

## SDN Next Generation Integrated Architecture (SDN-NGenIA) For HEP and Global Scale Science.



**PI: Harvey Newman, Caltech, [newman@hep.caltech.edu](mailto:newman@hep.caltech.edu), 626-395-6656**

The largest data- and network-intensive programs today, from the Upgraded High Luminosity LHC program, to the LSST and SKA astrophysics surveys, to the Joint Genome Institute and many other data-intensive emerging areas of growth<sup>1</sup>, face unprecedented challenges in global exascale data distribution, processing, access and analysis, and in the coordinated use of massive but still limited CPU, storage and network resources. In response to these challenges and opportunities for science, Caltech, working together with ESnet, Fermilab, Starlight/iCAIR, and other key laboratory, university and industry partners<sup>2</sup> proposes to design and develop the first stages of the **SDN Next Generation Integrated Architecture (SDN-NGenIA)** for HEP and global scale science.

The overarching goal is to maximize the discovery potential of scientific collaborations through the development of revolutionary open source products and methods in the SDN, virtualization and global system operation and optimization space. This will be accomplished by exploiting and contributing to the remarkable synergy emerging between: Deeply programmable software-defined agile and adaptive network infrastructures that are emerging as multi-service multi-domain network “operating systems” interconnecting next generation Science DMZs, and the systems developed by the data intensive science programs harnessing global workflow, scheduling and data management systems. While the initial focus will be on the challenging LHC use case, the products developed will be general, and apply to many fields of data intensive science. These will be informed by the LSST and bioinformatics/genomics use cases, which will be explored during the latter part of the project.

We will construct autonomous, intelligent site-resident services that dynamically interact with network-resident services, and with the science programs’ principal data distribution and management tools, to request or command network resources in support of high throughput petascale to exascale workflows, using:

- (1) smart middleware to interface to SDN-orchestrated data flows over network paths with guaranteed bandwidth all the way to a set of high performance end-host data transfer nodes (DTNs),
- (2) protocol agnostic SDN-based QoS and traffic shaping services at the site egress that will provide stable, predictable data transfer rates, and auto-configuration of data transfer nodes, and
- (3) host- and site agent systems coupled to machine learning methods

Specific work areas include: (1) deep site orchestration among virtualized clusters, storage subsystems and subnets to successfully co-schedule CPU, storage and network resources; (2) science-program designed site architectures, operational modes, and priorities adjudicated across multiple network domains and among multiple virtual organizations; (3) seamlessly extending end-to-end operation across both extra-site and intra-site boundaries through the use of next generation Science DMZs; (4) funneling massive sets of streams to DTNs at the site edge hosting petascale buffer pools configured for flows of 100 Gbps and up, exploiting state of the art data transfers where possible; and (5) unsupervised and supervised machine learning and modeling methods to drive the optimization of end-to-end workflow involving terabyte to multi-petabyte datasets.

The services developed will interface seamlessly to the adaptive SENOS network operating system being developed in companion projects by ESnet, Fermilab and others (including Caltech as a partner in some areas), to ensure that authorized applications achieve full throughput. The accumulated knowledge will serve to inform the design of the following generations of distributed petabit/sec systems, including continental scale instruments such as SKA, and real-time leadership computing systems of the next decade harnessing zettabyte datasets.

---

<sup>1</sup> See for example: [http://www.es.net/assets/pubs\\_presos/BER-Net-Req-Review-2012-Final-Report.pdf](http://www.es.net/assets/pubs_presos/BER-Net-Req-Review-2012-Final-Report.pdf)

<sup>2</sup> Note that this is a single institution proposal from Caltech. The partners mentioned here are at no cost to this proposal. They are funded from other sources, including the SENSE proposal funded by ASCR, and other proposals where applicable.