

# Researchers transfer 65 terabytes of data in under just 100 minutes

December 4 2014

---

A team of researchers from Argonne National Laboratory and DataDirect Networks (DDN) moved 65 terabytes of data in under just 100 minutes at a recent supercomputing conference. Typically, two days are needed to move this volume of data between sites with a 10 Gbps connection.

With help from Ciena, Brocade, and ICAIR, the team sustained data transfer rates in excess of 85 Gbps—with peaks at over 90 Gbps—between [storage](#) systems in Ottawa, Canada, and New Orleans, LA, over a 100 Gbps wide-area network (WAN) connection. The demonstration took place on November 19, 2014 at SC14, the leading international conference for high performance computing, networking, storage and analysis.

This unprecedented achievement required combining the embedded file system and virtual machine capabilities of the DDN storage controller, the high-speed wide-area data transfer capabilities of the Globus GridFTP server, and an advanced 100G wide-area network.

"Embedding the GridFTP servers in virtual machines on DDN's storage controller eliminates the need for external data transfer nodes and network adapters," explained Raj Kettimuthu, principal software development specialist at Argonne National Laboratory. "We sustained a data transfer rate of 85 Gbps for over 60 minutes—and occasionally for as long as 90 minutes—several times during the SC14 conference."

Achieving 90+ Gbps for memory-to-memory transfers using a benchmarking tool like iperf is straightforward and has been demonstrated several times in the past. Achieving similar rates for disk-to-disk transfers however, presents a number of challenges, including choosing the appropriate block size that works well for both disk I/O and network I/O, and selecting the appropriate combination of parallel storage I/O threads and parallel TCP streams for optimal end-to-end performance.

Network experts often claim that storage is the bottleneck in the end-to-end transfers on high-speed networks, while storage experts claim that the network is often the bottleneck on transfers between sites with high-performance parallel file systems. "This demonstration was aimed at bringing together the experts and latest developments in all aspects concerning disk-to-disk WAN data movement, including [network](#), storage, and data movement tools," said Kettimuthu.

The team expects that the approach can be used to achieve 100+ Gbps wide-area [transfer](#) rates between [storage systems](#) using multiple WAN paths and additional storage resources in the end systems.

Provided by Argonne National Laboratory

Citation: Researchers transfer 65 terabytes of data in under just 100 minutes (2014, December 4) retrieved 3 July 2024 from <https://phys.org/news/2014-12-terabytes-minutes.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.