White Paper

The NetApp Virtual Storage Tier

Paul Feresten, Larry Freeman, Mark Woods, NetApp
April 2011 | WP-7127

EXECUTIVE SUMMARY

Automated storage tiering has emerged as a means to help IT departments simplify data management and reduce storage costs while at the same time improving application performance and better aligning storage scaling capabilities with those of shared IT infrastructure. The NetApp® Virtual Storage Tier is an optimized approach that offers real-time promotion of hot data, the smallest level of granularity, simple deployment and management, and full integration with NetApp storage efficiency capabilities.

With the NetApp Virtual Storage Tier you can increase I/O performance by up to 75%, reduce the number of disk spindles required to accommodate I/O intensive workloads by 75%, and shift important workloads from FC or SAS disks to more economical, higher-capacity SATA disks.

This white paper examines automated storage tiering and highlights the advantages of the NetApp Virtual Storage Tier.

TABLE OF CONTENTS

1 INTRODUCTION

Matching the performance needs and lifecycle requirements of your data to the different types of available storage media has always been a bit of a headache. The goal is to make sure that data is on the right media, at the right time, for the right cost. The reality, however, is that it can be difficult to know what type of media will best address a particular need, and it can be hard to adjust your storage allocation and data placement as those needs change. Administrators and end users tend to hedge their bets by choosing faster, more capable storage than they really need. As a consequence the faster—and more expensive—storage tends to be oversubscribed, and that increases storage costs and potentially robs performance from applications that need it.

Flash-memory-based, solid-state drives (SSDs) and other forms of media have brought this problem to the fore. Flash-based media can complete 25 to 100 times more random read operations per second than even the fastest rotating media, but that performance comes at a premium of 15 to 20 times higher cost per gigabyte. This has created a strong need to reserve the use of Flash-based media for "hot" data—active data in high demand—as a way to maximize the benefits from those investments.

This is exactly where automated storage tiering fits in. Automated storage tiering is intended to identify and promote hot data to higher-performance storage media automatically, while leaving cold data in lower cost media. This approach promises to give you the maximum benefit from expensive, high-performance resources, while using less total resources to achieve a given level of performance.

Automated storage tiering has the potential to replace tedious manual data management and to significantly benefit advanced virtualization and cloud environments. This paper examines automated storage tiering capabilities in more detail and describes the design and performance characteristics of the NetApp Virtual Storage Tier, which can increase I/O operations by up to 75%[1], reduce the number of disk spindles required to accommodate I/O intensive workloads by 75%[2], and shift important workloads from FC or SAS disks to more economical, higher-capacity SATA disks.

2 STORAGE TIERING AND THE MOVE TOWARD AUTOMATION

It's common for a modern data center to use several types of storage media to address diverse needs of data usage and value. Assigning data to specific media types based on business value and application performance allows you to prioritize your storage costs on either a cost-per-IOPS or cost-per-capacity basis. In addition to the business value of data, media tiering decisions should be based on the actual access patterns that are driven by application workloads. The data that is most in demand can be thought of as "hot" while infrequently accessed data can be thought of as "cold". The challenge, of course, is how the alignment of media type and business value actually happens and how this alignment is managed over the lifetime of the data.

LIMITATIONS OF MANUAL TIERING

No matter what method you use to assign data to media type, as requirements change the need inevitably arises to move data between tiers. Manual data movement, however, can be time, resource and labor intensive. The first challenge is simply to identify the data that should be migrated (and when to move it), which can be a tedious manual process. Once that is accomplished there's a time-consuming data copy process that temporarily requires double the storage capacity. If the data is subject to compliance requirements the task may be even more complicated.

---

[1] NetApp *WP-7082: Using NetApp Flash Cache in Online Transaction Processing*
[2] NetApp *TR-3799: NetApp Performance Acceleration Module II in File Service Workloads*

The sheer quantities of data involved can make wholesale movement almost impossible for large data sets. Even when necessary storage and bandwidth are available, finding a large enough time window (hours or even days) when data can be taken offline can be problematic. Given the difficulties of manually managing multiple tiers of storage—compounded by the need to optimize the use of expensive Flash-based resources—the need for automation becomes clear.

The need to automate the tiering process is greatest when hot and cold data coexist on the same storage system and when the demand for specific subsets of that data is in a state of flux. Although it is important to promote hot data to the highest performance tier available, it is also important to do so in a granular manner so that cold data doesn't get moved with it. The move to an automated process should mean that performance goals are met in real-time and with maximum efficiency.

WHAT IS AUTOMATED STORAGE TIERING?

Traditionally, different types of storage media were isolated on different storage controllers. A key enabling technology was the ability to mix different types of media (SSD, SATA disk, FC disk, SAS disk) in the same storage system, a capability that was pioneered by NetApp. Automated storage tiering takes this a step further by enabling the use of these different media types for the same data set.

Automated storage tiering is designed to recognize hot data and "promote" that data to higher performance storage media as needed to maximize performance while at the same time keeping cold data out of the performance tier to minimize cost. The term *promote* is used here to encompass a variety of processes that can include wholesale migration of large chunks of data from one tier to another.

GOALS OF AUTOMATED STORAGE TIERING

Adopters of automated storage tiering hope to address a number of important requirements:

- **Minimize or eliminate the need for manual data management** by automating the process of data placement. The goal is to transform the painful process of performance management into a background activity that occurs automatically.
- **Reduce storage and data center costs** by optimizing the use of expensive storage resources and saving administrator time. For Flash-based media there's a strong desire to minimize the total amount needed. In addition, for many applications the number of spindles provisioned has more to do with achieving the necessary level of performance than it does with providing enough capacity. If hot data is automatically moved to more capable storage, this type of overprovisioning becomes unnecessary, reducing spindle count and allowing the majority of data to reside on higher capacity SATA disk drives. This results in savings on power and cooling cost while freeing up valuable rack space.
    - **Improve application performance** without overallocating resources**.** The ultimate goal for most data centers is to deliver the same or better application performance without permanently allocating expensive resources to each application. Many applications have predictable spikes in activity (quarter-end, holiday, and so on) with much lower activity at other times. Ideally, an automated storage tiering solution addresses these as well as unpredictable spikes that can also occur.
    - **Better align storage and server scaling capabilities** for shared IT infrastructure. Advanced server virtualization makes it possible to dynamically adjust server resources to adapt to changing needs. Automated storage tiering can provide a similar level of flexibility for storage, allowing I/O to be scaled dynamically.

CHOOSING THE RIGHT NUMBER OF TIERS

A key consideration in an automated storage tiering strategy is the number of tiers that are needed. The decision to introduce a distinct media type, in fact, should be based not only on the particular specifications of that media but on its overall impact on your performance, cost and efficiency goals. Although there may be clear price/performance differences among the various technologies, this does not mean that more physical tiers are necessarily better.

As suggested in Figure 1, there can actually be a system performance penalty as more tiers are added. This is because data movement across tiers consumes disk-drive I/O operations and other system resources which are precious commodities from a storage performance perspective. If a specific implementation of automated tiering requires wholesale data movement, then it will probably be necessary to perform the data migration during off peak hours. Delayed migration can reduce the value of automated storage tiering if the data is no longer in high demand by the time it reaches the performance tier. Additionally, this type of delay can affect the predictability of system performance.

The relatively small performance difference between hard disk types is also unlikely to justify the I/O overhead associated with the data movement between hard disk drives. The benefits of moving data between Flash media and hard disk drives are much greater. Flash-based media offers orders of magnitude improvement in performance per dollar versus hard disks. This, by itself, does not eliminate the cost of the data movement issue but the relative benefit in terms of improved access to “hot” data is much more substantial. As the cost of Flash media decreases over time and as Flash capacities increase, it will be increasingly possible for this tier to absorb most of the IOPS intensive load (hot data). It follows that the HDD tier is best optimized for lowest $/GB which again points to SATA as the logical choice.

An additional consideration is that each additional tier adds a level of sizing complexity and management overhead. If each storage tier is not sized correctly, there will be a direct impact on efficiency and overall cost.

Figure 1) I/O overhead increases as additional physical storage tiers are added

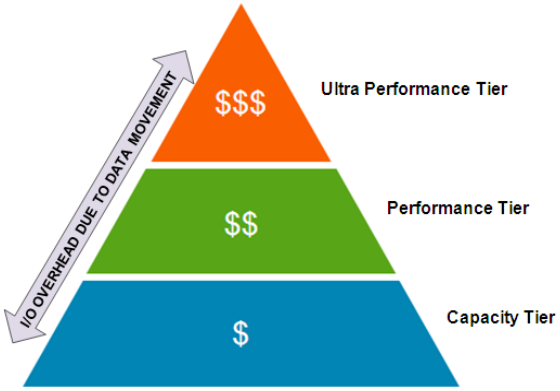

Because of these factors, NetApp believes that going forward the best tiering strategy in most cases will be two tiers: a tier for capacity and a tier for performance. The Virtual Storage Tier is designed with this principle in mind. You probably have a variety of disk drive types such as FC, SATA and SAS. Any of

these can serve as a capacity tier while the Virtual Storage Tier provides performance. As shown in Figure 2, the combination of a high performance tier (based on Flash and intelligent caching) and a single disk drive tier (based on SATA disk) is expected to make the most sense for the majority of applications going forward.

Figure 2) Evolution of storage tiers

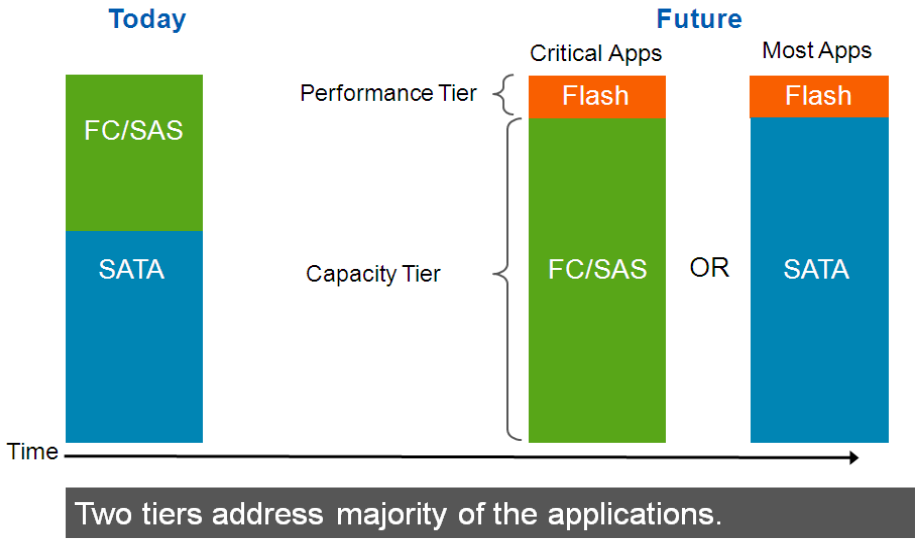

3 THE NETAPP VIRTUAL STORAGE TIER

The NetApp Virtual Storage Tier is NetApp's approach to automated storage tiering. We had several important goals in mind when we set out to design Virtual Storage Tier components:

- Use storage system resources as efficiently as possible, especially by minimizing I/O to disk drives
- Provide a dynamic, real-time response to changing I/O demands of applications
- Fully integrate storage efficiency capabilities so efficiency is not lost when data is promoted to the Virtual Storage Tier
- Use fine data granularity so that cold data never gets promoted with hot data thus making efficient use of expensive Flash media
- Simplify deployment and management

The Virtual Storage Tier is a self-managing, data-driven service layer for storage infrastructure. It provides real-time assessment of workload priorities and optimizes I/O requests for cost and performance without the need for complex data classification and movement.

The Virtual Storage Tier leverages NetApp's key storage efficiency technologies, intelligent caching, and simplified management. You simply choose the default media tier you want for a volume or LUN (SATA,

FC or SAS). Hot data from the volume or LUN is automatically promoted on demand to flash-based media.

Figure 3) NetApp Virtual Storage Tier

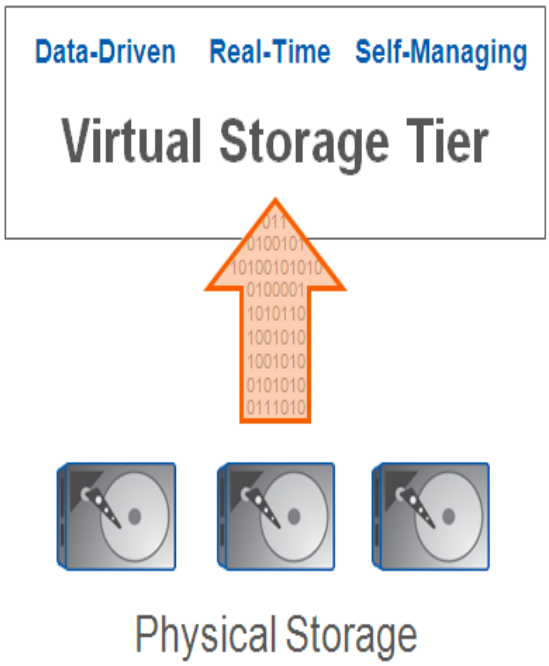

HOW THE VIRTUAL STORAGE TIER WORKS

The NetApp Virtual Storage Tier promotes hot data without the data movement overhead associated with other approaches to automated storage tiering. Any time a read request is received for a block on a volume or LUN where the Virtual Storage Tier is enabled, that block is automatically subject to promotion. Note that promotion of a data block to the Virtual Storage Tier is not data migration because the block remains on hard disk media when a copy is made to the Virtual Storage Tier.

With the Virtual Storage Tier, data is promoted to Flash media after the first read from hard disk drives. This approach to data promotion means that additional disk I/O operations are not needed to promote hot data. By comparison, other implementations may not promote hot data until it has been read from disk many times, and then additional disk I/O is still required to accomplish the promotion process.

Our algorithms distinguish high-value data from low-value data and then retain that data in the Virtual Storage Tier. NetApp metadata, for example, is always promoted when read for the first time. In contrast, sequential reads are normally not cached in the Virtual Storage Tier unless specifically enabled because they tend to crowd out more valuable data.You can change the behavior of the intelligent cache to meet the requirements of applications with unique data access requirements. For example, you can configure the Virtual Storage Tier to cache incoming random writes as they are committed to disk and to enable the caching of of sequential reads.

You can optionally create different classes of service by enabling or disabling the placement of data into the Virtual Storage Tier on a volume-by-volume basis.

BENEFITS OF THE VIRTUAL STORAGE TIER

The NetApp Virtual Storage Tier delivers on all of the promise of automated storage tiering by eliminating the need for manual data management, reducing storage costs, improving application performance, and addressing the needs of shared IT infrastructures. Specific benefits are described in the following subsections. Performance, as measured in laboratory and real-world environments, is described in section 4.

REAL-TIME PROMOTION OF HOT DATA WITH HIGH GRANULARITY

A data block typically enters the Virtual Storage Tier the first time it is read from disk. The performance benefit occurs in real time as subsequent reads are satisfied from the Virtual Storage Tier. The Virtual Storage Tier also identifies patterns of read behavior and will read blocks that are likely to be needed ahead of time, but it never does wholesale movement of data from one tier of storage to another. Caching may optionally be enabled for newly written blocks to further enhance performance of the first read after write. The Virtual Storage Tier contains temporary copies of blocks that remain in the underlying disk drive tier. The efficiency of this approach, combined with the ability to operate at the granularity of a single 4KB block, allows real-time promotion of hot data.

With automated storage tiering solutions from other vendors, hot data is migrated from one storage tier to another either as a background task or scheduled during off peak hours (to minimize the extra load on the storage system). Because these solutions typically operate at a level of granularity that is a minimum of 128X higher than the Virtual Storage Tier (ranging from 0.5MB up to 1GB or even an entire volume or LUN) data movement can take considerable time. Such approaches can miss important spikes of activity when those spikes have a shorter duration than the time needed to identify and promote hot data.

The fine granularity of the Virtual Storage Tier also means that it uses Flash-based media very efficiently. Solutions with coarser granularity are likely to include a lot of “cold” data along with each hot data block and are therefore likely to require a greater amount of expensive Flash media to achieve the same level of benefit.

ENHANCED STORAGE SYSTEM PERFORMANCE

The way that the Virtual Storage Tier handles data enhances storage system performance. Because only hot data blocks are promoted (not large chunks of data that may include both hot and cold blocks) and because data does not have to be moved back to the original tier when it becomes cold, the requirement for data movement is eliminated. This reduces the processing burden on system resources and reduces traffic in the disk subsystem. The result is better performance overall, even for volumes and LUNs where the Virtual Storage Tier is not enabled.

REDUCED STORAGE AND DATA CENTER COSTS

The NetApp Virtual Storage Tier can significantly decrease the cost of your disk purchases and make your storage environment more efficient. For example, testing in a Windows® file services environment [3] showed that:

- Combining Flash media with Fibre Channel or SAS disks can improve performance while using 75% fewer spindles and decreasing purchase price by 54%, while at the same time saving 67% on both power and space.

[3] NetApp *TR-3842 NetApp FlashCache in Windows File Services Workloads*

- Combining Flash media with SATA disks can deliver the same performance as Fibre Channel or SAS disks and more capacity while lowering cost per TB of storage by 57% and saving 66% on power and 59% on space.

See section 4 for information on application-specific benefits and savings.

EASY TO DEPLOY AND SIMPLE TO MANAGE

The Virtual Storage Tier works with your existing data volumes and LUNs. It requires no complicated or disruptive changes to your existing storage environment. There is no need to set policies, thresholds or time windows for data movement. To deploy the NetApp Virtual Storage Tier, you simply install Flash technology in your storage systems. Once this is accomplished, the Virtual Storage Tier becomes active for all volumes managed by the storage controller. You can then exclude user data for lower-priority volumes from the Virtual Storage Tier if desired.

Other automated storage tiering solutions may require incremental policy, data classification and structural changes to existing storage infrastructure such as the creation of dedicated storage pools and migration of data.

FULLY INTEGRATED

The Virtual Storage Tier is fully integrated with the NetApp Unified Storage Architecture, which means that you can use it with any NAS or SAN storage protocol with no changes. Other vendors may offer one automated storage tiering solution for NAS and a different solution for SAN.

In addition, other automated storage tiering solutions may not interoperate with storage efficiency features such as deduplication, compression, and so on. The NetApp Virtual Storage Tier works in conjunction with all NetApp storage efficiency features including thin provisioning, FlexClone® technology, deduplication, and compression, and this close integration works to your advantage and enhances the functioning of the Virtual Storage Tier over the lifecycle of your data.

For example, when you deduplicate a volume, the benefits of deduplication persist in the Virtual Storage Tier. A single block in the Virtual Storage Tier could have many metadata pointers to it, increasing the probability that it will be read again and thus increasing the value of promoting that block. With this *cache amplification* a single block in the Virtual Storage Tier can serve as several logical blocks. This can yield significant performance benefits for server and desktop virtualization environments (such as shortening the duration of boot storms) while at the same time reducing the amount of Flash media needed.

IDEAL FIT FOR ADVANCED VIRTUALIZATION AND CLOUD ENVIRONMENTS

The NetApp Virtual Storage Tier is a perfect complement to advanced virtualization and cloud environments, because it allows storage performance to scale to match the scaling in server performance.

In cloud environments, the Virtual Storage Tier can also be offered as a premium storage service that can be enabled on request without the need for any storage changes. By combining the Virtual Storage Tier with NetApp workload prioritization, multiple, graduated classes of service are possible, allowing many applications or users to share the same storage system.

4 NETAPP VIRTUAL STORAGE TIER PERFORMANCE

Because the elements of the Virtual Storage Tier have been tested in both laboratory and real-world environments, NetApp can provide significant guidance on the best ways to deploy the technology and on the benefits that you can expect to see.

SERVER AND DESKTOP VIRTUALIZATION AND CLOUD

Both server virtualization and virtual desktop infrastructure (VDI) create unique storage performance requirements that the Virtual Storage Tier helps to address. Any time you need to boot a large number of virtual machines at one time—for instance during daily desktop startup or, in the case of server virtualization, after a failure or restart—you create a significant storage load.

One regional bank with over 1000 VMware® View™ desktops was seeing significant storage performance problems with their previous environment despite having 300 disk spindles. When that environment was replaced with a NetApp solution using just 56 disks plus the Virtual Storage Tier, outages due to reboot operations dropped from 4-5 hours to just 10 minutes. Problems with non-responsive VDI servers simply went away and logins, which previously had to be staggered, routinely completed in just 4 seconds. The addition of NetApp Virtual Storage Tier gave the bank better performance at lower cost.

Because much cloud infrastructure is built on top of server virtualization, cloud environments experience many of the same benefits from the Virtual Storage Tier. In addition, you can fully define classes of service for different tenants of shared storage in a multi-tenant cloud environment, expanding your ability to deliver IT as a Service.

DATABASE

The Virtual Storage Tier offers significant benefits in online transaction processing environments. A NetApp white paper[4] examined two methods of improving performance in an I/O-bound OLTP environment: adding more disks or adding Flash Cache. Both approaches were effective at boosting overall system throughput. The Flash Cache configuration:

- Costs about 30% less than the same system with additional disk
- Reduces average I/O latency by 38%
- Consumes no additional power or rack space. (The configuration with additional disk increases both by a factor of more than 2.)

EMAIL

Email environments with large numbers of users are extremely data-intensive. The NetApp Virtual Storage Tier can significantly boost performance at a fraction of the cost of adding disks. For example, in recent NetApp benchmarking with Microsoft® Exchange 2010, addition of the Virtual Storage Tier with no other changes doubled the number of IOPS achievable and increased the supported number of mailboxes by 67%[5].

[4]NetApp *WP-7082: Using NetApp Flash Cache in Online Transaction Processing*.
[5]NetApp *TR-3867: Using Flash Cache for Exchange 2007*

5 CONCLUSION

Although automated storage tiering is emerging as an important mechanism for controlling storage costs and optimizing performance, there are substantial differences between implementations. It's important to understand your goals and then ask appropriate questions such as:

- How is hot data identified?
- How quickly is hot data promoted?
- What's the granularity?
- How difficult is the solution to deploy and manage?
- How robust is the solution?

The NetApp Virtual Storage Tier offers a unique approach to automated storage tiering. Enabled by our foundational strengths in storage efficiency and intelligent caching, the Virtual Storage Tier allows you to scale performance and capacity while achieving the highest level of storage efficiency in the industry. The result is maximum flexibility for your data center infrastructure.

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed

© Copyright 2011 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, xxx, and xxx are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. <<Insert third-party trademark notices here.>> All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. WP-7127-0411

www.netapp.com