Operating the Worldwide LHC Computing Grid: current and future challenges

Josep Flix, Alessandra Forti, Maria Girone, Andrea Sciabà

on behalf of the WLCG Operations, Coordination and Commissioning Team

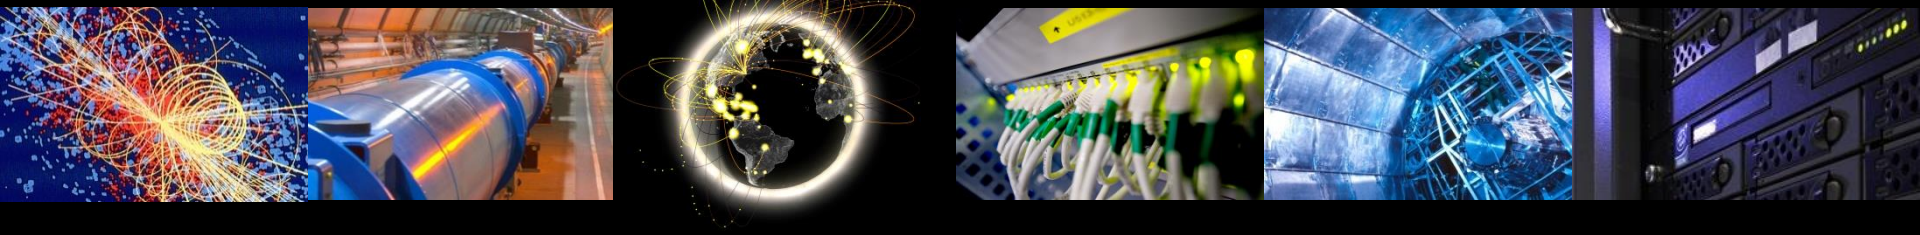

CHEP2013
October 14-18, 2013
Amsterdam

WLCG
Worldwide LHC Computing Grid

Outline

- Introduction
- Review of the WLCG operations
- The WLCG operations coordination working group
- The task forces
- Future challenges

WLCG

- The Worldwide LHC Computing Grid supports the distributed computing for the LHC experiments
  - Established in 2001, stable operations in 2010
  - Jointly operated with EGI, OSG and NorduGrid

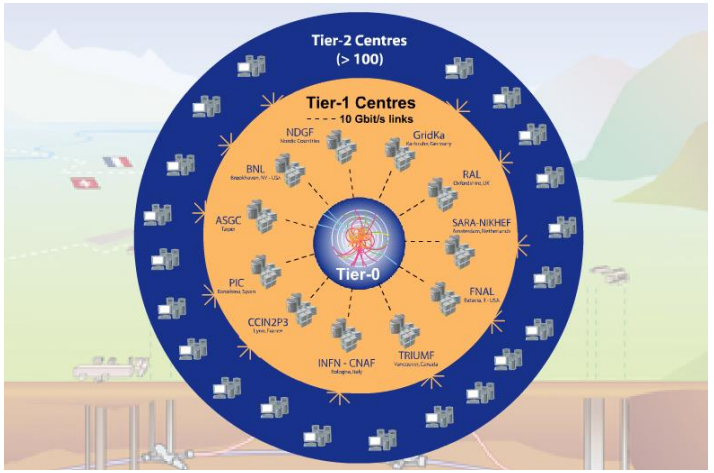

Pledged resources in 2013

	CPU (HS06)	Disk (PB)	Tape (PB)
1 Tier-0	360,000	29,000	71,000
12 Tier-1	680,000	75,000	110,000
143 Tier-2	990,000	89,000	
Total	**2,030,000**	**193,000**	**181,000**

Data distribution and processing

- Older model:
  - Tier-0 focuses on first pass reconstruction and long term archival
  - Tier-1 on reprocessing and archival
  - Tier-2 on simulation and analysis
- But rapidly evolving towards a full mesh model with little distinction between Tier-N roles
  - Mostly about hosting central services, tape systems and data custodiality

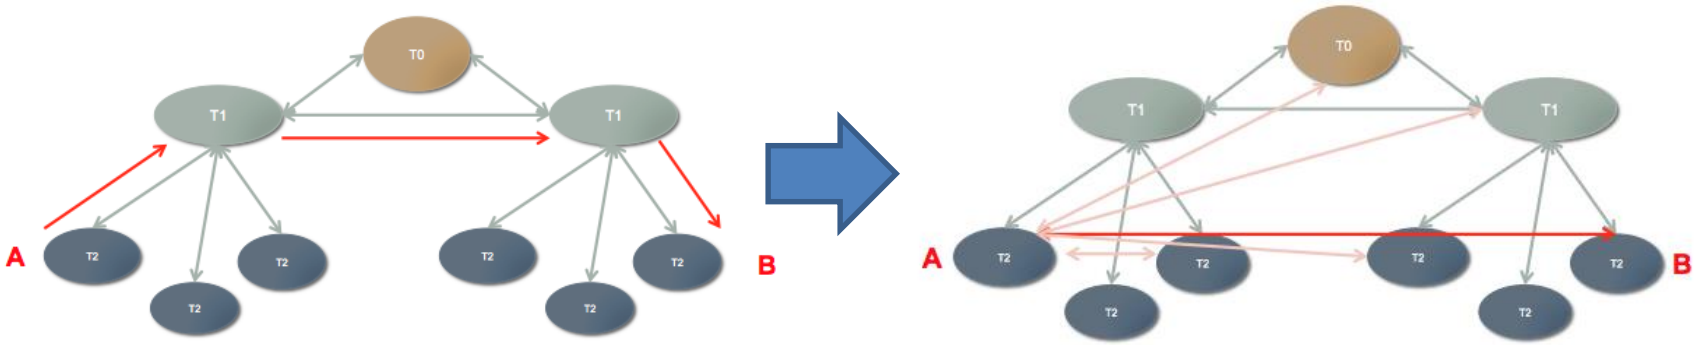

C. Grandi et al, CMS Computing Model Evolution, CHEP2013

WLCG operations

- Evolved from a decade long experience in prototypes, service and data challenges
  - Result of the effort of many individuals
  - Use also procedures and tools provided by the federated Grid projects in addition to its own
    - GGUS, GOCDB, OIM, EGI portals, etc.
  - Initial focus on delivering a stable service rather than sustainability
  - Very successful but at the expense of a high manpower cost
- In 2011 WLCG decided that it was the right time for an internal review
  - The Technical Evolution Groups
  - Not just for operations
- Main challenges were (still are)
  - Operate the system with less effort
  - Evolve towards new distributed computing models

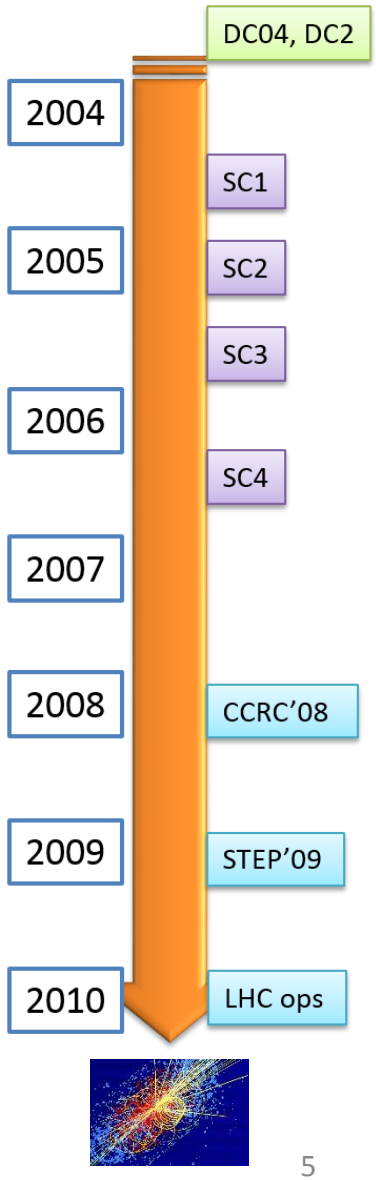

TEG operations review

- From October 2011 to April 2012 all aspects of WLCG operations were examined to describe the state of the art[(1)] and identify issues
  - Monitoring and site quality metrics
  - Operational procedures and support tools
  - Management of application software
  - Middleware validation and deployment
- The activity resulted in a series of recommendations, to be followed up by a new, dedicated coordination body

(1) Girone M and Templon J 2012 Final report of the operations & tools technical evolution group (TEG) https://espace.cern.ch/WLCG-document-repository/

Findings…

- WLCG operations reasonably smooth

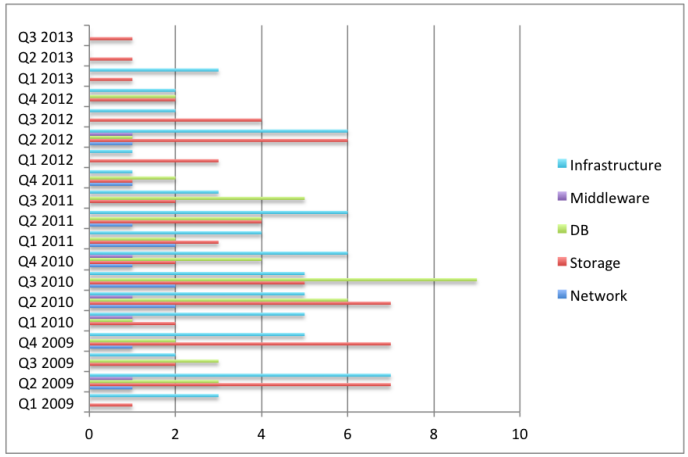

Service incidents per area

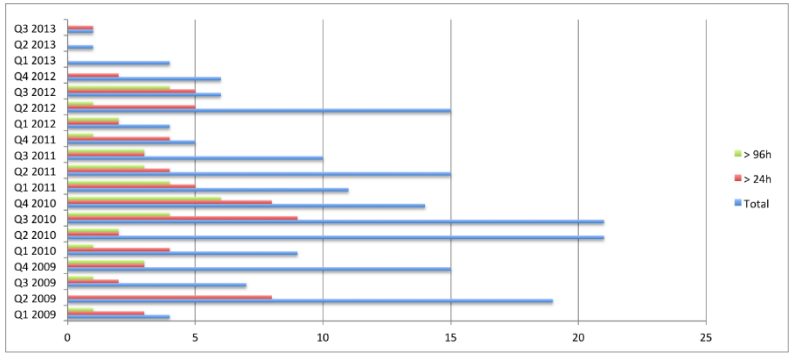

Service incidents per time to resolution

- The TEG highlighted some issues
  - Lack of effective communication with Tier-2's
  - Lack of a central WLCG operations team
  - Lack of a central body where experiments and sites can take operative decisions
  - Multiple (and sometimes fragile) experiment software installation and runtime configuration systems
  - Poor documentation and logging, difficult configuration for some services
  - Middleware validation with respect to experiments not centrally nor sufficiently organised

… and recommendations

- Several recommendations given
  - Here, only the most important are given
  - Recommendations concerning monitoring, network and information system went under the scope of other activities

Description
Establish a core team for coordinating WLCG operations
Expand the scope of existing meetings to fully involve Tier-2 sites
Adopt CVMFS to distribute experiment software (and middleware) at sites
Simplify the middleware stack and improve documentation and procedures
Improve middleware distribution and configuration mechanisms
Strengthen the participation of sites and experiments to the commissioning

The WLCG operations coordination working group

- Established in October 2012
- Acts as core operations and deployment coordination team
  - Manages operational issues, service deployment in synergy with EGI, OSG, NorduGrid
  - Discusses experiments plans and needs
  - Defines actions and work plans
  - Forms time-limited task forces on specific issues
  - Ensures communication among experiments, sites, projects
- All stakeholders are represented
  - LHC experiments, site regions, Tier-1's, Grid projects
  - Fortnightly meetings, quarterly planning meetings
  - Largely based on voluntary effort from the entire WLCG community

M. Girone, Operations Coordination Team, 11/07/2012, WLCG GDB

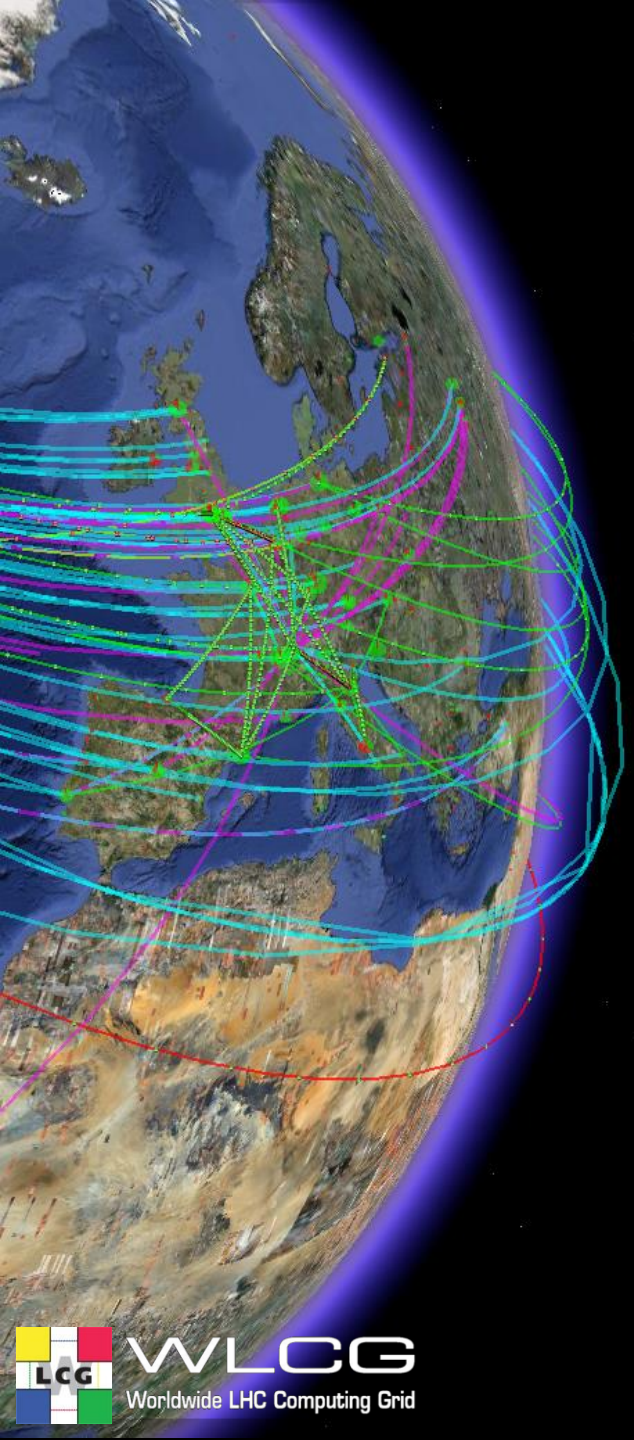

Task Forces review

- CVMFS
- perfSONAR
- SHA-2
- gLExec
- Tracking tools
- Squid monitoring
- FTS 3
- Xrootd
- SL6
- Machine/job features

WLCG
Worldwide LHC Computing Grid

CVMFS deployment

- A caching, HTTP based read-only filesystem
  - Initially developed by the CernVM project
  - Removes need for local software shared areas
  - Low load, highly scalable, little maintenance effort
  - Adopted by ATLAS in 2011, by LHCb and CMS in 2012, by ALICE this year
- The task force coordinated deployment on all WLCG sites
  - Central repository at CERN, replica at RAL
  - Almost all sites are now running it
  - 2/4 experiments support only CVMFS, the other two close to complete migration

perfSONAR deployment

- Evolution of WLCG from hierarchical to mesh model required all network paths to be monitored
  - Need for a WLCG-wide, experiment-independent end-to-end network monitoring to identify problems
  - perfSONAR-PS chosen to be deployed at all sites
- A task force to help sites
  - Installation procedures
  - Definition of tests and metrics
    - Latency, bandwidth, routing, packet loss
  - Centrally managed configurations of “full meshes” of tests

S. Campana et al, Deployment of a WLCG network monitoring infrastructure based on the perfSONAR-PS technology, CHEP2013

SHA-2 deployment

- The SHA-1 hash algorithm currently used to sign Grid certificates is increasingly risky due to known weaknesses
  - The agreed solution is to replace it with SHA-2
- All Grid projects and the International Grid Trust Federation (IGTF) agreed to start using SHA-2 by default from December 2013
- This requires all Grid middleware and experiment services to work with SHA-2
  - The task force coordinates the testing with experiments and projects
  - Almost everything has been verified to work

gLExec deployment

- Adoption of pilot job frameworks requires new mechanisms to allow traceability, user banning and "sandboxing"
  - gLExec is such agreed mechanism
  - Allows pilot jobs sent with "pilot proxies" to switch to user credentials at runtime
  - Needs to be supported by experiment frameworks

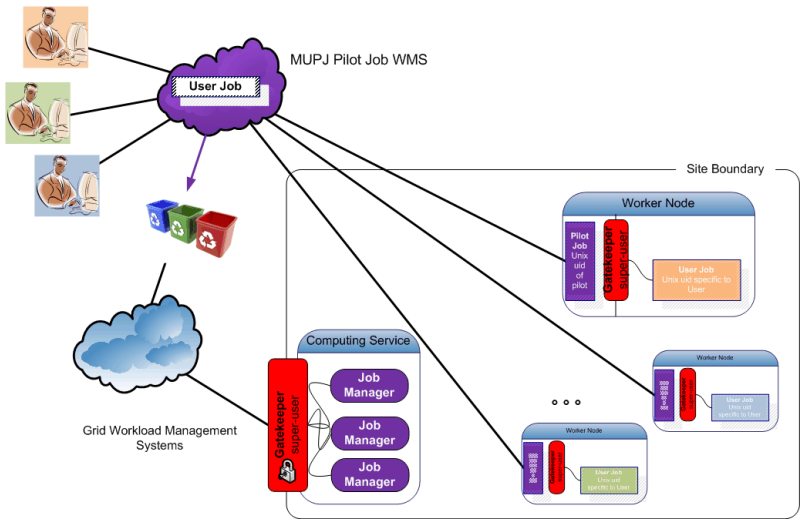

- The task force coordinates full scale deployment and adoption
  - To be finished by end 2013
    - Many sites will deploy it at the same time as SL6
  - Experiments committed to use it by 2014

Tracking tools evolution

- Several ticketing systems in use, with different (and sometime overlapping) scopes
  - GGUS, Savannah, Service Now, JIRA, etc.
- Crucial for effective communication, problem tracking, user support, statistics
- The task force includes experiments, sites, developers and coordinates
  - Commissioning/decommissioning of systems
  - Interfaces between different systems
  - Discussion of new features
- Specific examples
  - Savannah → JIRA migration
  - Savannah → GGUS migration for CMS

Squid monitoring

- Squid caches widely used in WLCG for efficient conditions data access via Frontier and for software distribution via CVMFS
- Needed to harmonise squid installations for different experiments
- The task force integrated Squid monitoring into WLCG common operations
  - Setup of central monitoring pages
  - Monitoring probes, functional tests
  - Registration procedures for GOCDB and OIM, etc.

FTS 3 deployment

- File Transfer Service drives the majority of data transfers in WLCG
- Version 3 builds on several years of experience
- Countless improvements
  - Simpler interface, better monitoring, multiple protocols, more DB backends, simpler configuration, …
  - Much simpler architecture and better scalability (1-2 instances could serve the whole WLCG)
- The task force does integration, testing, validation using pilot services
  - High volume tests already running for ATLAS, CMS and LHCb
  - Goal is delivery of a production service

FTS 3 production transfers

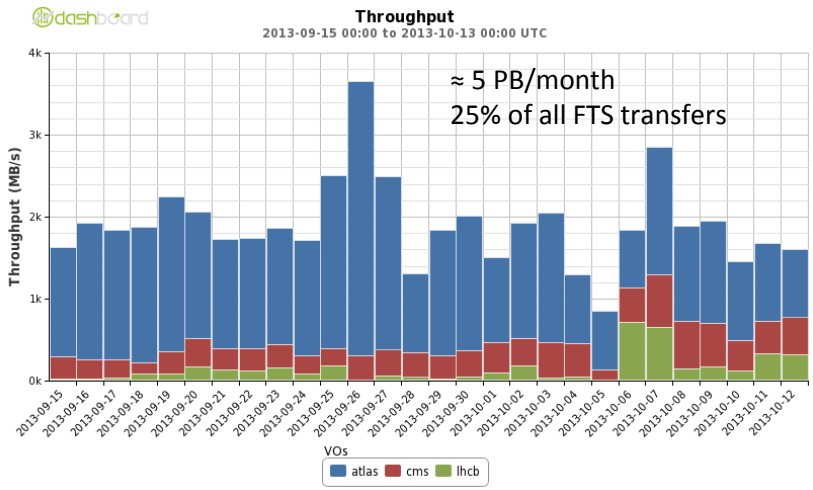

M. Salichos et al, FTS3 – Robust, simplified and high-performance data movement service for WLCG, CHEP2013

Xrootd deployment

- Main goals
  - provide support to the xrootd deployment - driven internally by AAA and FAX - and identify common needs
  - collect the monitoring requirements and liaise the monitoring effort
- Status of deployment strictly connected to the two experiments activities of AAA and FAX

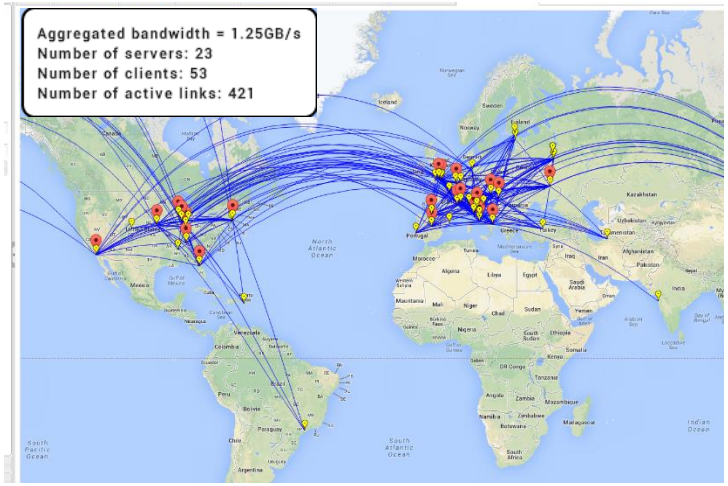

AAA (CMS)

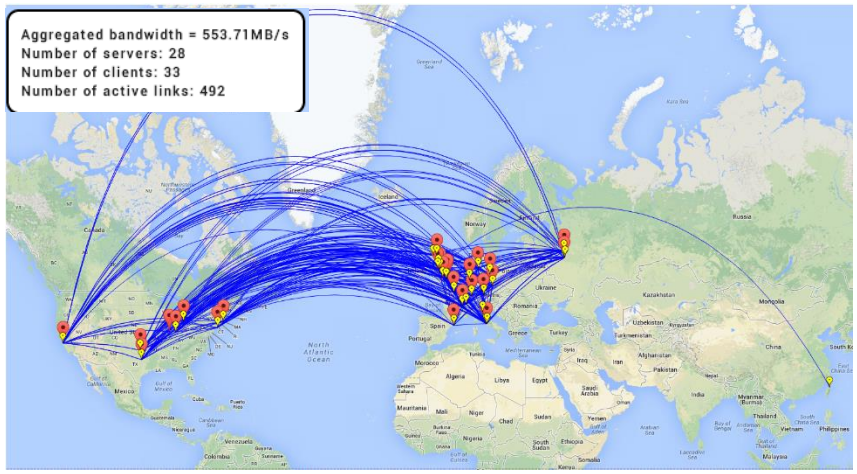

FAX (ATLAS)

Scientific Linux 6 migration

- Coordinate upgrade to SL6 by minimizing amount of unavailable resources
  - Set a timeline agreed by experiments and set procedures for sites to follow
    - "big bang" vs. rolling upgrades, mixed SL5/SL6 queues, etc.
    - HEP_Oslibs metapackage for experiment software dependencies
    - WLCG repository
  - Make sure experiments are ready
  - Follow and coordinate sites, track status
- Current status:
  - (proto-)Tier-0+1s: 10/16 (now), ≈ 15/16 (end October)
  - Tier-2's: 62% (now), ≈ 77% (end October)
  - In average, one site every 1.3 days since June without major disruptions to operations
  - Almost all remaining sites should migrate by November

Machine/job features

- Much needed functionality for running jobs
- Static information
  - HS06 rating, no. cores, scratch space, …
- Dynamic information
  - VM shutdown time, time limits, memory limits, allocated cores…
- The task force is charged with coordinating implementations and deployment in WLCG
  - Including monitoring of the correctness of the information
- First implementations available and tested on physical and virtual hosts

Future challenges (1/2)

- Medium term (during Long Shutdown 1)
  - Monitoring consolidation
    - By dedicated project
  - FTS 2 decommissioning
  - Expanding storage federations
  - Dynamic data placement

Future challenges (2/2)

- Long term (Beyond Long Shutdown 1)
  - Sustainable middleware validation and distribution
  - Full IPv6 compliance
  - New hardware technologies and architectures
  - Full integration of cloud resources

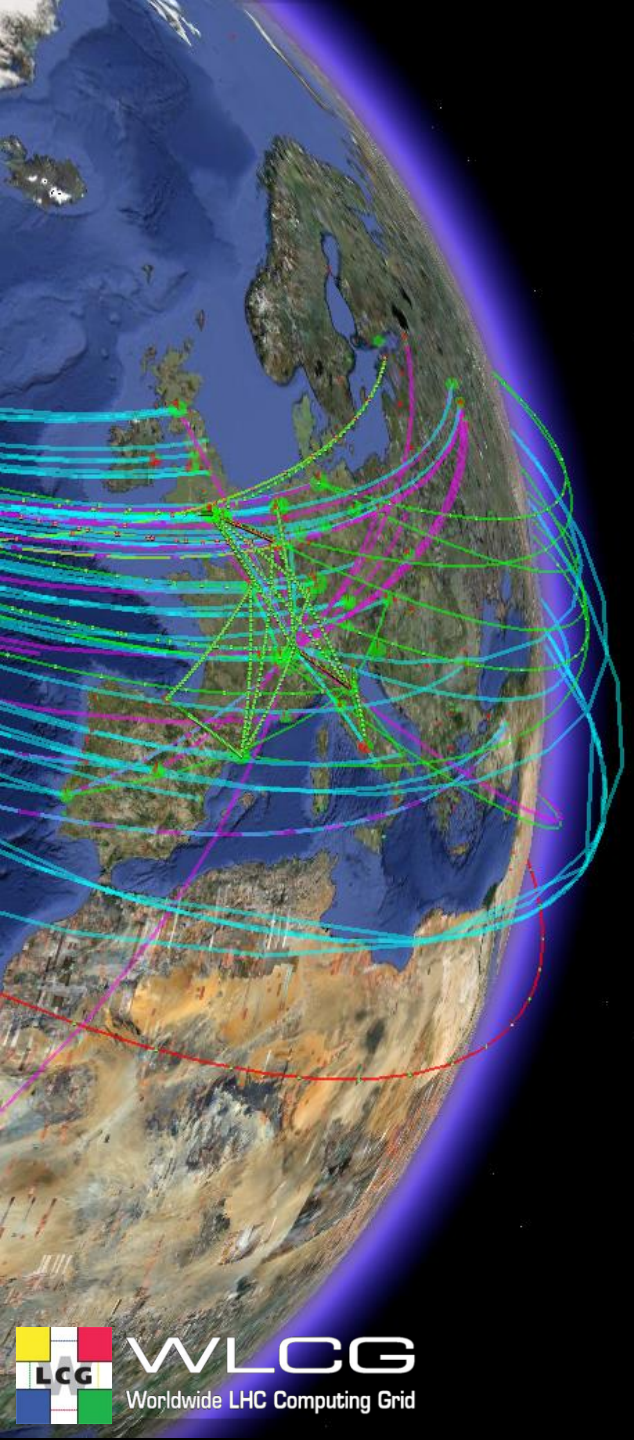

Conclusions

- WLCG Operations successfully work as a collaborative effort
- Significant progress on several areas
- Will need to follow the evolution of the infrastructure from an operational perspective in the coming years

Acknowledgements

- Thanks to the many people that contributed to this effort
  - C. Aiftimiei, M. Altunay, Y. Andreeva, J. Artieda, M. Barroso, D. Barberis, A. Beche, T. Bell, L. Betev, D. Benjamin, J. Blomer, B. Blumenfeld, B. Bockelman, D. Bonacorsi, P. Buncic, D. Bradley, S. Campana, A. Cavalli, T. Cartwright, M. Cattaneo, P. Charpentier, M. Cinquilli, J. Closier, J. Coles, I. Collier, B. Couturier, A. Da Silva, J. DeStefano, L. Dell'Agnello, A. De Salvo, A. Dewhurst, S. DeWitt, V. Diez, M. Dimou, A. Di Girolamo, M. Doidge, G. Donvito, A. Doria, H. Dres, D. Dykstra, M. Ernst, S. Fayer, G. Eulisse, T. Ferrari, I. Fisk, J. Flix, I. Gable, R. Gardner, C. Grandi, A. Forti, D. Giordano, M. Girone, I. Gonzalez, G. Grein, P. Gronbech, O. Gutsche, T. Hartmann, B. Hegner, N. Hoimyr, C. Hollowell, B. Holzman, J. Hover, D. Hufnagel, H. Ito, J. Iven, M. Jouvin, O. Keeble, B. Kersevan, A. Klimentov, A. Lahiff, F. Lee, L. Linares, S. Liu, R. Lopez, S. McKee, M. Litmaath, N. Magini, A. McNab, M. Martinez Pedreia, H. Meinhard, V. Mendez, J. Molina, E. Obreshkov, R. Quick, Y. Patois, A. Petzold, J. Pina, D. Rand, D. Qing, S. Roiser, M. Salichos, J. Sanchez, R. Santana, M. Schulz, U. Schwickerath, A. Sciabà, I. Sfiligoi, M. Sgaravatto, D. Da Silva Gomez, G. Smith, P. Solagna, J.E. Sundermann, A. Tiradani, S. Traylen, I. Ueda, A. Valassi, R. Walker, Y. Wei, T. Wenaus, C. Wissing, M. Zielinski, and many others
  - All site contacts and system administrators, and CERN IT