

The Data Superconductor: Demonstrating 100Gb WAN Lustre Storage Using OpenFlow Enabled Traffic Engineering

Project Lead: Stephen C. Simms, Manager, High Performance File Systems, Research Technologies, Indiana University

Project Staff: Matthew Davy, Matthew Link, Robert Henschel, David Hancock, Kurt Seiffert (Indiana University)

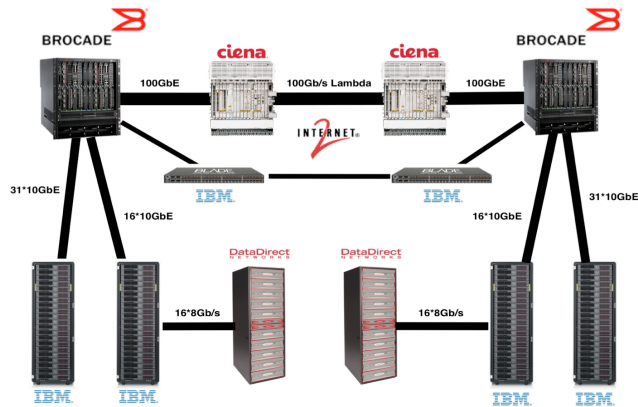


Figure 1. The above diagram shows the configuration that IU used to achieve data transfer results between Seattle and Indianapolis. IBM provided the 31 servers that functioned as compute nodes as well as providing the 16 Lustre storage servers that were attached to the Data Direct storage devices. Brocade and Ciena provided the 100 Gb routing equipment that permitted the data to flow between the cities.

One of the major concerns facing data-intensive research is the movement of massive data sets to and from supercomputing facilities. At SC11 using a demonstration of the Lustre file system mounted between Indianapolis, Indiana and Seattle, Washington, Indiana University set a record for fastest data transfer across a 100 Gigabit (Gb) network spanning thousands of miles. The 100Gb production network used was the first of its kind and permitted the IU researchers to push 6.5 gigabytes per second of data across 2,300 miles. To reframe this figure, you can think of the file system accepting over 43 DVDs worth of data in one minute.

This achievement shows that by using the open source Lustre file system and cutting edge network technology it is possible to transfer massive amounts of data very quickly. However, data transfer isn't the only possible benefit of this technology. During the demonstration IU was able to perform significant computation across distance utilizing a compute cluster and storage system separated by thousands of miles. IU was able to achieve bandwidth of 5.2 gigabytes per second while running scientific applications across the 100 Gigabit link. This means that it would be possible to achieve reasonable application performance across distance, permitting a single filesystem to bridge geographically distributed resources. This achievement has lasting implications for the fields of Climatology, Astrophysics, and Genome Analysis, to name a few.

This latest demonstration is a continuation of work involving the Lustre file system that IU began in 2006 along with the creation of the Data Capacitor, a large-capacity, high bandwidth data store for short- to mid-term storage of large research datasets (NSF MRI Award CNS-0521433). In 2007 a team led by IU won the SC07 bandwidth challenge, demonstrating that it was possible to saturate a 10Gb link by running multiple scientific applications across the wide area network, with data transferred to Reno, Nevada from as far as Dresden, Germany. In 2011 we see an order of magnitude increase in network capability and with IU's demonstration have shown the ability for researchers to utilize this new infrastructure.