Technical Report

NetApp Thin Provisioning
Increases Storage Utilization With On Demand Allocation

Carlos Alvarez, NetApp
May 2011 | TR-3563

ABSTRACT

This guide introduces NetApp® thin provisioning, describing the challenges of traditional approaches, the benefits of thin provisioning, and NetApp thin provisioning differentiators. It also discusses how to implement and manage thin provisioning in the environment.

This guide should prove useful to IT and storage decision makers seeking to understand how NetApp thin provisioning can reduce waste and increase storage utilization. Although it covers some implementation considerations, it does not include detailed deployment strategies or best practices.

TABLE OF CONTENTS

LIST OF FIGURES

1 INTRODUCTION

In today's fast-paced IT world, allocating the appropriate amounts of storage across a growing number of different projects is a constant challenge. As part of this process, storage administrators must consider different service-level agreements (SLAs), different data growth patterns, and limited or shrinking IT resources. Allocating storage has traditionally been a tedious and time-consuming task. Traditional provisioning approaches typically yield storage utilization rates that are fairly low.[1] Combined with the continuing explosion in data growth, it's understandable how storage has quickly become one of the largest IT budget line items.[2]

Thin provisioning is a powerful tool to address these challenges. It allows storage administrators to allocate auto-growing storage, while moving toward a pay-as-you-grow model. This model simplifies storage provisioning, and it also reduces high up-front storage costs because the storage can be purchased as needed.

NetApp introduced an enterprise-class thin provisioning storage solution in 2004, including the tools to effectively monitor and manage thin provisioning. NetApp thin provisioning is included at no charge with NetApp storage systems, and can be easily turned on for all areas of the storage system or for specific target areas, for both NAS and SAN.

This technical report discusses the problems associated with traditional storage provisioning and describes how NetApp thin provisioning resolves those problems. It also includes information about using thin provisioning in different configurations and discusses how to implement and manage thin provisioning in new and existing environments.

Although not covered in this document, NetApp thin provisioning can also be applied to non-NetApp storage through the use of the NetApp V-Series gateway system. For more information, see TR-3461 V-Series Best Practice Guide.

2 THE PROBLEM

Using a traditional storage provisioning approach, dedicated storage is purchased for each project request. These projects can be from the same group or from different groups in an organization. The storage is typically sized by guessing how much storage will be needed over the next 2 to 5 years, depending on the expected life of the project. It is common for the guessing process to include generous buffer space to avoid running out of storage at some point during the project. This approach results in a large amount of storage being purchased for future growth, but only a small fraction of that storage being used for the initial phase of the project. That translates to wasted storage, because most of the storage is not used on initial deployment.

Also, when using a traditional storage provisioning model, the allocated storage is dedicated to a particular volume or LUN at the time of its creation. That dedicated storage is tied to that particular volume or LUN, and is not available to any other projects, even though it is empty storage. This provisioning model, also referred to as thick provisioning or fat provisioning, is inefficient, with many visible and latent costs.

---

[1] Traditional storage provisioning typically results in low storage utilization, according to NetApp My AutoSupport analysis and the Silverton Consulting brief "NetApp Thin Provisioning – Doing More With Less."

[2] Storage is one of the top IT capital expenditures, according to a 2009 IDC Executive Interview Study.

2.1 HIGHER COSTS AND LOWER RETURN ON INVESTMENT

LOW STORAGE UTILIZATION

Since the storage will not be immediately filled to capacity upon its installation, it is common to have a significant amount of storage that goes unused for years.

HIGHER CAPITAL AND OPERATING EXPENSES

Buying all the storage up front can result in increased capital expenses. Storage requests are often inflated by the time they get to the storage team, because requestors often add a buffer to their needs assessments. This results in bloated storage provisioning.

More waste is evident when considering operational expenses. Although storage utilization may be low for long periods of time, all the hardware must still be racked up, powered, and cooled. This takes up space in the data center for storage that is not being actively used, and can cost lots of money for large amounts of dedicated storage that is sitting empty.

LOW RETURN ON INVESTMENT

In addition to the return on investment (ROI) for the storage being driven down by high capex and opex, there are additional expenses that further reduce the return of the storage. For instance, additional staff is needed to manage an oversized storage infrastructure. Also, a bloated storage infrastructure, where each allocation of storage is bound to a single project, can create a premature limit on the number of applications or users using the storage system. The ROI is further lowered in an environment that requires additional copies of the storage, causing the bloated infrastructure to be replicated to additional locations.

2.2 LACK OF FLEXIBILITY AND AGILITY

In a traditional, thick, provisioning model, once storage is provisioned for a particular volume or LUN, it is forever tied to that storage entity. Regardless of the fact that there may be free space in a particular volume or LUN, it cannot be shared with other volumes or LUNs in the system. This leads to rigid, silo-based storage configurations where free space cannot be shared on an as-needed basis across different silos.

DELAYS IN INITIATING NEW PROJECTS

In today's dynamic environment, with ever-changing priorities and requirements, rigid silo-based provisioning creates a bottleneck for the storage administrator. Silo-based storage does not allow storage to be quickly re-provisioned to keep up with the dynamic environment where project priorities can change quickly. Every time a new storage request is made, the administrator has to purchase and add new storage to the system, instead of leveraging existing unused storage from other entities.

Figure 1 below shows how the introduction of a new project requires an additional storage purchase, since thick provisioning does not allow storage from existing projects to be used.

Figure 1) Thick provisioning new projects.

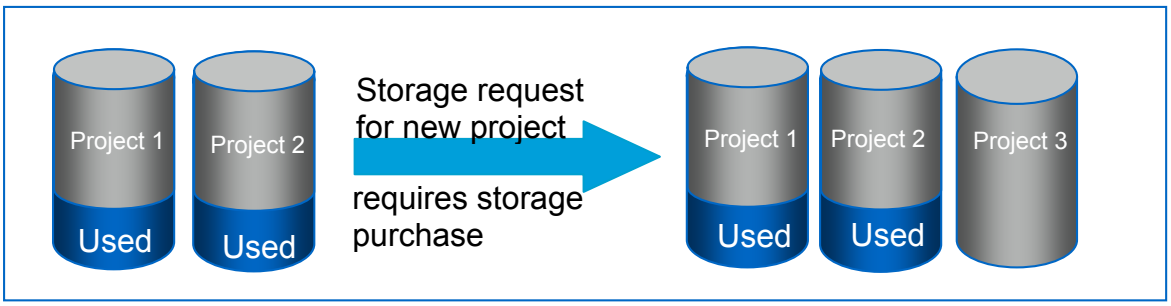

INEFFECIENCIES IN HANDLING CHANGES IN PRIORITY OF EXISTING PROJECTS

In environments where the priorities for different projects change over time, the traditional provisioning method does not let the storage administrator move storage from an underused, low- priority project to a higher-priority project without disruption to data access. To free storage from the low-priority project, access to its storage must be disabled while the data is migrated to a new, smaller storage allocation. Once the storage is freed, it can be combined with other free storage to create a large storage allocation. If there is not enough unused storage available to create a larger storage allocation, then additional storage must be purchased. And finally, access to the high-priority storage must be disabled while the data is migrated to the new storage allocation created for the high-priority project. This is an inflexible and highly inefficient approach from both a financial and a storage management perspective.

2.3 OPERATIONAL AND MANAGEMENT COMPLEXITY

Storage administrators are responsible for allocating storage in a productive and efficient manner. However, the limitations of a traditional storage provisioning approach make this very challenging. As the environment grows, the increasing number of projects and storage systems compounds the challenges.

OPERATIONAL COMPLEXITY

Operational responsibility in a growing environment is a very complex process with wide visibility and serious consequences if mistakes are made. If storage is not available for a project, then all production associated with that project comes to a halt. Because a traditional provisioning approach requires that all storage must be allocated for each project up front, project owners and storage administrators are forced to guess about the amount of storage that will be needed in the future. As previously mentioned, it has been a common practice to overestimate future needs, creating bloated storage. Although this approach helps reduce the initial complexities and risk, it creates long-term issues, including accelerated use of floors pace, power needs, and cooling requirements. If left unattended, these factors can grow into the need for a new data center.

MANAGEMENT COMPLEXITY

Storage headaches intensify as IT attempts to manage this bloated storage. Managing rapidly growing storage by using traditional methods, most of which is provisioned but never used, includes tasks such as monitoring, software maintenance, hardware maintenance, and data migration. Managing the management tasks can be a challenge in itself.

3 THE NETAPP SOLUTION

Included in storage systems as part of the NetApp Unified Storage Architecture, NetApp thin provisioning offers a way to safely and efficiently allocate all requested storage without the waste of bloated storage. NetApp thin provisioning creates all of the expected space logically, but it does not dedicate any storage to any specific project until data is written to the physical storage. The physical storage is available in a

shared storage pool and is dedicated to a project as it is used. The shared storage pool can be expanded as needed.

NetApp thin provisioning provides the capability to do the following:

- Present more logical storage to the hosts than the amount of physical storage present in the underlying physical storage pool
- Dynamically dedicate physical space to volumes or LUNs as the data is being written, instead of pre-dedicating the physical space
- In most configurations, automatically release freed space back to the shared storage pool when data in the volume or LUN is deleted
- Simultaneously use traditional storage provisioning and thin provisioning in the same storage pool
- Combine thin provisioning with other storage efficiency technologies, including NetApp deduplication, NetApp FlexClone® volumes, or NetApp data compression, to achieve even greater savings

Figure 2) Storage thin provisioning.

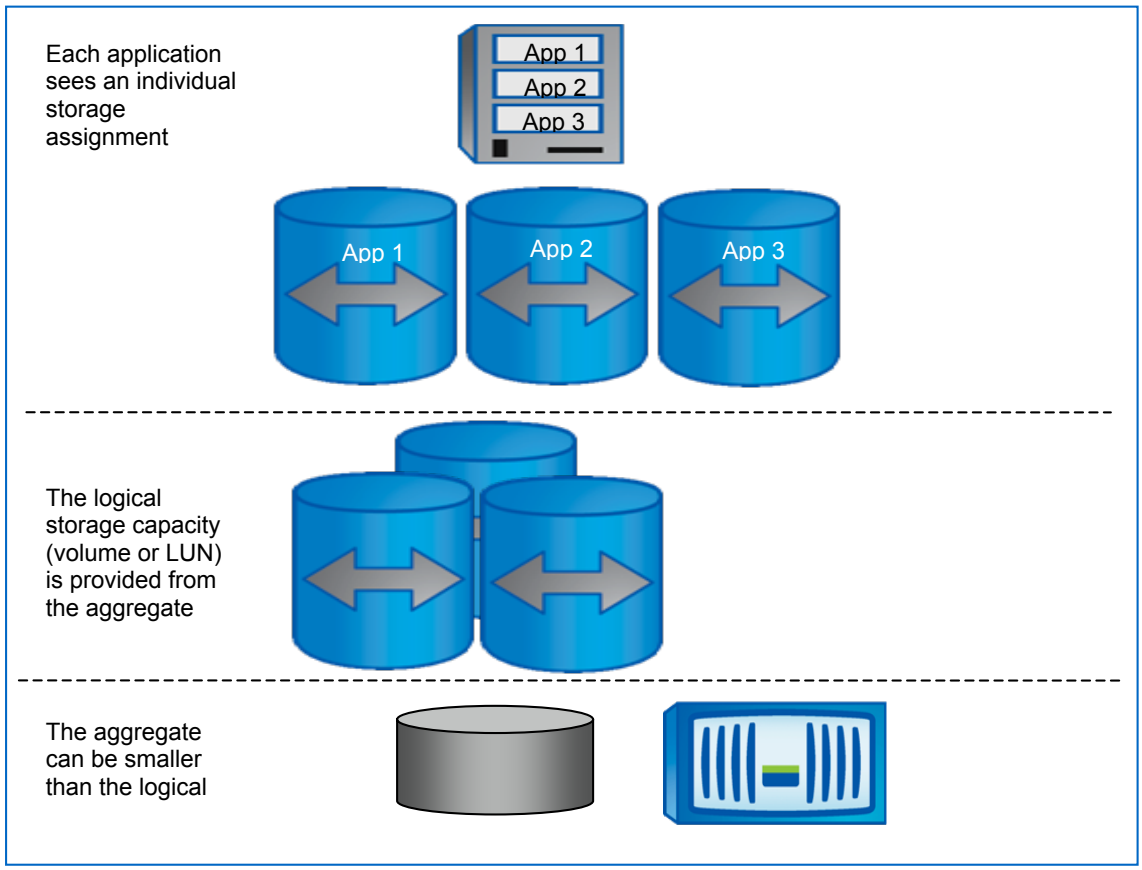

Figure 2 illustrates the basic concept of thin provisioning. The top portion of Figure 2 shows that the applications, app 1, app 2, and app 3, each see their individual allocation of storage, as expected. The middle portion shows that the three allocated spaces have been logically created on the storage system. The bottom portion shows a relatively small amount of storage, the gray cylinder, which is used to support all three applications, and also to provide the shared storage pool for future growth by any of the three applications. Essentially, the small amount of physical storage has been virtualized by the storage system and is being transparently shared by all three applications. This subject is covered in detail in section 5, "NetApp Thin Provisioning Model."

4 THE BENEFITS

The storage community is interested in thin provisioning because of the enormous benefits and flexibility it offers. In addition to providing an agile storage environment, it offers both financial and operational benefits.

4.1 FINANCIAL BENEFITS

Using thin provisioning allows storage to remain available to all applications until the application writes data to it. This approach means that the storage can be purchased over time as needed, rather than all at once up front.

HIGHER STORAGE UTILIZATION

The fact that new storage is added only as the existing physical storage is used reduces the amount of unused idle storage, maximizing storage utilization. It's no longer necessary to maintain unused and underused storage in the infrastructure.

Figure 3 compares traditional provisioning and thin provisioning. In both cases, the storage needed and the storage requested are the same. This reflects the fact that the user requests and uses storage in the same way, regardless of the approach that is used. The difference occurs when the storage administrator purchases the storage. Using the traditional provisioning approach, overprovisioning is the only method available to avoid running out of storage. This approach often ends up with a large amount of unused space. Using the thin provisioning approach, only a fraction of the requested storage is initially purchased (40% in this example). The storage utilization is then monitored, and additional storage is purchased later, as needed. The right column shows what the storage looks like after one year. With thin provisioning there is much less unused space, which translates to a much higher rate of storage utilization.

Figure 3) Thin provisioning can achieve higher storage utilization.

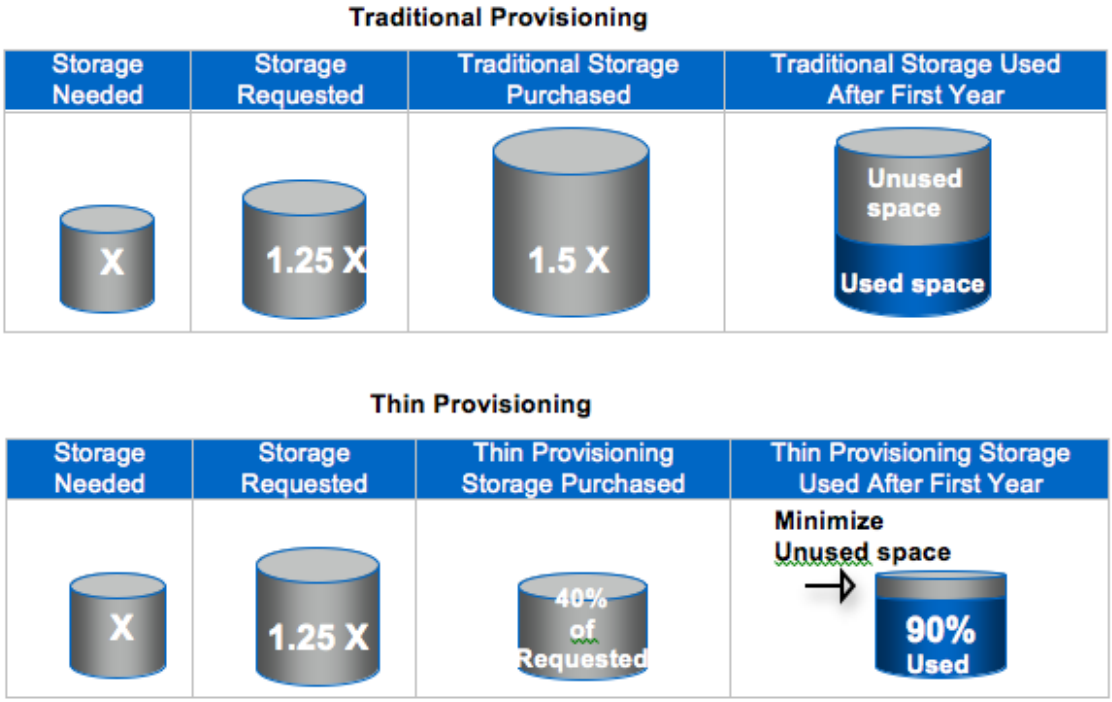

REDUCED CAPITAL AND OPERATING EXPENSES

Thin provisioning can reduce capital expenses. By removing the need to purchase all storage up front for each request, the initial cost of the storage is dramatically reduced and unneeded storage is never purchased. In addition, spreading the cost of the storage over time takes advantage of the lower cost of storage, because prices for storage decline over time. And finally, where multiple copies of the data must be maintained at different locations, the cost is reduced in multiples of the number of copies.

Thin provisioning also reduces operating expenses. Reducing the amount of storage that is purchased reduces the amount of space needed for unused storage, and also reduces the amount of associated power and cooling. Finally, allowing the storage to grow as it is used reduces the staff required to repeatedly review and re-provision storage.

FASTER AND HIGHER RETURN ON INVESTMENT

Higher storage use and lower costs result in a faster return on the investment for the storage infrastructure.

Figure 4 illustrates the advantages of thin provisioning over traditional provisioning. The figure compares the amounts of storage that would be purchased, and the resulting utilization levels, for the first and second years, using both traditional provisioning and thin provisioning. The final row shows the cumulative total costs for each approach in year 1 and year 2.

Figure 4) Financial benefits of thin provisioning.

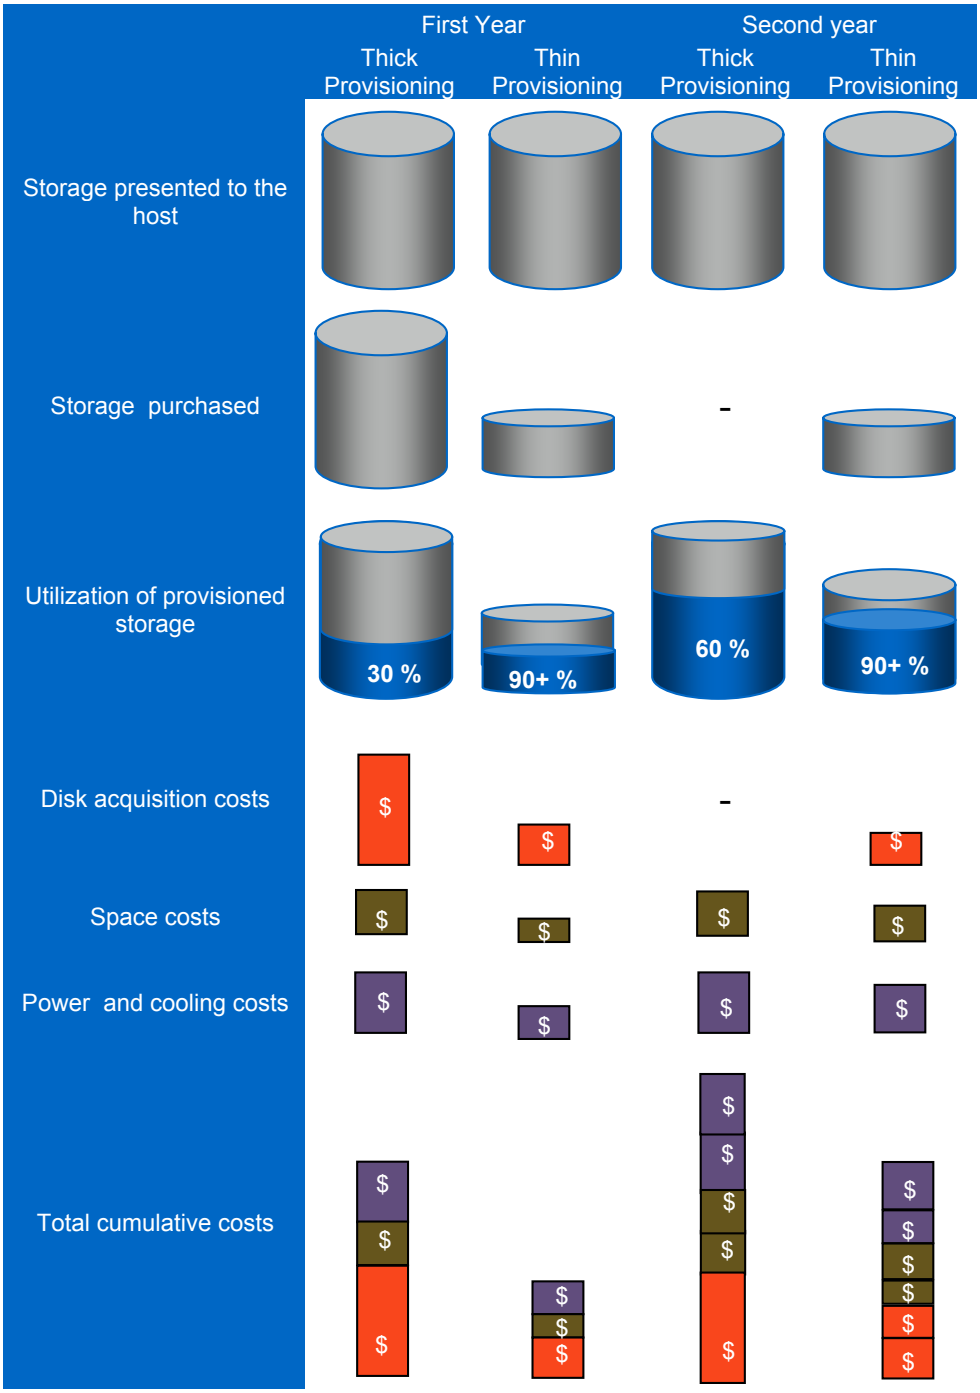

4.2 OPERATIONAL BENEFITS

Thin provisioning offers significant operational benefits that make your storage planning, designing, and purchasing processes flexible and agile.

SIMPLIFY THE STORAGE PLANNING PROCESS

NetApp thin provisioning can significantly simplify the storage capacity planning process by allowing IT administrators to focus their efforts on storage that is actually used. Not having to dedicate all the storage up front mitigates the risk of running out of storage, avoiding the need to make conservative storage-wasting guesses about how much storage will be needed.

Because thin provisioning dynamically dedicates storage for applications from a shared storage pool as needed, storage administrations can focus on managing the shared storage pool, rather than each individual storage assignment. Storage utilization and consumption trends can be easily monitored with NetApp's management software. NetApp Operations Manager monitors the storage for utilization levels and patterns, and sends alerts and alarms to the administrator when the storage reaches preset usage levels. Operations Manager can also be used to easily configure the storage to auto-grow as needed.

REACT TO CHANGE WITH SPEED, EFFICIENCY, AND AGILITY

It is common for the priorities of different activities and projects to change periodically. With NetApp thin provisioning, system administrators can respond with agility and speed to such changes. By reducing the time needed to respond to storage provisioning requests, thin provisioning helps speed up the launch of critical projects and activities. With thin provisioning, the amount of requested storage for high-priority projects can be allocated on demand. This is accomplished by using storage in the shared storage pool that is not being actively used by other projects. The shared storage pool can be adjusted as needed to accommodate growing storage needs, not simply based on a new storage request. This rapid response is possible because thin provisioning allows the total amount of allocated storage to be larger than the underlying physical storage. It's possible to reduce the delay in launching a new project because it's not necessary to wait until the new storage is purchased.

5 NETAPP THIN PROVISIONING MODEL

DATA ONTAP STORAGE ARCHITECTURE

Data ONTAP® is NetApp's proprietary operating system that runs on all NetApp storage systems. Data ONTAP virtualizes the storage by separating the physical storage pool present in the system from the logical storage pool presented to the connecting hosts.

Figure 5 shows the Data ONTAP storage architecture.

Figure 5) Data ONTAP storage model.

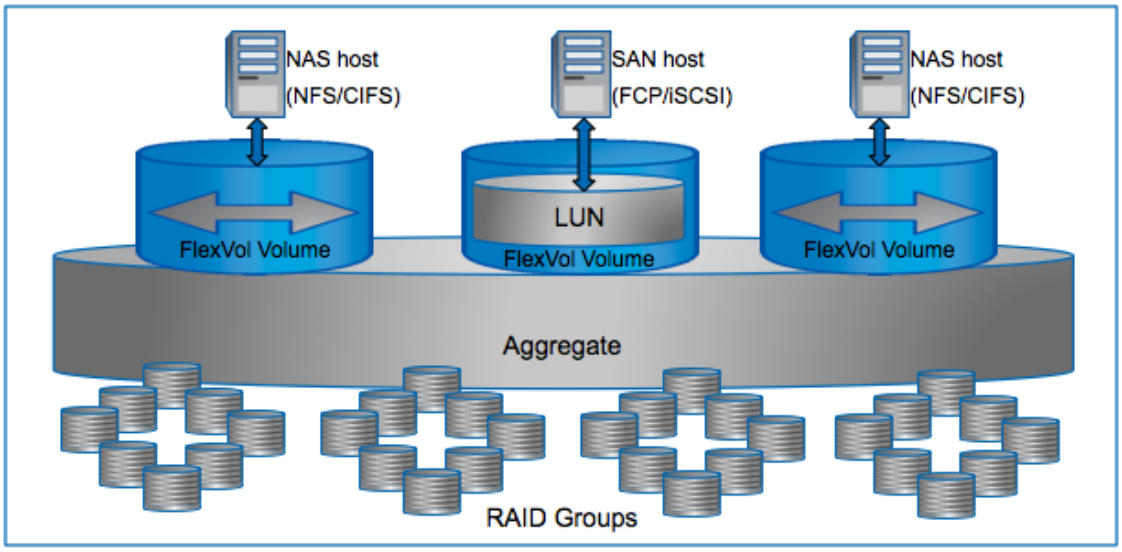

Data ONTAP combines physical disks into RAID groups, which are used to build a shared storage pool called an aggregate. The physical disks used to build the RAID groups consist of the data disks and the parity disks; they can be different disk types, such as SATA, FC, SAS, and SSD.

NetApp FlexVol® volumes are then created in the aggregate. FlexVol volumes are the logical containers that store user data. They are capable of growing and shrinking as needed. Hosts that connect to the storage system via NAS protocols (for example, CIFS and NFS) access and store data directly in FlexVol volumes. Hosts that connect to the storage system via SAN protocols (for example, FC, iSCSI, and FCoE) access LUNs that are created inside FlexVol volumes. The LUNs take advantage of the ability of FlexVol volumes to grow and shrink. Both NAS and SAN can be supported in a single aggregate.

The architecture just described abstracts out the logical data storage containers, which are FlexVol volumes and LUNs, from the underlying physical storage containers, the aggregate. This separation of logical and physical storage makes management easy and also enables NetApp thin provisioning technology.

PROVISIONING MODEL

Data ONTAP implements thin provisioning at the volume and LUN level. NetApp systems support unified protocol access, allowing hosts to access the same storage system using both SAN and NAS protocols.

The granularity of thin provisioning of the storage system depends on the protocol being used and the amount of flexibility required. FlexVol volumes can use thin provisioning in both NAS and SAN environments. LUNs exist as storage objects within the FlexVol volume, and can also use thin provisioning.

Figure 6 shows a few possible configurations when thin provisioning a NetApp system.

Figure 6) Thin provisioning configurations.

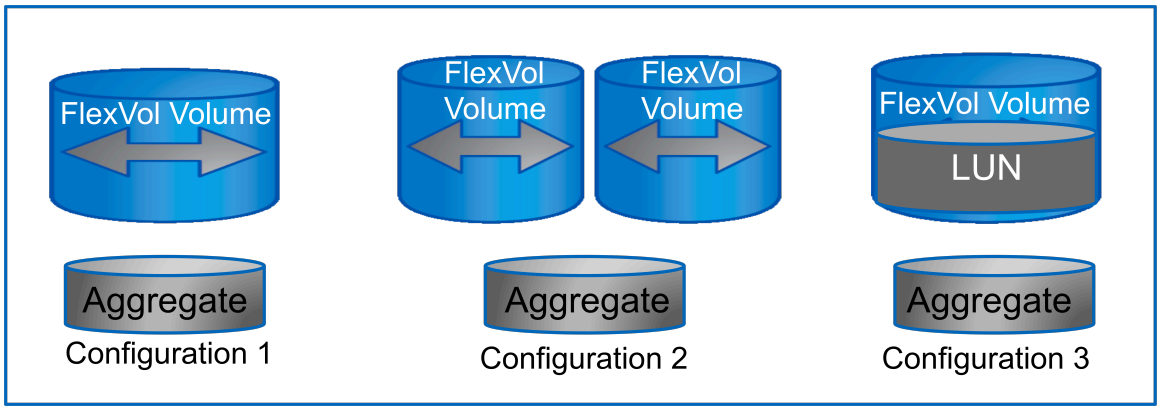

- In configuration 1, thin provisioning is used for the FlexVol volume, and the volume's logical size is bigger than the underlying aggregate. The host sees the logical size of the volume as the size of the volume. The space in the aggregate gets filled as data is written to the volume
- In configuration 2, multiple FlexVol volumes are using thin provisioning and the sum of the FlexVol volumes is bigger than the underlying storage space. Also, each of the volumes themselves have a bigger logical size than the aggregate. In this case, the free space in the underlying aggregate is used as a shared pool of free space by both volumes and they use that space as data gets written to the volumes. For either of these volumes, the host sees storage space equal to the size of the FlexVol volume.
- Configuration 3 shows a SAN environment where the FlexVol volume is using thin provisioning and has a LUN inside the FlexVol volume that is using thin provisioning. The logical space that the host sees in the LUN is much larger than the underlying physical storage space.

Figure 7 shows a basic workflow in a NetApp thin provisioning environment. Note that there is no disruption or any noticeable change to the host, even though storage is being added in the background.

Figure 7) Thin provisioning workflow.

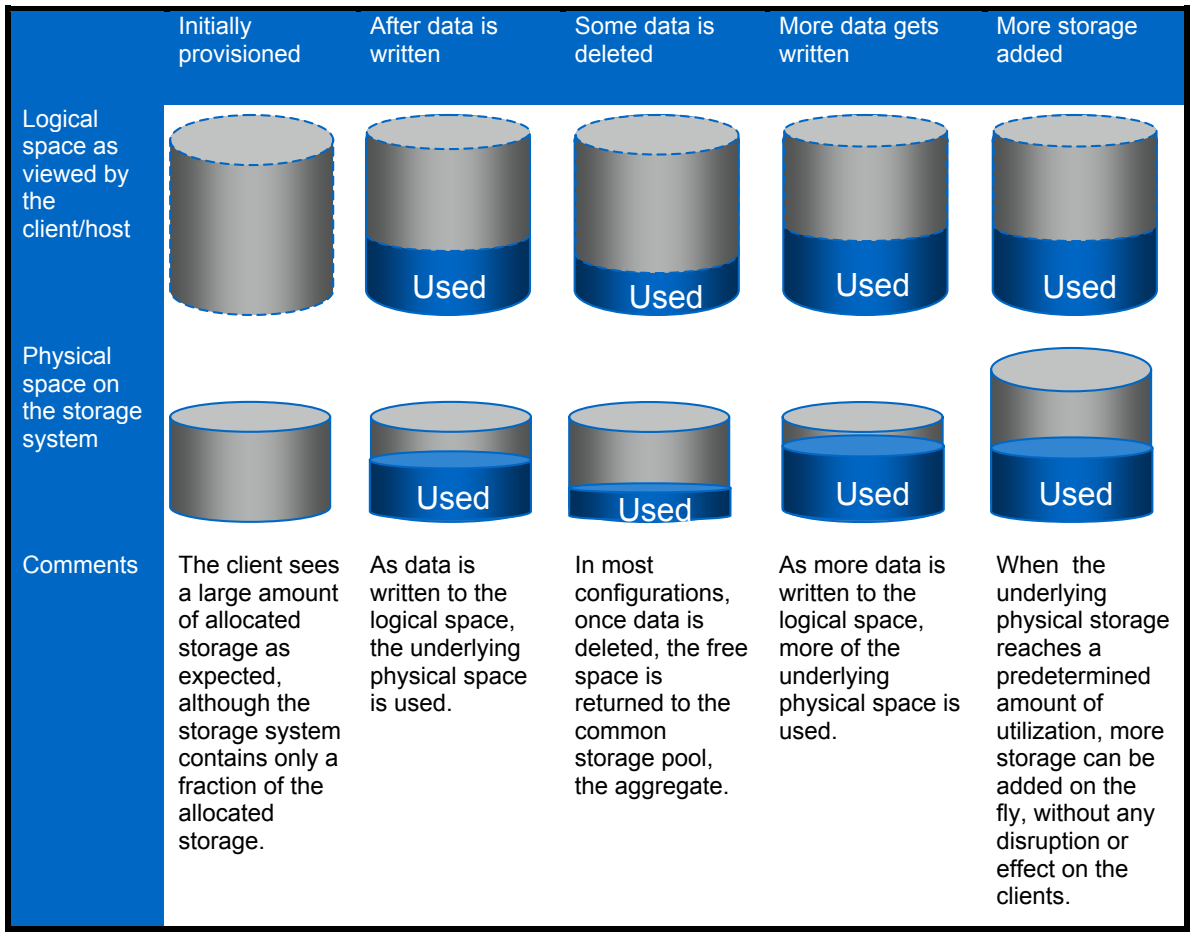

6 NETAPP THIN PROVISIONING IN ACTION

Many vendors tout their thin provisioning capabilities. However, not all thin provisioning vendor implementations are created equal. NetApp offers a powerful, efficient, and easy-to-manage thin provisioning solution.

Thin provisioning is a standard part of NetApp's unified architecture. Data ONTAP 7.0, released in 2004, was the first version of Data ONTAP to support thin provisioning. Since then, NetApp has continued to be a leader in thin provisioning technology by enhancing the feature set, providing tools to efficiently manage thin provisioning, and integrating thin provisioning with newer technologies to further lower the total cost of ownership of the NetApp storage system.

This section describes some of the additional benefits of implementing thin provisioning on a NetApp storage system.

6.1 INTEGRATION WITH NETAPP STORAGE EFFICIENCY TECHNOLOGIES

As a key component of the NetApp storage efficiency technologies, thin provisioning integrates seamlessly with other features to yield maximum space savings without the need for any extra work.

For example, consider a NetApp storage system that is using thin provisioning on a volume containing home directory data. In addition to high storage utilization for that volume, thin provisioning also offers the

option to turn on deduplication for that volume. Suppose that deduplication is turned on for that volume, for a typical savings of 30%. That 30% of freed storage is returned to the aggregate and made available to other projects as part of the shared storage pool. The same approach could be used to take advantage of NetApp data compression for the volume.

Another very popular combination is used for server and desktop virtualization. Consider a volume that is used for storing virtualized systems. NetApp FlexClone volumes are used to rapidly create virtual systems that take up minimal space. Using thin provisioning for the volume allows the volume to grow as needed, while maintaining a high level of storage utilization. Additionally, NetApp deduplication is often used to maintain maximum space savings within the volume.

The seamless integration of NetApp thin provisioning with other NetApp storage efficiency technologies allows different combinations to be used to meet different requirements.

6.2 THIN PROVISIONING EXISTING STORAGE INFRASTRUCTURE

One outstanding benefit of thin provisioning on NetApp systems is the ease of switching from traditionally provisioned storage, making it possible to instantly realize the benefits of thin provisioning on existing NetApp storage. This non-disruptive conversion process, offered in Data ONTAP, involves no data copy or movement. Benefits include the following:

- The conversion process is instantaneous and non-disruptive.
- Converting an existing volume or LUN from thick provisioning to thin provisioning is as easy as changing an option on the FlexVol volumes and LUNs.
- Traditional provisioning and thin provisioning can be used simultaneously, making it possible to turn on thin provisioning selectively for target areas.
- The conversion process can be staged for SAN environments, if necessary, by first thin provisioning the underlying FlexVol volume and later thin provisioning the LUN; or both can use thin provisioning at the same time.
- This process requires zero downtime, and hosts can be actively accessing data in the volume or LUN while the conversion is in process.
- The conversion is totally transparent to the host connected to the storage system. The host continues to see the volume or LUN at the original size.
- There is no difference in performance delivered by the storage system, either during the conversion to thin provisioning or after the conversion. The storage system delivers the same level of performance it was delivering before the conversion.
- No data movement or data copy is involved in the conversion process.
- Returning to a traditionally provisioned storage system is easy and non-disruptive, assuming that there is sufficient storage to be dedicated to the specific projects.

When thin provisioning is turned on for an existing FlexVol volume or LUN, the free space in the volume or LUN is released back into the free space pool of the aggregate and can be used by other volumes or LUNs in the aggregate. For the clients accessing the storage system, there is no noticeable change in the amount of free space, resulting in a totally transparent process.

Figure 8 shows how easy it is to turn on thin provisioning on an existing NetApp system.

Figure 8) Non-disruptive thin provisioning of existing storage.

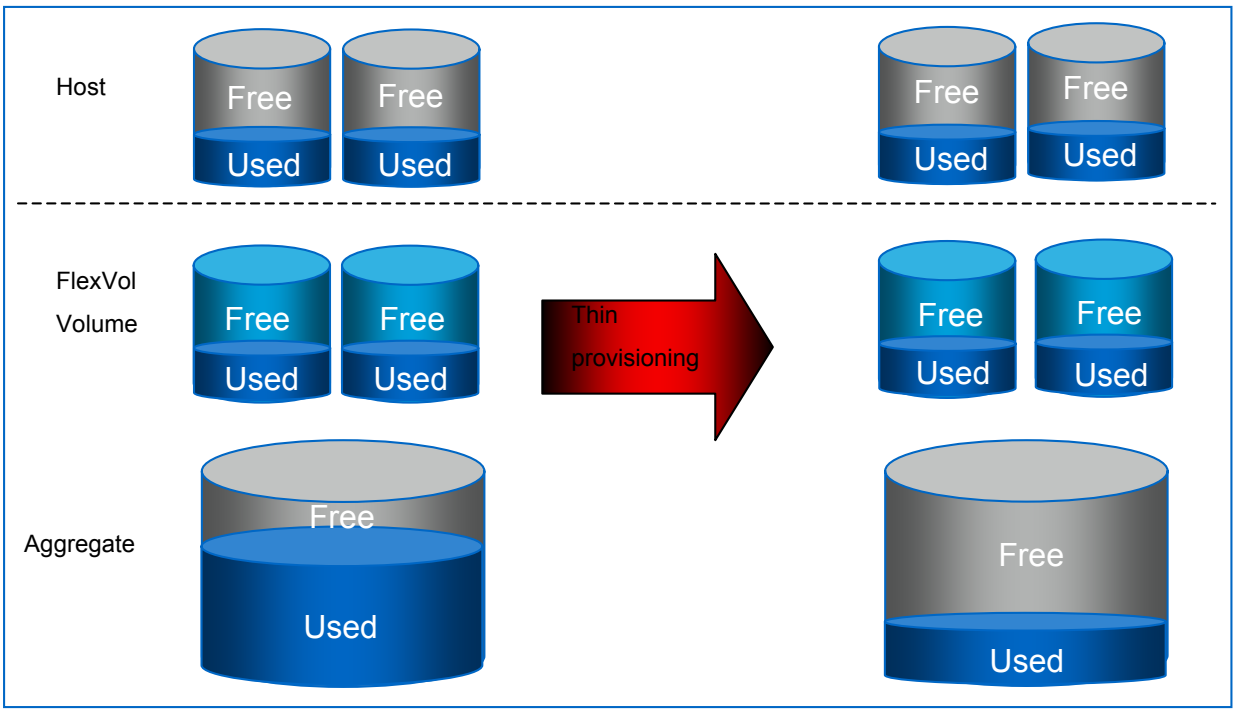

Although the host continues to see the same amount of used and free space, thin provisioning has freed up a significant amount of space in the aggregate, because it is no longer pre-dedicated to each volume or LUN. Instead, all unused storage is returned to the aggregate to be used as a shared storage pool available to all volumes and LUNs for increased flexibility.

6.3 IMPLEMENTATION AND MANAGEMENT

Implementing thin provisioning in either a new or an existing installation is a transparent, simple process with no downtime or disruption. NetApp management tools can be used to manage and monitor thin provisioning. These tools are easy to configure, and they offer a powerful way to manage and monitor the environment.

UNIFIED MULTIPROTOCOL AND MULTIGRANULAR SUPPORT

NetApp Data ONTAP supports unified multiprotocol access to the storage system. Thin provisioning is supported for all supported protocols at different granularities.

NAS protocols support thin provisioning on the FlexVol volume. SAN protocols support thin provisioning at two levels: the containing FlexVol volume level and the LUN level (contained in the FlexVol volume). This offers the flexibility to stage a thin provisioning roll-out based on requirements and planning. Figure 9 illustrates thin provisioning in a unified multiprotocol environment.

Figure 9) Thin provisioning a unified multiprotocol system at different granularities.

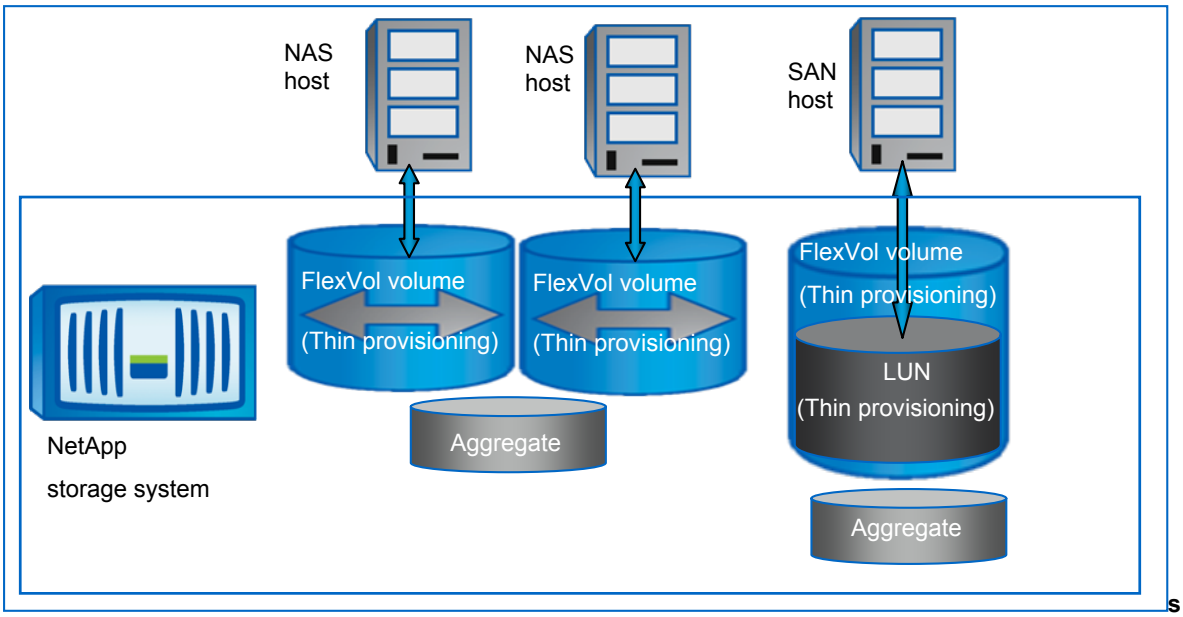

In Figure 9, the volumes to the left both use thin provisioning, allowing them to represent much more storage to the hosts than is physically present in the shared storage pool, the aggregate. The right volume contains a LUN, with thin provisioning applied at both the volume level and the LUN level. This configuration provides the maximum amount of storage utilization in a LUN configuration.

6.4 THIN PROVISIONING WITH DATA PROTECTION AND RETENTION

Data ONTAP uses NetApp Snapshot® technology to retain copies of data. The Snapshot technology enables the creation of point-in-time copies of file systems, which can be used to restore data—from a single file to a complete disaster recovery solution. Snapshot copies in Data ONTAP employ thin provisioning because they store only copies of blocks that were modified from the active file system. Snapshot copies occupy space from the storage system on an on-demand basis as data gets overwritten. This behavior is the same with volume thin provisioning, thus multiplying storage cost savings.

6.5 THIN PROVISIONING AND DATA REPLICATION

Volume SnapMirror® replicates an entire FlexVol volume, most often used for disaster recovery. It is an efficient data replication product that transfers only the data present on the source system to the destination system during the initial baseline transfer. All transfers from that point forward only transfer modified data blocks to the destination. Volume SnapMirror seamlessly integrates with thin provisioning to initially transfer and write to the destination only used data. Changed data is transferred and written later to minimize the amount of data transferred and space consumed on the destination. Thus the savings due to thin provisioning are inherited at the volume SnapMirror destination.

Figure 10 illustrates how the thin provisioning integration with volume SnapMirror provides high utilization at both the source and the destination.

Figure 10) Thin provisioning with data replication.

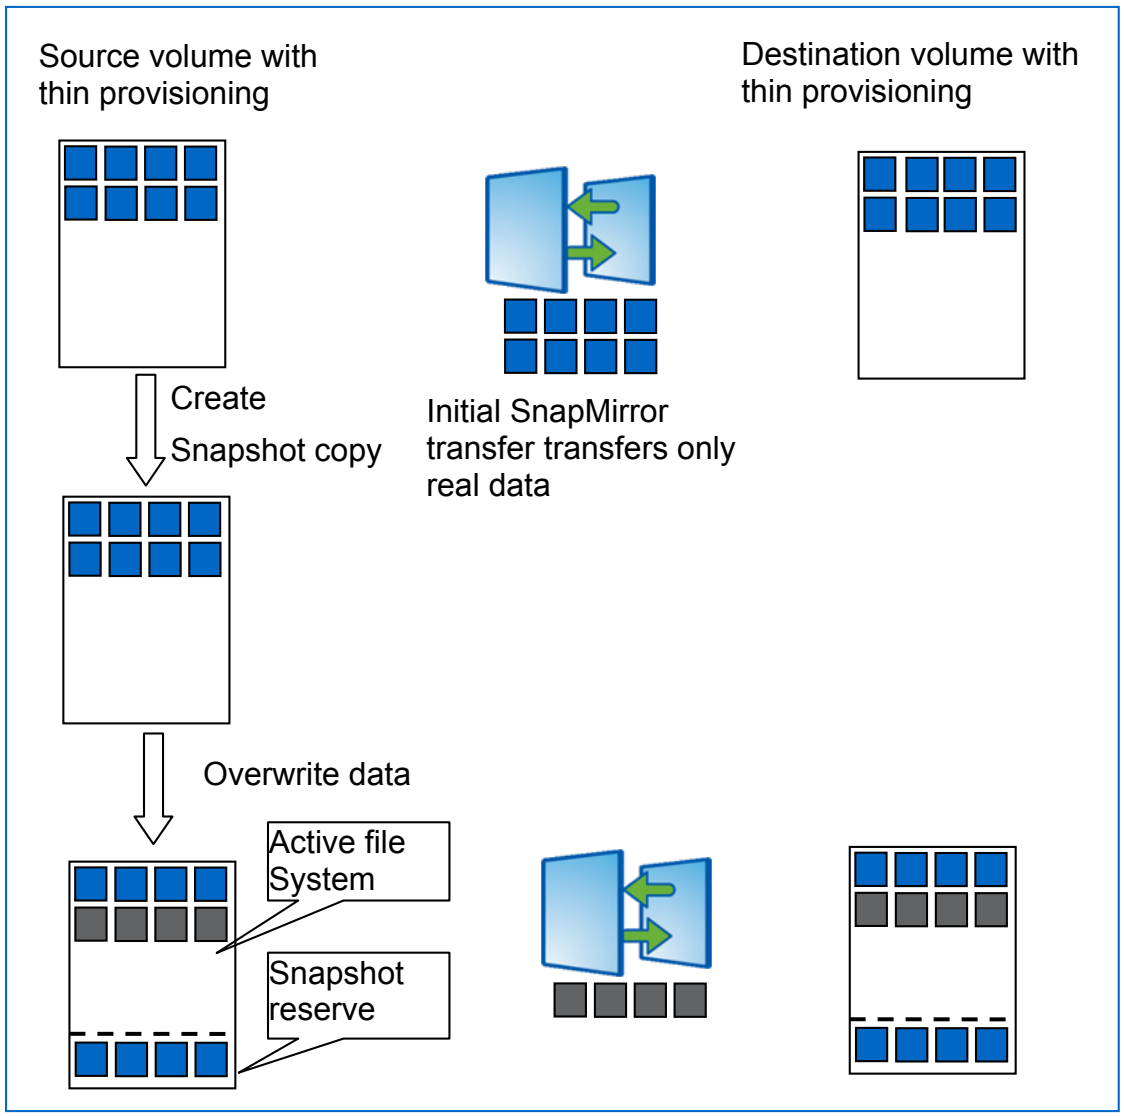

6.6 TIGHTLY INTEGRATED NATIVE FUNCTIONALITY

The ability to use thin provisioning on a FlexVol volume is a native functionality that is built into the Data ONTAP operating system. This has the following significant advantages over vendor implementations that use an add-on approach:

- The write allocation granularity for thin provisioning is 4KB blocks. This is the same as in traditionally provisioned environments. Unlike most other vendors, with NetApp there is no up-front page allocation or any other sort of space reservation for writes.
- There is no performance impact introduced when using thin provisioning. A system using thin provisioning delivers the same amount of performance as a similarly configured system using thick provisioning
- It is an instantaneous, seamless conversion from thick provisioning to thin provisioning on an existing storage system.
- In NAS environments, when data in a volume using thin provisioning is deleted, the freed blocks are returned to the shared free space pool of the aggregate. Therefore a FlexVol volume occupies only a

space equal to the real amount of data in it and returns free space to the aggregate as data in the volume is deleted.

6.7 SEAMLESS APPLICATION AND OPERATIONAL PROCESS INTEGRATION

Thin provisioning a NetApp storage system in a new installation or converting an existing system to thin provisioning is a simple and non-disruptive process. No special processes or tools are required to manage a NetApp storage system that uses thin provisioning. All the NetApp storage system management tools, host utilities like SnapDrive®, and application integration tools like the SnapManager® suite of products understand and support thin provisioning.

This integration makes it easy to roll out thin provisioning with no changes to the existing operational processes. The same tools are used to manage the storage systems. Thin provisioning also requires no changes to the existing storage requisition processes, because end users can still request as much storage as they need, including any buffer and future growth factored in. The storage administrator can still give users a logical container that satisfies their request, while the underlying physical storage pool is much smaller.

7 CONFIGURING NETAPP THIN PROVISIONING

Enabling thin provisioning basically means configuring the NetApp storage system to discontinue reserving storage for volumes or LUNs. That is, storage is not dedicated to the volume or LUN until data is actually written to it. Using this approach, It is possible to create multiple volumes or LUNs within a shared space (aggregate) and allow them to auto-grow as data is written. When additional space is required within the aggregate, additional disks can be added non-disruptively. Data ONTAP enables thin provisioning via options that are set on the volume or the LUN. To use volume thin provisioning, change the Guarantee option on the volume to None. To use LUN thin provisioning, set the No Reserve option to Yes. Using the NetApp management tools or the CLI, this is an instantaneous and non-disruptive process.

7.1 OVERPROVISIONING

As a result of being able to create logical containers (volumes or LUNs) whose combined size is larger than the aggregate, it is a common practice to offer more logical space to hosts than is available in the underlying physical storage. This results in an aggregate that is overprovisioned. Most Internet e-mail services employ this approach when they offer tens of gigabytes of space to all users. The e-mail storage capacity is configured to grow as the space is actually used.

The amount of overprovisioning used depends on the data growth characteristics of the volumes or LUNs within the aggregate. A pattern of multiple projects that do not grow as predicted typically means that less overprovisioning will be used. For example, for a moderate to strong amount of data growth predictability, 50%-100% overprovisioning is typical. Alternatively, for unknown data growth, or if data is expected to rapidly grow to fill the current volumes or LUNs, 0% overprovisioning is typical.

7.2 NEW INSTALLATION

To create a new volume that uses thin provisioning, set the Guarantee option to None while creating the volume from the CLI. When using the NetApp management tool, either directly select the space reservation on the volume, or select a provisioning profile that was configured to use thin provisioning, allowing the tool to auto-provision the storage system.

NetApp Operations Manager and Provisioning Manager tools can be used to create and manage storage using thin provisioning. Figure 11 shows NetApp Operations Manager being used to designate the policy that uses thin provisioning to create new a NAS environment. This policy approach greatly simplifies the implementation of thin provisioning.

Figure 11) Using the NAS thin provisioning policy.

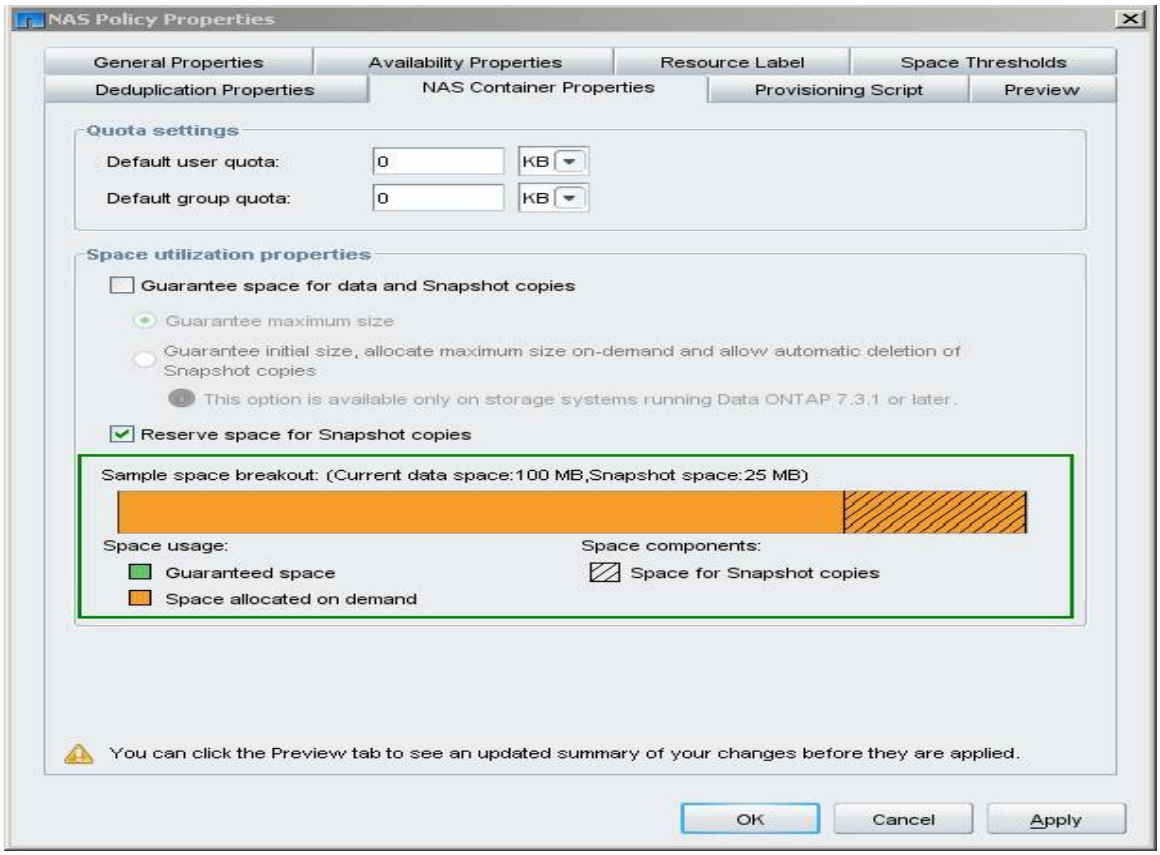

Figure 12 shows NetApp Operations Manager being used to specify the policy that uses thin provisioning for a new SAN deployment.

Figure 12) Using the SAN thin provisioning policy.

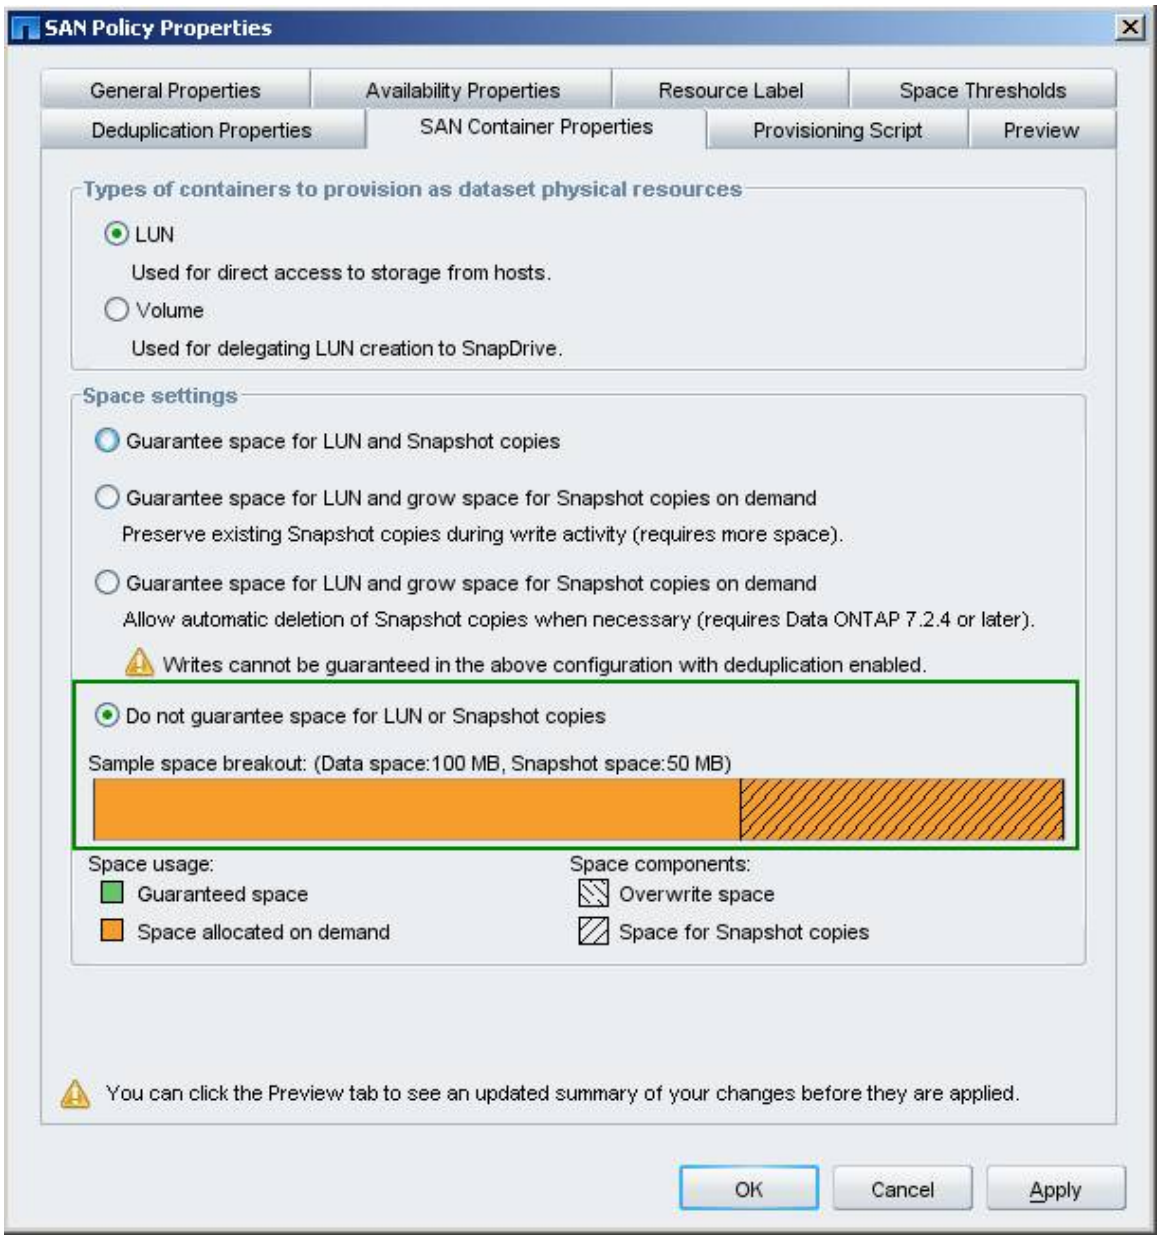

7.3 EXISTING INSTALLATION

To switch an existing volume from thick provisioning to thin provisioning, simply use either the NetApp management tools or the CLI to set the Guarantee option to None. This releases any unused reserved space from the volume to the aggregate.

8 MANAGING NETAPP THIN PROVISIONING

A critical requirement for thin provisioning is being able to efficiently monitor space utilization. Because thin provisioning often offers more logical space to their hosts than is available in the underlying physical storage, it is very important to monitor the underlying storage. The tools should be capable of generating appropriate messages and warnings based on utilization so that the storage administrator can react in a

timely manner to avoid running out of space. NetApp Operations Manager provides this capability and more.

Operations Manager allows the storage administrator to set alerts for different levels of utilization, making it possible to quickly respond to changes in storage consumption. Operations Manager can also be configured to alert with either warning or error messages when the aggregate is overprovisioned by a certain percentage, and also when the aggregate gets to a certain level of fullness. Figure 13 shows the configuration options.

Figure 13) Configuring alerts for aggregate overcommitment and aggregate fullness.

In addition to informing the storage administrator of what has already occurred, Operations Manager can also predict storage consumption by analyzing storage utilization trends. It reports the number of days it will take to fill an aggregate to 100%, if current consumption trends continue. This information can be used to take a proactive approach to thin provisioning and to plan storage purchasing based on the behavior of the environment, instead of the guessing approach used for traditional provisioning. Figure 14 shows how Operations Manager displays the number of days until space in an aggregate is completely consumed.

Figure 14) Aggregate fullness estimation analysis.

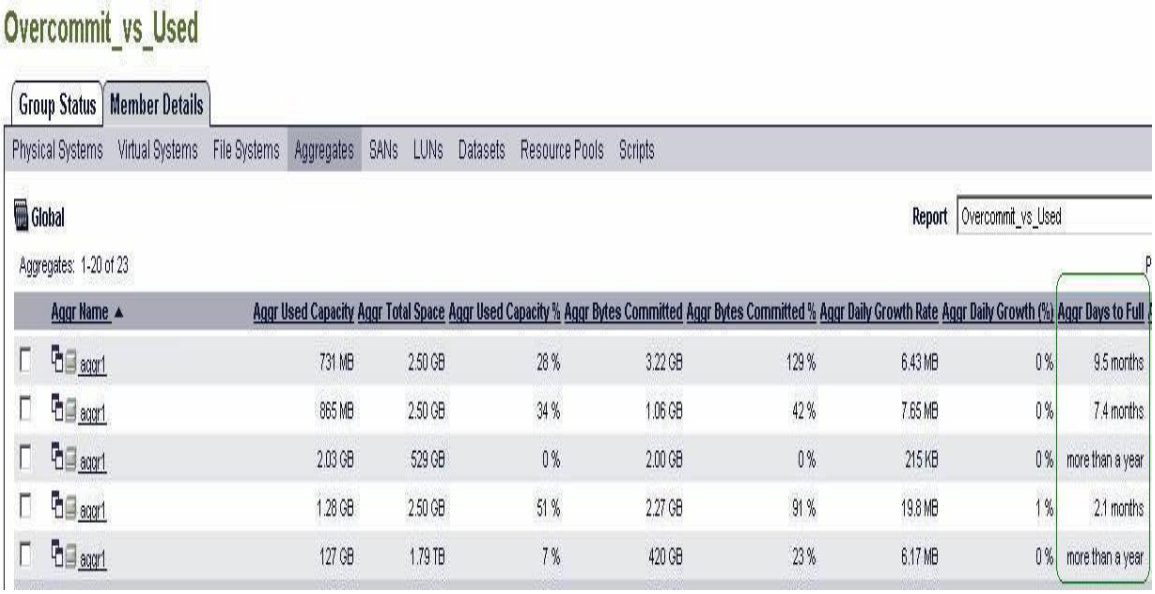

9 CONCLUSION

NetApp thin provisioning helps companies to realize huge financial savings in terms of both capex and opex, and it can also improve operational efficiencies. It offers storage administrators an easy way to enable thin provisioning in existing environments, the ability to enable thin provisioning on targeted storage, and the NetApp management tools to deploy and manage thin provisioning. Add to this the fact that thin provisioning is included at no cost in the NetApp storage system, and it makes sense to give it a try.

10 ADDITIONAL READING AND REFERENCES

- TR-3440: Operations Manager, Protection Manager, and Provisioning Manager Sizing Guide
  http://media.netapp.com/documents/tr-3440.pdf
- TR-3505: NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide.
  http://media.netapp.com/documents/tr-3505.pdf
- TR-3446: SnapMirror Async Best Practices Guide
  http://media.netapp.com/documents/tr-3446.pdf
- TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices
  http://media.netapp.com/documents/tr-3428.pdf
- TR-3824, Storage Efficiency and Best Practices for Microsoft Exchange Server 2010
  http://media.netapp.com/documents/tr-3824.pdf
- TR-3461, V-Series Best Practice Guide
  http://media.netapp.com/documents/tr-3461.pdf

11 ADDITIONAL ASSISTANCE

For additional support, contact one of the following:

- Your local account team
- A NetApp systems engineer
- NetApp Global Services
- NetApp Support site (was NOW™)
- 888.4.NETAPP (888.4.638277); United States and Canada
- 00.800.44.NETAPP (00.800.44.638277); EMEA
- +800.800.80.800; Asia/Pacific

12 VERSION TRACKING

Version 1	May 2007	Initial release
Version 2	May 2011	Updated with latest messaging, images, and terminology