



# **Storage Management Best Practices & Tips**

**Considerations for file & block storage provisioning**  
Anjan Dave, LSI Corporation

- ◆ 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.**

## ➤ File and Block Storage Deployment in Enterprise Environments

The non-stop growth in the amount of information that is generated and stored necessitates a review of storage networking best practices in medium to large sized storage environments. This tutorial provides some practical tips for planning, purchasing, and deployment of storage in areas of Network Attached Storage systems for unstructured data, deploying block storage for business critical applications, and managing Storage Area Networks in general.

The objective is present on how to plan the purchase and deployment of NAS storage in a large enterprise environment from storage capacity, performance and backups perspective. What methodologies to adopt for deploying block-storage for business critical applications for structured data, and what are some of the influencing factors in optimally managing the networked storage infrastructure for unstructured and structured data.

- Anytime you plan to purchase NAS storage or related solutions, plan for –
  - ◆ Storage Capacity Planning
    - › This is about the amount of storage you are dealing with
  - ◆ Storage Performance Planning
    - › This is about the performance of the amount of storage you are dealing with
  - ◆ File Storage Management
    - › Basic things to keep in mind when managing a large NAS environment
  - ◆ Also keep in mind data center power, cooling, floor space, LAN/WAN bandwidth, available admin-power, training, and future costs of supporting the hardware, software and licenses (not covered in this presentation)

## ➤ Storage Capacity Planning

- ◆ *How well do you know your data?*
  - › “Classify” your data (age, size, type, etc) so you know where to spend the money (figure 1)
  - › What % of your total data is “active?” 20% is not unusual...
  - › How will you or the solution deal with the remaining 80%
- ◆ Today’s active data will become inactive tomorrow
  - › Policy based data migration tools are there, but does a tool “fit” in your environment?
  - › Depending on your workflows and processes, how much of the inactive data can be “parked” for long term on secondary storage?
  - › You may NOT want to track, chart, report, backup, increase performance, provide high redundancy, occupy ports, maintain support (e.g., 4-hr support), buy a solution, or spin expensive disks for data that is hardly accessed

# Addressing Active Data

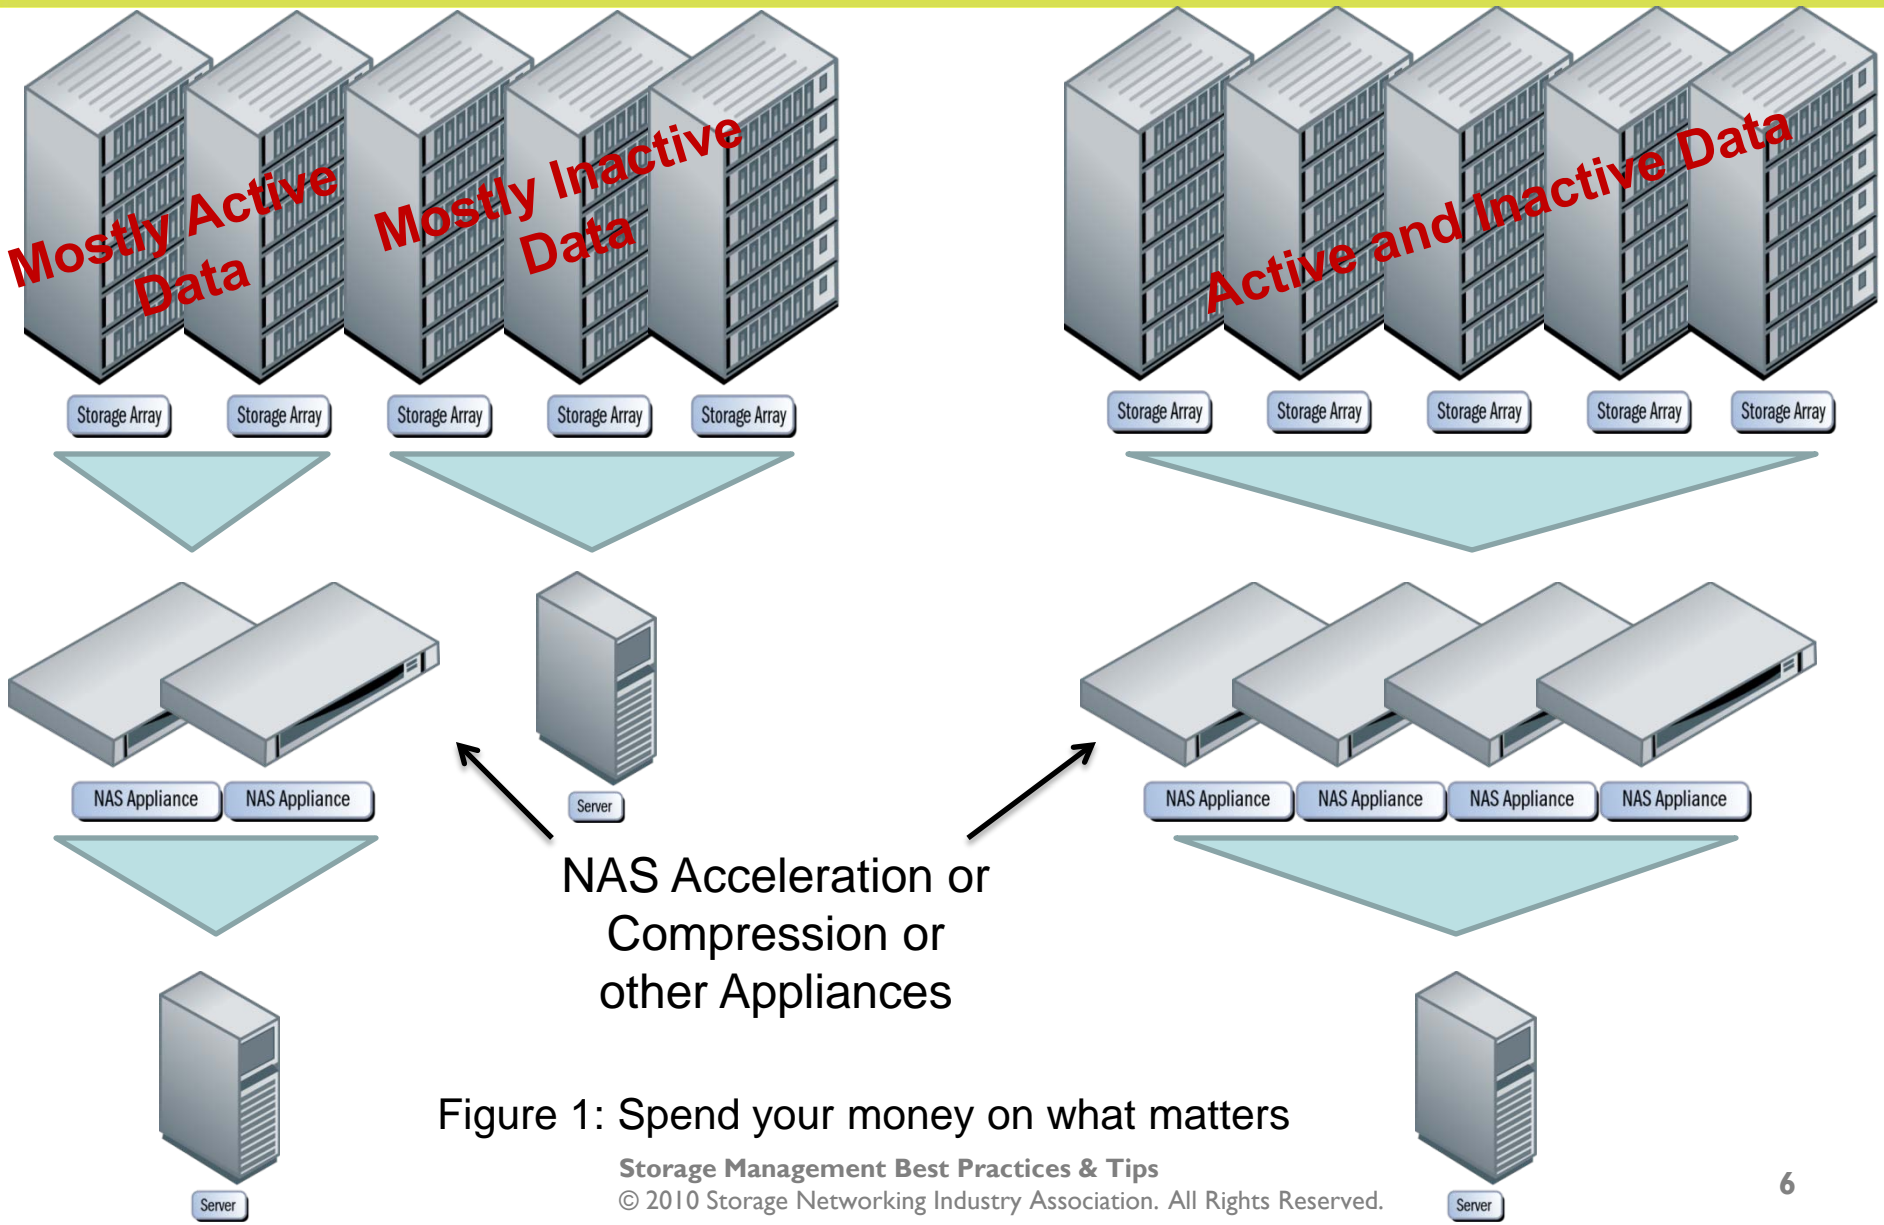


Figure 1: Spend your money on what matters

## ➤ Storage Capacity Planning

- ◆ Assuming you **HAVE** to buy more disks...
  - › One Storage Array could hold PetaBytes of storage. What happens when you need to do maintenance on it? Spreading the disks on more controllers may make more sense
  - › Know how you count your disk spaces – “Available” or “Usable” Disk Space can mean several things (see graphic)
  - › Consider disk overheads – formatted capacity, spares, parity disks, snapshots, etc. Don’t forget disks for replication
- ◆ You are not just buying disks!
  - › Do you have the additional ports on your fiber switch/director?
  - › Do you have additional ports on the LAN side?
  - › How about free ports on patch panels?
  - › Are electrical circuits available to power the new equipment?

# File Storage Purchase Planning

## What to count, where?


- Cold Capacity
- Formatted Disk Capacity



- Hot Spares
- RAID-Parity Space



- File System Overhead
- Snapshot/other reserves



- Quoted space
- Non-quoted space
- Archive space
- Compressed data
- Thin provisioning



## Disk Space Report Example

Purchased (TB)	Spinning Usable (TB)	Spinning Used (TB)	End-user Usable (GB)	End-user Used (GB)



## ➤ Storage Capacity Planning

- ◆ When was your last disk/node purchase?
  - › Will you be pairing the new (faster) controller with the old one?
  - › Can you accommodate a “new generation” node in an old cluster?
  - › How will you address the increasing disk drive densities? E.g., new 600GB drives Vs your existing 300GB drives
  - › Does the solution require you to grow disk capacity in specific capacities only? Check with your vendor
- ◆ Drive mixing
  - › Does the solution allow mixing SATA and FC drives in same cabinet, same disk module?
  - › How about SSDs?

## ➤ Storage Performance Planning

### ◆ Back to *know your data & end-users*

- › Do you know the life cycle of your data? Origin, growth, decay, expiry...
- › One file or a set of data will be accessed simultaneously? By how many users or applications? How many people will login simultaneously?
- › What is the average file size and how many? Dealing with lots of tiny files is far different than large files
- › Who is the customer? Online consumer, report generator, partner, developer, or a community? What is their workflow?

## ➤ Storage Performance Planning

### ◆ Know *your applications*

- › Understand Read-Only Vs Read-Write applications and what are the influencing factors (what, when and how it becomes RO).
- › Is any content static, such as graphics? Can it be compressed?
- › One file or a set of data will be accessed simultaneously? By how many users or applications?
- › Sequential data versus non-sequential data can be planned for, but can you aggregate data with similar characteristics?

### ◆ Other performance factors to consider

- › How many host ports, disk ports, Ethernet ports are available, what speeds, and what are the options for port aggregation?
- › If considering new storage networking technology, review your existing investment in the storage networking infrastructure

- Storage Performance Planning
  - ◆ Other performance factors to consider
    - › Incremental growth – mind the performance when adding disk space (concatenation Vs re-striping on the added space)
    - › Know what protocol(s) are strongly used in your env., and whether supporting that protocol is a strong trait of the product. CIFS is supported widely, but every vendor's implementation could be different

## ➤ Storage Performance Planning

### ◆ Other performance factors to consider

- › Understand how the file system works – a powerful hardware is not of much use with weak file system
- › Know how much flexibility you have in fine-tuning the system for performance. Ask what options are available, more the better. Out of box settings work only to some extent
- › To deliver highest performance, the vendor may want you to configure, operate and upgrade their product or solution in a specific manner. Ask for details and whether it is feasible for you
- › Does the system allow you to configure the back-end disk storage?
- › Make sure you understand the meaning of “front-end grows independently of back-end”

## ➤ File Storage Management

### ◆ Dealing with Data

- › Data migration – you may need to transfer lot of data on the new platform – what are your options, especially if downtime is hard to buy from users?
- › NFS and CIFS – coexistence – how well is this implemented?
- › Don't just know what protocols are to be supported, rather how widely and heavily each is used (e.g., 90% NFS, 10% CIFS)
- › Single-pane of management – if you consider deploying at multiple locations, you'll need it
- › Will the solution cut down the usage of in-house scripts/apps?
- › How does the system handle file system corruptions?
  - Does the product fix the file systems in mounted state?

## ➤ File Storage Management

### ◆ Reporting

- Reporting – can the system report the metrics across your full deployment landscape? Does reporting need it's own operating infrastructure?
  
- Is any of the following available?
  - Can the system report performance related numbers such as wait times, IOPs, throughput at the system and port level and device level?
  
  - Can the system report system and user level statistics for NFS/CIFS and other protocols?
  
  - Can the system report on it's CPU, Memory, Cache utilization?

## ➤ File Storage Management

### ◆ Reporting

- Can the system tell if it's overloaded in some way? Or it leaves you wondering?
- Can the system report 'hotspots' on disks?
- How intuitively does the system report on Snapshots, NDMP sessions

### > Logging – how well the product documents it's actions/issues?

- Does the system have internal hooks to capture vital data for troubleshooting?



## ➤ File Storage Management

- ◆ Understand the ‘Specs’ of the platform
  - › There’s a difference in published specs and recommended specs.
  - › What is the largest size of the volume/file system you can create?
  - › Practical max number of LUNs the system can address
  - › Max and recommended number of LUNs that make up the volume/file system
  - › Practical max number of volumes/file systems per “node”

## ➤ File Storage Management

- ◆ Understand the ‘Specs’ of the platform
  - › Practical max number of “nodes” in the cluster
  - › Practical max number of Tape Drives the system can address
  - › Maximum NFS and/or CIFS Ops/sec
  - › How much memory does the appliance/system have at various levels
  - › Max number of NFS/CIFS clients the system can handle
  - › Max number of NDMP sessions the system can handle per node

## ➤ Storage Management and SNIA's SMI-S

*SMI-S is a standard for storage management defining the communication between management applications (such as a Storage Resource Management application) and instrumentation (software/firmware components)*

- SMI-S can help you achieve a single pane of glass view into your storage landscape, managing, monitoring, and reporting via a single interface
- Verify following beforehand-
  - Are the existing arrays, new arrays and the SRM tools SMI-S compliant?
  - Understand what features do the arrays from different vendors support via SMI-S
  - Are you able to collect meaningful information (e.g., performance statistics) in a uniform way from all arrays? All arrays may not report the stats you need
  - Understand any limits imposed by a particular storage/SAN vendor (e.g., number of arrays, volumes, etc) that can affect your deployment
  - You may need to upgrade or install a specific version of firmware/software on arrays or other devices – this may affect your compatibility
  - In general, how much can you manage and/or monitor and report, end-to-end from host to SAN to Storage via SMI-S?

## ➤ File Storage Management

- ◆ What data will you migrate to FC/SAS, SATA, SSD areas?
- ◆ Distribute the data of I/O intensive apps across spindles to avoid hotspots...Know how the file system lays the data
- ◆ If you are allowed to build the back-end disk, keep in mind
  - the amount of data, the read Vs write, random Vs sequential, active Vs inactive data, snapshot reserve space, optimum use of disk space, disk types in the array, etc

## ➤ File Storage Management

- ◆ What other factors can you consider? – IOPs, controller cache settings, RAID level, meta-volumes, stripe size, number of spindles, controller ownership of lun, number of hot spares, etc
- ◆ What are your options to build the file system? Understand how the back-end design factors such as number of luns, lun size, stripe size, etc affect the performance
- ◆ Consider NFS mount options (udp/tcp), journaling and a-time updates, their impact to end-users

## ➤ File Storage Management

- ◆ The largest file system you can create (given backup windows, number/size of files, snapshots, file system checks, performance, etc)
- ◆ Test the network (client connectivity) configuration options. Port aggregation works differently depending on your LAN switches. Can you use Jumbo Frames?
- ◆ Consider number of snapshots, their location. Are snapshots based on volumes? Can snapshot creation and deletion be staggered?

- Know **all** the requirements **beforehand** –
  - ✓ Details of all disk spaces needed
  - ✓ Hardware details of all storage, SAN and host resources available
  - ✓ GROWTH in databases for next 1 to 2 yrs
  - ✓ List of all applications, databases, corresponding host names
  - ✓ Availability requirements, SLAs, DR (replication, clustering)
  - ✓ Performance requirements for each DB, and application

- Know **all** the requirements **beforehand** –
  - ✓ Processes needing storage team involvement such as DB upgrades, refreshes, patching, etc needing snapshots, or backup/restores
  - ✓ Don't forget impact of Disk/Tape backup, de-duplication especially for 5. and 6
  - ✓ Tiered storage, SSD, Virtualization at Host and Storage level, Caching, Compression – usage of these technologies and their expected benefits/impacts should be well understood

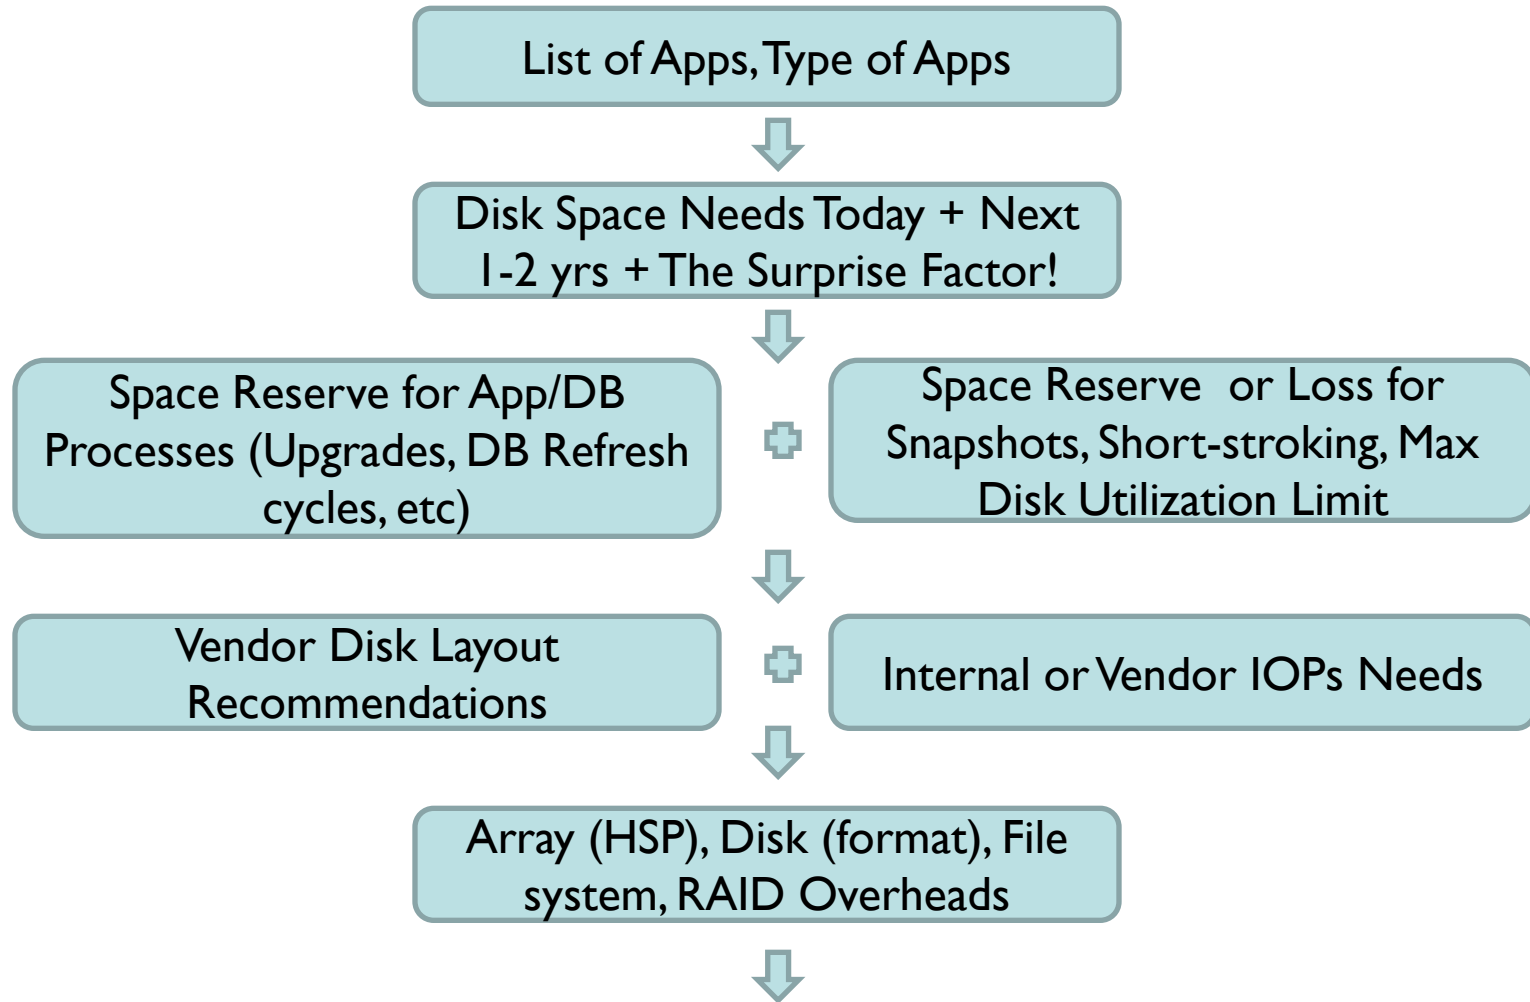


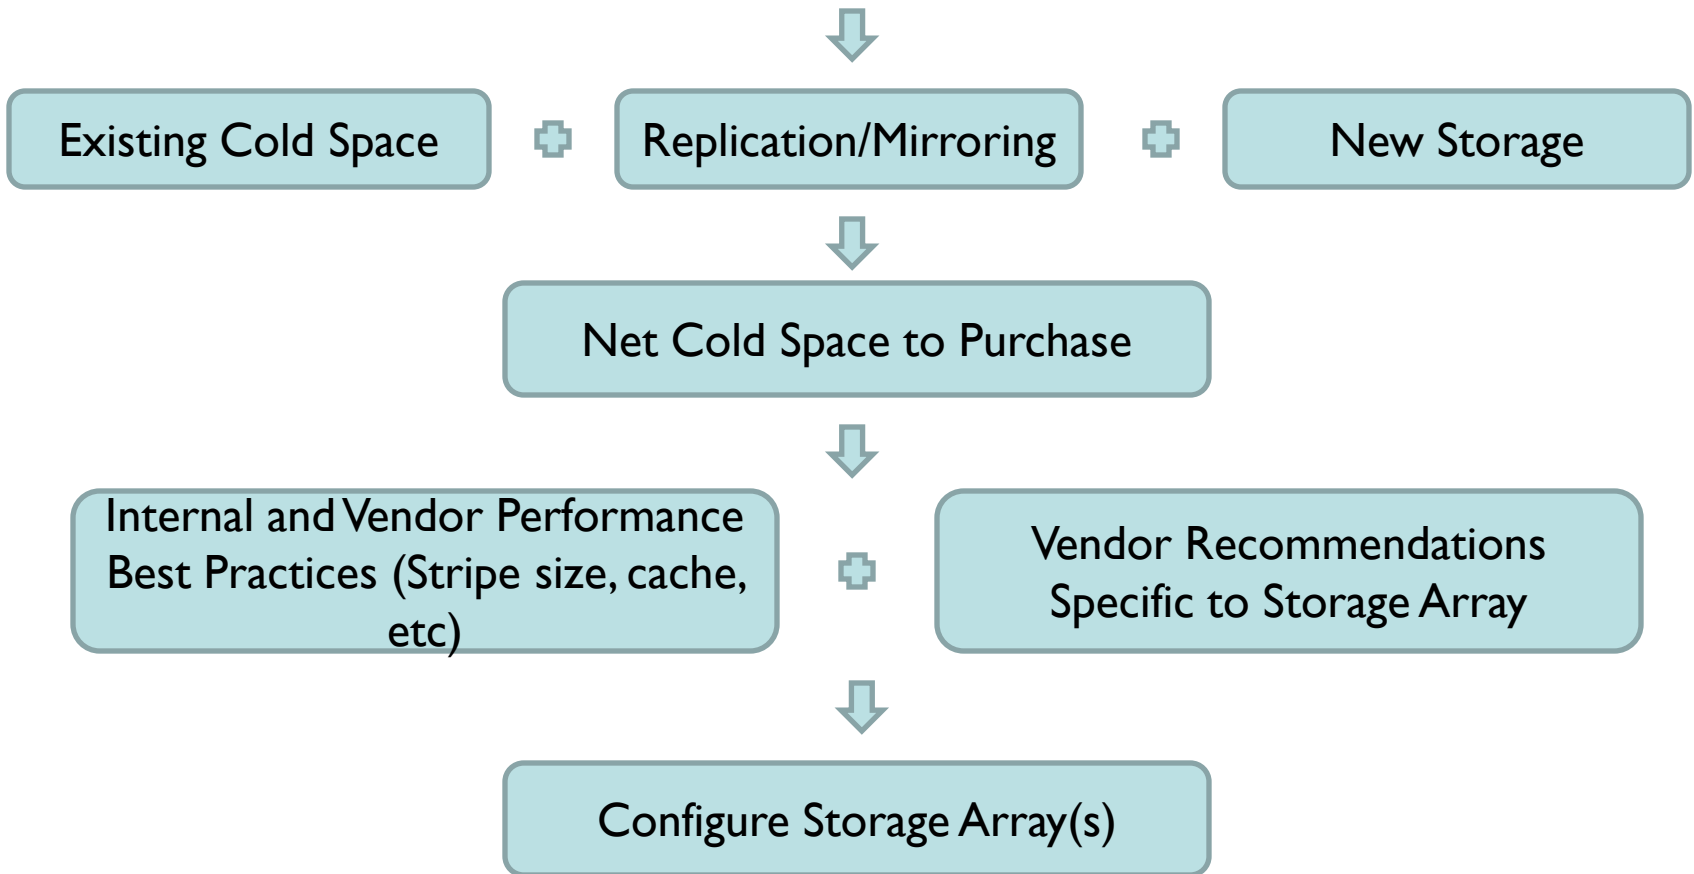
## ➤ Know **all** the requirements **beforehand** –

- ✓ For optimum disk layout for databases, check the vendor best practices in addition to internal standards
- ✓ Pick an application first. You can't club the best practices for a Report Generation application and an OLTP together
- ✓ Review DB tunable parameters first – they are usually tuned with Application in mind, not just storage
- ✓ Then match DB parameters with storage array settings (e.g., IO size pre-fetch unit could be a multiple of segment size)

- ✓ There may be specific parameters for hosting a particular RDBMS on a storage platform from certain vendor – get all parties in agreement, same for applications
- ✓ Understand the **COMPLETE** landscape of your block storage needs, not just the need at a time
- ✓ Pool apps with similar requirements on one array or a set of spindles
- ✓ Boot from SAN
  - › Don't put OS and Swap luns on same disks
  - › Segregate OS/Swap luns for different hosts on separate disks
  - › Test path failover thoroughly
  - › Put the OS related luns in beginning of the RAID-set

## Process for deploying new block storage for applications/databases





## Note

Virtualization, SSDs, Automated Tiered Storage can change the above process

- ◆ Inside the Array...
  - › Segregate like-minded applications on separate disk arrays/pools
  - › Set array and LUN level cache settings. Check with Vendor first
  - › Review array maintenance related background activities settings
  - › RAID Level
    - Review DB, Storage vendor best practices, and consult your DBAs
    - General Rule of Thumb - Choose RAID-10 over RAID-5 for write-heavy usage
    - Consider usage of RAID-5 Vs RAID-10 for table spaces, transaction logs, archive logs, Indexes, Temp space, Sort space, etc
    - For # of spindles, understand the IOPs requirements for reads/writes as well as the size of the IO
  - › Keep in mind the IOPs needs for backups + transactions
  - › Pick a suitable segment size based on each app/DB needs
  - › Can you stripe on top of a group of LUNs?

- ◆ Inside the Array...
  - › Standardize on 2 or 3 sizes of LUNs
    - Smaller LUNs for binaries, OS, Swap, etc
    - Larger luns for Database files
  - › Maintain a balance of ownership of luns among the controllers
  - › Allocate LUNs from different RAID-sets (i.e., spread the IO)
  - › If you're not short-stroking, then keep utilization below 80%
  - › For Redirect-on-write operation, make SURE the disks you allocate for the deltas are configured same as the original LUNs
  - › For Copy-on-write operation, make sure the disks used for Deltas are separate spindles, and not shared with original disks
  - › Standardize on Host/Host Group naming style (match it with zones)



**Check out SNIA Tutorial:  
Storage Performance 101**

- ◆ At the Host level...
  - › Set the optimum HBA driver settings
    - SAN Topology
    - Queue Depth
    - FiberChannel Speed
  - › Test path failover – break the path in every possible combination
    - Controller failure, just pull it out!
    - Manual LUN failover to alternate controller
    - Block port on the switch/director
    - Unplug fiber cable
  - › Know if LUN names are visible by the OS, it helps
  - › Make sure LUNs are ingested correctly – i.e., a LOG lun is not configured as a database lun
  - › Know the limits – Max # of luns, paths, LUN Zero requirement

- ◆ If Storage Virtualization is involved...
  - › Understand the new storage terminology against existing
    - LUN/Volume/Slice/Partition/Stripe/Pool/Reserved\_pool/disk\_group, etc, etc...
  - › Compatibility check will be more complex, but do it
  - › For new install, pick one OS platform and test all features/functions
  - › Test how you'll virtualize existing data/storage
  - › Test volume expansion/shrinking
  - › Document physical to logical (NPIV) mappings of the WWNs – do this end-to-end
  - › Test path failover – break the path in every possible combination
  - › Design several storage 'pools' with different characteristics
    - RAID-5, RAID-10, etc
    - Based on number of underlying disks
    - Based on disk type/capacity
    - Based on disk groups having hot spares
    - Based on array specs/type



**Check out SNIA Tutorial:  
Five Best Practices In  
Virtualization**



- ◆ If you plan to have SSDs in the environment
  - › Make sure (evaluate) that the specs meet your workload/workflow
    - Refer to the [SNIA SSS Performance Test Suite Specification](#)
  - › Know the costs for YOUR platform of choice, compared to disks
    - Refer to the [SNIA Enterprise TCO Calculator](#)
  - › Know your data! What to place on the SSDs
  - › What policies, processes and tools will you employ to –
    - Put data on SSDs (active and/or performance-critical)
    - Move data out of SSDs (to Tier1 or SATA)
    - What suits you? Move data at LUN level, file level or block level?
  - › Size the controller/CPU for # of SSD and Disk drives in the array
  - › Make sure you factor in the “write amplification” in your tests
  - › Know the “Erase Block Size” and align the partitions accordingly
  - › Make sure you can measure/quantify the expected performance gains

- Please send any questions or comments on this presentation to SNIA: [trackstoragemgmt@snia.org](mailto:trackstoragemgmt@snia.org)

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

**- SNIA Education Committee**

**Brandy Hammack  
Joseph White  
Anjan Dave**