Automated Storage Tiering on Infortrend's ESVA Storage Systems

White paper

Abstract

This white paper introduces automated storage tiering on Infortrend's ESVA storage arrays. Storage tiering can generate significant advantages, including optimization of storage resources and large cost reductions.

Table of Contents

Infortrend ESVA (Enterprise Scalable Virtualized Architecture) Storage Systems

The Infortrend ESVA (Enterprise Scalable Virtualized Architecture) Series is a leading storage solution designed for mid-range enterprise SAN. At affordable prices, it meets mission-critical storage demands for performance, scalability and reliability with advanced hardware designs and comprehensive data services. On the innovative Enterprise Scalable Virtualized Architecture, various features, including storage virtualization, thin provisioning, distributed load balancing, automated storage tiering, automatic data migration, prioritized volume accessibility, snapshot and replication, are consolidated to realize optimal business benefits. With ESVA systems, users can optimize ROI, simplify storage infrastructure, and maximize application productivity.

For more information about Infortrend's ESVA storage systems, please visit our website: http://esva.infortrend.com/

Storage Tiering Introduction

Storage tiering involves grouping data into different categories and assigning these categories to different types of storage media in order to optimize storage utilization. Data categories can be distinguished based on performance levels, usage frequency, cost/performance considerations and others.

Different types of architectures are used to implement storage tiering. Major differences between architectures revolve around factors such as whether tiering is implemented in a file-based or block-based structure, and whether tiering is either host-based or storage-based. Another important distinction is between software- or hardware-based implementations.

A rudimentary example of storage tiering is the use of tape storage media, whereby companies archive older data on tape while assigning their disk-based media to store current, high-performance applications. However, earlier rudimentary forms of storage tiering almost completely rely on manually initiated processes. These manual processes involve classifying data and determining the importance of data to an organization, as well as configuring different storage media so that data can be migrated between media and different storage tiers can be achieved.

More recently, automatic mechanisms have become a larger part of storage tiering architectures. Embedded algorithms in storage systems help to classify data and determine what type of priority should be assigned to each set of data. In addition, once data has been properly classified, automatic data migration processes help to move data between different types of storage media so that a tiered architecture can be fully optimized.

Two main purposes of storage tiering can be distinguished: ensuring that each application has access to its required performance levels, and migrating data that has already been stored to higher or lower tiers based on data usage patterns, in order to achieve the most efficient distribution of data within a storage system.

Automated Storage Tiering on Infortrend's ESVA Systems

Storage tiering on ESVA systems is block-based and implemented within the storage system hardware. Fully automatic features help users optimize storage utilization, improve performance, and lower operating costs. Besides automatic features, sufficient room is maintained for user input and configurations, ensuring that users can guide tiering operations to their maximum benefit. The fully automatic features complemented by manually-initiated configurations ensure that storage tiering achieves maximum benefits for ESVA users.

In the storage tiering architecture on ESVA systems, a maximum of four tiers (tier 0 to tier 3) can be deployed. The highest tier (tier 0) features the highest performance level, with performance gradually declining towards lower tiers.

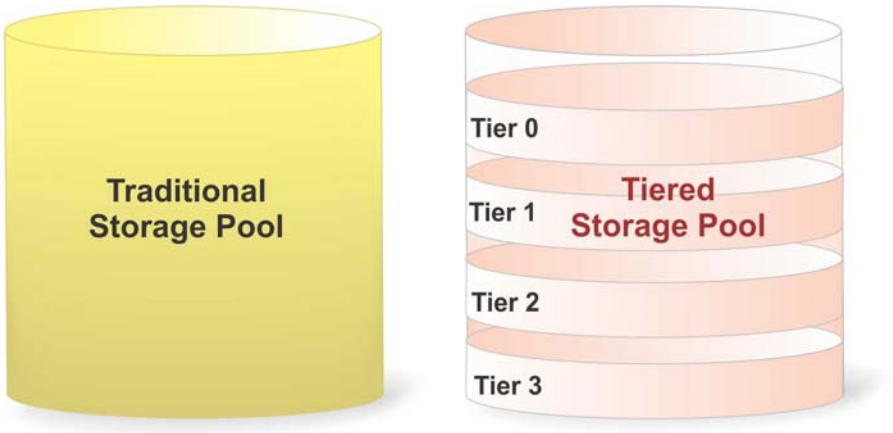

Figure 1: Traditional Storage Pool vs. Tiered Storage Pool

Storage tiers consist of different drive types and RAID configurations to provide different types of performance and utilization for applications. The following are general principles for drive types and RAID configurations:

Drive Type	Characteristics
SSD/SAS drives	Fast, but expensive and low capacity; ideal for high-performance requirements
SATA drives	Inexpensive and more capacity, but relatively slow; ideal for high-capacity requirements

Table 1: General Characteristics of Drive Types

RAID Level	Characteristics
RAID 0	High performance but lack of data protection
RAID 1	Excellent protection and better performance, but least amount of capacity
RAID 5 or RAID 6	Good capacity, protection and performance

Table 2: General Characteristics of RAID levels

ESVA storage tiering features default tier settings that are designed to help users get the most out of their tiered architecture. Default settings are listed below:

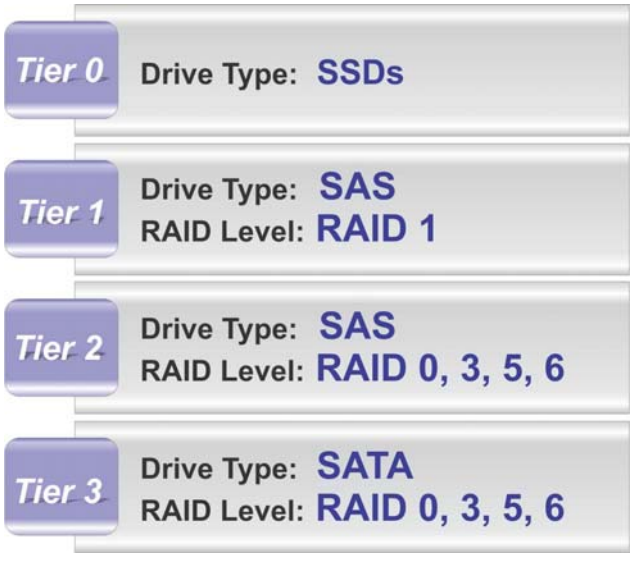

Table 3: Default Tier Settings on ESVA Systems

Leveraging Different Drive Types

ESVA systems support SSDs, SAS drives and SATA drives. Leveraging the characteristics of each of these drive types is an essential part of storage tiering:

- Supported SSDs on ESVA systems include 73 and 146GB STEC ZeusIOPS SSDs. In general, SSDs are ideal for high-performance, mission-critical applications. In a tiered structure, they are by default grouped in tier 0.

- Supported SAS drives include 300, 450 or 600GB 3.5" 15,000 RPM SAS drives. Despite being a step below SSDs in terms of performance, SAS drives are still generally identified as very suitable for high-performance applications, delivering a good balance between performance and capacity. In a tiered structure, SAS drives are mostly allocated to tier 1 and 2.

- Supported SATA drives include 1 or 2TB 3.5" 7,200 RPM SATA drives. SATA drives can be best used for secondary applications that require large storage capacity but do not need the performance levels of SSD or SAS drives.

Application Characteristics Suited for Higher Tiers (SSD/SAS drives)	**Application Characteristics Suited for Lower Tiers (SATA drives)**
Mission-critical applications	Non-mission-critical applications
High-performance applications	Low-performance applications
No significant cost considerations	Significant cost considerations
Low capacity requirements	High capacity requirements

Table 4: Application Characteristics Suited for Different Storage Tiers

Leveraging Different RAID Configurations

RAID levels refer to different ways drives are grouped together. The following are mainstream RAID levels that can be used in Infortrend's storage tiering architecture:

- RAID 0 (Disk striping): delivers maximum performance, without any data redundancy.

- RAID 1 (Disk mirroring): data is mirrored, or copied, on two different disks, offering maximized data redundancy and excellent reliability.

- RAID 3 (Disk striping with dedicated parity): delivers excellent performance, while one dedicated parity drive offers data redundancy.

- RAID 5 (Disk striping with interspersed parity): delivers excellent performance; data redundancy is same as RAID 3, with the difference that parity data is interspersed among all member drives.

- RAID 6 (Disk striping with two interspersed parity disks): A RAID 6 configuration is essentially an extension of a RAID 5 configuration with a second independent distributed parity scheme that offers even better data redundancy.

Allocating Storage Tiers

The aforementioned drive type and RAID configuration characteristics provide guidance for how tiers can be used. The next step is to assess the performance levels, workloads and cost considerations of applications. Applications that require high-performance drives, feature high workloads and do not have significant cost considerations, can be assigned to higher tiers. High-usage database applications could fall into this category.

Non mission-critical applications may feature lower workloads and not need the performance levels of SSDs. From a cost/performance standpoint, these applications could be assigned to lower tiers and SATA drives.

After the nature of applications is established, users can start planning storage tier capacity allocation. With ESVA systems, users have different options at their disposal to optimize this planning.

When creating a virtual volume in an ESVA storage pool, users can choose to which storage tier they wish to assign the volume based on the aforementioned considerations. A virtual volume can reside completely in one tier, or simultaneously reside in multiple tiers (maximum four tiers). If a virtual volume is assigned to reside in two different tiers, users can select specific ratios for these two tiers. The example below shows two different options.

Virtual volume 1 is assigned completely to tier 0, while virtual volume 2 resides in both tier 0 and tier 1 with a 50/50 ratio.

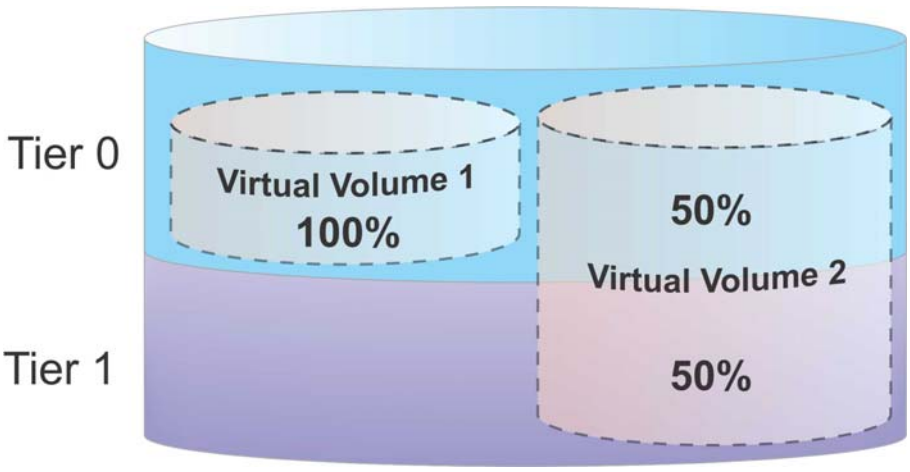

Figure 2: Virtual Volumes in a Tiered Storage Pool

Automatic Migration Processes Optimize Data Distribution

Automatic processes are at the heart of ESVA storage tiering. By default, data always enters the storage system through tier 0. After it has entered the storage system, data present in a storage pool automatically migrates between different tiers based on data activity. System-embedded algorithms drive automatic data migration. These algorithms are based on the age of the data (how long has data been present within the storage system) and the frequency with which the data is used. This algorithm yields a status of data blocks, determining whether the data blocks should be promoted (moved to a higher tier) or demoted (moved to a lower tier).

For example, if data blocks in tier 0 have been found to be idle for an extended period of time, the system will migrate these data blocks to lower tiers. By contrast, if the storage system detects high activity levels of data blocks in lower tiers, these data blocks will be promoted.

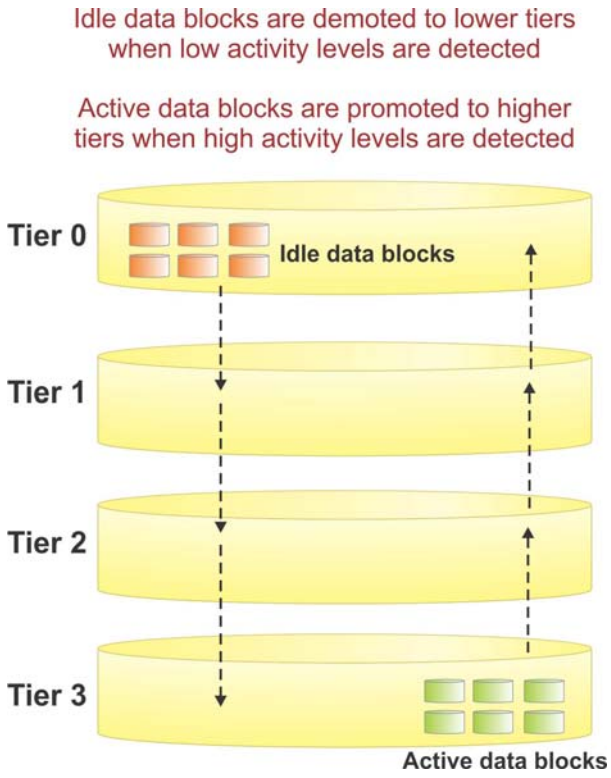

Figure 3: Automatic Data Migration

To make sure data migration conforms to the needs of system users, scheduling options are available. After setting an appropriate schedule, the system will take over and allow IT administrators to spend their valuable time on other important tasks in the datacenter.

Always Available Free Space in Each Tier Ensures Unimpeded Data Migration

Data migration between different storage tiers can only work efficiently if a certain amount of free storage space is available in each storage tier. If not, data blocks will be blocked from entering storage tiers and data migration will grind to a halt.

An exclusive feature of Infortrend's storage tiering is that by default a certain amount of free space is always available on each tier. When data migration is being conducted, users therefore do not need to worry about available storage space. This feature ensures that data migration is truly "automatic" and helps avoid data flow bottlenecks.

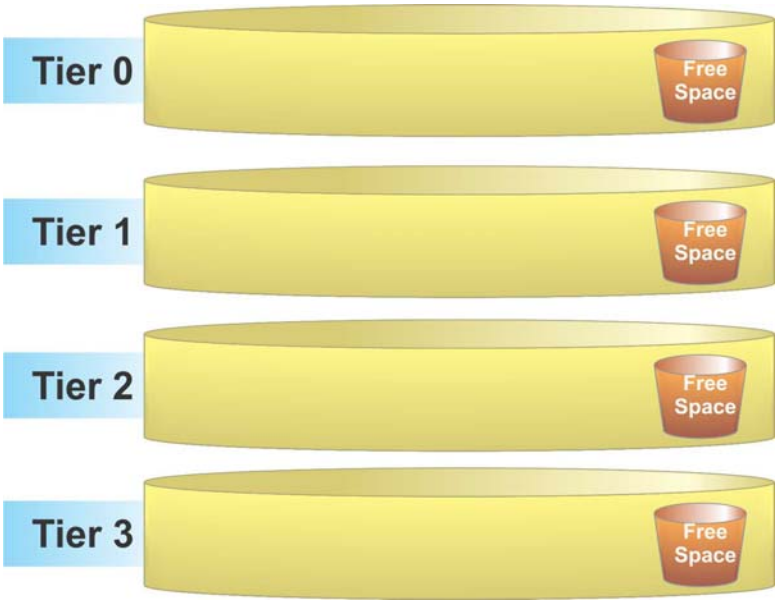

Figure 4: Always Available Free Space in Each Storage Tier

ESVA Storage Tiering Designed to Meet the Vast Majority of Needs

ESVA storage tiering thus consists of three key elements: user tier settings, automatic data migration based on embedded algorithms, and the always-available free space that ensures smooth migration between tiers. Storage tiering thus designed can meet the demand of the vast majority of users, helping generate significant business advantages. If needed, the storage tiering mechanism on ESVA systems can be fine-tuned for specific and complex requirements. This fine-tuning goes beyond the scope of this white paper.

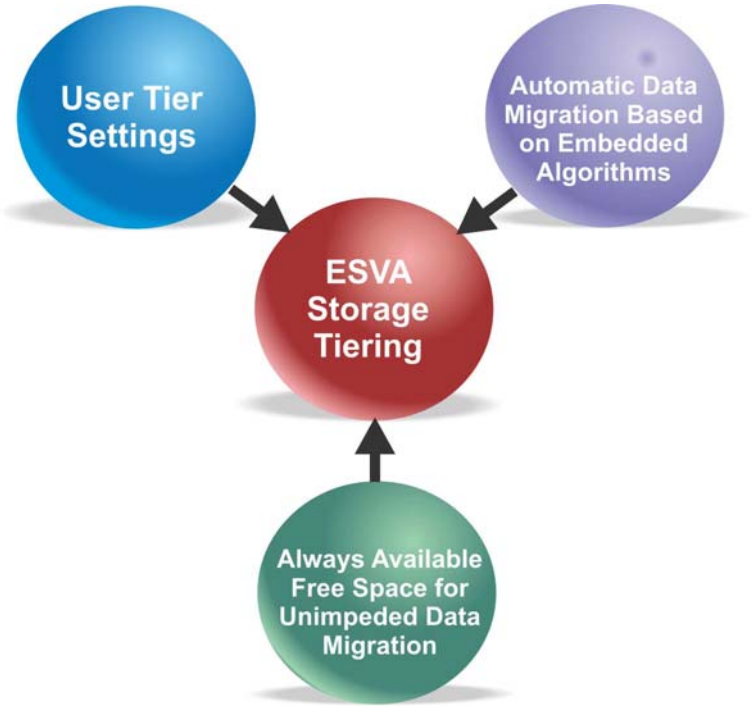

Figure 5: Key Elements in ESVA Storage Tiering Technology

Automated Storage Tiering Benefits

Improve Storage Performance and Utilization

Storage tiering on ESVA systems helps to optimize storage utilization and performance. Most importantly, the technology ensures that applications have access to the performance levels they need. High-performance applications can be assigned to high-performance tiers featuring SSDs or SAS drives, while applications requiring less performance can be assigned to lower tiers featuring SATA drives. Assigning applications in this way ensures that no storage resources are wasted and applications can function properly.

Another key element that helps improve storage utilization is automatic data migration. Without this feature, data would always remain on the same storage media it was first created on, unless manually moved to other storage media. Data usage can vary over time, and it might be no longer suitable for the data to be stored on their original storage media. For example, a certain application might have initially stored data on SSDs in high-performance tiers as it was heavily used and was of high importance. After a period of time, this data might become idle and does not need to be stored on the SSDs any longer, and without moving the data the SSDs might quickly reach their capacity. On traditional storage systems, however, moving data to other storage media can be very complex and time-consuming.

With ESVA storage tiering, data is automatically migrated based on data usage patterns. This ensures that SSDs and other high-performance drives have the necessary storage space to deal with high-performance applications and are not filled to capacity with older, less-frequently used data.

The automatic characteristics of storage tiering on ESVA systems also help reduce the workload of IT managers. Whereas in the past they would have had to manually configure many different settings to achieve a tiered storage environment, storage tiering on ESVA significantly reduces these manual processes and allows IT managers to spend their time on other important tasks in the datacenter.

Economic Advantages

The abovementioned tiering benefits can generate significant economic advantages. If storage space can be more efficiently used with the help of storage tiering, storage systems acquired by enterprises will have a longer lifespan,

increasing ROI. Furthermore, more efficient storage utilization also means more can be done with storage space, reducing storage footprint and lowering power consumption. Power consumption is further reduced by the use of SSDs in storage tiering architectures, as these drives consume only limited amounts of power when compared to other drive types.

In addition, the storage tiering embedded in ESVA storage systems eliminates the need for 3rd-party tiering or data migration solutions, further lowering storage acquisition costs.

Storage tiering simplifies and reduces management operations of storage systems. For traditional storage systems, IT administrators have to manually configure data migration operations and assign applications to specific storage media. With storage tiering, many of these operations are done automatically, reducing time administrators have to spend on storage management and thus lowering management costs.

ESVA Automated Storage Tiering Benefits	Details
Achieve ideal performance levels	Applications can achieve the performance they need when assigned to proper tiers
Data automatically distributed in storage system	Automatic data migration ensures that data is distributed to storage media that fits data usage patterns, reducing inefficiencies in data distribution
Reduce workload of IT administrators	Automatic features of tiering eliminate manual configurations and processes for data storage and migration
Lower management costs	Simplified and reduced management tasks help IT administrators save time and money

Improve storage acquisition ROI	Different tiers ensure that performance needs can be satisfied by a single system
	Better storage utilization improves lifespan of storage system
Lower power consumption	Better storage utilization reduces needed storage footprint and power consumption
	SSDs consume significantly less power than other drive types

Table 5: Benefits of ESVA Storage Tiering

Glossary of Terms

- **Mirroring**: RAID technology where two or more identical copies of data are kept on separate disk drives.
- **Parity**: In RAID, parity is a means of duplicating data by emulating it through parity information. Parity can be used to re-create the original data should it be destroyed by disk drive failure.
- **Pool element**: a group of hard disks logically combined to form a single large storage unit.
- **RAID**: Redundant Arrays of Independent Disks. The use of two or more disk drives to improve performance, error recovery, and fault tolerance.
- **Striping**: A method of distributing data evenly across all drives in an array by concatenating interleaved stripes from each drive.
- **Virtual volume**: a group of pool elements logically combined to form a single large storage unit.