Hewlett Packard Enterprise

Technical white paper

HPE MSA Remote Snap



Contents

Introduction 7
Array concepts background 7
Remote Snap product overview 7
What's new? 8
Benefits of Remote Snap 8
Disaster Recovery 8
Backup 8
Development 9
Components of Remote Snap 9
How it works 9
Replication process 10
Remote Snap requirements and limits 12
Limits 12
Restrictions and caveats 13
Peer connections 13
Replication sets 13
Snapshot space 14
Network requirements 15
Remote Snap basic functions 15
Using the SMU on HPE MSA Gen6 Storage and HPE MSA Gen7 Storage 15
Limitations and advantages using the SMU 17
License operations 18
Verifying the license 18
Installing the license 19
Peer connection operations 20
Checking for connectivity with another array 20
Creating a peer connection 23
Displaying peer connections and verifying links 24
Modifying peer connections 26
Deleting a peer connection 28
Peer connections and CHAP 30
Replication set and replication operations 34
Creating a replication set 34
Modifying a replication set 50
Scheduling replications 51
Initiating a replication 54
Suspending a replication set 55



Resuming a replication or replication set 57
Aborting a replication 59
Displaying a replication set 61
Showing replication snapshot history 64
Deleting a replication set 68
Accessing the data of the secondary volume 69
Creating a snapshot of a secondary volume 70
Using replication history snapshots 71
Disaster Recovery 71
Failover 72
Failback restore 72
Failback no restore 73
Reverse 74
Use cases 75
Multiple remote offices backup to one central office 75
Application consistent replications and VSS 75
Best practices 75
Fault tolerance 75
Volume size and policy 75
Scheduling 76
Application-consistent snapshots 76
Performance tips 76
Troubleshooting 76
Peer connections 76
Licensing 76
CHAP 76
Replications 77
Replication suspended 77
Unable to control or modify a replication set 77
Errors 77
The format for this section has a general heading for the component, subsections for specific errors with error messages seen in the SMU or CLI, as applicable, and how to resolve the error 77
Component heading 77
Licensing 77
Pools 77
CHAP 79
Peer Connection 79
Limits 81
Secondary volume 81

Volume groups 82
Replication set 83
Disaster recovery 86
Events 87
SNMP traps, email (SMTP) notification, and syslog 89
Frequently asked questions 89
Does Remote Snap support port failover and load balancing? 89
Can I expand the primary volume or secondary volume? 89
Can I map the secondary volume? 89
Why were replication history snapshots not taken? 89
Can the host ports run Remote Snap and I/O concurrently? 89
How does CHAP affect replication? 89
How do I delete a virtual replication set when the peer connection is down? 89
Why can't I delete a peer connection when no replication sets are present in the system? 90

Figures

Figure 1. Replication cycle 11
Figure 2. Areas of the SMU 16
Figure 3. Verifying the license 19
Figure 4. Querying a peer connection 22
Figure 5. Peer connection query results 23
Figure 6. Displaying peer connections 26
Figure 7. Modifying a peer connection 28
Figure 8. Deleting a peer connection 29
Figure 9. Add CHAP record – initial step 31
Figure 10. Add CHAP record – create the record 32
Figure 11. Enable CHAP 33
Figure 12. Reset host links 34
Figure 13. Accessing the Data Protection Configuration wizard (select volumes > Select Volume Action > Add Data Protection) 36
Figure 14. Accessing the Data Protection Configuration wizard (volume slide-over) 37
Figure 15. Accessing the Data Protection Configuration wizard (Replications > Add Data Protection) 38
Figure 16. Accessing the Data Protection Configuration wizard (Data Protection table view > Add Data Protection) 39
Figure 17. Data Protection Configuration wizard – Volumes – Data Protection Type (Remote Replication) and volume selection 40
Figure 18. Data Protection Configuration wizard – Connection – Data Protection Connection – existing or new 41
Figure 19. Data Protection Configuration wizard – Connection – Data Protection Connection – new – specify the Port Address 42



Figure 20. Data Protection Configuration wizard – Connection – Data Protection Connection – new – specify peer connection name and remote credentials 43

Figure 21. Data Protection Configuration wizard – Connection – Data Protection Connection – new – complete 44

Figure 22. Data Protection Configuration wizard – Connection – Data Protection Connection – choose existing 45

Figure 23. Data Protection Configuration wizard – Connection – Replication Settings – name, secondary name and pool, queue policy 46

Figure 24. Data Protection Configuration wizard – Connection – Replication Settings – replication history snapshots 47

Figure 25. Data Protection Configuration wizard – Schedules – Replication Schedule Options – Replication Schedule – name, start date and time, and repetition 48

Figure 26. Data Protection Configuration wizard – Schedules – Replication Schedule Options – Replication Schedule – end condition, date and time constraints 49

Figure 27. Data Protection Configuration wizard – Summary – Summary of Data Protection Settings 50

Figure 28. Modifying a replication set – Snapshot Count 51

Figure 29. Modifying a replication schedule – selecting the schedule to modify 53

Figure 30. Modifying a replication schedule – editing the schedule 54

Figure 31. Starting a replication 55

Figure 32. Suspending a replication set – no replications in progress 56

Figure 33. Suspending a replication set in progress 57

Figure 34. Resuming a replication set – no suspended replications 58

Figure 35. Resuming a suspended replication set 59

Figure 36. Aborting a running replication 60

Figure 37. Aborting a suspended replication 61

Figure 38. Displaying a replication set – volume table 63

Figure 39. Displaying a replication set – details 64

Figure 40. Displaying replication snapshot history – configuration 66

Figure 41. Displaying replication snapshot history – Data Protection table view 67

Figure 42. Displaying replication snapshot history – Snapshots tab 68

Figure 43. Deleting a replication set 69

Figure 44. Data Protection Configuration wizard – Volumes – Data Protection Type (Local Snapshots) and volume selection 70

Figure 45. Data Protection Configuration wizard – Schedules – Snapshot Schedule Options 71

Command Examples

Command Example 1. Verifying the license 18

Command Example 2. Querying a peer connection 21

Command Example 3. Creating the peer connection 23

Command Example 4. Displaying peer connections and verifying links 25

Command Example 5. Modifying a peer connection 27

Command Example 6. Deleting a peer connection using the local-only option 29



Command Example 7. Setting up CHAP 30
Command Example 8. Creating a replication set 35
Command Example 9. Modifying the name and snapshot count of a replication set 50
Command Example 10. Creating a replication task 52
Command Example 11. Creating a replication schedule 52
Command Example 12. Initiating a replication 54
Command Example 13. Suspending a replication set 56
Command Example 14. Resuming a replication set 57
Command Example 15. Aborting a replication 59
Command Example 16. Displaying a replication set 62
Command Example 17. Displaying replication snapshot history 65
Command Example 18. Deleting a replication set using the local-only option 69
Command Example 19. Creating a snapshot of a secondary volume 70
Command Example 20. Recover replication set – failover operation 72
Command Example 21. Mapping the secondary volumes 72
Command Example 22. Recover replication set – failback restore 73
Command Example 23. Recover replication set – failback restore – completion 73
Command Example 24. Recover replication set – failback restore – remapping the primary volumes 73
Command Example 25. Recover replication set – failback no restore 74
Command Example 26. Recover replication set – reverse 75

Tables

Table 1. Configuration limits 12
Table 2. Replication configuration limits for HPE MSA 1060 Storage 12
Table 3. Replication configuration limits for HPE MSA 2060/2062 Storage and HPE MSA 2070/2072 Storage 13
Table 4. Internal replication snapshots by Queue Policy 14
Table 5. Relevant icons in the SMU 16



Introduction

This document provides information for using HPE MSA Remote Snap (Remote Snap). The following topics are covered:

- Benefits
- Components
- How it works
- Requirements
- Basic functions
- Use cases
- Best practices
- Troubleshooting
- Frequently asked questions

This document covers replications on HPE MSA 1060/2060/2062 Storage (HPE MSA Gen6 Storage) and HPE MSA 2070/2072 Storage (HPE MSA Gen7 Storage) only. HPE MSA 1060/2060/2062 Storage also supports replications with HPE MSA 1050/2050/2052 Storage; for information about managing replications on HPE MSA 1050/2050/2052 Storage, see the previous version of the MSA Remote Snap Software technical white paper.

Array concepts background

HPE MSA Gen6 Storage and HPE MSA Gen7 Storage use virtual storage, a method of mapping logical storage requests to physical storage, abstracting the underlying storage subsystems. Physical disks are logically grouped into disk groups using one of the levels of RAID available. Disk groups are contained in one of two pools, and the storage in a pool is divided into pages, which are ranges of contiguous LBAs of a disk group.

Conceptually, a volume is a table of references to pages of data in the pool. These pages are distributed among the disk groups of the pool. The array can create a snapshot of a volume or another snapshot, providing an image of the data state of the volume. Creating the snapshot duplicates the base volume's table of references.

The snapshot uses redirect-on-write, which adds new references to pages of data to the table of references of the base volume, and is superior to copy-on-write, which copies the original data to the snapshot prior to writing the new data to the volume itself, incurring a penalty on performance for the copy.

For more information, see the HPE MSA Gen6 Virtual Storage Technical Reference Guide or the HPE MSA Gen7 Virtual Storage Technical Reference Guide.

Remote Snap product overview

Remote Snap, a licensed option, provides array-based asynchronous point-in-time replication from a primary array to a secondary array. Remote Snap replicates page-level data using the snapshot functionality of the array. The source might be a single volume or multiple volumes in a volume group; the destination matches the source – either a single volume or multiple volumes. The primary and secondary arrays might be at the same location or different locations. Remote Snap replicates efficiently by only replicating pages that have changed since the last replication. Remote Snap's fault-tolerant design allows replication to continue for array failures involving controllers, hard drives, I/O paths, or temporary power failure, based on the configuration of the arrays and the configuration of the replication connection. Remote Snap uses either FC or iSCSI for communication as well as I/O between arrays.

The following arrays support Remote Snap:

- HPE MSA 2070 32Gb Fibre Channel Storage
- HPE MSA 2070 16Gb Fibre Channel Storage
- HPE MSA 2070 10GbE iSCSI Storage
- HPE MSA 2070 10GBASE-T iSCSI Storage
- HPE MSA 2060 16Gb Fibre Channel Storage
- HPE MSA 2060 10GbE iSCSI Storage



- HPE MSA 2060 10GBASE-T iSCSI Storage
- HPE MSA 1060 16Gb Fibre Channel Storage
- HPE MSA 1060 10GBase-T iSCSI Storage

Visit HPE MSA Storage for more information on these controllers.

Note

Remote Snap is not supported on SAS models of HPE MSA Storage.

Remote Snap supports replication between arrays of the same generation, for example between two HPE MSA Gen7 Storage arrays, and between one generation and the previous generation, such as between an HPE MSA Gen7 Storage array and an HPE MSA Gen6 Storage array, with functionality limited to the more restrictive generation.

What's new?

This section describes new enhancements and support added with the release of the IN100 firmware for HPE MSA Gen6 Storage.

- Disaster recovery –HPE MSA Gen6 Storage and HPE MSA Gen7 Storage support the following disaster recovery operations:
  - Failover – Fails host access over to the secondary volume, allowing mapping the secondary volume to hosts, without changing primary volume access. The replication set will have a new status of *Failed Over*.
  - Failback restore – Used only when the replication set's status is *Failed Over*, this operation restores the original roles to the primary and secondary volumes, replicating the changes made to the secondary volume during failover back to the primary volume.
  - Failback no restore – Used only when the replication set's status is *Failed Over*, this operation restores the original roles of the primary and secondary volumes, discarding changes made to the secondary volume during failover.
  - Reverse – Used only when the replication set's status is *Failed Over*, this operation exchanges the roles of the primary and secondary volumes – the original primary volume becomes the secondary volume, and the original secondary volume becomes the primary volume.
- Automatic volume group creation when creating a replication set of multiple volumes in the graphical user interface Storage Management Utility (SMU).

Benefits of Remote Snap

Use Remote Snap for disaster recovery, backup and development. Remote Snap technology enables the following key data management and protection capabilities:

- Continuity of business systems in the event of a failure on the primary site
- Access to data at a remote site, for either dispersed operations or development activities
- Multiple recovery points using snapshots

Because Remote Snap supports either FC or Ethernet interconnects, businesses have the flexibility to use the technology that matches their current environment.

Disaster Recovery

Remote Snap provides access to data at a secondary site when the primary site experiences a critical failure. Replicating at regular intervals helps to protect the data. Recovery time is reduced because the data is available at the secondary site; applications can switch to the secondary site with minimal downtime using data from the last replication point. The data stored at the secondary site can then be used to restore the primary location once it is back online, or the secondary site can become the new primary site.

Backup

Remote Snap replicates volumes with marginal impact on server performance. It can be used by small businesses as a primary backup tool and by large businesses by replicating from remote offices to a central data center where backups occur. Remote Snap can be used as interim storage for backing up to removable media such as tape.



Development

Remote Snap enables the following development use cases:

- An application administrator can test patches or changes in the primary system by switching the applications to the secondary site. Once the testing of the patch update is completed, the administrator can switch the applications back to the primary site.
- A database application development team can have access to regularly scheduled snapshots of the replicated database volumes by taking traditional snapshots of the secondary volume or using replication history snapshots on the secondary array. When the traditional or replication history snapshot is no longer needed, it can be deleted.

Components of Remote Snap

- Primary volume—the volume that is the source of the replication. This volume can be mapped to hosts. The primary can be a volume, a snapshot, or a volume group.
- Secondary volume—the volume that is the destination or target of the replication. This can be either a volume or volume group, depending on the primary. Under normal operating conditions, this volume cannot be mapped to hosts, but it can be snapped, and the snapshot can be mapped to hosts using any access mode, including read-only or read-write. This snapshot can be reset to reflect the current contents of the secondary volume. If the replication set's status is *Failed Over*, the secondary volumes can be mapped.
- Secondary and primary volume replication snapshots—internal snapshots of the primary or secondary volume used in replication. There are three potential internal snapshots: the current snapshot, the previous snapshot, and a queued replication snapshot. Though these internal snapshots are not listed, and they do not count against license limits, they do consume space and count against volume count limits for snapshot trees and pools. They can be listed with their storage consumption using the `show snapshots type replication` command-line interface (CLI) command. The term "replication snapshot" when used without a modifier refers to one of these internal snapshots.
- Replication set—a primary and secondary volume pair connected for the purposes of replication, along with their associated replication snapshots. The replication set includes the peer connection used. See the following item for more information on the peer connection. If the replication set includes a volume group, you cannot add or remove volumes from the volume group. New for HPE MSA Gen6 Storage, the direction of a replication can be changed using the disaster recovery operations.
- Peer connection—the host ports used between two arrays involved in replications. These host ports are used for managing replications as well as transferring data for replication; the management ports of the array are not used for Remote Snap. Peer connections are bidirectional and can use FC or iSCSI host ports. All replications that share a peer connection use the same host ports.
- Replication snapshot history—snapshots taken either at the start (for the primary volume) or the end (for the secondary volume) of a replication to track the history of the replications of a replication set over time. The snapshots are taken of the secondary volume only, or both the secondary and primary volumes. Replication snapshot history applies only to a replication set of a volume, not to a replication set of a volume group. When referring to these snapshots, the modifier "history" will be included in the term.

How it works

Replication snapshots are used to track the data to be replicated. To perform a replication, the array takes a replication snapshot of the volume to be replicated, creating a point-in-time image of the data. The array copies this point-in-time image to the destination volume by copying only the differences in the data between the current replication snapshot and the previous one to the replication snapshot or snapshots on the secondary array. Once the differences have been copied, the secondary volume or volume group is rolled back to the replication snapshot or snapshots to complete the process.

Remote Snap functions independently of the disk group or pool configuration, so the secondary volume in a set might have different underlying tiers, RAID levels, drive counts, drive sizes or drive types than the primary volume, though the volume sizes are identical. Since Remote Snap functions at the raw page-level, it has no knowledge of the volume's operating system configuration, the file system, or any data that exists on the volume.

Remote Snap uses a push model, where the primary volume's array writes data to the appropriate replication snapshot on the secondary volume.



Replication process

Replication uses the following steps:

1. Save the current replication snapshots of the primary and secondary volumes to previous replication snapshots.
2. Reset the current replication snapshot of the primary volume; if the queue policy is set to *Queue Latest*, and a queued replication snapshot exists, copy the queued replication snapshot to the current primary replication snapshot. When replicating volume groups, replication snapshots of all volumes are reset simultaneously, ensuring consistent copies of the group of volumes. Create replication history snapshot of the primary volume, if enabled.
3. Compare the current primary replication snapshot to the previous primary replication snapshot.
4. Copy the differences to the current secondary replication snapshot. For the initial replication, all allocated pages are copied, but only allocated pages—no copying of empty unallocated pages. For subsequent replications, only the difference between the current primary replication snapshot and previous primary replication snapshot are copied.
5. Rollback the secondary volume to the current secondary volume snapshot. When replicating volume groups, snapshots of all volumes are rolled back simultaneously, after the replication of changes has completed for all volumes in the group. Create replication history snapshot of the secondary volume, if enabled.

When a new replication request occurs while a previous replication is in progress, if the queue policy is set to *Queue Latest* the array takes a snapshot of the volume and puts it in the queued replication snapshot, replacing any previous snapshot with the new one, thus queuing at most one snapshot. If the queue policy is set to *Discard*, the replication request is discarded. Once the replication that is in progress is complete, the queued snapshot is copied to the primary volume's current snapshot and the replication of the queued snapshot continues.

See Figure 1. Replication cycle for a visual representation of the replication topology and process.

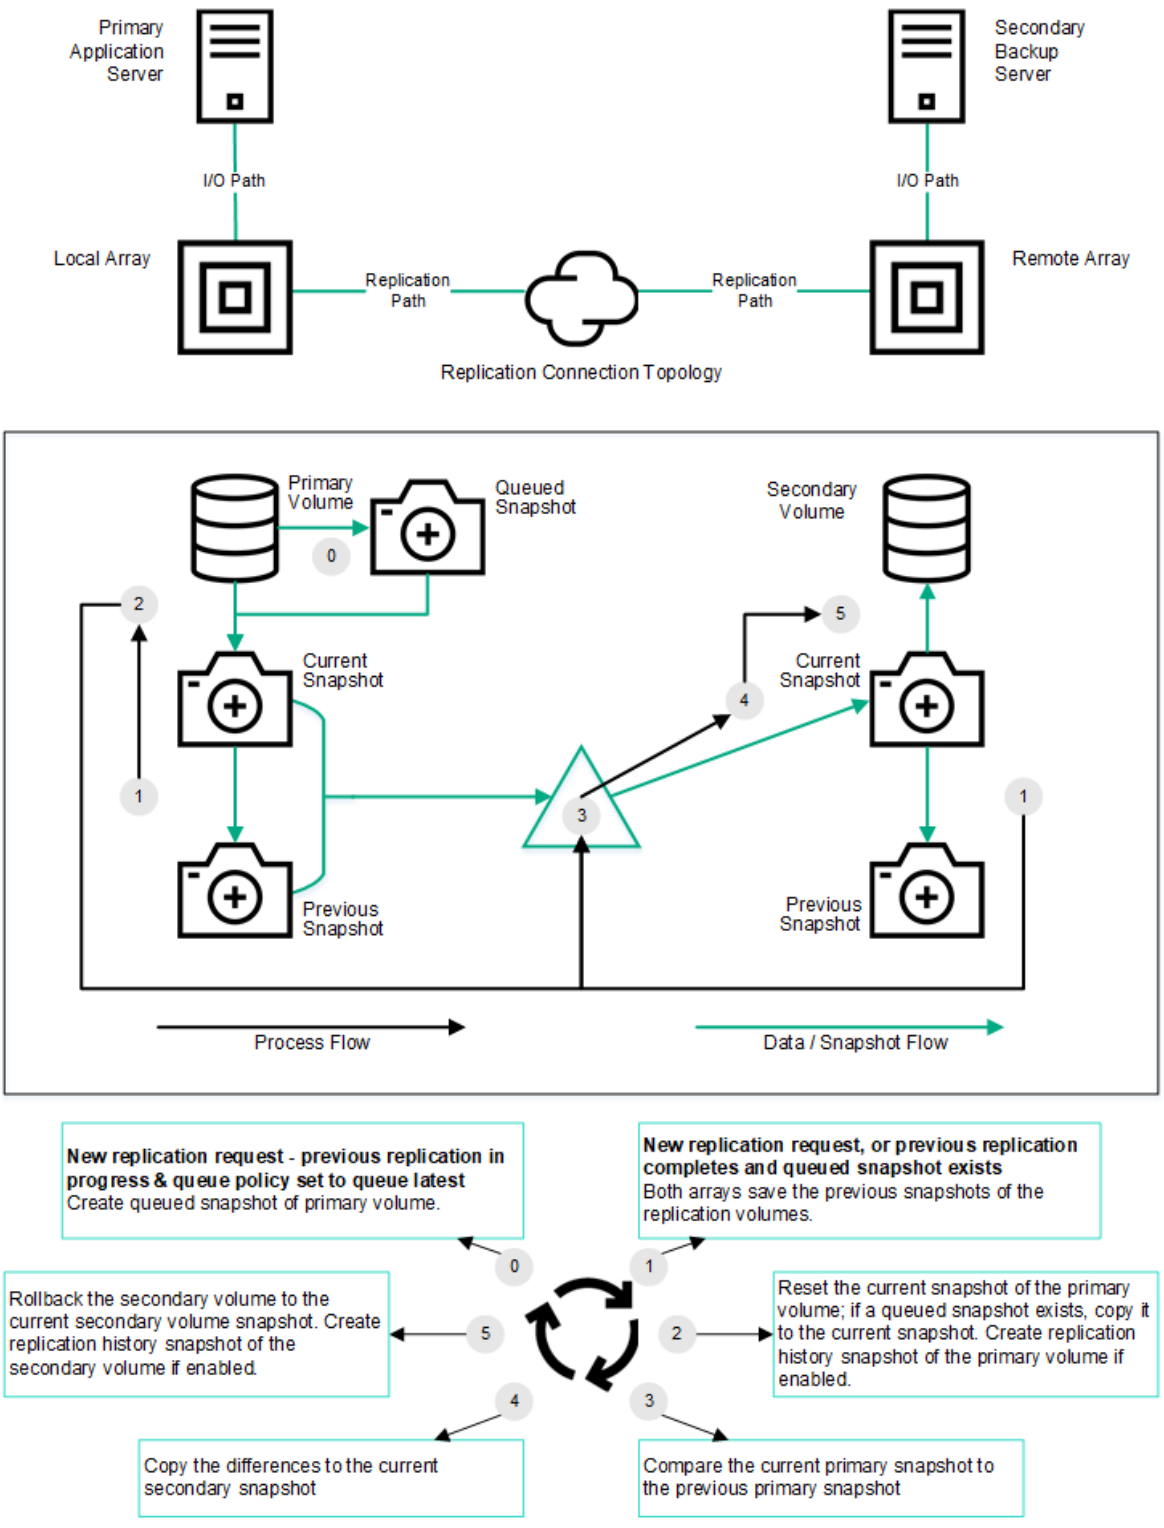


Figure 1. Replication cycle

Remote Snap requirements and limits

Remote Snap requires:

- Two arrays with all reachable ports configured and connected via a switch; direct connection between arrays is not supported
  - For iSCSI, the host ports of the two arrays do not have to be in the same subnet but must be connected to a network whose route tables and firewall or port-blocking settings allow iSCSI traffic (port 3260) to pass between the two arrays
  - For FC, the two arrays must be in the same zone as defined by the switches
- The Advanced Data Services license, which includes Remote Snap, must be installed on both arrays
- At least one pool on each array to create a peer connection between the two arrays
- At least one volume on the primary array to create a replication set

Limits

The two or three internal replication snapshots and all replication history snapshots count against the volumes per pool limit. Monitor the number of replication history snapshots created to avoid unexpectedly reaching this limit. Restrict the number of replication history snapshots using the `snapshot-count` parameter of the `set replication-set` command.

Table 1. Configuration limits

Property	Value
Maximum non-MSA-DP+ disk groups per pool / controller	16
Maximum MSA-DP+ disk groups per pool / controller	4
Maximum volumes	1024
Maximum volumes per pool	1024 (512 base volumes and 512 snapshots)
Maximum snapshots per volume	254

Table 2. Replication configuration limits for HPE MSA 1060 Storage

Property	Value
Peer connections	1
Replication sets	32
Replication volumes (primary and secondary)	32
Volumes per volume group replicated	16



Table 3. Replication configuration limits for HPE MSA 2060/2062 Storage and HPE MSA 2070/2072 Storage

Property	Value
Peer connections	4
Replication sets	32
Replication volumes (primary and secondary)	32
Volumes per volume group replicated	16

Restrictions and caveats

Peer connections

Remote Snap does not support SAS. Refer to the Remote Snap product overview in the Introduction (page 7).

Note

For more information on controller types and additional specifications, view the QuickSpecs for the product by visiting the HPE Support Center, enter the model in the search field, select the product from the list of results to further refine the results, click Parts/Specifications, and select the document with the QuickSpecs Subtype.

A peer connection can include any I/O path of the appropriate protocol available between the two arrays; at most 2 paths will be used in the peer connection; optimal paths will be used until depleted or until they fail – only then will nonoptimal paths be used.

If you want to use CHAP for authentication, set the CHAP secret to be the same on each array, and enable CHAP on each array; see Peer connections and CHAP on page 30.

If you want to use jumbo frames, verify that jumbo frames is enabled on all devices—arrays, routers and switches—that are in the replication datapath.

Remote array username and password authentication is required when creating or modifying peer connections.

Replication sets

A volume or snapshot may participate in one replication set only; you cannot create a replication set for a volume already participating in a replication set for a volume group.

For replication sets of volume groups, volumes cannot be removed from nor added to either the primary volume group or the secondary volume group once the replication set has been created. You cannot create a replication set for a volume group if a member of the volume group is already participating in a replication set. The maximum number of volumes for a group is 16. Replication snapshot history is not available for replication sets of volume groups. To create a replication set that includes a volume group in the SMU, select multiple volumes during volume selection when adding data protection; the volume group is created automatically.

Creating a replication set creates a matching volume or volume group on the secondary array automatically with the same volume sizes as the primary volumes or snapshots. If the primary is a volume or snapshot, a volume name other than the primary volume or snapshot name may be specified for the secondary volume name, otherwise, the primary name will be used by default. If the primary is a volume group, a volume group with the same name as the primary volume group, and volumes with the same names as those in the primary volume group, will be created on the secondary array. If a volume already exists on the secondary array with the same name as the volume intended to be created, or if a volume group exists on the secondary array with the same name as the primary volume group, creation of the replication set will fail.

If the secondary pool is not specified, the array will use the pool on the remote array with the same name as the pool containing the primary volume, snapshot or volume group on the local array. If the secondary pool does not exist, creation of the replication set will fail.

Creating a replication set creates two or three internal replication snapshots per volume in the replication set on HPE MSA Gen6 Storage or HPE MSA Gen7 Storage, depending on the queue policy. See Table 4. Internal replication snapshots by Queue Policy for the number of snapshots created based on Queue Policy.



Table 4. Internal replication snapshots by Queue Policy

Queue Policy	Discard	Queue Latest
Number of internal replication snapshots created	2	3

Once a replication set is created, none of the volumes in the replication set can be deleted.

Primary volumes or volumes in a primary volume group may be expanded; changes will be reflected in the secondary volumes upon the next replication; secondary volumes may not be expanded directly.

Volumes in a replication set will have new role, primary or secondary, as shown in the `show volumes details` CLI command.

Replication snapshot history is available only for a replication set of a volume, not for a replication set of a volume group. If enabled, when a replication occurs, a user accessible snapshot is taken of either the secondary volume only, or both the secondary and primary volumes (but not for the primary volume only). Primary volume snapshots are created when the replication begins, while secondary volume snapshots are created when the replication completes. You can set the number of snapshots to retain from 1 to 16 – the retention count applies to both the primary (if enabled) and secondary. Once the snapshot count is exceeded, the oldest unmapped history snapshot is discarded automatically.

Replication history snapshot names are of the form <basename>+"_nnnnnn", where nnnnnn is an integer, starting at 000000, that is incremented for each subsequent snapshot. If primary history snapshots are enabled, snapshots of the same name (including digits) will exist on both the primary and secondary arrays. If attempting to create a history snapshot, but a volume or snapshot with the name already exists, the history snapshot will not be created; however, the snapshot number is still incremented. In fact, all successful replication requests increment the number identifier of the history snapshot's name, even when replication snapshot history is *Disabled*. For example, if currently replicating with the associated history snapshot SrcDstRepHist_000004, the queue policy is set to *Queue Latest*, and four more replications are attempted before it is done, then when the replication is done, the next history snapshot will be SrcDstRepHist_000008 (SrcDstRepHist_000005 is replaced by SrcDstRepHist_000006 which is replaced by SrcDstRepHist_000007 which is replaced by SrcDstRepHist_000008 for the history snapshot name for the snapshot in the queue). Similarly, if a replication set is created without replication snapshot history enabled, five replications occur, then replication snapshot history is enabled with the snapshot basename SrcDstRepHist, the first replication history snapshot will be SrcDstRepHist_000006. If the queue policy is set to *Discard* and a replication request occurs while a replication is already in progress on the replication set, the replication request is *not* successful, and the snapshot number is not incremented.

Replication history snapshots count against the array-wide maximum licensed snapshots. Set the retention priority of replication history snapshots created to never-delete, high, medium or low—this specifies the retention priority in regard to snapshot space, not the snapshot history retention count; see the `snapshot-retention-priority` option of the `set volume` CLI command for more information on snapshot retention priority in the online help or the *HPE MSA 1060/2060/2062 CLI Reference Guide* or the *HPE MSA 2070/2072 CLI Reference Guide*. To view the guide, visit HPE Support Center, enter the model in the search field, select the product from the list of results to further refine the results, then click Reference Information from the Manuals tab.

Manually created snapshots of either the primary or the secondary do not count against the retention limit and are not managed by the replication snapshot history feature.

Snapshot space

Since Remote Snap compares snapshots to determine data to transfer for replication, care must be given to provide sufficient space for the snapshots, both internal and, if configured, replication history snapshots. Overall space requirements depend on the rate of data change and frequency of replication. To see the storage consumed by the replication snapshots, use the `show snapshots type replication` CLI command, and to see the storage consumed by the replication history snapshots, use the `show replication-snapshot-history` or `show snapshots` CLI commands.

If overcommit is disabled for the pool, snapshots are fully provisioned, including the replication snapshots; storage is allocated to the snapshots equal to the size of the volume. If overcommit is enabled, the space for all snapshots, including the replication snapshots, has a soft limit as a percentage of the pool's total space on the snapshot space; the default snapshot space limit is 10%. Set the snapshot space limit and the snapshot space limit policy using the `set snapshot-space` CLI command. The snapshot space limit policy can be set to *notify-only* (the default), or to *delete* snapshots. If the policy is set to *notify-only*, when the snapshot space is reached an event is generated. If the policy is set to *delete* snapshots, when the snapshot space exceeds the limit, an event is generated, and existing eligible snapshots are deleted until the snapshot space drops below the limit. No matter the setting of the limit policy, more snapshot space may be allocated for existing snapshots, if available.

Note that adding or removing disk groups from the pool changes the size of the snapshot space; if removing a disk group causes the snapshot space to exceed the limit, and the policy is set to *delete* snapshots, auto-deletion of snapshots occurs. Attempts to lower the limit below the current allocated size of the snapshots will fail, as this may otherwise cause auto-deletion of snapshots. Only a snapshot that is an unmapped leaf, that is, there are no snapshots of the snapshot and the snapshot is not mapped, will be considered for auto-deletion.

Retention priority is also considered when auto-deleting; the retention priority is inherited from the snapshot's parent, or can be reset on the snapshot itself, but changing the parent's priority does not propagate to existing children. Snapshot deletion is by priority first, then by date, oldest first. Snapshots are deleted one at a time until the snapshot space drops below the limit. If auto-deletion is in effect but no eligible snapshots exist, and snapshot space is still above the limit, no new snapshots can be created, but existing snapshots can continue to consume space; to allow new snapshots to be created, change the priority of one or more snapshots from *never-delete* to another priority, unmap a leaf snapshot, change the limit to a higher value, or add one or more disk groups to the pool. You can set thresholds for the snapshot space; the default for the low threshold is 75% of the snapshot space, the default for the middle threshold is 90%, and the default for the high threshold is 99%.

Network requirements

For FC peer connections, if your Fibre Channel fabric is partitioned into zones, at least one host port on each array must be in the same zone to allow connectivity. For iSCSI peer connections, the two arrays do not have to be in the same subnet but must be connected to a network whose route tables and firewall or port-blocking settings allow iSCSI traffic (port 3260) to pass between the two arrays. CHAP must be configured appropriately on both arrays, see Peer connections and CHAP (page 30). If jumbo frames are enabled on either array, it must be enabled on the other array and all network devices in the path between the arrays.

While there are no absolute minimal network performance requirements, your replication window and rate of data change will determine the network performance necessary. The replication window is simply the time it takes for a replication to complete – you'll want to finish the last replication before starting another. This is like a recovery point objective (RPO), which is the maximum acceptable loss of data measured in time. The minimum RPO should be roughly double the replication window, assuming the replication takes the entire window to complete. In other words, the replication window should be roughly half of the RPO. This is because data that changed right after a replication begins will not be replicated and will not be available at the secondary array until the end of the next replication.

You can estimate the rate of data change by using two snapshots. Create a snapshot of the volume to register the content, then after the time determined by the replication window, create another snapshot of the volume. Compare the sizes of the two snapshots – in the SMU, navigate to the **Data Protection** table view of the **Volumes** panel (**Provisioning** > **Volumes** > **Data Protection** table view) and expand the volume's entry using the expand icon (), and note the Used value for the two snapshots; using the CLI, use the `show snapshots` command and note the `Snap Data` values. The difference in size will be roughly the size of the replication, within a couple of percent. To continue estimating, reset the oldest snapshot when the amount of time determined by the replication window expires and compare sizes. See Using the SMU on HPE MSA Gen6 Storage and HPE MSA Gen7 Storage (page 15) for more on using the SMU.

Remote Snap basic functions

The following functions are highlighted in this section:

- Verifying and installing the license
- Checking and creating a peer connection
- Creating a replication set
- Scheduling replications
- Deleting a replication set
- Accessing the secondary volume's data
  - Creating a snapshot of a secondary volume
  - Using a replication history snapshot of a secondary volume
- Disaster recovery operations

For more information on Remote Snap functions, see the *HPE MSA 1060/2060/2062 Storage Management Guide* or the *HPE MSA 2070/2072 Storage Management Guide*. To view the guide, visit HPE Support Center, enter the model in the search field, select the product from the list of results to further refine the results, then click Reference Information from the Manuals tab.

Using the SMU on HPE MSA Gen6 Storage and HPE MSA Gen7 Storage

The SMU on HPE MSA Gen6 Storage and HPE MSA Gen7 Storage is task oriented (see Figure 2. Areas of the SMU for an example of the SMU). To perform a task, select the top-level option from the menu pane on the left side of the interface, outlined in purple in Figure 2, which expands a drop-down list of menu choices, then click the option from the list. Use the management pane on the right-hand side of the interface, outlined in gold in Figure 2, to select the volumes, settings, or actions to perform. Once the top level and secondary option from the menu pane is selected, the management pane becomes the "panel" for that option. Most tasks may be performed in the **Volumes** panel

(**Provisioning** > **Volumes**), while some actions are performed in the **Peer Connections** panel (**Settings** > **Peer Connections**), some are performed in the iSCSI panel (**Settings** > **iSCSI**), and licensing actions are performed in the **Install License** panel (**Maintenance** > **Support** > **Licensing**). This white paper will outline options to select in the menu pane in purple, and outline items to select in the management pane in gold. In textual descriptions, navigation to a panel will be in bold, e.g. **Settings** > **Peer Connections**, as are names of fields to fill and menu items and buttons to click.

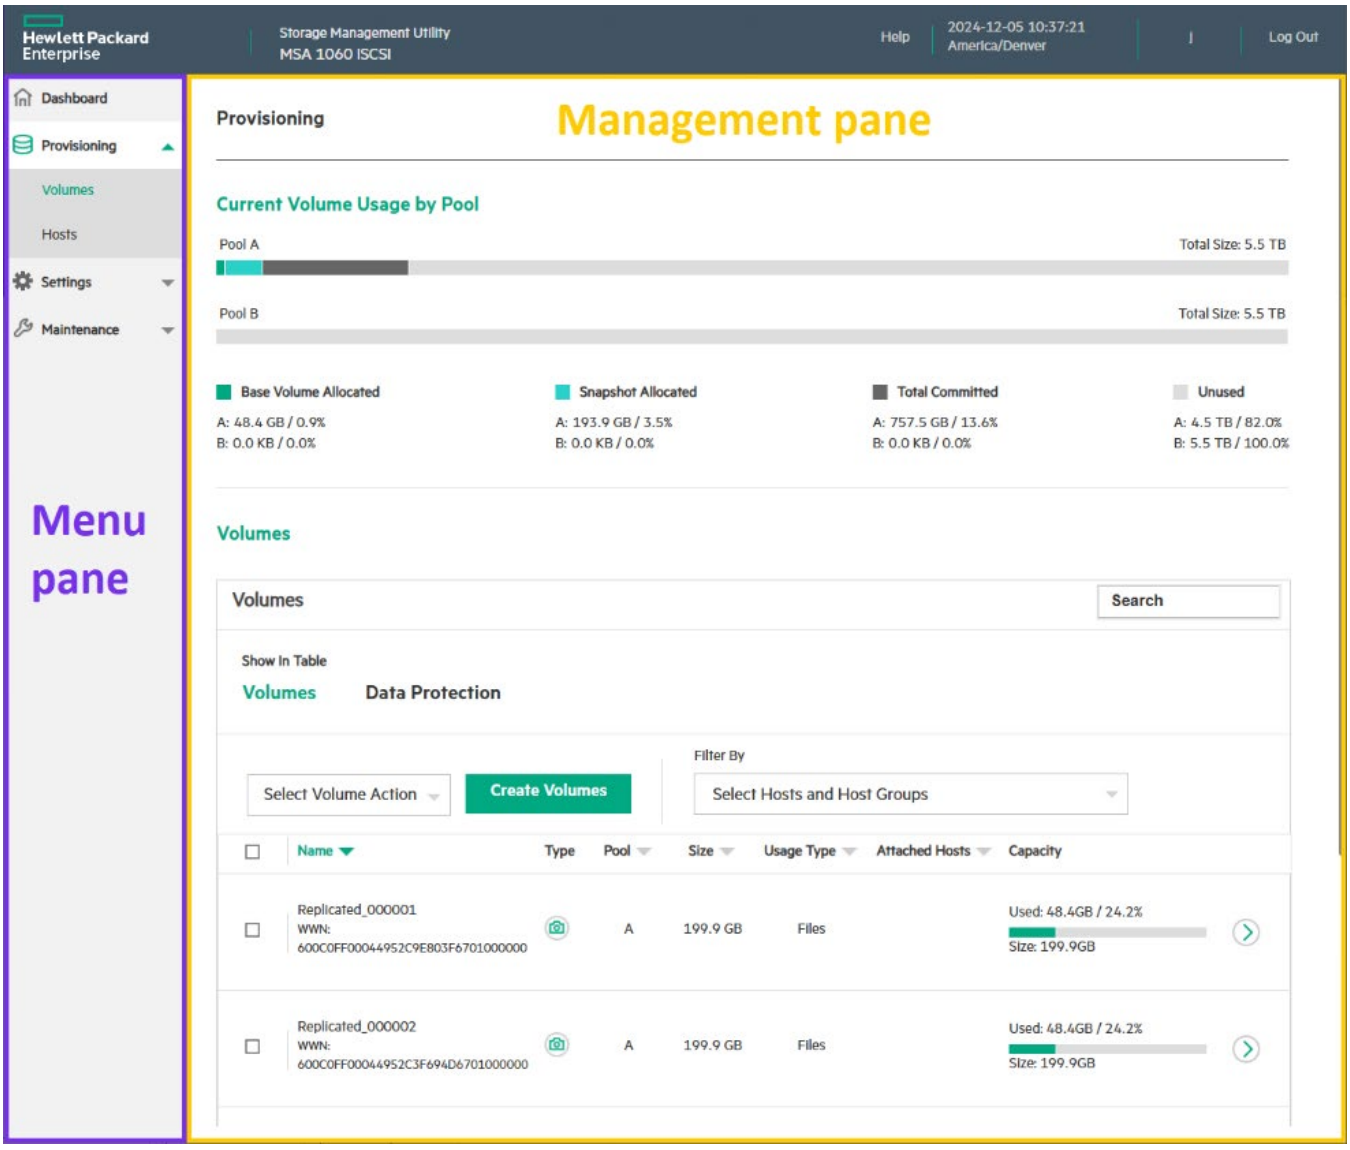


Figure 2. Areas of the SMU

There are several icons in the SMU for HPE MSA Gen6 Storage and HPE MSA Gen7 Storage that are associated with Remote Snap; see Table 5. Relevant icons in the SMU for a list of the icons and their usage:

Table 5. Relevant icons in the SMU

Icon	Name	Use
	Information	Opens a small window that defines or provides more information about a feature or option.
	Slide-over arrow (open)	Opens a panel that contains detailed information about an object

Icon	Name	Use
	Slide-over arrow (close)	Closes the panel previously opened using the slide-over arrow
	Expand	Expands a table row to provide more detail about an object
	Collapse	Collapses a table row to hide information about an object
	Edit	Edit a single value or options within an entire row or table
	Apply	Applies an edited operation or selection
	Cancel	Cancels an edited operation or selection
	Delete	Delete a value or object
	Primary Replication Volume	Identifies the primary replication volume
	Secondary Replication Volume	Identifies the secondary replication volume
	Schedule	Indicates that a specified task will take place at defined times
	Suspend	Suspends (pauses) an in-progress operation
	Resume	Resumes a suspended operation
	Abort / Cancel	Aborts or cancels an operation

Limitations and advantages using the SMU

The only time you can create a schedule for a replication set using the SMU is during the creation of the replication set—to create a schedule after creation of the replication set, create the task and schedule using the CLI; the schedule can be modified in the SMU. The advantage of using the SMU is that you can create both the schedule and its associated task when you create the replication set, rather than performing two separate steps in the CLI—one to create the task and one to create the schedule.

When manually initiating a replication, the SMU does not provide a means to use the last snapshot, or to specify a snapshot to use, of the primary volume, but will reset the current snapshot at the start of replication—if you need to use a specific snapshot, use the CLI.

The SMU does not differentiate between replication history snapshots and regular snapshots.

You cannot delete replication sets or peer connections in the SMU when the connection between the arrays is down or when one of the arrays is not functional.

To create a replication set of a volume group in the CLI, you must first create the volume group, then create the replication set; when using the SMU, you can simply select multiple volumes when creating the replication set—the volume group is created automatically.



License operations

A license is required on both the primary array and the secondary array to use Remote Snap.

Verifying the license

Using the CLI

To verify the license using the CLI, enter the `show license` command (see Command Example 1. Verifying the license), and look for the `Replication` field in the output (the field is underlined in this example):

```
# show license
License Key: 1c25a25517049d3caa73d6c2cf0bc98b
Licensing Serial Number: F5A440
…
Replication: Enabled
Replication Expires: Never
…
```

Command Example 1. Verifying the license



Using the SMU

Check if **License Installed** is **Yes** on the **Install License** panel (**Maintenance** > **Support** > **Licensing**) (see Figure 3. Verifying the license for an example with the license installed):

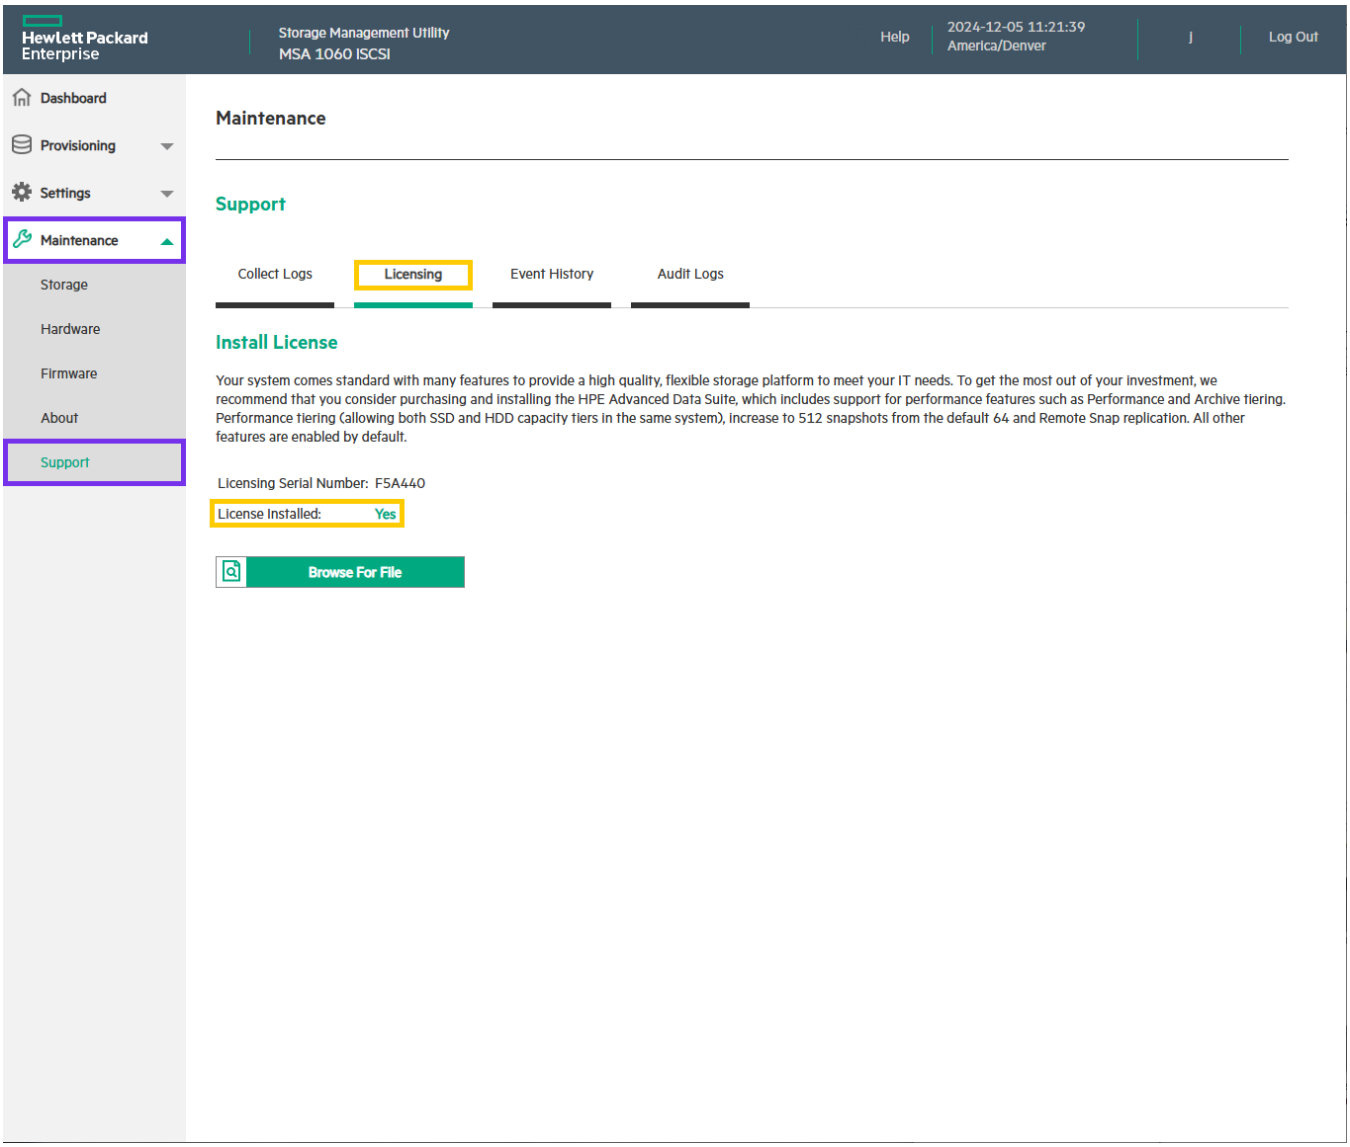


Figure 3. Verifying the license

Installing the license

Using SFTP or FTP

To install the license using SFTP (preferred) or FTP, verify the protocol is enabled, login to the SFTP/FTP interface, specifying the port if necessary, and enter the command `put <certificate.txt> license`.

Using the SMU

If the license is not installed, using the same **Install License** panel that you used to verify the license, click **Browse For File**, navigate to and select the license file, and when you return to the SMU, click **Upload File** (the button previously labeled **Browse For File**) to upload and install the license.

Peer connection operations

Replication requires a peer connection between the two arrays. The connection uses either SCSI or FC host ports for both communication and I/O. The connection includes all paths between the arrays—it uses up to two paths for load balancing. The connection is bidirectional—the concepts of primary and secondary do not apply to the connection. Also note that only one peer connection can exist between any two specific arrays.

Checking for connectivity with another array

To determine if there is connectivity to a remote array, provide some information about that array, and determine if a peer connection with that array is possible, use the peer connection query functionality. Using the CLI command provides more information than the SMU, specifically the licensing and possible port-to-port connections.



Using the CLI

Enter the `query peer-connection` CLI command (see Command Example 2. Querying a peer connection):

```
# query peer-connection 10.20.5.53
Info: This may take a few minutes to ping all port combinations...
Peer Connection Info
--------------------
System Name: Chicago
System Contact: System Admin (867-5309 sysadmin@hpe.com)
System Location: Rack5
System Information: Central Office MSA 2060 iSCSI
Midplane Serial Number: 00C0FFFB521E
Vendor Name: HPE
Product ID: MSA 2070 iSCSI
License Key: 0d92a4d54b9ba43afc4c73d6009448ef
Licensing Serial Number: FB521E
Maximum Licensable Snapshots: 1024
Base Maximum Snapshots: 64
Licensed Snapshots: 512
In-Use Snapshots: 0
Snapshots Expire: Never
Virtualization: Enabled
Virtualization Expires: Never
Performance Tier: Enabled
Performance Tier Expires: Never
Volume Copy: Enabled
Volume Copy Expires: Never
Replication: Enabled
Replication Expires: Never
VDS: Enabled
VDS Expires: Never
VSS: Enabled
VSS Expires: Never
SRA: Enabled
SRA Expires: Never

  Peer Controllers
  ----------------
  Controller: A
  Storage Controller Code Version: INS300R04-01
  Management Controller Code Version: IXM300R004-01
  IPv4 Address: 10.225.131.53
  Peer Host Name: Chicago-A
  IPv6 Address 1: Not Configured
  IPv6 Address 2: Not Configured
  IPv6 Address 3: Not Configured
  IPv6 Address 4: Not Configured

    Port  Type     Port Health  Port Address                                  Reachable Local Links
    ------------------------------------------------------------------------------------------------
    A1    iSCSI    Up           10.20.5.53                                    A1,B1
    A2    iSCSI    Up           10.30.5.53                                    A2,B2


  Peer Controllers
  ----------------
  Controller: B
  Storage Controller Code Version: INS300R04-01
  Management Controller Code Version: IXM300R004-01
  IPv4 Address: 10.10.5.53
  Peer Host Name: Chicago-B
  IPv6 Address 1: Not Configured
  IPv6 Address 2: Not Configured
  IPv6 Address 3: Not Configured
  IPv6 Address 4: Not Configured

    Port  Type     Port Health  Port Address                                  Reachable Local Links
    ------------------------------------------------------------------------------------------------
    B1    iSCSI    Up           10.20.5.153                                   A1,B1
    B2    iSCSI    Up           10.30.5.153                                   A2,B2
```

Command Example 2. Querying a peer connection



Using the SMU

Enter the iSCSI IP address or FC WWN in the **Peer System Port Address** field of the **Peer Connections** panel (**Settings** > **Peer Connections**) and click **Query Peer Connection** – see Figure 4. Querying a peer connection and Figure 5. Peer connection query results:

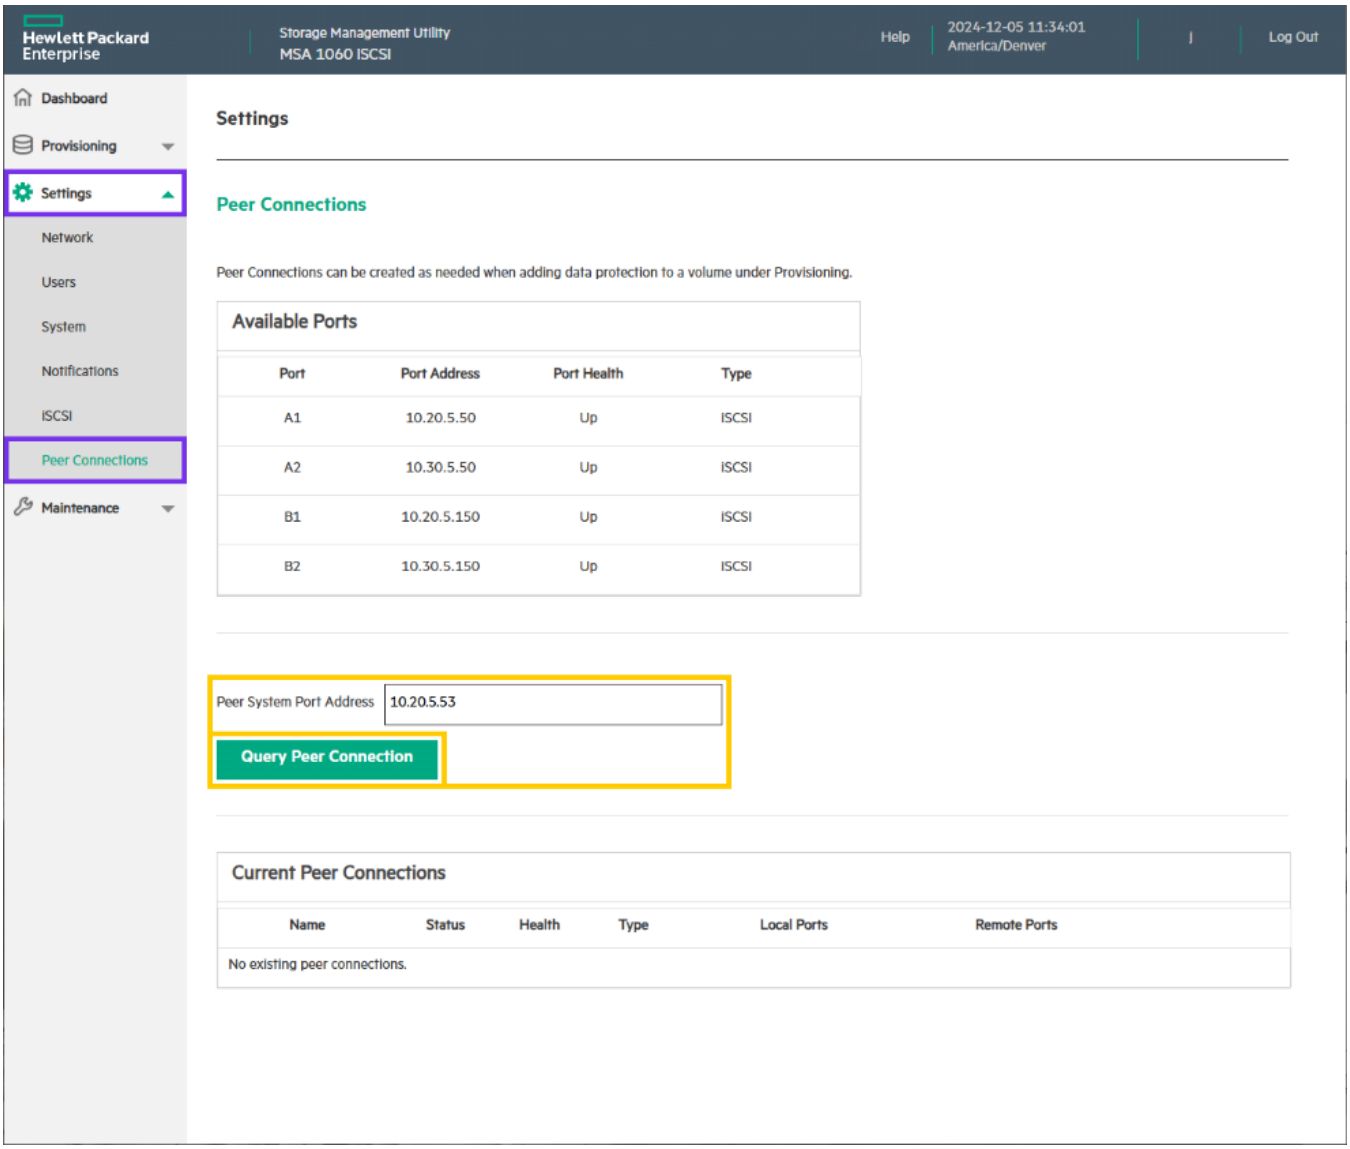


Figure 4. Querying a peer connection

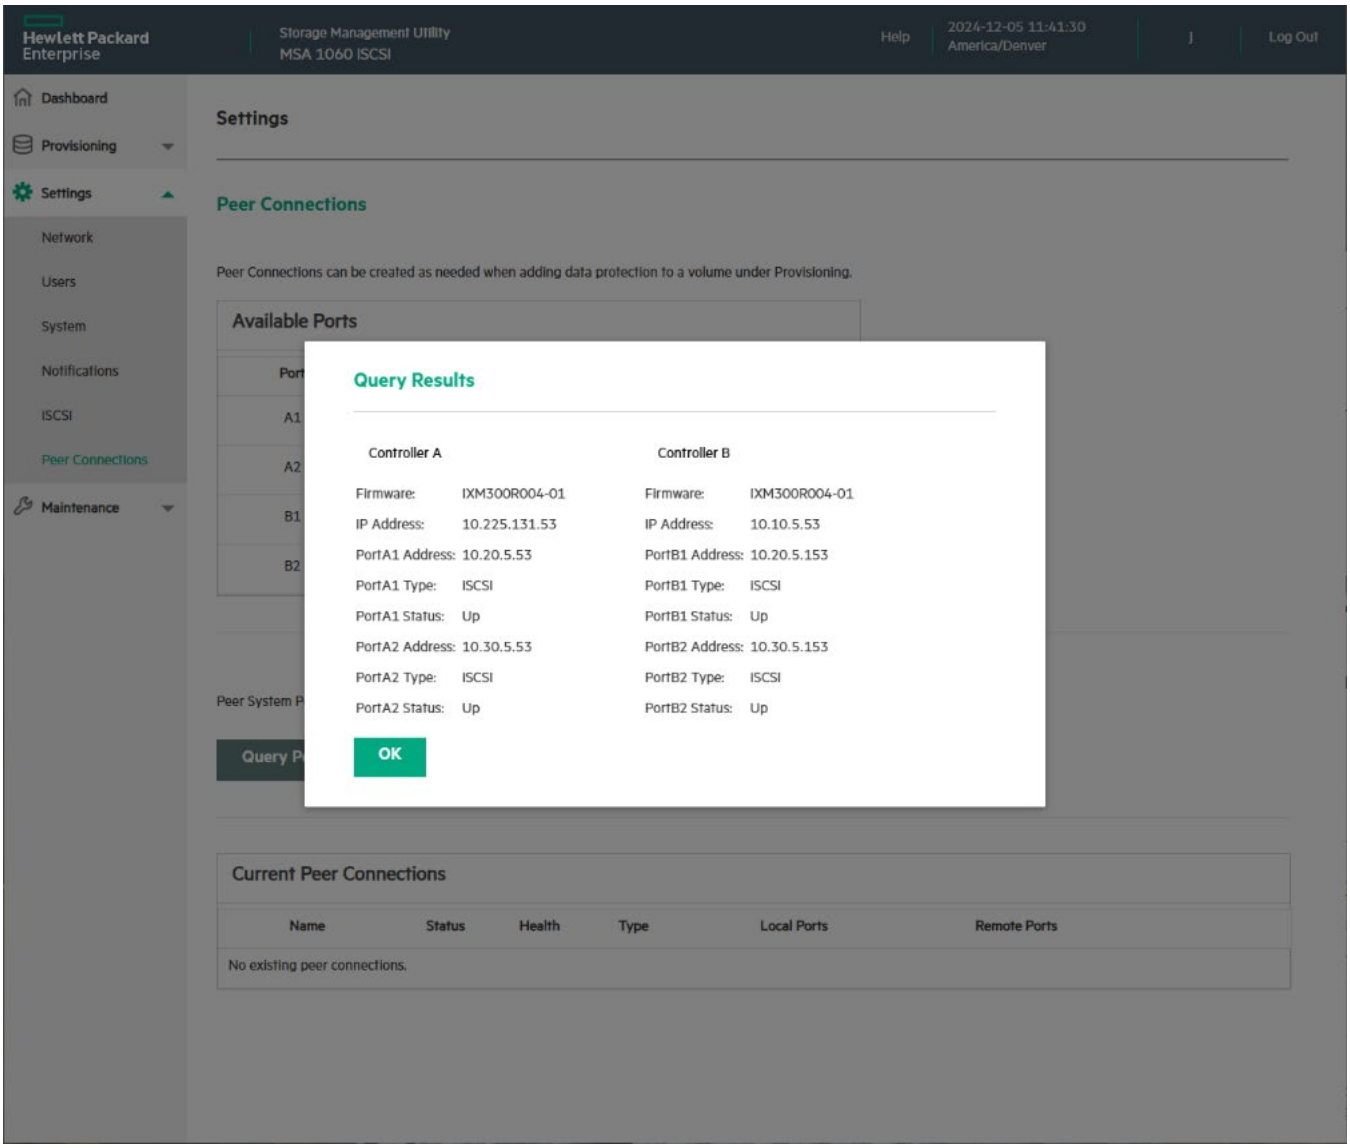


Figure 5. Peer connection query results

Note that while the ports on the remote array may be listed, actual connectivity between local and remote ports is not displayed—use the `query peer-connection` CLI command, or, if the peer connection already exists, the `show peer-connection` CLI command with the `verify-links` option to view port to port connectivity.

Creating a peer connection

Using the CLI

Use the `create peer-connection` CLI command (see Command Example 3. Creating the peer connection for an example that shows the remote-username parameter and remote password prompt):

```
# create peer-connection remote-port-address 10.20.5.53 remote-username manage Den-Chi
Info: This may take a few minutes to ping all port combinations...
Enter remote password: ***********
Re-enter password: ***********
Success: Command completed successfully.
```

Command Example 3. Creating the peer connection

Using the SMU

The SMU for HPE MSA Gen6 Storage and HPE MSA Gen7 Storage focuses on tasks, and since the peer connection is only needed for replication sets, and since a replication set requires a volume, the task of creating the peer connection is simply part of creating a replication set using an existing volume or group of volumes. See Using the SMU in Creating a replication set (page 34) for directions and examples for creating a replication set using the SMU and thereby creating a peer connection in the process.

Displaying peer connections and verifying links

Periodically check the status, health, and available ports of peer connections to achieve optimal replication performance. You can also verify the links used by the peer connections when troubleshooting or when changes have occurred in your environment.



Using the CLI

Use the `show peer-connections` CLI command to view the current status of the peer connections; use the `verify-links` option to update the currently available links. The following example (Command Example 4. Displaying peer connections and verifying links) shows two peer connections using the `verify-links` option – notice the message, "`Info: This may take a few minutes to ping all port combinations...`"; this indicates the availability of the links are being updated.

```
# show peer-connections verify-links
Info: This may take a few minutes to ping all port combinations...
Info: This may take a few minutes to ping all port combinations...
Peer Connections
----------------
Peer Connection Name: Den-Chi
Peer Connection Type: iSCSI
Connection Status: Online
Health: OK
Health Reason:
Health Recommendation:

  Local Port     Port Address                                  Reachable Remote Links
  ---------------------------------------------------------------------------------
  A1             10.20.5.53                                    A1,B1
  A2             10.30.5.53                                    A2,B2
  B1             10.20.5.153                                   A1,B1
  B2             10.30.5.153                                   A2,B2

  Remote Port    Port Address                                  Reachable Local Links
  ---------------------------------------------------------------------------------
  A1             10.20.5.50                                    A1,B1
  A2             10.30.5.50                                    A2,B2
  B1             10.20.5.150                                   A1,B1
  B2             10.30.5.150                                   A2,B2

Peer Connections
----------------
Peer Connection Name: Chi-Minn
Peer Connection Type: iSCSI
Connection Status: Online
Health: OK
Health Reason:
Health Recommendation:

  Local Port     Port Address                                  Reachable Remote Links
  ---------------------------------------------------------------------------------
  A1             10.20.5.53                                    A1,B1
  A2             10.30.5.53                                    A2,B2
  B1             10.20.5.153                                   A1,B1
  B2             10.30.5.153                                   A2,B2

  Remote Port    Port Address                                  Reachable Local Links
  ---------------------------------------------------------------------------------
  A1             10.20.5.52                                    A1,B1
  A2             10.30.5.52                                    A2,B2
  B1             10.20.5.152                                   A1,B1
  B2             10.30.5.152                                   A2,B2

Success: Command completed successfully.
```

Command Example 4. Displaying peer connections and verifying links



Using the SMU

To view the peer connections, use the **Peer Connections** panel (**Settings** > **Peer Connections**) (see Figure 6. Displaying peer connections), and view the **Status** and **Health** and port lists for the configured peer connections. The lists of **Local Ports** and **Remote Ports** are accurate as of the last time the peer connections were queried. Unlike the CLI command, the SMU only displays accessible ports, but not the actual links. Use the `show peer-connections` CLI command to show the port-to-port links available.

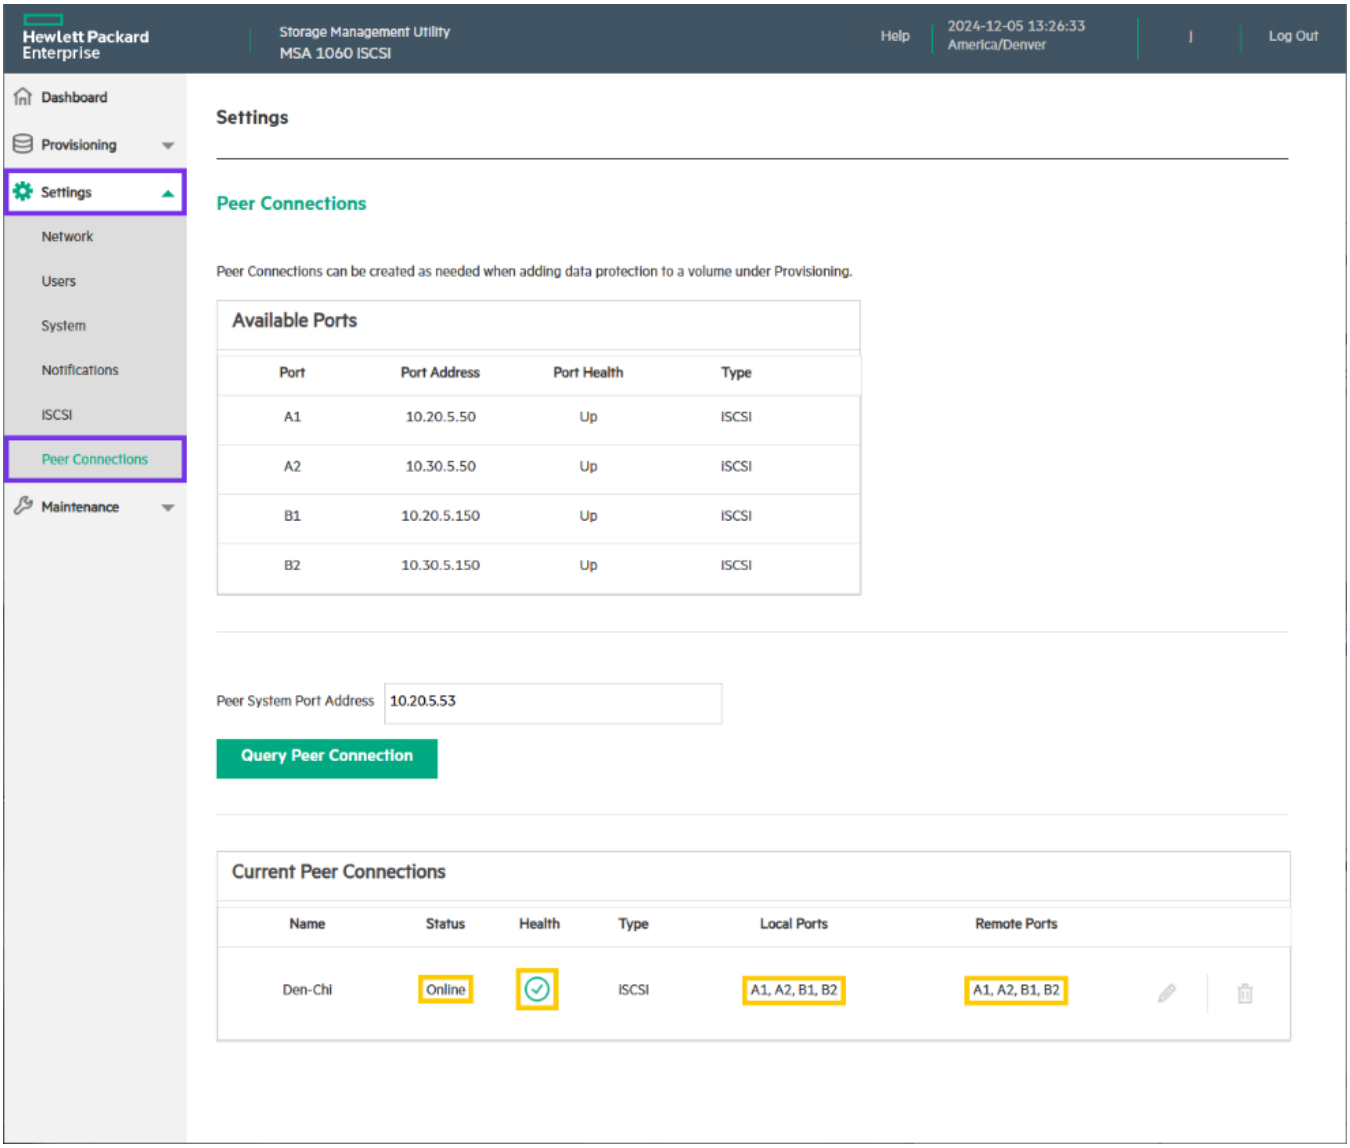


Figure 6. Displaying peer connections

Modifying peer connections

You can change the peer connection name or the peer connection's remote port address, but not both with the same operation. For example, if you move the remote array and change its IP address, you'll need to update the peer connection remote port address for replications to continue. If you change the host port addresses on both arrays, you only need set the peer connection's remote port address on one array. You must provide the name and password of a user on the remote array that has the *manage* or *standard* role. Since changing the peer connection remote port address modifies the connection, abort all replications in progress using the peer connection prior to changing the remote port address, and suspend all replication sets that use the peer connection, to avoid scheduled replications from attempting to start.

Using the CLI

Use the `set peer-connection` CLI command (see Command Example 5. Modifying a peer connection) to set either the name or the remote port address, but not both in the same command. The following example shows setting the remote port address after all port addresses on both arrays were changed – it took longer than expected, and so used an event to notify when the operation completed.

```
# set peer-connection remote-port-address 10.20.5.53 remote-username manage Den-Chi
Enter remote password: ***********
Re-enter password: ***********
Info: This may take a few minutes to ping all port combinations...
Warning: Command is taking longer to complete than anticipated. - The peer modification process has
started and may take several minutes to complete. When it is complete, event 584 will be logged.
View the event log to determine completion status.
# show events last 1
2024-12-05 12:20:14 [584]  #A1584:  MSA 1060 iSCSI Array SN#00C0FF44952C Controller A INFORMATIONAL
A peer connection was modified. (name: Den-Chi, SN: 00c0ff44952c000077e90e6701020000)
```

Command Example 5. Modifying a peer connection



Using the SMU

To modify a peer connection, use the **Peer Connections** panel (**Settings** > **Peer Connections**) (see Figure 7. Modifying a peer connection) and click the edit icon .

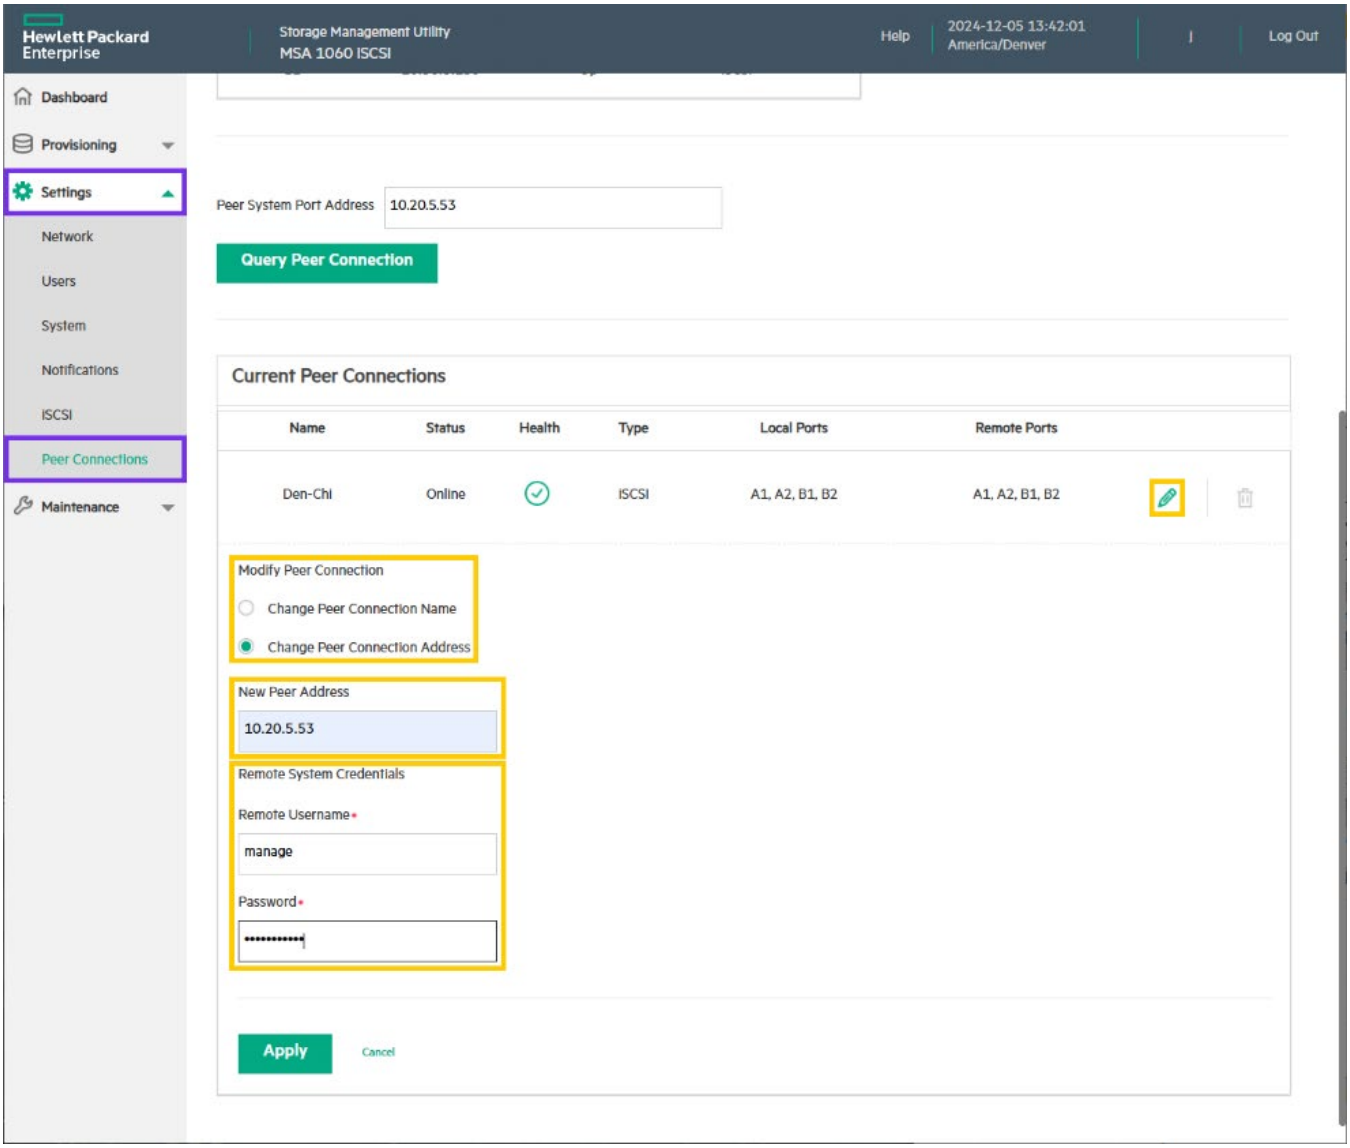


Figure 7. Modifying a peer connection

Deleting a peer connection

Deleting a peer connection is possible only when there are no replication sets using the peer connection. If no connection is available between the arrays, or if one of the arrays is not functional, you can use the `local-only` option of the `delete peer-connection` CLI command to delete the peer connection from the array you are logged into. If you've used the `local-only` option on one array to delete a peer connection, to delete the peer connection from the other array, you must use the `local-only` option on the second array as well, even if both arrays and the ports that were used in the connection are functioning. The `local-only` option is only available using the CLI, it is not available in the SMU.

Using the CLI

Use the `delete peer-connection` CLI command to delete the peer connection. See Command Example 6. Deleting a peer connection using the local-only option for an example that uses the `local-only` option.

```
# delete peer-connection local-only Den-Chi
WARNING: Do not use the local-only parameter unless other attempts to delete the peer connection
have failed. If you decide to run this command, make sure to run it on both primary and secondary
systems.
Do you want to continue? (y/n) y
Success: Command completed successfully.
```

Command Example 6. Deleting a peer connection using the local-only option

Using the SMU

Use the **Peer Connections** panel (**Settings** > **Peer Connections**) (see Figure 8. Deleting a peer connection) and click the delete icon for the peer connection you want to delete.

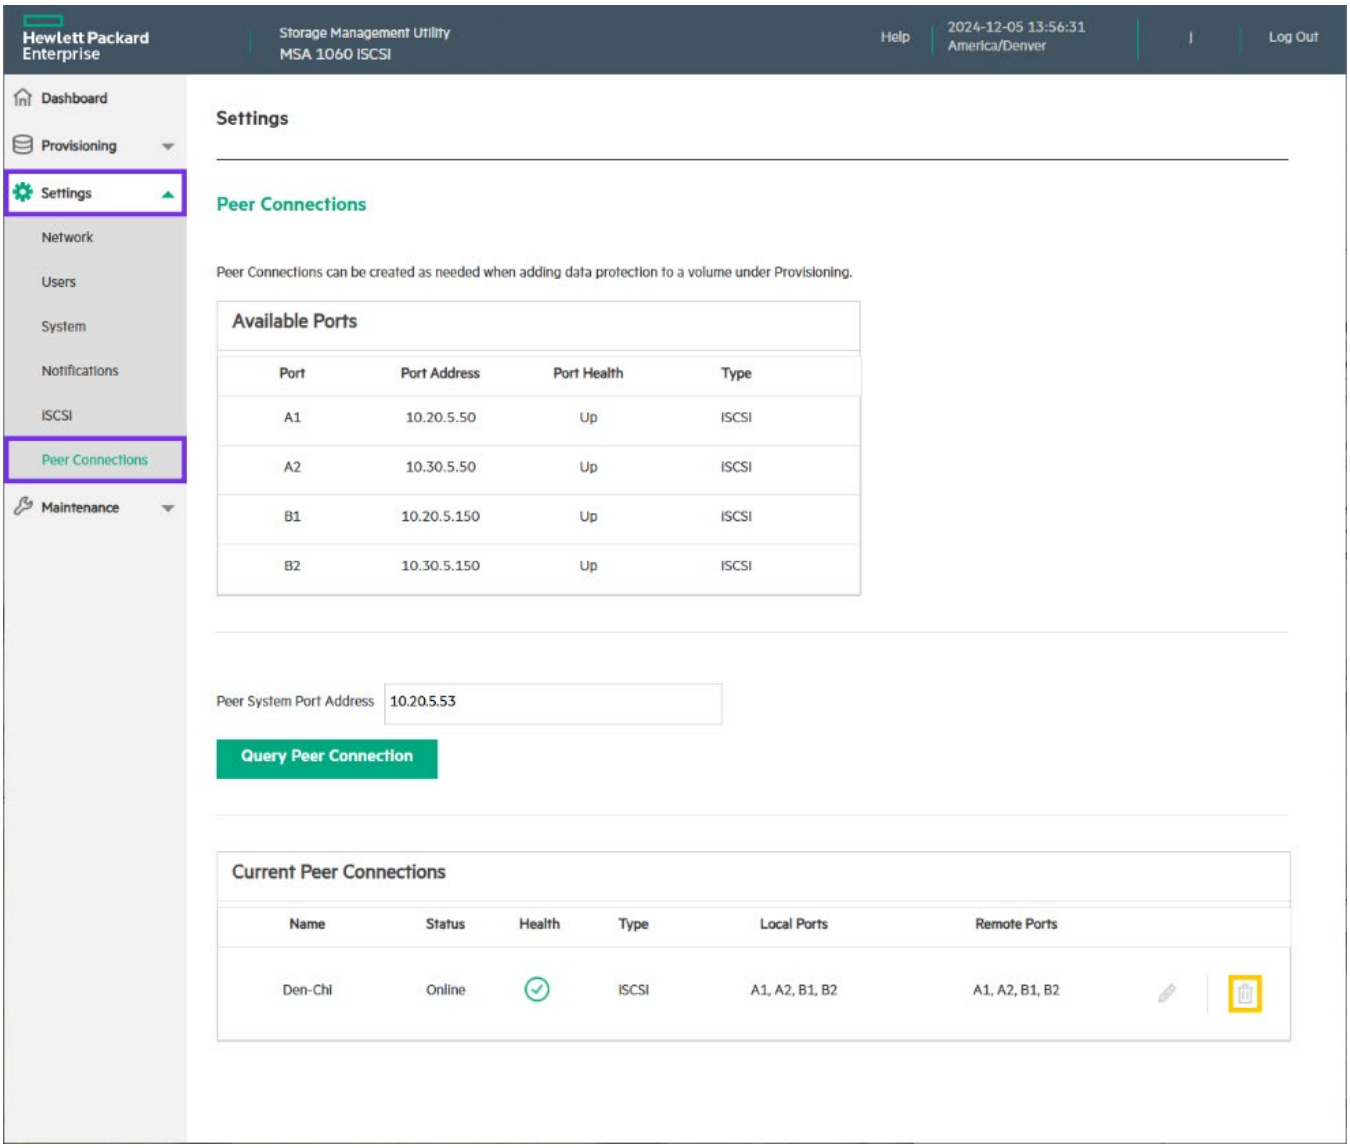


Figure 8. Deleting a peer connection

Peer connections and CHAP

The Challenge Handshake Authentication Protocol (CHAP) allows authentication between the initiator and the target of a login request. CHAP records specify either a single name-secret pair to authenticate the initiator only (one-way CHAP) or two name-secret pairs to authenticate both the initiator and the target (mutual CHAP). Remote Snap supports one-way CHAP only. The initiator supplies the secret from the record corresponding to the target's name to the target, while the target compares the secret supplied against the secret from the record corresponding to the initiator's name. Since both arrays will act as initiator or target, depending on the operation, use the same secret on both arrays.

If only one array has CHAP enabled, but both arrays have CHAP records for each other, the array with CHAP enabled can modify the peer connection but is unable to perform any other replication operations, such as creating or deleting replication sets, initiating replications, recovering a replication set, or suspending or resuming replications; the array without CHAP enabled is unable to perform any replication operations.

If the two arrays have CHAP records for each other with the same secret, they can perform all replication operations whether or not CHAP is enabled on either array. This is because they have the secret to supply or compare against in their CHAP records.

Note

If you enable CHAP, you enable it for all iSCSI connections, not just for the peer connection; you'll have to configure CHAP on all hosts that will access the array.

Using the CLI

First, create the CHAP record for the peer. Then, enable CHAP. To complete the process, reset the host port links to reset any existing connections and enforce the use of CHAP on subsequent logins. See Command Example 7. Setting up CHAP for an example of setting up CHAP.

```
# set chap-record name iqn.2015-11.com.hpe:storage.msa2070.baksz2402fb521e secret MSA-Remote-Snap
Info: CHAP record changes will not affect existing host port links until they are reset.
Success: Command completed successfully. (iqn.2015-11.com.hpe:storage.msa2070.baksz2402fb521e) - A
CHAP record was created.

# set iscsi-parameters chap enabled
Applying new parameters may disrupt access from connected hosts.
This command may take several minutes to complete.
Do you want to continue? (y/n) y
Success: Command completed successfully.

# reset host-link ports A1-2,B1-2
Resetting host link may cause lost connection to host.
Do you want to continue? (y/n) y
Success: Command completed successfully. - Reset host link(s) on requested port(s) from current
controller.
```

Command Example 7. Setting up CHAP



Using the SMU

Click **Add CHAP Record** on the **CHAP** panel (**Settings** > **iSCSI** > **CHAP**) (see Figure 9. Add CHAP record – initial step) to initiate adding a CHAP record.

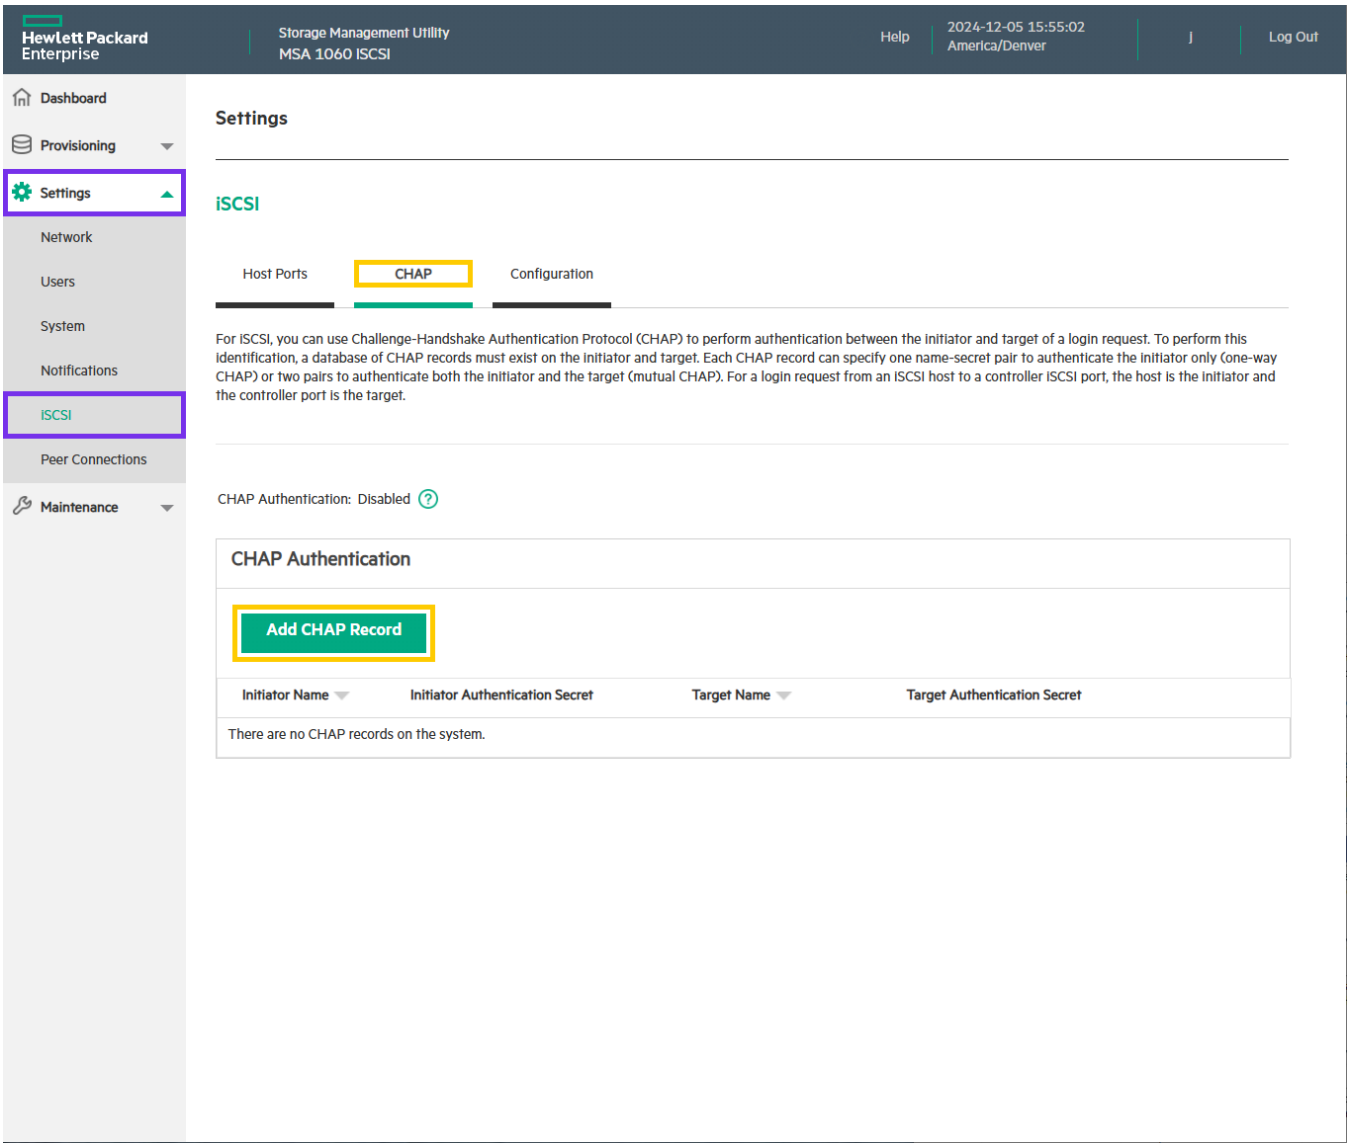


Figure 9. Add CHAP record – initial step

On the ensuing panel (see Figure 10. Add CHAP record – create the record), enter the **Initiator Name** and **Initiator Authentication Secret**, and click **Create CHAP Record** to complete creating the CHAP record.

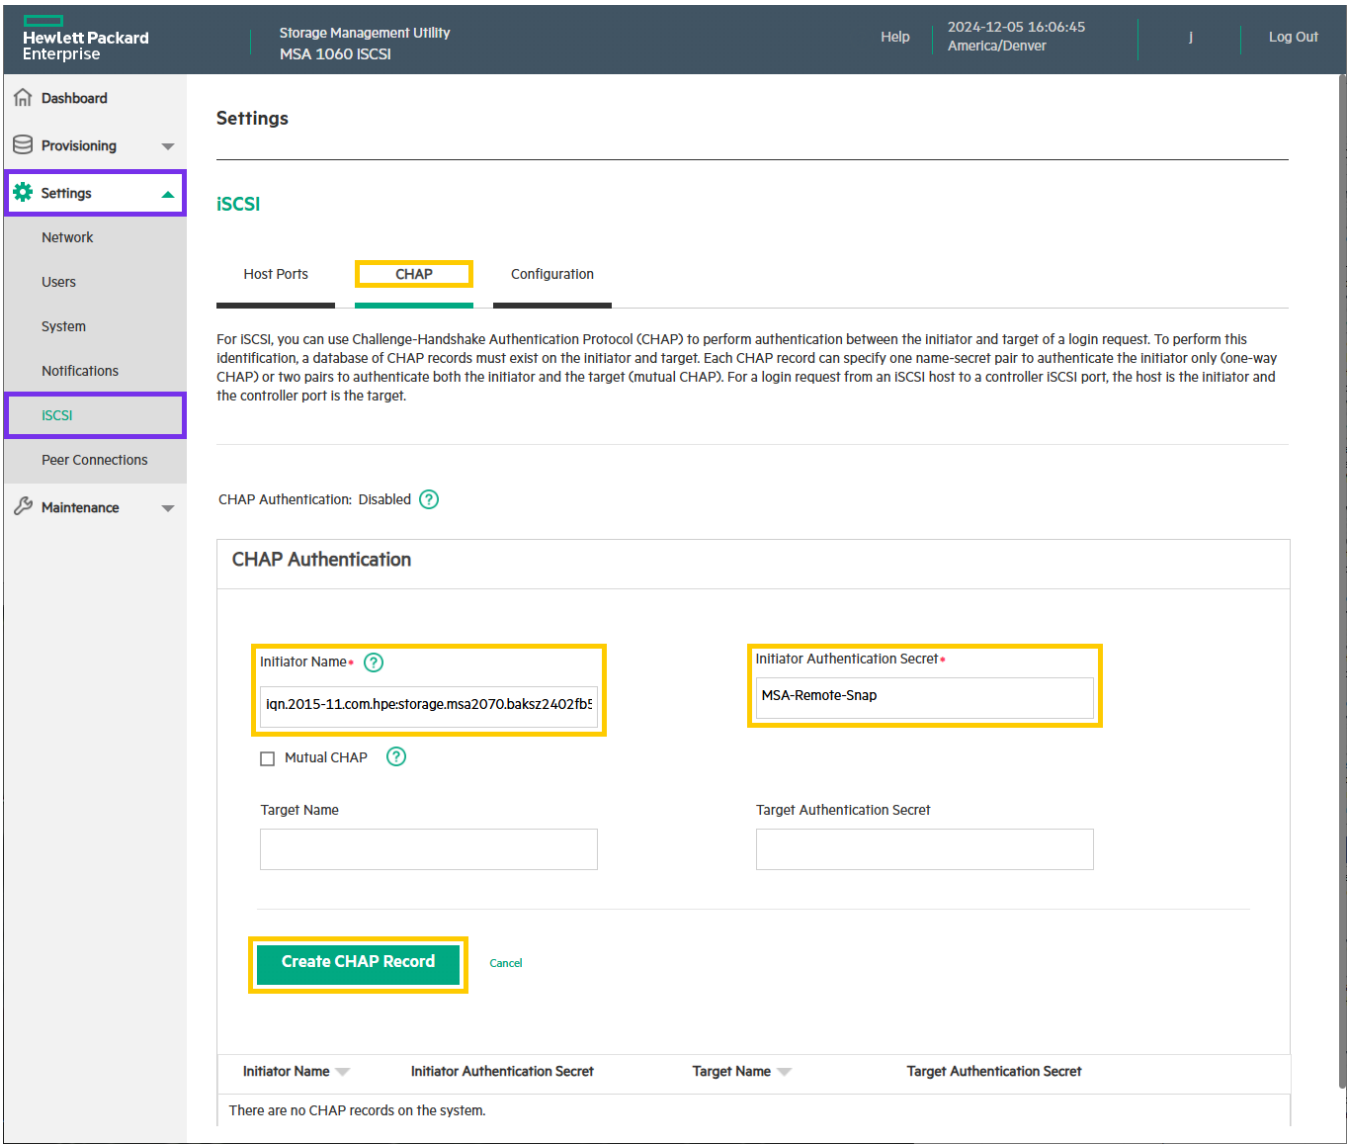


Figure 10. Add CHAP record – create the record

Use the **Configuration** tab of the **iSCSI** panel to enable **CHAP Authentication** and click **Set iSCSI Network**. See Figure 11. Enable CHAP for an example of enabling CHAP.

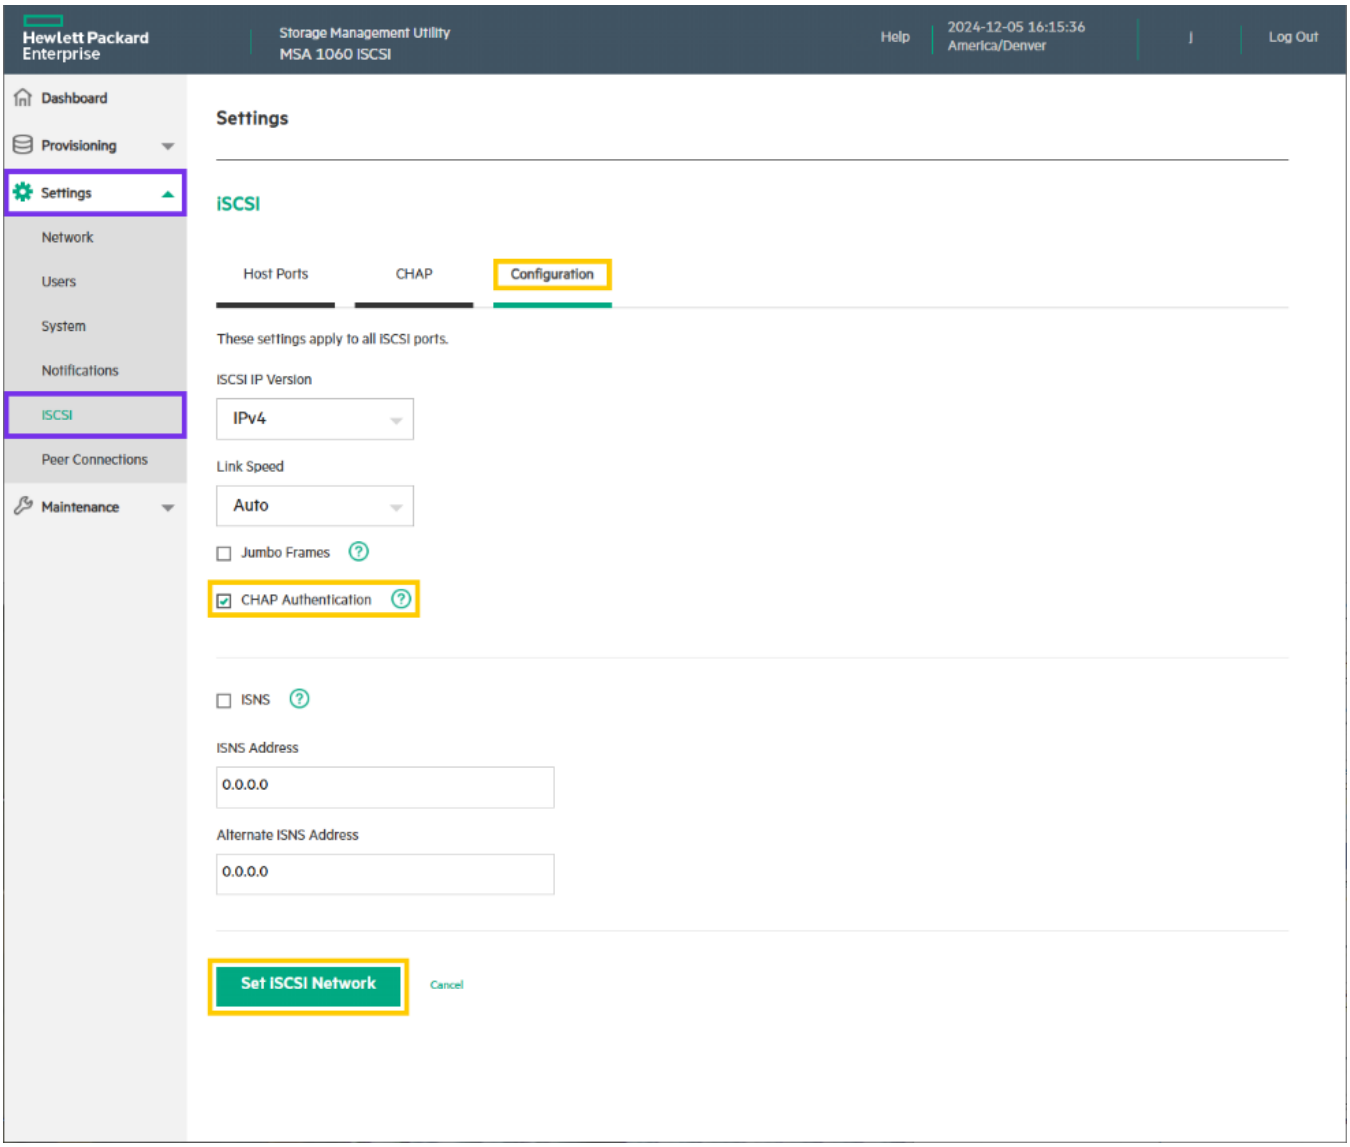


Figure 11. Enable CHAP

Finally, click **Reset Host Links** of the **Host Ports** tab of the **iSCSI** panel (**Settings** > **iSCSI** > **Host Ports**) (see Figure 12. Reset host links) to reset host links in order to reset the iSCSI sessions and enforce the CHAP configuration on subsequent logins. Click **OK** on the ensuing confirmation dialog.

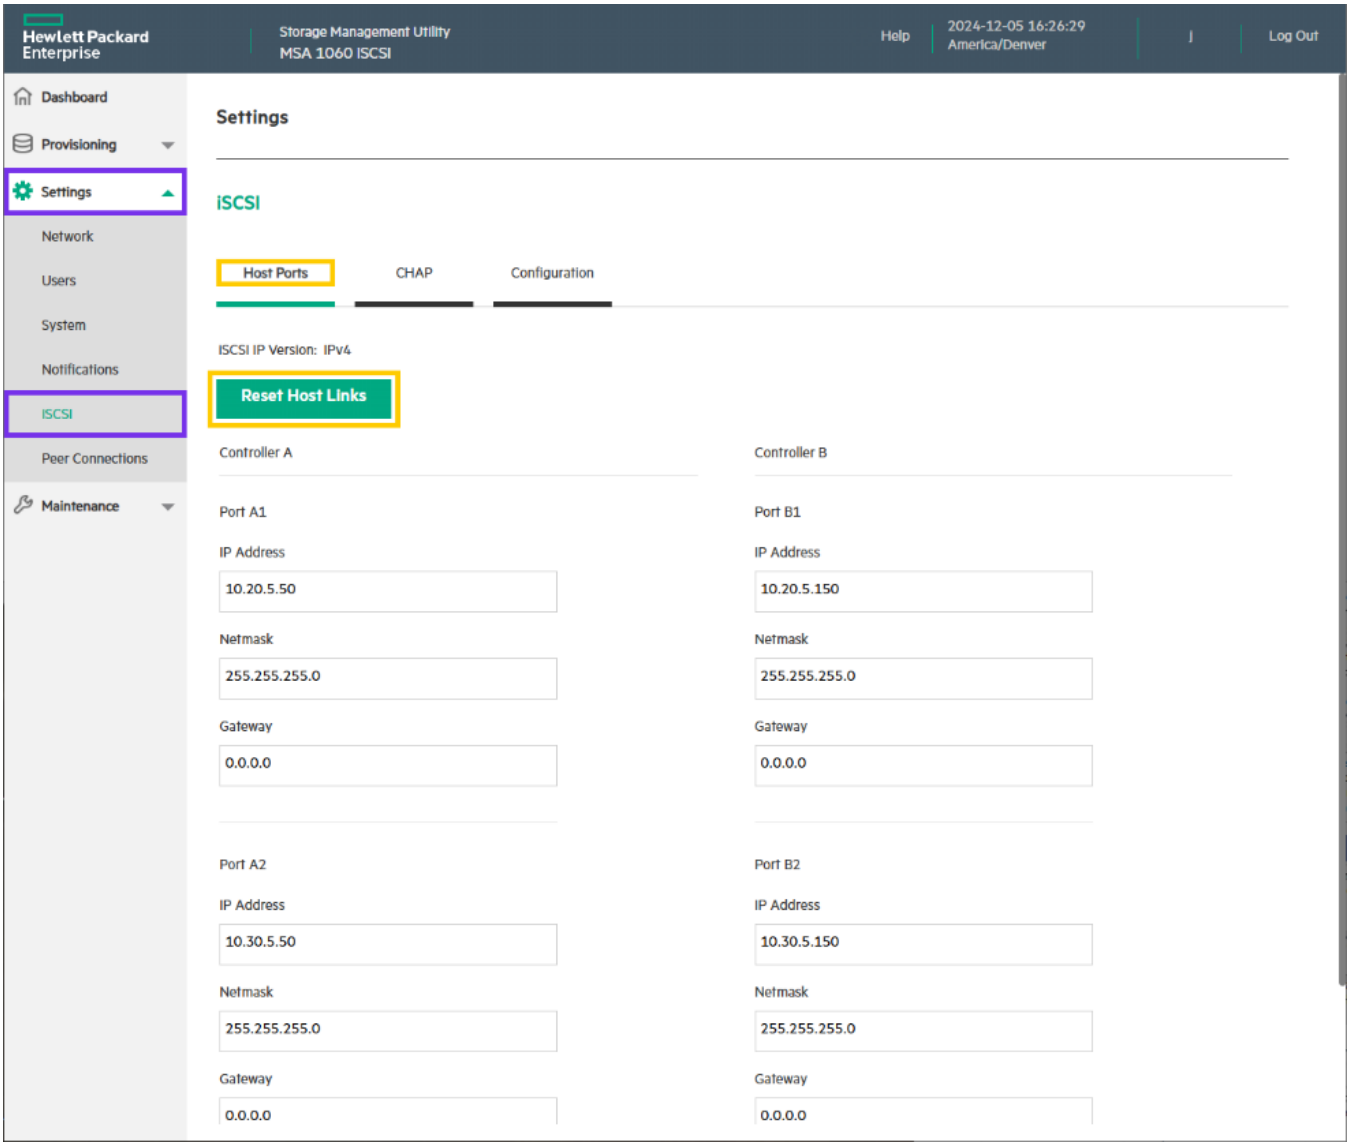


Figure 12. Reset host links

Replication set and replication operations

Creating a replication set

When creating a replication set using a primary volume or volume group, the secondary volume or volume group is created automatically; you cannot create a secondary volume or volume group manually. If the secondary pool name does not match the primary pool name, specify the secondary pool name manually.

Note

If using the SMU, creating the peer connection is part of creating the replication set, and will be covered in the Using the SMU in Creating a replication set (page 36).

Using the CLI

Use the `create replication-set` command (see Command Example 8. Creating a replication set) and specify the peer-connection, primary volume or volume group and replication set name, at a minimum. You may also choose the name of the remote pool (A or B – the default is the pool the primary belongs to) and change the remote volume name if replicating a volume rather than a volume group (the default is the primary volume name). You can also specify the queue policy (the default is to queue the latest replication snapshot if a replication is in progress) and replication snapshot history options (the default is that replication snapshot history is *Disabled*).

```
# create replication-set peer-connection Den-Chi primary-volume Source secondary-pool B secondary-volume-name Replica queue-policy queue-latest snapshot-history both snapshot-basename Replicated snapshot-count 4 Replicator
```

Command Example 8. Creating a replication set



Using the SMU

Use the **Data Protection Configuration** wizard to create the replication set. There are three paths to the wizard in the SMU. They all start with the **Volumes** panel (**Provisioning** > **Volumes**).

The first path to the wizard is to select one or more volumes, then select **Add Data Protection** from the **Select Volume Action** menu – see Figure 13. Accessing the Data Protection Configuration wizard (select volumes > Select Volume Action > Add Data Protection). Note that activating the **Select Volume Action** menu changes its background color and may no longer display its label; Figure 13 has been modified to show the label.

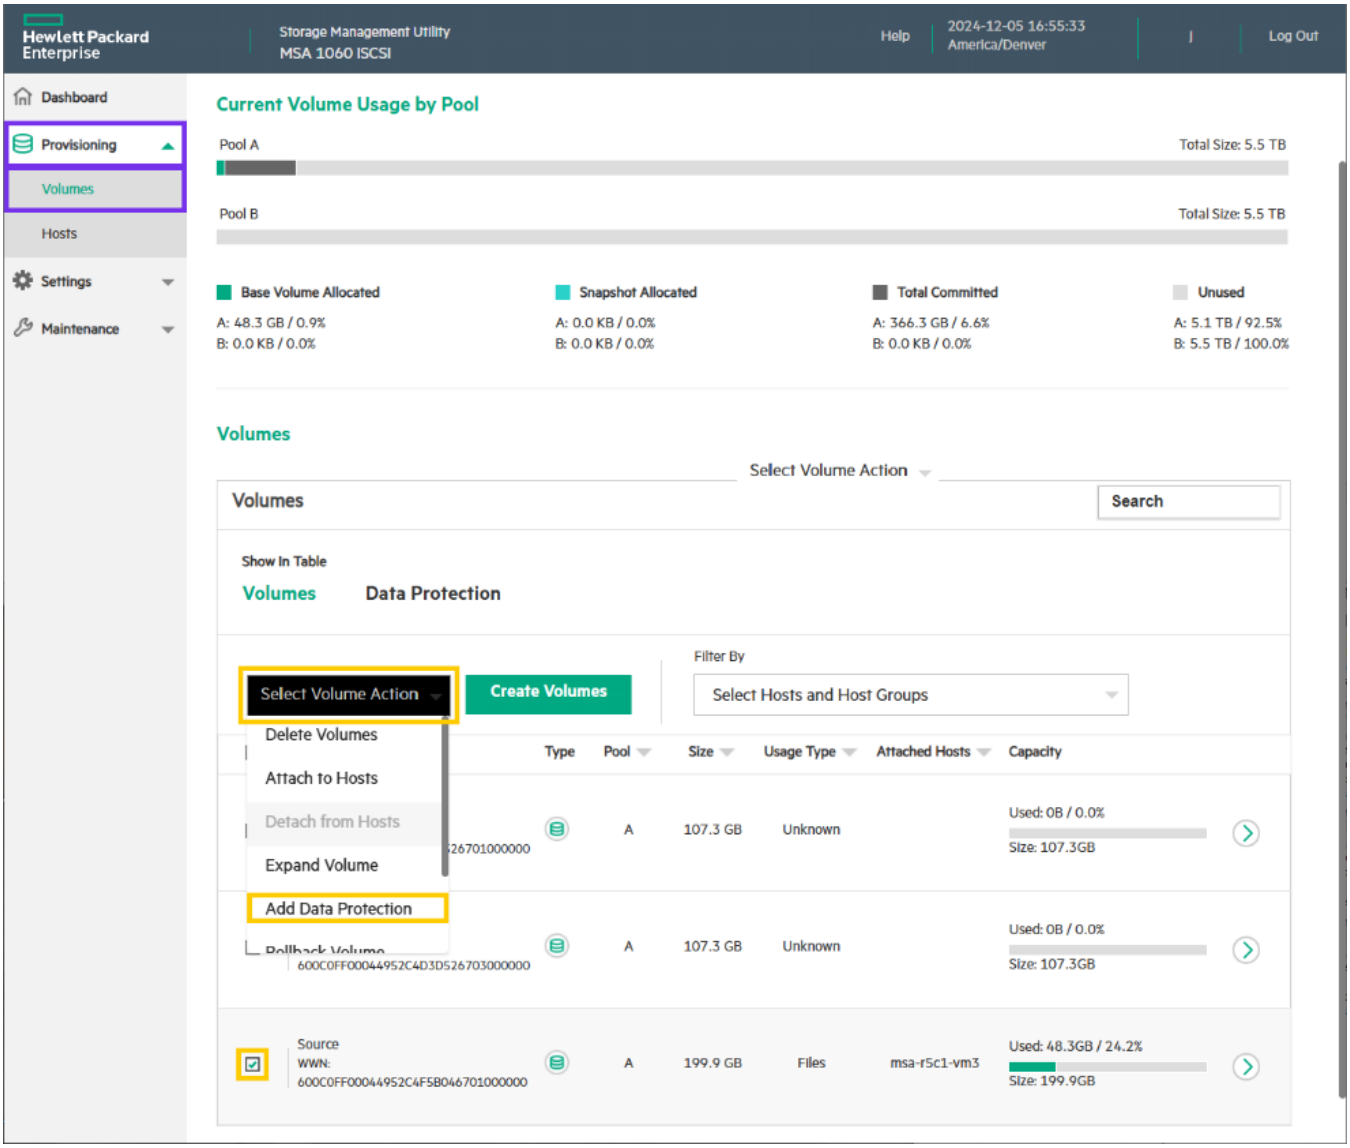


Figure 13. Accessing the Data Protection Configuration wizard (select volumes > Select Volume Action > Add Data Protection)

The second path to the wizard is to click the slide-over for a volume – see Figure 14. Accessing the Data Protection Configuration wizard (volume slide-over) – click the **Replications** tab, then click **Add Data Protection** – see Figure 15. Accessing the Data Protection Configuration wizard (Replications > Add Data Protection).

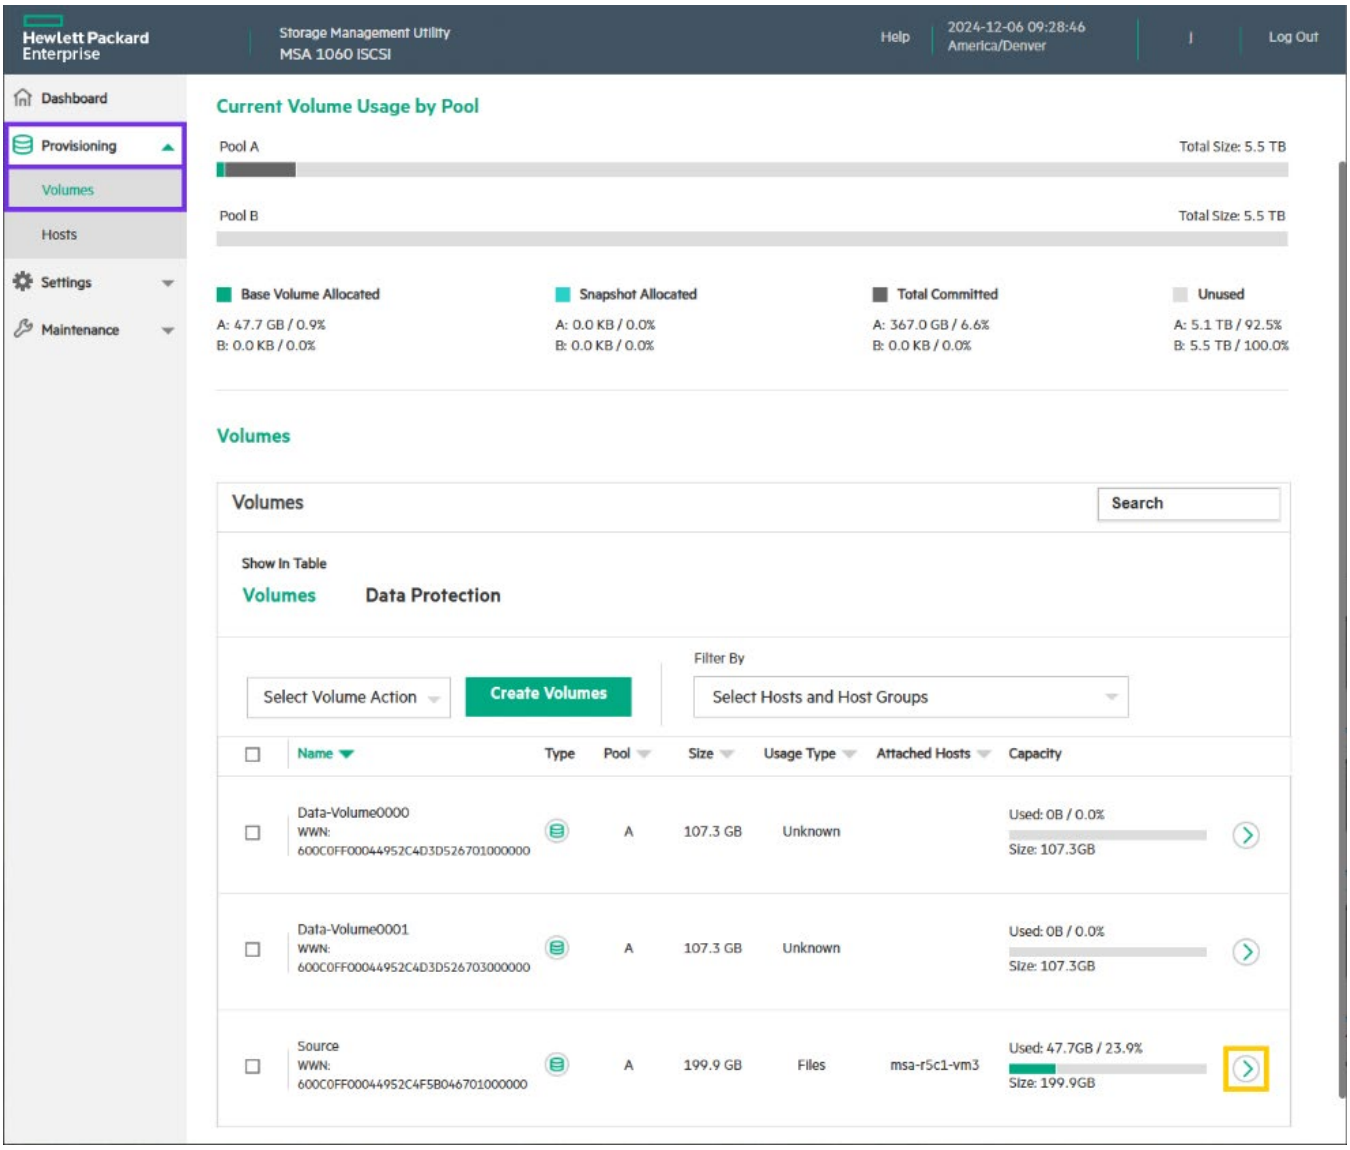


Figure 14. Accessing the Data Protection Configuration wizard (volume slide-over)

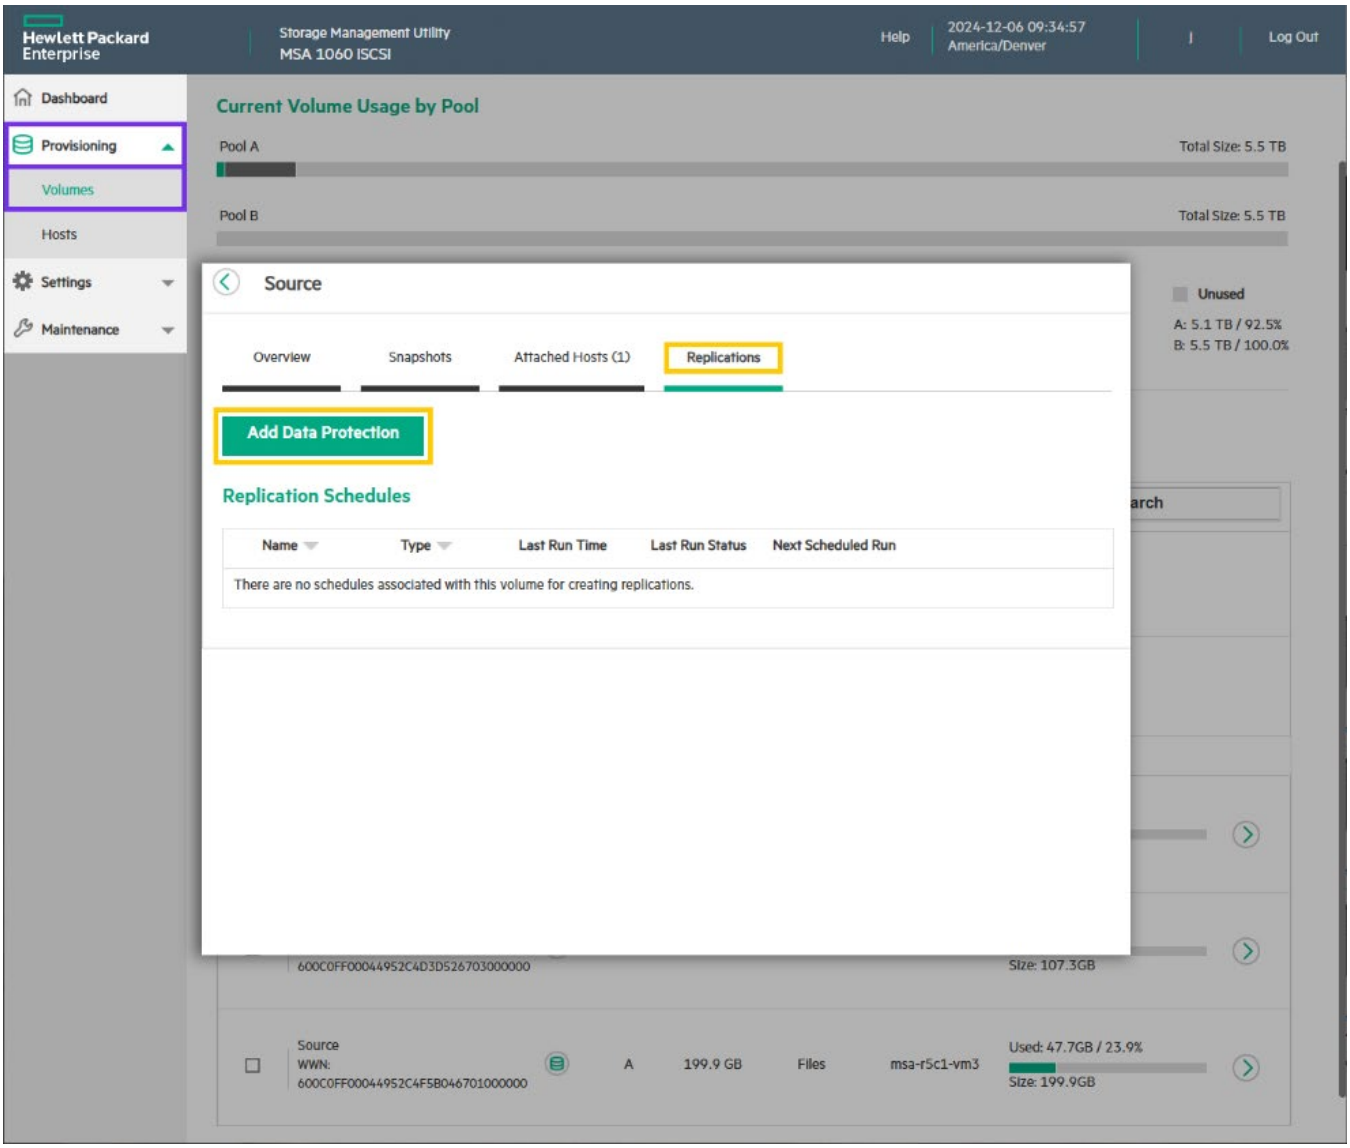


Figure 15. Accessing the Data Protection Configuration wizard (Replications > Add Data Protection)

The third path to the wizard is to select the **Data Protection** view of the table, and then click **Add Data Protection** – see Figure 16. Accessing the Data Protection Configuration wizard (Data Protection table view > Add Data Protection).

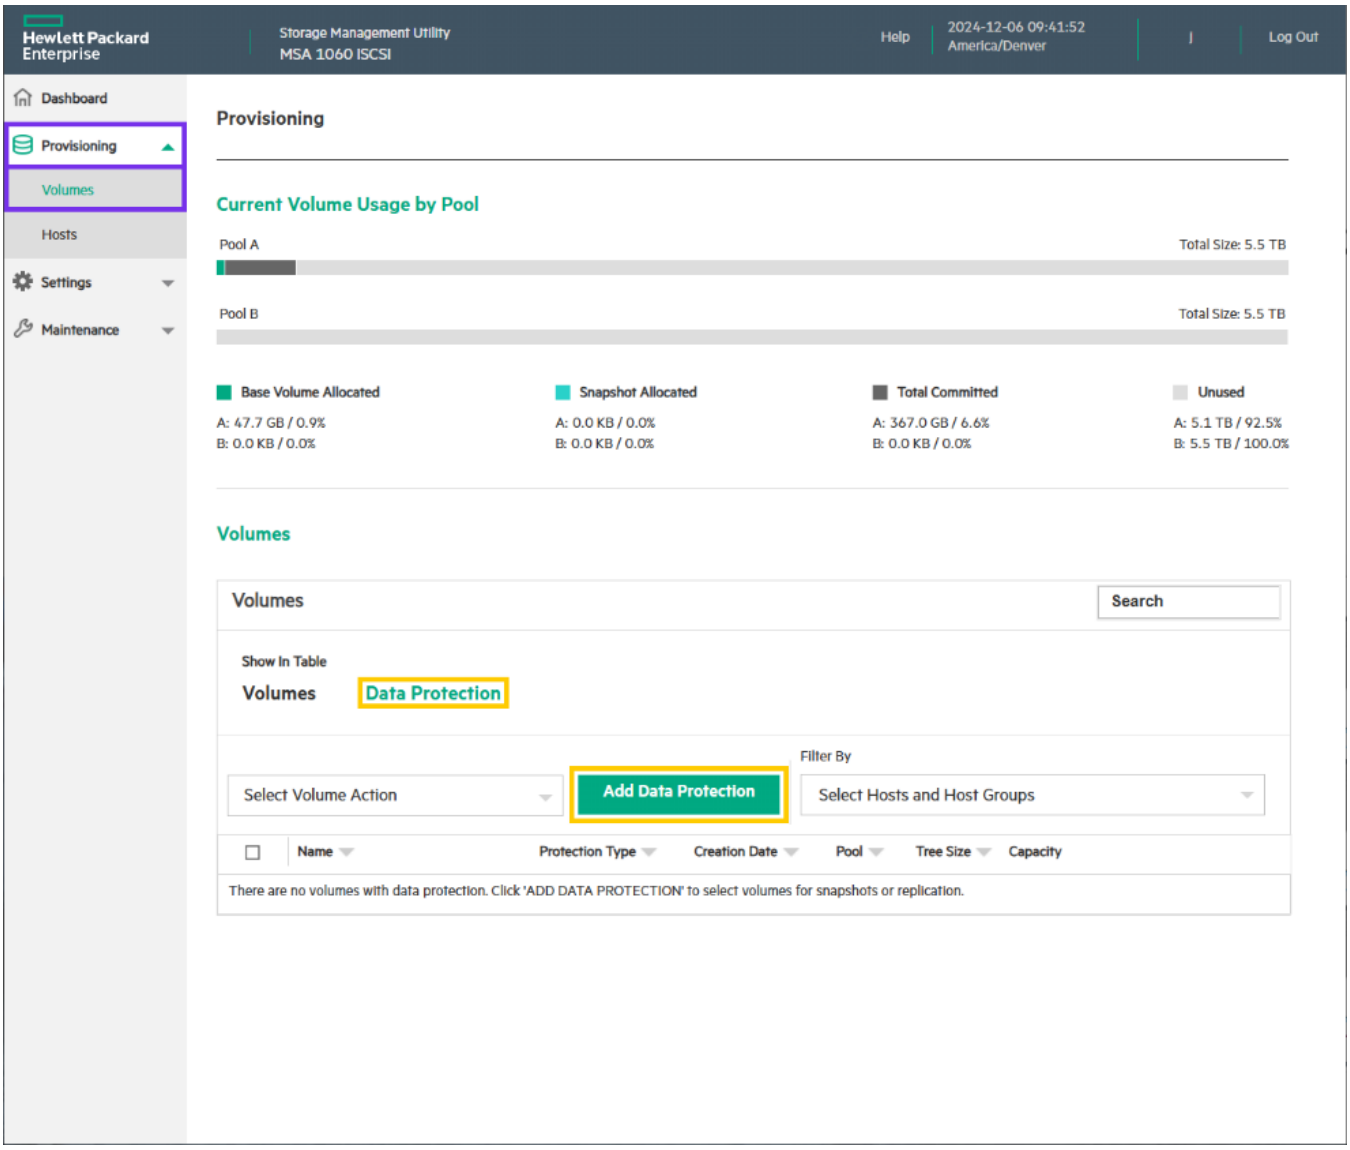


Figure 16. Accessing the Data Protection Configuration wizard (Data Protection table view > Add Data Protection)

Once the **Data Protection Configuration** wizard is accessed, the first step is to choose the data protection type and the volumes to include. Depending on the path taken to get to the wizard, the choice of data protection type and volumes may already be selected and cannot be changed as no other options are available.

Figure 17. Data Protection Configuration wizard – Volumes – Data Protection Type (Remote Replication) and volume selection shows both data protection type and volume selection options available; select **Remote Replication** (if offered) and select volumes to replicate (if offered) then continue to **Connection**.

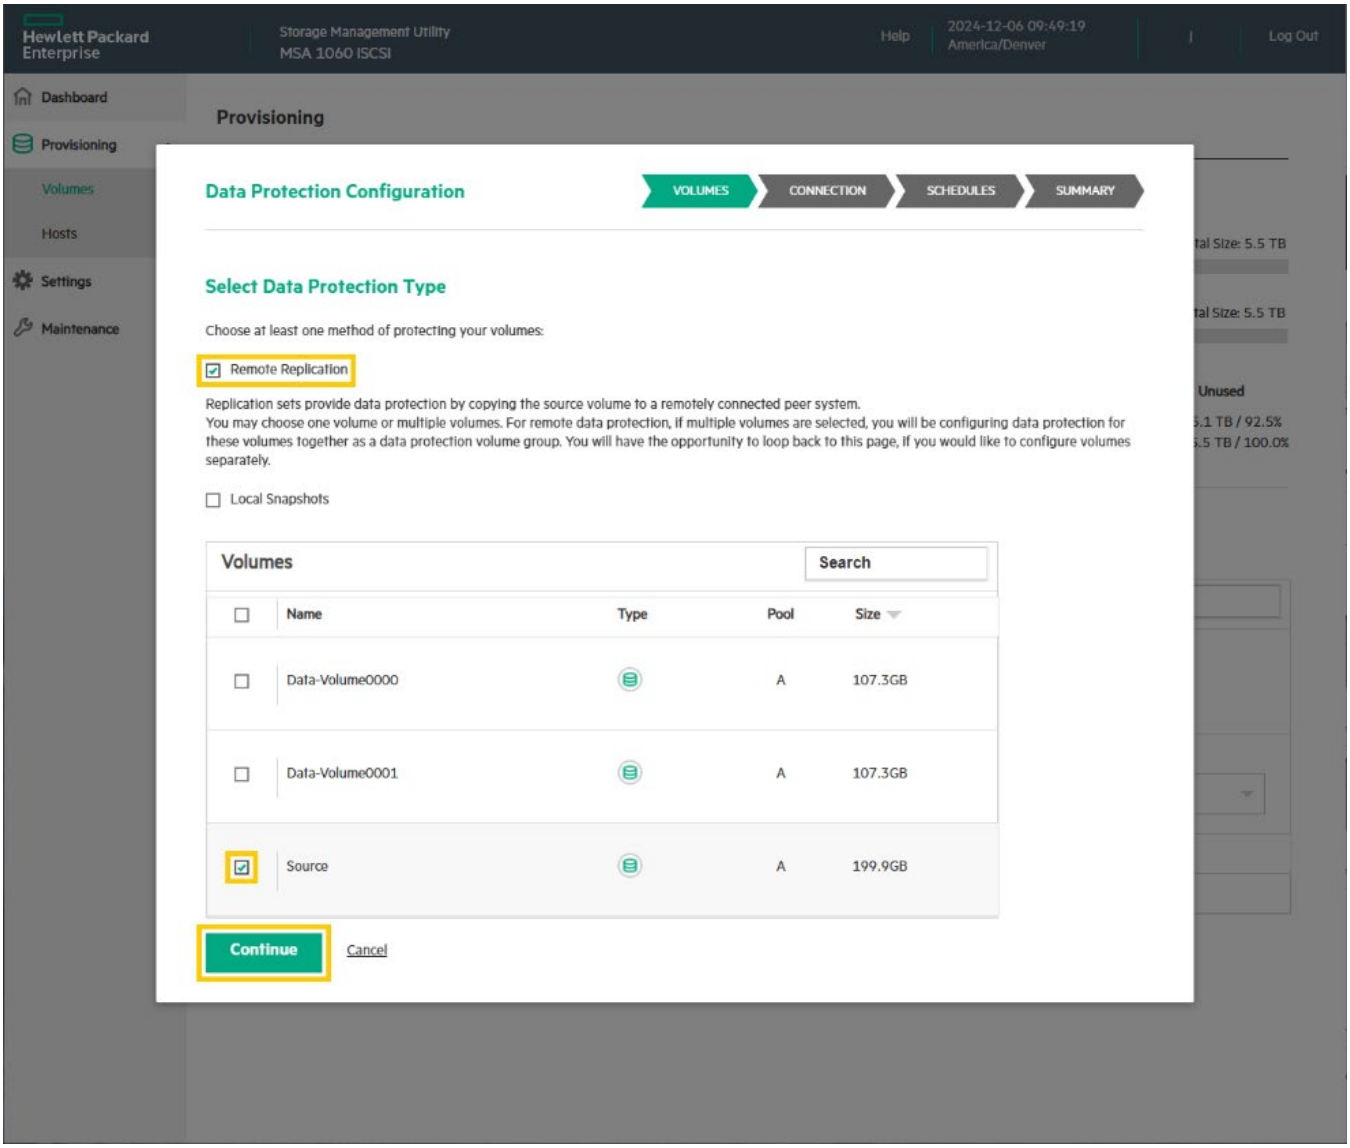


Figure 17. Data Protection Configuration wizard – Volumes – Data Protection Type (Remote Replication) and volume selection

Then select an existing peer connection or create a new peer connection (see Figure 18. Data Protection Configuration wizard – Connection – Data Protection Connection – existing or new).

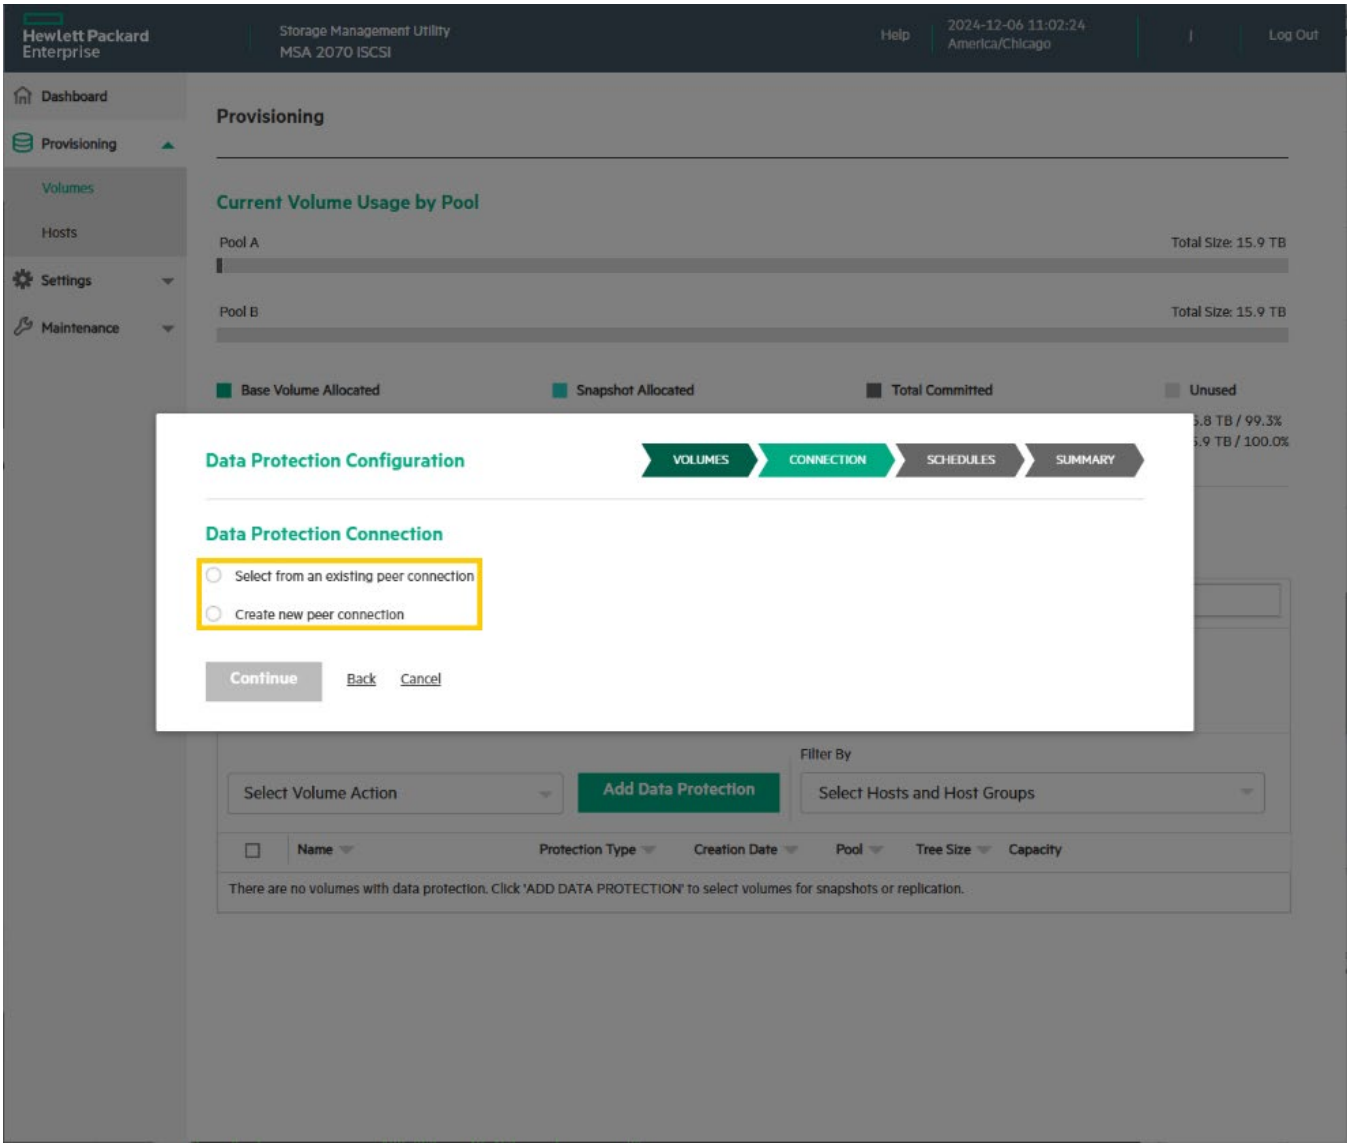


Figure 18. Data Protection Configuration wizard – Connection – Data Protection Connection – existing or new

If a peer connection does not exist, instead of the previous dialog, the wizard advances to the following dialog to initiate creating a peer connection.

If you chose to create a new peer connection, or the choice was pre-determined due to no existing peer connections, then use the following dialog to enter the port IP address or WWN of a host port on the remote array and click **Connect** (see Figure 19. Data Protection Configuration wizard – Connection – Data Protection Connection – new – specify the Port Address).

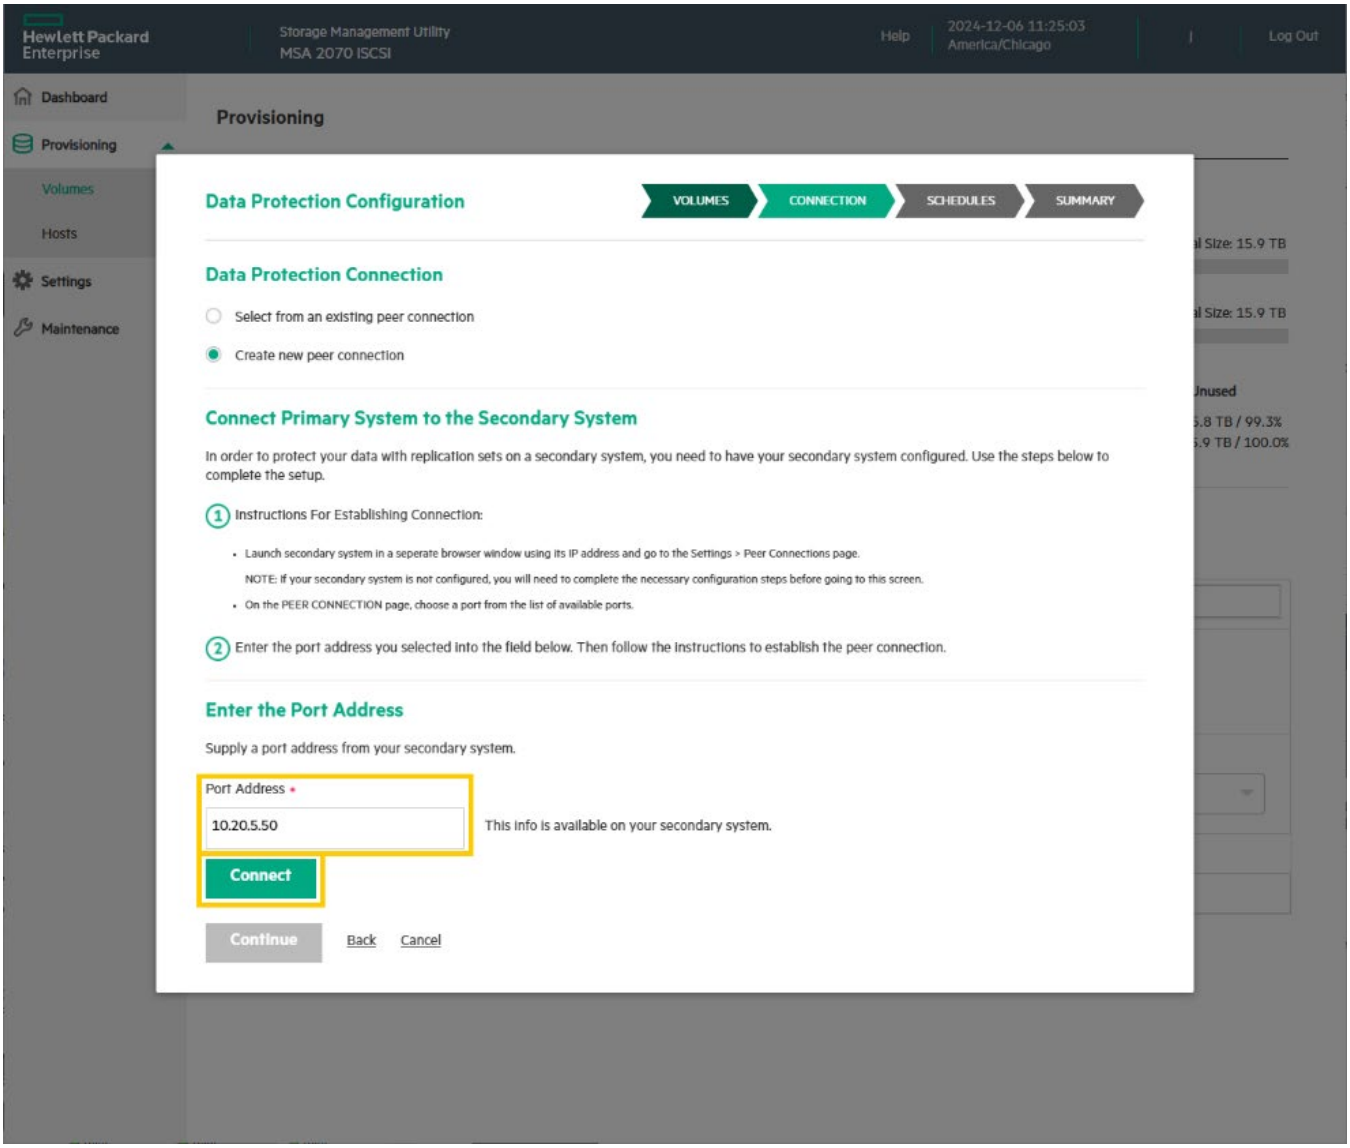


Figure 19. Data Protection Configuration wizard – Connection – Data Protection Connection – new – specify the Port Address

Once the remote array has responded, and the peer connection is viable, use the following dialog to create the peer connection, specifying the peer connection name, the name of a user on the remote array with the management role, and the password for the user, then click **Create Peer Connection** (see Figure 20. Data Protection Configuration wizard – Connection – Data Protection Connection – new – specify peer connection name and remote credentials).

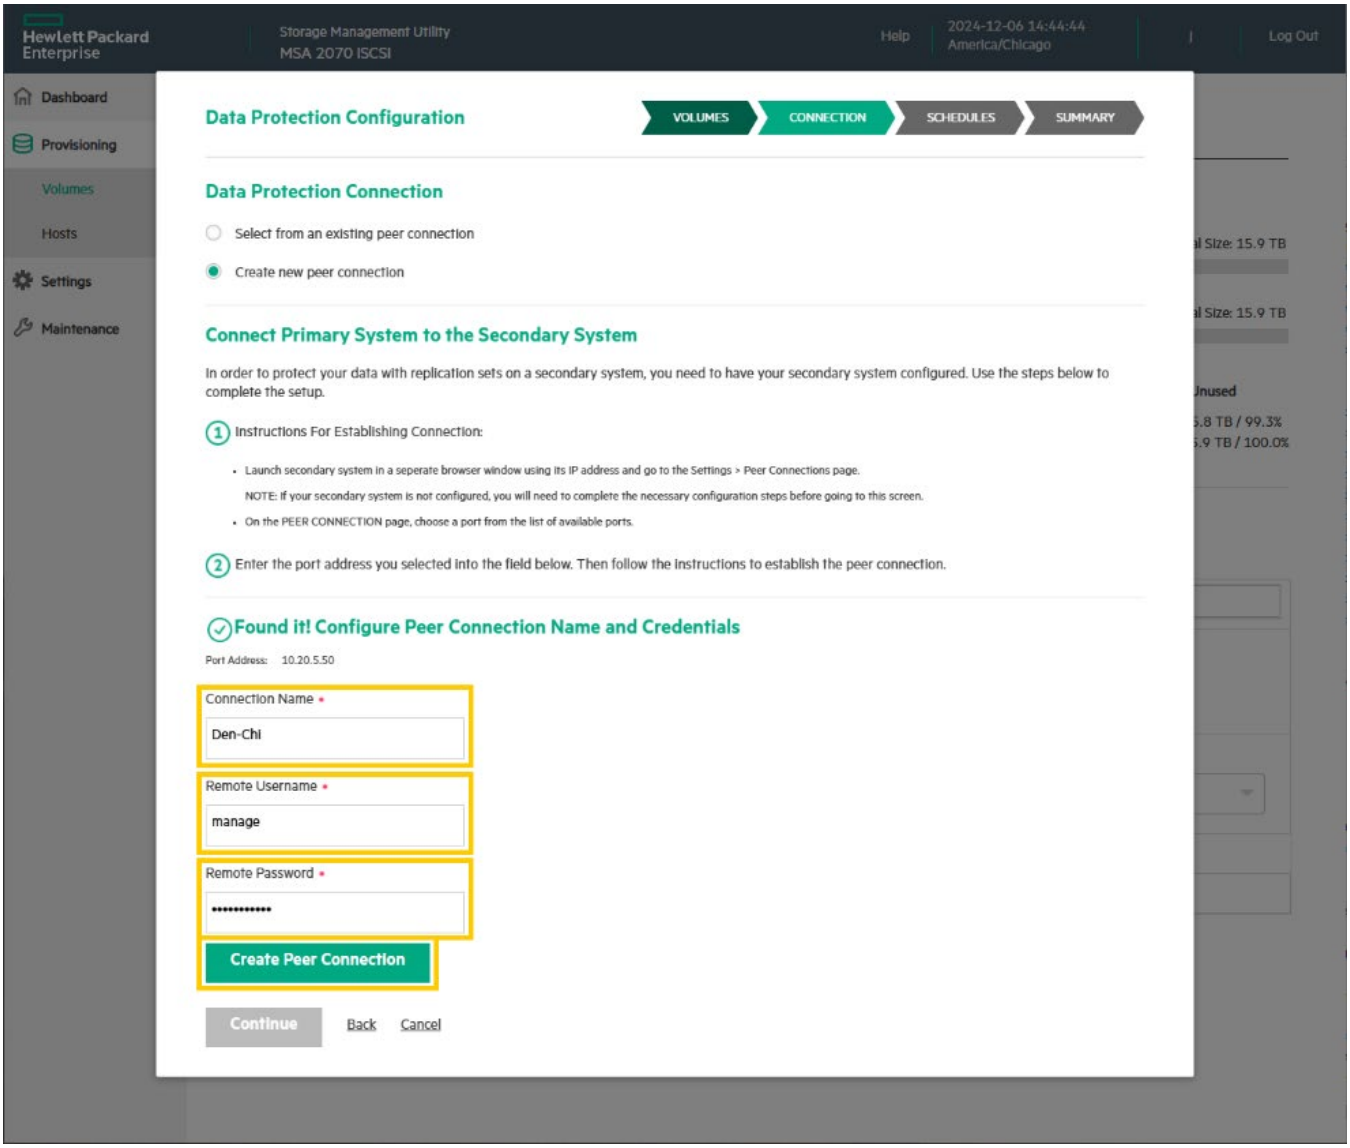


Figure 20. Data Protection Configuration wizard – Connection – Data Protection Connection – new – specify peer connection name and remote credentials

When the peer connection is created, the completion dialog shows the details – the name of the peer connection and the remote port address or WWN specified during the creation. Click **Continue** (see Figure 21. Data Protection Configuration wizard – Connection – Data Protection Connection – new – complete).

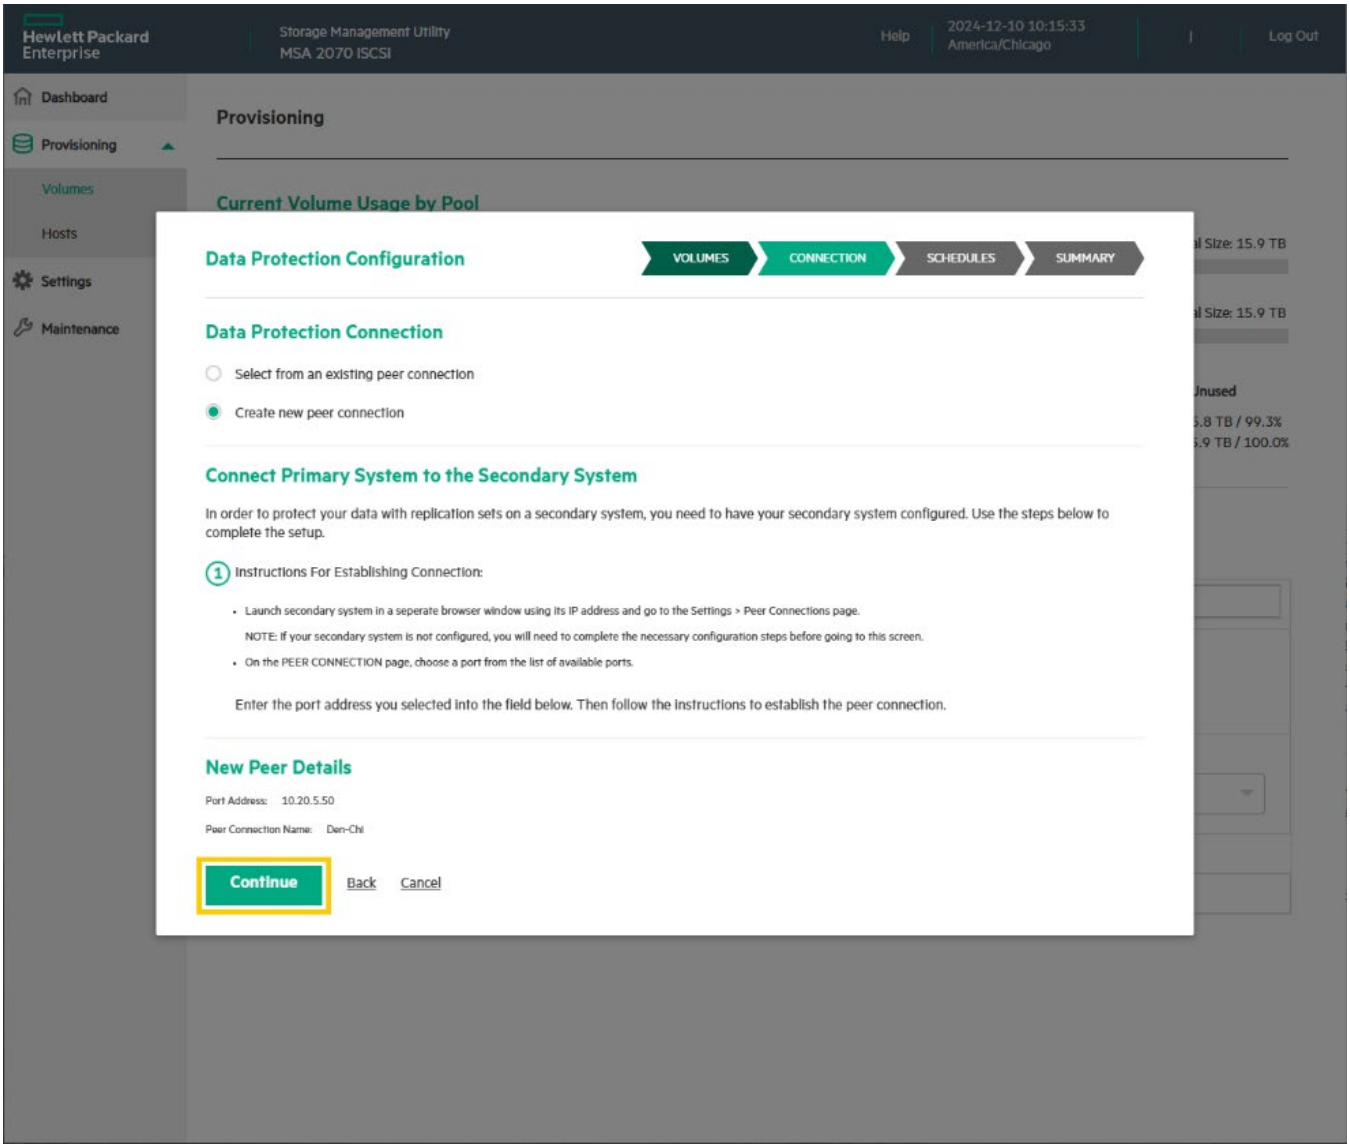


Figure 21. Data Protection Configuration wizard – Connection – Data Protection Connection – new – complete

Use the following dialog (see Figure 22. Data Protection Configuration wizard – Connection – Data Protection Connection – choose existing) if using an existing peer connection; select the peer connection to use from the **Choose Peer** drop-down list, then click **Continue**. Note that activating the **Choose Peer** menu changes its background color and may no longer display its label; Figure 22 has been modified to show the label.

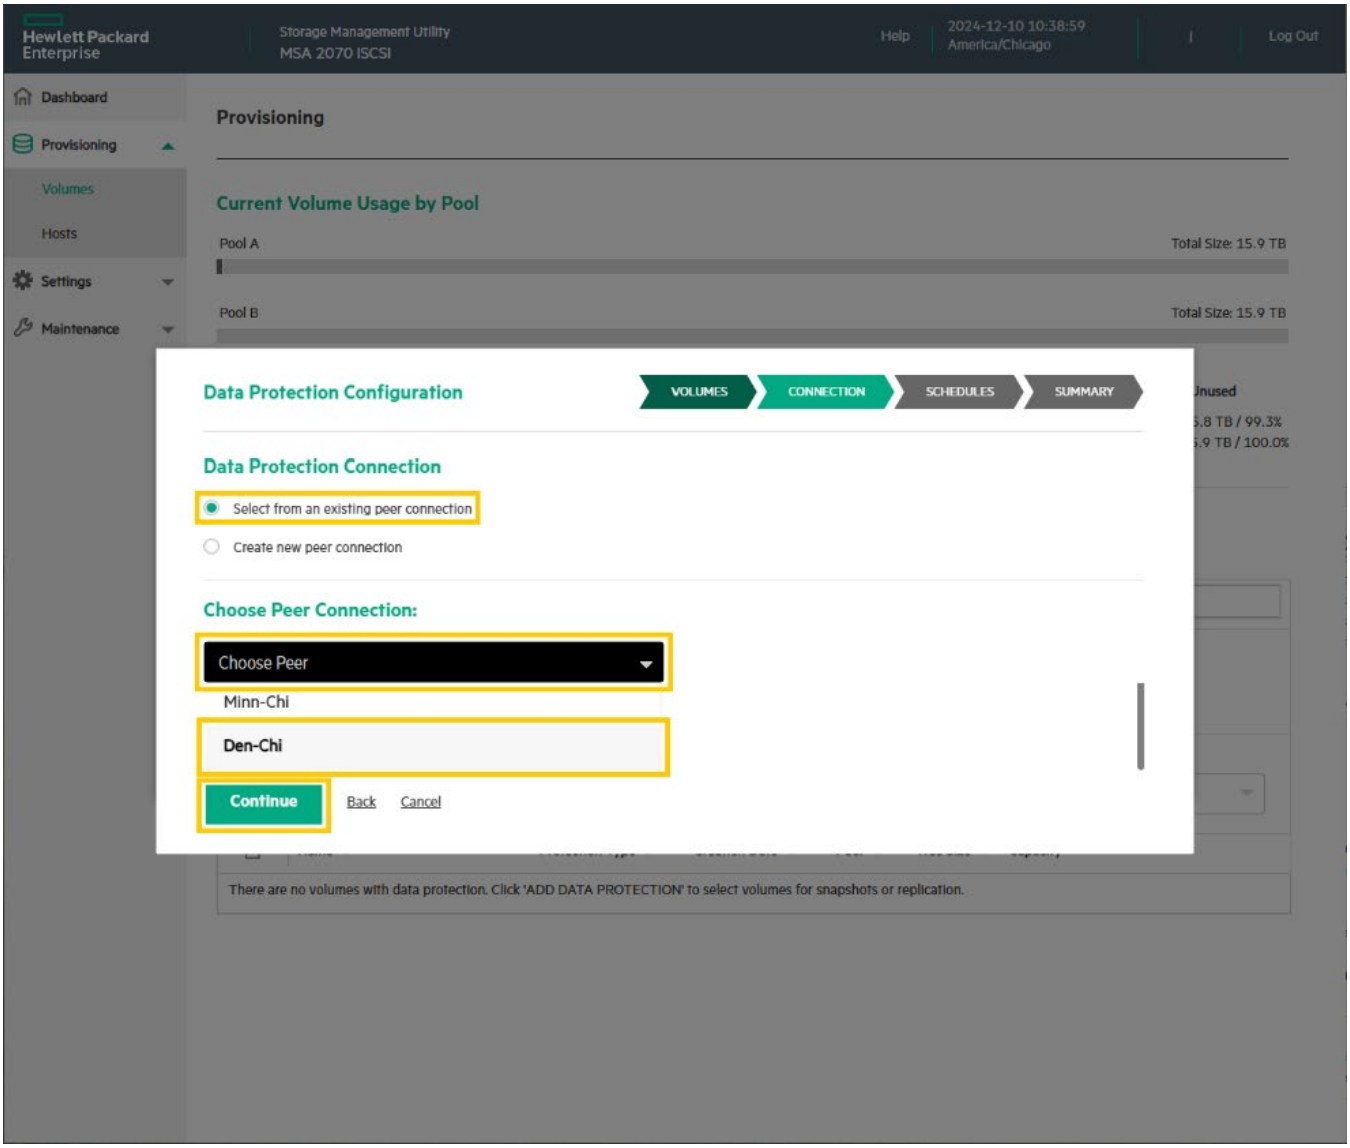


Figure 22. Data Protection Configuration wizard – Connection – Data Protection Connection – choose existing

Now that the peer connection has been established, it is time to set the remaining replication parameters. Provide the replication set name, set the secondary volume name (the primary volume name is the default), select the secondary pool, and select the queue policy in the top four sections of the dialog (see Figure 23. Data Protection Configuration wizard – Connection – Replication Settings – name, secondary name and pool, queue policy).

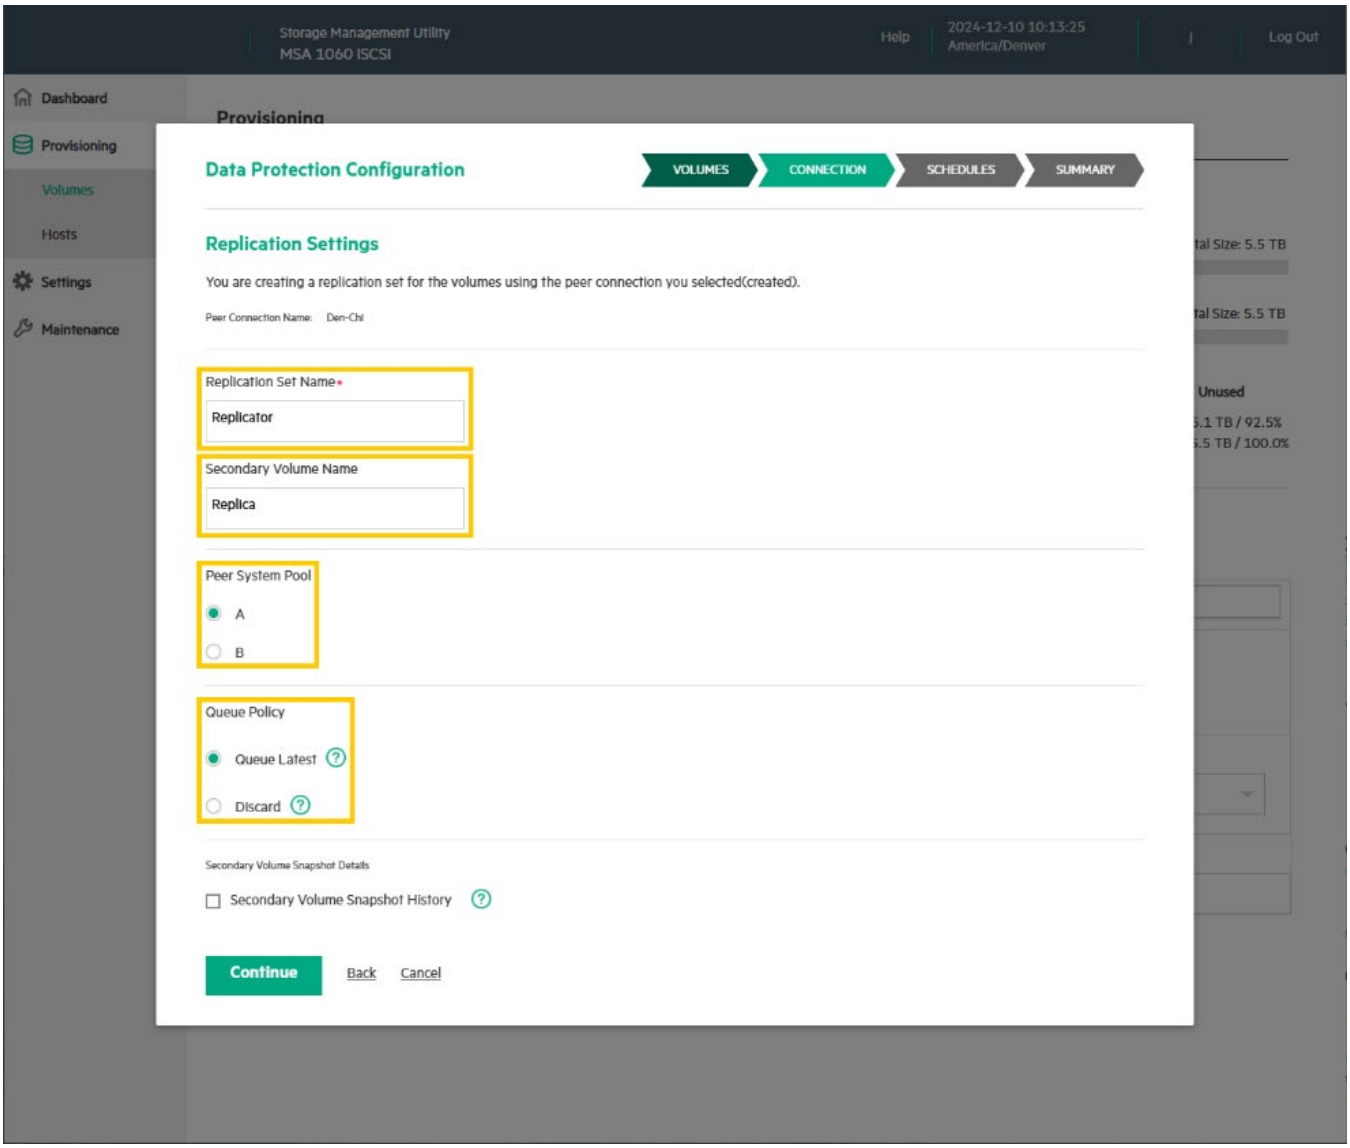


Figure 23. Data Protection Configuration wizard – Connection – Replication Settings – name, secondary name and pool, queue policy

Scroll down and continue with replication history snapshot configuration (not for a replication set of a volume group); select **Secondary Volume Snapshot History** to retain snapshot history on the secondary array and to enable setting the number of snapshots to retain (the default is 1), set the history snapshots' basename, select **Primary Volume Snapshot History** to retain snapshot history on the primary array, and choose the retention priority (the default is *Never-Delete*). When finished, click **Continue**. See Figure 24. Data Protection Configuration wizard – Connection – Replication Settings – replication history snapshots.

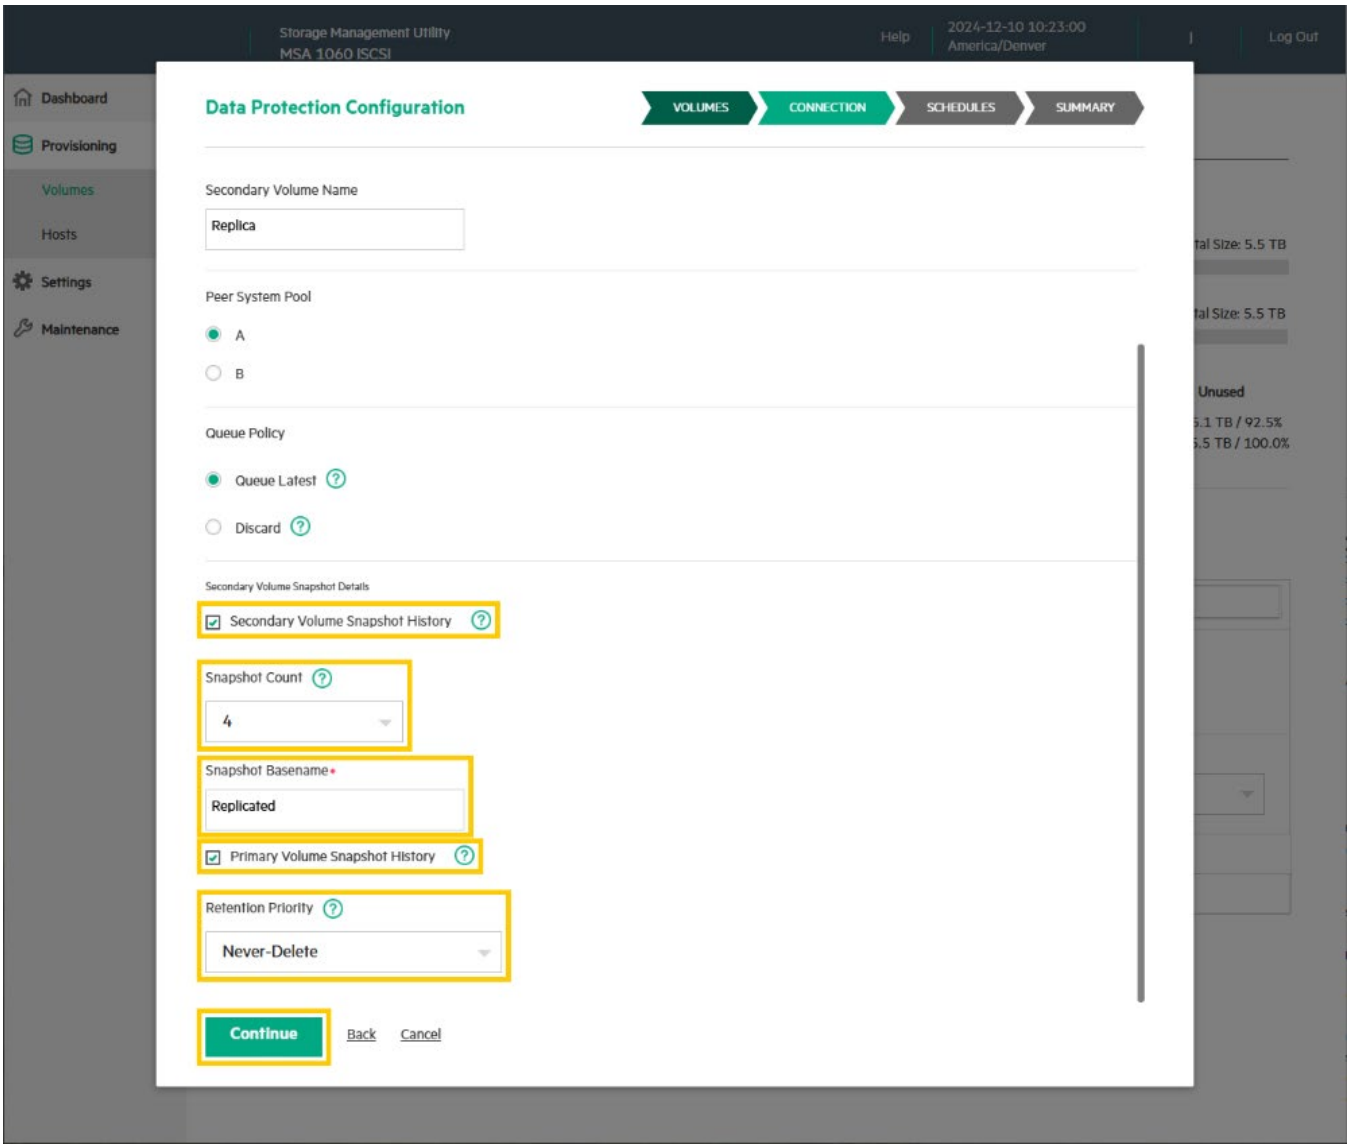


Figure 24. Data Protection Configuration wizard – Connection – Replication Settings – replication history snapshots

The last step is to optionally configure a replication schedule and optionally perform the initial replication upon replication set creation completion – *this is the only time you can create a schedule using the SMU* (see Figure 25. Data Protection Configuration wizard – Schedules – Replication Schedule Options – Replication Schedule – name, start date and time, and repetition) Click **Replicate Volume Immediately (upon configuration completion)** to perform the initial replication upon replication set completion and select **Create a Schedule** to create a schedule for replications. If you chose to create a schedule, the dialog opens up so that you can enter a name for the schedule, and optionally select **Last Snapshot** if the replication uses an existing snapshot; leave this unchecked to reset the current replication snapshot when replication begins. Enter the time and date to start the schedule and select **Every** to specify schedule repetition parameters.

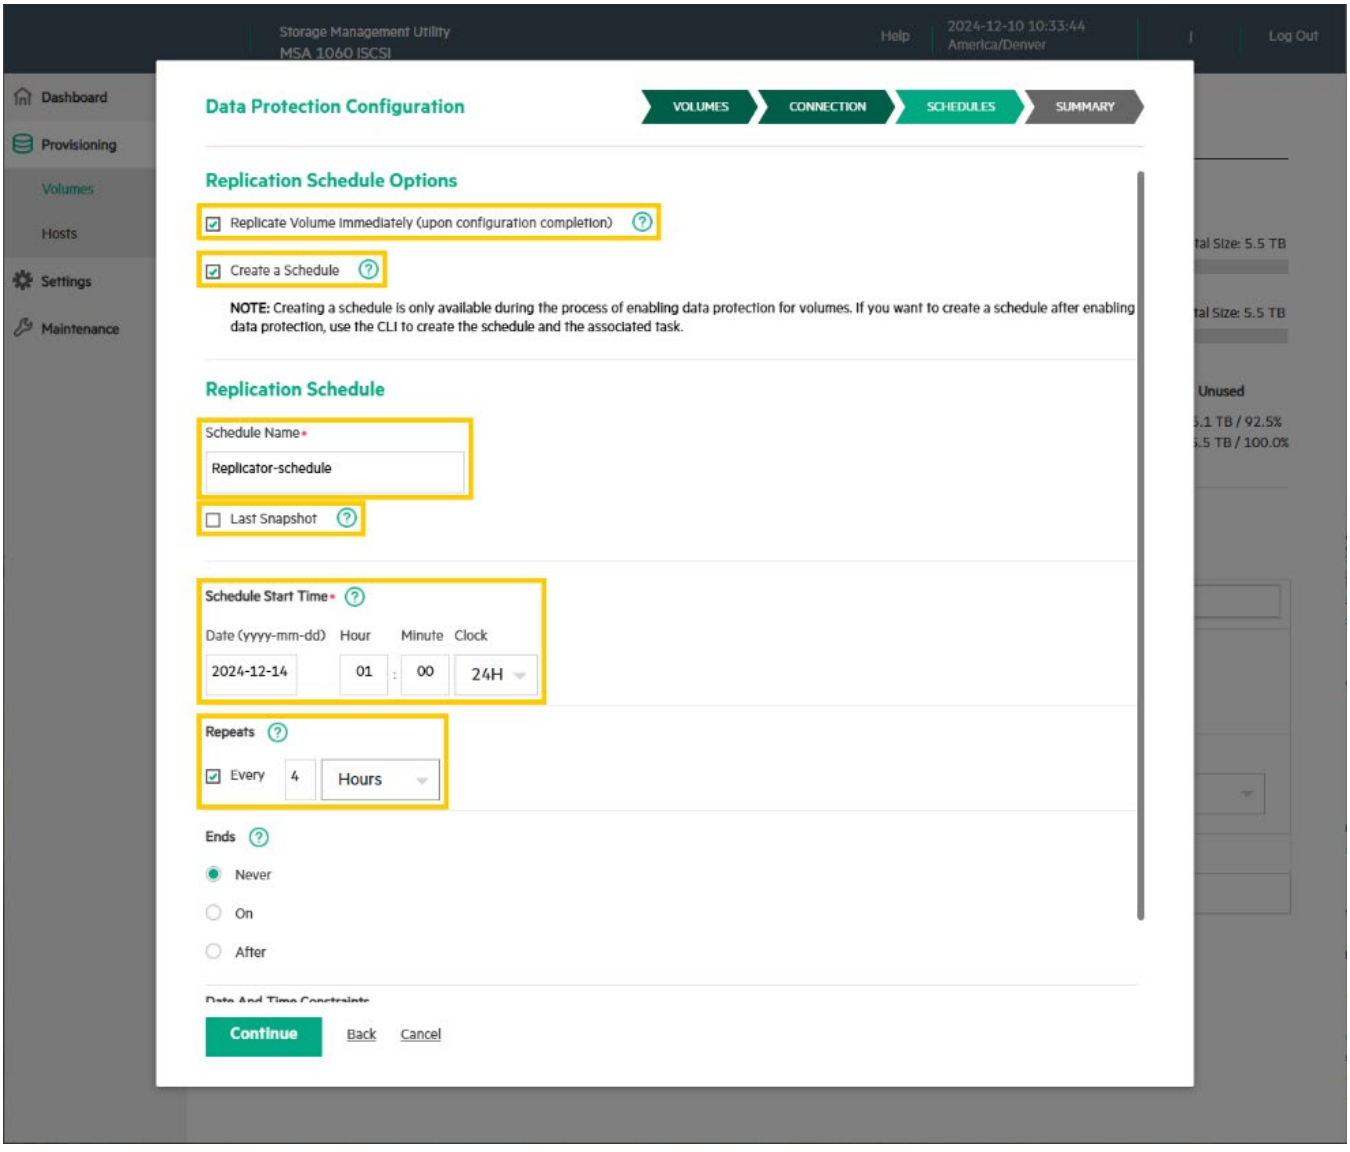


Figure 25. Data Protection Configuration wizard – Schedules – Replication Schedule Options – Replication Schedule – name, start date and time, and repetition

Scroll down and continue with the schedule ending condition (see Figure 26. Data Protection Configuration wizard – Schedules – Replication Schedule Options – Replication Schedule – end condition, date and time constraints), either to never end the schedule, or to end on a specified date and time, or end after a specified number of occurrences. Finally, you can specify date and time constraints under which to run the schedule.

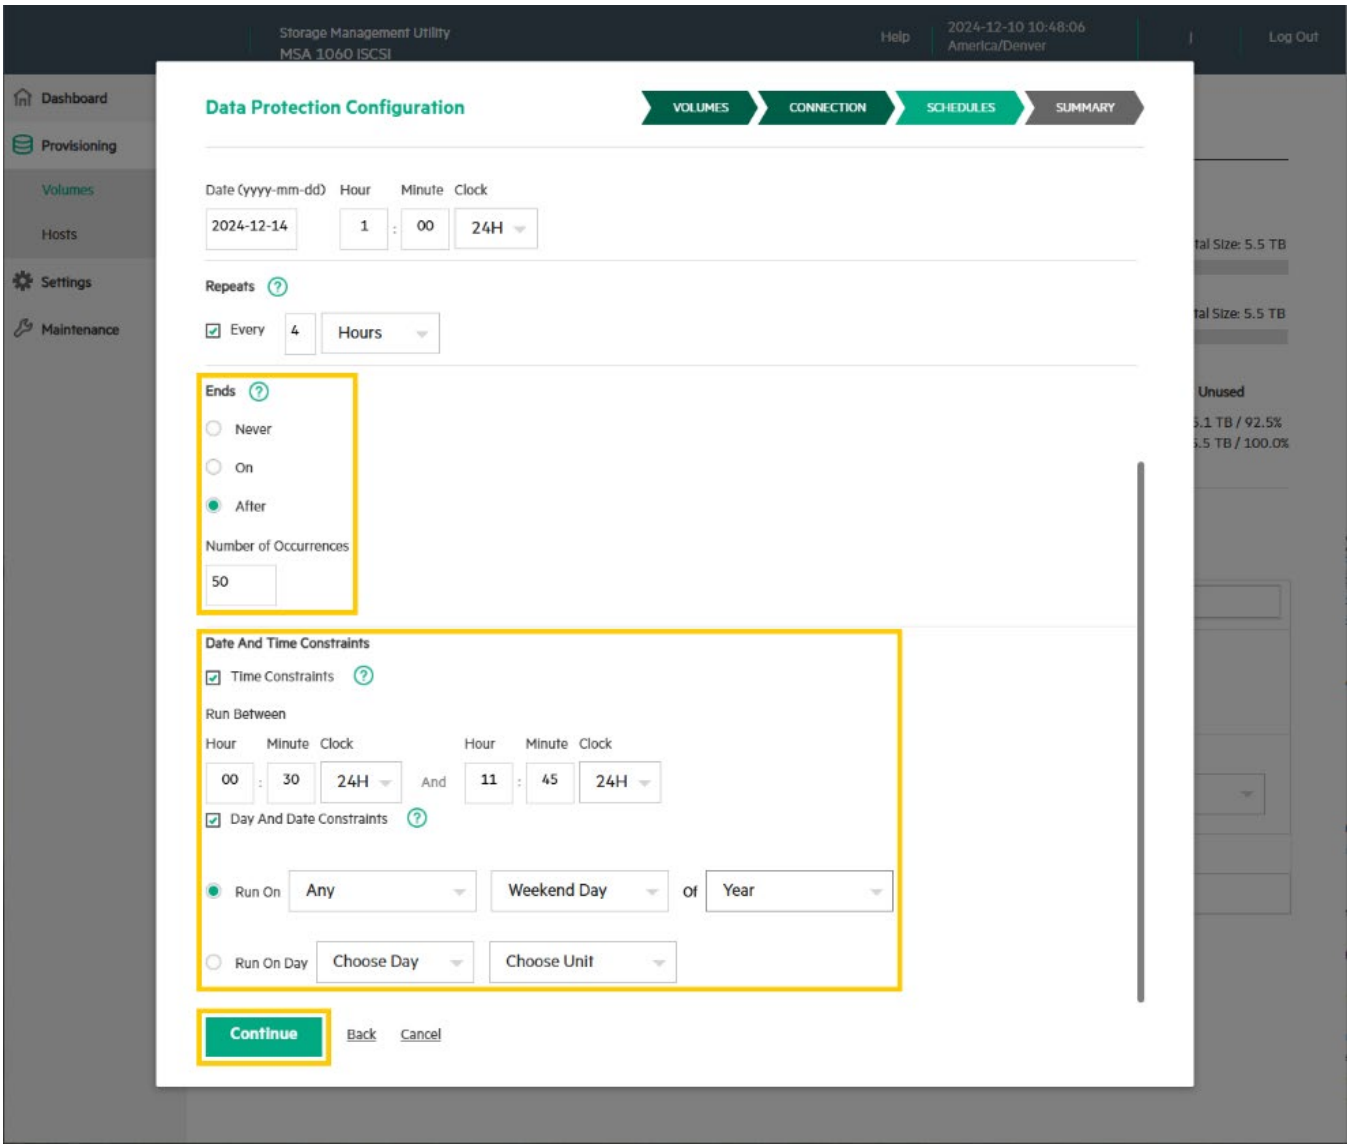


Figure 26. Data Protection Configuration wizard – Schedules – Replication Schedule Options – Replication Schedule – end condition, date and time constraints

The last page of the **Data Protection Configuration** wizard (see Figure 27. Data Protection Configuration wizard – Summary – Summary of Data Protection Settings) shows various parameters of the new replication set, such as primary volumes, secondary volumes and pool, whether or not a schedule will be created, and if so, the start date, the peer connection name, the queue policy, and replication snapshot history configuration. Select **Apply** to create the replication set.

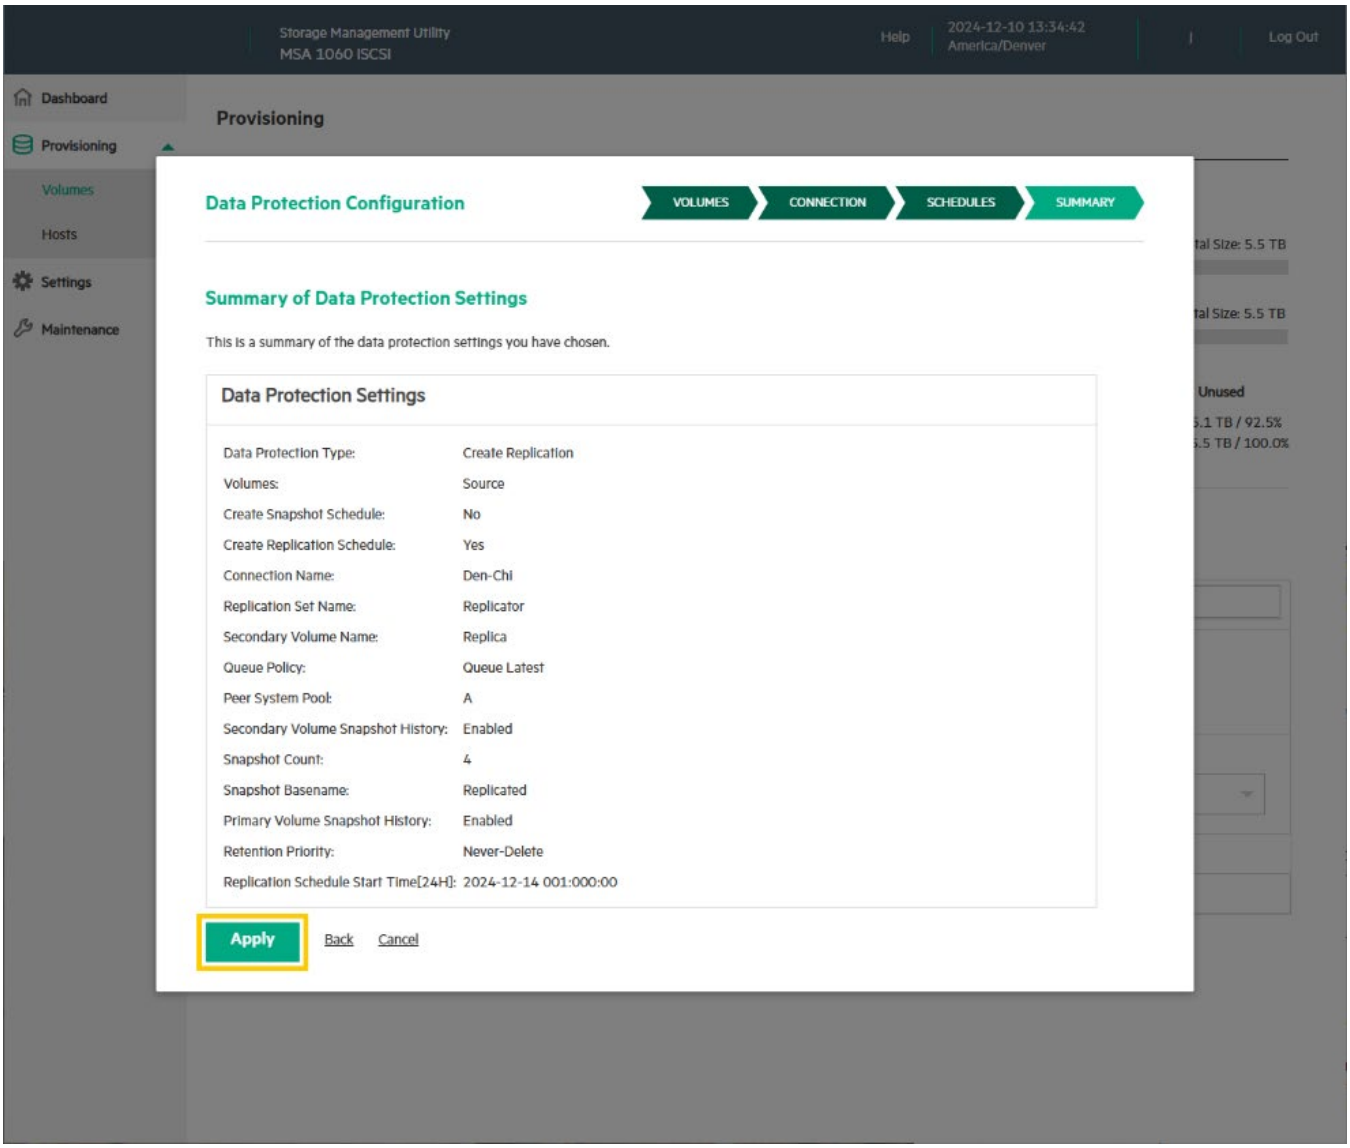


Figure 27. Data Protection Configuration wizard – Summary – Summary of Data Protection Settings

Modifying a replication set

You can change the name, replication snapshot history parameters, or queue policy of a replication set.

Using the CLI

Use the `set replication-set` command (see Command Example 9. Modifying the name and snapshot count of a replication set), with the replication set ID of the replication set to modify, and provide the new replication set name, replication snapshot history parameters and queue policy, as desired.

```
# set replication-set name Source-Replica snapshot-count 2 Replicator
```

Command Example 9. Modifying the name and snapshot count of a replication set

Using the SMU

From the **Volumes** panel (**Provisioning** > **Volumes**), in either the **Volumes** table or the **Data Protection** table, click the slide-over for a volume, select the **Replications** tab (see Figure 28. Modifying a replication set – Snapshot Count), and edit any item that has the edit icon . The following example shows editing the **Snapshot Count**. Use the apply button to apply the change, use the cancel button to discard the change.

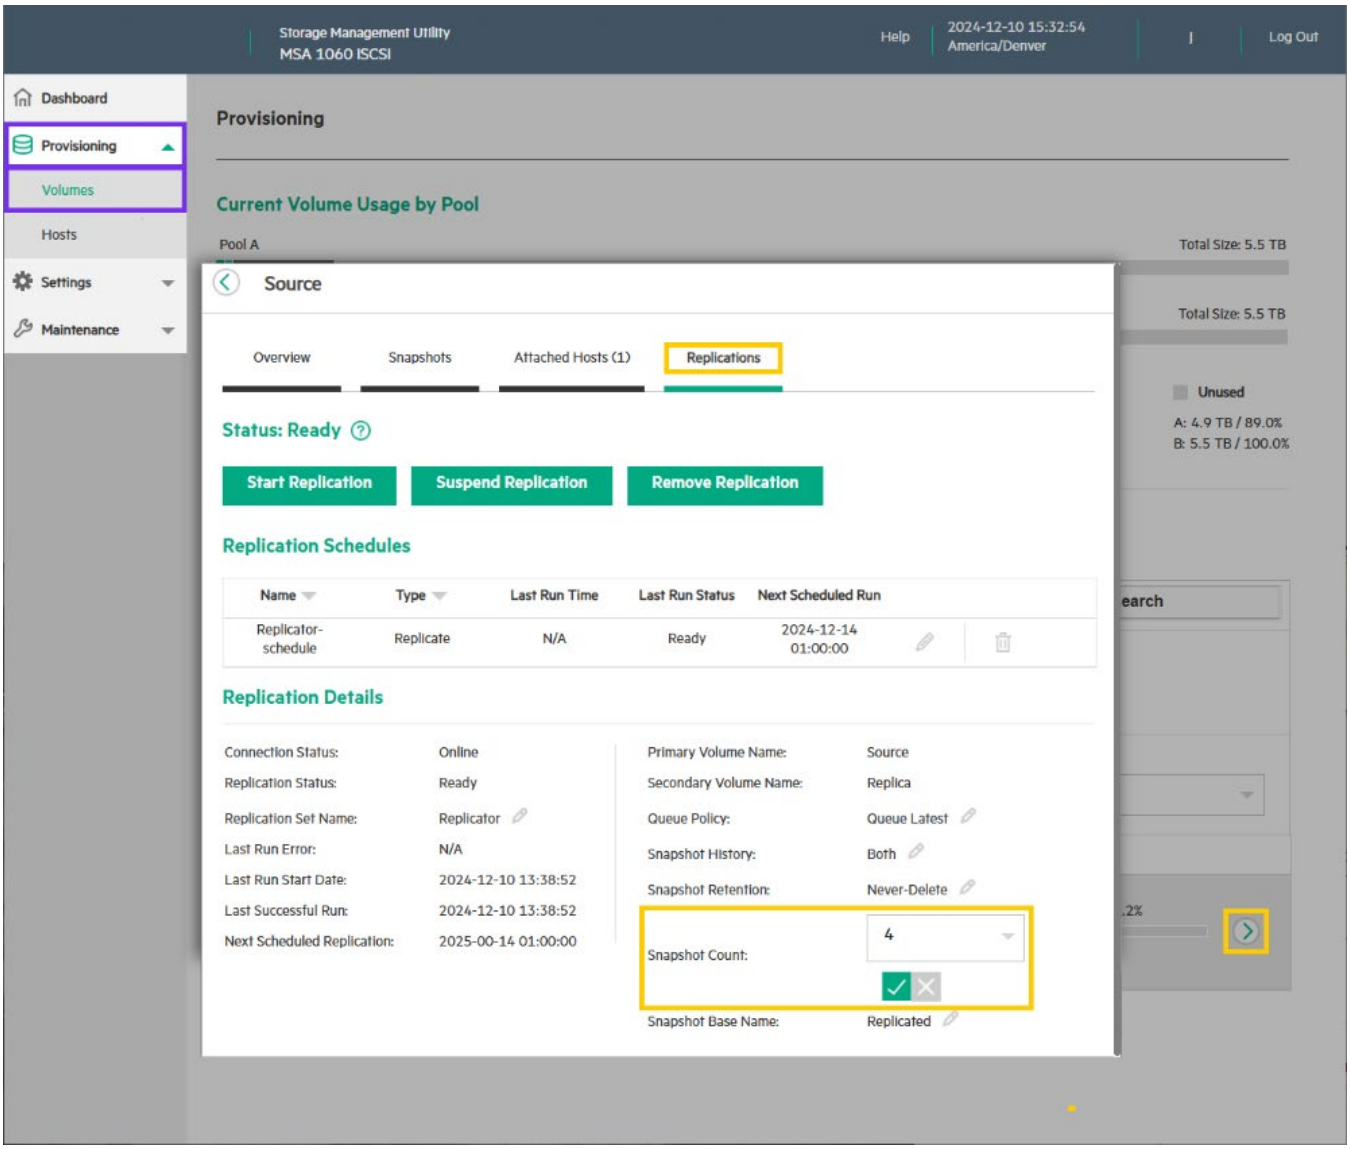


Figure 28. Modifying a replication set – Snapshot Count

Scheduling replications

The scheduler can be used to replicate data from primary volumes in regular intervals to secondary volumes. Creating a replication schedule consists of two parts; creating a replication task, which indicates the action of the task, in this case replication, and parameters associated with the task, and creating a schedule for running the task. The CLI requires two commands to perform this, while the SMU creates the task and schedule in one operation.

Common schedule parameters

Parameters common to all schedules are the start time and date, which must be a time and date in the future; recurrence or repetition interval—if not set or selected, the replication will occur only once; end time and date or count limit (the number of times to run the task); and time and date constraints, which only constrains when the task starts, and is not a window in which the task must complete.

Using the CLI

First, create the task (see Command Example 10. Creating a replication task). For replication tasks, you must specify the task type (*Replicate*), the replication set name or serial number, and the task name. Specify the last-snapshot parameter to use the existing last snapshot of the volume for replication, rather than resetting the replication snapshot when replication begins.

```
# create task type Replicate Replication-set Replicator ReplicatorTask
```

Command Example 10. Creating a replication task

Then, create the schedule (see Command Example 11. Creating a replication schedule). You must provide the task to run and the schedule's name; while you must also provide a schedule specification, only the start time is required. See Common schedule parameters (page 51) for more information. The minimum interval is 30 minutes.

```
# create schedule schedule-specification "start 2024-12-14 01:00 every 4 hours only any weekenddday
of year" task-name ReplicatorTask ReplicatorSchedule
```

Command Example 11. Creating a replication schedule

To change using an existing snapshot for a replication or resetting the current one, or to change the replication set associated with a task, use the `set task` CLI command.

Use the `set schedule` CLI command to modify the schedule parameters or to set the task associated with the schedule. See Common schedule parameters (page 51) for more information on the schedule specification.

Using the SMU

Note that the only time you can create a schedule using the SMU is when you create the replication set. You can modify an existing schedule in the SMU. Start in the **Volumes** panel (**Provisioning** > **Volumes**), then in either the **Volumes** table or the **Data Protection** table, click the slide-over for a volume, select the **Replications** tab (see Figure 29. Modifying a replication schedule – selecting the schedule to modify), and click the edit icon for a schedule.

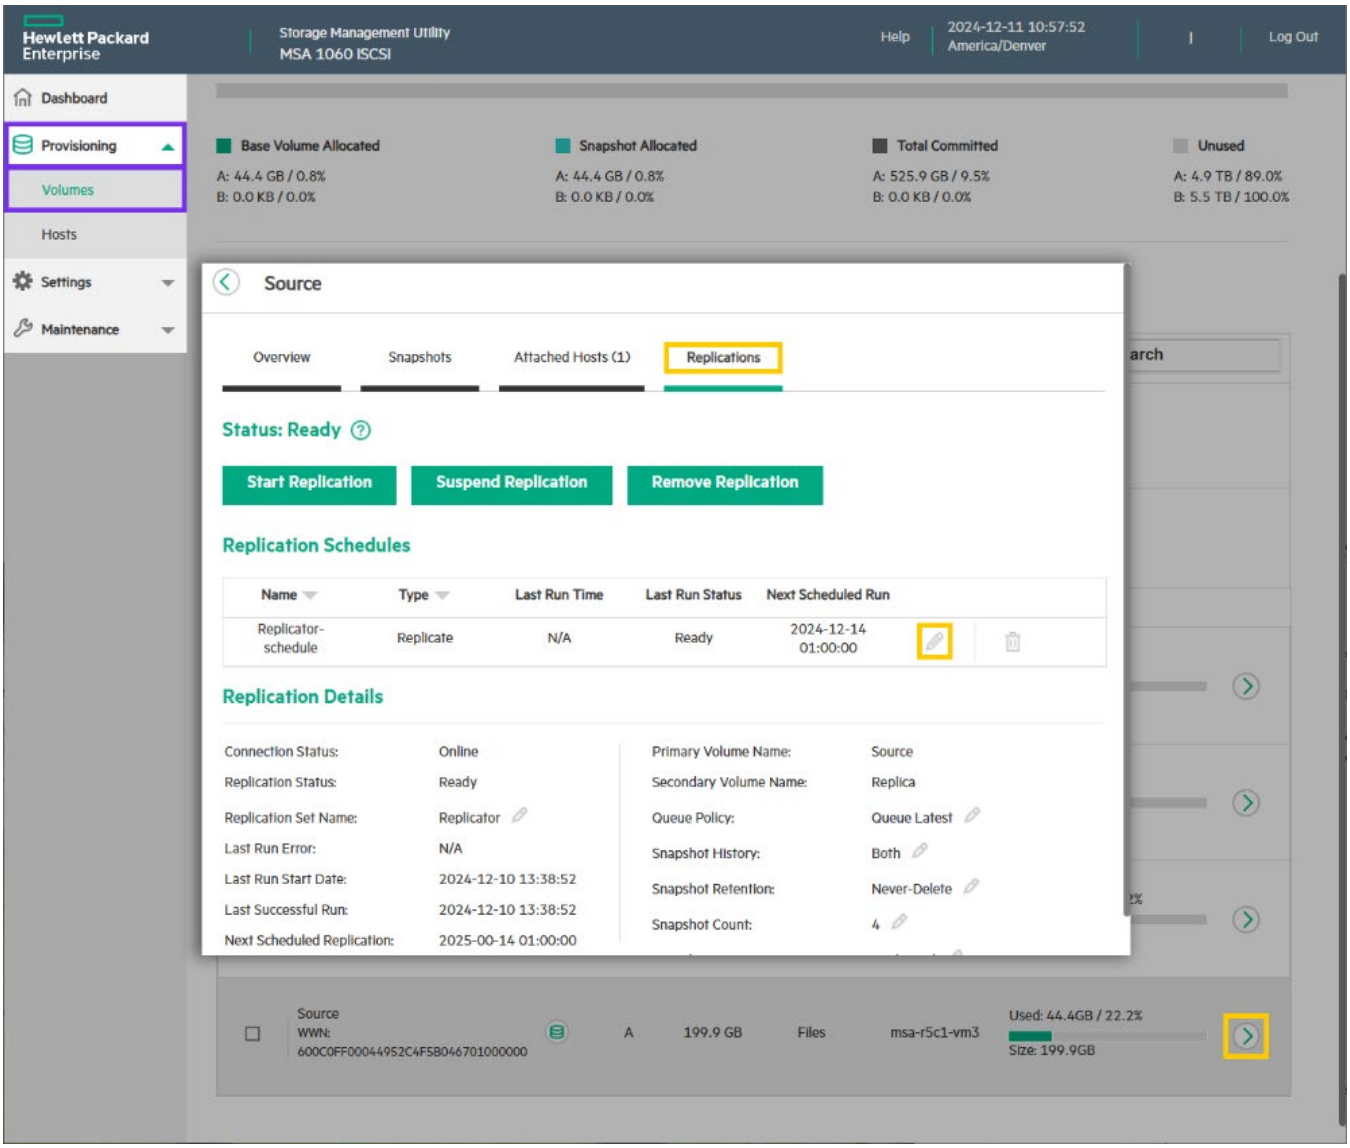


Figure 29. Modifying a replication schedule – selecting the schedule to modify

Then, modify the schedule similar to creating the schedule as described in Using the SMU section of Creating a replication set (page 36) – see Figure 30. Modifying a replication schedule – editing the schedule for a sample dialog.

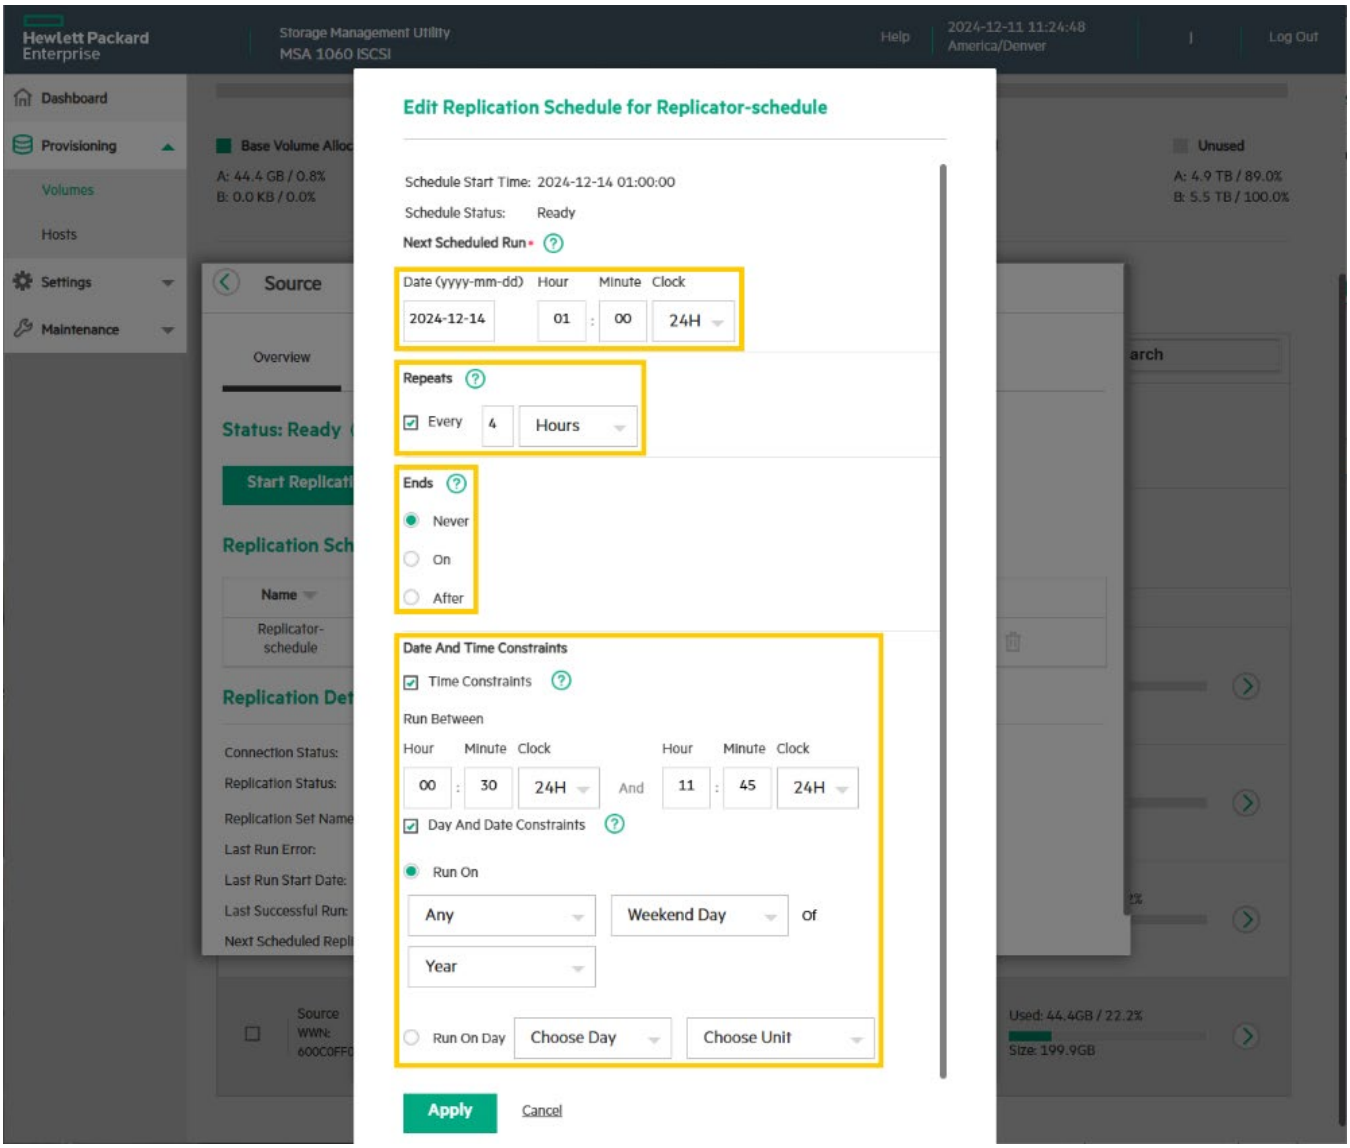


Figure 30. Modifying a replication schedule – editing the schedule

Initiating a replication

A replication may only be initiated on the array containing the primary volume. If a problem occurs, the array suspends the replication; it will attempt to resume the replication every 10 minutes for the first hour, then every hour until the replication resumes. A suspended replication set may be aborted. A running replication set may be manually suspended or aborted.

You may replicate the last snapshot or a specific existing snapshot of a volume rather than reset the current replication snapshot; this option is available using the CLI only, not the SMU, and applies only to a replication set of a volume, not to the replication set of a volume group. This might be used to replicate application created snapshots, such as those created by the Volume Shadow Copy Service (VSS) used in Microsoft operating systems. If one of these options is used, the snapshot's content is copied to the primary replication set's current replication snapshot, then the standard replication procedure progresses from there.

Using the CLI

Use the `replicate` CLI command – see Command Example 12. Initiating a replication for an example that shows replicating the Replicator replication set, using the VSS snapshot VSS170322170906953.

```
# replicate snapshot VSS170322170906953 Replicator
```

Command Example 12. Initiating a replication

Using the SMU

Start in the **Volumes** panel (**Provisioning** > **Volumes**), then in either the **Volumes** table or the **Data Protection** table, click the slide-over for a volume, select the **Replications** tab, and click **Start Replication** – see Figure 31. Starting a replication.

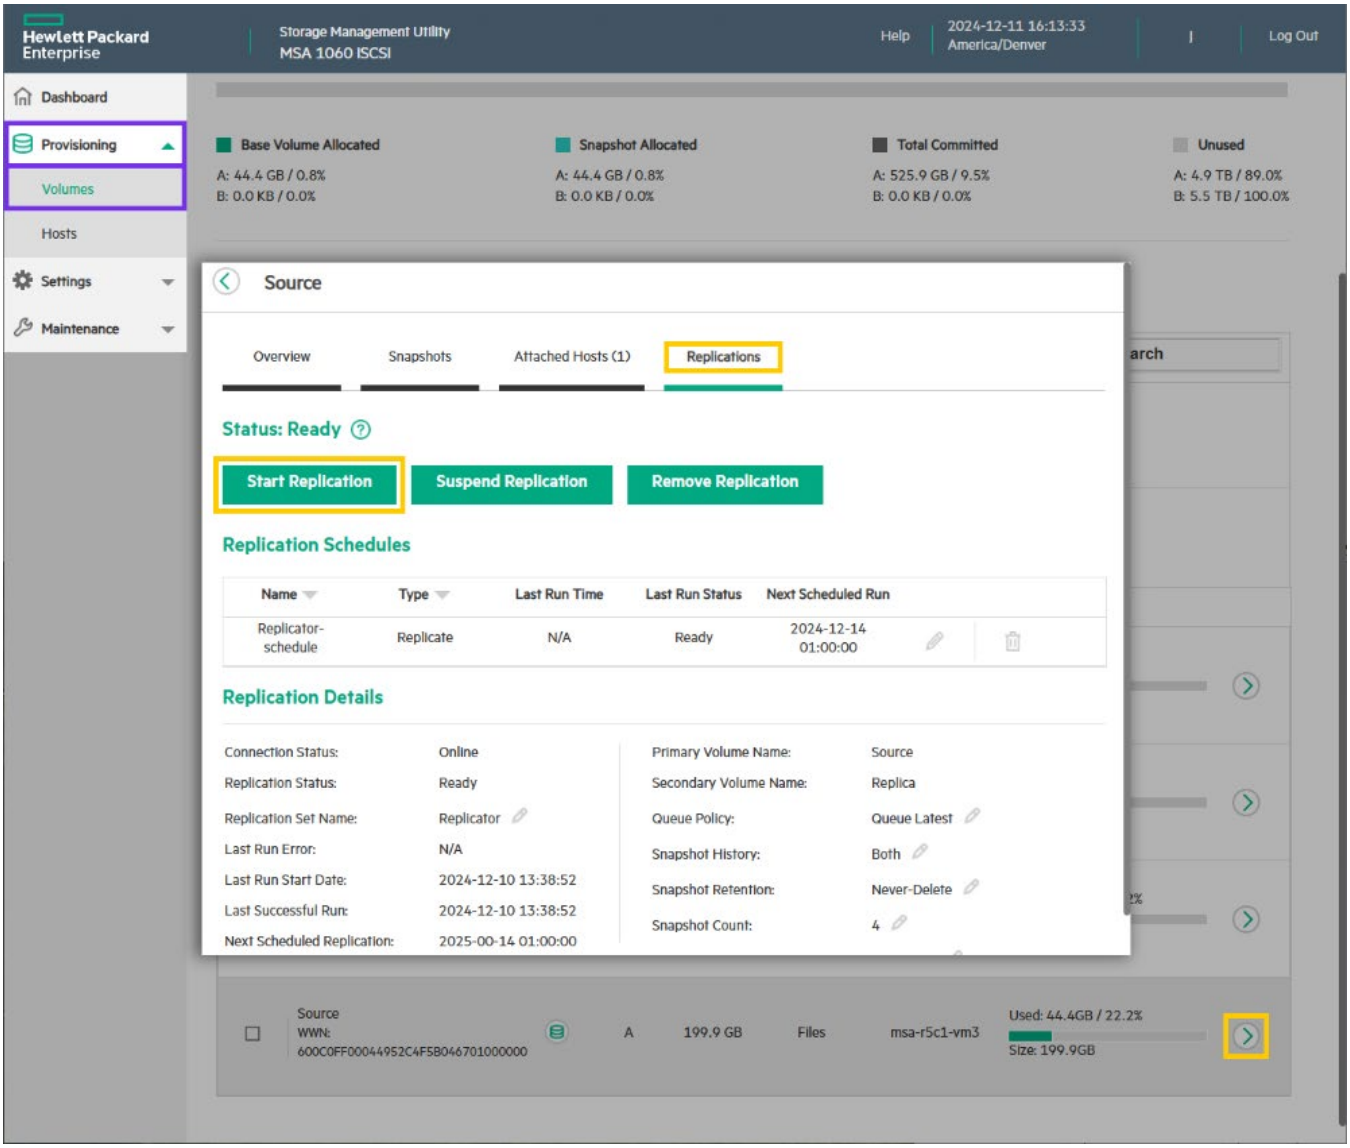


Figure 31. Starting a replication

Suspending a replication set

Suspending a replication set pauses any replication in progress and prevents new replications on the replication set. Replication sets may be suspended from the primary array only. Suspended replications can be aborted or resumed. Resume a suspended replication set to allow the replication set to resume replications that were in progress and not aborted, and to allow new replications to occur. Attempting to initiate a replication, including scheduled replications, fails while the replication set is suspended if the queue policy is set to *Discard*; if the queue policy is set to *Queue Latest*, the primary volume's queued replication snapshot is reset.

Aborting a replication while it is suspended prevents resuming the replication once the replication set is resumed.

Resetting a peer connection using a new remote address cannot be made while any associated replication set is in progress. Before modifying a peer connection, abort any replications in progress and suspend any associated replication sets.

Using the CLI

Use the `suspend replication` CLI command. See Command Example 13. Suspending a replication set for an example that suspends the replication set named Replicator and suspends any replication in progress on the replication set.

```
# suspend replication-set Replicator
```

Command Example 13. Suspending a replication set

Using the SMU

Use the **Replications** tab (**Provisioning** > **Volumes** > volume slide-over > **Replications**) or (**Provisioning** > **Volumes** > **Data Protection** table view > volume slide-over > **Replications**) to suspend a replication set, which also suspends a replication in progress.

Select **Suspend Replication** if no replications are currently in progress – see Figure 32. Suspending a replication set – no replications in progress.

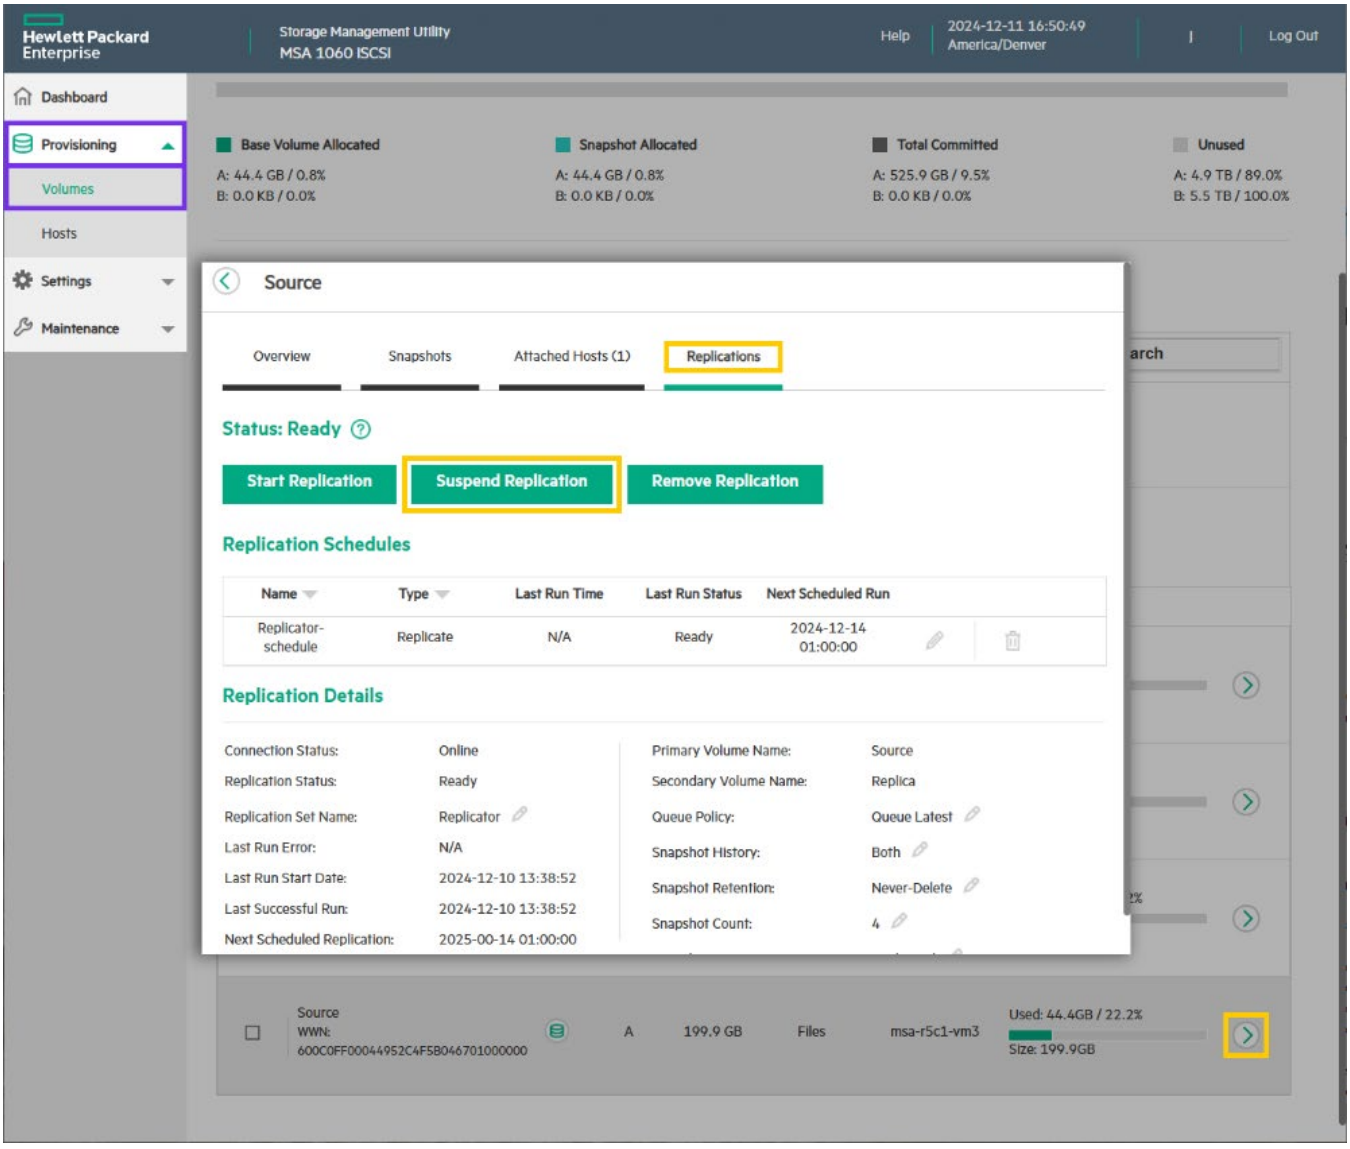


Figure 32. Suspending a replication set – no replications in progress

Select the suspend icon if a replication is in progress – see Figure 33. Suspending a replication set in progress.

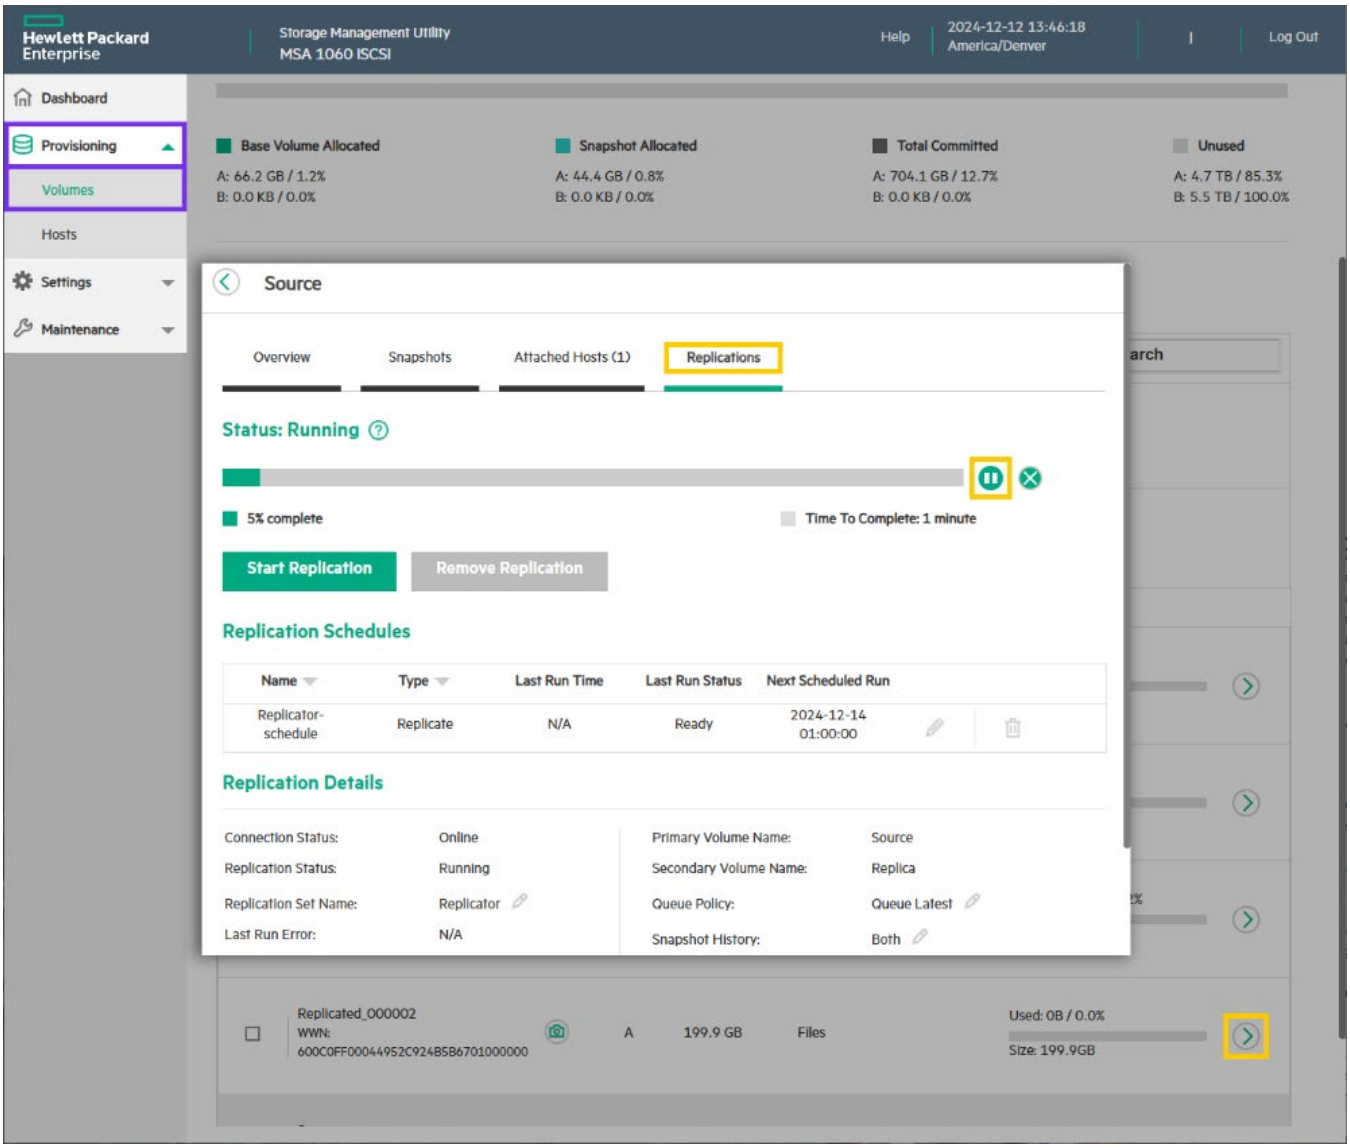


Figure 33. Suspending a replication set in progress

Resuming a replication or replication set

Resuming a suspended replication set resumes a suspended replication that was in progress and allows new replications to occur on the replication set. Replication sets may be resumed from the primary array only. Aborted replications cannot be resumed.

Using the CLI

Use the `resume replication-set` CLI command. See Command Example 14. Resuming a replication set for an example that resumes the replication set Replicator and resumes a suspended replication that was not aborted.

```
# resume replication-set Replicator
```

Command Example 14. Resuming a replication set

Using the SMU

Use the **Replications** tab (**Provisioning** > **Volumes** > volume slide-over > **Replications**) or (**Provisioning** > **Volumes** > **Data Protection** table view > slide-over > **Replications**) to resume a replication set, and to resume a suspended replication as long as it was not aborted.

Select **Resume Replication** if no replications are suspended – see Figure 34. Resuming a replication set – no suspended replications.

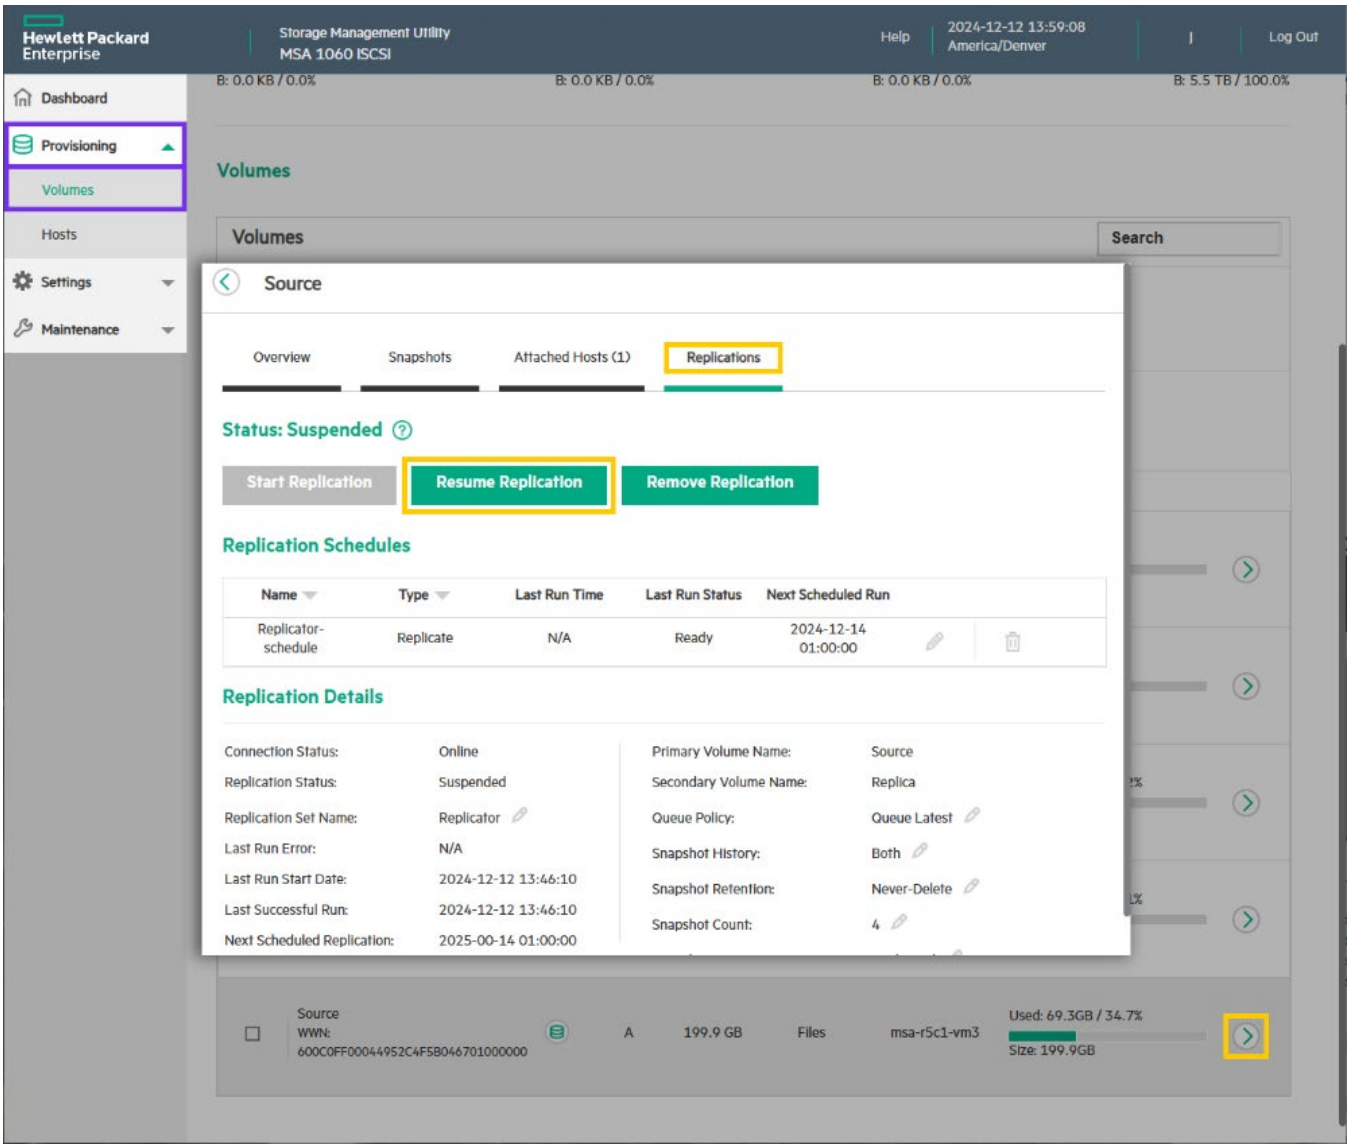


Figure 34. Resuming a replication set – no suspended replications

Select the resume icon if a replication was in progress, and is now suspended, but not aborted – see Figure 35. Resuming a suspended replication set.

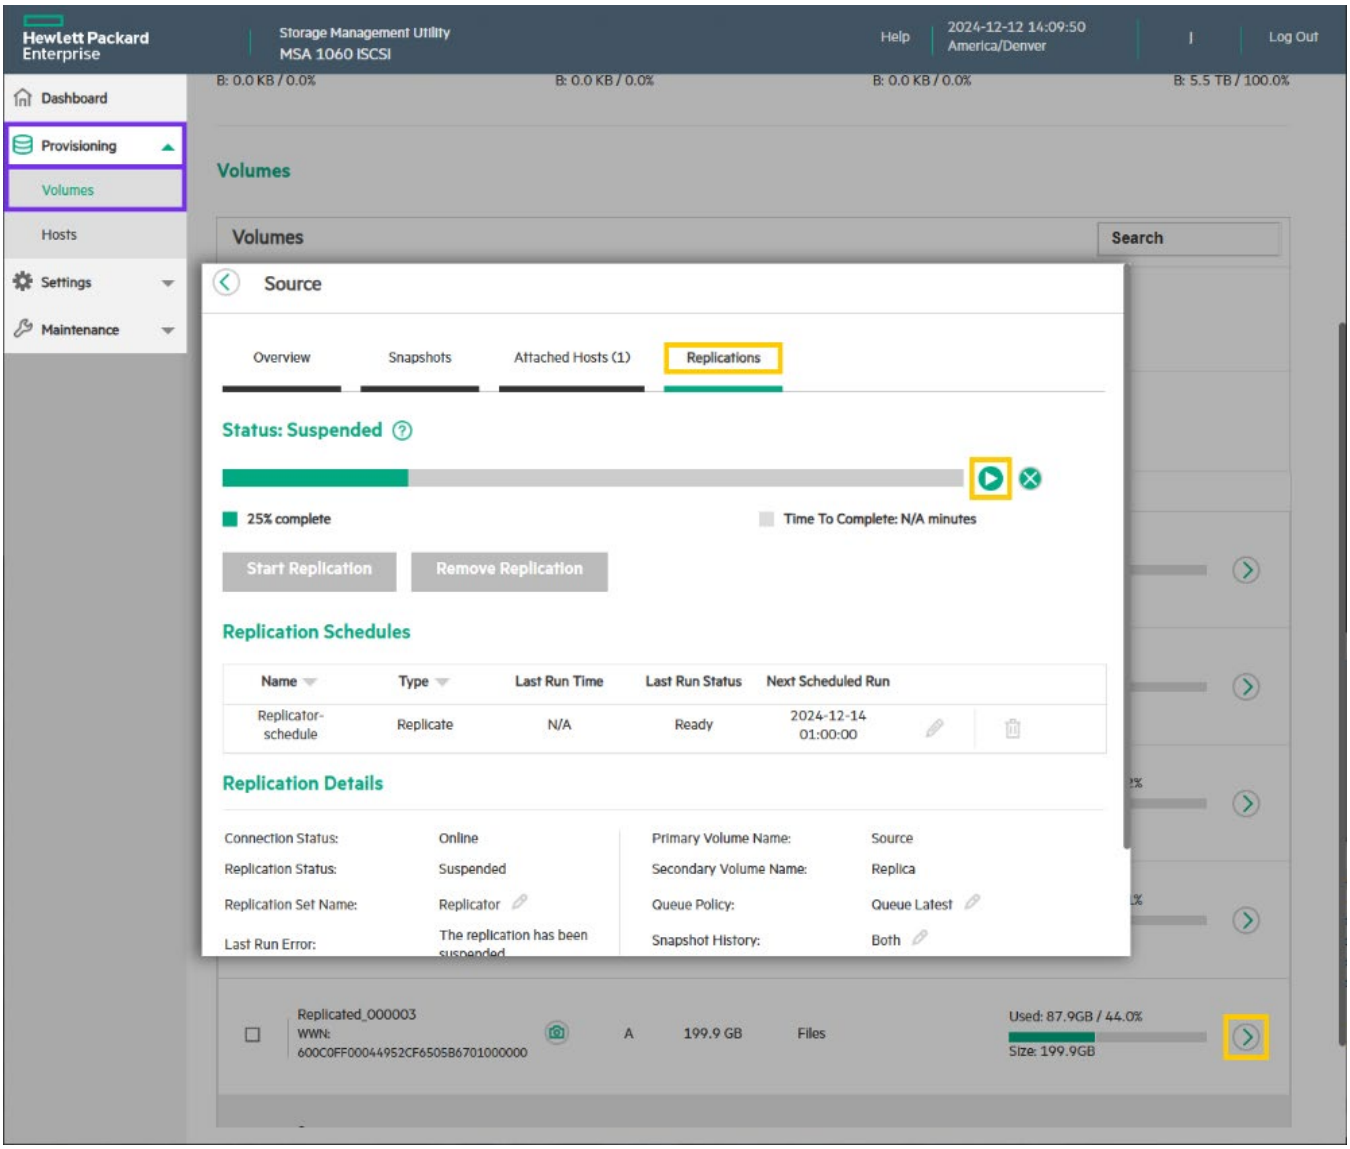


Figure 35. Resuming a suspended replication set

Aborting a replication

You can abort a replication that is either in a running or suspended state. Replications may be aborted from the primary array only. Aborted replications cannot be resumed.

Using the CLI

Use the `abort replication` CLI command. See Command Example 15. Aborting a replication for an example that aborts a running or suspended replication for the replication set named Replicator.

```
# abort replication Replicator
```

Command Example 15. Aborting a replication

Using the SMU

Use the **Replications** tab (**Provisioning** > **Volumes** > volume slide-over > **Replications**) or (**Provisioning** > **Volumes** > **Data Protection** table view > slide-over > **Replications**) and select the abort icon , which is only available if a replication is in progress or is suspended – see Figure 36. Aborting a running replication. If the replication is in progress, the status is Running.

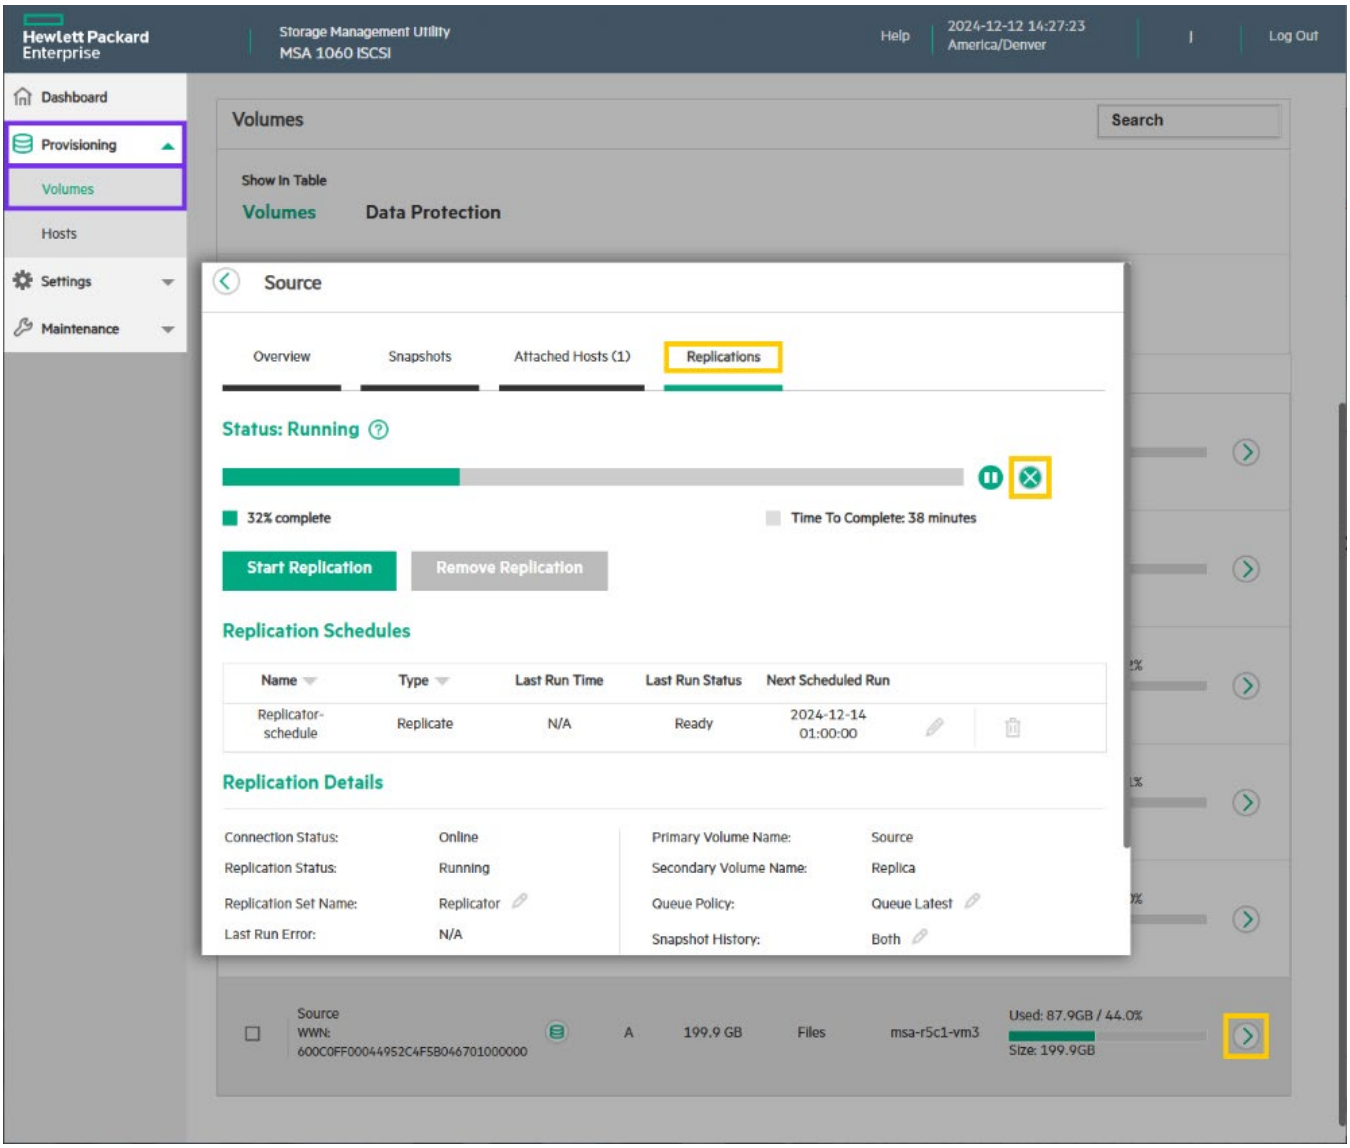


Figure 36. Aborting a running replication

Aborting a suspended replication uses the same tab, but shows the status as *Suspended* rather than *Running* – see Figure 37. Aborting a suspended replication.

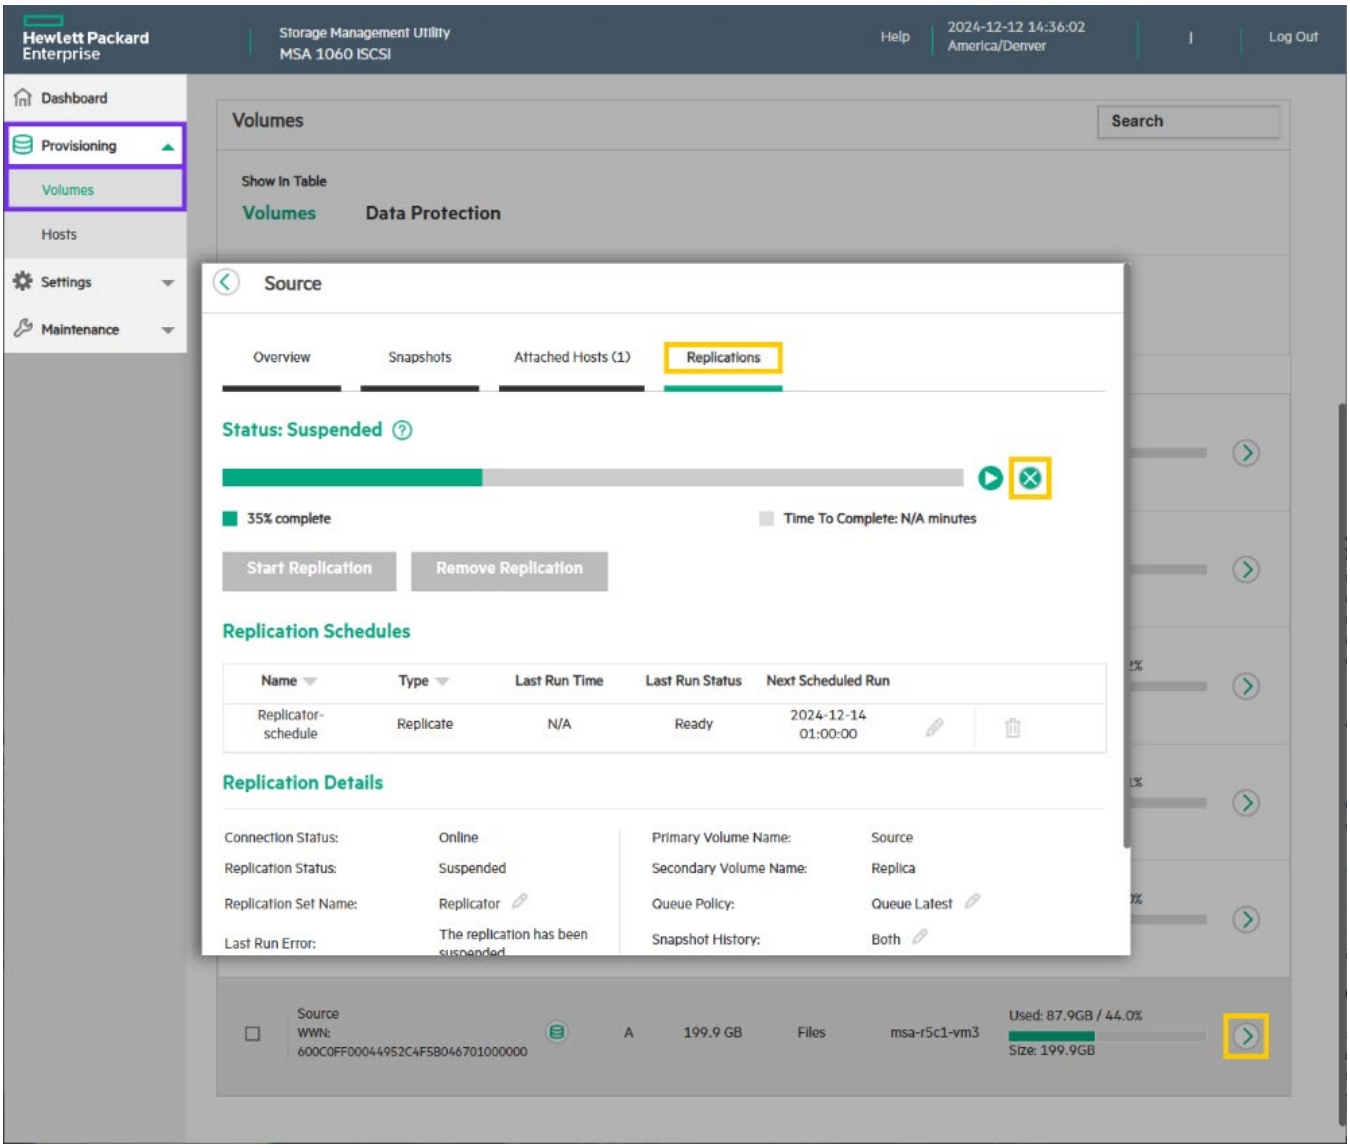


Figure 37. Aborting a suspended replication

Displaying a replication set

For each replication set, or for the replication set specified, you can view the name, if the replication set is for a group (indicated by the notation <group-name>.*), an indication if the primary volume is local or remote, the peer connection, the names of the primary and secondary volumes, the queue policy, the queue count (0 or 1), the status of the replication set, when the last successful replication ran, the status of the last attempted replication, and details on the current replication and the last replication run – progress, data transferred, start time, end time or estimated completion time, and whether an error occurred

Using the CLI

Use the `show replication-sets` CLI command. See Command Example 16. Displaying a replication set for an example that shows information about the replication set named Replicator. The header and columns have been adjusted for space.

```
# show replication-sets Replicator
Name        Group Primary Location  Peer     Primary Volume Secondary Volume Policy       Queue Count Status
  Failback In Progress Last Successful Run  Last Status
-------------------------------------------------------------------------------------------------------------
Replicator  No    local             Den-Chi  Source         Replica          Queue Latest 0           Ready
  False               2024-12-12 14:09:10  Success

   Replication Progress Data Transferred  Start Time           End Time             Run Error
   ---------------------------------------------------------------------------------------
   Last Run    100%     21.8GB            2024-12-12 13:46:10  2024-12-12 13:58:41  N/A


   Replication Progress Data Transferred  Start Time           Estimated Completion Run Error
   ---------------------------------------------------------------------------------------
   Current Run 100%     21.8GB            2024-12-12 14:09:10  N/A                  N/A

-------------------------------------------------------------------------------------------------------------
```

Command Example 16. Displaying a replication set

Using the SMU

Use the **Data Protection** table view of the **Volumes** panel (**Provisioning** > **Volumes** > **Data Protection** table view) – see Figure 38. Displaying a replication set – volume table. This table lists all volumes that are protected, either by a replication set or a snapshot. To see details of a replication set, click the slide-over for a volume that shows the Primary Replication Volume icon or Secondary Replication Volume icon.

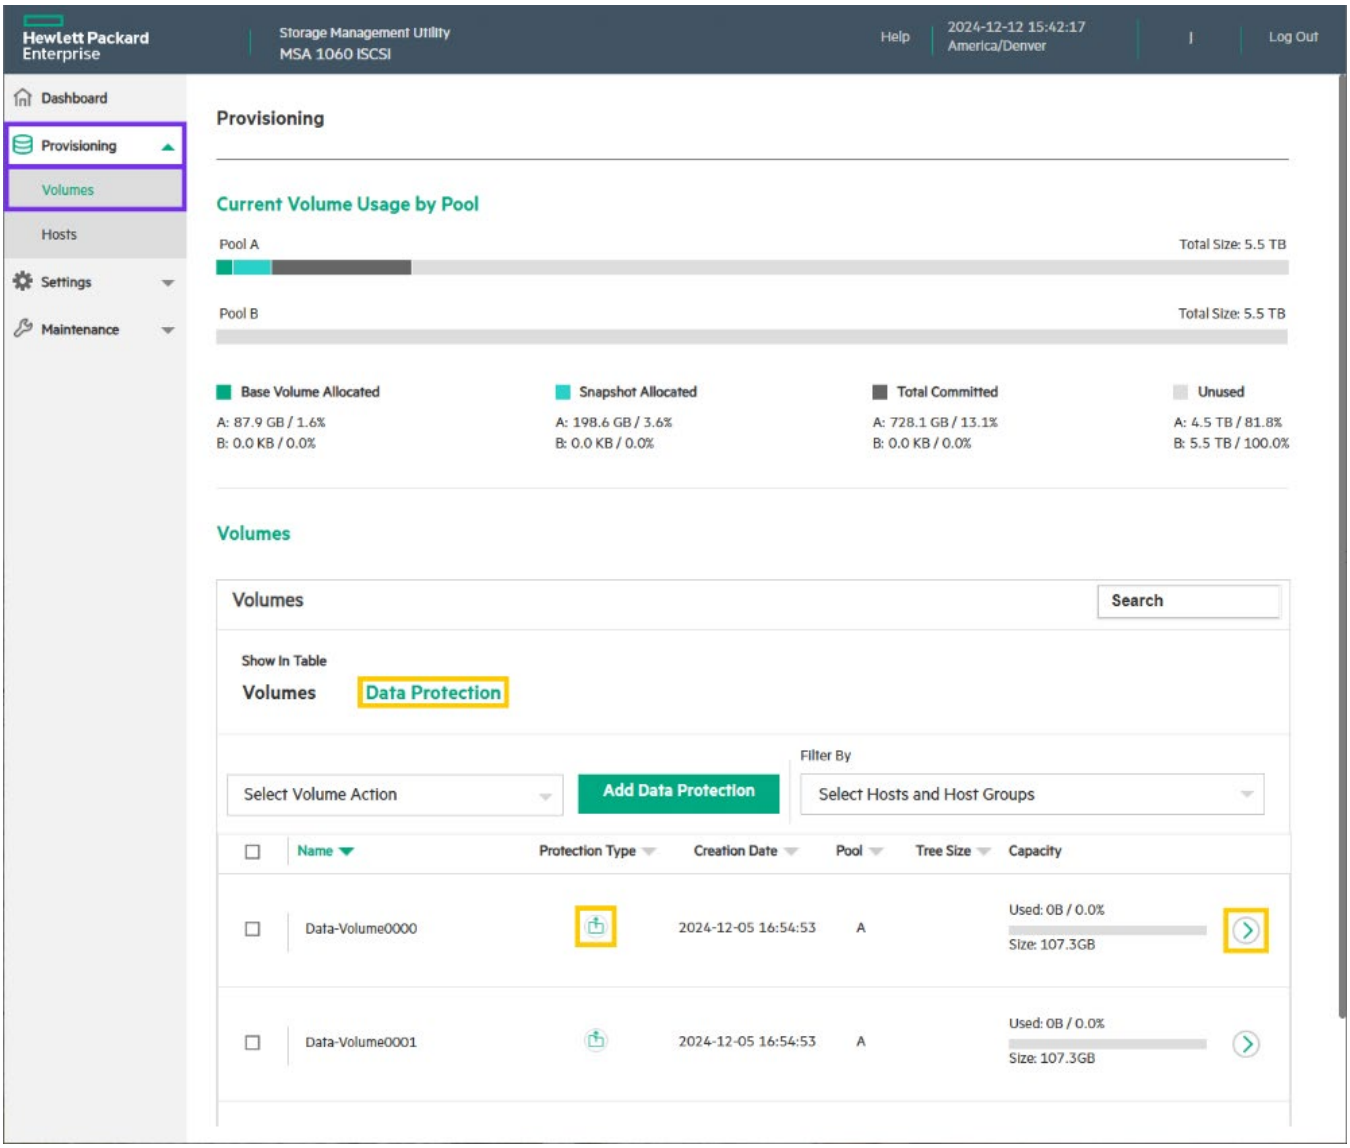


Figure 38. Displaying a replication set – volume table

From the detailed information panel of the volume, select the **Replications** tab. In the **Replications** tab, a volume group is designated using the <group-name>.* notation. See Figure 39. Displaying a replication set – details for an example that shows a replication set of a volume group.

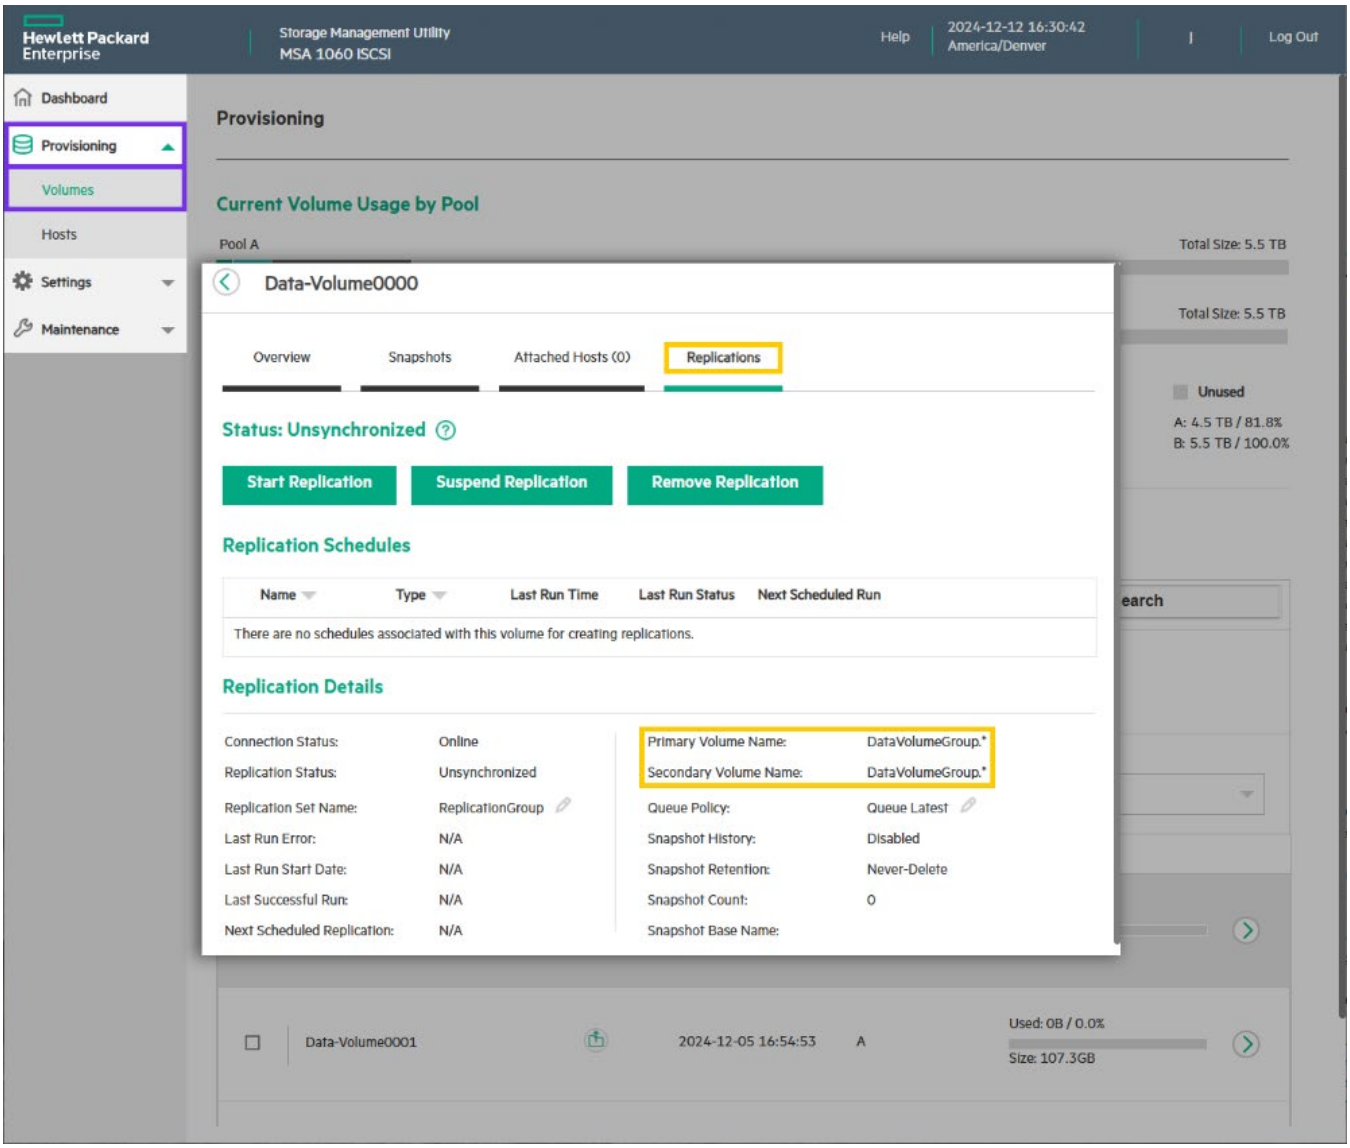


Figure 39. Displaying a replication set – details

Showing replication snapshot history

For each replication set, or for the replication set specified, show the replication set name, an indication of the replication snapshot history (*Disabled*, *Secondary* or *Both*), the retention count, the basename for the snapshots, the retention priority, and the list of replication history snapshots, each of which lists the snapshot name, the creation date and time, the amount of data associated with the snapshot, and the amount of data that is unique to the snapshot.

Using the CLI

Use the `show replication-snapshot-history` CLI command to display all the information listed above in one place. See Command Example 17. Displaying replication snapshot history for an example that shows the replication snapshot history for two replication sets, one of which, Replicator, has replication snapshot history enabled on both the primary and secondary volumes, the other, ReplicationGroup, does not have replication snapshot history enabled – it is actually a replication set of a volume group, and so replication snapshot history is not available for it.

```
# show replication-snapshot-history
Name              Snapshot History   Count      Snapshot Basename           Retention Priority
-------------------------------------------------------------------------------------------------
ReplicationGroup  disabled           0                                      never-delete
Replicator        both               4          Replicated                  never-delete

   Local Snapshot Name Creation Date/Time    Snap Data  Unique Data
   ---------------------------------------------------------------
   Replicated_000001   2024-12-10 13:38:52   44.4GB     742.3MB
   Replicated_000002   2024-12-12 13:46:10   66.2GB     0B
   Replicated_000003   2024-12-12 14:09:10   87.9GB     0B

-------------------------------------------------------------------------------------------------
```

Command Example 17. Displaying replication snapshot history

Using the SMU

To view replication snapshot history configuration, use the **Replications** tab of the detailed information panel of a volume that is a member of a replication set (**Provisioning** > **Volumes** > **Data Protection** table view > volume's slide-over > **Replications** tab); see Displaying a replication set (page 61) for information on selecting a volume that is a member of a replication set. The **Replications** tab (see Figure 40. Displaying replication snapshot history – configuration) shows the replication set name, an indication of the replication snapshot history (*Disabled*, *Secondary* or *Both*), the retention priority, the retention count, and the basename for the snapshots.



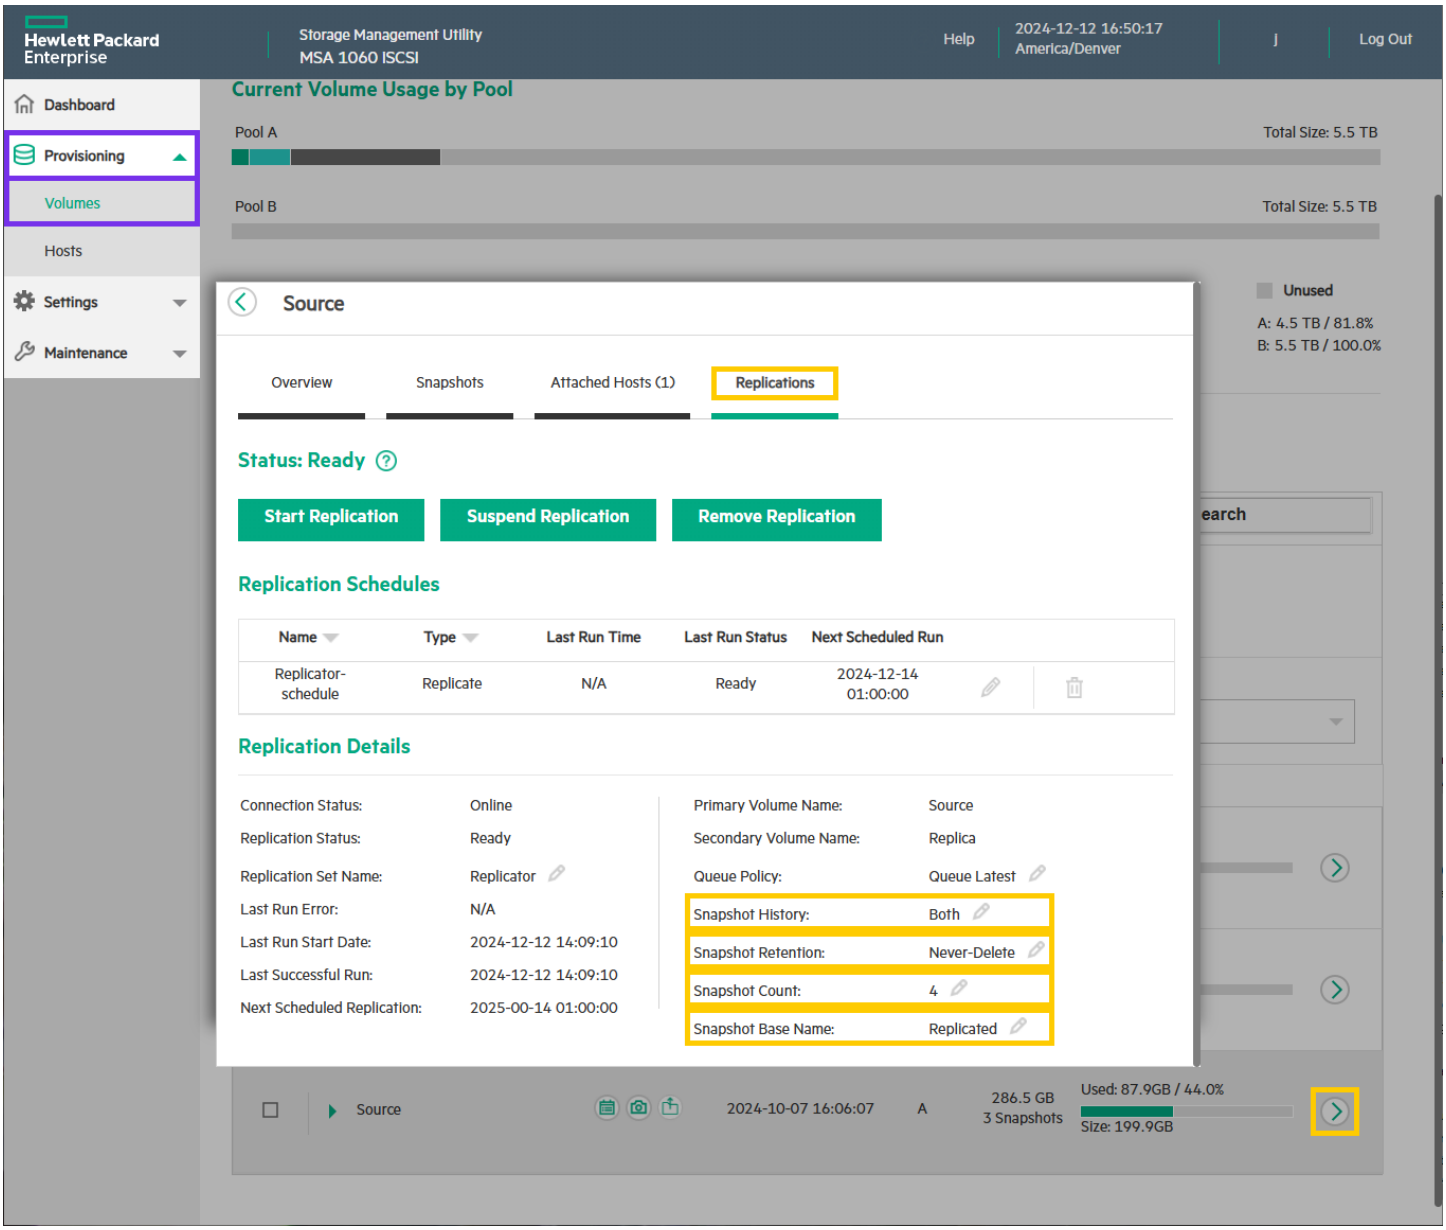


Figure 40. Displaying replication snapshot history – configuration

To view snapshots, including replication history snapshots, use the expand button ▶ to expand a volume row in the **Data Protection** table view of the **Volumes** panel (**Provisioning** > **Volumes** > **Data Protection** table view) – see Figure 42. Displaying replication snapshot history – Snapshots tab. All snapshots are listed, and there is no distinction between an ordinary snapshot and a replication history snapshot; to see only replication history snapshots, use the `show replication-snapshot-history` CLI command. The following example shows the volume row already expanded, with the collapse icon ▼ highlighted. Source-Snap is a regular snapshot, while Replicated_000001, Replicated_000002, and Replicated_000003 are replication history snapshots.

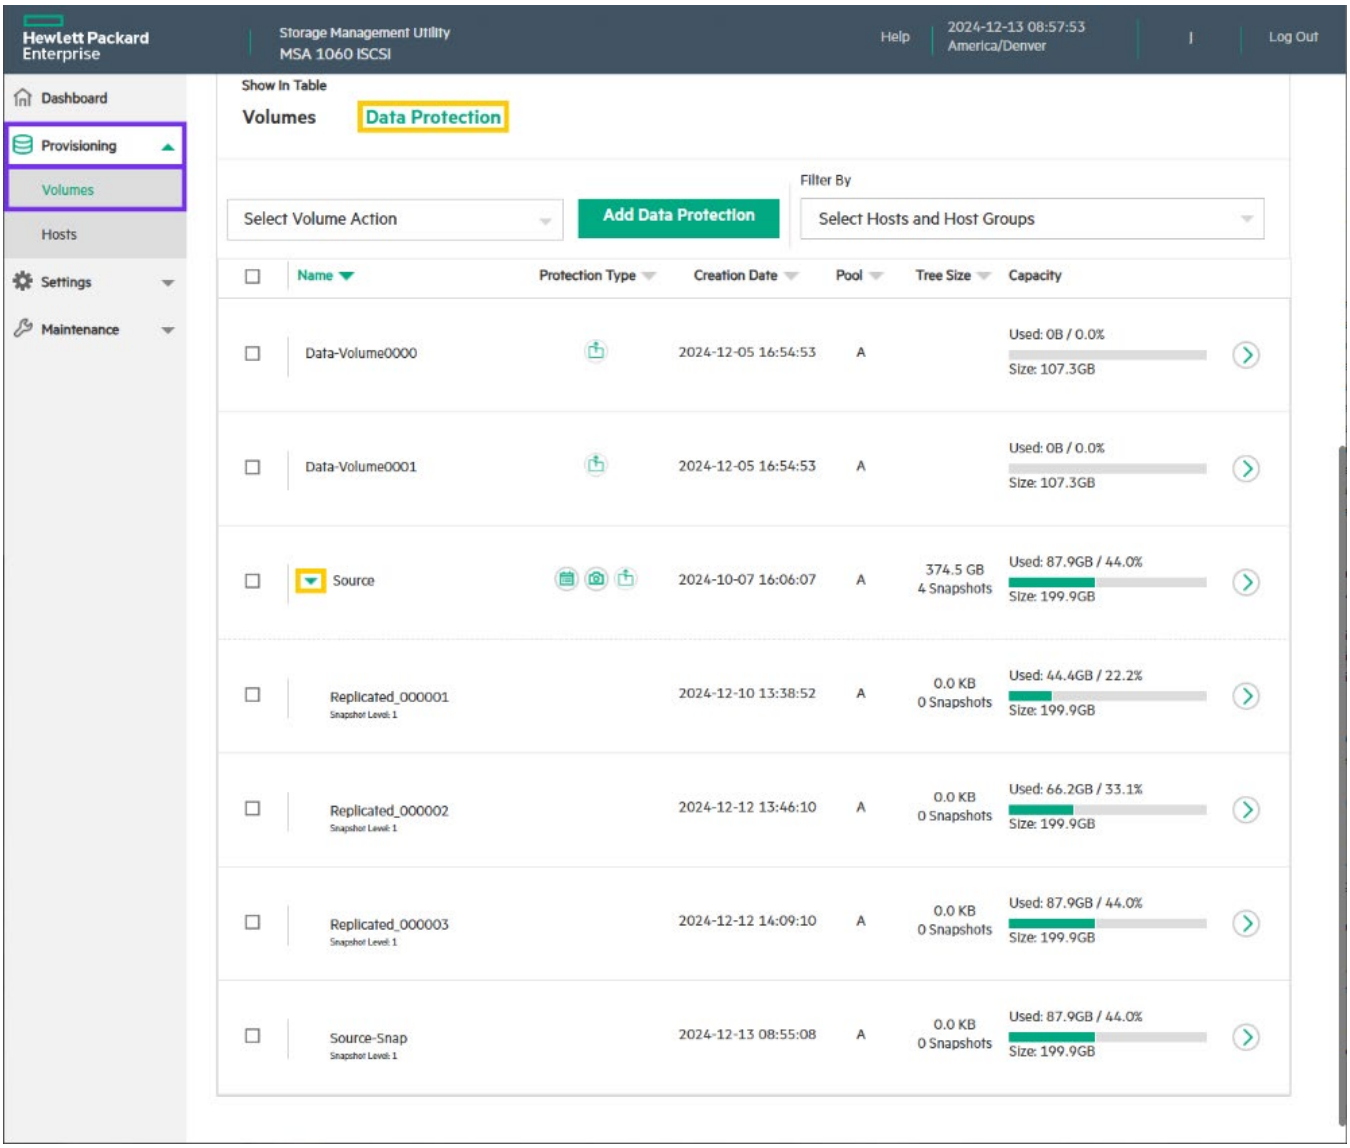


Figure 41. Displaying replication snapshot history – Data Protection table view

Alternatively, use the **Snapshots** tab of the detailed information panel of a volume that is a member of a replication set (**Provisioning** > **Volumes** > **Data Protection** table view > slide-over > **Snapshots** tab) – see Figure 42. Displaying replication snapshot history – Snapshots tab; see Using the SMU in Displaying a replication set (page 61) for information on selecting a volume that is a member of a replication set. If necessary, expand a volume's row to see the snapshots for the volume. Again, all snapshots are displayed, both ordinary snapshots and replication history snapshots. The following example shows the same snapshots as the previous example.

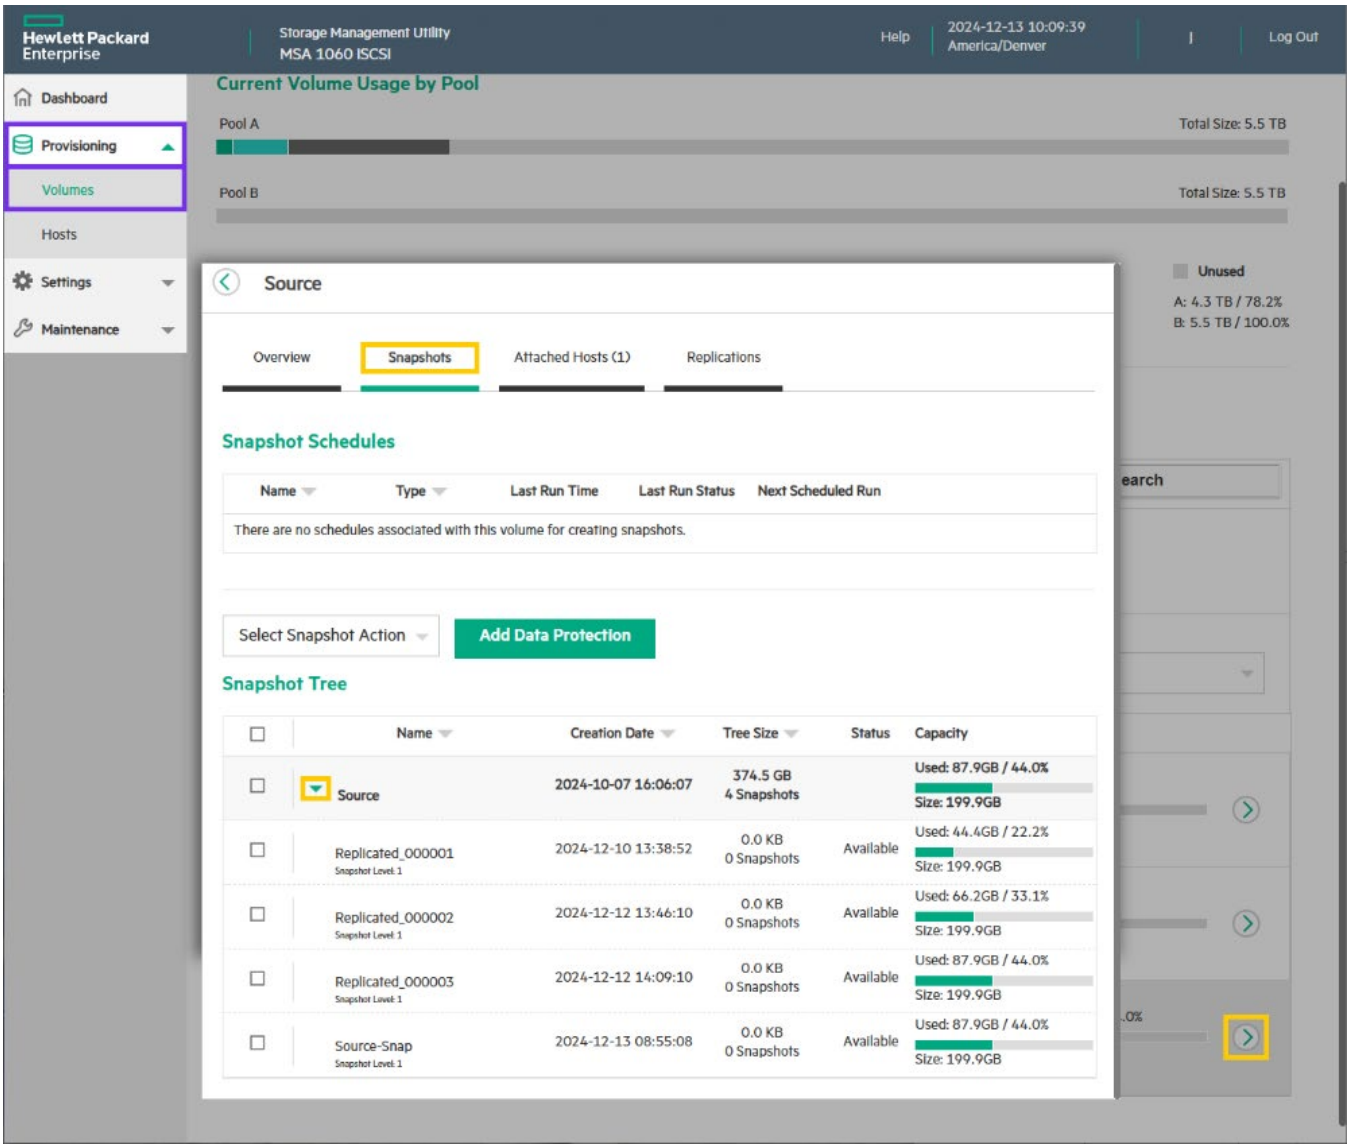


Figure 42. Displaying replication snapshot history – Snapshots tab

Deleting a replication set

Deleting a replication set deletes the replication snapshots used in the replication process but does not delete the volumes or replication history snapshots, or the peer connection used by the replication set; it does change the current secondary volume so that it can be mapped and accessed by hosts. If the peer connection used by the replication set is not functional, or if one of the arrays is not functional, you can use the `local-only` option of the `delete replication-set` CLI command to delete the replication set from the array you are logged into. If you've used the `local-only` option on one array to delete a replication set, when you delete the replication set from the other array, you must use the `local-only` option on the second array as well, even if both arrays and the peer connection that the replication set used are functioning. The `local-only` option is only available using the CLI, it is not available in the SMU.

Using the CLI

Use the `delete replication-set` CLI command. See Command Example 18. Deleting a replication set using the local-only option for an example that uses the `local-only` option.

```
# delete replication-set local-only ReplicationGroup
WARNING: Do not use the local-only parameter unless other attempts to delete the replication set
have failed. If you decide to run this command, make sure to run it on both primary and secondary
systems to completely remove the replication relationship between primary and secondary volumes.
Do you want to continue? (y/n) y
Success: Command completed successfully.
```

Command Example 18. Deleting a replication set using the local-only option

Using the SMU

Use the **Replications** tab (**Provisioning** > **Volumes** > **Volumes** table view > volume slide-over > **Replications**) or (**Provisioning** > **Volumes** > **Data Protection** table view > volume slide-over > **Replications**) and select **Remove Replication** – see Figure 43. Deleting a replication set.

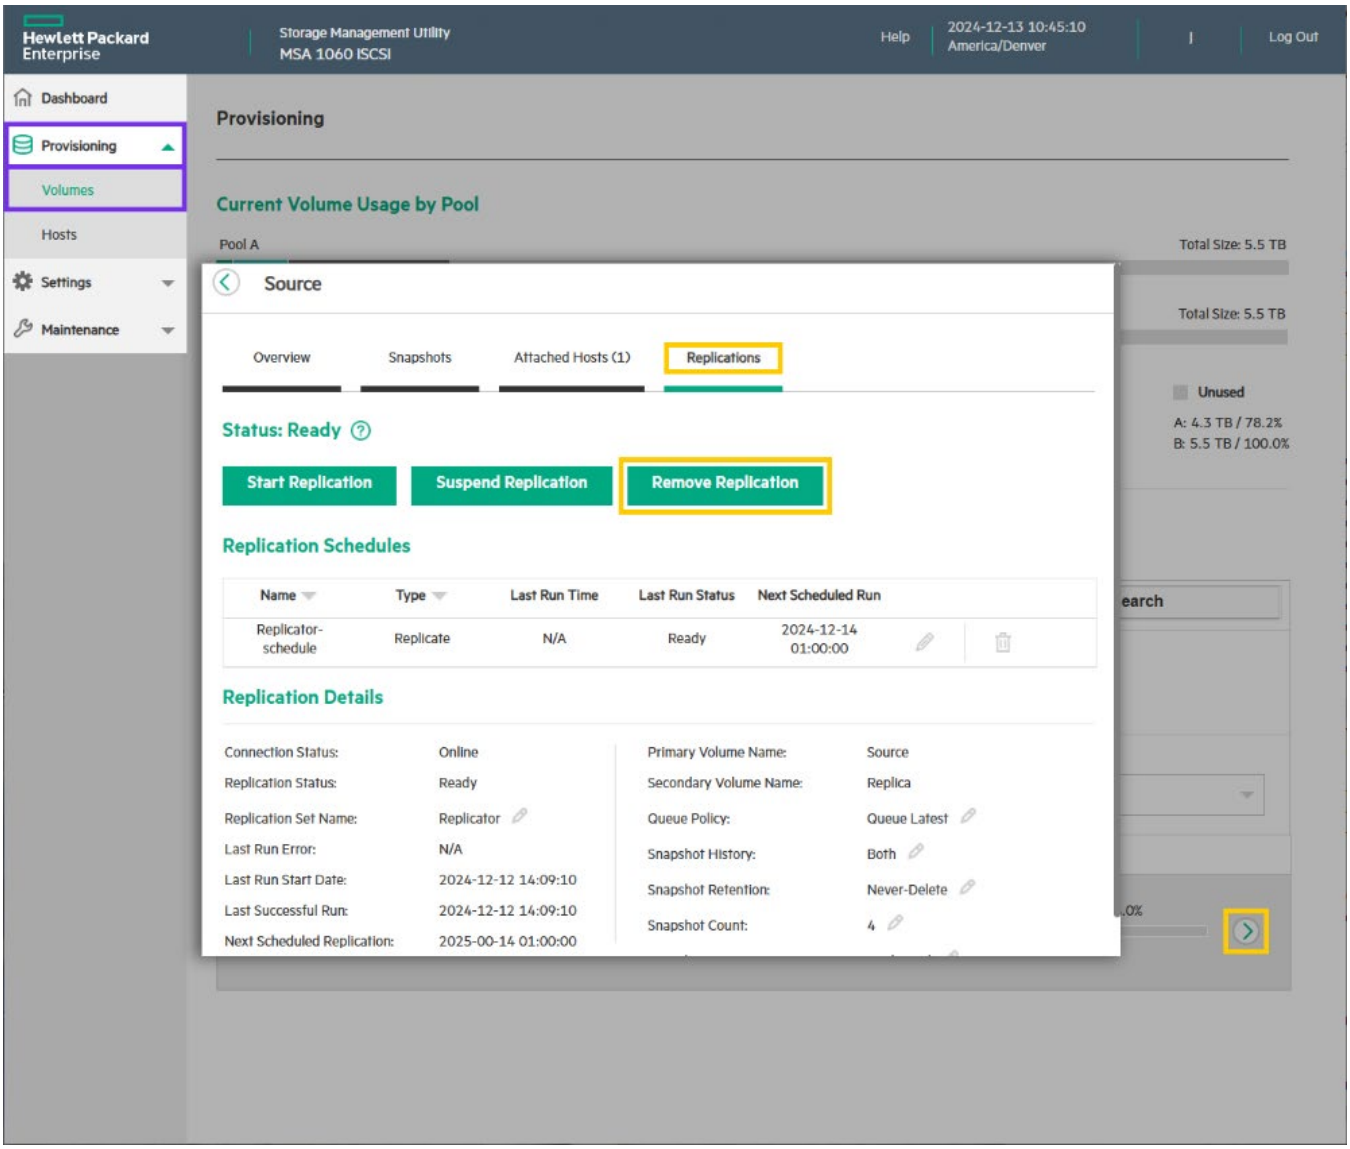


Figure 43. Deleting a replication set

Accessing the data of the secondary volume

The secondary volume cannot be accessed directly, unless the replication set's status is *Failed Over*. However, you can access the data of the secondary volume by accessing snapshots of the volume.

Creating a snapshot of a secondary volume

If you did not set up replication snapshot history for the replication set, then create a snapshot to access the data of the secondary volume.

Using the CLI

Use the `create snapshots` CLI command (see **Command Example 19.** Creating a snapshot of a secondary volume to create a snapshot of the secondary volume that you can map to hosts to access the data.

```
# create snapshots volumes Replica Replica-Access
Success: Command completed successfully. (Replica-Access) - Snapshot(s) were created.
```

Command Example 19. Creating a snapshot of a secondary volume

Using the SMU

Create snapshots using the **Data Protection Configuration** wizard. See Using the SMU in Creating a replication set (page 36) for instructions on how to access the **Data Protection Configuration** wizard; note that if using the slide-over to access the volume's details, select the **Snapshots** tab rather than the **Replications** tab. Once the wizard is available, select the **Local Snapshots** data protection type and the volumes to include – see Figure 44. Data Protection Configuration wizard – Volumes – Data Protection Type (Local Snapshots) and volume selection. Depending on the path used to access the wizard, the choice of data protection type and volumes may already be selected and cannot be changed as no other options are available.

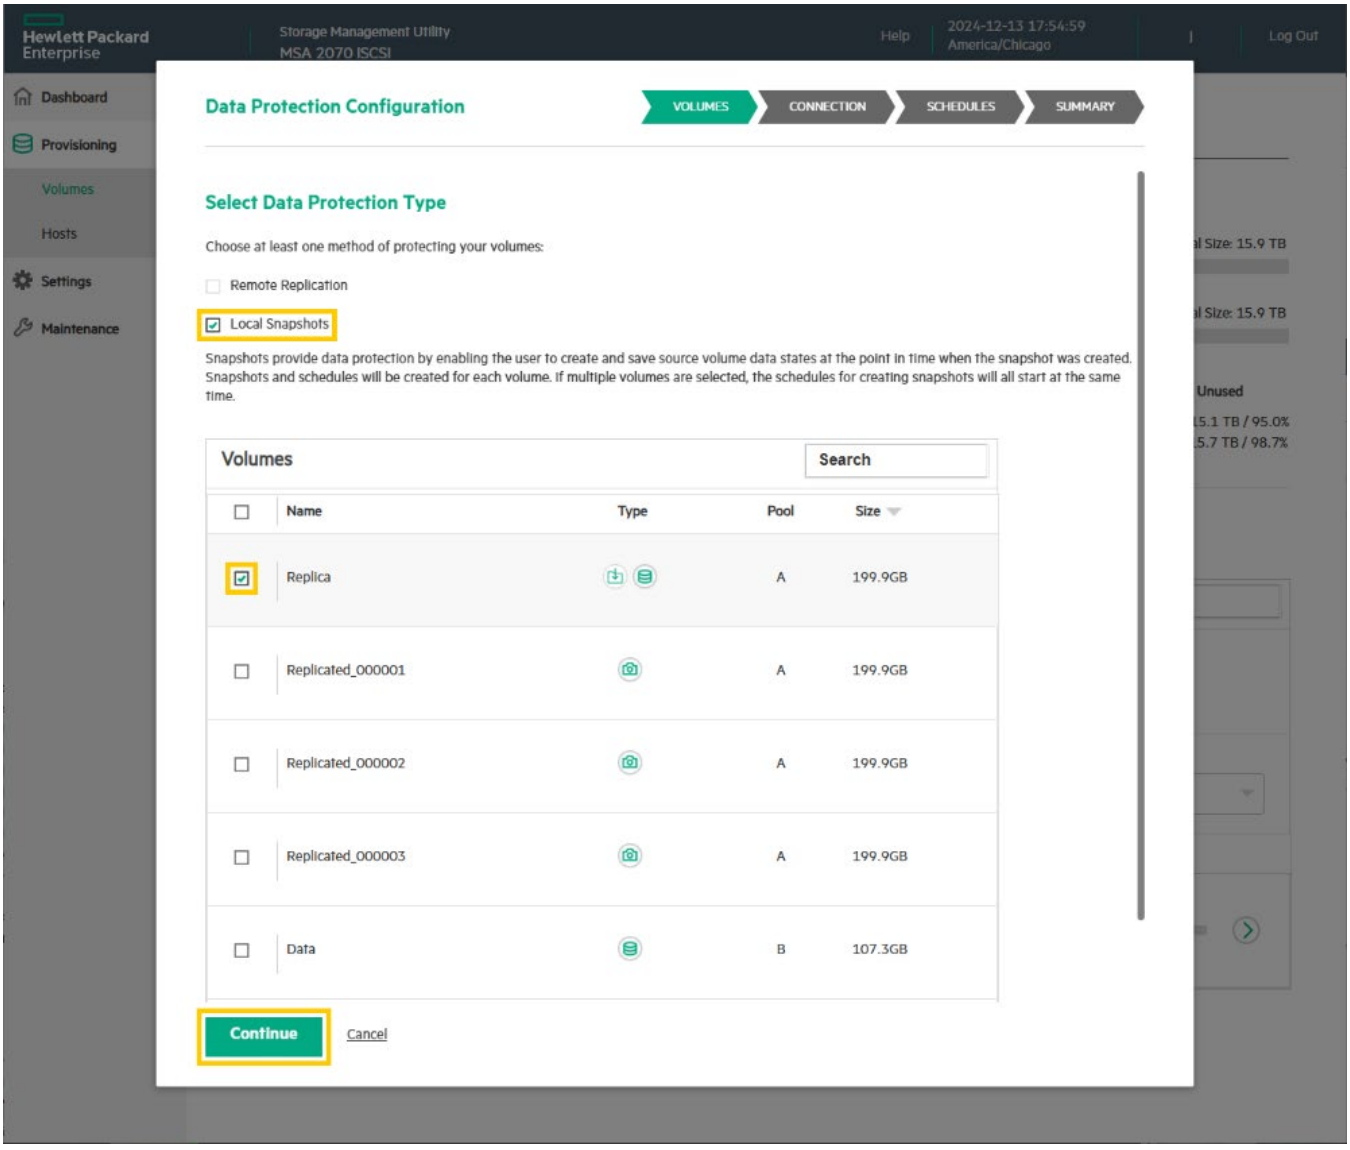


Figure 44. Data Protection Configuration wizard – Volumes – Data Protection Type (Local Snapshots) and volume selection

The next step allows you to create a snapshot immediately and configure a snapshot schedule; if creating a snapshot immediately, set the **Snapshot Name**; if creating a snapshot schedule, set the **Schedule Name**, **Retention Count** and **Snapshot Prefix**. The remainder of the schedule specification is the same as for *Remote Replication*, see the last steps of Using the SMU in Creating a replication set (page 36). Unless you need to schedule snapshots, choose **Create Snapshot Immediately (upon configuration completion)**. See Figure 45. Data Protection Configuration wizard – Schedules – Snapshot Schedule Options for an example that shows both creating a snapshot immediately and scheduling snapshots, and so all snapshot parameters are available for configuration in **Snapshot Settings**.

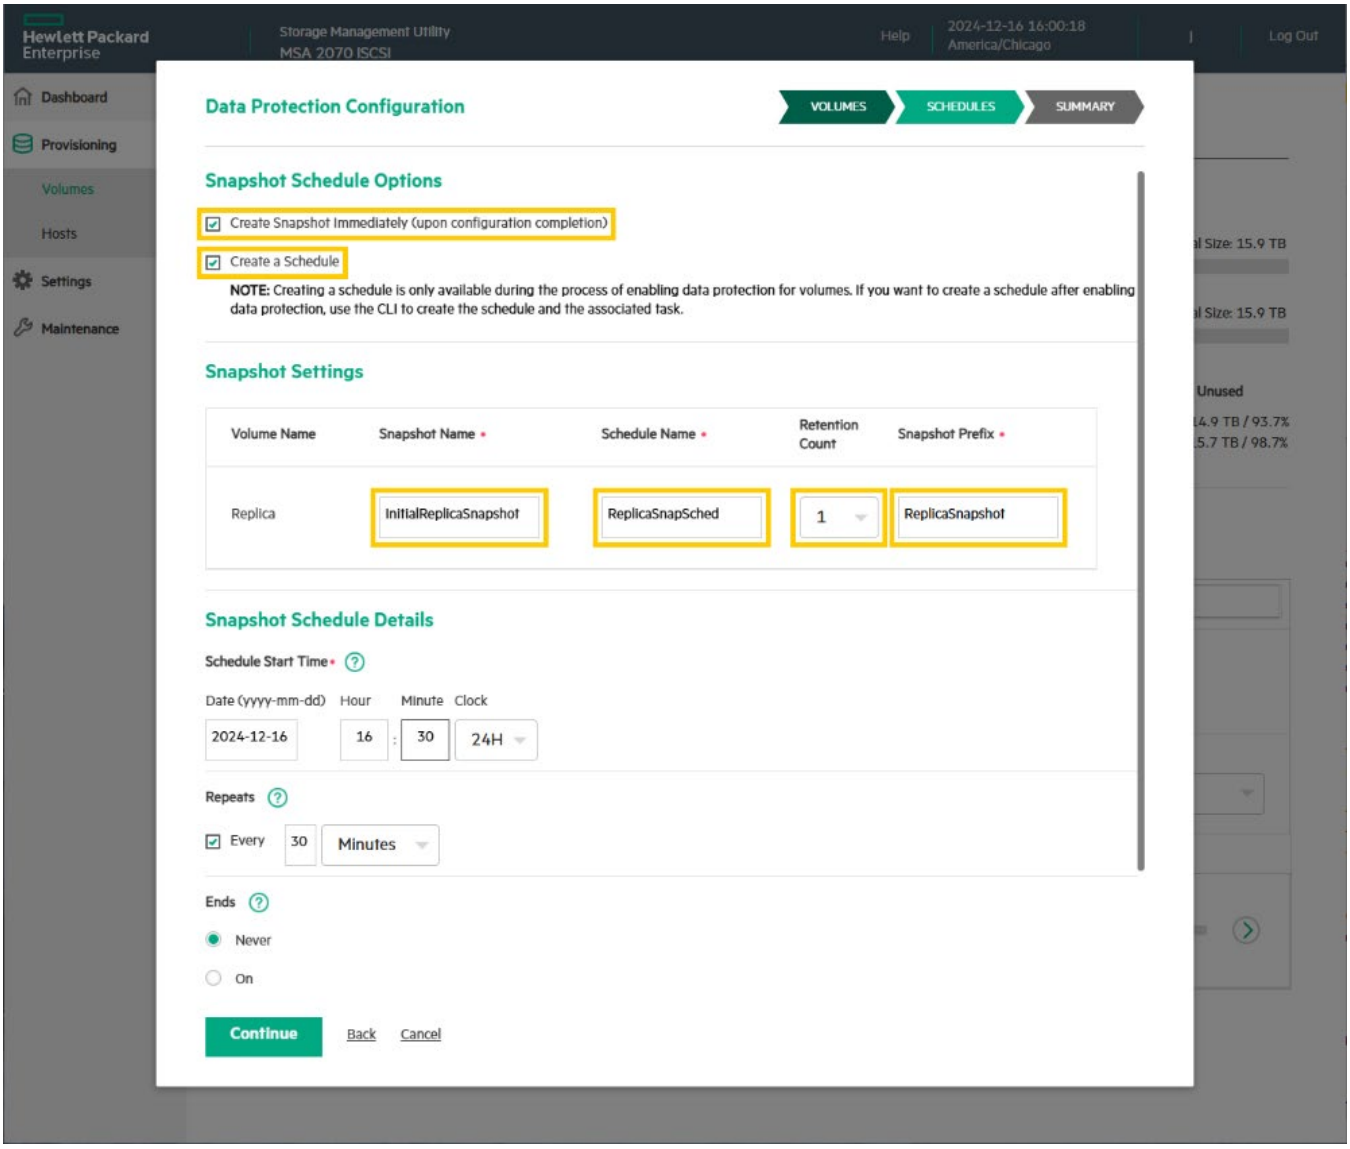


Figure 45. Data Protection Configuration wizard – Schedules – Snapshot Schedule Options

Using replication history snapshots

Setting up replication snapshot history is a great way to have access to the data of the secondary volume. Any replication history snapshot can be mapped to hosts with read-only or read-write access. This provides access to up to the number of snapshots specified by the replication set's snapshot count. Since these snapshots on the secondary volume are taken automatically at the end of the replication, they are available without the additional step of creating them.

Disaster Recovery

As mentioned in Disaster Recovery in Benefits of Remote Snap (page 8), Remote Snap provides access to data at a secondary site when the primary site experiences a critical failure. When the primary site becomes available, the modified volume data at the secondary site can be used to update the primary volumes, or the secondary volumes can become the primary volumes.

The first step in disaster recovery is planning. As mentioned in Network requirements (page 15), the rate of data change and frequency of replications will determine if the replication window can meet the recovery point objective (RPO). It bears repeating that the minimum RPO should be roughly double the replication window, assuming the replication takes the entire window to complete. In other words, the replication window should be roughly half of the RPO. This is because data that changed right after a replication begins will not be replicated and will not be available at the secondary array until the end of the next replication.

Failing the replication over to the secondary array is the first step in recovering from a disaster. Once failed over, there are three options for returning the replication set to a state where it can continue replicating. The choice of these options depends on whether the roles of primary and secondary change, and whether to replicate changes from the original secondary to the original primary volumes or discard them.

All the disaster recovery operations are performed using the CLI.

Failover

Failing over the replication set allows mapping the secondary volumes for access by hosts, while retaining the primary volume access. All current replications of the replication set cease, and new replications cannot be started. If not using replication history snapshots, take a snapshot of the secondary volume before failing the replication set over, in case you want to revert the changes made to the secondary volume during failover. Once the failover operation is complete, the replication has the status of *Failed Over*. With this status, additional replications cannot take place, as the roles of primary and secondary are in flux. Fail over the replication set on the secondary array.

Using the CLI

From the secondary array, use the `recover replication-set` command (see Command Example 20. Recover replication set – failover operation) specifying *failover* for the `operation` parameter.

```
# recover replication-set operation failover Replicator
Success: Command completed successfully.
# show replication-sets
Name        Group Primary Location  Peer     Primary Volume Secondary Volume Policy       Queue Count Status
  Failback In Progress Last Successful Run  Last Status
-------------------------------------------------------------------------------------------------------------------
Replicator  No    remote            Den-Chi  Source         Replica          Queue Latest 0           Failed Over
  False               2024-12-15 08:46:26  N/A

   Replication Progress Data Transferred  Start Time           End Time             Run Error
   ---------------------------------------------------------------------------------------
   Last Run    100%     0B                2024-12-15 04:46:24  2024-12-15 04:46:29  N/A


   Replication Progress Data Transferred  Start Time           Estimated Completion Run Error
   ---------------------------------------------------------------------------------------
   Current Run 100%     0B                2024-12-15 08:46:26  N/A                  N/A

-------------------------------------------------------------------------------------------------------------------
Success: Command completed successfully.
```

Command Example 20. Recover replication set – failover operation

Once failed over, map the secondary volumes to hosts – see Command Example 21. Mapping the secondary volumes.

```
# map volume Replica lun 43 initiator msa-r5c1-vm3
Info: Mapping succeeded. Initiator iqn.1994-05.com.redhat:msa-r5c1-vm3 was mapped for volume
Replica with LUN 43. (Replica)
Success: Command completed successfully. - The volume(s) were mapped successfully.
```

Command Example 21. Mapping the secondary volumes

Failback restore

The first option of taking the replication set out of failover copies changes made on the secondary while in failover back to the primary. Only available with a *Failed Over* status, this operation restores the original roles to the primary and secondary volumes. Perform this operation first on the secondary array to unmap both the primary and secondary volumes and replicate the modified data from the secondary volumes back to the primary volumes. Any data written to the primary volumes since the last replication is discarded.

Using the CLI

Use the `recover` CLI command with the *failback-restore* value for the `operation` parameter – see Command Example 22. Recover replication set – failback restore.



```
# recover replication-set operation failback-restore Replicator
WARNING: This may overwrite changes made to volumes during failover. It will unmap volumes associated with both the
primary and secondary volumes and may discard outstanding I/O from hosts. At the prompt, if you want to continue with
this operation enter "continue", or if you want to abort the operation and leave the replication set in the failover
state enter "abort".
> continue
If you continue with this operation, you risk losing data. Enter "accept" at the prompt if you intend to accept this
risk and continue with the operation or enter "abort" to abort the operation and leave the replication set in the
failover state.
> accept
Success: Command completed successfully.
# show replication-sets
Name        Group Primary Location  Peer     Primary Volume Secondary Volume Policy       Queue Count Status
  Failback In Progress Last Successful Run  Last Status
-------------------------------------------------------------------------------------------------------------
Replicator  No    remote            Den-Chi  Source         Replica          Queue Latest 0           Running
  True                2024-12-15 08:46:26  Success

   Replication Progress Data Transferred  Start Time           End Time             Run Error
   ---------------------------------------------------------------------------------------
   Last Run    100%     0B                2024-12-15 08:46:26  2024-12-15 08:46:31  N/A


   Replication Progress Data Transferred  Start Time           Estimated Completion Run Error
   ---------------------------------------------------------------------------------------
   Current Run 37%      1892.8MB          2024-12-20 12:57:24  2024-12-20 12:58:15  N/A

-------------------------------------------------------------------------------------------------------------
Success: Command completed successfully.
```

Command Example 22. Recover replication set – failback restore

Once the replication is complete, run the command again from the primary to take the replication back to the original roles of primary and secondary and allow mapping of the primary volumes – see Command Example 23. Recover replication set – failback restore – completion.

```
# show replication-sets
Name        Group Primary Location  Peer     Primary Volume Secondary Volume Policy       Queue Count Status
  Failback In Progress Last Successful Run  Last Status
-------------------------------------------------------------------------------------------------------------
Replicator  No    local             Den-Chi  Source         Replica          Queue Latest 0           Ready
  True                2024-12-20 14:50:34  Success

   Replication Progress Data Transferred  Start Time           End Time             Run Error
   ---------------------------------------------------------------------------------------
   Last Run    100%     0B                2024-12-15 10:39:36  2024-12-15 10:39:41  N/A


   Replication Progress Data Transferred  Start Time           Estimated Completion Run Error
   ---------------------------------------------------------------------------------------
   Current Run 100%     14.5GB            2024-12-20 14:50:34  N/A                  N/A

-------------------------------------------------------------------------------------------------------------
Success: Command completed successfully.
# recover replication-set operation failback-restore Replicator
This operation will finish the Failback Restore operation.
Are you sure you want to continue? (y/n) y
Success: Command completed successfully.
```

Command Example 23. Recover replication set – failback restore – completion

Once the replication set has been failed back, remap the primary volumes to hosts -see Command Example 24. Recover replication set – failback restore – remapping the primary volumes.

```
# map volume Source lun 23 initiator msa-r5c1-vm3
Info: Mapping succeeded. Initiator iqn.1994-05.com.redhat:msa-r5c1-vm3 was mapped for volume Source
with LUN 23. (Source)
Success: Command completed successfully. - The volume(s) were mapped successfully.
```

Command Example 24. Recover replication set – failback restore – remapping the primary volumes

Failback no restore

The second option for taking a replication set out of failover does not copy changes made to the secondary volumes while in failover back to the primary – the changes are discarded. Only available with a *Failed Over* status, this operation restores the original roles to the primary and secondary volumes. If the secondary volumes were mapped while the replication set was *Failed Over*, they are unmapped. Since this will not



copy changes made to the secondary volumes while failed over, no replication from the secondary to the primary occurs. Changes made to the secondary volumes remain until the next replication, at which time they are discarded.

Using the CLI

Use the `recover` CLI command with the *failback-no-restore* value for the `operation` parameter – see Command Example 25. Recover replication set – failback no restore.

```
# recover replication-set operation failback-no-restore Replicator
WARNING: This may overwrite changes made to volumes during failover. It will unmap volumes
associated with the secondary volume and may discard outstanding I/O from hosts. At the prompt, if
you want to continue with this operation enter "continue", or if you want to abort the operation
and leave the replication set in the failover state enter "abort".
> continue
If you continue with this operation, you risk losing data. Enter "accept" at the prompt if you
intend to accept this risk and continue with the operation or enter "abort" to abort the operation
and leave the replication set in the failover state.
> accept
Success: Command completed successfully.
```

Command Example 25. Recover replication set – failback no restore

Once the command completes, the replication set is ready to continue replicating in the original direction without further intervention. Note the differences in allocated sizes between the primary volume, Source, and the secondary volume, Replica. This is due to data written to the secondary volume when the replication set was failed over.

From the secondary array:

```
# show volumes
Pool Name     Total Size Alloc Size Type Health Reason Action
-------------------------------------------------------------
B    Data     107.3GB    0B         base OK
B    Logs     107.3GB    0B         base OK
A    Replica  199.9GB    116.9GB    base OK
-------------------------------------------------------------
Success: Command completed successfully.
```

From the primary array:

```
# show volumes
Pool Name    Total Size Alloc Size Type Health Reason Action
------------------------------------------------------------
A    Source  199.9GB    102.4GB    base OK
------------------------------------------------------------
Success: Command completed successfully.
```

Reverse

The third option of taking the replication set out of failover simply reverses the direction of the replication. Only available when the replication set's status is *Failed Over*, this operation reverses the roles of the primary and secondary volumes – the original primary volume becomes the secondary volume, and the original secondary volume becomes the primary volume. This operation also unmaps the original primary volumes (now secondary) and allows the original secondary volumes (now primary) to be mapped to hosts. Any data written to the primary volumes while the replication set was failed over is discarded on the next replication.

Using the CLI

Use the `recover` CLI command with the *reverse* value for the `operation` parameter – see Command Example 26. Recover replication set – reverse.



```
# recover replication-set operation reverse Replicator
WARNING: This may overwrite changes made to volumes during failover. It will unmap volumes
associated with the current primary volume and may discard outstanding I/O from hosts. At the
prompt, if you want to continue with this operation enter "continue", or if you want to abort the
operation and leave the replication set in the failover state enter "abort".
> continue
If you continue with this operation, you risk losing data. Enter "accept" at the prompt if you
intend to accept this risk and continue with the operation or enter "abort" to abort the operation
and leave the replication set in the failover state.
> accept
Success: Command completed successfully.
```

Command Example 26. Recover replication set – reverse

Use cases

Multiple remote offices backup to one central office

With the capability of up to 4 peer connections, one HPE MSA 2060/2062 Storage array or HPE MSA 2070/2072 Storage array can backup 4 remote HPE MSA Gen6 Storage arrays or HPE MSA Gen7 Storage arrays.

Application consistent replications and VSS

When snapshots are taken manually, no I/O quiescing is done at the array level. Use the `create snapshots` command to take time-consistent snapshots across multiple volumes once the associated applications have been quiesced. Then, replicate using the `last-snapshot` parameter of the `replicate` CLI command.

Software such as the Microsoft VSS framework enables quiescing applications and taking application—consistent snapshots.

Best practices

Fault tolerance

While the ports used for replication must be connected to a switch, connect the ports to multiple switches to avoid a single point of failure at the switch. For example, connect ports A1 and B1 to one switch, and A2 and B2 to another. In case of link failure, the array will attempt to resume the replication every 10 minutes for the first hour, then every hour until the replication resumes. You can also attempt to resume the replication manually, or abort, as desired.

Only dual controller mode is supported for HPE MSA Gen6 Storage and HPE MSA Gen7 Storage; this avoids a single point of failure at the controller level.

Volume size and policy

Review the Snapshot space section (page 14). When setting up the pools for the primary and secondary volumes, consider the impact of the overcommit flag for the pool.

- If overcommit is disabled, the volume and associated replication snapshots are fully provisioned, meaning the pool must be slightly over three or four times the size of the volume, depending on the queue policy for the replication set.
- If overcommit is enabled, then the size of the pool should be sufficient for the amount of data in the volume, and for two sets of data changes. The snapshot space can be managed using the `set snapshot-space` CLI command. The default snapshot space limit is 10% of the pool, and the default limit policy is *Notify Only*.

In addition to the overcommit flag, consider the impact of replication history snapshots on the pool. The replication history snapshot count applies to the secondary, and if enabled, to the primary array.

The size of the primary volume can be increased after creating the replication set, and the size of the pools that contain the volumes can change as well, increasing by adding disk groups, or decreasing by removing disk groups.



Scheduling

To replicate a manually taken snapshot or to replicate a snapshot taken by another process, for example VSS, select **Last Snapshot** when creating a schedule using the SMU, or specify `last-snapshot` when creating a task using the CLI. This helps in achieving application-consistent snapshots. This option is not available for replications of volume groups.

Schedules are only associated with the primary volumes. You can see the status of a schedule on the **Replications** tab of the primary volume's details (**Provisioning** > **Volumes** > **Volumes** table view > volume slide-over > **Replications** tab) or (**Provisioning** > **Volumes** > **Data Protection** table view > volume slide-over > **Replications** tab), or by using the `show schedules` CLI command. Use the same tab in the SMU to edit the schedule by clicking the edit icon, or use the `set schedule` CLI command. Replications can be scheduled at most once per half hour.

Consider the rate of data change, network capability, and host I/O when scheduling replications to minimize a replication attempting to begin before the previous replication has completed, and especially before the previously queued replication has had a chance to start.

Application-consistent snapshots

When snapshots are taken manually or taken outside of the array's management interfaces, no I/O quiescing is done at the array level – verify the application is quiesced before taking the snapshots. The VSS framework available in Microsoft operating systems provides quiescing the application and taking application-consistent snapshots.

If not employing a framework that quiesces applications and takes application-consistent snapshots, then use the `create snapshots` CLI command or the **Data Protection Configuration** wizard once the application is quiesced to create the application snapshots. See Using the SMU in Creating a replication set on page 36 for instructions on how to engage the wizard; once the wizard is available, select **Local Snapshots** rather than **Remote Replication**.

Select **Last Snapshot** in the SMU when creating the schedule or use the `last-snapshot` parameter in the `create task` CLI command to replicate the application-consistent snapshots rather than taking a new snapshot at the start of replication and replicating the new snapshot.

Performance tips

For a gain in replication and host I/O performance of up to 20%, enable jumbo frames on all infrastructure components in the path and on iSCSI controllers. Jumbo frames are disabled by default for the iSCSI host ports. You can enable them using either the SMU (**Settings** > **iSCSI** > **Configuration** tab) or the CLI (`set iscsi-parameters`).

Note

If your infrastructure does not support jumbo frames, enabling them only on your controllers may actually lower performance or even prevent the creation of replication sets or replications.

Troubleshooting

Peer connections

Licensing

Verify both the primary and secondary arrays have a license that enables replication. Use the CLI command `show license` or the **Install License** panel (**Maintenance** > **Support** > **Licensing**) in the SMU to verify replication is licensed.

CHAP

If using CHAP, enable it on both arrays and use the same secret on both arrays to avoid CHAP authentication failures. Use the CLI command `show peer-connections` with the `verify-links` parameter to check for issues with the peer connection, then use the `show chap-records show-secrets` and `show iscsi-parameters` CLI commands on each array, or use the **CHAP** panel (**Settings** > **iSCSI** > **CHAP**) in the SMU on each array and verify CHAP enablement and that the secrets match. Remember that the name of the CHAP record is the IQN of the other array, which can be obtained using the `show ports` CLI command on the other array.

Replications

Replication suspended

Check the peer connection using the `verify-links` parameter of the `show peer-connections` CLI command.

Check the CHAP settings – see the CHAP section of the Peer connections section of the Troubleshooting section (page 76).

If overcommit is enabled for the pool on either side of the replication, check if the high threshold has been exceeded. Use the `show pools` CLI command or the **Storage** panel (**Maintenance** > **Storage**) and expand the pool using the expand icon and expand the pool row to see if the pool is overcommitted and the percent the high threshold is set at, and the available space to see if there is sufficient space to continue. Add disk-groups to the pool or remove volumes or snapshots if necessary to increase the available size of the pool. This condition may occur on either the primary or secondary array, so check both sides of the replication.

Unable to control or modify a replication set

The primary array controls the replication set – initiating a replication, suspending a replication set, aborting a replication, or resuming a replication set must all be performed on the primary array, along with creating or modifying tasks and schedules associated with a replication set. Recovery of a replication set must begin on the secondary array and all failback or reverse operations are also initiated from the secondary array, with the second phase of failback restoring from the secondary to the primary occurring on the primary array. You can change the replication set parameters, such as the name, snapshot history settings, and queue policy from either the primary or secondary array.

Errors

The format for this section has a general heading for the component, subsections for specific errors with error messages seen in the SMU or CLI, as applicable, and how to resolve the error.

Component heading

Error

SMU error messages (if applicable)

CLI error messages (if applicable)

Resolution

Licensing

Replication not licensed

SMU

Failed to create replication set

Replication is not licensed.

CLI

```
Error: Replication is not licensed.
```

Resolution

Verify replication is licensed on both the primary and secondary arrays.

Pools

Pool nonexistent on peer

SMU

Failed to create replication set

The operation cannot be completed because the requested virtual pool does not exist on the peer array.

CLI

```
Error: The operation cannot be completed because the requested virtual pool does not exist on the
peer system.
```

Resolution

Specify a pool that does exist on the secondary array, or create the pool on the secondary array.

Remote pool insufficient space

SMU

Failed to create replication set

The requested size for a new volume exceeds the remaining physical space left in the pool. Failed to create the volume.

CLI

```
Error: The requested size for a new volume exceeds the remaining physical space left in the pool.
Failed to create the volume.
```

Resolution

Check that the available space in the remote pool is sufficient for the volume and two (if the queue policy is set to discard) or three snapshots (if the queue policy is set to queue latest).

Pool has exceeded high threshold

SMU

Failed to resume the replication.

The requested operation cannot be completed because the pool has exceeded its high threshold.

CLI

```
Error: The requested operation cannot be completed because the pool has exceeded its high
threshold.
```

Resolution

To resume the replication, add and configure disks to the array or remove unwanted snapshots or volumes. The high threshold is automatically calculated as the available capacity of the pool minus 400GiB — or if that result is below the mid threshold, then mid threshold capacity plus 1%.

Deleting a lone pool

SMU

You must delete all peer connections before the last remaining pool on the array may be deleted.

CLI

```
Error: You must delete all peer connections before the last remaining pool on the system may be
deleted.
```

Resolution

To create a peer connection, at least one pool must be available on the array. Thus, if a peer connection exists, a pool must exist, and the pool cannot be deleted. Delete the peer connection and retry deleting the last remaining pool on the array.

Deleting a pool containing a replication set

CLI

Error: The virtual pool cannot be deleted because it contains one or more replication sets.

Resolution

Delete the replication sets before deleting the pool.



Unhealthy pool

SMU

Unable to initiate replication.

The requested operation cannot be completed because the virtual pool is unhealthy.

CLI

```
Error: The requested operation cannot be completed because the virtual pool is unhealthy.
```

Resolution

Using the troubleshooting guide for your array, return the pool to a healthy, or, at a minimum, degraded, state.

CHAP

Authentication failure on remote array

CLI

```
Error: The operation cannot be completed due to a CHAP authentication failure on the remote system.
Either the remote system is not configured to allow access to this system, or this system is not
configured with the proper secret.
```

Resolution

Configure CHAP using the guidelines in Peer connections and CHAP (page 30).

Peer Connection

Invalid remote address

CLI

```
Error: Invalid remote address specified.
```

Resolution

Specify a valid iSCSI IP address or FC WWN when creating the peer connection.

Too many remote addresses specified

CLI

```
Error: Too many remote addresses were specified.
```

Resolution

While the `create peer-connection` CLI command accepts multiple addresses on the command line separated by spaces, only one address is actually allowed.

Invalid username or password

SMU

Failed creating <peer-connection-ID>

Invalid remote username or password specified.

CLI

```
Error: Invalid remote username or password specified.
```

Resolution

Verify the username and password are valid for the remote array.

Peer connection already exists

SMU

Failed creating peer <peer-connection-ID>

A duplicate name was specified.



CLI

```
Error: A duplicate name was specified.
```

Resolution

Select a name for the peer connection that does not already exist on both the primary and secondary arrays.

Peer connection not found

CLI

```
Error: The specified peer connection was not found.
```

Resolution

When creating a replication set, verify the peer connection is available using the `show peer-connections` CLI command.

Remote address already exists in a peer connection

SMU

Failed creating peer <peer-connection-name>

The remote address specified is already active in a peer connection.

CLI

```
Error: The remote address specified is already active in a peer connection.
```

Resolution

Either use the peer connection that already contains the address or, if the existing peer connection is not used, delete it and create a new peer connection with the remote address.

Cannot communicate with peer

SMU

Unable to initiate replication.

The operation cannot be completed because communication cannot be established with the peer.

CLI

```
Error: The operation cannot be completed because communication cannot be established with the peer.
```

Resolution

Verify peer links with the `show peer-connections verify-links` CLI command and repair the connections as necessary.

Peer connection currently unavailable

SMU

Could not modify peer connection <peer-connection-id>.

The operation cannot be completed because the specified peer connection is currently unavailable.

CLI

```
Error: The operation cannot be completed because the specified peer connection is currently
unavailable.
```

Resolution

You can only modify the peer connection's name if its connection status is online. You can however set the peer connection's remote port address to reflect changes in the network when the connection status is offline.



Peer connection unreachable by all local ports

CLI

```
Error: The remote ports of at least one peer connection are unreachable by all local ports. This includes local ports that are not participating in the peer connection. - Unreachable peer connections: <peer-connection-id>
```

Resolution

Repair the connections to the peer, and test with the `show peer-connections verify-links` CLI command.

Peer connection in use

CLI

```
Error: The operation cannot be completed because at least one side of the peer connection has a replication set that is using the peer connection.
```

Resolution

Verify no replication sets are using the peer connection before attempting to delete it.

Limits

Maximum replication sets

SMU

Failed to create replication set

Maximum number of replication sets are already in use.

CLI

```
Error: Maximum number of replication sets are already in use.
```

Resolution

The limit of 32 replication sets applies to both the primary and the secondary array. If the primary array this error occurred on does not have 32 replication sets, check if the secondary array has met this limit.

Maximum replication volumes

SMU

Failed to create replication set

The maximum number of volumes that can be replicated and/or copied has been exceeded on this system.

CLI

```
Error: The maximum number of volumes that can be replicated and/or copied has been exceeded on this system.
```

Resolution

The maximum number of replication volumes, both primary and secondary, are limited to 32. View the volumes already involved in replications via the `show volumes detail` CLI command and look for volumes whose `Role` is `Primary` or `Secondary`; or, from the **Data Protection** table of the **Volumes** panel (**Provisioning** > **Volumes**), look for the primary volume  or secondary volume  icons.

Secondary volume

Cannot specify secondary volume group name

CLI

```
Error: The secondary volume name cannot be specified when the primary volume is a group.
```

Resolution

When creating a replication set of a primary volume group, allow the remote array to create the secondary volume group with the same volume group name and create volumes with the same names as the primary volume group members.



Attempt to allow modification of secondary volumes

SMU

Example 1

Mapping failed for volume <volume-name> to host <host-name>.

This operation is not supported on the secondary volume in a virtual replication relationship.

Example 2

Unable to rollback volume <volume-name>.

This operation is not supported on the secondary volume in a virtual replication relationship.

CLI

```
Error: This operation is not supported on the secondary volume in a virtual replication
relationship.
```

Resolution

To keep track of changes made on the primary volume, changes are not permitted on the secondary volume except through replication; therefore, you cannot map or rollback the secondary volume unless the replication set is failed over.

Secondary volume exists

SMU

Failed to create replication set

The volume name already exists in the remote array.

CLI

```
Error: The volume name already exists in the remote system.
```

Resolution

When creating the replication set, specify a volume name not currently in use in the remote array for the secondary volume, if possible, or change the names of the primary volumes so that the new names do not already exist in the remote array.

Volume groups

Modifying volume group

CLI

Example 1

```
Error: This operation is not supported on a volume group active in a replication set. - Failed to
delete one or more volume groups.
```

Example 2

```
Error: This operation is not supported on a volume group active in a replication set.
(<replication-set-name) - The volume group member(s) were NOT added.
```

Example 3

```
Error: This operation is not supported on a volume group active in a replication set.
```

Example 4

```
Error: You cannot rename a volume group that is part of a replication set.
```

Resolution

You cannot delete, add members to, remove members from, or change the name of a volume group that is involved in a replication.



Membership in only one group

CLI

```
Error: Group members may only belong to a single group.
```

Resolution

A volume can be a member of only one group; check the currently existing volume groups for the volume.

Too many volumes for replication

CLI

```
Error: The operation cannot be completed because the number of volumes in the volume group exceed
the maximum allowable limit for virtual replication.
```

Resolution

Reduce the number of volumes in the volume group with the `remove volume-group-members` CLI command.

Cannot use last snapshot option for volume group

CLI

Example 1:

```
Error: Failed to create the task. (<task-id>) - Cannot use the last-snapshot option for a volume
group in a replication set.
```

Example 2:

```
Error: This operation is not supported on a volume group active in a replication set.
```

Resolution

Do not attempt to replicate the “last snapshot” of a volume group – you cannot create a snapshot as an atomic operation on the volume group, you would have to create individual snapshots of the individual volumes, and so cannot use the snapshots to dependably create a crash-consistent or application-consistent replication.

Replication set

Replication set name already exists

CLI

```
Error: The replication set name already exists in the system.
```

Resolution

Choose a different name for the replication set you are attempting to create.

Volume already in replication set

CLI

```
Error: The volume is already in a replication set.
```

Resolution

Specify a volume that is not currently in a replication set.

Replication set not found

CLI

```
Error: The replication set was not found.
```

Resolution

The array could not find a replication set with the name specified. Try executing the CLI command `show replication-sets` to verify the replication set name.



Basename too long

CLI

```
Error: The specified basename is too long.
```

Resolution

Specify a basename from 1 to 24 bytes in length.

Retention count too high

SMU

Failed to modify the replication set.

The retention count must be equal or greater than the current number of remotely retained snapshots. Delete snapshots on peer array in excess of new count.

CLI

Example 1:

```
Error: The retention count must be equal or greater than the current number of retained snapshots
for both the primary and secondary systems.  Delete snapshots in excess of new count.
```

Example 2:

```
Error: The retention count must be equal or greater than the current number of remotely retained
snapshots.  Delete snapshots on peer system in excess of new count.
```

Resolution

The array will not delete existing replication history snapshots on either the primary or secondary array to certify the count is less than that specified. Therefore, if fewer replication history snapshots are desired, manually delete snapshots you no longer need on both the primary and secondary arrays until the count is at or below the desired count, and then change the replication history snapshot retention count.

Retention count out of range

CLI

```
Error: The retention count value is not in the allowed range. Please check help for the allowed
range.
```

Resolution

Specify a retention count, `snapshot-count` in the CLI, between 1 and 16.

Replication operation on wrong array

CLI

Example 1

```
Error: This operation must be issued from the system that owns the primary volume in this
replication set.
```

Example 2

```
Error: Failed to create the task. (<task_name>) - This schedule must be created on the system that
owns the primary volume in this replication set.
```

Resolution

The array that owns the primary volume controls replications. Therefore, to alter a replication operation, run the command on the array that owns the primary volume. This includes starting, suspending, aborting, and resuming a replication set.

Operation not allowed on running replication

CLI

Example 1



```
Error: The operation cannot be completed on a replication set with an active replication.  You must abort the replication or wait until it completes to successfully issue this command.
```

Example 2

```
Error: The requested replication action cannot be completed on a replication set that is currently running a replication.
```

Resolution

You cannot delete a replication set with an ongoing replication – abort the replication or wait until it is complete.

Operation not allowed when replication is not running

CLI

```
Error: The operation cannot be completed because no replication is currently running on this replication set.
```

Resolution

You can only abort a replication that is currently running.

Replication set suspended

CLI

```
Error: The requested replication action cannot be completed on a suspended replication set.
```

Resolution

You cannot start a replication on a suspended replication set. The only operations allowed on a suspended replication set are abort or resume.

Replication queued

CLI

```
Success: Command completed successfully. - Warning: A new replication was requested while an existing replication is in progress. The replication has been queued and will run after the current replication completes.
```

Resolution

This is not an error; it is simply a warning notification to let you know the replication was queued. No action is required.

No queued replication

CLI

```
Error: There are no queued replications pending, so the requested operation cannot be completed.
```

Resolution

This is not necessarily an error, but more of a notice that the array is not in a state to perform the operation – in this case, attempting to clear the replication queue when there are no queued replications.

Replication set does not exist on remote array

SMU

Unable to initiate replication.

The operation cannot be completed because the replication set is not defined on the remote system.

CLI

```
Error: The operation cannot be completed because the replication set is not defined on the remote system.
```

Resolution

This is most likely due to removing the replication set from the remote array using the `local-only` parameter of the `delete replication-set` CLI command. Delete the replication set from the local array and recreate the replication set.



Deleting primary or secondary volume

CLI

```
Error: You cannot delete a volume that is acting as the primary or secondary volume in a
replication set.
```

Resolution

Delete the replication set before attempting to delete the primary or secondary volumes.

Disaster recovery

Operation issued on wrong array

CLI

```
Error: The operation must be issued from the secondary (destination) system rather than the primary
system.
```

Resolution

Run the command on the secondary array. If attempting a failover, the secondary array may be the only one running, so it must accept the operation, while the primary array will recognize the replication set is failed over if it tries to initiate a replication.

Operation not allowed in failover mode

CLI

```
Error: The operation cannot be completed because the replication set is in failover mode.
```

Resolution

Since the direction of a replication that is in failover mode is in flux, a replication cannot occur.

Volume mapping during failback restore

SMU

Mapping failed for volume <volume-ID> to host <host-ID>.

Cannot map a volume while its replication set is in the middle of a Failback Restore operation.

CLI

```
Error: Cannot map a volume while its replication set is in the middle of a Failback Restore
operation.
```

Resolution

Complete the failback restore operation before attempting to map the volume.

Task creation during failback restore

CLI

```
Error: Failed to create the task. (<task-ID>) - Cannot create a task targeted for a replication set
that is in the middle of a Failback Restore operation.
```

Resolution

Complete the failback restore operation before attempting to create a task for the replication set.

Replicating from original primary array during failback restore

SMU

Unable to initiate replication.

Because the affected replication set is in the middle of the Failback Restore operation, this command can only be issued from the secondary array.



CLI

```
Error: Because the affected replication set is in the middle of the Failback Restore operation,
this command can only be issued from the secondary system.
```

Resolution

Complete the failback restore operation to return the replication set to its original direction, then replicate from the original, and now current, primary volumes.

Operation allowed only in failover mode

CLI

```
Error: The operation cannot be completed because the replication set is not in failover mode.
```

Resolution

Failback and reverse operations are only allowed on replication sets in failover mode.

Initial replication not completed

CLI

```
Error: The operation is not allowed because the replication set has not completed an initial
replication.
```

Resolution

You cannot failover a replication set until the initial replication has completed.

Events

Code	Severity	Summary
214	Info	Create snapshots completed
215	Info	Commit snapshots completed
248	Info	License installed
249	Info	Licensed feature installed
250	Warning	License install failed
256	Info	Snapshot prepare completed
257	Info	Snapshot prepare and commit completed
258	Info	Snapshot commit completed
361	Critical/Error/Warning/Info	Scheduler event
362	Critical/Error/Info	Task event
473	Info	Volume threshold exceeded
474	Info	Volume threshold cleared



488	Info	Volume group create started
489	Info	Volume group create completed
490	Info	Volume group create failed
491	Info	Local volume group create started
492	Info	Local volume group removed
493	Info	Local volume group modified
571	Error/Warning/Info	Snapshot space exceeded a threshold
572	Info	Snapshot space dropped below threshold
574	Info	Peer connection created
575	Info	Peer connection deleted
576	Error/Info	Replication set created
577	Error/Info	Replication set deleted
578	Info	Replication job started
579	Warning/Info	Replication job completed
580	Info	Replication job cancelled
581	Warning/Info	Replication job suspended
582	Warning/Info	Another replication request queued
583	Error/Info	Replication set reversed
584	Info	Peer connection modified
585	Info	Replication set modified
586	Error/Info	Replication set resumed
587	Info	Pending replication removed from queue
588	Warning/Info	Replication set failed over



589	Info	Replication set restored
604	Warning	Replication snapshot failure

SNMP traps, email (SMTP) notification, and syslog

The array can send SNMP traps and email notifications for and remotely log the events described above. From the **Settings** > **Notifications** page, select the appropriate tab, **Email**, **SNMP** or **Syslog** to configure the notifications desired. For the CLI, use the `set snmp-parameters`, `set email-parameters` and `set syslog-parameters` commands.

Frequently asked questions

Does Remote Snap support port failover and load balancing?

Yes, all ports that are connected can be used for replication. For example, an array with a number of primary volumes owned by controller A has ports A1, A2, B1, and B2 connected and part of the peer connection. Both A1 and A2 will be used by all replications. If A1 fails, replications continue on A2. If A2 fails, replications continue on B1 and B2. Note that a replication will use up to two ports for load balancing – when one of the ports fails, it will continue with the second port; if the second port fails, then two ports will be chosen from the remaining available ports, preferably on the controller that owns the replication volume.

Can I expand the primary volume or secondary volume?

You can expand the primary volume in a replication set, but not the secondary volume. The secondary volume will not be expanded immediately but will be expanded on the next replication.

Can I map the secondary volume?

Secondary volumes cannot be presented to any hosts. You can create a snapshot of the secondary volume, then map the snapshot to hosts, or, if replication snapshot history is enabled, use one of the replication history snapshots.

Why were replication history snapshots not taken?

If a volume or a snapshot already exists with the name the replication history snapshot will use, the replication history snapshot will not occur. Verify the names of all existing volumes and snapshots do not conflict with the replication history snapshot naming convention, <basename>+"_nnnnnn", where <basename> is the name you specified when creating or modifying the replication set, and nnnnnn is a number with leading zeroes, starting at 0 and increasing with every replication of the replication set.

Can the host ports run Remote Snap and I/O concurrently?

You can use FC or iSCSI ports for both Remote Snap and I/O on all supported FW versions of HPE MSA Gen6 Storage and HPE MSA Gen7 Storage.

How does CHAP affect replication?

CHAP is just for initial authentication across nodes. Once a login is successful with another system, CHAP will not be involved in data transfer, so replication performance is not affected. Note however that changing the roles in the replication set, even temporarily, will cause an authentication from the current primary array to the current secondary array, and so the CHAP records and enablement should be set so as to succeed with the authentication. It is preferrable to have matching secrets on both the primary and secondary arrays; that way, if CHAP is enabled only on one array, both arrays, or neither array, it will not interfere with the connection.

How do I delete a virtual replication set when the peer connection is down?

Use the `local-only` option of the `delete replication-set` CLI command.



Why can’t I delete a peer connection when no replication sets are present in the system?

You may have deleted the replication set using the `local-only` option of the `delete replication-set` CLI command while the peer connection was done, and then the peer connection came back up. Delete the replication set from the other array to allow deleting the peer connection.



Learn more at

HPE.com

Chat now (sales)

© Copyright 2025 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft is either a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries. All third-party marks are property of their respective owners.

a50013554enw