Understand VMware VAAI with QNAP Turbo NAS

Technical White Paper
October 2013

Understand VMware VAAI with QNAP Turbo NAS

About VAAI

vStorage API for Array Integration (VAAI) is a technology that VMware makes available for storage vendors to use. The main goal of this technology is to offload the storage operations from the VMware server virtualization hardware to the storage array.

Audience
This document will treat the VAAI technology applied to QNAP users in a simple way to explain simply the way it works, the benefits, and why it should be used.

Supported hardware

- QNAP Turbo NAS, Business Series with firmware QTS 4.0
- VMware 5.x
- NAS Models : TS-x69 Series, TS-x70 Series, TS-x79 Series, TS-x79U-SAS Series, SS-x79U-SAS Series

Audience

This document is intended to help understanding VAAI technology on a technical point of view (requires some IT knowledge), and its benefits on a user point of view. It includes the description of a test environment to illustrate it and its usage in real situation other than labs and experiments.

Table of Contents

1. The VAAI technology explained and it's benefits

What is VAAI?

VMware vStorage APIs Integration (VAAI) is a set of APIs provides by vSphere that allows the ESXi hosts to offload data processing of certain storage-related services to storage systems. The QNAP Turbo NAS will process the data instead of the ESXi hosts.

Overall Benefits:
The integration of QNAP NAS with VMware VAAI offers customers the chance to deploy a virtualized environment with optimized performance. It will offload the data processing to the NAS and releases the ESXi hosts of certain tasks.

VAAI for iSCSI supports 4 features:

- Full copy (hardware-assisted copy)
- Block zeroing (hardware-assisted zeroing)
- Hardware-assisted locking
- Thin provisioning and space reclaim

VAAI for NAS (NFS) supports 3 features:

- Full File Clone
- Space Reserve
- Extended Statistic

VAAI for iSCSI

VAAI iSCSI - Block Copy

Block Copy enables the QNAP storage to make full copies of data within the NAS without needing to have the ESXi host read and write the data.

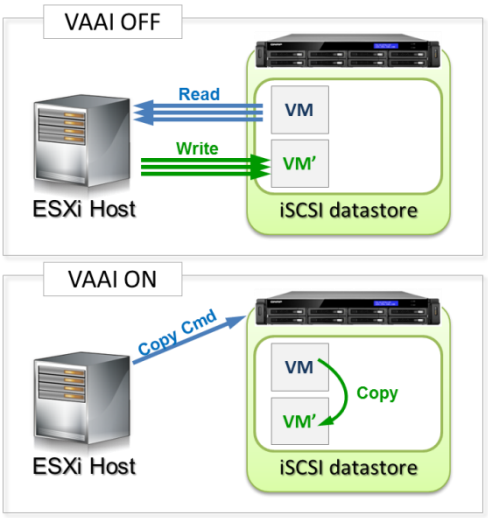

Without VAAI:
When an ESXi host copies data without VAAI, it will have to read the data from the NAS, and write back the data to the NAS to copy the data. This process consumes memory, network bandwidth, and CPU from the ESXi host, decreasing the resources available for the VMs.

With VAAI:
When an ESXi host copies data with VAAI, it will send a command to the NAS, and wait for the result. This process offloads the loading to the Turbo NAS, and leaves the resources allocated to the VM unchanged. The copy process is done in the NAS directly by the NAS. CPU, memory, and network bandwidth are not used in the copy procedure.

Benefits:

- Reduces the loading of ESXi hosts
- Faster virtual machine cloning
- Faster template deployment
- Faster storage vMotion

VAAI iSCSI - Block Zeroing

Block Zeroing allows the QNAP NAS to zero out a large number of blocks to speed up the provisioning of VMs.

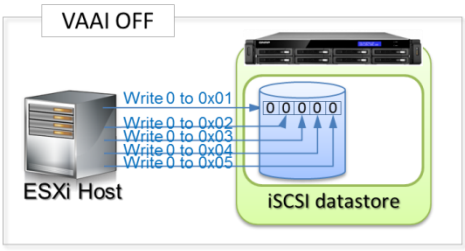

Without VAAI:
When an ESXi host creates and initializes a virtual disk without VAAI, it has to write numerous zeroes to the NAS to create the initialized file. This process consumes memory, network bandwidth, and CPU from the ESXi host, decreasing the resources available for the VMs.

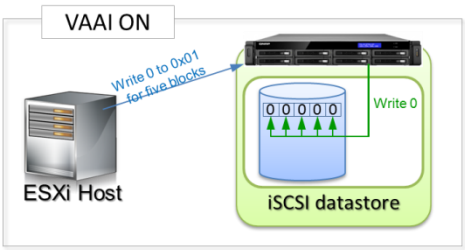

With VAAI:
When an ESXi host creates and initializes a virtual disk with VAAI, it will send a command to the NAS, and waits for the result. This process offloads the loading to the Turbo NAS, and leaves the resources allocated to the VM unchanged. The block initialization process is done by the NAS directly. CPU, memory, and network bandwidth are not used for that operation.

Benefits:

- Reduces the loading of ESXi hosts.
- Faster allocation of the capacity for virtual machines.

Effective usage:

- Create Thick Provision Eager Zeroed virtual disks.
- Create VM with Thick Provision Eager Zeroed virtual disks.

VAAI iSCSI - Hardware assisted Locking

Hardware Assisted Locking improves efficiency on large clusters by locking only required blocks instead of the entire LUN. When read/write operations occur from one ESXi server, it will keep the rest of the LUN available to be accessed to increase efficiency.

Without Hardware Assisted Locking:
When an ESXi host is accessing a datastore on an iSCSI LUN, the iSCSI LUN will be locked and other ESXi hosts have to wait the LUN to be released. The more ESXi hosts that share the same datastore, the lower the performance will be.

With Hardware Assisted Locking:
When an ESXi host is accessing a datastore on an iSCSI LUN, the required blocks are locked, and the iSCSI LUN remains available for the other ESXi hosts. More concurrent operations can be achieved while maintaining a good performance level.

Benefits:

- Permits the QNAP NAS to scale to more VMs
- Permits more ESXi hosts without performance penalty.

Effective usage:

- Large VMware cluster with shared datastores.

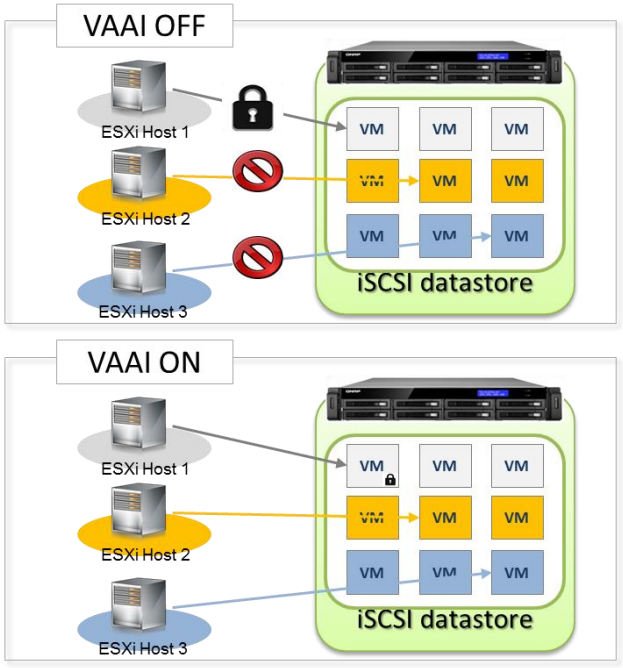

VAAI iSCSI - Thin Provisioning and Space Reclaim

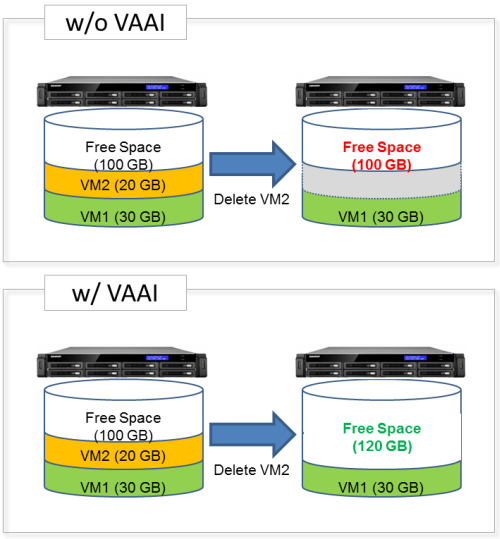

Thin Provisioning allows the allocation of capacity only when it's needed. It is possible to pre-allocate capacity for iSCSI LUNs without effectively using it. This capacity will be used only when data, Virtual Machines or virtual disks, will be stored in the LUN.

Space Reclaim allows the space to be released when VM or virtual disks are deleted or migrated to other LUNs. It avoids wasting capacity after data has been removed.

Without Thin Provisioning and Space Reclaim:
Without Thin Provisioning and Space Reclaim, the space allocation is done at LUN creation, and cannot be released until the LUN is deleted. In that case, if an administrator creates a 100 TB LUN, the 100 TB will be reserved for the LUN and cannot be used for any other applications. If some data are deleted from the LUN, the NAS will not have more free space, as the deleted data cannot be freed. The storage management is less flexible, requires more physical capacity (Hard Drives) and increases costs.

With Thin Provisioning and Space Reclaim:
With Thin Provisioning and Space Reclaim, the space allocation is done only when data is effectively written, and it can be released when data is deleted. In that case, if an administrator creates a 100 TB LUN, but will be consumed only when VM are created, and will be released when VMs are deleted. The free space on the NAS will dynamically increase and decrease as the VMs are created and deleted. The storage management is much more flexible and requires less physical capacity (Hard Drives). It increases efficiency and reduces costs of storage.

Benefits:

- More accurate reporting of real disk space usage.
- Better and more efficient storage management
- Reduce costs of storage.
- Use capacity only when needed
- Purchase Hard Drives only when needed

Effective usage:

- VM creation / deletion
- Snapshot creation/deletion

VAAI for NAS (NFS)

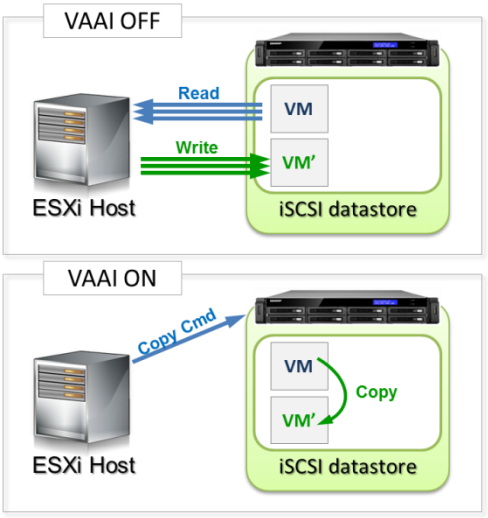

VAAI NAS - Full file clone

Similar to the Full Copy in VAAI iSCSI, it enables the QNAP storage to make full copies of data within the NAS without needing to have the ESXi host read and write the data.

Without Full File Clone:
When an ESXi host copies data without VAAI, it will have to read the data from the NAS, and write back the data to the NAS to copy the data. This process consumes memory, network bandwidth, and CPU from the ESXi host, decreasing the resources available for the VMs.

With Full File Clone:
When an ESXi host copies data with VAAI, it will send a command to the NAS, and wait for the result. This process offloads the loading to the Turbo NAS, and leaves the resources allocated to the VM unchanged. The copy process is done in the NAS directly. CPU, memory, and network bandwidth are not used for the

Benefits:

- Reduces the loading of ESXi hosts

Effective usage:

- vCenter Cloning
- Storage vMotion (cold clone)
- Deploy VMs from templates

VAAI NAS - Space Reserve

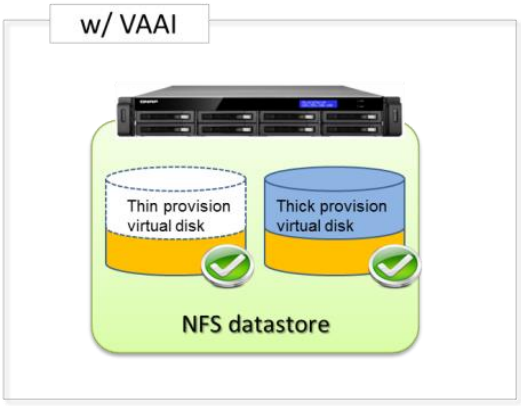

Historically, vSphere only supported building thin provisioned virtual disks on NFS. By supporting the Reserve Space, a space pre-allocated virtual disk (thick provision eager zeroed disk) is possible in QNAP NFS datastores, by offloading the file creation onto the NAS.

Benefits:

- Better virtual disk read/write performance (thin provision disks vs thick provision disks)

Without Space Reserve:

- Without Space Reserve, user cannot create Thick Provisioned virtual disks on NFS datastore and only Thin Provisioned is available.

With Space Reserve:

- With Space Reserve, user can create Thick Provisioned virtual disks on NFS datastores.

VAAI NAS - Extended statistic

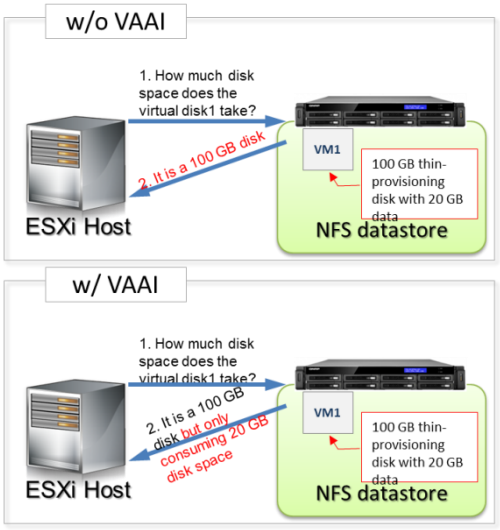

Extended Statistic enables vSphere to query space utilization details for virtual disks on QNAP NFS datastores. This includes the size of a virtual disk and the real space consumption of the virtual disk.

The extended statistics are used by VMware during some specific operations like Space Reserve and Clone.

Also, it can be retrieve by command line from the ESXi host:

```
/vmfs/volumes/86721f64-65fb714e/MyVM1 # vmkfstools --extendedstat MyVM1.vmdk
Capacity bytes: 17179869184
Used bytes: 1608630272
Unshared bytes: 1608630272
```

2. Installation and Verification

VAAI for iSCSI

VAAI iSCSI does not require any installation. It is enabled in VMware by default.
To Verify that the storage supports VAAI, simply open the storage configuration in vSphere Client, select the datastore. The column "Hardware Acceleration" (VAAI) will show supported.

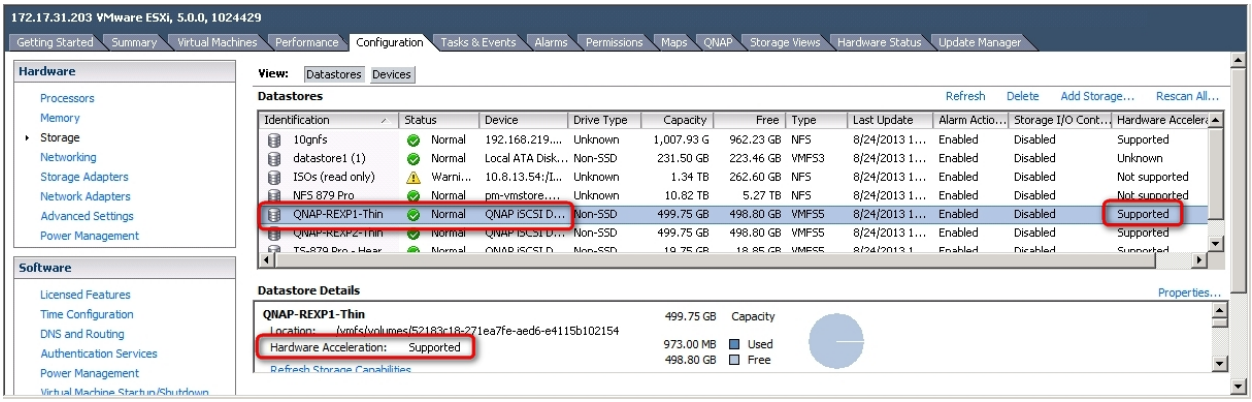

Enabling and disabling VAAI in VMware is possible. Please refer to VMware documentation for this operation:
Disabling the VAAI functionality in ESXi/ESX (1033665): http://kb.vmware.com/kb/1033665

Additionally, it can be verified by command line through SSH or ESXi Shell with the command:

```
~ # esxcli storage core device vaai status get
naa.6e843b689a4436ed8effd450ad9276d1
   VAAI Plugin Name:
   ATS Status: supported
   Clone Status: supported
   Zero Status: supported
   Delete Status: supported
```

We can see the 4 VAAI features supported by the QNAP iSCSI LUN.

VAAI for NAS (NFS)

VAAI NFS requires the installation of VAAI plugin on each ESXi hosts.

If the ESXi servers are not connected to internet:

- Download the VAAI NAS plugin from: http://www.qnap.com/utility
  - Or with the current direct link: http://eu1.qnap.com/Storage/Utility/QNAP_QNPNasPlugin_1.0.zip
- Install the plugin following VMware procedure:
  - Installing patches on an ESXi 5.x host from the command line: http://kb.vmware.com/kb/2008939
- Reboot each ESXi host.
- Verify from vSphere client that "Hardware Acceleration" is supported.

If the ESXi servers are connected to internet:

- Open the each ESXi console (or through SSH) and execute the command:

```
~ # esxcli software vib install -v http://eu1.qnap.com/Storage/Utility/QNAP_QNPNasPlugin_1.0.vib --no-sig-check
Installation Result
   Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
   Reboot Required: true
   VIBs Installed: QNAP_bootbank_QNPNasPlugin_1.0-1.0j
   VIBs Removed:
   VIBs Skipped:
~ #
```

  - –no-sig-check is set to allow non-signed VIB packages. The package will be signed when it will be VMware Ready (At this time, the certification process is in progress).
  - A reboot is required to activate the new VIB.
  - After each installation, it displays VIBs installed, removed, skipped.
- Reboot each ESXi host after installation
- Verify that VAAI NAS is supported on NFS datastores from vSphere Client :

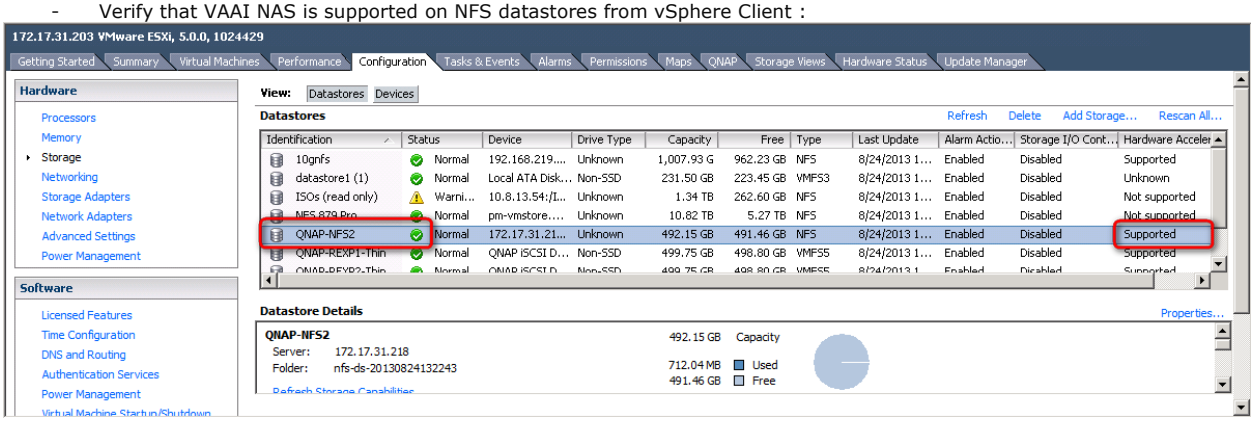

Note: During the installation, if a previous version of the plugin was installed, it will be listed as VIBs Removed:

```
~ # esxcli software vib install -v http://eu1.qnap.com/Storage/Utility/QNAP_QNPNasPlugin_1.0.vib --no-sig-check
Installation Result
   Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
   Reboot Required: true
   VIBs Installed: QNAP_bootbank_QNPNasPlugin_1.0-1.0j
VIBs Removed: QNAP_bootbank_QNPNasPlugin_1.0-1.0h
   VIBs Skipped:
```

3. Application scenarios and results

Difficulties that the storage has to face against virtualization storage

QNAP Turbo NAS provides administrators comfort and safety when using VMware, while saving storage budget for other usage. With the performance improvement, efficient storage management, and better storage usage reporting, QNAP Turbo NAS meet business needs by constantly adding and improving new features.

Performance for VMs management and movement

Virtual Machine can be compared as large files, between 10GB and 100GB. When an administrator needs to deploy additional virtual machines or copy/backup existing virtual machines, the process can be compared as copying those large files. In a traditional environment, it would involve and require resources from different parts of the environment such as storage, network, ESX servers, and CPU and would reduce the available resources for the end-user using those virtual machines.

Using QNAP Turbo NAS and VMware VAAI, the copy process will be offloaded to the storage so that the impact of running virtual machines will be minimized.

Performance for clusters

VMware Cluster allows connecting multiple ESXi host to the same datastore to provide high availability and easy live migration. However, the downside is that the more ESXi hosts belong the cluster performance will decrease when accessing the datastore, as explained in VAAI iSCSI - Hardware assisted Locking (page 5.) With traditional storage, it is advised to create multiple datastores with fewer virtual machines or to connect multiple storages.

Using QNAP turbo NAS and VAAI, the performance impact is largely reduced allowing for a larger number of VMs to be running.

Storage efficiency management

IP-SAN/iSCSI Storage provides capacity to servers. However, without thin provisioning physical storage is needed and the whole allocated capacity must be purchased before being provisioned. With thin provisioning, the space can be allocated, without having the real capacity available. The downside of standard thin provisioning is that when the content is deleted from the LUN it cannot be freed, and capacity management can become more complicated. Space Reclaim can solve that last issue by allowing compatible client to reclaim the capacity.
QNAP Turbo NAS with Thin Provisioning and Space Reclaim allows a more efficient storage management. It is possible to virtually allocate storage capacity but use it only when it is really needed. The administrators can now allocate capacity in advance with thin provisioning.

The following tests illustrate those issues and solution with QNAP Turbo NAS.

The Test Environment

The tests and results have been done with a non-official firmware before the official release. The performance result in this document may differ from the official firmware.

This application scenarios and tests have been realized using the following components:

- 2 ESXi hosts
  - HP ProLiant ML 110, with Xeon Processor and 24 GB of RAM
  - HP ProLiant ML 110, with Xeon Processor and 24 GB of RAM
- 1 QNAP Turbo NAS
  - TS-EC1279U-RP + expansion units
  - Turbo NAS firmware version 4.0 (supports VAAI NAS and VAAI iSCSI)
  - Expansion Units : REXP-1200U-RP x2
  - Hard Drives : SEAGATE ST9300653SS (SAS drives)
  - RAID configuration : RAID 6 with 9 hard drives + 1 Hot Spare on each expansion unit (total 10 HDD per unit)
- vSphere 5.0
  - VMware ESXi 5.0
  - Enterprise Licensed to support VAAI
  - 2 VMFS Datastores (iSCSI), block based, with thin provisioning
    - VMFS Datastore 1 : QNAP-REXP1-Thin
    - VMFS Datastore 2 : QNAP-REXP2-Thin
  - 2 NFS Datastores
    - NFS Datastore 1 : QNAP-NFS1
    - NFS Datastore 2 : QNAP-NFS2
    - ESXi host 1 with VAAI NAS disabled
    - ESXi host 2 with **VAAI NAS enabled**
- Standard gigabit network connection

VM template deployment: Offload data processing when copying data

Deploy a VM template to create a new VM- with VAAI iSCSI Full Copy or VAAI NAS Full File Clone

One possible application of VAAI is to deploy a new virtual machine from a template.
To deploy a template to a new virtual machine, click on the template and select "Deploy Virtual Machine from this Template".

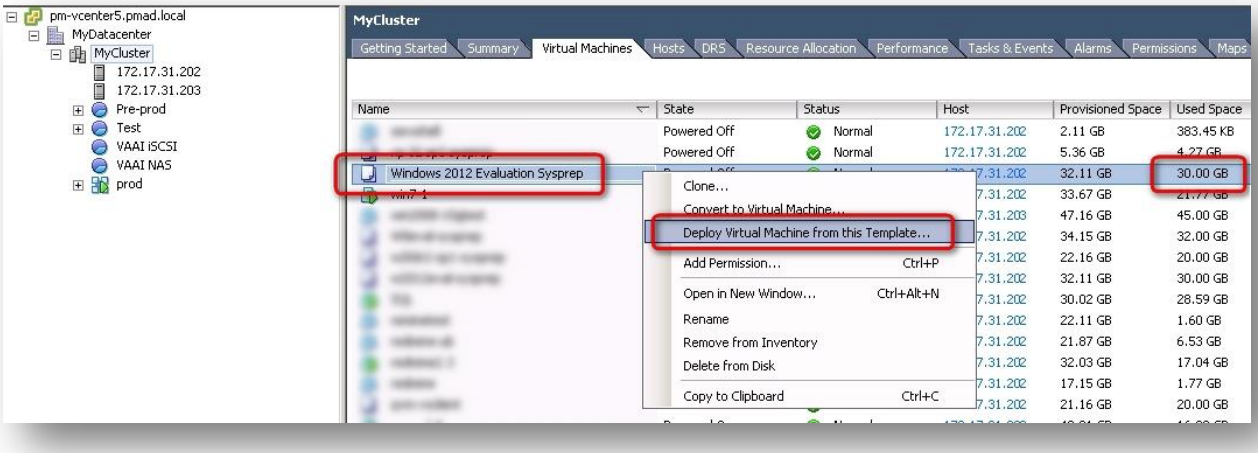

The Template is stored on the VMFS Datastore 1, on an iSCSI LUN with Thin Provisioning. On the above screenshot, we see how to deploy the template. It uses 30 GB of capacity.

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Ti...	Start Time	Completed Time
Clone virtual machine	Windows 2012 Evaluation Syspr...	Completed		PMAD\jaussa...	pm-vcenter5....	9/5/2013 10:21:32	9/5/2013 10:21:32	9/5/2013 10:22:39

Requested Start Ti...	Start Time	Completed Time
9/5/2013 10:21:32	9/5/2013 10:21:32	9/5/2013 10:22:39

From VMware logs, we can see that the time needed to deploy the template to create a new VM is only **1 minute and 7 seconds**! Because the template is 30 GB, this corresponds to a throughput of **458 MB/s**. Note that the speed of the deployment can vary depending on the environment, hardware, software and template location.

From the NAS UI, we can verify that the network has not been used. The copy operation has effectively been offloaded to the NAS.

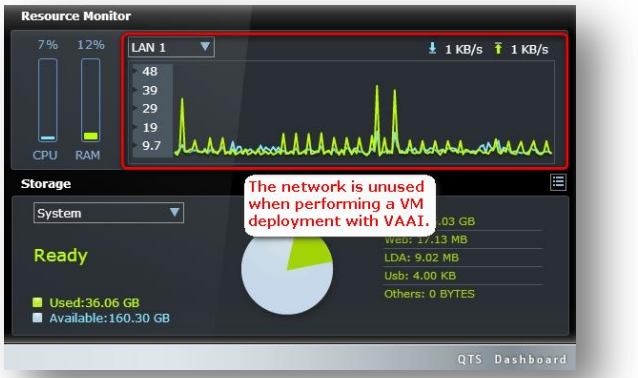

To compare, we repeat the procedure to another LUN without using VAAI.

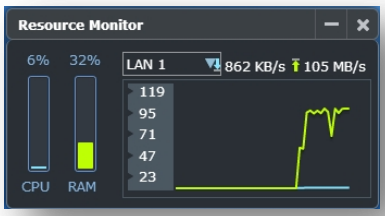

We can easily see the network usage from the NAS resource monitor, with a max throughput of 100 MB/s using 1 gigabit LAN. Because the NAS is fast enough, we can reach the maximum throughput allowed with 1 gigabit network, but it is still lower than using VAAI.

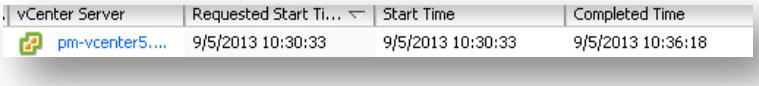

As a result we can conclude that in this test, there is a huge difference when deploying a template with VAAI and the benefit can easily be understood, in time, and resource usage:

	With VAAI	Without VAAI
Time to deploy 30 GB template	67 seconds	5 minutes 45 seconds
Average bandwidth to deploy the template	458 MB/s	89 MB/s

VAAI in this test is 5 times faster for deploying a template!

Clone a VM - with VAAI iSCSI Full Copy or VAAI NAS Full File Clone

The benefits obtained above when deploying a VM template can also be found when cloning a VM. VAAI provides this functionality for VMFS datastores (iSCSI) and for NAS (NFS).

- VAAI iSCSI provides "Full Copy"
- VAAI NAS provides "Full File Clone"

The behavior of cloning a virtual machine is the same as deploying a template. The process is offloaded to the NAS. The network will not be used to perform the copy and the performance needed time will be lower.

Note: VAAI will be used only with the same datastore type:

- VAAI will be used between 2 iSCSI LUNs
- VAAI will be used between 2 NFS shares.
- VAAI cannot be used between 1 iSCSI LUN and 1 NFS share.

Create new VMs and Offload data processing

Create a new VM on a iSCSI datastore with eager zeroed disks - with VAAI iSCSI Block Zeroing

During the creation of a new Virtual machine, the user as the choice to choose Thick provisioning (lazy or eager zeroed) and Thin Provisioning. If the user choose "Thick Eager Zeroed" disk, VMware will perform a virtual disk creation, reserve the capacity, and initialize the disk with zeros. This requires a lot of time if VAAI is not available, where it can take only a few second with VAAI.

In our test environment, we are equipped with only a gigabit LAN, so if we create a 100 GB Hard Drive, the time would be more than 15 minutes at 100 MB/s (single gigabit LAN average speed).
Using VAAI, the test shows that to create a 100 GB virtual machine with Thick Provisioned Eager Zeroed Disk, It took only 7 sec!

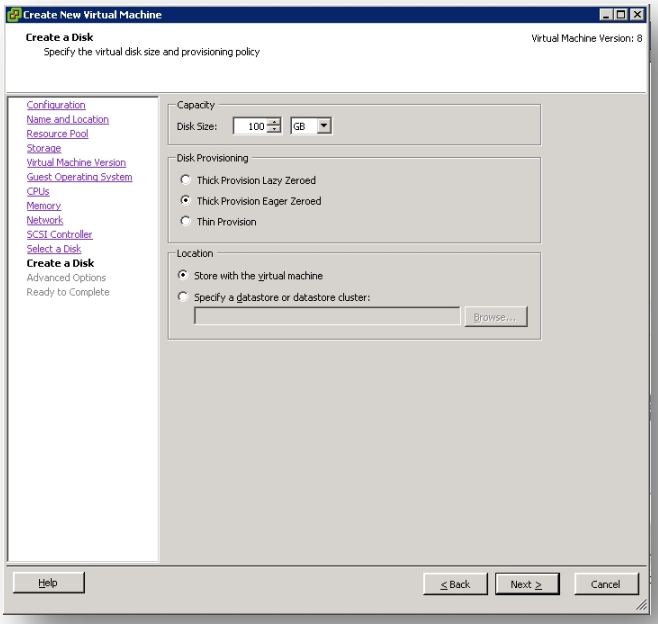

Thick Provisioned Eager Zeroed Disk on the QNAP Thick Provisioned iSCSI LUN takes **2 minutes 21 seconds**:

Name	Target	Status	Details	Initia..	vCenter Server	Requested Start Ti...	Start Time	Completed Time
Create virtual machine	MyDatacenter	Completed		PM...	pm-vcenter5....	9/5/2013 11:00:30	9/5/2013 11:00:30	9/5/2013 11:02:21

Thick Provisioned Eager ZeroedDisk on the QNAP Thin Provisioned iSCSI LUN takes **7 seconds**:

Name	Target	Status	Details	Initia..	vCenter Server	Requested Start Ti...	Start Time	Completed Time
Create virtual machine	MyDatacenter	Completed		PM...	pm-vcenter5....	9/5/2013 11:04:50	9/5/2013 11:04:50	9/5/2013 11:04:57

Create Thick Provision Disk on NFS Datastore – with VAAI NAS Space Reserve

VM Creation on NFS Datastore **withVAAI**

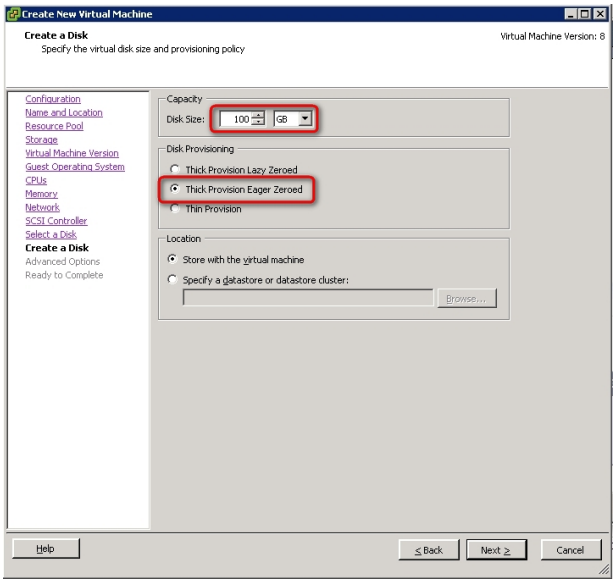

VAAI NAS allows Thick Provisioned Disk on NFS Datastore

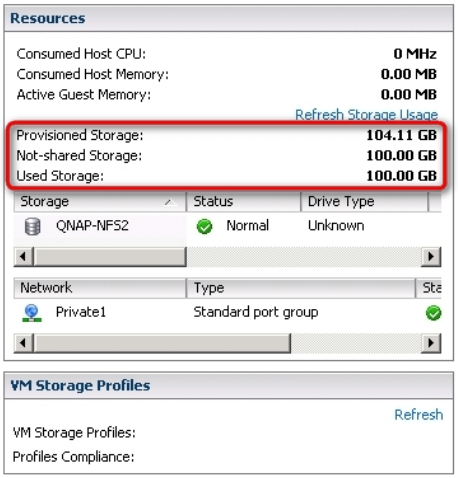

VAAI NAS allows for VMware to pre-allocate and reserve the capacity using Thick Provision Disk.

VM Creation on NFS Datastore **withoutVAAI**

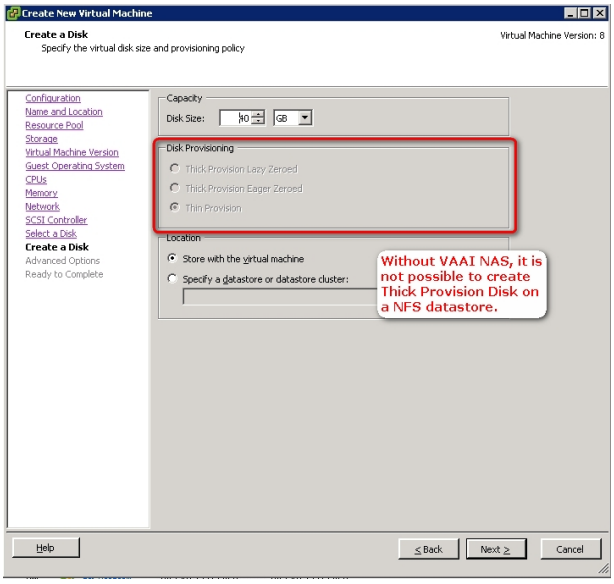

Without VAAI NAS, only Thin Provisioned Disk are allowed

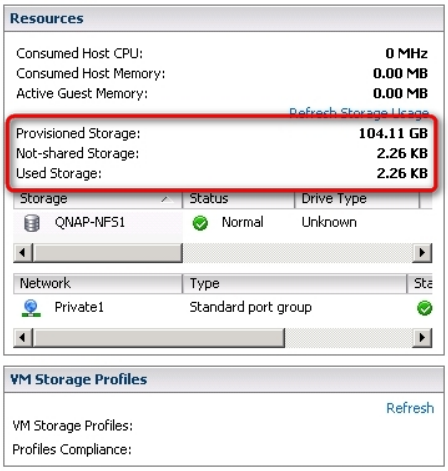

Only Thin Provision Disk on NFS datastore are allowed. Therefore it is not possible to reserve the capacity on NFS datastore without VAAI NAS.

Note about VAAI iSCSI Block Zeroing and VAAI NAS Space Reserve:

As explained previously, historically only Thin Provisioned Disk was possible on NFS datastores.
With the introduction of Space Reserve, it is possible to create Thick Provisioned VMDKs on NFS datastores.

However VAAI NAS Space Reserve is not like VAAI iSCSI block zeroing; it does not send the command to the storage to offload the operation, but simply allow the operation that was not possible in the past. When creating a VMDK on a VAAI NAS storage, Space Reserve will check the NAS storage to make sure that the space is available. Then, VMware has the possibility to create and reserve the capacity for the VMDKfile. This means that creating a Thick Provisioned Eager Zeroed Disk on a NFS datastore is a slow operation, and any write operations are sent over the network, and not offloaded as it is for iSCSI storage.

Efficient Storage Capacity Management with Thin Provisioning and Space Reclaim

As explained previously, Thin Provisioning LUN and Space Reclaim allow a better and more efficient capacity management.
This can be demonstrated simply.
The first step is to create a small datastore. "Small" so that it is easy to see the effect with a small example as the one below.

In this example, we are going to connect 2 hard drives and create a small Storage Pool in the QNAP NAS, only to demonstrate Thin Provisioning and Space Reclaim:

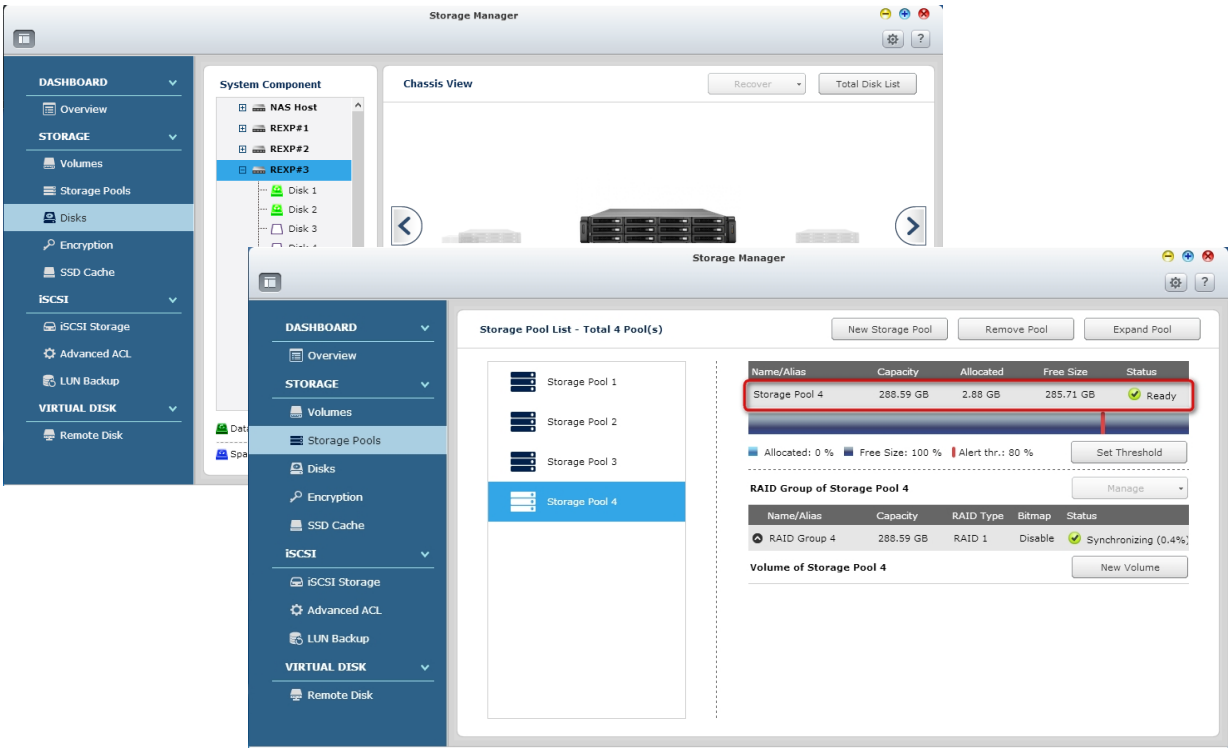

In the screenshot above, we added 2 disks in order to create a small storage pool for this test: 285 GB.

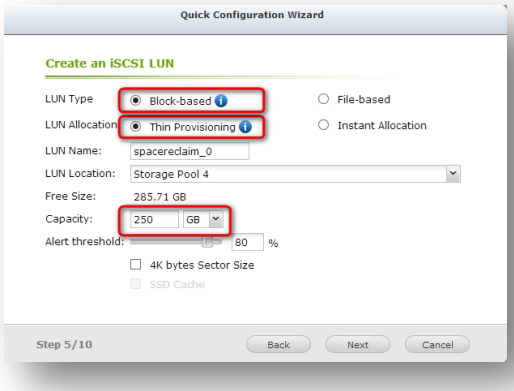

Now, we create a Thin Provisioned iSCSI LUN, which will be used as datastore for VMware.
iSCSI LUN size : 250 GB

Then we connect the newly created LUN to VMware and create a new datastore to test Thin Provisioning and Space Reclaim.

We can verify that VAAI is supported: "Hardware Acceleration Supported"

The free capacity is around 250 GB as we did not create or store any VM inside.

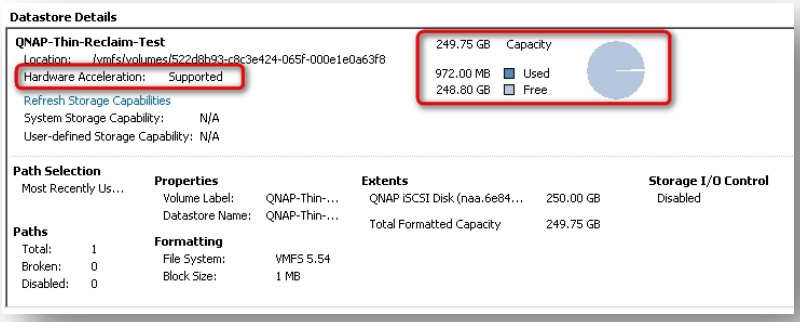

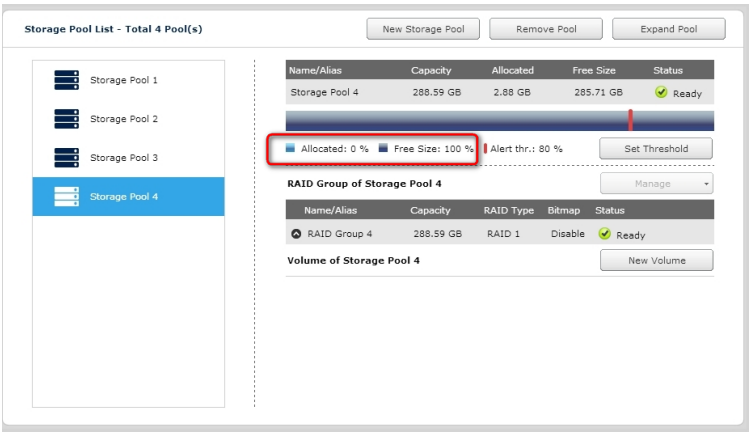

At this stage, we have 1 Storage Pool (285 GB Capacity), 1 datastore (250 GB):

On the NAS side, we can see below that only the real used data has been allocated. Because the datastore is empty, the Storage Pool on the NAS has almost no capacity allocated, and still has all capacity free that could be shared for other usages.

The next step is to store some data into the datastore.

Now that the test datastore is created, we are going to clone multiple VMs (7 in total) in the new datastore called "QNAP-Thin-Reclaim-Test", and see how the storage usage evolves.

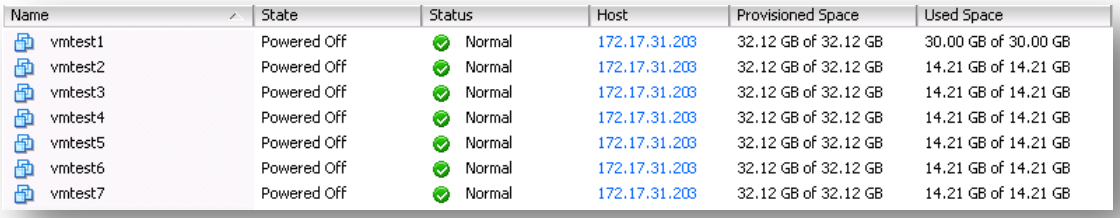

Name	State	Status	Host	Provisioned Space	Used Space
vmtest1	Powered Off	Normal	172.17.31.203	32.12 GB of 32.12 GB	30.00 GB of 30.00 GB
vmtest2	Powered Off	Normal	172.17.31.203	32.12 GB of 32.12 GB	14.21 GB of 14.21 GB
vmtest3	Powered Off	Normal	172.17.31.203	32.12 GB of 32.12 GB	14.21 GB of 14.21 GB
vmtest4	Powered Off	Normal	172.17.31.203	32.12 GB of 32.12 GB	14.21 GB of 14.21 GB
vmtest5	Powered Off	Normal	172.17.31.203	32.12 GB of 32.12 GB	14.21 GB of 14.21 GB
vmtest6	Powered Off	Normal	172.17.31.203	32.12 GB of 32.12 GB	14.21 GB of 14.21 GB
vmtest7	Powered Off	Normal	172.17.31.203	32.12 GB of 32.12 GB	14.21 GB of 14.21 GB

On VMware side:

- **115 GB used on the new datastore**
- 225 GB Provisioned on the datastore
- 133 GB free space in the datastore

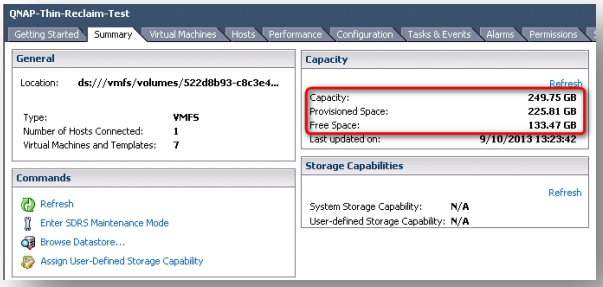

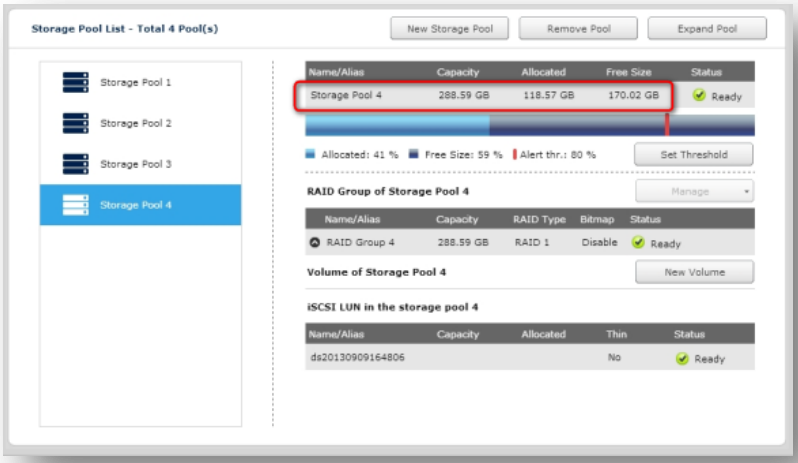

On the NAS:

- **118 GB allocated**
- 170 GB free space

At this stage, we have 1 Storage Pool (285 GB Capacity), 1 datastore (250 GB), and some VM in the datastore:

Thin Provisioning:

On NAS side, we can see that only the real used data has been allocated. 170 GB free capacity remains to be used for other services or LUN.

115 GB has been used in VMware and 118 GB allocated on the NAS (the small difference is due to metadata.)

If some data is removed within the datastore, by deleting 4 VMs for example, we can see below that the free space in the datastore increases and the used space decreases:

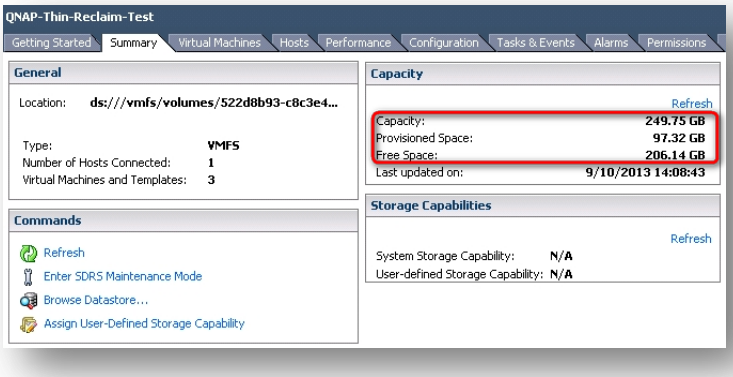

However, if we check the capacity on the NAS, the free capacity didn't change. The explanation is that thin provisioning allows dynamic capacity allocation, but, without any other function to get it back. With storages that supports only thin provisioning, once the capacity as been allocated, it cannot be freed. QNAP NAS provides a solution for this.

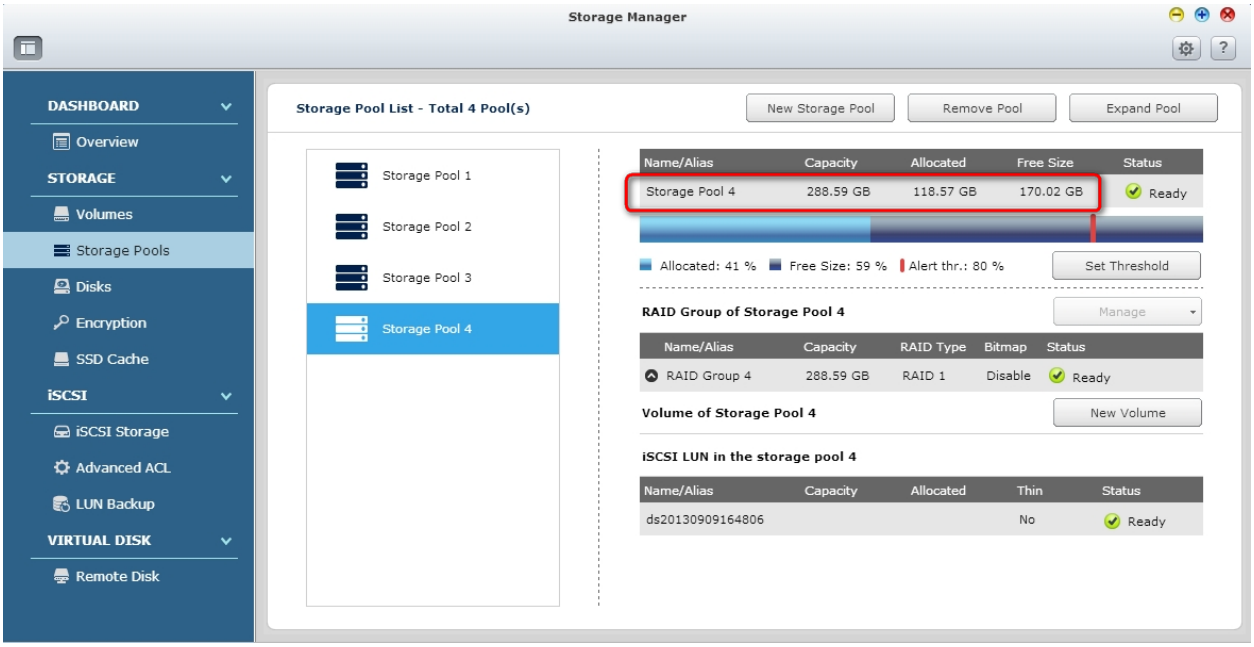

QNAP Turbo NAS with QTS 4.0 provides Space Reclamation function, supported by QTS and by VMware VAAI. Using this function it is possible to get back the capacity not used anymore, coming from deleted VMDK and deleted VMs from a datastore.

To trigger the space reclamation on iSCSI LUN with VAAI, it must be done by command line. vSphere Client do not provide any option from the graphical interface to trigger this function.
Moreover, it is not possible to estimate in advance the time needed to perform the space reclamation, as it depends on the number and size of deleted files, and the size of the datastore.
It is recommended to use space reclamation during maintenance timeframe, or out of business hours.

The space reclamation is triggered with the command "vmkfs", which takes an argument of the percentage of capacity to reclaim.
Examples:

```
~ # cd /vmfs/volumes/QNAP-Thin-Reclaim-Test/
/vmfs/volumes/522d8b93-c8c3e424-065f-000e1e0a63f8 # vmkfstools -y 99
Attempting to reclaim 99% of free capacity 206.1 GB (204.1 GB) on VMFS-5 file system 'QNAP-Thin-Reclaim-Test' with max file size
64 TB.
Create file .vmfsBalloon0kovyE of size 204.1 GB to reclaim free blocks.
Done.
/vmfs/volumes/522d8b93-c8c3e424-065f-000e1e0a63f8 #
```

Note: VMware recommend reclaiming 60% of the volume only.

After using the command above on the test datastore we can observe the result from the NAS UI:

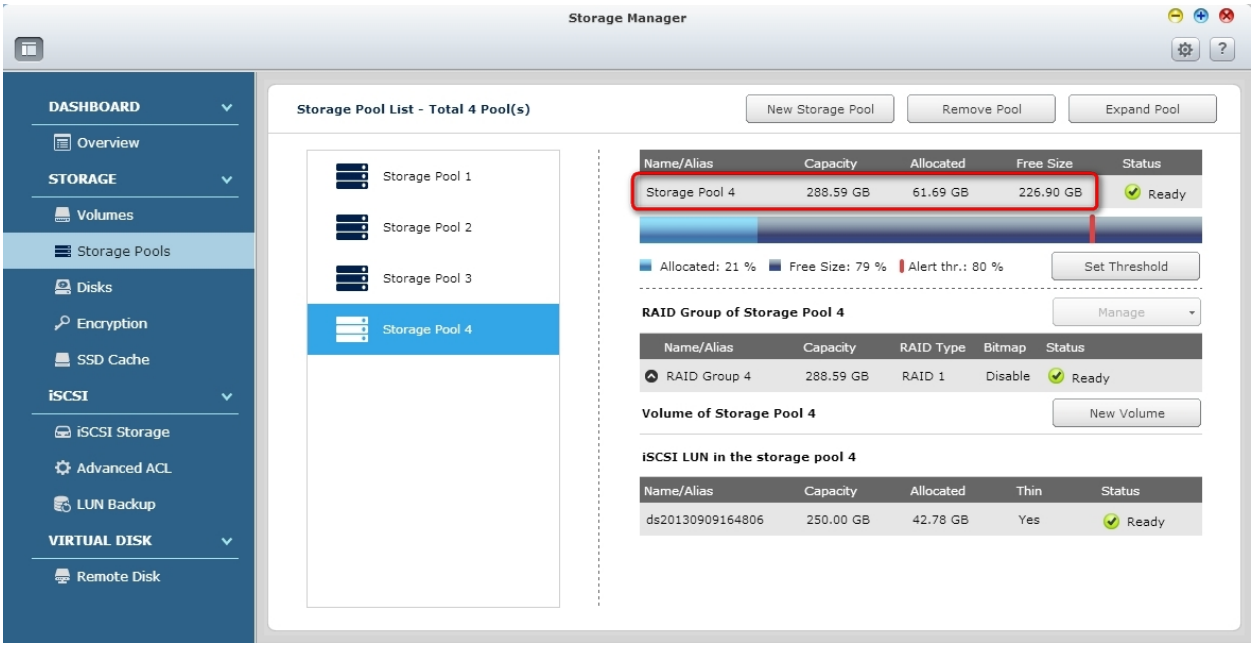

Before Space Reclaim :
41% allocated

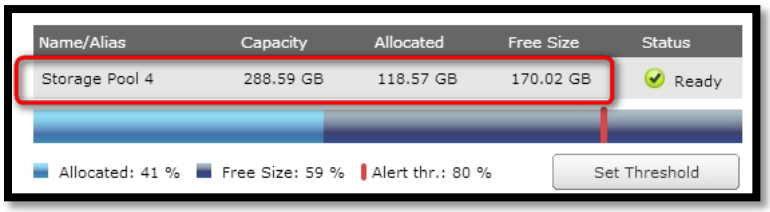

After Space Reclaim :
21 % allocated

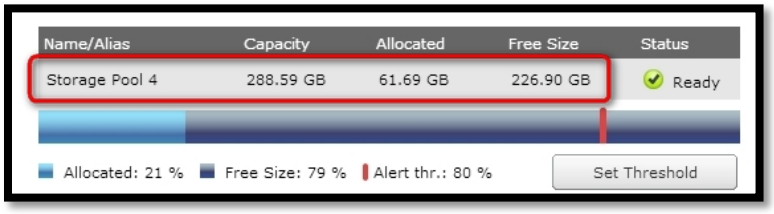

After using the Space Reclaim function, the capacity reclaimed in the NAS Storage Pool corresponds to the VM we have deleted from the datastore, around 56 GB (4 VMs of 14 GB).

Space Reclaim is not necessary when only 1 datastore is created on 1 NAS Storage Pool. It becomes useful when there are multiple datastores on a Storage Pool. All the datastores can now share the same storage pool. When needed, the Space Reclaim can be called to free the Storage Pool capacity. The Capacity that has been reclaimed can be used for other datastore if a new datastore needs to be created, can be used for new VM creation, or even other services if the QNAP NAS (service different than iSCSI).

One reason of using Thin Provisioning and Space Reclaim is to over-allocate the storage.
On a 285 GB Storage Pool, it is possible to create 2 or more iSCSI LUNs if they are using Thin Provisioning:

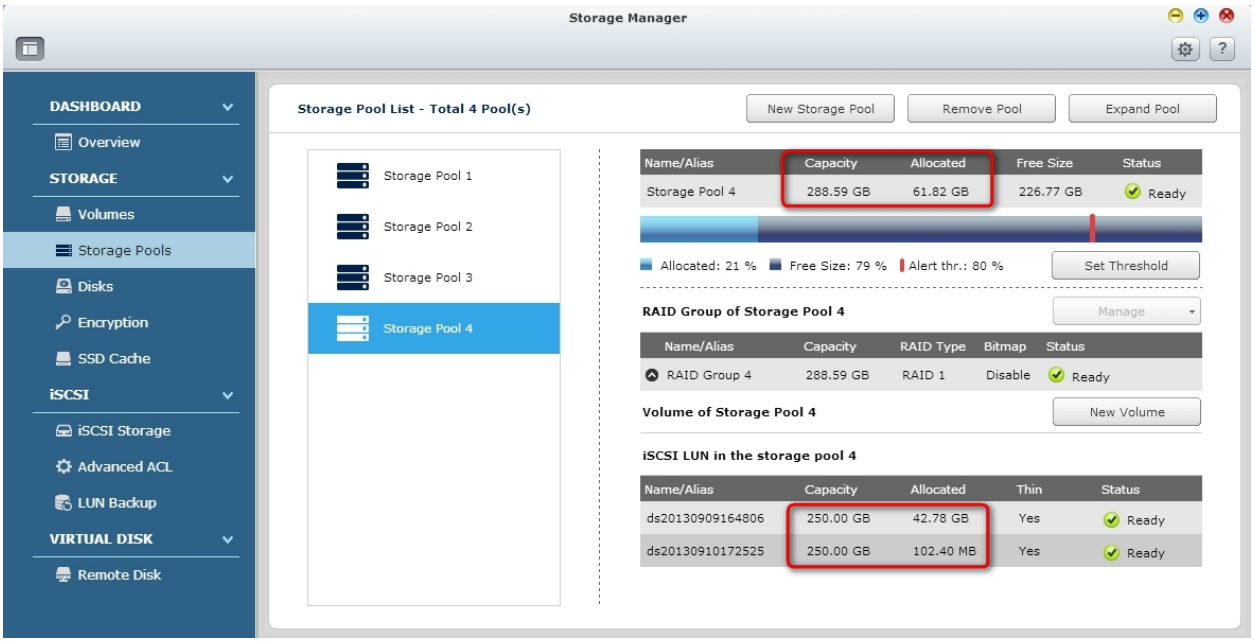

With Thin-Provisioning and Space Reclamation, the 2 LUNs above will be able to grow as data is stored.

In conclusion:

Thin Provisioning allows "over-allocation": It is possible to create iSCSI LUN larger than the real available capacity (provided by the Storage Pool).

Over-Allocation allows being flexible and scalable: With Thin Provisioning, the capacity will be used only when real data is stored in the LUN. Therefore, it is possible to create in advance iSCSI LUNs even if the capacity is not totally. When more capacity is needed because of data growth, it is possible to expand the Storage Pool online, with new hard drives for a scalable storage. It is also possible to create additional LUNs and datastores for temporary usage with Thin Provisioning, and remove them later.

Space Reclaim allows being flexible and efficient: When the Storage Pool is running out of space, it is also possible to reclaim the space with Space Reclaim, from the deleted data from VMware, and/or the deleted VMs. In that case, it may be not necessary to purchase additional hard drives. The storage management becomes more efficient.

But be careful of **not running out of space**! Over-allocation is a great, flexible and efficient way to manage the storage, but be aware that if not monitored correctly, the Storage Pool can run out of space. It is necessary to set-up an alert threshold on the storage pool with email notifications so that the administrator can know in advance the remaining capacity of the storage pool and decide what action to take.

4. Good to know

Thin volumes/LUNs offer better and more efficient capacity management.
Thick volumes avoid the out-of-space condition risk.
VAAI NAS and VAAI iSCSI are different. Clone a VM from NFS to iSCSI will not use VAAI, but ESXi resources and network.

5. References

Disabling the VAAI functionality in ESXi/ESX (1033665):
http://kb.vmware.com/kb/1033665
Installing patches on an ESXi 5.x host from the command line:
http://kb.vmware.com/kb/2008939
VMware vSphere® Storage APIs – Array Integration (VAAI):
http://www.vmware.com/files/pdf/techpaper/VMware-vSphere-Storage-API-Array-Integration.pdf
VMware vSphere Blog - VAAI Thin Provisioning Block Reclaim/UNMAP In Action:
http://blogs.vmware.com/vsphere/2012/04/vaai-thin-provisioning-block-reclaimunmap-in-action.html