

SC21 Network Research Exhibition: Demonstration Abstract

400 Gbps E2E 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, a long term trend that will accelerate with the advent of high luminosity research instrumentation. Consequently, the networking community must begin preparing for service paths beyond 100 Gbps, including 400 Gbps WAN and LAN services. Before 100 Gbps WAN/LAN services were widely deployed, it was necessary to develop techniques to effectively utilize that level of capacity. Today, 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 SC21 venue.

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 SC21 venue. The techniques being used for 100 Gbps do not directly scale to 400 Gbps. Consequently, multiple issues must be investigated and resolved 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 multiple 100 Gbps data transfer

over 400 Gbps WAN 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 WAN are a 400 Gbps SCinet path between McClean Va and the StarLight exchange, continuing to the StarLight booth at the venue, and a portion of the 1.2 Tbps Gbps requested from the StarLight facility in Chicago to the SC21 show floor.

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
- Rod Wilson, Ciena, rwilson@ciena.com
- Marc Lyonais, Ciena, mlyannai@ciena.com
- Scott Kohlert, Ciena, skohlert@ciena.com
- Gauravdeep Shami, Ciena, gshami@ciena.com
- Linden Mercer, Navel Research Lab, linden.mercer.ctr@nrl.navy.mil
- Basil Decina, Naval Research Lab, basil.decina@nrl.navy.mil
- Bill Fink NASA Goddard Space Flight Center, bill@sci.gsfc.nasa.gov,
- Paul Lang, NASA Goddard Space Flight Center, paul.a.lang@nasa.gov

- StarLight International/National Communications Exchange Facility and Consortium
- Metropolitan Research and Education Network (MREN)
- SCinet

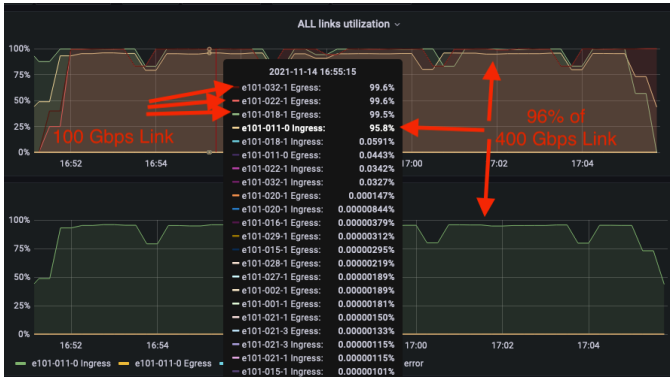



Figure 1. 400 Gbps transfer between StarLight and SCinet St. Louis. 380 Gbps over a 400 Gbps WAN is utilized, as well as other 3 x 100 Gbps is utilized at 99%.

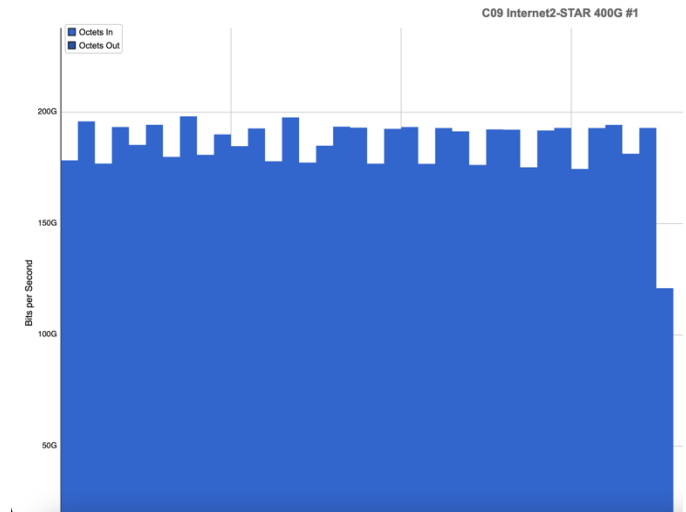
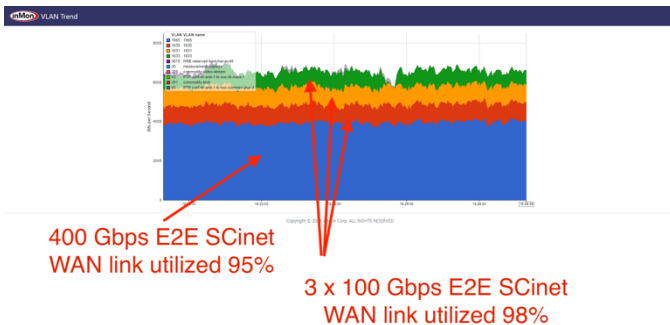


Figure 3. 200 Gbps E2E transfer with single interface



Figure 2. SCinet traffic monitor for 400Gbps E2E WAN NRE



400 Gbps E2E SCinet WAN link utilized 95%

3 x 100 Gbps E2E SCinet WAN link utilized 98%