QSAN Official Document

QReplica 3.0

Sync. and Async. Replication

White Paper

March 2023

ANNOUNCEMENT

Copyright

QSAN believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

Trademarks

- QSAN, the QSAN logo, QSAN.com, XCubeFAS, XCubeSAN, XCubeNXT, XCubeNAS, XCubeDAS, XEVO, SANOS, and QSM are trademarks or registered trademarks of QSAN Technology, Inc.
- Microsoft, Windows, Windows Server, and Hyper-V are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.
- Linux is a trademark of Linus Torvalds in the United States and/or other countries.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Mac and OS X are trademarks of Apple Inc., registered in the U.S. and other countries.
- Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
- VMware, ESXi, and vSphere are registered trademarks or trademarks of VMware, Inc. in the United States and/or other countries.
- Citrix and Xen are registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries.
- Other trademarks and trade names used in this document to refer to either the entities claiming the marks and names or their products are the property of their respective owners.

TABLE OF CONTENTS

FIGURES

TABLES

NOTICES

Information contained in this document has been reviewed for accuracy. But it could include typographical errors or technical inaccuracies. Changes are made to the document periodically. These changes will be incorporated in new editions of the publication. QSAN may make improvements or changes in the products. All features, functionality, and product specifications are subject to change without prior notice or obligation. All statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

PREFACE

Executive Summary

Data is usually one of the most valuable assets of an enterprise. This makes data a critical part of the organization's daily operations. Disruptions can occur at any time. Whether it is a planned or unplanned outage, ensuring that the data integrity to the enterprise is always available is a top priority. It is recommended to prepare a plan for critical data to prevent costly interruptions. Data replication solutions can ensure business continuity, high availability and data protection.

QSAN QReplica 3.0 (Remote Replication) provides synchronous and asynchronous replication solutions. These solutions enable you to replicate data to other systems, whether they are located at the same site or remote facilities. Having a remote copy of the data protects you from the failure of the primary system and allows you to recover quickly and easily on the target system. The solutions will help you protect your data and achieve your business goals in terms of data availability and protection.

In short, QReplica 3.0 can provide a complete backup solution for all business scales.

Audience

This document provides technical guidance for designing and implementing “Disaster Recovery” through remote replication feature, and it is intended for use by system administrators, SAN designers, storage consultants, or anyone who has purchased this product and is familiar with server and computer network, network administration, storage system installation and configuration, network attached storage management, and relevant protocols. Any settings which are configured with basic operations will not be detailed in this document. If there is any question, please refer to the user manuals of products, or contact QSAN support for further assistance.

Technical Support

Do you have any questions or need help trouble-shooting a problem? Please contact QSAN Support, we will reply to you as soon as possible.

- Via the Web: https://www.qsan.com/technical_support
- Via Telephone: +886-2-77206355
- (Service hours: 09:30 - 18:00, Monday - Friday, UTC+8)
- Via Skype Chat, Skype ID: qsan.support
- (Service hours: 09:30 - 02:00, Monday - Friday, UTC+8, Summer time: 09:30 - 01:00)
- Via Email: support@qsan.com

Information, Tip, and Caution

This document uses the following symbols to draw attention to important safety and operational information.

INFORMATION

INFORMATION provides useful knowledge, definition, or terminology for reference.

TIP

TIP provides helpful suggestions for performing tasks more effectively.

CAUTION

CAUTION indicates that failure to take a specified action could result in damage to the system.

1. INTRODUCTION

QSAN QReplica 3.0 (Remote Replication) provides synchronous and asynchronous replication solutions. With QReplica 3.0 feature, DR (Disaster Recovery) can be achieved with simple and easy orchestration between QSAN platforms. In this document, synchronous and asynchronous replication concepts and detail procedures will be introduced for having best practice of DR mechanism with QSAN platform.

1.1. Start with Disaster Recovery

DR (Disaster Recovery) is about preventing total failure of mission critical business systems and to recover within minimum time and impact. Preventing data loss requires a continuous data protection method. This includes preparation for and recovery from events of human error, software and hardware failure, network down, internal or external power failure and all other events. To beat this challenge, IT managers must plan for redundancy of one or more backup systems at different locations. This involves constant or periodically data duplication to infrastructures at different sites to ensure business continuity from constant availability.

1.1.1. Understanding RPO and RTO

RPO (Recovery Point Objective) and RTO (Recovery Time Objective) are the two most important parameters in a disaster recovery. To achieve continuous data protection, you must understand these 2 key indicators that measure acceptable business risk. These goals can guide companies in choosing the best data backup plan.

- **RPO (Recovery Point Objective):** RPO refers to the maximum amount of data loss that the system can tolerate. It refers to the time point corresponding to the recovered data after the business is restored. The RPO depends on the update level of the data recovery. This update level can be the last day. The backup data can also be 1 hour ago data. This is related to the frequency of data backup. In order to improve RPO, the frequency of data backup must be increased. RPO is an indicator that reflects the integrity of data restored.
- **RTO (Recovery Time Objective):** RTO is the duration of time the system can tolerate service interruption. For example, if the service needs to be restored within half an hour after the occurrence, the RTO value is 30 minutes. The RTO is only the time period between the occurrence of a failure, from the moment when the system goes down and the application

stops, to the time when it is restored to support the operation. RTO is an indicator reflecting the timeliness of business recovery. It indicates the time required for business to recover from interruption to normal. The smaller the RTO value, the stronger the data recovery capability of the disaster recovery system. The IT administrators can deploy many disaster recovery systems. Obtain the smallest RTO, but it means investing a lot of money.

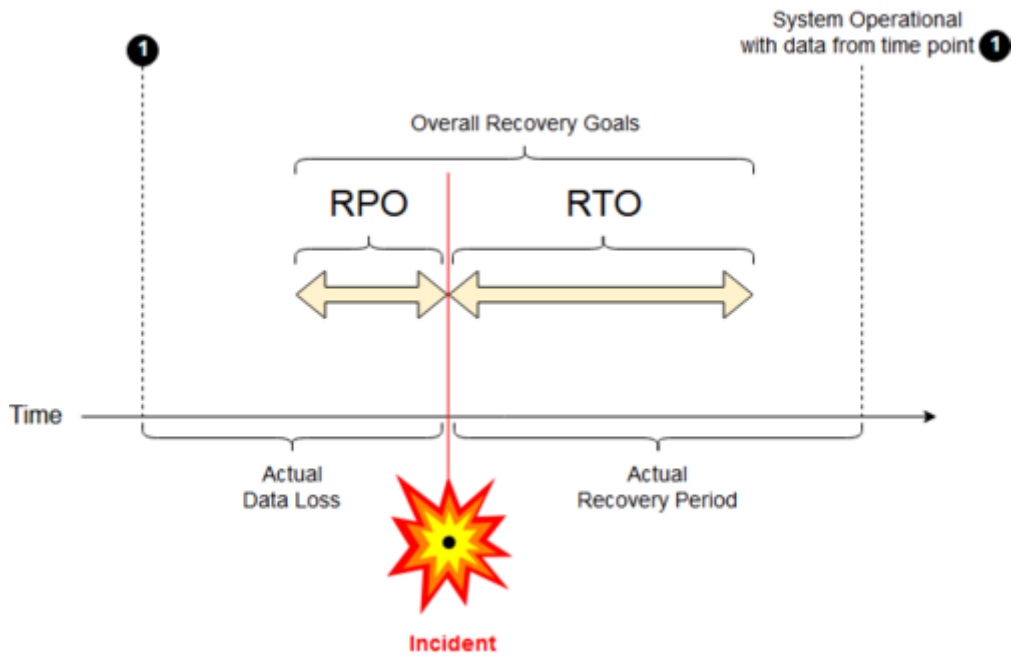

Figure 1-1 RPO and RTO Schematic Representation (from Wikipedia)

In Figure 1-1, the RPO indicator comes from before the failure occurs, and the RTO indicator comes from after the failure occurs. The smaller the value of the two is, the time interval can be effectively shortened from business transition period to normal.

INFORMATION

DR (Disaster Recovery) exercises must be performed regularly to achieve the target RPO (Recovery Point Objective) and RTO (Recovery Time Objective) according to your DR Plan.

1.1.2. How RPO and RTO Impact Business Continuity

In Figure 1-2, traditional backup and recovery techniques (such as tape backup or disk-to-disk) are measured in hours. It will take a long time for them to recover, such as several hours to several days.

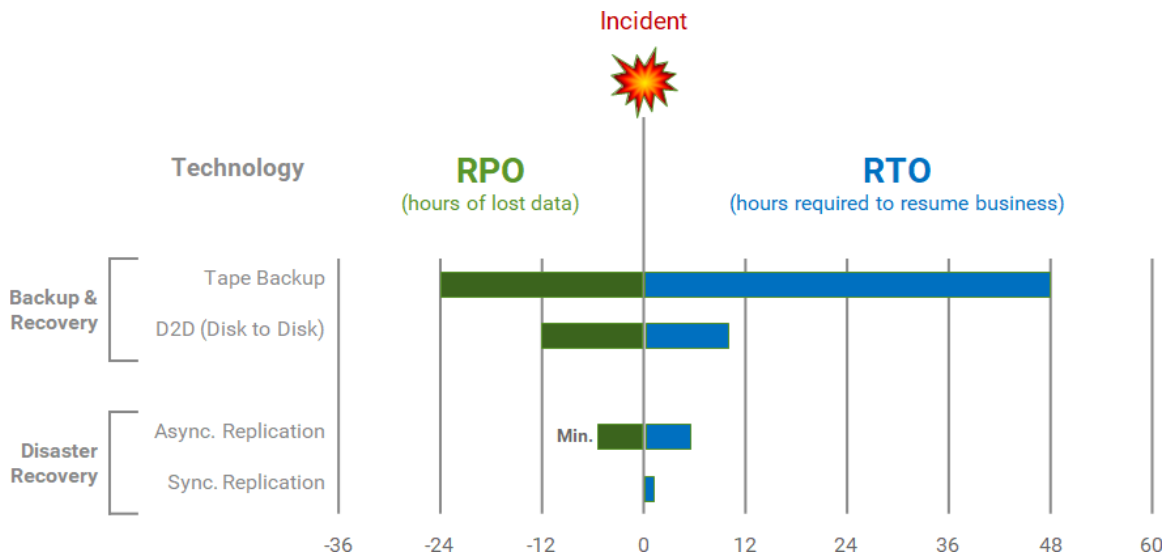

Figure 1-2 The Relationship between Backup Technologies and RPO / RTO

Disaster recovery techniques such as asynchronous replication are usually measured in minutes. In the event of an accident, the recovery time will be reduced to a few minutes. Another more stringent business continuity goal is measured in seconds, usually with no data loss at all. Synchronous replication technology can achieve this.

1.2. Introduce to Remote Replication

This section briefly introduces synchronous and asynchronous replication technologies.

1.2.1. Synchronous Replication

Synchronous replication can ensure data consistency between the replication source and target. This can be achieved by ensuring that every write I/O requested from the applications on the host will be completed after the host obtains the confirmation between the source and the target. If the write I/O cannot be committed on the source or target, the write operation will

not be committed to ensure consistency. In addition, write failures are sent back to its applications on the host. Then, application error handling will determine the next appropriate step for the pending transaction.

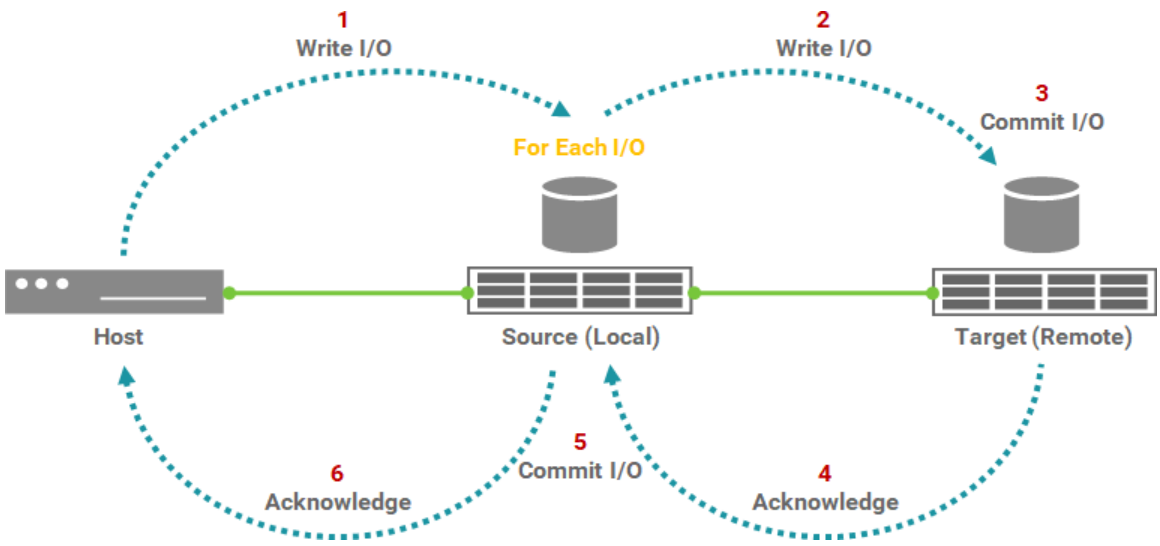

Figure 1-3 Write I/O Sequence of Synchronous Replication

Figure 1-3 describes the write I/O pattern sequence with synchronous replication:

1. The application on the host sends a write request to the source volume.
2. The write I/O copies to the target volume.
3. The write I/O at the target is committed to the target volume.
4. The write commit at the target is confirmed back to the source.
5. The write I/O is committed to the source volume.
6. Finally, the write acknowledgement will be sent to the application on the host.

1.2.2. Asynchronous Replication

Asynchronous replication can achieve the same data protection goals. However, the method and frequency of data replication are different from synchronous replication. Asynchronous replication commits writes at the source, and then send the acknowledgement to the application on the host. The accumulated writes committed on the source volume will be copied to the target volume in batches at scheduled intervals.

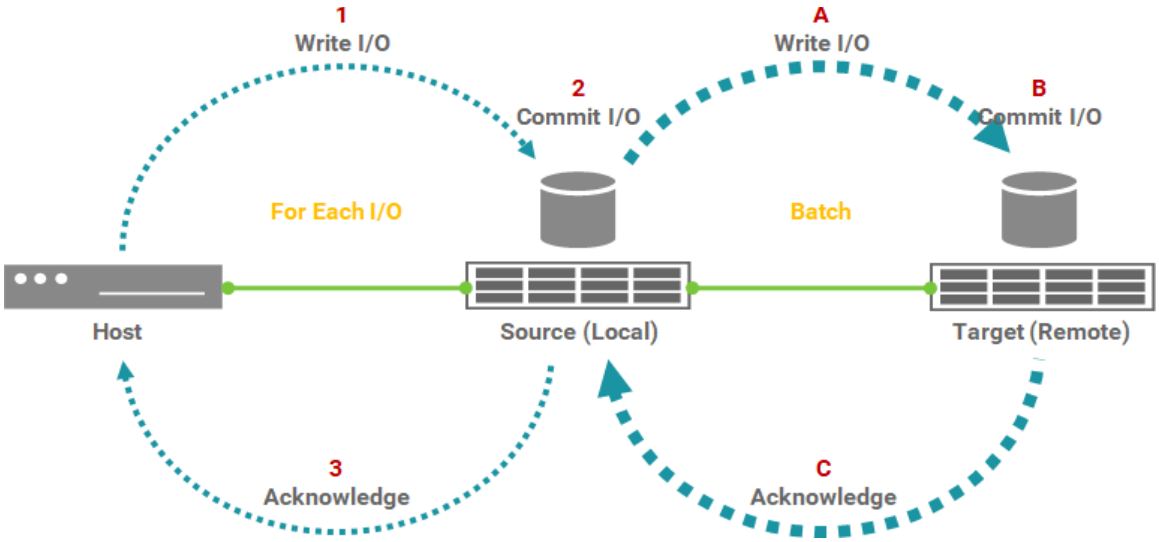

Figure 1-4 Write I/O Sequence of Asynchronous Replication

Figure 1-4 describes the write I/O sequence with asynchronous replication:

1. The application on the host sends a write request to the source volume.
2. The write I/O is committed to the source volume.
3. The write acknowledgement will be sent to the application on the host.

The process is repeated for each write I/O requested by the application on the host.

A. Periodically copy a batch of write I/O that has been submitted to the source volume to the target volume.
B. Commit to the target volume.
C. Batch confirmation is sent to the source.

1.3. Synchronous vs. Asynchronous Replication

Both synchronous and asynchronous replications have their advantages and disadvantages. Before choosing a backup method, you must understand the risks of disaster recovery.

Synchronous replication ensures that every data written to storage is saved locally and remotely. This can ensure zero data loss in the event of a disaster. So RPO is zero. In a synchronous replication solution, there are also trade-offs. Since each write needs to be saved locally and remotely, any problems affecting the source or target storage or the replication link between these two will affect the application in terms of latency and availability. As the distance between the source and target increases, this response time also increases. It is

recommended to ensure that the link delay between the source and target is less than 10 ms (millisecond). The converted distance is about 100 km (kilometer).

Synchronous replication is mainly used for mission critical applications that require instant failover when the source site fails. Although synchronous replication is usually used for disaster recovery, it is preferred for applications that cannot tolerate data loss and have a short recovery time.

Table 1-1 Synchronous and Asynchronous Comparison

ITEM	SYNCHRONOUS REPLICATION	ASYNCHRONOUS REPLICATION
RPO (Recovery Point Objective)	= 0	Last sync of incremental data
RTO (Recovery Time Objective)	0 ~ Very low	Time of next sync
Delay between Source and Target	< 10 ms	No limit, preferably < 1 sec.
Distance between Source and Target	100 km	No limit
Application	Mission critical	Business critical
License Required	Yes	No

Asynchronous replication is mainly used to replicate data over long distances. It does not affect host I/O latency, because once host writes are saved to the source storage, they are acknowledged. Since write operations are not immediately copied to the target, all writes will be tracked on the source. This data will be copied during the next synchronization. Therefore, RPO should be an acceptable amount of data, which may be lost due to failure. And RTO is equal to the period of the next synchronization. This time increment also affects the amount of data that needs to be replicated during the next synchronization and the amount of potential data loss in the event of a disaster.

Of course, the perfect solution is that both RTO and RPO are zero. This means that when a fault occurs, the system will respond immediately, and there will be no data loss. To achieve this goal,

the system design is extremely complicated and the cost is very expensive. So it should avoid falling into the vicious circle of simply pursuing the improvement of the two indicators, and combining the actual situation, appropriately improving the two indicators according to local conditions is the right way.

1.4. Support Firmware Version

The following summarizes the firmware versions supported by QReplica 3.0, including synchronous and asynchronous replication.

- SANOS
  - FW 2.0.0 supports synchronous replication with **license**
  - FW 2.0.1 supports local to remote of synchronous replication
- XEVO
  - FW 2.1.0 supports synchronous replication with **license**
  - FW 2.1.1 supports local to remote of synchronous replication

Table 1-2 Support Firmware Version

MODEL	OS	SYNCHRONOUS REPLICATION	ASYNCHRONOUS REPLICATION
XCubeSAN Series	SANOS FW 2.0.1	License	Built-in
XCubeFAS 2026 (XF2026)	XEVO 2.1.1	License	Built-in
XCubeFAS 3126 (XF3126)	XEVO 2.1.1	License	Built-in

2. Theory of Remote Replication Operations

This chapter will introduce the theory of remote replication operations in detail. Multiple backup solutions give you more choices to protect your data. With XEVO's unique automatic replication feature, you can easily deploy any synchronous or asynchronous remote replication without configuration. Through local-to-remote, you can transfer your local backup to remote sites without redoing full copy again. This technology gives you multiple ways to protect your digital assets safely.

Figure 2-1 Remote Replication Diagram

2.1. Synchronous Replication

Synchronous replication is a block-level volume backup function through LAN or WAN. It is a timely replication function that can always synchronize source and target data. It uses iSCSI function to establish a replication connection. The replication task can use the full bandwidth of the assigned network port to allow the best backup speed.

We assume that the synchronization source port must be greater than or equal to 10 Gbps to ensure normal operation. However, the synchronization source port in the source storage is not limited to the dedicated one. In order to balance replication traffic, it is still recommended to separate the port of incoming data from the host and the synchronization source port to the target storage.

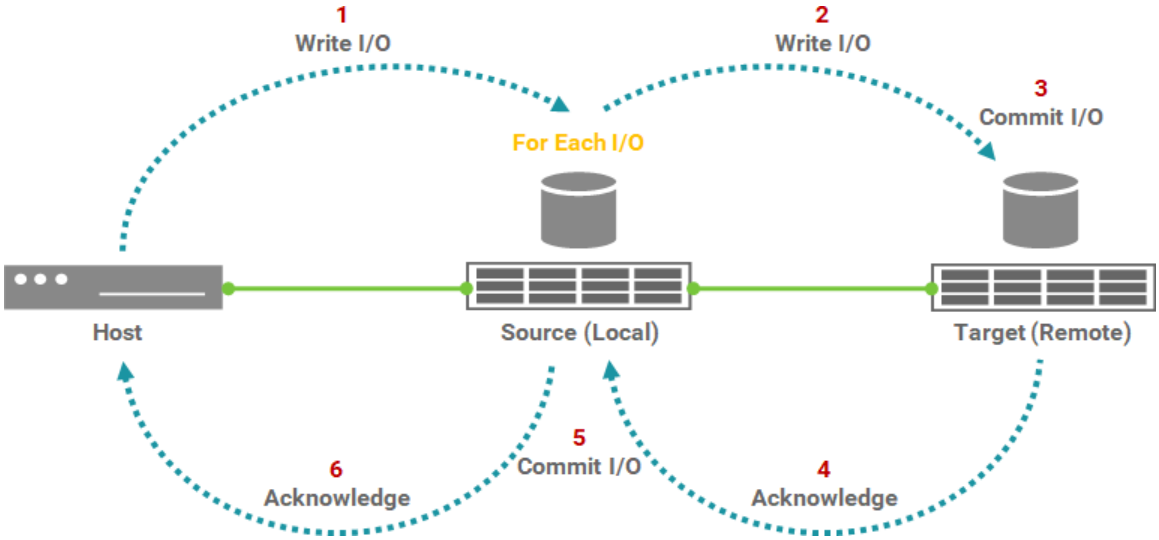

Figure 2-2 Synchronous Replication

TIP

The synchronous source port must be greater than or equal to 10 Gbps.

When a user creates a synchronous replication task, it will automatically copy all data from the source volume to the target at the beginning. This is also called **full replication**. At the same time, the source system will record a **fracture log** to maintain the block whether has been synchronized. The fracture log is a bitmap saved in the source storage to indicate which block of the source have been updated.

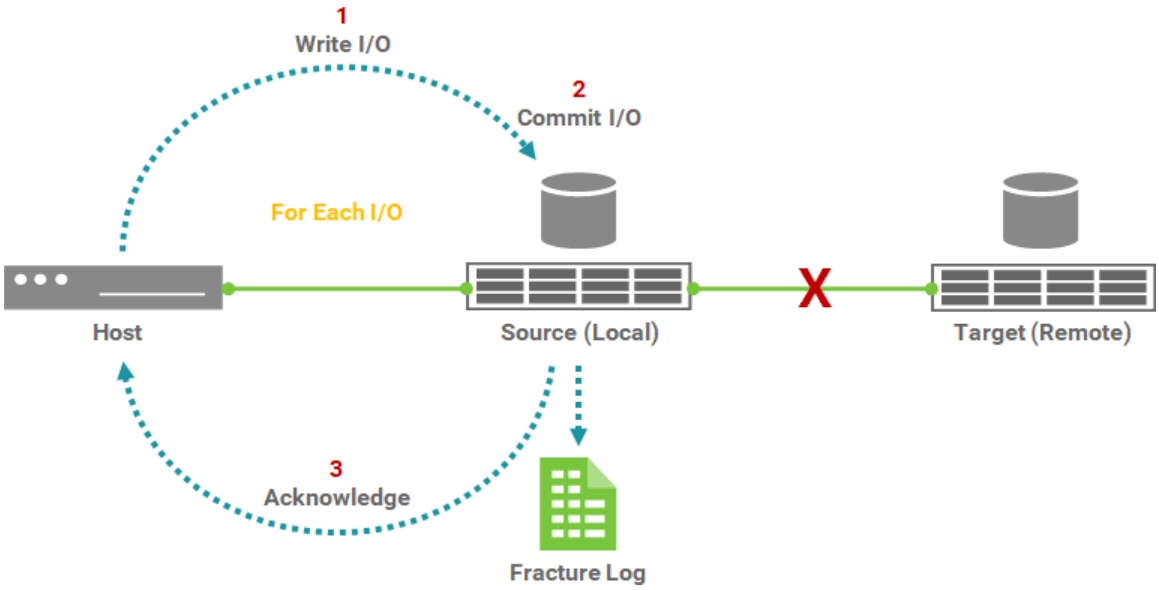

Figure 2-3 Fracture Log in Synchronous Replication

It is invoked when the target of a replication connection is lost for any reason and becomes out of sync. For example, when an incident such as a network interruption occurs, the source cannot replicate the data to the target, I/O from the host will continue to be written to Source. The fracture log is kept update until the incident is resolved. The source will learn about the unsynchronized blocks and replicate them to the synchronized state.

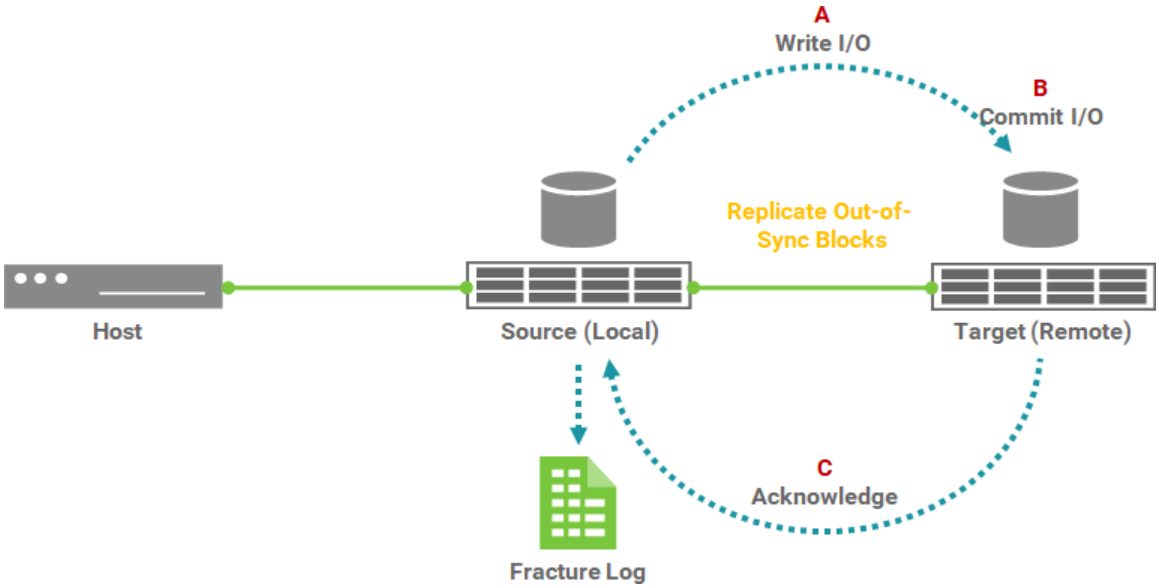

Figure 2-4 Replicate Out-of-Sync Blocks

If a network error still occurs and the replication is blocked, the system will issue a warning every 20 seconds and enter the failed state after 60 seconds. Therefore, the task will be stopped. At this point, if you want to continue the task, you must start it manually.

Another operation is that the user can manually stop the task until the user manually resumes the task, the situation is the same. If you are not sure whether the target has been touched, you can also perform a **full replication** first when restarting the task.

2.2. Asynchronous Replication

Asynchronous replication is a block-level, differential remote volume backup function through LAN or WAN. In the beginning, replication will copy all data from the source volume to the target. It is also called a **full replication**. After that, use snapshot technology to perform incremental replication.

TIP

Please be fully aware that incremental copies require snapshots to compare data differences. Therefore, it is very important to provide enough snapshot space for the volume.

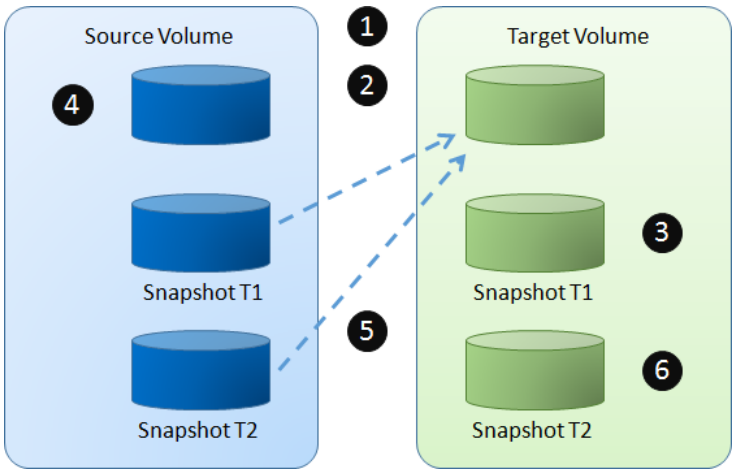

Figure 2-5 Asynchronous Replication Steps

Generally, asynchronous replication operates in the following manner.

1. Create an asynchronous replication task. Over time, a snapshot is taken on the Source Volume.
2. Copy the data from the Source Snapshot T1 to the Target Volume.
3. The Target Snapshot T1 is refreshed and becomes a base.
4. The host continuously writes new data to the Source Volume.
5. In the next synchronization, the Source Snapshot T2 will be refreshed and only the changes since the last synchronization will be copied.
6. The Target Snapshot T2 is refreshed and becomes the next base.

Asynchronous replication uses the iSCSI function to establish a replication connection. It can use the full bandwidth of the allocated network port to achieve the best backup speed. However, in order to balance replication traffic and non-replication traffic, the traffic shaping function can help reserve the necessary bandwidth for non-replication I/O.

2.3. Auto Configuration

It does easily deploy the remote replication without doing configurations. In a protection group, you only need to log in a remote system through management port. The source system will send the configurations of volumes to the remote system. The remote system will generate the corresponding volumes for the source system.

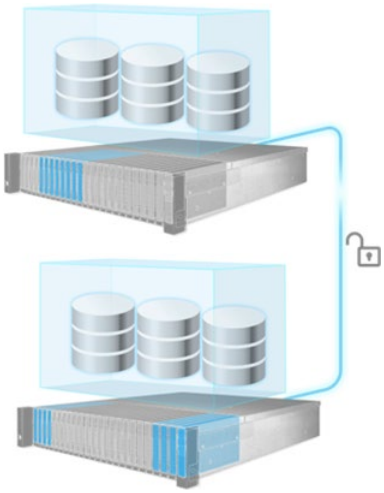

Figure 2-6 Auto Replication

The target system must prepare at least 3 disk drives and the capacity of the target pool should be larger than the total capacity of all source volumes in the protection group.

Generally, auto replication operates in the following manner.

1. From source unit, login to the management port of the target unit.
2. Source unit will send the configurations of volumes which user want to back up to the target unit.
3. The target unit will generate the corresponding volumes for source unit
4. The source unit will start the replication task.

2.4. Traffic Shaping

The traffic shaping function can help reserve the necessary bandwidth for non-replication I/O operations. Eight shaping groups can be set. In each shaping group, peak and off-peak time slots are provided for different bandwidths.

2.5. MPIO in Remote Replication Task

In the remote replication scenario, MPIO (Multi-Path I/O) supports redundancy. Usually, the remote replication task runs on the master controller (usually CTRL 1) of the source unit. The data is replicated from the master controller of source unit to the target unit, and the target unit is set with the target IP address (usually also CTRL 1) when creating the remote replication task. The second path from the source unit can be added to the CTRL 2 of the target unit. The maximum MPIO in a remote replication task is 2 paths. The following is a remote replication MPIO diagram.

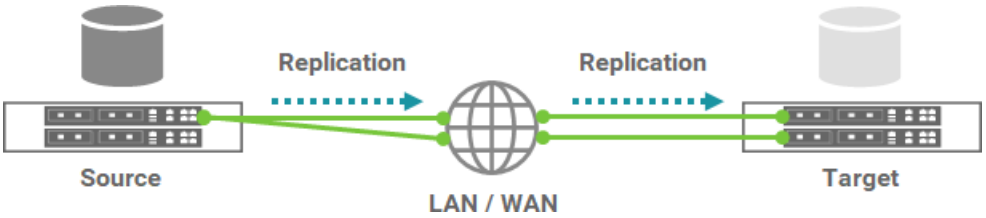

Figure 2-7 Remote Replication MPIO Diagram

How Redundancy Is Used with Remote Replication

If CTRL 1 fails on the source unit, the replication connection will be taken over by CTRL 2 and the replication task will continue running.

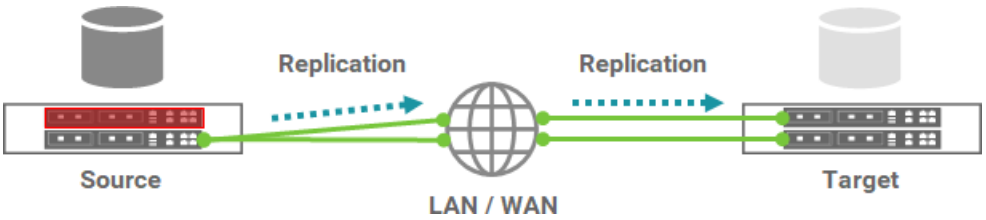

Figure 2-8 Remote Replication Source Controller Fail Diagram

In another case, when CTRL 1 fails on the target unit, the replication connection will fail over to the second path from CTRL 1 of the source unit to CTRL 2 of the target unit.

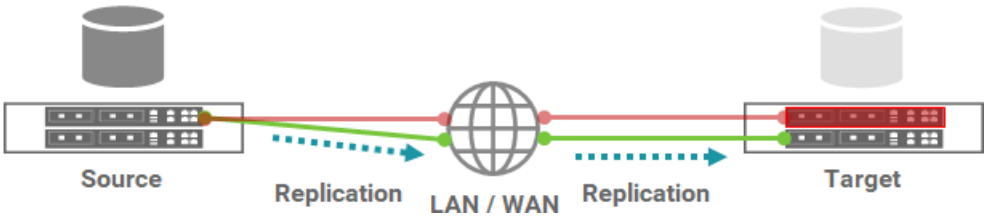

Figure 2-9 Remote Replication Target Controller Fail Diagram

2.6. MC/S in Remote Replication Task Path

MC/S (Multiple Connections per Session) is another feature in remote replication. If there are multiple iSCSI ports available on the source unit and they can be connected to other iSCSI ports on the target unit, you can add MC/S to increase the replication speed. The maximum MC/S for each task path is 4 connections.

2.7. Local Clone Transfers to Remote Replication

In order to handle full replication of large amounts of data, replication allows local cloning tasks to be converted into remote replication tasks. You can perform a local clone of the full copy first. Then, use the disk roaming function to physically transfer the disk drive containing the

cloned volume to the remote site. Finally, use the function of converting from a local clone task to a remote replication task.

When performing a replication task for the first time, full replication via LAN or WAN is always a problem. With limited network bandwidth, it may take days or weeks to copy data from the source to the target. We provide **local to remote** (local clone transfers to remote replication) method to help users shorten the time to execute a full copy.

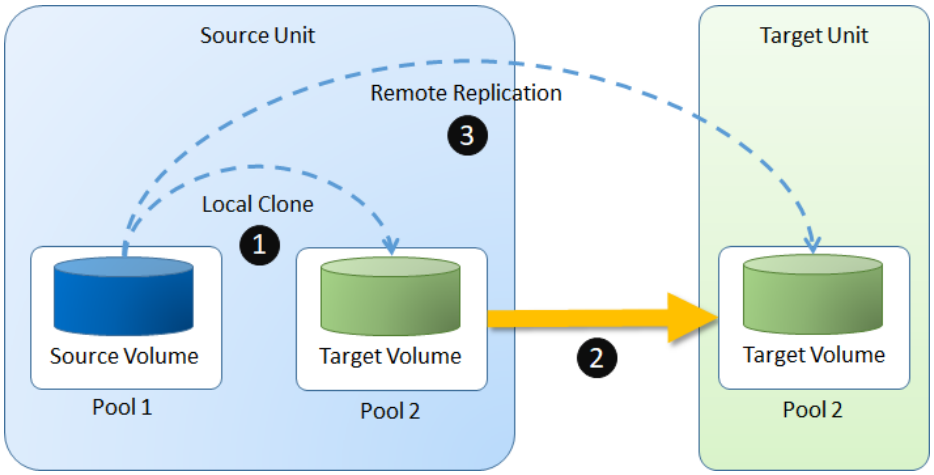

Figure 2-10 Local Clone Transfers to Remote Replication Diagram

1 Step Local-to-Remote

Through local-to-remote, transfer your local backup to remote sites without having to redo the full copy. On the remote side, you only need to insert all disk drives without any configurations. And then complete the local-to-remote process.

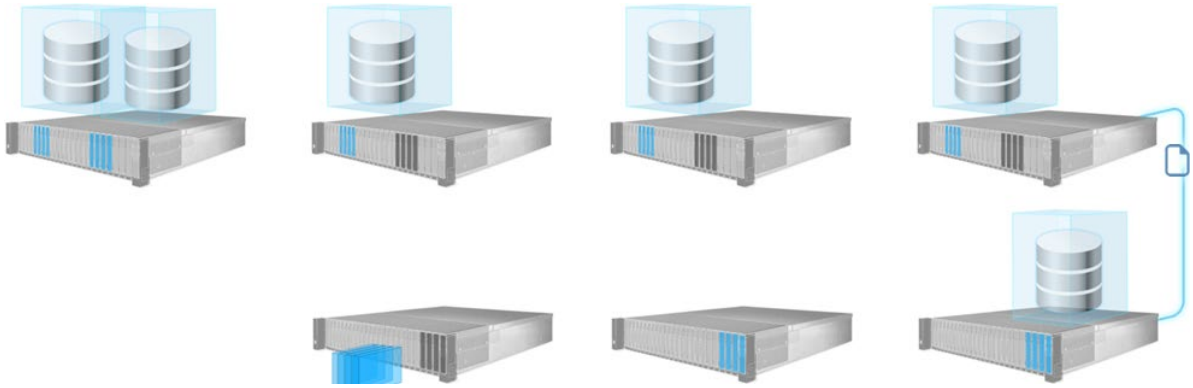

Figure 2-11 1 Step Local-to-Remote

INFORMATION

The **Auto Replication** and **1 Step Local-to-Remote** features are supported by XEVO only. It means that the source and target arrays are both running XEVO. In addition, the management ports and data ports of the source and target arrays must be connected to each other. Whether they are direct connections or through LAN switches.

Generally, local to remote operates in the following manner.

1. Create a Local Clone task from the Source Volume to Target Volume, and do it once.
2. After the local clone is completed, deactivate the target pool. Remove all disk drives of the target pool and deliver them to the Site B (Target Unit).

Figure 2-12 Local Clone Transfers to Remote Replication

3. Rebuild the replication connection.

2.8. Support Scenario

The replication scenarios supported by the models are summarized below. The XCubeSAN and XCubeFAS series can be the source and target of each other for remote replication through manual configuration. But only XEVO OS supports auto configuration, this feature is only available for XCubeFAS series. For the local clone to remote copy function, it is best that the source and target are the same model.

Table 2-1 Support Scenario

SOURCE	OS	MANUAL CONFIGURATION	AUTO CONFITURATION	LOCAL TO REMOTE
XCubeSAN Series	SANOS FW 2.0.1	XCubeSAN Series XF2026 XF3126	X	XCubeSAN Series
XCubeFAS 2026 (XF2026)	XEVO 2.1.1	XCubeSAN Series XF2026 XF3126	XF2026 XF3126	XF2026
XCubeFAS 3126 (XF3126)	XEVO 2.1.1	XCubeSAN Series XF2026 XF3126	XF2026 XF3126	XF3126.

3. Configure Remote Replication

This chapter will discuss the remote replication feature which allows users to create synchronous or asynchronous replication tasks for block storage between supported systems. All configuration and management operations in this section will be demonstrated.

3.1. Prerequisites

This section describes remote replication prerequisites. First, list the parameters of remote replication. The maximum replication task quantity per volume is 1, and the maximum replication task quantity per system is 32 which including both synchronous and asynchronous replication tasks.

Table 3-1 Synchronous Replication Parameters

ITEM	VALUE
Maximum remote replication task quantity per volume (Maximum remote replication pairs per source volume)	1
Maximum remote replication task quantity per system (include synchronous and asynchronous)	32
Maximum iSCSI multi-path quantity in a remote replication task	2
Maximum iSCSI multiple connection quantity per remote replication task path	4

Preliminary Preparation of Creating a Remote Replication Task

The following takes the creation of a remote replication task from the source volume to the target volume as an example. Suppose we defined the name and IP address in the example.

Figure 3-1 Example of Creating a Synchronous Replication Task

Site A (Source Unit) Configuration:

- Controller 1, Onboard LAN 1 IP Address: 10.10.1.1
- Controller 1, Onboard LAN 2 IP Address: 10.10.1.2
- Controller 2, Onboard LAN 1 IP Address: 10.10.1.3
- Controller 2, Onboard LAN 2 IP Address: 10.10.1.4
- Source Volume Name: Src_Vol_01, Src_Vol_02

Site B (Target Unit) Configuration:

- Controller 1, Onboard LAN 1 IP Address: 10.10.1.101
- Controller 1, Onboard LAN 2 IP Address: 10.10.1.102
- Controller 2, Onboard LAN 1 IP Address: 10.10.1.103
- Controller 2, Onboard LAN 2 IP Address: 10.10.1.104
- Target Volume: Name: Tgt_Vol_01, Tgt_Vol_02

3.2. License

The synchronous replication license must be purchased to enable the feature. There is also a 30-day free trial license for evaluation. Although the setting method is different from XEVO and SANOS, the following sections describe the configuration separately.

3.2.1. Enable Synchronous Replication License

The synchronous replication function is optional. Before using it, you must enable synchronous replication license. Select the **System** tab and the **Maintenance** subtab, and then click the **Licenses** pane to allow users to active licenses.

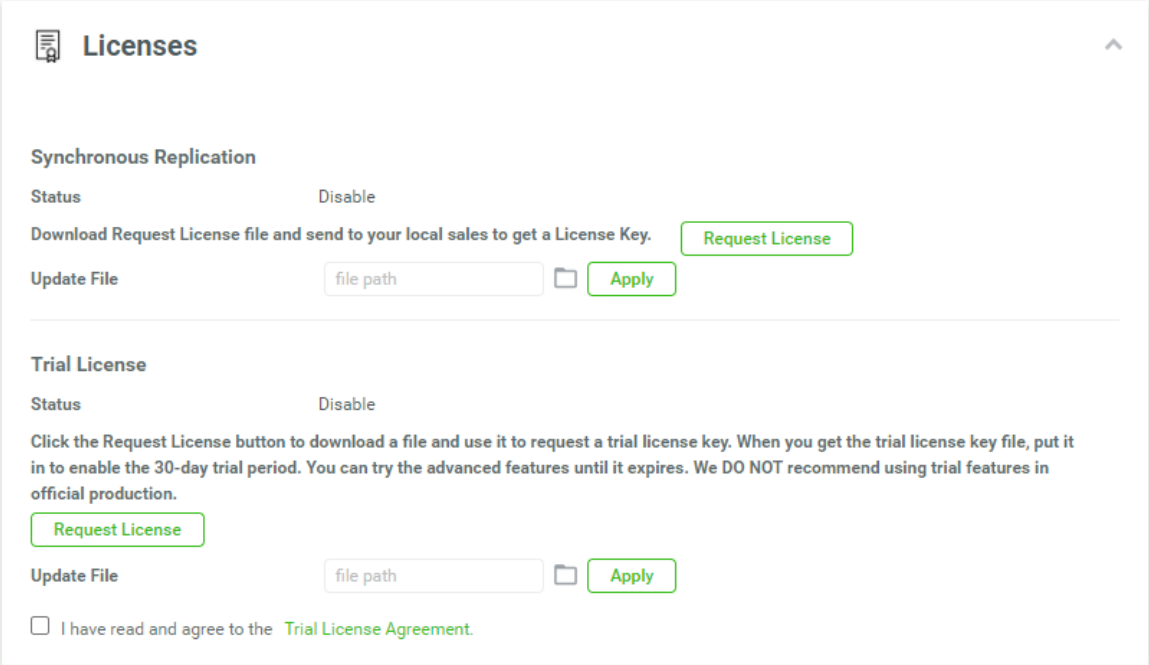

Figure 3-2 Enable Synchronous Replication License

Click the **Request License** button to download the file and send to your local sales to obtain a License Key. After getting the license key, click the icon to select it, and then click the **Apply** button to enable. Each license key is unique and dedicated to a specific system. If the license is active, the status will show as **Enable**. After enabling the license, the system must reboot manually to take effect.

3.3. Configure Synchronous Replication

This section describes how to create synchronous replication and its related functions. We provide two methods to setup the tasks, auto and manual configuration. With XEVO's unique automatic replication feature, you can easily deploy any remote replication without configuration. Manual configuration is the traditional way to complete the settings step by step.

3.3.1. Auto Configuration

Here is an example of creating a synchronous replication task with auto configuration.

Prepare in Site B (Target Unit)

The target unit has at least 3 free disks without any pool configuration, and the total capacity is greater than or equal to the source volumes. In addition, the data ports in the target unit are connected to the source unit and can be pinged.

TIP

The auto configuration can run in the same firmware version of the source and target units.

Configure Replication Plan-Remote in Site A (Source Unit)

1. Click the + icon in the **Protection Groups** pane to pop up a wizard.

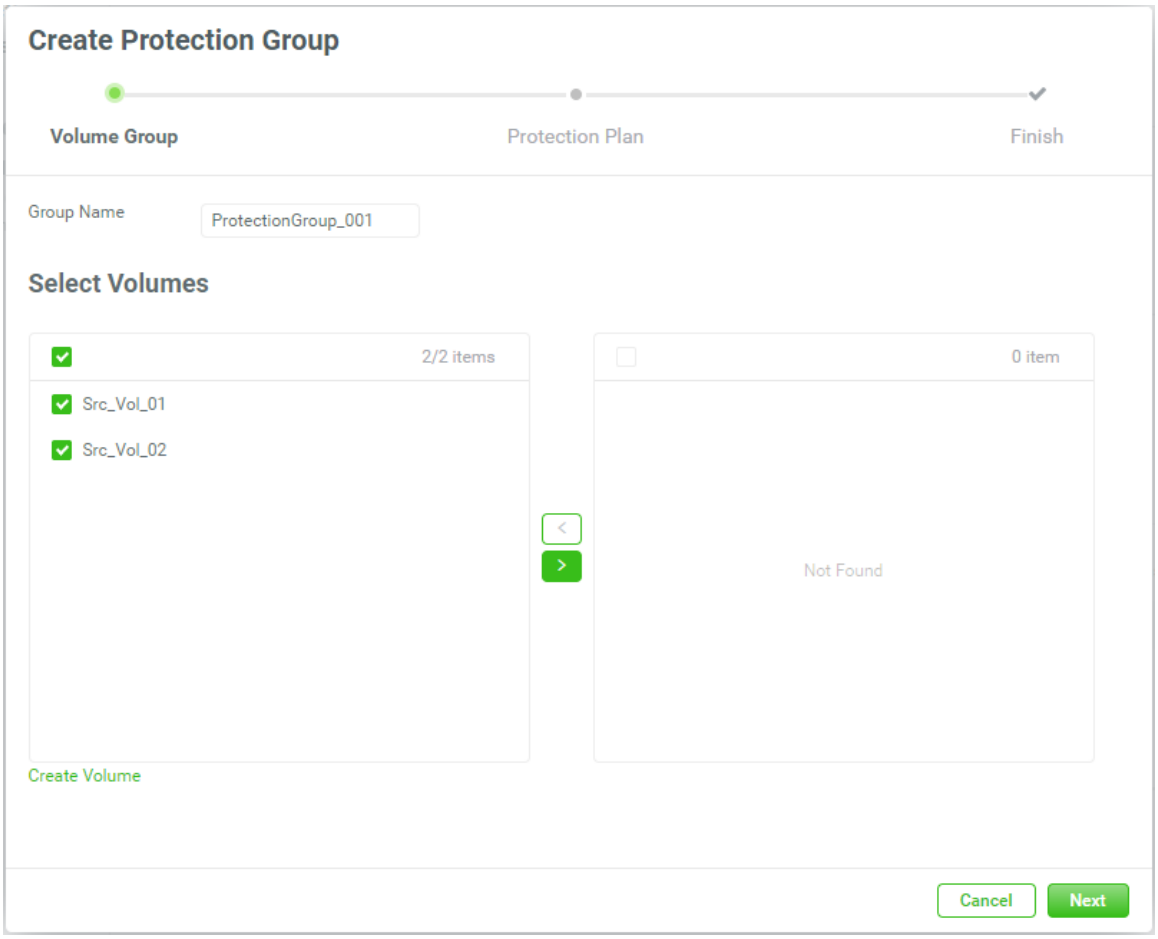

Figure 3-3 Synchronous Replication – Auto Configuration Step 1

2. Select source volumes, and then click the button to move them from the left side to the right side.
3. Click the **Next** button to continue.

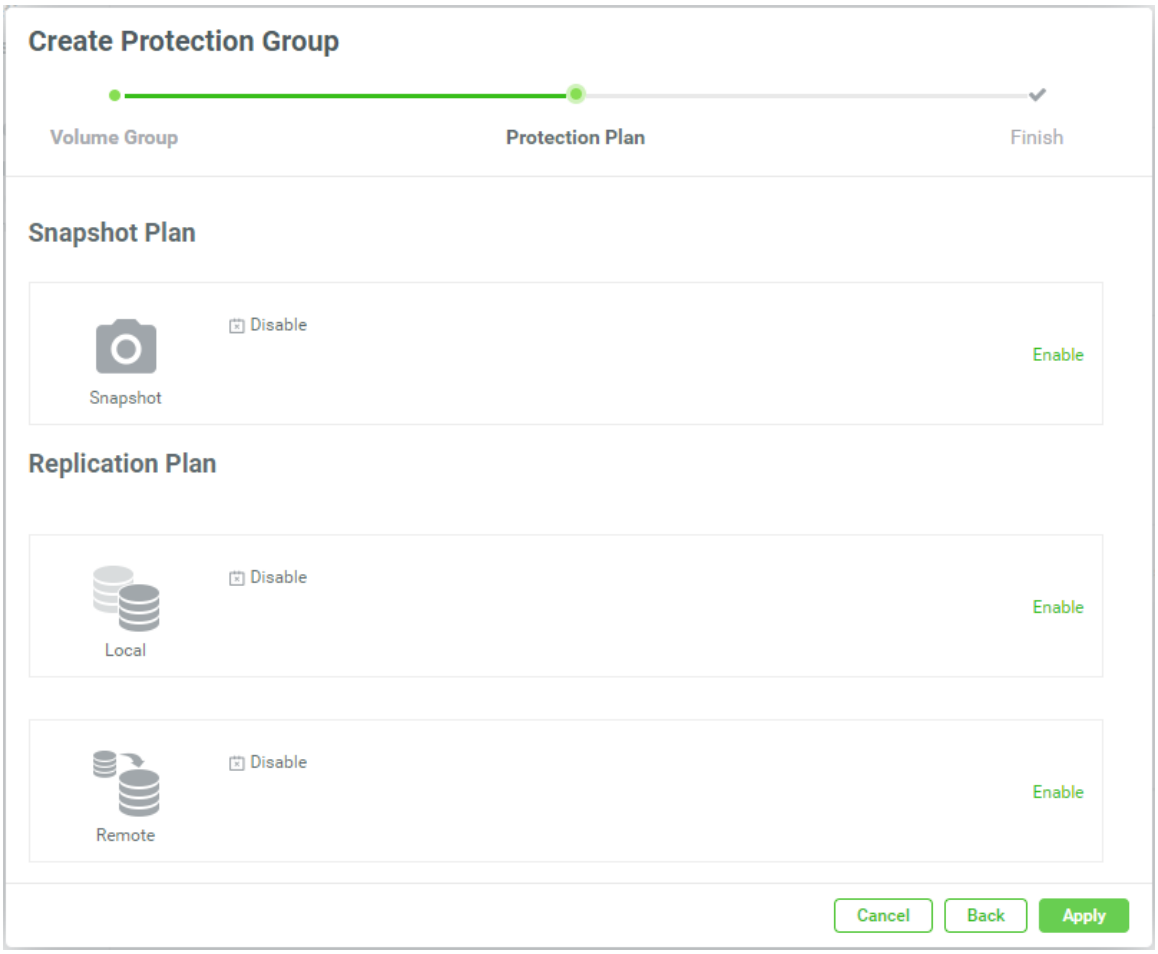

Figure 3-4 Synchronous Replication – Auto Configuration Step 2

4. Click the **Enable** text link in the **Remote** pane to enable the **Replication Plan - Remote**.

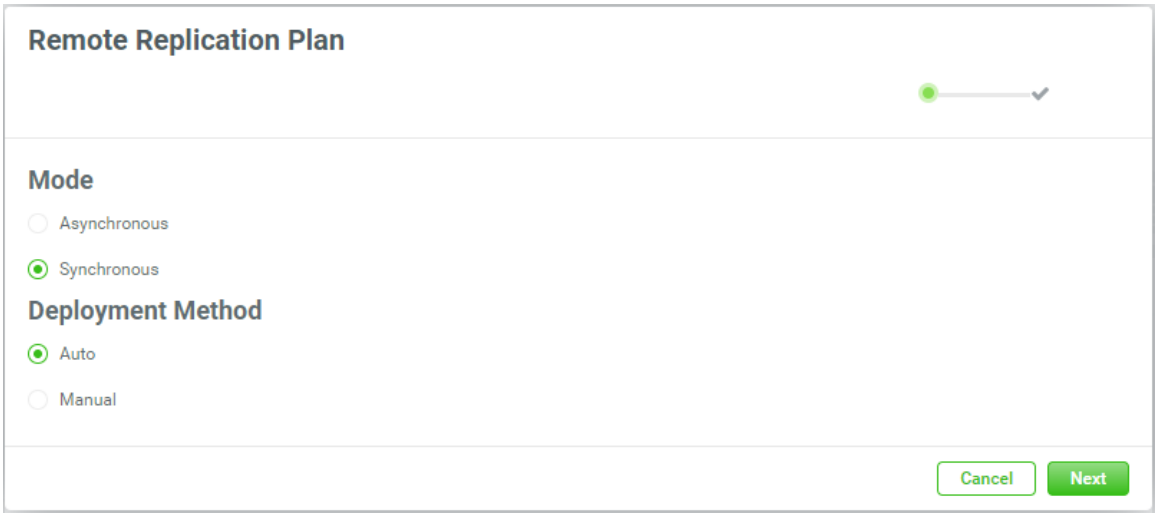

Figure 3-5 Synchronous Replication – Auto Configuration Step 3

6. Select the **Mode** as **Synchronous** and select the **Deployment Method** as **Auto**.
7. Click the **Next** button to continue.

INFORMATION

Synchronous replication requires a license. Please contact your local sales to obtain a license. If there is no license, the synchronization option is not visible.

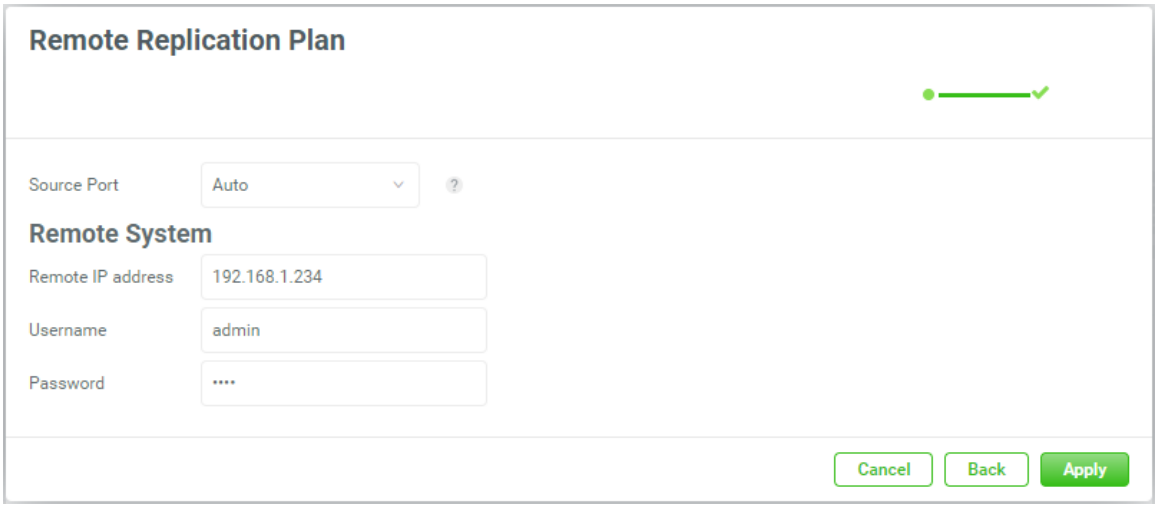

Figure 3-6 Synchronous Replication – Auto Configuration Step 4

8. If necessary, select the **Source Port** or leave it as Auto.
9. Login to the management port of the target unit. Enter the **Remote IP address** of the target management port.
10. Enter the **Username** and **Password** of the target administrator account.
11. Click the **Apply** button to finish the **Replication Plan - Remote** configuration.

TIP

The source port for synchronous must be greater than or equal to 10 Gbps. If you do not want to assign a fixed port, please leave the setting of **Source Port** as **Auto**. The system will try to connect to the target IP address automatically.

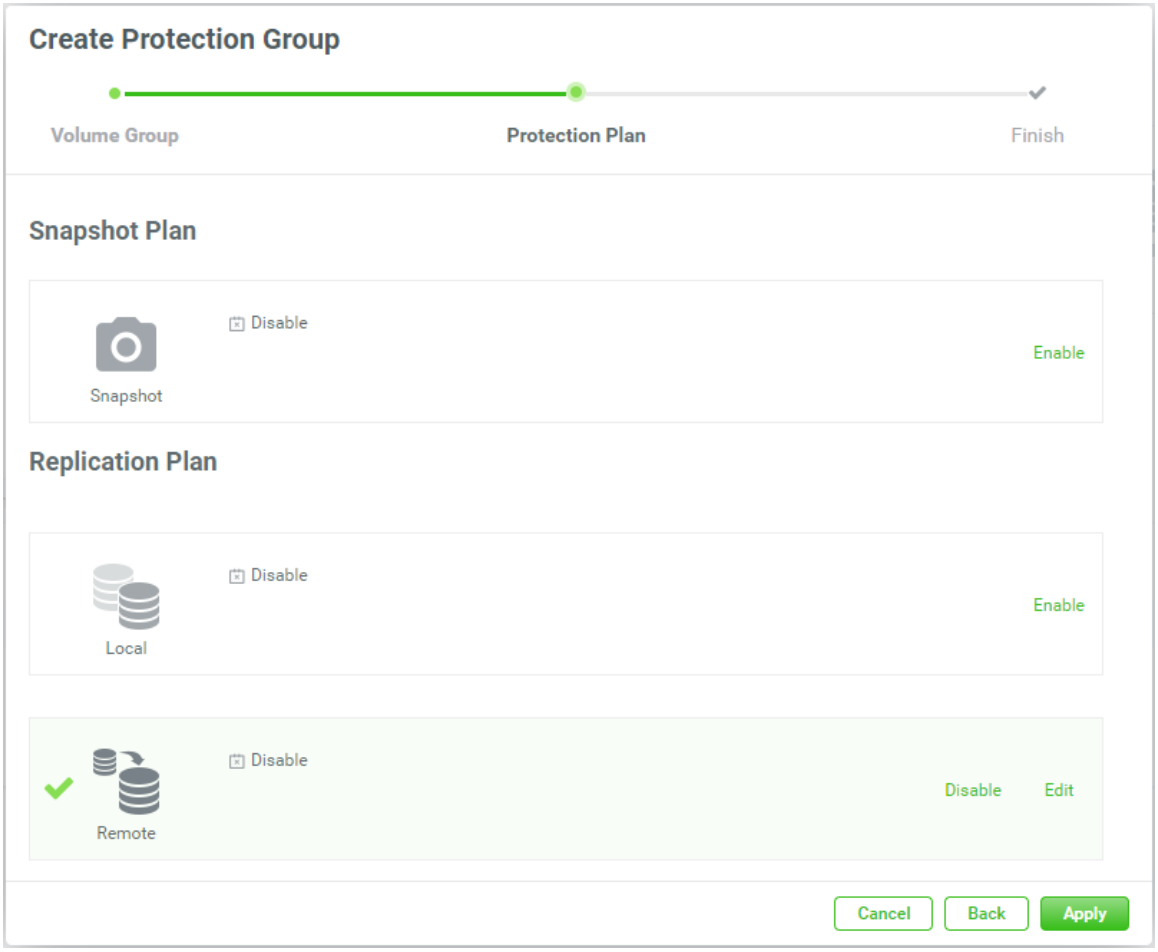

Figure 3-7 Synchronous Replication – Auto Configuration Step 5

12. Back to main screen of the wizard. Click the **Apply** button to create a replication task.

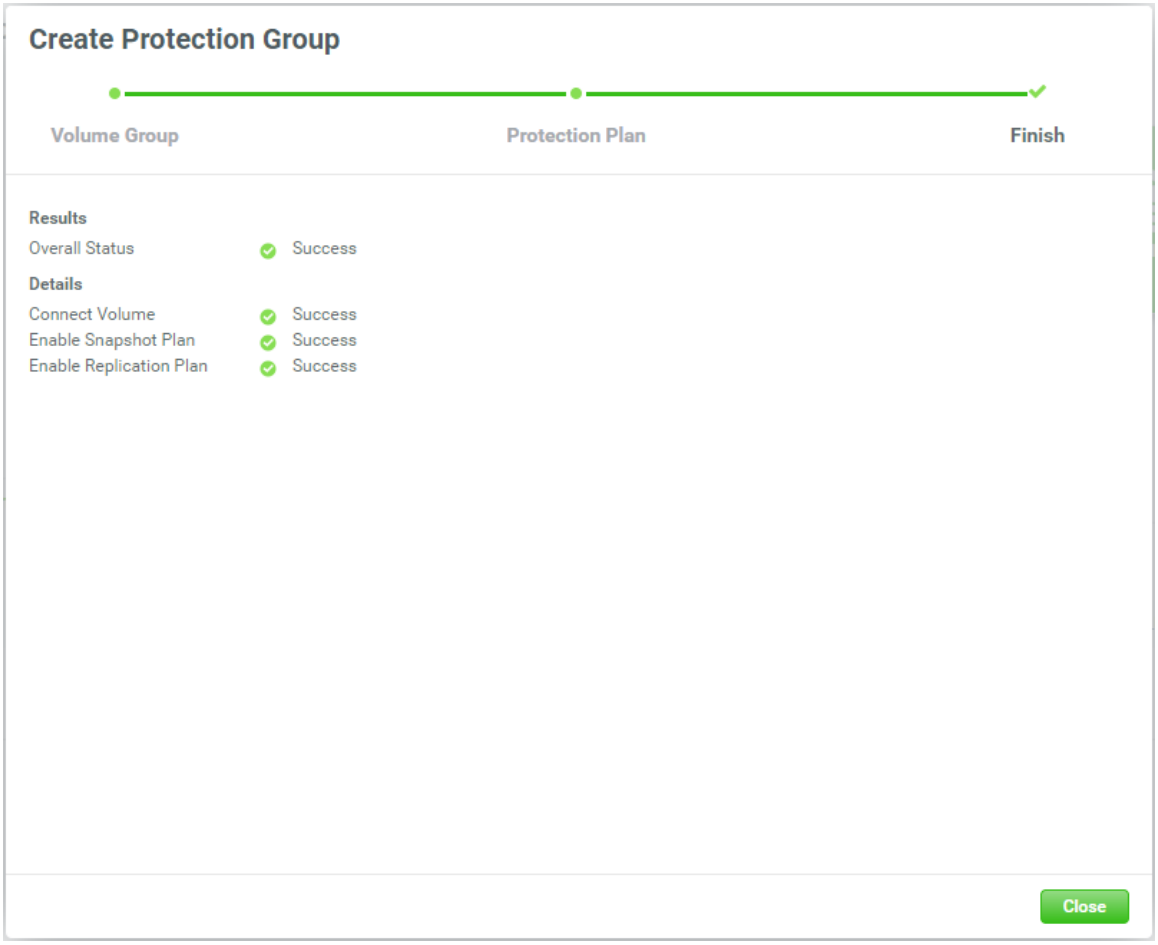

Figure 3-8 Synchronous Replication – Auto Configuration Step 6

13. There is a result page. Click the **Close** button to close the wizard.

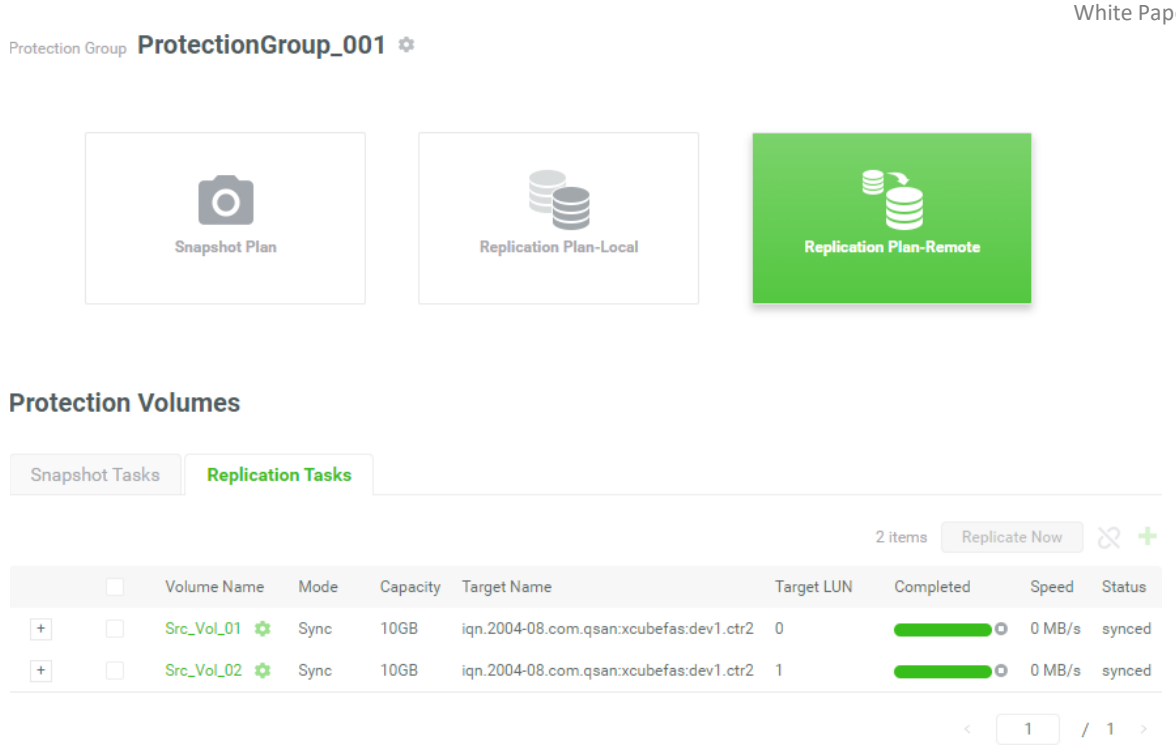

Figure 3-9 Synchronous Replication – Auto Configuration Step 7

14. Done. A protection group with synchronous replication task has been created.

TIP

When the synchronous replication task has been created, the full replication will start automatically.

3.3.2. Manual Configuration

Here is an example of creating a synchronous replication task with manual configuration.

Prepare Backup Volumes in Site B (Target Unit)

1. Create a pool and multiple volumes with capacity greater than or equal to the source volumes.

2. Click the icon beside each volume name to list the drop-down options, click the **Properties** option to change the volume properties.

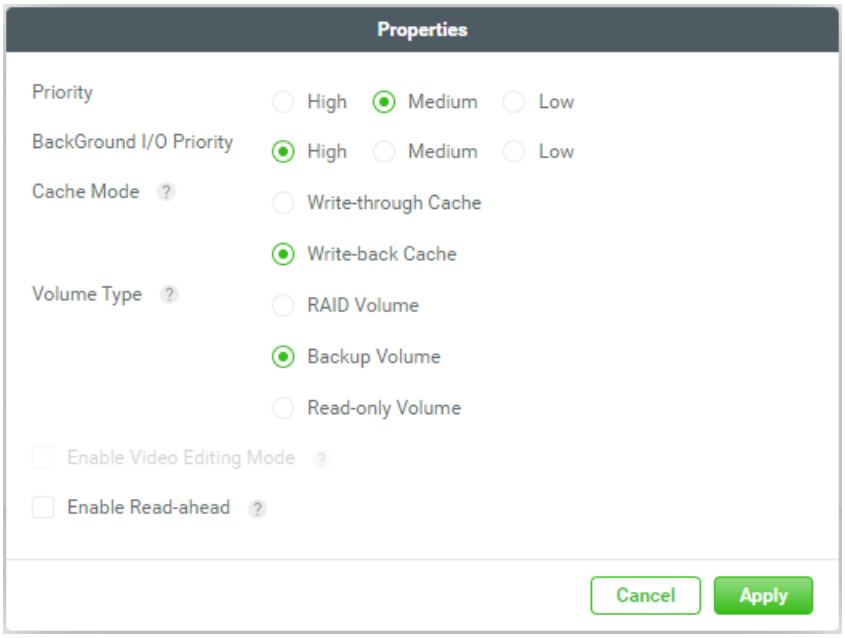

Figure 3-10 Synchronous Replication – Manual Configuration Step 1

3. Change the **Volume Type** to **Backup Volume**, and then click the **Apply** button to take effect.

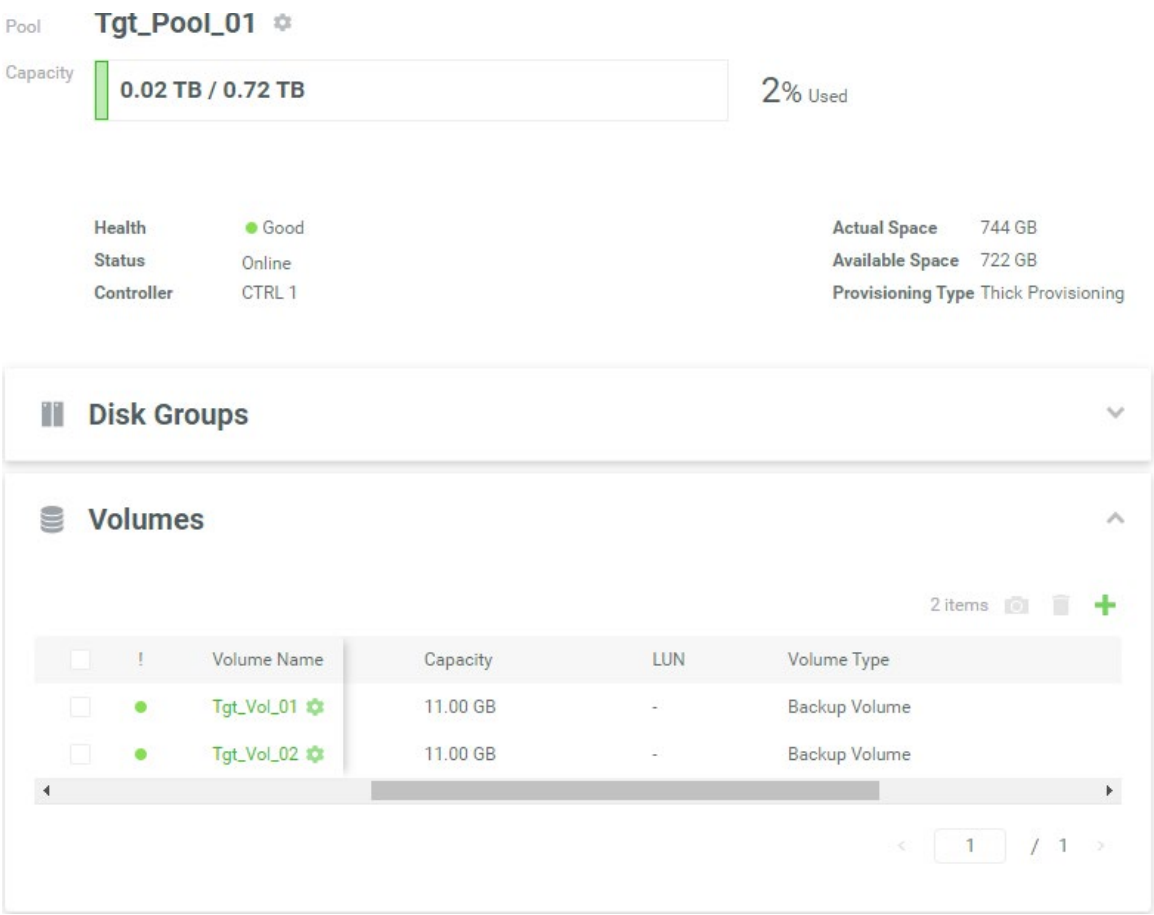

Figure 3-11 Synchronous Replication – Manual Configuration Step 2

4. Here are the target volumes.
5. Create a host group with the target volumes in the **Hosts** tab.

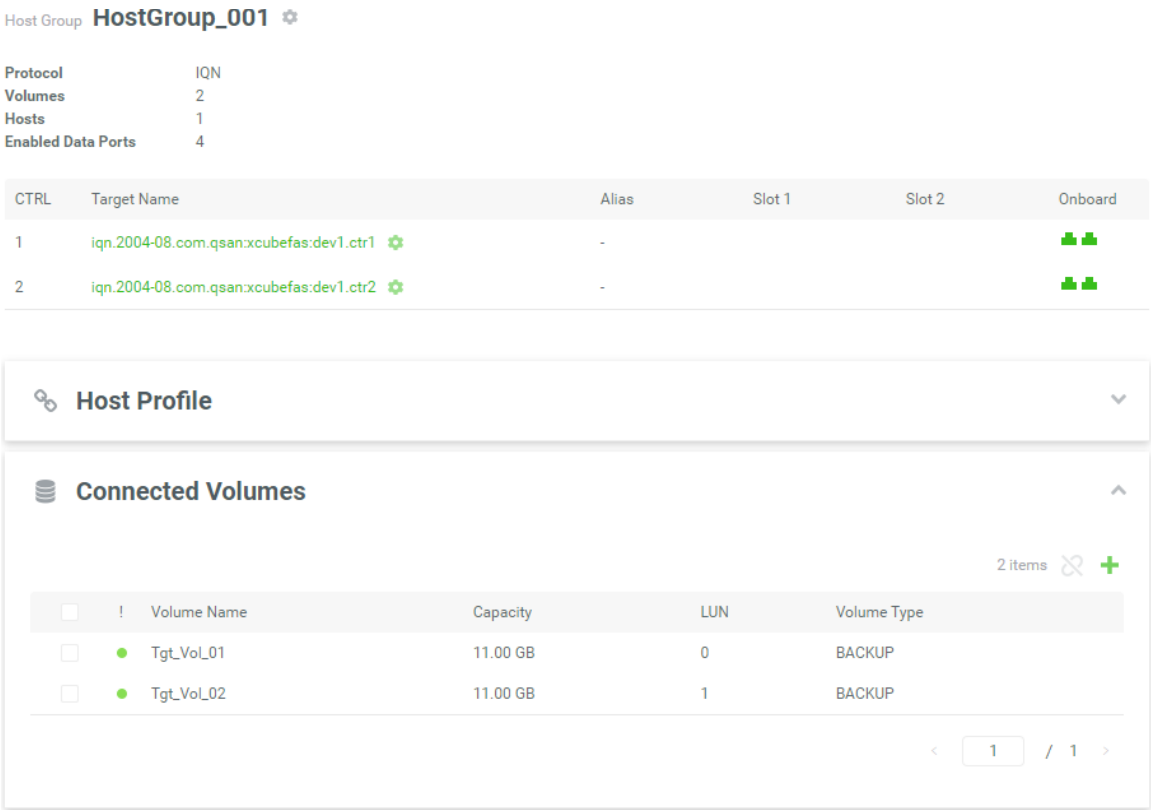

Figure 3-12 Synchronous Replication – Manual Configuration Step 3

6. Here is the host group for target volumes.

Configure Replication Plan-Remote in Site A (Source Unit)

7. Click the + icon in the **Protection Groups** pane to pop up a wizard.

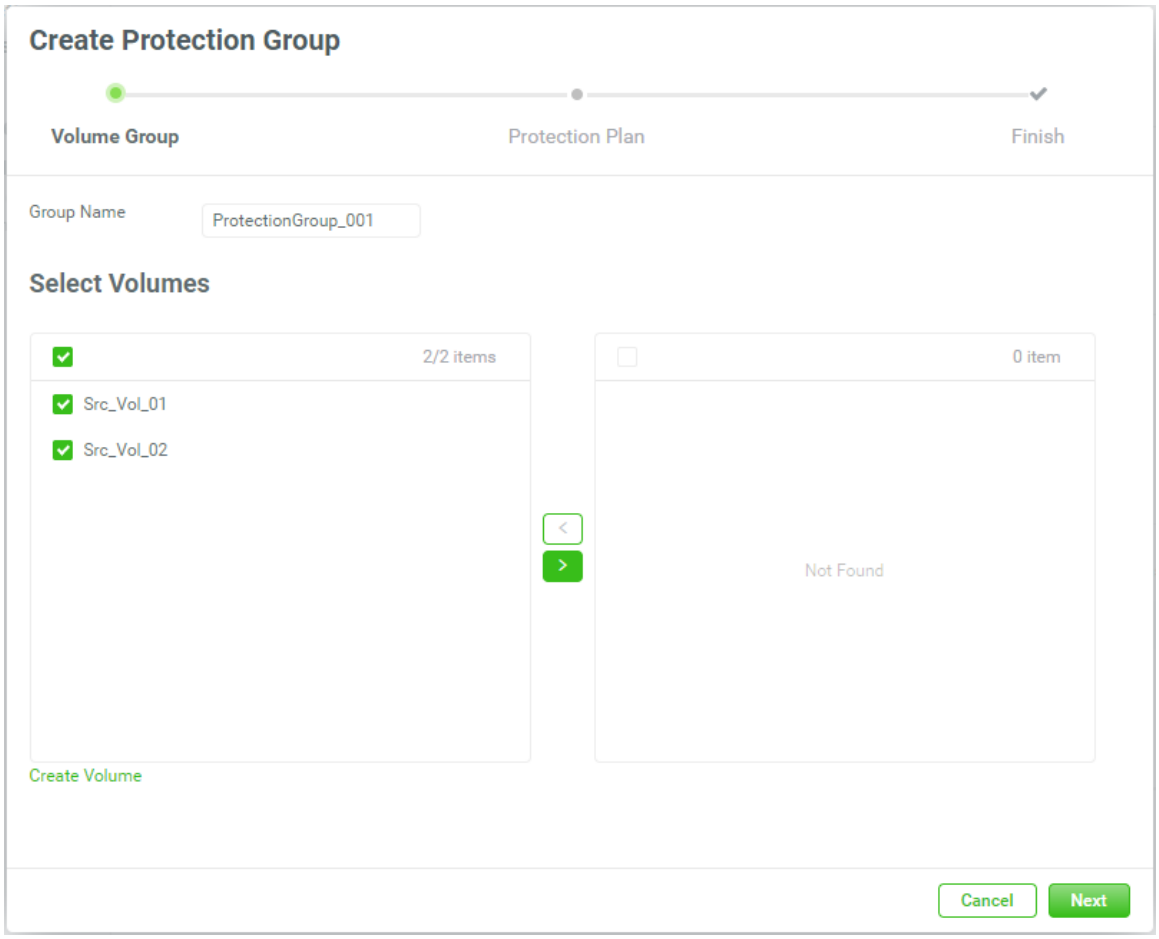

Figure 3-13 Synchronous Replication – Manual Configuration Step 4

8. Select source volumes, and then click the button to move them from the left side to the right side.
9. Click the **Next** button to continue.

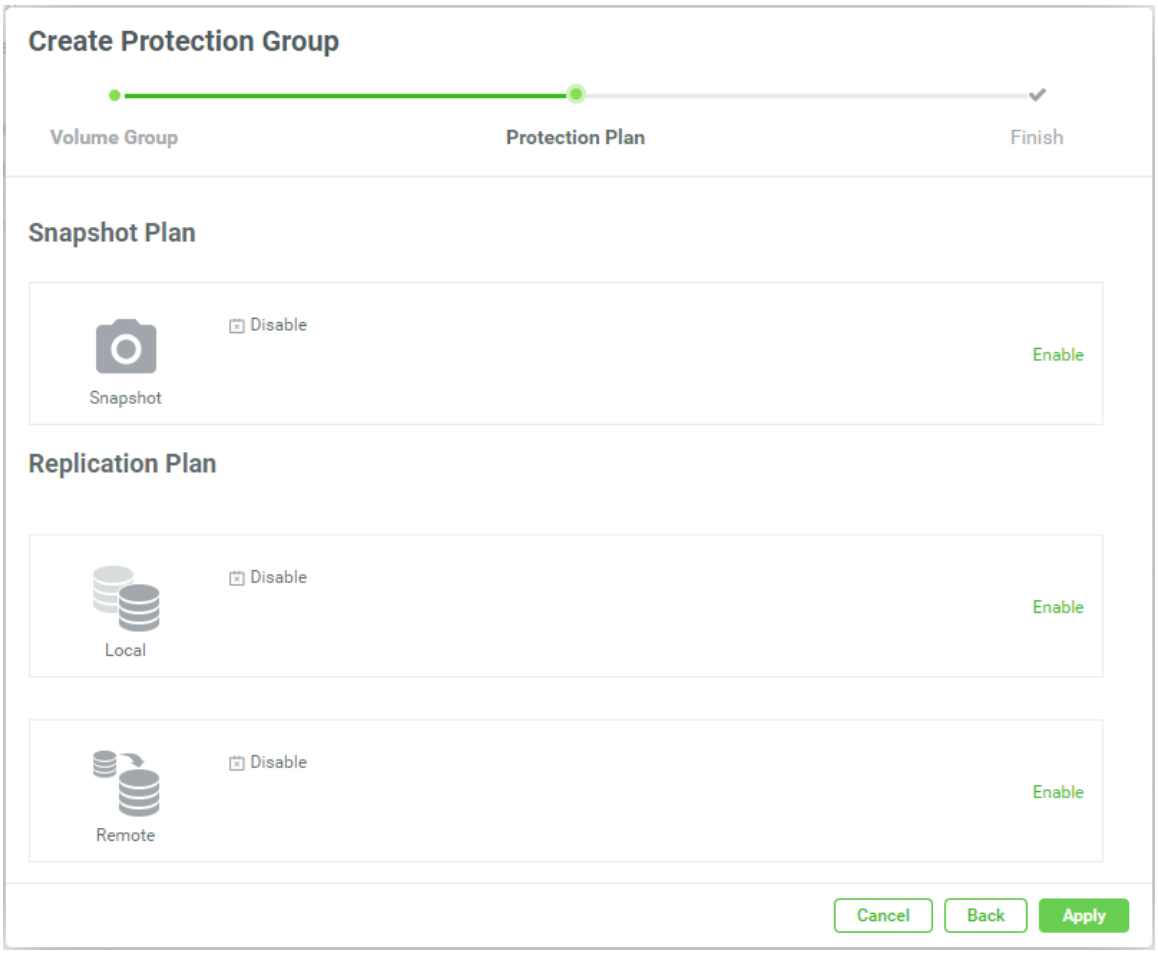

Figure 3-14 Synchronous Replication – Manual Configuration Step 5

10. Click the **Enable** text link in the **Remote** pane to enable the **Replication Plan - Remote**.

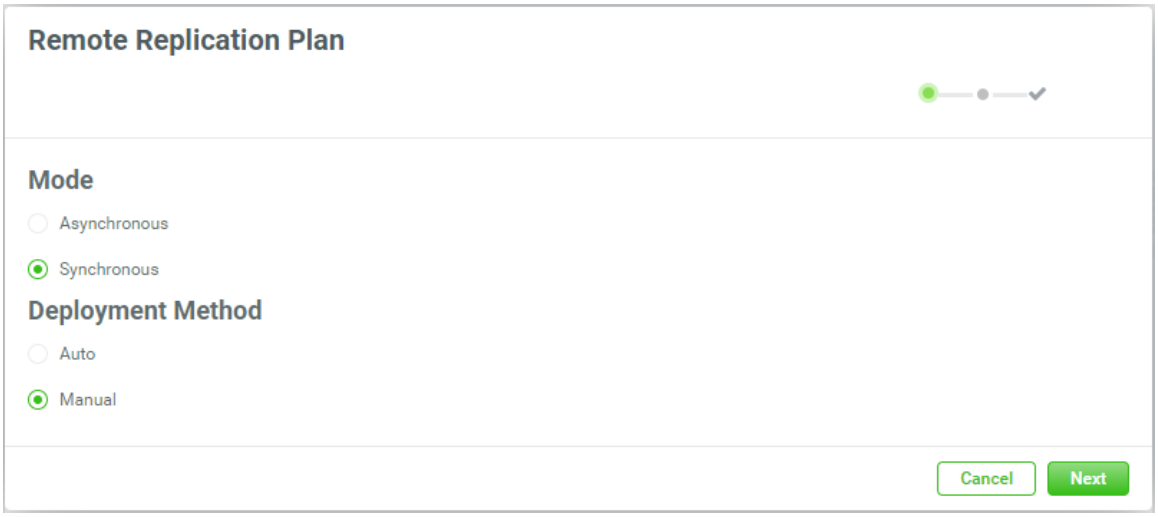

Figure 3-15 Synchronous Replication – Manual Configuration Step 6

11. Select the **Mode** as **Synchronous** and select the **Deployment Method** as **Manual**.
12. Click the **Next** button to continue.

INFORMATION

Synchronous replication requires a license. Please contact your local sales to obtain a license. If there is no license, the synchronization option is not visible.

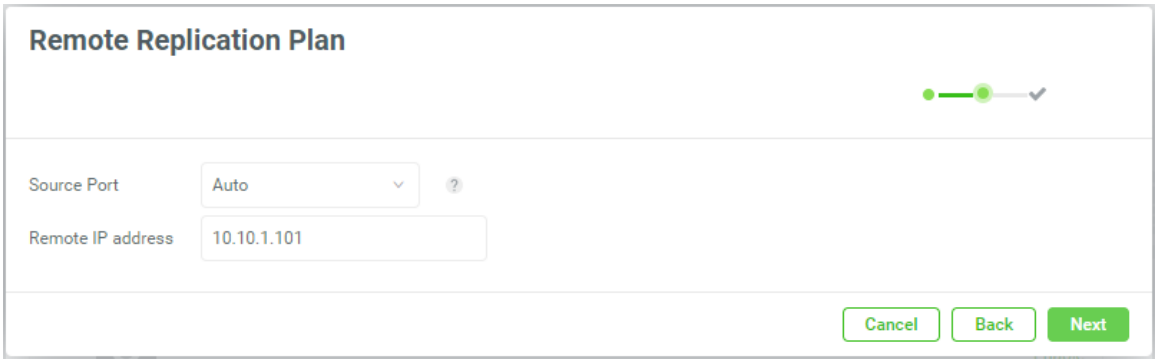

Figure 3-16 Synchronous Replication – Manual Configuration Step 7

13. If necessary, select the **Source Port** or leave it as Auto.

14. Enter the **Remote IP address** of the target data port.
15. Click the **Next** button to continue.

TIP

The source port for synchronous must be greater than or equal to 10 Gbps. If you do not want to assign a fixed port, please leave the setting of **Source Port** as **Auto**. The system will try to connect to the target IP address automatically.

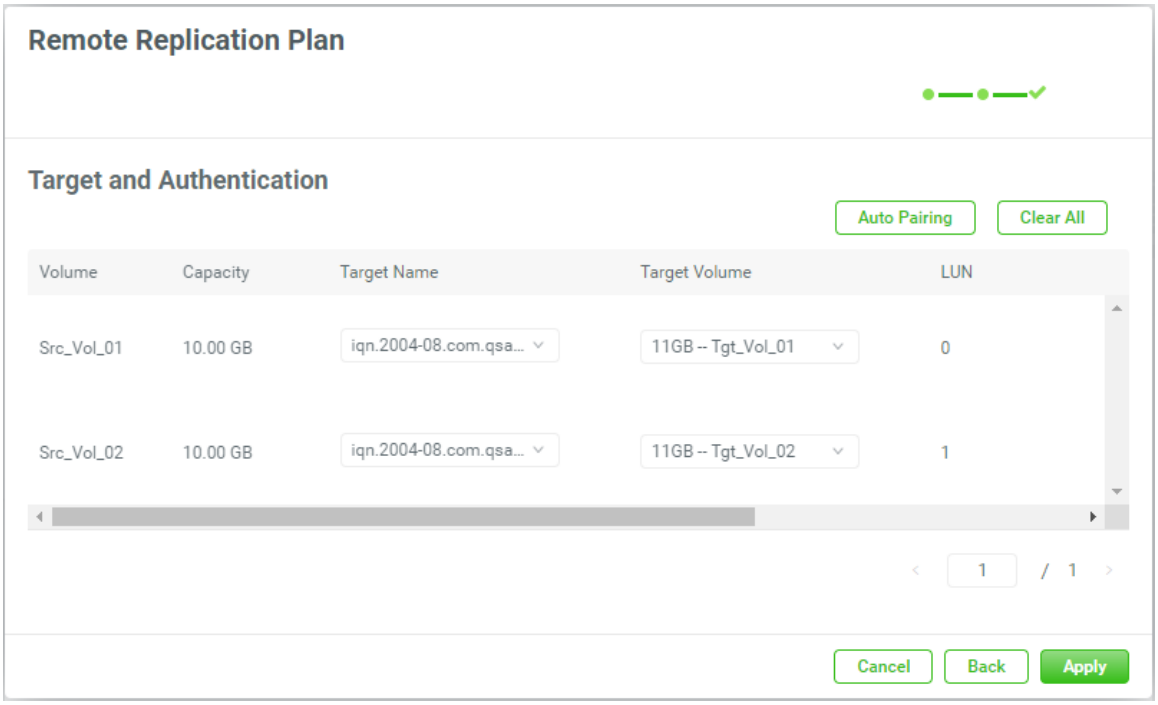

Figure 3-17 Synchronous Replication – Manual Configuration Step 8

16. Click the **Auto Pairing** button to automatically pair the backup volumes. If they are not what you want, manually select the **Target Name** and **Target Volume**.
17. If the target encrypts with CHAP, click the icon next to the Target Name, it will pop up a dialog box to enter the CHAP information.
18. Click the **Apply** button to finish the **Replication Plan - Remote** configuration.

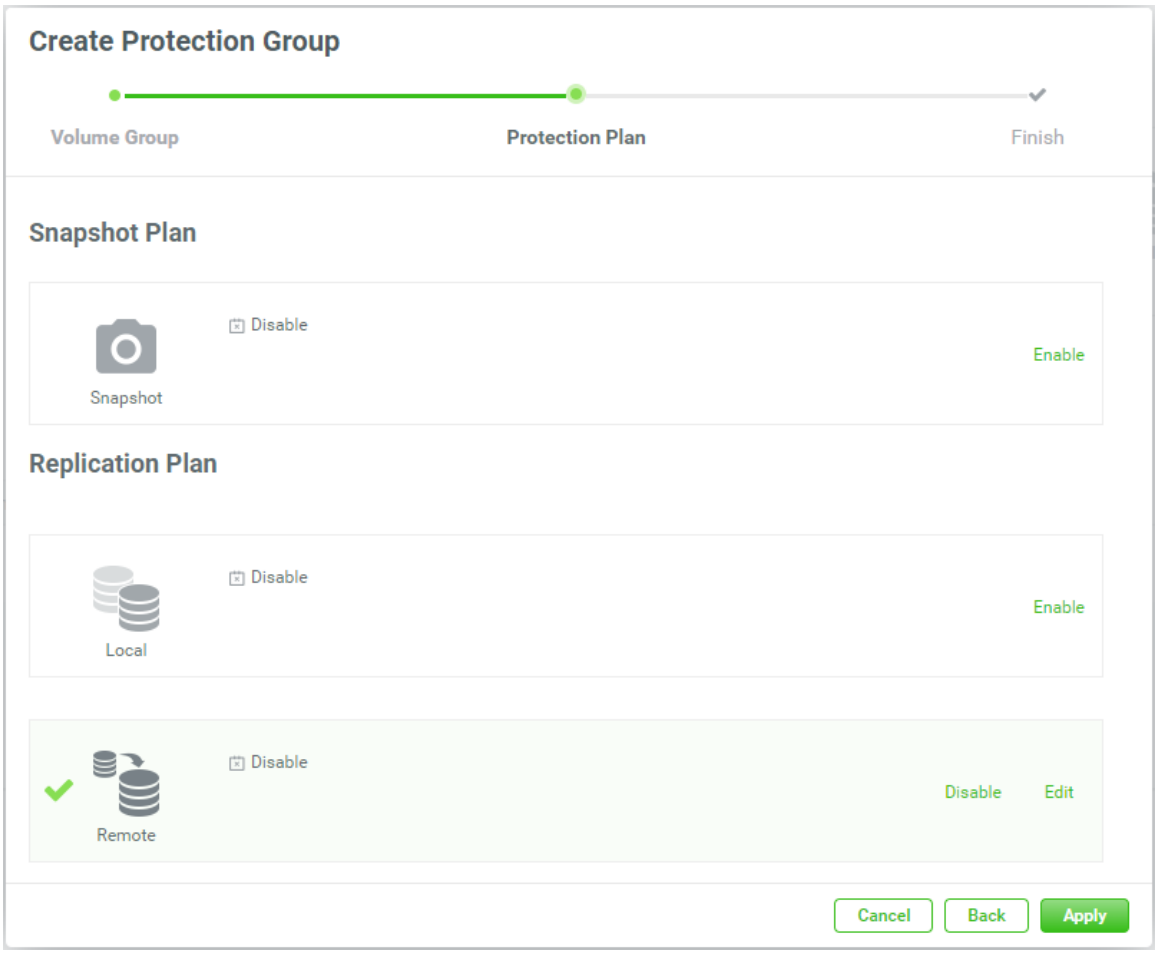

Figure 3-18 Synchronous Replication – Manual Configuration Step 9

19. Back to main screen of the wizard. Click the **Apply** button to create a replication task.

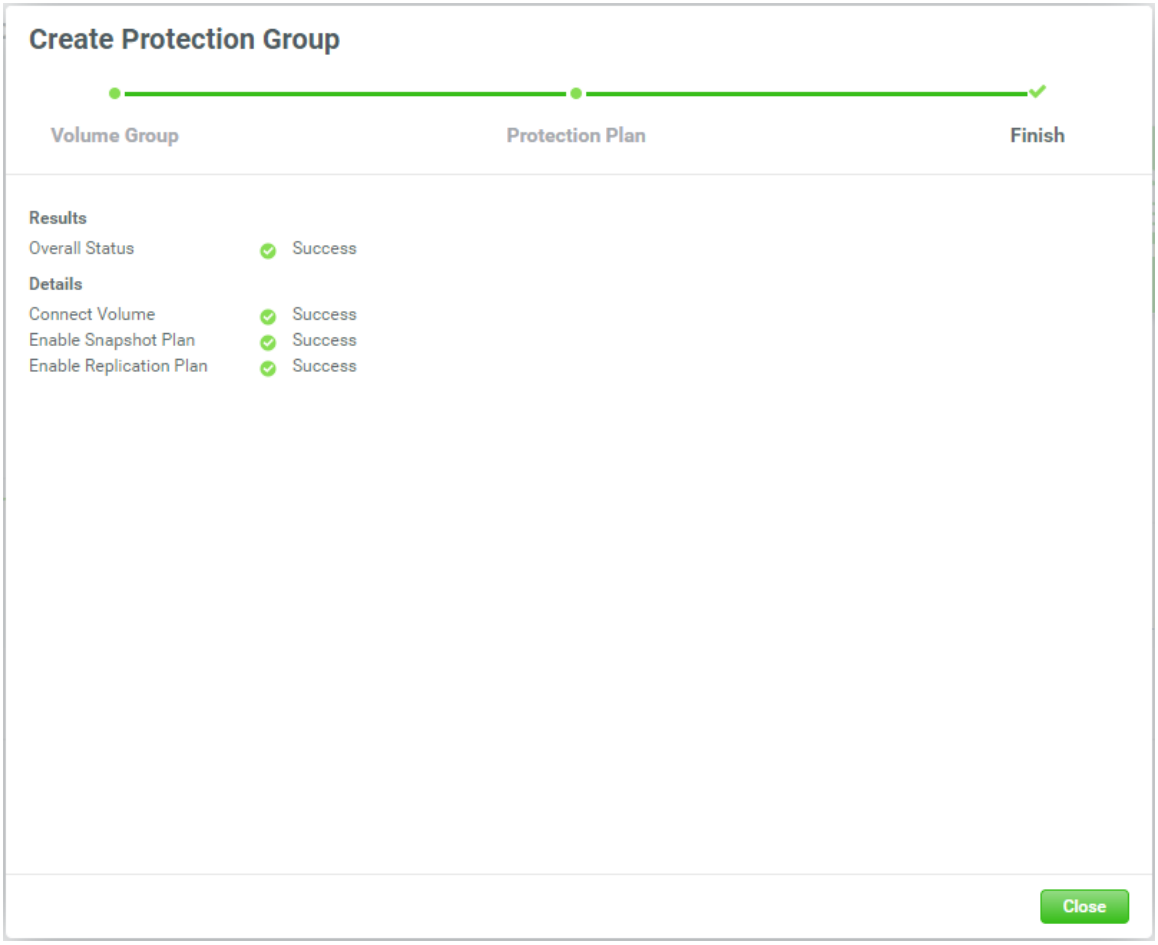

Figure 3-19 Synchronous Replication – Manual Configuration Step 10

20. There is a result page. Click the **Close** button to close the wizard.

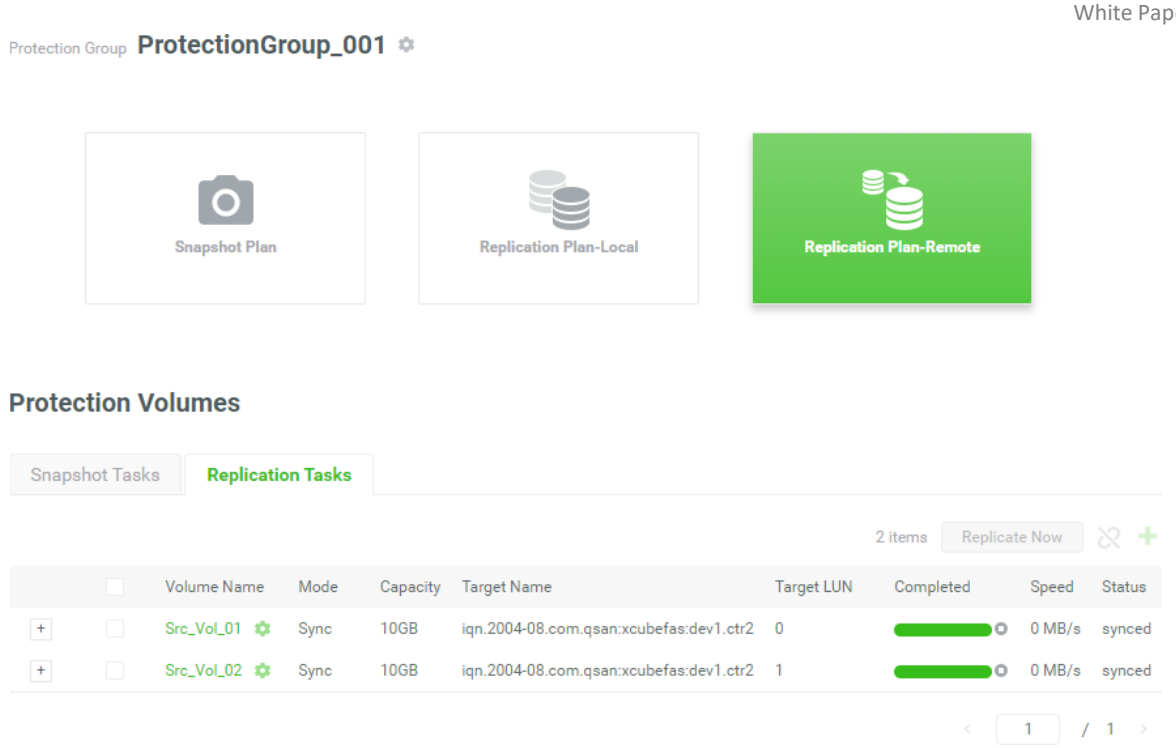

Figure 3-20 Synchronous Replication – Manual Configuration Step 11

21. Done. A protection group with synchronous replication task has been created.

TIP

When the synchronous replication task has been created, the full replication will start automatically.

3.3.3. Replication Options

The options are available in this tab.

Stop Replication Task

Click the icon of the volume to stop the replication task.

Start Replication Task

Select a volume, and then click the ▶ icon to replicate the volume immediately. It will pop up dialog box asking if you want to resume the task or perform a full replication. After selecting, click the **OK** button.

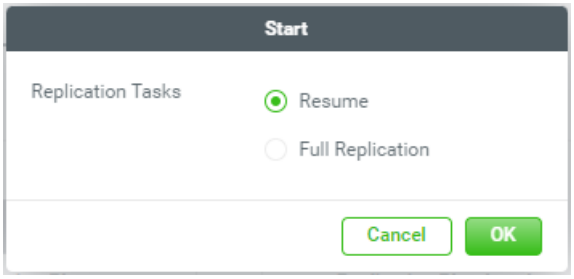

Figure 3-21 Synchronous Replication – Start Dialog

The option **Resume** will replicate the out-of-sync part. If you can not sure whether the target has been touched, you can also perform a **Full Replication**. The progress bar displays the current status. .

Ungroup Volume from Protection Group

To ungroup the volume from the protection group, check one or multiple volume names, and then click the button.

Delete Protection Group

To delete the protection group, click the icon beside the protection group name to list the drop-down options. Click the **Delete** option; it will pop up dialog box to confirm. Click the **Delete** button to delete the protection group.

3.3.4. Enable Multipath and Add Connections

Click the icon next to the volume name to list the drop-down options. Click the **Connection Properties** option to enable multipath and add connections for the replication task.

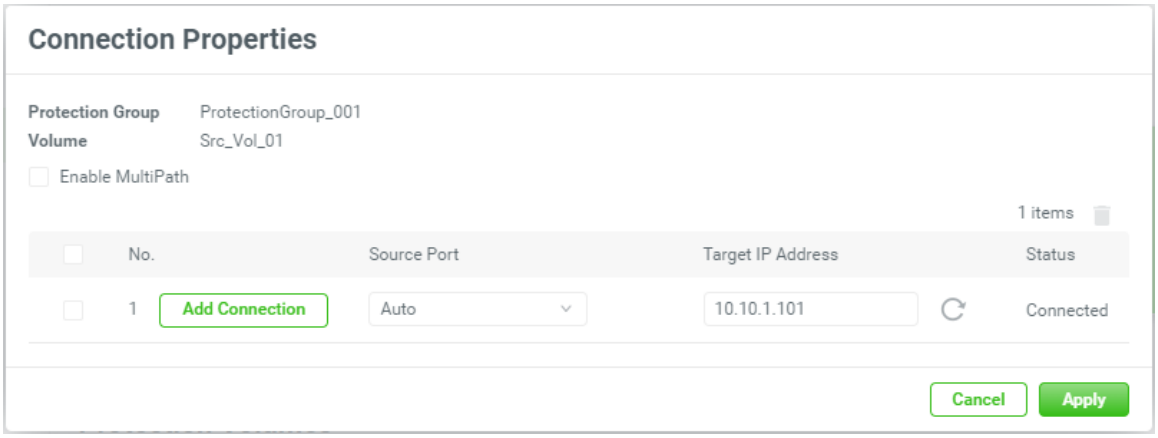

Figure 3-22 Synchronous Replication – Connection Properties

Enable Multipath

1. Check the **Enable MultiPath** checkbox to enable multipath for the replication task.

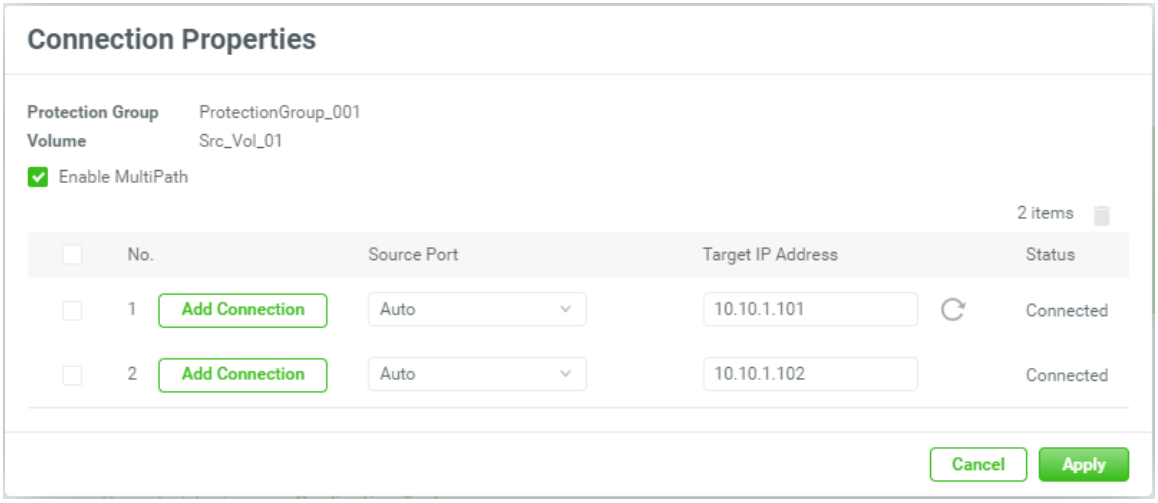

Figure 3-23 Synchronous Replication – Enable Multipath

2. If necessary, select the **Source Port** or leave it as Auto.
3. Enter the **Target IP address** of the target data port.
4. Click the **Apply** button to enable.

Add Connections

1. Click the **Add Connection** button to add another connection.

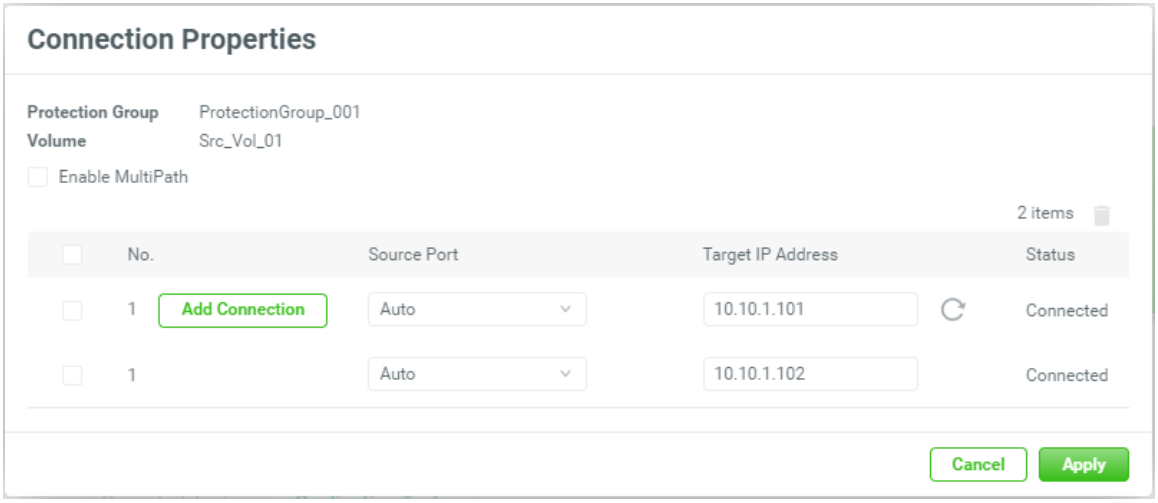

Figure 3-24 Synchronous Replication – Add Connection

2. If necessary, select the **Source Port** or leave it as Auto.
3. Enter the **Target IP address** of the target data port.
4. Click the **Apply** button to add.

3.3.5. Switch Local Clone to Remote Replication

Through local-to-remote, you can transfer local backup to remote sites without redoing full copy. You can plug and play when local-to-remote without any configurations.

Configure Replication Plan-Local in Site A (Source Unit)

1. Click the ✚ icon in the **Protection Groups** pane to pop up a wizard.

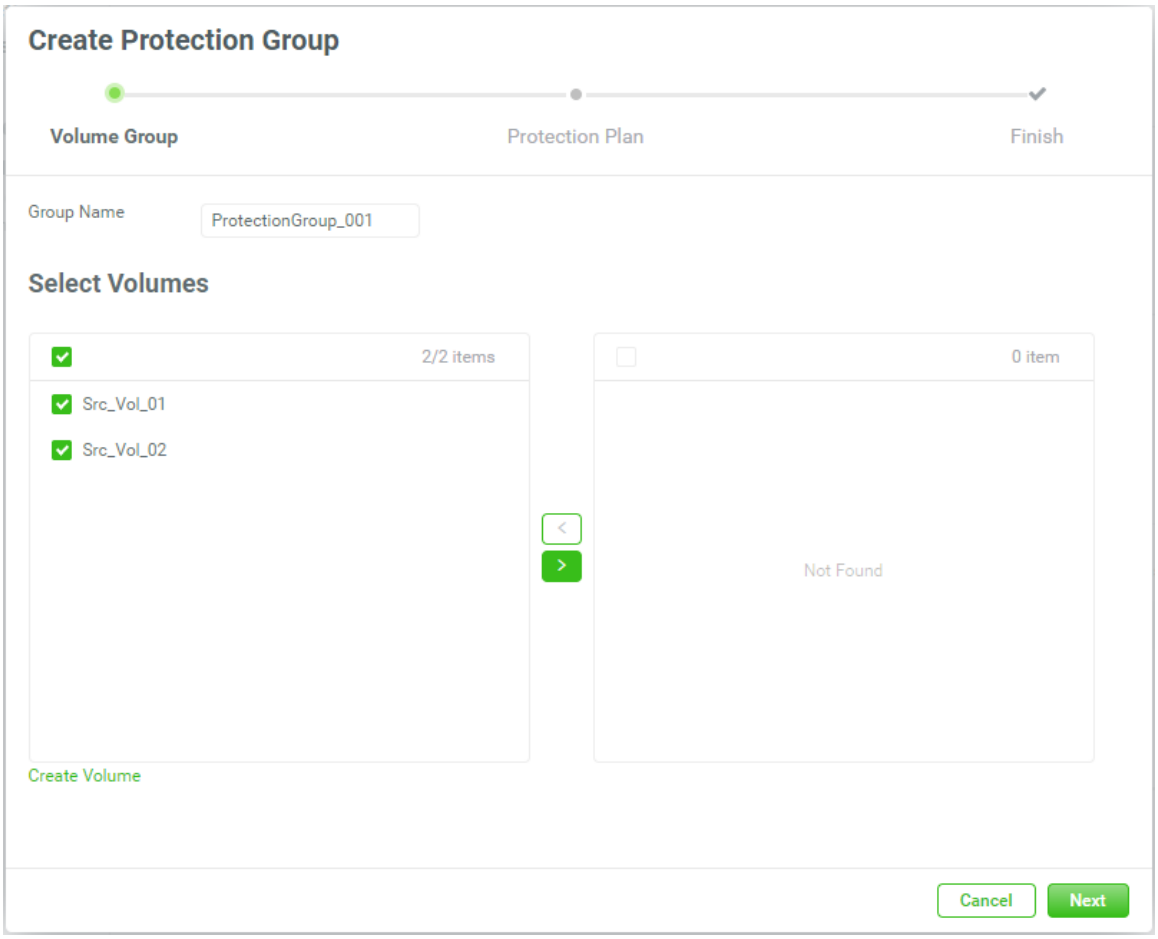

Figure 3-25 Synchronous Replication – Local to Remote Step 1

2. Select source volumes, and then click the button to move them from the left side to the right side.
3. Click the **Next** button to continue.

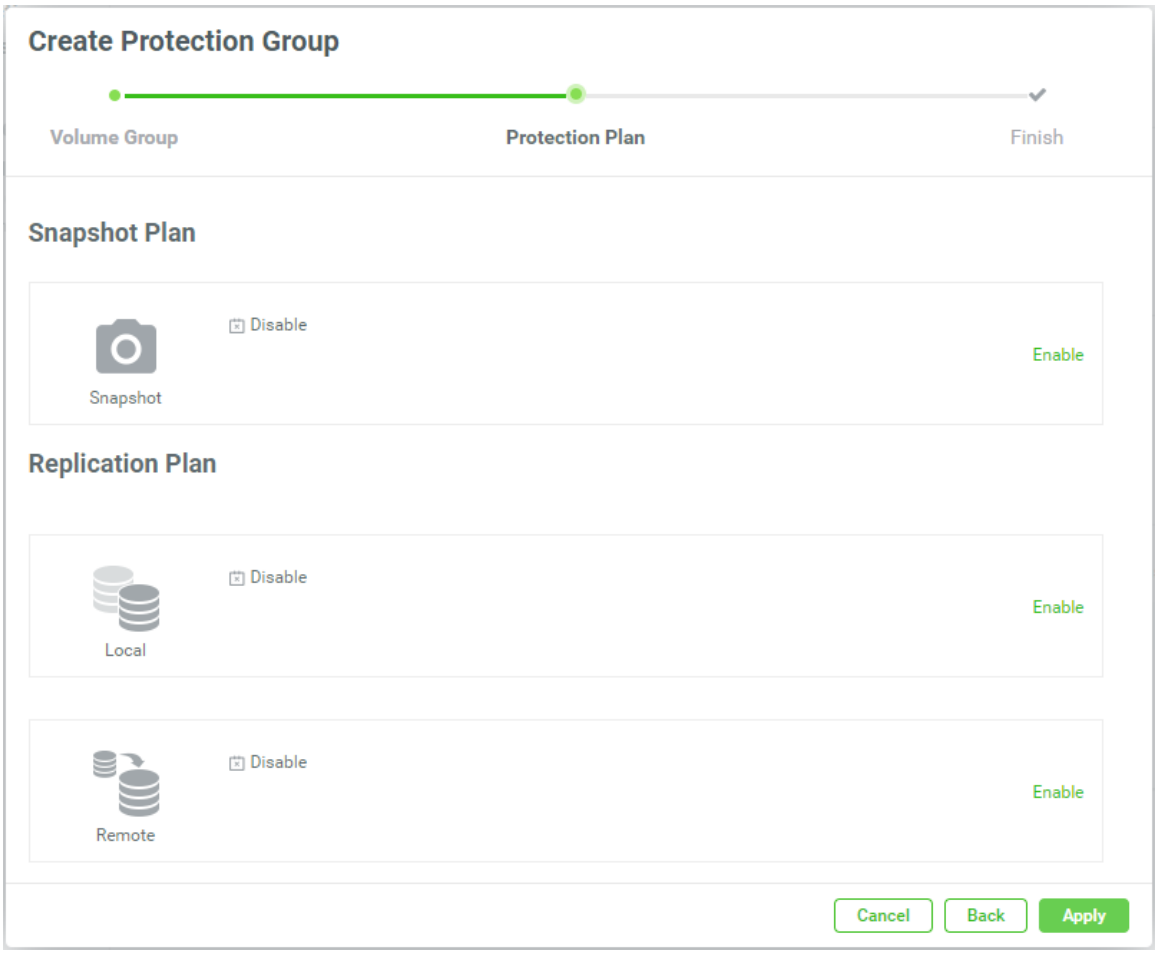

Figure 3-26 Synchronous Replication – Local to Remote Step 2

4. Click the **Enable** text link in the **Local** pane to enable the **Replication Plan - Local**.

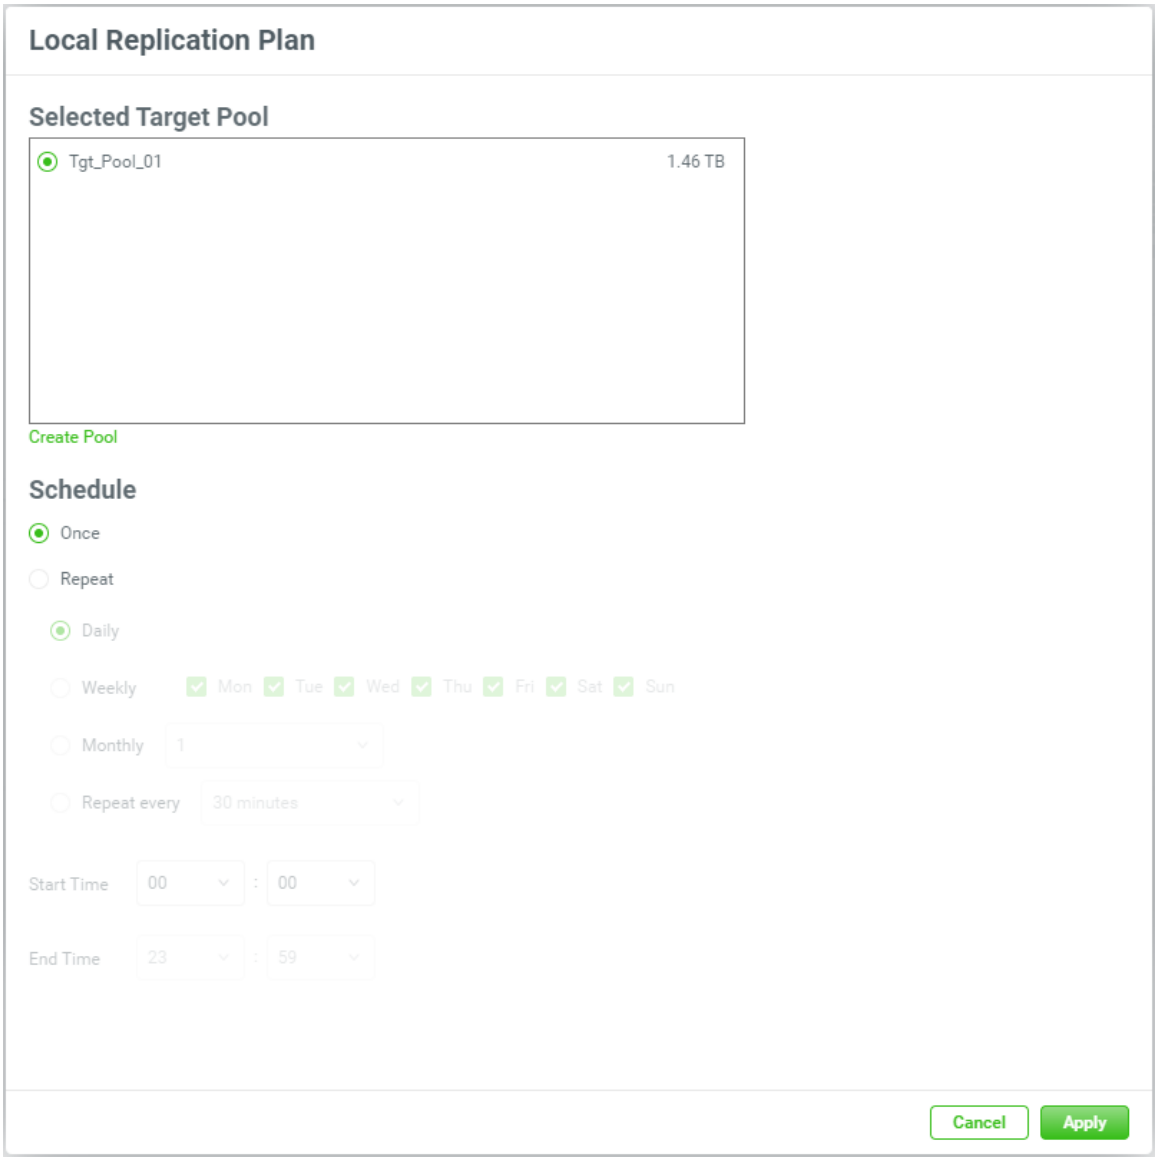

Figure 3-27 Synchronous Replication – Local to Remote Step 3

5. Select a target pool. If none can be selected, click the **Create Pool** link to create a pool.
6. Click the **Apply** button to create a local clone task. And then click the **Close** button to close the wizard.

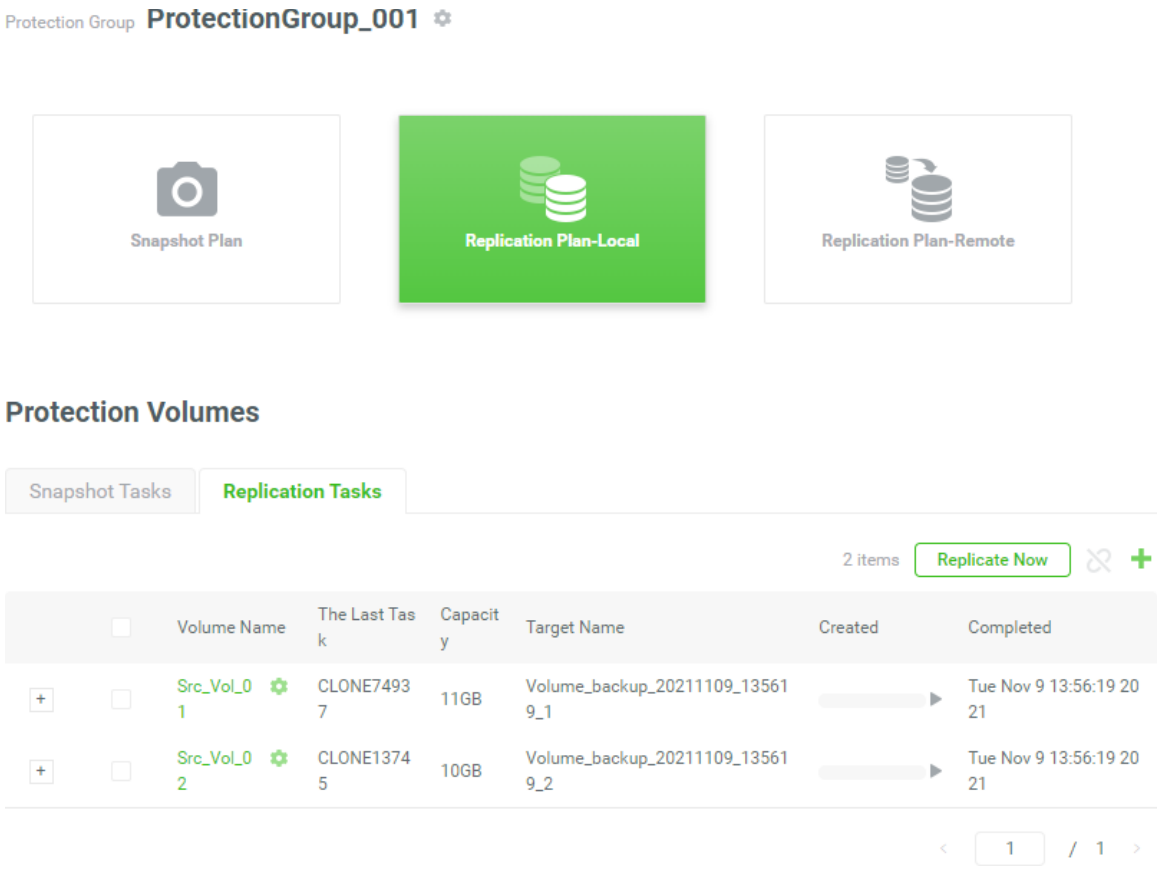

Figure 3-28 Synchronous Replication – Local to Remote Step 4

7. After the local clone is complete, click the icon next to the protection group name to list the drop-down options. And then click the **Protection Plan Setting** option.

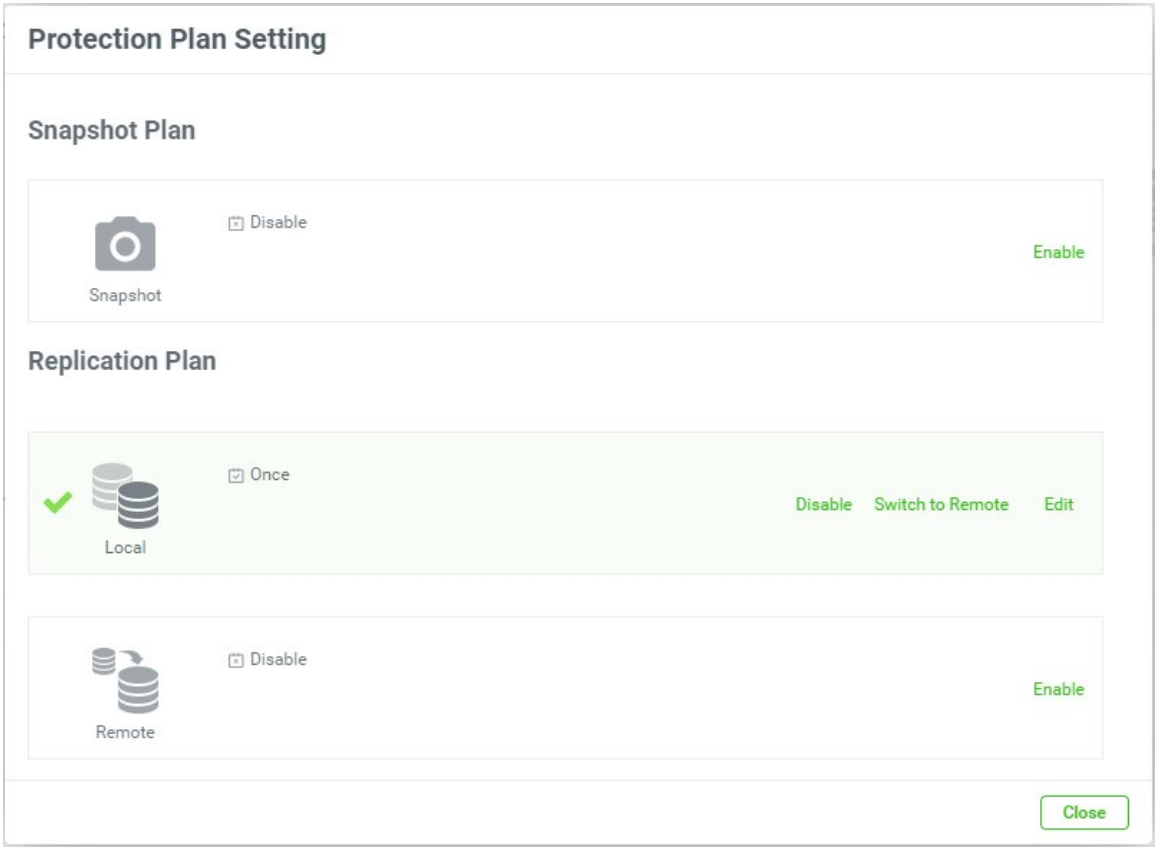

Figure 3-29 Synchronous Replication – Local to Remote Step 5

8. Click the **Switch to Remote** text to configure the replication plan.

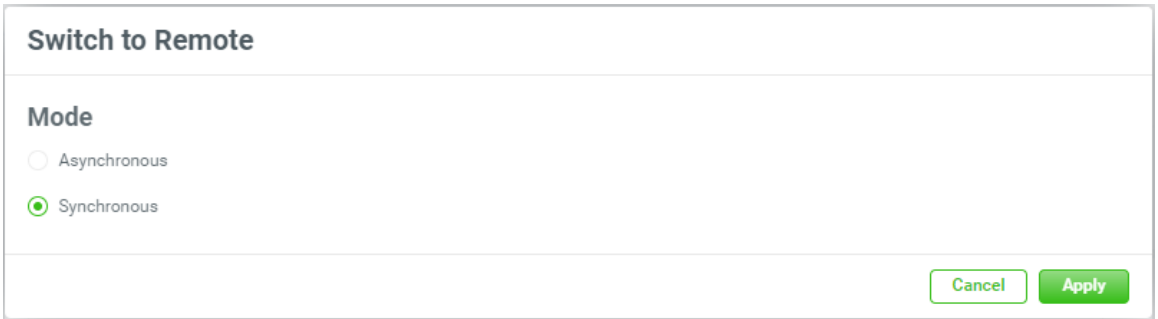

Figure 3-30 Synchronous Replication – Local to Remote Step 6

9. Select the **Mode** as **Synchronous**.
10. Click the **Apply** button and then click the **Close** button to close the wizard.
11. Deactivate the target pool.
12. Remove all disk drives of the target pool and deliver them to the Site B (Target Unit).

Prepare Backup Volumes in Site B (Target Unit)

13. Insert all disk drives of the target pool and activate them.
14. Create a host group to include the target volumes.

Configure Replication Plan-Remote in Site A (Source Unit)

15. Select the source volume of the protection group.

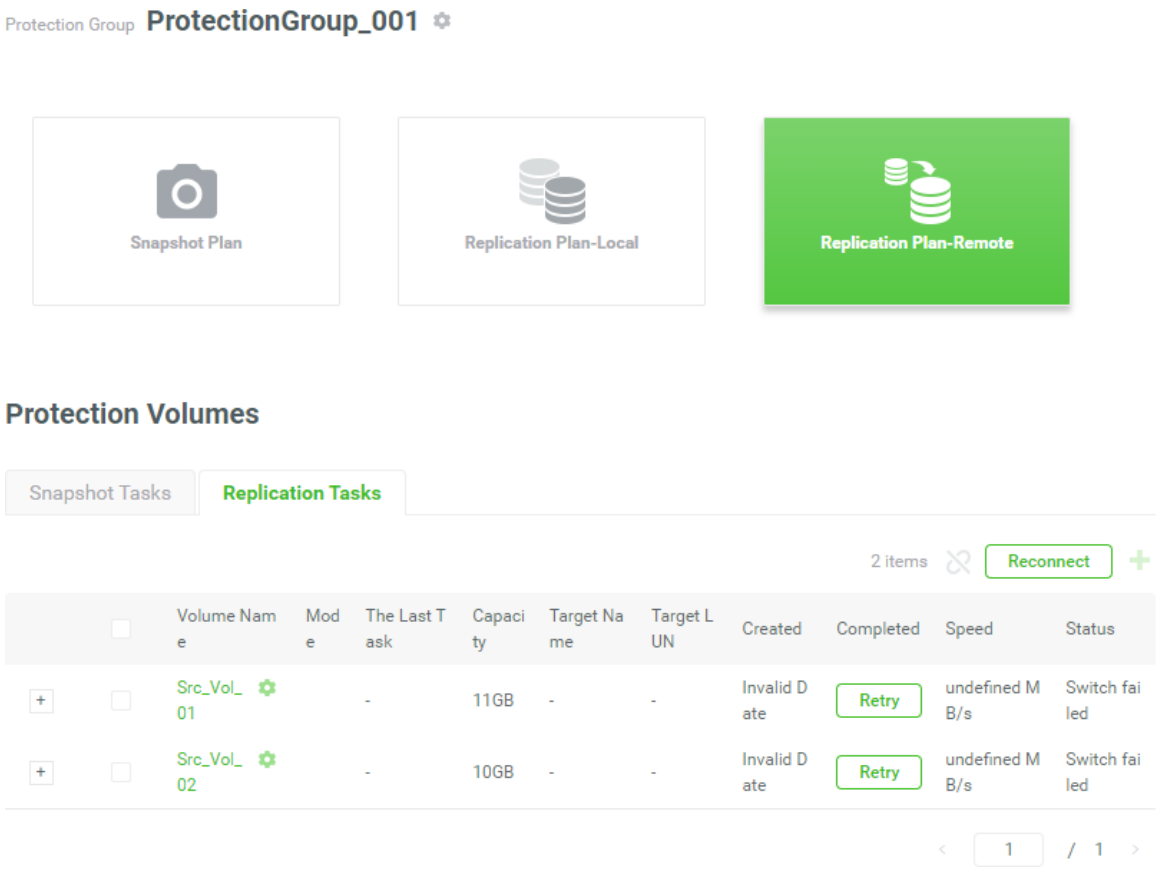

Figure 3-31 Synchronous Replication – Local to Remote Step 7

16. Click the **Reconnect** button to reconnect to the target volume. It will pop up a wizard.

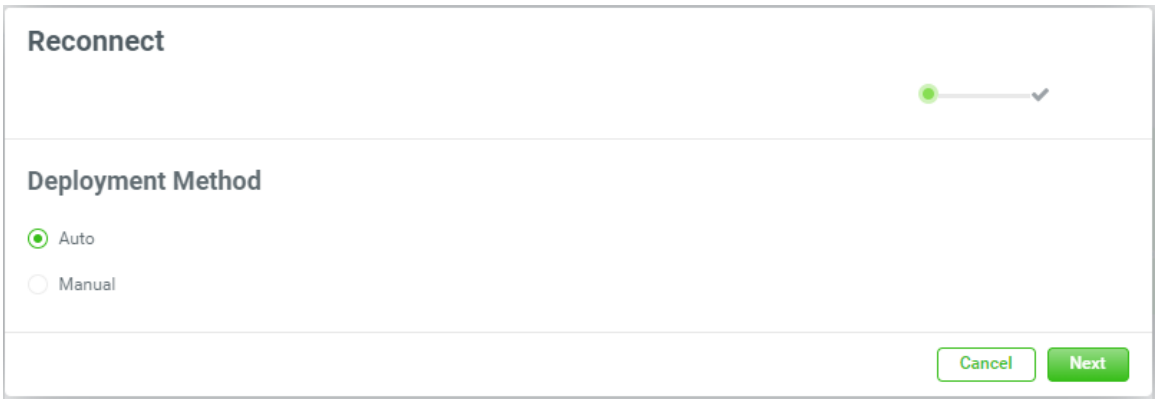

Figure 3-32 Synchronous Replication – Local to Remote Step 8

17. Select the **Deployment Method** as **Auto** or **Manual**.
18. Click the **Next** button to continue.
19. The following takes **Auto** as an example, if you choose **Manual**, please refer to steps 13 to 18 in the 3.3.2 Manual Configuration section.

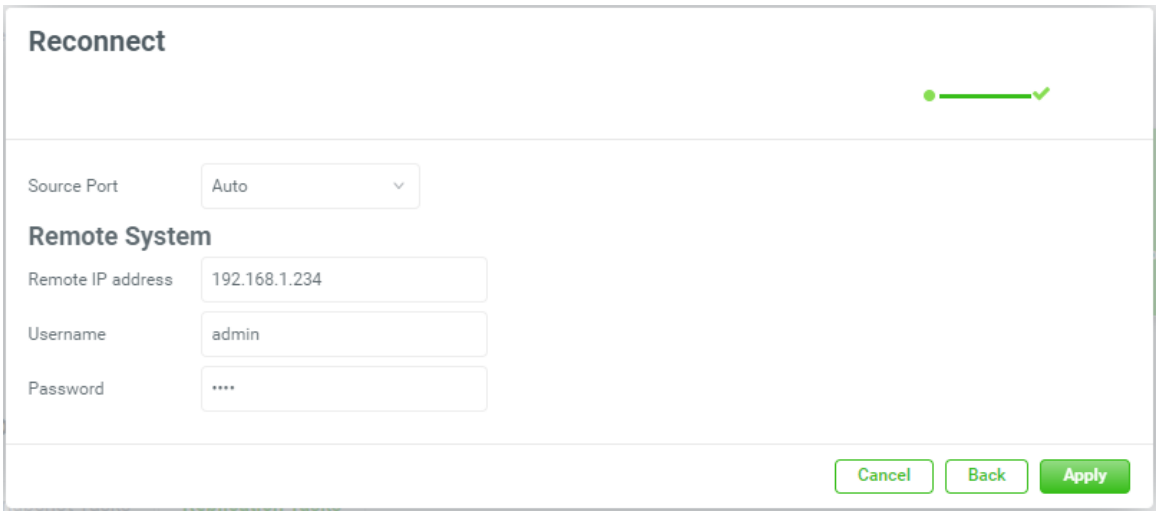

Figure 3-33 Synchronous Replication – Local to Remote Step 9

20. If necessary, select the **Source Port** or leave it as Auto.
21. Login to the management port of the target unit. Enter the **Remote IP address** of the target management port.
22. Enter the **Username** and **Password** of the target administrator account.
23. Click the **Apply** button to finish the **Reconnect** configuration.
24. Click the **Close** button to close the wizard.

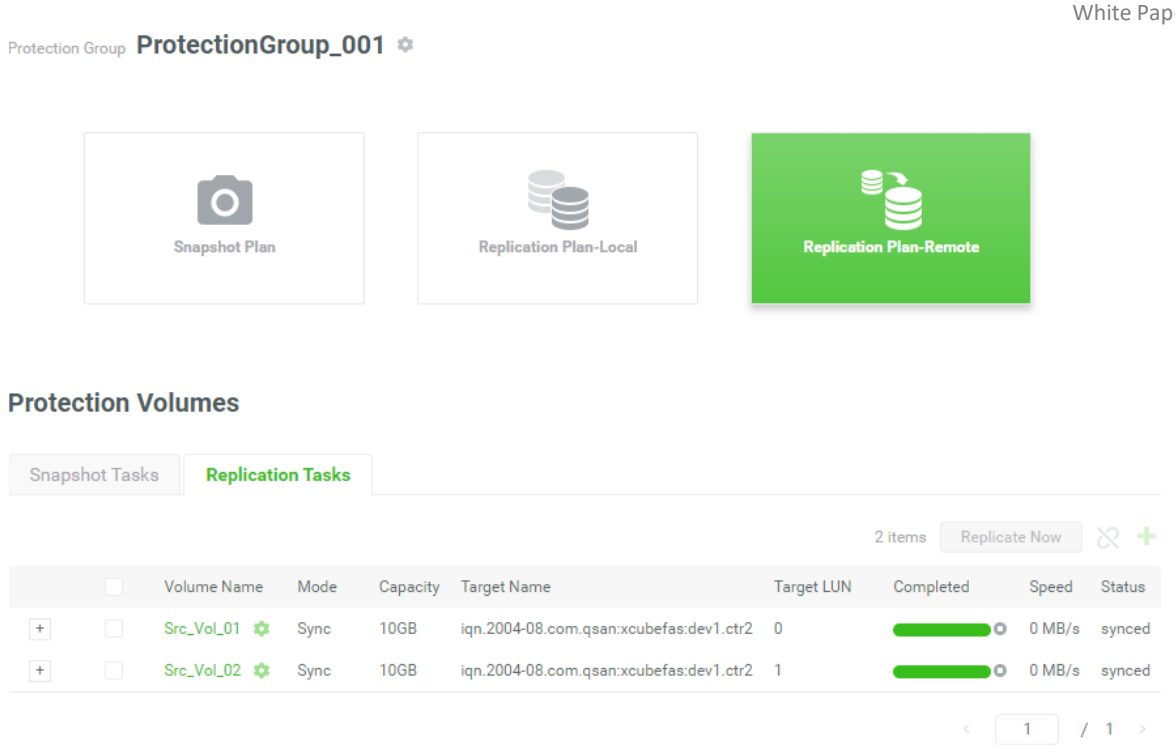

Figure 3-34 Synchronous Replication – Local to Remote Step 10

25. Done. A protection group with synchronous replication task has been reconnected.

3.4. Configure Asynchronous Replication

This section describes how to create asynchronous replication and its related functions. Just like configuring synchronous replication, we provide two methods to setup the tasks, auto and manual configuration.

3.4.1. Auto Configuration

Here is an example of creating a synchronous replication task with auto configuration.

Prepare in Site B (Target Unit)

The target unit has at least 3 free disks without any pool configuration, and the total capacity is greater than or equal to the source volumes. In addition, the data ports in the target unit are connected to the source unit and can be pinged.

TIP

The auto configuration can run in the same firmware version of the source and target units.

Configure Replication Plan-Remote in Site A (Source Unit)

1. Click the + icon in the **Protection Groups** pane to pop up a wizard.

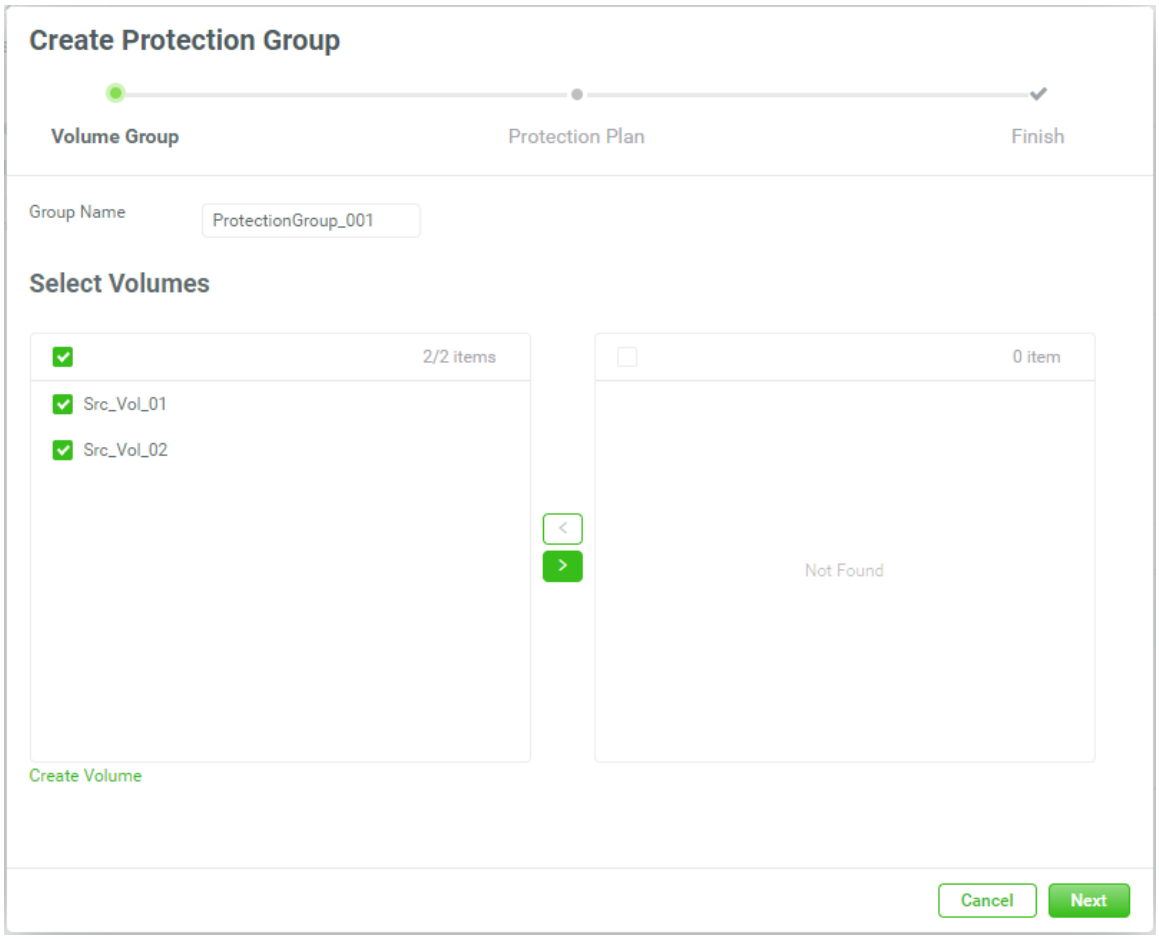

Figure 3-35 Asynchronous Replication – Auto Configuration Step 1

2. Select source volumes, and then click the button to move them from the left side to the right side.
3. Click the **Next** button to continue.

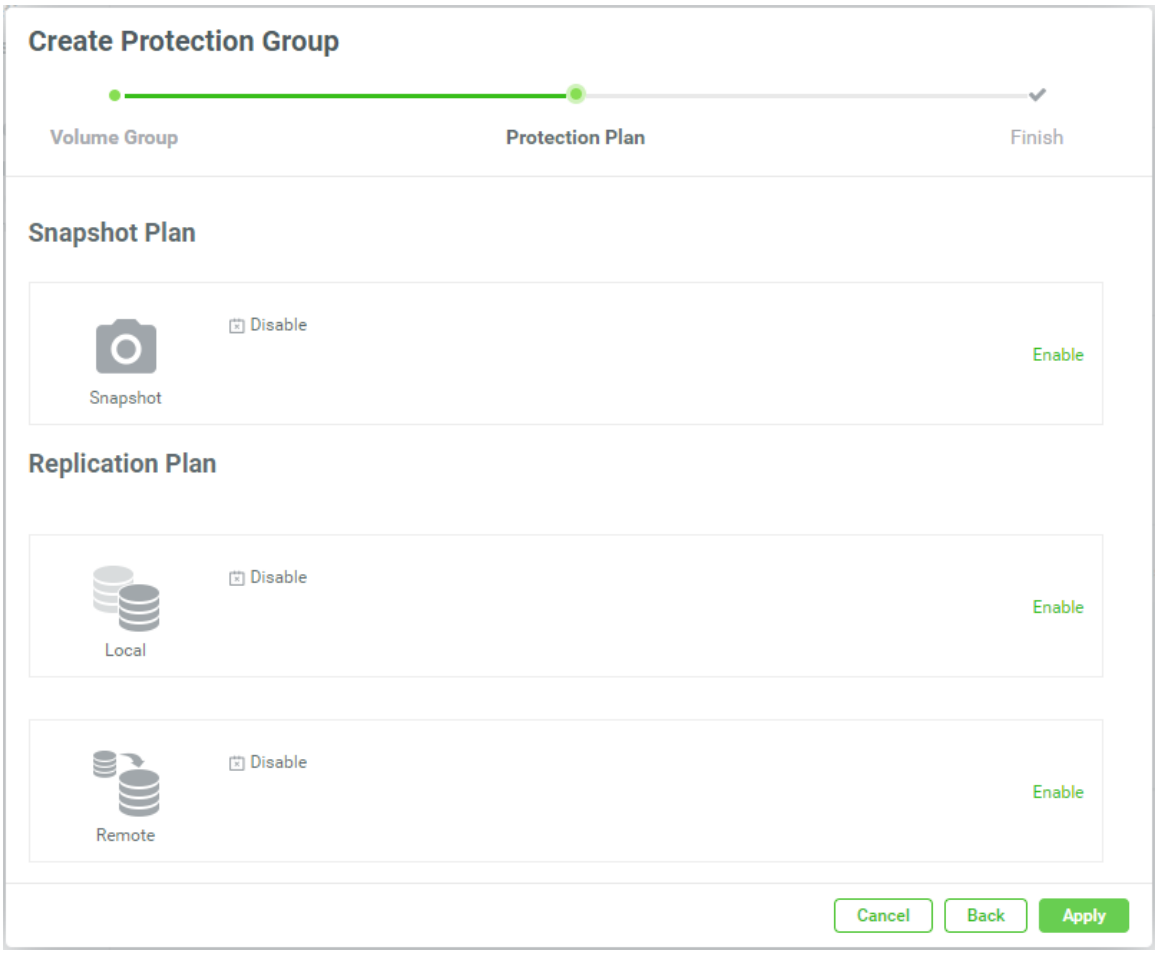

Figure 3-36 Asynchronous Replication – Auto Configuration Step 2

4. Click the **Enable** text link in the **Remote** pane to enable the **Replication Plan - Remote**.

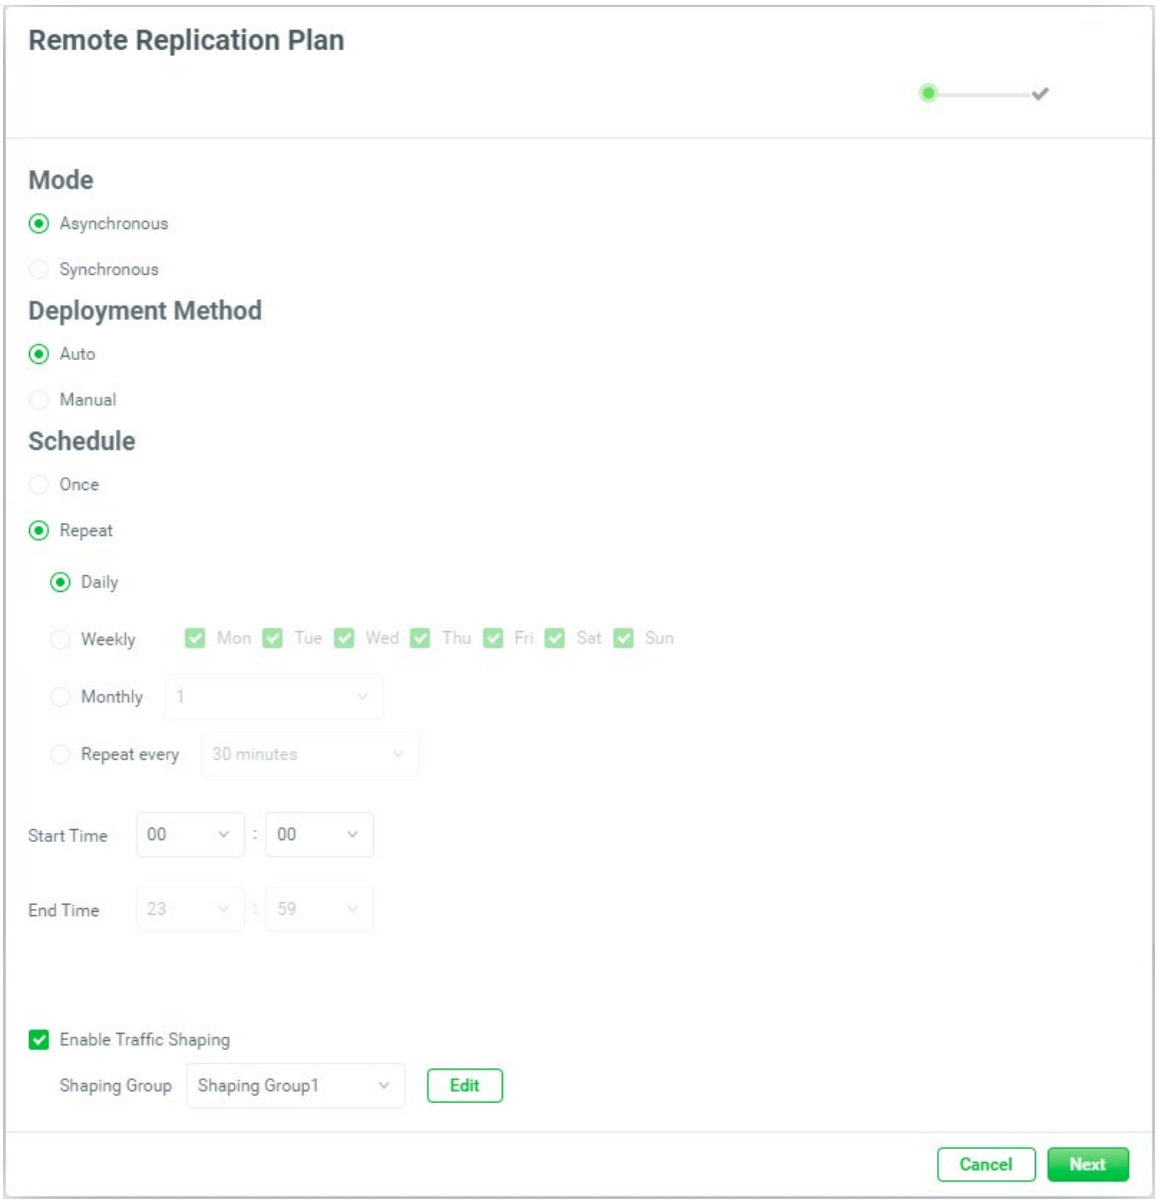

Figure 3-37 Asynchronous Replication – Auto Configuration Step 3

5. Select the **Mode** as **Asynchronous** and select the **Deployment Method** as **Auto**.
6. Select the **Once** option to execute the plan once. Or select the **Repeat** option to repeat the plan, and then check the repeat frequency.
7. If necessary, check the **Enable Traffic Shaping** and select the Shaping Group. Or click the **Edit** button to edit the shaping group.
8. Click the **Next** button to continue.

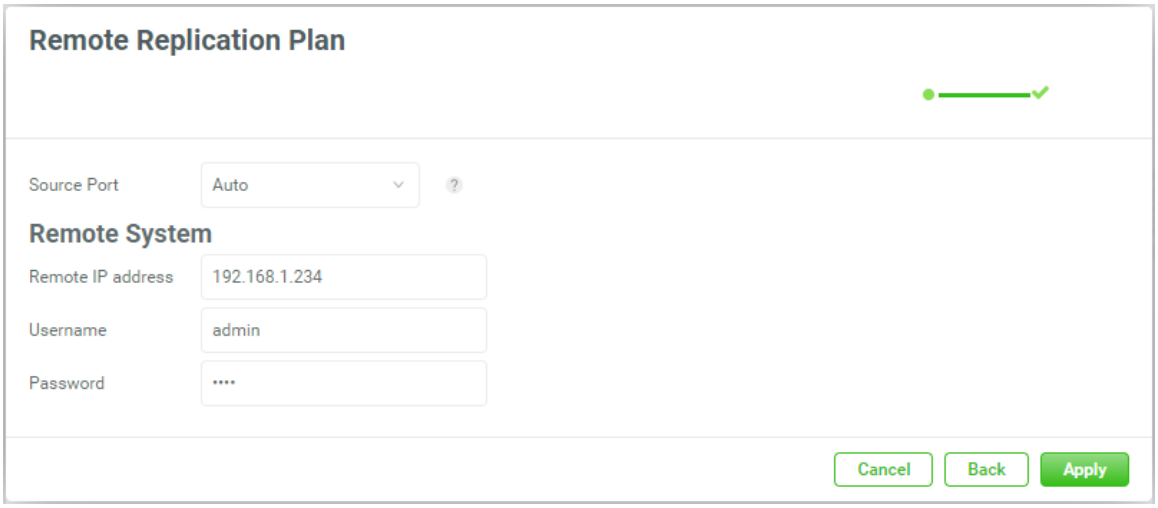

Figure 3-38 Asynchronous Replication – Auto Configuration Step 4

9. If necessary, select the **Source Port** or leave it as Auto.
10. Login to the management port of the target unit. Enter the **Remote IP address** of the target management port.
11. Enter the **Username** and **Password** of the target administrator account.
12. Click the **Apply** button to finish the **Replication Plan - Remote** configuration.

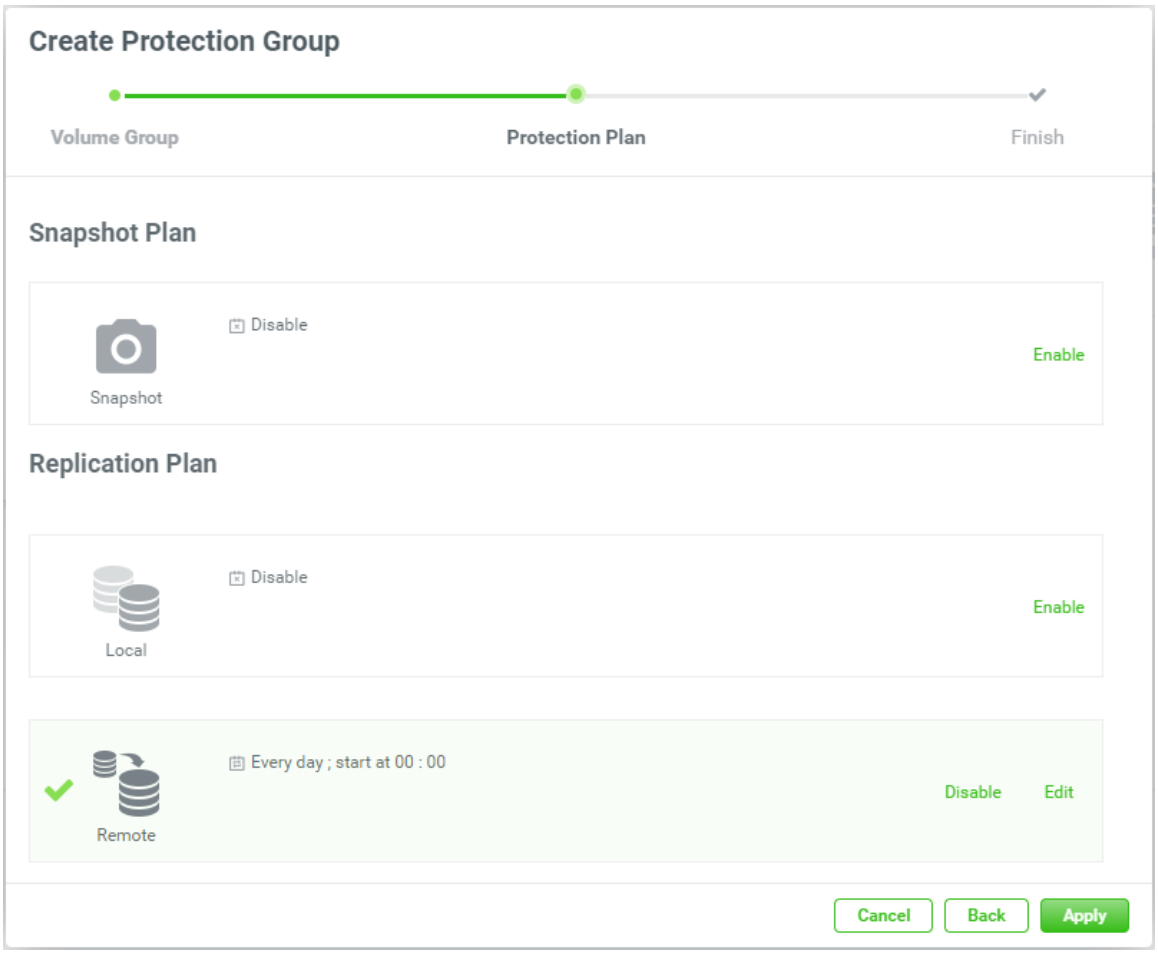

Figure 3-39 Asynchronous Replication – Auto Configuration Step 5

13. Back to main screen of the wizard. Click the **Apply** button to create a replication task.

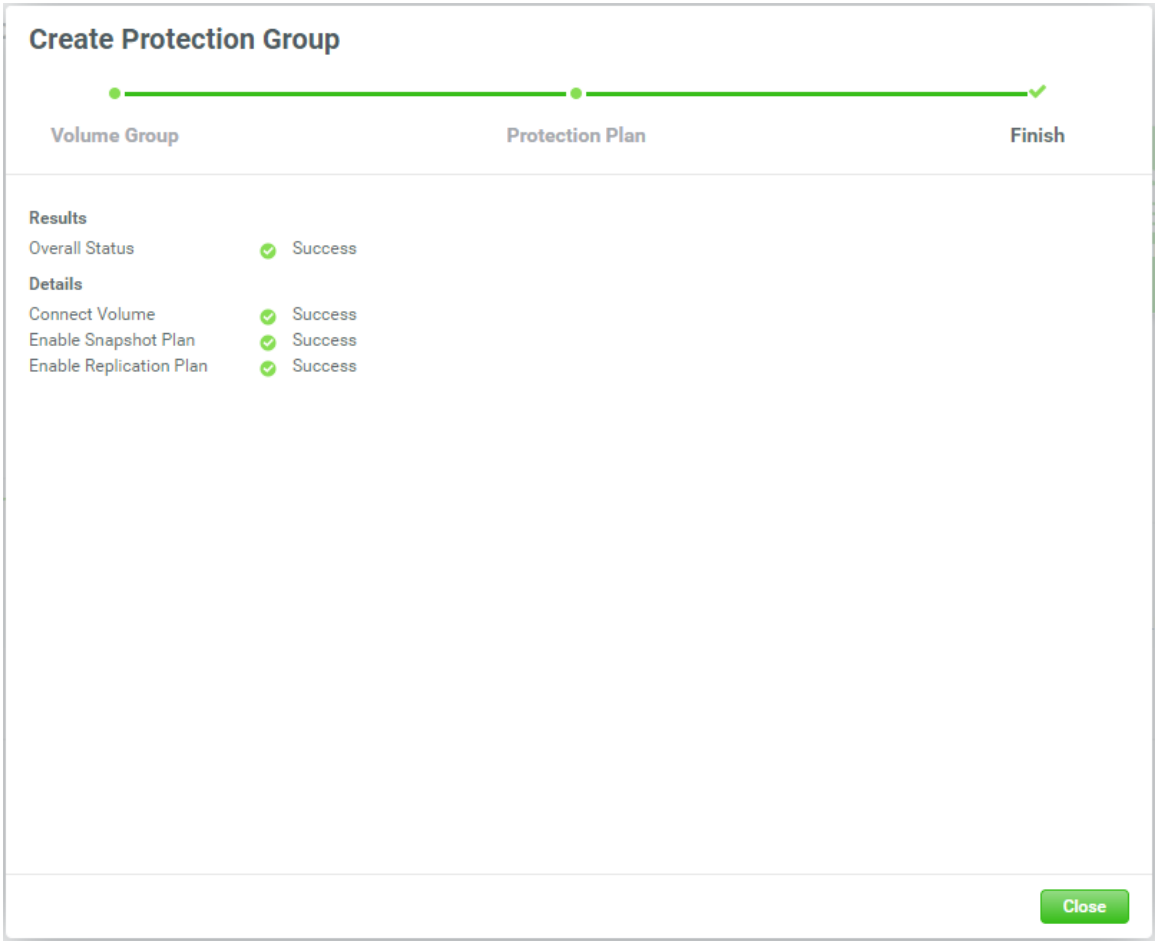

Figure 3-40 Asynchronous Replication – Auto Configuration Step 6

14. There is a result page. Click the **Close** button to close the wizard.

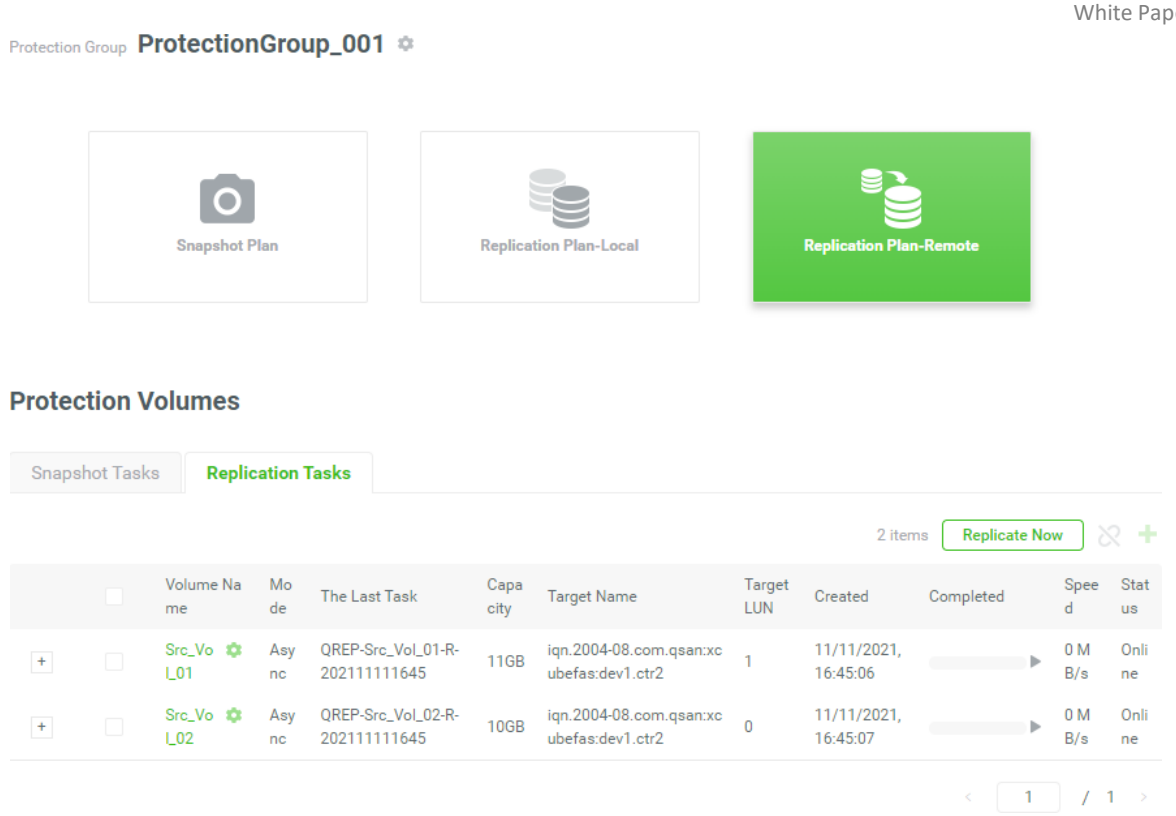

Figure 3-41 Asynchronous Replication – Auto Configuration Step 7

15. Done. A protection group with synchronous replication task has been created.

3.4.2. Manual Configuration

Here is an example of creating a synchronous replication task with manual configuration.

Prepare Backup Volumes in Site B (Target Unit)

1. Create a pool and multiple volumes with capacity greater than or equal to the source volumes.
2. Click the ⚙ icon beside each volume name to list the drop-down options, click the **Properties** option to change the volume properties.

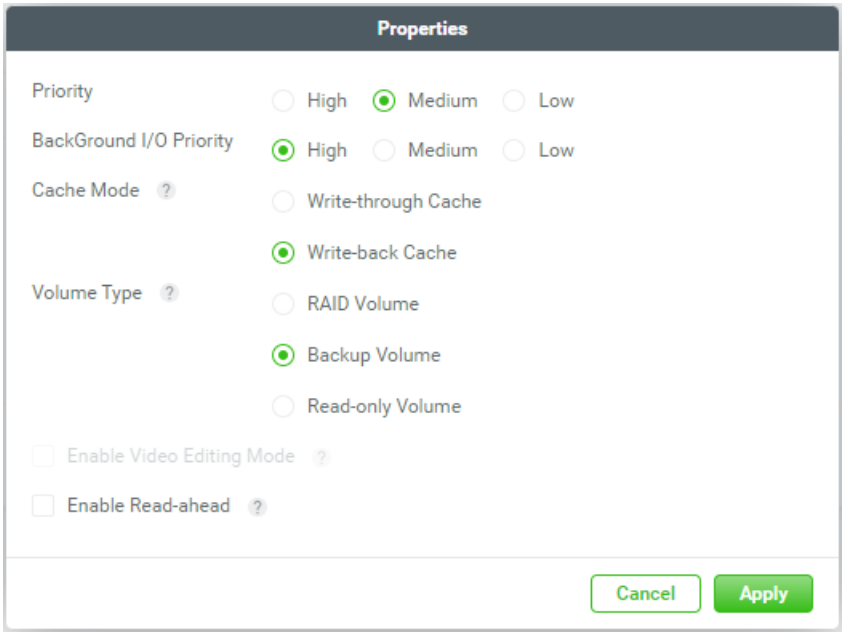

Figure 3-42 Asynchronous Replication – Manual Configuration Step 1

3. Change the **Volume Type** to **Backup Volume**, and then click the **Apply** button to take effect.
4. Click the icon beside each volume name to list the drop-down options, click the **Snapshot Center** option to enable snapshot space.

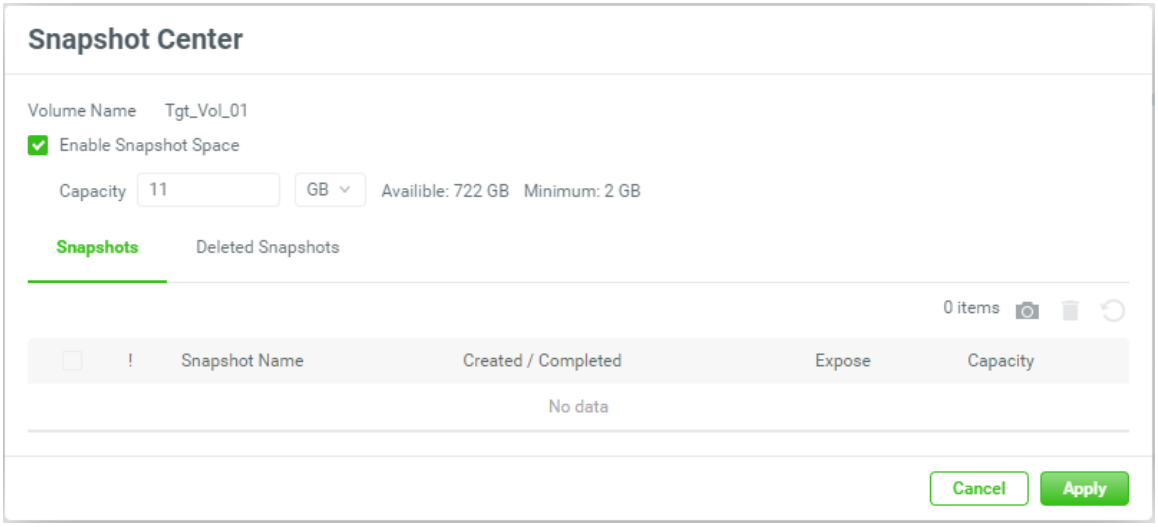

Figure 3-43 Asynchronous Replication – Manual Configuration Step 2

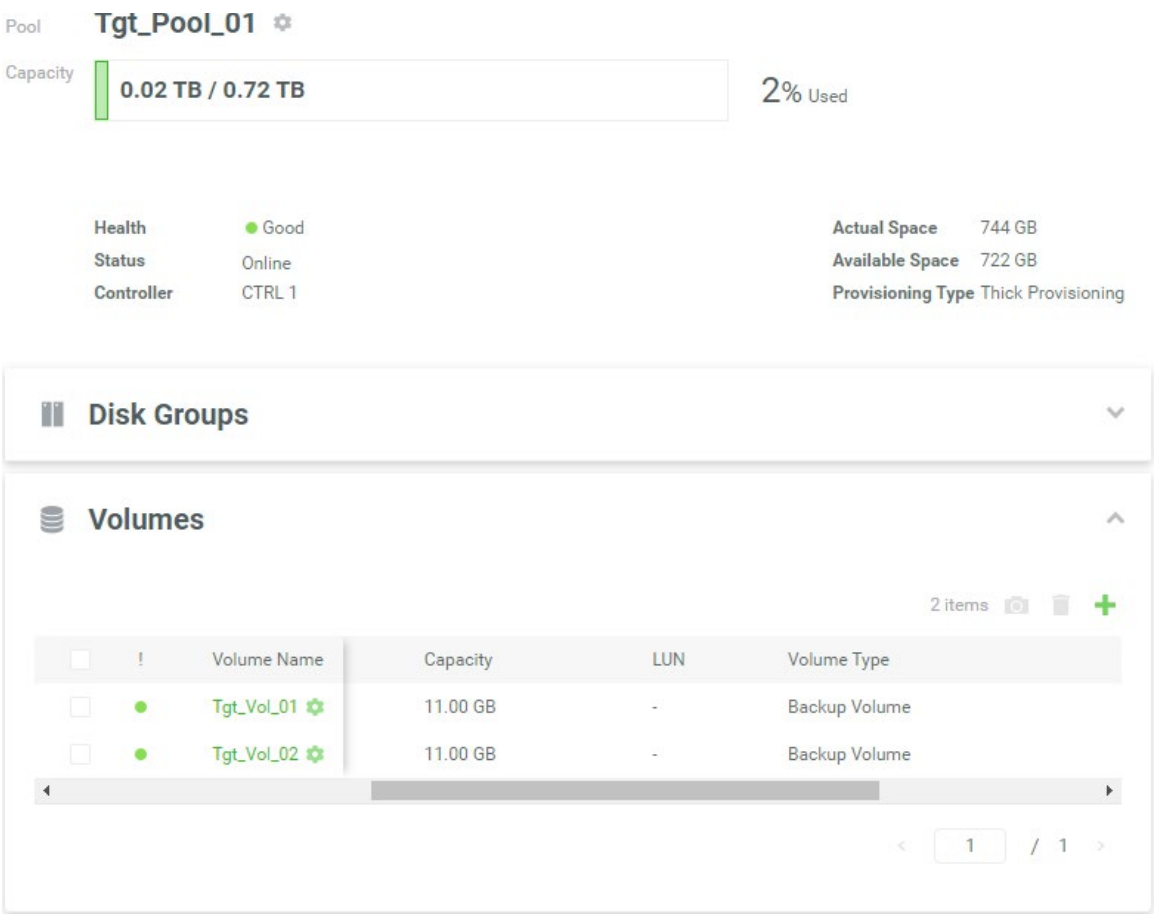

Figure 3-44 Asynchronous Replication – Manual Configuration Step 3

5. Here are the target volumes.
6. Create a host group with the target volumes in the **Hosts** tab.

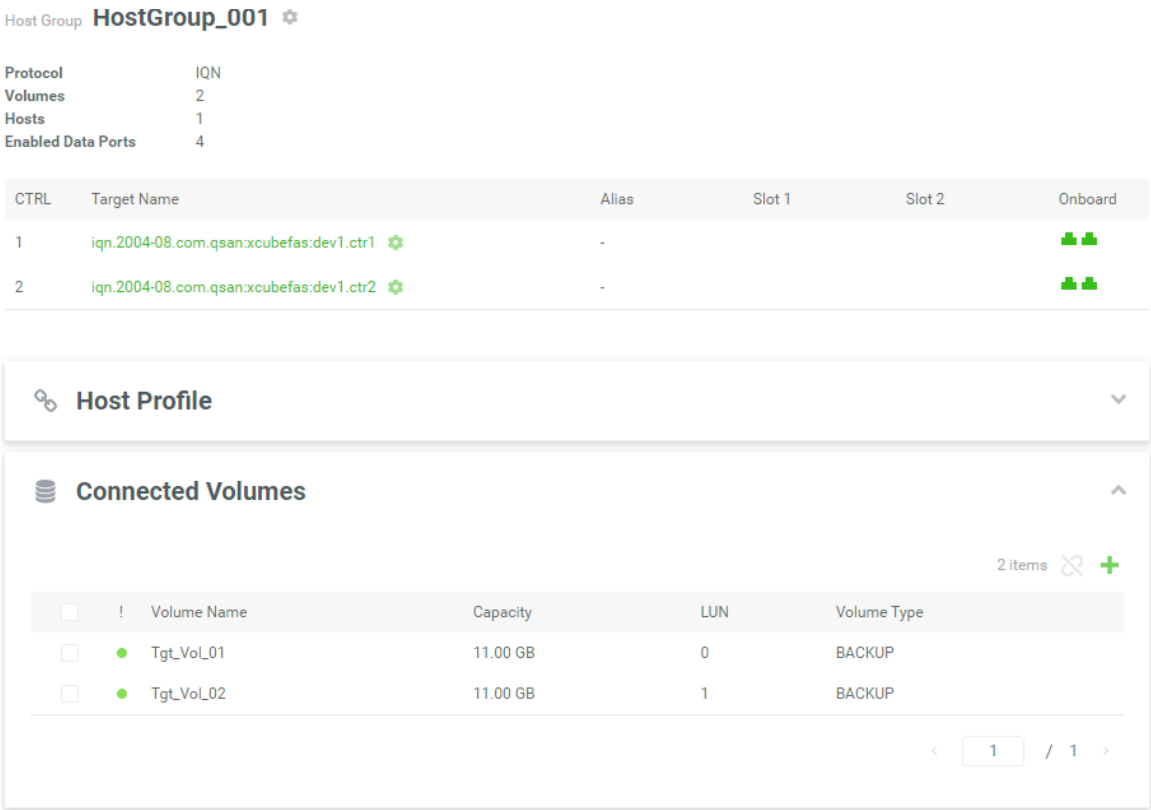

Figure 3-45 Asynchronous Replication – Manual Configuration Step 4

7. Here is the host group for target volumes.

Configure Replication Plan-Remote in Site A (Source Unit)

8. Click the + icon in the **Protection Groups** pane to pop up a wizard.

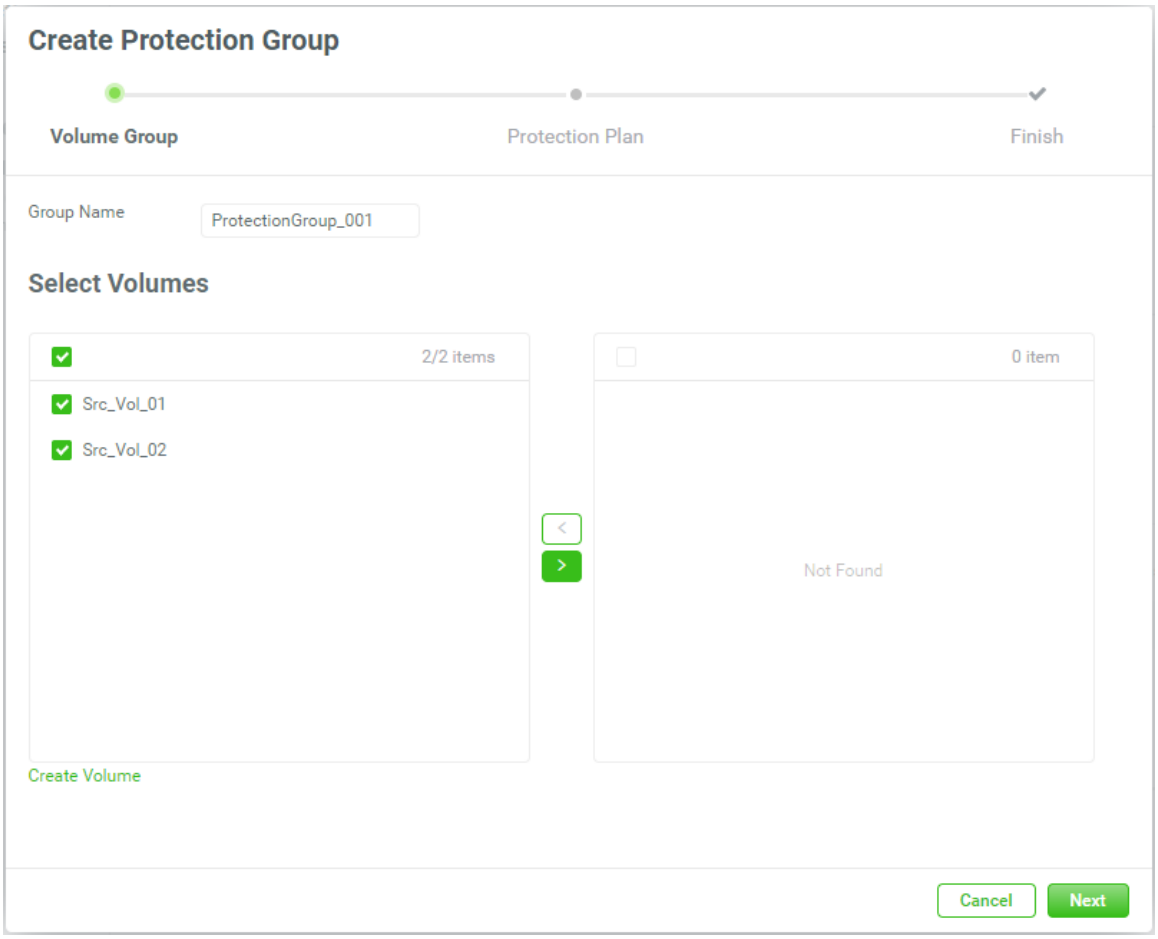

Figure 3-46 Asynchronous Replication – Manual Configuration Step 5

9. Select source volumes, and then click the button to move them from the left side to the right side.
10. Click the **Next** button to continue.

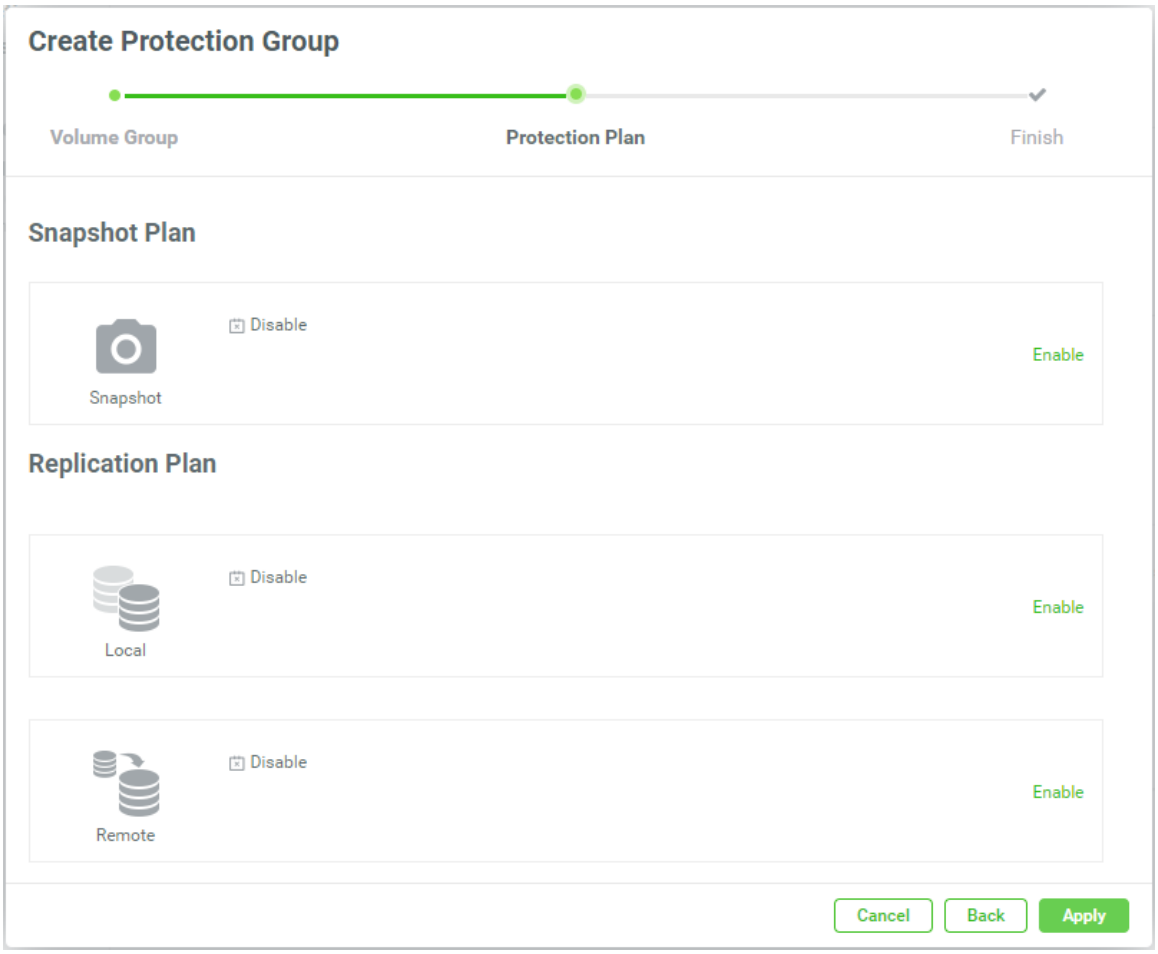

Figure 3-47 Asynchronous Replication – Manual Configuration Step 6

11. Click the **Enable** text link in the **Remote** pane to enable the **Replication Plan - Remote**.

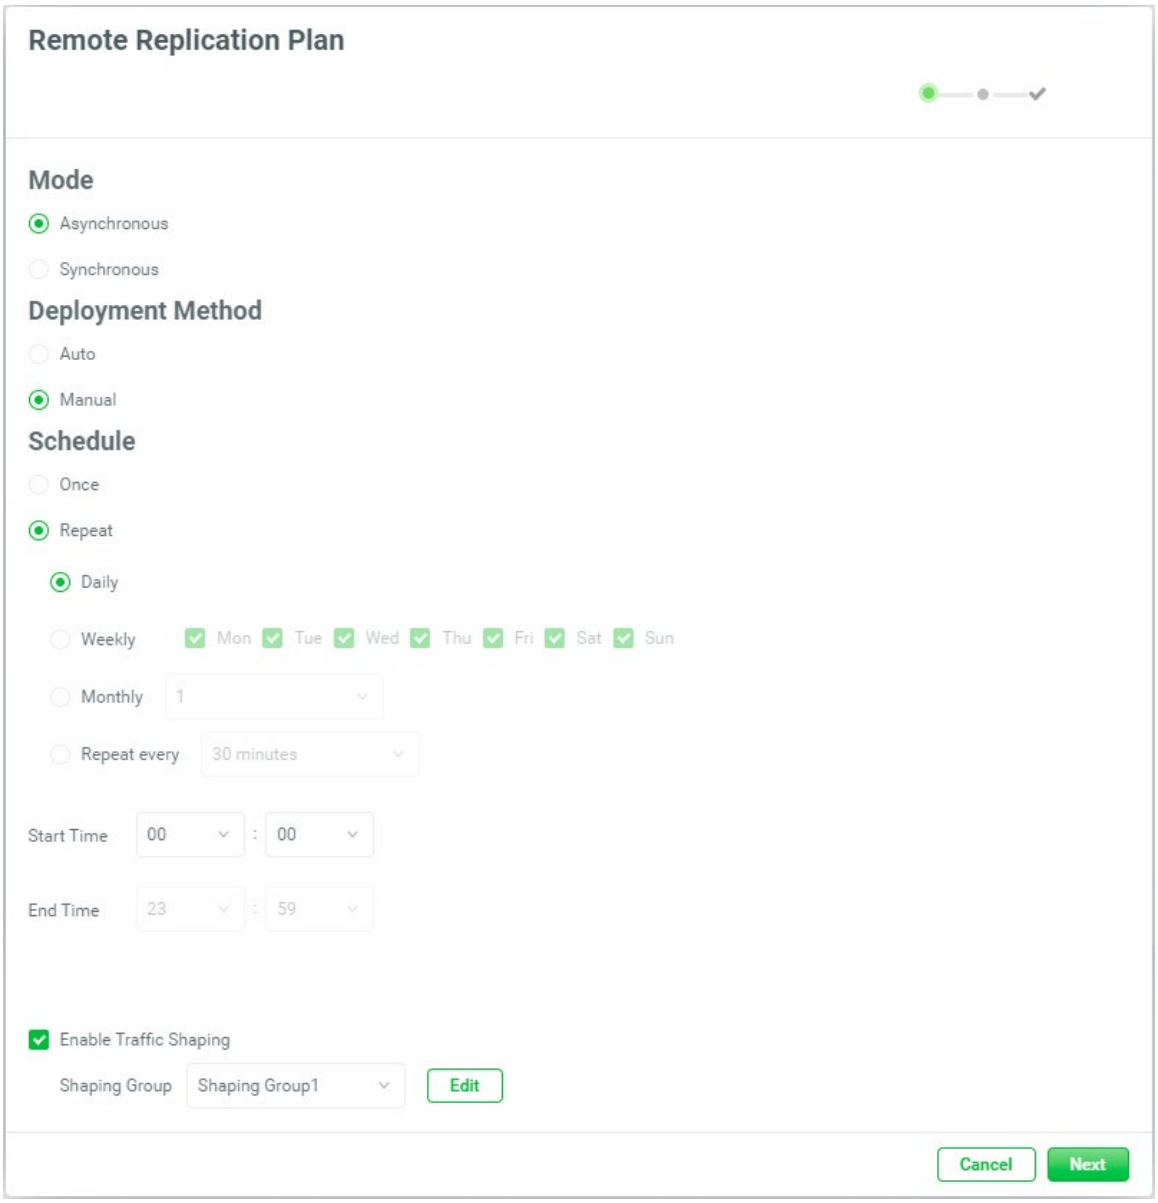

Figure 3-48 Asynchronous Replication – Manual Configuration Step 7

12. Select the **Mode** as **Asynchronous** and select the **Deployment Method** as **Manual**.
13. Select the **Once** option to execute the plan once. Or select the **Repeat** option to repeat the plan, and then check the repeat frequency.
14. If necessary, check the **Enable Traffic Shaping** and select the Shaping Group.
15. Click the **Next** button to continue.

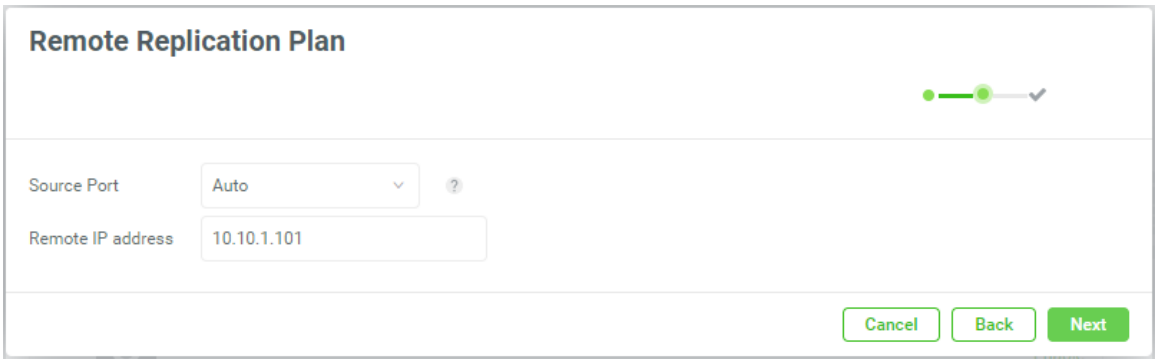

Figure 3-49 Asynchronous Replication – Manual Configuration Step 8

16. If necessary, select the **Source Port** or leave it as Auto.
17. Enter the **Remote IP address** of the target data port.
18. Click the **Next** button to continue.

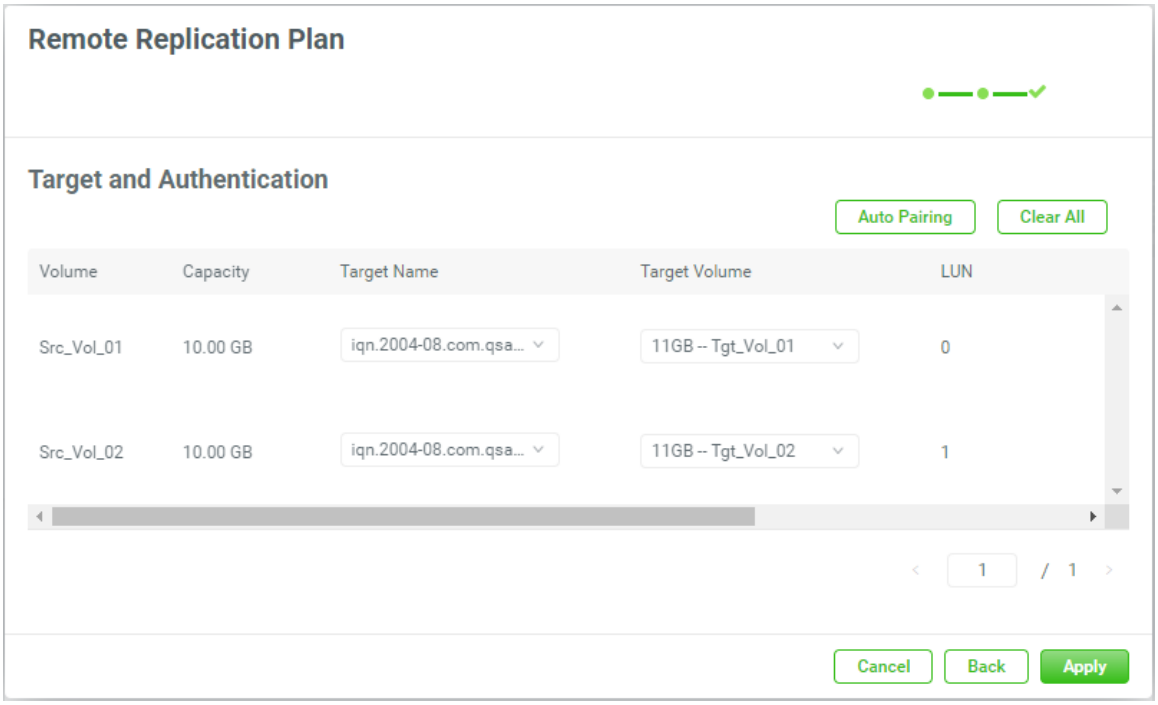

Figure 3-50 Asynchronous Replication – Manual Configuration Step 9

19. Click the **Auto Pairing** button to automatically pair the backup volumes. If they are not what you want, manually select the **Target Name** and **Target Volume**.
20. If the target encrypts with CHAP, click the icon next to the Target Name, it will pop up a dialog box to enter the CHAP information.

21. Click the **Apply** button to finish the **Replication Plan - Remote** configuration.

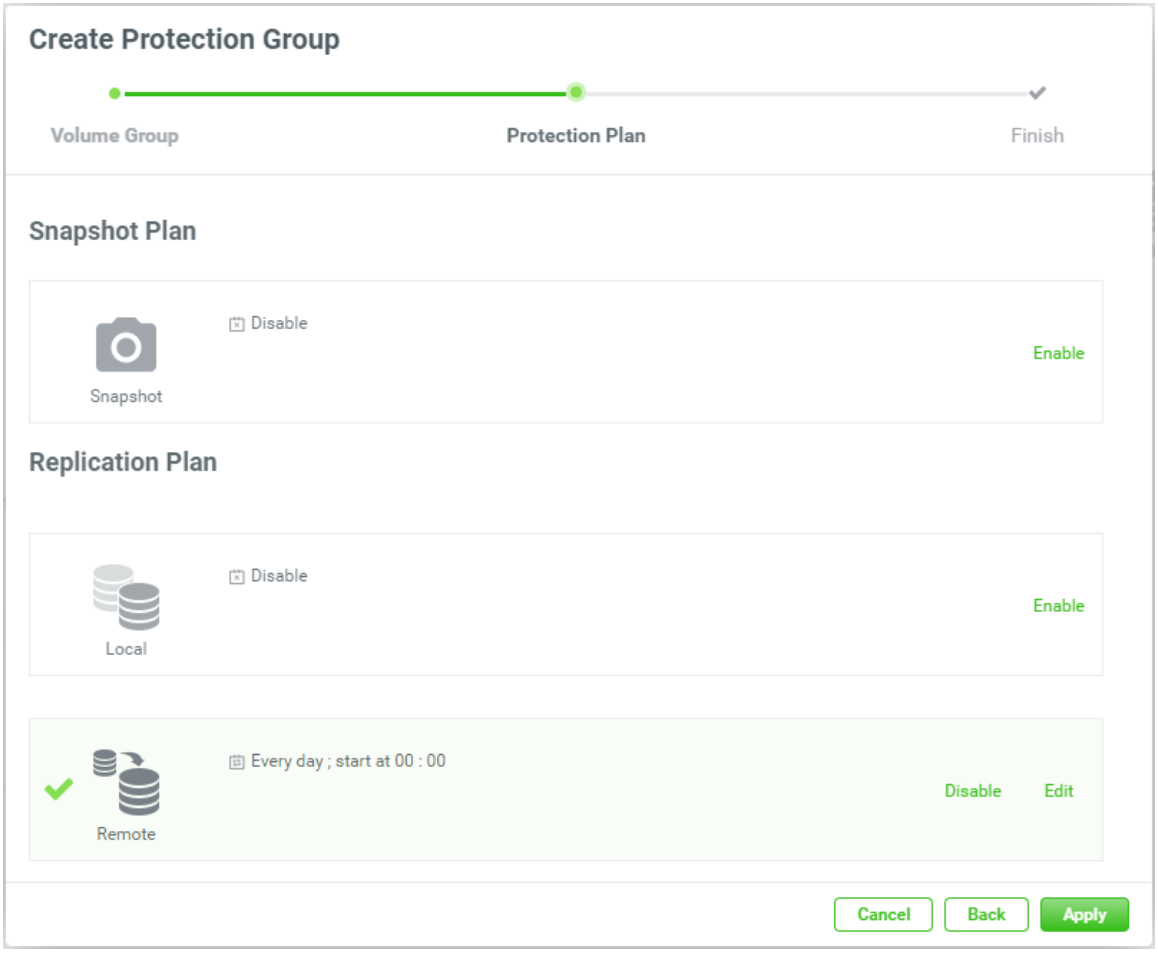

Figure 3-51 Asynchronous Replication – Manual Configuration Step 10

22. Back to main screen of the wizard. Click the **Apply** button to create a replication task.

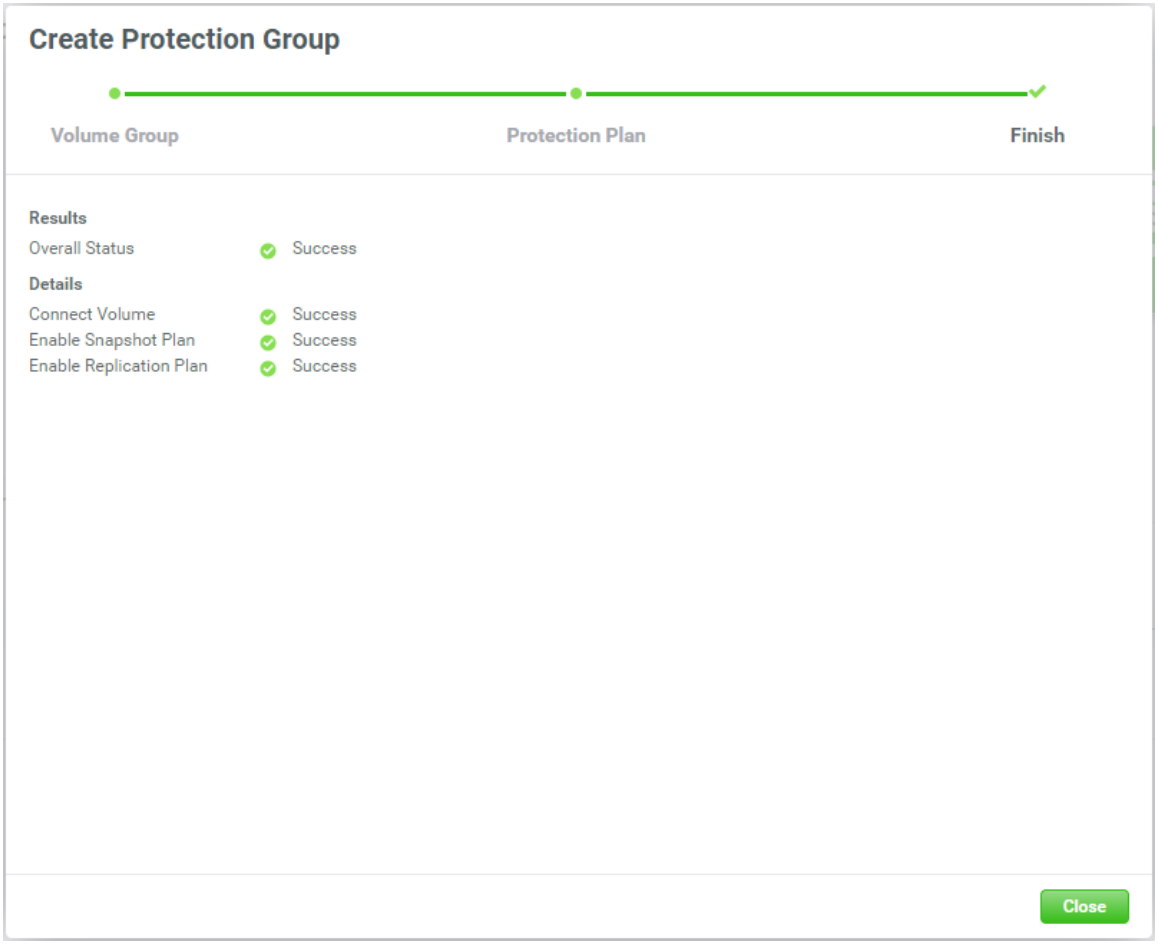

Figure 3-52 Asynchronous Replication – Manual Configuration Step 11

23. There is a result page. Click the **Close** button to close the wizard.

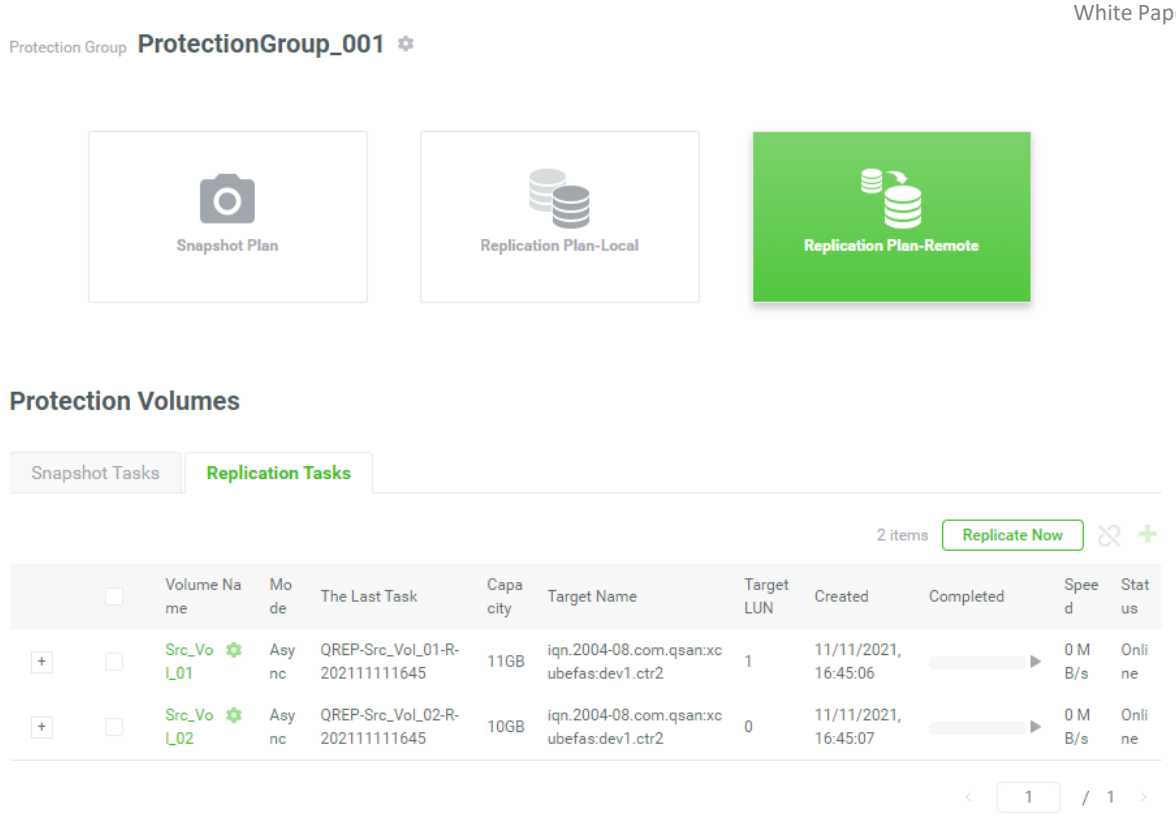

Figure 3-53 Asynchronous Replication – Manual Configuration Step 12

24. Done. A protection group with synchronous replication task has been created.

3.4.3. Replication Options

These options are available in the protection group. Click the ⚙ icon beside the protection group name to list the drop-down options.

Traffic Shaping

1. Click the **Traffic Shaping** option to edit the traffic shaping configurations.

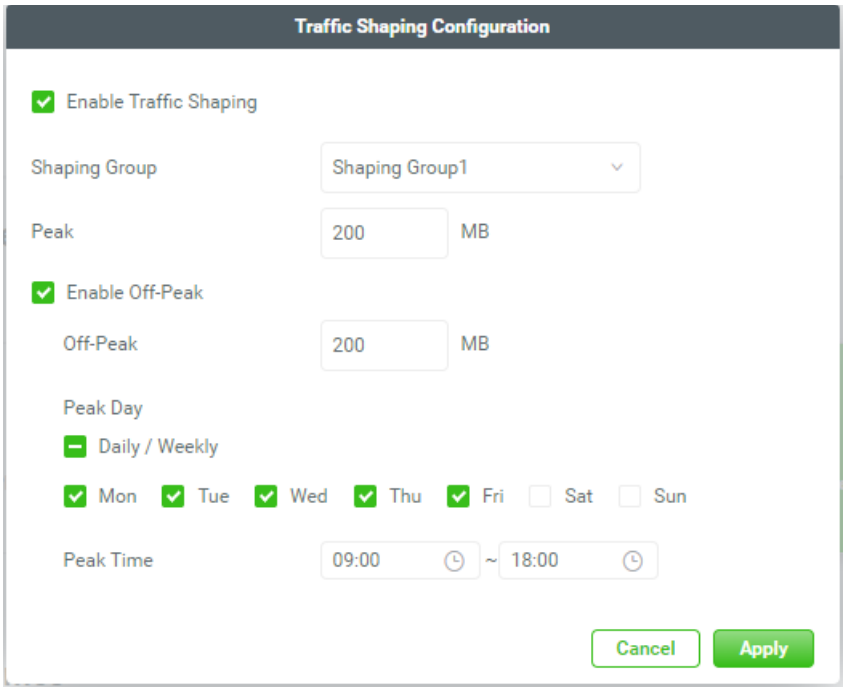

Figure 3-54 Asynchronous Replication – Traffic Shaping Configurations

2. Select a **Shaping Group** to be modified, enter a maximum throughput during **Peak** hours. Check the **Enable Traffic Shaping** option if necessary, and then enter a maximum throughput during **Off-Peak** hours and define the **Peak Day**.
3. Click the **Apply** button to continue.

The options are available in this tab.

Replicate Now

Click the **Replicate Now** button to replicate all volumes immediately. This option only applies to asynchronous replication tasks.

Start Replication Task

Select a volume, and then click the ▶ icon to replicate the volume immediately. The progress bar displays the current status.

Stop Replication Task

Click the icon of the volume to stop the replication task.

TIP

The **Replicate Now** button will replicate all volumes in the protection group. And click the ▶ icon of the volume will replicate the volume only.

Ungroup Volume from Protection Group

To ungroup the volume from the protection group, check one or multiple volume names, and then click the button.

Delete Protection Group

To delete the protection group, click the icon beside the protection group name to list the drop-down options. Click the **Delete** option; it will pop up dialog box to confirm. Click the **Delete** button to delete the protection group.

3.4.4. Enable Multipath and Add Connections

Click the icon next to the volume name to list the drop-down options. Click the **Connection Properties** option to enable multipath and add connections for the replication task.

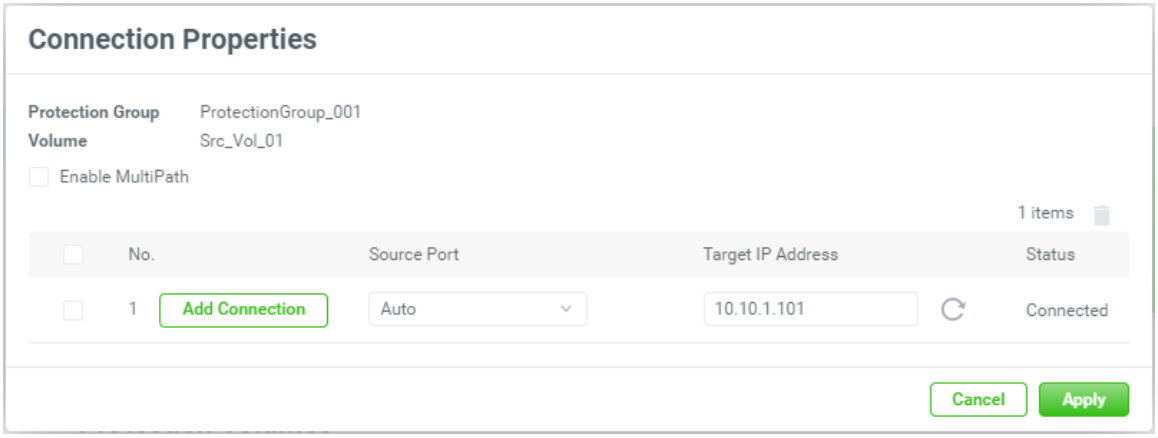

Figure 3-55 Asynchronous Replication – Connection Properties

Enable Multipath

1. Check the **Enable MultiPath** checkbox to enable multipath for the replication task.

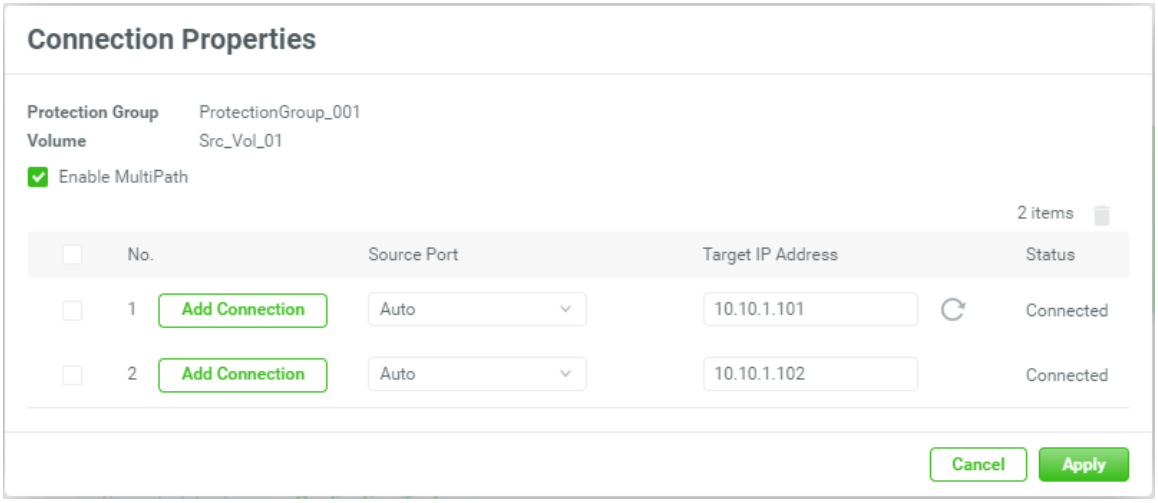

Figure 3-56 Asynchronous Replication – Enable Multipath

2. If necessary, select the **Source Port** or leave it as Auto.
3. Enter the **Target IP address** of the target data port.
4. Click the **Apply** button to enable.

Add Connections

1. Click the **Add Connection** button to add another connection.

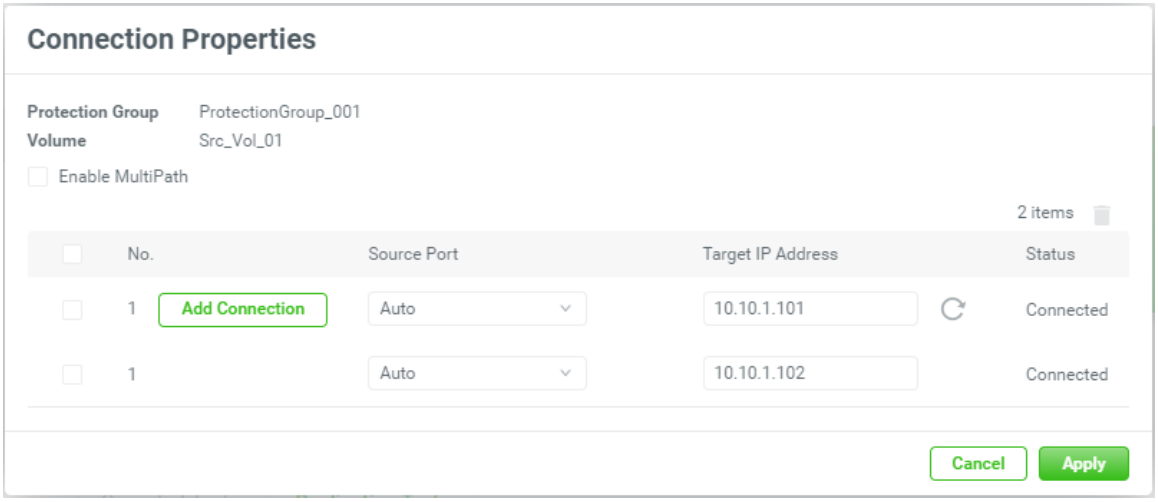

Figure 3-57 Asynchronous Replication – Add Connection

2. If necessary, select the **Source Port** or leave it as Auto.
3. Enter the **Target IP address** of the target data port.
4. Click the **Apply** button to add.

3.4.5. Switch Local Clone to Remote Replication

Through local-to-remote, you can transfer local backup to remote sites without redoing full copy. You can plug and play when local-to-remote without any configurations.

Configure Replication Plan-Local in Site A (Source Unit)

1. Click the ✚ icon in the **Protection Groups** pane to pop up a wizard.

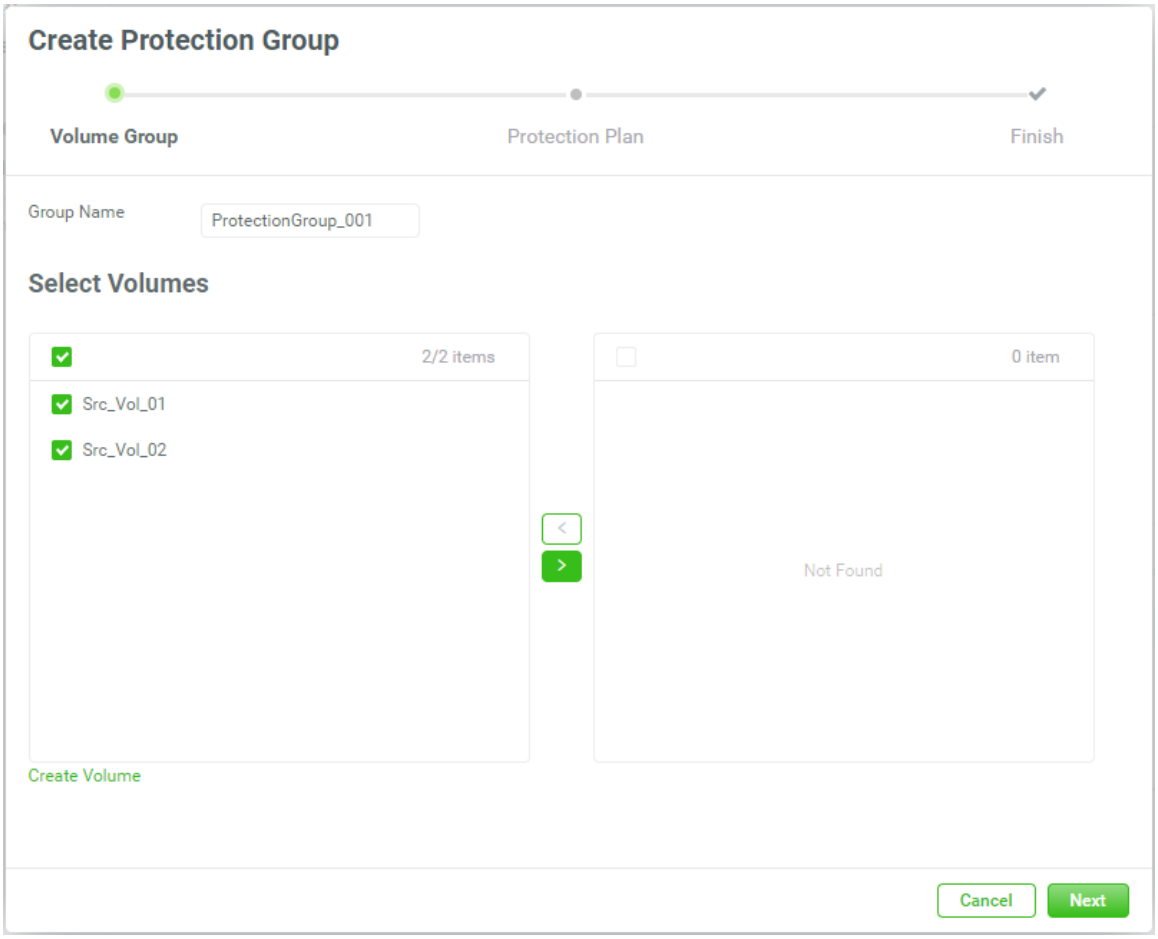

Figure 3-58 Asynchronous Replication – Local to Remote Step 1

2. Select source volumes, and then click the button to move them from the left side to the right side.
3. Click the **Next** button to continue.

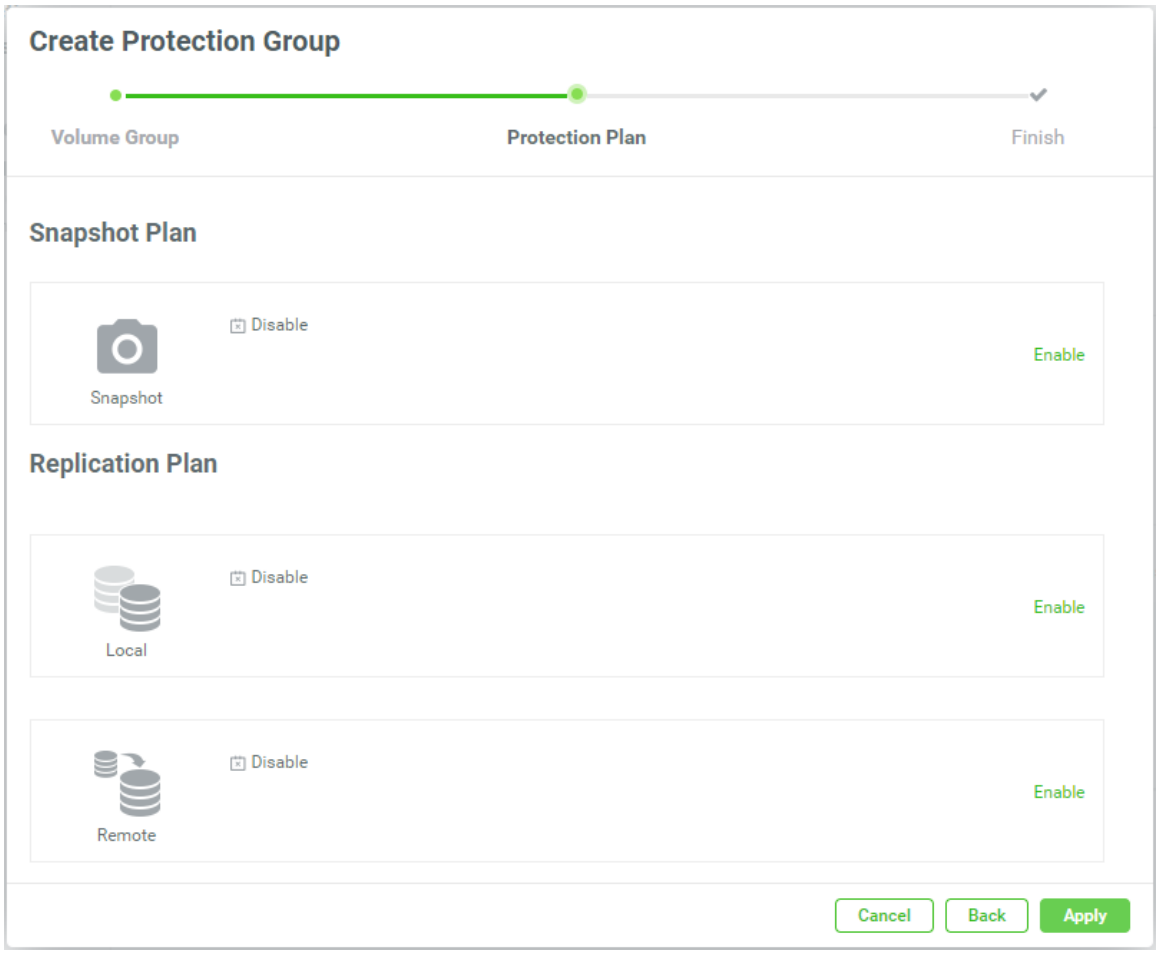

Figure 3-59 Asynchronous Replication – Local to Remote Step 2

4. Click the **Enable** text link in the **Local** pane to enable the **Replication Plan - Local**.

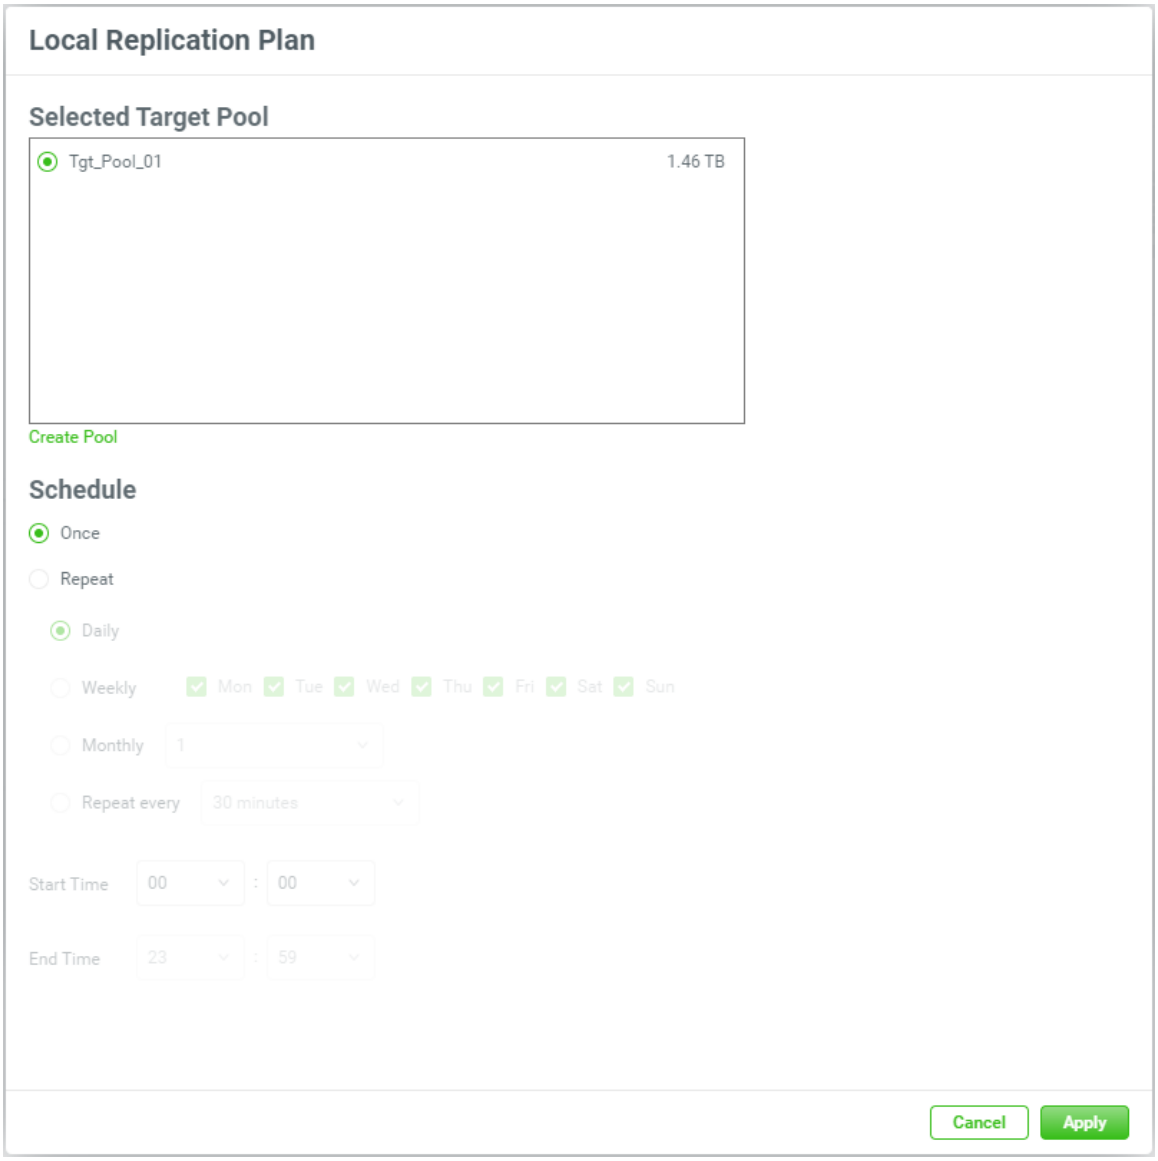

Figure 3-60 Asynchronous Replication – Local to Remote Step 3

5. Select a target pool. If none can be selected, click the **Create Pool** link to create a pool.
6. Click the **Apply** button to create a local clone task. And then click the **Close** button to close the wizard.

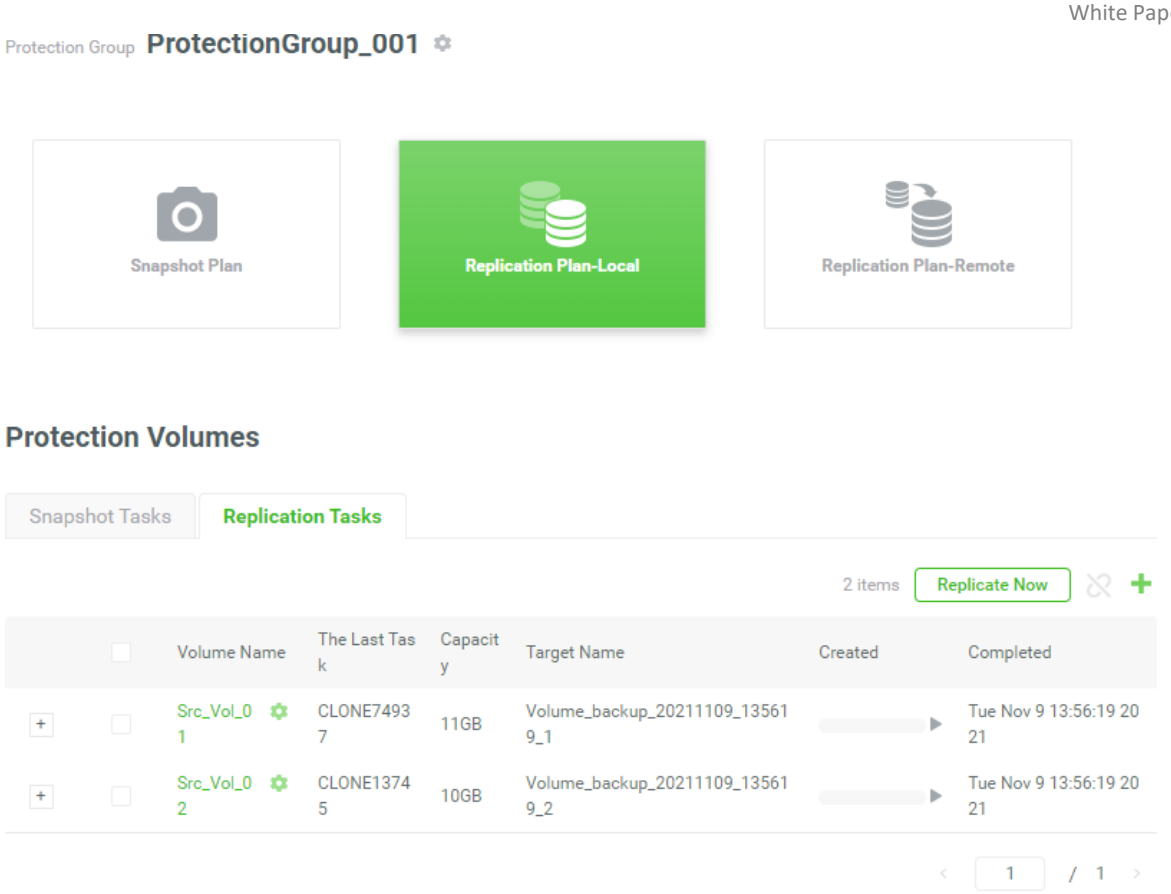

Figure 3-61 Asynchronous Replication – Local to Remote Step 4

7. After the local clone is complete, click the icon next to the protection group name to list the drop-down options. And then click the **Protection Plan Setting** option.

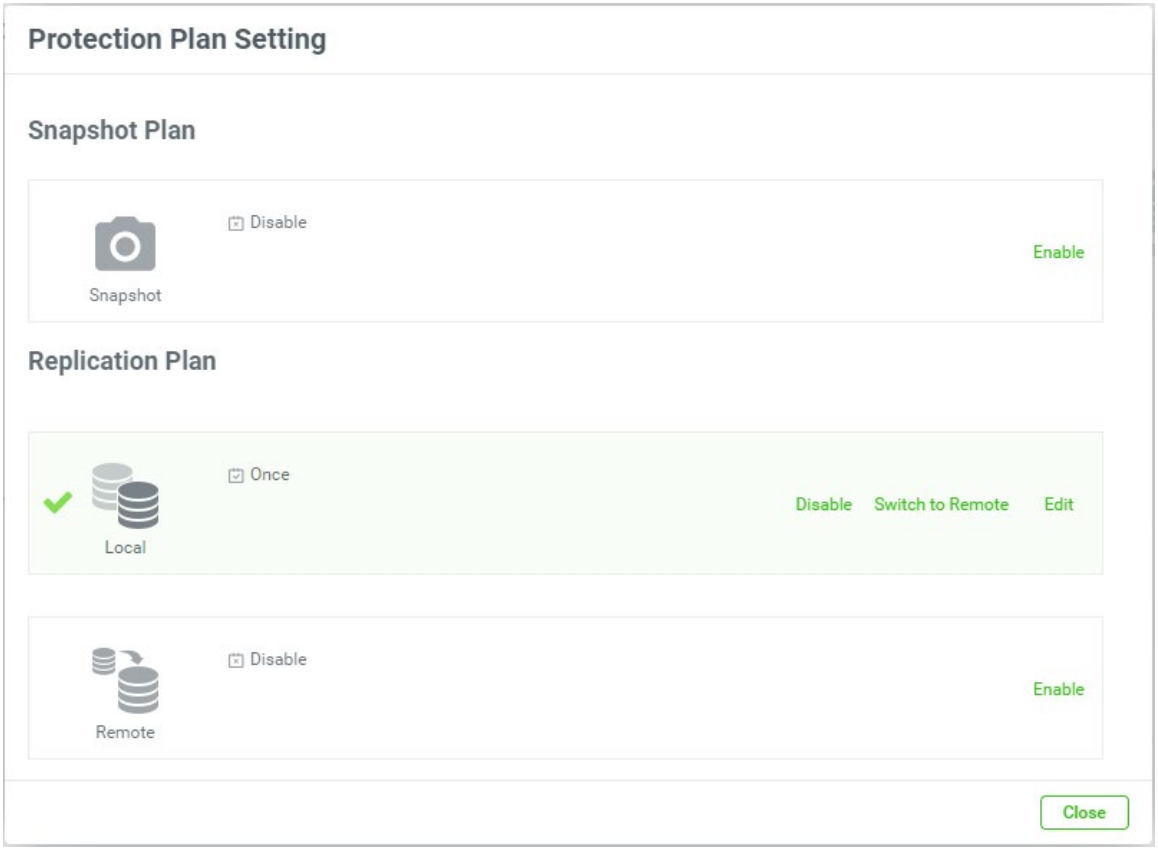

Figure 3-62 Asynchronous Replication – Local to Remote Step 5

8. Click the **Switch to Remote** text to configure the replication plan.

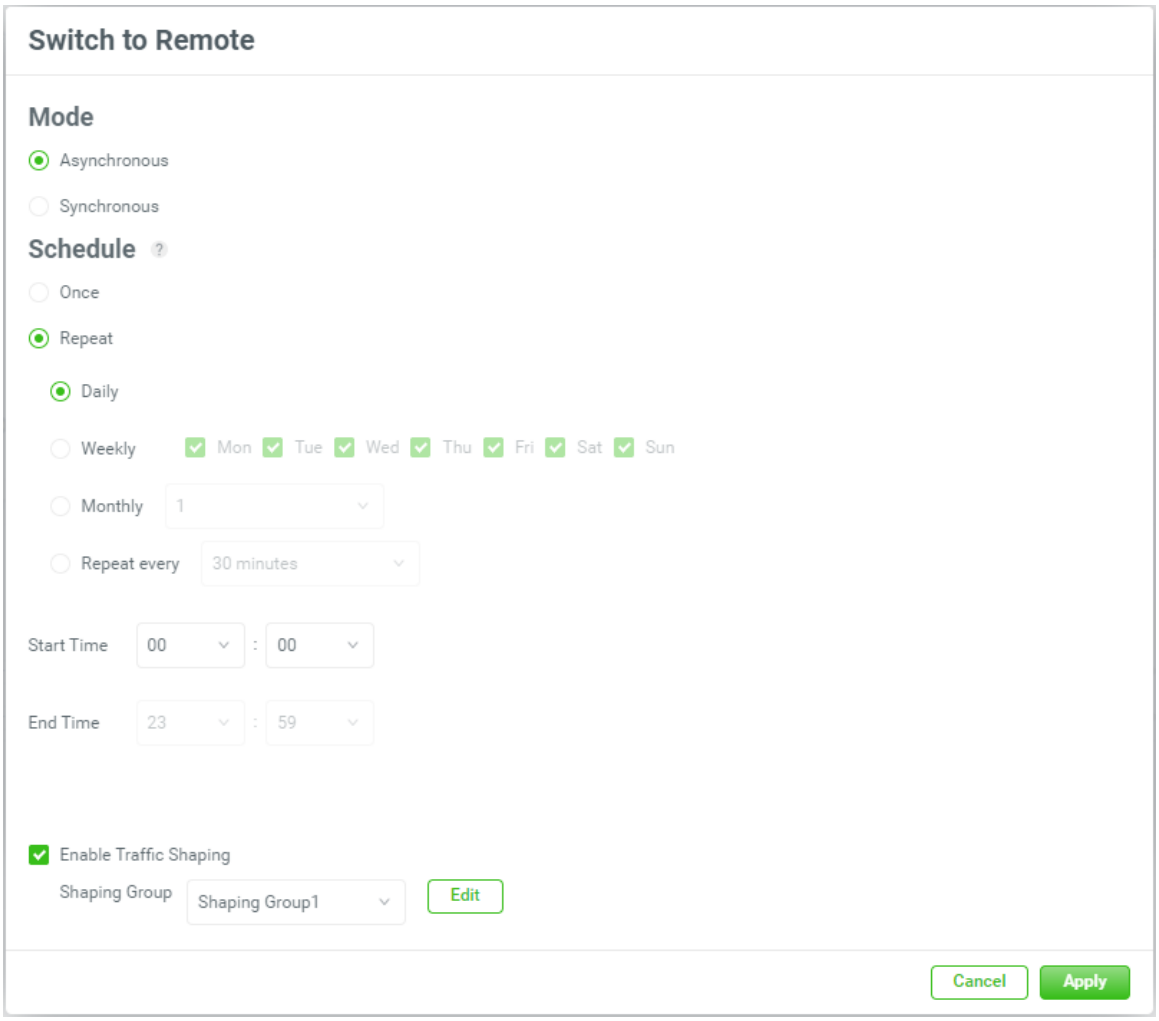

Figure 3-63 Asynchronous Replication – Local to Remote Step 6

9. Select the **Mode** as **Asynchronous**.
10. Select the **Once** option to execute the plan once. Or select the **Repeat** option to repeat the plan, and then check the repeat frequency.
11. If necessary, check the **Enable Traffic Shaping** and select the Shaping Group. Or click the **Edit** button to edit the shaping group.
12. Click the **Apply** button and then click the **Close** button to close the wizard.
13. Deactivate the target pool.
14. Remove all disk drives of the target pool and deliver them to the Site B (Target Unit).

Prepare Backup Volumes in Site B (Target Unit)

15. Insert all disk drives of the target pool and activate them.
16. Create a host group to include the target volumes.

Configure Replication Plan-Remote in Site A (Source Unit)

17. Select the source volume of the protection group.

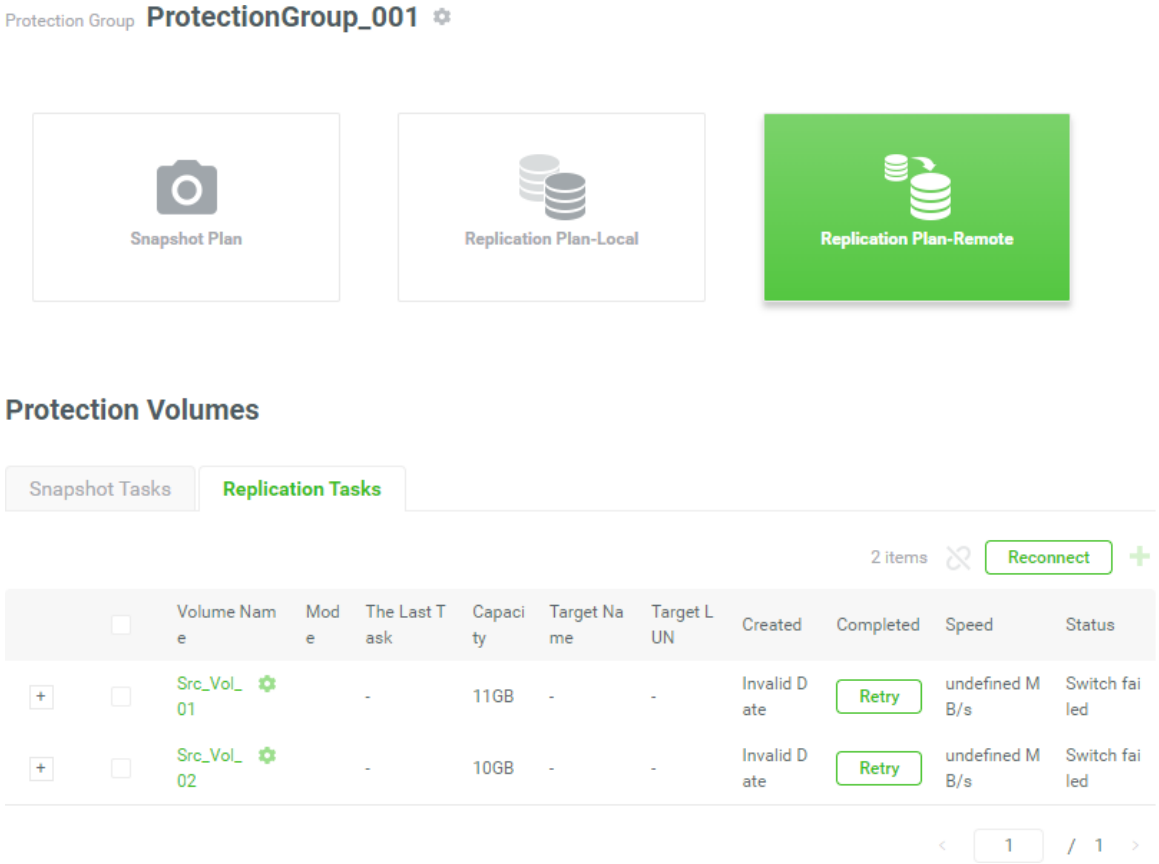

Figure 3-64 Asynchronous Replication – Local to Remote Step 7

18. Click the **Reconnect** button to reconnect to the target volume. It will pop up a wizard.

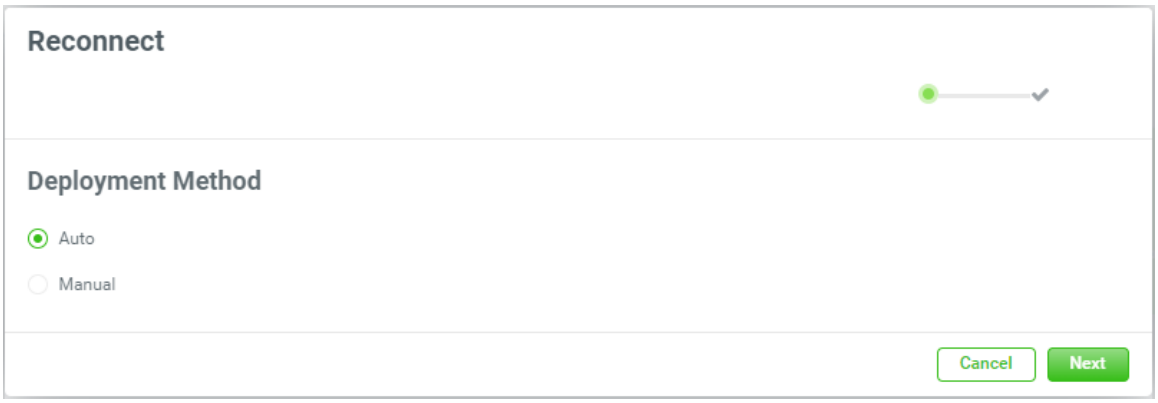

Figure 3-65 Asynchronous Replication – Local to Remote Step 8

19. Select the **Deployment Method** as **Auto** or **Manual**.
20. Click the **Next** button to continue.
21. The following takes **Auto** as an example, if you choose **Manual**, please refer to steps 16 to 21 in the 3.4.2 Manual Configuration section.

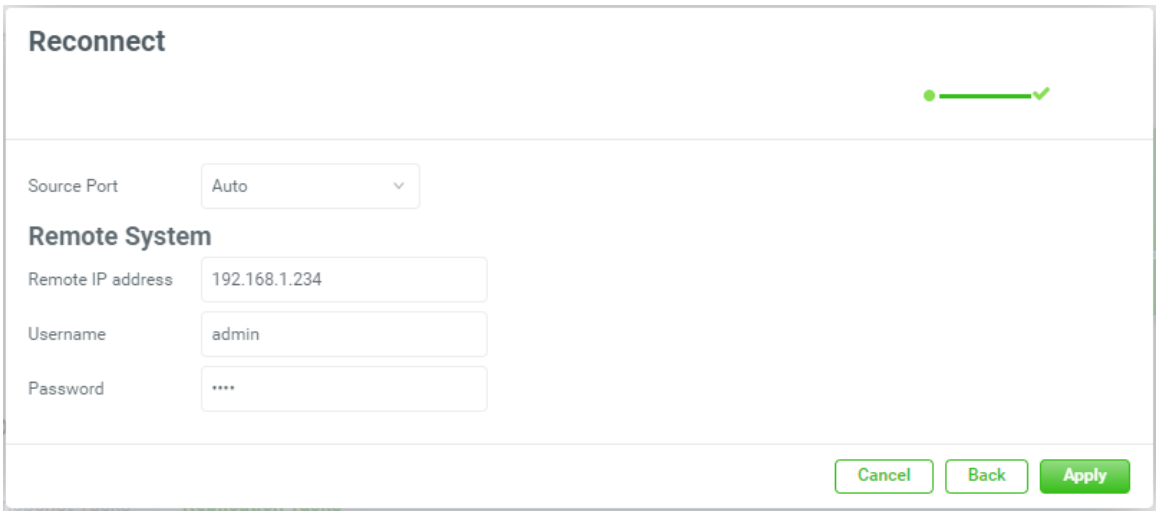

Figure 3-66 Asynchronous Replication – Local to Remote Step 9

22. If necessary, select the **Source Port** or leave it as Auto.
23. Login to the management port of the target unit. Enter the **Remote IP address** of the target management port.
24. Enter the **Username** and **Password** of the target administrator account.
25. Click the **Apply** button to finish the **Reconnect** configuration.
26. Click the **Close** button to close the wizard.

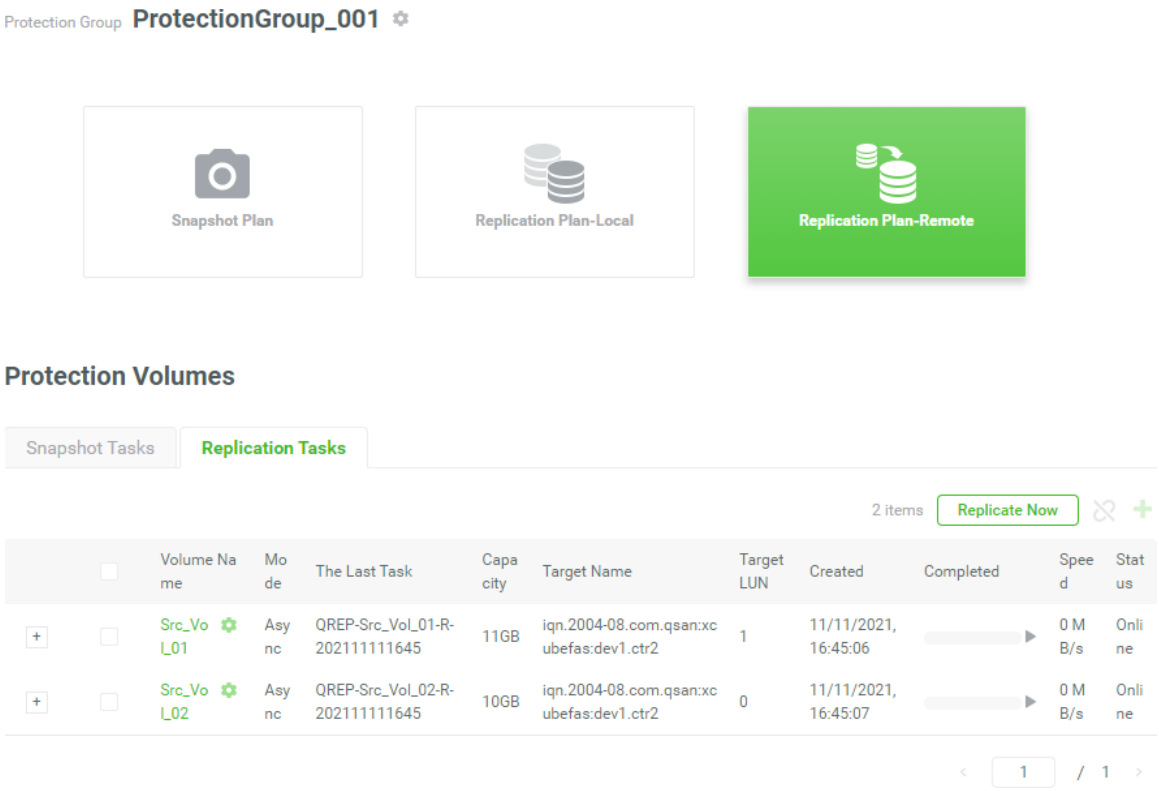

Figure 3-67 Asynchronous Replication – Local to Remote Step 10

27. Done. A protection group with synchronous replication task has been reconnected.

4. Use Cases

This chapter will introduce some use cases for remote replications.

4.1. Remote Replication Topologies

Remote replication supports multiple topologies to suit various disaster recovery configurations. They are one-directional, bi-directional, one-to-many, many-to-one, and many-to-many. Both the source volume and destination volume in a replication connection are exclusive to the pair. Either one can NOT be served as the source or destination volume of a different replication connection. Below are the supported topologies.

One-Directional

Figure 4-1 One-Directional Remote Replication

A Source Volume (S) in Site A is replicating to a Target Volume (T) in Site B. This is the most basic remote replication topology.

Bi-Directional

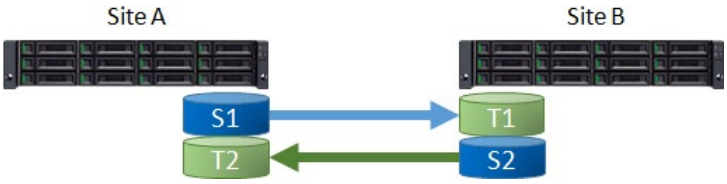

Figure 4-2 Bi-Directional Remote Replication

Each system in a two-system topology acts as a replication target for the other's production data. A Source Volume (S1) in Site A is replicating to a Target Volume (T1) in Site B. And a Source Volume (S2) in Site B is replicating to a Target Volume (T2) in Site A.

One-to-Many

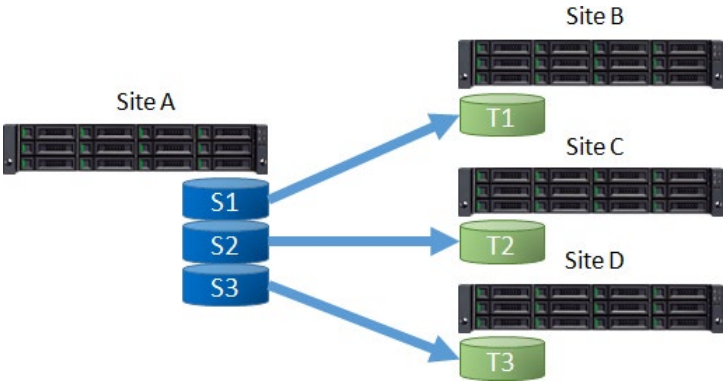

Figure 4-3 One-to-Many Remote Replication

A single source system replicates different storage resources to multiple target systems. A Source Volume (S1) in Site A is replicating to a Target Volume (T1) in Site B. At the same time, a Source Volume (S2) in Site A is replicating to a Target Volume (T2) in Site C. So does S3 in Site A to T3 in Site D.

Many-to-One

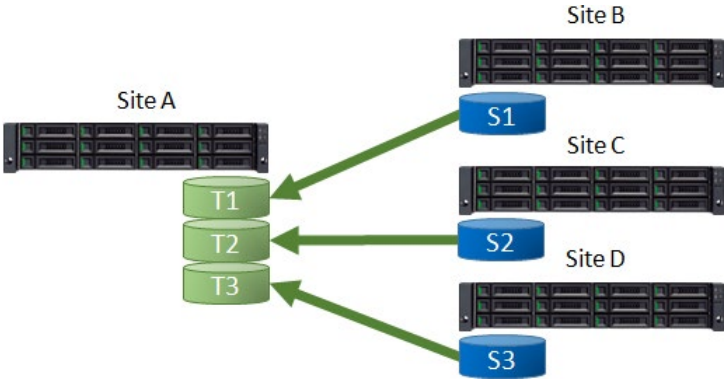

Figure 4-4 Many-to One Remote Replication

Multiple source systems replicate to a single target system. A Source Volume (S1) in Site B is replicating to a Target Volume (T1) in Site A. At the same time, a Source Volume (S2) in Site C is replicating to a Target Volume (T2) in Site A. So does S3 in Site D to T3 in Site A.

Many-to-Many

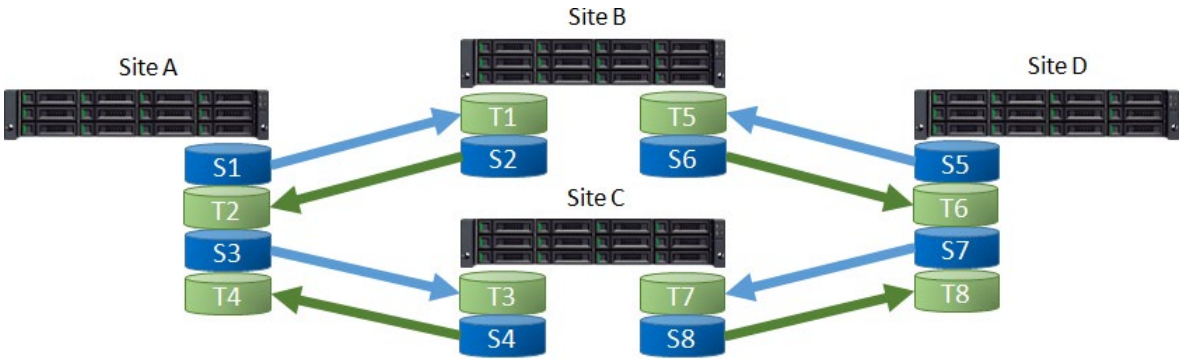

Figure 4-5 Many-to Many Remote Replication

Combination with bi-Directional, one-to-many, and many-to-one, remote replication also supports Many-to-Many topology. Multiple source systems replicate to multiple target systems. A Source Volume (S1) in Site A is replicating to a Target Volume (T1) in Site B. At the same time, a Source Volume (S2) in Site B is replicating to a Target Volume (T2) in Site A. And does S3 to T3, S4 to T4, …, S8 to T8.

TIP

Note that for each individual replication session in the topology, all supported topologies have a one-to-one configuration. For example, you cannot set a volume as a replication source and target at the same time. Similarly, a volume cannot be the source of synchronous or asynchronous replication at the same time.

4.2. Synchronous Replication Failover

Initially, you can set up normal usage of synchronous replication. The I/O sequence (blue lines) will go from Host to Source, and then Source to Target.

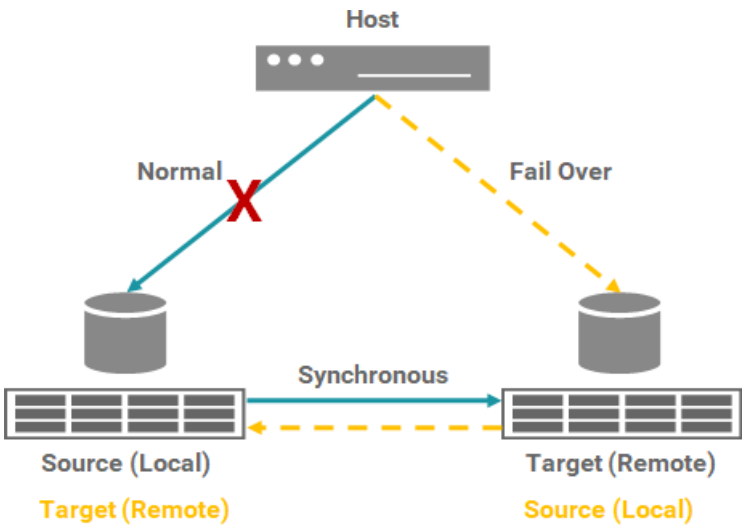

Figure 4-6 Synchronous Replication Failover

If an accident occurs and the Source is disconnected, you can manually failover to the target and let the I/O sequence (orange lines) go to the Target. At the same time, you can also switch Source and Target roles.

4.3. Synchronous Replication + C2F Solution

As we know, a volume with WT (Write-through) capability writes I/O to drives. This ensures that data is written to the drives. In contrast, a volume with WB (Write-back) capability temporarily writes I/O to cache and writes it to the drive-in batches as planned. Generally, a volume with write-back capability has better performance than a write-through volume.

If you use write-back parameter to set the volume, there are still risks. When the power is off unexpectedly, the data in the cache will disappear. C2F (Cache to Flash) solution can prevent this. Both Source and Target must add the C2F modules.

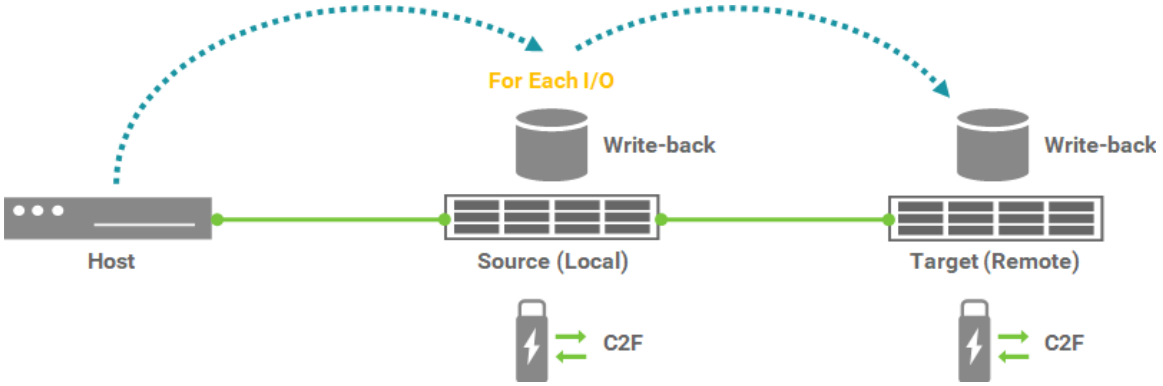

Figure 4-7 Synchronous Replication + C2F Solution

The best practice of volume write-back plus C2F solution can balance performance and data integrity.

5. Conclusion

This document discusses the various replication solutions. Configuring a data protection solution helps prevent unexpected situations, such as data loss. We provide local and remote data protection solutions to help easily recover in the event of a disaster. By using both synchronous and asynchronous replication solutions, data protection can be configured to meet the needs of enterprises.

Synchronous replication usually writes data to source unit and replicas at the same time. In this way, the primary copy and the replica are always in sync. In the event of a disaster, synchronous replication can provide maximum protection and ensure zero data loss.

Asynchronous replication first writes the data to the source unit and then copies the data to the replica, so there will be a certain delay in copying the data to the target unit. Asynchronous replication utilizes the snapshot technology to provide a consistent point-in-time copy that can be used in the event of a disaster. When using asynchronous replication, there will be no impact on host I/O because data will not be replicated when it enters the system. When data needs to be replicated over long distances, asynchronous replication can meet the needs of the enterprise.

6. APPENDIX

6.1. Apply To

Synchronous replication requires a **License** to enable the feature. We also provide a 30-day free **Trial License** for evaluation. The synchronous replication function is applicable to the following models.

- XEVO firmware 2.1.0 and later

6.2. Reference

Software Manuals

- XEVO Software Manual