QThin Introduction

P400Q Series
P600Q Series

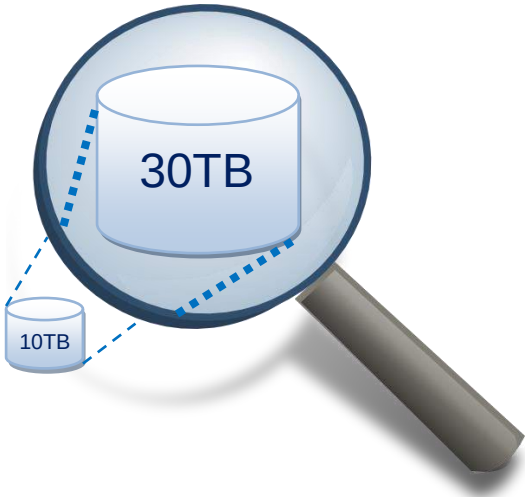

Version 1.2
March 2013

Qsan Technology, Inc.

4F., No.103, Ruihu St.,
Neihu Dist., Taipei City 114,
Taiwan (R.O.C.)

Tel: +886-2-7720-2118
Fax: +886-2-7720-0295

Email: Sales@QsanTechnology.com
Website: www.QsanTechnology.com

QThin - Qsan Thin Provisioning

Nowadays thin provisioning is a hot topic people talk about in IT management and storage industry. To make contrast to thin provisioning, it naturally brings to our minds with the opposite term - fat provisioning, which is the traditional way IT administrators allocate storage space to each logical volume that is used by an application or a group of users. When it comes to the point to decide how much space a logical volume requires for three years or for the lifetime of an application, it's really hard to make the prediction correctly and precisely. To avoid the complexity of adding more space to the volumes frequently, IT administrators might as well allocate more storage space to each logical volume than it needs in the beginning. This is why it's called "fat" provisioning. Usually it turns out that a lot of free space is sitting around idle. This stranded capacity is wasted, which equals to waste of investment and inefficiency. Various studies indicate that as much as 75% of the storage capacity in small and medium enterprises or large data centers is allocated but unused. And this is where thin provisioning kicks in.

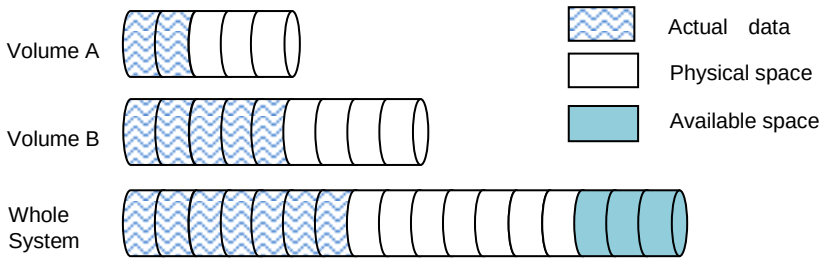

Traditional Fat Provisioning

Thin provisioning sometimes is known as just-in-time capacity or over allocation. As the term explains itself, it provides storage space by requests dynamically. Thin provisioning presents more storage space to the hosts or servers connecting to the storage system than is actually available on the storage system. Put it in another way. Thin provisioning allocates storage space that may or may not exist. The whole idea is actually another way of virtualization. Virtualization is always about a logical pool of physical assets and provides better utilization over those assets. Here the virtualization mechanism behind thin provisioning is storage pool. The capacity of the storage pool is shared by all volumes. When write requests come in, the space will be drawn dynamically from this storage pool to meet the needs.

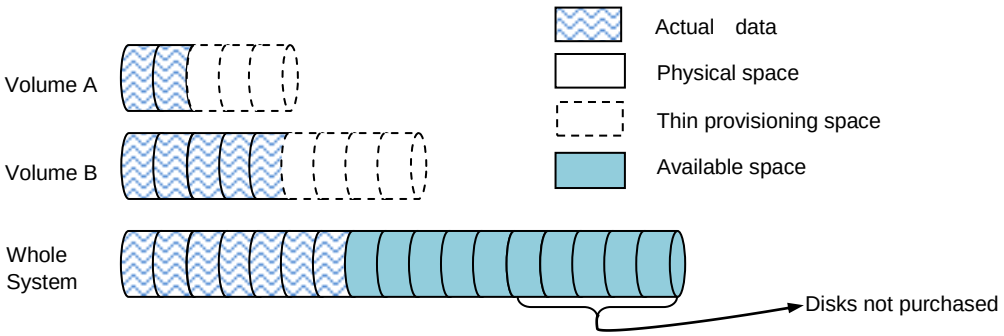

Qsan Thin Provisioning

The Benefits of QThin

- Less disk purchase is needed initially when setting up a new storage system.
  You don't need to buy more capacity to meet your future data growth at present time. Usually hard drive price declines as time progresses. You can buy the same hard drives with cheaper price at a later time. Why not save money upfront while you can?

- No stranded storage capacity, better utilization efficiency and lower total cost of ownership.
  QThin can make full use of the stranded capacity that traditional provisioning can't. All free capacity can be made available to other hosts. A single storage system can serve more hosts and servers to achieve high consolidation ratio. QThin can help you achieve the same level of services with less hard drives purchased upfront, which can significantly reduce your total cost of ownership.

- Scalability: storage pool can grow on demand.
  When the storage pool (RAID group) has reached the threshold you set before. Up to 32 RAID sets can be added to the RAID group to increase the capacity on demand without interrupting I/O. Each RAID set can have up to 64 physical disks.

- Automatic space reclamation mechanism to recycle unused blocks.
  The technology used here is called zero reclamation. When a thin RG is created, the initialization process will try to fill out all the storage pool space with zero. This process will run in background with low priority in order not to impact the I/O performance. This is the reason why when there is no I/O traffic from the hosts, the hard drive LED will keep blinking as if there are I/O activities. The purpose of zero reclamation is that when the actual user data happens to have all zero in a basic allocation unit (granularity), the storage system will treat it as free space and recycle it. Until the next time there is data update to this reclaimed

all zero basic unit, the storage system can swiftly return a basic unit from the free storage pool because it's already filled with zero.

- An eco-friendly green feature that helps to reduce energy consumption.
  Hard drive is the top power consumer in a storage system. Because you can use less hard drives to achieve the same amount of work, this translates directly to a huge reduction of power consumption and more green in your pocket.

Features Highlight

1. Downward firmware compatibility with existing Qsan array firmware.
   You can upgrade your current Qsan Q-series storage system to QThin-enabled firmware with no problem. Certain steps need to be followed to ensure a smooth transition to thin provision enabled environment.
2. Write on demand or allocate on demand.
   This is the most distinctive function in thin provisioning. You can see from the screenshots below that figure 1 shows there are two RAID groups created. "Fat-RG" is using traditional provisioning without QThin enabled and its size is 136GB. "Thin-RG" is QThin-enabled and its size is 272GB.

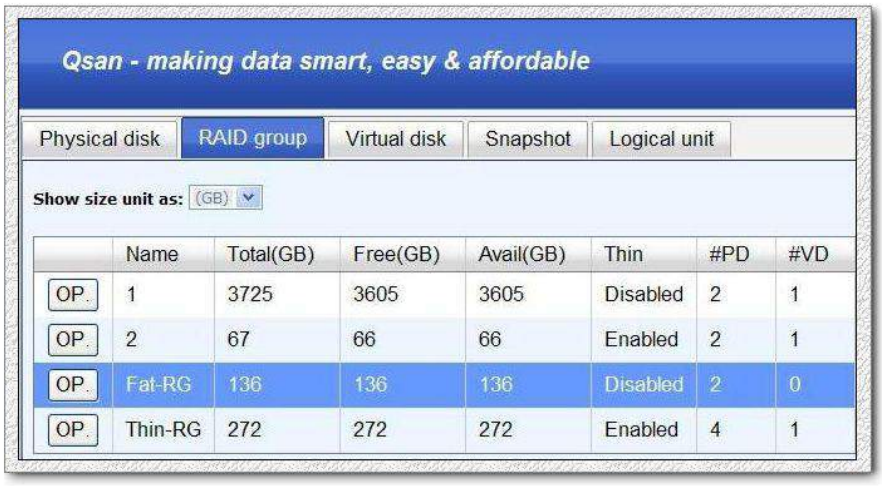

	Name	Total(GB)	Free(GB)	Avail(GB)	Thin	#PD	#VD
OP.	1	3725	3605	3605	Disabled	2	1
OP.	2	67	66	66	Enabled	2	1
OP.	Fat-RG	136	136	136	Disabled	2	0
OP.	Thin-RG	272	272	272	Enabled	4	1

Figure 1: No virtual disk is created

Let's create a Virtual Disk on each RAID group with the same size of 60GB respectively in figure 2 and see what happen.

Figure 2: Virtual disks are created.

In figure 3, the free space of "FAT-RG" immediately reduces to 76GB. 60GB is taken away by the virtual disk. However, the free space of "Thin-RG" is still 272GB even though the same size of virtual disk is created from the RAID group. Nothing is written to the virtual disk yet, so no space is allocated. The remaining 272GB can be used to create other virtual disks. This is storage efficiency.

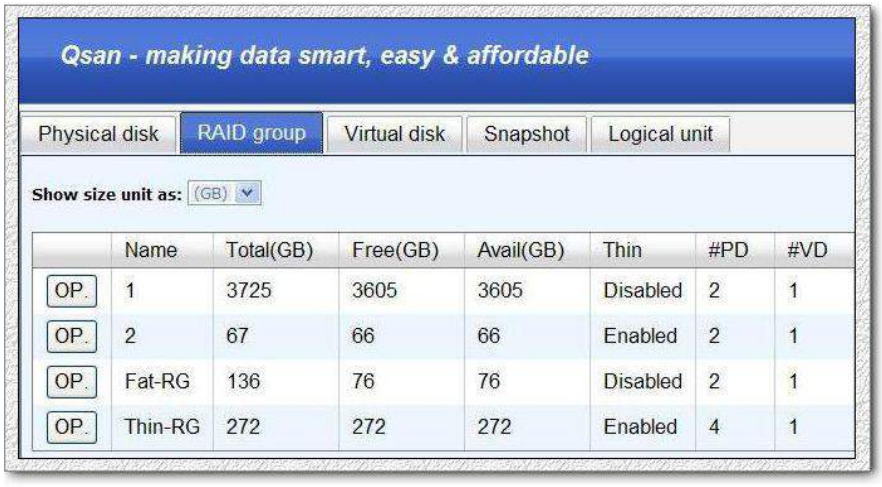

Figure 3: Write on demand

3. Expand capacity on demand without downtime.

   Extra RAID set can be added to the thin RAID group to increase the size of free storage pool. A thin RAID group can have up to 32 RAID sets with each RAID set containing up to 64 physical hard drives. The maximum size of each RAID set is 64TB. Figure 4 shows that "Thin-RG" consists of two RAID sets.

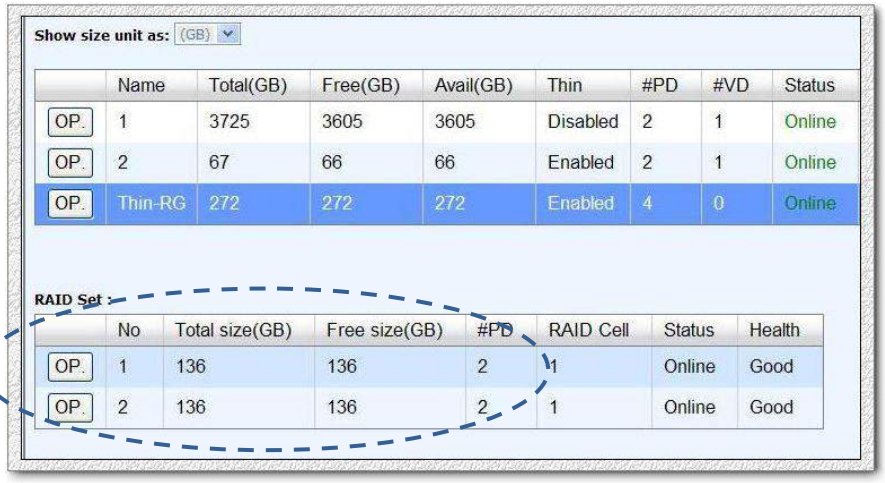

Figure 4: Scalable RAID group size

4. Allocation unit (granularity) is 1GB. This is a number that demands careful balance between efficiency and performance. The smaller it is, the better the efficiency and the worse the performance becomes, and vice versa.
5. Thin provisioned snapshot space and it is writeable.

   Snapshot space sits at the same RAID group of the volume that the snapshot is taken against. Therefore when you expose the snapshot into a virtual disk, it becomes a thin-provisioned virtual disk. It will only take up the just the right amount of space to store the data, not the full size of the virtual disk.
6. Convert traditional VD to QThin VD and vice versa.

   You can enjoy the benefits of QThin right now and right this moment. Upgrade your Qsan Q-series systems to QThin-enabled firmware. Move all your existing fat-provisioned virtual disks to thin-provisioned ones. VD cloning functions can be performed on both directions - fat to thin and thin-to-fat, depending on your application needs. Figure 5 shows cloning a fat virtual disk to a thin one.

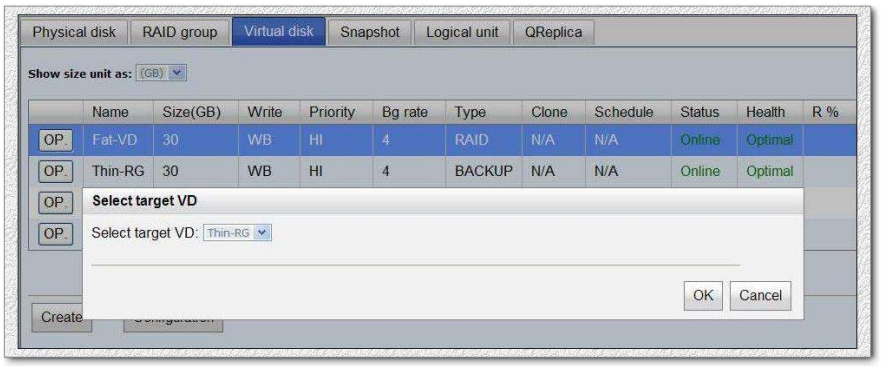

Figure 5: VD cloning between thin VD and fat VD

7. Threshold settings and capacity policies.

   These are designed to simplify the management and better monitoring the storage usage. You can set as many as 16 policies for each RAID group. When space usage ratio grows over the threshold set in the policy, the action will be taken and event log will be generated.

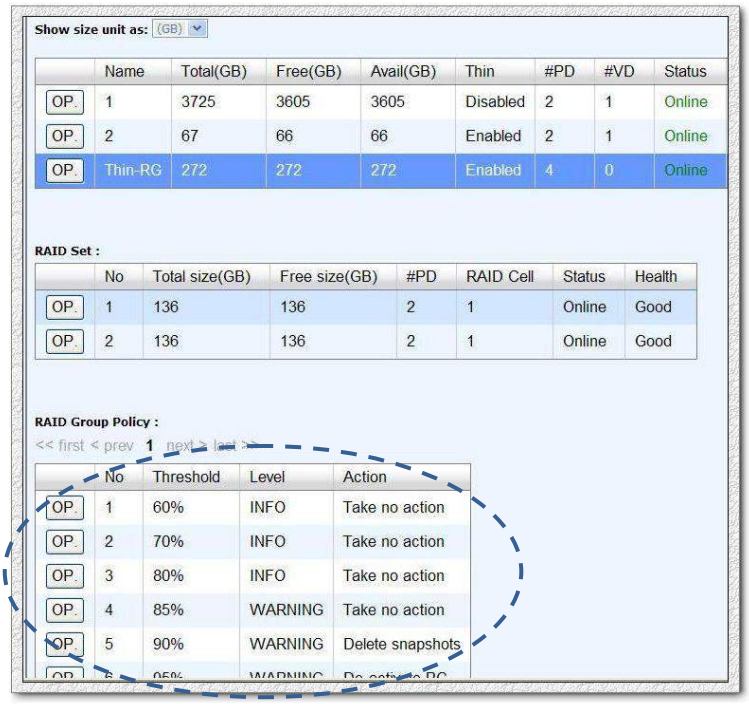

Figure 6: Capacity policy settings

8. Automatic space reclamation to recycle all zero basic units and increase utilization rate.

   Automatic space reclamation will be automatically activated in RAID group initialization process or it can be set manually through capacity policy. You can set as many as 16 policies. When space usage ratio grows over the threshold set in the policy, space reclamation will be enabled automatically at the background with the lowest priority or when the I/O is low. The resource impact is reduced to minimum.

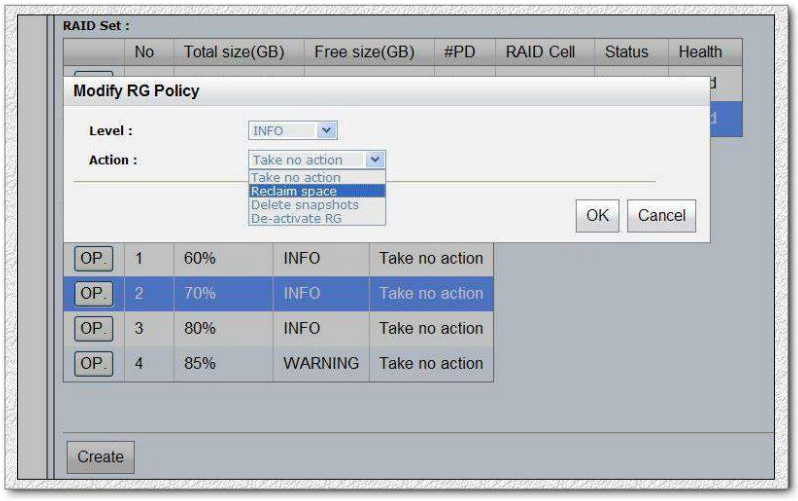

Figure 7: Space reclamation

What scenarios does Thin Provisioning fit well?

We suggest that you apply QThin to non-critical production applications first. Thin provisioning works well when the data written is thin-friendly, which means that the data written is not completely spread across the whole volume. Applications that spread metadata across the entire volume will obviate the advantages of thin provisioning. Some applications that expect the data to be contiguous at block level are not good candidates for thin provisioning as well.

QThin works well with email system, web-based archive, or regular file archive system. When the number of supported volumes grows larger, the benefits of QThin will become more apparent.

Summary

The introduction of thin provisioning in storage industry has been more than five years. It's proven to be a valuable and popular feature in designing storage array. As this technology matures over time, we will see more and more applications and new operating systems built to support thin provisioning. QThin delivers the highest possible storage utilization and allocates the right capacity at the right time with proper alert and quota management built in. Qsan believes it is now the perfect timing to add this great function to Qsan Q-series models and provide this powerful function to its enterprise customers to help them achieve higher level of storage efficiency, save a lot of electricity bills and produce smaller carbon footprints. We hope you will enjoy using the latest Qsan thin provisioning technology - QThin.

Reference

- Wikipedia - Thin provisioning
  http://en.wikipedia.org/wiki/Thin_provisioning