NETAPP WHITE PAPER

The 50% Virtualization Guarantee Program Technical Guide

Jack McLeod, NetApp
September 2008 |WP-7053-0908

Table of Contents

1 INTRODUCTION

With the cost of doing business rising every day, organizations are looking for ways to reduce overhead and consolidate resources into a single, easy to manage environment. With the amount of information and data doubling yearly in many organizations, storage has become a key component to doing business in today's marketplace. With the coming of the "virtualization age," the need and demand for storage has only increased, with virtualization condensing servers into nothing more than large files that act in every way like the physical servers they are based on. NetApp understands the challenges that organizations face and has designed business solutions to address the need to do more with less in a virtual environment. By using a combination of RAID-DP®, storage deduplication, thin provisioning, and Snapshot™ technologies, NetApp enables customers to achieve storage savings in a VMware® virtual infrastructure.

2 THE PROGRAM

NetApp is offering the 50% Virtualization Guarantee Program ("Program"), which is comprised of 50% space savings to NetApp customers, both direct and channel, who purchase certain new NetApp products and services, as described in this document and use such products and services in their VMware environments ("Customers"). If a Customer who participates in the Program achieves less than the space savings after complying with the terms of the Program Guide and Acknowledgement For 50% Virtualization Guarantee Program ("Program Guide'), best practices outlined in this document, TR 3428, and TR3505, NetApp will provide consulting and if needed additional capacity as described in the Program Guide.

3 PROGRAM REQUIREMENTS

The Program depends on the following conditions being met in a Customer's VMware environment.

3.1 BASELINE COMPARISON BASIS

Under the Program Customers will be able to fit their data requirement in the amount of raw storage they purchase from NetApp. Here's a definition of the terms used in this Program.

- Customer Data Requirement is defined as the amount of data the Customer needs to store in a system
- Storage Baseline is the amount (TB) of raw storage the Customer would have to purchase on a non-NetApp storage comparable performance and protection in order to accommodate Customer Data Requirement.
- The Storage Savings Target is 50% of the Storage Baseline. The Customer can meet the stated Customer Data Requirement by purchasing storage from NetApp at the Storage Savings Target

.

3.2 ELIGIBILITY REQUIREMENTS

To participate in the Program, a Customer must meet all of the following eligibility requirements.

- The Program shall apply only to the new purchase of the following NetApp Products to be used for primary storage only:
    - FAS2020, FAS2050
    - FAS3020, FAS3040, FAS3050, FAS3070, FAS3140, FAS3170
    - FAS6030, FAS6040, FAS6070, FAS6080
    - The Program excludes all NetApp® V-Series, S line, and VTL systems.
    - This Program is not applicable to the N series offered by IBM.

- The new storage system must use one or more of the following protocols: FC, iSCSI, or NFS
- The new storage system must be running Data ONTAP® 7.3 or later. Systems running Data ONTAP 10 are excluded.
- The NetApp FAS system under the Program must have at least 14 drives for the storage of data for virtual servers or virtual desktops.
- The Customer must agree to have the following features enabled on their NetApp storage system.
    - AutoSupport (ASUP)
    - RAID-DP
    - Thin provisioning without LUN reservation
    - Deduplication
    - NetApp Snapshot
- The Customer must follow the NetApp best practices described in the following technical reports:
    - This document
    - TR-3428: NetApp and VMware VI3 Storage Best Practices
    - TR-3505: Deduplication Implementation and Best Practices

Customer must comply with all of the requirements of the Program Guide

- Must purchase a minimum level of NetApp Professional Services (PS) deployment and implementation services, specified by the following table:

*Z[A-D] denotes the corresponding zone where purchased:
- ZA = Americas and US Public Sector
- ZB = EMEA
- ZC = APAC
- ZD = Japan

	NAS	SAN
SupportEdge Premium Purchased	o PS-STOR-IMPL-BASE-RDSP-Z[A-D]*, (Base Storage implementation service) o PS-VIRT-IMPL-SOW, (Virtualization Implementation Service)	o PS-STOR-IMPL-BASE-RDSP-Z[A-D]*, (Base Storage implementation service) o PS-VIRT-IMPL-SOW, (Virtualization Implementation Service) o PS-STOR-IMPL-SAN-RDSP-Z[A-D]*, (Base SAN Storage implementation
SupportEdge Premium Not Purchased	o Installation service from NetApp, or Customer may self-install o PS-STOR-IMPL-BASE-RDSP-Z[A-D]*, (Base Storage Implementation service) o PS-VIRT-IMPL-SOW, (Virtualization Implementation Service)	o Installation service from NetApp, or Customer may self-install o PS-STOR-IMPL-BASE-RDSP-Z[A-D]*, (Base Storage implementation service) o PS-VIRT-IMPL-SOW, (Virtualization Implementation Service) o PS-STOR-IMPL-SAN-RDSP-Z[A-D]*, (Base SAN Storage implementation

- No more than 10% of the following will be covered under the Program
    - Images and graphics,
    - XML data
    - Database data

    - Microsoft® Exchange data
    - Encrypted data
    - Scientific data
- Large database and Exchange deployments are excluded from this program.
- A deployment must have at least 10 similar virtual machines with identical OS versions per flexible volume.
- Workloads with high performance requirements that require spindles may be excluded. This will be determined by the Sales Engineer (SE) or by Professional Services (PS) during sizing.
- Only VMware data on a NetApp FAS system that meets these requirements is included in this Program. Other data that is not VMware specific that may coexist on the FAS system is not included in this Program.
- Placing data that is not covered under this program onto a separate aggregate excludes that aggregate from the program but allows any aggregate that solely contains data that is included in the Program to be covered.

Even though a Customer may not be eligible for the Program, substantial savings in a VMware environment can still be achieved. Therefore NetApp recommends that Customers follow best practices, whether they participate in the program or not.

4 THE NETAPP TECHNOLOGY THAT ALLOWS NETAPP TO OFFER THE PROGRAM

4.1 RAID-DP

RAID-DP provides performance that is comparable to that of RAID 10 with pricing comparable to RAID 4, with much higher resiliency than either. It provides protection against double disk failure as compared to RAID 5, which can only protect against one disk failure. Due to increased reliability and decreased cost compared to similar environments, RAID-DP offers businesses a compelling total cost of ownership storage option without putting their data at increased risk. For more information about RAID-DP, see TR-3298: RAID-DP: NetApp Implementation of RAID Double Parity for Data Protection.

4.2 THIN PROVISIONING

Thin provisioning is a method of storage virtualization that allows storage administrators to address and oversubscribe storage in the same way that server resources such as RAM and CPU are provisioned in a VMware environment, providing a level of “storage on demand.” Thin-provisioned storage is treated as a shared resource pool and is consumed only as each individual VM needs it. This sharing increases the total utilization rate of storage by eliminating the unused but provisioned areas of storage that are associated with traditional storage. By allowing as-needed provisioning and space reclamation, thin provisioning can result in better storage utilization and smaller capital expenditures on storage infrastructure.

For this Program, volume-level thin provisioning will be implemented. With the volume space guarantee set to None, Data ONTAP enables the administrator to create a large flexible volume (or multiple flexible volumes), which can be greater than the aggregates size. Space is allocated to a volume from the aggregate only when data is written to the volume. Additionally, the administrator can create a smaller volume with Data ONTAP and later resize the volume to accommodate growing data.

These approaches to provisioning and managing storage are extremely efficient because they start with less capacity than will eventually be needed and grow that capacity, manually or automatically, as needed. For example, the administrator could create home directories of size 100GB for 100 users, using a storage pool of 2TB available; without thin provisioning capabilities the administrator would require 10TB of storage space to satisfy the requirement.

For details about thin provisioning, see TR-3563: NetApp Thin Provisioning.

4.3 DEDUPLICATION

With NetApp FAS deduplication, VMware deployments can eliminate the duplicate data in their environment, enabling greater storage utilization. NetApp deduplication provides this space saving on primary storage by removing redundant copies of blocks within a volume. This process is transparent to the application and can be enabled or disabled on the fly. Deduplication virtualization technology enables multiple virtual machines to share the same physical blocks in a NetApp FAS system in the same manner that VMs share system memory. It can be seamlessly introduced into a virtual infrastructure without having to make any changes to VMware administration, practices, or tasks. Deduplication runs on the NetApp FAS system at scheduled intervals and does not consume any CPU cycles on the ESX Server. For more information about NetApp deduplication technology, see TR-3505: NetApp Deduplication for FAS, Deployment and Implementation Guide.

4.4 SNAPSHOT COPIES AND FLEXCLONE VOLUMES

A Snapshot copy is a frozen, read-only image of a traditional volume, a FlexVol® volume, or an aggregate that captures the state of the file system at a point in time. Snapshot copies are your first line of defense to back up and restore data. NetApp Snapshot technology can easily be integrated into VMware environments, where it provides crash-consistent versions of virtual machines for the purpose of full VM recovery, full VM cloning, or site replication and disaster recovery. This is the only snapshot technology that does not have a negative impact on system performance. VMware states that for optimum performance and scalability, hardware-based snapshot technology is preferred over software-based solutions.

A FlexClone® volume is a writable point-in-time image of a FlexVol volume or another FlexClone volume and is based on NetApp Snapshot technology. FlexClone volumes add a new level of agility and efficiency to storage operations. They take only a few seconds to create and are created without interrupting access to the parent FlexVol volume. FlexClone volumes use space very efficiently, leveraging the Data ONTAP architecture to store only data that changes between the parent and clone. The use of the FlexClone technology in a VMware environment offers significant savings in dollars, space, and energy. In addition to all these benefits, FlexClone volumes have the same high performance as other FlexVol volumes. For more information about FlexClone technology, see TR-3447: FlexClone Volumes: A Thorough Introduction.

5 CONCLUSION

By employing RAID-DP, Snapshot copies, deduplication, and thin provisioning in a FAS system, NetApp can offer Customers a 50% reduction in space in a VMware environment. NetApp is confident that Customers who follow the guidelines in this document will achieve the same high degree of success with our storage in a VMware environment that many other Customers have achieved.

6 CHECKLIST: REQUIREMENTS FOR VALIDATION OF VIRTUALIZATION CAPACITY

6.1 BEST PRACTICES

The best practices outlined in these technical reports must be followed by the Customer:

- TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices
- TR-3505: NetApp Deduplication for FAS Deployment and Implementation Guide

Note: Many of the items listed in this section are included in these best practices. If there are conflicts between these items and the technical reports, the items listed here take precedence.

6.2 DATA ONTAP VERSION

Controllers must be running Data ONTAP 7.3 or later. Data ONTAP 10 and related releases are excluded.

Note: NetApp reserves the right, as part of this Program, to require Customers to upgrade to a higher level of Data ONTAP to obtain any necessary bug fixes.

6.3 DISK DRIVES

Each controller must have a minimum of 14 drives for the storage of data for virtual servers or virtual desktops and use only full shelves. All configurations must also have enough capacity to include one additional shelf for future expansion.

A controller that is the primary owner of a set of drives can allocate up to the following number of spares (per drive type).

Number of Drives Per Controller	Number of Spares
14 - 27 drives	1 spare
28 - 100 drives	2 spares
Additional 84 drives	1 additional spare

Example: A controller with 184 drives (1 drive type) can have up to 3 spares.

Note: When using internal drives several external shelves may be required to meet data drive to parity drive ratios.

6.4 AGGREGATES

- Only aggregates or flexible volumes (no traditional volumes) can be used.
- All RAID groups must use RAID-DP:
  - Minimum 12 data drives for every 2 parity drives.
  - **Exception:** The minimum spare drives can be counted as data drives for this calculation. For example, a 28-drive controller with 2 spares and 2 aggregates (11 data drives: 2 parity drives) would be acceptable.
- No use of multiple plexes (SyncMirror®).

- **Note:** The root volume must be a flexible volume that shares an aggregate with other flexible volumes.

6.5 VOLUMES AND LUNS

- Maximum volume size:
    - No volume's physical size may exceed the maximum physical volume size (with deduplication) for the system type.
    - No volume's physical size may exceed 25% of the maximum data in a volume (with deduplication) for the system type.
    - See the Data Protection: Online Backup and Recovery Guide, "Deduplication," Table 10: Maximum Volume Size.
- Thin provisioning must be enabled with no use space reservations. Please also set the fractional space reserve to 0 for a LUN in an iSCSI and FC environment.
- The use of synchronous SnapMirror® deduplication is not supported.
- Don't use qtree SnapMirror or SnapVault®. Volume SnapMirror is the preferred means of replication for data that has been deduplicated.
- MetroCluster is not supported for this Program.

6.6 VOLUME CONTENT

- At least ten (10) Virtual machines running identical guest OS versions must be included in the same flexible volume.
    - Most VMs should be clustered into high-density groups of 10 or more with a identical guest OS in the same flexible volume.
    - Other VMs should be grouped to maximize the number of VMs with similar guest OSs in the same flexible volume.
    - At least 80% of the VMs must be in these high-density groups.
- All .vmdk files must be created at the proper alignments. See TR-3428, NetApp and VMware VI3 Storage Best Practices, section 4.1, for details.
- Volumes can be used only for VMware virtual machines. No other use, including other types of virtual machines, is permitted with this offer.
- Application data and home directories that are part of these virtual machines can be included in these volumes, provided that any data that doesn't deduplicate well—such as unique images, databases, and so on—doesn't exceed 10% of the total data in the aggregate.
- Placing data that is not covered under this program onto a separate aggregate excludes that aggregate from the program but allows any aggregate that solely contains data that is included in the program to be covered.

6.7 SNAPSHOT AND DEDUPLICATION MANAGEMENT

- Deduplication must be run to completion before taking any long-term Snapshot copies. Either:
    - Snapshot schedules should be turned off
- Snapshot data:
    - Snap reserve must be set to 20% or less.
    - Each volume should have 20% or less data in its Snapshot copies.
- Before validating space utilization, deduplication must be run successfully on each of the volumes within the past 2 to 3 days, or after a large data migration, if it has occurred more recently.
- In the aggregate containing Program data follow these guidelines:
    - Aggregate snapshots should not be taken
    - Snap reserve must be set to 0

7 APPENDIX: CALCULATIONS FOR VALIDATING SPACE SAVINGS

This appendix documents how to calculate the projected capacity of NetApp storage that is using deduplication.

For the Program:

- If the average percentage of deduplication for the FAS system is >= 50%, there is no need to do these calculations. The savings targets have already been exceeded.
- If the average percentage deduplication for the FAS system is < 50%, the following calculations will identify how much data the system is on track to store ("Projected Total Usable Capacity"). This number can be compared with the Customer Data Requirement as outlined in this document. If the numbers calculated here meet or exceed Customer Data Requirement, the Storage Savings Target has been met.

This section explains how the calculations are organized and how to perform them.

7.1 PROJECTED CAPACITY

Projected capacity is the amount of data that a NetApp FAS system is expected to hold before reaching maximum capacity, assuming that future storage savings match current trends. The following equations work from the top down to show how the projected capacity is calculated from its constituent parts.

The projected capacity of a FAS system is the sum of projected capacity of each aggregate:

(eq 1) `FAS System Projected Capacity = Sum (Aggregate Projected Capacity)`

The projected capacity of an aggregate is an extrapolation based on the amount of logical data that is stored in the aggregate and how much physical space it took to store this data:

(eq 2) `Aggregate Projected Capacity = Aggr Logical Data / Aggregate %Full`

The following subsections describe how to find these component values.

7.2 AGGREGATE %FULL

The raw data needed to calculate the aggregate %Full number can be found in the `aggr_show_space` command:

```
Filer> aggr show_space
Aggregate 'aggr0'

    Total space    WAFL reserve    Snap reserve    Usable space       BSR NVLOG
   1531673088KB     153167308KB      68925288KB    1309580492KB             0KB

Space allocated to volumes in the aggregate

Volume                          Allocated            Used       Guarantee
vol0                          113047212KB        361820KB          volume
ang_fcp                       106092896KB      76826496KB          volume
ang_nfs                       104933400KB      75543972KB          volume
ang_fcp_d                     105397104KB       9155220KB          volume
ang_nfs_d                     104961188KB       8716492KB          volume

Aggregate                       Allocated            Used           Avail
Total space                   534431800KB     170604000KB     776722260KB
Snap reserve                   68925288KB       1057636KB      67867652KB
WAFL reserve                  153167308KB        63620KB    153103688KB
```

First, to find the used space in an aggregate, add the three numbers in the Used column in the Aggregate section:

(eq 3) `Used Aggr Space = Total Space (Used) + Snap Res.(Used) + WAFL Res. (Used)`

Example: `Used Aggr Space = 170604000 + 1057636 + 63620 = 171725256`

The aggregate's %Full is then calculated as the fraction of used space versus the total capacity of the aggregate:

(eq 4) `Aggregate %Full = Used Aggregate Space / Total Space`

Example: `Aggregate %Full = 171725256 / 1531673088  = 11.2%`

7.3 LOGICAL DATA STORED

The logical data stored by a FAS system includes both data in the active file system and the snapshots:

(eq 5) `Aggr Logical Data = Active FS Logical Data + Snapshot Logical Data`

7.4 LOGICAL DATA [ACTIVE FILE SYSTEM]

(eq 6) `Active FS Logical Data = Sum (Vol Logical Data)`

To find the amount of logical data stored in the active file system, use the `df -s` command:

```
f32*> df -s
Filesystem               used      saved       %saved
/vol/vol0/             290632          0           0%
/vol/ang_fcp/        75591200          0           0%
/vol/ang_nfs/        75468172          0           0%
/vol/ang_fcp_d/       8615716   67612200          89%
/vol/ang_nfs_d/       8612904   67482728          89%
```

To calculate the logical data stored in a given volume, add the amount of space used to the space saved:

(eq 7) `Vol Logical Data = Used + Saved`

Example: `Vol Logical Data (ang_nfs_d) = 8612904 + 67482728 = 76095632`

7.5 LOGICAL DATA (SNAPSHOTS)

(eq 8) `Snapshot Logical Data = Sum (Vol Snapshot Logical Data)`

To estimate the amount of logical data in snapshots, combine the information in the `df -s` command with `df` commands:

```
Filer> df
Filesystem               kbytes       used      avail capacity  Mounted on
/vol/vol0/             90390400     290632   90099768       0%  /vol/vol0/
/vol/vol0/.snapshot    22597600      12076   22585524       0%
/vol/vol0/.snapshot
/vol/ang_fcp/         104857600   75591200   29266400      72%  /vol/ang_fcp/
```

```
/vol/ang_fcp/.snapshot          0          0          0     ---%
/vol/ang_fcp/.snapshot
/vol/ang_nfs/           104857600   75468172   29389428      72%  /vol/ang_nfs/
/vol/ang_nfs/.snapshot          0          0          0     ---%
/vol/ang_nfs/.snapshot
/vol/ang_fcp_d/         104857600    8615716   96241884       8%
/vol/ang_fcp_d/
/vol/ang_fcp_d/.snapshot        0          0          0     ---%
/vol/ang_fcp_d/.snapshot
/vol/ang_nfs_d/         104857600    8612904   96244696       8%
/vol/ang_nfs_d/
/vol/ang_nfs_d/.snapshot        0          0          0     ---%
/vol/ang_nfs_d/.snapshot
```

Using .snapshot information from `df` and the %saved information from `df -s`, do the following:

(eq 9) `Vol Snapshot Logical Data = Snapshot Used + Snapshot Saved`

(eq 10) `Snapshot Saved = Snapshot Used * Vol %Saved / (1 – Vol %Saved)`

Example: `Snapshot Save (vol0) = 12076 * 0% / (1 – 0%) = 0`

`Vol Snapshot Logical Data (vol0) = 12076 + 0 = 12076`

7.6 RECAP

To find the system's projected capacity (from the bottom up):

1. **Logical Data**
   a. Find Logical Data stored in each volume's snapshots (eqs 9 and 10).
      i. **Example:** Snapshot Logical Data (vol0) = 12076
   b. Find the Logical Data stored in each volume's active file system (eq 7).
      i. **Example:** Logical Data (ang_nfs_d) = 76095632
   c. Sum the data for the active file system (eq 6) and the snapshots (eq 8) together (eq 5).
      Example:
      1. Active File System Logical Data = 303673552
      2. Snapshot Logical Data = 12076
      3. Total Logical Data = 303685628

2. **Aggregate %Full**
   a. Find the Used Space in the aggregate (eq 3).
      Example: Used Space = 171725256
   b. Calculate the Aggregate %Full.
      Example: %Full = 171725256 / 1531673088 = 11.2%

3. **Find the Aggregate Projected Capacity**
   Divide the aggregate's logical data by its %Full (eq 2).
   Example: = 303685628 / 11.2% = 2711478821
   Note: you can divide the result by (1024 * 1024) to show the results in GBs: 2583.19.

4. **Find the System Projected Capacity**
   If you have more than one aggregate, add the results of each aggregate (eq 1).

   Note: you can divide the GB results by 1024 to show the final results in TBs: 2.52.

The guarantee and the related Program is limited to the terms set forth in [LINK], and is applicable only to prospective orders placed after the Program effective date and is dependent upon Customer's compliance with the terms and conditions set forth in this document and any of the instruction sets and specification set forth in the referenced documents. NetApp's sole and exclusive liability and Customer's sole and exclusive remedy associated with the terms of the guarantee and related Program is the provision by NetApp of the additional storage capacity as set forth in the Program.

www.netapp.com