



Technical Report

NetApp Thin Provisioning Deployment and Implementation Guide Data ONTAP 8.1 7-Mode

Carlos Alvarez, NetApp
September 2011 | TR-3965 | Rev. 1

ABSTRACT

This guide introduces NetApp[®] thin provisioning, describes in detail how to implement and use it, and provides information on best practices, operational considerations, and troubleshooting. Although the title clearly states this document applies to Data ONTAP 8.1, the general concepts will apply to Data ONTAP 7.3.x and 8.0.x. The main differences will occur in the feature availability of DataMotion. For details, see the included document references.

The guide should prove useful to NetApp customers and implementers who require assistance in understanding details and successfully deploying solutions that include thin provisioning.

TABLE OF CONTENTS

1	INTRODUCTION AND OVERVIEW OF THIN PROVISIONING	5
1.1	THIN PROVISIONING DEFINED	5
1.2	THIN PROVISIONING ARCHITECTURE	5
1.3	THIN PROVISIONING VOLUMES	7
1.4	THIN PROVISIONING LUNS	9
1.5	THIN PROVISIONING GENERAL CHARACTERISTICS	12
2	CONFIGURATION AND OPERATIONS	13
2.1	REQUIREMENTS	13
2.2	KEY DATA ONTAP PARAMETERS TO CONSIDER	14
2.3	INSTALLING AND LICENSING THIN PROVISIONING	20
2.4	COMMAND SUMMARY	20
2.5	THIN PROVISIONING QUICK START	22
2.6	THIN PROVISIONING POLICIES IN PROVISIONING MANAGER	23
2.7	MONITORING THIN PROVISIONING WITH OPERATIONS MANAGER	25
2.8	NOTIFICATION	29
2.9	PRESERVING STORAGE AVAILABILITY (MITIGATION ACTIVITIES)	31
3	BEST PRACTICES	34
3.1	DEFAULT (THICK) PROVISIONING	34
3.2	THIN PROVISIONING COMMON PRACTICES AND BEST PRACTICES	35
3.3	RECOMMENDED THIN PROVISIONING CONFIGURATIONS	36
3.4	THIN PROVISIONING IN PRACTICE	38
4	THIN PROVISIONING AND VMWARE	41
5	THIN PROVISIONING AND SHAREPOINT	41
5.1	USE-CASE SPECIFICS	41
5.2	BEST PRACTICES	42
6	THIN PROVISIONING AND SQL	45
6.1	USE-CASE SPECIFICS	45
6.2	BEST PRACTICES	45
7	THIN PROVISIONING AND EXCHANGE	48
7.1	USE-CASE SPECIFICS	48
7.2	BEST PRACTICES	50
8	THIN PROVISIONING AND ORACLE	53
8.1	USE-CASE SPECIFICS	53
8.2	BEST PRACTICES	53

9 THIN PROVISIONING WITH OTHER NETAPP FEATURES.....	56
9.1 THIN PROVISIONING AND MANAGEMENT/MONITORING TOOLS.....	56
9.2 THIN PROVISIONING AND DEDUPLICATION	58
9.3 THIN PROVISIONING AND FLEXCLONE VOLUMES	58
9.4 THIN PROVISIONING AND V-SERIES.....	59
9.5 THIN PROVISIONING AND SNAPMIRROR REPLICATION	59
9.6 THIN PROVISIONING AND SNAPVAULT.....	60
9.7 THIN PROVISIONING AND MULTISTORE (VFILER)	60
9.8 THIN PROVISIONING AND DATAMOTION FOR VOLUMES.....	60
10 TROUBLESHOOTING	60
10.1 NOT SEEING STORAGE SAVINGS.....	60
10.2 NOT GETTING ALERTS.....	60
10.3 PERFORMANCE SLOWDOWN.....	60
11 ADDITIONAL READING AND REFERENCES	60
12 ADDITIONAL ASSISTANCE.....	62
13 VERSION TRACKING.....	63

LIST OF TABLES

Table 1) Basic volume thin provisioning parameters.....	8
Table 2) Basic LUN thin provisioning parameters.	10
Table 3) Overview of thin provisioning technical requirements.	14
Table 4) Overview of thin provisioning environmental requirements.....	15
Table 5) Basic volume and LUN provisioning parameters.	21
Table 6) Provisioning commands and parameters.....	22
Table 7) Thin provisioning quick start.	23
Table 8) Mitigation activities for preserving space availability within aggregates.	34
Table 9) Mitigation activities for preserving space availability within volumes.....	35
Table 10) Default settings.	36
Table 11) Comparison of provisioning methods and requirements.	38
Table 12) Characteristics of the recommended thin provisioning configurations.	39
Table 13) Data ONTAP recommended volume and LUN thin provisioning parameters for SharePoint.....	43
Table 14) Data ONTAP recommended volume and LUN thin provisioning parameters for SQL.....	46

Table 15) Data ONTAP recommended volume and LUN thin provisioning parameters for Exchange.....	51
Table 16) Data ONTAP recommended volume and LUN thin provisioning parameters for Oracle.	54
Table 17) Provisioning functionality of NetApp tools.	57
Table 18) NetApp tools' effect on key settings.	58

LIST OF FIGURES

Figure 1) NetApp Data ONTAP storage architecture.	6
Figure 2) Volume thin provisioning.....	8
Figure 3) Multivolume thin provisioning, overcommitted.	8
Figure 4) LUN thin provisioning.....	9
Figure 5) Multi-LUN thin provisioning, overcommitted.	10
Figure 6) LUN filling with new files.	11
Figure 7) LUN filling with new files.	12
Figure 8) LUN overwrite reserve examples.....	17
Figure 9) LUN and volume thin provisioning space availability.	19
Figure 10) Volume provisioning with Provisioning Manager.....	24
Figure 11) LUN provisioning with Provisioning Manager LUN provisioning.	25
Figure 12) Operations Manager screen (partial) to configure thresholds on operational metrics.	26
Figure 13) Trending of aggregate data growth and days-to-full prediction in Operations Manager.....	28
Figure 14) Operations Manager storage efficiency dashboard.	29
Figure 15) Configuring an alarm based on the threshold, aggregate almost full.....	30
Figure 16) Example condition of aggregate tightness and action now needed.	40
Figure 17) Example mitigation of migrating a migration-candidate volum/LUN to a different aggregate.....	40
Figure 18) Compare thin and thick provisioning for Exchange 2010 with 100 mailboxes.	49
Figure 19) Storage thin provisioning for Exchange 2010 DAG.....	50

1 INTRODUCTION AND OVERVIEW OF THIN PROVISIONING

Although thin provisioning is a commonly used storage term, implementations will vary depending on the specific vendor, product, and even application. This section provides an overview of the NetApp thin implementation and its benefits. It includes key points, such as the ease of enabling thin provisioning, and how NetApp's thin provisioning architecture allows performance levels to be maintained. Thin provisioning is the most efficient way to provision storage, because the storage is not all preallocated up front. In other words, when a volume or LUN is created using thin provisioning, no space on the storage system is used. The space remains unused until data is written to the LUN or the volume, at which time only enough space to store the data will be used. To take full advantage of these and other benefits that thin provisioning has to offer, a mature storage management capability is necessary. For additional details, including thin provisioning differentiators and business cases, see [TR-3563: NetApp Thin Provisioning Increases Storage Utilization with on-Demand Allocation](#).

The alternative to thin provisioning is thick provisioning, which is the traditional approach in which storage is allocated at the time of volume or LUN creation. In this case, the space is dedicated to a specific LUN or volume, and it cannot be shared with other LUNs or volumes, even if the space is sitting empty and unused.

Although the title of this technical report clearly implies that this document applies to Data ONTAP 8.1, the general concepts will apply to Data ONTAP 7.3.x and 8.0.x. The main differences will occur in the feature availability of DataMotion. For more information about DataMotion, see the document references included within this document.

1.1 THIN PROVISIONING DEFINED

NetApp thin provisioning has two main abilities.

1. The ability to provision more logical storage space to the hosts connecting to the storage system than is actually available on the storage system
2. The ability to allocate storage on demand as the data comes in, instead of preallocating it

Consider an example in which there are 15 hosts, and each host estimates that he/she will need 500GB of storage space for the duration of the project over 3 years, totaling 7500GB. The storage administrator configures a storage system using thin provisioning in which the storage system only has 5000GB of usable storage. The resulting effect is that each host sees the capacity he/she requested and the storage system has not used any space for the initial creation of that storage. In contrast to thin provisioning, a thick-provisioned storage system would attempt to preallocate the storage required for each host and not be capable of servicing the requests for all projects. Thick provisioning would require the full 7500GB of storage to be available up front. Instead, with thin provisioning, 5000GB of storage remain available to all projects until a project writes data to the storage, at which time that specific storage is allocated to that project and no longer available to other projects. To avoid running out of space, the administrator monitors the storage system and adds storage as needed.

Thin provisioning is commonly used in many industries. Consider a water company. A water utility company's infrastructure does not have the capacity to maintain acceptable water pressure if all of its customers turn on their taps simultaneously. The company plans on having only a certain percentage of its customers simultaneously using water. This example illustrates the basic concept of thin provisioning.

1.2 THIN PROVISIONING ARCHITECTURE

NetApp FlexVol[®] technology provides the ability to create efficient flexible data containers that separate the content from the constraints of the physical storage. The FlexVol technology decouples the physical barrier between data containers and their associated physical disks: The result is a significant increase in flexibility and storage utilization. With the Data ONTAP[®] FlexVol technology, you can shrink or grow data

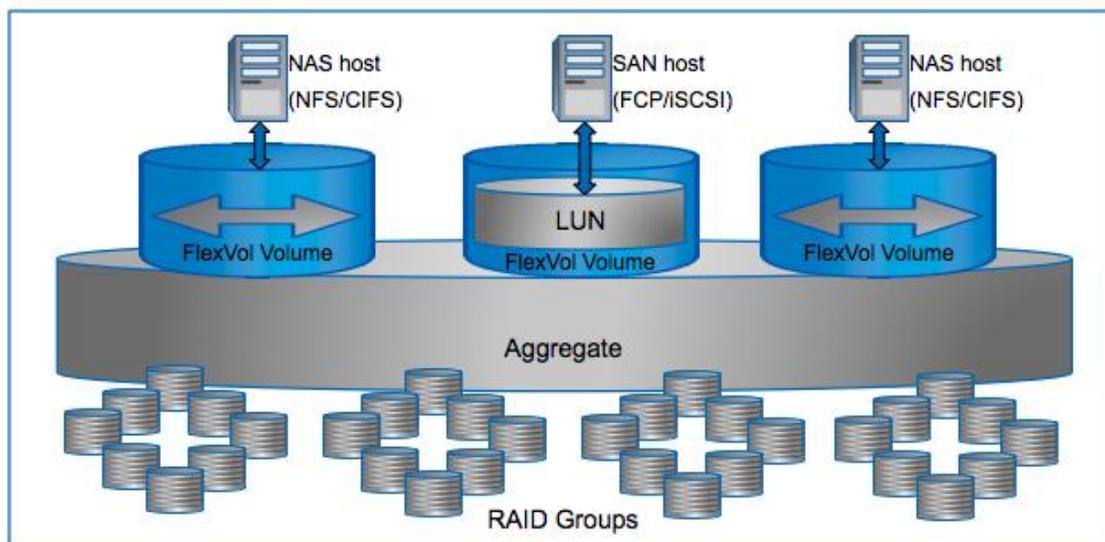
containers based on immediate needs. Adding disks can be done on the fly, without disrupting the system or the associated applications.

KEY COMPONENTS OF THIN PROVISIONING

1. **FlexVol volumes:** FlexVol volumes are a standard free component included with NetApp storage systems. They can be configured to automatically allocate space as data is written to them, and can be configured to work with file-level protocols (NAS).
2. **LUN:** LUNs are a standard free component included with NetApp storage systems. They provide support for block-level protocols (SAN). Within Data ONTAP, LUNs are simply objects within a volume, providing additional levels of flexibility and efficiency.
3. **Aggregate:** 32-bit aggregates and 64-bit aggregates are free components on NetApp storage systems, and are used as the storage pool for volumes and LUNs using thin provisioning. Multiple aggregates can be used simultaneously on a single system, and each can be expanded with additional storage if more space is needed in the storage pool.
4. **Management tools:** NetApp management tools, such as Provisioning Manager and Operations Manager, can be used to easily implement thin provisioning policies, report storage utilization, monitor growth trends, and trigger automatic actions and alerts based on different thresholds.

Figure 1 illustrates how these components can be combined to create a solution based on a single aggregate. In practice, it is common for a single system to contain multiple aggregates and various combinations of volumes and LUNs within the aggregates.

Figure 1) NetApp Data ONTAP storage architecture.



The NetApp Data ONTAP architecture uses aggregates to virtualize the physical storage into pools for logical allocation. The volumes and LUNs see the logical space, and the aggregate controls the physical space. This architecture provides the flexibility to create multiple volumes and LUNs that can exceed the physical space available in the aggregate. All volumes and LUNs in the aggregate will use the available storage within the aggregate as a shared storage pool. This will allow them to efficiently allocate the space available in the aggregate as data is written to it, rather than preallocating (reserving) the space. Preallocating space will often result in space that sits unused for long periods of time, resulting in lower storage utilization rates, also known as wasted storage. In the case that some volumes or LUNs do indeed have a valid reason for preallocating space, you have the flexibility to allow all other volumes and LUNs to carve out preallocated space in the aggregate. This can be achieved while simultaneously supporting thin provisioning within the same aggregate. More clearly stated, the aggregate is the shared

storage pool for volumes and LUNs using thin provisioning, and can also include the volumes and LUNs using preallocated storage, simultaneously. The aggregate can be expanded by simply adding more disks to it. Monitoring of the storage utilization takes place at the aggregate level to predict when the aggregate should be expanded.

Figure 1, above, shows that the aggregate’s physical space is made up of disks (SATA, FC, SAS, or SSD) and RAID type. The volumes and LUNs are created within the aggregate, achieving a level of abstraction between the logical volumes and LUNs and the physical storage. The separation of the logical storage container and the physical storage provides the flexibility to grow the aggregate as needed, independent of growing and shrinking the volumes and LUNs. This is all achieved transparent to the clients. The client hosts that are accessing the volumes and LUNs can see the logical size of those volumes and LUNs, but do not have direct visibility of the physical space size.

The next two sections provide more detail for applying this concept to NAS (volume) and SAN (LUN) environments.

1.3 THIN PROVISIONING VOLUMES

The Write Anywhere File Layout (WAFL[®]) file system provides the storage virtualization layer for Data ONTAP, enabling thin provisioning of the volume. To create a volume using thin provisioning, set the volume space guarantee option to “none.” You can also change the volume guarantee value to “none” for existing volumes without affecting performance. When a volume is created, it does not dedicate space in the aggregate for the volume, known as preallocation. Instead, the blocks are allocated from the aggregate when data is written to the volume. This allows provisioning more storage space than is actually physically present in the aggregate.

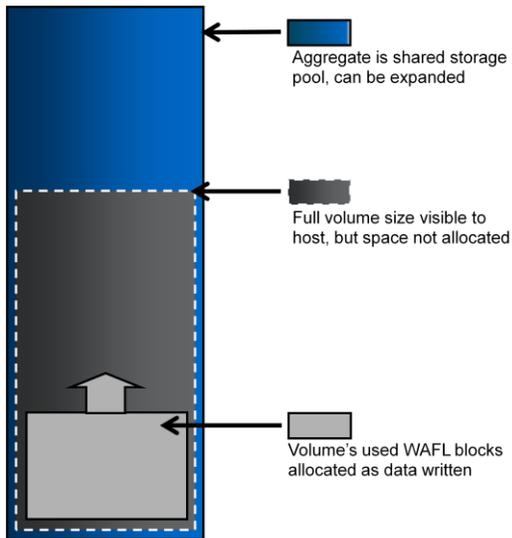
Table 1, below, shows the minimum parameters that must be considered when using volume thin provisioning. Setting a volume’s guarantee to “none” results in none of the space in the parent aggregate being reserved for that particular volume.

Table 1) Basic volume thin provisioning parameters.

Parameter	Value
Volume Options	
guarantee	none

Figure 2 represents volume thin provisioning in its simplest form. The only storage space that is allocated in this example is for data that has already been written to the volume. The unallocated space is available for use by this volume, or any other volume that is created within this aggregate.

Figure 2) Volume thin provisioning.



If more volumes are added into the aggregate, only used space will be consumed. This presents the option of overcommitting the aggregate. Overcommitment of the aggregate means that more space is presented to the hosts than is physically available within the aggregate.

Figure 3) Multivolume thin provisioning, overcommitted.

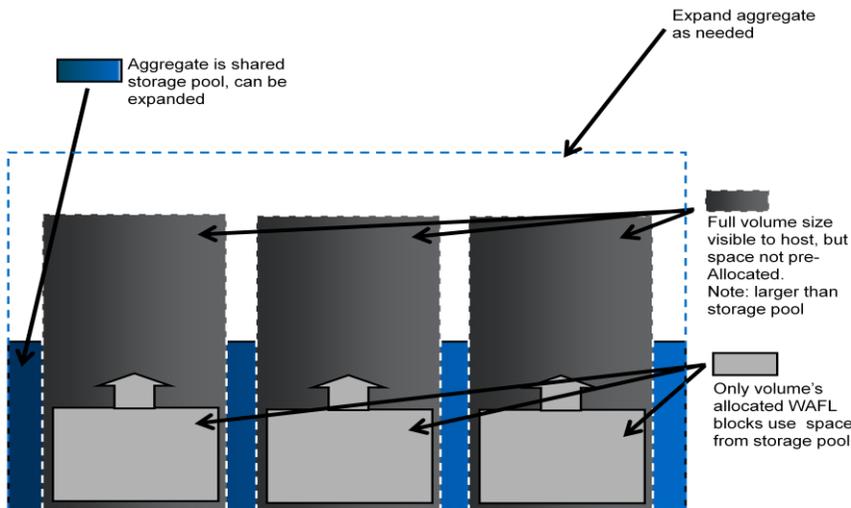


Figure 3 shows thin provisioning of multiple volumes within a single aggregate, where the total amount of storage provisioned is greater than the amount of physical storage space available in the aggregate. Because the storage is not allocated until the data is written to the volume, all unallocated storage can be used by any volume (or LUN) within the aggregate. As data is written to the volumes, the total physical storage space within the aggregate is allocated and the utilization level of the aggregate continues to increase. As the aggregate approaches being nearly full, a predetermined process, such as expanding the aggregate, should be activated to create more free storage space within the aggregate. Additional policies and parameters are considered within this document following the LUN thin provisioning overview.

1.4 THIN PROVISIONING LUNS

The NetApp unified architecture provides the ability to support both volumes and LUNs from the same storage system. When a LUN is created, it is simply an object within a FlexVol volume. This approach provides many benefits, including the ability to use thin provisioning on LUNs, just as can be done for volumes. The Write Anywhere File Layout (WAFL) file system provides the virtualization layer for Data ONTAP, enabling thin provisioning of the LUN. To create a LUN using thin provisioning, create a volume with no volume guarantee, and then create a LUN with LUN reservation turned off. You can also change the volume guarantee value and LUN reservation value for existing volumes and LUNs without affecting performance. When the LUN is created, it does not dedicate specific blocks in the aggregate for the volume or LUN, known as preallocation. Instead, the blocks are allocated from the aggregate when data is written to the LUN. This allows provisioning more storage space than is actually physically present in the aggregate.

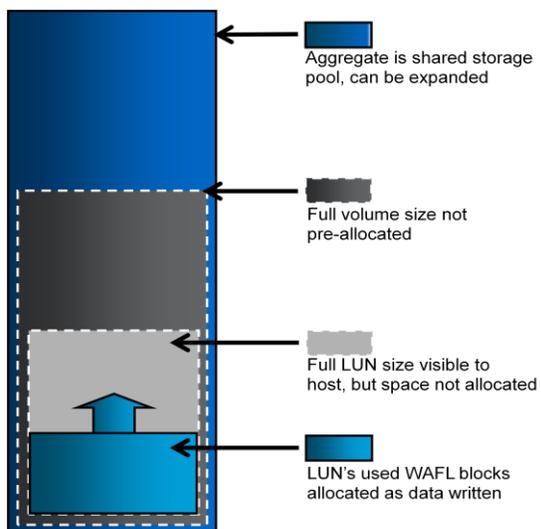
Table 2, below, shows the minimum parameters that must be considered when using LUN thin provisioning within a volume that is also using thin provisioning. Setting a volume's guarantee to "none" results in none of the space in the parent aggregate being reserved for that particular volume. Setting the LUN's space reservation to "disabled" results in none of the space within the parent volume being reserved for that particular LUN.

Table 2) Basic LUN thin provisioning parameters.

Parameter	Value
Volume Options	
guarantee	none
LUN Options	
Space reservation	disabled

Figure 4 represents LUN thin provisioning in its simplest form. The only storage space that is allocated in this example is for data that has already been written to the LUN. The unallocated space is available for use by this LUN, or any other LUN (or volume) that is created within this aggregate.

Figure 4) LUN thin provisioning.



Next, let's take a look at what happens if we add more volumes and LUNs into the aggregate, creating a level of overcommitment.

Figure 5) Multi-LUN thin provisioning, overcommitted.

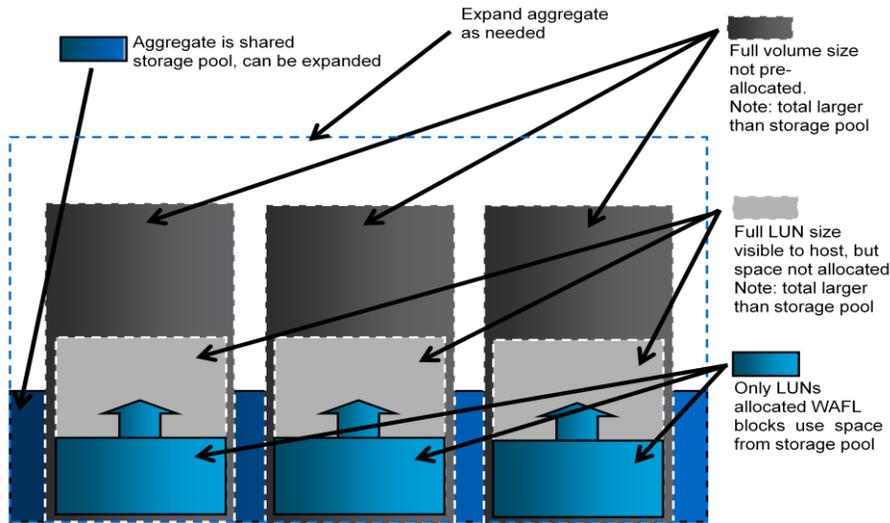


Figure 5 shows thin provisioning of multiple volumes and LUNs within a single aggregate, where the total amount of storage provisioned is greater than the amount of physical storage space available in the aggregate. Because the storage is not allocated until the data is written to the LUN, all unallocated storage is available to any LUN (or volume) within the aggregate. As data is written to the LUNs, the total physical storage space within the aggregate is allocated and the utilization level of the aggregate continues to increase. As the aggregate approaches being nearly full, a predetermined process, such as expanding the aggregate, should be activated to create more free storage space within the aggregate. Additional policies and parameters are considered within this document in following sections.

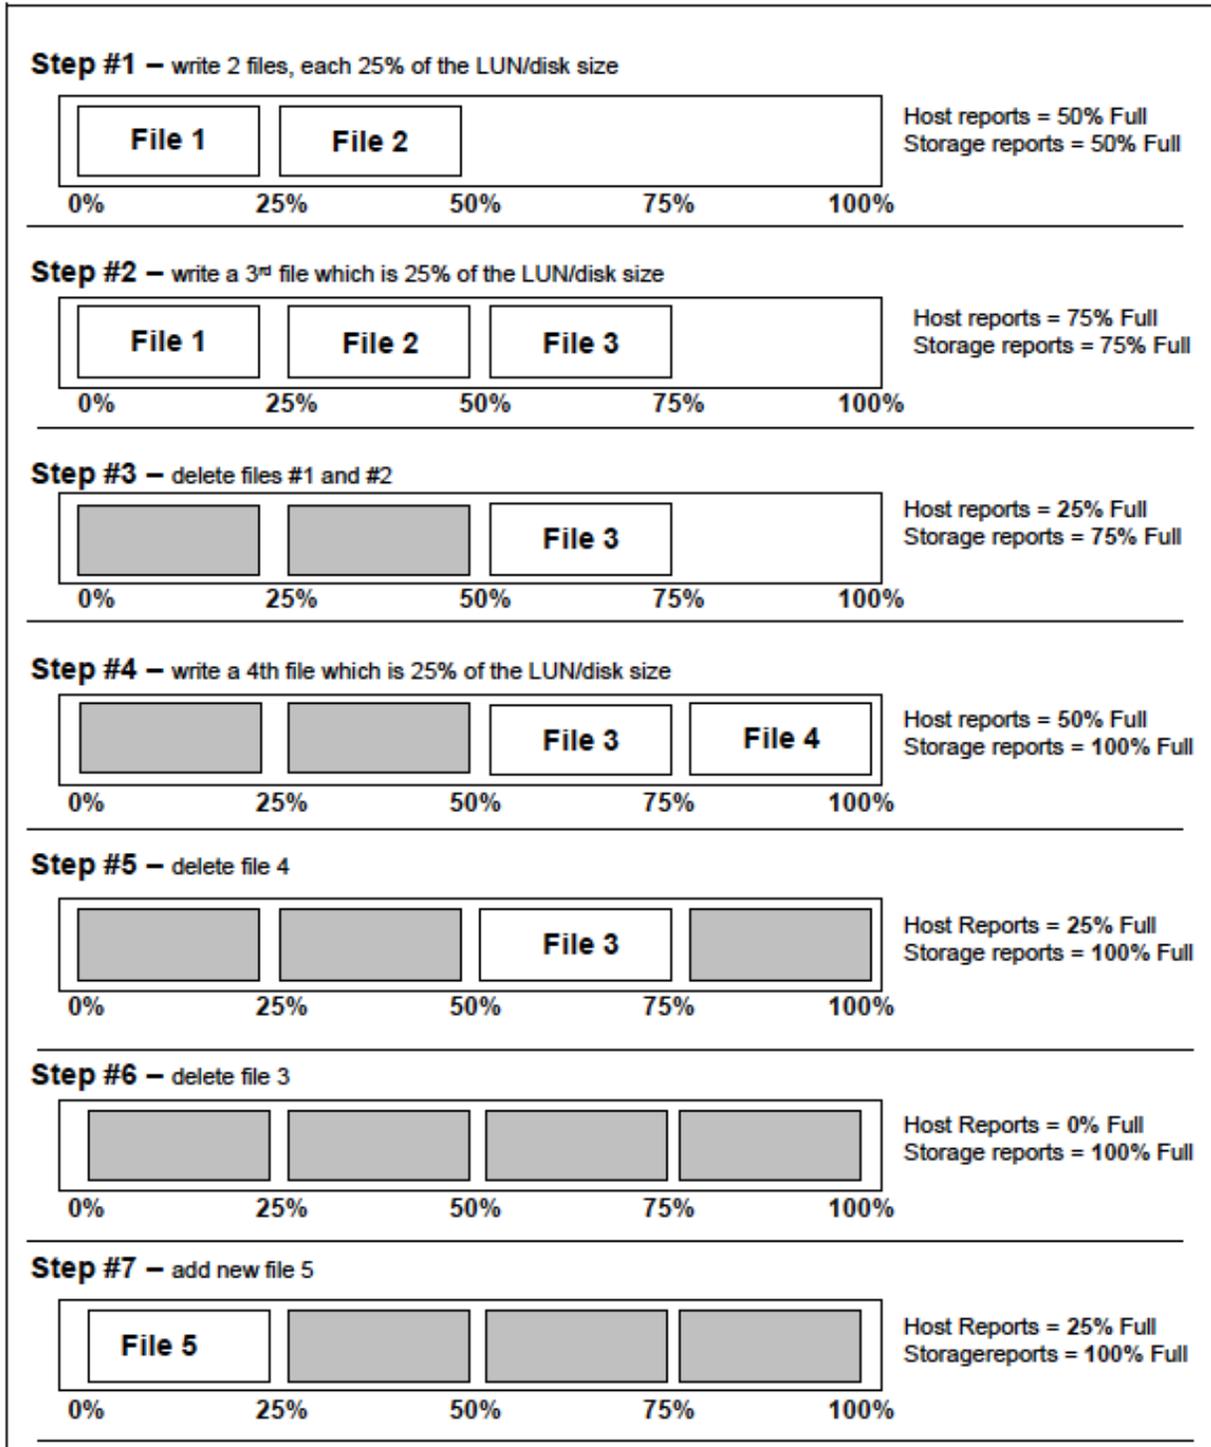
SPACE RECLAMATION

There is often confusion over the space usage inside a LUN. This confusion stems from the fact that the vast majority of LUNs have a file system installed on them, resulting in a possible discrepancy between how full the file system is and how much space has been used within the LUN. This scenario is not unique to NetApp. Assuming the file system is configured to have access to the complete LUN, the LUN will always be at least as full as the file system. From a storage system perspective, the percentage of a LUN used will always increase and never decrease unless space reclamation is used, which is discussed later.

Figure 6, below, shows an individual LUN's usage, and why space reclamation is needed. It depicts a theoretical case in which a file system and a LUN used as files are written and deleted from the file system. Steps 1 and 2 show, as expected, that the file system and the LUN both report the same percentage used as files are written to the file system. Step 3 is where a difference occurs between what the host, or file system, reports and what the storage reports. In this step, files 1 and 2 are deleted, meaning the blocks used by these files are available to the file system. The only file still visible in the file system is file 3, so the file system reports it is 25% full. However, the storage is still reporting that 75% of the LUN is used. The reason that the storage shows 75% full is because the storage has no way of knowing that the data written for files 1 and 2 is no longer needed. The space in the graphic is shaded because, from the storage perspective, those blocks still contain valid data. When a file is deleted the file system does not send any type of erase or delete command to the storage array. It simply changes a pointer in its tables that pointed to the file and makes those blocks available for future use. The fact that the blocks containing the files are not actually erased is what allows "undelete" utilities to recover files after the user has deleted them. In Step 4, file 4 is shown being written to the last 25% of the LUN. This

step depicts how a file system doesn't necessarily attempt to keep data at the beginning of a LUN. After Step 4, every block in the LUN has been written and is therefore full from the storage system's perspective. As mentioned, the fill percentage of a LUN will never go down without space reclamation, and Steps 5 through 7 show that, regardless of the host activity on the LUN, it will remain 100% full.

Figure 6) LUN filling with new files.

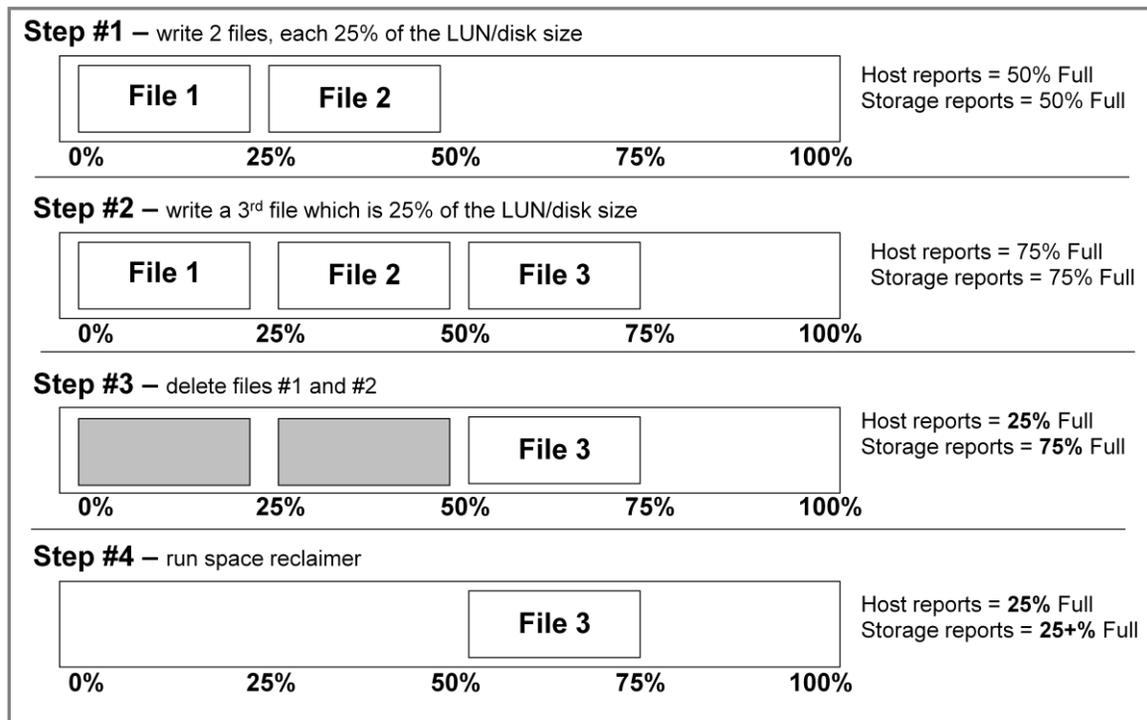


LUN space reclamation is the process of allowing the storage system to free the blocks no longer used by the host operating system. These blocks can then be used by the storage system for other purposes. SnapDrive® for Windows® version 5.0 introduced space reclamation for Windows NTFS file systems.

The example shown above in Figure 6 showed how space usage between the host file system and the LUN diverged over time. Space reclamation can in theory eliminate the difference between the capacity used from the LUN and the amount used within the host file system. In practice space reclamation cannot completely eliminate the difference because of the reoccurring wait periods prior to calling the reclamation process and the time to complete the reclamation process, but it can greatly reduce it.

The first three steps in Figure 7, below, show the same results as the example in Figure 6 without space reclamation. Step 4 in Figure 7 shows the results after space reclamation has run. Space reclamation is done by SnapDrive for Windows going through the host file system, locking a group of free blocks on the host file system, and sending a special command to the storage system for those blocks telling the storage system that those blocks are no longer holding used data. The blocks are temporarily locked to avoid an application trying to use those blocks during the reclamation process. The storage system then frees those blocks.

Figure 7) LUN filling with new files.



In addition to space reclamation built natively into NetApp products, NetApp has also integrated with Symantec's Thin Reclamation API. For additional details, see [WP-7111: Maximize Storage Efficiency with NetApp Thin Provisioning and Symantec Thin Reclamation](#).

1.5 THIN PROVISIONING GENERAL CHARACTERISTICS

Thin provisioning is enabled per volume and per LUN.

Thin provisioning can be enabled on creation of the volume or LUN. It can also be turned on at any time. There is no data migration necessary when enabling NetApp thin provisioning.

Thin provisioning can also be turned off at any time, assuming there is sufficient storage space available in the aggregate to turn the guarantees and reservations back on.

Thin provisioning can be enabled on any number and combination of volumes or LUNs in a storage system, as desired.

Thin provisioning of LUNs and volumes can be used transparently to client hosts and applications, but precautions should be taken so that storage space is available in the aggregate as needed. If an application supports thin provisioning, use the application best practices to properly configure thin provisioning at the application level and the NetApp storage level for efficiency.

When planning and using thin provisioning, consider other products and features that will increase or decrease the amount of storage space available in the aggregate. For example, deduplication and data compression can return storage space to the aggregate, while aggregate snapshot reserves and volume snapshot reserves will take storage space from the aggregate.

Details for using thin provisioning with other features and applications are covered in the following sections of this document.

2 CONFIGURATION AND OPERATIONS

This section discusses the requirements for configuring and managing aspects of thin provisioning. Although it discusses some basic concepts, this section assumes that the NetApp storage system is already installed and running and that the reader is familiar with basic NetApp administration.

2.1 REQUIREMENTS

Thin provisioning is well integrated into NetApp storage systems, and therefore does not have many limitations. Table 3 provides a simple list of technical requirements.

Table 3) Overview of thin provisioning technical requirements.

	Requirement
Hardware	All series Note: The V-Series systems and IBM N series systems are also supported
Minimum Data ONTAP version required	Data ONTAP 7.0, although this document is based on thin provisioning functionality that has been available since Data ONTAP 7.2.4
Licenses required	None
Volume types supported	FlexVol only, no traditional volumes
LUN support	Yes
Supported protocols	All

In addition to the technical requirements for thin provisioning are some very important environmental requirements. Table 4 lists these essential environment requirements that absolutely must be in place for a successful thin provisioning implementation.

Table 4) Overview of thin provisioning environmental requirements.

	Requirement
Monitoring	A monitoring infrastructure, such as NetApp Operations Manager, is required for determining storage growth, utilization, and availability.
Storage procurement	Predictable and known storage procurement cycles are required so that new storage will be available as storage is used, requiring more storage to be added to the shared storage pools.
Alert system	A predefined process must be in place to generate e-mail and/or paging alerts, triggered by nearly out-of-space or out-of-space conditions.

2.2 KEY DATA ONTAP PARAMETERS TO CONSIDER

The following subsection describes Data ONTAP parameters that affect thin provisioning.

AGGREGATE SNAPSHOT RESERVE

The aggregate snapshot reserve is space reserved within the aggregate for aggregate Snapshot™ copies, which are point-in-time, read-only images of the aggregate. The reserve is most often used in conjunction with SyncMirror®, and use cases in which the content of the entire aggregate should be recorded. This variable is an aggregate-level option and is set as a percentage of the aggregate. Data ONTAP removes the defined percentage of space from the aggregate. The default aggregate snapshot reserve for Data ONTAP versions up to 8.0.1 is 5%, but is not a requirement and can be changed to 0% unless SyncMirror is configured on the aggregate. It is included in this section since it will affect the amount of storage space available in the aggregate, and should be properly considered. If used, this space is reserved for aggregate Snapshot copies only, and is not available for use by volumes or LUNs, or for file usage. Aggregate snapshot reserve is beyond the scope of this document. For information about the aggregate snapshot reserve refer to the “Data ONTAP System Administration Guide.”

VOLUME SPACE GUARANTEE

Volume space guarantees define when space will be allocated or prereserved from the aggregate for a flexible volume. The main volume space guarantee settings to consider for thin and thick provisioning are `none` and `volume`. Setting the volume space guarantee to `none` will result in the volume using thin provisioning. Setting the volume space guarantee to `volume` will result in the volume using thick provisioning.

- **Volume**—A guarantee set to “volume” enables the amount of space required by the FlexVol volume to be always available from its aggregate. This is the default setting for FlexVol volumes. With the space guarantee set to “volume,” the space is subtracted, or reserved, from the aggregate’s available space at volume creation time. The space is reserved from the aggregate regardless of whether it is actually used for storing data or not.

The following example shows the creation of a 20GB volume with guarantee set to “volume.” The `df` command is used before and after the `volume create` command to show the space usage within the aggregate. It shows how the 20GB of space are removed from the aggregate as soon as the volume is created, even though no data has actually been written to the volume.

```
netapp1> df -A -g aggr0
Aggregate      total      used      avail      capacity
aggr0          85GB       0GB       85GB       0%
aggr0/.snapshot 4GB       0GB       4GB       0%
netapp1>
netapp1> vol create flex0 aggr0 20g
```

Creation of volume 'flex0' with size 20g on hosting aggregate 'aggr0' has completed.

```
netappl>
netappl> df -g /vol/flex0
File system      total      used      avail      capacity  Mounted on
/vol/flex0/      16GB       0GB       16GB       0%        /vol/flex0/
/vol/flex0/.snapshot 4GB       0GB       4GB       0%        /vol/flex0/.snapshot
netappl>
netappl> df -A -g aggr0
Aggregate        total      used      avail      capacity
aggr0            85GB      20GB      65GB      23%
aggr0/.snapshot 4GB       0GB       4GB       0%
```

Since the space has already been reserved from the aggregate, write operations to the volume will not cause more space from the aggregate to be used.

- **None**—A guarantee set to “none” reserves no space from the aggregate during volume creation. Space is not taken from the aggregate until data is actually written to the volume. The following example shows the creation of a 20GB volume with guarantee set to “none.” The `df` command is used periodically throughout the process to show how the space usage within the aggregate is affected along the way. The example also includes the creation of a LUN, which by default has LUN space reservation enabled. Note that the LUN space reservation does not cause space to be prereserved from the aggregate. Write operations to space-reserved LUNs in a volume with `guarantee=none` will fail if the containing aggregate does not have enough available space. The LUN reservation enables the LUN to have space in the volume, but setting the volume space guarantee to “none” doesn’t assure that the volume has space in the aggregate. LUN reservation is covered as a separate bullet item later in this section.

```
netappl> df -A -g aggr0
Aggregate        total      used      avail      capacity
aggr0            85GB      0GB      85GB      0%
aggr0/.snapshot 4GB       0GB      4GB       0%
netappl>
netappl> vol create noneflex -s none aggr0 20g
Creation of volume 'noneflex' with size 20g on hosting aggregate
'aggr0' has completed.
netappl>
netappl> df -g /vol/noneflex
File system      total      used      avail      capacity  Mounted on
/vol/noneflex/   16GB       0GB       16GB       0%        /vol/noneflex/
/vol/noneflex/.snapshot 4GB       0GB       4GB       0%        /vol/noneflex/.snapshot
netappl>
netappl> df -A -g aggr0
Aggregate        total      used      avail      capacity
aggr0            85GB      0GB      85GB      0%
aggr0/.snapshot 4GB       0GB      4GB       0%
netappl>
netappl> lun create -s 10g -t windows /vol/noneflex/foo
Mon Nov 24 15:17:28 EST [array1: lun.vdisk.spaceReservationNotHonored:notice]:
Space reservations in noneflex are not being honored, either because the volume
space guarantee is set to 'none' or the guarantee is currently disabled due to
lack of space in the aggregate.
lun create: created a LUN of size: 10.0g (10742215680)
netappl>
netappl> df -g /vol/noneflex
File system      total      used      avail      capacity  Mounted on
/vol/noneflex/   16GB      10GB      6GB       0%        /vol/noneflex/
/vol/noneflex/.snapshot 4GB 0GB 4GB 0% /vol/noneflex/.snapshot
netappl>
netappl> df -A -g aggr0
Aggregate        total      used      avail      capacity
aggr0            85GB      0GB      85GB      0%
aggr0/.snapshot 4GB       0GB      4GB       0%
```

- **File**—A guarantee set to “file” is a possible option, but NetApp does not recommend setting it. It should only be used if recommended by NetApp Technical Support or other NetApp personnel.

VOLUME SNAP RESERVE

The volume snap reserve is space reserved within the volume for volume Snapshot copies, which are point-in-time, read-only images of the volume. This variable is a volume-level option and is set as a percentage of the volume. Data ONTAP removes the defined percentage of space from the volume. The default volume snap reserve for Data ONTAP versions up to 8.0.1 is 20%, but it is not required and can be changed. If used, this space is reserved for Snapshot copies only, and it is not used for creating LUNs or for file usage. As Snapshot copies need space, they consume space in the snap reserve area. By default after the snap reserve area is filled, the Snapshot copies start to take space from the general volume.

LUN OVERWRITE RESERVE (FRACTIONAL_RESERVE OPTION)

The LUN overwrite reserve is the space reserved within a volume for LUNs to use for non-Snapshot data after all other storage space in the volume is used. When used, the LUN overwrite reserve is used for FlexVol volumes being used in SAN environments. The LUN overwrite reserve is not necessary in NAS environments, and therefore NetApp does not recommend it in NAS environments to avoid storage waste. In SAN environments, it is used as a safeguard so that there is enough space set aside in the volume for data that is being overwritten in the LUNs between Snapshot copies, especially if large amounts of new data are written between Snapshot copies, causing the Snapshot copies to use up excessive space in the volume.

The LUN overwrite reserve is a volume-level option and is set as a percentage. This percentage is calculated based on the current used blocks of all LUNs within the volume at the time of the Snapshot copy. Data ONTAP removes the defined percentage of space from the volume at the time of Snapshot copy creation. With the introduction of Snapshot autodelete functionality in Data ONTAP 7.1, NetApp generally recommends setting the LUN overwrite reserve to 0 and using the autodelete functionality. Autodelete is the option used to delete Snapshot copies based on a threshold, and is discussed in a subsequent section of this document.

If the LUN overwrite reserve is used, this space is reserved as a last-resort space for storing data written to the LUN when no other space is available in the volume. This reserve will never allow itself to be used to store Snapshot copies. Snapshot copies will take space from the volume snap reserve area until full, and then will take space from the volume’s free space. Snapshot copy creation will fail when there is no usable space, except for the LUN overwrite reserve. Adjusting the LUN overwrite reserve is allowed for FlexVol volumes in which the volume space guarantee is set to “volume,” and, beginning with Data ONTAP 7.3.3, in which the space guarantee is set to “none.” When the volume space guarantee is set to “none” or “file,” the LUN overwrite reserve is restricted to 100%. The amount of LUN overwrite space reserved can be observed using the “-r” option of the `df` command. The default LUN overwrite reserve, 100%, is the most conservative and least space-efficient value. When possible, it can be adjusted to a lower value that coincides with the amount of data that is being overwritten in the LUNs between Snapshot copies, while still providing a sufficient buffer in case the volume is filled with new data or Snapshot copies.

Figure 8) LUN overwrite reserve examples.

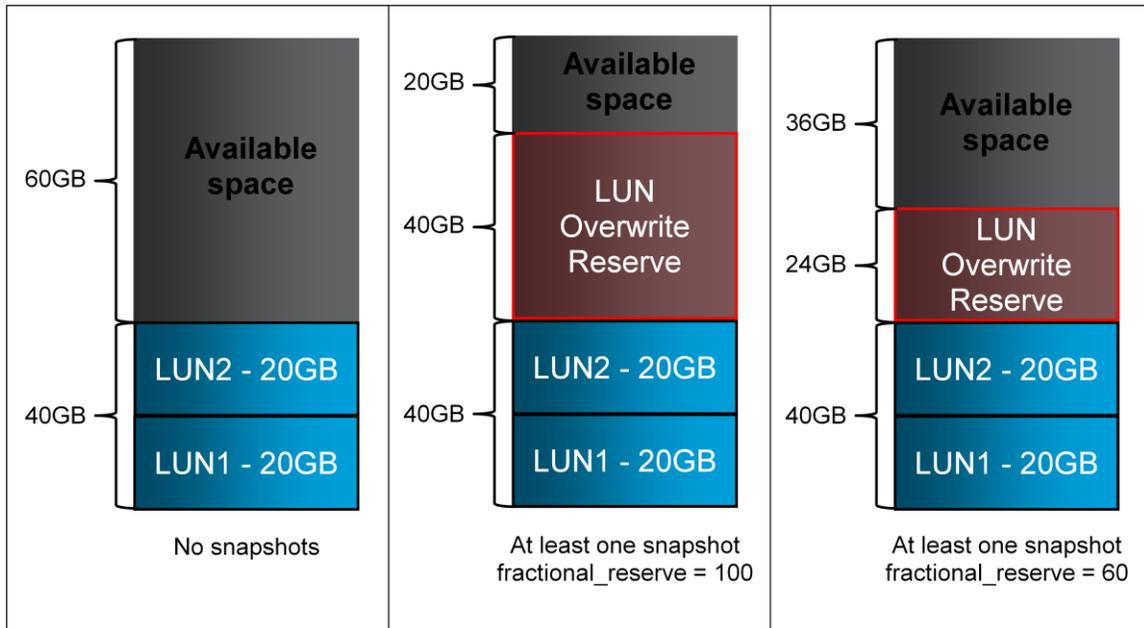


Figure 8 shows how different values for the LUN overwrite reserve will affect the available space within the thick-provisioned volume (guarantee = volume). The left-most part of Figure 8 shows a 100GB volume containing two 20GB LUNs. The center part of Figure 8 shows what would happen if the LUNs were full and the default LUN overwrite reserve size was being used. Triggered by the creation of a Snapshot copy, Data ONTAP would reserve 40GB (2 x 20GB) of space in the volume so that there was enough space for the data within the LUNs to be overwritten. Using 100% enables room for the data, even if both LUNs were to be completely overwritten with new data. The right-most part of Figure 8 shows what would happen in the same scenario if the LUN overwrite reserve was set to 60%. In this case, when the Snapshot copy was created, instead of reserving 40GB in the volume Data ONTAP would reserve 24GB (60% * [2 x 20GB]).

VOLUME AUTOSIZE

This volume setting (available in Data ONTAP 7.1 and later) defines whether a volume should automatically grow to avoid filling up to capacity. This option is available only for flexible volumes. Define the growth increments with the “-i” option. The default growth increment is 5% of the volume size at the time of creation. Define how large the volume is allowed to grow with the “-m” option. The default maximum size to grow to is 120% of the original volume size. Volume autosize is off by default.

VOLUME SNAPSHOT AUTODELETE

This volume setting (available as of Data ONTAP 7.1 and later) allows Data ONTAP to delete Snapshot copies if a threshold is met. This threshold is called a “trigger” and can be set so that Snapshot copies will be automatically deleted for one of the following trigger conditions:

- **volume**—The volume is nearly full and the volume’s snap reserve has been exceeded. This is reported in the first line reported for each volume in the `df` command. It should be noted that the volume can be full even though there might still be space in the “snap_reserve” areas, in which case autodelete will not be triggered.
- **snap_reserve**—The volume’s snap reserve space is nearly full.
- **space_reserve**—The space reserved in the volume is nearly full and the volume’s snap reserve has been exceeded. This is the space determined by the LUNs with space reservations enabled and the

LUN overwrite reserve (fractional_reserve) option. This option is not necessary, and NetApp does not recommend it, in a NAS environment.

NetApp strongly recommends that the trigger be set to volume.

The order in which Snapshot copies are deleted is determined by the following three options:

- `delete_order`—This option determines whether the oldest or newest Snapshot copies should be deleted first.
- `defer_deleted`—This option allows the user to define a group of Snapshot copies that should first be deleted when no other Snapshot copies are available. It is possible to defer the deletion of user-created Snapshot copies, scheduled Snapshot copies, or Snapshot copies beginning with a configurable prefix.
- `commitment`—This option determines how Snapshot copies used for SnapMirror® and dump operations should be handled. If set to “try,” this option will only delete these Snapshot copies if they are not locked. If this option is set to “disrupt,” these Snapshot copies will be deleted even if they are locked. These will always be the last Snapshot copies deleted. In many configurations, deleting the last SnapMirror Snapshot copy is not desired because a new full baseline copy will be required to resume mirroring operations. If, for example, the source and destination are at different sites, recreating this baseline can be a time-consuming and costly process. The recommended setting is “try.”

The algorithm to determine which Snapshot copy to delete will first look for a Snapshot copy that does not meet the “defer_delete” criteria and use the `delete_order` to determine whether to delete the oldest or more recent Snapshot copy. If no such Snapshot copy is found, the “defer_delete” criteria will be ignored in the selection process. If a Snapshot copy is still not available for deletion, then SnapMirror and dump Snapshot copies will be targeted depending on the “commitment” option.

Snapshot copies will stop being deleted when the free space in the trigger criteria reaches the value of the `target_free_space` variable, which defaults to 80%.

If both autosize and autodelete are enabled and the autodelete trigger is set to “volume,” the “try_first” volume option will determine whether a volume grow or Snapshot delete will be attempted first.

LUN RESERVATION

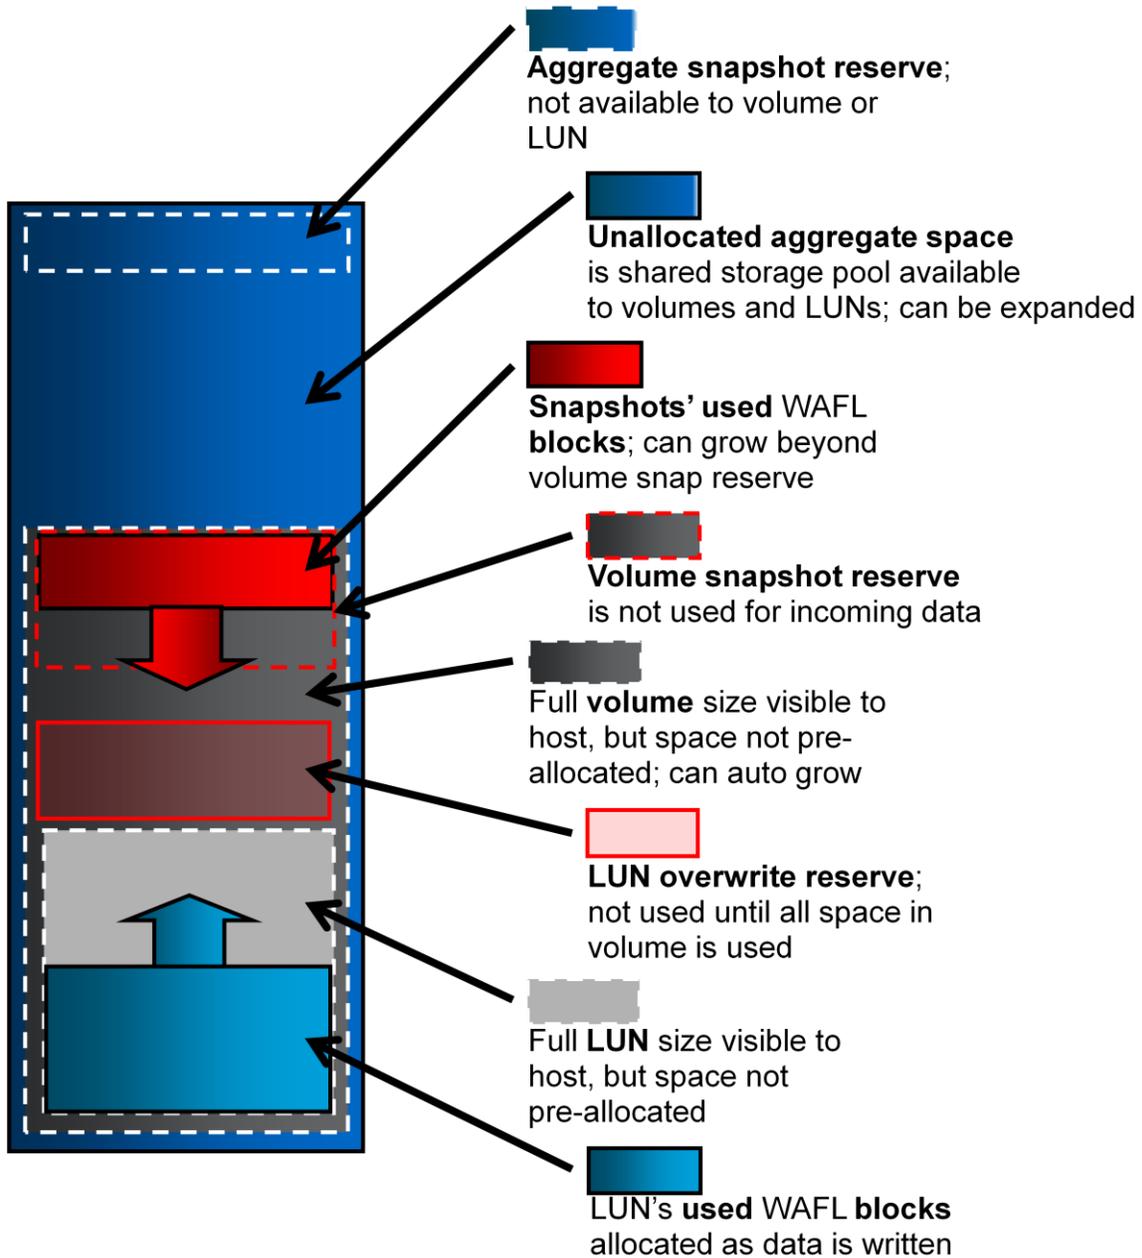
The LUN reservation determines when space is allocated for a LUN from its parent volume. If a LUN reservation is enabled, the space is subtracted from the parent volume when the LUN is created. For example, consider an aggregate with 100GB of space that contains a single space-guaranteed volume of size 80GB. This means that the aggregate has 20GB available, and 80GB was preallocated to the volume when it was created. Now consider that a 20GB LUN is created, with reservation enabled, in the 80GB volume. 20GB of space would be removed from the volume on creation of the LUN, reducing the available space in the volume to 60GB, even though the LUN is empty. The aggregate space would not be affected by the LUN reserve setting in this case, and would continue to show 80GB used (preallocated to the parent volume) and 20GB free.

If a LUN is created with reservation disabled, no space will be removed from the volume on creation of the LUN; instead, only space will be taken from the volume as data is written to the LUN. Using the previous example, if a 20GB LUN is created without LUN space reservation enabled, the free space in the volume will remain at 80GB and will only change as data is written to the LUN.

PARAMETERS AT A GLANCE

Figure 9 provides a visual overview of volume and LUN thin provisioning, including how the parameters previously described affect space availability.

Figure 9) LUN and volume thin provisioning space availability.



PARAMETER QUICK REFERENCE TABLE

Table 5 is a quick reference table of the parameters that were covered in this section, plus one additional parameter, "volume grow area." The "volume grow area" does not correlate to a specific setting, per se, but to how the aggregate size is maintained. A value of "yes" specifies that space in the aggregate must be maintained to allow the growth of a volume. This table is a simple way to display key settings for different use cases, and will be used as such within this document.

Table 5) Basic volume and LUN provisioning parameters.

Parameter	Value
Aggregate	
volume grow area	yes no
Volume Options	
guarantee	volume file none
fractional_reserve	W%
autosize	on off
autosize options	-m X% -i Y%
try_first	volume_grow snap_delete
Volume Snapshot Options	
reserve	Z%
autodelete	on off
autodelete options	volume snap_reserve space_reserve
LUN Options	
space reservation	Enabled disabled

2.3 INSTALLING AND LICENSING THIN PROVISIONING

Thin provisioning is a no-cost feature that comes with NetApp storage systems, and no special licenses are required to use thin provisioning. Using thin provisioning is a simple matter of system configuration.

NetApp thin provisioning is supported for NetApp FAS and V-Series models, as well as corresponding IBM N series systems.

2.4 COMMAND SUMMARY

This section describes the various commands and parameters that can be used in combination with thin provisioning. This information provides an overview of the commands, which are discussed in more detail later in this document.

Table 6) Provisioning commands and parameters.

Command	Summary
<p>Volume provisioning (new volume):</p> <pre>vol create flexvol_name aggr_name -s [volume file none] size [k m g t]</pre>	<p>Creates a FlexVol volume with the specified space guarantee type (volume, file, or none) and the desired size of the guaranteed space (KB, MB, GB, or TB).</p> <p>To use volume thin provisioning set guarantee to “none.”</p> <p>To use volume thin provisioning but allow individual LUNs to reserve space at the aggregate level (for LUN-level thick provisioning) set guarantee to “file.”</p> <p>To use volume thick provisioning set guarantee to “volume.”</p>
<p>Volume provisioning (existing volume):</p> <pre>vol options flexvol_name aggr_name guarantee [volume file none]</pre>	<p>Changes the space guarantee setting of an existing FlexVol volume.</p> <p>To use volume thin provisioning set guarantee to “none.”</p> <p>To use volume thin provisioning but allow individual LUNs to reserve space at the aggregate level (for LUN-level thick provisioning) set guarantee to “file.”</p> <p>To use volume thick provisioning set guarantee to “volume.”</p>
<p>LUN overwrite reserve (fractional_reserve):</p> <pre>vol options flexvol_name fractional_reserve percent</pre>	<p>Reserves space within a volume for LUNs to use for non-snapshot data after all other storage space in the volume is used. It is used as a safeguard so that there is enough space set aside in the volume for data that is being overwritten in the LUNs between Snapshot copies, even if large amounts of new data are written between Snapshot copies, causing the Snapshot copies to use up excessive space in the volume. (default = 100%)</p>
<p>Volume autosizing:</p> <pre>vol autosize flexvol name [-m <size>[k m g t]] [-i <size>[k m g t]] [on off reset]</pre>	<p>Determines if a volume should grow when nearly full (default = off)</p>
<p>Try_first:</p> <pre>vol options volname try_first [snap_delete volume_grow]</pre>	<p>Determines whether autovolume growing or auto Snapshot copy deleting is tried first</p>
<p>Volume snapshot reserve:</p> <pre>snap reserve -V flexvol_name percent</pre>	<p>Reserves space within the volume for storing only Snapshot copies by removing the defined percentage of space from the volume (default = 20%)</p>
<p>Volume snapshot autodelete:</p> <pre>snap autodelete flexvol_name [on off show reset help]</pre>	<p>Determines when (if) Snapshot copies will be automatically deleted (default = off)</p>

Command	Summary
<p>LUN provisioning (new LUN):</p> <pre>lun create -s size -t ostype [-o noreserve] lun_path</pre>	<p>Creates a LUN with specified space guarantee type, reserve or noreserve (default = reserve).</p> <p>To use LUN thin provisioning use “-o noreserve.”</p> <p>To use LUN thick provisioning omit “-o noreserve” and use default.</p> <p>NOTE: LUN thick provisioning enables space to be available in the volume for the LUN, but does not ensure that space is available in the aggregate for the volume. See volume provisioning for more information.</p>
<p>LUN provisioning (existing LUN):</p> <pre>lun set reservation lun_path [enable disable]</pre>	<p>Changes the reserve setting of an existing LUN, enable or disable (default = enable).</p> <p>To use LUN thin provisioning set reservation to “disable.”</p> <p>To use LUN thick provisioning, set reserve to “enable.”</p> <p>NOTE: Setting the value of “reserve” enables space to be available in the volume for the LUN, but does not ensure that space is available in the aggregate for the volume.</p>

2.5 THIN PROVISIONING QUICK START

This section describes the steps necessary to enable and configure thin provisioning on a FlexVol volume and LUN. For full syntax and available options, refer to man pages or the “Commands: Manual Page Reference” document available at <http://now.netapp.com>.

Table 7) Thin provisioning quick start.

Description	Commands
Aggregate-level settings	
Create an aggregate	<code>aggr create <aggr_name></code>
Expand the aggregate when more space is needed, as needed	<code>aggr add <aggr_name></code>
Volume-level settings	
Enable thin provisioning during volume creation (no space guarantee)	<code>vol create <vol_name> -s none <aggr_name> <size>g</code>
Enable thin provisioning on a preexisting volume (remove space guarantee)	<code>vol options <vol_name> guarantee none</code>
Set LUN overwrite reserve	<code>vol options <vol_name> fractional_reserve 0</code>
Enable volume autosizing	<code>vol autosize <vol_name> -m <max_size>g -I <increment_size>g on</code>

Description	Commands
Configure if volume growing or snapshot deleting will be tried first	<code>try_first [volume_grow snap_delete]</code>
Configure size of snapshot reserve	<code>snap reserve -V <vol_name> <percent></code>
Enable automatic snapshot deleting	<code>snap autodelete <vol_name> on</code>
Set when automatic snapshot deleting will begin	<code>snap autodelete <vol_name> trigger volume</code>
Configure automatic snapshot deleting behavior	<code>snap autodelete <vol_name> delete_order oldest_first</code>
LUN-level settings	
Enable thin provisioning during LUN creation (no space reservation)	<code>lun create -s <size>g -t <ostype> -o noreserve /vol/<vol_name>/<lun_name></code>
Enable thin provisioning on a preexisting LUN (remove space reservation)	<code>lun set reservation /vol/<vol_name>/<lun_name> disable</code>

2.6 THIN PROVISIONING POLICIES IN PROVISIONING MANAGER

Provisioning Manager provides a fast point-and-click approach to provision storage with provisioning policies. Using policies for NAS and SAN environments is easier than using scripts, and helps to minimize the risk of misconfigured storage.

Note: Versions of Provisioning Manager before 4.0 could not use autosize and autodelete; the policy template determined if these features were used and which options were selected. In order to implement the provisioning methods outlined, a provisioning script is required.

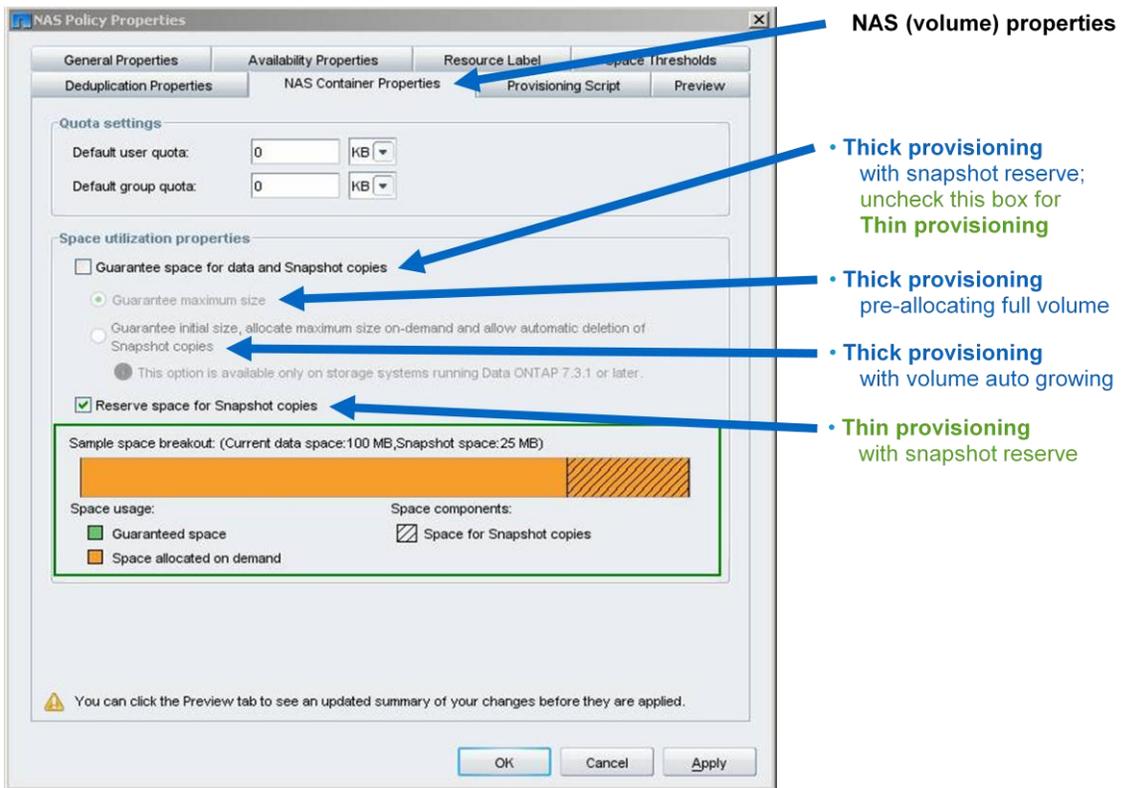
Note: The Provisioning Manager conformance tests will enforce its policies, which may not be in alignment with the intended policies. If Provisioning Manager is used to set up thin provisioning and then some parameters are changed with a postprocessing script, Provisioning Manager will revert the changes to match the original Provisioning Manager policy.

For more information, refer to [TR-3710: Operations Manager, Provisioning Manager, and Protection Manager Best Practices Guide](#).

VOLUME PROVISIONING WITH PROVISIONING MANAGER

Figure 10, below, shows a screenshot of the volume provisioning section of Provisioning Manager. Each option corresponds to a predefined volume policy, as shown. It allows the easy and consistent setting of NAS policy properties, including the guarantee space for data and Snapshot copies, automatic volume growing, and automatic deletion of Snapshot copies. Uncheck the first option to specify volume thin provisioning, which specifies that no space will be allocated for the data portion of the volume until data is written to it. Checking the box to Reserve space for Snapshot copies is allowed when using thin provisioning.

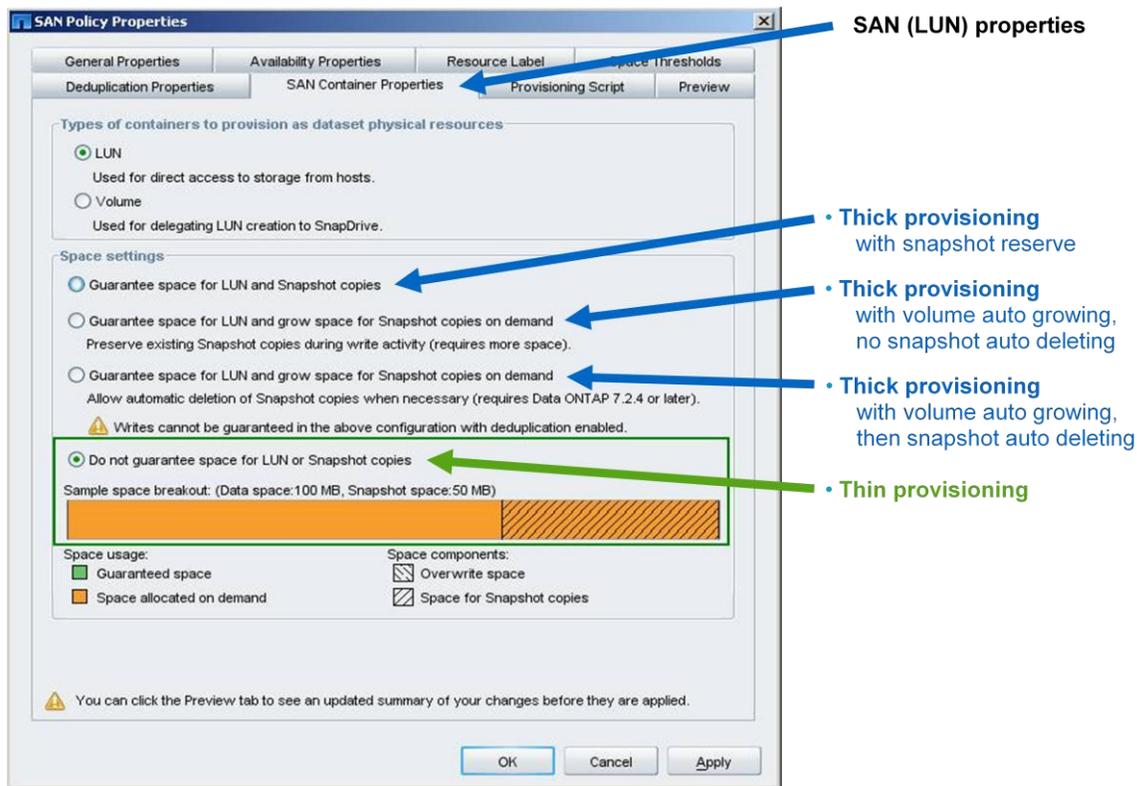
Figure 10) Volume provisioning with Provisioning Manager.



LUN PROVISIONING WITH PROVISIONING MANAGER

Figure 11, below, shows a screenshot of the LUN provisioning section of Provisioning Manager. Each option corresponds to a predefined LUN policy, as shown. It allows the easy and consistent setting of SAN policy properties, including the guarantee space for LUN and Snapshot copies, automatic volume growing, and automatic deletion of Snapshot copies. Use the last option to specify LUN thin provisioning, which specifies that no space will be allocated for the data portion of the LUN until data is written to it.

Figure 11) LUN provisioning with Provisioning Manager LUN provisioning.



2.7 MONITORING THIN PROVISIONING WITH OPERATIONS MANAGER

NetApp Operations Manager provides the capability to monitor and manage the shared storage pool in a thin-provisioned environment. Key components for thin provisioning are the ability to discover storage, the ability to set thresholds and alerts, the ability to report on utilization and trending, the ability to monitor performance, and the ability to configure the storage. Monitoring these components is critical for achieving storage efficiency while maintaining availability. Operations Manager is a client server application that delivers this key functionality through a single console. It also provides an application programming interface (API) for integration with third-party management and orchestration software vendors.

Operations Manager monitors the NetApp shared storage infrastructure and is able to alert on storage-related incidents. Events can be set to trigger an action when operational parameters are within a certain range and indicate a relevant situation. When the event triggers an alarm, notification can be sent by e-mail, pager, Simple Network Management Protocol (SNMP), or customized scripts.

LUNs can be a bit more challenging to monitor in certain instances. A simple approach for monitoring LUNs is to put LUNs in qtrees, and put a soft quota on the qtree. The qtree report will include key utilization information for each LUN. To interact with a monitoring tool that can receive SNMP traps, make the quota a threshold quota.

THRESHOLDS

Operations Manager monitors relevant parameters that indicate the presence of specified situations, such as volumes or aggregates becoming nearly full, or overcommitted levels approaching high levels. Operations Manager provides the ability to set thresholds that will trigger events when certain conditions are met. These thresholds can be set to trigger actions, for example, to notify the operational team that an

impending situation is occurring and should be reviewed. The thresholds can be set at the global level, or for individual aggregates.

In Operations Manager, one way the thresholds can be verified and set is by navigating to the Default Thresholds page at Setup→Options→Default Thresholds. Figure 12 shows a sample configuration page.

Figure 12) Operations Manager screen (partial) to configure thresholds on operational metrics.

Default Thresholds	
Host CPU Too Busy Threshold (%)	95
Host CPU Busy Threshold Interval	15 minutes
Aggregate Full Threshold (%)	90
Aggregate Nearly Full Threshold (%)	80
Aggregate Full Threshold Interval	0 seconds
Aggregate Overcommitted Threshold (%)	100
Aggregate Nearly Overcommitted Threshold (%)	95
Aggregate Snapshot Reserve Full Threshold (%)	90
Aggregate Snapshot Reserve Nearly Full Threshold (%)	80
Aggregate Nearly Over Deduplicated Threshold (%)	140
Aggregate Over Deduplicated Threshold (%)	150
Volume Full Threshold (%)	90
Volume Nearly Full Threshold (%)	80
Volume Full Threshold Interval	0 seconds

Monitoring the aggregates is very important. The aggregates are used as shared storage pools, physical containers of preallocated and growable storage objects that host application data. If an aggregate using thin provisioning is not monitored, it could have direct consequences for other applications using the aggregate.

The thresholds can be set at different levels according to the criticality of the event. There are warning thresholds and error thresholds that signal if a critical state is reached. When determining the values for the thresholds, it is imperative that the flexibility and dependability of the plans to address the conditions are taken into consideration. These plans are referred to as storage mitigation plans within this document. The storage mitigation plans describe the actions that are to be taken when specific conditions occur. Some instances will mandate conservative threshold settings, such as cases in which storage mitigation can only be scheduled during downtime windows. Other instances will allow more aggressive threshold settings, such as cases in which mitigation can occur as part of normal day-to-day operations.

Below are the thresholds that Operations Manager provides for aggregates. They represent absolute limits. Operations Manager alarms can be used to notify operational staff and raise awareness if these thresholds are reached.

- Aggregate full threshold: This threshold is set to designate when the aggregates will be considered full. This is a critical state.
- Aggregate nearly full threshold: This threshold is set to designate when the aggregates will be considered nearly full. This is a warning state.

- **Aggregate overcommitted threshold:** This threshold is set to designate when the aggregates will be considered fully overcommitted. This is a critical state. This metric refers to the amount of storage that is assigned to applications. A setting of 200% reflects that the aggregates can be overcommitted up to 100% beyond the amount of storage available in the aggregates.
- **Aggregate nearly overcommitted threshold:** This threshold is set to designate when the aggregates will be considered almost fully overcommitted. This is a warning state. This metric refers to the amount of storage that is assigned to applications. A setting of 150% reflects that the aggregates can be overcommitted up to 50% beyond the amount of storage available in the aggregates.

Volume space can also be monitored using thresholds. Volume space utilization can be a critical metric in certain cases, such as using a thick-provisioned volume as a storage pool for thin-provisioned LUNs. Similar to the aggregate thresholds, Operations Manager provides the following thresholds for volumes.

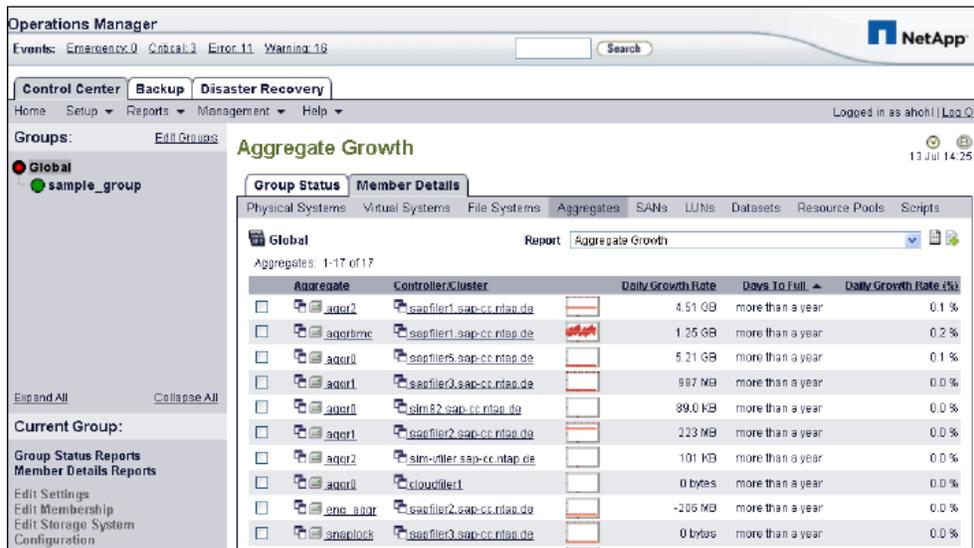
- **Volume full threshold:** This threshold is set to designate when the volume will be considered full. This is a critical state.
- **Volume nearly full threshold:** This threshold is set to designate when the aggregates will be considered nearly full. This is a warning state.
- **Volume growth event minimum change threshold:** This threshold is set to designate when the volumes have been extended by a minimum amount of change using the volume autosize functionality. This is an informative threshold. This metric can be used to track if volumes are being grown regularly, which could signify that volume space is being rapidly consumed by Snapshot copies or user data.

TRENDING

Operations Manager 4.0 and later support a variety of trending features for storage objects, such as aggregates and volumes. This is an important feature because it allows you to estimate the time frame in which storage objects are expected to fill, and when mitigation will need to be activated. The trend is calculated as a linear regression of up to 90 days in the past. For aggregates, Operations Manager calculates a trend on the daily growth rate. In Operations Manager, one way the aggregate trending information, including Daily Growth Rates and Days To Full, can be viewed is by navigating to the Aggregate Growth page at Reports→All. From there, select Aggregates from the Physical Objects section, then select Aggregate Growth under the Available Reports section and press the Show button. Click on any aggregate for more details. You can also select trending based on an interval of one day, one week, one month, three months, or one year, or you can modify the threshold settings for that aggregate. To see the effect of recent data activities, set the interval of a trend calculation to enclose this activity to investigate if growth rates calculated over different intervals deviate significantly.

Figure 13 shows the Aggregate Growth page where summary information can be sorted and quickly evaluated. Clicking on different items within this page will provide corresponding details, including access to multiple trending reports.

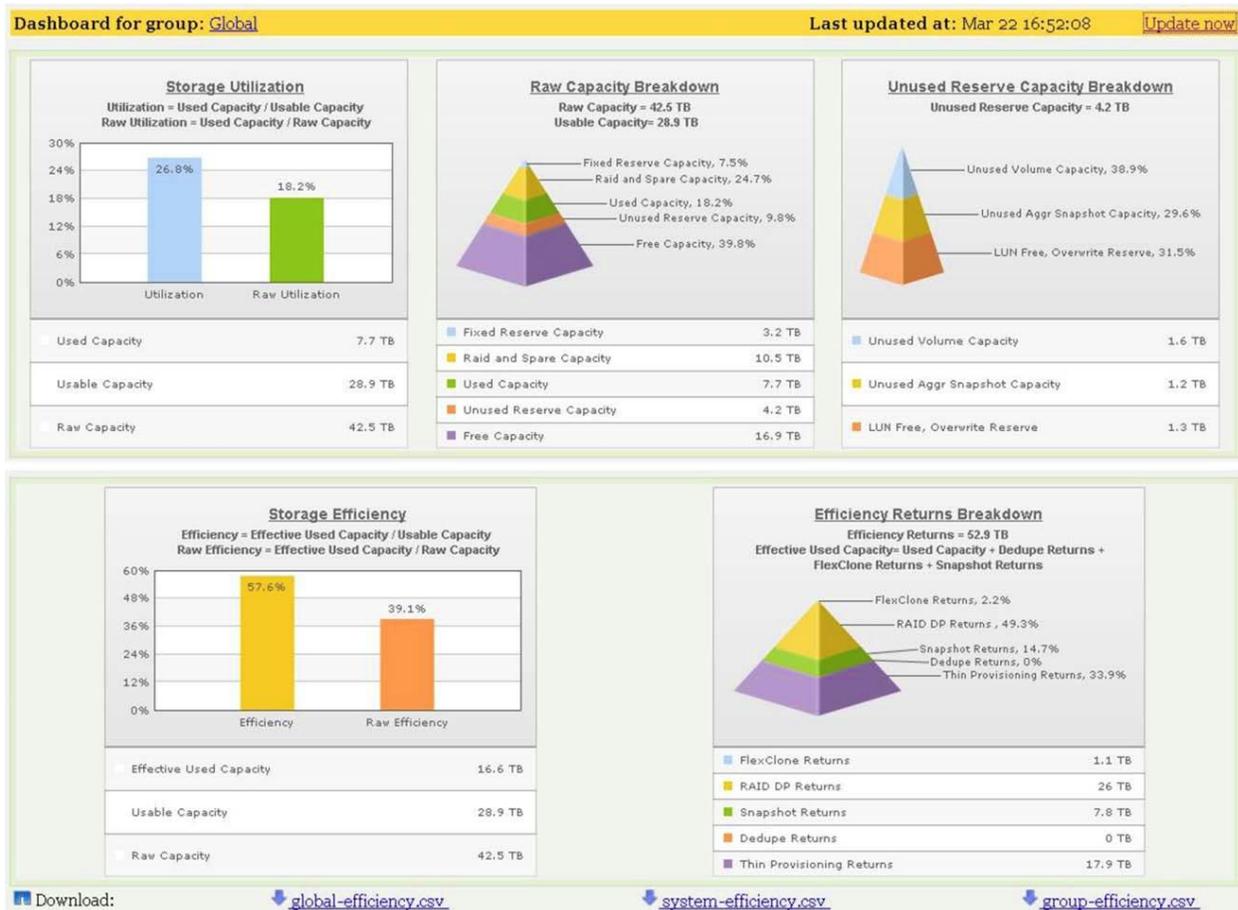
Figure 13) Trending of aggregate data growth and days-to-full prediction in Operations Manager.



MONITORING STORAGE EFFICIENCY RETURNS

NetApp Operations Manager provides a storage efficiency dashboard to visualize storage efficiency returns in the NetApp shared storage infrastructure. It includes important parameters that allow drill down into utilization, capacity, unused reserve capacity, storage efficiency, and efficiency return breakdown. This information shows the effectiveness of the NetApp storage efficiency technologies. Figure 14 provides a sample screenshot of the storage efficiency dashboard in NetApp Operations Manager. For more details, see [NetApp Operations Manager Efficiency Dashboard Installation and User Guide](#).

Figure 14) Operations Manager storage efficiency dashboard.



2.8 NOTIFICATION

Operational staff must be notified when situations are developing that require mitigation so that there is continued storage availability for users and applications. As previously mentioned, Operations Manager provides thresholds that will trigger events for certain conditions. These events can be used to trigger notifications to key personnel, such as operational staff, storage administrators, or storage capacity planners. Alarms are instrumental in reducing the management effort required for NetApp storage.

Depending on the organizational structure, the responsibilities to operate, plan, and administer the storage infrastructure can be separated into different groups, persons, or roles. Thus, it is common practice that the alerts and mitigation activities are characterized and assigned based on the required skill set and time to act. This allows easy alignment between required actions and a given organizational structure.

Operations Manager can send notifications using different methods. These notification methods can be used separately or in combination; for example, a notification can be sent by e-mail alone or by both e-mail and SNMP for a given condition.

NOTIFY BY E-MAIL

Operations Manager supports e-mail alerts through Alarms. An alarm can be sent to multiple e-mail addresses. Repeated notifications can be sent when the situation is not resolved. To set an alarm, access

the Alarm page at Setup→Alarms in Operations Manager. Clicking Advanced Version provides additional alarm options. Figure 15 below shows a screenshot for an alarm configuration including an e-mail address. The event designated by the Event Name on the Alarms page corresponds to a threshold, which can be adjusted as discussed in the thresholds section of this document.

NetApp recommends using distribution lists or aliases with meaningful names rather than addresses of individual persons. This avoids the need to make individual adjustments to Operations Manager alerts each time there is a change in responsibilities and roles.

NOTIFY BY SNMP

Operations Manager supports SNMP alerts/traps through Alarms. The SNMP traps can be directed to an SNMP trap host that will trigger appropriate actions outside of Operations Manager. SNMP is a widely used standard that is supported by most orchestration frameworks and ticketing systems. Using SNMP, Operations Manager can be integrated into existing ticketing systems. Figure 15 shows a screenshot for an alarm configuration that includes an SNMP trap host. The SNMP trap host is configured using hostname or IP address and the port on which the SNMP agent is listening. The event designated by the Event Name on the Alarms page corresponds to a threshold, which can be adjusted as discussed in the thresholds section of this document.

Figure 15) Configuring an alarm based on the threshold, aggregate almost full.

The screenshot shows the 'Configure Alarm' interface. It is divided into several sections:

- Select Event By:**
 - Severity: All (dropdown)
 - Event Name: Aggregate Almost Full (dropdown)
- Event Class:** (regular expression to select a class of events e.g. env\temp) [Empty text box]
- Recipients:**
 - Email Recipients - Admins (specify full login names of administrators) [Empty text box]
 - Email Recipients - Non-Admins (specify email addresses) [Empty text box]
 - Page Recipients - Admins (specify full login names of administrators) [Empty text box]
 - Page Recipients - Non-Admins (specify pager addresses) [Empty text box]
 - SNMP Trap Hosts: 10.68.93.113:1621 [Text box]
- Other Options:**
 - Time From: 00:00 [Time picker]
 - Time To: 23:59 [Time picker]
 - Repeat Notify: No (dropdown)
 - Repeat Interval: 30 [Text box]
 - Disable: No (dropdown)

An 'Update' button is located at the bottom right of the form.

Note: Remember that the SNMP event must be routed to the responsible groups or persons in the ticketing system for appropriate response(s).

NOTIFY BY SCRIPT

Operations Manager supports user-defined adapters, which can be used to provide notifications in highly customized integration scenarios. A user-defined adapter can be executed, delivering the information to

the infrastructure or system of choice. A script can be used to implement such an adapter and act as the glue between Operations Manager and the customized infrastructure. For example, to set an Operations Manager alarm that will trigger a script on the event Aggregate Almost Full, use the following command line:

```
dfm alarm create -s script_to_execute -h aggregate-almost-full
```

2.9 PRESERVING STORAGE AVAILABILITY (MITIGATION ACTIVITIES)

This section focuses on actions that are used as an integral part of thin provisioning in order to preserve storage availability. These actions, referred to as mitigation activities within this document, create shared storage space as it is needed in the aggregates and/or volumes. It is crucial to trigger the mitigation activities well in advance of running out of storage. Storage availability and utilization trends are continuously monitored and evaluated using tools as described in earlier sections of this document. Then, when appropriate, the mitigation activities will be applied and reapplied as needed for the lifetime of the volumes and aggregates. Proper planning and integration of these mitigation activities will avoid the creation of storage emergencies. When implemented appropriately, the mitigation activities will safely create additional shared space as it is needed, commonly referred to as “just-in-time” storage, resulting in higher storage utilization and efficiency.

Storage tightness is defined as a storage pool that no longer meets the minimum amount of available storage, as determined by the storage administrator. Storage tightness can occur in aggregates or volumes as its storage is used (allocated).

When volumes in an aggregate are configured to use thin provisioning, they use the unallocated space within the aggregate as a shared storage pool to deal with data growth. As the available space in the aggregate runs low or reaches a limit such as the “Aggregate Nearly Full Threshold,” a mitigation activity at the aggregate level is necessary.

When LUNs in a thick-provisioned volume are configured to use thin provisioning, they use the unallocated space within the volume as a shared storage pool to deal with data growth. As the available space in the volume runs low or reaches a limit, such as the “Volume Nearly Full Threshold,” a mitigation activity at the volume level is necessary.

The next two subsections list mitigation activities that can be applied to aggregates and volumes.

MITIGATING RISK OF RUNNING OUT OF SPACE

Aggregates are the coarsest storage object within a NetApp storage controller. When used, thin provisioning is applied on finer-grained storage objects such as FlexVol volumes and LUNs. While the volumes and LUNs might grow with demand, because they live within an aggregate of physically limited size, the growth of the volumes and LUNs are limited. The thin provisioning plan should include one or more actions, or mitigation activities, listed below so that additional space is adequately provided in the aggregate for continued availability. Alternatively, the mitigation activities listed below could also be used to address storage emergencies, such as an unplanned event causing storage allocation within the aggregate to spike faster than storage can be made available. Such a storage emergency would obstruct the ability of any contained volumes and LUNs to grow, and should be addressed immediately.

1. **Expand the aggregate:** Disk drives can be added to aggregates during operation without interruption. You can repeat this mitigation activity until the maximum size of the aggregate is reached. The maximum aggregate size depends on the Data ONTAP version, the type of aggregate, and the type of storage controller. 64-bit aggregates are supported with Data ONTAP 8, and have very high aggregate size limits. Additional drives can be used immediately; however, remember to consider the procurement cycle so that the drives will be available in time to meet the demand. Rebalancing data between existing and new drives results in a uniformly distributed use of the drives.
2. **Decrease aggregate Snapshot copy reserve:** This reserve is needed in MetroCluster™ and for SyncMirror® configurations. In other configurations, you can decrease this reserve or set it to zero.

3. Shrink preallocated volumes: If volumes with unused, preallocated space exist (guarantee = volume), consider reducing their size to make the unused space available to other volumes and LUNs in the aggregate. Remember to consider the utilization trends and expected application behavior when determining if space in the volume should be unallocated.
4. Use deduplication on volumes: Enable deduplication on volumes that will realize savings. If the volume is using thin provisioning, the deduplication savings will be returned to the aggregate. If the volume is using thick provisioning, the deduplication savings will be returned to the volume and available to the LUNs within that volume. The volume can also be reduced in size to return the deduplication savings within the volume to the aggregate.
5. Online migration of a MultiStore vFiler: Migrate a migration candidate storage container online to a different aggregate, where a migration candidate is a storage container belonging to a vFiler that has been previously designated as movable. MultiStore[®] allows a single controller to provide secure multi-tenancy where multiple separate storage systems are hosted on a single controller. MultiStore includes the ability to transparently migrate storage containers between controllers of different high-availability pairs with no downtime, which can be used to achieve this mitigation activity. VMware also provides mobility with VMotion[™], allowing transparent migration of storage containers with no downtime, but this is outside the scope of this document. Adequate free space must be available in the target aggregate. This mitigation activity is not limited in its applicability. For additional details about DataMotion[™] for vFiler[®], refer to [TR-3814: NetApp DataMotion for vFiler](#).
6. Online migration of a LUN from one aggregate to another: Starting with Data ONTAP 8.1 7-Mode, NetApp DataMotion for LUNs provides the ability to transparently migrate LUNs (FCP, FCoE, or iSCSI) between aggregates with no downtime, which can also be used to achieve this mitigation activity. For additional details about NetApp DataMotion for Volumes, refer to [TR-3873: NetApp DataMotion for Volumes](#), and [TR-3881: DataMotion for Volumes for Enterprise Applications](#).
7. Migrate a volume, with minimal downtime, from an aggregate to another aggregate within the same or different controller: Migrate a migration candidate volume to a different aggregate with minimal downtime, where a migration candidate is a volume that has been previously designated as movable. To accomplish this requires a SnapMirror license. If one of the controllers is a non NetApp controller a V-Series license is also required. SnapMirror replicates the data while it is still served. To switch over to the replicated data, the client needs to detach from the source and reattach to the replica. After completion, the replica is considered the new source. This operation has an impact on client downtime. Typically, inter-data-center bandwidth allows you to synchronize the source and the replica within the range of a few minutes. Adequate free space must be available in the target aggregate. This application is limited to applications and scenarios that will allow a short downtime period.
8. Stop application for mitigation activities: If none of the mitigation activities listed above can be used, the application must be stopped to achieve a consistent state.

Table 8 summarizes the list of mitigation activities that can be used to aggregate tightness. Note that Provisioning Manager performs mitigation alternative 3 to 6 for secondary storage online.

Table 8) Mitigation activities for preserving space availability within aggregates.

No.	Mitigation Activity	Repeatability	SLA Impact	Preparation Time	Time to Show Effect
1	Increase aggregate capacity by adding disks (Data ONTAP 7.x)	Low limits	None	Hardware procurement	Immediate
	Data ONTAP 8	High limits			
2	Decrease the aggregate's Snapshot copy reserve area if configured and not used	One time	None	None	Immediate

No.	Mitigation Activity	Repeatability	SLA Impact	Preparation Time	Time to Show Effect
3	Shrink other volumes in the aggregate if they have enough free space	One time	Low	None	Immediate
4	Run deduplication and shrink volumes	Repeatable	Low/possible performance impact	Time to execute dedupe	Immediate
5	Migrate volumes to a different aggregate; same or different controller (online)	Repeatable	Low	None	Minutes: DataMotion for vFiler
6	Migrate LUNs to a different aggregate(online)	Repeatable	Low	None	Minutes: DataMotion for LUNs
7	Migrate volumes to a different aggregate (offline)	Repeatable	Med-high	Next planned downtime window	Minutes: Volume switch-over time
8	Prevent application data loss and stop the application, then migrate (offline)	Repeatable	Low-high	Coordinate with application owner	Minutes: Migration time

PRESERVING VOLUME SPACE AVAILABILITY

It is possible to configure LUNs to use thin provisioning and for the containing volume to use thick provisioning. In this case, LUN thin provisioning will cause the LUNs to use the unallocated space within the thick-provisioned volume as a shared storage pool. While the LUNs might grow with demand, because they live within a volume of physically limited size, the growth of the LUNs are limited. The thin provisioning plan should include one or more actions, or mitigation activities, listed below so that additional space is adequately provided in the volume for continued availability. Alternatively, the mitigation activities listed below could also be used to address storage emergencies, such as an unplanned event causing storage allocation within the volume to spike faster than storage can be made available. Such a storage emergency would obstruct the ability of any contained LUNs to grow, and should be addressed immediately.

Note: Some of these mitigation alternatives depend on and affect used capacity in the aggregate.

Table 9) Mitigation activities for preserving space availability within volumes.

No.	Mitigation Activity	Repeatability	SLA Impact	Prep. Time	Time to Show Effect
1	Reduce the volume's snapshot reserve if configured and not used	One time	Low	None	Immediate
2	Increase the volume size if possible and there is free space in the aggregate	One time	Low	None	Immediate
3	Delete unneeded Snapshot copies, which may include those skipped by the autodelete function	Limited	Low	None	Immediate

4	Activate deduplication for the volume	One time	Low/possible performance impact	Time to execute dedupe	Immediate
5	If the volume contains more than a single LUN, migrate those objects to another volume or aggregate	Repeatable	High	Next planned downtime window	Minutes: Migration time
6	Stop application and migrate data	Repeatable	High	Coordinate w/app owner	Minutes: Migration time

3 BEST PRACTICES

This section provides thin provisioning best practices. The objective of this section is to provide guidance for implementing thin provisioning in the most common scenarios. See the additional sections within this document for additional information about using thin provisioning with specific applications, such as VMware®, Oracle®, SharePoint®, and others.

3.1 DEFAULT (THICK) PROVISIONING

Before considering the thin provisioning best practice considerations, consider some pros and cons of using the default traditional settings when provisioning storage.

If the default values are used during storage provisioning, the end result is storage that is using thick provisioning with extra reserve space for Snapshot copies and LUN overwrites. Table 10, below, shows the results of using all default values during storage provisioning.

Table 10) Default settings.

Parameter	Value
Aggregate	
volume grow area	NA
Volume Options	
guarantee	volume
fractional_reserve	100%
autosize	off
autosize options	NA
try_first	NA
Volume Snapshot Options	
reserve	20%*
autodelete	Off
autodelete options	NA
LUN Options	
space reservation	enabled

The objectives of these values are to enable space to always be available for incoming data and to leave Snapshot copies intact. A lot of extra space is reserved (LUN overwrite reserve) just in case there is a big data growth (or data change) spike that could cause the consumption of all the space in the LUN, the volume, and the snapshot reserve. More specifically, 100% of the size of the data that is stored is set aside for emergencies. This traditional thick provisioning approach is appropriate for certain solutions and conditions (for example, see the SQL Server® and Exchange sections of this document).

The disadvantage of using the default settings is that it is very expensive to have lots of extra storage available just in case of a data growth spike. The extra emergency reserve (LUN overwrite reserve) cannot be repurposed, and typically remains unused. Using this approach requires a good amount of guessing during the storage sizing, since this approach lacks flexibility.

So as IT is asked to do more with less storage, the question quickly moves away from “if” thin provisioning should be used and moves toward “when” thin provisioning should be used.

3.2 THIN PROVISIONING COMMON PRACTICES AND BEST PRACTICES

Before considering the recommended thin provisioning configurations, consider the following recommendations based on successful customer implementations and internal testing. These recommendations will be used to help formulate the recommended practices for thin provisioning.

- Use Storage Services when setting up thin provisioning policies so that proper considerations and best practices are used.
- Unless otherwise recommended, using separate aggregates for thin provisioning and thick provisioning, when possible, will simplify thin provisioning management.
- Aggregates can be expanded by adding additional disks until the maximum aggregate size limit for the system is reached, at which time an action must be taken to make space available in the parent aggregate (for example, volume deduplication, volume/LUN migration to a different aggregate).
- Use thick provisioning for root volumes.
- If no trending information is available at the time of sizing primary volumes that are using thin provisioning, use an approach that is similar to sizing for thick provisioning, avoiding big overcommitments. The amount of overcommitment can be adjusted at a later time when trending information is available.
- When sizing backup (for example, with SnapVault[®]) or DR (with volume SnapMirror) destination volumes, use thin provisioning and use a volume size equal to the maximum possible volume sizes for the platform, or the size of the parent aggregate. This will automatically handle volume autosizing of the primary storage.
- Thin provisioning is commonly used with FlexClone[®] for maximum space efficiency as clones are created.
- VMware provides mobility with VMotion, allowing mission-critical applications to handle out-of-space conditions with no downtime. This makes thin provisioning a possibility for mission-critical applications.
- NetApp DataMotion and MultiStore provide the capability to handle out-of-space conditions with no downtime, making thin provisioning a possibility for mission-critical applications.
- In general, VMware VMotion and NetApp DataMotion have zero downtime. When using critical applications that will use thin provisioning without a zero-downtime solution, be more cautious with monitoring, setting monitoring thresholds to a lower value.
- Solutions using thin provisioning should have predictable data growth rates.
- Using volume autosize is common practice for some applications, as discussed later. While effective at automatically providing additional space, the resizing of volumes can make space accounting and tracking challenging for other applications.
- Snapshot autodelete can be used to automatically create more space in a volume that is running low, but is often not used in order to maintain service-level agreements (SLAs). One possible alternative is using scripts with SLA logic to delete specific Snapshot copies.

3.3 RECOMMENDED THIN PROVISIONING CONFIGURATIONS

There are many possible thin provisioning configurations. Based on the best practices and common practices described in the previous section, the list can be reduced to a few key configurations.

Note: If Provisioning Manager will be used to configure thin provisioning, review the “Thin Provisioning Policies in Provisioning Manager” section of this document.

Table 11 shows the parameter settings for each of the following recommended thin provisioning configurations:

- Volume thin provisioning with no reserves
- Volume thin provisioning and LUN thin provisioning with no reserves
- Volume thick provisioning and LUN thick provisioning with volume autosize first
- Volume thick provisioning and LUN thick provisioning with autodelete first

Table 11) Comparison of provisioning methods and requirements.

Parameter	Vol ThinProv No LUNs No Reserves	Vol ThinProv LUN ThinProv No Reserves	Vol ThickProv LUN ThinProv Autosize First	Vol ThickProv LUN ThinProv Autodelete First
Aggregate				
volume grow area	NA	NA	Yes	Yes
Volume Options				
guarantee	None	None	Volume	Volume
fractional_reserve	0%	0%	0%	0%
autosize	Off	Off	On	On
autosize options	NA	NA	trigger=volume	Trigger= Snap_reserve Space_reserve
try_first	NA	NA	Volume_grow	Snap_reserve
Volume Snapshot Options				
reserve	0%	0%	0%	5%
autodelete	Off	Off	(optional) ¹	On ¹
autodelete options	NA	NA	delete_order= oldest_first	delete_order= oldest_first
LUN Options				
space reservation	NA	Disabled	Disabled	Disabled

Table 12 provides an overview of some key characteristics for each of the four recommended thin provisioning configurations, where:

- X is the size of the primary data = sum of all user data (files and directories) within the volume
- Δ is the amount of space needed to hold Snapshot data
- N is the traditional provisioning impact = amount of blocks logically allocated but not used

¹Do not use autodelete if there is SnapManager data in the volume.

²X is the size of the primary data = sum of all user data (files and directories) within the volume; Δ is the amount of space needed to hold Snapshot data; N is the traditional provisioning impact = amount of blocks logically allocated but not used.

Table 12) Characteristics of the recommended thin provisioning configurations.

Characteristics	Vol ThinProv No Reserves	Vol ThinProv LUN ThinProv No Reserves	Vol ThickProv LUN ThinProv Autosize First	Vol ThickProv LUN ThinProv Autodelete First
Space consumption ²	X – N	X - N	X	X + Δ
Space efficient	Yes	Yes	Partially, LUNs and snapshots	Partially, LUNs and snapshots
Monitoring	Required at aggregate level	Required at volume and aggregate level	Required at volume and aggregate level	Required at volume and aggregate level
Notification/Mitigation process required	Yes	Yes	Optional for most cases	Optional for most cases
Dedupe returns free space to:	Aggregate free space area	Aggregate free space area	Volume free space area	Volume free space area
Risk of an out-of- space condition on primary data	Yes, when monitoring and notification processes are missing	Yes, when monitoring and notification processes are missing	No, as long as autodelete is able to delete any snapshot	No, as long as autodelete is able to delete any snapshot
Typical use cases	Shared storage infrastructure, home dirs, DR destination, FlexClone destination, test/dev	Shared storage infrastructure, some DB data volumes, SharePoint, storage pools for virtualized servers, test/dev	Large database environments	Large database environments where OK to delete snapshots

HOW SHOULD A VOLUME BE SIZED?

Because the physical allocation of data within a volume using thin provisioning is done on demand, theoretically the volume size can be set to a very high value that can easily fit all application data and Snapshot copies. Since the unallocated space in the volume is not exclusively reserved from the aggregate, all other applications using storage within the aggregate can benefit from the unallocated (unused) space in the shared storage pool. However, NetApp recommends that you size the volume to the expected size of its containing objects and use the autosize option to let it grow on demand, if needed.

Note: Remember, if using deduplication, the maximum deduplication volume size limits should be taken into account.

3.4 THIN PROVISIONING IN PRACTICE

APPLICATION RECOMMENDATIONS

Thin provisioning is most effective when applications use data that is committed to them in increments. When applications preformat or preinitialize their storage, the immediate effect of thin provisioning is lost and only deduplication may reclaim sharable blocks. Because thin provisioning has no performance penalty, the general recommendation is to provision using thin provisioning, assuming the previously described requirements are met.

For SAN-attached storage, NetApp recommends that you use file systems supporting space reclamation technologies such as the `SCSI UNMAP` and `SCSI WRITESAME` commands. These commands pass information through the storage stack to communicate that a particular block is no longer used.

THIN PROVISIONING WITH MIGRATION CANDIDATES USING NETAPP DATAMOTION

Thin provisioning with migration candidates involves simply architecting the storage solution in a way that includes volumes or LUNs that are classified as movable to a different aggregate or controller, if needed as part of a mitigation activity. This approach helps automate the mitigation process, while increasing the utilization levels of NetApp shared storage. When the NetApp online migration features are used, mitigation to address a storage tightness scenario is relatively simple. It does not depend on planned downtime windows. Since storage is classified into potential migration candidates, migration candidates can be migrated away from an aggregate tightness situation to a different aggregate or controller while enabling continued accessibility. It provides an elegant solution to manage aggregate space availability while maintaining accessibility.

Note: DataMotion for LUNs is available with Data ONTAP 8.1 7-mode.

NetApp MultiStore provides online migration abilities through the vFiler abstraction. NetApp recommends the use of MultiStore vFiler units for online migration.

Migration candidates can be selected based on different criteria. It can be as simple as designating the oldest volumes as candidates, if it is expected that they will be used less as time goes on and can be migrated to lower-performing storage. Or a volume may be designated for migration if it goes beyond a predetermined size, or based on the fact that it has the lowest priority based on SLAs. Whatever the determining factor, it will allow an automatic response to a storage tightness scenario. Figures 16 and 17 illustrate different phases of an example scenario using the migration candidate concept, including storage tightness and volume migration from one aggregate to another. The nonmigration candidates indicate that data that will not be moved during its lifetime. Since the nonmigration candidates will never move, they may or may not use MultiStore vFiler units. Doing so could greatly simplify operation and hardware maintenance of the storage controller, but vFiler units are not a requirement for the nonmigration candidates. The migration candidate is considered movable, and thus must make use of vFiler units. The ratio between the size of the no-migration candidates and the migration candidates depends on the growth rate and lifetime of the data in the nonmigration candidates. Assuming that the aggregate size is constant over the life of the candidates, the aggregate will have been filled with nonmigration candidate data and migration candidate data. Over the data lifetime, more and more migration candidates will have been migrated away. At the end of the storage lifetime, only the nonmigration candidate data will be left.

It is irrelevant whether the data growth happens in the nonmigration candidates or migration candidates. When a migration candidate is migrated away, the resource situation on the aggregate is relaxed because more space is available in the aggregate. It is preferable to provision several migration candidates of different sizes. This allows you to:

- React to different growth scenarios of the data.
- Quickly migrate smaller migration candidates if time is limited or network bandwidth between controllers is limited.
- Operate the aggregate in its operational sweet-spot corridor over a long time frame; by dividing the storage into appropriate, smaller migration candidates, you can help enable the aggregate to operate in a predefined use interval.

Figure 16) Example condition of aggregate tightness and action now needed.

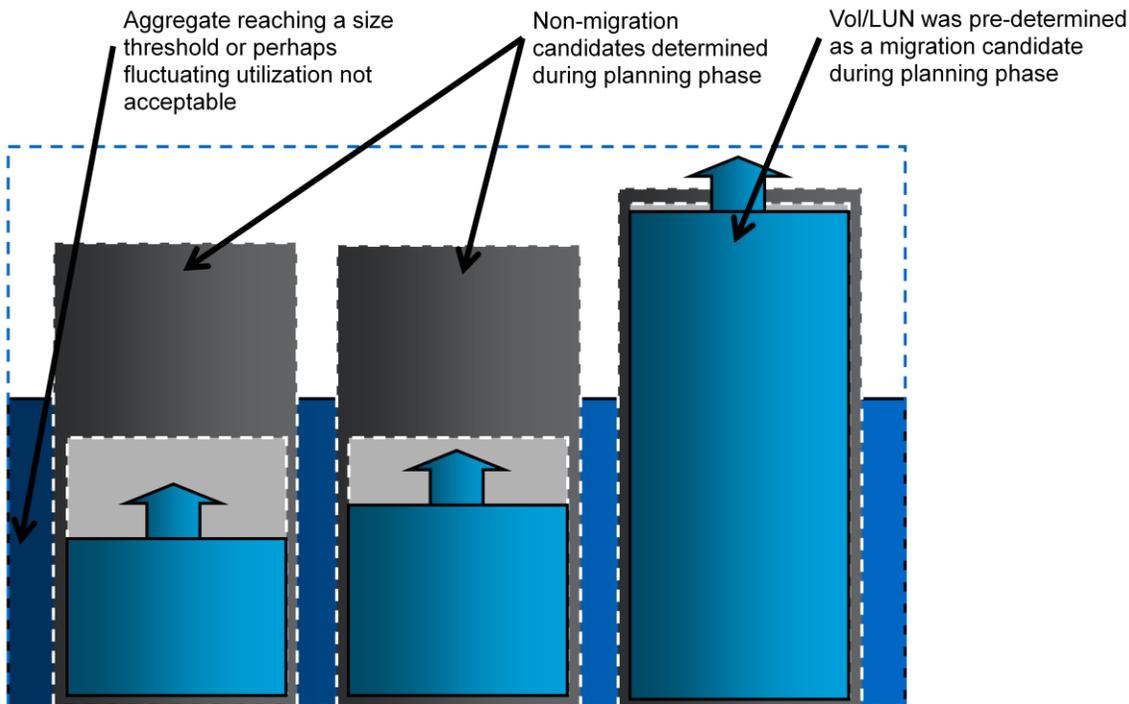
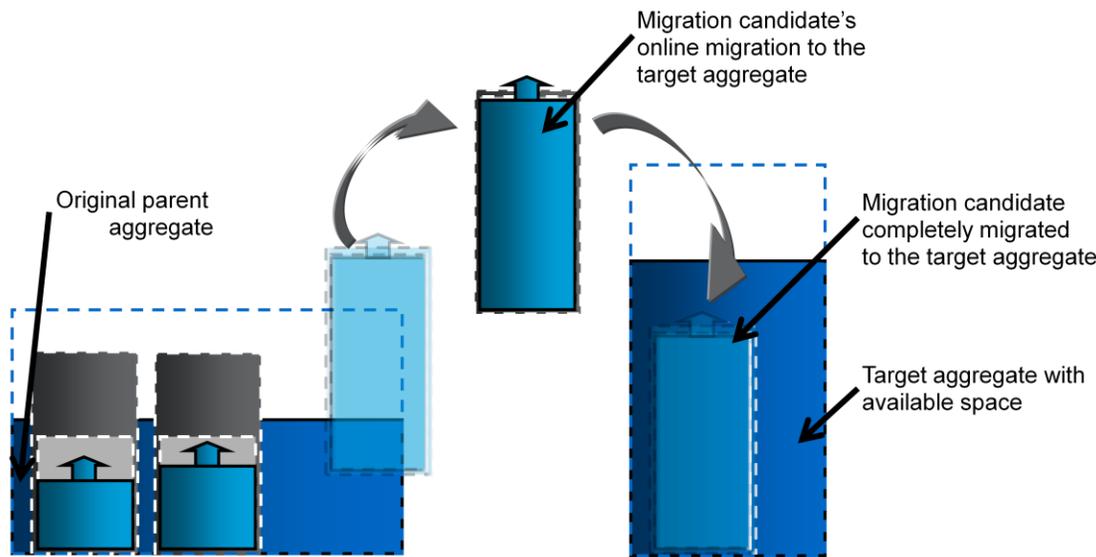


Figure 17) Example mitigation of migrating a migration-candidate volume/LUN to a different aggregate.



PERFORMANCE TEST CASE: THIN VERSUS THICK PROVISIONING

Thin provisioning can be implemented while provisioning new storage or it can be enabled seamlessly on storage that is already in production, without affecting performance. Internal testing and customer experiences have verified similar performance when comparing thin provisioning and traditional thick provisioning.

The following test scenario exemplifies the type of performance similarities typically experienced when comparing thin provisioning and thick provisioning. The test environment in this case used Microsoft[®] Exchange, and followed the thin provisioning best practices as described in the Exchange section later in this document.

The following are some testing specifics:

- The Exchange workload was simulated using Jetstress 2010
- Number of mailboxes (users): 1,000
- Mailbox size: 200MB
- IOPS/Mailbox: 1 (heavy load)
- Run time: 2 hours

The results confirmed that there is no significant performance degradation between thin provisioning and thick provisioning for an Exchange workload. The average IOPS performance degradation was only 3% between thin and thick provisioning configurations on the storage system.

PERFORMANCE IMPACT

Thin provisioning does not cause any major performance impact on applications.

4 THIN PROVISIONING AND VMWARE

Testing of thin provisioning with VMware is in progress. Test results and best practices will be provided in a later revision of this document. In the meantime, please see the documents in the “References” section below.

REFERENCES

For additional information please refer to the following:

- NetApp and VMware vSphere Storage Best Practices
<http://media.netapp.com/documents/tr-3749.pdf>
- NetApp and VMware Virtual Infrastructure 3 Storage Best Practices
<http://media.netapp.com/documents/tr-3428.pdf>

5 THIN PROVISIONING AND SHAREPOINT

This section provides guidance for using thin provisioning with Microsoft SharePoint. The objective of the information is to show where thin provisioning can be considered within this solution and key considerations. The guidance is based on results collected from customer experiences and internal testing.

5.1 USE-CASE SPECIFICS

The guidance in this section was developed with the following application and OS versions.

- Microsoft SharePoint Server 2010
 - Remote BLOB storage
- Microsoft SQL Server 2008 R2
- Microsoft Server 2008 R2 x64
- NetApp System Manager 1.1

- NetApp SnapDrive for Windows 6.3
- NetApp SnapManager for SharePoint 6.0
- NetApp Data ONTAP 7.3.3

Some adjustments to the following recommendations may be needed if you are using different versions of the applications and operating systems.

5.2 BEST PRACTICES

This section provides best practices and considerations specific to using NetApp thin provisioning with Microsoft SharePoint. If you are new to thin provisioning, NetApp highly recommends that NetApp or a NetApp authorized partner be contacted for assistance. Table 13, below, shows the recommended NetApp thin provisioning parameter settings.

CONFIGURING DATA ONTAP FOR SHAREPOINT

Table 13) Data ONTAP recommended volume and LUN thin provisioning parameters for SharePoint.

Parameter	Value
Aggregate	
volume grow area	yes
Volume Options	
guarantee	none
fractional_reserve	0%*
autosize	on
autosize options	-m Xg* -l Yg*
try_first	volume_grow
Volume Snapshot Options	
reserve	Z%*
autodelete	off
autodelete options	NA
LUN Options	
space reservation	disabled

Note: Although Microsoft SQL is part of the SharePoint solution, recommendations regarding the SQL Server database are not included in this section. Instead, they are available in the “Thin Provisioning and SQL” section of this document.

IMPLEMENTING SHAREPOINT

- Thin provisioning is recommended for the SharePoint RBS Archived Data Store (RBS CIFS share).
- Thin provisioning is not recommended for the RBS Archived Index Store. Since the Index Store does not require much storage space, it should not pose any special challenges to fully provision storage for the Index.
- NetApp best-practice recommendations suggest using a NetApp LUN for the Archived Index/Extender LUN and an SMB (CIFS) share for the Archived Data device. For additional details, see [TR-3877: SharePoint Server 2010 and SnapManager 6.0 for SharePoint Best Practices Guide](#).

* This value is specific to the environment, and is to be determined as part of the sizing process.

- NetApp recommends a CIFS share for the BLOB, but if LUN thin provisioning is used for the BLOB, it is required that the Quick Format option be used during the creation of the NTFS LUN.
- By default, System Manager 1.1 creates LUNs using thick provisioning. When System Manager has been used to create a LUN, the space-reserved properties of the LUN will need to be changed to “none” (thin provisioning) as a separate step.
- By default, System Manager 1.1 creates all NAS volumes with autodelete set to OFF and autogrow set to OFF. Values will need to be changed as a separate step.
- By default, System Manager 1.1 creates all SAN volumes with autodelete set to ON and autogrow set to ON. Values will need to be changed as a separate step.

MONITORING SHAREPOINT

- By default, SNMP alerts are not configured. SNMP alerts (events and messages) must be configured as a separate step for Operations Manager and other systems.
- Appropriate monitoring, accompanied by an effective action plan, is required when using thin provisioning to avoid application downtime.
- Make sure to set alerts and thresholds so that the proper people are contacted with enough time to react as storage is filled.
- Make sure to set the SnapManager for SharePoint alerts for any job monitoring. For example, if any jobs related to the Archiver or Extender fail due to not enough space in any CIFS shares, SNMP needs to send alerts to the administrator.
Note: The alerts in SnapManager for SharePoint 2010 differ from those in SnapManager for SharePoint 6. Make sure the appropriate documentation and best practices are referenced during implementation.
- Be careful about monitoring the space consumed in the SQL Log LUNs. If the transaction logs are not truncated regularly, then the Log files can grow very large, and you need to increase the space in the Log LUNs.

GROWING STORAGE FOR SHAREPOINT

- If the Archived Datastore runs out of space, uploading files to SharePoint fails.
 - The following error message is generated:
"The URL 'Shared Documents/test.doc' is invalid. It may refer to a nonexistent file or folder, or refer to a valid file or folder that is not in the current Web."
 - The existing files in the SharePoint farm will still be available for read operations.
 - Once the space of the Archived Datastore is increased, the upload of new files will once again be possible without needing any additional steps.
 - NetApp also recommends setting up multiple CIFS shares for the BLOB storage device, so that if one CIFS share fills up, then SnapManager for SharePoint can move onto the next CIFS share without disrupting the document upload process.
 - CIFS Share: If the size of the volume containing the CIFS share is increased using the Data ONTAP CLI or System Manager GUI, the additional space is automatically recognized and usable by the server.
- NetApp SnapDrive for Windows (SDW) can be used to make more space available for the LUN.
 - NTFS LUN: If the size of the LUN is increased using the CLI or GUI, the additional space is not automatically recognized. Use SnapDrive for Windows (SDW) to increase the size of the LUN. Using SDW allows the additional space to be automatically recognized and usable by the server.
Note: Microsoft recommends dismounting the LUN before increasing the storage space. For details, refer to the KB article <http://support.microsoft.com/kb/959613>.
- For volume autogrowing to work, it is mandatory that the containing aggregate has enough space (at least 1.2 times the volume size).

SPECIAL CONSIDERATIONS FOR USING THIN PROVISIONING WITH SHAREPOINT

- When storage consumption is unpredictable or highly volatile, it is best to reduce the level of storage overcommitment so that storage is available for any growth spikes. Consider limiting storage commitment to 100%—no overcommitment—and using the trending functionality to determine how much overcommitment is acceptable, if any.
- Overcommitment of storage must be carefully considered and managed for mission-critical applications in which even a minimal outage is not tolerable. In such a case, it is best to monitor storage consumption trends to determine how much overcommitment is acceptable, if any.
- If the time required to procure new storage is very long, adjust storage overcommitment thresholds accordingly. The overcommitment threshold should alert administrators early enough to allow new storage to be procured and installed.

REFERENCES

For additional information please refer to the following:

- SharePoint Server 2010 and SnapManager 6.0 for SharePoint Best Practices Guide:
<http://media.netapp.com/documents/tr-3887.pdf>

6 THIN PROVISIONING AND SQL

This section provides guidance for using thin provisioning with Microsoft SQL. The objective of the information is to show where thin provisioning can be considered within this solution, and key considerations. The guidance is based on results collected from customer experiences and internal testing.

6.1 USE-CASE SPECIFICS

The guidance in this section was developed with the following application and OS versions.

- Microsoft SQL Server 2008 R2
 - Single Instance Database
 - OLTP
- Microsoft Server 2008 R2 x64
- NetApp System Manager 1.1
- NetApp SnapDrive for Windows 6.3
- NetApp SnapManager for SQL 6.0
- NetApp Data ONTAP 7.3.3

Some adjustments to the following recommendations may be needed if you are using different versions of the applications and operating systems.

6.2 BEST PRACTICES

This section provides best practices and considerations specific to using NetApp thin provisioning with Microsoft SQL. If you are new to thin provisioning, NetApp highly recommends that NetApp or a NetApp authorized partner be contacted for assistance. Table 14, below, shows the recommended NetApp thin provisioning parameter settings.

CONFIGURING DATA ONTAP FOR SQL

Table 14) Data ONTAP recommended volume and LUN thin provisioning parameters for SQL.

Parameter	Value
Aggregate	
volume grow area	yes
Volume Options	
guarantee	none
fractional_reserve	0%*
autosize	on
autosize options	-m Xg* -l Yg*
try_first	volume_grow
Volume Snapshot Options	
reserve	Z%*
autodelete	off
autodelete options	NA
LUN Options	
space reservation	disabled

IMPLEMENTING SQL

- Thin provisioning is recommended for SQL Server database files.
- Thick provisioning is recommended for all other files and file types.
 - System databases, such as master, model, msdb, and tempdb
 - User database logs
- Use separate volumes for each SQL database as a best practice for SnapManager backup and recovery. For additional details, see [TR-3768: SnapManager 5.0 for SQL Server: Best Practices Guide](#).
- When using LUN thin provisioning with Windows, the Quick Format option must be used during the creation of the NTFS LUN.

* This value is specific to the environment, and is to be determined as part of the sizing process.

- By default, the sizing tool within SnapManager for SQL assumes the usage of the default thick provisioning settings. You do not have to use those default settings, although doing so is a common practice. For additional details, see Appendix C of [SnapManager 2.1 for Microsoft SQL Server](#).
- For SQL, NetApp recommends that the snapshot autodelete functionality in SnapManager for SQL be used, and not the Data ONTAP autodelete feature.
- By default, System Manager 1.1 creates LUNs using thick provisioning. When the System Manager has been used to create a LUN, the space-reserved properties of the LUN will need to be changed to "none" (thin provisioning) as a separate step.
- By default, System Manager 1.1 creates all NAS volumes with autodelete set to OFF and autosize set to OFF. Values will need to be changed as a separate step.
- By default, System Manager 1.1 creates all SAN volumes with autodelete set to ON and autosize set to ON. Values will need to be changed as a separate step.
- When using SQL with volume thin provisioning, SQL Server Instant file initialization must be enabled using the "perform volume maintenance task" credentials under the local security policy. For additional details, see Microsoft article MS175935 at <http://msdn.microsoft.com/en-us/library/ms175935.aspx>.
- Both autodelete and autosize work at the volume level, and not on individual LUNs. This means that LUNs will not automatically grow, and must be handled separately with different commands. Consider the following:
 - Autosize may fail due to space constraints in the aggregate, and must be properly monitored using Operations Manager.
 - NetApp does not recommend autodelete because automatically deleting snapshots is typically not an accepted practice in critical environments. If you do not have a critical environment, autodelete could be used to make space available as a second resort when autosize is not possible.
 - Autosize and autodelete are typically used in combination in place of setting a LUN overwrite reserve. As discussed in a previous section, the LUN overwrite reserve preallocates space for a very safe approach, but greatly increases the space requirements for a reserve that is only used as a last resort.

MONITORING SQL

- By default, SNMP alerts are not configured. SNMP alerts (events and messages) must be configured as a separate step for Operations Manager and other systems.
- Appropriate monitoring, accompanied by an effective action plan, is required when using thin provisioning to avoid application downtime.
- Make sure to set alerts and thresholds so that the proper people are contacted with enough time to react as storage is filled.

GROWING STORAGE FOR SQL

- NetApp SnapDrive for Windows (SDW) can be used to make more space available for the LUN.
 - CIFS Share: If the size of the volume containing the CIFS share is increased using the CLI or GUI, the additional space is automatically recognized and usable by the server.
 - NTFS LUN: If the size of the LUN is increased using the CLI or GUI, the additional space is not automatically recognized. Use SnapDrive for Windows to increase the size of the LUN. Using SDW allows the additional space to be automatically recognized and usable by the server.
- **Note:** Microsoft recommends dismounting the LUN before increasing the storage space. For details, refer to the KB article <http://support.microsoft.com/kb/959613>.
- For volume autosizing to work, the containing aggregate must have enough space (at least 1.2 times the volume size).

SPACE RECLAMATION WITH SQL

- Space reclamation must be initiated from time to time to recover the unused space in a LUN. NetApp recommends that the space first be reclaimed at the database level using SQL Server Management studio → Shrink database/files option, then at the storage level using the SnapDrive → Start space Reclaimer option.

SPECIAL CONSIDERATIONS FOR USING THIN PROVISIONING WITH SQL

- When storage consumption is unpredictable or highly volatile, it is best to reduce the level of storage overcommitment so that storage is available for any growth spikes. Consider limiting storage commitment to 100%—no overcommitment—and using the trending functionality to determine how much overcommitment is acceptable, if any.
- Overcommitment of storage must be carefully considered and managed for mission-critical applications in which even a minimal outage is not tolerable. In such a case, it is best to monitor storage consumption trends to determine how much overcommitment is acceptable, if any.
- If the time required to procure new storage is very long, adjust storage overcommitment thresholds accordingly. The overcommitment threshold should alert administrators early enough to allow new storage to be procured and installed.

REFERENCES

For additional information please refer to the following:

- SnapManager 5.0 for SQL Server: Best Practices Guide:
<http://media.netapp.com/documents/tr-3768.pdf>
- Best Practice Guide for Microsoft SQL Server on NetApp Storage:
<http://media.netapp.com/documents/tr-3821.pdf>

7 THIN PROVISIONING AND EXCHANGE

This section provides guidance for using thin provisioning with Exchange. The objective of the information is to show where thin provisioning can be considered within this solution, and key considerations. The guidance is based on results collected from customer experiences and internal testing.

7.1 USE-CASE SPECIFICS

The guidance in this section was developed with the following application and OS versions.

- Exchange Server 2010
- Microsoft Server 2008 R2 x64
- NetApp System Manager 1.1
- NetApp SnapDrive for Windows 6.3
- NetApp SnapManager for Exchange 6.0
- NetApp Data ONTAP 7.3.3

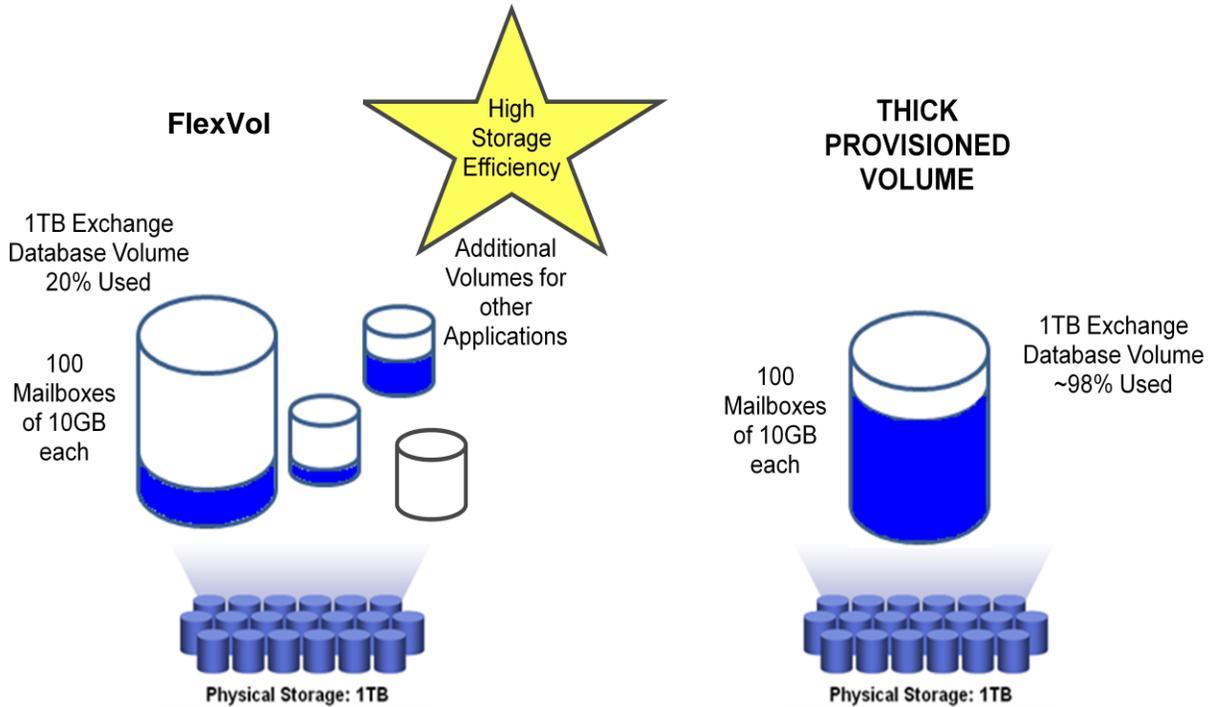
Some adjustments to the following recommendations may be needed if you are using different versions of the applications and operating systems.

MIGRATION TO EXCHANGE SERVER 2010

Significant space savings can be achieved using thin provisioning during the migration to Exchange Server 2010. The default mailbox size on Exchange 2010 is now a larger 10GB, compared to 2GB on

earlier versions. The larger mailbox size will contribute to rapid storage growth as mailboxes are created. Thin provisioning can be used to greatly reduce the amount of storage that will be preallocated for mailboxes, and instead allow the storage to be allocated as e-mails are stored. This avoids the creation of lots of empty mailbox space, instead allocating and using storage on demand, increasing storage efficiency. Figure 18 displays storage allocation sizes for thin provisioning and thick provisioning in an Exchange 2010 solution.

Figure 18) Compare thin and thick provisioning for Exchange 2010 with 100 mailboxes.

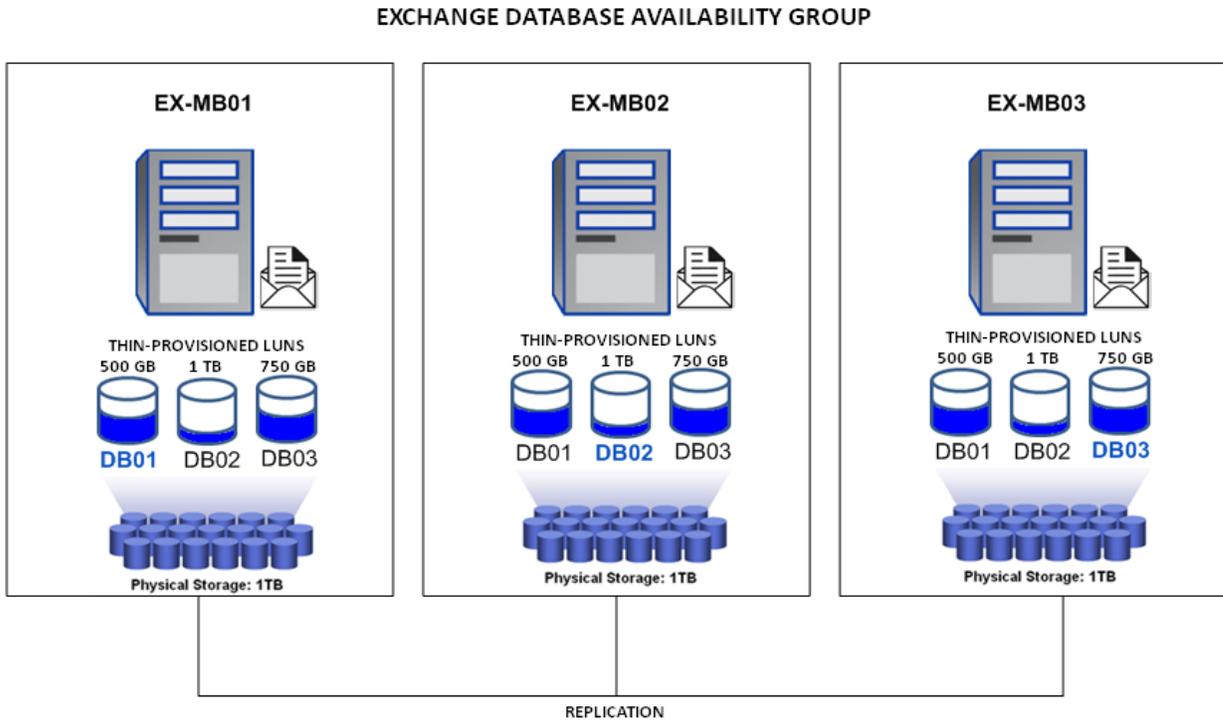


DATABASE AVAILABILITY GROUP (DAG)

When replacing the server and data resiliency options available in earlier versions of Exchange, Microsoft implemented the Database Availability Group (DAG) feature. A DAG consists of 2 or more mailbox servers; however, there can be no more than 16 servers in the DAG. Each mailbox server can hold one or more active or passive copies of the database. Each database maintains separate status, so one server can host copies of multiple databases and have just some of those copies active at one time.

Figure 19 shows a solution with multiple DAGs. There are three servers and three copies of each database, one on each server. The active database copy for each server is designated with a blue font. The passive copies of each database are updated by way of log shipping. Using NetApp thin provisioning with the Exchange databases will help reduce storage consumption at multiple replication sites by removing the preallocation of blocks at the time of creation, delaying storage allocation until data is written to the copies.

Figure 19) Storage thin provisioning for Exchange 2010 DAG.



7.2 BEST PRACTICES

This section provides best practices and considerations specific to using NetApp thin provisioning with Exchange. If you are new to thin provisioning, NetApp highly recommends that NetApp or a NetApp authorized partner be contacted for assistance. Table 19, below, shows the recommended NetApp thin provisioning parameter settings.

CONFIGURING DATA ONTAP FOR EXCHANGE

Table 15) Data ONTAP recommended volume and LUN thin provisioning parameters for Exchange.

Parameter	Value
Aggregate	
volume grow area	yes
Volume Options	
guarantee	none
fractional_reserve	0%*
autosize	on
autosize options	-m Xg* -l Yg*
try_first	volume_grow
Volume Snapshot Options	
reserve	0%
autodelete	off
autodelete options	NA
LUN Options	
space reservation	disabled

IMPLEMENTING EXCHANGE

- Thin provisioning is recommended for both Exchange logs and the Exchange database volumes.
- NetApp deduplication typically provides storage savings between 10% and 30% for Exchange 2010. Deduplication can be an effective way to make space available when needed in an Exchange solution using thin provisioning. For additional details, see [Storage Efficiency and Best Practices for Microsoft Exchange Server 2010](#).
- By default, the sizing tool within SnapManager for Exchange assumes the usage of the default thick-provisioning settings. You do not have to use those default settings, although they often are used as a common practice.
- For optimal performance, NetApp strongly recommends that the thin-provisioned volumes be spanned across multiple disks. This substantially decreases the bottlenecks at the spindle level.
- When autodelete is used, NetApp recommends that the Snapshot autodelete functionality in SnapManager for Exchange be used, and not the Data ONTAP autodelete feature.

* This value is specific to the environment, and is to be determined as part of the sizing process.

MONITORING EXCHANGE

- By default, SNMP alerts are not configured. SNMP alerts (events and messages) must be configured as a separate step for Operations Manager and other systems.
- Appropriate monitoring, accompanied by an effective action plan, is required when using thin provisioning to avoid application downtime.
- Make sure to set alerts and thresholds so that the proper people are contacted with enough time to react as storage is filled.

GROWING STORAGE FOR EXCHANGE

- Use NetApp SnapDrive for Windows (SDW) to make growable space available for the LUN.
 - NTFS LUN: If the size of the LUN is increased using the CLI or GUI, the additional space is not automatically recognized. Use SnapDrive for Windows (SDW) to increase the size of the LUN. Using SDW allows the additional space to be automatically recognized and usable by the server.
- Note:** Microsoft recommends dismounting the LUN before increasing the storage space. For details, refer to the KB article <http://support.microsoft.com/kb/959613>.

SPACE RECLAMATION WITH EXCHANGE

- Space reclamation must be initiated from time to time to recover the unused space in a LUN. Storage space can be reclaimed at the storage level using the SnapDrive -> Start Space Reclaimer option.

SPECIAL CONSIDERATIONS FOR USING THIN PROVISIONING WITH EXCHANGE

- When storage consumption is unpredictable or highly volatile, it is best to reduce the level of storage overcommitment so that storage is available for any growth spikes. Consider limiting storage commitment to 100%—no overcommitment—and using the trending functionality to determine how much overcommitment is acceptable, if any.
- Overcommitment of storage must be carefully considered and managed for mission-critical applications in which even a minimal outage is not tolerable. In such a case, it is best to monitor storage consumption trends to determine how much overcommitment is acceptable, if any.
- If the time required to procure new storage is very long, storage overcommitment thresholds should be adjusted accordingly. The overcommitment threshold should alert administrators early enough to allow new storage to be procured and installed.

REFERENCES

For additional information please refer to the following:

- Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment:
<http://media.netapp.com/documents/tr-3483.pdf>
- Storage Efficiency and Best Practices for Microsoft Exchange Server 2010:
<http://media.netapp.com/documents/tr-3824.pdf>
- Maximize Storage Efficiency with NetApp Thin Provisioning and Symantec Thin Reclamation:
<http://media.netapp.com/documents/wp-7111.pdf>
- SnapDrive 6.2 for Windows Best Practices:
<http://media.netapp.com/documents/tr-3828.pdf>

8 THIN PROVISIONING AND ORACLE

This section provides guidance for using thin provisioning with Oracle. The objective of the information is to show where thin provisioning can be considered within this solution, and key considerations. The guidance is based on results collected from customer experiences and internal testing.

8.1 USE-CASE SPECIFICS

The guidance in this section was developed with the following application and OS versions.

- Oracle
 - Single Instance Database
 - OLTP 11g
- NetApp System Manager 1.1
- NetApp SnapDrive for Windows 6.3
- NetApp SnapManager for Oracle 6.0
- NetApp Data ONTAP 7.3.3

Some adjustments to the following recommendations may be needed if you are using different versions of the applications and operating systems.

8.2 BEST PRACTICES

This section provides best practices and considerations specific to using NetApp thin provisioning with Oracle. If you are new to thin provisioning, NetApp highly recommends that NetApp or a NetApp authorized partner be contacted for assistance. Table 16, below, shows the recommended NetApp thin provisioning parameter settings.

CONFIGURING DATA ONTAP FOR ORACLE

Table 16) Data ONTAP recommended volume and LUN thin provisioning parameters for Oracle.

Parameter	Value
Aggregate	
volume grow area	yes
Volume Options	
guarantee	none
fractional_reserve	0%*
autosize	on
autosize options	-m Xg* -l Yg*
try_first	volume_grow
Volume Snapshot Options	
reserve	0%
autodelete	off
autodelete options	NA
LUN Options	
space reservation	disabled

IMPLEMENTING ORACLE

- Thin provisioning is recommended for Oracle datafile volumes (.dbf file types).
- Thick provisioning is recommended for all other (nondatafile) file types.
 - Log files: Oracle records all transactions in the form of redo buffers, and flushes them to the disk from time to time in the form of redo logs. These logs are extremely important for the database and should be available without any interruption.
- Oracle Automatic Storage Management (ASM) provides excellent synergy with thin provisioning at the storage level. It creates an ASM header in each of the LUNs used in the ASM group. However, preinitialization is not performed in any of the storage space.
- In a scenario in which the free space of an ASM group is exhausted, the group can be expanded by just adding more ASM disks, followed by automatic rebalancing of the previously existing ASM data.

* This value is specific to the environment, and is to be determined as part of the sizing process.

- As a best practice during Oracle implementations, it is advisable to have a database and all of its associated files under a single volume. Maintaining separate volumes for each database makes backup extremely convenient.
- Oracle requires that you explicitly set the autoextend option to ON while creating tablespaces residing on thin-provisioned volumes.
- When overprovisioning, take note of the following:
 - As the space is occupied within a volume using thin provisioning, the other overprovisioned volumes within the same parent aggregate will only see the remaining unallocated space available in the aggregate. This is important where NFS is used since the application will see much less space than provisioned because the space has already been allocated to other volumes.
- By default, System Manager 1.1 creates LUNs using thick provisioning. When System Manager has been used to create a LUN, the space-reserved properties of the LUN will need to be changed to “none” (thin provisioning) as a separate step.
- By default, System Manager 1.1 creates all NAS volumes with autodelete set to OFF and autosize set to OFF. Values will need to be changed as a separate step.
- By default, System Manager 1.1 creates all SAN volumes with autodelete set to ON and autosize set to ON. Values will need to be changed as a separate step.
- By default, System Manager 1.1 uses a Snapshot reservation of 20% for all NAS volumes. The Snapshot reserve option can be unchecked if desired when considering best practices. For additional details, see TR-3633: “NetApp Best Practice Guidelines for Oracle Database 11g.”
- Both autodelete and autosize work at the volume level, and not on individual LUNs. This means that LUNs will not automatically grow, and must be handled separately with different commands. Consider the following:
 - Autosize may fail due to space constraints in the aggregate, and must be properly monitored using Operations Manager.
 - NetApp does not recommend autodelete because automatically deleting Snapshot copies is typically not an accepted practice in critical environments. If this is not the case, autodelete could be used to make space available as a second resort when autosize is not possible.
 - Autosize and autodelete are typically used in combination in place of setting a LUN overwrite reserve. As discussed in a previous section, the LUN overwrite reserve preallocates space for a very safe approach, but greatly increases the space requirements for a reserve that is only used as a last resort.

MONITORING ORACLE

- By default, SNMP alerts are not configured. SNMP alerts (events and messages) must be configured as a separate step for Operations Manager and other systems.
- Appropriate monitoring, accompanied by an effective action plan, is required when using thin provisioning to avoid application downtime.
- Set alerts and thresholds so that the proper people are contacted with enough time to react as storage is filled.

SPECIAL CONSIDERATIONS FOR USING THIN PROVISIONING WITH ORACLE

- When storage consumption is unpredictable or highly volatile, it is best to reduce the level of storage overcommitment so that storage is available for any growth spikes. Consider limiting storage commitment to 100%—no overcommitment—and using the trending functionality to determine how much overcommitment is acceptable, if any.

- Overcommitment of storage must be carefully considered and managed for mission-critical applications in which even a minimal outage is not tolerable. In such a case, it is best to monitor storage consumption trends to determine how much overcommitment is acceptable, if any.
- If the time required to procure new storage is very long, storage overcommitment thresholds should be adjusted accordingly. The overcommitment threshold should alert administrators early enough to allow new storage to be procured and installed.

REFERENCES

For additional information please refer to the following:

- TR-3633: NetApp Best Practice Guidelines for Oracle Database 11g
Contact your local NetApp account team or NetApp authorized partner for access to this NetApp technical report.
- SnapManager 3.0 for Oracle Best Practices:
<http://media.netapp.com/documents/tr-3761.pdf>

9 THIN PROVISIONING WITH OTHER NETAPP FEATURES

This section provides some pointers for using thin provisioning with other NetApp features.

9.1 THIN PROVISIONING AND MANAGEMENT/MONITORING TOOLS

If thin provisioning is to be employed, it is imperative that the free space on the storage system be well monitored. This is necessary so that additional space can be added quickly enough to meet the application's demands.

Table 17 shows which tools can be used for storage provisioning. Some of the listed tools are not full management tools, but were included for completeness since they can be used for some management processes.

Table 17) Provisioning functionality of NetApp tools.

	Ontap CLI	FilerView	System Manager	SnapDrive for Windows	SnapDrive for Unix	Provisioning Manager
Configure Data ONTAP	✓					
Create Aggregate	✓	✓	✓			
Create Volume	✓	✓	✓			✓
Configure Volume	✓		✓			✓
Create Qtree	✓	✓	✓			✓
Create LUN	✓	✓	✓	✓	✓	✓
Configure LUN	✓	✓	✓	✓	✓	✓
Map LUN	✓	✓	✓	✓	✓	✓

Table 18 provides an overview of how each tool will affect current settings if used for storage provisioning. It is important to note that the (x) with yellow background represents settings that will be returned to a default value if configured with the designated tools. It is a good practice to check the final settings of the parameters when storage provisioning is complete.

Table 18) NetApp tools' effect on key settings.

(x) = will be changed by the tool to a fixed setting

		Ontap CLI	FilerView	System Manager	SnapDrive for Windows	SnapDrive for Unix	Provisioning Manager
System	autodelete thresholds	x					
Volume	guarantee	x	x	x			x
Volume	fractional_reserve	x		x			(x)
Volume	autosize	x		x			(x)
Volume	autosize options	x					(x)
Volume	try_first	x		(x)			(x)
Snapshot	reserve	x	x	(x)			(x)
Snapshot	autodelete	x		(x)			x
Snapshot	autodelete options	x					(x)
LUN	space reservation	x	x	x	x	x	x

9.2 THIN PROVISIONING AND DEDUPLICATION

- Deduplication works seamlessly with thin provisioning for both LUNs and volumes.
- Deduplication savings will make more space available within the LUN, volume, or aggregate, depending on the configuration.
- Deduplication savings may not survive a data migration scenario, depending on the type of data migration that is used.

For additional details, see [TR-3505: NetApp Deduplication Deployment and Implementation Guide](#).

9.3 THIN PROVISIONING AND FLEXCLONE VOLUMES

- FlexClone files are not automatically space reserved after creation. No matter what the space reservation on the source file is, the clone file has no space reservation. To enable space reservation on a file clone, use the `file reservation` command.
- FlexClone LUNs inherit space guarantee settings from the source. If there is not enough space in a volume to create a clone with the same space guarantee as the source, the cloning process will fail. Note that the source and clone LUN will share blocks on the disk even with space guarantee enabled.
- At clone creation, Data ONTAP creates metadata for the new instance of the data. It allocates data for storing changes to the cloned copy or new data on request. Thus, the overcommitment of the aggregate containing the cloned data increases when creating the clone. However, this does not affect the space used in the aggregate. When data in the clone is rendered and new data is added by the application, the aggregate use will grow.

9.4 THIN PROVISIONING AND V-SERIES

- Thin provisioning is supported with V-Series.
- V-Series systems can be used as a source or target for thin provisioning mitigation activities in storage tightness scenarios.

9.5 THIN PROVISIONING AND SNAPMIRROR REPLICATION

BEST PRACTICES

Best practices must be used when using thin provisioning in a disaster recovery (DR) configuration so that the destination site is able to keep up with the automatic storage adjustments as they occur at the primary site.

- If volumes are being mirrored with SnapMirror and the primary volume grows in size, the SnapMirror relationship will fail if the target volume is not at least as large as the primary. The following approach enables the DR destination to be at least as large as the primary, even when the primary destination is automatically grown using `autosize` or by manual process:
 - Consider primary volumes with space guarantee set to “volume” and with `autosize` enabled.
 - At the DR destination site, create your target volumes in an aggregate that is the same size as the aggregate of the DR primary site.
 - Set the DR destination guarantee to “none,” and set all the destination volume sizes to just smaller than the size of the aggregate (that is, if your aggregate is 10TB, create all of your target volumes with size 9.9TB). The total size of the volumes can be larger than the containing aggregate.
 - DO NOT turn off `FS_SIZE_FIXED` at the destination.
 - Now the DR target volumes are using thin provisioning, with a max size equal to the destination aggregate size or, more precisely, very close to it. This allows the automatic growing of source volumes using `autosize` while preventing the source volumes from exceeding the max size of the target volumes.
 - Keep in mind that beginning with Data ONTAP 7.3, when a DR failover occurs and the SnapMirror relationship is broken, the target volumes will maintain their original characteristics, and not inherit the characteristics of the source volumes (that is, space guarantee = none will persist through the failover). Therefore, if the destination will continue to be used as the primary for a significant amount of time, it may be desirable to change the volume settings of the new primary volumes to the standard settings used for primary storage systems (that is, change the space guarantee of the new primary from “none” to “volume” in this example). With Data ONTAP 7.2, a SnapMirror break would cause the destination volume to inherit the volume guarantee of the source volume, and should be planned for appropriately.
 - Later, when the original primary system is brought back online, the volume settings should be reviewed to determine that they are set appropriately at the DR primary and destination sites. For instance, in this example, determine that the DR primary volumes have the space guarantee set to “volume” and that the DR destination volumes have the space guarantee set to “none,” as originally configured. These settings should be reviewed, regardless of whether the original primary system is once again used as the primary site, or the new primary system is left to run as the primary site.
 - If resync is not expected between the DR primary and the DR destination and volume autosizing is desired, set `FS_SIZE_FIXED` to off.

For additional details, see [TR-3446: SnapMirror Async Overview and Best Practices Guide](#). If using the synchronous or semi-synchronous mode of SnapMirror, see [TR-3326: SnapMirror Sync and SnapMirror Semi-Sync Overview and Design Considerations](#).

9.6 THIN PROVISIONING AND SNAPVAULT

At the writing of this document, the best practice for a SnapVault secondary is to set space guarantees of the volume to “volume.”

9.7 THIN PROVISIONING AND MULTISTORE (VFILER)

- MultiStore lends itself very nicely to the thin provisioning mitigation activities.
- Beginning with Data ONTAP 7.3.3, DataMotion for vFiler is available for both FAS and V-Series systems, and provides online migration capabilities.
- MultiStore also provides offline data migration, which can be used in cases in which the environment can handle small amounts of downtime.

BEST PRACTICES

- MultiStore vFiler units are not required for nonmigration candidate volumes, but the vFiler units must be used for migration candidates if they are to use the MultiStore online migration capability.

9.8 THIN PROVISIONING AND DATAMOTION FOR LUNS

- DataMotion for LUNs is available beginning with Data ONTAP 8.1 operating in 7-Mode, and can be used as a mitigation activity to address a storage tightness scenario.

10 TROUBLESHOOTING

This section provides some initial guidance to help troubleshoot some high-level issues.

10.1 NOT SEEING STORAGE SAVINGS

The most convenient tools to use for viewing storage savings attributed to thin provisioning are My AutoSupport™ and the Operations Manager Storage Efficiency Dashboard.

10.2 NOT GETTING ALERTS

Check that the alerts are properly configured in Operations Manager. Remember that you can set the alarms at both the default level and for each individual aggregate. Also check that the correct e-mail address, pager number, or SNMP trap host name has been entered in Operations Manager. If an e-mail alias is being used, make sure it has not been modified. If the SNMP trap host hostname is being used, try using the IP address in case the name is not being properly resolved.

10.3 PERFORMANCE SLOWDOWN

There should not be any significant performance effects due to using NetApp thin provisioning. If there is a performance slowdown, try to isolate the issue. If additional assistance is required, please contact NetApp Technical Support. The contact information is provided later in this document.

11 ADDITIONAL READING AND REFERENCES

- Operations Manager Administration Guide
http://now.netapp.com/NOW/knowledge/docs/DFM_win/rel40/html/software/opsmgr/frameset.html
- Operations Manager Efficiency Dashboard Installation and User Guide
http://now.netapp.com/NOW/download/tools/omsed_plugin/InstallUserGuide.pdf

- Performance Study of VMware vStorage Thin Provisioning
http://www.vmware.com/pdf/vsp_4_thinprov_perf.pdf
- RA-0007: Storage Efficiency Every Day: How to Achieve and Manage Best-in-Class Storage Use
<http://media.netapp.com/documents/RA-0007.pdf>
- SnapManager 2.1 for Microsoft SQL Server
http://now.netapp.com/NOW/knowledge/docs/SnapManager/relsmsql_21/html/software/admin/cfgappl9.htm
- TR-3263: WORM Storage on Magnetic Disks Using SnapLock Compliance and SnapLock Enterprise
<http://media.netapp.com/documents/tr-3263.pdf>
- TR-3326: SnapMirror Sync and SnapMirror Semi-Sync Overview and Design Considerations
<http://media.netapp.com/documents/tr-3326.pdf>
- TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices
<http://media.netapp.com/documents/tr-3428.pdf>
- TR-3440: Operations Manager, Protection Manager, and Provisioning Manager Sizing Guide
<http://media.netapp.com/documents/tr-3440.pdf>
- TR-3446: SnapMirror Async Best Practices Guide
<http://media.netapp.com/documents/tr-3446.pdf>
- TR-3450: Active-Active Controller Configuration Overview and Best Practice Guidelines
<http://media.netapp.com/documents/tr-3450.pdf>
- TR-3466: Open Systems SnapVault Best Practice Guide
<http://media.netapp.com/documents/tr-3466.pdf>
- TR-3483: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment
<http://media.netapp.com/documents/tr-3483.pdf>
- TR-3487: SnapVault Best Practice Guide
<http://media.netapp.com/documents/tr-3487.pdf>
- TR-3548: MetroCluster Design and Implementation
<http://media.netapp.com/documents/tr-3548.pdf>
- TR-3563: NetApp Thin Provisioning Increases Storage Utilization with on-Demand Allocation
<http://media.netapp.com/documents/tr-3563.pdf>
- TR-3584: Microsoft Exchange 2007 Disaster Recovery Model Using NetApp Solutions
<http://media.netapp.com/documents/tr-3584.pdf>
- TR-3608: Data ONTAP SNMP Trap Management
<http://media.netapp.com/documents/tr-3608.pdf>
- TR-3651: Microsoft Exchange 2007 SP1 Continuous Replication Best Practices Guide
<http://media.netapp.com/documents/tr-3651.pdf>
- TR-3694: NetApp and Citrix XenServer 4.1: Building a Virtual Infrastructure from Server to Storage
<http://media.netapp.com/documents/tr-3694.pdf>
- TR-3702: NetApp Storage Best Practices for Microsoft Virtualization
<http://media.netapp.com/documents/tr-3702.pdf>
- TR-3705: NetApp and VMware VDI Best Practices
<http://media.netapp.com/documents/tr-3705.pdf>
- TR-3712: Oracle VM and NetApp Storage Best Practices Guide

- <http://media.netapp.com/documents/tr-3712.pdf>
- TR-3729: Simplified SAN Provisioning and Improved Space Utilization Using NetApp Provisioning Manager
<http://www.netapp.com/us/library/technical-reports/tr-3729.html>
- TR-3742: Using FlexClone to Clone Files and LUNs
<http://media.netapp.com/documents/tr-3742.pdf>
- TR-3747: Best Practices for File System Alignment in Virtual Environments
<http://media.netapp.com/documents/tr-3747.pdf>
- TR-3749: NetApp and VMware vSphere Storage Best Practices
<http://media.netapp.com/documents/tr-3749.pdf>
- TR-3770: 2,000-Seat VMware View on NetApp Deployment Guide Using NFS: Cisco Nexus Infrastructure
<http://media.netapp.com/documents/tr-3770.pdf>
- TR-3814: DataMotion Best Practices
<http://media.netapp.com/documents/tr-3814.pdf>
- TR-3824: Storage Efficiency and Best Practices for Microsoft Exchange Server 2010
<http://media.netapp.com/documents/tr-3824.pdf>
- TR-3827: Storage Efficiency—If You Are Doing This, Then Your Storage Could Be Underutilized
<http://www.netapp.com/us/library/technical-reports/tr-3827.html>
- TR-3828: SnapDrive 6.2 for Windows Best Practices
<http://media.netapp.com/documents/tr-3828.pdf>
- TR-3873: NetApp DataMotion for Volumes
<http://www.netapp.com/us/library/technical-reports/tr-3873.html>
- TR-3881: DataMotion for Volumes for Enterprise Applications
<http://media.netapp.com/documents/tr-3881.pdf>
- TR-3886: Understanding and Using vStorage APIs for Array Integration with NetApp Storage
<http://media.netapp.com/documents/tr-3886.pdf>
- WP-7053: The 50% Virtualization Guarantee* Program Technical Guide
<http://media.netapp.com/documents/wp-7053.pdf>
- WP-7111: Maximize Storage Efficiency with NetApp Thin Provisioning and Symantec Thin Reclamation
<http://media.netapp.com/documents/wp-7111.pdf>

12 ADDITIONAL ASSISTANCE

For additional support, contact one of the following:

- Your local account team
- Systems engineer
- Account manager
- [NetApp Global Services](#)
- [NOW](#)[®] (NetApp on the Web)
- 888.4.NETAPP (United States and Canada)
- 00.800.44.NETAPP (EMEA/Europe)

13 VERSION TRACKING

Version 1 September 2011 Initial release

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.



www.netapp.com

© 2011 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, AutoSupport, DataMotion, Data ONTAP, FilerView, FlexClone, FlexVol, MetroCluster, MultiStore, NOW, SANscreen, SnapDrive, SnapLock, SnapManager, SnapMirror, Snapshot, SnapVault, SyncMirror, vFiler, and WAFL are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Windows, Microsoft, SQL Server, and SharePoint are registered trademarks of Microsoft Corporation. VMware is a registered trademark and vCenter, vSphere, and VMotion are trademarks of VMware, Inc. Oracle is a registered trademark and Oracle11g is a trademark of Oracle Corporation. Symantec is a trademark of Symantec Corporation. UNIX is a registered trademark of The Open Group. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-3965-0911