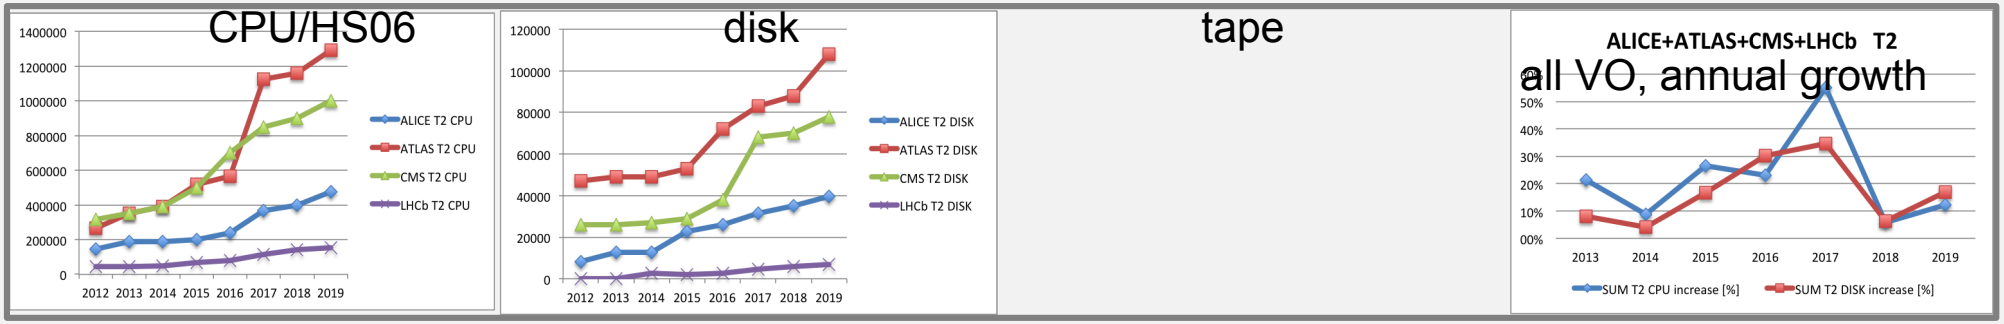


As of November 20017 (autumn C-RRB)

# Flat budget

## Constraints

- Funding agencies will support LHC computing at best within a flat budget.
- Experiments require already at Run 2 a large increase of resources.
  - More events registered
  - More complex events

## Models

- Some models for an annual growth of resources within a flat budget are being developed:

| Modèle | CPU | DISK | TAPE |
|---|---|---|---|
| CERN [1] | 20% | 15% | 15% |
| CC-IN2P3 [2] | 25% | 18% | 38% |

[1] CERN-LHCC-2014-04, édité par WLCG (2014)

[2] R. Vernet, CHEP 2015

# Operations - storage

## Tier-1

- handling dCache + HPSS

## Tier-2/3

- Historically, T2/T3 sites in France were operating DPM
  - Subatech (ALICE-only) moved to EOS ; ALICE storage is mostly xrootd (native)
  - Sites are happy with DPM (good performances)
  - Our biggest DPM site is IRFU with 2.9 PB (ALICE, ATLAS and CMS)
- Activities in France around DPM (rep. Andréa Sartirana LLR)
  - FR is an active member of the DPM collaboration since it was created (2013)
    - 2018 tasks: test on the removal of the gridftp redirection & muli-site DPM
  - CPPM: the very first site to go to DOME
  - LLR: part of the MW readiness WG (CMS)
  - LAPP: offers a test bed to ATLAS

# Operations - compute

## EGI/WLCG recommendations

- EGI (2014): move away from Maui since it is no more maintained [EGI-SVG-2013-5545]
- WLCG: modern batch system allow a better handling of jobs (e.g. cgroups, containers)
  https://twiki.cern.ch/twiki/bin/view/LCG/WLCGWorkloadRecommendations
  - Batch systems : HT-Condor for HTC loads ; SLURM for HPC loads
  - CEs: HT-Condor works with HT-Condor (accounting non std yet) ; ARC-CE works with SLURM and HT-Condor, also used as lightweight CE (local file staging solution)
  - e. g. CERN is HT-Condor + ARC-CE

## CC-IN2P3

- Univa Grid Engine
- HT-Condor as a possible replacement

## FR Tier-2/3

- Historically, FR T2/T3 sites were operating Torque/maui + CREAM-CE
- A growing interest for HT-Condor (only IRFU_GRIF has an ARC-CE in prod.)
  - Already moved: LAL_GRIF and LLR_GRIF
  - Under test: IPNO_GRIF, CPPM (+ ARC-CE CentOS 7 WN)

| Number of sites | Torque/Maui | SLURM | HTCondor | USGE/SoGE | LSF |
|---|---|---|---|---|---|
| 2014-04-13 | 101 | 10 | 10 | 15 | 7 |
| 2016-01-26 | 92 | 14 | 17 | 14 | 6 |
| 2017-01-17 | 84 | 15 | 21 | 10 | 5 |
| 2018-02-26 | 78 | 18 | 25 | 8 | 5 |

Ref. GDB March 2018

# Operations – moving to EL7

- WLCG/EGI:
  - No push from WLCG to go to EL 7
  - Nothing organised at the EGI level
- Some sites wish to migrate
  - other communities have interest in EL 7
  - HS07 benchmarking gives +15% on CC7
- VO SW readiness:
  - ALICE and LHCb: can run on EL6 or EL7
  - ATLAS can run on EL7, no push for all sites to migrate (user analysis is a bottleneck)
  - CMS: requires both EL6 and EL7 on every site
- Migration at site is mainly a site decision in coordination with the hosted VOs

- In France:
  - CC-IN2P3 is migrating,
  - CPPM plans to deliver an HT-Condor+ARC-CE site soon (new WNs installed in CC7)
  - IPHC envision to  move

# Operations – containers (1)

- Container solutions are very popular
  - Flexibility & no performance loss compared to VM
  - Singularity: interesting new engine from the HPC world, very lightweight
- CMS required every each site to install Singularity (SL6 &CC7) ; their needs are:
  - To process the 2018 data with a CentOS 7-only release
  - To continue to analyse existing data with SL 6
  - A dead-line has been set on March 1st 2018 to include the singularity test in their critical profile, i.e. a site without singularity would see its availability go to null.
- In France, we have had some discussions about this plan:
  - the security of this deployment – no official assessment in this subject
    - Only the privileged mode is available on SL 6
  - the commonalties of this solution – what about other VO needs ?
    - No overlayfs is available in the « un-privileged mode » (one should know all the mounting point prior the creation of the image) → not suitable to ATLAS
    - In the mean time: the WLCG working group on Containers did not conclude

# Operations – containers (2)

- Our conclusion was to go for the deployment on CMS sites (CC-IN2P3, IPHC, IPNL, IRFU, LLR, LPNHE, LAL)
  - Mode privileged: no security bottleneck identified (CERN, CC)
  - This solution is not suitable to other VO but CMS needs it
- To add a layer of security (CMS recipe) : one should forbid the user to be able to:
  - create new images → only “singularity-runtime” to be installed
  - mount volumes (not suitable to ATLAS) → disable the “loop-back based” image mount
- Status in France as of today:
  - Each CMS site deployed singularity and the CMS SAM test is OK
  - The CC-IN2P3 did deploy directly on CentOS7
- Singularity is quite a new product and it is moving fast (bug reports)
  - We refer to the WLCG baseline table for the minimal version to be deployed https://twiki.cern.ch/twiki/bin/view/LCG/WLCGBaselineTable
  - We are following the WLCG containers WG wlcg-containers@cern.ch

# Operations – storage & compute

## Storage

- CC-IN2P3 uses dCache, xrootd and HPSS
- T2/T3 sites in France were operating DPM historically (+xrootd for ALICE)
  - FR has developed expertise and are happy with DPM (good performances)
  - We are active in the DPM collaboration (A. Sartirana) and in the MW readiness WG
  - Our biggest DPM site is IRFU with 2.9 PB (ALICE, ATLAS and CMS)
- Some sites moved to EOS (Subatech for ALICE, IPNL for non LHC VO)

## Compute

- Historically, FR T2/T3 sites operates Torque/maui + CREAM-CE ; T1 UGE
- EGI recommends to move away from torque/maui
- WLCG recommend new sites to go for a modern batch system (e.g. cgroups, containers)
  https://twiki.cern.ch/twiki/bin/view/LCG/WLCGWorkloadRecommendations
  - HT-Condor or SLURM ; HT-Condor-CE (limitation) or ARC-CE
- HT-Condor in FR: LAL&LLR in production, IPNO&CPPM under tests, possibly CC-IN2P3
- ARC-CE in FR: IRFU_GRIF in production, CPPM under test (CC7 migration)

# Operations – moving to EL7

- WLCG: no push from WLCG to go to EL 7 (no recommandation at the EGI level)
- Some sites wish to migrate (non LHC communities)
- VO SW readiness:
  - ALICE and LHCb: can run on EL6 or EL7
  - ATLAS: can run on EL7, no push for all sites to migrate (user analysis is a bottleneck)
  - CMS: requires both EL6 and EL7 on every site (2018 data proc. will need EL7)
- Migration at site is mainly a site decision in coordination with the hosted VOs
- In France: CC-IN2P3 is migrating, CPPM plans to deliver an HT-Condor+ARC-CE site soon (new WNs installed in CC7)

**A big push on containers deployment**

- CMS required every each site to install Singularity (SL6 &CC7) before March 1st; they need to process 2018 data with EL7 releases and previous data with SL6 releases
- In France, we have had some discussions about this plan: security ? Commonalties ?
- Status in France as of today: each CMS site deployed singularity (CCIN2P3 on CC7)
- Worries: singularity is new and moving fast, we follow the WLCG containers WG wlcg-containers@cern.ch

# French landscape

- LCG-France sites are part of the French NGI France Grilles
  - They are EGI certified
  - All but one are open to non-LHC VO

- Within France Grilles
  - Five sites are not be IN2P3
  - Large move of non-LCG sites to cloud (Lille, Toulouse, Montpellier)
    - To serve tail of sciences
  - A national working group on cloud deployment and operations (federated cloud)
  - Regarding grid, France Grilles operates basic structure deployment (EGI specifics), monitoring and security

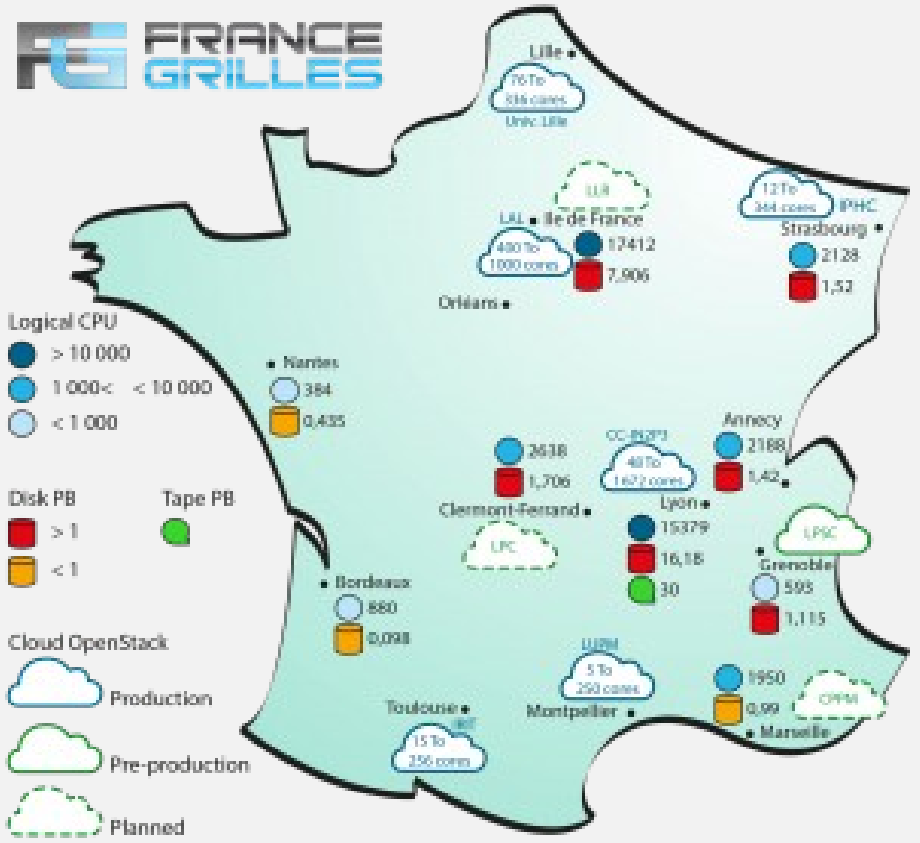

# T2s - Pledge evolution

- On any given year, pledge increases depends on local situation and how much hardware is to be renewed. Overall all sites show an increase over time.

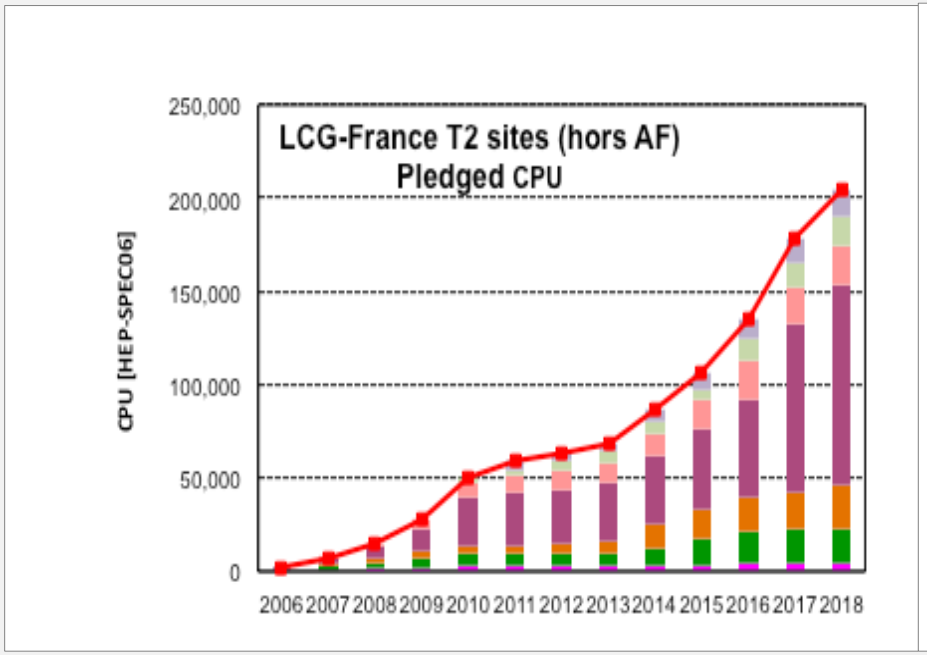


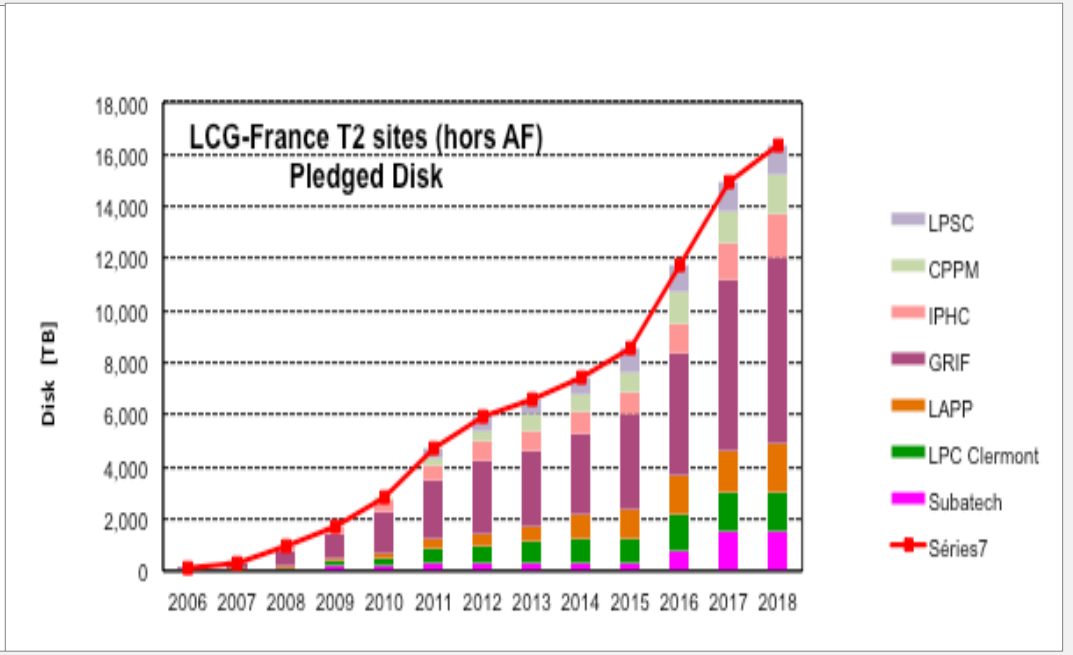

# T2-FR vs exp. requests (pledges)

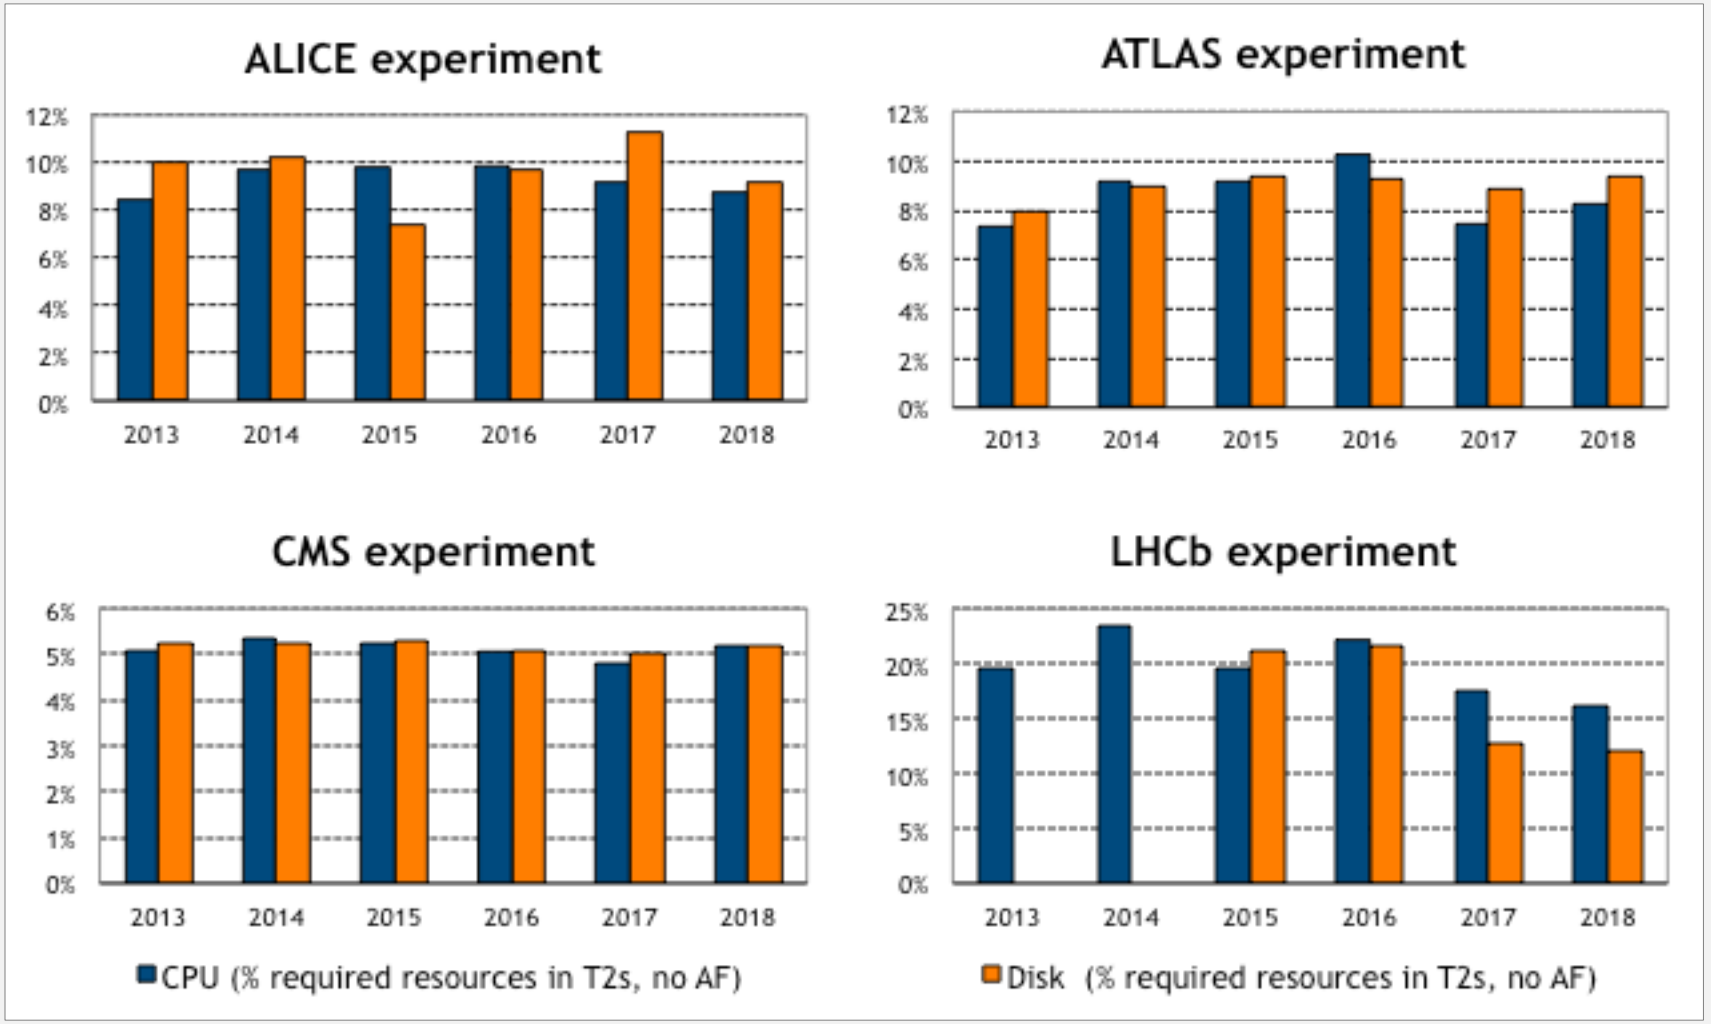

# Tier-1

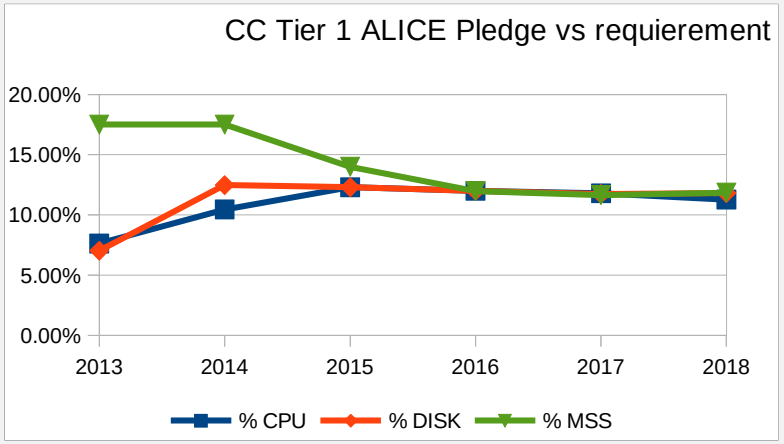


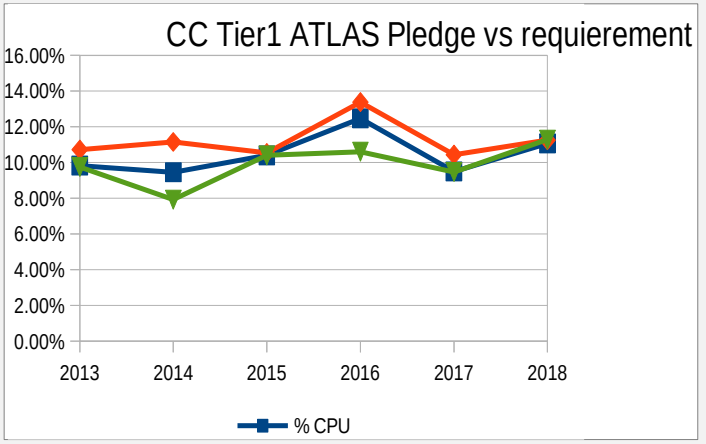


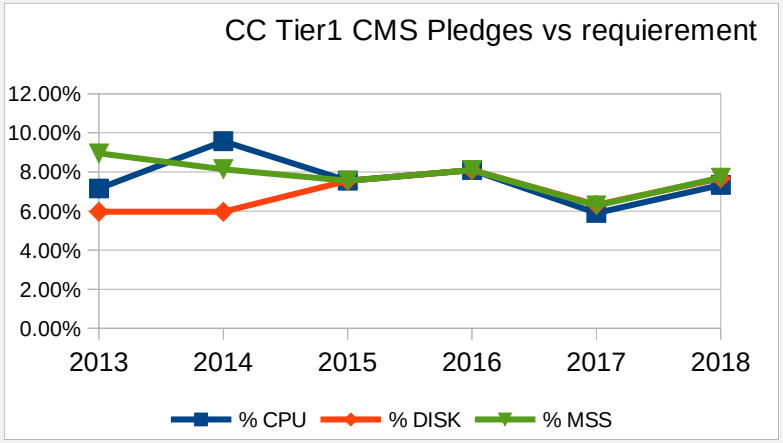


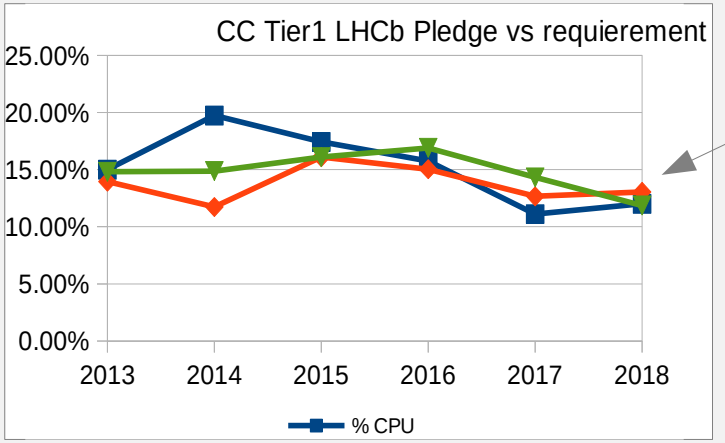


LHCb 2018: less needs than forseen (not in REBUS yet)

# Other communities

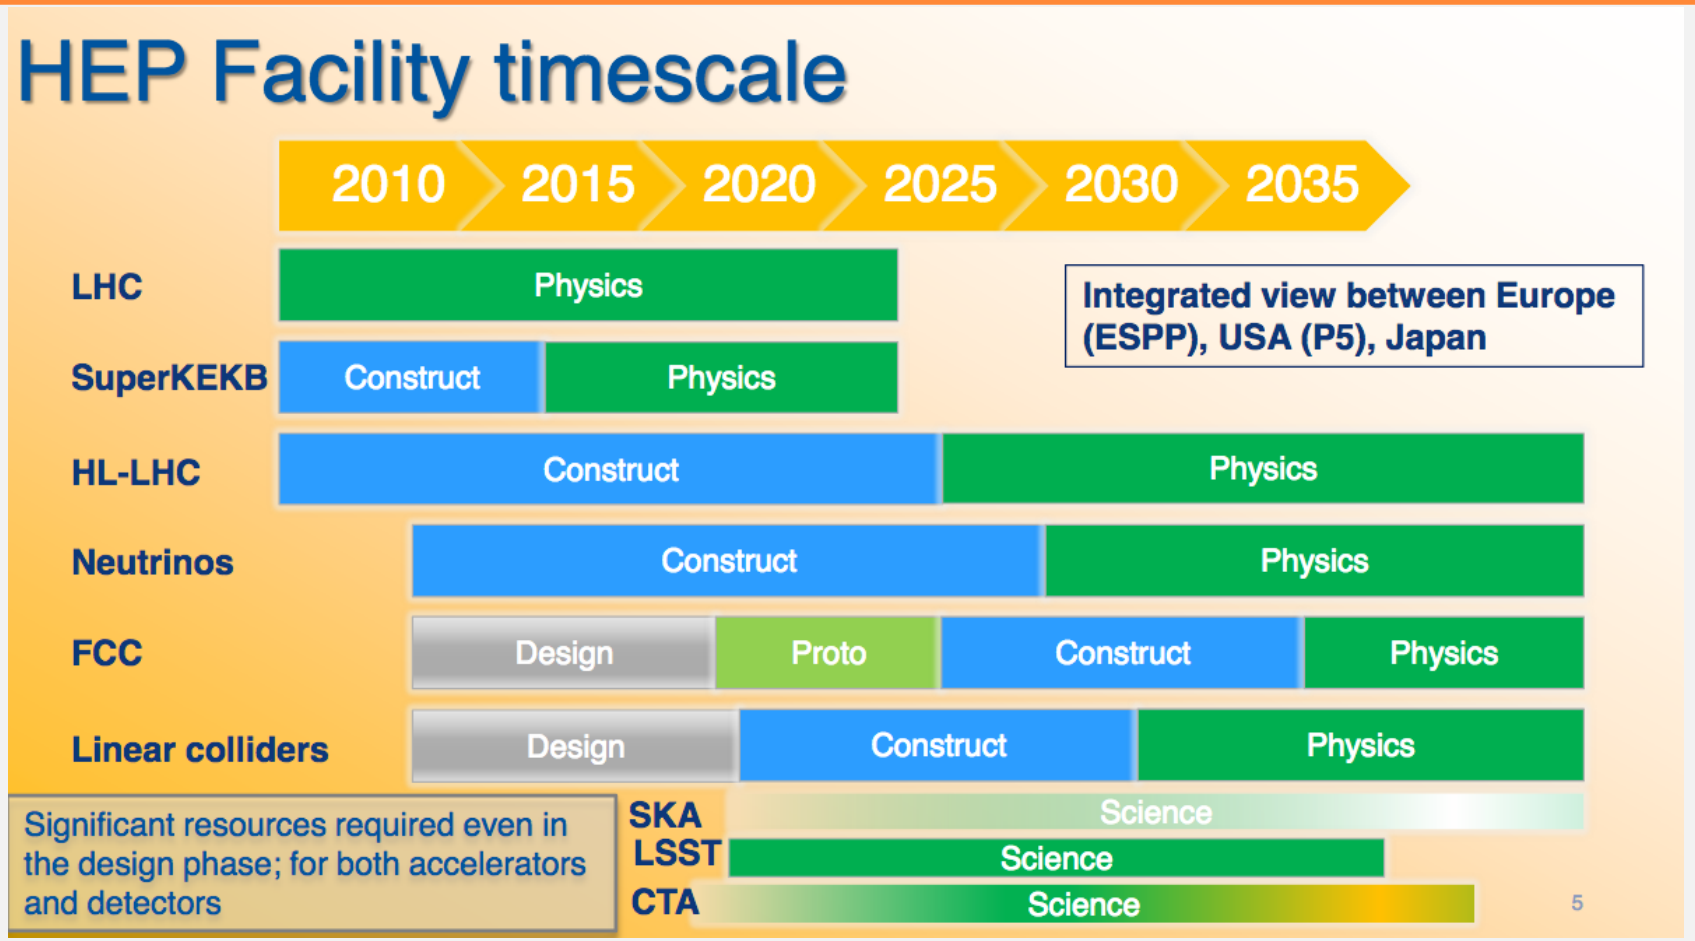


Ref.: 1st Scientific Computing Forum (February 2017), Ian Bird, https://indico.cern.ch/event/581096/contributions/2450212/attachments/1404826/2161766/Computing-Strategy-Feb2017.pdf