

Thursday, April 15, 2004

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*
Demyn 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, 2004

8:00 am – 11:00 am – Continental Breakfast
8:00 am – 12:00 pm – Conference Registration
10:00 pm – 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 – Conference Adjourns



*Long-Term Stewardship
of Globally-Distributed
Storage*

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

The Inn and Conference Center
University of Maryland
Adelphi, Maryland, USA

April 13 – 16, 2004



The Baltimore Oriole is the
State Bird of Maryland.

Tuesday, April 13, 2004

Parallel Tutorials

8:30 am - 12:00 pm

Chair, Alan Montgomery

Data Grid Management Systems

Arun Jagatheesan,
Reagan Moore, *UCSD*

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 and Conference Registration

Wednesday, April 14, 2004

8:00 am - 11:00 am - Continental Breakfast
8:00 am - 4:00 pm - Conference Registration
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*

4:00 pm - 5:00 pm — Short Papers

Chair, Fynette 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 Unitree 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, 2004

8:00 am - 11:00 am - Continental Breakfast
8:00 am - 4:00 pm - Conference Registration
10:00 pm - 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 Ferrelli

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*