WLCG Technical Evolution toward HL-LHC

LHCOPN - LHCONE meeting #55

James Letts (UCSD) and Ale Di Girolamo (CERN)

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.

(Some of) the Challenges (not only for WLCG)

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.

WLCG (FTS) transfers - last 5 years

- Networking is one of the main pillars from which WLCG is built.
- See the next presentation about the technical roadmap for networking in WLCG.

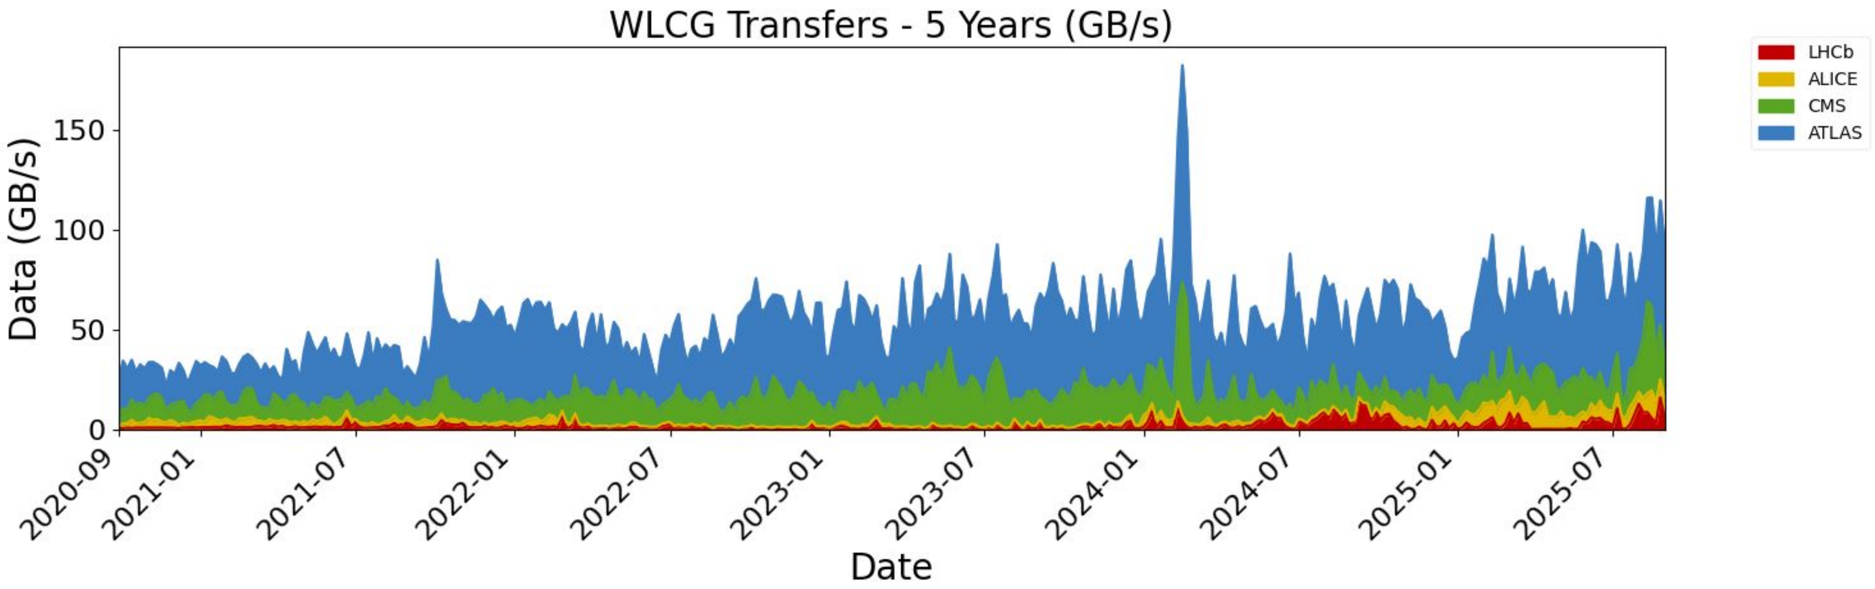

(Some of) the Challenges (not only for WLCG)

Worldwide LHC Computing Grid

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

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

.

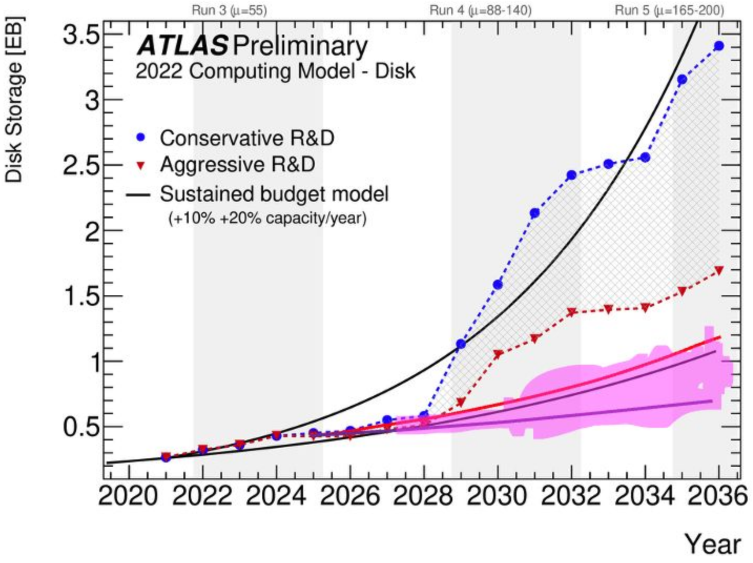

Tackling the Challenges

- Identify the gaps
- Develop a concrete
- Facilitate the coordi

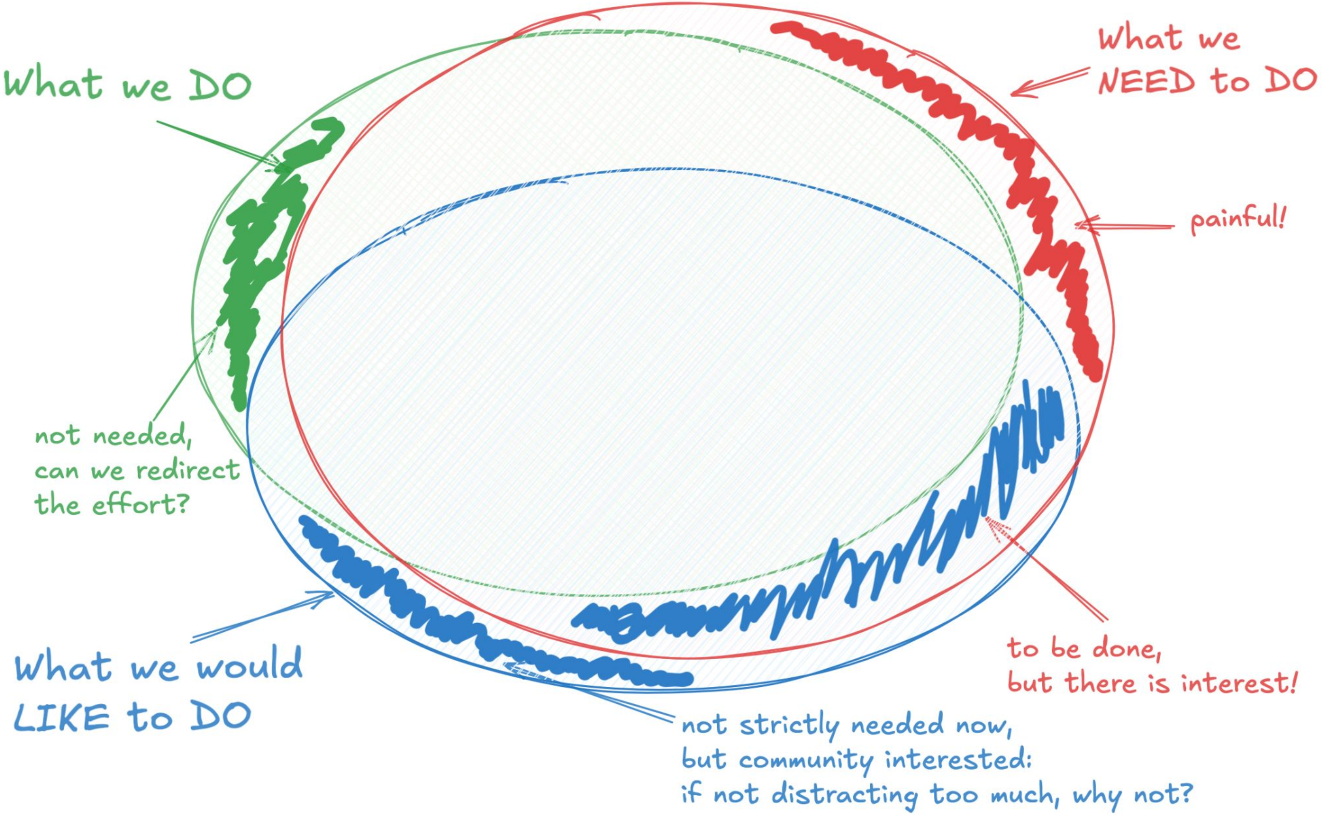

Focus on: Capabilities, Interests and Needs

The TCB will aim to:

- Identify the gaps
- Develop the Technical Roadmap
- Facilitate the coordination of stakeholders to execute the Roadmap.

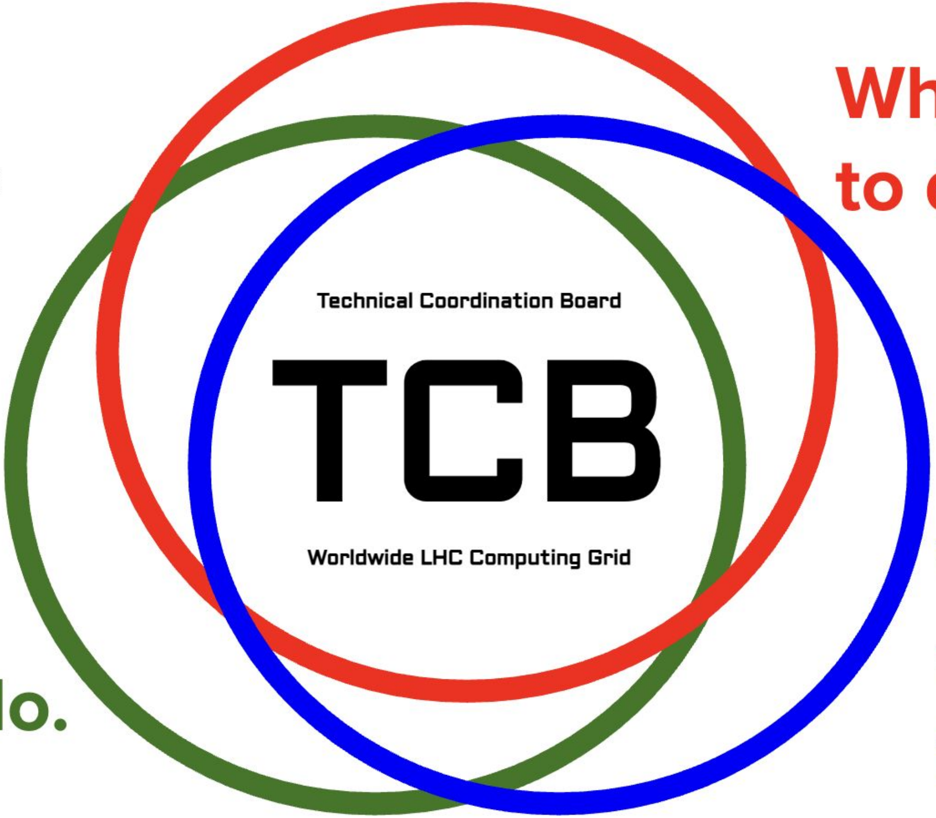

What we need to do.

What we would like to do.

What we do.

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).

Technical Roadmap - the outline

WLCG
Worldwide LHC Computing Grid

- Flow begins with the Experiments' requirements, which drive the entire process.
- Following are the Facilities, which bring their own requirements, which are essential to fulfill those of the experiments.
- New processing architectures and infrastructures might have longer term milestones
- Sustainability (both technical and environmental)

Table of Contents

Outlook

- Networking is one of the building blocks connecting our infrastructure.
- It has been working very reliably
- For HL-LHC we are going to scale up - and we (WLCG) are not going to be the only network-hungry science.
  - Monitoring is going to be even more important than now
- We fully trust this community to keep on supporting WLCG in this journey.

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

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

WLCG
Worldwide LHC Computing Grid

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