

Smart Hybrid Storage Using Intelligent Data Access Classification

Dr. M. K. Jibbe, Technical Director
NetApp, Wichita, Ks USA

Mahmoud.jibbe@netapp.com

316 636 8810

Senthil Kannan, Test Architect
NetApp, Bangalore, India

Objective Relevant to Storage Industry

- ❑ Understand various data access demands
- ❑ Realize the challenges involved to satisfy the increase of performance requirements
- ❑ Find out possible methods that can help overcome the growing performance requirements

In SAN, not all drive groups are equally accessed. The data access on the drive group depends on application need as well as user requests. Broadly we can classify the need for data access in to three major categories.

- ❑ **Mission-critical data or High Performance or Sensitive data**

Mission-critical data is used in the key business processes or customer applications. Typically has a very fast response time requirement. These are transactional data which require high I/O operation (IOP)/High performance with optimal /moderate reliability.

- ❑ **Reliable data**

Reliable data may be classified as company secret and doesn't require instantaneous recovery for the business to remain in operation. The redundancy of data is a must and should be always available.

- ❑ **Reliable & Sensitive Data**

This is the requirement where the data requires a high IOP and highly reliable storage technology.

Current Challenges in data access

Key challenges we face in SAN today

- ❑ Conventional method to achieve best I/O performance in SAN is to use RAID 0. This is suitable for high performance data access category. The best reliable RAID is RAID 6.
 - ❑ Both these methods have their own constraints.
 - ❑ For example, RAID 0 with hard drives are not as fast as RAID 0 with Solid State Drives (SSD). But using SSD for entire drive groups is not a cost effective solution and RAID 0 does not provide reliability.
 - ❑ For example, RAID 6 option has reliability but it does not have good performance. Again using SSD for the entire drive group is not a cost effective solution.
- ❑ I/O demand is not constant and varies based on
 - ❑ Different user access load and
 - ❑ Application need.
- ❑ Predefined Drive group/ Volume Group parameters will not suite the dynamic demands of user/host I/O.

In both these cases, having a hybrid storage with built-in intelligence that can respond to dynamics in SAN will be a suitable answer

A Solution to data access dynamics

Adding a **hybrid storage** to a **suitable RAID** algorithm that also has the **intelligence** to learn and adopt to the I/O demand → Reliable and high performance RAID solution

- ❑ Hybrid storage: Co-exist Hard Drives and Solid State Drives in a disk drive group/ volume group. This enhances the performance of the storage subsystem dynamically based on the I/O (Input/output) loads.
- ❑ Identify the suitable RAID: The most suitable combination that can provide Reliability/Performance based on the data access category are
 - ❑ Reliable Storage – RAID 51 or RAID 61
 - ❑ Sensitive/high performance Data Storage – RAID 50 or RAID 60
- ❑ Introduce dynamics with Hybrid storage → Dynamic Drive Provisioning
 - ❑ Developing intelligence by learning and classifying data access pattern with above identified enhancements.
 - ❑ This is done by the Intelligent data pattern learn logic engine with smart data access classification in Storage array Controller Firmware.
 - ❑ This layer is introduced in to the I/O path. This helps in analyzing and classifying the data access need.
 - ❑ SSD Drives can then be attached /detached to the Volumes groups or Disks based on the frequently accessed I/O pattern that helps in boosting performance by acting as drive group level cache.

Conventional RAID approach




Figure 3: The implementation for Reliable storage. RAID 5 shown in all the figures is for illustration purpose. This can be any Redundant RAID levels


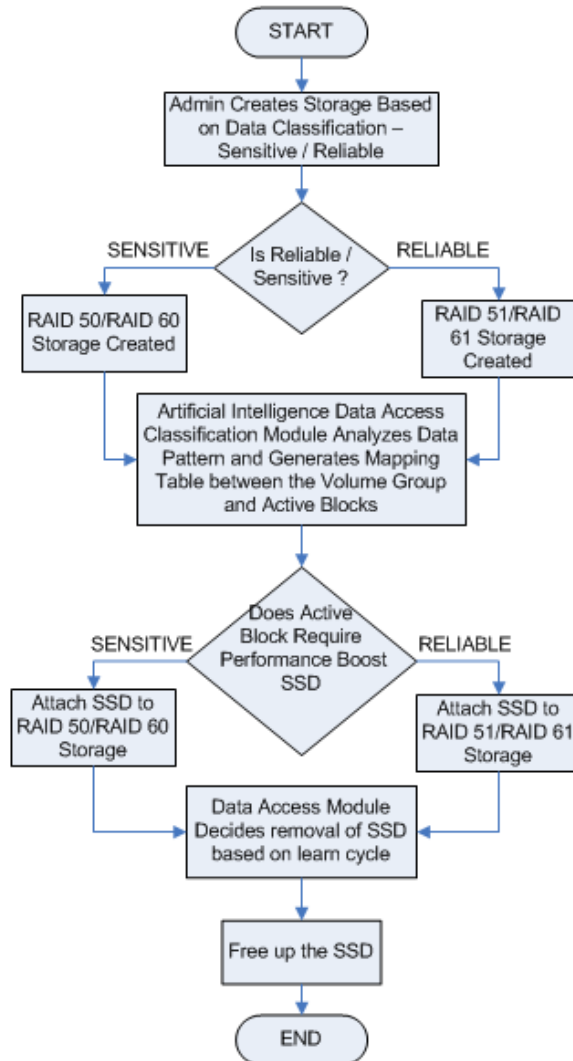


Figure 4 :The Implementation for Sensitive /High performance oriented data storage. The RAID 5 shown is for illustration purpose. This can be any Redundant RAID level

Hybrid Storage with intelligence



This flow chart just illustrates how user can select the data access requirements. And how the SSDs are allocated for existing drive groups based on the learn cycles from artificial intelligent Data access engine (aka Artificial Intelligence Data Access Classification Module)

Visual representation I/O path

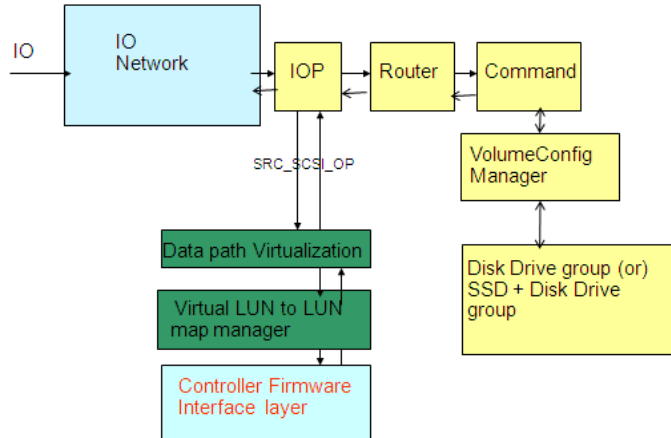


Figure 1: This figure shows I/O transaction through the I/O path virtualization layer, this layer updates the SCSI IOP with vendor unique bits on CONTROL BYTES based on data access need criteria

7	6	5	4	3	2	1	0
Vendor Unique	Reserved			Flag	Link		

Figure 2: This figure shows the vendor unique bits in the CONTROL BYTE of a SCSI CDB

Need for controlled IOPs ?

- ❑ Controlled IOPS (I/O Operations) is to add flexibility for controlling the data access learning logic.
 - ❑ For example, Controlled IOP can be used for monitoring,
 - ❑ The reliable drive group or
 - ❑ The sensitive drive group as needed by users.
- ❑ Helps the Intelligent data pattern learning logic to analyze the I/O pattern based on the control byte classification and start monitoring the demand for faster I/O access.
- ❑ Leverage the Hybrid Storage technology:
 - ❑ Addition of SSD drive Pools which are kept as global reserved cache drive group
Determine the performance requirement of drive group
 - ❑ Decide the improvement needed for a Drive group/ volume group as
 - ❑ A whole or just
 - ❑ Any particular drives in a Drive group/Volume group.
 - ❑ Allocate SSD drive or a set of SSD drive to be mapped to the existing drive groups.

- ❑ Having the capability to learn and adopt to incoming I/O patterns.
- ❑ Introducing Virtualized data path from the Host/Initiator to the drive groups or volume groups.

An example to identify and learn data access need :

- ❑ Have a data path virtualization layer to receive SCSI IOP from initiators and update the IOPs with vendor unique bit information on the CONTROL BYTE of SCSI CDB before passing to I/O router module in the firmware which will route the I/O to actual Drive or Volume groups.
- ❑ The vendor unique bits can be used for identifying the data access need as taken in the below example;
 - ❑ The vendor unique bits are set as “zeros” for I/Os to “Data Reliable” only drive groups or volume groups
 - ❑ The vendor unique bits are set as “Ones” for I/Os to “Data Sensitive” only drive groups or volume groups
 - ❑ The vendor unique bits are set as “zeros and Ones” for I/Os to “Data Reliable + Data sensitive” drive groups or volume groups -> this need is dynamically set based on the pattern learn logic and the need for I/O bandwidth or load.

Dynamics of Hybrid Storage with intelligence

- ❑ SSD are allocated dynamically based on the decision taken by Hybrid storage Intelligence.
- ❑ SSD are allocated to particular physical drive or single drive group or multiple drive groups during I/O access and further boosts the performance.



Figure 5 : Shows I/O transaction using SSD for single Disk drive groups performance boost as on need

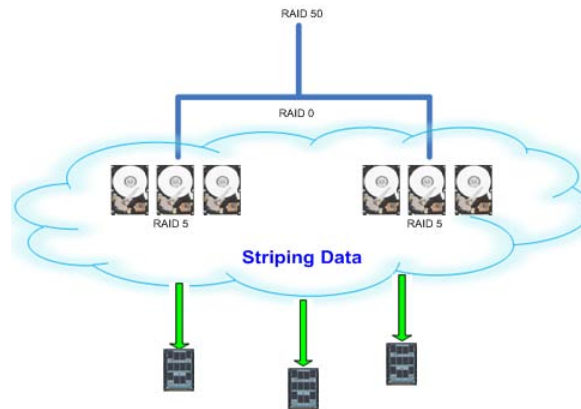


Figure 6 : Shows I/O transaction using SSD for the whole or multiple Disk drive groups for performance boost)

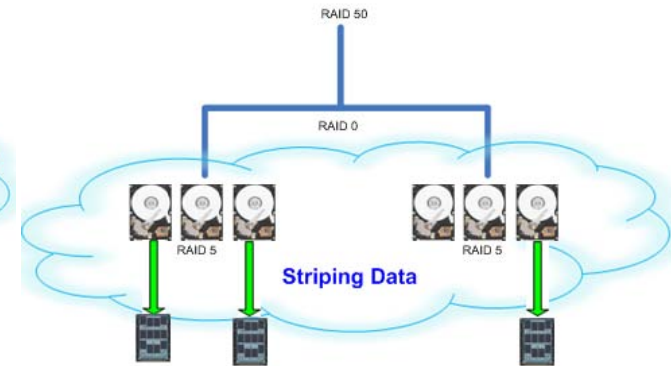


Figure 7: Shows I/O transaction using SSD for individual disk drives in the Disk drive groups for Performance boost

Benefits of intelligent hybrid storage

- ❑ Enables the users to achieve efficient storage allocation on the fly and eliminates the overhead of storage administrators to monitor the data access and change the storage allocation frequently.
- ❑ Helps in Easy and Efficient Storage Planning based on data access needs
- ❑ Better reliability.
- ❑ Cost Vs Performance advantage.
- ❑ Dynamic Performance boost.
 - ❑ Usage of Hybrid drives with NAND flash integrated for disk caching can boost the performance further

References

- ❑ This paper is a result of the Patent Idea that has been filed on Oct 07 2010 – Reference L10-0513 NetApp/LSI Invention Disclosure ID.
- ❑ “Trends in Solid State Storage” – Jeff Kimmel (NetApp) presented @ SNIA – Santa Clara 2010.

Thank you

