

Storage Technologies

IT403

Prepared by
Dr. Ahmed El-Douh

Lecture 7



Storage Technologies

Lecture 7

IT403

Common Advanced Storage Technologies

Prepared by
Dr. Ahmed El-Douh

Objectives

Upon completion of this module, you will be able to:

- Understand the principles of advanced storage optimization technologies such as SmartThin and SmartTier.
- Describe the mechanisms of Smart QoS and SmartPartition for performance control and resource isolation.
- Explain the role of Snapshot in data protection and recovery scenarios.
- Understand how SmartQuota manages and allocates storage resources effectively.

Outlines


- SmartThin
- SmartTier
- Smart QoS
- SmartPartition
- Snapshot
- SmartQuota

Chapter
07

SmartThin




Overview of Thin Provisioning



Thin LUN

- It is a logical disk accessible by hosts.
- It is referred to as Thin LUN and its storage space are dynamically allocated from the storage pool based on actual user storage capacity in use.


Allocated capacity equals to the actual user capacity in use.



Main Features of SmartThin


- The main features of SmartThin are as following:
 - Supports Thin LUN capacity virtualization. SmartThin allows hosts to detect the storage space larger than the actual storage space consumed by the Thin LUN.
 - Supports resource allocation on write. SmartThin allows the dynamic allocation of actual storage resource to the Thin LUN when the host is writing data. The amount of storage space allocated is the same as the amount of data written.
 - Supports online expansion of Thin LUN. SmartThin allows 2 types of online expansion method which are indirect storage expansion by the storage pool and the direct storage expansion on the Thin LUN.
 - Supports Thin LUN space reclamation. SmartThin supports 2 types of space reclamation methods which are standard SCSI space reclamation command and zero data space release reclamation method.

Thin LUN Capacity Virtualization




Capacity - on - write

- Write I/O request to the Thin LUN will trigger the storage capacity allocation.
- When the currently allocated physical storage is lower than the threshold, it will apply for new storage space from the storage pool.



Direct - on - time

- Due to the adoption of the capacity-on-write technology, the actual storage area for the data is not fixed, thus when there is a read/write I/O request for the Thin LUN, it need to be directed on time to the actual storage area.



Application Scenario


- For core business services that has high requirements for business continuity, using SmartThin configuration allows online expansion of the system without interruption to business service. For example: bank transaction systems.
- For businesses that has high application data growth speed that is unable to be accurately evaluated, using SmartThin configuration allows on demand allocation of physical storage space and avoid wastage in utilization rate. For example: Email services and Web Disk services.
- For business with complex systems and has different requirements for storage for different services and applications, using SmartThin configuration allows different business services to compete for the physical storage space, allowing optimum configuration for physical storage space. For example: Carrier Service Providers.

Chapter
07

SmartTier




Overview of Storage Tiers




Levels of Storage Tier

Storage Tier	Disk Type	Disk Characteristics	Data Characteristics
High Performance	SSD	SSD has very short response time, each unit of storage capacity has high costs.	High Performance Tier is suitable for data with high frequency of access.
Performance	SAS	SAS disks has short response time, and each unit of storage capacity has medium costs.	Performance Tier is suitable for data with medium frequency of access.
Capacity	NL-SAS	NL-SAS disks has relatively longer response time, and each unit of storage capacity has low cost.	Capacity Tier is suitable for data with low frequency of access.


Working Principles



Key Technologies



Initial Capacity Allocation




Automatic Allocation:

Allocation sequence is from Performance Tier, Capacity Tier and to High Performance Tier. Only when there is no free space in the previous tier then the allocation to the next tier occurs.

Allocation Based on Specified Storage Tier:

Allocation preference is based on the specified storage tier. If the space is insufficient, then it will allocate from the next tier. The allocation sequence is: Performance Tier, Capacity Tier, High Performance Tier.



Migration Policies

No Migration

No migration operation on the data in the LUN.

Migrate Towards Higher Performance Tier

The LUN data will migrate towards higher performance tier.


Migrate Towards Lower Performance Tier

The LUN data will migrate towards lower performance tier.

Automatic Migration


LUN data will migrate based on data activity level.

I/O Monitoring and Analysis



- Based on the I/O statistics information during the monitoring period of the monitored storage pool, such as Read/Write I/O, Average I/O size, the I/O distribution in the storage pool can be determined periodically. Thus, it can be used to determine the data activity level (hot/cold) and then figure out the data blocks to be migrated through analysis.

Data Migration



HPT : High Performance Tier
PT : Performance Tier
CT : Capacity Tier


- In a multi controller environment, SmartTier will link all the controllers together for an unified data analysis to determine the data activity level, and perform cross controller data migration into corresponding storage tiers.

Chapter
07

SmartQoS



Overview



Working Principles of SmartQoS

- **I/O Priority Scheduling:**

- Classifies services by importance to define response priority. High-priority services receive resources first, especially when system resources are limited. Supports three levels: High, Medium, and Low.


- **I/O Flow Control:**

- Uses the token bucket mechanism to regulate data flow based on performance metrics such as IOPS and bandwidth. Prevents heavy-load applications from affecting other services.


- **I/O Performance Assurance:**

- Uses weighted scheduling to guarantee minimum performance levels such as minimum IOPS, minimum bandwidth, and maximum latency for high-priority services, ensuring they meet required thresholds as much as possible.

I/O Priority Scheduling




I/O Flow Control



Application Scenario - Ensure Service Performance of High Priority Users

- To satisfy the high priority users, we can configure the I/O priority level on the LUN to ensure the operations of the service.
- Configure user A with higher priority, to ensure that the services for user A will operate as normal.
- Configure user B with lower priority, to ensure that services for user B will not impact the normal operations of user A's services.



Chapter
07


SmartPartition



Overview of SmartPartition

- **SmartPartition** is a cache partitioning technology designed to guarantee performance for key applications by allocating dedicated cache resources. Administrators can assign different cache sizes to different applications, ensuring exclusive access to cache within each partition and dynamic resource adjustment based on real-time needs.
- It follows the industry-standard cache partitioning concept, similar to solutions from vendors like EMC and HDS, where cache is divided into logical regions. The main differences between vendors lie in their internal algorithms and adjustment strategies.


The Storage Requirements of the Converging Trend of IT Architecture



Working Principles

- **Cache partitioning technologies** ensure the service quality of key services through isolating and partitioning the required cache resources for the different services.
- In storage systems, the amount of cache that can be used by one service is the main factor of whether it can impact the service quality of other services.
- The size of the cache that a service can occupy on the storage system is the most important factor that affects the performance of the storage system. A large cache capacity translates into:
 - Higher write combination rate, higher write hit ratio, and better disk access sequence for an application that delivers write I/Os.
 - Higher read hit ratio for an application that delivers read I/Os.
 - For an application that delivers sequential I/Os, its cache capacity does not need to be large but must meet I/O requirements.
 - Higher access rate and better performance for an application that delivers random I/Os.

Scheduling of Multiple Applications



Application Scenario - Ensuring Core Service Performance In Multi - Service Scenario

- The service characteristics of storage system that runs both Production System and Test System:

Service Type	Service Characteristic
Production System	Frequent Read and Write I/O
Test System	Frequent Read I/O, Lesser Write I/O

- Based on the Read/Write I/O frequency of the production and testing system, configure a reasonable read and write cache, which will improve the I/O performance of the Production System while at the same time not affect the normal operation of the Test System.
 - Create SmartPartition1 for Production System (e.g. 20GB Read Cache, 10GB Write Cache), which fully satisfies the frequent Read/Write I/O requirements of the Production System.
 - Create SmartPartition2 for Testing System (e.g. 15GB Read Cache, 8GB Write Cache), which has smaller cache capacity that ensures normal operation of the Test System while not affecting the performance of Production System.


Chapter
07

Snapshot



LUN Snapshot Overview

- **Definition:** **Snapshot** refers to consistent data copies from the source data at a certain point of time. After the snapshot is generated, it can be accessed by the host or used as a backup of data at a certain point of time.
- **The main features of Snapshot includes:**
 - **Instant Generation:** Storage system can generate a snapshot in seconds, and obtain a consistent copy of the source data.
 - **Takes up less storage space:** The generated snapshot is not a full complete copy of the physical data, it does not takes up much of the storage space. Thus, even if the source data is huge, snapshots only takes up very less amount of storage space.




Related Concepts


Terms	Definition/Description
source volume	The volume where the source data resides for snapshot operation, it is represented to the host as LUN. Source Volume includes Meta Volume and Data Volume. -Meta Volume: Records the location of source data. -Data Volume: Records the data stored in the source volume.
COW data space	COW data space is dynamic storage allocated from the storage pool where the source LUN resides. This storage is used to store COW data following snapshot generation and activation. All snapshot LUNs of a source LUN share a single COW data space.
snapshot volume	It is the logical data copy generated after the snapshot is created from the source volume. It is represented as a LUN to the users.
snapshot rollback	Replicate the data from the snapshot LUN to the source LUN, which restores the source LUN data to the point where the snapshot was made.
mapping table	It is used to record the changes of data in the source and snapshot LUN at a certain point of time and the changed storage location of data. It is divided into shared and private mapping tables.
inactive	It is a status of the snapshot. In this status, snapshot are unable to be used, and requires activation operation before snapshots can be used as normal.

Mapping Table


- Used for represent the mapping relationship between the snapshot data and the source data:
 - The left portion of the mapped entry refers to the source data location, which serves as the lookup value.
 - The right portion refers to the block pointer, which records the locations of data blocks.
 - Can add or remove the entries within the table.
 - Stored in the form of B+ tree.




Copy On Write (COW)



Snapshot Working Principles



Application Scenario - Continuous Protection of Data




Chapter
07

SmartQuota



Overview of Smart Quota




- Restrict the resource usage amount of a single directory, user, and user group.
 - Prevent individual users from occupying too much of storage space and affecting other users.
- Alert users on resource usage through alarms and event notifications.

Related Terms

Terms	Descriptions
Quota Tree	Quota tree is a special directory of a file system. In a quota tree, you can set directory quotas, file limit quotas and manage storage space occupied by files in the directory.
Root Quota Tree	The root quota tree is actually the root directory of the file system. User quotas and group quotas can be set on the Root Quota Tree to limit the resources available to users.
Soft Quota	Soft quotas refers to a quota value that can be configured. Once the resources used by users exceed this value, the system generates alarms. Alternately, if the used resources are changed from higher to lower than the soft quotas, the alarms are eliminated.
Hard Quota	Hard quota refers to the maximum value of the resources that can be consumed by the users. The amount of resource used by users cannot exceed this value.

Application Scenario - Flexible Restriction of the User Resources



- Shared Directory for R&D Department (Quota Tree 0):
 - Configure the directory quota for Quota Tree 0 to limit the total available resources of the whole R&D department.
 - Configure the private user quota of manager A to individually limit the available resources for manager A.
 - Configure the user group G and E private user group quota to limit the available resources for each of the user groups.



Thank You