

Key Facts and Figures – CERN Data Centre

As of 1 December 2019_V1

General Information

Websites

- Latest numbers: <https://cern.ch/go/datacentrebynumbers> and <https://monit-grafana.cern.ch/d/000000024/cern-openstack-overview?orgId=3&refresh=15m> (special access rights are required)
(Or if you are not on CERN site: <http://go.web.cern.ch/go/datacentrebynumbers>)
- 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

	Meyrin Data Centre*	Wigner Extension**	TOTAL
Servers	14 738	1 286	16 026 => 16 K
Processor cores	257 356	20 576	277 932 => ≈ 280 K
Disks	112 284	16 060	128 344 => ≈ 130 K (units) (about 280 PB capacity but data is copied twice on disks)
Tape Cartridges			29 500 => ≈ 30 K (units) (about 400 PB capacity)
Virtual Machines			30 849 => 31 K

*Including two containers lent by LHCb to the IT department until LS3 and 2nd network hub.

**The CERN IT equipment currently at the Wigner Data Centre is progressively brought back to the Meyrin Data Centre. This transfer will be finalised by the end of 2019. In the meantime the numbers given above are more up to date than the ones here <https://cern.ch/go/datacentrebynumbers>.

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.

A remote extension of the CERN Data Centre used to be hosted at the Wigner Research Centre for Physics in Budapest, Hungary, 1200 km away. The CERN IT equipment currently at the Wigner Data Centre is progressively brought back to the Meyrin Data Centre and to the two containers lent by LHCb to the IT department until Long Shutdown 3 (LS3).

About 280 000 processor cores and 16 000 servers run 24/7.

As of the beginning of December 2019, there are over 31 000 virtual machines currently running in the CERN Data Centre.

CERN runs a private OpenStack cloud with about 260 000 cores (mix of hyper threaded cores and physical cores so this number can be, in some circumstances, different from the number of physical processors cores in the data centre mentioned on page 1 of this document) and more than 7000 hypervisors. (<https://monit-grafana.cern.ch/d/000000024/cern-openstack-overview?orgId=3&refresh=15m>)

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, 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).

STORAGE:


Websites:

www.cern.ch/eos

www.cern.ch/castor

<https://filer-carbon.cern.ch/grafana/dashboard/db/castor-dashboard?orgId=1>

Data recorded on tapes at CERN on a monthly basis in TB



This plot shows the amount of data recorded on tape generated by the LHC experiments, the other experiments, various back-ups and users. In 2018, over 115 petabytes of data in total (including about 88 petabytes of LHC data) were recorded on tape, with a record peak of 15.8 petabytes in November.

Magnetic tapes are used as the main long-term storage medium.

The tapes are stored in tape libraries and are retrieved by robotic arms.

- The amount of data recorded on tape at CERN is steadily increasing over time, with this trend accelerating over time (+40 % data stored in 2016 compared to what had been accumulated by the end of 2015 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 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 over 130 000 disks in the CERN Data Centre with a capacity of about 315 PB. 10-15% of the 130 000 of the disk drives are SSD ones (but it represents far less than 10% of the capacity itself). For the high capacity storage (~300 PB) used for the physics data in CASTOR 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.
- **As of the end of April 2019**, we had **339 PB on tapes** and the total capacity on tapes reached 800 PB (this can be extended as much as needed though).
- **As of the end of June 2019**, **173 PB were stored on disks** for a total capacity of 280 PB on disks. 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' 86 PB of 'real' data is on disks at the moment and the 'real' capacity is about 140 PB.

- The CERN storage system, EOS, was created for the extreme LHC computing requirements. In 2018, [EOS instances at CERN exceed three billions files](#), matching the exceptional performances of the LHC machine and experiments.

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 servers are water-cooled using active or passive heat exchangers in their rear doors, providing a higher cooling capacity per rack.

WLCG – Worldwide LHC Computing Grid

Websites:

www.cern.ch/wlcg
www.cern.ch/wlcg-public


Key facts and numbers:

- More than 170 data centres in more than 40 countries
- Close to 1 million CPU cores and 1 Exabyte of storage available to the LHC experiments
- Allows more than 12 000 physicists to access LHC data

- >300 000 jobs run concurrently on the Grid
- In 2018, global transfer rates regularly exceeded 60 GB/s

Evolution of the global processor time delivered by the Worldwide LHC Computing Grid (WLCG)

As seen on the graph, the global central processing unit (CPU) time delivered by WLCG (expressed in billions of HS06 hours per month, HS06 being the HEP-wide benchmark for measuring CPU performance) shows a continual increase. In 2018, WLCG combined the computing resources of about 900 000 computer cores.



Short overview:

- Today, the Worldwide LHC Computing Grid (WLCG) combines the computing resources from more than 170 data centres in more than 40 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. Close to 1 million CPU cores and 1 Exabyte of storage available to the LHC experiments.
The Worldwide LHC Computing Grid runs over 2 million tasks per day and in 2018, global transfer rates regularly exceeded 60 GB/s.
These numbers will increase as time goes on and as computing resources and new technologies become ever more available across the world.

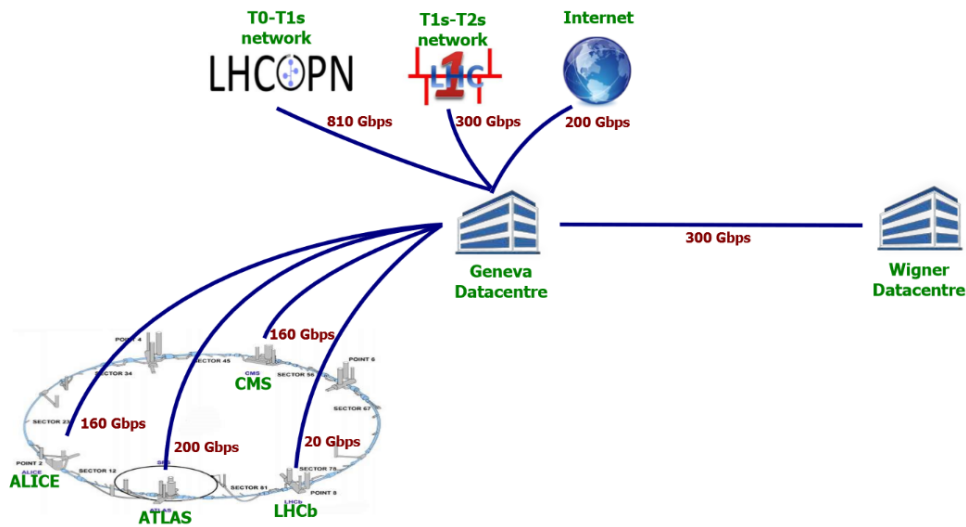
NETWORK:

Webpage:

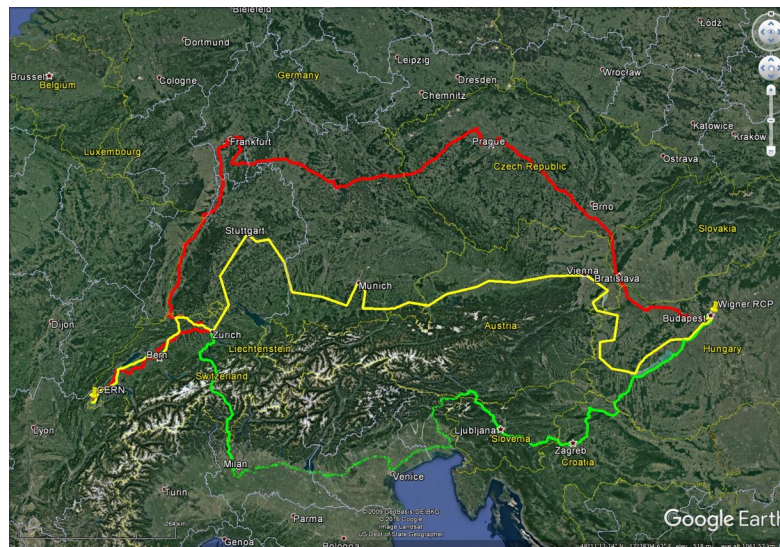
<https://cixp.net>

Key facts and numbers:

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



- Three 100 Gigabit per second circuits used to connect the Meyrin Data Centre to its extension in Hungary (cf. below).



- Since the beginning of the second four-year-long LHC running period, the data transfer rates around the globe reached new peak rates – about 60 gigabytes per second rates, around a factor of three 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.
- 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

Open Source:

Cf. computing chapter of the CERN annual report 2016: https://e-publishing.cern.ch/index.php/Annual_Report/article/view/496/354

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 companies and research institutes.

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>

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>

THE PAST:

- *Work in progress, a guide for guides about the history of the Data Centre is being prepared.*
- *The first section of this video gives a brief overview of the history of the Data Centre (but the rest of the video is not anymore up to date in terms of numbers, etc.): CERN IT in 8 minutes (2013) <https://cds.cern.ch/record/1604210>*

THE FUTURE:

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 2026. By this time, the total computing capacity required by the experiments is expected to be 50-100 times greater than today, with data storage needs expected to be in the order of exabytes.

In order to address this, 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.

The large-scale data challenges of HL-LHC are also shared with other science communities. WLCG and CERN have been working with the Square Kilometre Array (SKA) radio astronomy community to explore some of these common challenges. In June 2017, CERN and the SKA Organisation signed a Collaboration Agreement for joint work in computing and data management.

CERN openlab has also 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.

To know more:

<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>

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).