WLCG, status and plans

Julia Andreeva (CERN)

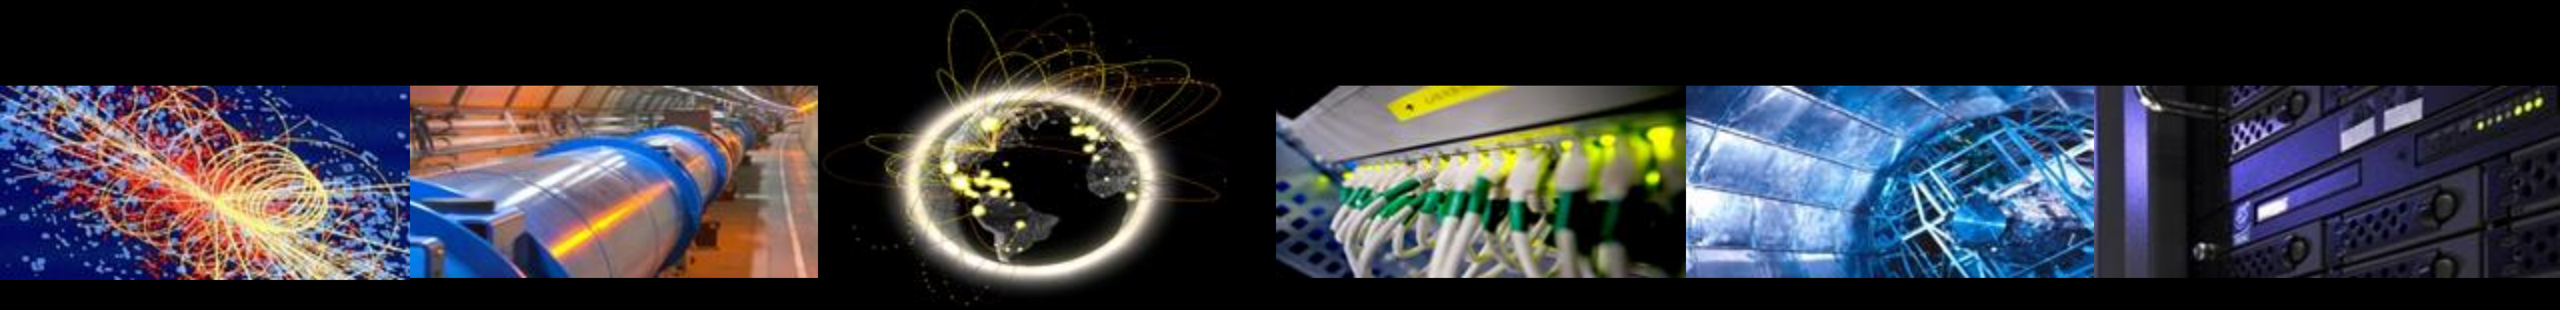

Why a distributed GRID model for LHC computing?

The distributed grid model is a natural fit for:

🎯 **National Strategies (Funding Agencies)**

- Supports the development of national **cyberinfrastructure**
- Encourages investment in **data centers, networks, and expertise**
- Aligns with long-term goals for scientific and technological advancement

🌍 **The Structure of the HEP Community**

- Mirrors the **globally distributed** nature of High-Energy Physics (HEP)
- Engages **hundreds of institutes** in international collaborations
- Facilitates participation in large-scale experiments like the LHC

Certainly, there were also technical reasons:

- Unprecedented Data Volume of LHC
- Resilience and Redundancy

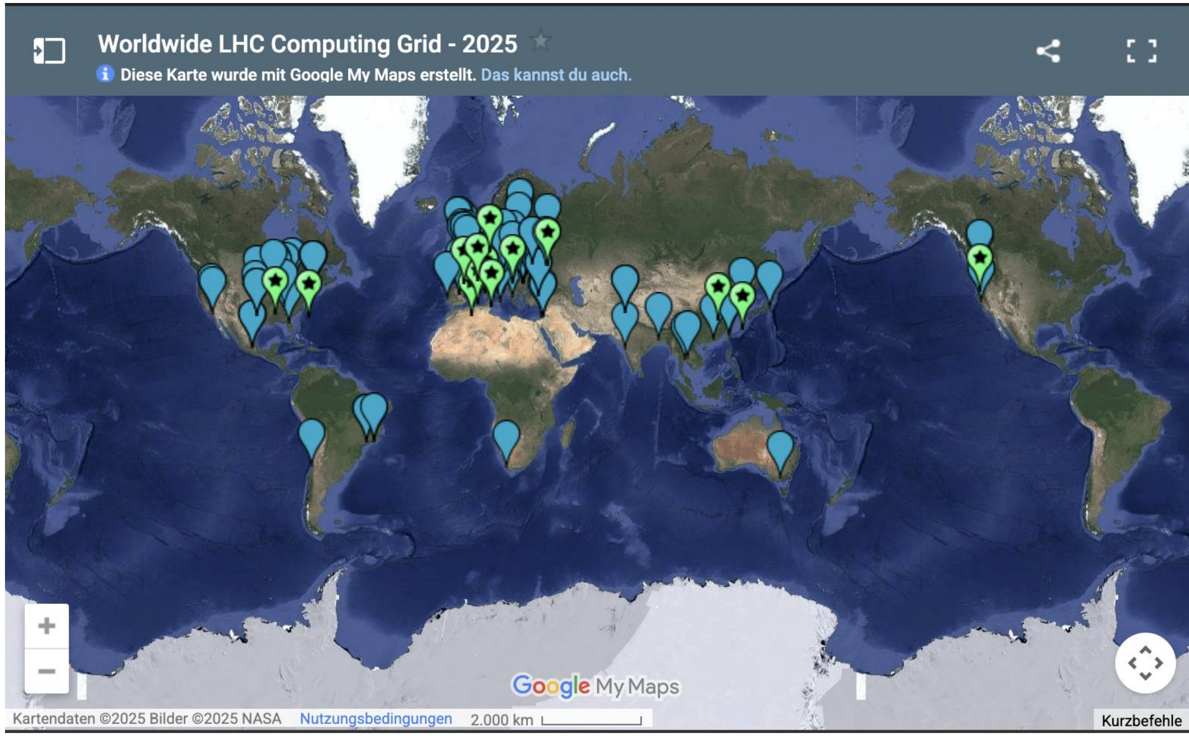

Worldwide LHC Computing Grid (WLCG)

A global collaboration that enables the **storage, distribution, processing, and analysis** of data from the Large Hadron Collider (LHC) experiments.

Brings together computing centres, LHC experiments, GRID infrastructure projects (EGI, OSG, NorduGRID, GridPP), funding agencies, software developers and technology providers

Since 2006, we have ensured continuous global data processing and transfers—24/7/365—with increasing scale over time. Our systems offer exceptional availability and reliability, supported by a robust and sustainable funding model.

In numbers

65 MoUs, 40+ countries, 164 sites

The success: meets the needs of the LHC experiments

- by the integration of globally distributed Exascale resources with services and software
- within a trust framework that transcends site and national boundaries.

WLCG Resources 2025	
CPU (cores)	~1.3M
Disk (PB)	~1EB
Tape (PB)	~2EB

WLCG
Worldwide LHC Computing Grid

A bit of history

The initial Computing Model

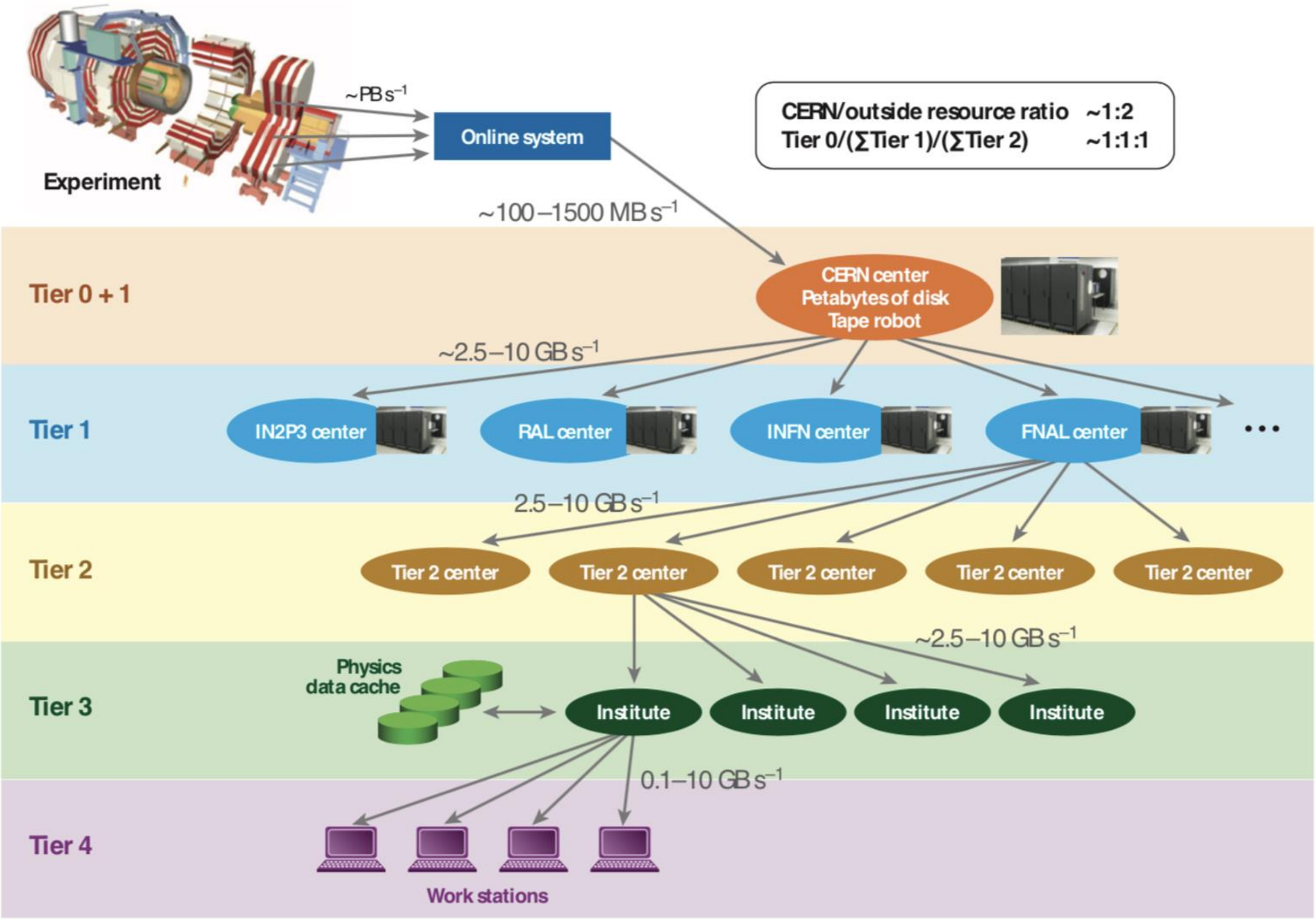

This was the initial computing model (1999)

Uncertainty over network performance, reliability

Focus on distributing data *globally* to compute resources in a hierarchical structure

Every center runs a compute and a storage service

No concept of data remote from compute

Predefined roles of Tiers: T2s for simulation and analysis, T1s for reconstruction and archive

Quickly evolved …

Adapted from a slide by Simone Campana

WLCG
Worldwide LHC Computing Grid

And evolution (1)

Heterogeneous Grid Infrastructure

Organisation of WLCG facilities in Tiers: simple but effective way to define the expected quality-of-service and target metrics specified in the WLCG MoU. But ..

- Network connectivity not a limiting factor (WAN connectivity increased X10 times in last 10 years)
- Relaxing the Tier model increases flexibility

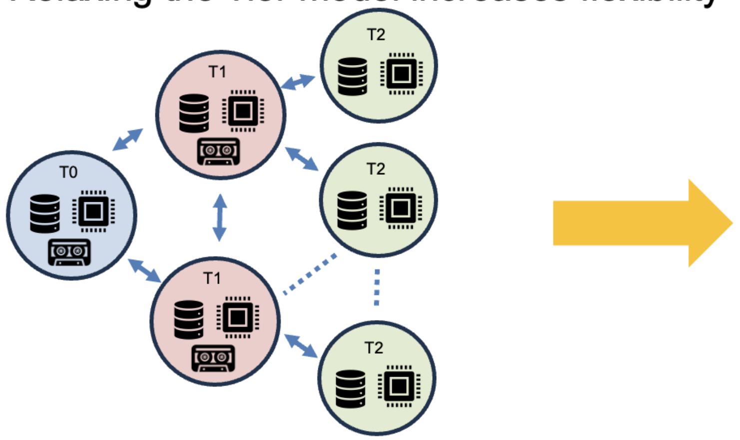

- Heterogeneous facilities (HPCs and Clouds) offer additional opportunities

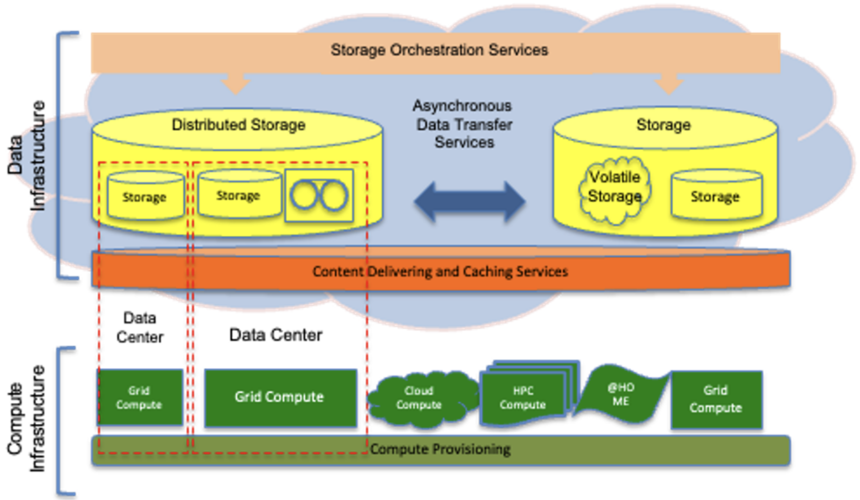

Adapted from a slide by Simone Campana

And evolution (2)

Evolution to a more flexible Computing Model

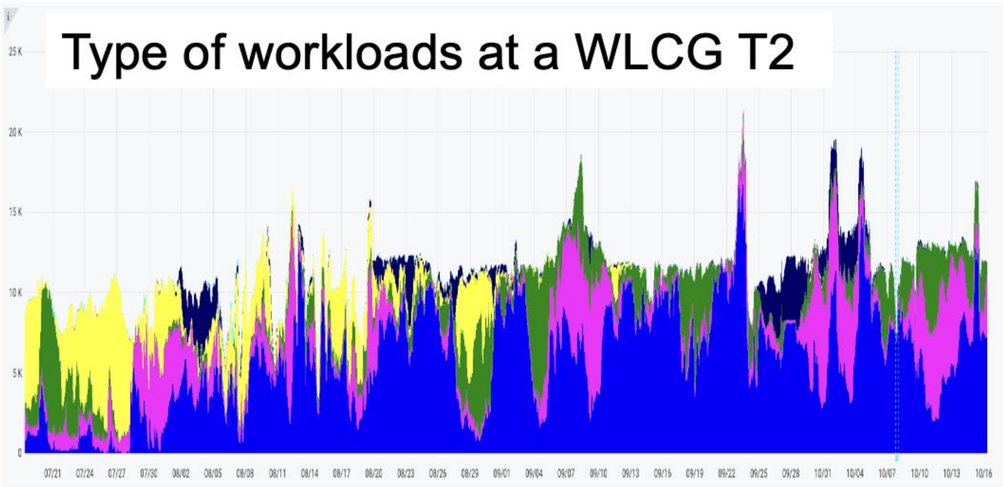

T1s and T2s run flexibly a mixture of workloads depending on their capabilities and the needs of the experiments

Data is transferred in a full mesh rather than hierarchically

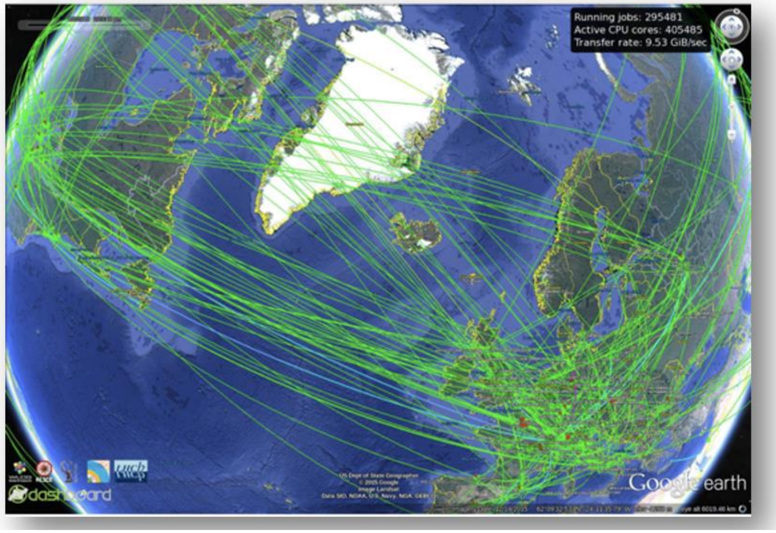

Adapted from a slide by Simone Campana

WLCG
Worldwide LHC Computing Grid

Data processing, simulation and analysis

CPU usage is constantly growing, even during shutdown periods

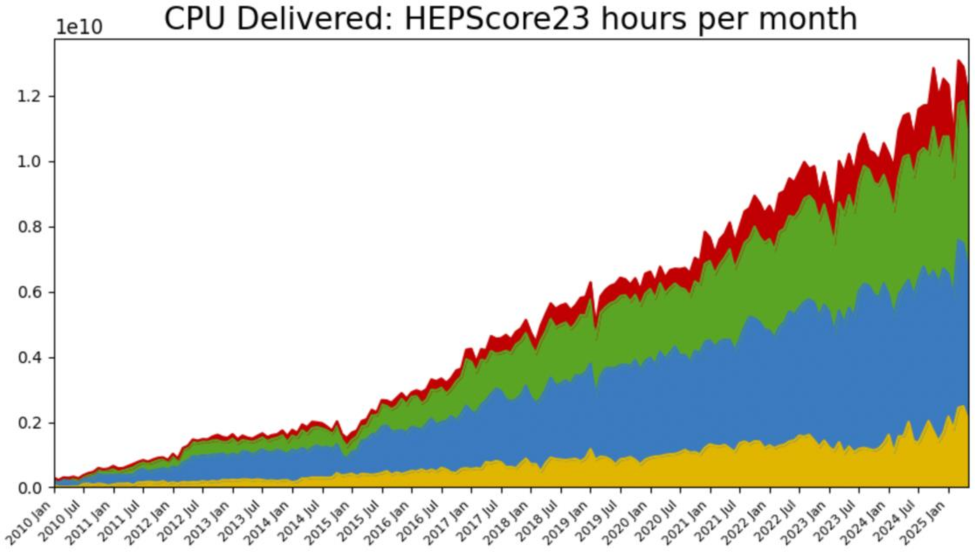

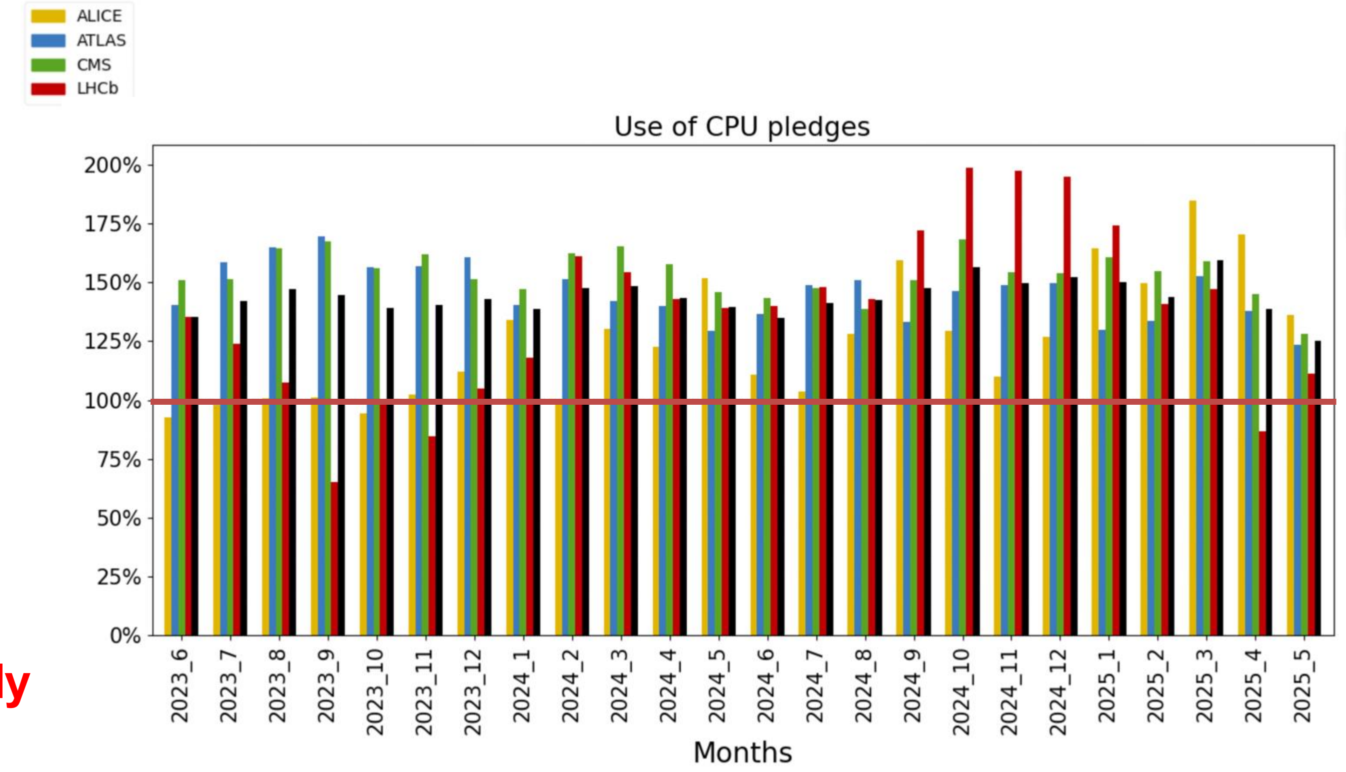

The delivered CPU usage consistently exceeds the pledged amount.

Data transfer

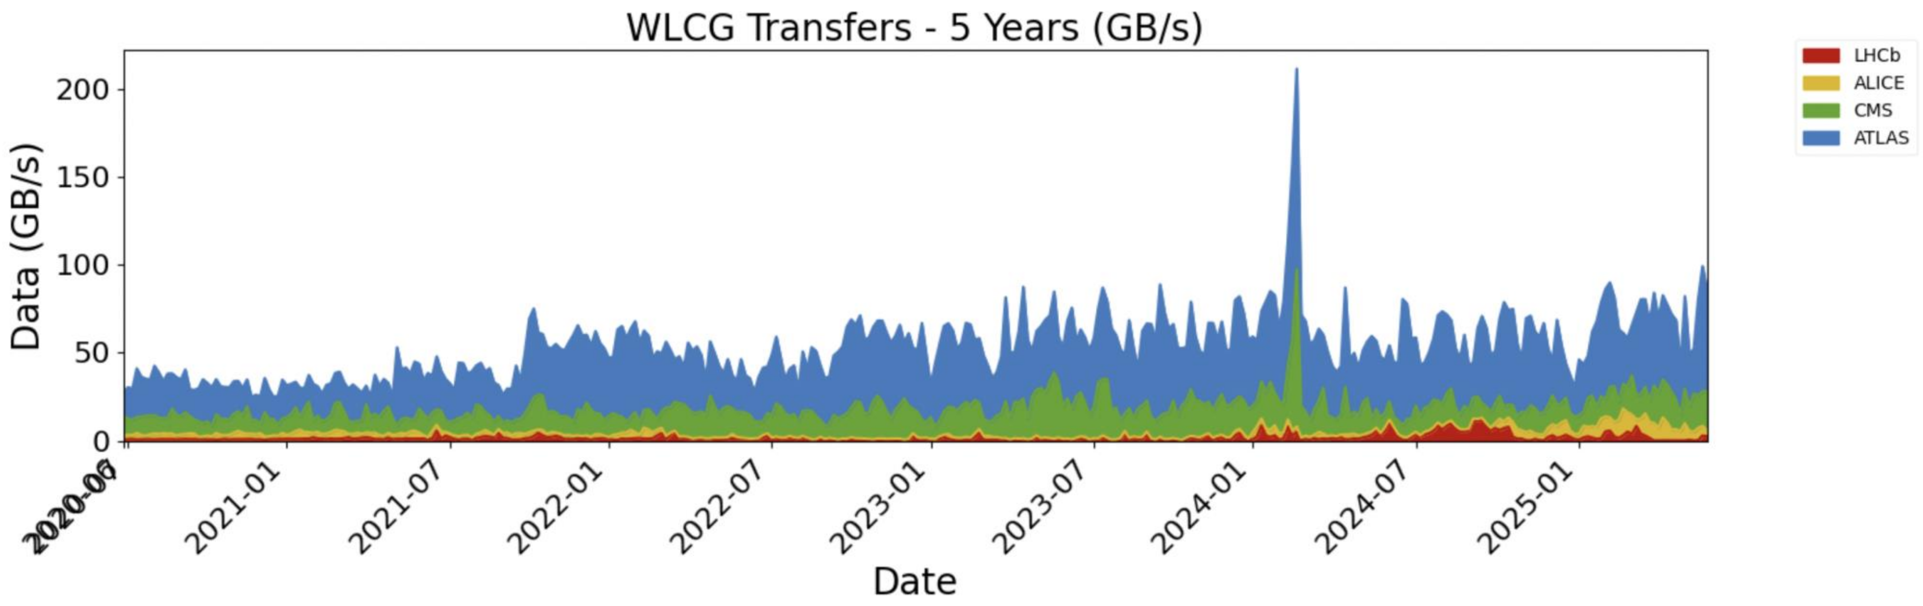

Transfer rate is also steady growing

WLCG
Worldwide LHC Computing Grid

WLCG networks

WAN connectivity increased x10 in the last 10 years

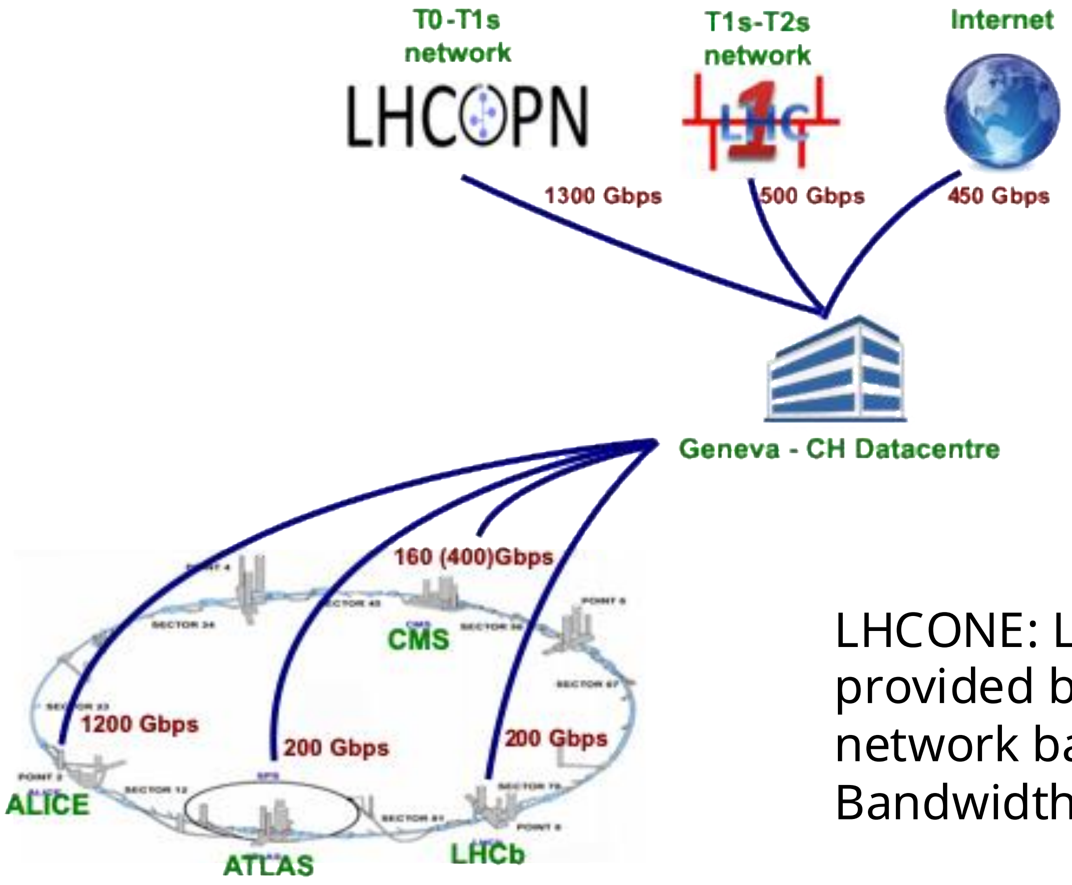

WLCG networks are an opportunity, not a limitation as initially expected

LHCOPN: Private network connecting Tier0 and Tier1s. Dedicated to LHC data transfers and analysis

LHCONE: Layer3 (routed) Virtual Private Network provided by the R&E network providers Worldwide network backbone connecting Tier1s and Tier2s. Bandwidth dedicated to High Energy Physics

Slide by Simone Campana

WLCG Worldwide LHC Computing Grid

WLCG operations

WLCG is a highly complex, distributed infrastructure:
•~170 sites worldwide
•Over 1.2 million cores running on average
•5 to 10 million data transfers per day
•Sustained throughput averaging 60 GB/s

The volume of resources in WLCG (CPU, storage, network) has increased by more than one order of magnitude since the start of the LHC, with no degradation of the service efficiency and performance.

Operating at this scale is a significant challenge.

Operations are shared between sites, experiments and WLCG OPS coordination team

Site availability monitor

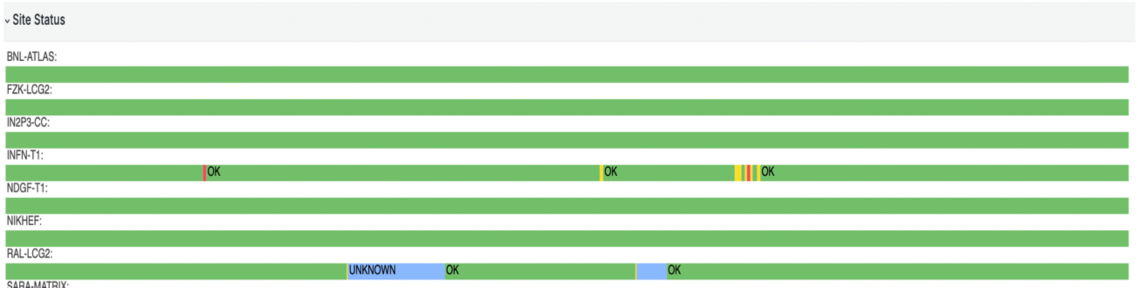

Accounting reports

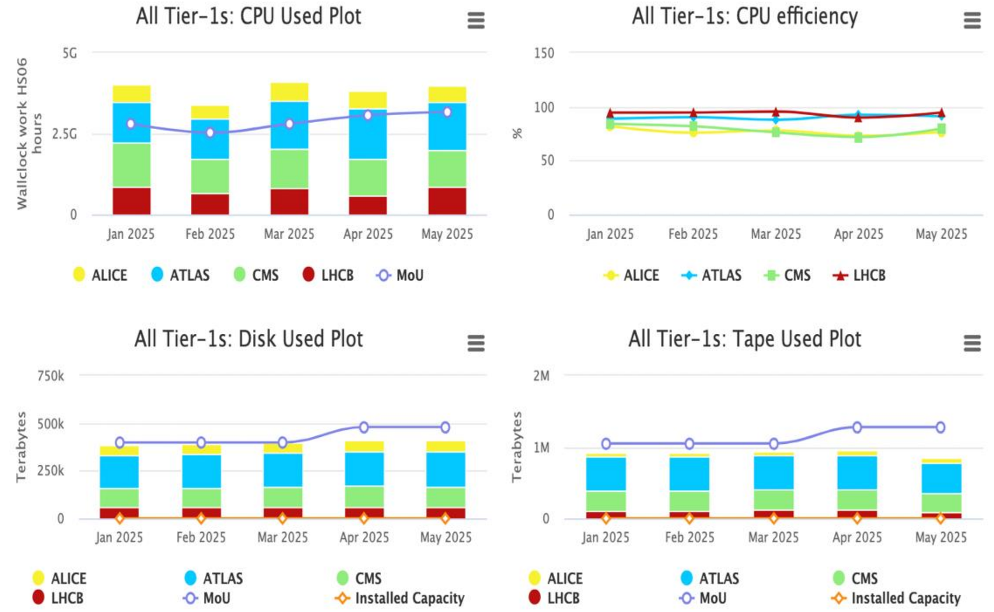

WLCG
Worldwide LHC Computing Grid

WLCG Technical Coordination Board

The TCB is responsible for guiding the **technical evolution** of WLCG service, aligning with:

- The needs of the LHC experiments
- The capabilities of infrastructure providers

Its core mandate is to develop and maintain the **WLCG Technical Roadmap**, and to:

- Facilitate, coordinate, and monitor its implementation
- Ensure coherence and long-term sustainability of WLCG services

- The process is consensus-driven, engaging all major stakeholders
- Collaboration is fostered through the **Open Technical Forum** — a series of inclusive, topic-focused meetings open to the community

WLCG Worldwide LHC Computing Grid

JINR is a part of WLCG infrastructure

- **JINR is a CMS T1 site and T2 site for all 4 LHC experiements**

- CMS perspective:
  - JINR is one of seven Tier-1 sites of CMS
  - **The Tier-1 contribution of JINR is highly valued in CMS !**
  - The site provides around 19,000 cores, 14 PBytes of disk, and 25 PBytes of tape storage to the CMS computing grid. This is about 10% of the CMS Tier-1 needs
  - JINR also provides valuable Tier-2 services to CMS
  - Services are robust and both Tier-1 and 2 are reliable sites
  - Over the past four years, JINR contributions have remained stable, while the computing requirements for CMS have increased with the progression of Run 3.

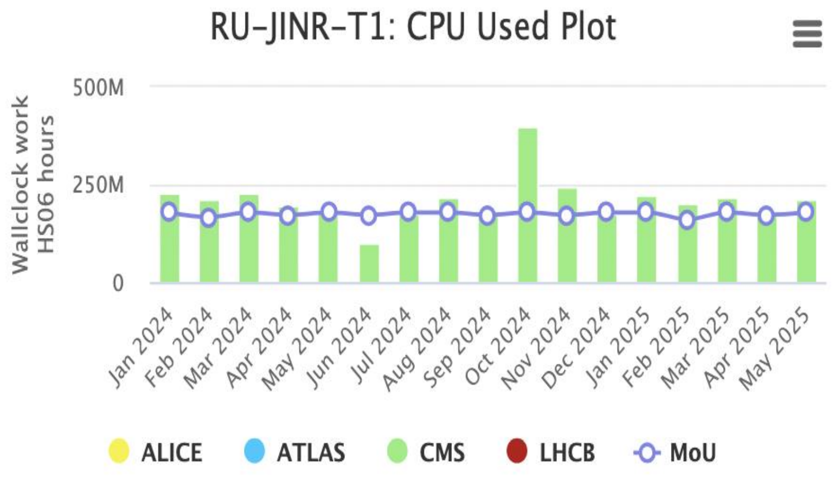

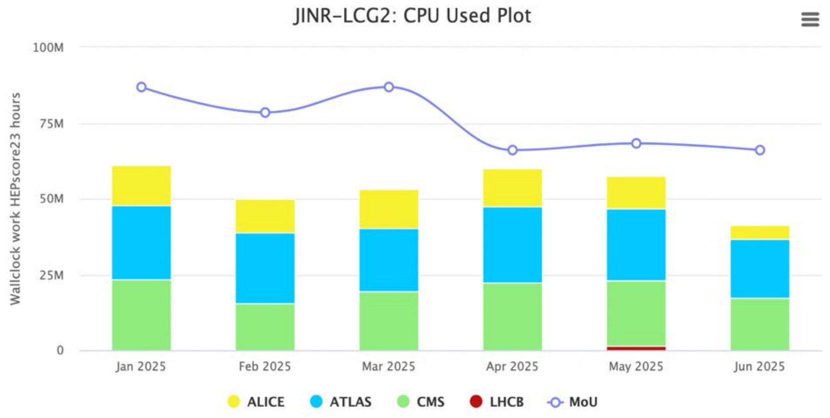

WLCG
Worldwide LHC Computing Grid

Looking ahead – data volumes

Data volume and complexity of HL-LHC data

Several other significant HEP experiments are also underway or coming online in 2020s : Belle II, JUNO, DUNE

And other sciences with similar computing requirements.
SKA is the main example: expect 700 PB/year of data products

Looking ahead –hardware trends

- Cost of hardware decreasing vs time 😄 … but not as steeply as before 😳
- In general, trends driven by market (revenues) rather than technology

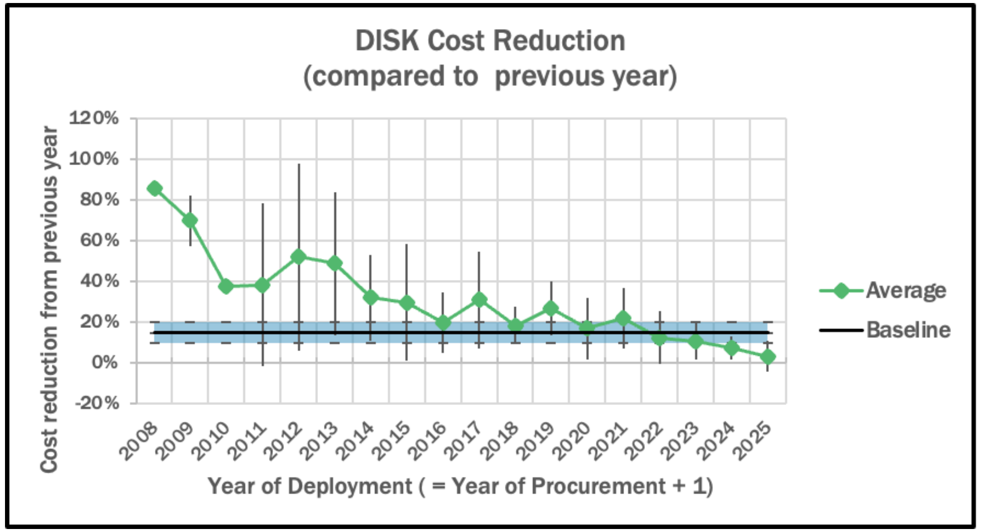

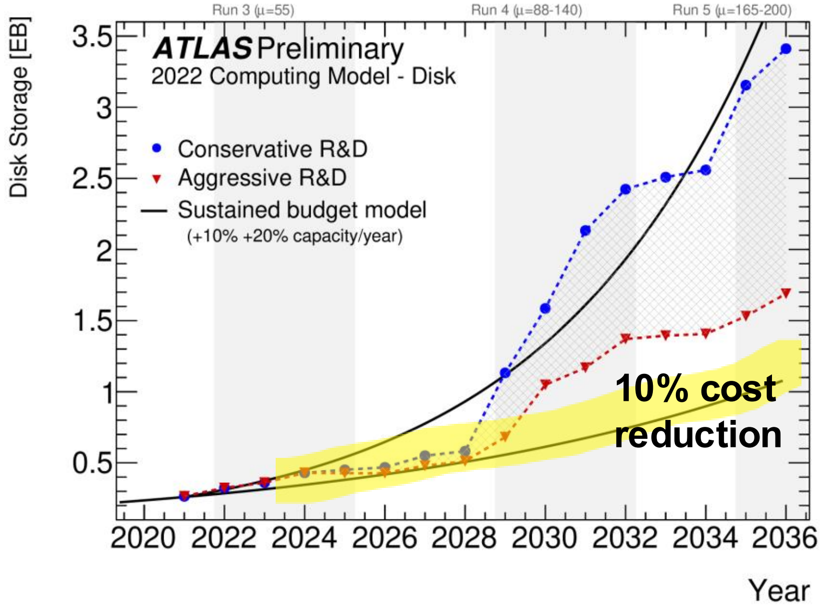

ATLAS Phase-2 computing model (2022)

Adapted from a slide by Simone Campana

Main HEP computing challenges

Exponential Growth in Demand: The HL-LHC will generate 10 times more data per year and 5 times more complex—making it the primary driver of increased resource needs.
Flat Funding Outlook: No significant increase in funding for HEP computing is anticipated in the foreseeable future.
Challenging Tech Trends: Developments in hardware technology and cost trends are not aligned with the needs of HEP, making it harder to scale affordably.
Global Instability: A volatile geopolitical environment poses risks to globally distributed infrastructure and international collaboration.
Outdated Systems: Many existing HEP computing services were designed in a different era, and the ICT landscape has evolved dramatically since then.
Pressure for Broader Impact: Funding agencies increasingly expect their investments to support multiple scientific disciplines, not just HEP.

WLCG Worldwide LHC Computing Grid

How do we progress?

Reliable Operations

- First and foremost, WLCG **must continue to operate reliably** without interruption.
  - Disruptive changes should be implemented in the production environment **without causing downtime or disruption**.
  - Significant progress is required in **automation** to improve efficiency and stability.

Increasingly commission the WLCG infrastructure to the HL-LHC scale

- Example: Data challenges programme

Middleware Modernization

- Original middleware designed before mature distributed computing. Modernize software stack to exploit current technologies
- Example: WLCG DOMA (Data Organization, Management, Access) initiative

Integration & Performance

- Integrate heterogeneous resources (HPS, clouds)
- Enchance software performance & portability to meet future HEP computing challenges

Innovation and Collaboration: Cornerstones of the WLCG Strategy

- **Innovation:** Modernize software & services with cutting-edge tech
- **Collaboration:** Leverage synergies across HEP & other sciences

WLCG
Worldwide LHC Computing Grid

WLCG Data Challenge 2024 (DC24)

Advancing Readiness for HL-LHC

Challenge Overview

- Second in a series of scale-up challenges
- Targeted 25% of HL-LHC throughput
- Involved LHC experiments, Belle II, and DUNE
- Tested: data transfers, tools (FTS, Rucio), token-based auth, and network capabilities

Key Results

- **Performance targets met**: 1.2 Tbps (Minimal), 2.4 Tbps (Flexible)
- ~50% of transfers used tokens; first large-scale operational test
- Most sites stable; a few identified local bottlenecks for future upgrades
- Monitoring improved
- Successful evaluation of advanced network technologies

Looking Ahead

- Insights critical for future scalability and resilience

•**Next Challenge: Autumn 2026**, targeting 50% of HL-LHC traffic

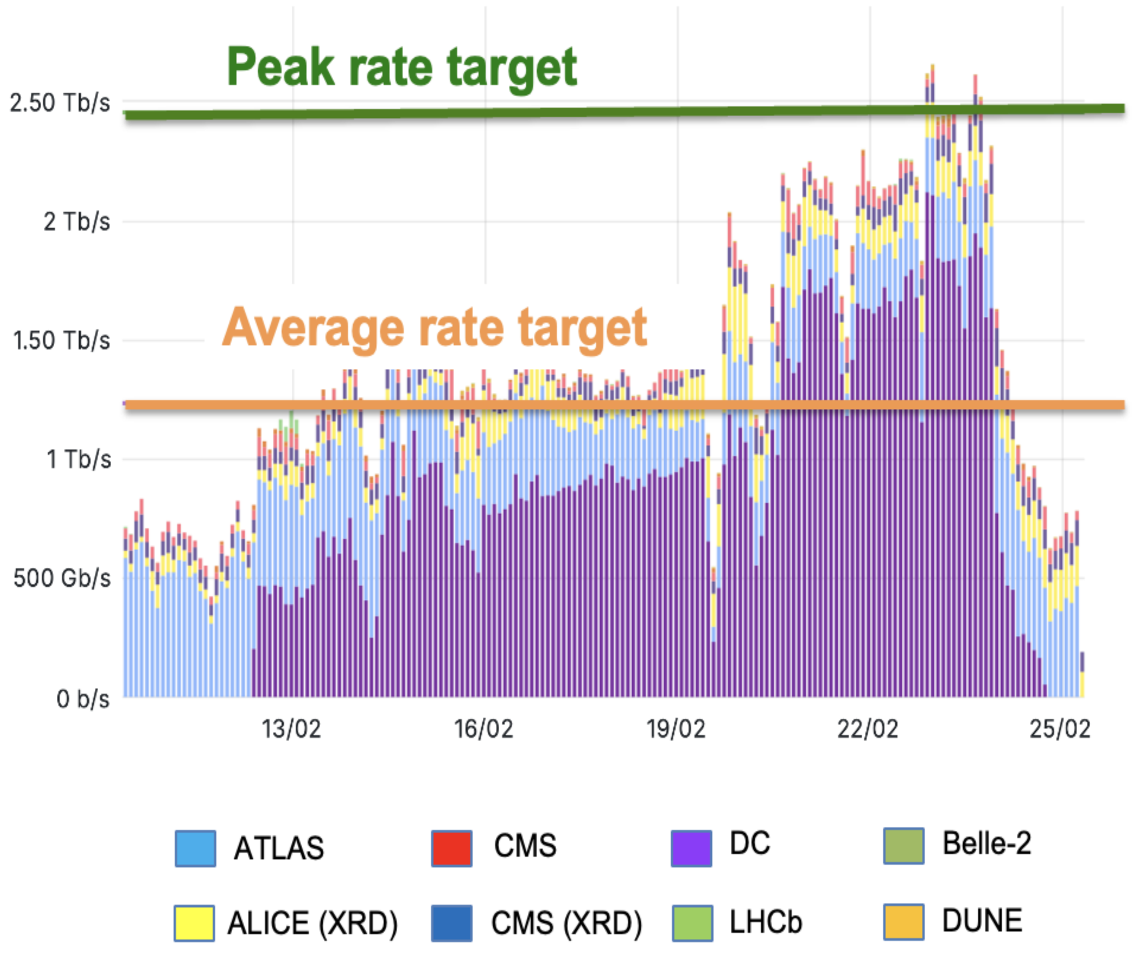

WLCG
Worldwide LHC Computing Grid

Transition to tokens

- Since 2017, WLCG has been working towards enabling token-based authentication and authorization throughout its entire middleware stack
- A lot of progress has been done:
    - VOMS-Admin legacy services were decommissioned and IAM services deployed in production
    - Majority of WLCG services dealing with job submission, transfers and data access were modified
- This is ongoing process and technical roadmap with tentative milestones has been presented and discussed at the WLCG WS in spring 2025

Technical roadmap – tentative milestones

- 2025-Q2 Specification of token usage for tape operations
- 2025-Q3 Release of WLCG Token Profile v2.0
- 2025-Q4 First production usage of Rucio tokens in ATLAS
- 2025-Q4 First usage of DIRAC tokens in LHCb
- 2025-Q4 First usage of tokens in ATLAS jobs
- 2025-Q4 Operational Risks: Service risk document detailing the effects of outages – possibly with actual downtime tests
- 2026-Q3 CMS grid jobs with only tokens
- 2026-Q4 Token Grand Challenge, with use of tokens by jobs
- 2027-Q1 Data Challenge 2027
- 2028-Q1 Completion of the X509 / VOMS phaseout

WLCG
Worldwide LHC Computing Grid

Leveraging Heterogeneous Facilities - HPCs

HPCs have been integrated as WLCG resources since ~ 20 years

Some HPC are very challenging to use: limited connectivity, special harware (GPUs, non-X86 CPUs) and configs (RAM, local disk), different AAI.

The benefits are considerable in many cases. Attempts to move the collaboration with HPC centers into consolidated partnerships

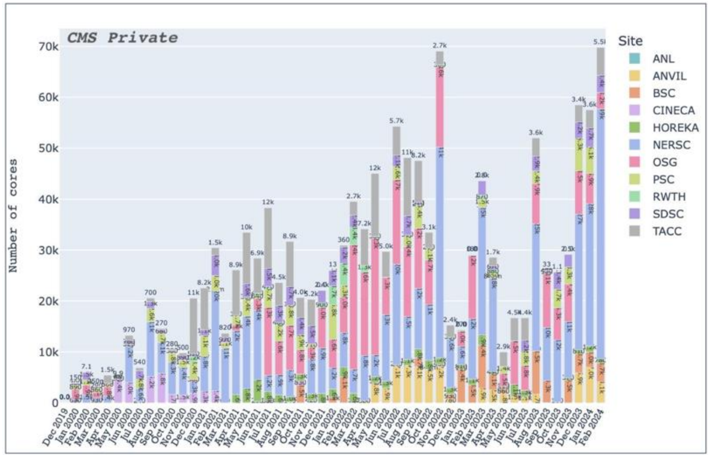

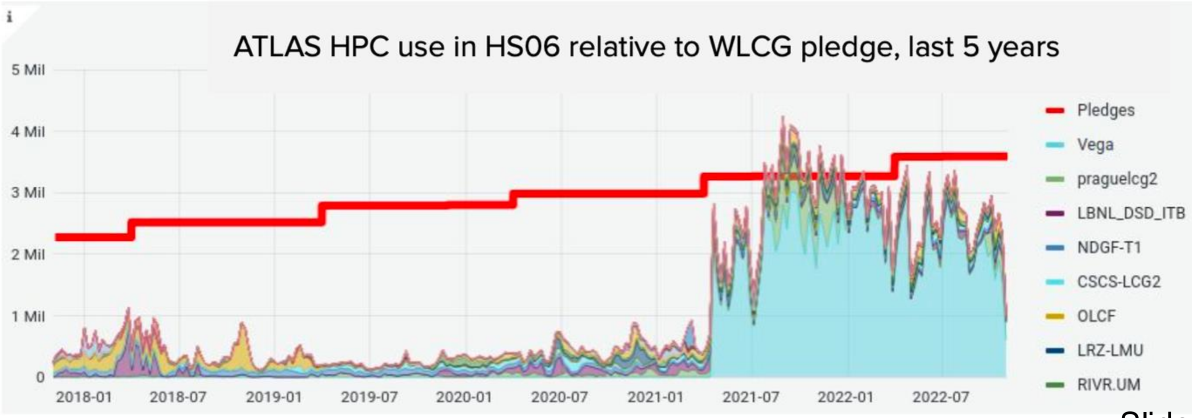

Slide by Simone Campana

WLCG Worldwide LHC Computing Grid

Software Performance and Portability

Software performance, modernization, innovation and portability is one of the key areas to address the future HEP computing challenge. It allows to:

- access available resources otherwise not useful (e.g GPUs, non-X86 CPUs)
- optimize procurement and deployment to guarantee the best value for money
- support different experiments and sciences on the same resources
- leverage the most modern hardware features (e.g vectorization)
- reduce the resource needs/event (compute and storage)

It needs upfront large investments and forward-looking solutions

- E.g. LHC experiments have a 20+ years codebase and rewriting from scratch is not an option

Table 6. SLOCCount measured lines of source code for ATLAS and CMS.

Experiment Type	Source Lines of code (SLOC)	Development effort (person-years)	Total estimated cost to develop
ATLAS	5.5M	1630	220 M$
CMS	4.8M	1490	200 M$

Linux Kernel is: 15M sloc, 4800 FTEy, 650M$;

Geant4 is: 1.2M sloc, 330 FTEy, 45 M$

Slide by Simone Campana

WLCG Worldwide LHC Computing Grid

Event Generators on GPUs

A very good candidate for GPU acceleration with benefits for many experiments

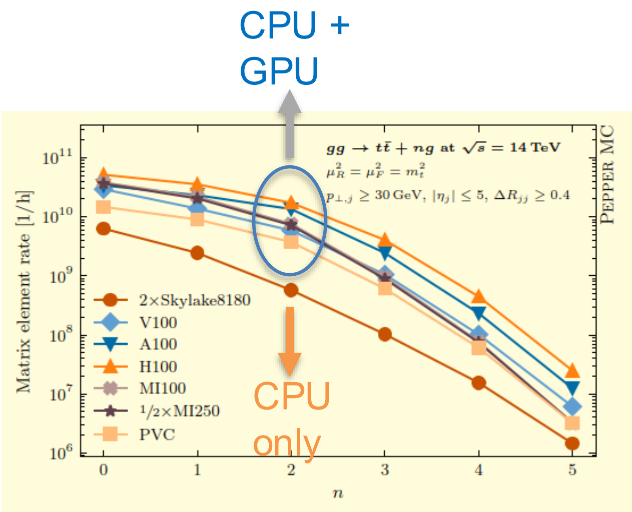

Sherpa gg->tt+ng

Matrix Element event throughput: up to x10 gain when using GPUs

Available for production

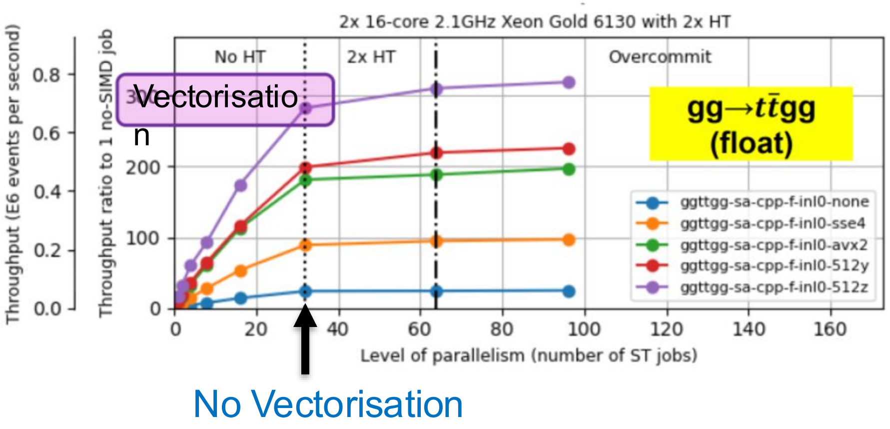

Madgraph gg->tt+ng (n=2)

GPU-enabled Leading Order: being released to production. By-product: enabling of CPU vectorisation: up to x8 gain in ME event throughput (x6 global). Note: all CPUs in WLCG provide vectorisation

GPU-related work brings immediate benefits also on CPUs

Slide by Simone Campana

WLCG Worldwide LHC Computing Grid

AI in HEP: Unlocking New Capabilities at Scale

- The use of AI in High Energy Physics is not new
- **What's Game-Changing Now**
  - Breakthroughs in AI models:
    - Transformers & LLMs: Capable of learning complex patterns from high-dimensional data
    - Normalizing Flows: Powerful tools for generative modeling and fast simulation
    - Graph Neural Networks (GNNs): Naturally suited to model particle interactions and detector geometry
  - Availability of large-scale AI infrastructure:
    - Access to GPU-accelerated clusters, AI-specific hardware (e.g. TPUs), and HPC centers
    - Emerging AI-as-a-service platforms and optimized software stacks

To fully leverage this potential, **WLCG must advance its support for AI-driven applications**, enabling scalable, efficient, and integrated workflows across the HEP community.

WLCG
Worldwide LHC Computing Grid

User analysis evolution

Adapting to Evolving User Expectations

- **Past:** Users primarily programmed via ROOT and C++ macros.
- **Present:** The landscape has shifted toward Python, Jupyter notebooks, browser-based frontends, and a variety of modern data science tools.

Modern Analysis Facilities

- Need to provide tailored analysis ecosystems, offering integrated tools and services that support interactive and scalable workflows.
- Looking ahead, these facilities may also evolve to become centers for large-scale AI training, equipped with specialized hardware and infrastructure

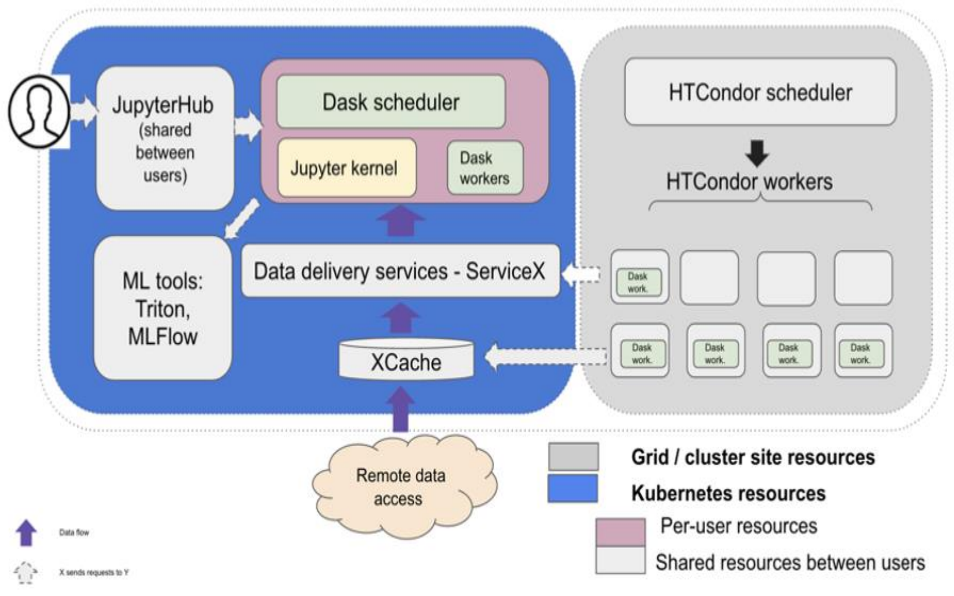

WLCG
Worldwide LHC Computing Grid

Conclusions

- For over **20 years**, WLCG has provided a shared, distributed computing infrastructure supporting the LHC experiments.
- Over time, both the computing models and the supporting infrastructure and services have evolved—adapting to new technologies, operational experience, and funding realities.
- The upcoming High-Luminosity LHC (HL-LHC) represents an unprecedented challenge in terms of scale, complexity, and long-term sustainability.
- At the same time, other data-intensive scientific projects will operate alongside LHC:
  - With similar computational and data access patterns
  - On comparable timelines
  - Often relying on the same physical infrastructure
- This presents a unique opportunity to share tools, services, and infrastructure, align policies and best practices and collaborate across scientific domains for mutual benefit

WLCG
Worldwide LHC Computing Grid

WLCG 20th anniversary

On the 8th of December 2025 WLCG will celebrate its 20th anniversary !

WLCG
Worldwide LHC Computing Grid