Key Facts and Figures – CERN Data Centre

General Information

Websites

- Latest numbers: http://go.web.cern.ch/go/datacentrebynumbers and https://monit-grafana.cern.ch/d/000000473/it-overview?orgId=1 (special access rights are required)
- How to use the Data Centre visit point (please sign in to access that page): http://www.cern.ch/information-technology/about/computer-centre/visits/visitpoint
- The general content about computing is available on the CERN website at this url: https://home.cern/fr/science/computing

Key facts and numbers

	CERN Data Centre	
Servers	9 500	≃ 10 K
Processor cores	444 100	≃ 450 K
Disks	103 500	≃ 100 K (units)
Tape Cartridges	41 540 *(about 424 PB used)*	≃ 40 K (units)
Virtual Machines	13 100	≃ 13 K

Short overview

The CERN Data Centre is the heart of CERN's entire scientific, administrative and computing infrastructure. All services, including e-mail, scientific data management and videoconferencing use equipment based in the Data Centre.

Most of the CERN IT equipment is hosted in the Meyrin Data Centre. However, a second network hub has been inaugurated in 2017 and is located in Prévessin (https://home.cern/news/news/computing/inauguration-second-cern-network-hub). Two containers have also been lent by LHCb (https://home.cern/news/news/computing/alice-and-lhcb-upgrade-their-data-centres) to the IT department until Long Shutdown 3 (LS3) which may start at the end of 2025.

About 450 000 processor cores and 10 000 servers run 24/7.

As of the beginning of June 2022, there are about 13 000 virtual machines currently running in the CERN Data Centre.

CERN IT runs a private OpenStack cloud on about 9 000 physical nodes with about 482 000 cores, provisioned to the users either via virtual machines or directly as bare metal servers. The use of overcommit as a measure to make efficient use of the available physical

resources can, in some circumstances, result in a core number which is higher than the number of physical processor cores in the data centre mentioned on page 1 of this document. (https://monit-grafana.cern.ch/d/8f4TgzF7z/cern-openstack-overview?orgId=3)

In October 2018, the CERN Data Centre passed the milestone of 300 petabytes of data permanently archived in its tape libraries. At the end of 2018, 330 PB of data were permanently archived on tapes in the CERN Data Centre.

New record in November 2018 for data taking over a single month: 15.8 petabytes of data (from all sources) were written on tape that month.

Within one year, when the LHC is running, more than one exabyte (the equivalent to 1000 petabytes) of data is being accessed (read or written).

The LHC data are aggregated in the CERN Data Centre, where initial data reconstruction is performed, and a copy is archived to long-term tape storage. Another copy is sent to several large-scale data centres around the world. Subsequently hundreds of thousands of computers from around the world come into action: harnessed in a distributed computing service, they form the Worldwide LHC Computing Grid (WLCG).

The new CERN Data Centre in Prévessin will come online during the final quarter of 2023 (https://home.cern/news/news/computing/building-work-cerns-new-data-centre-prevessin-begins), while the CERN Data Centre in Meyrin will continue to run in parallel, with a particular focus on data storage.

The LHC experiments plan to collect more data during LHC Run 3 than they did in the first two runs combined.

It means that the computing challenge during LS2 was to get ready to store and analyse more than 600 petabytes of data (600 million gigabytes), equivalent to over 20 000 years of 24/7 HD video recording.

Computing presents huge challenges for CERN in both the near and far terms. When the High-Luminosity LHC, HL-LHC, comes online in 2029, the storage and processing requirements are expected to be ten times greater than today.

STORAGE:
Websites:
www.cern.ch/eos
www.cern.ch/cta
www.cern.ch/castor
www.cern.ch/fts
https://monit-grafana.cern.ch/d/kEbF3aZGz/namespace-castor-cta?orgId=29
https://filer-carbon.cern.ch/grafana/d/000000036/eos-control-tower?orgId=1&from=now-12h&to=now-1h&refresh=5m&var-instance=All

Magnetic tapes are used as the main long-term storage medium. The tapes are stored in tape libraries and are retrieved by robotic arms. At the moment there are 5 tape libraries in building 513 and 2 tapes libraries in building 613.

- In June 2022, we had 424 PB of data on tapes.
- As of June 2022, the CERN Data Centre had a total storage capacity of 634 PB on disks. Around 42%, i.e. 369 PB, were currently used, leaving enough space for new data expected with the upcoming begining of Run 3. However, as disks are less reliable than tapes (a thousand times less, 30 disks fail each week in the data centre), we always copy twice the data on disk so that if a disk fails we do not lose data => so 'only' 184.5 PB of 'real' data is on disks at the moment.
- The CERN storage system, EOS, was created for the extreme LHC computing requirements. EOS in 2020 has served 2.5 exabyte of physics data to the experiments and at the end of 2019, EOS instances at CERN exceeded five billions files, matching the exceptional performances of the LHC machine and experiments.
- The CERN Tape Archive (CTA) entered production in 2020, progressively replacing CASTOR before LHC Run 3 for tape storage: https://home.cern/news/news/computing/cern-data-storage-gets-ready-run-3
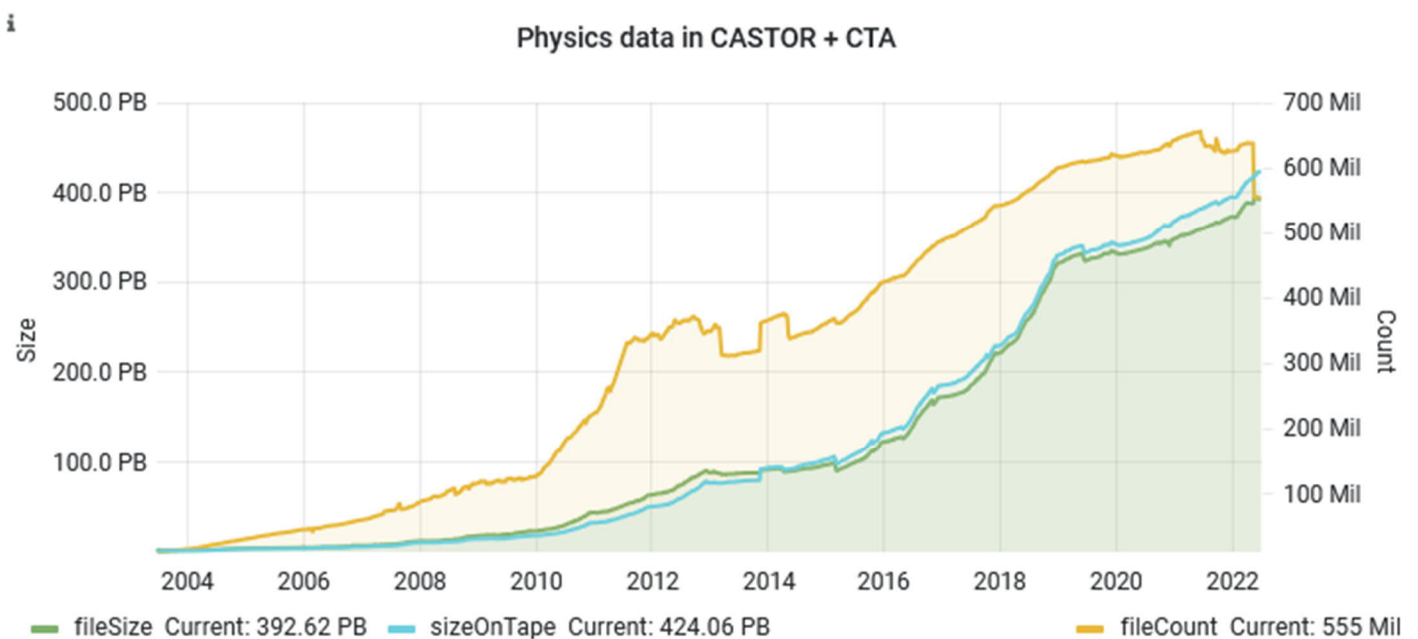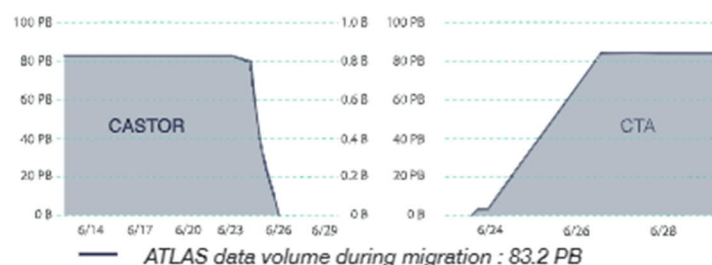
On 29 June 2020, the CERN Tape Archive (CTA) officially entered production after 83 PB of ATLAS data initially stored in CASTOR were successfully migrated to CTA.

- FTS, the File Transfer Service, has moved in one year a billion file in the WLCG for around 1.3 ExaByte of data
- The amount of data recorded on tape at CERN is steadily increasing over time, with this trend accelerating over time (+53 % data stored in 2018 compared to what had been accumulated between the start of CERN's operations and the end of 2017 for instance).

- 15.8 petabytes of data (from all sources) and 13.47 PB of LHC data alone were written to tape in November 2018, two new records.

- In 2018, over 115 PB of new data (out of which about 88 PB of LHC data) were written to tape. In addition, over 140 PB of data were migrated ("repacked") to higher-density cartridges.

- In 2017, 40 PB of LHC data were recorded at the CERN Data Centre on tapes. In total 72 PB of data were recorded on tapes (from LHC and non LHC experiments, etc.) with a striking 12.3 PB in the month of July alone.

- In 2016, more than 49 petabytes of LHC data were recorded at the CERN Data Centre on tapes. In total, 73 PB of data were recorded on tapes (from LHC and non LHC experiments, etc.) with a striking 11 PB in the month of July alone.

- While tape drives are faster than disk drives, latency for accessing tape data is relatively high. It takes about 1-3 minutes from tape being located and mounted on a tape drive before data can be read. Often physicists need to access the latest data immediately, so it is also made available on disk servers, where latency delays are significantly lower.

- There are currently about 100 000 disks in the CERN Data Centre with a capacity of about 634 PB. 25% of the 100 000 disk drives are SSD ones (but this represents less than 10% of the capacity itself). For the high capacity storage used for the physics data in CASTOR, CTA, and EOS, CERN uses normal mechanical (spinning) disk drives since the price per terabyte is significantly lower (5-10x).

- On 9 October 2018, the CERN Data Centre passed the milestone of 300 petabytes of data permanently archived in its tape libraries. Where do these data come from? Particles collide in the Large Hadron Collider (LHC) detectors approximately 1 billion times per second, generating about one petabyte of collision data per second. The detectors can be compared to digital cameras with 100 million electronics channels that would be taking over 40 million pictures per second (which currently corresponds to 1 billion proton-proton interactions per second).

- However, such quantities of data are impossible for current computing systems to record and they are hence filtered by the experiments, keeping only the most "interesting" ones. The filtered LHC data are then aggregated in the CERN Data Centre, where initial data reconstruction is performed, and where a copy is archived to long-term tape storage. Even after the drastic data reduction performed by the experiments, the CERN Data Centre processes on average one petabyte of data per day. If we keep the digital camera analogy, that would mean that we only keep a 1000 pictures per second out of the 40 million we initially had.

ELECTRICITY:

Websites:

https://monit-grafana.cern.ch/d/000000265/data-center-infrastructure-monitoring-homepage?orgId=12

Key facts and numbers:

- 3 megawatt computing power consumption (cooling not included and represents roughly one additional megawatt) from a maximum of 4 megawatt. The two containers lent by LHCb to the IT department are on their side having a computing power consumption of about 0.5 megawatt.

- The electrical infrastructure is a vital element of the data centre. Strategies for increasing power efficiency are permanently investigated to be able to maximise the computing power

serving the CERN's infrastructure and scientific programme whilst staying within the 3.5 megawatt electrical capacity envelope available in the data centre.

- In the case of a major electrical cut, Uninterruptible Power Supplies (UPS) provide time for all the non-critical systems to be properly shutdown and a combination of UPS systems and diesel generators ensure that the critical services keep working.

COOLING:

- About 1 megawatt is dedicated to the cooling of the data centre.
- Main Machine room: chilled air via silver ducts, into false floor and then into closed server aisles. We also use water-cooled racks in some areas (basement: 'vault').

Efficient cooling is a key element of the data centre. When the outside temperature is low, air can be used to cool the servers, otherwise chiller systems are used to cool the data centre air. Cold air is distributed via the silver ducts on the sides of the data centre room. It then goes under the false floor and into the closed server aisles through the perforated floor tiles, to be drawn finally through the servers to cool them. Some racks are water-cooled using passive heat exchangers in their rear doors, providing a higher cooling capacity per rack.

Articles recently published on this matter:

https://home.cern/news/news/cern/environmental-awareness-challenges-cerns-it-infrastructure

https://home.cern/news/news/cern/environmental-awareness-reducing-energy-consumption-cerns-it-infrastructure

WLCG – Worldwide LHC Computing Grid

Websites:
       www.cern.ch/wlcg
       www.cern.ch/wlcg-public
Documents:
       Cf. dedicated section in the yearly IT Department groups and activities reports:
       https://cds.cern.ch/record/2631468

Key facts and numbers:

- More than 170 data centres in more than 42 countries
- Close to 1.4 million CPU cores and 1.5 Exabytes of storage available to the LHC experiments
- Allows more than 12 000 physicists to access LHC data
- Over 2 million tasks daily
- The total amount of storage used globally is around 1.5 Exabytes: 40% is provided by disks and 60% is provided by tapes.

Evolution of the overall core processor time delivered by WLCG

As the graph shows, the total CPU time delivered by WLCG (expressed in billions of HS06 hours per month, HS06 being the HEP-wide benchmark for measuring CPU performance) is continually increasing. In 2020, WLCG pooled the computing resources of about one million cores.
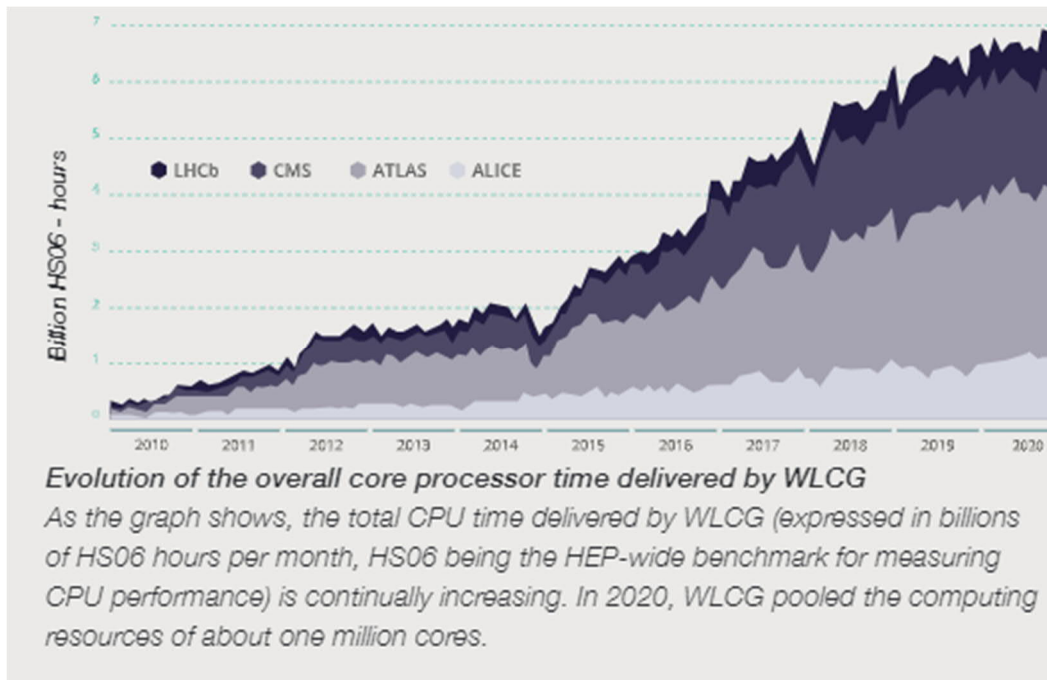
Short overview:

- Today, the Worldwide LHC Computing Grid (WLCG) combines the computing resources from more than 170 data centres in more than 42 countries, producing a massive distributed computing infrastructure that provides more than 12 000 physicists around the world with near real-time access to LHC data, and the power to process it. About 1.4 million CPU cores and 1.5 Exabytes of storage available to the LHC experiments.
The Worldwide LHC Computing Grid runs over 2 million tasks per day.
These numbers will increase as time goes on and as computing resources and new technologies become ever more available across the world.
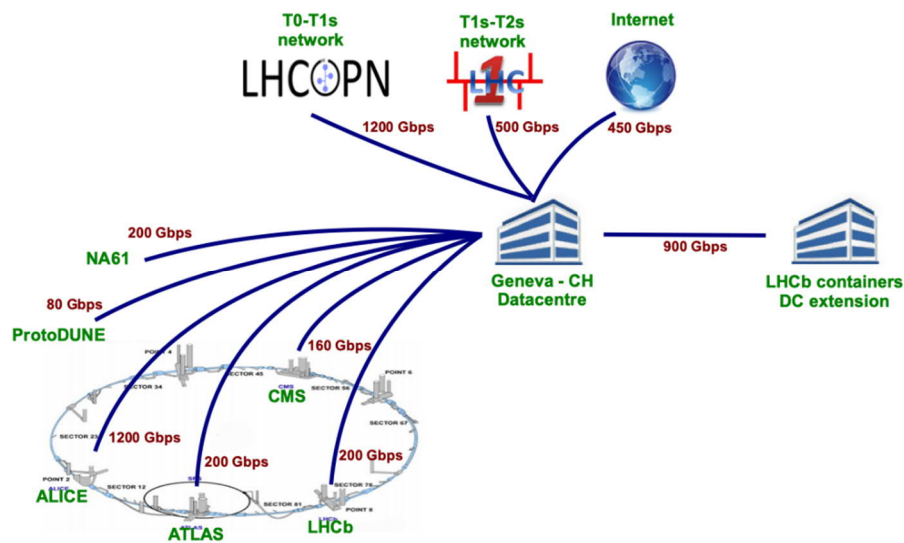
NETWORK:

Webpage:

https://cixp.net
https://twiki.cern.ch/twiki/bin/view/LHCOPN/LhcopnStats
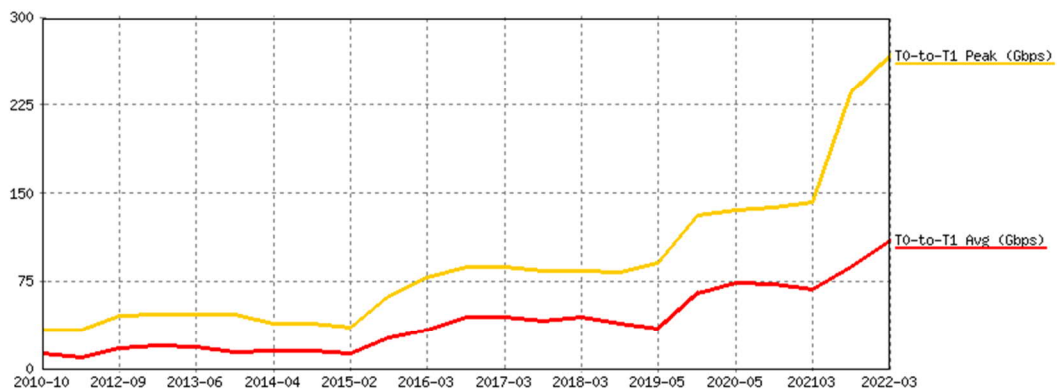
Key facts and numbers:

- Over 50,000 km of optical fibre providing network connectivity throughout the CERN sites.
- Network connections for LHC data:

LHCOPN long term traffic statistics

LHCOPN
LHCOPN network traffic from CERN Tier0 to all the aggregated Tier1s

Run1: 2010-12 LS1:2013-14 Run2: 2014-2018 LS2: 2019-21 Run3: 2022
Y-Axis: Gbps (Giga bit per second)
Out: direction Tier0 to all Tier1s
Avg: average network bandwidth on the previous 12 months
Peak maximum peak network bandwidth on the previous 12 months

- The networking infrastructure went through significant upgrades during LS2 to get ready for Run 3, and the data transfer rates around the globe reached new peak rates – up to 33 gigabytes per second (267 gigabits per second) rates, around a factor of eight higher than what had been typical during the LHC Run 1.

- CIXP: https://cixp.net/, The CERN Internet eXchange Point (CIXP) is a carrier-neutral exchange point based in the CERN Data Centre. Our partners are telecom operators and ISPs in Switzerland and France, as well as national and international research network operators. The service is provided jointly by CERN and Equinix's data-centres in Geneva and Zurich.

- A second network hub was inaugurated on the Prévessin site in 2017: https://home.cern/news/news/computing/inauguration-second-cern-network-hub).

- CERN played a central role in the history of the internet's development: in 1991, 80% of the internet capacity in Europe for international traffic was installed in the CERN Data Centre. To know more, please check: https://home.cern/cern-people/opinion/2013/06/how-internet-came-cern

Data Preservation:

Cf. computing chapter of the CERN annual report 2016: https://e-publishing.cern.ch/index.php/Annual_Report/article/view/496/354
As well as the dedicated section in the yearly IT Department groups and activities reports: https://cds.cern.ch/record/2631468

Open Source:

Cf. computing chapter of the CERN annual report 2018: https://e-publishing.cern.ch/index.php/Annual_Report/article/view/894/719 and the CDA group report in the yearly IT Department groups and activities reports: https://cds.cern.ch/record/2631468

LHC@home:
www.cern.ch/lhcathome
Volunteer computing for the LHC

CERN openlab:
www.cern.ch/openlab
CERN openlab is a unique public-private partnership that accelerates the development of cutting-edge solutions for the worldwide LHC community and wider scientific research. Through CERN openlab, CERN collaborates with leading ICT (Information Communication Technologies) companies and research institutes. CERN openlab began a new three-year phase in 2021, with its work split into three main technology areas: exascale technologies, artificial intelligence (AI) technologies and quantum computing.
CERN openlab annual report 2020: https://zenodo.org/record/5793362#.YkrIzihBwuU

CERN Quantum Technology Initiative:
https://quantum.cern/
In 2021, CERN QTI published its first roadmap, which sets out a research programme divided into four main areas: quantum computing and algorithms; quantum theory and simulation; quantum sensing, metrology and materials; and quantum communication and networks. In 2021, CERN QTI established 20 research projects to explore how quantum technologies can best benefit the particle physics research community and beyond.
To know more:
https://cerncourier.com/a/cern-and-quantum-technologies/
https://home.cern/news/news/computing/cern-meets-quantum-technology
https://cds.cern.ch/record/2789149
https://home.cern/news/press-release/knowledge-sharing/cern-quantum-technology-initiative-unveils-strategic-roadmap

EU projects:
www.cern.ch/information-technology/about/projects/eu/eu-funded-projects
List of the current EU projects http://information-technology.web.cern.ch/about/projects/eu/current/projects
Cf. dedicated section in the yearly IT Department groups and activities reports: https://cds.cern.ch/record/2631468

High Performance Computing Collaboration:
Members: CERN; SKAO, the organisation leading the development of the Square Kilometre Array radio-telescope; GÉANT, the pan-European network and services provider for research and education; and PRACE, the Partnership for Advanced Computing in Europe.

To know more:
https://home.cern/news/news/computing/cern-skao-geant-and-prace-collaborate-high-performance-computing
http://cern.ch/go/Nf8R

CERN and experiments' Computing against Covid-19:
https://againstcovid19.cern/articles/cern-and-lhc-experiments-computing-resources-global-research-effort-against-covid-19
https://againstcovid19.cern/actions

UNOSAT:
https://unitar.org/unosat/
CERN's powerful IT infrastructure is useful for fields other than fundamental research. For 15 years, UNOSAT has been using the Laboratory's computing centre infrastructure for the purposes of its humanitarian work.
To know more:
http://home.cern/about/updates/2016/10/unosat-15-years-humanitarian-mapping
http://cds.cern.ch/record/2223516?ln=en
https://home.cern/fr/news/news/knowledge-sharing/unosat-cern-satellite-mapping-good-humanity

SAFER
https://www.safer-trust.org/
In 2021, CERN cofounded a global trust group called SAFER, aimed at protecting research-and-education facilities from cyberthreats.
To know more:
https://home.cern/news/news/computing/computer-security-safer-teamwork

THE PAST:
- *10 panels presenting the history of the data centre will be installed in February 2021 in the corridor (building 49) linking building 31 to building 513 as part of the exhibition circuit. Their content is also available to you here:* https://cds.cern.ch/record/2750600
- *This video gives a brief overview of the history of the Data Centre:* https://videos.cern.ch/record/2731052
- Many documents and resources are available to you on this page: https://information-technology.web.cern.ch/about/cern-computing-history

THE FUTURE:

The LHC experiments plan to collect more data during LHC Run 3 than they did in the first two runs combined. It means that the computing challenge during LS2 was to get ready to store and analyse more than 600 petabytes of data (600 million gigabytes), equivalent to over 20 000 years of 24/7 HD video recording.

A continuing programme of upgrades and consolidations to the LHC and the experiments at CERN will result in hugely increased ICT demands in the coming years, exceeding what is expected to be achievable with a constant investment budget by several factors in both storage and computing capacity. The High-Luminosity LHC, the successor to the LHC, is planned to come online in around 2029. By this time, the experiments will collect 10 times more physics events than before (in the order of exabytes per year) and those events will be more complex to process. Storing and analysing

those events with the available computing hardware investments will be very challenging. In this context, a call for tender will be prepared in 2020 for the creation of a new CERN Data Centre. Several other solutions, listed below, have also been investigated and are developed.

In 2017 the High Energy Physics community produced a [roadmap Community White Paper (arXiv:1712.06982)](#) that explores these software challenges and describes a number of potential mechanisms by which the community can address these problems of capacity and efficiency that will be faced during the next decade.

CERN openlab, the unique public-private partnership through which CERN collaborates with leading ICT companies and other research organisations, has published in 2017 a white paper identifying the major ICT challenges that face CERN and other 'big science' projects in the coming years in order to collaborate with the industry and research institutes to tackle them. In 2019, CERN openlab launched four projects in 2019 related to quantum computing technologies, which offer great potential to provide solutions for CERN's future ICT challenges.

The WLCG team is leading the DOMA project, which aims to prototype the concepts of a "data lake", where data can be streamed on demand to processing centres rather than being pre-placed there. The CERN team is leading a work package, within the EU-funded ESCAPE project, to demonstrate this "data lake" technology for high energy physics (HEP), as well as for astronomy, astroparticle physics and related fields. It started in early 2019 and is already operating a set of prototypes that will form the basis for future development work in collaboration with these sciences.

As the computational challenges posed to HEP by the HL-LHC and the upgrades of the detectors might not be solved entirely through the use of traditional central processing units (CPUs), the LHC experiments, WLCG and CERN openlab also started investigating novel approaches to accommodate the large amount of computation that will be required. They invested in R&D efforts to leverage GPUs for traditional HEP data processing and analysis. The ALICE experiment had already pioneered the use of GPUs for its high-level trigger (HLT) during Run 2. After a preliminary study in 2015, ATLAS resumed investigations into the potential use of GPUs for data reconstruction and analysis. The CMS experiment started R&D that has demonstrated that code which accounts for about one third of the time needed to run the HLT event filtering sequence can be offloaded to GPUs. The LHCb collaboration demonstrated the feasibility of porting on GPUs the first stage of the software dedicated to its newly developed trigger system that is able to decide whether an event contains physics signatures relevant for further processing. GPU resources have also been made available in the CERN Data Centre through a batch system and significantly accelerate certain applications.

A new CERN Data Centre in Prévessin will come online during the final quarter of 2023 (https://home.cern/news/news/computing/building-work-cerns-new-data-centre-prevessin-begins). It will provide computing resources up to a total electrical power requirement of 12 megawatts. These resources will be delivered in three phases. Each phase corresponds to one of the three floors of the new data centre, with the first phase set to run from 2023 to 2025. It will see computing resources requiring up to 4 megawatts of electrical power installed; this is approximately the same as the power of the current CERN Data Centre in Meyrin for computing (excluding cooling). The CERN Data Centre in Meyrin will continue to run in parallel, with a particular focus on data storage.

To know more:
https://home.cern/news/news/computing/building-work-cerns-new-data-centre-prevessin-begins
https://cerncourier.com/time-to-adapt-for-big-data/
https://home.cern/news/news/computing/cern-openlab-tackles-ict-challenges-high-luminosity-lhc

https://home.cern/news/news/computing/lhc-pushing-computing-limits
https://computing-blog.web.cern.ch/2020/02/computing-for-next-generations-accelerator-graphics-processing-units/
https://home.cern/news/news/computing/cern-data-storage-gets-ready-run-3
https://cds.cern.ch/record/2631468 : CERN IT department groups and activities annual reports (2016, 2017, 2018, and 2019)

Thank you to all the people who have produced material that has been included/copy pasted in this document.

Please do not hesitate to modify this document with the track changes feature on and to send it back to Mélissa Gaillard so that the modifications are taken into account in the next version (this document is updated on a monthly basis and is made available on the CERN guides' website, the IT website and as part of the ITMM minutes on a monthly basis).