WLCG toward HL-LHC

Meeting Annuale ATLAS Italia Computing
16-17 December 2025

Ale Di Girolamo (CERN)

WLCG
Who, Where, and What

The WLCG Collaboration

WLCG is the collaboration providing global computing resources to store, distribute, process and analyze the LHC experiments data

WLCG brings together:

- **CERN** - Operates Tier-0 and provides core services and coordination.
- **LHC Experiments (ATLAS, CMS, ALICE, LHCb)** - Define computing needs, develop experiment-specific tools and workflows, and provide operational support
- **Computing Centers (Tier-1, Tier-2, Tier-3)** - Contribute storage, processing, and operational expertise.
- **Networking providers (NRENs & partners)** - connecting our sites and laboratories world-wide
- **Grid infrastructure projects** - national and regional collaborations, e.g. OSG, GridPP, NorduGrid, EGI, etc.
- **Software developers and technology providers** - Build and maintain middleware, data management, workflow systems, and other tools.
- **Funding Agencies** - Provide financial support for resources, infrastructure, and staff.

How does WLCG work

- WLCG is a loosely coupled collaboration, it operates through collaboration and consensus, rather than top-down decisions.
- Policies in Memoranda of Understanding (MoUs), activities through Working Groups, Task Forces and technical boards.
- Key coordinating bodies:
  - Collaboration Board (CB), Overview Board (OB), Management Board (MB), Ops Coordination, Technical Coordination Board (TCB), Open Technical Forum (OTF)
- WLCG Reporting structure:
  - WLCG sessions at LHCC and RRB (through C-RSG)
- Experiments are an integral part of WLCG:
  - WLCG is not something external - it includes the experiments and relies on their active participation.
- Preparation for future challenges:
  - HL-LHC requires evolving WLCG's technology, infrastructure, and collaboration models.

WLCG Computing Centers

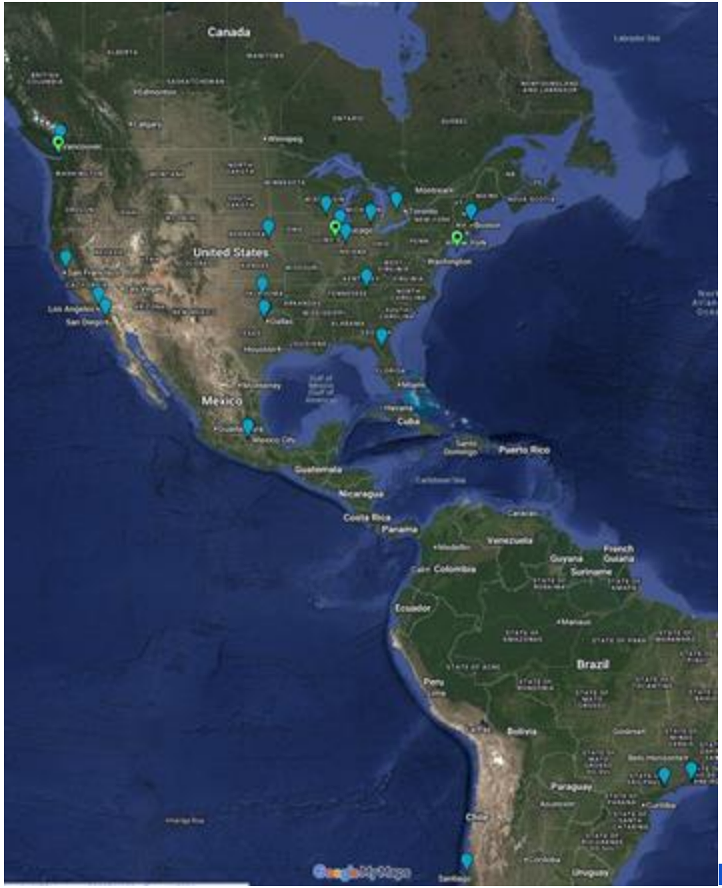

WLCG in numbers

65 MoUs, 40+ countries, 150+ sites

- meets the needs of the LHC experiments
- by the integration of globally distributed Exascale resources with services and software
- establishing a trust framework that transcends site and national boundaries.

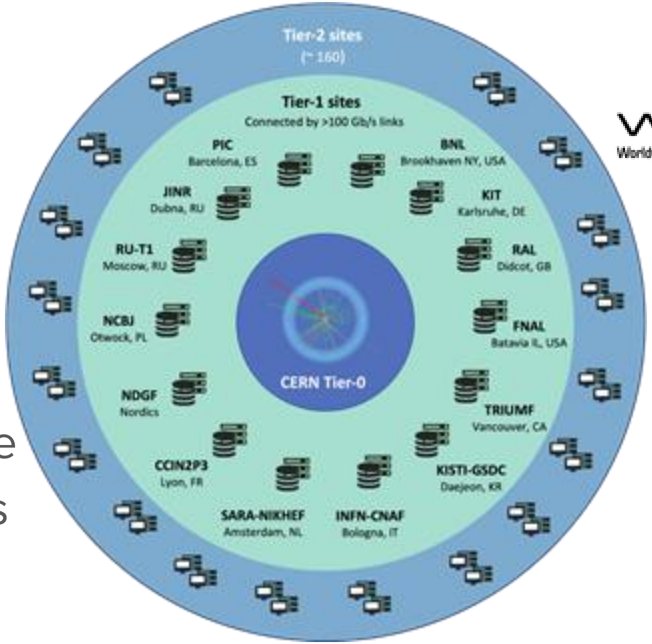

WLCG Resources 2025	
CPU (cores)	~1.3M (~13M HS23)
Disk (PB)	~1.3EB
Tape (PB)	~2.4EB

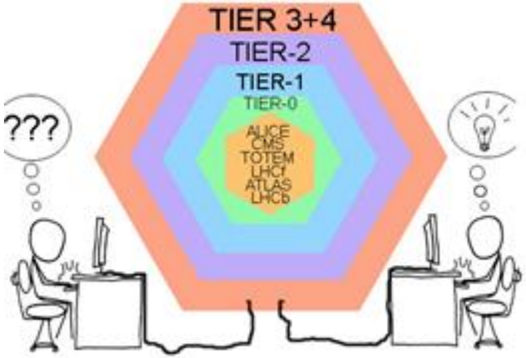

Running Cores - last 5 years - all facilities (Grid, HLTs, HPC)

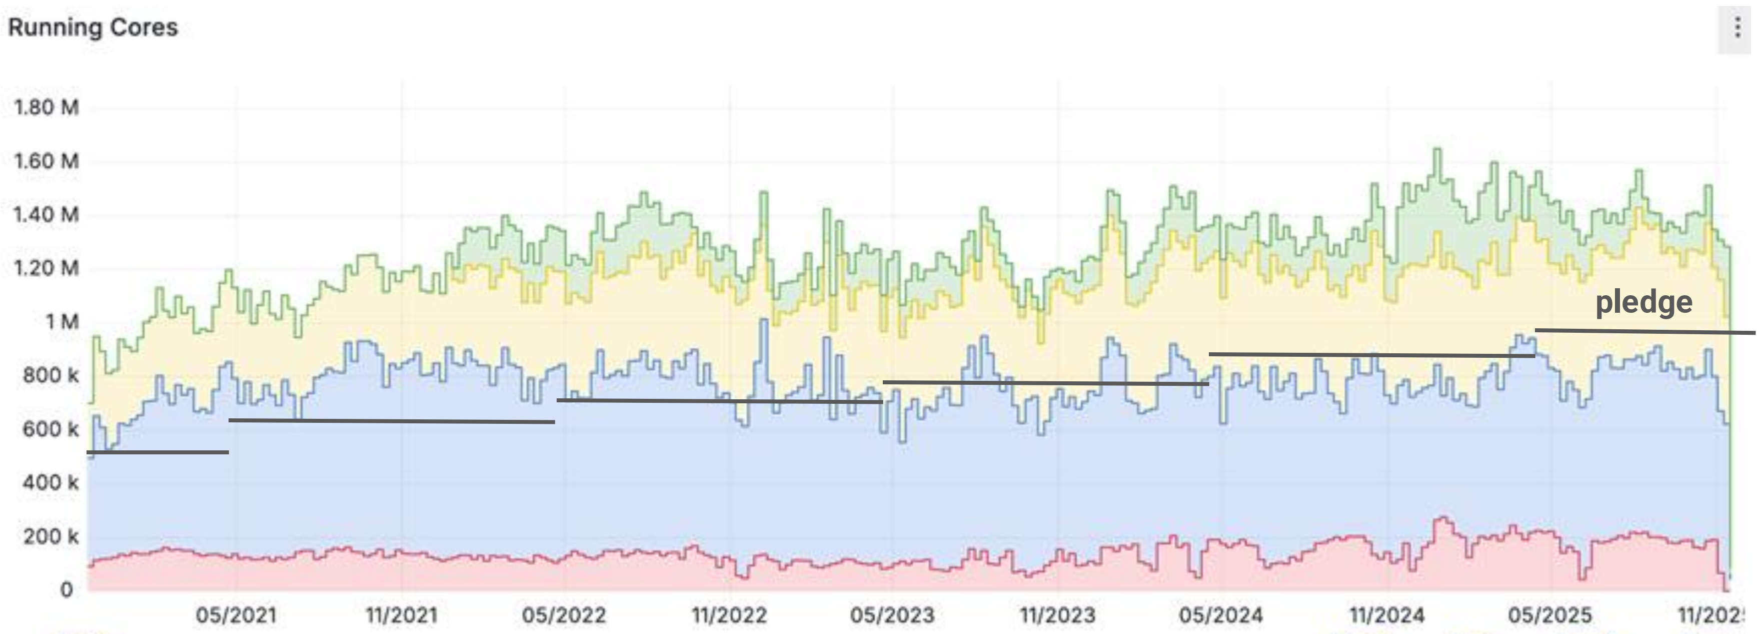

Name	Mean	Max	Last *	Total
ALICE	140 k	276 k	68.8 k	36.3 M
ATLAS	636 k	878 k	42.9 k	167 M
CMS	384 k	553 k	23.0 k	101 M
LHCb	143 k	315 k	18.7 k	29.1 M

WLCG Site-Site (FTS - WAN) Transfers - last 5 years

- Networking is one of the main pillars from which WLCG is built.

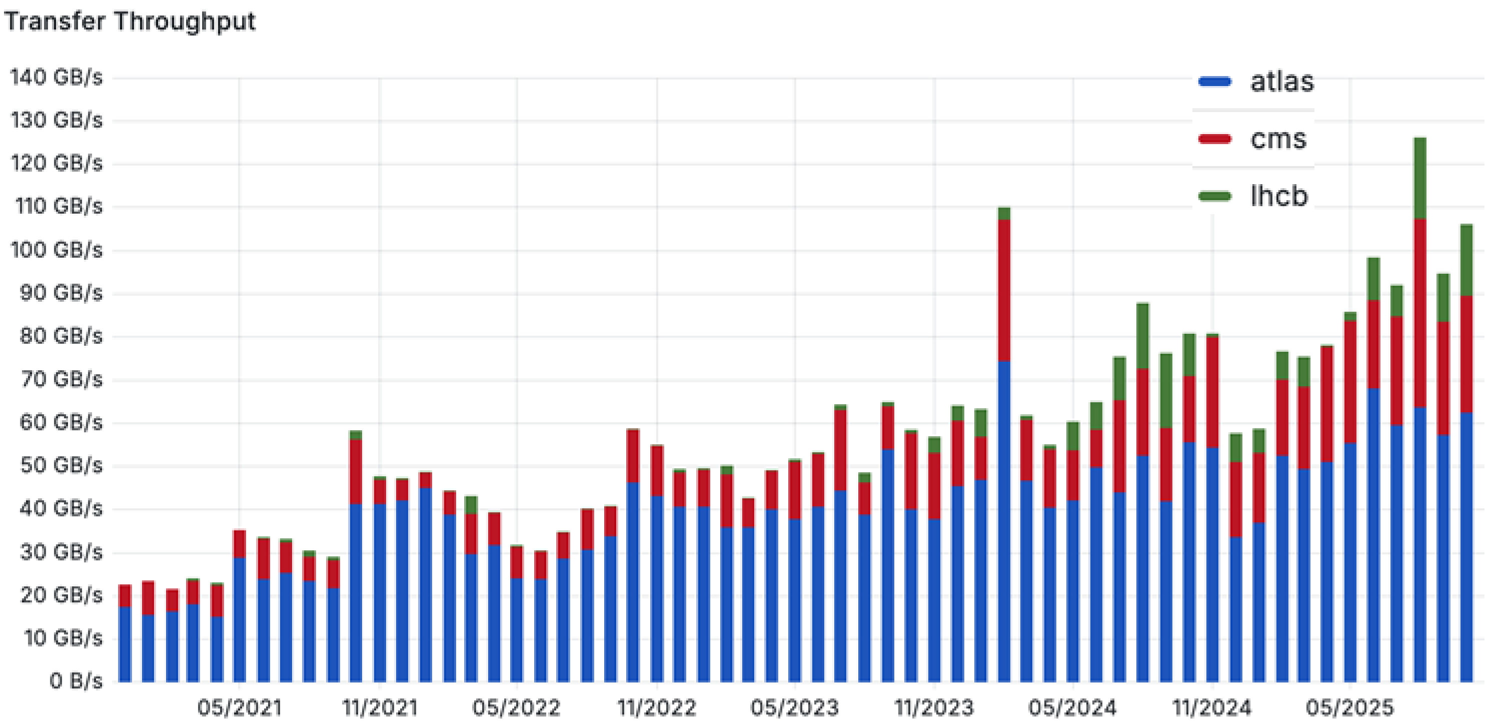

Storage from/to WN Data Transfers - a random site

- **Mostly - but not only!! - LAN**

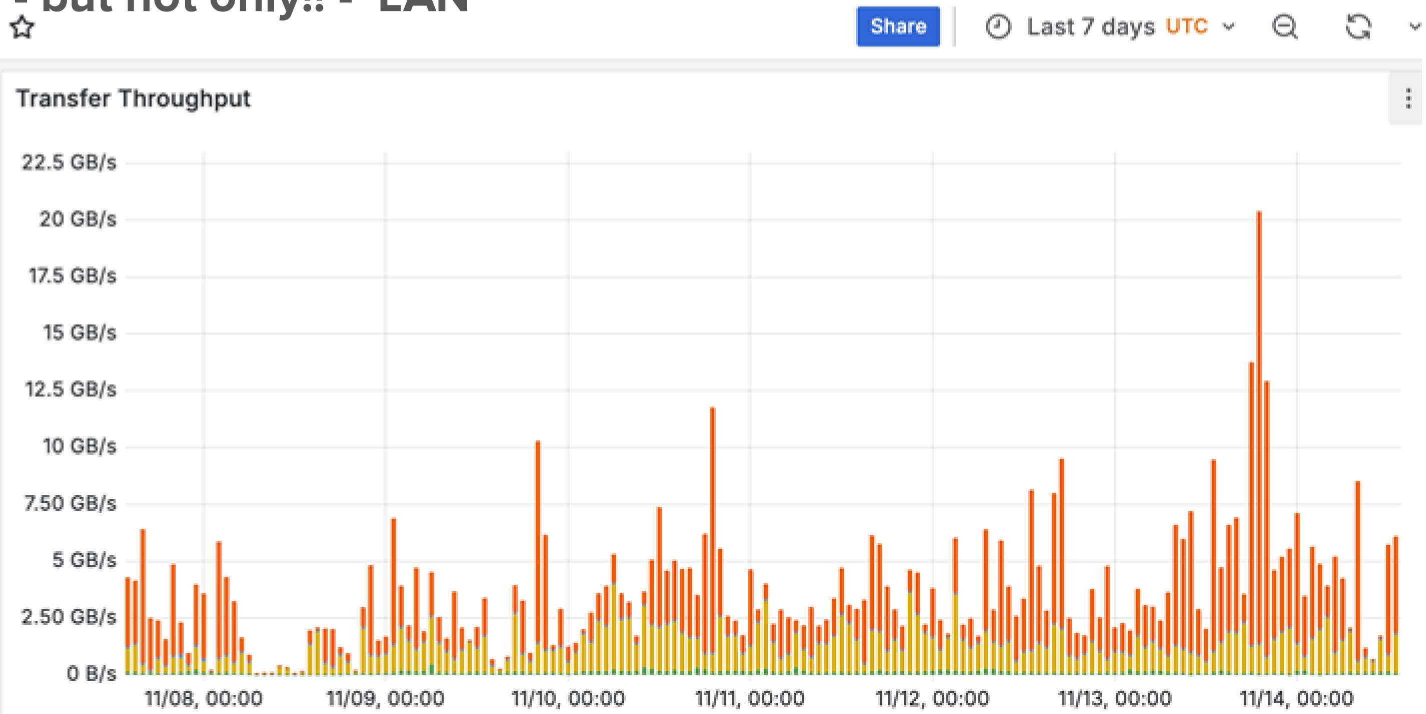

Name	Mean	Last *	Max	Min
Analysis Download	2.34 GB/s	4.26 GB/s	19.0 GB/s	7.08 MB/s
Production Download	1.21 GB/s	1.74 GB/s	3.81 GB/s	16.4 MB/s
Production Upload	123 MB/s	90.2 MB/s	449 MB/s	13.2 MB/s
Analysis Upload	19.9 MB/s	10.1 MB/s	136 MB/s	455 kB/s

Storage from/to WN Data Transfers - a random site

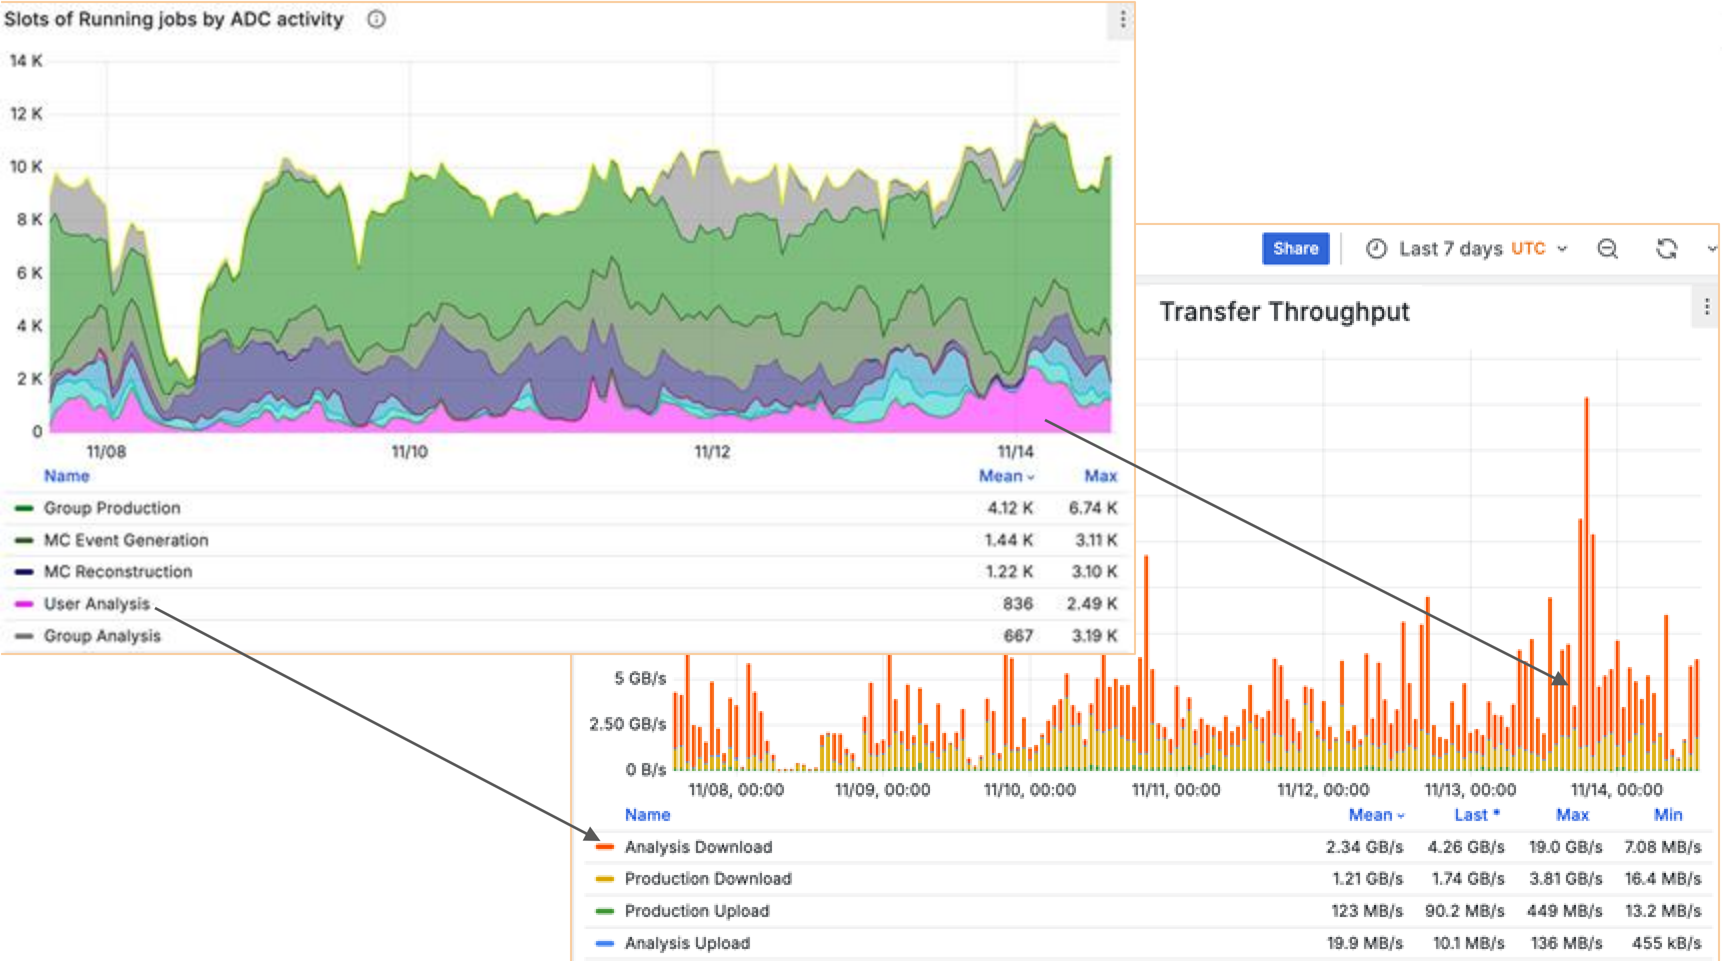

HL-LHC Scalability and Technical Challenges

The HL-LHC Challenges - Data Volume and Complexity

- Complexity of events $\propto$ pileup
- Data Volume $\propto$ integrated luminosity

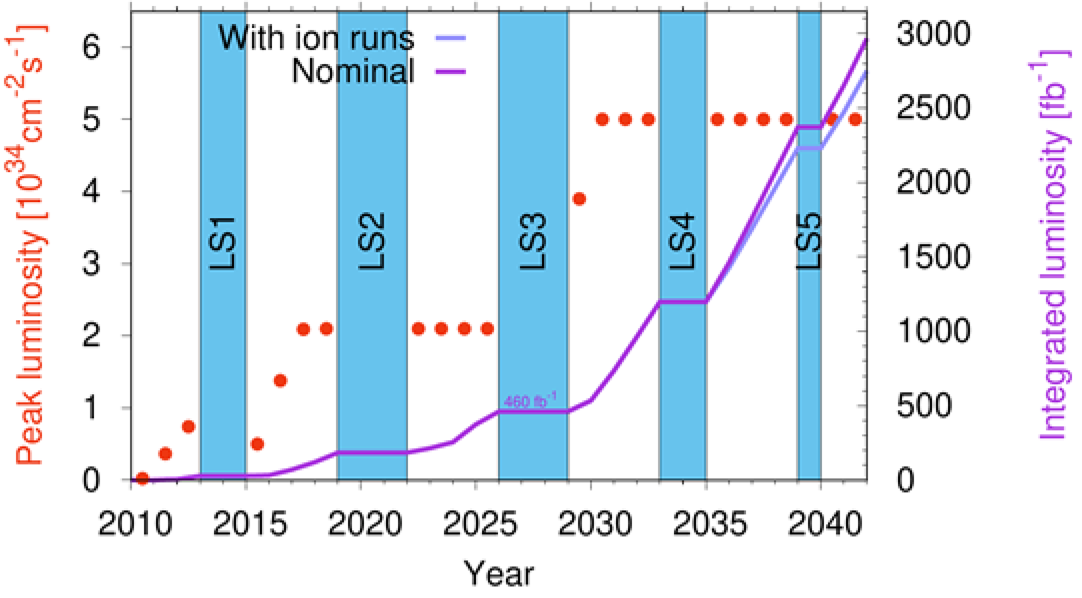

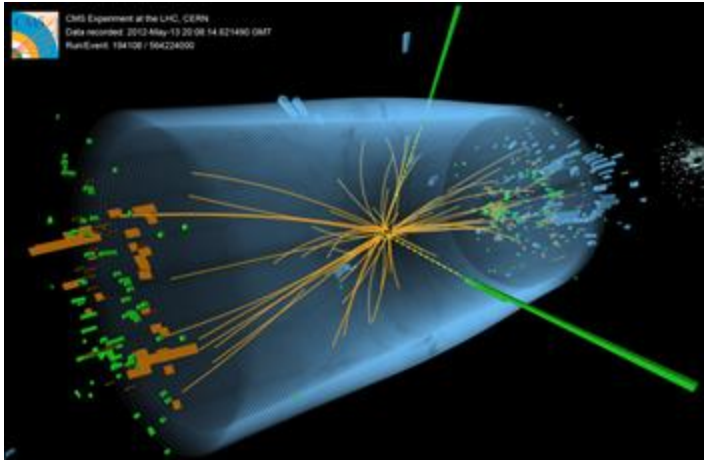

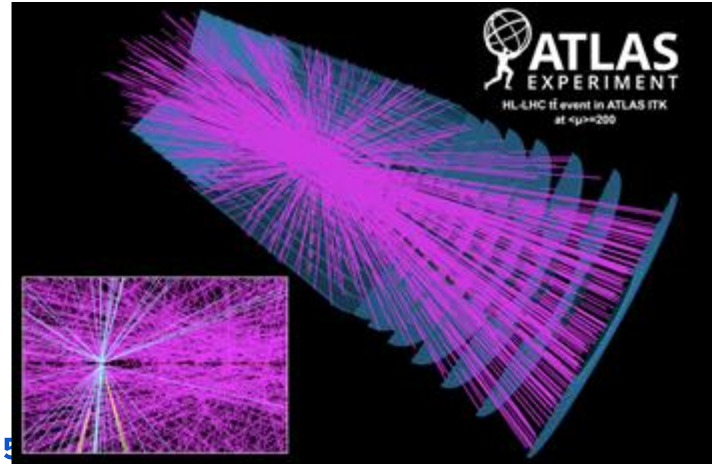

… some more Challenges

Technical and Operational

- **Facilities evolution**: planning for future architectures, storage technologies, and operational models
- **Hardware cost trends**: slower cost reductions than expected ➔ pressure on flat budgets
- **Data Management evolution**: evolving services to tackle the future needs.
- **Workflow management**: adapting (re-thinking?) orchestration and job scheduling to new workflows and new technologies (like K8s)
- **Heterogeneous resources**: GPUs, HPCs, Cloud - technical, operational, and cost
- **AI**: leveraging AI for operations (and physics) while managing complexity and skills gaps
- **Cybersecurity**: evolving threats and the complexity of modern security models (e.g. token transition)
- **Networking:** scaling bandwidth and resilience across LHCOPN/LHCONE, plus flow/packet tagging to improve traffic identification, monitoring, and optimization.

… some more Challenges

Organizational and Sustainability

- **Personpower and expertise**: attracting and retaining skilled people is harder; our environment is less unique and “sexy” than 20 years ago
- **Environmental sustainability**: reducing environmental impact
- **Adapting to rapid change**: WLCG’s size and complexity make it slower to adapt, while today’s tech world is fast-evolving ➔ balancing innovation with stability and operational reliability
- **Coordinating a loosely coupled global collaboration**: achieving consensus and common plans across diverse stakeholders

Data Archived at Tier-0 - looking into the future

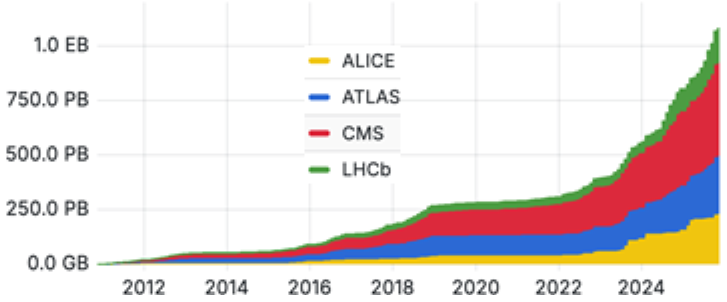

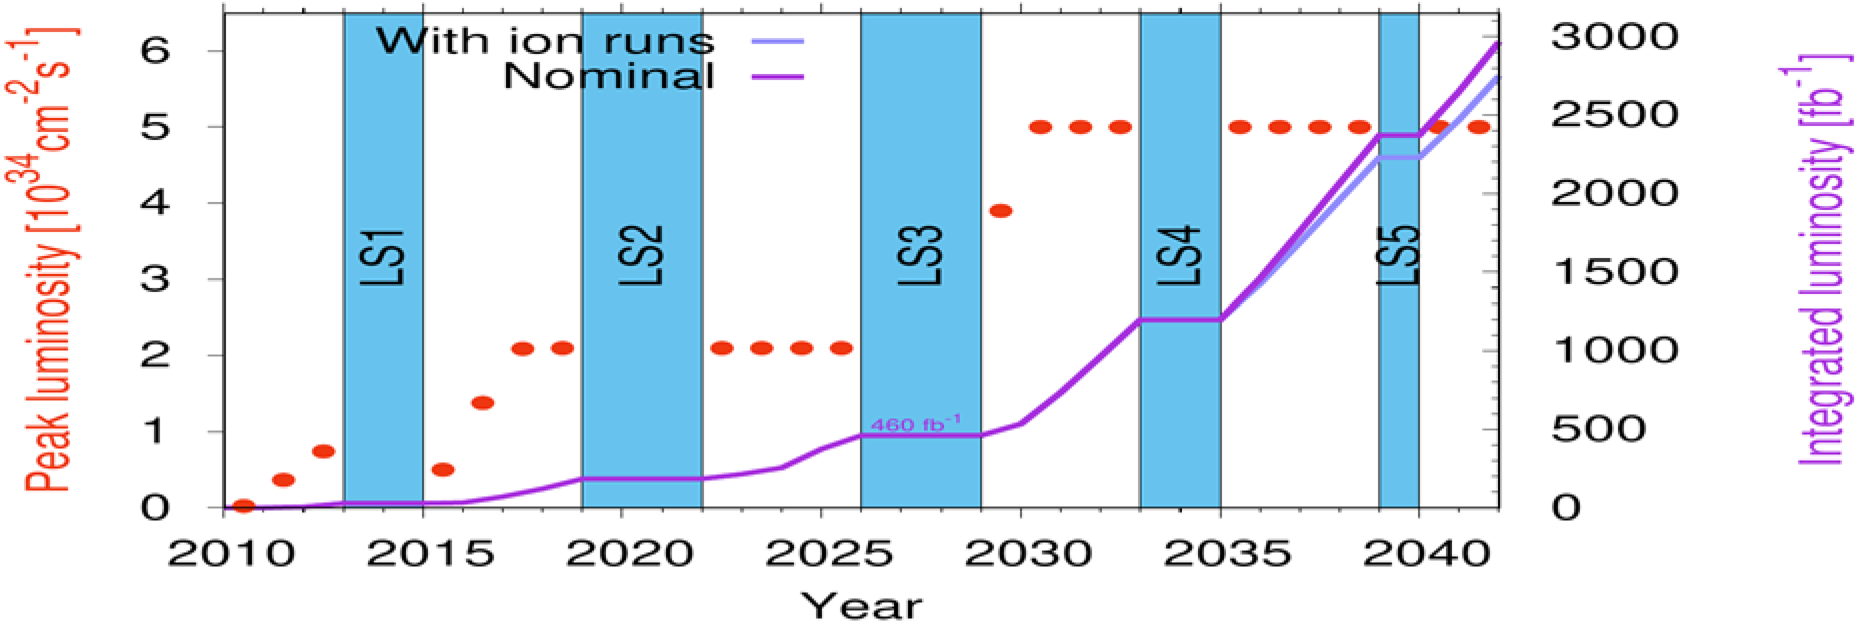

ATLAS and CMS resource needs up to HL-LHC

Sources: ATLAS and CMS

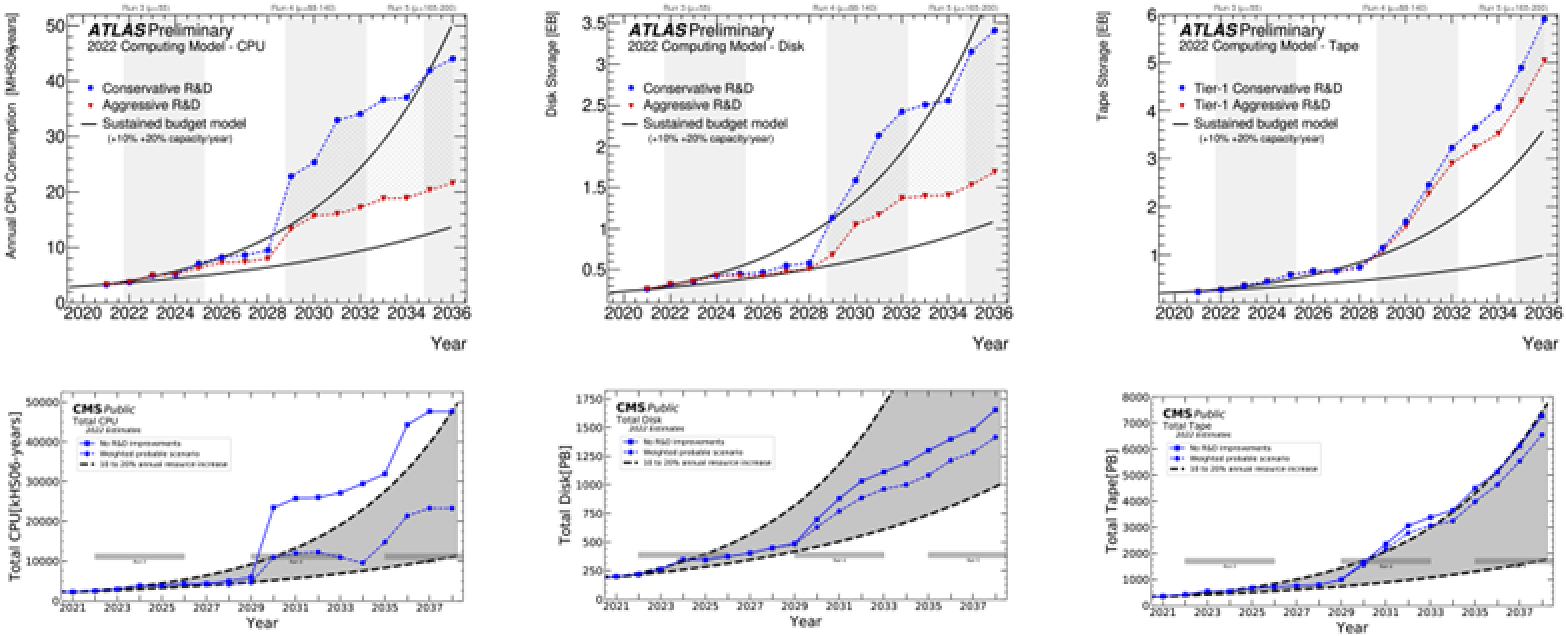

- From the experiment plots, with a flat budget, a 15%/year price reduction is required even in the most optimistic case (for CPU and disk)
- WLCG already decided to review the expected cost decrease to (10 ± 5)% / year for CPU and tape and (5 ± 5)% / year for disk

CMS new CPU resource estimates

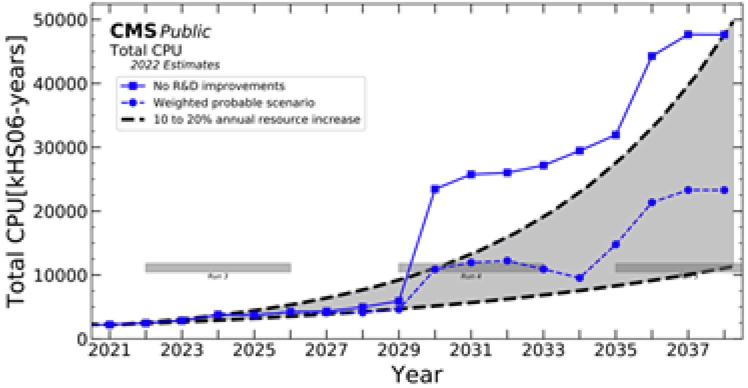

...we improve CPU usage in CMS?

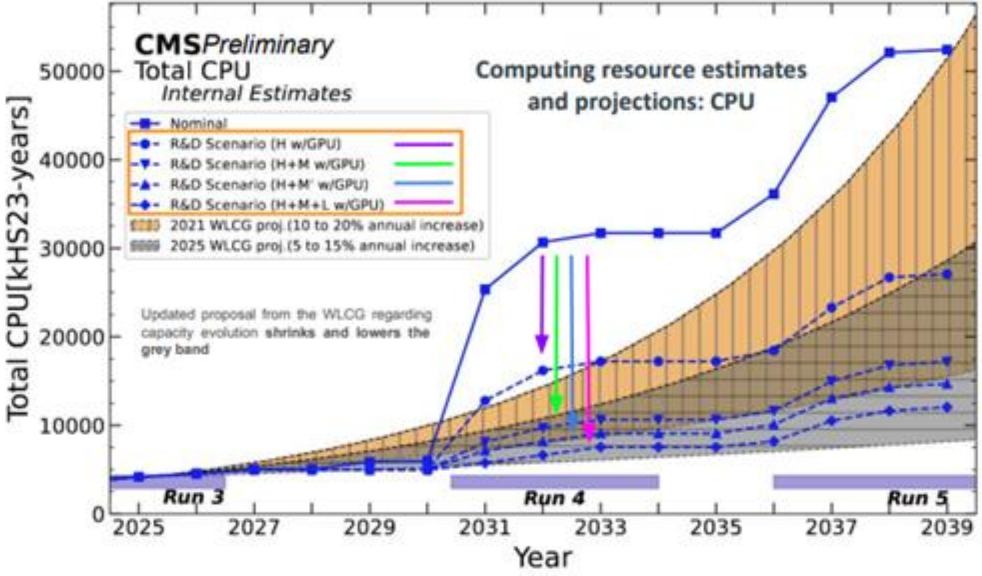

To keep up with the needs in processing power, R&D is being carried out including offloading of several tasks on hardware accelerators

Different scenarios are being considered for the fraction of code that could run on GPUs

* CMS will soon publish a **CDR** detailing its **Offline Software and Computing model for HL-LHC,** where many of the points mentioned today are described in detail

Hardware Cost - trend estimate

The cost per unit of processing power (CPU) or storage capacity (Disk) is still decreasing, but at a slower rate than in the past and significantly below the (now old) baseline expectation of around 15% per year.

Based on the numbers collected in the past years (more info [A.Sciaba @ WLCG workshop 2025](#)), the [official numbers for the WLCG estimates](#) of yearly cost reductions in a flat budget model were agreed to:

- CPUs (HS23/CHF): (10% ± 5%) yearly increase
- Disk (TB/CHF): (5% ± 5%) yearly increase
- Tape (TB/CHF): (10% ± 5%) yearly increase

.

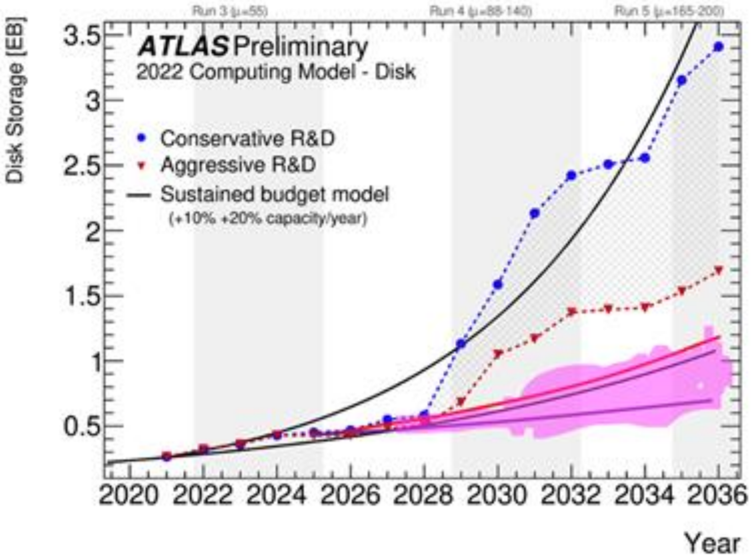

WLCG
Evolution

Tackling the Challenges

- Identify the gaps
- Develop a concrete plan
  ➜ Technical Roadmap
- Facilitate the coordination of stakeholders to execute the plan

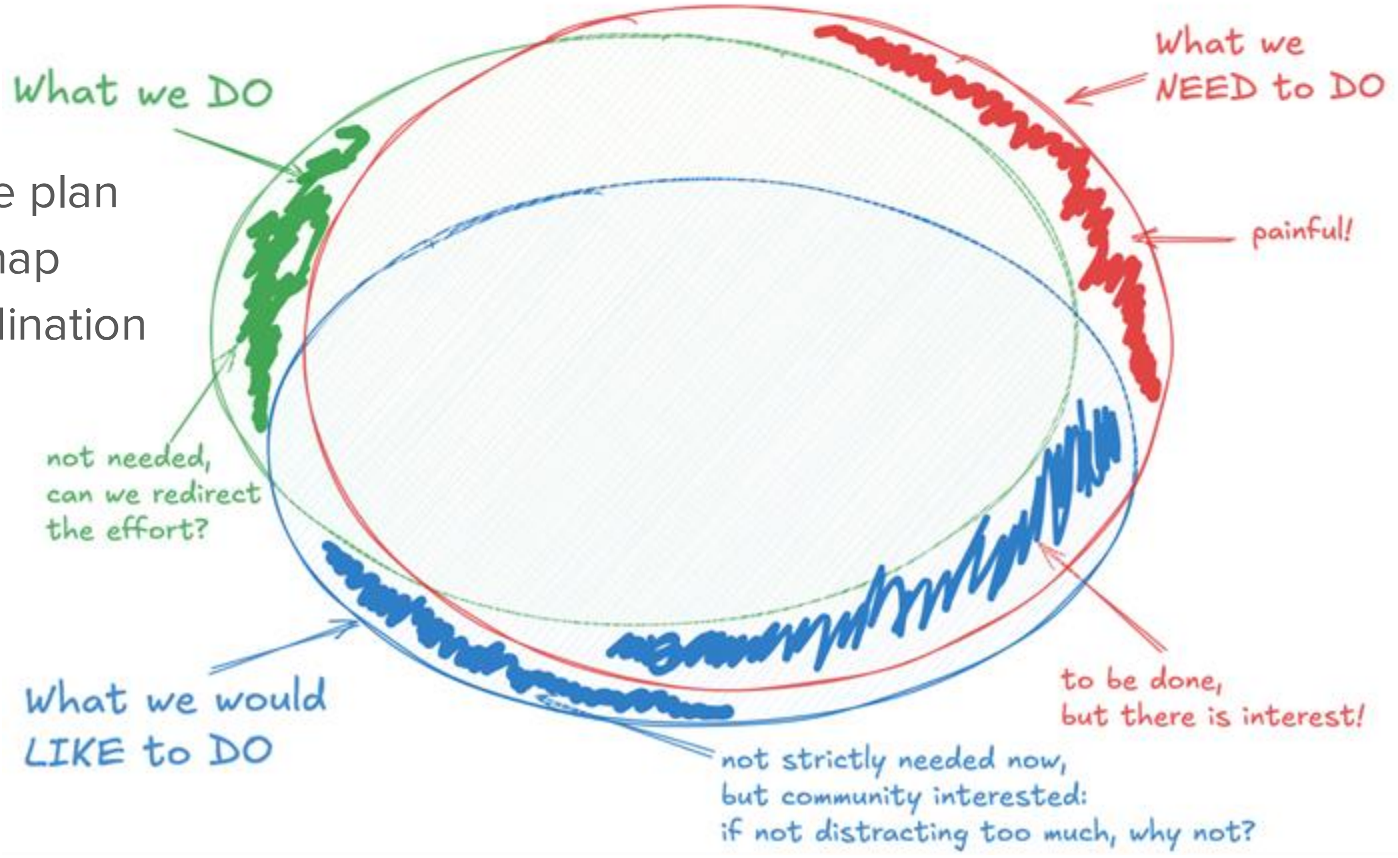

The WLCG Technical Roadmap

- Rationale of the WLCG TR:
  - To define and agree on our technical ambitions and priorities within the WLCG community, and
  - Demonstrate to reviewers the solidity of our plans and the progress
- A few considerations:
  - The chapters are mostly oriented around technical areas, but within each one have a problem-focused approach
  - High-level milestones for each area (important to have also the fine-grained milestones, but should be tracked internally)
  - Consideration of major risks, prioritization and any foreseen gaps in effort (no overall personpower evaluation).
  - TR is meant to highlight the gaps (and provide a plan to address them), not to detail the status quo ➔ avoid repetition of already existing documents (reference to them)
- Outline of the chapters discussed and agreed @ the WLCG Workshop May 2025.
  - Presented to LHCC in June - positive feedback on the process and the structure.
- Timeline as requested from LHCC: end of 2025 for a draft of some of the chapters. Before Summer 2026 for a full draft.
  - Agreed that this Technical Roadmap will be a living document.
  - Milestones will be reviewed once per year (e.g. in the WLCG workshop).

Operations and Innovation - not an easy balance!

- Balancing Operations and Innovation is a difficult but critical challenge for long-term success
  - hyper-focus on either can lead to stagnation or inefficiency

- “What if” we do not have (we are not moving toward) a reasonable balance?
  - Over-prioritizing operations may clip creativity and stall progress
  - Over-prioritizing innovation may cause instability and neglect of core responsibilities
- How can we balance Innovation and Operations?
  - Align (and prioritize) innovation with our strategic goals,
  - Adopt phased implementation
  - Foster cross-functional collaboration
  - Dedicate time and resources to innovation!

Data Challenges

- DOMA is the place where the (mini) Data Challenges are planned and discussed
  - Bottom up approach! https://twiki.cern.ch/twiki/bin/view/LCG/DomaMiniChallenges
  - CNAF participated in July (CMS stress test with CNAF new disks)
  - Next “big one” foreseen is DC27 (details still to be discussed)
- Very interesting presentation from CMS at the last DOMA General (end of Oct)

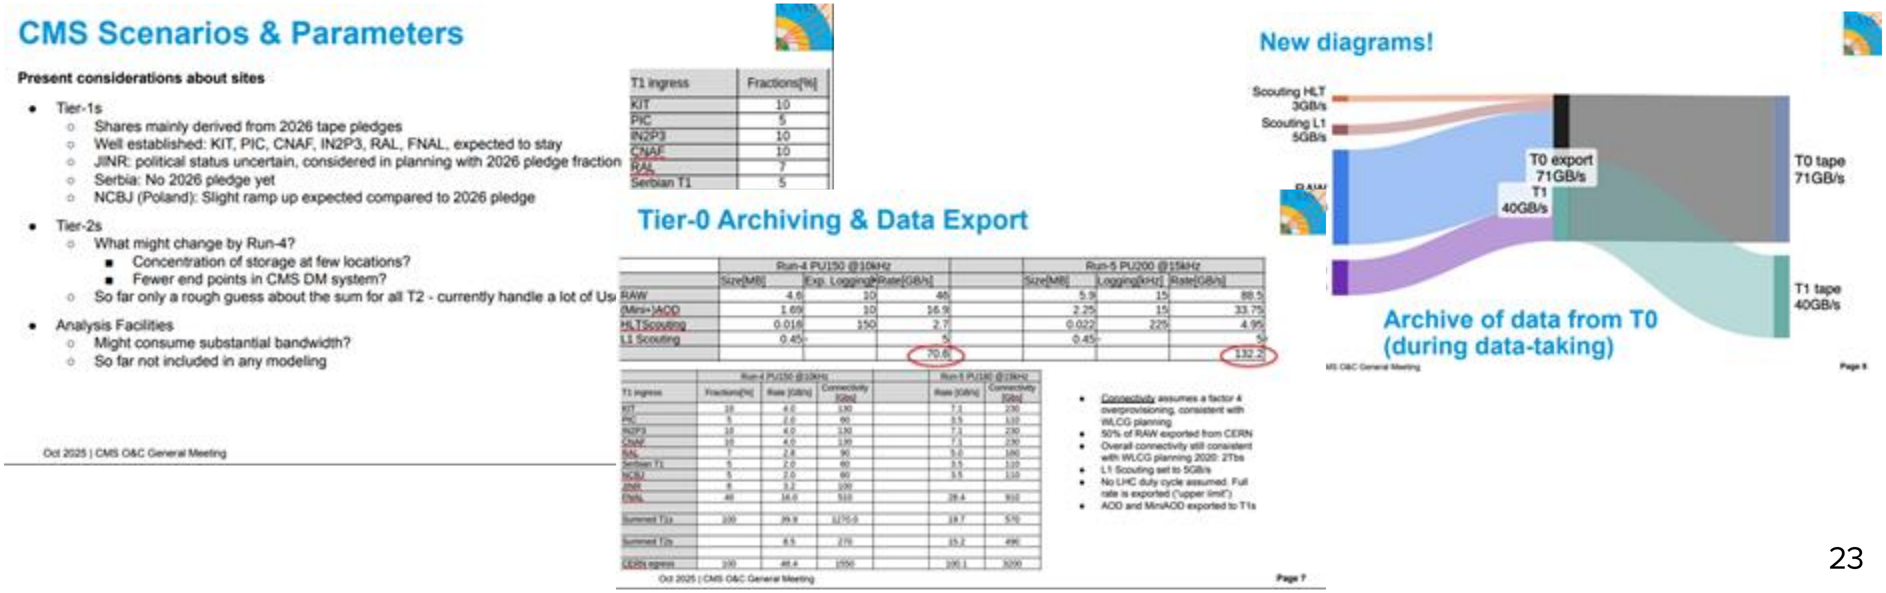

Heterogeneous Architecture

- "Nice" Workshop just a few days ago: summary in today's WLCG MB
- Lots of work ongoing: *"The Workshop concluded that while GPU acceleration is a vital R&D area across all major experiment workflows"… "no experiment is yet in a position to request or accept formal pledges of GPU resources from WLCG federations".*

9. **WLCG Information System**: Discovery and advertising of heterogeneous resources as well as an agreed-on description schema was seen as needing a proposal for an improved WLCG information system. The inherent heterogeneity itself (GPU models, CUDA driver versions, etc.) was seen as a challenge with respect to job failures, compilation difficulties, portability of code bases, etc.
10. **Monitoring:** Understanding how to improve monitoring the occupancy and performance of GPUs.
11. **Accounting**: Based on improved monitoring, Improving the accounting of GPU performance and utilization.
12. **Cost Estimates:** *N.B.* presented but not discussed further during the Workshop: we believe that it would be useful to establish shared cost predictions for GPUs, which could be done in the context of the HEPiX TechWatch working group. This can be very useful to guide WLCG federations in optimizing multi-year funding requests and procurements and the Experiments in their technology choices.
13. **Total Cost of Ownership**: It is important to understand the "total cost of ownership" of GPUs including power consumption, GPU lifecycle, failure rates, and embedded carbon. *N.B.* While not discussed during the Workshop, the work of both the WLCG Sustainability Forum and the HEPiX TechWatch Working Group could be very important here.

"I would like to help ..." non exhaustive list 1/2

- DOMA - planning (a-la CMS!) and running (mini)DataChallenges
- Support the work to define the useful metrics needed to monitor and account GPU. Benchmarking Working Group
- TechWatch Working Group: cost estimates, including GPUs
- Heterogeneous resources diagram. What are the analysis facilities for you?? I am still not sure!

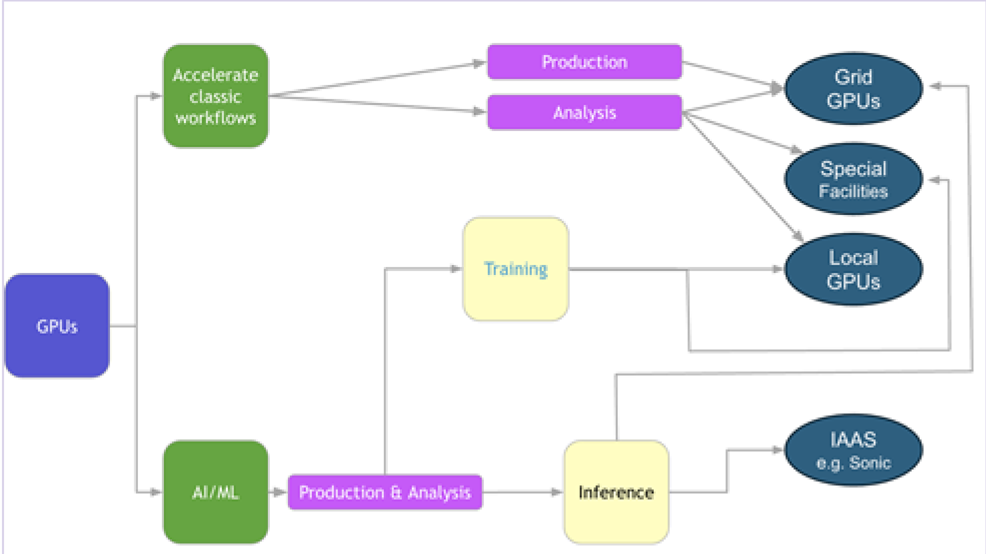

"I bet if that was a hard drive in a computer you could change it."

“I would like to help …” non exhaustive list 1/2

- WLCG Technical Roadmap:
    - Facility Evolution (chapter 2) , Workflow management (chapter 5) and new Architecture (chapter 8) would benefit of more engagement
- Facility Evolution: sharing the experience is very important - HEPiX
- IAM: CNAF is playing a leading role in IAM development, thanks! Still the Transition to Tokens is not yet done.

- Sometimes They Come Back: [Operational Intelligence](). Can this be imp… us, exploiting the tools now available?

“I would like to help - what can I do?”

A few open points (not exhaustive list!) - 1/2

- When is it more effective to rely on caching versus pre-placing files for different use cases (analysis, simulation, reconstruction), and how do we measure the gains?
  - Do caches provide enough resilience to network issues and load spikes, or does pre-placement still offer reliability that caching cannot guarantee?
- Storageless sites: while reducing local storage maintenance, does this shift operational effort elsewhere, and how do we manage new complexities for remote sites and network resources?

A few open points (not exhaustive list!) - 2/2

- For remote data reads, where's the balance point where latency, network dependency (or site strain) outweighs the savings from more centralized storage?
- Monitoring remote traffic: what could ROOT and XRootD partners do together to establish robust, detailed, and actionable monitoring of data access?
  - Can improved visibility help us optimize resource placement, discover bottlenecks, and support both cost and performance goals?

It’s going to be fun!

Grazie mille per l’invito!

EXTRA

E-Groups & Indico Categories

There are 3 relevant CERN e-groups:

- **`wlcg-technical-coordinators`** (to contact Ale & James)
- **`wlcg-technical-coordination-board`**
- **`wlcg-open-technical-forum`** (**open to self-subscription, please join!**)

and two new Indico categories (both public):

- WLCG Technical Coordination Board: https://indico.cern.ch/category/18888/
- WLCG Open Technical Forum: https://indico.cern.ch/category/18889/

The GDB e-groups will eventually be archived but for now we are also using project-lcg-gdb to publicize the OTF meetings.

TR - Chapters 1 & 2

Chapter 1 - Experimental Computing Requirements

Scope: As best they are known at this time:

- LHC parameters and experiment resource requirements for HL-LHC [TCB#4], including the overall network bandwidth and I/O requirements.
- Major similarities and differences of the experimental computing models (focusing on differing usage patterns of computing, storage, and network),
- Plans from experiments to leverage HLT resources, heterogeneous architectures, HPCs and Cloud resources, again focusing on major similarities and differences of the experimental computing models rather than repeating information found in other documents.

Chapter 2 - Facility Evolution

Scope:

- Technologies for management of disk storage (e.g., dCache, StoRM, EOS)
- Technologies for management of tape storage (e.g., CTA)
- Storage performance vs. capacity [OTF#2]
- Technologies for management of computing resources (e.g, (k8s, containers, batch systems) [OTF#5 TBC]
- Facility Deployment Plans [OTF#4, 2025 WLCG/HSF Workshop]

TR - Chapters 3, 4 & 5

WLCG
Worldwide LHC Computing Grid

Chapter 3 - Data Management

Scope:

- Rucio, FTS, DIRAC, plugins (for dCache, XRootD/EOS, StoRM, gfal)
- DOMA and Data Challenges [DOMA].
- Data Lakes and data "reservoirs", caches, origin edge-filtering services.

Chapter 4 - Network Infrastructure and Management

Scope:

- Details on network bandwidth and I/O requirements (wrt what is available now)
- Visibility (SciTags, packet marking)
- Optimization (e.g., jumbo frames, IPv6, etc.)
- Orchestration (SENSE/Rucio, NOTED) [RNT WG]

Chapter 5 - Workflow Management

Scope:

- Technologies for computing resource allocation e.g., HTCondor, ARC-CE, Kubernetes, ALICE VOBox, Panda/GlideInWMS/Dirac
- Late binding vs advance allocation (pilot vs HPC), future of pilot frameworks
- Entry point for workflow management: are we going to keep on having this way of working? Instead of sending jobs, do we see the need of sending "workflows"?
- Implications of new types of workflows (e.g., interactive analysis or ML training) on infrastructure design - and will we need specialized services to access new types of facilities?

TR - Chapters 6 & 7

Chapter 6 - Security and Authentication & Authorization Infrastructure

Scope:

- All aspects of the transition to token-based authentication: [OTF#1, 2025 WLCG/HSF Workshop]
- Security incident response workflows
- Security: infrastructure policy set: review and rejuvenation of existing security policies

Chapter 7 - Services

Scope:

- Software and Conditions Distribution (e.g., CVMFS, Varnish, Squid, Frontier)
- Monitoring and Accounting
- Information systems
- Helpdesk & Support

TR - Chapters 8 & 9

Chapter 8 - New Architectures and Infrastructures

Scope:

- Heterogeneous Resources: Cost effectiveness and integration challenges for GPU relative to CPU, as well as access via IRI, EuroHPC, HEPCloud, et al.
- Benchmarking and Accounting (e.g., HEPscore)
- Analysis Infrastructures and Analysis Facilities: Defining their role and purpose
- HPC Centers: Cost-Effectiveness, Technical Requirements, and Political Challenges
- Cloud Computing: Cost-Effectiveness, Technical Requirements, and Political Challenges
- Security aspects, risk identification and management of new infrastructures (e.g., k8s) and evolution of existing ones [OTF#3]
- AI and WLCG

Chapter 9 - Sustainability

Scope:

- Technical Sustainability
- People & Community (e.g. ensuring expertise retention)
- Environmental sustainability

Conclusions

Focus on: Capabilities, Interests and Needs

The WLCG Technical Coordination Board

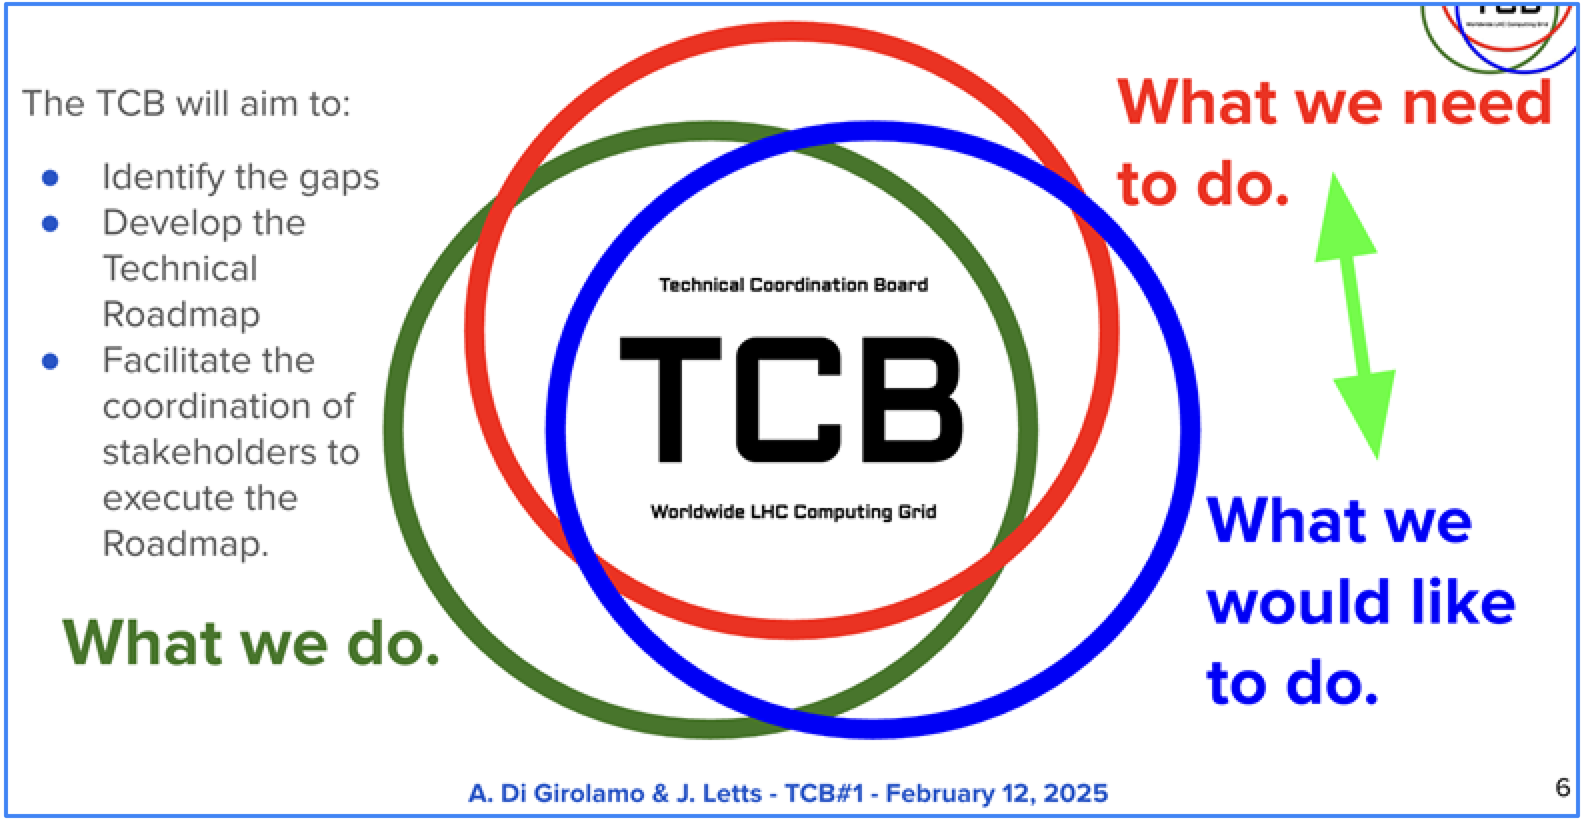

- Where do we need more help
- Mini data challenges. Take Katy Ellis presentation can someone engage in this specific activity? It requires punctual help, give large breadth of understanding, but yes, it’s work. DOMA meeting, twiki
- Evolution of facilities (chapter 2 facility evolution). We are a bit in trouble here. It is not clear how and where all the different FAs will go, which per-se is not an issue
- Root workshop some R&D ideas
- Heterogeneous Arch Workshop, notes, draft report to MB
- Analysis infrastructure: PhD half and half? Federated Identity as entry point for Analysis Infra? Sharing GPUs?