SDN Next Generation Integrated Architecture For HEP and Global Science.

H. Newman (*), M. Spiropulu, J. Balcas, D. Kcira, I. Legrand,
A. Mughal, J.R. Vlimant, R. Voicu

High Energy Physics, California Institute of Technology
Pasadena, CA 91125

(*) Presenter

Abstract: I describe a software-defined global system under development by Caltech and partner network teams in support of the LHC and other major science programs that coordinates workflows among hundreds of multi-petabyte data stores and petascale computing facilities interlinked by 100 Gbps networks, and the Exascale systems needed by the next decade.

1. Introduction

We are entering a new era of exploration and discovery in many fields, from high energy physics and astrophysics to climate science, genomics, seismology and biomedical research, each with its own complex workflow requiring massive computing, data handling and networks. The continued cycle of breakthroughs in each of these fields depends crucially on our ability to extract the wealth of knowledge, whether subtle patterns, small perturbations or rare events, buried in massive datasets whose scale and complexity continue to grow exponentially with time. "Big Data" today ranges from hundreds of terabytes to hundreds of petabytes, exabyte science datasets are just on the horizon, and the first zettabyte science datasets may be less than a decade away.

In spite of technology advances, the largest data- and network-intensive programs supported by the DOE and partner agencies, including the Upgraded High Luminosity LHC program, the LSST and SKA astrophysics surveys, photon-based sciences, the Joint Genome Institute applications, the Earth System Grid and any other data-intensive emerging areas of growth[1], will continue to face unprecedented challenges: in global data distribution, processing, access and analysis, in the coordinated use of massive but still limited computing, storage and network resources, and in the coordinated operation and collaboration within global scientific enterprises each encompassing hundreds to thousands of scientists.

2. Systems for Next Generation Science

The long term strategy is based on co-design of the methods that make best use of the network and computing and storage infrastructures, together with data structures and real-time adaptive algorithms. Rather than writing code for a distributed system assumed to be static and rigid, the success of these programs will depend on the efficient interplay of software with an elastic and diverse set of resources – CPU, storage, and network. In finding an overall optimal solution, new modes of steering, use (and reuse) of data products produced and consumed at many locations, new modes of propagating information on data product availability and the cost of delivery versus re-computation in real time, and interactions among user groups, end sites and the network as a system will need to be developed.

The crux of the solution to this generational challenge lies in the remarkable synergy emerging between:

- Deeply programmable, agile software-defined network (SDN) infrastructures which are evolving towards multi-service multi-domain network "operating systems" interconnecting science teams across regional, national and global distances, and
- Worldwide distributed systems developed by the data intensive science programs, harnessing global workflow, scheduling and data management systems they have developed, which are

[1] See for example: http://www.es.net/assets/pubs_presos/BER-Net-Req-Review-2012-Final-Report.pdf

978-1-943580-07-1/16/$31.00 ©2016 Optical Society of America

enabled by distributed operations and security infrastructures riding on high capacity (but still-passive) networks.

As in many revolutions, the groups working at the intersection of their domain science and computational science and technology will have a crucial role.

A new overarching concept is one of *"consistent operations"*, where the experiments' workflow management systems will be deeply network aware, reactive and proactive, responding to moment-to-moment feedback on actual versus estimated task progress, state changes of the networks and end systems, and a holistic view of workflows with diverse characteristics and requirements serving many fields. This will enable the major science programs to develop a new more efficient operational paradigm based on software-driven bandwidth allocation, load balancing, flow moderation and topology reconfiguration on the fly where needed, leading to full use of the available network as well as computing and storage infrastructures while avoiding saturation and blocking of other network traffic.

While the systems to be developed should be targeted at many programs, taking diverse "process of science" paradigms into account, one fertile area for development (as well as progressive large scale field testing) is the LHC program, which is now on the cusp of its second three year run, anticipated to yield a new round of groundbreaking discoveries, as well as a new level of "global data and network intensity". This is complemented by the very different but equally challenging real-time workflows in diverse fields, including bioinformatics, computational astrophysics, radio astronomy, and oceanic and atmospheric sciences.

3. Leadership CSN Ecosystem for Next-Gen Data Intensive Science

To gauge the great opportunity in terms of CPU resources for the HEP program (using the CMS example at the LHC) one only has to recall that the CPU requirements are expected to grow by 65 to 200 times between now and the HL LHC, while the affordable CPU power obtainable within a fixed budget, including Moore's law and possible code improvements, is estimated to be an order of magnitude less.

For HEP and the Argonne and Oak Ridge LCFs as well as other major HPC facilities such as NERSC, key issues to develop this vision include:

From the client site and science Virtual Organization side (using the HEP example):

- Recasting HEP's generation, reconstruction and simulation codes, case by case, to adapt to the emerging HPC architectures, addressing issues of memory, dataflow versus CPU etc.
- Identifying and matching the units of work in HEP's workflow to the specific HPC resources or sub-facilities well-adapted to the task (after the code recasting step)
- Building dynamic and adaptive "just in time" systems that respond rapidly (on the required timescale) to offered resources as they occur.
- Developing algorithms that effectively co-schedule CPU, memory, storage, IO port, local and wide area network resources
- Developing an appropriate security infrastructure, and corresponding system architectures in hardware and software, that meet the security needs of the LCF
- Applying "machine learning" to optimize the workflow of the HEP experiments, using self-organizing system methods which are well-adapted to such problems;
  while also taking the special parameters, conditions, and restrictions of LCFs into account as part of the workflow
- Exploiting the intense ongoing development of virtualized computing systems, networks and services in the research community and in industry: in the data center, campus and wide area network space aimed at coherent distributed system operations [including software defined networking, network function virtualization, and service chaining, along with emerging higher level concepts]

The key issues for the LCF and other HPC facilities such as NERSC mirror several of the elements, and include

From the HPC facility side:

- Identifying and matching the units of work in HEP's workflow to the specific HPC resources or sub-facilities well-adapted to the task (after the recasting step)
- Building dynamic and adaptive "just in time" systems that respond rapidly (on the required timescale) to offered demands as they occur; including client-side/server-side coordination for a consistent outcome
- Developing algorithms that effectively co-schedule CPU, memory, storage, IO port, local and wide area network resources; with the necessary coordination as above
- Developing an appropriate security infrastructure, and corresponding system architectures in hardware and software, that meet the security needs of the LCF. For the LCFs this means adopting a new mode of ongoing service to a major client in quasi-real time, in a way that can be adapted to meet the LCF's requirements.
- Applying "machine learning" (loosely defined) to optimizing the workflow of the HEP experiments, using self-organizing system methods which are well-adapted to such problems; while also taking the special parameters, conditions, and restrictions of LCFs into account as part of the workflow.
- Exploiting the intense ongoing developments of virtualization of computing systems and services in the research community and in industry: in the case of the LCFs, the recent developments of "site orchestration" of virtualized resources, and even newer concepts of secure ways to bridge the site edge, such next generation Science DMZs or similar edge-bridging methods are relevant.

4. LCF-Edge Data Intensive System Operational Model

A promising direction centers on the use of a new class of LCF-Edge Data Intensive Systems. The use of secure systems at the site perimeter means that security (both human and AI) and countermeasures where needed can be focused on a limited number of subsystems and entities (proxies), so that the manpower burden may be acceptable.

The operational concept is that HEP data be brought into the edge systems in chunks (a petabyte per chunk was mentioned), far enough in advance so that the data is always waiting and ready when the corresponding jobs are scheduled to start. Multiple chunks for different stages of the overall workflow are foreseen, with each chunk identified to have a certain provenance and certain attributes (such as the ratio of CPU to I/O requirements) so that clusters of chunks is matched to an HPC subsystem configured to match the attributes while working with high efficiency of utilization. At a later stage, once can also foresee dynamic restructuing of the HPC resources, especially if they are virtualized in logical "sectors".

Considering the parameters in this problem yields interesting consequences. As of today, a 1 petabyte chunk would occupy a 100 Gbps link if used to 100% capacity for a full 24 hour day. Given the 300 petabytes currently stored by the LHC experiments and the fact that approximately 100 petabytes flowed over the networks in and out of the US in the past year, the 1 petabyte chunks each represent a relatively small "data transaction" compared to the whole task at hand, and so one would like to transport many chunks to and from the LCF. A typical configuration today would thus preferably include several 100 Gbps links today, migrating to several 400 Gbps links within approximately 5 years and several 1 Tbps links by the startup of the High Luminosity LHC a decade from now, depending on the demand evolution and the cost evolution during this period.

As a result, the use and network requirements of such LEDIS facilities will no doubt present a significant challenge and opportunity for the conception and development of the next generation of intelligent networked systems supporting data intensive science.