



Qsan Document - White Paper

How to use QReplica 2.0

Version 2.0
March 2015

Copyright

Copyright@2004~2015, 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, Ruihu St.,
Neihu Dist., Taipei City 114,
Taiwan (R.O.C.)

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

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

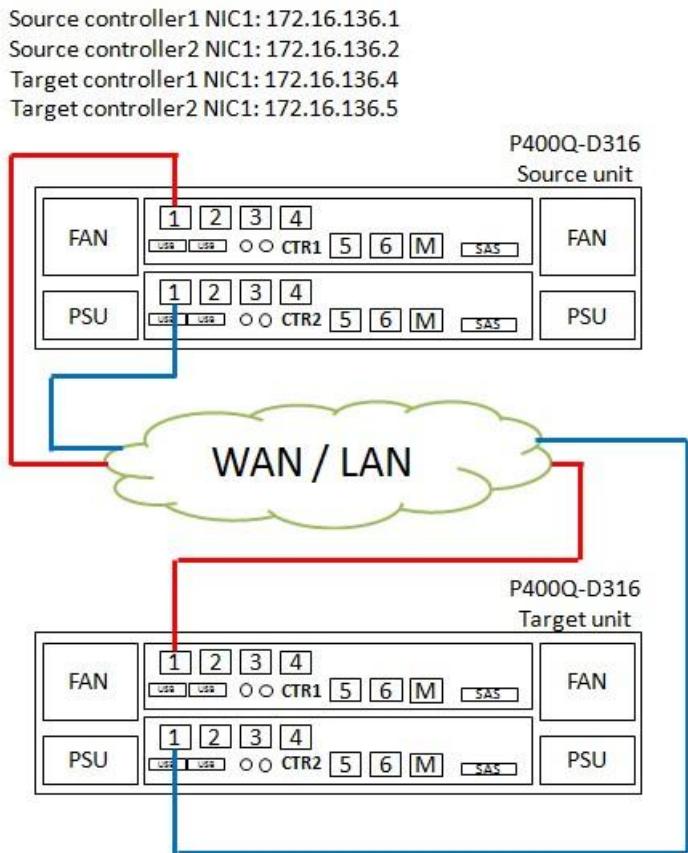
Introduction

QReplica2.0 is a new feature which added in the latest controller firmware. It helps user to replicate data from source SAN storage to target SAN storage via LAN or WAN more rapidly than previous QReplica function. QReplica2.0 has more powerful abilities, including unlimited transferring speed, traffic shaping during peak time, multiple connections for one replication job and etc. In this document, we will demonstrate how this feature will improve users' daily backup jobs.

Environment

Network Diagram

Below is a standard configuration of QReplica which provides the redundancy of replication job through controllers and network.



Prerequisites

Here are the prerequisites.

1. Make sure the controller firmware version of source and target SAN storage is above than V3.2.2.
2. Make sure there is enough space on the RIAD group of the **Source** virtual disk so that the snapshot space can be successfully allocated automatically after the replication job created, otherwise user has to allocate the snapshot space manually before creating the replication job.
3. Make sure the snapshot space on the RIAD group of **Target** virtual disk has been allocated before executing any replication job. The target SAN storage will automatically take a snapshot on the BACKUP VD after each replication job finishes.

Configuration

How to Create a Replication Job

The QReplica function copies data from the virtual disk of source SAN storage to the target over LAN or WAN. You may follow the steps as listed below to finish the creation of first replication job.

1. Make sure which virtual disk on the source SAN storage is the one you want to backup.
2. Create a virtual disk on the target SAN storage as a backup virtual disk.

3. The backup virtual disk will need to be attached with an iSCSI LUN ID so that it could be discovered by the source SAN storage and the replication job could be established successfully.

Attach LUN

Attach a logical unit number (LUN) to a virtual disk.

Virtual Disk: targetVD(200 GB) ▼

Allowed Hosts: * Add Host

Target: 0 ▼

LUN: -LUN 1- ▼

Permission: Read-only Read-write

the permission for this LUN will be read-only

OK Cancel

- Go to **Volume Configuration -> QReplicas** on the web interface of source SAN storage and create the first replication job via clicking **Create** button.

Physical Disks RAID Groups Virtual Disks Snapshots Logical Units **QReplicas**

Create Rebuild QReplica Options Shaping Setting Configuration

- In the pop-up window, select the virtual disk you want to backup data and then click **Next** button.

QReplica - Select a Source Virtual Disk


<< first < prev 1 next > last >>

Name	Size (GB)	Status	RAID	LUN #	Snapshot #	RAID Group
<input checked="" type="radio"/> sourceVD	100	Online	RAID 5	0	0	sourceRG

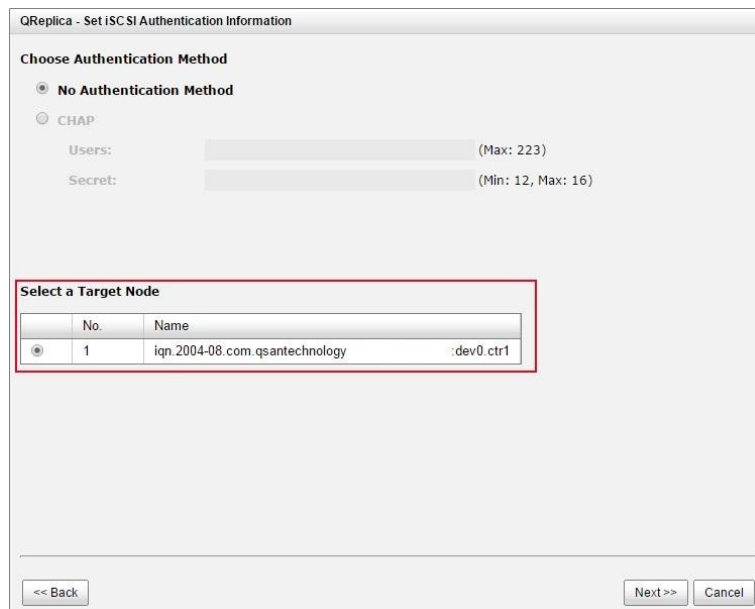
<< first < prev 1 next > last >>

Next>> Cancel

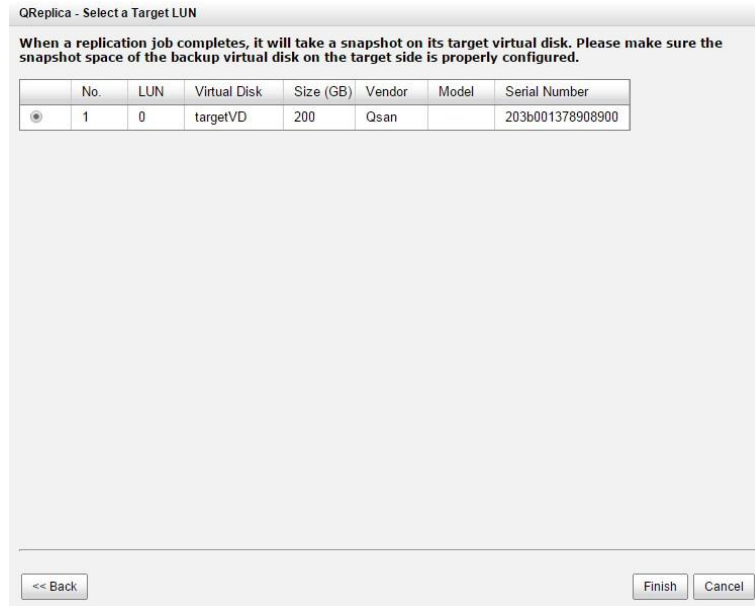
- Select the portal which you would like to use as the transferring port from the source SAN storage, and specify the IP address of the portal that you want to use as receiving port on the target SAN storage. If **Auto** is selected, system will choose one available portal automatically.



- Select the iSCSI target that just attached on the target SAN storage.



- Select the available virtual disk which attached to this target, and finish the creation of replication job.



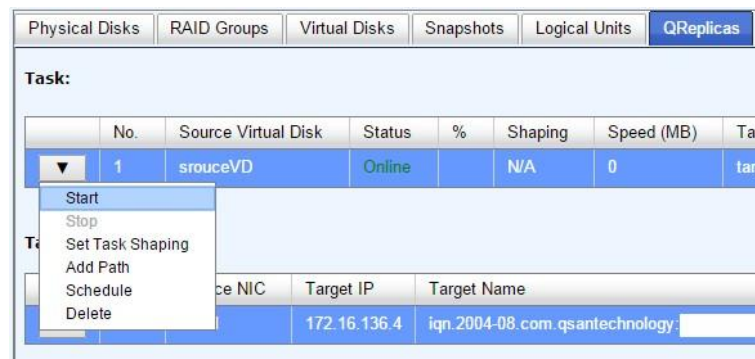
- The task will be listed and the path of session will be displayed for monitoring the status of replication job.

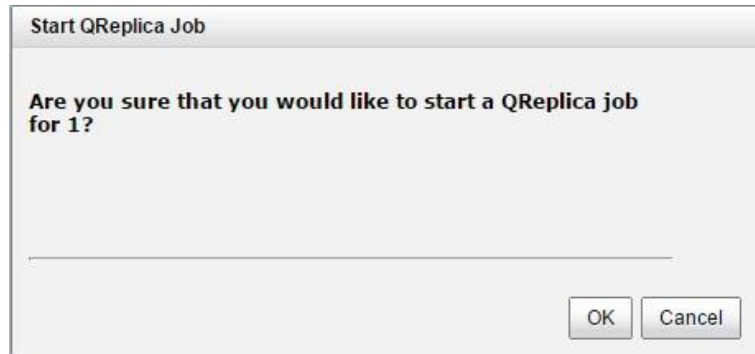


The Operation of Replication Job

Start the job manually

- Click on the ▼ -> **Start** to execute the replication job.





- You may see the status of task displays as **Queued** when the percentage is stay in **0%**, but actually the data transferring process has already begun at the moment.

Physical Disks								RAID Groups								Virtual Disks								Snapshots								Logical Units								QReplicas							
Task:																																															
No.	Source Virtual Disk	Status	%	Shaping	Speed (MB)	Ta																																									
▼ 1	srouceVD	Queued	0	N/A	0	ta																																									

- While the percentage becomes **4%** or more, the status will be displayed as **Replicating**, and there is transferring speed of the task displayed in the **Speed** column.


Physical Disks								RAID Groups								Virtual Disks								Snapshots								Logical Units								QReplicas							
Task:																																															
No.	Source Virtual Disk	Status	%	Shaping	Speed (MB)	Ta																																									
▼ 1	srouceVD	Replicating	4	N/A	99	ta																																									

- At the meantime there will be a snapshot taken on the source virtual disk when the replication job starts.

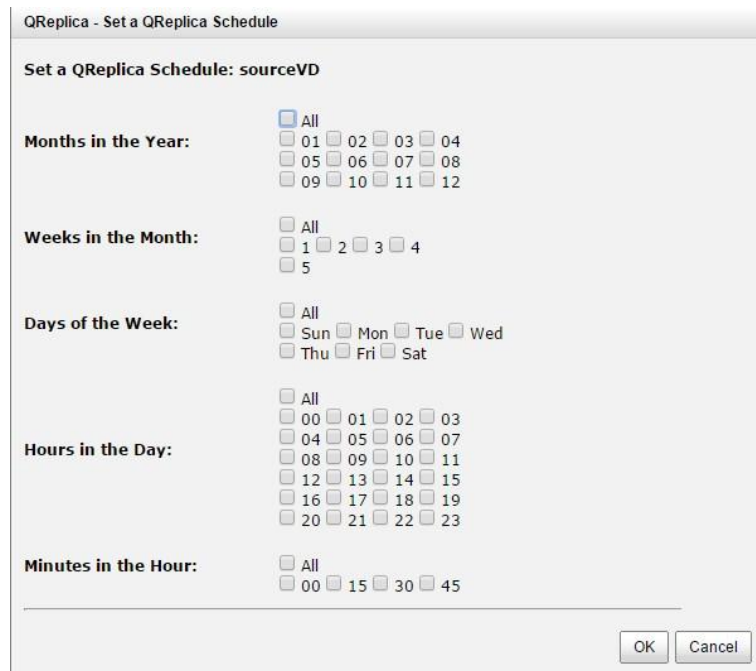
Physical Disks								RAID Groups								Virtual Disks								Snapshots								Logical Units								QReplicas							
Show snapshots for virtual disk: - srouceVD -																Show disk size in: GB																															
No.	Name	Used (GB)	Status	Health	Exposure	Cac																																									
▼ 1	QREP164103	0	Replicated	Good	No	N/A																																									
Set Snapshot Space				Schedule Snapshots				Take a Snapshot				Cleanup Snapshots																																			

Run by schedule

- Click on the ▼ -> **Schedule** to arrange the time for executing replication job.



- And select the time point you want to execute the replication job



- Please note that if you would like to set each single day at 00:00 to execute the replication job, check the hours in the **Hours in the Day** column will be enough. For example if you selected like below:




- This replication job will execute three times at 00:00 every day, and execute once for each hour.

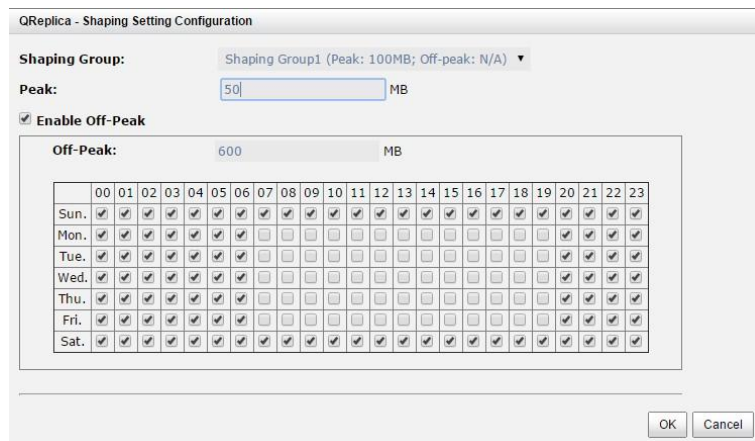
The Operation of Replication Job

The QReplica2.0 provides a new feature to limit the replicating speed via **shaping** which helps customers who have to replicate the data during office hours but don't want to affect the service traffic or loading.

- First of all, modify the shaping setting according to your requirement.




- Choose the **Shaping Group** you want to use, set the maximum replicating speed in **Peak** column (for P400Q system, it can be set from 10MB to 600MB).

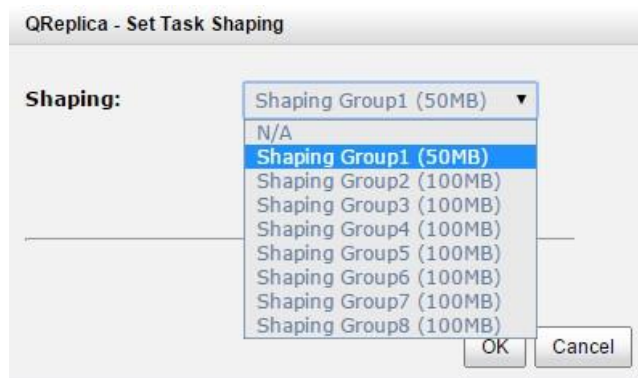


In this example, the [Shaping Group1] is selected, and from Monday to Friday 07:00 ~ 19:00, the replicating speed will be limited within 50MB/s in this group. But for the rest of time, the replicating speed will not be limited. Click OK to finish.

3. Click ▼ -> **Set Task Shaping** on the task.



4. Select the **Shaping Group1** in the drop-down menu and click **OK** button.



5. There will be an additional column showed and indicating the current applied Shaping group for this replication job.




TIP:

There are up to 8 shaping groups in the system for users to apply. If two replication jobs apply to the same shaping group, the limited replicating speed will be shared. For example, if there are two replication jobs created, applied to the same shaping group which the limited replicating speed is 100MB/s, and they will be running in the same time, the maximum speed for each replication


job can only reach to 50MB/s at the moment.

How Redundancy Works with QReplica


In normal case, the replication job is running on the controller 1 of source SAN storage. The data is replicated from the controller 1 of source SAN storage to the controller 1 of target SAN storage. The only MPIO policy that QReplica supports is fail-over. It means the data will be transmitted through one network path at a time.



When controller1 failed on source SAN storage, the replication job will be taken over by the controller 2 of source SAN storage and the replication will continue running.

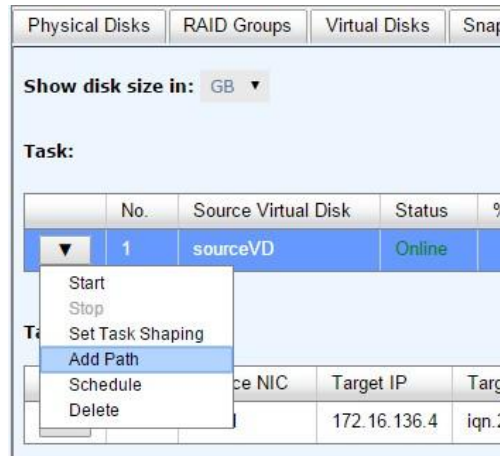


When controller1 failed on target SAN storage, the network path will be failed over by the second path from the controller 1 of source SAN storage to the controller 2 of target SAN storage.




Add Path to Setup the Redundancy of Replication Job

1. Click ▼ -> **Add Path** on the existed replication job.



2. Select the source port which is the same port that used to create this replication job, and enter the IP address of controller2 on the target SAN storage.



3. Follow the wizard to finish the creation of second path, and you may see the target name is the same one but comes from the controller2.

QReplica - Set iSCSI Authentication Information

Choose Authentication Method

No Authentication Method

CHAP

Users: (Max: 223)

Secret: (Min: 12, Max: 16)

Select a Target Node

No.	Name
<input checked="" type="radio"/> 1	iqn.2004-08.com.qsantechology .dev0.ctr2

<< Back Next >> Cancel

QReplica - Select a Target LUN

When a replication job completes, it will take a snapshot on its target virtual disk. Please make sure the snapshot space of the backup virtual disk on the target side is properly configured.

No.	LUN	Virtual Disk	Size (GB)	Vendor	Model	Serial Number
<input checked="" type="radio"/> 1	0	targetVD	200	Qsan		2041001378908900

<< Back Finish Cancel

Show disk size in: GB

Task:

No.	Source Virtual Disk	Status	%	Shaping	Speed (MB)	Target Virtual Disk	Size (GB)	Sched
1	sourceVD	Online		N/A	0	targetVD	200	

Task 'sourceVD' Path:

No.	Source NIC	Target IP	Target Name	LUN	Status
1	LAN1	172.16.136.4	iqn.2004-08.com.qsantechology	dev0.ctr1	Connected
2	LAN1	172.16.136.5	iqn.2004-08.com.qsantechology	dev0.ctr2	Connected

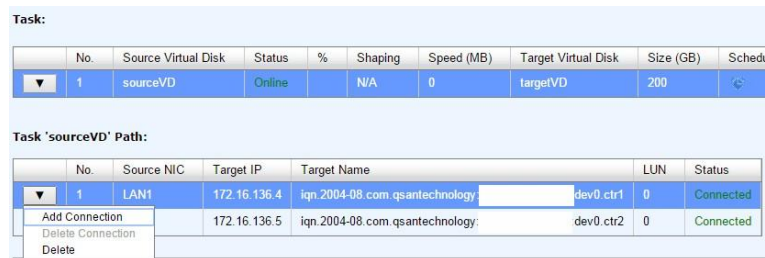
You may see there are two paths for one replication job to replicate data from source to target VD. One is from controller1 of the source SAN storage to the controller1 of target SAN storage, the

other is to the controller2 of target SAN storage. According to the mechanism we mentioned above for multipath policy of QReplica, the replicating speed will still be within one path at a time.


Add Connection in One Path to Increase the Performance of Replication Job

This is a new added feature in the QReplica2.0, and this feature is just like the MC/S feature when using the iSCSI initiator on Windows series OS. If there is one or more iSCSI port available on the same controller of source SAN storage, and it may ping to the other iSCSI IP address of controller on the target SAN storage, you may consider to add more connections for increasing the replicating speed.

1. Click ▼ -> **Add Connection** on the existing path.



2. Choose another available port on the controller1, and enter the IP address of other available iSCSI port on target SAN storage.



3. There will be an additional connection added under this path.

Task 'sourceVD' Path:

No.	Source NIC	Target IP	Target Name	LUN	Status	
1	LAN1	172.16.136.4	iqn.2004-08.com.qsantechology:	dev0_ctr1	0	Connected
	LAN2	172.16.136.7				
2	LAN1	172.16.136.5	iqn.2004-08.com.qsantechology	dev0_ctr2	0	Connected

In theory, with two connections under one path for one replication job, the replicating speed can reach two times comparing to the condition of one connection. But just to remind that if the outgoing transmitting speed of your location is limited, the replicating speed might not be increased at the moment.

Clone Transfers to QReplica

It is always being a problem that to do full copy over LAN or WAN when the replication task is executed at the first time. It may take days or weeks to replicate data from source to target within limited network bandwidth. We provide two methods to help user shorten the time of executing full copy.

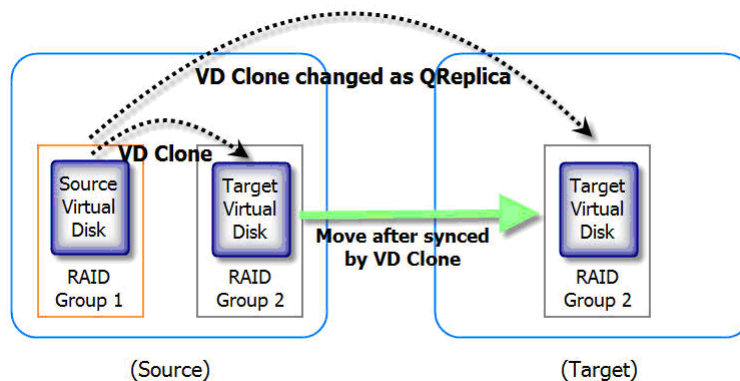
1. One is to skip full copy on a new, clean virtual disk. The term “clean” means that the virtual disk has never been written data since created. For a new created virtual disk which has not been accessed, the system will recognize it and skip full copy automatically when the replication task is created on this virtual disk at the first time.



TIP:

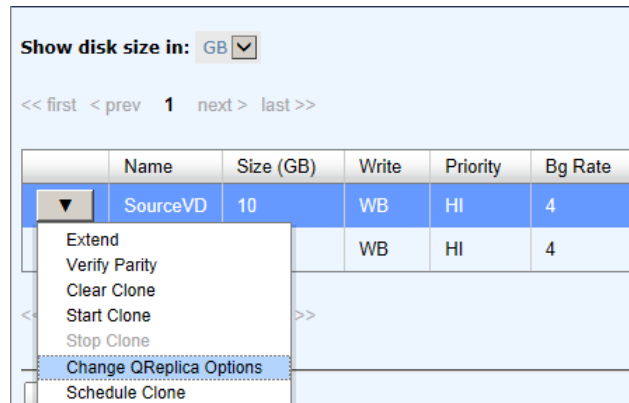
Any I/O access to the new created virtual disk will make it as “not clean”, even though executing “Erase” function when a virtual disk is created. The full copy will take place in such a case.

2. The other way is to use virtual disk clone function, which is a local data copy function between virtual disks to execute full copy at the first time. Then move all the physical drives of the target virtual disk to the target system and then turn the cloning job into replication task with differential copy afterward.



To do that virtual disk clone transfers to QReplica, please follow the procedures.

1. Create a clone job on an existing virtual disk with data stored already.
2. It is better that there is no host connected to the source virtual disk. Then run **Set Clone**, **Start Clone** to synchronize the data between source and target virtual disks.
3. After the data is synchronized, change the cloning job to a QReplica task. Select the source virtual disk, and then click ▼ -> **Change QReplica Options**.



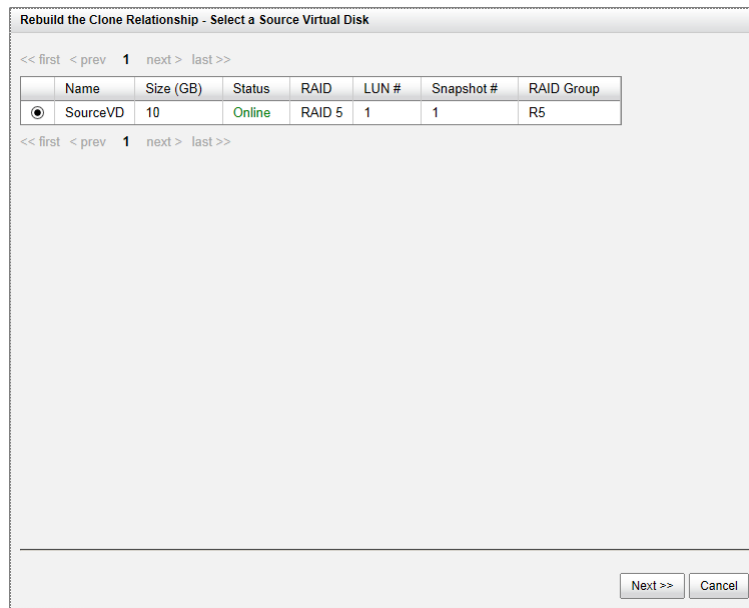
4. The **Clone** column of the source virtual disk will be changed from the name of the target virtual disk into **QRep**.



CAUTION:

Changing a cloning job to a replication task is only available when the cloning job has been finished. This change is irreversible.

5. Deactivate the RAID group which the target virtual disk resides in and move all physical disks of the RAID group to the target system. Then activate the RAID group in the target system. Remember to set snapshot space for the target virtual disk. And then attach the target virtual disk to a LUN ID.
6. In **QReplicas** tab of the source side, click **Rebuild** button to rebuild the replication task which is changed from a cloning job formerly.
7. Rebuild the clone relationship, select a source virtual disk.



8. Next steps are the same as the procedure of creating a new replication task.
9. If a wrong target virtual disk is selected when rebuilding the replication task, there will be an alert and the system stops the creation.

The Comparison of QReplca 1.0 and 2.0

The following table is the comparison with QReplca 1.0 and 2.0. QReplca 2.0 starts from AegisSAN LX FW 3.2.0 and AegisSAN Q500 FW 1.1.0.

	V2.0	V1.0
Replication port	Any iSCSI data port	Specified iSCSI port only
Tasks per controller	AegisSAN LX: 32 tasks AegisSAN Q500: 16 tasks	8 tasks
MPIO	2 paths per task	2 paths per task
MC/S	4 connections per session	N/A
Traffic shaping	8 traffic shapers can apply to each task	N/A, max 20 MB/s
Traffic shaping scheduling	Peak and off-peak to apply different traffic shaping	N/A
Snapshot	Auto rollback to latest image if any accident during replication	Manually

Conclusion

The QReplica is a block-level copy which can provide more efficient way in differential copy than file-level copy. Because the changes on a volume are calculated by block, it copies the modified blocks only but not whole file.

It is very important to plan the snapshot space usage. Snapshot can be used for regular snapshot, VSS hardware provider, virtual disk clone ...etc. Each service shares the same snapshot space. Users have to be careful to prevent running out of snapshot space, otherwise it will cause any service failed. So user should arrange the proper snapshot space. The most conservative way is reserve two or three times larger than virtual disk space for snapshot.

Applies To

- AegisSAN LX FW 3.2.0
- AegisSAN Q500: FW 1.1.0

Obsolete

- Qsan White Paper
QWP201005-P300H-How_to_use_the_QReplica.v1.5.pdf