SC19 Network Research Exhibition: Demonstration Preliminary Abstract

SANDIE: SDN-Assisted NDN for Data Intensive Experiments

Submitted by the SANDIE Team

Abstract

The SANDIE (SDN-Assisted NDN for Data Intensive Experiments) project is based on the novel yet well-founded Named Data Networking (NDN) architecture, supported by advanced Software Defined Network (SDN) services to meet the challenges facing data-intensive science programs such as the Large Hadron Collider (LHC) high energy physics program. The SANDIE project has developed a new and highly effective approach to data distribution, processing, gathering and analysis of results to accelerate the workflow for the CMS experiment at the LHC. This demonstration will exhibit improved performance of the SANDIE system by leveraging three key components, the high-speed NDN-DPDK forwarder, the VIP jointly optimized caching and forwarding algorithm, and an NDN-based filesystem plugin for XRootD. The demonstration results strongly support the overall SANDIE project goal of providing more rapid and reliable data delivery, with varying patterns and granularity over complex networks, progressing in scale from the Terabyte to eventually the Petabyte range in support of the LHC physics and other data-intensive programs.

Goals

This demonstration is designed to exhibit improved performance of the SANDIE system for workflow acceleration within large-scale data-intensive programs such as the LHC high energy physics program. To achieve high performance, the demonstration will leverage three key components, the recently developed high-speed NDN-DPDK forwarder, the VIP jointly optimized caching and forwarding algorithm, and an NDN-based filesystem plugin for XRootD that allows XRootD servers to query NDN via the Open Storage System (OSS) plugin player. The demonstration activities will take place over a network testbed connecting the SANDIE participant institutions (Northeastern, Caltech and Colorado State) and the SC19 Caltech booth.

1. NDN-DPDK Forwarder and XRootD OSS plugin

To achieve high throughput, we adopt the NDN-DPDK forwarder recently developed at the National Institute of Standards and Technology (NIST).[1] The NDN-DPDK forwarder follows the NDN [1] protocol and leverages concurrency to significantly improve packet processing performance and throughput. In addition, our applications utilize DPDK to produce NDN packets which are encapsulated in Ethernet packets, thereby eliminating the overhead introduced by the overlay NDN stack. With both the NDN-DPDK forwarder and DPDK-based applications running over the testbed, we will demonstrate significantly improved throughput compared with that of our demonstration at SC18.

We will also demonstrate a XRootD OSS plugin implementation based on NFD (NDN Networking Forwarding Daemon), by using the xrdcp XRootD client to copy files of different sizes. This implementation will not showcase high throughput. Rather, it will demonstrate that the high-throughput NDN-DPDK consumer-based application can and will be integrated into an XRootD file system plugin.

2. VIP Jointly Optimized Caching and Forwarding Algorithm

We will demonstrate the performance of an adaptive, distributed jointly optimal caching and forwarding algorithm, namely the VIP algorithm [2], over the SANDIE testbed. The VIP algorithm uses the metric of *Virtual Interest Packets (VIP)* to indicate the level of local demands for different data contents in the network. The VIP metrics are used to determine caching and forwarding strategies to maximally fulfill user demands. The VIP algorithm has been implemented with the NDN-DPDK forwarder. Specifically, each node keeps a *VIPTable* which contains the VIP statistics of its neighbors and itself. A periodic function is designed to update such statistics in a distributed manner. VIP information is formatted, packetized and treated as data in the network, and each node periodically generates Interest Packets to fetch such data and updates its local VIPTable. The caching and forwarding decisions are then determined based on the VIPTable. We plan to operate the VIP algorithm on a set of popular datablocks based on observations of real CMS traffic, and demonstrate the advantage in terms of throughput, delay, and cache hit rates and content distribution patterns.

3. Demonstration Topology

The figure below outlines our demo topology both at layer 2 and layer 3. The solid lines between the Caltech booth, Scinet, Cenic, Esnet, Internet2, Northeastern University (NEU), Colorado State University (CSU), and Caltech, represent the

---

[1] Codebase available at https://github.com/usnistgov/ndn-dpdk .

layer 2 topology. The layer 2 topology is provided by SCinet in collaboration with Internet2 and the regional providers to NEU, CSU, and Caltech. The nodes at the supercomputing booth are connected to all other nodes (Caltech, NEU, and CSU) at layer 2 with the use of VLANs. The NDN forwarder we will use for this demo, NDN-DPDK, runs on top of this layer 2 topology as the layer 3 forwarder. The NDN topology at layer 3 is shown by the red and the green dotted lines - the red dotted lines represent direct layer 3 paths between the data consumer (at the booth) and the data sources (where the data producers reside, one at Caltech and one at NEU). We will also have indirect (layer 3, NDN) paths to these data producers through CSU. When the direct paths do not work, the NDN requests will go through the CSU node.

CMS jobs on LHC-generated data make use of our system, as well as the maximum throughput level we can reach. In addition, we plan to show network-based load-balancing using the adaptive optimized caching and forwarding algorithm over our NDN network.

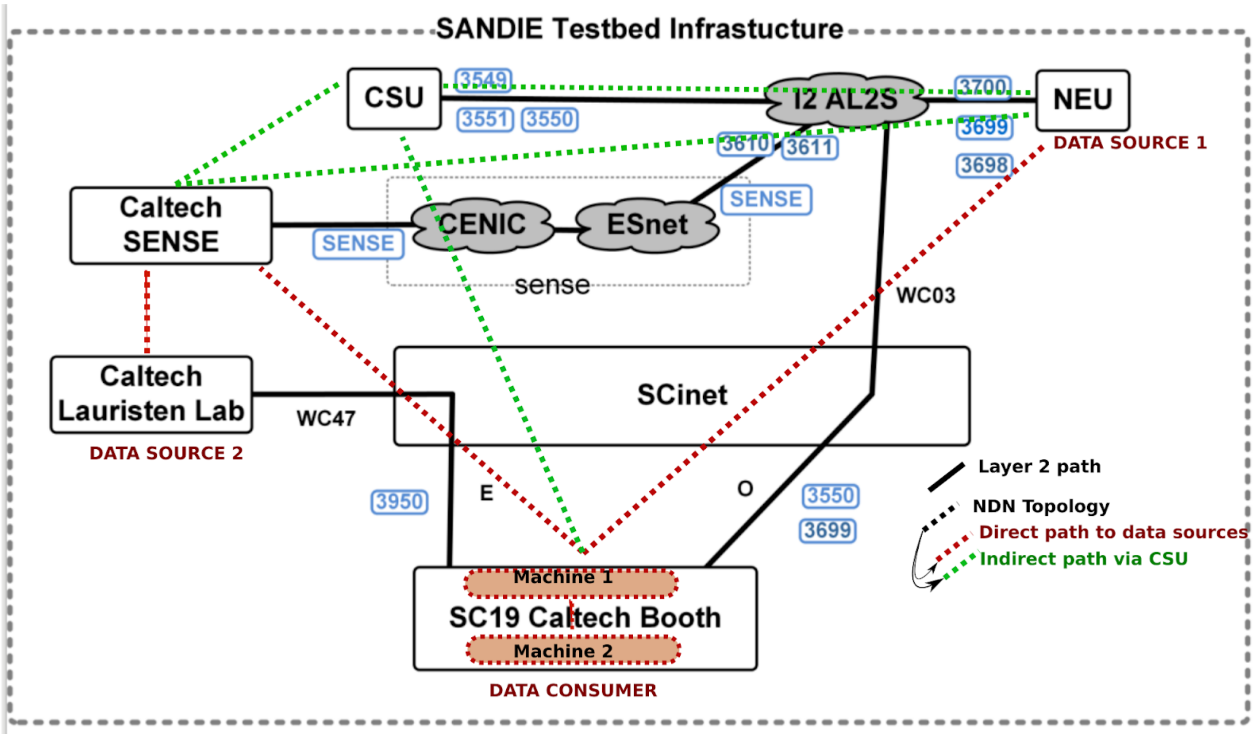

4. Demonstration Activities

We will demonstrate the functionality and speed of the newly developed OSS NDN-based XRootD plugin and the complementary NDN producer by using XRootD clients (xrdcp) to copy files over the NDN network, and also by running scientific applications developed CMS Software components (CMSSW) on byte ranges from different files in a dataset. The client will run on the server at the Caltech booth on the SC19 floor, which will be sending requests (Interest packets) for Data over NDN. If the Data is not available on the local cache or any other cache along the path to a producer in the network, the request forwarded by NDN forwarders will reach the producer that publishes the data. By using this method, the XRootD redirectors are no longer needed. Instead, we require only a local XRootD server, on the client machine using our OSS plugin. Our testbed will leverage the DPDK-based forwarder to improve throughput. We will demonstrate how files are copied with integrity over the (high-speed) NDN network, how

Involved Parties

- Edmund Yeh, Northeastern Univ., eyeh@ece.neu.edu
- Ran Liu, Northeastern Univ., rliu1@ece.neu.edu
- Yuanhao Wu, Northeastern Univ., wu.yuanh@husky.neu.edu
- Volkan Mutlu, Northeastern Univ., fvmutlu@ece.neu.edu
- Yuezhou Liu, Northeastern Univ., liu.yuez@husky.neu.edu
- Harvey Newman, Caltech, newman@hep.caltech.edu
- Catalin Iordache, Caltech, catalinn.iordache@gmail.com
- Raimondas Sirvinskas, Caltech, raimis.sirvis@gmail.com
- Justas Balcas, Caltech, justas.balcas@cern.ch
- Susmit Shannigrahi, Tennessee Tech, sshannigrahi@tntech.edu
- Chengyu Fan, Colorado State Univ., chengyu.fan@colostate.edu
- Craig Partridge, Colorado State Univ., Craig.Partridge@colostate.edu

References

[1] Zhang, et al. "Named data networking (ndn) project.", technical report, October 2010.
[2] Yeh, et al. "VIP: A framework for joint dynamic forwarding and caching in named data networks." Proceedings of the 1st ACM Conference on Information-Centric Networking. ACM, 2014.