


# NetApp Thin Provisioning - Doing More With Less

Silverton Consulting, Inc.  
StorInt™ Briefing

## Introduction

Reducing or deferring a datacenter's capital and/or operational expense has become a paramount concern in today's challenging economic climate. Many storage vendors' promote functionality and features that can purportedly provide potential cost savings but few actually deliver on this promise. However, one proven feature offered in the open system marketplace since the early 2000's is thin provisioning. This storage feature can readily provide immediate and ongoing capital and operational expense savings. The potential cost savings of this feature are significant and are being more and more recognized. Indeed, as can be seen in Figure 1 below from the TheInfoPro (TIP) Wave 13 survey of storage professionals and CIOs, thin provisioning has emerged as the second most important virtualization functionality for storage systems. Consequently, thin provisioning will have a prominent place in storage system deployments over the near future.

## Virtualization Functionalities' – Importance Rank



(12/14/09): Full Sample, Wave 12, n=33; Wave 13, n=92.

Figure 1 TheInfoPro Wave 13 Virtualization Functionality Importance survey results

### What is storage thin provisioning?

Storage thin provisioning is a feature which allows systems to provision more storage than is physically present for users or applications. Graphically, thin provisioning can be likened to a set of uninflated balloons in an empty room. The balloons are LUNs or storage volumes while the room is the actual physical storage capacity of the storage system. As balloons are inflated or volumes are written, they expand and consume the actual, physical room. All the balloons can expand at their own independent consumption rate as long as the total aggregate space demand does not exceed the capacity of the room.

In contrast, fixed or non-thinly provisioned storage is like constructing rooms for a multi-floor building. Each room or LUN/storage volume must be allocated some portion of the floor space prior to construction. The actual usage of the room or LUN does not affect the storage room dimensions; rather the dimensions are predetermined and are often used inefficiently. Changing the room's original space configuration is very disruptive and usually expensive.

In storage terms, non-thinly provisioned storage systems can result in poor space utilization rates. In fact, most administrators with non-thinly provisioned storage systems typically use only a small percentage of actual physical space. The remainder of the storage is reserved, yet empty space. Thin provisioning this same poorly used space can effectively allow an administrator to regain 50 to 70 percent of the empty space. This increase in utilization is achieved by making better use of the provisioned storage. With thin provisioning, an administrator essentially creates a pool of free space that all applications can share from, greatly improving utilization rates.

### Immediate cost savings of thin provisioning

The benefits and cost saving of thin provisioning are significant, multi-faceted and include:

- **Higher capacity utilization generates both immediate and ongoing capital outlay savings.** With thin provisioning, utilization can be increased by making better use of physical capacity present. Unused space is not locked up by individual applications. This upfront contraction of the initial physical space needed significantly reduces the initial capital expense outlay. Subsequent capital expenses are also deferred since capacity expansion is not required until the storage is actually used.
- **Powering and cooling less storage generates savings.** By deploying less storage with thin provisioning, less power and cooling is required, again generating immediate and ongoing savings. However, for some data centers, thin provisioning not only generates significant operations savings but also is critical in sustaining ongoing operations. For example, some data centers have restricted service limits from their power providers and so, cannot plug-in additional equipment without renegotiating power contracts. In some situations even renegotiating may not be an option making the ability to thin provision

## NetApp Thin Provisioning – Doing More With Less

storage even more paramount. Other industries, where government imposed environmental constraints are critical, may also benefit from thin provisioning.

- **Easing space constraints leads to less administration headache.** Doing more work with less space allows administrators to approve new service requests more often and more easily. For example, installation of a new application and its corresponding storage requirements can be done by configuring the storage needs even though total storage configured exceeds the actual physical space. The only absolute constraint is that the actual application data does not totally consume the available physical storage space.

### A Caveat about thin provisioning

Thin provisioning can undoubtedly help many data centers improve their abysmal real space utilization rates. However, with thin provisioning in use, monitoring of storage space consumption becomes even more critical. Lack of diligence here could have dire consequences because actual storage space may be entirely consumed. On the other hand with thin provisioning, given proper diligence, capacity additions can be timely made with little disruption of ongoing operations and immediate expansion of the thinly provisioned storage pool.

### Current state of thin provisioning

Most major storage vendors have introduced some sort of thin provisioning capability today. However, as expected many of these vendors have named their process something other than thin provisioning. While NetApp refers to thin provisioning, newer vendors have touted “virtual” or “dynamic” provisioning. All of these products attempt to accomplish the same thing – allow administrators to over-provision physical storage and increase effective capacity utilization.

The following is a list of criteria that should be considered when evaluating thin provisioning technologies, with NetApp’s implementation model noted:

- **Storage pool or pools.** - Thin provisioning normally requires some sort of homogeneous pool of storage. A few vendors aggregate their total physical storage into one storage pool, which can be carved up and allocated to any thinly provisioned LUN that needs storage. Other vendors segregate their storage into multiple pools, each of which may only contain thin or non-thin provisioned storage. Still other vendors support multiple pools of storage, each of which can contain both thin and non-thin storage.
  - **NetApp supports intermixed thin and non-thin pools** – intermixed pools provide more flexibility.
- **Page size.** – Storage pools are typically divided into separately addressed storage units often called pages. For some vendors page size can be measured in MB, for others in KB. Page size is a trade-off between storage utilization and management overhead. For example, larger pages cause more storage to be wasted but require less overhead to allocate and manage. In contrast, smaller page sizes require more management overhead but waste less space.

## NetApp Thin Provisioning – Doing More With Less

- **NetApp's page size is 4KB** - a smaller page size optimizes the use of storage space. However, NetApp needs no additional overhead because they employ this page size for all storage allocation, thin or not.
- **New space allocation.** – When new data is written to thinly provisioned volumes, a physical page of storage must be available to store the write. How much storage is allocated and the time of its availability varies by vendor. Some vendors reserve multiple storage pages for volume writes to eliminate interrupting write operations. Then, as pages are consumed the storage system replenishes this reserved capacity in the background. Other vendors only allocate a page of storage to a volume as required to satisfy data write. Similar considerations also apply to the initial configuration of a volume, i.e., some vendors reserve space to a volume when first configured, while others do not.
  - **NetApp pages are allocated at time of write** – allocation at write time optimizes use of storage. In addition, no space is allotted when a volume is first configured only when it is written.
- **High watermark allocation.** – When a new high block number is written to a LUN, a few vendors allocate enough storage to write all blocks up to this new highest block number, even though they may never be written. Other vendors support sparse space allocation, and as such, only allocate storage to blocks that are actually written regardless of block number.
  - **NetApp pages are only allocated to blocks/files that are written** – high or low block numbers make little difference to NetApp thin provisioned storage. For example, no more storage is allocated to a LUN for a block write to block number 10,000 than is allocated for one to block 0.
- **Storage space contraction.** – Advanced file systems can release storage that is no longer needed. Some vendors can even release storage when a block of storage is written as all 0's. Still other vendors keep a page, once written, as long as the volume is defined.
  - **NetApp pages are freed up only when a file, LUN or volume is deleted** – regardless of the data pattern written, once written it's retained until deleted.
- **Thin snapshots/clones/mirrors.** – Thinly provisioned volumes are often snapshot, cloned, or mirrored to other volumes. Some vendor snapshots, clones, and mirrors automatically inherit thin provisioning attributes of an original source volume. For those mirrors and/or clones, only data that has been written to the original volume need be copied, other data can be ignored. Other vendors do not support thin mirrors, thin clones, and/or thin snapshots. In this case, even though a source volume is thin provisioned, the target of a snapshot, clone and/or mirror must be non-thin provisioned and all volume data will be copied, whether written or not.

## NetApp Thin Provisioning – Doing More With Less

- **NetApp clones and mirrors can inherit thin provisioning attributes of source volumes.** Furthermore, NetApp snapshots only record changed data so they are inherently thinly provisioned.
- **Space warning thresholds and response.** – As writes consume real physical storage, free space can become scarce and as such, operations need to be notified by alerts, emails or SNMP traps to add more storage. Most vendors provide multiple warning thresholds to alert operations at predetermined low space conditions. Also, some vendors offer automated responses to free space after a system warning. For example, some systems offer the option to automatically delete the oldest snapshots when space runs low.
  - **NetApp supports multiple warning thresholds and automated responses** – with NetApp’s Operations Manager and Provisioning Manager software, numerous warning thresholds and automated space handling capabilities are available.
- **Thin provisioning pricing.** – Some vendors charge separately for thinly provisioned storage space. Often this surcharge is a function of how much storage is thinly provisioned. Other vendors offer thin provisioning at no additional charge.
  - **NetApp thin provisioning functionality comes free with the purchase of any NetApp storage** – there are no additional charges for thin provisioning.

### Specifics of using NetApp thin provisioning

NetApp utilizes storage in pools made up of physical capacity not exceeding 16TB and consisting of one or more RAID groups of storage devices. With Data ONTAP 8.0 coming in early 2010, NetApp removes this size constraint and depending on the platform, storage pools can reach 100TB. These storage pools are called **aggregates** by NetApp and are divided into one or more **volumes**. Each volume can contain one or more **files, Qtrees** or **LUNs**. Files are CIFS or NFS files, Qtrees are quota controlled file aggregations and LUNs are SAN or IP/SAN attached volumes. An aggregate storage space also holds Snapshots for volumes and space reserves.

A NetApp volume can be configured as a **traditional volume (TradVol)** or a **flexible volume (FlexVol)**. TradVols cannot be expanded, cannot be thinly provisioned and reserve all their space at configuration time. On the other hand, FlexVols can be extended and thinly provisioned resulting in little space being allocated to it initially.

When configuring a NetApp FlexVol for thin provisioning, a **space guarantee** attribute is specified. FlexVol space guarantees can be one of three types:

- **Space guarantee=volume** – The default setting of a volume where provisioning is fixed; that is, not thin provisioned.

## NetApp Thin Provisioning – Doing More With Less

- **Space guarantee=none** – The volume is thinly provisioned.
- **Space guarantee=file** – The volume can hold files and Qtrees that are thinly provisioned and LUNs that are fixed or non-thinly provisioned. As such, space guarantee = file supports a mixture of thin provisioned files and non-thin provisioned LUNs.

The specific NetApp command to create a FlexVol is as follows:

```
vol create vol_name [-l language_code] [-s {volume|file|none}] aggr_name  
size{k|m|g|t}
```

where **s {volume|file|none}** - indicates the **space guarantee** for the volume.

Under this overall versatile design, NetApp's aggregate can contain any combination of FlexVols and TradVols. Moreover, FlexVols with and without space guarantees can co-exist within the same aggregate.

### Using Operations Manager to monitor NetApp space utilization

Even with this simplified architectural design, NetApp has further simplified the thin provisioning process by developing software to help administrators more easily and more effectively manage space utilization of their aggregates, collectively or independently. With **Operations Manager**, NetApp's add-on software package, an administrator can easily specify two distinct space thresholds for any aggregate. Specifically, Operations Manager thresholds are:

- **Aggregate Nearly Full (%)** – specifies a threshold of aggregate space consumption that will trigger an aggregate-nearly-full event. This value defaults to 80%.
- **Aggregate Full (%)** – specifies a threshold of aggregate space consumption that will trigger an aggregate-full event. This value defaults to 90%.

These thresholds can be defined differently to account for the projected installation time of new equipment and other critical variables. However, even more significantly, meeting these thresholds triggers multiple, varied alarms like emails, pager calls or SNMP traps. In addition, scripts can be invoked to reclaim storage space. As such, NetApp's Operation Manager software streamlines and automates much of the more mundane, yet critical information flow regarding a potential storage space (aggregate) problem as well as initiate steps to avert the problem.

# NetApp Thin Provisioning – Doing More With Less

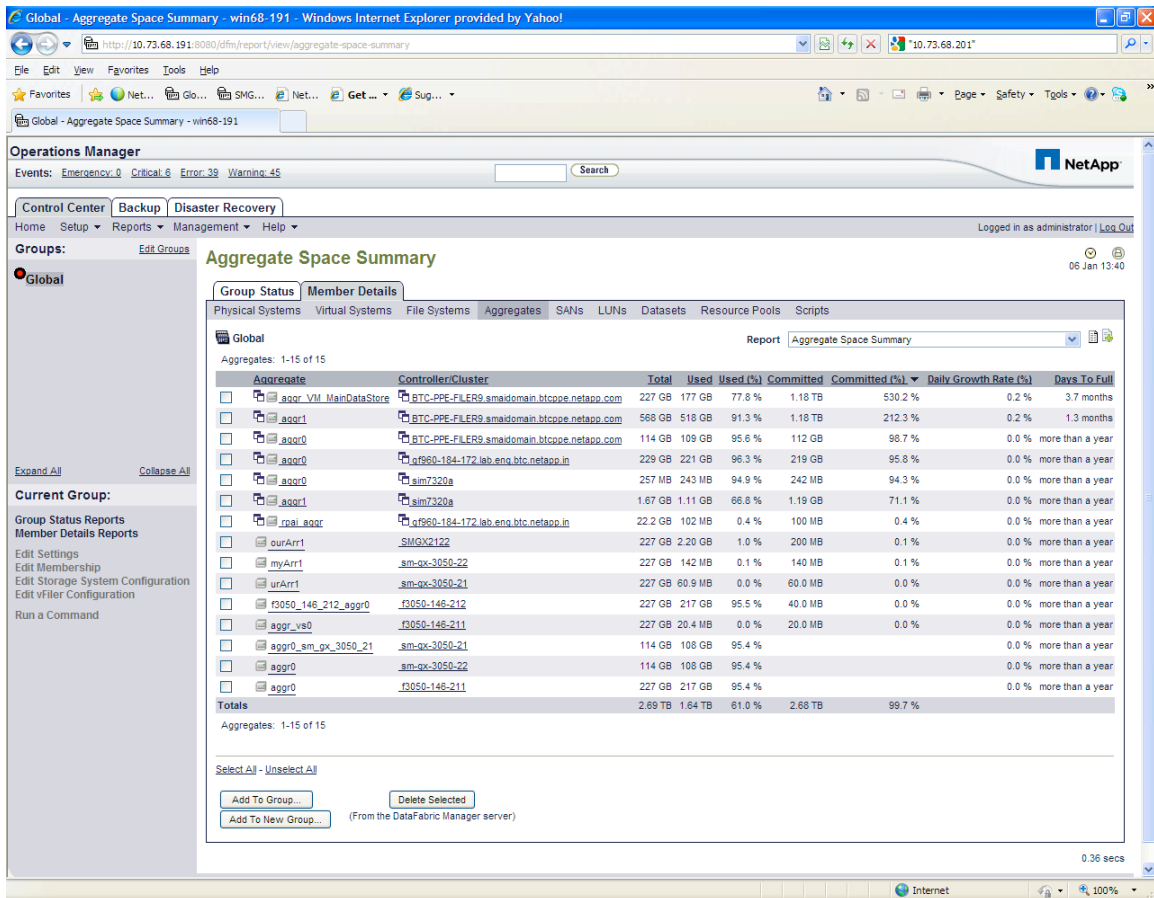


Figure 2 Operations Manager - Aggregate Summary Report

## Using Provisioning Manager to complement NetApp thin provisioning

Another NetApp software program complementing and easing the use of thin provisioning is its **Provisioning Manager**. This sophisticated software can automate many predefined responses to predefined events. Furthermore, this software can execute the responses over a data center's total storage **resource pool**, which may include multiple aggregates, volumes or even storage systems. For example, given the use of Provisioning Manager and a storage system with only one type of RAID storage, a new volume will be assigned by default to the aggregate with the most free space and the least amount of IO activity. In effect, the software is automating the thin provisioning allocation process.

Provisioning Manager can also be used to specify policies to invoke automated activities when free space becomes low in a resource pool. Automated responses can be as simple as deleting the oldest snapshots to as sophisticated as migrating volumes.

### Considerations when using NetApp thin provisioning

#### Good thin provisioning candidates

Some situations where NetApp's thin provisioning offering is particularly useful include:

#### Large quantity of volumes

Thinly provisioned storage, including NetApp's product, works best when used across lots of storage, i.e., across a large quantity of volumes. For example, with 10TB of NetApp physical capacity, it's generally better to thinly provision 20 volumes at 1TB each than thin provision 2 volumes at 10TB each. With only two volumes sharing 10TB of physical space, one quickly growing volume could easily impact the storage available to the other. In contrast, with 20 volumes a few extreme consumers of storage are commonly offset by a few with much smaller consumption. In addition, average utilization estimates over a large number of volumes are more predictive of actual space usage.

#### Predictable storage consumption

NetApp thin provisioning also works well with predictable storage consumption. For example, applications writing storage at a predictable rate would be easy to monitor for usage trends. In fact, monitoring thresholds could be set more aggressively and storage added only when needed. Some prime candidates for NetApp's thin provisioning might include:

- **User directories** – Generally large numbers of small volumes are deployed, and in total, storage consumption is predictable and proportional to time in use.
- **Virtual machine storage** – Similar to user directories, many VMs are usually deployed and in total, storage consumption is predictable.
- **Data warehouse** – Such data is usually read intensive and once created, infrequently written at best. Also, the rate of change (or write) is very predictable. Moreover, some database data can be sparsely written, i.e., much of the database is empty of data.
- **Secondary storage** - Secondary storage is read often and written infrequently making for very predictable consumption.

#### Unlikely thin provisioning candidates

Storage thin provisioning may encounter interruption of service when space is depleted. As such, thin provisioning may not be suitable for mission critical storage where out of space events could not be successfully tolerated.

Thin provisioned storage may also be undesirable with unpredictable storage consumption. This volatility can make the determination of warning thresholds overly conservative and restrictive.

## NetApp Thin Provisioning – Doing More With Less

### NetApp thin provisioning and various O/S

#### Using NetApp thin provisioning with Windows

External storage on Windows can be based on networked file shares, NTFS file systems mapped to storage LUNS, or devices that are application controlled. NetApp **networked file shares** (NFS) are always good candidates for thin provisioning as the NetApp file server controls space consumption. Carefully configured CIFS file systems are also good candidates for thin provisioning because of their slow growth. In contrast, NTFS volumes are not good candidates for thin provisioning because storage on them is not re-used until the storage space is consumed. **Application controlled storage** presents a mixed case. Some applications utilize sparse writing or efficiently reuse space, thus being good candidates for thin provisioning. Other applications automatically format the whole volume space and/or do not reuse storage space making them poor candidates for thin provisioning.

#### Using NetApp thin provisioning with VMware

VMware defines virtual machine storage as VMDKs which can be hosted on **native NFS file systems, VMFS storage using FCP or iSCSI LUNS, or raw devices**. With NetApp controlled storage space, NFS file systems are always good candidates for thin provisioning. However, non-thinly provisioned VMFS VMDKs make poor candidates even though they could be initially minimized. VMDKs hosted on raw devices could be either a good or poor candidate for thin provisioning depending on how the VM actually uses the raw storage.

#### Using NetApp thin provisioning with Unix and Linux

Symantec storage foundation file systems are available for many Unix and Linux distributions. Storage foundation is very NetApp thin provisioning friendly because space is re-used before new space is consumed.

On the other hand, Unix or Linux systems not using Storage Foundation can be either good or poor candidates for thin provisioning. If the O/S supports NFS file servers, such storage can be a good candidate. For other native file systems, space reuse and/or pre-formatting varies considerably, and as such, may or may not be good candidates for thin provisioning.

## Thin Provisioning Case Study: The College of St. Rose

### NetApp CIFS/iSCSI/DR Design

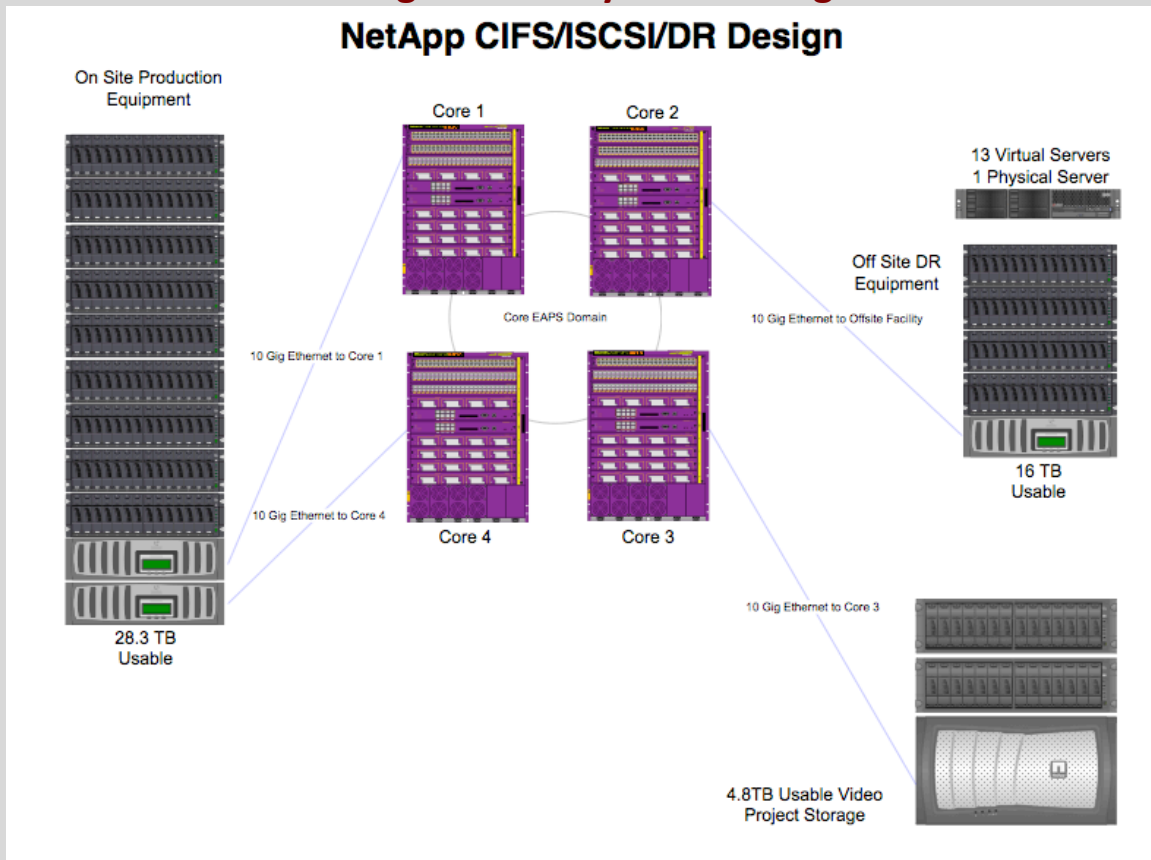


Figure 3 St. Rose storage configuration

### St. Rose's data center

St. Rose is a college in the Eastern United States with an approximate enrollment of 5,000 students. St. Rose's datacenter is in an older building on campus and has space, powering and cooling constraints. St. Rose uses VMware to virtualize their servers to reduce physical infrastructure. Before virtualization, St. Rose had 50-60 servers, now they have 15 physical servers with about 80 VMs running today. They also support about 2000 Exchange mailboxes and ~1000 user desktops.

For storage St. Rose uses NetApp systems exclusively (see Figure 2 above). They have one NetApp FAS3050 cluster with 44TB of usable space for production storage, another FAS3140 with 4TB dedicated to student video activity and a remote hot DR site using one FAS3020 with 16TB. Some production data is mirrored to the hot DR site. The NetApp storage is configured as FC, iSCSI or CIFS storage.

Aleks Pavlinik, Senior Network Administrator – Communication Systems and one other person manage the data center infrastructure (non-desktop). This includes

## NetApp Thin Provisioning – Doing More With Less

servers, storage, and wired and wireless networking for the campus and DR sites with over 200 access points across campus.

### **St. Rose and thin provisioning**

About a year ago St. Rose was running low on storage. At that time they were running at 90% storage consumption but actually had very low actual space utilization. They either had to buy more storage or try something else. At 90% consumption, Mr. Pavlinik and the other administrator were constantly scrambling to shift storage around for applications and users.

Given their economic constraints at the time, St. Rose decided to try thin provisioning on some of their user CIFS storage. Pavlinik had noticed previously that CIFS users consumed little actual storage, typically only 15-20% of allocated space and so, would be a good place to start. He further reasoned that the NetApp thin provisioning capability carried no additional cost and was easy to implement with a single command.

When Pavlinik first enabled thin provisioning for CIFS storage, he noticed their FAS3050 storage consumption went down from 90% to 86% then to 80%. At that point, he knew he was onto something that could solve St. Rose's storage constraints.

### **St. Rose today**

Given their success with CIFS storage, Pavlinik is now starting to thinly provisioning their VM storage. The problem with VM storage is that once provisioned as fixed (or thick) storage, it has to be moved to be thinly provisioned. Such data migration takes more administrator effort and time. Pavlinik believes using NFS to host future VMDK storage requests should eliminate that problem.

---

**“Nowadays, I no longer have to say ‘no’ to users requesting storage.”**

**Aleks Pavlinik, Senior Network Administrator**

---

Today, about 13TB of their 44TB production storage is thinly provisioned. For this storage, they are only physically consuming 9TB of storage, a 44% reduction. Consequently, one year after implementing thin provisioning, St. Rose runs at 76% storage utilization and uses NetApp's Operations Manager to monitor free space consumption. They have various SNMP traps/e-mail alerts to warn them when free space gets low. Also they use Operations Manager scripts to drive auto-deletion of oldest snapshots whenever free space becomes limited.

Pavlinik reported that it could take 2 months from the storage need identification to the actual deployment of the additional capacity. He believes that the alerts and his continued monitoring of free space will give him plenty of warning to get more storage, if needed.

## NetApp Thin Provisioning – Doing More With Less

### Benefits from using thin provisioning

The most significant benefit from thin provisioning for St. Rose was their ability to indefinitely defer buying additional storage. Further, power consumption remained the same. Thus, thin provisioning and virtualization has given St. Rose the breathing space to better utilize their current infrastructure by deferring new equipment acquisition to a future date.

Another, less measurable benefit Pavlinik noted was that he and the other administrator are less harried today. With 76% storage utilization they have a lot more breathing room and can deploy more applications and users without worry.

---

**“We are now able to deploy more virtual machines. We are really not limited by CPU or memory constraints, it really comes down to storage space.”**

**Aleks Pavlinik, Senior Network Administrator**

---

### Summary

Thin provisioning is a mature technology that has been deployed by 1,000's of data centers today for multiple diverse applications and across a multitude of O/S environments. For the right environments, with proper vigilance, thin provisioning can provide immediate, significant economic benefits. These tangible economic benefits are magnified and cumulative over the long-term and allow data centers to make more effective use of all current and future storage purchases. Indeed, given today's challenging economic environment, thin provisioning capabilities such as NetApp's offering should become a top consideration for a majority of data centers.

---

***Silverton Consulting, Inc. is a Storage, Strategy & Systems consulting services company, based in the USA offering products and services to the data storage community.***