WLCG evolution towards the High-Luminosity LHC

from an Ops Coordination perspective

M. Litmaath, CERN

June 23, 2026

What is the Worldwide LHC Computing Grid?

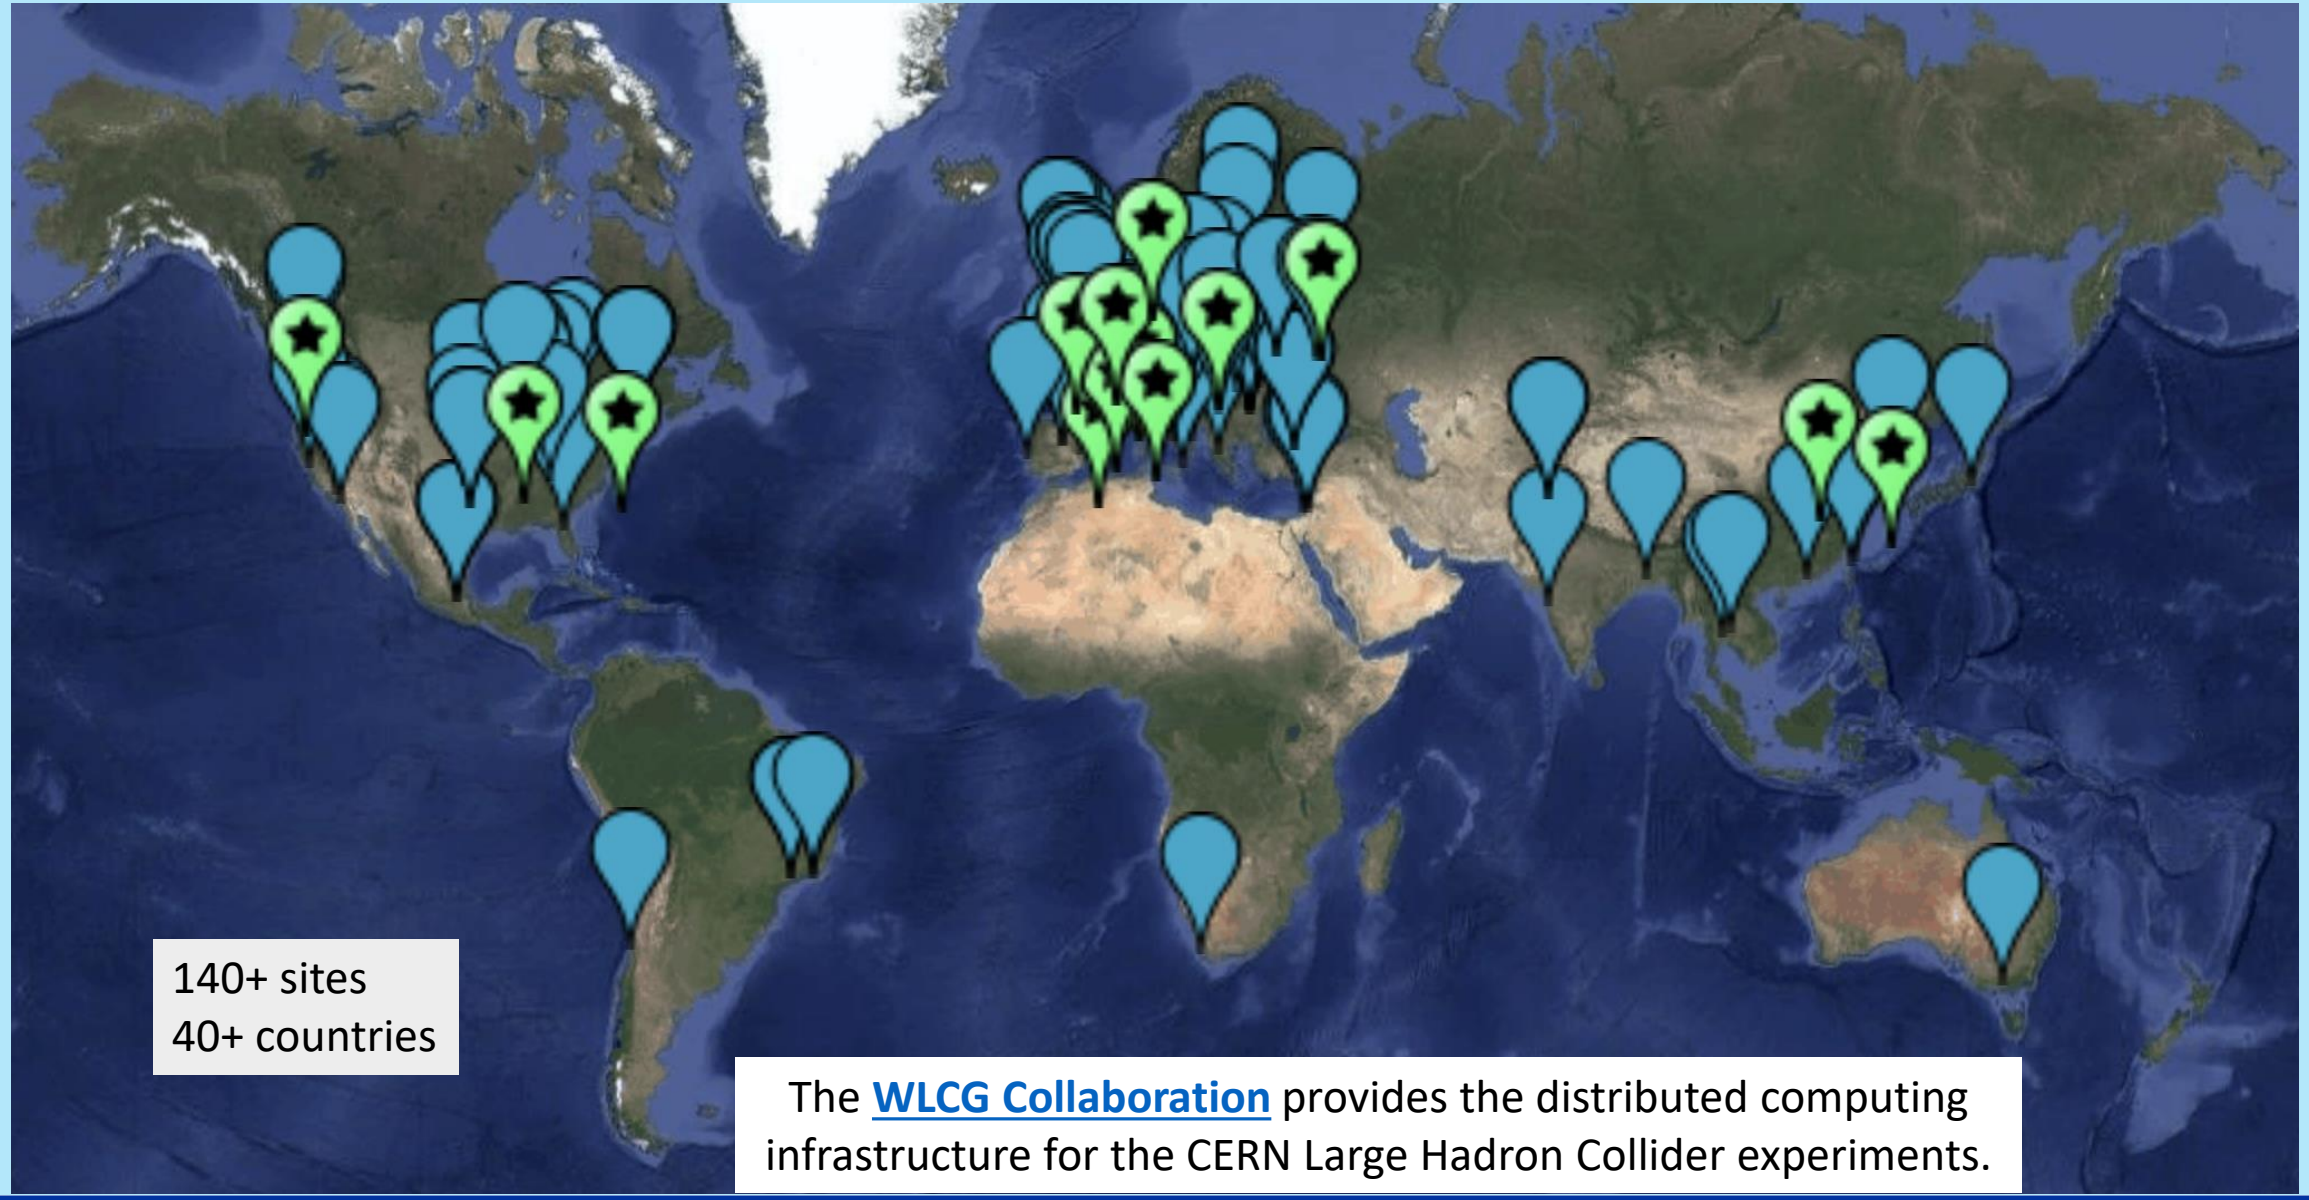

The **WLCG Collaboration** provides the distributed computing infrastructure for the CERN Large Hadron Collider experiments.

Worldwide data storage and processing 24/7

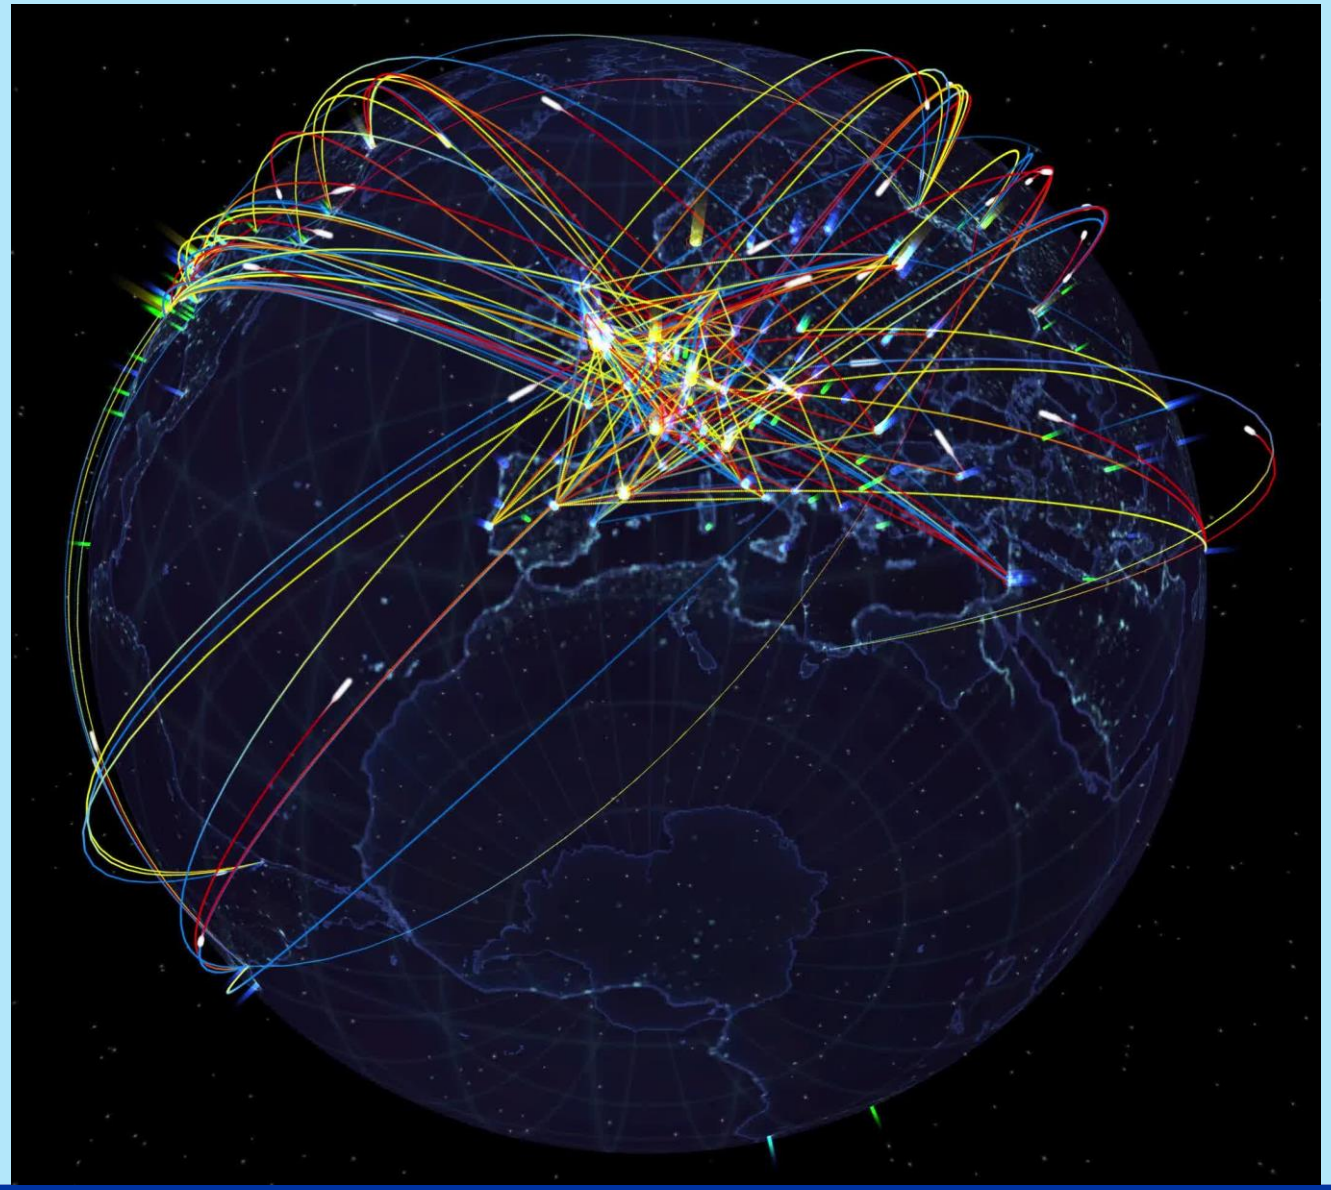

LHC experiment data on tape at CERN

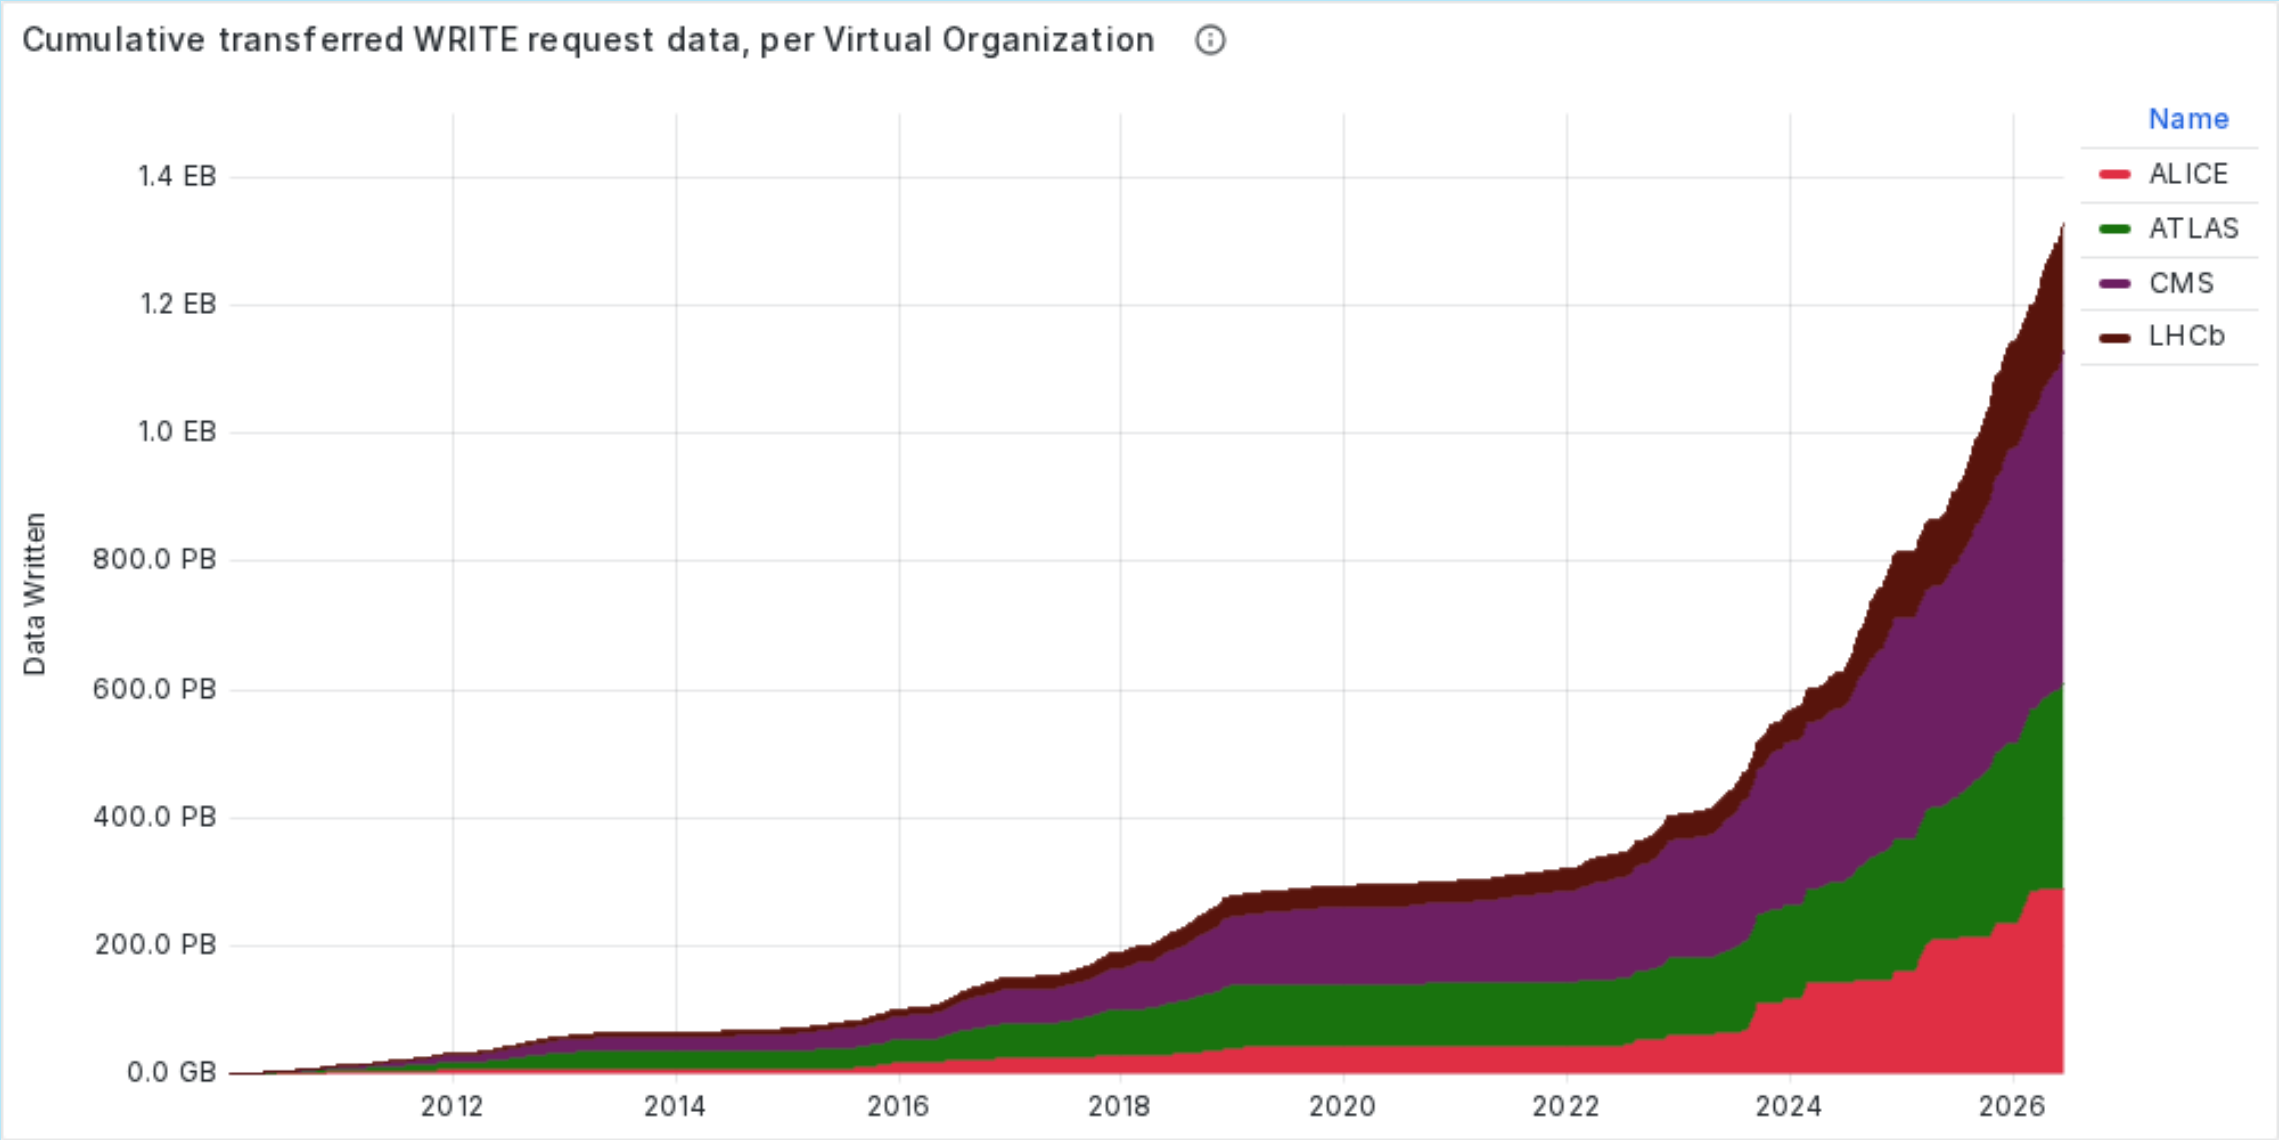

Disk pledge history

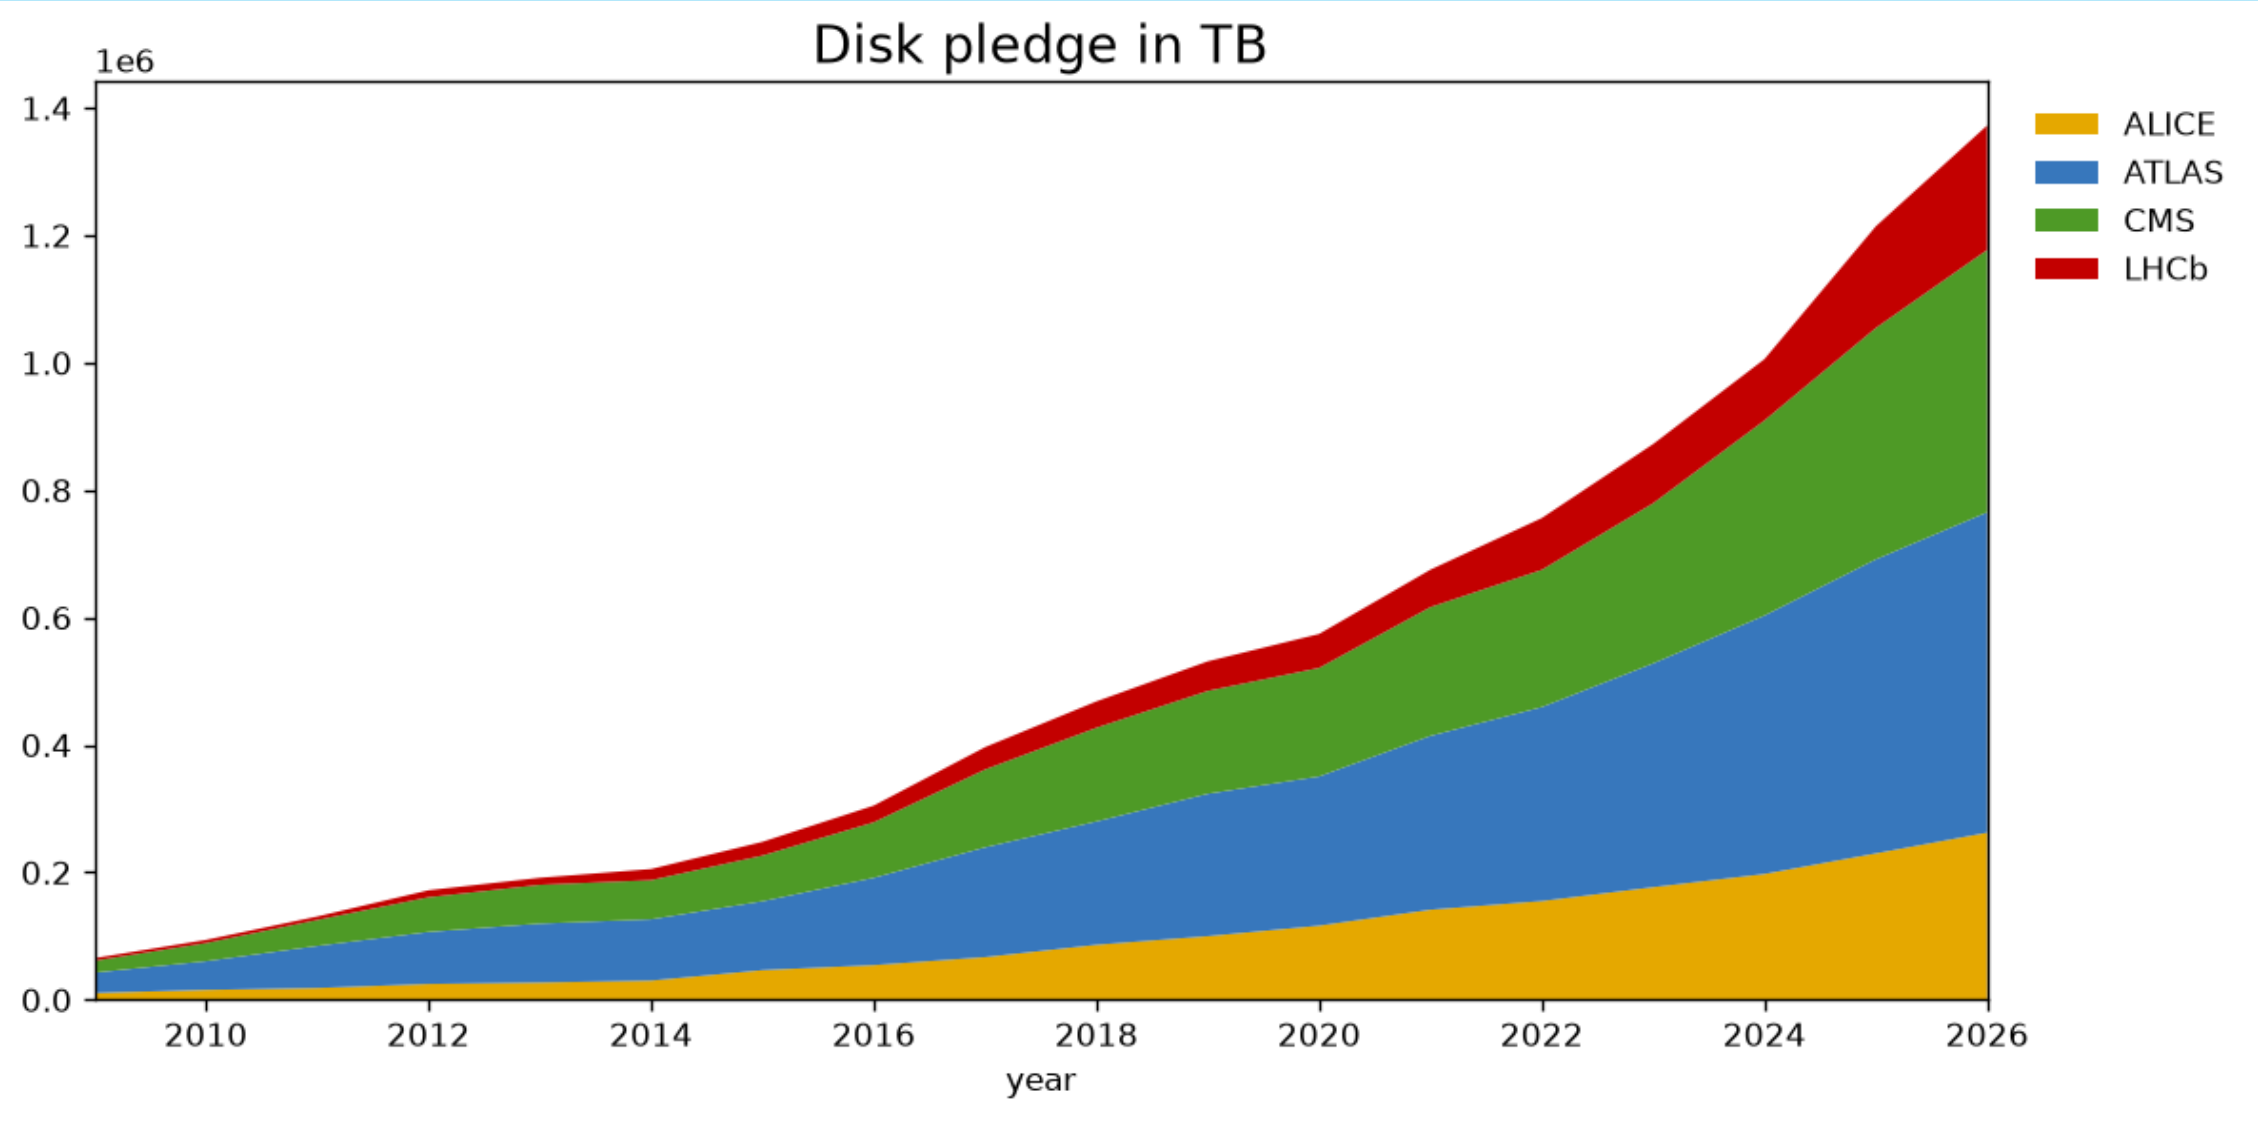

Corresponding increases in computing

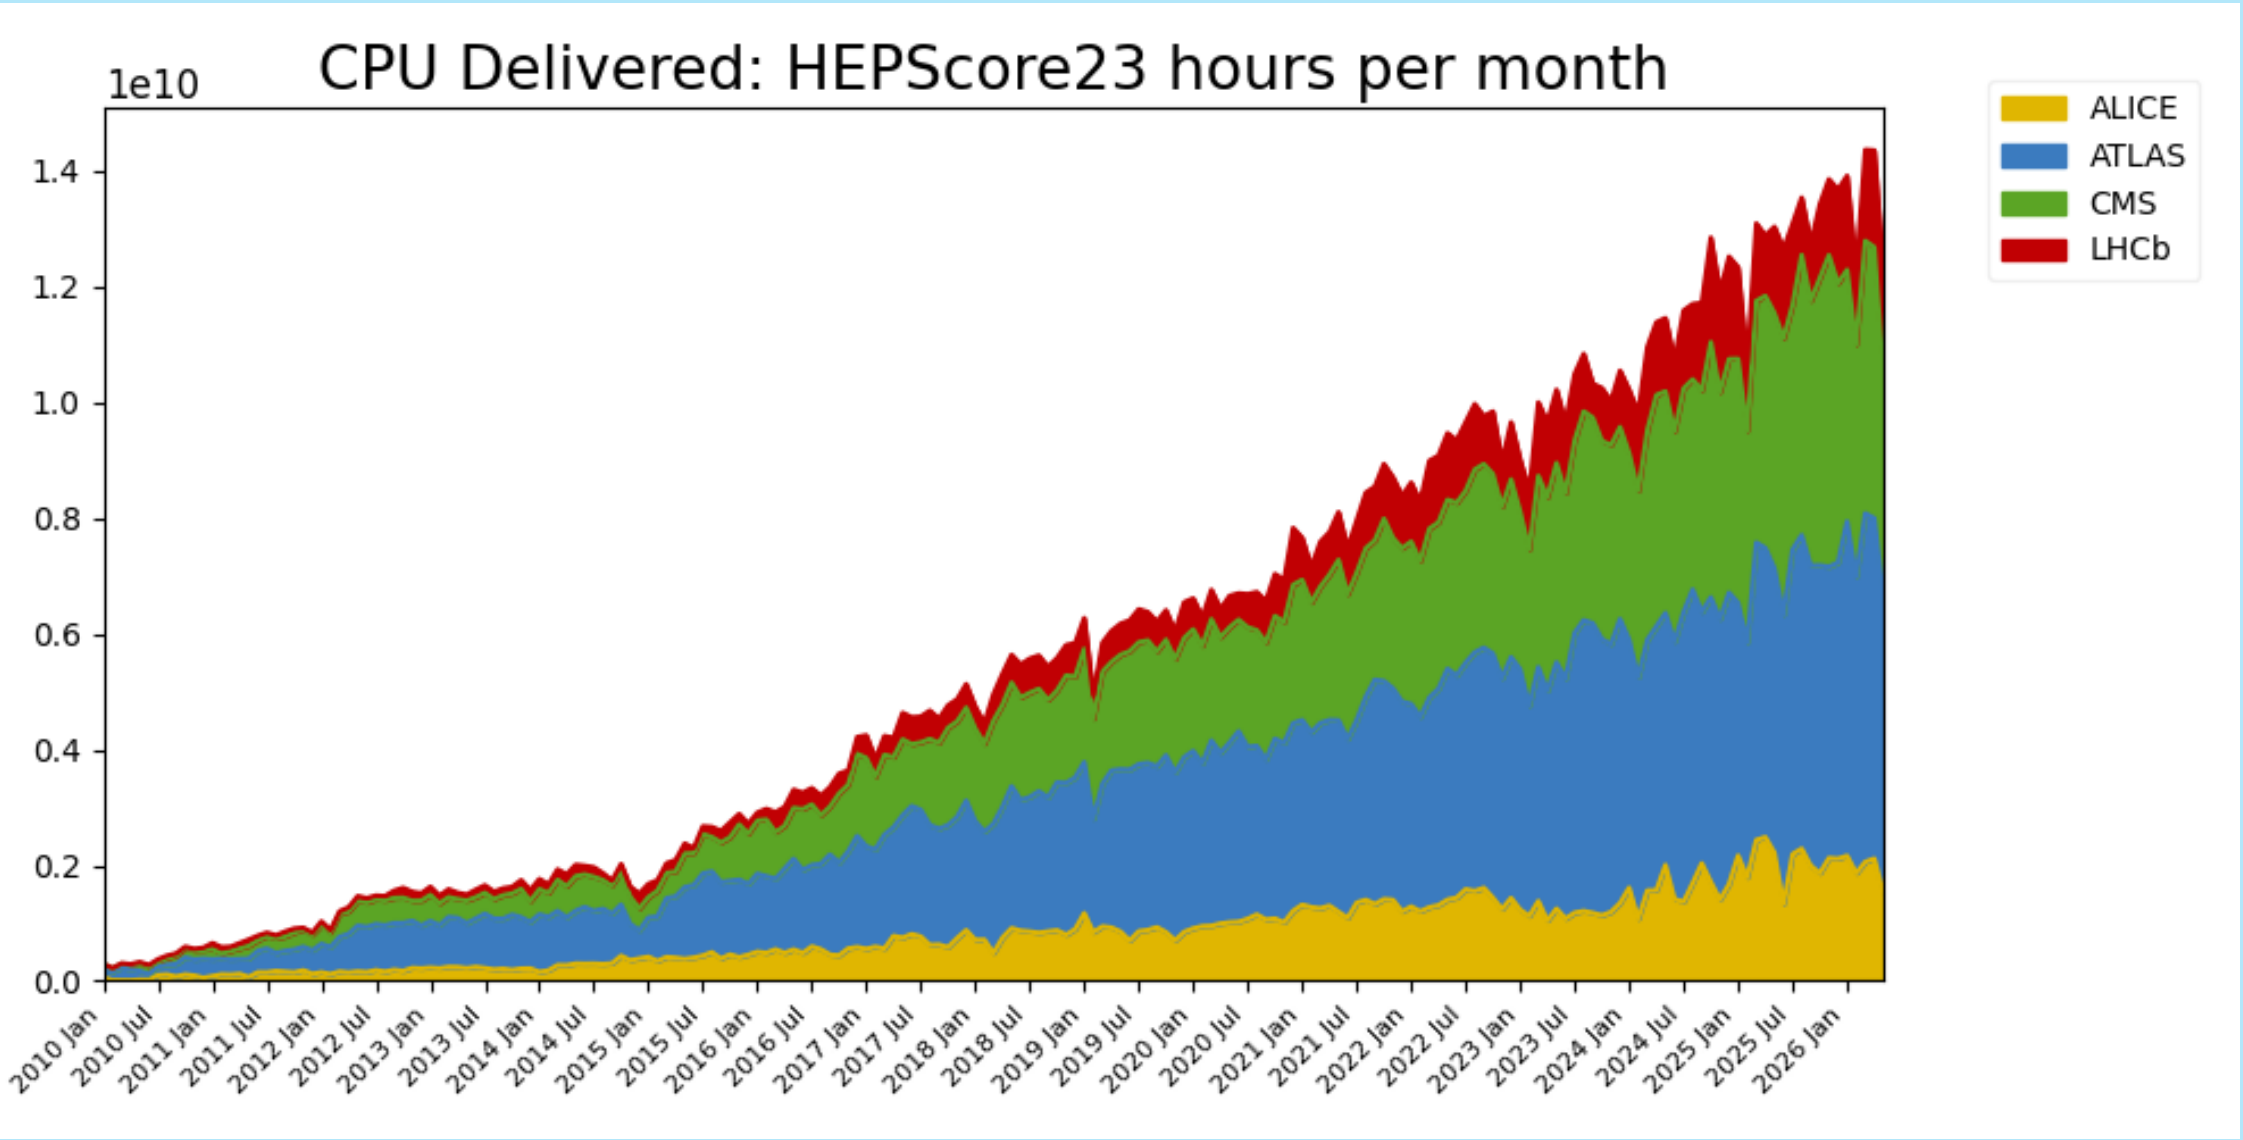

Ever increasing luminosity and pileup …

Focusing on ATLAS and CMS

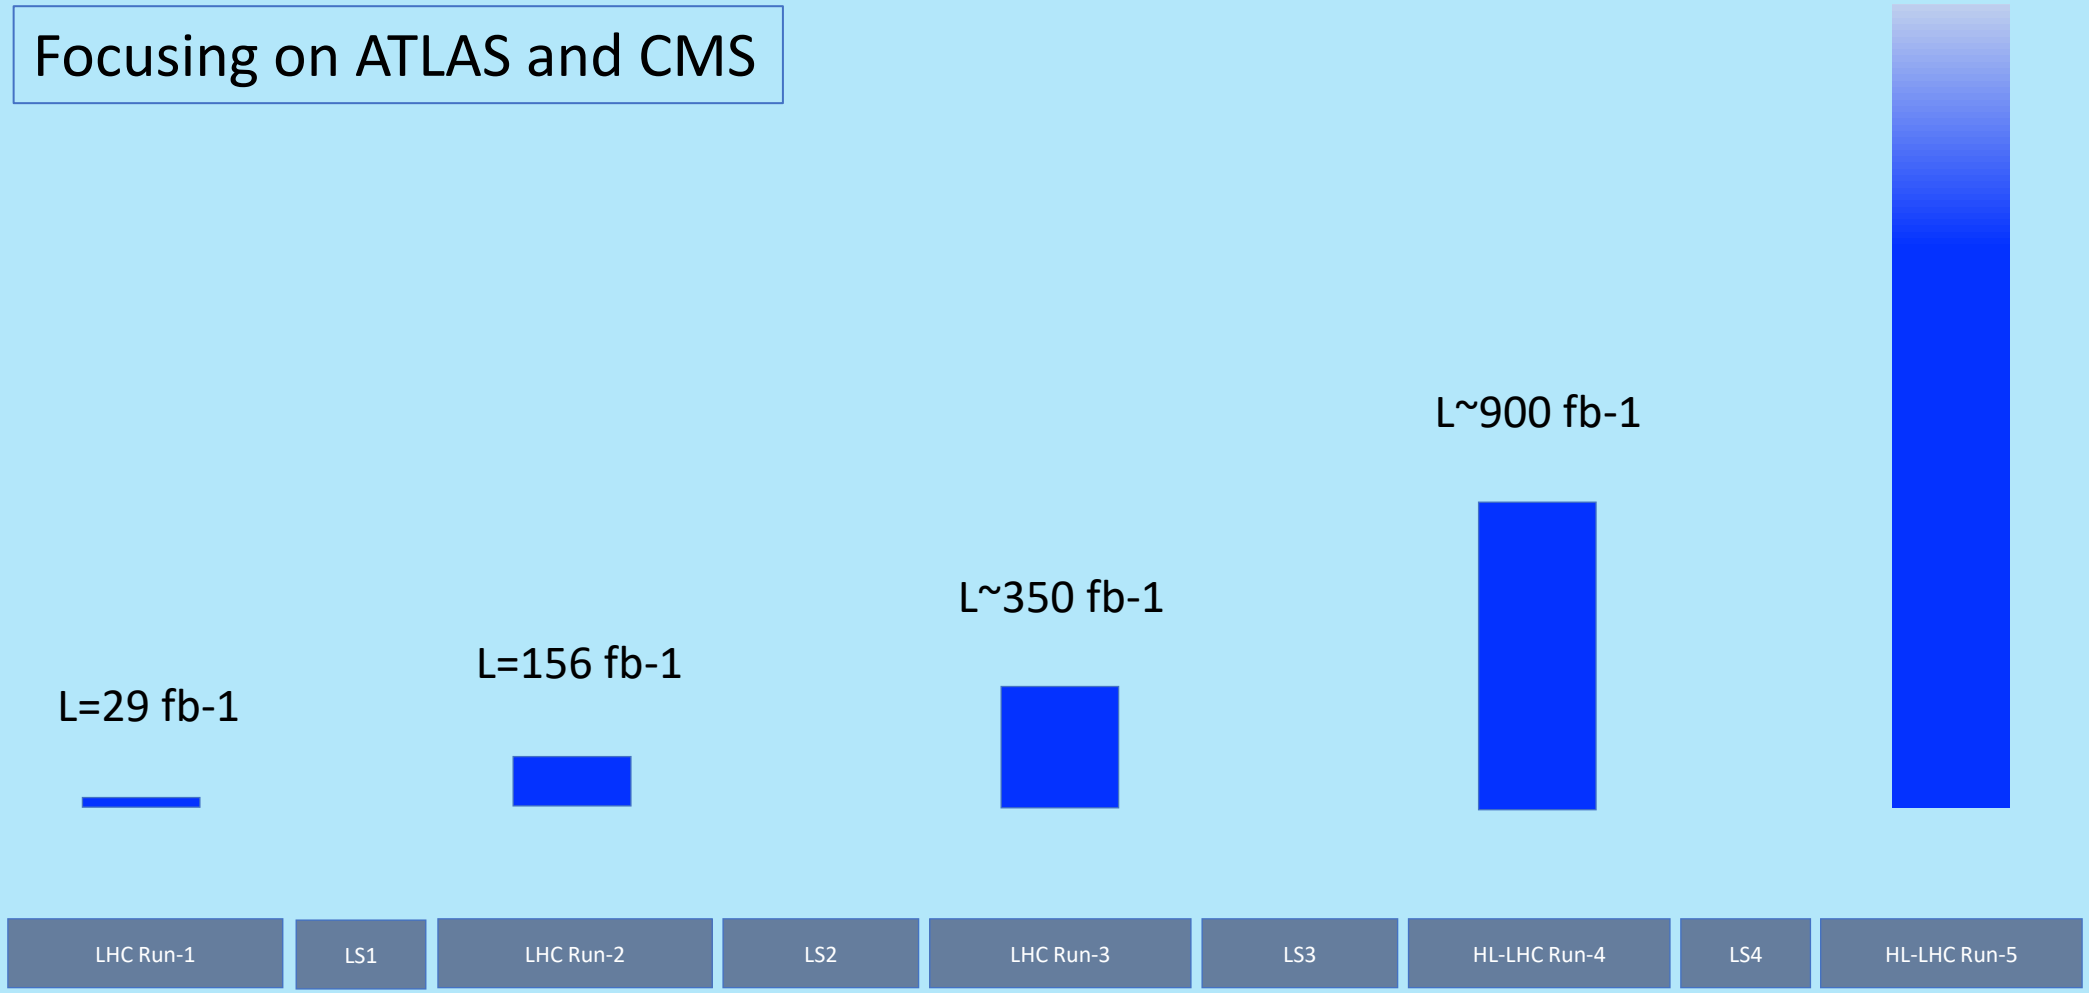

… leading to ever increasing processing needs

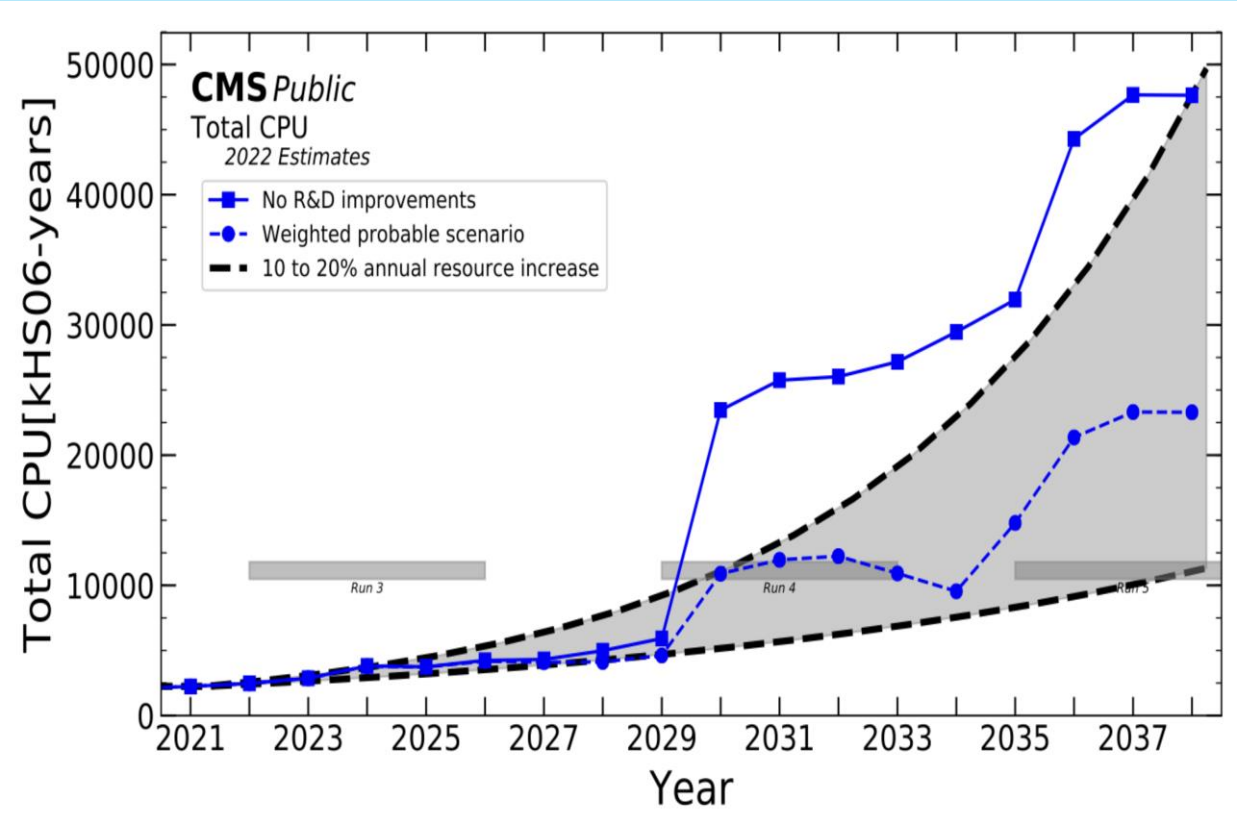

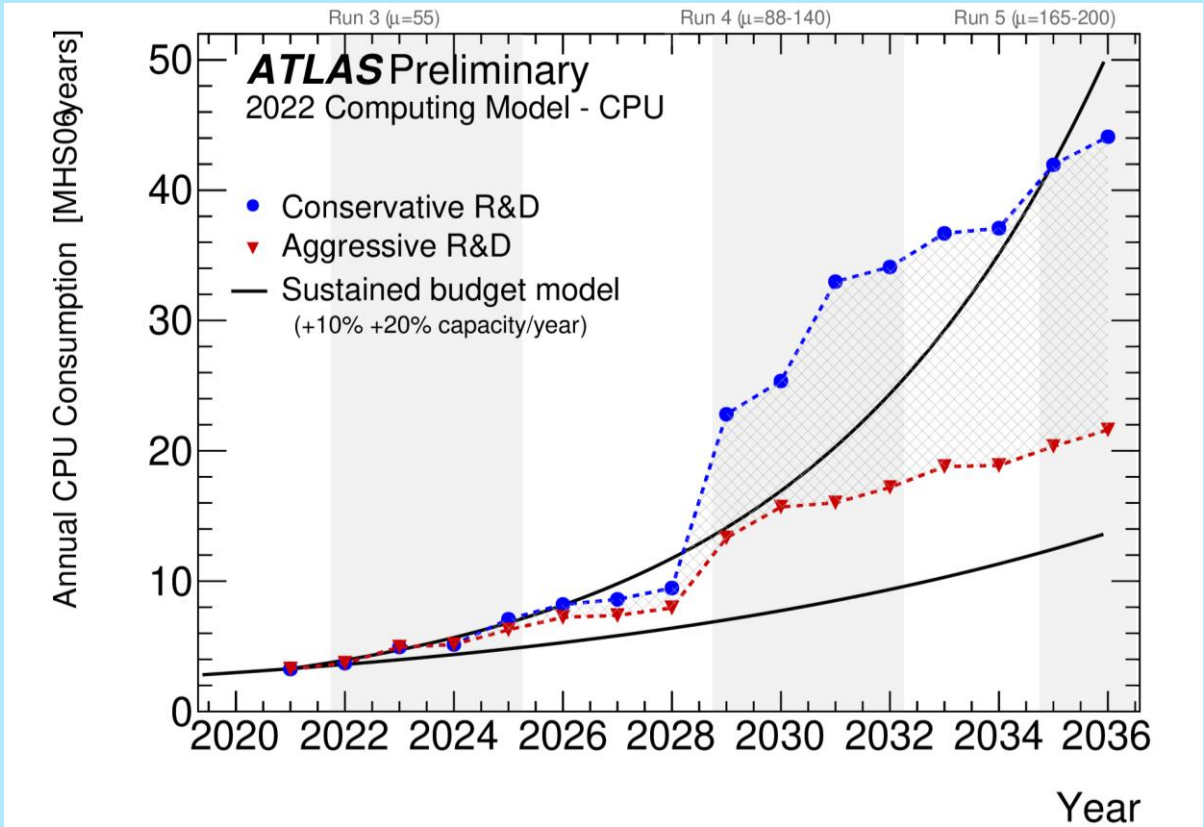

Estimates from before the AI boom!

… that require improvements in SW & services

Technical Evolution Roadmap

- A community-driven roadmap for technical evolution of WLCG is being prepared
- The document itself will also evolve in the next years, to take account of progress as well as external changes in the big-science computing landscape and beyond
- The first version will be published this autumn
- The audience includes everyone involved in WLCG in some way
  - In particular our LHCC referees and our funding agencies
- Please see this CHEP 2026 talk for further details

Computing model & SW evolution

- Improved data formats for less space and faster processing
  - RNTuple

- Alternative architectures for cheaper and/or faster processing
  - ARM
  - GPU
  - After very extensive physics validation!

Quantum computing?
→ Expected to ***complement*** classical systems, some day!

- SW development groups, foundations and projects foster R&D
  - SFT
  - HSF
  - IRIS-HEP
  - SWIFT-HEP
  - EVERSE
  - …

WLCG facility perspectives

- HW cost trends – badly hit by AI factories!
- Energy costs may also rise, depending on the country
  - Power *and* cooling
- *Alternative architectures* may be better in one or more ways
  - Work per second per **monetary** / **energy** / **sustainability** budget
- Energy may also be saved at peak periods
  - CPU frequency reduction
  - Active capacity reduction
- Future: climate-neutral data centers…
- Now: **UPB in Bucharest** is a new T1 center for ALICE !

Dec 2025 Workshop on Heterogeneous Architectures summary

Sustainability Forum

Analysis Facilities

- HL-LHC computing models foresee Analysis Facilities as a new type of resources
- Some WLCG facilities may want to offer such resources instead of traditional set-ups, or even in addition
  - Proper recognition needed in pledges and accounting
- Input data must not too quickly become the bottleneck
  - Fast local storage and network are needed (cf. HPC *burst buffers*)
  - Possibly in the form of caches, i.e. local data losses can be tolerated
- Profiting again from alternative architectures
  - ARM
  - GPU
- Users may access such services through notebooks
  - Auto-scaling from the laptop to the grid…

Benchmarking

- HEPScore is a flexible, modern benchmarking suite
  - Based on containers and a configurable variety of reference workloads
- HEPScore23 is a specific benchmark constructed from 7 modern HEP workloads
  - Released for x86_64 **and** ARM
  - Adopted by WLCG **and** EGI
- Under construction: **GPU** support, **sustainability** metrics, …
- CHEP 2026 talk

Accounting

- To support *new architectures*, *sustainability metrics*, and different *types of resources* (grid, HPC, cloud), the APEL legacy framework will be replaced with PEAR (Portal for Enhanced Accounting Records) at RAL and AUDITOR at EGI sites
  - Supported by the ENSURE EU project (September 2026 – 2029)
  - The legacy Accounting Portal will already be replaced in July!

- CHEP 2026 talk

High Performance Computing centers

- HPC centers are used by the experiments since many years
  - Mostly opportunistically so far
  - Also part of NDGF-T1 pledges
- Some countries may intend to pledge HPC rather than HTC resources
- Various barriers have been reduced over the years, some still exist
  - Operator access controls
  - Network access limitations
  - Lack of CVMFS
  - …
- Some can be worked around
  - ARC CE data handling
  - …
- EU and US projects help make HPC look more like our HTC sites

Cloud computing

- Experiments have been using commercial cloud resources through special projects since many years
  - Sponsored by big providers and EU projects
- So far, resources on our own premises have been less expensive
  - That may well change in the coming years, though
- For computing, cloud procurements can make a lot of sense
  - Potential accounting issue: no standard HS23 ratings for the underlying HW…
- For storage, we must be able to *abandon* our data!
  - All **custodial** storage must remain with our T0 and T1 sites
  - Data replicas in the cloud help speed up processing and analyses
- Integrating such resources has security implications today
  - Their CAs are not in IGTF
  - Perhaps we won't be needing our own set of CAs anymore in the future

Containers

- Pilot jobs of the LHC experiments run user tasks in Apptainer containers distributed via CVMFS
  - Provide legacy payloads with their expected environments
  - Isolate payloads to limit interference and potential fallout
    - With additional help from *cgroups* v2 where available

- Containers are also used increasingly for service deployments
  - Can come largely pre-configured
  - Pods with micro-services can be restarted or scaled up as needed
    - If the service design allows for that
  - *May* (!) help simplify the lives of site admins

- Victoria T2 – all services already on Kubernetes except storage
  - R&D project: EOS deployment via K8s on CephFS

Storage

- Erasure coding has matured and may allow sites to use a lot less disk space to provide good resilience against HW failures
- As an alternative to expensive tape *infrastructure*, T1 sites could consider looking into custodial *disk* storage
  - KISTI have pioneered such a service for ALICE, **UPB** is using it as well
- Industry has been looking into new custodial storage technologies (e.g. optical) for many years, but progress has been slow
  - Meanwhile tape is still looking OK – see this spring 2026 HEPiX talk
- Object stores and access protocols are becoming more popular
  - Ceph
  - S3
  - Experiment data model adjustments desired or required to take full advantage
    - RNTuple

Networks

- Dual-stack deployment of storage is at ~100% while computing is at 75%
- IPv4 will need to remain supported at least a few more years for *legacy* workflows that cannot handle IPv6
- IPv6 networks allow new features to be taken advantage of
  - Dynamic provisioning through *Software Defined Networks*
    - Allowing better use of available capacity instead of relying on overprovisioning
  - *Packet marking* and *flow labeling*
    - Allowing activities to be monitored separately, both between *and* within VOs
- New L3 VPNs like LHCONE are expected to be set up for partner projects
  - Our MW and network monitoring tools will also be shared where feasible
- Such networks also need to be interfaced with HPC centers and cloud providers
  - To avoid overloading general-purpose networks connecting sites to the internet

Monitoring

- Though many systems are used by experiments and sites, complete overviews at the WLCG level are tricky and incomplete

- For example, data traffic monitoring improvements, in particular for the *Xrootd* protocol (CHEP 2026 talk), concern many stakeholders
  - Experiments
  - SE and FTS developers
  - CERN MONIT
  - WLCG, OSG and EGI Operations

- MONIT is foreseen to be used even more for consolidation
  - With CRIC used as WLCG **information system**

- New systems keep emerging and gaining popularity for monitoring data center activities in general
  - OpenTelemetry
  - …

Authentication & Authorization (1)

- WLCG is progressing in transitioning from X509 user certificates and VOMS proxies towards using *WLCG tokens* instead
- Steered via the Authorization WG, Bulk Data Transfers WG and Token Trust and Traceability WG
- Pilot submission – HTCondor CE done, ARC CE in progress
- Data – in production for large fractions of FTS traffic
  - To be done: tape transfers, grid job payloads, user workflows
- End game (2028): users no longer need X509 certificates!
  - And do not need to know about tokens either!

Authentication & Authorization (2)

- As more experience was acquired, the WLCG Token Profiles have been updated accordingly: v1.1 (Sep), v1.2 (Dec), v1.3 (June)

- Accompanied by related publications
  - WLCG standard system JWKS cache
  - WLCG Token Usage Risk Assessment

- IAM v1.14.0 is the first version that allows access tokens **not** to be stored in the service DB
  - Should soon allow ATLAS to increase token usage for disk-to-disk transfers from 50 to 100%

- IAM v2 (ETA autumn) is based on a modern stack that will allow further requirements and enhancements to be implemented

Data Challenge 2027

- DC27 (Feb 22 – March 5) will allow us to test the infrastructure at the next level toward HL-LHC conditions
  - The aim was 50% → may not quite be feasible with the available HW

- To test improvements implemented after lessons learned from DC24 – some were tested in mini challenges at smaller scale
  - FTS4 will replace FTS3
  - Tokens will be used on the largest scale so far
    - Also planned for tape

- Next stop DC29
  - Aiming for 70 – 100% ?

Where can WLCG use AI / GPUs?

- Computing
  - Fast and full simulation workflows
  - Offline reconstruction: GPU-compatible fraction will increase
  - End-user high-level analysis → Analysis Factories?

- Operations – experiments, sites, central services
  - Anomaly / problem detection and, to some extent, automatic recovery by agents
    - Data / workflow management systems and services, …
  - Chatbots, expert systems
    - Ditto, GGUS, CRIC, EOS, …
  - Agents for routine use cases → what might we see in the next years?
    - Inspiration can be taken from this Belle II CHEP 2026 talk

- Coding assistants for SW development, testing, defect analysis

Operations Coordination

- Coordinating the **evolution** of the WLCG infrastructure is part of the mandate of the Operations Coordination team, in close **collaboration** with stakeholders
  - Experiments
  - Infrastructure provider projects
  - Sites
  - Middleware projects

- Various activities serve in different capacities
  - Weekly operations meetings
  - Monthly Operations Coordination meetings
  - Task forces
  - Working groups
  - Open Technical Forum and Management Board meetings
  - Workshops

Partners & Projects

- WLCG users and admins will rely more on *federated identities* to have access to vital auxiliary services through EGI Check-in
- For example: GGUS
- Current and new infrastructure and SW **partnerships** and **projects** will continue playing important roles
  - EGI, NT1, OSG, ASGC, IGTF, HEPiX, GÉANT, ESnet, ATCF, perfSONAR, …
  - EOSC Beyond, PATh, EuroHPC JU, PRACE, …
  - *And all the **middleware** and **experiment software** projects!*
- Infrastructure and middleware will continue being shared with ***other communities***, in particular through MoUs with various collaborations
  - Belle II, DUNE, JUNO, Virgo, SKAO, …

Not so fast…

- Security incidents may affect part of our infrastructure at times
  - Security experts from our infrastructure provider projects and sites **collaborate** to minimize the fallout
    - For example through proactive mitigation against SW vulnerabilities
  - WLCG sites can **collaborate** on incident response and prevention through the Security Operations Center WG
  - Sites can **contribute** through passive DNS SOC services
  - **Training** opportunities are provided in **EGI Conferences** (2026 session) and the CERN School of Computing (2025 timetable)

- We may be affected by upstream licensing or support changes
  - Case in point: the RHEL source code policy for the last 3 years
  - Which had consequences for AlmaLinux and Rocky Linux

Conclusions & Outlook

- Already for 20+ years, the Worldwide LHC Computing Grid has successfully provided the distributed computing infrastructure for the CERN LHC experiments.
- During that time, the WLCG has seen evolution in technologies as well as growth, to deal with ever increasing data rates.
- Those trends need to be made to continue, to allow the WLCG to take on High-Luminosity LHC data volumes as of 2030.
- Improvements across a wide range of services and software have been described, some of which already bring benefits as of today.
- Thanks to many **partners and projects**, not only do the LHC experiments profit, but many **other communities** as well !

Merci de votre aimable attention!