WHITE PAPER: BEST PRACTICES



# *Start Thin*, *Get Thin*, and *Stay Thin* with Thin Provisioning Technology from Symantec and 3PAR

*February, 2009*

Symantec Technical Network White Paper

White Paper: Symantec Best Practices

# *Start Thin*, *Get Thin*, and *Stay Thin* with Thin Provisioning Technology from Symantec and 3PAR

**Contents**

## Introduction

Balancing growing storage requirements *and* limited IT resources is the challenging and often thankless job of IT management today. In the face of this challenge, IT organizations are increasingly turning to *thin provisioning* technologies to save significant amounts of capital. Thin provisioning is a key virtualization technology that has the potential to dramatically reduce capacity requirements and simplify operations in the datacenter. This is made possible by allowing the creation of virtual disks or LUNs that, to hosts and applications, appear much larger than the actual physical storage allocated to them. However, not all "thin" approaches available today use the same technologies or offer the same benefits.

3PAR® Thin Provisioning software enables 3PAR Utility Storage arrays to automatically allocate physical storage only when actual data is written to the virtual disk. This approach enables companies to slowly scale their storage investments rather than having to estimate their storage needs and purchase all of their storage up front. Purchasing and allocating storage up front leads to low utilization rates and results in the allocation of more capacity than a given project may ever need for actual written data. Simply put, 3PAR Thin Provisioning eliminates this waste. However, not all thin technologies available today can make this claim. Together, 3PAR and Symantec offer organizations an end-to-end thin provisioning strategy that employs 3PAR Thin Provisioning and Veritas Storage Foundation from Symantec to allow customers to not only *start thin*, but to *get thin* and to *stay thin*.

This white paper discusses the fundamental problem that thin provisioning was designed to address—that of *allocated-but-unused* storage. It explains the concepts behind thin provisioning and describes key benefits in detail. It examines 3PAR Thin Provisioning software and also discusses how, for applications already deployed on legacy storage, IT administrators can successfully migrate existing volumes to a thin environment by leveraging Thin Provisioning 2.0 technology within 3PAR Utility Storage and Veritas Storage Foundation. Finally, it looks to the future of thin provisioning, dubbed "Thin Provisioning 3.0," and how 3PAR and Symantec are partnering together to enable users to stay thin over time.

## Veritas Storage Foundation Overview

Veritas Storage Foundation from Symantec is an integrated software suite that provides core storage management capabilities for servers. Major Storage Foundation components include the Veritas File System (VxFS), a high-performance file system; Veritas Volume Manager (VxVM), which creates and manages logical volumes comprised of LUNs or physical disks; Dynamic Multipathing (DMP), to provide load balancing and I/O error recovery on SANs; and Storage Foundation Manager, providing central management and task workflow for the other Storage Foundation components.

## 3PAR InServ Storage Server Overview

3PAR Utility Storage is a highly virtualized, tightly clustered, and dynamically tiered storage platform built for utility computing. The 3PAR InServ® Storage Server family is the hardware foundation of 3PAR Utility Storage. Unlike modular and monolithic (or cache-centric) storage arrays, 3PAR InServ Storage Servers utilize a cluster-based approach.

3PAR InServ arrays are the first storage systems to build thin technologies into storage hardware to enable increased capacity utilization while maintaining high service levels. As a pioneer of thin provisioning—a green technology developed to address storage underutilization and inefficiencies—3PAR offers products to reduce the costs of allocated storage capacity, administration, and SAN infrastructure while increasing adaptability and resiliency.

## The Problem: Allocated but Unused Capacity

Discussions regarding poor capacity utilization have traditionally revolved around the inherent limitations of direct-attached storage. Most IT managers today recognize the value of network storage and have already shifted to a SAN-based approach. Through SANs, organizations have addressed the problem of storage “islands” and large amounts of unallocated storage and, in doing so, improved storage utilization.

However, while network storage has indeed provided significant advantages in reducing unallocated storage capacity, the question remains: are we using allocated storage as efficiently as possible? What percentage of the storage made available to applications contains actual

written data? For some organizations, this is a troubling question. They suspect, or, worse yet, they *know* that there is a problem. Others have not yet even begun to consider this question.

Whether they recognize it or not, for most companies the problem of unallocated capacity is alarming. In research conducted by Glasshouse Technologies, enterprises reported that, on average, only 25% of their allocated capacity was actually used by applications for written data. Even this figure can be high when considering capacity that appears to be written in databases but in reality is not. As a result, operating resources are wasted on the housing, powering, and cooling of countless disk drives that effectively go unused. In addition, human resources—whether internal or contracted via service agreements—are consumed in the maintenance of these exaggerated infrastructures that are potentially four times larger that they need to be. Needless to say, this is an extraordinary waste of operational resources.

The "allocated-but-unused" problem occurs for a number of reasons, both organizational and technical. These include:

**Traditional storage is built on "dedicate-on-allocation" technology.** Traditional storage uses a dedicate-on-allocation approach, which means that customers must purchase all their storage up front, prior to actually allocating this storage.

**Over-requesting capacity.** When requesting capacity, administrators need to account for current needs as well as anticipated growth. To the extent that anticipated growth is uncertain, capacity estimates are increased beyond anticipated needs to create a buffer that ensures enough capacity will be available in the future.

**Applications are less successful than planned.** Sometimes applications are simply not utilized as originally anticipated so little growth occurs. Meanwhile, significant storage was provisioned and allocated up front and cannot be easily reclaimed even though it is never actually required.

**Replication of unwritten data.** For backup and recovery purposes, copies of data volumes are often kept online. Since most data volume copy technologies deployed today do not distinguish between written versus unwritten portions of the volume, they physically replicate the entire base volume. Any "waste" in a base volume is then perpetuated throughout all of its copies.

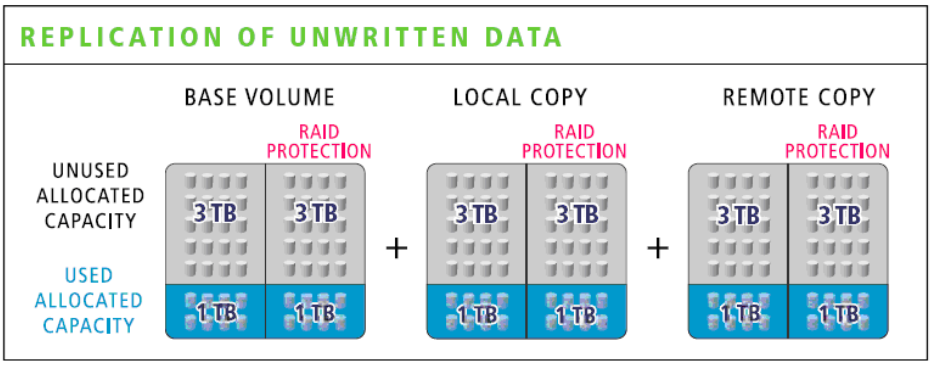


**Figure 1.** Unwritten Data Consumes Additional Storage for Each Copy

**Over-provisioning to avoid future allocation complexity and risk.** Since storage provisioning on legacy storage is complex and painful, generous provisioning is often employed up front to avoid adding capacity in the future. For traditional storage systems, provisioning is a manual process requiring careful planning and coordination by IT management, storage administrators, system administrators, and application administrators.

These many factors contribute to allocated storage capacity figures that routinely exceed the actual capacity required for written data and produce significant datacenter waste.

## Start Thin: Thin Provisioning 1.0

Thin provisioning offers a simple solution to the problem of unused allocated capacity. It is a virtualization technology that allows IT departments to safely allocate up front as much logical capacity as is conceivably required by an application over that application's lifetime. This approach breaks the connection between what the application "sees" as physical capacity and the amount of physical capacity that is actually purchased and deployed. More likely than not, what the application "sees" is much greater than the actual physical storage capacity of the system.

### 3PAR Thin Provisioning Software

3PAR Thin Provisioning is the pioneering thin software application that has enabled countless organizations to "start thin" by safely de-coupling "allocated" storage from "used" storage, enabling just-in-time delivery of storage to applications and maximizing capacity utilization. 3PAR Thin Provisioning uses *dedicate–on–write* technology, as opposed to *dedicate–on–*

*allocation* to break the connection between allocated and physical capacity (Figure 2). Physical capacity is drawn only as needed from a small buffer of free capacity. That is, physical capacity is only allocated and consumed when an application writes actual data to the logical volume. With 3PAR InServ arrays, physical capacity can be dynamically and non-disruptively added to this capacity buffer at any time.

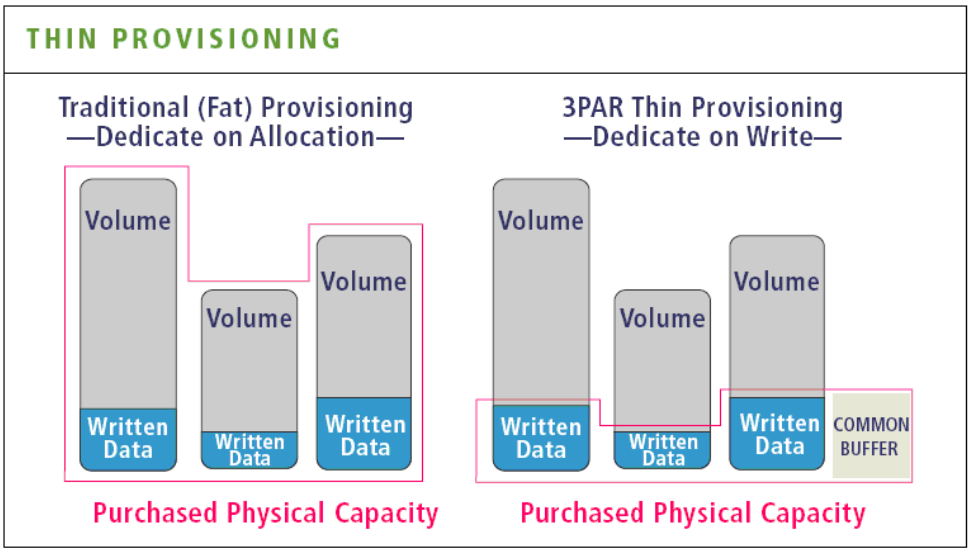


**Figure 2**. 3PAR Thin Provisioning versus Traditional "Fat" Provisioning

With 3PAR Thin Provisioning, allocated storage is presented to host servers using Thin Provisioning Virtual Volumes (TPVVs). Unlike traditional volumes—which are pre-mapped to underlying logical disks and ultimately to specific blocks of physical capacity—TPVVs are mapped to a logical common provisioning group that serves as a template for provisioning capacity to applications. When writes are made to a TPVV, the common provisioning group creates the mapping to underlying logical disks and free space is autonomically provisioned and consumed in fine-grained 16-KB increments.

By breaking the traditional link between allocated and purchased capacity, 3PAR Thin Provisioning allows application costs to grow in direct accordance with application use and growth. Users can continue to request and receive desired amounts of capacity but capacity purchases are now determined by actual utilization, so costs are less dependent on the ultimate

"success" of the application and more tightly aligned with actual business growth rather than estimations and projections.

For IT, 3PAR Thin Provisioning brings with it an important additional benefit: the pain and complexity of provisioning is dramatically reduced. By allocating generous amounts of logical capacity to applications up front, follow-on provisioning tasks and workflow are eliminated. Administrators allocate capacity just once—safely and economically. Built-in mechanisms for monitoring and controlling planned and unplanned data growth provide administrators with the necessary tools and safety features to confidently deploy TPVVs without having to actively manage them. For example, 3PAR Thin Provisioning allows administrators to set capacity thresholds flexibly, so that when an administrator-set threshold is reached, the 3PAR InServ array generates the desired alerts. Over time, as TPVVs utilize capacity within the array and as utilization approaches these specified thresholds, the system automatically generates several types of warnings to provide ample time for the administrator to plan for and add necessary capacity. In the unlikely scenario that the hard limit is reached, the InServ naturally prevents new writes from occurring until more capacity is added.

3PAR Thin Provisioning software was introduced with the first InServ Storage Server in 2002 and today is supported by all 3PAR arrays. For more details on 3PAR Thin Provisioning and the 3PAR architecture, please read *3PAR Thin Provisioning: Eliminating Allocated but Unused Storage and Accelerating ROI*, a white paper available for download from www.3PAR.com.

## Not All Thin Provisioning Is Created Equal

As thin provisioning technologies are adopted by various storage vendors, different methods of implementing the concepts behind thin provisioning have emerged. The following checklist of thin provisioning criteria is useful in differentiating between the various thin provisioning implementations and testing their quality:

**Allocation unit size.** How much physical capacity is consumed when a write is received? Some thin provisioning implementations dedicate megabytes of physical storage upon even the smallest write. A coarse unit of space allocation can mean that the simple creation of a file system on a thin provisioned volume can completely fill the volume, eliminating any value of thin provisioning before the first piece of application data is ever written to the file.

**Reservation-less versus reservation-based implementation.** Is physical capacity configured into and reserved in a specific pool on which thin provisioning works? With reserved, dedicate-on-write implementations, physical capacity is pre-configured and

reserved, placed up front into specific pools, which represents a tremendous waste. This approach commits silos of capacity before any data is written, working against the benefits of true thin provisioning. In addition, the pools themselves require ongoing management and provisioning, adding both administration time and risk.

**Autonomic versus manual provisioning.** Do administrators have to manually configure storage and RAID groups into pools to keep them replenished? When pool replenishment is a manual activity, organizations forego thin provisioning or compensate by over-replenishing—materially decreasing capacity utilization and efficiency. Conversely, when thin provisioning is implemented autonomically, capacity is dedicated and configured naturally, without the need for human intervention.

**Dual-controller versus clustered architectures.** Can you aggregate lots of distinct workloads? Fundamentally, thin provisioning is about making capacity promises that may have to be fulfilled in the future, so the scalability of an array platform is very important. Dual-controller architectures, which usually operate on an active-passive basis, are disadvantaged in this respect. Conversely, highly scalable, clustered systems are ideal for thin provisioning deployments since they are purposely designed for storage consolidation and growth within a single system.

**Built-in versus "bolt-on" virtualization.** Are virtualization features designed into the array, or are they added after the fact? The special demands that thin provisioning can place on an array that is not specifically architected for it can mean unhappy tradeoffs. For example, users may expect diminished performance, an inability to replicate thin provisioned volumes in thin or snapshot form, and an inability to thin provision all available physical capacity. In other cases, traditional reporting capabilities are inadequate, imposing a heavy monitoring burden on administrators.

## "Thin-Friendly" File Systems

Because thin provisioning employs dedicate-on-write technology, whenever an operating system or application writes data to a TPVV, physical storage is allocated and consumed. As a result, some file systems prove to be more "thin-friendly" than others. This is achieved in two ways: first, by limiting the number of blocks written during the file system creation, and second, through re-using blocks during the normal course of creating and deleting files.

Because thin storage environments allocate physical storage upon writes in a fixed size, often referred to as a *chunk*, some file systems require more chunks up front during the file system's initial installation than others. File systems like the Unix File System (UFS) can result in small writes being written in large steps across the entire underlying volume to support the file system metadata (inodes).

Figure 3 illustrates how a non-"thin-friendly" file system that uses a large (MB-size) allocation unit size results in a large amount of physical storage being allocated for just a small amount of

data. As depicted in Figure 3, each small write from the file system results in the allocation of a *chunk* of physical storage from the underlying thin LUN.

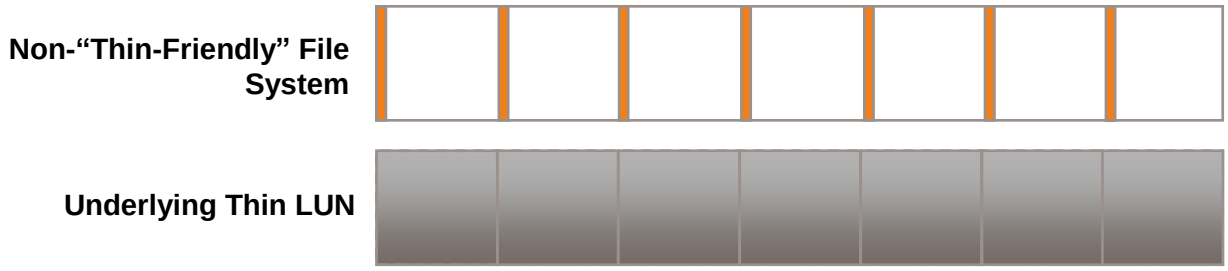


**Figure 3.** Non-“Thin-Friendly” File system and Large Allocation Unit Size Produce Inefficient Allocation of Physical Storage

The Veritas File System (VxFS) is an extent-based file system which uses thin provisioning very efficiently and is able to minimize unnecessary storage allocation over time. Contrary to the file systems like the one illustrated above in Figure 3, VxFS creates metadata only as needed. More importantly, the VxFS inodes are allocated at the beginning of the volume, thereby minimizing the amount of storage allocation triggered on the thin LUN. As a result, VxFS uses thin provisioning very efficiently, even with thin provisioning implementations requiring large allocation units.

Figure 4 shows how the “thin-friendly” VxFS only requires a single “chunk” from the common buffer to be allocated compared to the seven chunks in the non-“thin-friendly” example from Figure 3.

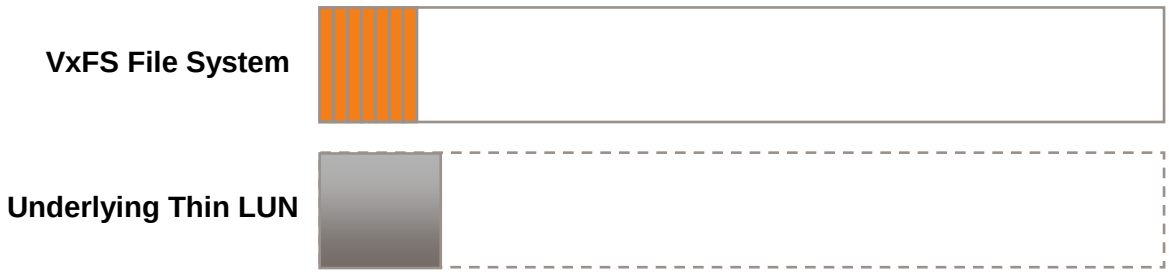


**Figure 4.** “Thin-Friendly” File System Minimizes Allocated Storage, Even with Large Allocation Units on the Thin LUN

Another benefit of VxFS is that it minimizes wasted space—not only when it is initially installed (as shown above), but also when deleting and creating files. When new files are created, the associated I/O writes trigger the allocation of physical storage on the thin LUN. A file deletion on the other hand typically does nothing other than causing the file system to update metadata. As a result, the physical space on the thin LUN that supported a file that got deleted stays allocated on the thin LUN. This physical space is effectively wasted until new file system writes occur on it. VxFS minimizes the amount of wasted space on a thin LUN by aggressively reusing freed space. When a file is deleted, the already allocated space is re-used for subsequent writes rather than requesting that a new block of storage be allocated.

Thin technology represents the potential for tremendous savings by solving the challenge that faces businesses with a growing quantity of underutilized storage assets. By deploying thin provisioning with a thin-friendly file system, companies can reduce capital expenditures, lower operating expenses, and accelerate return on investment. The next section examines how to deploy thin technologies in your environment, whether starting with a new storage system or looking to efficiently store existing data by migrating from legacy storage.

## Get Thin: Thin Provisioning 2.0

In an ideal world, all your storage would start thin. 3PAR Thin Provisioning software makes it extremely easy to quickly provision a new LUN that is thin provisioned. With thin LUNs, the physical storage required up front is significantly less than that required with traditional storage implementations. Moreover, additional physical capacity is only required as new data is generated and stored. The IT administrator can immediately see the savings and, since thin provisioned volumes grow as needed without active management, this approach requires little other volume maintenance moving forward.

But for many, starting thin has not been an option with their existing storage vendors. What about the petabytes of storage that currently reside on expensive and underutilized legacy storage? While the growing enthusiasm for and adoption of thin provisioning benefits new applications, until recently the issue of how to efficiently and cost-effectively migrate from “fat” legacy storage over to thin storage and immediately reclaim storage has been overlooked.

The development of Thin Provisioning 2.0 capabilities by industry leaders 3PAR and Symantec has enabled organizations to take the next step of migrating existing "fat" volumes onto thin storage, simply and affordably. Administrators now have the option of either using host-based or array-based approaches to migrate from traditional, fat storage to more efficient thin storage.

### Overcoming Legacy Migration Challenges with Veritas SmartMove

To help IT managers migrate their existing legacy storage to a thin environment, Symantec has announced Veritas SmartMove™, available on Unix and Linux in Storage Foundation 5.0 MP3, and on Windows in Storage Foundation for Windows 5.1. SmartMove analyzes data during online migrations and ensures that only the necessary blocks of data are moved. This new feature maximizes the value of using a thin provisioned storage system by eliminating unused space during migration and allowing organizations to immediately and easily reclaim unused space. Considering the fact that storage utilization typically averages around the 30% to 40% mark, there are tremendous amounts of unused storage that stand to be reclaimed in a typical legacy datacenter.

Traditional logical volume managers cannot move volumes from legacy arrays to thin provisioned arrays without moving every block. As a result, the migration completely fills up the new, thin LUN (Figure 5)—meaning that physical capacity must be allocated for the entire virtual volume, and defeats the purpose of thin provisioning.

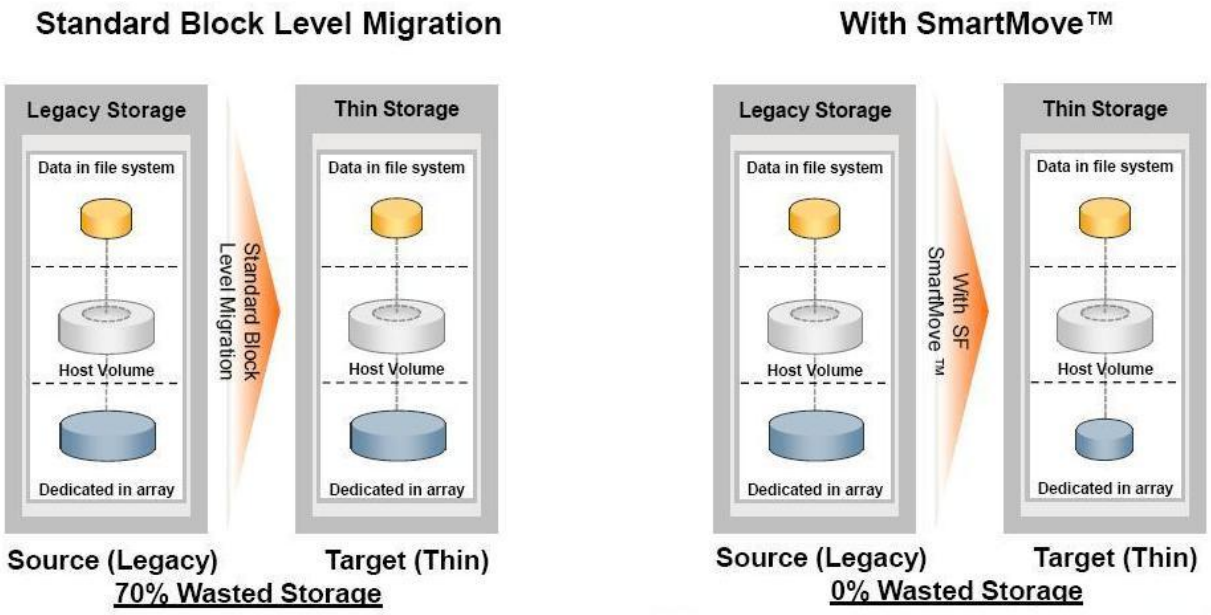


**Figure 5.** Standard Block Level Migration vs. SmartMove

The solution to this problem is to move up to a software level that knows whether a block is occupied or not, which is where the file system comes in. The basic job of a file system is to remember which blocks make up each file. They also track which blocks are free in order to make new files or grow existing ones. As a result, file systems know which blocks are free and which are not.

Veritas Storage Foundation uses this knowledge to enable online migrations from legacy storage to thin storage, thus reclaiming unused free space. The method used employs the basic storage migration capability of Veritas Volume Manager, which has been in the product since its first release in 1990.

When Veritas Volume Manager (VxVM) begins a migration, it requests the free space map from the Veritas File System (VxFS). This is transparent to users and administrators, and the migration steps are identical to those performed in traditional array migrations. This facility is known as Veritas SmartMove. When creating the new mirror on the thin LUN, Storage Foundation only copies the occupied blocks. It does not copy the free blocks (Figure 6), regardless of whether they contain zeroes or the old data of deleted files. By only copying the used space, SmartMove

enables faster migrations and ensures that there are no unnecessary storage allocations on the thin LUN.

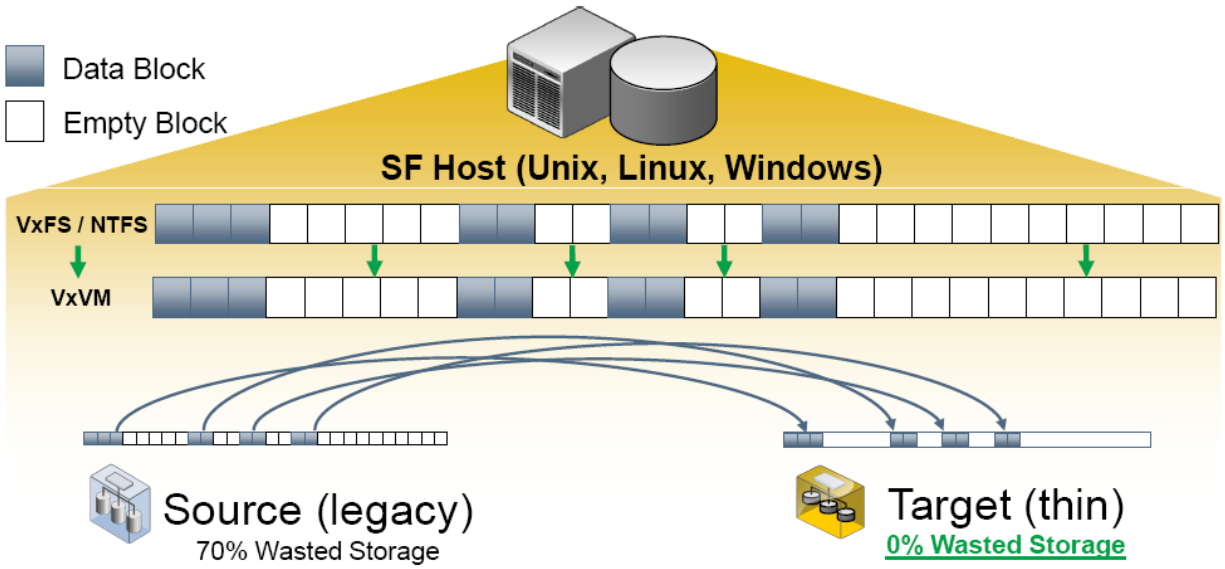


**Figure 6.** Veritas SmartMove Only Migrates Occupied Blocks

SmartMove is a simple and fast mechanism for migrating onto thin provisioned environments. On Unix and Linux, Storage Foundation 5.0 MP3 DMP also has the ability to automatically discover when a device connected to a host is thin. If any type of copy operation is performed on a device that is known to be thin, SmartMove automatically takes over. SmartMove is compatible with any storage array offering thin provisioning, including the 3PAR InServ, and supports Windows, Unix, and Linux platforms.

### 3PAR Zero-Detection Capability and Thin Built In™ ASIC

3PAR is the first in the industry to offer storage systems with thin processing technology known as *zero-detection* built right into the silicon of its arrays. When combined with the appropriate version of the 3PAR InForm® Operating System, the zero-detection-capable 3PAR Gen3 ASIC supports wire-speed “fat-to-thin” conversions and is compatible with any host volume. This Thin Provisioning 2.0 technology can effectively thin a heterogeneous datacenter to one-quarter of its original size or less, making a technology refresh more affordable and providing space and power consumption relief for datacenters approaching maximum capacity.

With 3PAR's hardware-based zero-detection capabilities, migration of "fat" volumes from other storage platforms to new "thin" volumes on the InServ can be achieved quickly and without the application disruption that can accompany software-based implementations. With fat-to-thin volume conversions taking place at the hardware level and offloaded from the storage controller processors, more parallel memory transactions are possible and controller performance is not impacted.

## Stay Thin: Thin Provisioning 3.0

Given most companies' limited IT resources and growing storage requirements, thin provisioning clearly delivers value by enabling applications to be deployed with minimal physical storage outlay. Organizations can add physical storage as applications grow and as actual data requires more physical space. Thin provisioned arrays provide immediate benefits to new or migrated applications during the initial storage allocation, but what happens to those volumes when a "transient data" application is employed that regularly deletes or moves data off of the thin volume?

Currently, users cannot reclaim and reallocate storage space when large amounts of data is deleted from a thin volume, thus transient data applications are not as well suited to take advantage of thin provisioning products available on the market today.

What good is being thin if you cannot stay thin? To address this challenge, 3PAR and Symantec—recognized as leaders in storage industry innovation—have integrated 3PAR Thin Provisioning technology with Veritas Storage Foundation to allow granular, automated, non-disruptive thin reclamation within the 3PAR storage array. Together, 3PAR and Symantec are working to bring to market a Thin Provisioning 3.0 solution that leverages the autonomic use of file system-level intelligence to continually optimize storage utilization and deliver infrastructure automation.

### Large Deletions and Data Movement Reduce Thin Provisioning Benefits

As users or applications delete files—whether at the file system or the database layer—logical space is freed, but physical storage in the thin volume remains allocated. Consider a thin volume that claims to be 10 GB, but with actual physical storage at only 2 GB. That 2 GB represents the "high water mark" for data used.  Imagine this system had two files, each 1 GB in size, as in

Figure 7. A user deletes one file, reducing the actual logical usage at the file system layer to 1 GB; however, the size of the thin volume remains 2 GB (Figure 8).

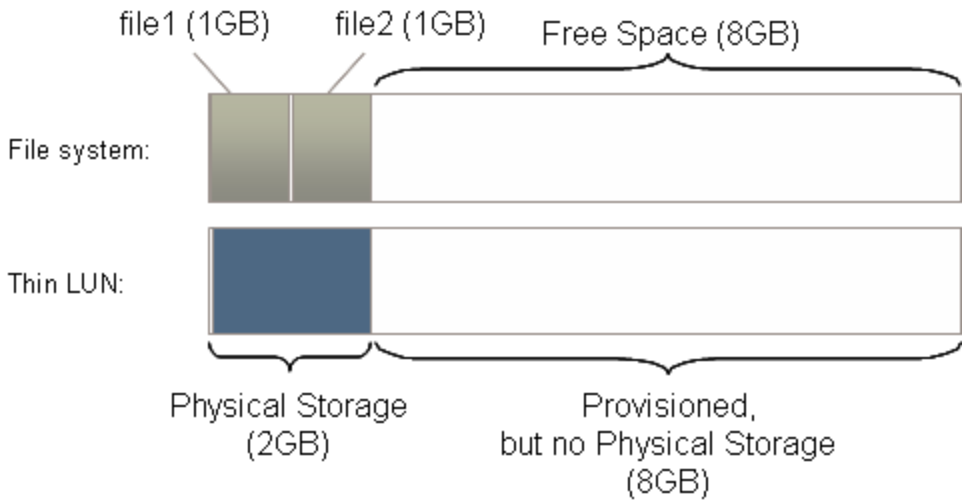


**Figure 7.** Thin Volume Before Any Deletion Occurs

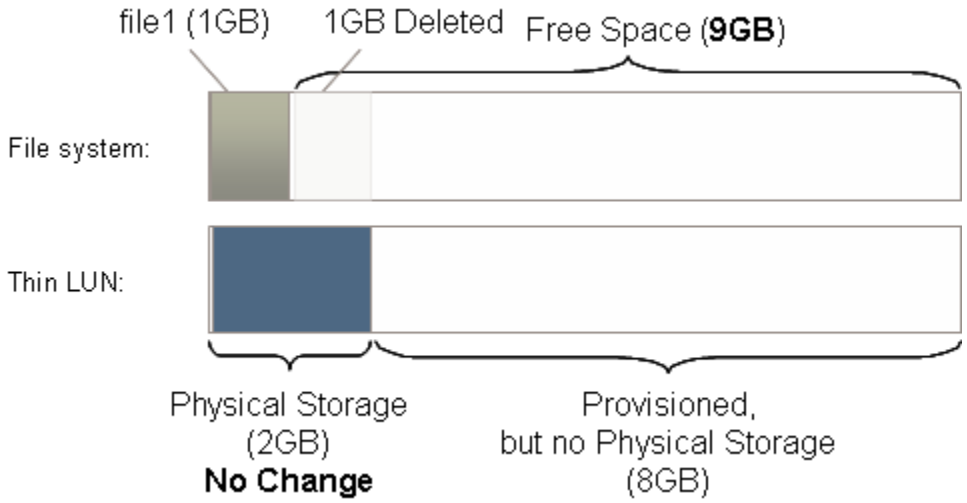


**Figure 8.** Deleting Files Does Not Reclaim Thin Provisioned Storage

Some file systems, like VxFS, partially reduce this challenge by efficiently using the allocated storage. When any file system receives a request for more storage—when an application or user creates a new file or appends data to an existing file—that file system uses internal algorithms to

choose the free blocks to assign to that request. The Veritas File System's algorithm allocation is thin provisioning-friendly, designed specifically to aggressively reuse freed space, thus minimizing unnecessary storage allocation on thin LUNs.

While file systems that efficiently use already allocated space help maintain a high level of storage utilization over time, this approach still does not offer a way to truly reclaim allocated storage, an issue which is particularly acute in transient data environments where large file deletions are common and/or data is retained for only a short period of time. Consider the example illustrated in Figure 8. If half of the originally allocated 2 GB of data is deleted, the application may not write 1 GB worth of data to disk anytime soon, if ever. In this case, it would be extremely advantageous to reclaim the unused space—1 GB in this example—for one of the many other applications that can run on the same thin provisioned array.

### Online Thin Reclamation with 3PAR and Symantec

3PAR and Symantec recognized the need to not only enable customers to *get thin*, but also to *stay thin* over the life of their applications and storage systems. The information necessary to optimize thin storage is in the host file system. Only the host file system knows which blocks are used and which blocks are free. Enabling online thin reclamation requires integration between the host and the storage hardware. As a result, 3PAR and Symantec have teamed together to usher in the next wave of thin provisioning, Thin Provisioning 3.0. We already looked at Thin Provisioning 2.0 technology, which allows users to *get thin* with array-based or host-based fat to thin migration solutions. The core feature of Thin Provisioning 3.0 is Online Thin Reclamation technology, which gives organizations the ability to automatically and efficiently reclaim allocated storage space that is no longer used (Figure 9).

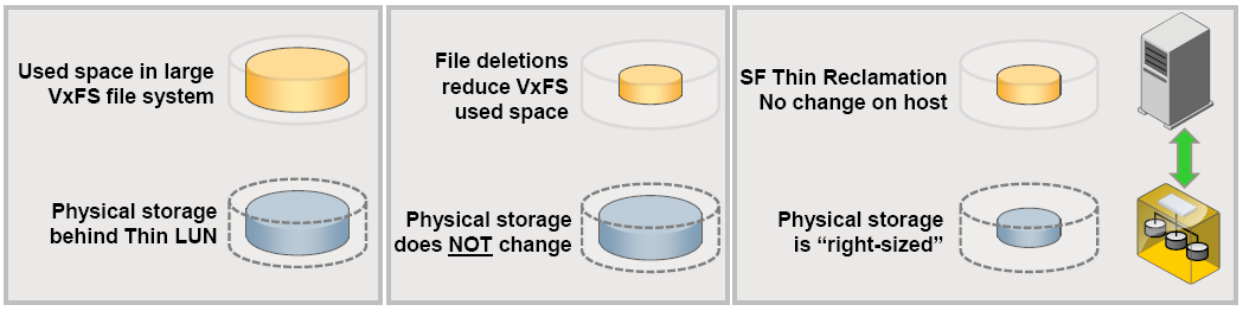


**Figure 9.** Right-sizing thin storage with Veritas Storage Foundation Thin Reclamation

3PAR and Symantec are the first vendors to develop an automated mechanism to enable IT administrators to reclaim their unused storage capacity and allow the system to automatically allocate that storage to other projects and applications.

The joint efforts between 3PAR and Symantec give Veritas Storage Foundation from Symantec the ability to track the mapping of block-level capacity of the Veritas File System to thin provisioned volumes on 3PAR InServ Storage Servers. This capability gives the Veritas File System within Storage Foundation the ability to periodically communicate the location of free blocks to the 3PAR array as file deletions and changes happen within the file system. 3PAR arrays then use this granular file system-level information to autonomically reclaim unused space within thinly provisioned virtual volumes, thus minimizing storage waste without impacting host applications.

In order to enable the new Online Thin Reclamation functionality, enhancements must be made to both the storage array and the host file system. These enhancements leverage industry-standard SCSI commands and are therefore not based on a proprietary interface and are available for any file system or storage system vendor to implement. They are referred to as the Thin Reclamation API in the following paragraphs.

The following breakdown describes the communication between the array and Veritas file system that enables customers to stay thin with Thin Reclamation technology:

1. The Veritas file system determines whether it is communicating with a thin provisioned volume or a traditional volume. 3PAR arrays and Veritas Storage Foundation use the standard SCSI commands to identify whether or not a volume is thin, and whether or not the thin volume supports Thin Reclamation API.

2. For devices that are known by Storage Foundation to support the Thin Reclamation API, after files are deleted, the amount of wasted storage in the array (storage that is allocated but unused on the host) increases. To reduce the amount of wasted storage, a thin reclamation operation can be triggered on VxFS or on a VxVM disk. During a thin reclamation operation, the file system communicates with the array to indicate which blocks were deleted and are therefore available for reclamation. Storage Foundation uses the Thin Reclamation API to notify the 3PAR array of all the block ranges that are free and should be reclaimed by the array. Once notified by Storage Foundation, the 3PAR storage array marks the corresponding blocks as “free” in the thin provisioned volume. This is done with 16-KB page granularity. Once identified by the storage system as empty, these 16-KB blocks become available for reuse.

The complete thin reclamation solution is still under development by 3PAR and Symantec and is expected to be widely available in 2009. The host side of the Thin Reclamation API by Symantec is available today starting with Storage Foundation 5.0 MP3 on Unix and Linux.

## Conclusion

Modern datacenters are constantly asked to do more with less, especially during times of tightening IT budgets. Thin provisioning enables them to deliver exactly that. It solves the allocated-but-unused problem, saving millions of dollars in base and replicated capacity, as well as related costs for housing, powering, cooling, and software licenses.

It is a mistake to stop there, however. As we have discussed, datacenters can not only *start thin*, but can also *get thin* and *stay thin* when Veritas Storage Foundation from Symantec and 3PAR Utility Storage are used together. Available today, IT administrators can leverage Storage Foundation and SmartMove to simply and quickly migrate traditional storage to a thin provisioned environment and reclaim all unused space. Storage Foundation delivers all the capabilities needed to support the migration itself and handle operations in a thin environment in the most efficient manner available today. And with the next release of the InForm Operating System, 3PAR will offer an additional option for fat-to-thin volume conversion with the first application available for its thin processor, the Gen3 ASIC. With this new capability, 3PAR Utility Storage will support wire-speed fat-to-thin conversions that are compatible with any supported host environment.

As we look to the future with Thin Provisioning 3.0, 3PAR and Symantec are working together to expand thin provisioning even further by enabling organizations to stay thin through Online Thin Reclamation capabilities. This functionality will provide administrators with an easy and automated method for maximizing storage capacity utilization by reclaiming fine-grained physical storage blocks that are no longer being used by the file system. With a comprehensive thin provisioning strategy encompassed by products from Symantec and 3PAR, customer can for the first time *start thin, get thin*, and *stay thin*.

**About Symantec**

Symantec is a global leader in providing security, storage and systems management solutions to help businesses and consumers secure and manage their information. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries. More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation
World Headquarters
20330 Stevens Creek Boulevard
Cupertino, CA 95014 USA
+1 (408) 517 8000
1 (800) 721 3934
www.symantec.com



02/08