

IEEE / NASA MSST2003
Twentieth IEEE / Eleventh NASA Goddard
Conference on Mass Storage Systems and Technologies

Global Access to Distributed Storage

April 7-10, 2003
Paradise Point Resort
San Diego, California, USA

Sponsored by:



<http://storageconference.org/2003>

Conference Overview

The IEEE Mass Storage Systems Technical Committee will hold its 20th Conference in cooperation with the 11th NASA Goddard Conference April 7–10, 2003 in San Diego, California. The conference combines technical papers, posters and a vendor exposition in a unique forum tailored for managers, designers, users and vendors of large-scale storage systems and technologies. The conference reflects the trend toward globalization of storage management with a focus on data management, storage systems, media, long term retention, and global access to distributed data.

Papers and presentations address requirements for data grids, iSCSI access to remote data, Storage Area Networks, Network Attached Storage, file systems, and storage architectures. One session is devoted to experiences from sites running very large archives (hundreds of terabytes to petabytes). A vendor session provides information on commercially available systems. Individual papers cover topics ranging from new approaches toward the creation of mass storage systems based on data grid technology, design and performance of the iSCSI protocol, reliability mechanisms for very large storage, and others. Tutorials will be given on Monday, April 7th.

A Round Table on Emerging Technologies, a hallmark of these Symposia, chaired by Dick Watson and Paul Rutherford, will include a panel of experts discussing promising new technologies. All sessions are informal with audience participation encouraged. A Short Paper Session will be held Tuesday evening, April 8th, and a Vendor Expo will be held Tuesday through Thursday, offering additional opportunities for informal information exchange. Additional details can be obtained on the web at:

<http://storageconference.org/2003>

We look forward to seeing you in San Diego.

Conference Chairs:

Ben Kobler (*NASA Goddard Space Flight Center*)
Jack Cole (*US Army Research Laboratory*)

Program Committee Chairs:

Reagan Moore (*SDSC*)
Tom Ruwart (*I/O Performance*)

Vendor Chair:

Gary Sobol

IEEE Computer Society Sponsor:

Merritt Jones (*MITRE*)



Conference Agenda

Monday, April 7

Tutorial

Tutorial Chair: James Hughes

7:00 am – 8:30 am – Breakfast, Sunset Ballroom Deck

8:30 am – 10:00 am – LAN, SAN, MAN, WAN - Making an Intelligent Choice for your Storage
Silvano Gai (*Andiamo*)

10:00 am – 10:30 am – Break, Sunset Ballroom Deck

10:30 am – 12:00 pm – Management Interoperability and Extensibility in Storage Area Networks
Roger Reich (*VERITAS Software*)

12:00 pm – 1:30 pm – Lunch, Paradise Lawn

1:30 pm – 3:00 pm – Encrypted Storage
James Hughes (*StorageTek*)

3:00 pm – 3:30 pm – Break, Sunset Ballroom Deck

3:30 pm – 5:00 pm – Shared Storage Systems
Paul Rutherford (*ADIC*)

5:00 pm – 8:00 pm – Welcoming Reception, Dockside and Bayview Rooms

Tuesday, April 8

12:00 pm – 8:00 pm – Vendor Expo

7:00 am – 8:30 am – Breakfast, Sunset Ballroom Deck

8:30 am – 10:00 am – Welcome & Keynote
Chair: Ben Kobler

The New (Old) Challenges for Technology Infrastructure Today
Greg Bishop (*America Online*)

10:00 am – 10:30 am – Break, Sunset Ballroom Deck

10:30 am – 12:00 pm – Grids
Chair: Jack Cole

Data Grids, Collections, and Grid Bricks
Arcot Rajasekar, Michael Wan, Reagan W. Moore,
George Kremenek, Tom Guptill
(*San Diego Supercomputer Center*)

Correct Modeling of Cache Replacement Policies in a Data-Grid
Ekow Otoo, Arie Shoshani
(*Lawrence Berkeley National Laboratory*)

A Centralized Data Access Model for Grid Computing
Phil Andrews, Tom Sherwin, Bryan Bannister
(*San Diego Supercomputer Center*)

12:00 pm – 1:30 pm – Lunch, Paradise Lawn

1:30 pm – 3:00 pm – Site Reports
Chair: Gene Harano

Building Cost-Effective Remote Data Storage Capabilities for NASA's EOSDIS
Stephen Marley (*Raytheon*)
Mike Moore (*NASA GSFC*)
Bruce Clark (*EDS*)

Wednesday, April 9

8:00 am – 5:00 pm – Vendor Expo

7:00 am – 8:30 am – Breakfast, Sunset Ballroom Deck

8:30 am – 10:00 am – iSCSI and SAN

Chair: P. C. Hariharan

Design of the iSCSI Protocol

Kalman Z. Meth, Julian Satran (*IBM Haifa Research Lab*)

A Performance Analysis of the iSCSI Protocol

Stephen Aiken, Dirk Grunwald, Andy Pleszkun
(*University of Colorado*)

Architectural Considerations and Performance Evaluations Of Shared Storage Area Networks at NASA Goddard Space Flight Center

Hoot Thompson (*Patuxent Technology Partners*)
Curt Tilmes, Robert Cavey, Bill Fink, Paul Lang, Ben Kobler (*NASA GSFC*)

10:00 am – 10:30 am – Break, Sunset Ballroom Deck

10:30 am – 12:00 pm – Panel

Co-chairs: Richard Watson, Paul Rutherford

Emerging Object/Active-Storage Technologies

Representatives from user, vendor, and research organizations will discuss the many critical challenges in achieving location-transparent, secure, and high-performance access to distributed storage resources.

Reagan Moore (*SDSC*) – GRID Perspective

Gary Grider (*LANL*) - ASCII user Perspective

Garth Gibson (*CMU, Panassas*) - Academic and Panasas Perspectives

Ethan Miller (*UCSC*) - Academic Perspective

Julian Satran (*IBM*) - IBM and Standards Perspectives

David Black (*EMC*) - EMC Perspective

Dave Anderson (*Seagate*) - Seagate and Standards Perspectives

12:00 pm – 1:30 pm – Lunch, Paradise Lawn

1:30 pm – 3:00 pm – Reliability

Chair: Jean-Jacques Bedet

Reliability Mechanisms for Very Large Storage Systems

Qin Xin, Ethan L. Miller, Darrell D. E. Long, Scott A. Brandt (*UCSC*)

Thomas Schwarz (*Santa Clara University*)

Witold Litwin (*Université Paris 9 Dauphine*)

Software-based Erasure Codes for Scalable Distributed Storage

Joseph A. Cooley, Jeremy L. Mineweaser, Leslie D. Servi,

Eushuan T. Tsung (*MIT Lincoln Laboratory*)

Towards an Object Store

Alain Azagury, Vladimir Dreizin, Michael Factor, Ealan Henis,

Dalit Naor, Noam Rinetzky, Ohad Rodeh, Julian Satran,

Ami Tavori, Lena Yerushalmi (*IBM Haifa Research Laboratories*)

3:00 pm – 3:30 pm – Break, Sunset Ballroom Deck

3:30 pm – 4:45 pm – Vendor Solutions

Chair: Ben Kobler

Design and Implementation of a Storage Repository Using Commonality Factoring

Eric Olsen, Jim Hamilton (*Avamar Technologies*)

Design and Implementation of Block Storage Multi-Protocol Converter

Irina Gerasimov, Alexey Zhuravlev, Mikhail Pershin,

Dennis V. Gerasimov (*TechnoMages*)

IP SAN – From iSCSI to IP-addressable Ethernet Disks

Peter Wang, Robert Gilligan, Henry Green, Jeff Raubitschek
(*Intransa*)

Enabling Advanced Data Management with Sun StorEdge™ QFS

Harriet Coverston, (*Sun Microsystems*)

4:45 pm – 5:45 pm – Extemporaneous Talks

Chair: James Hughes

Archive Management: The Missing Component

Howard J. Diamond, John J. Bates, David M. Clark,
and Robert L. Mairs
(*National Environmental Satellite, Data, and Information Service,
National Oceanic and Atmospheric Administration*)

An overview of a large-scale data migration

M. Lübeck, A. Valassi
(*European Organization for Nuclear Research, CERN*)

3:00 pm – 3:30 pm – Break, Sunset Ballroom Deck

3:30 pm – 4:30 pm – Short Papers

Chair: Robert Chadduck

Towards Optimal I/O Scheduling for MEMS-based Storage

Hailing Yu, Divy Agrawal, Amr El Abbadi
(*Department of Computer Science
University of California, Santa Barbara*)

A Case for the Global Access to Large Distributed Data Sets using Data Webs Employing Photonic Data Services

Robert L. Grossman, Yunlong Gu, Dave Hanley,
Xinwei Hong, Jorge Levera, Marco Mazzucco
(*Laboratory for Advanced Computing, University of Illinois at Chicago*)

Design and Implementation of Multiple Addresses Parallel Transmission Architecture for Storage Area Network

Bin Meng, Patrick B. T. Khoo, T. C. Chong
(*MCSA Group, NST Division - Data Storage Institute
Affiliated with the National University of Singapore*)

Storage Devices, Local File System and Crossbar Network File System Characteristics, and 1 Terabyte File IO Benchmark on "Numerical Simulator III"

Naoyuki Fujita, Hirofumi Ookawa
(*National Aerospace Laboratory of Japan*)

Views, Objects, and Persistence for accessing a High Volume Global Data Set

Richard T. Baldwin
(*US DOC NOAA National Climatic Data Center*)

NAS Switch: A Novel CIFS Server Virtualization

Wataru Katsurashima, Satoshi Yamakawa, Takashi Torii,
Jun Ishikawa, and Yoshihide Kikuchi
(*Internet Systems Research Laboratories, NEC Corporation*)

NSM: A Distributed Storage Architecture for Data-Intensive Applications

Zeyad Ali, Qutaibah Malluhi
(*Jackson State University*)

Implementing and Evaluating Jukebox Schedulers Using JukeTools

Maria Eva Lijding, Sape Mullender, and Pierre Jansen
(*University of Twente*)

Media Stability and Life Expectancies of Magnetic Tape for Use with IBM 3590 and Digital Linear Tape Systems

J. S. Judge, R. G. Schmidt, R. D. Weiss, (*Arkival Technology*),
G. Mille, (*University Of New Hampshire*)

The Fermilab Data Storage Infrastructure

Jon Bakken, Eileen Berman, Chih-Hao Huang, Alexander
Moibenko, Donald Petravick, Michael Zalokar
(*Fermi National Accelerator Laboratory*)

High Bandwidth Scientific Data Management using Storage Area Networking: Lessons Learned at the Starfire Optical Range

Terry S. Duncan
(*Kirtland Air Force Base*)

NCDC the "One Stop Shop" for all NEXRAD

WSR-88D Level II Data Services
Stephen A. Del Greco
(*National Climatic Data Center*)

5:00 pm – 8:00 pm – Reception and Poster Session,
Dockside and Bayview Rooms

Thursday, April 10

8:00 am – 2:00 pm – Vendor Expo

7:00 am – 8:30 am – Breakfast, Sunset Ballroom Deck

8:30 am – 10:00 am – Network-Attached Storage

Chair: Ethan Miller

A Scalable Architecture for Clustered Network Attached Storage

Jonathan Bright (*Bright Consulting Group*)
John A. Chand, (*University of Connecticut*)

zFS - A Scalable distributed File System using Object Disks

O. Rodeh, U. Schonfeld, A. Teperman
(*IBM Labs, Haifa University*)

The Concept and Evaluation of X-NAS: a Highly Scalable NAS System

Yoshiko Yasuda, Shinichi Kawamoto, Atsushi Ebata,
Jun Okitsu, Tatsuo Higuchi
(*Hitachi, Ltd.*)

10:00 am – 10:30 am – Break, Sunset Ballroom Deck

10:30 am – 12:00 pm – File Systems

Chair: Curtis Anderson

**Using Multiple Predictors to Improve the Accuracy
of File Access Predictions**

Gary A. S. Whittle, Jehan-François Pâris (*University of Houston*)
Ahmed Amer (*University of Pittsburgh*)
Darrell D. E. Long (*University of California, Santa Cruz*)
Randal Burns (*Johns Hopkins University*)

Peabody: The Time Traveling Disk

Charles B. Morrey III, Dirk Grunwald
(*Dept. of Computer Science,
University of Colorado*)

**SPIRAL: A Client-Transparent Third-Party Transfer
Scheme for Network Attached Disks**

Xiaonan Ma
(*IBM*)

12:00 pm – 1:30 pm – Lunch, Paradise Lawn

1:30 pm – 3:00 pm – Storage Architectures

Chair: Merritt Jones

**Effective Management of Hierarchical Storage
Using Two Levels of Data Clustering**

Ratko Orlandic
(*Department of Computer Science,
Illinois Institute of Technology*)

A Simple Mass Storage System for the SRB Data Grid

Michael Wan, Arcot Rajasekar, Phil Andrews, Reagan Moore
(*San Diego Supercomputer Center*)

**Efficient Metadata Management in Large
Distributed Storage Systems**

Scott A. Brandt, Lan Xue, Ethan L. Miller, Darrell D. E. Long
(*Storage Systems Research Center,
University of California, Santa Cruz*)

Program Committee

Nabil Adam (*Rutgers University*)

Curtis Anderson (*Universal Network Machines*)

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*)

Sam Coleman (*LLNL, Retired*)

Bob Coyne (*IBM*)

Bruce K. Haddon (*Sun Microsystems*)

Gene Harano (*NCAR*)

P C Hariharan (*SES*)

Jim Hughes (*StorageTek*)

Merritt Jones (*MITRE*)

Ben Kobler (*NASA Goddard Space Flight Center*)

Ethan Miller (*University of California, Santa Cruz*)

Matt O'Keefe (*Sistina Software*)

Paul Rutherford (*ADIC*)

Donald Sawyer (*NASA Goddard Space Flight Center*)

Jamie Shiers (*CERN*)

Rodney Van Meter (*Nokia*)

Hotel and Logistics Information

Paradise Point Resort, San Diego April 7-10, 2003

The Twentieth IEEE / Eleventh NASA Goddard Conference on Mass Storage Systems and Technologies will be held at the Paradise Point Resort in San Diego, California, April 7 - 10, 2003.

Rates/Reservations for the Paradise Point Resort

The group rate is \$165.00/day plus tax and \$10.00/day for parking. Please make your reservations directly with the hotel by March 7, 2003:

Paradise Point Resort
1404 West Vacation Road
San Diego, CA 92109-7905
+1-858-274-4630
+1-800-344-2626
<http://www.paradisepoint.com>

Give the group name of "IEEE Computer Society MSS" to receive the conference rate. Cancellations must be 48 hours from the date of arrival.