

Proceedings

**24th IEEE Conference on Mass  
Storage Systems and Technologies**  
***(MSST 2007)***

**San Diego, California, USA**

**September 24–27, 2007**

**Sponsored by**

IEEE Computer Society Mass Storage Systems Technical Committee  
(MSSTC)



Los Alamitos, California  
Washington • Tokyo



# Contents

## Proceedings of the 24th IEEE Conference on Mass Storage Systems and Technologies

---

**MSST 2007**

Message from the Chairs .....	viii
Conference and Program Committees .....	ix

---

### Architectures

Preservation DataStores: Architecture for Preservation Aware Storage.....	3
<i>Michael Factor, Dalit Naor, Simona Rabinovici-Cohen, Leeat Ramati, Petra Reshef, Julian Satran, and David L. Giarretta</i>	
GreenStor: Application-Aided Energy-Efficient Storage .....	16
<i>Nagapramod Mandagere, Jim Diehl, and David Du</i>	
Modeling the Impact of Checkpoints on Next-Generation Systems.....	30
<i>Ron A. Oldfield, Sarala Arunagiri, Patricia J. Teller, Seetharami Seelam, Maria Ruiz Varela, Rolf Riesen, and Philip C. Roth</i>	

### Data Grids

Storage Resource Managers: Recent International Experience on Requirements and Multiple Co-Operating Implementations.....	47
<i>Lana Abadie, Paolo Badino, Jean-Philippe Baud, Ezio Corso, Matt Crawford, Shaun De Witt, Flavia Donno, Alberto Forti, Ákos Frohner, Patrick Fuhrmann, Gilbert Grosdidier, Junmin Gu, Jens Jensen, Birger Koblitz, Sophie Lemaitre, Maarten Litmaath, Dmitry Litvinsev, Giuseppe Lo Presti, Luca Magnoni, Tigran Mkrtchan, Alexander Moibenko, Rémi Mollon, Vijaya Natarajan, Gene Oleynik, Timur Perelmutov, Don Petravick, Arie Shoshani, Alex Sim, David Smith, Massimo Sponza, Paola Tedesco, and Riccardo Zappi</i>	
Grid-Enabled Standards-Based Data Management.....	60
<i>Lana Abadie, Paolo Badino, Jean-Philippe Baud, James Casey, Ákos Frohner, Gilbert Grosdidier, Sophie Lemaitre, Gavin Mccance, Rémi Mollon, Krzysztof Nienartowicz, David Smith, and Paolo Tedesco</i>	
Quota Enforcement for High-Performance Distributed Storage Systems.....	72
<i>Kristal T. Pollack, Darrell D. E. Long, Richard A. Golding, Ralph A. Becker-Szendy, and Benjamin Reed</i>	

## Securing Storage

Implementing and Evaluating Security Controls for an Object-Based Storage System .....	87
<i>Zhongying Niu, Ke Zhou, Dan Feng, Hong Jiang, Frank Wang, Hua Chai, Wei Xiao, and Chunhua Li</i>	
Trustworthy Migration and Retrieval of Regulatory Compliant Records.....	100
<i>Soumyadeb Mitra, Marianne Winslett, Windsor H. Hsu, and Xiaonan Ma</i>	
Capability Based Secure Access Control to Networked Storage Devices .....	114
<i>Michael Factor, Dalit Naor, Eran Rom, Julian Satran, and Sivan Tal</i>	

## Object-Based Storage

Enabling Database-Aware Storage with OSD .....	129
<i>Aravindan Raghuv eer, Steven W. Schlosser, and Sami Iren</i>	
Design and Implementation of a Network Aware Object-Based Tape Device.....	143
<i>Dingshan He, Nagapramod Mandagere, and David Du</i>	
Providing Quality of Service Support in Object-Based File System .....	157
<i>Joel C. Wu and Scott A. Brandt</i>	

## Making Storage More Reliable

Efficient Logging and Replication Techniques for Comprehensive Data Protection .....	171
<i>Maohua Lu, Shibiao Lin, and Tzi-cker Chiueh</i>	
Partial Disk Failures: Using Software to Analyze Physical Damage .....	185
<i>Hai Huang and Kang G. Shin</i>	
RAIF: Redundant Array of Independent Filesystems .....	199
<i>Nikolai Joukov, Arun M. Krishnakumar, Chaitanya Patti, Abhishek Rai, Sunil Satnur, Avishay Traeger, and Erez Zadok</i>	

## Short Papers

TPT-RAID: A High Performance Box-Fault Tolerant Storage System.....	215
<i>Yitzhak Birk and Erez Zilber</i>	
Tornado Codes for MAID Archival Storage.....	221
<i>Matthew Woitaszek and Henry M. Tufo</i>	
Cryptographic Security for a High-Performance Distributed File System .....	227
<i>Roman Pletka and Christian Cachin</i>	
Implementation and Evaluation of a Popularity-Based Reconstruction Optimization Algorithm in Availability-Oriented Disk Arrays.....	233
<i>Lei Tian, Hong Jiang, Dan Feng, Qin Xin, and Xing Shu</i>	
Performance Evaluation of Multiple TCP Connections in iSCSI .....	239
<i>Bhargava K. Kancherla, Ganesh M. Narayan, and K. Gopinath</i>	
E-Beam Hard Disk Drive Using Gated Carbon Nano Tube Source and Phase Change Media .....	245
<i>William S. Oakley</i>	
The RAM Enhanced Disk Cache Project (REDCAP) .....	251
<i>Pilar González-Férez, Juan Piernas, and Toni Cortes</i>	

Inter-Node Communication in Peer-to-Peer Storage Clusters .....	257
<i>André Brinkmann and Sascha Effert</i>	
Attribute Storage Design for Object-Based Storage Devices .....	263
<i>Ananth Devulapalli, Dennis Dalessandro, Pete Wyckoff, and Nawab Ali</i>	
Cost Analysis of the X-Code Double Parity Array .....	269
<i>Alexander Thomasian and Jun Xu</i>	
CASTOR: A Distributed Storage Resource Facility for High Performance Data Processing at CERN.....	275
<i>Giuseppe Lo Presti, Olof Barring, Alasdair Earl, Rosa María García Rioja, Sébastien Ponce, Giulia Taurelli, Dennis Waldron, Miguel Coelho Dos Santos</i>	
Using Comprehensive Analysis for Performance Debugging in Distributed Storage Systems .....	281
<i>Andrew W. Leung, Eric Lalonde, Jacob Telleen, James Davis, and Carlos Maltzahn</i>	
Hybrid Host/Network Topologies for Massive Storage Clusters.....	287
<i>Asha Andrade, Ungzu Mun, Dong Hwan Chung, and Alexander E. Mohr</i>	
Early Experiences in Managing Inter-Site Storage Area Networks Using Secure Web Services .....	293
<i>Ben Kobler, Fritz McCall, Mike Van Opstal, Hoot Thompson, and Kirk Hunter</i>	
Flushing Policies for NVCache Enabled Hard Disks.....	299
<i>Timothy Bisson and Scott A. Brandt</i>	
<b>Author Index</b> .....	<b>305</b>