

Enterprise Storage

Leah Schoeb, Vice Chair, SNIA Technical Council

SNIA Emerald[™] Training

SNIA Emerald Power Efficiency Measurement Specification, for use in EPA ENERGY STAR[®]

June 24-27, 2013









Day 1 of this course presents an overview of storage technology, efficiency, and performance.



Course Information



Who should attend?

- Information Technology professionals
- Engineers
- Consultants

Objectives – what you will learn

- Basics of enterprise storage technology
- What are the initiatives for optimizing the data center
- Current efficiency technologies used in storage
- Understand Storage Performance basics
- IO Generation tools are not all created equal





- Enterprise storage overview
- Break
- Lab Tour
- Enterprise storage efficiency
- Enterprise Storage Performance and load generation
- ♦ Q&A



What is a Storage?



Storage is more than a device

- In general, we think of a box
- We look at the device attributes
 - > Capacity
 - > Performance (speeds and feeds)
- It's really about storing and retrieving information with integrity, availability, data protection, cost
- Need a strategy and efficiency to address how we deal with information





Source: IBM



Storage Technology BLOCK VS. FILE VS. OBJECT



SNIA Emerald[™] Training ~ June 24-27, 2013



File I/O vs. Block I/O vs. Object I/O

- Applications can do block I/O or file system I/O or object I/O
- File systems turn file I/O into block I/O
- Block I/O goes to specific device and reads or writes a block from/to that device
 - > Linear address space of blocks
 - > May do multiple blocks in single operation
 - > Typically fixed length blocks
- File I/O is represented by a file with file name and some offset into the file
 - > Read or writes data in the file
 - > Some number of bytes involved in the operation
- Object I/O is storing data as objects with new control/metadata information





Block I/O

- Application writes data block
- Block goes to HBA and over storage interface
- Storage controller receives block
- Data written to device as data block





Tape I/O

- Backup application writes data block to tape driver
- Block converted to tape image and goes to HBA and over storage interface
- Tape controller receives tape block
- Data written to tape as tape image block



Tape I/O



Tape Library I/O

- Backup application writes data
- Block converted and tape volume identified
- Tape library receives tape block and volume information
- Data written to selected tape as tape image block





Virtual Tape I/O

 Block converted to tape image

- Tape image written to disk controller
- Depending on controls,
 VTL reads tape image from disk and writes to tape library
- A few products can go direct to tape





Block I/O vs. File I/O





File I/O

- Application writes data block to a mounted file system
- Block goes to HBA and over storage interface
- Storage controller receives block
- Data written to device as data block
- Many Protocols



File System I/O







- Application writes file data to mounted file system
- File redirected over network
- NAS / File server takes file and converts to block I/O
- Data written to device as data block





Object I/O

- Application writes object information
- Object file system creates attributes and sends object to HBA / NIC
- Storage controller receives object
- Data written to device as data block







Scale out vs scale up



SNIA Emerald[™] Training ~ June 24-27, 2013

Scale Out vs. Scale Up







Implemented using distributed file system usually

- Clustered hardware
- Global namespace across nodes

Key implementation points

- Linkage between nodes InfiniBand, Ethernet:10Gb
- I/O balancing between nodes
- Capacity balancing between nodes
- Coherency across distance
- Switching requirements



Scale Out Storage: Block Storage



- Two approaches used
- Multiple controller cards to common backend storage device pools
 - Backplane connected typically
 - Normally associated with high-end enterprise systems
- Federation of separate controller nodes (included with NAS systems in some cases as integrated unified storage)
 - Complexities in cache coherency and I/O routing
 - Vendor differences





Storage Technology REPLICATION



SNIA Emerald[™] Training ~ June 24-27, 2013





Migrate data at same or different location

- New technology or added capacity
- Consolidations
- Mergers
- Workload balancing
- Application testing
- Data exchange for warehousing / mining
- Business continuance, Disaster Recovery, Data Protection
 - Motivated by:
 - > Business economics
 - > RPO and RTO



> Legal requirements

SNIA Emerald[™] Training ~ June 24-27, 2013

Replication Technologies



Point-in-Time Copy – Snapshot / Flashcopy

- Space efficient copy only changed data
- Variations in storage system implementations

Remote replication

- Synchronous continuous transfer of all data changed, operations wait until transfer completes
- Asynchronous data sent but operation continues without waiting for transfer completion
- Asynchronous periodic data sent periodically, most often is incremental snapshot





Snapshot

- Copy-on-write only changed data is copied
- Redirect-on-write writes to new location
- Various implementations use pointer manipulations, side files, or dynamically allocated space from storage pool

Cloning

- Makes a complete copy of data
- Clone may be continuous or split and later resynced
- Variations are differentiation points for vendors



Remote Copy Approaches



Synchronous

- Storage system-based or through fabric appliance
- Data must be stored at remote site before application can resume
 - > Significant impact on performance limits distance
- Application stops if I/O can't complete
 - > May stall application

Asynchronous

- Initially host software based, but system based is now available
- Application continues before data is stored on remote site
- Delayed write represents a window of "risk"
- continuous or periodic transmission of data

Multi-hop or cascading

 Combination of synchronous and asynchronous with intermediate storage system





Storage Technology



SNIA Emerald[™] Training ~ June 24-27, 2013

Archiving Definition



- Different than backup
- One time operation
- Protection is done at time of archiving
 - Multiple copies are made
- Indexing for searching done at time of archive
- File or object based element not a transformation by backup software
- Backup does a restore, Archive does a retrieve



Archiving Benefits



- Economic efficiency improvements
- Reduction in primary capacity demand
 - Continuing economic benefit
- Reduced amount of data for regular backups
 - Fewer operational expenses for data protection
 - Reduction in backup capacity demand
- Reduced capacity-based licensing charges for software and data management services
- Reduction in power/space/cooling based on different type of archival storage



Archiving Overview







Storage Technology VIRTUALIZATION



SNIA Emerald[™] Training ~ June 24-27, 2013

IO Performance Needs Monitoring at Every Level





Server Virtualization Impact on Storage

- Increased workload demand
- Over 80% of storage related performance problems stem from misconfigured storage hardware
- Tendency to over provision storage
- Mixed workloads







VMware APIs for Array Integration - VAAI

- Full copy Enables the storage system to make full copies of data within the storage system without having the ESX host read and write the data.
- Hardware-assisted locking Provides an alternative means to protect the metadata for VMFS cluster file systems
- **Block zeroing** Enables storage systems to zero out a large number of blocks to speed provisioning of virtual machines.
- Thin Provisioning support
 - > TP-STUN error code to report out of space for thin volumes
 - > UNMAP zero page reclaim to maintain "thinness"
- NAS support
 - > Full file system clone uses NAS system to make copies
 - Reserve space creation of thick virtual disk on NAS
 - > Native snapshot support

4.1

5.0

VAAI Performance Affects





SNIA Emerald[™] Training ~ June 24-27, 2013

SNIA Emerald

VMware vSphere APIs for Storage Awareness - VASA



- Administrator can see capabilities of storage system
- Allows creation of VM Storage Profiles for Profile Driven Storage
- API for storage systems to communicate their capabilities to vSphere
 Storage Capabilities & Profile-Driven Storage capabilities & Profile-Driven Storage
 API for storage systems to communicate their compliant
 - RAID protection
 - Thin provisioned LUN
 - Deduplication capability
 - Tiered storage capability
 - Capacity
 - Storage health

Vendor plugins – Storage Providers





os



VM Storage Profile referencing Storage Capabilities

Storage Capabilities surfaced by VASA or user-defined



SMB Persistence

 Ability to create a continuously available persistence file share without changes to a way a file is opened

SMB Scale-Out

- In a cluster configuration this new ability shows a share in all nodes of a cluster.
- Active/Active configuration makes it possible to scale file server clusters (not to be combined with other features)
- Cluster Shared Volumes (CSV) A volume that shows on every cluster node simultaneously so all nodes can access the volume
- Dynamic Network Name (DNN) Single name to connect
- Less IP addresses needed for the cluster

SMB Encryption

- No PKI or certificates required
- Accelerated Encryption optimized their algorithm





SMB Multichannel

- Transparent Failover for network paths Allows for an SMB client to continue working when there is a failure in a SMB file server cluster node.
- Interface Arrival Mulit-channel will re-establish connections automatically
- Link-local These are special IP address that are assigned automatically when there are no manual IP addresses configured or no DHCP server is found on the network

SMB Direct

- Discovery over TCP/IP and Fail back to TCP/IP, shift from TCP/ IP to RDMA
- the ability to use RDMA network interfaces, high bandwidth, low latency



SNIA Emerald™ Training ~ June 24-27, 2013



Volume Shadowcopy Services (VSS) for SMB File Share

- Same model as block VSS
- Stream snapshots from file server creates application consistent snapshots of data volumes.
- Server Manager
 - SMB file shares can be created with a set of wizards
- SMB PowerShell
 - Used to manage a significant number of file shares
 - Automatically checks the cluster type based on disk type
 - Combined view of the cluster with all nodes





VSA - Virtual Storage Appliance





Storage Technology on the Horizon



SNIA Emerald[™] Training ~ June 24-27, 2013

Big Data and Data Analytics Overview



Big data

- Storage for large amounts of data
- Data Analytics
 - Analytics against very large amounts of data
- Usually from machine-tomachine data
 - Called pervasive computing
- So, what does this mean for storage and IT?







- Distributed computing for analytics (Hadoop, for example) is moving from science experiment to mission-critical
- As this happens, data encompassed by these applications becomes the responsibility of people who worry about:
 - Security
 - Data protection/disaster recovery/business continuance
 - Data governance and compliance
 - Digital records management and archiving



Distributed, Shared-Nothing Architectures for Big Data Analytics





SNIA Emerald[™] Training ~ June 24-27, 2013

Storage-As-A-Service





Software Defined Storage



Software-Defined Storage





Technology Summary



Storage is critical to a business

- Having a storage strategy can mean real economic savings
- Potential to impact the corporation without adequate plans and procedures

Technology changes quickly

- Need to plan for changes
- Represent opportunities
 - > Costs
 - > Competitive advantage



Technology Summary



Storage needs to be dealt with on an economic value basis

- Factor in elements of availability, administration costs, acquisition cost, facility cost, energy cost. etc.
- Decisions based on economics
 - > Vendor offerings
 - > Long-term implications





Thank You

Leah Schoeb leah@evaluatorgroup.com Twitter:@vLeahSchoeb



SNIA Emerald[™] Training ~ June 24-27, 2013