

Product Brief

Symantec's Broad Approach to Thin Provisioning

Date: June 2010 Author: Mark Peters, Senior Analyst

Abstract: Thin provisioning is a smart solution when users want to manage storage efficiently, but it's not a storage management panacea. Using well-integrated tools to target wasted storage capacity and put it back to work is a sound strategy to significantly reduce storage costs. <u>Symantec</u>'s Veritas Storage Foundation provides a well balanced and integrated management solution that yields broad storage management efficiencies that embrace—but do not rely solely upon—thin provisioning.

Why Thin Provisioning is Important

The Big Picture

The recent market contractions and contortions have pushed the economic pressure for IT performance to new extremes; to re-coin an old phrase, we in IT have never seen so much pressure to do so much with so little. This partly explains the rapid adoption of virtualization to help allow more to be done with less. Even so, in storage, it is common to find a vast amount of capacity effectively underutilized and beneath the management radar. It is a best practice for organizations to optimize their existing storage infrastructure before procuring new storage—this is where "thin provisioning" comes in. Thin provisioning is an elegant solution that increases storage allocation efficiency and thereby more effectively leverages a shared pool of storage capacity. In doing so, there is further opportunity to do more with less and to answer the call of economic necessity from business leadership.

From Fat Storage to Thin Provisioning

In recent years, studies have shown that as much as 70% of physical storage capacity is often wasted by what can only be described as mismanagement. This happens partly as a result of disconnected storage management tools and processes; the problem is sometimes aggravated by a lack of compatibility and interoperability between storage management solutions. From a management efficiency point of view, this can create chaos, leaving allocation and utilization efficiency measurements low. Rather than fixing the larger issue of the storage management solution jumble by using a product (or products) that is/are well integrated with internal IT processes for managing capacity, administrators often buy cavernous amounts of new capacity to meet the never-ending demands for storage capacity. Thin provisioning is invariably a great efficiency enhancer, but used alone or out of management context, it can only help, but not solve, the broader need for improved storage efficiency. IT management is complex and so new technologies must be easy to operate and integrate, just as much as they are financially attractive.

Thin provisioning is comprised of two major capabilities. It aggregates multiple volumes, which can span arrays, into what appears to be a single reserve or free pool of capacity. This allows for a singular view of a defined storage pool spanning multiple arrays. Secondly, thin provisioning uses a technical approach known as "allocate-on-write" to only allocate capacity if blocks are actually written. If, for example, a volume is created for 500 GB, but the application only writes and uses 100 GB, then the true physical storage allocated is just the 100 GB. The remaining 400 GB are free to be shared and used among other operating systems tasks, applications, or users. Aside from the efficient "just-in-time" usage of actual capacity, storage administrators can over-allocate "just-in-case" capacity as a buffer for applications to have room to grow into. Each application, user, or system administrator can have (at least conceptually) as much storage allocated as is desired—the storage manager needs to monitor only the <u>real</u> physical consumption of storage across the entire system.



The "Allocated, but Unused" Problem

Put directly, the objective of thin provisioning is to reduce the total amount of storage necessary and it achieves this by allocating and using shared storage capacity on demand. In the scope of mismanaged capacity, a standard provisioning scenario typically has a very significant percentage of the total available capacity (often 25% or even more) that is *allocated but unused* or *allocated but orphaned*—either way, it is stranded. This can happen from a request for storage capacity from a user or application owner who never actually creates a file system or stores a single piece of data. The challenge here is that traditional storage reclamation methods are invariably cumbersome, expensive, and sub-optimal. Multiple teams—such as the application, database administration, system administration, and storage administration teams—have to cooperate to perform a series of time-intensive tasks, such as shrinking the applications, file systems and host volumes; reducing the relevant LUN space; and finally returning the freed space to the shared pool. In a nutshell, the reclamation process is highly disruptive to applications, requires cross-team communication, and often ends up being sub-optimal even if it is completed successfully (because of the greater cost and complexity of this reclamation method, it often simply does not get done).

Rather than solving the tough problem of identifying and reclaiming space that has been allocated but is unused or stranded, administrators will often continue to over-purchase (or over-provision) storage capacity as a way of masking the limitations of existing storage management tools or processes, and believe that they will generally not get caught running out of capacity. Indeed, some storage administrators will even refer to the allocated, but unused, space as their "emergency free pool." If they run out of capacity, they'll claim that they can simply "find" this space and return it to the free pool of storage. All of this is a very ad-hoc and unstructured methodology that can eventually lead to application issues from unexpectedly running out of storage capacity and crashing or costly emergency storage hardware purchases. Once again, this is not a blame-game—the technology to do things better did not always exist. Today it does, but it would be naive to assume that simply purchasing new hardware will solve all problems.

Not all Thin Provisioning is Created Equal

Symantec's Veritas Storage Foundation: a Broad View

Certainly, even a marginal use of thin provisioning is invariably a good thing, but leveraging the technology can be done even more efficiently when the thin provisioning enabled hardware is integrated with the host —specifically, the file system and volume manager. Symantec's Veritas Storage Foundation file system can gaze across multiple arrays, regardless of brand, and offers unique capabilities when it comes to leveraging thin provisioning. The primary benefit of these capabilities is that an organization can address the two major elements involved in the successful long term adoption of thin provisioning: as with any form of diet (this being a capacity diet!) organizations first need to get thin and then to stay thin. And they need to do both of these in a manner that is non-disruptive and seamless to the business.

Getting Thin. The first challenge for IT organizations is how to migrate to a thin environment while eliminating waste. Unfortunately, most online storage migration tools perform block-level copy, which means that every block on the destination is identical to the source. The reason for this limitation is that most migration tools are unable to differentiate between a used and an unused block of data. As a result, these tools treat all blocks equally and copy all blocks during a migration. This means that even in the process of getting into a thin environment, there can be a lot of waste generated in the new environment.

Veritas Storage Foundation solves this problem by uniquely identifying allocated, but unused, blocks and copies only used data blocks to the target storage device. A few other unique capabilities of Veritas Storage Foundation are that in addition to migrating only the blocks that matter, the migration is hardware independent (that is, from any vendor to any vendor) and is performed via a simple web-based interface while applications stay online. This facilitates an easy, hardware-independent, non-disruptive transition to a thin environment combined with ongoing bandwidth and storage savings at any other time that a migration is performed.



Staying Thin. The second challenge IT organizations face is how to keep an environment thin over time. While this may seem like a rational expectation, hardware vendors often don't make it easy to stay thin. Specifically, blocks of data which are erased by an application or user are typically non-reclaimable even by the most advanced thin provisioning enabled array. The reason for this is simple: when applications or users delete data, the respective blocks of data are not "zeroed" out. This means that the array is not informed that those blocks are available for reuse.

Veritas Storage Foundation provides the host side intelligence to notify the array and actively trigger reclamation of previously used but now available data blocks, returning them to the free pool. In this way, the reserve free pool can remain as capacious as possible, providing optimum value to the user by elongating the amount of time the available (real) storage capacity can stretch its effective use before new storage capacity must be purchased. This capability is important. Other file systems and volume managers use a "write-once, allocate-forever" approach that precludes the ability to repurpose unused blocks.

The Broader View. When thin provisioning is implemented *along with* a powerful suite of storage management tools (such as Symantec's Veritas Storage Foundation), not only can 40-50%—or sometimes even more—efficiency in usable storage capacity be gained, there are also other immediate and measurable improvements:

- Centralized management of data centers can result in reduced management complexity.
- Higher availability and performance of applications, databases, and data center assets.
- Significant reduction in the capital and operating expenses for storage infrastructure.

The point of the link between storage management and overall storage costs is an important one—simply stated, inefficient storage management is a platform from which storage vendors can oversell capacity. Certainly, if storage were free of cost, any issues with the growing complexity of managing resources could be hidden behind vast amounts of capacity. However, not only is this expensive from a CAPEX point of view, it is disastrous from an OPEX point of view; neither of which is consistent with the issues of tight[ened] budgets and management objectives for bottom line efficiency in IT.

Optimization Leads to Value

There are some interesting lessons here. The first is that the combination of orphaned, unused, and inappropriately used space can easily represent some 30% of total storage capacity; in other words, it is storage capacity—money—that is wasted. Perhaps another 40% of the capacity is previously used data blocks and thus ripe for migration. The recovery of all this "wasted" storage could provide a year of data growth—free! There is value to be had, but how should users get at and release it? The issue is often not addressed at all and yet these are IT manifestations of low-hanging fruit. The best way to really identify the opportunities is to avoid the use of non-integrated storage management solutions that transfer the burden of storage usage accounting to the storage array.

Instead, best practice is to use an integrated storage management platform, such as Veritas Storage Foundation, which offers simplified management, higher performance and higher availability, and works in a heterogeneous fashion with existing and new storage assets. When combined with thin provisioning enabled hardware, storage administrators can effectively provision and manage their physical storage use. Using Veritas Storage Foundation, storage administrators can painlessly move to a thin environment and, moreover, stay thin over time. This combination of effortless migration and automated identification and reclamation of wasted storage is what sets Storage Foundation apart in effectively addressing what can otherwise—and all too commonly—represent around 60% or 70% of wasted and/or wrongly used storage capacity.



An Example of the Financial Impact of Effective Thin Provisioning

From a business point of view, the potential gains of thin provisioning can be significant. Imagine an organization needs roughly 150 TB of usable (real data) storage capacity. Here are the other assumptions:

- Calculated over three years.
- Depreciation over five years.
- \$2 per gigabyte.
- Price erosion rate 2%.
- Data growth rate 40%.
- Energy costs \$.12kWh.
- No FTEs included in OPEX.

Table 1 highlights the financial benefits that could be achieved using Veritas Storage Foundation in this environment.

Table 1: Financial Results			
	Traditional Storage Environment	Thin Provisioning Enabled Environment	Optimized Thin Provisioning with Veritas Storage Foundation
Total Usable Space	510 TB	510 TB	510 TB
Space Used by Application Data	153 TB	153 TB	153 TB
Allocated Physical Storage	510 TB	306 TB	153 TB
Wasted Physical Storage (Allocated, but Unused)	357 TB	153 TB	ОТВ
САРЕХ	\$2,302,347	\$1,381,408	\$690,704
OPEX	\$120,848	\$72,509	\$36,254
TOTAL:	\$2,423,195	\$1,453,917	\$726,959
SAVINGS:		\$969,278	\$1,696,237

Source: Symantec / Enterprise Strategy Group, 2010.



The Bigger Truth

The success of any business and its IT department—in good times or bad—is dependent on a nimble and efficient infrastructure. An efficient infrastructure is, in turn, dependent on IT professionals getting to the root of waste and minimizing or eliminating it. Tools that can identify and automatically reclaim waste are therefore critical. Veritas Storage Foundation offers an integrated solution that facilitates efficient storage management. Its storage virtualization and management platform enables storage administrators to truly leverage the power of thin provisioning by allowing them to effortlessly migrate to a thin storage infrastructure. Storage Foundation can then continually optimize the environment by reclaiming wasted storage in an automated fashion. This can actually be a majority of the installed storage capacity in many sites! It is an uncomfortable truth, but also an opportunity for IT managers to be heroes by returning bottom line value to their organizations.

Smaller amounts of well managed capacity can deliver the same value —and at a significantly lower cost—as larger amounts of poorly managed storage. Putting all of an organization's wasted capacity back to work has frequently proven to provide millions of dollars in net present value. Sure, there is an upfront cost to buying any software, along with the effort to have IT administrators trained. But, the amount saved by efficient storage management is almost always much greater than the cost. Veritas Storage Foundation has proven for over a decade that it can be the cornerstone of a sound storage management practice. Although optimizing thinly provisioned storage is just one part—one might say a thin slice—of the much larger value that Veritas Storage Foundation offers, it is a pragmatic solution that any businessoriented IT manager should consider implementing.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482-0188.