LHC activities and HL-LHC R&D

The LCG-France and DOMA-FR community

v1.0 25 mai 2022

Table of Contents

In the context of the Scientific Council examination, we have focussed particularly on the organisation, main achievements, R&D and prospectives. LGC-France primarily task is to operate computing production for the LHC experiments but we left out the technical details and reporting that are examined yearly by the "Calcul et Données" IN2P3 DAS. Similarly, we will concentrate on the results of activities in the DOMA project and omit the yearly reporting informations.

It did not seem important to report on all activities in the more than 15 years existence of the LCG-France project, we have decided to concentrate on the period 2015 – 2022.

Description

The WLCG (*Worldwide LHC Computing Grid*) [2] is an international collaboration of 42 countries in charge of organising resources (CPU and storage) for the LHC experiments. Driven by the specific needs of the experiments, it brought together a hierarchical structure of sites using the European (EGEE then EGI), Nordic (NORDUGRID) and North-American (OSG) grid computing organisation :

- a Tier 0 : located at CERN for the data collection, initial reconstruction and data distribution;
- a dozen of Tier 1 : large national computing centres with first class network connectivity, availability and quality of service for data distribution and storage, reconstruction, simulation and analysis;
- about 170 Tier 2: local sites for simulation and analysis;
- Tier 3: institute clusters or large facilities dedicated to analysis, often co-located with a Tier 2 or Tier 1

The organisation was formally put together with an MoU [3] between institute providing a Tier 1 or Tier 2 and CERN with corresponding quality of service requirements and more general agreement to pledge enough computing resources to experiments compared to the overall needs. Tier 3 sites are not part of the MoU. Typically, experiment requests are examined by the C-RSG (Computing Resources Scrutiny Group) and approved by the RRB (Resources Review Board) then sites express their pledges towards these requests for the following year.

WLCG was originally designed and operated as a strict hierarchical system at least for data distribution, with data flow T0 ↔ T1 ↔ T2. The performances of networks and the quality of service provided by the larger Tier 2 sites was such that this scheme rapidly evolved [4] towards site to site data exchange with, for some of the experiments, selected Tier 2 being used as data consolidation sites.

LCG-France is the French national organisation created by CEA/DRF and IN2P3/CNRS for the coordination of the LHC computing in France within WLCG with a Tier 1 centre (CC-IN2P3), seven[1] Tier 2 centres (GRIF is formally a single entity from the grid point of view but distributed on several sites) and one Tier 3 centre as shown in Figure 1. LCG-France implements the technical coordination of computing and ensure that the French contributions correspond to the need of the experiments and in particular physicists from French Institutes. It also manages, in coordination with CEA, IN2P3 and the director of CC-IN2P3, a yearly budget of 1.8-2M€ that is part of the CNRS IR « Centre de Calcul de l'IN2P3 ».

1 The last Tier 2 co-located with the Tier 1 which was serving CMS has been decommissioned in 2021.

The French sites are certified sites of the EGI grid infrastructure. Most of them offer services to non LHC and even non HEP experiments. Each Tier 2 and Tier 3 decides on the experiment(s) it supports and on the shares between them, mostly based on the laboratory level of effort in the experiment. Details can be found in Table 1. Note that although IPNO and LAL have merged into IJCLAB, these are still two different sites from the grid point of view. This might evolve once the underlying technical infrastructure has converged.

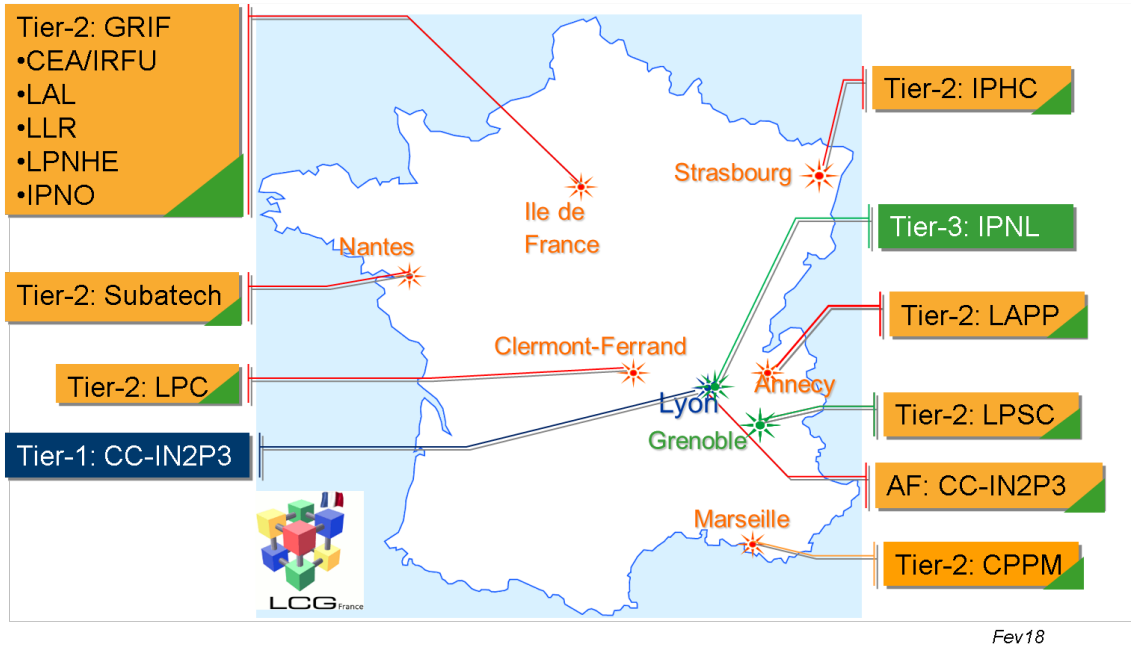

Figure 1: LCG-France sites

The French Tier 1 at CC-IN2P3 supports all experiments with shared defined in a "Protocole d'Accord" (see later). CC-IN2P3 no longer runs co-located Tier 2 site but supports analysis by offering both grid and off-grid resources (details depending on the experiment) labelled as Analysis Facilities (AF).

In the period 2015 – 2017 it became more and more obvious that the computing models and technical solutions in use would not scale to HL-LHC, in particular for ATLAS and CMS, and that a more common effort would be needed to tackle these challenges given the person-power available inside the collaborations. This lead to the creation of the HEP Software Foundation [9] to address shortcomings of the full software spectrum (generators, simulation, reconstruction, analysis) and in particular the usage of accelerator such as GPU or FPGA. On the computing side, WLCG created several working groups, one of which being DOMA (Data Organisation Management and Access), the IN2P3 DOMA-FR master project being used to coordinate the French contributions.

Site	ALICE	ATLAS	CMS	LHCb
T1 CC-IN2P3	✔	✔	✔	✔
T2 CPPM		✔		✔
T2-GRIF	✔	✔	✔	✔
T2 IPHC	✔		✔	
T2 LAPP		✔		✔
T2 LPC	✔	✔		✔
T2 LPSC	✔	✔		
T2 Subatech	✔			
T3 IPNL	✔		✔	

Site	ALICE	ATLAS	CMS	LHCb
GRIF_IPNO	✔			
GRIF_IRFU	✔	✔	✔	
GRIF_LAL		✔	(✔)	✔
GRIF_LLR		(✔)	✔	✔
GRIF_LPNHE		✔	(✔)	✔

Table 1 *Experiments supported at each of the Tier 2 / Tier 3 sites (left) and breakdown for the federated GRIF site (right). The symbol (✔) represents experiments for which the site is not officially a Tier 2 but that runs jobs for that experiment since the batch system is shared been sites.*

Organisation

As a distributed computing infrastructure for production in a constantly changing landscape of LHC computing and WLCG, LCG-France needs constant coordination both on the technical side and on the policy / prospective side, with a technical and a scientific coordinator for day to day operations. The coordinators represent France at the various WLCG meetings (Operations (and France-Grilles Operations), Management Board, Grid Deployment Board, Overview Board, Collaboration Board).

A monthly technical coordination meeting relays informations from WLCG, EGI and the experiment computing operations, organises middleware migrations, setup specific interest groups, etc… The direction, experiment contact persons and site technical contact persons are members of this technical committee. Since 2022, this meeting has merged with the corresponding France-Grilles meeting.

The monthly steering board (“Conseil de Direction”) focuses more on the computing policies and organisation, budget, WLCG and experiment policy, long term evolutions of computing models, R&D, French and European political context, etc… The board is composed of the direction, the CC-IN2P3 director, experiment and site scientific contacts.

An Executive Board meets in practice very infrequently. It is composed of the IN2P3 director and DAS involved in HEP, computing and technical coordination, the IRFU director and computing contact, the

WLCG coordinator, the LHC experiment project leaders for France, the CC-IN2P3 director, a representative of the Tier 2 / Tier 3 sites and the LCG-France direction.

The membership of these last two committees are detailed on the LCG-France web site [5] with their role defined in this document [6].

The technical and steering board committees meet monthly via videoconference but can't replace face-to-face meetings for more in-depth discussions and also more informal discussions and coordination. For that we organise two four-half-day workshops [7] (obviously this turned out to be impossible in 2020 and 2021 due to the COVID-19).

For the last 10 years or so, a non binding MoU ("Protocole d'Accord") [8] between hosting laboratories, the LHC projects and IN2P3 organises the funding (of the IN2P3 sites for the current version) and objectives of LCG-France. More precisely, the 2018-2022 version specifies:

- The goal is being to pledge 8-10% of the global CPU and storage pledges;
- Provide budget for replacement of out-of-warranty capacities, covering 70% of this cost for Tier 2;
- A share per experiment at CC-IN2P3 corresponding to the overall French contributions to the experiments, namely 45% for ATLAS, 25% for CMS and 15% each for ALICE and LHCb;
- Priorities in assigning funds:
    1. maintain capacities at Tier 1 and AF
    2. minimal growth of 10% for Tier 1
    3. contribute to Tier 2 capacity renewal
    4. additional growth for Tier 1 and AF

In practice, we never had to restrict funding to Tier 2 sites below the 70% contribution listed in the document. Note that Tier 2 sites have to provide the funding for the remaining part of the renewal, infrastructure (data centre, UPS, cooling, service machines, ...) and growth.

The DOMA-FR master project has a more informal organisation with general meetings organised as needed (three dedicated workshops and additional sessions during LCG-France workshops).

The success of this project was made possible by the implication and dedication of people (noticeable in particular during the difficult years 2020-2021 with the pandemic). The overall human resource has stabilised since 2016 (see Figure 2 for IN2P3 to which about 1 FTE from IRFU can be added) but the Tier 2 IT fraction slowly but steadily decreases with time. We consider 1-1.5 FTE as a minimum to maintain a Tier 2 grid site if this relies on more than one person. Several of the French Tier 2 have reached this minimal level, in particular LPSC is below and SUBATECH is anticipated to go below with the retirement of a key site admin. As a consequence, these two laboratories have decided to stop

their grid site by the end of Run 3. More retirements are expected to take place in the next few years, certainly before HL-LHC, in three other labs so recruitment planning is essential. Even more worrying is that a position opened at IRFU a year ago has not been filled, so recruiting might be difficult even with a position opened.

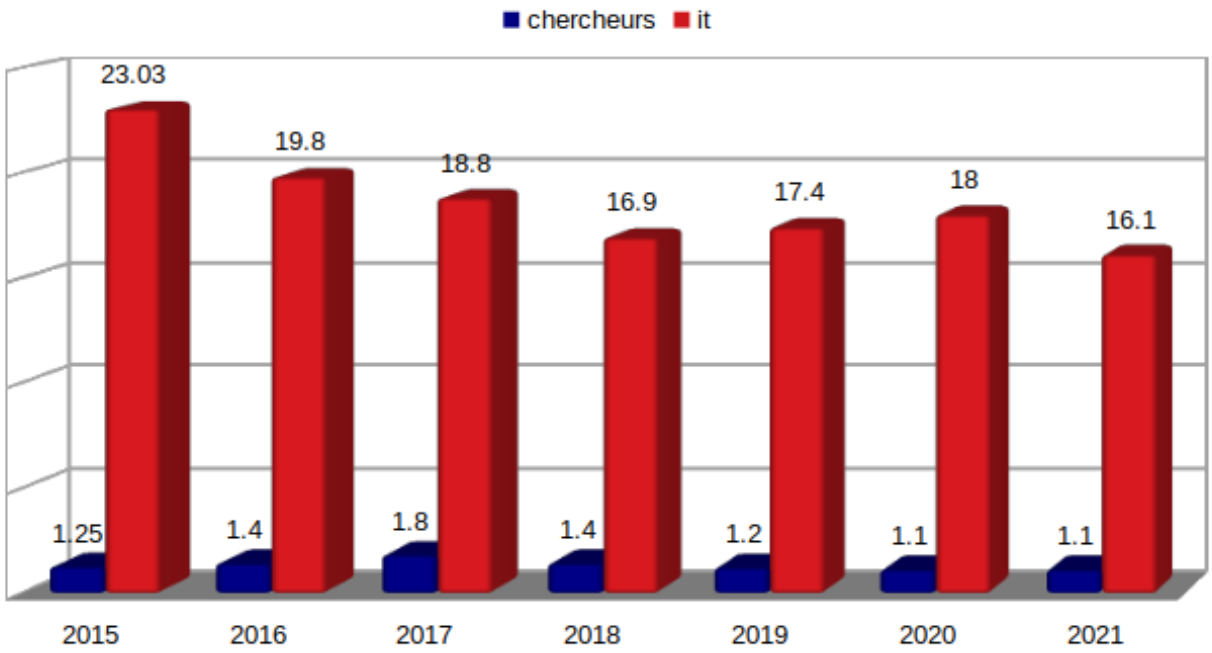

Figure 2: FTE from IN2P3 in LCG-France for physicists (blue) and engineers (red)

Networking

The availability of reliable, high capacity networks is at the core of today's experiment computing models and HL-LHC models will even more rely on networks. WLCG uses two dedicated networks:

- LHCOPN composed of dedicated circuits between CERN and the Tier 1
- LHCOne implemented as overlay circuits, to which CERN, the Tier 1 and a large fraction of the Tier 2 are connected. In the past few years non LHC experiments have joined.

In France these two networks are provided by RENATER. CC-IN2P3 have redundant (North path and South path) connectivity at 100 Gb/s for LHCOPN and LHCOne. RENATER had initially experienced lots of difficulties to implement LHCone due to problem with hardware from vendor that won the bid but the situation has been very stable in the past few years. As a consequence we are also late in increasing the capacities of connections to several Tier 2 sites (CPPM, LPC, ...)[2] and the SUBATECH bandwidth is reduced. The Tier 2 data centres are connected to LHCone via a campus or regional

2 The LPNHE connectivity has been upgraded in May.

network to the nearest RENATER Point-of-Presence with the exception of LAPP that is connected to CC-IN2P3 via the regional network operator Amplivia.

The traffic in the HL-LHC era is expected to be much higher than the capacities installed. In order to ramp-up during Run 3 and to follow the programmed Data Challenges (see later), we have produced a document specifying our needs in the next years to RENATER.

RENATER currently has severe financial difficulties. It is unclear as of now if the entire upgrade program for HL-LHC can be achieved in a timely manner.

Main results

Due to space constraints we will concentrate here on the main overall results of LCG-France and not go in detailed descriptions site by site. Results from the DOMA project are listed in the next section.

Availability and reliability

Availability and reliability metrics are indication of how often a site is up and running and how often it is in unplanned downtime. These are important metrics of the quality of service of a site and as such are part of the MoU agreement between WLCG and Tier 1 and Tier 2 sites. French sites are among the most reliable sites. Some unplanned downtime occurred due to external causes (flood, fiber cut during construction work,…) but some are due to ageing data centre infrastructure like cooling systems.

LCG-France contributions to WLCG

Overall the budget has proven to be adequate to provide computing capacities in steady overall growth year after year. As explained above, Tier 2 sites must find financial support for 30% of the cost of renewal, as weel as for infrastructure maintenance and growth. A sizeable growth is in practice achievable only if the site is part of a project that brings external funding, mostly LABEX, CPER and FEDER (CPPM, IPHC, SUBATECH)[3], local projects and/or hosting (LAPP and, to a lesser extend, LPSC). On average over this represents 250-300k€ invested in hardware per year. Some of the site are hosted by an institution that covers electrical expenses, saving about 150k€ on budget.

A global illustration of the success of the project is to show the overall share of the France with respect to the overall WLCG effort. The next four figures shows the evolution of the shares of CC-IN2P3 resources in CPU, disk and tape for the four experiments compared to all Tier 1 sites.

3 IJCLAB (IPNO and LAL at that time) has benefited from CPER funds to build the current data centre.

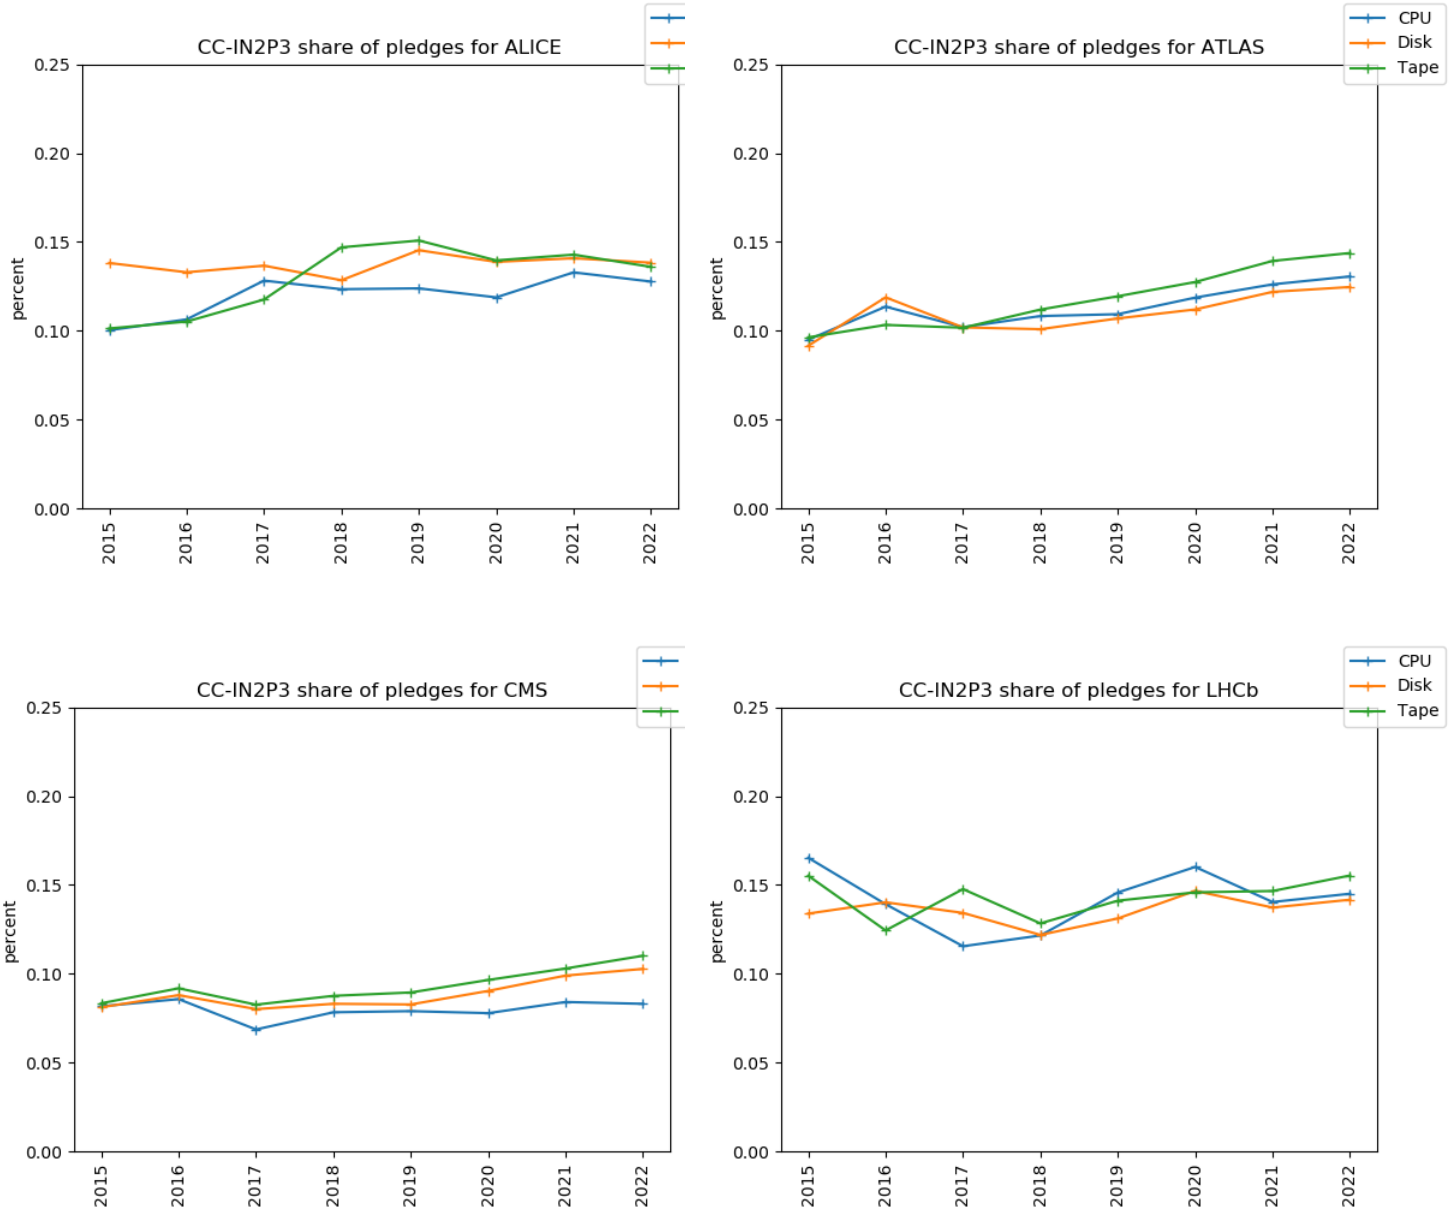

One can see that we have maintained our share of overall Tier 1 resources at a level compatible with the expectations listed in the “Protocole d’Accord”. The next four figures also shows that the same applies to Tier 2 sites. Each Tier 2 site is in control of the share of its resources between experiments.

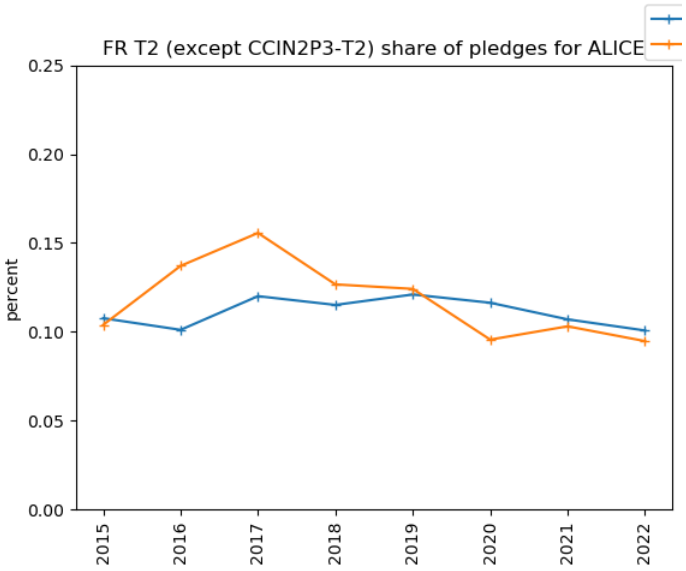
FR T2 (except CCIN2P3-T2) share of pledges for ALICE
percent
0.25
0.20
0.15
0.10
0.05
0.00
2015
2016
2017
2018
2019
2020
2021
2022

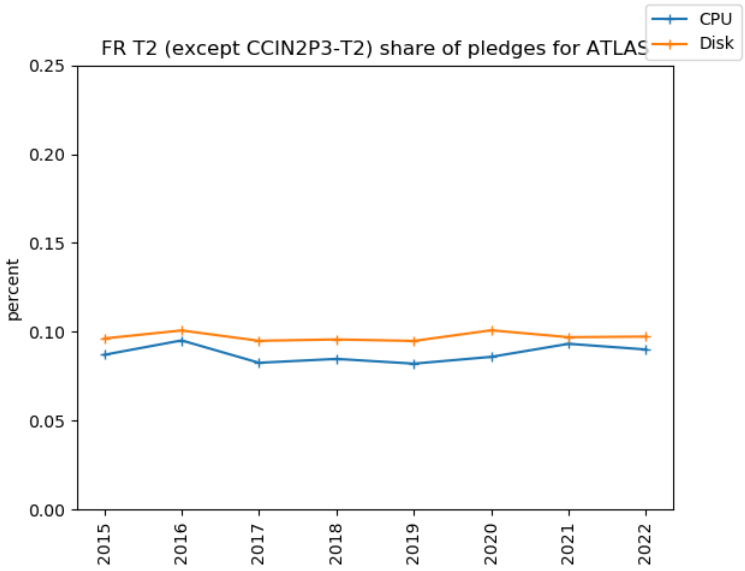
CPU
Disk
FR T2 (except CCIN2P3-T2) share of pledges for ATLAS
percent
0.25
0.20
0.15
0.10
0.05
0.00
2015
2016
2017
2018
2019
2020
2021
2022

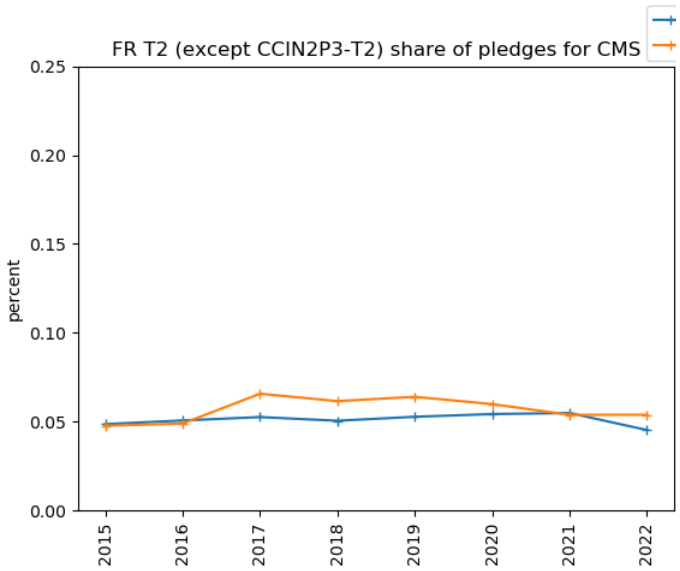
FR T2 (except CCIN2P3-T2) share of pledges for CMS
percent
0.25
0.20
0.15
0.10
0.05
0.00
2015
2016
2017
2018
2019
2020
2021
2022

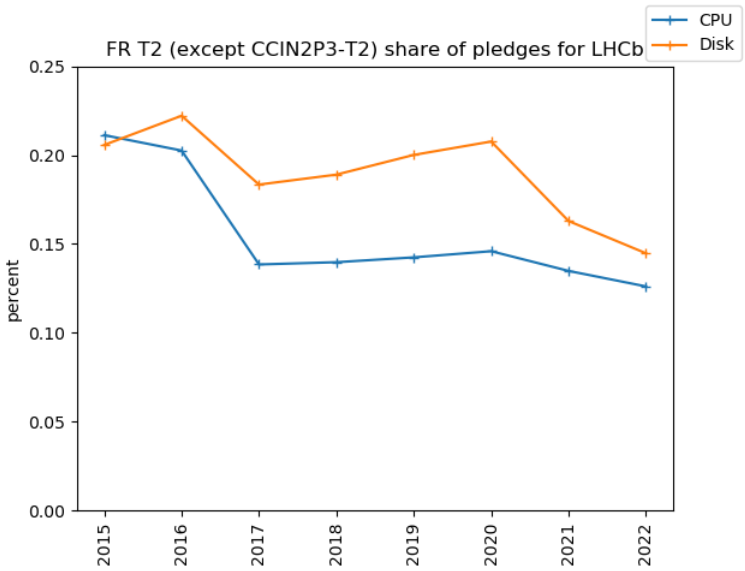
CPU
Disk
FR T2 (except CCIN2P3-T2) share of pledges for LHCb
percent
0.25
0.20
0.15
0.10
0.05
0.00
2015
2016
2017
2018
2019
2020
2021
2022

Working Group participations

We concentrate here only on the R&D efforts (WLCG, DOMA) with significant French contributions. Due to limited person-power, we could not participate actively in all working groups. We also list French participations in other working groups related to WLCG activities, some purely on operations and some with a mix of operations and R&D.

System Performance and Cost Modelling Working Group

The goal of the working group was to evaluate the requirements (memory, CPU, storage) of each major HEP workflow (generation, simulation, digitisation, reconstruction) and to estimate the Total Cost of Ownership of the computing infrastructure so as to develop a model to evaluate the cost of a given computing model for HL-LHC.

A simulator was produced, initially based on the CMS simulator that was further refined and made more generic. Several workflow were measured. The actual cost of computing hardware, infrastructure and running cost (in particular electricity) was found to vary significantly from one country to the other.

Two of the main contributors of the group were French and made presentation at international conferences (WLCG workshop 2018 and CHEP 2019), and several other French people made contributions.

DOMA – ACCESS

During the first phase of DOMA project the ACCESS working group was charged at exploring solutions of data distribution and access. One of the convenors was French. The idea of a Data Lake[4], a mix of data storage sites and data processing sites (accessing data remotely, possibly through caches), emerged. The storage being identified as the most administratively heavy task compared to processing, it was thought that CPU-only sites would save a significant amount of person-power. The group studied in particular the performance of applications with local, remote and remote-with-cache data access.

The main French contribution was the implementation of a test bed processing infrastructure called ALPAMED federating storage and computing from four sites with a single entry point. That in particular implied that a given job can process data from any of the sites. ALPAMED was included in the ATLAS and ESCAPE computing. Initial tests were made with specific jobs in order in particular to measure the efficiency of specific applications when the data is accessed remotely. ALPAMED was then included as an ATLAS production site. This work was presented in an international conference (CHEP 2019).

4 Unfortunately Data Lake in this context differ from the definition used more broadly in computing.

Even if implementation, just starting, was not done during the DOMA-ACCESS group time, the evolution of the storage of the sites under the GRIF federation follows a similar philosophy: unifying the storage of four sites with a technology that will present the geographically separated storage hardware as a single entity from the point of view of the user.

DOMA – Third Party Copy

Under this relatively cryptic name, the working group covered the activities related to the protocol (providing the new Third Party Copy feature) and authentication - authorisation for data transfer. One of the first task was to deal with the obsolescence of a particular toolkit for data transfer and its replacement by a more standard-based tool (gridftp to http-based transfer). The second task was the replacement of the authentication – authorisation system. This work is still ongoing in the second phase of DOMA and its Bulk Data Transfer group. People from CC-IN2P3 actively participated to DOMA-TPC.

Archiving – Data Carousel – Tape Challenges

CC-IN2P3 is a long standing member of the HEPiX “Archival” working group but during Run 2 of the LHC the computing models have shifted toward a more dynamic use of tapes in data processing activities. In particular, processing using a “Data Carousel” model (where data is staged from tape to disk buffers in order to be processed and quickly replaced by the next chunk of data from tape) has been commissioned and then integrated as standard workflows. This required a significant amount of work to optimise the tape systems for this kind of much more dynamic workflows, in coordination within WLCG.

In its second phase, the WLCG DOMA activities have switched from pure R&D to implementation of some of the solutions which lead to the definition of “Challenges” as demonstrators and stress tests, one of which being the Tape Challenge as a follow-up to the Data Carousel activities. So far two Challenges have been organised, in the autumn of 2021 and spring of 2022, with the specific goal of testing the readiness of the tape systems in Tier 0 and Tier 1 for Run 3. CC-IN2P3 has shown good performances, exceeding the requirements. Further Challenges will be organised in the coming years, raising progressively the bar to build the infrastructure for HL-LHC.

Data Challenge

Similarly to the Tape Challenge, DOMA is organising bulk data transfer challenges in order to test networks and disk storage systems. The first one took place in autumn 2021 and involved CC-IN2P3 and several French Tier 2 with good performances. As for the Tape Challenge, further Challenges are foreseen, roughly every other year, with increasing goals to pave the way towards an HL-LHC ready infrastructure. This schedule has driven in particular our requests for upgrade of the French network infrastructure for WLCG mostly operated by RENATER.

Benchmarking and accounting

Properly measuring the performances of the hardware in terms of processing power and translating it into an agreed-upon unit is necessary both for accounting of each site contributions but also for them to pledge a given level of resources and for experiments to express their need. Several years ago, a set of benchmarks based on the SPEC generic toolkit were defined and turned into a global performance measurement named HS06.

Along the years, some HEP workflows have shown different scaling than the SPEC benchmark. The Benchmarking Working Group has been charged to define and implement a different solution that would reflect better the actual experiment workflows and that could be extended to measure the performances for non CPU-only hardware, e.g. CPU+GPU. The working group has developed a toolkit based on containers encapsulating actual workflows from the LHC experiments (generation, simulation, digitisation and reconstruction). The working group is now in the process of defining the set of tools that would be combined in a new official overall performance index named HEPSCORE. A French person is a member of the steering group and other French people have participated in the implementation and test of the toolkit.

DPM Collaboration

DPM is the storage solution used by most French Tier 2. It’s a product developed mostly by CERN people around which a collaboration was formed. The French contribution consisted into providing test beds for testing new versions of the product and migrations. Our two representatives to the Collaboration Board have coordinated the writing of DPM Community Whitepaper of 2019-2020. The development for this product has stopped in 2020.

Other working groups

We have participated in other working groups, some of them formed for a specific task like transition from one solution (or version of a production) to another:

- Middlware Readiness WG
- Network and Transfer Metrics WG
- CREAM CE migration task force

and created a Perfsonar LCG-France Task Force (following a significant effort in investigating network monitoring).

French Context

France-Grilles

First and foremost, our natural partner project in France is the GIS France-Grilles, acting as the French National Grid Infrastructure against EGI. WLCG has been using EGI solutions and services since its inception (although this is less and less the case now) and there is a large overlap between France-Grilles sites and LCG-France sites.

We have organised one common workshop and had France-Grilles presentations during LCG-France workshops. We have encouraged participation to the Journées SUCCES (co-organised by France-Grilles, Groupe Calcul and “coordination mésocentres”) that became JCAD (Journées Calcul et Données, co-organised by the same partners with the addition of GENCI), where we have presented LCG-France or DOMA topics four times.

The participation to Journées SUCCES and JCAD was also an attempt to reach out to other communities in the French computing landscape, in particular the mésocentres[5] (some of our sites are mésocentres) and the computer scientists. At one of DOMA workshop we had people from CNES, INSERM and Nantes. We had contacts with INRIA scientists but this did not turn into a concrete collaboration.

As mentioned earlier, the France-Grilles operations meeting and LCG-France technical meeting have merged this year since there was a large overlap between the two (as LCG-France technical meetings include a report on WLCG operations).

High-Performance Computing

As the LHC computing is constantly starving for resources, there have been efforts to enable access to HPC centres from all over the world for several years. One of the difficulties is that the access restrictions and technical situations vary a lot from one HPC site to another (operation of border services, internal and external network access and bandwidth, allocation policies, etc…).

A first attempt to use the IDRIS machines was made with people from CC-IN2P3 and ATLAS. There were technical difficulties but the demonstration of feasibility was made, although the service would have required permamnent human interventions. Another hurdle was to accommodate the policy of resource requests which were designed for larger computations over a short time period while we would like a constant minimum allocation during the entire year to make the human investment worthwhile.

CC-IN2P3 leads for FITS project, which involve IDRIS and GENCI (the actual proprietary of IDRIS supercomputer), aiming to favor usage of both facilities by relevant users and the possibility for long duration requests. We hope that this will turn into an opportunity for LCG-France to give WLCG access to those resources.

5 Mésocentres are regional HPC resource centres (hardware resources, support team and developers).

In the last couple of years, France has express the intent to host one of the PRACE hexascale HPC computers. In particular a working group has been put together in order to reach out to all the communities that could use such a facility and assess their readiness to use a heavily GPU based architecture. IN2P3 is one of the communities represented there and LHC computing one of the topics. The requirements of LHC computing is quite different from the usual requirements of an HPC machine, in particular in terms of wide area network access at high speed. One specific project detailed in the working group was the usage of such a facility for LHCb trigger processing: it would require Tb/s networks. The working groups were also the opportunity to reach out to other communities in need of high data volume processing and in particular remote processing. The working group is now closed and the report being written but we hope that the discussions, in particular around distributed data, will continue as promised in a different forum.

Feuille de route du numérique

In the last few years, the French ministry of Higher Education Research and Innovation and CNRS have developed roadmaps for their computing activities, from hosting services or hardware to scientific computing. One item of potential direct consequence for LCG-France was the goal of drastically reducing the number of Data Centres (DC). The arguments were that small local DC were proliferating, not very energy efficient and in constant demand of human resources for their operation. The main idea was that a single DC per French region would be labelled (with possible exceptions e.g. for larger or more populated regions like Paris – Île de France) and that national agencies would then only fund computing projects with hardware hosted in a labelled DC. That would directly affects most of our external sources of funding, in particular CPER.

This caused great concern in our community since, except for CC-IN2P3 (already considered as a National Centre), essentially none of our DC would likely be labelled. We wrote a short document summarising the problems that a move to a different DC would cause. For example we argued that Tier 2 sites have signed an MoU with requirements on the availability of the site which implies that quick and easy access to the DC is needed. We estimated that Tier 2 admins would need physical access to the hardware once or twice a month at least for repair and installation of hardware. While this could be doable if the DC is located in the same town, it would be problematic if it is located dozen of kilometres away (or more, given the size of regions in France).

The process of site selection has since been very slow and the “one DC per region” not strictly followed. There are labelled DC in Marseille and Strasbourg, so not too far away from our DC. In Marseille, the DC is even one that was involved in common projects with the CPPM site. The IPHC site anticipates having to host at least part of their new hardware to the labelled DC. The situation is less clear for the other LCG-France Tier 2 sites.

International Context

As a production infrastructure within WLCG, LCG-France is constrained by the overall model and strategies and part of the technical solutions decided at that level although it is very much part of the decision process. The overall weight of CERN in the decision process is very large although in practice the management seek for a wide consensus. So far no controversial decision have been taken[6] with the possible exception of the withdrawal of CERN from development of the DPM storage solution, effectively bringing it to an end.

Other HEP and non HEP experiments will be faced with large volume of data, and it seems important that the LHC experiments and WLCG share their experience and, as much as possible, tools. In many countries, these experiments will be hosted in the same data centres (e.g. Tier 1) and having to support different tools for different experiments would be very costly human wise. One example is that the LHCOne network, initially designed for the LHC experiments, has been opened to more and more experiments that use the same data centres (Belle II, Pierre Auger Observatory, NOvA, XENON, JUNO), perhaps a bit too enthusiastically as it becomes difficult to identify the source of traffic in case of congestion. WLCG and CERN have also started high level discussions with SKA. Likewise, some EU projects like XDC and ESCAPE where French labs from WLCG were/are involved use tools developed in WLCG. The ALPAMED test bed has been used in WLCG for ATLAS and also in ESCAPE. We anticipate that such tools would also be used in other projects related to EOSC.

Perspectives and Risks

We are now in the process of defining a new agreement (“Protocole d’Accord”) between the IN2P3 direction, sites and experiments project leaders for the next 4-5 years. One of the main discussion points is how the budget is shared between experiments at CC-IN2P3: the current fractions are based in the IN2P3 participation in LHC experiments but the needs of the experiments have changed for Run 3 (with a huge increased of request by LHCb due to their Phase I upgrade) and will change again approaching HL-LHC (where ATLAS and CMS needs will dominate).

As already explained two of our Tier 2 already announced that they would stop by the end of Run 3 so they already do not receive funding for hardware renewal and won’t be part of the new agreement. For the other Tier 2 sites, continued participation make sense if they intend to run during the HL-LHC era. Most laboratory directors indicated that they rely on external funding for their Tier 2 sites and that funding is not guaranteed beyond the next few years, which makes them hesitant to commit.

The current funding model rely on the budget coming from CNRS IR for support to hardware renewal at Tier 2 and renewal and growth at CC-IN2P3. The target budget has been raised by 10% a few years

6 At the time of writing, no decision has been taken year on possible exclusion of Russia and Belarus.

ago but an additional increase would probably be necessary to follow the huge increase in requested resources for HL-LHC.

In addition to budget concerns, there are other sources of uncertainties for the next few years and even more at medium term like the start of HL-LHC. First, the DOMA R&D did not uncover any technical solution that would significantly reduce the cost of disk storage. It seems the baseline now for ATLAS and CMS is to achieve reduction of the analysis format tier to ~ 10kB per event in order to be able to store it permanently on disk. Second, the trend of constant reduction of cost of a unit of disk or CPU is no longer guaranteed or at least the reduction is less that what was common a few years ago. On a short term basis, it is expected that the current “silicon crisis” will continue to cause delays in delivery and shortage of supplies.

The ATLAS and CMS computing models rely heavily on the use of tapes as a cheaper storage than disk. They are also attractive in the current context of steeply rising cost of electricity. However, prices are driven by the global market and the market for tape usage in enterprise computing is shrinking (or at least not strongly growing). As a consequence key vendors have left the tape market so we are at risk of relying on single vendors with less investment in R&D and no incentive for price drops.

Another uncertainty we have already mentioned is due to the financial difficulties of RENATER, the French Research and Education Network provider. It is currently unclear how the cost of upgrading the capacities to the level of our anticipated needs will be funded.

As already mentioned, the success of our projects is built on the efforts and dedication of our people. We indicated that some sites have reached what we consider the minimum level of effort to maintain a healthy Tier 2 site. In addition we know of several retirements of key site administrators in the next few year. We must carefully plan for their replacement and this needs strong support for the laboratory directions and from the IN2P3 direction.

Bibliography

2: http://wlcg.web.cern.ch
3: http://wlcg.web.cern.ch/collaboration/mou
4: I. Bird, Computing for the Large Hadron Collider, Annual Review of Nuclear and Particle Science, Vol. 61: 99-118 (2011)
5: http://lcg.in2p3.fr/wiki
6: LCG-France - Définition du périmètre du projet (2011), ATRIUM-75646 (accès restreint)
7: https://indico.in2p3.fr/category/64/
8: LCG-France - Protocole d'accord pour le financement des sites IN2P3 de LCG France de 2018 à 2022 (protected)