

FlashDisk

Automated Storage Tiering

Contents

Product models covered by this document 3
Automated architecture increases storage efficiency..... 4
Data behavior in a tiered volume..... 5
Automated data migration ensures most efficient distribution of data 6
FlashDisk tiering advantage comparison 9
Optimized storage performance and enhanced ROI 12
Simplified storage management..... 15

Products covered by this document

- **FlashDisk FX series storage systems**
- **FlashDisk VX series storage systems**

Automated Architecture Increases Storage Efficiency

Enterprises tend to use many different storage resources to manage their data. However, assigning applications to different storage resources and migrating data when application requirements change can be a time consuming process that lacks efficiency and drives storage costs up.

Enterprises therefore need a solution that enables them to use resources more efficiently and easily achieve variable performance levels that can meet all service requirements at data centers and elsewhere. By adding a significant degree of automation to resource allocation and data migration, FlashDisk automated storage tiering offers an ideal solution to help meet these enterprise needs.

Enabling users to flexibly assign applications to tiers distinguished by different drive types and RAID levels, FlashDisk automated storage tiering provides an architecture that fully leverages the advantages of different storage media, including SSDs for high performance, and helps users more easily accommodate and meet different service level requirements. As mentioned earlier, users can tap greatly optimized storage performance for increased ROI.

Automated storage tiering on FlashDisk systems is block-based and implemented within the storage solution hardware. Four tiers can be deployed, with the highest tier (tier 0) featuring the highest performance level.




Figure 1. FlashDisk tiering enables users to flexibly assign applications to tiers distinguished by different drive types and RAID levels

Data Behavior in a Tiered Volume

1. Host I/O priority:
For performance considerations, data from the host is always written to its upper tier if space is available.
2. Storage services:
For space consideration, data used for storage services is allocated to lower tiers with minor performance consideration.
3. How to trigger re-allocation of data through ratio settings:
User can trigger re-allocation of data by executing a tier migration operation to proper tier levels manually or automatically through scheduled tasking.

Automated Data Migration Ensures Most Efficient Distribution of Data

Automated data migration ensures data is dynamically stored based on actual usage without complex and manually-driven processes. The system automatically analyzes data hotness (based on access frequency and age, or tenure) and moves data to the optimal storage tiers.

The typical behavior of tiering and migration involves data passing through logical volumes starting with the highest tier first, and users can trigger tier migration to move data between different tiers to re-arrange data blocks to the most suitable tier. Migration is based on a proprietary algorithm that takes into account the age of data (how long it has been present within the storage system) and the frequency with which the data is used. This algorithm yields the status of data blocks and determines which ones should be promoted (moved to a higher tier) or demoted (moved to a lower tier).

With sub-volume tiering on FlashDisk systems, a partition (LUN) can simultaneously reside in multiple tiers, allowing data to move in a highly granular fashion. Ratios for different tiers can be dynamically and intelligently adjusted by the system based on migration results. If two tiers are in one thin-provisioned partition, the default reside ratio of each tier is averaged (ex. a thin-provisioned partition across tier 0 and tier 1 allocates 50% to each tier). If a thin or full-provisioned partition only has one tier, the default reside ratio for that tier is 100%. Reside ratios can also be modified intelligently by the system to achieve optimization, and users cannot modify these reside ratio themselves. A full-provisioned partition can only be created and resided on one specific tier.

Tier migration is triggered by user-defined scheduling, automatically, or manually without scheduling by users, after which the system takes over and conducts migration based on the aforementioned algorithm. To make sure tier migration is conducted based on user requirements, scheduling options are available.

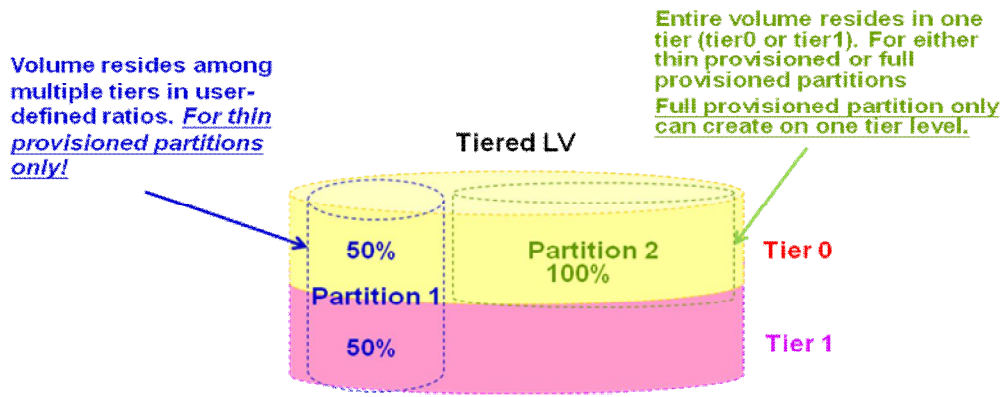


Figure 2. The default reside ratio is averaged for each tier

What kind of data blocks should be promoted? Those we refer to as “hot data” (active data), are defined as follows:

1. All data enters the storage pool though highest tier first.
2. Hotness prioritization is weighed by data residence age and access frequency.
3. Newly written and heavily accessed data is marked as hotness data.
4. Hotness data has higher priority for residing in higher performance tiers when migration is conducted.




Figure 3. All data enters the storage pool though highest tier first

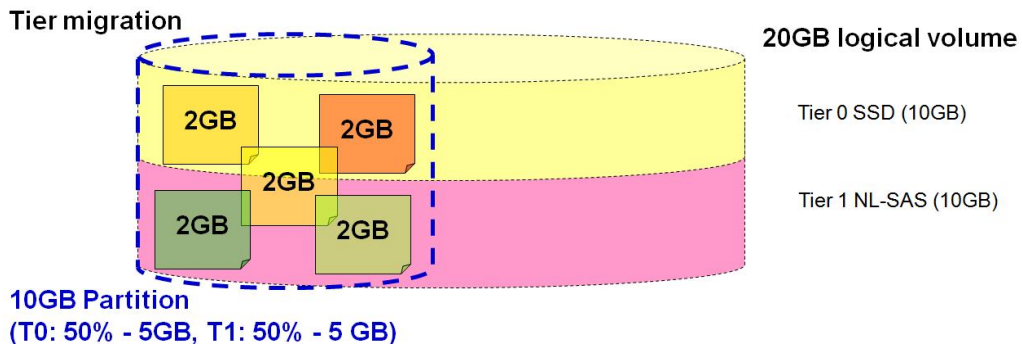


Figure 4. Hot data has higher priority for residing in higher performance tiers when migration takes place (darker color indicates hot data)

How can we make sure new data is written into the top tier for higher performance? We already have a reserve ratio, set by default to 10% for each tier (excluding the lowest tier). When tier migration is triggered, each tier follows the reserve ratio to flush 10% of its data downward.

As shown in the below migration behavior schematic, if a user creates a thin-provisioned partition across two tiers (tier 0 and tier 1), all data is written to tier 0 first if it is still available, and then the remaining data is written to tier 1. After tier migration is triggered, tier 0 flushes 10% of its allocated space to tier 1, which means tier 0 has 10% high speed capacity available for new data.




Figure 5. Reserving space for a tier

FlashDisk Tiering Advantage Comparison

FlashDisk systems provide an easy way to meet customer requirements. Please refer to the below comparison table for more about FlashDisk tiering advantages:

WinSys	IBM	HP	EMC	DELL	
Automated Tiering	Easy Tier	Adaptive Optimization Software	FAST (Fully Automated Storage Tiering)	EqualLogic Tiered Storage	Compellent Data Progression
tier a volume up to 4 tiers	2 tiers	3 tiers	3 tiers	3 tiers	3 tiers
define tiers by disk type and RAID level	only disk types	yes	only disk types	yes, but only support mixing SSD and SAS	yes
define tier reside ratio automatically or manually	no	no	yes	no	yes
schedule data migration amongst tiers	no, moves data every 24 hours	yes	yes	no	yes

1. Tier a logical volume with up to four tiers:

Suitable for diverse applications and requirements, such as ERP/SAP, email, and file sharing, which can be consolidated in a single logical volume.

For FlashDisk systems without storage tiering support, we recommend users create a logical volume with one drive type and only one RAID level. If users want to use different drive types in the same data set, it is better to create another logical volume. With FlashDisk automated storage tiering support, users can put two drive types together and leverage all their different benefits.




Figure 6. Meeting high performance and large capacity demands with two tiers

- Defining tiers by different disk types and RAID levels:
For example, FlashDisk systems can use RAID 1 SSD and RAID 6 NL-SAS as two tiers in one partition.

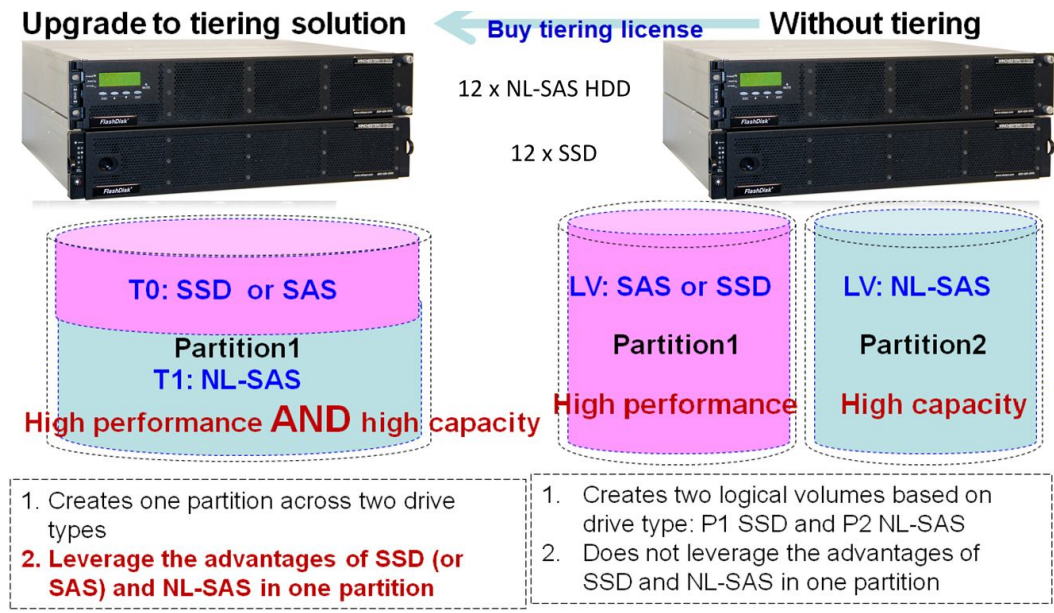


Figure 7. Leveraging advantages of each drive type and RAID level with tiering

- Defining the tier reside ratios of volumes using system intelligence or default settings:
Reside ratios for different tiers can be dynamically and intelligently adjusted by systems after tier migration is triggered. For a two-tier thin-provisioned partition, the default reside ratio is 50% each for T0 and T1. One-tier partitioning can be applied to thin or full-provisioned data, and the default value is 100%.

- 4. Triggering data movement manually or using scheduled timing:
Users can trigger data movement (tier migration) by scheduling it in advance or manually after work hours to minimize performance impact. Customers can also choose tier migration priority (low, normal, high) to suit their application environment.

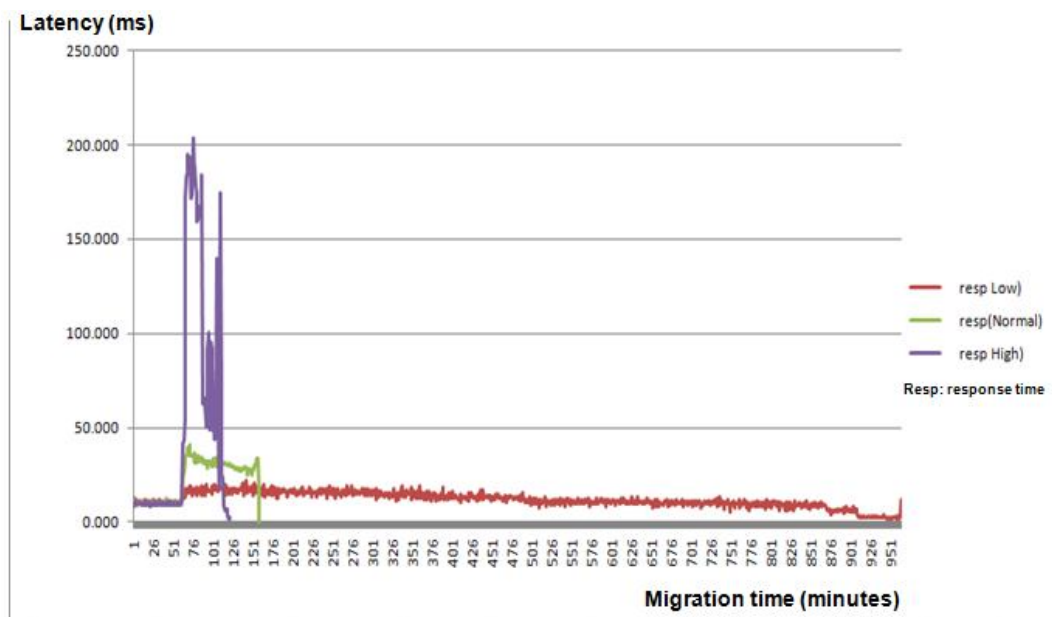


Figure 8. Deciding migration priority (low, normal, high) to suit application environment

Optimized Storage Performance and Enhanced ROI

Automated storage tiering on FlashDisk systems helps users significantly optimize storage performance and increase ROI. Performance is optimized by efficiently integrating SSDs, which deliver performance levels far beyond that which SAS or near line drives are able to offer. Adopting SSDs in a tiered architecture enables users to meet the stringent demands of mission-critical applications in a cost effective way.

When compared to a traditional FlashDisk configuration consisting of one RAID array and one expansion enclosure equipped with 32 x 600GB SAS 10K drives, a FlashDisk configuration with automated storage tiering that consists of only one RAID array equipped with four 400GB SSDs and twelve 4TB near line drives can offer 19.1 times better write performance and 11.72 times better read performance plus around 2.5 times bigger capacity, all at the same cost of a traditional FlashDisk configuration.

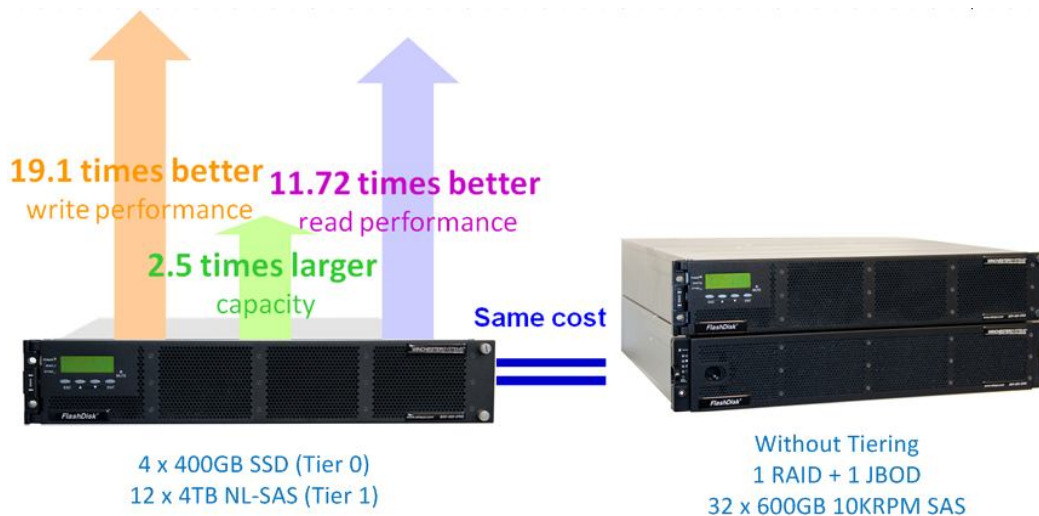


Figure 9. Comparing tiered and traditional configurations with the same purchase cost

For example, if a customer already has one RAID system with 16 x 600GB 10K SAS drives and would like to improve current performance by 15 times, they would need to purchase 15 x JBOD enclosures and 224 x 10K SAS drives, which is very expensive in terms of hardware cost.

With tiering, the user simply needs to purchase a tiering license, one JBOD, four SSDs and 12 x 600GB 10K SAS drives, and then get around 38 times better write performance and approx. 23 times better read performance. More importantly, this setup is also around 8 times cheaper compared to a traditional FlashDisk configuration of the same performance level.

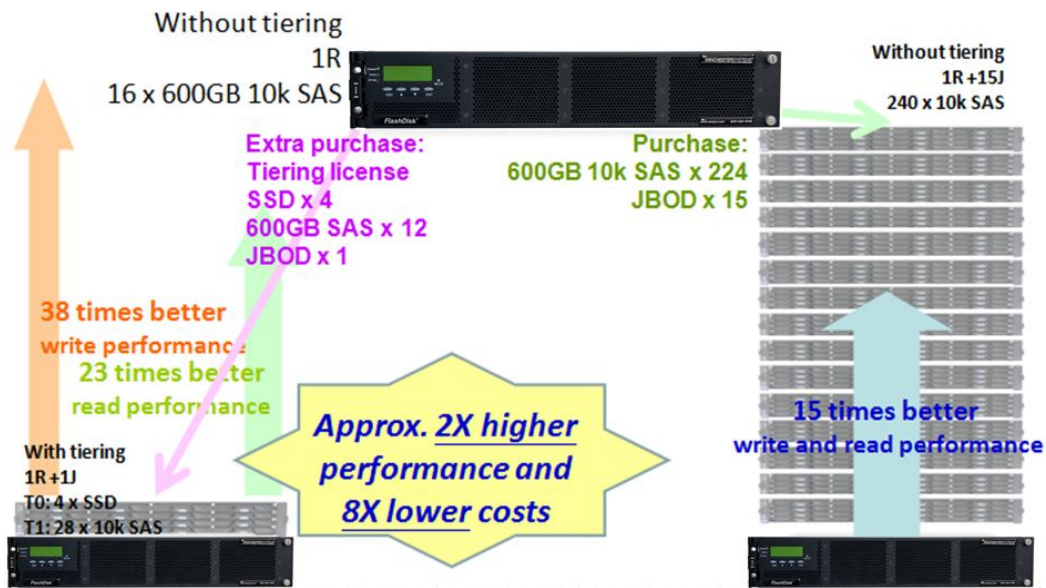


Figure 10. Comparing tiered and traditional configurations for improving performance

In another example, a customer wants to get five times more capacity from one RAID array with 16 x 600GB 10K SAS drives. To do that, they need to purchase five JBOD enclosures and 96 x 10K SAS drives. With a tiering solution, they just need to purchase a tiering license, 16 x 4TB NL-SAS drives and one JBOD. The benefit is clear: the user gets five times the capacity as desired, plus spends around 3.6 times less money.

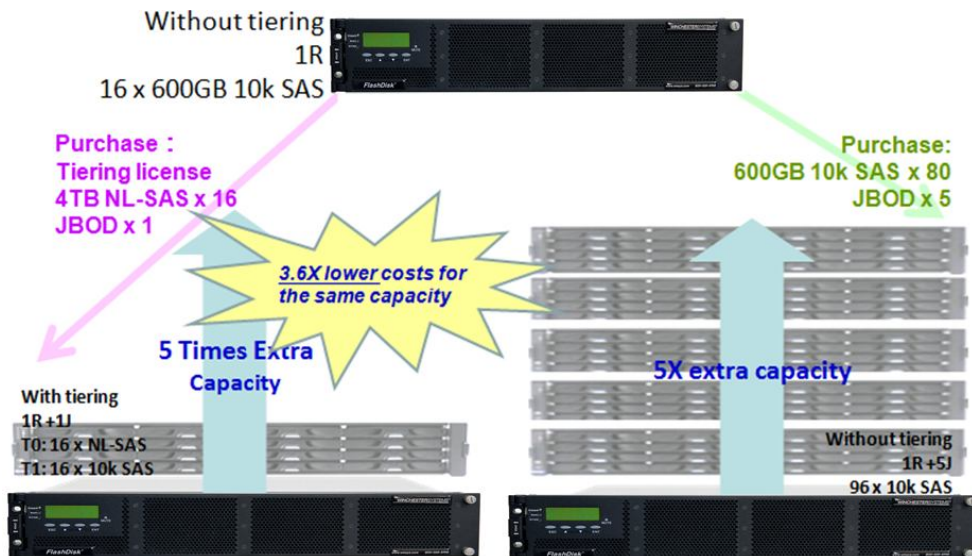


Figure 11. Comparing tiered and traditional configurations for expanding capacity.

With storage resources utilized more efficiently, users can also reduce their physical storage footprint and achieve lower power consumption. Power consumption is further reduced by the use of SSDs, as these drives consume relatively limited amounts of energy. Users can also reduce management overheads using tier migration to move data to optimal locations.

Please check the below table for a broad summary of the advantages of FlashDisk automated storage tiering.

	Non-tiering	Tiering
Cost for performance	High	Low
Cost for capacity	High	Low
Management overhead	High	Low
Space efficiency	Low	High
Power consumption	High	Low

Simplifying Storage Management

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

Automated storage tiering can be easily managed through the FlashDisk software suite, FlashDisk Global Manager. Tiering configurations are fully integrated into the FlashDisk Global Manager user interface, enabling users to setup and maintain tiering operations quickly and efficiently.