SC25 Network Research Exhibition: Demonstration Publishable Abstract

NRE 106 Multi-Resource Cyberinfrastructure Services for Science Domain Workflows via SENSE

Tom Lehman, ESnet, tlehman@es.net

Abstract

The Software-defined network for End-to-end Networked Science at Exascale (SENSE) system enables the automated provision of Cyberinfrastructure services with integrated network, compute, and storage resources. These services can be Application Programming Interface (API) driven, which allows science workflow agents to initiate, monitor, and optimize through the lifecycle of their workflow.

SENSE includes a model-based orchestration system which operates between the SDN layer controlling the individual networks/end-sites, and science workflow agents/middleware. The key SENSE components are the "Orchestrator", "Network Resource Manager" and "Site Resource Manager" which together enable advanced features in the areas of multi-resource services, automated control, real time optimizations, and full life-cycle monitoring. Figure 1 shows this SENSE architecture.

This demonstration will show SENSE in operation across a variety of infrastructures including; Networks and Exchange Points (ESnet, Internet2, GNA-G AutoGOLE, CENIC, PacificWave, AMPATH/AMLight, StarLight, NetherLight Exchange), Edge Systems (National Research Platform Kubernetes Cluster, Data Transfer Nodes, LHC XRootD/dCached systems, Amazon Web Services), and the FABRIC infrastructure which includes network embedded compute/storage resources.

SENSE services and features which benefit multiple science workflows will be demonstrated including the following:

- Rucio/FTS/XRootD/dCache LHC CMS/ATLAS
- DOE Integrated Research Infrastructure (IRI)
- GNA-G AutoGOLE Applications
- National Research Platform (NRP) Integration
- Realtime workflow specific monitoring and workflow topology visualization

Goals

This demonstration intends to show how workflow-driven multi-domain, multi-resource service provisioning, with full life-cycle monitoring and adjustments, greatly enhances domain science workflow performance. The example workflows will be used to automate the workflow specific infrastructure setup, and to compare performance with and without SENSE managed cyberinfrastructure services.

In addition, this demonstration aims to show how network operators can use SENSE services to facilitate network infrastructure configuration and user services provisioning.

Impacts

Data transfers and associated network services typically do not have the same level of deterministic resource allocation as in place for compute resources. As a result data transfers can introduce a degree of uncertainty in workflow operations, and the associated lack of network information does not allow for either the workflow operations or the network use to be optimized.

The SENSE functionally demonstrated at SC25 is intended to show how network resources and data transfer management can be elevated to the same level as compute and storage in the context of services access, scheduling, and life cycle management for science workflow orchestration. A key next step is to show how this functionality can be easily deployed at multiple sites and networks across the R&E ecosystem. The SCinet infrastructure enables demonstration of this capability at high performance, geographic distribution, and resource heterogeneity levels.

Figure 2 shows an example workflow based on the LHC data movement stack of Rucio / File Transfer System / XRootD / dCache which has been enhanced to use SENSE for deterministic network services.

Resources

SENSE controlled network elements will be deployed within the SCinet NOC. This will include an UMASS/MGHPCC provided Juniper QFX5130 switch/router will be dynamically controlled to build demonstration specific topologies from wide area connected locations, local SCInet National Research Platform (NRP) cluster, and direct connections to Caltech, StarLight, and Ciena booths. Figure 3 shows this demonstration infrastructure over which SENSE will operate.

SENSE will provision a variety of demonstration topologies which combine Layer 2, Layer 3 network services and end system configurations tailored to specific workflow requirements.

Involved Parties

- Tom Lehman, ESnet, tlehman@es.net
- Xi Yang, ESnet, xiyang@es.net
- Justas Balcas, ESnet, jbalcas@es.net
- Alberto Jimenez, ESnet, ajimenez@es.net
- Kiran Vasu, ESNet, kvasu@es.net
- John MacAuley, ESnet, macauley@es.net
- Chin Guok, ESnet, chin@es.net
- Inder Monga, ESnet, inder@es.net
- Mohammad Firas Sada, UCSD, mfsada@ucsd.edu
- John Graham, UCSD, jjgraham@ucsd.edu
- Frank Würthwein, UCSD, fkw@ucsd.edu
- Tom Hutton, UCSD, hutton@ucsd.edu
- Chris Hughes, UCSD, crhughes@ucsd.edu
- Aashay Arora, UCSD, aarora@ucsd.edu
- Diego Davila, UCSD, didavila@ucsd.edu
- Harvey Newman, Caltech, newman@hep.caltech.edu
- Preeti Bhat, Caltech, preeti@caltech.edu
- Andres Moya, Caltech, amoya@caltech.edu
- Marcos Felipe Schwarz, RNP, marcos.schwarz@rnp.br
- Everson Scherrer Borges, IFES, everson@ifes.edu.br
- Christopher Bruton, cbruton@cenic.org
- Phil Demar, FNAL, demar@fnal.gov
- Asif Shah, FNAL, sshah@fnal.gov
- Dmitry Litvintsev, FNAL, litvinse@fnal.gov
- Eduardo Bach, UMASS, eduardobach@umass.edu
- Rafael Coelho Lopes de Sa, UMASS, rclsa@umass.edu
- Jessa Westclark, UMASS, jessicaw@umass.edu
- Mazahir Hussain, KISTI, mazahir.hussain@kisti.re.kr
- Buseung Cho, KISTI, bscho@kisti.re.kr
- Fei Yeh, Northwestern University, fyeh@northwestern.edu
- Jim Hao Chen, Northwestern University, jim-chen@northwestern.edu
- Joe Mambretti, Northwestern University, j-mambretti@northwestern.edu

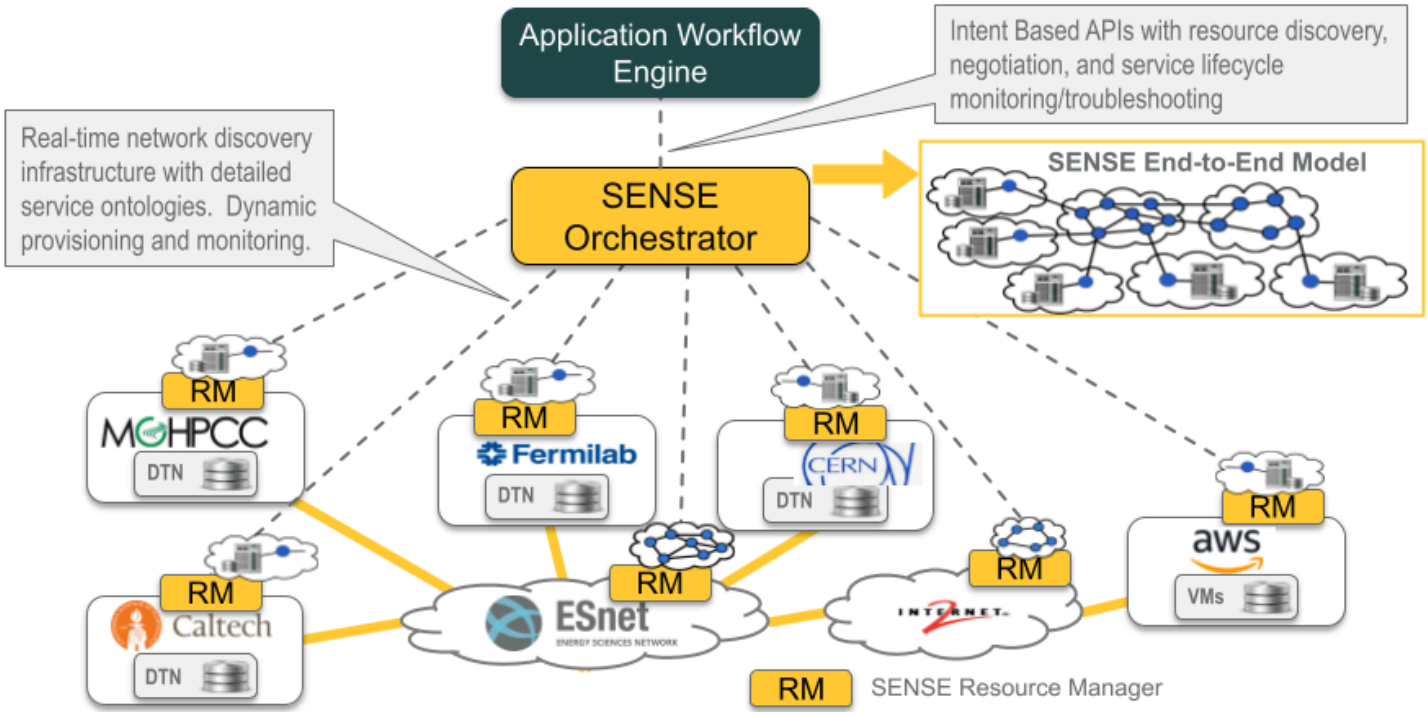

Figure 1 SENSE Architecture

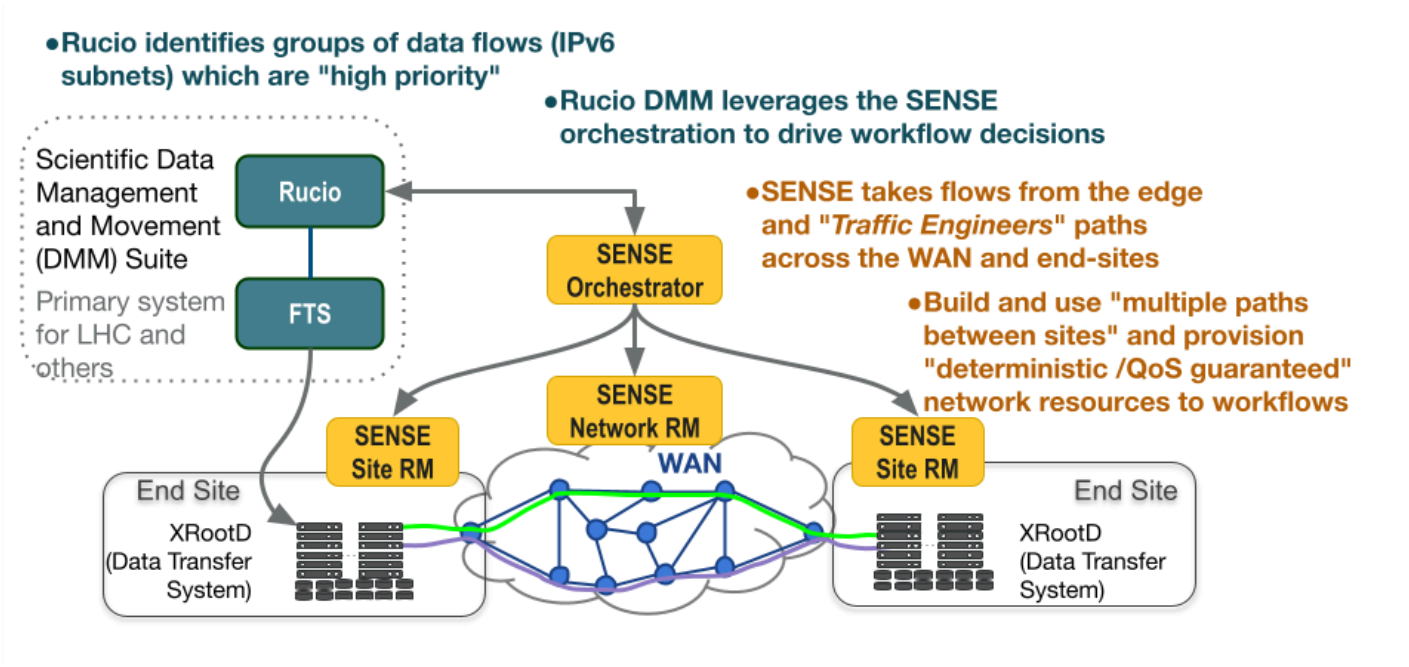

Figure 2 RUCIO/File Transfer System/XRootd/dCache Data Management an Movement Workflows with SENSE Priority Network Services

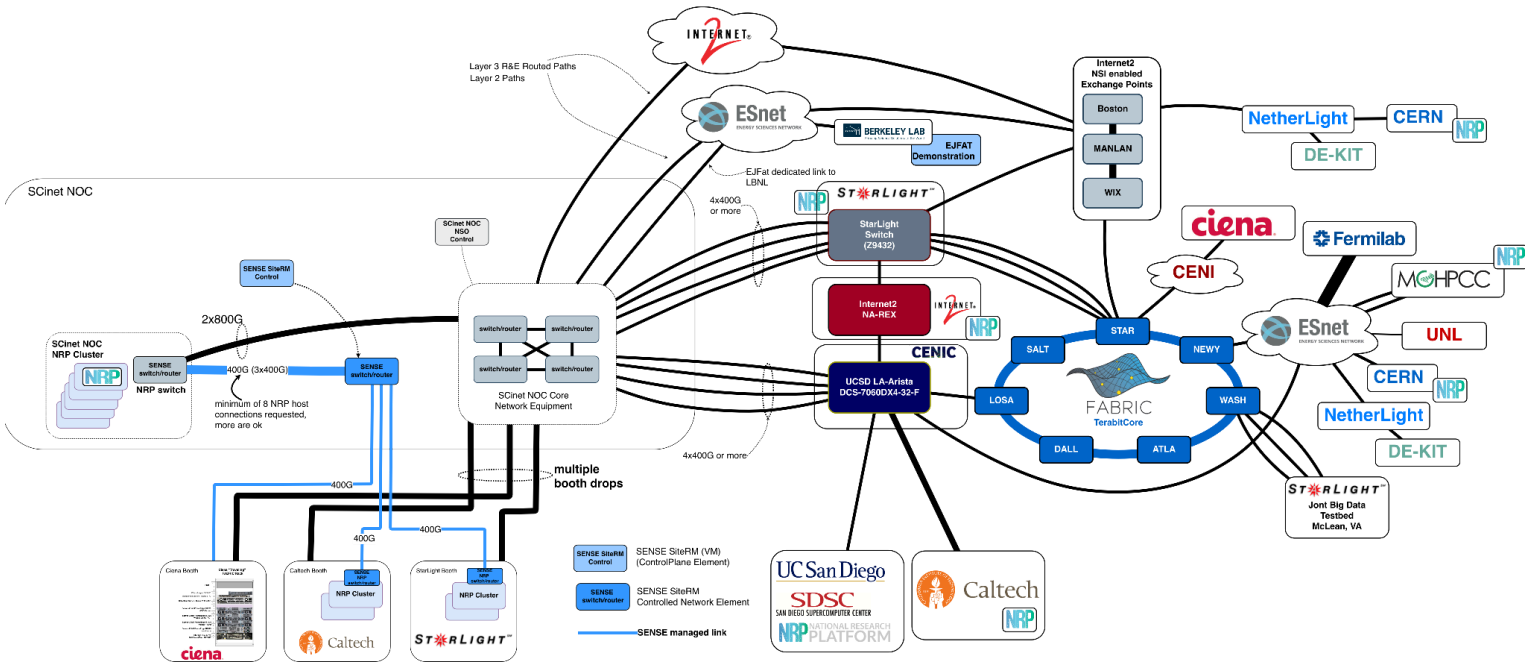

Figure 3 SC25 SCinet NRE Infrastructure with SENSE Controlled Network Element and NRP Cluster Integration