

HP 3PAR Thin Technologies



Table of contents

| | |
|---|----|
| Introduction | 3 |
| Overview of HP 3PAR Thin Suite | 3 |
| HP 3PAR Thin Provisioning Software..... | 3 |
| HP 3PAR Thin Persistence Software..... | 4 |
| HP 3PAR Thin Conversion Software..... | 4 |
| HP 3PAR Thin Copy Reclamation Software..... | 4 |
| HP 3PAR ASIC with Thin Built In..... | 4 |
| The benefits of thin provisioning | 5 |
| Avoid frequent storage capacity addition to servers | 5 |
| Accelerate time to market | 5 |
| Chargeback model in utility storage | 5 |
| Using fully provisioned volumes..... | 6 |
| Licensing..... | 6 |
| HP 3PAR Volume Manager | 7 |
| Physical Disks..... | 7 |
| Logical Disks..... | 7 |
| Common Provisioning Groups..... | 7 |
| Virtual Volumes | 7 |
| Common Provisioning Groups..... | 8 |
| CPGs and workloads..... | 8 |
| Availability level..... | 8 |
| Reasons to create multiple CPGs | 9 |
| CPG automatic growth..... | 9 |
| Changing LD characteristics within a CPG..... | 10 |
| Snapshot space within CPGs | 10 |
| Allocating CPG space to TPVVs | 10 |
| TPVV sizing guidelines..... | 11 |
| Thin provisioning and applications | 11 |
| Oracle..... | 11 |
| Microsoft SQL Server..... | 11 |
| Microsoft Hyper-V | 12 |
| SAP | 12 |
| VMware vSphere | 12 |

| | |
|---|----|
| Managing thin storage | 13 |
| Allocation warnings | 13 |
| Allocation limits | 14 |
| Remaining physical capacity alerts | 15 |
| Remaining CPG free space alerts | 15 |
| View logged warning and limit alerts | 15 |
| User recommendations | 16 |
| Tracking volume space usage | 16 |
| TPVVs | 16 |
| CPGs | 17 |
| System space | 17 |
| Total used space | 18 |
| Running out of space | 19 |
| Running out of space in a TPVV | 19 |
| Running out of space in a CPG | 19 |
| Running out of space on the system | 19 |
| Usage reporting and trend analysis | 19 |
| Reclaiming Unused Space | 19 |
| Reclaiming Unmapped Logical Disk Space from CPGs | 20 |
| Automatically Reclaiming Unused Snapshot Space from Volumes | 20 |
| Deleted Volume Snapshot Space | 20 |
| Logical Disks and Chunklet Initialization | 20 |
| Free page consolidation | 20 |
| HP 3PAR Thin Persistence software | 21 |
| Thin persistence methods | 21 |
| Thin Reclamation for Microsoft Windows Server 2012 | 21 |
| Thin Reclamation for Microsoft Windows Server 2003 and 2008 | 22 |
| Thin Reclamation for VMware vSphere | 22 |
| Thin Reclamation for Linux | 23 |
| Thin Reclamation for UNIX | 24 |
| Thin Reclamation for Veritas Storage Foundation | 24 |
| Thin Reclamation for Oracle databases | 24 |
| Migration to thin-provisioned volumes | 26 |
| Migrate from full to thin provisioned volumes across arrays | 26 |
| Migrate between full and thin provisioned volumes on the same array | 26 |
| Conclusion | 27 |

Introduction

Balancing the storage needs of new projects or of unpredictable workloads against limited resources is the prime challenge of IT managers today. Of the many proposed solutions for improving storage efficiency, few are actually implemented. Of those implemented, fewer still achieve demonstrable success.

Thin provisioning has achieved widespread adoption as it dramatically increases capacity efficiencies. It has become a data center “must have” for its ability to break the connection between logical and physical capacity. However, not all thin-provisioning implementations deliver the same results. Some are complex to deploy, while others use coarse allocation units and cannot deliver the required space savings.

Thin provisioning allows a volume to be created and made available as a logical unit number (LUN) to a host without the need to dedicate physical storage until it is actually needed. HP 3PAR Thin Provisioning software has long been considered the gold standard in thin provisioning for its simplicity and efficiency. Unlike other “bolt-on” implementations, HP 3PAR Thin Provisioning software is simple and efficient, helps your organization start new projects more quickly and on demand and save millions of dollars. HP 3PAR Thin Provisioning leverages the dedicate-on-write approach of HP 3PAR StoreServ Storage, allowing enterprises like yours to purchase only the disk capacity they actually need. HP 3PAR Thin Provisioning integrates seamlessly with VMware vSphere, Windows® Server 2012, Red Hat Enterprise Linux, and Symantec Veritas Storage Foundation—greatly enhancing the operative and administrative efficiency of these platforms.

While HP 3PAR Thin Provisioning software is extremely simple to deploy and use, a certain amount of planning is advantageous to maximize its benefits. This paper documents best practices on thin provisioning on HP 3PAR StoreServ Storage and is intended for administrators looking to get the most out of their HP 3PAR StoreServ deployment. In addition, it describes other HP 3PAR thin technologies that you can use in conjunction with HP 3PAR Thin Provisioning software to maximize its effectiveness. Unique to HP 3PAR StoreServ, HP 3PAR Thin Conversion software enables you to reduce capacity requirements by 50 percent or more¹ by deploying HP 3PAR StoreServ in place of legacy storage.

HP 3PAR Thin Persistence software and other thin-reclamation solutions enable thin-provisioned storage on HP 3PAR StoreServ arrays to stay thin over time by ensuring that unused capacity is reclaimed for use by the array on an ongoing basis.

Overview of HP 3PAR Thin Suite

HP 3PAR Thin Provisioning Software

HP 3PAR Thin Provisioning is the most comprehensive thin provisioning software solution available. Since its introduction in 2002, HP 3PAR Thin Provisioning has been widely considered the gold standard in thin provisioning. It leverages HP 3PAR StoreServ’s dedicate-on-write capabilities to make storage more efficient and more green, allowing customers to purchase only the disk capacity they actually need and only as they actually need it.

With Thin Provisioning, there is no more up-front capacity allocation. No more dedicating resources for each application. No more paying to house, power, and cool disks that might not be needed for months or years to come, or may never be needed at all.

Product Highlights

- HP 3PAR Thin Provisioning is completely automated
- HP 3PAR Operating System uses a reservationless, dedicate-on-write approach to thin provisioning that draws and configures capacity in fine-grained increments from a single free space reservoir without pre-dedication of any kind.
- HP 3PAR Thin Provisioning uses an allocation unit size of just 16 KB, so you don’t have to worry about small writes consuming megabytes or even gigabytes of capacity.
- HP 3PAR StoreServ is a storage platform built from the ground up to support thin provisioning by eliminating the diminished performance and functional limitations that plague bolt-on thin solutions.

¹ Refer to the HP 3PAR Thin Guarantee for details: hp.com/storage/getthin.

HP 3PAR Thin Persistence Software

HP 3PAR Thin Persistence software is an optional feature that keeps Thin Provisioned Virtual Volumes (TPVVs) and read/write snapshots of TPVVs small by detecting pages of zeros during data transfers and not allocating space for the zeros. This feature works in real-time and analyzes the data before it is written to the source TPVV or read/write snapshot of the TPVV. Freed blocks of 16 KB of contiguous space are returned to the source volume, freed blocks of 128 MB of contiguous space are returned to the CPG for use by other volumes. With HP 3PAR Thin Persistence, HP 3PAR StoreServ Storage customers can now leverage next generation space reclamation technology to minimize storage TCO.

HP 3PAR Thin Conversion Software

With HP 3PAR Thin Conversion, a technology refresh no longer requires a terabyte-for-terabyte replacement, but instead offers the opportunity to eliminate up to 75%² of your legacy capacity, simply and rapidly. This savings alone can save you up to 60% on the cost of a technology refresh.

With HP 3PAR Thin Conversion, you can quickly shrink your storage footprint, reduce storage TCO, and meet your Green IT targets. HP 3PAR Thin Conversion software makes this possible by leveraging the zero-detection capabilities within the HP 3PAR ASIC and Thin Engine (a unique virtualization mapping engine for space reclamation) to power the simple and rapid conversion of inefficient, “fat” volumes on legacy arrays to more efficient, higher-utilization “thin” volumes. Getting thin has never been so easy.

HP 3PAR Thin Conversion Software is an optional feature that converts a fully-provisioned volume to a Thinly-Provisioned Virtual Volume (TPVV). Virtual volumes with large amounts of allocated but unused space are converted to TPVVs that are much smaller than the original volume. During the conversion process, allocated but unused space is discarded and the result is a TPVV that uses less space than the original volume.

HP 3PAR Thin Copy Reclamation Software

An industry first, Thin Copy Reclamation keeps your storage as lean and efficient as possible by reclaiming unused space resulting from deleted virtual copy snapshots and remote copy volumes. Thin Copy Reclamation is built on a virtualization mapping engine for space reclamation called HP 3PAR Thin Engine, which is included as part of the HP 3PAR OS. HP 3PAR Thin Copy Reclamation Software is an optional feature that reclaims space when snapshots are deleted from a system. As snapshots are deleted, the snapshot space is reclaimed from a Thinly-Provisioned Virtual Volume (TPVV) or fully-provisioned virtual volume and returned to the CPG for reuse by other volumes. Deleted snapshot space can be reclaimed from virtual copies, physical copies, or remote copies.

HP 3PAR ASIC with Thin Built In

At the heart of every HP 3PAR StoreServ node there is a HP 3PAR ASIC with Thin Built In which features an efficient, silicon-based zero-detection mechanism. This unique hardware capability gives HP 3PAR StoreServ Storage the power to remove allocated but unused space inline and non-disruptively without impacting performance.

The zero detect capability will recognize an incoming write request of 16 KB of zeros and either not allocate space for the zero block or free the space if it was already allocated for that block. All this happens in cache and therefore no zeroes are written to the backend. When a read request comes in for a block that is unallocated the HP 3PAR StoreServ will immediately return zeros back to the host.

Many other storage arrays do not detect blocks of zeroes on write. Instead, the zeros are written to disk and a scrubbing process later detects these zeroed blocks and discards them. With this approach the zeroed blocks consume space until they're scrubbed and therefore they may not be available for use by other volumes when needed. Also there is increased load placed on the storage as the scrubbing process examines the block contents on the physical storage.

The HP 3PAR StoreServ built-in zero detection capability can be controlled per virtual volume and it is enabled by default.

² Based on documented client results that are subject to unique business conditions, client IT environment, HP products deployed, and other factors. These results may not be typical; your results may vary.

The benefits of thin provisioning

Thin provisioning can be used with nearly all applications available in the market to improve storage utilization dramatically. The following use cases illustrate the superior value from thin provisioning on HP 3PAR StoreServ.

Avoid frequent storage capacity addition to servers

Storage administrators often allocate and present a large amount of storage to a server at the start of a project to accommodate its long-term growth requirements. This practice reduces the number of times that LUNs mapped to the server have to be expanded, an operation that can be complex, time consuming, and cause server downtime. With the HP 3PAR Thin Provisioning software, it is possible to allocate large amounts of storage to a particular server but only consume physical space on the array as used. This fits environments where:

- The addition or expansion of storage provisioned to servers is not desired, for example in mission-critical environments.
- Relatively slow growth rates or unpredictable growth over time occurs. This can happen on large file systems used for group shares, mailbox databases, or general database space.

Accelerate time to market

In an application's development stage, there is sometimes a requirement for the application storage to be in place and ready before the application goes live. With the HP 3PAR Thin Provisioning software, it is possible to present the full amount of storage immediately so that it is ready for the developers to work on without the requirement for the full physical capacity being in place. Because thin provisioning gives storage administrators the ability to place limits on TPVVs and CPGs, the administrator can make sure that the developer's work does not affect other production systems by using the free space within the HP 3PAR StoreServ. After the application is ready to go live, these limits can be removed.

Chargeback model in utility storage

HP 3PAR Thin Provisioning is ideal for enterprise customers and service providers wishing to deploy a storage offering where usage chargeback is an important component of the service. Thin provisioning offers these customers the following benefits:

- Deploy quickly
- Decouple the charges for storage from the limitations of actual presented storage
- Remove the exposure to disruption of service during future capacity expansions
- Increase profit margins based on the savings resulting from the difference in cost for provisioned storage compared to actual physically used storage
- Collect detailed charging data at the individual VV level (usage shown by the `showvv -p -prov [full|cpw|tpw|snp]`) or at the CPG level (usage shown by the `showcpg -r <cpg>` or `showvv -p -cpg <cpg>`)

When planning to collect charging data, it is recommended that the names chosen for objects like CPGs, VVs, snapshots, and domains contain a meaningful prefix or suffix referring to the project, application, line of business or department the objects belong to. This enables the grouping of objects in the chargeback report. The HP 3PAR OS allows up to 31 characters for the name of an object.

Using fully provisioned volumes

The use of thin provisioning has minimal performance impact and has the significant operational benefit of reducing storage consumption. However, there are certain workloads and applications thin provisioning may not be of benefit such as:

- **Systems with a high file system utilization.** File systems on TPVVs that are nearly full offer reduced benefits of thin provisioning. This relates to the fact that, for thin volumes, the thin-provisioning license charge is incurred in addition to the cost for physical disk capacity used. In the case of file system utilization rates of 80 percent or higher, it may be more cost-efficient to use fat-provisioned volumes to hold the data.
- **Applications that write continuously to new space.** An example of this is Oracle redo log files.
- **Databases not in “auto-extend” mode (or equivalent).** Some databases initialize their assigned storage at creation time by writing markers at regular intervals over the entire volume. This has the same effect as provisioning file systems with a high utilization on thin volumes.
- **Small capacity requirements.** Thin provisioning is ideal for large-scale volumes. For small size VVs (256 MB up to a few tens of GB) even the minimum growth increment of the CPG may mean that minimal benefit is realized. Use care in the selection of the CPG growth increment in this case.
- **Environments that require host encrypted volumes.** Writing blocks of zeros to a host encrypted volume on a newly created HP 3PAR StoreServ thin-provisioned volume will cause space to be allocated on the TPVV because the encryption alters the content of the blocks. Applying encryption to thin-provisioned volumes that already contain data or rekeying them also inflates the zero blocks, making the volume consume space as if it was fully provisioned. Attempting to rethin the volume by writing zeros to allocated but unused space is not possible as well. As a result, host encryption and thin provisioning do not cooperate well.
- **Environments that require SAN encrypted volumes.** Like host based encryption, encryption by a device in the data path (e.g. SAN switch) will also alter the data stream so that blocks of zeros written by the host are not passed onto the storage. A notable exception is Brocade SAN switches. With the introduction of Fabric OS 7.1.0, the Fabric OS encryption switch can automatically detect if a disk LUN is a thin provisioned LUN. If a LUN is detected as being thin provisioned, then first-time encryption and rekey are done on the allocated blocks only. This thin provision LUN support requires no action by the user.
- **Copy-on-Write file systems.** File systems that write to new blocks rather than overwrite existing data are not suitable for thin provisioning as every write will allocate new storage until the volume is fully allocated. An example of a CoW file system is Solaris ZFS.

Licensing

The HP 3PAR Thin Provisioning software license is needed to create TPVVs. Thin provisioning is licensed based on the number of terabytes of raw capacity written to thin provisioned volumes. Space used by snapshots is also included in this figure. Because of the “raw capacity written” metric, no fraction of the license is consumed by creating thin provisioned volumes and not writing to them. The rate of consumption of a thin provisioning license will depend on the RAID level of the TPVVs. Assume the array has a thin provisioning license for 1 TB and a 500 GB TPVV in a RAID 1 was created. When the user writes 300 GB of data to it, the thin provisioning license is consumed for $300 \text{ GB} * 2 = 600 \text{ GB}$. The multiplier of 2 originates from the RAID 1 overhead factor. If the TPVV was created in RAID 5 (3D+1P), $300 * 4/3 = 400 \text{ GB}$ of space would have been consumed from the thin provisioning license.

When doing a remote copy of a thin volume, the destination array needs a thin provisioning license as well. Otherwise the copy of a TPVV volume on the source array will create a full provisioned volume at the target.

The HP 3PAR Thin Conversion software, HP 3PAR Thin Persistence software and HP 3PAR Thin Copy Reclamation software also require licenses. The functionality offered by each license is summarized in table 1.

Table 1. Summary of the thin storage technology licenses

| License | Functionality |
|-----------------------|---|
| Thin Provisioning | The creation of Thinly-Provisioned Virtual Volumes |
| Thin Persistence | Zero detection to free previously allocated space |
| Thin Conversion | Zero detection to prevent allocation of space and enables UNMAP support |
| Thin Copy Reclamation | Reclaiming unused space resulting from deleted virtual copy snapshots and remote copy volumes |

For the HP 3PAR StoreServ 7000, 7450, and 10000, the licenses for HP 3PAR Thin Provisioning, Thin Conversion, Thin Persistence, and HP 3PAR Thin Copy Reclamation are included in the HP 3PAR Operating System Software Suite.

HP 3PAR Volume Manager

The HP 3PAR Operating System has a logical volume manager that provides the virtual volume abstraction. It is comprised of several layers with each layer being created from elements of the layer below.

Physical Disks

Every physical disk (PD) that is admitted into the system is divided into 1 GB³ chunklets. A chunklet is the most basic element of data storage of the HP 3PAR OS. These chunklets form the basis of the RAID sets, depending on the sparing algorithm and system configuration some chunklets are allocated as spares.

Logical Disks

The logical disk (LD) layer is where the RAID functionality occurs. Multiple chunklet RAID sets, typically from different PDs, are striped together to form a LD. All chunklets belonging to a given LD will be from the same drive type. LDs can consist of all NL, FC, or SSD chunklets. There are no mixed type LDs with the exception of Fast Class (Fibre Channel or SAS) where an LD may consist of mixed 10K and 15K drive chunklets.

There are three types of logical disk:

1. User (USR) logical disks provide user storage space to fully provisioned virtual volumes.
2. Snapshot data (SD) logical disks provide the storage space for snapshots, virtual copies or thinly provisioned virtual volumes (TPVV).
3. Snapshot administration (SA) logical disks provide the storage space for snapshot and TPVV administration. They contain the bitmaps pointing to which pages of which SD LD are in use.

The LDs are divided into "regions", which are contiguous 128 MB blocks. The space for the virtual volumes is allocated across these regions.

Common Provisioning Groups

The next layer is the common provisioning group (CPG) which defines the LD creation characteristics, such as RAID type, set size, disk type for chunklet selection, plus total space warning and limit points. A CPG creates a virtual pool of LDs that allows Volumes to share resources and to allocate space on demand. A volume created from a CPG can automatically be grown on demand by mapping new regions from the LDs associated with the CPG. Common provisioning groups and their importance to thin provisioning will be discussed in more detail in the next section.

Virtual Volumes

The top layer is the virtual volume (VV) which is the only data layer visible to hosts. Virtual volumes draw their resources from CPGs, and the volumes are exported as Logical Unit Numbers (LUNs) to hosts.

A VV is classified by its type provisioning which can be one of the following:

- full. Fully provisioned VV, either with no snapshot space or with statically allocated snapshot space.
- tpvv. Thin Provisioned VV, with space for the base volume allocated from the associated user CPG and snapshots space allocated from the associated snapshot CPG (if any).
- cpvv. Commonly Provisioned VV. The space for this VV is fully provisioned from the associated user CPG and the snapshot space allocated from the associated snapshot CPG.

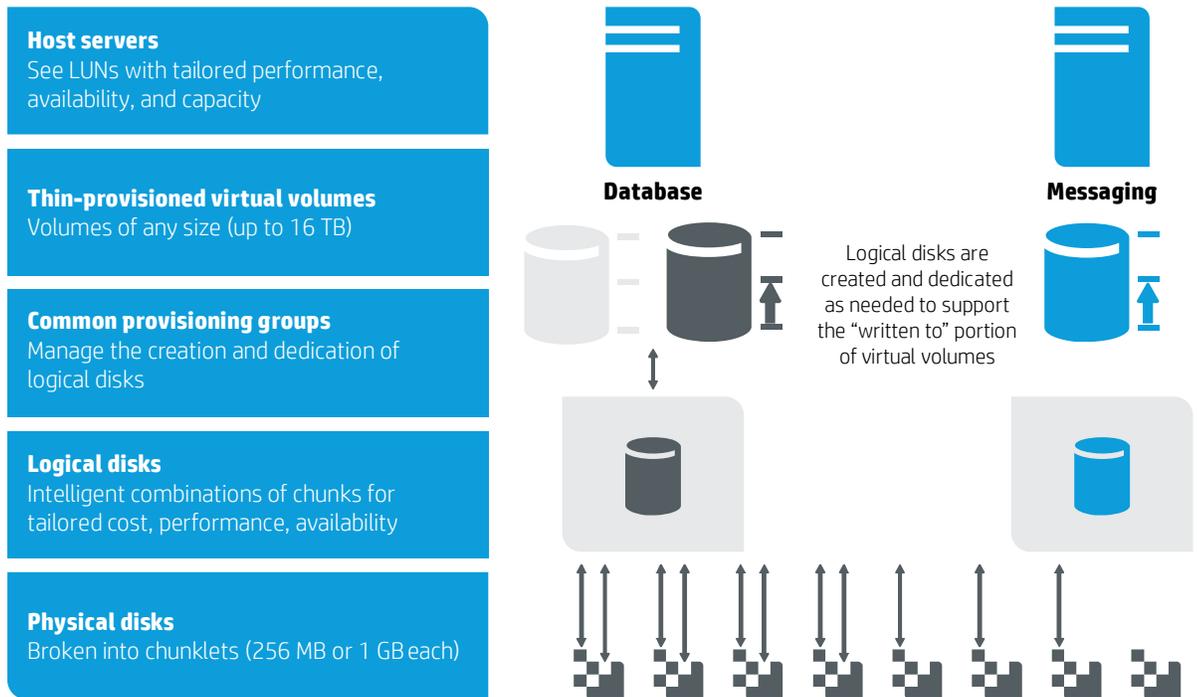
On creation a TPVV the size of the VV is specified, but no storage is allocated. Storage is allocated on demand in the Snapshot Data area as required by the host operation being performed. The Snapshot Admin area contains the metadata indexes that point to the user data in the SD area. Since the SA metadata needs to be accessed to locate the user data the indexes are cached in policy memory to minimize the performance impact of the lookups.

TPVVs associated with the same CPG share the same LDs and draw space from that pool as needed, allocating space on demand in small increments for each controller node. As the volumes that draw space from the CPG require additional storage, the HP 3PAR OS automatically creates increases the logical disk storage until the CPG reaches the user-defined growth limit which restricts the CPG's maximum size. The size limit for an individual TPVV volume is 16 TB.

³ Legacy systems (S/E/FT-class) use 256 MB chunklets.

The relationship between the HP 3PAR OS abstraction layers is illustrated in figure 1.

Figure 1. Overview of thin-provisioned virtual volumes



Common Provisioning Groups

Common Provisioning Groups (CPGs) are policies for how free chunklets within the HP 3PAR StoreServ should be used when creating volumes. A CPG policy contains parameters such as disk type, disk speed, RAID type, growth increment, chunklet radial placement, and availability level. CPGs automatically grow the underlying LD storage, according to the stated growth parameters, on demand to store data in a TPVV. No administrator intervention is needed with this provisioning operation.

CPGs and workloads

HP 3PAR StoreServ performs optimally for any type of workload, and different workloads can be mixed on the same array. These different workloads may need different types of service levels to store their data. For example, for high-performance mission-critical workloads it may be best to create volumes with RAID 5 protection on Solid State Drives (SSD) or RAID 1 protection on Fast Class (Fibre Channel or SAS performance HDDs). For less-demanding projects, RAID 5 on FC drives or RAID 6 on Nearline (NL) drives may suffice. For each of these workloads, you can create a CPG to serve as the template for creating VVs for the workload. It is most efficient to keep the number of CPGs low as each CPG reserves its amount of growth space. TPVVs can be moved between CPGs with the HP 3PAR Dynamic Optimization (DO) software command tuneuv, thereby changing their underlying physical disk layout and hence their service level.

The following sections discuss what to consider when planning CPGs for thin provisioning and virtual copy snapshots and recommend a set of best practices.

Availability level

To provide high availability, chunklets from the same RAID set should be distributed across multiple components.⁴ There are three levels of availability that can be selected with HP 3PAR StoreServ.

- **HA CAGE** means that no two members of the same RAID set can be in the same drive enclosure. For example, to support RAID 5 3+1 (set size 4), four drive chassis **connected** to the same node pair are required. This helps ensure that data is still available in the event that access to an entire drive cage is lost. This applies to drive chassis that are point-to-point connected to the nodes (no daisy chain).

⁴ It is important to understand that drive magazines consist of four drives for HP 3PAR StoreServ 10000. Drive magazines consist of only a single drive in the HP 3PAR StoreServ 7000 and 7450 Storage systems.

- **HA MAG** means that no two members of the same RAID set are in the same drive magazine. This allows a wider stripe with fewer drive chassis; for example, a RAID 5 stripe size of 7+1 (set size 8) would be possible with only four drive chassis, provided each chassis had at least two drive magazines.
- **HA PORT** applies only to daisy-chained drive chassis. When this level of availability is selected, no two members of the same RAID set can be in drive chassis that are dependent on one another for node connectivity. For example, in a system in which there are eight drive chassis with four of the drive chassis connected to another drive chassis for node access, HA PORT would only allow RAID 5 3+1 (set size 4) in order to prevent the loss of one drive chassis from causing a loss of data access. On systems that do not have daisy chained cages, such as the HP 3PAR StoreServ 10000, setting HA PORT is the same as setting HA CAGE.

When creating CPGs it is strongly recommended to always select HA cage availability.

Reasons to create multiple CPGs

All TPVVs associated with a CPG allocate space from a shared pool of LDs. This means that VVs associated with a particular CPG all have identical LD characteristics. VVs that require different characteristics must use a different CPG.

Reasons to create multiple CPGs include:

- To define different service levels e.g. RAID 1, RAID 5, or RAID 6 with FC, NL, or solid-state drives (SSDs)
- To map TPVVs belonging to different lines of business, departments, or customers onto particular CPGs for reporting and management purposes; creating CPGs with exactly the same characteristics but a different name is possible; this allows a logical separation of resources and may help with chargeback models, as chargeback could be based on CPG space usage rather than usage at an individual VV level
- When virtual domains are used, because a CPG can only belong to one virtual domain
- When HP 3PAR Adaptive Optimization (AO) is used, because a CPG can only belong to one Adaptive Optimization policy

While there are several reasons to create multiple CPGs it is recommended that the number of CPGs be kept low as each CPG will reserve its own growth space.

CPG automatic growth

By default CPGs dynamically allocate storage in increments specified by the CPG's growth increment. This on-demand allocation unit determines the automated response to the growth demands of TPVVs. The growth increment should be large enough to ensure wide striping across most or all physical disks that are part of the CPG. To grow the TPVV, the HP 3PAR OS may expand existing LDs according to the CPGs growth increment or create additional ones. Growth is triggered when the CPG's available space falls below 85 percent of the growth increment value. The CPG growth increment can be changed at any time, which also changes the threshold for the next growth increment to happen. A mechanism with warnings and limits can be configured on the array to control the growth of a CPG. When the growth increment is set to zero, the CPG does not grow automatically. If the CPG cannot grow then when the free space in the CPG is exhausted I/Os that require space to be allocated will fail.

The default and the minimum growth increment for a CPG depend on the number of nodes in the system. Table 2 lists the default growth increment and its limits for different numbers of node pairs in a system. The maximum growth increment for a CPG is 2,047.750 GB being 2 TB minus 256 MB.

Table 2. Default and limits for the growth increment per node pair

| Number of nodes | Default (GB) | Minimum (GB) | Maximum (GB) |
|-----------------|--------------|--------------|--------------|
| 2 | 32 | 8 | 2,047.750 |
| 4 | 64 | 16 | 2,047.750 |
| 6 | 96 | 24 | 2,047.750 |
| 8 | 128 | 32 | 2,047.750 |

Considerations when selecting a CPG growth increment for a CPG include:

- The CPG growth increment can be changed at any time.
- To determine the optimal growth increment for your environment, review the following factors influencing the value for it:
 - With the default growth increment of 32 GB per node pair and 1 GB chunklets the new CPG growth will be spread across 16 disks on each node.
 - If there are less than 16 drives of the device type associated with the CPG connected to each node then type then consideration should be given to reducing the CPG growth increment to minimize the amount of reserved
 - Growth space. This is of particular importance for CPGs created on SSDs, the space of which is usually limited within the array.
 - If there are more than 16 drives of the device type associated with the CPG connected to each node then type then it may seem tempting to increase the growth increment to spread the data across more disks. This is not necessary as the next growth will use chunks from the remaining drives. It is therefore recommended to keep the growth at the default value.
 - However, if the environment is write-intensive, the rate of consumption of CPG space might be significant. In this case, it is recommended that the growth increment be set to a value above the default value listed in table 2.

Changing LD characteristics within a CPG

LD space within a CPG is automatically allocated to TPVVs as needed. Storage administrators have no direct control over which LDs are used for a particular allocation. Although the capability exists to modify the LD characteristics for new LDs within a CPG, its use is advisable only in rare instances, for example when different LD characteristics would enable the CPG to fully utilize the remaining available physical capacity.

Snapshot space within CPGs

A TPVV can be associated at creation time or at a later stage with Snapshot Data (SD) space, which is used to store data changed in the original volume. The space for one or more snapshots or virtual copies can originate from the same CPG as for the user space in the TPVV or from a different one. The CPG for the snapshot space of a TPVV can be changed anytime later using the `tunevv` command. The snapshot space grows on-demand with growth increments as defined by the CPG in which it was created. This association eliminates the planning required to estimate and dedicate capacity upfront for copy-on-write snapshots of a volume.

Allocating CPG space to TPVVs

Each TPVV associated with a CPG allocates space from the following sources in the listed order:

- Unallocated LD space in CPGs. LD space is shared by all the TPVVs associated with the CPG. This space can grow to a considerable amount as the SA and SD space of any VV associated with the CPG is returned to the CPG as unallocated LD space upon VV removal or when the `free space` command is issued for a TPVV. The returned, unallocated LDs are remapped to the SD space of associated TPVVs or their snapshots as needed over time. LDs that are entirely unused will be removed by the system if the available free space is more than the growth increment. Unallocated LD space for a CPG can be displayed using the command `showld -cpg <CPG>`. Newly created LDs to accommodate TPVV growth will show as unused initially.
- Free chunklets. Free chunklets available to a CPG for creating new LDs may be limited by the LD creation parameters for the CPG. It is important to understand that, if different CPGs can draw from the same pool of chunklets, the chunklets allocated to one CPG will reduce the pool of storage available to the other CPGs. It is recommended that storage administrators implement the following strategies to stay abreast of available free space:
 - Set the CPG allocation warnings (and limits, if necessary)
 - Monitor the rate of free space reduction using the command `showspace -cpg <CPG> -hist`
 - Set the free space capacity warning with the `setsys RawSpaceAlertXX` command (as detailed in [“Allocation warnings”](#) section later in this paper)
 - Monitor available free space alerts set by the `showalert` command
 - Maintain a buffer of physical capacity

TPVV sizing guidelines

What is the ideal size for a TPVV? There is no definite answer to this, but you should consider the following when deciding on the size for thin volumes:

- The minimum size for a TPVV is 256 MB, the maximum size is 16 TB for all HP 3PAR OS versions on all types of HP 3PAR StoreServ systems.
- Thin provisioning demonstrates the most value in situations involving large-scale consolidation. For small virtual volumes (256 MB up to a few tens of GB), the growth increment of the CPG of the TPVV may be many times higher than the TPVV size meaning that minimal benefit is realized after one growth increment was applied.
- It is possible to increase the size of a thin volume. This provides a significant amount of flexibility. However, the impact of growing volumes for host-based OS needs to be considered. It is not possible to shrink a TPVV.
- When faced with a choice, it is preferable to make volumes larger than needed over making them too small. If the volumes that are presented are too small for the ongoing requirements, then TPVV growth or additional volumes will be required in the future, which is something that needs to be managed.
- The TPVV license capacity applies to allocated raw space, therefore creating larger TPVVs will not consume more license units.

Thin provisioning and applications

All applications including mission-critical ones can use thin provisioned volumes. The use of thin volumes may need some consideration upfront to enable them to utilize the volume space thin-efficiently. The following sections include recommendations for installing a few well-known applications on thin-provisioned volumes.

Oracle

In an Oracle environment, HP 3PAR Thin Provisioning provides the optimal storage space utilization when used in the following cases:

- **Autoextend-enabled tablespaces.** Oracle ASM-managed datafiles with the autoextend feature grow as the need for tablespace grows. Using such auto-extendable datafiles on thin provisioned volumes allows the system to allocate disk space to a database as it grows.

However, Oracle's process of extending a tablespace is I/O intensive and can affect the performance of the database during the file extension. Lessen the frequency of the performance impact by increasing the increment of space that Oracle adds (the `AUTOEXTEND ON NEXT` parameter of the `CREATE TABLESPACE` command). Set autoextend to a larger value for rapidly growing datafiles, and a smaller value if the data is growing at a slower rate.

- **ASM disk groups as archive log destinations.** The combination of Oracle ASM and HP 3PAR Thin Provisioning enables an expanding ASM archive log destination. Oracle databases become inaccessible when the archive log destinations fill up. Put Oracle archive logs on thin provisioned ASM disk groups, which allows the underlying storage to self-tune, accommodating unexpected increases in log switch activity. After a level 0 backup, remove the archive logs and use the Oracle ASM Storage Reclamation Utility (ASRU, described later) to free up the storage that was allocated to the old logs.

HP 3PAR Thin Provisioning may not be the best option for the following:

- **Datafiles residing on file systems.** When an Oracle datafile is first created, Oracle initializes the data blocks with a block headers and other metadata. Due small block size of most databases, there are usually no contiguous ranges of zero blocks for space reclamation. This causes a TPVV to provision all of its space, nullifying the value of thin provisioning.
- **Systems with high file system utilization.** If the file systems or ASM disk groups are full, then the benefits of thin provisioning are reduced. Consider any ASM disk group or file systems with utilization rates above 80 percent as inappropriate for use with thin provisioning. In this case, it may be more efficient to use fully provisioned virtual volumes to hold this data.
- **Oracle datafiles that are not in "autoextend" mode.** Oracle databases write format information to datafiles during tablespace creation. This has the same effect as provisioning file systems with high utilization and may be inefficient depending upon the ratio of provisioned storage to database size.

Microsoft SQL Server

Microsoft® SQL Server transaction log files are formatted when a database is created. The pages written mostly contain contiguous zeros and therefore the HP 3PAR Thin Persistence Software will keep the allocation of the TPVV to a minimum. Over time, as transactions are written to the log files, the log file TPVV allocation will grow so the CPG associated with log files should be appropriately sized.

As of Microsoft SQL Server 2005 the database creation phase no longer initializes datafiles if the account under which the SQL Server service is running, has the “Perform volume maintenance tasks” user rights assignment under the local security policy. The Instant File Initialization process creates a sparse file which is thin friendly and does not allocate storage space in advance.

Microsoft SQL Server 2012 is able to take advantage of the SCSI UNMAP support in Windows Server 2012 to automatically reclaim space after shrink/drop datafile operations. In addition when datafiles and databases are deleted on Windows Server 2012, the system automatically reclaims the space without administrator intervention. Prior to Windows Server 2012 space reclamation required a zerofile tool to write zeros to the free space in the file system.

Microsoft Hyper-V

The virtual hard disk (VHD) and the new virtual hard disk (VHDX) formats of Microsoft Hyper-V are fully compatible with HP 3PAR Thin Provisioning. Both formats offer a choice of fixed or dynamic sizing. A fixed VHD has an allocated size that does not change whereas a dynamic VHD will expand as data is written to it.

Although it may seem natural to use a full provisioned VV with a fixed VHD, contiguous zeros are written to the space allocated so they are ideal candidates for thin provisioning as the zeros will be reclaimed by the HP 3PAR Thin Persistence software.

In the past fixed VHDs were recommended for production instead of dynamic VHDs as their I/O performance was higher. With the new VHDX format introduced in Windows Server 2012 performance of dynamic VHDs has been significantly improved and there is now little difference between the two VHD types.

In addition, VHDX disks report themselves to the guest operating systems as being “thin provision capable”. This means that if the guest operating system is UNMAP capable it will be able to send UNMAPs to the VHDX file, which will then be used to ensure that block allocations within the VHDX file are freed up for subsequent allocations as well as forwarding the UNMAP requests to the physical storage.

SAP

By migrating data from traditional arrays to HP 3PAR StoreServ via Thin Conversion, legacy SAP systems can reduce up to 80 percent of the capacity in the storage environment. In an SAP system, data gets moved around or deleted within system storage volumes. HP 3PAR Thin Persistence enables that the thin volumes used by SAP systems stay efficient as possible by reclaiming unused space associated with deleted data.

Typically, SAP databases store data in the form of a matrix (tables and indexes) consisting of rows and columns. Most of the columns in the tables are of fixed length and there are often leading or trailing zeroes in these columns. The HP 3PAR ASIC is able to detect these zeros if they form a contiguous 16 KB block and prevent storage from being allocated. In testing with SAP ERP6 IDES a LUN with HP 3PAR Thin Persistence software enabled consumed 10% less storage space than a traditional thin LUNs after the initial database creation. See the [HP 3PAR StoreServ Storage for SAP Systems White Paper](#) for further details.

VMware vSphere

VMware vSphere has its own thin provisioning capabilities and as HP 3PAR StoreServ also supports thin provisioning, the question often arises on which level to implement thin provisioning when them together. The recommendation from HP is to use both vSphere and HP 3PAR Thin Provisioning.

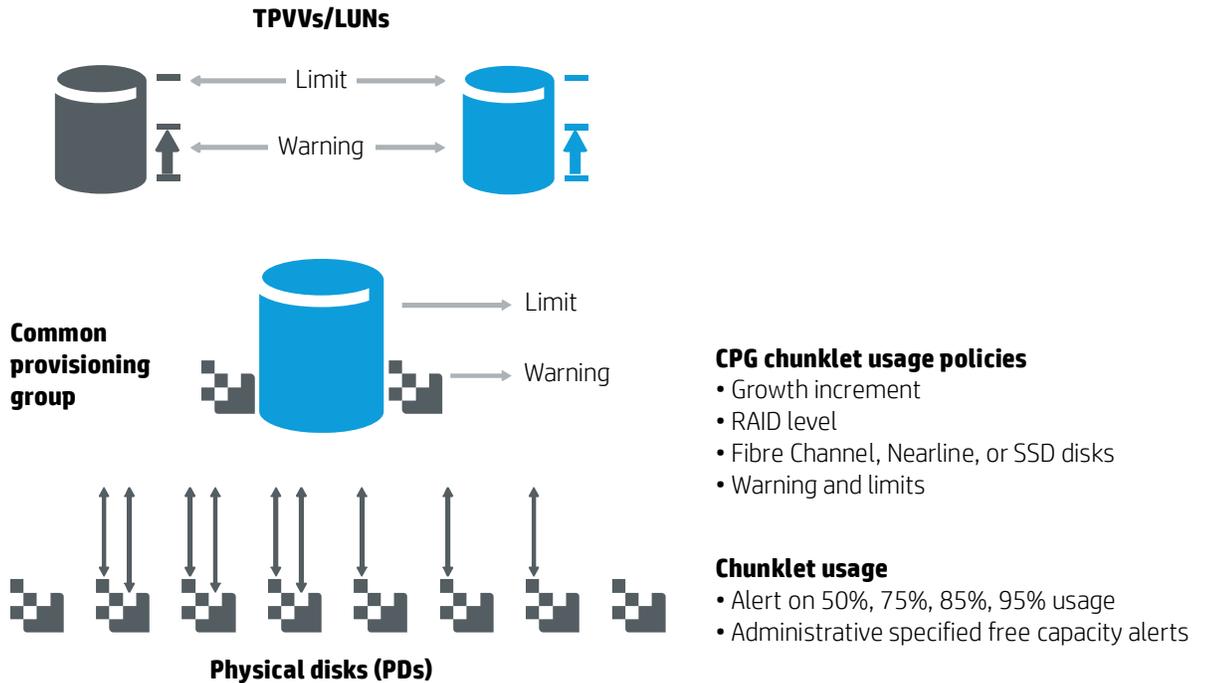
When creating VMs, there are several options for the file system layout of the VMDK files. By default VMDK files are created with the “Thick Provision Lazy Zeroed” option which means the VMDK is sparsely populated so not all the blocks are immediately allocated. When a guest VM reads from the unallocated areas of the VMDK the ESX server detects this and returns zeros rather than reading from disk. This VMware thin provisioning capability enables the oversubscription of the sizes of the VMDK files within the datastore.

For performance-intensive environments, VMware recommends using “Thick Provision Eager Zeroed” (EZT) virtual disks. These EZT disks have lower runtime overhead but require zeros to be written across all of the capacity of the VMDK at the time of creation. On traditional arrays, this VMDK format would negate all the benefits of thinly provisioned LUNs as all of the physical storage is allocated when the volume is zero-filled during creation. However, HP 3PAR Thin Persistence software allows clients to retain thin provisioning benefits when using EZT VMDKs without sacrificing any of the performance benefits offered by this VMDK option. In this case, Thin Persistence ensures that, when a new EZT volume is created, the entire volume is not allocated from physical storage since all zeros that have been written to the VMDK were intercepted and discarded by the HP 3PAR ASIC.

Managing thin storage

Overprovisioning capacity to hosts is one of the prime reasons why people choose to create thin provisioned volumes. Natural growth or a process writing excessive amounts of data to VVs can cause a CPG or the entire array to run out of space potentially causing an interruption of service for applications consuming space from that CPG or array. While HP 3PAR Thin Provisioning software is simple to use, a certain level of management is required to maximize its benefits. The HP 3PAR OS provides multiple categories of alerts that notify storage administrators of important events to ensure the smooth running of an HP 3PAR StoreServ. Figure 2 illustrates the categories for which warnings can be set.

Figure 2. HP 3PAR Thin Provisioning alert options



These include:

- Allocation warnings and limits for TPVVs
- Growth warnings and limits for CPGs
- Used physical capacity alerts
- Free physical capacity alerts

Allocation warnings

Allocation warnings provide a mechanism for informing storage administrators when a specific capacity threshold is reached. An allocation warning can be specified independently for each TPVV and each CPG. It is recommended that allocation warnings be used, at least on the CPG level, and acted upon when they get triggered.

The relevant CLI commands for setting allocation and growth warnings are:

- `setvv -usr_aw <percent> <TPVV>:` sets the allocation warning for the user space of the TPVV as a percentage of the TPVV size
- `setvv -snp_aw <percent> <TPVV>:` sets the allocation warning for the snapshot space of the TPVV as a percentage of the TPVV size
- `setcpg -sdgw <num> <CPG>:` sets the growth warning for the CPG in MB (append to the value num "g" or "G" for GB or "t" or "T" for TB)

These warnings can be changed at any time and are effective immediately. The CLI commands `showvv -alert` and `showcpg -alert` lists the allocation warnings that were set per TPVV and CPG.

Allocation limits

Applications sometimes get into an abnormal state writing data continuously to the storage device. Allocation limits provide a mechanism to prevent such “runaway” applications from consuming disk capacity beyond a specified threshold. Allocation limits can be specified independently for each TPVV and each CPG. For a TPVV, after the allocation limit is reached, the capacity allocated to the TPVV stops growing and new writes by the application fail. Similarly, for a CPG, after the allocation limit is reached, the automatic creation of new LDs, if configured, is disabled.

The relevant CLI commands related to setting allocation and growth limits are:

- `setvv -usr_al <percent> <TPVV>`: sets the allocation limit for the user space of the TPVV as a percentage of the TPVV size
- `setvv -snp_al <percent> <TPVV>`: sets the allocation limit for the snapshot space of the TPVV as a percentage of the TPVV size
- `setcpg -sdgl <num> <CPG>`: sets the growth limit of the CPG in MB (append to the value num “g” or “G” for GB, or “t” or “T” for TB)

These alerts can be changed at any time and are effective immediately. The CLI commands `showvv -alert` and `showcpg -alert` list the allocation limits that were set per TPVV and CPG.

The VV allocation limits and warnings can be set with the HP 3PAR Management Console by selecting Show advanced options checkbox when creating or editing a VV as shown in figure 3.

Figure 3. VV allocation limit and warning options

The CPG growth limits and warnings can be set with the HP 3PAR Management Console by selecting Show advanced options checkbox when creating or editing a CPG as shown in figure 4.

Figure 4. CPG allocation limit and warning options

The screenshot displays two sections: Allocation Settings and Growth Settings. Allocation Settings includes fields for Device Type (FC), RAID Type (RAID 1), Device RPM (<Default>), Set Size (2 data), Availability (Cage (Default)), Step Size (256 KIB), and Preferred Chunklets (Fast). Growth Settings includes Growth Increment (set to <Default> and GIB), Growth Limit (checkbox 'Enabled' and a GIB input field), and Growth Warning (checkbox 'Enabled' and a GIB input field). Each input field has a help icon (question mark).

It is important to note that the growth limit for a CPG is a hard limit and the CPG will not grow beyond it. Any VVs that require more space from the CPG, after the hard limit is reached they will not be able to grow and will eventually present write errors to host systems until the CPG allocation limit is raised. Because of this, it is recommended that TPVV, CPG, and free space warnings and limits be set to sensible levels and managed when they are triggered. As an example, the CPG warning limit should be set sufficiently below the CPG allocation limit so that it alerts the storage administrator with ample time to react before the CPG allocation limit is reached.

Remaining physical capacity alerts

As available physical capacity across the HP 3PAR StoreServ Storage gets consumed by either traditional VVs or TPVVs, or both, preconfigured alerts are generated at 50 percent, 75 percent, 85 percent, and 95 percent of physical capacity in use per disk type (FC, NL, or SSD). Furthermore, the storage administrator can use the CLI command `set sys` as follows to set another warning level when the available space within the system falls below a custom-defined capacity point:

- `set sys RawSpaceAlertFC <value>` where `value` is the remaining capacity on FC disks in GB
- `set sys RawSpaceAlertNL <value>` where `value` is the remaining capacity on NL disks in GB
- `set sys RawSpaceAlertSSD <value>` where `value` is the remaining capacity on SSD disks in GB

These serve as array-wide, advance warnings to the storage administrator to plan for and add necessary physical capacity. The alerts generated should be monitored and promptly acted upon to prevent all free space of a particular disk type be consumed.

Remaining CPG free space alerts

HP 3PAR OS samples the space available to CPGs once per day. The history of used and free CPG space is stored in an internal table and can be displayed using the `-hist` option in the `showspace` and `showcpg` commands. An alert is automatically generated if the available free space for a CPG falls below the CPG warning limit or the CPG allocation limit.

View logged warning and limit alerts

All warning and limit alerts mentioned above can be viewed in several ways:

- The CLI commands `showalert` and `showeventlog` list the alerts in various formats and with various options
- The HP 3PAR Management Console shows the alerts in the “Events” section
- SMI-S integrated in HP 3PAR OS provides asynchronous notification of events for changes in the elements managed by the CIM server. A CIM client can subscribe to selected CIM Indications to receive event notifications from the CIM server.
- The SNMP agent within HP 3PAR OS allows for retrieving the alerts by remote SNMP clients.

Alerts can be forwarded (`set sys RemoteSysLogHost`) to a Log Host for viewing them in an Enterprise Management application like HP OpenView.

User recommendations

The monitoring of alerts for available capacity by storage administrators and internal business processes are a critical component of a successful HP 3PAR Thin Provisioning management and administration strategy. You should nominate a primary and if possible a backup storage administrator for each site with HP 3PAR StoreServ equipment. The storage administrator's roles include:

- Proactively monitor free space availability per TPVV and CPG
- Proactively monitor consumption rates for TPVVs and CPGs
- Proactively monitor consumed TPVV capacity and compare to licensed thin provisioning capacity
- Proactively monitor physical capacity thresholds for each disk type and for the entire array
- Ensure adequate purchasing and installation of additional physical disk capacity buffer and thin-provisioning license upgrades in a timely manner
- Nominate an escalation contact who has proper authority to drive the customer responsibilities outlined in this document if the nominated storage administrators fail to carry out their responsibilities.

If you have a network connection with HP 3PAR Central via the Service Processor, the health of the HP 3PAR StoreServ will be proactively monitored for CPG growth problems and you can request to receive thin provisioning and other alerts by mail or via phone. You retain responsibility for managing the thin-provisioning capacity and CPGs; HP is not responsible for any failure when thresholds are met or exceeded.

Tracking volume space usage

TPVVs consume user space, admin space, and possibly snapshot space on the disk array. The following sections provide the CLI commands needed to determine how much space of every type is being consumed within the array. The output of these CLI commands show the reserved and the raw reserved space. The reserved space is what is offered by the array as usable space to the host. This value is also shown in the HP 3PAR Management Console in the Reserved User Size column for a TPVV and in the pie chart for the Logical option in the Summary tab for the virtual volume details screen. The raw reserved space is calculated from the reserved space by multiplying the latter by its RAID overhead factor. For example, this factor has a value of 2 for RAID 1 and 8/7 for RAID 5 with a set size equal to 8. The HP 3PAR Management Console shows the Raw Reserved space in the pie chart for the Raw option in the Summary tab for the virtual volume details. In chargeback models, most IT Departments bill their customers on the amounts of raw reserved space consumed.

TPVVs

Use the `showvv -s -p -prov tpvv` command to see how much admin, user and snapshot space is used by each TPVV. The reserved totals show how much space has been allocated whereas the used totals show how much of the space is currently in use by the VV. A significant difference between the space in use and the reserved space would indicate that space reclaim has been initiated on the VV and the reserved space will decrease over time as the space is reclaimed in the background. This is an example of the `showvv -s` output:

```
cli% showvv -s -p -prov tpvv
      ---Adm---  -----Snp-----  -----Usr-----
      --(MB)---  --(MB)---  -(% VSize)--  ----(MB)-----  -(% VSize)--  -----(MB)-----

  Id Name  Prov Type Rsvd Used Rsvd Used Used Wrn Lim  Rsvd  Used Used Wrn Lim  Tot_Rsvd  VSize
1167 vol101 tpvv base 256 127  0  0 0.0 -- -- 190464 103319 50.4  0  0  190720 204800
1262 vol102 tpvv base 128  9  0  0 0.0 -- -- 13312 11980  5.8  0  0  13440 204800
1265 vol103 tpvv base 128  1  0  0 0.0 -- --  512  65  0.0  0  0  640 204800
1266 vol104 tpvv base 128 15  0  0 0.0 -- -- 25088 22320 10.9  0  0  25216 204800
1267 vol105 tpvv base 128  1  0  0 0.0 -- --  512  1  0.0  0  0  640 204800
1268 vol106 tpvv base 256 16  0  0 0.0 -- -- 25600 22495 11.0  0  0  25856 204800
1269 vol107 tpvv base 128  1  0  0 0.0 -- --  512  1  0.0  0  0  640 204800
-----
      7 total 1152 170  0  0 256000 160181 257152 1433600
```

The same information can be displayed in the HP 3PAR Management Console by selecting Used User Size column when viewing the VV provisioning as shown in figure 5.

Figure 5. VV allocation sizes

| Name | State | Provisioning | RAID | Virtual Size (GiB) | Reserved User Size (GiB) | Reserved User Size (% Virtual) | Used User Size (GiB) | Reclaimed User Space (GiB) |
|-------|--------|--------------|--------|--------------------|--------------------------|--------------------------------|----------------------|----------------------------|
| vol01 | Normal | Thin | RAID 5 | 200,000 | 186,000 | 93% | 100,889 | 15,000 |
| vol02 | Normal | Thin | RAID 5 | 200,000 | 13,000 | 6% | 11,717 | 0,000 |
| vol03 | Normal | Thin | RAID 5 | 200,000 | 0,500 | 0% | 0,500 | 0,000 |
| vol04 | Normal | Thin | RAID 5 | 200,000 | 24,500 | 12% | 21,950 | 0,000 |
| vol05 | Normal | Thin | RAID 5 | 200,000 | 0,500 | 0% | 0,500 | 0,000 |
| vol06 | Normal | Thin | RAID 5 | 200,000 | 25,000 | 12% | 21,968 | 0,000 |
| vol07 | Normal | Thin | RAID 5 | 200,000 | 0,500 | 0% | 0,500 | 0,000 |

Note that the bar shows the Reserved User Size as percentage whereas the showvv command shows the Used User Size as a percentage. The Management Console view is more accurate as it indicates how much data is actually allocated.

CPGs

Space in use on the array can be tracked per CPG. The showcpg -r command shows the user, snapshot, and admin space in Used and Raw Used amounts. You can work out the unallocated space within the CPGs by subtracting the used space from the Totals listed.

In addition to showing the CPG usage the showspace -cpg command will also show how much LD space may still be created, given the amount of free chunklets in the system and the CPG parameters (e.g. RAID level, HA level, device types, etc.):

```
cli% showspace -cpg *
----- (MB) -----
CPG -----EstFree----- -----Usr----- -----Snp----- -----Adm-----
Name RawFree LDFree Total Used Total Used Total Used
SSD_r5 0 0 3531136 3531136 17024 0 34816 33248
FC_r5 176078848 154068992 21951616 21951616 104320 3072 32768 8480
NL_r6 54263808 40697856 27969280 27969280 96512 69632 23552 21248
```

System space

Not all space on the physical disks can be used for storing your data. A small portion of the space on the array is dedicated to volumes with an administrative function.

There is a fully provisioned volume named admin that is used to store system administrative data such as the System Event Log. The logging LDs, starting with the name log, are used to store data temporarily during physical disk failures and disk replacement procedures. There are also preserved data space logical disks (PDSLDS) which start with the name pdsld. Preserved data is the data moved from the system’s cache memory to the PDSLDS space in the eventuality of multiple disk or cage failures. On HP 3PAR StoreServ systems running HP 3PAR OS 3.1.2, the HP 3PAR System Reporter software is integrated into the OS and executes on the controller nodes and the database files are stored in a fully provisioned volume called .srdata.

The amount of raw disk space consumed by these system functions is summarized in table 3.

Table 3. System space usage⁵

| Number of nodes | Logging LDs | Admin LDs | Preserved data ⁶ | Srdata |
|-----------------|-------------|-----------|-----------------------------|--------|
| 2 | 240 GB | 20 GB | 128 GB | 60 GB |
| 4 | 480 GB | 20 GB | 256 GB | 80 GB |
| 6 | 720 GB | 20 GB | 384 GB | 100 GB |
| 8 | 960 GB | 20 GB | 512 GB | 100 GB |

Total used space

The `showsys -space` command gives the total raw space in use on the system with a breakdown of how the space is allocated. This is an example of `showsys` output:

```
cli% showsys -space
-----System Capacity (MB)-----
Total Capacity          : 5588480
  Allocated              : 1009920
    Volumes              : 180224
      Base Volumes      : 0
        User             : 0
          Copy           : 0
            Admin        : 0
      CPGs (TPVVs & CPVVs) : 180224
        Copy             : 131072
          Used           : 1024
            Unused      : 130048
          Admin          : 49152
            Used         : 384
              Unused    : 48768
        Unmapped        : 0
      System             : 829696
        Internal         : 114688
        Spare            : 715008
          Used           : 0
            Unused      : 715008
      Free               : 4578560
        Initialized     : 4578560
        Uninitialized    : 0
        Unavailable     : 0
        Failed          : 0
```

The value in the System section of the output is the raw space in use by the admin and the `.srdata` volume, the PDSLDS, the logging LDs and the spare chunklets. The internal entry lists the space associated with the local boot disks of the nodes so this should be treated separately from the space allocated from the regular drives.

⁵ These are maximum values. The exact values will vary depending on the HP 3PAR StoreServ model and disk configuration.

⁶ The total capacity of the PDSLDS is equal to the sum of all data cache memory located in the controller nodes of the system.

Running out of space

The following sections describe the severity of the situation when space limitations are met, from the least critical (for example, running out of space on a single TPVV), to the more severe (for example, running out of space on a CPG or on the entire system).

Running out of space in a TPVV

If an allocation limit was set for a TPVV, the array stops all writes to that TPVV after the limit is reached—returning SCSI error codes if more writes come in. Depending on the OS and application, the effect of this can range from the host application stalling to the host server crashing. Reaching a TPVV limit purely affects that volume and its resident application(s).

TPVV limits can be changed at any time with the `setvv` command. After an appropriate change, writes to the TPVV continue to be processed.

Running out of space in a CPG

Reaching a CPG limit has the same effect as reaching a TPVV limit; however, many more TPVVs and other types of VVs may be impacted. CPG limits can be changed at any time with the `setcpg -sdg1` command. After an appropriate change, writes to the volumes in the affected CPG continue to be processed.

Running out of space on the system

Running out of space on the system will affect TPVVs and all other types of volumes being written to in all CPGs. In the unlikely scenario that all physical capacity becomes used, the system prevents new writes from occurring until more capacity is added.

Usage reporting and trend analysis

The CLI command `showcpg -hist <CPG>` gives a daily account of CPG usage split into user, snapshot, and admin space. The command `showspace -cpg <CPG> -hist` also shows this information.

You can also use the `srcpgspace` and `srvvspace` commands to query the internal System Reporter database.⁷ Additionally, the optional HP 3PAR System Reporter software has the ability to track the CPG and TPVV usage for comprehensive usage and trend analysis.

Reclaiming Unused Space

The HP 3PAR OS space consolidation features allow you to change the way that virtual volumes are mapped to logical disks in a Common Provisioning Group (CPG). Deleting TPVVs or moving virtual volume regions from one logical disk to another enables you to compact logical disks, and free up disk space so that it can be reclaimed for use by the system.

Mapping is the correspondence of Logical Disk (LD) regions to the virtual volume regions. Virtual volumes are made up of multiple logical disks, and each logical disk contains regions that are mapped to the virtual volume. All types of volumes are created by mapping data from one or more logical disks to the virtual volume.

Logical disks can be shared by multiple virtual volumes. As volumes are deleted or as volume copy space grows and then shrinks, logical disks can use space less efficiently. When logical disks do not efficiently use space, the unused space consumes regions on the LD that are not available for use by the system when creating new logical disks. The space management features enable you to consolidate used space onto fewer fully-used logical disks so that unused regions are forced onto one or more logical disks that are then deleted. Deleting these logical disks frees the unused space for general use by the system. You can also truncate LDs to free up space. The LD's used regions are compacted by moving them to the beginning of the LD and then the LD is shortened so that unused space can be returned to the system's free chunklet pool.

⁷ The System Reporter functionality has been included in HP 3PAR OS 3.1.2. Prior versions of HP 3PAR OS will require the optional HP 3PAR System Reporter software on an external host.

Reclaiming Unmapped Logical Disk Space from CPGs

Common Provisioning Groups (CPGs) provide a shared pool of logical disk capacity for use by all virtual volumes that draw space from that pool. If volumes that draw from a CPG are deleted, or if copy space for these volumes grows and then shrinks, the underlying logical disks in the CPG pool can become less efficient in space usage. One or more logical disks in the CPG pool may have only a small portion of their regions mapped to existing virtual volumes. However, the logical disk's unused regions are not available for use by the volumes mapped to the CPG.

Compacting the logical disk regions mapped to these volumes may recover and free logical disk space. Compacting a CPG allows you to reclaim space from a CPG that has become less efficient in space usage from creating, deleting, and relocating volumes. Compacting consolidates logical disk space in CPGs into as few logical disks as possible. Compacting CPGs can be performed with both the HP 3PAR Command Line Interface (CLI) and the HP 3PAR Management Console.

When using thin provisioning it is recommended to schedule regular CPG compactions during periods of low activity. To setup a daily schedule of a CPG compaction at midnight, execute the following CLI command:

```
cli% createsched "compactcpg -f -pat "*" "0 0 * * *" compactcpg
```

Note:

If using Adaptive Optimization on all the CPGs, scheduling a CPG compaction may not be required. Automatic CPG compaction is a configurable option of Adaptive Optimization.

Automatically Reclaiming Unused Snapshot Space from Volumes

The HP 3PAR OS automatically reclaims unused snapshot and administration space from Thinly-Provisioned Virtual Volumes (TPVVs) and fully-provisioned virtual volumes and returns the space to the LDs. The system examines the snapshot and administration space for large areas of unused space. The identified areas are unmapped from the corresponding LD regions and the space is returned to the LDs.

Deleted Volume Snapshot Space

HP 3PAR Thin Copy Reclamation is an optional feature that reclaims space when a snapshot of a virtual copy, full copy, or remote copy volume is deleted from the system. The deleted snapshot space is reclaimed from either a standard or thin provisioned volume and is automatically returned to its CPG for reuse by other volumes. The HP 3PAR Thin Copy Reclamation feature works on all HP 3PAR StoreServ and requires that an HP 3PAR Thin Provisioning software license be present on the array.

Logical Disks and Chunklet Initialization

After Logical Disks are deleted or space is reclaimed by the thin persistence software, the underlying chunklets are reinitialized before their space is available to be used by other LDs. The initialization process for chunklets happens in the background and at a low priority so as not to impact system performance. To see chunklets that are currently in the process of being initialized, issue the `showpd -c` command. Chunklets that are uninitialized are listed in the `Uninit` column.

Free page consolidation

As individual 16 KB pages are freed the space in the LDs can become fragmented and therefore the HP 3PAR OS implements a defragment process similar to an operating system that may run on a file system. There is a thread that periodically examines the VVs to determine whether defragmentation should be initiated. The data from partially allocated regions is then reorganized to create larger contiguous blocks of free space that can be reclaimed.

HP 3PAR Thin Persistence software

Traditionally when data is deleted on a host the operating system will report that space has been freed but the storage is not informed that the data is no longer in use. With fully provisioned volumes this is not an issue but with thin provisioned volumes the unused space will remain allocated on the array causing the volumes to grow over time. This creates a hidden utilization penalty that can significantly reduce the space savings of thin provisioning.

HP 3PAR Thin Persistence software is able to maintain the benefits of HP 3PAR Thin Provisioning software and HP 3PAR Thin Conversion software by enabling “thin reclamation” of allocated but unused capacity so the thin volumes are able to stay as lean and efficient as possible. It leverages the HP 3PAR OS support for the WRITE_SAME and UNMAP commands of the SPC-3 standard and the unique zero-detection capabilities of the HP 3PAR ASIC to give HP 3PAR StoreServ Storage the power to reclaim unused space associated with deleted data simply, quickly, and non-disruptively.

UNMAP is a SCSI command that a host can issue to tell the storage that blocks no longer need to be allocated. This is particularly important in thinly provisioned environments as it allows the storage array to recognize these blocks are not used and to return them to the free capacity pool for reuse by other volumes.

The HP 3PAR ASIC features an efficient, silicon-based zero-detection mechanism. This unique hardware capability gives HP 3PAR StoreServ Storage the ability to remove allocated but unused space as small as 16 KB on the fly without impacting performance.

In addition, the benefits of HP 3PAR Thin Persistence are available to read/write snapshots of TPVVs. The mechanism for reclamation is the same as for the parent TPVV: writing zeros to the allocated but unused space in a read/write snapshot will trigger the ASIC to initiate reclamation of the deleted space. To benefit from thin reclamation, the zero-detect policy needs to be enabled on each read/write snapshot.

To use HP 3PAR Thin Persistence software, the array also needs a license for HP 3PAR Thin Provisioning.

Thin persistence methods

The most optimal thin persistence method is for the host operating system to issue an UNMAP for the unwanted blocks when the data is deleted. Typically this would be done by the file system although if the operating system offered a suitable API an application could directly communicate that it no longer needed some data. This method provides continual reclaiming of space to allow the storage volumes to stay thin. The disadvantage is it requires significant changes to the storage stack and only the most modern operating systems have implemented native UNMAP support.

Another method is to have an operating system utility that will examine the blocks of a volume and issue UNMAPs for those that are not being used. This type of utility also requires host operating system UNMAP support but to a lesser extent than the continual method and they are specific to a file system or volume manager type. Most of these utilities can be run when the data is online but as they generate the UNMAP requests for all the unused blocks in one go they are generally run manually during an outage window or scheduled to run during a quiet period so the reclaims do not adversely impact other workloads on the storage.

The final method is to reclaim space using a “zerofile” utility that writes zeros to all allocated but unused space on a file system. On HP 3PAR StoreServ systems the zero-detection capability of the HP 3PAR ASIC intercepts the blocks of zeros being written and automatically triggers the reclamation of the space. The advantage of this method is that it does not require any special operating system support and the utilities to generate zerofiles are often supplied with the base operating system distribution. To achieve the good reclaim rates the utilities need fill the majority of the available free space so they are generally run manually during an outage window or scheduled to run during a quiet period to avoid applications failing due to a lack of space.

For the zerofile utilities to work the zero-detect policy needs to be set for each TPVV. Blocks of 16 KB of contiguous zeros are freed and returned for reuse by the VV; if 128 MB of space is freed it is returned to the CPG for use by other volumes.

Thin Reclamation for Microsoft Windows Server 2012

Windows Server 2012 integrates very well with thin provisioned volumes on HP 3PAR StoreServ. It identifies thinly provisioned volumes on HP 3PAR StoreServ systems, writes entries in the Windows Event Log file when storage thresholds are reached on the CPG and the TPVV level and supports active reclaim by issuing UNMAPs upon file deletion or file shrinking on thin provisioned volumes on NTFS formatted volumes. The standard Defrag scheduled task also automatically reclaims storage.

In addition, Windows Server 2012 extends UNMAP support to the virtual layer. The Hyper-V VHDX disks report themselves to the guest operating systems as being “thin provision capable”. This means that if the guest operating system is UNMAP capable it will be able to send UNMAPs to the VHDX file, which will then be used to ensure that block allocations within the VHDX file are freed up for subsequent allocations as well as forwarding the UNMAP requests to the physical storage.

There is also a File TRIM API which is mapped to the TRIM command for ATA devices and the UNMAP command for SCSI devices. TRIM hints allow the application to notify the storage that blocks that previously were allocated are no longer needed and can be reclaimed.

In summary, the following operations trigger storage space reclamation in Windows Server 2012:

- Deletion of a file from a file system on a thin provisioned volume
- Running Storage optimization, a new feature of Windows Server 2012 disk management
 - You can use manual or automatic scheduling of the Optimize operation utility
 - The standard Defrag scheduled task automatically runs Optimize
- UNMAP requests from a Hyper-V guest operating system
 - Deleting a file from the file system of a UNMAP capable guest operation system sends UNMAP requests to the driver stack of the Hyper-V host
- UNMAP requests from applications using the TRIM API

The default behavior of issuing UNMAPs can be disabled on a Windows 2012 server by setting the “disabledeletenotify” parameter of the `fsutil` command. This will prevent reclaim operations from being issued against all volumes on the server.

To disable reclaim operations run the following PowerShell command:

```
fsutil behavior set DisableDeleteNotify 1
```

Thin Reclamation for Microsoft Windows Server 2003 and 2008

Windows Server versions prior to 2012 do not implement UNMAP support and therefore to reclaim thinly provisioned storage you must leverage the zero detection capabilities of HP 3PAR StoreServ.

Microsoft provide a Sysinternals advanced system utilities suite that includes the SDelete (Secure Delete) application that can used to overwrite a deleted file's on-disk data to make disk data unrecoverable.

As well as overwriting a single files data space, SDelete can indirectly overwrite free space by allocating the largest file it can and then performing a secure overwrite to ensuring that all the disk space that was previously free becomes securely cleansed. You can utilize this feature of SDelete to perform thin reclamation on zero-detect enabled HP 3PAR StoreServ volumes by specifying the `-z` flag when running SDelete to writes zeros to the free space.

One disadvantage of SDelete is that it does not support mount points so a volume must have a drive letter associated with it before the utility can be run. Using the `subst` command one can temporarily attach a drive letter to a mount point before running SDelete.

It is recommended that applications are shutdown before running SDelete as it can cause a `file_system_full` condition due to consuming all “free” space. An alternative solution is to create a PowerShell script that uses `fsutil` to create a balloon file that is limited to a certain percentage of the “free” space.

This is an example of using `fsutil` to zero 100MiB of space:

```
fsutil file createnew zerotemp.txt 104857600
fsutil file setvaliddata zerotemp.txt 104857600
fsutil file setzerodata offset=0 length=104857600 zerotemp.txt
del zerotemp.txt
```

Thin Reclamation for VMware vSphere

vSphere 5.0 introduced automatic space reclamation upon deletion of a VMDK. However, due to varied response times from the storage devices, it was found that the UNMAP command could result in poor performance and VMware have recommended that this feature should be disabled on ESXi 5.0 hosts. As an alternative ESXi 5.0 Update 1 included an update to the `vmkfstools` command that provided a new option (`-y`) to issue UNMAPs to reclaim unused space on a thin provisioned volume. This new implementation is not an automated process and will need to be run manually. The command should be run after changing to the root directory of the desired VMFS volume:

```
# cd /vmfs/volumes/<volume-name>
# vmkfstools -y <percentage of free space to reclaim>
```

The `vmkfstools -y` command creates a temporary “balloon” file in the datastore which can be up to the size of the free space available in the datastore and then it issues UNMAPs for the blocks of the balloon file. However, it is recommended that the reclaim percentage is not more than 60% as the resulting balloon file temporarily uses space on the datastore which could cause the deployment of new virtual disks to fail while the `vmkfstools` command is running. This is an important consideration as there is no way of knowing how long a `vmkfstools -y` operation will take to complete. It can be anywhere from few minutes to several hours depending on the size of the datastore and the amount of space that needs to be reclaimed.

There is an alternative method that can be deployed that takes advantage of the ASIC based zero-detect capability of HP 3PAR StoreServ. The `vmkfstools -y` command writes zeros to the balloon file and therefore if the volume has zero-detect enabled the blocks will have already been freed so the UNMAP phase is redundant. Instead run the `vmkfstools` command with the `-d eagerzeroedthick` option to create a zerofile which can then be removed:

```
# cd /vmfs/volumes/<volume-name>
# vmkfstools -c <size of space to reclaim> -d eagerzeroedthick zerofile
# rm zerofile-flat
```

The VMware hypervisor does not report disks as being thin provisioning capable to the guest OS, when using VMDKs therefore to reclaim thinly provisioned storage you must leverage the zero detection capabilities of HP 3PAR StoreServ. This means using standard file system tools (such as SDelete in Microsoft Windows, `dd` in UNIX®/Linux) to write zeros across deleted and unused space in a VM's file system. The zeros will be autonomically detected by the HP 3PAR ASIC and the disk space they were consuming will be freed up and returned to the thin provisioned volume.

If running ESX 4.1, it is strongly recommended to install the HP 3PAR VAAI plug-in before creating EZT VMDKs as it will speed up their creation by 10x to 20x.

Thin Reclamation for Linux

Linux uses the term discard for the act of informing a storage device that blocks are no longer in use. This is because it uses the same mechanism to support both TRIM commands on ATA SSDs and UNMAP commands on SCSI devices so discard was chosen as a protocol neutral name. Red Hat Enterprise Linux 6 was the first major distribution to support discards and offers both real-time and batched support.

Real-time support is offered by a ext4 file system when mounted with the `-o discard` option, the default is not to issue discards. When data or files are deleted on a discard enabled ext4 file system UNMAPs are generated to free up space on the thinly-provisioned virtual volume on the storage. The logical volume manager (LVM) and the device-mapper (DM) targets also support discards so space reclaim will also work on file systems created on LVM and/or DM volumes.

For example to mount the DM device `tpvv_lun` on `/mnt` with discards enabled run:

```
# mount -t ext4 -o discard /dev/mapper/tpvv_lun /mnt
```

This will cause the RHEL 6.x to issue the UNMAP command, which in turn causes space to be released back to the array from the TPVV volumes for any deletions in that ext4 file system. This is not applicable for fully provisioned virtual volumes.

In addition the `mke2fs`, `e2fsck`, and `resize2fs` utilities also support discards to ensure the TPVV volumes are optimized when administration tasks are performed.

There is also batched discard support available using the `fstrim` command. This can be used on a mounted file system to discard blocks which are not in use by the file system. It supports ext3 and ext4 file systems and will also rethin a file system not mounted with the discard option.

For example to initiate storage reclaim on the `/mnt` file system run:

```
# fstrim -v /mnt
/mnt: 21567070208 bytes were trimmed
```

The Linux swap code will also automatically issue discard commands for unused blocks on discard-enabled devices and there is no option to control this behavior.

Thin Reclamation for UNIX

On UNIX systems or Linux distributions that do not support discard you will need to reclaim space using a “zerofile” script that writes zeros to all allocated but unused space on a file system. The zero-detection capability of the HP 3PAR ASIC will intercept the blocks of zeros being written and automatically trigger the reclamation of the space.

The script would use the `dd` command to copy zero data blocks from the `/dev/zero` device to a file in the file system. However, it is recommended that the size of the space to zero is not more than 70% of the free space as a very large zerofile could cause the creation of new files to fail while the zerofile script is running.

Thin Reclamation for Veritas Storage Foundation

Hosts with VxFS file systems managed by the Veritas Volume Manager (VxVM) software can also reclaim space associated with file deletions or file shrinking on HP 3PAR StoreServ storage.

The space reclamation is not automatic, the VxFS file system informs the VxVM volume manager about deleted blocks and a VxVM command has to be manually run which send WRITE SAME SCSI commands with the UNMAP bit turned on to the HP 3PAR StoreServ. No tool to write zeros to the deleted space in the file systems is required for reclaiming the space.

The list of disks whose allocated but unused space can be reclaimed is given by the command executed on the host:

```
# vxdisk -o thin, fssize -u m list
```

This will display the VV allocated space and the file system usage.

The space reclamation is initiated by the VxVM `vxdisk` command:

```
# vxdisk reclaim [<disk>|<dg>|<encl>]
```

By default, the reclamation does not affect unmarked space, which is the unused space between subdisks. If a logical unit (LUN) has a lot of physical space that was previously allocated, the space between the subdisks might be substantial. Use the `-o full` option to reclaim the unmarked space.

```
# /opt/VRTS/bin/fsadm -V vxfs -R /<VxFS_mount_point>
```

To monitor the reclamation status, run the following command:

```
# vxtask list
```

| TASKID | PTID | TYPE/STATE | PCT | PROGRESS |
|--------|------|------------|--------|--------------------------------------|
| 171 | | RECLAIM/R | 00.00% | 0/41875931136/0 RECLAIM vol100 dg100 |

The `vxrelocd` daemon tracks the disks that require reclamation. The schedule for reclamation can be controlled using the `vxdefault` command. The `reclaim_on_delete_wait_period` parameter specifies the number of days after a volume or plex is deleted when VxVM reclaims the storage space. The default value is 1, which means the volume is deleted the next day. A value of -1 indicates that the storage is reclaimed immediately and a value of 367 indicates that the storage space can only be reclaimed manually using the `vxdisk` `reclaim` command. The `reclaim_on_delete_start_time` parameter specifies the time of day when VxVM starts the reclamation for deleted volumes and this defaults to 22:10.

To completely disable thin-reclaim operations, the add the parameter `reclaim=off` to the `/etc/vxdefault/vxdisk` file.

Thin Reclamation for Oracle databases

During an Oracle Automatic Storage Management (ASM) database lifecycle, the utilization of allocated storage capacity in a thin provisioned volume can decrease as changes are made to the database through common operations such as:

- Dropping of a tablespace or database upon deletion of transient data
- Resizing of an Oracle datafile upon shrinking a tablespace
- Addition of new disks to an ASM disk group to accommodate growth or load balance performance

These changes result in the creation of unused ASM disk space that can build up over time and although this space is available for reuse within ASM it remains allocated on the storage array. The net result is that the storage utilization on the array eventually falls below desirable levels.

To solve this problem, HP and Oracle have partnered to improve storage efficiency for Oracle Database 10g and 11g environments by reclaiming unused (but allocated) ASM disk space in thin provisioned environments. The Oracle ASM Storage Reclamation Utility (ASRU) is a stand-alone utility that works with HP 3PAR Thin Persistence software to reclaim storage in an ASM disk group that was previously allocated but is no longer in use. Oracle ASRU compacts the ASM disks,

writes blocks of zeroes to the free space, and resizes the ASM disks to original size with a single command, online and non-disruptively. The HP 3PAR StoreServ, using the zero-detect capability of the HP 3PAR ASIC, will detect these zero blocks and reclaim any corresponding physical storage.

You can issue a SQL query to verify that ASM has free space available that can be reclaimed:

```
SQL> select name, state, type, total_mb, free_mb from v$asm_diskgroup
where name = 'LDATA';
```

| NAME | STATE | TYPE | TOTAL_MB | FREE_MB |
|-------|---------|--------|----------|---------|
| LDATA | MOUNTED | EXTERN | 1023984 | 610948 |

Run the Oracle ASRU utility as the Oracle user with the name of the disk group for which space should be reclaimed:

```
# bash ASRU LDATA
```

```
Checking the system ...done
```

```
Calculating the new sizes of the disks ...done
```

```
Writing the data to a file ...done
```

```
Resizing the disks...done
```

```
/u03/app/oracle/product/11.2.0/grid/perl/bin/perl -I /u03/app/oracle/product/
11.2.0/grid/perl/lib/5.10.0 /home/ora/zerofill 5 /dev/oracleasm/disks/LDATA2
129081 255996 /dev/oracleasm/disks/LDATA3 129070 255996 /dev/oracleasm/disks/
LDATA4 129081 255996 /dev/oracleasm/disks/LDATA1 129068 255996
```

```
126928+0 records in
```

```
126928+0 records out
```

```
133093654528 bytes (133 GB) copied, 2436.45 seconds, 54.6 MB/s
```

```
126915+0 records in
```

```
126915+0 records out
```

```
133080023040 bytes (133 GB) copied, 2511.25 seconds, 53.0 MB/s
```

```
126926+0 records in
```

```
126926+0 records out
```

```
133091557376 bytes (133 GB) copied, 2514.57 seconds, 52.9 MB/s
```

```
126915+0 records in
```

```
126915+0 records out
```

```
133080023040 bytes (133 GB) copied, 2524.14 seconds, 52.7 MB/s
```

```
Calculating the new sizes of the disks ...done
```

```
Resizing the disks...done
```

```
Dropping the file ...done
```

Migration to thin-provisioned volumes

Traditionally, data migrations associated with a storage technology refresh have carried forward the poor utilization rates from one system to the other. Administrators had to procure at least the same amount of capacity on a new array as on the legacy systems from which data was being migrated, even if the migration involved a large amount of allocated but unused capacity. The HP 3PAR Thin Conversion software leverages the zero-detection capabilities built into the HP 3PAR ASIC to drive the conversion of inefficient, traditional storage volumes on legacy arrays to more efficient, higher-utilization thin provisioned volumes on any model of HP 3PAR StoreServ.

This “thinning” of volumes enables data centers to meet green IT targets, eliminate wasted space, and increase utilization rates. Because the fat-to-thin conversion mechanism is built into the array hardware, volume conversions take place inline and at wire speeds, while preserving service levels, and without causing disruption to production workloads. The use of HP Thin Conversion software is supported across all HP 3PAR StoreServ Storage systems and requires that an HP 3PAR Thin Provisioning license is present on the destination array. The zero-detect capability, required for thin conversion, has to be enabled on every destination TPVV.

Migrate from full to thin provisioned volumes across arrays

There are several approaches to consider when migrating data from a source array to thin provisioned volumes on an HP 3PAR StoreServ. Using file-based migration tools such as rsync in UNIX/Linux, Quest in Windows, or backup-and-restore, the target TPVV will grow only to the capacity of data contained within the file system, therefore making efficient use of thin provisioning on the target HP 3PAR StoreServ. A block-based migration tool such as a host-based logical volume manager or a SAN-based virtualization appliance will copy all blocks of the source volume to the TPVV, whether they are used or not. This means that the customer may not benefit at all from thin provisioning when using block-based migration tools. It is therefore recommended to run a zerofile tool in advance so that the allocated but unused space in the fat volume on the source array was filled with zeros, the HP 3PAR ASIC enables fat-to-thin block-based array-to-array migrations to reap the benefits of thin volumes immediately and under all circumstances.

Migrate between full and thin provisioned volumes on the same array

All HP 3PAR StoreServ systems running HP 3PAR OS 3.1.2 have the ability to make a convert from a thin-provisioned volume to a fully-provisioned volume (or vice versa) without requiring an offline transition.⁸

Converting from full to thin provisioning saves space for volumes that are sparsely consumed whereas thin to full provisioning saves on thin provisioning license usage for volumes that are mostly or completely allocated.

The `tunevv` command is used to convert between fully-provisioned and thinly-provisioned Virtual Volumes. In the following example, the logical disks used for user space are moved to CPG FC_r5 for virtual volume vol01 and the virtual volume is converted to a TPVV:

```
cli% tunevv usr_cpg FC_r5 -tpvv vol01
```

When the `-tpvv` or `-full` options for the `usr_cpg` subcommand are specified, the tune will automatically rollback on a failure. These options do not support virtual volumes with remote copy. These options will only convert virtual volumes using snapshots if the `-keepvv` option is used, but the snapshots will reside in the virtual volume specified by the `-keepvv` option.

During the thin conversion the HP 3PAR ASIC will assist in reducing the amount copied by using its zero-detect capability to remove the need to copy blocks of zeros. To make optimal use of this feature, it is advantageous to write zeros to the allocated but unused space on the fully provisioned volume prior to the conversion.

⁸ For earlier versions of HP 3PAR OS, this “thinning” operation does not complete online: a brief disruption of service is required to change the host mapping from the full to the thin provisioned volume. An HP 3PAR Full Copy (or Physical Copy) operation of the source volume is required for HP 3PAR OS 3.1.1 and prior, the license for which is included in HP 3PAR OS.

Conclusion

HP 3PAR StoreServ Storage is the only platform that offers a comprehensive thin strategy that not only allows storage to start thin, but to get thin and stay thin. HP 3PAR Thin Provisioning software is a simple yet powerful tool for improving the efficiency of storage. Following the best practices outlined in this paper will allow IT staff to maximize the benefit of HP 3PAR Thin Provisioning and do more with less. To supplement the dramatic savings of thin provisioning, HP 3PAR StoreServ features a unique HP 3PAR ASIC with thin capabilities built in and a range of software offerings that can save enterprises 50 percent or more on the cost of a storage technology refresh.

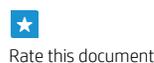
Together, HP 3PAR Thin Provisioning and the accompanying HP 3PAR thin technologies not only minimize upfront and ongoing storage costs, but also the cost of housing, powering, cooling, and managing storage so that enterprises can shift resources away from operations and apply them to innovation.

Fine grained virtualization is what separates HP 3PAR StoreServ from many other storage offerings. Having such tight control of resources decreases wasted capacity, increases both performance and utilization without tradeoffs. It also is why our storage array can be so flexible, allowing administrators the freedom to change drive types, service levels, and raid configurations non-disruptively.

When combined with Thin Provisioning, Thin Conversion is so powerful that HP even guarantees the ability to reduce storage capacity requirements by at least 50 percent with the Get Thin Guarantee. With this program, customers deploying HP 3PAR StoreServ as part of a storage technology refresh are guaranteed to reduce storage capacity requirements by 50 percent or more when using HP 3PAR Thin Conversion software to non-disruptively convert legacy storage from fully allocated, traditional volumes to new thin volumes on the HP 3PAR StoreServ Storage system. More information on this Get Thin Guarantee program is available at hp.com/storage/getthin.

Learn more at
hp.com/go/3PARStoreServ

Sign up for updates
hp.com/go/getupdated



© Copyright 2012–2013 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation. Oracle is a registered trademark of Oracle and/or its affiliates. UNIX is a registered trademark of The Open Group.

