

# I D C   E X E C U T I V E   B R I E F

## **Thin Provisioning: Using Intelligent Storage Virtualization Technology for More Efficient Use of Storage Assets**

---

*October 2010*

*Adapted from **Worldwide Enterprise Storage Systems 2010–2014 Forecast: Recovery, Efficiency, and Digitization Shaping Customer Requirements for Storage Systems** by Natalya Yezhkova, Richard L. Villars, Liz Conner, and Benjamin Woo, IDC #223234*

**Sponsored by Symantec**

---

### **Introduction**

IT departments today face a huge challenge with data storage. The amount of digital content being created is growing exponentially, but the corporate budgets to store and manage that information are not.

The slowdown in the economy has not impacted, nor will it impact, this creation of content or the requirement to retain it. For example, IDC predicts annual enterprise storage system shipments will increase to 76.9 exabytes by 2014 — a 49% compound annual growth rate (CAGR), with external and internal capacity shipments growing at 52% and 40.1% CAGRs, respectively.

Given the nature of today's distributed online economy, IDC estimates that 70% of the data is generated by consumers and corporate workers. Globalization, Web access, new techniques for data gathering and analysis, digital communications, conversion from paper to digital processes, and increased regulatory and legal requirements all have added to an information explosion.

While some of this data comes from traditional IT applications — business processes, customer information, etc. — much of it comes from nontraditional sources, such as VoIP packets, surveillance videos, and real-time sensor information. Also, because this data is often unstructured and file based, it consumes more space and must be retained longer. This is true especially in regulated industries, where storing documentation of compliance is just as important as storing the original data.

Due to the economic downturn and higher energy costs, the traditional method of solving data storage problems — increase storage capacity — is no longer a viable option. As both data and

storage grow, the number of available administrators to manage increasing capacity does not scale at the same rate.

Additionally, overall storage utilization rates remain low — averaging 25–35%. For a number of years now, organizations simply have bought storage to support data growth — but have not addressed low storage utilization rates. As a result, total storage costs are increasingly becoming a larger part of the IT budget, especially the operational costs of floor space, energy consumption, and other expenses.

IDC sees a simple solution to the storage issue: use existing resources more efficiently and increase utilization while reducing administration costs/time. IDC believes that most enterprises misallocate storage resources to internal stakeholders and as a result have significant excess storage capacity spread throughout the organization.

Simply slowing the growth of storage footprint — or stopping it for a period of time — not only saves capital expenditures but also reduces operational costs associated with real estate and energy consumption. This involves improvements in provisioning and allocation of resources through storage virtualization or resource aggregation.

Most important, enterprises should take a hard look at an approach that not only saves resources immediately but also sets the stage for organizations to prepare for the paradigm shift in IT from compute-centric organizations to data-centric organizations.

## **The Thin Provisioning Approach**

Thin provisioning is storage virtualization technology that allows storage administrators to create, manage, and optimize a shared pool of storage by using a distributed on-demand model of storage.

In traditional approaches to storage provisioning, application developers or database managers estimate how much storage capacity they will need. Typically, they overestimate in order to give themselves a buffer of space. This results in storage capacity being reserved and never utilized, creating the vast amounts of unused resources often found in traditional, or "fat," provisioning.

Thin provisioning challenges the traditional approach to data storage by automating the process of allocating physical storage resources. This technology virtually assigns storage capacity to applications and databases, but the physical space is utilized only when actually needed. As a result, IT can eliminate the vast amounts of reserved but never used capacity.

With thin provisioning, administrators can overallocate usable capacity to an application and consume physical capacity only as applications actually write data. In effect, thin provisioning provides just-in-time storage and allows administrators to simplify their storage provisioning and increase their storage utilization rate.

Since physical storage is consumed only when needed, provisioned storage capacity space combined across all hosts and applications can safely exceed the amount of purchased capacity in the system.

As a result, users do not have to overrequest capacity in order to avoid future disruption. Similarly, thin provisioning eliminates the need for administrators to overprovision physical storage to avoid downtime as applications and associated data grow over time.

## **Thin Provisioning Benefits**

Thin provisioning is like other virtualization technologies in that it provides exceptional benefits with relatively low incremental costs. Thin provisioning eliminates excess capacity purchasing, enabling a firm to purchase only the disk capacity needed at a point in time. It eliminates the need for up-front allocation of capacity, which may sit unused for months or years, and also eliminates requirements for dedicating resources (power, cooling, and space) needed at an unknown future point in time.

This approach saves both capital and operational expenditures. IDC sees no reduction in the exponential data growth that businesses are witnessing today. The ability, therefore, to better utilize existing storage devices, and the associated overhead, represents a direct benefit to an organization's bottom line.

## **Challenges**

However, as a newer technology approach, thin provisioning does present some challenges.

The first key challenge is that organizations must plan thin provisioning adoption while eliminating waste and minimizing downtime. The critical challenge here is that traditional migration tools perform block-level copy, which means that all data blocks are duplicated at the destination during a migration process. The reason for this is that most migration tools are unable to differentiate between a used and an unused block of data. As a result, traditional migration treats all blocks equally and copies all blocks, generating waste in the new environment. An intelligent migration solution should be used that reclaims wasted space during the migration to a thin environment while causing minimal or no disruption to the business.

The second key challenge is that organizations need to keep their storage environment thin over time. Once an organization has migrated into a thin environment, and as applications begin writing and deleting data to underlying storage, the storage environment no longer stays thin. Maintaining a thin storage environment requires consistent reclamation so that physical storage can be freed and reused again as required. Unfortunately, storage hardware by itself is unable to "rightsize" storage use because it is unable to distinguish between used and unused physical storage blocks. Intelligent host-based software should be used that seamlessly works with the underlying hardware to free up physical storage when it is no longer in use.

The challenges of thin provisioning are minor compared with the benefits. In some cases, IDC has seen storage utilization efficiency increase by as much as 30%. In other instances, companies have reduced capital and operational expenditures, often eliminating the need to purchase additional storage capacity for as much as a year.

## **Conclusion and Essential Guidance**

As the amount and value of information assets have grown, so too have IT operational costs. Traditionally, enterprises have met storage challenges by adding capacity. But the economic downturn has forced IT departments to contain capital and operational expenditures by focusing more on increased efficiency.

Costs are only part of the scenario, however. Enterprises also are paying more attention to the environmental aspects of the increasing demand for storage capacity and its power demands, so more companies are looking for ways to minimize the environmental impact of their datacenter operations.

Organizations now face the difficult challenge of increasing storage capacity with fewer resources. Thin provisioning provides a way to effectively use storage resources by allocating space as it is needed rather than setting aside a predetermined amount of capacity specified by application developers, database managers, or other administrators. Thin provisioning enables resources to be shared and excess capacity to become a working asset as opposed to an unused drain on the company.

With thin provisioning, administrators can overallocate usable capacity to an application, which then consumes physical capacity only as the application actually writes data. In effect, thin provisioning provides just-in-time storage and allows administrators to simplify their storage provisioning, thereby increasing their storage utilization factor.

These advantages of thin provisioning can be leveraged even further if enterprise servers are "thin provisioning aware." This means that thin provisioning is not just a storage administrator's decision but an additional view of valuable assets from the top down (server to disk).

Organizations that expect the capacity needs of a group of hosts or applications to change over time are the best candidates for thin provisioning. IDC recommends that when selecting a thin provisioning solution, organizations look for the following capabilities to gain the maximum benefit from thin provisioning:

- Efficient writes of data across a thin provisioned volume, whereby data is written methodically instead of haphazardly across the volume
- Nondisruptive migration of used data blocks when transitioning to a thin environment
- Automated recapturing and/or reclaiming of space once data has been deleted

Increasingly, thin provisioning is becoming a part of new storage architectures, strategies, and plans. Adopters of thin provisioning include forward-looking organizations for which low storage utilization is a significant source of inefficiency and cost.

While less advanced organizations share many of the same concerns in terms of their storage infrastructure, they tend to deploy new technology when it has conclusively proven itself in the market, as well as when well-established vendors begin providing solutions. This is happening now, and IDC believes that thin provisioning will continue to be a valuable technology with significant benefits for enterprises.

#### COPYRIGHT NOTICE

The analyst opinion, analysis, and research results presented in this IDC Executive Brief are drawn directly from the more detailed studies published in IDC Continuous Intelligence Services. Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from IDC. Contact IDC Go-to-Market Services at [gms@idc.com](mailto:gms@idc.com) or the GMS information line at 508-988-7610 to request permission to quote or source IDC or for more information on IDC Executive Briefs. Visit [www.idc.com](http://www.idc.com) to learn more about IDC subscription and consulting services or [www.idc.com/gms](http://www.idc.com/gms) to learn more about IDC Go-to-Market Services.

Copyright 2010 IDC. Reproduction is forbidden unless authorized.