



Saving capacity by using Thin provisioning, Deduplication, and Compression In Qsan Unified Storage

**U400Q Series
U600Q Series**

**Version 1.0.0
November 2012**

Copyright

Copyright@2004~2012, Qsan Technology, Inc. All rights reserved. No part of this document may be reproduced or transmitted without written permission from Qsan Technology, Inc.

Trademarks

All products and trade names used in this manual are trademarks or registered trademarks of their respective companies.

Qsan Technology, Inc.

4F, No.103, Ruei Hu St
Neihu Dist, Taipei 114
Taiwan, R.O.C.

Tel: +886-2-7720-2118
Fax: +886-2-7720-0295

Email: sales@QsanTechnology.com
Website: www.QsanTechnology.com

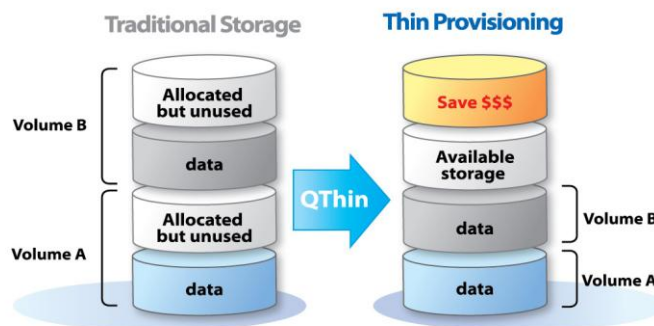
Qsan storage efficiency technology

As all kinds of digital and mobile devices as well as social media become more popular, companies and individuals keep generating more and more data. How to store these massive amount of data efficiently to reduce cost and energy footprint is the top priority for most SMB companies and international enterprises.

Qsan technology brings the sophisticated and expensive functions such as deduplication, compression, and thin provisioning that would normally be offered in mid-range enterprise storage to the SMB price segment with NO extra charges. This application note will show you how to use these features and help you achieve the storage efficiency you need.

Thin provisioning

Thin provisioning uses the concept of “allocate-on-demand” or “just-in-time capacity”. There is no pre-allocated space. Space is only allocated when data is copying to the storage. No more stranded space exists. You can use exactly the same amount of space to store the data.



How to use thin provisioning with file-system

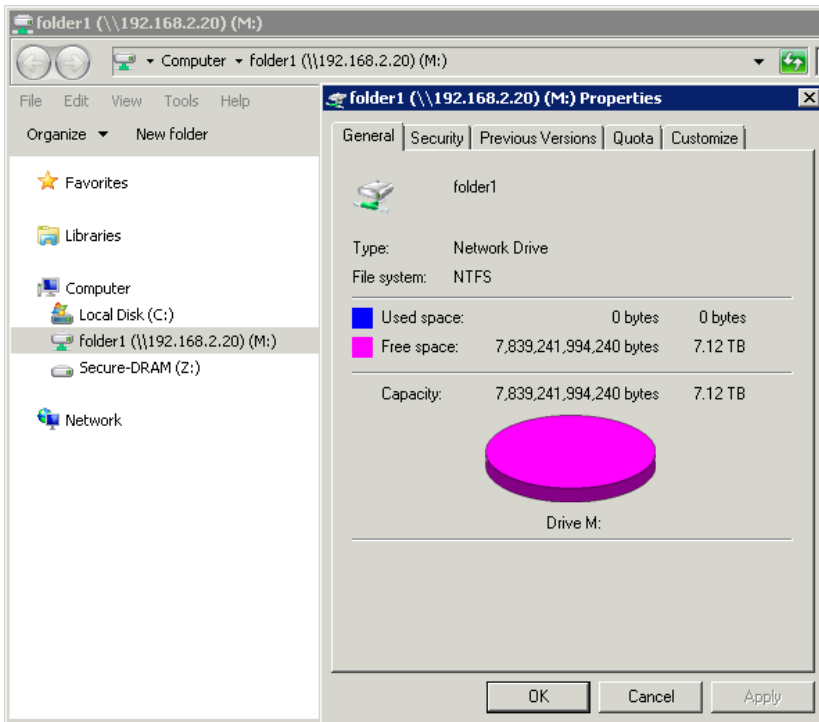
The way to set thin provisioning for file-system and volume is different. To create a file-system with thin provisioning, there is no need to specify the size. The size input will grey out if thin provisioning is checked.

The size of the file-system is the same as the size of the storage pool, which means it can use up to the full capacity of the storage pool if necessary.

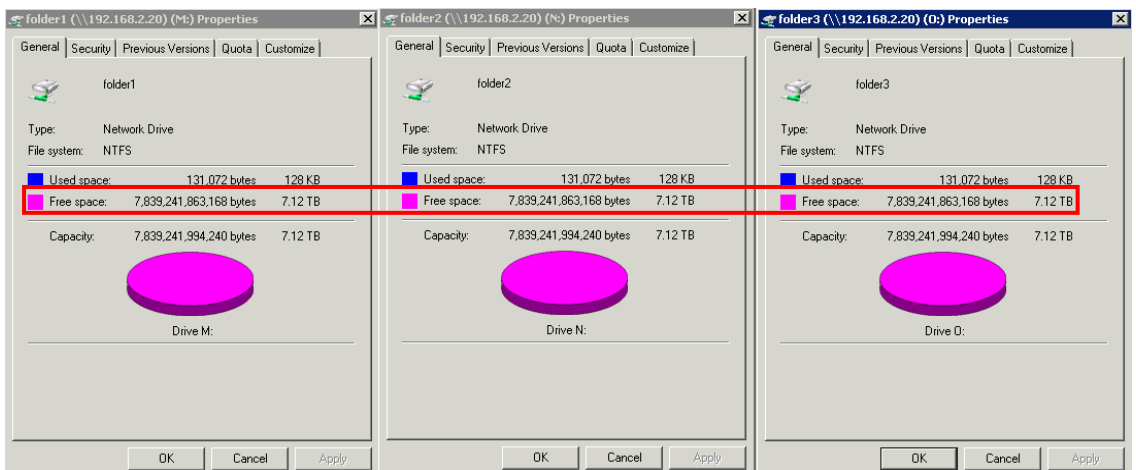
Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Total (GB)	Used (GB)	Free (GB)	Capacity	Dedup	Status
Pool1	7300.86	0	7300.86	0%	1.00x	Online

Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Type	Quota (GB)	Reserved (GB)	Used (GB)	Free (GB)	
Pool1/folder1	filesystem	none	none	0	7300.86	

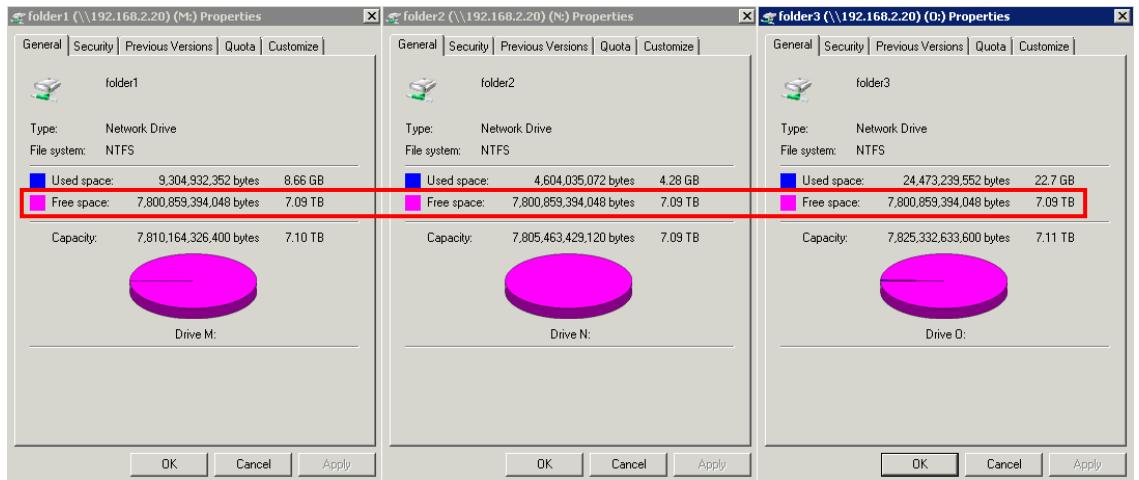
If you share out this file-system as a network drive, the size of this network drive is the same as that of the storage pool.



You may create several file-systems for different applications with thin provisioning. All of them will share the free space of the storage pool based on a first-come-first-served policy. When the storage pool is used up, you can expand the storage pool capacity by adding extra RAID sets without disrupting any data services.

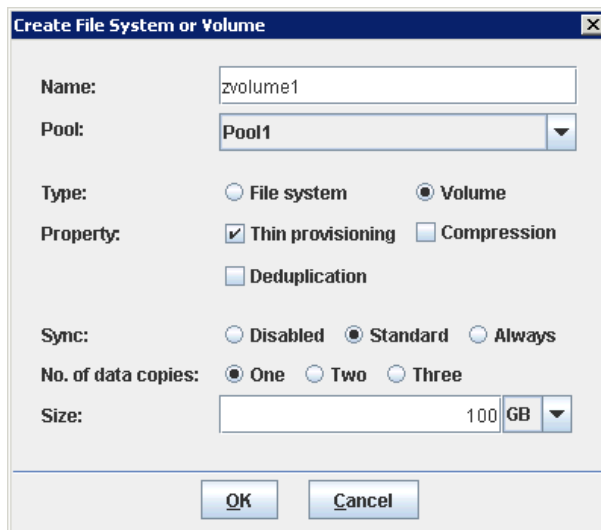


For example, copy different amount of data to each network drive and verify the result. You can see that all of them have the same “Free space” size, although the “Capacity” part is different.



How to use thin provisioning with iSCSI volume

When creating an iSCSI volume with thin provisioning, you need to specify the size of the volume. This number can be larger than the size of the free space left in the storage pool. After the volume is created, there is no actual space set aside. Space is allocated by on-demand basis.

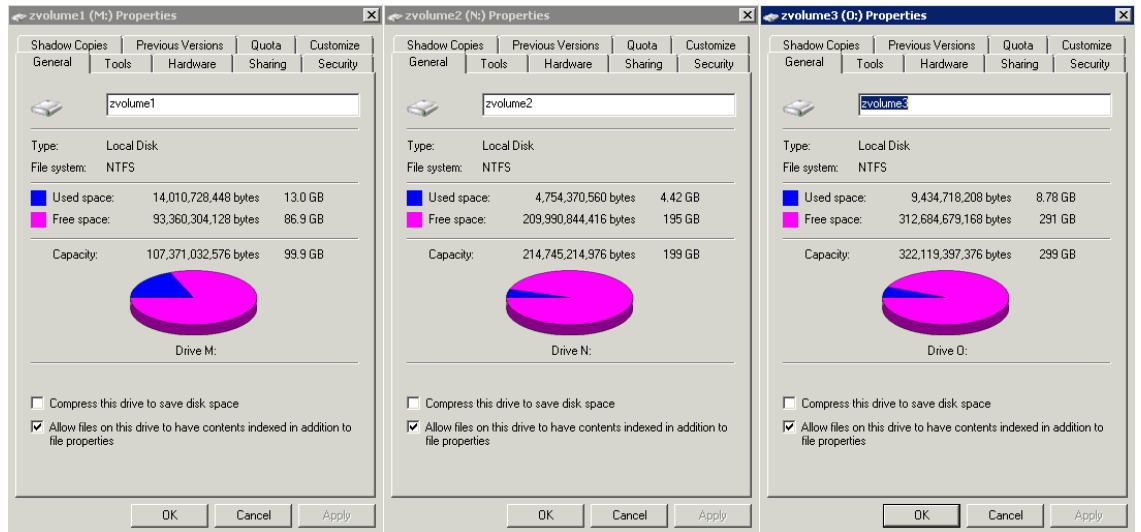


Let’s create 3 iSCSI volumes with thin provisioning and see how the space is allocated.

Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Total (GB)	Used (GB)	Free (GB)	Capacity	Dedup	Status
Pool1	7300.86	0	7300.86	0%	1.00x	Online

Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Type	Quota (GB)	Reserved (GB)	Used (GB)	Free (GB)	Dedup
Pool1/zvolume1	volume	100.00	none	0	7300.86	off
Pool1/zvolume2	volume	200.00	none	0	7300.86	off
Pool1/zvolume3	volume	300.00	none	0	7300.86	off

Let's use iSCSI initiator to mount each volume as a drive and copy different amount of data to each drive.



Checking the actual "Used space". It's exactly the total amount of data copied to each drive about 26GB. Although we "allocate" 600GB space for three drives, but the available free space in the storage pool is still around 7274GB as it would have been 6700GB without using thin provisioning. So more available space can be allocated to other applications.

Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Total (GB)	Used (GB)	Free (GB)	Capacity	Dedup	Status
Pool1	7300.86	26.18	7274.67	0%	1.00x	Online





TIP

The examples shown here are using Windows 7 or Windows operating system as client machine connecting Qsan unified storage. The same behaviors and functions will apply to other operating systems such as Linux, Mac and Unix.

Deduplication

Qsan deduplication technology provides inline, block-level, efficient redundancy removing function. Block level deduplication is generally more efficient than file level one. Inline means that as the data is copying into the storage, the deduplication is processing the data in real time.





The greater the data redundancy is, the more the deduplication can help.

How to verify deduplication in action?

Qsan deduplication technology is enabled on a per file-system and volume basis. It can be turned on and off on the fly depending on user's needs. The space that is saved by deduplication will NOT be reflected in the share or iSCSI volume directly. This saved space will be accumulated in the free space of the storage pool. So you can use these "growing" space for other applications.

Let's check the fresh start of a new storage pool. The total space is 7300GB as shown in the screenshot below.

Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Total (GB)	Used (GB)	Free (GB)	Capacity	Dedup	Status
Pool1	7300.86	0	7300.86	0%	1.00x	Online

Let's create a 100GB file-system and share it out. Without using thin provisioning, 100GB is reserved for this file-system share and the free space of the pool becomes 7200GB.

Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Total (GB)	Used (GB)	Free (GB)	Capacity	Dedup	Status
Pool1	7300.86	100.00	7200.86	0%	1.00x	Online

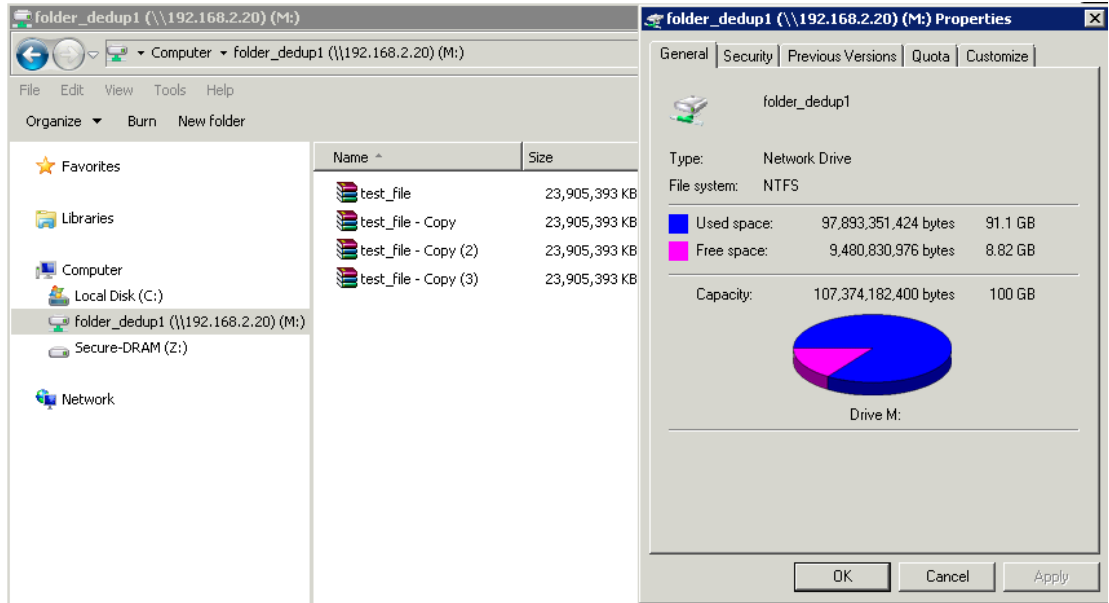
Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Type	Quota (GB)	Reserved (GB)	Used (GB)	Free (GB)	Dedup
Pool1/folder_dedup1	filesystem	100.00	100.00	0	99.99	on

Let's copy the same file repeatedly 4 times to this share. With deduplication enabled, the second and third copies will not take up space. But where does the saved space go? The total space of the storage pool becomes 7368GB. The saved space goes to the storage pool as free space.


Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Total (GB)	Used (GB)	Free (GB)	Capacity	Dedup	Status
Pool1	7368.12	100.08	7268.03	0%	4.00x	Online

Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Type	Quota (GB)	Reserved (GB)	Used (GB)	Free (GB)	Dedup
Pool1/folder_dedup1	filesystem	100.00	100.00	91.16	8.82	on

The free space of the storage pool grows 68GB larger, which is the saved space by using deduplication or the size of 3 times the copy file. In fact, the physical space doesn't "grow". Because in the file-system table it shows that 91GB space is used. Use this 68GB to offset the 91GB used space. The net effect is the size of ONE copy file instead of four.

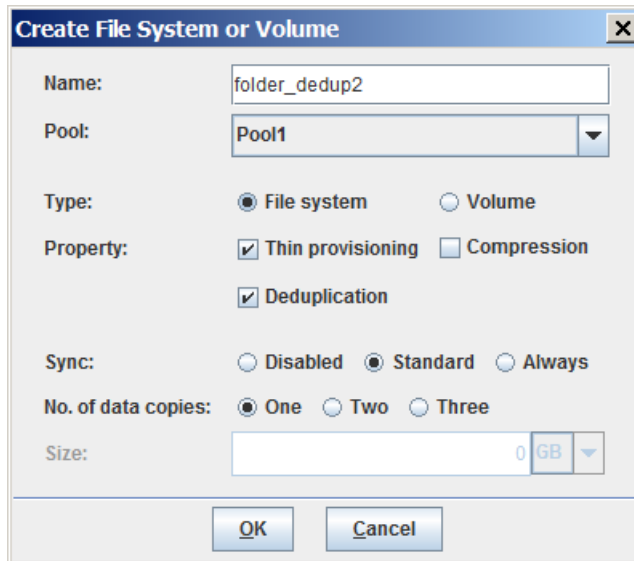


From the client OS viewpoint, you don't feel the deduplication in action. You still see 4 files that take up 96GB space as it normally would do. The free space doesn't get any bigger. Try to duplicate the file again and you will get insufficient space warning. Apparently you don't get the immediate benefit from using deduplication this way.



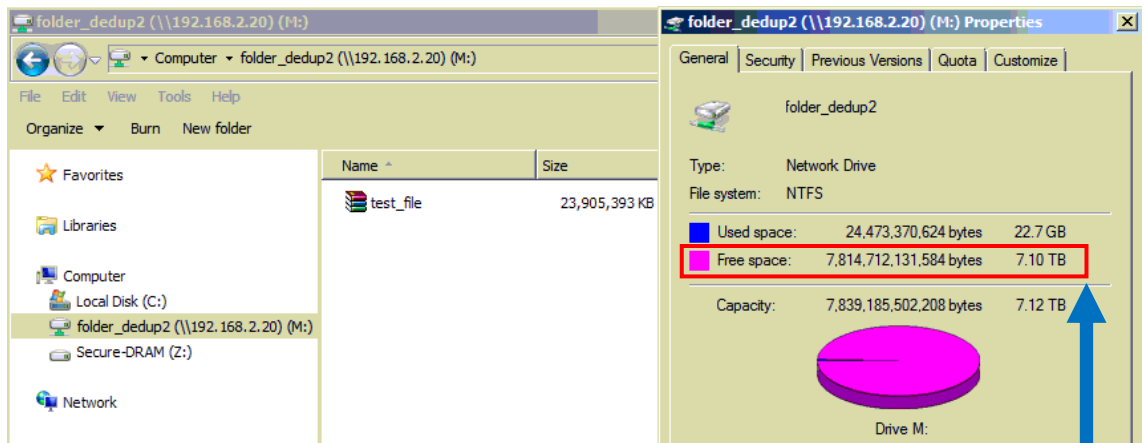
How to use deduplication with file-system?

If you want to use a share with deduplication, please use it with thin provisioning. Create a file-system with both deduplication and thin provisioning. There is no need to specify the size as explained earlier.

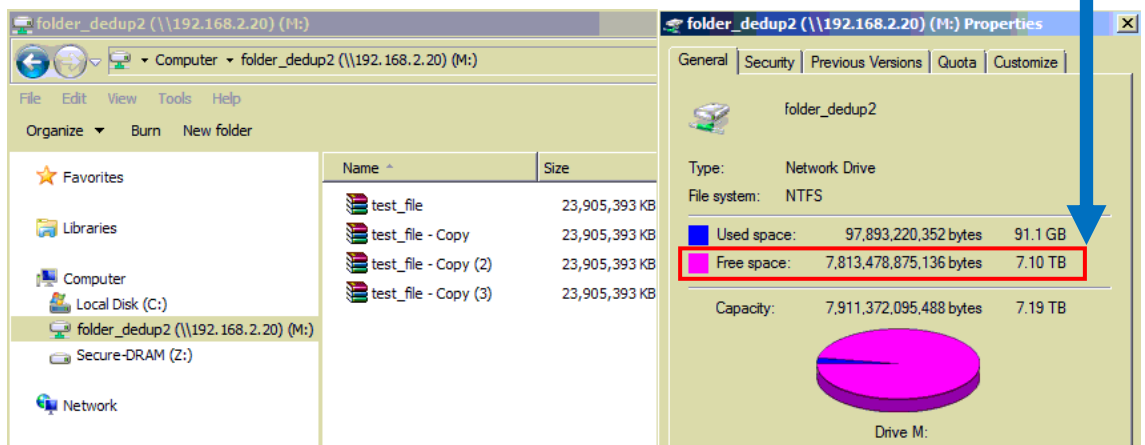


Name	Type	Quota (GB)	Reserved (GB)	Used (GB)	Free (GB)	Dedup
Pool1/folder_dedup2	filesystem	none	none	0	7300.86	on

Again, copy the same file over. Check the result in client OS.



Duplicate the file three times and check the result again.



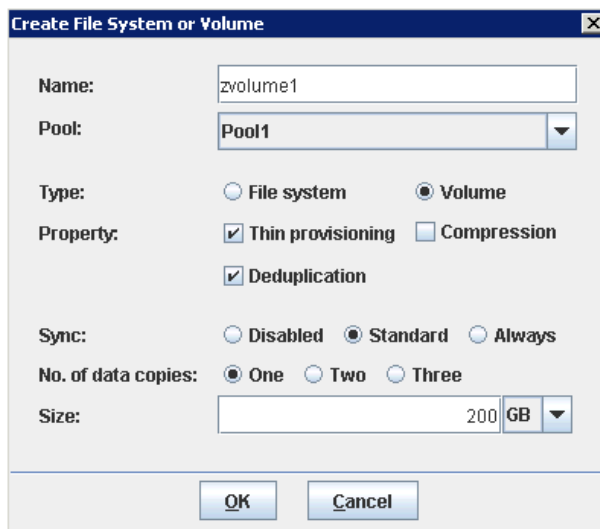
The size of free space doesn't decrease. It means that those three duplicate copies do NOT take up too much space. They are deduplicated!

How to use deduplication with iSCSI volume

When using iSCSI volume with deduplication enabled, you will not see dedup ratio going up immediately after duplicating the file several times. This is because the iSCSI volume presents itself as a raw drive to the client OS. Client OS will format the raw drive to the client file system (such as EXT3, EXT4, NTFS, FAT32) before it can store files onto the iSCSI volume.

Because the data saved in iSCSI volume is processed by client file system. Although the original files are the same, when they are saved in the storage pool in raw block format, the block content become different. This is why the dedup ratio doesn't increase when you duplicate data copies in the same iSCSI volume.

To use storage pool more efficiently, enabling both deduplication and thin provisioning for iSCSI volume is recommended. Please refer to earlier paragraph – How to use thin provisioning with iSCSI volume. The way it works is basically the same.



The screenshot shows a Windows-style dialog box titled "Create File System or Volume". It contains the following fields and options:

- Name:** zvolume1
- Pool:** Pool1 (dropdown menu)
- Type:** File system Volume
- Property:** Thin provisioning Compression Deduplication
- Sync:** Disabled Standard Always
- No. of data copies:** One Two Three
- Size:** 200 GB (dropdown menu)

Buttons for "OK" and "Cancel" are at the bottom.



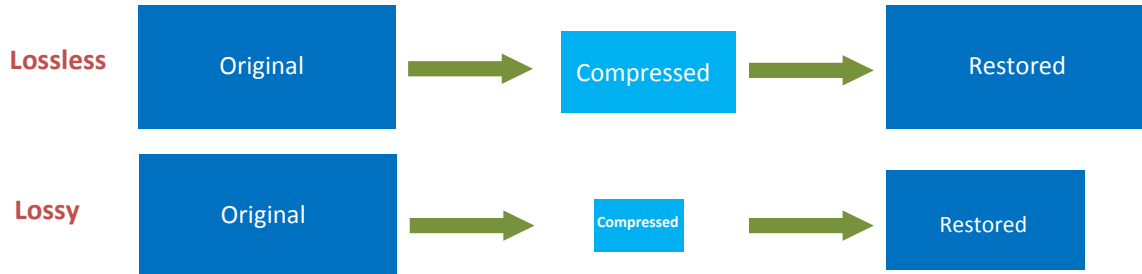
TIP

The examples shown here are using Windows 7 or Windows operating system as client machine connecting Qsan unified storage. The same behaviors and functions will apply to other operating systems such as Linux, Mac and Unix.

Compression

The compression algorithm used in Qsan unified storage is LZJB, a lossless, lightweight and fast data compression algorithm. The compression function is running real time, on the fly to compress

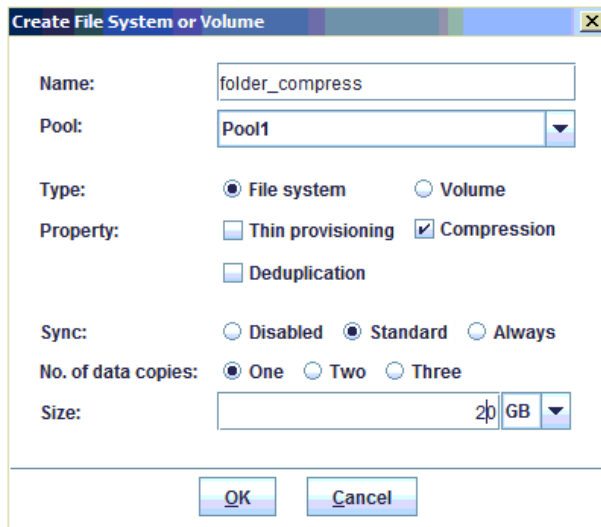
the data written into the storage. You may turn off the compression function at a later time. The new data will not be compressed and this won't affect the old data that remain in compressed format. The lossless algorithm means the compression process is reversible and all original data content will be kept intact. The primary purpose of compression is to save space. Generally speaking, raw data like pure text file, log file or image bitmap file have the best compression ratio.



In the demonstration below, we will use bitmap file (*.bmp) and iobw.tst file with good compression ratio to show you how the compression is in action.

How to use compression with file-system

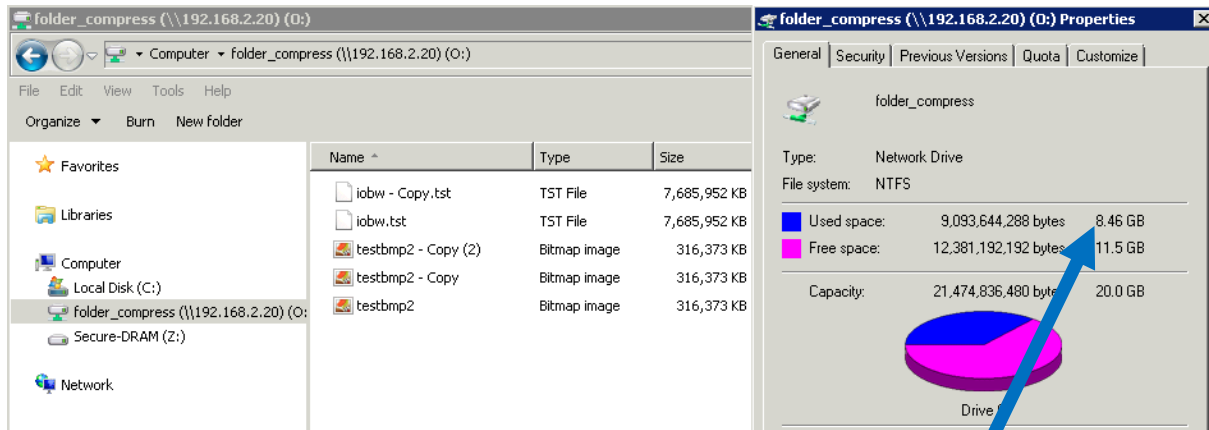
Using compression with a share is probably the most intuitive way to see and get the space saving benefits. To demonstrate this, we create a 20GB file-system and share it out.



Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Total (GB)	Used (GB)	Free (GB)	Capacity	Dedup	Status
Pool1	7300.86	20.00	7280.86	0%	1.00x	Online

Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule	
Name	Type	Quota (GB)	Reserved (GB)	Used (GB)	Free (GB)	Dedup	Compress
Pool1/folder_compress	filesystem	20.00	20.00	0	19.99	off	on

Copy some files with good compression ratio and check the result below.



Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Total (GB)	Used (GB)	Free (GB)	Capacity	Dedup	Status
Pool1	7300.86	20.00	7280.86	0%	1.00x	Online

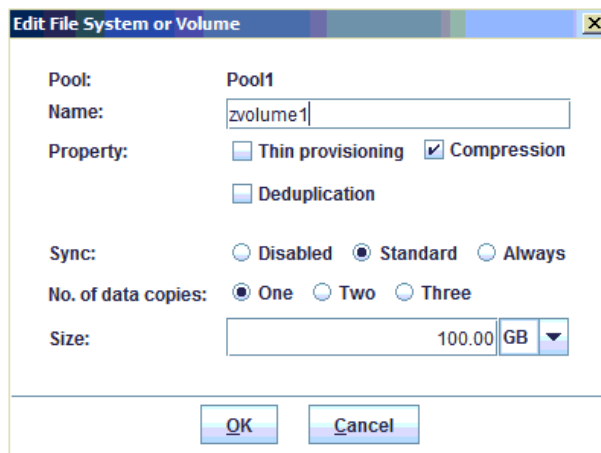
Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule	Used (GB)	Free (GB)	Dedup	Compress
Name	Type	Quota (GB)	Reserved (GB)	Used (GB)	Free (GB)	Dedup	Compress			
Pool1/folder_compress	filesystem	20.00	20.00	8.46	11.53	off	on			

The nominal size of all the files is about 16GB. However, the client OS reports only 9GB are used. In the end, user can store more data than the network drive is allowed to.

You may use both compression and thin provisioning in creating a file-system and share it out. All the files copied to the share will be compressed and the saved space goes to the storage pool that can be shared by other applications. Both storage efficiency and utilization rate are increased.

How to use compression with iSCSI volume

Create an iSCSI volume of 100GB in size with only compression enabled.



Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Total (GB)	Used (GB)	Free (GB)	Capacity	Dedup	Status
Pool1	7300.86	100.00	7200.86	0%	1.00x	Online

Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule	
Name	Type	Quota (GB)	Reserved (GB)	Used (GB)	Free (GB)	Dedup	Compress
Pool1/zvolume1	volume	100.00	100.00	0.00	7300.85	off	on

Let's copy iobw.tst five times to the volume and check the result.

Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Total (GB)	Used (GB)	Free (GB)	Capacity	Dedup	Status
Pool1	7300.86	100.00	7200.86	0%	1.00x	Online

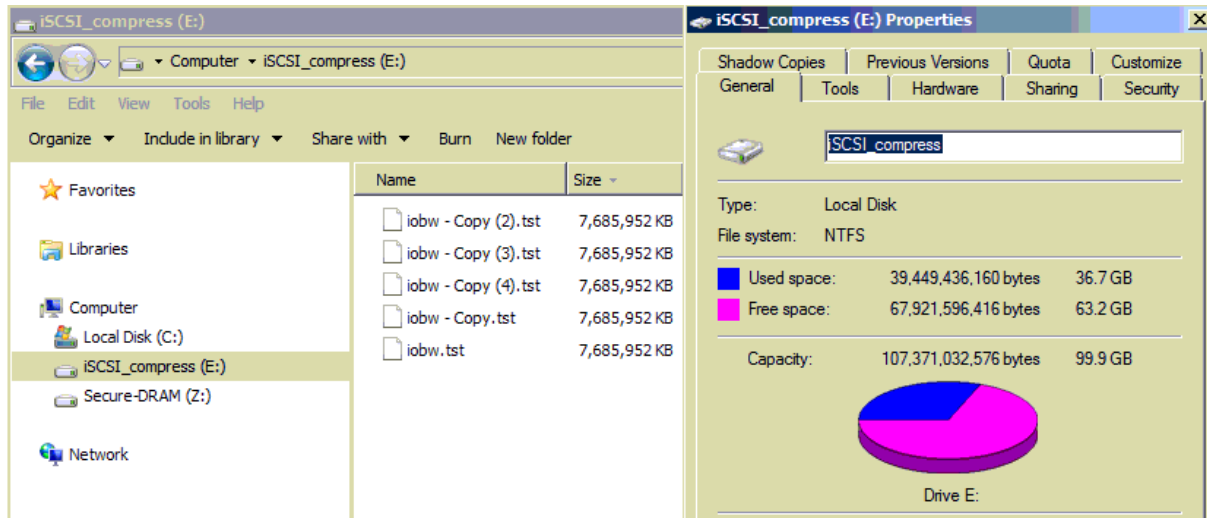
Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule	
Name	Type	Quota (GB)	Reserved (GB)	Used (GB)	Free (GB)	Dedup	Compress
Pool1/zvolume1	volume	100.00	100.00	20.67	7280.18	off	on

The client OS doesn't recognize any benefit from compression. The overall used space reports 36.7GB, which is exactly 5 times the size of iobw.tst without compression. However, if you check the used space of iSCSI volume in NAS, it only takes 20.6GB. This means compression is working and space is saved. But it's just that the saved space is confined in the volume and we have NO ACCESS to it. Because if you check the free space in storage pool (Pool1), it is still the same as 7200GB.

Therefore, we DO NOT recommend you use iSCSI volume with compression this way. So, how can we move the saved space to the free space of the storage pool? You are correct – using thin provisioning!

Use compression and thin provisioning with iSCSI volume

Repeat the same thing above to create an iSCSI volume (100GB in size) with both compression and thin provisioning. And then copy the file (iobw.tst) five times to the volume.



Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Total (GB)	Used (GB)	Free (GB)	Capacity	Dedup	Status
Pool1	7300.86	20.67	7280.18	0%	1.00x	Online

Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule	
Name	Type	Quota (GB)	Reserved (GB)	Used (GB)	Free (GB)	Dedup	Compress
Pool1/zvolume1	volume	100.00	none	20.67	7280.18	off	on

By adding thin provisioning to the equation, now the saved space is accumulated in the free space and can be shared with other application usage. The idea behind this is the same as what we've described earlier in creating iSCSI volume using deduplication with thin provisioning.



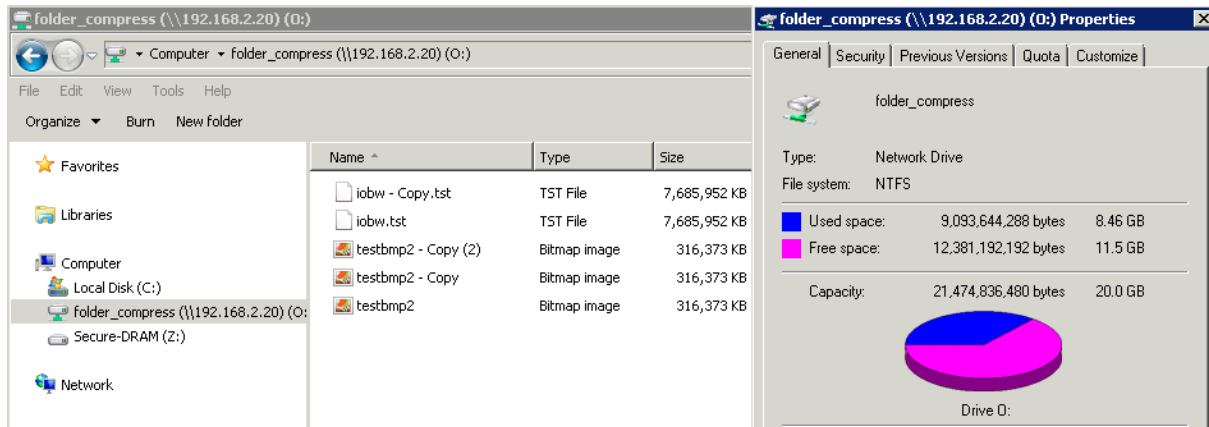
TIP

The examples shown here are using Windows 7 or Windows operating system as client machine connecting Qsan unified storage. The same behaviors and functions will apply to other operating systems such as Linux, Mac and Unix.

Deduplication plus Compression

Using both deduplication and compression at the same time can give you combined effect. You can get the most storage saving out of the equation. Let's use the first example in compression section above.

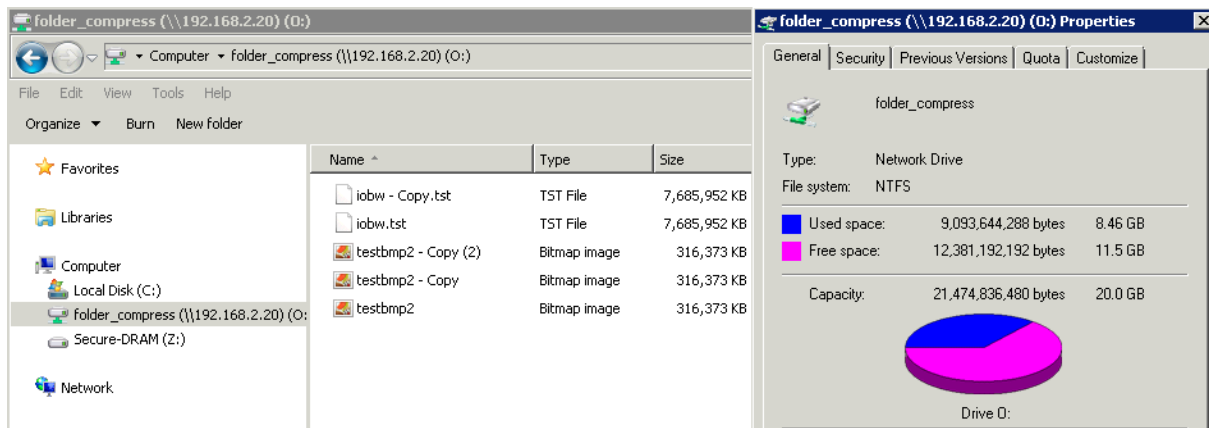
Repeat the same thing with only compression (without deduplication). You get the following data.



Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Total (GB)	Used (GB)	Free (GB)	Capacity	Dedup	Status
Pool1	7300.86	20.00	7280.86	0%	1.00x	Online

Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule	
Name	Type	Quota (GB)	Reserved (GB)	Used (GB)	Free (GB)	Dedup	Compress
Pool1/folder_compress	filesystem	20.00	20.00	8.46	11.53	off	on

Restart the whole thing with both compression and deduplication. Below is the result.



Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule
Name	Total (GB)	Used (GB)	Free (GB)	Capacity	Dedup	Status
Pool1	7309.00	20.00	7288.99	0%	43.27x	Online

Physical disk	S.M.A.R.T.	Pool	ZFS	LUN	Snapshot	Snapshot schedule	
Name	Type	Quota (GB)	Reserved (GB)	Used (GB)	Free (GB)	Dedup	Compress
Pool1/folder_compress	filesystem	20.00	20.00	8.46	11.53	off	on

By enabling deduplication at the same time, extra 8GB space is saved. These space saving features don't come cheap without any cost and sacrifice. The performance will drop a little in exchange to use these enterprise storage features. If performance is not the first priority, using all three features (deduplication, compression and thin provisioning) at the same time can give you the most storage saving and achieve the highest storage efficiency.

Summary

Examples and demonstrations shown so far are using small amount of data for easy understanding and comparison. When running a full scale project, both deduplication and compression can save substantial space depending on the applications and the data types. Applications such as full backup, server virtualization, or Exchange server database will be suitable to use deduplication. For storing huge amount of pure text files, raw image files, or raw multimedia files, using compression is a better choice given that those data types have a higher compression ratio.

In a nutshell, we summarize the following rules for your deployment reference.

- Using iSCSI volume, the size of the volume is limited by the ZFS quota. From the client OS viewpoint, the volume itself can not recycle or benefit from the saved space by using deduplication and compression.
- Always use thin provisioning along with deduplication and compression, the saved space will go to the free space of the storage pool and can be shared with other applications.
- Without thin provisioning, if you want to recycle and benefit from the saved space, a shared file-system with compression is the ONLY option.
- Using deduplication and compression to save storage space will have the consequence of lower performance. Please be aware!

Applications	Data services	Best configuration	Performance impact
File backup	CIFS/NFS/AFP	Thin provision +Compression +Deduplication	Medium
Mail server	CIFS/NFS	Thin provision +Compression +Deduplication	Medium
VM image	NFS	Thin provision +Deduplication	Medium

A simple comparison table to summarize the rules above.

To store **more** data in **current** share or volume using the saved space from dedupl and compression.

	Dedup	Compress	Dedup +Compress	Thin +Dedup	Thin +Compress	Thin+Dedup +Compress
Share(file-system)	No	Yes	Yes	Yes	Yes	Yes
Volume (iSCSI)	No	No	No	No	No	No

To increase **the free space of the storage pool** due to the saved space from dedup and compression.

	Dedup	Compress	Dedup +Compress	Thin +Dedup	Thin +Compress	Thin+Dedup +Compress
Share(file-system)	Yes	No	Yes	Yes	Yes	Yes
Volume (iSCSI)	Yes	No	Yes	Yes	Yes	Yes