PLATINUM SPONSORS

EMC²

NetApp

rackspace HOSTING

GOLD SPONSORS

EMULEX

ORACLE

SCALITY

CONTRIBUTING SPONSORS

NexiTech

OrangeFS

SILVER SPONSORS

ALTIOR elevated performance

anobit add another bit

IDEMA

Likewise

SNIA Solid State Storage Initiative

SNIA Storage Management Initiative

SwiftTest

CIFS/SMB/SMB2 PLUGFEST UNDERWRITER

Microsoft

IN PARTNERSHIP WITH:

NetApp

CLOUD PLUGFEST UNDERWRITER

SNIA Cloud Storage Initiative

MEDIA PARTNERS

BusinessSolutions

COMPUTERWORLD

Data Storage Connection Storage Networking News And Solution

Extension MEDIA

LINUX Magazine

sns EUROPE

StorageMojo

ASSOCIATION PARTNERS

CLDS

DMTF

FCIA

STA SCSI Trade Association

USENIX

Sponsors as of 9/15/11

Don't just be a cloud builder. Be a rainmaker.

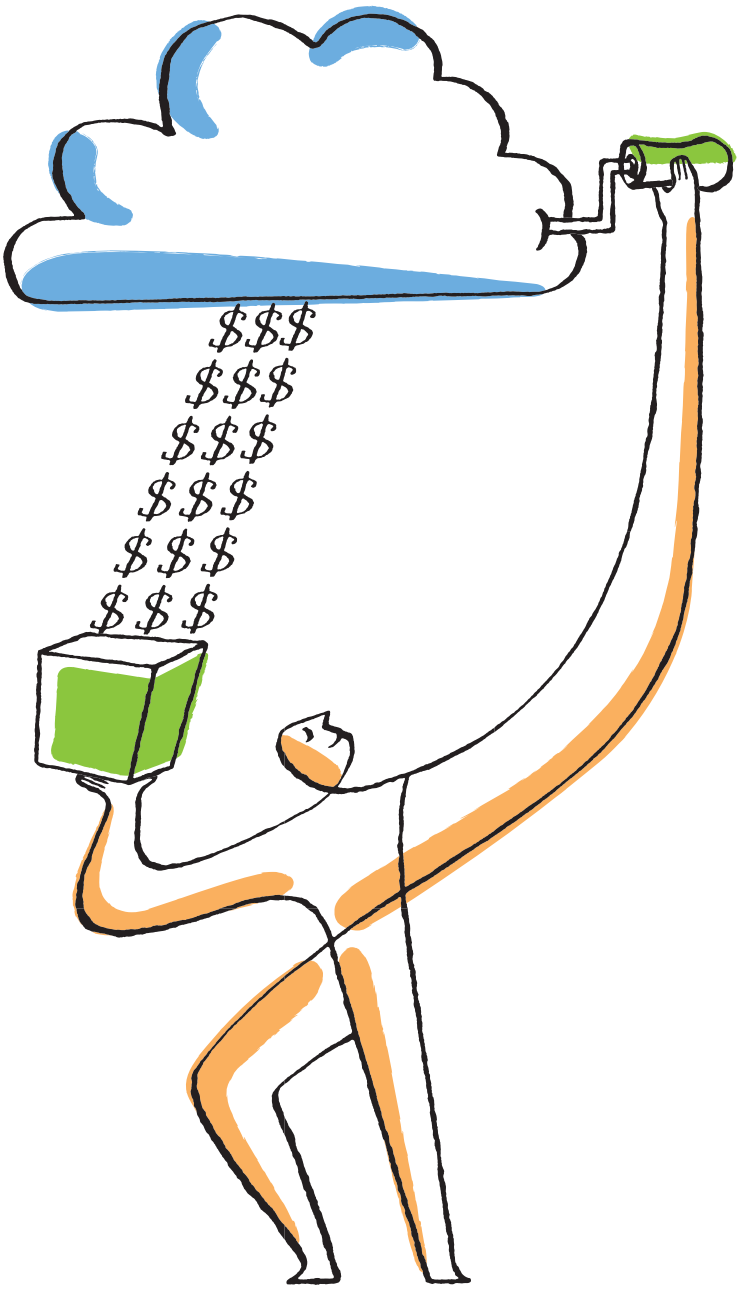

IT thought leaders and over 1 billion end users profit from clouds built on a NetApp storage foundation. To make sure your storage architecture is designed to deliver all the rewards the cloud has to offer, visit NetApp.com/BuiltOn.

Go further, faster®

Profitable clouds are built on **NetApp**®

Scan to learn how to reap the full rewards of the cloud. Get the app at 2dscan.com.

WELCOME

Greetings SDC Attendees,

On behalf of SNIA's Technical Council, Board of Directors and Staff, welcome to the 2011 Storage Developer Conference (SDC). This marks SNIA's 8th annual SDC and we have expanded and improved our conference to keep you engaged and learning while providing an important event for you to network with your peers and industry leaders.

We have a rich and varied agenda featuring speakers including Eric Baldeschwieler, CEO of Hortonworks as our SDC keynote. We have panel discussions, as well as multiple technical session tracks, our SNIA CIFS/SMB/SMB2 Interoperability Plugfest, and new this year, the SNIA CDMI and OCCI Cloud Plugfest. We encourage you to explore all facets of SDC and expect that you will find the entire experience to be valuable and productive.

Please remember to visit the Sponsor Showcase. You won't want to miss the opportunity to meet with SDC sponsors and several of our association and media partners. While you're there, please thank them for supporting the event.

Note that on Thursday the SNIA Cloud Storage Initiative will deliver the 2nd annual Cloud Burst Summit. Consider participating—you can register at no charge.

Finally, please take a moment to complete the SDC evaluations which will be e-mailed to you at the end of each day. We take your input very seriously and rely on it as we plan for next year's event. SNIA leaders will be available throughout the event to talk with you about your SNIA membership experience, the SNIA ROI on technical standards work, and ideas to take SNIA to new heights for 2012.

Thank you for joining us this week. Enjoy the Conference!

Regards,

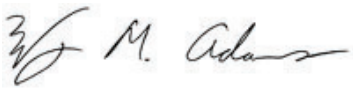

Richelle Ahlvers
Chairperson,
SNIA Technical Council

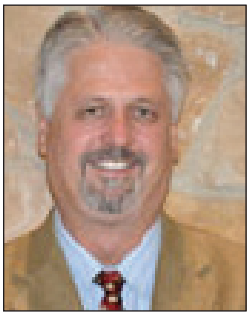

Wayne M. Adams
Chairman,
SNIA Board of Directors

TABLE OF CONTENTS

HOTEL & EVENT MAP

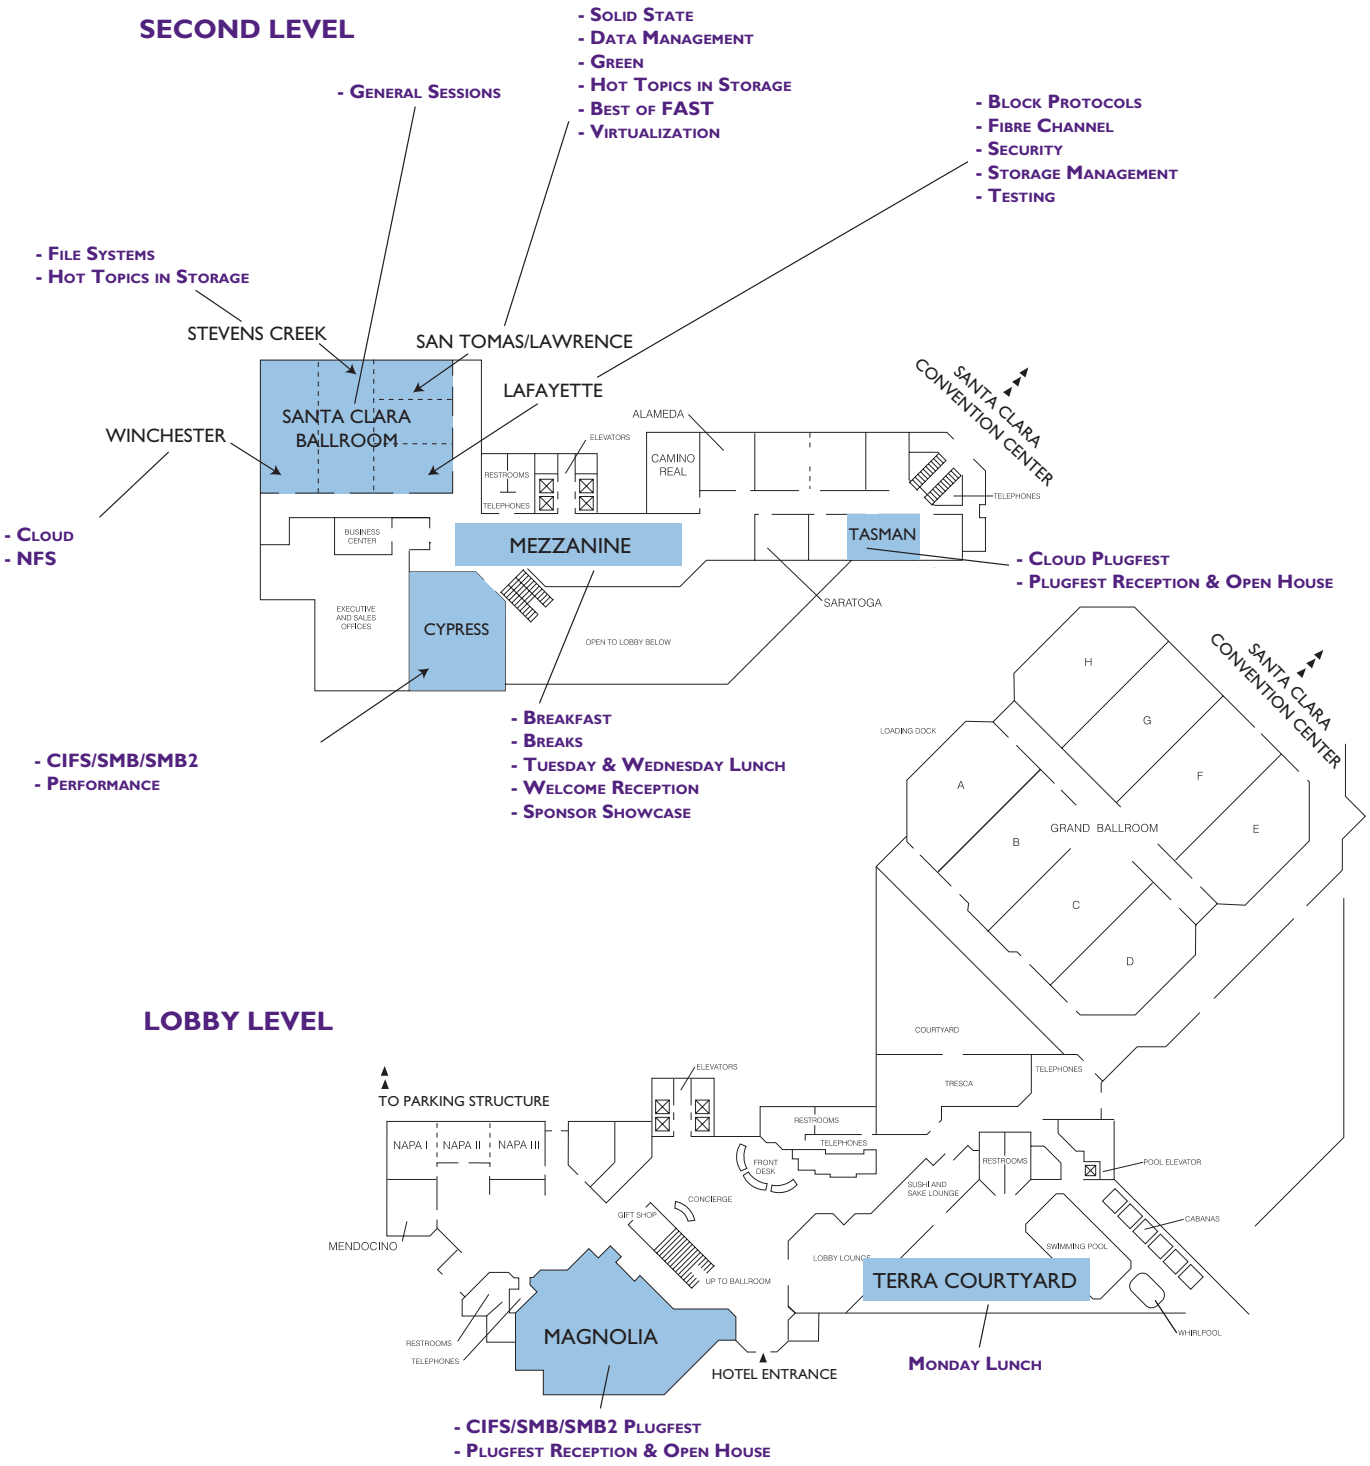

GENERAL INFORMATION

THE SNIA TECHNICAL COUNCIL

The Storage Developer Conference is brought to you by the Storage Networking Industry Association (SNIA), the leading association for education, interoperability, and standards programs for the storage networking industry.

The technical program for the conference was developed by the SNIA Technical Council, a select group of acknowledged industry experts who work to guide the SNIA's technical efforts. The Technical Council oversees and manages SNIA Technical Work Groups, reviews architectures submitted by Work Groups, and is the SNIA's technical liaison to standards organizations.

WIRELESS INTERNET ACCESS

During the conference complimentary wireless Internet access will be available. Just log in to the **sniawifi network**, using the **40-bit WEP key dc00112233**. Please be mindful that the Internet is shared with all your attendees; any abuse may result in revocation of your connection rights.

ONLINE ACCESS TO PRESENTATIONS

SDC EVALUATIONS

As an SDC attendee, you will have the opportunity to take part in many different areas of the conference whether it be attending a presentation, participating in the Plugfests, or attending a Birds of a Feather session. Your feedback on our agenda and program is invaluable to us.

We will be sending daily evaluations via email each night for that day's events. Please take a moment to complete these brief evaluations. The information you provide will help us in the planning and development for future conferences.

As an incentive, we will be offering a $100 gift card to a randomly selected individual that completes the evaluations. Remember, each time you complete an evaluation, your chances of winning go up!

SCHEDULE AT-A-GLANCE

Monday, September 19, 2011

Continental Breakfast	7:30 am - 8:45 am
Registration	7:30 am - 5:00 pm
Technical Sessions	8:30 am - 5:25 pm
Lunch	12:30 pm - 1:30 pm
Reception and Plugfest Open House	6:30 pm - 7:30 pm

Tuesday, September 20, 2011

Continental Breakfast	7:30 am - 8:45 am
Welcome Remarks	8:45 am - 9:00 am
General Sessions	9:00 am - 12:00 pm
Lunch	12:00 pm - 1:00 pm
Track Sessions	1:00 pm - 4:55 pm
Reception and Sponsor Showcase	5:00 pm - 7:00 pm
Birds of a Feather Sessions	7:00 pm - 9:00 pm

Wednesday, September 21, 2011

Continental Breakfast	7:30 am - 8:45 am
Introduction to Day	8:45 am - 9:00 am
General Sessions	9:00 am - 10:40 am
Food Fight Panel	11:10 am - 12:00 pm
Lunch/Sponsor Showcase	12:00 am - 1:00 pm
Track Sessions	1:00 pm - 4:55 pm
SNIA Cloud Plugfest Concludes	5:30 pm
Birds of a Feather Sessions	7:00 pm - 9:10 pm

Thursday, September 22, 2011

Continental Breakfast	7:30 am - 8:30 am
Introduction to Day	8:15 am - 8:30 am
Keynote Speaker	8:30 am - 9:20 am
Track Sessions	9:30 am - 12:15 pm
Sessions Conclude	12:15 pm
SNIA CIFS/SMB/SMB2 Plugfest Concludes	1:00 pm

MONDAY AGENDA

7:30 - 8:45	Continental Breakfast & Networking (Mezzanine)				
7:30	Registration Opens (Mezzanine)				
	Cloud (Winchester)	**File Systems (Stevens Creek)**	**Solid State (San Tomas/Lawrence)**	**CIFS/SMB/SMB2 (Cypress)**	**Block Protocols (Lafayette)**
8:30 - 9:20	Programming the Cloud John Kilroy Principal Software Engineer, EMC Fleur Dragan Consultant Software Engineer, EMC	NFSv4 Protocol Development Tom Haynes, Ph.D. Sr. Engineer, NetApp	HDDs and Flash Memory: A Marriage of Convenience Thomas Coughlin President, Coughlin Associates Jim Handy President, Objective Analysis	Hidden Gems in the NAS Protocols James Cain Principal Software Architect, Quantel Limited	SAS Standards and Technology Update Harry Mason Director Industry Marketing, LSI Marty Czekalski Program Manager, Seagate
9:30 - 10:20	CDMI for Cloud IPC David Slik Technical Director, Object Storage NetApp	"The Impossible Takes Longer": Emulating Windows File System Semantics on POSIX Jeremy Allison Google	Enhance NAND to Expand Enterprise SSD Market Esther Spanjer SMART Modular	Through the Looking Glass; Debugging CIFS/SMB/SMB2 Robert Randall Senior Software Architect, Micron Technologies, Inc Christopher Hertel Storage Architect and CIFS Geek, ubiqx Consulting, Inc	File Systems and Thin Provisioning Frederick Knight Standards Technologist NetApp
10:20 - 10:35	Break (Mezzanine)				
10:35 - 11:25	Open Source Droplet Library with CDMI Support Giorgio Regni CTO, Scality	Leveraging Btrfs Transactions Sage Weil Co-founder, New Dream Network	Emerging Performance Tests for Solid State Storage Devices Eden Kim CEO, Calypso Systems, Inc	Lessons Learned Implementing a Multi-threaded SMB2 Server in OneFS Aravind Velamur Srinivasan Senior Software Engineer, Isilon Systems	Storage Data Movement Offload Frederick Knight Standards Technologist, NetApp
11:35 - 12:25	CDMI Federations, Year 2 David Slik Technical Director, Object Storage, NetApp	Ceph Distributed Storage Sage Weil Co-founder, New Dream Network	PCIe Solid State Storage Devices Robert Randall Senior Software Architect, Micron Technologies, Inc	Implementing SMB 2.1 in Likewise Storage Services Gerald Carter CTO, Likewise Software	Data Integrity from Application to Storage William Martin Engineer Consultant, Emulex Frederick Knight Standards Technologist, NetApp
12:30 - 1:30	Lunch (Terra Courtyard)				

PLUGFEST INFORMATION

Plugfest Open House and Reception - Monday 6:30 - 7:30 PM • Magnolia & Tasman

CLOUD PLUGFEST FEATURING CDMI/OCCI

Sunday 1:00 PM - Wednesday 5:30 PM • Tazman Room • Underwritten by SNIA Cloud Storage Initiative

The purpose of this Plugfest is for vendors to bring their implementations of Cloud standards based products to test, identify, and fix bugs in a collaborative setting with the goal of providing a forum in which companies can develop interoperable products. This year's Plugfest will include the testing of client and server products using CDMI and/or OCCI.

2011 Cloud Plugfest Participants as of 9/12/11

NetApp SCALITY R2AD TATA

MONDAY AGENDA

	Cloud (Winchester)	File Systems (Stevens Creek)	Solid State (San Tomas/Lawrence)	CIFS/SMB/SMB2 (Cypress)	Fibre Channel (Lafayette)
1:30 - 2:20	Smart Hybrid Storage based on Intelligent Data Access Classification M. K. Jibbe Technical Director Arun Rajendran Software Engineer, NetApp	The Design and Evolution of the Apache Hadoop Distributed File System Dhruba Borthakur Project Lead Apache Hadoop DFS, Facebook	Hybrid Redundancy System New Approach to SSD Redundancy Avraham Meir CTO, Anobit	Samba Status Report Volker Lendecke Samba Team / SerNet	16GFC Sets the Pace in Storage Networks Scott Kipp Senior Technologiest, FCIA Mark Jones Director Technical Marketing, Emulex
2:30 - 3:20	CDMI Conformance and Performance Testing David Slik Technical Director, Object Storage, NetApp	GPFS: Scale-Out File Storage John Palmer IBM	Speeding Up Cloud/Server Applications Using Flash Memory Sudipta Sengupta Research Scientist, Microsoft	CTDB Status: Clustered Samba Growing Up Michael Adam Senior Software Engineer, Samba Team / SerNet	Fibre Channel over Ethernet (FCoE) John Hufferd Owner, Hufferd Enterprises
3:20 - 3:35	Break (Mezzanine)				
3:35 - 4:25	Use of Storage Security in the Cloud David Dodgson Software Engineer, Unisys	GPFS-SNC: A Scalable File System for Analytics and Clouds Prasenjit Sarkar Computer Science Researcher and Master Inventor, IBM	SSDs in the Cloud Dave Wright CEO, SolidFire	Experiences in Clustering CIFS for IBM Scale Out Network Attached Storage Dr. Jens-Peter Akelbein, IBM Germany	FCoE: The Next Generation Michael Ko CTO Office, Huawei Symantec
4:35 - 5:25	Authenticating Cloud Storage with Distributed Keys Jason Resch Senior Software Engineer, Cleversafe	Windows 8 File System Performance and Reliability Enhancements in NTFS Neal Christiansen Principal Development Lead, Microsoft	How Scale-Up and Scale-Out Flash-Based Databases Dr. John Busch Founder and CTO, Schooner Information Technology	A CIFS Geek in Exile: What I did on my Holiday Christopher Hertel Storage Architect and CIFS Geek, ubiqx Consulting, Inc	Open-FCoE Software Initiator(s): Architecture, Management and Performance Prafulla Deuskar Storage Networking Architect, Intel
5:30 - 6:30				Introduction to the SNIA CIFS/SMB/SMB2 Plugfest	
6:30 - 7:30	Plugfest Open House and Reception (Magnolia & Tasman)				

PLUGFEST INFORMATION

Plugfest Open House and Reception - Monday 6:30 - 7:30 PM • Magnolia & Tasman

SNIA CIFS/SMB/SMB2 PLUGFEST

Sunday 8:30 AM - Thursday 1:00 PM • Magnolia Room • Underwritten by Microsoft in partnership with NetApp

The purpose of this Plugfest is for vendors to bring their implementations of CIFS/SMB/SMB2 to test, identify, and fix bugs in a collaborative setting with the goal of providing a forum in which companies can develop interoperable products. The participants of the Plugfest work together to define the testing process, assuring that objectives are accomplished.

2011 CIFS/SMB/SMB2 Plugfest Participants as of 9/12/11

TUESDAY AGENDA

7:30 - 8:45	Continental Breakfast & Networking (Mezzanine)				
8:45 - 9:00	Welcome Remarks: Wayne Adams, Chairman, SNIA Board; EMC Richelle Ahlvers, Chairperson, SNIA Technical Council; HP				
9:00 - 9:50	 General Session Speaker: Ric Wheeler, File System Team Manager and Architect, Red Hat Linux File & Storage Systems: Enabling the Latest Storage Hardware in Linux				
9:50 - 10:40	  General Session Speakers: Dr. Thomas Pfenning, General Manager, File Server Technologies, Microsoft Jim Pinkerton, Partner Architect, File Server Technologies, Microsoft The Future of File Protocols: SMB 2.2 in the Data Center				
10:40 - 11:10	Break (Mezzanine)				
11:10 - 12:00	 General Session Speaker: Garth Gibson, Professor, Carnegie Mellon University, and CTO, Panasas, Inc Scalable Table Stores: Tools for Understanding Advanced Key-Value Systems for Hadoop				
12:00 - 1:00	Lunch & Sponsor Showcase (Mezzanine)				
	Cloud (Winchester)	File Systems (Stevens Creek)	Data Management (San Tomas/Lawrence)	CIFS/SMB/SMB2 (Cypress)	Security (Lafayette)
1:00 - 1:50	Resilience at Scale in the Distributed Storage Cloud Alma Riska Consultant Software Engineer, EMC	Storage Stack Evolution in Windows Shiv Rajpal Principal Development Lead, Microsoft	Long Term Information Retention Sam Fineberg Information Management Chief Technologist, HP Software Simona Rabinovici-Cohen Research Staff Member, IBM Research - Haifa	SMB 2.2: Bigger. Faster. Scalier: Part 1 David Kruse Principal Development Lead, Microsoft	Using Protocol Fuzzing to Harden Storage Systems and to Protect Them from 0-day Attacks Mikko Varpiola Senior Security Expert, Codenomicon
2:00 - 2:50	Changing Requirements for Distributed File Systems in Cloud Storage Wesley Leggette Cleversafe, Inc	A Lightweight Layered Compressed File System with Hardware Acceleration Shirish H. Phatak Vice President of Technology, Altior Inc.	Open Unified Data Protection and Business Continuity Dr. Anupam Bhide CEO/Founder, Calsoft	SMB 2.2: Bigger. Faster. Scalier: Part 2 Mathew George Sr. Software Development Engineer, Microsoft	Adding Role Based Access Control onto a Unix Storage Platform Steven Danneman Senior Software Developer, Isilon Systems
2:50 - 3:05	Break (Mezzanine)				
	Cloud (Winchester)	File Systems (Stevens Creek)	Data Management (San Tomas/Lawrence)	CIFS/SMB/SMB2 (Cypress)	/etc (Lafayette)
3:05 - 3:55	Best Practices in Designing Cloud Storage Based Archival Solution Jim Rice Principal Engineer, EMC	Linear Tape File System (LTFS) Dr. David Pease Senior Technical Staff Member, Manager, IBM	A Centralized Data Protection Application for Cross Vendor Storage Systems Nishi Gupta Tata Consultancy Services Prateek Sinha Storage Solutions Developer, Tata Consultancy Services	Advancements in Backup to Support Application Storage on a File Server Molly Brown Principal Development Lead, Microsoft	A Case Study: Unique NAS Issues and Solutions at The MathWorks Ira Cooper Senior Systems Software Engineer, The MathWorks, Inc.
4:05 - 4:55	Tape's Role in the Cloud Chris Marsh Market Development Manager, Spectra Logic	TBA	TBA	SMB 2.2 over RDMA Thomas Talpey Software Architect, Microsoft Greg Kramer, Ph.D. Software Development Engineer, Microsoft	Deep Dive into CIM Client Development with SBLIM Brian Mason Staff Software Engineer, NetApp
5:00 - 7:00	Reception & Sponsor Showcase (Mezzanine)				

WEDNESDAY AGENDA

7:30 - 8:45	Continental Breakfast & Networking (Mezzanine)				
8:45 - 9:00	Welcome Remarks: Wayne Adams, Chairman, SNIA Board; EMC Richelle Ahlvers, Chairperson, SNIA Technical Council; HP				
9:00 - 9:50	 General Session Speaker: Andy Walls, Distinguished Engineer and Technical Lead, IBM Solid State Architectures in the Modern Data Center				
9:50 - 10:40	 General Session Speaker: Bret Piatt, Director of Corporate Development, Rackspace Leveraging the Cloud for Your Storage Needs				
10:40 - 11:10	Break (Mezzanine)				
11:10 - 12:00	"Food Fight" Panel: NFSv4.1 Implementation/Adoption: The Challenges Ahead				
12:00 - 1:00	Lunch & Sponsor Showcase (Mezzanine)				
	NFS (Winchester)	File Systems (Stevens Creek)	Data Management (San Tomas/Lawrence)	CIFS/SMB/SMB2 (Cypress)	Storage Management (Lafayette)
1:00 - 1:50	NFS High Availability in Windows Roopesh Battepati Principal Development Lead, Microsoft	Implementing Alternate Data Streams in Likewise Storage Services Wei Fu Software Design Engineer, Likewise Software Gerald Carter Director of Engineering, Likewise Software	Understanding Primary Storage Optimization Options Jered Floyd Chief Technology Officer and Founder, Permabit	SMB 2.2 – Advancements for WAN Molly Brown Principal Development Lead, Microsoft Mathew George Sr. Software Development Engineer, Microsoft	Implementing a SMI-S Provider from Checkbox to Industrial Strength Steve Peters Storage Management Software, PMC-Sierra
2:00 - 2:50	NFSv3 and SMB/SMB2 Interoperability in Likewise Storage Services Evgeny Popovich Senior Software Engineer, Likewise Software	TBA	Etracker: Track Files on Your Laptop and Enhance Your Storage Using Email Uttam Kaushik Manager, Engineering, EMC	Moving an Enterprise Database Platform to run on CIFS/SMB/SMB2 File Access Protocols Kevin Farlee Storage Engine Program Manager, SQL Server, Microsoft	Proxy Providers versus Embedded Providers (SMI-S) Srinivasa Reddy Gandlaparthi Software Architect, NetApp
2:50 - 3:05	Break (Mezzanine)				
	NFS (Winchester)	Hot Topics in Storage (Stevens Creek)	Green (San Tomas/Lawrence)	CIFS/SMB/SMB2 (Cypress)	Storage Management (Lafayette)
3:05 - 3:55	IETF NFSv4 Working Group: What's Next? Spencer Shepler Performance Architect, Microsoft	Advanced Format in Legacy Infrastructures – Disruptive or Transparent? Curtis Stevens Western Digital	A Method to Vary the Host Interface Signaling Speeds in a Storage Array Driving Towards Greener Storage Dr. M. K. Jibbe Technical Director Arun Rajendran Software Engineer, NetApp	Thinking Inside the Box: Embedded Active Directory / Storage Appliances Based on Samba Kai Blin Embedded Developer, Samba Team	"Windows Server 8" and SMB 2.2 - Advancements in Management Jose Barreto Principal Program Manager, Microsoft
4:05 - 4:55	Scale-out NAS with NFS Referrals and pNFS Dmitry Yusupov Vice President Software, Founder, Nexenta Systems	Programmable I/O Controllers as Data Center Sensor Networks: Build and Deliver High-Performance Network and Storage Solutions Sanjeev Datla Senior Director Engineering, Emulex Shaun Walsh VP Marketing, Emulex	Vibration Management System for Storage Performance Gus Malek-Madani CTO and Founder, Green Platform Corporation	Accelerating SMB2 Mark Rabinovich R&D Manager, Visuality Systems	Microsoft SMI-S Roadmap Update Jeff Goldner Principal Architect, Microsoft
5:30	Cloud Plugfest Concludes				

THURSDAY AGENDA

7:30 - 8:30	Continental Breakfast & Networking (Mezzanine)			
8:15 - 8:30	Introduction to the Day and Housekeeping Announcements Wayne Adams, Chairman, SNIA Board of Directors; EMC Richelle Ahlvers, Chairperson, SNIA Technical Council; HP			
8:30 - 9:20	 Keynote Speaker: Eric Baldeschwieler, CEO, Hortonworks Apache Hadoop Today and Tomorrow			
	Virtualization (Lawrence)	**Hot Topics in Storage (San Tomas)**	**Testing (Lafayette)**	**Performance (Cypress)**
9:30 - 10:20	Advancements in Hyper-V Storage Todd Harris Sr. Software Development Engineer Microsoft Senthil Rajaram Senior Program Manager, Microsoft	The Role of InfiniBand and Automated Data Tiering in Achieving Extreme Storage Performance Cynthia Mcguire, Software Engineering Director, Oracle Sun Storage Group	RESTful Fault Injector Jim Rice Principal Engineer, EMC	SMB 2.2: Advancements in Server Application Performance Dan Lovinger Principal Software Architect, Microsoft
10:20 - 10:30	Break (Mezzanine)			
	Virtualization (Lawrence)	**Best of FAST (San Tomas)**	**Testing (Lafayette)**	**Performance (Cypress)**
10:30 - 11:20	Supporting Virtualization and Large workloads on NAS Storage Dennis Chapman Senior Technical Director, NetApp	Emulating Goliath Storage Systems with David Leo Prasath Arulraj Software Development Engineer, Amazon Award winning paper from the USENIX File & Storage Technology Conference	Challenges of Testing Unified Storage Richard Sharpe Architect, Scale Computing	Performence Analysis of iSCSI & iSER in MPIO Environment Seikh Basiruddin Member Technical Staff, NetApp
11:25 - 12:15	Benefits of ARI support in Virtualization Sivakumar Subramani Senior Project Leader, Wipro Technologies	A Study of Practical Deduplication Dutch Meyer Ph.D. Student, University of British Columbia, Canada Award winning paper from the USENIX File & Storage Technology Conference	An Extensible Open-Source Synthetic iSCSI/SCSI Initiator Peter Murray Product Manager, SwiftTest	TBA
12:15	Sessions Conclude			
1:00	SNIA CIFS/SMB/SMB2 Plugfest Concludes			

SNIA™

Not yet a member of the SNIA? Contact Marty Foltyn to find out about the benefits you receive from SNIA membership at marty.foltyn@snia.org or 619.992.2195.

Already a member and have a question about member benefits or joining a SNIA Forum or Initiative? Contact Lisa Mercurio at lisa.mercurio@snia.org or 781.293.9860.

Session descriptions are listed chronologically.

MONDAY SESSIONS

CLOUD TRACK

PROGRAMMING THE CLOUD
MONDAY 8:30 - 9:20

This talk will describe how cloud-based applications interact with stored data. The traditional semantics of file systems are often not applicable or relevant. We will discuss how language-specific idioms from several common development frameworks are mapped into the stored data abstractions present in the cloud. We will also outline some of the practical implications of deploying such applications.

CDMI FOR CLOUD IPC
MONDAY 9:30 - 10:20

In addition to providing storage services to end users, cloud storage systems enable cloud-aware programs to communicate between themselves in a distributed and asynchronous manner by using the cloud as a platform for Inter-Process Communications (IPC). This session discusses ways that CDMI enables cloud IPC, and the various use cases enabled by the use of a cloud in this manner. Special emphasis is given to CDMI Queues, which provide a first-in-first-out storage object, and are valuable for fan-in, fan-out and buffering messages and data exchanged between programs.

OPEN SOURCE DROPLET LIBRARY WITH CDMI SUPPORT
MONDAY 10:35 - 11:25

A year ago we started the Scality Open Source Program here at SDC by opening our Droplet project, a BSD licensed cloud storage client library. Now Droplet has been extended to support CDMI in addition to the S3 protocol and it's the only cross cloud compatible C client library. Developers around the world have contributed and worked on very promising projects including a cloud migration tool, an incremental backup agent with data deduplication as well as virtual machine image targets.

CDMI FEDERATIONS, YEAR 2
MONDAY 11:35- 12:25

In addition to standardizing client-to-cloud interactions, the SNIA Cloud Data Management Interface (CDMI) standard enables a powerful set of cloud-to-cloud interactions. Federations, being the mechanism by which CDMI clouds establish cloud-to-cloud relationships, provide a powerful multi-vendor and interoperable approach to peering, merging, splitting, migrating, delegating, sharing and exchange of stored objects. In last years SDC presentation, the basics of CDMI federation were discussed. For year two, we will review further refinements to making federations interoperable, demonstrate common use cases enabled by federation, and discuss the ongoing work within the SNIA Cloud Storage Technical Working Group to add federation as a formal part of the CDMI 1.1 that is standard currently under development.

CDMI RETENTION IMPROVEMENTS
MONDAY 1:30 - 2:20

CDMI is the functional interface that applications may use to discover the capabilities of cloud and manage data. Retention is an important feature to secure data. CDMI uses time criteria to determine retention on CDMI object. Accuracy and integrity of retention time depends on integrity of clock. Relative time differences can lead to undesirable retention management behavior. This presentation discusses an alternative that leverages the existing retention capability of backend storage and associating it with CDMI objects. Having a feedback mechanism from the backend to the application would help to make retention period consistent and helps to avoid relying on local clock time. This can be achieved in collaboration with CDMI interface.CDMI application should maintain database of retention ID and respective backend time which will be queried to find the lifetime of data after frontend retention expires.

CDMI CONFORMANCE AND PERFORMANCE TESTING
MONDAY 2:30 - 3:20

In order to deliver on the multi-vendor interoperability promise of the Cloud Data Management Interface (CDMI) standard, conformance testing and performance benchmarking tools are an essential part of the development and user community ecosystem. This session reviews the goals of conformance and performance testing, and provides an overview of the open source confCDMI and perfCDMI tools released by NetApp to assist in the validation and performance characterization of CDMI storage systems.

USE OF STORAGE SECURITY IN THE CLOUD
MONDAY 3:35- 4:25

Everyone is concerned with the security of their storage in the cloud, however; security in any particular case depends on what the user is trying to accomplish. Someone storing pictures of their children in the cloud will have a different idea of security than someone who is generating payroll information. Storage security needs to be implemented with an understanding of the different needs of different users. Enterprises will want to use secure private clouds that are customized to their individual security requirements, while individuals will want public clouds to address their needs. The most important security requirements will be those that satisfy the needs of the greatest number of users in a particular cloud environment.

AUTHENTICATING CLOUD STORAGE WITH DISTRIBUTED KEYS
MONDAY 4:35 - 5:25

Cloud storage is different from traditional systems. Typically the storage provider is not fully trusted. Passwords are often reused, easy to crack, difficult to remember, and depend on availability of the authentication service. Private keys offer a more resilient and secure method, but migrating, using and maintaining such keys is burdensome. A new technique will be discussed which combines the best of both words: the ease of use of passwords with the security properties of keys. Using this method in cloud storage systems user experience, security, and robustness can be greatly improved. This presentation provides an overview of current INCITS T11.3 FC-BB-6 working group activities. Items that will be discussed include new terminology, the ENode functional model, VN2VN_Port functionality (i.e., point-to-point and point-to-multipoint), and distributed FCF functionality (i.e., Controlling FCF and FDF).

FILE SYSTEMS TRACK

NFSV4 PROTOCOL DEVELOPMENT
MONDAY 8:30 - 9:20

The NFSv4 protocol undergoes a lifecycle of definition and implementation. We'll examine the lifecycle, what goes into the selection of new features, how these features are refined, and the impact these features will have on end users. We'll also look at how implementation experience will feed back into the protocol definition.

"THE IMPOSSIBLE TAKES LONGER" : EMULATING WINDOWS FILE SYSTEM SEMANTICS ON POSIX
MONDAY 9:30 - 10:20

Over the years Samba has moved from a thin layer of Windows emulation on top of POSIX to implementing something similar to the Windows "File System Algorithms" layer. If you have to emulate Windows completely on the wire, you need to emulate it completely on top of your platform. As most new storage platforms are Linux-based, learn how Samba manages to create the illusion of Windows on POSIX, and about some of the things that are really impossible to get right.

SESSION DESCRIPTIONS

LEVERAGING BTRFS TRANSACTIONS
MONDAY 10:35 - 11:25

Btrfs is a relatively new file system for Linux built on top of a copy-on-write btree abstraction. Unlike most other file systems, btrfs stores all metadata (and some data) in the btree, and uses a common transaction commit framework to ensure that the file system image is consistent on disk at all times. The Ceph distributed file system switched to using btrfs as the underlying storage for each object storage node because it could hook into the transaction framework to keep its data and metadata consistent at all times. The mechanism used to do this has evolved over the past few years. This talk will provide a btrfs design overview, including its copy-on-write and writable snapshot features, and then talk more specifically about how btrfs's architecture can be leveraged by applications.

CEPH DISTRIBUTED STORAGE
MONDAY 11:35 - 12:25

As the size and performance requirements of storage systems have increased, file system designers have looked to new architectures to facilitate system scalability. Ceph is a distributed object store, network block device, and file system designed for reliability, performance, and scalability from terabytes to exabytes.

Ceph's architecture consists of two main components: an object storage layer, and a distributed file system that is constructed on top of this object store. The object store provides a generic, scalable storage platform with support for snapshots and distributed computation. This storage backend is used to provide a simple network block device (RBD) with thin provisioning and snapshots, or an S3 or Swift compatible RESTful object storage interface. It also forms the basis for a distributed file system, managed by a distributed metadata server cluster, which similarly provides advanced features like per-directory granularity snapshots, and a recursive accounting feature that provides a convenient view of how much data is stored beneath any directory in the system.

This talk will describe the Ceph architecture and then focus on the current status and future of the project. This will include a discussion of Ceph's relationship with btrfs, the file system and RBD clients in the Linux kernel, RBD support for virtual block devices in Qemu/KVM and libvirt, and current engineering challenges.

THE DESIGN AND EVOLUTION OF THE APACHE HADOOP DISTRIBUTED FILE SYSTEM
MONDAY 1:30 - 2:20

This talks describes the architecture of the Apache Hadoop Distributed File System (HDFS). It analyzes the evolution of HDFS by discussing why certain design decisions are made, what features are deemed more important than others and the type of applications that use HDFS. It contends that HDFS has been a creative but disruptive force in the world of general purpose file-systems.

GPFS - SCALE-OUT FILE STORAGE
MONDAY 2:30 - 3:20

Parallel file systems, once peculiar to the world of engineering and scientific computing, are expanding their reach into mainstream information technology. Industry analysts predict the overwhelming portion of data will be in files as opposed to raw block storage within the next few years, and the majority of that will be in scale-out systems. This talk will present GPFS, IBM's parallel file system, and discuss how it is broadening its reach from supercomputing to the mainstream commercial world. This entails not only growing new features and supporting new workloads, but also subsuming many of the functions of block storage like copy services and disaster recovery.

GPFS-SNC: A SCALABLE FILE SYSTEM FOR ANALYTICS AND CLOUDS
MONDAY 3:35 - 4:25

GPFS-SNC is a scale out file system that leverages locally attached disk to provide high bandwidth to data parallel applications. In this talk, I will give an overview of the internals of GPFS-SNC and show examples of its application in MapReduce applications, data warehousing and cloud systems.

FILE SYSTEM EVOLUTION IN WINDOWS
MONDAY 4:35 - 5:25

Architecture/approach delivering major benefits expected of modern file systems for Windows.

SOLID STATE TRACK

HDDS AND FLASH MEMORY: A MARRIAGE OF CONVENIENCE
MONDAY 8:30 - 9:20

This talk will be based upon research by Coughlin Associates and Objective Analysis on ways in which flash memory and HDDs can be combined to create computer architectures with performance approaching that of SSDs and storage costs approaching those of HDDs. The talk is based upon a report from Coughlin Associates and Objective Analysis released in 2011 called "HDDs and Flash Memory: A Marriage of Convenience" and will explore the ways that flash memory and HDDs can be combined in computers, including tablets. The talk will give our projections for growth of these hybrid and paired storage products and future developments.

ENHANCE NAND TO EXPAND ENTERPRISE SSD MARKET
MONDAY 9:30 - 10:20

While early enterprise adoption of SSDs has been focused primarily on extreme performance and endurance using SLC NAND flash, proliferation into the enterprise mainstream will be fueled by the value proposition offered by lower-cost MLC NAND-based SSD solutions. Off the shelf, MLC NAND flash is not capable of meeting the endurance and retention requirements of enterprise server and storage applications. In order to meet life time and reliability requirements, SSD designers must use advanced technologies, have access to internal flash features, and incorporate system level design techniques to enhance the native capabilities of MLC NAND flash. This presentation talks about methods to achieve high endurance in MLC NAND flash, including a multi-faceted approach that incorporates advanced signal processing, optimization and adapt ion algorithms, and SSD architectural solutions to create compelling storage solutions that are bound to change how designers look at storage architectures.

EMERGING PERFORMANCE TESTS FOR SOLID STATE STORAGE DEVICES
MONDAY 10:35 - 11:25

Discussion of emerging Performance tests for NAND Flash based solid state storage devices used in Client and Enterprise applications. This discussion focuses on device level synthetic tests that are tuned to more closely reflect workload characteristics observed in real world use cases. Difficulties in successful capture and replay of user workloads (IO trace capture and playback) leads to the need for more finely tuned synthetic device level tests that can provide repeatable and comparable performance tests.

PCIE SOLID STATE STORAGE DEVICES
MONDAY 11:35 - 12:25

The rapid pace of the evolution of storage devices is creating many new challenges in systems design and architecture. PCIe GEN2 devices with MSI-X interrupt processing can deliver staggering throughput which can stress the traditional storage stacks on today's operating systems. There are also new standards, NVM-Express, T10 SOP and PQI, which seek to standardize the interface between the operating system and device (much like AHCI did for ATA) for PCIe attached solid state storage devices. Early adopters face the traditional trade-offs between legacy compatibility and new architectures which deliver higher performance but diverge from conventional thinking.

HYBRID REDUNDANCY SYSTEM NEW APPROACH TO SSD REDUNDANCY
MONDAY 1:30 - 2:20

Solid State Drives (SSDs) utilize a large number of flash devices. Flash redundancy might be needed to ensure high data reliability and availability. Meanwhile, legacy Redundant Array of Independent Disks (RAID) implementations are often employed to ensure greater overall reliability and performance at the array system level. By combining SSD-level redundancy with RAID-level redundancy, higher overall reliability and performance can be achieved than when utilizing these techniques separately.

SPEEDING UP CLOUD/SERVER APPLICATIONS USING FLASH MEMORY
MONDAY 2:30 - 3:20

Flash is a non-volatile memory technology that sits conveniently in the huge gap between RAM and hard disk in terms of both cost and performance. With its properties of low power consumption, physical ruggedness, and small size, flash has enabled new experiences with many consumer electronic devices. However, it is only recently that flash is seeing widespread adoption in desktop and server applications, in the form of Solid State Drives (SSDs). The new applications of flash involve different storage access patterns (vs. in typical consumer devices) and pose new challenges to flash, due to its device properties, to deliver sustained high throughput and low latency. We advocate that innovation at the system/application software layer when using flash memory can lead to several factors of improvement in performance over simply using it as a drop-in hardware replacement for existing storage technologies. The key to deploying flash in the data center lies in (i) designing the software in a flash-aware manner so as to exploit its unique properties and work around its constraints, and (ii) identifying applications that can utilize the sweet spot between cost and performance. As an example of (i), we will present FlashStore, a high throughput, low latency persistent key-value store, that illustrates some guiding principles for designing software for flash -- exploiting fast random reads, minimizing random writes, using RAM space efficient techniques to index data on flash, and recognizing its non-volatile property. As concrete examples for (ii), we will present and evaluate two cloud/server applications that can benefit from a flash-based key-value store – (a) game state backend for Xbox LIVE online multiplayer gaming, and (b) ChunkStash, a flash-assisted inline data deduplication system.

SSDs IN THE CLOUD
MONDAY 3:35 - 4:25

This session will cover three different methods of using solid state drives to provide persistent, high-performance primary storage within the cloud. It will explain the use of solid state as cache, as a storage tier, and as a full data storage solution, covering the advantages and disadvantages of each method. The speaker will also discuss how advances in SSD technology are enabling strides in storage efficiency, as well as performance.

HOW SCALE-UP AND SCALE-OUT FLASH-BASED DATABASES
MONDAY 4:35 - 5:25

We present emerging storage and database software technologies providing optimal scale-up through ultra-high flash and multi-core parallelism and optimal scale-out through synchronous replication, exploiting commodity hardware advances to yield 10x performance and 90% reduction in downtime, and providing key new building blocks for greatly improving data center QOS and TCO. The certification of a 10 Gb iSCSI and a 10 Gb FCOE RAID Storage System on the same network elicits a lot of challenges at the development level and the Test / Quality Assurance level. The challenges are due to the fact that both 10 Gb iSCSI and 10Gb FCOE are newly deployed host interfaces in the RAID Storage environment. As a result, development module test for both interfaces on the same network should be designed very carefully to establish test coverage beyond basic operations with a RAID Storage system, standard RAID testing, or an Ethernet plug fest.

CIFS/SMB/SMB2 TRACK

HIDDEN GEMS IN THE NAS PROTOCOLS
MONDAY 8:30 - 9:20

Having spent the last few years implementing SMB and SMB2 servers, the presenter has discovered that there are parts of these protocols that seem to offer untapped semantic richness. This session will propose theories and demonstrate practical working examples that test these theories. Examples will include: dynamically offering different representations of the same resource, extracting provenance from running software and avoiding NAS head state to support dynamic failover in a clustered file-system. All these examples will be built up from theoretical principals and demonstrated with working prototypes.

THROUGH THE LOOKING GLASS; DEBUGGING CIFS/SMB/SMB2
MONDAY 9:30 - 10:20

While protocol suite tests are quite useful, there are other ways to understand how Windows is interacting with your SMB client or server. Built into the kernel of Windows is a treasure trove of telemetry which provides a rich context and clear complaints when redirectors are interacting with another end point. Watch as the treasure is revealed through simple examples that demonstrate the value of knowing what the Windows kernel can tell you about how your client or server is behaving. Leave with step by step instructions on how to use these valuable tools. The tools are freely available from Microsoft's web site.

LESSONS LEARNED IMPLEMENTING A MULTI-THREADED SMB2 SERVER IN ONEFS
MONDAY 10:35 - 11:25

This talk will examine the lessons learned implementing SMB2 in the OneFS operating system and also highlight performance optimizations made in a multi-threaded SMB server implementation. In addition to these, the talk will also compare SMBv1 vs SMB2 in OneFS operating system to stress the performance benefits of using SMB2 over SMBv1. SMB2 offers a server-side credit mechanism to throttle greedy clients. Different credit algorithms can cause weird client behavior in certain scenarios. We'll examine some common mistakes to avoid. The multi-threaded SMB servers have their own advantages and disadvantages. This presentation will throw light on the performance optimizations that can be made in a multi-threaded server implementation. In addition to the aforementioned points, the presentation will also highlight the inherent performance benefits achieved by using SMB2 protocol over SMBv1, by presenting some performance numbers of using SMBv1 vs SMB2 in OneFS operating system.

IMPLEMENTING SMB 2.1 IN LIKEWISE STORAGE SERVICES
MONDAY 11:35 - 12:25

After completing support for SMB2.0 and MS Vista clients, server implementers must turn their focus to the additional SMB2 protocol features utilized by Windows 7 and Windows 2008 R2 clients. This session will focus on experiences and knowledge gained from implementing SMB2.1 feature support in the Likewise Storage Services platform. Topics covered will include protocol dialect negotiation beyond SMB 2.0, concurrent support for Windows 7 leases and legacy oplocks, multi-credit I/O support, and persistent file handles.

SAMBA STATUS REPORT
MONDAY 1:30 - 2:20

Samba is a rapidly evolving project that is part of the basis for many NAS vendors. This talk will give an overview of the current development of Samba.

SESSION DESCRIPTIONS

CTDB STATUS - CLUSTERED SAMBA GROWING UP
MONDAY 2:30 - 3:20

CTDB is a highly specific clustered database and management software sitting between Samba and a cluster file system. It allows to create scaling CIFS/NFS clusters on Linux. An early self-contained implementer of all-active service clustering, CTDB now slowly finds its way into the Linux distributions as a managed resource of the Linux cluster stack. Initially, the problems that Samba is facing, when running on a file system cluster, are recalled as well as the design and history of the CTDB software. An overview is given of the past year's bigger changes in CTDB, especially transaction handling and vacuuming. The various modes in which CTDB can be run are described, and how major Linux distributions start to integrate CTDB in their cluster products.

EXPERIENCES IN CLUSTERING CIFS FOR IBM SCALE OUT NETWORK ATTACHED STORAGE
MONDAY 3:35 - 4:25

Clustering the CIFS and SMB2 protocol is enabling managing large scale data in a single name space scaling the bandwidth of access as well. IBM SONAS uses clustering across various nodes while scaling capacity indepently by a second tier of nodes. Experiences with clustering CIFS including the underlying clustered file system the past years led to improvements in regards of performance and stability. Utilizing SMB2 as the protocol leads to improvements beyond CIFS capabilities. Compared to traditional active-passive configurations larger clusters provide active-active configurations allowing flexible maintenance and management. This talk should give an insight on resolved performance challenges in applying clustered CIFS with different installations and workloads including improvements being made or currently applied to Samba and CTDB for being used in an Enterprise product.

A CIFS GEEK IN EXILE: WHAT I DID ON MY HOLIDAY
MONDAY 4:35 - 5:25

BranceCache is a distributed caching system implemented by Windows SMB2 servers. BITS, according to at least one Microsoft Blog, is the "Earth's most widely used file transfer service". This presentation covers a new Open Source implementation of both BITS and BranchCache.

INTRODUCTION TO THE SNIA CIFS/SMB/SMB2 PLUGFEST
MONDAY 5:30 - 6:30

Every year at the Storage Developers Conference, a group of elite engineers hides out in a darkened room with long rows of tables, lots of equipment, plenty of caffeinated beverages, and a guard at the door. What's with that?

The annual SNIA CIFS/SMB/SMB2 Plugfest is an opportunity for CIFS/SMB/SMB2 implementers to test their products for compatibility, exchange results, and to work together to develop interoperable solutions. If your datacenter has a mix of products that all work together, this Plugfest is one reason why.

This session will explain how the Plugfest works, who is there, what they are testing, and how your organization can participate next year. It will also prepare you to ask questions at the Plugfest Open House reception, which follows immediately after this brief talk.

BLOCK PROTOCOLS TRACK

SAS STANDARDS AND TECHNOLOGY UPDATE
MONDAY 8:30 - 9:20

SAS has become the backbone of enterprise storage deployments. SAS has rapidly evolved by adding new features, capabilities, and performance enhancements. This talk will include an up-to-the-minute recap of the latest additions to the SAS standard and roadmaps. It will focus on areas of enhanced connectivity solutions, MultiLink SAS, status on 12Gb SAS development status, and a new transport investigation of SOP (SCSI over PCIe).

FILE SYSTEMS AND THIN PROVISIONING
MONDAY 9:30 - 10:20

New operations to manage Thin Provisioning have been added to or updated in the ATA and SCSI standards recently. This session will explain these capabilities and their storage APIs so that file system developers and application developers will gain an understanding of how and when to use these new features to enhance the efficient use of their storage subsystems.

STORAGE DATA MOVEMENT OFFLOAD
MONDAY 10:35 - 11:25

SCSI Operations that allow storage devices to offload data movement by the host into the storage device have recently been updated and enhanced to provide new capabilities. This session will explore these new capabilities, and how they are being used today. File system and application developers will gain an understanding of how these features operate and how they can be used to improve performance

DATA INTEGRITY FROM APPLICATION TO STORAGE
MONDAY 11:35 - 12:25

Data integrity failures to high-visibility applications have prompted vendors to add data integrity mechanisms to databases, file systems, and storage devices. The Data Integrity model being developed by SNIA's Data Integrity Technical Working group (DITWG) presents a model of how data can be protected from the application to the storage device. The Data Integrity model defines the building blocks used in operating systems for protecting data and providing true end-to-end data integrity protection. This model utilizes the "protection Information" feature (also known as DIF) defined in the SCSI block device command set (SBC) standard as one of the forms of protection provided in the complete stack. However, additional interfaces are necessary to extend this protection all the way to the application.

FIBRE CHANNEL TRACK

16GFC SETS THE PACE IN STORAGE NETWORKS
MONDAY 1:30 - 2:20

Storage area networks based on 16 Gigabit/second Fibre Channel (16GFC) will be deployed in 2011 and double the throughput of 8GFC SANs. This presentation will give historical perspectives on the 7 generations of Fibre Channel and how 16GFC is different from other Fibre Channel speeds. The presentation will investigate applications that need 16GFC and the benefits of using 16GFC. The presentation will also discuss the latest developments of 32GFC and hardware that is driving the need for more speed.

FIBRE CHANNEL OVER ETHERNET (FCOE)
MONDAY 2:30 - 3:20

The Fibre Channel (T11.3) standards committee developed a Standard called Fibre Channel over Ethernet (FCoE) The FCoE standard specifies the encapsulation of Fibre Channel frames into Ethernet Frames and the amalgamation of these technologies into a network fabric that can support Fibre Channel protocols and other protocols such as TCP/IP, UDP/IP etc. A "Direct End-to-End" FCoE variant has been accepted for the next version of the Standard The tutorial will show the Fundamentals of these FCoE concepts and describe how they might be exploited in a Data Center environment and its position with regards to FC and iSCSI. The requirements on the Ethernet Fabric for support of FC protocols will also be shown.

FCoE: THE NEXT GENERATION
MONDAY 3:35 - 4:25

The FCoE standard allows host servers on the Ethernet network to access storage on the Fibre Channel SAN. A bridging element known as an FCoE Forwarder forwards frames between the two dissimilar fabric. As currently defined, all FCoE control and data frames must pass through the FCoE Forwarder. The resulting bottleneck problem is unavoidable when the bridging function is required. This is not the case when Fibre Channel devices are gradually phased out and replaced with FCoE ones. This talk will describe the different ways currently being standardized to resolve the bottleneck problem. In addition, this talk will describe an alternative solution which allows customers to retain the use of the FCoE Forwarder for the control plane but bypass it in the data plane.

OPEN-FCOE SOFTWARE INITIATOR(S) – ARCHITECTURE, MANAGEMENT AND PERFORMANCE
MONDAY 4:35 - 5:25

Open-FCoE Software initiator(s) have been released for Windows, Linux and ESX (in progress). Open-FCoE stack has been certified with EMC E-labs and NetApp storage certification. This presentation takes a look at the architecture of the Open-FCoE software stack for different Operating Systems, how it inter-operates with DCB and how it plugs into management frameworks. We will also take a look at performance of the stack and how it compares with HW initiators especially with benchmarks which mimic real world applications like JetStress, SQL-IOSim etc. integrate CTDB in their cluster products.

TUESDAY SESSIONS

CLOUD TRACK

RESILIENCE AT SCALE IN THE DISTRIBUTED STORAGE CLOUD
Tuesday 1:00 - 1:50

The cloud is a diffuse and dynamic place to store both data and applications, unbounded by data centers and traditional IT constraints. However, adequate protection of all this information still requires consideration of fault domains, failure rates and repair times that are rooted in the same data centers and hardware we attempt to meld into the cloud. This talk will address the key challenges to a truly global data store, using examples from the Atmos cloud-optimized object store. We discuss how flexible replication and coding allow data objects to be distributed and where automatic decisions are necessary to ensure resiliency at multiple levels. We will discuss the impact of using a virtualized infrastructure inside the cloud, noting where this does and does not change the resiliency characteristics of the complete system and how it changes the design reasoning compared to purely physical hardware. Automatic placement of data and redundancy across a distributed storage cloud must ensure resiliency at multiple levels, i.e., from a single node to an entire site. System expansion must occur seamlessly without affecting data reliability and availability. All these features together ensures data protection while fully exploiting the geographic dispersion and platform adaptability promised by the cloud.

CHANGING REQUIREMENTS FOR DISTRIBUTED FILE SYSTEMS IN CLOUD STORAGE
Tuesday 2:00 - 2:50

File systems typically have centralized metadata servers that present performance bottlenecks as concurrent users and system size increase. These are unique challenges for distributed file systems. Cloud storage systems often store large unstructured content, and the streaming write access patterns typical of such systems allows for optimizations that cannot be made in traditional file systems. A new technique that adapts principals from NOSQL and object storage paradigms - and uses information dispersal for both underlying storage and metadata - provides a viable solution for streaming write access patterns. This technique allows for distributed writes, no single point of failure, scalability of both system size and concurrent clients, and limits performance bottlenecks.

BEST PRACTICES IN DESIGNING CLOUD STORAGE BASED ARCHIVAL SOLUTION
Tuesday 3:05 - 3:55

Cloud storage facilitates the use case of digital archiving for long periods of time by transparently providing scalable storage resources. With ever increasing amount of data to be preserved for legal and compliance reasons, cloud storage when designed correctly, can provide a low cost solution in a geographically distributed environment. This presentation highlights the key considerations while developing an archive product using cloud storage based on REST interface. It goes on to highlight the design choices while developing a file based archiving solution to cloud storage using EMC Atmos as an example. The aspects covered in the slides are – security, performance, using vendor neutral APIs, developing portable application irrespective of the backend cloud supported, taking advantage of geographically spread cloud storage nodes, faster searches and an efficient disaster recovery mechanism.

TAPE'S ROLE IN THE CLOUD
Tuesday 4:05 - 4:55

There is no doubt cloud storage is having a profound impact on IT and how technologies are deployed and consumed. Tape is the strong, silent partner to the cloud – very much present and in use, but completely transparent to the end-user. Chris will discuss how cloud storage's consumption model is built around ease of use, flexibility and cost savings, and why tape is one of the most logical and cost effective tiers for storing data in the cloud; particularly as the cost difference between tape and disk increases as data sets grow. He will review the key benefits of tape, reveal why it is quickly becoming the media of choice for cloud providers, and provide real-world examples of tape's role in the cloud.

FILE SYSTEMS TRACK

STORAGE STACK EVOLUTION IN WINDOWS
Tuesday 1:00 - 1:50

Architecture/approach delivering major benefits for scalable deployments using Windows.

A LIGHTWEIGHT LAYERED COMPRESSED FILE SYSTEM WITH HARDWARE ACCELERATION
Tuesday 2:00 - 2:50

We will discuss a lightweight layered compressed file system that can be layered over either a Linux or Windows native file store. The Compressed File System can achieve 3:1 storage efficiencies using hardware-accelerated data compression while preserving the native file system syntax and semantics. Intelligent application of hardware acceleration enables the compressed file system to run in real time without adversely impacting system throughput. By using sparse file allocation on the native file system, the compressed file system generates no additional meta data.

LINEAR TAPE FILE SYSTEM (LTFS)
Tuesday 3:05 - 3:55

While there are many financial and practical reasons to prefer tape storage over disk for various applications, the difficultly of using tape in a general way is a major inhibitor to its wider usage. We present a file system that takes advantage of a new generation of tape hardware to provide efficient access to tape using standard, familiar system tools and interfaces. The Linear Tape File System (LTFS) makes using tape as easy, flexible, portable, and intuitive as using other removable and sharable media (such as a USB drive).

SESSION DESCRIPTIONS

DATA MANAGEMENT TRACK

LONG TERM INFORMATION RETENTION
Tuesday 1:00 - 1:50

As more of the world's information is digital throughout its entire lifecycle, we are faced with the age old issues of record keeping and preservation, applied to devices and formats that were never intended to last. Long-term digital information suffers from issues that didn't exist in short-term or paper world, such as media and format obsolescence, bit-rot, and loss of metadata. The SNIA Long Term Retention (LTR) TWG has taken the lead on this issue for the storage industry. Working with key stakeholders in the preservation field, the LTR TWG is developing the Self-contained Information Retention Format (SIRF). SIRF is a low-level data format that enables applications to interpret stored data, no matter what application stored it originally. SIRF will be examined in a new European Union integrated research project, called ENSURE – Enabling kNowledge, Sustainability, Usability and Recovery for Economic Value. ENSURE creates a preservation infrastructure for commercial digital information built upon cloud storage and virtualization enablement technologies. It explores issues such as evaluating cost and value for different digital preservation solutions, automation of preservation processes, content-aware long term data protection, new and changing regulations, and obtaining a scalable affordable solution by leveraging cloud technology. The presentation will cover use cases, requirements, and the proposed architecture for SIRF as well as its potential usage in ENSURE storage services.

OPEN UNIFIED DATA PROTECTION AND BUSINESS CONTINUITY
Tuesday 2:00 - 2:50

A unified model for Data Protection and Business Continuity in complex enterprise systems Today's data center administrators are faced with challenges of managing and protecting complex enterprise systems comprising of physical and virtual components composed of heterogeneous hardware and software with complex interconnects. A typical enterprise system consists of Applications – SAP, Exchange, Share Point more Middleware – Databases, Web Servers, Operating Systems – Windows, Linux, UNIX & Hypervisors Servers – Servers (Physical), Blades, Servers (Virtualized) Network – Switches Physical & Virtual, Routers, Firewalls Storage – SAN, NAS, DAS, Cloud All components in the enterprise system uses persistent storage in the form of DAS, NAS or SAN. The unified model presented and implemented by Calsoft explores a open independent framework for data protection using storage level snapshots. The framework uses SMI-S to interact with various enterprise system components and ensures a consistent state to perform data protection and disaster recovery.

A CENTRALIZED DATA PROTECTION APPLICATION FOR CROSS VENDOR STORAGE SYSTEMS
Tuesday 3:05 - 3:55

Data centers consist of Storage Products from different vendors to meet the storage needs of multiple servers hosting various applications like: Databases, MailServers. These mixed environments pose a challenge for administrators as they have to use different Data Protection Applications (DPAs) for different combinations of Applications, Operating Systems and Storage Systems. Here we are presenting an approach to resolve the mentioned issue. A centralized DPA can be developed with a single user interface, which uses storage API's in the backend to perform data protection operations on cross-vendor storage systems. It consists of pluggable modules for different functionalities in areas example Backup/Recovery, Archival, Compliance, Deduplication. We will be sharing the results of the PoC and the challenges that were faced and proposing the need of Standards in Storage Data Protection APIs.

SMR/SSD HYBRID DISK ARRAY: TAKING ADVANTAGE OF HIGH CAPACITY SHINGLED MAGNETIC DRIVES WITH HIGH PERFORMANCE SSDS
Tuesday 4:05 - 4:55

Shingled Magnetic Recording (SMR) drive is one of newly emerging disk drive technologies that could potentially exceed the super-paramagnetic limit imposed on growth of areal density in conventional hard disk drives. The potential advantages of high-capacity drives such as SMR drives can be best utilized only when overall performance metric such as response time or throughput can also improve for storage subsystem as a whole. In this presentation we show that a SMR/Solid State Drive (SSD) hybrid storage subsystem can address this overall performance gain by exploiting complementary advantages of high performance SSDs and high-capacity SMR drive with standard RAID configurations.

CIFS/SMB/SMB2 TRACK

SMB 2.2: BIGGER. FASTER. SCALIER - (PARTS I AND 2)
Tuesday 1:00 - 1:50 and Tuesday 2:00 - 2:50

This session comprises two parts that will take a detailed look at new extensions to the SMB 2 protocol. These new developments target improving file server availability and client-server performance. The sessions will give you an overview of what is new in SMB 2.2, and then examine in detail specific areas of the protocol.

SMB 2.2 Multichannel adds new levels of network scalability and support for modern interconnects. The 'persistent handles' feature adds fault tolerance and continuous availability to the protocol. We will also discuss auxiliary protocols which live side by side with the SMB 2.2 protocol to provide end-to-end reliability and manageability.

ADVANCEMENTS IN BACKUP TO SUPPORT APPLICATION STORAGE ON A FILE SERVER
Tuesday 3:05 - 3:55

There are many compelling reasons for server applications, such as Hyper-V, to store their data on a file share, but this cannot be done if it compromises the application's data backup and recovery strategy. This session will describe the new protocol, MS-FSRVP, designed for Windows Server 8 that allows an application server to drive the required coordination with a file server so that existing backup and recovery strategies continue to work as the application data moves from local to remote storage.

SMB 2.2 OVER RDMA
Tuesday 4:05 - 4:55

A new protocol, SMB2 Direct, has been created which allows SMB 2.2 to operate over Remote Direct Memory Access (RDMA) transports such as iWARP, Infiniband and RoCE. This layering enables significant increases in performance for all SMB 2.2 file-based workloads and dramatically broadens the applicability of SMB 2.2. The presentation will outline the goals and motivations for the new approach, and will make a deep dive into the SMB2 Direct protocol itself, including early performance results.

SECURITY TRACK

USING PROTOCOL FUZZING TO HARDEN STORAGE SYSTEMS AND TO PROTECT THEM FROM 0-DAY ATTACKS
Tuesday 1:00 - 1:50

Protocol fuzzing is a proactive method for discovering previously unknown flaws in software. Defects discovered through fuzzing, unless fixed, have a potential of exposing affected systems to Denial of Service (DoS) situations and Zero Day Attacks, which could increase liability, damage business reputation and cripple sales. This presentation explains how fuzzing can be used to harden interfaces of the modern storage system with hands-on examples of protocols such as SMB2 and NFSv4.

ADDING ROLE BASED ACCESS CONTROL ONTO A UNIX STORAGE PLATFORM
Tuesday 2:00 - 2:50

The traditional Unix authorization model defines an all powerful root user who can perform any system task, modify any file, and change any system configuration. This simple model produces several fundamental problems for a storage platform. The root user, whether maliciously or accidentally, can cause catastrophic data loss. They can also view and undetectably modify the contents of any file and thus need to be an extremely trusted individual. Solving these problems requires partitioning the traditional root administrative rights among many different users and limiting within the file system the ability of any one user to view and modify all files. This can be accomplished with Role Based Access Control.

/etc TRACK

A CASE STUDY: UNIQUE NAS ISSUES AND SOLUTIONS AT THE MATHWORKS
Tuesday 3:05 - 3:55

The MathWorks is not the first company that comes to mind when one thinks of heavy NAS users. However, our testing environment relies heavily on NAS, and our needs are very different from those of most NAS users. As such, we face a unique set of issues and challenges. This case study will trace the progression of the MathWorks' NAS implementation -- from our start with off-the-shelf vendors to our current homegrown solution. We will detail the decisions we made, why we made them, and what ultimately drove us to develop our own solution. We hope that, by the end of this talk, you will have a better idea of what your clients are thinking and why.

DEEP DIVE INTO CIM CLIENT DEVELOPMENT WITH SBLIM
Tuesday 4:05 - 4:55

TCIM XML client development can be painful and slow without a good client side toolkit. Some developers choose to make their own, but why? SMBLIM is an open source client side library that simplifies client creation. It has libraries for: C++, Java and with a little help .NET. The presentation covers querying Object, Classes and Association using the programmatic API as well as WBEM Query Language (WQL). Function calls with intrinsic and user defined method and a deep dive into performance tuning queries. As a bonus a brief tutorial on how to use the Java version of the library in a .NET environment is included. All examples are done in Java.

WEDNESDAY SESSIONS

NFS TRACK

NFS HIGH AVAILABILITY IN WINDOWS
Wednesday 1:00 - 1:50

This session covers advancements in high availability for the NFS file services provided in Windows Server. This discussion is centered around using multiple NFS file servers in a failover cluster. The talk will cover briefly cover Failover Cluster Resource model, NFS Resource DLL and NFS Server HA architecture.

NFSV3 AND SMB/SMB2 INTEROPERABILITY IN LIKEWISE STORAGE SERVICES
Wednesday 2:00 - 2:50

IT professionals are continually striving to reduce the management costs of storage systems and to provide seamless cross-protocol access. NFS and SMB/SMB2 deployments struggle with three common interoperability problems: how to deal with multiple directories (NIS/LDAP/AD), cross-protocol access control, and differences in file locking semantics. Likewise Identity Services provide administrators with a means to interact with directory services and the addition of an NFSv3 server to the Likewise Storage Services architecture makes it simple to solve the other two. The result is a storage layer that does not require users mapping, provides equal access to the same user accessing files from different protocols, works around certain protocol limitations such as the 16 groups AUTH_SYS limitation, and provides support for cross-protocol locking.

IETF NFSV4 WORKING GROUP -- WHAT'S NEXT?
Wednesday 3:05 - 3:55

With the delivery of NFSv4.1 (RFC 5661) in January of 2010, the NFS version 4 community has been busy building and delivering NFSv4.1 products. Fresh from that experience and emerging application requirements, the IETF NFSv4 Working Group has been busy identifying features and building out a proposal for the NFSv4.2 protocol. The attendee will be provided an insider's view of the proposed feature set, the timeline involved, and up-to-the-minute status of where the working group is headed. And most importantly, the attendee will learn what these new NFSv4.2 features will provide for the end application.

SCALE-OUT NAS WITH NFS REFERRALS AND PNFS
Wednesday 4:05- 4:55

This talk describes the referral features added to the NFS standard and how to take advantage of the features for scaling out NAS deployments. This feature will be compared with Microsoft's referral-based system, DFS. The talk will also cover pNFS, providing parallel data access for NFS clusters.

FILE SYSTEMS TRACK

IMPLEMENTING ALTERNATE DATA STREAMS IN LIKEWISE STORAGE SERVICES
Wednesday 1:00 - 1:50

Modern SMB/SMB2 clients make use of alternate data streams for a variety of application purposes such as desktop UI enhancements, additional document properties, and location information for files downloaded from untrusted networks. Expectations from end-users and client machines make support for data streams a highly desirable, if not required, feature in today's storage devices. Likewise-CIFS is the SMB/SMB2 file server component of the Likewise Open project's Active Directory integration effort and is part of the larger Likewise Storage Services platform. This session will present both an architectural overview of the Likewise PVFS driver's data stream implementation and as a case study about the effort required to add stream support into a pre-existing file server.

HOT TOPICS IN STORAGE TRACK

ADVANCED FORMAT IN LEGACY INFRASTRUCTURES – DISRUPTIVE OR TRANSPARENT?
Wednesday 3:05 - 3:55

Since the launch of Advanced Format (AF) technology on hard disk drives in July 2010, many drives have been shipped and integrated using 512 byte emulation standards (AF 512e). As the industry prepares to introduce AF into long-standing legacy infrastructures, including enterprise systems, additional concerns for in compatibility and data loss have be raised. How real are these concerns? Do AF 512e and its 4K native derivative (AF 4Kn) have negative implications to legacy infrastructures? Take this opportunity to learn about the disruptive or transparent nature of AF from industry expert, Curtis Stevens.

SESSION DESCRIPTIONS

PROGRAMMABLE I/O CONTROLLERS AS DATA CENTER SENSOR NETWORKS: BUILD AND DELIVER HIGH-PERFORMANCE NETWORK AND STORAGE SOLUTIONS
Wednesday 4:05- 4:55

In this session, we will present a brief introduction to the evolution of the next-generation programmable I/O controllers along with a framework of tools and best practices for building, monitoring, managing and deploying host and embedded applications that maximize your network and storage capabilities. We will discuss the key areas that can sap your performance across your application configuration, driver stack, network integration and storage back-end. Additionally, we will present two case studies from our target developer program for storage partners that focus on the tools, tricks and APIs required to make you network and storage applications sing for cloud content delivery, network and storage appliances.

DATA MANAGEMENT TRACK

UNDERSTANDING PRIMARY STORAGE OPTIMIZATION OPTIONS
Wednesday 1:00 - 1:50

Selecting the right primary data optimization technology for your storage platform and integrating it into your existing software can be much less painful than it may at first sound. While compression and data deduplication are both now standard features for backup, few storage vendors have integrated either into primary storage. There are many challenges when trying to leverage technology designed for backup within primary storage, but newer technologies designed specifically for primary storage are much simpler to implement. Combining deduplication with compression provides additive savings, allowing storage vendors to drive cost savings to their users while still maintaining data safety. In this session, Jered Floyd, CTO of Permabit, will compare the two technologies using real world case studies and will explore which is a better fit for different data types and system architectures.

ETRACKER – TRACK FILES ON YOUR LAPTOP AND ENHANCE YOUR STORAGE USING EMAIL
Wednesday 2:00 - 2:50

As humans we rely a lot on computers for doing day-to-day activities. This is driving personal computer data growth at a phenomenal pace – from photos to videos, movies to songs; we want to keep it all "online". To make space, data is burnt on DVDs or addition of external storage. This results in multiple copies of data at multiple places making it difficult to track files when needed. Small home business (SOHO) imperatively require to backup critical data and from time-to-time need to archive legal documents with proper tracking of changes done to such documents. This presentation describes an application that can help in versioning, back-up and archive of files residing in a laptop to a public email system, like Gmail or yahoo mail, but can still be accessed from laptops seamlessly. It explains how to leverage Windows NTFS features for managing life-cycle of laptop files.

GREEN TRACK

A METHOD TO VARY THE HOST INTERFACE SIGNALING SPEEDS IN A STORAGE ARRAY DRIVING TOWARDS GREENER STORAGE
Wednesday 3:05- 3:55

This paper describes a method which we can effectively alter the signaling speeds of a Host Interface based on set performance criterion or user defined time of day criterion that are user definable. The end goals are considerable power savings by changing the signaling speeds to a lower supported speed as our background study indicates the same .Bringing down the MTBF of the components by operating them at nominal speeds and improving the operable life span of the system move towards Greener Storage, low power operation, minimize Heat dissipation and emission reduction. We are also achieving demand based Host interface bandwidth allocation to balance the throughput requirements of the application.

VIBRATION MANAGEMENT SYSTEM FOR STORAGE PERFORMANCE
Wednesday 4:05- 4:55

Gus Malek-Madani, Founder and CTO, Green Platform Corporation, will share 3 sets of test results that demonstrate how normal levels of data center vibration can degrade IOPS and throughput performance in HDDs by as much as 2/3. These tests also show how this lost storage performance can be restored by mitigating vibration.

CIFS/SMB/SMB2 TRACK

SMB 2.2 – ADVANCEMENTS FOR WAN
Wednesday 1:00 - 1:50

The SMB 2.2 protocol continues to evolve to better support deployments over Wide Area Networks (WAN). In this session, we discuss the improvements made to allow for higher retrieval rates from BranchCache peers and the optimizations made to support increased caching of file and directory metadata.

MOVING AN ENTERPRISE DATABASE PLATFORM TO RUN ON CIFS/SMB/SMB2 FILE ACCESS PROTOCOLS
Wednesday 2:00 - 2:50

There are a lot of considerations to go through when converting a performance-sensitive enterprise app designed to run against direct-attached or SAN hardware, and running it against Network Attached Storage. The Microsoft SQL Server Storage Engine team faced this problem when redesigning SQL Server to run over CIFS/SMB/SMB2 protocols to use NAS. I will discuss some of key issues we tackled: I'll walk you through each of these stages, and discuss how this progressed in real life.

THINKING INSIDE THE BOX: EMBEDDED ACTIVE DIRECTORY / STORAGE APPLIANCES BASED ON SAMBA
Wednesday 3:05- 3:55

In many SOHO setups, a central storage server or NAS device is already in use. Existing Open Source software makes it very easy to also move the Active Directory domain controller to the same machine, providing easy-to-use user management and file/print services to SOHO customers. This talk will describe a proof-of-concept implementation of an embedded Active Directory DC and SMB/CIFS file/print server for SOHO setups that can be administrated using a web interface or existing AD management tools. While the proof-of-concept implementation is limited to about a dozen clients, the same system is useable on more powerful hardware for bigger networks.

ACCELERATING SMB2
Wednesday 4:05- 4:55

Global networks introduce a significant challenge when it comes to CIFS traffic, making it virtually unbearable for an end user. SMB2 is no less chatty than SMB, hence - it may be accelerated. We will show how to improve SMB/SMB2 traffic using various acceleration techniques. This presentation emphasizes on challenges introduced by SMB2 as in comparison with SMB. I will share our CIFS acceleration experience and performance statistics for accelerated WANs with SMB2 traffic.

STORAGE MANAGEMENT TRACK

IMPLEMENTING A SMI-S PROVIDER FROM CHECKBOX TO INDUSTRIAL STRENGTH
Wednesday 1:00 - 1:50

Data storage continues to grow at a rapid pace and managing the data becomes increasingly challenging. Complying with SNIA's SMI-S must be more than a check box. The presentation will chronicle the development and evolution of a full featured SMI-S provider and a web based GUI that manages a PCIe RAID card.

PROXY PROVIDERS VERSUS EMBEDDED PROVIDERS (SMI-S)
Wednesday 2:00 - 2:50

The implementation of SMI-S providers for managing Storage Arrays or Controllers often involves selection of type of providers (proxy or embedded). This presentation compares the advantages and dis-advantages of Proxy and embedded providers, design considerations for selecting any one of them along with various methods of implementing embedded and proxy providers. Issues in managing large number of objects, association traversal related issues and typical capabilities required in the providers to overcome this issue are also discussed. The Management client design consideration while managing proxy and embedded providers is also discussed.

"WINDOWS SERVER 8" AND SMB 2.2 - ADVANCEMENTS IN MANAGEMENT
Wednesday 3:05- 3:55

This session covers advancements in SMB2 file services management. This includes details on specific implementations of industry standards like Web-Based Enterprise Management (WBEM), Common Information Model (CIM) and Web Services-Management (WS-Man). It also includes discussions on management of Windows and Non-Windows systems providing SMB2 file services.

MICROSOFT SMI-S ROADMAP UPDATE
Wednesday 4:05- 4:55

Microsoft has been working to add SMI-S support to our products. Recent progress is evident in the announcement that System Center Virtual Machine Manager (VMM) 2012 will use industry-standard SMI-S storage providers for active management of storage arrays for configuring virtualized environments. This session will detail the progress around SMI-S support by Microsoft and discuss further work to integrate SMI-S into Microsoft's management infrastructure.

THURSDAY SESSIONS

VIRTUALIZATION TRACK

ADVANCEMENTS IN HYPER-V STORAGE
Thursday 9:30 - 10:20

Hyper-V is a virtualization solution included as part of Windows Server 2008 and Windows Server 2008 R2. It provides the ability to expose virtual storage to a virtual machine in a number of different ways, including the use of Virtual Hard Drive (VHD) files. The talk will include discussions on different storage configuration options, workloads, and performance for Hyper-V.

SUPPORTING VIRTUALIZATION AND LARGE WORKLOADS ON NAS STORAGE
Thursday 10:30 - 11:20

This presentation examines the hosting of enterprise level hypervisor and application workloads on storage provided by NAS servers. It will present a brief overview of the two main file protocols NFS & CIFS/SMB. Then a discussion of the use of NAS storage by a hypervisor and its guests. Next the use of NAS storage by a large database. A discussion on configuring namespace to more efficiently support the hypervisor and application workloads. Discussion on the use of array value-add such as snapshots, dedup or cloning with NAS. Finally a discussion of future trends in this area.

BENEFITS OF ARI SUPPORT IN VIRTUALIZATION
Thursday 11:25 - 12:15

As per PCI specification, a single physical adapter can support only up to eight individual functions as only three bits are allotted for identifying a function in BDF (BUS / DEVICE / FUNCTION) value used to refer any PCI device. PCI SIG group has come up with a new method called ARI (Alternative Routing ID) to interpret the Device Number and Function Number fields within Routing IDs, Requester IDs, and Completer IDs, thereby increasing the number of functions that can be supported by a single Device. Alternative Routing-ID Interpretation (ARI) enables next generation I/O implementations to support an increased number of concurrent users of a multi-function device while providing the same level of isolation and controls found in existing implementations. While ARI obviously benefits the virtualized operating environments where each Function can be uniquely assigned to a guest OS, ARI also benefits non-virtualized environments where, e.g. due to increased process improvements, a large number of I/O Functions can be integrated into a single Device. This ARI is used in both Multi function adapters and SRIOV (Single Root I/O virtualization) functionalities to support more number of functions on a single physical function. This paper will analyze the benefits that can be achieved by using ARI in Multifunction and SRIOV (in virtualized environments like KVM, VMware).

HOT TOPICS IN STORAGE TRACK

THE ROLE OF A INFINIBAND AND AUTOMATED DATA TIERING IN ACHIEVING EXTREME STORAGE PERFORMANCE
Thursday 9:30 - 10:20

Extreme storage performance demands a unique and innovative approach to balancing system performance and data integrity. During this presentation, we will discuss the architectural trade-offs and advantages associated with implementing extremely fast storage systems. Including in this presentation will be a discussion on the use of InfiniBand in contemporary storage and system designs, and the associated software mechanisms necessary to take advantage of very low latency protocols.

BEST OF FAST TRACK

EMULATING GOLIATH STORAGE SYSTEMS WITH DAVID
Thursday 10:30 - 11:20

Benchmarking file and storage systems on large file-system images is important, but difficult and often infeasible. Typically, running benchmarks on such large disk setups is a frequent source of frustration for file-system evaluators; the scale alone acts as a strong deterrent against using larger albeit realistic benchmarks. To address this problem, we develop David: a system that makes it practical to run large benchmarks using modest amount of storage or memory capacities readily available on most computers. David creates a "compressed" version of the original file-system image by omitting all file data and laying out metadata more efficiently; an online storage model determines the runtime of the benchmark workload on the original uncompressed image. David works under any file system as demonstrated in this paper with ext3 and btrfs. We find that David reduces storage requirements by orders of magnitude; David is able to emulate a 1 TB target workload using only an 80 GB available disk, while still modeling the actual runtime accurately. David can also emulate newer or faster devices, e.g., we show how David can effectively emulate a multi-disk RAID using a limited amount of memory.

SESSION DESCRIPTIONS

A STUDY OF PRACTICAL DEDUPLICATION
Thursday 11:25 - 12:15

We collected file system content data from 857 desktop computers at Microsoft over a span of 4 weeks. We analyzed the data to determine the relative efficacy of data deduplication, particularly considering whole-file versus block-level elimination of redundancy. We found that whole-file deduplication achieves about three quarters of the space savings of the most aggressive block-level deduplication for storage of live file systems, and 87% of the savings for backup images. We also studied file fragmentation finding that it is not prevalent, and updated prior file system metadata studies, finding that the distribution of file sizes continues to skew toward very large unstructured files.

TESTING TRACK

RESTFUL FAULT INJECTOR
Thursday 9:30 - 10:20

With the advent of cloud storage, REST (Representational State Transfer) is becoming the common method for access and utilization of cloud storage systems. This presentation explains a product development accelerator that can help in reducing the cycle time in development of REST based cloud storage and ISV application's by using a fault injector. A RESTful fault injector helps ISV to develop reliable applications for cloud storage. This tool facilitates simulation of error scenarios by injecting REST and HTTP error codes to validate handling of those errors by the application. This presentation explains the high level approach, the implementation details in Windows, tuning to enable usage of this accelerator across various cloud storage systems and configurations to simulate different error scenarios. The details of the implementation are elaborated by taking EMC Atmos cloud storage as an example.

CHALLENGES OF TESTING UNIFIED STORAGE
Thursday 10:30 - 11:20

Testing iSCSI targets can be difficult because teams often have to rely on existing, OS-based initiators. Testers could be more productive if an OS-agnostic synthetic initiator was available. This talk presents an extensible and open-source framework for building synthetic SCSI initiators based on iSCSI. It provides ways to build iSCSI-based tests as well as SCSI-based tests. At the lowest layers it uses Ronnie Sahlberg's iscsilib toolkit and uses C++ to build an iSCSI transport layer along with SCSI Request classes that can be easily extend. Examples of extending the basic SCSIRequest class are given along with test code that uses it. The source code is available under the GPL. It could be extended to SAS and FibreChannel.

AN EXTENSIBLE OPEN-SOURCE SYNTHETIC ISCSI/SCSI INITIATOR
Thursday 11:25 - 12:15

Storage systems that support multi-protocol file, block and object storage are challenging to test. These system are more complex to test than single access method systems. Both functional and load testing require new strategies to ensure correct operation and performance.

PERFORMANCE TRACK

SMB 2.2: ADVANCEMENTS IN SERVER APPLICATION PERFORMANCE
Thursday 9:30 - 10:20

This session discusses SMB2 file services performance, focused on scenarios where the SMB2 client is running an application server workload like SQL Server. This includes extensive comparative analysis of different configurations and specific optimizations for application server workloads.

PERFORMANCE ANALYSIS OF ISCSI & ISER IN MPIO ENVIRONMENT.
Thursday 10:30 - 11:20

iSCSI is an emerging storage network technology that allows block-level access to storage devices, such as disk drives, over a computer network. Since iSCSI runs over the ubiquitous TCP/IP protocol, it has many advantages over its more proprietary alternatives. Due to the recent movement toward 10 gigabit Ethernet, storage vendors are interested to see the benefits this large increase in network bandwidth could bring to iSCSI. In order to make full use of the bandwidth provided by a 10 gigabit Ethernet link, specialized Remote Direct Memory Access hardware is being developed to offload processing and reduce the data-copy-overhead found in a standard TCP/IP network stack.This analysis will cover the performance benefit of using RDMA in iSCSI environment over the normal software iSCSI stack. This presentation will also cover the benefit in a single path environment as well as multipaths environment.

ENGINEERING DYNAMIC OPTIMAL I/O PERFORMANCE ORACLE ASM WITH IDP
Thursday 11:25 - 12:15

Oracle ASM Intelligent Data Placement (IDP) provides an API to dynamically set ASM data files associate with the most actively used database objects in hot regions, where the actual disk speed and physical bandwidth is better and more appropriate for them. The engineering of hybrid array pools using flash and NAND technologies is also presented. The implementation of algorithmic strategies that will further support IDP beyond a native intelligent design can lead to performance improvements of up to 25% regardless of the storage technology used, although it is expected that best results will be widely better in solid state machines. The focus is to provide a set of best practices based on the recognition of specific scenarios and the resilient application of strategic algorithms accordingly.

HOT SPARES

Note: Back-up sessions will be presented if openings occur during SDC. The slides are posted on the SDC website.

ARCHIVING AND MANAGING BIG DATA IN THE CLOUD
Greg Arnette, CTO and Founder, Sonian

IP STORAGE PROTOCOLS: ISCSI
John Hufferd, Owner, Hufferd Enterprises

FCOE DIRECT END-NODE TO END-NODE (AKA FCOE VN2VN)
John Hufferd, Owner, Hufferd Enterprises

SMART HYBRID STORAGE BASED ON INTELLIGENT DATA ACCESS CLASSIFICATION
M. K. Jibbe, Technical Director, Arun Rajendran Software Engineer, NetApp

INTEROPERABILITY TOOLS FOR CIFS/SMB/SMB2 AND MORE
Paul Long, Technical Evangelist, Microsoft

PLATINUM

EMC Corporation, a Fortune 500 company, is the world's leading developer and provider of information infrastructure technologies. We provide customers of all sizes and across all industries with the products, services, and solutions they need to manage data growth and get the most value from their information. We help customers with solutions for information storage and management; backup, recovery, and archiving; governance, risk, and compliance; business continuity, information security, virtualization, and more.

EMC leads customers on the journey to the private cloud—a dramatically more efficient and flexible way to manage, deliver, and consume IT services for reduced costs and increased business agility. We're uniquely positioned to provide fully virtualized, next-generation architectures enabling complete control over data and applications, faster and higher ROI—and the ultimate in efficiency, control, and choice. Discover how EMC can help your organization control data growth and get more value from your information at: www.EMC.com.

NetApp creates innovative products—storage systems and software that help customers around the world store, manage, protect, and retain one of their most precious corporate assets: their data. They are recognized throughout the industry for continually pushing the limits of today's technology and consistently find ways to enable their customers to do things they couldn't do before at a speed they never thought possible. NetApp partners with industry leaders to create the most efficient and cost-effective solutions optimized for their IT needs, and to deliver and support them worldwide.

Rackspace® Hosting is the world's leading specialist in the hosting and cloud computing industry, and the founder of OpenStack™ , an open source cloud platform. The San Antonio-based company provides Fanatical Support® to its customers, across a portfolio of IT services, including Managed Hosting and Cloud Computing. In 2010, Rackspace was recognized by Bloomberg BusinessWeek as a Top 100 Performing Technology Company and listed on the InformationWeek 500 as one of the nation's most innovative users of business technology. The company was also positioned in the Leaders Quadrant by Gartner Inc. in the "2010 Magic Quadrant for Cloud Infrastructure as a Service and Web Hosting." For more information, visit www.rackspace.com.

GOLD

Emulex, the leader in converged networking solutions, provides enterprise-class connectivity for servers, networks and storage devices within the data center. The Company's product portfolio of Fibre Channel host bus adapters, network interface cards, converged network adapters, controllers, embedded bridges and switches, and connectivity management software are proven, tested and trusted by the world's largest and most demanding IT environments. Emulex solutions are used and offered by the industry's leading server and storage OEMs including, Apple, Cisco, Dell, EMC, Fujitsu, Hitachi, Hitachi Data Systems, HP, Huawei, IBM, LSI, Lenovo, NEC, Oracle, NetApp, Samsung and ZTE. Emulex is headquartered in Costa Mesa, Calif., and has offices and research facilities in North America, Asia and Europe. More information about Emulex (NYSE:ELX) is available at http://www.Emulex.com.

For more than a decade, Oracle has been a leader in the areas of grid computing, clustering, server virtualization, SOA shared services, and large-scale management automation, with hundreds of thousands of customer successes and a high level of investment. Now having added Sun server, storage, operating system and virtualization technology to our industry leading applications, middleware, and database solutions, Oracle offers a complete technology stack in which every layer is engineered to work together as a single system. Oracle's open, standards-based architecture gives our customers unmatched benefits including excellent systems availability, scalability, energy efficiency, powerful performance, and low total cost of ownership. See for yourself why more customers are choosing Oracle to implement their cloud infrastructure, be it private, public, and/or hybrid clouds. Give us a call at 1-800-633-0738 or visit www.Oracle.com.

Scality's RING Organic StorageTM is the ultimate solution for the storage of unstructured data for large scale Cloud Services such as webmail, photo, video and file sharing, SaaS, cloud backup, and Cloud Storage providers (STaaS).

RING Organic Storage is the very first object-based storage technology to deliver the performance required for primary application storage. It easily scales up to exabytes using x86 commodity server hardware with direct attached storage. The organic design creates a system with distributed intelligence, which has no single point of failure. As a result, RING Organic Storage is resilient and self-healing, and it implements technology refresh with zero downtime and no data migration.

RING solves a multitude of common administrative and operational problems associated with SAN, helping cloud services providers to deliver consistently excellent service, while reducing infrastructure costs. For more information, please visit www.scality.com.

SPONSOR INFORMATION

SILVER

Founded in 2004 as CebaTech, Altior™ offers a broad range of industry-leading hardware and software solutions for the networking and storage markets. Altior products deliver high performance for increased throughput, significant energy savings, and effective network andstorage optimization. For more information, please visit http://www.altior.com.

Anobit provides high performance flash storage solutions for client and cloud applications. Anobit products are comprised of embedded flash controllers for Smartphones and Tablet Computers, and SSDs (Solid State Drives) for enterprise storage. Anobit's products are used by world leading flash manufacturers, consumer electronics vendors and storage system providers.

The Advanced Format technology committee was formed by IDEMA in collaboration with leading data storage companies to educate and support the industry as hard disk drives transition to larger physical sector sizes using multiples of 512-byte logical sectors of user data. Today, the 4K sector format has gained broad adoption and is setting the groundwork for future breakthroughs in storage technology.

Likewise Software is an open source company that provides audit and authentication solutions designed to improve security, reduce operational costs and help demonstrate regulatory compliance in mixed network environments. Likewise Open is the first open source enterprise-class solution to address the authentication needs of organizations with mixed networks, and easily joins virtually all enterprise platforms to common directories.

SwiftTest products deliver high-capacity, highly-configurable traffic for all of the major storage protocols — CIFS/SMB, NFS, and iSCSI, and has an in-depth knowledge of storage and testing, and experience working with leading vendors in the storage industry, so that your testing problems will be solved by products and expertise unmatched by anything on the market today.

The SNIA Solid State Storage Initiative (SSSI) fosters the growth and success of solid state storage in commercial and consumer environments. The SSSI focuses on providing authoritative information on solid state storage, educating the vendor and user communities about solid state storage, performing market outreach that accurately highlights the advantages of solid state storage, guiding and pursuing standards for solid state storage, collaborating with other industry associations related to the success of solid state storage, enabling worldwide adoption of solid state storage.

Member-driven and vendor-neutral, the SMI is chartered to promote multi-vendor storage management interoperability for the products created by member companies through consensus building. Through its committees and programs, like SMI-Lab and the Conformance Test Program (CTP), the SMI is directing the movement to create standards for storage networking manage¬ment. Based on SMI's research and roadmap, SMI Technical Work Groups cooperatively develop the Storage Management Technical Specification known as SMI-S.

CIFS/SMB/SMB2 PLUGFEST UNDERWRITER

Founded in 1975, Microsoft (Nasdaq "MSFT") is the worldwide leader in software, services and solutions that help people and businesses realize their full potential.

CLOUD PLUGFEST UNDERWRITER

The mission of the SNIA Cloud Storage Initiative (CSI) is to foster the growth and success of the market for what is generally referred to as cloud storage and more generally the use of data storage resources and services in the Cloud. The CSI will also engage in cross-industry collaboration activities with relevant industry associations and standards development groups to further develop and promote all facets of Cloud services and standards along with Cloud Storage (e.g. compute, security, best practices).

GO WHERE UNIQUE TALENTS CONVERGE

At EMC, innovative thinking is sustained through diverse perspectives. The EMC workforce is a world of more than 50,000 thought leaders working together to drive the future of cloud computing and information management solutions.

Through innovative products and services, EMC accelerates the journey to cloud computing, helping IT departments to store, manage, protect and analyze their most valuable asset— information— in a more agile, trusted and cost-efficient way. To learn more about EMC, visit www.emc.com. To learn about working at EMC, visit www.emc.com/careers.

EMC²