IBM XIV® Storage System:

Thin Provisioning Reinvented

White Paper

September 2008

info@xivstorage.com

www.xivstorage.com

Contents

Introduction 1
IT Without Thin Provisioning: Costly and Complex 1
Capacity Gone to Waste 1
The Real Cost of Unused Storage 2
Volume Resizing 2
The Solution: Thin Provisioning 3
Thin Provisioning: Scenario #1 4
Thin Provisioning: Scenario #2 5
Managing Thin Provisioning 5
Monitoring Capacity Usage 6
Limiting Thin Provisioning to Specific Applications 6
Reclaiming Logically Unused Capacity 7
Thin Provisioning and Snapshots 7
Summary 7

Introduction

Managing IT storage infrastructure is an endless balancing act of providing enterprise-class functionality and services while controlling costs. One area that is a considerable component in capital and operational costs is unused storage space that has been allocated to applications but is not being used.

"Thin provisioning" is the practice of allocating storage to applications on a "just-in-time" and "as needed" basis by defining a logical capacity that is larger than the physical capacity. This common practice is a means to reducing or postponing storage costs.

The IBM XIV® Storage System allows a highly flexible, cost-optimizing approach to thin provisioning. It enables XIV system users to allocate capacity based on total space actually consumed. The result is improved storage utilization rates, leading to greatly reduced capital and operational costs.

This document explores the benefits of thin provisioning, the XIV approach to implementing it, and other XIV capabilities that ease management and increase ROI for IT managers and storage administrators.

IT Without Thin Provisioning: Costly and Complex

The cost of storage is one of the main concerns of IT managers and storage administrators, who are faced with burgeoning organizational demands for storage capacity yet expected to curb storage expenses.

Let's examine how traditional approaches to storage allocation and thin provisioning actually inflate costs.

Capacity Gone to Waste

One of the disturbing facts about storage costs is that 70 percent or even more of the storage capacity in enterprise settings goes unused. Storage space is allocated to an application, which only actually uses a portion of it. This "waste" is caused by several factors:

- **Pre-allocating for Future Requirements.** A common practice is to pre-allocate storage space for the application based on future requirements. This practice is essential because resizing volumes is a difficult and time-consuming IT task for both storage and applications.
- **Over-estimating Needs.** Storage space allocation is typically based on a prediction of future use. If this prediction is inaccurate and the allocated storage goes unused, there is no practical way to reclaim the allocated storage for other applications.

- **Full Copies of Unwritten Data.** Many applicative needs, such as testing and data mining, require copies of the original data, either at the application level (file copy and so on) or the storage level (snapshots and full copies). Traditional architectures copy used and unused data, thus increasing the amount of unused space.

These reasons help explain why storage, one of IT's most expensive and complex resources, is under-utilized to a large extent. The cause is not inappropriate management practice but the very nature of the pre-allocation of storage and inability to easily reclaim this storage if it goes unused.

The Real Cost of Unused Storage

Although unused storage is costly unto itself, it involves many additional hidden and indirect costs, made up of several elements:

- **Capital Expense.** The direct capital expense of storage can be reduced to the storage required for actual usage. For example, a storage project that costs $1M in storage capital expenses, with a 70 percent non-utilization rate, can be implemented using a thin-provisioned system for only $300K.
- **Environmental Factors.** The unused capacity consumes electricity and generates heat, adding to power and cooling expenses. A storage project with a 70 percent non-utilization rate would cost only 30 percent of the power and cooling if implemented with thin provisioning.
- **Floor Space.** The same analysis applies to savings in floor space, which is another very expensive resource (involving building, raised floor, UPS systems, cooling systems and so on) and is often in short supply.
- **Declining Storage Costs.** Without thin provisioning, storage must be purchased at the time of allocation — sometimes years before it is used. This scenario involves a doubling of costs: not only are power, cooling, and floor space wasted during the years of unuse, but the organization is unable to leverage yearly declines in the cost of storage. Also, the power and space efficiency of storage systems improves over time; these advances go untapped when storage is bought upfront.

Volume Resizing

A partial solution to some storage capacity problems is to re-allocate more capacity to an application when needed. On the storage side, this is done by resizing the volume. This operation is complex and time consuming, due to several obstacles:

- From the application point of view, volume resizing can only be done, if at all, with downtime, which is costly and does not support the IT paradigm of non-stop service
- The process of monitoring usage and performing resizing requires significant, on-going operational attention and work, even if eventually no resizing is needed

- Reconfiguring the operating system and applications for volume resizing is a complex and costly process
- Backup systems need to be re-configured when volumes are resized

The Solution: Thin Provisioning

The idea behind thin provisioning is simple. Users define volumes with any logical size, and acquire and install only the physical capacity needed for data that is actually written.

This simple idea is similar to others used widely in other IT areas, for example:

- Virtual memory in a computer gives each process the perception of a huge memory space, while physical memory is much smaller
- File systems allocate user quotas that may total more than the available space
- Communication infrastructure paradigms have switched from dedicating capacity per user or application to statistical multiplexing of multiple users or applications

With the XIV system, thin provisioning functions similarly to the above concepts. Thin provisioning principles are as follows:

- For a system, two types of capacities are defined:
  - **Hard Capacity.** The physical disk capacity available to applications (net capacity, after taking into account redundancy and spares)
  - **Soft Capacity.** The logical capacity of the entire system. Volumes are defined out of this capacity.
- When defining a volume, its size is only limited by the remaining soft capacity of the system.
- At the time of provisioning, the new volume is defined as **formatted**, meaning all zeros, and does not consume any hard capacity. As applications start to write data, the written areas consume hard capacity.

Note: Any system must potentially be expanded to its soft capacity. When defining the soft capacity, the customer must pre-plan how storage capacity will be added, reserve the required floor space, and ensure that the future power and cooling requirements can be met.

Thin Provisioning: Scenario #1

A telecom company has a Call Detail Record (CDR) archiving storage project for storing historical records that is planned to consume 180 TB over a three-year period (continuous stream of five TB per month). Capacity planning for this project can be done in several ways:

- Using traditional architectures, a storage solution of 180 TB can be built, with all volumes pre-defined. This ensures smooth operation of the project without downtime, but also means that:
  - The capital expense is paid upfront for three years.
  - On average, half of the power and cooling expenses are wasted on unused storage capacity.
  - Floor space is required from Day One of the initial installation.
  - Technological improvements in storage density and power reduction during these three years are not leveraged.
- Using traditional architectures, a storage facility of 60 TB can be built and expanded with 60 TB each year. This alternative ensures the following:
  - Power, cooling, upfront capital payment, and use of out-dated technology are reduced from three years of lack of use to just one year. Yet, significant resource costs remain.
  - Each year, as new storage is provisioned, the system needs to undergo a complex resizing process related to the reconfiguration of storage, hosts and applications. This process involves high operational costs and scheduled downtime that, in some environments, is unacceptable.
- Using thin provisioning, a 15 TB system is installed, with a logical size of 180 TB, and quarterly capacity upgrades of 15 TB. This alternative ensures the following:
  - Capital expense related to storage procurement is spread over time
  - Power, cooling, and floor space consumption are only for required storage
  - New storage solutions with higher density and lower power consumption can be deployed as they emerge
  - No downtime is required for resizing volumes

As one can see, thin provisioning provides the best of both worlds: savings in direct and indirect costs, and higher availability.

Thin Provisioning: Scenario #2

A financial institution is evaluating its future storage needs. A requirement for an additional 200 TB of capacity is needed, serving 15 different storage projects, each requiring 10 to 30 TB.

In the previous year, more than half of their projects used only 50 percent of their allocated storage, while 25 percent required capacity upgrades within the first year. Eventually, more storage had to be procured to supplement the projects with capacity requirements, even though total aggregate usage was below the forecast. It is quite possible that the same patterns will repeat this year.

Capacity planning for this financial institution can be done in one of two ways:

- Using traditional storage architectures, allocate according to the forecast. A storage solution of 200 TB will be purchased, and to accommodate the needs of the projects requiring upgrades, storage will have to be expanded an additional 50 TB throughout the year. Aggregate capacity in use is only 150 TB. Using this alternative means that:
  - Total storage purchased is 250 TB, together with power, cooling, and space overhead expenses.
  - Capacity will have to be added to projects that require additional capacity, which means administration work and application downtime.
- With thin provisioning, only a 100 TB solution will be purchased upfront, while users give an aggressive prediction of the storage needs. Overall, 300 TB will be defined (compared to the original 200 TB) and physical storage can be added when needed. Eventually, 150 TB will be enough to satisfy all project requirements. This means that:
  - Storage spending makes up only 60 percent of the alternative; power, cooling, and floor space expenses make up an even smaller percentage (since storage is added over time).
  - No volume resizing is needed and therefore no administrative effort, and there is no application downtime.

Managing Thin Provisioning

Like other computing strategies that allocate shared resources, thin provisioning can potentially reach a situation where demands for additional storage capacity cannot be met. In some cases, such as communication or virtual memory, a resource problem would translate into a performance problem. In other cases, such as file system quotas, a failure of the statistical paradigms may have more severe implications.

A failure to provide enough storage capacity through thin provisioning must be avoided. If physical capacity is not available, the relevant volumes will not accept any new writes. Such a locking of the system for any write command is equivalent to shutting down the application, a scenario that must be avoided at all costs.

Monitoring Capacity Usage

The IBM XIV Storage System provides “thin provisioning monitoring,” preventing the risk of running out of physical capacity. It constantly monitors storage consumption and allows the storage administrator to configure the system to send notifications when space utilization exceeds a user-defined threshold. These notifications can be sent to various destinations, either as SNMP traps, or e-mail or text messages. The severity and destination of the notification depends on the actual consumption threshold passed.

The XIV system can be configured so that if a given event notification is not handled within a certain time period, another event notification is sent to the same addressees or to a broader distribution list.

Limiting Thin Provisioning to Specific Applications

Another important feature of the XIV system is its ability to manage thin provisioning per “storage pool.” The concept of a storage pool is unique to the XIV system: a storage pool is a logical entity, defined in the system, which contains a group of volumes and their snapshots. The total capacity consumed by volumes and snapshots that belong to the same pool is limited by the size of this pool.

In the XIV system, a thin provisioning policy can be defined per storage pool. Each pool can have its own hard capacity (which limits the actual disk space it can consume) and soft capacity (which limits the total size of volumes defined in it).

The separation of thin provisioning per storage pool is essential in limiting the effect of running out of physical disk space. Thin provisioning management is performed per pool and running out of space in one pool does not impact other pools.

An example of this is when a data center has backup-to-disk and ERP applications running in its environment. The backup-to-disk uses thin provisioning with 100 TB of logical volumes and only 50 TB of physical space. The ERP application needs only 10 TB and thin provisioning is not used. With a poorly-managed thin provisioning implementation, backup-to-disk storage could potentially impact the ERP application. Such a scenario is unacceptable; its very existence may rule out implementing thin provisioning.

With the XIV system, an end-user can use two storage pools: one for the backup-to-disk one for the ERP. The ERP storage pool will have identical soft and hard capacities and thus never be locked as a result of having run out of physical space. On the other hand, the storage pool used for the backup-to-disk application can be configured with thin provisioning (that is, with soft capacity larger than hard capacity) and provide all the advantages of thin provisioning.

Administrating and managing Storage Pools in the XIV system is extremely easy since pools are a purely logical entity. One can always resize a storage pool, move storage capacity between storage pools, or move volumes between storage pools. Pools are not associated with any physical entity, such as a disk drive or module. In fact, all storage pools are distributed equally in the XIV system among all hardware components.

Reclaiming Logically Unused Capacity

The XIV system also has a sophisticated mechanism for reclaiming unused areas of the volume, even after they have been defined as “used.” The effect of this feature is that used capacity can actually be decreased if used data is logically erased by writing all zeros. This capability can be leveraged to reclaim space in several scenarios:

- When migrating volumes from legacy storage equipment to the XIV system, parts of the volume which are all zeros are marked as “unused” and will not consume physical space.
- When an application writes a long sequence of zeros, the relevant part is marked as unused, even if previous information was not zero.
- As a background process, the XIV system scans volumes and searches for long areas that are all zeros, marking these areas as unused.

This reclaiming process provides huge potential savings, especially when migrating storage from legacy systems. The operational cost reduction associated with the XIV system is not the only advantage of thin provisioning (power, cooling and floor space per TB). An additional advantage is the ability to take an existing volume on a legacy system, migrate it to the XIV system, and reclaim what was once used space. The migrated volume may now require only 30 percent of the original capacity.

Thin Provisioning and Snapshots

The XIV architecture seamlessly integrates thin provisioning with snapshots. The system automatically thin-provisions snapshots, consuming physical space only when a delta exists between the snapshot and master volume, or between subsequent snapshots of the same volume. For more information on the XIV system’s snapshot implementation and features, see the XIV white paper “Snapshots Reinvented.”

Summary

Thin provisioning is an important technological improvement that can solve many storage problems. The ability to separate the logical view of the system and the actual capacity provisioned and procured provides a significant cost savings, while enhancing ease of management.

However, traditional thin provisioning comes with several management challenges. Thin-provisioned systems need to be carefully planned and managed, since new risks for storage availability exist when thin provisioning is deployed.

The XIV system provides a leading thin provisioning architecture, which is tightly coupled with advanced snapshot and snapshot management technologies. With the XIV system, you can:

- Use thin provisioning to limit storage expenses to the actual capacity used.
- Reduce management overhead and the operational expense of resizing volumes.

- Limit the effect of thin provisioning to specific applications, safe-guarding your most critical applications from running out of space.
- Get notifications on thin provisioning usage.
- Reclaim capacity that is no longer in use.

With the XIV system, thin provisioning can be deployed easily and managed effectively, and result in significant storage cost savings within an enterprise IT environment.