AccuRAID iSCSI Auto-Tiering Best Practice

Overview

From the perspective of storage features, the performance of SSDs are high, but the cost is also high per GB. Relatively speaking, the cost of a traditional hard drive is low, so as performance is relatively poor. If we follow the 80/20 rule to configure storage systems, all-SSD configurations are unreasonable for all but the most intensive applications. In fact, SSD will be needed in only a small part for most typical applications, regardless of whether or not a critical application, thus giving SSD resources for general storage needs is hugely cost-prohibitive. Although traditional hard disk performance is enough for general applications which I/O requirements are not high, the traditional all-hard-drive configuration is also gradually been inadequate.

On the other hand, the data itself has a lifecycle. Since the data in the course of its life cycle, it has experienced different levels of activity. In common usage, when creating the data, it is usually used. As the age of the data increases, it is accessed less often.

The Solution

Therefore, to balance performance and cost factors, adapting hybrid storage architecture with a mixture of SSDs and traditional HDDs seem to be the most reasonable approach for modern IT environments. Generally, SSD-based storage capacity in 10 to 15% of the total storage capacity should be enough to fulfill the requirements of critical high I/O applications. An automated tiering pool is a simple and elegant solution for dynamically matching storage requirements with changes in the frequency of data access.

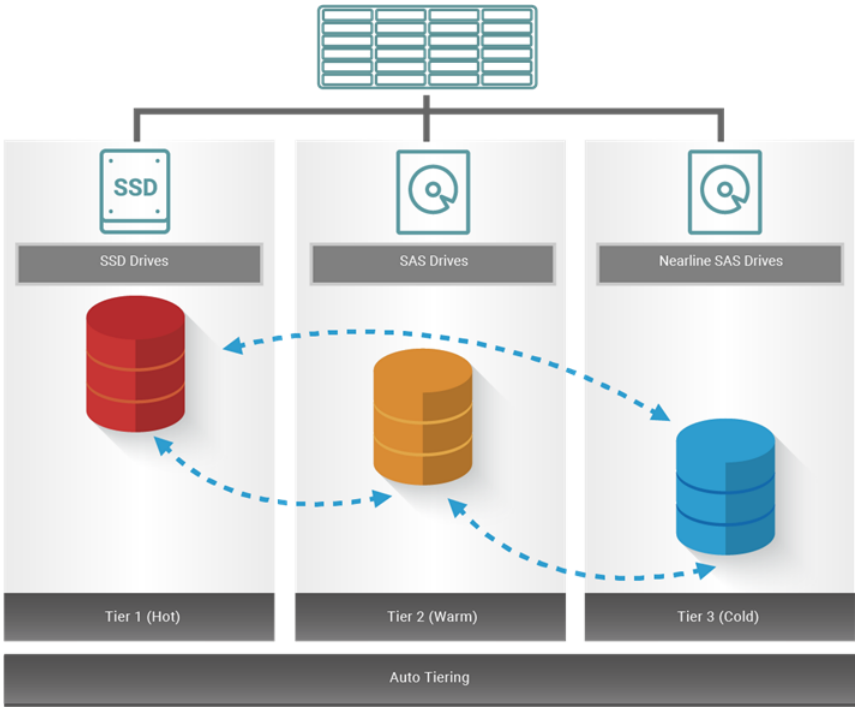

Auto Tiering Pool

Tier Categories

As the name suggestion, auto tiering must have two tiers at least. Automated tiering pool segregated disk drives into three categories for dual controllers and four for single controller.

- Tier 1: SSD drives for extreme performance tier
- Tier 2: SAS drives (15K or 10K RPM SAS HDD) for performance tier
- Tier 3: Nearline SAS drives (7.2K or lower RPM SAS HDD) for capacity tier
- Tier 4: SATA drives for capacity tier (for single controller only, not recommended)

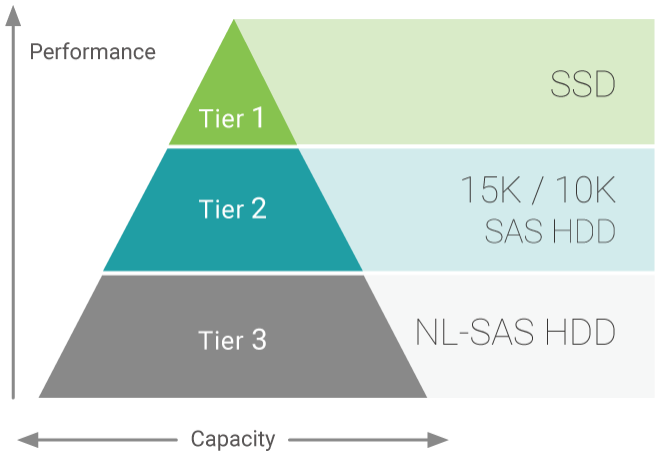

3 Levels of Tiered Storage

Tier 1 / SSD Tier / Extreme Performance Tier

Use the SSD tier when response time and performance are the most important criteria for storage. This tier uses flash technology that does not contain moving parts. This revolutionary technology eliminates the rotation latencies and can improve performance and save energy significantly.

Compared to traditional spinning drives, SSD drives have higher cost per gigabyte, but lower per IO cost. For the best practice, use the SSD drive to get data that requires fast response time and high IOPS. Auto tiering enables you to optimize the use of these high-performance resources because it automatically relocates "hot" data to the SSD tier.

Tier 2 / SAS HDD Tier / Performance Tier

Use the SAS HDD tier to achieve a combination of performance and capacity. The SAS HDD tier provides high levels of performance, reliability, and capacity. SAS HDD stores data on a series of fast rotating disks based on mechanical hard disk drive technology.

This tier includes 15K and 10K RPM spinning drives, which are valuable because it provides a high level performance with consistent response time, high throughput and good bandwidth at moderate price.

Tier 3 / NL-SAS HDD Tier / Capacity Tier

Use the NL-SAS HDD tier to reduce the cost per GB of data. This tier consists of 7.2K or lower RPM SAS HDD which is designed to achieve the maximum capacity at an appropriate performance level. While NL-SAS HDDs have slower speeds than SAS

HDDs, NL-SAS HDDs can significantly reduce power consumption and extend capacity in more expensive and higher performance storage tiers.

In a typical system, most of the application data has very little I/O activity. Because NL-SAS HDDs cost less per GB, they are the most appropriate media type for the "cold" data. NL-SAS HDDs consume less power than SAS HDDs and provide total cost of ownership improvement that take into purchase cost.

Flexible RAID and Disk Configurations

Auto Tiering 2.0 supports flexible RAID and disk configurations. You can create each tier (disk group) with different RAID levels and different a quantity of disks. For example, SSD tier uses 4 disks with RAID 10 for extreme performance, SAS tier uses 6 disks with RAID 6, and NL-SAS tier uses 8 disks with RAID 5 for capacity. This feature is very important for IT administrators to arrange storage plans flexibly.

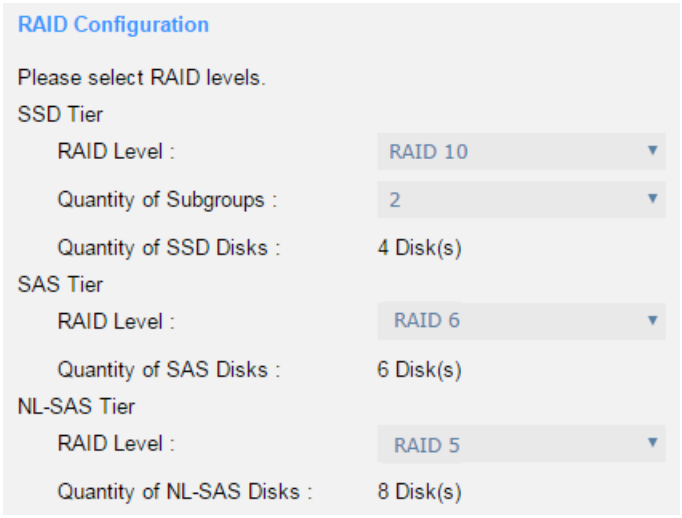

Flexible RAID and Disk Configurations

Theory of Operation

Auto tiering is the automated progression or demotion of data across different tiers (types) of storage devices and media. The movement of data takes place in an automated way with the help of software and is assigned to the ideal storage media according to performance and capacity requirements. It also includes the ability to define rules and policies that dictate if and when data can be moved between the tiers,

and in many cases provides the ability to pin data to tiers permanently or for specific periods of time.

Auto Tiering Architecture

A newly created auto tiering pool is based on thin provisioning technology. Each tier works based on one or more disk groups. The following is the storage architecture of an auto tiering pool.

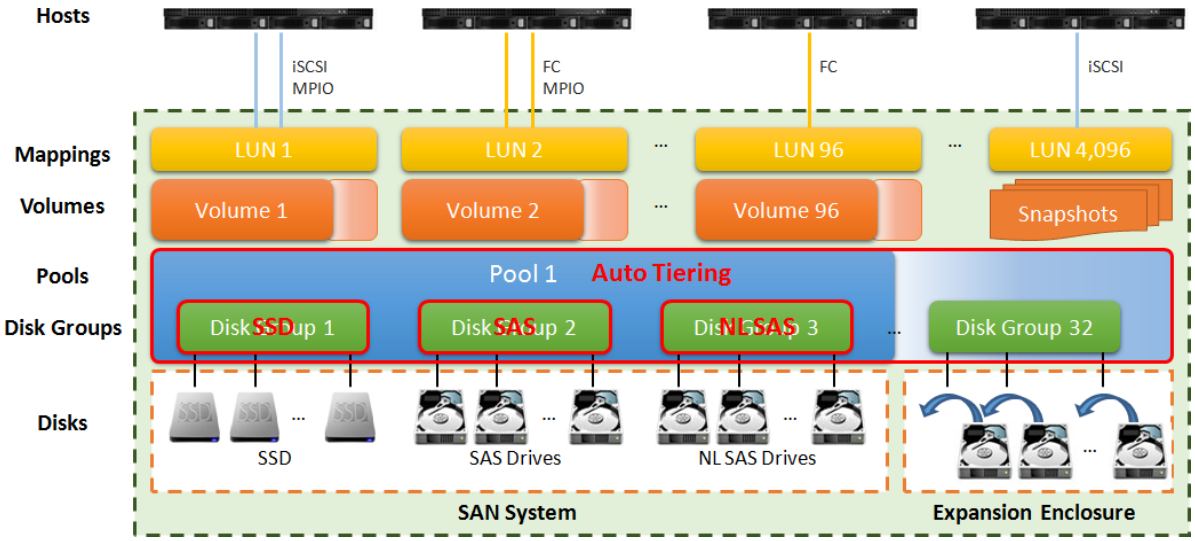

Storage Architecture of Auto Tiering Pool

To increase the capacity of an auto tiering pool, any tier (disk group) which contains either one tier of SSDs, SAS HDDs, or NL-SAS HDDs can be added to the pool any time. An auto tiering pool can have up to 32 disk groups with each disk group contains up to 64 disk drives. And the maximum disk drive quantity in a pool is 256. The maximum addressable capacity of each disk group is 64TB. So the maximum capacity in a system is 256TB.

Auto Tiering Pool Parameters

Item	Value
Maximum disk group quantity in a pool	32
Maximum disk drive quantity in a disk group (include dedicated spares)	64

Maximum disk drive quantity in a pool (include dedicated spares)	256
Maximum pool quantity per system	64
Maximum dedicated spare quantity in a pool	8
Maximum tiers (include SSD, SAS HDD, NL-SAS HDD)	3
Maximum addressable capacity of a disk group	64TB
Maximum addressable capacity of an auto tiering pool	256TB
Maximum addressable capacity of total auto tiering pools (include thin provisioning pools)	1,024TB
Provisioning granularity	1GB

By design, the auto tiering feature allows selecting policies that define how data are moved between different tiers, and in many cases provides the ability to pin data to tiers permanently or for specific periods of time.

Auto tiering storage is the assignment of different categories of data to different disk types. It operates based on relocating the most active data up to the highest available tier and the least active data down to the lowest tier. Auto tiering works based on an allocation unit (granularity) of 1GB and relocates data by moving the entire unit to the appropriate tier, depending on the tiering policy selected for that particular volume.

In order to ensure sufficient space in the higher tiers, 10% of the space is reserved in each higher tier to prepare for the data allocation for those tiering policies which would allocate initial space in highest available tiers. By reclaiming this 10% headroom, the least active units within each tier move to lower tiers. The whole mechanism of auto tiering contains three steps, statistic collection by accessed counts, ranking hotness data by the statistic collection, and then relocation data via ranking.

Intelligent Auto Tiering Mechanism

Auto tiering storage management system manages the data relocation and monitors the data hotness ratio using half-life coefficient and advanced ranking algorithm. It operates on three major functions.

Statistics Collection

The volume space is divided into units of equal size in which the hotness is collected and analyzed per hour. This is also called sub LUN. Activity level of a sub LUN is determined by counting the quantity of read and write access on the sub LUN. Logical volume manager maintains a cumulative I/O count and weights each I/O by how recently it arrived. The new coming I/O is given a full weight. After approximately 24 hours, the weight of this IO is nearly cut in half and continues to decrease. The reduction weight is processing per hour by our precision algorism. This statistics collection occurs continuously in the background for auto tiering pool.

Ranking

This analysis produces a rank ordering of each sub LUN within the pool. Note that the policies of volumes would affect how sub LUNs are ranked.

After analysis, the system would generate following information for each tier:

- The amount of data to be moved up
- The amount of data to be moved down
- The amount of data to be moved into a tier.

TIP:
The hotness analysis process which includes statistics collection and ranking may take minutes to complete.

Relocation

According to the hotness analysis, relocation is processed during the user-defined relocation window, which is the number of minutes given to the relocation process. When the window closes, the relocation process would stop relocating data. The other parameter is relocation rate which controls speed of the relocation process. Valid value of relocation rate is Fast, Medium, and Slow.

Auto tiering promotes sub LUNs according to the candidate list that it created in the analysis stage. During relocation, it prioritizes relocating sub LUNs to higher tiers. At the same time, sub LUNs are only relocated to higher tiers if the space they occupy is required for a higher priority. Using the mechanism, auto tiering makes sure that the higher performing drives are always used.

During I/O, as data is written to a pool, auto tiering attempts to move it to the higher tiers if space is available and the tiering policy allows for it. As we describe before, the relocation process will keep 10% of the free space in all tiers. This space is reserved for any new allocations of higher priority sub LUNs before the next relocation. Lower tiers are used for capacity when needed. The entire relocation process is complete automatically based on the user-defined relocation schedule, or manually if user triggers by himself. The following figure provides an illustration of how auto tiering can improve sub LUN placement in a pool.

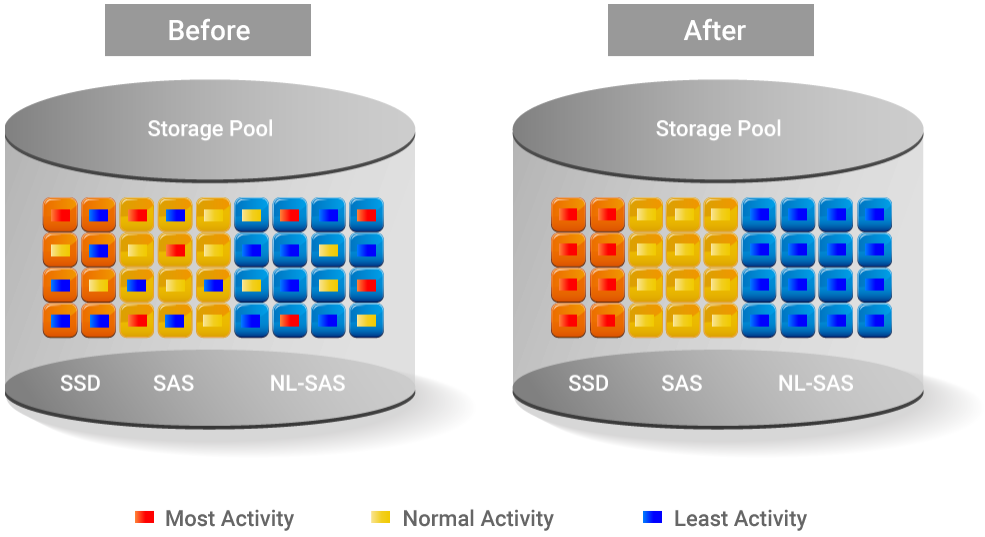

Auto Tiering Relocation

Tiering Policies

For the best performance in various environments, auto tiering has a completely automated feature that implements a set of tiering polices. Tiering policies determine how new allocations and ongoing relocations should apply within a volume for those requirements. Auto tiering uses an algorithm to make data relocation decisions based on the activity level of each unit. It ranks the order of data relocation across all volumes within each separate pool. The system uses this information in combination with the tiering policy per volume to create a candidate list for data movement. The following volume policies are available:

Auto Tiering (Default)

It allows moving a small percentage of the “hot” data to higher tiers while maintaining the rest of the data in the lower tiers. This policy automatically relocates data to the most appropriate tier based on the activity level of each data. Sub LUNs are relocated

based on the highest performance disk drives available and its hotness. Although this setting relocates data based on the performance statistics of the volume, the volume sets with “Highest available Tier” take precedence. Initial space is allocated in the tier which is healthier and has more free capacity than other tiers, then relocated according to hotness of the data. This is the recommended policy and it is the default policy for each newly created volume.

Start Highest then Auto Tiering

This takes advantage of the both “Highest Available Tier” and “Auto Tiering” policies. “Start Highest then Auto Tiering” sets the preferred tier for initial data allocation to the highest performing disks with available space, and then it relocates the volume’s data based on the performance statistics and the auto-tiering algorithm. With this tiering policy, less active data is moved to lower tiers, making room for more active data in the higher tiers. Initial space is allocated in highest available tier first, then relocated according to hotness of the data.

Highest Available Tier

Use this policy when quick response times are a priority. This tier is effective for volumes which require high levels of performance whenever they are accessed. The policy starts with the “hottest” first and places them in the highest available tier until the tier’s capacity or performance capability limit is hit. Then it places the sub LUNs into the second higher tier. Initial space is allocated in highest available tier. Auto tiering would prioritize sub LUNs with highest available tier selected above all other settings.

Lowest Tier

Use this policy when cost effectiveness is the highest priority. With this policy, data is initially placed on the lowest available tier with capacity. Select this policy for volumes that are not performance sensitive or response-time sensitive. Regardless of their activity level, all sub LUN of these volumes will remain on the lowest storage tier available in their pool. Data of volumes with “Lowest tier” policy would always reside in the lowest tier. Changing policy of a volume with data in higher tiers to ‘”Lowest tier” would cause all its data in higher tier to be relocated down to the lowest tier.

No Data Movement

If a volume is configured with this policy, no sub LUN provisioned to the volumes is relocated across tiers. Data remains in its current position, but can still be relocated within the tier. The system still collects statistics on these sub LUNs after the tiering policy is changed. Initial space is allocated in the tier which is healthier and has more

free capacity than other tiers. No relocation would be performed in a volume which selects “No data movement” tiering policy.

The following table summarizes the tiering policies.

Summary of Tiering Policies

Tiering Policy	Description
Auto Tiering	Sets the initial data placement to the optimized tier (disk group) and then relocates the data based on the statistics such that data is relocated among tiers according to the I/O activity.
Start Highest then Auto Tiering	First sets the preferred tier for the initial data placement to the highest tiers with available space, then relocates the data based on the statistics and the auto tiering algorithm.
Highest Available Tier	Sets the preferred tier for the initial data placement to the highest tiers with available space, and so as the succeeding data relocation.
Lowest Tier	Sets the preferred tier for the initial data placement to the lowest tiers with available space, and so as the succeeding data relocation.
No Data Movement	Sets the preferred tier for the initial data to the optimized tier, and retains the data without movement.

SSD Cache vs. Auto Tiering

The SSD cache and auto tiering solutions can work together and complement each other. A key difference between tiering and cache is that tiering moves data to SSD instead of simply caching it. Tiering can also move data both from slower storage to faster storage and vice versa. However, SSD cache is essentially a one-way transaction. When the cache is done with the data it was accelerating it simply nullifies it instead of copying it back to HDD. The important difference between moves and copies is that a cache does not need to have the redundancy that tiering does. Tiering stores the only copy of data for potentially a considerable period of time so it needs to have full data redundancy like RAID or mirroring.

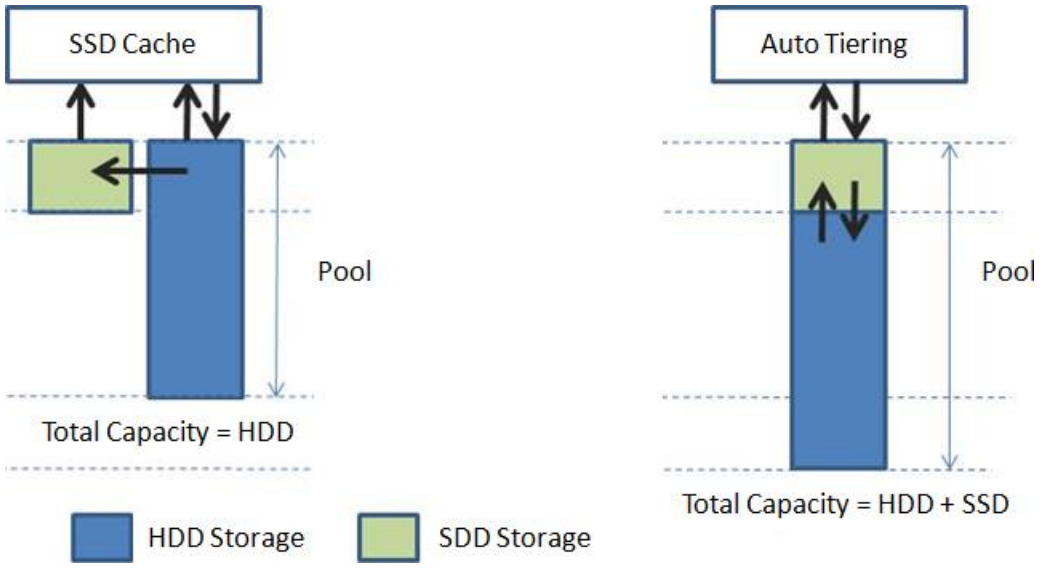

SSD Cache vs. Auto Tiering

Total storage capacity in auto tiering is a sum of all individual tier capacities whereas in cache, the cache capacity does not add to the overall slower storage capacity. This is one of the key differences. In addition, SSD cache affects faster than auto tiering because auto tiering will be affected by relocation the data in a period of time. So SSD cache warm-up timeframe is usually minutes/hours whereas tiering warm-up is usually days.

SSD cache is used for highly frequent data access environments and is effective short term, such as virtualization or video editing applications. However, auto tiering is used for predictable I/O workloads and is effective in long term. It's suitable for web, file, or email server applications.

SSD Cache vs. Auto Tiering

	SSD Cache	Auto Tiering
Total Capacity	HDD	HDD + SSD
When SSD is Damaged	Pool Works Fine	Pool Fails
Performance	Effective in Short Term	Effective in Long Term

Best Practice

Auto tiering technology provides a solution to achieve optimal storage efficiency and improved performance, making it the most cost effective storage solution for data center environments with dynamic workload changes.

If your applications are belongs to sequential I/O from beginning to end, such as surveillance or backup, or their access profiles are very random in the large address range, a homogeneous pool is recommended for your applications. In a homogeneous pool, only one drive type (SSD, SAS, or NL-SAS) is selected during pool creation. If using auto tiering technology in these applications, the data will move up and down frequently without any benefit.

TIP:
Homogeneous pool is suitable for the application of sequential I/O from beginning to end or very random in the large address range. In addition, auto tiering is suitable for the data which has a lifecycle.

Configuration Planning Advice

SSD / SAS / NL-SAS Tier RAID Level and Capacity Ratio

The following is a general guide to the auto tiering pool planning. The user can fine-tune according to the actual situation.

- SSD Tier ($$$)

  Suggest SSD tier using at least 4 disks with RAID 10 (better) or 2 disks with RAID 1 for extreme performance. Prepare SSD storage capacity in 10% to 15% of the total pool capacity to fulfill the requirements of critical high I/O applications.

- SAS Tier ($$)

  Suggest SAS HDD tier configuring with RAID 6 (better) or RAID 5. Prepare about 30% of the total storage capacity.

- NL-SAS Tier ($)

  For capacity tier, suggest NL-SAS HDD using RAID 5 level to store cold data. This tier occupies the rest of the storage capacity.

Best Practice of Auto Tiering

Take an example for reference. First, you can estimate the total capacity used, and estimate how much hot data or high I/O your application uses every day. Assuming 666GB per day, the recommended SSD tier capacity is at least 1.5 times, 1.5 x 666GB = 1TB, as a conservative estimate. Then, calculate the SAS HDD tier capacity about 3 times of the SSD tier capacity, 3 x 1TB = 3TB, as if the SSD tier full of the buffer, so that the performance does not drop too much. This tier is optional. The remaining space is left for NL-SAS HDD tier. The following table is the summary for reference.

Tier RAID Level and Capacity Ratio

Tier	Capacity per Drive	Quantity	RAID Level	Capacity per Tier	Capacity Ratio
SAS SSD Tier	500GB	4	RAID 10	(4/2) x 500GB = 1TB	10%
SAS HDD Tier	1TB	5	RAID 6	(5-2) x 1TB = 3TB	30%
NL-SAS HDD Tier	3TB	3	RAID 5	(3-1) x 3TB = 6TB	60%

This is a rough planning proposal. Whether to meet the requirements also requires users to calculate the performance and necessary capacity. Of course, if more capacity is needed, you can also add a disk group to any tier.

Relocation and Its Effect

In the Intelligent Auto Tiering Mechanism section, we introduce there are three major functions in auto tiering technology. Statistics collection and ranking operate automatically, but relocation can be configurable manually. We would like to suggest that users can set the schedule relocation at midnight every day (**Daily 00:00**), the relocation period sets to 8 hours (**08:00**), and the relocation rate sets to **Fast**. So you can ensure that the performance at working hours will not be affected.

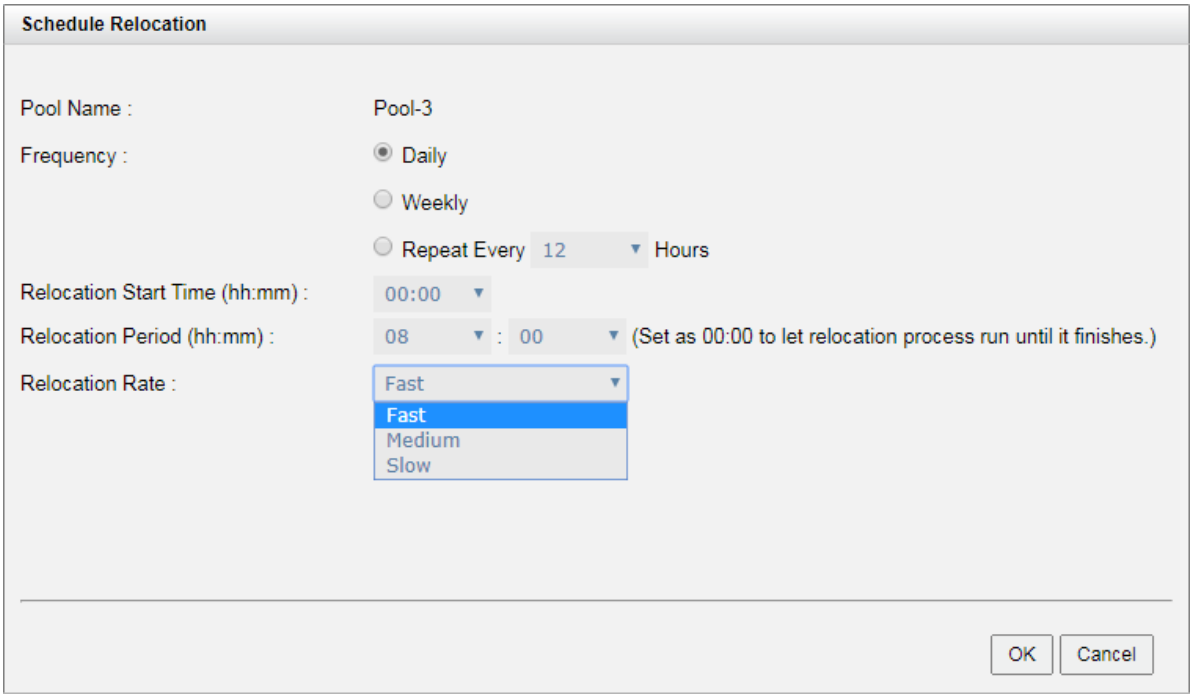

Schedule Relocation Setting

TIP:

If the storage needs to provide 7 x 24 hours of data access services, may or may not find a long period without data access, please try to find a time frame with a slight I/O of inbound and outbound data flow, execute the relocation rate with **Medium** or **Slow** by either schedule or manual for eliminating the possible performance impact.

Also note that performance improvements may not be obvious when using a relocation rate with **Medium** or **Slow** compared to **Fast**, as the execution time is the same, since relocation may not be completed.

Auto Tiering Policies and Their Effect

In the Tiering Policies section, there are five policies described, each policy has a suitable situation.

- **Auto Tiering (**Default**)**

  This can be used in a large volume of storage structure. Usually the user does not know how to put the data to the right tier; it is entirely handled by the storage system. By default, the data will be relocated at midnight. At this case, hot data calculations take a long time to accumulate and move up, and a few fixed blocks require extreme high performance (but usually the user does not understand the situation). Using this policy will have a significant effect.

- **Start Highest then Auto Tiering**

  This can be used for hot data in a short time, such as video editing. The new coming films are often edited at the beginning. After the editing is complete, the files are not always used and eventually moved to the archive. In this scenario, you need to understand the capacity of the hot data and prepare the capacity of the SSD tier. Then this policy can maximize the efficiency.

- **Highest Available Tier**

  This allows users to allocate resources in a timely manner. Assuming that some volumes will be frequently accessed tomorrow, the IT administrator can manually adjust to this policy. As a result, the data will be relocated to the highest available tier at midnight. In this case, you can get better efficiency under the same resources. Of course, the premise is that the capacity of the volume needs to be controlled.

- **Lowest Tier**

  It is for the purpose of data backup, for those volumes which do not need the performance, and the need for large capacity storage of data. It can be set to this policy.

- **No Data Movement**

  This should be least used. The data in the volume using this policy will not operate any hotness analysis. It is suitable for infrequently used data.

As mentioned above, you can choose the right policy based on your application. Or you are unsure, it is recommended to use **Auto Tiering** policy when creating a volume, and the relocation schedule remains in daily. Then observe the usage of every volume via the performance monitor for a while. And then set the required policy for each volume.

Case 1: Video Editing

We assume that video editing has the characteristics of focus data over a period of time. When users edit a new video, the video remains at the SSD tier and performs extreme performance. After the editing is complete, the video moves to the HDD tier and leaves the space for the next video. Therefore, we recommend setting the auto-tiering policy to **Start Highest then Auto Tiering**.

Test Equipments and Configurations

- Server
  - Model: ASUS RS700-E6/ERS4 (CPU: Intel Xeon E5620 2.4GHz / RAM: 24GB)
    10GbE HBA: Broadcom BCM57810 NetXtreme || 10 GigE
    OS: Windows Server 2012 R2
- Storage
  - Model: AccuRAID iSCSI 5-Series 3U 16 bay
    Memory: 8GB (1 x 8GB in bank 1) per controller
    Firmware 1.2.1
    SAS SSD: 4 x HGST Ultrastar SSD800MH.B, HUSMH8010BSS200, 100GB, SAS 12Gb/s
    SAS HDD: 4 x HGST Ultrastar C15K600, HUC156030CS4200, 300GB, SAS 12Gb/s
    NL-SAS HDD: 4 x Seagate Constellation ES, ST500NM0001, 500GB, SAS 6Gb/s
  - Auto Tiering Pool: 2.09TB
    SSD Tier: RAID 10 with 4 x SAS SSD, 185GB
    SAS Tier: RAID 6 with 4 x SAS HDD, 558GB
    NL-SAS Tier: RAID 5 with 4 x NL-SAS SSD, 1.36TB
  - Volume: 1 x 2.09TB in Auto Tiering Pool
  - Auto Tiering Policy: **Start Highest then Auto Tiering**
- Simulate Video Files
  - 12 x 100GB files

Test Scenario and Result

1. Create an auto tiering pool with the following configurations.
   - Auto Tiering Pool: 2.09TB
     SSD Tier: RAID 10 with 4 x SAS SSD, 185GB
     SAS Tier: RAID 6 with 4 x SAS HDD, 558GB

NL-SAS Tier: RAID 5 with 4 x NL-SAS SSD, 1.36TB

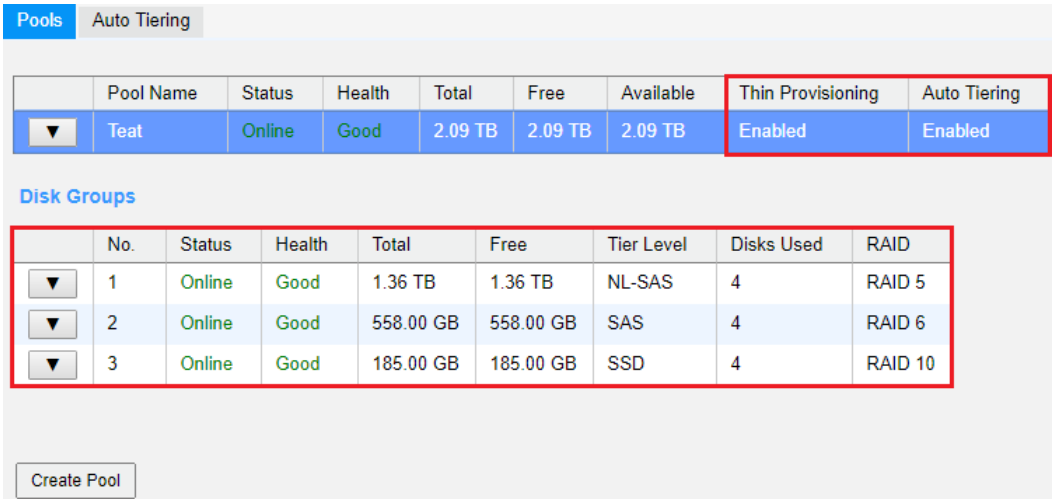

Create an Auto Tiering Pool

2. Create a volume of the capacity 2.09TB, and set the tiering policy as **Start Highest then Auto Tiering**.

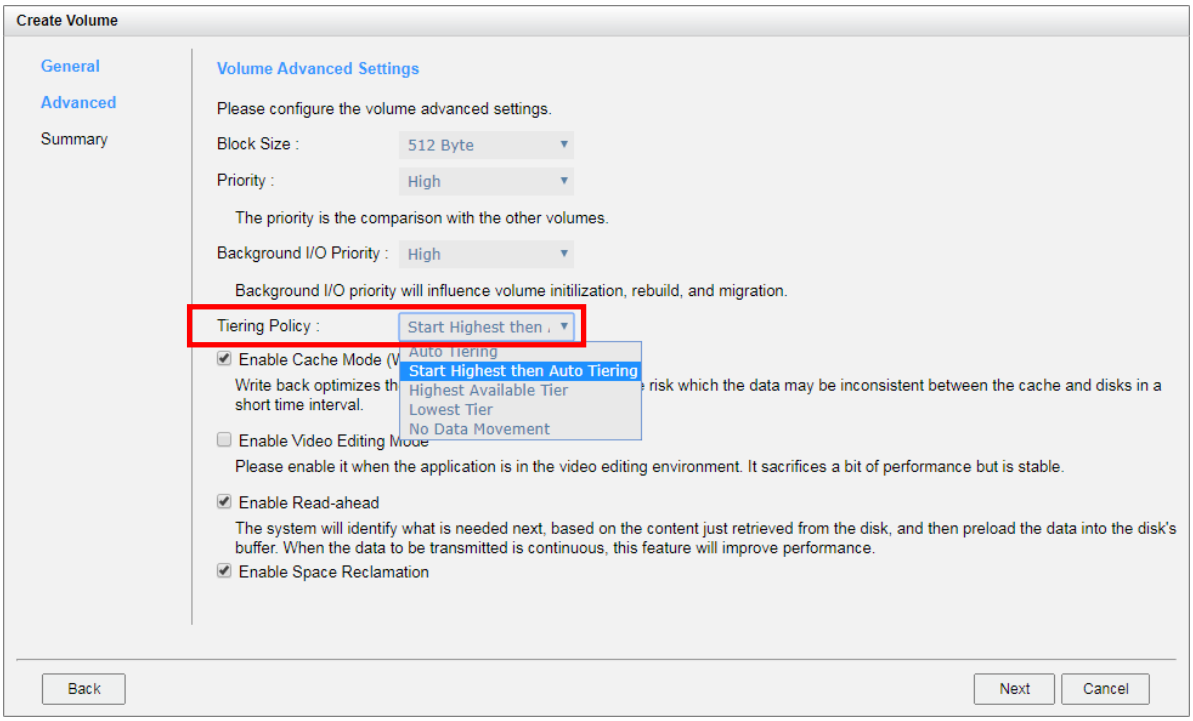

Create a Volume and Set the Tiering Policy as Start Highest then Auto Tiering

3. Copy a 100GB file into the volume. It spends 2 minutes to complete and the transmission speed is around 780 ~ 830 MB/s. The figure shows that the SSD tier is being used.

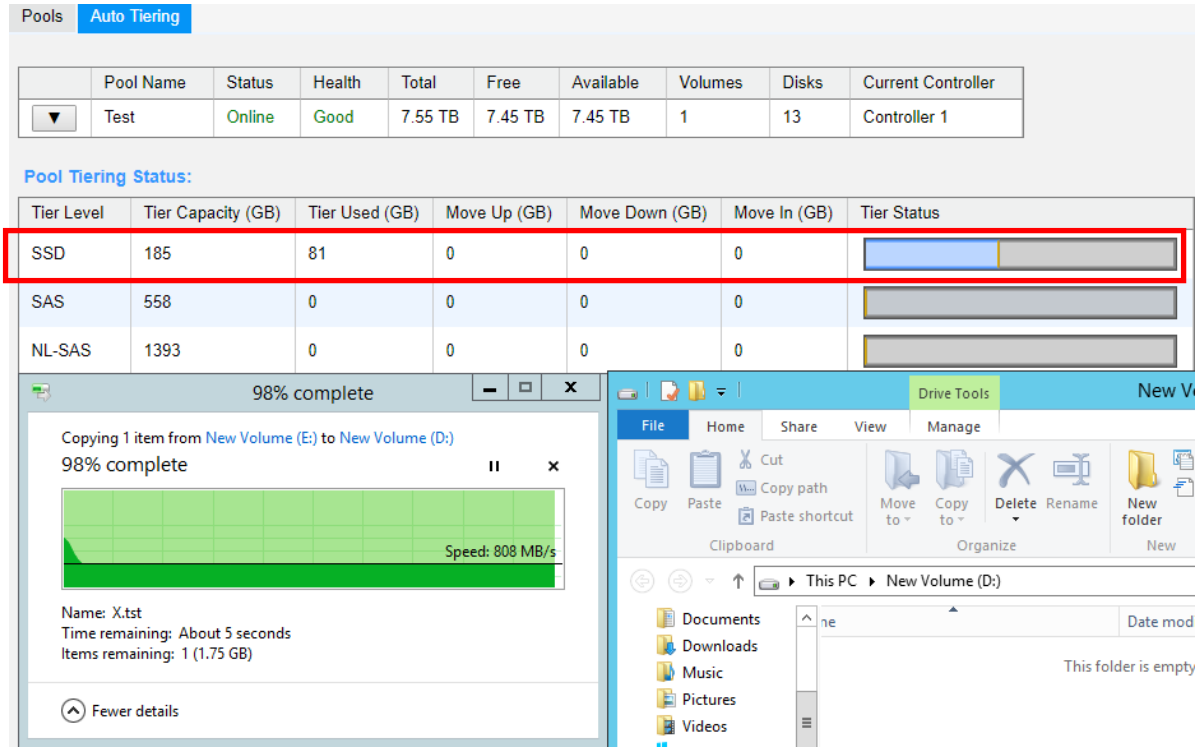

Copy a 100GB File into the Volume

4. The first coming file is located in SSD tier because the tiering policy is set as **Start Highest then Auto Tiering**.

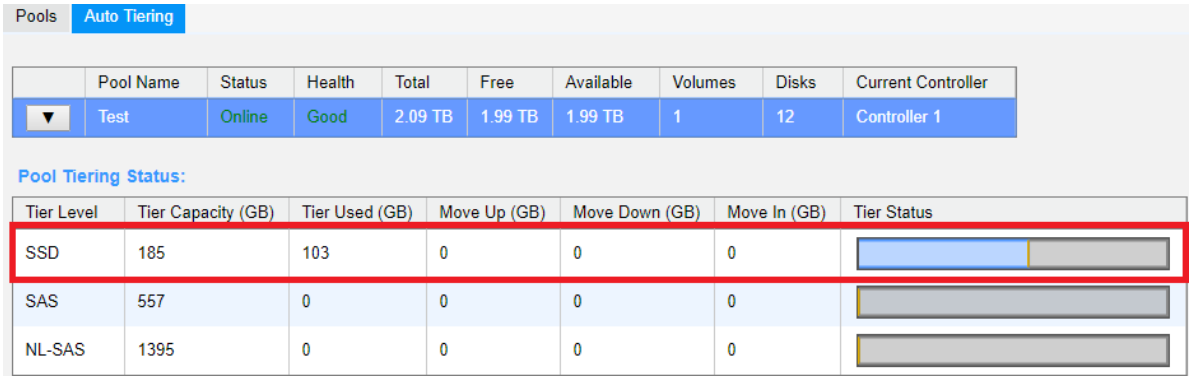

The File is Located in the SSD Tier

5. Copy another 100GB file into the volume. Since the capacity of SSD tier is full, the system will save the data at the next tier. So it spends 2 minutes and 20 seconds to complete. The transmission speed is around 460 ~ 830 MB/s.

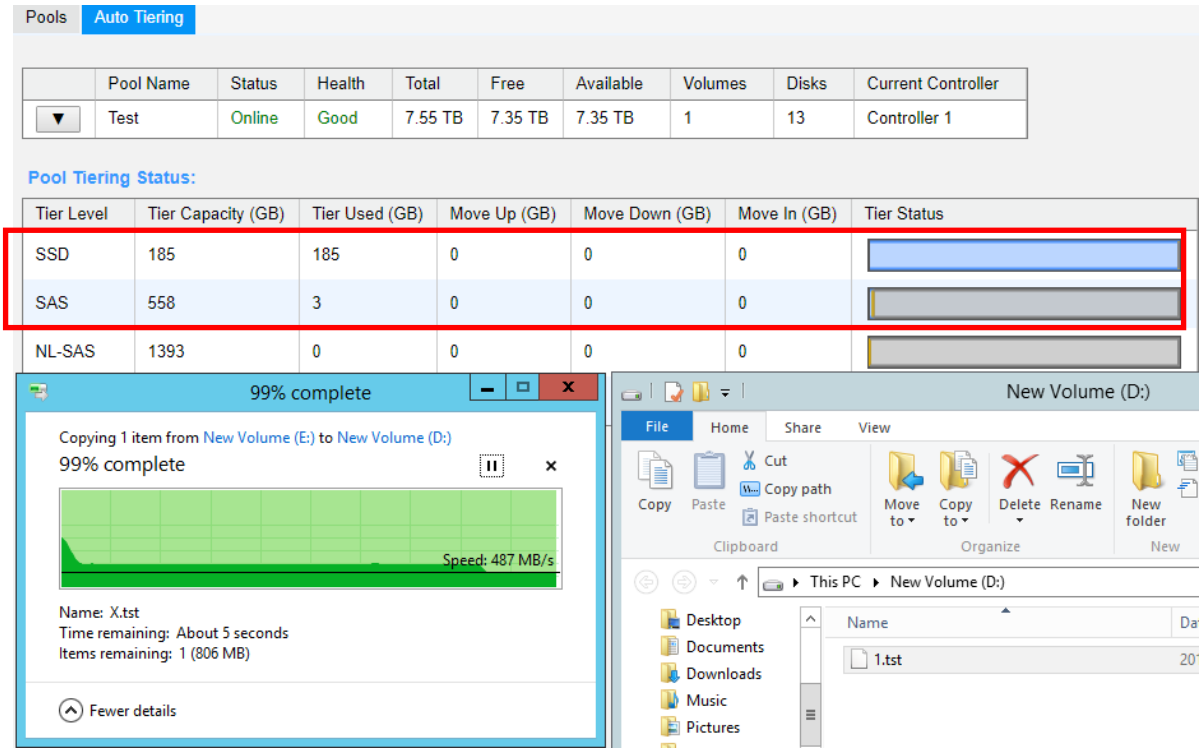

Copy the Second 100GB File into the Volume

6. The second file is distributed in the SSD tier and SAS tier.

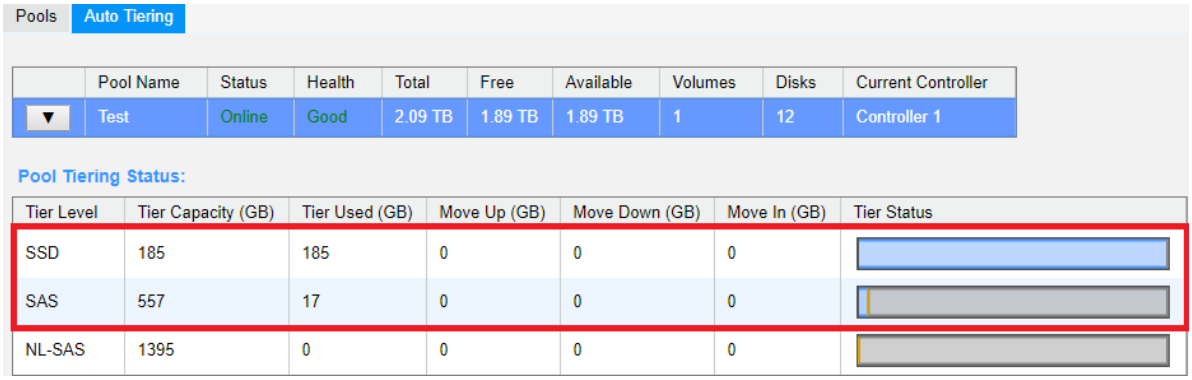

The File is Distributed in the SSD Tier and SAS Tier

7. After an hour, the system analyzes the data automatically, and the data will be relocated at midnight. The figure shows that 18GB data in SSD tier will be moved down to the SAS tier.

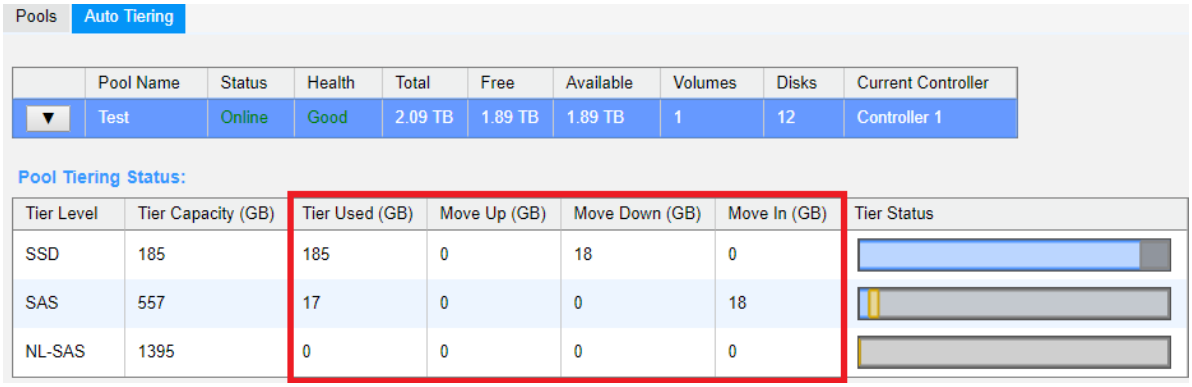

Statistic Collection and Ranking

8. At the next day, 18GB data in SSD tier has been moved down to the SAS tier. And the event log records how much data is moved. You can see that SSD tier reserved about 10% of the capacity for incoming data.

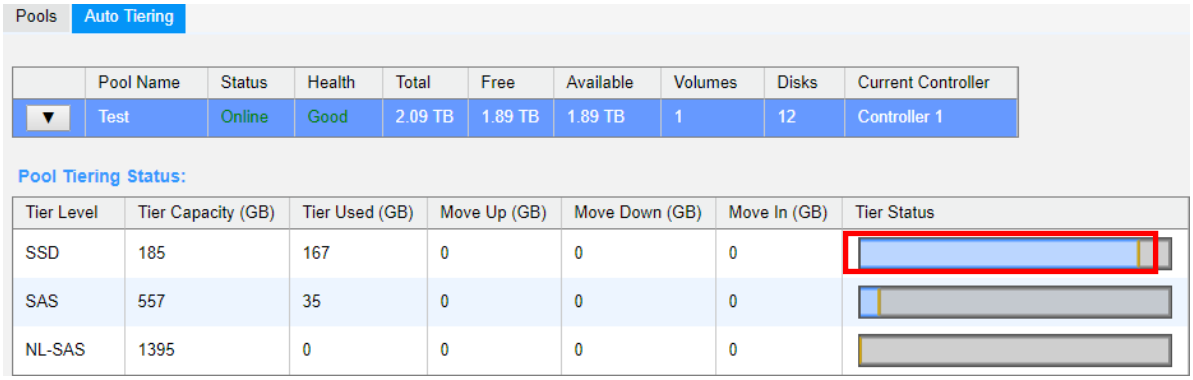

Complete Relocation

9. Continue copying the third 100GB file into the volume. It spends 3 minutes and 8 seconds to complete. The transmission speed is around 460 ~ 500 MB/s. The file is copied to the SAS tier.

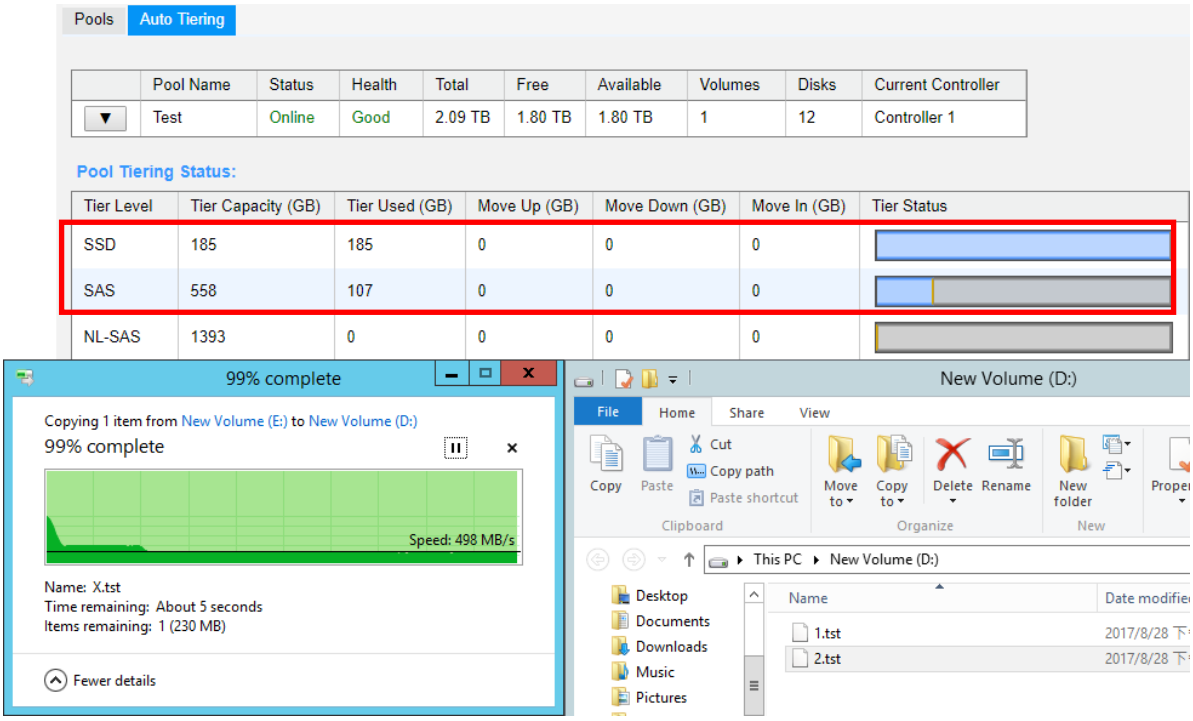

Copy the Third 100GB File into the Volume

10. Again, the system analyzes the data automatically after an hour, and the data will be relocated at midnight. The figure shows that 19GB data in SSD tier will move down to the SAS tier, and 1GB data in SAS tier will move up to the SSD tier.

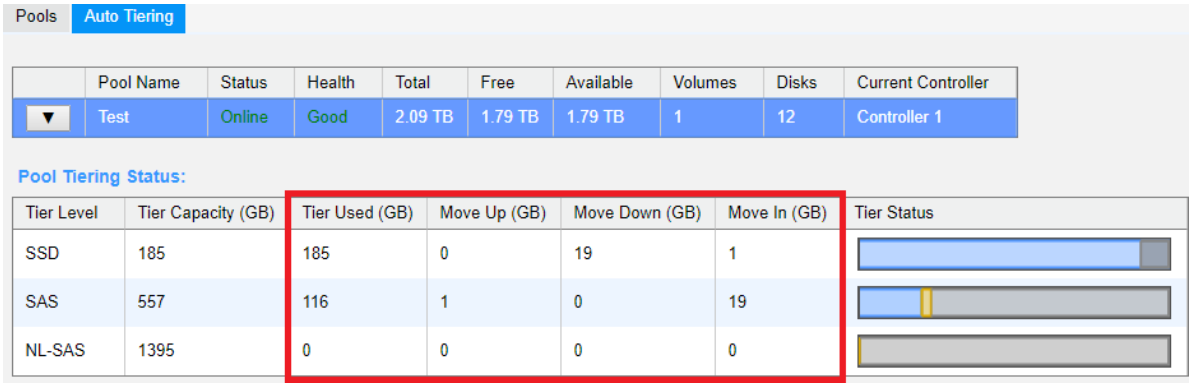

Statistic Collection and Ranking

11. At the next day, the relocation completes.

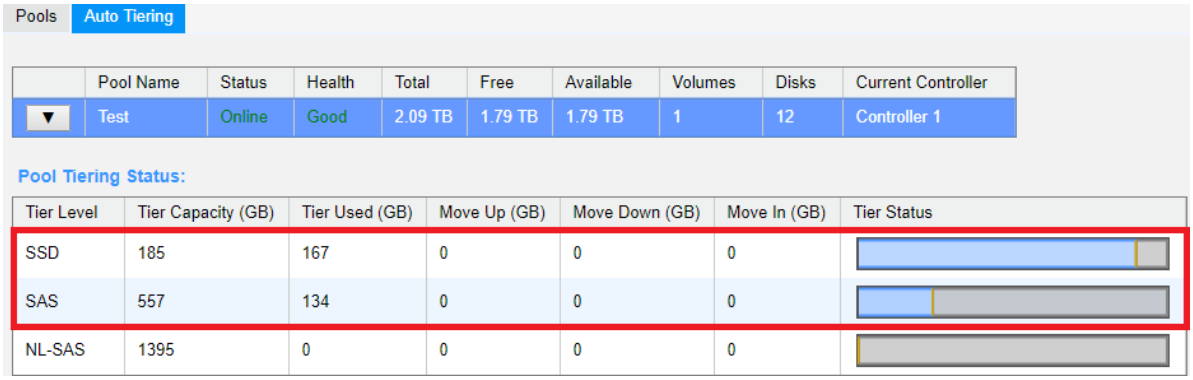

Complete Relocation

12. Repeat several times until SSD tier and SAS tier are full of data. The hot data will be moved up to the higher tier and the cold data will be moved down to the lower tier.

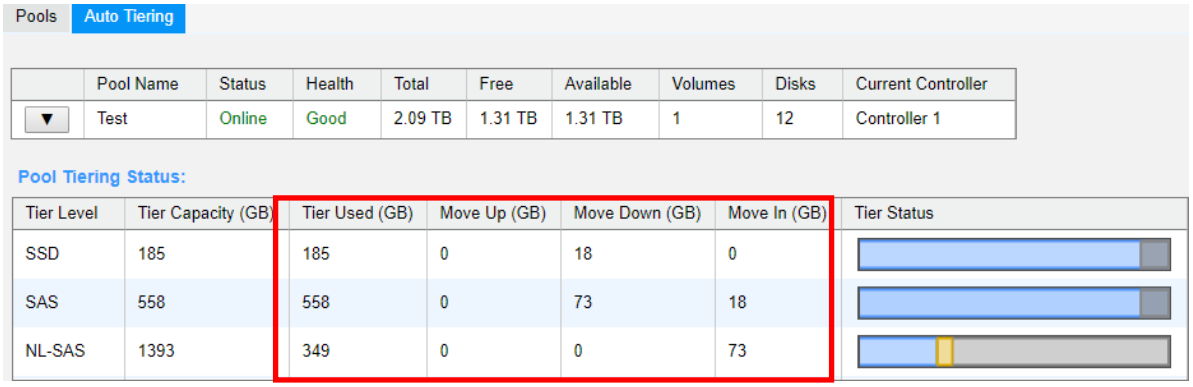

Statistic Collection and Ranking

13. The relocation completes.

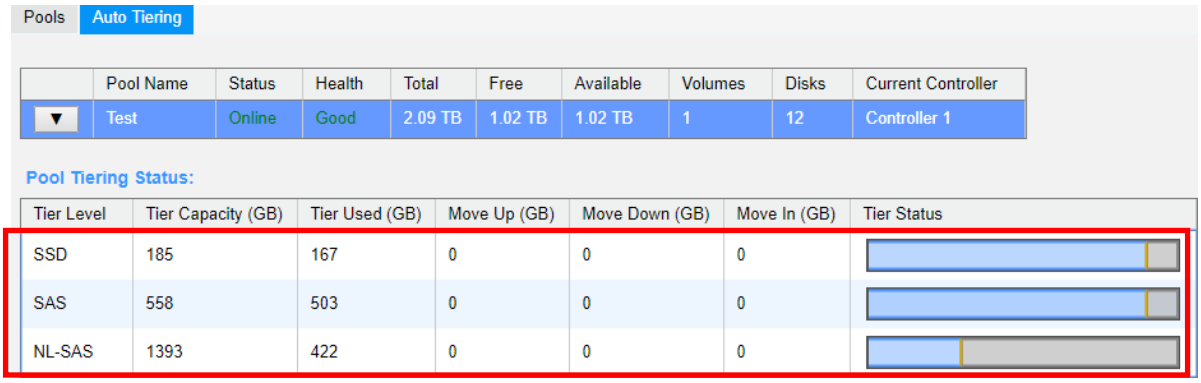

Complete Relocation

14. Last, copy the first file back to the source volume and observe the transmission speed. You can also compare the performance monitor of disks in the web UI and observe which tier the data is located.

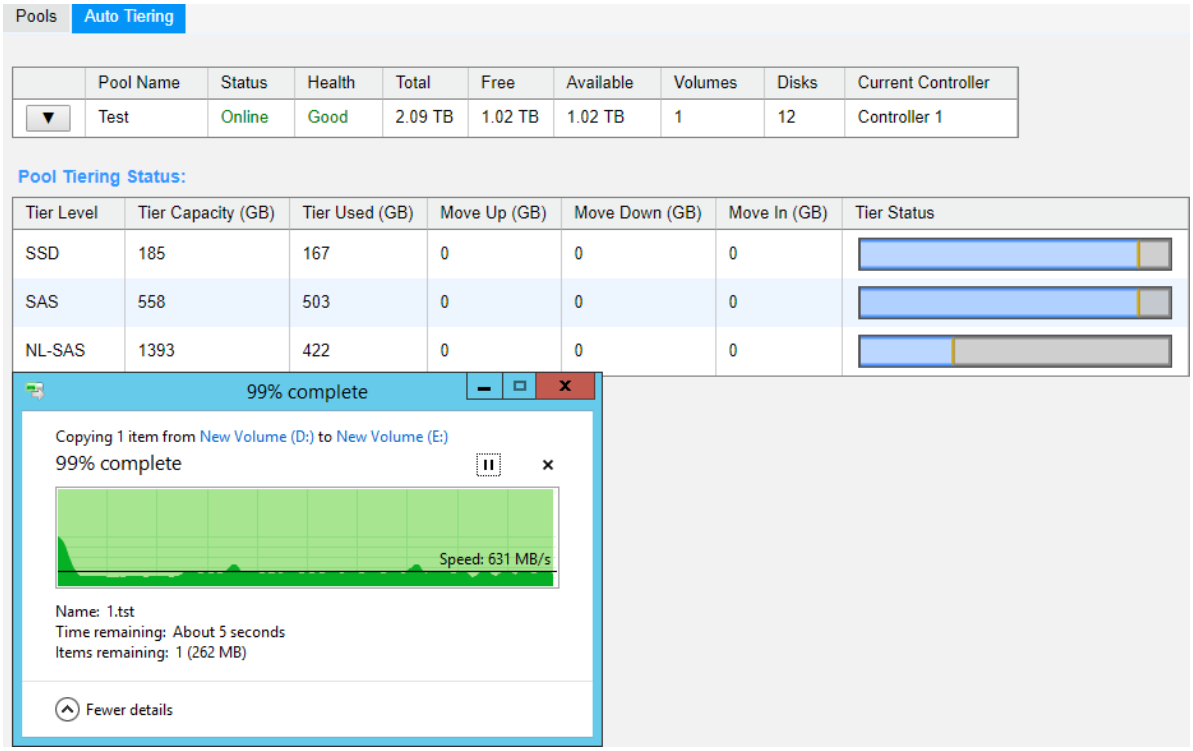

Test the Performance

Summary

In case 1, the data locates at the SSD tier first because we set the auto-tiering policy as **Start Highest then Auto Tiering**. When user edits a new video, the video remains at the SSD tier and performs extreme performance. After the editing is complete, the video moves to the HDD tier and leaves the space for the next video. The scenario meets the expectations of video editing.

Case 2: VMware

We simulate 8 VMs (Virtual Machines) running on a server, assume that they have different I/O queue depths and possess intensive I/O flows. We recommend setting the auto-tiering policy as **Auto Tiering**. After working a while, we assume that the data with heavy I/O will be relocated to the higher tier for better performance.

Test Equipments and Configurations

- Server
  - Model: ASUS RS700-E6/PS4 (CPU: Intel Xeon E2620 2.0GHz / RAM: 20GB)
    10GbE HBA: Intel Ethernet CNA X710-DA4 FH
    OS: VMware ESXi 6.5
- Storage
  - Model: AccuRAID 3-series 4U 24 bay
    Memory: 8GB (2 x 4GB in bank 1 & 3) per controller
    Firmware 1.2.1
    SAS SSD: 4 x HGST Ultrastar SSD800MH.B, HUSMH8010BSS200, 100GB, SAS 12Gb/s
    SAS HDD: 4 x HGST Ultrastar C15K600, HUC156030CS4200, 300GB, SAS 12Gb/s
    NL-SAS HDD: 4 x Seagate Constellation ES.3, ST1000NM0023, 1TB, SAS 6Gb/s
  - Auto Tiering Pool: 3.45TB
    SSD Tier: RAID 10 with 4 x SAS SSD, 185GB
    SAS Tier: RAID 6 with 4 x SAS HDD, 558GB
    NL-SAS Tier: RAID 5 with 4 x NL-SAS SSD, 2.73TB
  - Volume: 1 x 3.45TB in Auto Tiering Pool, 8 x VMs in the Volume
  - Auto Tiering Policy: **Auto Tiering**
- I/O Pattern
  - Tool: IOmeter V1.1.0
  - Workers: 1
  - Access Specifications:
    VM1: 256KB, 100% Write, 100% Random, Outstanding **128**, Maximum Disk Size 10GB
    VM2: 256KB, 100% Write, 100% Random, Outstanding **16**, Maximum Disk Size 20GB
    VM3: 256KB, 100% Write, 100% Random, Outstanding **32**, Maximum Disk Size 10GB
    VM4: 256KB, 100% Write, 100% Random, Outstanding **48**, Maximum Disk Size 20GB
    VM5: 256KB, 100% Write, 100% Random, Outstanding **64**, Maximum Disk Size 10GB
    VM6: 256KB, 100% Write, 100% Random, Outstanding **80**, Maximum Disk Size 20GB

VM7: 256KB, 100% Write, 100% Random, Outstanding **96**, Maximum Disk Size 10GB
VM8: 256KB, 100% Write, 100% Random, Outstanding **112**, Maximum Disk Size 20GB

Test Scenario and Result

1. Create an auto tiering pool with the following configurations.
   - Auto Tiering Pool: 3.45TB
     SSD Tier: RAID 10 with 4 x SAS SSD, 185GB
     SAS Tier: RAID 6 with 4 x SAS HDD, 558GB
     NL-SAS Tier: RAID 5 with 4 x NL-SAS SSD, 2.73TB

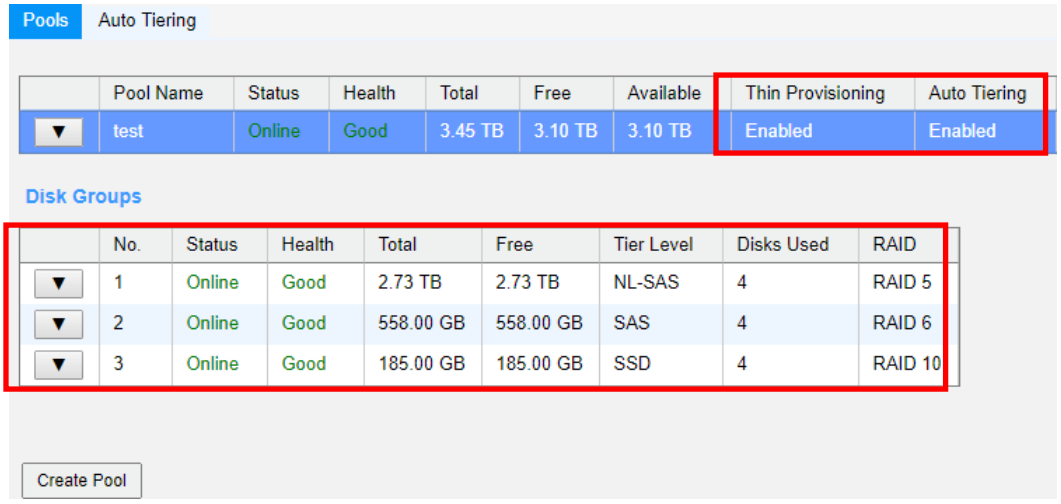

Create an Auto Tiering Pool

2. Create a volume of the capacity 3.45TB, and sets the tiering policy as **Auto Tiering**.

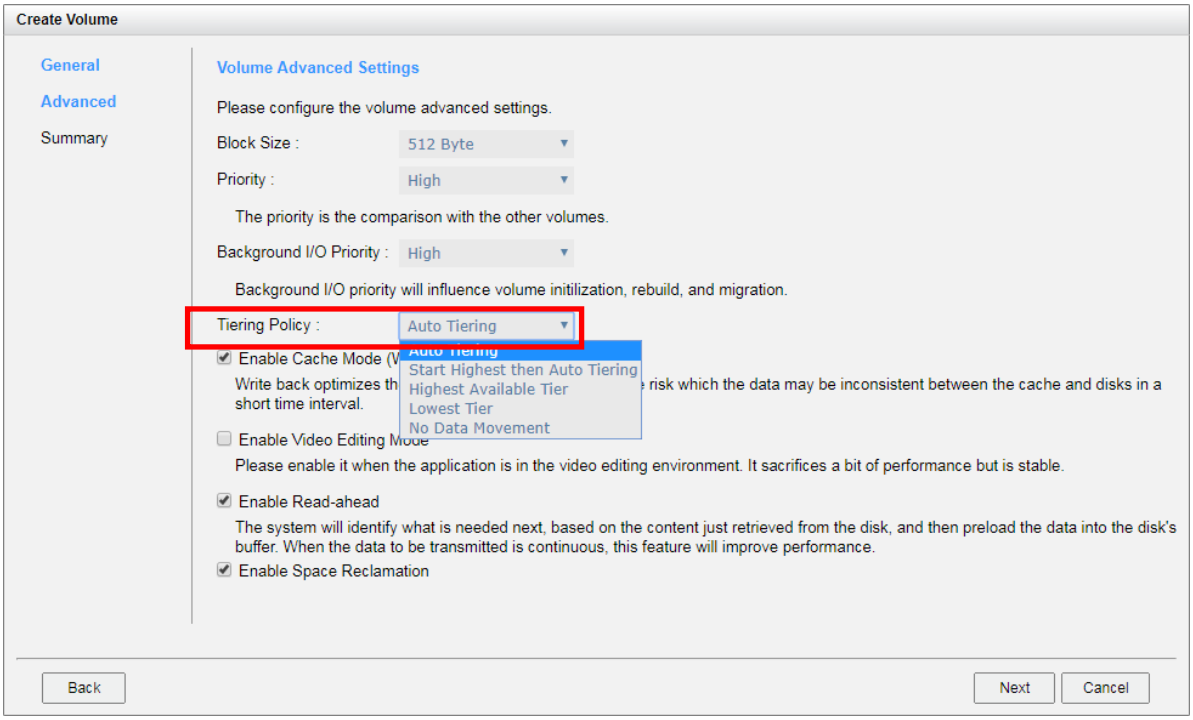

Create a Volume and Set the Tiering Policy as Auto Tiering

3. Create eight VMs and save their datastores in the volume. When they are ready, run IOmeter on each VM to observe the performance. Because the tiering policy is set as **Auto Tiering**, the initial space is allocated in the tier which is healthier and has more free capacity than other tiers. The data is located in the NL-SAS tier with RAID 5.

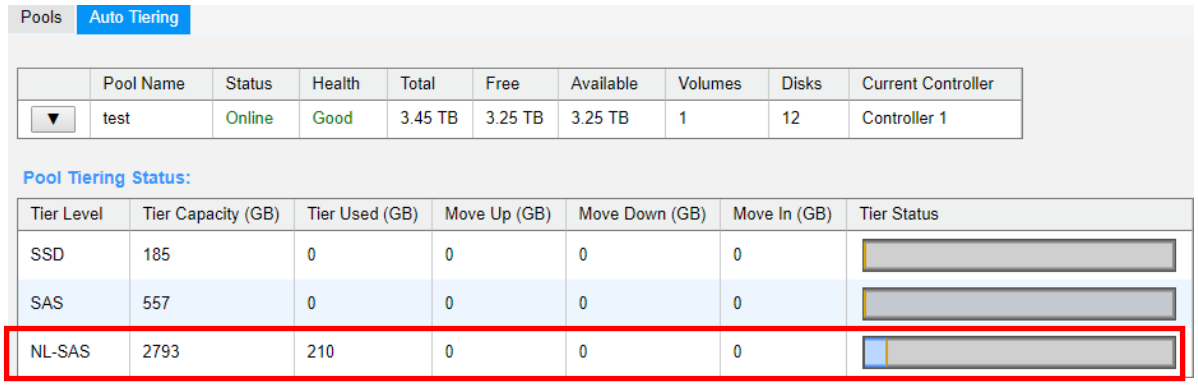

The Data is Located in the NL-SAS Tier

4. The followings are the throughput of VMs running by IOmeter at the beginning.
   - VM1: 256KB, 100% Write, 100% Random, Outstanding **128**, Maximum Disk Size 10GB, the throughput is 9.96 MB/s

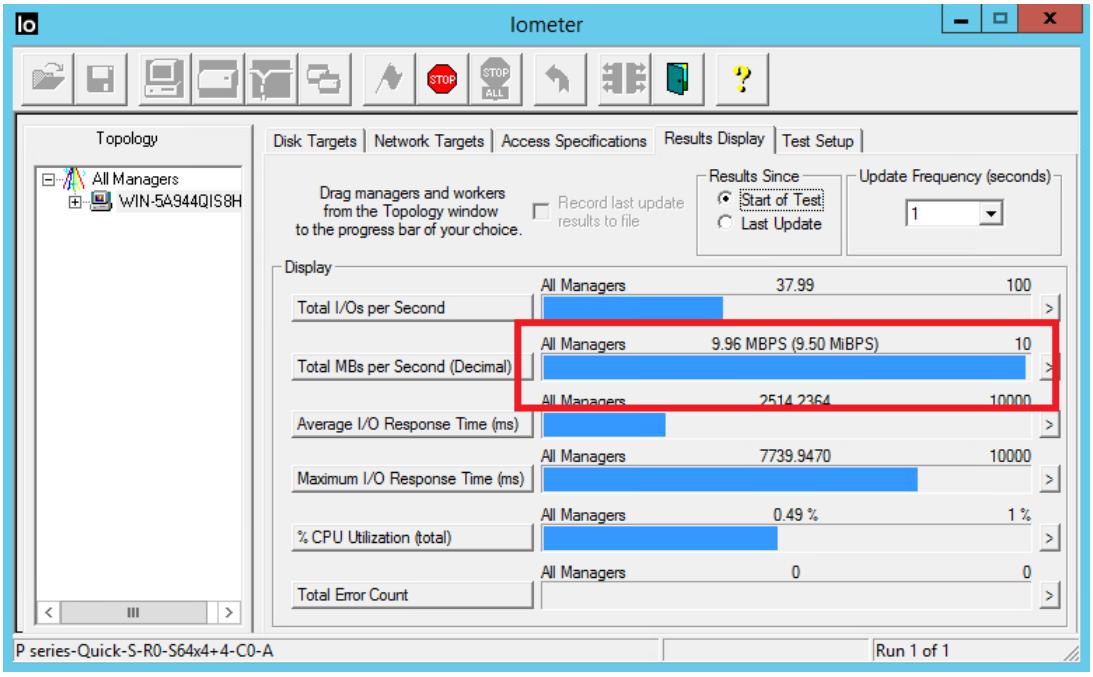

Throughput of VM1 at the Beginning

- VM2: 256KB, 100% Write, 100% Random, Outstanding **16**, Maximum Disk Size 20GB, the throughput is 4.78 MB/s
- VM3: 256KB, 100% Write, 100% Random, Outstanding **32**, Maximum Disk Size 10GB, the throughput is 4.41 MB/s
- VM4: 256KB, 100% Write, 100% Random, Outstanding **48**, Maximum Disk Size 20GB, the throughput is 4.13 MB/s
- VM5: 256KB, 100% Write, 100% Random, Outstanding **64**, Maximum Disk Size 10GB, the throughput is 3.98 MB/s
- VM6: 256KB, 100% Write, 100% Random, Outstanding **80**, Maximum Disk Size 20GB, the throughput is 3.79 MB/s
- VM7: 256KB, 100% Write, 100% Random, Outstanding **96**, Maximum Disk Size 10GB, the throughput is 3.70 MB/s
- VM8: 256KB, 100% Write, 100% Random, Outstanding **112**, Maximum Disk Size 20GB, the throughput is 3.61 MB/s

5. Stop VM2~VM8 I/O but keep VM1 running I/O, the throughput of VM1 is up to 40.98 MB/s.

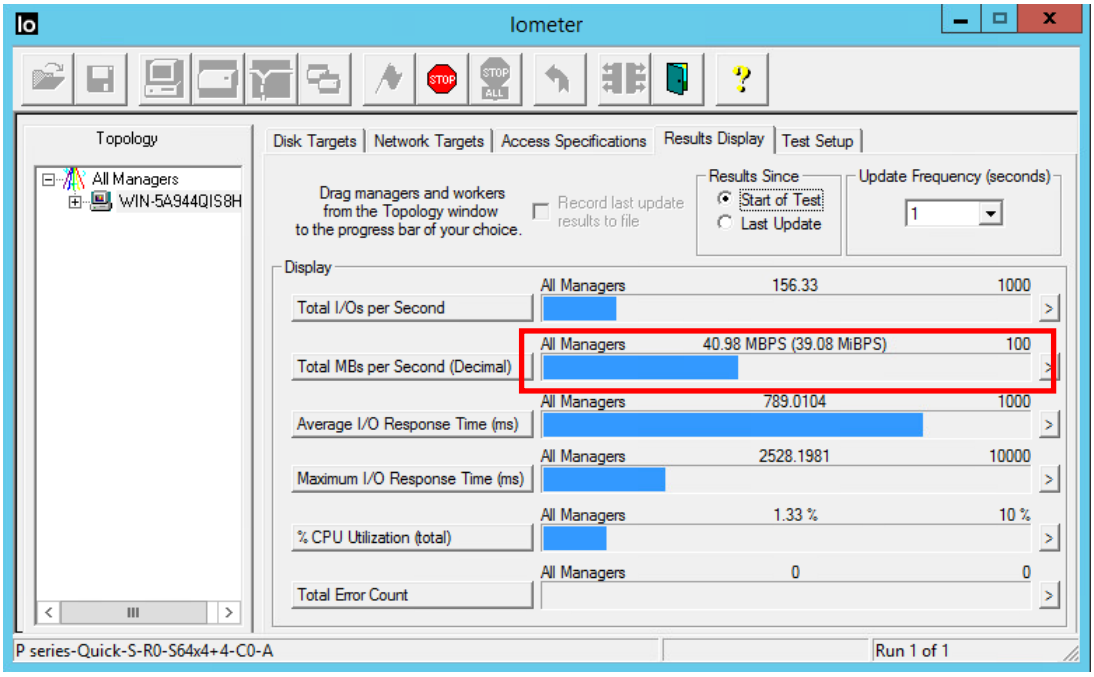

Throughput of VM1 when Stop VM2~VM8 I/O

6. Because VM1 keeps I/O, the data in VM1 will be accessed more frequently than others. After analysis and relocation by auto tiering mechanism, the data in VM1 has been moved to a higher tier. We check the performance of VM1 again; the throughput is up to 465.86 MB/s.

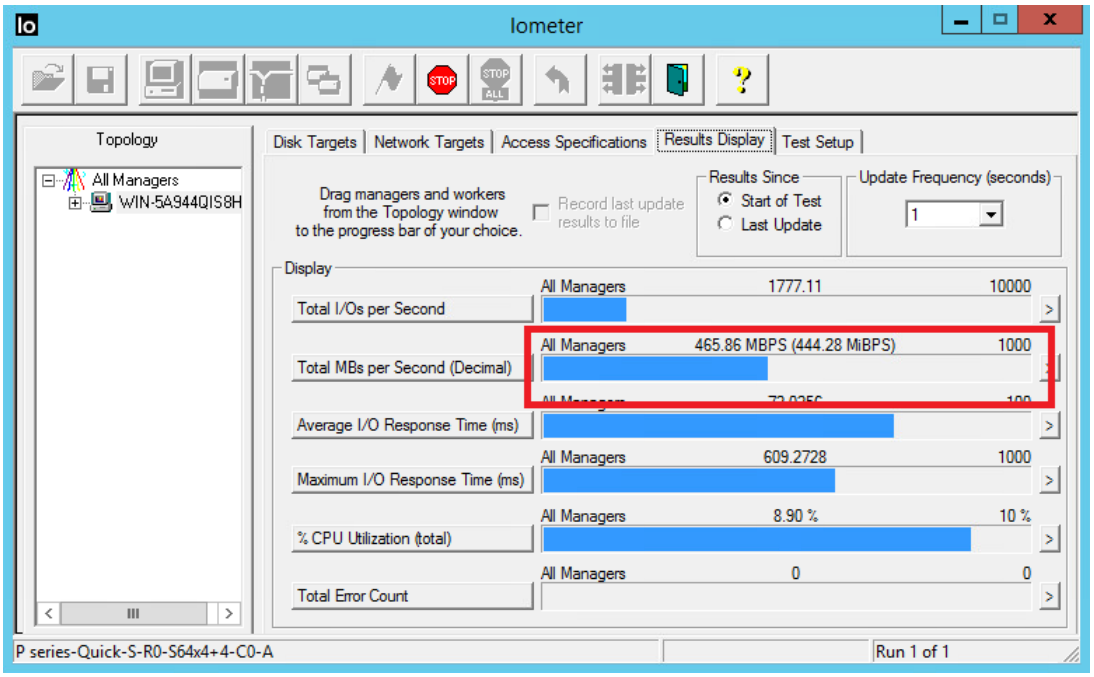

Throughput of VM1 after Analysis and Relocation

7. Run VM2~VM8 I/O again, check performance. The followings are the throughput of VMs running by IOmeter.
    - VM2: 256KB, 100% Write, 100% Random, Outstanding **16**, Maximum Disk Size 20GB, the throughput is 74.75 MB/s

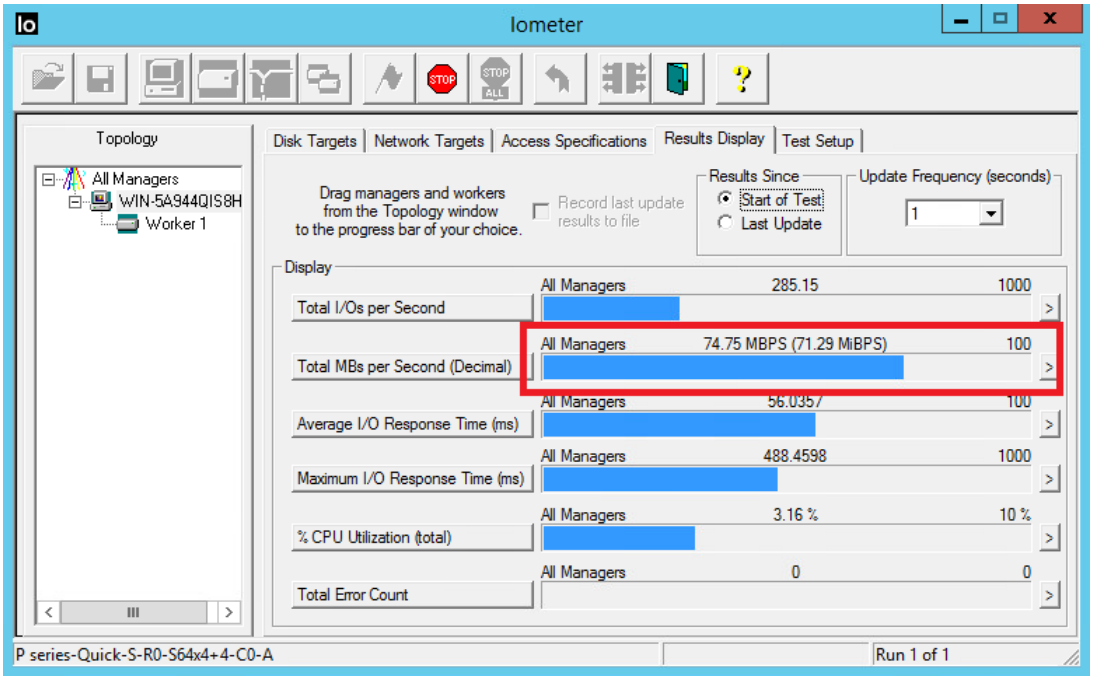

Throughput of VM2

- VM3: 256KB, 100% Write, 100% Random, Outstanding **32**, Maximum Disk Size 10GB, the throughput is 68.78 MB/s
- VM4: 256KB, 100% Write, 100% Random, Outstanding **48**, Maximum Disk Size 20GB, the throughput is 63.59 MB/s
- VM5: 256KB, 100% Write, 100% Random, Outstanding **64**, Maximum Disk Size 10GB, the throughput is 60.03 MB/s
- VM6: 256KB, 100% Write, 100% Random, Outstanding **80**, Maximum Disk Size 20GB, the throughput is 57.12 MB/s
- VM7: 256KB, 100% Write, 100% Random, Outstanding **96**, Maximum Disk Size 10GB, the throughput is 54.90 MB/s
- VM8: 256KB, 100% Write, 100% Random, Outstanding **112**, Maximum Disk Size 20GB, the throughput is 54.18 MB/s

Summary

In case 2, although the auto-tiering policy sets to **Auto Tiering**, the data is allocated in the tier which is healthier and has more free capacity than other tiers at the beginning.

Then the data with frequently accessed I/O will be relocated to the higher tier for better performance. The following table summarizes the throughput before and after the relocation and an improvement percentage calculation as a reference. This verifies the scenario and meets the expectations of VMware.

Summarize the Throughput Before and After the Relocation

VM Name	Throughput Before Relocation	Throughput After Relocation	Improved
VM1	9.96 MB/s	465.86 MB/s	4,577%
VM2	4.78 MB/s	74.75 MB/s	1,464%
VM3	4.41 MB/s	68.78 MB/s	1,460%
VM4	4.13 MB/s	63.59 MB/s	1,440%
VM5	3.98 MB/s	60.03 MB/s	1,408%
VM6	3.79 MB/s	57.12 MB/s	1,407%
VM7	3.70 MB/s	54.90 MB/s	1,384%
VM8	3.61 MB/s	54.18 MB/s	1,401%

Case 3: Sudden Reaction

In order to cope with an expected sudden event, IT administrators can move the required data to the SSD tier in advance. In general, we recommend setting the auto-tiering policy to **Lowest Tier**. The day before the activity, IT administrator manually set the volume containing the required data to **Highest Available Tier** and then performs **Relocation Now** manually to force relocating data.

Test Equipments and Configurations

- Server
  - Model: ASUS RS700-E6/ERS4 (CPU: Intel Xeon E5620 2.4GHz / RAM: 24GB)
    10GbE HBA: Intel Ethernet CNA X710-DA4 FH
    OS: Windows Server 2012 R2
- Storage
  - Model: AccuRAID 5-Series 3U 16 bay
    Memory: 16GB (2 x 8GB in bank 1 & 3) per controller
    Firmware 1.2.1

SAS SSD: 4 x HGST Ultrastar SSD800MH.B, HUSMH8010BSS200, 100GB, SAS 12Gb/s
SAS HDD: 4 x HGST Ultrastar C15K600, HUC156030CS4200, 300GB, SAS 12Gb/s
NL-SAS HDD: 4 x Seagate Constellation ES, ST500NM0001, 500GB, SAS 6Gb/s

- Auto Tiering Pool: 2.09TB
  SSD Tier: RAID 10 with 4 x SAS SSD, 185GB
  SAS Tier: RAID 6 with 4 x SAS HDD, 558GB
  NL-SAS Tier: RAID 5 with 4 x NL-SAS SSD, 1.36TB
- Volume: 1 x 2.09TB in Auto Tiering Pool
- Auto Tiering Policy: **Lowest Tier** then **Highest Available Tier**

- I/O Pattern
  - Tool: IOmeter V1.1.0
  - Workers: 1
  - Outstanding (Queue Depth): 128
  - Maximum Disk Size: 50GB
  - **Access Specifications: 4KB, 100% Write, 100% Random**

Test Scenario and Result

1. Create an auto tiering pool with the following configurations.
   - Auto Tiering Pool: 2.09TB
     SSD Tier: RAID 10 with 4 x SAS SSD, 185GB
     SAS Tier: RAID 6 with 4 x SAS HDD, 558GB
     NL-SAS Tier: RAID 5 with 4 x NL-SAS SSD, 1.36TB

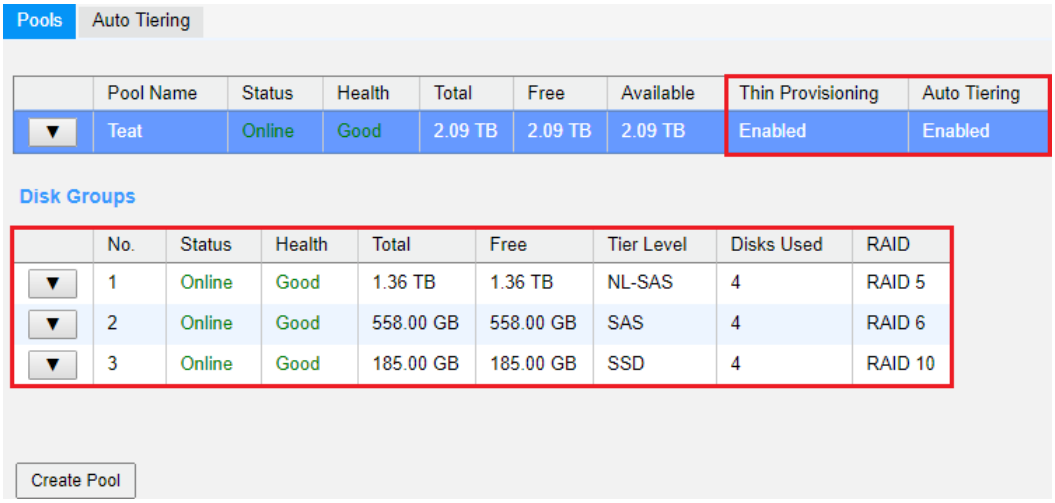

Create an Auto Tiering Pool

2. Create a volume of the capacity 2.09TB, and the tiering policy sets as **Lowest Tier**.

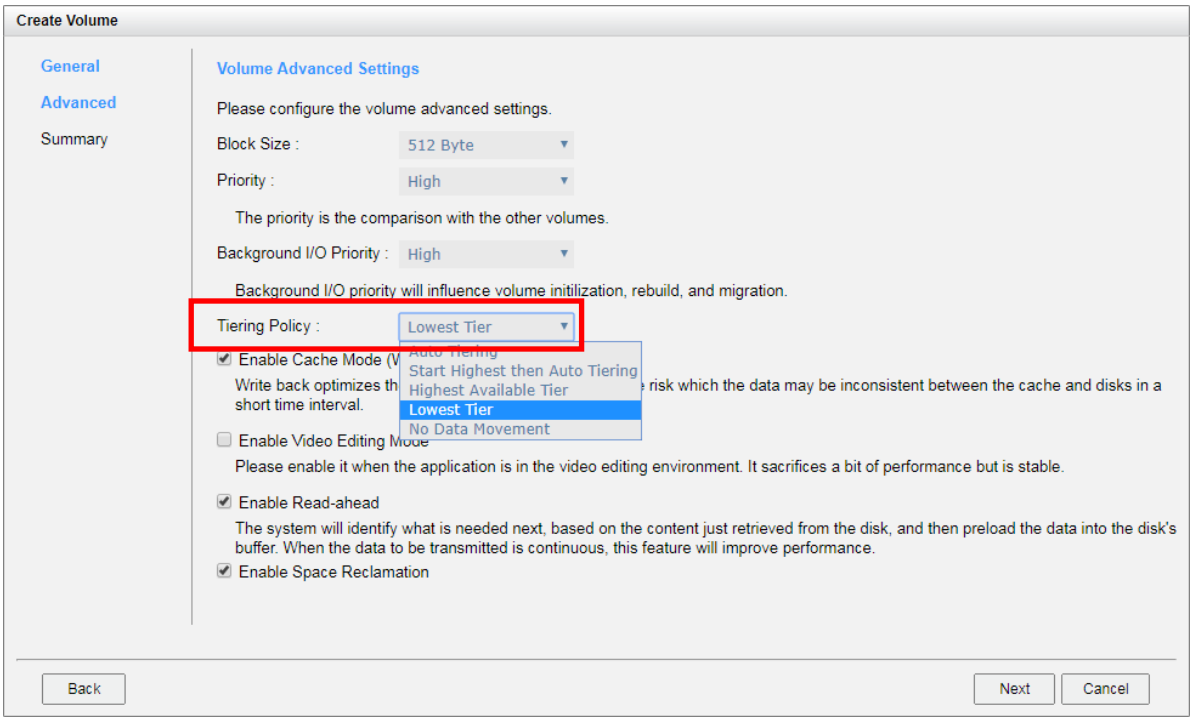

Create a Volume and Set the Tiering Policy as Lowest Tier

3. Run IOmeter to observe the performance. IOmeter parameters are on the following.
   - Tool: IOmeter V1.1.0
   - Workers: 1
   - Outstanding (Queue Depth): 128
   - Maximum Disk Size: 50GB
   - **Access Specifications: 4KB, 100% Write, 100% Random**

   Because the tiering policy sets as **Lowest Tier**, the I/O file is located in the NL-SAS tier, and the IOPS is 341.28.

Pools | Auto Tiering

	Pool Name	Status	Health	Total	Free	Available	Volumes	Disks	Current Controller
▼	Teat	Online	Good	2.09 TB	2.04 TB	2.04 TB	1	12	Controller 1

Pool Tiering Status:

Tier Level	Tier Capacity (GB)	Tier Used (GB)	Move Up (GB)	Move Down (GB)	Move In (GB)	Tier Status
SSD	185	0	0	0	0	
SAS	557	0	0	0	0	
NL-SAS	1395	52	0	0	0	

The I/O File is Located in the NL-SAS Tier

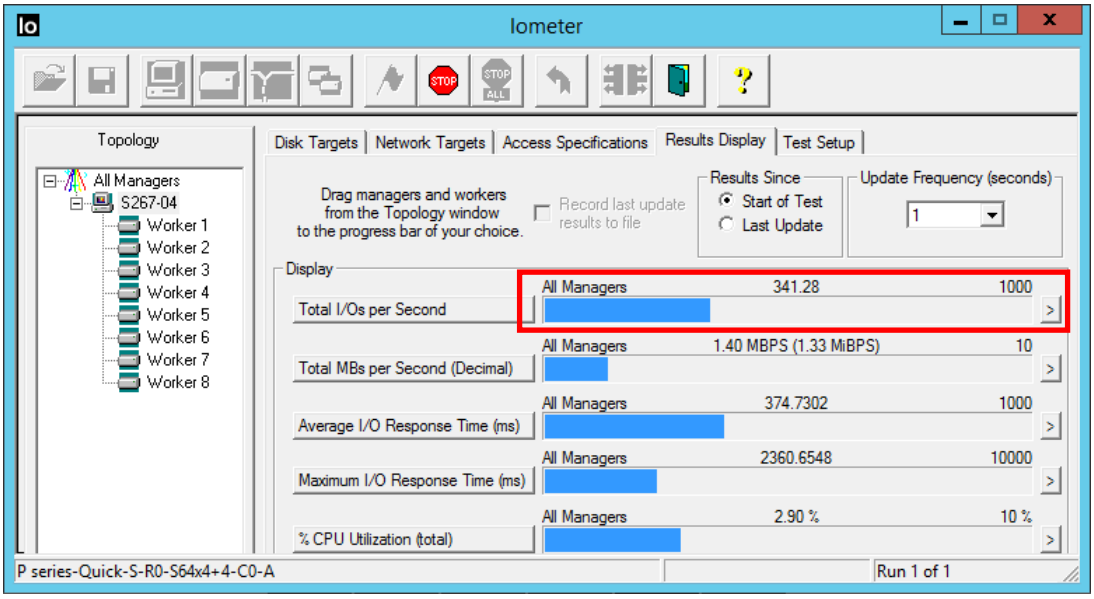

IPOS of the Volume

4. Assume that the data in this volume will be used frequently tomorrow; manually change the tiering policy to **Highest Available Tier**.

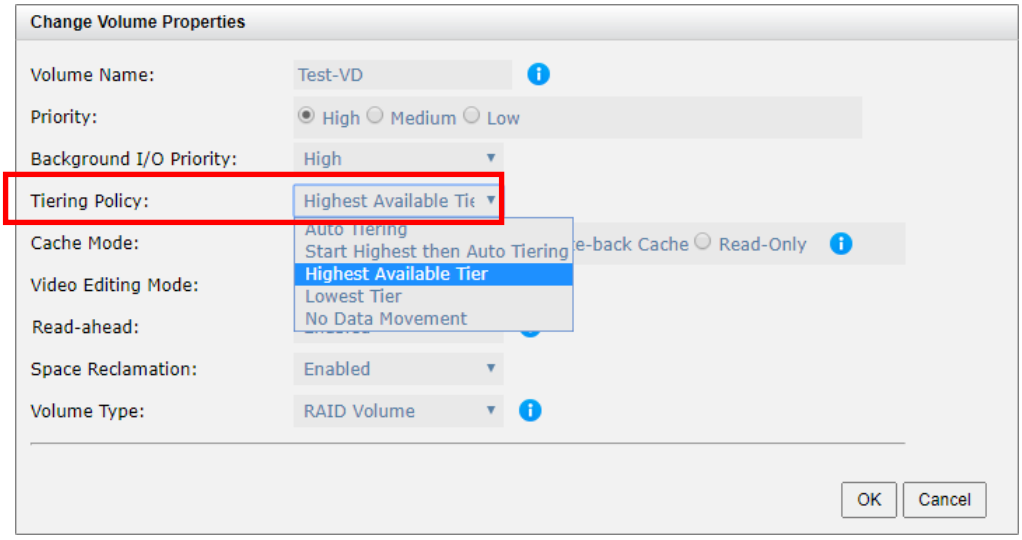

Change the Tiering Policy to Highest Available Tier

5. After an hour, the system analyzes the data automatically, and it will be relocated at midnight or manually execute relocation via the function **Relocation Now**. You can also set the relocation rate as **Medium** or **Slow** to eliminate the possible performance impact. The figure shows that 52GB data in NL-SAS tier will be moved up to the SSD tier.

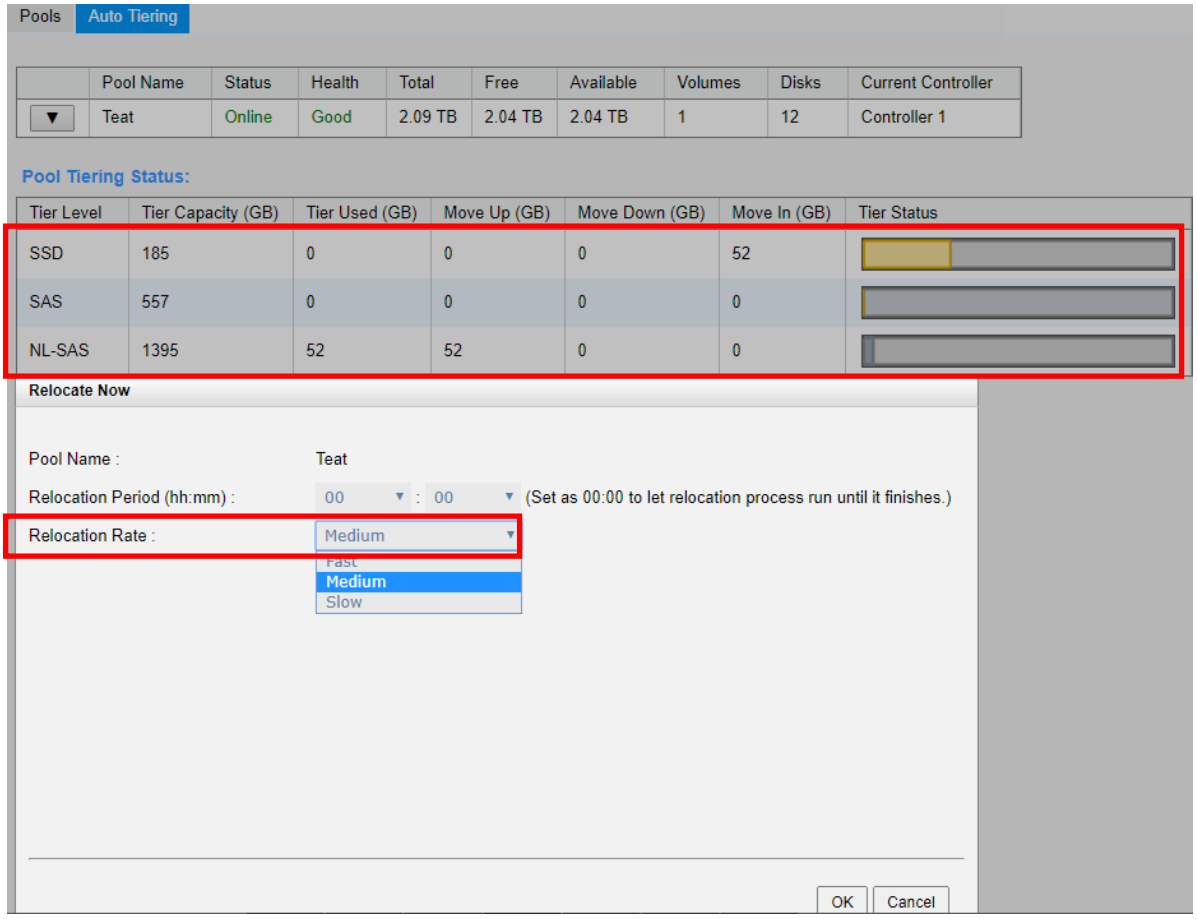

Execute Relocation Now Manually

6. The relocation completes. The data has been moved to the SSD tier.

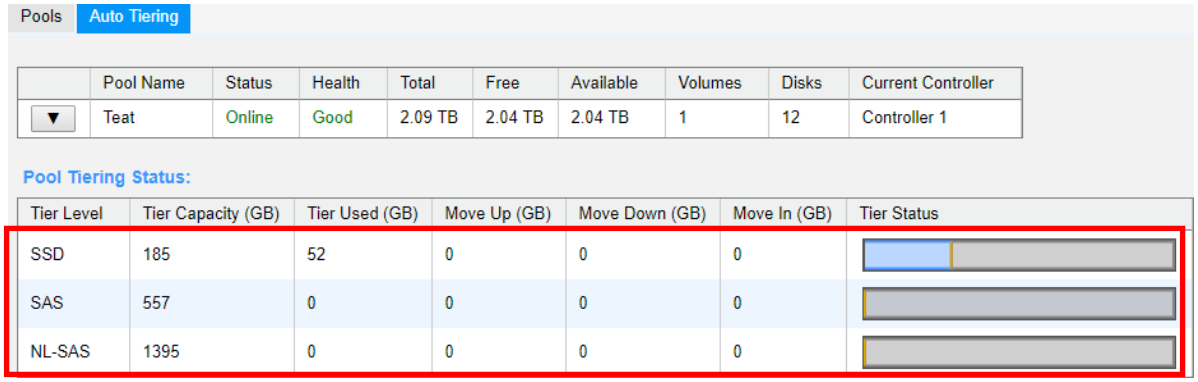

Complete Relocation

7. The IOPS of this volume increases to 44170.28.

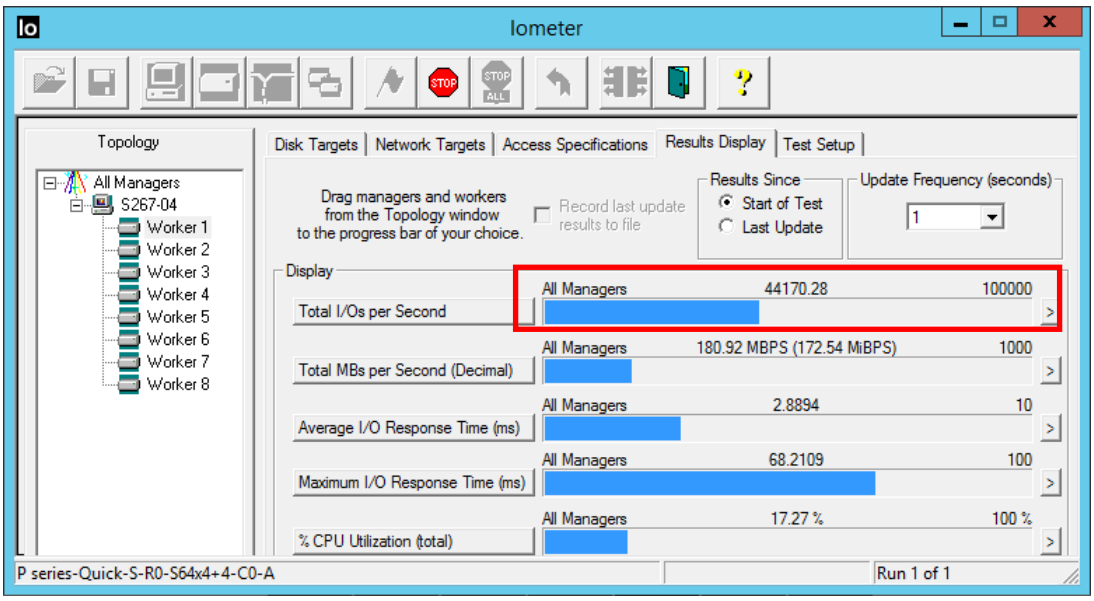

IPOS of the Volume after Relocation

Summary

In case 3, IT administrator can manually control the data into the higher or lower tier in advance. The scenario meets the expectations of an expected sudden event.

Auto Tiering Notes

There are some notices about auto tiering.

- In our design, the snapshot data will be located at the lowest tier in order to obtain economic benefits, and retain the highest space for performance usage. If an auto tiering pool enables snapshots, the performance may be limited to the HDDs at the lowest tier.
- If using SATA SSDs in dual controller system, the performance of each SSD is limited to 270MB/s per SSD due to the MUX board.
- In the SSD Cache vs. Auto Tiering section, we know that the effectiveness of SSD cache can be seen in a short term, and auto tiering is effective in a long term. Both functions can be used at the same time and achieve complementary effects. Be notice that the quantity and the capacity of SSDs which SSD cache and auto tiering use, and IT administrator should adjust via the performance monitor at any time to get better.

Conclusion

With auto tiering technology, the AccuRAID iSCSI series can help you put the right data in the right place at the right time for optimal use of all storage tiers and allow you to reduce storage costs and management overhead while increasing performance and capacity.

Intelligent algorithm behind auto tiering manages the data relocation and monitors the data hotness ratio using half-life coefficient and advanced ranking mathematics. Relocations can occur on the user-defined relocation schedule, making auto tiering a truly automated offering.