WLCG Strategy 2024-2027

D. Britton (Glasgow/GridPP), S. Campana (CERN),
T. Boccali (INFN)

Process

The input for this strategy

- Interviews with many of the stakeholders
- Recommendations/comments from the LHCC and RRB
- Trends from experiments and sites discussed at GDB/MB
- Input from ESPP, Snowmass, ICFA, ECFA (JENA)

Summarized in the attached document, executive summary presented here (with a lot of backup slides)

WLCG today

The Worldwide LHC Computing Grid (WLCG) is the global collaboration established to provide the computing infrastructure for the processing and analysis of data from the LHC experiments

- It extends across 165 sites in 42 countries and is underpinned by 66 MoUs signed by Funding Agencies across the world.

The success: meet 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.

- confidence of the community and funders to commit to the endeavour.

The future brings new challenges in terms of scale, technology, funding, sustainability, and the growth of other related communities

Objectives: WLCG needs to ...

- Support the core use cases (LHC exp.) whilst evolving the infrastructure for HL-LHC
- Evolve services to benefit from modern technologies with a careful evaluation of in-house vs commodity
- Optimize the person-power needed to run the infrastructure.
- Explore solutions that constrain the cost of our computing or attract new resources
- Increasingly embrace heterogeneous architectures and include diverse facilities
- Develop methods to reduce our energy footprint and/or reduce CO2 emission
- Collaborate with other communities representing sciences, software, and digital infrastructure, to share R&D and operational aspects where mutually beneficial

Strategy pillars: WLCG plans to …

[1] Ensure continuity and evolution through a programme of data challenges and improved monitoring. Our strategy is to re-risk the project by avoiding technical debt without introducing disruptive changes

[2] Modernize tools and services, benefiting from standards that are less HEP-specific where applicable. Our strategy is to profit where we can from well supported software that is shared by other groups.

[3] Minimise cost and energy footprint of computing and/or seek new opportunities. Our strategy is to build agility so that opportunities can be seized:

- use a wide range of resources including Grid, public or private clouds, HPC

- expand the architectures from the original x86 to include ARM, Power, GPUs, and other technologies that may become available.

Strategy: WLCG plans to …

[4] Increase collaboration with other communities, including formalizing the status of observers in WLCG and increasing synergies with the national and international initiatives developing large scale computing infrastructures. Our strategy is to benefit where we can from other relevant work

[5] Review the structure and governance of WLCG; set up bodies, boards, and task forces to address problems and issues. Our strategy is to focus and optimise our existing strengths

To implement the strategy

We Identified 11 areas of consideration for the future of WLCG

[OPS]	WLCG service operations
[TECH]	WLCG technical evolution
[FIN]	Financial sustainability of WLCG services and infrastructure
[INFRA]	Heterogeneous Grid infrastructure
[REL]	Relationship with other HEP and non-HEP communities
[REL]	Relationship between WLCG and OSG, EGI and other (inter)national grid initiatives
[SOFT]	(Offline) software and middleware development:
[GOV]	Engagement and Governance
[ENV]	Environmental sustainability: energy efficiency, green computing, carbon footprint
[RISK]	Identify and review Risks and Mitigations
[IMPACT]	Impact of WLCG on science and society, training and outreach

Each area was analyzed. Successes, considerations and strategic actions are detailed in the backup slides.

Example: WLCG technical evolution

Successes

- Demonstrated capability to introduce innovation while running operations
- Successful programs (e.g. Data Challenges) commissioning short-term improvements as part of the long-term goals

Considerations

- Natural caution about the introduction of disruptive innovation into a system that works needs to be balanced with risk of accumulating technical debt

Example: WLCG technical evolution

Strategic Actions

[TECH-1] Develop plans to broaden the scope of modernization of tools and services as an essential ingredient for the future technical sustainability of the WLCG infrastructure

[TECH-2] Develop/expand a structure for testing innovation, possibly based on the successful model of data challenges

[TECH-3] Set up a technical coordination body (e.g. a Technical Coordination Board) to oversee evolution and innovation of services

[TECH-4] Produce a Technical Design Report in time to be useful for the start of HL-LHC

WLCG in four years from now …

Make sure it works

- WLCG is ready to support the experiment needs for HL-LHC. The excellent track record operating the infrastructure is measurably maintained with a healthy turnaround of expertise

Make it sustainable - financially, technically, environmentally

- The WLCG infrastructure blends non-HEP and HEP specific technologies to provide a service for scientific computing. WLCG should engage with the open-source communities to maximise benefit

- The WLCG financial model evolves to facilitate long term planning of experiments and Funding Agencies.

WLCG in four years from now …

Make it sustainable (cont)

- The WLCG infrastructure evolves into a grid of heterogeneous systems (HEP institutes and labs, HPCs, clouds, ..). New facilities with different resource provisioning model and different hardware can be integrated efficiently and provide value to the LHC experiments
- More sciences see the value of the WLCG model and its services. WLCG collaborates with experiments of similar scale sharing services and infrastructure
- WLCG becomes a model for environmental sustainability: do more science impacting less on the environment
- WLCG is seen as an opportunity for science beyond HEP and society
- WLCG is capable of communicating all of the above...

Next Steps

At this Overview Board:

- we would like to have a first round of feedback
- we would like to agree about the 5 pillars of the strategy

In the next couple of months:

- Add what is missing in the various areas, e.g. more on networks, more on security, FAIR data …
- Share the strategy document with the WLCG collaboration and collect feedback

At the WLCG workshop in May 2024:

- Conclude the discussion and agree on the implementation roadmap

BACKUP

WLCG service operations

Successes

- Excellent record providing services for the LHC experiments computing needs
- Strong collaboration experiments/sites/middleware providers. Effective role of the Operations Coordination team
- Demonstrated capability in reacting to changing conditions

Considerations

- Reduction of the effort available for operations at sites and in the experiments

Strategic Actions

[OPS-1] investigate ways to further optimize the operational effort needed to run the infrastructure, services and workflows

WLCG technical evolution

Successes

- Demonstrated capability to introduce innovation while running operations
- Successful programs (e.g. Data Challenges) commissioning short-term improvements as part of the long-term goals

Considerations

- Resistance introducing disruptive innovation in a system that works. Risk to accumulate technical debt

WLCG technical evolution

Strategic Actions

[TECH-1] Develop plans to broaden the scope of modernization of tools and services as an essential ingredient for the future technical sustainability of the WLCG infrastructure

[TECH-2] Develop/expand a structure for testing innovation, possibly based on the successful model of data challenges

[TECH-3] Set up a technical coordination body (e.g. a Technical Coordination Board) to oversee evolution and innovation of services

[TECH-4] Produce a Technical Design Report in time to be useful for the start of HL-LHC

Financial sustainability of services and infrastructure

Successes

- Continuous support of the Funding Agencies, seeing value for money. Good dialog between experiments, FAs, sites, WLCG mgt
- Pledges ~ inline with experiment needs, deployed in time
- Considerable resources available beyond pledge (opportunistically). FAs responsive in case of unexpected/unprecedented needs

Considerations

- HEP as a science not expanding. Expect current “flat” budget to remain available but not to increase significantly.
- In several places, the same budget cover costs increasing over time (e.g. infrastructure and electricity)
- The hardware technology evolution is favorable, but markets are not (last 5 years)

Financial sustainability of services and infrastructure

Strategic Actions

[FIN-1] Identify the proper structure to monitor hardware and market trends, globally and at national level

[FIN-2] Develop and maintain multi-year resource planning (e.g. 3 to 5 years)

[FIN-3] Monitor pledges against experiments authorship levels

[FIN-4] Enable finer-grain pledges with respect of current annually flat pledges, without impacting the experiment core business

[FIN-5] More formally recognize the commitments (or list the contributions) of the Funding Agencies for middleware development and support

Heterogeneous Grid Infrastructure – WLCG Tiers

Successes

- The WLCG tiered structure proved to be simple and effective for pledges and definition of service levels
- Very reliable networks offering capabilities beyond the original hierarchical tiered model

Considerations

- The original tiered model considerably evolved
- Different capabilities provided by sites in the same tier. Experiments would benefit if capabilities could be more exposed.
- Discussing commitments in term of capabilities would however complicate the process

Strategic Actions

[INFRA-1] Collect, organize and expose information about site capabilities (compute, storage, network)

Heterogeneous Grid Infrastructure – Clouds

Successes

- Well established techniques integrating compute and storage resources with WLCG services through native cloud interfaces
- Many successful examples of cloud resources use in production for the LHC experiments .

Considerations

- The opportunities offered by cloud resources/interfaces are not pursued holistically in WLCG
- Some FAs are willing to provide resources as part of a cloud infrastructure
- Little understanding in WLCG of the TCO when using commercial cloud resources
- Non-HEP communities relying more on cloud resources than WLCG (and building experience)
- Moderate experience in WLCG with cloud storage. No cloud usage policy or risk management

Heterogeneous Grid Infrastructure – Clouds

Strategic Actions

[INFRA-2] Identify and document the technical solutions to integrate cloud resources with WLCG

[INFRA-3] Produce a blueprint for the integration of cloud resources and launch activity tracking the costs in different scenarios

[INFRA-4] Establish channels to follow progress of other communities exploiting clouds

[INFRA-5] Define a cloud storage usage policy

Heterogeneous Grid Infrastructure – HPCs

Successes

- Very successful use of HPC resources in WLCG, despite HEP use case needs little HPC capabilities
- Experiments and sites very adaptive to HPC policies and interfaces. Many solutions developed to adapt to these
- Measurable progress adapting experiment sw and common libraries to HPC hw architectures

Considerations

- HPC policies and interfaces still prevent efficient use in many cases
- Offline sw not supported on some HPC heterogeneous hardware
- Opportunistic use of HPCs today for a limited set of workflows. Little “multi year allocations”.
- Some (more) FAs expect to commit HPC resources in future as part of the pledges

Heterogeneous Grid Infrastructure – HPCs

Strategic Actions

[INFRA-6] Document existing solutions to integrate HPC centres and organize a knowledge base. When possible, propose one or more reference implementations via blueprint documents.

[INFRA-7] Construct a dialogue with the funding agencies and the relevant global bodies to drive the future allocation policies at HPC sites

[INFRA-8] Leverage the relationship with the HPC centers and with the Funding Agencies to influence the architecture of future HPCs also via interactions with the relevant international bodies

[INFRA-9] WLCG to monitor the projects on national and global scale designing collaboration strategies for HPCs, and foster the participation of its members

Heterogeneous Grid Infrastructure – Analysis Facilities

Successes

- Well funded projects to prototype tools and services for facilities specialized in future physics analysis.
- Interest of many WLCG sites, engaged in the discussion and prototyping

Considerations

- The future development of analysis facilities will likely impact WLCG sites (co-hosting or services, specialized hardware, ..)
- Role of specialized facilities for ML/AI only superficially discussed in WLCG
- How analysis will evolve for HL-LHC has still many open options

Strategic Actions

[INFRA-9] Explore how facilities supporting future analysis models could be hosted synergically with WLCG services

Heterogeneous Grid Infrastructure – general

Considerations

- The hardware landscape is now very heterogeneous: non X86 CPUs, GPUs. Need to evolve the offline sw to use the most beneficial and also to be able to anticipate trends. Need to prepare WLCG services for the use of heterogeneous architectures
- WLCG services and protocol are rather HEP specific. Benefit in moving towards more modern standards, while retaining the HEP core functionalities. Legacy is hard to decommission

Strategic Actions

[INFRA-10] Prepare for heterogeneous compute architectures: facilitate the development of the offline software and progress in the area of benchmarking and accounting

[INFRA-11] Establish a process for adopting modern, non-HEP specific standards where appropriate and decommission legacy. Include risk management for external dependencies

[INFRA-12] Accommodate the national plans consolidating facilities. Monitor design, implementation, deployment of the datalake and its components as solutions for the HL-LHC era. Engage the service managers in the transition and retain expertise

Relationship with other communities

Successes

- WLCG successfully positioned itself as a science-driven organization: it provides services for HEP, creates an ecosystem useful to many other sciences
- Engaged other HEP and non-HEP experiments w/o a strong governance model
- Many success stories: contribution of the “observers” in the evolution of WLCG; WLCG services adopted and customized by many experiments/sciences

Considerations

- The collaborating experiments/sciences have today smaller scale requirements than LHC, but some novel use cases.
- The process to become “observer” is not very formal and could benefit of more structure

Strategic Actions

[REL-1] WLCG to better formalize the process to become “partner” (a.k.a. “observer”): ratification moved to a higher body than the MB, clarification of the criteria, evaluation of the resource impact

Relationship with (inter) national grid initiatives

Successes

- Fruitful partnership and collaboration between WLCG, OSG and EGI in many areas: cybersecurity coordination, provisioning of federates services, support of middleware distributions, support of innovation, facilitate dialog with middleware providers
- Success of OSG and EGI supporting the long tail of science

Considerations

- Strategic interest of WLCG continuing the collaboration
- More coordination would be beneficial to build a common roadmap
- Risk of OSG and EGI loosing funding and reducing services is low but exists
- Other actors are appearing in the landscape of scientific computing e.g. EOSC, EuroHPC Joint Undertaking and WLCG has less direct relationship with them

Relationship with (inter) national grid initiatives

Strategic Actions

[REL-2] Take a leading role reviewing regularly the interaction and interdependency with OSG and EGI and strengthening coordination across the three parties

[REL-3] Ensure that operations are secured in case one of the partner initiative faces funding issues

[REL-4] Promote the benefits in being part of EGI (or OSG) while ensuring that ultimately countries can join the collaboration and provide resources even if not part of EGI (or OSG)

[REL-5] Build synergies with the national and international initiatives developing large scale computing infrastructures

(Offline) Software and middleware developement

Successes

- Middleware development sustained thanks to commitment and investment of HEP institutes and labs
- Excellent track record of offline software improvements (experiment specific and common libraries). This allowed a better use of the resources in the WLCG infrastructure
- Important role of the HSF making progress in several areas of common software

Considerations

- Many non-HEP-specific large open-source projects providing solutions that could and are being adopted by WLCG.
- Progress in offline software has a large impact on WLCG but. Responsibility is very distributed: experiments, HEP labs and institutes
- HSF loose structure optimal for community building, less for decision making

(Offline) Software and middleware developement

Strategic Actions

[SOFT-1] WLCG to continue encouraging the necessary software improvements to make the best use of the infrastructure, while not being in charge of such a development

[SOFT-2] The role and structure of the HSF needs to be re-discussed based on recent years of experience. WLCG to facilitate this process

[SOFT-3] WLCG to consider creating a mechanism for which PIs of current large projects or influential visionaries for new projects can discuss with the experiments and generally WLCG on the main strategic directions

Engagement and Governance

Successes

- The success of the WLCG Collaboration is based on building consensus between the parties: experiments, sites, middleware and services providers, national and international scientific computing initiatives

Considerations

- The WLCG governance structure needs to support the decision-making process based on consensus
- The governance needs to evolve together with the Collaboration and its level of maturity

Strategic Actions

[GOV-1] Review its governance to better support this strategy, while preserving the “principle of consensus”. Areas to consider include project leadership, OB, MB, GDB

Environmental sustainability

Successes

- The process to establish an environmental sustainability plan started in WLCG: long term estimate of energy needs, progress in using more energy efficient hardware
- Several examples of WLCG facility upgrades improving energy efficiency

Considerations

- Environmental sustainability is a strategic area of priority for many of the WLCG Funding Agencies. WLCG efforts ramped-up recently

Strategic Actions

[ENV-1] WLCG to agree metrics and provide a framework to collect information related to energy efficiency

[ENV-2] WLCG to enable the use of more energy-efficient hardware where possible

[ENV-3] WLCG to develop and promote a sustainability plan to improve energy efficiency, covering software, computing models, facilities, and hardware technology and lifecycle.

Risk Management

Successes

- WLCG produced a Risk Register in 2020 and reviewed it periodically. The process in in place

Considerations

- The structure of WLCG does not allow a formal risk management process due to the loose relationship between the partners
- Somehow CERN accepts the risk that computing is organized through WLCG rather than being uniquely a host lab responsibility
- There is value in having a Risk Register and keeping it up-to-date

Strategic Actions

[RISK-1] Continue reviewing the risk register at least once per year and present it to the Overview Board for endorsement

Science and sociaty, training, outreach

Successes

- The role of WLCG in scientific computing allowed various countries to develop a research and education cyberinfrastructure for the benefit of other sciences
- Effective national structures are in place, to support new centers. Experiments complement these structures
- Rich training program supported by the HSF and the experiments

Considerations

- Not always WLCG has been effective communicating the value it provides, beside supporting LHC computing
- Strategic and social interest in WLCG supporting countries with limited computing expertise. This support today is limited
- Career opportunities for software and computing experts in HEP remain a problem

Science and sociaty, training, outreach

Strategic Actions

[IMPACT-1] Establish a proper structure for communication and develop a communication strategy that is reviewed periodically by the collaboration. Form a communication team

[IMPACT-2] Define a strategy for supporting new countries willing to volunteer with limited resources

[IMPACT-3] Continue pursuing the objective to establish proper career opportunities for scientific software development, creating synergies with other sciences and leveraging existing organizations such as ECFA and national training programs

[IMPACT-4] Review the format of the WLCG Workshop, offering the opportunity for the community to present new ideas, for example through lightning talks and poster sessions

[IMPACT-5] Maintain a list of job opportunities at institutes willing to advertise the openings. This will hopefully allow us to retain expertise in the community and attract new ones