



Technical Report

If You're Doing This, Then Your Storage Could Be Underutilized

John Tyrrell, NetApp
April 2010 | TR-3827

TABLE OF CONTENTS

IF YOU'RE DOING THIS, THEN YOUR STORAGE COULD BE UNDERUTILIZED 1
 JOHN TYRRELL, NETAPP APRIL 2010 | TR-38271

1 INTRODUCTION 3

2 STORAGE UTILIZATION – STORAGE EFFICIENCY TERMINOLOGY 3
 2.1 STORAGE UTILIZATION5

3 STORAGE POOLS 5
 3.1 UTILITY MODEL STORAGE POOL5
 3.2 IF YOU'RE DOING THIS, THEN YOU COULD BE UNDERUTILIZED!7

4 NETAPP MANAGEMENT CONSTRUCTS AND SOFTWARE 9
 4.1 NETAPP MANAGEMENT CONSTRUCTS9
 4.2 NETAPP MANAGEMENT SOFTWARE.....12

5 GAINING HIGH STORAGE UTILIZATION WITH NETAPP® 13
 5.1 MAXIMIZE THE NETAPP ARRAY STORAGE13
 5.2 MAXIMIZE THE “FREE CAPACITY (# 7) IN THE RESOURCE POOL.....13
 5.3 DEFINING RESOURCE POOLS14
 5.4 DEFINING MANAGEMENT POLICIES AND STORAGE SERVICES15
 5.5 SETTING DFM GLOBAL OPTIONS16
 5.6 PROVISIONING LUNS AND VOLUMES16
 5.7 MONITORING THE RESOURCE POOL.....17
 5.8 CHARGEBACK CONSIDERATIONS17

6 NETAPP STORAGE EFFICIENCY TERMINOLOGY DEFINITIONS AND CALCULATIONS18

1 INTRODUCTION

The NetApp® storage system is a truly unique array in the industry with a number of innovative concepts, functionality, and storage objects. The first thing to realize is that it is not a simple block array. It is really a sophisticated storage system front-ended with a proprietary operating system called Data ONTAP® (Data “on tap”) and a file system called WAFL® (“Waffle” - Write Anywhere File Layout). This unique architecture allows NetApp the ability to offer some clever technology for functionality, flexibility, and significant storage efficiencies within customer data centers.

Because of this architecture, NetApp® brings a robust type of “Unified Storage” to the industry. The following list is a set of key technology that a combination of NetApp “Unified Storage” and NetApp management software brings to the table:

- **Multiple Protocols** – FC SAN, FCoE SAN, IP SAN (iSCSI), NAS (CIFS® and/or NFS® File Access) – all in the same storage system
- **Multiple Transports** – Fibre channel, Ethernet (up to 10 GbE), FCoE
- **Storage Efficiency** – Deduplication, Cloning, Snapshots™, etc.
- **Storage Utilization** – Thin Provisioning, Provisioning Manager, Resource Pools
- **Efficient Data Protection** – RAID-DP™, Snapshot, SnapRestore, SnapMirror, SnapLock, SnapVault, SnapValidator, MetroCluster
- **Unique Storage Objects** – FlexVols, Aggregates, FlexClones, qtrees, LUN/File clones, Snapshots, Multi-Store (vfilers)
- **Performance** – Aggregate Striping, FlexCache, FlexShare, HA Clusters, Storage Acceleration Appliances
- **Multi-Tenancy** – Multi-Store
- **Common APIs** – Commands, APIs, SMI-S, CLIs, and SNMP that apply to entire family of arrays

2 STORAGE UTILIZATION – STORAGE EFFICIENCY TERMINOLOGY

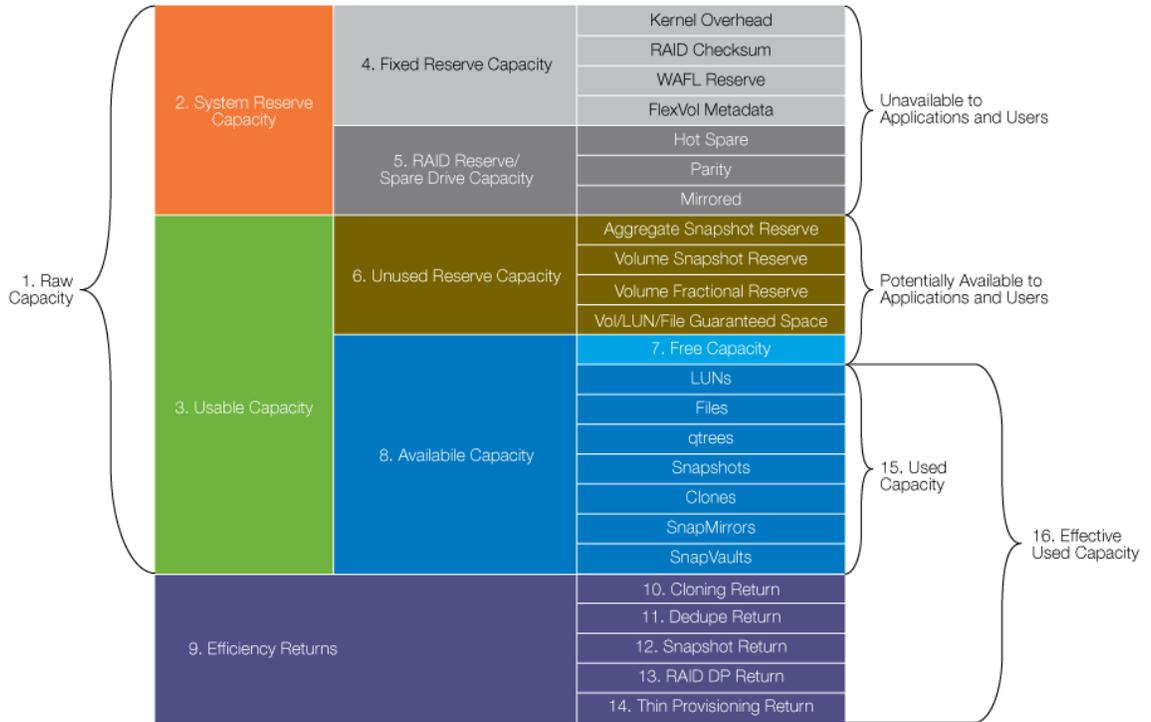
This paper will focus on how the combination of NetApp’s unified storage technology and management software can be used to help reduce not only the capital expenditures for storage but also the operating expenditures associated with the management of that storage. There are two aspects to these savings – Storage Utilization and Storage Efficiency.

Storage Utilization answers the question, “Am I effectively using the storage capacity available to my applications. The picture in Figure 1) Storage Utilization and Storage Efficiency Breakdown below shows the entire category breakdown of NetApp storage. All of the terminology and calculations for this figure are described in section 6 on page 18.

The “Raw Capacity” (# 1) represents all of the storage installed in your data center. Not all of the raw storage in your shop is available for application data. There are certain overheads with every vendor’s storage that you have no control over. This would include things like disk right-sizing, boot kernel, checksums, etc. This is shown as “Fixed Reserve Capacity” (# 4). Then, there is a set of storage where you have some control over like RAID configuration with both parity disks and hot spares (shown as “RAID Reserve/Spare Drive Capacity” (# 5)). Collectively, these two categories are called the “System Reserve Capacity” (# 2). When you remove the “System Reserve Capacity” (# 2), you are left with potentially usable storage for application data, called the “Usable Capacity” (# 3).

The “Unused Reserve Capacity” (# 6) in yellow represents areas of either “guaranteed” or “reserved” space. Every storage vendor has their own example of this category (e.g., pre-carved LUNs). This space is potential free space to all applications but, if left as reserved or guaranteed, this space becomes used but contains no real data. The “Free Capacity” (# 7) is space that is currently unassigned and available to all applications. Finally, the “Used Capacity” (# 15) is the storage that actually contains actual user data.

Storage Utilization (# 17) is calculated by dividing the “Used Capacity” (# 15) by the “Usable Capacity” (# 3). To calculate the Raw Storage Utilization (# 18), divide the “Used Capacity” (# 15) by the “Raw Capacity” (# 1).



Storage Utilization Calculations

- 17. Storage Utilization = Used Capacity / Usable Capacity
- 18. Raw Storage Utilization = Used Capacity / Raw Capacity

Storage Efficiency Calculations

- 19. Storage Efficiency = Effective Used Capacity / Usable Capacity
- 20. Raw Storage Efficiency = Effective Used Capacity / Raw Capacity

Figure 1) Storage Utilization and Storage Efficiency Breakdown

Storage Efficiency answers the question, “How much more real storage would I have needed to install without some key storage saving technologies like deduplication, cloning, snapshots, etc.?” For example, if you have a 50 GB volume or LUN and you create 20 clones of that object using NetApp technology, the clones basically take up zero space. The “Cloning Return” (# 10) of 1 TB (20 x 50 GB volume) represents the additional storage you would have needed without the NetApp cloning technology (FlexClones®, LUN clones, and File clones). Similarly, suppose you deduped a set of volumes resulting in 9500 GB of “Dedupe Returns” (# 11), that is 9.5 TB of additional storage you would have needed without that technology.

There are several NetApp Storage Efficiency technologies such as deduplication, cloning, snapshots, etc. that allow customers to avoid the procurement of additional storage. The “Efficiency Returns” (# 9) represents the sum of those savings. The combination of the “Efficiency Returns” (# 9) and the “Used Capacity” (# 15) is called the Effective Used Capacity (# 16).

The Storage Efficiency (# 19) is then calculated by dividing the Effective Used Capacity (# 16) by the “Usable Capacity” (# 3). To calculate the Raw Storage Efficiency (# 20), divide the Effective Used Capacity (# 16) by the “Raw Capacity” (# 1).

Obviously, both Storage Utilization and Storage Efficiency are extremely important components in reducing your storage costs and the goal should be to make both percentages as high as possible. This white paper will concentrate solely on the Storage Utilization aspect. The Storage Efficiency side of the equation is completely covered in an existing NetApp white paper (WP-7022-1208, “The NetApp Storage Efficiency Guide”, Larry Freeman, December, 2008).

2.1 STORAGE UTILIZATION

Most of the open systems world of storage is sitting at very low Storage Utilization numbers, some as low as 30%. This means that only 30% of the true storage available to applications actually contains real data. Many data centers would probably deny this. Regardless of what we think, it is important to actually determine your true utilization. It is where you really are and you can then decide if change is necessary. To compute a true utilization, you have to determine how much real data (not reserved, carved up or guaranteed) is there, and divide by the total real data capacity available.

If you carve up an array into LUNs, let's say 50 GB each, your storage reports will show 50 GB in use for any of those LUNs. But when you look at the host view where this LUN was mounted as a Windows “D” drive, it could show that only 30 GB is in use, say by Oracle. When you look into the view from Oracle, maybe only 15 GB is in use. In this example, you are only 30% utilized (15 GB/50 GB).

In the example above, the extra 35 GB contains no application data, thus it is not part of the “Used Capacity” (# 15). It is, however, allocated to that LUN in case more data gets written but it is not available to any other user of the system. Any area that is allocated as guaranteed space or reserved space to any container is called the “Unused Reserve Capacity” (# 6). Why do we highlight this – because this space could be significant (and usually is) and it represents potential free space in the storage pool for all users.

What if you could free up 70% of your storage? Is that really possible? Running at 100% Storage Utilization would be great but probably not practical unless you are not growing at all. But how much free space do you need to run your business every day? You certainly don't need 70%, probably more like 15% to 20% depending on your growth rate and size of your data. Section 3 will address how you can greatly improve your storage utilization. Before we describe how to improve your Storage Utilization with NetApp, it is important to discuss the concept of storage pools and some of the current techniques you may be using on your existing storage resources.

3 STORAGE POOLS

There are a number of definitions for Storage Pools. The “Utility Model Storage Pool” presented in this paper is by far the most space efficient type in the storage industry and has proven to be so for almost three decades now. This is the model that this paper will promote and describe how to implement it using NetApp technology.

3.1 UTILITY MODEL STORAGE POOL

We have had an interesting history of storage in the past 25 years from the heyday of the mainframe to open systems with direct attached storage to SAN/NAS networked storage to complete virtualization of both storage and servers and now cloud storage, but we may have dropped some important lessons along the way. In these difficult economic times, we are struggling to minimize not only the capital expenses for storage but also the ongoing operational costs of storage in both the people management costs and the environmental costs as well. Doing it with less is best for both of these.

Let's take a quick look at the storage lessons learned on the mainframe. At one point in the 1980s, users and groups were given individual volumes to put their data on. Back in those days, the Storage Utilization was about 30%, right about where we are today.

If I use laptops as an example, each of us has our own hard drive storage. My friend, Charlie, is almost completely out of space. I have plenty but he can't use it. In fact, the 1000 employees with the same laptops collectively have enough free space for an additional 400 employee laptops. The management cost is high in that every individual has to monitor his/her own space and, when they run out, it is painful. Of course, laptops are what they are and it is not possible to share the total hard drive storage as a single pool.

In the mainframe environment, all of the storage was connected to a JES-plex (JES stood for Job Entry System and the JES-plex was a cluster of large processors attached to huge farms of storage). But, like the laptop example above, the storage was carved up into private individual volumes being managed as thousands of little pools by individuals and the overall Storage Utilization was only 30%.

During the 1980s, IBM conducted over 100 studies of large enterprise shops and concluded that the more "pools" of storage (i.e. volumes), the more space problems, performance problems and people to manage the storage. The studies also showed that 80% of the management cost was spent on the individuals using the storage. When an individual runs out of space, jobs fail, things need to be deleted or re-provisioned with more space, application jobs need to be rerun, etc. This may be only a couple of hours for each individual but multiplied by thousands of them adds up real quick to a lot of lost productivity.

As a result, IBM engaged in a huge development effort to build SMS (System Managed Storage). SMS introduced the notions of storage pools (Storage Groups), policies (for retention, data protection, performance, etc.) at the file level and automatic assignment of these policies at file creation time. The physical notion of a volume basically disappeared. The volume names could still be in scripts (CLISTs or REXX scripts) or job control language (JCL) but they were ignored for provisioning. IBM's SMS basically turned the physical container into a dotted line and the only space consuming event was the creation of real data, not artificial volumes with artificial pockets of "locked" free space. The entire notion of provisioning changed from the carving up of individual private volumes to simply the dynamic allocation of storage from the available free space in the storage pool.

As a result of SMS, most customers ended up with one or two pools over the entire data center, Storage Utilization went to 85% or more, and mainframe storage basically ran on auto-pilot. Customers did not have to add people as data and storage grew, and most of the performance problems went away by the natural balancing of space over the pool. The pool does have to be managed but the management under SMS was far simpler than monitoring the thousands of little pools. You simply monitored the available free space and the growth of data. SMS also had automatic space management which would delete any unnecessary space, expire old data, migrate less used data, etc.

Another interesting aspect of SMS high-utilization and cheaper management was the effect on chargeback. Prior to SMS, many companies would do chargeback based on the volume containers handed out to individuals. The cost of managing the SMS pool was so much lower than managing physical puddles that many customers offered the choice. Each business unit could keep their own physical volume containers at \$X per GB or they could go in the SMS pool for 1/5 of that and be charged only by what they actually used. This may have been one of the reasons why over 3500 customers moved to SMS in less than two years.

People worry about giving users an "infinite" view of storage. The users will go nuts and take every byte available. The fact of the matter is that most corporate employees have better things to do than just sit around and eat up the storage resources. Furthermore, you always have to monitor for violations no matter how you dole out space.

By having a large storage pool, you create a huge buffer of free space for growth, gain significantly higher Storage Utilization, balance both space and performance over the entire pool, and use far less people to manage it. The "storage utility" model also has enough resources available to get through things like month end processing or any other high period of space activity. You are not artificially tying up huge amounts of storage in tiny little pools (volumes, LUNs), and you are not creating artificial space outages during high peak activity because of pre-carved puddles of space. The "used" storage space contains real data, not pockets of free space for each user and locked out for the rest of the corporation.

This model makes absolute sense if you think about household utilities and our own lives. We don't call the electric utility and ask them for more kilowatt-hr provisioning because we will be doing some heavy usage of the radial arm saw this weekend!

Another example – imagine the water company giving individual homes their own water tanks. Your sister, Julie, returns from a local playground with her 3 little kids. They are filthy and need a bath. Unfortunately, you just ran out of your provisioning supply. Bummer! It's going to be a long Fourth of July weekend for you and your family. Furthermore, at least 5 of your neighbors have an ample supply of water they won't be using because they went away this week. Unfortunately, this available water is not in your tank.

We don't live this way and we don't want to live this way. We view the electricity grid and the water supply as infinite. It isn't infinite and it needs to be managed, which it is. But the bottom line here is that the most efficient Storage Utilization model is the utility model with the fewest number of managed storage pools. IBM proved this with thousands of SMS customers that allowed them to achieve the highest Storage Utilizations in the computing industry.

Somehow, we forgot this model. As we cruised down the new SAN/NAS storage highway, we forgot to look in the rearview mirror. Now, we are in uncertain and difficult economic times, looking for cost savings, and living with underutilized storage farms. What's wrong with this picture?

3.2 IF YOU'RE DOING THIS, THEN YOU COULD BE UNDERUTILIZED!

Unlike the jokes made famous by comedian Jeff Foxworthy, these techniques aren't funny at all. In fact, in these times, they are really quite important because they represent real financial loss in capital expense, environmental ("green") expense, and operational expense. Here are the top ten obvious ones that are draining corporate wallets.

3.2.1 CREATING VOLUMES BASED ON PHYSICAL DISKS

This is truly a bad idea because as disk sizes have increased dramatically over the years, this made the size of volumes being handed out to applications an enormous waste of space. NetApp learned that back in the early 2000s when they replaced Traditional Volumes with FlexVols and Thin Provisioning in Data ONTAP 7 and above. The transition from Traditional Volumes to FlexVols has been unsurpassed in acceptance with NetApp customers.

3.2.2 PRE-CARVING UP STORAGE LUNS

This technique breaks every law in the books of the storage utility model. Someone decides that there will be standard sized LUNs to be doled out to users requesting storage. So you decide 50 GB is the standard and you hand them out to 1000 users. You now have 50 TB of storage initially sitting empty (to allow for future growth).

This is the most underutilized environment you can create – thousands of little underutilized pools. You can try to manage this mess by installing host agents and software to look at the real utilization from the host view and going around constantly shrinking and increasing these zillion little puddles. In the utility industry, this would be like having the water company build individual tanks for each home or the electric company building home owner sub-stations. The management cost for this is high, and you have most likely created many artificial performance bottlenecks by the pre-carving of these LUNs. NetApp has had thin provisioning since the introduction of ONTAP 7G in 2003 and allows volumes to auto-grow as space is needed.

3.2.3 COMBINING MULTIPLE LUNS FOR PERFORMANCE AND/OR CAPACITY

The technique of combining pre-carved LUNs into a "super" LUN for striping or extended capacity is another bad Storage Utilization technique built on top of the pre-carving technique above. Volumes should be automatically striped across the underlying disk storage. NetApp volumes are automatically striped across the Aggregate of real disks, and they are allowed to auto-grow. There are no special artifacts that one needs to create for performance or capacity.

3.2.4 PURCHASING SEPARATE STORAGE FOR SAN, NAS, AND TRANSFER PROTOCOLS

This popular technique is a very quick way to open the drain on corporate wallets because it automatically increases capital expenditure for new hardware and management software, new environmental costs for floor space, power, cooling, etc., and extra people costs to manage the separate towers of storage. Wow, cut me a slice of that!

NetApp believes in "Unified Storage" (SAN/NAS [CIFS/NFS]) with multiple transfer protocols (FCP, iSCSI, FCoE), a single platform (ONTAP/WAFL) that works across the entire storage family with the same APIs and data protocols. Unified Storage puts you in control of your destiny without requiring the financial overhead of bringing in multiple towers of storage and multiple management software.

3.2.5 KEEPING DUPLICATE BLOCKS ON STORAGE VOLUMES

We all know that today's data growth rates continue to stress the storage environment and the associated storage budgets. This being said, a large amount of this data contains duplicate blocks of information. This notion becomes obvious when you think of your own experiences with things like email chains, attachments to multiple users, common graphic, logos, pictures, golden boot images, multiple backups, data archives, unstructured data, etc. These practices will continue but they don't have to eat up storage unnecessarily.

NetApp customers have already achieved from 30% to 90% storage savings with our deduplication technology. Not only can the technology help produce significant savings for your primary storage but it will propagate to secondary and tertiary storage because of NetApp's inherent Thin Replication with SnapVault® and SnapMirror®.

3.2.6 MAKING MULTIPLE PHYSICAL COPIES OF VIRTUAL STORAGE OBJECTS

Most companies today are ending "server sprawl" and moving to virtualized server environments. This activity creates a huge cost savings for IT in that it replaces hundreds of physical servers with virtualized servers. Many of these servers use the same "golden boot image" and share the same "data stores". Rather than making duplicate copies of these storage objects, one can gain huge savings by using a single physical copy and hundreds of zero space clones. NetApp has the ability to clone logical volumes (FlexVols), LUNs and files. NetApp clones are also writeable but the shared common blocks with the parent will only have one image regardless of the number of clones created. This capability would easily allow application testing or data mining applications to run against real data bases by cloning the inactive backup images, enabling even more space savings without any impact to the online data.

3.2.7 TAKING FULL D2D BACKUPS OF VOLUMES

If you are taking full copies of volumes for backup, archive, disaster/recovery, records retention, business continuance, regulatory compliance, or other reasons, you could be wasting huge amounts of storage. NetApp snapshot technology is unsurpassed in that it only takes copies of the changed blocks on volumes. If there are no changed blocks, then no storage is consumed at all. This technology works for both block technology (SAN) and file technology (NAS), even running from the same storage array. The NetApp Snapshot technology creates compelling space savings not only for the primary backups, but these savings get propagated to the NetApp SnapMirror and NetApp SnapVault copies as well.

3.2.8 CREATING MULTIPLE SMALL STORAGE POOLS

Several vendors now have the notion of Thin Storage Pools, similar in concept to NetApp Aggregates. Although this is a better technique than LUN carving, it still creates multiple storage pools. The IBM studies mentioned earlier showed that there was a one-to-one correspondence between the number of pools and the number of space and performance problems, and the number of people needed to manage the storage farm.

NetApp introduced a technology in January of 2007 called Resource Pools. Resource Pools allow you to combine thousands of Aggregates over hundreds of NetApp arrays to be managed as a single pool for storage using NetApp's Provisioning Manager software. It's one thing being able to collect multiple pools of storage as a single pool. It's another thing to effectively manage that pool. Later in this paper, we will discuss how Provisioning Manager will dynamically balance both space and performance over the pool and provide a simple monitoring and trend analysis for storage growth. Not only does this greatly improve your Storage Utilization, but it reduces the number of failures, and people to manage it.

3.2.9 DOING MANUAL MANAGEMENT OF THE STORAGE POOL

Many of the storage management techniques in use today are highly resource intensive and not very effective. Putting quotas and limits on storage has never been effective in corporate environments, other than universities and controlled educational environments.

Whether I am being managed by quotas or a pre-carved sized LUN, when I run out of space, I will go to my manager and ask him if he still wants the project completed by November and he/she will call IT to either lift the quota or re-provision me. Not only is this labor intensive but the fact remains that many of the people in my business unit have plenty of space left on their pre-carved LUNs but I can't use it. So, we are spending people dollars on my side and unused storage dollars in other parts. This is not the way to run a railroad.

The Storage Utility model presents an “infinite” view of storage to the consumer just as the real life utilities like the electric company. We don’t think twice about plugging electronic and electric devices into our home receptacles as to how much power is really left in the power grid. It is our power pool and it looks infinite for our purposes. But it isn’t. The electric utility company is constantly monitoring the grid, peak periods, overall demand growth, outages, severe heat or cold, etc. It has to be monitored and managed to work. This is also true of storage.

NetApp’s Provisioning Manager[®] will not only let the storage administrator set up large Resource Pools over hundreds of physical arrays, but also automated policies for thin provisioning and automated management like snapshot auto-delete, volume auto-grow, dedupe, etc. Provisioning Manager will balance both capacity and performance over hundreds of NetApp arrays. Each Resource Pool will be monitored for growth trends along with automated triggers for thresholds and other alerts. Not only is this done for the primary Resource Pool, but it works hand in hand with NetApp’s Protection Manager to manage the secondary and tertiary pools of storage used for SnapVaults and SnapMirrors as well. This complete package of NetApp space management technologies will help customers achieve significant Storage Utilization percentages.

3.2.10 MANUALLY SEPARATING APPLICATIONS FOR SPACE/PERFORMANCE

This particular technique will increase the people cost and it often creates both artificial space and performance problems. If you only have 2 possible routes for your work commute versus 10 possible routes for your co-worker’s commute, then you are 5 times likely to hit more bottlenecks than he/she will. Furthermore, the isolation of free space on an application basis gets back to the puddle management versus the pool management.

With NetApp’s Provisioning Manager, nothing is manual and the real free space is shared in the Resource Pool. The provisioning selection is automatic, based on the Provisioning Policy for that set of application data (called Datasets) and the storage will be selected accordingly. Both space and performance will be evenly balanced by the automatic selection of storage over the entire set of arrays in the Resource Pool. This results in high Storage Utilization, balanced throughput, and negligible people management resources.

4 NETAPP MANAGEMENT CONSTRUCTS AND SOFTWARE

Before describing how to get the highest utilization using NetApp Technology, we need to define some basic NetApp Management Constructs, Management Software, and the relationship between them.

4.1 NETAPP MANAGEMENT CONSTRUCTS

4.1.1 RESOURCE POOLS

In the simplest definition, a Resource Pool is a collection of Aggregates or entire NetApp arrays, including V-Series. Aggregates are a collection of RAID Groups. NetApp storage arrays can contain multiple sets of Aggregates. There really is no hard and fast limit on the size of a Resource Pool. It can contain hundreds of arrays with thousands of Aggregates. Resource Pools will have standard properties (all of which are optional) like the following:

- Name – a user-specified name
- Description – a user-specified description of this Resource Pool
- Contact information for the Resource Pool owner(s)
- Multiple space thresholds for monitoring, alerting and trending
- Time zone (for pools selected for geographical reasons)
- Tags – user-controlled labels for provisioning control/monitoring (e.g., gold, silver, bronze, etc.), management control/monitoring, etc. (e.g., geographic location, business unit, etc.)

4.1.2 POLICIES

Policies are collection of management options, controls, and specifications for directing the automated management of data. Currently, there are two types of policies:

- Provisioning Policies
- Protection Policies

4.1.2.1 PROVISIONING POLICIES

Provisioning Policies not only help the selection of the appropriate storage for provisioning, but also the management controls for the sets of data being provisioned. These policies include things like:

- SAN/NAS provisioning specifications
- Thin provisioning, auto-grow controls
- Snapshot reserved space controls
- Auto-deletion controls (e.g., old snapshots)
- Deduplication controls
- Space utilization thresholds alert controls
- Automated offline migration controls for data mobility, policy correction, data migration, etc. for NFS exports, CIFS shares, LUN mappings, etc.
- Export protocol access (FC, iSCSI)
- Multi-tenancy access controls (e.g., CIFS ACLS access control)
- Host access controls
- Storage protection controls (RAID level, mirroring, external RAID (via V-Series), controller protection (e.g., active/active clusters, etc.)

4.1.2.2 PROTECTION POLICIES

Protection Policies set the automation controls for scheduling, monitoring, alerts, etc. on any set of data in terms of normal backup, offsite backup, disaster/recovery backup, regulatory copies, etc. This would include things like:

- Schedules for backups, mirror copies, mirror copy transfers
- Data transfer controls
- Backup data retention controls
- Warning and error thresholds
- Disaster/recovery controls for failover from primary to secondary or tertiary nodes

4.1.3 DATASETS

A Dataset is a collection of user data (volumes, qtrees, directories) plus all replications of that data to be managed by the same set of policies. It is possible to include objects in a Dataset by indirect reference – for example, and Aggregate, storage system, or OSSV¹ host. All members of the Dataset must have the same Protection Policy but different nodes of the Dataset could have a different Provisioning Policy. For example, you might have one policy for all of your Windows Exchange applications, but the Provisioning Policies for your primary, backup, and mirror nodes certainly could be different.

The attributes of a Dataset would include things like:

- Dataset name, owner, email contact, description, timezone
- Prefix for names of any objects created for this Dataset
- Protection Policy for this Dataset
- Provisioning Policy for this Dataset
- Specific Application Management for this Dataset (e.g., SnapManager for Oracle, SnapManager for Exchange, SnapManager for SharePoint, etc) – this would be used when the specific management application would be in control of the primary backup creation, and possibly retention.

¹ Open Systems SnapVault, or OSSV, is NetApp software, which when installed on an Open System makes it possible to back up file data on that system to a storage system or NearStore using the NetApp SnapVault[®] technology.

4.1.4 STORAGE SERVICES

A Storage Service is a combination of a Protection Policy and a Provisioning Policy that defines a level of service to be applied to a set of data, known as a Dataset. A cloud provider or a service provider could define a portfolio of Storage Services to represent a storage catalog of offerings to their clients. As an example, one might create a “Gold” Storage Service that is a combination of a “Gold” Protection Policy and a “Gold” Provisioning Policy. The “Gold” Storage Service could be applied to different Datasets (e.g., ERP Production, SAP Production, etc.).

4.2 NETAPP MANAGEMENT SOFTWARE

Your data continues to grow. You are being asked to increase your storage efficiency, utilize your storage better, save on both storage purchases, and environmental costs. You will have to do this without increasing your staff to handle the additional workload. The only way to scale the management of your data growth and, at the same time, reduce your cost, is to manage the storage and data using the higher level objects and automation policies described above. There are two NetApp software products that work hand-in-hand with these constructs, Provisioning Manager and Protection Manager.

4.2.1 PROVISIONING MANAGER

Provisioning Manager simplifies and automates the tasks of provisioning and managing NetApp storage for NAS and SAN access, and it can help you greatly improve your Storage Utilization. Some of its basic capabilities include:

- Create/modify Datasets, Storage Services, Resource Pools and Provisioning Policies
- Storage selection based on policy specifications, array I/O and CPU workload, required licenses and/or ONTAP versions, protocol service status, free space and thresholds
- Periodic checking to confirm that the storage conforms to the customer's policy
- Offline automatic migration of Multi-Tenant (vFilers) Datasets
- Automatic scheduling of Volume deduplication

4.2.2 PROTECTION MANAGER

Protection Manager simplifies and automates the tasks of managing backup and mirror relationships and provides easy interfaces for backup, restore, and failover operations. Furthermore, it uses Provisioning Manager to provision and manage its storage in the secondary and tertiary Resource Pools. Some of its basic capabilities include:

- Create/modify Protection Policies to manage primary data and the backup and mirror relationships for all of its replication copies and the scheduling of all backup functions
- Provision secondary storage for backups and SnapVault copies based on the assigned Protection Policy
- Provide a disaster/recovery (option) to quickly enable your offsite storage to provide data access until your primary storage has been re-enabled
- Validate the backup and disaster/recovery configurations with a conformance checker
- Overall monitoring with status dashboards, alerts, status value triggers.
- Data transfer reporting

The diagram in Figure 2) Overview of Provisioning Manager, Protection Manager with Datasets, Policies, and Resources Pools below shows the relationships between the constructs and the NetApp management software.

Each Dataset is assigned a Storage Service, which is the combination of Protection and Provisioning Policies. The Protection Policy defines all of the processing requirements for backup copies, offsite, disaster/recovery, schedules, etc. The Provisioning Policy defines the level of service required for provisioning that data in the Primary Pool. If the Protection Policy for that Dataset required the SnapVault of backups to a second site for compliance, offsite testing, archival retention, etc., then there would be a Secondary Resource Pool and the Dataset would have a Provisioning Policy for that pool.

Furthermore, if the NetApp Disaster/Recovery option was being employed to SnapMirror backup volumes, there would be a Tertiary Resource Pool at yet another geographic location and a Protection Policy describing the remote mirroring criteria and a Provisioning Policy defining the provisioning criteria for the tertiary pool.

The combination of these management constructs and the NetApp Management Software will completely automate the storage provisioning for all pools, dynamically balancing both the space and performance across the all pools through Provisioning Manager. It will also totally automate all of the backup processing through Protection Manager. The Storage Administrator need only monitor the free space decline over time

and threshold alerts in the Resource Pools along with the backup failures via automated alerts to manage the entire storage farm. This is the only way to automate and scale.

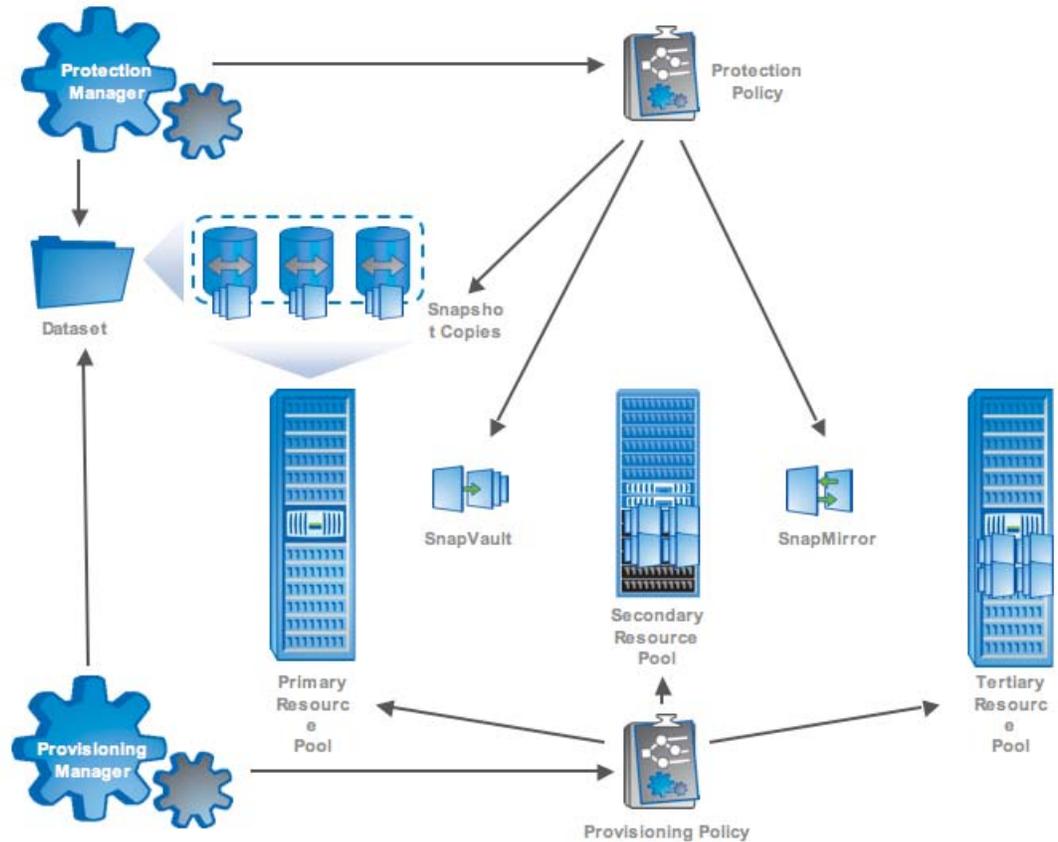


Figure 2) Overview of Provisioning Manager, Protection Manager with Datasets, Policies, and Resource Pools

5 GAINING HIGH STORAGE UTILIZATION WITH NETAPP®

5.1 MAXIMIZE THE NETAPP ARRAY STORAGE

There is a very simple list of things one can do to gain high space utilization in NetApp arrays. For each Aggregate you build, consider the following:

- Use the same size raw disks to avoid loss due to right-sizing (this also applies to the LUNs mounted from the OEM backend storage to the V-Series)
- Use RAID-DP™ for maximum disk failure protection and allow more data disks for Aggregate striping. To get the equivalent level of protection with the same performance, you would have to go with a RAID 10 scheme, which doubles the number of data disks.
- Do not use SyncMirror® unless the application absolutely demands that level of protection – it will double the number of raw disks in each Aggregate
- Ensure that you have enough Hot Spare Disks but not an excessive amount – you need at least two per controller to support a multiple failure situation in a RAID-DP™ configuration.
- Maximize the number of Data Disks in the Aggregate for the largest possible capacity available and for striping

5.2 MAXIMIZE THE “FREE CAPACITY (# 7) IN THE RESOURCE POOL

Resource Pools are the physical unit of storage and should be considered as very large containers for all data. These represent the “Usable Capacity” (# 3) in Figure 1) Storage Utilization and Storage Efficiency

Breakdown on page 4. To achieve the highest Storage Utilization, these Resource Pools should not contain any reserved or guaranteed space as described in “Unused Reserved Capacity” (# 6), only “Free Capacity” (# 7) and “Used Capacity” (# 15).

When the data center had only one or two storage arrays, we were more cautious and our white papers suggested very conservative defaults in the spirit of minimizing failures. Many things have changed since those days (Snapshot auto-delete, volume auto-grow, deduplication, Provisioning Manager, Resource Pools, thin provisioning, etc.).

It is not unusual in large enterprise shops today to find hundreds of storage arrays. But the fundamental management and Storage Utilization problems apply to customers with ten arrays and customers with a thousand arrays. It is just a matter of degree – the more storage systems, the bigger the problem. Lastly, the economic conditions of the early twenty-first century have totally changed the storage utilization equation. Saving money is no longer optional. It is critical to the business.

Here is a simple list of things you can do to easily gain at least 25% free space back in your storage pool:

- Turn off Aggregate Snapshot Reserve (approximate gain of 5%) – it is hardly ever used (or needed) today
- Turn off Snapshot Reserve (approximate gain of 20%) – when snapshots run out of space in the Snapshot Reserve area, they immediately take space out of the Volume Guaranteed Space area anyway. So, why have an artificial fence?
- Set Fractional Reserve to zero (default is 100%) for LUN and file overwrites

Note: Another way to turn off all reserved space is to ensure that the “Reserve space for snapshot copies” box is unchecked when using the Provisioning Manager UI Wizard to define Provisioning Policies.

The notion of turning off all guaranteed space and all reserved space may initially sound like a careless idea. You put those little fences there to ensure the application had enough free space in LUNs and volumes for both data and free space. So, let’s look at that closely. When the volume snapshots fill the Volume Snapshot Reserve space, it starts eating into the Volume Guaranteed Space of the volume so what is the point of having the boundary? You probably have hundreds, even thousands of volumes. When you pre-carve up space for LUNs, Volumes, and Snapshots, you have chosen the poorest space utilization technique because all of that storage appears as used capacity. No one else can use it. Furthermore, you have created an environment where you now have to manage and monitor thousands of little pools.

So, if you turn off all reserved and guaranteed space, what will happen if the application needs space? The answer is simple – from the free space. When you turn off all reserved and guaranteed space, you are truly using “Thin Provisioning” and all of the “Unused Reserved Capacity” (# 6) gets turned into “Free Capacity” (# 7). The management of this environment is so simple because you grow as you go and monitor only the free space in the Resource Pool.

5.3 DEFINING RESOURCE POOLS

The more Resource Pools you create, the more objects you have to monitor and manage so the goal should be to have as few as possible. There are three areas where you may want to define Resource Pools:

- Primary Pool
- Secondary Pool (for SnapVault copies of data to a second site)
- Tertiary Pool (for SnapMirror copies to yet a third remote location)

5.3.1 PRIMARY POOL

For your Primary Pool, how many storage Tiers will you have. Let’s say that you have a number of FAS6000 type arrays with FC drives, active-active clusters, RAID-DP™ protection as your Tier 1 storage and FAS3000 type arrays with ATA drives with RAID-DP™ as your Tier 2 storage. In this case, you could build two primary Resource Pools, one for Tier 1 storage and one for Tier 2 storage or you could build one Resource Pool and label your Tier 1 as “gold” and your Tier 2 as “silver”. Based on the Provisioning Policy associated with a particular Dataset, Provisioning Manager will automatically choose the proper Resource Pool for any provisioning requests for that set of data.

5.3.2 SECONDARY POOL

With Protection Manager, you will be defining Protection Policies for Datasets. These policies may include the use of SnapVault® to copy the backups to a remote secondary site. In this case, you will need to define a secondary Resource Pool for that location so Protection Manager can automatically provision storage (via Provisioning Manager) to the secondary location. Volume SnapVaults are perfect candidates for auto-deduplication turned on in the Provisioning Policy.

5.3.3 TERTIARY POOL

If your SLAs require additional storage sites for things like Disaster/Recovery or other requirements, then the Protection Policies for those Datasets requiring these additional backup requirements, would also have a Provisioning Policy for a tertiary Resource Pool.

5.4 DEFINING MANAGEMENT POLICIES AND STORAGE SERVICES

The whole point of moving to Storage Services and automated policies is to simplify the storage management problem. You certainly don't want to create a different management problem by the artificial proliferation of management constructs. How many of each should there be? My recommendation is think single digits, not hundreds, and certainly not thousands. Make sure that when you look at the name of each management construct, you can describe the meaning of it without looking up the definition.

5.4.1 DEFINING PROVISIONING POLICIES

The Provisioning Policies should be the easiest to define because they are closely related to the Resource Pool definitions. For your primary Resource Pool, you may have decided on a small set of Storage Tiers based on the level of service they could deliver (Tier 1, Tier 2, etc.) and/or transfer protocols (e.g., FCP, iSCSI, FCoE). If you have both SAN and NAS, you will need a separate Provisioning Policy for each. Lastly, if you have secondary and/or tertiary Resource Pools, you will need a Provisioning Policy for each. As you define these policies, you should be identifying the set of data these would apply to, things like home directories, Exchange mail boxes, production Oracle data, application development data, etc.). This information will be important to building Datasets.

To maximize your Storage Utilization, do the following in the Provisioning Manager UI Wizard when defining Provisioning Policies:

- Make sure the "Space guaranteed for dataset" box is unchecked when defining Provisioning Policies
- Ensure that the "Auto-grow Capacity" box is checked on the policy to allow auto-grow of the storage object on out of space conditions.
- Check the "Auto-delete Snapshot copies" box is on the policy to allow auto-deletion of older snapshots when free space is getting critical.
- There is a deduplication setting on Provisioning Policies called "Automated" – this will start a deduplication of any volume enabled for deduplication if more than 20% of new data has been added to the volume. Consider turning this option on to spread the deduplication load over the Resource Pool as new data is created and, in turn, creating new free space efficiency returns in the Resource Pool.

5.4.2 DEFINING PROTECTION POLICIES

How many protection SLAs (Service Level Agreement) do you have in your environment for local backups, remote backups, disaster/recovery, compliance, etc. What are your RTO and RPO criteria associated with them? How do you want to schedule the backup workloads to meet your objectives? The first thing you need to do is to decide on the set of Protection Policies that define each of the levels of data protection in your environment.

Each Protection Policy defines the schedule for backups and mirrors on the primary storage, SnapVaults of backups to secondary storage locations, SnapMirrors to tertiary locations, etc. For each unique set of Protection criteria, you need to define a unique Protection Policy. As you are defining these, mark down the sets of data these would apply to as it will be important to defining Datasets.

5.4.3 DEFINING STORAGE SERVICES

How many Storage Services do you need? Again, the number should be small. For example, you might have six Protection Policies and four Provisioning Policies. This doesn't mean you need all of the permutations of those policies, twenty-four in this case. It may turn out that the valid permutations for your

installation is only eight, which is a reasonable number. If your Service Catalog contains too many Storage Services, they will quickly become indiscernible.

5.4.4 DEFINING DATASETS

The next task is to decide which sets of data each Storage Service is to be applied to. These are called Datasets. You should have already identified these as you were defining your Provisioning Policies and your Protection Policies. A Dataset can have only one Protection Policy. In the primary Resource Pool, a Dataset can have only one Provisioning Policy. If the Dataset has a secondary and/or a tertiary Resource Pool based on its Protection Policy, you can assign an additional Provisioning Policy for each of those storage pools.

You don't want to arbitrarily create Datasets to the point where they become unmanageable and you should name them so they are meaningful to you. For example, suppose you have five applications (App1 through App5) and they each have a production set of data and a development set. If all of the production applications used the same Storage Service (e.g., SSProd), you could choose one Dataset called AppProd. Likewise, if all of the development applications used the same Storage Service (e.g., SSDev), you could create a second Dataset called AppDev.

The tiebreaker to choosing Datasets is the level of granularity you want/need for management. For example, if you absolutely need to do chargeback to individual applications, you could create ten Datasets (AppProd 1 through AppProd5, and AppDev1 through AppDev5).

Cloud providers and service providers have the biggest challenge here in that they have to separate charging for many customers, where each customer would need different levels of service (Storage Service). In this case, you might want Datasets like Cust1Prod through CustNProd and Cust1Dev through CustNDev.

If your granularity of management requires more than a relative handful of Datasets, then give some serious thought to the naming of them so you will understand the meaning just by looking at the name. Don't create a different kind of management problem by an artificial proliferation of constructs.

5.5 SETTING DFM GLOBAL OPTIONS

There are some DFM (Data Fabric Manager) global options that can be set using the dfm option set command that can also help improve Storage Utilization in the Resource Pools. The following sections describe these options.

5.5.1 DFM OPTION – SECONDARY VOLUME SIZING

You should ensure that the DFM dpDynamicSecondarySizing option is enabled. When the option is enabled, Protection Manager sizes new secondary volumes to a size capable of containing the primary data plus replication copies. It also resizes existing secondary volumes when necessary to provide enough space for a data transfer or when a new backup or mirror relationship is added.

5.5.2 DFM OPTION – DELETING ORPHAN RELATIONSHIPS

Set the dpReaperCleanupMode=Orphan to cause deletion of orphan non-imported relationships that are no longer in the Dataset. This setting also prevents the cleanup of any redundant relationships.

5.6 PROVISIONING LUNS AND VOLUMES

There are two ways to look at Thin Provisioning. One view is that we are really over provisioning the storage, like overbooking airline seats or hotel rooms to fill them. NetApp thinks the more positive view is that of the "go as you grow" – storage on demand, use only what you use. This is the highest utilization possible because you are only using what you need and managing the total free space required to run your business.

The NetApp Resource Pools can definitely handle the "go as you grow" but some operating systems get surprised by LUNs mounted on them when they grow over their initial size. To avoid this, use NetApp's Operations Manager to look at your current LUN sizes and their growth over time. From this, you can easily choose LUN sizes for your different Datasets that will last for a long time, maybe forever. The LUN will actually be thinned provisioned so it really won't take any space until data gets written to it. The containing

FlexVols can automatically grow using the auto-grow volume option. No real space will be taken for their provisioning until data gets written to them.

Some enterprise applications already support some form of thin provisioning and the autosize feature. Specifically, Microsoft® Exchange, Microsoft SQL Server® 2005, and Oracle® support the following options:

- Provision storage without initialization (that is, without zero-fill)
- Autosize storage for on-demand space consumption

5.7 MONITORING THE RESOURCE POOL

Using full “thin provisioning” with a huge Resource Pool will certainly return a significant amount of free storage to be used by applications and it will appear as an infinite utility. But, nothing is infinite. It must be monitored and managed. The task of doing this is now much simpler and less prone to failure because it is a single pool of free space to monitor.

Monitoring a single pool of free space is far easier, especially with Provisioning Manager. Here is a simple list of things you can do with Provisioning Manager to easily monitor and manage the huge Resource Pool for both SAN and NAS from a single console:

- The Resource Pool UI Wizard allows the settings of multiple threshold trigger events (e.g., normal and critical). These values should be set accordingly based on the normal amount of free space you need to run your business. These events could auto-page the storage admin, email the storage manager, trigger a dedupe of volumes, auto-delete snapshots, add hot spares to the Aggregate, etc. in order to keep the Resource Pool running in emergency situations.
- Provisioning Manager keeps track of both the used space and the free space in the Resource Pool so you can track the growth trends and determine based on current growth rates when the amount of remaining free space will hit critical stages. This would be the normal method of monitoring the Resource Pools.
- The Provisioning Policy UI Wizard allows the Storage Administrator to specify a script to be auto-invoked both before and after provisioning. This could be used for multiple purposes like chargeback, authorization, local provisioning rule verification, etc.

5.8 CHARGEBACK CONSIDERATIONS

There is no doubt that Resource Pools and full thin provisioning will create large amounts of new free space for future provisioning and growth. As you fill this free space, you will be able to run with a much higher utilization than you are experiencing today. The management of both growth and space availability will be infinitely simpler than managing all of the self-imposed boundaries of things like guaranteed and reserved space. Furthermore, using the management constructs of Storage Services, Datasets, Provisioning Policies, and Protection Policies will allow full automation of storage management to meet service level objectives freeing up significant people time. This will enable data centers to handle large data growth without the need to hire more personnel.

So, given these benefits, one might want to consider the SMS model mentioned earlier in the document – give users a choice of paying a premium price to keep things as they now are or a greatly reduced price in the managed Resource Pool.

6 NETAPP STORAGE EFFICIENCY TERMINOLOGY DEFINITIONS AND CALCULATIONS

Table 1) NetApp Storage Efficiency Terminology Bridge

| Definition | Calculation |
|---|--|
| 1. Raw Capacity: Total amount of addressable blocks on physical disk drives. Calculated by multiplying the number of disk drives by the labeled capacity of those disk drives. For V-Series platforms it's the size of the array LUNs | Raw Capacity = Disk Drives x Labeled Disk Drives Capacity |
| 2. System Reserve Capacity: Capacity required for fixed system reserves, RAID parity, Mirroring, and Spare Drives | System Reserve Capacity = Fixed Reserve + RAID Reserve + Spares |
| 3. Usable Capacity: Capacity available to applications and users | Usable Capacity = Raw Capacity - System Reserve Capacity |
| 4. Fixed Reserve Capacity: Capacity consumed by Kernel, Checksum, WAFL Reserve and FlexVol Metadata. These are not adjustable by the user | Fixed Reserve Capacity = Kernel Overhead + RAID Checksum, + WAFL Reserve + FlexVol Metadata |
| 5. RAID Reserve/Spares Capacity: Capacity consumed by Parity, Mirrors and Spares | RAID Reserve/Spares Capacity = Parity + Mirrors + Spares |
| 6. Unused Reserve Capacity: Capacity allocated but unused by Aggregate Snapshot Reserve, Volume Snapshot Reserve, Volume Fractional Reserve, Vol/LUN/File Guaranteed Space. These reserves are adjustable by the user | Space Reserve Capacity = Aggregate Snapshot Reserve + Volume Snapshot Reserve + Volume Fractional Reserve + Vol/LUN/File Guaranteed Space |
| 7. Free Capacity: Capacity that is available but not currently used | Free Capacity = Usable Capacity - Used Capacity - Space Reserve |
| 8. Available Capacity: Remaining Usable Capacity after Space Reserved Capacity has been assigned | Available Capacity = Usable Capacity - Space Reserve Capacity |
| 9. Used Capacity: Usable capacity that has been provisioned for applications or users | Used Capacity = Usable Capacity - Free Capacity - Space Reserve |
| 10. Effective Used Capacity: Used space that is visible and reported as used by applications and users, taking into account the efficiency returns | Effective Used Capacity = Used Capacity + Efficiency Returns |
| 11. Efficiency Returns: Effective capacity return of NetApp efficiencies when compared to traditional storage techniques | Efficiency Returns = Capacity saved by: Cloning + Deduplication + Snapshot + RAID-DP + Thin Provisioning |
| 12. Cloning Return: Capacity savings resulting from FlexClones and LUN Clones | Cloning Return = Number of clone copies x size of parent clone |
| 13. Dedupen Return: Capacity savings resulting from deduplication and sub-volume cloning (File, LUN, sub-File clones) introduced as sis-clone in ONTAP 7.3.1 | Dedupe Return = dedupe savings (includes the sis-clone savings) |
| 14. Snapshot Return: Capacity savings of Snapshot copies when compared to full volume copies | Snapshot Return = Volume capacity - Snapshot capacity (compares periodic full volume copies [default of 1] vs the total of all snapshots – future calculation will allow user specified number of days between full volume copies for fairer comparison) |
| 15. RAID-DP Return: Capacity savings of RAID-DP when compared to disk mirroring (RAID 10) to get the equivalent protection and performance | RAID-DP Return = Number of disk drives required for disk mirroring - Number of disk drives required for RAID-DP |
| 16. Thin Provisioning Return: Capacity savings due to the over-provisioning of volume storage | Thin Provisioning Return = Provisioning size requested for the FlexVol – actual size of the FlexVol |

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

© 2010 NetApp. All rights reserved. Specifications are subject to change without notice. NetApp, the NetApp logo, "Go further, faster," Data ONTAP, WAFL, FlexClone, SnapMirror, SyncMirror, SnapVault, Provisioning Manager and FlexClone are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Oracle® is a registered trademark of the Oracle Corporation. Microsoft® is a registered trademark of the Microsoft Corporation. TR-3827