

SC19 Network Research Exhibition: Demonstration Preliminary Abstract

400 Gbps WAN Services: Architecture, Technology and Control Systems

Joe Mambretti, Jim Chen, Fei Yeh, Se Young Yu

International Center for Advanced Internet Research - Northwestern University

j-mambretti, jim-chen, fyeh@northwestern.edu, young.yu@northwestern.edu

Abstract

Data production among science research collaborations continues to increase. Recently, LHC planners estimated that the annual rate of increase for the foreseeable future will be 50% per year. Consequently, the networking community must begin preparing for 400 Gbps WAN services. 100 Gbps WAN services, which have become ubiquitous, have been implemented for over ten years. Before they were widely deployed, it was necessary to develop techniques to effectively utilize that level of capacity. Similarly, the requirements and implications of 400 Gbps WAN services must be explored at scale. These demonstrations showcase large scale E2E 400 Gbps WAN services from the StarLight International/National Communications Exchange Facility in Chicago to the SC19 venue in Denver.

Goals

With its research partners including the SCinet WAN group, the International Center for Advanced Internet Research (iCAIR) at Northwestern University is designing, implementing and demonstrating an E2E 400 Gbps WAN service from the StarLight International/National Communications Exchange Facility in Chicago to the SC19 venue in Denver. Just as techniques used for utilizing 10 Gbps paths did not directly scale to 100 Gbps. The techniques being used for 100 do not directly scale to 400 Gbps. Consequently, multiple issues must be addressed to enable to transition to 400 Gbps WAN services.

1. At both ends of the E2E path, iCAIR will implement 400 Gbps switches.
2. Those switches will be connected to optimized 100 Gbps Data Transfer Nodes (DTNs).
3. These demonstrations will leverage iCAIR's experimental research into the optimal design, configuration, components, and integration technologies for 100 Gbps (DTNs), including techniques for kernel bypass using

zero-copy for memory and disk copy to avoid bottlenecks in 100 Gbps data transfer and optimal affinity bindings for NUMA architecture for higher resource utilization

4. The demonstration will also showcase middleware for reliable high-speed network data transfer to orchestrate infrastructure resources for optimal high performance transfers.
5. We will also show measurement techniques for real-time monitoring, benchmarking and evaluation at 400 Gbps.

Resources

Required resources from SCinet are a portion of the 10*100 Gbps circuits requested from the SCinet WAN group from the StarLight facility in Chicago to the StarLight booth on the SC19 showfloor.

Involved Parties

- Joe Mambretti, iCAIR, j-mambretti@northwestern.edu
- Jim Chen, iCAIR, jim-chen@northwestern.edu
- Fei Yeh, iCAIR, fyeh@northwestern.edu
- Se Young Yu, iCAIR, young.yu@northwestern.edu
- StarLight International/National Communications Exchange Facility and Consortium
- Metropolitan Research and Education Network (MREN)
- SCinet

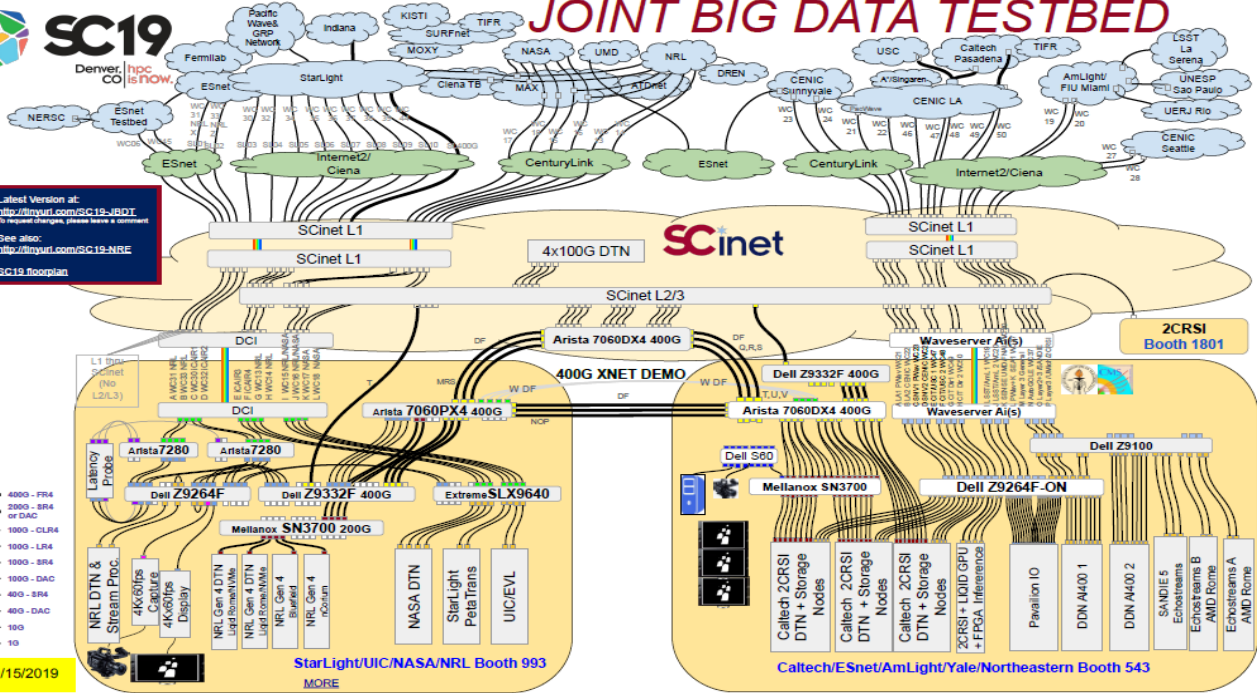


SC19
Denver, CO **hpc is how.**

JOINT BIG DATA TESTBED

Latest Version at:
<http://tinyurl.com/SC19-JBDT>
To request changes, please leave a comment
See also:
<http://tinyurl.com/SC19-NRE>
SC19 floorplan

- 400G - FR4
 - 200G - SR4 or DAC
 - 100G - CLR4
 - 100G - LR4
 - 100G - SR4
 - 40G - SR4
 - 40G - DAC
 - 10G
- 10/15/2019



b