SDC 13

Conference Guide

SEPTEMBER 16-19, 2013
HYATT REGENCY SANTA CLARA

Gold Sponsors

Platinum Sponsors

Silver Sponsors

Contributing Sponsors

Media Partners

Association Partners

1 WELCOME

Greetings SDC Attendees,

On behalf of SNIA's Technical Council, Board of Directors and Staff, welcome to the 2013 Storage Developer Conference (SDC). This marks SNIA's 10th annual SDC and we have expanded and enhanced 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 technology luminaries such as Intel's Andy Rudoff and storage infrastructure Best Practices shared by Go Daddy and Netflix along with the perspective of Addison Snell, CEO, Intersect360 Research as our SDC closing keynote. We have panel discussions, as well as multiple technical session tracks, our SNIA SMB2/SMB3 Interoperability Plugfest, the Cloud Interoperability Plugfest, and new this year, the SNIA iSCSI 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 Analytics and Big Data Summit will be happening onsite. Consider participating—you can register and participate at no charge.

Finally, please take a moment to complete the SDC evaluation surveys 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 2014.

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

Regards,

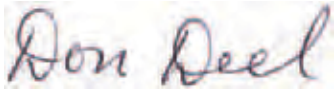

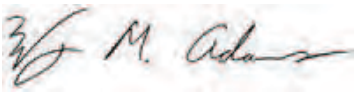

Don Deel
Chairman,
SNIA Technical Council

Wayne M. Adams
Chairman,
SNIA Board of Directors

CONTENTS

GENERAL INFORMATION

THE SNIA TECHNICAL COUNCIL

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.
Login Information:
SSID: SNIA-SDC
Key: sdc-2013

Wireless Internet
Sponsored by:

ONLINE ACCESS TO PRESENTATIONS

To access all the presentations from the SDC conference, go to:

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 16, 2013

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

Tuesday, September 17, 2013

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
Breakout 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 - 8:00 pm

Wednesday, September 18, 2013

Continental Breakfast ..... 7:30 am - 8:45 am
Introduction to Day ..... 8:45 am - 9:00 am
General Sessions ..... 9:00 am - 11:20 am
Software Defined Storage Roundtable ..... 11:20 am - 12:15 pm
Lunch/Sponsor Showcase ..... 12:15 pm - 1:15 pm
Breakout Sessions ..... 1:15 pm - 5:10 pm
Birds of a Feather Sessions ..... 7:00 pm - 9:00 pm

Thursday, September 19, 2013

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
Breakout Sessions ..... 9:30 am - 12:15 pm
Cloud Interoperability Plugfest Concludes ..... 12:00 pm
Sessions Conclude ..... 12:15 pm
SNIA SMB2/SMB3 and iSCSI Plugfests Conclude ..... 1:00 pm

3 MONDAY AGENDA

	Cloud (Winchester)	File Systems (San Tomas / Lawrence)	SMB2/SMB3 (Cypress)	Block Storage (Lafayette)	Distributed Storage (Stevens Creek)
7:30	Registration Opens (Mezzanine)				
7:30 - 8:45	Continental Breakfast (Mezzanine)				
8:30 - 9:20	**CDMI, The Key Component of Scality Open Cloud Access** Giorgio Regni, CTO, Scality Philippe Nicolas, Director of Product Strategy, Scality	**Advancements in Storage and File Systems in Windows 8.1** Andy Herron, Principal Software Developer, Microsoft	**SMB3 Meets Linux: The Linux Kernel Client** Steven French, Senior Engineer SMB3 Architecture, IBM	**SCSI Standards and Technology Update** Marty Czekalski, President, SCSI Trade Association	**Getting the Most Out of Erasure Codes** Jason Resch, Lead Software Engineer, Cleversafe, Inc.
9:30 - 10:20	**CDMI Federations, Year 4** David Slik, Technical Director, NetApp	**HDFS - What is New and Future** Sanjay Radia, Co-founder, Hortonworks Suresh Srinivas, Hortonworks	**Mapping SMB onto Distributed Storage** Christopher R. Hertel, Senior Principal Software Engineer, Red Hat José Rivera, Software Engineer, Red Hat	**SCSI and FC Standards Update** Fred Knight, Standards Technologist, NetApp	**LRC Erasure Coding in Windows Storage Spaces** Cheng Huang, Researcher, Microsoft Research
10:20 - 10:35			Break (Mezzanine)		
10:35 - 11:25	**LTFS and CDMI - Tape for the Cloud** David Slik, Technical Director, NetApp	**Multiprotocol Locking and Lock Failover in OneFS** Aravind Velamur Srinivasan, Senior Software Engineer, EMC, Isilon Systems	**Pike - Making SMB Testing Less Torturous** Brian Koropoff, Consulting Software Engineer, EMC Isilon	**Extending SAS Connectivity in the Data Center** Bob Hansen, Storage Architecture Consultant, LHP Consulting Group	**Distributed Storage on Limping Hardware** Andrew Baptist, Lead Architect, Cleversafe
11:35 - 12:25	**Profile Based Compliance Testing of CDMI: Approach, Challenges & Best Practices** Sachin Goswami, Solution Architect, and Storage COE Head@HiTech, Tata Consultancy Services Udayan Singh, Head SPE-Storage & Platform and Lead Innovate@HiTech, Tata Consultancy Services	**Snapshots for IBRIX - Highly Distributed Segmented Parallel FS** Boris Zuckerman, Distinguish Technologist, HP	**Exploiting the High Availability Features in SMB 3.0 to Support Speed and Scale** James Cain, Principal Software Architect, Quantel	**FCoE Direct End-Node to End-Node (aka FCoE VN2VN)** John Hufferd, Hufferd Enterprises	**Method to Establish a High Availability and High Performance Storage Array in a Green Environment** M. K. Jibbe, Director of Quality Architect Team for All APG Products, NetApp Marlin Gwaltney, Quality Architect, NetApp
12:30 - 1:30	Lunch (Terra Courtyard)				

MONDAY AGENDA

4

	Cloud (Winchester)	File Systems (San Tomas / Lawrence)	SMB2/SMB3 (Cypress)	Object Storage (Lafayette)	Distributed Storage (Stevens Creek)
1:30 - 2:20	Open-source CDMI-compliant Proxy: "Stoxy" Ilja Livenson KTH Royal Institute of Technology	Snapshot Cauterization Sandeep Joshi Manager, EMC	SMB3 Update David Kruse Development Lead, Microsoft	Huawei SmartDisk based Object Storage UDS Qingchao Luo Cloud Storage Architect, Huawei	Transforming PCIe-SSDs and HDDs with **Infiniband into Scalable** Enterprise Storage Dieter Kasper Principal Architect, Fujitsu
2:30 - 3:20	Lessons Learned Implementing Cross-protocol Compatibility Layer Scott Horan Integration Engineer, Cleversafe	Cluster Shared Volumes Vladimir Petter Principal Software Design Engineer, Microsoft	Implementation of SMB3.0 in Scale-Out NAS Kalyan Das Chief Architect, Storage Protocols, Huawei Technologies Jun Liu Software Architect - Storage and NAS, Huawei Technologies	Architecting Block and Object Geo-replication Solutions with Ceph Sage Weil, Founder & CTO, Inktank	**Software Defined X** (Stevens Creek) **Defining Software Defined Storage** Lazarus Vekiarides Entrepreneur and Technology Executive
3:20 - 3:35			Break (Mezzanine)		
	Cloud (Winchester)	**File Systems** (San Tomas / Lawrence)	**SMB2/SMB3** (Cypress)	**Hardware** (Lafayette)	**Software Defined X** (Stevens Creek)
3:35 - 4:25	COSBench: A Benchmark Tool for Cloud Storage Yaguang Wang Sr. Software Engineer, Intel	Balancing Storage Utilization Across a Global Namespace Manish Motwani Lead Software Developer, Cleversafe	1 S(a) 2 M 3 B(a) 4 Michael Adam SerNet GmbH	PCI Express and Its Interface to Storage Architectures Ron Emerick Principal HW Engineer, Oracle	Hosting Performance-Sensitive Applications in the Cloud with Software **Defined Storage** Felix Xavier Co-Founder and CTO, CloudByte
4:35 - 5:25	Architecting An Enterprise Storage Platform Using Object Stores Niraj Tolia Chief Architect, Maginatics	Scale-out Storage Solution Mahadev Gaonkar Technical Architect, iGATE	Samba 4.0 released: What Now for the Open Source AD Domain Controller? Andrew Bartlett Samba Developer, Samba Team	PCI Express IO Virtualization Overview Ron Emerick Principal HW Engineer, Oracle	**Software Defined** Network Technology and the Future of Storage Stuart Berman Chief Executive Officer, Jeda Networks
6:30 - 7:30	Plugfest Open House and Reception (Magnolia Room)				

SNIA iSCSI PLUGFEST

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

Presentations are only part of what is going on at the SNIA's 2013 Storage Developer Conference. This year we are pleased to announce the new SNIA iSCSI Plugfest made possible through the support of Microsoft, our 2013 iSCSI Plugfest underwriter. The purpose of this Plugfest is for vendors to bring their implementations of iSCSI to test, identify, and fix bugs in a collaborative setting with the goal of providing a forum in which companies can develop interoperable products.

2013 SNIA iSCSI Plugfest Participants (As of 9/10/13)

Chelsio Communications · CODENOMICON · ixia · Microsoft · ORACLE · SwiftTest

5 TUESDAY AGENDA

7:30 - 8:45	Continental Breakfast (Mezzanine)				
8:45 - 9:00	Welcome Remarks Wayne Adams, SNIA Board Chairman Don Deel, SNIA Technical Council Chairman				
9:00 - 9:45	The Impact of the NVM Programming Model Andy Rudoff, Intel, SNIA NVM Programming TWG				
9:45 - 10:25	Windows Azure Storage – Scaling Cloud Storage Andrew Edwards, Principal Architect, Windows Azure Storage, Microsoft				
10:25 - 10:40	Break (Mezzanine)				
10:40 - 11:20	Optical Storage Technologies: The Revival of Optical Storage Ken Wood, CTO – Technology & Strategy Office of Technology and Planning, Hitachi Data Systems				
11:20 - 12:00	Hypervisors and Server Flash Satyam Vaghani, CTO, PernixData				
12:00 - 1:00	Lunch and Sponsor Showcase (Mezzanine)				
	Cloud (Lafayette)	**File Systems** (San Tomas / Lawrence)	**SMB2/SMB3** (Cypress)	**Hardware** (Stevens Creek)	**Solid State** (Winchester)
1:00 - 1:50	Windows Azure Storage - Speed and Scale in the Cloud Joe Giardino Senior Development Lead, Microsoft	A Brief History of the BSD Fast Filesystem Dr. Marshall Kirk McKusick Computer Scientist Author and Consultant	SMB Direct update Greg Kramer Sr. Software Engineer, Microsoft Tom Talpey Architect, Microsoft	Addressing Shingled Magnetic Recording Drives with Linear Tape File System Albert Chen Western Digital Jim Malina Technologist, Western Digital	SNIA NVM Programming Model Paul von Behren Software Architect, Intel
2:00 - 2:50	Resilience at Scale in the Distributed Storage Cloud Alma Riska Consulting Software Engineer, EMC	**Testing Methodologies** (San Tomas / Lawrence) Testing iSCSI / SCSI Protocol Compliance Using Libiscsi Ronnie Sahlberg Site Reliability Engineer, Google	A Status Report on SMB Direct (RDMA) for Samba Richard Sharpe Samba Team Member, Panzura	**InfiniBand Architectural** Overview David Deming President, Solution Technology	NVMe based PCIe SSD Validation – Challenges and Solutions Apurva Vaidya Technical Architect, iGATE Sachin Gade Technical Specialist, iGATE
2:50 - 3:05			Break (Mezzanine)		

CLOUD INTEROPERABILITY PLUGFEST

Monday, 8:00 AM - Thursday 12:00 PM • Central Room • Underwritten by SNIA Cloud Storage Initiative

A part of the Cloud Interoperability Week, the 10th Cloud Interoperability Plugfest is open to participation by other standards organizations, individual developers, vendors and software framework projects. The plugfest is designed to promote interchange of methods, tools and techniques that are defined by the plugfest participants as being suitable to test features such as portability, interoperability, ease of use, self-consistency and security when used in a cloud setting.

TUESDAY AGENDA

	Cloud (Lafayette)	Testing Methodologies (San Tomas / Lawrence)	SMB2/SMB3 (Cypress)	Hardware (Stevens Creek)	Solid State (Winchester)
3:05 - 3:55	CDMI and Scale Out File System for Hadoop Philippe Nicolas Director of Product Strategy, Scality	Message Analysis and Visualization in Heterogeneous Environments Paul Long Senior Program Manager, Microsoft	Samba Scalability and Performance Tuning Volker Lendecke Samba Team / SerNet	**Infiniband Verbs and** Memory Management - RDMA David Deming President, Solution Technology	**TBF: A Memory-Efficient** Replacement Policy for Flash-based Caches Biplob Debnath Research Staff Member, NEC Laboratories America
4:05 - 4:55	Transforming Cloud Infrastructure to Support Big Data Dr. Ying Xu Principle Engineer, Aspera	A Software Based Fault Injection Framework for Storage Server Vinod Eswaraprasad Lead Architect, Wipro Technologies	Scaled RDMA Performance & Storage Design with Windows Server SMB 3.0 Dan Lovinger Principal Software Engineer, Microsoft Spencer Shepler Principal Architect, Microsoft	Introduction to HP Moonshot Tracy Shintaku Distinguished Technologist, Server Engineering R&D, Hewlett Packard	Delivering Nanosecond-Class Persistent Storage Steffen Hellmold Vice President of Marketing, Everspin Technologies
5:00 - 7:00	Welcome Reception - Sponsor Showcase (Mezzanine)				
7:00 - 8:00	"Birds of a Feather" Meetings				

SNIA SMB2/SMB3 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 SMB2/SMB3 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.

2013 SNIA SMB2/SMB3 Plugfest Participants (As of 9/10/13)

Apple, catalyst IT Limited, Citrix, Codenomicon, Dell, EMC², HP, HvNAS Networks, IBM, IT Access, ixia, MathWorks, Microsoft, NetApp, nexenta, Oracle, Quantel, Quantum, redhat, Samba, suse, SwiftTest, Symas Corporation, Visuality Systems

7 WEDNESDAY AGENDA

7:30 - 8:45	Continental Breakfast and Networking (Mezzanine)
8:45 - 9:00	Introduction to the Day and Housekeeping Announcements Don Deel, SNIA Technical Council Chairman
9:00 - 9:45	Storage Infrastructure Performance Validation at Go Daddy – Best Practices from the World's #1 Web Hosting Provider Julia Palmer, Manager of Virtual Storage and Backup Teams, Go Daddy Justin Richardson, Senior Storage Engineer, Go Daddy
9:45 - 10:25	**Migrating to Cassandra in the Cloud, the Netflix Way** Jason Brown, Senior Software Engineer, Netflix
10:25 - 10:40	Break (Mezzanine)
10:40 - 11:20	Platform as a Service and the Newton: One of These Things is Just Like the Other Gerald Carter, Senior Consulting Engineer, EMC, Isilon Storage Division
11:20 - 12:15	**Software Defined Storage Roundtable** Moderator: Leah Schoeb, Senior Partner, Evaluator Group Panelists: Alex McDonald, Standards & Industry Associations Group, CTO Office, NetApp Alan Yoder, Senior Manager, Storage Standards, Huawei Technologies Dale Degan, Worldwide Senior Marketing Manager, Hewlett-Packard Lazarus Vekiarides, Entrepreneur and Technology Executive Rick Walsworth, Director Product Marketing, EMC Satyam Vaghani, CTO, PernixData
12:15 - 1:15	Lunch and Sponsor Showcase (Mezzanine)

	Long Term Retention (San Tomas / Lawrence)	NFS (Cypress)	Hot Topics (Winchester)	New Thinking (Stevens Creek)	Performance (Lafayette)
1:15 - 2:05	Combining SNIA Cloud, Tape and Container Format Technologies for the Long Term Retention of Big Data Sam Fineberg, Distinguished Technologist, Hewlett Packard Simona Rabinovici-Cohen, Research Staff Member, IBM	NFS on Steroids: Building Worldwide Distributed File System Gregory Touretsky, Solutions Architect, Intel	Can Your Storage Infrastructure Handle the Coming Data Storm? Amritam Putatunda, Technical Marketing Engineer, Ixia	Screaming Fast Galois Field Arithmetic Using Intel SIMD Instructions Ethan Miller, Solutions Architect, Professor, Director of the NSF Industry/University Cooperative Research Center, University of California	High-Throughput Cloud Storage Over Faulty Networks Yogesh Vedpathak, Software Developer, Cleversafe
2:15 - 3:05	Best Practices, Optimized Interfaces, API's Designed for Storing Massive Quantities of Long Term Retention Data Stacy Schwarz-Gardner, Strategic Technical Architect, Spectra Logic	Implementing NFSv3 in Userspace: Design and Challenges Tai Horgan, Software Engineer, EMC	OpenStack Cloud Storage Dr. Sam Fineberg, Distinguished Technologist, Hewlett-Packard	NV-Heaps: Making Persistent Objects Fast and Safe with Next-Generation Non-Volatile Memories Joel Coburn, Software Engineer, Google/UCSD	Forget IOPS: A Proper Way to Characterize & Test Storage Performance Peter Murray, Senior Product Specialist, SwiftTest

"BIRDS OF A FEATHER" MEETINGS

NVM Programming Model - Next Steps • Tuesday, 7:00 - 8:00 PM • Winchester Room
The NVM Programming TWG recently celebrated its first birthday and is finalizing its first publication. We looking for suggestions from TWG members and non-members on NVM software extensions for future publications. The BOF includes a short overview of the TWG and NVM Programming Model specification, followed by a round-table suggestion of future work items.

Microsoft File Server Protocol Patent Licensing • Tuesday, 7:00 - 8:00 PM • Stevens Creek Room
The Microsoft SMB2, SMB3, and File Server Protocols specifications are freely available on the Microsoft website, but a patent license may be required for any commercialization. Come to this BoF to get more information.

WEDNESDAY AGENDA

	Big Data (San Tomas / Lawrence)	NFS (Cypress)	Data Management (Winchester)	New Thinking (Stevens Creek)	Performance (Lafayette)
3:05 - 3:20			Break (Mezzanine)		
3:20 - 4:10	Can Storage Fix Hadoop? John Webster Senior Partner, Evaluator Group	pNFS, NFSv4.1, FedFS and Future NFS Developments Alex McDonald Standards & Industry Associations Group, CTO Office, NetApp	Data Deduplication as a Platform for Virtualization and High Scale Storage Adi Oltean Principal Software Design Engineer, Microsoft Sudipta Sengupta Sr. Researcher, Microsoft	LazyBase: Trading Freshness for Performance in a Scalable Database Brad Morrey Senior Research Scientist, HP Labs	Improvements in Storage Energy **Efficiency via Storage** Subsystem Cache and Tiering Chuck Paridon Storage Performance Architect, Hewlett Packard Herb Tanzer Storage Hardware Architect, Hewlett Packard
4:20 - 5:10	Hadoop: Embracing Future Hardware Sanjay Radia Co-founder, Hortonworks Suresh Srinivas Hortonworks	pNFS Directions Matt Benjamin Founder, Cohort LLC Peter Honeyman Founder, CohortFS LLC	Virtual Machine Archival System Parag Kulkarni VP Engineering, Calsoft Dr. Anupam Bhide CEO, Co-Founder, Calsoft	GraphChi: Large-Scale Graph Computation on Just a PC Aapo Kyrola Ph.D. Student, CMU Computer Science Department	Lessons Learned Tuning HP's SMB Server with the FSCT Benchmark Bret McKee Distinguished Technologist, Hewlett Packard Vinod Eswaraprasad Lead Architect, WiPro Technologies
7:00 - 9:00			"Birds of a Feather" Meetings		

"BIRDS OF A FEATHER" MEETINGS (CONTINUED)

Green Storage - The Big Picture - Tuesday, 7:00 - 8:00 PM • Cypress Room

"The most expensive storage purchased is that which causes the deployment of another Data Center."
- George Crump, President & Founder Storage-Switzerland

In a world of more, more, more, using 'less' to store all of it, is a crucial skill, which translates to a real competitive advantage for an organization. Join us at the MESS (Media/Entertainment & Scientific Storage) meetup as our panel of experts discuss the key techniques for reducing the power, cooling, space, and networking impact of storage, using new paradigms like: IO density metrics; Geo-dispersal of data; Next-generation storage pods; Self-healing protection algorithms …together contributing heartily to a simple goal 'A Lower COST Footprint.' Improve your company's bottom line by attending the September MESS meet up!

Building a Linux Storage Appliance with Data Optimization • Wednesday, 7:00 - 8:00 PM • Winchester Room

Data deduplication and compression are no longer storage optimizations relegated to backup. They have become mainstream in primary and high performance (flash) storage. In this BOF session, we will discuss how to build a Linux storage appliance using standard Linux components (XFS, LVM2, and Linux iSCSI) and Permabit Albireo Virtual Data Optimizer (VDO). Whether you are designing cloud storage, backup solutions, or high performance flash arrays, this discussion will show you how to build a storage-optimized product in matter of hours.

Cloud Application Management for Platforms (CAMP) • Wednesday, 7:00 - 8:00 PM • Stevens Creek Room

There are multiple commercial PaaS offerings in existence using languages such as Java, Python and Ruby and frameworks such as Spring and Rails. Although these offerings differ in such aspects as programming languages, application frameworks, etc., there are inherent similarities in the way they manage the applications that are deployed upon them. Cloud Application Management for Platforms (CAMP) specifies deployment artifacts and a RESTful API designed to both ease the task of moving applications between PaaS platforms as well as provide an interoperable mechanism for managing PaaS-based applications in a way that is language, framework, and platform neutral.

pNFS Open Discussion • Wednesday, 8:00 - 9:00 PM • Stevens Creek Room

The last two years have been busy ones for pNFS. This BOF will provide an opportunity for NFSv4 and pNFS implementors, users, and interested parties to come together for open discussion. Potential topics for discussion include details of current pNFS implementations, pNFS scalability, and future directions for NFSv4 and pNFS.

THURSDAY AGENDA

7:30 - 8:15	Continental Breakfast (Mezzanine)			
8:15 - 8:30	Introduction to the Day, Housekeeping and Other Announcements Wayne Adams, SNIA Board Chairman Don Deel, SNIA Technical Council Chairman Gilda Foss, SNIA Analytics and Big Data Committee Chairperson			
8:30 - 9:20	Closing Keynote: Worlds Colliding: Why Big Data Changes How to Think about Enterprise Storage Addison Snell, CEO, Intersect360 Research			
	Security (San Tomas)	**End User Viewpoint** (Lawrence)	**Storage Utilities** (Lafayette)	**Workloads** (Cypress)
9:30 - 10:20	Trusted Computing Technologies for Storage Dr. Michael Willett Storage Security Strategist, Samsung	Demand for Storage Systems from a Customer Viewpoint in Japan Satoshi Uda Assistant Professor, Japan Advanced Institute of Science and Technology	Building the Public Storage Utility Wesley Leggette Storage Architect, Cleversafe	Tunneling SCSI over SMB: **Shared VHDX files for Guest** Clustering in Windows Server 2012 R2 Jose Barreto Principal Program Manager, Microsoft Matt Kurjanowicz Software Development Engineer, Microsoft
10:20 - 10:30	Break (Mezzanine)			
	Security (San Tomas)	**Development Methodologies** (Lawrence)	**Storage Management** (Lafayette)	**Workloads** (Cypress)
10:30 - 11:20	Multi-vendor Key Management – Does It Actually Work? Tim Hudson Technical Director, Cryptsoft	A Method to Establish Concurrent Rapid Development Cycle and High Quality in a Storage Array System Environment M. K. Jibbe Director of Quality Architect Team NetApp APG Products, NetApp Kuok Hoe Tan QA Architect, NetApp	Event Trees for Storage Monitoring and Management Rajesh Radhakrishnan Chief Architect and Principal Consultant, Zygous	An SMB3 Engineer's View of Windows Server 2012 Hyper-V Workloads Gerald Carter Sr. Consulting Software Engineer, EMC
11:25 - 12:15	Matching Security to Data Threats – More is Not Better, but Less Can be Bad Chris Winter Director Product Management, SafeNet	Code Coverage as a Process Aruna Prabakar Software Engineer, EMC Niranjan Page Engineering Manager, EMC	A Method to Backup and **Restore Configuration** Settings for Each and Every Component in the SAN Environment Using SMI-S Dhishankar Sengupta Test Architect, NetApp Krishanu Dhar Test Lead, NetApp	Direct NFS – Design Considerations for Next-gen NAS Appliances Optimized for Database Workloads Gurmeet Goindi Principal Product Manager, Oracle Akshay Shah Principal Software Engineer, Oracle
12:00	Cloud Interoperability Plugfest Concludes			
12:15	Sessions Conclude			
1:00	SNIA SMB2/SMB3 and iSCSI Plugfests Conclude			

TUESDAY

The Impact of the NVM Programming Model - Andy Rudoff, Intel, SNIA NVM Programming TWG

As exciting new Non-Volatile Memory (NVM) technologies emerge, the SNIA NVM Programming Technical Workgroup (TWG) has been working through the applicable programming models. Andy will talk about the impact these programming models will have on the industry, focusing especially on the more disruptive areas of NVM like Persistent Memory.

Windows Azure Storage – Scaling Cloud Storage - Andrew Edwards, Principal Architect, Windows Azure Storage, Microsoft

In today's world that is increasingly dominated by mobile and cloud computing application developers require durable, scalable, reliable, and fast storage solutions like Windows Azure Storage. This talk will cover the internal design of the Windows Azure Storage system, how it is engineered to meet these ever growing demands, and lessons learned from operating at scale.

Optical Storage Technologies: The Revival of Optical Storage

Ken Wood, CTO – Technology & Strategy, Office of Technology and Planning, Hitachi Data Systems

Optical storage is seeing a resurgence in new industry verticals for it's improved and unique preservation and environmental qualities. Recent developments have increased capacities and functionality while maintaining decades of backwards compatibility. This is due to the wide range of industries and markets that support this medium.

Hypervisors and Server Flash - Satyam Vaghani, CTO, PernixData

Hypervisors and server flash is an important but inconvenient marriage. Server flash has profound technology and programming implications on hypervisors. Conversely, various hypervisor functions make it challenging for server flash to be adopted in virtualized environments. In this talk, we will present specific hypervisor design areas that are challenged by the new physics of storage presented by server flash, and possible solutions. We will discuss the motivation and use cases around a software layer to virtualize server-flash and make it compatible with clustered hypervisor features like VM mobility, high availability, distributed VM scheduling, data protection, and disaster recovery. Finally, we will present some empirical results from one such flash hypervisor (FVP) implemented at PernixData, and its potential long term impact on data center storage design.

WEDNESDAY

Storage Infrastructure Performance Validation at Go Daddy – Best Practices from the World's #1 Web Hosting Provider

Julia Palmer, Manager of Virtual Storage and Backup Teams, Go Daddy

Justin Richardson, Senior Storage Engineer, Go Daddy

Infrastructure is evolving rapidly these days, especially for storage professionals. A flurry of new technologies such as SSDs and tiering promise faster, cheaper, and more cost-effective storage solutions. Storage-as-a-service offers a new blueprint for flexible, optimized storage operations. Go Daddy is taking full advantage of these opportunities with continual innovation. Attend this presentation to hear how Go Daddy utilized an innovative new approach to storage infrastructure validation that enabled them to accelerate the adoption of new technologies and reduce costs by nearly 50% while maintaining 99.999% uptime for their 28 PB of data. The new process empowers Go Daddy with the insight they need to optimize both service delivery and vendor selection. Audience members will also learn how to evaluate storage workloads and identify potential performance and availability problems before they are experienced by end users.

Migrating to Cassandra in the Cloud, the Netflix Way - Jason Brown, Senior Software Engineer, Netflix

Netflix grew up using the traditional enterprise model for scaling: monolithic web application on top of a monolithic database in a single datacenter, buying bigger boxes, stuffing more user data into session memory. It all worked great when Netflix had less than 1 million customers (and rapidly growing). Then one day that model failed us, miserably. The single-point-of-failure bug hit us hard, and we were hobbled for days. Since then, Netflix has reinvented it's technology stack from top to bottom - abandoning the single, monolithic web application for tiered distributed services, as well as moving beyond our SPOF database to more resilient architectures. In this talk Jason will be discussing my involvement with Cassandra at Netflix, first as a user of this new system, then as a developer of it. He will discuss how we migrated from our traditional datacenter to the cloud, how we store and backup data, and the problems of rapidly scaling out a persistence layer under a burgeoning distributed architecture.

THURSDAY

Platform as a Service and the Newton: One of These Things is Just Like the Other

Gerald Carter, Senior Consulting Software Engineer, EMC, Isilon Storage Division

Platform as a Service (PaaS) offerings, like the Apple Newton, launched at a time when technology had not matured to the point necessary to cross the chasm to the early majority and into mass markets. Successes did exist, but were limited to specialized applications already targeted at a vendor's existing platform. Infrastructure as a Service (IaaS), and its successor of software defined entities, are necessary intermediate steps towards decoupling application development from operational overhead. This talk will explore what the future will look like when developers can once again focus solely on applications and interfaces and turn a blind eye to operations

Worlds Colliding: Why Big Data Changes How to Think about Enterprise Storage - Addison Snell, CEO, Intersect360 Research

Addison Snell of Intersect360 Research will present an overview of how Big Data trends have changed some fundamental drivers in acquiring, architecting, and administering enterprise storage. With the majority of Big Data implementations coming from in-house development -- Hadoop is just the tip of the iceberg -- storage developers will find themselves taking on new roles that are defined by performance and scalability as much as reliability and uptime. Learn why high performance computing technologies like parallel file systems and InfiniBand could cross the Rubicon into enterprise, while an IT darling like Cloud might not play.

Breakout descriptions are listed chronologically.

MONDAY SESSIONS

CLOUD TRACK

CDMI, THE KEY COMPONENT OF SCALITY OPEN CLOUD ACCESS

Monday 8:30 - 9:20

Announced during SNIA SDC 2012, Scality Open Cloud Access, aka OCA, was the first converged access methods between file and object mode, from local or remote site. Dewpoint, the Scality CDMI server, demonstrated during previous 3 CDMI annual plugfests during SDC, continues to be the pivotal component of this strategy. During last year, Scality leverage CDMI to build the Scality solution for Hadoop and plan to announce a few other innovations.

CDMI FEDERATIONS, YEAR 4

Monday 9:30 - 10:20

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 year's SDC presentation, bi-directional federation between two CDMI-enabled clouds was discussed and demonstrated. For year four, we will discuss how CDMI federation when combined with CDMI versioning, enables mobile and web-based applications to synchronize data with clouds, effectively allowing clients to create "mini-clouds" local to the client. This architectural approach allows clients to easily store, cache, and merge cloud-resident content, provide disconnected operation, and provides a foundation for application-specific conflict resolution.

LTFS AND CDMI - TAPE FOR THE CLOUD

Monday 10:35 - 11:25

LTFS tape technology provides compelling economics for bulk cloud storage and transportation of data. This session provides an overview of the use cases identified by the joint LTFS and Cloud Technical Working Group, including when tape provides lower-cost alternatives to network and disk-based transportation, and when tape provides lower-cost alternatives to disk-based storage and archiving. This session will introduce standardization efforts underway to allow for simple tape-based bulk data transport to, from, and between clouds, and standardized of how to store rich object data on standard LTFS tapes.

PROFILE BASED COMPLIANCE TESTING OF CDMI: APPROACH, CHALLENGES & BEST PRACTICES

Monday 11:35- 12:25

Cloud Data Management Interface Specifications are now moving towards profile based categories. This is due to the increased focus where organizations are planning to adopt profile based CDMI in their products, example Service, ID and Self Storage Management profiles. TCS has been focusing on implementing 'CDMI Automated Test Suite' and with new developments is focusing towards incorporating profiling to the same. In this proposal we will share the approach and challenges for testing of profile based scenarios towards CDMI profile based compliance, of the cloud products. Also, we will share additional challenges / learning towards testing of CDMI Products for compliance. These learning's will serve as a best practice and a ready reference for other organizations in developing their own CDMI product suit.

OPEN-SOURCE CDMI-COMPLIANT PROXY: "STOXY"

Monday 1:30 - 2:20

We present our on-going effort in development of a open-source server (Stoxy - STOrage proXY), which exposes CDMI-compliant interface on the frontend and allows to store and manage data on several public and private cloud backends, incl. AWS and MS Azure. The work is based on a CDMIProxy prototype developed in EU VENUS-C project with a continuos development effort sponsored by EGI Inspire project as well as commercial companies. Presentation will cover the architecture of Stoxy and highlight certain key components, esp. related to the data streaming. In addition, we shall give initial experience from integrating Stoxy with the other products (cloud managers, storage servers) inside the EGI Federated cloud task force.

LESSONS LEARNED IMPLEMENTING CROSS-PROTOCOL COMPATIBILITY LAYER

Monday 2:30 - 3:20

Over the past year, we have integrated our storage solution with a number of cloud and object storage APIs, including Amazon S3, WebDAV, OpenStack, and HDFS. While these protocols share much commonality, they also differ in meaningful ways which complicates the design of a cross-protocol compatibility layer. In this presentation, we detail how the various storage protocols are the same, how they differ, and what design decisions were necessary to build an underlying storage API that meets the requirements to support all of them. Further, we consider the lessons learned and provide recommendations for developing cloud storage APIs such as CDMI.

COSBENCH: A BENCHMARK TOOL FOR CLOUD STORAGE

Monday 3:35- 4:25

With object storage services becoming increasingly accepted as one new offering comparing to traditional file or block systems, it is important to effectively measure the performance of these services. Thus people can compare different solutions or tune their systems for better performance. However, little has been reported on this specific topic as yet. To address this problem, we developed COSBench (Cloud Object Storage Benchmark), a benchmark tool that we are currently working on inside Intel for cloud object storage services. In addition, we will share the status for CDMI supporting, and share results of the experiments we have performed so far.

ARCHITECTING AN ENTERPRISE STORAGE PLATFORM USING OBJECT STORES

Monday 4:35 - 5:25

While object storage systems such as S3 and Swift are exhibiting rapid growth, there is still an impedance mismatch between their feature set and enterprise requirements. This talk dives into the design and architecture of MagFS: a strongly consistent and multi-platform distributed file system that layers itself on top of multiple object storage systems. In particular, it covers the challenges of using eventually consistent object stores, optimizing both data and metadata traffic for wide-area network communication and mobile devices, and how MagFS delivers an on-premises security model while still being able to leverage off-premises storage. This talk also discusses how specific enterprise requirements have influenced the technical design of MagFS and some of the surprises we encountered during our design and implementation.

FILE SYSTEMS TRACK

ADVANCEMENTS IN STORAGE AND FILE SYSTEMS IN WINDOWS 8.1

Monday 8:30 - 9:20

There are some advances, refinements, and improvements in Windows File Systems coming, which we'll be able to talk about at SDC 2013.

HDFS - WHAT IS NEW AND FUTURE

Monday 9:30 - 10:20

Hadoop 2.0 offers significant HDFS improvements: new append-pipeline, federation, wire compatibility, NameNode HA, performance improvements, etc. We describe these features and their benefits. We also discuss development that is underway for the next HDFS release. This includes much needed data management features such as Snapshots and Disaster Recovery. We add support for different classes of storage devices such as SSDs and open interfaces such as NFS; together these extend HDFS as a more general storage system. As with every release we will continue improvements to performance, diagnosability and manageability of HDFS.

MULTIPROTOCOL LOCKING AND LOCK FAILOVER IN ONEFS

Monday 10:35 - 11:25

This talk will examine the details on how multiprotocol locking is implemented in a distributed clustered file system such as Isilon's OneFS and also looks into the existing lock failover implementation in OneFS for NFS and how it can be extended for implementing lock failover for SMB3. A clustered file system such as Isilon's OneFS can have multiple clients accessing the server using different protocols such as SMB and NFS. A robust and efficient distributed lock manager is necessary is necessary to achieve both protocol correctness and data consistency in the presence of multi-protocol access to data/files. We also need a failover mechanism to implement the failover semantics of these protocols so that the locks are not lost even when a node in the cluster goes down. This talk will examine the details of such a locking mechanism in OneFS.

SNAPSHOTS FOR IBRIX - HIGHLY DISTRIBUTED SEGMENTED PARALLEL FS

Monday 11:35 - 12:25

This presentation explores designing 'native snapshots' for scale-out segmented parallel file systems (Ibrix). An appropriate model of snapshots requires flexibility and fluidity to allow easy selection of objects, reliability to assure logical unity of such subsets. We scale linearly adding servers and segments fundamentally by limiting the number of objects participating in operations and de-centralizing control over meta-data. With snapshots, associated state transition has to affect not only directly referenced objects, but has to be immediately propagated to all the descendant nodes controlled by a large number of other servers. We also look into recovery, achieving quick rollback logically resetting state of the subspace to a desired point in time and allowing corresponding longer running cleanup processes to finish in the background.

SNAPSHOT CAUTERIZATION

Monday 1:30 - 2:20

Snapshot cauterization and MetaSnaps When a user takes a snapshot of a file-system, it captures all the data, some of which the user may not want to retain. Presently there are no mechanisms for a user to delete this data without deleting the whole snapshot. We present methods to cauterize such unwanted data from a snapshot, to reclaim space. This technique can be used to build more features which will be useful for file system analytics.

CLUSTER SHARED VOLUMES

Monday 2:30 - 3:20

Cluster Shared Volumes is a cluster file system for the Windows Hyper-V and File Server workloads. It enables concurrent access to volumes and files from any node in a Windows Server Failover Cluster. In this session, we will describe how Cluster Shared Volumes leverages and extends existing Windows technology, such as NTFS for metadata and storage allocation, SMB 3.0 for high-speed interconnect, Volume Snapshot Service for distributed backups, oplocks for cache coherency, and failover clusters for multi-node coordination.

BALANCING STORAGE UTILIZATION ACROSS A GLOBAL NAMESPACE

Monday 3:35 - 4:25

This talk provides a taxonomy of filesystem and storage development from 1979 to the present with the BSD Fast Filesystem as its focus. It describes the early performance work done by increasing the disk block size and by being aware of the disk geometry and using that knowledge to optimize rotational layout. With the abstraction of the geometry in the late 1980's and the ability of the hardware to cache and handle multiple requests, filesystems performance ceased trying to track geometry and instead sought to maximize performance by doing contiguous file layout. Small file performance was optimized through the use of techniques such as journaling and soft updates. By the late 1990's, filesystems had to be redesigned to handle the ever growing disk capacities. The addition of snapshots allowed for faster and more frequent backups. The increasingly harsh environment of the Internet required greater data protection provided by access-control lists and mandatory-access controls. The talk concludes with a discussion of the addition of symmetric multi-processing support needed to utilize all the CPUs found in the increasingly ubiquitous multi-core processors.

SCALE-OUT STORAGE SOLUTION

Monday 4:35 - 5:25

Today, data is growing at an exponential rate and the need to provide an efficient storage mechanism has become more critical than ever. In this presentation, we will discuss about a scale out storage solution intended to address small and medium businesses in a cost effective manner. This is a Linux based software-only solution that works on commodity hardware. It is a POSIX compliant solution and provides file storage through CIFS/NFS interfaces. The entire solution is designed to have small footprint and easy installation on available Linux machines. This paper presents technical details of the solution and implementation challenges. In addition, the paper will also discuss about tools and techniques used to test scale out storage product.

SMB2/SMB3 TRACK

SMB3 MEETS LINUX: THE LINUX KERNEL CLIENT

Monday 8:30 - 9:20

SMB3 support has been merged into Linux ever since the 3.8 kernel. What have we learned? And Why use SMB3? Learn what SMB3 features are available to Linux users, and when you should consider using SMB3 instead of other protocols when running Linux, and how to configure optional features and improve performance. In addition, there will be demonstration of some of the newer features, and a description of what improvements are expected in the coming months.

13 BREAKOUT SESSIONS

MAPPING SMB ONTO DISTRIBUTED STORAGE

Monday 9:30 - 10:20

The SMB protocol is, and always has been, an extension of the Operating Systems and File Systems on which it was designed to run. Yes, that means DOS/FAT, OS2/HPFS, and Windows/NTFS. Building an SMB server on any other platform requires a lot of special handling, sort of like the special pounding of a square peg into a round hole with a finely tuned sledgehammer. The design of many newer file systems, particularly object, cluster, and distributed storage systems, complicates matters even more by "relaxing" adherence to standard semantics. This presentation will highlight a number of ways in which these "relaxed fit" file systems clash with SMB expectations, and will provide some examples of ways to bridge the gap.

PIKE - MAKING SMB TESTING LESS TORTUROUS

Monday 10:35 - 11:25

Pike is a new Python protocol testing library for SMB2 and SMB3. It will be made publicly available along with a collection of tests under an open-source license. Pike has a simple and extensible architecture. It aims to make common case scenarios concise while still allowing deep control over message construction and dispatch when necessary. This talk will explore the core architecture of Pike and how to extend it with new features and tests. Along the way, we'll see how the dynamism and expressiveness of Python make it a great environment for protocol testing.

EXPLOITING THE HIGH AVAILABILITY FEATURES IN SMB 3.0 TO SUPPORT SPEED AND SCALE

Monday 11:35 - 12:25

Microsoft have made massive changes in version 3.0 of SMB protocol, many of which contribute towards SMB 3.0 offering high availability (HA) for use in the data centre. This talk will present the results of investigations into how these innovations can be exploited for improved I/O speed & architectural scale. The presentation will first look at the needs of HA in a NAS protocol. It will then offer an insight into why features were added to SMB 3.0, providing technical analysis of the protocol itself and live demonstrations of noted features using the authors own implementation of an SMB server.

SMB3 UPDATE

Monday 1:30 - 2:20

The past year has seen multiple companies and teams release SMB3 solutions, and many customers deploy them into production. This talk will look at some upcoming minor adjustments to SMB3 based on lessons learned, and cast forward for what might come next.

IMPLEMENTATION OF SMB3.0 IN SCALE-OUT NAS

Monday 2:30 - 3:20

SMB 3.0 features several improvements over the CIFS protocol on which it is based. Continuous Availability achieves high storage availability through transparent failover. Copy Offload and RDMA boost performance dramatically in Windows Server 2012. These features, along with Multi-channels, are quite attractive to customers. However, these features together pose implementation challenges in a scale-out NAS context. We discuss our experience implementing SMB3 on a clustered system without compromising functionality or performance.

1 S(a) 2 M 3 B(a) 4

Monday 3:35 - 4:25

Samba 4.0 has been released in December 2012. It is is the first release of Samba that features the Active Directory Compatible Domain Controller. But it is also a very important file server release: Samba 4.0 ships with SMB 3 enabled by default. In this talk, I will describe the subset of SMB3 that Samba 4.0 already offers, and report about the work in progress in the implementation of more SMB 3 features, giving an outlook of what to expect in upcoming Samba 4.1.

BLOCK STORAGE TRACK

SCSI STANDARDS AND TECHNOLOGY UPDATE

Monday 8:30 - 9:20

SCSI continues to be the backbone of enterprise storage deployments and 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 the status of 12Gb/s SAS staging, advanced connectivity solutions such as MultiLink SAS™ and cover SCSI Express, a new transport of SOP (SCSI over PCIe). Presenters will also provide updates on new SCSI feature such as atomic writes and remote copy.

SCSI AND FC STANDARDS UPDATE

Monday 9:30 - 10:20

This session will examine the ongoing work of the T10 Standards committee that defines SCSI and SAS; the T11 Standards committee that defines Fibre Channel (including FCoE); and the IETF working group that defines iSCSI. It will provide an introduction to the new capabilities and features being proposed for SCSI, iSCSI, SAS, and Fibre Channel, including SCSI Atomic operations, Scatter/Gather operations, Profiles, Time Limited operations, Large Scale virtualization enhancements, SMR technology, 12Gb SAS, new iSCSI features, 32GFC, 128GFC, FCoE VN2VN, Automatic FC Zoning, and Energy Efficient FC.

EXTENDING SAS CONNECTIVITY IN THE DATA CENTER

Monday 10:35 - 11:25

Serial Attached SCSI (SAS) is the connectivity solution of choice for disk drives and JBODs in the data center today. SAS connections are getting faster while storage solutions are getting larger and more complex. Data center configurations and disaster recovery solutions are demanding longer cable distances. This is making it more and more difficult or impossible to configure systems using passive copper cables. This presentation discusses the application, limitations and performance of passive copper, active copper and optical SAS cabling options available today and those likely to be available in the next few years.

FCoE DIRECT END-NODE TO END-NODE (AKA FCoE VN2VN)

Monday 11:35 - 12:25

A new concept has just been accepted for standardized in the Fibre Channel (T11) standards committee; it is called FCoE VN2VN (aka Direct End-Node to End-Node). The FCoE standard which specifies the encapsulation of Fibre Channel frames into Ethernet Frames is being extended to permit FCoE connections Directly between FC/FCoE End-Nodes. The tutorial will show the Fundamentals of the extended FCoE concept that permits it to operate without FC switches or FCoE Switches (aka FCF) and will describe how it might be exploited in Small, Medium or Enterprise Data Center environments -- including the "Cloud" IaaS (Infrastructure as a Service) provider environments.

OBJECT STORAGE TRACK

HUAWEI SMARTDISK BASED OBJECT STORAGE UDS

Monday 1:30 - 2:20

Huawei UDS (Universal Distributed Storage) is a massive object-storage system that offers enterprises and service providers a comprehensive solution to data explosion difficulties. Through its proprietary technology, Huawei UDS system in one Data Center is able to scale out to 25k SmartDisk(s), which is composed of one ARM chip and one hard drive. As a result, it provides competitive scalability, reliability, and cost.

ARCHITECTING BLOCK AND OBJECT GEO-REPLICATION SOLUTIONS WITH CEPH

MONDAY 2:30 - 3:20

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 fully open source distributed object store, network block device, and file system designed for reliability, performance, and scalability from terabytes to exabytes. The Ceph architecture was initially designed to accommodate single data center deployments, where low latency links and synchronous replication were an easy fit for a strongly consistent data store. For many organizations, however, storage systems that span multiple data centers and geographies for disaster recovery or follow-the-sun purposes are an important requirement. This talk will give a brief overview of the Ceph architecture, and then focus on the design and implementation of asynchronous geo-replication and disaster recovery features for the RESTful object storage layer, the RBD block service, and Ceph's underlying distributed object store, RADOS. The fundamental requirements for a robust georeplication solution (like point in time consistency) and the differing requirements for each storage use-case and API and the implications for the asynchronous replication strategy will be discussed.

HARDWARE TRACK

PCI EXPRESS AND ITS INTERFACE TO STORAGE ARCHITECTURES

MONDAY 3:35 - 4:25

PCI Express Gen2 and Gen3, IO Virtualization, FCoE, SSD, PCI Express Storage Devices are here. What are PCIe Storage Devices – why do you care? This session describes PCI Express, Single Root IO Virtualization and the implications on FCoE, SSD, PCIe Storage Devices and impacts of all these changes on storage connectivity, storage transfer rates. The potential implications to Storage Industry and Data Center Infrastructures will also be discussed.

PCI EXPRESS IO VIRTUALIZATION OVERVIEW

MONDAY 4:35 - 5:25

PCI Express IO Virtualization Specifications working with System Virtualization allowing multiple operating systems running simultaneously within a single computer system to natively share PCI Express Devices. This session describes PCI Express, Single Root and Multi Root IO Virtualization. The potential implications to Storage Industry and Data Center Infrastructures will also be discussed.

DISTRIBUTED STORAGE TRACK

GETTING THE MOST OUT OF ERASURE CODES

MONDAY 8:30 - 9:20

Erasure codes are a recent addition to many storage technologies. They provide increased reliability with less overhead. Yet, they are not without downsides. Selecting the best parameters for erasure coding is a complex optimization problem. As one varies the threshold, write threshold, number of pieces, system capacity, and site count, there may be drastic effects on reliability, availability, storage overhead, rebuilding cost, and CPU expense. Selecting erasure code parameters without weighed consideration may have catastrophic results. In this presentation we present the techniques we developed and use for designing erasure coded systems with the best combination of storage overhead, computational efficiency, reliability, availability, and rebuilding cost for any given system constraints. Finally, we introduce some advanced techniques for reducing rebuilding cost.

LRC ERASURE CODING IN WINDOWS STORAGE SPACES

MONDAY 9:30 - 10:20

RAID is the standard approach for fault tolerance among multiple disk drives and has been around for decades. However, new hardware trends, including the advent of hard disk drives (HDDs) with huge capacity, widely adoption of solid state drives (SSDs) with fast I/O, etc., have created new opportunities to optimize fault tolerance schemes. Windows now introduces a new fault tolerance scheme in its Storage Spaces technology. The new scheme is developed based on a novel erasure coding technology, called Local Reconstruction Code (LRC). Compared to RAID under same durability metric, LRC significantly reduces rebuild time, while still keeping storage overhead very low. In addition, LRC offers much more flexibility in balancing rebuild time and storage overhead. The presentation will provide an overview of the Windows Storage Spaces technology, cover the design of its fault tolerance mechanism, discuss the implementation of LRC in detail and share experiences learned from real-world workloads.

DISTRIBUTED STORAGE ON LIMPING HARDWARE

MONDAY 10:35 - 11:25

It is easy to design storage systems that assume nothing bad ever happens. It is marginally harder to design one that assumes nodes are either available or not. What is difficult is designing storage systems that handle how nodes fail in the real world. Such "limping nodes" may respond slowly, occasionally, or unpredictably; they are neither entirely failed nor entirely healthy. This presentation covers the mechanisms we developed for dealing with limping nodes in a distributed storage system. These techniques allow limping nodes to be tolerated with negligible impact on performance, latency, or reliability. We introduce some of the intelligent writing techniques we created for this purpose, which include: write thresholds, impatient writes, optimistic writes, real-time writes, and lock-stealing writes.

METHOD TO ESTABLISH A HIGH AVAILABILITY AND HIGH PERFORMANCE STORAGE ARRAY IN A GREEN ENVIRONMENT

MONDAY 11:35 - 12:25

Shift Left is a combination of a change in development and validation approaches and key engineering framework improvements to ensure that each phase of the release process provides a solid foundation for the subsequent phase till final product release. As the name implies, the goal is to move development and validation earlier into the release cycle to ensure content design, development, validation and bug fixes are occurring when the bulk of the Engineering resources are engaged and available. Each release phase is focused on delivering the building blocks for a successful and high quality release. We have adopted industry standard best practices and methodologies to bolster our engineering framework and process to support this transformation. Agile forms the cornerstone of our new content development scrum teams with the goal of maintaining a potentially shippable product early and consistently.

TRANSFORMING PCIE-SSDS AND HDDS WITH INFINIBAND INTO SCALABLE ENTERPRISE STORAGE

MONDAY 1:30 - 2:20

Developers and Technologists are fascinated by the low latency and high IOPS of PCIe-SSDs. But, customers expect a healthy balance between performance and enterprise features such as high availability, scalability, elasticity and data management. The open source distributed storage solution Ceph, designed for reliability, performance, and scalability in combination with commodity hardware as Infiniband, PCIe-SSDs and HDDs will merge into a perfect team, if the best I/O parameters and the right interconnect protocols are used and tuned.

SOFTWARE DEFINED X TRACK

DEFINING SOFTWARE DEFINED STORAGE

Monday 2:30 - 3:20

The notion of data services being comprised of software has natural appeal, but what exactly does it mean? Given a huge portfolio of software and hardware that is available for a datacenter today, it is difficult to make sense of what "software defined storage" truly is and what benefits it could provide. While there is some truth to the idea that it is about reducing reliance on costly hardware, many see it as a way to bring new flexibility to datacenter operations. In this discussion, we will propose a set of requirements and benefits, while walking through some examples of various software technologies with the goal of producing a crisp definition.

HOSTING PERFORMANCE-SENSITIVE APPLICATIONS IN THE CLOUD WITH SOFTWARE-DEFINED STORAGE

Monday 3:35 - 4:25

Hosting performance-sensitive enterprise applications requires the delivery of guaranteed quality of service (QoS), which has been Achilles' heel of large cloud service providers. So, what stops legacy solutions from delivering guaranteed QoS? Noisy neighbors! Within a shared storage platform, legacy solutions cannot isolate and dedicate a specific set of resources to any application. As a result, applications are in a constant struggle for the shared storage resources. This session will look at different storage options with a focus on software-defined storage solutions that help solve the noisy neighbor problem and guarantee QoS to every application in a shared storage environment.

TUESDAY SESSIONS

CLOUD TRACK

WINDOWS AZURE STORAGE - SPEED AND SCALE IN THE CLOUD

Tuesday 1:00 - 1:50

In today's world that is increasingly dominated by mobile and cloud computing application developers require durable, scalable, reliable, and fast storage solutions like Windows Azure Storage. This talk will cover the internal design of the Windows Azure Storage system and how it is engineered to meet these ever growing demands. This session will have a particular focus on performance, scale, and reliability. In addition, we will cover patterns & best practices for developing performant solutions on storage that optimize for cost, latency, and throughput. Windows Azure Storage is currently leveraged by clients to build big data and web scale services such as Bing, Xbox Music, SkyDrive, Halo 4, Hadoop, and Skype.

RESILIENCE AT SCALE IN THE DISTRIBUTED STORAGE CLOUD

Tuesday 2:00 - 2:50

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. 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 shamelessly 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.

CDMI AND SCALE OUT FILE SYSTEM FOR HADOOP

Tuesday 3:05 - 3:55

Scality leverages its own file system for Hadoop and replaces HDFS while maintaining HDFS API. Scality Scale Out File System aka SOFS is a POSIX parallel file system based on a symmetric architecture. This implementation addresses the Name Node limitations both in term of availability and bottleneck with the absence of meta data server with SOFS. Scality leverages also CDMI and continues its effort to promote the standard as the key element for data access. Scality capitalizes on 2 data protection techniques - Replication and Erasure Coding with Scality ARC - to boost data access, improve data durability and reduce hardware footprint and costs.

TRANSFORMING CLOUD INFRASTRUCTURE TO SUPPORT BIG DATA

Tuesday 4:05 - 4:55

Cloud systems promise virtually unlimited, on-demand increases in storage, computing, and bandwidth. As companies have turned to cloud-based services to store, manage and access big data, it has become clear that this promise is tempered by a series of technical bottlenecks: transfer performance over the WAN, HTTP throughput within remote infrastructures, and size limitations of the cloud object stores. This session will discuss principles of cloud object stores, using examples of Amazon S3, Microsoft Azure, and OpenStack Swift, and performance benchmarks of their native HTTP I/O. It will share best practices in orchestration of complex, large-scale big data workflows. It will also examine the requirements and challenges of such IT infrastructure designs (on-premise, in the cloud or hybrid), including integration of necessary high-speed transport technologies to power ultra-high speed data movement, and adoption of appropriate high-performance network-attached storage systems.

FILES SYSTEMS TRACK

A BRIEF HISTORY OF THE BSD FAST FILESYSTEM

Tuesday 1:00 - 1:50

This talk provides a taxonomy of filesystem and storage development from 1979 to the present with the BSD Fast Filesystem as its focus. It describes the early performance work done by increasing the disk block size and by being aware of the disk geometry and using that knowledge to optimize rotational layout. With the abstraction of the geometry in the late 1980's and the ability of the hardware to cache and handle multiple requests, filesystems performance ceased trying to track geometry and instead sought to maximize performance by doing contiguous file layout. Small file performance was optimized through the use of techniques such as journaling and soft updates. By the late 1990's, filesystems had to be redesigned to handle the ever growing disk capacities. The addition of snapshots allowed for faster and more frequent backups. The increasingly harsh environment of the Internet required greater data protection provided by access-control lists and mandatory-access controls. The talk concludes with a discussion of the addition of symmetric multi-processing support needed to utilize all the CPUs found in the increasingly ubiquitous multi-core processors.

TESTING METHODOLOGIES TRACK

TESTING iSCSI / SCSI PROTOCOL COMPLIANCE USING LIBISCSI

Tuesday 2:00 - 2:50

Presenting the libiscsi iSCSI/SCSI test suite. There are many independent iSCSI/SCSI software targets on the market today but no real iSCSI/SCSI test suites. My experience from writing tests and testing many popular iSCSI/SCSI implementations has shown me there is a real need for a good test suite. The libiscsi userspace initiator library comes with a comprehensive iSCSI/SCSI test suite. This is the most complete and comprehensive open test suite available for SCSI target implementors today. This presentation will cover the structure of libiscsi, the testsuite, how to add more tests. We will look at source code and we will run a short demo against a software target. And we will talk about how this test suite will help your target be better. This presentation targets iSCSI target developers. It aiming to show a test suite that can be easily extended and applied in an automated regression test framework and the benefits it will bring to improve target quality.

MESSAGE ANALYSIS AND VISUALIZATION IN HETEROGENEOUS ENVIRONMENTS

Tuesday 3:05 - 3:55

Microsoft Message Analyzer is the next generation tool for analyzing messages from almost any source. Diagnosis of heterogeneous systems has continued to evolve as we explore new ways to visualize information for any type of trace data, be it a text log file, comma or tab separated data, network capture, or ETW component. Discover how to import Samba debug logs directly or define Text Log adapters, then inspect, filter, and organize as structured data. Learn how to analyze your file systems interoperability with Windows without having to read documentation. Expand your understanding of the interactions by including Windows component-specific information to gain insight into deep protocol and system behaviors.

A SOFTWARE BASED FAULT INJECTION FRAMEWORK FOR STORAGE SERVER

Tuesday 4:05 - 4:55

With increasing complexity of storage systems, the ability to gracefully handle errors at all layers of a storage server (array firmware, driver, file system, protocols) has become a key challenge to developers. This is crucial in scalable storage environment, where error handling has to be synchronized across multiple nodes. This makes software fault injection at various layers of the stack more important in storage development and testing. Currently there is no single infrastructure that allows selective injection of faults in a typical storage server implementation. While investigating this problem, we have studied available options and designed a framework that uses combination of Kprobe, frysk, at system and protocol layer, and custom firmware fault injection mechanism that can simulate transient and hard errors at various layers.

SMB2/SMB3 TRACK

SMB DIRECT UPDATE

Tuesday 1:00 - 1:50

This talk will explore upcoming changes to the SMB 3 protocol that increase SMB Direct performance for high IOP workloads. The protocol changes will be motivated by performance analyses, including updated SMB Direct performance results for a variety of IO workloads.

A STATUS REPORT ON SMB DIRECT (RDMA) FOR SAMBA

Tuesday 2:00 - 2:50

Since Microsoft announced SMB Direct there has been interest in providing support for SMB DIrect under Samba. This presentation will describe the current state of the project to provide that support. It will discuss the process that we have undertaken, the players, and what we have working todate.

SAMBA SCALABILITY AND PERFORMANCE TUNING

Tuesday 3:05 - 3:55

During the last few months a lot of work has been spent on making Samba scale well for large numbers of clients. In particular together with ctdb a few bottlenecks were discovered that were surprising for the developers. Luckily, most of these bottlenecks could be fixed. This talk will present the details of our improved scalability in Samba.

SCALED RDMA PERFORMANCE & STORAGE DESIGN WITH WINDOWS SERVER SMB 3.0

Tuesday 4:05 - 4:55

This session will present a summary of the performance of the Windows Server 2012 File Server's Remote Direct Memory Access capabilities over SMB 3.0. Systems presented will range from production "Windows Server Cluster in a Box" from EchoStreams to rack-scaled storage systems, across multiple RDMA solutions. Design considerations for highly scaled systems and their tradeoffs will be discussed. With RDMA the processor cost of bulk data access on remote file systems has the potential to approach the range of local storage. This provides a novel build option for deploying high speed and highly efficient consolidated storage solutions.

HARDWARE TRACK

ADDRESSING SHINGLED MAGNETIC RECORDING DRIVES WITH LINEAR TAPE FILE SYSTEM

Tuesday 1:00 - 1:50

Shingled Magnetic Recording is a disruptive technology. It increases the capacity of the drive at the expense of not supporting random writes. This limits the adoption of SMR devices in traditional systems with write in place file systems. We can address the write expectations of SMR through various layers of abstraction, from application to firmware. A high abstraction layer provides more room for innovation and a more consistent performance guarantee. Thus, one potential implementation is through the familiar POSIX/Unix file system interface which provides a stable and familiar abstraction for both the storage vendor and user. In this presentation we would like to share some of the thoughts, lessons and experiences that we went through in making Linear Tape File System work with WD SMR drives.

INFINIBAND ARCHITECTURAL OVERVIEW

Tuesday 2:00 - 2:50

This session will provide an overview of the entire InfiniBand Architecture including application, transport, network, link, and physical layers. This session is meant to update the student on current and future enhancements to the IB architecture including 8 Gbps links and RoCE.

INFINIBAND VERBS AND MEMORY MANAGEMENT - RDMA

Tuesday 3:05 - 3:55

This session overviews the verb interface and RDMA protocol including how memory regions and windows are used for inter-processor communication. NVM and SCSI Express both utilized similar programming interfaces (queue pairs) to communicate between host RAM and either another hosts' RAM or to a non-volatile storage device.

INTRODUCTION TO HP MOONSHOT

Tuesday 4:05 - 4:55

HP's Moonshot represents a series of products designed to ease and expedite the onramp of emerging low-power, low-cost, high-density, high-volume technologies in the data center. HP's first Moonshot System breaks new ground in terms of power efficiency and compute density with a flexible cartridge-based form factor. Learn about the capabilities of HP Moonshot and emerging technologies as we explore the genesis of the platform, where it could go and what it could mean for storage in the low power, highly efficient data center.

SOLID STATE TRACK

SNIA NVM PROGRAMMING MODEL

Tuesday 1:00 - 1:50

Upcoming advances in Non-Volatile Memory (NVM) technologies will blur the line between storage and memory, creating a disruptive change to the way software is written. The new SNIA NVM Programming Model describes behavior provided by operating systems enabling applications, file systems, and other software to take advantage of new NVM capabilities. This tutorial describes four programming modes. Two modes address NVM extensions for NVM emulating hard disks; block mode (as use by file systems) and file mode (as use by most applications). There are also two modes for Persistent Memory (PM); kernel extensions (as used by PM-aware file systems) and PM file mode (as used by PM aware applications). The tutorial also addresses some broader NVM software issues, such as strategies for storing pointers in persistent memory.

NVMe BASED PCIe SSD VALIDATION – CHALLENGES AND SOLUTIONS

Tuesday 2:00 - 2:50

PCI express (PCIe) solid state drives (SSDs) provide significant performance benefits in enterprise applications as compared to traditional HDDs and SSDs with a legacy interface. The existing protocols [SAS, SATA] pose architectural limitations that prohibit them to deliver much desired throughput for SSD. The ideal solution to this problem is to move these devices closer to PCIe space which would provide the optimum speed without adding the overheads posed by protocols like SAS and SATA. Emergence of non-volatile memory express (NVMe), a scalable host controller interface specifically developed for PCIe SSDs, and a supporting ecosystem, will allow SSD suppliers to transition to NVMe based PCIe SSD products. It is essential to understand the product validation challenges to reduce time to market for PCI-SSD vendors. This paper highlights the challenges in validating queue configuration, handling of outstanding IO's, queue arbitration, interrupt coalescing, etc and provides solutions to address these challenges.

TBF: A MEMORY-EFFICIENT REPLACEMENT POLICY FOR FLASH-BASED CACHES

Tuesday 3:05 - 3:55

TBF presents a RAM-frugal cache replacement policy that approximates the least-recently-used (LRU) policy. It uses two in-RAM Bloom filters to maintain the recency information and leverages an on-flash key–value store to cache objects. TBF could be easily integrated with any key-value stores to provide caching functionalities. TBF requires only one additional byte of RAM per cached object while providing similar performance as LRU and its variants, thus makes it suitable for implementing a very large flash-based cache.

DELIVERING NANOSECOND-CLASS PERSISTENT STORAGE

Tuesday 4:05 - 4:55

NAND flash solves many problems in storage with its non-volatility and high IOPS performance. Designers can always deliver more IOPS with NAND assuming unlimited power and space. However, designers can't deliver nanosecond-class response times with NAND because the medium isn't fast enough. Spin Torque MRAM complements NAND flash and forms a persistent replacement for battery or capacitor-backed DRAM, delivering higher IOPS/$ and IOPS/W than NAND flash with nanosecond-class response times. In this presentation, the speaker will discuss how ST-MRAM is enabling a latency revolution in storage, just as NAND flash delivered an IOPS revolution.

WEDNESDAY SESSIONS

LONG TERM RETENTION TRACK

COMBINING SNIA CLOUD, TAPE AND CONTAINER FORMAT TECHNOLOGIES FOR THE LONG TERM RETENTION OF BIG DATA

Wednesday 1:15 - 2:05

Generating and collecting very large data sets is becoming a necessity in many domains that also need to keep that data for long periods. Examples include astronomy, atmospheric science, genomics, medical records, photographic archives, video archives, and large-scale e-commerce. While this presents significant opportunities, a key challenge is providing economically scalable storage systems to efficiently store and preserve the data, as well as to enable search, access, and analytics on that data in the far future. Both cloud and tape technologies are viable alternatives for storage of big data and SNIA supports their standardization. The SNIA Cloud Data Management Interface (CDMI) provides a standardized interface to create, retrieve, update, and delete objects in a cloud. The SNIA Linear Tape File System (LTFS) takes advantage of a new generation of tape hardware to provide efficient access to tape using standard, familiar system tools and interfaces. In addition, the SNIA Self-contained Information Retention Format (SIRF) defines a storage container for long term retention that will enable future applications to interpret stored data regardless of the application that originally produced it. This tutorial will present advantages and challenges in long term retention of big data, as well as initial work on how to combine SIRF with LTFS and SIRF with CDMI to address some of those challenges. SIRF for the cloud will also be examined in the European Union integrated research project ForgetIT – Concise Preservation by combining Managed Forgetting and Contextualized Remembering.

BEST PRACTICES, OPTIMIZED INTERFACES, API'S DESIGNED FOR STORING MASSIVE QUANTITIES OF LONG TERM RETENTION DATA

Wednesday 2:15 - 3:05

The growth, access requirements and retention needs for data in a mass storage infrastructure for HPC, life sciences, media and entertainment, higher education and research are becoming unmanageable. Organizations continue to utilize legacy methodologies to manage Big Data Growth of today and it is not working. Traditional storage tiering and backups do not solve the problem and create additional cost and overhead. Redefining the term "Archive" as an online,

accessible, affordable data management platform decoupled from infrastructure will be required to solve data growth and retention challenges going forward. Leveraging new optimized interfaces and API's for disk, tape, and cloud will be required to fully enable the Active Archive experience.

CAN STORAGE FIX HADOOP?

Wednesday 3:20 - 4:15

Survey data shows that at least half of all enterprise data center Hadoop projects are stalled and that only 20% are actually making into production. This presentation looks at the problems with Hadoop that enterprise data center administrators encounter and how the storage environment can be used to fix at least some of these problems.

HADOOP: EMBRACING FUTURE HARDWARE

Wednesday 4:20- 5:15

This talk looks at the implications to Hadoop of future server hardware - and to start preparing for them. What would a pure SSD Hadoop filesystem look like, and how to get there via a mixed SSD/HDD storage hierarchy? What impact would that have on ingress, analysis and HBase? What could we do do better if network bandwidth and latency became less of a bottleneck, and how should interprocess communication change? Would it make the graph layer more viable? What would massive arrays of WIMPy cores mean -or a GPU in every sever. Will we need to schedule work differently? Will it make per-core RAM a bigger issue? Finally: will this let us scale Hadoop down?

NFS TRACK

NFS ON STEROIDS: BUILDING WORLDWIDE DISTRIBUTED FILE SYSTEM

Wednesday 1:15 - 2:05

Intel R&D environment spans dozens of locations across the Globe. It includes over 50,000 compute servers running Linux, and 10s of PBs of centralized NFS storage. NFS provides good solution for data sharing within data center. However, it doesn't necessarily give an answer for cross-site access over high latency links. The presentation will share Intel IT experience with development, implementation and adoption of the NFS-based federated distributed secure storage infrastructure, where every file is accessible from any client worldwide. It will describe multiple steps required to enable global multi-level on-demand caching infrastructure, including new caching and manageability solutions, environment standardization and more. Improvements made in data transfer and storage synchronization protocols will also be covered.

IMPLEMENTING NFSv3 IN USERSPACE: DESIGN AND CHALLENGES

Wednesday 2:15 - 3:05

NFS in usermode is an uncommon challenge, one which necessitates unique design features not present in most UNIX implementations. EMC Isilon's NFS team has been tasked with moving the kernel-based NFS server previously implemented in OneFS with one in userspace in order to take advantage of a new protocol-agnostic access auditing framework. This talk will serve as a post-mortem discussion and case study of EMC Isilon's new usermode NFSv3 server. We will discuss some of the challenges we encountered while redesigning the server for its new environment, future-proofing our design, and maintaining common code with SMB, NFSv4, and other protocols.

pNFS, NFSv4.1, FedFS AND FUTURE NFS DEVELOPMENTS

Wednesday 3:20- 4:10

The NFSv4 protocol undergoes a repeated life cycle of definition and implementation. The presentation will be based on years of experience implementing server-side NFS solutions up to NFSv4.1, with specific examples from NetApp and others. We'll examine the life cycle from a commercial implementation perspective; what goes into the selection of new features (including FedFS and NFSv4.2 and NFSv4.3), the development process and how these features are delivered, and the impact these features have on end users. We'll also cover the work of Linux NFS developers and provide suggestions for file system developers based on these and vendor experiences; and finally, we'll discuss how implementation and end-user experience feeds back into the protocol definition, along with an overview of expected NFSv4.2 features.

pNFS DIRECTIONS

Wednesday 4:20 - 5:10

This session features leading pNFS developers describing ways in which pNFS is adapting to meet new challenges in distributed storage. New features are under discussion by the IETF NFS Working Group and new implementation platforms are emerging as pNFS adapts to market adoption. Some of these include: Metadata scaling proposals; New front-and back-ends, including Ceph and Ganesha; Defining a software-defined storage layer; Integration with cloud storage APIs; Defining workloads for pNFS performance and stability measurement

HOT TOPICS TRACK

CAN YOUR STORAGE INFRASTRUCTURE HANDLE THE COMING DATA STORM?

Wednesday 1:15 - 2:05

In day-to-day operations, a storage infrastructure must effectively perform unique tasks, like data storage, backups, access validations, edits, deletes, analysis, etc. Any delay introduced at the storage level impacts user quality of experience (QoE). To ensure effective storage infrastructure, you must evaluate and optimize the system's ability to perform under extreme environments. Strong and resilient storage not only must handle today's data storm – business-critical financial transactions, the fire hose of big data, on-demand video and gaming, etc. – but also stores and protects the most precious artifacts of modern-world data.

OPENSTACK CLOUD STORAGE

Wednesday 2:15 - 3:05

OpenStack is an open source cloud operating system that controls pools of compute, storage, and networking. It is currently being developed by thousands of developers from hundreds of companies across the globe, and is the basis of multiple public and private cloud offerings. This presentation will outline the storage aspects of OpenStack including the core projects for block storage (Cinder) and object storage (Swift), as well as the emerging shared file service. It will cover some common configurations and use cases for these technologies, and how they interact with the other parts of OpenStack. The talk will also cover new developments in Cinder that enable a variety of storage devices and storage fabrics to be used

19 BREAKOUT SESSIONS

DATA MANAGEMENT TRACK

DATA DEDUPLICATION AS A PLATFORM FOR VIRTUALIZATION AND HIGH SCALE STORAGE

Wednesday 3:20 - 4:10

The primary data deduplication system in Windows Server 2012 is designed to achieve high deduplication savings at low computational overhead on commodity storage platforms. In this talk, we will build upon that foundational work and present new techniques to scale primary data deduplication on both the primary data serving and optimization pathways. This will include hardware accelerated performance improvements for hashing and compression, better file system integration to reduce write path overheads and optimize live files, and deduplication aware caching to mitigate disk bottlenecks. We will show how this enables deduplication to be leveraged as a platform for storage virtualization.

VIRTUAL MACHINE ARCHIVAL SYSTEM

Wednesday 4:20- 5:10

Popular server virtualization vendors have enabled integration with backup and recovery solutions, but not with virtual machine archival systems. Server virtualization system should have knowledge of various storage systems attached to it such as SSD, HDD, Object Storage, Tape library and Cloud. For instance, VMWare ecosystem. We propose:

Virtual Machine Archival System with following functionality:

- Decision on which type of storage should be used as destination
- Labeling locations of VMs data
- Discovery interface and VM archival policies

Server virtualization system facilitates following functionality:

- Storage types available around
- Archival link creation in file system containing VM data
- Passing on archival and restore request to archival system
- GUI integration for archival system

NEW THINKING TRACK

SCREAMING FAST GALOIS FIELD ARITHMETIC USING INTEL SIMD INSTRUCTIONS

Wednesday 1:15 - 2:05

Galois Field arithmetic forms the basis of Reed-Solomon and other erasure coding techniques to protect storage systems from failures. Most implementations of Galois Field arithmetic rely on multiplication tables or discrete logarithms to perform this operation. However, the advent of 128-bit instructions, such as Intel's Streaming SIMD Extensions, allows us to perform Galois Field arithmetic much faster. This talk outlines how to leverage these instructions for various field sizes, and demonstrates the significant performance improvements on commodity microprocessors. The techniques that we describe are available as open source software.

NV-HEAPS: MAKING PERSISTENT OBJECTS FAST AND SAFE WITH NEXT-GENERATION NON-VOLATILE MEMORIES

Wednesday 2:15 - 3:05

Persistent, user-defined objects present an attractive abstraction for working with non-volatile program state. However, the slow speed of persistent storage (i.e., disk) has limited their performance. Fast, byte-addressable, non-volatile technologies, such as phase change memory, will remove this constraint and allow programmers to build high-performance, persistent structures in non-volatile storage that is almost as fast as DRAM. However, existing persistent object systems are ill-suited to these memories because the assumption that storage is slow drives many aspects of their design. Creating structures that are flexible and robust in the face of application and system failure, while minimizing software overheads, is challenging. The system must be lightweight enough to expose the performance of the underlying memories, but it also must avoid familiar bugs such as dangling pointers, multiple free()s, and locking errors in addition to unique types of hard-to-find pointer safety bugs that only arise with persistent objects. These bugs are especially dangerous since any corruption they cause will be permanent.

We have implemented a lightweight, high-performance persistent object system called NV-heaps that prevents these errors and provides a model for persistence that is easy to use and reason about. We implement search trees, hash tables, sparse graphs, and arrays using NV-heaps, BerkeleyDB, and Stasis. Our results show that NV-heap performance scales with thread count and that data structures implemented using NV-heaps out-perform BerkeleyDB and Stasis implementations by 32x and 244x, respectively, when running on the same memory technology. We also quantify the cost of enforcing the safety guarantees that NV-heaps provides and measure the costs for NV-heap primitive operations.

LAZYBASE: TRADING FRESHNESS FOR PERFORMANCE IN A SCALABLE DATABASE

Wednesday 3:20- 4:10

The LazyBase scalable database system is specialized for the growing class of data analysis applications that extract knowledge from large, rapidly changing data sets. It provides the scalability of popular NoSQL systems without the query-time complexity associated with their eventual consistency models, offering a clear consistency model and explicit per-query control over the trade-off between latency and result freshness. With an architecture designed around batching and pipelining of updates, LazyBase simultaneously ingests atomic batches of updates at a very high throughput and offers quick read queries to a stale-but-consistent version of the data. Although slightly stale results are sufficient for many analysis queries, fully up-to-date results can be obtained when necessary by also scanning updates still in the pipeline. Compared to the Cassandra NoSQL system, LazyBase provides 4X--5X faster update throughput and 4X faster read query throughput for range queries while remaining competitive for point queries. We demonstrate LazyBase's tradeoff between query latency and result freshness as well as the benefits of its consistency model. We also demonstrate specific cases where Cassandra's consistency model is weaker than LazyBase's.

GraphChi: LARGE-SCALE GRAPH COMPUTATION ON JUST A PC

Wednesday 4:20- 5:10

In "GraphChi: Large-Scale Graph Computation on Just a PC" at OSDI '12, we proposed Parallel Sliding Windows (PSW), a novel method for efficiently processing large graphs from external memory (disk). Based on PSW, we designed and implemented a complete system, GraphChi, for vertex-centric graph computation. We demonstrated that GraphChi is capable of solving even the biggest graph computation problems on just a single PC, with performance often matching distributed computation frameworks. This talk will discuss the motivations for single-computer computation, present the GraphChi system and its design and talk about some recent work for extending and improving GraphChi, including a novel random walk engine DrunkardMob (to be presented in ACM RecSys'13). The speaker will also talk about challenges of graph computation on general level and discuss future directions of his research.

PERFORMANCE TRACK

HIGH-THROUGHPUT CLOUD STORAGE OVER FAULTY NETWORKS

Wednesday 1:15 - 2:05

Storage systems increasingly rely on the Internet as the medium of data transfer. The Internet, as a high-bandwidth, high-latency, high-packet loss connection is very different from the clean networks of typical SANs. Under such conditions, TCP's capabilities are often stretched to their breaking point. In this presentation we detail the methods we used to overcome random network slowdowns, packet drops, congestion control, and other challenges. Our result: we achieved storage throughputs on the Internet that were 80% that of the same test on a low-latency, zero packet loss LAN.

FORGET IOPS: A PROPER WAY TO CHARACTERIZE & TEST STORAGE PERFORMANCE

Wednesday 2:15 - 3:05

Storage workloads are changing. Applications stress storage infrastructure in different ways – like workloads generated by virtualized applications or ones containing high amounts of meta-data. Such dynamics make it difficult to confidently predict how storage systems will behave in the real world for both end users and vendors. How can storage performance be better understood and defined?

IMPROVEMENTS IN STORAGE ENERGY EFFICIENCY VIA STORAGE SUBYSTEM CACHE AND TIERING

Wednesday 3:20- 4:10

The energy efficiency of storage subsystems in terms of Idle Capacity/Watt, IOPs/Watt, and MB/s/Watt can be significantly improved through the deployment of Capacity Optimization Methods (COMs). These features affect the apparent capacity, IO rate and throughput (MB/s) and therefore also the target "green" metrics cited above. This paper describes a case study of the compound effect of two features, that of storage subsystem cache and tiered storage on the primary metrics of the SNIA Emerald Power Efficiency Specification using both the former random workloads and the recently adopted "Hot Band" Workload as the comparative test stimuli. Also described is the potential energy efficiency benefit of several additional COM types.

LESSONS LEARNED TUNING HP'S SMB SERVER WITH THE FSCT BENCHMARK

Wednesday 4:20- 5:10

As part of an ongoing effort to increase performance of our NAS system, we have been tuning the system's performance on the FSCT benchmark. As part of this effort, we have learned a number of things about the benchmark which might be of general interest. Areas which will be discussed include an overview of the benchmark, setup and configuration steps required to run the benchmark and insight into understanding the results and errors generated by running FSCT. The intention is to discuss "what we learned about the benchmark while using it to make our system faster", including what SMB protocol elements it uses, what impact newer SMB features like leases have, etc.

THURSDAY SESSIONS

SECURITY TRACK

TRUSTED COMPUTING TECHNOLOGIES FOR STORAGE

Thursday 9:30 - 10:20

The Trusted Computing Group (TCG) has created specifications for trusted computing, with a focus on ease of use, transparency, robust security functions in hardware, integration into the computing infrastructure, and inexpensive; including Self-Encrypting Drives (SED). TCG technologies will be described, including application to the design of trusted storage.

MULTI-VENDOR KEY MANAGEMENT – DOES IT ACTUALLY WORK?

Thursday 10:30 - 11:20

A standard for interoperable key management exists but what actually happens when you try to use products and key management solutions from multiple vendors? Does it work? Are any benefits gained? Practical experience from implementing the OASIS Key Management Interoperability Protocol (KMIP) and from deploying and interoperability testing multiple vendor implementations of KMIP form the bulk of the material covered. Guidance will be provided that covers the key issues to require that your vendors address and how to distinguish between simple vendor tick-box approaches to standard conformance and actual interoperable solutions.

MATCHING SECURITY TO DATA THREATS – MORE IS NOT BETTER, BUT LESS CAN BE BAD

Thursday 11:25 - 12:15

A chain is only as strong as its weakest link and adding more links doesn't make it any stronger. The same is true for securing critical data with encryption – just adding more encryption doesn't necessarily make critical data more secure. The challenges facing most organizations are twofold: 1) understanding which threats and vulnerabilities apply to them and their data, and 2) knowing when they have sufficient data encryption to protect them from the threats, but not so much that their costs and management resources are strained. It is additionally important to understand that not all threats can be addressed by data encryption and that some threats may have to be rationalized by an organization in terms of the cost of the remedial work.

END USER VIEWPOINT TRACK

DEMAND FOR STORAGE SYSTEMS FROM A CUSTOMER VIEWPOINT IN JAPAN

Thursday 9:30 - 10:20

We are providing storage services more than 20 years with large-scale NAS systems in JAIST for central data managing whole over our institute. Recently, construction of the private cloud environment based on virtualization technology complicate the dependency structure of systems, and also makes it increasingly difficult to operate our storage systems. In this presentation, we talk about case study and demands for storage systems from a customer viewpoint, based on our knowledge from deploying and operating our storage systems. Furthermore, the non-technical matter is also important in operating storage systems, i.e. we need to get good support for trouble shooting. We mention the consideration from this point with story in Japan. Note: This presentation will be jointly-conducted with SNIA-J (SNIA Japan Branch).

DEVELOPMENT METHODOLOGIES TRACK

A METHOD TO ESTABLISH CONCURRENT RAPID DEVELOPMENT CYCLE AND HIGH QUALITY IN A STORAGE ARRAY SYSTEM ENVIRONMENT

Thursday 10:30 - 11:20

Shift Left is a combination of a change in development and validation approaches and key engineering framework improvements to ensure that each phase of the release process provides a solid foundation for the subsequent phase till final product release. As the name implies, the goal is to move development and validation earlier into the release cycle to ensure content design, development, validation and bug fixes are occurring when the bulk of the Engineering resources are engaged and available. Each release phase is focused on delivering the building blocks for a successful and high quality release. We have adopted industry standard best practices and methodologies to bolster our engineering framework and process to support this transformation. Agile forms the cornerstone of our new content development scrum teams with the goal of maintaining a potentially shippable product early and consistently.

CODE COVERAGE AS A PROCESS

Thursday 11:25 - 12:15

There are many tools to get code coverage for different languages, but the data is of no use if not used to improve the quality of the product through testing. In this presentation I will be sharing our EMC/DataDomain successful process/infrastructure. After this presentation the audience would be able to start thinking about code coverage from both Development and QA perspective if they don't already have one.

STORAGE UTILITIES

BUILDING THE PUBLIC STORAGE UTILITY

Thursday 9:30 - 10:20

The holy grail of cloud storage is to make storage into a utility. That is, an ubiquitous, standard, public resource, in the same sense that electricity and tap water are today. What makes utility storage difficult is that unlike water or electricity each user's data is unique and private. In this presentation we propose a solution to this problem. Our proposed solution enables a global anonymous public storage service where the storage system has no knowledge about users, or the data or metadata they store. Yet each user has their own private and secure storage space. Further, we consider some of the payment options that exist within a fully anonymous storage utility.

STORAGE MANAGEMENT TRACK

EVENT TREES FOR STORAGE MONITORING AND MANAGEMENT

Thursday 10:30 - 11:20

Event Tree, Fault Tree and Decision Tree Analysis are powerful techniques for storage resource and service management. This paper and presentation discusses a framework for designing, developing and management a known event tree database for the storage domain and application of different sets of analytic on known event trees and real time event trees (as they occur) for better event, incident and problem management as well as better service management.

A METHOD TO BACKUP AND RESTORE CONFIGURATION SETTINGS FOR EACH AND EVERY COMPONENT IN THE SAN ENVIRONMENT USING SMI-S

Thursday 11:25 - 12:15

As of today all disaster recovery solutions getting shipped communicates with APIs from different vendor devices to retrieve information and perform management oriented operations on it. To perform these operations it is very important the solution is interoperable with the APIs provided by the vendors. The method proposed in this paper overcomes the issues arising from dependency on the vendor APIs. The solution is a software stack that backs up the configuration on devices participating in a SAN and upon failure of a site or a device in a site the solution bears the capability of replicating the same configuration that existed on the previous site/device on the new site/device. The solution is based on SNIA standards using the SMI providers developed by each of the device vendors. Since the SMI providers are built on a standard CIM model to which all devices seek compliance to; the interoperability factor between the individual devices is overcome and developing a solution comprising of different vendor products becomes lot easier, enhanced request/response time and cost effective.

WORKLOADS TRACK

TUNNELING SCSI OVER SMB: SHARED VHDX FILES FOR GUEST CLUSTERING IN WINDOWS SERVER 2012 R2

Thursday 9:30 - 10:20

Windows Server Failover Clustering is a well-established technology for increasing application availability in the Microsoft platform. For Hyper-V virtualized workloads, you can also create a Failover Cluster comprised of a set of virtual machines and some shared storage in the form of iSCSI LUNs, virtual Fibre Channel LUNs or SMB file shares. In this session, we'll start by describing the overall Guest Clustering scenario and then dive into the new Windows Server 2012 R2 option to use Shared VHDX files. We'll then introduce the "Remote Shared Virtual Hard Disk Protocol" (the new protocol behind the Shared VHDX feature) and explain how this protocol allows SCSI commands to be tunneled over the SMB protocol. We'll also cover how non-Windows SMB implementations can offer the Shared VHDX capability via this new protocol.

AN SMB3 ENGINEER'S VIEW OF WINDOWS SERVER 2012 HYPER-V WORKLOADS

Thursday 10:30 - 11:20

Many traditional physical storage workloads are well understood. Quantifying how access patterns change when a hypervisor in inserted in between server applications and physical storage requires rethinking what is optimal for a NAS configuration. This presentation will examine several Hyper-V workloads from the perspective of an SMB3 implementer.

DIRECT NFS – DESIGN CONSIDERATIONS FOR NEXT-GEN NAS APPLIANCES OPTIMIZED FOR DATABASE WORKLOADS

Thursday 11:25 - 12:15

NAS appliances have traditionally been a popular choice for shared storage as they support a standardized and mature NFS protocol and leverage inexpensive Ethernet networking. However, the NFS protocol and traditional NAS appliances are designed for general purpose file system storage. Database workloads are very unique in the kind of requirements they place on a storage system. Different database workloads can have very different response time or bandwidth requirements. Along with the traditional database requirements of atomicity

and consistency; critical database systems also have strict uptime and high availability requirements. Database workloads have the ability to convey this information to the storage. This session will explore some novel ideas to help design the next generation of NAS appliances and integrate newer NFS protocols and breakthrough technologies such as flash storage for optimizing database performance.

BACKUP SESSIONS

EMPOWERING ENGINEERING COMPUTE WORKLOADS WITH pNFS

Bikash Roy Choudhury, Solutions Architect, NetApp

The amount of job processing done in compute farms that uses Engineering application\tools like Electronic Device Automation(EDA)\Semi-Conductor manufacturers has been scaling over time as the complexity of the the end products has been constantly on the rise. Time to market is very critical for these workloads. While the compute farms and storage scale; will pNFS as a protocol help to improve the job completion times as a pararrel file-system.

CLOUD GATEWAY – CACHE CONSIDERATIONS

Devendra Vaja, Technical Architect, iGATE

Nilesh Wilanker, Technical Architect, iGATE

With the evolution of Cloud storages supporting various cloud storage APIs based on REST or SOAP, apparent need of Cloud storage gateway has emerged. Cloud Gateway is a protocol translator which converts the existing block based protocols such as iSCSI, or Fiber Channel or File based interfaces such as NFS or CIFS to cloud storage APIs. Cloud Storage Gateway adoption is being discussed but there is not much information about the technical challenges involved. This paper presents one of the Cloud Gateway implementation challenges, the Caching, a performance enhancement technique, and critical criteria to the selection of Cloud Gateway service.

A METHOD FOR IMPROVING VMWARE INFRASTRUCTURE (VI) PERFORMANCE ON A FLASH STORAGE ARRAY WITH VAAI

Dr. M. K. Jibbe, Director of Quality Architecture Team for all NetApp APG Products, NetApp

Bernard Chan, Quality Architect, NetApp

VAAI is a standard provided by VMware to enhance performance to VM solutions by allowing the ESX to offload operations involving data copy and data initialization to the storage hardware. This paper explains how the NetApp EF540 flash storage array can further enhance the performance of the VM solutions when VAAI is enabled. When data copy and data initialization operations are offloaded to the storage hardware, the speed of which the storage controllers can read from and write to the disk drives has a significant impact on how fast the operation will take. This paper provides an overview of the algorithm the EF540 storage controller takes to accomplish the VAAI (i.e. Write Same and Extended Copy) requests and will show the Flash Storage performance improvement.

PREQUEL: DISTRIBUTED CACHING AND WAN ACCELERATION

Jose Rivera, Software Engineer, Red Hat, Inc.

Christopher Hertel, Senior Princical Software Engineer, Red Hat, Inc.

Prequel is an open-source implementation of PeerDist, a de-facto standard protocol for wide-area distributed caching developed by Microsoft. PeerDist is more commonly known as Microsoft's BranchCache feature. Prequel seeks to bring PeerDist into the open source world and integrate its capabilities with data access projects like Samba. This talk will be held in two parts: history and demo. In the history portion you will learn about Prequel's development, where its been in the last year, and the various components of the project. In the demo portion, we'll demonstrate a working Prequel installation and walk through the various steps of the protocol. Questions will be taken all through the presentation.

SMB2 ACCELERATION STORIES

Mark Rabinovich, R&D Manager, Visuality Systems Ltd.

The CIFS/SMB protocol suffers from signal latency. SMB2 and its successor – SMB3, being declared to solve this issue still does not provide sufficient requirements. The SMB2 traffic may be improved, though, with proxy techniques. During this session we will discuss the challenges introduced by SMB2 in comparison with SMB. Then we will continue on the methods of improving the traffic with charts and numbers illustrating each of them. As any proxy, SMB2 acceleration challenges cache and context coherency. This will be discussed in the last part of the presentation.

OPEN SOURCE STORAGE OPPORTUNITIES AND LIMITATIONS

Thomas Coughlin, President, Coughlin Associates

This talk will look at current and forthcoming developments in open source storage systems using commodity hardware. We will examine some case studies for open storage systems such as those in the Open Computer Project, and related projects, to create new methods of leveraging commodity hardware to provide sophisticated enterprise functionality at lower overall costs and with less reliance on a single hardware vendor. We shall discuss how this approach has served as the backbone of many cloud storage systems but also examine where open source storage has made less progress. The talk will also examine future anticipated developments in open source storage and how these approaches will impact the future of storage device as well as system architectures.

SDC SPONSORS

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: me.emc.com

NetApp creates innovative storage and data management solutions that deliver outstanding cost efficiency and accelerate business breakthroughs. Our dedication to the principles of simplicity, innovation, and customer success has made us one of the fastest-growing storage and data management providers today. Customers around the world choose us for our "go beyond" approach and broad portfolio of solutions for business applications, storage for virtual servers, disk-to-disk backup, and more. Our solutions and virtualized storage technologies can provide nonstop availability of critical business data and speed product development, so you can deploy new capabilities with confidence and get to revenue faster than ever before. You can rely on our industry-leading solutions to lower the cost of protecting your data, business, and reputation.

We bring together the industry's leading solution partners and technology to deliver services that help you maximize what you get from your infrastructure. Our collaborative approach, working as one team with one goal, means that you get a solution that is just right for you—on time and on budget. To support your global business, we provide highly responsive support in local languages. From London to Austin, Bangalore to Tokyo, we'll work with you to solve your technical issues and help you reach your goals.

GOLD

Huawei is a leading global solutions provider. Through 25 years of dedication to customer solutions and partner relationships, Huawei has established an end-to-end strategy across the carrier network, enterprise and cloud computing space, and consumer products. We are committed to creating maximum value for telecom carriers, enterprises and consumers by providing complete and competitive CT and ICT solutions and services. Our products have been deployed in over 140 countries, serving more than one third of the world's population.

Huawei entered the enterprise storage market in 2002. We provide a wide range of storage products, including the HVS-series (SAN storage), T-series (Unified storage), N-series (NAS storage), Dorado (SSD all flash storage) and UDS (Cloud storage) with integrated backup, data deduplication and protection. Huawei embraces the concept of On-demand Services with Converged Infrastructure. This concept helps customers address the challenges posed by exploding data growth and continuously changing demands on the storage by applications in areas of virtualization, data analytics and cloud computing. Huawei's line of storage products provides its customers with a complete storage solution that dramatically reduces TCO.

HP creates new possibilities for technology to have a meaningful impact on people, businesses, governments and society. With the broadest technology portfolio spanning printing, personal systems, software, services and IT infrastructure, HP delivers solutions for customers' most complex challenges in every region of the world. More information about HP (NYSE: HPQ) is available at http://www.hp.com.

GOLD

Ixia provides comprehensive converged IP network validation and visibility solutions. Manufacturers, service providers, enterprises and government agencies use Ixia to design, verify and monitor communications equipment and networks. Ixia's performance test solutions emulate realistic media-rich traffic and network conditions to optimize and validate the design and performance of storage networks/devices.

SerNet is the leading service provider for the open source software SAMBA. Since 1996 members of the international SAMBA core team work with SerNet and offer support for vendors, system integrators and consulting companies. The SAMBA core team is a group of individuals that work for a variety of companies all over the world. SerNet is the only company that offers free developer and support resources from those core team members world wide and in service level agreements 24/7.

In 2013 more than 150 customers world wide rely on SAMBA knowledge from SerNet. With team members in almost every major time zone SAMBA coding never sleeps. SerNet takes care that vendors and integrators can cope with this fast development and get the most out of SAMBA.

SerNet hosts "samba eXPerience" - the SAMBA conference every year in May since 2002, where the team meets with users and customers to share knowledge about current developments and future plans: Refer to http://sambaXP.org for more information.

SerNet offers "Enterprise SAMBA", a special release of SAMBA in ready to run binaries for a variety of systems including all Enterprise Linux platforms: http://enterpriseSAMBA.com

For more information refer to: https://www.sernet.de/en/samba

As the defacto standard in storage infrastructure testing and performance validation, SwiftTest empowers storage system design and QA professionals with the insight needed to optimize the performance, availability, and cost of their storage products. SwiftTest enables comprehensive modeling of application workloads and their associated I/O profiles. By accurately emulating real-world application behavior and being able to create massive amounts of traffic with our purpose-built load generation appliances, SwiftTest enables storage developers to comprehensively test networked storage systems to ensure performance scalability and reliability under load. With support for all leading file, block and object storage protocols, SwiftTest appliances offer the unique ability to cost-effectively stress today's most complex storage systems, protocols, and storage technologies to their limits.

SILVER

Founded in 2001, Berkeley Communications is a leading systems integrator providing technology products and services to a wide range of customers worldwide. Founded by datacenter and storage experts, Berkeley Communications focus is on designing and implementing highly available, scalable, and secure datacenter and network infrastructures to its customers.

Chelsio is a leading technology company focused on solving high performance networking and storage challenges for virtualized enterprise data centers, cloud service installations, and cluster computing environments. Chelsio is delivering hardware and software solutions including Unified Wire Ethernet network adapter cards, unified storage software, high performance storage gateways and unified management software.

SDC SPONSORS

SILVER

IBM System Storage helps you maximize return on your investments, reduce complexity and drive innovation with storage hardware, software and services. For more information, visit us on the web at: ibm.com/storage or www.storagecommunity.org.

Intel is a world leader in computing innovation. The company designs and builds the essential technologies that serve as the foundation for the world's computing devices. Additional information about Intel is available at www.intel.com.

Mellanox Technologies is a leading supplier of end-to-end InfiniBand and Ethernet interconnect solutions and services for servers and storage. Mellanox interconnect solutions increase data center efficiency by providing the highest throughput and lowest latency, delivering data faster to applications and unlocking system performance capability. Additional information available at www.mellanox.com.

Scality is an industry leader in petabyte-scale storage. Scality's RING software offers an award winning, scale-out storage solution that operates seamlessly on any commodity server hardware. It provides outstanding scalability and data persistence, while the end-to-end parallel architecture provides unsurpassed performance. For more information, visit: www.scality.com or @Scality on Twitter.

Western Digital Corp. (NASDAQ: WDC), Irvine, Calif., is a global provider of products and services that empower people to create, manage, experience and preserve digital content. Its companies design and manufacture storage devices, networking equipment and home entertainment products under the WD, HGST and G-Technology brands. Visit the Investor section of the company's website (www.westerndigital.com) to access a variety of financial and investor information.

SNIA SMB2/SMB3 AND iSCSI PLUGFESTS 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 INTEROPERABILITY 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).

CONTRIBUTING

SurDoc™ Corporation is an innovation leader in document processing and backup technology. Seeking to redefine what consumers should expect from a free backup service, SurDoc offers users on all platforms this new open document sharing capability – a feat made possible through UOML (Unstructured Operation Markup Language) technology. UOML is an OASIS global standard that defines instructions on how to process documents, i.e., any information that can be printed on paper.