

NDM 2014

Network-aware Data Management Workshop

Room 274



1:30 pm – 5:30 pm

November 16th, 2014 - Sunday

Welcome to NDM



- The Fourth Workshop on Network-aware Data Management

<http://ndm-meeting.org>

Mehmet Balman,
Surendra Byna,
Brian L. Tierney



Scope of the workshop

- Discuss emerging trends in use of networking for data management
- Create new collaborations between network and data management communities

Agenda

1:30 - 1:40 Opening Remarks

1:40 - 2:15 Smart Cyber Infrastructure for Big Data Processing
Prof. Cees de Laat (University of Amsterdam)

2:15 - 3:05 **Paper Session I** (25min each - 2 papers)

3:05 - 3:30 *Refreshment Break*

3:30 - 4:05 RDMA in the Cloud: Enabling high-bandwidth, low-latency communication in virtualized environments for HPC,
Josh Simons (VMware Inc. Office of the CTO)

4:05 - 5:20 **Paper Session II** (25min each - 3 papers)

5:20 - 5:30 Closing Remarks and Open Discussion

Technical Papers

- **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 Flows**, Nathan Hanford (UC Davis)
- **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)

Program Committee

- **Ilya Baldin**, RENCI/UNC Chapel Hill, USA
- **Amitabha Banerjee**, VMware Inc., USA
- **Jerry Chou**, National Tsing Hua University, Taiwan
- **Zhihui Du**, Tsinghua University, China
- **Renato Figueiredo**, University of Florida, USA
- **Dipak Ghosal**, University of California, Davis, USA
- **Zhiyi Huang**, University of Otago, New Zealand
- **Alexandru Iosup**, Delft University of Technology, The Netherlands
- **Raj Kettimuthu**, ANLand University of Chicago, USA
- **Jinoh Kim**, Texas A&M University-Commerce, USA
- **Siva Kulasekaran**, Texas Advance Computing Center, USA
- **Malathi Veeraraghavan**, University of Virginia, USA
- **Chen Wu**, ICRAR, The University of Western Australia, Australia
- **Wenji Wu**, Fermilab, USA
- **Lei Xia**, VMware Inc., USA
- **Esma Yildirim**, Fatih University, Turkey

Keynote I:

Smart Cyber Infrastructure for Big Data Processing, Prof. 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.

Keynote II:**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.