

NASA / IEEE MSST2004
Twelfth NASA Goddard / Twenty-First IEEE
Conference on Mass Storage Systems and Technologies

*Long-Term Stewardship of
Globally-Distributed
Storage*



April 13 – 16, 2004
The Inn and Conference Center
University of Maryland University College
Adelphi, Maryland, USA

*Final
Program*

Sponsored by:



<http://storageconference.org/2004>



Long-Term Stewardship of Globally-Distributed Storage



The NASA Goddard Space Flight Center will hold its 12th Conference on Mass Storage Systems and Technologies in cooperation with the IEEE Mass Storage Systems Technical Committee's 21st Conference. The conference, April 13-16, 2004 in Adelphi, Maryland, combines technical papers, posters, and a vendor exposition in a unique forum tailored for users, designers, managers, and vendors of large-scale storage systems and technologies.

With the deployment of large, globally-distributed mass storage systems, systems scaling to thousands of sites, each holding billions of objects, including terabyte or larger objects, will soon be common. As solutions are found that provide global access and delivery, problems of long-term stewardship are exacerbated. The conference will examine solutions for these problems, with papers and presentations addressing preservation and stewardship, performance, infrastructure resources, and file systems. A vendor session provides information on commercially available systems. Parallel tutorials will be given on Tuesday, April 13th. A Round Table on Emerging Technologies, a hallmark of these conferences, will include a panel of experts discussing promising new technologies. A short-paper session will be held Wednesday afternoon, April 14th. All sessions are informal with audience participation encouraged.

A Vendor Expo will be held Wednesday through Friday, offering additional opportunities for informal information exchange. Additional details can be obtained on the web at:

<http://storageconference.org/2004>

Agenda

Tuesday, April 13

Parallel Tutorials

8:30 am – 12:00 pm

Chair, Alan Montgomery

Data Grid Management Systems
Arun Jagatheesan, Reagan Moore,
University of California, San Diego

Chair, Jim Hughes

Trials and Tribulations of Shared File Systems in a Heterogeneous Environment
Henry Newman, Nathan Schumann, *Instrumental*

12:00 pm – 1:00 pm – Lunch

1:00 pm – 3:00 pm

Chair, Ahmed Amer

Fibre Channel and IP SAN Integration
Henry Yang, *McData*

Chair, Curtis Anderson

Long-Term Stewardship of Globally-Distributed Representation Information
David Holdsworth, Paul Wheatley, *Leeds University*

3:30 pm – 5:30 pm

Chair, Ahmed Amer

Emerging Data Replication Technology
R. F. Kern, V. T. Peltz,
IBM Systems Group

Chair, Curtis Anderson

Challenges in Long-Term Data Stewardship
Ruth Duerr, *University of Colorado*

5:30 pm – 8:00 pm – Welcoming Reception

Wednesday, April 14

12:00 pm – 7:00 pm – Vendor Expo

9:00 am – 10:30 am – Welcome and Keynote Address

Chair, Ben Kobler

Perfect Devices: The Amazing Endurance of Evolving Hard Disk Drives
Dr. Giora J Tarnopolsky, *Tarnotek*

11:00 am – 12:30 pm – Long Term Preservation, Stewardship

Chair, Jack Cole

NARA's Electronic Records Archive (ERA) - The Electronic Records Challenge
Mark Huber, *American Systems*, Alla Lake, *Linfo Systems*, Robert Chadduck, *NARA*

Preservation Environments

Reagan Moore, *San Diego Supercomputer Center*

Data Management as a Cluster Middleware Centerpiece

Jose Zero, *Halcyon Systems*, Dr. Daniel Duffy, *Computer Sciences*

12:30 pm – 2:00 pm – Lunch

2:00 pm – 3:30 pm – Performance

Chair, Randal Burns

Regulating I/O Performance of Shared Storage with a Control Theoretical Approach

Han Deok Lee, Young Jin Nam, Chanik Park,
Pohang University, Republic of Korea

SAN and Data Transport Technology Evaluation at the NASA Goddard Space Flight Center

Hoot Thompson, *Patuxent Technology Partners*

File System Workload Analysis For Large Scientific Computing Applications

Feng Wang, Qin Xin, Bo Hong, Scott A. Brandt, Ethan L. Miller, Darrell D. E. Long, *University of California, Santa Cruz*, Tyce T. McLarty, *Lawrence Livermore National Laboratory*

Wednesday (continued)

4:00 pm – 5:00 pm – Short Papers

Chair, Fynnette Eaton, NARA

V:Drive - Costs and Benefits of an Out-of-Band Storage Virtualization System

André Brinkmann, Michael Heidebuer, Friedhelm Meyer auf der Heide, Ulrich Rückert, Kay Salzwedel, Mario Vodisek, Paderborn University, Germany

Identifying Stable File Access Patterns

Purvi Shah, Jehan-François Pâris, University of Houston
Ahmed Amer, University of Pittsburgh
Darrell D. E. Long, University of California, Santa Cruz

An On-line Backup Function for a Clustered NAS System (X-NAS)

Yoshiko Yasuda, Shinichi Kawamoto, Atsushi Ebata, Jun Okitsu, Tatsuo Higuchi, Hitachi, Japan

dCache, the Commodity Cache

Patrick Fuhrmann, DESY, Germany

Hierarchical Storage Management at the NASA Center for Computational Sciences: From Unirete to SAM-QFS

Ellen Salmon, Adina Tarshish, Sanjay Patel, Marty Saletta, Mike Rouch, Robert Caine, Jeff Paffel, Lisa Burns, Ed. Vanderlan, Nancy Palm, Dr. Daniel Duffy, NASA Goddard Space Flight Center

Parity Redundancy Strategies in a Large Scale Distributed Storage System

John Chandy, University of Connecticut

Reducing Storage Management Costs via Informed User-Based Policies

Erez Zadok, Jeffrey Osborn, Ariye Shater, Charles Wright, Kiran-Kumar Muniswamy-Reddy, Stony Brook University
Jason Nieh, Columbia University

A Design of Metadata Server Cluster in Large Distributed Object-based Storage

Jie Yan, Yao-Long Zhu, Hui Xiong, Renuga Kanagavelu, Feng Zhou, So Lih Weon, Data Storage Institute, Singapore

An iSCSI Design and Implementation

Hui Xiong, Renuga Kanagavelu, Yaolong Zhu, Khai Leong Yong, Data Storage Institute, Singapore

Quanta Data Storage: A Cross Layer Architecture for the Storage Networks

Prabhanjan Gurumohan, Sai S. B. Narasimhamurthy, Joseph Y. Hui, Arizona State University

Rebuild Strategies for Redundant Disk Arrays

Alexander Thomasian, Gang Fu, Chunqi Han, New Jersey Institute of Technology

Evaluation of Efficient Archival Storage Techniques

Lawrence You, University of California, Santa Cruz
Christos Karamanolis, Hewlett-Packard Labs

An Efficient Data Sharing Scheme for iSCSI-Based File Systems

Dingshan He, David H. C. Du, University of Minnesota

Using DataSpace Archives to Support Long-Term Stewardship of Remote and Distributed Data

Robert L. Grossman, Dave Hanley, Xinwei Hong, Parthasarathy Krishnaswamy, University of Illinois at Chicago

5:00 pm – 8:00 pm – Reception and Poster Session

Thursday, April 15

10:00 am – 5:00 pm – Vendor Expo

9:00 am – 10:30 am – Infrastructure Resources

Chair, Steve Louis

Promote-IT: An Efficient Real-Time Tertiary-Storage Scheduler

Maria Eva Lijding, Sape Mullender, Pierre Jansen, University of Twente, The Netherlands

The Data Services Archive

Rena Haynes, Wilbur R. Johnson, Sandia National Laboratories

Interconnection Architectures for High-Performance Object-Based File Systems

Andy Hospodor, Ethan L. Miller, University of California, Santa Cruz

11:00 am – 12:30 pm – Round Table on Emerging Technologies

Chair, P. C. Hariharan

Panel Moderator, Mark Ferelli

Panel

Dr. Kevin Daly, Avamar

Professor Garth Gibson, Carnegie Mellon University and Panasas

Dr. Alope Guha, COPAN Systems

R. B. Hooks, StorageTek

Professor Kai Li, Princeton University and Data Domain

12:30 pm – 2:00 pm – Lunch

2:00 pm – 3:30 pm – File Systems

Chair, Curtis Anderson

OBFS: A File System for Object-Based Storage Devices

Feng Wang, Scott A. Brandt, Ethan L. Miller, Darrell D. E. Long, University of California, Santa Cruz

Duplicate Data Elimination in a SAN File System

Bo Hong, Darrell D. E. Long, University of California, Santa Cruz
Demyan Plantenberg, Miriam Sivan-Zimet, IBM Almaden Research Center

Clotho: Transparent Data Versioning at the Block I/O Level

Michail Flouris, University of Toronto

Angelos Bilas, Foundation for Research and Technology, Greece

4:00 pm – 5:00 pm – Work in Progress

Chair, Ethan Miller

Utilization of Storage Resource Managers by Compact Muon Solenoid Large Hadron Collider Experiment

Donald Petravick, Timur Perelmutov, Fermi Nat. Accelerator Lab.

Shared Rapid File System on Ethernet

Naoyuki Fujita, Japan Aerospace Exploration Agency

Virtualized RAID for High Performance Shared File Systems

Randy C. Seamans, 3PARdata

Computational Methods of Signal Decoding for Recovering Recorded Data

Glenn Sakaguchi, Laslo Orto, ActionFront Data Recovery Labs,
Chuck Sobey, Channel Science

Low-level Simulation of Back-end Storage Networking Using HASE

Tim Courtney, Xyratex Technology, Franck Chevalier, Univ. of Edinburgh

Storage Grid and Distributed Block Systems

Geoff Hayward, YottaYotta

5:00 pm – 6:00 pm – Extemporaneous

Chair, Jim Hughes

A forum for short presentations of last-minute thoughts and ideas. Presenters will sign up for speaking slots during the conference.

6:00 pm - Dinner

7:00 pm - After Dinner Speaker

Pharmacogenomics and Supercomputing

John Reynders, VP Informatics, Celera Genomics

Friday, April 16

10:00 am – 2:00 pm – Vendor Expo

9:00 am – 10:30 am – Site Reports

Chair, Gene Harano

U.S. National Oceanographic Data Center Archival Practices and the Open Archives Information System Reference Model

Donald Collins, *National Oceanographic Data Center*

Storage Resource Sharing with CASTOR

Olof Barring, Ben Couturier, Jean-Damien Durand, Emil Knezo, Sebastien Ponce, *CERN, Switzerland*

GUPFS: The Global Unified Parallel File System Project at NERSC

Greg Butler, Rei Lee, Mike Welcome, *Lawrence Berkeley National Laboratory*

11:00 am – 12:30 pm – iSCSI and SAN

Chair, Jean-Jacques Bedet

SANSim - A Platform for Simulation and Design of a Storage Area Network

Yao-Long Zhu, Chao-Yang Wang, Wei-Ya Xi, Feng Zhou, *Data Storage Institute, Singapore*

Cost-Effective Remote Mirroring Using the iSCSI Protocol

Ming Zhang, Yinan Liu, and Qing (Ken) Yang, *University of Rhode Island*

Simulation Study of iSCSI-based Storage System

Yingping Lu, Farrukh Noman, David H.C. Du, *University of Minnesota*

12:30 pm – 1:30 pm – Lunch

1:30 pm – 3:00 pm – Vendor Solutions

Chair, Bruce Rosen

Comparative Performance Evaluation of iSCSI Protocol over Metro, Local, and Wide Area Networks

Ismail Dalgic, Kadir Ozdemir, Rajkumar Velpuri, Jason Weber, *Intransa, Umesh Kukreja, Atrica, Helen Chen, Sandia Nat. Lab.*

H-RAIN: An Architecture for Future-Proofing Digital Archives

Andres Rodriguez, Jack Orenstein, *Archivas*

A New Approach to Disk-Based Scalable Mass Storage System

Aloke Guha, *COPAN Systems*

Multi-Tiered Storage

Gil VanSchoor, *BlueArc*


Managing Scalability in Object Storage Systems for HPC Linux Clusters

Brent Welch, Garth Gibson, *Panasas*

The Evolution of a Distributed Storage System

Norman Margolus, *Permanit*

3:00 pm – Adjourn



The Baltimore Oriole is the State Bird of Maryland.

People

Conference and Program Committee Chair

Ben Kobler, *NASA Goddard Space Flight Center*

Keynote and Invited Papers Chair

P C Hariharan, *SES*

Tutorial Chair

Jim Hughes, *StorageTek*

Vendor Expo Chair

Gary Sobol, *StorageTek*

Publications Chairs

P C Hariharan, *SES*

Nabil Adam, *Rutgers University*

Work In Progress Chairs

Ethan Miller, *University of California, Santa Cruz*

Randal Burns, *Johns Hopkins University*

Publicity Chairs

Jack Cole, *US Army Research Laboratory*

Sam Coleman, *Lawrence Livermore National Laboratory, Retired*

IEEE Computer Society Liaison

Merritt Jones, *MITRE*

Program Committee

Ahmed Amer, *University of Pittsburgh*

Curtis Anderson, *Mendocino Software*

Jean-Jacques Bedet, *SSAI*

John Berbert, *NASA Goddard Space Flight Center*

Randal Burns, *Johns Hopkins University*

Robert Chadduck, *NARA*

Jack Cole, *US Army Research Laboratory*

Bob Coyne, *IBM*

Jim Finlayson, *Department of Defense*

Dirk Grunwald, *University of Colorado*

Bruce K. Haddon, *Sun Microsystems*

Gene Harano, *National Center for Atmospheric Research*

P C Hariharan, *SES*

Jim Hughes, *StorageTek*

John Jensen, *NOAA*

Merritt Jones, *MITRE*

Ben Kobler, *NASA Goddard Space Flight Center*

Steve Louis, *Lawrence Livermore National Laboratory*

Ethan Miller, *University of California, Santa Cruz*

Alan Montgomery, *Department of Defense*

Reagan Moore, *San Diego Supercomputer Center*

Bruce Rosen, *NIST*

Paul Rutherford, *ADIC*

Tom Ruwart, *University of Minnesota*

Julian Satran, *IBM Haifa Research Laboratory, Israel*

Donald Sawyer, *NASA Goddard Space Flight Center*

Rodney Van Meter, *Keio University, Japan*

Hotel Information

The Inn and Conference Center

University of Maryland University College

3501 University Boulevard East

Adelphi, Maryland, 20783

The room rate is \$129 per day plus tax; parking is available at the conference center parking lot at no additional fee.

Please make your reservations directly with the hotel by March 22, 2004. For reservations, call:

+1-301-985-7300

Give the group name "NASA Mass Storage Systems and Technologies" to receive the conference rate.