

DAFT --- Results From
NASA High End Computing (HEC)
WAN File Transfer Experiments/Demonstrations
Super Computing 2016 (SC16)
• Bill Fink

- Computational And Information Science and Technology Office (CISTO)
 - High End Computer Networking Team (HECN), Code 606.1
 - NASA Goddard Space Flight Center
 - November 14-17 2016

Overview

- These SC16 Network Research Exhibition (NRE) demonstrations were designed to showcase capabilities for transporting extremely large scale data for petascale scientific research using two 100 Gbps long distance Wide Area Network (WAN) and Local Area Network (LAN) circuits.
- For the SC16 Super Computing Conference in Salt Lake City, a consortium of researchers implemented a national optical network testbed consisting of multiple 100 Gbps optical circuits using ESnet and exchange facilities (see attached diagram).
- This testbed is an extension of an existing testbed that was established to develop advanced services and technologies for next generation data-intensive petascale science, under the GSFC High End Computing Program.
- These demonstrations build on earlier efforts related to demonstrations and experiments on a persistent HEC testbed connecting the GSFC and Starlight, and on national WAN and LAN testbeds implemented for SC10 thru SC15.

Collaborating Organizations

- NASA Partners in “Near 200G Disk-to-Disk Network Data Transfers” Special SC16 Demonstration/Evaluation Experiments
- Organizations: Energy Science Network (Esnet), International Center for Advanced Internet Research, Northwestern University (ICAIR), Mid-Atlantic Crossroads, Maryland University (MAX), StarLight International/National Communications Exchange Facility Consortium, Metropolitan Research and Education Network (MREN), Open Cloud Consortium (OCC), Laboratory for Advanced Computing, University of Chicago (LAC), Large Scale Networking Coordinating Group of the Networking and Information Technology Research and Development (NITRD) program.
- Corporations providing loaner equipment include: Arista, Brocade, Ciena, Edgecore
- On-site SC16 support from Brocade (Matt Lowe and Wilbur Smith).

Pre SC16 Successes and Issues

- On visit to CalTech lab, saw Supermicro NVMe server and various Non-Volatile Memory PCI Express (NVMe) drives which were very fast but too expensive for our budget.
- Found NVMe U.2 to M.2 adapter that helped reduce the price of the server to half the price it would have been if we use U.2 drives.
- On Oct 19, received NVMe drives after a long delay, only to troubleshoot a PCIe riser issue and get replacement riser on Oct 25th. Completed build and testing of this NVME server in time to send to Starlight on Oct 28th.
- Local testing between two of the HECN custom built 200G NVMe servers resulted in 181 GigaBit per second disk-to-disk transfer (2 processors, 16 NVMe drives, 16 simultaneous 64 GB transfers (total 1TB) transferred in less than 50 seconds.

SC16 Successes and Issues

- Worked through a few setup issues as WAN circuits were moved due to errors found on circuit.
- Successfully obtained 178 Gbps disk-to-disk transfer consisting of sixteen simultaneous transfers of 64 GigaByte datasets, for a total of one TeraByte in about 50 seconds. The transfers were performed both to and from a server on the exhibit floor at SC15, from or to a server located either at NASA Goddard, or one that was deployed at the Starlight facility.

NASA HEC WAN File Transfer

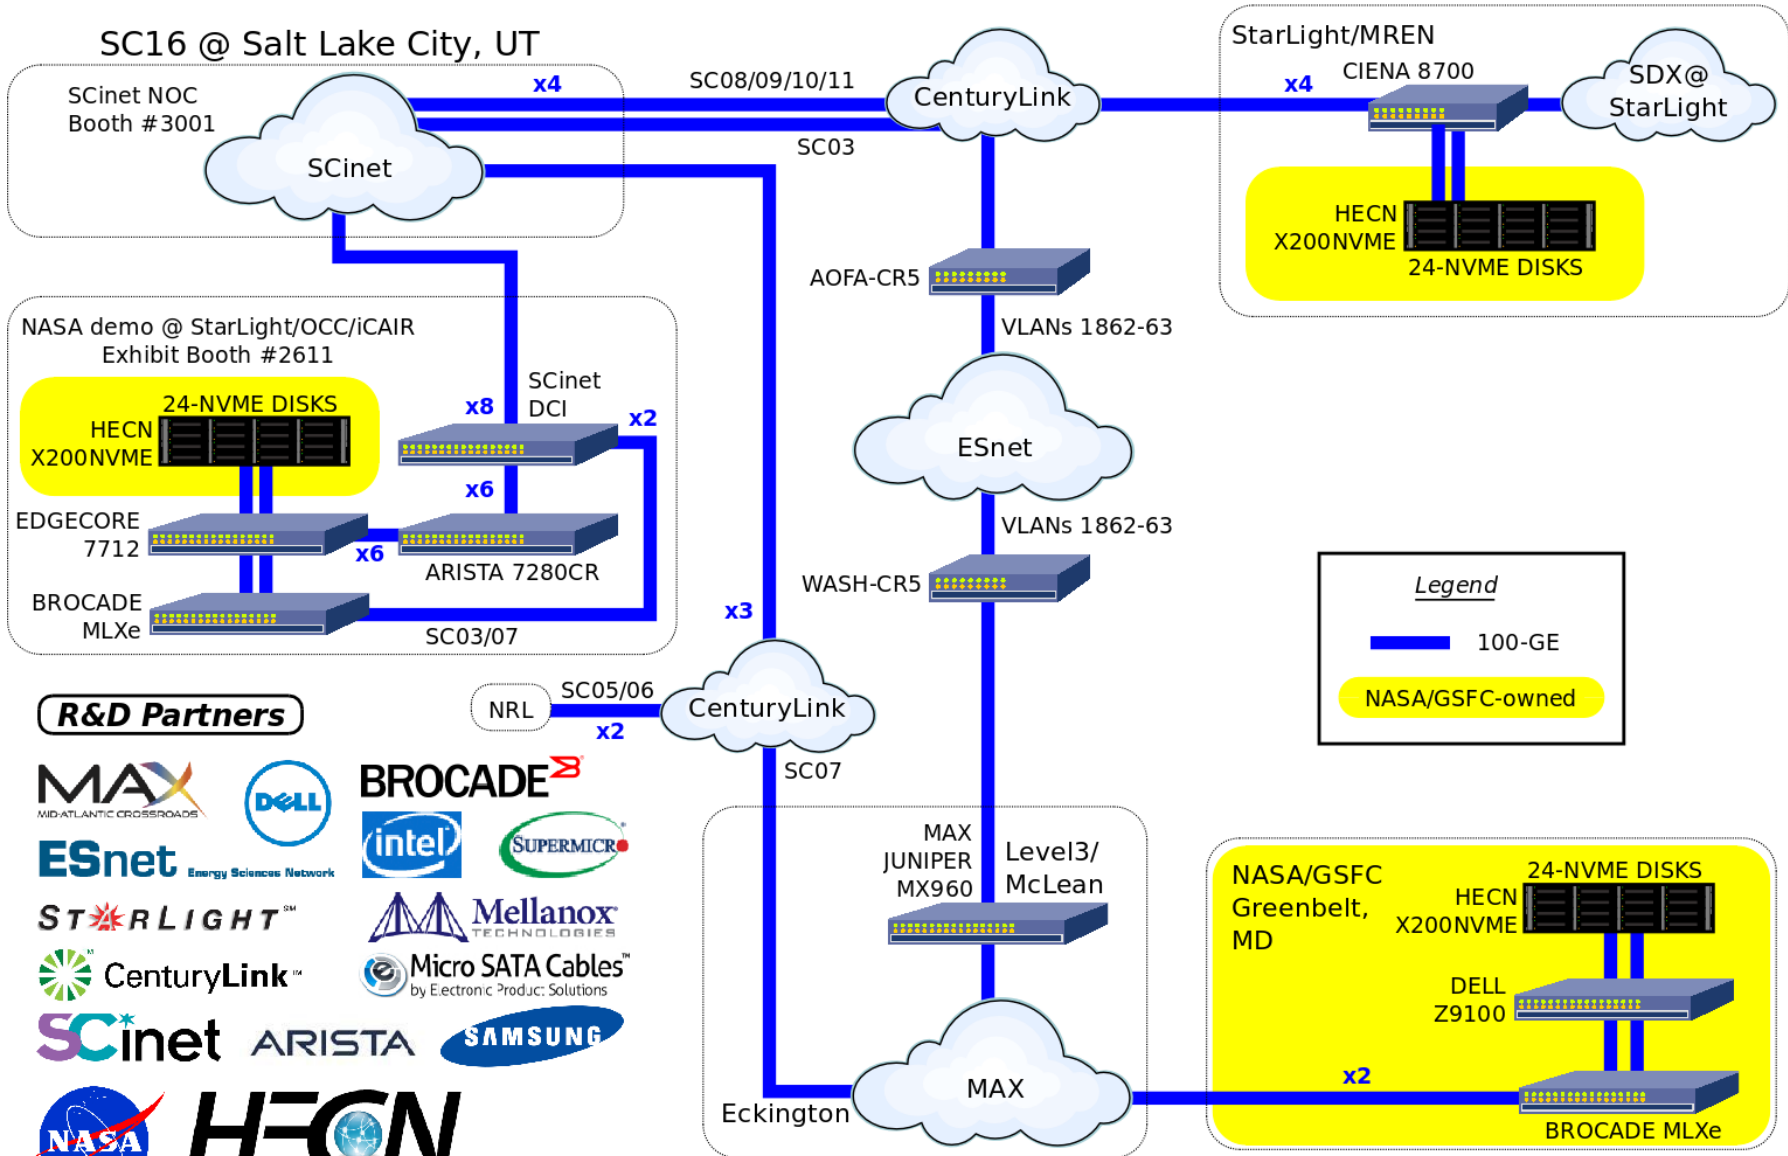
Experiments/Demonstrations At SC16

- Special SC16 Demonstration/Evaluation Experiments
- Use of three custom NASA/HECN Team built network-testing-raid-servers deployed into the LAC/iCAIR booth, Starlight International and National Communications Exchange facility in Chicago, and at GSFC, capable of:
- 181 Gbps back-to-back uni-directional disk-to-disk file copies, using two 100G interfaces and 16 NVMe drives per server.
- Demonstrate/Evaluate interoperability between multiple vendor 100G products from Arista, Brocade, Ciena, Dell/Force10, Edgecore, Fujitsu, Juniper, over SCinet, Century Link, ESnet, Starlight, and MAX/DRAGON
- Achieved 178 Gbps TCP WAN disk-to-disk from SC16 in Salt Lake City to or from either Starlight International and National Communications Exchange in Chicago or NASA Goddard Space Flight Center in Greenbelt, MD

SC16

Demonstrations of 200 Gbps Disk-to-Disk WAN File Transfers using Parallelism across NVMe Drives


An SC16 Collaborative Initiative Among NASA and Several Partners



R&D Partners



SC16 – Bill Fink and Paul Lang touching 200G NVME server



High End Computer Networking (HECN) Team



Bill Fink
Acting Project Lead
NASA/GSFC



Paul Lang
Network Engineer
NASA/ADNET Systems



Aruna Muppalla
Network Engineer
NASA/ADNET Systems



Jeff Martz
Network Engineer
NASA/ADNET Systems



Mike Stefanelli
Network Engineer
NASA/ADNET Systems



Pat Gary
(In Memoriam)
43 Years NASA/GSFC

HECN SC16 200G NVME server

Description	Model	Price	Qty	Subtotal
24bay NVMe server	Supermicro SYS-2028U-TN24R4T+	\$3,600.00	1	\$3,600.00
8-core processor	Intel E5-2667 V4	\$2,057.00	2	\$4,114.00
CPU cooler	Dynatron R14	35.00	2	\$70.00
8GB 2400MHz DDR4	Kingston ECC 2400 MHz DDR4 memory	\$53.00	8	\$424.00
100G NIC	Mellanox MCX415A-CCAT	\$990.00	2	\$1,980.00
NVMe drive	Samsung 512GB 950 Pro	\$321.00	24	\$7,704.00
U.2 to M.2 Adapt.	Micro SATA Cables PN# CASE-994-U2	\$50.00	24	\$1,200.00

				\$19,092.00

The server can be purchased with processors and memory.

from MicroSataCable

<http://www.microsatacables.com/u2-sff8639-to-m2-pcie-if-with-25-inch-95mm-case-case-994-u2>

(or search for 994-u2 on Amazon).