

TUTORIAL DESCRIPTIONS

TUTORIALS: STANDARDS TO SOLUTIONS

FCOE TECHNOLOGY UPDATE

MONDAY 9:00 - 9:50

Many new FCoE architectural features have been defined in the FC-BB-5 standard in recent months. Dave Peterson, the editor of the FC-BB-5 standard will provide a technical overview of the FCoE technology with focus on the new architectural features of FCoE. Such topics include port models, FIP discovery, virtual link instantiation, virtual link maintenance, and VLAN discovery.

FCOE - FROM HYPE TO REALITY

MONDAY 10:05 - 10:55

With FCoE very close to an approved T11 standard, production deployments have begun. This presentation will look at some of those deployments and show how FCoE integrates with your existing FC environment and how FCoE will move your organization into the future Fibre Channel world of Ethernet.

NFSV4.1 AND BEYOND

MONDAY 11:05 - 11:55

Abstract coming soon.

DIGITAL STORAGE FOR PROFESSIONAL MEDIA AND ENTERTAINMENT

MONDAY 1:00 - 1:50

Overview of digital storage demand and performance for all elements of the digital media and entertainment (M&E) industry. Many types of storage devices are used including capacity hard disk drives, performance hard disk drives, optical discs and magnetic tape. In addition flash memory has emerged as a strong storage contender, particularly in rugged content capture devices as well as in high performance content delivery systems.

SHIFTING GEARS WITH SSDS

MONDAY 2:00 - 2:50

Abstract coming soon.

XAM: THE NEW STANDARD FOR DATA MANAGEMENT

MONDAY 3:15 - 5:05

The XAM Interface provides for the encapsulation of application data, user meta-data, and data system metadata as part of the 'content record,' thus permitting interoperability between fixed content storage systems and applications and allowing the automation of ILM-based practices. This session will bring together XAM visionary technologists and application developers to understand the exciting new opportunities a standard access method to fixed content storage devices provides.

XAM SDK HOLD: GET YOUR HANDS ON XAM

MONDAY 5:15 - 6:30

XAM developers will help you install the XAM SDK, show you code examples including the winners of the TopCoder competition and answer questions about XAM and the software that is available for free.

TUTORIALS: EMERGING TECHNOLOGIES

ZFS BASIC

MONDAY 9:00 - 9:50

Abstract coming soon.

ZFS ADVANCED

MONDAY 10:05 - 10:55

Abstract coming soon.

ZFS WHAT'S NEW

MONDAY 11:05 - 11:55

Abstract coming soon.

LINUX ADVANCED FILESYSTEMS

MONDAY 1:00 - 1:50

The latest generation of the ext2/ext3 file systems is the ext4 file system, which recently left the development status of 2.6.28. With extents, delayed allocation, multiblock allocation, persistent preallocation, and its other new features, it is substantially faster and more efficient compared to the ext3 file system. This presentation will discuss the history of the ext4 file system, its features and advantages, and how best to use it.

INTRODUCING THE NEW CLOUD DATA MANAGEMENT INTERFACE

MONDAY 2:00 - 2:50

The use of the term cloud in describing new models for storage and computing arose from architecture drawings that typically used a cloud as the dominant networking icon. The cloud conceptually represented any to any connectivity in a network, but also an abstraction of concerns such as the actual connectivity and the services running in the network that accomplish that connectivity with little manual intervention.

GREEN STORAGE PRODUCTS: EFFICIENCY WITH ENERGY STAR AND BEYOND

MONDAY 3:15 - 4:05

This talk will outline the storage-specific topics related to energy-efficiency, including metrics for measuring, managing and designing for power. We will overview ongoing efforts in the SNIA Green Storage TWG and in partnership with the EPA ENERGY STAR program, The Green Grid, the DMTF and other industry groups.

GREEN 2

MONDAY 4:15 - 5:05

Abstract coming soon.

GREEN WORKSHOP/HOLD

MONDAY 5:15 - 6:30

Abstract coming soon.

SESSION DESCRIPTIONS

TECHNICAL SESSIONS: CIFS/SMB/SBM2

REDOCUMENTING HISTORIC CIFS

MONDAY 9:00 - 9:50

The history of the SMB/CIFS protocol impacts current implementations and even casts its shadow on the new SMB2 protocol. This presentation covers an effort sponsored by Microsoft to update and correct all earlier SMB documentation, and produce a new CIFS specification covering the Windows NT implementation.

HOW TO SQUEEZE >700MB/SEC OUT OF SMB1, WHY SMB2?

MONDAY 10:05 - 10:55

With Samba 3.2, the smbclient command properly makes use of the multiplex ID field provided in the SMB1 protocol to speed up SMB file transfer. We have been able to demonstrate more than 700 Megabytes per second using 10GigE, so the question should be asked why SMB2 is required. This talk is aimed at client implementors and will point out critical aspects we found out during the implementation of the async SMB client library.

CIFS ACCELERATION TECHNIQUES

MONDAY 11:05 - 11:55

Global networks introduces a significant challenge when it comes to CIFS traffic, making it virtually unbearable for an end user. We will show how to improve CIFS traffic using various acceleration techniques. This presentation covers CIFS scenario acceleration versus file caching, proxy and client solutions, SMB2 acceleration. A special emphasis on VPN networks. We will share the Visuality's CIFS acceleration experience and performance statistics for accelerated WANs.

BRANCHCACHE: NEXT GENERATION BRANCH OFFICE OPTIMIZATIONS

MONDAY 1:00 - 1:50

With the IT management trend towards data centralization in a data center, more users are connecting to their data through a thin, expensive WAN link from a branch office. To help decrease WAN usage, Windows 7 and Windows Server 2008 R2 introduce BranchCache™ to enable access to data from within the branch when available — either from other clients that have already downloaded the data, or from a centrally provisioned server in the branch office. BranchCache by virtue of being an end to end solution can even optimize SSL, IPsec and SMB signing. This presentation will give an overview of the BranchCache architecture and describe how the SMB and HTTP protocol implementations have been extended to make use of this functionality.

ARCHITECTURAL DEEP DIVE OF THE LIKewise SMB SERVER

MONDAY 2:00 - 2:50

The Likewise CIFS server is a new introduction to the Likewise Open project sponsored by Likewise Software. This new code base is released under a combination of the GNU LGPLv2.1+ and GPLv2+ licenses. This session explain the architecture of the Likewise IO manager and related processes, provide code walk through of the relevant drivers, and explain how the IO paths can be customized for specific storage platforms. More information can be found at <http://www.likewiseopen.org/>. The source code is available from <http://git.likewise.com/>

DEVELOPING SCALABLE AND PORTABLE CIFS SERVER

MONDAY 3:15 - 4:05

Desktops, embedded devices, mobile devices, NAS solutions apply various and contradicting requirements for. CIFS implementation. This presentation discusses a scalable CIFS Server architecture which considers all of these; various target platforms introduce yet another challenge. This presentation introduces a flexible porting approach with examples from our CIFS solutions on more than 20 platforms.

SAMBA ARMED AND READY: RUNNING AN AD DC ON 2 WATTS

MONDAY 4:15 - 5:05

For small networks, a big beefy server is usually overpowered if it just needs to provide authentication and authorization services for a few users. This talk will describe how to get Samba4 running on an OMAP3530 ARM system to provide Active Directory services at 2 Watts. The talk will also give some benchmark figures to illustrate the performance of this embedded setup.

INTRODUCTION TO THE CIFS/SMB/SMB2 PLUGFEST

MONDAY 5:15 - 6:30

Abstract coming soon.

DATA MANAGEMENT TRACK

USING DATA CLASSIFICATION TO MANAGE FILE SERVERS

TUESDAY 1:00 - 1:50

One of the traditional challenges within file storage has always been understanding what kinds of data were consuming the storage of one's infrastructure. Furthermore, administrators have no easy way of implementing business policies (security, data management) based on the data. In this talk we will be presenting the 'File Classification Infrastructure' which Microsoft is shipping within Windows Server 2008 R2. FCI is a pluggable software framework intended to help storage developers create complete classification solutions that enable customers to act upon their data more intelligently. In this session we offer an overview of the extensibility model of FCI and we will present the functionality available to storage developers. We will also be demonstrating what some storage developers are already building upon this architecture. Developers — come learn what opportunities are available for you to build data management and security solutions that leverage a extensible data classification infrastructure. This session is intended to make you aware of the business and IT challenges around classification and equip you on how you can develop innovative data management solutions based on the File Classification Infrastructure.

LONG TERM INFORMATION RETENTION FORMAT

TUESDAY 2:00 - 2:50

As the world becomes digital, we may enter a digital "Dark Ages" in which business, public and personal assets are in ever greater danger of being lost. While preserving those assets includes various stakeholders and technologies, the storage component has a key role in this spectrum as it is where the data is located for most of its lifecycle. Recognizing this role, the storage industry initiated work on Self-contained Information Retention Format (SIRF) to help its customers interpret preservation objects in the future even by applications that do not exist today.

REAL WORLD EXPERIENCES CONVERTING TO XAM

TUESDAY 3:05 - 3:55

Clearpace undertook a redesign in 2008 to make use of the new XAM standards and virtualize access to the file system layer in our products. This session will cover the impact it has had on the engineering team and the company as a whole and give advice for others considering embarking on using the XAM apis.

SESSION DESCRIPTIONS

EFFECTIVE STORAGE TIERING FOR DATABASES

TUESDAY 4:05 - 4:55

Most databases occupy several files. Typical access patterns result in some heavily accessed files and some not. Moreover, these access characteristics change over time. This presentation describes a tiering mechanism for dynamically identifying hot and cold files and relocating them to suitable storage media, transparently to database managers and other applications. For example, when the “temperature” (access rate) of a file exceeds a threshold, the files might move to solid state disk. When the files “cool,” they return to lower-cost storage media. File relocation makes storage tiering for databases possible, without requiring that system or database administrators micromanage data placement.

CLOUD TRACK

CLOUD SEEDING: SOME PRACTICAL METHODS FOR MOVING DATA BETWEEN THE ENTERPRISE AND THE CLOUD

TUESDAY 1:00 - 1:50

This presentation describes some recent experiments in using commercial cloud storage in conjunction with customer-owned storage. Multiple techniques for data transfer will be discussed, along with functional and performance characteristics of each. Some possible use cases are to provide recoverability of data from site disasters and migration to or from cloud. Service level agreements, the economics of using cloud storage as a data replica and some implementation alternatives will be explored.

BUILDING A STORAGE SERVICE IN A PRIVATE CLOUD

TUESDAY 2:00 - 2:50

Storage in a private cloud is the storage that sits within a particular enterprise security domain and the data is considered more secure. This presentation discusses how we can build your own private cloud storage using off the shelf servers and open source components. This presentation walks through how we can leverage the different open source components that are available, the advantages and suitability of these components in building a private cloud storage solution with respect to (i) Types of Storage [NFS/CIFS/Block Storage] (ii) Backup solutions (iii) Scalability and (iv) Multi-Site Recovery, for building the storage solution within a private cloud. The presentation concludes with how a storage service was built in a private cloud using the open source components.

MANAGING STORAGE OF VIRTUAL MACHINES WITH MICROSOFT SYSTEM CENTER VIRTUAL MACHINE MANAGER (SC-VMM)

TUESDAY 3:05 - 3:55

Modern day Data Centers depend on unified management of physical and virtual machines, centralized resource optimization, consolidation of underutilized physical servers, and rapid provisioning of virtual machines (VM). Running IT infrastructure that is distributed across multiple sites and dependent on heterogeneous storage configurations poses some unique challenges including discovering this infrastructure and making sense out of it, reliably configuring VM\Host dependencies on the remote storage, finding creative ways to effectively assist with the typical administrative tasks (VM host patching, monitoring load, upgrading storage, etc), coping with the vendor specific interpretations of standard interfaces as well as the complete lack of standardization in certain aspects. VMM and SANs: Microsoft System Center Virtual Machine Manager (SC-VMM) solves these problems by enabling SAN-based creation, migration and backup of VM images. VMM auto-detects existing SAN infrastructure on the hosts, libraries or cluster configurations that VMM manages and orchestrates

iSCSI \ FC \ NPIV based VM migrations. Today VMM relies on VDS software and hardware provider APIs as well as the new NPIV interface defined in conjunction with the Windows team that received support from major HBA vendors who mapped the NPIV functionality of their drivers to this standard interface. Storage VM Migrations: In order to migrate VMs, VMM always chooses the best available method tuned to specific circumstances and environment - whether it is a LIVE / QUICK cluster-based migration or one of the SAN-based migrations or the QUICK STORAGE MIGRATION that copies virtual disks underneath of a running VM or a LAN-based option.

A SIMPLE AND SCALABLE VIRTUALIZATION PLATFORM

TUESDAY 4:05 - 4:55

Abstract coming soon.

SECURITY

KEY MANAGEMENT INTEROPERABILITY PROTOCOL - A STORAGE DEVELOPER'S APPROACH TO THE STANDARD

TUESDAY 1:00 - 1:50

Data in networked storage environments is significantly more vulnerable to unauthorized access, theft, or misuse than data stored in direct-attached storage. Encrypting data-at-rest can mitigate threats and allow data security, but widespread use of encryption is complicated by inconsistencies and duplication in key management systems supporting encryption environments. An emerging standard, Key Management Interoperability Protocol (KMIP), promises to greatly reduce OPEX and improve compliance visibility, and this session provides a developer's perspective on its implementation.

TCG - TRUSTED PERIPHERAL SIMULATOR FOR DATA-AT-REST

TUESDAY 2:00 - 2:50

Storage Security has become an important aspect in an enterprise due to the cost incurred due to data loss/theft. Trusted Computing Group (TCG) has defined a data-at-rest security standard for peripheral devices including storage disks. This presentation explains a product development accelerator that can help in reducing the cycle time in development of TCG enabled disks and ISV application's by using test qualifiers and simulators. A trusted peripheral (TPer) simulator helps ISV to develop applications that are compliant to TCG specification even in the absence of a storage device that is compliant to TCG. Similarly, a TPer qualifier quickens the testing of TCG compliant disk. This presentation explains the high level approach, the implementation details in Linux, interface with test automation framework and features of TCG storage specification that have been implemented.

HOST BASED STORAGE SECURITY - PROVIDING ADVANCED STORAGE SECURITY FOR VIRTUALIZATION-BASED DYNAMIC INFRASTRUCTURE ENVIRONMENTS

TUESDAY 3:05 - 3:55

Virtualization and converged networks are disruptive technologies changing how data centers are architected, managed, and secured. Virtual workloads are dynamic requiring dynamic storage access across any network. This makes current storage security models inadequate since traditional SAN security advantages are lost. This new environment demands an improved virtualization-aware host-centric solution to storage security — the Host based Storage Security architecture. In this discussion, IBM and Emulex will present a real world proof-of-concept of an HbSS solution.

SESSION DESCRIPTIONS

SELF-ENCRYPTING HARD DRIVES: FROM LAPTOPS TO THE DATA CENTER

TUESDAY 4:05 - 4:55

Learn how to deploy and manage encrypting drives, the newest tool in data center and client-side security. This session will review basic concepts of encrypting hard drives and underlying industry specifications that enable these drives to secure data in an interoperable fashion. Attendees will also learn about basic key management concepts related to managing these drives and how to test software to remotely provision drives, manage keys and decommission drives that are out of use.

CIFS/SMB/SMB2 TRACK

UNDERSTANDING WINDOWS FILE SYSTEM TRANSACTIONS

TUESDAY 1:00 - 1:50

Windows Vista introduced the concept of user controlled file system transactions into the NTFS file system. This allows applications to group logical units of work together into transactions which can be committed or rolled back atomically and are isolated from the rest of the system. This talk will cover the design philosophy behind file system transactions, and describe how they can be used effectively in applications. We will discuss previous solutions for crash recovery for applications and compare and contrast transactions as an alternative mechanism.

SMB v2.1

TUESDAY 2:00 - 2:50

SMB v2.1 is a minor revision to the SMB protocol that looks to improve the applicability of the protocol in both branch office and enterprise networks. This talk will examine the individual changes in the revision and the benefits/challenges they provide with a goal of assisting implementers in adding this functionality to their solutions. Time permitting, we will also talk about behavioral changes your client or server implementation may see when interacting with Windows 7.

SMB2 MODEL BASED TESTING

TUESDAY 3:05 - 3:55

The SMB2 test suite was developed by Microsoft to validate the MS-SMB2 protocol documentation and behaves as a synthetic client. This session provides an overview of the test suite that will be available during CIFS/SMB/SMB2 Plugfest to test participants' SMB2 implementations. The adapter architecture is based on a complementary framework, required to generate synthetic protocol traffic and to observe traffic coming back from the server. During the talk, the SMB2 test suite protocol model will be explained and Spec Explorer will be used to explore states in a particular scenario. The test suite will be executed to demonstrate how the scenarios are turned into code that tests a server implementation. We will analyze the test suite log data resulting from these tests.

FILE SERVER CAPACITY TOOL (FSCT) FOR CIFS/SMB/SMB2

TUESDAY 4:05 - 4:55

This presentation describes "FSCT", a CIFS/SMB/SMB2 file server tool for capacity planning and bottleneck identification. This tool was created by Microsoft after analyzing the workload of production file servers. "FSCT" can simulate that typical file server workload in a lab environment. The presentation will include a demonstration of how to use the tool, including some sample results.

SOLID STATE STORAGE TRACK

GETTING THE MOST OUT OF SSDS - IT SYSTEM OPTIMIZATION BEST PRACTICES

WEDNESDAY 1:30 - 2:20

SSD adoption in enterprise data centers will increase dramatically over the next few years, due to SSD's significant advantages compared with HDDs: dramatically higher I/O performance, smaller footprint and reduced energy consumption (power/cooling). This presentation will provide guidelines and best practices to help server and storage system developers architect future IT systems to take full advantage of these performance and efficiency gains, while simultaneously meeting the increasing enterprise environmental and efficiency requirements.

DIFFERENTIATED STORAGE SERVICES

WEDNESDAY 2:30 - 3:20

LSI & Intel are working on a project to examine the proposition that that a storage device can better optimize performance if given information concerning the type of data it is handling. While SSD are becoming more affordable, they are still expensive. Can the storage system more efficiently utilize the performance advantages of SSD if applications and file systems provide information to the storage system about which type of data is associated with each block access command?

METHODOLOGIES FOR CALCULATING SSD USEABLE LIFE

WEDNESDAY 3:35 - 4:25

Due to shrinking process geometries and storing more bits per cell on NAND flash, the primary media in solid-state drives (SSDs), the net effect has been a reduction in endurance and data retention specifications. SSDs are still a very robust, reliable storage solution for many embedded system OEM applications. This tutorial will arm designers with tools, guidelines and metrics to accurately measure and predict SSD useable life for a wide range of application usage models.

FROM 512 TO 4K:

A CASE STUDY IN SUPPORTING LARGE SECTOR SIZE SSDS IN SOLARIS

WEDNESDAY 4:35 - 5:25

The advent of SSD represents a sea change in storage system which brings exceptional performance and reliability compared to the HDD. The optimal sector size of SSD is typically 4KB, it's different from traditional 512 bytes ones. We should make changes to accommodate this new trait. In this presentation, The speaker will give a case study of supporting large sector size SSDs in Solaris, including the motivation and background to change from 512 to 4K, the design and implementation in Solaris host driver stack and the performance gain introduced by this new design.

SESSION DESCRIPTIONS

STORAGE MANAGEMENT TRACK

THIN PROVISIONING AND RECLAMATION OF UNUSED STORAGE

WEDNESDAY 1:30 - 2:20

Thin provisioning has proven so useful in improving storage utilization that most storage vendors offer the capability in some form. But every storage vendor manages “thin” LUNs differently. Thus, it becomes the responsibility of host-based applications to manage “thin” LUNs, and in particular, to detect unused storage capacity and reclaim it for deployment elsewhere. This presentation describes a host-based file system and volume manager that cooperate to detect space unused by file systems and reclaim it transparently to applications using the storage. The solution is hardware-independent in that it supports major storage vendors’ thin provisioning capabilities, and is readily adaptable to others. Moreover, the solution accommodates advanced data management capabilities such as data migration between fully-provisioned and thin LUNS, and snapshots.

SMI-S RECIPE INTERPRETER

WEDNESDAY 2:30 - 3:20

A SMI-S Recipe Interpreter is provided in the presentation. SMI-S standard uses the Recipe describe typical sample operation to the profile. As parts of the certificate compliance tests any SMI-S implementation need pass the recipe validation. Today’s validation tools like CPT require manually coding the recipe logical into JAVA to implement the test. In the presentation, a recipe interpreter is provided which can understand and generate validation points from Recipe directly. The method can make the recipe tests more effective and efficient.

CHALLENGES IN SMI-S PROVIDER DEVELOPMENT

WEDNESDAY 3:35 - 4:25

This presentation describes the process of SMI-S proxy provider development for a performance sensitive multi-device environment. It explains the concepts involved by taking into consideration an actual provider development scenario for a RAID solution. It highlights the decision process in choosing between proxy provider and embedded provider, development tools like CIMPLE vs. CMPI framework and choosing a suitable design while developing the provider. The presentation goes on to explain approaches to improve provider response time by having an efficient approach for caching device details, multi-threading support etc. The presentation concludes by throwing light on resolving issues in a multi-vendor scenario during provider development specifically around triage and defect turnaround times.

MICROSOFT STORAGE MANAGEMENT UPDATE

WEDNESDAY 4:35 - 5:25

A year ago Microsoft presented our perspective on SMI-S in the Windows environment. We have made a lot of progress over the year and we are working to integrate SMI-S into the Windows management infrastructure. The System Center Cross-platform Storage Management Team will be presenting our vision and updating the community on our work.

ISCSI TRACK

iSCSIsim

WEDNESDAY 1:30 - 2:20

iSCSIsim is a powerful generic open-source test tool for iSCSI targets. Written entirely in Python, its capabilities include SCSI command injection, error insertion, and recovery testing: all with built-in validation of test results. Command sequencing and task management are also supported, along with a new proxy function which simulates loss-of-connection. iSCSIsim has shown itself to be useful for developer testing and could be the basis of an iSCSI conformance test suite for your organization. This presentation will discuss the architecture of iSCSIsim, how to install and use the tool, and how to create custom test suites for testing your own iSCSI target tailored to the specific needs of your organization. The presentation will finish with a discussion of areas of ongoing development for iSCSIsim, as well as how you can get involved.

iSCSI TESTING: WHAT ARE THE TEST CHALLENGES UNDER THE HOOD OF A 10 GB iSCSI STORAGE PRODUCT CERTIFICATION?

WEDNESDAY 2:30 - 3:20

The certification of a 10 Gb iSCSI RAID Storage System elicits a lot of challenges at the development level and the Test / Quality Assurance level. The challenges are due to the fact that a 10 Gb iSCSI is a newly deployed iSCSI host interface in the RAID Storage environment. As a result the size of a development module level test should be designed very carefully to establish a test coverage beyond basic implementation verification, standard RAID testing, or the iSCSI plug fest. These module level tests must tackle the test time windows associated with the following iSCSI characteristics: 1.Device discovery, 2.10 GB switch traffic control and congestion, 3. Security mechanisms with different Operating systems, 4.Operational parameters associated with I/O retries and recovery 5.Management, Administration, and Integration with Storage products 6.Design For Testability “DFT” mechanisms 7. Diagnostics, problem Isolations 8.IPV4 vs. IPV6 However a number of the module tests above can be leveraged from the certification a 1 Gb iSCSI RAID products. There are specific features such as backup, snapshot, remote mirroring, and cluster application compatibility that must be supported by the RAID product and must be verified during the testing of the RAID controller host interface.

CIFS/SMB/SMB2 TRACK

BUILDING A CLUSTERED NAS SYSTEM

WEDNESDAY 1:30 - 2:20

This presentation will describe the design of a cost-effective clustered NAS system using readily available off-the-shelf hardware and software components. The presentation will describe the underlying cluster technology, and show how it was adapted to provide scalable NFS and CIFS access to one or more file systems. Design choices throughout the development and the rationale for them will be presented, including a novel method for delivering scalable CIFS service that uses DFS technology. Performance, availability, installability, supportability, and the challenges of delivering a “software storage appliance” will be touched upon.

SESSION DESCRIPTIONS

DFS-N OVERVIEW AND SCALABILITY

WEDNESDAY 2:30 - 3:20

DFS Namespace (DFS-N) is a storage management solution that gives administrators a more flexible way to centrally manage their distributed file server resources. An overview of DFS-N will be presented and results and learned lessons on scalability and performance comparing different versions of DFS-N will be discussed. The focus will be on a technical understanding of the applicability and scalability aspects, enabling storage administrators to take advantage of the technology.

SMB/CIFS ACCESS CONTROL AND IDENTITY MAPPING IN OPENSOLARIS

WEDNESDAY 3:35 - 4:25

The OpenSolaris CIFS service leverages ZFS access control and idmapd identity mapping to create a unified, ubiquitous access management interface for all OpenSolaris users, whether accessing the operating system locally in a shell or via SMB or NFS. This presentation will cover details of ZFS ACLs, idmapd identity mapping and how these components support Windows interoperability with OpenSolaris.

PERMISSIONS MAPPING BETWEEN CIFS ACLs, NFSv4 ACLs AND POSIX MODE BITS IN ISILON'S ONEFS FILE SYSTEM

WEDNESDAY 4:35 - 5:25

There are a few authorization domains used widely today: POSIX mode bits and CIFS/NFSv4 Access Control Lists. Creating a file system which intelligently maps between different sets of file permissions is tricky. I will explain how we accomplish this in Isilon's OneFS operating system. By mapping ACLs from one authorization domain to the other, our clustered file server provides one unified permissions model for CIFS, NFSv3 and NFSv4 clients.

MANAGEMENT TRACK

DMTF PROFILES FOR STORAGE

THURSDAY 8:30 - 9:20

DMTF has several standards and technologies that apply to storage management, many of which are leveraged by the SNIA SMI-S standard. This talk will cover some of the basic profiles that can be leveraged by implementations of standards beyond SMI-S, such as the recent SMASH and DASH standards that the DMTF has created. Attendees will leave with a good understanding of how DMTF technologies can be leveraged in their own products.

PROFESSIONAL DEVELOPMENT TRACK

BUILDING TRUST WITHIN THE TEAM

THURSDAY 9:30-10:20

Building Trust will provide you with tools that will raise your level of awareness around trust with yourselves and others. Trainees will have an opportunity to look at what trust means to them and develop skills around building trust. This training is about the single most significant factor in everyone's life! Every single relationship that we have is build upon trust. This seminar directly addresses trust in our families, friends, community, business, and the entire world. In high trust-based environments everything moves more quickly and productivity is vastly increased. This session will provide practical techniques and practices to increase the teams trust which will directly impact and decrease development costs.

GLOBAL TEAM MANAGEMENT BEST PRACTICES

THURSDAY 10:45-11:45

Geographically distributed global engineering teams is now a norm in the IT, Software, and Storage industries. Yet operating such teams in an optimal manner continues to be a challenge. This workshop is developed and delivered in collaboration by Abhinav Jawadekar, founder of Sound Paradigm (www.soundparadigm.com) and Tom Hickman, founder of Inside Outsource Consulting (www.isosconsulting.com). Abhinav has substantial experience as a service provider in India and Tom has worked as an outsourcing manager from the US; these perspectives combine to offer a unique understanding of global teams, and the talk shares best practices derived from hands on experience in building and managing global software engineering teams over a combined 30+ years of experience in the industry.

FILE SYSTEMS TRACK

RELIABLE LOCKING FOR CLUSTERED NFS SERVERS

THURSDAY 8:30 - 9:20

NFS servers implement the Network Lock Manager protocol so that multiple clients can share access to files and directories without risking corruption due to competing updates. In a single NFS server, implementing NLM is relatively straightforward, but in clustered servers, things become more complicated during normal operation and even more so when failure of a cluster node necessitates lock recovery. This presentation briefly describes NLM protocol semantics and operation, outlines the complications introduced with clustered NFS servers, presents some typical solutions and issues with them, and concludes by describing a novel solution which both outperforms and is able to deliver more robust guarantees of correctness than conventional approaches.

NFS ON THE FAST TRACK - FINE TUNING AND FUTURES

THURSDAY 9:30-10:20

This talk highlights client side performance tuning over NFS, with Linux and NetApp storage best practices for particular workloads. We also discuss (a) gathering application level requirements and choosing the right backend storage for different workloads, (b) setting up the right mount options for NFS clients and determining the right network switch port requirements. We conclude with a summary of the current state of the NFSv4 standard, including implementation progress for several platforms.

OPENSOLARIS PARALLEL NFS (pNFS): BLENDING PERFORMANCE AND MANAGEABILITY

THURSDAY 10:45-11:45

Parallel NFS or pNFS is an extension to NFSv4 minor version 1 to support scalable, parallel I/O by providing open protocol support to separate a file's data from its metadata. The goal of the pNFS implementation in OpenSolaris is to provide scalable performance as well as novel data management features which draw from the numerous benefits of data, metadata separation.

PERFORMANCE TRACK

IO PERFORMANCE IMPROVEMENT IN A VIRTUALIZED ENVIRONMENT

THURSDAY 8:30 - 9:20

IO performance in virtualized environment has always been a challenge owing to “in-direct” access to the hardware. This has been addressed using “pass-thru” techniques to directly access the hardware; however this does not allow migration of tasks. This presentation introduces a virtual interface technique that is used to route IO via this interface before the migration is initiated. This brings the IO performance in virtualized environment very close to the native driver performance along with support for all key features of virtualization. The presentation goes on to explain how this is achieved using a 10Gbe adapter in a Xen environment.

“BEST OF FAST” TRACK

CA-NFS: A CONGESTION-AWARE NETWORK FILE SYSTEM

THURSDAY 9:30-10:20

Award winning paper from the USENIX File & Storage Technology Conference!

We develop a holistic framework for adaptively scheduling asynchronous requests in distributed file systems. The system is holistic in that it manages all resources, including network bandwidth, server I/O, server CPU, and client and server memory utilization. It accelerates, defers, or cancels asynchronous requests in order to improve application-perceived performance directly. We employ congestion pricing via online auctions to coordinate the use of system resources by the file system clients so that they can detect shortages and adapt their resource usage. We implement our modifications in the Congestion-Aware Network File System (CA-NFS), an extension to the ubiquitous network file system (NFS). Our experimental result shows that CA-NFS results in a 20% improvement in execution times when compared with NFS for a variety of workloads.

GENERATING REALISTIC IMPRESSIONS FOR FILE-SYSTEM BENCHMARKING

THURSDAY 10:45-11:45

Award winning paper from the USENIX File & Storage Technology Conference!

The performance of file systems and related software depends on characteristics of the underlying file-system image (i.e., file-system metadata and file contents). Unfortunately, rather than benchmarking with realistic file-system images, most system designers and evaluators rely on ad hoc assumptions and (often inaccurate) rules of thumb. Furthermore, the lack of standardization and reproducibility makes file system benchmarking ineffective. To remedy these problems, we develop Impressions, a framework to generate statistically accurate file-system images with realistic metadata and content. Impressions is flexible, supporting user-specified constraints on various file-system parameters using a number of statistical techniques to generate consistent images. In this paper we present the design, implementation and evaluation of Impressions, and demonstrate its utility using desktop search as a case study. We believe Impressions will prove to be useful for system developers and users alike.

WINDOWS 7 TRACK

DELETE NOTIFICATION IN WINDOWS 7

THURSDAY 8:30 - 9:20

Microsoft's Windows 7 will include support for the T13 Data Set Management TRIM attribute. This allows storage device to be notified when files are deleted and space is freed up. This talk will cover the details of how devices get notified. We will also discuss how and when the file system stack issues this command and assumptions the file system stack is making.

OPEN SOURCE TRACK

ORCHESTRATING OPEN SOURCE COMPONENTS FOR HOME SANs

THURSDAY 9:30-10:20

Homes already have huge amounts of data and multiple networked computers/devices. The need for consolidation and central management is pressing. While a number of open source components are available to put together cheap SANs, their orchestration remains a challenge. This session discusses how different Linux components (MD - Software RAID, LVM, iSCSI Enterprise Target) can be put together and made usable as an iSCSI SAN storage array (or appliance) along with live demonstrations based on virtual machines running on the presenter's laptop.

INTEGRATED OPEN SOURCE & PROPRIETARY STORAGE SOLUTIONS

THURSDAY 10:45-11:45

Enterprises constantly face the challenge of choosing solutions (including storage solutions) that fulfill their present and future requirements, while curtailing their Capital Expenditure. CIOs and their teams are cautious about using open source products due to concerns around integration with other open source and proprietary Storage products. We are proposing a Architecture using optimized combination of Open Source Applications (Zimbra MailServer), Proprietary Storage products and an integration 'glue' code. This architecture encompasses Disaster Recovery, Data Protection (from online/offline) and Hierarchical Storage Management aspects as well. The proposed solution has been deployed in a production environment.