Education

Understanding Enterprise NAS

Anjan Dave, Principal Storage Engineer
LSI Corporation

Author: Anjan Dave, Principal Storage Engineer, LSI Corporation

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.
  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Abstract

Understanding Enterprise Network Attached Storage (NAS)

- With the continuous growth of unstructured data, it has become paramount for enterprise storage stakeholders to understand the features and benefits of enterprise NAS solutions, and differentiate between the values of scale-out and scale-up NAS storage. This tutorial will help the audience gain insight into the usefulness and effectiveness of some of the key differentiating features of today's enterprise NAS offerings

Related Tutorials

Check out SNIA Tutorial:

What's Old is New Again - Storage Tiering

Check out SNIA Tutorial:

The File Systems Evolution

Check out SNIA Tutorial:

pNFS & NFSv4.2; a Filesystem for Grid, Virtualization and Database

Agenda

- NAS Overview
- Enterprise NAS Flavors
- Enterprise NAS Requirements – Performance, Scalability
- Caching in NAS
- Storage Tiering in NAS
- Scale-Up NAS Solutions
- Scale-Out NAS Solutions
- About De-duplication

NAS Overview

- Traditional NAS – what is it?
  - Storage accessible over the network, usually over IP
  - From simple partition sharing to traditional file servers to dedicated NAS appliances
  - Well suited for workgroups/departments, SMB environments
  - Can be multi-purpose (i.e. combine with email, print, other services)
  - Not scalable in performance or capacity

NAS Overview

- Enterprise NAS – what is it?
  - Purpose built - to serve structured & unstructured data over one or more protocols
  - Scalable in capacity and performance
  - Higher performance
  - Can support large number of clients
  - Better redundancy
  - Enterprise features - tiering, caching, de-duplication, multi-tenancy, replication, multi-protocol support, etc
  - Well suited for high-performance needs, large datasets, large number of clients

NAS Overview

- Types of Enterprise NAS Solutions
  - Traditional Scale up (add disks and upgrade NAS "head")
    - Appliance based (with integrated or 3rd party storage)
    - Proprietary solutions (native file processing and storage)
    - Open-Source software (with commodity hardware & storage)
    - Clustered NAS
  - Scale-Out
    - Software based (or software on appliance) using commodity hardware (Commercial and Open-Source)
    - Proprietary Integrated File Processing and Storage
  - Special purpose
    - Some software programmed into Integrated Circuits
    - NAS Cloud Gateway

NOTE: NAS Head refers to NAS Processing Unit, wherever applicable

NAS Overview - Examples

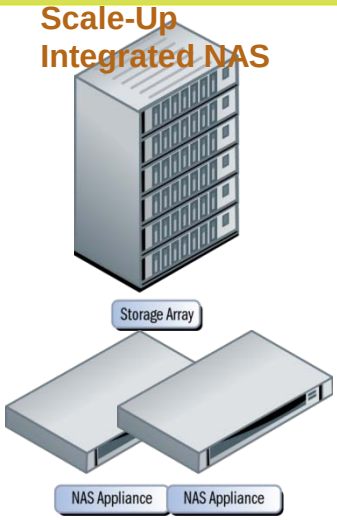

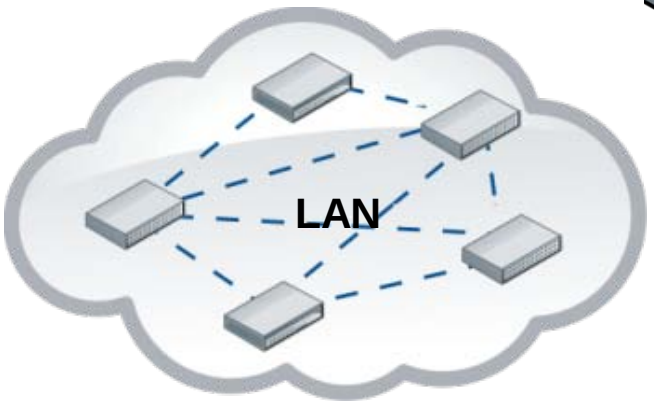

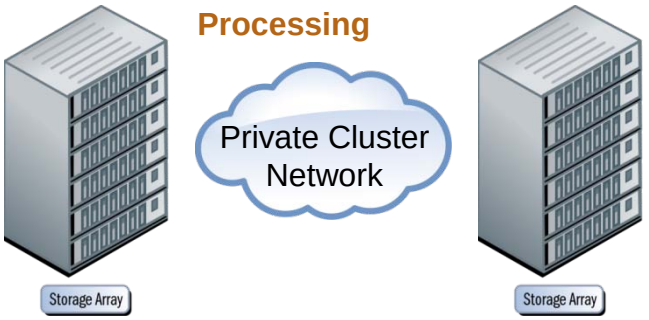

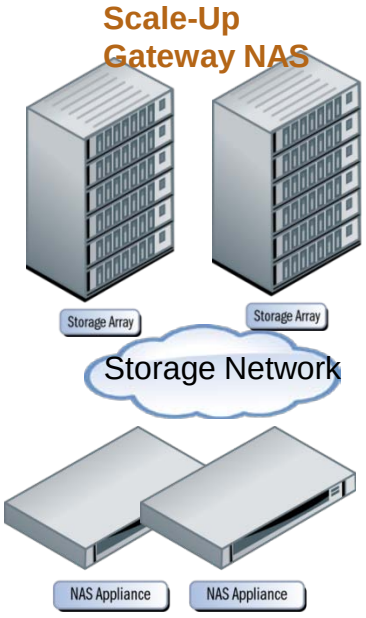

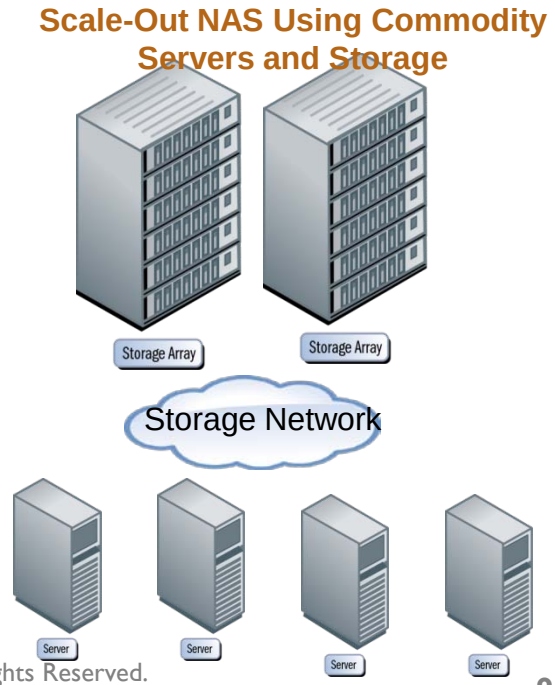

NAS Overview

- Enterprise NAS Features – Today
  - Clustered NAS (scale-up or scale-out)
  - PetaByte scale, 1000s of disk drives
  - Support for drive mixing
  - Storage Tiering
  - Enhanced Caching
  - De-duplication (file or block level, or both) and compression
  - Single Name Space
  - Multi-Tenancy

Enterprise NAS Requirements

- Performance Requirements
  - Types of workloads
    - High Disk IOPs driven applications (disk driven)
    - Very High NFS (or CIFS) Operations ('head' driven)
    - Cache friendly or unfriendly workloads
    - Small files Vs large files and Sequential Vs Random workflows
    - Replication and Backup workloads
    - De-duplication and other overheads

Enterprise NAS Requirements

- Performance Areas to Watch
  - "Head" Vs "Disk" contention
    - Dedup, replication, backups could cause contention on head or disks or both
  - Head memories - how much, where?
    - File Processing, FC, CPU, Network processing, NDMP
  - Caching (more on this in other slides)
    - 'Front-end' Cache to accelerate performance (comes in various forms) could become an overhead
    - Are the workloads cache friendly?
    - Does the system allow cache utilization reporting?
    - Is it tunable?

Enterprise NAS Requirements

- Performance Areas to Watch
  - Contention on the LAN
    - Load balancing across multiple 1GigE trunks using Link Aggregation
    - Use 10GigE if possible, combine with 1GigE for failover

Enterprise NAS Requirements

- How to gauge performance?
  - Create your own performance benchmarks to find a baseline
  - Find worst case and best case scenarios based on system configurations
  - Put it to test for worst real-world workloads
  - Do this in "production" configuration, not in test configuration
  - Test for performance during a component failure scenario
- Combine it with backups, replication, dedup, etc
- Understand the penalties of failover of "head"

Enterprise NAS Requirements

- Enterprise NAS Requirements - Scalability
- Long Term Concerns –
  - What is your growth rate?
  - Can you visualize your NAS/Storage infrastructure in next 5 years?
  - How do you plan to do Data Management across the entire NAS footprint?
  - How would you manage your NAS Infrastructure today and in 5 years?

Enterprise NAS Requirements – Scale-Up NAS Silos

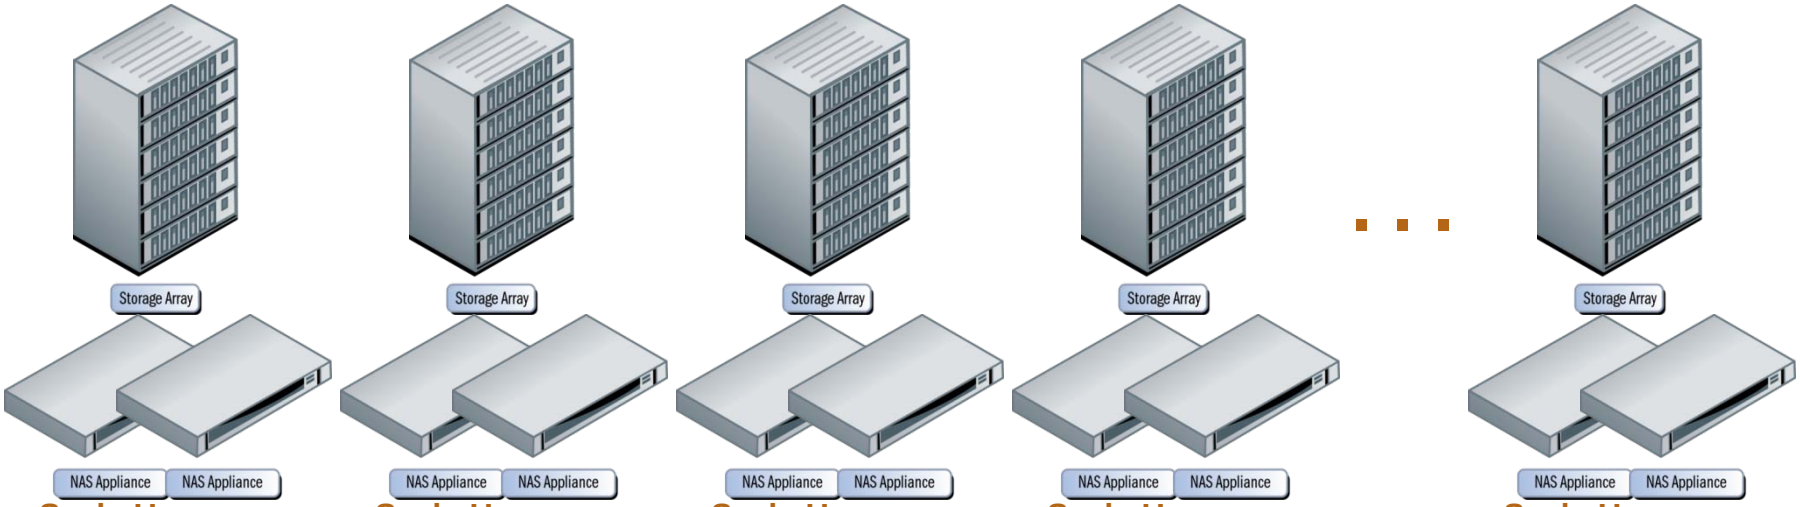

Scale-Up NAS system 1
25TB
SAS Only
Feature set 1

Scale-Up NAS system 2
30TB
SAS+NL-SAS
Feature set 2

Scale-Up NAS system 3
35TB
SAS Only

Scale-Up NAS system 4
60TB

Scale-Up NAS system 255
30TB
Feature set N

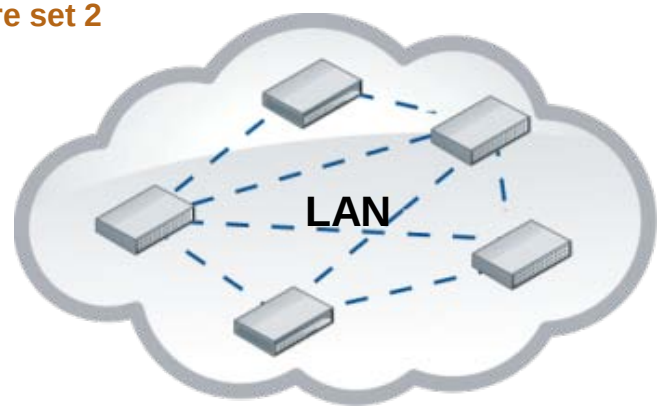

Separate NAS Entities resulting in storage inefficiencies and sprawl over time
(Example capacities)

Enterprise NAS Requirements

- Short Term Concerns –
  - Are you adding to or replacing existing solution?
    - How much performance and disk capacity are you planning for?
  - How would you add more performance?
  - How would you add more capacity? Independently or tied to performance?
  - How much capacity is too much on a single system?

Caching in NAS

- About "Front-End" Cache Implementations
  - Examples -
    - External Caching Device - between clients and NAS
    - Internal, based on PCIe SSD adapter
    - Internal as a set of SSDs
    - External Metadata Acceleration
    - External Extended NAS configuration (hub and spoke)

Caching in NAS

- "Front-End" Cache – Considerations
  - Usually, 'learning' or adaptation period is involved
    - Can take minutes to weeks depending on workloads
    - Size of the cache is a factor
  - Sizing front-end cache (if it's an option) against spinning disks can be challenging - due to costs
  - Determine Workload suitability for "hit" rates
    - Known or predictable Vs unknown workloads can impact cache efficiency
    - Structured and unstructured data
    - Reads Vs writes, random Vs sequential
  - Serves Reads, Writes, or both?

Caching in NAS

- "Front-End" Cache – Considerations
  - PCIe and SSD implementations have difference in performance/scalability
    - PCIe based usually is faster, but not scalable
  - Is it volatile to power or system exceptions?
    - Re-learning in case of power failures (can take a long time to re-learn, volatile Vs non-volatile)
  - Data promotion/eviction or maintaining multiple cache copies across cluster nodes could become an overhead for certain workloads

Caching in NAS

- About Storage Tiering and Caching
  - Both could be SSD implementations
  - Cache is reactive, tiering can be both reactive and proactive based on policies
  - Caching is to Accelerate Data, Tiering is to Manage Data
  - Caching is workload dependent, Tiering is time dependent
  - Caching and Tiering complement each other, not compete

Storage Tiering In NAS

- Can be block or file based, automatic or manual
- Beneficial if system intelligently places data for you
- Even better if it works across discrete systems - automatically!
- Far better if tiered storage is file-system aware
- Backup Performance in block-based tiered Storage can be a challenge
- Sizing the disk tiers for balance of performance and cost is the key - be conservative initially

Storage Tiering In NAS

- Tier0 in NAS (Tier0 = SSD for this discussion)
  - Best if Tier0 is space is managed automatically
  - Estimating Tier0 capacity is the key
    - Estimate against lower tiers
    - Between 1 to 10% being standard
    - Significant cost differences between 3% and 10% Tier0
  - Sizing active data appropriately is important
  - Getting numbers from backup related stats (or replication deltas) could be useful

Storage Tiering In NAS

- Reporting & Monitoring of Automatic Tiering
  - How effective is Tier0 (since it's most expensive)?
    - Consider Hits and Capacity Utilization
  - Which workloads are being benefitted?
  - Can you establish a correlation between blocks and files?
  - Are there levers to fine-tune the process and policies?

Scale-Up NAS

- Scale-up or Vertically Scalable (typically only disks)
  - For high performance – cannot add large amount of disks
    - Overruns the Storage/NAS Processor
  - Storage Processor and/or NAS Heads do not scale
    - Could be upgraded though
  - Creates a big sprawl of independent NAS systems
    - Results in disk space inefficiencies
- Good all-in-one/multi-purpose solutions
- Multi-protocol support is good
- Not focused on one type of workload
  - Could be designed for specific workloads

Scale-Up

- Feature rich - snapshots, dedup, replication, VMware integration, etc
- Multiple storage buckets, not all in one solution
- Flexibility with tiered storage
- Block and file can be unified
- Can be open-source, commercial or a combination
- Tiering across individual NAS systems is a challenge

Scale-Out NAS

- Scale-out or Horizontally scalable (typically capacity and performance) - usually are clustered NAS
  - Multiple redundant nodes performing file processing and/or disk functions
- Scale-out NAS origins
  - Typical Use Cases in - Media, Entertainment, Oil & Gas exploration, Imaging (high bandwidth)
  - Usually associated with High-Performance Computing
- Usually distributed filesystem
- Performance independent redundancy
  - Data striped across multiple nodes

Scale-Out NAS

- Data and metadata could be separated
- One large storage bucket
- Could use commodity hardware and storage
- Can be open-source, commercial, or a combination
- Could provide higher aggregate performance
  - Single-client performance usually not impacted much
  - High-performance, mainly benefiting large files
- Massive scalability for disk capacity
- Usually serves a single purpose

Scale-Out NAS

- May increase networking requirements (due to more nodes)
- May increase power/cooling needs depending on the configuration
- Could get better disk utilization as the system grows

Scale-Up Vs Scale-Out NAS

- When investigating scale-out Vs scale-up, check your requirements and concerns –
  - What type of performance are you looking for (small file, large file workloads)
  - Snapshot capabilities
  - Uptime requirements (SLAs), and vendor support
  - Are you combining file and block on the same system
  - Replication, Mirroring
  - Protocol support (NFS, CIFS, FTP, HTTP, etc)
  - Backing up the data to tape

Scale-Up Vs Scale-Out NAS

- Existing Infrastructure and protocols in place for LAN/SAN, and Application layer) – the ones that are in use and IT teams are familiar with - or not in place, i.e., additional costs to support it
- Multi-tenancy support
- De-duplication, compression, Virtualization environment integration

De-duplication on NAS

Education SNIA

- De-duplication & Compression
  - Know the differences
    - Compression - removal of redundant bits
    - File Dedup - removal of redundant files
    - Block/Data Dedup - removal of redundant blocks
- Check what is offered
  - Dedup only (file or block) - good
  - Dedup (file or block) + compression - better
  - Dedup (File or block) + compression - best

De-duplication on NAS

- De-duplication not for everyone
  - May put performance penalty on CPU and memory
  - May impact user response times on active filesystems
  - Storage space gains can be realized for less active filesystems, but could be at the cost of additional data management

De-duplication on NAS

- Challenging to track on/off/status/gains for 100s of active filesystems
- Under what scenarios (backups, mirroring, snapshots) and workloads hydration is triggered
- Space gains from file Vs block dedup could vary – additional 2% gain in a large environment is significant

In Summary…

- Many Enterprise NAS flavors available today
- Know your requirements to find the right solution
- Plan for long term - scalability, stability, data management
- Scale-up and Scale-out both have their use cases

Q&A / Feedback

Many thanks to…

Many thanks to the following individuals for their contributions to this tutorial.

- SNIA Education Committee

Dr. Joseph White

Simon Gordon

Send any questions or comments on this presentation to SNIA: tracktutorials@snia.org