WLCG Operations: evolution and challenges on the way to HL-LHC

Julia Andreeva, Maarten Litmaath, Panos Paparrigopoulos

Outline

- WLCG Operations, what is it and how does it work?
- What will WLCG operations look like in 5 years?
- What have we achieved so far?
- What are our future plans?

What is “Operations” today?

Evolving WLCG while running WLCG

- WLCG is a complex infrastructure
  - ~170 loosely coupled sites
    - Different choices on tech and ops models.
  - More than 1.2 million running cores on average
  - 5 to 10 million transfers per day
    (average throughput 60GB/s)

- Operations on this scale is a challenge!
  - Experiments are doing a lot of operations internally.
  - Site teams are also invaluable and there are contributions by EGI and OSG ops.
  - In contrast to the amount of resources, the central ops team is less than a handful of people, concerned with central tasks and services that are common to several experiments:
    - Accounting, auth, monitoring, deployment campaigns…

How do we operate the GRID?

- Coordinating 170 sites, supporting 4 different experiments and using software from a very diverse set of teams:
  - Upgrade campaigns are far from trivial in such a landscape
- Incidents will always happen:
  - Making sure the many tickets and any Service incident reports are being followed up as needed.
- Developing tools needed for central operations: topology system, accounting platform etc..
- Creating an environment where operational related tasks can be discussed and move forward:
  - Operations coordination meetings, weekly and ad-hoc ops meetings etc…
- The many scientific papers published every year by LHC experiments imply we are succeeding!

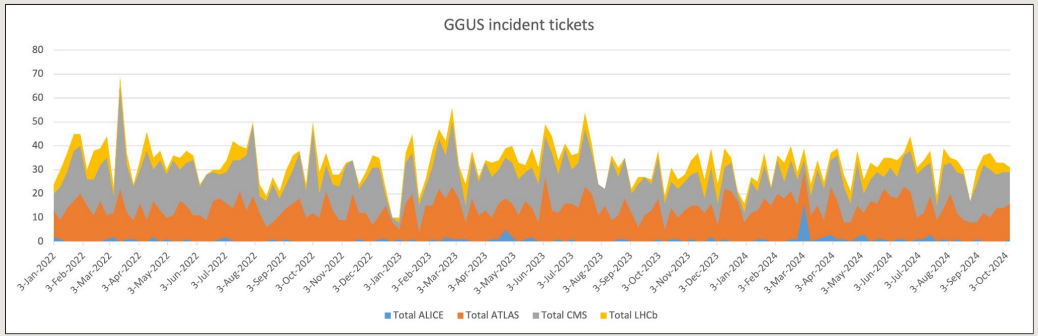

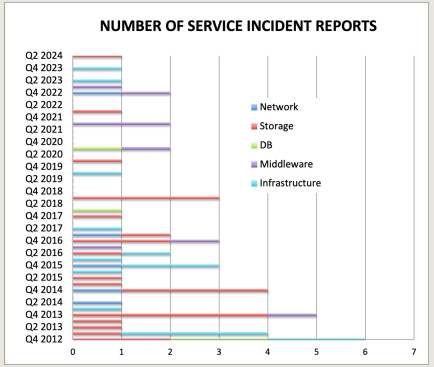

HL-LHC, just around the corner

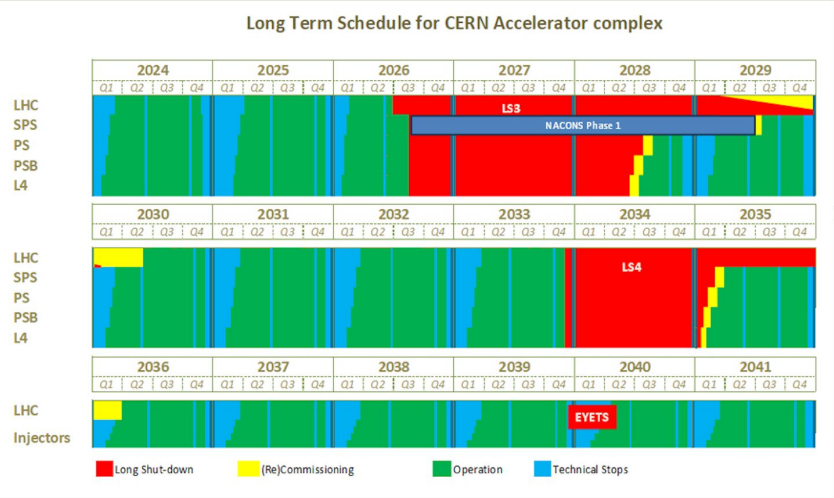

Updated LHC schedule (Oct 2024)

- The scale and the complexity of the infrastructure will grow massively.
- The available effort will probably won't follow this trend.
- We will continue to **evolve while running**.

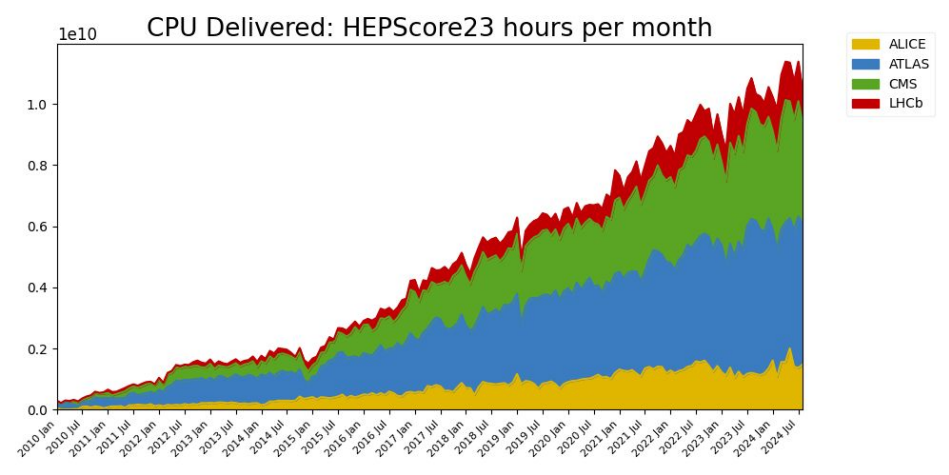

Plot from WLCG CRIC

What will operations be?

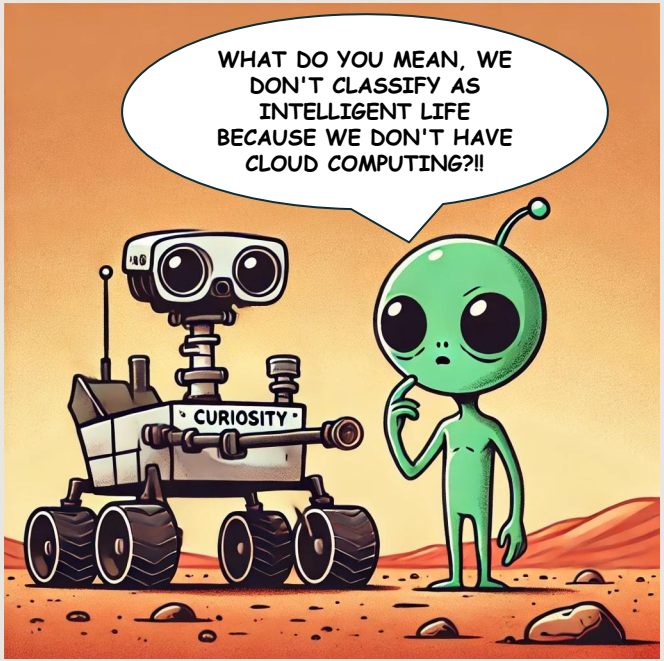

- What will we be operating in five years?
  - Increased Hybridization:
    - x86_64, ARM, GPUs, HPCs, and Clouds are now a reality—how do we integrate them effectively?
  - AI and Automation:
    - We would like to take advantage of anomaly detection and follow the recent AI related trends
  - Sustainability:
    - How do we balance performance with energy efficiency and sustainability at scale?

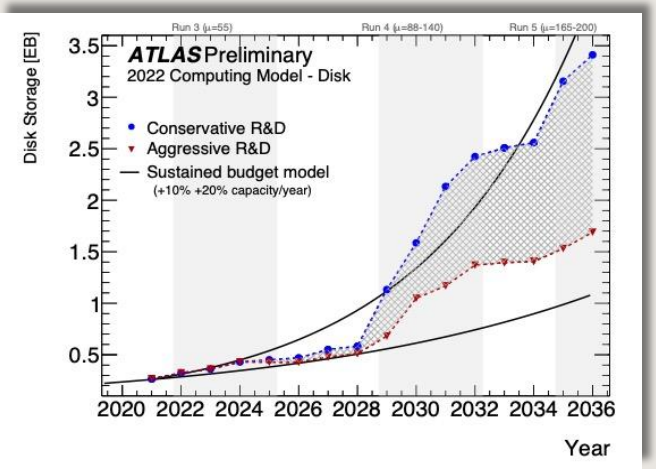

ATLAS Software and Computing HL-LHC Roadmap

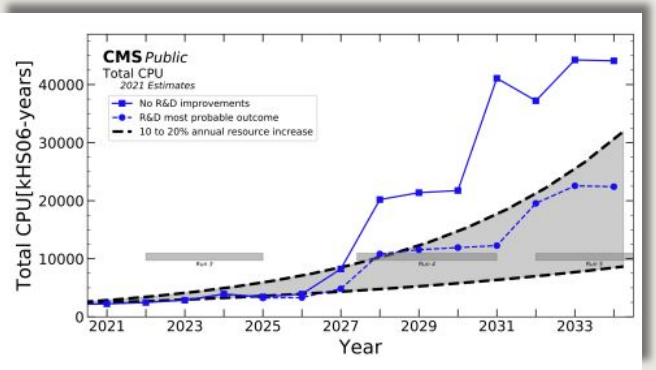

CMS Phase-2 Computing Model

What are the operational challenges

- It's paramount to leverage the **technology evolution**:
  - Incorporation of modern resources such as ARM, HPCs, GPUs and Cloud.
  - Migrate completely out of x509 and into tokens.
  - Support new types of sites

- Operating a larger and more complex infrastructure with **same (or even less) effort and personpower.**
  - WLCG is continuously growing, with increasing commitments and frequent new sites that we help onboard.
  - While not all sites have the same level of resources and expertise, everyone fits into our inclusive environment.

- **Common** tools and models:
  - In a such heterogeneous environment it's hard to reach consensus on tools and approaches.

- Identify where we need to **evolve** our current tools and policies and where we need a more radical **change** developing new modern tools.

Current operational tasks

- New questions and ideas are always emerging—we need a plan to address them!!
  As WLCG operations coordination, we work exactly on that.

- Operations must manage current challenges while planning for the future.
  We keep things running smoothly as we develop technologies to meet future demands:

  - Transitioning to tokens and adapting the operational model for this new era
  - Solidifying WLCG accounting practices and integrating the new benchmark
  - Evolving our topology description tools
  - Enhancing monitoring capabilities
  - Integrating new technologies
  - Running massive deployment and upgrade campaigns

CRIC

- **CRIC is the main topology system for WLCG and several experiments for many years now**

- CRIC is evolving in many directions.
  - Network topology data models
  - New pledge management tools
  - Supporting the service is always an ongoing task

- As CRIC has been invaluable to running various WLCG campaigns, we could enhance its functionalities to use it as proper Campaign management tool:
  - Support of the deployment and upgrade campaigns implies integration with new WLCG Helpdesk

- We are always open to exploring the possibility of deploying CRIC for other communities and projects.

Tokens

- Tokens are the new reality! While the infrastructure is ready and we successfully deployed IAM, migrating all relevant VOMS Admin functionality to it, the final operational model is still being developed.
- We will try to avoid fragmentation in operational models, from the very beginning:
  - How much security is necessary? Currently, proxy certificates last for multiple days. Any token with a shorter lifespan is inherently more secure..
  - What granularity we should use in terms of time and scopes?
    - There are different answers for different workflows!
    - 1h lifetime is secure but it would create a huge challenge for operation teams!
  - **What is the balance between operability, security and performance?**
- **WLCG is leading the charge**, bringing all relevant parties together to address the technical details for data transfers, job submissions and user experience.

Accounting

WLCG accounting has been running for two decades(!) and is now more accurate than ever.
However, we still need further improvements: more flexibility, anomaly detection, automation.

- Improving data quality:
  - Importing experiment data and comparing it with central WLCG accounting.
  - Integrating the new HS23 benchmark and using benchmarking tests to validate benchmarking factors. (more in ~40 minutes)

- We need to make sure we keep up with developments of new accounting tools:
  - AUDITOR: A new tool to manage accounting data with an extendable schema, allowing multiple customers (including APEL) to be served.

- We are already working on integrating new technologies:
  - Ongoing effort on benchmarking GPU resources.
  - Dynamic resources (Clouds, HPCs, …) which come with dynamic allocations need to be properly accounted.

Networking

- To support physics in 2030, networks must be significantly improved. **R&D** and **data challenges** will play a key role in achieving this.

- WLCG Ops support the Network team providing network topology and organizing deployment and upgrade campaigns
  - A LHCONE tagging campaign on 87 sites, powered by CRIC, aims to fully enable this functionality by spring 2025.

- Data challenges are organized to make sure that the HL-LHC targets are met and we take an active part in this activity:
  - Data Challenge 24 was very successful, the next one is already being planned.

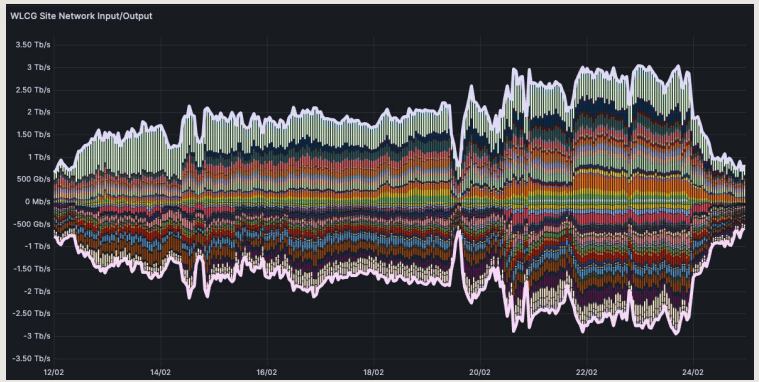

Network perspective of DC24

Monitoring

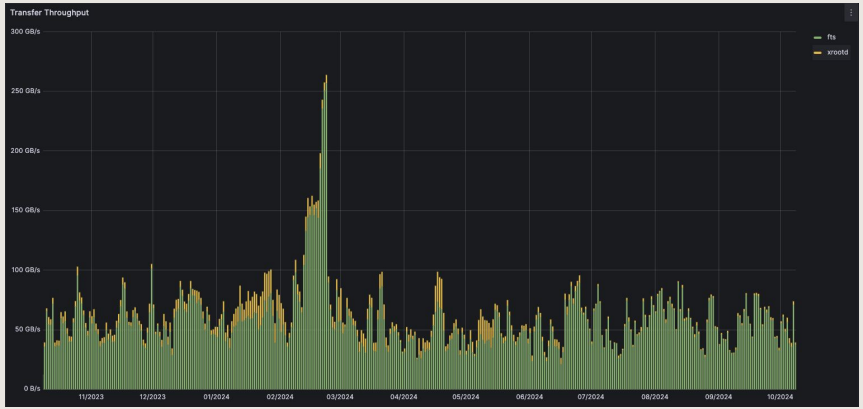

- Over the years there has been a lot of progress towards having a unified monitoring picture, not limited to each experiment's scope.
  - Job processing, data transfers and remote data accesses

- There is a significant ongoing effort to gather CPU metrics from the experiments and finally have a unified picture.
  - See this morning's presentation

- The WLCG monitoring task force has made significant strides toward establishing reliable XRootD monitoring (now in validation phase).
  - Yesterday's presentation

- Reliable and complete monitoring is a prerequisite for **problem detection** and **automation of the operational tasks** . We want to move forward from passive monitoring towards 'proactive' operations

OpInt

- We already attempted to introduce AI into our operations. With technology evolution and enhancements of WLCG monitoring it is the time to make real progress in this area.
  - AI is everywhere around us.
  - Emerging technologies like generative AI, are too promising to ignore.
  - An AI assistant would help with operations relieving precious brainpower for more tricky tasks.

- As our needs and resources grow, implementing anomaly detection is unavoidable.
  - In order to run a more sustainable infrastructure we need to spot and fix errors in a timely manner.

- It's crucial that WLCG operations remain proactive and ensure we take advantage of the technological shift.

New WLCG Help-Desk

- A new ticketing system has been developed by KIT to address WLCG's operational needs for the coming years.
- The new WLCG Helpdesk, based on the open-source platform Zammad, is set to replace the GGUS legacy system by the end of the year.
- Early adopters will be testing the system as more functionality will become available.
- In the end, GGUS will enter read-only mode, and all new tickets will be created in the Helpdesk.
    - Old tickets will remain available in the old system, read-only for at least 1 year.

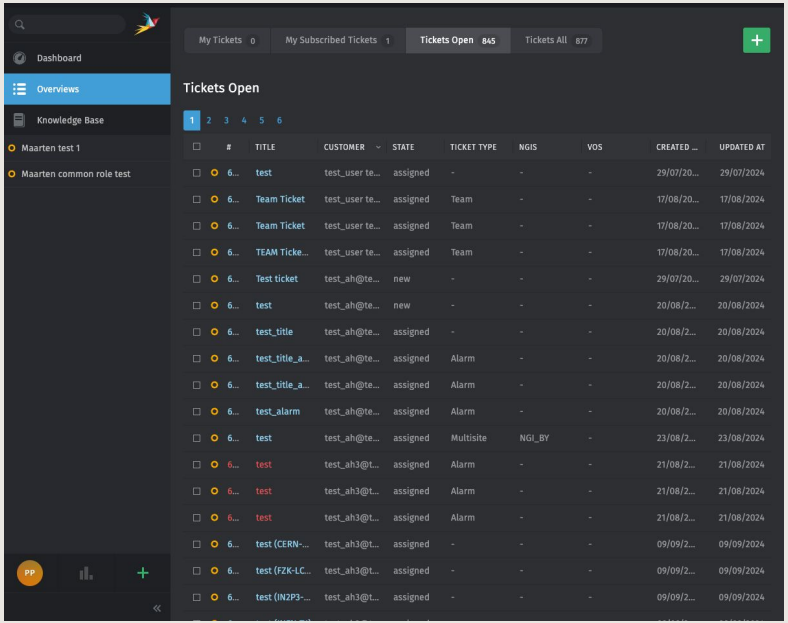

Screenshot taken from the Helpdesk dev instance

Conclusions

- **Adapting to Change**
  - WLCG operations are evolving to meet the growing demands of new technologies, such as tokens, HPCs, and hybrid cloud infrastructures.
- **Collaboration and Consistency**
  - Strong collaboration across all different teams remains key, while we continue to use common tools and approaches to ensure consistent operations.
- **Innovation is Essential**
  - From AI and anomaly detection to new network functionalities, ongoing R&D is crucial for addressing future challenges.
- **Looking Ahead**
  - The introduction of new tools like the WLCG HelpDesk and advancements in accounting, monitoring, and operational intelligence will position WLCG to succeed in HL-LHC era
- **People Make the Difference**
  - Ultimately, it's the teams and individuals behind WLCG operations that will drive our success!

Thank you