

The 4th International Workshop on
Network-aware Data Management

in cooperation with ACM [SIGHPC](#), in conjunction with [SC'14](#)
International Conference for High Performance Computing, Networking, Storage and Analysis

Sunday, November 16th, 2013 (1:30pm - 5:30pm) Room **274**
Ernest N. Morial Convention Center, New Orleans, Louisiana

Scope: Data sharing and resource coordination among distributed teams are becoming significant challenges every passing year. Networking is one of the most crucial components in the overall system architecture of a data centric environment. Many of the current solutions both in industry and scientific domains depend on the underlying network infrastructure and its performance. There is a need for efficient use of the networking middleware to address increasing data and compute requirements. Main scope of this workshop is to promote new collaborations between data management and networking communities to evaluate emerging trends and current technological developments, and to discuss future design principles of network-aware data management. We will seek contribution from academia, government, and industry to address current research and development efforts in remote data access mechanisms, end-to-end resource coordination, network virtualization, analysis and management frameworks, practical experiences, data-center networking, and performance problems in high-bandwidth networks.



<http://ndm-meeting.org>

Invited Talks:

Smart Cyber Infrastructure for Big Data Processing

Cees de Laat, University of Amsterdam

Abstract: The landscape of cyberinfrastructure for research is rapidly changing. There is a move towards virtualized and programmable infrastructure. The cloud paradigm enables applications to use computing resources at different places and optimize workflows in either bringing computing to the data or the other way around. Programmable networks allow networks to be utilized in unprecedented ways to create application specific Internets. This talk presents the latest developments in the Research and Education Networks to support Big Data sciences.

Bio: Professor de Laat chairs the System and Network Engineering (SNE) research group in Informatics Institute of the Faculty of Science at University of Amsterdam. Research in his group ranges from optical and switched networking and workflows for processing of big data in PetaScale e-Science applications, Semantic Web to describe e-infrastructure resources, information complexity, Authorization architectures and Systems Security & privacy of information in distributed environments. Prof. de Laat serves on the Lawrence Berkeley Laboratory Policy Board on matters regarding ESnet, is co-founder of the Global Lambda Integrated Facility (GLIF) and founding member of CineGrid.org. His group is/was part of EU projects SWITCH, CYCLONE, ENVRI, EuroBrazil, Geysers, NOVI, NEXTGRID, EGEE, and others. He is a member of the Advisory Board Internet Society Netherlands and Scientific technical advisory board of SURF Netherlands. A snapshot of his scientific career is available here: <http://delaat.net/>.

RDMA in the Cloud: Enabling high-bandwidth, low-latency communication in virtualized environments for HPC,

Josh Simons, VMware Inc. Office of the CTO

Abstract: While throughput and embarrassingly parallel HPC applications run with very little (<5%) performance degradation in modern virtualized environments, very latency sensitive applications present more of a challenge due to I/O overheads introduced by virtualization. Consequently, enabling use of RDMA-capable interconnects like InfiniBand and RoCE is an important aspect of broadening acceptance and use of cloud computing in HPC. This talk will present our latest performance results using InfiniBand and RoCE in passthrough mode as well as our early experiences with InfiniBand SR-IOV, an important feature that enables multi-VM access to RDMA-connected parallel file systems like Lustre and GPFS.

Bio: Josh Simons currently leads an effort from the Office of the CTO at VMware to bring the value of virtualization to HPC. Previously, he was a Distinguished Engineer at Sun Microsystems working on HPC and prior to Sun he worked at Thinking Machines Corporation, a supercomputer manufacturer. Josh has worked on tools for parallel computing, including language and compiler design and parallel debugger design and development. He has a degree in Engineering from Harvard College and a Masters in Computer Science from Harvard University. He is currently serving as Chairman of the OpenMP Board of Directors.

The 4th International Workshop on

The 4th International Workshop on Network-aware Data Management

Sunday, November 16th, 2014 (1:30pm - 5:30pm)

Room **274**



<http://ndm-meeting.org>



Workshop Program

13:30 - 13:40 Opening Remarks

14:40 - 14:15 *Keynote I (35min)*

Smart Cyber Infrastructure for Big Data Processing, Cees de Laat (University of Amsterdam)

14:15 - 15:05 *Paper Session I (25min each)*

Adaptation and Policy-Based Resource Allocation for Efficient Bulk Data Transfers in High Performance Computing Environments, Ann L. Chervenak (USC/ISI)

Analysis of the Effect of Core Affinity on High-Throughput, Nathan Hanford (UC Davis)

15:05 - 15:30 Refreshment Break

15:30 - 16:05 *Keynote II (35min)*

RDMA in the Cloud: Enabling high-bandwidth, low-latency communication in virtualized environments for HPC, Josh Simons (VMware Inc. Office of the CTO)

16:05 - 17:20 *Paper Session II (25min each)*

Flexible Scheduling and Control of Bandwidth and In-transit Services for End-to-End Application Workflows, Manish Parashar (Rutgers University)

Towards Energy Awareness in Hadoop, Krish K.R. (Virginia Tech)

Towards Managed Terabit/s Scientific Data Flows, Artur Barczyk (Caltech)

17:20 - 17:30 Closing Remarks