

SDN-NGenIA A Software Defined Next Generation integrated Architecture for HEP and Data Intensive Science

Tuesday, 11 October 2016 16:30 (15 minutes)

The SDN Next Generation Integrated Architecture (SDN-NGeNIA) program addresses some of the key challenges facing the present and next generations of science programs in HEP, astrophysics, and other fields whose potential discoveries depend on their ability to distribute, process and analyze globally distributed petascale to exascale datasets.

The SDN-NGenIA system under development by the Caltech and partner HEP and network teams is focused on the coordinated use of network, computing and storage infrastructures, through a set of developments that build on the experience gained in recently completed and previous projects that use dynamic circuits with bandwidth guarantees to support major network flows, as demonstrated across LHCONe and in large scale demonstrations over the last three years, and recently integrated with CMS' PhEDEx and ASO data management applications.

The SDN-NGenIA development cycle is designed to progress from the scale required at LHC Run2 (0.3 to 1 exabyte under management and 100 Gbps networks) to the 50-100 Exabyte datasets and 0.8-1.2 terabit/sec networks required by the HL LHC and programs such as the SKA and the Joint Genomics Institute within the next decade. Elements of the system include (1) Software Defined Network (SDN) controllers and adaptive methods that flexibly allocate bandwidth and load balance multiple large flows over diverse paths spanning multi-domain networks, (2) high throughput transfer methods (FDT, RDMA) and data storage and transfer nodes (DTNs) designed to support smooth flows of 100 Gbps and up, (3) pervasive agent-based real-time monitoring services (in the MonALISA framework) that support coordinated operations among the SDN controllers, and help triggering re-allocation and load-balancing operations where needed, (4) SDN transfer optimization services developed by the teams in the context of OpenDaylight, and (5) machine learning coupled to prototype system modeling, to identify the key variables and optimize the overall throughput of the system, and (6) a "consistent operations" paradigm that limits the flows of the major science programs to a level compatible with the capacity of the campus, regional and wide area networks, and with other network usage.

In addition to general program goal of supporting the network needs of the LHC and other science programs with similar needs, a recent focus is the use of the Leadership HPC facility at Argonne National Lab (ALCF) for data intensive applications. This includes installation of state of the art DTNs at the site edge, specific SDN-NGenIA applications and the development of prototypical services aimed at securely transporting, processing and returning data "chunks" on an appropriate scale between the ALCF, and LHC Tier1 and Tier2 sites: from tens of terabytes now, to hundreds of petabytes using 400G links by 2019, and a petabyte at a time using terabit/sec links when the first exaflop HPC systems are installed circa 2023.

Tertiary Keyword (Optional)

Data processing workflows and frameworks/pipelines

Primary Keyword (Mandatory)

Network systems and solutions

Secondary Keyword (Optional)

Distributed data handling

Primary author: Prof. NEWMAN, Harvey (Caltech)

Co-author: KCIRA, Dorian (California Institute of Technology (US))

Presenter: KCIRA, Dorian (California Institute of Technology (US))

Session Classification: Posters A / Break

Track Classification: Track 9: Future directions