WLCG Network WG (WNWG) - Mandate, Scope, Membership …

Michael Ernst
Brookhaven National Laboratory

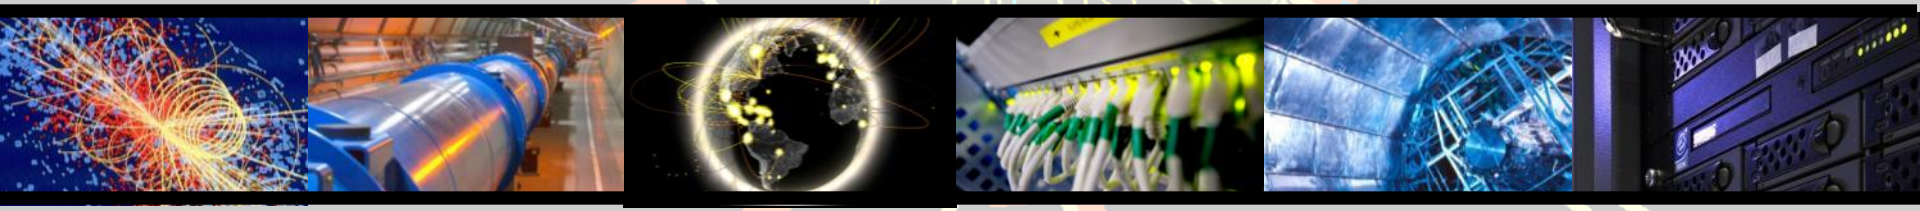

From the April C-RRB Report

"The CRSG encourages close collaboration of the different centres with the experiments to continue the implementation of intelligent storage management policies to allow efficient and cost-effective access to data. In particular the implications on network bandwidth for best-use of resources should be considered. We consider this issue very relevant for the operation of the LHC experiments after 2014."

http://indico.cern.ch/getFile.py/access?contribId=16&sessionId=2&resId=0&materialId=slides&confId=174803

WLCG Facilities

- Today in WLCG our primary concern regarding resources is CPU and storage capacities
  - As of today the network is essentially outside of WLCG's scope – even the OPN is not part of the capacity planning process
  - All LHC experiments are increasingly relying on networks to provide connectivity between distributed facility elements across all levels
    - E.g. ATLAS is ramping up the usage of T2 storage for primary datasets
  - Wide Area Network infrastructure is complex
    - Besides the (mostly) predictable service of the OPN the vast majority of traffic (i.e. for ATLAS and CMS) has been using the GPN
    - With LHCONE there is a 3rd piece of network infrastructure
      - Meant to serve T1 ⇔ T2/T3, T2 ⇔ T2, T2 ⇔ T3 needs
      - Evolution & Deployment largely driven by network service providers
        - » LHC experiments not properly represented at planning meetings
  - ➢ With vanishing T1/T2 differences Convergence of LHCOPN and LHCONE could (should?) be considered
    - ➢ Should conduct a feasibility study

Issues

- Global Networking between sites has evolved into a service that is vital to the success of LHC Computing
  - Traditionally, when looking at networking at all, we were focusing on the OPN but we/WLCG should make planning & operations aspects of networking in general part of our agenda, as we do for CPU & storage
  - Network providers are guessing how experiments are using and making assumptions as to how they would be using the service in the future
    - At LHCONE planning meetings I am hearing, frequently, requests for input but there is no authoritative/coherent information provided by the WLCG community
    - LHC Experiments realize that recent developments in networking go well beyond packet exchange through a “passive” pipe
      - Network service providers looking to engage us into traffic engineering
      - (Dynamic) Circuit technology (p-2-p service) w/ additional b/w for high priority data replication requires stakeholders’ active participation
      - Evolving service technologies are engaging applications in actively managing data flows (OpenFlow)
        - » Very appropriate response to realizing the importance of the network in the global facility architecture
        - » Internet2, Google, Industry, others are working on the standard/specs

Motivation

- Rather than the Communities, e.g. ATLAS, which are viewing the currently evolving deployment of “LHCONE” production-grade, the Network Providers understand LHCONE as a first phase of a new network infrastructure to improve global T1 <=> T2, T2 <=> T2, T3 <=> T2, … communication
  - The service was not stable in the beginning (now much improved) and network providers reserve the right to make changes, and mistakes …
  - Sites are joining LHCONE whenever they want
  - ➢ This has/could cause(d) disruptions to the experiments’ workflows
- Having access to network performance data (bandwidth, RTT) allows the application to optimize usage of CPU & Storage provided by distributed facilities.
- Immediate goal is to optimize the usage of networks for centrally managed transfers (DDM/FTS for ATLAS) and task brokerage. If storage federations become popular, random/chaotic access across countries & continents will be even more challenging.

WLCG Network Working Group (WNWG): Scope

- Focusing on network infrastructure and network services in support of distributed applications of the LHC Experiments
  - Aims at finding common solutions across the experiments
  - Includes definition & implementation of the interface between LHC community and network service providers
    - Meant to address high(er) level issues (vs. daily operational ones)
    - Not to interfere with procedures used by site/local network providers
    - Purpose is to address functionality & performance issues all along the vertical stack (from layer 2 to application)
    - Voice LHC requirements in consolidated & coherent fashion
  - Technical representatives for planning process and liaison on behalf of stakeholders for communication with providers and developers
    - Includes definition, setup and tracking of R&D activities
    - ➢ Foundation for informed choices
    - ➢ Develop a strategic plan

Proposed Composition of the WNWG

- Experts from the following areas
  - Applications (includes analysis models, workflows) and Distributed Data Management
    - 1-2 people per experiment
  - Facilities and Operations Coordination
    - Operations interface between WLCG entities and experiment services and network service providers
    - Provider representation could be delegate(s) from LHCOPN/LHCONE WG
    - 3-4 people representing T1/2/3 and the WLCG Ops Coordination Team
  - Network Services and Technologies
    - Up to 3 people representing Network Providers and emerging technologies
      - Provider representation could be delegate(s) from LHCOPN/LHCONE WG
      - Could be the same individuals as above
    - Representation from R&D community/organizations

WNWG Charge (1/2)

- Collect, document and update requirements.
- Propose and develop technical approaches and technical specifications to meet the functional, performance and operational requirements of the LHC Experiments
  - Common solutions as guiding principle
- Provide recommendations to WLCG management as to how the requirements w.r.t. functionality, commissioning and operations of network services could be met by network providers and sites
- Estimate amount of effort to implement them and ask WLCG Management to provide resources if they are adopted
- Establish a timeline of critical milestones and deliverables and track their implementation

WNWG Charge (2/2)

- Communicate and collaborate with network service and technology providers and sites to provide:
  - Resolution to problems (daily operations issues are handled by the WLCG Operations Coordination Team (OCT))
    - » E.g. the deployment of perfSONAR monitoring at sites will be handled by the OCT
  - Quarterly reviews of the performance of network services relative to the original requirements and the actual evolving needs
    - » Actual Performance as measured to be compared against Metrics established by Experiments
  - Quarterly evaluations of the adequacy of the system and the performance of the providers
  - Recommendations for adjustments, improvements and modifications to be undertaken in the future

WNWG and R&D

- The WNWG shall actively be looking at results of R&D to identify, prioritize, and to promote initiatives that advance the usage and usability of network services as part of the WLCG computing infrastructure. Specific activities include
  - Opportunities & Risks
  - Adoption of new network technologies
  - Establish WGs for R&D
    - WGs will consult broadly and will work closely with experiments, network service & technology providers and sites

Network Technology has evolved – Great potential to improve LHC application performance

We need flexible definitions of a flow

- Unicast, multicast, waypoints, load-balancing
- Different aggregations

We need direct control over flows

- Flow as an entity we program: To route, to make private, to move, …

Exploit the benefits of packet switching

- It works and is universally deployed
- It's efficient (when kept simple)

Software Defined Networking (SDN) – This looks like a promising Technology

- SDN Paradigm - Network control by applications; provides an API to externally define network functionality
  - Enabler for applications to fully exploit available network resources
  - OpenFlow, a SDN implementation, widely adopted by Industry (Network Equipment Manufacturers, Google, Cloud Computing)

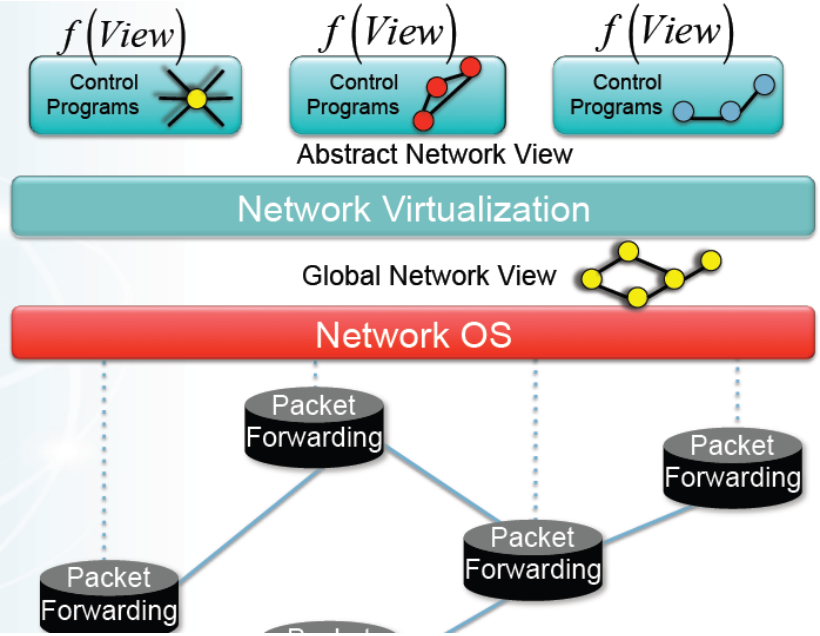

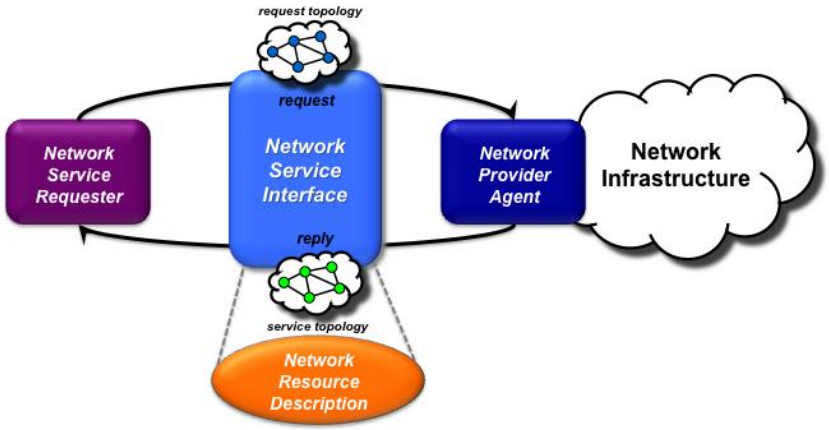

Network Services Interface Framework

Open Networking Foundation: https://www.opennetworking.org/index.php

Example: ATLAS Workload Management
(works likewise w/ other WLM Systems)

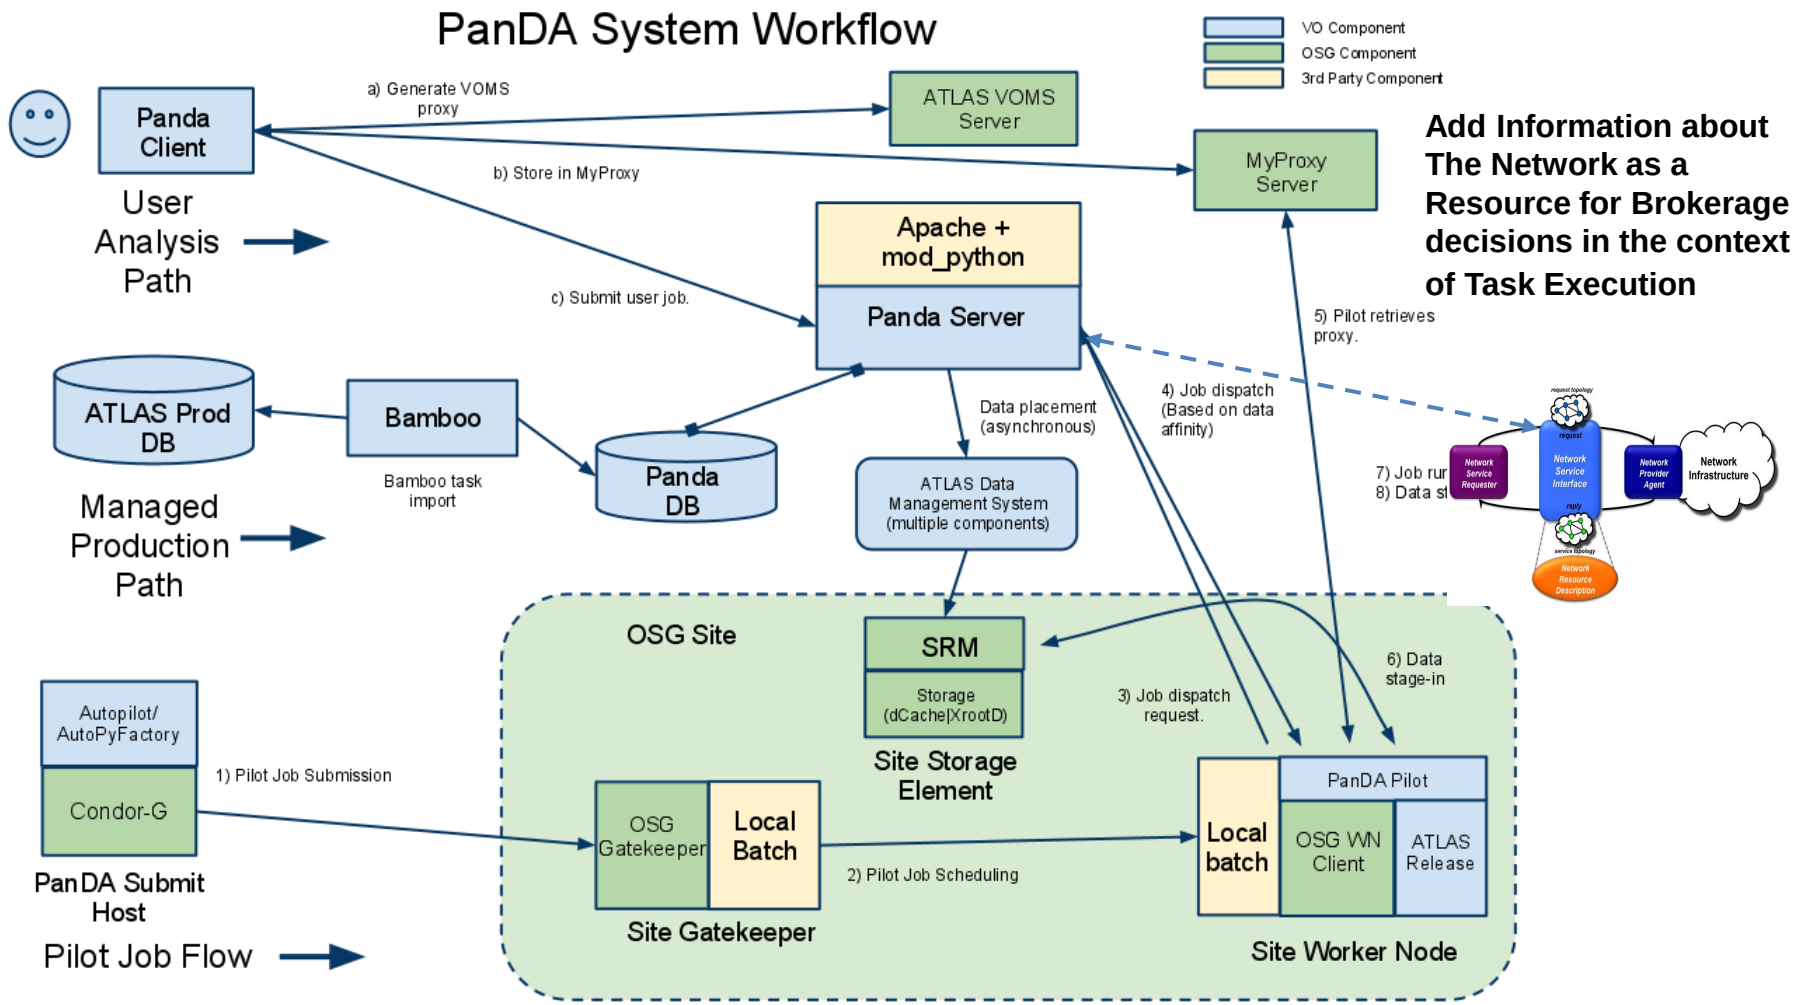

Example: ATLAS Distributed Data Management
(works likewise w/ other DM Systems)

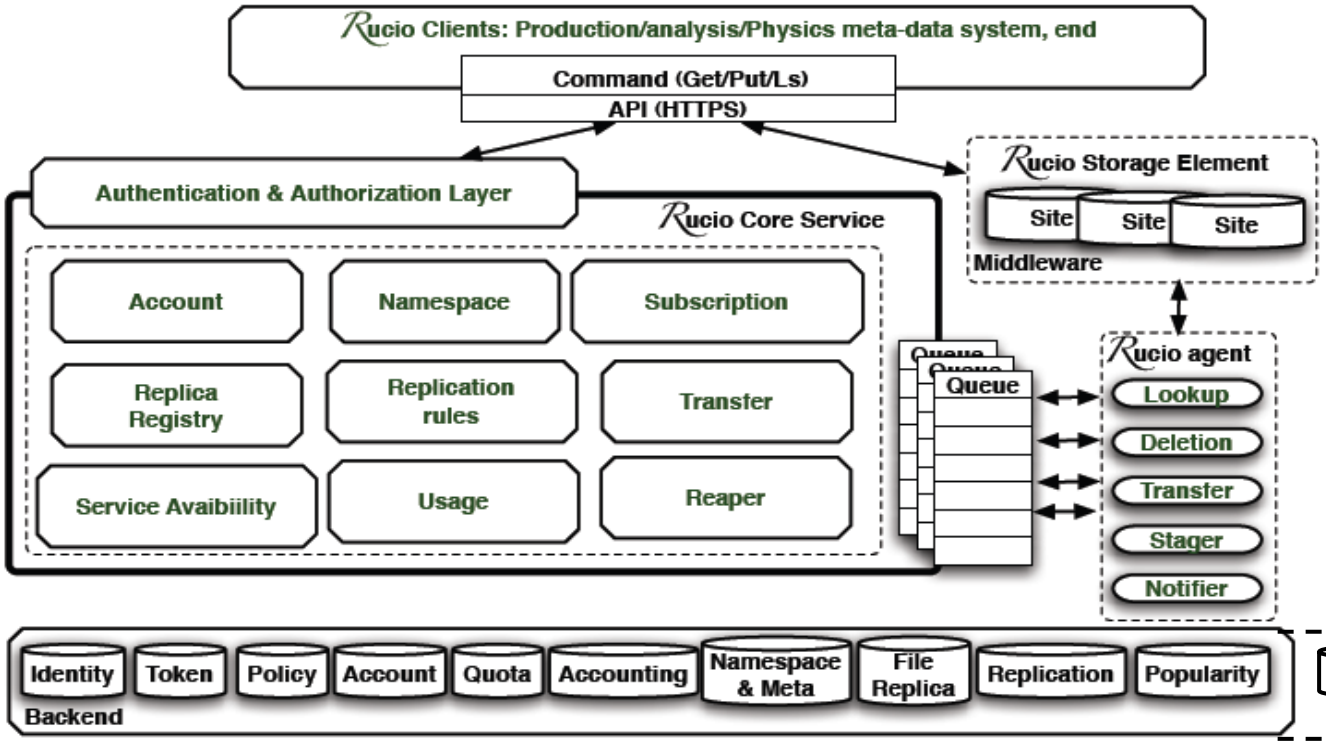

Add Information about the Network as a Resource for Scheduling Decisions in the context of Data Placement and large scale Data Replication activities

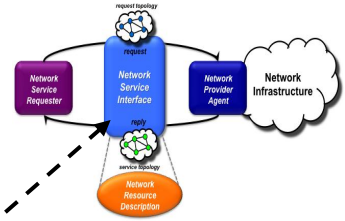

SDN – A Reality Check

- Research in using OpenFlow for SDN has a great potential
  - Different directions for solutions pursued, no stable definitions yet
  - How can the LHC Community participate in the development or should we better wait?
- Network abstractions is a key research topic
  - No stable interface definition at this point
  - How long will it take to create such an interface?
- A programming language for networking primitives is needed
  - Is there a place for LHC application developers to participate prior to the arrival of a (quasi) stable version?
    - Feasible to just rely on Network Engineers?

Virtualized Network Infrastructure for Domain-Specific Applications (i.e. LHC Computing)

- Domain-specific networks as an overlay on top of shared infrastructure
  - Predictable Performance characteristics

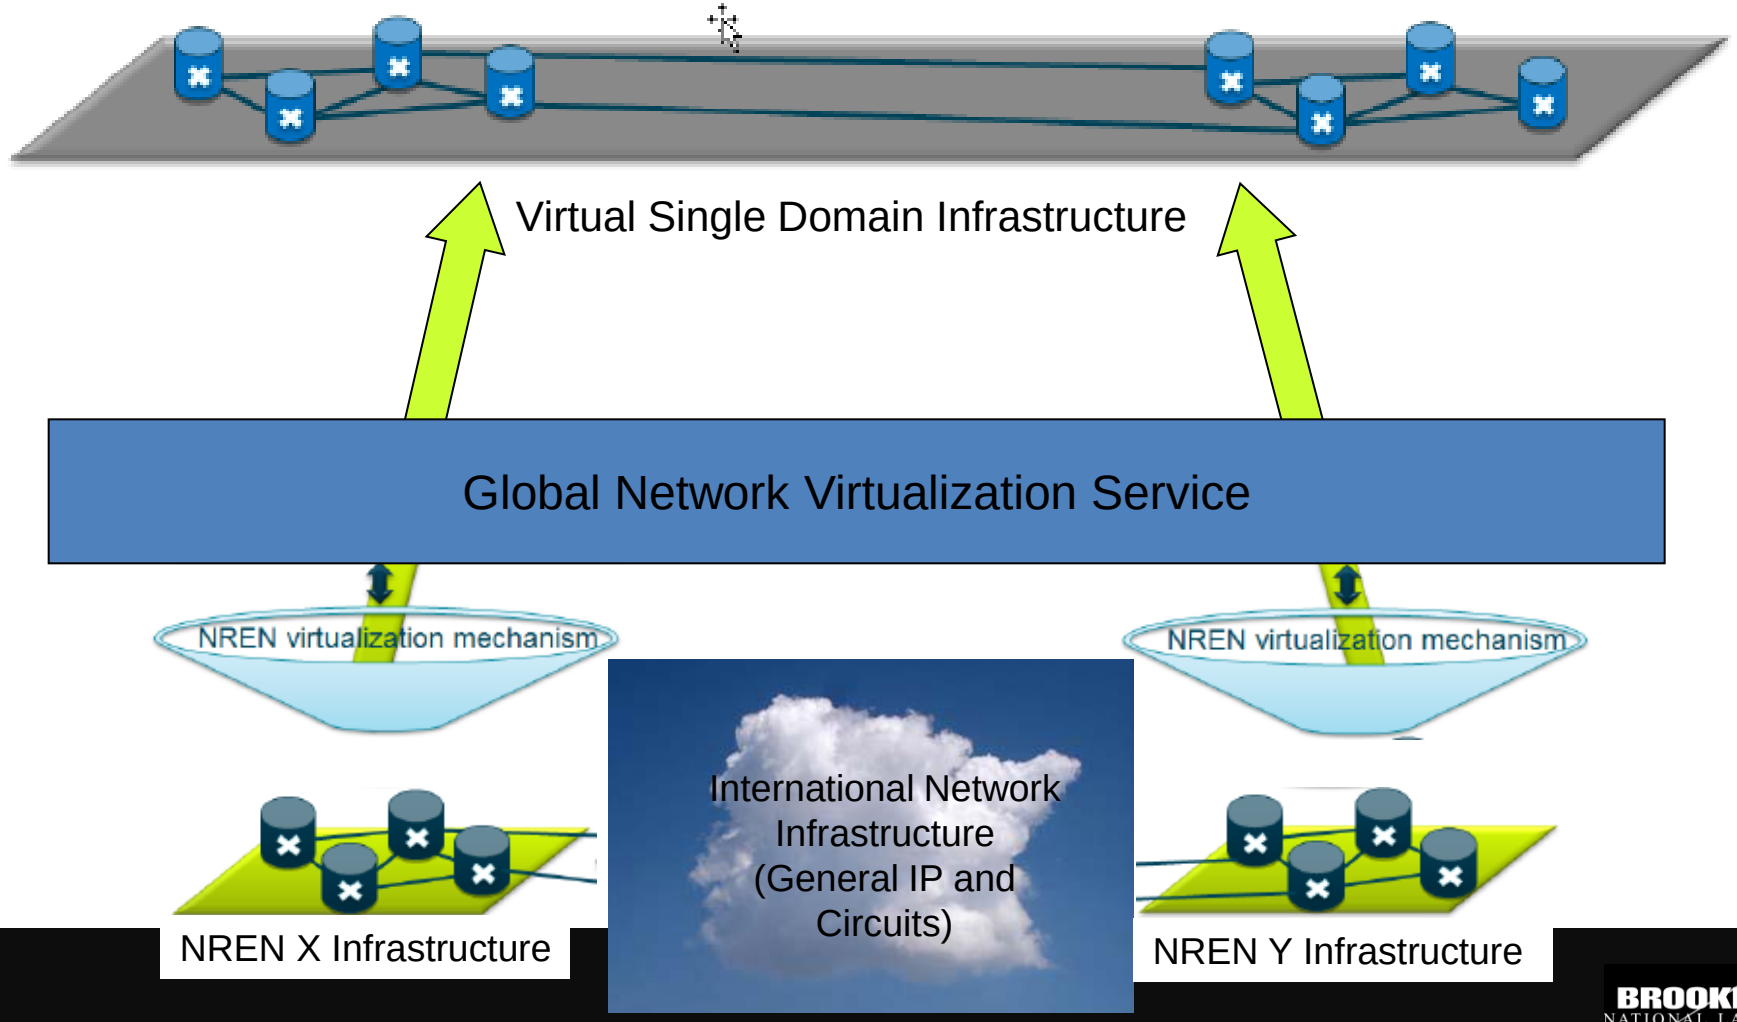

WNWG Relations

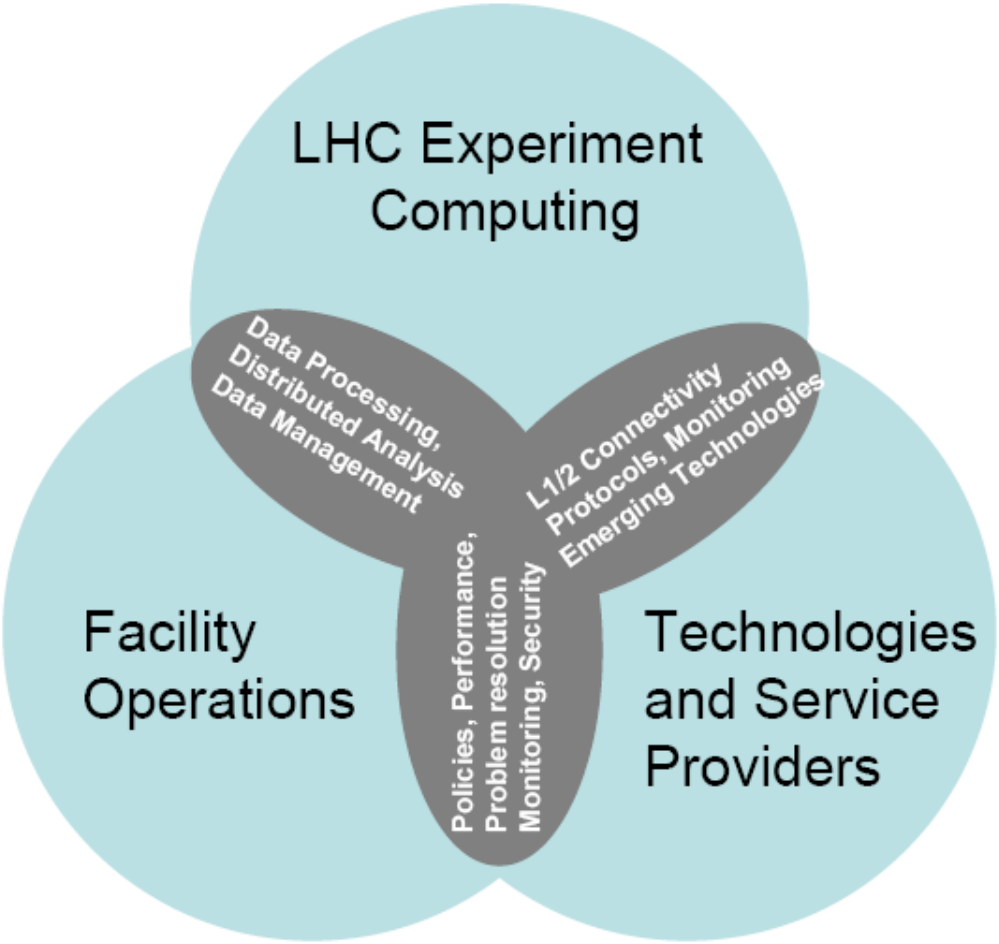

Communication & Documentation

- Email/virtual meetings and Twiki
- Active participation in meetings like this and network-related workshops
- Quarterly (or whenever necessary) reports to WLCG MB
  - And, if appropriate, GDB and C-RRB (e.g. in case of implication to funding)

A possible path forward (1)

- We are close to a 2-year shutdown of the LHC machine
  - As we have developed a better understanding of our computing needs, and how we can best benefit from the middleware that supports flexible & nimble operations, we want to use the time to review the situation and participate in shaping the environment, including the network as a fully integrated facility component
  - By and large we understand we have your full attention
    - Define & impl. Operational Interface between Users and Providers
      - E.g. the Working Group on LHCONE Operations is inactive, and we need something that covers OPN, LHCONE and the GPN anyway
    - Besides providing production services most of the organizations you represent participate in pilot and research activities
      - with the motivation of providing new types of services to the user communities, and it's our role to contribute applications and the associated experts/expertise

A possible path forward (2)

- We aim at making the network with its operational and development components part of the global resource planning and operations process that integrates worldwide distributed LHC computing facilities and the associated network requirements
  - The LHCOPN, LHCONE (as well as the GPN infrastructure we depend on) is part of the scope. In the larger context the network's role and implementation will be reviewed and integrated into the planning process.

- Instead of standing aside WLCG, with the LHC applications, will participate in research activities conducted by network providers, which should be ambitious in its approaches on the timeframe of the next n years. It should aim to bring an alternative to "traditional" large-scale IP infrastructures or demonstrate that it is not feasible, economic or advantageous to do so.

- WLCG will engage as an early adopter user community in a managed way with you, the network service providers, to demonstrate that new infrastructures can provide production services. This should bring new capabilities to us and not impact the current production services. This is a coordination task and should be run as a joint effort.

Status and Next Steps

1) WLCG MB to decide whether to instantiate WG - DONE
2) Refine the proposed Scope, Charge, Composition etc, if needed – Your Input is most welcome
3) Find & appoint the WG Chair - DONE
4) Find and appoint WG Members
   - Experiments to nominate WG Members - DONE
     - Have received nominations from ALICE, ATLAS, CMS and LHCb
     - Once done WG could start work on requirements – In preparation
   - Solicit nomination for Site Representative(s) – In preparation
   - Solicit nomination for Representatives from Network Service and Technology Provider community – proceed according to 5)
5) Present WG at the next LHCOPN/LHCONE meeting in late September in Oslo
   - That's what I'm doing right now
6) Full first WG Meeting – Soon