

Hitachi Ops Center Protector

7.10

User Guide

MK-99PRT002-11
March 2025

© 2016, 2024 Hitachi Vantara LLC. All rights reserved

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., Hitachi Vantara, Ltd., or Hitachi Vantara Corporation (collectively “Hitachi”). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. “Materials” mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara LLC at https://support.hitachivantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara LLC.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, DB2, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, OS/390, PowerHA, PowerPC, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z15, z16, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, Microsoft Edge, the Microsoft corporate logo, the Microsoft Edge logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found in the product documentation, at https://www.hitachivantara.com/en-us/company/legal.html.

Contents

Preface....16
Software version....16
Intended audience....16
Related documents....17
Document conventions....17
Conventions for storage capacity values....19
Accessing product documentation....20
Getting help....20
Comments....20
Chapter 1: Before you begin....21
Prerequisites....21
Hitachi Block prerequisites....21
Hitachi Software Defined Storage prerequisites....24
Generation 1 Hitachi Content Platform prerequisites....24
Generation 2 Hitachi Content Platform prerequisites....24
Amazon S3 prerequisites....26
HCP Cloud Scale prerequisites....26
Upgrading Hitachi Ops Center Protector....26
Upgrading to Protector 7.10....27
Chapter 2: Concepts....28
Ops Center Protector Overview....28
About Ops Center Protector....28
Features and benefits....29
Data protection terminology....29
Architecture....30
Policy orientated data management....32
Real time monitoring....32
Storage hardware orchestration framework....32
Hitachi Remote Ops monitoring....34
Data security....34
Unified Backup Infrastructure concept....35
Security with Ops Center Protector....37
Protection from an external attack....37

Protection from a Compromised OS....38
Protection from a Compromised Account....39
TLS configuration for NGINX web server....39
Universal Tags....39
Node Concepts....47
About nodes....47
About node groups....49
About multi-tenancy for Hitachi block storage....50
Data Flow Concepts....50
About data flows....50
About two-step teardown....51
About data flow implementation....52
About many-to-one data flow topologies....53
About one-to-many data flows....54
About cascading data flows....55
About parallel versus serial data flows....55
About best practices for drawing data flows....57
About best practices for naming objects....59
About Repository and HCP based backups....60
About Repository based batch backup....60
Host based granular file I/O data capture....61
About Repository based source side deduplication....61
About Repository to Repository backup....61
About tiering Gen1 Repositories to HCP....62
About Hitachi Block based backup technologies....63
About mover types used with Hitachi Block operations....64
About Thin Image and Thin Image Advanced differential and refreshed snapshots....66
About ShadowImage replication....70
About TrueCopy replication....74
About Universal Replicator....75
About Global-Active Device replication....76
About three datacentre cascade (3DC)....78
About three datacentre multi-target....79
About three datacentre multi-target with delta....80
About static clones of a clone....81
About clone with replication....81
About snapshot with replication....82
About snapshot of a clone....83
About replication of a clone....83
About remote snapshot....84
About local and remote snapshots....84

About remote clone....85
About local & remote clones....86
About local snapshot and remote clones....86
About Hitachi Block replication adoption....87
Policy Concepts....90
About policies....90
About policy classifications....93
About VMware policy classifications....95
About Hyper-V policy classifications....96
About Microsoft SQL Server Policy Classifications....97
About path macros....98
About application quiescing....98
About synchronization groups....99
About the automated Mount operation....100
About the repurposing mount sequence....100
About the proxy backup mount sequence....101
Schedule Concepts....101
About Schedules....101
Monitor Concepts....106
About monitoring....107
Log Concepts....107
About logs....107
Notification Concepts....108
About notifications....108
Job Concepts....109
About jobs....109
Restore Concepts....110
About restoring data....110
About Hitachi block based replication swapping (takeover/takeback)....111
About restoring legacy application data....113
Restoring legacy Microsoft Exchange Server backups....113
Restoring legacy Microsoft SQL Server backups....113
Restoring legacy SAP HANA backups....113
Restoring legacy CDP and Live backups....114
Storage Concepts....114
About storage....115
About Repositories....115
Hardware and OS recommendations....116
About Cloud Nodes (Amazon S3, HCP and HCP Cloud Scale)....117
Report Concepts....118
About reports....119

About Global Replication reports....120
Access Control Concepts....122
About Role Based Access Control (RBAC)....122
About Access Permissions....122
About Ops Center Protector's implementation of RBAC....123
Selecting an Authentication Service....124
Authorising users and restricting access....126
License Concepts....128
About licenses....128
UI Concepts....129
UI Timeout when Protector is linked to Ops Center....129
Chapter 3: Data Protection Workflows....130
Data Protection Workflows Overview....130
How to protect your data....130
How to restore your data....132
How to protect Ops Center Protector....133
How to recover Ops Center Protector....134
How to revert to the original Master....138
How to protect Ops Center applications....138
How to restore Ops Center application backups....139
Host Based and HCP Workflows....140
How to batch backup a file system path to a repository....141
How to tier a file system path to HCP via a repository....144
How to restore a repository snapshot of a file system path....147
How to backup an onsite repository to an offsite repository....149
How to seed an offsite repository....152
Hitachi Block Workflows....153
How to snapshot an Hitachi Block LDEV with Thin Image....154
How to snapshot a file system with Thin Image....157
How to replicate an Hitachi Block LDEV with ShadowImage....160
How to teardown the S-VOLs of a replication removed from a data flow....164
How to reactivate a replication operation that has been accidently deactivated....164
Case 1: Replication operation has not been removed but the data flow is deactivated....164
Case 2: Replication operation has been removed and the data flow has been reactivated....165
Case 3: The data flow containing the replication has been deleted....165
How to create and use an Hitachi Block Host node....165
How to replicate a file system with ShadowImage....167
How to replicate a file system with Refreshed Thin Image....171

How to replicate a file system with Universal Replicator..... 175
How to replicate a file system with TrueCopy.....179
How to replicate a file system with Global-Active Device.....182
How to implement 3DC multi-target with delta UR replication.....186
How to synchronize snapshots with a replication.....191
How to automatically mount a snapshot or replication 196
How to mount an Hitachi block snapshot or replication.....199
How to revert a file system path from a snapshot or local replication..... 201
How to swap (takeover/takeback) a replication.....202
How to expand a journal.....204
How to delete a journal.....205
How to adopt a replication into Protector..... 206
How to dissociate a replication from Protector.....209
How to snapshot disks of a Windows Server Failover Cluster (WSFC)..... 210
Chapter 4: Tasks.....213
Access Control Tasks.....213
How to configure basic role based access control..... 213
How to configure advanced role based access control.....215
How to create a resource group.....219
How to create a role..... 219
How to create an access control profile..... 220
How to clone an access control profile.....220
How to edit an access control profile.....221
How to create an Authentication Space..... 221
How to configure an LDAP authentication space.....222
How to create an Access Control Profile Association..... 226
How to view the access control settings summary.....227
How to edit object permissions.....227
Data Flow Tasks.....228
How to create a data flow.....228
How to connect nodes on a data flow..... 231
How to apply a policy to nodes on a data flow..... 231
How to activate a data flow..... 232
General Tasks.....232
How to navigate to a page using the breadcrumbs.....232
How to select a single item in an inventory..... 233
How to select all items on an inventory page.....233
How to set a filter on an inventory..... 234
How to change the UI settings..... 234
Installation Tasks.....234
Settings required when using a virus detection program..... 235

How to install/upgrade Protector on Windows and Linux or AIX..... 235
How to upgrade Protector’s PostgreSQL databases from version 12 to version 16..... 240
How to verify the SSL/TLS fingerprint of a client node prior to authorising 242
How to install a Protector master or client on a Windows cluster..... 243
How to install Protector on a node in a cluster..... 244
How to add the hub service to a cluster..... 245
How to add licenses to a clustered master..... 245
How to configure authentication for a cluster..... 246
How to remotely upgrade Protector client nodes..... 247
How to configure a third party firewall for Protector..... 249
How to configure addresses for nodes on multiple networks 249
How to configure a server-side SSL certificate using a UI with Ops Center..... 251
How to configure a server-side SSL certificate using a UI..... 251
Configure server-side SSL certificate using Windows..... 251
Configure server-side SSL certificate using Linux..... 252
How to configure a server-side SSL certificate manually..... 252
How to add our CA certificate to the certificate store..... 254
Using the CLI version of the certificate tool..... 254
How to unregister the current certificate..... 254
How to uninstall Protector..... 254
How to setup an existing Protector master for Ops Center..... 256
Removing Protector from Ops Center..... 257
Job Tasks..... 257
How to view and control running jobs..... 258
License Tasks..... 258
How to Add a License..... 258
How to view a license..... 259
Log Tasks..... 259
How to view logs..... 259
Monitor Tasks..... 260
How to trigger an operation from an active data flow..... 260
How to deactivate an active data flow..... 261
Node Tasks..... 261
How to add a node..... 261
How to authorize a node..... 262
How to stop the services..... 263
How to start the services..... 263
How to enable or disable encryption on a node..... 263
How to change the account credentials for a Block Device node..... 264
Notification Tasks..... 264

How to create a notification....264
How to configure email settings for notifications....265
How to configure SNMP Notifications....266
How to test a notification....270
How to customize alert notifications....270
About the notifications log message file....271
Policy Tasks....272
How to create a policy....273
How to edit a policy....273
How to delete a policy....274
How to add a classification to a policy....274
How to edit a classification in a policy....275
How to delete a classification from a policy....275
How to add a filter to a classification....276
How to edit a filter on a classification....276
How to delete a filter from a classification....276
How to add an operation to a policy....277
How to edit an operation in a policy....277
How to delete an operation from a policy....278
Report Tasks....278
How to view a report....278
Restore Tasks....279
How to view available backups....279
Schedule Tasks....279
How to create a schedule....280
Storage Tasks....280
How to view the status of a repository....281
How to view the contents of a snapshot in a repository store....281
How to view the status of a Hitachi Block storage device....282
Chapter 5: User Interface Reference....283
User Interface Overview....283
User Interface Structure....283
User Interface Page Layout....285
Main Banner....286
Settings Wizard....287
About dialog....289
Online Help dialog....289
Navigation Breadcrumbs....290
Navigation Sidebar....291
Default Dashboard....292
Inventory Page....293

Tile Control..... 295
Button Icons.....296
Universal Tags.....301
UI Timeout when Protector is linked to Ops Center..... 309
Access Control UI Reference.....309
Login Page..... 309
Access Control Dashboard.....310
Access Control Profile Associations Inventory.....312
Access Control Authentication Spaces Inventory..... 319
Access Control Profiles Inventory.....330
Access Control Roles Inventory.....336
Access Control Resource Groups Inventory.....344
Access Control Summary.....348
Access Control Permissions Inventory.....349
Common Controls UI Reference.....350
Path Dialog.....350
Date Time Picker.....351
Date Time Range Picker..... 352
Hitachi Block Host Resize Dialog.....354
Data Flows UI Reference.....355
Data Flows Inventory.....355
Activate Data Flow Dialog.....357
Data Flow Wizard.....361
Data Flow Details.....491
Jobs UI Reference.....493
Jobs Inventory.....493
Purge Jobs Dialog.....496
Edit Job Purge Schedule..... 497
Acknowledge Jobs Dialog.....497
Export Jobs Dialog.....498
Job Details..... 499
Tasks Inventory.....501
Export Tasks Dialog..... 503
Task Details.....504
Licenses UI Reference..... 506
Licenses Inventory.....506
Activate License Wizard.....507
License Details.....509
Logs UI Reference.....509
Logs Inventory.....509
Purge Logs Dialog..... 513

Log Purge Schedule..... 514
Acknowledge Logs Dialog.....515
Export Logs Dialog.....516
Session Log Details..... 517
Log Attachments Dialog.....519
Monitor UI Reference.....520
Monitor Inventory.....520
Monitor Details..... 521
Node Groups UI Reference..... 531
Node Groups Inventory..... 531
Node Group Wizard..... 532
Node Group Details.....535
Nodes UI Reference..... 536
Nodes Inventory..... 536
Node Type Wizard.....539
Node Details.....650
Node Type Icons..... 657
Node Status Indicators.....659
Notifications UI Reference..... 659
Notifications Inventory..... 659
Email Notifications Settings Wizard..... 660
Notification Wizard..... 662
Policies UI Reference..... 670
Policies Inventory..... 670
Policy Wizard..... 671
Policy Details.....737
Reports UI Reference.....739
Reports Dashboard..... 739
RPO Report.....741
Repository Usage Report.....743
Backup Report..... 745
Pool Usage Report.....747
Journal Usage Report..... 749
Network Transfer Report.....750
Command Control Interface (CCI) Count Report.....752
HCP Usage Report..... 753
HCP Cloud Scale Usage Report.....754
S3 Usage Report.....755
Reposotory Usage Report..... 756
Global Replication Reports..... 757
Report Filters.....758

Export Report Dialog....761
Restore UI Reference....763
Restore Inventory....763
Block Restore Options....766
Host Based Backup Restore Options....826
Schedules UI Reference....849
Schedules Inventory....849
Schedule Wizard....850
Schedule Details....857
Storage UI Reference....858
Storage Inventory....858
Hitachi Block Device Details....859
Hitachi Block Host Groups Inventory....861
Hitachi Block Journals Inventory....863
Hitachi Block Pools Inventory....866
Hitachi Block Logical Devices Inventory....868
Hitachi Block Snapshots Inventory....869
Hitachi Block Replications Inventory....878
Hitachi Block Remote Path Groups Inventory....892
Hitachi Block Device Advanced Settings Dialog....893
Hitachi Block Device Scheduled Cache Refresh....894
Hitachi Block Quorums Inventory....895
Hitachi Block Resource Groups Inventory....896
Hitachi Block NVMe Subsystem Details....898
Hitachi Software Block Device Details....901
Hitachi Software Block Snapshots Inventory....902
Hitachi Software Block Replications Inventory....905
Gen1 Repository Stores Inventory....912
Resynchronize Repository Store Dialog....914
Create Repository User Snapshot Dialog....914
Gen1 Repository Store Details....915
Generation 1 Repository Details....921
Generation 2 Repository Details....922
Gen2 Repository Stores Inventory....924
Cloud Storage Details....928
Cloud Stores Inventory....929
Cloud Store Details....930

Chapter 6: Command Line Interface Reference....933
Gathering diagnostic information with **`diagdata`**....933
Gathering diagnostic information with ‘support manager cli’....938
Triggering policies and operations with **`hdidcmd`**....951

Archiving master node setting with **mastersettings** 955
About the mastersettings.cfg file 958
Reducing repository size 960
Generating log messages with **sendlog** 960
Changing a node's profile with **setconfig** 962
Installing and upgrading Protector from the command line 963
Listing Oracle RMAN channel configurations with schedulershow 966
Query Storage Command Line Utility 967
Commands used to generate GRR 968

Chapter 7: Troubleshooting 969

Troubleshooting General 969
Performance issues with the Log Database with very busy Protector instances 969
Error copying file during Protector upgrade 970
How to check if the link between nodes 970
Protector backups failing on some Linux distributions 973
Logging in with multiple tabs open causes circular log in request 973
Moving system time backwards causes problems 974
Session ID and access failure messages using IE 974
Unicode characters not displayed correctly in web UI 974
Troubleshooting Hitachi Block 974
SSB error code information 974
Cannot create any more snapshots 975
Cannot create snapshots in the specified pool 975
Cannot make a mounted snapshot multipathed 975
Cannot revert snapshot if source machine is unavailable 975
Cannot tear down adopted replication having other pairings 976
Clean-up actions not performed if replication setup fails 976
Error 10360 activating UR data flow with grouped source nodes 976
Error when removing replications from Protector 976
ISM crashed, shutdown or rebooted whilst activating a replication or snapshot data flow 977
Policy returns 'Could not start HORCM instance [...]' 977
Proxy node cannot access the LDEV's resource group 977
Renaming an LDEV fails due to unresolved variables 978
Windows 2008 Server doesn't show changes to reverted disk 978
Snapshot fails with error "Cannot create snapshots: specified pool does not have a platform type of "OPEN"" 979
Replication fails with error "Replication policy specifies a pool that is not of type Thin Provisioning or Smart Tiers, cannot create secondary logical devices" 979

Error received stating "Failed to start HORCM instance for command device"....979
Troubleshooting HCP....979
Changing the namespace size on HCP node has no effect....979
Cannot connect to HCP or HTTP 503 Service Unavailable....980
Couldn’t resolve HCP host name....980
HCP connection returns HTTP status 403 Forbidden....980
HCP returns Request Entity Too Large error....980
HCP restore finished with minor issues....981
HCP Node stays offline after node creation....981
Troubleshooting Oracle Database....981
An online backup or mount fails....981
Failed to discover Oracle environment when creating application node....982
Oracle database snapshot fails to mount....982
RAC lock on database failed....982
Cannot find Oracle database metadata files on mount....983
Error when reverting on ASM in normal/high redundancy mode....983
Warning during backup if data/redo files in the same directory....983
Oracle RAC RPO based policy creates more snapshots than expected...983
Listing Oracle RMAN channel configurations with schedulershow....984
Troubleshooting VMware....985
VM MAC Conflict alarm when restoring cloned VM....985
Restoring VMs to original location fails with 'Restore failed to recover all the required VMs'....985
SAN transport message logged for non-SAN datastore....986
Host Based SAN Recovery fails with error writing to Virtual Disk....986
SRM recovery fails with 'Cannot process consistency group [...] expected [...] role target'....988
vRO Ad Hoc Backup fails with '[...] Tag 'HDID/Protector Ad Hoc' already in use [...]'....989
vRO Ad Hoc Backup fails with 'Cause: VMwareException[Tagging cardinality violation]'....990
Troubleshooting Amazon S3....991
Amazon S3 Node stays offline after node creation....991
Troubleshooting Oracle RMAN....991
ORA-27211: Failed to load Media Management Library....992
ORA-19511: ... sbtinit2: initialize communication failed....992
ORA-19511: ... sbtbackup: write to store not allowed....992
ORA-19511: ... Data Folder not found....993
ORA-19511: ... no Root Object....993
ORA-19511: ... database <dbname> not allowed....994
ORA-19511: ... Store exists but cannot be accessed....994

Troubleshooting Hyper-V..... 995
A virtual machine is not included in a backup..... 995
A virtual machine is no longer included in backups after it has been restored..... 996
Restore to original fails stating the node does longer exists..... 996
Restore to Hyper-V cluster fails as virtual machine already exist..... 996
Verification of Hyper-V credential fails.....996
Virtual FC adapters, pass through disks or mapped ISOs are missing after a restore..... 997
VM is skipped during backup because config version is too small.....997

Glossary..... 998
Conventions for storage capacity values..... 1002

Preface

This document describes how to develop scripts and applications for Hitachi Ops Center Protector.

Software version

This document revision applies to Ops Center Protector version 7.10. Please refer to the accompanying Release Notes for information on what's changed in this release.

Intended audience

This document is intended for system administrators, backup administrators and other users who are responsible for installing, configuring, and operating Hitachi Ops Center Protector.

Tip: If you are installing or upgrading Protector, please refer to "Installation Tasks" in User Guide or the accompanying Quick Start Guide.

The User Guide is organized into the following chapters:

- Refer section "Concepts" in User Guide that describes each feature, its purpose, theory of operation and best practices.
- Refer section "Data Protection Workflows" in User Guide that provides high level, step-by-step guides on how to set up and use Protector to protect, manage and restore your data.
- Refer section "Tasks" in User Guide that describes how lower level activities are performed step-by-step via the user interface.
- Refer section "User Interface Reference" in User Guide that details the entire graphical user interface, describing every web page, wizard and dialog. Hyperlinks are provided to enable you to navigate through the UI in a similar way to the actual product.
- Refer section "Command Line Interface Reference" in User Guide that details the command line interface, describing commonly used commands and associated options.
- Refer section "Troubleshooting" in User Guide that lists commonly encountered issues and their solutions.

Related documents

Main product guides:

- *Hitachi Ops Center Protector Software Release Notes*.
- *Hitachi Ops Center Protector Quick Start Guide*.
- *Hitachi Ops Center Protector User's Guide*.
- *Hitachi Ops Center Protector Oracle Application Guide*.
- *Hitachi Ops Center Protector VMware Application Guide*.
- *Hitachi Ops Center Protector Hyper-V Application Guide*.
- *Hitachi Ops Center Protector Microsoft SQL Application Guide*

Programming guides:

- *Hitachi Ops Center Protector REST API User Guide*.
- *Hitachi Ops Center Protector REST API Reference Guide*.
- *Hitachi Ops Center Protector REST API Change Log*.

Document conventions

This document uses the following typographic conventions:

Convention	Description
Bold	▪ Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example: Click **OK**. ▪ Indicates emphasized words in list items.
Italic	▪ Indicates a document title or emphasized words in text. ▪ Indicates a variable, which is a placeholder for actual text provided by the user or for output by the system. Example: `pairdisplay -g group` (For exceptions to this convention for variables, see the entry for angle brackets.)
`Monospace`	Indicates text that is displayed on screen or entered by the user. Example: `pairdisplay -g oradb`

Convention	Description
< > angle brackets	Indicates variables in the following scenarios: ▪ Variables are not clearly separated from the surrounding text or from other variables. Example: `Status-<report-name><file-version>.csv` ▪ Variables in headings.
[ ] square brackets	Indicates optional values. Example: [a | b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a | b } indicates that you must choose either a or b.
| vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a | b] indicates that you can choose a, b, or nothing. { a | b } indicates that you must choose either a or b.

This document uses the following icons to draw attention to information:

Icon	Label	Description
	Note	Calls attention to additional information.
	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
	Important	Highlights information that is essential to the completion of a task.
	Caution	Warns the user of adverse conditions and/or consequences (for example, disruptive operations, data loss, or a system crash).
	CAUTION	Warns the user of a hazardous situation that, if not avoided, could result in major or minor injury.
	WARNING	Warns the user of a hazardous situation which, if not avoided, could result in death or serious injury.

Conventions for storage capacity values

Physical storage capacity values (for example, disk drive capacity) are calculated based on the following values:

Physical capacity unit	Value
1 kilobyte (KB)	1,000 (10^3) bytes
1 megabyte (MB)	1,000 KB or $1{,}000^2$ bytes
1 gigabyte (GB)	1,000 MB or $1{,}000^3$ bytes
1 terabyte (TB)	1,000 GB or $1{,}000^4$ bytes
1 petabyte (PB)	1,000 TB or $1{,}000^5$ bytes
1 exabyte (EB)	1,000 PB or $1{,}000^6$ bytes

Logical capacity values (for example, logical device capacity, cache memory capacity) are calculated based on the following values:

Logical capacity unit	Value
1 block	512 bytes
1 cylinder	Mainframe: 870 KB Open-systems: ▪ OPEN-V: 960 KB ▪ Others: 720 KB
1 KB	1,024 (2^{10}) bytes
1 MB	1,024 KB or $1{,}024^2$ bytes
1 GB	1,024 MB or $1{,}024^3$ bytes
1 TB	1,024 GB or $1{,}024^4$ bytes
1 PB	1,024 TB or $1{,}024^5$ bytes
1 EB	1,024 PB or $1{,}024^6$ bytes

Accessing product documentation

Product user documentation is available on: https://docs.hitachivantara.com. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting help

The Hitachi Vantara Support Website is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to the Hitachi Vantara Support Website for contact information: https://support.hitachivantara.com/en_us/contact-us.html.

Hitachi Vantara Community is a global online community for Hitachi Vantara customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to community.hitachivantara.com, register, and complete your profile.

Comments

Please send comments to doc.feedback@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara LLC.

Thank you!

Chapter 1: Before you begin

Prerequisites

It is important that the following prerequisites are met before you implement any of the policies described in this guide.

To ensure that your hardware and software environment is fully supported, please refer to https://compatibility.hitachivantara.com/assets/ops-protector.

For detailed information on installing the Ops Center Protector Master, and Client components, refer to the *Hitachi Ops Center Protector User's Guide*.

Hitachi Block prerequisites

The following prerequisites are mandatory if you plan to perform snapshotting and/or replication of Hitachi Block volumes. Please also refer to the Protector support matrices at https://compatibility.hitachivantara.com/assets/ops-protector .

- A machine (known as an ISM) must be assigned that controls the Block storage device. This node must be a supported Windows or Linux machine with the Protector Client software installed.

  Caution: ISM nodes and their associated CMDs used to control storage devices must not be shared with applications other than those forming part of the Ops Center suite.

- All primary data (paths or application data) to be snapshotted or replicated within the same data protection policy must exist on the same storage device.
- The block storage hardware must:
  - Support the data protection technologies you intend to use
  - Have the correct firmware version installed
  - Have the correct SVOS version installed
- For all replication types the P-VOLs must be setup in the host group
- In order to resize logical devices represented by the Block Host node that are part of a replication or snapshot pair the array account must have the Support Personal permission.
- For UR, journals must be set up, although for HM800 and later arrays Protector can create journals
- For GAD, the quorum disks or quorum-less disks must be provided. In most cases, one quorum device between participating storage systems is satisfactory. Please refer to the *Global-Active Device User Guide* for best practices

- For GAD, the array should be configured to allow virtualized LDEVs if they are required (where supported by the array)
- Port security must be enabled.
- Primary volumes must be set up using other Hitachi tools prior to selection in Protector
- For application consistent snapshots, the application must be installed and configured to use P-VOLs on a storage array, see the relevant application guide for details on the application configuration
- The password for authorizing a Block Device node must contain only useable CCI command characters: `A-Za-z0-9'-./:@\_`
- The device must have adequate shared memory (see Provisioning and Technical Guides)
- Pools must be created using Storage Navigator prior to selecting the Target Storage in Ops Center Protector:
  - For standard mode (non-cascading) TI the TI Pools must be set up
  - For cascade mode TI the Dynamic Provisioning Pools must be set up to also be a hybrid pool, otherwise a TI pool will also be required
  - For SI, TC, UR and GAD the Dynamic Provisioning Pools must be set up
- The following licensed features may be required depending on the features being used:
  - Dynamic Provisioning
  - Storage Navigator
  - Thin Image (for TI snapshot and RTI replication scenarios)
  - ShadowImage (for SI replication scenarios)
  - TrueCopy (for TC replication scenarios)
  - Universal Replicator (for UR replication scenarios)
  - Global-Active Device (for GAD replication scenarios)
  - Remote Replication Extended (for 3DC scenarios)
- The Protector ISM node controlling the block storage device must have:
  - The correct version of Hitachi CCI installed.
  - If CCI is not installed in the default location there are two options:
    1. Add a symbolic link from the default location to the install directory
    2. Configure Protector to use CCI in the custom location using the following instructions:
       a. Stop the Protector services on the ISM node
       b. Go to the directory <Protector home>\db\config
       c. Make the change to all files matching hitachivirtualstorageplatform*.cfg

d. Change the <BinDirectory> value from C:/HORCM/etc to the correct installation path

```
<!-- Install directory of CCI, override to change
installation directory. -->

<BinDirectory>C:/HORCM/etc</BinDirectory>
```

e. Ensure the change has been made to all files at per 3 including the default one.

f. Start the Protector services on the ISM node

- Access to a dedicated Command Device (CMD) on the storage device, set up as follows:

  WARNING: When running the Analyzer probe server, API Configuration Manager, and Protector ISM Client on the same VM, all components share the same command device, but API Configuration Manager and Protector ISM Client must access the storage systems using different credentials. This means that API Configuration Manager and Protector ISM client must use different login accounts when accessing the storage system.

  - Security disabled
  - User authentication enabled
  - Device group definition disabled
  - The CMD must be visible to the host OS where the Protector proxy resides
  - The CMD must be offline
  - The CMD must be added to the meta_resource only.
  - Multiple active command devices may be visible to a Protector proxy as long as each one represents a different block storage device. Behaviour is undefined if multiple active command devices represent the same block storage device, unless these are configured in the Protector proxy node fail-over priority list.
  - Fibre channel and IP command devices are supported.
  - Multipath for Command Devices is supported

- A dedicated user (specified when creating the Hitachi Block Device node) for Protector must be created on the storage device with at least the following roles:
  - Storage Administrator (Provisioning)
  - Storage Administrator (Local Copy)
  - Storage Administrator (Remote Copy)
  - Security Administrator (View & Modify).

  The user must also have access to Resource Group 0 on the storage device.
- Fibre connectivity (including zoning) or iSCSI connectivity and pre-configured RCU paths between arrays for remote replication technologies
- If physical and software block devices are being configured in a single Protector environment it is essential that they do not share an ISM node.

Hitachi Software Defined Storage prerequisites

- If physical and software block devices are being configured in a single Protector environment it is essential that they do not share an ISM node.
- Proxy nodes for software block devices are only supported on Linux operating systems. Windows proxies are not supported at this time.

Generation 1 Hitachi Content Platform prerequisites

If you plan to use the Generation 1 Hitachi Content Platform (HCP) with Ops Center Protector, then the following Search Facility Settings must be selected within the Settings tab of the HCP Management Console, for the Metadata Query Engine (MQE):

- Enable indexing
- Enable indexing of custom metadata

Click Update MQE Settings to reflect the changes.

- A Protector HCP node can only be created if the Protector Master node can directly connect to the HCP web interface.
- The tenant must have Enable management through APIs turned on in the **HCP Tenant Management Console**.
- The user should have at least the following enabled in the **HCP Tenant Management Console**:
  - Roles:
    - Administrator
    - Compliance
  - Permissions:
    - Read
    - Write
    - Delete
    - Privileged
    - Search

Generation 2 Hitachi Content Platform prerequisites

If you plan to use the Generation 2 Hitachi Content Platform (HCP) with Ops Center Protector, then the following prerequisites must be met:

Configure the following system global settings:

HCP UI Location	Setting Name	Enabled
Security	Enable the management API	✔

HCP UI Location	Setting Name	Enabled
Services	<MQE service is running>	✓
Services	Enable metadata query API	✓
Services	Enable indexing	✓
Services	Enable indexing of custom metadata	✓

- Create and configure the tenant(s)
  - The Tenants *must be manually created* by the administrator account
  - After a tenant is created the following tenant settings *must be manually configured*by the administrator account

HCP UI Location	Setting Name	Enabled
Tenants → <tenant name> → Settings	Search	✓
Tenants → <tenant name> → Settings	Versioning	✓

  - These settings are verified prior to creating a new namespace. If these settings are not enabled the namespace will not be created
- Configure each tenants' global settings
  - After a tenant is created the following global tenant settings *must be manually configured* by an administrator account
-

HCP UI Location	Setting Name	Enabled
Security	Enable the management API	✓

- Create the tenant accounts to be used by the HCP Cloud Connector
  - These accounts must be created from the HCP Tenant Management Console
  - Each HCP account must have, as a minimum, the following roles enabled:

HCP Role Name	Enabled
Administrator	✓
Compliance	✓

Amazon S3 prerequisites

If you plan to use the Amazon S3 with Ops Center Protector, then the following prerequisites must be met:

- The Amazon S3 proxy node must have an internet connection and be able to communicate with Amazon S3
- The Protector Master must have an internet connection and be able to communicate with Amazon S3
- An AWS account with an Access Key ID and a Secret Access Key.
- The AWS account must have *AmazonS3FullAccess* and *CloudWatchReadOnlyAccess* permissions.
- Protector can only use the bucket it creates, it can not be configured to use a different bicket.

HCP Cloud Scale prerequisites

If you plan to use HCP for cloud scale with Ops Center Protector, then the following prerequisites must be met:

- The HCP for cloud scale proxy node must be able to communicate with the HCP for cloud scale device.
- The Protector Master must be able to communicate with the HCP for cloud scale device.

Upgrading Hitachi Ops Center Protector

This section addresses the issues when upgrading to this release of Protector

Note: Before upgrading, unmount any mounted snapshots.

Do not upgrade while replications are pairing. Any active replications must be in the paired state before you start the upgrade.

For full details on upgrade paths, see the upgrade path matrix by clicking on "View Full Specifications" here https://compatibility.hitachivantara.com/assets/ops-protector

Upgrading to Protector 7.10

Only v7.1 and later can be upgraded to v7.10. If running an earlier version, you must upgrade to version 7.1 or later before upgrading

If upgrading from a version prior to Version 7.5.0. due to changes to the Jobs system introduced in 7.5 it is necessary to activate the data flows rules so that the jobs have their correct titles. Additionally, the Jobs history will be reset on upgrade.

If the current version of Protector is lower than Version 7.4.0 and push upgrades are required it is necessary to push upgrade the clients first before the master. Protector 7.4.0 introduces increased security of its inter-node communications. Due to this it's not possible for Version 7.4.0 to communicate to older versions of the software. So when crossing the 7.4.0 boundary its important to first push upgrade any clients to 7.4.0 before upgrading the master node. Any clients that are push upgraded to 7.4.0 or later will show as offline until the master is upgraded. Once the master node is upgraded to any older clients will show offline and will need to be installed by running the installer on the client directly.

Version 7.1 uses a new backend database to store logs and recovery points. If upgrading from a pre v7.0 version, the recovery point database is automatically built. However, the logs database is empty after upgrade, which affects both the logs and some historic reports. If required, you can import the old logs into the new system using a CLI tool. Importing logs may take some time, however the import runs at the same time as the services so that daily operations can continue. For details, see the user guide.



Note: If importing logs older than 13 weks into the system, it is important to change the log purge settings before importing the logs so that the older logs are not automatically purged and are kept for the required time period.

Chapter 2: Concepts

This chapter describes Ops Center Protector's functionality and features.

Ops Center Protector Overview

This section provides an overview of Ops Center Protector.

For further information, refer to:

- Tasks (on page 213)
- User Interface Reference (on page 283)
- Command Line Interface Reference (on page 933)
- Troubleshooting (on page 969)

About Ops Center Protector

Ops Center Protector provides a modern, holistic approach to data protection, recovery and retention. It has a unique workflow based policy engine, presented in an easy-to-use whiteboard-style user interface that helps map copy-data management processes to business priorities. A wide range of fully integrated hardware storage-based and host-based incremental-forever data capture capabilities are included that can be combined into complex work flows to automate and simplify copy-data management. With these you can:

- Choose the right technology for each workload, based on service level requirements, but manage them from one place.
- Drag-and-drop a range of protection, retention and repurposing capabilities to easily create and manage complex work flows.
- Automate and orchestrate Hitachi storage-assisted snapshots, clones and replications to eliminate backup windows.
- Automate the mounting of Hitachi storage based snapshots, clones and replications for proxy backup and repurposing.

Protector supports a wide range of data storage targets, including repository, block, and file based storage.

Features and benefits

Ops Center Protector enables you to:

- Meet operational recovery, disaster recovery and long-term recovery challenges in a single, unified platform.
- Drive the backup window, recovery point objective (RPO) and recovery time objective (RTO) to near zero.
- Employ incremental-forever data capture at the block level helping to reduce secondary storage requirements by 90% or more.
- Use intuitive, drag-and-drop workflow creation, node and policy management wizards for administrative agility.
- Use Hitachi Thin Image snapshot orchestration for your critical data and applications.
- Easily create or adopt ShadowImage, TrueCopy®, Universal Replicator, Global-Active Device and File Replication without scripting, on a scheduled or ad-hoc basis.
- Orchestrate application-consistent snapshot and clone management for Oracle.
- Orchestrate application-consistent snapshots and clones across other applications and file systems.
- Combine automated local snapshot and off-site replication for an end-to-end modern data protection and recovery solution.
- Automatically trigger snapshots and clones of remote replica data for secondary operations, such as test and development.
- Mount and unmount snapshots and clones automatically as part of an orchestrated policy workflow.
- Further reduce network and storage costs through source and target data deduplication.
- Protect VMware hypervisors and virtual machines.
- Support a broad range of storage types, including repositories and block storage.

Data protection terminology

The following terminology is frequently used when describing data protection in this guide:

Host based backup involves backing up data over the network to a storage device such as a repository. The backup and restore operations are implemented by Protector software processes that utilise the smallest possible amount of the computing, memory and disk resources on the node being backed up or restored to.

Repositories are data stores implemented by Protector software processes that hold backup data. Repositories utilise disk volumes mounted on the host where the repository is located. They have built-in capabilities for data reduction, such as block differentials and de-duplication, as well as the capability to record point-in-time snapshots of captured data. Backup and restore operations may utilise significant amounts of the computing, memory and disk resources of the node where the repository is located.

Proxies interface with physical storage devices and hypervisors, and assist in coordinating the backup and restore processes. Proxy nodes have Protector software installed on them and in some cases require additional third party software to interface with the devices they control. Examples of devices requiring proxy nodes include:

- Block storage arrays
- VMware vCenter servers

Block hardware storage devices use snapshotting and replication technologies implemented by the storage device itself. Proxies running Protector software orchestrate backup and restore functions on these devices but are not directly involved in capturing or transferring backup data. Examples of these devices include Hitachi Block storage arrays.

File hardware storage devices use snapshotting and replication technologies implemented by the storage device itself. Proxies running Protector software orchestrate backup and restore functions on these devices but are not directly involved in capturing or transferring backup data.

Application nodes are existing application servers or clusters that require Protector and additional third party software to be installed. Protector software discovers application data paths and coordinates backup and restore operations. Applications can be placed in the correct state to ensure consistent backup and restore processing. Example application nodes include:

- Oracle Database

Architecture

At the heart of Protector is the *Master*; a software component installed on a dedicated server that acts as a central hub, communicating with users via a REST interface and web based UI. The *Master* monitors and controls the activity of a set of *Clients*; software components installed on servers in the protected environment. Normally, the *Master* does not actively take part in data protection data flows; *Clients* do this themselves according to *rules* distributed from the *Master*.

A Protector installation thus consists of:

- A *Master* that controls the UI and the actions of all other nodes on the environment.

  Master nodes can coexist on the same network, allowing multiple Protector environments to coexist.

  Passive standby master nodes can be installed on a Windows failover cluster so that the standby can be brought online in the event that the active master fails.

  Note: It is recommended that the *Master* server not be assigned any other roles.

- Multiple *Clients* that participate directly in data protection data flows, acting as sources of data to be protected (e.g. database servers) and destinations capable of receiving backup data (e.g. repositories).

  Clients implement various different roles within a data flow and must be configured to perform these roles. This consists of ensuring the appropriate prerequisites are installed, and configuration via the UI.

  Client nodes can be installed on application clusters such as Oracle RAC.

  Each *Client* can only be authorized and controlled by one master.

The figure below shows a possible installation scenario where two completely independent Protector environments coexist on one network.

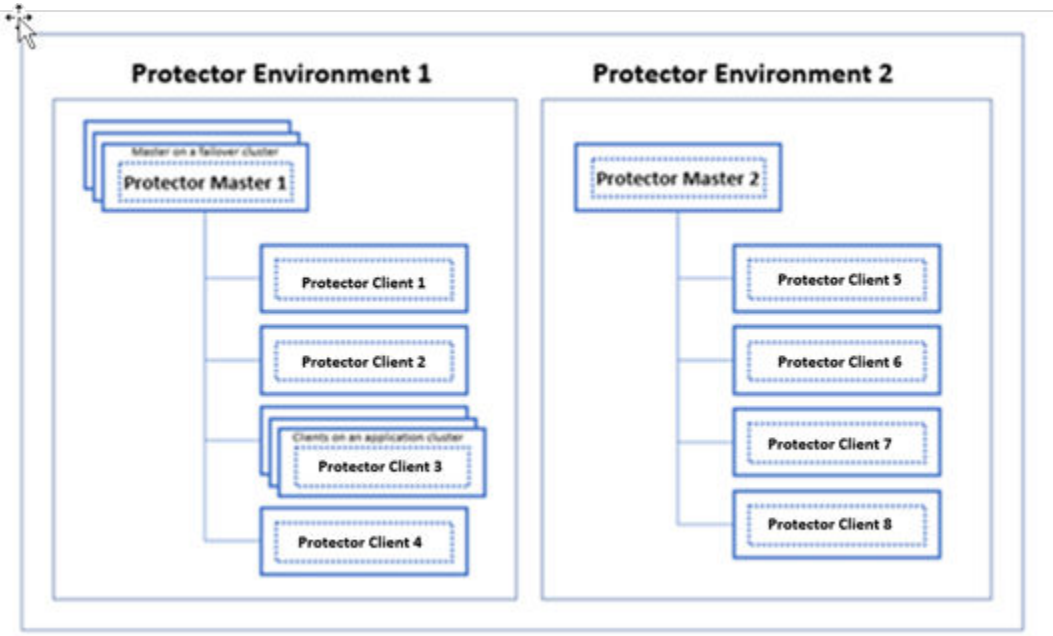


Figure 1 Possible installation scenario

Policy orientated data management

Ops Center Protector is designed to allow your site to implement a business data protection process by using policies. Policies are mapped to machines or groups, and to the data management agents that implement those policies. A large, complex environment can be configured in a matter of minutes rather than days through the application of two workflows:

- *Policy Definition* defines the business objective in terms of data classification and the associated backup operations.
- *Data Flows* define where those policies are applied and the flow of protected data, reflecting the information life cycle.

Real time monitoring

Data protection activity can be monitored via a central dashboard and mimic displays of active data flows. Data flow diagrams containing job queues, log messages, meters and status icons provide real-time visual feedback on:

- *Outgoing/Incoming Network Traffic*: The amount of data being sent/received over the network.
- *Node Status*: An alert badge indicating a node needs attention.
- *Jobs*: Currently active and recently completed data protection tasks along with their details.
- *Logs*: Information, warning and error messages logged by the various participating nodes in the data flow.

Storage hardware orchestration framework

Ops Center Protector is able to control the snapshotting and replication of data stored on Hitachi block and file storage hardware.

Ops Center Protector's hardware orchestration framework is able to create snapshots and replications in an application consistent state by interacting with production applications (including Oracle) and the storage array hardware to ensure that data is in a consistent state during snapshotting or replication; it does this with minimal disruption to the production application. The orchestration framework discovers the underlying hardware paths for application data and creates the required snapshots and replications when the policy is triggered.

The orchestration framework is extensible, allowing other managed storage technologies and applications to be integrated in future. These capabilities are all easily accessible using Ops Center Protector's intuitive, policy and data flow definition UI.

The following figure shows the Protector nodes and processes involved in replicating application data on a block storage device at the primary site to a block storage device at the secondary site, and potentially creating local and remote snapshots of that data.

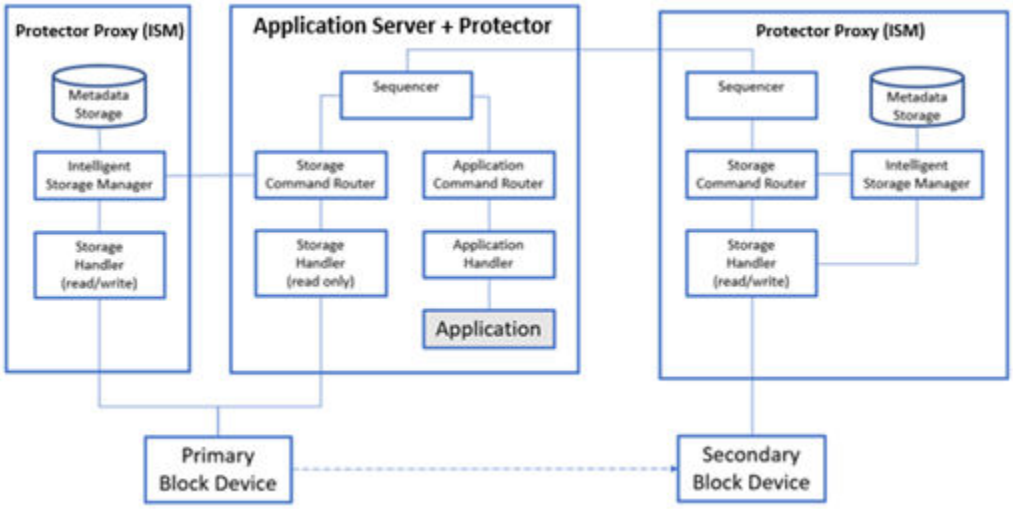


Figure 2 Hardware orchestration process model (remotely replicated block based application)

ISMs (Protector Clients with CCI software installed) are assigned to control the primary and secondary Block Storage Devices. The Application Server that is to be protected has the Protector Client software installed on it, along with any prerequisites required to enable Protector to interact with the application. Rules that define when and what application data is to be replicated/snapshotted are distributed from the Protector Master to the Application Server.

The Sequencer process on the Application Server coordinates a series of actions involving the primary and secondary ISMs when these rules are triggered. This includes:

- Querying the Application for the file system paths that comprise the application data to be replicated. A list of LUNs is then compiled based on the application data paths obtained.
- Querying the Primary Block Device to find the underlying LDEVs that represent the application's LUNs.
- Configuring the Primary and Secondary Block Devices to replicate/snapshot the application's LDEVs.
- Quiescing the Application to ensure application consistency (assuming the policy has requested it).
- Triggering and synchronizing the replication/snapshot operations between the Primary and Secondary Block Devices. For applications such as Oracle this involves multi-stage backup operations to capture databases, logs and redo-logs.

A Sequencer process on the destination ISM acts as a counterpart to the application Sequencer. It manages the destination storage sequences.

The Application Command Router process routes commands from the Sequencer to the appropriate Application Handler, enabling multiple applications to be controlled on the same node.

The Application Handler process(es) present a generic interface for each application being controlled by Protector.

The Storage Command Router process routes commands from the Sequencer to the appropriate Storage Handler, enabling multiple storage devices to be controlled from the same node.

The Storage Handler Process(es) present a generic interface for each storage device being controlled by Protector. Note that the Storage Handler on an Application node is read-only; this is a security feature that restricts modification of the storage device to authorised ISM nodes only.

The ISM process maintains a list of snapshot and replication records along with their status and manages the retention of those records.

Hitachi Remote Ops monitoring

Hitachi Remote Ops is a centralised monitoring service which tracks the operational status of Hitachi devices within a customer's site. Hitachi uses information collected by Hitachi Remote Ops for statistical analysis, diagnostics, and pre-emptive error detection. Protector collects data from each of its nodes and uploads this information, via a locally installed monitoring agent, to Hitachi Remote Ops which then interprets it. Protector includes Hitachi Remote Ops integration capabilities as standard, but the monitoring agent must be obtained and installed separately.

The Hitachi Remote Ops data centre is managed by Hitachi and cannot be accessed by customers. A single monitoring agent, connected to the Hitachi Remote Ops data centre via secure SSL/TLS link, is installed at the customer's site and the customer's Hitachi devices are manually registered with it. A bespoke data schema is collected from each device containing a *Site ID* unique to the customer site. For Protector, the *Redact* option is permanently enabled to instruct the monitoring agent to redact all sensitive data in the returned payload. Protector uses Role Based Access Control (RBAC) whereby clients must authenticate themselves before they are allowed to interact with it. The Protector *Telemetry Admin* role and *Default Telemetry* ACP exist within Protector for use with Hitachi Remote Ops. This role and associated ACP restricts what actions the monitoring agent is able to perform on Protector. The monitoring agent must login to Protector with credentials that enable the *Default Telemetry* ACP prior to requesting data.

Protector is queried daily for its current operational status by the monitoring agent and that status is uploaded to the Hitachi Remote Ops data centre. If any errors are received within Protector's status data, then the monitoring agent can increase its polling frequency to every 2.5 minutes until the errors are cleared. This behaviour is optional and a device does not have to support this higher rate reporting mechanism. The monitoring agent may also be manually triggered to request data at any time; such a request may contain additional diagnostic data which may not be sent as part of the standard request.

Data security

Ops Center Protector allows the administrator to secure the data that the system protects with encryption in a number of ways:

1. If it is known that the node will be operating over a nonsecure network, such as the Internet, then enabling the Internet connected node option is recommended. This will require the Protector ‘over-the-wire’ encryption licence: This will encrypt transmitted data as an extra security precaution. With the appropriate licence this uses the SHA-256 cryptographic hash function to encode data and allow the data integrity to be verified. In the non-licenced case it is used to verify the integrity of the certificate for the secure channel. Protector uses a NIST validated implementation of the SHA-256 algorithm.
2. With the appropriate ‘at-rest’ data encryption licence (where available) Protector uses the SHA-256 hash algorithm to encode the data that is stored on an encrypted repository. The data remains encrypted and secure even when the repository is not being used.

Unified Backup Infrastructure concept

UBI is a new infrastructure that allows applications and storage to interface to each other in a standard and uniform manner. Previously, the storage handlers needed some understanding of the data they were storing requiring dedicated logic for each supported application. This complicated and lengthened the development of both applications and storage handlers.

UBI however provides a standard way for applications and storage work together. This required the development of new UBI compliant application and storage handlers.

UBI Applications includes:

- Filesystem
- VMware
- Oracle RMAN

UBI Storage includes:

- Gen2 Repository
- Gen2 HCP
- Amazon S3

Any UBI compliant application can store data in any UBI compliant storage. UBI compliant storage can forward an instance of data to any other UBI compliant storage.

Data within UBI Storage is structured as follows:

Datastore

A Datastore is the UBI abstraction for an entire storage location (e.g. an AWS Bucket or an HCP namespace)

Store

This is a logical construct. Each instance of an application has its own store. If the application is distributed across many nodes, it would still share the same store.

Instance

This is the instance of data within the store. Each store will have a live instance and historical instances. Consider this analogous to a filesystem that you can read and write but also has snapshots. The historical instances facility is available on all UBI storage, even cloud storage such as S3 which doesn’t have native snapshots. If an

instance of data is removed, data unique to that instance is also removed, thereby providing efficient automated space reclamation.

Objects

Objects are place holders for data. They can be hierarchical, and you can attach metadata to them. They do not however store data directly, data is stored in streams within the object.

Streams

Streams are what objects use to store their data. Typically, there will be a “data” stream. Applications may also use other streams to store data relevant to that object, e.g. security stream, thumbnail, etc.

The UBI standard key features:

Parallelism

To facilitate performance, the UBI standard allows a datastore to be accessed by multiple processes on the same machine and/or multiple processes across multiple machines. This allows for a clustered application to backup from multiple nodes to one store.

Recovery Point

This allows the user to see that there is a point in time recoverable data. Typically a Recovery Point maps to an *instance* within a *datastore*, but more complicated applications could in theory have multiple instances to make recoverable data. It is the application that determines whether a Recovery Point is created. Some applications, like Oracle RMAN, don’t have Recovery Points as the RMAN application manages the life cycle of individual portions of a backup without using *instances*. Recovery Points are stored in a database on the storage but can be searched from the Master.

Item Index

An item index is a method to add a layer of indexing more granular than the recovery point index. It is primarily used to index large objects. A good example is a VMware backup creates an Item Index record for each VM. This allows the user to look for instances of a VM by doing a search for a VM rather than browsing Recovery Point Index. Like Recovery Point Index, these are stored on the storage.

Source Side Block Deduplication

UBI has built-in source side block deduplication as part of its standard. The infrastructure calculates the checksum of a block for a given stream of data prior to transmitting it to the storage and will not send that data if it is already on the storage. This reduces both storage and network transfer overhead.

Access Control

Many applications across many OS nodes can access the storage. Each application however can access the same storage but only their data within that storage. This includes applications that are using cloud storage. With Amazon S3, for instance, an application can only access its own data even if it’s in the same Amazon Bucket as other applications.

Cloud Credential Security

The UBI infrastructure assumes that because applications servers are not under control of the backup / storage administrator, they can be compromised. So, although multiple application may access the same cloud storage node, to ensure security, an application never receives the cloud credentials.

Application aware data browsing

When performing a partial recovery of a dataset, it is required to be able to browse the contents of a Recovery Point. To do this an application specific agent on the master is used to interpret the data retained on the storage device. This ensures the recovery interface can display recovery information that is optimized for that application.

Security with Ops Center Protector

This document the built-in measures Protector employs to protect the system from an external attack and if a system is compromised either through a bad actor or hack, the way the way systems are further protected.

Protection from an external attack

An external attack is when an attempt is made to gain access or compromise the operation of the system without having OS access to the system. To protect against such eventuality Protector employs the following:

- Secure node to node communication - Communication between nodes is performed through a private protocol on a fix network port. Protector uses TLS connection between nodes on that port and utilises digital signatures that were generated during installation time for node identity verification.
- Secure REST API - All user requests are performed through the Role Based Access Control (RBAC) secured REST API to the Master Node. Making any changes to the configuration requires credentials and permissions at the right levels, no mater whether accessed via GUI or API.
- There is no concept of a default user or password. The initial user and all subsequent users must be specifically defined.
- Regular update of internal libraries to be update to date with the latestinformation from National Vulnerability Database (NVD).

Protection from a Compromised OS

There may be many systems that make up part of a backup ecosystem. Some of these systems include the source machines, applications and file servers that are not under the backup admin's tight control and are more prone to security breaches, whether by malware or hacks. To protect against the compromised systems:

- Regardless of the copy technology used (storage array replication or backup):
  - Nodes can only communicate with each other where a predetermined communication path has been specified by a dataflow.
  - Nodes within Protector that can communicate to each other, can only issue requests for predetermined tasks. This is done via the policy / dataflow that are compiled into rules that are propagated to both the source and destination and enforced by both nodes so that if either attempts operations outside of the rules then the operations are not permitted.
  - Source Nodes within Protector cannot delete or alter their recovery points using the internal node to node protocol.
  - Protector' Windows clients use Protector's own VSS provider. This provider has no capability to delete its own snapshots and the client does not have visibility of the backup catalog.
- For storage array replication technologies only:
  - Source Nodes within Protector do not have direct access to control a hardware array. Only Protector's storage proxy nodes require array Command Devices. Array Command Devices do not need to be allocated to source nodes. This means that a source node cannot manipulate its own or other system's storage, replications and snapshots.
  - Use of hardware Data Retention Utility (DRU) to lock down backups for a given period of time. DRU protected storage cannot be modified and is akin to be putting data in a time locked safe. This allows a company's assets to be protected for a timeframe suitable for the detection of an intrusion.
- For backups only:
  - Nodes within Protector that perform backup to cloud (Amazon S3 / HCP) do not have access to the credentials of that cloud service.
  - Backups to HCP, a WORM object store, where you can archive lock data is another suitable solution for host based backups.

Protection from a Compromised Account

It is possible that an account within the backup system could be compromised. E.g.through social engineering , a bad actor, etc. To mitigate this the following best practices can be employed:

- Protector has a very granular RBAC system. It not only defines the what roles as user can have but also roles against resources. So a user that may have complete control of certain assets within a system and may have no control or even visibility of other assets within Protector's inventory. By limiting access to the minimum level that is required for someone who has an account within Protector, it is possible to limit the scope of what that account can do if it is compromised. You can for instance give a user access to use an array within Protector but not give them administrative rights.
- Use of hardware Data Retention Utility (DRU) to lock down backups for a given period of time. DRU protected storage cannot be modified and is akin to be putting data in a time locked safe. This allows a company's assets to be protected for a timeframe suitable for the detection of an intrusion.
- Backups to HCP, a WORM object store, where you can archive lock data is another suitable solution for host based backups.

TLS configuration for NGINX web server

The Protector GUI and REST API is served through an NGINX web server on the master. The cryptographic cyphers and TLS version can be configured by editing `<installdir>/runtime/nginx/conf/nginx.conf`. Instructions on how to do so can be found at http://nginx.org/en/docs/http/ngx_http_ssl_module.html. This will not affect communications between nodes, that is configured by `<intalldir>/db/config/hub.cfg` however since it's only used to interface with Protector and does not need to be accessed by third party software such as a web browser Protector uses strictest settings it can thus there's typically no need to change it.

Note: Any configuration changes made to nginx.conf will be overwritten on upgrade.

Note: Incorrect changes to these configuration files can prevent Protector from functioning.

Universal Tags

As the number of managed objects increases within a system, it becomes progressively more difficult to manage large lists using only the object’s native metadata. Universal tags provide a method for users to add extra metadata in the form of Tags to enhance the search and filter abilities of lists.

Tag-able objects

Tags can be applied to the following objects:

- Nodes
- Node groups
- Policies
- Policy Operations
- Data Flows
- Recovery Points

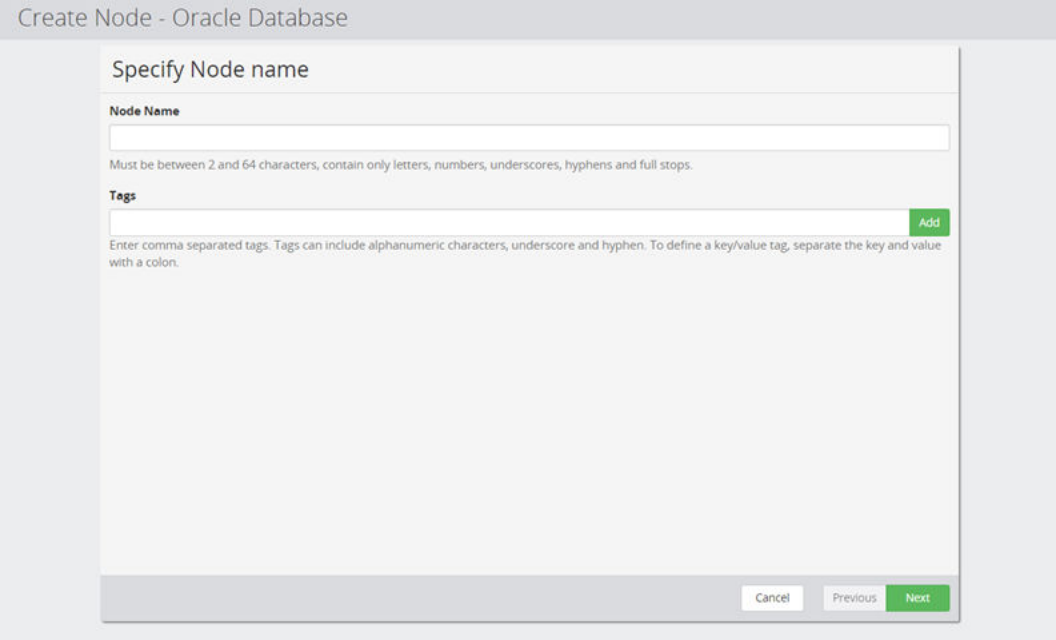


Figure 3 Create Node - Oracle Database

Control	Description
Node Name	Enter a name for the Node.
Tags	Add the tags to be associated with the object being created.
Cancel	Discards all changes made to the settings and reverts to those used prior to opening the current page.
Previous	Takes the user to the previous screen in the wizard.
Save	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

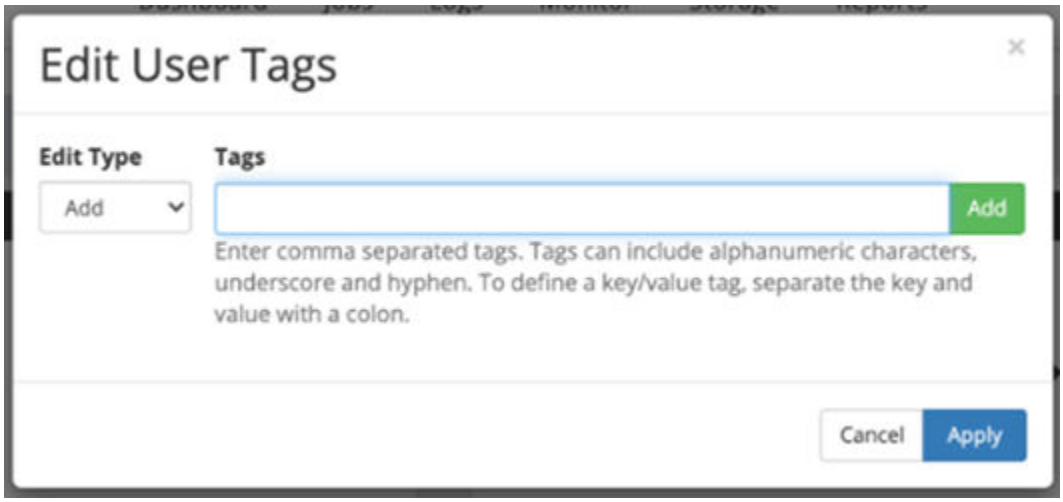


Figure 4 Add Tags - Edit User Tags

Control	Description
Edit Type	Enter the Edit Type.
Tags	Add the tags to be associated with the object being created.
Cancel	Discards all changes made to the settings and reverts to those used prior to opening the current page.
Apply	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

Tag inheritance and Propagation

Jobs and recovery points are objects that are generated by the system after a backup / snapshot operation.

To assist the manageability of these, they inherit the tags from the objects that were involved in creating them. Jobs and recovery points will inherit the tags from:

- The source node they represent in the data flow
- The node group the source node is in
- The storage node that stores the recovery point
- The data flow that created the recovery point
- The policy that created the recovery point
- Custom tags that were included if manually triggered

Inherited tags are assigned at the point that the jobs or recovery points are created and stored as independent tags against those objects. The tags assigned due to inheritance persist unchanged even if the originated object's tags later change, e.g. if a recovery point has inherited the tag "Production" from the Application Node, later removing that tag from that Application Node will not affect existing recovery point.

Note also that new recovery points and jobs will also only inherit any tags changes once the affected data flow have been re-activated. So if you were to alter the tags on a node, new recovery points will not reflect those changes until you re-activate the data flow that have those nodes.

You can add and remove tags on recovery points. You can modify tags on Recovery Points by either selecting the Recovery Point on the Restore Screen and selecting the "Edit Tag" icon, or viewing the Recovery point record and selecting the "Edit Tag" icon.

Types of Tags and Format

There are two types of tags; simple tags and key value pair tags.

Simple tags are straight forward labels. You may for instance tag all your production source machines with the tag “production”. You could then go to the RPO report and check the RPO only for machines that are marked “production”.

There are also key value pair tags. They are in the form of *key:value* . The concept is the ability to categorise a tag. For instance you can assign nodes a location tag in the form of *location:paris* or *location:London*. If you lost connection to a data center, you would be easily see that all the nodes at a particular location are down.

Tags can only contain alphanumeric characters, underscores and spaces. The “:” in a tag will designate it to be a key value pair. Capitalisation of tags is preserved but from a search and filter point of view ignored.

Search Filter using User Tags Query box

Many inventory screens have the ability to perform filter queries using tags. On those screens you will be provided with an entry box for adding a tag to search on. The search will return results that contains any tags that are in the search query, e.g. this search is an "or" operation on the selected tags. The tag needs to be an exact match the tag or the key value of the tag. The search however is not case sensitive.

So if you had an entry tagged: "Location:London" and "production" the following is the behaviour of the searches:

- Searching for tag: "Production" will return the entry because it is a case insensitive match.
- Searching for tag: "Prod" will not return the entry as this is a partial tag
- Searching for tag: "location:" will return the entry because the key value matches. "London" is not required
- Searching for tag: "location" will return the entry because the key is treated as a simple tag as well as key value pair tag
- Searching for tag: "Location:Paris" will not return the entry as the value doesn't match

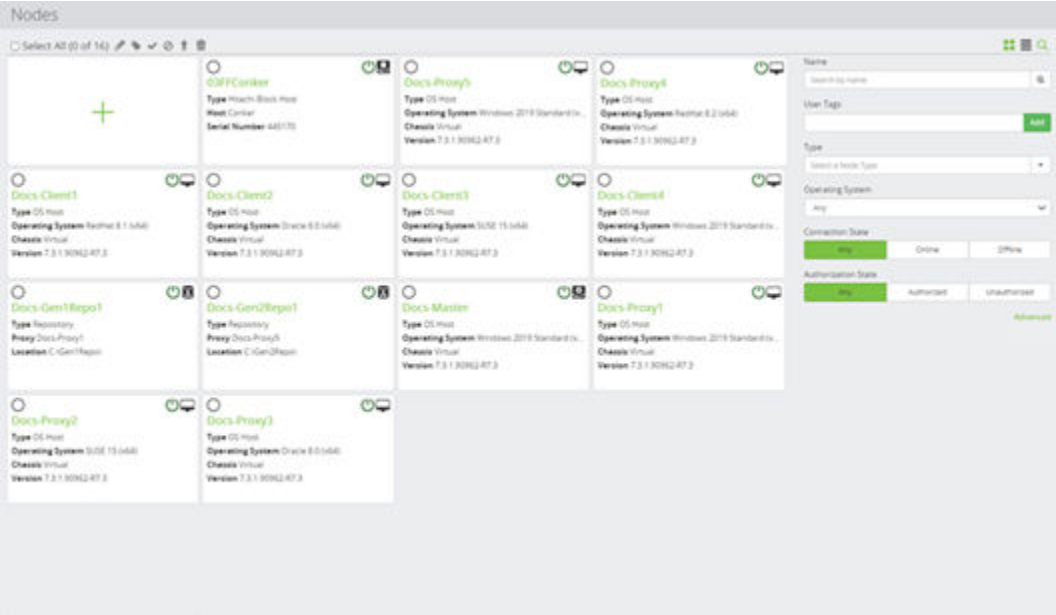


Figure 5 Nodes Inventory

Control	Description
Edit	Edits an existing node in the inventory. The Node Type Wizard (on page 539) is launched to enable the node's attributes to be changed.
Tags	Modifies the tags of an existing object from the either the inventory screen or the details screen of the object.
Authorize	Enabled only if one or more unauthorized nodes is selected in the inventory. Attempts to authorize the selected nodes with the Master node. Only nodes that have been authorized by the Master node may perform Protector functions. **Note:** If an attempt is made to authorize an inactive or unknown node, or if the master node fails to communicate with the node, an error log is generated and the node remains unauthorized.
Deauthorize	Enabled only if one or more authorized nodes is selected in the inventory. Attempts to deauthorize the selected nodes. Nodes that have been deauthorized cannot perform Protector functions.

Control	Description
	Note: ▪ Before deauthorizing nodes that are no longer required, they must first be deleted from the current data flow definitions, then any modified active data flows must be reactivated. If any attempt is made to deauthorize a node that is implementing rules in the currently active data flows, an error message is generated for each such node and the node remains authorized. ▪ It is possible to de-authorize a node with active mounts. This should be avoided because once de-authorized it is not possible to unmount.
 Upgrade Clients	Enabled only if one or more authorized nodes is selected in the inventory. Attempts to remotely upgrade the Protector software installed on the selected nodes. The required upgrade installer and configuration files must be present in the `C:\Programs Files \Hitachi \Protector\runtime\updater` folder on the Master node. **Note:** ▪ Only *OS Host* nodes can be upgraded. It is recommended to upgrade nodes in batches of 20. ▪ It is recommended to manually upgrade the master node.
 Delete Node	Enabled only if one or more nodes is selected in the inventory. The node is deleted from the inventory. **Note:** If an *OS Host* node is still running the Protector hub process and is configured to use the current *Master* node, then the node will re-appear as an unauthorized node as it periodically reconnects to the *master* node. Protector should be uninstalled from the node to stop this periodic reconnection.
 Create a new item	Adds a new node to the inventory. The Node Type Wizard (on page 539) is launched to guide you through the process.
Existing Node(s)	Nodes on which Protector has been installed are automatically detected and listed here along side those that have been added by the user. The Node Details (on page 650) is displayed to enable the node's details to be viewed and edited.

Control	Description
	Note: ▪ DHCP renewal can cause temporary disconnection of a node. ▪ If the host of a virtual machine(s) creates a Windows restore point, then the virtual machine(s) can temporarily disconnect. ▪ If a node has been assigned to another master, it will not appear in the inventory.
Filter on Node Name	Filters the displayed results based on Node Name.
Filter on User Tags	Filters the displayed results based on Tags contained in the data flow.
Filter on Node Type	Filters the displayed results based on Node Type.
Filter on Operating System	Filters the displayed results based on Operating System.
Filter on Connection State	Filters the displayed results based on Connection State.
Filter on Authorization State	Filters the displayed results based on Authorization State.

Search Filter using advance Query

Using the Advance Query String you can get more control of how the queries work.

While the Tag Query box does an exact match of any of the tags. So in the Tag Query box you entered tags "vmware" and "location:new york" it would be equivalent to:

```
((userTagsKeys = "vmware") OR (userTags = "location:new york"))
```

There are however extra search methods available using the Advance Query. These include

Value	Description
Exact Match Include	entry = value
Partial Match Include	entry ~ value
Boolean OR	condition OR condition
Boolean AND	condition AND condition

If you wanted only to match any nodes with the string "vm" in the tag but only in location New York, you can use this custom query:

```
((userTagsKeys ~ "vm") AND (userTags = "location:new york"))
```



Figure 6 Nodes Inventory - Search Filter

Trigger Operation

Triggering Operations from the Monitor screen now have the ability to add a tag to the trigger. In a situation for instance where you wanted to create a backup of a system prior to the upgrade, you could do a trigger with a tag “Pre upgrade”. It will be clear then in the restore screen the purpose of this adhoc backup and its timing.

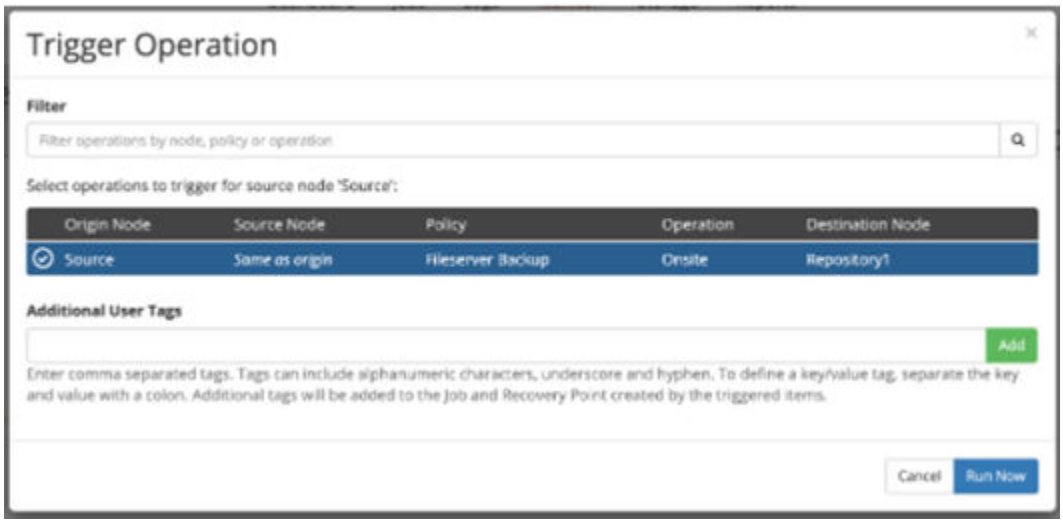


Figure 7 Trigger Operation

Permissions

If you have permission to modify an object e.g. Node, data flow, etc. as governed by the role based access rules, you will have the ability to add / remove tags. Tags belong to the object so any tags are visible to anyone who can view the object.

Tags Limitation

- There is no central inventory of existing tags. So currently there is no pick list or hint that a tag is in use.
- Generation 1 of the Repository does not fully support tags.

Node Concepts

This section describes Ops Center Protector's node management features.

For further information, refer to:

- Node Tasks (on page 261)
- Nodes UI Reference (on page 536)
- Node Groups UI Reference (on page 531)

About nodes

The figure below shows how nodes created within the Protector UI are implemented by the Protector *Client* software installed on servers in the protected environment (see Installation Tasks (on page 234)). These servers communicate directly with physical hardware storage devices and hypervisors that are to be protected or used as backup destinations. One Protector *Client* can communicate with more than one storage device, however when doing this, many factors must be considered, including hardware performance and system availability.

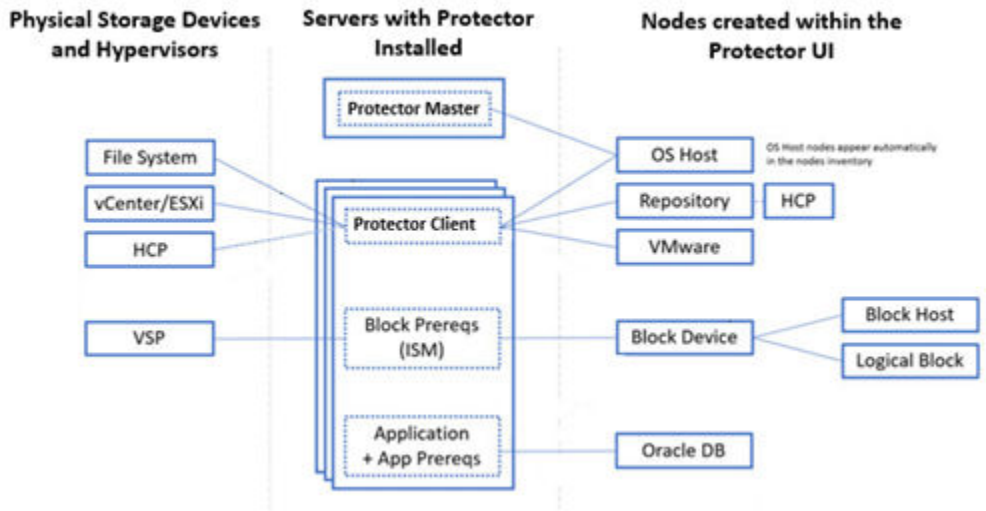


Figure 8 Hardware device nodes, server nodes and logical nodes in a Protector environment

Protector *Client* software must be installed on each server participating in the data protection environment. One server has the Protector *Master* software installed on it and acts as the central controller and the connection point for the web based user interface. Active data protection policies will continue to function with or without the Master being available, because the participating Clients operate autonomously using rules distributed from the Master.

Physical storage devices and hypervisors interface with Protector via servers having Protector *Client* software installed on them. One Protector *Client* can be used to communicate with more than one storage device, however when doing this, many factors must be considered, including server hardware performance and the impact on system availability should that server go down.

The physical environment is represented within Protector as nodes created within the Nodes Inventory (on page 536), these represent data sources and destinations as follows:

- *OS Host* nodes are automatically added to the inventory. These represent every server that has Protector Master or Client software installed on it. A basic OS Host node can only be used as a file system data source on a data flow and can only support path and disk type classifications.
- *Repository* nodes require a Protector Client with access to suitable disk storage capacity to act as a host. Repository nodes act as data destinations on data flows performing backup and continuous data protection operations. Repositories are implemented entirely by software processes within the Protector Client. Depending on the backup data change rate and number of policies, servers that host repositories can be subject heavy network, processor and disk I/O loading, so it is important to specify the correct hardware and monitor performance (see About Repositories (on page 115) for further guidance).
  - HCP nodes do not require a Protector Client to be specified. They act as data destinations on data flows, but must be accessed via a *Repository* acting as a cache when performing tiering operations.
- *VMware* nodes require a Protector Client with network access to the vCenter. These nodes act as data sources on data flows, enabling users to specify VMware specific policies and backup operations to repository or block storage destination nodes.
- *Block Device* nodes require a Protector Client with block prerequisites installed to act as an ISM (see Hitachi Block prerequisites (on page 21)). Block Device nodes act as both data sources and destinations on data flows, enabling users to specify block specific policies with snapshot and replication operations to other Block Device nodes. Note that Block Host and Logical Block nodes can only be used as data sources.
  - *Block Host* and *Logical Block* nodes can only be created by specifying a *Block Device* node upon which they are based. These nodes represent a subset of the resources available on the *Block Device* node upon which they are based.
- *Oracle DB* node requires a Protector Client to be installed on the application server. Appropriate prerequisites (detailed in the relevant Protector Application Guide) must also be installed to enable Protector to interact with the application. These nodes act as data sources on data flows, enabling users to specify application specific policies and backup operations to repository or block storage destination nodes.



WARNING: Do not decommission hardware or virtual machines before removing the nodes representing them.

Removing a node from Protector may require access to the entity that it represents. For example, removing a Block Device node from Protector requires access to the storage system it represents. Therefore, it is not recommended to decommission an entity from the environment until all nodes representing it have been removed from Protector. Removing the nodes from Protector without the hardware may require manual intervention from support.

About node groups

Ops Center Protector allows you to assign machines to groups. A machine can belong to more than one group.

Policies can be assigned to individual machines or to a node group. This feature provides an easy method to manage large sets of machines, eliminating the need to create and assign policies for individual machines.

Once policies have been assigned to a node group on a data flow and activated, the whole group can be monitored easily. You can expand a group to see activity within that group.

About multi-tenancy for Hitachi block storage

Ops Center Protector implements multi-tenancy by limiting which resources each tenant has access to on a block storage device that it shares with other tenants. Each tenant operates within its own environment that is defined using RBAC Resource Groups and their associated Source and Logical Block nodes. Nodes within each tenant's environment are unable to be seen or accessed by other tenants, despite the fact that they share a common block storage device.

Each tenant accesses the common block storage device using its own Logical Block Device node and the associated RBAC Access Control Profiles control which resources are available to each tenant. Configuration of the Logical Block Device nodes is performed by the storage administrator having jurisdiction over that device, using the Hitachi Logical Block Device Node Wizard (on page 609). These Logical Block Device nodes are then used by the backup administrator for each tenant to gain access to the resources allocated to them on the block storage device. The master node is thus able to present only those block storage device resources that it is authorized to use when constructing data flows and policies.

Data Flow Concepts

This section describes Ops Center Protector's data flow management features.

For further information, refer to:

- Data Flow Tasks (on page 228)
- Data Flows UI Reference (on page 355)

About data flows

A *Data Flow* is a diagrammatic representation of the nodes involved in a data protection scenario where each node is represented by an icon. Data flow diagrams identify both physical and logical entities and the connections between them. The data that is to be protected flows from a *Source Node* to a *Destination Node* during the data protection process by way of a *Mover*; the direction of movement being indicated by an arrow on the connector between nodes. Data is transferred in scheduled batches indicated by a solid *Batch* mover. For host based backups, data transmitted across a network can be compressed to reduce bandwidth utilization and bandwidth throttling schedules can be applied to movers, to ensure that data protection activity does not degrade normal network performance.

Node Groups can be placed on data flows so that multiple nodes having common properties can be treated as one entity.

Each node in a data flow plays a part in implementing the data protection scenario by having a *Policy* assigned to it (see About policies (on page 90)). Once a data flow is constructed, it must be *Activated* before it becomes operational. An active data flow can be *Deactivated* to stop that data protection process.

Note: Deactivating a hardware storage dataflow marks replications within it as eligible for being torn down. The actual teardown process must be initiated by the user via the user interface. See About two-step teardown (on page 51).

The process of compiling a data flow performs validity checks on the data flow and assigned policies, then generates a set of rules for each node in the data flow. The compiled rules are distributed to the affected nodes and activated; the participating nodes use these rules to act autonomously. The operation of a data flow can be monitored in real-time using the same data flow diagram rendered as a mimic display (see About monitoring (on page 107)).

Data flow topologies generally fall into one of the following groups (although combinations of these are also possible):

- One-to-one - data from a single source is backed up to a single destination.
- Many-to-one - data from multiple sources is backed up to a single destination.
- One-to-many - data from a single source is backed up to multiple destinations.
- Many-to-many - data from multiple sources is backed up to multiple destinations.
- Cascaded - data from a source is backed up to one destination then forwarded on to a second destination.

About two-step teardown

Two-step teardown reduces the possibility of inadvertently tearing down Block Storage replications due to accidental deactivation of a dataflow, or activation of an erroneous dataflow.

With two-step teardown, on deactivation of a dataflow in which a replication operation appears, or reactivation of a data flow where a replication operation has been removed, the replication is now flagged as being eligible for teardown in the **Restore** and **Storage** inventories. When a replication operation is deactivated, the underlying replication on the hardware will continue to operate as normal, except that any further batch resynchronizations will not be scheduled.

The final teardown operation must be explicitly initiated by the user, via the **Storage** inventory.

If a user initiated teardown operation fails, it is not automatically retried, so needs to be re-initiated by the user. Teardown failure may occur in the case of GAD 3DC dataflows, where teardowns must be performed in a specific order. Prior to the introduction of two-step teardown, automatic retries were performed indefinitely until successful.

Note: A dataflow that is deactivated and then subsequently reactivated without the replication operation having been removed, re-instantiated or manually torn down before reactivation will, in effect, be re-adopted. This re-adoption removes the *eligible for teardown* flag from the replication, as if the dataflow had never been deactivated.

It is possible to disable two-step teardown on a per ISM basis via a configuration file, so different teardown policies can be implemented within the same environment.

About data flow implementation

Data flows drawn in the Protector UI are abstractions that hide the underlying hardware and software implementation. Before a data flow can be constructed, Protector Client software must be installed on the physical nodes connected to the storage devices, and the equivalent nodes must be created in the Nodes Inventory. The relationship between storage devices, Clients and nodes appearing in the Nodes Inventory is described in About nodes (on page 47).

As an example of how a data flow is implemented, consider the backup of a file path on a server to a repository. In this case, the source and destination storage devices are file system volumes mounted on separate servers. Protector Client software must be installed on both servers. These Protector Clients are automatically detected by the Protector Master and appear in the Nodes Inventory as *OS Hosts*. The source *OS Host* can be used 'as-is' to identify the volume and file path in the backup policy. The *Repository* node is created via the UI; the destination *OS Host* node being selected as the proxy when configuring the node. When the rules for this data flow are activated, the source server will transmit the files identified in the policy over the network to the repository on the destination server. The figure below shows the data flow as it appears on the UI with the underlying implementation shown beneath.

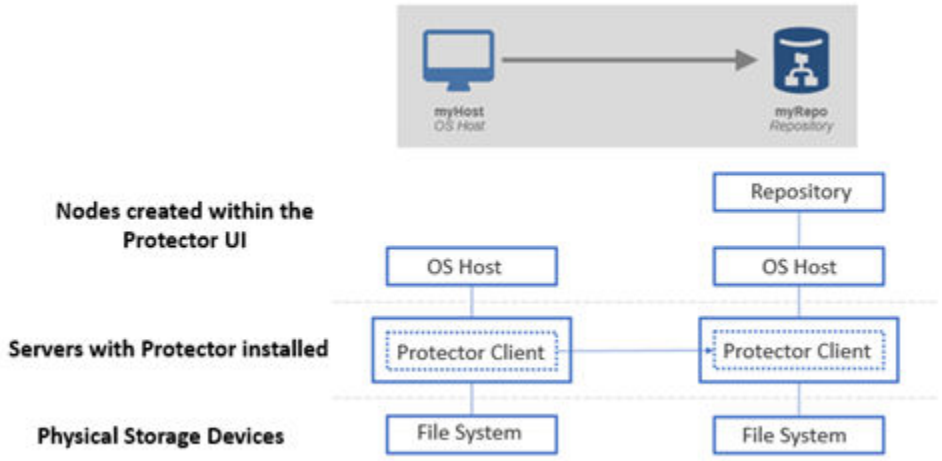


Figure 9 Files on a server backed up to a repository

A more complex example is a data flow representing application data, stored on a source block storage device, being replicated to a destination block storage device. The source block device has LDEVs mounted to an application server machine that are being used to store the application data. Protector Client software must be installed on the application server and on servers at the source and destination sites that are designated to control the block devices. The application server must have prerequisites installed to enable Protector to interact with the application software. The servers (ISMs) at the source and destination sites must have prerequisites installed to enable Protector to interact with the block devices. All three Protector Clients will be automatically detected by the Protector Master and appear in the Nodes Inventory as *OS Hosts*. Source and destination *Block Device* nodes are created via the UI; the source and destination *OS Host* nodes representing the ISMs being selected as proxies when configuring the nodes. When the rules for this data flow are activated, the source (VSP 1) will replicate the LDEVs identified in the policy over the data link to the destination (VSP 2). The figure below shows the data flow as it appears on the UI with the underlying implementation shown beneath. Note that the source *Block Device* node is required by Protector to control the replication, but does not appear on the data flow.

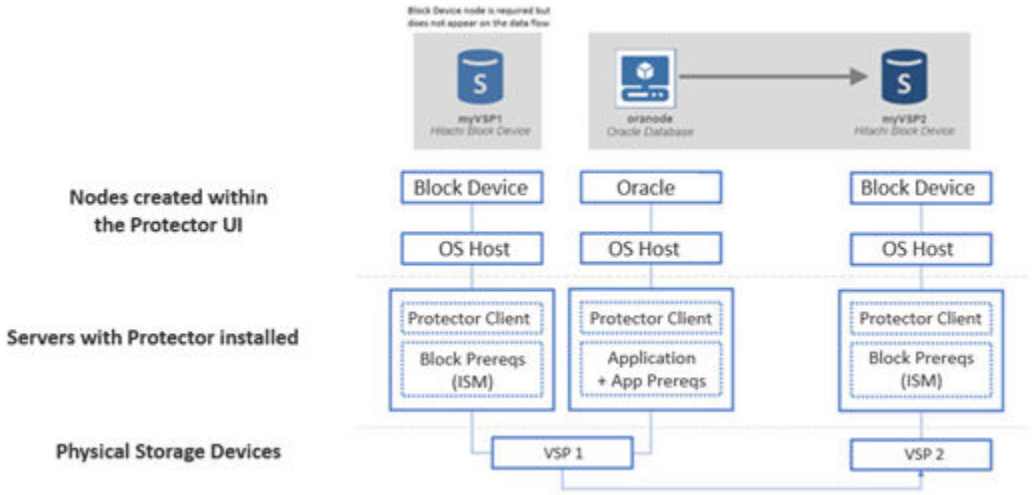


Figure 10 Block based application replication

About many-to-one data flow topologies

Many-to-one topologies should be used only where absolutely necessary. These topologies allow data from multiple nodes to be stored on a single destination node. For example, two or more source nodes replicate to a common backup server.

Consider whether the source nodes should be treated as a cohesive *node group* rather than individually.

If the source nodes are not functionally related then consider drawing separate data flows for each one; this will allow rules activation/deactivation for each flow to be decoupled.

Note: Many-to-one policies preclude the use of the *Mount* operation (for proxy back-up or re-purpose) on a destination Hitachi Block device. A *Mount* operation is only valid when one replication is applied to the destination node.

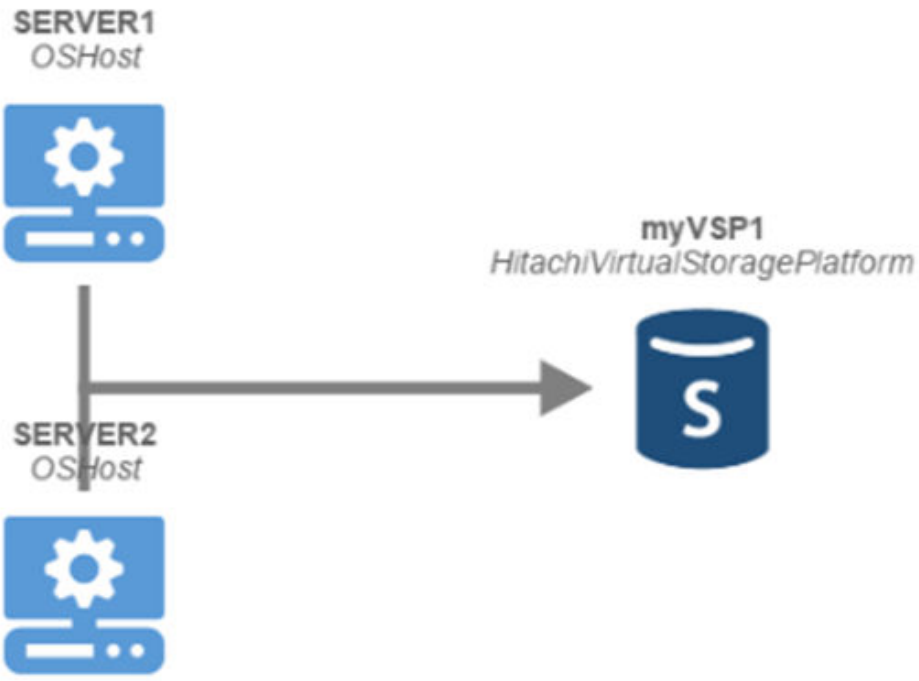


Figure 11 Many-to-one data flow

About one-to-many data flows

One-to-many topologies are one of the standard building blocks for creating complex data flows. These topologies allow data from one node to be stored on multiple destination nodes. For example, one source node replicates to a DR server and a repurposing server. A policy is applied to the source node and all destination nodes.

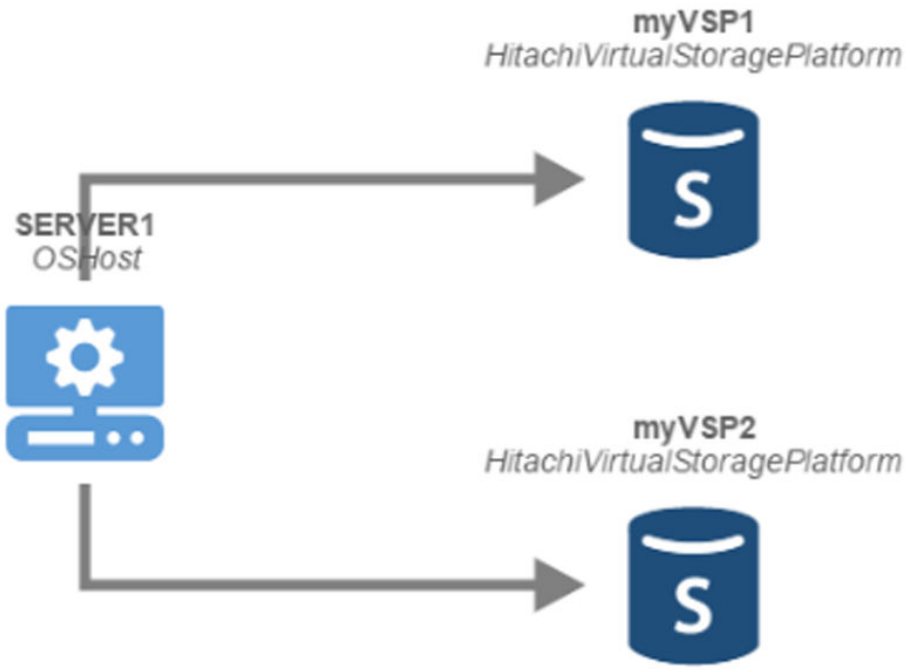


Figure 12 One-to-many data flow

About cascading data flows

In a cascading topology, a source node sends data to a primary destination node and then on to a secondary destination node. The primary use case is to copy the data to an intermediary site and then onto a remote site. The secondary use case is for local repurposing, where multiple copies of the same data can be mounted to different servers and used by different sets of users. The primary and secondary destination nodes both store the data. By sending all data to local and remote sites, the disaster recovery options are improved. A policy is applied to the source node, the local and the remote destination nodes.



Figure 13 Cascading data flow

About parallel versus serial data flows

Ops Center Protector is designed to allow data to be filtered according to defined policies, and for data to be either sent directly to a node or routed through one or more intermediate nodes.

Typically, a set of policies can be implemented by using parallel or serial data flows with different tradeoffs. Consider a scenario where the following backup policies are required:

- The entire contents of a source machine need to be backed up to a destination machine for disaster recovery purposes.
- Only databases from that same source machine need to be backed up to a different destination machine for testing purposes.

About the parallel data flow solution

For the parallel solution the source node has two policies assigned; one for the whole machine (replicating all attached logical devices to myVSP1) and one for the databases (replicating only a subset of those logical devices to myVSP2). Because there are two movers in parallel, data that is common to both policies must be transmitted by the source node twice. This could be considered wasteful as it doubles the computing load and network bandwidth usage. However:

- This topology is more fault tolerant since one destination node does not depend on the other destination node to implement its policy.
- For *Host based* data flows, very fine grain data classification filters, separate bandwidth throttling and/or scheduling windows can be applied to each branch of the data flow.

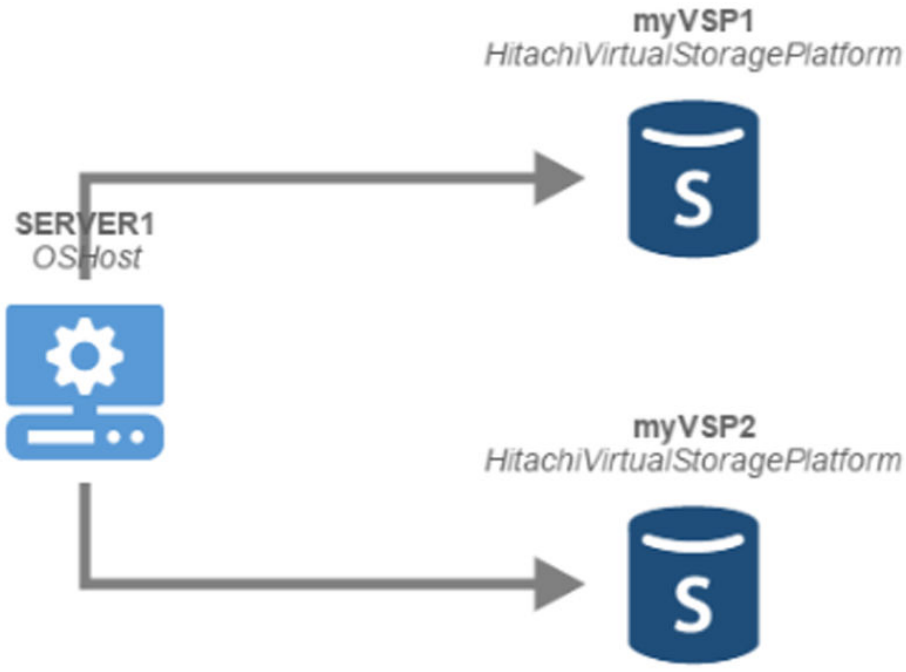


Figure 14 Parallel solution

About the serial data flow solution

For the serial solution the same policies are applied as in the parallel solution, the difference is that the data going to the second destination node is routed through the first destination node. This reduces the data traffic on the source node since it only has to transmit it once. However the policy on the second destination node now requires the first destination to be running and available on the network.



Figure 15 Serial solution

About best practices for drawing data flows

It is important to create data flows that are well formed, understandable, manageable and maintainable. To achieve these goals it is advisable to adhere to a number of guidelines when constructing them.

The Data Flow Wizard (on page 361) has an intentionally small workspace area. This is done to dissuade users from drawing overly complex data flows. It is possible to pan and zoom the workspace to aid working on lower resolution displays. This feature should not be used to create a large drawing area.

Data flows should always be drawn with data flowing in the reading direction appropriate for the locale. The general convention for such diagrams is left to right. Although top to bottom data flows can be drawn, mover labels will not be displayed in this format.

Position nodes on diagrams so that adequate space is left between them for the operation name, operation type and mover label to be drawn below the movers.

Make full use of the ability to name and describe nodes, node groups, policies, operations, data flows, movers and schedules. Use consistent naming conventions and descriptions that convey intent; this will greatly enhance understanding of the purpose of a data flow for other users and will aid ongoing maintenance. Devise and adhere to a common naming convention based on the guidelines described in About best practices for naming objects (on page 59).

A central principle for determining which nodes and policies should be placed on each data flow is to consider the granularity with which you want to be able to activate and deactivate individual data protection policies. If you combine too many flows and policies into a single diagram, then it will not be possible to deactivate a single policy without deactivating many others. For this reason it is recommended that you:

- Consider how you may want to activate and deactivate separate policies.
- Aim, wherever possible, to have only one flow and one policy per diagram.
- Only place related nodes(1) and policies on the same data flow. Note that a policy can contain multiple operations, and for hardware orchestration data flows, all operations on a PVOL should be contained within the same data flow.
- Separate unrelated policies sharing the same data flow into separate diagrams.
- Only use one-to-many topologies if all the participating nodes and policies are related.
- Consider placing multiple instances of the same node on one diagram to improve presentation.
- Use node groups for source nodes that share identical policies.
- Use separate data flows in place of many-to-one topologies, unless absolutely necessary.

1. Related nodes and policies typically form part of the same business process.

About best practices for naming objects

Providing meaningful names for the objects you create in Protector is one of the best ways to ensure that your data protection strategy is understandable, manageable and maintainable. The following guidelines will help you achieve this:

- Adopt a naming convention that results in similar concepts being listed together when sorted alphabetically. For example `VSPCorporate` and `VSPUKOffice` are preferable to `CorpVSP` and `UKOfficeVSP`.
- In general, Protector names should use nouns to describe objects and verbs to describe operations that are specific to the business processes in which they are used rather than Protector terminology.
- Names should be unique enough so as to avoid misunderstanding and confusion.

  For example: `Site1Broadstone` and `Site11London` are preferable to `Site1` and `Site11`.
- *Node* names should describe the node's role in the organization and possibly its location. Avoid describing its type or other generic attributes that can be determined from the Node Details (on page 650).

  For example: `VSPTexasSalesCorp` is preferable to `PrimaryVSPG400`.
- *Node Group* names should describe the node group's role in the organization and possibly its location. Avoid describing its type or other generic attributes that can be determined from the Node Group Details (on page 535).

  For example: `ClusterDbLegal` is preferable to `OracleRAC`.
- *Policy* names should describe what the policy does, specifically in the context of the organization's data protection strategy. Avoid describing its classifications, filters, operations, scheduling or other generic attributes that can be determined from the Policy Details (on page 737).

  For example: `OracleDRMajorHeadOffice` is preferable to `UniversalRepOraclePolicy`.
- *Operation* names should describe what the operation does, specifically within the context of the policy in which it is used. Avoid describing its type, scheduling or other generic attributes that can be determined from the Policy Operation Details (on page 739).

  For example: `ReplicateLocal` and `ReplicateOffsite` are preferable to `Replicate1` and `Replicate2`.
- *Data Flow* names should describe the data protection policy that is being implemented, specifically within the context of the organization's data protection strategy. Avoid describing attributes that can be determined from the Data Flows Inventory (on page 355), Data Flow Details (on page 491) or Monitor Details (on page 521).

  For example: `OracleFailOverToSubOffice` is preferable to `ReplicateOracleDataFlow`.
- *Mover* labels should describe the purpose of the mover and/or the association between the nodes it connects. Avoid describing the operation name or operation type since these already appear in the Data Flow Details (on page 491).

  For example: `Fine Grain Protection` is preferable to `Weekly backup`.

- *Schedule* names should describe the purpose of the schedule. If the schedule is related to a specific policy then name it after that policy. Avoid describing its type or other attributes that can be determined from the Schedule Details (on page 857).

  For example: `BackupWindowOvernight` is preferable to `Weekdays10PMTo6AM`.



Note: The naming conventions used elsewhere in this guide are designed to aid understanding of Protector concepts, tasks and workflows. They do not represent good naming conventions when using the product in a production environment.

About Repository and HCP based backups

Repository-based backups are created by replicating data to a disk-based store. Restore point snapshots are created in the repository as required. The repository snapshot is assigned a retention period and is available for full or partial restore until it is retired (deleted from the store). A Generation 1 repository store contents, created using a batch mover, can also be tiered to HCP for long term retention. For Generation 2 HCP, there is no need for a repository between the source node and HCP, backups can go straight from the source node to HCP. However, tiering is not supported on Generation 2 HCP.

The concept of Generation 1 and Generation 2 HCP nodes applies to Protector only and differentiates between an older HCP interface implementation and the current one. Generation 1 HCP is deprecated and should not be used for new configurations. In addition, you cannot upgrade Generation 1 implementations to Generation 2. There is difference between a backups to Generation 2 HCP and to HCP Cloudscale.

The concept of Generation 1 and Generation 2 repositories is also a Protector concept. Again. Generation 1 repositories are deprecated and should not be used for new configurations and cannot be upgraded to Generation 2 repositories. Generation 2 repositories have many advantages over Generation 1 repositories especially around parallelism.

About Repository based batch backup

The repository is updated periodically using scheduled resynchronizations. This method involves a scan of the source machine’s file system. With batch resynchronization, the changed blocks are transferred. By default, the block size that is transferred is 2 MB; it is 16 KB if fine change detection is configured in the store template for Gen 1 repositories or in the Policy definition for Gen 2 repositories. An entire file is transferred if it is less than one block. Batch backup is useful for data that does not change often, such as data contained on the operating system disk.

Fine change detection should be used with the batch mover if there are frequent, widely distributed changes to files. Typically however, content on file servers does not benefit from fine change detection, because Microsoft Office documents are automatically compressed by the applications. This means that the smallest of modification causes a completely different file to be created, negating any benefit from detecting changed blocks.

Host based granular file I/O data capture

The data capture system is highly granular, with the option to tightly define the type of data that is transferred within the policy classification. Unlike systems that replicate all contents on the volume, Ops Center Protector's *Host Based* data protection technologies replicate only specified data, thereby saving bandwidth and storage space.

About Repository based source side deduplication

Source side deduplication is a mechanism within Protector that improves network bandwidth utilization by avoiding sending the same data multiple times.

In simple terms, as soon a block on the source machine is detected as changed during post-scan in a batch policy, those blocks are transferred to the repository. Source side deduplication uses Single Instance Store (SIS), a one-by-one method of scanning/ transferring data from source to repository during initial synchronization and subsequent resynchronizations. If, for example, twenty 5TB source nodes (each having much of the data duplicated across them) are backed up, the repository will ingest data from the first source node, go to the second source node and compare its data with what is already ingested, then transfer only non-duplicate data. This process is repeated with the remaining source nodes.

Even if some duplicate files get through, with source side dedupe enabled, SIS post processing will remove those duplicates. SIS post processing is applied to the dataset once per day during repository cleanup.

The changed block transfer method is far more efficient and has less impact than traditional backup systems.

About Repository to Repository backup

Note: Repository to Repository backups are only supported within the same generation of repository i.e. Generation 1 to Generation 1 or Generation 2 to Generation 2.

Ops Center Protector has the ability to replicate data being sent to an on-site repository, to an off-site repository without needing to gather the data from the source machine(s) again. This minimises the load on the source machines and quickly and efficiently transfers the data off-site.

Using Smart Repository Sync, the secondary repository creates backups from the primary repository by taking the incremental changes required to create a new snapshot; the source node is not involved in this process. The secondary repository has the following capabilities:

- The user selects which backups are sent to the secondary repository. You may choose to only send a subset of the primary.
- You do not have to use the same schedule. You may for instance back up to the primary repository every hour but to the secondary once a day.
- You can have different retention. You may for example keep backups on the first repository for a week and on the secondary repository for 6 months.

Repository to repository backups should be scheduled such that the secondary (off-site) repository takes the latest completed backup from the primary (on-site) repository. The secondary repository backup is therefore scheduled to run on completion of the primary's backup.

When the policy is first triggered, the on-site repository will be resynchronised with the source. The empty off-site repository will then be synchronised with the on-site repository. Depending the amount of backup data and the bandwidth of the network between the on-site and off-site repositories, this initial synchronisation process can take a considerable time (many hours) to complete. To overcome this, a technique called 'repository seeding' (see How to seed an offsite repository (on page 152)) may be used for efficiently setting up an off-site data store, reducing the time and bandwidth required to load the initial backup into the secondary repository. Once seeded, the amount of data transferred between the on-site and off-site repositories is much reduced and dependent only on the data change rate of the source machine(s).



Note: If the on-site repository contains large amounts of non-critical or legacy data that does not require additional off-site protection, then it is recommended that you review your local backup policies and repository architecture prior to replicating data to the off-site repository. This will allow you to identify and replicate only your critical data to the off-site location.

About tiering Gen1 Repositories to HCP



Caution: Data that is tiered from an encrypted repository will not be encrypted on HCP. The use of encrypted repositories for tiering is not recommended.



Note: Tiering file system data to HCP is backwards compatible with Protector 5.x. However the following features are not yet supported in Protector:

- Tiering from a live data store (only batch data stores are supported).
- Stubbing and removing data from the source.
- Setting retention on HCP objects.

Setting DPL (Data Protection Level) on HCP namespaces is no longer done through Protector. This is now done at the tenant level via HCP's Tenant Management Console.

File system data from any supported OS, that is backed up to a Protector repository using a batch mover, can be tiered to HCP cloud storage platform. A repository store is tiered to an associated HCP namespace. Each tiered repository stream (a file can consist of one or more streams) is saved as an HCP object within its related namespace, along with metadata that enables it to be restored back to the original source node. If a previously tiered repository stream is modified, then the entire stream is tiered to HCP again. Once a stream is tiered to HCP it is removed from the repository.

Repository ingestion rate throttling helps to constrain repository data store size growth. If a repository's tiering queue gets too large, the repository will stop receiving data from the source until the tiering queue length reduces. High and low watermarks control the growth of a repository, allowing newly ingested files to occupy the space that tiered files previously occupied. By default up to 50 streams can be tiered concurrently to HCP, with another 50 queued awaiting tiering, before repository ingestion is paused. If performance outweighs repository growth in your environment, please contact Protector product support to adjust the throttling behaviour.

Protector does not set a retention on HCP objects. When a tiered stream is no longer referenced by any repository snapshots, the corresponding object will be deleted from HCP. A repository must be present in an active data flow and configured to tier to the HCP to be able to delete objects. When a repository store is deleted, the corresponding namespace for that store will be deleted from HCP.

Deleting a repository node will not remove the repository streams or the data tiered to HCP. To remove all Protector's data tiered to HCP after deleting a repository node, use the **`HCPDM`** utility provided by HCP.



Tip:

Protector repository stores are mapped to HCP namespaces, with the UUID of the Protector store being used as the HCP namespace name.

NTFS files typically consist of at least two streams containing data and security information respectively. Each repository stream is mapped to an HCP object, with an incremental hexadecimal number being assigned by Protector as the object name. Objects are stored in HCP in their native format, so it possible to view them in HCP. Object naming is designed to distribute the load optimally across HCP cluster nodes.

Repository stream objects are described by an HCP content class to enable them to be indexed and subsequently searched using the HDIM Content Class and its associated properties in an HCP Structured Query. Notice the use of the legacy product name HDIM here for backward compatibility; Protector was once named HDIM.

Protector communicates with HCP using the native REST API over HTTPS by default. HTTP can also be used to increase tiering speed if performance needs outweigh security.

About Hitachi Block based backup technologies

Ops Center Protector supports the following Hitachi Block based snapshot and replication technologies, allowing data flows to be constructed graphically without the user needing to construct HORCM files:

- Thin Image (TI)
- ShadowImage (SI)
- TrueCopy (TC)
- Universal Replicator (UR)
- Global-Active Device (GAD)

These technologies can be combined in numerous ways to create complex, block based, data protection and repurposing scenarios.

Ops Center Protector is used to create these data protection policies and manage them after replications and snapshots are created. It offers the ability to view the current state of replications and control their activities. It fully manages the lifecycle of snapshots, keeping an index of their existence and removing them from the system once the user designated retention has been reached.

Tip: When Protector creates fully provisioned snapshots and replications, it does so in S-VOLs mapped into one or more host groups that it creates on a port, so that these S-VOLs can have LUNs assigned. For GAD replications, the user is additionally able to specify one or more host groups that will have a LUN assigned for the S-VOL.

Caution: Anything created by or imported into Protector is managed by Protector and should not be modified or deleted without doing so through Protector. If a replication is removed in Protector then it will be removed from the hardware storage device too.

About mover types used with Hitachi Block operations

When constructing a data flow containing Hitachi Block based storage devices, it is necessary to use the correct combination of mover type (*Batch* or *Live/Continuous*) in conjunction with a given snapshot or replication technology (Thin Image, Refreshed Thin Image, ShadowImage, TrueCopy, Universal Replicator or Global-Active Device).

The term *Differential Snapshot* means that a new target volume(s) is created each time the operation is triggered. The target volume(s) remain suspended after creation. The target records the deltas so the source is also required to reconstruct the full data set.

The term *Refreshed Snapshot* means that a new static target volume(s) is created the first time the operation is triggered. The same target volume(s) is refreshed with new deltas from the source on subsequent triggers. The target volume(s) remain suspended after creation. The target records the deltas and thus the source is also required to reconstruct the full data set.

The term *Batch Replication* means a static target volume(s) is created the first time the operation is triggered. After initial synchronization, the replication is suspended. Subsequent triggers result in a resynchronization from the source, after which the target volume(s) are suspended.

The term *Live Replication* means a static target volume(s) is created the first time the operation is triggered. The source and target volume(s) are continuously kept in sync (paired) either by copy-on-write (COW) or copy-after-write CAW) data transfer mechanisms. The replication pairs can be paused and resumed as required.

The following table lists typical block based data flow scenarios along with the move types that can be used.

Table 1 Hitachi Block based scenarios and associated mover types

Scenario	Mover	Description
Thin Image Snapshot (TI)	Batch	About Thin Image and Thin Image Advanced differential and refreshed snapshots (on page 66)
ShadowImage Replication (SI)	Batch or Live	About ShadowImage replication (on page 70)
TrueCopy Replication (TC)	Live	About TrueCopy replication (on page 74)
Universal Replicator (UR)	Live	About Universal Replicator (on page 75)
Global-Active Device (GAD)	Live	About Global-Active Device replication (on page 76)
Three Data Centre Cascade (TC+UR)	Live TC + Live UR	About three datacentre cascade (3DC) (on page 78)
Three Data Centre Multi Target (TC+UR)	Live TC + Live UR	About three datacentre multi-target (on page 79)
Clones of a Clone (SI+SI)	Batch SI(1) + Batch SI	About static clones of a clone (on page 81)
Clone with Replication (SI +GAD)	Batch SI(1) + Live GAD	About clone with replication (on page 81)
Snapshot with Replication (TI+GAD)	Batch TI + Live GAD	About snapshot with replication (on page 82)
Snapshot of a Clone (SI+TI)	Batch SI(1) + Batch TI	About snapshot of a clone (on page 83)
Replication of a Clone (SI +TC or SI+UR)	Batch SI(2) + Batch TC or UR	About replication of a clone (on page 83)
Remote Snapshot (TC+TI or UR+TI)	Live TC or UR + Batch TI	About remote snapshot (on page 84)
Local and Remote Snapshots (TC+TI/TI, UR +TI/TI or GAD+TI/TI)	Live TC, UR or GAD + Batch TI	About local and remote snapshots (on page 84)

Scenario	Mover	Description
Remote Clone (TC+SI or UR +SI or GAD+SI)	Live TC, UR or GAD Batch SI(1)	About remote clone (on page 85)
Local and Remote Clones (TC+SI/SI or UR+SI/SI or GAD+SI/SI)	Live TC, UR or GAD Batch SI(1)	About local & remote clones (on page 86)
Local Snapshot and Remote Clones (TC+TI/SI or UR+TI/ SI or GAD+TI/SI)	Live TC, UR or GAD + Batch TI + Batch SI(1)	About local snapshot and remote clones (on page 86)
Remote Snapshot and Local Clones (TC+SI/TI or UR+SI/ TI or GAD+SI/TI)	Live TC, UR or GAD + Batch SI(1) + Batch TI	About local snapshot and remote clones (on page 86) (reversed)

(1) Continuous SI can also be used in these topologies to vary the use case.

(2) Continuous SI cannot be used in these topologies since it is not possible to chain a remote replication from Continuous SI.

About Thin Image and Thin Image Advanced differential and refreshed snapshots

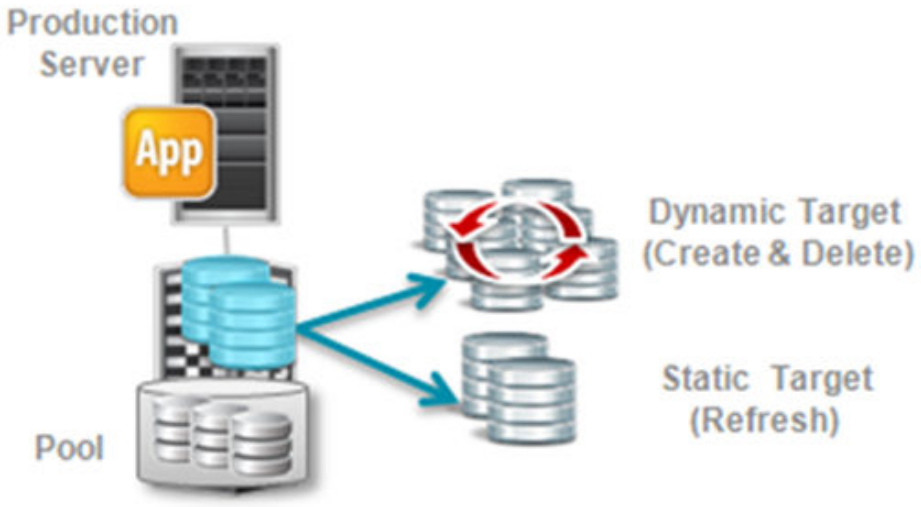


Figure 16 Differential and Refreshed Thin Image Snapshot

Thin Image and Thin Image Advanced enables rapid creation of in-system, space efficient, read/write, volume-consistent snapshots and subsequent rollback of entire volumes.

When handling multiple primary volumes, the storage system takes a snapshot of each volume sequentially. This means that a slight difference can be seen across the snapshot timestamps. If exactly the same timestamp is required among the snapshot set (for crash-consistent backup), the snapshot set should be created using Consistency Groups (CTGs).

When taking a snapshot, there are two options for the target volumes:

- Differential Snapshot: Creates a new snapshot for each backup and deletes it when the retention time expires. This simplifies the management of numerous snapshots and is suitablefor backup operations. Data Retention Utility (DRU) protection can be applied to the snapshot's LDEV so that it can be used as a read-only volume or to protect it against both read and write operations.
- Refreshed Snapshot: Creates a new snapshot for the first backup, and then resynchronizes it on the following backups. This enables static target volumes (Port, Host Groupand LUN, although the LUN may not remain constant depending on the mount host OS) and is suitable for repurpose operations.

Floating Device is an improved snapshot capability, used in conjunction with Thin Image andThin Image Advanced, that simplifies snapshot management. With this capability snapshots can be created without creating target volumes upfront. This means that the limit on the number of snapshots in the entire storage system is increased (the number of snapshots of a specific primary volume is 1024). To revert the snapshot it is only necessary to select the required timestamp. To mount the snapshot for re-purposing, it must be mapped to a specific LDEV/LUN. After the snapshot is un-mounted, the volumes will be deleted by Protector as part of the unmount process.

The storage system supports two types of snapshots. Thin Image and Thin Image Advanced, the storage system will automatically decide the type of snapshot used. For Data Reduction Shared (DRS) volumes Thin Image Advanced snapshots will always be used and for all other types of volumes Thin Image snapshots will be used.

About Thin Image Snapshots:

Snapshots of a volume in the storage system are stored in a dedicated area called the Thin Image pool. For floating snapshots, where no LDEV is assigned until it is mounted, no data movement occurs and thus creation of a snapshot is near instantaneous. For non-floating snapshots, the auxiliary tasks of creating an LDEV and a LUN will take some time.

Once a snapshot is created, subsequent updates to the primary data causes the storage system to move the data blocks being updated to the TI pool. TI pool usage thus increases as the primary data changes and snapshots are retained.

Caution:

Filling a Thin Image pool to capacity will invalidate all snapshot data contained within that pool. All snapshots in the pool will have to be deleted before snapshotting can be resumed.

When accessing a snapshot, the storage system presents the virtualized contents by merging the primary volume with its differentials held in the TI pool.

When deleting a snapshot, the storage system releases the differentials held in the TI pool. When multiple snapshots are involved, the delete operation may take some time. The differentials are reference counted and will only be deleted when no remaining snapshot requires them.

Note: Once snapshots have been released back into the pool, that space can only be reused by the same primary volume. This hardware limitation means that this space cannot be used by a different volume. The only way to completely free the space to the pool for any volume is to delete all the snapshots on that primary volume.

When handling multiple primary volumes, the storage system takes a snapshot of each volume sequentially. This means that a slight difference can be seen across the snapshot timestamps. If exactly the same timestamp is required among the snapshot set (for crash-consistent backup), the snapshot set should be created using Consistency Groups (CTGs).

When taking a snapshot, there are two options for the target volumes:

- Differential Snapshot: Creates a new snapshot for each backup and deletes it when the retention time expires. This simplifies the management of numerous snapshots and is suitable for backup operations. Data Retention Utility (DRU) protection can be applied to the snapshot's LDEV so that it can be used as a read-only volume or to protect it against both read and write operations.
- Refreshed Snapshot: Creates a new snapshot for the first backup, and then resynchronizes it on the following backups. This enables static target volumes (Port, Host Group and LUN, although the LUN may not remain constant depending on the mount host OS) and is suitable for repurpose operations.

Floating Device is an improved snapshot capability, used in conjunction with Thin Image, that simplifies snapshot management. With this capability snapshots can be created without creating target volumes upfront. This means that the limit on the number of snapshots in the entire storage system is increased (the number of snapshots of a specific primary volume is 1024). To revert the snapshot it is only necessary to select the required timestamp. To mount the snapshot for re-purposing, it must be mapped to a specific LDEV/LUN. After the snapshot is un-mounted, the volumes will be deleted by Protector as part of the unmount process.

About Thin Image Advanced Snapshots:

Snapshots of a volume in the storage system are stored in the same pool as the primary. When a snapshot is taken the data mapping of the primary volume is copied. Once a snapshot is created, subsequent updates to the primary data cause the storage system to redirect the new data blocks to new storage in the primary pool. Primary pool usage will increase as the primary data changes and snapshots retain the old data blocks.

Note:

- When using a snapshot, intensive read/write access to the snapshot may impact the performance of the primary volume, due to the way the snapshot volume is virtualised. If this is of concern then consider using ShadowImage or ShadowImage-Thin Image in cascade, where the Thin Image primary volume becomes the ShadowImage secondary volume instead of the original source.
- Long-term retention increases the number of differentials held in the TI pool once writes are distributed across all data blocks. Thus, for long-term backups, it is recommended to use ShadowImage.
- Thin Image requires the primary data in order to present a virtualised snapshot volume. This means that snapshots will be lost if a disk failure occurs on the primary volumes. To protect the data from such hardware failure use ShadowImage to create a clone and take snapshots of the clone instead.

About cascaded Thin Image and Thin Image Advanced snapshots

Thin Image and Thin Image Advanced snapshots of a P-VOL can be cascaded, the first layer being referred to as L1 S-VOLs. Cascading can be recursive so as to form a snapshot tree to a depth of up to 64 layers (L64 S-VOLs), consisting of the *root* P-VOL, intermediate *node* S-VOLs and terminal *leaf* S-VOLs. The total number of S-VOLs in a tree is limited to 1024.

To create cascadable snapshots, the L1 snapshot volumes must be created in cascade mode as a floating device or fully provisioned. Cascade mode snapshots must be provisioned, either at creation or mount time from a dynamic pool. From Protector 6.5 onwards, L1 snapshots are created in cascade mode by default, and are dynamically provisioned, although standard mode can still be specified if the storage device does not support cascading.

For Thin Image Advanced snapshots Protector always uses the pool of the source P-VOL for the cascade mode snapshot, however for Thin Image snapshots it uses a number different of pools for cascade mode snapshots as follows:

- *Snapshot Pool* - a Thin Image or hybrid pool where the P-VOL/L1 and L1/L2 snapshot pair data is held. If a hybrid pool is specified then Protector may also create the snapshot S-VOLs here.
- *Cascade Pool* - a dynamic or hybrid pool where Protector creates snapshot S-VOLs if they are fully provisioned.
- *Mount Pool* - a dynamic pool where Protector creates the snapshot S-VOLs if the *Snapshot Pool* is a Thin Image pool or if a floating device was specified for the snapshot operation. If a *Mount Pool* is specified as an option then it will be used in preference.

For Thin Image both standard and cascade mode snapshots require a *Snapshot Pool* to be specified regardless of the mode. If fully provisioned cascade mode is selected then a *Cascade Pool* must be specified when configuring the operation.

When mounting a cascade mode snapshot, Protector provides the option to mount the original (L1) or a duplicate (L2) snapshot. The duplicate (L2) snapshot can be modified without changing the original (L1) snapshot's data. When the duplicate (L2) snapshot is unmounted, it is deleted and any changes made to it are lost. For Thin Image snapshots original and duplicate mount modes for cascade mode snapshots may or may not require a *Mount Pool* to be specified, as per the following table:

Snapshot Pool Type	Provisioning Type	Mount Mode	Specify Mount Pool?
Thin Image	Floating Device	Original (L1)	Required
Thin Image	Floating Device	Duplicate (L2)	Required
Thin Image	Fully Provisioned	Original (L1)	N/A
Thin Image	Fully Provisioned	Duplicate (L2)	Optional
Hybrid	Floating Device	Original (L1)	Optional
Hybrid	Floating Device	Duplicate (L2)	Optional
Hybrid	Fully Provisioned	Original (L1)	N/A
Hybrid	Fully Provisioned	Duplicate (L2)	Optional

About ShadowImage replication

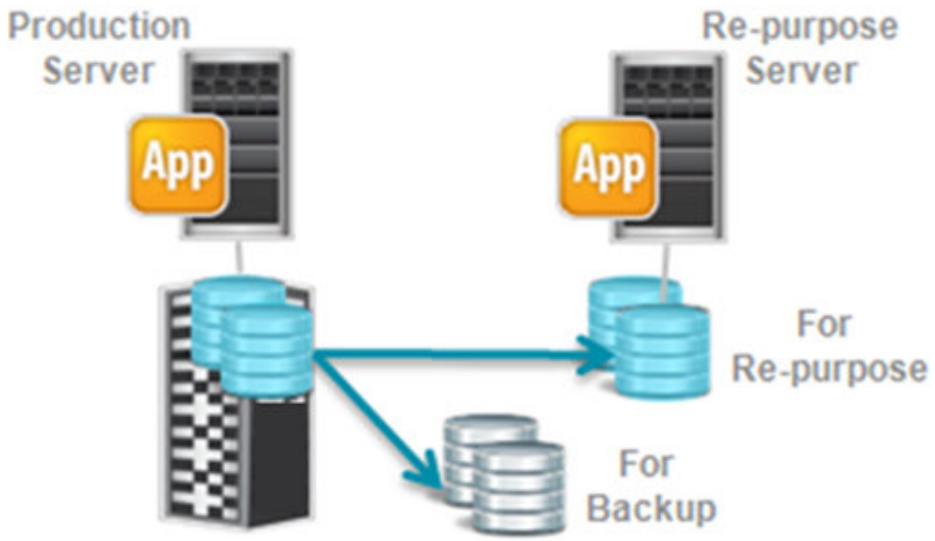


Figure 17 Full clone using batch mode ShadowImage

ShadowImage enables the creation of in-system, RAID-protected, read/write, volume-consistent, full clones.

As with TI snapshots, consistent clones can be created using Consistency Groups (CTGs).

Note: ShadowImage has a limitation on the maximum number of clones that can be created at one time. There can be up to three 1st level (L1) clones and then two L2 clones per L1 clone, giving a potential total of six L2 clones. Including the L1 clones, the potential total is nine clones. If more copies are required beyond this then use Refreshed Thin Image snapshots.

When taking a clone of a primary volume the storage system copies all of the data to the secondary volume. The point at which this is done depends on the split type selected:

- *Quick Split* - copying from primary to secondary is performed in the background so that the secondary is immediately available for reading/writing. The performance of the primary may be affected if access the secondary references data that has not yet been copied from the primary. In this case, on-demand copying of that data from the primary is required.
- *Steady Split* - copying from primary to secondary is performed in the foreground before the secondary is made available for reading/writing. The creation of the secondary takes time depending on volume size.

If using Dynamic Provisioning (DP) volumes for both primary and secondary volumes, the copy is applied only for the allocated area; the unallocated area is ignored.

Once the clone is created, the storage system updates the bitmap for the primary, which records which blocks have been modified. Pair resynchronization can be performed in one of the following ways:

- *Quick Resync* - resynchronization is performed in the background and on-demand. The secondary is briefly made read only (for less than 1 second), after which it becomes available for reading/writing (i.e it enters the PAIR state in less than 1 second). The performance of the primary may be affected if access to the secondary requires on-demand resyncing from the primary.
- *Normal Copy* - the secondary is made unavailable while the resynchronization is performed. The resync takes time depending on the size of differentials between the primary and secondary.

When the secondary is accessed, behaviour depends on the mode of operation as follows:

- *Steady Split* and *Normal Copy* - the storage system presents the actual contents of the secondary volume. This is in contrast to TI snapshots, where a merging process is required between the primary and secondary volumes to reconstruct the data.
- *Quick Split* and *Quick Resync* - the storage system presents the actual contents of the secondary volume. However a merging process may be required between the primary and secondary volumes to reconstruct the accessed block of data, if the background copy of that block has not yet been performed.

The following table shows how *Quick Split* and *Quick Resync* (indicated by the suffix q) are affected by upstream and downstream operations in an SI data flow:

Data Flow	Behaviour
SIq	The SI is performed using quick operations. The SI secondary is immediately available for manually mounting.
SIq with auto-mount of secondary	The SI is performed using quick operations. The SI secondary is auto-mounted immediately.
SIq with downstream replications/snapshots	SI is performed using quick operations. However: ▪ The downstream replications/snapshots will wait until the SI secondary has been completely copied. ▪ The SI secondary will only be available for manually mounting once it is completely copied. See note below.
SIq with auto-mount and downstream replications/snapshots	SI is performed using quick operations. However: ▪ The downstream replication/snapshot will wait until the SI secondary has been completely copied. ▪ The SI secondary will only be auto-mounted once it is completely copied. See note below.
Upstream replications with downstream SIq	SI is performed using quick operations. The SI secondary is immediately available for manually mounting. There is no impact on upstream replications.

Note: In cases where the SI replication must be fully evaluated, *Quick Resync* and *Quick Split* will take as long as *Normal Copy* and *Steady Split*. However the use of *Quick Resync* will have a beneficial affect on when and for how long the production application is quiesced.

When the clone is deleted the storage system releases the bitmap.

Protector supports *Steady Split*/*Normal Copy* and *Quick Split*/*Quick Resync*.



Figure 18 Full clone using continuous ShadowImage with protected, isolated, multiple TI snapshots

Continuous ShadowImage can be used to:

- Protect access to local TI snapshots if the production volumes fail.
- Isolate production volumes from performance impacts caused by heavy I/O on local TI snapshots.
- Allow multiple scheduled mount operations (beyond the limits imposed by SI mirror counts) without affecting the original backup, through the use of Refreshed TI.

Note:

Continuous SI can be combined with all hardware operations (i.e. TI, RTI, SI, TC, UR or GAD), with the exception that a continuous SI S-VOL cannot also be the P-VOL of a remote replication (i.e. TC, Universal Replicator or GAD).

i.e. It is not possible to chain a remote replication from a continuous SI target.

The typical use cases for continuous ShadowImage include:

- Repurpose on Demand - using continuous SI, keeps a close copy of the primary volume and allows pause and mount for repurposing.
- Protected Backup - using continuous SI to TI snapshots, retains snapshots in the event that the primary volume fails.
- DRU Protected Backup - using continuous SI to TI snapshots with DRU, retains snapshots with DRU lock in the event that the primary volume fails.
- Repurposing (TI) - using continuous SI to RTI snapshots, provides multiple repurposing copies, possibly in excess of the SI limit.
- Repurposing (SI) - using continuous SI to batch SI, provides a repurposing copy.
- Repurposing (SI) with Backup - using continuous SI to batch SI to TI snapshots, provides a repurposing copy with snapshots for protection.
- Repurposing (SI) with DRU Backup - using continuous SI to batch SI to TI snapshots with DRU, provides a repurposing copy with snapshots for protection with DRU lock.

About TrueCopy replication

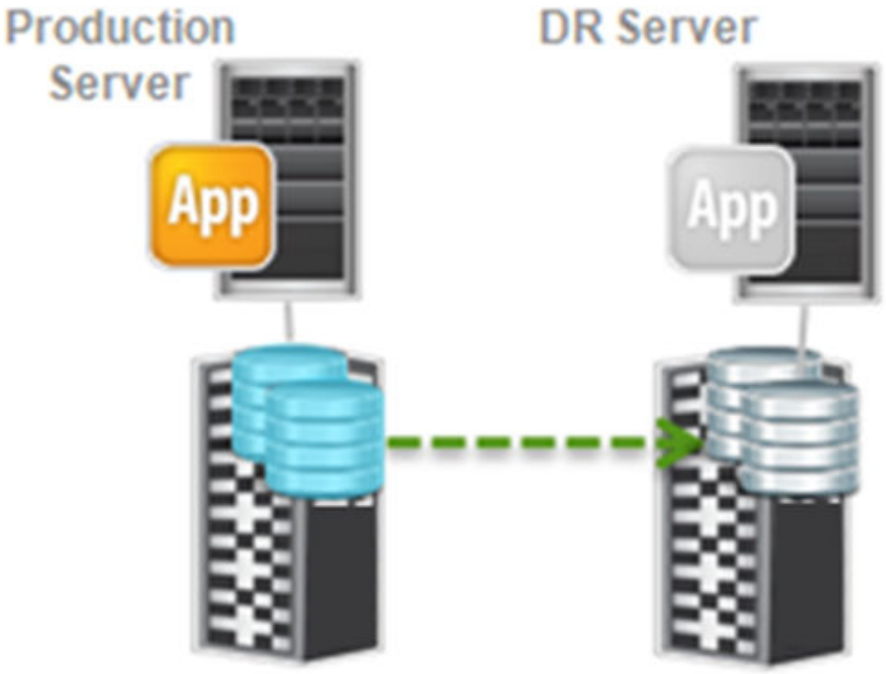


Figure 19 TrueCopy Replication

TrueCopy provides remote, volume consistent, synchronous replication.

When establishing a replication between primary and secondary volumes, the storage system copies all the data from the primary to the secondary volume. Depending on the volume size, the creation of replicas takes time. As with ShadowImage, data copy can be optimized by using DP volumes.

After creation of the replicas, the storage system maintains the replica on the secondary volume by synchronously transferring each write made to the primary. In synchronous (copy on write) mode, the storage system signals write I/O completion only when it has been transferred to the secondary volume. The write order is completely guaranteed so that the secondary volume is crash-consistent at any point in time.

In the *COPY* state, no read/write operation is permitted to the secondary volume. To access the secondary volume, the replication must be paused (i.e. the pair is placed in the suspended state). As with ShadowImage, a bitmap is maintained for later pair synchronization. When the replication is deleted the storage system releases the bitmap.

Note:

- Fence Level determines behaviour when an update to the secondary volume fails. This option should be set based on the business priority (i.e. keeping replications consistent versus keeping production data available):
  - Data – prevents writes to the primary volume if updates to the secondary volume fail. This setting is appropriate for volumes that are critical to recovery.
  - Status – prevents writes to the primary volume if the secondary volume's status cannot be set to 'suspended' in the event of a failure. This setting enables rapid resynchronisation once a failure is resolved.
  - Never – allows continued writes to the primary volume even if updates to the secondary volume fails. This setting is appropriate for volumes that must remain available.
- To maintain synchronised data transfer, sufficient bandwidth must be provided for the remote link, otherwise performance problems may be encountered on the production volume. When a replication is required over remote links with poor bandwidth or over long distances, use Universal Replicator.

About Universal Replicator

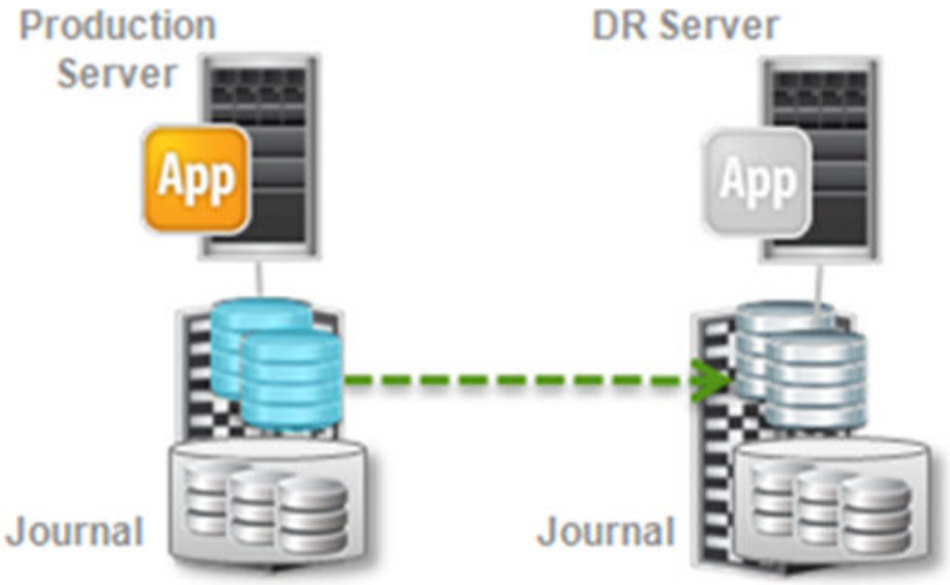


Figure 20 Universal Replicator

Universal Replicator performs volume consistent, asynchronous replication.

When establishing a replication to a target secondary volume, the storage system copies all the data to the secondary volume. Depending on the volume size, the creation of the replica takes time. As with TrueCopy, the copying can be optimized using DP volumes.

After the replication is created, the storage system maintains the replica on the secondary volume, transferring each write to the secondary volume. In asynchronous mode, the storage system signals each write completion as soon as it is performed on the primary volume, it then transfers it to the secondary volume (copy after write). Journalling ensures the write order is completely guaranteed so that the secondary volume is crash-recoverable at any point in time.

In the replicating state, no read/write operations are permitted to the secondary volume. To access to the secondary volume, the replication must be paused (placed in the suspended state). Universal Replicator maintains a journal of the primary volume for later pair synchronization. In contrast to TrueCopy, long duration pausing can be tolerated by configuring sufficiently large UR journals. When the replication is deleted the storage system releases the UR journals.

The use of Consistency Groups (CTGs) is mandatory for Universal Replicator.

Note: Long duration pausing, with sufficiently large UR journals, may cause service level violation with increased RPO. To avoid this, RPOs must be monitored to ensure that they are satisfying the SLA.

About Global-Active Device replication

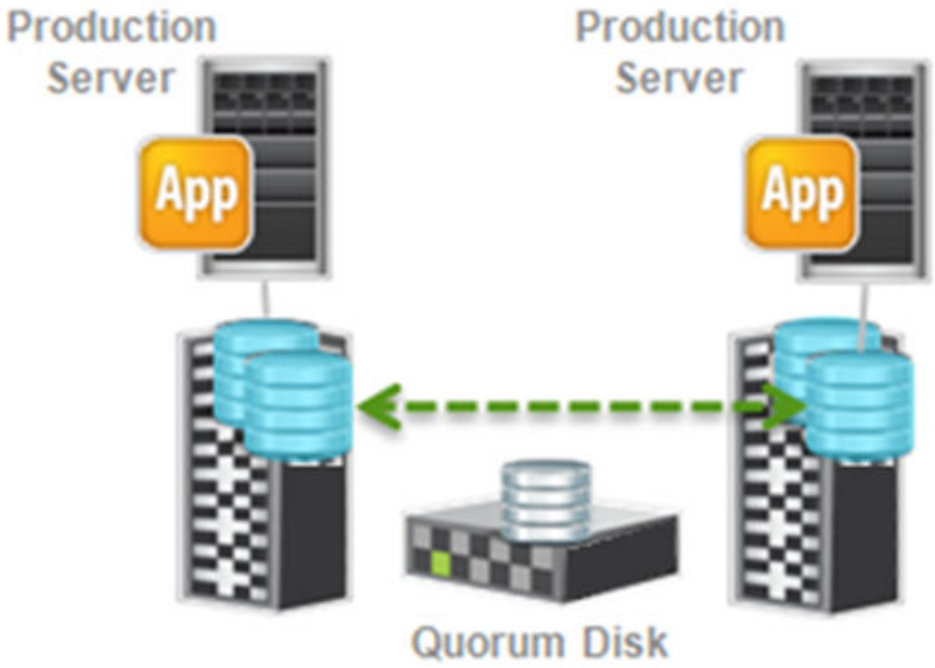


Figure 21 Global-Active Device Replication

Global-Active Device allows volume consistent, remote, active-active replication.

When establishing a replication, the storage system copies all the data to the secondary volume. Depending on the volume size, the creation of the replica takes time. As with TrueCopy, you can optimize the process using DP volumes.

After creating the replication, the storage system maintains the replica on the secondary volume. The copy mechanism and data consistency is the same as for TrueCopy.

Unlike TrueCopy, read/write operations are permitted on both the primary and secondary volume even in the replicating state, hence both sides of the replication pair are said to be active. All updates to the secondary volume are also transferred back to the primary volume. When the replication is paused (i.e. placed in the suspended state), the storage system determines the owner volume using the quorum disk and prohibits any read/write access to the non-owner volume.

As with TrueCopy, a bitmap is maintained for pair re-synchronization. When the replication is deleted, the storage system releases the bitmap.

Note: To fully handle a failure scenario, it is recommended that the quorum disk is located at a secure third site.

Protector is not be able to operate any dataflow containing a GAD replication if the primary and secondary have been swapped externally, since path resolution will see the original secondaries as the new primaries. The swap must be reversed externally to its normal direction.

About Global-Active Device Cross Path

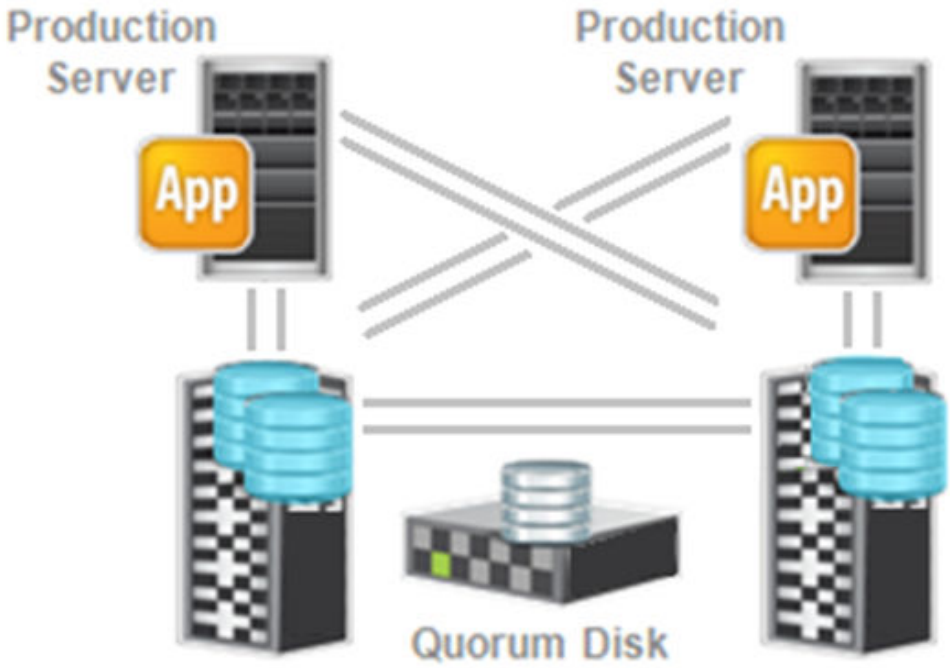


Figure 22 Fully Redundant GAD Cross-Path and Multi-Path Scenario

In a GAD cross-path and multi-path environment, the application servers may have one or more LUN paths to the PVOL, and also one or more LUN paths to the S-VOL. The GAD pair may also have one or more LUN paths between them.

Protector is capable of configuring and adopting hardware path, file system path and application replications for GAD cross-path and multi-path scenarios.

In situations where the *Application Host* or *OS Host* has a LUN path to a GAD P-VOL and also has a LUN path to the S-VOL of that replication, Protector will resolve application paths including that P-VOL when the host has one or more LUN paths to both volumes involved in the replication.

Protector can:

- Adopt the replication
- Re-evaluate the replication if it already exists in Protector
- Create a snapshot of the replication's P-VOL
- Swap the replication

This feature requires **`raidscan`** (version 01-41-03/03 or later), which, when issued against the secondary volumes of a remote replication, includes the serial number of the array hosting the primary volumes.

This feature only considers cross paths to the primary application. Failover support for secondary applications is not supported.

About three datacentre cascade (3DC)



Figure 23 Three datacentre cascade

Three datacentre (3DC) cascade using TrueCopy and Universal Replicator provides the maximum level of data protection by combining synchronous replication between the primary and local secondary site, cascaded with asynchronous replication between the local secondary and remote tertiary site.

If a component failure or power failure occurs at the primary site, production can be handed over to the local site with no data loss.

If a site failure or localised natural disaster occurs, both the primary and local site may be lost. In this case production can be handed over to the remote site with minimal data loss.

Consider the following when using 3DC cascade:

- Due to the cascading topology, a failure at the local site will prevent both the synchronous asynchronous replications, leaving the primary site unprotected. To avoid this situation, it is recommended to use the 3DC multi-target configuration instead.

About three datacentre multi-target

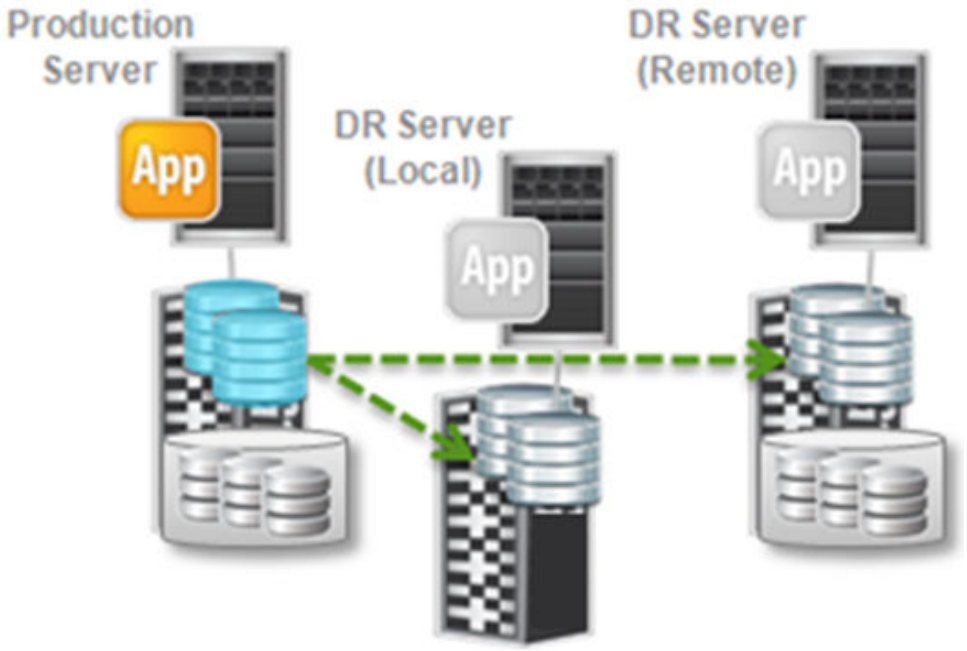


Figure 24 Three datacentre multi-target

Three datacentre (3DC) multi-target using Global-Active Device, TrueCopy and Universal Replicator provides the same level of data protection, equivalent to 3DC cascade, but improves upon it by solving the issue of local secondary site failure leaving the primary site unprotected. If the local secondary site fails, replication between the primary and remote secondary site can continue uninterrupted.

3DC multi-target may also be configured using Universal Replicator instead of TrueCopy. Symmetric configuration and operation simplifies component failure and site failure since both are handled with a single set of operations. Consider this option if simplicity of operation is the first priority.

About three datacentre multi-target with delta



Figure 25 Three datacentre multi-target with delta

Three datacentre (3DC) multi-target with delta is an improvement on 3DC multi-target that provides on-going protection even in the event of a failure at the primary site.

A replication from the production site to the local secondary site is configured using Universal Replicator. A replication from the production site to the remote secondary site is configured using Global-Active Device or TrueCopy. Additionally, a suspended, asynchronous Universal Replicator (delta-UR) replication is established between the local and remote secondary sites. The local and remote secondary sites will be near identical once pairing with the primary site is complete. Differences that appear between the secondary sites over time, due to a number of factors, are tracked by the suspended delta-UR replication.

Failure of the local or remote secondary site is handled in the same way as for 3DC multi-target, in that the primary site remains protected by the surviving secondary site.

In the event that the production site fails, the local secondary site can take over. The local secondary site is then rapidly brought under the protection of the remote secondary site by resuming the suspended delta-UR replication. Because only the deltas between the local and remote secondary sites need to be resynchronized, this pairing takes only a short time to achieve, meaning that the local site only remains unprotected for a brief period. Without the pre-existing, suspended delta UR between the local and remote secondary sites, it could take hours or even days to establish a replication between the secondary sites, leaving the local secondary site vulnerable for an extended period while this takes place.

About static clones of a clone



Figure 26 Static clones of a clone

Static clones of a clone enables multiple repurposing with heavy workloads (read/write I/O).

A 1st level clone (golden image) is created based on a schedule. The golden image can then be manually replicated as 2nd level clones for repurposing. The 2nd level of clones need to be static because the re-purpose servers require static access points (Port, Host Group and LUN although the LUN may not remain constant, depending on the mount host OS) for mounting the volumes.

If re-purposing beyond the ShadowImage limitation of two 2nd level clones is required, use Thin Image on the 2nd level.

About clone with replication

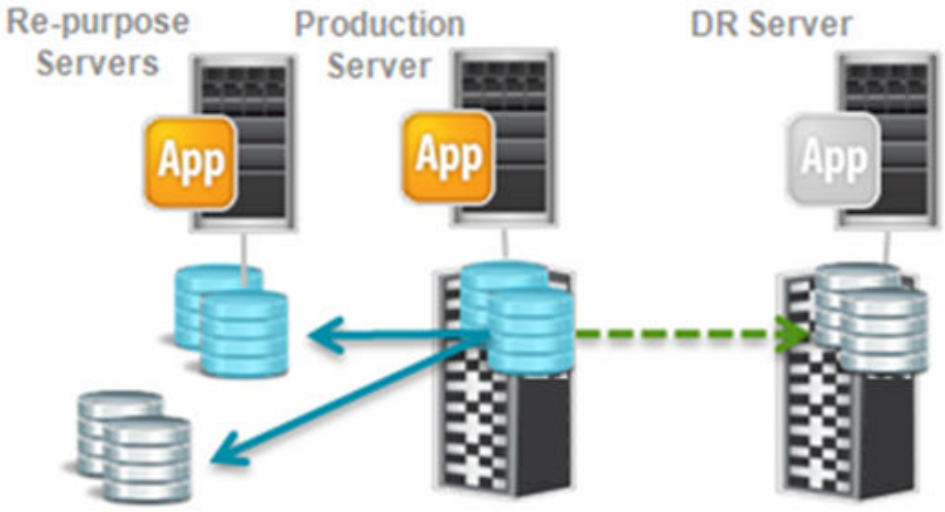


Figure 27 Clone with replication

Clone with replication enables disaster recovery and local backup and/or local re-purposing.

A replica of the production data is maintained on the remote site using TrueCopy, Universal Replicator or Global-Active Device. The remote replication is used for disaster recovery.

Local clones are created based on a schedule using ShadowImage. These clones can be used for fast operational recovery or repurposing.

About snapshot with replication

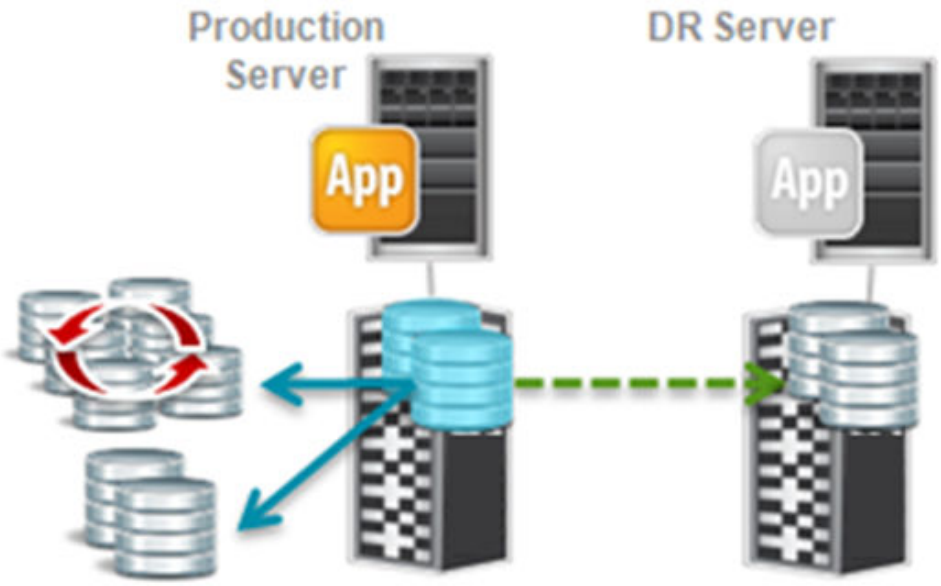


Figure 28 Snapshot with replication

Snapshot with replication enables disaster recovery and local backup.

A replica of the production data is maintained on the remote site using TrueCopy, Universal Replicator or Global-Active Device. The remote replication is used for of disaster recovery.

Local snapshots are created based on a schedule using Thin Image. These clones can be used for fast operational recovery.

About snapshot of a clone

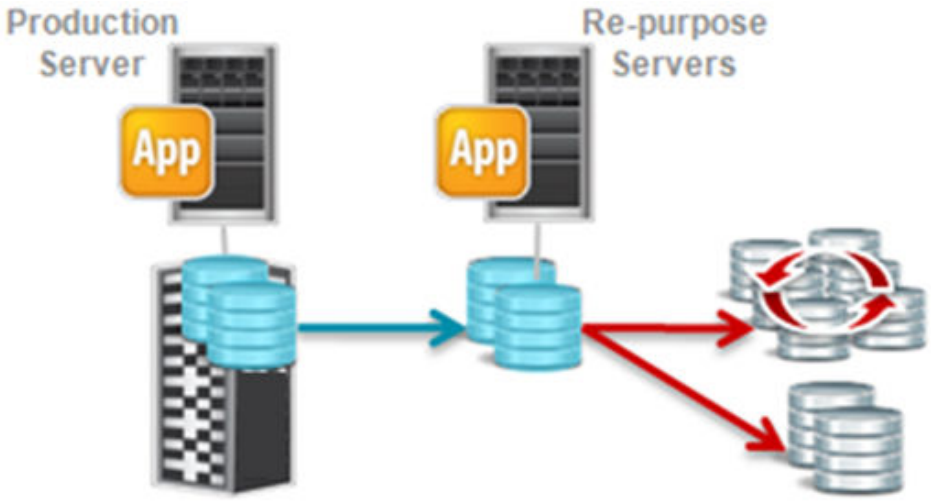


Figure 29 Snapshot of a clone

Snapshot of a clone enables ad-hoc backup for repurposing.

A clone is created using batch or continuous ShadowImage and is made available for repurposing.

A snapshot is taken so that the clone can be reverted quickly if required. Reversion may be needed when performing recurrent testing or configuration change/patching (that may fail) on the repurposed clone.

About replication of a clone

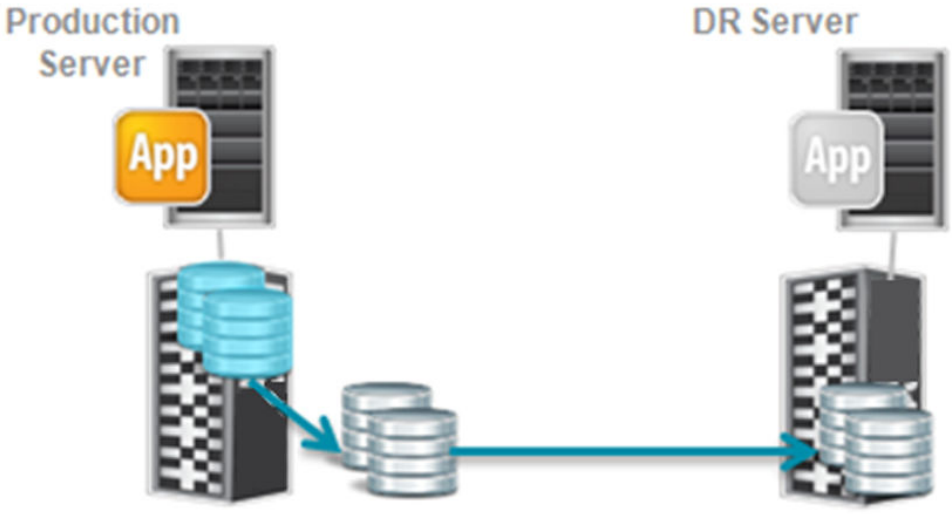


Figure 30 Replication of a clone

Replication of a clone enables remote backup without DR.

An in-system clone is created using ShadowImage based on a schedule. The cloned point-in-time image is then replicated using TrueCopy or Universal Replicator. Having an intermediate clone means that the replication process does not have any impact to the production volume. The replication is performed as a batch copy, so limited RPO is achievable (typically a few hours). For this reason, this technique is not common for high end storage systems.

About remote snapshot

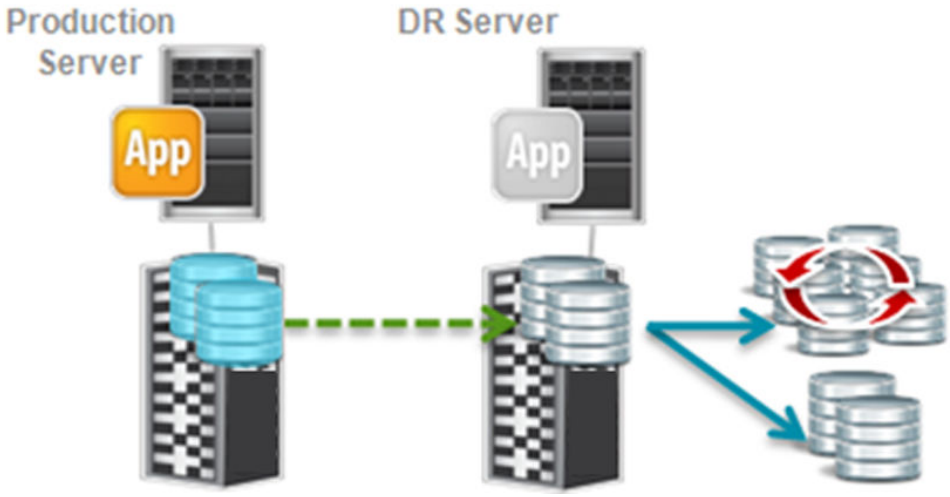


Figure 31 Remote snapshot

Remote snapshot enables disaster recovery with remote backup.

A replication is created using TrueCopy, Universal Replicator or Global-Active Device to maintain a replica image of production data on the remote site.

A snapshot is created from the remote replica based on a schedule.

This achieves a high level of disaster recovery, while remote site snapshots enable quick operational recovery even during disaster recovery.

About local and remote snapshots

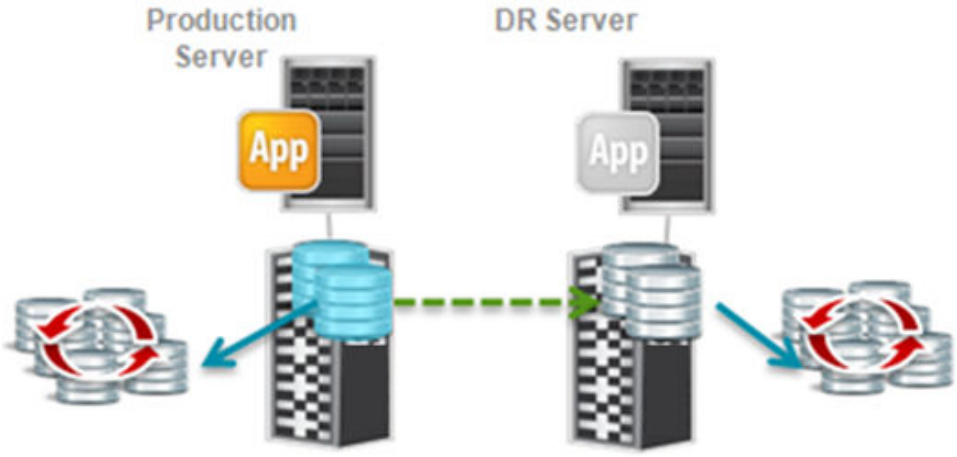


Figure 32 Local and remote snapshots

Local and remote snapshots enable disaster recovery with backups.

A replication is created using TrueCopy, Universal Replicator or Global-Active Device to keep a replica image of production data on the remote site.

Snapshots of the local production data and remote replica are created based on the same schedule.

This keeps the data consistent on the both sites, simplifying the process of operational recovery during site failover, contributing to a better RTO.

About remote clone

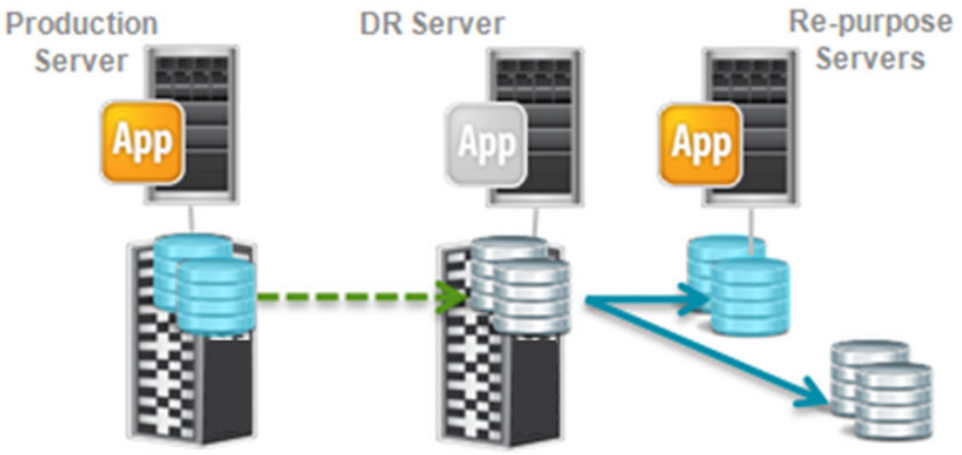


Figure 33 Remote clone

Remote clone enables disaster recovery with remote backup and/or remote repurposing.

A replication is created using TrueCopy, Universal Replicator or Global-Active Device to maintain the replica of production data on the remote site.

Clones of the remote replica are created using ShadowImage based on a schedule.

Replication achieves the highest level of disaster recovery, while remote snapshots/clones enable quick operational recovery even during disaster recovery. Also, repurposing with remote snapshots/clones avoids any performance impact to the production site.

About local & remote clones

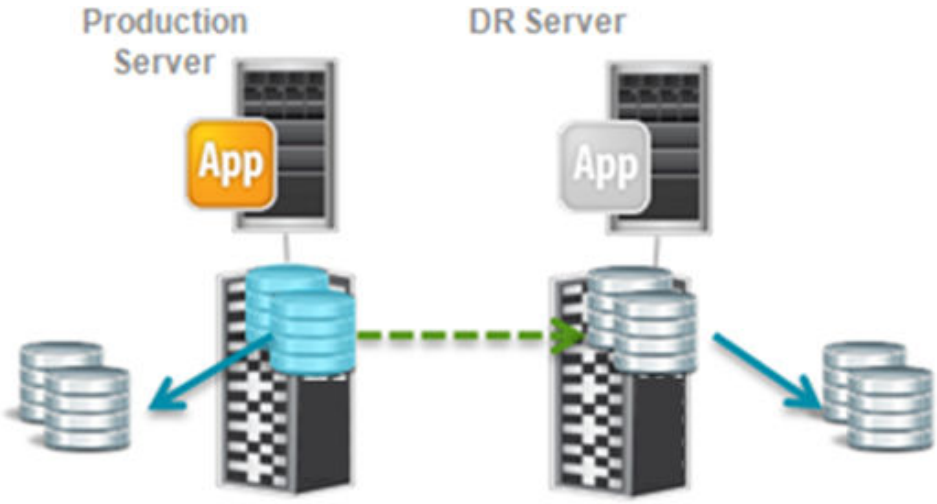


Figure 34 Local & remote clones

Local and remote clones offer the same benefits as local and remote snapshots, with the additional benefit of protecting the data from physical failure of the production or replica volumes.

Note the local and remote clones do not allow the user to recover to a point-in-time. The clone is completely refreshed so that it is equal to the primary or secondary volume. Operational recovery is therefore more limited as a result.

About local snapshot and remote clones

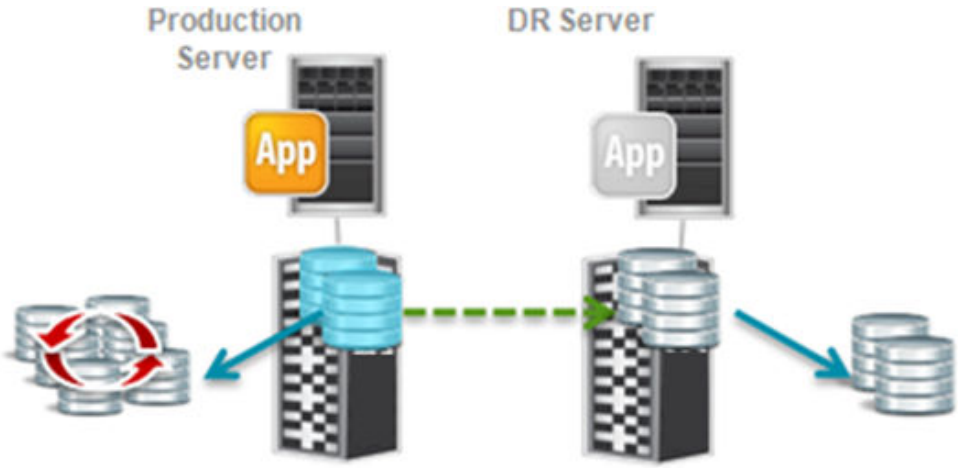


Figure 35 Local snapshot and remote clones

Local snapshot and remote clones enables disaster recovery with tiered backup.

A replication is created using TrueCopy, Universal Replicator or Global-Active Device to keep a replica image of production data on the remote site.

Local snapshots of the production data and remote clones of the replica are created using Thin Image and ShadowImage respectively, based on different schedules.

Keeping different data on each site enables quick recovery on the production site, while satisfying the long-term protection on the remote site.

This scenario can be flipped (i.e. local clone and remote snapshot) so that the clone is maintained at the local site and the snapshots at the remote site. Protector also supports snapshots and clones on the local and remote site concurrently.

About Hitachi Block replication adoption

Ops Center Protector can adopt and manage ShadowImage, TrueCopy, Universal Replicator and Global-Active Device replications that already exist on Hitachi Block storage hardware. An adopted replication can then be managed via the GUI.

To adopt a replication, the user must specify a replication policy that identifies the required source LDEV(s) or Host Group, draw an appropriate data flow that identifies the source and destination storage devices, replication type (and mirror number for SI and UR), mark the policy for adoption and then activate the rules.

Note: Any classification can be used to specify the source LDEV(s), including *Application* and Filesystem *Path*. Protector will attempt to resolve and adopt them. This is subject to existing limitations.

Adopted replications can be augmented with user defined replication and snapshot operations.

A replication can be dissociated from Ops Center Protector without being removed from the storage hardware.

Caution: Be aware of the difference in semantics between dissociating and removing replications from data flows:

- Dissociating a replication will leave that replication intact on the hardware.
- Removing a replication from a data flow and redistributing the rules will cause that replication to be torn down on the hardware.

Note: The following apply when adopting replications:

- Replications can be adopted from any supported block storage systems.
- In-system, 2DC and 3DC replications are supported.
- All valid replication data flows including cascades and multi target are supported.
- The user must understand and create the data flow prior to adopting.
- The user needs to know the type of replication that is to be adopted in addition to the Mirror Unit Number. The remaining properties will be discovered from storage.
- At least one existing replication pair must exist on the selected mirror.
- The replication being adopted must be in the in the same direction as the one defined in the data flow, i.e. it cannot be in the reversed flow state.
- Refreshed Thin Image replications cannot be adopted.
- The source and destination journals specified must match those of the Universal Replicator replication being adopted.
- Primary volumes replicating on different mirror unit numbers are not supported.
- Adopting by copy groups or device groups is not supported.
- There is no check for attempting to adopt the same hardware pairs on multiple, active, coexisting data flows.
- Limitations for other features still apply if they are relevant to the adopted replication.

When adopting replications Ops Center Protector will behave in the following ways depending on the policy and data flow attributes supplied by the user when attempting to perform the adoption process:

Replication Policy and Data flow Configuration	Ops Center Protector's Behaviour
Any	**1.** Primary volumes that are not being replicated on the specified mirror will have a secondary volume provisioned and that pair is added to the replication set while respecting the replication type, journal, CTG, and fence level options.

Replication Policy and Data flow Configuration	Ops Center Protector's Behaviour
	2. Adopted replications are flagged as such in the hardware resource information along with their CTG ID, journals, mirror unit numbers, and fence levels.
Any	▪ If the replication is found to be in PAIR and the Protector mover type is *Batch* then the replication will be suspended. ▪ If the replication is found to be in PSUS/SSUS and the Protector mover type is *Continuous* then the replication will be resumed.
The user does not select a pool, but does select a mirror unit number	**1.** If the mirror unit is assigned, adopt the replication pairs on the mirror. **2.** If the mirror unit is assigned but there are one or more P-VOLs not replicating on that mirror, create S-VOLs for those P-VOLs, in the pool used by the existing S-VOLs (or error if the S-VOLs exist in more than one pool). **3.** If the mirror unit is unassigned, log the error `"Cannot provision, no pool selected"`.
The user selects a mirror number that is not supported for the selected replication type. Valid mirrors numbers are: ▪ ShadowImage: 0, 1 or 2 ▪ TrueCopy: 0 only ▪ Universal Replicator: 0, h1, h2 or h3 ▪ Global-Active Device: 0, 1, h1 or h2	Log the error `"Could not determine pool/journal/quorum to use. Ensure that there is at least one existing pair of matching replication type"` **Note:** When Protector attempts to adopt a TC or GAD replication, it will detect UR pairs on mirror number 0 and, if present, log the error: `Handler 'HitachiVirtualStoragePlatform' call failed: [TrueCopy|GAD] mirror for one or more adopted pairs already in use by UR.` See above for more limitations relating to mirror unit and replication type combinations.
The user changes the mirror number after initial data flow activation	The user will be warned via the GUI that the following actions will be taken before they reactivate the rules: **1.** Volumes and relationships that have been adopted will be unadopted **2.** Volumes and relationships that have been created by the user will be destroyed

Replication Policy and Data flow Configuration	Ops Center Protector's Behaviour
	3. Replications will be re-adopted and created based on the new mirror number

Policy Concepts

This section describes Ops Center Protector's policy management features.

For further information, refer to:

- Policy Tasks (on page 272)
- Policies UI Reference (on page 670)

About policies

A *Policy* consists of *Classifications* that specify what data is to be protected and *Operations* that specify how that data is to be protected.

Physical classifications specify the data by directly naming the path, logical device/volume or disk type to be protected, whereas *Application classifications* specify the data indirectly by naming the application instance such as databases or virtual machines to be protected, Ops Center Protector then discovers the volumes that the application's data resides on. The following table shows which classifications can be used in conjunction with a particular type of source node.

Table 2 Classifications vs. Source Node Types

Source Node Type	Policy Classifications				
	Oracle Database	Oracle RMAN	Disk Type	Block	Path
OS Host			Yes		Yes
Oracle	Yes	Yes			
VMware		Yes			
Hitachi Block Device				Yes	
Hitachi Logical Block Device					
Hitachi Block Host				Yes	
Generation 1 Repository	A Generation 1 Repository can act as a source node in a cascaded data flow, where one generation 1 repository is backed up to another generation 1 repository. The policy classification is effectively the same as the one applied to the original data source in the cascade.				
Generation 2 storage nodes (HCP, Amazon S3 and Repository)	Can act as a source node in a cascaded data flow, where one generation 2 node is backed up to another generation 2 node. The policy classification is effectively the same as the one applied to the original data source in the cascade				

The *Operations* in a policy define the methods to be used to create backups of the primary data. Operations can be implemented using Protector software methods (e.g. Backup), or by orchestrating the hardware storage devices (e.g. Hitachi Block) to implement operations in hardware (e.g. Thin Image and Universal Replicator). The following table shows which operations can be used in conjunction with a particular type of source or destination node.

For example: from the table below we can see that an *OS Host* node can act as the *Source (S)* for a *H/W* based *Replicate* operation, and that a *Block Device* node can act as the *Destination (D)*.

Table 3 Node Types vs. Operation Types. (Key: S=Source node for the operation, D=Destination node for the operation)

Node Type	Policy Operations						
	Backup	Access	Replicate		Snapshot		Mount
	S/W	S/W	S/W(1)	H/W	S/W(2)	H/W(3)	H/W
OS Host	S			S		S	D
Oracle		S		S		S	D
VMware	S			S		S	
Hitachi Block Device				S or D		S	
Hitachi Logical Block Device				S or D		S	
Hitachi Block Host				S		S	
Repository	S or D						
HCP	S or D						
Amazon S3	S or D						
(1)	(2) Software Replication is not currently supported. (3) Software Snapshot is not currently supported. (4) Snapshots are local operations, so the source is always the same as the destination on a data flow.						

Operations defined in a *Policy* work in conjunction with *Movers* placed on a *Data Flow*. When a policy is assigned to a source node, all the operations in that policy are assigned to that node. The policy is then routed by a mover and the contained operations assigned to one or more destination nodes. The following table shows which operations can be used in conjunction with a particular type of mover:

Table 4 Operation Types vs. Mover Types

Policy Operation	Implementation	Batch Mover	Continuous Mover
Backup	Software	Scheduled or on-demand incremental backups are created on a repository HCP or Amazon S3 node, based on RPO.	N/A
Replicate	Software	Not supported.	N/A.
	Hardware	Hitachi Block: Scheduled or on-demand replicas of the P-VOL are created on an S-VOL using Refreshed Thin Image, ShadowImage, TrueCopy or Universal Replicator.	Hitachi Block: Live changes to the P-VOL are replicated to an S-VOL using ShadowImage, Universal Replicator or Global-Active Device.
Snapshot	Software	Windows: Not currently supported.	N/A
		Linux/AIX: Not currently supported.	
	Hardware	Hitachi Block: Scheduled or on-demand local snapshots are created using Thin Image. The batch mover is implied but not shown on the data flow since the operation is local to the storage device.	N/A
Mount	Hardware	Automatically mounts a Hitachi Block LDEV based on a schedule. The mounted LDEV can then be used for repurposing or proxy backup.	N/A

About policy classifications

A Classification specifies what data is to be protected either by directly identifying files or volumes, or indirectly by specifying application level objects such as databases or datastores.

Protector resolves a classification down to a unit of data appropriate to the underlying storage device type.

- For a host based backup, the classification is resolved down to disk data blocks where the data being protected resides and it is changed disk data blocks that are recorded in the repository.
- For block storage based backup, the classification is resolved down to LDEVs and it is entire LDEVs that are replicated or snapshotted.

A classification is resolved when a data flow that implements a policy is first activated, when it is re-activated while currently active and possibly when re-triggered. It is important to understand the behaviour when classification resolution results in changes to the set of underlying data units being backed up.

For a classification that is applied to a host based backup operation, when resolution alters the files included in the backup then:

- Any files that remain in the backup are unaffected and continue to be backed up as before.
- Any files that are newly included are added to the backup.
- Any files that are no longer included remain in the backup, but changes to them are no longer recorded in the repository.

For a classification that is applied to a block storage device backup operation (snapshot or replication), when resolution alters the LDEVs included in the backup set then:

- Any LDEVs that remain in the backup set are unaffected and continue to be backed up as before.
- Any LDEVs that are newly included are added to the backup set, pairings are created with the respective S-VOLs.
- Any LDEVs that are no longer included are removed from the backup set, their pairings are torn down and the respective S-VOLs are deleted.



Note: This behaviour differs from the case when deactivating a data flow, where the entire backup set is torn down but the S-VOLs are left in place as static copies.



Note: For Global-Active Device, when a Thin Image snapshot, Refreshed Thin Image, ShadowImage or Universal Replicator replication is taken from the GAD S-VOL, the GAD S-VOL cannot be deleted.

About VMware policy classifications

When items are added to the inclusion or exclusion lists displayed in the VMware Classification Wizard (on page 700), the VMware Resource Selection Wizard (on page 701) is launched. This wizard enables virtual machines and templates to be selected based on the VMware inventories in which they appear in vSphere, or by pattern matching of VMware container object name, virtual machine name or template name. The list of VMs and templates included in the classification is evaluated at different times depending on how they are specified:

- Evaluation is done only once (i.e. when the data flow implementing the policy is compiled), if VMs and templates are:
  - Explicitly selected from a list or inventory tree.
  - Specified using their full name (i.e. without using wildcards, e.g. `Sales_SQLServer`).
- Evaluation is done every time the operation is triggered, if VMs and templates are:
  - Implicitly selected using a container object (folder, host, cluster, datastore, resource pool, datacenter, or vApp).
  - Selected using a tag defined in vSphere.
  - Specified using a name pattern (i.e. using wildcards, e.g. `Sales_Client*`).



Tip: With this method of classification, VMs will be automatically added to the backup (without reactivating the data flow) when they are added to a container, assigned the appropriate tag or given a name that matches the defined pattern. For continuous replications it will be necessary to trigger the relevant operation to cause re-evaluation.

Every VMware object selected in the classification is resolved to a list of VMs and templates. For example, when selecting a datastore, all the VMs and templates that are in that datastore are selected. If any included VMs and templates reference VMDKs located in another datastore, these will be selected too. This ensures that VMs and templates that are backed up can be fully restored.



WARNING: When vCenter objects like VMs and Datastores change their names or IDs, this will cause the backup to fail, to avoid backing up incorrect objects. Updating the policy to replace these with the new instances will be required.

Backup behaviour differs depending on the type of operation the VMware classification is combined with in a policy.

For host based *Backup* operations, the VMware files that record each selected VM's state (system configuration, virtual hard disk configuration and virtual hard disk data) are backed up as dictated by the policy's operation(s). If a VM contains RDM storage then:

- Physical compatibility mode RDM disks are not backed up because they are not included in a VMware snapshot.
- Virtual compatibility mode RDM disks are backed up.

For block based *Snapshot* and *Replicate* operations, the datastores that contain the selected VMs are identified. Those datastores that reside on Hitachi Block storage are then resolved down to their underlying LDEVs and are snapshotted/replicated as dictated by the policy's operation(s).

- If the VM contains physical or virtual compatibility mode RDM storage, the backup operation will continue without backing up the RDM storage and the following warning will be logged:

  ```
  VM: <VM_NAME>. Contains a RawDiskMapping (RDM). This RDM storage won't be backed up.
  ```

- If the VM contains Passthrough storage, the backup operation will continue without backing up the passthrough storage and the following warning will be logged:

  ```
  VM: <VM_NAME>. Contains <TYPE> Passthrough storage. This Passthrough storage won't be backed up.
  ```

- If the VM has a dependency on a non-VMFS datastore (i.e. one that is not located on a block storage device), then:
  - If no VMDKs for the VM are present in the non-VMFS datastore, the backup operation will continue and the following warning will be logged:

    ```
    VM contains non-VMFS datastore '<DATASTORE_NAME>', which won't be backed up.
    ```

  - If VMDKs for the VM are present in the non-VMFS datastore, the backup operation will be aborted and the following error will be logged:

    ```
    The following non-VMFS datastores contain VM disks which won't be backed up: <LIST OF NON-VMFS DATASTORE NAMES WITH VMDKS>.
    ```

Tip: Any RDM storage that cannot be protected by a *VMware* classification can be backed up using a separate *Physical* classification if appropriate.

About Hyper-V policy classifications

The Hyper-V policy classification defines which virtual machines will be protected to which level of consistency as part of a policy.

When items are added to the inclusion or exclusion lists displayed in the Hyper-V Classification Wizard (on page 692), the Hyper-V Resource Selection Wizard (on page 694) is launched. This wizard enables virtual machines to be selected based via browsing, or by pattern matching of virtual machine name, configuration path or host running the VM.

The list of VMs included in the classification is evaluated at different times depending on how they are specified:

Evaluation is only performed once, when browsing and selecting the virtual machine directly.

Evaluation is done every time the operation is triggered if VMs are:

- Implicitly selected using a virtual machine location (path) or virtual machine host (hypervisor)
- Specified using a name pattern (i.e. using wildcards, e.g. Sales*)

Tip: With this method of classification, VMs will be automatically added to the backup (without reactivating the data flow) when they are added to a selected path or host or given a name that matches the defined pattern. For continuous replications it will be necessary to trigger the relevant operation to cause re-evaluation

Every Hyper-V object selected in the classification is resolved to a list of VMs. For example, when selecting a virtual machine location, all VMs with configurations in that location are selected. If any included VMs utilize additional paths, these will be added too. This ensures that backed up virtual machines can be fully restored.

Tip: Use the preview selection functionality of the Hyper-V Classification Wizard (on page 692) to preview which VMs will be included for a given Hyper-V Node based on the provided inclusion and exclusion lists.

Virtual machine consistency

The consistency of virtual machine backups differs based on the consistency level chosen in the policy. Before the block-based backup is performed Protector will create a checkpoint for every virtual machine.

Application consistent checkpoints will use Hyper-V integration services to quiesce the data inside the VM. In case of Windows VMs, Protector will utilize VSS within the virtual machine to save application and filesystem data to disk. For Linux, the Hyper-V integration will request a quiesce, however usually this will only cause the filesystem buffers to be saved to disk. If an application consistent VM checkpoint cannot be created, Protector will create a crash consistent checkpoint instead.

Crash consistent checkpoints will create a checkpoint of the virtual disks as they are.

About Microsoft SQL Server Policy Classifications

The Microsoft SQL Server policy classification defines which databases will be protected and with what impact to the differential base.

Adding the classification starts the Microsoft SQL Server Classification Wizard (on page 681). Once the source Microsoft SQL node has been selected, databases can be added using the Microsoft SQL Server Database Selection Wizard (on page 683).

The Backup mode selection defines if the snapshot or clone will be recognized as a full backup by the SQL Server or if it should just be a copy which does not affect the differential base. In both cases, Protector will protect the complete database as well as the current log files and will represent the database at the point in time of the snapshot or clone creation.

It is not possible to backup only the transaction logs. Use the “BACKUP LOG” SQL command or a 3rd party tool to protect your transaction logs. Protector’s file system backups can help to store the files safely in the datacenter or cloud.

Tip: Keep in mind Microsoft SQL Server requires that at least a single full backup before transaction log backups can be performed.

About path macros

Path macros are used to define paths when constructing a policy.

There are occasions where it is desirable to define paths logically when constructing a policy, knowing that only the source machines that will receive the policy can actually determine where those paths should really be on the local system. For example, all the fixed drives or the Exchange directory, as these can differ from machine to machine.

Ops Center Protector facilitates this requirement by allowing the use of macros when defining paths. These macros are placed in a path name using the syntax `$[MACRO_NAME]`, and are resolved during initial data flow activation and on-going rules updates on each source machine. They are resolved by looking up the macro name in a `macros.cfg` file stored on each source machine in `C:\Program Files\Hitachi\Protector\db\config\`.

Entries within `macros.cfg` are given in the form:

keyword=path

For example:

```
files=c:\files
```

To use this macro, `$[files]` would be entered in either the Include or Exclude list for a *Path* classification. Note that macros can be combined with normal text, so that `$[files]\subdirectory` would be a valid entry. If the macro is defined with a trailing slash, ensure it is not combined with another slash when combining it with plain text.

Caution: Macro substitution is case sensitive. Ensure the case of macros in each `macros.cfg` file matches the case of macros when used to define a policy.

Multiple paths can be assigned to a single macro by delimiting each path with a semi-colon (;), as so:

keyword=path_1;path_2;path_3

About application quiescing

Ordinarily, when a backup is performed, the state of the backup consists of the exact state of the file system at the point in time at which it was triggered. This can be problematic when the backup is eventually restored, since some applications that were running at the time cannot, for example, have flushed all their in-memory file buffers, and their data files on disk could therefore be in an inconsistent state.

To resolve this, it is possible to request certain applications to quiesce into a safe state prior to a backup being triggered, and then to restore their previous active state after the backup is triggered. This ensures their state during the backup is consistent and that they can run as expected if the backup is restored. The quiesce option applies to classifications containing:

- File system data - where those policies are implemented on Windows source nodes.
- Applications - where the application type is supported by Protector.
- VMware VMs - where those virtual machines are running Windows, are online and running VMware tools.

If an application cannot be quiesced then a snapshot of its files is said to be in a crash-consistent state, as if the application had crashed while running. Crash consistent backups can, depending on the application, be used to restore state to some point prior to the backup.

Ops Center Protector includes built-in classifications that enable specific application types to quiesce into a consistent state prior to a Ops Center Protector Backup being triggered. This is done by selecting the Quiesce configured applications before backup check box in the backup operation's attributes. When a backup is triggered in a data flow that makes use of this policy, the resulting backup will contain a quiesced version of the application’s files.

If other applications, not directly supported by Protector, need to put themselves into a certain state prior to a backup operation, pre and post scripts can be run during the backup cycle. This is done by selecting the Pre script and Post script check boxes in the backup operation's attributes, and then supplying the names of appropriate script files. This script file must be present on the source node that implements the policy and must be placed in the Protector scripts directory (`C:\Program Files\Hitachi\Protector\scripts` is the Windows default).

About synchronization groups

The purpose of synchronization groups is to ensure that, if required, multiple hardware snapshots or batch hardware replications of source data can be created on multiple nodes using the same schedule and yet remain consistent with each other and with the source data.

To this end any two operations in the same policy using the same schedule name will be placed in the same synchronization group. When a hardware snapshot/replication data flow is created, any nodes implementing operations using that synchronization group will be scheduled all as one.

Note:

The RPO of all hardware snapshot operations in the same synchronization group must be identical. The rules compiler checks the RPO settings and if inconsistent generates error 10356: `Policy policy contains multiple hardware snapshot operations with different RPO settings but all use schedule schedule`

In a simple case of a cascaded batch data flow (i.e. Node A --> Node B --> Node C), where Nodes B and C both implement a hardware replication operation within the policy, then if both those operations share the same schedule name, they will both be part of the same synchronization group and be scheduled as one. This ensures that the replica of Node B that is created on Node C is created only when the replica of Node A created on Node B has been completed and so the replicas on Nodes B and C remain identical.

When mixing live and batch data flows or implementing different synchronization groups (schedules) at different points of the data flow, then the basic rule for these is that any operation is traced back towards the source, and if an operation sharing the same synchronization group is implemented on an immediately preceding node, then they will be scheduled together by the first preceding node that does not use that synchronization group. Note that live replications are considered to be in all synchronization groups since they should always match the source.

Where synchronization groups are not in use, replicate operations are scheduled on the sending node and snapshot operations are scheduled on the destination node.

About the automated Mount operation

The automated *Mount* policy operation provides a mechanism for mounting a Hitachi Block based snapshot or replication to a host machine so that the mounted volume can be used for *repurposing* or *proxy backup* to another medium (e.g. a Protector repository). The mount operation can be scheduled to occur at the desired phase relative to other operations in the policy. The computing load of repurposing or backup processing is offloaded to the host where the volume is mounted. The sequencing of the mount operation is dependent on the type of mount selected when the policy is assigned on the data flow:

- Repurpose - the volume remains mounted but is periodically unmounted, re-evaluated, then remounted. See About the repurposing mount sequence (on page 100).
- Proxy Backup - the volume remains unmounted but is periodically mounted, the backup process run, then the volume is unmounted. See About the proxy backup mount sequence (on page 101).

About the repurposing mount sequence

Note: Repurposing is not valid for continuous replications.

In this scenario the replication destination remains mounted on the mount host and is available to the repurposing application. The replication is, depending on the schedule, periodically unmounted, re-evaluated then remounted. The actual sequence is as follows:

1. The Before Unmount script (specified in the Mount Operation Wizard (on page 721)) is invoked on the mount host. The script would typically stop any applications from accessing the repurposed volumes while they are being re-evaluated.
2. The replicated volumes are unmounted from mount host.
3. The replication is resumed, re-evaluated and then paused again.
4. The updated volumes are remounted to the mount host.
5. The After Mount script is invoked on the mount host. The script would typically restart any applications accessing the repurposed volumes.

When repurposed volumes are subsequently unmounted by the user, the replication operation is automatically resumed.

About the proxy backup mount sequence

Proxy backup is valid for live and batch backup. In this scenario the replication destination is mounted, work performed on it and then unmounted. The actual sequence is as follows:

1. The replication is re-evaluated.
2. If the replication is live then it is paused.
3. The updated replicated volumes are mounted on the mount host.
4. The After Mount script (specified in the Mount Operation Wizard (on page 721)) is invoked on the mount host. This script would typically cause the proxy backup to run, calling out to a third party backup product or invoke a Protector backup using a trigger schedule.
5. The Before Unmount script is invoked on the mount host. This script may be omitted.
6. The replicated volumes are unmounted from mount host.
7. If the replication is live then it is resumed.

Schedule Concepts

This section describes Ops Center Protector's schedule management features.

For further information, refer to:

- Schedule Tasks (on page 279)
- Schedules UI Reference (on page 849)

About Schedules

A *Schedule* defines when an operation should start or a period within which an operation is allowed or prevented from occurring.

The Schedule Wizard (on page 850) can be accessed from the Schedules Inventory (on page 849) or by clicking Manage Schedules in the Policy Wizard (on page 671) when configuring a Backup, Snapshot, Replicate or automated Mount operation.

Central to the operation of the Schedule Manager is the Recovery Point Objective (RPO). This setting defines how regularly a backup should be triggered. An RPO of 12 hours tells the Schedule Manager to schedule a backup at that interval, whilst also obeying any other user-defined constraints imposed on the schedule. For example, a backup can be allowed to occur only between 12am and 6am, and this restriction must be adhered to even if the RPO requires backups to occur every hour. See the Policy Wizard (on page 671) for details about RPOs.

Backup events are triggered by one of the following:

- A policy indicating that the RPO is reached.
- A schedule indicating that a Trigger Time is reached.
- A user manually injecting a backup event by clicking the Trigger operation button on the UI.
- A user (or script) manually injecting a backup event via the CLI.

Permitted backup times depend on the constraints used to build the schedule:

- RPO Backup Window – An RPO backup event can only occur during a backup window. By default, an implicit backup window is used that operates all day, every day. If a custom backup window element is created, then the default backup window is ignored.
- RPO Exclude Window – An RPO or scheduled backup event can never occur during an exclusion period.

Trigger Time, RPO Backup Window and RPO Exclude Window can be used in any combination within a single schedule, such that a backup will occur if one of the following is true:

- The backup has reached its RPO
- AND an RPO Backup Window is active
- AND an RPO Exclude Window is not active

OR:

- A Trigger Time schedule event occurs
- AND an RPO Exclude Window is not active

OR:

- A backup event is injected manually via the UI or CLI



Note: It is advised that scheduling features are used carefully. If you define too complex a mix of triggers, backup windows and exclude windows then it may be difficult to predict when backups will occur and thus achieve the desired RPO.

While backups occur either when a policy's Recovery Point Objective dictates, a scheduled Trigger Time or a manually injected backup event, the precise time at which they occur is subject to a number of factors:

- At the most basic level, the time at which the rules are activated determines when the initial and successive RPO based backups occur, unless affected by other factors.

  Caution: For host based and hardware storage based backups, the first backup occurs automatically after the rules distribution is completed, irrespective of RPO.

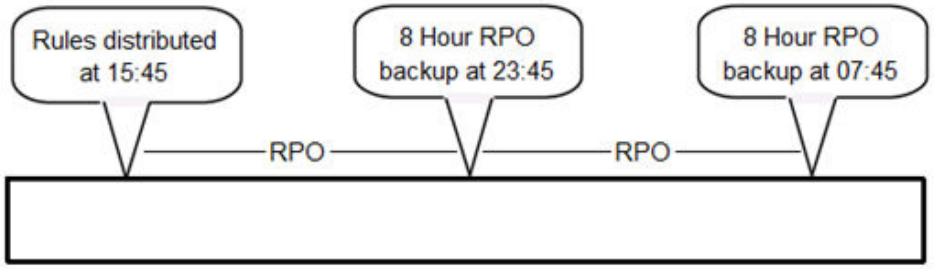


Figure 36 Data flow activation time (15:45) determines RPO backup times

- The RPO and/or Trigger Time defined in the policy determines when the next RPO backup will occur, relative to the previous RPO or Trigger Time backup.

  Note: A backup event injected manually via the UI or CLI does not affect the timing of RPO based backups.

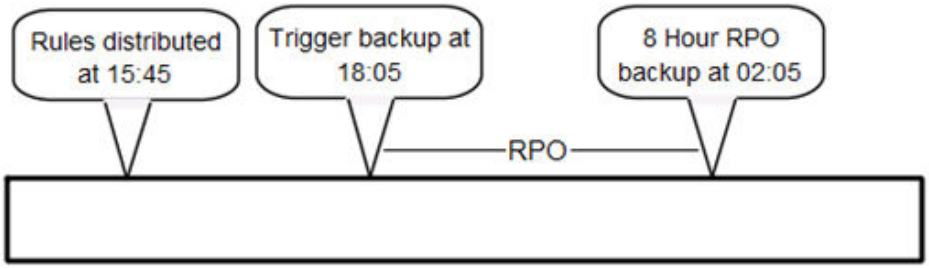


Figure 37 RPO backups occur relative to the time of the previous RPO or Trigger Time backup (18:05 + 8)

- If a backup is scheduled to occur Weekly or Monthly then:
  - If the Time is set to All Day then the backup will recur at the time of day that the rules were activated.

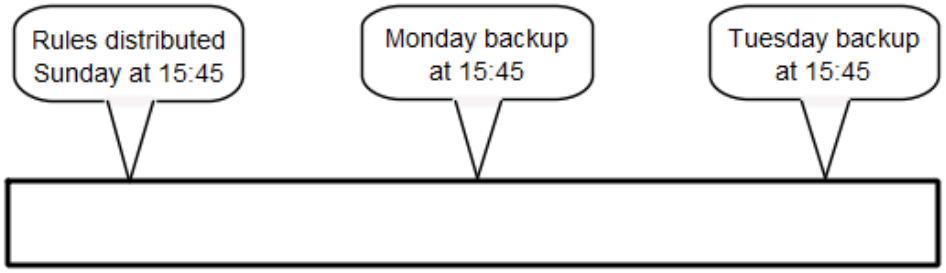


Figure 38 Data flow activation time (15:45) determines 'All Day' backup times

  - If the Time is set to Scheduled Time then the backup will recur at the Start time or, if this can't be achieved, at the earliest opportunity within the Duration specified after the Start time.

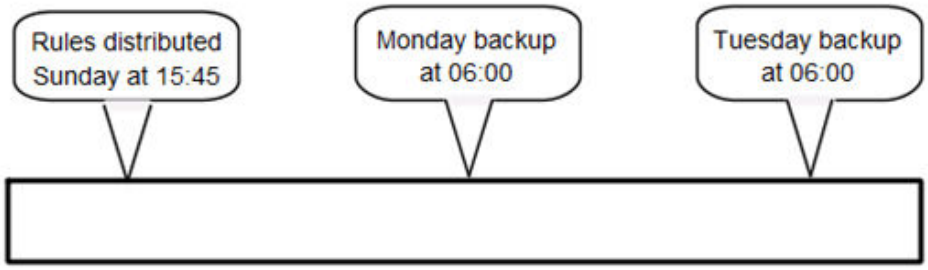


Figure 39 Start time (06:00) determines 'Trigger Time' backup times

- If a backup event is injected manually via the UI or CLI then a backup will occur at the time the event is received, but will not affect the timing of existing RPO backups.

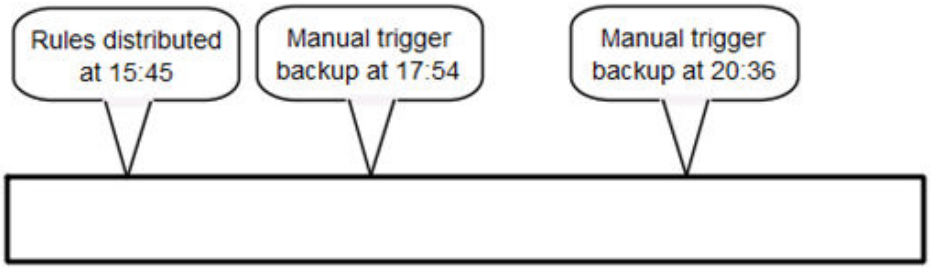


Figure 40 Manual injected backup events occur at the time the event is received (17:54, …)

- The existence of an RPO Backup Window will prevent RPO based backups from occurring outside the defined window. Any RPO based backups that are scheduled outside of the backup window will be deferred until immediately after the next backup window opens (the operation is greyed on the diagram to indicate deferral). Successive RPO based backups will be rescheduled to occur relative to the deferred backup time.

Note: Backup events injected manually via the UI or CLI are not affected in any way by backup windows.

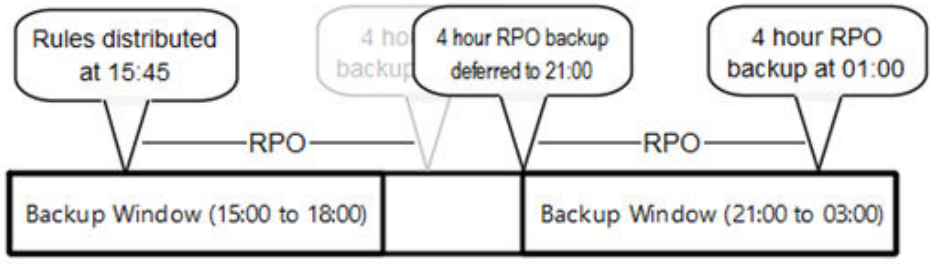


Figure 41 RPO based backups are deferred and rescheduled (21:00, …) if outside a backup window

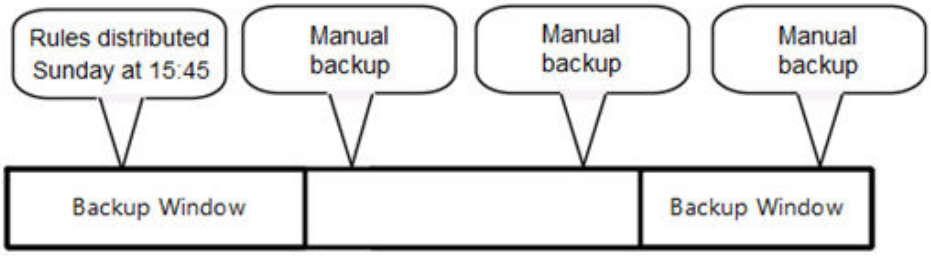


Figure 42 Manually injected backup events are not affected by backup windows

- The existence of an RPO Exclude Window period will prevent both Trigger Time and RPO based backups from occurring within the exclusion period. Any Trigger Time or RPO based backups that occur inside the exclusion period will be deferred until immediately after the exclusion period ends (the operation is greyed on the diagram to indicate deferral). Successive RPO based backups will be rescheduled to occur relative to the deferred backup. Trigger Time based backups proceed at their normally scheduled times outside of an RPO Exclude Window.

  Note: Backup events injected manually via the UI or CLI are not affected in any way by exclusion windows.

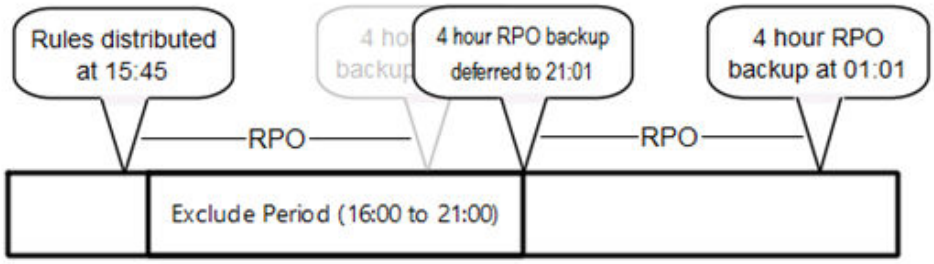


Figure 43 RPO backups are rescheduled (21:01, …) if they occur inside an exclude period

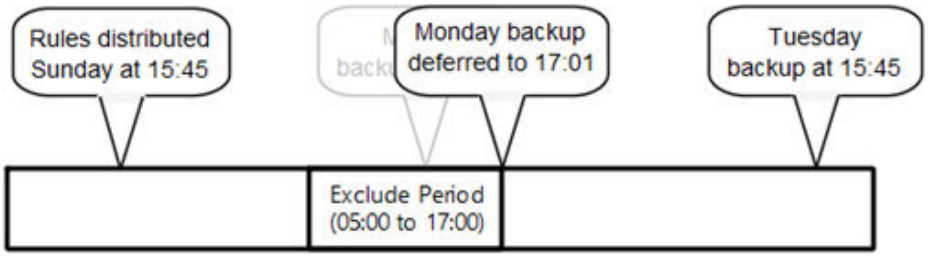


Figure 44 A Trigger Time backup is deferred (17:01) if it occurs inside an exclude period

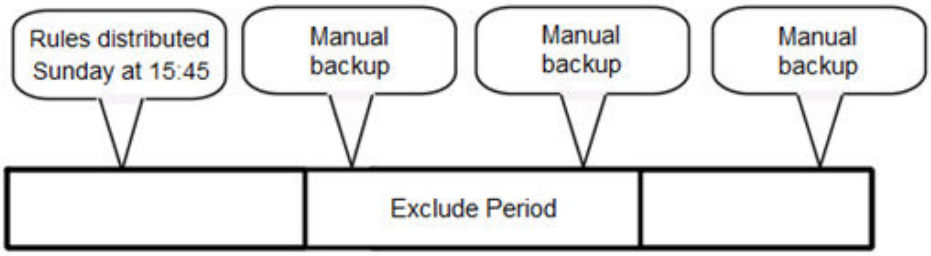


Figure 45 Manually injected backup events are not affected by exclusion periods

Monitor Concepts

This section describes Ops Center Protector's monitoring features.

For further information, refer to:

- Monitor Tasks (on page 260)
- Monitor UI Reference (on page 520)

About monitoring

Ops Center Protector provides real-time displays for all active data flows. The same data flow diagrams created when configuring data protection policies are used to show the status of participating nodes in real-time. The data flow diagrams are annotated with status icons on nodes and connectors, augmented with status information and various actions can be performed that are relevant to the selected node or connector.

When a node on the data flow is selected:

- The Jobs and Logs relating to that node are displayed alongside the data flow.
- Policy operations relating to that node can be triggered.
- If a status warning icon is displayed adjacent to the node then a description of the node status is provided.
- Depending on the node type selected, additional performance and status information pertaining to that node is displayed.
- If the node uses network bandwidth or on-board cache, then the current utilization is displayed.

Log Concepts

This section describes Ops Center Protector's logging features.

For further information, refer to:

- Log Tasks (on page 259)
- Logs UI Reference (on page 509)

About logs

Ops Center Protector maintains a list of log messages generated by the various subsystems running on Protector nodes. All log messages and are sent to and stored on the Master node, where they are made available via the user interface.

Each log entry includes the following information in addition to a short textual description of the event:

- A unique ID assigned to every log entry, useful when discussing logs with colleagues and support engineers.
- Audit Flag indicating a change to the system status. These cannot be modified or deleted, for compliance purposes. The identity of the user initiating the action is recorded.
- Date and time recording when the log was generated on the originating node and received by the master node.
- Level identifying the importance of the log entry.
- Attachment enabling additional context information to be attached to the log entry.
- An ID identifying the position within the Protector code where the message was generated. The ID can be used to reference an engineering support database that provides additional description and identifies possible causes and solutions.

Numerous filters can be applied to the displayed logs to enable the user to locate relevant entries.

Log entries that relate to a particular activity, such as a repository resynchronization, are marked with a session marker. This session marker can be used to filter the logs so that only the log entries relating to that activity are displayed.

Log entries can be acknowledged by the user and marked with a comment so that it is clear that the event has been noticed and addressed.

The entire log data-base can be exported in various file formats so that it can be analyzed offline or presented in a report.

If a specific event or set of events occurs in the logs then notifications can be generated via email, SNMP or the OS System Event Log.

Notification Concepts

This section describes Ops Center Protector's notification features.

For further information, refer to:

- Notification Tasks (on page 264)
- Notifications UI Reference (on page 659)

About notifications

Ops Center Protector supports Windows System Event Logging, email, SNMP and custom alerting mechanisms.

Whenever a log message is received from a Protector node by the logging process on the Master node, it is checked against a user defined, priority ordered list of notification conditions. Each condition in this list can issue an alert depending on whether the log message meets certain criteria and, if necessary, stop any further notifications lower down the list being triggered.

If your site needs to receive alert notifications based on some other mechanism, then you can create a custom configuration file with instructions for which script/program Ops Center Protector should run, and the arguments to pass to it.

Job Concepts

This section describes Ops Center Protector's job management features.

For further information, refer to:

- Job Tasks (on page 257)
- Jobs UI Reference (on page 493)

About jobs

Ops Center Protector monitors jobs that are being performed in the system and provides progress and status information. Jobs are broken down into tasks being processed by the system, each job will have one or more tasks. They are either system initiated tasks (e.g. a scheduled Backup or Report creation) or user initiated tasks (e.g. a Restore or a Repository Analysis). Jobs are executed by various subsystems on various Protector nodes and are tracked by the master node so that it can provide progress and status information to users.

Jobs are split into three types:

- A Job can be in one of the following states:
  - Succeeded – All Tasks succeeded
  - In Progress – One or more tasks are still running
  - Failed – One or more Tasks Failed
  - Paused – One or more Tasks Paused
  - Cancelled – One or more Tasks Cancelled
  - Unknown – The master is unable to communicate with the node on which the job is running. The node may be genuinely offline, or it could be that the network is down between the master and the node. This status will be updated when the nodes can communicate again.
- A Task can be in one of the following states:
  - Succeeded – All good
  - In Progress – Still running
  - Failed – All not so good
  - Paused – User has paused the Task
  - Cancelled – User has cancelled the Task
  - Unknown – The master is unable to communicate with the node on which the Task is running. The node may be genuinely offline, or it could be that the network is down between the master and the node.
- Additionally, a Task/Job can complete with:
  - Warnings – the state is SUCCEEDED but something was amiss (e.g., Rules distribution with an offline node)

The Jobs user interface presents the user with a list of currently active jobs and jobs that have recently completed. It enables some jobs to be paused or cancelled.

Restore Concepts

This section describes Ops Center Protector's restore management features.

For further information, refer to:

- Restore Tasks (on page 279)
- Restore UI Reference (on page 763)

About restoring data

The restore dashboard provides access to all the available backups created or managed by Ops Center Protector.

Depending on the type of backup created, the following actions can be performed:

- Repository, HCP and Amazon S3 backups - can be restored to their original or new location.
- Hitachi Block snapshots and ShadowImage replications can be reverted or mounted to a new location.
- Hitachi Block replications can be paused, resumed or have their direction swapped[1].

(1). Swapping applies only to continuous TrueCopy, Universal Replicator and Global-Active Device.

About Hitachi block based replication swapping (takeover/takeback)

Protector allows the user to swap the direction of TrueCopy, Universal Replicator and Global-Active Device replications. When a replication is swapped, the S-VOL takes-over the role of the primary volume and the P-VOL takes over the role of the secondary volume. A swapped replication can, of course, be swapped back to its normal state with the P-VOL as the primary and S-VOL as the secondary. The swap operation can also be used to restore consistency between Protector and storage where the direction of a replication in Protector does not match the current actual direction of the replication on storage.

A swap operation is typically performed either because maintenance is required, an application failure has occurred, a storage device has failed or disaster has befallen the primary site. Fail-over to the secondary site is therefore necessary.

For active-passive replications (TC and UR):

- If both P-VOL and S-VOL are operable and the link between the two sites is available then the replication will be reversed. This is called a swap takeover.
- If the replication cannot be re-established once in the swapped state (S-VOL to P-VOL), then the pair is placed in the *suspend for swapping (SSWS)* state until the problem is resolved. Once the problem is resolved, the swap must be completed by retrying the operation. This is called an S-VOL takeover.
- If a swap fully succeeds, the volume that is now the P-VOL will accept writes, and the volume that is now the S-VOL will not.

For active-active replications (GAD):

- A swap operation may be performed to move array processing load from the primary to the secondary device. If both P-VOL and S-VOL are operable and the link between the two sites is available, the secondary array will assume the higher processing load.
- If the replication pair has entered an *error* or *suspended* state, then once the problem is resolved, the site with the most recent data must be used to re-establish the replication. Because the replication is active-active and cross-path set-ups are possible, depending on the nature of the fault, the P-VOL or S-VOL could contain the most recent data:
  - If the P-VOL contains the most recent data, no swap is required:
    1. If necessary, unsuspend and unpause the replication.
    2. Resynchronize the replication (via manual trigger or data flow reactivation).
  - If the S-VOL contains the most recent data:
    1. Swap the replication to copy the data from the S-VOL to the P-VOL.
    2. Swap the replication again to restore the original direction. This is optional, but highly recommended.
- The swap operation will result in the both P-VOL and the S-VOL remaining writable.

The following limitations apply when performing a swap operation:

- The replication must be implemented using an active live mover; paused or batch replications cannot be swapped.
- The replication must be implemented using TrueCopy, Universal Replicator or Global-Active Device.
- Snapshot (TI) and local replication (SI and RTI) operations can be assigned to the P-VOL or S-VOL, however remote replication (TC, UR, GAD and Failover) operations cannot.
- The state of the takeover is shown on the Monitor page and via the Storage page.
- The destination proxy node (ISM) must be available and active since it holds the required metadata for the replication.
- Volumes cannot be added or removed while the replication is reversed or in the SSWS state.
- When a GAD replication is swapped, the original P-VOL is marked as reserved. If the replication is torn-down whilst in the swapped state, the reserved flag on the original P-VOL will prevent it being reused. By design Protector never modifies P-VOL properties, so the reserved flag must be manually reset. It is only safe to do this if the replication is no longer under Protector's control; as is the case when tear-down is complete.
- Takeover is not integrated with any application swap automation (high availability fail-over or cluster management) software. Therefore additional steps may be required at the OS and/or application level to complete the process.
- It is important that Protectors view of the replication direction matches that of the storage. Replications should not be swapped outside of Protector. If this has happened, it is important that the replication direction is corrected to match Protectors direction before any non-swap operations are performed within Protector. This can be done using the swap operation to either alter the replication direction on the storage to match Protector or vice versa depending on the direction chosen in the wizard.

About restoring legacy application data

During the life cycle of Ops Center Protector, we add new functionality in each release. However, overtime some functionality had to be deprecated and removed from the product. To ensure your backups are usable, Ops Center Protector aims to enable restores of previously created backups even if the capability to create new backups of the same type is no longer available.

The following section details how to Restore data for backups created with legacy features

Restoring legacy Microsoft Exchange Server backups

Host-based

Microsoft Exchange Server backups can be restored as files to a user-specified location. You can then use Exchange Server's management tools and/or PowerShell commands to use them.

Block-based

Block-based backups can be mounted to a Windows client which enables access to the files. You can then use Exchange Server's management tools and/or PowerShell commands to use them.

It is not possible to revert block-based Exchange Server backups.

Restoring legacy Microsoft SQL Server backups

Host-based

Microsoft SQL Server backups can be restored as files to a user-specified location. The backups can be attached to an existing SQL Server instance using the SQL Server's management tools.

It is not possible to apply transaction logs to the restored databases.

Block-based

Microsoft SQL Server backups can be mounted to a Windows client, which provides access to the files. The backups can be attached to an existing SQL Server instance using the SQL Server's management tools.

It is not possible to apply transaction logs to the restored databases. It is also not possible to revert block-based SQL Server backups.

Restoring legacy SAP HANA backups

Host-based

Not applicable.

Block-based

SAP HANA backups can be mounted to a Linux client which makes the files available at a user-specified path. In addition, the configuration files which are part of the backup are restored to the following location:

```
$INSTALL_DIR/saphana_backup/<SID>/<backup_id>
where:
$INSTALL_DIR is the Protector installation dir,
<SID> is the SAP HANA SID,
<backup_id> is the snapshot ID
```

Restoring legacy CDP and Live backups

Host-based

There are no changes to the restore behavior, CDP and Live backups can be restored normally.

Block-based

Not applicable.

Storage Concepts

This section describes Ops Center Protector's storage management features.

For further information, refer to:

- Storage Tasks (on page 280)
- Storage UI Reference (on page 858)

About storage

The storage inventory provides access to all the available data storage types managed by Ops Center Protector as follows:

- Generation 1 and Generations 2 Repository - batch backups of applications, VMs and filesystems.
- Hitachi Content Platform - batch backup of applications, VM's and filesystems.
- Amazon S3 - batch backup of applications, VM's and filesystems.
- Hitachi Block and Logical Block - snapshots and replications of applications, VMs, filesystems and LDEVs.

For each device type, selecting an entry in the inventory will display the details screen where the following are displayed:

- Logs relating to the device
- Jobs relating to the device
- Device configuration
- Storage and Store Details

Depending on the type of storage selected, the following actions can be performed:

- Repositories:
  - Mount/Unmount
- Hitachi Block:
  - View all host groups, journals, logical devices
  - View pool usage
  - Mount, unmount, revert, adjust expiry, transfer ownership or delete snapshots
  - Mount, unmount, pause, resume, swap, dissociate or delete replications

About Repositories

About Gen 2 Repository

The Protector Repository Gen 2 is a performance and capacity optimised object store. It conforms to the Unified Backup Infrastructure storage specification. Refer to the Unified Backup Infrastructure concept (on page 35) for more details.

The Repository differs from other native object storage in that it supports snapshots natively utilising re-direction on write. This allows for highly efficient management of shared data blocks and space reclamation is both fast and automatic.

There are no practical limits to the size of objects or number of objects in a repository.

To achieved optimal data reduction, it retains both object level and block level check sums.

For performance optimization, the Gen 2 Repository can be setup to be split across different storage. See Repository Storage Node Wizard (on page 632) for more details.

- Data directory – this is the where the bulk of the data is stored. Mass storage devices with performance characteristics of fast write streaming storage is ideal.
- Metadata Directory and Checksum Directory holds the metadata and the repository structure information. These benefits from fast random access such as SDD and NVMe storage. Typically these should be around 5% of the Data Directory capacity.

Hardware and OS recommendations

This section describes the ideal additional specification information for a repository machine. For more information see the product support matrices at https://www.hitachivantara.com/en-us/products/data-protection/ops-center-protector.html#tech-specifications .

Hardware:

As usual in most scenarios more is better but a point of diminishing returns will be reached. The most beneficial upgrade is extending the available memory. The repository storage group processes themselves never use a large amount of memory, but increasing memory beyond 16GB will ensure there is plenty of head room.

More important when specifying CPU, is the generation of CPU technology. Because moving mass data is at the core of what the repository does, faster memory will have a greater impact than for a normal application so the faster memory the better.

File System:

In addition to optimising the disk layout, it is also important to use the correct block size for NTFS. NTFS volumes should be formatted using 16k block size with a repository storage group default setting. It should under no circumstances be larger than the repository `DataBlockShift` setting which is 16k by default. If the NTFS block size is larger, then it will force the filesystem to read the block before it performs a write.

E.g. with a 64k NTFS block size, if Ops Center Protector needs to update a single repository block of 16k, it will write out the 16k block to the file system. However the file system will need to read the 64k NTFS block, change the relevant 16k section and then write back the full 64k block. This badly degrades performance, but it may not be noticeable when the repository is initially populated since these full backups would likely be large 2MB writes across several blocks. The degradation in performance would only be noticed over time when smaller, incremental updates start occurring.

Shutdown Time:

Windows only allows a few seconds for all services to shutdown before forcibly terminating them. An Ops Center Protector repository storage group can take a significant amount of time to close especially if it is busy synchronizing many machines at the point of a system shutdown.

The time Windows allows for process shutdown can be changed via the Windows registry. The registry entry to change this is:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control
\WaitToKillServiceTimeout
```

The value is in milliseconds; by default it is either 12000 or 20000 depending on OS version. It is important for the repository to shutdown correctly so we highly recommend changing this value to 5 minutes (300000 milliseconds).

Note: Incorrect editing of the Windows registry can cause system-wide problems that may require you to re-install Windows to correct them.

The knock on effect is that your server may take longer to shutdown, this setting applies to all services, so if a service is badly behaved and does not shutdown when asked, then the machine will take the full 5 minutes (plus OS cleanup time) to shutdown.

About Cloud Nodes (Amazon S3, HCP and HCP Cloud Scale)

Cloud storage nodes, Amazon S3, HCP and HCP cloud scale are used to transfer application data to cloud storage. All use the S3 protocol for communication.

Cloud storage nodes represent the storage in Amazon S3 cloud or on the HCP. When creating the node, you will define the necessary credentials to access the S3 service. These include:

- Access Key ID
- Secret Access Key
- Bucket Name
- Region (Amazon S3 only)

Additionally, you will define a Proxy node that will be the intermediary between the application and the S3 service. The proxy node can be any Windows or Linux OS nodes.

The role of the proxy node is to:

- Provide a secure system that can access the S3 service. Although the proxy node can be the same OS host as an application, it is recommended for security reason that this is an independent system to that of the application. The benefit of this is that this system can be given limited access ensuring that the application servers do not have access to the S3 credentials. This S3 node can then be used by multiple applications without the danger of compromising security or data leakage. It also means that the application servers do not need direct access to the internet.
- Provide a cache of the recovery points stored on the S3 server. This cache is only a copy of the index information and metadata, it is not the data itself and can be rebuilt from the S3 content. Having the cache however improves responsiveness of the system and decreases some of the S3 access, thereby saving cost.
- Provide an UBI compliant application interface. This allows you to use S3 storage with any of the next generation applications in the same fashion that you can use Generation 2 Repository. For a list of UBI compliant applications see XYZ. As UBI compliant storage, the S3 node can also receive data from or send data to other UBI compliant storage, such Generation 2 Repository or Generation 2 Hitachi Content Platform. This means it can be used to form a layered data retention solution.

Cloud storage node key features include:

Search friendly native format

A key feature of S3 support is to keep data stored on S3 in a native format as much as practicable. Any files less then 4GB in size, has its data stored in a single object. This means that 3rd party search, indexing and analytics tools can process the data in S3 directly. By default, files bigger than 4GB are split into 4GB segments that can easily be stitched together if required.

Delta block differencing

To reduce bandwidth and improve transfer speeds, by leveraging S3 partial transfers only the block changes within a file are sent to S3.

Virtual snapshots

A system whereby data is retained and indexed in S3 in such a way that recovery points use objects that are whole and not need rehydration but will share data that is common between other recovery points, saving storage space and providing optimal recovery speed.

Multi-factor parallelism

To further optimise data transfer speed in addition to delta block differencing, multiple objects are sent to S3 at the same time and multiple data blocks within those objects are sent separately to S3 in parallel.

Self-consistent store

Although some metadata is cached on the proxy, the data stored in S3 is self-consistent. Should a site experience complete data loss, a new instance of Protector can be spun and an S3 node can be configured to import the data.

Report Concepts

This section describes Ops Center Protector's reporting features.

For further information, refer to:

- Report Tasks (on page 278)
- Reports UI Reference (on page 739)

About reports

Ops Center Protector is able to generate a number of reports viewable both within the UI and exportable in various formats: The reports are split in to 3 different groups:

- **Jobs**:
  - RPO Report - Generated when accessing the report page and also generated on a daily basis, lists operations that met and/or failed to meet their specified RPO.
- **Block Storage**:
  - HORCM Count Report - An ongoing report listing the actual temporary and current HORCM instances for each Block Storage Device.
  - Pool Usage Report - An ongoing report listing used and available pool capacities and thresholds for each Block Storage Device.
  - Journal Usage Report- An ongoing report listing used and available journal capacities for each replication policy on each Block Storage Device.
  - Network Transfer Report - An ongoing report listing last, in-flight and last-received Q markers and transfer times for each replication policy on each Block Storage Device.
  - All Pairs - This report provides a list of all the replication pairs. There is no grouping. If multiple LDEVs are part of replication setup in Protector or they were together in a HRCM file, each individual LDEV will be shown as a replication. See About Global Replication reports (on page 120) for more information.
  - Pairs by Storage Record - This report shows replication grouped by Protector record groups. This would be the replications shown inside a Protector dataflows. See About Global Replication reports (on page 120) for more information.
  - Pairs by Consistency Group - This report shows replications grouped by Array Consistency Group. See About Global Replication reports (on page 120) for more information.
  - Pairs by Copy Group - This report shows replications grouped by Array Copy Groups. See About Global Replication reports (on page 120) for more information.
  - Pairs by Snapshot Group - This report shows snapshots grouped by Snapshot Groups. See About Global Replication reports (on page 120) for more information.
- **Host Based Storage**:
  - Repository usage report - An ongoing report listing used size and total size, deduplication queue and tier queue lengths for each repository.
  - HCP usage report - An ongoing report listing used size, capacity and object count for each HCP node.
  - HCP cloud scale report - An ongoing report listing used size, capacity and object count for each HCP cloud scale node.
  - Amazon S3 usage report - An ongoing report listing used size for each Amazon S3 node.
  - Backup report - An ongoing report listing many parameters relating to repository, HCP and Amazon S3 backups. Each backup’s session logs can be viewed from the summary report.

About Global Replication reports

The Global Replication Reports provide a method of giving a complete list of all replications on the arrays including replications that are not managed by Protector. This feature works by interfacing to Ops Center CM-REST which provides an inventory of replications on individual arrays. Protector caches the result from CM-REST and aggregates the results.

Key features

- Report on all replications on arrays
- Includes both replications that are managed by Protector and replications that are not
- Provide a list of individual replications pairs or grouped replication report Group replication reports include
  - Pairs by Storage Record report - List of replications and snapshots that are managed by Protector and have a replication or snapshot record
  - Pairs by Consistency Group report - List of replications and snapshots that have a common consistency group
  - Pairs by Copy Group report - List of replications that have a common copy group
  - Pairs by Snapshot Groups report - List of snapshots that have a common snapshot group
- Reports have predefined filters and an advanced filter.
- Ability to drill down on a replication / snapshot listing and view further details (Detail View).

Pre-requisite and Setup

To provide the reports, Protector queries Ops Center CM-REST. Consequently it is a requirement that CM-REST is configured for every array where the report is required. When setting up the Protector Block Node, there is now an additional section where you can specify the connection and credentials details for the REST interface on CM-REST.

When setting up a Block Node, it is optional to add the CM-REST information. You cannot however have a Block Node configured to use the CM-REST interface only.

Ops Center CM-REST can be configured to manage multiple arrays in the same way that a Protector node might manage multiple array. However it is not a requirement that Protector and CM-REST are configured symmetrically each dividing up the management the same way. It is also possible for CM-REST and Protector be installed on the same machine or on different systems.

Note that if CM-REST and Protector reside on the same array and are configured to use In-band Command device, they must be configured to uses different array accounts. In the configuration where CM-REST and Protector both reside on the same host, if Protector and CM-REST share the same array account credentials there will be unpredictable behaviour and failures will occur.

Architecture

Each host that services a Protector Block Node, eg array, has a database to store the replication information for that array. Protector queries CM-REST every 1 hour by default to get the inventory of all replications for a configured array and stores the result in the database. When the user creates a report, the Protector Master will query each of the nodes that manage the block host to query their databases. The results will then be merged to provide the federated results into a single report. Because replications are frequently across two arrays, it is possible to get the result from both sides of the replications (eg from the P-VOL and from the S-VOL array). In that situation, Protector will choose the result that had the newest refresh. In the extended view of the report, the field "Cache Updated" gives an indication of how recently the database had been updated for that record.

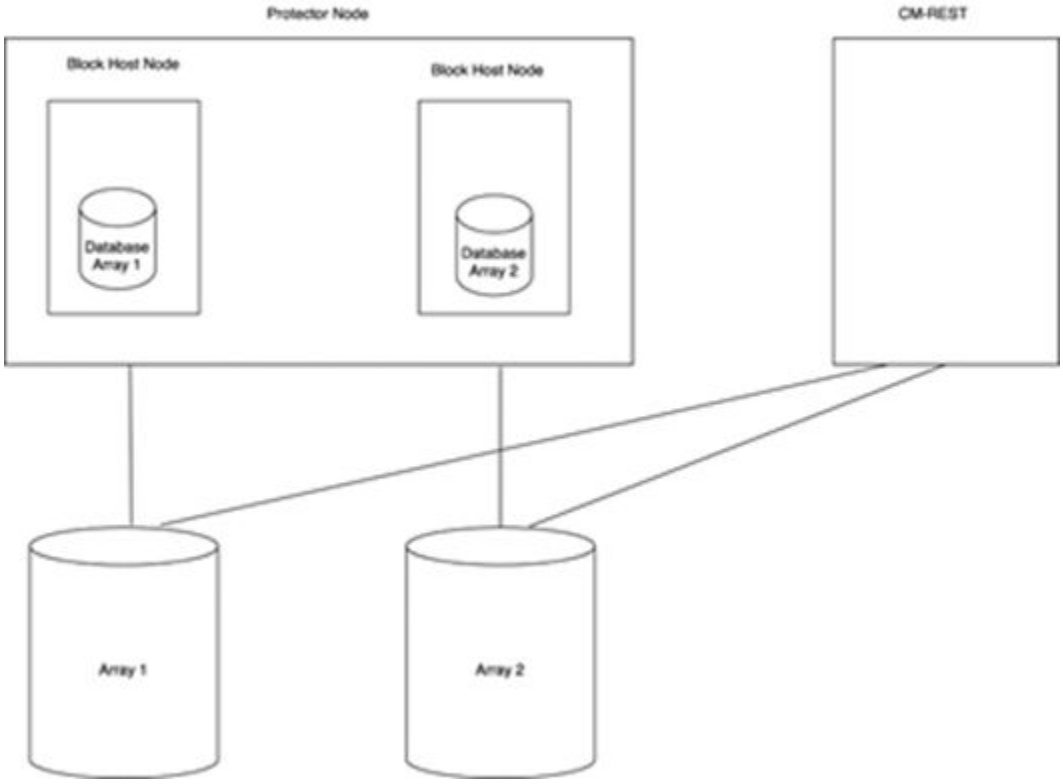


Figure 46 Global Replication Reports - Architecture

Report Choices

There are multiple choices of reports split into how the replications should be grouped. By default the reports will report across all replications on all arrays. Each report has extensive filtering capabilities that allow the reports to be more manageable in size and targeted towards the users needs. Note that the maximum number of return values is capped by default to 1000 rows. When returning to a report in the UI, the last search result will automatically be cached.

All Pairs

This report provides a list of all the replication pairs. There is no grouping. If multiple LDEVs are part of replication setup in Protector or they were together in a HRCM file, each individual LDEV will be shown as a replication.

Pairs by Storage Record

This report shows replication grouped by Protector record groups. This would be the replications shown inside a Protector dataflows.

Pairs by Consistency Group

This report shows replications grouped by Array Consistency Group.

Pairs by Copy Group

This report shows replications grouped by Array Copy Groups.

Pairs by Snapshot Group

This report shows snapshots grouped by Snapshot Groups.

Access Control Concepts

This section describes Ops Center Protector's access control features.

For further information, refer to:

- Access Control Tasks (on page 213)
- Access Control UI Reference (on page 309)

About Role Based Access Control (RBAC)

Role Based Access Control (RBAC) is a framework for defining what a user can see and do within Protector. Users are only allowed to perform activities that are contained within the roles they are assigned. Furthermore they can only perform those activities on the resources to which they have access.

When setting up access control, the following must be defined:

- How a user's credentials will be verified (*Authentication*).
- What rights of access a user will be granted (*Authorisation*).

About Access Permissions

Data flows, policies, schedules and notifications are created in Protector to configure its behaviour. Access to these configuration objects is controlled by granting *Permissions*. Permissions determine if a particular user or group has *Read Only Access* or *Read Write Access* to a configuration object.

Permissions are a secondary access check within Protector.

Taking policies as an example:

The system will first check if a user has the RBAC activity *View Policies*. If they don't then they won't be able to see anything. If they do then for all policies that exist, the system looks at the permission on each one. E.g:

Policy1 - created by Paul@Contoso.com has READ/WRITE access (he can see Policy1)

Policy2 - created by Ian@Contoso.com has READ/WRITE access (Paul cannot see Policy2; Ian can)

Policy3 - created by Ian@Contoso.com has READ/WRITE access, shared with Paul@Contoso.com with READ access. paul can see Policy3 but he cannot edit it.

Policy4 - created by Ian@Contoso.com has READ/WRITE access, shared with Contoso.com \Managers group with READ access. Paul is not in the managers group so he can't see Policy4. However, Simon@Contoso.com is, so he can see Policy4 but can't change it.

About Ops Center Protector's implementation of RBAC

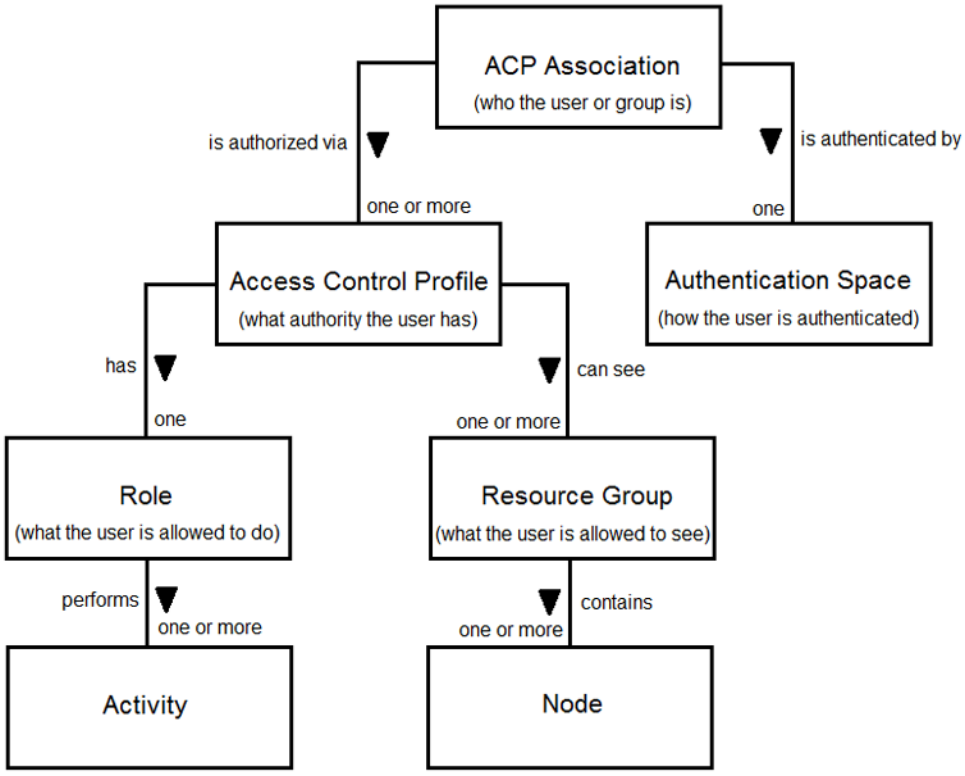


Figure 47 Ops Center Protector's RBAC Object Structure

An *Activity* is an operation (typically of the type: create, read, update or delete) that is performed on a resource, backup or other object within Protector. For example:

- a node may be created and have its properties viewed or modified.
- a data flow may be created, viewed, edited or deleted.
- a snapshot may be created, browsed, restored or retired.

Activities are arranged into *Activity Groups* that represent a cohesive set of activities relating to features such as *Nodes*, *Data Flows* and *Logs*. Activity groups make it easier to organise, locate and assign activities.

Roles contain a set of activities. For example, the role *Backup Administrator* is assigned to a user tasked with managing nodes, creating backup policies and restoring backups, whereas the role *Security Administrator* is assigned to a user responsible for authentication and authorisation of other users.

Resource Groups define a collection of physical and logical computing or storage resources that a user is permitted to access (such as groups of servers, storage devices, repositories etc.). The visibility of any given *Resource* within the Protector system is determined by whether that resource appears in the ACPs assigned to the user at authorization time.

Access Control Profiles combine *Roles* and *Resource Groups* to define the activities that can be carried out on groups or resources.

Authentication Spaces specify which authentication service to use for a specific user or group of users. Authentication services supported include Active Directory, LDAP, RADIUS and Local system logon. Protector enables you to simultaneously configure multiple authentication services. Protector also supports OpenID Connect to enable integration with Hitachi Ops Center's single sign-on functionality.

ACP Associations link individual users and groups of users to one or more *Access Control Profiles* so as to grant them the required level of authority. Authentication is performed when a user or group attempt to log on, by passing a user's credentials to the authentication service specified by the *Authentication Space*.

Backup data sets (such as the volumes and files included in a snapshot or a replicated file system) are created, archived, restored and retired by Protector as the result of executing backup policies. These backup objects store information about the resource from which the backup data originated. The visibility of the backup object (in some storage location such as a repository) is then generally governed by whether the user has access to the node from which the data originated.

To enable RBAC to be configured quickly and easily, Protector is shipped with a number of predefined Activity Groups, Roles, a default Resource Group, Access Control Profiles and an administrator level ACP Association (refer to How to configure basic role based access control (on page 213)). These predefined RBAC objects can, if required, be tailored to suit each customer’s specific environment, as described in How to configure advanced role based access control (on page 215).

Selecting an Authentication Service

Protector can be configured to authenticate users and groups against the following:

- Local accounts on a Protector node running any supported OS.
- Accounts in Active Directory via a Protector proxy node running Windows.
- Accounts on an LDAP Server via a Protector proxy node running Linux or AIX.
- Accounts on a RADIUS server via a Protector proxy node running any supported OS.
- Accounts in Hitachi Ops Center via OpenID Connect to support single sign-on.

When a user logs in via the Protector web UI they must provide either their:

- Username – to identify themselves to Protector.
- Authentication Space - to tell Protector which authentication service to use.
- Password – to authenticate themselves.

or, if integrated with Hitachi Ops Center, redirect to the single sign-on page to provide their Ops Center credentials.

Authentication spaces

Protector uses the concept of *Authentication Spaces* when authenticating users. The Authentication Space tells the Protector master node which authentication service is able to validate a specific user's login credentials. Authentication requests are routed, by the master, to the proxy node responsible for communicating with the authentication service in question.

Although it is possible to nominate the master node as the proxy, it is often the case that the master node is either not in the required space or is of the wrong OS type. Consider the following scenario:

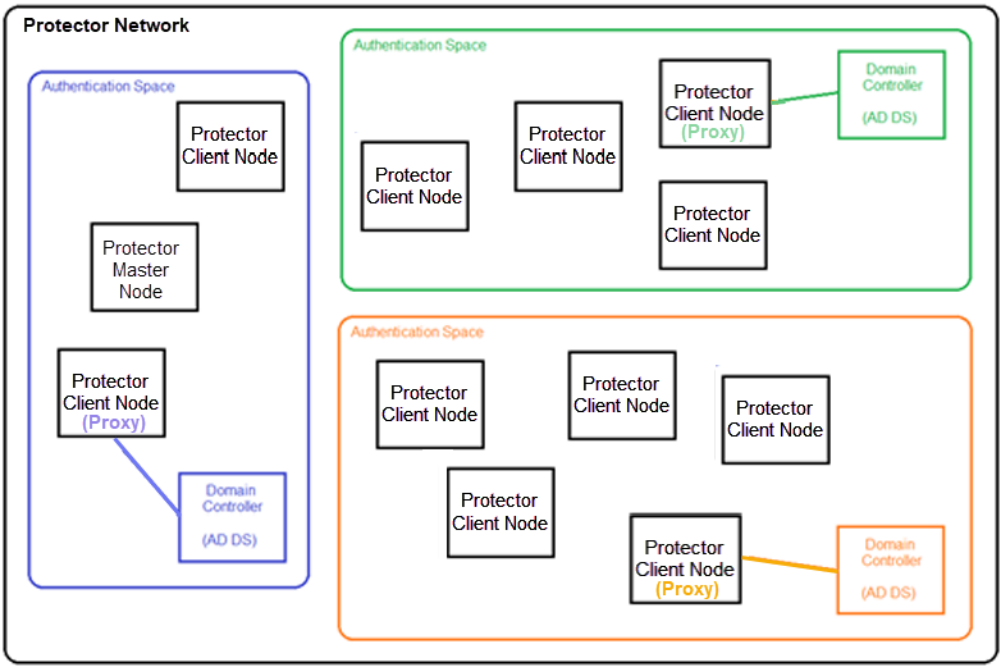


Figure 48 Authentication in a multi-space Protector environment

Here the Protector master node is responsible for nodes in three separate spaces, so a separate proxy is required to communicate with each authentication service. It is possible to have different authentication service types for each space.

Each *Authentication Space* requires configuration within Protector. The name of the authentication space forms part of the user's login. For example:

- An *Active Directory* authentication space could be configured to point to the AD server `ContosoDC.HV.local` and could be named `Contoso`. Users in this AD domain must authenticate with Protector using `<username>@Contoso`. Here `Contoso` is used as an alias for `ContosoDC.HV.local`. If you want to use the full domain name then the authentication space must also be named using the full domain name.
- A *Local Machine Accounts* authentication space is configured to point to the computer `TBell-Win7PC` and could be named `TBW7`. Users with local accounts on the machine `TBell-Win7PC` must authenticate with Protector using <username>@TBW7. Here `TBW7` is used as an alias for `TBell-Win7PC`. If you want to use the full computer name then the authentication space must also be named using the full computer name.

About Single Sign-On with Hitachi Ops Center

Protector can be integrated with Hitachi Ops Center and supports Ops Center's OpenID Connect based single-sign-on mechanism.

When Protector is installed with Ops Center, it is registered with Ops Center's OpenID Connect Provider. A corresponding OpenID Connect authentication space is also automatically added to Protector during the regsitration process.

When registered with Ops Center, Protector displays a button on it's login page that enables users to be redirected to Ops Center's authentication page. Once authenticated by Ops Center's authorization server, users are redirected to the Protector UI and granted access as normal.

The OpenID Connect authentication space added to Protector can be used, in the same way as any other type of authentication space, to create the required ACP Associations that grant the required permissions.

Note: Only Group ACP Associations are supported for Ops Center OpenID Connect authentication spaces.

Authorising users and restricting access

Authorisation is granted to a user, group or entire authentication space by associating it with one or more ACPs. Each ACP defines a role that the user is assigned and one or more resource groups:

- The role defines what activities the user is allowed to perform.
- The resource groups define which resources the user is able to interact with.

In practice, an ACP is constructed by considering the following:

- What responsibilities will the users be given?
- What should the users be allowed to do?
- Which nodes should be visible to the users?

Bear in mind that if a resource is not included in a user's resource group then it will not appear in any Protector view for that user, be it node lists, data flow diagrams, monitor views, reports, storage inventories, logs etc. So consider whether users should be prevented from seeing resources or simply restricted in what they can do with those resource

Default access control configuration

Access control is configured to default setting immediately after installing the Protector master for the first time. The Protector administrator must log on via the web UI using the *<username>* credentials specified in the Master installation. This user is granted full access to everything within Protector allowing them to configure access control and any other aspect of Protector. The following access control objects are automatically defined at installation:

- The *Master* authentication space is used to direct authentication to the OS on the master node.
- The *<Username>@Master* ACP association grants the *<username>@master* user the built-in *Default Administrator* ACP.
- The *Default Administrator* ACP is given the role Protector *Admin* which allows all activities to be performed. This ACP also grants access to all nodes and all backups on any storage location.
- The *default* resource group includes all Protector nodes that identify themselves to the master node, be they authorised or unauthorised.

Note: It is recommended that: the top echelon of administrator accounts are created first, allowing lower level administrators to configure their resources and access rights independently.

Planning roles

Roles tend to follow reasonably consistent patterns across organisations, so in a multi-tenant environment for example, it is worth considering how roles can be defined so as to be reusable across each tenant's environment. When roles are defined at a general level they can be reused in ACPs for defining specific privileges. Thus a role such as *Backup Administrator* could be reused by multiple ACPs such as *Accounts Backup Admin*, *Legal Backup Admin* and *Production Backup Admin*. What differs between these ACPs are the accessible resources, not the activities that will be performed on them.

Protector ships with a number of pre-defined roles that can be cloned and modified or used as-is.

Planning resource groups

Resource groups are the mechanism for restricting the visibility of nodes in Protector. When a user is associated with an ACP, they will be restricted to viewing only those nodes listed in the resource groups included in that ACP.

Note: Resource groups control visibility of nodes. What the user is able to do with those nodes is dependent on the activities they are allowed to perform on them.

Resource groups are typically defined based on tenancy, organisational, divisional and departmental hierarchies. These hierarchies are likely to exist already in the IT infrastructure and can thus be reused as a basis for creating resource groups.

Applying access levels

An *Access Level* is attached to each resource group in an ACP. The access level controls which backups (including any logs or reports relating to that backup) are visible, and can be set to one of the following:

- FULL – All backups in a given storage location (e.g. a repository) are visible to the user irrespective of where the data originated from.
- LIMITED – Backups are visible if they originated from any of the nodes contained in the ACPs assigned to the user.

  For LIMITED to work the storage node needs to be in the same resource group as the source nodes. For example, if a resource group is created with only a storage node in it and a user has FULL access to it, they can see all of the backups on that node. If the user has LIMITED access they can see no backups even if they have access to other resource groups.



Note: The access level only controls the visibility. What a user is able to do with visible backups is dependent on the activities they are allowed to perform on them.

A user given LIMITED access to a storage destination included in their resource group will only see log messages for that storage destination if they pertian to their backups. A user given FULL access level will see all log messages for that destination.

License Concepts

This section describes Ops Center Protector licensing.

For further information, refer to:

- License Tasks (on page 258)
- Licenses UI Reference (on page 506)

About licenses

Before using Ops Center Protector, you must enter one or more license keys using the user interface. You can obtain perpetual license keys that will give unlimited access to the features of the product by contacting your Hitachi Sales representative and providing the machine ID listed in the Protector License screen. The perpetual license keys are locked to the machine ID, which is locked to the hardware of the machine it is running on.



Note: If the MAC address of the Master nodes network card changes this will invalidate the Protector license. When running on Microsoft Hyper-V it is necessary to configure the Master node to have a Static MAC address so that it is not changed dynamically if the VM changes Host.

A five week trial license is included with Protector, this will give unlimited access to the features of the product with the exception of encrypted repositories.

Note: Upgrading protector will invalidate the trial license.

UI Concepts

UI Timeout when Protector is linked to Ops Center

When Protector is linked to an instance of Ops Center the UI timeout will now be based on the user's actual activity on the UI (e.g. mouse / keyboard movement) rather than whether the UI has sent an API call to the back end.

There are two modes of operation which are governed by Ops Center settings:

1. Timeout after user inactivity regardless of which screens the user is on.
2. Timeout after user inactivity only if the UI does not refresh its information. This would mean a status or log page that periodically updates the screen would prohibit the UI from ever timing out. This is useful if the user wishes to leave screen up as a dashboard but may have unwanted security implication.

When the user is timed out, the UI will return the Protector logon screen. From the logon screen the user can elect to "Login with Ops Center" whereby the user would login depending on the current login state within Ops Center.

Chapter 3: Data Protection Workflows

This chapter describes typical workflows for backing up and restoring data using Ops Center Protector.

Data Protection Workflows Overview

Each data protection workflow described here is a high level task that references of a series of subtasks.

Subtasks are detailed, lower level descriptions that explain how to do something via the UI or CLI.

Workflows and subtasks both refer out to user interface reference topics that describe the details of a particular web page, wizard or dialog that the user interacts with.

When following a workflow for the first time, it is recommended that you study the related subtasks and reference topics before carrying out the steps. This will ensure that you are aware of any notes, warnings and cautions that may apply.

Once you are familiar with the workflows and subtasks, it may only be necessary to refer to the reference topics, since all the important notes, warnings and cautions are located within the reference topics at the relevant point.

Tip: Refer to Button Icons (on page 296) to find the icons associated with the control names used in the step-by-step guides.

How to protect your data

Before you begin

This generic workflow describes the steps for protecting your data with Ops Center Protector.

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the source nodes and/or proxies.
- Any applications being backed up have been installed as per the Protector requirements and prerequisites. Refer to the Protector support matrices at https://www.hitachivantara.com/en-us/products/data-protection/data-instance-director.html#tech-specifications and the appropriate Protector Application Guides listed in Related documents (on page 17).
- The Protector Client software and prerequisites have been installed on the destination nodes and/or proxies.
- Any destination storage devices have been set up as per the Protector requirements and prerequisites. Refer to: .
  - Hitachi Block prerequisites (on page 21)
  - Generation 1 Hitachi Content Platform prerequisites (on page 24)
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. Refer to How to configure basic role based access control (on page 213) or How to configure advanced role based access control (on page 215).

This task describes the general steps to follow when protecting your data:

Procedure

1. Create the required source nodes and then check that they are authorized and online.

   Source nodes represent the places where the data to be protected resides in your system. See How to add a node (on page 261) and How to authorize a node (on page 262).

   Note: Source Hardware Storage Device nodes will need to be created even if they don't appear on the data flows for snapshot and replication operations.

2. Create the required destination nodes and then check that they are authorized and online.

   Destination nodes represent the places where the data is to be backed up to. See How to add a node (on page 261) and How to authorize a node (on page 262).

3. Define the data protection policies.

   Policies define the data sets to be protected and the methods to be used to protect them. See How to create a policy (on page 273).

4. Draw the backup data flows.

   Data flow diagrams show the participating source and destination nodes and the data paths interconnecting them. See How to create a data flow (on page 228).

5. Assign the policies to the participating nodes on the data flows.

   Policy assignments define precisely how the data flows from each source node to the respective destination nodes. See How to apply a policy to nodes on a data flow (on page 231).

6. Compile and activate the data flows.

   The source and destination nodes work autonomously by implementing rules locally. These rules are generated by the master node and disseminated to the participating nodes. See How to activate a data flow (on page 232).

7. Trigger the policies if required.

   Policies are invoked according to a user defined schedule and/or RPO. In some cases it may be necessary or desirable to trigger policies manually. See How to create a schedule (on page 280), How to trigger an operation from an active data flow (on page 260) and Triggering policies and operations with hdidcmd (on page 951).

8. Monitor the data flows, logs etc. to ensure policies are operating as expected.

   A number of other tools are provided that enable the correct functioning of data protection policies to be confirmed, including the:

   - Jobs Inventory (on page 493) - lists each backup job as it is scheduled.
   - Logs Inventory (on page 509) - lists information, warning and error messages.
   - Session Log Details (on page 517) - list related log message for each session.
   - Reports Dashboard (on page 739) - enables reports to be viewed and exported.
   - Notifications Inventory (on page 659) - allows notification of events via email etc. (see How to create a notification (on page 264).

9. Review the storage nodes to ensure backups are being created.

   The various storage devices require ongoing surveillance to ensure that they are operating correctly and sufficient resources are available to store your data securely. These are shown in the Storage Inventory (on page 858).

How to restore your data

Before you begin

This generic workflow describes the steps for restoring data that is protected with Ops Center Protector.

It is assumed that a data protection policy has been implemented and that at least one complete backup data set has been created on the designated storage device. See How to protect your data (on page 130).

This task describes the general steps to follow when restoring your data:

Procedure

1. Identify the destination where the data set is to be restored.

   Depending on the scenario, you may want to restore data to its original location or a different location

2. Ensure that the restore location is prepared to receive the restored data set. In the case of a replication, it will be necessary to prepare both sides of the replication pair.

   Applications that access the data will need to be placed in the correct state and sufficient resources must be available to receive the data set.

> **Caution:** The process of restoring data can result in the destruction of some original data that exists on the restore target. Ensure that any critical data is copied to a safe location and/or is included in the data set being restored.

3. Depending on the scenario, it may be necessary to suspend any backup policies currently active on the location where the data set is being restored.

   Data flows can be deactivated via the Monitor Inventory (on page 520) or Data Flows Inventory (on page 355).
4. Locate the data set to be restored.

   The Restore Inventory (on page 763) and Storage Inventory (on page 858) tracks all backups maintained by Protector.
5. Select the restore method and start the restore process.

   It may be possible to restore the data to the original location, a different location, mount a volume to a new host or swap the direction of a replication. The Restore Inventory (on page 763) and Storage Inventory (on page 858) provide functions to restore, mount, revert and swap replications and snapshots.
6. Once the restore process is complete, further steps may be needed to fix-up the data set before using it.

   For supported applications, these additional steps are described in the appropriate Application Guide (see Related documents (on page 17)). For other applications, consult the vendor's documentation.
7. Restart any applications that access the restored data.

   For supported applications, these additional steps are described in the appropriate Application Guide (see Related documents (on page 17)). For other applications, consult the vendor's documentation.
8. Resume any existing backup policies for the restored data set. If a new restore location was chosen, consider if it is necessary to introduce new backup policies.

   Data flows can be reactivated via the Data Flows Inventory (on page 355).

How to protect Ops Center Protector

Before you begin

In order for Ops Center Protector Master Node backups to be available following the total loss of the master node backups should be made to remote storage. Such storage may be a repository that is proxied by a client node HCP; HCP for cloud scale or Amazon S3.

When using HCP; HCP for cloud scale or Amazon S3 for backup either a client node or the master itself may be the proxy for the storage node.

When using a repository that repository must be hosted (i.e. proxied) by a client node. Note that only generation 2 repositories should be used.

Backing up the Protector Master Node Settings

Procedure

1. Create a policy. Refer How to create a policy (on page 273)
2. Add classification of type "Ops Center". Refer How to add a classification to a policy (on page 274)

3. For that classification select the master node, then under "Include Applications" select either "Include all Ops Center applications" or press the Add button and select "Ops Center Protector Master". Refer Ops Center Classification Wizard (on page 685) .
4. Add a backup operation to the policy with suitable schedule and/or RPO, etc
5. Create a dataflow. Refer How to create a data flow (on page 228) .
6. Drag the master node onto the data flow, the master node will be the source node.
7. Select the policy created above on that master. Refer How to apply a policy to nodes on a data flow (on page 231).
8. Drag the target storage node onto the data flow and connect the master source node to the repository. Refer How to connect nodes on a data flow (on page 231).
9. Select the backup operation in the above policy on the storage node. Refer How to apply a policy to nodes on a data flow (on page 231).
10. Activate the dataflow. Refer How to activate a data flow (on page 232).

    This complete the backup of the protector settings, they will be automatically backed up in accordance with the schedule and/or RPO selected within the policy.

    An additional backup may be performed by selecting the data flow from within the "Monitor" screen. Refer How to trigger an operation from an active data flow (on page 260).

How to recover Ops Center Protector

Before you begin

Ops Center Protector master node backups will be required when attempting to rebuild that master node whether that node is to be hosted on the original hardware or on a new system

When recovering to a new system that system must not be in used as a Ops Center Protector.

Where cloud storage has been used for backups this new system must have an internet connection and be able to communicate with required cloud storage system.

Where a repository hosted (proxied) by a client node has been used for backups you will need to know the directory path for that repository. This is the directory path that was specified when the repository was created.

In the event that the Master node is lost or otherwise permanently unavailable all settings for that master node may be recovered to a new master as described below. Note that if testing this procedure while the original master node is still available that master **MUST** be stopped prior to the commencement of the following actions.

Procedure

1. **Create a new Master Node**
   - Install Protector as a recovery master node on a new system (see How to install/ upgrade Protector on Windows and Linux or AIX (on page 235) . Where the system has previously had Ops Center Protector installed ensure it is completely uninstalled before re-installing (i.e. ensure the install directory has been removed)
   - Give the new master a temporary name (e.g. Recovery_Master) This will be overwritten later in this procedure with the original master node name.
   - Note the IP address of the new master.
2. **Preparing a Repository Storage Proxy Node for use by the new Master Node**

   Note: This section can be ignored if the Ops Center Protector Master Node backups are not stored on a repository node:

   On the client node that is proxy for the repository containing the protector settings back ups:
   - Stop the protector services. Refer How to stop the services (on page 263).
   - To prepare the repository containing the Ops Center Master node backups for use by the 'recovery master' run the command.

     ```
     setconfig --recoverymaster <ip_address> --repository <repository_path>
     ```

     where <ip_address> is the address of the new master and <repository_path> is the base directory of the repository containing the backup to be recovered. When successful the following response will be issues:

     `Repository`located at <repository_path> prepared for use by a recovery master
   - Start the protector services. Refer How to start the services (on page 263).
   - Add Client and Repository nodes to the new recovery master Node, on the UI Nodes screen authorize the client node when it appears. Refer How to authorize a node (on page 262).
   - Create new node. Refer How to add a node (on page 261).
   - Select 'Storage' node type "Repository". Refer Node Type Wizard (on page 539).
   - Specify a temporary name to be used for the repository node.
   - Select the client node as the "Proxy Node"
   - For creation mode select "Import existing repository", giving the path of the repository on the client when asked to "Select existing repository directory".
   - When the node has been created the imported repository will appear in the nodes list and its will appear as 'online' once it has automatically mounted,
3. **Preparing HCP/Cloud Storage for use by the new Master Node**

   This section can be ignored if the Ops Center Protector Master Node backups are not stored on a cloud node:

The new recovery master will be used as the proxy for the storage node. This storage node should be created as follows:

- Create new node. Refer How to add a node (on page 261).
- Select 'Storage' node type and then select the storage type that contains the Ops Center Protector Master back to be recovered. Refer Node Type Wizard (on page 539).
- Specify a temporary name to be used for the node.
- For creation mode select "Import an existing node".
- For proxy node select the recovery master itself .
- For the metadata cache directory specify a temporary directory. Be sure that this directory does not already exist.
- On the configure node screen specify the details of the HCP/Cloud storage.

When the node has been created the imported storage node will appear in the nodes list and its should appear as 'online' once it has automatically mounted

4. **Recover the Original Master Settings**
   - On the restore screen (see How to view available backups (on page 279)) hit the 'Search' button. Unfortunately, it will be difficult to filter the results (other than by specifying the temporary name used for the storage node) because until the master settings have been recovered the node names, policy names, data flow names, etc, used for filtering are not known.
   - Select the most recent master node settings backup from the recovery points listed and hit the 'restore' tool icon (see Repository Snapshot Details (Restore) – Ops Center (on page 846).
   - Select "Ops Center Protector Master" from the list of applications to restore, and hit 'Next'
   - Select the current (master) node as the destination node and hit the 'Finish' button.
   - The progress of the restore screen may be monitored from the 'Jobs' screen (see How to view and control running jobs (on page 258)) or 'Logs' screen (see How to view logs (on page 259)).

- When the restore completes it will be shown as 'Completed with warnings'. This is because the process of changing the master node to use the recovered settings cannot be performed while the master node is active. A message will be sent to the Log screen as the job completes detailing the Windows DOS command/Linux Shell command required to complete the recovery. Take a copy of this command string. For e.g.



Figure 49 Log Manager - Attachments

- Stop the protector services. Refer How to stop the services (on page 263).
- From a Windows DOS command prompt or Linux shell command run the command noted in the previous step to restore the original master node settings.
- Start the protector services. Refer How to start the services (on page 263).

  Where using a repository node the proxy node for that repository will appear on the UI as "offline". To return it to use by the recovered master:

- Stop the protector services. Refer How to stop the services (on page 263).
- To return the repository used for the recovery for use by the 'recovered master' run the command within a Windows DOS/Unix shell window change directory to the Protector 'bin' directory

```
setconfig --recoveredmaster <ip_address> --repository <repository_path>
```

  where <ip_address> is the address of the recovered master and <repository_path> is the base directory of the repository. When successful the following response will be issues:

  ```
  Repositorylocated at <repository_path> returned for use by the recovered master
  ```

- Start the protector services. Refer How to start the services (on page 263)
- The master node UI should now be showing the client node and the repository node as 'online'.

Note: Although all client nodes will still have copies of their active rules the recovered master does not know what rules are active on each client, and so the UI Monitor screen shows no data flows. All required data flows should now be compiled and activated refer How to activate a data flow (on page 232).

How to revert to the original Master

It is possible, following successful recovery of the master node settings to a new 'Recovered Master', to subsequently revert back to the original master system.

If any changes have been made to the master settings following recovery then a fresh backup of the Ops Center Protector Master Node settings should be performed and then the same procedure should be followed to restore this backup to the original system as was used to create the recovered master.

If no changes have been made to the master settings following recovery then the following steps may be followed to revert to using the original master node without the need for any new backup or recovery:

- Stop the protector services (see How to stop the services (on page 263) on the recovered master.
- Ensure the protector services are not running on the original master .
- Within a Windows DOS/Unix shell window change directory to the Protector 'bin' directoryon the original master enter the CLI command `setconfig --forceprobe`
- Start the cofiohub service on the original master.

How to protect Ops Center applications

You can use Protector to back up all of your Ops Center products. You can then restore your backups if you must recover one or more of your Ops Center applications.

Backing up Ops Center applications

Before you begin

To ensure that backups are available when you must restore a product, always back up to remote storage. The remote storage can be a repository that is proxied by an HCP node, HCP for cloud scale, or Amazon S3. When using any of these remote storage options for backup, the repository must be hosted (proxied) by a client node, which can be a standard client node or the master node.

Note: You must use a generation 2 repository.

Procedure

1. Create a policy. Refer to How to create a policy (on page 273).
   a. Add a classification type of **Ops Center**. Refer to How to add a classification to a policy (on page 274).

b. For the classification, select the node acting as the proxy, and then under **Include Applications** select either **Include all Ops Center applications** or click **Add** to select individual products. Refer to Ops Center Classification Wizard (on page 685) .

c. Add a backup operation to the policy and specify the schedule or RPO that you want.

2. Create a data flow. Refer to How to create a data flow (on page 228) .

   a. Drag the node acting as the proxy onto the data flow so that this node will be the source node.

   b. Select the policy that you created on the master and apply it to the master node. Refer to How to apply a policy to nodes on a data flow (on page 231).

   c. Drag the target storage node into the data flow and connect the master source node to the repository. Refer to How to connect nodes on a data flow (on page 231).

   d. Select the backup operation in the policy on the storage node. Refer to How to apply a policy to nodes on a data flow (on page 231).

   e. Activate the data flow. Refer to How to activate a data flow (on page 232).

   The backup will occur automatically based on the configured schedule or RPO selected within the policy. You can run a manual backup by selecting the data flow from within the Monitor screen. Refer to How to trigger an operation from an active data flow (on page 260).

How to restore Ops Center application backups

You must restore your Ops Center backup archives before you can recover your Ops Center products. Some products require that you restore to the same host, while others let you restore to a different host.

Before you begin

If you used cloud storage for your backups, the system must have an internet connection and be able to communicate with required cloud storage system. You also need the directory path of the repository on the client node that acted as a proxy. This is the directory path you specified when creating the repository.

This task describes the steps to follow to restore your data:

Procedure

1. Identify the destination where you want to restore the data set.

   Depending on the scenario, you may want to restore data to the original location or a different location.

2. Ensure that the restore location is prepared to receive the restored data set.

   Applications that access the data must be placed in the correct state and sufficient resources must be available to receive the data set.

   Caution: The process of restoring data can result in the destruction of some original data that exists on the restore target. Ensure that any critical data is copied to a safe location or is included in the data set being restored.

3. Depending on the scenario, it may be necessary to suspend any backup policies currently active on the location where the data set is being restored.

   Data flows can be deactivated from either the Monitor Inventory (on page 520) or Data Flows Inventory (on page 355).
4. Locate the data set to restore.

   The Restore Inventory (on page 763) and Storage Inventory (on page 858) track all backups maintained by Protector.
5. Select the restore method and start the restore process.

   It may be possible to restore the data to the original location, a different location, or mount a volume.

   The Restore Inventory (on page 763) and Storage Inventory (on page 858) provide functions to restore snapshots.
6. After the restore process is complete, you can recover your applications by following the steps described in the *Hitachi Ops Center Installation and Configuration Guide*.
7. After you recover and restart your applications, resume any existing backup policies for the restored data set. If a new restore location was chosen, you must create new backup policies. Data flows can be reactivated using the Data Flows Inventory (on page 355).

Host Based and HCP Workflows

This section describes high level workflows for repository based data protection including tiering to HCP and backup to cloud. These workflows focus on basic data protection scenarios involving files located on an OS Host. For guidance on protecting supported applications, refer to the relevant Protector Application Guide listed in Related documents (on page 17).

If you encounter problems, please refer to the following:

- Troubleshooting General (on page 969)
- Troubleshooting HCP (on page 979)
- Troubleshooting VMware (on page 985)

How to batch backup a file system path to a repository

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the source node where the file system path resides.
- The Protector Client software has been installed on the destination node where the Repository will reside.
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).

This task describes the steps to follow when protecting data that resides on a file system, to a repository, using batch mode backup. The data flow and policy are as follows:



Figure 50 Batch Data Flow

Table 5 Path Backup Policy

Classification Type	Parameters	Value
Path	Include	C:\testdata

Operation Type	Parameters	Value	Assigned Nodes
Backup	RPO	10 mins	Repository
	Retention	1 hour	
	Run Options	Run on RPO	

Procedure

1. Locate the source node in the Nodes Inventory (on page 536) and check that it is authorized and online. This node is where the production data to be backed up resides.

   For a file system backup using a *Path* classification, a basic *OS Host* node is required. It is not necessary to create the source node in this case since all Protector client nodes default to this type when installed. See How to authorize a node (on page 262).
2. Locate the destination node in the Nodes Inventory (on page 536) and check that it is authorized and online.

   This node is where the repository will be hosted and is identified as the **Proxy Node** when creating the repository in the next step.
3. Create a new Repository node using the Repository Storage Node Wizard (on page 632) and check that it is authorized and online.

   The *Repository* node type is grouped under **Storage** in the Node Type Wizard (on page 539). You can direct data from multiple nodes to a single repository so there is no need to create a new repository if a suitable one already exists. See How to add a node (on page 261), and How to authorize a node (on page 262).
4. Define a policy as shown in the table above using the Policy Wizard (on page 671), Path Classification Wizard (on page 713) and Backup Operation Wizard (on page 718).

   The *Path* classification is grouped under **Physical** in the Policy Wizard (on page 671). See How to create a policy (on page 273).
5. Draw a data flow as shown in the figure above, that shows the *OS Host* source node connected to the *Repository* destination node via a *Batch* mover, using the Data Flow Wizard (on page 361).

   See How to create a data flow (on page 228).
6. Assign the *Path-Backup* policy to the *OS Host* source node and the *Backup* operation to the *Repository* destination node on the data flow.

   Select the *Standard Store Template* when assigning the operation to the repository. See How to apply a policy to nodes on a data flow (on page 231).
7. Compile and activate the data flow, checking carefully that there are no errors or warnings.

   See How to activate a data flow (on page 232).
8. Locate the active data flow in the Monitor Inventory (on page 520) and open its Monitor Details (on page 521).

   The policy will be invoked automatically to create an initial backup and then repeatedly according to the RPO specified in the policy. The policy can also be manually triggered from the source node in the monitor data flow. See How to trigger an operation from an active data flow (on page 260).
9. Watch the active data flow via the Monitor Details (on page 521) to ensure the policy is operating as expected.
   For a healthy data flow you may periodically see:

   - An animated resynchronization icon appear above the batch mover each time the RPO is reached.
   - Transient **Node Status** icons appear over nodes and associated information messages displayed to the right of the data flow area.

- Backup jobs appearing in the **Jobs** area below the data flow that cycle through stages and ending in *Progress - Completed*.
- Information messages appearing in the **Logs** area below the data flow indicating rules activation, replication and resynchronization events.

For a problematic data flow you may see:

- Permanent **Node Status** icons appear over nodes and associated warning messages displayed to the right of the data flow area.
- Backup jobs appearing in the **Jobs** area below the data flow that cycle through stages and terminating in *Progress - Failed*.
- Warning and error messages appearing in the **Logs** area below the data flow indicating failed events.

10. Review the status of the *Repository* via the relevant Generation 1 Repository Details (on page 921) and the stores via the relevant Gen1 Repository Store Details (on page 915), to ensure backup snapshots are being created.

Repositories require ongoing surveillance to ensure that they are operating correctly and sufficient resources are available to store your data securely. See How to view the status of a repository (on page 281).

New snapshots will appear in the Gen1 Repository Store Details (on page 915) periodically as dictated by the *RPO* of the policy. Old snapshots will be removed periodically as dictated by the *Retention Period* of the policy. The retention period of individual snapshots can be modified here if required.

How to tier a file system path to HCP via a repository

Before you begin

Note: Tiering repository data to HCP is only available when using a generation 1 repository node and a generation 1 HCP node.

Note: In order to tier data from a repository store to HCP, a new tiering data flow, using a new, unpopulated repository store is required. Adding a tiering mover and HCP node to an existing repository data flow will not work.

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the source node where the file system path resides.
- The Protector Client software has been installed on the destination node where the Repository will reside.
- HCP generation 1 node has been set up as per the Protector requirements and prerequisites. Refer to Generation 1 Hitachi Content Platform prerequisites (on page 24).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).

This task describes the steps to follow when tiering data that resides on a file system, to HCP. These files are first ingested by a repository, using batch mode backup, and then immediately moved from the repository to an associated namespace within a tenant on HCP. Once data is tiered to HCP, the repository retains only the metadata describing the backed up files system. The data tiered to HCP can be located and restored by following the same workflow as that used for restoring repository data (see How to restore a repository snapshot of a file system path (on page 147)). The data flow and policy are as follows:



Figure 51 Tiering Data Flow

If you need frequent backups available in a local repository as well as long term backups on HCP, then implement a repo-to-repo data flow (see How to backup an onsite repository to an offsite repository (on page 149)) and tier to HCP from the second repository. The first repository will hold frequent local backups while the second manages long term retention on HCP.

Table 6 Path Backup Policy

Classification Type	Parameters	Value
Path	Include	C:\testdata

Operation Type	Parameters	Value	Assigned Nodes
Backup	RPO	1 Week	Repository (this controls the retention of data on HCP)
	Retention	5 Years	
	Run Options	Run on RPO	
Tier	None	N/A	Hitachi Content Platform

Procedure

1. Locate the source node in the Nodes Inventory (on page 536) and check that it is authorized and online. This node is where the production data to be backed up resides.

   For a file system backup using a *Path* classification, a basic *OS Host* node is required. It is not necessary to create the source node in this case since all Protector client nodes default to this type when installed. See How to authorize a node (on page 262).
2. Locate the intermediate node in the Nodes Inventory (on page 536) and check that it is authorized and online.

   This node is where the repository will be hosted and is identified as the **Proxy Node** when creating the repository in the next step.
3. Create a new generation 1 Repository node using the Repository Storage Node Wizard (on page 632) and check that it is authorized and online.

   The *Repository* node type is grouped under **Storage** in the Node Type Wizard (on page 539). You can direct data from multiple nodes to a single repository so there is no need to create a new repository if a suitable one already exists. See How to add a node (on page 261), and How to authorize a node (on page 262).
4. Create a new generation 1 Hitachi Content Platform node using the Hitachi Content Platform Storage Node Wizard (on page 600) and check that it is authorized and online.

   The Hitachi Content Platform node type is grouped under **Storage** in the Node Type Wizard (on page 539). You can direct data from multiple repository stores to a single HCP node (each repository store maps to a separate namespace within the HCP tenant), so there is no need to create a new HCP node if a suitable one already exists.

   Note: Each HCP node in Protector represents a single tenant, so if you need to strictly segregate repository data then create separate HCP nodes for each tenant and consider placing them in separate RBAC resource groups.

   See How to add a node (on page 261), and How to authorize a node (on page 262).

5. Define a policy as shown in the table above using the Policy Wizard (on page 671), Path Classification Wizard (on page 713), Backup Operation Wizard (on page 718) and Tier Operation Wizard (on page 726).

   The *Path* classification is grouped under **Physical** in the Policy Wizard (on page 671). See How to create a policy (on page 273).

6. Draw a data flow as shown in the figure above, using the Data Flow Wizard (on page 361), that shows the *OS Host* source node connected to the *Repository* intermediate node via a *Batch* mover, then to the Hitachi Content Platform destination node via a second *Batch* mover.

   See How to create a data flow (on page 228).

7. Assign the *Path-Backup-Tier* policy to the *OS Host* source node, the *Backup* operation to the *Repository* node and the *Tier* operation to the HCP node on the data flow.

   Select the *Standard Store Template* when assigning the operation to the repository. There is no value in selecting a template that performs source-side or repository-side deduplication in this situation. See How to apply a policy to nodes on a data flow (on page 231).

8. Compile and activate the data flow, checking carefully that there are no errors or warnings.

   See How to activate a data flow (on page 232).

   Note: Do not deactivate tiering data flows unless they are longer required. Subsequent reactivation will force the source and repository to undergo resynchronization, leading to all files being re-tiered to HCP.

9. Locate the active data flow in the Monitor Inventory (on page 520) and open its Monitor Details (on page 521).

   The policy will be invoked automatically to create an initial backup and then repeatedly according to the RPO specified in the policy. The policy can also be manually triggered from the source node in the monitor data flow. See How to trigger an operation from an active data flow (on page 260).

10. Watch the active data flow via the Monitor Details (on page 521) to ensure the policy is operating as expected.
    For a healthy data flow you may periodically see:

    - An animated resynchronization icon appear above the batch mover into the repository each time the RPO is reached.
    - An animated tiering icon appear above the batch mover into the HCP node each time the repository tiers data.
    - **Repository Statistics - Queues - Tier** values changing, indicating objects queued and actively being tiered to HCP.

      Tip: Check the tier queue if RPO is not being met for tiered data flows.

    - Transient **Node Status** icons appear over nodes and associated information messages displayed to the right of the data flow area.
    - **Network/Cache Utilization** fluctuations within normal limits if large amounts of data are being backed up to the repository.

- Backup jobs appearing in the **Jobs** area below the data flow that cycle through stages and ending in *Progress - Completed*. Note there is no *Tiering* job since this process takes place on an ad hoc basis.
- Information messages appearing in the **Logs** area below the data flow indicating rules activation, HCP namespace creation, resynchronization and ingestion throttling events.

For a problematic data flow you may see:

- Permanent **Node Status** icons appear over nodes and associated warning messages displayed to the right of the data flow area.
- **Local/Remote Memory Cache** constantly at excessively high levels if large amounts of data are being backed up, indicating data transfer issues.
- Backup jobs appearing in the **Jobs** area below the data flow that cycle through stages and terminating in *Progress - Failed*.
- Warning and error messages appearing in the **Logs** area below the data flow indicating failed events.

11. Review the status of the *Repository* via the relevant Generation 1 Repository Details (on page 921) and the stores via the relevant Gen1 Repository Store Details (on page 915), to ensure backup snapshots are being created. Also monitor the health of HCP via its **Tenant Management Console**, especially **Namespaces - Usage**.

    The UUID of the repository store (used to name the corresponding HCP namespace) can be found in the Gen1 Repository Store Details (on page 915).

    Repositories require ongoing surveillance to ensure that they are operating correctly and sufficient resources are available to store your data securely. See How to view the status of a repository (on page 281). A repository store that has been completely tiered to HCP will report a zero size.

    The space reported by HCP to tier a repository may appear much larger than that reported by the source node's file system. HCP allocation size is 8KB minimum per object, and each file requires at least 2 HCP objects (one per stream). When tiering many small files, the size reported HCP usage will appear larger than expected. Add to this the fact that HCP will create 2 or more replicas depending on the DPL setting.

    New snapshots will appear in the Gen1 Repository Store Details (on page 915) periodically as dictated by the *RPO* of the policy. Old snapshots will be removed periodically as dictated by the *Retention Period* of the policy. The retention period of individual snapshots can be modified here if required.

How to restore a repository snapshot of a file system path

Before you begin

It is assumed that a file system path policy that creates repository snapshots has been implemented and that at least one snapshot has been created in the designated repository store. See How to batch backup a file system path to a repository (on page 141) for an example of how to do this.

This task describes the steps to follow when restoring a file system path snapshot from a repository store to a node other than the one from which the backup originated, as shown:

Procedure

1. Identify the destination where the data set is to be restored. Here we will restore the snapshot to a different machine to the one from which the data originated.

   Depending on the scenario, you can restore data to its original node and directory, its original node in a different directory or to a different node entirely.
2. Ensure that the restore location is prepared to receive the restored data set by locating the node in the Nodes Inventory (on page 536) and checking it is authorized and online.

   The restore location must have Protector Client software installed. We will assume that no applications are accessing the restore location since the restored data doesn't yet exist.
3. In this example we will assume that there are no backup policies currently active on the location where the data set is being restored, so there is no need to suspend any data flows. The existing backup policy can continue to run while we perform our restore.
4. Locate the data set to be restored by navigating to the Repository Snapshot Details (Storage) - File System (on page 917) for the repository store snapshot in question.

   See How to view the contents of a snapshot in a repository store (on page 281).
5. Check that the target restore location (identified in the previous steps) has enough free space to accommodate the restored snapshot.

   The **Logical Size** of the snapshot is shown in the **Analysis Details** area of the Repository Snapshot Details (Storage) - File System (on page 917). It may be necessary to click **Analyze** if these details have not yet been evaluated.
6. Click **Restore Snapshot** to open the Restore Repository Snapshot Wizard - File System (on page 826).

   > **Caution:** The process of restoring data may result in the overwriting some of the original data that exists on the restore location.
   >
   > Ensure that any critical data is copied to a safe location or is included in the data set being restored.
   >
   > The Restore Repository Snapshot Wizard - File System (on page 826) - **Select restore options** page provides numerous **File Name Collision Policy** options to help manage potential file overwrite situations.

   In this example we will restore to a different destination node but use the original file paths. No routing is required since the Repository and Restore target are connected to the same LAN.

   a. From the Restore Repository Snapshot Wizard - File System (on page 826), choose whether to restore the **Entire Snapshot** or a **User Selection** of files. Click **Next**.
   b. If restoring a **User Selection**, select the files and folders to be restored. Click **Next**.
   c. Set the **Destination Node** to the one identified in the previous steps above.
   d. Set **Restore To** to **Original Location** so that the files are placed on the same path as the originals.

e. Set **File Name Collision Policy** to **Rename any colliding files** so that any existing files of the same name are preserved.

   This has no effect when initially creating the restored files in a new location, but if the restore is repeated then it will preserve any existing files from previous restore jobs.

f. Review the restore options carefully to ensure that everything has been specified correctly, then click **Finish** to initiate the restore job.
   A `Processing` message will appear briefly, then the wizard will close and the Jobs Inventory (on page 493) will be displayed. A new *Restore Job* will appear at the top of the Jobs list, with the *Progress* entry initially indicating processing and finally indicating successful completion.

7. Once the restore process is complete, further steps may be needed to fix-up the data set before using it. In this example we will assume that no additional work is required other than inspecting the restored data on the target machine.

   The amount of fix-up work required depends on the applications accessing the restored data.

   Note: This example restores data created using a *Path* classification. If you are backing up one of the application types directly supported by Protector, then you should use one of the *Application* classifications and refer to the appropriate Application Guide listed in Related documents (on page 17)).

8. Restart any applications that access the restored data.

   For supported applications, these additional steps are described in the appropriate Application Guide (see Related documents (on page 17)). For other applications, consult the vendor's documentation.

9. Resume any backup policies for the restored data set. If you have restored data to a new location for repurposing (test and development work for example), you should consider if it is necessary to implement a new backup policy to protect this new instance.

   Data flows can be reactivated via the Data Flows Inventory (on page 355).

How to backup an onsite repository to an offsite repository

Before you begin

Note: When performing repository to repository backups repository generations can't be mixed.

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the source node where the file system path resides.
- The Protector Client software has been installed on the destination nodes where the repositories will reside.
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).

This task describes the steps to follow when protecting data backed up to a primary repository by copying it to a secondary repository. The secondary repository would normally be located offsite to provide protection from local catastrophic failures. The data flow and policy are as follows:



Figure 52 Repository to Repository Data Flow

Table 7 Repository Backup Policy

Classification Type	Parameters	Value
Path	Include	C:\testdata

Operation Type	Parameters	Value	Assigned Nodes
Onsite Daily (Backup)	RPO	1 day	Onsite Repository
	Retention	1 week	
	Run Options	Run on RPO	
Onsite Weekly (Backup)	RPO	1 week	Onsite Repository
	Retention	6 months	
	Run Options	Run on RPO	
Offsite Daily (Backup)	RPO	N/A	Offsite Repository

Operation Type	Parameters	Value	Assigned Nodes
	Retention	1 week	
	Run Options	Run on completion of operation *Onsite Daily*	
Offsite Weekly (Backup)	RPO	N/A	Offsite Repository
	Retention	6 months	
	Run Options	Run on completion of operation *Onsite Weekly*	

Procedure

1. Locate the source node in the Nodes Inventory (on page 536) and check that it is authorized and online. This node is where the production data to be backed up resides.

   See How to authorize a node (on page 262).
2. Locate the destination nodes in the Nodes Inventory (on page 536) and check that they are authorized and online.

   These nodes are where the onsite and offsite repositories will be hosted and are identified as the **Proxy Node** when creating the repositories in the next step.
3. Create two new Repository nodes using the Repository Storage Node Wizard (on page 632) and check that they are authorized and online.

   You can direct data from multiple nodes to a single repository so there is no need to create new repositories if suitable ones already exists. See How to add a node (on page 261), and How to authorize a node (on page 262).
4. Define a policy as shown in the table above using the Policy Wizard (on page 671), Path Classification Wizard (on page 713) and Backup Operation Wizard (on page 718).

   Note: Both onsite operations are triggered according to their RPOs; the offsite operations are then triggered on completion of their respective onsite operations. The RPOs for the offsite operations do not need to be specified; they have no effect.

   See How to create a policy (on page 273).
5. Draw a cascaded data flow as shown in the figure above, that shows the *OS Host* source node connected to the *Onsite Repository* and then to the *Offsite Repository* destination nodes via *Batch* movers, using the Data Flow Wizard (on page 361).

   See How to create a data flow (on page 228).

6. Assign the *Repository-Backup* policy to the *OS Host* source node, both *Onsite Backup* operations to the *Onsite Repository* and both *Offsite Backup* operations to the *Offsite Repository* destination nodes on the data flow.

   If using generation 1 repositories select the *Standard Store Template* when assigning the operation to the repositories. See How to apply a policy to nodes on a data flow (on page 231).
7. Compile and activate the data flow, checking carefully that there are no errors or warnings.

   See How to activate a data flow (on page 232).
8. Locate the active data flow in the Monitor Inventory (on page 520) and open its Monitor Details (on page 521).

   The policy will be invoked automatically to create an initial backup and then repeatedly according to the RPO specified in the policy. The policy can also be manually triggered from the source node in the monitor data flow. See How to trigger an operation from an active data flow (on page 260).
9. Watch the active data flow via the Monitor Details (on page 521) to ensure the policy is operating as expected.

   Note: If the onsite repository contains a large amount of backup data and/or the network connection between the onsite and offsite repositories has limited bandwidth, the process of synchronizing the two repositories can be expedited by seeding the offsite repository as described in How to seed an offsite repository (on page 152).
10. Review the status of the *Repositories* via the relevant Generation 1 Repository Details (on page 921) and the *Stores* via the relevant Gen1 Repository Store Details (on page 915), to ensure backup snapshots are being created.
    New snapshots will appear in the Gen1 Repository Store Details (on page 915) periodically as dictated by the *RPO* of the policy. Old snapshots will be removed periodically as dictated by the *Retention Period* of the policy. The retention period of individual snapshots can be modified here if required.

How to seed an offsite repository

Before you begin

It is assumed that you have an onsite repository containing a large amount of backup data, and you want to create a secondary offsite backup of this data. You may have attempted to set up an offsite data flow as described in How to backup an onsite repository to an offsite repository (on page 149), but data volume and/or bandwith restrictions between sites are preventing the initial synchronization from completing in an acceptable timeframe.

Ensure that you have a removable disk available with enough spare capacity to hold the onsite repository's data files. This disk will be required to physically transport the data to the offsite location.

The initial synchronization can be completed more rapidly by seeding the offsite repository store with data from the onsite store. Once seeded, the offsite repository may request additional, lower volume, differential updates to ensure the two sites are completely synchronized. The process, known as seeding is performed as follows:

Procedure

1. In the Monitor Details (on page 521) for the relevant data flow, resynchronise the onsite repository by selecting the source nodes feeding into it, clicking **Trigger Operation** then selecting the operations that backup to the onsite repository.
2. Wait for the onsite repository to become synchronised with the source machine(s) at the primary site by confirming that the corresponding backup job has completed.
3. For both the onsite and offsite repository, note down the **Mount Directory** displayed in the **Configuration** area of the Generation 1 Repository Details (on page 921).
4. For both the onsite and offsite repository, unmount the repository by clicking **Unmount**.

   The *Offline* status icon will appear on the respective repository tiles in the Storage Inventory (on page 858) when the stores have unmounted.
5. Using Windows File Explorer, navigate to the repository root directory for the onsite repository (noted above) and copy the folder to a removable disk, then safely eject the disk from the machine.
6. For the onsite repository only, navigate to the Generation 1 Repository Details (on page 921) and mount the repository by clicking **Mount**.

   The *Online* status icon will appear on the onsite repository tile in the Storage Inventory (on page 858) when the stores have mounted. The onsite repository will now resume backup of your source nodes according to the onsite policies on force.
7. Take the removable disk containing the copy of the onsite repository to the remote site, insert it into the machine hosting the offsite repository and replace the files in the directory of the offsite repository (noted above) with those on the removable disk.

   Physically transferring the data between the onsite and offsite nodes circumvents the bandwidth restrictions of the network between the two sites.
8. For the offsite repository, navigate to the Generation 1 Repository Details (on page 921) and mount the repository by clicking **Mount**.

   The *Online* status icon will appear on the offsite repository tile in the Storage Inventory (on page 858) when the stores have mounted.
9. In the Monitor Details (on page 521) for the relevant data flow, resynchronise the repositories by selecting the onsite repository, clicking **Trigger Operation** then selecting the operations that backup to the offsite repository.

   Because the seeding process has been performed, only a minimal amount of data transfer will be required between the two sites.

Hitachi Block Workflows

This section describes high level workflows for Hitachi Block based data protection. These workflows focus on basic data protection scenarios involving primary and secondary LDEVs located on block storage devices. For guidance on protecting supported application data located on a block storage device refer to the relevant Protector Application Guide listed in Related documents (on page 17).

If you encounter problems, please refer to the following:

- Troubleshooting General (on page 969)
- Troubleshooting Hitachi Block (on page 974)
- Troubleshooting VMware (on page 985)

How to snapshot an Hitachi Block LDEV with Thin Image

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the source node that will act as a proxy for the Hitachi Block storage device. Note that for a Thin Image snapshot, the source and destination LDEVs are located on the same device.
- The storage device has been set up as per the Protector requirements and prerequisites. Refer to Hitachi Block prerequisites (on page 21).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).

This task describes the steps to follow when protecting an LDEV allocated from a Hitachi Block storage device. This is useful when Protector has no way of interacting with the application or OS that is using the LDEV. The snapshot will be crash consistent, because Protector is not able to orchestrate the snapshot operation in conjunction with applications using the LDEV. Thin Image hardware snapshots of the P-VOL are created as S-VOLs residing within the same storage device. For more information, refer to About Thin Image and Thin Image Advanced differential and refreshed snapshots (on page 66). The data flow and policy are as follows:



Figure 53 Hitachi Block Snapshot Data Flow

Table 8 Hitachi Block Snapshot Policy

Classification Type	Parameters	Value
Hitachi Block	Specify additional selections	Selected
	Logical Devices	10323/10

Operation Type	Parameters	Value	Assigned Nodes
Snapshot	Mode	Hardware	Hitachi Block Device
	Hardware Type	Hitachi Block	
	RPO	10 mins	
	Retention	1 hour	
	Run Options	Run on RPO	
	Quiesce...	Not selected	

Procedure

1. Locate the node in the Nodes Inventory (on page 536) that will control the Hitachi Block Device via a CMD (Command Device) interface and check that it is authorized and online.

   This node is used by Protector to orchestrate snapshot creation and is identified as the **Proxy Node** when creating the Hitachi Block Device node in the next step. This node is known as an ISM node. The ISM node does not appear in the data flow.

2. Create a new Hitachi Block Device node (unless one already exists) using the Hitachi Block Device Node Wizard (on page 583) and check that it is authorized and online. This node is where the production LDEV to be snapshotted is located.

   For a snapshot using an Hitachi *Block* classification, an Hitachi *Block Device* node is required.

   The Hitachi *Block Device* node type is grouped under **Storage** in the Node Type Wizard (on page 539). See How to add a node (on page 261) and How to authorize a node (on page 262).
3. Define a policy as shown in the table above using the Policy Wizard (on page 671), Hitachi Block Classification Wizard (on page 709) and Snapshot Operation Wizard (on page 724).

   The Hitachi *Block* classification is grouped under **Physical** classifications. See How to create a policy (on page 273).
4. Draw a data flow as shown in the figure above, that shows only the Hitachi *Block Device* source node, using the Data Flow Wizard (on page 361).

   At this stage the snapshot icon is not shown. See How to create a data flow (on page 228).
5. Assign the *Snapshot* operation to the Hitachi *Block Device* source node. The *Block-Snapshot* policy will then be assigned automatically.

   See How to apply a policy to nodes on a data flow (on page 231).

   The Hitachi Block Snapshot Configuration Wizard (on page 384) is displayed.
6. Select the **Snapshot Pool** by selecting one of the available Thin Image or hybrid pools.
7. Leave the remaining **Advanced Options** at their default settings, then click **OK**.

   The snapshot icon is now shown superimposed over the source node.
8. Compile and activate the data flow, checking carefully that there are no errors or warnings.

   See How to activate a data flow (on page 232).

   Note: If the **Quiesce configured applications before backup** option was not deselected in the Snapshot Operation Wizard (on page 724), then a compiler warning message will be generated because Protector will not be able to quiesce applications using the LDEV.
9. Locate the active data flow in the Monitor Inventory (on page 520) and open its Monitor Details (on page 521) page.

   The policy will be invoked repeatedly according to the RPO specified. The policy can also be manually triggered from the source node in the monitor data flow. You may want to manually trigger to create an initial snapshot. See How to trigger an operation from an active data flow (on page 260).
10. Watch the active data flow via the Monitor Details (on page 521) to ensure the policy is operating as expected.

For a healthy data flow you will periodically see:

- Snapshot jobs appearing in the **Jobs** area below the data flow that cycle through stages and ending in *Progress - Completed*.
- Information messages appearing in the **Logs** area below the data flow indicating rules activation, storage handler and sequencer events.
- Attachments to storage handler log events confirming which volumes are being snapshotted.

For a problematic data flow you may see:

- Permanent **Node Status** icons appear over nodes and associated warning messages displayed to the right of the data flow area.
- Backup jobs appearing in the **Jobs** area below the data flow that cycle through stages and terminating in *Progress - Failed*.
- Warning and error messages appearing in the **Logs** area below the data flow indicating failed events.

11. Review the status of the Hitachi *Block Device* via the relevant Hitachi Block Device Details (on page 859) and snapshots via the Hitachi Block Snapshots Inventory (on page 869), to ensure snapshots are being created.

    Hitachi Block Devices require ongoing surveillance to ensure that they are operating correctly and sufficient resources are available to store your data securely. See How to view the status of a Hitachi Block storage device (on page 282). The retention period of individual snapshots can be modified here if required.

    New snapshots will appear in the Hitachi Block Snapshots Inventory (on page 869) periodically as dictated by the *RPO* of the policy. Old snapshots will be removed periodically as dictated by the *Retention Period* of the policy.

How to snapshot a file system with Thin Image

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the source node where the Hitachi Block LDEV is mounted. Note that the LDEV is actually located on the Hitachi Block storage device.
- The Protector Client software has been installed on the destination node that will act as a proxy for the Hitachi Block storage device. Note that for a Thin Image snapshot, the source and destination LDEVs are located on the same device.
- The storage device has been set up as per the Protector requirements and prerequisites. Refer to Hitachi Block prerequisites (on page 21).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).

This task describes the steps to follow when protecting data that resides on a file system created on an LDEV allocated from a Hitachi Block storage device. Thin Image hardware snapshots of the P-VOL are created as S-VOLs residing within the same storage device. For more information, refer to About Thin Image and Thin Image Advanced differential and refreshed snapshots (on page 66). The data flow and policy are as follows:



Figure 54 Hardware Snapshot Data Flow

Table 9 Path Snapshot Policy

Classification Type	Parameters	Value
Path	Include	E:\testdata (E: is where the Hitachi Block LDEV is mounted)

Operation Type	Parameters	Value	Assigned Nodes
Snapshot	Mode	Hardware	OS Host
	Hardware Type	Hitachi Block	
	RPO	10 mins	
	Retention	1 hour	
	Run Options	Run on RPO	

Procedure

1. Locate the source node in the Nodes Inventory (on page 536) and check that it is authorized and online. This node is where the production LDEV to be snapshotted is mounted.

   For a file system snapshot using a *Path* classification, a basic *OS Host* node is required. It is not necessary to create the source node in this case since all Protector client nodes default to this type when installed. See How to authorize a node (on page 262).

2. Locate the node in the Nodes Inventory (on page 536) that will control the Hitachi Block Device via a CMD (Command Device) interface and check that it is authorized and online.

   This node is used by Protector to orchestrate snapshot creation and is identified as the **Proxy Node** when creating the Hitachi Block Device node in the next step. This node is known as an ISM node. The ISM node does not appear in the data flow.
3. Create a new Hitachi Block Device node (unless one already exists) using the Hitachi Block Device Node Wizard (on page 583) and check that it is authorized and online.

   The Hitachi *Block Device* node type is grouped under **Storage** in the Node Type Wizard (on page 539). See How to add a node (on page 261) and How to authorize a node (on page 262). Note that this node does not appear in the snapshot data flow diagram, but is identified when assigning the snapshot policy.
4. Define a policy as shown in the table above using the Policy Wizard (on page 671), Path Classification Wizard (on page 713) and Snapshot Operation Wizard (on page 724).

   The *Path* classification is grouped under **Physical** classifications. See How to create a policy (on page 273).
5. Draw a data flow as shown in the figure above, that shows only the *OS Host* source node, using the Data Flow Wizard (on page 361).

   At this stage the snapshot icon is not shown. See How to create a data flow (on page 228).
6. Assign the *Snapshot* operation to the *OS Host* source node. The *Path-Snapshot* policy will then be assigned automatically.

   See How to apply a policy to nodes on a data flow (on page 231).

   The Hitachi Block Snapshot Configuration Wizard (on page 384) is displayed.
7. Select the **Snapshot Pool** by selecting one of the available Thin Image or hybrid pools.

   Caution:

   Filling a Thin Image pool to capacity will invalidate all snapshot data contained within that pool. All snapshots in the pool will have to be deleted before snapshotting can be resumed.
8. Leave the remaining **Advanced Options** at their default settings, then click **OK**.

   The snapshot icon is now shown superimposed over the source node.
9. Compile and activate the data flow, checking carefully that there are no errors or warnings.

   See How to activate a data flow (on page 232).
10. Locate the active data flow in the Monitor Inventory (on page 520) and open its Monitor Details (on page 521) page.

    The policy will be invoked repeatedly according to the RPO specified. The policy can also be manually triggered from the source node in the monitor data flow. An initial snapshot will be taken shortly after rules distribution has completed. See How to trigger an operation from an active data flow (on page 260).
11. Watch the active data flow via the Monitor Details (on page 521) to ensure the policy is operating as expected.

For a healthy data flow you will periodically see:

- Snapshot jobs appearing in the **Jobs** area below the data flow that cycle through stages and ending in *Progress - Completed*.
- Information messages appearing in the **Logs** area below the data flow indicating rules activation, storage handler and sequencer events.
- Attachments to storage handler log events confirming which volumes are being snapshotted.

For a problematic data flow you may see:

- Permanent **Node Status** icons appear over nodes and associated warning messages displayed to the right of the data flow area.
- Backup jobs appearing in the **Jobs** area below the data flow that cycle through stages and terminating in *Progress - Failed*.
- Warning and error messages appearing in the **Logs** area below the data flow indicating failed events.

12. Review the status of the Hitachi *Block Device* via the relevant Hitachi Block Device Details (on page 859) and snapshots via the Hitachi Block Snapshots Inventory (on page 869), to ensure snapshots are being created.

    Hitachi Block Devices require ongoing surveillance to ensure that they are operating correctly and sufficient resources are available to store your data securely. See How to view the status of a Hitachi Block storage device (on page 282). The retention period of individual snapshots can be modified here if required.

    New snapshots will appear in the Hitachi Block Snapshots Inventory (on page 869) periodically as dictated by the *RPO* of the policy. Old snapshots will be removed periodically as dictated by the *Retention Period* of the policy.

How to replicate an Hitachi Block LDEV with ShadowImage

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the destination node that will act as a proxy for the Hitachi Block storage device. Note that for a ShadowImage replication, the source and destination LDEVs are located on the same device.
- The storage device has been set up as per the Protector requirements and prerequisites. Refer to Hitachi Block prerequisites (on page 21).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).

This task describes the steps to follow when protecting an LDEV allocated from a Hitachi Block storage device. This is useful when Protector has no way of interacting with the application or OS that is using the LDEV. The replication will be crash consistent, because Protector is not able to orchestrate the replication operation in conjunction with applications using the LDEV. A ShadowImage hardware replication of the P-VOL is created as an S-VOL residing within the same storage device. For more information, refer to About ShadowImage replication (on page 70). The data flow and policy are as follows:



Figure 55 ShadowImage Replication Data Flow

Table 10 Hitachi Block Replication Policy

Classification Type	Parameters	Value
Hitachi Block	Specify additional selections	Selected
	Logical Devices	10323/10

Operation Type	Parameter	Value	Assigned Nodes
Replicate	Run Options	Run on Schedule (See below)	Hitachi Block Device (source), Hitachi Block Device (destination)
	Quiesce...	Not selected	

Schedule Item Type	Parameter	Value	Policy Operations
Trigger	Days	Select All	Replicate (See above)
	Weeks	Select All	
	Time	Scheduled Time	
	Start Time	15:00	
	Duration	00:00	

Procedure

1. Locate the node in the Nodes Inventory (on page 536) that will control the Hitachi Block Device via a CMD (Command Device) interface and check that it is authorized and online.

   This node is used by Protector to orchestrate replication and is identified as the **Proxy Node** when creating the Hitachi Block Device node in the next step. This node is known as an ISM node. The ISM node does not appear in the data flow.
2. Create a new Hitachi Block Device node (unless one already exists) using the Hitachi Block Device Node Wizard (on page 583) and check that it is authorized and online. This node is where the production LDEV to be replicated is located.

   For a replication using an Hitachi *Block* classification, an Hitachi *Block Device* node is required.

   The Hitachi *Block Device* node type is grouped under **Storage** in the Node Type Wizard (on page 539). See How to add a node (on page 261) and How to authorize a node (on page 262). This node appears in the replication data flow as both the source and the destination node.
3. Define a policy as shown in the table above using the Policy Wizard (on page 671). See How to create a policy (on page 273)
   a. Define an Hitachi Block classification using the Hitachi Block Classification Wizard (on page 709).

      The Hitachi *Block* classification is grouped under **Physical** in the Policy Wizard (on page 671).
   b. Define a *Replicate* operation using the Replicate Operation Wizard (on page 722).

      In this example, ShadowImage replication will run as a batch operation based on a *Trigger* schedule. Continuous ShadowImage could also be implemented by using a continuous mover on the dataflow.
   c. Define a *Trigger* schedule using the Schedule Wizard (on page 850); accessed by clicking on **Manage Schedules**.

      See How to create a schedule (on page 280).
4. Draw a data flow as shown in the figure above using the Data Flow Wizard (on page 361), that shows the Hitachi *Block Device* source node connected to the same Hitachi *Block Device* via a *Batch* mover.

   ShadowImage is an in-system replication technology, so the Hitachi *Block Device* node is where both the source (P)VOL) and destination (S)VOL) volumes are located. See How to create a data flow (on page 228).
5. Assign the *Block-Replicate* policy to the Hitachi *Block Device* source node.

   See How to apply a policy to nodes on a data flow (on page 231).
6. Assign the *Replicate* operation to the Hitachi *Block Device* destination node.
   The Hitachi Block Replication Configuration Wizard (on page 398) is displayed.
7. Set the replication type to **In System Clone**, then choose a **Pool** from one of the available *Dynamic Pools*. Leave the remaining parameters at their default settings and click **OK**.
8. Compile and activate the data flow, checking carefully that there are no errors or warnings.

   See How to activate a data flow (on page 232).

> **Note:** If the **Quiesce configured applications before backup** option was not deselected in the Replicate Operation Wizard (on page 722), then a compiler warning message will be generated because Protector will not be able to quiesce applications using the LDEV.

9. Locate the active data flow in the Monitor Inventory (on page 520) and open its Monitor Details (on page 521).

   The policy will be invoked automatically to create a replication according to the schedule specified in the policy. The policy can also be manually triggered from the source node in the monitor data flow.

   > **Note:** No replication will be created until it is first triggered manually or by the schedule.

   See How to trigger an operation from an active data flow (on page 260).

10. Watch the active data flow via the Monitor Details (on page 521) to ensure the policy is operating as expected.
    For a healthy data flow you will periodically see:

    - Replication jobs appearing in the **Jobs** area below the data flow that cycle through stages and ending in *Progress - Completed*.
    - Information messages appearing in the **Logs** area below the data flow indicating rules activation, storage handler and sequencer events.
    - Attachments to storage handler log events confirming which volumes are being replicated.

    For a problematic data flow you may see:

    - Permanent **Node Status** icons appear over nodes and associated warning messages displayed to the right of the data flow area.
    - Backup jobs appearing in the **Jobs** area below the data flow that cycle through stages and terminating in *Progress - Failed*.
    - Warning and error messages appearing in the **Logs** area below the data flow indicating failed events.

11. Review the status of the Hitachi *Block Device* via the relevant Hitachi Block Device Details (on page 859) and replications via the Hitachi Block Replications Inventory (on page 878), to ensure the replication is being created and refreshed.

    Hitachi Block Devices require ongoing surveillance to ensure that they are operating correctly and sufficient resources are available to store your data securely. See How to view the status of a Hitachi Block storage device (on page 282).

    The replication process can be paused and resumed from here if required.

    A new ShadowImage replication will appear in the Hitachi Block Replications Inventory (on page 878) and be updated periodically as dictated by the schedule for the policy operation. The previous replication will be overwritten upon each refresh.

How to teardown the S-VOLs of a replication removed from a data flow

Before you begin

Either:

- The corresponding replication operation must be removed from the dataflow where it is defined and that dataflow must be reactivated

or:

- The dataflow defining the replication operation must be permanently deactivated.

An Hitachi block replication that was defined within or adopted by Protector must be explicitly removed from the underlying hardware as follows:

Procedure

1. Locate the replication record (corresponding to the replication operation that has been removed) in the Hitachi Block Replications Inventory (on page 878).

   Replications that are eligible for teardown are marked with an ✖ in the top right corner of the tile.
2. Select the replication record to tear down, then click **Teardown** from the context menu.
3. The **Teardown** Hitachi **Block Replication Dialog** is displayed. If you are sure you want to proceed, type the word 'TEARDOWN', then click **Teardown**.
4. Go to the Jobs Inventory (on page 493) to ensure that a teardown job has been initiated and wait for it to complete.

   The replication entry is not removed from the replications inventory until the teardown operation is completed successfully. If the teardown is unsuccessful, review the Logs Inventory (on page 509) to find out why. The teardown operation must be re-initiated by the user once the problem is resolved.

How to reactivate a replication operation that has been accidently deactivated

An Hitachi block replication that was defined within or adopted by Protector may have been accidently deactivated. One of following three cases will apply:

Case 1: Replication operation has not been removed but the data flow is deactivated

Protector considers a replication to have been removed from a dataflow only if the link between the source and destination has been removed or the source and/or destination node has been removed.

Note: It is possible to edit the replication parameters as long as any changes are supported by the hardware for that replication type. Protector will still consider the replication instance to be the same.

Procedure

1. If none of the above have occurred then the data flow can simply be reactivated by the user via the Data Flows Inventory (on page 355).

   Because the replication has not been torn down, Protector will effectively re-adopt the corresponding replication from the storage hardware.

Case 2: Replication operation has been removed and the data flow has been reactivated

Protector considers a replication to have been removed from a dataflow if the link between the source and destination has been removed or the source and/or destination node has been removed.

Procedure

1. If this is the case, then the data flow must have a new replication operation added back in and then be reactivated by the user via the Data Flows Inventory (on page 355). Because Protector considers this new replication operation as an entirely new instance, the replication pair must be created from scratch on the storage array. The old replication becomes a static copy.

Case 3: The data flow containing the replication has been deleted

Protector considers the replication to have been removed.

Procedure

1. If this is the case, then a new data flow must be created containing a new replication operation and then be reactivated by the user via the Data Flows Inventory (on page 355).

   Because Protector considers this new replication operation as an entirely new instance, the replication pair must be created from scratch on the storage array. The old replication becomes a static copy.

How to create and use an Hitachi Block Host node

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the source node that will act as a proxy for the Hitachi Block storage device.
- The storage device has been set up as per the Protector requirements and prerequisites. Refer to Hitachi Block prerequisites (on page 21).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).

A block host node can be used in a data flow to represent a host machine that has LDEVs mounted on it that require protection. The block host can be used as a convenient alternative to specifying the hardware paths in the Hitachi Block Classification Wizard (on page 709). Typically the block host node is used to represent an application server that is not directly supported by a Protector classification.

Procedure

1. Identify the existing Hitachi Block Storage node where the LDEVs that require protection reside and ensure it is authorised and online.
2. Create a new Hitachi Block Host node using the Hitachi Block Host Node Wizard (on page 559). This node will represent the host where the production LDEVs to be protected are mounted.
    a. Specify a **Node Name** that reflects the name or purpose of the host it represents. This is where the LDEVs are mounted.
    b. Allocate the block host node to the same Access Control Resource Group as that of the block device node specified in the next step.
    c. Select the **Hitachi Block Device**. This is the storage device where the LDEVs reside.
    d. Optionally, specify the LDEVs mounted on the host that are to be protected. Normally you would specify these in the Hitachi Block Classification Wizard (on page 709), but it is often more logical to capture this information here.
3. Check that the newly created Hitachi Block Host node is authorised and online in the Nodes Inventory (on page 536).
4. Place the block host node on data flows in the same way that you would a block device source node.

    Note: You can only use a block host node as a source node. You cannot, for example, replicate to a block host node.

5. Create a policy that includes a snapshot and/or replicate operation, and an associated block classification using the Hitachi Block Classification Wizard (on page 709).
    a. Either: Select **Use Hitachi Block Host selections** to indicate that you want only the LDEVs specified in the Hitachi Block Host Node Wizard (on page 559) to be protected.
    b. Or: Select **Specify additional selections** to indicate that you want the LDEVs specified in the Hitachi Block Host Node Wizard (on page 559) to be protected in addition to any specified in the **Logical Devices** field of the classification.
6. Assign the snapshot or replicate operation to the block host node in the data flow. See How to apply a policy to nodes on a data flow (on page 231).
7. Activate the data flow as normal. See How to activate a data flow (on page 232).

How to replicate a file system with ShadowImage

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the source node where the Hitachi Block LDEV is mounted. Note that the LDEV is actually located on the Hitachi Block storage device.
- The Protector Client software has been installed on the destination node that will act as a proxy for the Hitachi Block storage device. Note that for a ShadowImage replication, the source and destination LDEVs are located on the same device.
- The storage device has been set up as per the Protector requirements and prerequisites. Refer to Hitachi Block prerequisites (on page 21).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).

This task describes the steps to follow when protecting data that resides on a file system created on an LDEV allocated from a Hitachi Block storage device. A ShadowImage hardware replication of the P-VOL is created as an S-VOL residing within the same storage device. For more information, refer to About ShadowImage replication (on page 70). The data flow and policy are as follows:



Figure 56 ShadowImage Replication Data Flow

Table 11 Path Replication Policy

Classification Type	Parameter	Value
Path	Include	E:\testdata (E: is where the Hitachi Block LDEV is mounted)

Operation Type	Parameter	Value	Assigned Nodes
Replicate	Run Options	Run on Schedule (See below)	OS Host, Hitachi Block Device

Schedule Item Type	Parameter	Value	Policy Operations
Trigger	Days	Select All	Replicate (See above)
	Weeks	Select All	
	Time	Scheduled Time	
	Start Time	15:00	
	Duration	00:00	

Procedure

1. Locate the source node in the Nodes Inventory (on page 536) and check that it is authorized and online. This node is where the production LDEV to be replicated is mounted.

   For a file system replication using a *Path* classification, a basic *OS Host* node is required. It is not necessary to create the source node in this case since all Protector client nodes default to this type when installed. See How to authorize a node (on page 262).
2. Locate the node in the Nodes Inventory (on page 536) that will control the Hitachi Block Device via a CMD (Command Device) interface and check that it is authorized and online.

   This node is used by Protector to orchestrate replication and is identified as the **Proxy Node** when creating the Hitachi Block Device node in the next step. This node is known as an ISM node. The ISM node does not appear in the data flow.
3. Create a new Hitachi Block Device node (unless one already exists) using the Hitachi Block Device Node Wizard (on page 583) and check that it is authorized and online.

   The Hitachi *Block Device* node type is grouped under **Storage** in the Node Type Wizard (on page 539). See How to add a node (on page 261) and How to authorize a node (on page 262). This node appears in the replication data flow as the destination node.
4. Define a policy as shown in the table above using the Policy Wizard (on page 671). See How to create a policy (on page 273)
   a. Define a Path classification using the Path Classification Wizard (on page 713).

      The a *Path* classification is grouped under **Physical** in the Policy Wizard (on page 671).
   b. Define a *Replicate* operation using the Replicate Operation Wizard (on page 722).

      In this example, ShadowImage replication will run as a batch operation based on a *Trigger* schedule. Continuous ShadowImage could also be implemented by using a continuous mover on the dataflow.
   c. Define a *Trigger* schedule using the Schedule Wizard (on page 850); accessed by clicking on **Manage Schedules**.

      See How to create a schedule (on page 280).
5. Draw a data flow as shown in the figure above using the Data Flow Wizard (on page 361), that shows the *OS Host* source node connected to the Hitachi *Block Device* via a *Batch* mover.

   ShadowImage is an in-system replication technology, so the Hitachi *Block Device* node is where both the source (P-VOL) and destination (S-VOL) volumes are located. See How to create a data flow (on page 228).
6. Assign the *Path-Replicate* policy to the *OS Host* source node.

   See How to apply a policy to nodes on a data flow (on page 231).
7. Assign the *Replicate* operation to the Hitachi *Block Device* node.
   The Hitachi Block Replication Configuration Wizard (on page 398) is displayed.
8. Set the replication type to **In System Clone**, then choose a **Pool** from one of the available *Dynamic Pools*. Leave the remaining parameters at their default settings and click **OK**.

9. Compile and activate the data flow, checking carefully that there are no errors or warnings.

   See How to activate a data flow (on page 232).

10. Locate the active data flow in the Monitor Inventory (on page 520) and open its Monitor Details (on page 521).

    The policy will be invoked automatically to create a replication according to the schedule specified in the policy. The policy can also be manually triggered from the source node in the monitor data flow.

    Note: No replication will be created until it is first triggered manually or by the schedule.

    See How to trigger an operation from an active data flow (on page 260).

11. Watch the active data flow via the Monitor Details (on page 521) to ensure the policy is operating as expected.
    For a healthy data flow you will periodically see:

    - Replication jobs appearing in the **Jobs** area below the data flow that cycle through stages and ending in *Progress - Completed*.
    - Information messages appearing in the **Logs** area below the data flow indicating rules activation, storage handler and sequencer events.
    - Attachments to storage handler log events confirming which volumes are being replicated.

    For a problematic data flow you may see:

    - Permanent **Node Status** icons appear over nodes and associated warning messages displayed to the right of the data flow area.
    - Backup jobs appearing in the **Jobs** area below the data flow that cycle through stages and terminating in *Progress - Failed*.
    - Warning and error messages appearing in the **Logs** area below the data flow indicating failed events.

12. Review the status of the Hitachi *Block Device* via the relevant Hitachi Block Device Details (on page 859) and replications via the Hitachi Block Replications Inventory (on page 878), to ensure the replication is being created and refreshed.

    Hitachi Block Devices require ongoing surveillance to ensure that they are operating correctly and sufficient resources are available to store your data securely. See How to view the status of a Hitachi Block storage device (on page 282).

    The replication process can be paused and resumed from here if required.

    A new ShadowImage replication will appear in the Hitachi Block Replications Inventory (on page 878) and be updated periodically as dictated by the schedule for the policy operation. The previous replication will be overwritten upon each refresh.

How to replicate a file system with Refreshed Thin Image

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the source node where the Hitachi Block LDEV is mounted. Note that the LDEV is actually located on the Hitachi Block storage device.
- The Protector Client software has been installed on the destination node that will act as a proxy for the Hitachi Block storage device. Note that for a Refreshed Thin Image replication, the source and destination LDEVs are located on the same device.
- The storage device has been set up as per the Protector requirements and prerequisites. Refer to Hitachi Block prerequisites (on page 21).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).

This task describes the steps to follow when protecting data that resides on a file system created on an LDEV allocated from a Hitachi Block storage device. A Refreshed Thin Image hardware replication of the P-VOL is created as an S-VOL residing within the same storage device. For more information, refer to About Thin Image and Thin Image Advanced differential and refreshed snapshots (on page 66). The data flow and policy are as follows:



Figure 57 Refreshed Thin Image Replication Data Flow

Table 12 Path Replication Policy

Classification Type	Parameter	Value
Path	Include	E:\testdata (E: is where the Hitachi Block LDEV is mounted)

Operation Type	Parameter	Value	Assigned Nodes
Replicate	Run Options	Run on Schedule (See below)	OS Host, Hitachi Block Device

Schedule Item Type	Parameter	Value	Policy Operations
Trigger	Days	Select All	Replicate (See above)
	Weeks	Select All	
	Time	Scheduled Time	
	Start Time	15:00	
	Duration	00:00	

Procedure

1. Locate the source node in the Nodes Inventory (on page 536) and check that it is authorized and online. This node is where the production LDEV to be replicated is mounted.

   For a file system replication using a *Path* classification, a basic *OS Host* node is required. It is not necessary to create the source node in this case since all Protector client nodes default to this type when installed. See How to authorize a node (on page 262).
2. Locate the node in the Nodes Inventory (on page 536) that will control the Hitachi Block Device via a CMD (Command Device) interface and check that it is authorized and online.

   This node is used by Protector to orchestrate replication and is identified as the **Proxy Node** when creating the Hitachi Block Device node in the next step. This node is known as an ISM node. The ISM node does not appear in the data flow.
3. Create a new Hitachi Block Device node (unless one already exists) using the Hitachi Block Device Node Wizard (on page 583) and check that it is authorized and online.

   The Hitachi *Block Device* node type is grouped under **Storage** in the Node Type Wizard (on page 539). See How to add a node (on page 261) and How to authorize a node (on page 262). This node appears in the replication data flow as the destination node.
4. Define a policy as shown in the table above using the Policy Wizard (on page 671). See How to create a policy (on page 273)
   a. Define a Path classification using the Path Classification Wizard (on page 713).

      The *Path* classification is grouped under **Physical** in the Policy Wizard (on page 671).
   b. Define a *Replicate* operation using the Replicate Operation Wizard (on page 722).

      Refreshed Thin Image replication runs as a batch operation based on a *Trigger* schedule.
   c. Define a *Trigger* schedule using the Schedule Wizard (on page 850); accessed by clicking on **Manage Schedules**.

      See How to create a schedule (on page 280).
5. Draw a data flow as shown in the figure above using the Data Flow Wizard (on page 361), that shows the *OS Host* source node connected to the Hitachi *Block Device* via a *Batch* mover.

   Refreshed Thin Image is an in-system replication technology, so the Hitachi *Block Device* node is where both the source (P-VOL) and destination (S-VOL) volumes are located. See How to create a data flow (on page 228).
6. Assign the *Path-Replicate* policy to the *OS Host* source node.

   See How to apply a policy to nodes on a data flow (on page 231).
7. Assign the *Replicate* operation to the Hitachi *Block Device* node.
   The Hitachi Block Replication Configuration Wizard (on page 398) is displayed.
8. Set the replication type to **Refreshed Snapshot**, then choose a **Pool** from one of the available Thin Image *Pools*. Leave the remaining parameters at their default settings and click **OK**.

9. Compile and activate the data flow, checking carefully that there are no errors or warnings.

   See How to activate a data flow (on page 232).

10. Locate the active data flow in the Monitor Inventory (on page 520) and open its Monitor Details (on page 521).

    The policy will be invoked automatically to create a replication according to the schedule specified in the policy. The policy can also be manually triggered from the source node in the Monitor Details (on page 521).

    Note: No replication will be created until it is first triggered manually or by the schedule.

    See How to trigger an operation from an active data flow (on page 260).

11. Watch the active data flow via the Monitor Details (on page 521) to ensure the policy is operating as expected.
    For a healthy data flow you will periodically see:

    - Replication jobs appearing in the **Jobs** area below the data flow that cycle through stages and ending in *Progress - Completed*.
    - Information messages appearing in the **Logs** area below the data flow indicating rules activation, storage handler and sequencer events.
    - Attachments to storage handler log events confirming which volumes are being replicated.

    For a problematic data flow you may see:

    - Permanent **Node Status** icons appear over nodes and associated warning messages displayed to the right of the data flow area.
    - Backup jobs appearing in the **Jobs** area below the data flow that cycle through stages and terminating in *Progress - Failed*.
    - Warning and error messages appearing in the **Logs** area below the data flow indicating failed events.

12. Review the status of the Hitachi *Block Device* via the relevant Hitachi Block Device Details (on page 859) and replications via the Hitachi Block Replications Inventory (on page 878), to ensure the replication is being created and refreshed.

    Hitachi Block Devices require ongoing surveillance to ensure that they are operating correctly and sufficient resources are available to store your data securely. See How to view the status of a Hitachi Block storage device (on page 282).

    The replication process can be paused and resumed from here if required.

    A new Refreshed Thin Image replication will appear in the Hitachi Block Replications Inventory (on page 878) and be updated periodically as dictated by the schedule for the policy operation. The previous replication will be overwritten upon each refresh.

How to replicate a file system with Universal Replicator

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the source node where the Hitachi Block LDEV is mounted. Note that the LDEV is actually located on the primary Hitachi Block storage device.
- The Protector Client software has been installed on the nodes that will act as proxies for both primary and secondary Hitachi Block storage devices. Note that for a Universal Replicator replication, the source and destination LDEVs are located on different devices.
- The primary and secondary storage devices have been set up as per the Protector requirements and prerequisites. Refer to Hitachi Block prerequisites (on page 21).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).

This task describes the steps to follow when protecting data that resides on a file system created on an LDEV allocated from a Hitachi Block storage device. A Universal Replicator hardware replication of the P-VOL is created as an S-VOL residing within a different storage device. For more information, refer to About Universal Replicator (on page 75). The data flow and policy are as follows:



Figure 58 Universal Replicator Replication Data Flow

Table 13 Path Replication Policy

Classification Type	Parameter	Value
Path	Include	E:\testdata (E: is where the Hitachi Block LDEV is mounted)

Operation Type	Parameter	Value	Assigned Nodes
Replicate	Refresh Options	Select a schedule for 'Refresh on schedule'	OS Host, Secondary Hitachi Block Device
	Source Options	Do not Quiesce configured applications before backup. See note in the 'compile and activation' step below.	

Procedure

1. Locate the source node in the Nodes Inventory (on page 536) and check that it is authorized and online. This node is where the primary LDEV to be replicated is mounted.

   For a file system replication using a *Path* classification, a basic *OS Host* node is required. It is not necessary to create the source node in this case since all Protector client nodes default to this type when installed. See How to authorize a node (on page 262).
2. Locate the nodes in the Nodes Inventory (on page 536) that will control the primary and secondary Hitachi Block Devices via CMD (Command Device) interfaces and check that they are authorized and online.

   These nodes are used by Protector to orchestrate replication of the primary LDEV to the secondary and are identified as the **Proxy Node** when creating the primary and secondary Hitachi Block Device nodes in the next step. These nodes are known as ISM nodes. The ISM nodes do not appear in the data flow.
3. Create new primary and secondary Hitachi Block Device nodes (unless they already exist) using the Hitachi Block Device Node Wizard (on page 583) and check that they are authorized and online.

   The Hitachi *Block Device* node type is grouped under **Storage** in the Node Type Wizard (on page 539). See How to add a node (on page 261) and How to authorize a node (on page 262). The secondary Hitachi Block Device node appears in the replication data flow as the destination node. The primary Hitachi Block Device node is represented in the data flow by the *OS Host* node where the primary LDEV is mounted.
4. Define a policy as shown in the table above using the Policy Wizard (on page 671). See How to create a policy (on page 273).
   a. Define a Path classification using the Path Classification Wizard (on page 713).

      The a *Path* classification is grouped under **Physical** in the Policy Wizard (on page 671).
   b. Define a *Replicate* operation using the Replicate Operation Wizard (on page 722).

      When an application is used to automatically select the PVOLs used in this continuous replication a trigger schedule can be defined that invokes the application to re-evaluate the PVOLs involved.
5. Draw a data flow as shown in the figure above using the Data Flow Wizard (on page 361), that shows the *OS Host* source node connected to the secondary Hitachi *Block Device* via a *Continuous* mover.

   Universal Replicator is a remote replication technology, so the Hitachi *Block Device* node shown on the data flow is the where the destination (S-VOL) volume is located. See How to create a data flow (on page 228).
6. Assign the *Path-Replicate* policy to the *OS Host* source node.

   See How to apply a policy to nodes on a data flow (on page 231).
7. Assign the *Replicate* operation to the Hitachi *Block Device* node.
   The Hitachi Block Replication Configuration Wizard (on page 398) is displayed.
8. Set the replication type to **Asynchronous Remote Clone**, then:
   a. Choose a **Pool** from one of the available *Dynamic Pools*.
   b. Select a **Source Journal** on the primary Hitachi *Block Device* node.

c. Select a **Destination Journal** (the secondary Hitachi *Block Device* node is selected implicitly).

d. Leave the remaining parameters at their default settings and click **OK**.

9. Compile and activate the data flow, checking carefully that there are no errors or warnings.

   See How to activate a data flow (on page 232).

   Note: Because Universal Replicator cannot guarantee that the quiesce period constraint for Microsoft VSS can be met, the Rules Compiler will generate a warning if **Quiesce configured applications before backup** is selected in the policy's *Replicate* operation.

10. Locate the active data flow in the Monitor Inventory (on page 520) and open its Monitor Details (on page 521).

    The policy will be invoked automatically to create and then maintain the replication according to the policy.

11. Watch the active data flow via the Monitor Details (on page 521) to ensure the policy is operating as expected.
    For a healthy data flow you will see:

    - An initial replication job appearing in the **Jobs** area below the data flow that cycle through stages and ending in *Progress - Completed*.
    - Information messages appearing in the **Logs** area below the data flow indicating rules activation, storage handler and sequencer events.
    - Attachments to storage handler log events confirming which volumes are being replicated.

    For a problematic data flow you may see:

    - Permanent **Node Status** icons appear over nodes and associated warning messages displayed to the right of the data flow area.
    - Backup jobs appearing in the **Jobs** area below the data flow that cycle through stages and terminating in *Progress - Failed*.
    - Warning and error messages appearing in the **Logs** area below the data flow indicating failed events.

12. Review the status of the Hitachi *Block Device* via the relevant Hitachi Block Device Details (on page 859) and replications via the Hitachi Block Replications Inventory (on page 878), to ensure the replication is being created and maintained.

    Hitachi Block Devices require ongoing surveillance to ensure that they are operating correctly and sufficient resources are available to store your data securely. See How to view the status of a Hitachi Block storage device (on page 282).

    The replication process can be paused and resumed from here if required.

    A Universal Replicator replication will appear in the Hitachi Block Replications Inventory (on page 878) and will be updated as and when writes to the primary are made.

How to replicate a file system with TrueCopy

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the source node where the Hitachi Block LDEV is mounted. Note that the LDEV is actually located on the primary Hitachi Block storage device.
- The Protector Client software has been installed on the nodes that will act as proxies for both primary and secondary Hitachi Block storage devices. Note that for a TrueCopy replication, the source and destination LDEVs are located on different devices.
- The primary and secondary storage devices have been set up as per the Protector requirements and prerequisites. Refer to Hitachi Block prerequisites (on page 21).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).

This task describes the steps to follow when protecting data that resides on a file system created on an LDEV allocated from a Hitachi Block storage device. A TrueCopy hardware replication of the P-VOL is created as an S-VOL residing within a different storage device. For more information, refer to About TrueCopy replication (on page 74). The data flow and policy are as follows:



Figure 59 TrueCopy Replication Data Flow

Table 14 Path Replication Policy

Classification Type	Parameter	Value
Path	Include	E:\testdata (E: is where the Hitachi Block LDEV is mounted)

Operation Type	Parameter	Value	Assigned Nodes
Replicate	Refresh Options	Select a schedule for 'Refresh on schedule'	OS Host, Secondary Hitachi Block Device

Procedure

1. Locate the source node in the Nodes Inventory (on page 536) and check that it is authorized and online. This node is where the primary LDEV to be replicated is mounted.

   For a file system replication using a *Path* classification, a basic *OS Host* node is required. It is not necessary to create the source node in this case since all Protector client nodes default to this type when installed. See How to authorize a node (on page 262).

2. Locate the nodes in the Nodes Inventory (on page 536) that will control the primary and secondary Hitachi Block Devices via CMD (Command Device) interfaces and check that they are authorized and online.

   These nodes are used by Protector to orchestrate replication of the primary LDEV to the secondary and are identified as the **Proxy Node** when creating the primary and secondary Hitachi Block Device nodes in the next step. These nodes are known as ISM nodes. The ISM nodes do not appear in the data flow.

3. Create new primary and secondary Hitachi Block Device nodes (unless ones already exists) using the Hitachi Block Device Node Wizard (on page 583) and check that they are authorized and online.

   The Hitachi *Block Device* node type is grouped under **Storage** in the Node Type Wizard (on page 539). See How to add a node (on page 261) and How to authorize a node (on page 262).The secondary Hitachi Block Device node appears in the replication data flow as the destination node. The primary Hitachi Block Device node is represented in the data flow by the *OS Host* node where the primary LDEV is mounted.

4. Define a policy as shown in the table above using the Policy Wizard (on page 671). See How to create a policy (on page 273)

   a. Define a Path classification using the Path Classification Wizard (on page 713).

      The a *Path* classification is grouped under **Physical** in the Policy Wizard (on page 671).

   b. Define a *Replicate* operation using the Replicate Operation Wizard (on page 722).

      When an application is used to automatically select the PVOLs used in this continuous replication a trigger schedule can be defined that invokes the application to re-evaluate the PVOLs involved.

5. Draw a data flow as shown in the figure above using the Data Flow Wizard (on page 361), that shows the *OS Host* source node connected to the secondary Hitachi *Block Device* via a *Continuous* mover.

   TrueCopy is a remote replication technology, so the Hitachi *Block Device* node shown on the data flow is the where the destination secondary volume (S-VOL) is located. See How to create a data flow (on page 228).

6. Assign the *Path-Replicate* policy to the *OS Host* source node.

   See How to apply a policy to nodes on a data flow (on page 231).

7. Assign the *Replicate* operation to the Hitachi *Block Device* node.
   The Hitachi Block Replication Configuration Wizard (on page 398) is displayed.

8. Set the replication type to **Synchronous Remote Clone**, then choose a **Pool** from one of the available *Dynamic Pools*. Leave the remaining parameters at their default settings and click **OK**.

9. Compile and activate the data flow, checking carefully that there are no errors or warnings.

   See How to activate a data flow (on page 232).

10. Locate the active data flow in the Monitor Inventory (on page 520) and open its Monitor Details (on page 521).

    The policy will be invoked automatically to create and then maintain the replication according to the policy.

11. Watch the active data flow via the Monitor Details (on page 521) to ensure the policy is operating as expected.

For a healthy data flow you will periodically see:

- An initial replication job appearing in the **Jobs** area below the data flow that cycle through stages and ending in *Progress - Completed*.
- Information messages appearing in the **Logs** area below the data flow indicating rules activation, storage handler and sequencer events.
- Attachments to storage handler log events confirming which volumes are being replicated.

For a problematic data flow you may see:

- Permanent **Node Status** icons appear over nodes and associated warning messages displayed to the right of the data flow area.
- Backup jobs appearing in the **Jobs** area below the data flow that cycle through stages and terminating in *Progress - Failed*.
- Warning and error messages appearing in the **Logs** area below the data flow indicating failed events.

12. Review the status of the Hitachi *Block Device* via the relevant Hitachi Block Device Details (on page 859) and replications via the Hitachi Block Replications Inventory (on page 878), to ensure the replication is being created and maintained.

    Hitachi Block Devices require ongoing surveillance to ensure that they are operating correctly and sufficient resources are available to store your data securely. See How to view the status of a Hitachi Block storage device (on page 282).

    The replication process can be paused and resumed from here if required.

    A TrueCopy replication will appear in the Hitachi Block Replications Inventory (on page 878) and will be updated as and when writes to the primary are made.

How to replicate a file system with Global-Active Device

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the source node where the Hitachi Block LDEV is mounted. Note that the LDEV is actually located on the primary Hitachi Block storage device.
- The Protector Client software has been installed on the nodes that will act as proxies for both primary and secondary Hitachi Block storage devices. Note that for a Global-Active Device replication, the source and destination LDEVs are located on different devices.
- The primary and secondary storage devices have been set up as per the Protector requirements and prerequisites. Refer to Hitachi Block prerequisites (on page 21).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).

This task describes the steps to follow when protecting data that resides on a file system created on an LDEV allocated from a Hitachi Block storage device. A Global-Active Device hardware replication of the P-VOL is created as an S-VOL residing within a different storage device. For more information, refer to About Global-Active Device replication (on page 76). The data flow and policy are as follows:



Figure 60 Global-Active Device Replication Data Flow

Table 15 Path Replication Policy

Classification Type	Parameter	Value
Path	Include	E:\testdata (E: is where the Hitachi Block LDEV is mounted)

Operation Type	Parameter	Value	Assigned Nodes
Replicate	Refresh Options	Select a schedule for 'Refresh on schedule'	OS Host, Secondary Hitachi Block Device

Procedure

1. Locate the source node in the Nodes Inventory (on page 536) and check that it is authorized and online. This node is where the primary LDEV to be replicated is mounted.

   For a file system replication using a *Path* classification, a basic *OS Host* node is required. It is not necessary to create the source node in this case since all Protector client nodes default to this type when installed. See How to authorize a node (on page 262).

2. Locate the nodes in the Nodes Inventory (on page 536) that will control the primary and secondary Hitachi Block Devices via CMD (Command Device) interfaces and check that they are authorized and online.

   These nodes are used by Protector to orchestrate replication of the primary LDEV to the secondary and are identified as the **Proxy Node** when creating the primary and secondary Hitachi Block Device nodes in the next step. These nodes are known as ISM nodes. The ISM nodes do not appear in the data flow.

3. Create new primary and secondary Hitachi Block Device nodes (unless ones already exists) using the Hitachi Block Device Node Wizard (on page 583) and check that they are authorized and online.

   The Hitachi *Block Device* node type is grouped under **Storage** in the Node Type Wizard (on page 539). See How to add a node (on page 261) and How to authorize a node (on page 262). The secondary Hitachi Block Device node appears in the replication data flow as the destination node. The primary Hitachi Block Device node is represented in the data flow by the *OS Host* node where the primary LDEV is mounted.

4. Define a policy as shown in the table above using the Policy Wizard (on page 671). See How to create a policy (on page 273)

   a. Define a Path classification using the Path Classification Wizard (on page 713).

      The *Path* classification is grouped under **Physical** in the Policy Wizard (on page 671).

   b. Define a *Replicate* operation using the Replicate Operation Wizard (on page 722).

      When an application is used to automatically select the PVOLs used in this continuous replication a trigger schedule can be defined that invokes the application to re-evaluate the PVOLs involved.

5. Draw a data flow as shown in the figure above using the Data Flow Wizard (on page 361), that shows the *OS Host* source node connected to the secondary Hitachi *Block Device* via a *Continuous* mover.

   Global-Active Device is a remote replication technology, so the Hitachi *Block Device* node shown on the data flow is the where the destination (S-VOL) volume is located. See How to create a data flow (on page 228).

6. Assign the *Path-Replicate* policy to the *OS Host* source node.

   See How to apply a policy to nodes on a data flow (on page 231).

7. Assign the *Replicate* operation to the Hitachi *Block Device* node.
   The Hitachi Block Replication Configuration Wizard (on page 398) is displayed.

8. Set the replication type to **Active-Active Remote Clone**, then:

   a. Choose a **Pool** from one of the available *Dynamic Pools*.

   b. Choose a **Target Quorum** from one of those listed.

   c. Leave the remaining parameters at their default settings and click **OK**.

9. Compile and activate the data flow, checking carefully that there are no errors or warnings.

   See How to activate a data flow (on page 232).

10. Locate the active data flow in the Monitor Inventory (on page 520) and open its Monitor Details (on page 521).

    The policy will be invoked automatically to create and then maintain the replication according to the policy.

11. Watch the active data flow via the Monitor Details (on page 521) to ensure the policy is operating as expected.
    For a healthy data flow you will periodically see:
    - An initial replication job appearing in the **Jobs** area below the data flow that cycle through stages and ending in *Progress - Completed*.
    - Information messages appearing in the **Logs** area below the data flow indicating rules activation, storage handler and sequencer events.
    - Attachments to storage handler log events confirming which volumes are being replicated.

    For a problematic data flow you may see:
    - Permanent **Node Status** icons appear over nodes and associated warning messages displayed to the right of the data flow area.
    - Backup jobs appearing in the **Jobs** area below the data flow that cycle through stages and terminating in *Progress - Failed*.
    - Warning and error messages appearing in the **Logs** area below the data flow indicating failed events.
12. Review the status of the Hitachi *Block Device* via the relevant Hitachi Block Device Details (on page 859) and replications via the Hitachi Block Replications Inventory (on page 878), to ensure the replication is being created and maintained.
    Hitachi Block Devices require ongoing surveillance to ensure that they are operating correctly and sufficient resources are available to store your data securely. See How to view the status of a Hitachi Block storage device (on page 282).
    The replication process can be paused and resumed from here if required.
    A Global-Active Device replication will appear in the Hitachi Block Replications Inventory (on page 878) and will be updated as and when writes to the primary or secondary are made.

How to implement 3DC multi-target with delta UR replication

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the source node where the Hitachi Block LDEV is mounted. Note that the LDEV is actually located on the primary Hitachi Block storage device.
- The Protector Client software has been installed on the nodes that will act as proxies for primary, secondary and tertiary Hitachi Block storage devices. Note that for GAD, TC and UR replications, the source and destination LDEVs are located on different devices.
- The primary, secondary and tertiary storage devices have been set up as per the Protector requirements and prerequisites. Refer to Hitachi Block prerequisites (on page 21).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).

This task describes the steps to follow when protecting data that resides on a file system created on an LDEV allocated from a Hitachi Block storage device. A GAD or TC replication of the P-VOL at the primary site is created as an S-VOL at the secondary site. A UR replication of the P-VOL is created as an S-VOL at the tertiary site. A Delta UR replication is created between the S-VOLs at the secondary and tertiary sites (this remains suspended unless primary site failure occurs). For more information, refer to About three datacentre multi-target with delta (on page 80).

Note: Protector currently only supports the setup of *3DC Multi-target with Delta Replication*. In the event of a primary, secondary or tertiary site failure, the Monitor Details (on page 521) data flow will display notifications indicating any problems with the corresponding movers, and appropriate messages will appear in the Logs Inventory (on page 509).

- For primary site failure:
  1. The *Delta* UR failover link will be invoked automatically by the underlying hardware storage devices to provide near immediate protection of the secondary site.
  2. The data flow should be dissociated from Protector before the hardware storage devices are recovered, following procedures defined in the relevant storage device operating manuals. See How to dissociate a replication from Protector (on page 209).
  3. The data flow for the recovered replication should be re-adopt into Protector and re-activated. See How to adopt a replication into Protector (on page 206)
- For secondary or tertiary site failure, the data flow should remain actived in Protector. Once the hardware storage devices are recovered, Protector will clear its notifications and resume.

The data flow and policy are as follows:

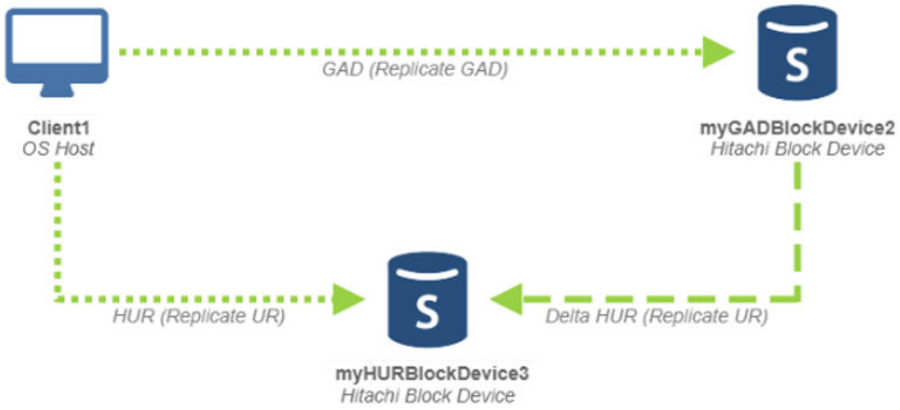


Figure 61 3DC Multi-target with Delta Replication Data Flow

Table 16 Path Replication Policy

Classification Type	Parameter	Value
Path	Include	E:\testdata (E: is where the Hitachi Block LDEV is mounted)

Operation Type	Parameter	Value	Assigned Nodes
Replicate	Run Options	N/A (GAD is a continuous replication, so the Run option is ignored)	Secondary Hitachi Block Device (from the primary)
Replicate	Run Options	N/A (UR is a continuous replication, so the Run option is ignored)	Tertiary Hitachi Block Device (from the primary)
Replicate	Run Options	N/A (Delta UR is a continuous replication, so the Run option is ignored)	Tertiary Hitachi Block Device (from the secondary)

Procedure

1. Locate the source node in the Nodes Inventory (on page 536) and check that it is authorized and online. This node is where the primary LDEV to be replicated is mounted.

   For a file system replication using a *Path* classification, a basic *OS Host* node is required. It is not necessary to create the source node in this case since all Protector client nodes default to this type when installed. See How to authorize a node (on page 262).
2. Locate the nodes in the Nodes Inventory (on page 536) that will control the primary, secondary and tertiary Hitachi Block Devices via CMD (Command Device) interfaces and check that they are authorized and online.

   These nodes are used by Protector to orchestrate replication of the primary LDEV to the secondary and tertiary sites, and are identified as the **Proxy Node** when creating the primary and secondary Hitachi Block Device nodes in the next step. These nodes are known as ISM nodes. The ISM nodes do not appear in the data flow.
3. Create new primary, secondary and tertiary Hitachi Block Device nodes (unless ones already exists) using the Hitachi Block Device Node Wizard (on page 583) and check that they are authorized and online.

   The Hitachi *Block Device* node type is grouped under **Storage** in the Node Type Wizard (on page 539). See How to add a node (on page 261) and How to authorize a node (on page 262). The secondary and tertiary Hitachi Block Device nodes appear in the replication data flow as the destination nodes. The primary Hitachi Block Device node is represented in the data flow by the *OS Host* node where the primary LDEV is mounted.
4. Define a policy as shown in the table above using the Policy Wizard (on page 671). See How to create a policy (on page 273)
   a. Define a Path classification using the Path Classification Wizard (on page 713).

      The *Path* classification is grouped under **Physical** in the Policy Wizard (on page 671).
   b. Define three *Replicate* operations (these represent the primary to secondary GAD or UR, primary to tertiary UR and secondary to tertiary Delta UR replications) using the Replicate Operation Wizard (on page 722).

      GAD and UR replications run as continuous operations and thus no schedule needs to be defined.
5. Draw a data flow as shown in the figure above using the Data Flow Wizard (on page 361), that shows the *OS Host* source node connected to the secondary and tertiary Hitachi *Block Device*s via *Continuous* movers, and the secondary connected to the tertiary Hitachi *Block Device* via a *Failover* mover.

   GAD and UR are remote replication technologies, so the Hitachi *Block Device* nodes shown on the data flow are where the secondary and tertiary destination (S-VOL) volumes are located. See How to create a data flow (on page 228).
6. Assign the *Path-Replicate* policy to the *OS Host* source node.

   See How to apply a policy to nodes on a data flow (on page 231).
7. Assign the first *Replicate* operation to the secondary Hitachi *Block Device* node. The Hitachi Block Replication Configuration Wizard (on page 398) is displayed.
8. Set the replication type to **Active-Active Remote Clone**, then:
   a. Choose a **Pool** from one of the available *Dynamic Pools*.

b. Choose a **Target Quorum** from one of those listed.
c. Leave the remaining parameters at their default settings and click **OK**.

9. Assign the second *Replicate* operation to the tertiary Hitachi *Block Device* node.
The Hitachi Block Replication Configuration Wizard (on page 398) is displayed.
10. Set the replication type to **Asynchronous Remote Clone**, then:
    a. Choose a **Pool** from one of the available *Dynamic Pools*.
    b. Choose a **Source Journal** from one of those listed for the primary node.
    c. Choose a **Destination Journal** from one of those listed.
    d. Leave the remaining parameters at their default settings and click **OK**.
11. Assign the third *Replicate* operation to the tertiary Hitachi *Block Device* node.
The Hitachi Block Replication Configuration Wizard (on page 398) is displayed.
12. Set the replication type to **Asynchronous Remote Failover**, then:
    a. Choose a **Pool** from one of the available *Dynamic Pools*.
    b. Choose a **Source Journal** from one of those listed for the tertiary node.
    c. Leave the remaining parameters at their default settings and click **OK**.

    Protector will automatically use the same **Destination Journal** as selected for the **Asynchronous Remote Clone** replication configured in the preceding steps.

    Note: If you specify a **Mirror Unit** for this **Asynchronous Remote Failover** replication, then it must differ from the one selected for the **Asynchronous Remote Clone** replication in the preceding steps.
13. Compile and activate the data flow, checking carefully that there are no errors or warnings.

    See How to activate a data flow (on page 232).
14. Locate the active data flow in the Monitor Inventory (on page 520) and open its Monitor Details (on page 521).

    The policy will be invoked automatically to create and then maintain the replication accordingly.
15. Watch the active data flow via the Monitor Details (on page 521) to ensure the policy is operating as expected.
For a healthy data flow you will periodically see:
    - Initial replication jobs appearing in the **Jobs** area below the data flow that cycle through stages and ending in *Progress - Completed*.
    - Information messages appearing in the **Logs** area below the data flow indicating rules activation, storage handler and sequencer events.
    - Attachments to storage handler log events confirming which volumes are being replicated.

For a problematic data flow you may see:

- Permanent **Node Status** icons appear over nodes and associated warning messages displayed to the right of the data flow area.
- Backup jobs appearing in the **Jobs** area below the data flow that cycle through stages and terminating in *Progress - Failed*.
- Warning and error messages appearing in the **Logs** area below the data flow indicating failed events.

16. Review the status of each Hitachi *Block Device* via the relevant Hitachi Block Device Details (on page 859) and replications via the Hitachi Block Replications Inventory (on page 878), to ensure the GAD and UR replications are being created and maintained.

    Hitachi Block Devices require ongoing surveillance to ensure that they are operating correctly and sufficient resources are available to store your data securely. See How to view the status of a Hitachi Block storage device (on page 282).

    The replication processes can be paused and resumed from here if required.

    There will be three replication records in the Hitachi Block Replications Inventory (on page 878) corresponding to the GAD, the active UR and the suspended failover UR replication.

How to synchronize snapshots with a replication

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the source node where the Hitachi Block LDEV is mounted. Note that the LDEV is actually located on the primary Hitachi Block storage device.
- The Protector Client software has been installed on the nodes that will act as proxies for both primary and secondary Hitachi Block storage devices. Note that for a TrueCopy replication, the source and destination LDEVs are located on different devices.
- The primary and secondary storage devices have been set up as per the Protector requirements and prerequisites. Refer to Hitachi Block prerequisites (on page 21).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).

This task demonstrates how snapshot and replication operations can be synchronized with one another to ensure that all are performed at the same point in time and thus capture the identical state of the source data. For more information refer to About synchronization groups (on page 99).

Here, a TrueCopy replication of the P-VOL is created as an S-VOL residing within a different storage device. Synchronized TI snapshots are then created on the primary and secondary Block Storage devices. For more information, refer to About local and remote snapshots (on page 84). The data flow and policy are as follows:

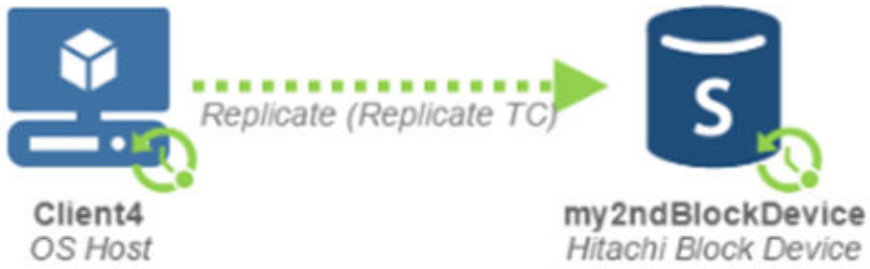


Figure 62 TrueCopy Replication with Local and Remote Thin Image Snapshots Data Flow

Table 17 Path Replication Policy

Classification Type	Parameter	Value
Path	Include	E:\testdata (E: is where the Hitachi Block LDEV is mounted)

Operation Type	Parameter	Value	Assigned Nodes
Replicate	Run Options	N/A (TrueCopy is a continuous replication, so the Run option is ignored)	OS Host, Secondary Hitachi Block Device
Snapshot (on local device)	Mode	Hardware	OS Host
	Hardware Type	Hitachi Block	
	Retention	2 hour	
	RPO	10 mins	
	Run Options	Run on Schedule (see synch group schedule below)	

Operation Type	Parameter	Value	Assigned Nodes
Snapshot (on remote device)	Mode	Hardware	Secondary Hitachi Block Device
	Hardware Type	Hitachi Block	
	Retention	1 hours (this can differ from the local snapshot)	
	RPO	10 mins (this must match the local snapshot)	
	Run Options	Run on Schedule (see synch group schedule below)	

Table 18 Synchronization Group Schedule

Schedule Item Type	Parameter	Value	Policy Operations
Trigger	N/A (this schedule defines a synchronization group name for local and remote snapshots. All parameters are ignored.)	N/A	Snapshot (local), Snapshot (remote)

Procedure

1. Locate the source node in the Nodes Inventory (on page 536) and check that it is authorized and online. This node is where the primary LDEV to be replicated is mounted.

   For a file system replication using a *Path* classification, a basic *OS Host* node is required. It is not necessary to create the source node in this case since all Protector client nodes default to this type when installed. See How to authorize a node (on page 262).

2. Locate the nodes in the Nodes Inventory (on page 536) that will control the primary and secondary Hitachi Block Devices via CMD (Command Device) interfaces and check that they are authorized and online.

   These nodes are used by Protector to orchestrate replication of the primary LDEV to the secondary and are identified as the **Proxy Node** when creating the primary and secondary Hitachi Block Device nodes in the next step. These nodes are known as ISM nodes. The ISM nodes do not appear in the data flow.

3. Create new primary and secondary Hitachi Block Device nodes (unless ones already exists) using the Hitachi Block Device Node Wizard (on page 583) and check that they are authorized and online.

   The Hitachi *Block Device* node type is grouped under **Storage** in the Node Type Wizard (on page 539). See How to add a node (on page 261) and How to authorize a node (on page 262). The secondary Hitachi Block Device node appears in the replication data flow as the destination node. The primary Hitachi Block Device node is represented in the data flow by the *OS Host* node where the primary LDEV is mounted.

4. Define a policy as shown in the table above using the Policy Wizard (on page 671). This policy contains operations for the replication, local and remote snapshots. See How to create a policy (on page 273)

   a. Define a Path classification using the Path Classification Wizard (on page 713).

      The *Path* classification is grouped under **Physical** in the Policy Wizard (on page 671).

   b. Define a *Replicate* operation using the Replicate Operation Wizard (on page 722).

      TrueCopy replication runs as a continuous operation and thus no schedule needs to be defined.

   c. Define a local *Snapshot* operation using the Snapshot Operation Wizard (on page 724).

      Thin Image snapshots run based on the RPO. However we also want to synchronize the local and remote snapshots. This is done by defining a trigger schedule that is applied to both the local and remote snapshot operations.

   d. Define a *Trigger* schedule using the Schedule Wizard (on page 850); accessed by clicking on **Manage Schedules** in the Snapshot Operation Wizard (on page 724) for the local snapshot.

      Only the trigger schedule name is required; the parameters are not relevant here since the RPO of the local snapshot dictates when the local and remote snapshot operations are triggered. See How to create a schedule (on page 280).

   e. Define a remote *Snapshot* operation using the Snapshot Operation Wizard (on page 724).

      To synchronize the local and remote snapshots, apply the same trigger schedule to this snapshot operation that was applied to the local snapshot operation.

      Note: The local and remote snapshots must have the same RPO, otherwise a rules compiler error will be generated.

5. Draw a data flow as shown in the figure above using the Data Flow Wizard (on page 361), that shows the *OS Host* source node connected to the secondary Hitachi *Block Device* via a *Continuous* mover.

   TrueCopy is a remote replication technology, so the Hitachi *Block Device* node shown on the data flow is the where the destination (S-VOL) volume is located. See How to create a data flow (on page 228).
6. Assign the *Path-Replicate-Snaphot-Snapshot* policy to the *OS Host* source node.

   See How to apply a policy to nodes on a data flow (on page 231).
7. Assign the local *Snapshot* operation to the *OS Host* source node.
   The Hitachi Block Snapshot Configuration Wizard (on page 384) is displayed.
8. Select the **Snapshot Pool** by selecting one of the available Thin Image or hybrid pools.
9. Leave the remaining **Advanced Options** at their default settings, then click **OK**.

   The snapshot icon is now shown superimposed over the source node.
10. Assign the remote *Snapshot* operation to the remote Hitachi *Block Device* node.
    The Hitachi Block Snapshot Configuration Wizard (on page 384) is displayed.
11. Select the **Snapshot Pool** by selecting one of the available Thin Image or hybrid pools.
12. Leave the remaining **Advanced Options** at their default settings, then click **OK**.

    The snapshot icon is now shown superimposed over the source node.
13. Assign the *Replicate* operation to the remote Hitachi *Block Device* node.
    The Hitachi Block Replication Configuration Wizard (on page 398) is displayed.
14. Set the replication type to **Synchronous Remote Clone**, then choose a **Pool** from one of the available *Dynamic Pools*. Leave the remaining parameters at their default settings and click **OK**.
15. Compile and activate the data flow, checking carefully that there are no errors or warnings.

    See How to activate a data flow (on page 232).
16. Locate the active data flow in the Monitor Inventory (on page 520) and open its Monitor Details (on page 521).

    The policy will be invoked automatically to create and then maintain the replication according to the policy. Snapshot operations will be triggered synchronously on the source and destination nodes according to the RPO.
17. Watch the active data flow via the Monitor Details (on page 521) to ensure the policy is operating as expected.
    For a healthy data flow you will periodically see:

    - An initial replication job appearing in the **Jobs** area below the data flow that cycles through stages and ends in *Progress - Completed*.
    - Repeated replication and snapshot jobs appearing for the source node in the **Jobs** area triggered according to the RPO.
    - Repeated snapshot jobs appearing for the destination node in the **Jobs** area synchronized to the local snapshot.

- Information messages appearing in the **Logs** area below the data flow indicating rules activation, storage handler and sequencer events.
- Attachments to storage handler log events confirming which volumes are being replicated.

For a problematic data flow you may see:

- Permanent **Node Status** icons appear over nodes and associated warning messages displayed to the right of the data flow area.
- Backup jobs appearing in the **Jobs** area below the data flow that cycle through stages and terminating in *Progress - Failed*.
- Warning and error messages appearing in the **Logs** area below the data flow indicating failed events.

18. Review the status of the local and remote Hitachi *Block Devices* via the relevant Hitachi Block Device Details (on page 859) to ensure the replication and snapshots are being created and maintained.

    Hitachi Block Devices require ongoing surveillance to ensure that they are operating correctly and sufficient resources are available to store your data securely. See How to view the status of a Hitachi Block storage device (on page 282).

    The replication process can be paused and resumed from here if required.

    A TrueCopy replication will appear in the remote Hitachi Block Replications Inventory (on page 878) and will be updated as and when writes to the primary are made. New snapshots will appear in the local and remote Hitachi Block Snapshots Inventory (on page 869) periodically as dictated by the RPO of the policy. Old snapshots will be removed periodically as dictated by the *Retention Period* of the policy.

How to automatically mount a snapshot or replication

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the source node where the production Hitachi Block LDEV is mounted. Note that the LDEV is actually located on the primary Hitachi Block storage device.
- The Protector Client software has been installed on the nodes that will act as proxies for both primary and secondary Hitachi Block storage devices. Note that for a TrueCopy replication, the source and destination LDEVs are located on different devices.
- The primary and secondary storage devices have been set up as per the Protector requirements and prerequisites. Refer to Hitachi Block prerequisites (on page 21).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).

This task describes the steps to follow when replicating data on an a Hitachi Block storage LDEV, then using the replication for repurposing or performing a proxy backup. A TrueCopy replication of the P-VOL is created as an S-VOL and the S-VOL is then automatically mounted on a designated host machine. For more information, refer to About TrueCopy replication (on page 74) and About the automated Mount operation (on page 100). The data flow and policy is as follows:



Figure 63 TrueCopy Replication and Automated Mount Data Flow

Table 19 Path Replication Policy

Classification Type	Parameter	Value
Path	Include	E:\testdata (E: is where the Hitachi Block LDEV is mounted)

Operation Type	Parameter	Value	Assigned Nodes
Replicate	Run Options	N/A (TrueCopy is a continuous replication, so the Run option is ignored)	OS Host, Secondary Hitachi Block Device
Mount	Run Options	Run on Schedule (see schedule below)	Secondary Hitachi Block Device
	Source Options	Pre Script Post Script (scripts are user defined and	

Operation Type	Parameter	Value	Assigned Nodes
		application dependent)	

Schedule Item Type	Parameter	Value	Policy Operations
Trigger Time	Days	Select All	Mount (See above)
	Weeks	Select All	
	Time	Scheduled Time	
	Start Time	13:00	
	Duration	00:00	

Procedure

1. Follow the steps shown in How to replicate a file system with TrueCopy (on page 179) to create a TrueCopy replication, but with the addition of the Mount operation as follows:
2. Add the Mount operation to the *Path-Replicate* policy as shown in the table above using the Policy Wizard (on page 671). See How to create a policy (on page 273)
   a. Define a *Mount* operation using the Mount Operation Wizard (on page 721).

      The Mount operation is initiated using a *Trigger* schedule.
   b. Define the *Trigger* schedule shown in the table above, using the Schedule Wizard (on page 850) which is accessed by clicking on **Manage Schedules**.

      See How to create a schedule (on page 280).
3. Assign the *Mount* operation to the Hitachi *Block Device* node.
   The Hitachi Block Mount Configuration Wizard (on page 485) is displayed.
4. Set the mount operation type set to **Proxy Backup** and click **Next**.

   Proxy Backup can be used in conjunction with a continuous replication because it only pauses the replication while the proxy backup script is running.
5. Set the mount level to **OS** and click **Next**.
6. Set the **Host** to the node where the replication S-VOL is to be mounted and click **Next**.
7. Set the **Mount Location** to **Drive starting at letter**, select an available drive letter, then click **Finish**.
8. Compile and activate the data flow, checking carefully that there are no errors.

   A compiler warning (10366) will be generated stating that the replication will stop copying data to the destination while performing the proxy backup. This is expected behaviour. See How to activate a data flow (on page 232).

9. Locate the active data flow in the Monitor Inventory (on page 520) and open its Monitor Details (on page 521).

   The policy will be invoked automatically to create and then maintain the replication according to the policy. The proxy backup operation will run based on the specified schedule, pausing the replication while it executes.

10. Watch the active data flow via the Monitor Details (on page 521) to ensure the policy is operating as expected.
    For a healthy data flow you will periodically see:

    - An initial replication job appearing in the **Jobs** area below the data flow that cycle through stages and ending in *Progress - Completed*.
    - Information messages appearing in the **Logs** area below the data flow indicating rules activation, storage handler and sequencer events.
    - Attachments to storage handler log events confirming which volumes are being replicated.

    For a problematic data flow you may see:

    - Permanent **Node Status** icons appear over nodes and associated warning messages displayed to the right of the data flow area.
    - Backup jobs appearing in the **Jobs** area below the data flow that cycle through stages and terminating in *Progress - Failed*.
    - Warning and error messages appearing in the **Logs** area below the data flow indicating failed events.

11. Review the status of the Hitachi *Block Device* via the relevant Hitachi Block Device Details (on page 859) and replications via the Hitachi Block Replications Inventory (on page 878), to ensure the replication is being created and maintained.

    Hitachi Block Devices require ongoing surveillance to ensure that they are operating correctly and sufficient resources are available to store your data securely. See How to view the status of a Hitachi Block storage device (on page 282).

    A TrueCopy replication will appear in the Hitachi Block Replications Inventory (on page 878) and will be updated as and when writes to the primary are made.

How to mount an Hitachi block snapshot or replication

Before you begin

It is assumed that a file system path policy that creates hardware snapshots or a replication has been implemented and that at least one snapshot or replication has been created on the designated Block storage device. See How to snapshot a file system with Thin Image (on page 157) or How to replicate a file system with ShadowImage (on page 167) for an example of how to do this.



Note: It is not possible to mount the S-VOL of a GAD replication, paused or otherwise.

This task describes the steps to follow when mounting a file system path snapshot or replication from a Block storage device to the node other than the one from which it originated:

Procedure

1. Identify the destination where the data set is to be restored. Here we will *mount* a snapshot or replication to a destination machine and volume.

   Depending on the scenario, you can mount the snapshot or replication to its original node as a different volume or to a different node entirely. You can control the level of the mount operation so that the snapshot is added to a host group, through to mounting to the host OS.

2. Ensure that the restore location is prepared to receive the snapshot or replication data set by locating the node in the Nodes Inventory (on page 536) and checking it is authorized and online.

   Note:
   - For **Host** and **OS** level mounting, the mount location must have Protector Client software installed.
   - **SAN** level mount does not specify a host so Protector Client software does not need to be installed.

3. Locate the data set to be mounted by navigating to the Hitachi Block Snapshots Inventory (on page 869) or Hitachi Block Replications Inventory (on page 878) for the Hitachi Block storage device in question.

   See How to view the status of a Hitachi Block storage device (on page 282).

4. Select the snapshot or replication that you want to mount, then click **Mount** to open the Hitachi Block Snapshot or Replication Mount Wizard (on page 772).
   a. Select the mount level (**SAN**, **Host** or **OS**).
   b. Choose **Automatically discover** or a **Selected Host Group**.
   c. For **SAN** mount click **Finish**, for other mount types click **Next**.
   d. Specify the **Host** (i.e. the target machine).
   e. For **Host** mount click **Finish**, for **OS** mount click **Next**.
   f. Specify the **Mount Location**.
   g. For **OS** mount click **Finish**.

   The Jobs Inventory (on page 493) is displayed and a mount job will appear that cycles through stages and ending in *Progress - Completed*.

5. Once the mount process is complete, further steps may be needed to fix-up the data set before using it. In this example we will assume that no additional work is required other than inspecting the restored data on the target machine.

   The amount of fix-up work required depends on the applications accessing the restored data.

   Note: This example mounts data created using a *Path* classification. If you are backing up one of the application types directly supported by Protector, then you should use one of the *Application* classifications and refer to the appropriate Application Guide listed in Related documents (on page 17)).

6. Mounted snapshots or replications have a mount icon displayed on the corresponding tile in the Hitachi Block Snapshots Inventory (on page 869) or Hitachi

Block Replications Inventory (on page 878). Once the mounted snapshot or replication is finished with, click **Unmount** to unmount it.

How to revert a file system path from a snapshot or local replication

Before you begin

It is assumed that a file system path policy that creates snapshots (TI) or a local replication (RTI or SI) has been implemented and that at least one snapshot or replication has been created on the designated Block storage device. See How to snapshot a file system with Thin Image (on page 157), How to replicate a file system with Refreshed Thin Image (on page 171) or How to replicate a file system with ShadowImage (on page 167) for an example of how to do this.

This task describes the steps to follow when reverting a file system path snapshot or local replication from a Block storage device to the node from which the snapshot originated:

Procedure

1. Identify the destination where the data set is to be restored. You can only *revert* the snapshot or local replication to its original node and volume, so the destination will be the same machine and volume from which the data originated.
2. Ensure that the restore location is prepared to receive the reverted data set by locating the node in the Nodes Inventory (on page 536) and checking it is authorized and online.

   The revert location must have Protector Client software installed.
3. Stop any applications that access the revert location and ensure the filesystem is unmounted from the OS.

   For supported applications, these additional steps are described in the appropriate Application Guide (see Related documents (on page 17)). For other applications, consult the vendor's documentation.
4. The existing snapshot or local replication operation (and any replications immediately up and downstream of it) should be paused while revert is performed.

   > **Caution:** If snapshots and replications are combined on a data flow it is not possible to deactivate the scheduling of snapshots without also tearing down any replications on that data flow.

   See How to deactivate an active data flow (on page 261).
5. Locate the data set to be reverted by navigating to the Hitachi Block Snapshots Inventory (on page 869) or Hitachi Block Replications Inventory (on page 878) for the Hitachi Block storage device in question.

   See How to view the status of a Hitachi Block storage device (on page 282).
6. Select the snapshot or local replication that you want to revert to, then click **Revert** to open the Hitachi Block Revert Wizard (on page 813). The word '`REVERT`' must be typed in to enable the revert operation to proceed.

> **Caution:** The process of reverting data will result in the overwriting of all of the original data that exists on the revert location.
>
> Ensure that any critical data is copied to a safe location or is included in the data set being restored.

7. Once the revert process is complete, further steps may be needed to fix-up the data set before using it. In this example we will assume that no additional work is required other than inspecting the restored data on the target machine.

   > **Note:** When reverting a volume on a Windows machine, it is necessary to perform a reboot to ensure the volume is remounted correctly. For dynamic disks, if the reverted volume's status is indicated as *'Healthy (At Risk)'* then it will be necessary to **Offline** then **Online** the volume via the Windows **Disk Management** console.

   The amount of fix-up work required depends on the applications accessing the restored data.

   > **Note:** This example restores data created using a *Path* classification. If you are backing up one of the application types directly supported by Protector, then you should use one of the *Application* classifications and refer to the appropriate Application Guide listed in Related documents (on page 17)).

8. Restart any applications that access the restored data.

   For supported applications, these additional steps are described in the appropriate Application Guide (see Related documents (on page 17)). For other applications, consult the vendor's documentation.

9. Resume any backup policies for the reverted data set.

   Data flows can be reactivated via the Data Flows Inventory (on page 355).

How to swap (takeover/takeback) a replication

Before you begin

It is assumed that you have implemented a simple file system path, active-passive (i.e. TC or UR) replication policy with production applications accessing the primary volume(s) of the replication pair. See How to replicate a file system with TrueCopy (on page 179) for an example of how to do this.

If you are swapping an active-active (GAD) replication then additional steps may be required, especially if using a cross-path setup (refer to About Hitachi block based replication swapping (takeover/takeback) (on page 111) and About Global-Active Device Cross Path (on page 77)).

In the case of primary site maintenance, application failure, primary volume failure or disaster recovery, it may be necessary to move production to the secondary site, resolve the issue at the primary site and then move production back to the primary site.

Procedure

1. Move production from the primary to the secondary site by performing a *Swap* as follows:

a. Stop any applications that access the primary volumes to be taken over and unmount the filesystem from the OS.

b. Locate the replication to be taken over by navigating to the Hitachi Block Replication Details (Storage) (on page 882) on the secondary device. See How to view the status of a Hitachi Block storage device (on page 282).

   You will see that the replication's *Type* (displayed on the corresponding tile in the Hitachi Block Replications Inventory (on page 878)) is `Active Full Copy` and the *Swapped* state (displayed on the Hitachi Block Replication Details (Storage) (on page 882)) is `No`.

c. From the Hitachi Block Replication Details (Storage) (on page 882) click the **Swap** button.

   The Hitachi Block Replication Swap Wizard (on page 818) will appear warning that swapping can cause data loss, however as long as access to the primary volumes has been stopped, it will be safe to proceed.

d. Select a direction from the *"Direction"* dropdown. This is the intended final direction of the replication once the swap operation is complete.

e. Type the word `'SWAP'` into the **Confirm Swap** field and click **OK**.

   The Jobs Inventory (on page 493) is displayed and a new job will appear indicating that a *Swap Replication* operation is in progress. Click on the *Job Type* in the table to open the Job Details (on page 499) which list the log messages relating to the swap operation.

f. Return to the Hitachi Block Replication Details (Storage) (on page 882) and review the replication's status:

   - If the swap is successful then the *Swapped* state is set to `Yes`, indicating that the replication is now reversed (S-VOL to P-VOL) and is back in PAIR status. A *Swapped* status badge (see Monitor Status Badges (on page 530)) will also appear above the replication's mover on the Monitor Details (on page 521).
   - If the swap cannot be completed due to a P-VOL or data link fault then the *Swapped* state is set to `No` and *Suspend for Swap* state is set to `Yes`, indicating that the swap is not yet complete and is in SSWS status. Further action will be required on the primary block storage device or data link before the replication process can be re-established, but the secondary will be writeable.

   Note: The flow direction of a replication pair should ONLY be determined by referring to the **Summary** - *Swapped* field on the Hitachi Block Replication Details (Storage) (on page 882) for the secondary Block storage device. Primary and secondary volume information shown in the replications Session Log Details (on page 517) and associated Log Attachments Dialog (on page 519) should not be used to infer the flow direction following a swap.

g. Start any applications that access the secondary volumes that have been taken over and resume production at the secondary site.

2. Perform any maintenance and recovery tasks at the primary site, resolve any faults with the data link between sites, then go back to the Hitachi Block Replication Details (Storage) (on page 882) for the secondary to determine the status of the S-VOLs. Perform one of the following actions as appropriate:

a. If the replication is *Swapped* (S-VOL status = PAIR) then proceed with moving production back to the primary site when ready, as detailed in step 3 below.
b. If the replication is *Suspended for Swap* (S-VOL status = SSWS) then click the **Unsuspend** button. The swap operation will be completed as described above. Production at the secondary site can now continue with replication to the primary site in operation.
c. If the S-VOL status is some value other than PAIR or SSWS then you will need to run the following CCI command sequence from outside Protector to recover the replication pairing: **`pairsplit -R, pairsplit -S, paircreate`**

3. Move production back to the primary site when ready to resume normal operations by performing a *Swap* as follows:
   a. Stop any applications that access the secondary volumes to be taken back.
   b. Locate the replication to be taken back by navigating to the Hitachi Block Replication Details (Storage) (on page 882) on the secondary device. See How to view the status of a Hitachi Block storage device (on page 282).

      You will see that the replication's *Type* (displayed on the corresponding tile in the Hitachi Block Replications Inventory (on page 878)) is `Active Full Copy` and the *Swapped* state (displayed on the Hitachi Block Replication Details (Storage) (on page 882)) is `Yes`.
   c. From the Hitachi Block Replication Details (Storage) (on page 882) click the **Swap** button.

      The Hitachi Block Replication Swap Wizard (on page 818) will appear warning that swapping can cause data loss, however as long as access to the secondary volumes has been stopped, it will be safe to proceed.
   d. Type the word `'SWAP'` into the **Confirm Swap** field and click **OK**.

      The Jobs Inventory (on page 493) is displayed and a new job will appear indicating that a Swap VSP Replication operation is in progress. Click on the *Job Type* in the table to open the Job Details (on page 499) which list the log messages relating to the swap operation.
   e. Return to the Hitachi Block Replication Details (Storage) (on page 882) and review the replication's status:

      When the swap (takeback) is completed the *Swapped* state will be set to `No`, indicating that the replication is now normal (P-VOL to S-VOL) and is back in PAIR status. The *Swapped* status badge will disappear from above the replication's mover on the Monitor Details (on page 521).
   f. Start any applications that access the primary volumes that have been taken back and resume production at the primary site.

How to expand a journal

Before you begin

It is assumed that you have one or more journals defined on the storage system and visible in the Journals inventory. (If the journal is not visible, then it may be necessary to manually refresh the inventory.)

If the journal is currently used by one or more replication pairs in PAIR/PAIR status, then it will be necessary to pause those replication pairs for the expansion to take full effect. Once expanded, the pairs can be resynchronized to PAIR/PAIR status, if desired.

Only journals composed of DP-VOLs can be resized by the product.

Procedure

1. If necessary, pause any replications using the journal to be expanded.
   - Replications managed by the product can be paused as described in Hitachi Block Replications Inventory (on page 878) and Hitachi Block Replication Details (Storage) (on page 882).
   - Replications not managed by the product will need to be manually paused using an external tool.
2. Locate the journal to be expanded, by navigating to one of the following locations:
   - the Hitachi Block Journals Inventory (on page 863) for the storage system that hosts the journal
   - the Hitachi Block Journal Details (on page 864) for the journal
   - Hitachi Block Replication - Journals (on page 888) for the replication that uses the journal
3. From any of these locations, click the Expand button. The Figure 575 Hitachi Block Journal Expansion Dialog (on page 864) will appear.
4. Enter a new size for the journal into the New Journal Size field, and click OK.
5. Go to the Jobs Inventory (on page 493) to ensure that a journal expansion job has been initiated, and wait for it to complete.
6. If the journal expansion is unsuccessful, review the Logs Inventory (on page 509) to find out why. The journal expansion operation must be re-initiated by the user once the problem is resolved.
7. If desired, resume any replications using the journal that was expanded.
   - Replications managed by the product can be resumed as described in Hitachi Block Replications Inventory (on page 878) and Hitachi Block Replication Details (Storage) (on page 882).
   - Replications not managed by the product will need to be manually resynchronized using an external tool.

How to delete a journal

Before you begin

It is assumed that you have one or more journals defined on the storage system and visible in the Journals inventory. (If the journal is not visible, then it may be necessary to manually refresh the inventory.)

If the journal is currently used by one or more replication pairs, then it cannot be deleted.

Only journals composed of DP-VOLs can be deleted by the product.

Procedure

1. Locate the journal to be expanded, by navigating to one of the following locations:
   - the Hitachi Block Journals Inventory (on page 863) for the storage system that hosts the journal
   - the Hitachi Block Journal Details (on page 864) for the journal
2. From any of these locations, click the Delete button. A confirmation dialog will appear.
3. Review the information presented and click OK.
4. Go to the Jobs Inventory (on page 493) to ensure that one journal deletion job has been initiated per selected journal, and wait for them to complete.
   - If a journal deletion is unsuccessful, review the Logs Inventory (on page 509) to find out why. Each journal deletion operation must be re-initiated by the user once the problem is resolved.

How to adopt a replication into Protector

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258).
- The Protector Client software has been installed on the nodes that will act as proxies for both primary and secondary Hitachi Block storage devices.
- The primary and secondary storage devices have been set up as per the Protector requirements and prerequisites. Refer to Hitachi Block prerequisites (on page 21).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).
- Read About Hitachi Block replication adoption (on page 87) to understand how adoption works, its prerequisites, limitations and behaviour.

This task describes the steps to follow when adopting a replication that has been set up on the underlying hardware, outside of Protector. For more information, refer to About Hitachi Block based backup technologies (on page 63). The data flow and policy in this example are as follows:



Figure 64 Adopted TrueCopy Replication Data Flow

Table 20 Hitachi Block Replication Policy

Classification Type	Parameter	Value
Hitachi Block	Logical Devices	212418/100 212418/101

Note: If you want to add source volumes to a replication policy after it has been adopted, then the Adopt existing replication option in the Hitachi Block Replication Configuration Wizard (on page 398) must remain selected when you subsequently reactivate the data flow with the modified policy settings.

Operation Type	Parameter	Value	Assigned Nodes
Replicate	Run Options	N/A (TrueCopy is a continuous replication, so the Run option is ignored)	Primary Hitachi Block Device, Secondary Hitachi Block Device

To adopt a replication perform the following steps:

Procedure

1. Locate the nodes in the Nodes Inventory (on page 536) that will control the primary and secondary Hitachi Block Devices via CMD (Command Device) interfaces and check that they are authorized and online.

   These nodes are used by Protector to orchestrate replication of the primary LDEV(s) to the secondary LDEV(s) and are identified as the **Proxy Node** when creating the primary and secondary Hitachi Block Device nodes in the next step. These nodes are known as ISM nodes. The ISM nodes do not appear in the data flow.

2. Create new primary and secondary Hitachi Block Device nodes (unless ones already exists) using the Hitachi Block Device Node Wizard (on page 583) and check that they are authorized and online.

   The Hitachi *Block Device* node type is grouped under **Storage** in the Node Type Wizard (on page 539). See How to add a node (on page 261) and How to authorize a node (on page 262). The primary and secondary Hitachi Block Device nodes appear in the replication data flow as the source and destination nodes.

3. In the Policies Inventory (on page 670), create a new policy. See How to create a policy (on page 273)
   a. Add an Hitachi Block classification, select **Specify additional selections** and specify the LDEV(s) or Host Group of the primary volume(s) in the **Logical Devices** field.
   b. Add a **Replicate** operation. Select **Run on Schedule** and define a suitable schedule if a batch replication is being adopted.
   c. Click **Finish** to create the policy.

4. In the Data Flows Inventory (on page 355), create the replication data flow, corresponding to the one you want to adopt. See How to create a data flow (on page 228).
   a. Place the corresponding Hitachi *Block* source and destination nodes in the Data Flow workspace.
   b. Connect the two nodes using a **Batch** or **Continuous** mover, as appropriate to the replication type being adopted.
   c. Select the source node and assign the Hitachi *Block-Replicate* policy defined above.
   d. Select the destination node and assign the *Replicate* operation.
   The Hitachi Block Replication Configuration Wizard (on page 398) is displayed.
   e. Select the type of replication on the left hand side of the dialog and then the **Adopt existing replication** option on the right.

      Note: Refreshed Snapshot replications using Thin Image cannot by adopted.

      Only the fields required for identifying the adopted replication are enabled. Those that have been disabled will be populated automatically once the replication has been adopted. Refer to the table in About Hitachi Block replication adoption (on page 87) to understand how the policy and data flow attributes are interpreted during the adoption process.

   f. Enter the parameters required to identify the adopted replication on the hardware.

5. Compile and activate the data flow, checking carefully that there are no errors or warnings.

   See How to activate a data flow (on page 232).

6. Locate the active data flow in the Monitor Inventory (on page 520) and open its Monitor Details (on page 521).
   a. If you are adopting a batch replication, select the source node and click **Trigger Operation**.
   The Trigger Operation Dialog (on page 529) is displayed.

b. Select the replication operation and click **OK** to trigger it.

7. Go to the Logs Inventory (on page 509), identify the corresponding session and open the Session Log Details (on page 517) by clicking **View Session** to the left of one of the related message.
   If the adoption did not complete successfully, you will see one or more of the log messages listed in the table in About Hitachi Block replication adoption (on page 87). Make the necessary changes to the data flow and/or policies, recompile and activate the rules and try again.
8. Once the adoption process has completed successfully, go to the Hitachi Block Replication Details (Storage) (on page 882) for the corresponding replication and review the information to ensure that the desired replication has been adopted. See How to view the status of a Hitachi Block storage device (on page 282).

   The secondary LDEVs can be found under **Record Details** for the selected **Record Id.** which is listed under **Phase -** `Logical Devices`. The `Adopted` attribute will be set to *true*.

   If the wrong replication was adopted because the wrong *Mirror Unit* number was specified, it can be changed in the Data Flow Wizard (on page 361). Recompile and activate the rules and adopt the correct replication. The previously (erroneously) adopted replication will be left intact on the hardware but discarded by Protector.

How to dissociate a replication from Protector

An Hitachi block replication that has been defined within Ops Center Protector can be dissociated without it being removed from the underlying hardware.

Procedure

1. Go to the Hitachi Block Replications Inventory (on page 878) or Hitachi Block Replication Details (Storage) (on page 882) and locate the adopted replication that you want to dissociate.
2. Select the replication(s) and select **Dissociate** from the context menu.
3. A warning dialog is displayed. If you are sure you want to proceed then type the word `'DISSOCIATE'` then click **OK**.

   The replication entry is immediately removed from the list. However the dissociated data flow and policy definition will remain and must also be removed.
4. Go to the Data Flows Inventory (on page 355) and delete the dissociated data flows (assuming they are not involved in other policies still being managed by Protector).
5. Go to the Policies Inventory (on page 670) and delete the policies for the dissociated replications (assuming they are not involved in other policies still being managed by Protector).

How to snapshot disks of a Windows Server Failover Cluster (WSFC)

Before you begin

It is assumed that the following actions have been performed:

- The Protector Master software has been installed and licensed on a dedicated node. See Installation Tasks (on page 234) and License Tasks (on page 258)
- The Protector Client software has been installed on all nodes of the Microsoft Windows Cluster setup, which should be protected, and the nodes have been authorized in Protector.
- The Protector Client software has been installed on the node that will act as a proxy for the Hitachi storage device where the Windows Server Failover Cluster data is located.
- The block storage device has been set up as per the Protector requirements and prerequisites. Refer to Hitachi Block prerequisites (on page 21).
- The primary storage devices have been set up as per the Protector requirements and prerequisites. Refer to Hitachi Block prerequisites (on page 21).
- The disks are mapped to all cluster nodes of the Windows Server Failover Cluster
- All disks that should be snapshotted together should be same member of the same cluster resource group.

  Note: You cannot take a snapshot of both traditional clustered disks and cluster shared volumes at the same time. You can only take a snapshot of either one or more clustered disks or one or more Cluster Shared Volumes (CSV).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups. Refer to How to configure basic role based access control (on page 213).

This task describes the steps to follow when snapshotting clustered disks that reside on a Hitachi Block storage device. The data flow and policy are as follows:



Classification Type	Parameters	Value
File System	Node	File System application node hosting the cluster disks
	Cluster Resource Group	The cluster resource the clustered disks belong to
	Included Paths	The paths or mount points of the clustered disks

Operation Type	Parameters	Value
Snapshot	Mode	Hardware
	Hardware Type	Hitachi Block
	Run Options	Run on RPO
	RPO	1 day
	Retention	5 days

Procedure

1. Locate the source OS Host nodes for your Microsoft Windows Server Failover Cluster (WSFC) environment in the Nodes Inventory and check that the nodes are authorized and online. These nodes represent the Protector Clients installed on your Microsoft Windows Server Failover Cluster environment.
2. Create a new Fileserver node (unless a suitable one already exists) using the File System Node Wizard (on page 544). The File System node type is grouped under **Application** in the **Node Type Wizard**.
   a. Select one node which is part of your Windows Server Failover Cluster environment.
   b. Select which configuration you want the File System node to represent. In case of a standalone system, you will only have one option. In the case of a clustered system, you can choose if the application node should represent just this host or a clustered instance.
3. Locate the node in the Node Inventory that will control the Hitachi Block Device via a CMD (Command Device) where the Microsoft Windows Server Failover Cluster data is located. Check that the node is authorized and online.

   This node is used by Protector to orchestrate snapshot creation and is identified as the Proxy Node when creating the Hitachi Block Device node in the next step. This node is known as an ISM (Intelligent Storage Manager) node. The ISM node does not appear in the data flow.

4. Create a new Hitachi Block Device node (unless one already exists) using the Hitachi Logical Block Device Node Wizard (on page 609) and check that it is authorized and online.

   The Hitachi Block Device node type is grouped under **Storage** in the **Node Type Wizard**. Note that this node does not appear in a Failover Cluster snapshot data flow diagram but is identified when assigning the snapshot policy

5. Define a policy as shown in the table above using the Policy Wizard (on page 671), File System Classification Wizard (on page 679) and Snapshot Operation Wizard (on page 724).

   The File System classification is grouped under Application in the Policy Wizard.

6. Draw a data flow as shown in the figure above, that shows only the File System source node, using the Data Flow Wizard. At this stage, the snapshot icon is not shown.
7. Assign the Snapshot operation to the File System source node. The File System Snapshot policy will then be assigned automatically. The Block Snapshot Configuration Wizard is displayed.
8. Select the Storage Node corresponding to the Hitachi Block Storage device where the File System data is located. Then select a Snapshot Pool from one of the available snapshot or hybrid pools.
9. Leave the remaining Advanced Configuration options at their default settings, then click OK.

   Tip: Using Cascade mode (the default setting) will allow the data to be used for reverts after it has been mounted. The snapshot icon is now shown superimposed over the source node.

10. Compile and activate the data flow, checking carefully that there are no errors.
11. Locate the active data flow in the Monitor Inventory and open its Monitor Details. The policy will be invoked automatically to create a snapshot repeatedly according to the RPO specified in the policy. The policy can also be manually triggered from the source node in the monitor data flow.
12. Watch the active data flow via the Monitor Details to ensure the policy is operating as expected.

    You should periodically see:

    - Backup jobs appearing in the Jobs area below the data flow that show progress percentage, ending in Progress - Completed.
    - Information messages appearing in the Logs area below the data flow indicating rules activation and snapshot events.

13. Review the status of the Hitachi Block Device to ensure snapshots are being created. New snapshots will appear in the Block Snapshot Inventory periodically as dictated by the RPO of the policy. Old snapshots will be removed periodically as dictated by the Retention Period of the policy.

Chapter 4: Tasks

This chapter describes, step by step, the tasks that users will perform to install, configure and operate Ops Center Protector.

Tip: Refer to Button Icons (on page 296) to find the icons associated with the control names used in the step-by-step guides.

Access Control Tasks

This section describes access control tasks that users will perform with Ops Center Protector.

For further information, refer to:

- Access Control Concepts (on page 122)
- Access Control UI Reference (on page 309)

How to configure basic role based access control

Before you begin

You will need to have:

- A Protector account with *Default Administrator* ACP authority. You will already have a *<Username>@Master* login with this authority if you installed Protector on the Master node. If you do not have an account with this authority then you will need to request one from your Protector administrator.
- Knowledge of users and user groups who require access to Protector and their data protection roles and responsibilities.
- The details of any authentication services that you intend to use to authenticate Protector users (e.g. Active Directory, LDAP, RADIUS etc.)

Refer to Access Control Concepts (on page 122) and Access Control UI Reference (on page 309) for further information.

Protector implements RBAC to control what actions users can perform on which resources. The RBAC implementation is extremely flexible and can be configured to be as open or restrictive as an organization demands.

This procedure will allow you to get up and running quickly, however to fully utilize RBAC's features you will need to setup a more advanced RBAC implementation. Refer to How to configure advanced role based access control (on page 215) for details on how to do this.

Protector includes the following built-in access control objects:

- The *'default'* Resource Group that all Protector nodes are a member of by default.
- The Roles:
  - Protector *Admin* that can perform all activities.
  - Protector *Security Manager* that can perform all access control activities.
  - Protector *Operator* that can view all resources and perform restore activities.
- The Access Control Profile:
  - *Default Administrator* that can perform all activities on all (*default*) resources.
- The *'Master'* Authentication Space that represents the local authentication service on the Master node's OS.
- The following Access Control Profile Association (depending on the *UserName* of the account on the Master node specified when Protector was installed):
  - <*UserName*>@*Master* that represents a user that has *Default Administrator* privileges.

This topic explains how to implement a basic RBAC policy:

Procedure

1. Use a web browser to log on to the Protector user interface at: `https://<Master>`, where <*Master*> is the IP address or DNS name of the Master node.
   The Login Page (on page 309) will be displayed.
2. Enter the username `<UserName>@master` and the associated password to log in with *Default Administrator* privileges.
3. Click the **Access Control** link on the Navigation Sidebar (on page 291) to open the Access Control Dashboard (on page 310).
4. Create an Authentication Space that represents your organization's existing authentication service (see How to create an Authentication Space (on page 221)).
5. Create ACP Associations for each user, group or entire authentication space that requires access to Protector, using one of the built-in Access Control Profiles (see How to create an Access Control Profile Association (on page 226)).

   You can create your own ACPs or clone an existing ACP and make changes to the clone (see How to create an access control profile (on page 220) and How to clone an access control profile (on page 220)).
6. It is recommended that the default ACP Association <*UserName*>@*master* is replaced with your own ACP associations, using dedicated usernames created in your organization's domain.

   The default ACP Association cannot be deleted, but can be rendered unusable when the corresponding local Windows account is disabled. In the event that administrators are locked out from Protector due to access control configuration issues, this ACP Association is available as a way back in, by re-enabling the local Windows account.

Caution:

- The default ACP Association is generated automatically when Protector is installed, to enable initial configuration of access control features. This is based on the local Windows account specified during installation.
- The default *<Username>@Master* ACP association should be assigned to a user with the specific responsibility as the primary Protector administrator, to ensure security is not compromised.
- Access to the Master node should be strictly controlled to prevent malicious access to the Protector executables and associated configuration data.

How to configure advanced role based access control

Before you begin

You will need to have:

- A Protector account with *Default Administrator* authority. You will already have an *<Username>@master* login with this authority if you installed Protector on the Master node. If you do not have an account with this authority then you will need to request one from your Protector administrator.
- A good understanding of your organization's computing resources and the way they are managed and grouped into departments.
- Knowledge of where computing resource will be backed up to (i.e. the storage devices to be used).
- Knowledge of users and user groups who require access to Protector and their data protection roles and responsibilities.
- The details of any authentication services that you intend to use to authenticate Protector users (e.g. Active Directory, LDAP, RADIUS etc.)

Refer to Access Control Concepts (on page 122) and Access Control UI Reference (on page 309) for further information.

Protector implements Role Based Access Control (RBAC) to ensure that only those users with sufficient privileges can view or modify resources. The RBAC implementation is extremely flexible and can be configured to be as open or restrictive as an organization demands.

If you require only a basic RBAC implementation then refer to How to configure basic role based access control (on page 213)

Alternatively, custom roles and resource groups can be created that precisely control the nodes that are visible and the operations that can be performed on them. This topic explains how to plan and implement a custom RBAC policy:

Procedure

1. Identify the computing resources within your organisation, based on geographical, divisional, departmental, functional and project groupings.

These resources may be managed locally and/or centrally and this will also dictate how they are grouped together, for the purposes of access control when:

- Designing data protection policies and data flows
- Monitoring and reporting Protector performance
- Allocating and monitoring backup storage resources
- Auditing for compliance
- Administrating security and access controls
- Repurposing data for test and development
- Executing restore and disaster recovery procedures

For example, you might need to create the following resource groups in addition to the built-in *default* group:

(The names in this example are designed only to help illustrate how the RBAC objects are assembled into hierarchies).

- groupAccountsGlobal
- groupLegalGlobal
- groupHumanResourcesUS
- groupHumanResourcesUK
- groupDevelopmentUK
- groupProductionPrimaryUS
- groupProductionSecondaryUS

For guidance, refer to How to create a resource group (on page 219).

2. Identify the generic roles (not the individuals) required within your organisation for administering computing resources and the associated data protection processes.

   For example, you might need to create the following roles based on, or in addition to, the built-in roles:

   - roleBackupAdmin
   - roleComplianceAuditor
   - roleSecurityAdmin
   - roleDevelopmentLead
   - roleProtectorUser

3. Define precisely what activities each role should and should not be able to perform.

   Protector defines numerous Activity Groups that are cohesive collections of Activities. Normally a role would be assigned all activities in a group, however individual activities can be assigned if fine grain control is required. For example, the *Logs* Activity Group contains the following Activities which can be granted to a role en-mass or individually:

   - View Logs
   - Manage Logs

- Purge Audit Logs
- Verify Audit Logs

For guidance on configuring Roles and their associated Activities, refer to How to create a role (on page 219).

4. Create Access Control Profiles based on the Resource Groups and Roles identified above. These ACPs combine a Role with one or more Resource Groups.

   For example, it might be necessary to create the following ACPs, in addition to the built-in *Default Administrator* ACP:

   - acpAccountsBackupAdmin

     - to allow *roleBackupAdmin* access to *groupAccountsGlobal*

   - acpLegalBackupAdmin

     - to allow *roleBackupAdmin* access to *groupLegalGlobal*

   - acpDevelopmentBackupAdmin

     - to allow *roleBackupAdmin* access to *groupDevelopmentUK*

   - acpProductionBackupAdmin

     - to allow *roleBackupAdmin* access to *groupProductionPrimaryUS* and *groupProductionSecondaryUS*

   - acpDevelopmentUser

     - to allow *role*Protector*User* access to *groupDevelopmentUK*

   For guidance on associating Roles with Resource Groups, refer to How to create an access control profile (on page 220).

5. Identify how users will be authenticated by Protector.

   Protector supports a number of authentication protocols. If your organization has an established AD, LDAP or RADIUS authentication service or uses local accounts, then these can be used.

For example, it might be necessary to create the following user and group accounts:

- Donald McPhee has a UID (User ID) of *donald.mcphee* in the Active Directory authentication service *global.widgetdev.com*.

  An Authentication Space is created named *widgetdev* that refers to that AD service. He logs into Protector with the UPN (User Principal Name) *donald.mcphee@widgetdev*.

- Pete Traynor has a UID of *traynorp* in the local OS Account on the Protector node *WIN7-PCEA45*.

  An Authentication Space is created named *WIN7-PCEA45* that refers to that node. He logs in with the UPN *traynorp@WIN7-PCEA45*.

- Sarah Dean has a UID of *svpdean* in the RADIUS service *uk.widgetdev.com*.

  An Authentication Space is created named *uk.widgetdev* that refers to that RADIUS service. She logs into Protector with the UPN *svpdean@uk.widgetdev*.

- The contract development team members are in a user group that has a UID of *devteam* in the LDAP authentication service *datadevs.biz*.

  An Authentication Space is created named *datadevs* that refers to that LDAP service. They log in using the UPN *devteam@datadevs*

For guidance, refer to How to create an Authentication Space (on page 221).

6. Associate authenticated users and user groups with Access Control Profiles (i.e. Roles and Resource Groups) so that those users are able to log on to Protector, access the resources they need and carry out the activities their roles allow.

   An individual user can be associated with more than one ACP, and an ACP can be assumed by more than one user.

   For example the following ACP Associations are required:

   - *donald.mcphee@widgetdev* and *svpdean@uk.widgetdev* are authorized to perform the activities defined by *acpDevelopmentBackupAdmin* on its associated resources.
   - *svpdean@uk.widgetdev* is, in addition, authorized to perform the activities defined by *acpProductionBackupAdmin* on its associated resources.
   - The entire database development team *devteam@datadevs* are authorized to perform the activities defined by *acpDevelopmentUser* on its associated resources.
   - *traynorp@WIN7-PCEA45* is authorized to perform the activities defined by *acpAccountsBackupAdmin* and *acpLegalBackupAdmin* on its associated resources.

   For guidance on authorizing users with their respective Roles and Resource Groups, refer to How to create an Access Control Profile Association (on page 226).

7. It is recommended that the default ACP Association *<username>@master* is replaced with your own ACP associations, using dedicated usernames created in your organization's domain.

Caution:

- The default ACP Association is generated automatically when Protector is installed, to enable initial configuration of access control features. This is based on the local Windows account specified during installation. Best practice states that local accounts should be disabled on the Master to reduce security vulnerabilities.
- The default <*Username*>@*Master* ACP association should be assigned to a user with the specific responsibility as primary Protector administrator, to ensure security is not compromised.
- Access to the Master node should be strictly controlled to prevent malicious access to the Protector executables and associated configuration data.

How to create a resource group

Before you begin

Refer to How to configure advanced role based access control (on page 215) which describes how resource groups are used in configuring access control.

To create a resource group:

Procedure

1. From the Access Control Dashboard (on page 310) click **Manage Resource Groups** to open the Access Control Resource Groups Inventory (on page 344).
2. Click the **Create new item** tile to open the Access Control Resource Group Wizard (on page 345).
3. Enter a **Name** for the resource group and a **Description**, then click **Next**.
4. Select the resources to be included in the resource group from the left-hand list by clicking on each resource.
   Each resource selected is added to the right-hand list. Resources can be removed from the right-hand list by selecting them there.
5. Click **Finish** to close the wizard and return to the inventory.

How to create a role

Before you begin

Refer to How to configure advanced role based access control (on page 215) which describes how roles are used in configuring access control.

To create a role:

Procedure

1. From the Access Control Dashboard (on page 310) click **Manage Roles** to open the Access Control Roles Inventory (on page 336).
2. Click the **Create new item** tile to open the Access Control Role Wizard (on page 337).

3. Enter a **Name** and **Description** for the role, then click **Next**.
4. Select the activity groups to apply to the role by clicking the checkbox to the left of the activity group names.
5. To apply individual activities, click the **+** to the left of the activity group name to expand the group, then select the required activities by clicking the checkbox to the left of the activity names.
6. Click **Finish** to close the wizard and return to the inventory.

How to create an access control profile

Before you begin

Refer to How to configure advanced role based access control (on page 215) which describes how access control profiles are used in configuring access control.

Ensure that the necessary resource groups and roles have been created (see How to create a resource group (on page 219) and How to create a role (on page 219)).

To create an access control profile:

Procedure

1. From the Access Control Dashboard (on page 310) click **Manage ACPs** to open the Access Control Profiles Inventory (on page 330).
2. Click the **Create new item** tile to open the Access Control Profile Wizard (on page 331).
3. Enter a **Name** and **Description** for the role, then click **Next**.
4. Select th required **Role** from the menu, then click **Next**.
5. Select the resource groups to be included in the access control profile from the left hand list by clicking on each resource group.
   Each resource group selected is added to the right hand list. Resource groups can be removed from the right-hand list by selecting them there.
6. For each resource group included in the ACP, set the **Access Level** in the dropdown control to the right of the resource group in the right hand list.

   The access level controls the visibility of backups of nodes in the resource group.
7. Click **Finish** to close the wizard and return to the inventory.

How to clone an access control profile

Before you begin

Refer to How to configure advanced role based access control (on page 215) which describes how access control profiles are used in configuring access control.

Note: A cloned access control profile is a point-in-time copy as opposed to inheritance of a parent access control profile. If you modify the original, none of those changes will be reflected in the clone.

Cloning is a way of creating a new ACP, based on an existing ACP. To clone an access control profile:

Procedure

1. From the Access Control Dashboard (on page 310) click **Manage ACPs** to open the Access Control Profiles Inventory (on page 330).
2. Select the tile for the ACP you want to clone, then click **Clone** in the menu above. A clone of the ACP will be added to the inventory having the same name but with (clone) appended.
3. To rename the clone refer to How to edit an access control profile (on page 221).

How to edit an access control profile

You can make changes to an existing access control profile as follows:

Procedure

1. From the Access Control Dashboard (on page 310) click **Manage ACPs** to open the Access Control Profiles Inventory (on page 330).
2. Click on the name of the ACP you want to edit. The Access Control Profile Details (on page 335) opens, showing the associated role and resource groups.
3. Click on the Edit button in the top right corner of the page. The Access Control Profile Wizard (on page 331) opens.
4. Edit the parameters as required, clicking **Next** and **Previous** to locate the items to edit.
5. When you have finished editing, go to the final page of the wizard and click **Finish**. The wizard is closed and the details page is displayed showing the updated parameters.

How to create an Authentication Space

Before you begin

Refer to How to configure basic role based access control (on page 213) or How to configure advanced role based access control (on page 215) which describes how Authentication Spaces are used in configuring access control.

Protector communicates with an authentication server via a single proxy node, which is specified when the access control Authentication Space is created. Ensure that the following prerequisites are met before you configure an AD, RADIUS or LAPD Authentication Space in Protector:

- The Protector proxy (which can be a Client or Master node) connecting to the authentication server is registered with it and any prerequisites listed in the authentication server documentation are met.
- The authentication server is not blocked by any firewalls.
- The configuration parameters for type of authentication server selected are known. See Access Control Authentication Space Wizard (on page 321) for what is required for each server type.

To create an Authentication Space:

Procedure

1. From the Access Control Dashboard (on page 310) click **Manage Authentication Spaces** to open the Access Control Authentication Spaces Inventory (on page 319).
2. Click the **Create new item** tile to open the Access Control Authentication Space Wizard (on page 321).
3. Enter a **Name** and **Description** for the Authentication Space, then click **Next**.

   Note: For Active Directory, the **Name** must be the AD Domain Name.

4. Select the type of Authentication Space you require from the list on the left of the wizard. The parameters appropriate to the Authentication Space type selected are displayed on the right of the wizard. All Authentication Space types require a **Proxy** to be selected, (except **OS Accounts**, which require an **Authentication Node**) that actually holds the account information.
5. Enter the parameters required for the selected Authentication Space type, then click **Finish**.

How to configure an LDAP authentication space

Before you begin

Ensure the LADPv3 server is correctly configured as per the instructions supplied with the LDAP software.

Configure a Linux based Protector (Master or Client) node with a connection to the LDAP server to act as a proxy. If you have a Windows Master, then you must select a Linux Client as a proxy. In this example the node `Client5RHEL` will be nominated as the proxy.

If using LDAP over TLS, place the TLS CA certificate file on the Protector proxy node.

Note: Protector supports LDAP and LDAP over TLS (LDAPS) protocols. We recommend that initial communication checks are performed without TLS (using LADP protocol). Once correct operation has been confirmed, change to TLS (using LADPS protocol).

This is an illustrative example only. LDAP configurations vary considerably between organisations so the output for your environment may be quite different to that shown here. It is assumed that the person performing this configuration is well versed in LDAP and the way it is configured in your organization:

Procedure

1. Examine the configuration of the LDAP server, using one of the following methods to ensure you can log into the LDAP server (preferably via the Protector proxy node to confirm the connection is working). Make a note of the *Base DN* and *User/Group DNs* listed in the output:
   a. Either connect to the LDAP server via a web based interface.

b. Or connect via a command shell using the following Linux command. Consult the Linux man page for full syntax:

```
ldapsearch
    -D "uid=admin,dc=mydomain,dc=com"
    -w pa55w0rd
    -H ldap://mydomain.com
    -b "dc=mydomain,dc=com"
    -s sub "(objectClass=*)"
```

Where the `mydomain.com` LDAP server's administrator *UID* is `admin` and the password is `pa55w0rd`.

As an example, the output from **`ldapsearch`** should include the following configuration information. The ***highlighted*** parts will be required in the steps that follow:

- The *Base DN*:

```
# mydomain.com
dn: dc=mydomain,dc=com
objectClass: top
objectClass: dcObject
objectClass: organization
o: mydomain.com
dc: mydomain
```

- The Administrator's *DN* (used as the *Bind DN*):

```
# admin, mydomain.com
dn: cn=admin,dc=mydomain,dc=com
objectClass: simpleSecurityObject
objectClass: organizationalRole
cn: admin
description: LDAP administrator
```

- User *DN*s and *UID*s:

```
# Joe Bloggs, mydomain.com
dn: cn=Joe Bloggs,dc=mydomain,dc=com
givenName: Joe
sn: Bloggs
cn: Joe Bloggs
uid: jbloggs
uidNumber: 1000
gidNumber: 500
homeDirectory: /home/users/jbloggs
loginShell: /bin/bash
objectClass: inetOrgPerson
objectClass: posixAccount
objectClass: top
```

- Group *DN*s and *memberUid*s:

```
# Managers, Groups, mydomain.com
dn: cn=Managers,ou=Groups,dc=mydomain,dc=com
gidNumber: 501
objectClass: posixGroup
objectClass: top
cn: Managers
memberUid: jbloggs
memberUid: tsmith
memberUid: mjones
...
```

2. Configure the parameters in the Access Control Authentication Space Wizard (on page 321) as follows:



Note: The values entered are dependent on the particular LDAP configuration. Be sure to check the output generated by **`searchldap`** for your configuration to obtain the correct values.

   a. On the **Configure authentication type** page, select **LDAP** authentication, then enter the following parameters:

      - **Proxy**: `Client5RHEL`
      - **Server URI**: `ldaps://mydomain.com`
      - **Server Port**: If not using the default value, enter a port number.
      - **Base DN**: `dc=mydomain,dc=com` from the **`searchldap`** output:

```
# mydomain.com
...
dn: dc=mydomain,dc=com
...
```

      - Select **Bind using specified account**

- **Bind Account DN**: `cn=admin,dc=mydomain,dc=com` from the **`searchldap`** output:

```
# admin, mydomain.com
...
dn: cn=admin,dc=mydomain,dc=com
...
```

- **Bind Account Password**: `pa55w0rd`

b. If using the LDAPS protocol, click **TLS Configuration** and configure the **TLS Request Certificate Check** method, **TLS CA Certificate Directory** and **TLS CA Certificate File** to use.

c. Click **Advanced Configuration** and enter the following parameters based on the given **`searchldap`** output:

- **Person Filter**: `(objectClass=inetOrgPerson)`

```
# Joe Bloggs, mydomain.com
...
objectClass: inetOrgPerson
...
```

- **Group Filter**: `(objectClass=posixGroup)`

```
# Managers, Groups, mydomain.com
...
objectClass: posixGroup
...
```

- **Group Strategy**: select **Groups know users**

```
# Managers, Groups, mydomain.com
...
memberUid: jbloggs
memberUid: tsmith
memberUid: mjones
...
```

- **Group Member Attribute**: `memberUid`

```
# Managers, Groups, mydomain.com
...
memberUid: jbloggs
...
```

- **Group Member Type**: select **Member value contains a UID**

```
# Managers, Groups, mydomain.com
...
memberUid: jbloggs
...
```

```
# Joe Bloggs, mydomain.com
...
uid: jbloggs
...
```

- The following attribute values are evident from the output:
  - **CN Attribute**: `cn`
  - **DN Attribute**: `dn`
  - **UID Attribute**: `uid`

3. Click **Finish** to close the wizard.

How to create an Access Control Profile Association

Before you begin

Refer to How to configure basic role based access control (on page 213) which describes how access control profile associations are used in configuring access control.

Ensure that the necessary Access Control Profiles and Authentication Spaces have been created (see How to create an access control profile (on page 220) and How to create an Authentication Space (on page 221)).

To create an access control profile association:

Procedure

1. From the Access Control Dashboard (on page 310) click **Manage ACP Associations** to open the Access Control Profile Associations Inventory (on page 312).
2. Click the **Create new item** tile to open the Access Control Profile Association Wizard (on page 313).
3. Enter a **Name** and **Description** for the ACP Association.
4. Select the type of association you require from the list on the left of the wizard:
   - **User** - associates the specified user with the selected ACPs.
   - **Group** - associates all users in the specified group with the selected ACPs.
   - **Authentication Space** - associates all users in the specified Authentication Space with the selected ACPs.

   The parameters appropriate to the ACP association type selected are displayed on the right of the wizard.
5. Enter the parameters required for the selected ACP association type, then click **Next**.

6. Select the ACPs to be included in the ACP association from the left-hand list by clicking on each ACP.
   Each ACP selected is added to the right-hand list. ACPs can be removed from the right-hand list by selecting them there.
7. Click **Finish** to close the wizard and return to the inventory.

How to view the access control settings summary

Before you begin

Refer to How to configure basic role based access control (on page 213).

Ensure that the necessary access control profile associations have been created (see How to create an Access Control Profile Association (on page 226)).

To view a summary of the current access control settings for each Protector user or group:

Procedure

1. From the Access Control Dashboard (on page 310) click **Manage ACP Associations** to open the Access Control Profile Associations Inventory (on page 312).
   You can also access the summary information from the Access Control Authentication Spaces Inventory (on page 319), Access Control Profiles Inventory (on page 330), Access Control Roles Inventory (on page 336) or Access Control Resource Groups Inventory (on page 344).
2. Open the drop down menu in the Navigation Breadcrumbs (on page 290) by clicking the  button and select **Summary** from the menu.
   The Access Control Summary (on page 348) is displayed.
3. Click on the **[>]** to the left of the User or Group of interest to view its related ACPs, Role, Activity Groups,Activities, Resource Groups and Resources.

How to edit object permissions

Permissions control if an object (e.g. a data flow, schedule etc.) is visible to, or modifiable by specific users.

Note:

For normal creation of objects (e.g. policies, dataflows, schedules and store templates), the creating user is given Read/Write access, allowing that user to see and change the object. Users having the RBAC *Override Ownership Permissions* privilege can also see and edit the object. Nobody else will be able to view the object unless granted access.

Normal users (i.e. those without the RBAC *Override Ownership Permissions* privilege) are prevented from removing all permissions, although they can still remove their own access rights. Only users with the RBAC *Override Ownership Permissions* privilege can remove all permissions.

To edit the permissions for an object:

Procedure

1. Go to the *Details* or *Inventory* page of the object for which you want to edit the permissions.
2. Click **Edit Permissions** in the top right of the page.
   The Access Control Permissions Inventory (on page 349) will be displayed, showing the users and groups that have read and write access to the object.
3. You can then do one of the following:
   - Add a new user or group permission by clicking the **Create New Item** tile.
   - Edit an existing permission by clicking the user or group name on a tile.
   - Remove an existing permission by selecting a tile and clicking **Remove**.

   Editing does not allow the user or group name to be changed. To do this, remove the existing permission, then add a new one.

   When you edit or create a permission, the Access Control Permissions Inventory (on page 349) opens.
4. For new permissions, Select the type of permission you require from the list on the left of the wizard:
   - User - to grant a single user permission
   - Group - to grant a group of users permission

   The parameters appropriate to the permission type selected are displayed on the right of the wizard.
5. Check **Write Access** if you want the user or group to be able to modify the object. Read access is automatically granted to any user or group added to the permissions inventory for that object.
6. Click **Finish** to close the wizard and return to the inventory.

Data Flow Tasks

This section describes data flow configuration tasks that users will perform with Ops Center Protector.

Refer to Data Protection Workflows (on page 130) for detailed descriptions of specific Repository, Hitachi Block data protection scenarios.

For further information, refer to:

- Data Flow Concepts (on page 50)
- Data Flows UI Reference (on page 355)

How to create a data flow

Before you begin

Ensure that the policies that you want to assign have been defined, see How to create a policy (on page 273).

The following procedure describes how to create a simple one-to-one data flow. More complex data flows involving one-to-many and cascaded topologies can be constructed by following the same general approach:

Procedure

1. Click the **Data Flows** link on the Navigation Sidebar (on page 291) to open the Data Flows Inventory (on page 355).
2. Click the **Create new item** tile to open the Data Flow Wizard (on page 361).
3. Enter a **Name** and **Description** for the data flow, then click **Next**.
   The next wizard page is displayed with a blank workspace.
4. Drag a source node from the **Nodes** or **Node Groups** list onto the data flow workspace.
   The node is displayed on the workspace with a grey box around it showing it is selected. The available **Policies** appear next to the workspace. If a policy contains operations that can be performed locally to the node, without the need for a separate destination node (e.g. local snapshot operations), then these will be displayed directly below the policy.
5. Select the policies and/or local operations that you want to assign to the source node by checking all those that apply in the **Policies** listed to the right of the workspace.
   - If a policy is selected that requires a destination node and corresponding operation assignment to complete it, then a warning triangle icon will appear next to the source node, indicating that the node has an incomplete policy assigned to it. Completing the policy assignment is described in the steps that follow.
   - If an operation is selected, then an operation properties dialog will be displayed. Enter the required operation properties in the dialog and click **OK**.

     If you choose not to define the operation properties now (by clicking **Cancel**), they can be configured later by clicking the **Configure Operation Properties** button displayed below the respective operation in the **Policies** area to the right of the data flow workspace.

     After the operation properties have been applied, they can be edited by clicking the **Edit Operation** button in the operation summary box next to the operation name.

     If a node has a snapshot operation assigned to it, a snapshot icon will appear in the bottom right corner of the node.
6. Now, place and connect the destination node. There are two methods for connecting nodes:
   - Drag the destination node from the **Nodes** list, passing over the source node that you want it to connect to, then drop the destination node where you want to place it. A connection is created between the destination and source node. Now select the destination node.
   - Place the destination node on the workspace then refer to How to connect nodes on a data flow (on page 231).

   The destination node is displayed on the workspace with a grey box around it showing it is selected. The available **Policies** appear to the right of the workspace. If a policy contains operations that can be performed by the destination node (e.g. remote replication operations), then these will be displayed below the policy.

7. With the destination node selected, choose the operations that you want to assign to it by checking all those that apply in the **Policies** area to the right of the workspace. Note that the policy checkbox cannot be selected by the user. When an operation is selected, an operation properties dialog will be displayed. Enter the required properties in the dialog and click **OK**.
   If an operation is selected that completes a policy previously selected on the source node, then the warning triangle icon will be removed from the source node, indicating that the node now has a completed policy assigned to it. The image below shows a source node with a remote policy assigned (*myReplication*) that is completely specified (the operation *Mirror (Replicate)* has been assigned to the destination node). A local operation (*mySnapshot*) has also been assigned to the source node.

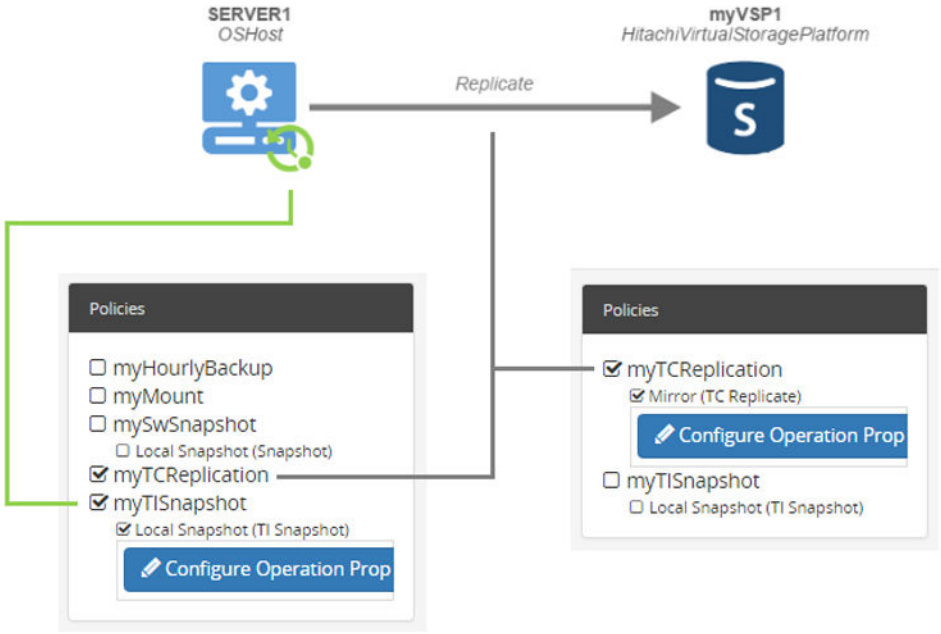


Figure 65 Source node with a remote policy assigned (completed assignment)

8. Select the connection between the source and destination nodes to display the **Routed Policies** and **Mover Settings** to the right of the workspace.
   a. The **Routed Policies** area lists the policies being routed along the selected connector.
   b. Select the **Type** of mover to be used in the **Mover Settings**.

      Note: The Data Flow Wizard (on page 361) only prevents some incorrect mover and operation combinations from being constructed. The *Rules Compiler* will however generate warnings or errors for incorrect combinations. Ensure the correct mover type is used with a given operation when creating data flows.

   c. Optionally enter a **Label** for the connection.
   d. Turn network compression on or off with **Enable network data compression**.

e. For *Host Based* policies only, click **Bandwidth Settings** to open the Mover Bandwidth Settings Dialog (on page 375), then set the times and days for **Default Speed**, **High Speed** and **Low Speed** network utilization by clicking the required cells.

9. When you have finished drawing the data flow and assigning policies, click **Finish**.

How to connect nodes on a data flow

Before you begin

Create a data flow as described in How to create a data flow (on page 228).

Nodes can be connected on a data flow as follows:

Procedure

1. Drop the two nodes that are to be connected on the data flow canvas.
2. Select the node where data will flow from.
3. Click the **Connect To** button in the top left of the canvas.
   A dashed line will appear connected to the selected node at one end and the mouse cursor at the other.
4. Move the mouse cursor to the node where data will flow to and click to connect the two nodes.
   A line will be drawn from the first node, to the second node with an arrowhead indicating the direction of data flow.
5. If the mover is not already selected, click on it to view and set the **Routed Policies** and **Mover Settings**.

How to apply a policy to nodes on a data flow

Before you begin

Create a data flow as described in How to create a data flow (on page 228).

Policies are applied to nodes on a data flow as follows:

Procedure

1. Select the source node on the data flow canvas.
2. In the **Policies** area to the right of the canvas select each policy that needs to be applied to the source node.
3. Click on the mover that routes the policy to the destination node to view the **Routed Policies**.
4. Select the destination node on the data flow canvas.
5. In the **Policies** area to the right of the canvas select each operation that needs to be applied to the destination node.

   Only the individual operations can be applied; not the policy.

   An **Operation Properties** dialog appropriate to the destination node and operation type will be displayed. For example the Hitachi Block Replication Configuration Wizard (on page 398) is displayed when applying a *Replication* operation to a Hitachi *Block* node.

6. Configure the operation properties as required then click **OK**.
7. Finally click **Finish** to progress to the next page of the Data Flow Wizard (on page 361).

How to activate a data flow

Before you begin

Ensure the data flows that you want to compile have been correctly defined (see How to create a data flow (on page 228)), that the required polices have been assigned and that no significant warning icons[1] are displayed on nodes in the data flow diagrams.

1. There is no reason why all policy operations must be applied in all cases. Warning icons may therfore be present, but they may indicate a warning, not an error.

To compile a data flow and activate the resulting rules:

Procedure

1. Click the **Data Flows** link on the Navigation Sidebar (on page 291) to open the Data Flows Inventory (on page 355).
2. Select the data flows that are to be compiled by clicking in the selection icon in top left corner of the corresponding tiles.

   Although it is possible to compile multiple data flows in one go, it may be easier to initially compile one at a time and rectify any compilation errors, before compiling all data flows and distributing rules in one operation.

   Note: Activate data flows in batches not exceeding 20 data flows at a time. Activating more than this simultaneously can result in longer activation times.
3. Click **Activate** above the inventory.
   The Activate Data Flow Dialog (on page 357) is displayed and the selected data flow(s) start compiling. After a short time the results of the compilation process are displayed with a message indicating that the compilation process succeeded or failed.
4. If the compilation succeeds then click **Activate** to update the rules on the affected nodes.
5. If the compilation fails then the **Activate** button will remain disabled. Examine the compiler output to locate the cause of the failure, rectify the data flow and/or policy then recompile.

General Tasks

This section describes general tasks that users can perform with Ops Center Protector.

How to navigate to a page using the breadcrumbs

You can find your location within the hierarchy of web pages in the Protector user interface at any time and navigate quickly to a parent or related page (refer also to User Interface Page Layout (on page 285)). To navigate using the breadcrumbs:

Procedure

1. Locate the Navigation Breadcrumbs (on page 290) just below the Main Banner (on page 286) at the top of the current web page.
   The path to the current page is displayed with each parent page name separated by '/' . The current page is displayed in black. If related pages are available then a dropdown menu button will be displayed to the right of the page name.
2. To navigate to a parent page, click on the page name in the path.
3. To navigate to a related page, click on the button to show a menu of related pages, then click on the page you require.

How to select a single item in an inventory

Items displayed in an inventory must be selected before an action can be applied to them using the buttons displayed above (see Inventory Page (on page 293)). To select individual items:

Procedure

1. Locate the selection button for the inventory item you want to select.
   - In tile view the selection button is displayed in the top left corner of each tile (see Tile Control (on page 295)).
   - In table view the selection button is displayed to the left of each item.
2. Click the selection button so that a check mark appears.
   You can select one or more items at once, although some actions can only be applied to a single item.
   The available action buttons above the inventory will be enabled or disabled as appropriate to the number of items selected.

How to select all items on an inventory page

On some inventory pages it is possible to quickly select all tiles or table entries (see Inventory Page (on page 293)). To select all items on a page:

Procedure

1. Go to the menu items just above the inventory tiles or table and click the **Select All** check box.
   A check mark will appear against all tiles or table items in the inventory.
2. Click on the required command to apply the action to all items simultaneously.



Note: Select All only selects the items on the currently displayed page. If multiple pages are available, items on the non-displayed pages will not be selected.

How to set a filter on an inventory

Many inventory pages have a set of filters that can be applied to limit the items being displayed. The filter controls are normally displayed to the right of the inventory items (see Inventory Page (on page 293)). To filter an inventory:

Procedure

1. If the filter terms are not displayed to the right of the inventory page, click the **Show Search** button in the top right corner of the page.
2. Select the filter terms in the filter controls by:
   - Selecting the term from a dropdown menu
   - Typing in text directly, or
   - Selecting one or more items from a list

   Most inventory pages will update in real time as the filter terms are selected, however the inventories accessed from the Restore Inventory (on page 763) require the **Search** button to be clicked before the filter is applied.
3. All filter terms will be cleared when you navigate away from the page, unless you return to it using your browser's **Back** button.
4. Individual filter terms can be cleared by:
   - Selecting the Any/All filter term if available
   - Manually deleting the text in the filter control, or
   - Deselecting the items in a list

How to change the UI settings

The locale, language and other factors that affect the way in which Protector displays data can be configured via the Settings Wizard (on page 287) (refer also to User Interface Page Layout (on page 285)). To adjust UI settings:

Procedure

1. On the Main Banner (on page 286), click the **Settings** icon.
   The Settings Wizard (on page 287) opens.
2. Adjust the user interface settings to those required and click **Save**.
   The settings will take effect immediately.

Installation Tasks

This section describes the software installation and initial configuration tasks.

Settings required when using a virus detection program

Before you begin

If a virus detection program accesses database-related files used by Protector, an error might occur for reasons such as I/O delays or file locks. To prevent such errors while Protector is running, exclude the following directories from the targets scanned by the virus detection program:

- <install_path>/Protector, if any advanced configuration has been carried out to split trace or logs to different locations this should be added also.
- The HORCM installation directory C:\HORCM or /HORCM
- The location of the ISM metadata, this is defined at creation time by the user.
- The location of the Repository folders, this is defined at creation time by the user

For details on the directories of other Ops Center products to exclude from the scanning targets, see the manual for each product.

How to install/upgrade Protector on Windows and Linux or AIX

Before you begin

Caution: Thin Image snapshot operations default to *Provisioned using floating device*. In versions prior to 6.5, this setting would automatically fallback to *Fully provisioned* if the hardware storage device did not support floating devices. With the introduction of *Cascade mode* snapshots in version 6.5 (now the default mode), this automatic fallback has been removed. Consequently:

- Data flows created in versions prior to 6.5 that contain Thin Image snapshot operations, with cascade mode is enabled, will now fail if the underlying hardware does not support *floating device*. These data flows must be manually reconfigured to use the correct provisioning type. Please note that if these data flows also contain replication operations, then these will be re-evaluated when the modified data flows are re-activated.
- Data flows created in versions prior to 6.5 that contain Thin Image snapshot operations, with cascade mode enabled, will now fail if the P-VOLs have any pre-existing non-cascade mode snapshots. If cascade mode is to remain enabled then any non-cascade mode snapshots must be deleted.
- For data flows created in version 6.5 or later, if the hardware storage device on which Thin Image snapshot operations are performed does not support *Floating device* and *Cascade mode*, then Thin Image operations will fail, unless the appropriate settings are selected in the respective data flows.

Caution:

- Refer to the Protector support matrices at https://www.hitachivantara.com/en-us/products/data-protection/data-instance-director.html#tech-specifications before attempting installation, to ensure that you understand the infrastructure requirements, available functionality and upgrade paths for your particular environment.
- If you intend to use Protector with Hitachi storage hardware then refer to the following before proceeding:
  - Hitachi Block prerequisites (on page 21)
  - Refer section "Generation 1 Hitachi Content Platform prerequisites" in User Guide.

Note:

When upgrading, it is highly recommended (and often necessary) to upgrade all nodes to the same version. The recommended order in which to upgraded is:

1. Internet Connected Nodes. If ICNs are not upgraded first, they may not be able to be 'push upgraded' through the UI. ICNs may appear to go offline until the master is upgraded.
2. Master Node(s).
3. Clients acting as ISMs, controlling Hitachi Block and File storage devices.
4. Clients acting as Data Destinations, such as Repositories.
5. Clients acting as Data Sources, such as application servers.

Note:

- For a new installation, the Master node must be installed before any Client nodes.
- When installing a new Master, you must select the account to use for initial Protector log-on. This account must be a local OS account which is able to successfully log in to the machine. These credentials will be input to the Protector UI to enable the initial access control configuration to be performed.
- Read the release notes shipped with the installer to ensure that the currently installed version can be upgraded to the new version. It may be necessary to upgrade to an intermediate version first and perform additional actions prior to and after upgrading.
- You may need to create an exception in your anti-virus software when installing.
- DO NOT upgrade while replications are in the process of pairing. Any active replications must be in the paired state before upgrade is carried out.
- It is recommend that you perform an upgrade only after currently active backups have been completed.
- Before upgrading, unmount any mounted snapshots and replications.

- It is recommended to backup /protector folder to avoid any issue during OS upgrade.
- If you uninstall a Client node, you will need to delete it from the Nodes Inventory before subsequently reinstalling it. This ensures that the Master node regenerates new identifiers for that node. If you do not do this then the reinstalled Client will not be recognized.
- Operating System Specific Behaviour:

OS	Note
Linux	We recommend that Linux source nodes have a Logical Volume Manager (LVM) on each volume group.
Linux and AIX	Ensure you have execute permissions using the command **`chmod 755`**, before running the installer. A minimum of 10 GB of free space is required in the 'unused' portion of the volume that is to be backed up. This is in addition to the space required for the allocated storage area. For example, if 100 GB of usable storage is required, then the total disk size will be 110 GB (100 GB of usable storage and 10 GB of unused storage for snapshot administration).
AIX	It is only possible to install the Protector Client on an AIX node.

All nodes that will participate in a backup data flow need to have Ops Center Protector installed. A node used only to access the web based user interface does not need to have any components installed on it.

Procedure

1. Locate and run the installer appropriate for the target OS and hardware architecture.

   Tip: Protector can be installed directly from the command line if required. Refer section "Installing and upgrading Protector from the command line" from **User Guide** for details of the installer's command line options.

   The installer filename has the following format:

   Protector-**R***m.n-m.n.n.nnnnn-ostype-proctype*

where:

- *m.n-m.n.n.nnnnn* - is the version and build number
- *ostype* - is the target operating system type:
  - **`Win`**
  - **`Linux`**
  - **`AIX`**
- *proctype* - is the target processor type:
  - **`x64`**

Examples:

- Windows: `Protector-7.x.x.x+20231124-094611-d960e0b3-windows-x64.exe`
- Linux: `Protector-7.x.x.x+20231124-094611-d960e0b3-linux-x64.run`
- AIX: `Protector-7.x.x.x+20231124-094611-d960e0b3-aix-x64.run`

The **Setup** wizard will be launched if a GUI shell is available. If not then the same information will be presented using the text mode shell.

2. If a previous installation of Ops Center Protector is found, the installer will prompt you to upgrade or abort the installation.
   - Click **Yes** to upgrade the existing Protector installation on this node.

     Tip: When upgrading along a supported version upgrade path, any existing data flows, policies, schedules etc. will be preserved. If any further actions are required post upgrade, then these will be described in the Release Notes shipped with the new version.
   - Click **No** to exit the installer wizard, in which case no changes will be made to the current installation.
3. When the **Setup** wizard appears, a welcome message is displayed. Click **Next** to begin the installation.
4. When prompted, read the License Agreement.
   - Select **I accept the agreement** if happy to proceed and click **Next**.
   - Select **I do not accept the agreement** if not happy to proceed and click **Next**. The installation will be aborted and no changes will be made to the machine.

   License keys are entered once the installation is complete via the License Inventory in the UI (See "How to Add a License" in User Guide).
5. To install in a non-default location, enter the path in the **Installation Directory** field or use the folder browser. Click **Next**.

6. Select the type of installation then click **Next**.
   - **Master** Select this option to install the Master node in your network. If this is a new installation then the Master node must be installed before any Client nodes. The Master node is the central controller for all other nodes and serves as the connection point for the Web UI and REST API.

     Note: The Master node automatically has all the capabilities of a Client node. It is however recommended that the Master node functions only as a Master. Multiple Master nodes can coexist on the same network, however Client nodes can only be authorized and controlled by one Master node.
   - **Client** - Select this option to install all other node types. The specific roles assumed by Client nodes are defined via the Nodes Inventory once installation is completed (See section "How to add a node" in User Guide.). These roles include:
     - Data Sources (basic hosts, VMs, application servers, etc.)
     - ISMs (for controlling Hitachi Block and File storage hardware, etc.)

       Caution: ISM nodes and their associated CMDs used to control storage devices must not be shared with applications other than those forming part of the Ops Center suite.
     - Repositories (acting as host based backup storage destinations)
7. Specify a node name to be used within Ops Center Protector then click **Next**.

   Note: Node names are limited to a maximum of 64 characters. By default, the name is set to the machine's host name. This name is only used by Protector, and will not change the name set by the operating system.
8. If the **Master** is being installed:
   a. Select a local **User Account** for logging on to the Protector web user interface immediately after installation is complete. You must know the password associated with the selected account. The account used can be changed after installation of the Master. Click **Next**.
   b. Accept or edit the **Protector User Interface Port** used to connect to the web based UI. By default it is set to 443, but can be changed if you already have another web server running on that port.
9. If a **Client** is being installed:
   a. Enter the **Master Hostname or IP** address or a DNS resolvable name of the Master node. If it is known that this node will be operating over a non-secure network, then we recommend enabling the **Internet connected node** option. This will encrypt transmitted data as an extra security precaution. Over-the-wire encryption requires a license and may not be available in some territories.
10. When the wizard indicates that it is ready to begin the installation, click **Next**.
    Ops Center Protector files are copied to the designated directories and the necessary components installed.

11. When the wizard indicates that the installation is complete, you will have the option to **Start Hitachi Ops Center Protector User Interface Now** in a web browser. Click **Finish**.

    Note: You do not need to restart the machine. The installer starts all the necessary Ops Center Protector components on the system.

    If a third party firewall is installed on the network, Protector will generate firewall warnings when it starts running. See How to configure a third party firewall for Protector (on page 249).

12. If this is a new installation of a Master node, then use a web browser to log on to the user interface at: `https://<Master>/#/Login`, where *<Master>* is the IP address or DNS name of the Master node. Refer to How to configure a server-side SSL certificate using a UI (on page 251) to prevent security warnings being displayed by the web browser.

    Note: Initial logon must be done with the username specified for the **User Account** during the Master installation. The username must be qualified with the local domain name `master` as follows:

    `<username>@master`

    or

    `master/<username>`

    or

    `master\<username>`

13. Set the locale and time zone appropriate the location of the nodes you are working with.

    See "How to change the UI settings" in User Guide

14. If you are upgrading, the Dataflows should now be reactivated and all destination nodes should be manually resynchronized with their sources. The upgrade process is then complete.

How to upgrade Protector's PostgreSQL databases from version 12 to version 16

Before you begin

Protector uses several PostgreSQL databases to store data. These databases are not automatically upgraded along with the rest of the product and need to be done manually. Protector will still function with an older database version after an upgrade but it is recommended to receive the latest PostgreSQL performance improvements and security fixes. PostgreSQL was updated from version 12 to 16 in Protector 7.8 so databases from Protector 7.7 and earlier can be upgraded. There can be several databases on the Master and client nodes: RBAC, Logs, Repository Gen1, Repository Gen2, HCP Gen2, S3, Cloud Scale, and ISMmetadata these databases will be discovered as part of the upgrade process.

1. **Upgrade Protector service from 7.7.x to 7.8**.

2. **List all the PostgreSQL databases used on all nodes**. The List --all option, supported on a master node only, will display database name, database directory, and the node on which each database is located. Use the –local option if running on a client node.

   On a Windows Master node run:

   ```
   <install_path\bin>postgres-list--all
   ```

   On a Linux Master node run:

   ```
   <install_path/bin>./postgres-list--all
   ```

   Example output:

Database Type	Host Node Name	Location
RBAC	AM-EDC-Win	C:/Program Files/Hitachi/Protector/db/authorization
Logs	AM-EDC-Win	C:/Program Files/Hitachi/Protector/db/saves/logs
HWO RPMDS	AM-EDC-Win-2	C:/metadata
HWO Block Cache (GRR)	AM-EDC-Win-2	Cannot verify database location(s). Please run "postgres-list --local" on the specificied host node

3. **Upgrade PostgreSQL Databases:**

   Upgrading the PostgreSQL needs to be done on each node identified as containing databases.

   a. Login to the respective node.

   b. Stop the Protector Service on the respective node:

   ```
   diagdata --stop
   ```

   c. List the database location and credentials file using the command

   ```
   postgres-list--local
   ```

   Example output:

Database Type	Database Path	Credentials File	Database Version
RBAC	C:/Program Files/Hitachi/Protector/db/authorization	C:/Program Files/Hitachi/Protector/db/config/authorization-credentials.cfg	12

Database Type	Database Path	Credentials File	Database Version
Logs	C:/Program Files/ Hitachi/ Protector/db/ saves/logs	C:/Program Files/ Hitachi/ Protector/db/ config/log-db-user.cfg	12
HWO RPMDS	C:/metadata/ ISMMetadata/db	C:/metadata/ ISMMetadata/ config/ postgresql.cfg	12

d. Upgrade the Postgres database using the specified database location and credential files with command ***'post-migrator'***

```
postgrs-migrator –src-db <db_location> --src-creds
```

<credentials_file> --stop-dst-db

Replace **<Protector install>/bin**, **<db_location>**, and **<credentials_file>** with the actual paths and file names used in your setup.

How to verify the SSL/TLS fingerprint of a client node prior to authorising

Before you begin

For security conscious customers a mechanism exists be used to ensure a node connecting to the system is that node and that there is no man in the middle (MITM) attack.

The Node Details screen for the relevant node displays the calculated SSL/TLS fingerprint for the node:



Figure 66 Node Details - calculated SSL/TLS fingerprint for the node

This displayed fingerprint can then be compared to the one generated on the client machine itself.

To generate the fingerprint on the client run the following cli command from a command prompt in the <install_path>bin directory:

```
sslfingerprint
```

The fingerprint for the node will be displayed and can be used for comparison.

How to install a Protector master or client on a Windows cluster

Before you begin

WARNING: This section only applies when clustering the master or protecting data without one of Protector's application integrations.

Do not use this method of installation in case you want to protect supported applications like Microsoft SQL Server or Microsoft Hyper-V. For these applications, install the client as you would for any regular standalone system (see How to install/upgrade Protector on Windows and Linux or AIX (on page 235)).

The Protector *Master* or *Client* capability can be installed on servers in a Windows cluster. Before starting the installation, review the following:

- Ensure that all servers meet the software prerequisites that are outlined in How to install/upgrade Protector on Windows and Linux or AIX (on page 235).
- Confirm that a cluster environment is set up and working correctly.
- For *Master* installations only, ensure that you have all the required license keys for Protector. A separate license key is required for each node in the cluster.

Protector installation for a cluster environment involves performing the following sub-tasks:

Procedure

1. Install Protector on each node in the cluster (see How to install Protector on a node in a cluster (on page 244)).
2. Add the *Cofio Hub* service (see How to add the hub service to a cluster (on page 245)).
3. For *Master* installations only:
   a. Add Protector licenses for each node (see How to add licenses to a clustered master (on page 245)).
   b. Add an authentication domain for the cluster (How to configure authentication for a cluster (on page 246)).

How to install Protector on a node in a cluster

To install Protector on a Windows Failover Cluster node:

Procedure

1. From the Windows **Failover Cluster Manager** console, identify the cluster server node that is currently in control (i.e. the node which has access to the shared disks).
2. Install Protector on the controlling node by running `Protector-m.n-m.n.n.nnnnn-WIN-ppp.exe`.
   The Protector **Setup** dialog is displayed. Click **Next**.
3. Read and accept the license agreement, then click **Next**.
4. Change the **Installation Directory** to the shared drive where you want install Protector. If Protector is already installed on another node in the cluster then select that install location. Click **Next**.
   For example, if your shared drive is `E:\` then the installation path might be on the shared drive: `E:\Hitachi\Protector`.

   Note: The Protector **Installation Directory** must be the same for all nodes in the cluster.
5. Select either the **Master** or **Client** installation type, then click **Next**.
6. Change the default **Node Name** to one that represents the entire cluster. If Protector is already installed on another node in the cluster then use that name. Click **Next** to start the installation.

   Note: The Protector **Node Name** must be the same for all nodes in the cluster.
7. If the **Master** is being installed:
   a. Select a local **User Account** for logging on to the Protector web user interface immediately after installation is complete. You must know the password associated with the selected account. The account used can be changed after installation of the Master. Click **Next**.
   b. Accept or edit the **Protector User Interface Port** used to connect to the web based UI.
8. If a **Client** is being installed:

a. Enter the **Master Hostname or IP** address or a DNS resolvable name of the Master node.

9. When the wizard indicates that it is ready to begin the installation, click **Next**.
10. When the wizard indicates that the installation is complete, click **Finish**.
11. Open the Windows **Services** console, stop the *Cofio Hub* service and set its *Start Up Type* to *Manual*.

    The *Cofio Hub* service will be controlled by the Failover Cluster Manager in a clustered configuration.
12. From the **Failover Cluster Manager** console, invoke failover to the next node in the cluster, then repeat the above procedure.

How to add the hub service to a cluster

Complete the following steps on the active node in the cluster.

Procedure

1. In the **Failover Cluster Manager** console, configure a *Generic Service* for the cluster using the **High Availability Wizard**.
2. Select the *Cofio Hub* service from the list.
3. Name the service (e.g. *CofioHubSvc*) and assign an IP address, then click **Next**.
4. Select the shared disk that Protector is installed on.
5. Do not replicate any registry settings.
6. Confirm that all settings are correct and then proceed with configuring the service for high availability.
7. Review the report to ensure that the service was configured correctly, then click **Finish**.
8. If a **Master** is being installed, make a note of *Cofio Hub* service's IP address shown in the **Failover Cluster Manager** console. This will be used for accessing the Protector web UI.

   Note: If the web UI is not accessible at this address, check your firewall for the following:

   a. Ensure that Protector is not blocked by the firewall.
   b. Enable port 30304 (Protector) and port 443 (HTTPS).
9. If a **Client** is being installed, authorize the node in the Nodes Inventory.
10. Force a fail-over, to confirm that the Protector node remains online in the Nodes Inventory.

How to add licenses to a clustered master

Before you begin

Install the Protector Master on each of the nodes in the Windows cluster as described in How to install a Protector master or client on a Windows cluster (on page 243).

Because each machine in the cluster has a different machine ID, a separate license is required for each node:

Procedure

1. Navigate to the Protector web UI using the IP address for the *Cofio Hub* service.
2. Log in using the local machine credentials of `administrator@master`.
3. Add the license key for this machine following the procedure described in How to Add a License (on page 258).
4. To add the licenses for the remaining cluster nodes, browse to the installation directory, then navigate to the subdirectory `db\config`
5. Right click `License.xml` and open the file with **WordPad**.
   The following XML is displayed showing the license key for the active node:

```
<licenses>
    <entry>ADYFKAE9PMM9BM2KXDNWCEO8PSABE244U8GA</entry>
</licenses>
```

6. Copy line 2 (the license key for the currently active cluster node) and insert it as a new line below line 2.
7. Change the license key on line 3 to match the one provided for the second cluster node.
   The following XML should be displayed:

```
<licenses>
    <entry>ADYFKAE9PMM9BM2KXDNWCEO8PSABE244U8GA</entry>
    <entry>DFGA54AGFDFHDJK675HH86453GHTFGD553DR</entry>
</licenses>
```

8. Repeat for each additional node.
9. Save and close the file.

How to configure authentication for a cluster

Before you begin

Install Protector on each of the nodes in the Windows cluster as described in How to install a Protector master or client on a Windows cluster (on page 243).

During installation, a local administrator account is used. To complete a cluster installation, a domain administrator account must be set up within Protector that is available to all nodes in the cluster:

Procedure

1. Navigate to the Protector web UI using the IP address for the *Cofio Hub* service.
2. Log in using the local machine's `administrator@master` credentials.
3. Add an *Authentication Space* that will perform authentication for the cluster.
   a. Specify the Authentication Space **Name**.
   b. Select the required authentication type (e.g. **Active Directory**).
   c. Select the cluster node as the **Proxy**.
   d. Enter the **Active Directory Domain Name**.

4. Add an *ACP Association* that will provide administrator level access to Protector (see "How to create an Access Control Profile Association" (Refer User Guide).
   a. Specify the ACP Association **Name**.
   b. Select the **User** *ACP Association* type.
   c. **Browse** for the required **User Name** from the **Authentication Space** specified in the previous step.
   d. Add the *Default Administrator* from the **Available Profiles** listed.
5. Log out of the web UI.
6. Log back into the web UI using the new `User@Domain` credentials specified in the above steps.
7. Force a fail-over, then log in again using the new credentials to confirm that Protector is still accessible.
8. Finally, refer to one of the following topics to setup accounts for those users that required access to the Protector user interface.
   - "How to configure basic role based access control" (Refer User Guide) to configure the minimum required access control functionality.
   - "How to configure advanced role based access control" (Refer User Guide) to configure full access control functionality.

How to remotely upgrade Protector client nodes

Before you begin

Note:

- The *Manage Software Updates* RBAC activity must be assigned to users who perform upgrades. It is recommended that this activity is restricted to administrative users only.
- Upgrade of a Microsoft failover cluster node needs to be done locally, following a similar procedure to that described in How to install a Protector master or client on a Windows cluster (on page 243). It's not possible to upgrade remotely because the standby nodes in the cluster are not available. Attempting a remote upgrade will place the clustered Protector service in a failed state, taking it offline.
- Before upgrading unmount any mounted snapshots.
- DO NOT upgrade while replications are in the process of pairing. Any active replications must be in the paired state before upgrade is carried out.

An upgrade can be performed if the existing installation has become corrupted or a newer version of Ops Center Protector is available.

If upgrading, obtain the upgrade installer files from your Hitachi Vantara support representative.

When an upgrade starts, the Ops Center Protector services are shutdown, causing the upgrading *OS Host* node and any nodes for which it serves as a proxy, to go offline in the Nodes Inventory. Any active data flows using those nodes will be temporarily interrupted. These nodes will come back online again when the services are automatically restarted on the *OS Host* node, after the upgrade process is completed, and the affected data flows will resume operation.

Note: We recommend upgrading Ops Center Protector only after current backups have been completed.

Note:

When upgrading, it is highly recommended (and often necessary) to upgrade all nodes to the same version. The recommended order in which to upgraded is:

1. Internet Connected Nodes. If ICNs are not upgraded first, they may not be able to be 'push upgraded' through the UI. ICNs may appear to go offline until the master is upgraded.
2. Master Node(s).
3. Clients acting as ISMs, controlling Hitachi Block and File storage devices.
4. Clients acting as Data Destinations, such as Repositories.
5. Clients acting as Data Sources, such as application servers.

Procedure

1. Locate the installers and accompanying configuration files appropriate for the target OSs and hardware architectures in your Protector environment.

   The installer's executable and configuration filenames have the following format:

   Protector-**R***m.n-m.n.n.nnnnn-ostype-proctype*

   where:

   - *m.n-m.n.n.nnnnn* - is the version and build number
   - *ostype* - is the target operating system type:
     - **Win**
     - **Linux**
     - **AIX**
   - *proctype* - is the target processor type:
     - **x32**
     - **x64**
     - **PPC**

2. Copy both the installer and configuration files to the `C:\Programs Files\Hitachi\Protector\runtime\updater` folder on the Master node.

   Copy both the installer and configuration files to the “/opt/hitachi/protector/runtime/updater” folder on the Master node.

   This folder will need to be created manually if this is the first time an update has been applied.

3. Click **Nodes** on the Navigation Sidebar to open the Nodes Inventory
4. Select the nodes to be upgraded (Master first, then Clients), then click **Upgrade Clients** to start the upgrade process.

   Note:

   - Only *OS Host* Client node types can be upgraded remotely
   - It is recommended to push out the upgrade to clients in batches of up to 20 nodes at a time.

   When upgrading the Master node, the UI will logout when the node's services are stopped by the installer. Wait a few minutes, then log back in again and complete the upgrade of the remaining nodes.

   When upgrading Client nodes, each one (and any nodes for which that Client acts as a *Proxy*) will go offline temporarily while the upgrade is applied.

   When the nodes come back online, the *Version* shown on the respective tile will be updated accordingly.

5. All active data flows should now be reactivated and all destination nodes should be resynchronized with their sources.

How to configure a third party firewall for Protector

If a third party firewall is installed within your network, when processes are started as part of the Ops Center Protector software installation, firewall warnings might be generated.

Configure the firewall to allow communication between all Protector nodes on one open port: 30304 (TCP).

Additional ports that might need to be opened are:

- Protector User Interface, on port 443 (HTTPS), or the alternative specified during Master installation.
- Protector vCenter proxies, also on port 443 (HTTPS), use the VDDK and vSphere web API to talk to the vCenter.
- Protector VMware proxies, on port 902, are instructed by the vCenter to talk to the ESXi host for virtual disk transfer.

How to configure addresses for nodes on multiple networks

If your environment includes nodes connected to more than one network, e.g. a production network, backup network and the internet, then you will need to configure Ops Center Protector after installation in order to:

- Use one or more of those networks in preferential order depending on availability.
- Prevent (bar) the use of certain networks.
- Connect to remote nodes over the internet by defining an external IP addresses where those nodes may be contacted.

The following procedure must be performed on every node that is connected to the network in question:

Procedure

1. Open a command prompt as administrator.
2. Change directory to the `<Ops Center Protector install>\bin` folder.
3. Stop the Protector services by entering the command **`diagdata --stop hub`**.
4. Check the node's current configuration by entering the command **`setconfig -l`**

   Refer to "Changing a node's profile with **`setconfig`**" in User Guide.
5. To configure a preferred IP address for the node, enter the command **`setconfig -p`** *nnn.nnn.nnn.nnn* where *nnn.nnn.nnn.nnn* is the IP address for the node on the preferred network.

   Additional addresses may be added by listing them in order of preference separated by semicolons. Multiple addresses may need to be quoted depending on Operating System.
6. To configure a barred IP address for the node, enter the command **`setconfig -b`** *nnn.nnn.nnn.nnn* where *nnn.nnn.nnn.nnn* is the IP address for the node on the barred network.

   Additional addresses may be added by listing them in order of preference separated by semicolons. Multiple addresses may need to be quoted depending on Operating System.
7. To configure an external (internet) IP address for the node, enter the command **`setconfig -x`** *nnn.nnn.nnn.nnn* where *nnn.nnn.nnn.nnn* is the IP address for the node on the internet.

   Additional addresses may be added by listing them in order of preference separated by semicolons. Multiple addresses may need to be quoted depending on Operating System.
8. If you need to remove an address then enter the command **`setconfig -r`** *nnn.nnn.nnn.nnn* where *nnn.nnn.nnn.nnn* is the IP address.

   Additional addresses may be added by listing them in order of preference separated by semicolons. Multiple addresses may need to be quoted depending on Operating System.

   Note: The address(es) are removed from the preferred, barred and fixed list.
9. Review the node's new configuration by re-entering the command **`setconfig -l`**
10. Restart the Protector services by entering the command **`hub start`**.

11. Wait a few minutes, then reauthorise the node on the Master (See "How to authorize a node" in User Guide).

How to configure a server-side SSL certificate using a UI with Ops Center

If Protector is configured as part of an Ops Center setup, the csslsetup command is available. By using the csslsetup command, you can configure SSl communication for Hitachi Ops Center products installed on the same management server using a common secret key and server certificate. For more information on the usage and support scope of the csslsetup command, refer to "Configuring SSL communications by using the csslsetup command" in the Hitachi Ops Center Installation and Configuration Guide.

How to configure a server-side SSL certificate using a UI

There are two versions of the tool:

- `Certificatetool` This is a UI based tool and requires a windowing manager to be installed on the Master
- `Certificatetoolcli` This is CLI tool and can be run on a ‘windows core installation’ or a non-X11 linux installation

Before you begin

For Protector installations on Windows and Linux, there is a certificate tool which enables you to:

- Create a self-signed certificate
- Create a certificate signing request
- Install existing certificate

You will need the following Distinguished Names (DN) to create a CSR or self-signed certificate:

- Common Name (CN) - The fully qualified domain name (FQDN) to be secured.
- Organisation (O) - The legal incorporated name of your company.
- Organisational Unit (OU) - The department administering Protector.
- City/Locality (L)
- State/Country/Region (S)
- Country (C)
- Email Address - A point of contact in your Protector administrative team (optional)

Configure server-side SSL certificate using Windows

Windows

Procedure

1. On the Master node, open the folder `<Protector home>\bin`
2. Locate and run the program `certificatetool.exe` with administrator privileges.
3. Follow the instructions in the **Certificate Tool** wizard to:
   - **Create a self-signed certificate**
   - or **Create a certificate signing request**

     Tip: The `.csr` file to be passed to your CA is created in the folder `<installation path>\db\ssl\webui\certificates`

     The corresponding `.key` file is created in `C:\Program Files\Hitachi\Protector\db\ssl\webui\private`
   - **Install Certificate**

Configure server-side SSL certificate using Linux

Linux

Procedure

1. On the Master node, open the folder `<installation path>/bin/`
2. Locate and run the program `certificatetool.exe`
3. Follow the instructions in the **Certificate Tool** wizard to:
   - **Create a self-signed certificate**
   - or **Create a certificate signing request**

     Tip: The `.csr` file to be passed to your CA is created in the folder `<installation path>/bin/db/ssl/webui/certificates`

     The corresponding `.key` file is created in `/opt/hitachi/protector/bin/db/ssl/webui/private`
   - **Install Certificate**

     Note: The Certificate tool can be run as either a UI version or CLI text version, thus a windowing manager is no longer required for Linux or Windows installations

How to configure a server-side SSL certificate manually

Before you begin

Obtain a valid SSL certificate file (`.crt` or `.cer`) and private key file (`.key`) for the Protector Web Server from your organization's Certificate Authority.

To configure server-side SSL for the Web UI:

Procedure

1. Copy your SSL certificate file to the folder `<installation path>/bin/db/ssl/webui/certificates`
2. Copy your SSL private key file to the folder `/opt/hitachi/protector/bin/db/ssl/webui/private`
3. On the Master node, open a command prompt at `<installation_path>/bin`.
4. Stop the Protector hub service by entering the command **`diagdata --stop hub`**.
5. Open the configuration file `<installation_path>/db/config/uic-system-config.cfg` in a text editor and do the following:
   a. Set `SSL Certificate` to the location of your SSL certificate file:

```
<item argtype="single" name="SSL Certificate">
    <value type="string"> <installation path>/bin/db/ssl/webui/certificates/
your_certificate.crt</value>
</item>
```

   b. Set `SSL Certificate Key` to the location of your SSL private key file:

```
<item argtype="single" name="SSL Certificate Key">
    <value type="string">/opt/hitachi/protector/bin/db/ssl/webui/private/
your_key.key</value>
</item>
```

   c. Set the `Webserver SSL` value to true:

```
<item argtype="single" name="Webserver SSL"
        <value type="bool">true</value>
<\item>
```

   You may need to insert the above XML within the following section of the config file:

```
<cofioxmllist>
 ...
<\cofioxmllist>
```

   Note: If you cut and paste these code fragments please ensure that line ends are correctly positioned.

6. Save the changes to the `uic-system-config.cfg` file.
7. Restart the Protector hub service on the Master node by entering the command **`diagdata --start`**.

How to add our CA certificate to the certificate store

Before you begin

In order to satisfy security scanner software the CA certificate used by Protector for inter-node communications needs to be trusted by the tool. To do this:

- Create a copy of `<installdir>/db/ssl/certificates/cacert.pem`
- Rename it to **`cacert.crt`**
- Propagate **`cacert.crt`** to all the relevant machines via your preferred method. This varies by OS and by configuration so is not documented here.

Using the CLI version of the certificate tool

As mentioned above, there is also a CLI version of the certificate tool. This was introduced in 7.3 to solve the problem of a non-X11 Linux installation or where the user is using a text based remote session. Also, where Windows is a 'Core-Installation' and thus a CLI version only.

The CLI version is not as rich in options as its UI counterpart, but it does have a set of useful options along with a set of examples.

To run the CLI version, simply enter the 'exe' name followed by the required arguments.

To see a list of the arguments and examples, like with all the Protector based CLIs enter:

`certificatetoolcli --help`

How to unregister the current certificate

There will always be situations where removal of the certificate and its components are required. This can be completed using the CLI version by using the 'u' argument:

`certificatetoolcli -u`

Note: It might be necessary to clear the browser cookies/Cached images and files for this change to take effect.

How to uninstall Protector

Before you begin

Caution:

- Uninstalling the Ops Center Protector Master does not deactivate active rules on Client nodes. To do this, you must deactivate all data flow(s) before uninstalling, or uninstall the Clients individually.
- Before uninstalling the Protector application the node (and any nodes proxied from it) must be removed from all dataflows, be unauthorized and deleted from the UI.

Operating System Specific Behaviour:

OS	Note

Linux and AIX	Completely removing Protector from a Linux or AIX machine is straight forward. Enter the following commands and follow the on-screen instructions:
	`/opt/hitachi/protector/uninstall`
	`rm -rf /opt/hitachi/protector`

Windows	Completely removing Protector from a Windows machine involves significantly more work, due partly to the additional functionality available on Windows nodes. Follow the steps listed below.

Procedure

1. Make a note of the path for the install directory and any repositories or ISMs. These will be required later if you wish to completely remove all traces of the installation and do not want to retain backup data stored in the repositories.
2. From the **Start** menu run **Uninstall Hitachi Ops Center Protector**. Alternatively, navigate to the installation directory and execute the command **`uninstall.exe`**.
   A dialog will be displayed asking you to confirm that you wish to proceed. Running the command **`uninstall.exe --mode unattended`** will cause the uninstall process to proceed without requiring user interaction; although popups may briefly appear before being automatically dismissed.
3. Click **Yes** to proceed with the uninstall process (or **No** to abort the process).
   A **Setup** dialog is displayed containing an **Uninstall Status** bar and information about what actions are currently being performed.
4. At some point during the uninstallation process you may be presented with a **Warning** dialog telling you that any repositories must be manually deleted. The uninstall process will pause until you click **OK**.
   Manual removal of repositories is covered later on.
5. Eventually an **Info** dialog will be displayed stating that the uninstall process is complete. Click **OK**.
6. Delete any repository directories noted down before uninstalling. Ensure that there are no files or windows open at these locations or attempts at deletion may fail.
   Protector and its repositories are now uninstalled and you can end the procedure at this point.

   Some artefacts will still remain on your machine. These items are left in place either to facilitate easier reinstallation or because they remained after upgrading from an earlier version or from a failed install. If you want to remove these artefacts, then continue with the following steps:
7. Delete the Protector installation directory tree noted down before uninstalling. Ensure that there are no files or windows open at this location or attempts at deletion may fail.
8. A hidden directory `Protector-RecycleBin` may remain in the root directory. This can be removed.

9. Delete the Bitrock installer log files `bitrock_installer.log` and `bitrock_installer_<nnnn>.log` normally located in `C:\Users\Administrator\AppData\Local\Temp`
10. Open the Windows **Registry Editor** from the **Start** > **Run...** dialog by entering `regedit.exe` in the **Open:** field.
    The **Registry Editor** window is displayed. Alternatively use the **`reg delete`** command from within a **Command Prompt**.
11. Carefully delete the following registry keys:

    Caution: Incorrectly editing the Registry can cause the operating system to become unstable or stop working. Exercise extreme caution when performing this step.

    - `HKEY_LOCAL_MACHINE\SOFTWARE\Cofio Software`
    - `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Cofio Software`
    - `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CofioHub`
    - `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\dcefltr`
    - `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\sdrefltr`
    - `HKEY_LOCAL_MACHINE\SOFTWARE\Hitachi\Hitachi Ops Center Protector`
12. Confirm that the `dcefltr` and `sdrefltr` filter drivers are no longer installed on your system by entering the commands **`sc query dcefltr`** and **`sc query sdrefltr`**. You should see the following output in response to both commands:

    ```
    The specified service does not exists as an
    installed                service.
    ```

Next steps

Occasionally an uninstallation may fail to complete successfully. In this situation there are a number of things that can be done before retrying:

- Ensure any files or folders that can block the uninstallation are closed.
- Reboot the machine to return it to a known, stable state.
- Reinstall Protector over the existing partially uninstalled version. In most cases this will allow an uninstallation to be successfully re-performed.

How to setup an existing Protector master for Ops Center

Adding Protector to Ops Center can be achieved using the *setupcommonservice* CLI tool located in the Protector bin directory. This will add Protector to as an Ops Center Instance and enable SSO functionality.

Run the setupcommonservice CLI tool with the appropriate arguments. You will be prompted for a username (if not provided as an argument) and password for an Ops Center administrator account. This will create an opscenter Ops Center space and an ACP Association for Ops Center administrators. Protector should now be listed as a product in Ops Center titled with the hostname provided.

Abbreviated Argument	Full Argument	Description
-h	--help	Display help
--cs-uri arg		CS base uri e.g. https://myopscenter.com:443/portal
--cs-username arg		CS username (optional). If this option is not provided it will be requested interactively.
--app-scheme arg		Protector protocol scheme (http/https)
--app-hostname arg		Protector hostname It can be the hostname or IP address.
--app-port arg		Protector port (usually 443 if installed outside of Ops Center)
--ignore-ssl-errors arg		Ignore SSL certificate errors (on/off) Default is on.

Removing Protector from Ops Center

The operation for removing Ops Center Protector that was registered in Common Services is performed using the Hitachi Ops Center Portal.

Ops Center ships with a Protector master preconfigured. If an existing Protector instance is being used it may be useful to remove this. To remove Protector, navigate to the 'Products' page in Ops Center and click the delete button (bin icon) on the Protector tile. This will unregister Protector and remove the Ops Center space.

Note: Any ACP associations for Ops Center will **not** be removed when deleting Protector from Ops Center.

Job Tasks

This section describes job control tasks that users will perform with Ops Center Protector.

For further information, refer to:

- Job Concepts (on page 109)
- Jobs UI Reference (on page 493)

How to view and control running jobs

Ops Center Protector tracks background asynchronous operations through jobs which can be viewed as follows:

Procedure

1. Click **Jobs** on the Main Banner (on page 286) to open the Jobs Inventory (on page 493).

   Jobs can also be viewed from the Monitor Details (on page 521).

   All currently running and recently completed jobs are displayed in a table with the most recently started job listed first.
2. From the Jobs Inventory (on page 493) select one or more jobs you want to control, by clicking the radio buttons to the left.
3. Click on the **More Actions** button above the table to see a menu of the available actions.

   Depending on the type of job selected and the phase it has reached, some or all of the actions may not be available.
4. To view the details of a specific job, click on the **Job Type** in the table.
   The Job Details (on page 499) are displayed.

License Tasks

This section describes licensing tasks that users will perform with Ops Center Protector.

For further information, refer to:

- License Concepts (on page 128)
- Licenses UI Reference (on page 506)

How to Add a License

You must provide Ops Center Protector with a valid license key in order to use its features as follows:

Procedure

1. Click **Licenses** on the Navigation Sidebar to open the License Inventory.
2. Copy the **Machine ID** displayed at the top of the License Inventory and include this in your license request email to your Hitachi Vantara support representative.
3. When you have received your license key, click **Add** to open the Activate License Wizard.
4. Copy the license key provided by your vendor and paste it into the **Licence Key** field.
5. Click **Finish**.
   The License Details will be displayed. Review it to ensure that its expiry data is correct and that it provides the required features.

How to view a license

Ops Center Protector features are enabled via license keys which can be viewed as follows:

Procedure

1. Click **Licenses** on the Navigation Sidebar (on page 291) to open the Licenses Inventory (on page 506).
2. Click on the name of the license you want to view.
   The License Details (on page 509) will be displayed.

Log Tasks

This section describes log review tasks that users will perform with Ops Center Protector.

For further information, refer to:

- Log Concepts (on page 107)
- Logs UI Reference (on page 509)

How to view logs

Ops Center Protector generates log messages as it performs activities. These can be viewed as follows:

Procedure

1. Click **Logs** on the Main Banner (on page 286) to open the Logs Inventory (on page 509).
   All log messages are displayed in a table with the most recently generated log message listed first.
2. Click on the **Condensed View** or **Extended View** buttons to view fewer or more columns in the table.
3. Select one or more log entries you want to perform an action on, by clicking the check box buttons within the table.
   You can then **Acknowledge Selected Logs** or **Create Log Notification**s for a message by clicking on the buttons above the table.
4. To view the attachments for a specific log entry, click on the **Attachments** icon in the table next to the log entry.
   The Log Attachments Dialog (on page 519) is displayed.
5. To view only those logs that are associated with the same session (a job logs to a session), click the **View Session** icon in the table next to the log entry.
   The Session Log Details (on page 517) are displayed, listing all logs for the session.

Monitor Tasks

This section describes data flow monitoring tasks that users will perform with Ops Center Protector.

For further information, refer to:

- Monitor Concepts (on page 106)
- Monitor UI Reference (on page 520)

How to trigger an operation from an active data flow

Before you begin

A data flow and associated policy has been defined, compiled and activated. For a typical example see How to batch backup a file system path to a repository (on page 141).

Tip: It is also possible to trigger operations from the RPO Report (on page 741).

Note: If a replication is paused, and subsequently evaluated due to trigger, schedule, or dataflow activation, the job must succeed, without actually resynchronizing the replication pairs of the paused replication. If the job were to fail in such circumstances, then jobs for downstream operations (other replication or snapshots) could not be executed.

An operation can be triggered manually from a monitor data flow as follows:

Procedure

1. Click **Monitor** on the Main Banner (on page 286) to open the Monitor Inventory (on page 520).
2. Locate the tile for the active data flow containing the nodes and policy operations you want to trigger, then click on the name of the data flow to open its Monitor Details (on page 521).
3. Select the **source node** on the data flow that implements the operation you want to trigger then click the **Trigger Operation** button in the top left of the data flow canvas. The Trigger Operation Dialog (on page 529) is displayed, listing all the available policies that can be triggered from that node.
4. Select the policy or policies that you want to trigger then click **OK**.
   The dialog closes and a pop-up message is displayed indicating that the operation has been initiated. An animated icon appears above the related mover(s) on the data flow and a new job appears at the top of the Jobs list below the data flow. Once the operation has finished, the *Job Progress* will be set to *Completed* if the operation was successful or *Failed* if not.

How to deactivate an active data flow

Before you begin

A data flow and associated policy has been defined, compiled and activated. For a typical example see How to batch backup a file system path to a repository (on page 141).

Tip: Once deactivated, the data flow will be removed from the Monitor Inventory (on page 520). To reactivate the data flow, follow the steps described in How to activate a data flow (on page 232).

An active data flow can be deactivated as follows:

Procedure

1. Click **Monitor** on the Main Banner (on page 286) to open the Monitor Inventory (on page 520).
2. Locate the tile for the active data flow that you want to deactivate, then click on the name of the data flow to open its Monitor Details (on page 521).

   You can also deactivate the data flow from the Monitor Inventory (on page 520), however opening the Monitor Details (on page 521) enables you to confirm that you deactivating the correct data flow.
3. Click the **Deactivate** button.
   A dialog is displayed before the data flow is deactivated. The Monitor Details (on page 521) closes and the Monitor Inventory (on page 520) is displayed with the deactivated data flow no longer listed.

Node Tasks

This section describes node configuration tasks that users will perform with Ops Center Protector.

For further information, refer to:

- Node Concepts (on page 47)
- Nodes UI Reference (on page 536)
- Node Groups UI Reference (on page 531)

How to add a node

Before you begin

The node must have the Protector Client software component installed (see How to install/upgrade Protector on Windows and Linux or AIX (on page 235)).

Note: *OS Host* nodes do not need to be added to the Nodes Inventory (on page 536) because they are automatically detected by the Master when the Client software component is installed.

The specific role of a node is determined when it is added to the Nodes Inventory (on page 536) as follows:

Procedure

1. Click **Nodes** on the Navigation Sidebar (on page 291) to open the Nodes Inventory (on page 536).
2. Click **Create new item** to open the Node Type Wizard (on page 539).
3. Select the type of node to be added then click **Start**.
   The Node Type Wizard (on page 539) will guide you through creating a node of the required type.
4. Finally click **Finish** and the new node will appear in the Nodes Inventory (on page 536).

How to authorize a node

Before you begin

The node must appear in the Nodes Inventory (on page 536) (see How to add a node (on page 261)).

Nodes must be authorized by the Master node, before they are allowed to participate in a data flow, as follows:

Procedure

1. Click **Nodes** on the Navigation Sidebar (on page 291) to open the Nodes Inventory (on page 536).
2. Select the tile for the node you want to authorize by clicking in its top left corner.

   Unauthorized nodes display a icon in the top right corner
3. Click **Authorize Node** above the inventory area.
4. Wait a few moments for the node to be authorized.

   Authorized nodes display a icon in the top right corner.

> **Note:** For extra security it is possible to verify the SSL/TLS fingerprint of the client node prior to authorization see How to verify the SSL/TLS fingerprint of a client node prior to authorising (on page 242)

How to stop the services

Before you begin

Sometimes it is necessary to stop the Protector services on a node, this needs to be done locally on the node itself.

- Open a console/command prompt window (On Windows select "Run as administrator", On Linux ensure you are running with root privileges using sudo))
- Within the command window change directory to the Protector 'bin' directory
- To stop the hub services – run the command:

```
diagdata --stop
```

How to start the services

Before you begin

If the Protector services have been stopped the services need to be restarted locally on the node itself.

- Open a console/command prompt window (On Windows select "Run as administrator", On Linux ensure you are running with root privileges using sudo))
- Within the command window change directory to the Protector 'bin' directory
- To stop the hub services – run the command:

```
diagdata --start
```

How to enable or disable encryption on a node

Before you begin

The node must appear in the Nodes Inventory (on page 536) (see How to add a node (on page 261).

Nodes must be authorized by the Master node see How to authorize a node (on page 262).

Procedure

1. Click **Nodes** on the Navigation Sidebar (on page 291) to open the Nodes Inventory (on page 536).
2. Select the node you want to set the encryption state by clicking in its node name green text. The Node Details (on page 650) is displayed to enable the node's details to be viewed and edited.
   a. Click **Enable\Disable Encryption** above the inventory area.
   b. A popup message will be displayed indicating the encryption has been enabled or disabled as appropriate.
   c. The **Encryption** entry in the **Configuration** table will be updated with the current state

How to change the account credentials for a Block Device node

Before you begin

It is assumed that there is already a block device configured and connected to a storage array.

Procedure

1. Click **Nodes** on the Navigation Sidebar (on page 291) to open the Nodes Inventory (on page 536).
2. Select the tile for the Block Device Node you want to change the credentials for by clicking in its top left corner.
3. Click **Edit Node** above the inventory area.
4. Click through the Hitachi Block Device Node Wizard (on page 583) until you reach the 'Specify credentials for device' screen
5. Enter the new username/password as required.
6. Click through the rest of the wizard and click finish at the end. The credentials will now be updated for this Block Device node.

Notification Tasks

This section describes notification configuration tasks that users will perform with Ops Center Protector.

For further information, refer to:

- Notification Concepts (on page 108)
- Notifications UI Reference (on page 659)

How to create a notification

Before you begin

If you intend to use email as the notification method, you must configure the email notification settings first. Refer to How to configure email settings for notifications (on page 265).

Any message that appears in the Logs Inventory (on page 509) can be configured to generate a notification via email, SNMP or the System Event Log. A notification can either be created from the Notifications Inventory (on page 659) (in which case the notification parameters must be supplied by the user) or from the Logs Inventory (on page 509) (in which case the notification parameters are provided automatically from the selected log message:

Procedure

1. EITHER:
   a. From the Logs Inventory (on page 509) select the log message that you want to be notified about by clicking the selection icon to the left of the message.

b. Click the **Create Log Notification** button at the top of the log message table.
The Notification Wizard (on page 662) opens and the parameters are automatically populated based on the selected log message.

2. OR:
   a. Open the Notifications Inventory (on page 659) by clicking on **Notifications** in the Navigation Sidebar (on page 291)
   b. Click **Create New Item** at the top of the notification table.
   The Notification Wizard (on page 662) opens so that the parameters describing the log message can be entered.
   c. Select the type of message for which a notification will be generated and enter the required parameters:
      - **Logs** - creates a notification based on a generic log message.
      - **Hitachi Block Device Monitoring** - creates a notification based on a specific type of log message associated with Hitachi Block Storage device parameters.
3. Click **Next**, then select the **Match Behaviour** and **Notification Method(s)**
4. Click **Finish** to close the wizard and display the Notifications Inventory (on page 659).
5. Create additional conditions by clicking **Create New Item** at the top of the notification table. Once you have finished creating conditions, it is recommended that you test them (see How to test a notification (on page 270)).

How to configure email settings for notifications

Before you begin

Set up an email account that can be used by Protector to send out notification emails. You will need to know the following details for this account:

- The email Account Provider - e.g. Google, Hotmail, etc.
- Account Name - e.g. dataprotection@company.com
- Host Name - e.g. smtp.company.com
- IP Port Number - e.g. 25 for SMTP
- Encryption Type - TLS, START_TLS or NONE
- Authentication Method - ON, OFF or NTLM
- Username
- Password

Note: It is necessary to lower the security levels for Yahoo and Google email providers. Refer to your provider's account security settings to allow applications with less secure sign in technology.

To configure the email address for Protector notifications:

Procedure

1. Click the **Notifications** option in the Navigation Sidebar (on page 291) to open the Notifications Inventory (on page 659).
2. Click the **Configure Email Settings** button in the top right corner of the page to open the Email Notifications Settings Wizard (on page 660).
3. Enter the details for the **Sender Account** and **Receiver Account**.
4. Click **Test Email Settings** and check that a test email is received from Protector as expected.
5. Click **Finish** to close the wizard.

How to configure SNMP Notifications

Before you begin

Protector supports SNMPv2c & SNMPv3 and sends Notification 'TRAPS' based on the following events:

Notification	Description
Logs	Generated by: Message, Category, Log Id, Actioned by 'node', level
Pool Monitoring	Generated by: Ser No., Pool Id, Pool Name, Pool Type, Pool Capacity
Journal Monitoring	Generated by: Ser No., Journal Id, Policy, Journal Capacity
HUR Monitoring	Generated by: Ser No., Policy, Source, Dest, HUR Network Latency
Custom	Generated by: Postgres SQL query

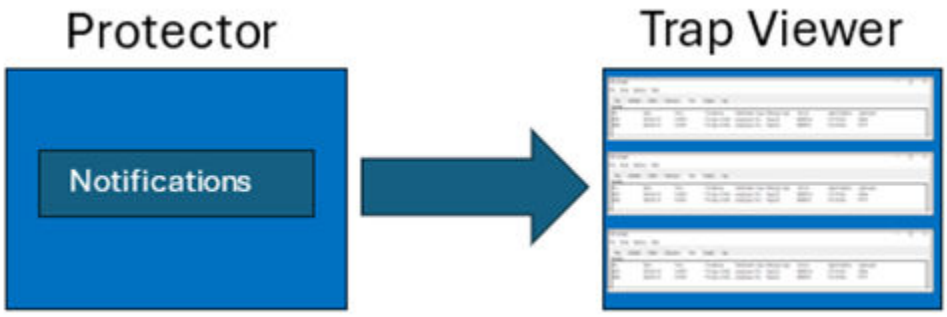


Figure 67 SNMP Viewer - Generating Notification

Notifications are generated on the Protector Master node (The Agent) and captured by an SNMP viewer (The Manager).

Protector is shipped with its own MIB file which will be required by the SNMP Manager to interpret the OIDs (Object Identifiers).

To configure Protector to send SNMP traps, the following steps need to be carried out:

1. Configure the Protector UI
2. Edit the Protector configuration files
3. Firewall configuration
4. Restart the Protector services
5. Copy the MIB file

Configure the Protector UI

Specify Name:

From the LHS pane select 'Notifications' and then a name:

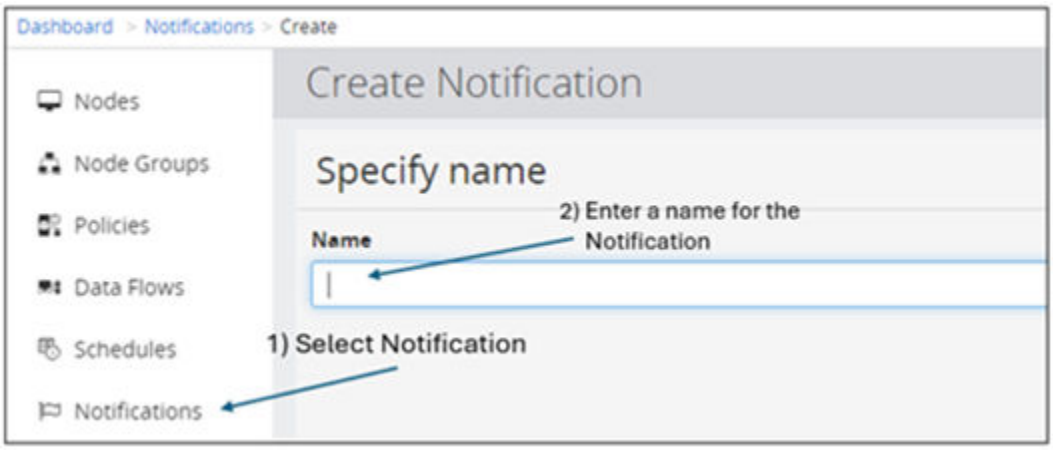


Figure 68 Notification - Create Notification

Specify Notification criteria:

Now, in the Notification Criteria, select where it will come from and then fill in the required field(s).

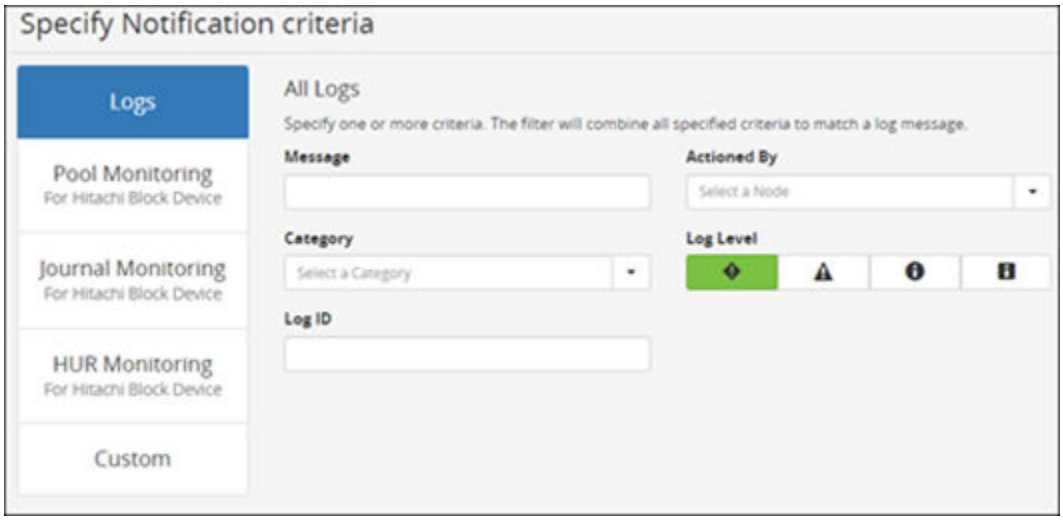


Figure 69 Notification - Specify Notification Criteria

Select Notification Options:

Finally, select the SNMP version, which can be either SNMPv2c/SNMPv3 or both. Once done, click Finish.

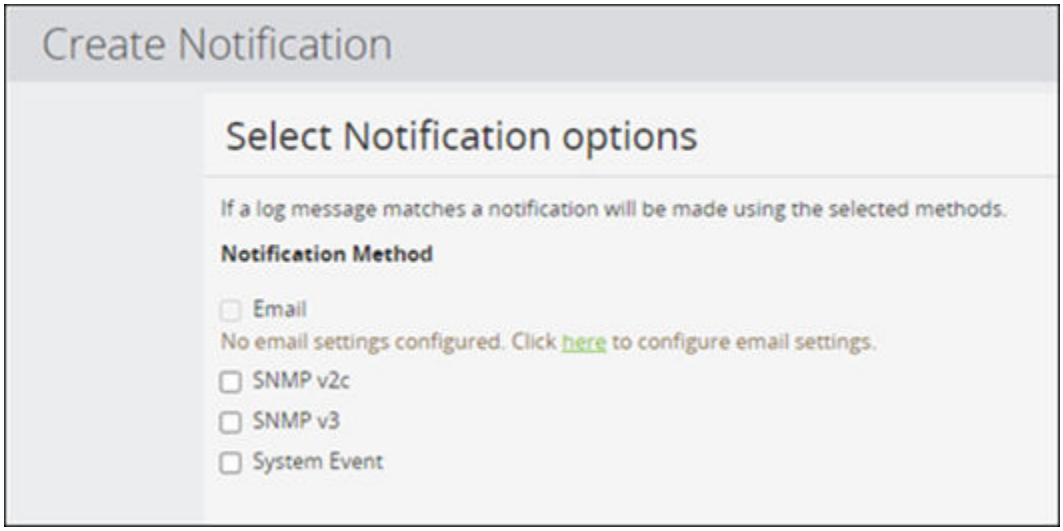


Figure 70 Notification - Select Notification Options

Edit the Protector configuration files

On the Master host, there are two files that need to be edited to complete the SNMP configuration, dependent on the version of SNMP required.

1. For SNMPv2c – `edit db/config/ notification/snmpv2c.cfg`
2. For SNMPv3 – `edit db/config/notification/snmpv3.cfg`

SNMPv2

1. In an editor open db/config/ notification/snmpv2c.cfg
2. Edit the appropriate HandlerExecutable section by inserting the desired values for the flags

Flag	Options
-c	community string default = public
-trap_host	The DNS/IP of the SNMP Manager (Trap host)
-trap_port	The port that the SNMP Manager will listen on

SNMPv3

1. In an editor open db/config/ notification/snmpv3.cfg
2. Edit the appropriate HandlerExecutable section by inserting the desired values for the flags

Flag	Options
-a	authentication protocol (MD5|SHA|SHA-224|SHA-256|SHA-384|SHA-512), default:SHA
-A	authentication protocol pass phrase, default:mypassword
-x	privacy protocol (DES|AES), default:AES
-X	privacy protocol pass phrase, default:mypassword
-l	security level (noAuthNoPriv|authNoPriv|authPriv), default: authPriv
-u	security name, default:traptest
-trap_host	The DNS/IP of the SNMP Manager (Trap host)
-trap_port	The port that the SNMP Manager will listen on default:162

For more information see: http://www.net-snmp.org/

Firewall configuration

The firewall on the Protector Master must to be opened to allow data on port 162 through

Restart the Protector services

The services on the Protector Master need to be restarted each time the files snmpv2c.cfg and/or snmpv3.cfg are modified.

Copy the MIB file

Locate the Protector MIB which will be here: db/mibs/ HITACHI-Protector-MIB.txt and copy it to the SNMP Manager.

Protector agent files

Protector is an SNMP agent, and these are the files associated with the notification traps:

snmpv2c.cfg	db/config/notification	SNMPv2c details
snmpv3.cfg	db/config/notification	SNMPv3 details
HITACHI-Protector-MIB.txt	db/mibs	Protector's MIB file

How to test a notification

Before you begin

Create one or more notifications (see How to create a notification (on page 264)).

Use the **`sendlog`** command (see Generating log messages with sendlog (on page 960)) to create logs to test that a notification event handler (NEH) functions properly.

When a notification event is triggered, a log message file is created that exists for the duration of the handler execution, after which it is deleted. To view the contents of a log message file, an event handler can be created that dumps the file to standard output. In Windows the following command could be placed in a batch file that acts as the NEH:

```
more < %1
```

A log message entry will appear on the Logs Inventory (on page 509) with an attachment containing the above batch command, followed by the XML formatted log message that triggered it.

How to customize alert notifications

If your site must receive alert notifications based on real-time or audit log events, then you can create a custom configuration file with instructions for which script/program Ops Center Protector must run as an event handler, and the arguments to pass to it.

Custom event handlers are added to the Notification Wizard (on page 662) as checkboxes in the Notification Method area.

By default, all Ops Center Protector notification handler configuration files are stored relative to the installation directory:

`\db\config\notification`

The path to the handler executable must be relative to the `\bin` directory where Ops Center Protector is installed. The default locations are as follows:

For Windows:

`C:\Program Files\Hitachi\Protector\bin`

For Linux:

`/opt/hitachi/protector/bin`

The built-in Email and System Event notification methods are configured by `email.cfg` and `systemlog.cfg` and implemented by **`nehemail.exe`** and **`neheventlog.exe`**, respectively (we recommend using the prefix `neh` to identify notification event handlers).

The first parameter passed to the handler is always the path to a log message file (About the notifications log message file (on page 271)). You can also specify your own additional arguments for Windows and Linux that are stored separately, so the same configuration file can be reused on both operating systems.

The custom hander is specified in the format shown in the following example (`customeventhandler.cfg`):

```
<cofioxmllist>
    <!-- text for custom notification method check box -->
    <item name="HandlerName" argtype="single" >
        <value type="string" >Custom</value>
    </item>

    <!-- default state of custom notification method check box -->
    <item name="Default" argtype="single" >
        <value type="bool" >False</value>
 </item>

    <!-- handler executable or script name -->
    <item name="HandlerExecutable" argtype="list" >
        <item name="WinExecutable" argtype="single" >
            <value type="string" >nehcustom.exe</value>
        </item>
        <item name="UnixExecutable" argtype="single" >
            <value type="string" >nehcustom</value>
        </item>
    </item>

    <!-- arguments 2, 3, ..., n passed to custom handler -->
    <item name="WinHandlerArgs" argtype="list" >
  <item name="HandlerArg" argtype="single" >
      <value type="string">windows custom arg</value>
  </item>
 </item>
 <item name="UnixHandlerArgs" argtype="list" >
  <item name="HandlerArg" argtype="single" >
      <value type="string">unix custom arg</value>
  </item>
 </item>
</cofioxmllist>
```

Any text printed to the standard output device by the handler will be included as an attachment to an error log message:

handler name `notifier failed. (Attachment 1) *** Attachment count 1 ***"`

The exit code of the handler is ignored.

About the notifications log message file

The log message file name is passed as the first parameter to all NEH handlers. It consists of a list of value pairs named:

`NotificationDetailFieldName` – a string identifying the parameter

and

`NotificationDetailFieldValue` – a value for the parameter

By reading these name-value pairs, the handler is able to analyze the log entry that caused it to be invoked and respond accordingly.

Log message files have a variable number of fields and can contain additional fields, such as a `SessionID` or a `Data Source`. The following table lists common fields.

Immutable	The log message will not be cleared when erasing the logs.
SequenceNumber	All logs have an ascending sequence number.
MasterDate	A 32 bit timestamp of when the master received the log message.
LocalDate	A 32 bit timestamp of when the log message was sent.
Source	The node from which the log message originated.
Category	The log message category.
Level	A value between 1 and 4 with 1 being *Detail* and 4 being *Error*.
Attachments	The text of any attachment to the log message.
Log	The body of the log message.
Acked	Indicates if the log message is acknowledged (0 or 1).
BuildNumber	The version of Ops Center Protector on the node that sent the log message.

Policy Tasks

This section describes policy configuration tasks that users will perform with Ops Center Protector.

Refer to Data Protection Workflows (on page 130) for detailed descriptions of specific Repository and Hitachi Block data protection scenarios.

For further information, refer to:

- Policy Concepts (on page 90)
- Policies UI Reference (on page 670)

How to create a policy

Before you begin

Ensure that the nodes acting as sources and destinations for your backup data have been added to Protector, see How to add a node (on page 261).

To create a policy:

Procedure

1. Click the **Policies** link on the Navigation Sidebar (on page 291) to open the Policies Inventory (on page 670).
2. Click the **Create new item** tile to open the Policy Wizard (on page 671).
3. Enter a **Name** and **Description** for the policy, then click **Next**.
4. From the **Classifications** inventory page of the Policy Wizard (on page 671) you can either:
   - Add filters to an existing classification. Refer to How to add a filter to a classification (on page 276).
   - Click the **Create new item** tile to add another classification to the policy. Refer to How to add a classification to a policy (on page 274).
   - Edit an existing classification in the policy. Refer to How to edit a classification in a policy (on page 275).
   - Delete an existing classification from the policy. Refer to How to delete a classification from a policy (on page 275).
5. At least one classification must be added to a new policy. When you have finished adding classifications and filters, click **Next**.
   The **Operation Selection** page of the Policy Wizard (on page 671) is displayed.
6. Select the required operation for the policy. At least one operation must be added to a new policy. Refer to How to add an operation to a policy (on page 277) before proceeding further.
7. From the **Operations** inventory page of the Policy Wizard (on page 671) you can either:
   - Click the **Create new item** tile to add another operation to the policy. Refer to How to add an operation to a policy (on page 277).
   - Edit an existing operation in the policy. Refer to How to edit an operation in a policy (on page 277).
   - Delete an existing operation from the policy. Refer to How to delete an operation from a policy (on page 278).
8. When you have finished adding operations, click **Finish**.

How to edit a policy

Before you begin

It is recommended that all data flows where the policy is used are reviewed so that the effect of editing the policy is understood.

To edit a policy:

Procedure

1. Click the **Policies** link on the Navigation Sidebar (on page 291) to open the Policies Inventory (on page 670).
2. Select the required policy tile.
   The Policy Details (on page 737) opens.
3. Click **Edit** above the inventory to open the Policy Wizard (on page 671).
4. Editing the policy is similar to creating one. Refer to How to create a policy (on page 273).
5. If required, you can now recompile and reactivate all data flows where the edited policy is used. Refer to How to activate a data flow (on page 232).

How to delete a policy

Before you begin

Note: It is not possible to delete a policy which has been applied to a node within a data flow (active or not). All data flows where the policy is used must therefore be reviewed and edited so that the policy to be deleted is unassigned from all nodes.

To delete a policy:

Procedure

1. Click the **Policies** link on the Navigation Sidebar (on page 291) to open the Policies Inventory (on page 670).
2. Select the policy to be deleted by clicking in the top left corner of the corresponding tile.
3. Click **Delete** above the inventory.
4. Recompile and reactivate the rules for all data flows where the deleted policy was used.

How to add a classification to a policy

Before you begin

It is recommended that all data flows where the policy is used are reviewed so that the effect of adding the classification is understood.

To add a classification to a policy:

Procedure

1. Navigate to the **Classifications** inventory page of the Policy Wizard (on page 671) either by creating a new policy (How to create a policy (on page 273)) or editing an existing one (How to edit a policy (on page 273)).
2. Click the **Create new item** tile to add a classification to the policy.
   The **Classification Selection** page of the wizard is displayed.

3. Select the classification category you require on the left of the wizard.
   The available classification types are displayed on the right of the wizard.
4. Select the required classification from the list on the right of the wizard, then click **Next**.
   One of the following classification wizards is displayed:
   - Application classifications:
     - Oracle Database Classification Wizard (on page 686)
     - Oracle RMAN Classification Wizard (on page 688)
   - Hypervisor classifications:
     - VMware Classification Wizard (on page 700)
   - Physical classifications:
     - Disk Type Classification Wizard (on page 708)
     - Hitachi Block Classification Wizard (on page 709)
     - Path Classification Wizard (on page 713)
5. Enter the parameters required for the selected classification, then click **Apply**.

How to edit a classification in a policy

Before you begin

It is recommended that all data flows where the policy is used are reviewed so that the effect of editing the classification is understood.

To edit a classification:

Procedure

1. Navigate to the **Classifications** inventory page of the Policy Wizard (on page 671) by editing an existing policy.
2. Select the classification to be edited by clicking on the **<classification name>** of the corresponding tile.
   The classification wizard is displayed.
3. Make the required changes to the classification, then click **Apply**.

How to delete a classification from a policy

Before you begin

It is recommended that all data flows where the policy is used are reviewed so that the effect of deleting the classification is understood.

To delete a classification:

Procedure

1. Navigate to the **Classifications** inventory page of the Policy Wizard (on page 671) by editing an existing policy.

2. Select the classification to be deleted by clicking in the top left corner of the corresponding tile.
3. Click **Delete** above the inventory.

How to add a filter to a classification

Before you begin

It is recommended that all data flows where the policy is used are reviewed so that the effect of adding the filter is understood.

To add a filter to a classification:

Procedure

1. Navigate to the **Classifications** inventory page of the Policy Wizard (on page 671) either by creating a new policy (How to create a policy (on page 273)) or editing an existing one (How to edit a policy (on page 273)).
2. Click the **Edit Filters** link on the required classification tile.
   The Classification Filters Wizard (on page 727) is displayed.
3. Click the **Create new item** tile.
   The **Classification Filter Selection** page of the wizard is displayed.
4. Select the required filter from the list displayed by the wizard, then click **Next**.
5. Enter the parameters required for the selected filter, then click **Apply**.

How to edit a filter on a classification

Before you begin

It is recommended that all data flows where the policy is used are reviewed so that the effect of editing the filter is understood.

To edit a filter on a classification:

Procedure

1. Navigate to the **Classifications** inventory page of the Policy Wizard (on page 671) by editing an existing policy.
2. Select the filter to be edited by clicking on **Edit Filters** of the corresponding classification tile.
   The Classification Filters Wizard (on page 727) is displayed.
3. Select the filter to be edited by clicking on the **<filter name>** of the corresponding tile.
4. Make the required changes to the filter, then click **Apply**.

How to delete a filter from a classification

Before you begin

It is recommended that all data flows where the policy is used are reviewed so that the effect of deleting the filter is understood.

To delete a filter from a classification:

Procedure

1. Navigate to the **Classifications** inventory page of the Policy Wizard (on page 671) by editing an existing policy.
2. Select the filter to be deleted by clicking on **Edit Filters** of the corresponding classification tile.
   The Classification Filters Wizard (on page 727) is displayed.
3. Select the filter to be deleted by clicking in the top left corner of the corresponding tile.
4. Click **Delete** above the inventory.

How to add an operation to a policy

It is recommended that all data flows where the policy is used are reviewed so that the effect of adding the operation is understood. Adding an operation may cause a compiler warning if the operation is not applied to all data flows using the policy.

To add an operation to a policy:

Procedure

1. Navigate to the **Operations** inventory page of the Policy Wizard (on page 671) either by creating a new policy (How to create a policy (on page 273)) or editing an existing one (How to edit a policy (on page 273)).
2. Click the **Create new item** tile to add an operation to the policy.
   The **Operation Selection** page of the wizard is displayed.
3. Select the required operation from the list displayed by the wizard, then click **Next**.
   One of the following operation wizards is displayed
   - Mount Operation Wizard (on page 721)
   - Backup Operation Wizard (on page 718)
   - Replicate Operation Wizard (on page 722)
   - Snapshot Operation Wizard (on page 724)
   - Tier Operation Wizard (on page 726)
   - Access Operation Wizard (on page 732)
4. Enter the parameters required for the selected operation. For some operations it is possible to select **Run on Schedule** under **Run Options** to specify when the operation will be run. To create a schedule click **Manage Schedules** then refer to How to create a schedule (on page 280).
5. Once the operation's parameters have been entered, click **Apply**.

How to edit an operation in a policy

Before you begin

It is recommended that all data flows where the policy is used are reviewed so that the effect of editing an operation is understood.

To edit an operation:

Procedure

1. Navigate to the **Operations** inventory page of the Policy Wizard (on page 671) by editing an existing policy.
2. Select the operation to be edited by clicking on the **<operation name>** of the corresponding tile.
   The operation wizard is displayed.
3. Make the required changes to the operation, then click **Apply**.

How to delete an operation from a policy

Before you begin

It is recommended that all data flows where the policy is used are reviewed so that the effect of deleting the operation is understood.

To delete an operation:

Procedure

1. Navigate to the **Operations** inventory page of the Policy Wizard (on page 671) by editing an existing policy.
2. Select the operation to be deleted by clicking in the top left corner of the corresponding tile.
3. Click **Delete** above the inventory.

Report Tasks

This section describes reporting tasks that users will perform with Ops Center Protector.

For further information, refer to:

- Report Concepts (on page 118)
- Reports UI Reference (on page 739)

How to view a report

Ops Center Protector generates a number of different reports which can be viewed as follows:

Procedure

1. Click **Reports** on the Main Banner (on page 286) to open the Reports Dashboard (on page 739).
2. Select one the report type you want to view by clicking the corresponding button on the dashboard.
   The selected report is displayed.

3. From a report page, click on the **Export** button in the top right to open the Export Report Dialog (on page 761) which enables the report to be exported in various formats.

Restore Tasks

This section describes data restore tasks that users will perform with Ops Center Protector.

Refer to Data Protection Workflows (on page 130) for detailed descriptions of specific Repository and Hitachi Block data restore scenarios.

For further information, refer to:

- Restore Concepts (on page 110)
- Restore UI Reference (on page 763)

How to view available backups

Tip: All the operations available from the Restore UI are also available from the Storage UI (see the Storage UI Reference (on page 858)). The Storage UI also enables additional operations to be performed.

Ops Center Protector maintains an inventory of all available backups that can be viewed as follows:

Procedure

1. Click **Restore** on the Navigation Sidebar (on page 291) to open the Restore Inventory (on page 763).
2. Modify the filter options in the Filters section of the Inventory page to narrow down the search results as required.
3. Click **Search** in the Filters section of the inventory page.
   An inventory of all backups for that storage device type is displayed.
4. To view more details, click on the **<date and time>** of a backup tile to open the details page for a backup.
5. To perform a restore action, click on the **Restore**, **Mount**, **Unmount**, **Revert** or **Swap** button as appropriate to the backup type selected.

Schedule Tasks

This section describes schedule configuration tasks that users will perform with Ops Center Protector.

For further information, refer to:

- Schedule Concepts (on page 101)
- Schedules UI Reference (on page 849)

How to create a schedule

To create a schedule:

Procedure

1. Click the **Schedules** link on the Navigation Sidebar (on page 291) to open the Schedules Inventory (on page 849).
2. Click the **Create new item** tile to open the Schedule Item Wizard (on page 853).
   The **Specify name and description** page of the wizard is displayed.
3. Enter a **Name** and **Description** for the schedule, then click **Next**.
   The **Add one or more Schedules** page of the wizard is displayed.
4. From the **Add one or more Schedules** inventory page of the Schedule Wizard (on page 850), click **Create new item** to add an item to the schedule.
   The **Select Schedule item type** page of the Schedule Item Wizard (on page 853) is displayed.
5. Select the item **Type** from the list, then click **Next**.
   The **Configure Schedule item** page of the Schedule Item Wizard (on page 853) is displayed.
6. Select **Weekly** or **Monthly** on the left and edit the parameters displayed on the right.
7. Click **Apply** to add the item to the schedule.
   The Schedule Wizard (on page 850) is displayed, showing the list of items in the schedule.
8. From the **Add one or more Schedules** inventory page of the Schedule Wizard (on page 850) you can either:
   - Click **Create new item** to add another item to the schedule.
   - Select an existing item by clicking the selection icon to the left and click **Delete** to remove the item from the schedule.
   - Click on the item description (or click the selection icon to the left and click **Edit**) to edit an operation.
9. When you have finished adding items, click **Finish**.

Storage Tasks

This section describes storage management tasks that users will perform with Ops Center Protector.

For further information, refer to:

- Storage Concepts (on page 114)
- Storage UI Reference (on page 858)

How to view the status of a repository

Before you begin

It is assumed that a policy that creates repository snapshots has been implemented.

To view the status of a repository:

Procedure

1. Click **Storage** on the Main Banner (on page 286).
   The Storage Inventory (on page 858) is displayed.
2. In the filters area, set the **Type** to *Repository*.
   The inventory now only lists Repository storage nodes.
3. Click on the title of the tile corresponding to the required Repository.
   The Generation 1 Repository Details (on page 921) for the that Repository are displayed.
4. Click one of the following buttons on the Generation 1 Repository Details (on page 921):
   - **View Node Details** to view the Node Details (on page 650).
   - **Analyze** or **Refresh** to analyze the repository sizes.
   - **Stores** to view the Gen1 Repository Stores Inventory (on page 912).

How to view the contents of a snapshot in a repository store

Before you begin

It is assumed that a policy that creates repository snapshots has been implemented and that at least one snapshot has been created in the designated repository store.

To view the contents of a repository snapshot:

Procedure

1. Click **Storage** on the Main Banner (on page 286).
   The Storage Inventory (on page 858) is displayed.
2. In the filters area, set the **Type** to *Repository*.
   The inventory now only lists Repository storage nodes.
3. Click on the title of the tile corresponding to the required Repository.
   The Generation 1 Repository Details (on page 921) for that Repository are displayed.
4. Click the **Stores** button on the Generation 1 Repository Details (on page 921).
   The Gen1 Repository Stores Inventory (on page 912) is displayed.
5. Click on the title of the tile corresponding to the required Store.

   Each Store tile indicates the *Source Node* and *Policy* it is associated with.

   The Gen1 Repository Store Details (on page 915) for the that Store are displayed.
6. Click on the title of the tile corresponding to the required *Snapshot*.

   Each Snapshot tile indicates the *Date and Time* when it was created.

   The Repository Snapshot Details (Storage) - File System (on page 917) for the Snapshot are displayed.

How to view the status of a Hitachi Block storage device

Before you begin

It is assumed that a policy that creates Hitachi Block hardware snapshots has been implemented and that at least one snapshot has been created in the designated pool on the device.

To view the hardware status:

Procedure

1. Click **Storage** on the Main Banner (on page 286).
   The Storage Inventory (on page 858) is displayed.
2. In the filters area, set the **Type** to Hitachi Block Device or Hitachi *Logical Block Device*.
   The inventory lists only Hitachi Block storage nodes.
3. Click on the title of the tile corresponding to the required device.
   The Hitachi Block Device Details (on page 859) for the device are displayed.
4. Click one of the following buttons on the Hitachi Block Device Details (on page 859):
   - **Host Groups** to view the Hitachi Block Host Groups Inventory (on page 861)
   - **Journals** to view the Hitachi Block Journals Inventory (on page 863)
   - **Pools** to view the Hitachi Block Pools Inventory (on page 866)
   - **Logical Devices** to view the Hitachi Block Logical Devices Inventory (on page 868)
   - **Snapshots** to view the Hitachi Block Snapshots Inventory (on page 869).
   - **Replications or Clones** to view the Hitachi Block Replications Inventory (on page 878)
5. Click on the title of a tile in the selected inventory to view the details for that item.

Chapter 5: User Interface Reference

This chapter describes the Ops Center Protector based user interface.

User Interface Overview

This section provides a general overview of the Protector user interface.

User Interface Structure

The Ops Center Protector user interface follows a general pattern for the hierarchical layout of pages. Understanding this hierarchy will help you navigate the UI more effectively, enabling you to locate, create and edit resources quickly and easily.

Note: Before you can access the Protector UI, you must be granted access by a Protector administrator who will provide you with a user name and password. These credentials must be entered via the Login Page (on page 309). Depending on the roles assigned to you, your credentials may not allow you access to all areas or functions of the UI.

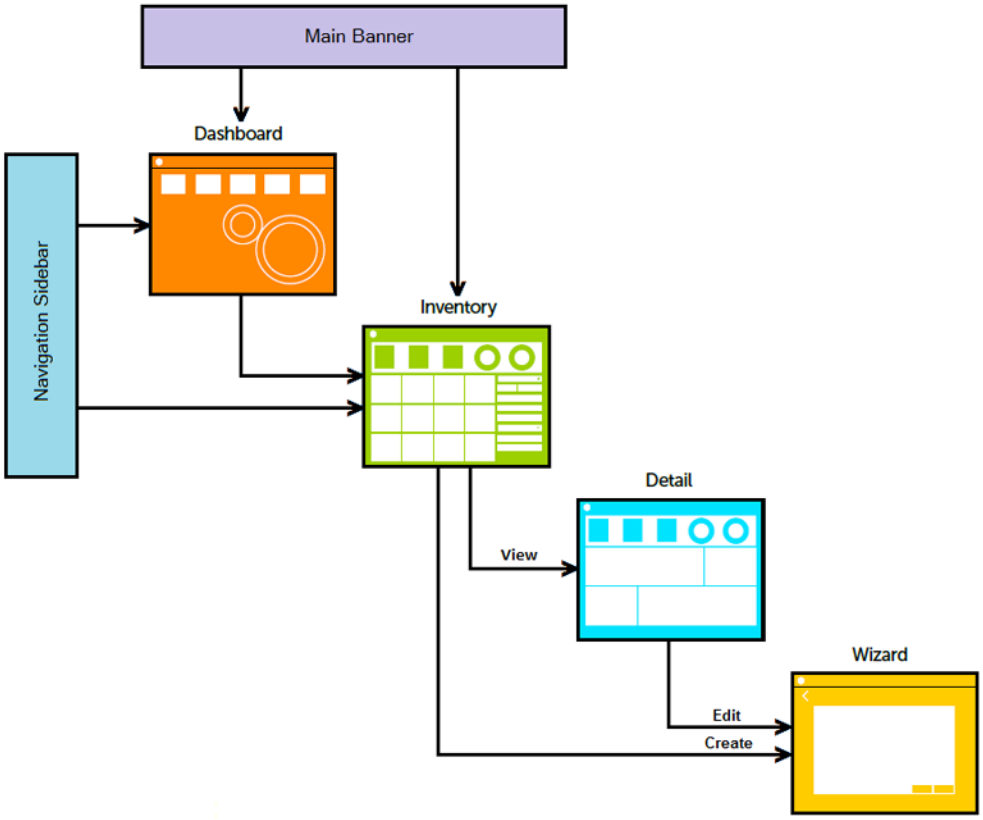


Figure 71 User interface general hierarchy

Page Type	Description
Main Banner	The Main Banner (on page 286) appears at the top of all pages in the user interface and provides links to all the frequently accessed monitoring and administration dashboards and inventories.
Navigation Sidebar	The Navigation Sidebar (on page 291) appears at the left of all pages in the user interface and provides links to all the frequently accessed management, configuration and design dashboards and inventories.
Dashboard	Dashboard pages (e.g. Default Dashboard (on page 292)) are used when information needs to be grouped together into related sets of resources and functions. From a dashboard it is possible to view the inventories of resources.
Inventory	Inventory pages (e.g. Nodes Inventory (on page 536)) list collections of resources of the same type. They can be filtered to display the entire set or a specific subset of those resources. From an inventory page it is possible to create, view and delete resources and perform other available actions on them.

Page Type	Description
Detail	Detail pages (e.g. Node Details (on page 650)) display all the properties of a particular resource. If a resource is a composite of other resources (e.g. policies are a composite of classifications and operations) then its detail page may contain inventories of those sub-resources. From a detail page it is possible to edit the resource and perform other available actions on them.
Wizard	Wizard pages (e.g. Node Type Wizard (on page 539)) are an ordered collection of dialogs that guide the user through the creation or modification of a resource. A dedicated wizard is provided for creating and editing each type of resource. If a resource is a composite of other resources then its wizard may invoke wizards for creating those sub-resources.

User Interface Page Layout

The Ops Center Protector user interface is displayed in a web browser window and is arranged in the following way:

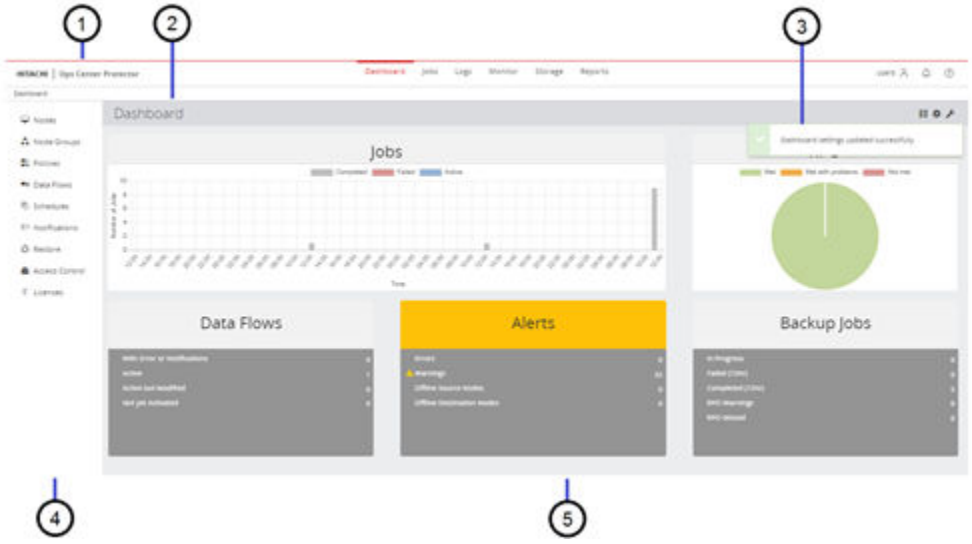


Figure 72 User interface page layout

Control	Description
1. Banner	The banner remains visible and active at all times. It contains links to frequently used monitoring and administrative functions. Refer to the Main Banner (on page 286) description for more details.

Control	Description
2. Breadcrumbs	The breadcrumbs remain visible and active at all times. They provide an indication of where you are in the hierarchy of pages and allow easy navigation to related pages. Refer to the Navigation Breadcrumbs (on page 290) description for more details.
3. Pop-up Session Notifications	One or more color coded session notification messages may appear down the right hand side of the page, these messages are displayed in response to user interaction via the UI: ▪ Green - Indicates that the command has been executed successfully. ▪ Blue - Indicates that the command has been actioned but the outcome is not yet known. ▪ Red - Indicates the command failed. The session messages are removed after a few seconds, but can be viewed in the Session Notifications list displayed via the Main Banner (on page 286). When a session notification message appears, check for entries in the Logs Inventory (on page 509) if further information is required.
4. Sidebar	The sidebar is available at all times. It contains links to management, configuration and design functions. Refer to the Navigation Sidebar (on page 291) description for more details.
5. Page Area	All pages relating to the Protector UI appear in this area. Refer to the Inventory Page (on page 293) description for more details. Following logon, the Default Dashboard (on page 292) is displayed. It presents an overview of the current system status and contains links to the associated pages.

Main Banner

The main banner appears at the top of all pages in the user interface and provides links to all the frequently accessed monitoring and administration functions.

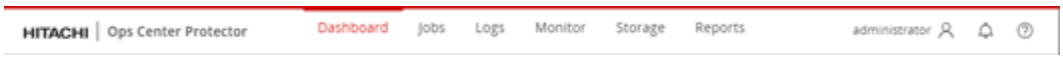


Figure 73 Main Banner

Control	Description
Dashboard	Opens the Dashboard (on page 292)
Jobs	Opens the Jobs (on page 493)

Control	Description
Logs	Opens the Logs (on page 509)
Monitor	Opens the Monitor (on page 520)
Storage	Opens the Storage Inventory (on page 858)
Reports	Opens the Reports (on page 739)
	The name of the currently logged in user is displayed here. Opens the menu from which the Settings Wizard (on page 287) Dialog and the logout option can be accessed.Click to logout.
 Session Notifications	Click to view the list of session notification messages. The messages listed here are a cache of the pop-up messages displayed in the web UI. The number of unread error notifications is displayed to the right of the button. All unread session notifications are considered to have been read when the list is opened then subsequently closed. All notifications can be cleared by selecting Clear at the top of the popup list.  **Note:** If you have a number of browser tabs open simultaneously, the notifications displayed here are local to the tab in which you are currently working.
 Help	Displays a menu from which the Online Help dialog (on page 289) and About dialog (on page 289) can be accessed.

Settings Wizard

Controls the presentation of information within the user interface.

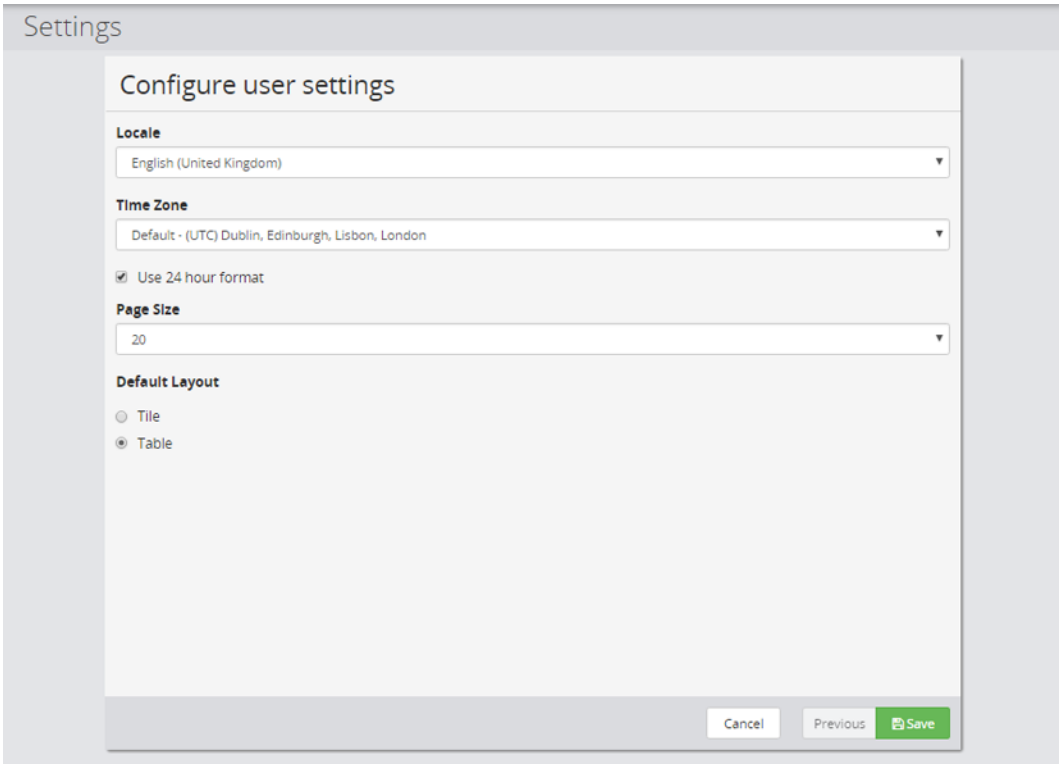


Figure 74 User Interface Settings

Control	Description
Locale	Selects the locale so that information is presented in accordance with the conventions of the associated culture.
Time Zone	The time zone to be used to display dates and times on the UI. The Master node's time zone is not changed.
Use 24 hour format	Displays time in HH:MM:SS format when checked, otherwise AM/PM format is used.
Page Size	Sets the number of tiles or table entries displayed per page in the Inventory Page (on page 293).
Default Layout	Selects tile or table view as the default in all Inventory Page (on page 293)s.
Cancel	Discards all changes made to the settings and reverts to those used prior to opening the Setting page.
Previous	Takes the user to the previous screen in the wizard.
Save	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

About dialog

This dialog is displayed when the Help > About option is selected from the Main Banner.



Figure 75 About Dialog

This dialog shows the build number for the Master node. You may need this when communicating with Hitachi Vantara Support.

Online Help dialog

This dialog is displayed when the Help > Online Help option is selected from the Main Banner.



Figure 76 Online Help Dialog

Click the link to open online help versions of the Protector User Guide in separate browser tabs. This also includes Application specific information.

Control	Description
	Link to open the online help versions of below documents in separate browser tabs. This also includes Application specific information. ▪ User Guide ▪ API Reference Guide

Navigation Breadcrumbs

The breadcrumbs appear at the top of all pages in the user interface and show the hierarchical location of the current page. This allows the user to navigate back to a parent page by clicking on the page name.

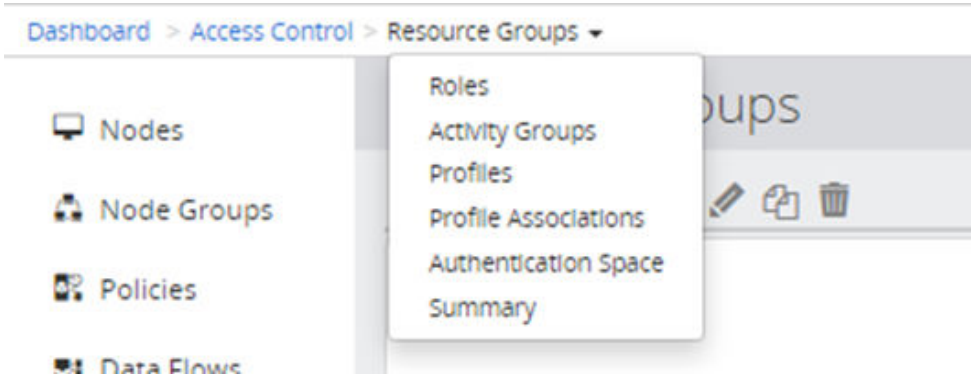


Figure 77 Breadcrumbs

Control	Description
Dashboard	The home page in the hierarchy (i.e. the Default Dashboard (on page 292)). Click here to navigate directly to the home page.
1st Level Sub-page	The 1st level sub-page (e.g. **Access Control**) Click here to navigate directly to this page.
2nd Level Sub-page	The 2nd level sub-page (e.g. **Resource Groups**). Click here to navigate directly to this page. If other related pages are available, these are shown in a pop-up menu by clicking the icon to the right of the parent page name. Click on one of the menu items to navigate directly to that page.
Current Page	The currently displayed page (e.g. **Create**).

Navigation Sidebar

The sidebar is displayed to the left of all pages in the user interface and provides links to all the frequently accessed management, configuration and design functions.

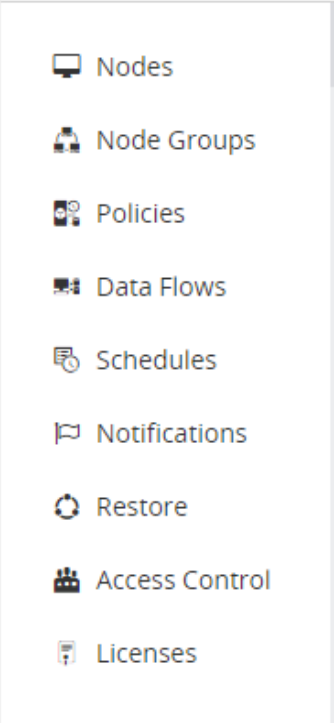


Figure 78 Navigation Sidebar

Control	Description
Nodes	Opens the Nodes Inventory (on page 536)
Node Groups	Opens the Node Groups Inventory (on page 531)
Policies	Opens the Policies Inventory (on page 670)
Data Flows	Opens the Data Flows Inventory (on page 355)
Schedules	Opens the Schedules Inventory (on page 849)
Notifications	Opens the Notifications Inventory (on page 659)
Restore	Opens the Restore Inventory (on page 763)
Access Control	Opens the Access Control Dashboard (on page 310)
Licenses	Opens the Licenses Inventory (on page 506)

Default Dashboard

The Default Dashboard is the home page for the Protector user interface. It displays summary information about system status including jobs, RPO, data flows, alerts and backups. It also provides links to the underlying pages where more detailed information can be found.

Note: When accessing the underlying pages from the Dashboard, a filter is automatically applied to the target page so that only the items indicated are displayed.

For example, if 5 Errors are indicated on the Alerts tile, clicking the corresponding text will open the Logs page from the Dashboard with a filter applied such that only error level logs are displayed.

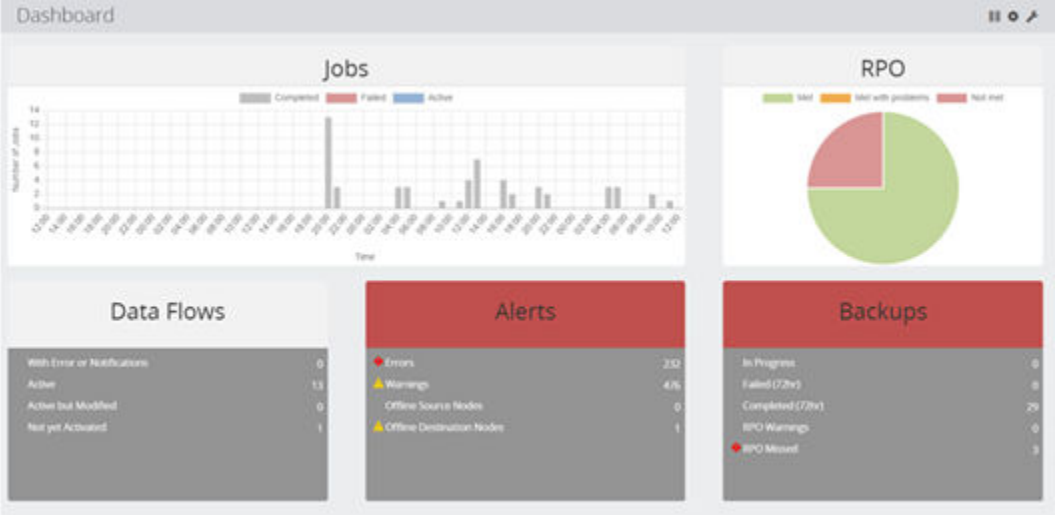


Figure 79 Main Dashboard

Control	Description
Jobs	Plots the number of completed, failed and active jobs by the hour. The chart can be filtered by clicking the legends shown at the top of the tile. Clicking on the coloured part of a bar in the time series opens the Jobs Inventory (on page 493) with the appropriate filter applied. **Note:** The data displayed here will be reset if the hub service is restarted on the master node.
RPO	Plots the RPO status for all backups. The chart can be filtered by clicking the legends shown at the top of the tile. Clicking on a segment of the pie chart opens the RPO Report (on page 741) with the appropriate filter applied.
Data Flows	Clicking on a counter opens the Monitor Inventory (on page 520) or Data Flows Inventory (on page 355) with the appropriate filter applied.

Control	Description
Alerts	Indicates the number of unacknowledged error and warning level log messages, and the number of offline nodes. Clicking on a counter opens the Logs Inventory (on page 509) or Nodes Inventory (on page 536) with the appropriate filter applied.
Backups	Indicates the status of backup jobs. Clicking on a counter opens the Tasks Inventory (on page 501) or RPO Report (on page 741) with the appropriate filter applied.
Pause/Resume Dashboard Transition	When additional, user defined dashboards are added to the dashboard set, this button pauses or resumes the periodic transition between them.
Configure User Dashboard Set	Opens the **Configure Dashboard Set** wizard, enabling user dashboards to be added to or removed from the dashboard set.
Manage Dashboards	Opens the **Dashboard Inventory**, enabling new dashboards to be created, edited or deleted. **Note:** Dashboard creation requires an in-depth understanding of the Protector REST API. Please contact professional services if you need to add new dashboards.

Inventory Page

Inventory pages are used frequently within the Protector UI to manage and access resources. They follow a common layout with controls grouped in a consistent way as illustrated below.

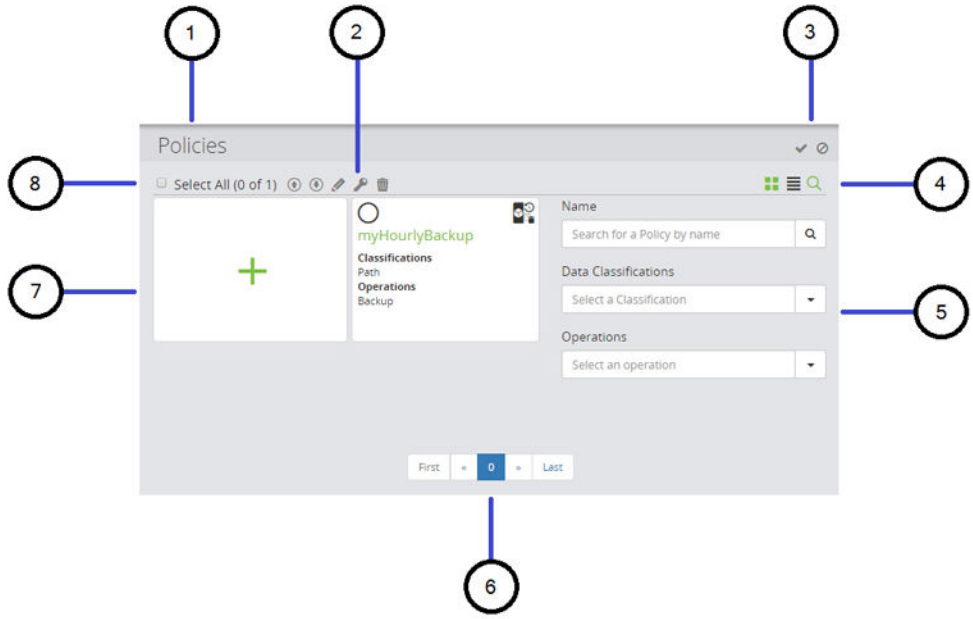


Figure 80 Inventory Page general arrangement (for illustration only)

Item	Description
1. Page Title	Identifies the function or scope of the page.
2. Tile Related Controls	Buttons displayed here represent commands that can be applied to items selected in the Tile/Table Area immediately below them. The buttons may be enabled (displayed in black) or disabled (displayed in grey) depending on the number and type of tiles selected.
3. Page or Resource Related Controls	Buttons displayed here represent commands that apply to the page in general (e.g. email sender account settings for notifications).
4. Layout Controls	 Tile View Information about the resources displayed in an inventory are presented in tile mode.
	 Table View Information about the resources displayed in an inventory are presented in table mode.
	 Filters Click this button to toggle the filter controls (5) between hidden or shown.
5. Filter Controls	The controls displayed here allow the results displayed in the Tile/Table Area to the left to be filtered based on the criteria selected. The filter controls can be hidden or displayed (4). Each inventory has an associated set of filters. In the majority of cases the filters are applied as you select them. For Restore inventories the filters are applied once the Search button is clicked.  **Note:** If the display area is limited, the filter controls will be positioned above the Tile/Table Display Area (7).
6. Tile/Table Page Controls	When the number of items returned for the given filter criteria exceeds one page, these buttons are displayed to allow more results to be loaded.

Item	Description
7. Tile/Table Display Area	The items returned for the given filter criteria are displayed here as separate tiles or rows in a table (depending on the display mode selected). Refer to Tile Control (on page 295) for more details about the controls displayed on tiles
8. Select All Items Checkbox	All inventory items passing the applied filter criteria and displayed on the current page only are selected/deselected.

Tile Control

Tiles are used to display the status and details of a resource in an inventory. They also contain a link to the details page for the resource in question and a checkbox to enable the resource to be selected.

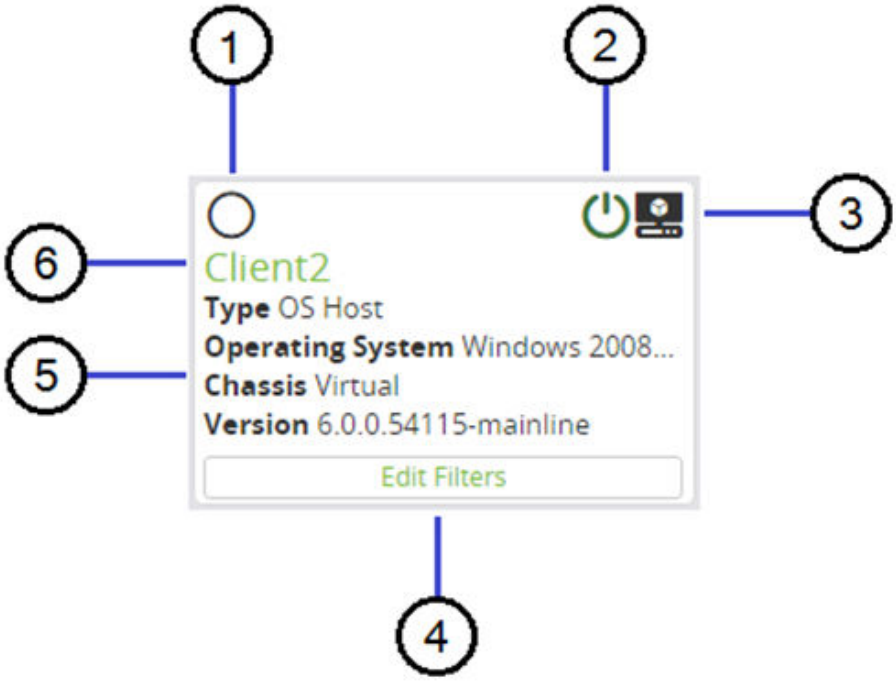


Figure 81 Tile general arrangement (for illustration only)

Item	Description
1. Select/Deselect Checkbox	This control is used to select/deselect tiles. Selected items display a check mark in this control and will have the commands (selected in the control bar above the tile area) applied to them.

Item	Description
2. Supplementary Icons	For some tile types, supplementary icons may be displayed here to convey additional information (e.g. status).
3. Item Type Icon	Each tile will have a specific icon displayed here to indicate the type of item to which it corresponds.
4. Additional Hyperlink	For tiles that contain other items (e.g. Classifications may contain Filters) additional links to those items may me displayed here. Click the link to open the corresponding page.
5. Item Description and Parameters	The key properties of the resource are displayed here.
6. Item Name and Hyperlink	The name of the item is displayed here. Click on the name to open the corresponding details page.

Button Icons

The following table lists all button icons used in the UI, their control names and their associated functions. The control names are used in the step by step guides described in Tasks (on page 213) and Data Protection Workflows (on page 130).

Control Name (Listed Alphabetically)	Button Icon	Example
Acknowledge All Error Logs		See Logs Inventory (on page 509).
Acknowledge Selected Logs		See Logs Inventory (on page 509).
Activate		See Data Flows Inventory (on page 355).
Add to additional Host Groups		Hitachi Block Replications Inventory (on page 878)
Analyze		See Generation 1 Repository Details (on page 921).
Authorize		See Nodes Inventory (on page 536).

Control Name (Listed Alphabetically)	Button Icon	Example
Back		See Monitor Details (on page 521).
Clone		See Access Control Profiles Inventory (on page 330).
Condensed View		See Logs Inventory (on page 509). (Green when selected, black when deselected)
Configure Email Settings		See Notifications Inventory (on page 659).
Connect To		See Data Flow Wizard (on page 361).
Create a new item		See Inventory Page (on page 293).
Create Log Notification		See .Logs Inventory (on page 509)
Deactivate		See Data Flows Inventory (on page 355).
Deauthorize		See Nodes Inventory (on page 536).
Decrease Priority		See Notifications Inventory (on page 659).
Delete		See Data Flows Inventory (on page 355) etc.
Dissociate		See Hitachi Block Replications Inventory (on page 878).
Edit Permissions		See Data Flows Inventory (on page 355) etc.
Edit		See Data Flows Inventory (on page 355) etc.

Control Name (Listed Alphabetically)	Button Icon	Example
Tags		See Universal Tags (on page 39)etc.
Export		See Logs Inventory (on page 509).
Extended View		See Logs Inventory (on page 509). (Green when selected, black when deselected)
Help		See Main Banner (on page 286).
Hide/Show Search		See Inventory Page (on page 293). (Green when selected, black when deselected)
Host Groups		See Hitachi Block Device Details (on page 859).
Increase Priority		See Notifications Inventory (on page 659).
Journals		See Hitachi Block Device Details (on page 859).
Logical Devices		See Hitachi Block Device Details (on page 859).
Manage		See Node Details (on page 650).
More actions		See Jobs Inventory (on page 493).
Mount		See Hitachi Block Snapshots Inventory (on page 869).
Pause		See Hitachi Block Replications Inventory (on page 878).

Control Name (Listed Alphabetically)	Button Icon	Example
Pools		See Hitachi Block Device Details (on page 859).
Purge Logs.		See Logs Inventory (on page 509).
Refresh		See Generation 1 Repository Details (on page 921).
Remove from Host Groups		Hitachi Block Replications Inventory (on page 878)
Replications or Clones		See Hitachi Block Device Details (on page 859).
Restore		See Gen1 Repository Store Details (on page 915).
Resume		See Hitachi Block Replications Inventory (on page 878).
Resynchronize		See Gen1 Repository Store Details (on page 915).
Revert		See Hitachi Block Snapshots Inventory (on page 869).
Select/Deselect item		See Tile Control (on page 295).
Session Notifications		See Main Banner (on page 286).
Set Expiry Date / Change Retention		See Hitachi Block Snapshots Inventory (on page 869).
Settings		See Main Banner (on page 286).
Snapshots		See Hitachi Block Device Details (on page 859).

Control Name (Listed Alphabetically)	Button Icon	Example
Stores		See Gen1 Repository Stores Inventory (on page 912).
Swap		See Hitachi Block Replications Inventory (on page 878).
Transfer RBAC Permissions to Another Node		See Hitachi Block Replications Inventory (on page 878) .
Trigger Operation		See Monitor Details (on page 521).
Unmount		See Restore Inventory (on page 763).
Unsuspend		See Hitachi Block Replications Inventory (on page 878).
Upgrade Clients		See Nodes Inventory (on page 536).
User Actions		See Main Banner (on page 286).
User Snapshot		See Gen1 Repository Stores Inventory (on page 912).
Validate		See Gen1 Repository Stores Inventory (on page 912).
View as list		See Inventory Page (on page 293).
View as tiles		See Inventory Page (on page 293).
View Attachments		See Logs Inventory (on page 509).
View Group Nodes		See Monitor Details (on page 521).

Control Name (Listed Alphabetically)	Button Icon	Example
View Policy etc.		See Data Flow Wizard (on page 361). (Opens the associated item in a new tab)
View Session		See Logs Inventory (on page 509).
View The <node name> Node Details		See Generation 1 Repository Details (on page 921).
Expand		See Hitachi Block Journals Inventory (on page 863)
Remote Path Group		See Hitachi Block Remote Path Groups Inventory (on page 892)
Enable encryption		See Node Details (on page 650)

Universal Tags

As the number of managed objects increases within a system, it becomes progressively more difficult to manage large lists using only the object's native metadata. Universal tags provide a method for users to add extra metadata in the form of Tags to enhance the search and filter abilities of lists.

Tag-able objects

Tags can be applied to the following objects:

- Nodes
- Node groups
- Policies
- Policy Operations
- Data Flows
- Recovery Points

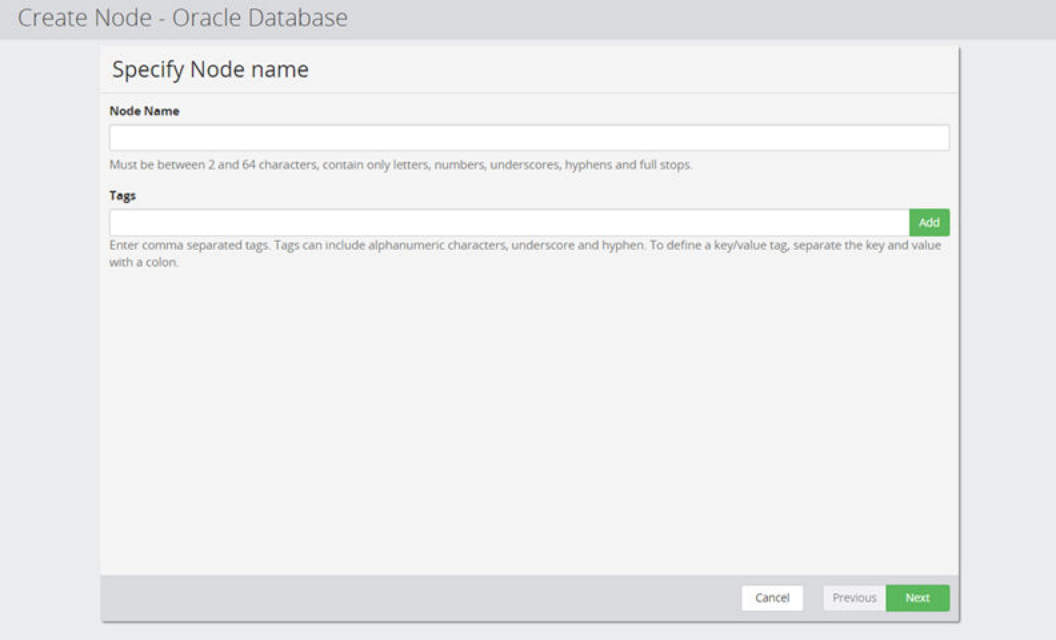


Figure 82 Create Node - Oracle Database

Control	Description
Node Name	Enter a name for the Node.
Tags	Add the tags to be associated with the object being created.
Cancel	Discards all changes made to the settings and reverts to those used prior to opening the current page.
Previous	Takes the user to the previous screen in the wizard.
Save	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

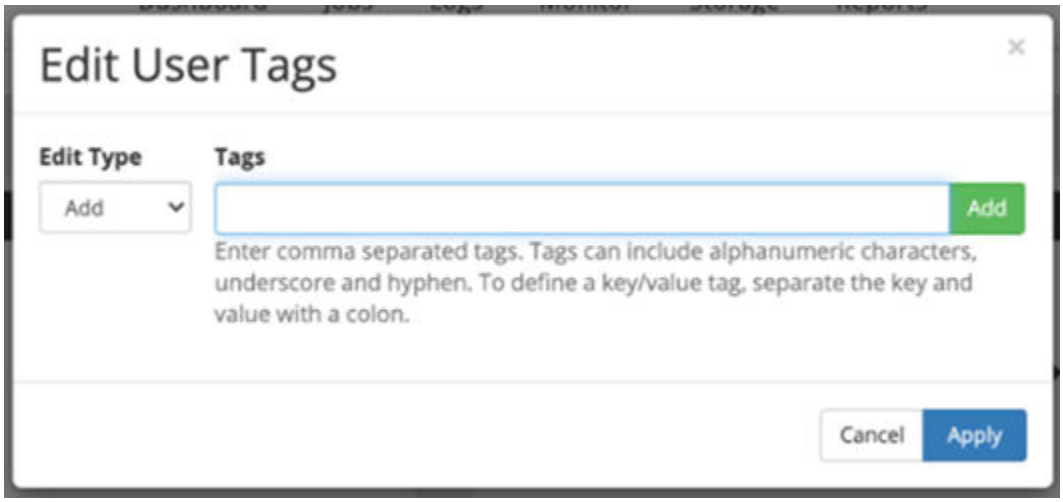


Figure 83 Add Tags - Edit User Tags

Control	Description
Edit Type	Enter the Edit Type.
Tags	Add the tags to be associated with the object being created.
Cancel	Discards all changes made to the settings and reverts to those used prior to opening the current page.
Apply	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

Tag inheritance and Propagation

Jobs and recovery points are objects that are generated by the system after a backup / snapshot operation.

To assist the manageability of these, they inherit the tags from the objects that were involved in creating them. Jobs and recovery points will inherit the tags from:

- The source node they represent in the data flow
- The node group the source node is in
- The storage node that stores the recovery point
- The data flow that created the recovery point
- The policy that created the recovery point
- Custom tags that were included if manually triggered

Inherited tags are assigned at the point that the jobs or recovery points are created and stored as independent tags against those objects. The tags assigned due to inheritance persist unchanged even if the originated object's tags later change, e.g. if a recovery point has inherited the tag "Production" from the Application Node, later removing that tag from that Application Node will not affect existing recovery point.

Note also that new recovery points and jobs will also only inherit any tags changes once the affected data flow have been re-activated. So if you were to alter the tags on a node, new recovery points will not reflect those changes until you re-activate the data flow that have those nodes.

You can add and remove tags on recovery points. You can modify tags on Recovery Points by either selecting the Recovery Point on the Restore Screen and selecting the "Edit Tag" icon, or viewing the Recovery point record and selecting the "Edit Tag" icon.

Types of Tags and Format

There are two types of tags; simple tags and key value pair tags.

Simple tags are straight forward labels. You may for instance tag all your production source machines with the tag “production”. You could then go to the RPO report and check the RPO only for machines that are marked “production”.

There are also key value pair tags. They are in the form of *key:value* . The concept is the ability to categorise a tag. For instance you can assign nodes a location tag in the form of *location:paris* or *location:London*. If you lost connection to a data center, you would be easily see that all the nodes at a particular location are down.

Tags can only contain alphanumeric characters, underscores and spaces. The “:” in a tag will designate it to be a key value pair. Capitalisation of tags is preserved but from a search and filter point of view ignored.

Search Filter using User Tags Query box

Many inventory screens have the ability to perform filter queries using tags. On those screens you will be provided with an entry box for adding a tag to search on. The search will return results that contains any tags that are in the search query, e.g. this search is an "or" operation on the selected tags. The tag needs to be an exact match the tag or the key value of the tag. The search however is not case sensitive.

So if you had an entry tagged: "Location:London" and "production" the following is the behaviour of the searches:

- Searching for tag: "Production" will return the entry because it is a case insensitive match.
- Searching for tag: "Prod" will not return the entry as this is a partial tag
- Searching for tag: "location:" will return the entry because the key value matches. "London" is not required
- Searching for tag: "location" will return the entry because the key is treated as a simple tag as well as key value pair tag
- Searching for tag: "Location:Paris" will not return the entry as the value doesn't match

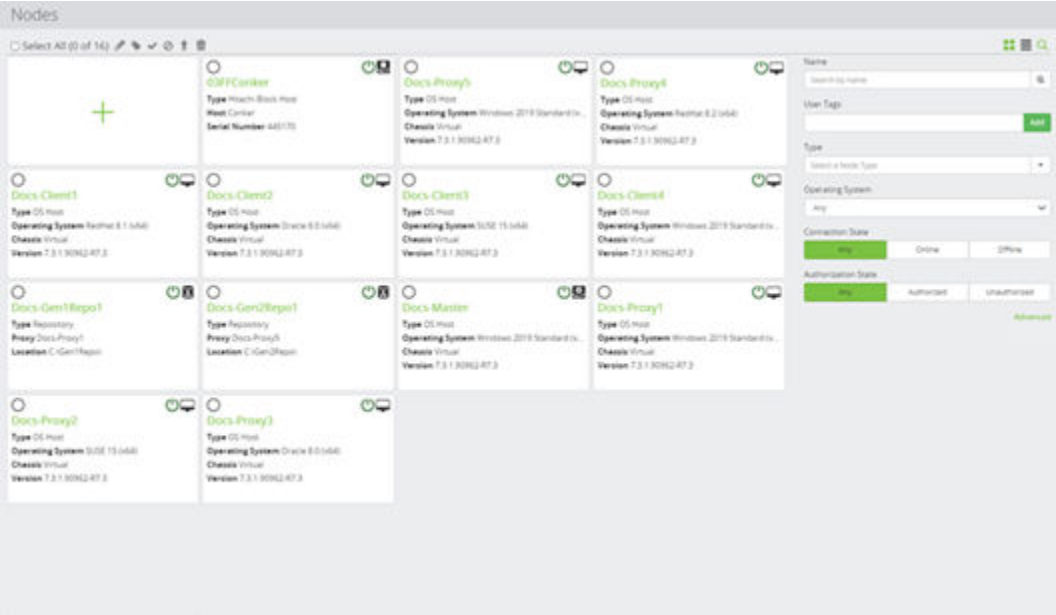


Figure 84 Nodes Inventory

Control	Description
Edit	Edits an existing node in the inventory. The Node Type Wizard (on page 539) is launched to enable the node's attributes to be changed.
Tags	Modifies the tags of an existing object from the either the inventory screen or the details screen of the object.
Authorize	Enabled only if one or more unauthorized nodes is selected in the inventory. Attempts to authorize the selected nodes with the Master node. Only nodes that have been authorized by the Master node may perform Protector functions. **Note:** If an attempt is made to authorize an inactive or unknown node, or if the master node fails to communicate with the node, an error log is generated and the node remains unauthorized.
Deauthorize	Enabled only if one or more authorized nodes is selected in the inventory. Attempts to deauthorize the selected nodes. Nodes that have been deauthorized cannot perform Protector functions.

Control	Description
	Note: ▪ Before deauthorizing nodes that are no longer required, they must first be deleted from the current data flow definitions, then any modified active data flows must be reactivated. If any attempt is made to deauthorize a node that is implementing rules in the currently active data flows, an error message is generated for each such node and the node remains authorized. ▪ It is possible to de-authorize a node with active mounts. This should be avoided because once de-authorized it is not possible to unmount.
 Upgrade Clients	Enabled only if one or more authorized nodes is selected in the inventory. Attempts to remotely upgrade the Protector software installed on the selected nodes. The required upgrade installer and configuration files must be present in the `C:\Programs Files \Hitachi \Protector\runtime\updater` folder on the Master node. **Note:** ▪ Only *OS Host* nodes can be upgraded. It is recommended to upgrade nodes in batches of 20. ▪ It is recommended to manually upgrade the master node.
 Delete Node	Enabled only if one or more nodes is selected in the inventory. The node is deleted from the inventory. **Note:** If an *OS Host* node is still running the Protector hub process and is configured to use the current *Master* node, then the node will re-appear as an unauthorized node as it periodically reconnects to the *master* node. Protector should be uninstalled from the node to stop this periodic reconnection.
 Create a new item	Adds a new node to the inventory. The Node Type Wizard (on page 539) is launched to guide you through the process.
Existing Node(s)	Nodes on which Protector has been installed are automatically detected and listed here along side those that have been added by the user. The Node Details (on page 650) is displayed to enable the node's details to be viewed and edited.

Control	Description
	Note: ▪ DHCP renewal can cause temporary disconnection of a node. ▪ If the host of a virtual machine(s) creates a Windows restore point, then the virtual machine(s) can temporarily disconnect. ▪ If a node has been assigned to another master, it will not appear in the inventory.
Filter on Node Name	Filters the displayed results based on Node Name.
Filter on User Tags	Filters the displayed results based on Tags contained in the data flow.
Filter on Node Type	Filters the displayed results based on Node Type.
Filter on Operating System	Filters the displayed results based on Operating System.
Filter on Connection State	Filters the displayed results based on Connection State.
Filter on Authorization State	Filters the displayed results based on Authorization State.

Search Filter using advance Query

Using the Advance Query String you can get more control of how the queries work.

While the Tag Query box does an exact match of any of the tags. So in the Tag Query box you entered tags "vmware" and "location:new york" it would be equivalent to:

```
((userTagsKeys = "vmware") OR (userTags = "location:new york"))
```

There are however extra search methods available using the Advance Query. These include

Value	Description
Exact Match Include	entry = value
Partial Match Include	entry ~ value
Boolean OR	condition OR condition
Boolean AND	condition AND condition

If you wanted only to match any nodes with the string "vm" in the tag but only in location New York, you can use this custom query:

```
((userTagsKeys ~ "vm") AND (userTags = "location:new york"))
```



Figure 85 Nodes Inventory - Search Filter

Trigger Operation

Triggering Operations from the Monitor screen now have the ability to add a tag to the trigger. In a situation for instance where you wanted to create a backup of a system prior to the upgrade, you could do a trigger with a tag “Pre upgrade”. It will be clear then in the restore screen the purpose of this adhoc backup and its timing.

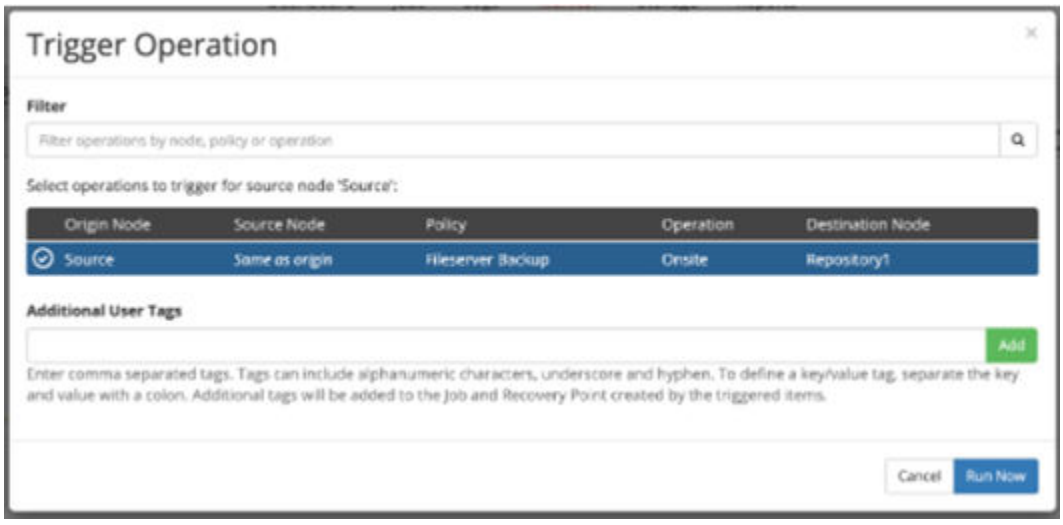


Figure 86 Trigger Operation

Permissions

If you have permission to modify an object e.g. Node, data flow, etc. as governed by the role based access rules, you will have the ability to add / remove tags. Tags belong to the object so any tags are visible to anyone who can view the object.

Tags Limitation

- There is no central inventory of existing tags. So currently there is no pick list or hint that a tag is in use.
- Generation 1 of the Repository does not fully support tags.

UI Timeout when Protector is linked to Ops Center

When Protector is linked to an instance of Ops Center the UI timeout will now be based on the user's actual activity on the UI (e.g. mouse / keyboard movement) rather than whether the UI has sent an API call to the back end.

There are two modes of operation which are governed by Ops Center settings:

1. Timeout after user inactivity regardless of which screens the user is on.
2. Timeout after user inactivity only if the UI does not refresh its information. This would mean a status or log page that periodically updates the screen would prohibit the UI from ever timing out. This is useful if the user wishes to leave screen up as a dashboard but may have unwanted security implication.

When the user is timed out, the UI will return the Protector logon screen. From the logon screen the user can elect to "Login with Ops Center" whereby the user would login depending on the current login state within Ops Center.

Access Control UI Reference

This section describes the Access Control UI, accessed via the Navigation Sidebar (on page 291).

For further information, refer to:

- Access Control Concepts (on page 122)
- Access Control Tasks (on page 213)

Login Page

The user enters their credentials on this page in order to gain access to the web based user interface.



Figure 87 Login Page

Control	Description
Username@AuthenticationSpace	Enter the username and Authentication Space in the format `Username@AuthenticationSpace`. Authentication Spaces are configured via the Access Control Authentication Spaces Inventory (on page 319).
Password	Enter the password for the given username
Log in	Click to login. If authentication succeeds, the Default Dashboard (on page 292) will be displayed. If authentication fails for any reason, a message will be displayed just below the Username and Password fields.
Login with	When Protector is integrated with Hitachi Ops CenterOps Center, this button is displayed to enable users to log in using SSO (Single Sign-On). The name of the OpenID Connect authentication space is displayed on the button. In this case, the Username@AuthenticationSpace and Password fields should be left unpopulated; the user being redirected to the appropriate SSO authentication page.

Access Control Dashboard

This dashboard enables the configuration of role based access control (RBAC) for users and groups who interact with Protector.

Tip: When configuring RBAC, start by defining Roles and Resource Groups first, then ACPs and Authentication Spaces. Finally define ACP Associations.

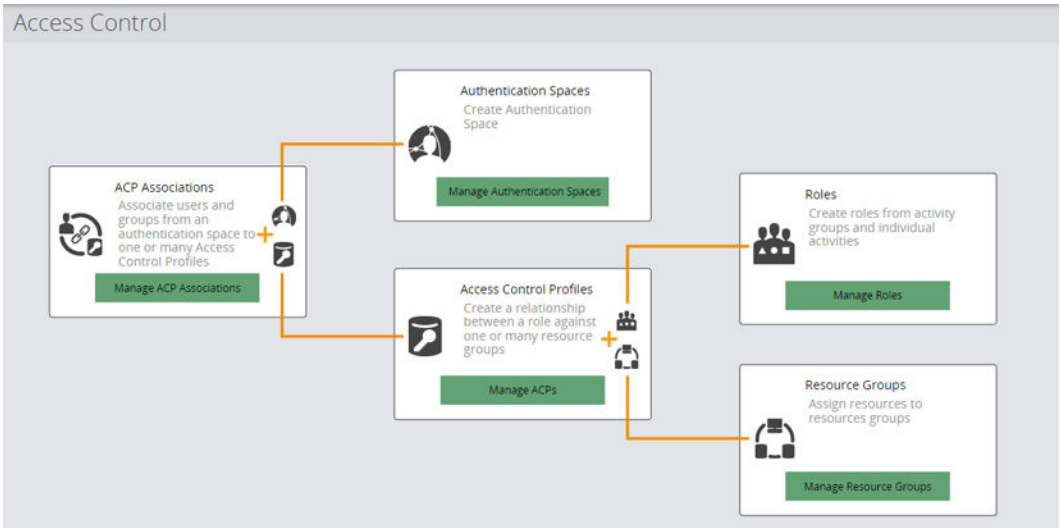


Figure 88 Access Control Dashboard

Control	Description
Manage ACP Associations	Opens the Access Control Profile Associations Inventory (on page 312)
Manage Authentication Spaces	Opens the Access Control Authentication Spaces Inventory (on page 319)
Manage ACPs	Opens the Access Control Profiles Inventory (on page 330)
Manage Roles	Opens the Access Control Roles Inventory (on page 336)
Manage Resource Groups	Opens the Access Control Resource Groups Inventory (on page 344)

Access Control Profile Associations Inventory

This inventory details all defined ACP Associations. ACP Associations link individual users, groups of users, or all users in an entire authentication space to one or more Access Control Profiles. This in turn governs what activities users are able to perform within Protector, and on which resources.

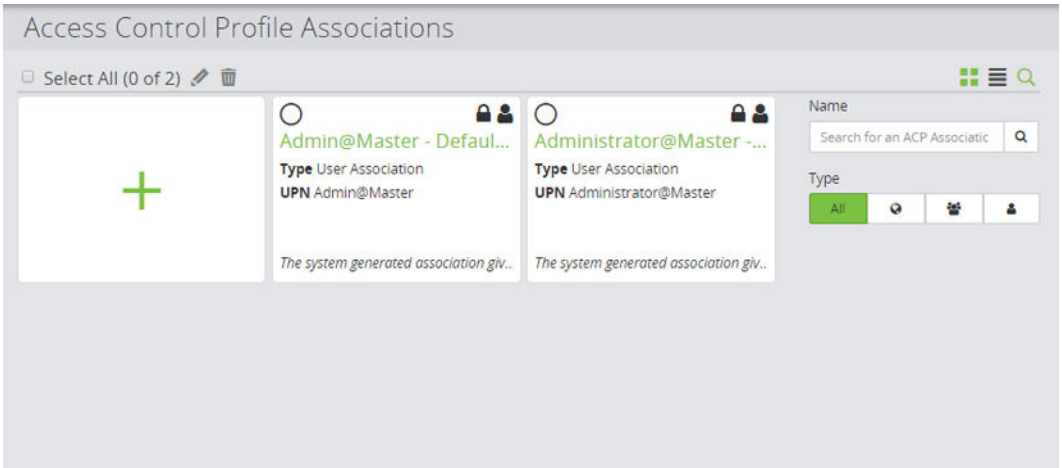


Figure 89 Access Control Profile Associations Inventory

Control	Description
Summary	Select the Summary option from the drop down menu in the Navigation Breadcrumbs (on page 290) to view the Access Control Summary (on page 348).
Edit	Edits an existing ACP Association in the inventory. The Access Control Profile Association Wizard (on page 313) is launched to enable the ACP Association's attributes to be changed.
Delete	Enabled only when one or more ACP Association is selected. Deletes the selected item from the inventory. The associated Authentication Space and ACP are not deleted.
Add	Creates a new ACP Association. The Access Control Profile Association Wizard (on page 313) is launched to guide you through the process.
System generated ACP Association(s)	At least one system generated ACP association is available when the product is installed. It associates an account chosen at installation time with a built-in ACP that provides the Protector Administrator role. This association cannot be deleted since it is necessary for initial login and configuration of access control by the application installer. System generated ACP Associations are marked with a icon to indicate that they cannot be modified. The Access Control Profile Association Details (on page 319) is displayed to enable the association to be viewed.

Control	Description
User defined ACP Association(s)	Any number of user defined ACP Associations can be created. These are displayed in the inventory and are marked with an Authentication Space, Group or User icon depending on whether the ACP is associated with an entire Authentication Space, a group or an individual user. ACP Associations must be defined in order to grant users access to the web and command line interfaces. The Access Control Profile Association Details (on page 319) is displayed to enable the association to be viewed and edited.
Filter on ACP Association Name	Filters the displayed results based on the name.
Filter on Type	Filters the displayed results based on the association type: ▪ All ▪ Authentication Space Association ▪ Group Association ▪ User Association

Access Control Profile Association Wizard

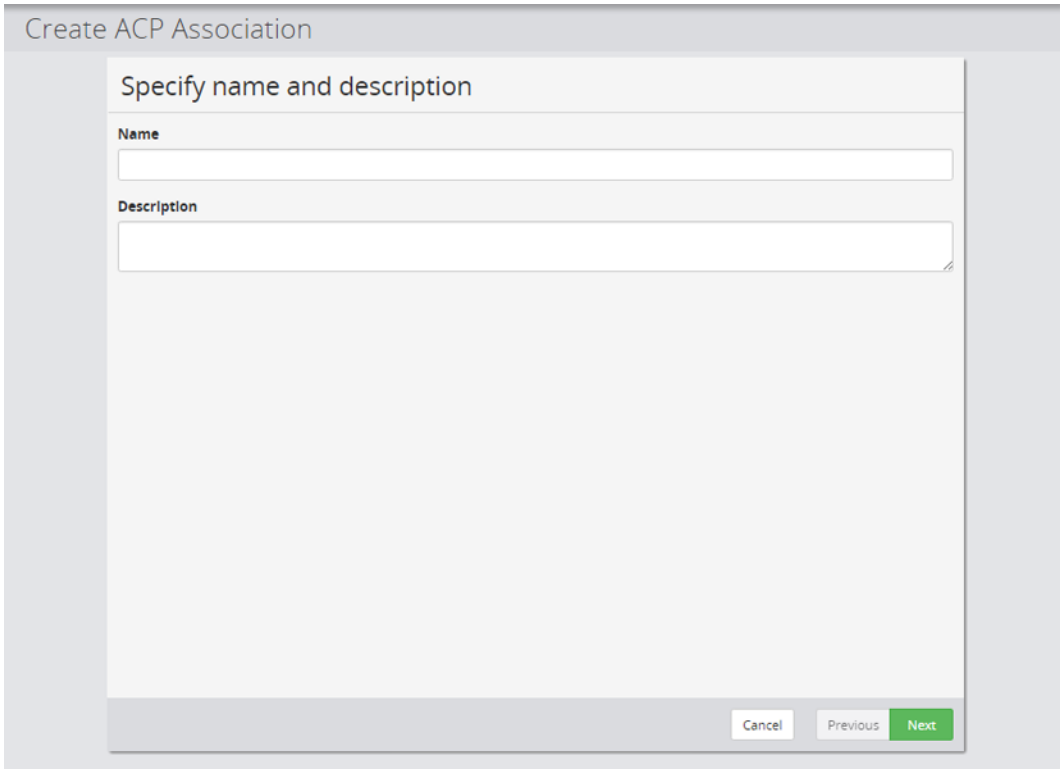


Figure 90 ACP Association Wizard - Specify name and description

Control	Description
Name	Enter the name of the ACP association.
Description	Optional. Enter a short description of the ACP association.
Cancel	Discards all changes and reverts to the previous page.
Finish	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

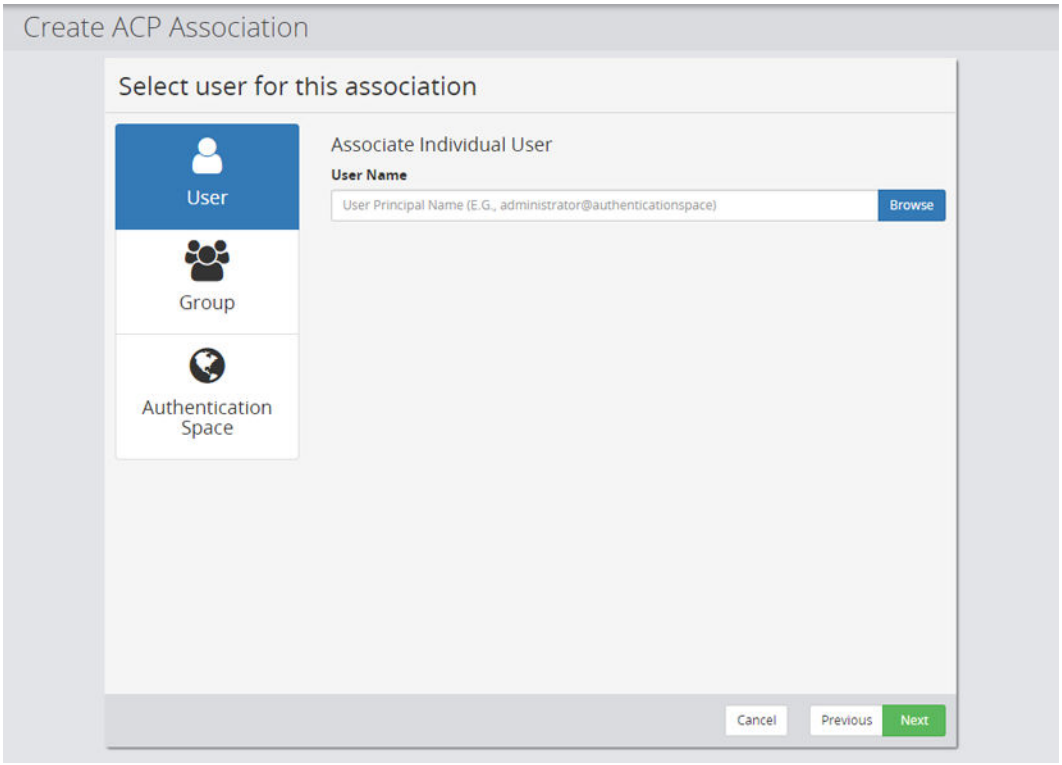


Figure 91 ACP Association Wizard - Select user for this association (User)

Tip: A *User ACP Association* would typically be created to give individual users the highest level of access, e.g. Administrator or Security Manager ACP.

Control	Description
User Name	The name of the user to associate with the ACP.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.

Control	Description
Finish	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

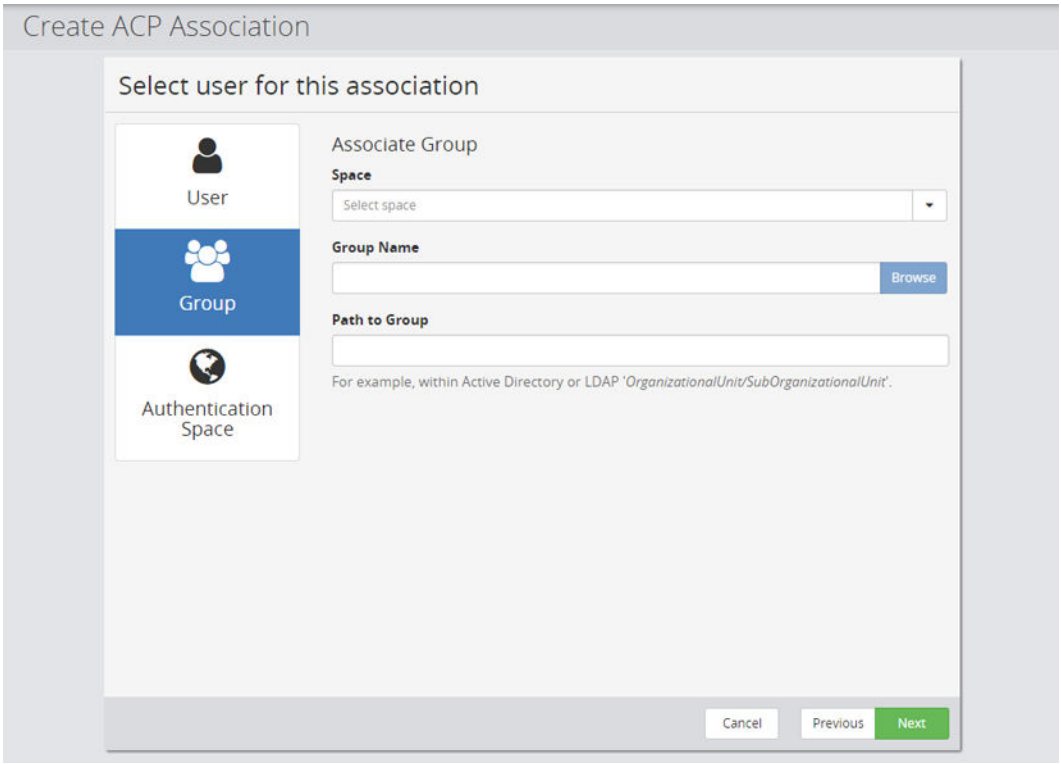


Figure 92 ACP Association Wizard - Select user for this association (Group)

Tip: A *Group ACP Association* would typically be created to give management groups a high level of access.

Control	Description
Authentication Space	The name of the Authentication Space from which the group is to be selected. Enter or select an Authentication Space from the drop down list, then click Browse to view a list of Groups for an Authentication Space. The selected Group will be automatically entered in the Group Name field below.
Group Name	The name of the group to associate with the ACP.
Path to Group	Provide the path to the required OU, using '/' as the path delimiter.

Control	Description
	For example, the following AD structure defines a *Managers* group within three different OUs: ▪ Contoso.com • Accounts ▪ Managers (Group) ▪ Goods In (Group) • Engineering ▪ Design ▪ Managers (Group) ▪ CAD Users (Group) ▪ Test ▪ Managers (Group) ▪ QA Engineers (Group) Here Path to Group must be set to one of the following, depending on which *Managers* group is required: ▪ *Accounts* ▪ *Engineering/Design* ▪ *Engineering/Test*
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Finish	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

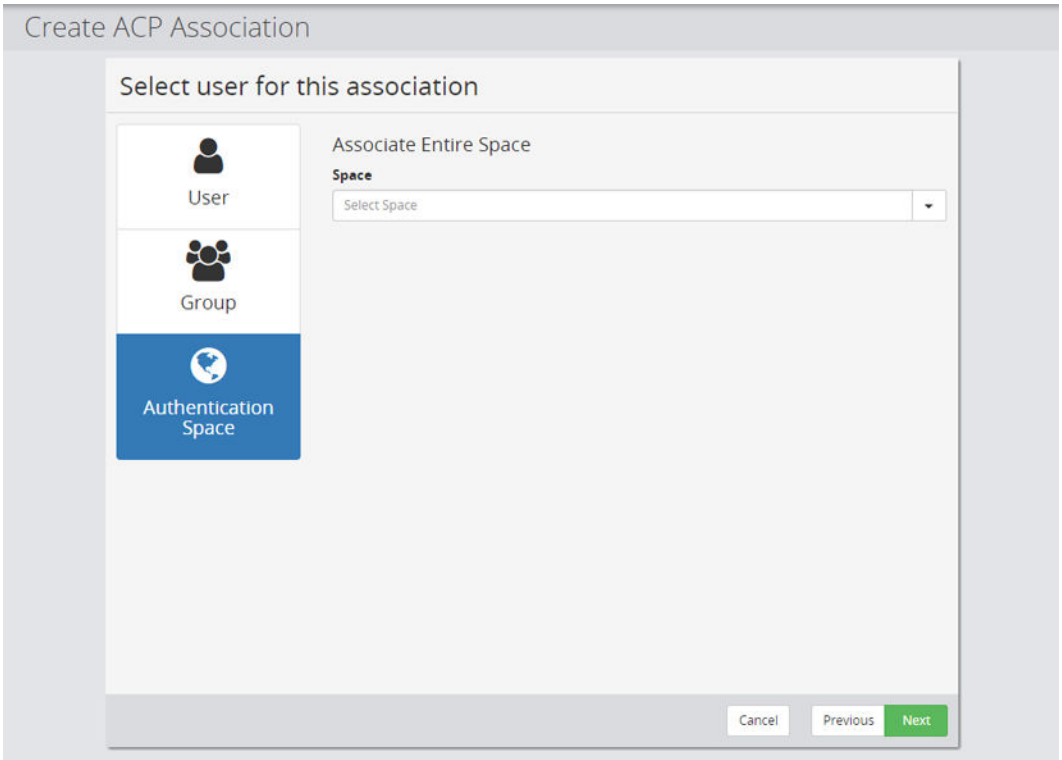


Figure 93 ACP Association Wizard - Select user for this association (entire Authentication Space)

Tip: A *Authentication Space ACP Association* would typically be created to give everyone the most limited level of access, i.e. Operator ACP.

Control	Description
Authentication Space	The name of the Authentication Space to associate with the ACPs. **Note:** All users and groups within the Authentication Space will be associated with the chosen ACPs.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Finish	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

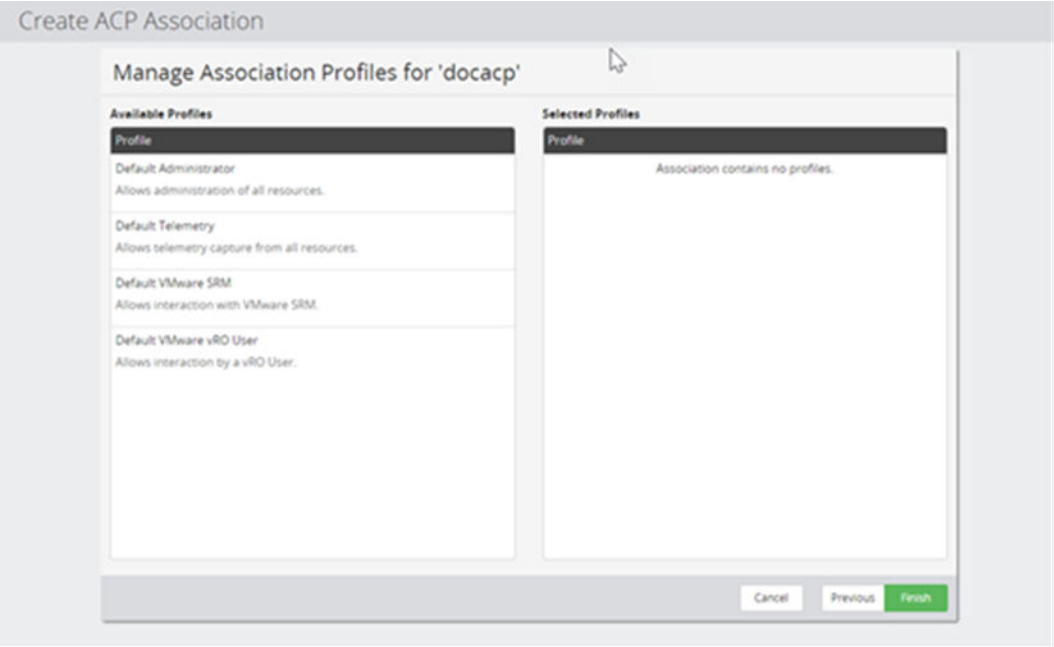


Figure 94 ACP Association Wizard - Manage Association Profiles

Control	Description
Available Profiles	List of available Access Control Profiles listed. Click on one or more of the available ACPs to add them to the ACP Profile.
Selected Profiles	List of selected Access Control Profiles listed. Click on one or more of the selected ACPs to remove them from the ACP Profile.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Finish	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

Access Control Profile Association Details

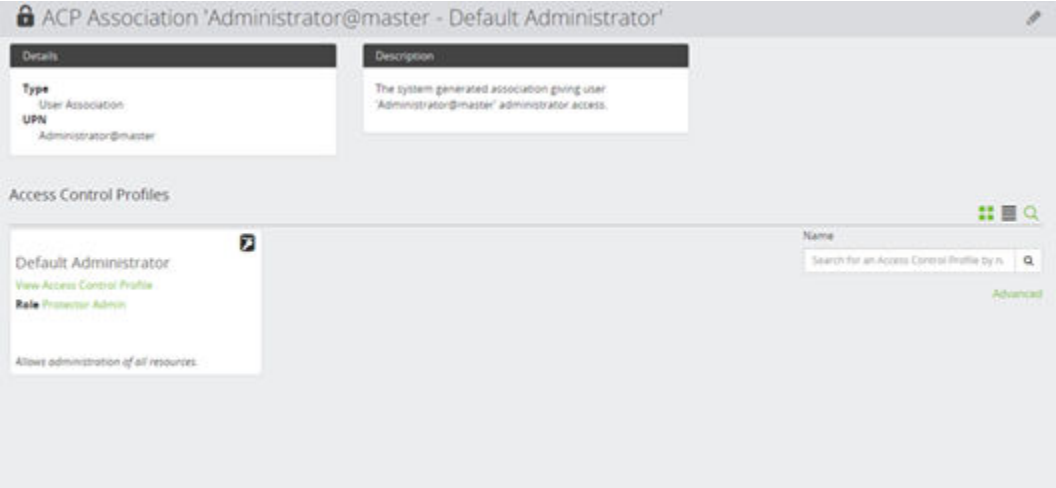


Figure 95 ACP Association Details

Control	Description
Edit	Launches the Access Control Profile Association Wizard (on page 313) to enable the group to be edited.
View Access Control Profile	Click on the link on the ACP tile to open the Access Control Profile Details (on page 335) to enable the ACPs to be viewed and edited.
Role	Click on the Role link on the ACP tile to open the Access Control Role Details (on page 340) to enable the Role to be viewed and edited.
Filter on Access Control Profile Name	Filters the displayed ACPs based on the name.

Access Control Authentication Spaces Inventory

This inventory list all defined Authentication Spaces. These specify authentication services that Protector uses to authenticate users when they login.

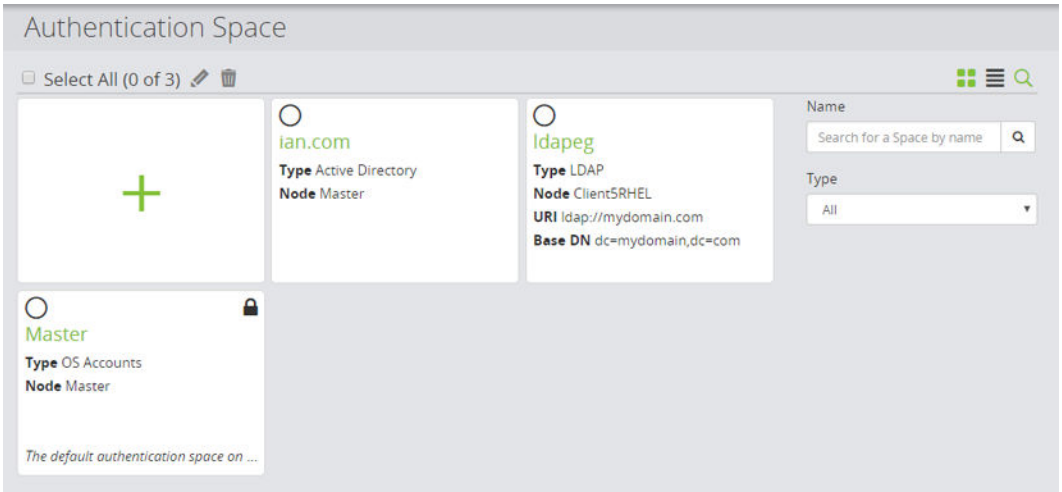


Figure 96 Authentication Spaces Inventory

Control	Description
Summary	Select the Summary option from the drop down menu in the Navigation Breadcrumbs (on page 290) to view the Access Control Summary (on page 348).
Edit	Edits an existing Authentication Space in the inventory. The Access Control Authentication Space Wizard (on page 321) is launched to enable the Authentication Space's attributes to be changed.
Delete	Enabled only when one or more Authentication Spaces is selected. Deletes the selected item from the inventory.
Add	Creates a new Authentication Space. The Access Control Authentication Space Wizard (on page 321) is launched to guide you through the process.
System generated Authentication Space	A system generated Authentication Space is available when the product is installed. It defines where the local administrator account(s) on the master machine is authenticated. This Authentication Space cannot be deleted since it is necessary for initial login and configuration of access control by the application installer. The system generated Authentication Space is marked with a icon to indicate that it cannot be modified. Click on the name of the Authentication Space to open the Access Control Authentication Space Details (on page 329).
OpenID Connect Authentication Space(s)	When Protector is integrated with Hitachi Ops Center, an OpenID Connect Authentication Space is automatically created to support single sign-on. Click on the name of the Authentication Space to open the Access Control Authentication Space Details (on page 329). **Note:** OpenID Connect Authentication Spaces cannot be created via the UI.

Control	Description
User defined Authentication Space(s)	Any number of user defined Authentication Spaces can be created. These are displayed in the inventory and can be based on Active Directory, Local Machine, RADIUS or Stand-alone LDAP authentication servers. Click on the name of the Authentication Space to open the Access Control Authentication Space Details (on page 329).
Filter on Authentication Space Name	Filters the displayed results based on the Authentication Space name.
Filter on Type	Filters the displayed results based on the Authentication Space server type.

Access Control Authentication Space Wizard

This wizard is launched when a new Authentication Space is added to the Authentication Spaces Inventory.

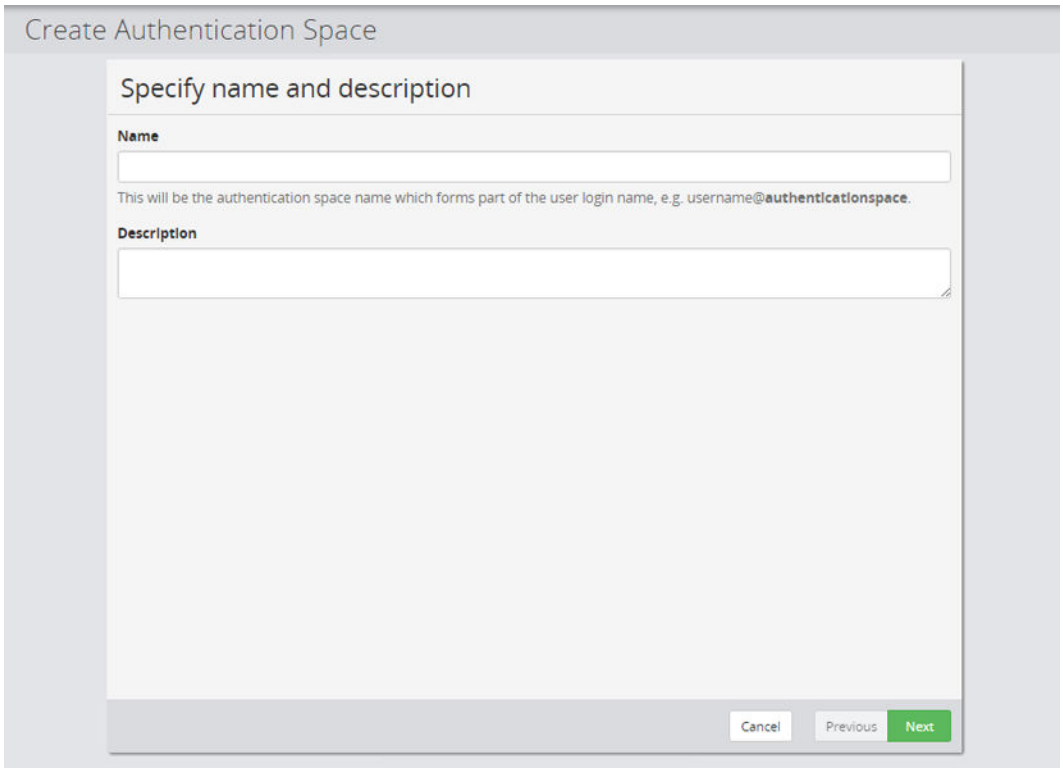


Figure 97 Authentication Space Wizard - Specify name and description

Control	Description
Name	Enter a name for the Authentication Space.
Description	Optional. Enter a short description of the Authentication Space.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Save	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

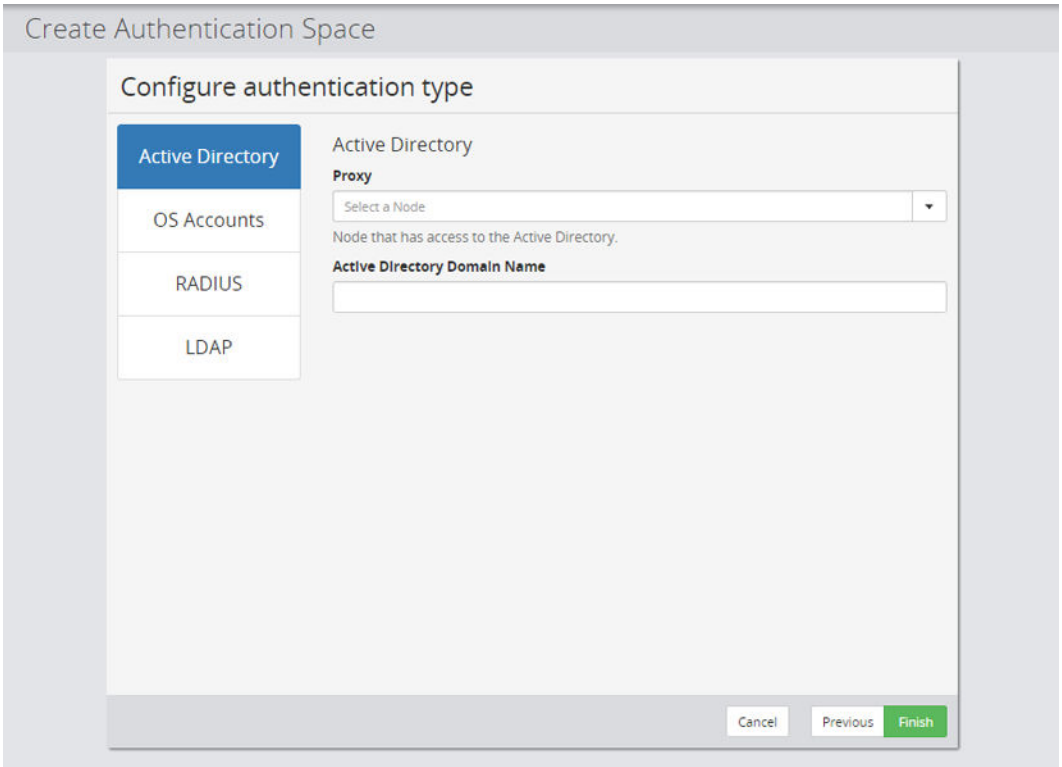


Figure 98 Authentication Space Wizard - Active Directory

Control	Description
Proxy	Type or select a Protector node that has a connection to the required Active Directory Service. **Note:** The AD Proxy must be a Windows node. **Tip:** To avoid the proxy becoming a single point of failure for authentication, select a clustered node where possible.
Active Directory Domain Name	Enter the AD domain name, e.g. `Contoso.com`
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Save	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

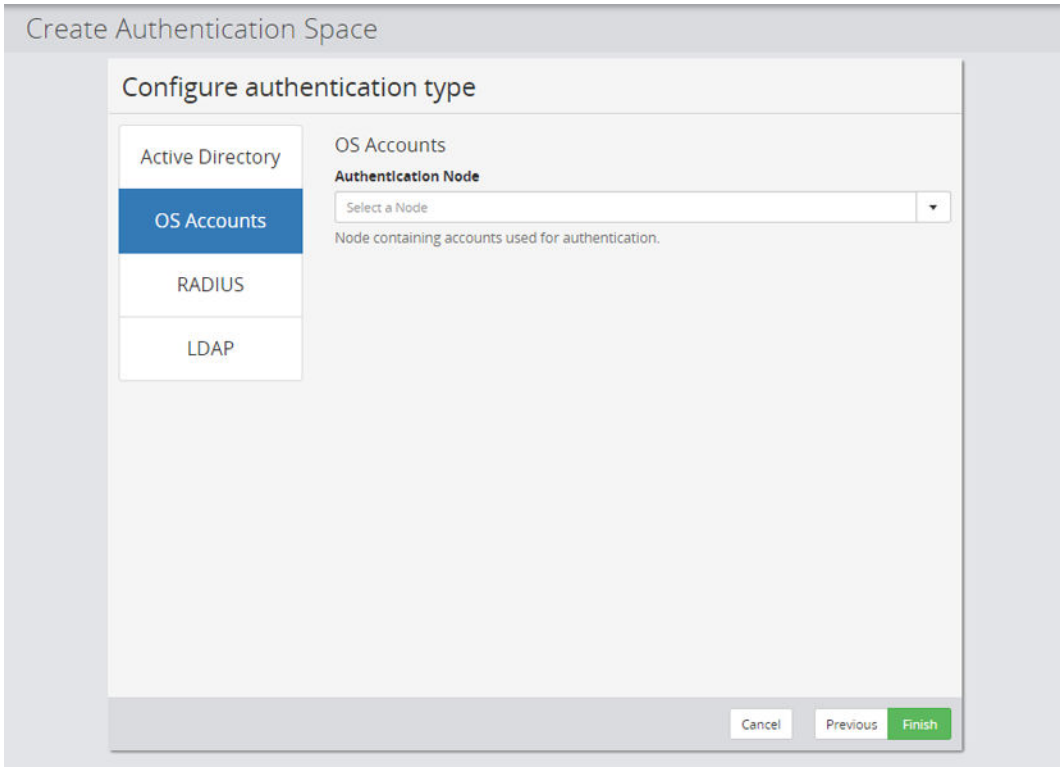


Figure 99 Authentication Space Wizard - OS Accounts

Control	Description
Authentication Node	Type or select a Protector node that will provide local authentication using the OS's authentication service.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Save	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

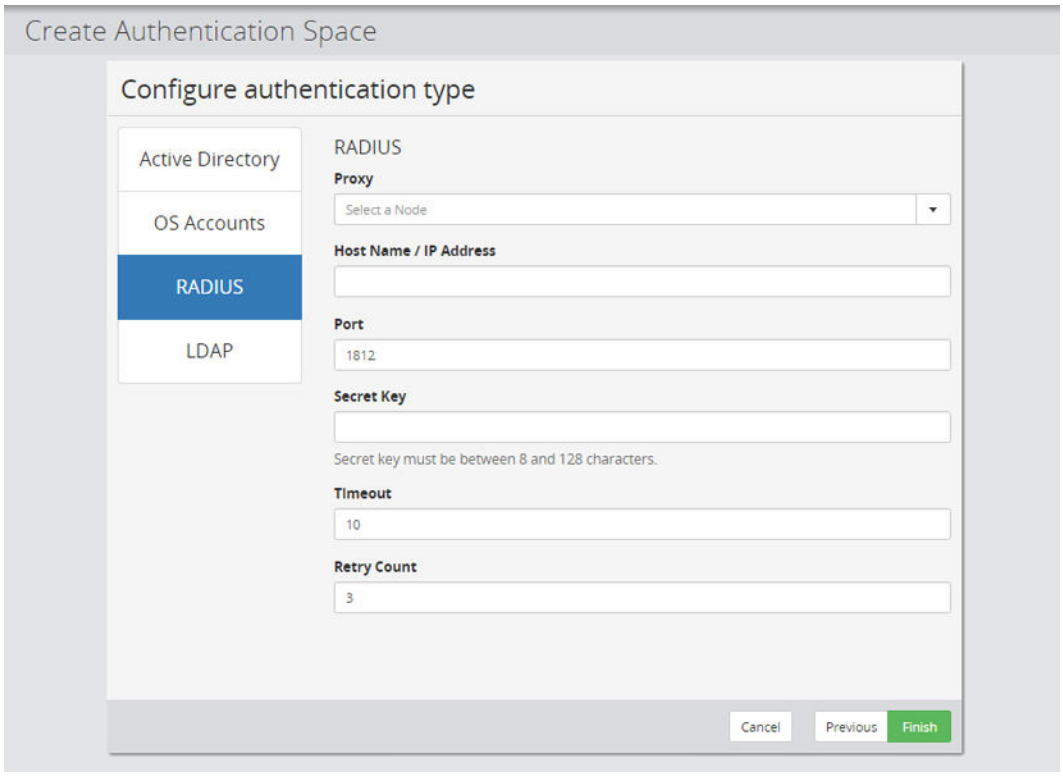


Figure 100 Authentication Space Wizard - RADIUS

Control	Description
Proxy	Type or select a Protector node that has a connection to the required RADIUS server.
Host Name / IP Address	Specify the IP address or DNS resolvable name of the required RADIUS server.
Port	Specify the IP port number or use the default port number (1812).
Secret Key	Specify the Secret Key for the RADIUS server.

Control	Description
Timeout	Specify the timeout period in seconds.
Retry Count	Specify the number of times a retry should be performed.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Save	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

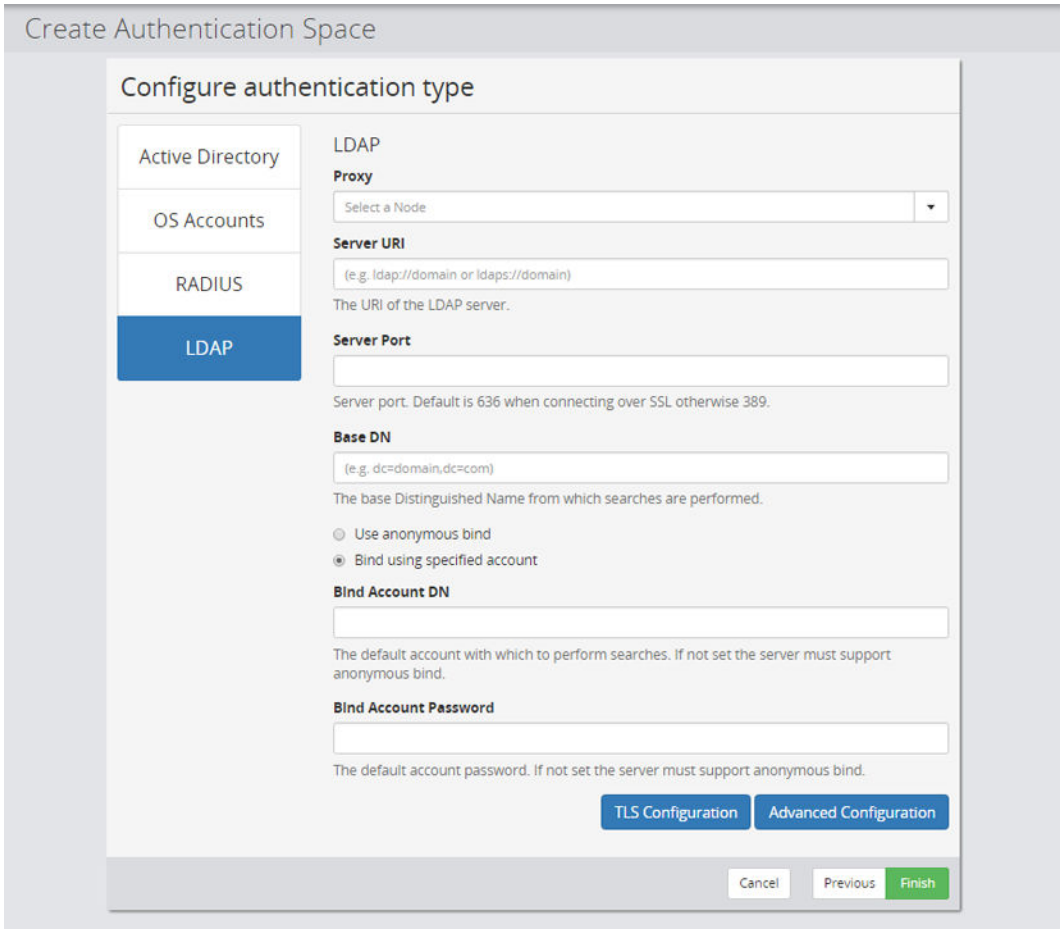


Figure 101 Authentication Space Wizard - LDAP

Control	Description
Proxy	Type or select a Protector node that has a connection to the required LDAP server.

Control	Description
	Note: ▪ The LDAP Proxy must be a Linux node. ▪ An LDAP Authentication Space cannot be used to authenticate via Microsoft Active Directory, despite similarities between the two technologies.
Server URI	Enter the URI of the required LDAP server in the format `ldap://domain` or `ldaps://domain`
Server Port	Enter the LDAP server port number if different from the default value. The default is 636 when connecting over SSL, otherwise it is 389.
Base DN	Enter the base Distinguished Name from which searches are performed. The default account distinguished name (DN) with which to perform initial searches. This is in LDAP DN format (e.g. `cn=Admin`, `ou=Users`, `dc=mydomain`, `dc=com`) The default account is needed to perform a lookup of a user's DN from their UID. Users log into Protector with a UID (e.g. `bmortimer@mydomain.com`) but the user's DN is needed for the LDAP bind and it is found using this account. If this value is not supplied the LDAP server must support anonymous bind.
Bind Using	Select how to bind to the server: ▪ Use anonymous bind ▪ Bind using specified account - enter the credentials below
Bind Account DN	Enabled only if Bind using specified account is selected. Enter the default account with which to perform searches.
Bind Account Password	Enabled only if Bind using specified account is selected. The default account password. This will be stored in an encrypted form within Protector until needed. If not set the server must support anonymous bind.
TLS Configuration	Click to specify TLS configuration options. See below.
Advanced Configuration	Click to specify advanced configuration options. See below.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.

Control	Description
Save	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

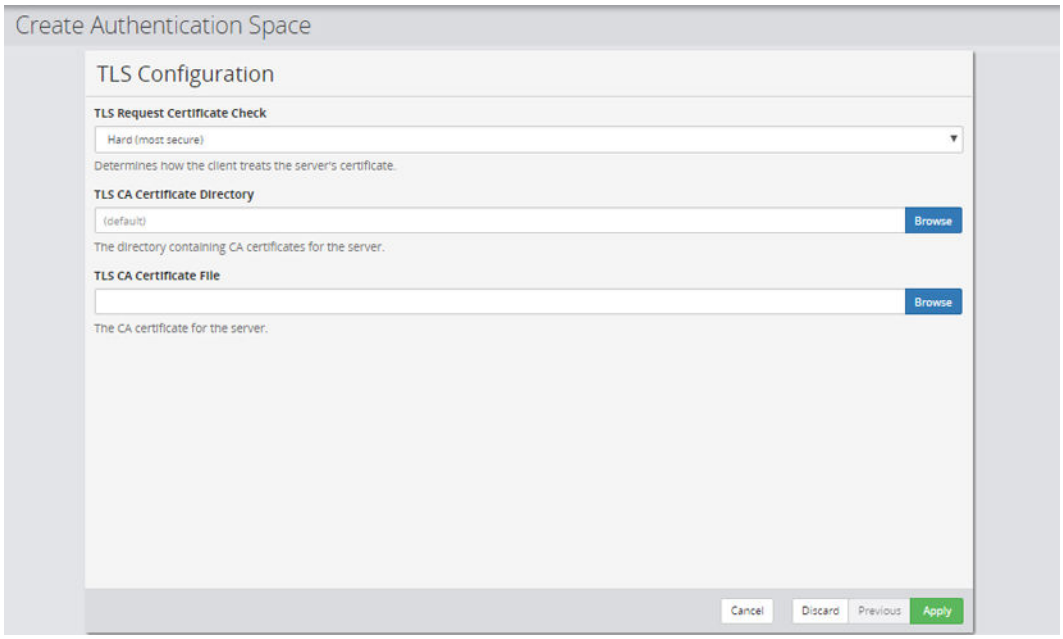


Figure 102 TLS Configuration

Control	Description
TLS Request Certificate Check	The TLS configuration in LDAP validates the way the Protector client treats the server certificate: ▪ Never - Least Secure. No certificate is requested. The Protector client will allow the use of self-signed certificates from the server. ▪ Allow - A certificate is requested. The TLS session completes normally even if a bad certificate or no certificate is provided. ▪ Try - A certificate is requested. The TLS session completes normally even if no certificate is provided. However if a bad certificate is provided, TLS session terminates. ▪ Demand - Hard Secure. This is the default setting. A certificate is requested. The TLS session terminates if a bad certificate or no certificate is provided. The LDAP server's certificate must be valid and signed by a trusted CA.

Control	Description
TLS CA Certificate Directory	Specifies path to the directory containing CA certificate files for the server.
TLS CA Certificate File	Specifies the CA certificate file for the server.
Cancel	Cancels all changes and reverts to the previous page.
Discard	Discards all changes and reverts to the Access Control Authentication Space Wizard (on page 321)
Previous	Takes the user to the previous screen in the wizard.
Save	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

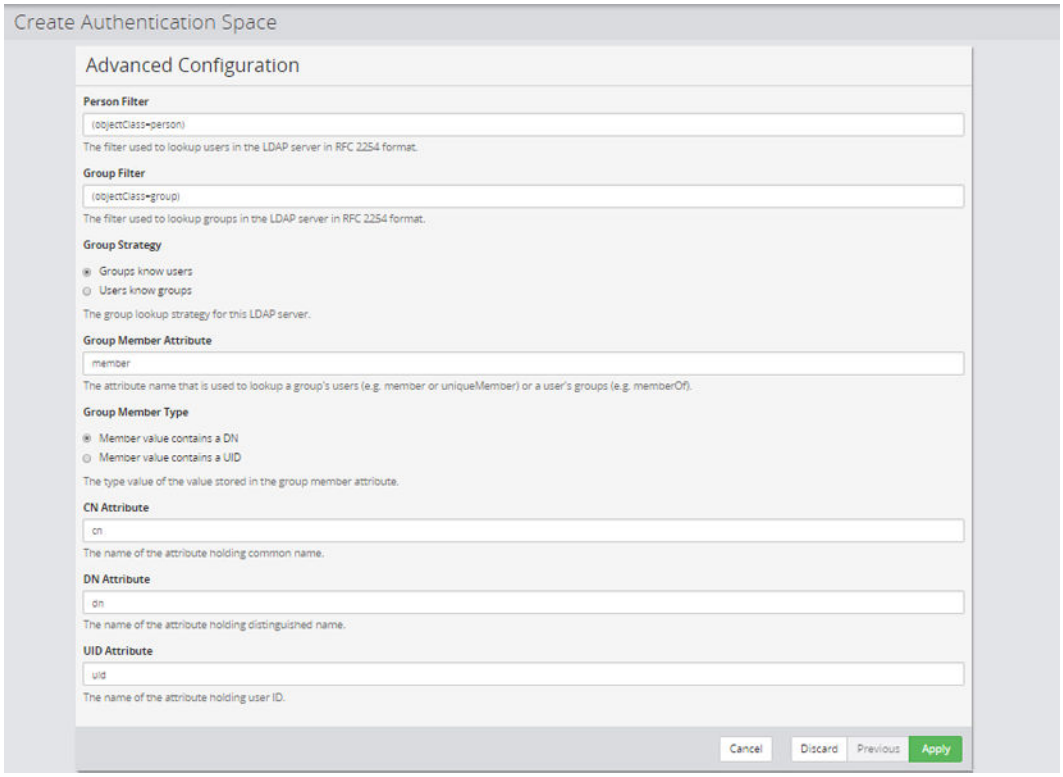


Figure 103 Advanced Configuration

Control	Description
Person Filter	Enter a search filter in RFC 2254 format to look up users.

Control	Description
Group Filter	Enter a search filter in RFC 2254 format to look up groups.
Group Strategy	Specifies the group look up strategy of the LDAP server: ▪ Groups know users - Each group knows its members, which are found using member attributes. ▪ Users know groups - Each user knows its groups.
Group Member attribute	Used to look for a group's users or a user's groups.
Group Member Type	Specifies value type stored in Group Member attribute: ▪ Member value contains a DN ▪ Member value contains a UID
CN Attribute	Name of the attribute holding Common Name (CN).
DN Attribute	Name of the attribute holding Distinguished Name (DN).
UID attribute	The number of the attribute holding the user ID with the default UID.
Cancel	Cancels all changes and reverts to the main page.
Discard	Discards all changes and reverts to the Figure 98 Authentication Space Wizard - Active Directory (on page 322) page.
Previous	Takes the user to the previous screen in the wizard.
Save	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

Access Control Authentication Space Details

This page displays the details of an Access Control Authentication Space and enables you launch the wizard to edit them.



Figure 104 Authentication Space Details

Control	Details
Edit	Launches the appropriate Access Control Authentication Space Wizard (on page 321) to enable the Authentication Space to be edited

Access Control Profiles Inventory

This inventory lists all defined Access Control Profiles (ACPs). These associate Roles to Resource Groups thus controlling what activities are allowed on each resource.

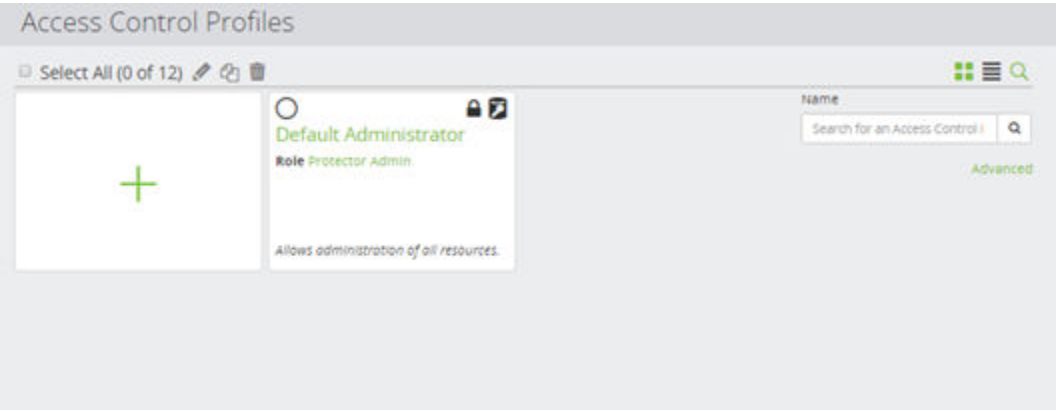


Figure 105 Access Control Profiles Inventory

Control	Description
Summary	Select the Summary option from the drop down menu in the Navigation Breadcrumbs (on page 290) to view the Access Control Summary (on page 348).

Control	Description
 Edit	Edits an existing ACP in the inventory. The Access Control Profile Wizard (on page 331) is launched to enable the ACP's attributes to be changed.
 Clone	Enabled only when one or more ACPs is selected. Creates a clone of the selected item which can then be modified. The clone is a shallow copy, in that it references the same Role and Resource Groups as the original.
 Delete	Enabled only when one or more ACP Association is selected. Deletes the selected item from the inventory. The associated Role and Resource Groups are not deleted.
 Add	Creates a new ACP. The Access Control Profile Wizard (on page 331) is launched to guide you through the process.
 System generated ACPs	At least one system generated ACPs are available when the product is installed. This associates system generated role Protector Administrator, with the default resource group. These ACPs cannot be deleted since they provide a basic level of access control. System generated ACPs are marked with a  icon to indicate that they cannot be modified. The Access Control Profile Details (on page 335) is displayed to enable the permissions to be viewed.
 User defined ACP(s)	Any number of user defined ACPs can be created. These are displayed in the inventory. ACPs should be defined in order to grant the required level of access to groups of resources as appropriate to the users' responsibilities. The Access Control Profile Details (on page 335) is displayed to enable the permissions to be viewed and edited.
Filter on Access Control Profile Name	Filters the displayed results based on the Access Control Profile Name.

Access Control Profile Wizard

This wizard is launched when a new ACP is added to the Access Control Profiles Inventory.

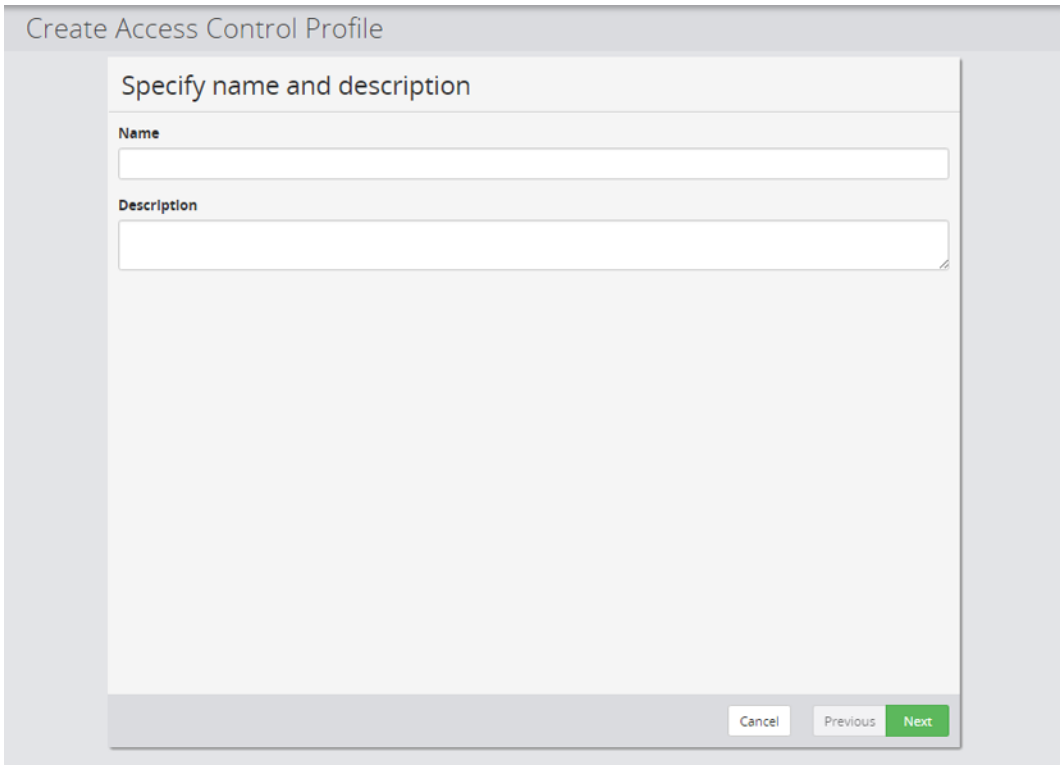


Figure 106 Access Control Profile Wizard - Specify name and description

Control	Description
Name	The name of the Access Control Profile.
Description	Optional. A short description of the ACP.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

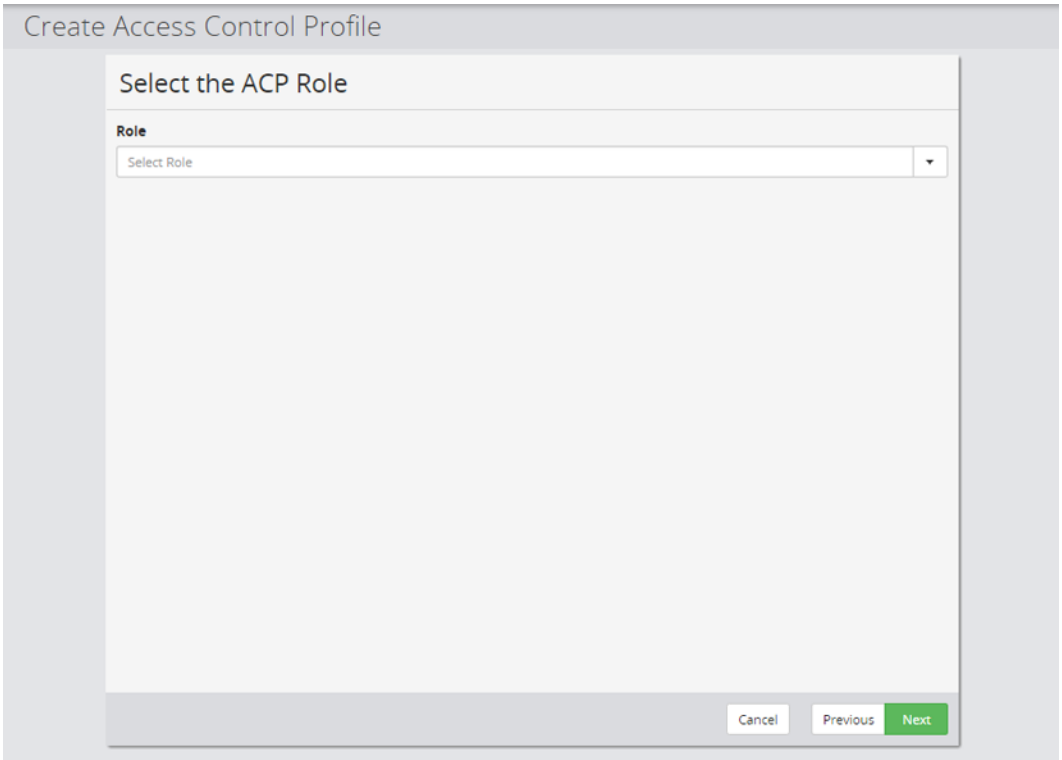


Figure 107 Access Control Profile Wizard - Select the ACP Role

Control	Description
Role	Select the role from the drop down list. The chosen role determines what activities owners of this ACP will be able to perform.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

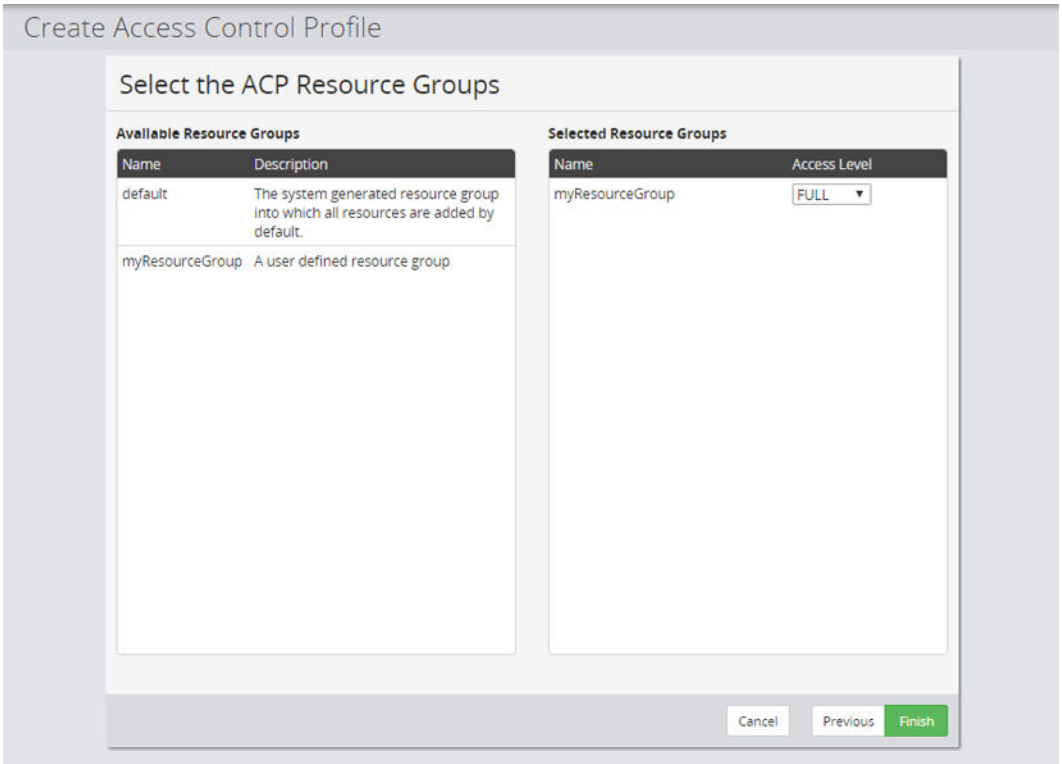


Figure 108 Access Control Profile Wizard - Select the ACP Resource Groups

Control	Description
Available Resource Groups	List of available Resources Groups. Click on one or more of the available Resource Groups to add them to the ACP.
Selected Resource Groups	List of selected Resources Groups. The role chosen in the previous page of the wizard can be performed on these resources. Click on one or more of the Resource Groups to remove them from the ACP.
Access Level	Select the required Access Level from the dropdown list to the right of the Selected Resource Groups entry. The access level controls visibility of backups (including any logs or reports relating to that backup) on storage nodes within the resource group as follows: ▪ LIMITED - a backup is visible if it originated from a node in the resource group. ▪ FULL - all backups, regardless of where they originated, are visible on any storage nodes within the resource group.

Control	Description
	Note: The access level only controls the visibility. What a user is able to do with visible backups is dependent on the activities they are allowed to perform on them. A user given LIMITED access to a storage destination included in their resource group will only see log messages for that storage destination if they pertian to their backups. A user given FULL access level will see all log messages for that destination.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Finish	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

Access Control Profile Details

This page displays the details of an Access Control Profile and enables you launch the wizard to edit them.

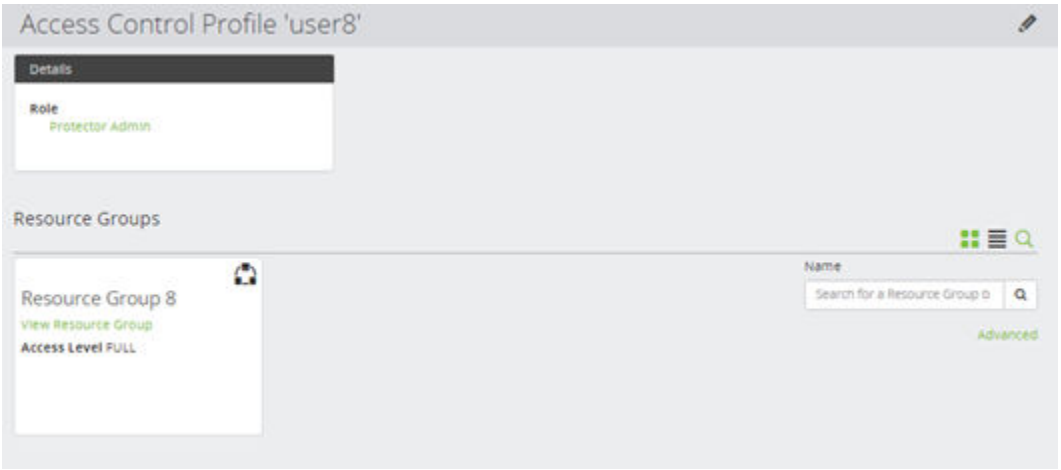


Figure 109 Access Control Profile Details

Control	Description
Edit	Launches the Access Control Profile Wizard (on page 331) to enable the ACP to be edited.

Control	Description
Role	Click on the role name link to open the Access Control Role Details (on page 340).
Resource Groups	Click on View Resource Group link on a tile to open the Access Control Resource Group Details (on page 347).
Filter on Resource Name	Filters the displayed resources groups.

Access Control Roles Inventory

This inventory lists all defined Roles. These roles define what activities are allowed to be performed.

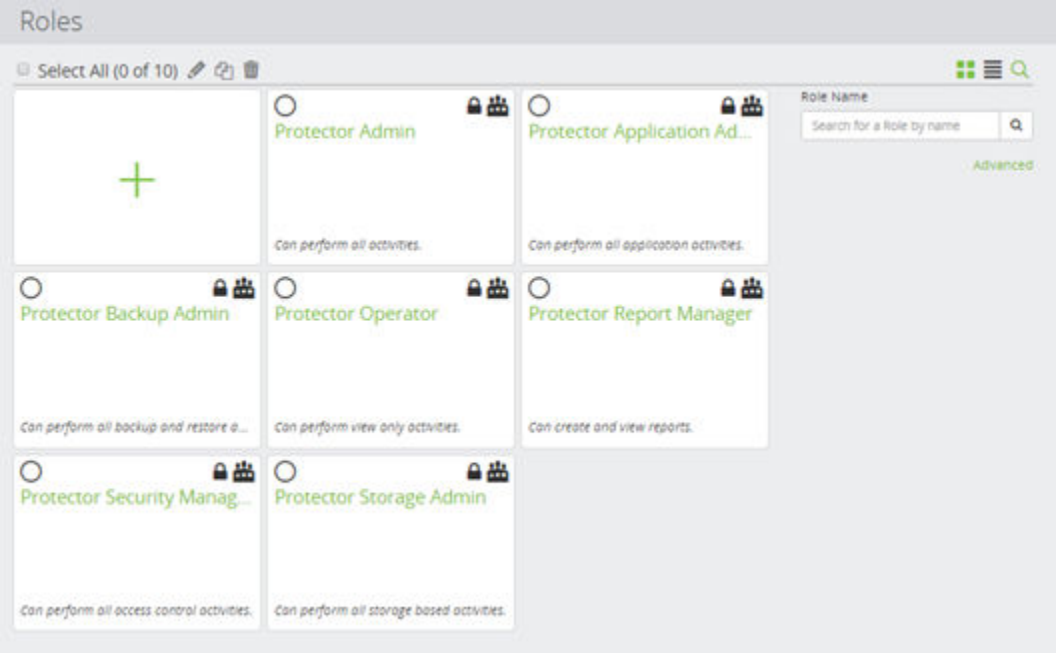


Figure 110 Roles Inventory

Control	Description
Summary	Select the Summary option from the drop down menu in the Navigation Breadcrumbs (on page 290) to view the Access Control Summary (on page 348).

Control	Description
 Edit	Edits an existing Role in the inventory. The Access Control Role Wizard (on page 337) is launched to enable the Role's attributes to be changed.
 Clone	Enabled only when one Role is selected. Creates a clone of the selected item which can then be modified.
 Delete	Enabled only when one or more Roles are selected. Deletes the selected item from the inventory.
 Add	Creates a new Role. The Access Control Role Wizard (on page 337) is launched to guide you through the process.
 System generated Roles	At least three system generated Roles are available when the product is installed. They define default administrator, security manager and user roles. These roles cannot be deleted since they provide a basic level of access control. System generated roles are marked with a  icon to indicate that they cannot be modified. The Access Control Role Details (on page 340) is displayed to enable the permissions to be viewed.
 User defined Role(s)	Any number of user defined Roles can be created. These are displayed in the inventory. Roles should be defined in order to grant the required level of functionality appropriate to the users' responsibilities. The Access Control Role Details (on page 340) is displayed to enable the permissions to be viewed and edited.
Filter on Role Name	Filters the displayed results based on the Role Name.

Access Control Role Wizard

This wizard is launched when a new Role is added to the Roles Inventory.

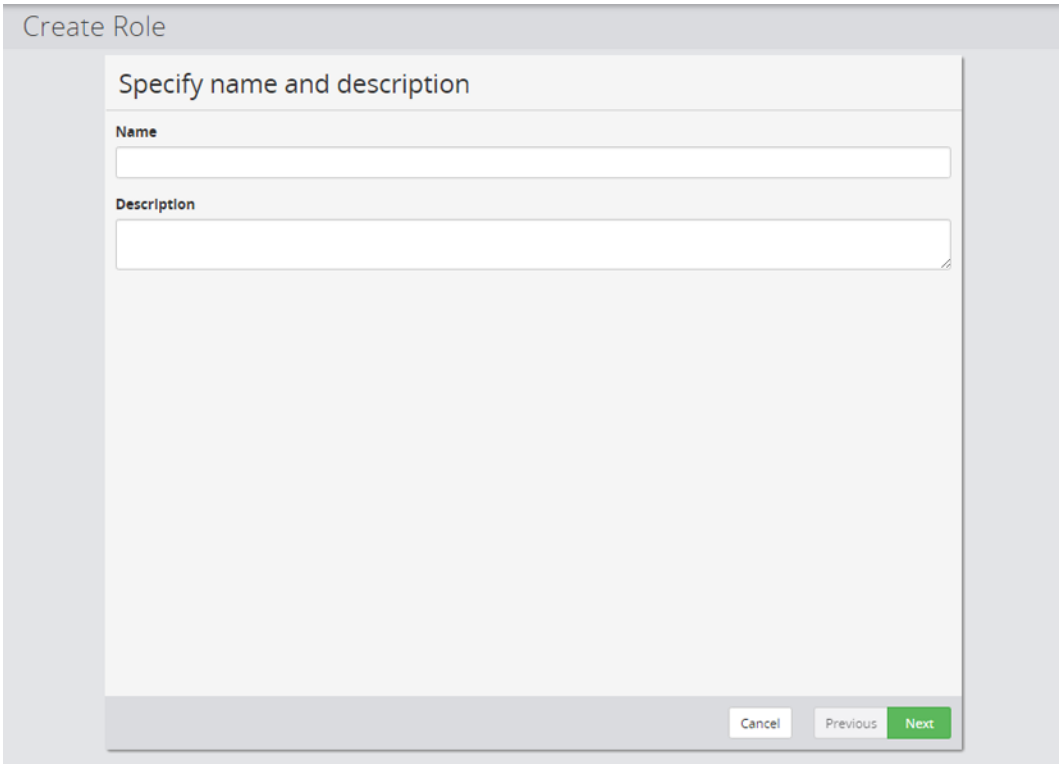


Figure 111 Role Wizard - Specify name and description

Control	Description
Name	The name of the Role.
Description	Optional. A short description of the Role.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

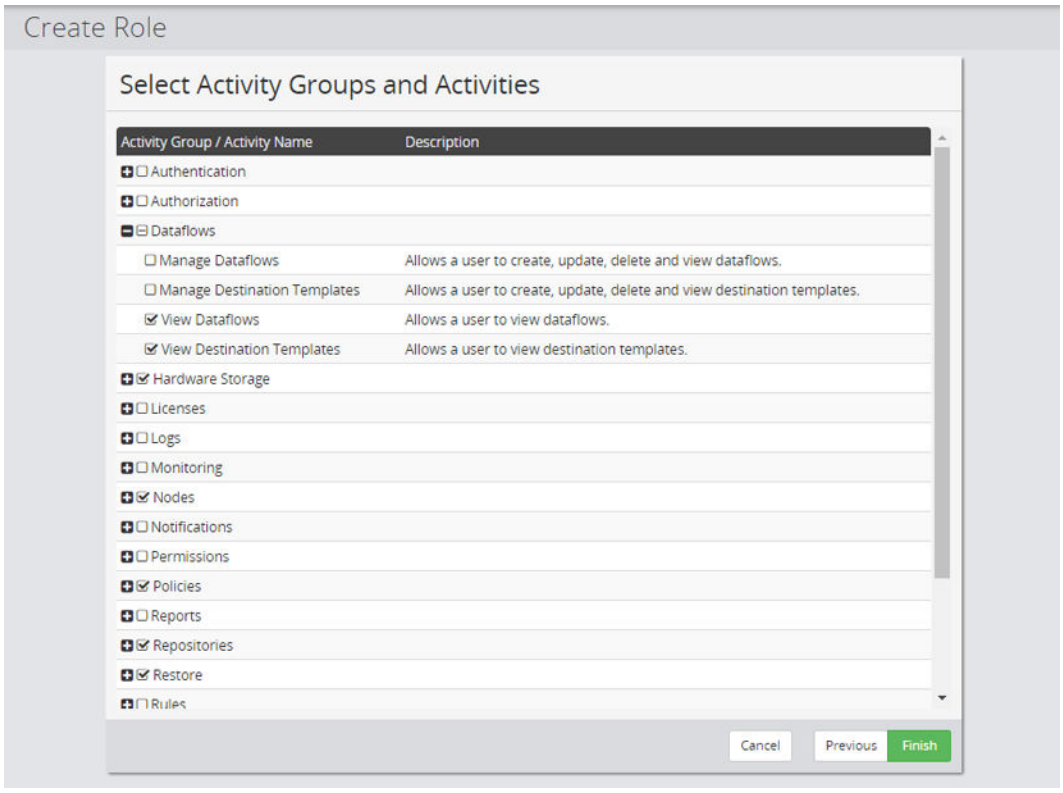


Figure 112 Role Wizard - Select Activity Groups and Activities

Control	Description
Activity Groups	Activity Groups contain a set of functionally cohesive Activities that are typically applied to a Role en-mass. Click the + button to the left of an Activity Group to expand and view the activities within a group. Click the checkbox to the left of an Activity Group to apply or remove all Activities within that group for the Role. Click the checkbox to the left of an Activity to apply or remove that Activity for the Role. The check box to the left of an Activity Group displays a '-' instead of a tick if only some of the Activities in the group have been applied. Refer to Controlling access to UI features with Activities and Activity Groups (on page 341) for details on how each activity affects access to the UI.

Control	Description
	Note: The *Override Ownership Permissions* activity within the *Permissions* activity group allows users with this activity to view Policies, Dataflows, Destination Templates and Schedules regardless of who created them or who they are assigned to. Enable this permission with care.
Activities	Activities define what a user can do within Protector via the UI and via the REST API. Click the checkbox to the left of an Activity to apply or remove the activity for the Role.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Finish	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

Access Control Role Details

This page displays the details of a Role and enables you to launch the wizard to edit them.

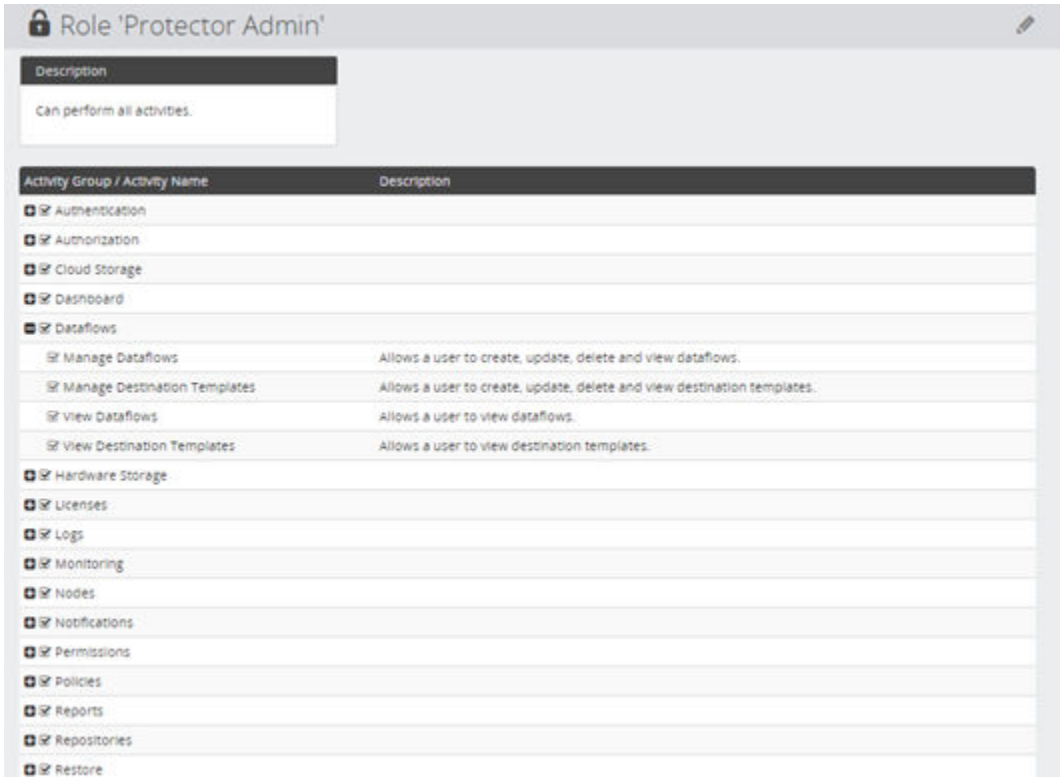


Figure 113 Role Details

Control	Description
 Edit	Launches the Access Control Role Wizard (on page 337) to enable the role to be edited.

Controlling access to UI features with Activities and Activity Groups

Access to Protector features is controlled by the *Activity Groups* and *Activities* assigned to the user. Access to a feature can be:

- Denied completely by disabling all activities
- Set to read-only by enabling only *View* activities (e.g. *View Policies*)
- Set to full control by enabling *Manage* activities (e.g. *Manage Policies*)
- Set to enable specific functionality within some features (e.g. *Trigger Operations* enables policies to be triggered from the **Monitor** page)

The following table lists the features available in Protector, along with the *Activity Groups* that control general access to those features. Fine grain access control to specific functions within a feature can be achieved by enabling or disabling specific *Activities*.

User Interface Page	Activity Group	Activity
Pages accessible from the Main Banner:		
Dashboard	Refer to Logs, Jobs, Nodes, Monitor, Policies, Data Flows and Licenses below.	
Jobs	Monitoring	View Jobs
		Manage Jobs
Logs	Logs	View Logs
		Manage Logs
		Purge Audit Logs
		Verify Audit Logs
	Notifications	View Log Notifications
		Manage Log Notifications
Monitor	Monitoring	View Node Statistics
	Triggering	Trigger Operations

User Interface Page	Activity Group	Activity
Storage	Hardware Storage	Manage Hardware. **Note:** Enabling this option will automatically enable all other activities in this activity group regardless of their current state.
		Manage Hardware Snapshots and Clones
		Mount Hardware Snapshots and Clones
		Pause Hardware Replications
		Revert Hardware Snapshots
		Swap Hardware Replications
		View Hardware
	Repositories	View Repositories
		Use Repositories
		Manage Repositories
Reports	Reports	View Reports
		Manage Reports
Pages accessible from the Sidebar:		
Nodes	Nodes	View Source Nodes
		Manage Source Nodes
		View Storage Nodes
		Manage Storage Nodes
	Software Updates	View Software Updates
		Manage Software Updates
Node Groups	Nodes	View Source Nodes
		Manage Source Nodes
		View Storage Nodes
		Manage Storage Nodes
Policies	Policies	View Policies

User Interface Page	Activity Group	Activity
		Manage Policies
Data Flows	Dataflows	View Dataflows
		Manage Dataflows
		View Destination Templates
		Manage Destination Templates
	Rules	Manage Rules
Schedules	Schedules	View Schedules
		Manage Schedules
Notifications	Notifications	View Notification Settings
		Manage Notification Settings
Restore	Restore	Perform Restores
Access Control	Authentication	View Authentication Spaces
		Manage Authentication Spaces
	Authorization	View RBAC Configuration
		Manage RBAC Configuration
Licences	Licences	View Licenses
		Manage Licenses
<Item> Permissions	Permissions	Override Ownership Permissions

Note:

If a user does not have the required *Activity Group* or *Activity* assigned to them via a *Role* then the user interface will prevent the user performing the activity or viewing information in one or more of the following ways:

- Suppressing display of the associated controls.
- Overlaying a warning triangle icon within the associated controls.
- Displaying an `Access Denied` hover hint when the user moves the cursor over the associated controls.
- Displaying an `Access Denied` message where the associated information would normally appear on a page, wizard or dialog.
- Displaying an `Access Denied` pop-up *Session Notification* when the request is denied by the back-end handler code.

Access Control Resource Groups Inventory

This inventory lists all defined Resource Groups. Resource Group are created to define logical groups of computing resources in the context of access control. They are distinct from Node Groups which are created to help define policies.

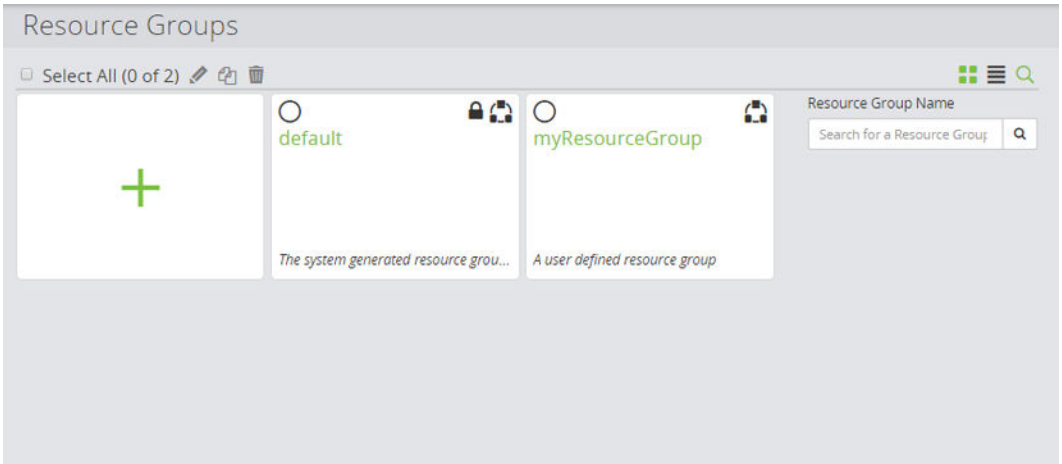


Figure 114 Resource Groups Inventory

Controls	Description
Summary	Select the Summary option from the drop down menu in the Navigation Breadcrumbs (on page 290) to view the Access Control Summary (on page 348).
Edit	Edits an existing Resource Group in the inventory. The Access Control Resource Group Wizard (on page 345) is launched to enable the Resource Group's attributes to be changed.
Clone	Enabled only when one Resource Group is selected. Creates a clone of the selected item which can then be modified.
Delete	Enabled only when one or more Resource Groups is selected. Deletes the selected item from the inventory. The associated Nodes are not deleted.
Add	Creates a new Resource Group. The Access Control Resource Group Wizard (on page 345) is launched to guide you through the process.
System generated Resource Group	One system generated Resource Group is available when the product is installed. All nodes that are listed in the Nodes Inventory (on page 536) are automatically added to this default resource group. This resource cannot be deleted since it provides a basic level of access control. System generated Resource Groups are marked with a icon to indicate that they cannot be modified. The Access Control Resource Group Details (on page 347) is displayed to enable the permissions to be viewed.

Controls	Description
User defined Resource Group(s)	Any number of user defined Resource Groups can be created. These are displayed in the inventory. Resource Groups should be defined in order to restrict access to nodes. The Access Control Resource Group Details (on page 347) is displayed to enable the permissions to be viewed and edited.
Filter on Resource Group Name	Filters the displayed results based on the name.

Access Control Resource Group Wizard

This wizard is launched when a new Resource Group is added to the Resource Groups Inventory.

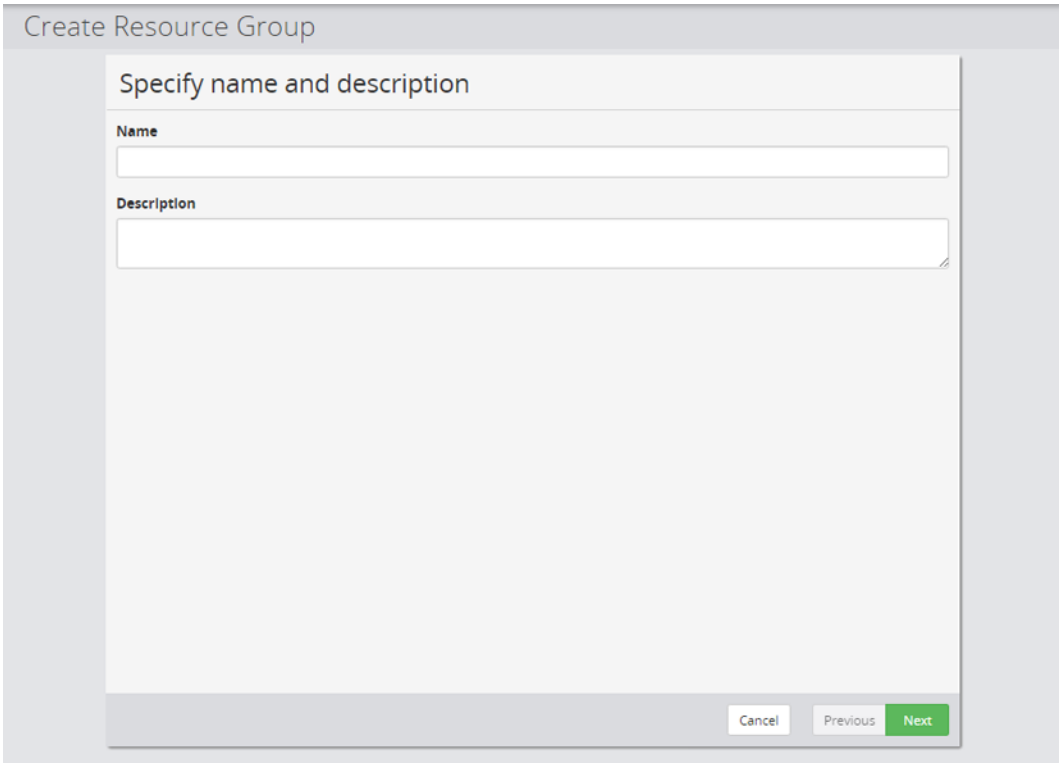


Figure 115 Resource Group Wizard - Specify name and description

Control	Description
Name	Enter the name of the resource group.
Description	Optional. Enter a short description of the resource group.
Cancel	Discards all changes and reverts to the previous page.

Control	Description
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

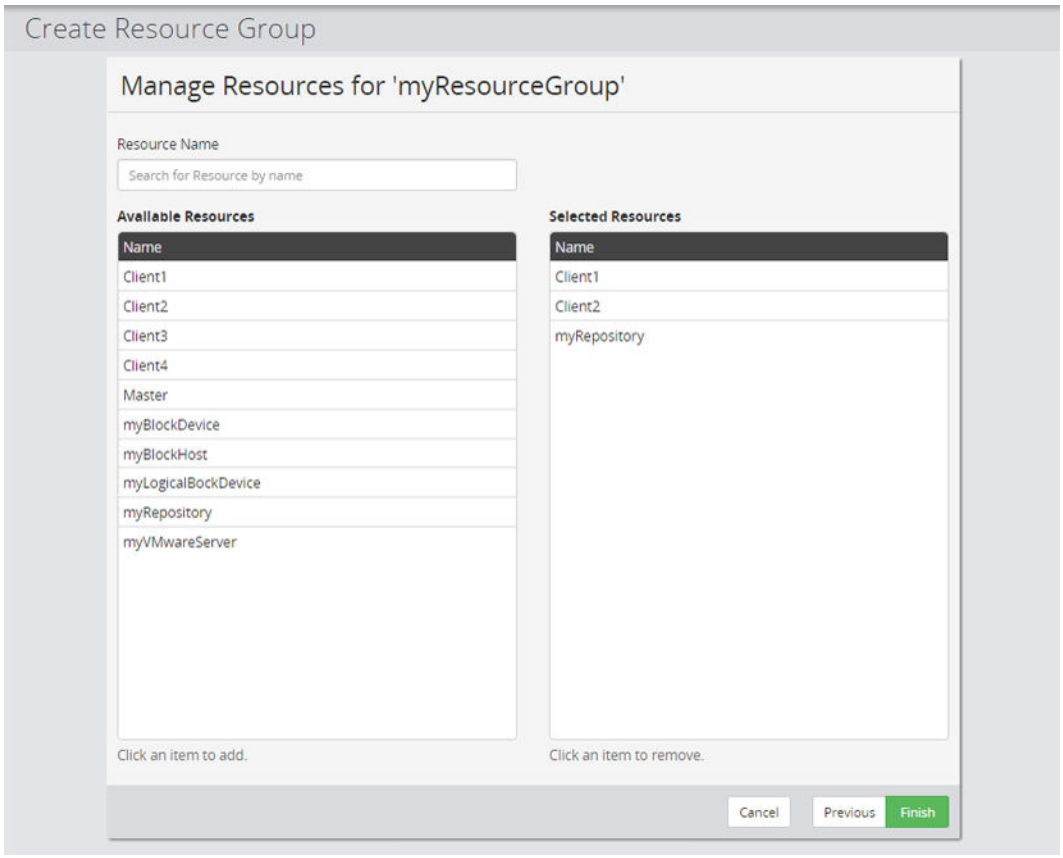


Figure 116 Resource Group Wizard - Manage Resources

Control	Description
Resource Name	Searches for the resource by name.
Available Resources	Lists the available resources. Click on the name of the resource to add to the selected resource list. **Note:** If the Master node is included in a resource group then users with access to that group will also have access to administrative log messages. Access to the Master node should only be granted to administrative users.
Selected Resources	Lists the selected resources. Click on the name of the resource to remove it from the selected resource list.

Control	Description
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Finish	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

Access Control Resource Group Details

This page displays the details of a Resource Group and enables you launch the wizard to edit them.

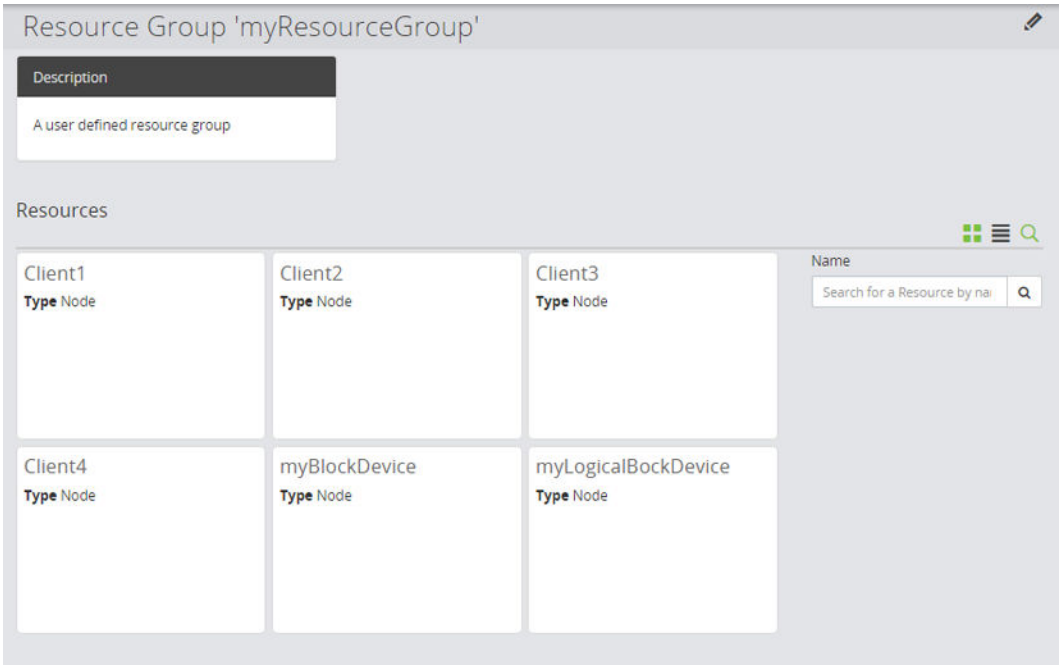


Figure 117 Resource Groups Details

Control	Description
Edit	Launches the Access Control Resource Group Wizard (on page 345) to enable the group to be edited.
Filter on Resource Name	Filters the displayed results based on name.

Access Control Summary

This page displays Access Control settings for each user or group configured within Protector.

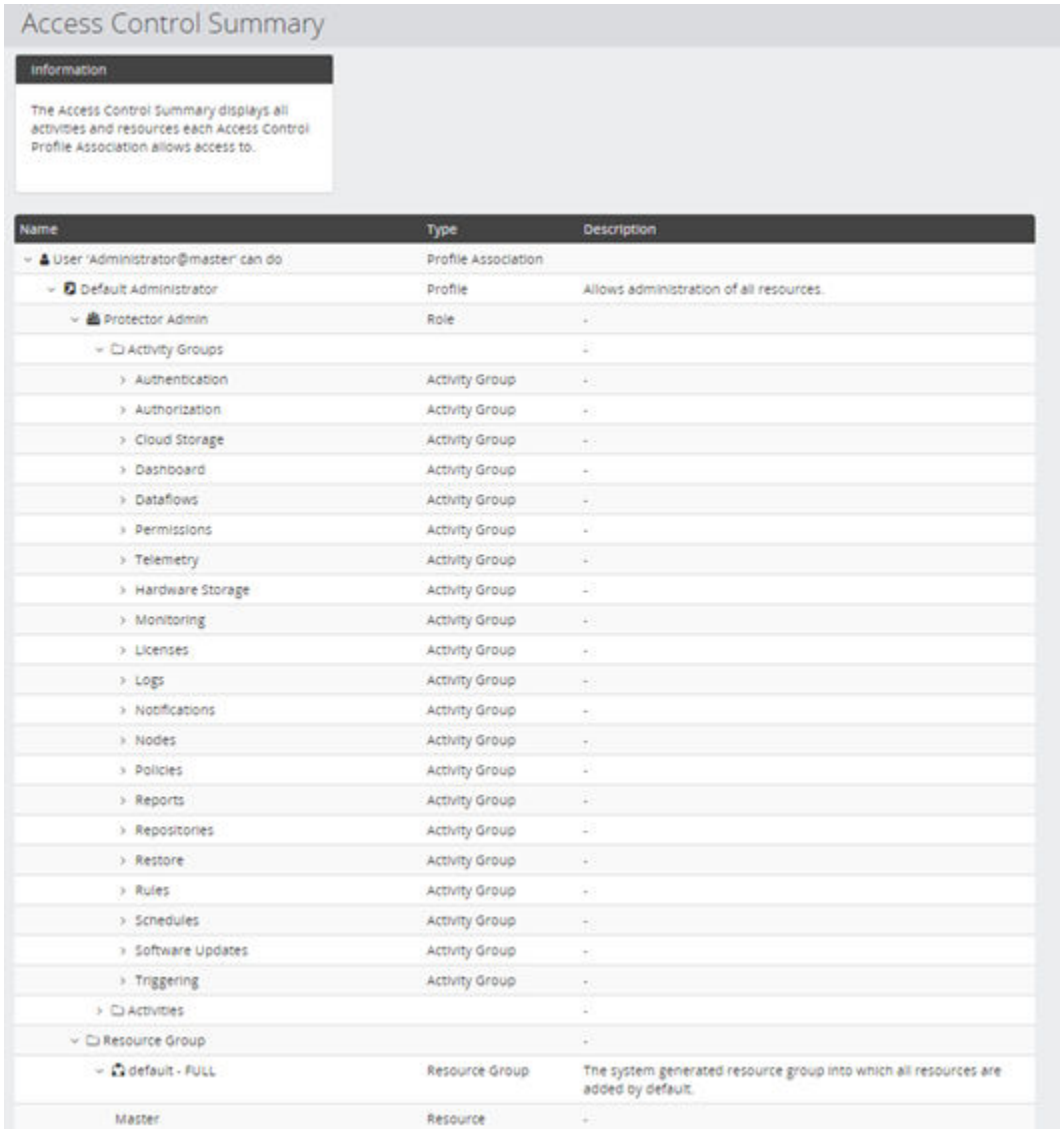


Figure 118 Access Control Summary

Tree Node	Description
ACP Association Name	Each ACP Association defined within Protector is listed by name. Click [+] to view the related ACPs.
ACP Name	Shows the related ACPs below the ACP Association. Click [+] to view the related Role and Resource Groups.
Role Name	Shows the related Role below the ACP. Click [+] to view the related Activity Groups and Activities

Tree Node	Description
Activity Group Name	Shows the related Activity Groups below the Role. Click [+] to view the related Activities.
Activity Name	Shows the related Activities below the Activity Group.
Resource Group Name	Shows the related Resource Groups below the ACP. Click [+] to view the related Resources.
Resource	Shows the related Resources below the Resource Group.

Access Control Permissions Inventory

The Permissions Inventory is accessed via the View Permissions button on various items within the Web UI including policies, data flows and schedules. It allows to view the access permissions for those items granted to specific users and groups.

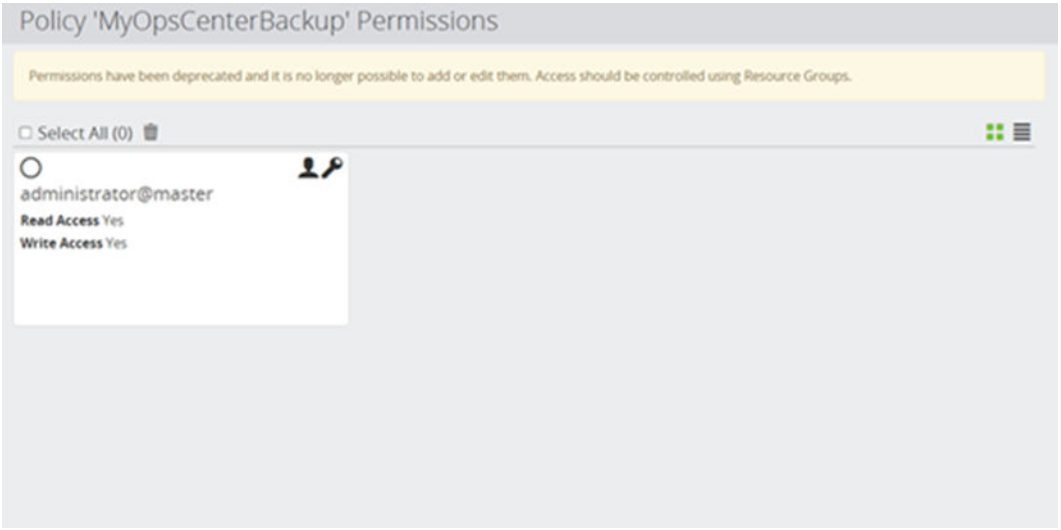


Figure 119 Permissions Inventory

Control	Description
Delete	Enabled only when one or more permissions is selected. Deletes the selected item from the inventory.
Default Permission	By default the system adds the *administrator@master* user permission to controlled items and grants READ/WRITE access. The default permission provides a basic level of access control.

Common Controls UI Reference

This section describes common controls found throughout the UI.

Path Dialog

This dialog is displayed when a path is selected using the Browse button to select a path.

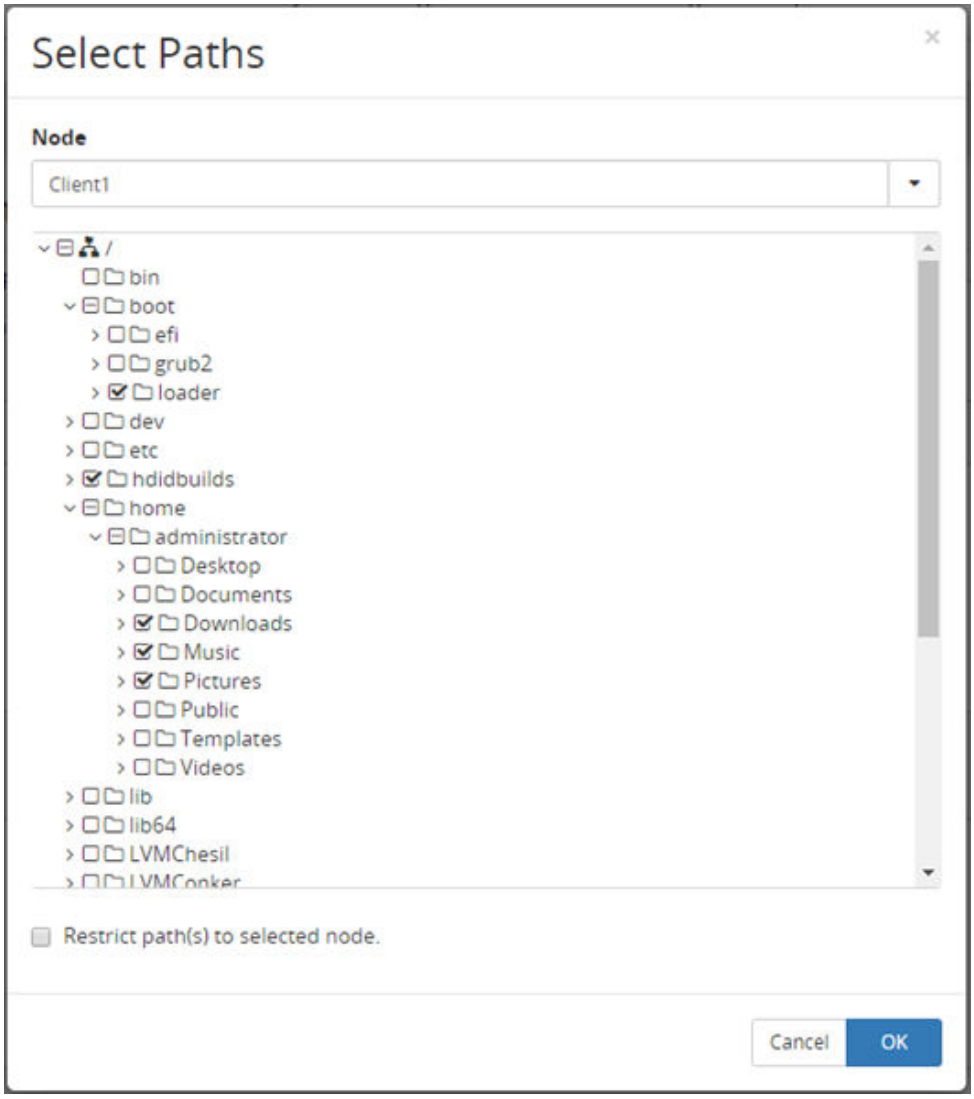


Figure 120 Path Dialog

Control	Description
Node	Select the node where the path resides. This control may be disabled if the node is already known or cannot be changed.

Control	Description
Directory Tree	Expand the directory tree and select the required folder to specify the path. **Note:** If the directory does not exist, it can be created by closing this dialog and typing the required path into the edit box adjacent to the Browse button.
Restrict path(s) to selected node.	Appears only when defining a *Path Classification* for a *Policy*. The path will be used only in conjunction with the specific Node specified above, rather than being applied to any node.

Date Time Picker

The date time picker control allows a calendar date and time to be selected.

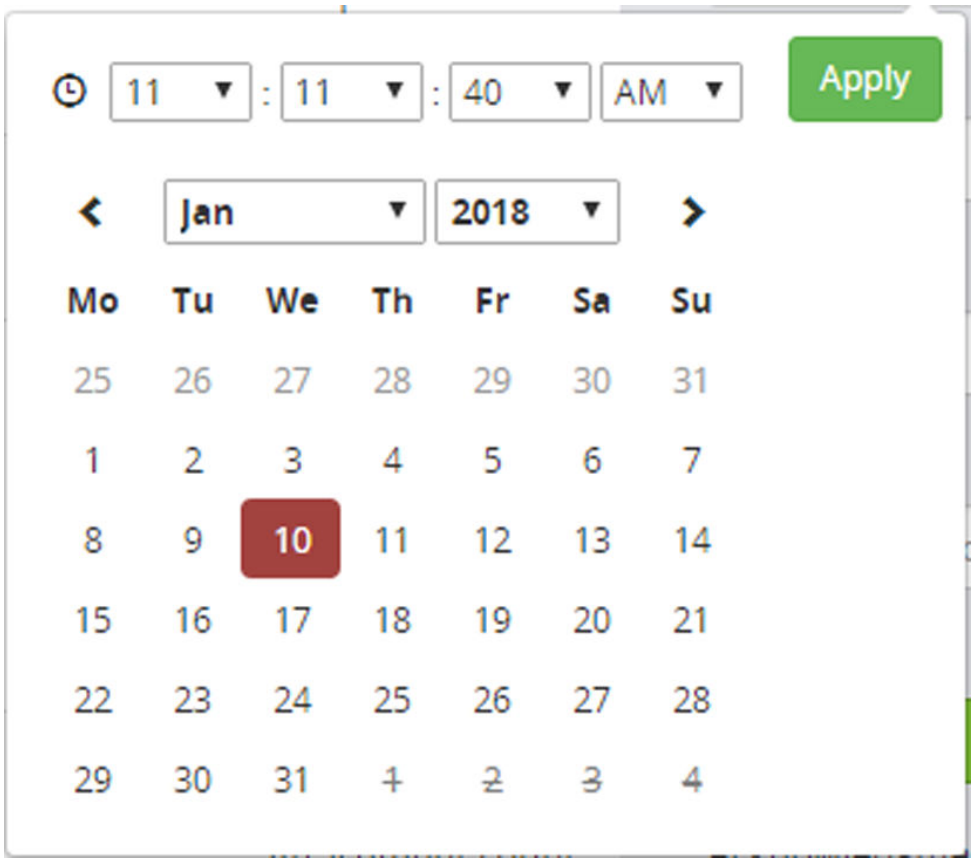


Figure 121 Date Time Picker

Control	Description
Hour	Select the hour from the dropdown list.

Control	Description
Minute	Select the minute from the dropdown list.
Second	Select the second from the dropdown list.
AM/PM	Select AM or PM (not available if Use 24 hour format is selected in the Settings Wizard (on page 287)).
<	View previous month.
>	View next month.
Month	Select the month from the dropdown list.
Year	Select the year from the dropdown list.
Day	Select the day in the calendar (unavailable days are struck through).

Note: It is necessary to select a Day when setting the Date/Time, if a day is not selected the selection will not take effect.

Date Time Range Picker

The date time range picker allows a calendar date and time range to be selected.

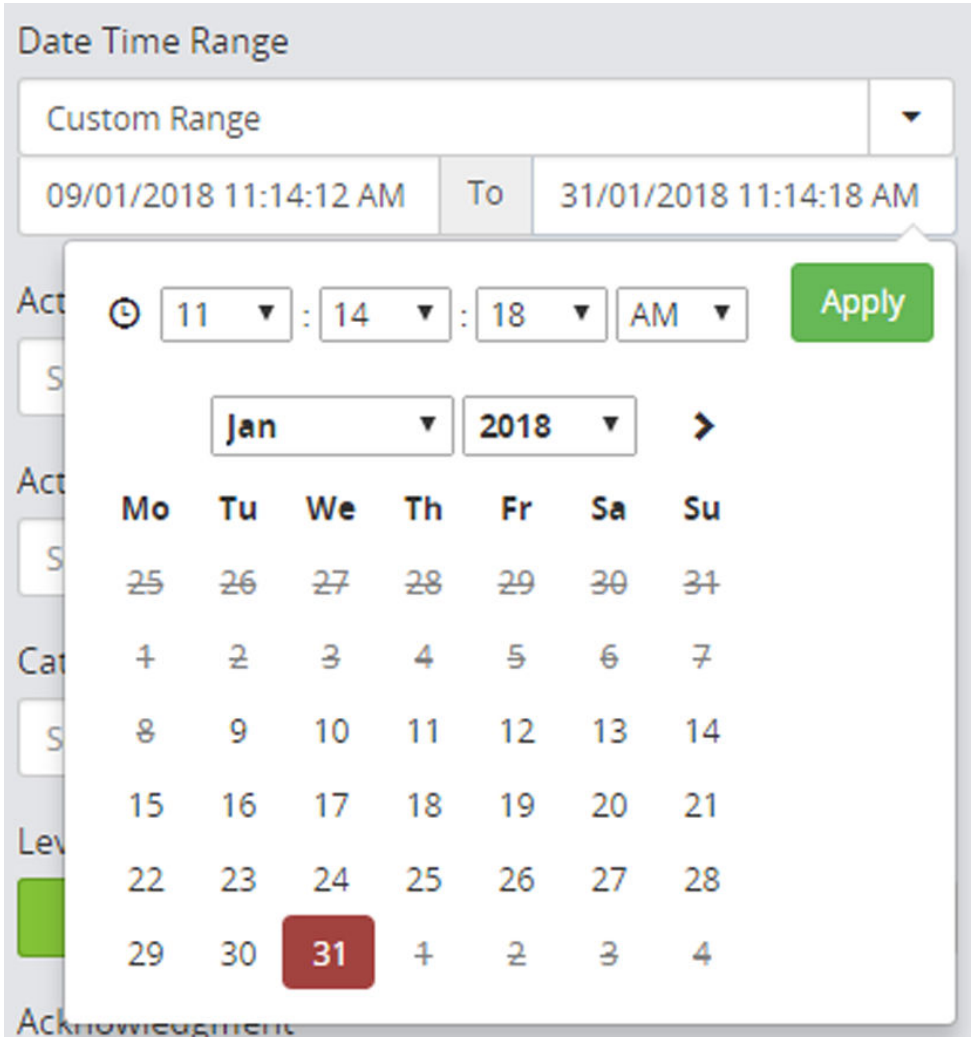


Figure 122 Date/Time Range Picker

Control	Description
Presets	Select one of the following predefined ranges from the top dropdown list: ▪ Today ▪ Yesterday ▪ Last 7 Days ▪ Last 14 Days ▪ This Month ▪ Last Month ▪ Custom Range - enables the calendar controls described below.

Control	Description
From	If Custom Range is selected, click the lower left field to open a Date Time Picker (on page 351) to select the 'From' date time.
To	If Custom Range is selected, click the lower right field to open a Date Time Picker (on page 351) to select the 'To' date time.

Hitachi Block Host Resize Dialog

The Logical Device resize dialog can be accessed from the Block Host screen using the ⤢ tool bar item. This dialog can be used to resize logical devices represented by the Block Host node that are part of a replication pair. For active live replications both the PVOLS and SVOLS are expanded to the same size. For paused live replications and batch replications only the PVOLS are expanded to the requested size the SVOLS will be expanded when the replication is resumed or resynchronized, see Hitachi Block Device Advanced Settings Dialog (on page 893) for the array settings required to enable this functionality.

Note: Expansion is not supported for:

- TI replications.
- Replications with non-floating snapshots.

Note: To expand replications which have a live cascaded replications the cascaded replications must paused prior to expanding the initial replication.

Note: To expand a block host which has multiple live replications all except one of the replications must be paused prior to expanding the block host.

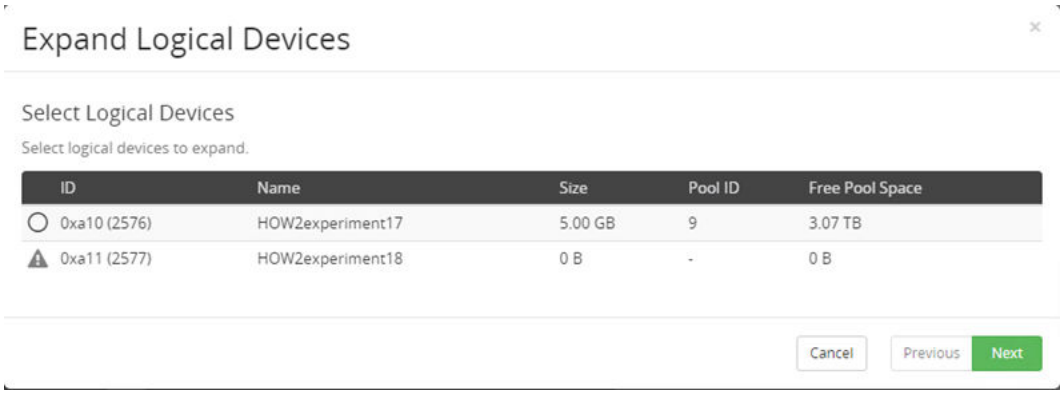


Figure 123 Expand Logical Devices - Select Logical Devices

First, select all logical devices that require expansion and click Next

Note: Any entries in the table with the ⚠ are LDEVs which are not available for expansion. Hovering the mouse pointer over the warning triangle will give more information as to why the LDEV can not be expanded.

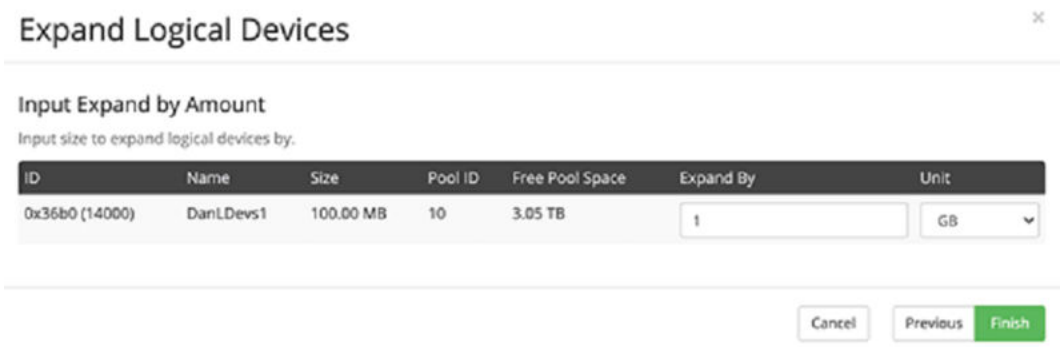


Figure 124 Expand Logical Devices - Input Expand by Amount

Then for each selected logical device specify how much to expand the volume by in units of GB, TB or Blocks.

Note: During the expansion of a live GAD replication, the expansion will be attempted while the replication is active. If this fails, or for any other live replication, the replication will be paused, then the PVOL(s) and SVOL(s) will be expanded and finally the replication will be resumed.

However,for swapped replications the expansion is done in the reverse order expanding the active SVOL first and then expanding the active

Data Flows UI Reference

This section describes the Data Flows UI, accessed via the Navigation Sidebar (on page 291).

For further information, refer to:

- Data Flow Concepts (on page 50)
- Data Flow Tasks (on page 228)

Data Flows Inventory

This inventory lists all defined Data Flows whether they are active, inactive or under construction.

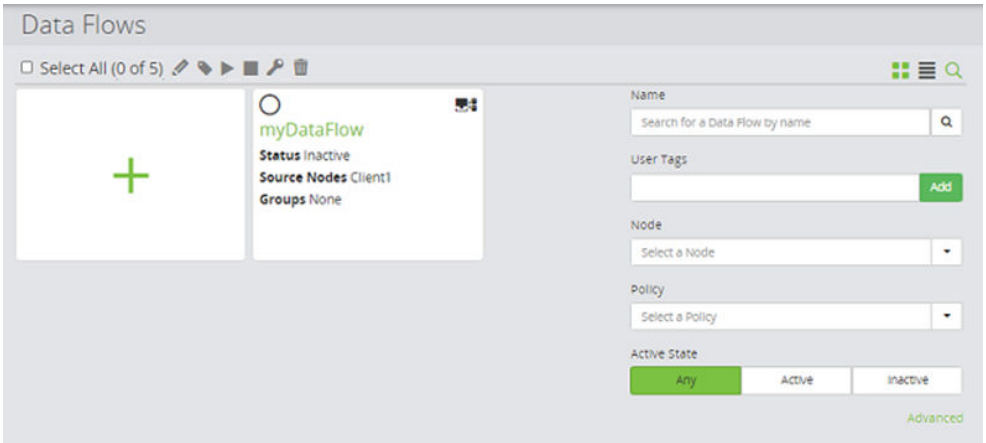


Figure 125 Data Flows Inventory

Control	Description
Edit	Edits an existing data flow in the inventory. The Data Flow Wizard (on page 361) is launched to enable the data flow's attributes to be changed. **Note:** If an active data flow has been modified, but has not been reactivated, a warning triangle will be displayed in the *Status* field on the corresponding tile.
Tag	Modifies the tags of an existing object from the either the inventory screen or the details screen of the object.
Activate	Enabled only when one or more data flows are selected. Displays the Activate Data Flow Dialog (on page 357) and attempts to compile the rules for the selected data flows. If compilation is successful then the rules can be activated. **Caution:** Activate data flows in batches not exceeding 20 data flows at a time. Activating more than this simultaneously can result in longer activation times.
Deactivate	Enabled only when one or more data flows are selected. Deactivates the selected data flows. **Caution:** If the deactivated data flows contain storage hardware based operations, replications will be placed in the eligible for tear down state.
Delete	Enabled only when one or more data flows are selected. Deletes the data flow from the inventory. Active data flows cannot be deleted.

Control	Description
 Add	Creates a new Data Flow. Launches the Data Flow Wizard (on page 361) to guide you through the process.
 Existing Data Flow	Click the data flow name to open the Data Flow Details (on page 491) to enable you to view and edit the data flow.
Filter on Name	Filters the displayed results based on Data Flow Name.
Filter on User Tags	Filters the displayed results based on Tags contained in the data flow.
Filter on Node	Filters the displayed results based on Node Name contained in the data flow.
Filter on Policy	Filters the displayed results based on Policy Name contained in the data flow.
Filter on Active State	Filters the displayed results based on the active state of the data flow.

Activate Data Flow Dialog

This dialog is displayed when one or more data flows are activated.

Rules files control the operation of the Ops Center Protector components on each node in the data flow definition. Rules files are generated by the *Rules Compiler* after policies have been created, the data flows have been defined and policies assigned to source and destination nodes.



Caution: If you modify a policy or data flow that is currently active, then the data flow must be reactivated before your changes will take effect.

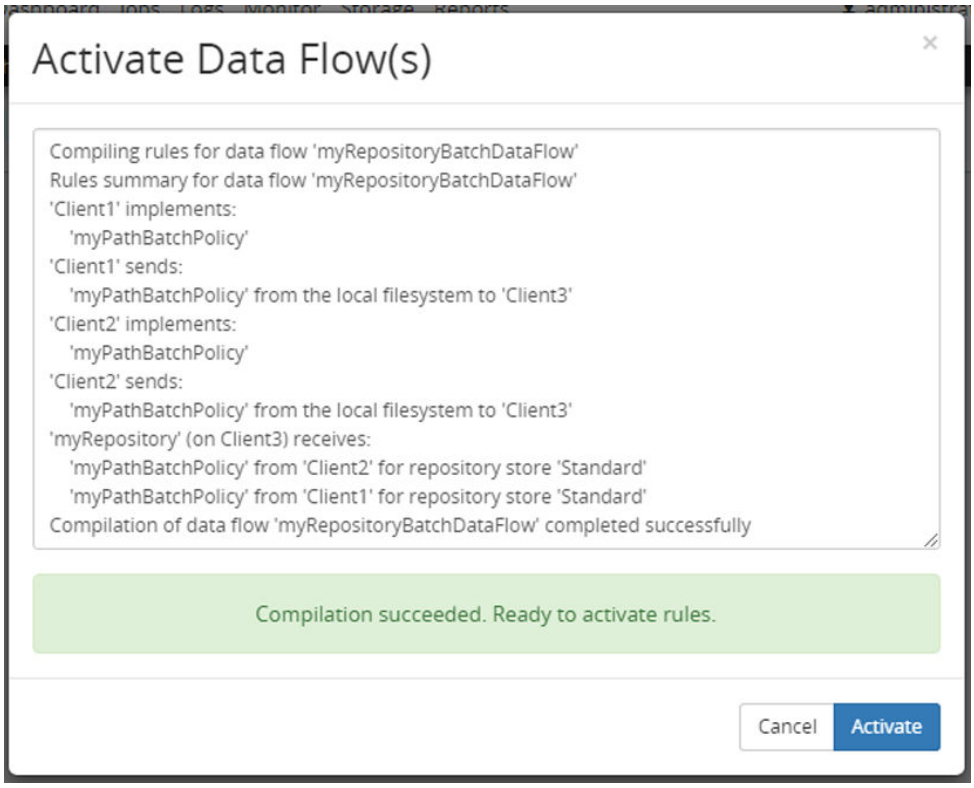


Figure 126 Activate Data Flow(s) Dialog (Compilation succeeded)

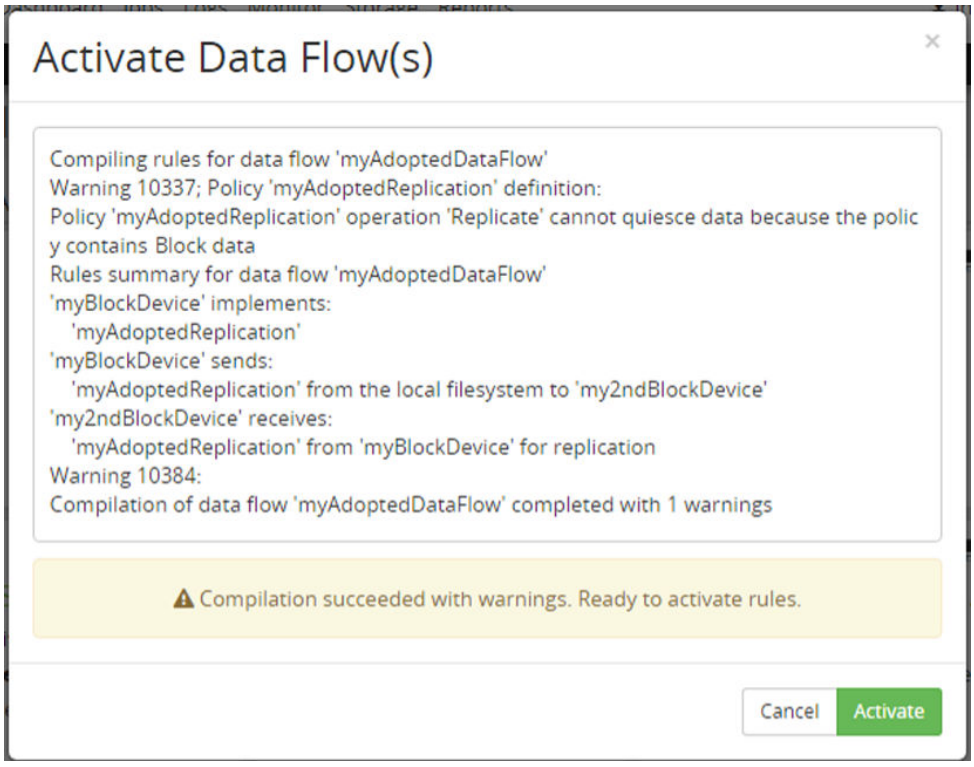


Figure 127 Activate Data Flow(s) Dialog (Compilation succeeded with warnings)

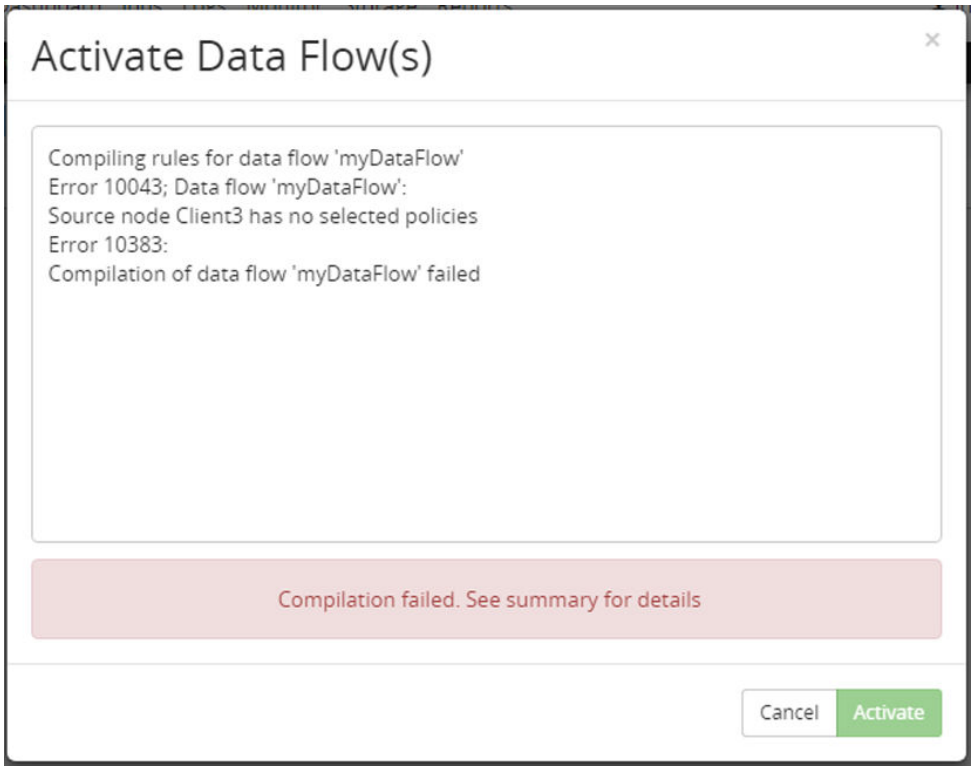


Figure 128 Activate Data Flow(s) Dialog (Compilation failed)

Control	Description
Compilation Details	This lists the output from the rules compiler and shows error, warning and information messages. If the compilation is successful, it includes a summary of what policies each node in the data flow will be implementing, forwarding or processing. The data flow can be activated. If the compilation is successful but contains warnings, then the information contained here will assist you in resolving potential issues with the data flow and/or policies. The data flow can be activated. If the rules fail to compile then the information contained here will assist you in resolving issues with the data flow and/or policies. You may need to iterate through several compilation cycles to remove all compilation errors. The data flow cannot be activated. **Caution:** Always inspect the compiler output. A successful compilation may contain warning messages which you should review. Successful compilation only indicates that the rules are valid, it does not imply that they are verified against your data protection requirements. Please regularly inspect the Default Dashboard (on page 292) and related status screens to ensure that your data is being protected in the manner you intended.

Control	Description
	Compiler errors are commonly caused by: ▪ incorrect or incomplete policies ▪ incomplete data flow definitions ▪ incorrect or incomplete data flow item attributes ▪ incomplete or ambiguous routes for policies ▪ missing parameters on destination nodes ▪ adding, editing or deleting policy classifications or operations without updating data flows
Compilation Outcome	Indicates one of the following: ▪ Blue - the rules are being compiled. Please wait. ▪ Green - the rules have been compiled successfully. The data flows and policies are valid (but not necessarily correct) and can be activated. ▪ Amber - the rules have been compiled successfully with warnings. The data flows and policies are valid (but not necessarily correct) and can be activated. The data flows and/or policies contain warnings that may need to be rectified. ▪ Red - the rules have not been compiled. The data flows and or policies contain errors that must be rectified before the data flows can be activated.
Activate	This button is only enabled once the rules have compiled successfully. Click to distribute the rules to the affected nodes and activate them. The policies will then take effect depending on the availability of affected nodes, operations, trigger conditions or schedules defined. The Logs Inventory (on page 509) can be used to watch and review the progress of data flow activation.

Control	Description
	Note: Rules activation methods depend on the type of policy: ▪ Host based: An initial 2 minute rules settlement period is applied after activation to allow rules to reach all participating nodes. If any operation is triggered within the settlement period it will be deferred until after it has expired. ▪ Block based: If the current rules are null or they predate the new rules, the ISM will wait up to 2 minutes for a new rule set to arrive. An information level log will be generated if this delay is invoked: `Storage proxy node does not have required rules, waiting up to <DELAY> seconds.` If, after waiting, the rules are still out of date, the ISM will attempt to use the rules it currently has. ▪ All: If a node in the data flow definition is unavailable, then the rules files cannot be sent to that node; activation continues for the available nodes. When the activation is complete, the new rules are activated on all available nodes. When the absent node becomes available and reconnects to the master, the new rules are sent automatically to that node and activated immediately after they are received. If any nodes are deleted from a previously activated data flow definition, then these nodes are deactivated when the new rules are activated. If any such node is not currently available, it is deactivated as soon as it reconnects to the master.

Data Flow Wizard

This wizard is displayed when a new data flow is being created.

The Data Flow Wizard performs two principal functions:

- Defines the routing of policies that move data from source nodes to destination nodes.
- Assign policies to nodes on your network, defining what data to backup and the methods of protection to employ.

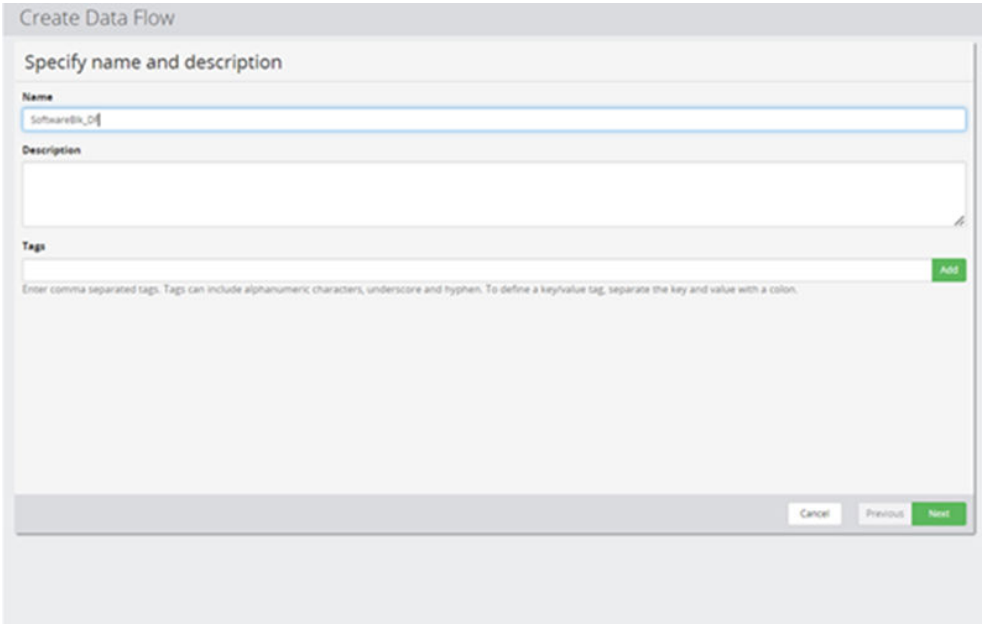


Figure 129 Data Flow Wizard - Name

Control	Description
Name	Enter a name for the Data Flow.
Description	Optional. Enter a short description of the Data Flow.
Tags	Add the tags to be associated with the object being created.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

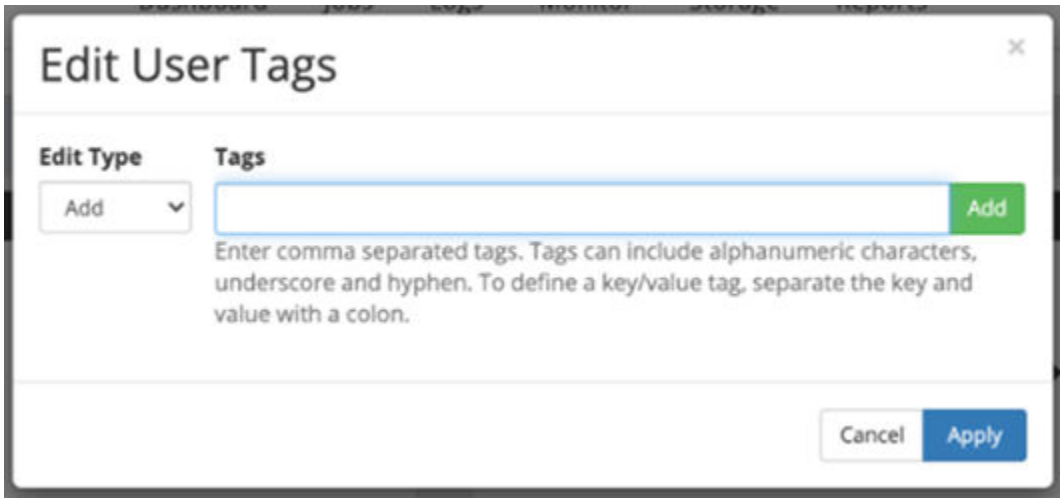


Figure 130 Add Tags - Edit User Tags

Control	Description
Edit Type	Enter the Edit Type.
Tags	Add the tags to be associated with the object being created.
Cancel	Discards all changes and reverts to the previous page.
Apply	Commits the new changes. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

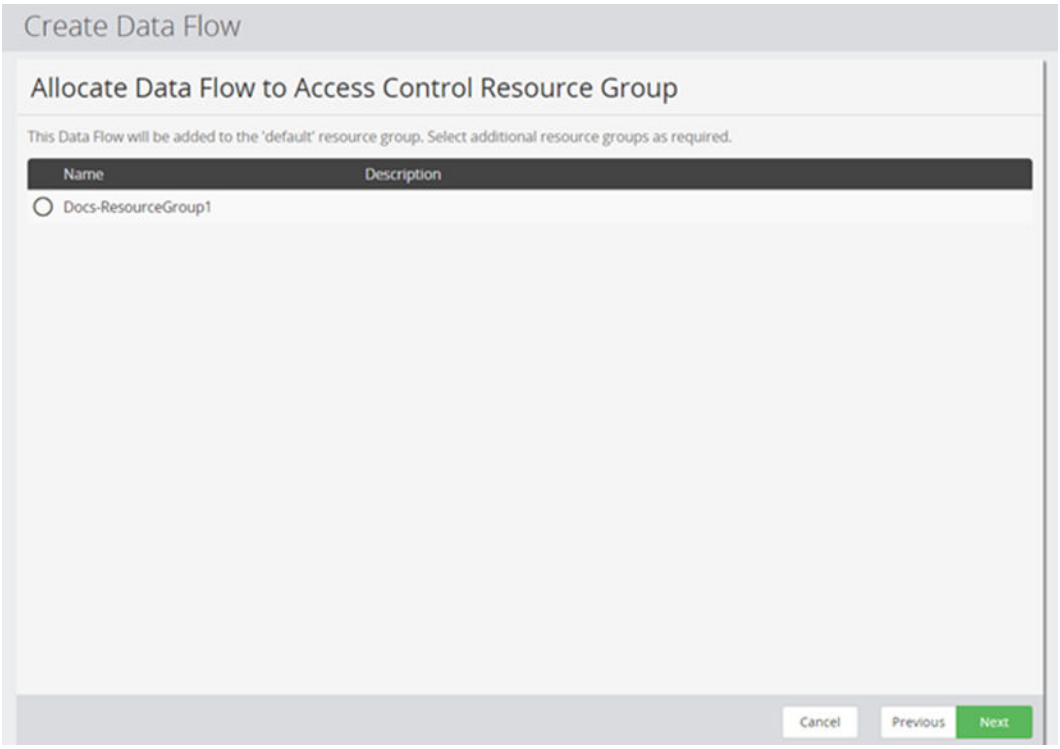


Figure 131 Dataflow Wizard - Allocate Dataflow to Resource Group

Control	Description
Resource Groups	It allows the user to view the access permissions for those items granted to specific users and groups. **Note:** A single Data Flow can be assigned to multiple resource groups.

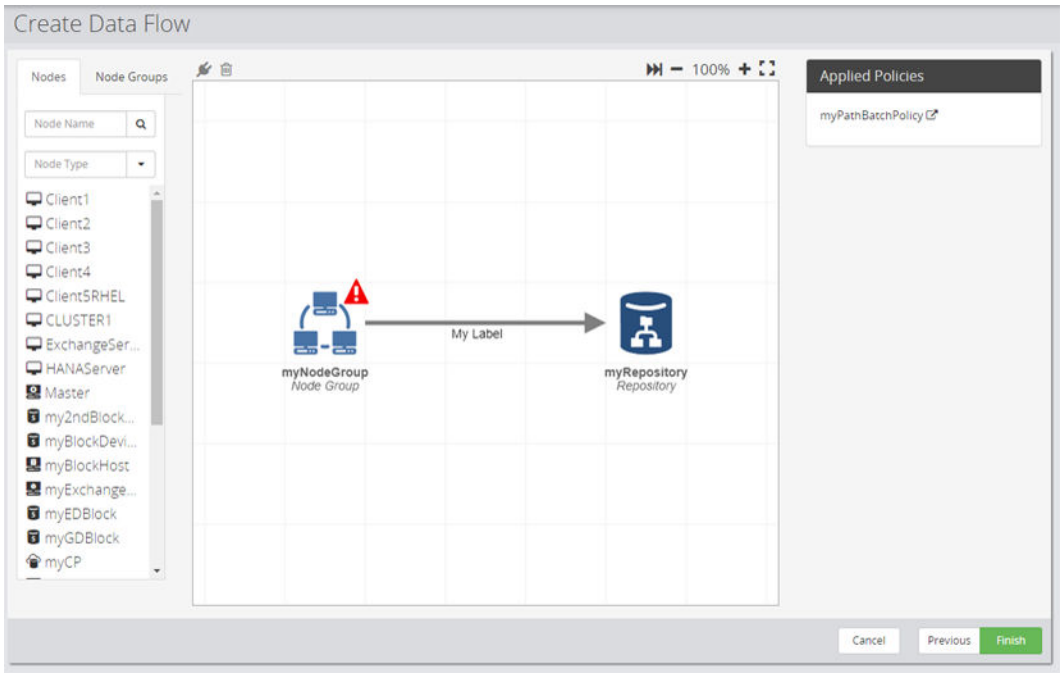


Figure 132 Data Flow Wizard

Table 21 Nodes Tab

Control	Description
Filter on Node Name	Filters the displayed results based on Node Name.
Filter on Node Type	Filters the displayed results based on Node Type.
Nodes List	List the available source and destination nodes that can be dragged onto the Data Flow Workspace.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Finish	Commits the new changes. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

Table 22 Node Groups Tab

Control	Description
Filter on Group Name	Filters the displayed results based on Node Group Name.
Node Group List	List the available node groups that can be dragged onto the Data Flow Workspace.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Finish	Commits the new changes. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

Table 23 Data Flow Workspace and Applied Policies

Control	Description
Connect To	Enabled only when a single node is selected on the diagram. Click to create a connection from the selected node to another node on the diagram. When the connector is displayed, click again on the destination node to complete the connection. By default the Mover Type is set to Batch; the properties associated with each connection can be edited via the Mover Settings controls (see 'Mover Settings' controls below). Only compatible elements can be connected. **Tip:** Connections can also be created by dragging a destination node from the node list, over a source node to pick up the connector and then dropping it at the desired location. Alternatively you can drag a source node already on the diagram, over a destination node on the diagram and drop it there.
Delete	Enabled only when one or more nodes or movers are selected on the diagram. Removes the selected node(s) or mover(s) from the diagram. You can also press the **Delete** key.
Node or Node Group Icon	Each node type is represented by a different icon (see Node Type Icons (on page 657)). The currently selected node is enclosed in a box. **Tip:** If a warning condition is detected whilst constructing a data flow (e.g. a policy has only been assigned to a source node) then a red warning triangle will appear next to the offending node.

Control	Description
Multiple Selection	Click and drag the cursor over nodes on the workspace to select multiple nodes and movers. Alternatively press and hold **CTRL** and click multiple nodes.
Select All	Click on the workspace then press **CTRL+**+**A** to select all nodes and connectors.
Drag and Drop	Select and drag one or more nodes or movers to reposition them on the workspace. If the user moves a node, the icon will appear green if the node is dragged over another node it can be connected to, or red if not. If the user moves multiple items, the icons will not change color as creating connections from/to multiple items at once is not supported. If the user attempts to drop multiple items over the top of another item, the action will be canceled.
⏭ Next Node	Click to move the focus to the next node on the data flow.
+ − Zoom In/Out	Click the buttons next to the workspace, press +/- on the keypad or hold down the **CTRL** key whilst using the mouse wheel to zoom in and out. The current zoom level is displayed between the zoom buttons. Click the zoom level to reset to 100% zoom.
⛶ Fit to Screen	Click the button next to the workspace or press the **HOME** key to select a zoom level that allows the entire data flow to fit within the bounds of the workspace.
Pan	Right click anywhere on the workspace and drag the cursor.
Applied Policies (all)	If no nodes or movers are selected on the data flow workspace, then the area to the right of the workspace lists all the policies that have been applied in the data flow. Click the ⧉ button next to the policy name to open the Policy Details (on page 737) for that policy in a new browser tab.

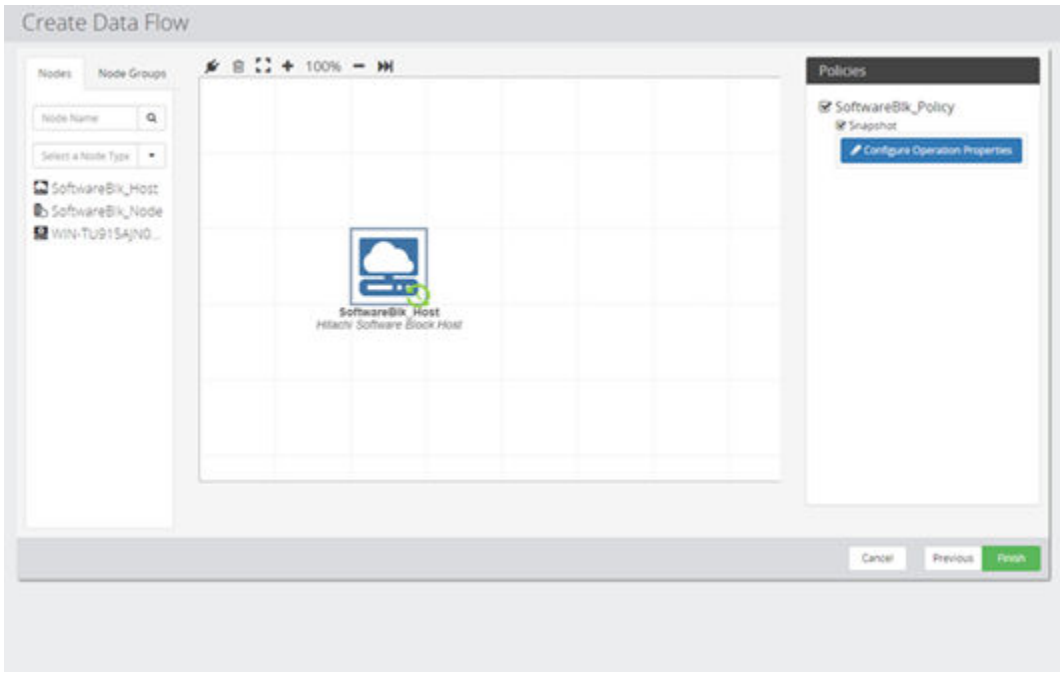


Figure 133 Data Flow Wizard - Source Node Policies

Table 24 Policies

Control	Description
Policies (source node selected)	If a single source node is selected on the data flow canvas, then the area to the right of the canvas lists all the policies that can be or are applied to that node. Apply policies to the node by clicking the required policy names. **Note:** If the selected policy requires a destination node to complete the policy definition, the source node displays a warning symbol to indicate that the policy definition is incomplete.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Finish	Commits the new changes. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

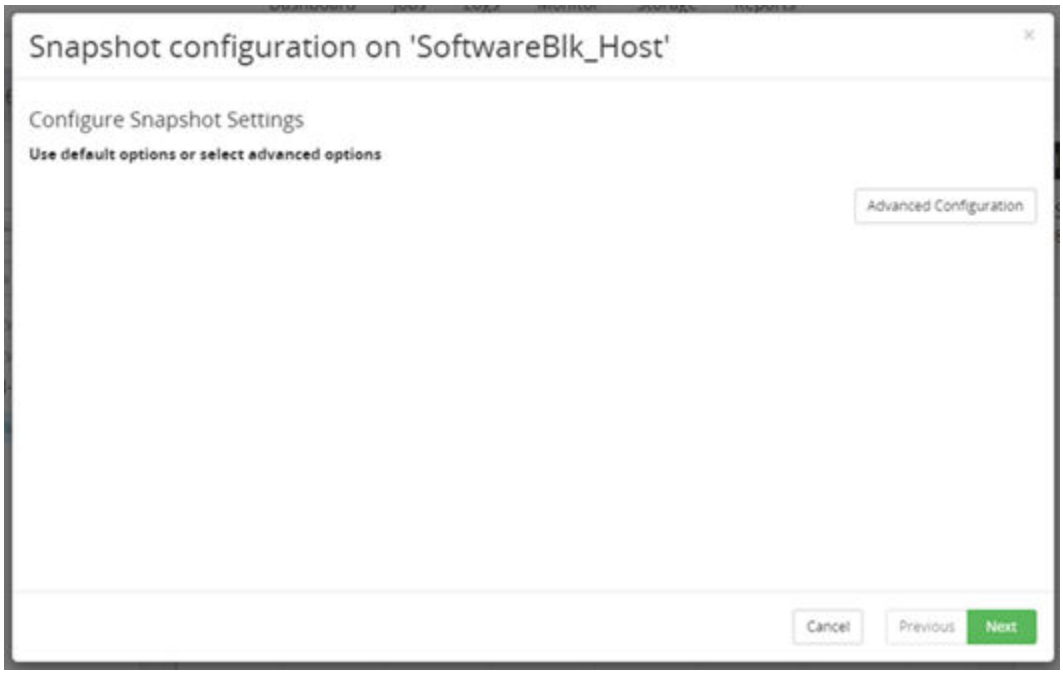


Figure 134 Data Flow Wizard - Configuration Snapshot Settings

Table 25

Control	Description
Advanced Configuration	Select Advanced Configuration to configure advanced settings for the Data Flow

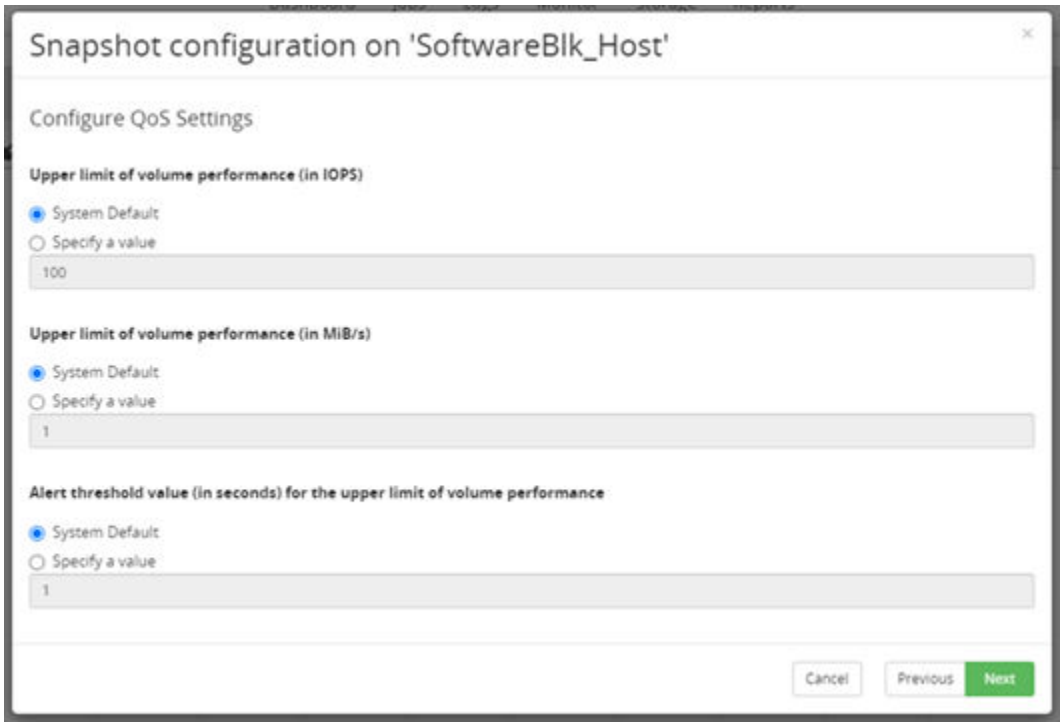


Figure 135 Data Flow Wizard – Snapshot Configuration Settings

Table 26

Control	Description
Upper limit of volume performance	Specify the upper limit of the volume performance. Select system default or specify a value in ▪ IOPS ▪ MiB/s
Alter threshold value	Specify threshold value in seconds for the upper limit of performance. Select system default or specify a value .

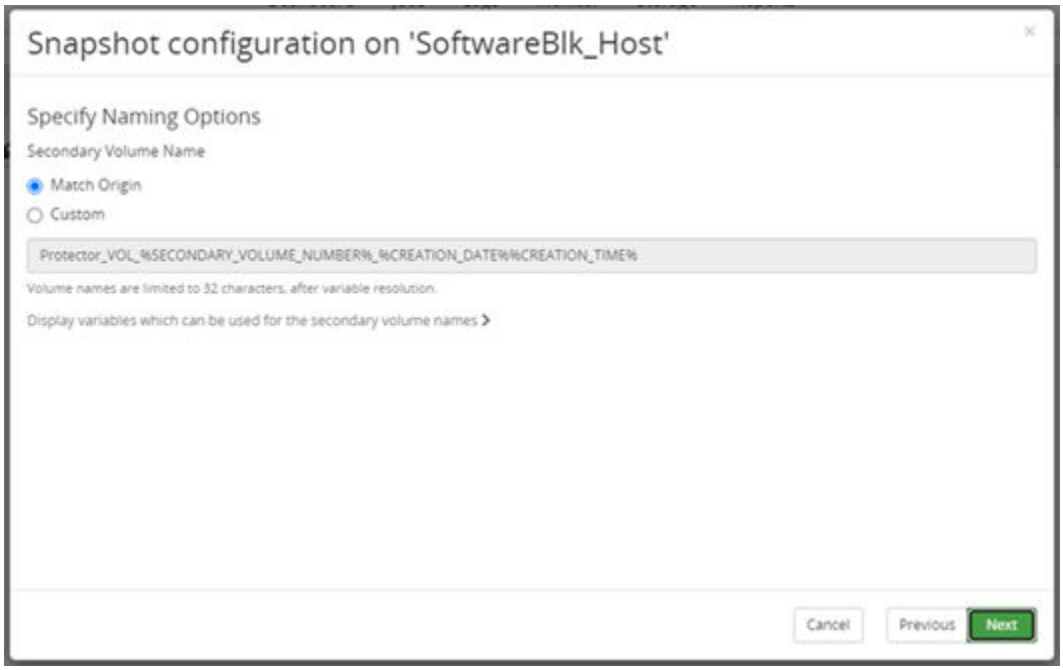


Figure 136 Data Flow Wizard – Snapshot Configuration – Specify Naming Options

Table 27

Control	Description
Specify Secondary Volume Name	Specify secondary volume name as ▪ Match Origin ▪ Custom

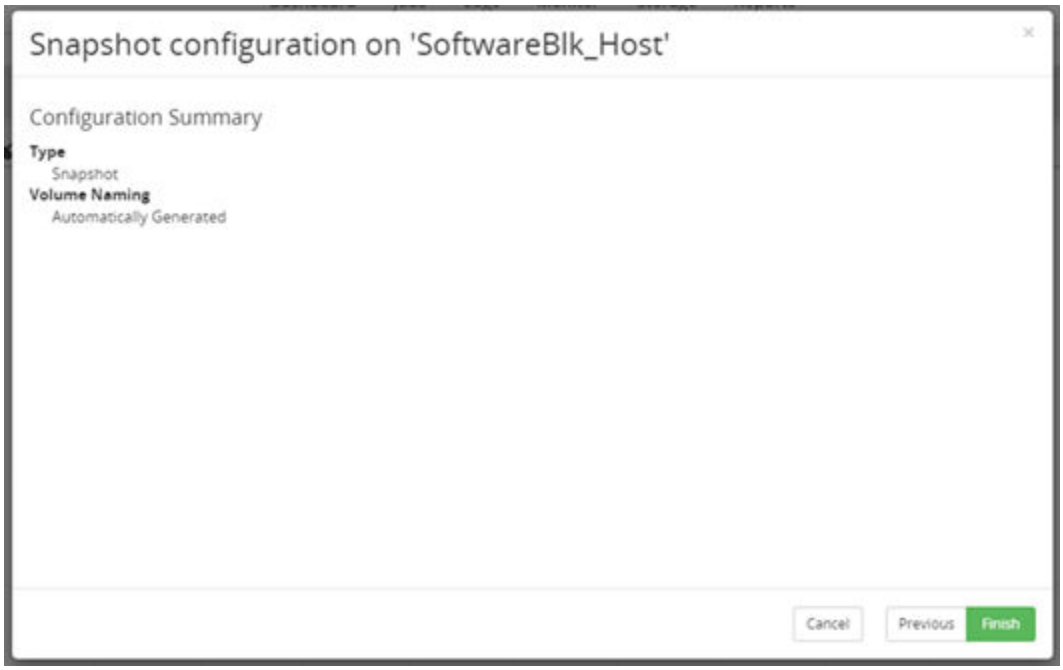


Figure 137 Data Flow Wizard – Snapshot Configuration Summary

Table 28

Control	Description
Summary	Summary of the settings entered.

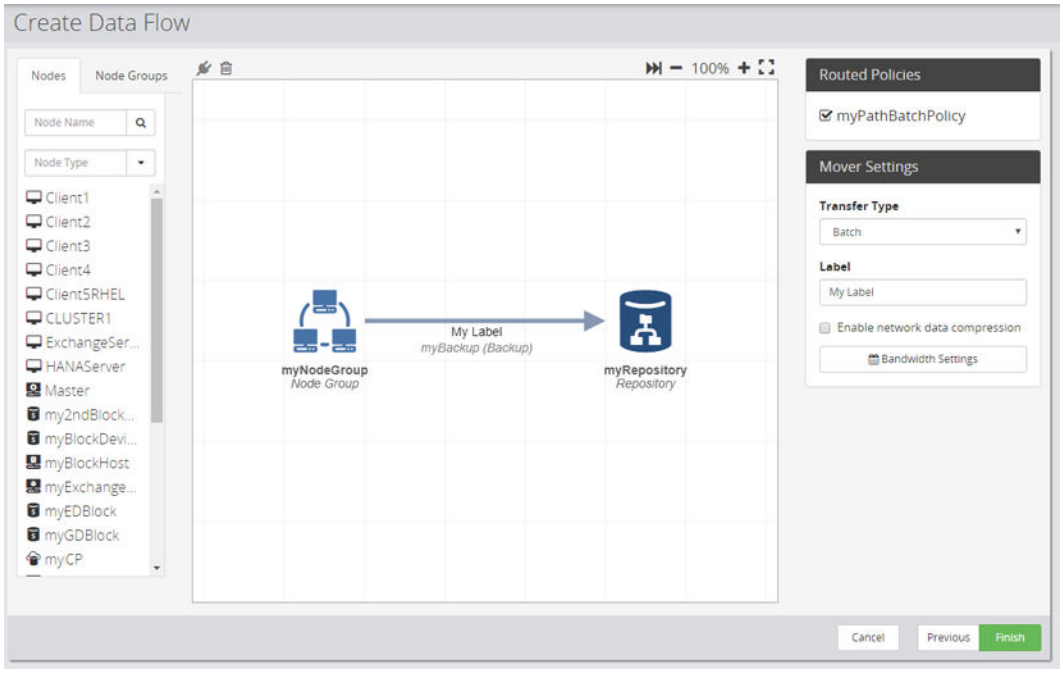


Figure 138 Data Flow Wizard - Mover Settings

Movers are connectors that represent how data is transferred from source node to destination node.

The mover provides options on how to route data for policies. If a source node is configured to implement a policy, then any movers connected to it automatically route that policy. If the mover is not configured to implement a policy, then the nodes further downstream are not able to either configure or implement the policy.

Table 29 Mover Settings

Control	Description
 Routed Policies (mover selected)	If a mover is selected on the data flow canvas, then the area to the right of the canvas lists all the policies that are routed by that mover.
Transfer Type	Changes the mover type to one of the following: ▪ Batch (solid line) - data is moved in batches based on a schedule or trigger event. ▪ Continuous (dotted line) - used only for Hardware operations. ▪ Failover (dashed line) - used only for 3DC delta resync data flows. Defines a suspended replication path, between the secondary and tertiary site, that can be invoked if the primary site fails. The arrowhead indicates the direction of data flow during normal backup operations. **Note:** It is import to use a mover type that is compatible with the operation type specified in the Policy. This requires an understanding of the snapshot and replication technologies being used. If an incompatible mover type is used, then you will only be notified about the error when you attempt to compile the rules. 1. Remove the replication from the data flow. 2. Reactivate the data flow. 3. Replace the replication on the data flow using the required mover type.
Label	By default the mover label displayed on the connector is empty. This can be replaced with a label describing the connection on the data flow. This feature can be useful in situations where there is more than one mover connected to a node.

Control	Description
Enable network data compression	For Host based operations only (ignored for Hardware based operations). Turns on data compression on the datalink between two nodes. **Note:** Compression decreases network utilization, but increases CPU usage.
Bandwidth Settings	For Host based operations only (ignored for Hardware based operations). Opens Mover Bandwidth Settings Dialog (on page 375) to enable you to control the amount of bandwidth that data transfers use so that a network connection can be used simultaneously with other applications in the environment.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Finish	Commits the new changes. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.



Figure 139 Data Flow Wizard - Destination Node Policies

Table 30 Policies

Control	Description
 Policies (destination node selected)	If a single destination node is selected on the data flow canvas, then the area to the right of the canvas lists all the policies and their contained operations that can be or are applied to that node. Apply operations to the node by clicking the required operation names. A dialog opens, appropriate to the destination node type, that allows those properties to be specified (see Configure Operation Properties). The policies listed here are restricted to those that have been applied to upstream nodes and those that the destination node has the appropriate capabilities to receive.  **Note:** A single policy can have several operations within it for a given data classification. The destination indicates which of those operations it can implement by displaying only the supported operation check boxes. The policy name check box cannot be selected by the user, the individual operations within the policy must be selected.
Configure Operation Properties	When an operation is applied to a destination node you may need to specify properties for that operation. Click this button to open a dialog, appropriate to the destination node type, that allows those properties to be specified: ▪ Generation 1 Repository Backup Configuration Dialog (on page 377) ▪ Hitachi Block Snapshot Configuration Wizard (on page 384) ▪ Hitachi Block Replication Configuration Wizard (on page 398) ▪ Hitachi Block Mount Configuration Wizard (on page 485) Once the properties are specified, the button is replaced with a summary of the properties. Click the Edit button next to the summary to reopen the dialog to change any properties. The operation type is displayed in italics next to the mover along with the user defined label.  **Note:** For Hitachi Block Replication Operation Properties: Once the operation properties have been configured and the data flow has been activated, the properties cannot be changed. Redistributing rules for the active data flow with edited properties will not change them. To change the operation properties, the existing data flow must first be deactivated then reactivated with the new properties. Deactivating the data flow will mark the replication eligible for tear down.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.

Control	Description
Finish	Commits the new changes. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

Mover Bandwidth Settings Dialog

This dialog is used to specify when and how much bandwidth can be used by Protector to transfer data across the datalink between two nodes.

You can assign specific bandwidth constraints for defined time windows so that a network connection can be used simultaneously with other applications in the environment. For instance, you can decide to decrease bandwidth during the workday when bandwidth is required for production and increase it at night to allow backups to utilize maximum bandwidth.



Note: If the amount of data that must be replicated continuously exceeds the current bandwidth quota, then data is cached on the source until the bandwidth increases. Cache size on the Source Node can be adjusted to avoid any problems with non-transferred data reaching the cache limit. The configuration file `distributor.cfg` contains the values *MaxDiskCache* and *MaxMemoryCache* which can be increased if the cache limit becomes a problem.

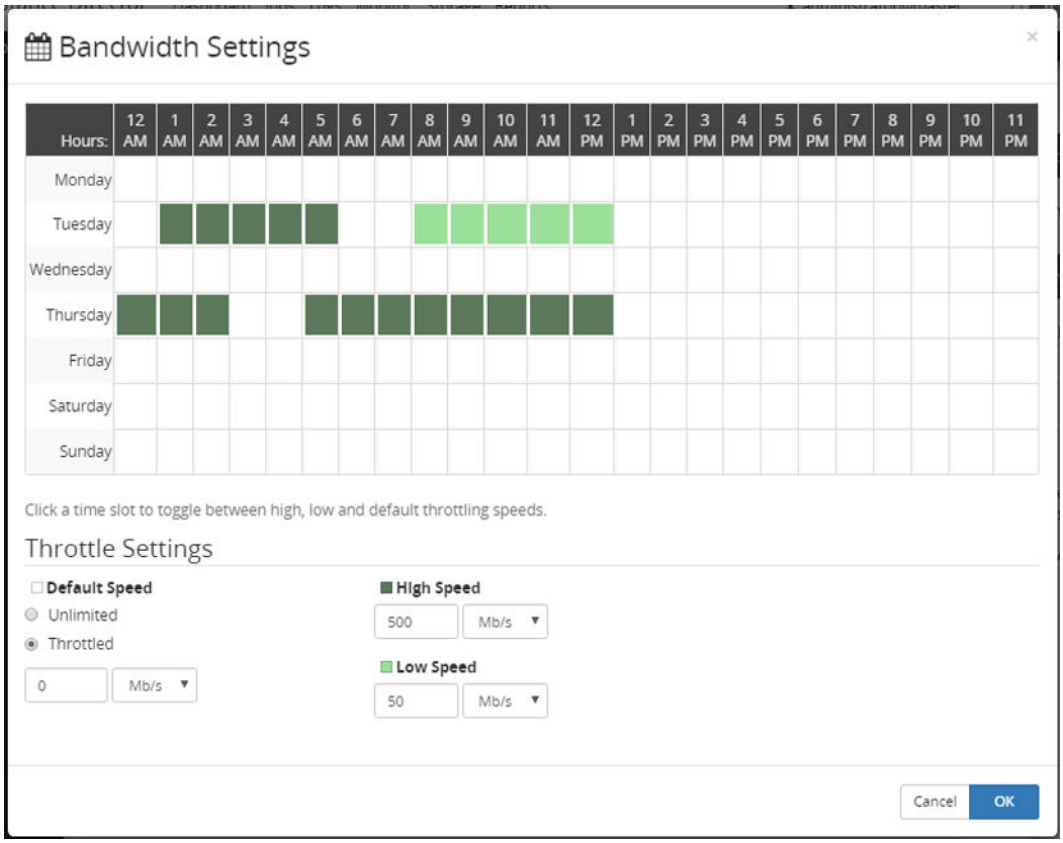


Figure 140 Bandwidth Settings Dialog

Control	Description
Week Grid	Click the cells in the grid, corresponding to the hour-long periods where you want to throttle network bandwidth used by Protector, to one of the following levels (cells cycle through the three states each time they are clicked): ▪ High - (Dark Green) High Speed throttling at the setting defined below. ▪ Low - (Light Green) Low Speed throttling at the setting defined below. ▪ None - (White) Default Speed throttling at the setting defined below.
Unlimited/ Thottled	Sets the default speed to the maximum allowed (Unlimited) by the network or to a predefined level (Throttled).
Default Speed	Specifies the default throttle level. This may be set to any value including 0.
High Speed	Specifies the high throttle level. This must be set to a non-zero value and greater than the Low Speed level.

Control	Description
Low Speed	Specifies the low throttle level. This must be set to a non-zero value and less than the High Speed level.

Generation 1 Repository Backup Configuration Dialog

This dialog is displayed when you assign an operation to a Generation 1 Repository node on a data flow.

When using a repository as a destination for a data flow a number of configuration options are available. These destination options for a repository node are contained within Store Templates. Each repository node can be associated with multiple stores templates. There are two default store templates (Standard and Deduplicated), however additional templates can be created.

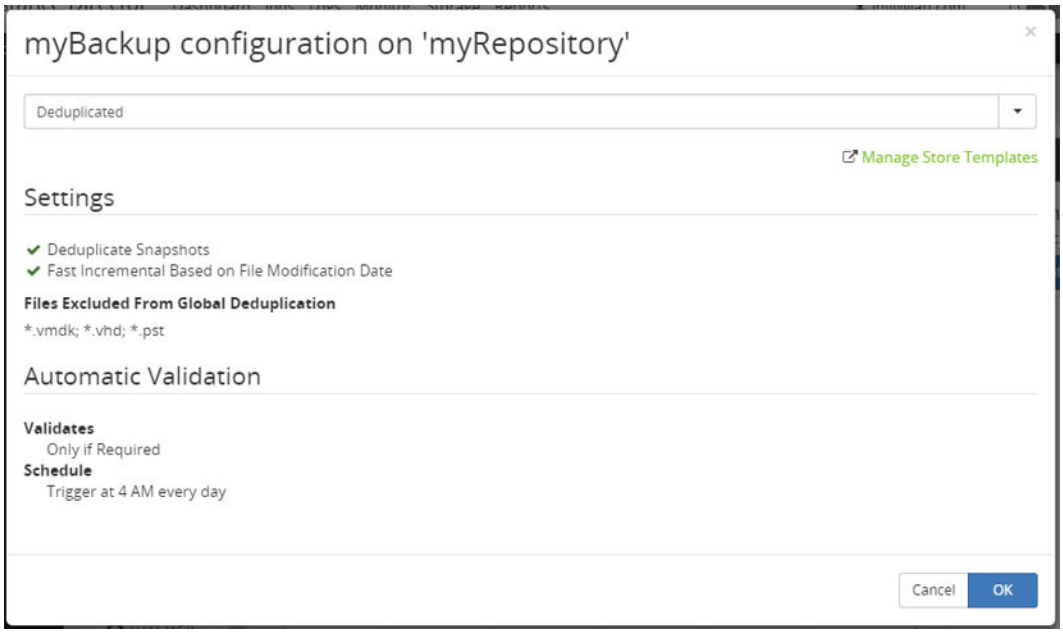


Figure 141 Repository Backup Configuration Dialog

Control	Description
Select Destination Template	Enter or select a Destination Template from the dropdown list. Once selected, the settings of the store template are displayed below.

Control	Description
Manage Store Templates	Click this link to add to or edit the available destination templates. The Generation 1 Repository Destination Templates Inventory (on page 378) is opened in a new browser tab. **Note:** Once you have finished adding or editing templates in the new tab, simply close the tab and continue working with the Operation Properties Dialog. Any changes you have made to the templates will be applied, although they may not appear in the dialog unless you re-select the template.
Settings	Lists the settings for the selected template.

Generation 1 Repository Destination Templates Inventory

This inventory is displayed when managing Destination Templates for a Generation 1 Repository node.

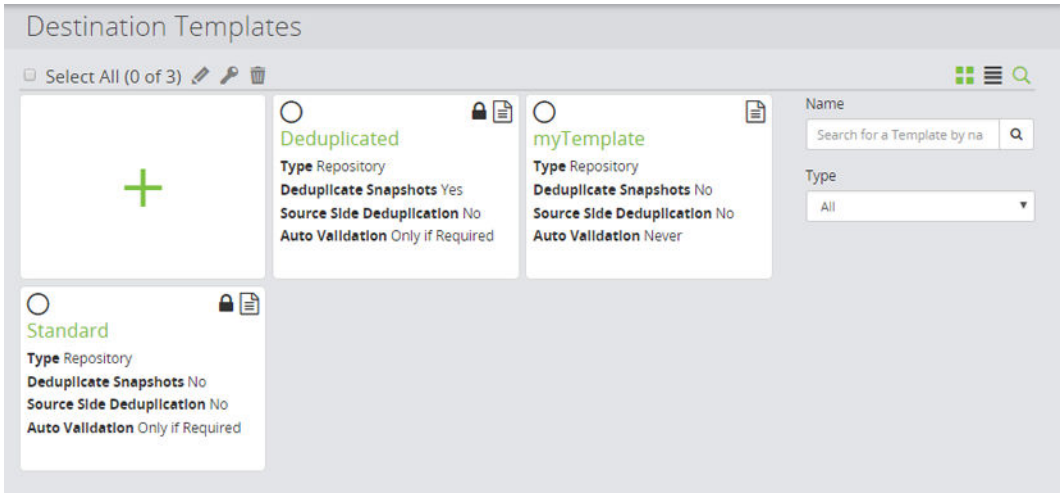


Figure 142 Destination Templates Inventory

Control	Description
Edit	Edits an existing template in the inventory. The Destination Template Wizard (on page 379) is launched to enable the template's attributes to be changed.
Edit Permissions	Edits an existing template's access permissions. The Access Control Permissions Inventory (on page 349) is launched to enable the template's access permissions to be changed.

Control	Description
 Delete	Enabled only when one or more Templates is selected. Deletes the selected item from the inventory.
 Add	Creates a new Template. The Destination Template Wizard (on page 379) is launched to guide you through the process.
 System Generated Templates	At least two system generated Templates are available when the product is installed. These Templates cannot be deleted since they provide basic functionality. System generated Templates are marked with a  icon to indicate that they cannot be modified. Click the template's name to open the Destination Template Details (on page 383) to enable the parameters to be viewed.
 User Defined Template(s)	Any number of user defined Templates can be created. These are displayed in the inventory. Click the template's name to open the Destination Template Details (on page 383) to enable the parameters to be viewed.
Filter on Template Name	Filters the displayed results based on the template name.
Filter on Template Type	Filters the displayed results based on the template type.

Destination Template Wizard

This wizard is displayed when a new Destination Template is being created.

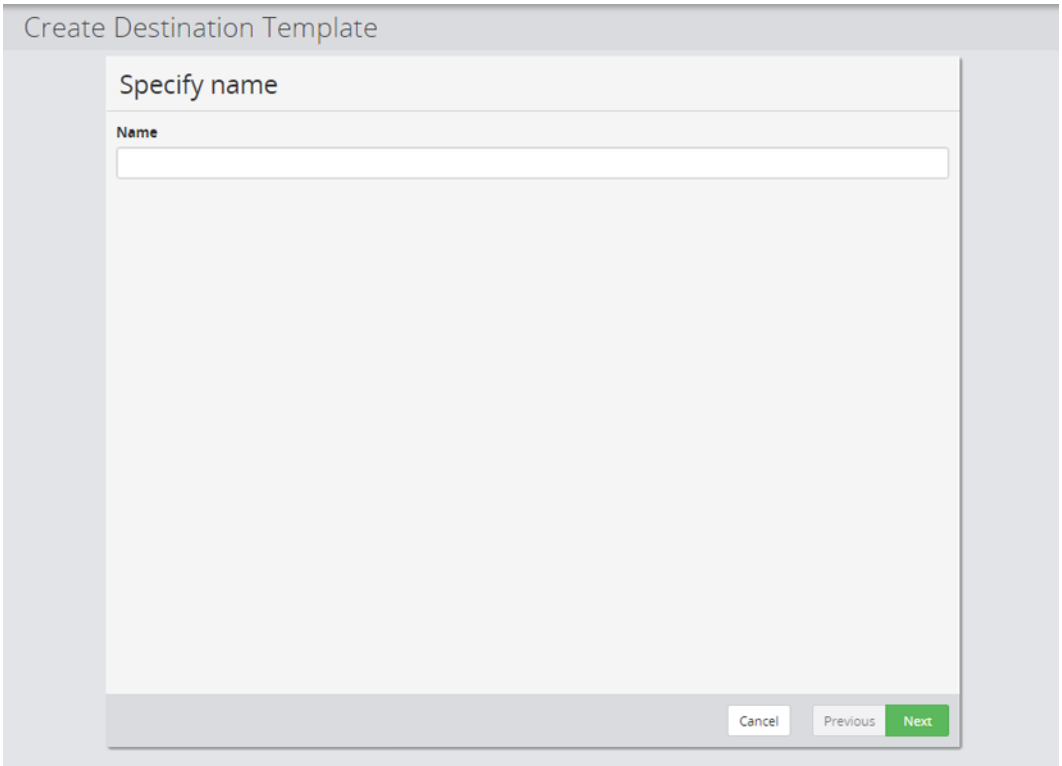


Figure 143 Destination Template Wizard - Specify name

Control	Description
Name	Enter a name for the template.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

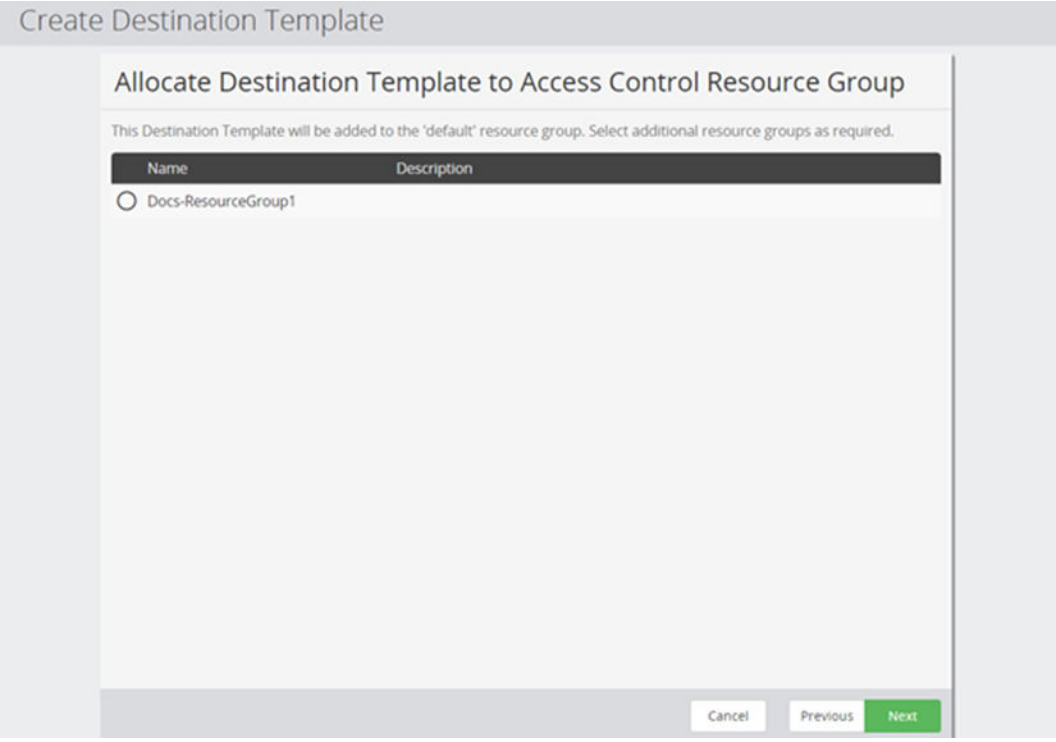


Figure 144 Destination Template Wizard - Allocate Destination Template to Resource Group

Control	Description
Resource Groups	It allows the user to view the access permissions for those items granted to specific users and groups. **Note:** A single Destination Template can be assigned to multiple resource groups.

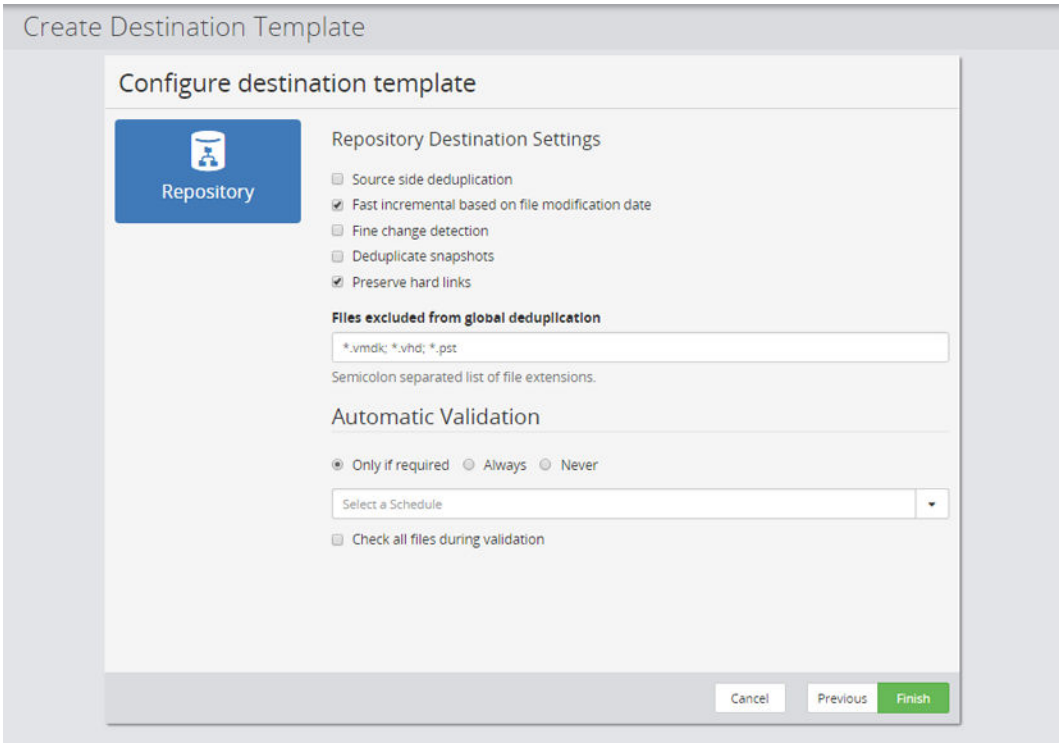


Figure 145 Destination Template Wizard - Configure destination template

Control	Description
Source side deduplication	On occasions where many machines have identical roles within the data flow and also contain very similar data (for instance, OS data and installed software on machines on a corporate network) then a network speedup can be obtained by avoiding sending the same data multiple times. Check this option to make use of this.
Fast incremental based on file modification date	If this is checked then Ops Center Protector decides what files need resynchronizing based on whether the modification date has changed. This reduces the time taken to resynchronize, but can be disabled if it is known that software is installed that will modify files without updating their size or modification date. **Note:** If only file metadata changes between batch backups (e.g. file ownership or file permissions), then the changes are not captured. These changes are only captured when the file data changes.
Fine change detection	Reduces the amount of data transferred and stored during a resynchronization. An entire file is transferred if it has changed and is less than one block in size. (This option should be used sparingly as there is a processing overhead.)

Control	Description
Deduplicate snapshots	Enables the storage group to deduplicate data across snapshots so the storage group only stores a single copy of the data. This option has a processing overhead.
Preserve hard links	Performs checks so that only one instance of the data is stored, regardless of the number of links pointing to it. This option increases the file system scan time during a resynchronization or batch backup. **Note:** As of Protector 6.5, this option is enabled by default. This is not retrospective and will only apply when new stores are created.
Files excluded from global deduplication	A semicolon-separated list of file extensions to be excluded from the store group's duplication detection can be entered here.
Automatic Validation	Automatically resynchronize source to the destination based on a schedule. One of the following options can be selected: ▪ Only if required – Will trigger resynchronization only if the destination is out of sync with the source nodes. ▪ Always – Will always perform a resynchronization with the source nodes based on the schedule. ▪ Never – An automatic validation will never be performed. Checking this option will disable the remaining options.
Select a Schedule	Select a schedule to trigger automatic validation. **Note:** Schedules are created for automatic validation the same way as schedules for a policy. A schedule for automatic validation must use a Trigger, by default a pre created schedule named Trigger at 4 AM every day is selected.
Check all files during validation	Ignores file modification date and checks every file for changes.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Finish	Commits the new settings. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

Destination Template Details

This page displays the details of a Destination Template and enables you launch the wizard to edit them.

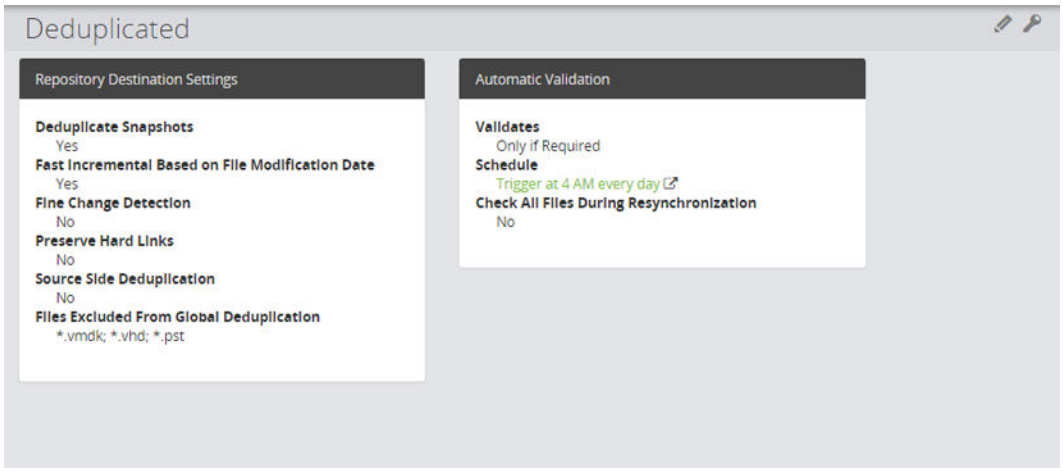


Figure 146 Destination Template Details

Control	Description
Edit	Launches the Repository page of the Destination Template Wizard (on page 379) to enable you to edit the template.
Permissions	Displays the Access Control Permissions Inventory (on page 349) to enable you to view and edit the template's permissions.
Repository Destination Settings	These are the settings entered via the Repository page of the Destination Template Wizard (on page 379) when the Template was created.
Automatic Validation	These are the settings entered via the Repository page of the Destination Template Wizard (on page 379) when the Template was created. Click on the Schedule name to open the Schedule Details (on page 857) in a separate browser tab.

Hitachi Block Snapshot Configuration Wizard

This dialog is displayed when you assign a snapshot operation to a node that hosts data stored on an Hitachi Block device.

Caution: When a snapshot operation runs, Protector locks the *meta_resource* of the Block device until the operation completes. Block operations are queued waiting to get the resource lock; this can impact RPO.

Note:

All replication and snapshot S-VOLs must be created using free LDEV IDs that are mapped to the *meta_resource* group, and have virtual LDEV IDs matching their corresponding physical LDEV IDs.

For fully provisioned snapshots and all replications, this applies to the operation that creates that snapshot or replication.

For floating device snapshots and snapshots mounted using cascade mode, this applies to the mount or restore operation.

For fully provisioned snapshots mounted using cascade mode, this applies both to the operation that creates that snapshot and to the mount or restore operation.

If an operation tries to create one or more LDEVs, that operation will fail if there are not enough free LDEV IDs that meet the above conditions.

Note:

- The classifications that are applicable to Hitachi Block Snapshot operations are Path, Application, Hypervisor and Hitachi Block.
- Protector does not support any Hitachi Block storage LDEV with more than one partition or volume.
- If Protector is subsequently uninstalled, the existing snapshots and mounted volumes are left in place.

Tip:

- The actions that are performed by Protector as it executes a backup policy are captured by the Log Manager. Detailed logs and their attachments can be exported as a text file. (See Export Logs Dialog (on page 516) for more information.)
- Details of the hardware resources (LDEVs) used for a particular snapshot can be found in the Hitachi Block Device Details (on page 859).

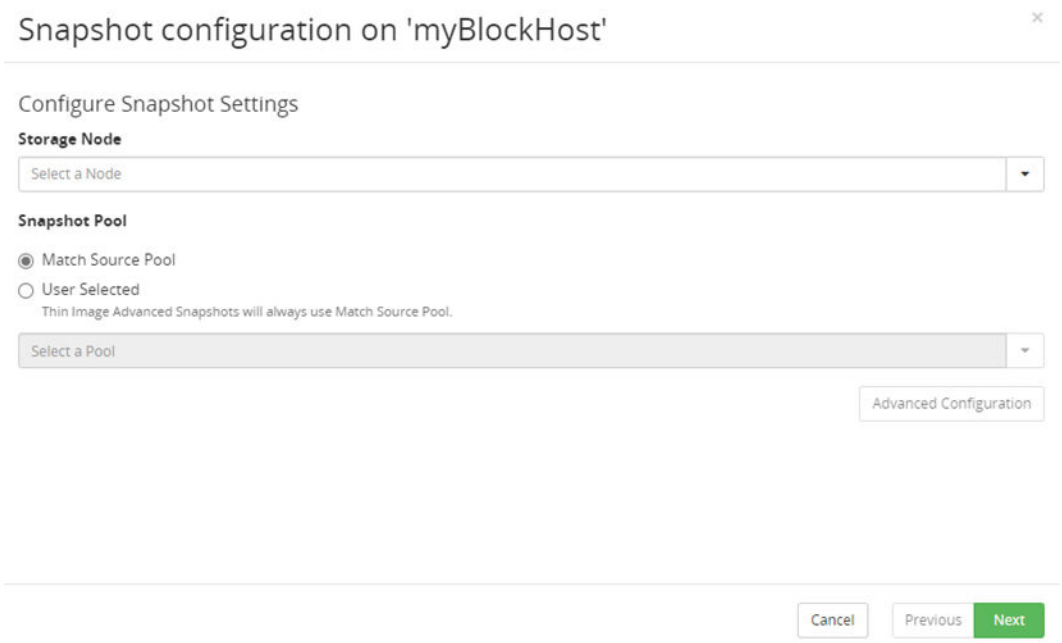


Figure 147 Snapshot Configuration Wizard - Differential snapshot (using Thin Image)

Control	Description
Storage Node	Specifies the target storage node where the P-VOL and snapshot S-VOLs are allocated.
Snapshot Pool	Specifies the target storage pool from one of the following options to which snapshot S-VOLs are allocated. ▪ Match Source Pool – Allocates the snapshot in the same pool as the source volumes. ▪ User Selected – Select a Thin Image Pool or a Dynamic Provisioning Pool. Select a Thin Image Pool or a Dynamic Provisioning Pool. **Caution:** Filling a Thin Image pool to capacity will invalidate all snapshot data contained within that pool. All snapshots in the pool will have to be deleted before snapshotting can be resumed. **Note:** Thin Image and Dynamic Provisioning Pools must be created using Storage Navigator prior to selecting the Target Storage in Protector.

Control	Description
	Tip: If the P-VOL has an association with a VSM or you select a target storage pool associated with a VSM then Protector will attempt to make the virtual and the physical ID of the S-VOL holding the snapshot identical. For this to work, an LDEV with a physical ID within the defined virtual LDEV range for the selected VSM must available for use.
Advanced Configuration	Click to step through the advanced configuration option pages of the wizard, described below.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

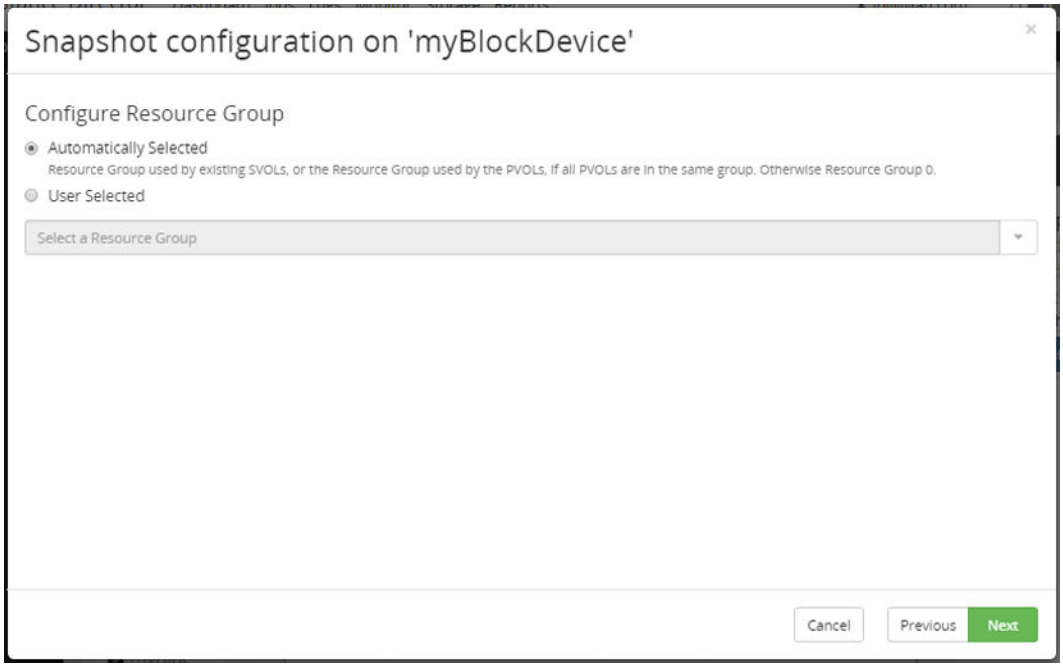


Figure 148 Snapshot Configuration Wizard - Configure Resource Group

Specifies the resource group to be used for S-VOLs, in order to support snapshots and replications from VSM volumes (adding volumes to a VSM is performed by adding the volumes to the correct resource group).

Control	Description
Automatically Selected	Allows Protector to automatically select a resource group in the following order of priority:

Control	Description
	1. If there are existing S-VOLs, then the resource group used by those will be selected. 2. The resource group used by the P-VOLs, if the replication is in-system and the P-VOLs are all in one resource group. 3. Resource group 0. **Note:** If existing S-VOLs are in multiple resource groups, then the operation will fail with an error.
User Selected	Specify the Resource Group in the associated combo-box. **Note:** If there are existing S-VOLs, then the resource group used by those will be selected. If the existing S-VOLs are in multiple resource groups or in a resource group that contradicts the user specification, then the operation will fail with an error.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

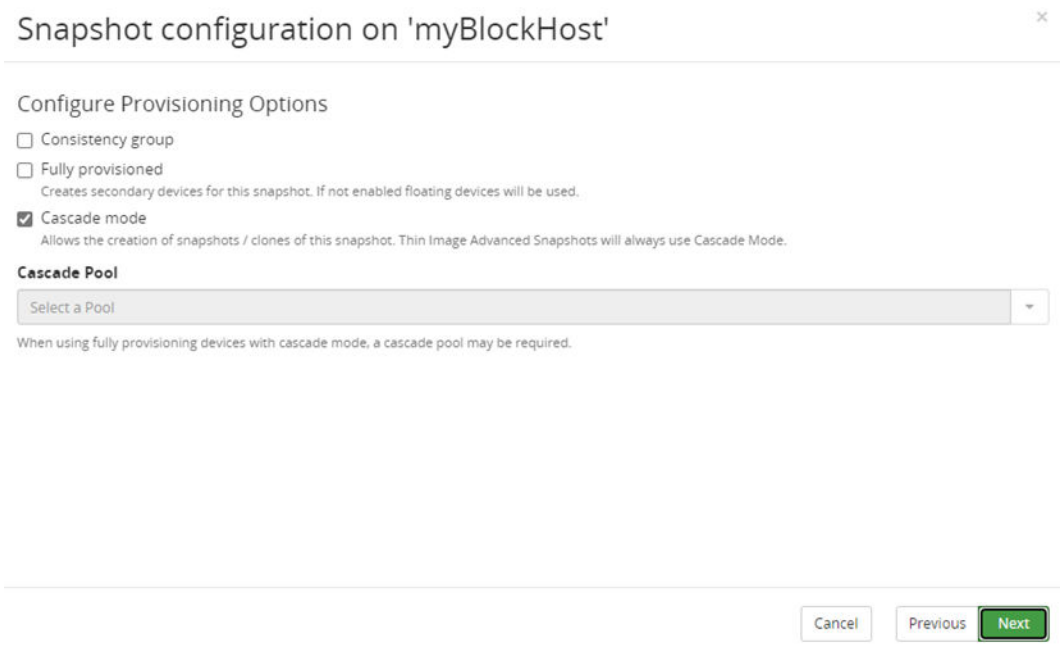


Figure 149 Snapshot Configuration Wizard - Provisioning Options

Specifies how the snapshot is provisioned and its mode of operation. Snapshots are created as differential, in-system snapshots using Thin Image.

Caution: Block storage has a limit of 1024 snapshots per LDEV. Ensure that the RPO and Retention periods are set such that this limit is not exceeded (i.e. Retention / RPO is less than or equal to 1024).

Control	Description
Consistency group	When performing a crash consistent snapshot (for example: using the hardware path classification), use this option to make it truly crash consistent. The hardware has a limited number of consistency groups available so they should be used sparingly.
Fully Provisioned	By default, this option is deselected and the snapshot will be provisioned using floating devices. Where supported by the array the storage device can store a larger number of snapshots using floating device.
Cascade Mode	The snapshot can be cascaded. This allows Protector to mount either the original snapshot or a duplicate of the original. Mounting a duplicate enables modifications to be made without affecting the original snapshot. Refer to About cascaded Thin Image and Thin Image Advanced snapshots (on page 69) for further guidance on configuring and mounting cascade mode snapshots.
Cascade Pool	If Fully Provisioned is selected above, it may be necessary to specify a dynamic or hybrid pool where Protector can create snapshot S-VOLs.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

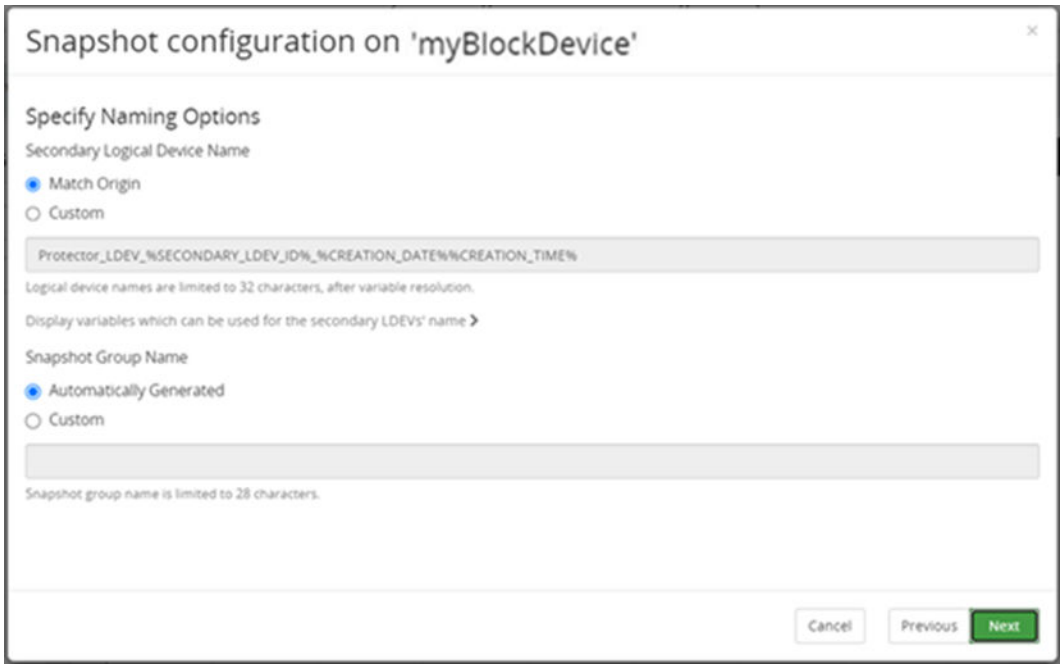


Figure 150 Snapshot Configuration Wizard - Naming Options

Specifies how secondary LDEVs and snapshot groups will be named.

Control	Description
Secondary Logical Device Name	Specifies how S-VOLs will be named: ▪ Match Origin - The S-VOL will be given the same name as that used for the origin P-VOL (i.e. the left-most volume in the data flow). ▪ Custom - The S-VOL will be named using the naming rule provided. The naming rule can consist of literal strings and/or one or more substitution variables listed. Click Display variables which can be used for the secondary LDEVs' name to view the available substitution variables: • `%ORIGIN_SERIAL%` - S/N of leftmost array in data flow. E.g. output string: `210613` • `%ORIGIN_LDEV_ID%` - ID of leftmost LDEV in data flow. E.g. output string: `00:3A:98` • `%ORIGIN_LDEV_NAME%` - name of leftmost LDEV in data flow. • `%PRIMARY_SERIAL%` - S/N of primary array in this operation. E.g. output string: `442302` • `%PRIMARY_LDEV_ID%` - ID of primary LDEV in this operation. E.g. output string: `00:4C:EB` • `%PRIMARY_LDEV_NAME%` - name of primary LDEV in this operation. • `%SECONDARY_SERIAL%` - S/N of secondary array in this operation. E.g. output string: `356323` • `%SECONDARY_LDEV_ID%` - ID of secondary LDEV in this operation. E.g. output string: `01:F4:35` • `%CREATION_DATE%` - date secondary LDEV was created by this operation. E.g. output string: `20180427` • `%CREATION_TIME%` - time secondary LDEV was created by this operation. 1130
Snapshot Group Name	Specifies how the snapshot group will be named when the Provisioning Options - Consistency group option is not selected. ▪ Automatically Generated - The snapshot group name is generated by Protector based on the rules context ID and policy name. ▪ Custom - The snapshot group is named using the string provided (limited to 28 characters). An '@' separator followed by a unique ID is then automatically appended to this name. The unique ID is composed of 3 base 36 characters and is required to enable Protector to manage the groups.
Cancel	Discards all changes and reverts to the previous page.

Control	Description
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

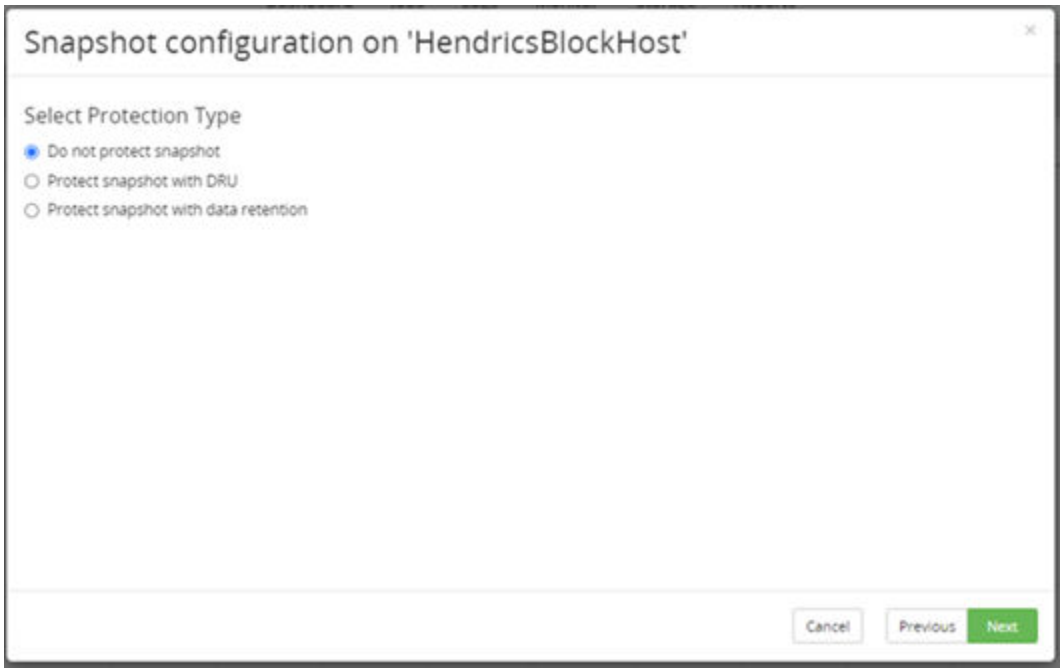


Figure 151 Snapshot Configuration Wizard - Protection Type

Control	Description
Select Protection Type	Select Between the following options: ▪ Do not protect snapshot ▪ Protect snapshot with DRU ▪ Protect snapshot with data retension
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

If Protect snapshot with DRU is selected the following screen is displayed:

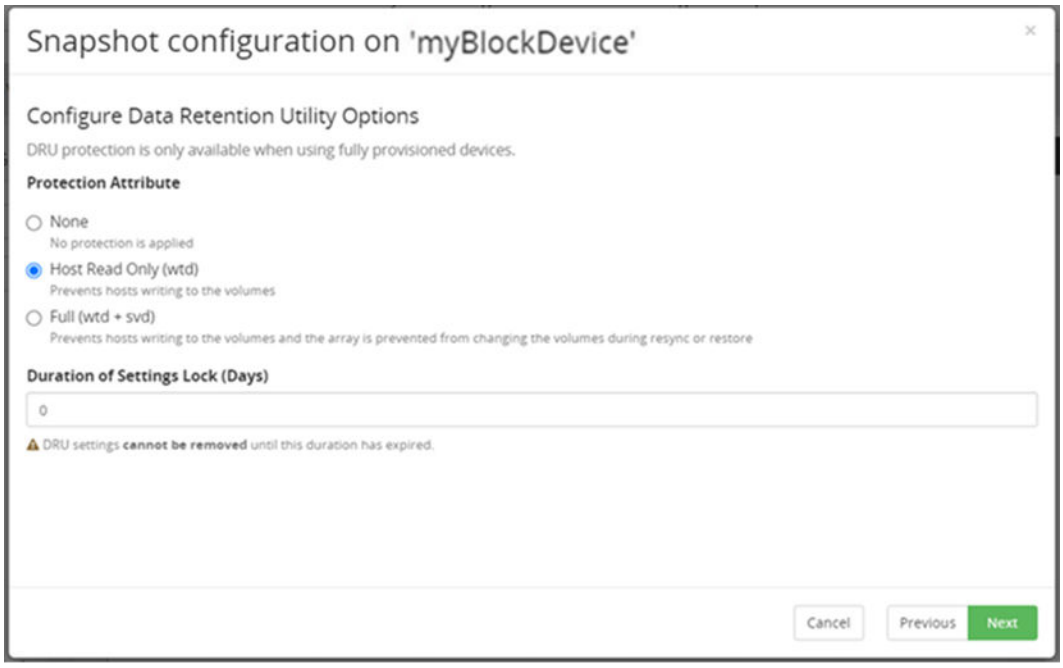


Figure 152 Snapshot Configuration Wizard - DRU Options

If Fully Provisioned is selected on the **Provisioning Options** page on this wizard, these options are enabled to specify Data Retention Utility (DRU) protection parameters.

Caution: Protector cannot mount a DRU protected snapshot. However a cascaded snapshot can be created and mounted if the original has Cascade mode enabled.

Control	Description
Protection Type	One of the following can be selected: ▪ None - DRU protection is not applied the snapshot ▪ Host Read Only (wtd) - Prevents hosts writing to the snapshot LDEV. ▪ Full (wtd + svd) - As for wtd, plus the array is prevented from changing the contents of a snapshot LDEV using resync or restore. It also prevents deletion, mapping and unmapping of the snapshot from its LDEV. **Note:** DRU can prevent mounting if the OS rejects a read-only volume (mount via cascade snapshot is unaffected, and therefore recommended).

Control	Description
Duration Of Settings Lock (Days)	**Caution:** DRU protection cannot be removed while the lock is active. Specify a duration in days during which the applied DRU protection cannot be removed. Once this duration has expired the protection is not automatically removed; it must be removed manually.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

If Protect snapshot with data retension is selected the following screen is displayed:



Figure 153 Snapshot Configuration Wizard - Configur Snapshot Data Retention

Control	Description
Data retention lock period	**Caution:** Protection cannot be removed while the lock is active. Specify a duration in hours, days, weeks, months or years during which the applied protection cannot be removed. Once this duration has expired the protection is not automatically removed; it must be removed manually.

Control	Description
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

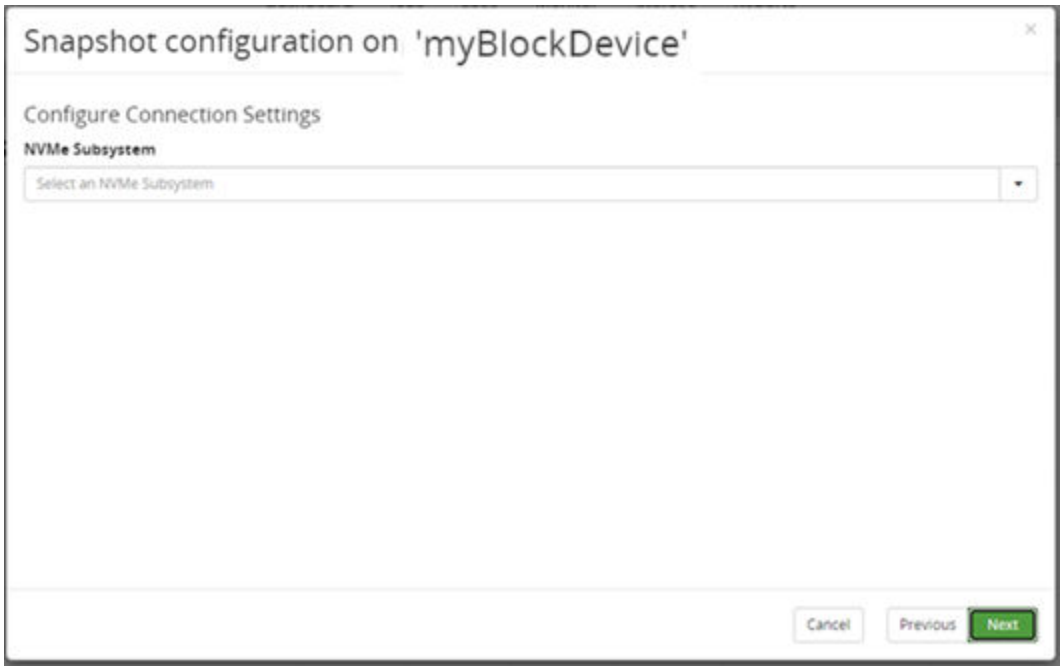


Figure 154 Snapshot Configuration Wizard - Configure Connection Settings

Control	Description
NVMe Subsystem	**Note:** If required select an NVMe subsystem from the dropdown list.

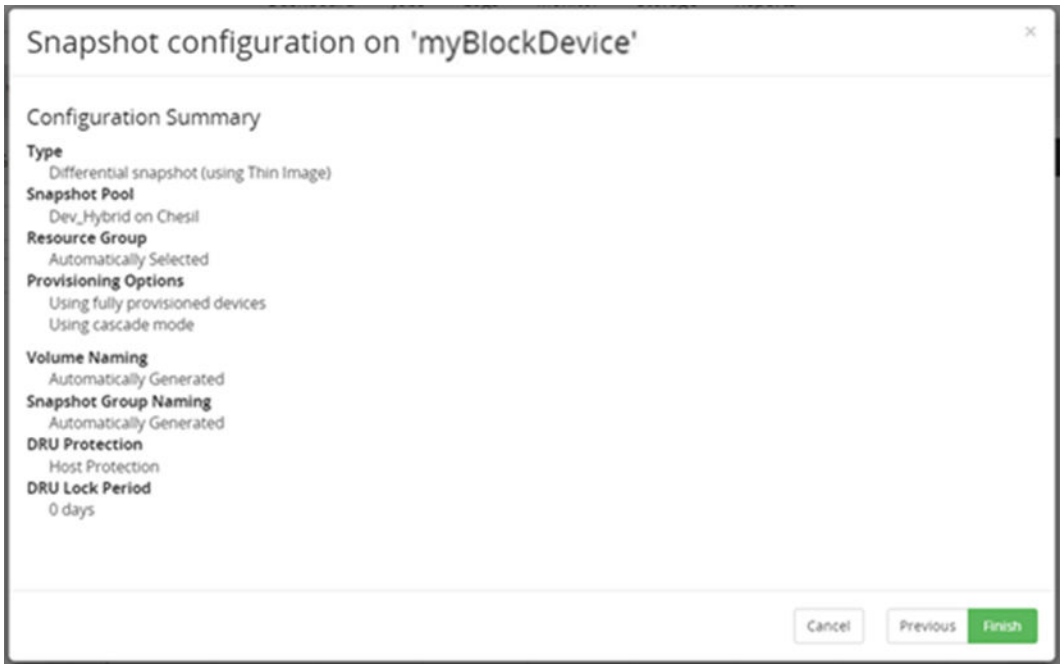


Figure 155 Snapshot Configuration Wizard - Summary

Summarizes the configuration settings made by the user.

Hitachi Software Block Snapshot Configuration Wizard

This dialog is displayed when you assign a snapshot operation to a node that hosts data stored on an Hitachi Software Block device.

Tip:

- The actions that are performed by Protector as it executes a backup policy are captured by the Log Manager. Detailed logs and their attachments can be exported as a text file. (See Export Logs Dialog (on page 516) for more information.)
- Details of the hardware resources (LDEVs) used for a particular snapshot can be found in the Hitachi Software Block Device Details (on page 901).

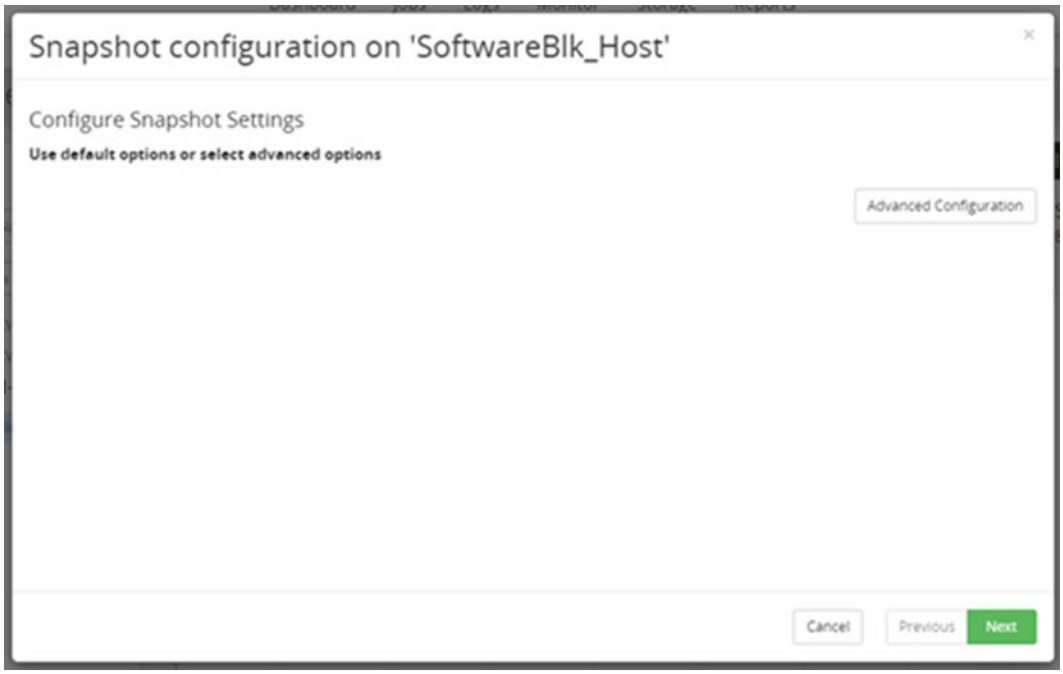


Figure 156 Data Flow Wizard - Software Block Configuration Snapshot Settings

Control	Description
Advanced Configuration	Select Advanced Configuration to configure advanced settings for the Data Flow.



Figure 157 Data Flow Wizard – Snapshot Configuration – Specify Naming Options

Control	Description
Specify Secondary Volume Name	Specify secondary volume name as ▪ Match Origin ▪ Custom

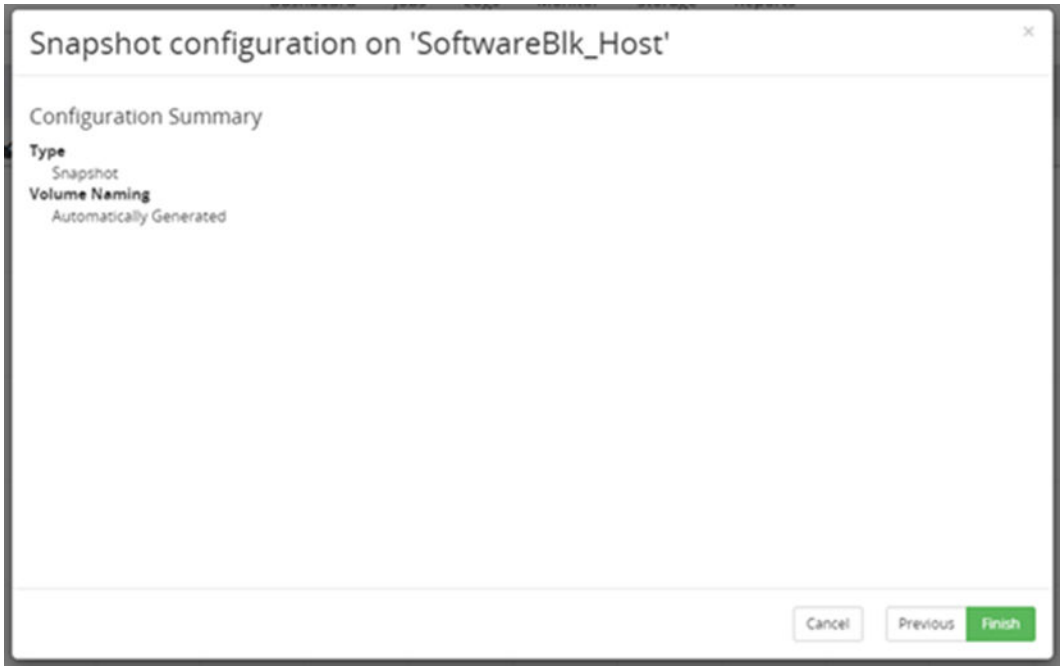


Figure 158 Data Flow Wizard – Snapshot Configuration Summary

Control	Description
Summary	Summary of the settings entered.

Hitachi Block Replication Configuration Wizard

This wizard is displayed when you assign a replication operation to a Hitachi Block Device node on a data flow.

Caution: When a replication operation runs, Protector locks the *meta_resource* of the Block Device until the operation completes. Block operations are queued waiting to get the resource lock; this can impact RPO.

Note:

All replication and snapshot S-VOLs must be created using free LDEV IDs that are mapped to the *meta_resource* group, and have virtual LDEV IDs matching their corresponding physical LDEV IDs.

For fully provisioned snapshots and all replications, this applies to the operation that creates that snapshot or replication.

For floating device snapshots and snapshots mounted using cascade mode, this applies to the mount or restore operation.

For fully provisioned snapshots mounted using cascade mode, this applies both to the operation that creates that snapshot and to the mount or restore operation.

If an operation tries to create one or more LDEVs, that operation will fail if there are not enough free LDEV IDs that meet the above conditions.

Note:

- The classifications that are applicable to Hitachi Block Replication operations are File Path, Application, Hypervisor and Hitachi Block LDEV.
- Protector does not support any Hitachi Block storage LDEV with more than one partition or volume.
- If Protector is subsequently uninstalled, the existing replications and mounted volumes are left in place.

Note: Protector will attempt to match destination LDEV IDs with those used by the source (except for in-system SI and RTI replications where this is impossible). The LDEV ID must be within the range configured for the target storage system and the LDEV ID must not be in use. If the LDEV ID cannot be matched on the target then the first available in-range LDEV ID will be selected.

Note:

Once the operation properties have been configured and the data flow has been activated, the properties cannot be changed. Reactivating rules for the active data flow with edited properties will not change them.

To change the operation properties, the existing data flow must first be deactivated then reactivated with the new properties. Deactivating the data flow will tear down the current replication.

Tip:

- The actions that are performed by hardware orchestration as it executes a block backup policy are captured by the Log Manager. Detailed logs and their attachments can be exported as a text file. (See Export Logs Dialog (on page 516) for more information.)
- Details of the hardware resources (LDEVs) used for a particular replication can be found in the Hitachi Block Device Details (on page 859).

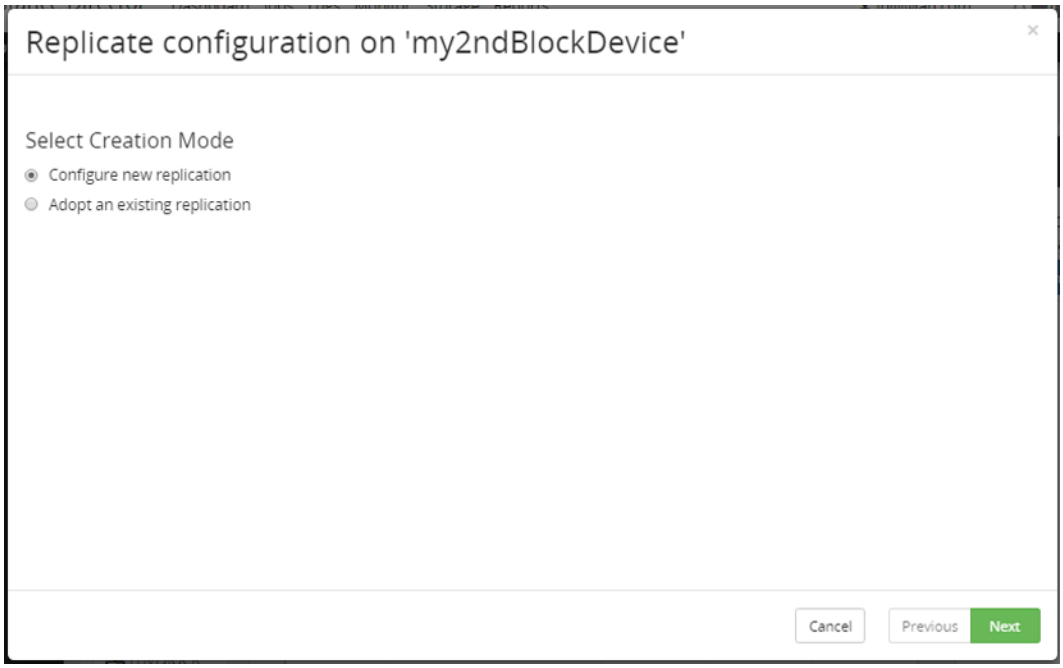


Figure 159 Replicate Configuration Wizard - Configure New or Adopt Existing Replication

Control	Description
Configure new replication	Creates a new replication defined by the user defined parameters. The **Configure New Replication** page of the wizard is displayed next.
Adopt existing replication	Adopts the replication defined by the user defined parameters from the matching pre-existing one on the hardware. The **Adopt Existing Replication** page of the wizard is displayed next.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

Note: To change the Create / Adopt configuration of an existing dataflow it is necessary to deactivate the dataflow, dissociate the replication record, then edit the dataflow to create/adopt and then activate the dataflow.

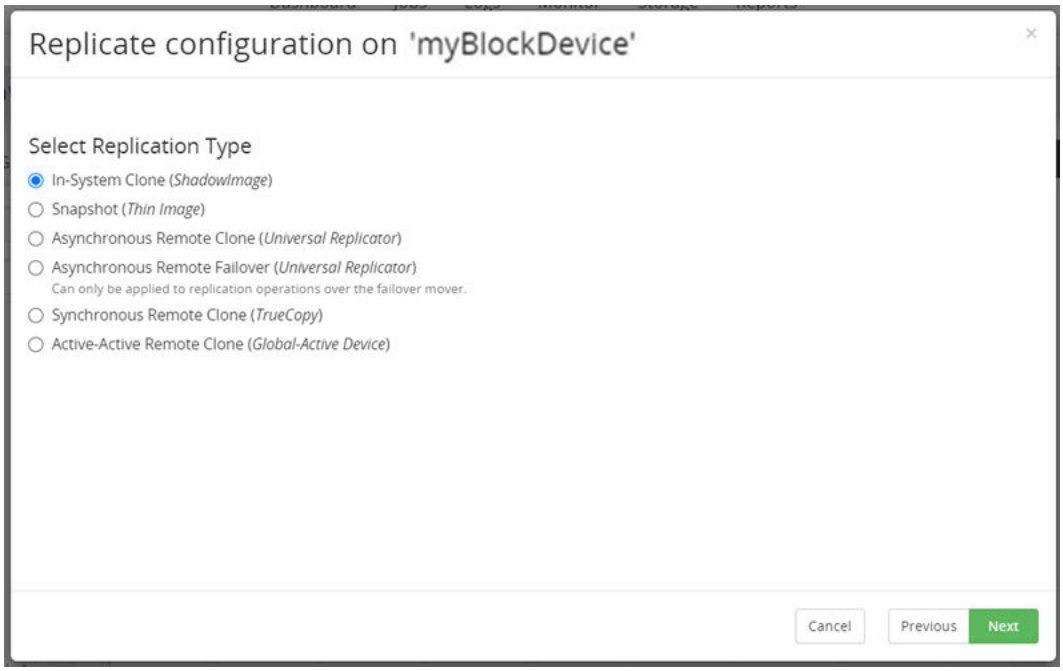


Figure 160 Replicate Configuration Wizard - Configure New Replication

Control	Description
Replication Type	Select the type of replication to create: ▪ In System Clone - opens the Replication Configuration Wizard - In-System Clone (on page 404). ▪ Refreshed Snapshot - opens the Replication Configuration Wizard - Refreshed Snapshot (on page 416). ▪ Asynchronous Remote Clone - opens the Replication Configuration Wizard - Asynchronous Remote Clone (on page 423). ▪ Asynchronous Remote Failover - opens the Replication Configuration Wizard - Asynchronous Remote Failover (on page 436). ▪ Synchronous Remote Clone - opens the Replication Configuration Wizard - Synchronous Remote Clone (on page 443). ▪ Active-Active Remote Clone - opens the Replication Configuration Wizard - Active-Active Remote Clone (on page 454). **Note:** When the target of the replication is a Software Block Device the only supported option is Asynchronous Remote Clone all other replication types are disabled.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.

Control	Description
Next	Takes the user to the next screen in the wizard.

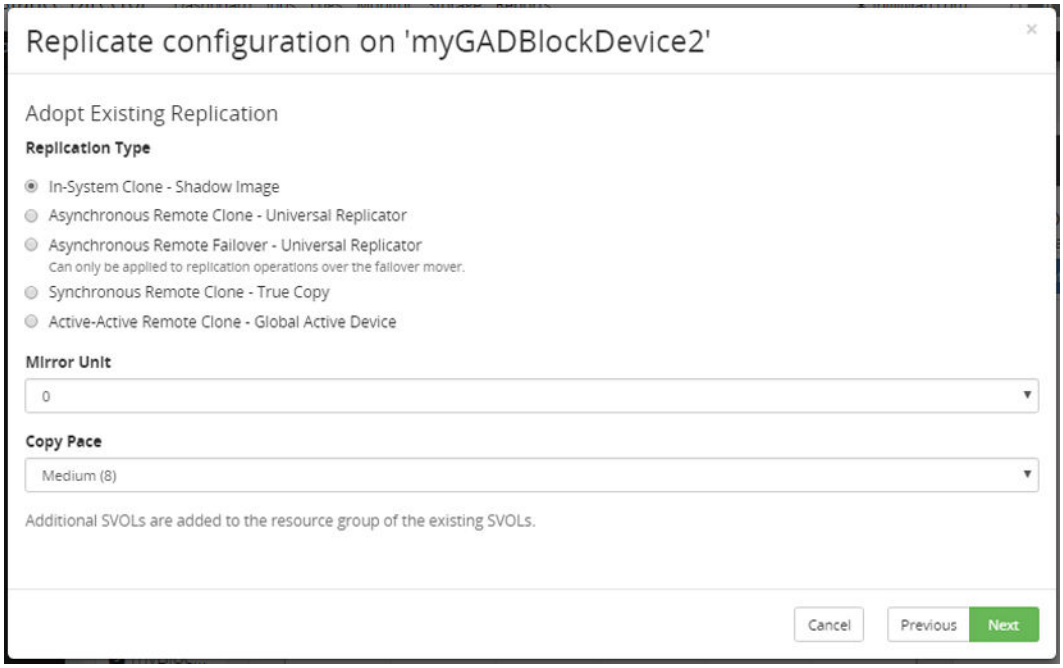


Figure 161 Replicate Configuration Wizard - Adopt Existing Replication

Note: Adoption requires the rules to be activated and the relevant policy to be triggered. Refer to About Hitachi Block replication adoption (on page 87) before using this feature.

Control	Description
Replication Type	Select the type of replication to adopt: ▪ In System Clone - see About ShadowImage replication (on page 70) ▪ Refreshed Snapshot - see About Thin Image and Thin Image Advanced differential and refreshed snapshots (on page 66) ▪ Asynchronous Remote Clone - see About Universal Replicator (on page 75) ▪ Asynchronous Remote Failover - see About three datacentre multi-target with delta (on page 80) ▪ Synchronous Remote Clone - see About TrueCopy replication (on page 74) ▪ Active-Active Remote Clone - see About Global-Active Device replication (on page 76)

Control	Description
Mirror Unit	Identify the mirror unit number of the replication to be adopted. Adoption requires at least one existing pair on the selected mirror unit. **Note:** If the mirror unit of an active replication is changed after initial data flow activation then: 1. S-VOLs and pairing relationships for the replication will be destroyed (or dissociated if previously adopted). 2. The replication will then be recreated (or readopted if previously adopted) on data flow reactivation. A warning is issued by the rules compiler prior to activation.
Copy Pace	Determines how quickly the storage array will be told to copy data for the adopted replication. The array's default is Slow (3), Protector defaults to Medium (8).
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

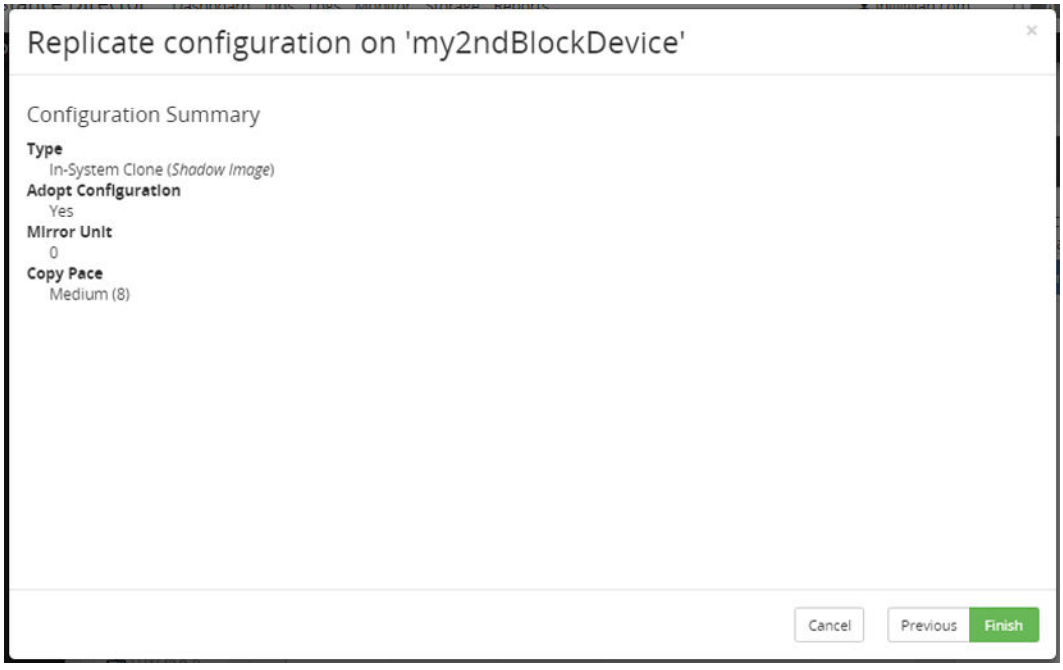


Figure 162 Replicate Configuration Wizard - Adoption Summary

Shows a summary of the replication configuration specified by the user. All other parameters defining the replication are obtained from the hardware when the data flow is activated and triggered.

Replication Configuration Wizard - In-System Clone

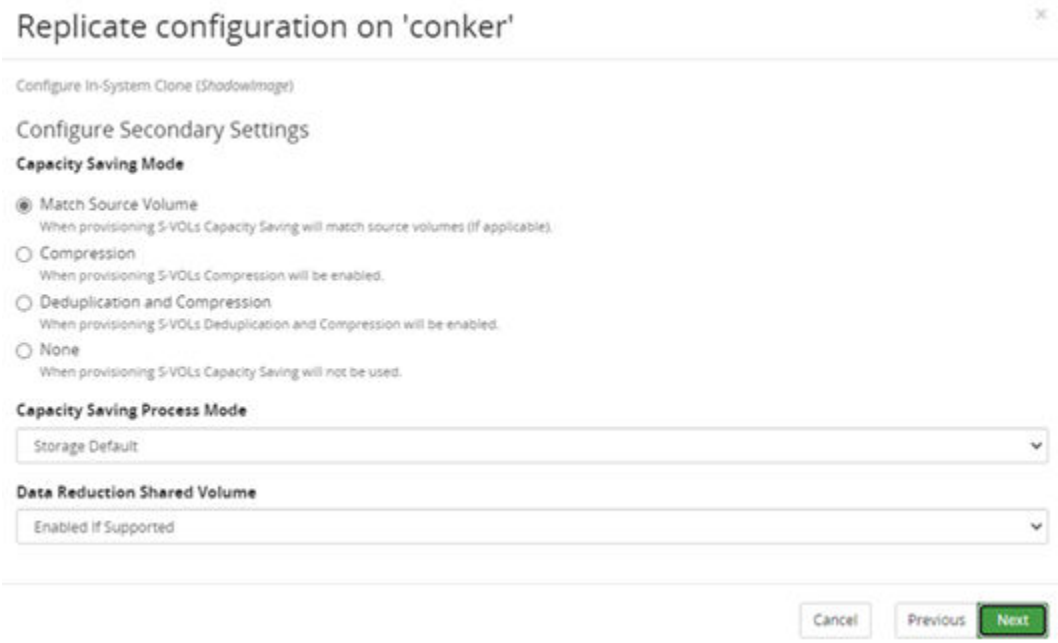


Figure 163 Replicate Configuration Wizard - In System Clone (ShadowImage) - Configure Capacity Savings

Control	Description
Capacity Saving Mode	One of the following options: ▪ Match Source Volumes – When provisioning S-VOLs Capacity Saving will match the settings of the source volumed. ▪ Compression - When provisioning S-VOLs Compression will be enabled, the data compression function utilizes the LZ4 compression algorithm to compress the data. ▪ Deduplication and Compression - When provisioning S-VOLs Deduplication and Compression will be enabled. The data deduplication function deletes duplicate copies of data written to different addresses in the same pool and maintains only a single copy of the data at one address. ▪ None - When provisioning S-VOLs Capacity Saving will not be used

Control	Description
Capacity Saving Process Mode	Only available when a Capacity Saving Mode other than None is selected. Can be one of the following: ▪ Inline - When you apply capacity saving with the inline mode the compression and deduplication processing are performed synchronously for new write data. The inline mode minimizes the pool capacity required to store new write data but can impact I/O performance more than the post-process mode. ▪ Post Process - When you apply capacity saving with the post-process mode the compression and deduplication processing are performed asynchronously for new write data. ▪ Storage Default – match the default option set on the storage array. ▪ Match Source Volume – match the settings of the source volume.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

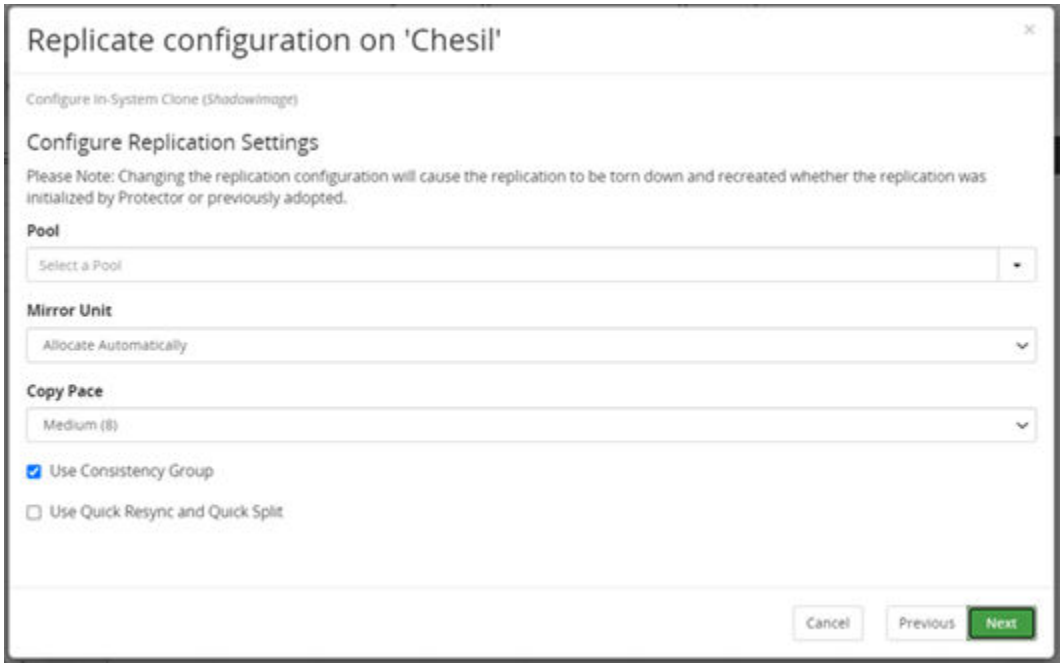


Figure 164 Replicate Configuration Wizard - In System Clone (ShadowImage) - Configure Pool etc.

Control	Description
Pool	Specifies the target storage pool from which replication LDEVs are allocated. Provides a list of available pools giving name and available space. **Note:** All replication types have pool except Asynchronous Remote Failover (Universal Replicator). Select a Dynamic Provisioning Pool for all replication types. **Note:** Dynamic Provisioning Pools must be created using Storage Navigator prior to selecting the Target Storage/Pool in Protector.
Mirror Unit	The mirror unit number for the replication can be set to 0, 1 or 2. Select Allocate Automatically to allow Protector to choose one. **Note:** If the mirror unit of an active replication is changed after initial data flow activation then: 1. S-VOLs and pairing relationships for the replication will be destroyed (or dissociated if previously adopted). 2. The replication will then be recreated (or readopted if previously adopted) on data flow reactivation. A warning is issued by the rules compiler prior to activation.
Copy Pace	Determines how quickly the storage array copies data. The array's default is Slow (3), Protector defaults to Medium (8).
Use Consistency Group	All P-VOLs in a replication are, by default, placed in the same consistency group to ensure consistency of data across all volumes. This option allows this behavior to be disabled.
Quick Resync/ Split	If selected, then *Quick Split* and *Quick Resynch* operations are performed by the storage hardware in the background, so that the secondary is available for reading/writing almost immediately after the replication is paused or resynchronized (depending on downstream data flow operations). If deselected then *Steady Split* and *Normal Copy* operations are performed in the foreground and the secondary is made available only once the operation is completed.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.



Figure 165 Replicate Configuration Wizard - In System Clone - ShadowImage - DRU Options

If Fully Provisioned is selected on the **Provisioning Options** page on this wizard, these options are enabled to specify Data Retention Utility (DRU) protection parameters.

Caution: Protector cannot mount a DRU protected snapshot. However a cascaded snapshot can be created and mounted if the original has Cascade mode enabled.

Control	Description
Protection Type	One of the following can be selected: ▪ None - DRU protection is not applied the snapshot ▪ Host Read Only (wtd) - Prevents hosts writing to the snapshot LDEV. ▪ Full (wtd + svd) - As for wtd, plus the array is prevented from changing the contents of a snapshot LDEV using resync or restore. It also prevents deletion, mapping and unmapping of the snapshot from its LDEV. **Note:** DRU can prevent mounting if the OS rejects a read-only volume (mount via cascade snapshot is unaffected, and therefore recommended).

Control	Description
Duration Of Settings Lock (Days)	**Caution:** DRU protection cannot be removed while the lock is active. Specify a duration in days during which the applied DRU protection cannot be removed. Once this duration has expired the protection is not automatically removed; it must be removed manually.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

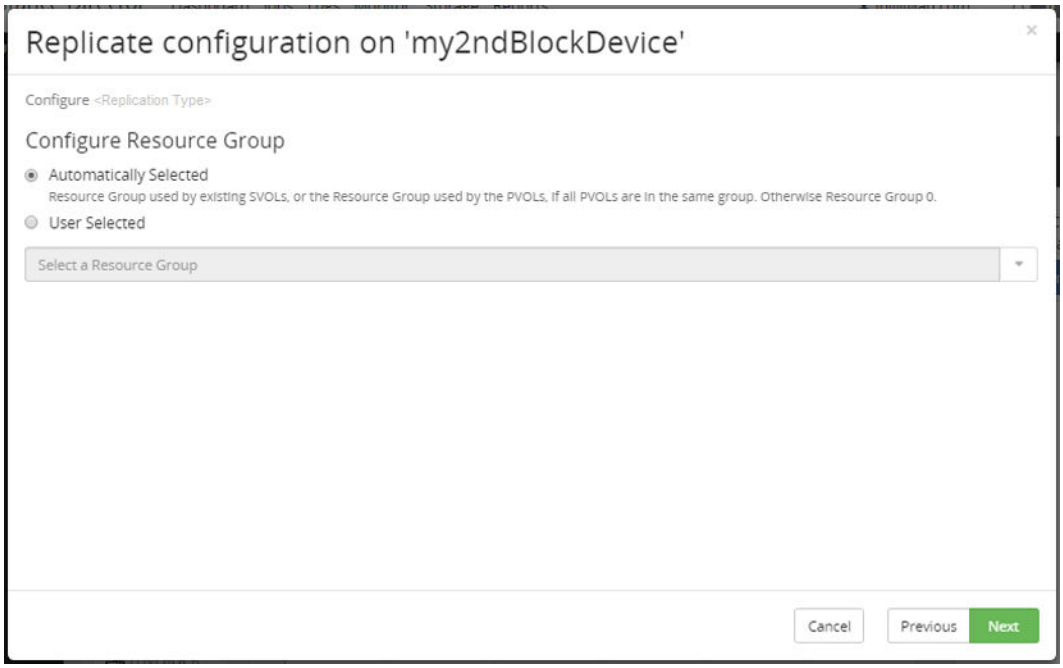


Figure 166 Replicate Configuration Wizard - Configure Resource Group

Control	Description
Configure Resource Group	Specifies the resource group to be used for S-VOLs, in order to support snapshots and replications from VSM volumes (adding volumes to a VSM is performed by adding the volumes to the correct resource group).

Control	Description
	Note: If there are existing S-VOLs, then the resource group used by those will be selected. If the existing S-VOLs are in multiple resource groups or in a resource group that contradicts the user selection, then the operation will fail with an error. This setting should not be modified for existing replications. ▪ Automatically Selected - Allows Protector to automatically select a resource group in the following order of priority: 1. If there are existing S-VOLs, then the resource group used by those will be selected. 2. The resource group used by the P-VOLs, if the replication is in-system and the P-VOLs are all in one resource group. 3. Resource group 0. **Note:** If existing S-VOLs are in multiple resource groups, then the operation will fail with an error. ▪ User Selected - The user specifies the Resource Group.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

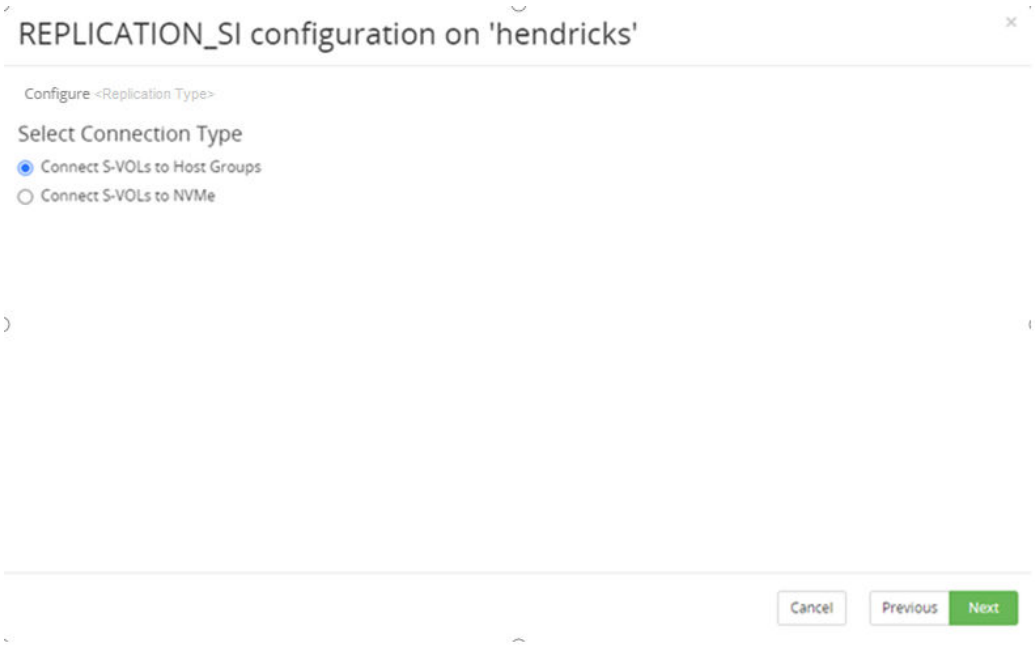
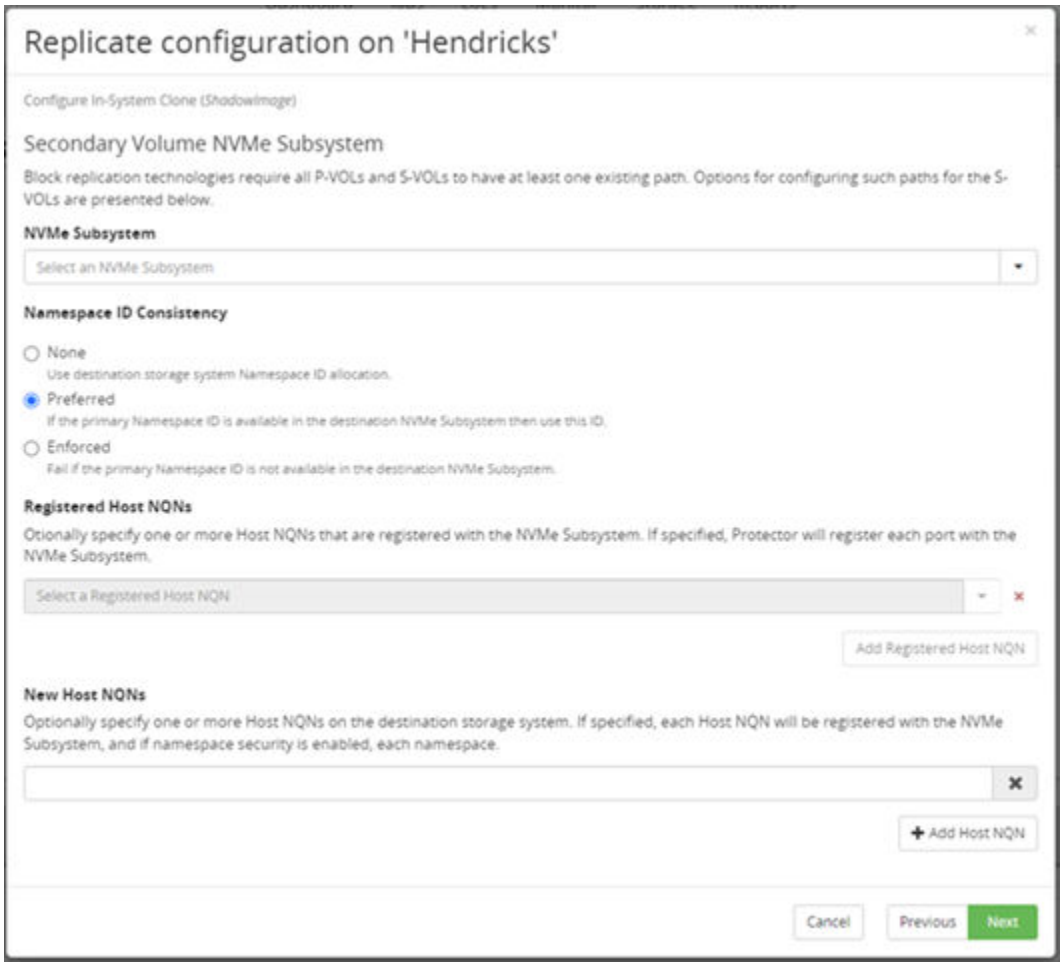


Figure 167 Replicate Configuration Wizard - Select Connection Type

Control	Description
Select protection type	Specifies whether to configure the replication via Host Groups or through an NVMe subsystem.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

If the NVMe option is selected the following screen is displayed:



Note: This screen is not displayed for an adopt dataflow.

Control	Description
NVMe Subsystem	Pre-existing NVMe Subsystem to create namespaces in for each secondary volume that is created.
Namespace ID Consistency	How the namespace ids consistency between primary and secondary volumes should be managed ▪ None – Use the automated storage system namespace id allocation, by default this incrementally counts up for each created namespace in an NVMe subsystem. ▪ None – Use the automated storage system namespace id allocation, by default this incrementally counts up for each created namespace in an NVMe subsystem. ▪ Enforced – March the namespace ids for the primary and secondary volumes and fail if this is not possible.
Registered Host NQNs	Select one or more host NQNs that are registered with the NVMe subsystem. A namespace path for each host nqn and namespace will be created as required if namespace security is enabled.
Add Registered Host NQN	Adds another registered host NQN selection.
New Host NQNs	Enter one or more new host NQNs. The host NQN will be registered with the NVMe subsystem. A namespace path for each host nqn and namespace will be created as required if namespace security is enabled.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

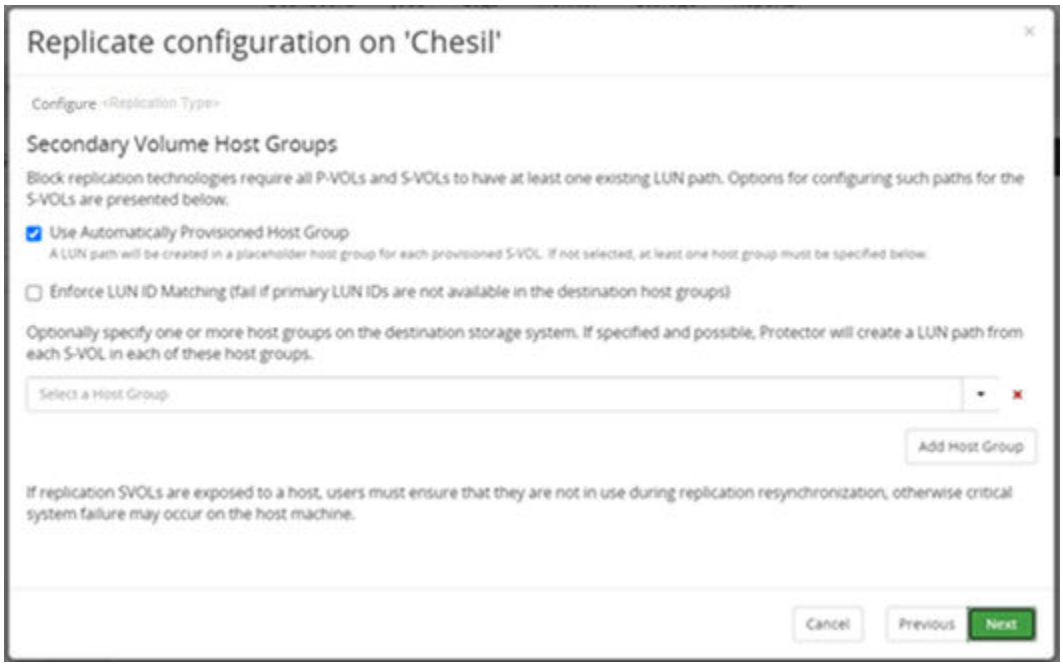


Figure 168 Replicate Configuration Wizard - Secondary Volume Host Groups (In-System Clone)

Caution: If replication S-VOLs are exposed to a host, the user is responsible for ensuring they are not in use during replication resynchronization. Failing to do so may result in a critical failure of the host.

Note: Protector analyses the LUN IDs of the P-VOLs in all host group paths and, if consistent and available, uses these LUN IDs for all S-VOL host group paths either during initial set-up, mounting, or when adding additional LUNs on demand.

To provide a graceful fallback:

- If the P-VOL LUN IDs are not consistent in all host group paths Protector will still use a consistent ID for S-VOL mappings, but these will not necessarily match any of the P-VOL LUN IDs.
- If a P-VOL has a consistent ID in all host group paths, but this LUN ID is not available in all S-VOL host group paths then Protector will choose a different, consistent LUN ID for the S-VOL for all host group paths.
- When not able to match LUN IDs with those used by P-VOLs and/or the S-VOLs, Protector chooses an unused LUN ID. In order to keep LUN ID's compatible with VMware and Hyper-V when selecting LUN ID's Protector will first attempt to select ID's at or below 255. If this is not it will then attempt at or below 1024, followed by at or below 2048. Finally, at or below the array maximum. A warning will be displayed when an ID can not be found in a range.

Some systems expect LUN ID 0 to be used only as a boot volume. Protector will therefore only use LUN ID 0 for the S-VOL host group paths if the P-VOL LUN ID is 0.

Control	Description
Use Protector Provisioned Host Group	Protector will create a LUN path from each S-VOL it provisions in a placeholder host group. If this option is not selected, at least one Secondary Host Group must be specified below.
Enforce LUN ID Matching	For environments where LUN ID consistency is mandatory, selecting this option will cause the replication to fail if: ▪ The P-VOL does not have a consistent LUN ID in all host group paths. ▪ The P-VOL LUN ID is not available in all S-VOL host group paths.
Secondary Host Groups	Specify zero or more host groups that Protector will configure to provide access to the S-VOL(s) when configuring replication scenarios. If no host groups are specified here then Protector will place the S-VOL(s) in it's dummy host group. Click the Add Host Group button to insert another Host Group selection control. Click the Remove button next to a Host Group selection control to delete it. **Note:** If a LUN path to be created already exists, Protector will not attempt to add it again, or to change its ID. The specified host groups must be in the same resource group as the secondary volumes. For GAD replications, if the host group names and port IDs match between primary and secondary storage nodes, Protector will attempt to match the LUN IDs used for the S-VOLs with those of the respective P-VOLs. If this cannot be achieved then a warning will be logged and the next available LUN ID will be used.

Control	Description
	Tip: Use this option to specify host groups that enable access to S-VOL(s) when configuring GAD cross-path and multi-path and other replication scenarios where the S-VOL(s) will need to be accessed (e.g. during failover).
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

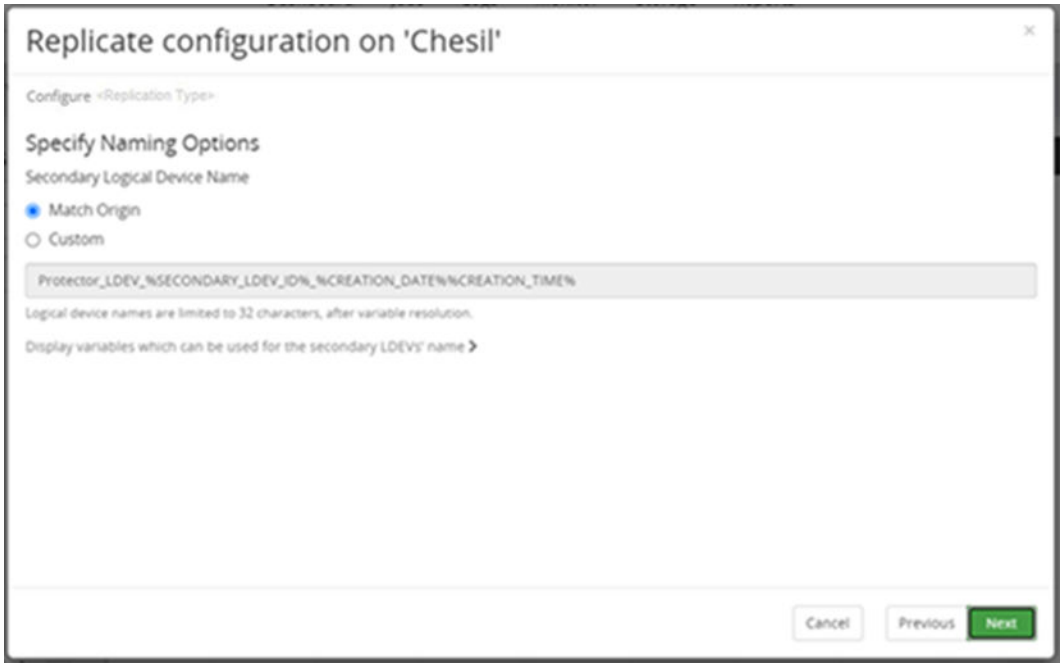


Figure 169 Replicate Configuration Wizard - Naming Options

Control	Description
Secondary Logical Device Name	Specifies how S-VOLs will be named: ▪ Match Origin - The S-VOL will be given the same name as that used for the origin P-VOL (i.e. the left-most volume in the data flow). ▪ Custom - The S-VOL will be named using the naming rule provided. The naming rule can consist of literal strings and/or one or more substitution variables listed. Click Display variables which can be used for the secondary LDEVs' name to view the available substitution variables: • `%ORIGIN_SERIAL%` - S/N of leftmost array in data flow. E.g. output string: `210613` • `%ORIGIN_LDEV_ID%` - ID of leftmost LDEV in data flow. E.g. output string: `00:3A:98` • `%ORIGIN_LDEV_NAME%` - name of leftmost LDEV in data flow. • `%PRIMARY_SERIAL%` - S/N of primary array in this operation. E.g. output string: `442302` • `%PRIMARY_LDEV_ID%` - ID of primary LDEV in this operation. E.g. output string: `00:4C:EB` • `%PRIMARY_LDEV_NAME%` - name of primary LDEV in this operation. • `%SECONDARY_SERIAL%` - S/N of secondary array in this operation. E.g. output string: `356323` • `%SECONDARY_LDEV_ID%` - ID of secondary LDEV in this operation. E.g. output string: `01:F4:35` • `%CREATION_DATE%` - date secondary LDEV was created by this operation. E.g. output string: `20180427` • `%CREATION_TIME%` - time secondary LDEV was created by this operation. 1130
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

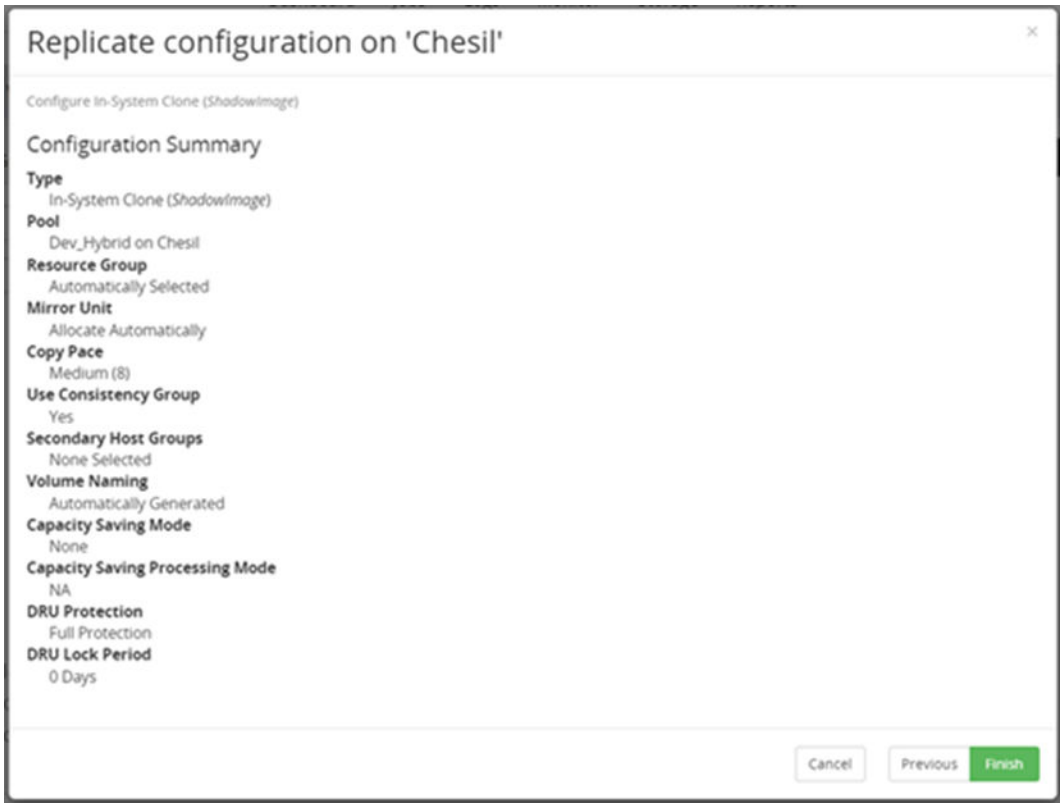


Figure 170 Replicate Configuration Wizard - In System Clone (ShadowImage) - Summary

Shows a summary of the replication configuration specified by the user.

Replication Configuration Wizard - Refreshed Snapshot

Note: Refreshed Thin Image must be used with a Batch mover on the data flow.

The replication is created as a single, differential, in-system snapshot using Thin Image. The snapshot is refreshed on each batch resync rather than creating a new snapshot for each resync.

When an RTI data flow is deactivated, the refreshed snapshot is deleted.

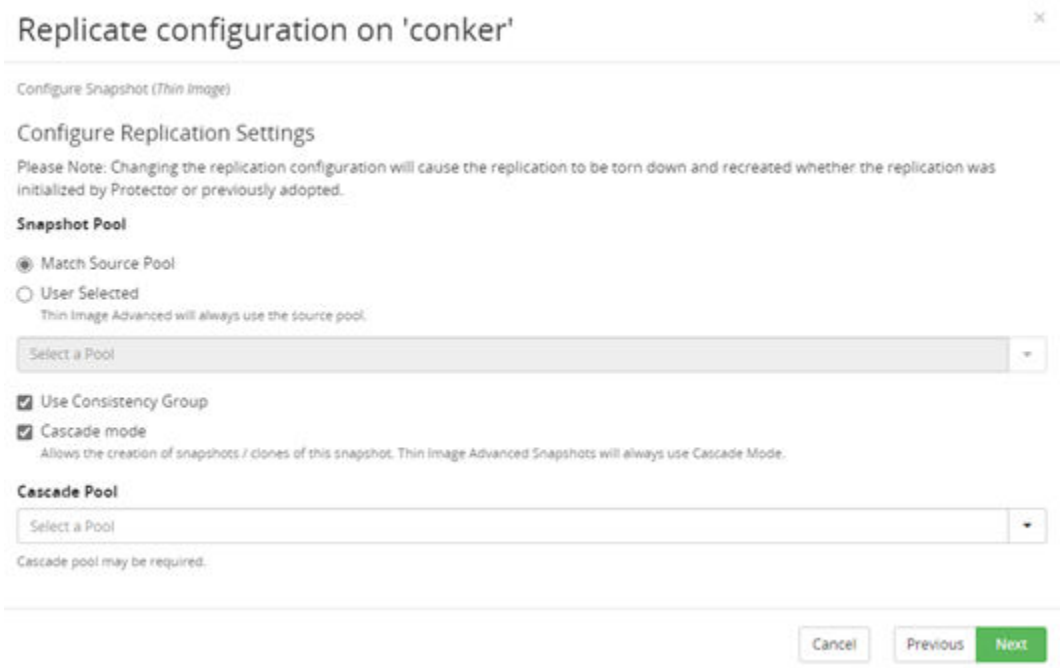


Figure 171 Replicate Configuration Wizard - Refreshed Snapshot (Thin Image) - Configure Pool etc.

Control	Description
Snapshot Pool	Specifies the target storage pool from one of the following options to which replication LDEVs are allocated. ▪ Match Source Pool – Allocates the snapshot in the same pool as the source volumes. ▪ User Selected – Select a Thin Image Pool or a Dynamic Provisioning Pool. **Note:** Thin Image Advanced Snapshots which are RoW snapshots will match source pool. Provides a list of available pools giving name and available space. **Note:** All replication types have pool except Asynchronous Remote Failover (Universal Replicator). **Note:** Thin Image and Dynamic Provisioning Pools must be created using Storage Navigator prior to selecting the Target Storage/Pool in Protector.

Control	Description
Cascade Mode	The snapshot can be cascaded. This allows Protector to mount either the original snapshot or a duplicate of the original. Mounting a duplicate enables modifications to be made without affecting the original snapshot. Refer to About cascaded Thin Image and Thin Image Advanced snapshots (on page 69) for further guidance on configuring and mounting cascade mode snapshots.
Use Consistency Group	All P-VOLs in a replication are, by default, placed in the same consistency group to ensure consistency of data across all volumes. This option allows this behavior to be disabled.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

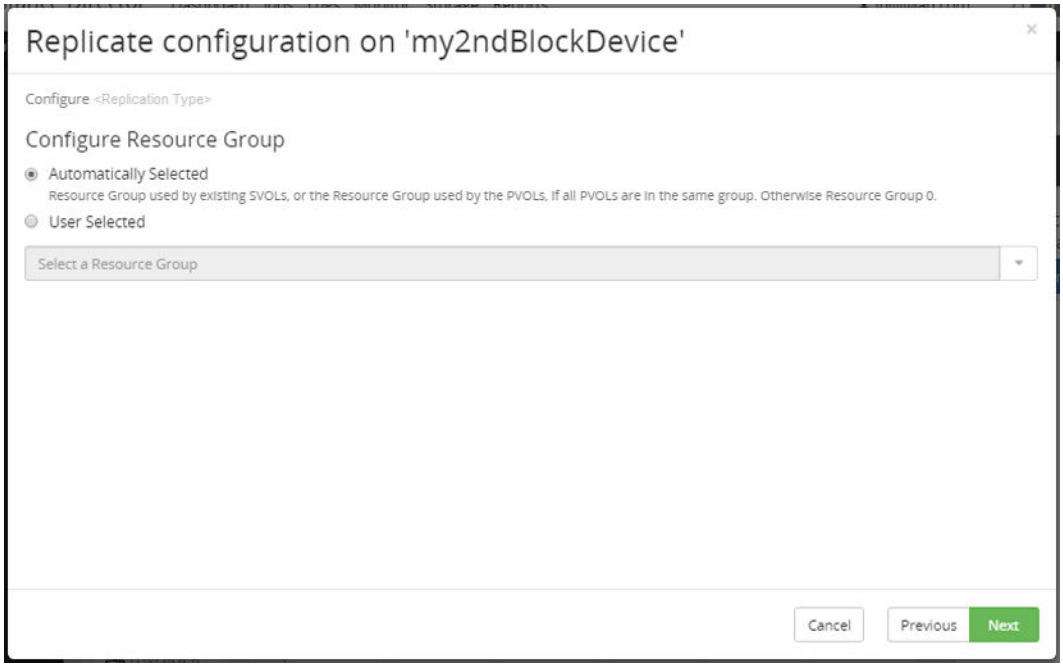


Figure 172 Replicate Configuration Wizard - Configure Resource Group

Control	Description
Configure Resource Group	Specifies the resource group to be used for S-VOLs, in order to support snapshots and replications from VSM volumes (adding volumes to a VSM is performed by adding the volumes to the correct resource group).

Control	Description
	Note: If there are existing S-VOLs, then the resource group used by those will be selected. If the existing S-VOLs are in multiple resource groups or in a resource group that contradicts the user selection, then the operation will fail with an error. This setting should not be modified for existing replications. ▪ Automatically Selected - Allows Protector to automatically select a resource group in the following order of priority: 1. If there are existing S-VOLs, then the resource group used by those will be selected. 2. The resource group used by the P-VOLs, if the replication is in-system and the P-VOLs are all in one resource group. 3. Resource group 0. **Note:** If existing S-VOLs are in multiple resource groups, then the operation will fail with an error. ▪ User Selected - The user specifies the Resource Group.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

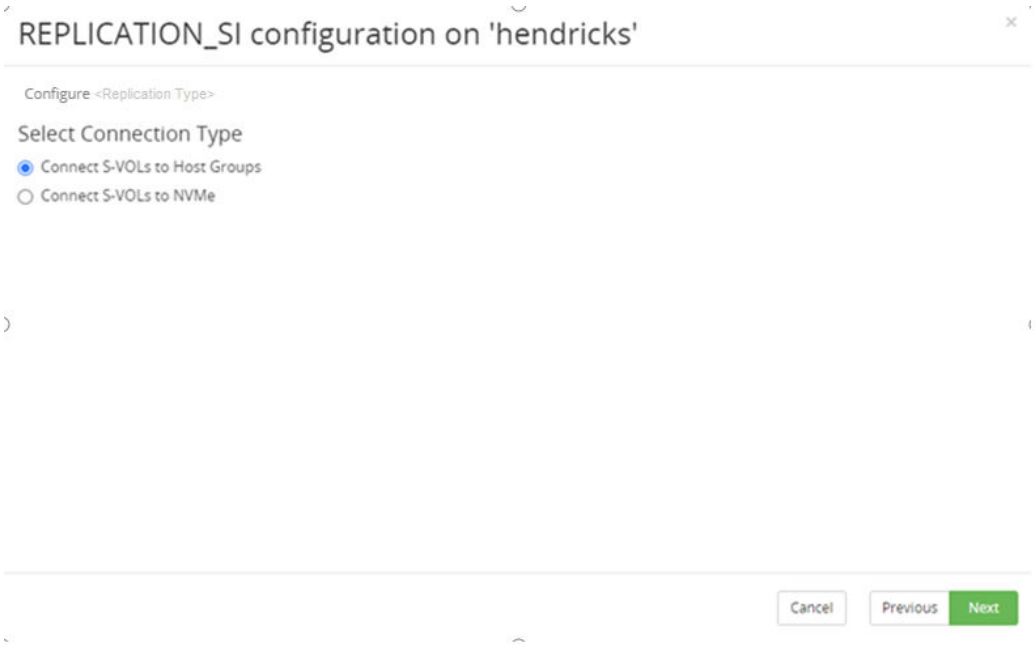


Figure 173 Replicate Configuration Wizard - Select Connection Type

Control	Description
Select protection type	Specifies whether to configure the replication via Host Groups or through an NVMe subsystem.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

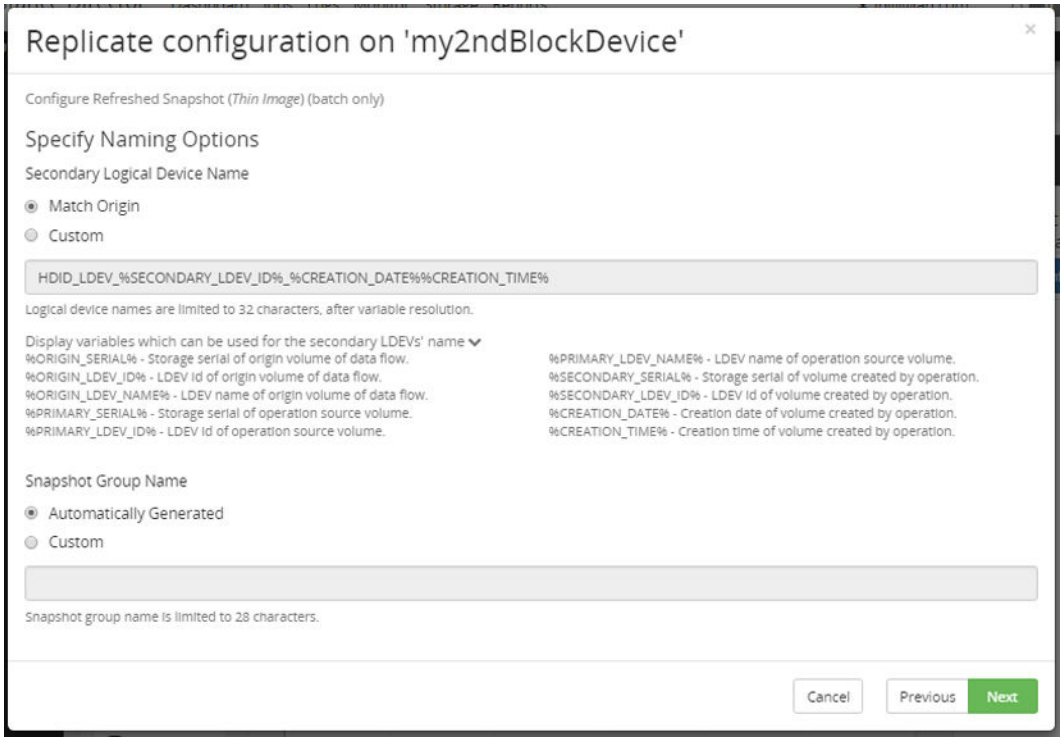


Figure 174 Replicate Configuration Wizard - Refreshed Snapshot (Thin Image) - Naming Options

Control	Description
Secondary Logical Device Name	Specifies how S-VOLs will be named: ▪ Match Origin - The S-VOL will be given the same name as that used for the origin P-VOL (i.e. the left-most volume in the data flow). ▪ Custom - The S-VOL will be named using the naming rule provided. The naming rule can consist of literal strings and/or one or more substitution variables listed. Click Display variables which can be used for the secondary LDEVs' name to view the available substitution variables: • `%ORIGIN_SERIAL%` - S/N of leftmost array in data flow. E.g. output string: `210613` • `%ORIGIN_LDEV_ID%` - ID of leftmost LDEV in data flow. E.g. output string: `00:3A:98` • `%ORIGIN_LDEV_NAME%` - name of leftmost LDEV in data flow. • `%PRIMARY_SERIAL%` - S/N of primary array in this operation. E.g. output string: `442302` • `%PRIMARY_LDEV_ID%` - ID of primary LDEV in this operation. E.g. output string: `00:4C:EB` • `%PRIMARY_LDEV_NAME%` - name of primary LDEV in this operation. • `%SECONDARY_SERIAL%` - S/N of secondary array in this operation. E.g. output string: `356323` • `%SECONDARY_LDEV_ID%` - ID of secondary LDEV in this operation. E.g. output string: `01:F4:35` • `%CREATION_DATE%` - date secondary LDEV was created by this operation. E.g. output string: `20180427` • `%CREATION_TIME%` - time secondary LDEV was created by this operation. 1130
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

Control	Description
Snapshot Group Name	Specifies how the snapshot group will be named. ▪ Automatically Generated - The snapshot group name is generated by Protector based on the rules context ID and policy name. ▪ Custom - The snapshot group is named using the string provided (limited to 28 characters). An '@' separator followed by a unique ID is then automatically appended to this name. The unique ID is composed of 3 base 36 characters and is required to enable Protector to manage the groups.

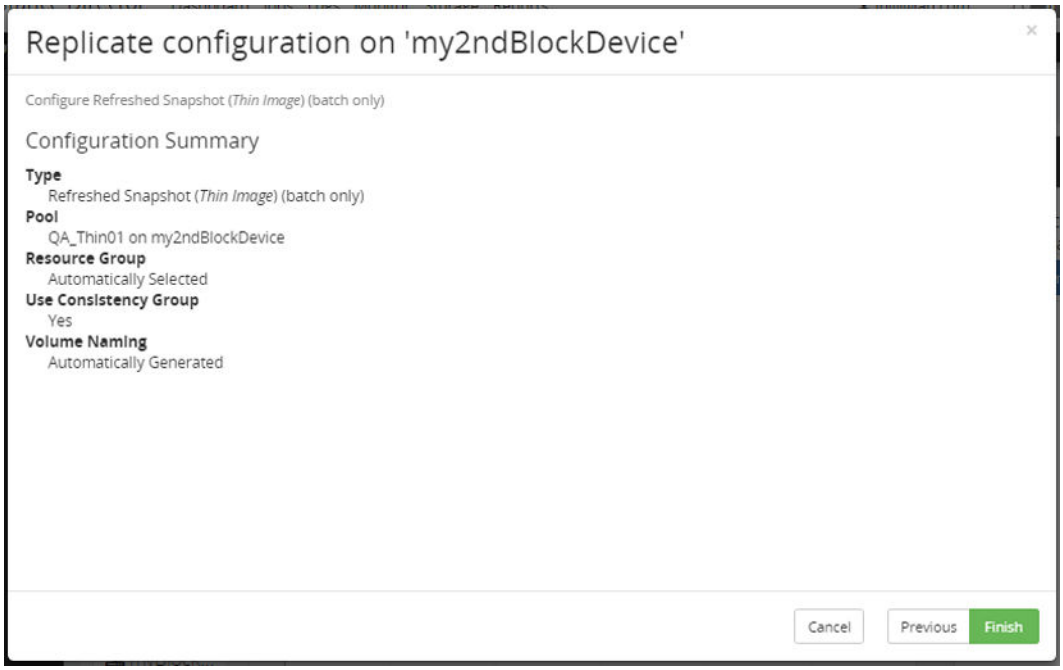


Figure 175 Replicate Configuration Wizard - Refreshed Snapshot (Thin Image) - Summary

Shows a summary of the replication configuration specified by the user.

Replication Configuration Wizard - Asynchronous Remote Clone

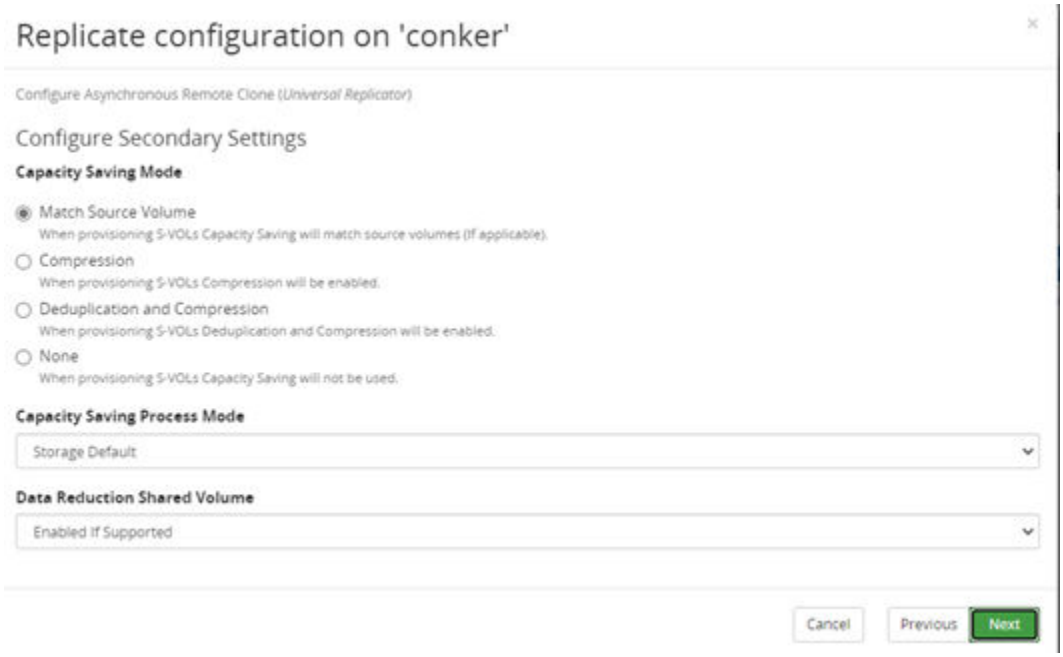


Figure 176 Replicate Configuration Wizard - Asynchronous Remote Clone (Universal Replicator) - Configure Capacity Savings

Control	Description
Capacity Saving Mode	One of the following options: ▪ Match Source Volumes – When provisioning S-VOLs Capacity Saving will match the settings of the source volumed. ▪ Compression - When provisioning S-VOLs Compression will be enabled, the data compression function utilizes the LZ4 compression algorithm to compress the data. ▪ Deduplication and Compression - When provisioning S-VOLs Deduplication and Compression will be enabled. The data deduplication function deletes duplicate copies of data written to different addresses in the same pool and maintains only a single copy of the data at one address. ▪ None - When provisioning S-VOLs Capacity Saving will not be used

Control	Description
Capacity Saving Process Mode	Only available when a Capacity Saving Mode other than None is selected. Can be one of the following: ▪ Inline - When you apply capacity saving with the inline mode the compression and deduplication processing are performed synchronously for new write data. The inline mode minimizes the pool capacity required to store new write data but can impact I/O performance more than the post-process mode. ▪ Post Process - When you apply capacity saving with the post-process mode the compression and deduplication processing are performed asynchronously for new write data. ▪ Storage Default – match the default option set on the storage array. ▪ Match Source Volume – match the settings of the source volume.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

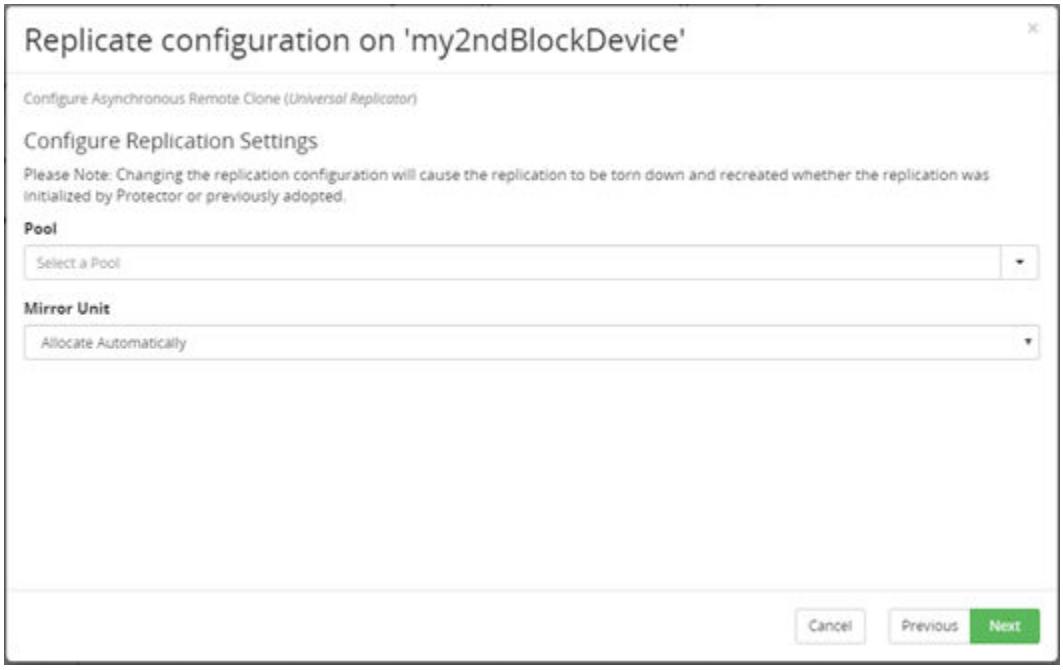


Figure 177 Replicate Configuration Wizard - Asynchronous Remote Clone (Universal Replicator) - Configure Pool and Mirror Unit

Control	Description
Pool	Specifies the target storage pool from which replication LDEVs are allocated. Provides a list of available pools giving name and available space. **Note:** All replication types have pool except Asynchronous Remote Failover (Universal Replicator). Select a Dynamic Provisioning Pool. **Note:** Dynamic Provisioning Pools must be created using Storage Navigator prior to selecting the Target Storage/Pool in Protector.
Mirror Unit	The mirror unit number for the replication can be set to 0, h1, h2 or h3. **Note:** If the replication target is a software defined storage array then the only supported Mirror Unit is 0, as such this setting will default to 0 and can not be changed. Select Allocate Automatically to allow Protector to choose one. **Note:** If the mirror unit of an active replication is changed after initial data flow activation then: 1. S-VOLs and pairing relationships for the replication will be destroyed (or dissociated if previously adopted). 2. The replication will then be recreated (or readopted if previously adopted) on data flow reactivation. A warning is issued by the rules compiler prior to activation.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

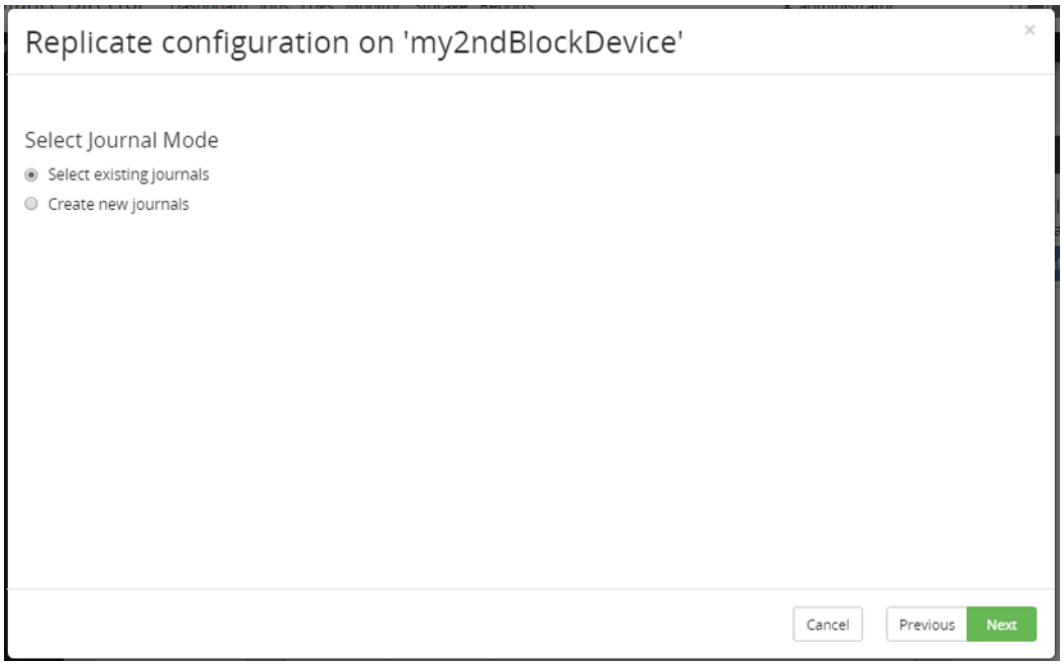


Figure 178 Replicate Configuration Wizard - Asynchronous Remote Clone (Universal Replicator) - Select Journal Mode

Control	Description
Select existing journals	Select this option to use journals that already exist on the source and destination storage arrays. The **Select existing journals** wizard page is displayed next.
Create new journals	Select this option to have Protector configure new journals on the source and destination storage arrays. The **Create journals** wizard page is displayed next.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

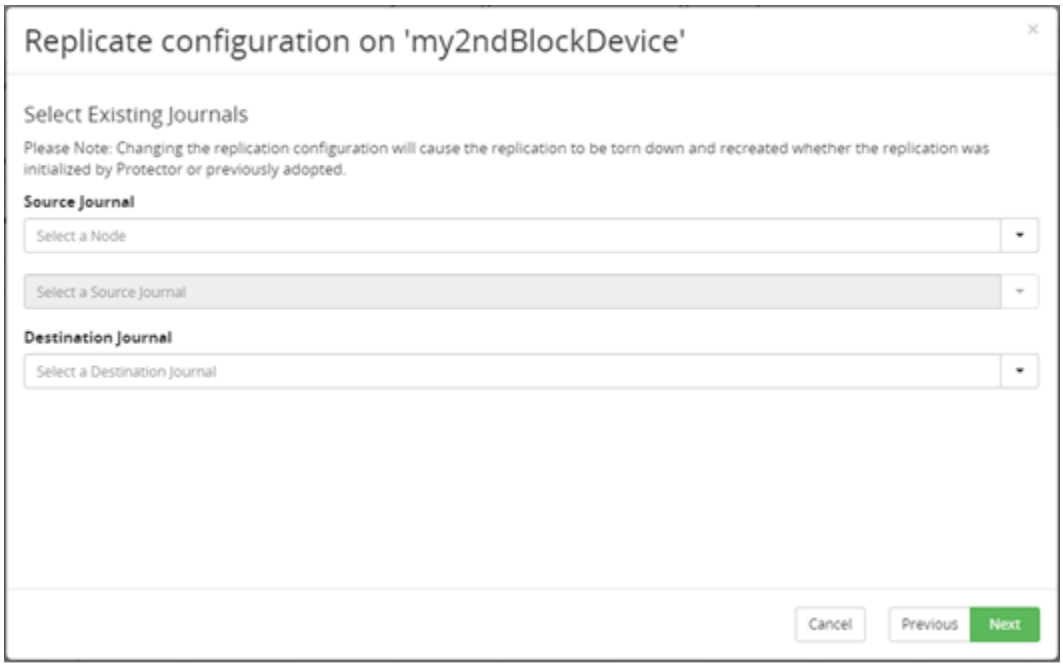


Figure 179 Replicate Configuration Wizard - Asynchronous Remote Clone (Universal Replicator) - Select existing journals

Note: The journals must be unique to each operation and policy in the data flow. The journal should also be used exclusively by Protector.

Control	Description
Source Journal	Specifies the node and journal on the source side of the replication.
Destination Journal	Specifies the journal on the destination side of the replication.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

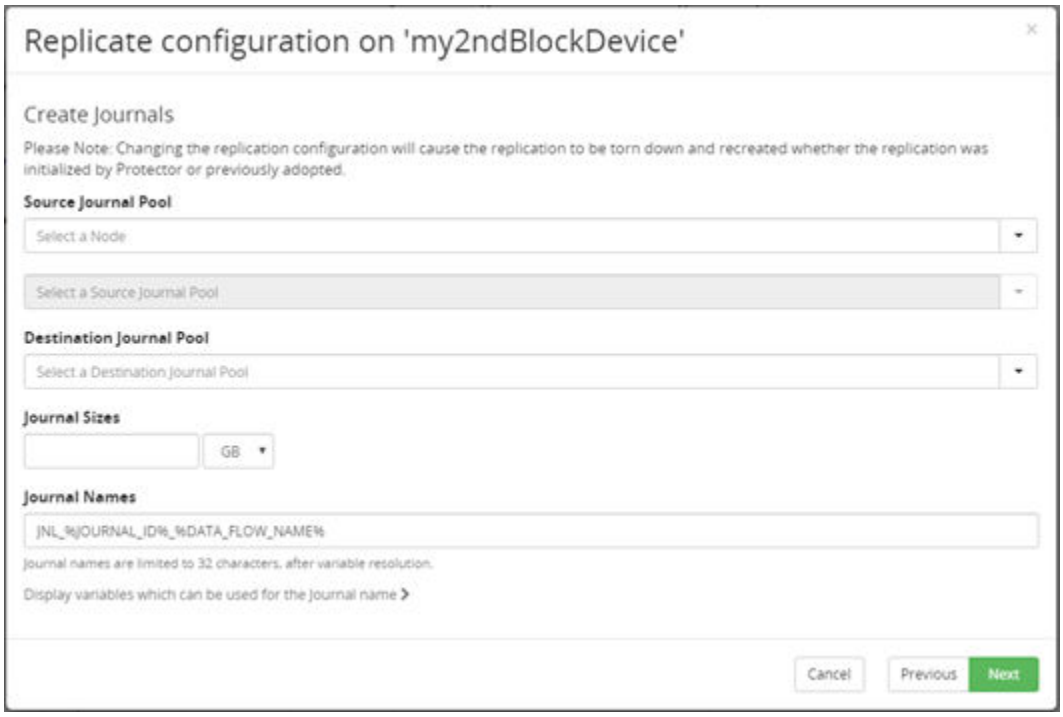


Figure 180 Replicate Configuration Wizard - Asynchronous Remote Clone (Universal Replicator) - Create journals

Note: The journals must be unique to each operation and policy in the data flow. The journal should also be used exclusively by Protector.

Control	Description
Source Journal Pool	Specifies the node and pool where the source side journal will be created.
Destination Journal Pool	Specifies the pool where the destination side journal will be created.
Journal Sizes	Specifies the size of the source and destination journal. **Note:** Journal sizing is based on data change rate and link bandwidth. Refer to your storage array documentation for details.
Journal Names	The journals will be named using the naming rule provided. The naming rule can consist of literal strings and/or one or more substitution variables listed. Click Display variables which can be used for the journal name to view the available substitution variables.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.

Control	Description
Next	Takes the user to the next screen in the wizard.



Figure 181 Replicate Configuration Wizard - Asynchronous Remote Clone - Select Remote Path Group

Control	Description
Select Remote Path Group	Specifies the Remote Path Group to be used for the replication. ▪ Automatically Selected - Allows Protector to automatically select a Remote Path Group **Note:** For GAD it is recommended the user specifies a group to avoid sharing with other replications. ▪ User Selected - The user specifies the Remote Path Group. **WARNING:** You cannot specify the "User Selected" option and select a path group with an ID of 0. To use a path group with an ID of 0, specify the "Automatically Selected" option. The path group with the lowest ID will be selected (which will be ID 0, if a path group with that ID exists). Remote path groups are listed in the format: Path Group Id: 0x51 Port Mappings: 5E <-> 3E

Control	Description
	The arrow depicts the direction of the path, either left to right or bidirectional. The arrows can also have a line through them depicting the path is currently broken. Only remote path groups that are suitable for the replication are displayed, for example for GAD only bidirectional path groups are listed.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

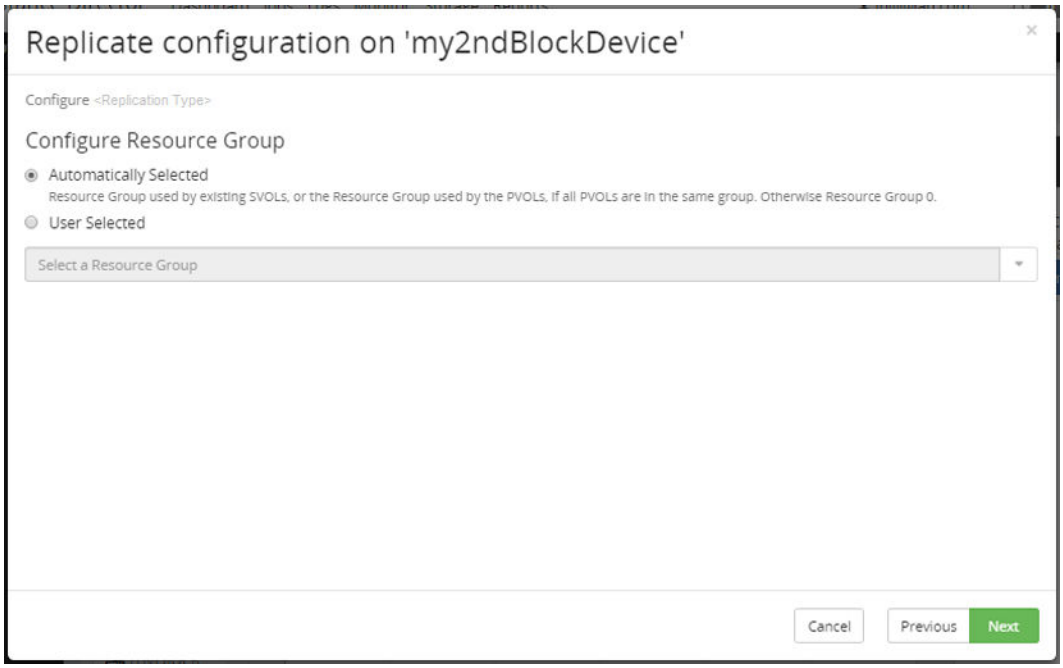


Figure 182 Replicate Configuration Wizard - Configure Resource Group

Control	Description
Configure Resource Group	Specifies the resource group to be used for S-VOLs, in order to support snapshots and replications from VSM volumes (adding volumes to a VSM is performed by adding the volumes to the correct resource group).

Control	Description
	Note: If there are existing S-VOLs, then the resource group used by those will be selected. If the existing S-VOLs are in multiple resource groups or in a resource group that contradicts the user selection, then the operation will fail with an error. This setting should not be modified for existing replications. ▪ Automatically Selected - Allows Protector to automatically select a resource group in the following order of priority: 1. If there are existing S-VOLs, then the resource group used by those will be selected. 2. The resource group used by the P-VOLs, if the replication is in-system and the P-VOLs are all in one resource group. 3. Resource group 0. **Note:** If existing S-VOLs are in multiple resource groups, then the operation will fail with an error. ▪ User Selected - The user specifies the Resource Group.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

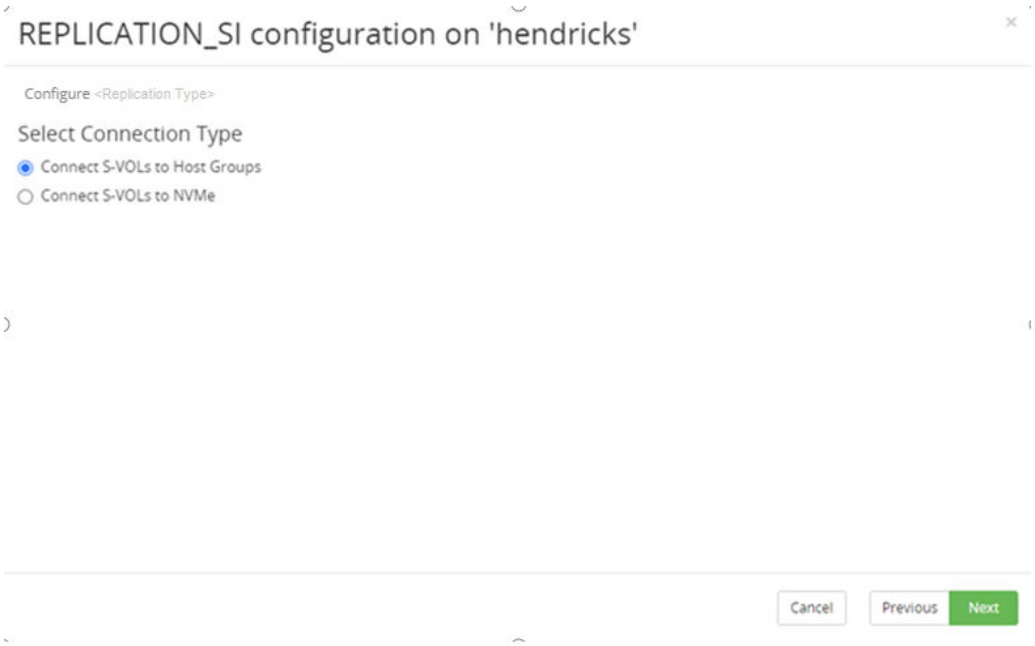


Figure 183 Replicate Configuration Wizard - Select Connection Type

Control	Description
Select protection type	Specifies whether to configure the replication via Host Groups or through an NVMe subsystem.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

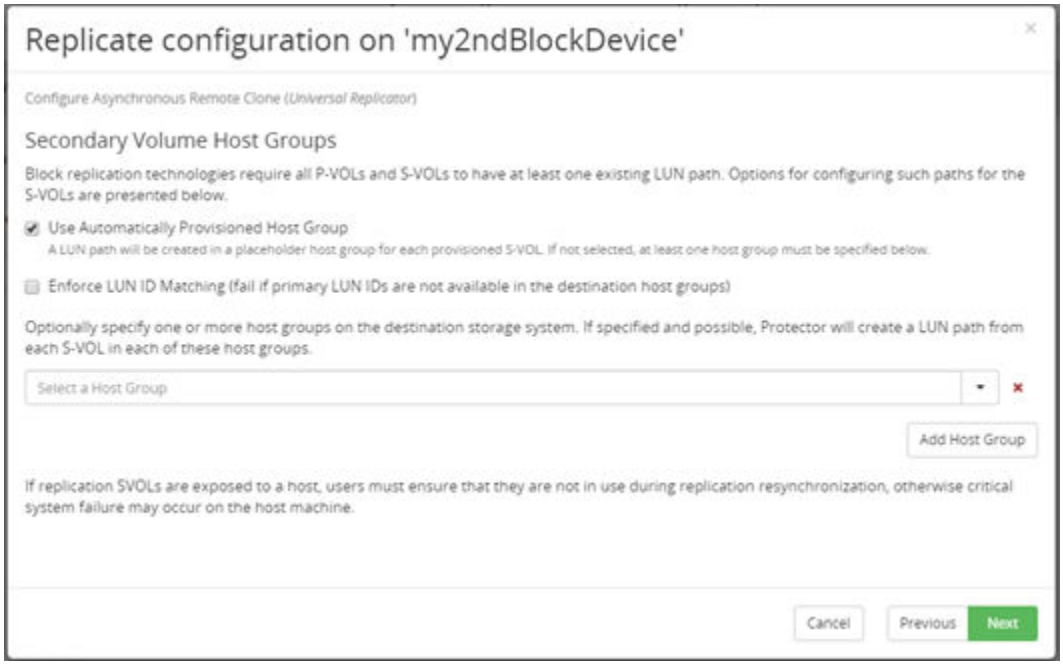


Figure 184 Replicate Configuration Wizard - Secondary Host Groups (Async. RC)

Caution: If replication S-VOLs are exposed to a host, the user is responsible for ensuring they are not in use during replication resynchronization. Failing to do so may result in a critical failure of the host.

Note: Protector analyses the LUN IDs of the P-VOLs in all host group paths and, if consistent and available, uses these LUN IDs for all S-VOL host group paths either during initial set-up, mounting, or when adding additional LUNs on demand.

To provide a graceful fallback:

- If the P-VOL LUN IDs are not consistent in all host group paths Protector will still use a consistent ID for S-VOL mappings, but these will not necessarily match any of the P-VOL LUN IDs.
- If a P-VOL has a consistent ID in all host group paths, but this LUN ID is not available in all S-VOL host group paths then Protector will choose a different, consistent LUN ID for the S-VOL for all host group paths.
- When not able to match LUN IDs with those used by P-VOLs and/or the S-VOLs, Protector chooses an unused LUN ID. In order to keep LUN ID's compatible with VMware and Hyper-V when selecting LUN ID's Protector will first attempt to select ID's at or below 255. If this is not it will then attempt at or below 1024, followed by at or below 2048. Finally, at or below the array maximum. A warning will be displayed when an ID can not be found in a range.

Some systems expect LUN ID 0 to be used only as a boot volume. Protector will therefore only use LUN ID 0 for the S-VOL host group paths if the P-VOL LUN ID is 0.

Control	Description
Use Protector Provisioned Host Group	Protector will create a LUN path from each S-VOL it provisions in a placeholder host group. If this option is not selected, at least one Secondary Host Group must be specified below.
Enforce LUN ID Matching	For environments where LUN ID consistency is mandatory, selecting this option will cause the replication to fail if: ▪ The P-VOL does not have a consistent LUN ID in all host group paths. ▪ The P-VOL LUN ID is not available in all S-VOL host group paths.
Secondary Host Groups	Specify zero or more host groups that Protector will configure to provide access to the S-VOL(s) when configuring replication scenarios. If no host groups are specified here then Protector will place the S-VOL(s) in it's dummy host group. Click the Add Host Group button to insert another Host Group selection control. Click the Remove button next to a Host Group selection control to delete it.

Control	Description
	Note: If a LUN path to be created already exists, Protector will not attempt to add it again, or to change its ID. The specified host groups must be in the same resource group as the secondary volumes. For GAD replications, if the host group names and port IDs match between primary and secondary storage nodes, Protector will attempt to match the LUN IDs used for the S-VOLs with those of the respective P-VOLs. If this cannot be achieved then a warning will be logged and the next available LUN ID will be used. **Tip:** Use this option to specify host groups that enable access to S-VOL(s) when configuring GAD cross-path and multi-path and other replication scenarios where the S-VOL(s) will need to be accessed (e.g. during failover).
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

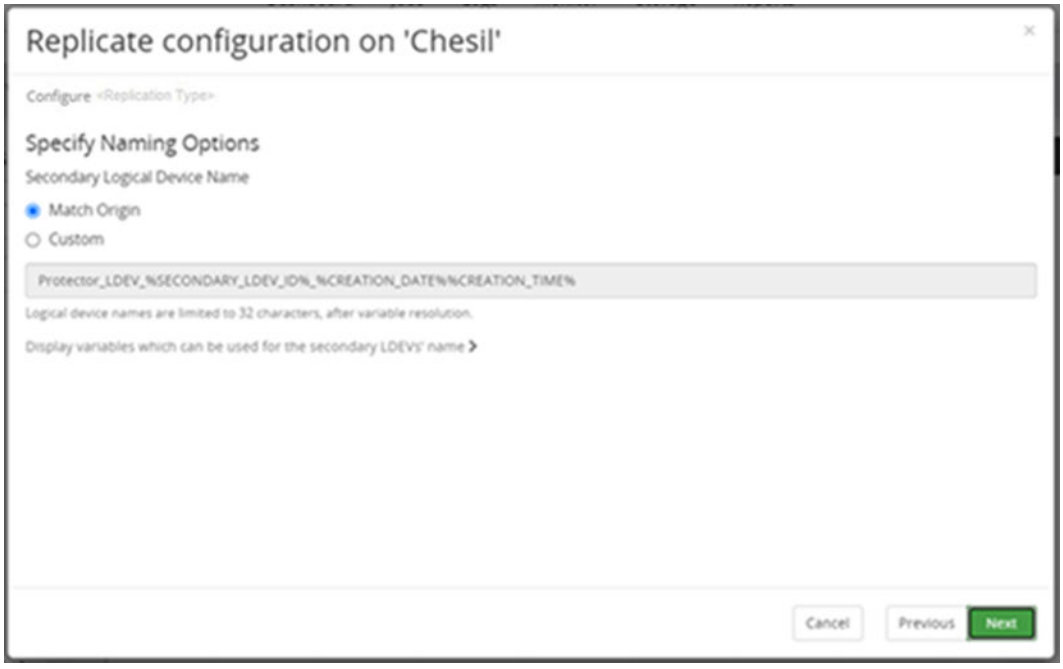


Figure 185 Replicate Configuration Wizard - Naming Options

Control	Description
Secondary Logical Device Name	Specifies how S-VOLs will be named: ▪ Match Origin - The S-VOL will be given the same name as that used for the origin P-VOL (i.e. the left-most volume in the data flow). ▪ Custom - The S-VOL will be named using the naming rule provided. The naming rule can consist of literal strings and/or one or more substitution variables listed. Click Display variables which can be used for the secondary LDEVs' name to view the available substitution variables: • `%ORIGIN_SERIAL%` - S/N of leftmost array in data flow. E.g. output string: `210613` • `%ORIGIN_LDEV_ID%` - ID of leftmost LDEV in data flow. E.g. output string: `00:3A:98` • `%ORIGIN_LDEV_NAME%` - name of leftmost LDEV in data flow. • `%PRIMARY_SERIAL%` - S/N of primary array in this operation. E.g. output string: `442302` • `%PRIMARY_LDEV_ID%` - ID of primary LDEV in this operation. E.g. output string: `00:4C:EB` • `%PRIMARY_LDEV_NAME%` - name of primary LDEV in this operation. • `%SECONDARY_SERIAL%` - S/N of secondary array in this operation. E.g. output string: `356323` • `%SECONDARY_LDEV_ID%` - ID of secondary LDEV in this operation. E.g. output string: `01:F4:35` • `%CREATION_DATE%` - date secondary LDEV was created by this operation. E.g. output string: `20180427` • `%CREATION_TIME%` - time secondary LDEV was created by this operation. 1130
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

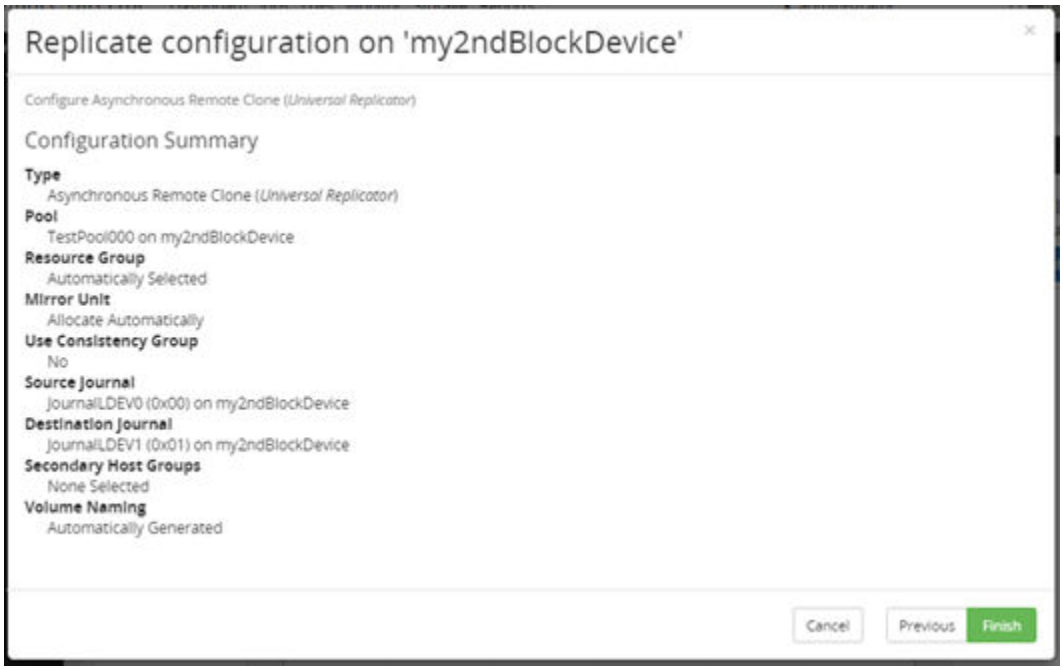


Figure 186 Replicate Configuration Wizard - Asynchronous Remote Clone (Universal Replicator) - Summary

Shows a summary of the replication configuration specified by the user.

Replication Configuration Wizard - Asynchronous Remote Failover

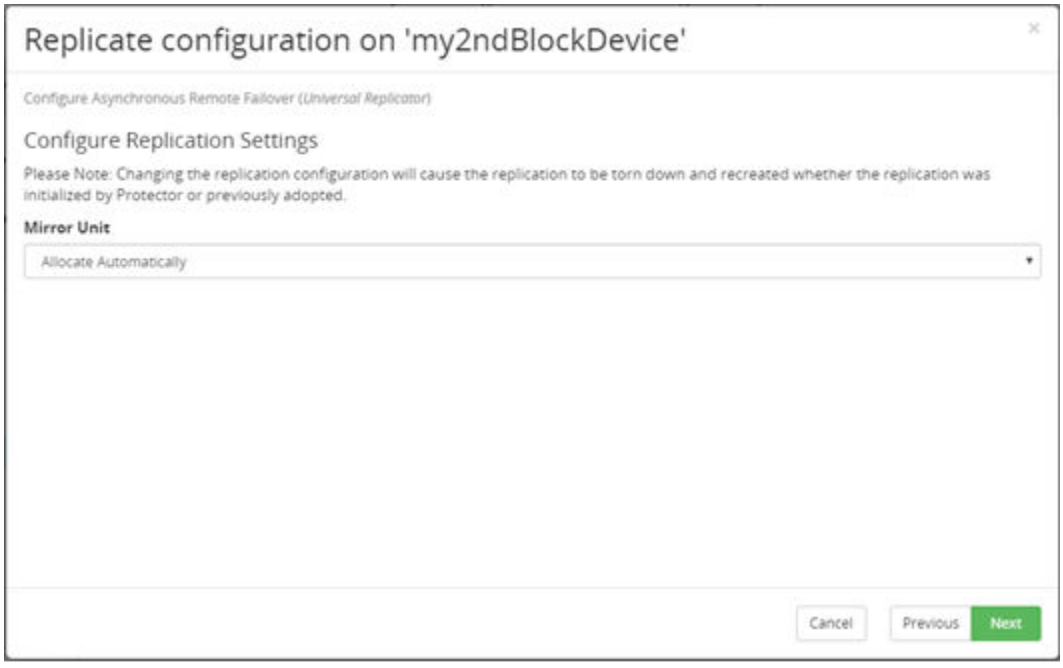


Figure 187 Replicate Configuration Wizard - Asynchronous Remote Failover (Universal Replicator Failover) - Configure Mirror Unit

Control	Description
Mirror Unit	The mirror unit number for the replication can be set to 0, h1, h2 or h3. Select Allocate Automatically to allow Protector to choose one. **Note:** If the mirror unit of an active replication is changed after initial data flow activation then: 1. S-VOLs and pairing relationships for the replication will be destroyed (or dissociated if previously adopted). 2. The replication will then be recreated (or readopted if previously adopted) on data flow reactivation. A warning is issued by the rules compiler prior to activation.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

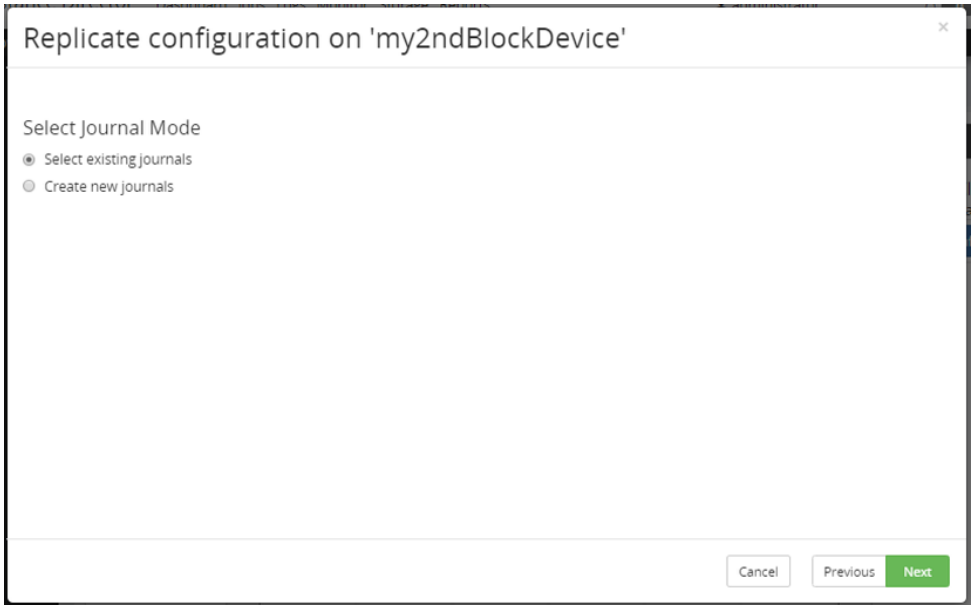


Figure 188 Replicate Configuration Wizard - Asynchronous Remote Failover (Universal Replicator Failover) - Select Journal Mode

Control	Description
Select existing journals	Select this option to use journals that already exist on the source and destination storage arrays. The **Select existing journals** wizard page is displayed next.

Control	Description
Create new journals	Selecf this option to have Protector configure new journals on the sournce and destination storage arrays. The **Create journals** wizard page is displayed next.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

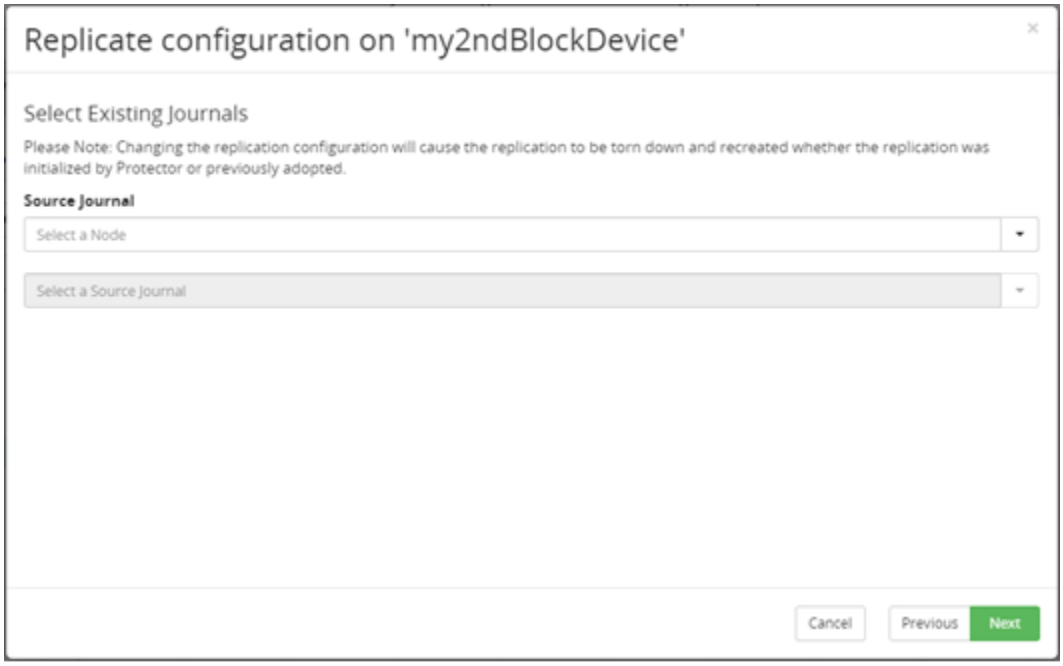


Figure 189 Replicate Configuration Wizard - Asynchronous Remote Failover (Universal Replicator Failover) - Select existing journals

Note: The journal must be unique to each operation and policy in the data flow. The journal should also be used exclusively by Protector.

Control	Description
Source Journal	Specifies the node and journal on the source side of the replication.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

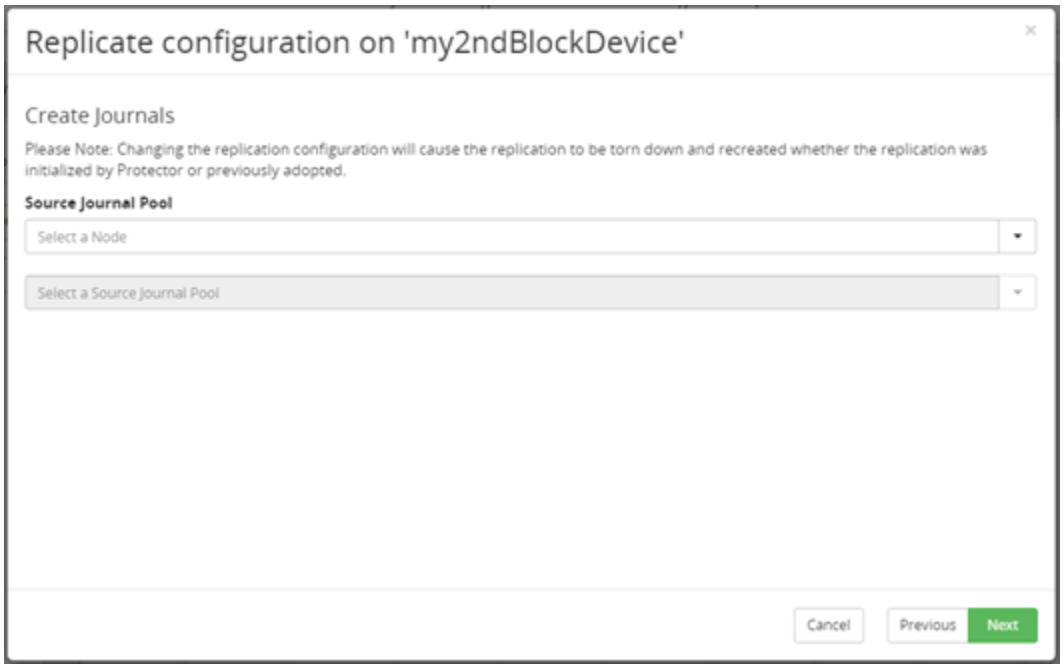


Figure 190 Replicate Configuration Wizard - Asynchronous Remote Failover (Universal Replicator Failover) - Create journals

Note: The journal must be unique to each operation and policy in the data flow. The journal should also be used exclusively by Protector.

Control	Description
Source Journal Pool	Specifies the node and pool where the source side journal will be created.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.



Figure 191 Data Flow Wizard HPE Block Replication Configuration - Select Remote Path Group

Control	Description
Select Remote Path Group	Specifies the Remote Path Group to be used for the replication. ▪ Automatically Selected - Allows Protector to automatically select a Remote Path Group **Note:** For GAD it is recommended the user specifies a group to avoid sharing with other replications. ▪ User Selected - The user specifies the Remote Path Group. **WARNING:** You cannot specify the "User Selected" option and select a path group with an ID of 0. To use a path group with an ID of 0, specify the "Automatically Selected" option. The path group with the lowest ID will be selected (which will be ID 0, if a path group with that ID exists). Remote path groups are listed in the format: Path Group Id: 0x51 Port Mappings: 5E <-> 3E The arrow depicts the direction of the path, either left to right or bidirectional. The arrows can also have a line through them depicting the path is currently broken. Only remote path groups that are suitable for the replication are displayed, for example for GAD only bidirectional path groups are listed.

Control	Description
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

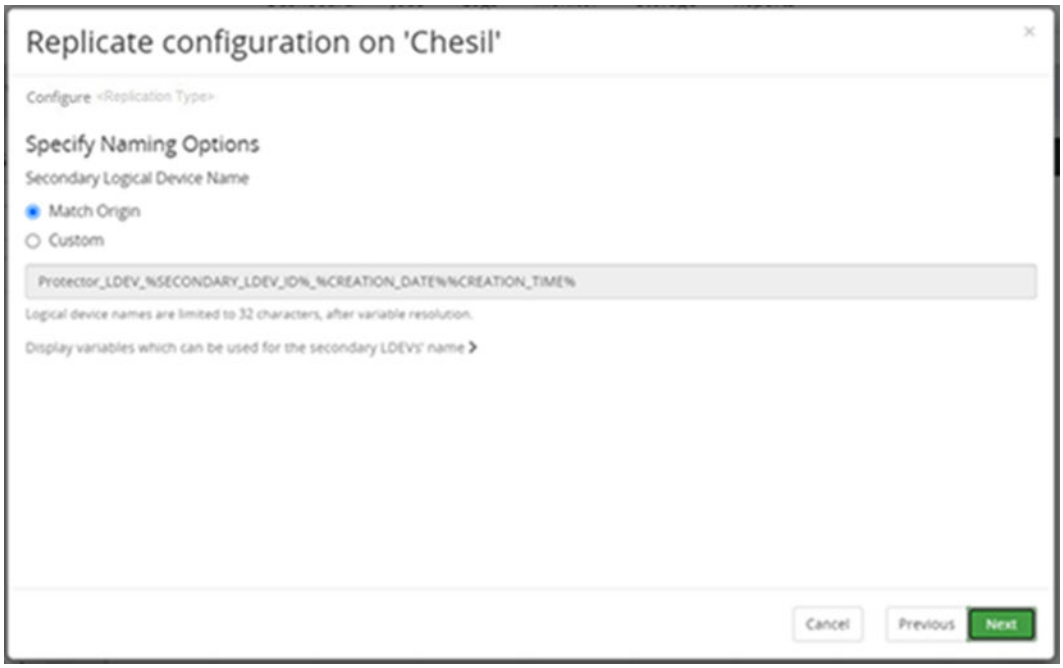


Figure 192 Replicate Configuration Wizard - Naming Options

Control	Description
Secondary Logical Device Name	Specifies how S-VOLs will be named: ▪ Match Origin - The S-VOL will be given the same name as that used for the origin P-VOL (i.e. the left-most volume in the data flow). ▪ Custom - The S-VOL will be named using the naming rule provided. The naming rule can consist of literal strings and/or one or more substitution variables listed. Click Display variables which can be used for the secondary LDEVs' name to view the available substitution variables: • `%ORIGIN_SERIAL%` - S/N of leftmost array in data flow. E.g. output string: `210613` • `%ORIGIN_LDEV_ID%` - ID of leftmost LDEV in data flow. E.g. output string: `00:3A:98` • `%ORIGIN_LDEV_NAME%` - name of leftmost LDEV in data flow. • `%PRIMARY_SERIAL%` - S/N of primary array in this operation. E.g. output string: `442302` • `%PRIMARY_LDEV_ID%` - ID of primary LDEV in this operation. E.g. output string: `00:4C:EB` • `%PRIMARY_LDEV_NAME%` - name of primary LDEV in this operation. • `%SECONDARY_SERIAL%` - S/N of secondary array in this operation. E.g. output string: `356323` • `%SECONDARY_LDEV_ID%` - ID of secondary LDEV in this operation. E.g. output string: `01:F4:35` • `%CREATION_DATE%` - date secondary LDEV was created by this operation. E.g. output string: `20180427` • `%CREATION_TIME%` - time secondary LDEV was created by this operation. 1130
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

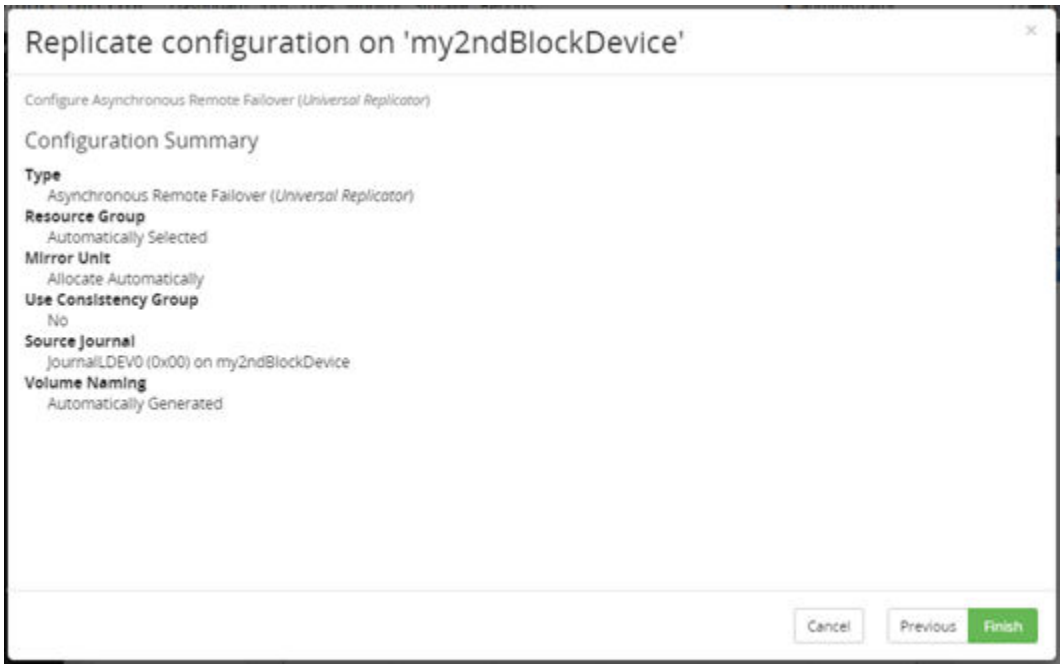


Figure 193 Replicate Configuration Wizard - Asynchronous Remote Failover (Universal Replicator Failover) - Summary

Shows a summary of the replication configuration specified by the user.

Replication Configuration Wizard - Synchronous Remote Clone

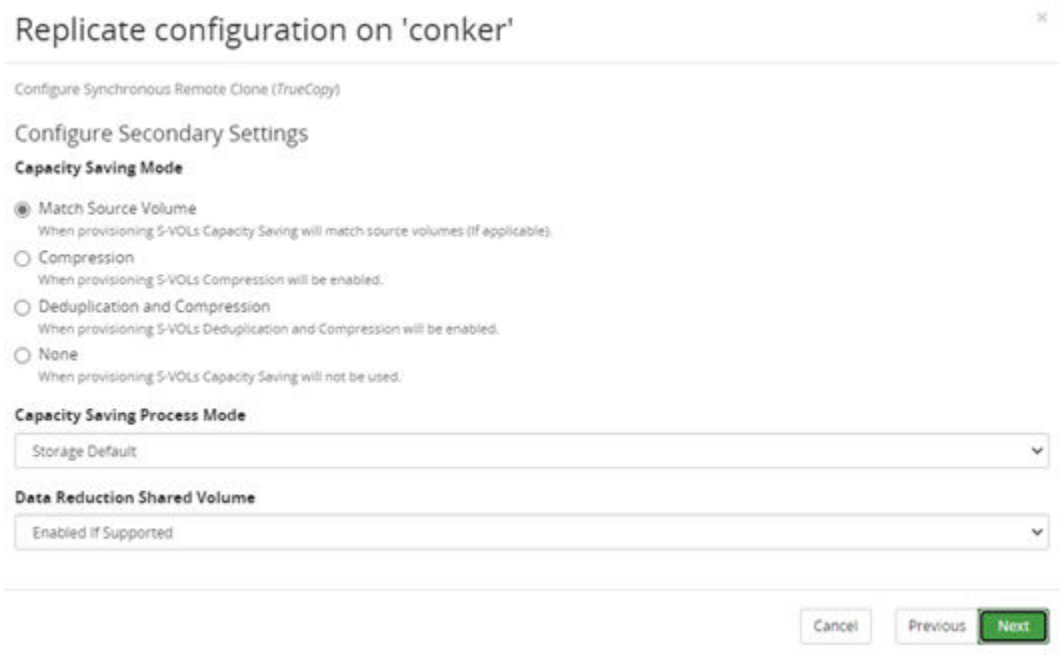


Figure 194 Replicate Configuration Wizard - Synchronous Remote Clone (TrueCopy) - Configure Capacity Savings

Control	Description
Capacity Saving Mode	One of the following options: ▪ Match Source Volumes – When provisioning S-VOLs Capacity Saving will match the settings of the source volumed. ▪ Compression - When provisioning S-VOLs Compression will be enabled, the data compression function utilizes the LZ4 compression algorithm to compress the data. ▪ Deduplication and Compression - When provisioning S-VOLs Deduplication and Compression will be enabled. The data deduplication function deletes duplicate copies of data written to different addresses in the same pool and maintains only a single copy of the data at one address. ▪ None - When provisioning S-VOLs Capacity Saving will not be used
Capacity Saving Process Mode	Only available when a Capacity Saving Mode other than None is selected. Can be one of the following: ▪ Inline - When you apply capacity saving with the inline mode the compression and deduplication processing are performed synchronously for new write data. The inline mode minimizes the pool capacity required to store new write data but can impact I/O performance more than the post-process mode. ▪ Post Process - When you apply capacity saving with the post-process mode the compression and deduplication processing are performed asynchronously for new write data. ▪ Storage Default – match the default option set on the storage array. ▪ Match Source Volume – match the settings of the source volume.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

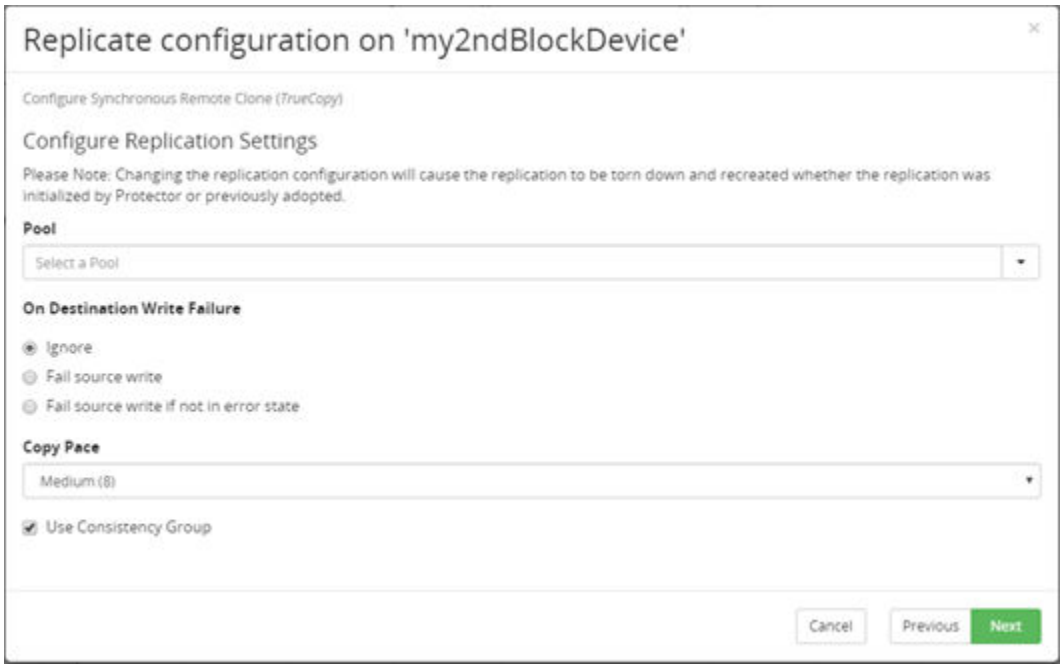


Figure 195 Replicate Configuration Wizard - Synchronous Remote Clone (TrueCopy) - Configure Pool etc.

Control	Description
Pool	Specifies the target storage pool from which replication LDEVs are allocated. Provides a list of available pools giving name and available space. **Note:** All replication types have pool except Asynchronous Remote Failover (Universal Replicator). Select a Dynamic Provisioning Pool. **Note:** Dynamic Provisioning Pools must be created using Storage Navigator prior to selecting the Target Storage/Pool in Protector.
Ignore	When writing to the primary volumes, confirm the write regardless of whether the data has been copied successfully to secondary volumes (Fence Level Never) .
Fail source write	Only confirm primary volume writes if the data is successfully copied to the secondary volume, generate a write error if not (Fence Level Data).
Fail source write if not in error status	Only generate a write error if the data is not successfully copied to the secondary volume and the replication has not been put into an error status PSUE (Fence Level Status).

Control	Description
Copy Pace	Determines how quickly the storage array copies data. The array's default is Slow (3), Protector defaults to Medium (8).
Use Consistency Group	All P-VOLs in a replication are, by default, placed in the same consistency group to ensure consistency of data across all volumes. This option allows this behaviour to be disabled.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

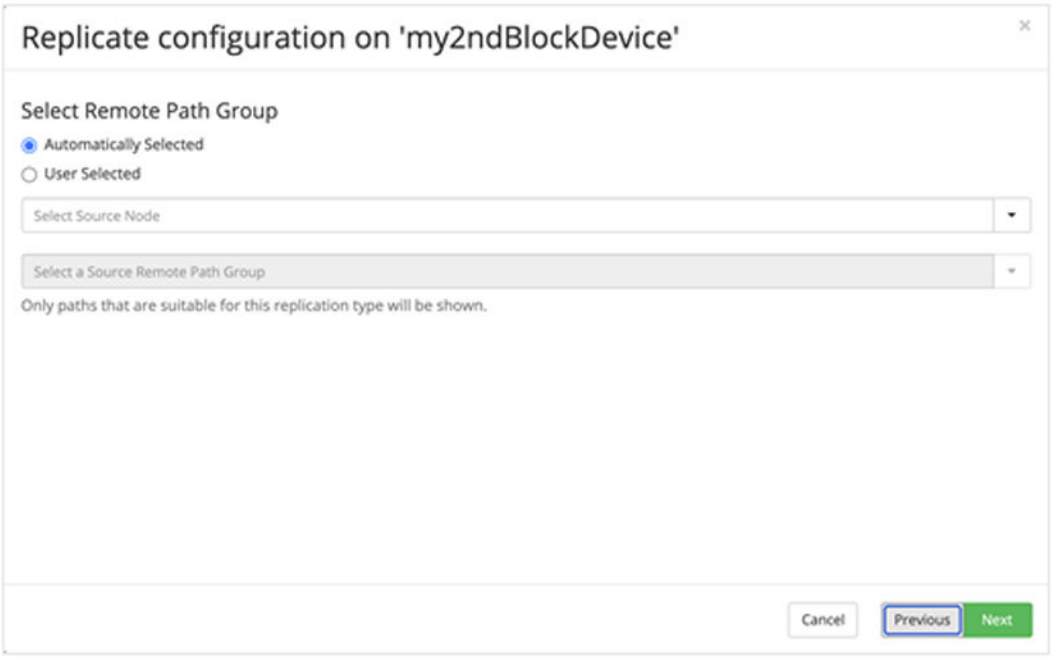


Figure 196 Data Flow Wizard Hitachi Block Replication Configuration - Select Remote Path Group TC Pool

Control	Description
Select Remote Path Group	Specifies the Remote Path Group to be used for the replication. ▪ Automatically Selected - Allows Protector to automatically select a Remote Path Group **Note:** For GAD it is recommended the user specifies a group to avoid sharing with other replications. ▪ User Selected - The user specifies a Fibre path or iSCSI path for the Remote Path Group.

Control	Description
	WARNING: You cannot specify the "User Selected" option and select a path group with an ID of 0. To use a path group with an ID of 0, specify the "Automatically Selected" option. The path group with the lowest ID will be selected (which will be ID 0, if a path group with that ID exists). Remote path groups are listed in the format: Path Group Id: 0x51, Port Type: Fibre, Port Mappings: 5E <-> 3E The arrow depicts the direction of the path, either left to right or bidirectional. The arrows can also have a line through them depicting the path is currently broken. Only remote path groups that are suitable for the replication are displayed, for example for GAD only bidirectional path groups are listed.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

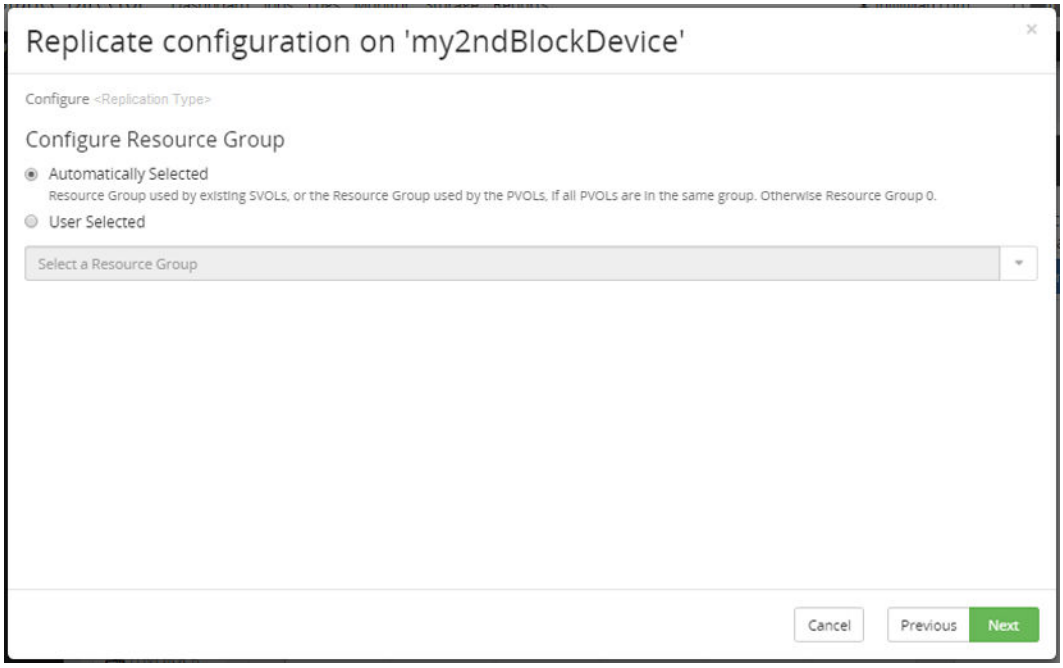


Figure 197 Replicate Configuration Wizard - Configure Resource Group

Control	Description
Configure Resource Group	Specifies the resource group to be used for S-VOLs, in order to support snapshots and replications from VSM volumes (adding volumes to a VSM is performed by adding the volumes to the correct resource group). **Note:** If there are existing S-VOLs, then the resource group used by those will be selected. If the existing S-VOLs are in multiple resource groups or in a resource group that contradicts the user selection, then the operation will fail with an error. This setting should not be modified for existing replications. ▪ Automatically Selected - Allows Protector to automatically select a resource group in the following order of priority: 1. If there are existing S-VOLs, then the resource group used by those will be selected. 2. The resource group used by the P-VOLs, if the replication is in-system and the P-VOLs are all in one resource group. 3. Resource group 0. **Note:** If existing S-VOLs are in multiple resource groups, then the operation will fail with an error. ▪ User Selected - The user specifies the Resource Group.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

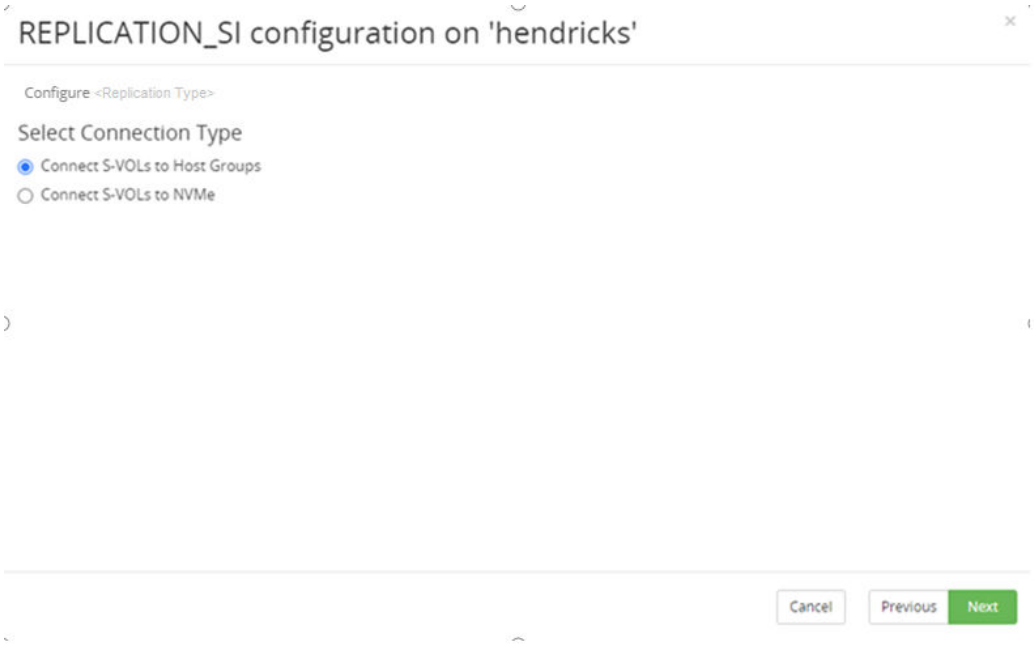


Figure 198 Replicate Configuration Wizard - Select Connection Type

Control	Description
Select protection type	Specifies whether to configure the replication via Host Groups or through an NVMe subsystem.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

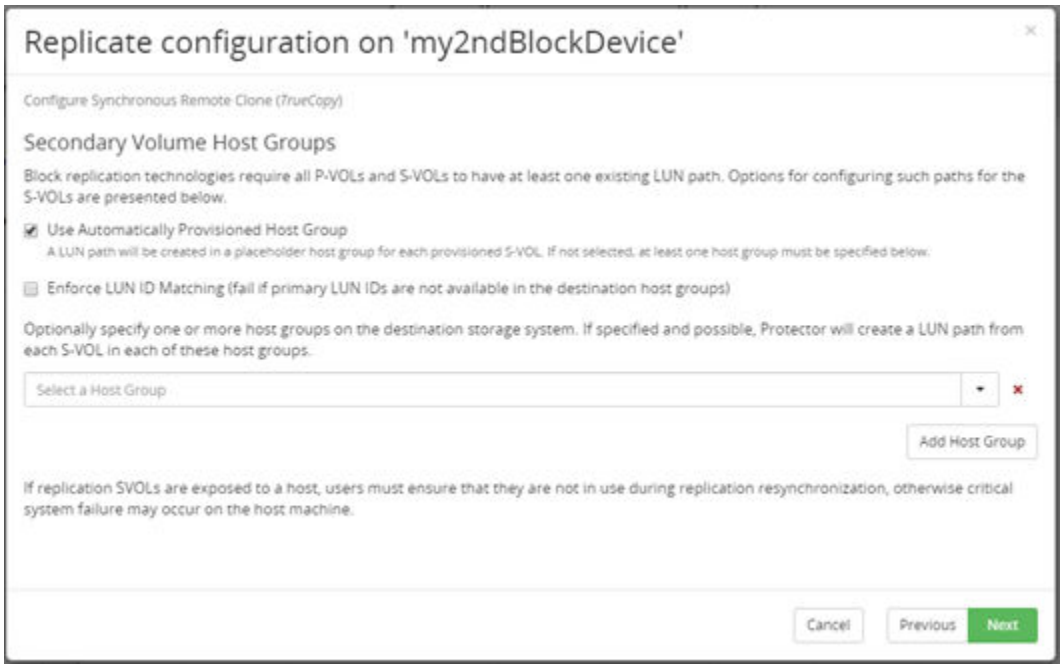


Figure 199 Replicate Configuration Wizard - Secondary Host Groups (Sync. RC)

Caution: If replication S-VOLs are exposed to a host, the user is responsible for ensuring they are not in use during replication resynchronization. Failing to do so may result in a critical failure of the host.

Note: Protector analyses the LUN IDs of the P-VOLs in all host group paths and, if consistent and available, uses these LUN IDs for all S-VOL host group paths either during initial set-up, mounting, or when adding additional LUNs on demand.

To provide a graceful fallback:

- If the P-VOL LUN IDs are not consistent in all host group paths Protector will still use a consistent ID for S-VOL mappings, but these will not necessarily match any of the P-VOL LUN IDs.
- If a P-VOL has a consistent ID in all host group paths, but this LUN ID is not available in all S-VOL host group paths then Protector will choose a different, consistent LUN ID for the S-VOL for all host group paths.
- When not able to match LUN IDs with those used by P-VOLs and/or the S-VOLs, Protector chooses an unused LUN ID. In order to keep LUN ID's compatible with VMware and Hyper-V when selecting LUN ID's Protector will first attempt to select ID's at or below 255. If this is not it will then attempt at or below 1024, followed by at or below 2048. Finally, at or below the array maximum. A warning will be displayed when an ID can not be found in a range.

Some systems expect LUN ID 0 to be used only as a boot volume. Protector will therefore only use LUN ID 0 for the S-VOL host group paths if the P-VOL LUN ID is 0.

Control	Description
Use Protector Provisioned Host Group	Protector will create a LUN path from each S-VOL it provisions in a placeholder host group. If this option is not selected, at least one Secondary Host Group must be specified below.
Enforce LUN ID Matching	For environments where LUN ID consistency is mandatory, selecting this option will cause the replication to fail if: ▪ The P-VOL does not have a consistent LUN ID in all host group paths. ▪ The P-VOL LUN ID is not available in all S-VOL host group paths.
Secondary Host Groups	Specify zero or more host groups that Protector will configure to provide access to the S-VOL(s) when configuring replication scenarios. If no host groups are specified here then Protector will place the S-VOL(s) in it's dummy host group. Click the Add Host Group button to insert another Host Group selection control. Click the Remove button next to a Host Group selection control to delete it. **Note:** If a LUN path to be created already exists, Protector will not attempt to add it again, or to change its ID. The specified host groups must be in the same resource group as the secondary volumes. For GAD replications, if the host group names and port IDs match between primary and secondary storage nodes, Protector will attempt to match the LUN IDs used for the S-VOLs with those of the respective P-VOLs. If this cannot be achieved then a warning will be logged and the next available LUN ID will be used. **Tip:** Use this option to specify host groups that enable access to S-VOL(s) when configuring GAD cross-path and multi-path and other replication scenarios where the S-VOL(s) will need to be accessed (e.g. during failover).
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

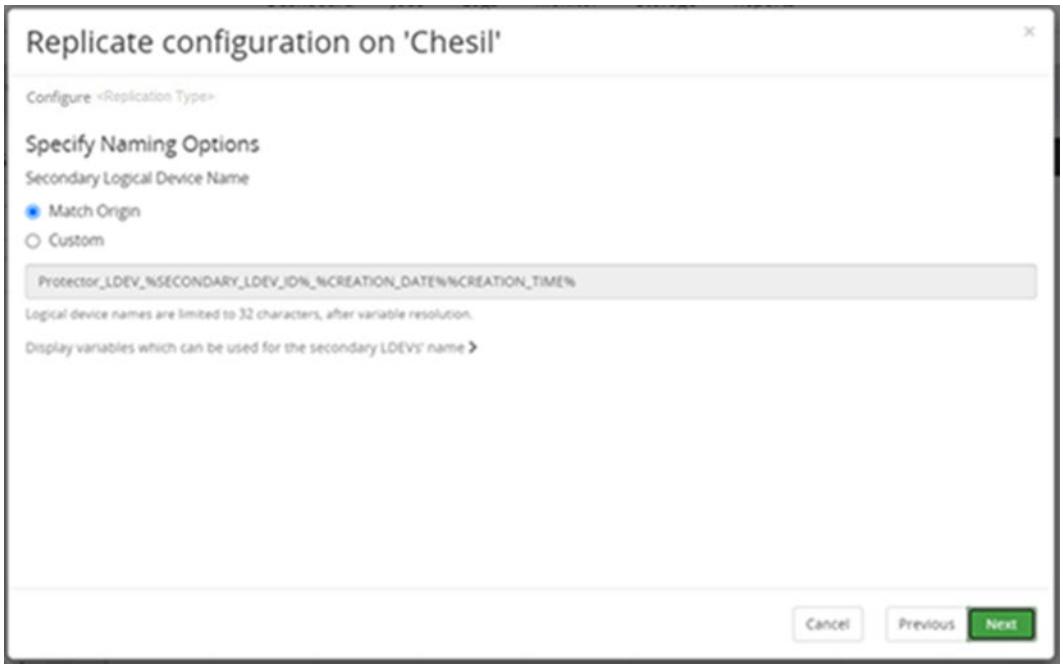


Figure 200 Replicate Configuration Wizard - Naming Options

Control	Description
Secondary Logical Device Name	Specifies how S-VOLs will be named: ▪ Match Origin - The S-VOL will be given the same name as that used for the origin P-VOL (i.e. the left-most volume in the data flow). ▪ Custom - The S-VOL will be named using the naming rule provided. The naming rule can consist of literal strings and/or one or more substitution variables listed. Click Display variables which can be used for the secondary LDEVs' name to view the available substitution variables: • `%ORIGIN_SERIAL%` - S/N of leftmost array in data flow. E.g. output string: `210613` • `%ORIGIN_LDEV_ID%` - ID of leftmost LDEV in data flow. E.g. output string: `00:3A:98` • `%ORIGIN_LDEV_NAME%` - name of leftmost LDEV in data flow. • `%PRIMARY_SERIAL%` - S/N of primary array in this operation. E.g. output string: `442302` • `%PRIMARY_LDEV_ID%` - ID of primary LDEV in this operation. E.g. output string: `00:4C:EB` • `%PRIMARY_LDEV_NAME%` - name of primary LDEV in this operation. • `%SECONDARY_SERIAL%` - S/N of secondary array in this operation. E.g. output string: `356323` • `%SECONDARY_LDEV_ID%` - ID of secondary LDEV in this operation. E.g. output string: `01:F4:35` • `%CREATION_DATE%` - date secondary LDEV was created by this operation. E.g. output string: `20180427` • `%CREATION_TIME%` - time secondary LDEV was created by this operation. 1130
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

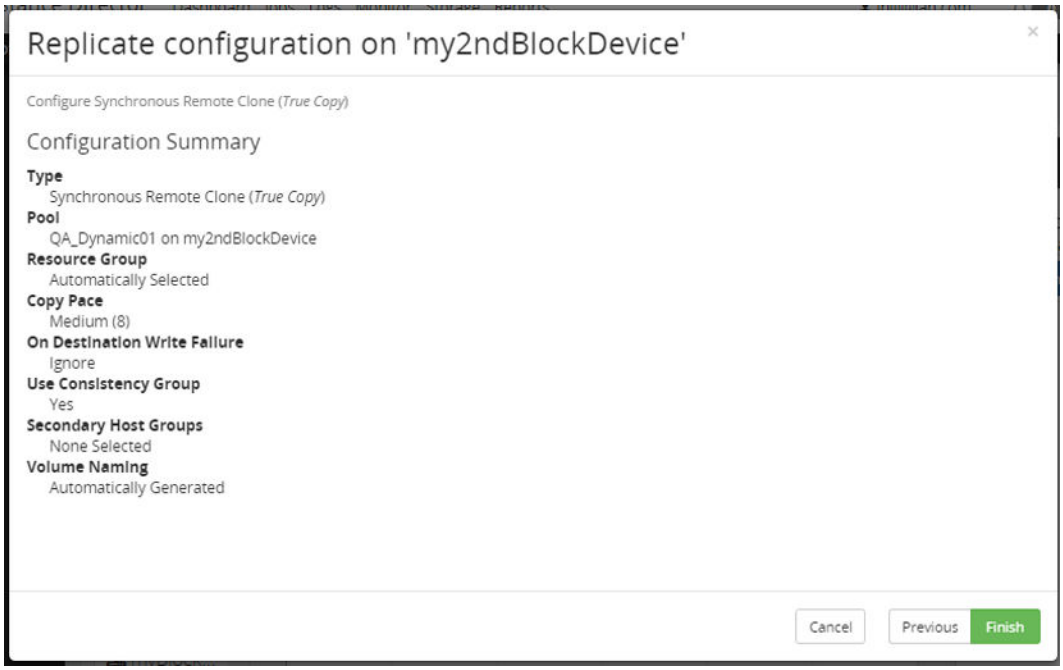


Figure 201 Replicate Configuration Wizard - Synchronous Remote Clone (TrueCopy) - Summary

Shows a summary of the replication configuration specified by the user.

Replication Configuration Wizard - Active-Active Remote Clone

Note:

- GAD is only available on VSP G series because virtualized LDEVs are required.
- Configuration or adoption of GAD cross-path scenarios requires CCI version 01-41-03/03 or greater to be installed on the ISM node.
- Primary volumes must be set up within a host group prior to configuring a GAD replication using Protector.
- GAD replications require the P-VOL and S-VOL to have matching virtual serial numbers and virtual LDEV IDs. Selecting Automatically Selected in the **Configure Resource Group** page of the wizard ensures this is done.
- Any LUN path to the secondary volume will become inaccessible if the GAD replication is deleted from Protector (i.e. torn down). This occurs because the virtual LDEV ID is automatically deleted from the S-VOL, causing host I/Os to be rejected. To recover from this, either recreate the GAD pair (if the pair was deleted unintentionally) or assign a new virtual ID to the clone S-VOL

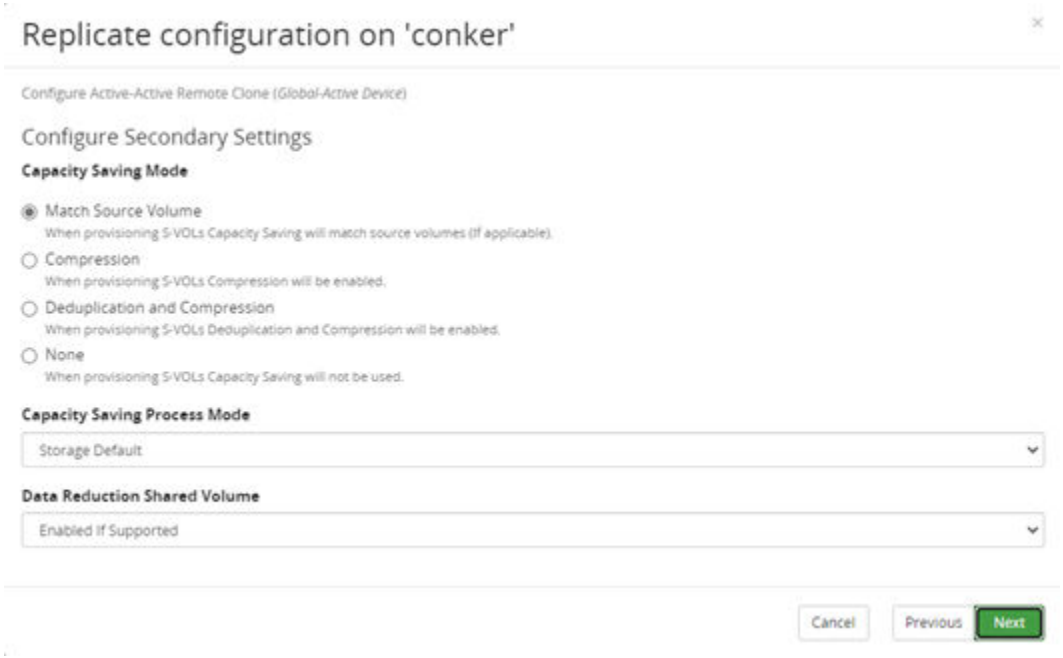


Figure 202 Replicate Configuration Wizard - Active-Active Remote Clone (Global-Active Device) - Configure Capacity Savings

Control	Description
Capacity Saving Mode	One of the following options: ▪ Match Source Volumes – When provisioning S-VOLs Capacity Saving will match the settings of the source volumed. ▪ Compression - When provisioning S-VOLs Compression will be enabled, the data compression function utilizes the LZ4 compression algorithm to compress the data. ▪ Deduplication and Compression - When provisioning S-VOLs Deduplication and Compression will be enabled. The data deduplication function deletes duplicate copies of data written to different addresses in the same pool and maintains only a single copy of the data at one address. ▪ None - When provisioning S-VOLs Capacity Saving will not be used

Control	Description
Capacity Saving Process Mode	Only available when a Capacity Saving Mode other than None is selected. Can be one of the following: ▪ Inline - When you apply capacity saving with the inline mode the compression and deduplication processing are performed synchronously for new write data. The inline mode minimizes the pool capacity required to store new write data but can impact I/O performance more than the post-process mode. ▪ Post Process - When you apply capacity saving with the post-process mode the compression and deduplication processing are performed asynchronously for new write data. ▪ Storage Default – match the default option set on the storage array. ▪ Match Source Volume – match the settings of the source volume.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

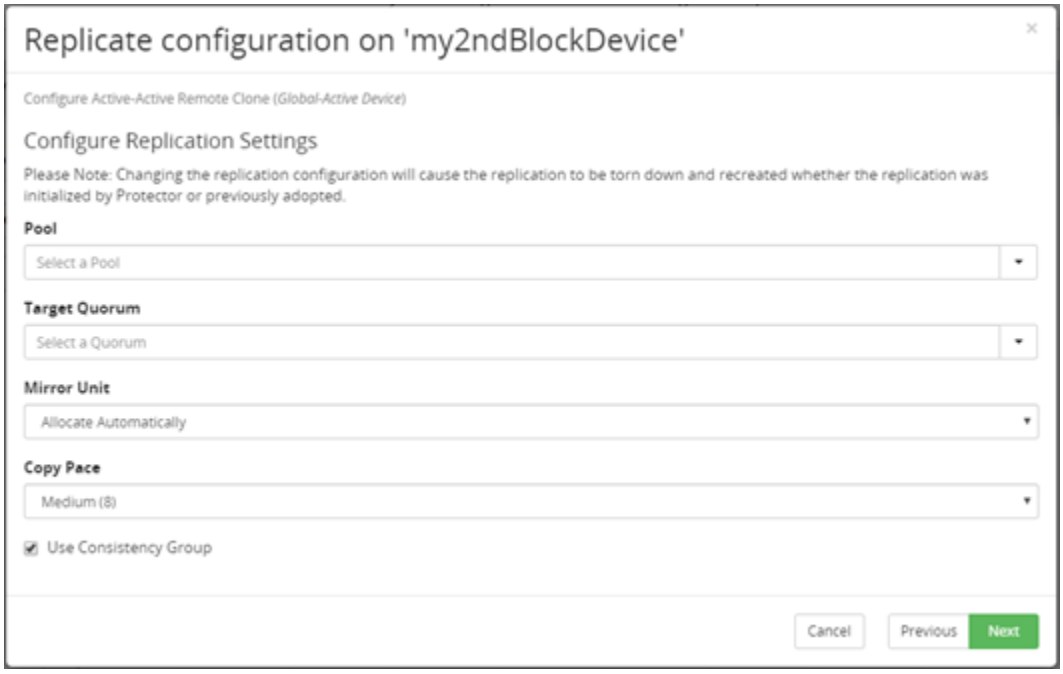


Figure 203 Replicate Configuration Wizard - Active-Active Remote Clone (Global-Active Device) - Configure Pool etc.

Control	Description
Pool	Specifies the target storage pool from which replication LDEVs are allocated. Provides a list of available pools giving name and available space. **Note:** All replication types have pool except Asynchronous Remote Failover (Universal Replicator). Select a Dynamic Provisioning Pool. **Note:** Dynamic Provisioning Pools must be created using Storage Navigator prior to selecting the Target Storage/Pool in Protector.
Target Quorum	Selects the volume to use as the quorum disk. **Note:** ▪ A Quorum disk is required to manage each GAD replication pair. It is best practice to allocate a separate Quorum disk for each pair. ▪ Both Quorum and Quorum-less disks are available for selection.
Mirror Unit	The mirror unit number for the replication can be set to 0, 1, h1 or h2. Select Allocate Automatically to allow Protector to choose one. **Note:** If you add pairs to an existing Dataflow with a different mirror unit ID Protector may delete the pair relationship of the newly added pairs. You must ensure that the mirror IDs are the same for the newly added pairs. If you change the mirror ID for the current dataflow, Protector may destroy the existing pair relationship.
Copy Pace	Determines how quickly the storage array copies data. The array's default is Slow (3), Protector defaults to Medium (8).
Use Consistency Group	All P-VOLs in a replication are, by default, placed in the same consistency group to ensure consistency of data across all volumes. This option allows this behavior to be disabled.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

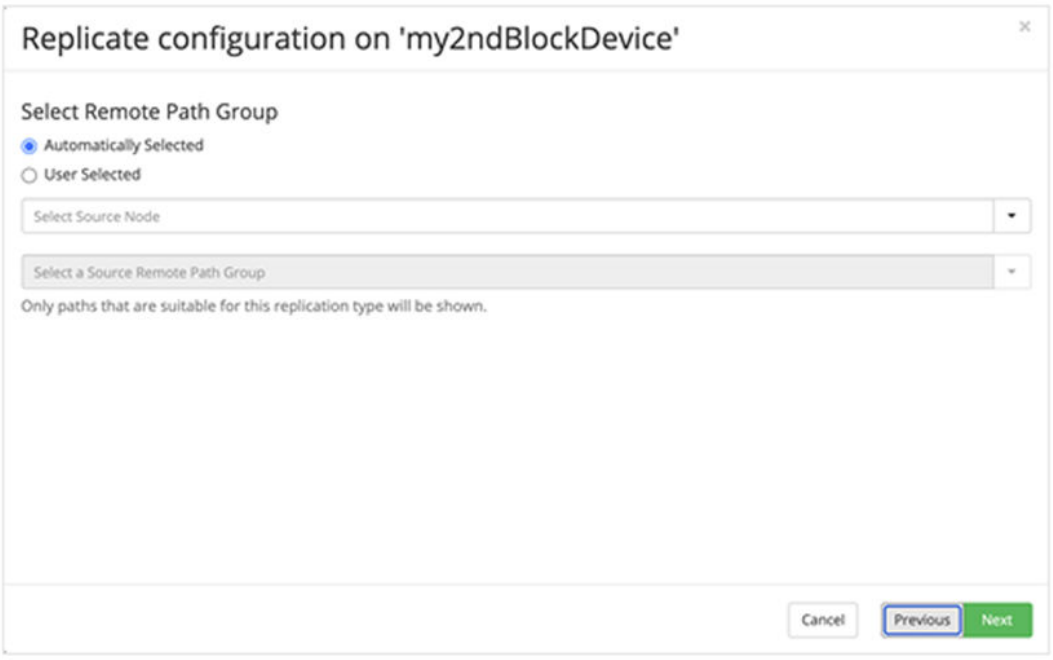


Figure 204 Data Flow Wizard Hitachi Block Replication Configuration - Select Remote Path Group

Control	Description
Select Remote Path Group	Specifies the Remote Path Group to be used for the replication. ▪ Automatically Selected - Allows Protector to automatically select a Remote Path Group **Note:** For GAD it is recommended the user specifies a group to avoid sharing with other replications. ▪ User Selected - The user specifies the Remote Path Group. **WARNING:** You cannot specify the "User Selected" option and select a path group with an ID of 0. To use a path group with an ID of 0, specify the "Automatically Selected" option. The path group with the lowest ID will be selected (which will be ID 0, if a path group with that ID exists). Remote path groups are listed in the format: Path Group Id: 0x51 Port Mappings: 5E <-> 3E The arrow depicts the direction of the path, either left to right or bidirectional. The arrows can also have a line through them depicting the path is currently broken. Only remote path groups that are suitable for the replication are displayed, for example for GAD only bidirectional path groups are listed.

Control	Description
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

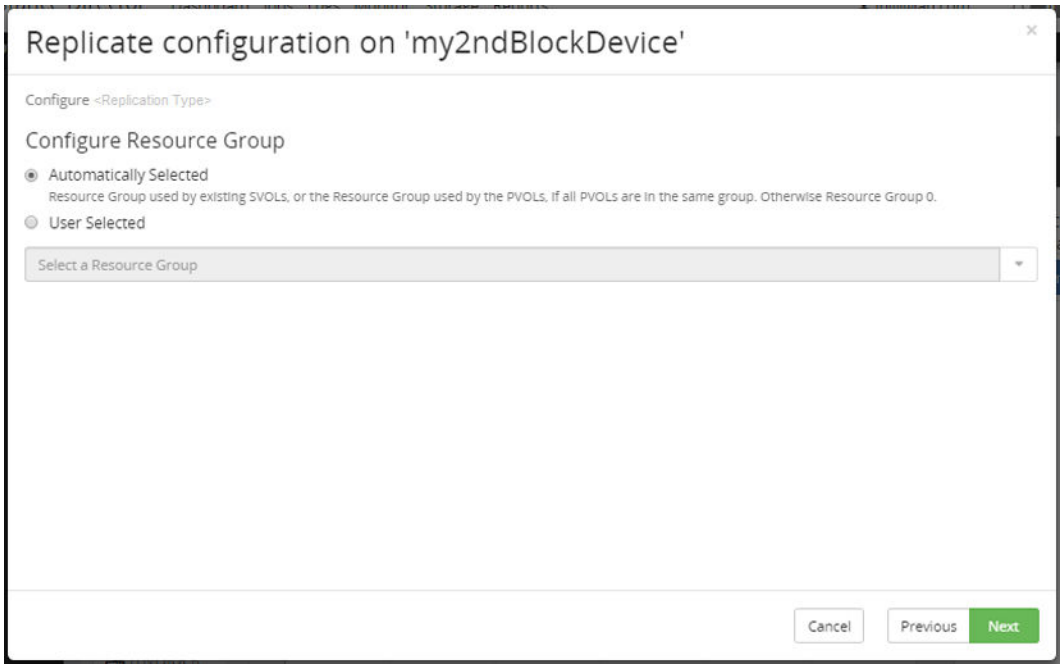


Figure 205 Replicate Configuration Wizard - Configure Resource Group

Control	Description
Configure Resource Group	Specifies the resource group to be used for S-VOLs, in order to support snapshots and replications from VSM volumes (adding volumes to a VSM is performed by adding the volumes to the correct resource group).

Control	Description
	Note: If there are existing S-VOLs, then the resource group used by those will be selected. If the existing S-VOLs are in multiple resource groups or in a resource group that contradicts the user selection, then the operation will fail with an error. This setting should not be modified for existing replications. ▪ Automatically Selected - Allows Protector to automatically select a resource group in the following order of priority: 1. If there are existing S-VOLs, then the resource group used by those will be selected. 2. The resource group used by the P-VOLs, if the replication is in-system and the P-VOLs are all in one resource group. 3. Resource group 0. **Note:** If existing S-VOLs are in multiple resource groups, then the operation will fail with an error. ▪ User Selected - The user specifies the Resource Group.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

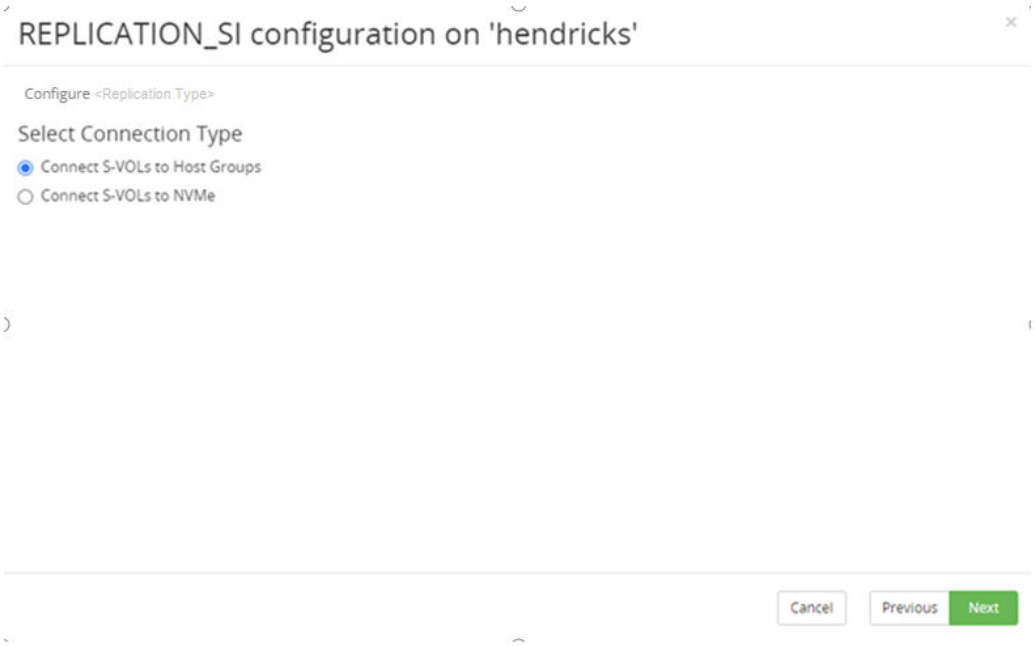


Figure 206 Replicate Configuration Wizard - Select Connection Type

Control	Description
Select protection type	Specifies whether to configure the replication via Host Groups or through an NVMe subsystem.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

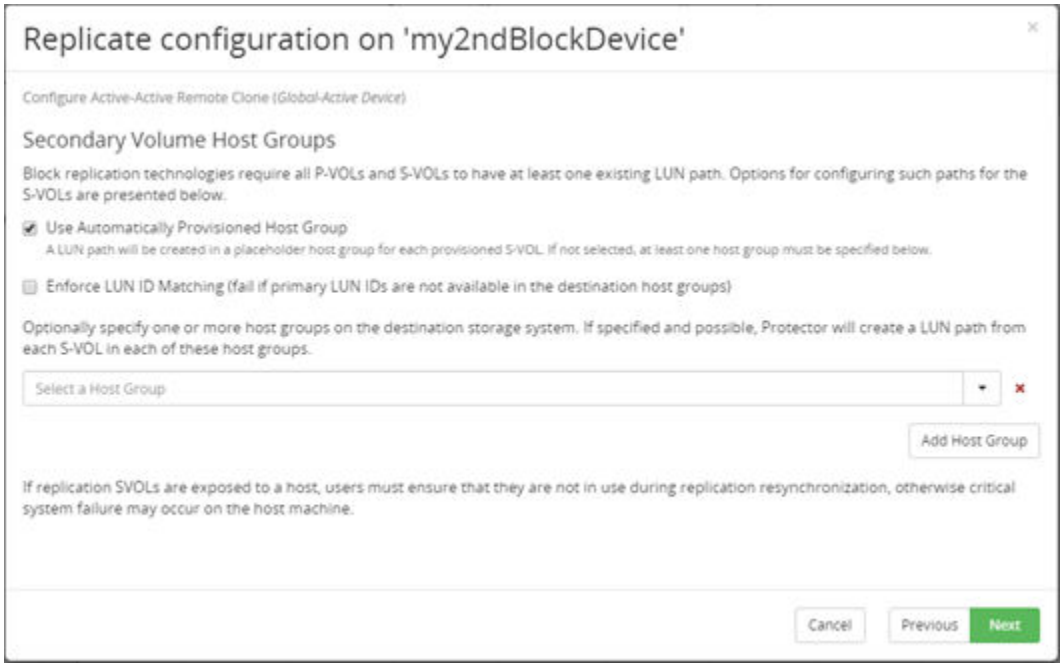


Figure 207 Replicate Configuration Wizard - Secondary Host Groups (GAD)

Caution: If replication S-VOLs are exposed to a host, the user is responsible for ensuring they are not in use during replication resynchronization. Failing to do so may result in a critical failure of the host.

Note: Protector analyses the LUN IDs of the P-VOLs in all host group paths and, if consistent and available, uses these LUN IDs for all S-VOL host group paths either during initial set-up, mounting, or when adding additional LUNs on demand.

To provide a graceful fallback:

- If the P-VOL LUN IDs are not consistent in all host group paths Protector will still use a consistent ID for S-VOL mappings, but these will not necessarily match any of the P-VOL LUN IDs.
- If a P-VOL has a consistent ID in all host group paths, but this LUN ID is not available in all S-VOL host group paths then Protector will choose a different, consistent LUN ID for the S-VOL for all host group paths.
- When not able to match LUN IDs with those used by P-VOLs and/or the S-VOLs, Protector chooses an unused LUN ID. In order to keep LUN ID's compatible with VMware and Hyper-V when selecting LUN ID's Protector will first attempt to select ID's at or below 255. If this is not it will then attempt at or below 1024, followed by at or below 2048. Finally, at or below the array maximum. A warning will be displayed when an ID can not be found in a range.

Some systems expect LUN ID 0 to be used only as a boot volume. Protector will therefore only use LUN ID 0 for the S-VOL host group paths if the P-VOL LUN ID is 0.

Control	Description
Use Protector Provisioned Host Group	Protector will create a LUN path from each S-VOL it provisions in a placeholder host group. If this option is not selected, at least one Secondary Host Group must be specified below.
Enforce LUN ID Matching	For environments where LUN ID consistency is mandatory, selecting this option will cause the replication to fail if: ▪ The P-VOL does not have a consistent LUN ID in all host group paths. ▪ The P-VOL LUN ID is not available in all S-VOL host group paths.
Secondary Host Groups	Specify zero or more host groups that Protector will configure to provide access to the S-VOL(s) when configuring replication scenarios. If no host groups are specified here then Protector will place the S-VOL(s) in it's dummy host group. Click the Add Host Group button to insert another Host Group selection control. Click the Remove button next to a Host Group selection control to delete it.

Control	Description
	Note: If a LUN path to be created already exists, Protector will not attempt to add it again, or to change its ID. The specified host groups must be in the same resource group as the secondary volumes. For GAD replications, if the host group names and port IDs match between primary and secondary storage nodes, Protector will attempt to match the LUN IDs used for the S-VOLs with those of the respective P-VOLs. If this cannot be achieved then a warning will be logged and the next available LUN ID will be used. **Tip:** Use this option to specify host groups that enable access to S-VOL(s) when configuring GAD cross-path and multi-path and other replication scenarios where the S-VOL(s) will need to be accessed (e.g. during failover).
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

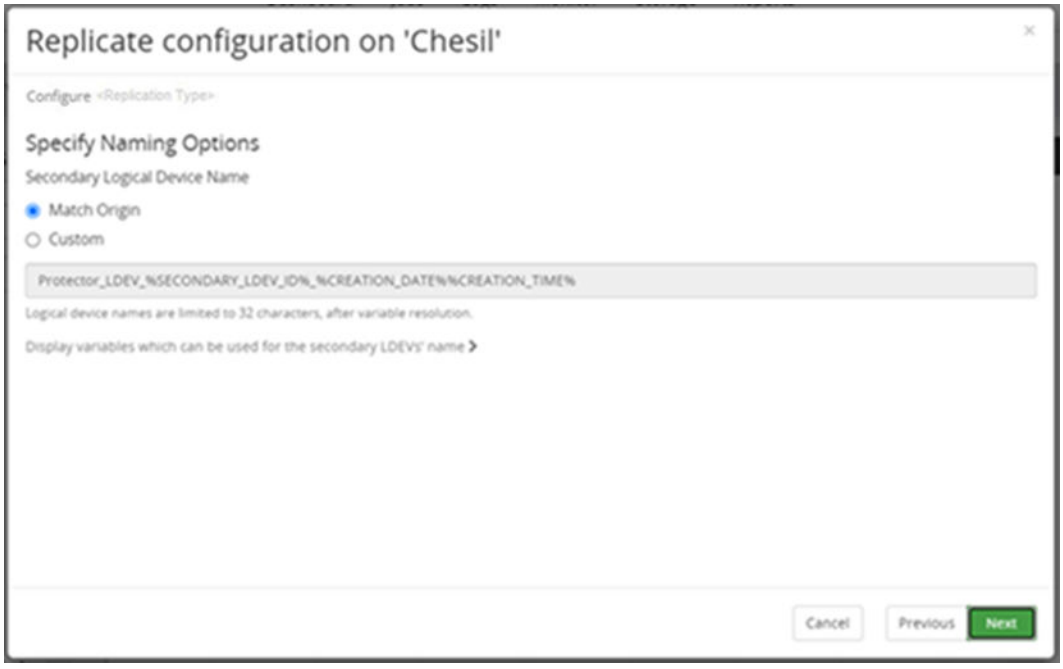


Figure 208 Replicate Configuration Wizard - Naming Options

Control	Description
Secondary Logical Device Name	Specifies how S-VOLs will be named: ▪ Match Origin - The S-VOL will be given the same name as that used for the origin P-VOL (i.e. the left-most volume in the data flow). ▪ Custom - The S-VOL will be named using the naming rule provided. The naming rule can consist of literal strings and/or one or more substitution variables listed. Click Display variables which can be used for the secondary LDEVs' name to view the available substitution variables: • `%ORIGIN_SERIAL%` - S/N of leftmost array in data flow. E.g. output string: `210613` • `%ORIGIN_LDEV_ID%` - ID of leftmost LDEV in data flow. E.g. output string: `00:3A:98` • `%ORIGIN_LDEV_NAME%` - name of leftmost LDEV in data flow. • `%PRIMARY_SERIAL%` - S/N of primary array in this operation. E.g. output string: `442302` • `%PRIMARY_LDEV_ID%` - ID of primary LDEV in this operation. E.g. output string: `00:4C:EB` • `%PRIMARY_LDEV_NAME%` - name of primary LDEV in this operation. • `%SECONDARY_SERIAL%` - S/N of secondary array in this operation. E.g. output string: `356323` • `%SECONDARY_LDEV_ID%` - ID of secondary LDEV in this operation. E.g. output string: `01:F4:35` • `%CREATION_DATE%` - date secondary LDEV was created by this operation. E.g. output string: `20180427` • `%CREATION_TIME%` - time secondary LDEV was created by this operation. 1130
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

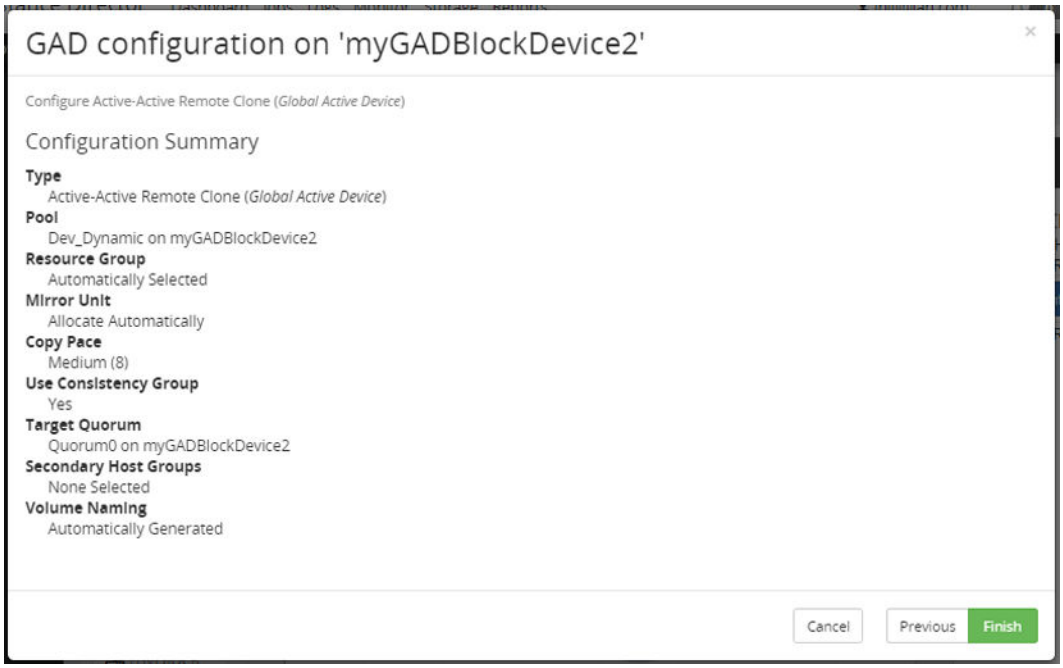


Figure 209 Replicate Configuration Wizard - Active-Active Remote Clone (Global-Active Device) - Summary

Shows a summary of the replication configuration specified by the user.

Hitachi Software Block Replication Configuration Wizard

This wizard is displayed when you assign a replication operation to a Hitachi Software Block Device node on a data flow.

Caution: When a replication operation runs, Protector locks the *meta_resource* of the Block Device until the operation completes. Block operations are queued waiting to get the resource lock; this can impact RPO.

Note:

All replication and snapshot S-VOLs must be created using free LDEV IDs that are mapped to the *meta_resource* group, and have virtual LDEV IDs matching their corresponding physical LDEV IDs.

For fully provisioned snapshots and all replications, this applies to the operation that creates that snapshot or replication.

For floating device snapshots and snapshots mounted using cascade mode, this applies to the mount or restore operation.

For fully provisioned snapshots mounted using cascade mode, this applies both to the operation that creates that snapshot and to the mount or restore operation.

If an operation tries to create one or more LDEVs, that operation will fail if there are not enough free LDEV IDs that meet the above conditions.

Note:

- The classifications that are applicable to Hitachi Software Block Replication operations are Hitachi Block LDEV.
- Protector does not support any Hitachi Block storage LDEV with more than one partition or volume.
- If Protector is subsequently uninstalled, the existing replications and mounted volumes are left in place.

Note:

Once the operation properties have been configured and the data flow has been activated, the properties cannot be changed. Reactivating rules for the active data flow with edited properties will not change them.

To change the operation properties, the existing data flow must first be deactivated then reactivated with the new properties. Deactivating the data flow will tear down the current replication.

Tip:

- The actions that are performed by hardware orchestration as it executes a block backup policy are captured by the Log Manager. Detailed logs and their attachments can be exported as a text file. (See Export Logs Dialog (on page 516) for more information.)
- Details of the hardware resources (LDEVs) used for a particular replication can be found in the Hitachi Software Block Device Details (on page 901).



Figure 210 Software Block Hitachi Replicate Configuration Wizard - Configure New or Adopt Existing Replication

Control	Description
Configure new replication	Creates a new replication defined by the user defined parameters. The **Configure New Replication** page of the wizard is displayed after the Adopt workflow is described..
Adopt existing replication	Adopts the replication defined by the user defined parameters from the matching pre-existing one on the hardware. The **Adopt Existing Replication** page of the wizard is displayed next.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

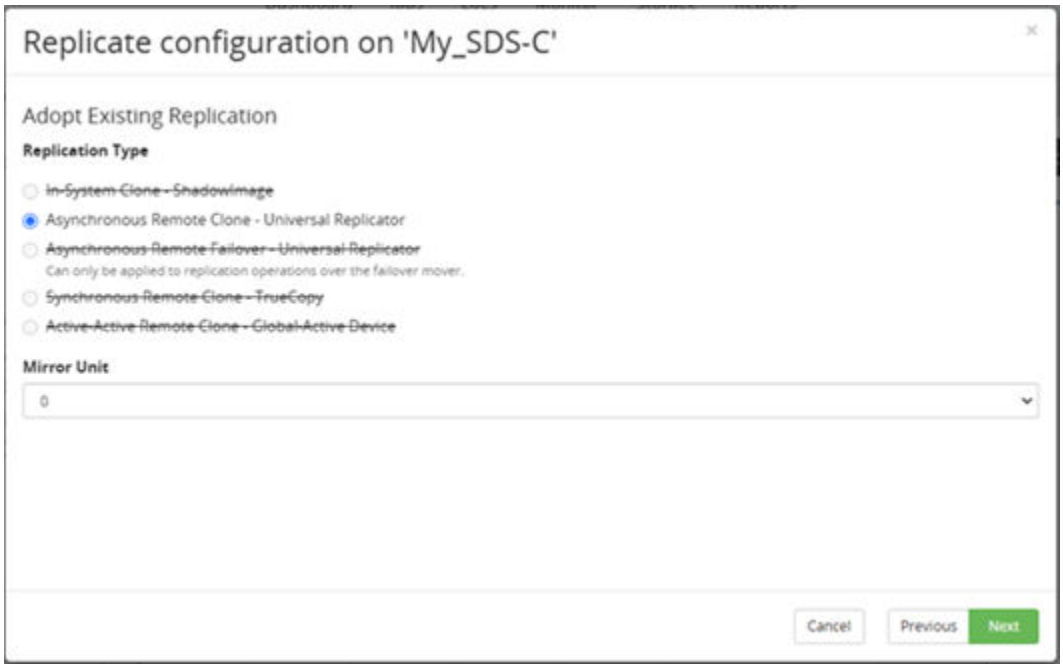


Figure 211 Software Block Hitachi Replicate Adoption Wizard - Adopt Existing Replication

Note: Adoption requires the rules to be activated and the relevant policy to be triggered. Refer to About Hitachi Block replication adoption (on page 87) before using this feature.

Control	Description
Replication Type	Select the type of replication to adopt: ▪ In System Clone - not supported for SDS-B or SDS-C ▪ Refreshed Snapshot - not supported for SDS-B or SDS-C ▪ Asynchronous Remote Clone - see About Universal Replicator (on page 75) ▪ Asynchronous Remote Failover - not supported for SDS-B or SDS-C ▪ Synchronous Remote Clone - not supported for SDS-B or SDS-C ▪ Active-Active Remote Clone - not supported for SDS-B or SDS-C
Mirror Unit	Identify the mirror unit number of the replication to be adopted. Adoption requires at least one existing pair on the selected mirror unit. **Note:** If the mirror unit of an active replication is changed after initial data flow activation then: 1. S-VOLs and pairing relationships for the replication will be destroyed (or dissociated if previously adopted). 2. The replication will then be recreated (or readopted if previously adopted) on data flow reactivation.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

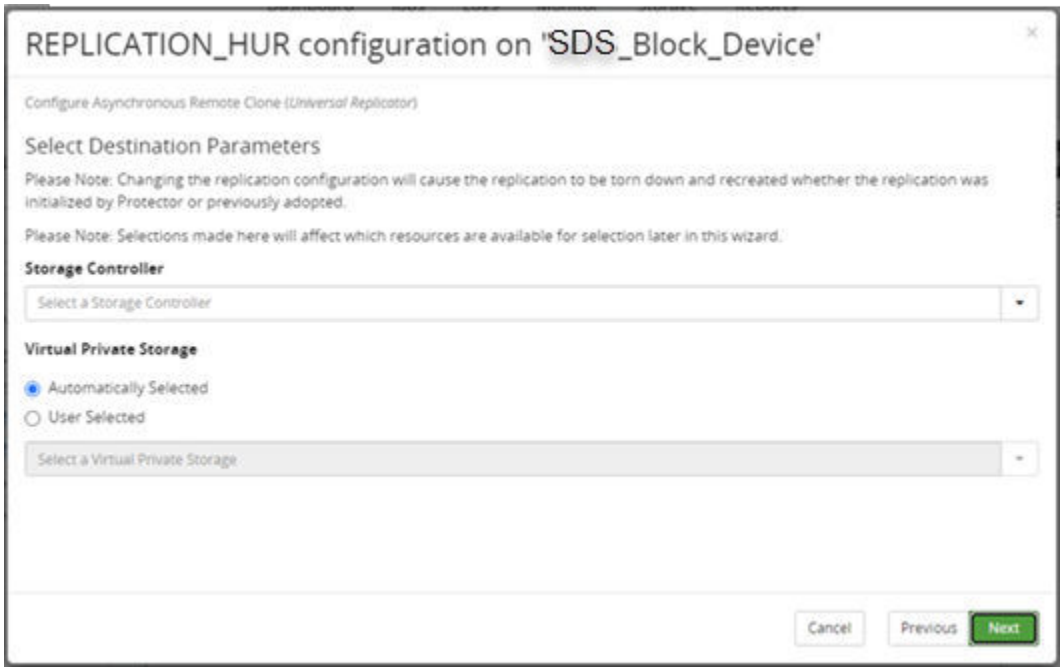


Figure 212 Software Block Hitachi Replicate Adoption Wizard - Asynchronous Remote Clone (Universal Replicator) –Destination Parameters

Control	Description
Storage Controller	Select the relevant storage controller from the drop down list.
Virtual Private Storage	If left to automatically selected Protector will use the VPS that has been configured on the Software Block Device.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

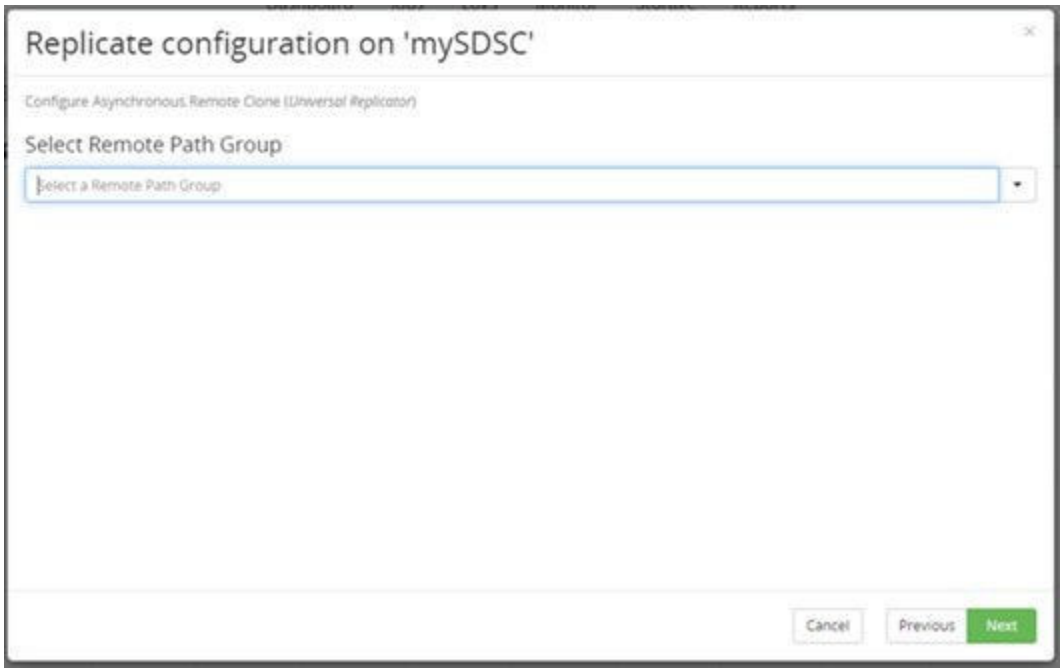


Figure 213 Software Block HitachiReplicate Adoption Wizard - Asynchronous Remote Clone - Select Remote Path Group

Control	Description
Select Remote Path Group	Specifies the Remote Path Group to be used for the replication. Remote path groups are listed in the format: <Path_Grup_ID>(RemoteSerial: <Array_Serial_No>, remote Type:<Array_Type>)
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

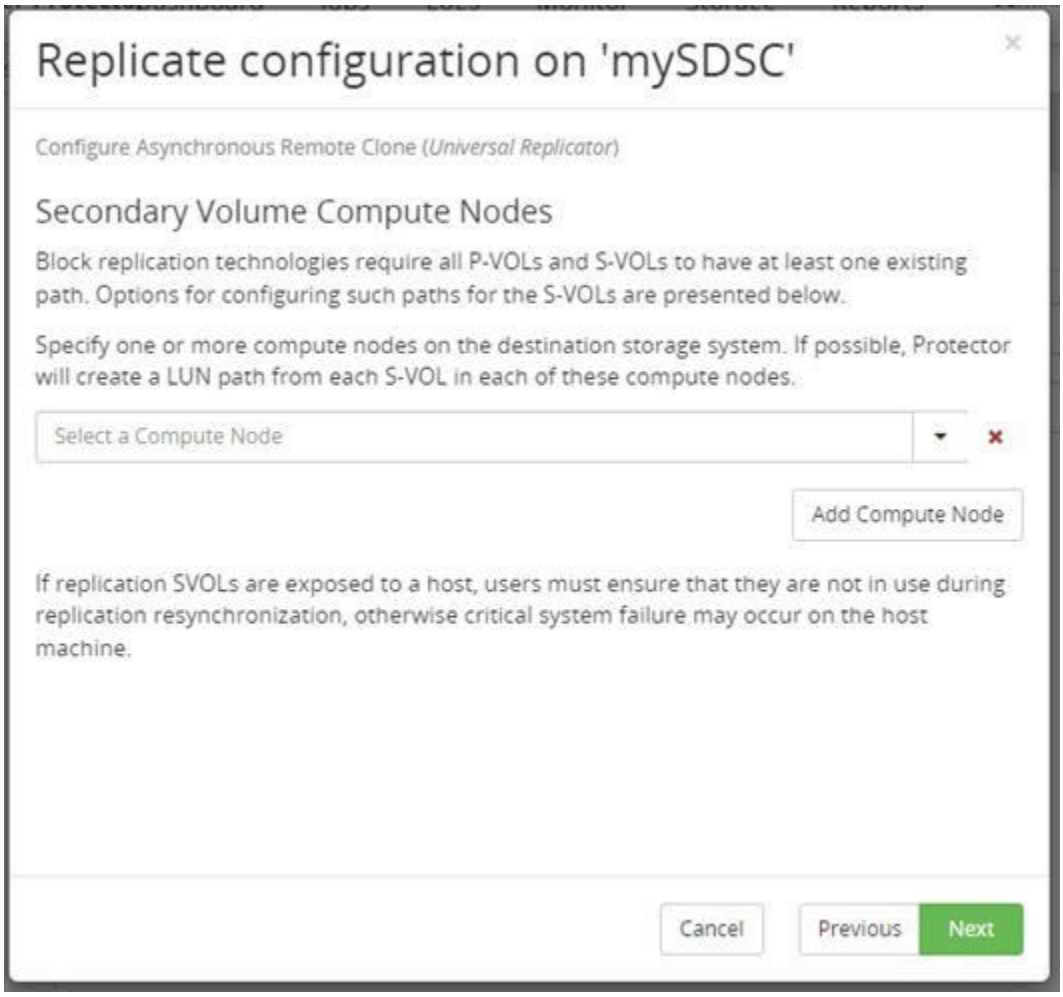


Figure 214 Software Block Hitachi Replicate Adoption Wizard - Configure Asynchronous Remote Clone

Control	Description
Secondary Volume Compute Nodes	Select one or more compute nodes for the secondary volumes. Adding one node will automatically display the option to add more
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.



Figure 215 Software Block Hitachi Replicate Adoption Wizard - Naming Options

Control	Description
Secondary Logical Device Name	Specifies how S-VOLs will be named: ▪ Match Origin - The S-VOL will be given the same name as that used for the origin P-VOL (i.e. the left-most volume in the data flow). ▪ Custom - The S-VOL will be named using the naming rule provided. The naming rule can consist of literal strings and/or one or more substitution variables listed. Click Display variables which can be used for the secondary LDEVs' name to view the available substitution variables: • `%ORIGIN_SERIAL%` - S/N of leftmost array in data flow. E.g. output string: `210613` • `%ORIGIN_LDEV_ID%` - ID of leftmost LDEV in data flow. E.g. output string: `00:3A:98` • `%ORIGIN_LDEV_NAME%` - name of leftmost LDEV in data flow. • `%PRIMARY_SERIAL%` - S/N of primary array in this operation. E.g. output string: `442302` • `%PRIMARY_LDEV_ID%` - ID of primary LDEV in this operation. E.g. output string: `00:4C:EB` • `%PRIMARY_LDEV_NAME%` - name of primary LDEV in this operation. • `%SECONDARY_SERIAL%` - S/N of secondary array in this operation. E.g. output string: `356323` • `%SECONDARY_LDEV_ID%` - ID of secondary LDEV in this operation. E.g. output string: `01:F4:35` • `%CREATION_DATE%` - date secondary LDEV was created by this operation. E.g. output string: `20180427` • `%CREATION_TIME%` - time secondary LDEV was created by this operation. 1130
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.



Figure 216 Software Block Hitachi Replicate Configuration Wizard - Adoption Summary

Shows a summary of the replication configuration specified by the user. All other parameters defining the replication are obtained from the hardware when the data flow is activated and triggered.

Software Block Replication Configuration Wizard - Asynchronous Remote Clone

The following describes the process for setting up an Asynchronous Remote Clone replication to an SDS-B or SDS-C node.

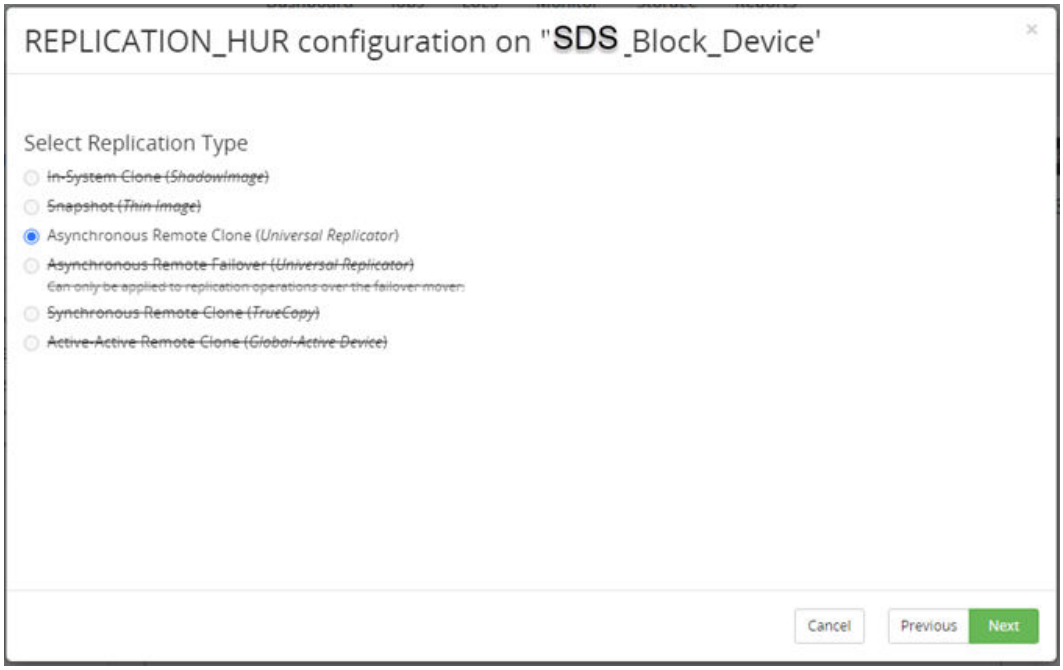


Figure 217 Replicate Configuration Wizard - Configure New Replication

Control	Description
Replication Type	Select the type of replication to create: ▪ Asynchronous Remote Clone.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

This wizard is displayed when you assign a replication operation to a Hitachi Software Block Device node on a data flow.

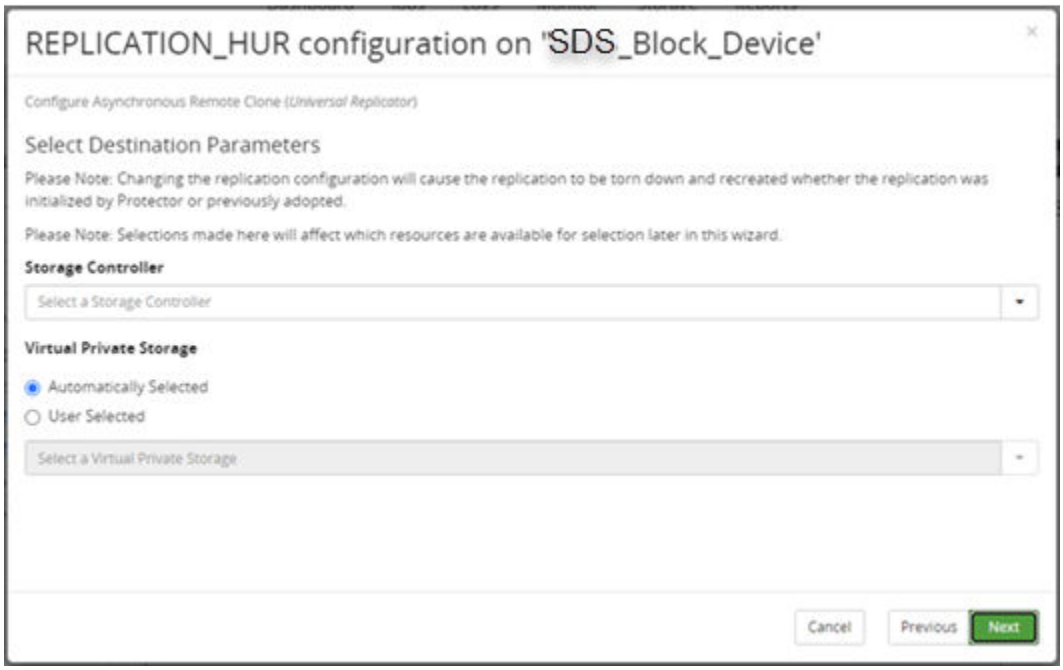


Figure 218 Software Block Replicate Configuration Wizard - Asynchronous Remote Clone (Universal Replicator) – Destination Parameters

Control	Description
Storage Controller	Select the relevant storage controller from the drop down list.
Virtual Private Storage	If left to automatically selected Protector will use the VPS that has been configured on the Software Block Device. If desired, click User Selected to manually choose the VPS.
Cancel	Discards all changes and reverts to the previous page.

Control	Description
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

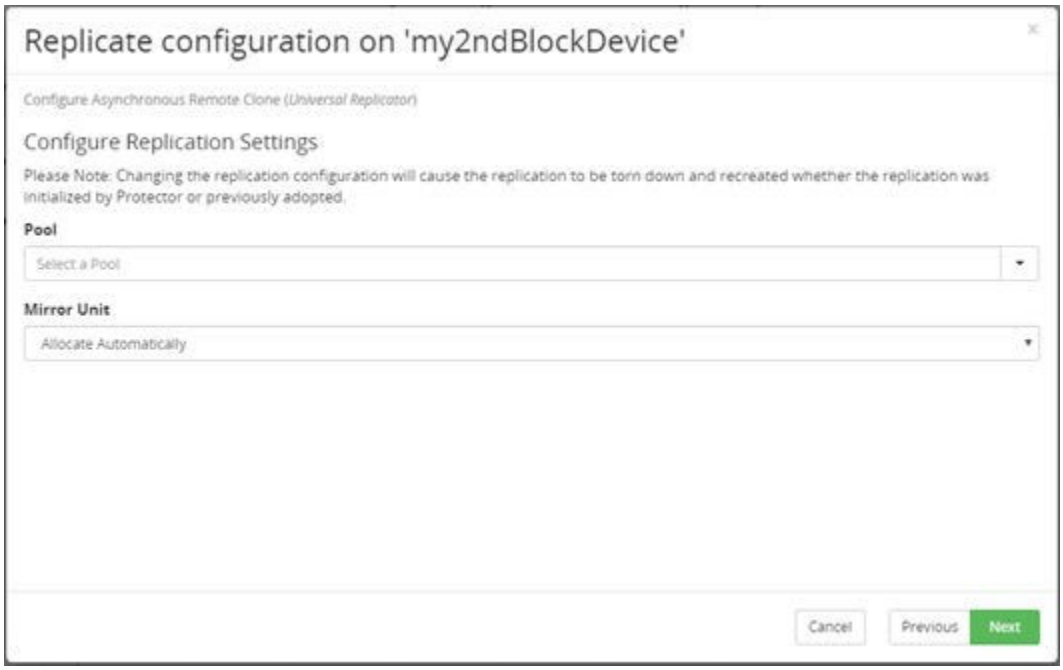


Figure 219 Software Block Replicate Configuration Wizard - Asynchronous Remote Clone (Universal Replicator) - Configure Pool and Mirror Unit

Control	Description
Pool	Specifies the target storage pool from which replication LDEVs are allocated. Provides a list of available pools giving name and available space. **Note:** All replication types have pool except Asynchronous Remote Failover (Universal Replicator). Select a Dynamic Provisioning Pool. **Note:** Dynamic Provisioning Pools must be created using Storage Navigator prior to selecting the Target Storage/Pool in Protector.
Mirror Unit	For a software defined storage array as a replications target the only supported Mirror Unit is 0, as such this setting will default to 0 and can not be changed.

Control	Description
	Select Allocate Automatically to allow Protector to choose one. **Note:** If the mirror unit of an active replication is changed after initial data flow activation then: 1. S-VOLs and pairing relationships for the replication will be destroyed (or dissociated if previously adopted). 2. The replication will then be recreated (or readopted if previously adopted) on data flow reactivation. A warning is issued by the rules compiler prior to activation.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

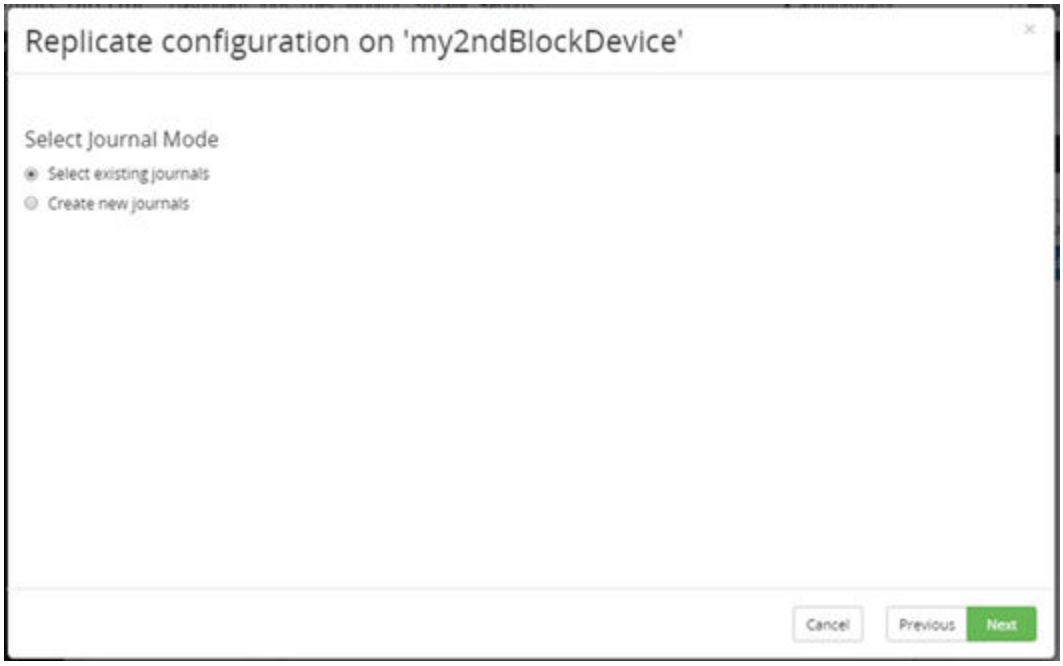


Figure 220 Software Block Replicate Configuration Wizard - Asynchronous Remote Clone (Universal Replicator) - Select Journal Mode

Control	Description
Select existing journals	Select this option to use journals that already exist on the source and destination storage arrays. The **Select existing journals** wizard page is displayed next.

Control	Description
Create new journals	Select this option to have Protector configure new journals on the source and destination storage arrays. The **Create journals** wizard page is displayed next.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

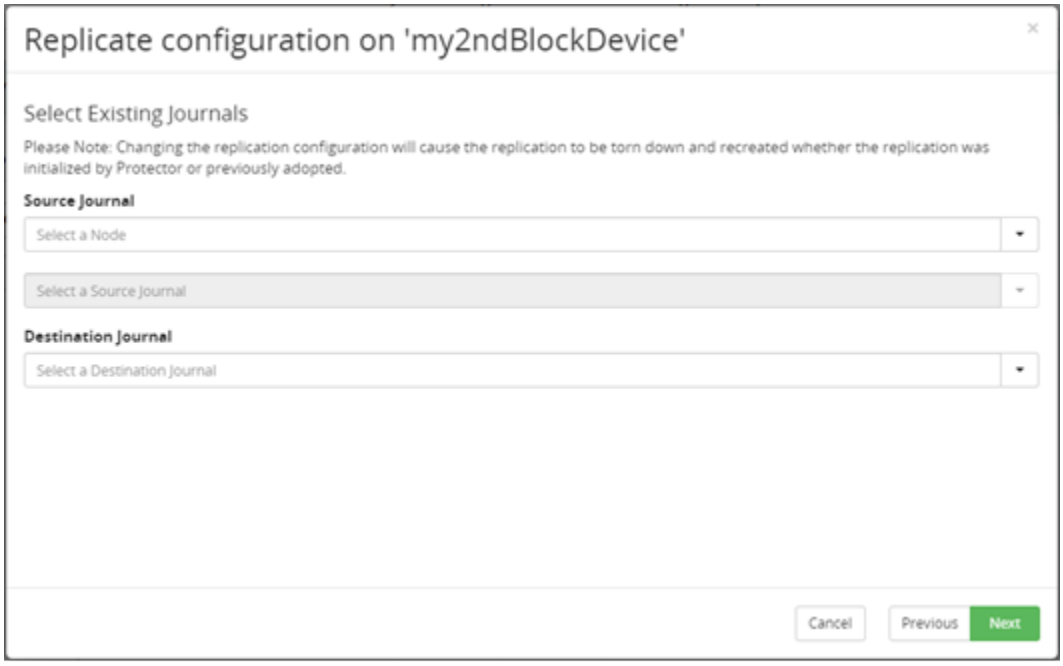


Figure 221 Software Block Replicate Configuration Wizard - Asynchronous Remote Clone (Universal Replicator) - Select existing journals

Note: The journals must be unique to each operation and policy in the data flow. The journal should also be used exclusively by Protector.

Control	Description
Source Journal	Specifies the node and journal on the source side of the replication.
Destination Journal	Specifies the journal on the destination side of the replication.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.

Control	Description
Next	Takes the user to the next screen in the wizard.

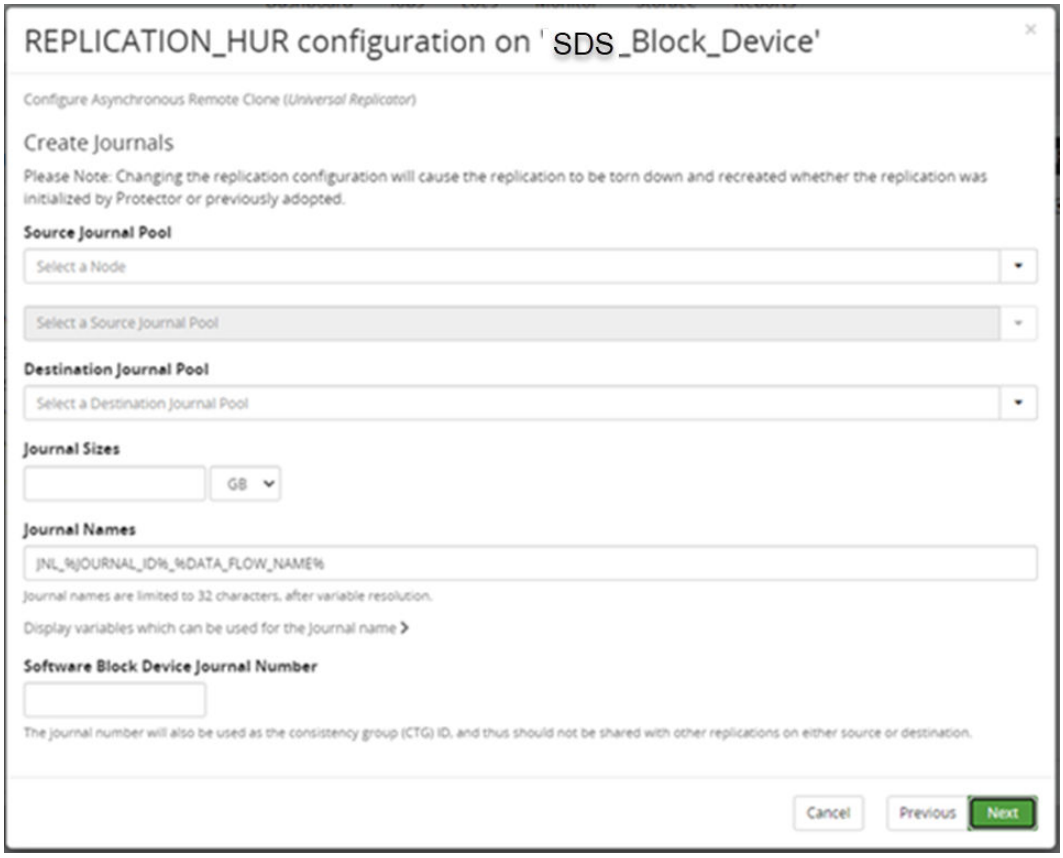


Figure 222 Software Block Replicate Configuration Wizard - Asynchronous Remote Clone (Universal Replicator) - Create journals

Note: The journals must be unique to each operation and policy in the data flow. The journal should also be used exclusively by Protector.

Control	Description
Source Journal Pool	Specifies the node and pool where the source side journal will be created.
Destination Journal Pool	Specifies the pool where the destination side journal will be created.

Control	Description
Journal Sizes	Specifies the size of the source and destination journal. **Note:** Journal sizing is based on data change rate and link bandwidth. Refer to your storage array documentation for details.
Journal Names	The journals will be named using the naming rule provided. The naming rule can consist of literal strings and/or one or more substitution variables listed. Click Display variables which can be used for the journal name to view the available substitution variables.
Software Block Journal Number	Specify a journal number. **Note:** The journal number will also be used as the Consistency Group (CTG) ID and so should not be shared with other remote replications on either source or destination. To determine an unused CTG ID on the array and the correct CTG ID/Journal ID on the SDS please use the ocp-query-storage tool. Please see Query Storage Command Line Utility (on page 967).
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

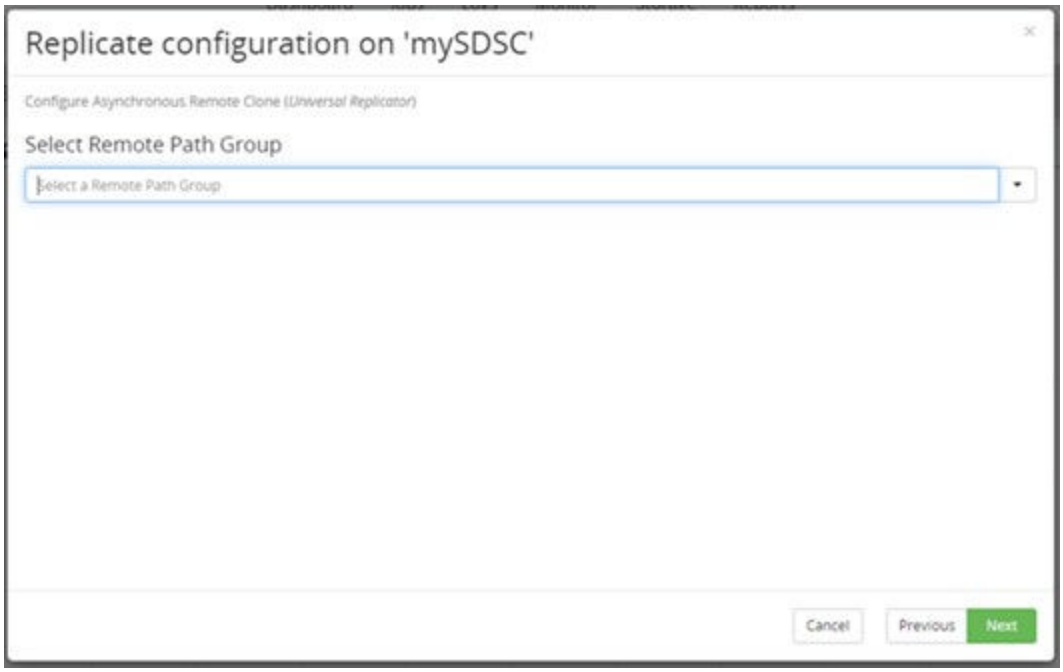


Figure 223 Replicate Configuration Wizard - Asynchronous Remote Clone - Select Remote Path Group

Control	Description
Select Remote Path Group	Specifies the Remote Path Group to be used for the replication. Remote path groups are listed in the format: <Path_Grup_ID>(Remote Serial: <Array_Serial_No>, remote Type: <Array_Type>)
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

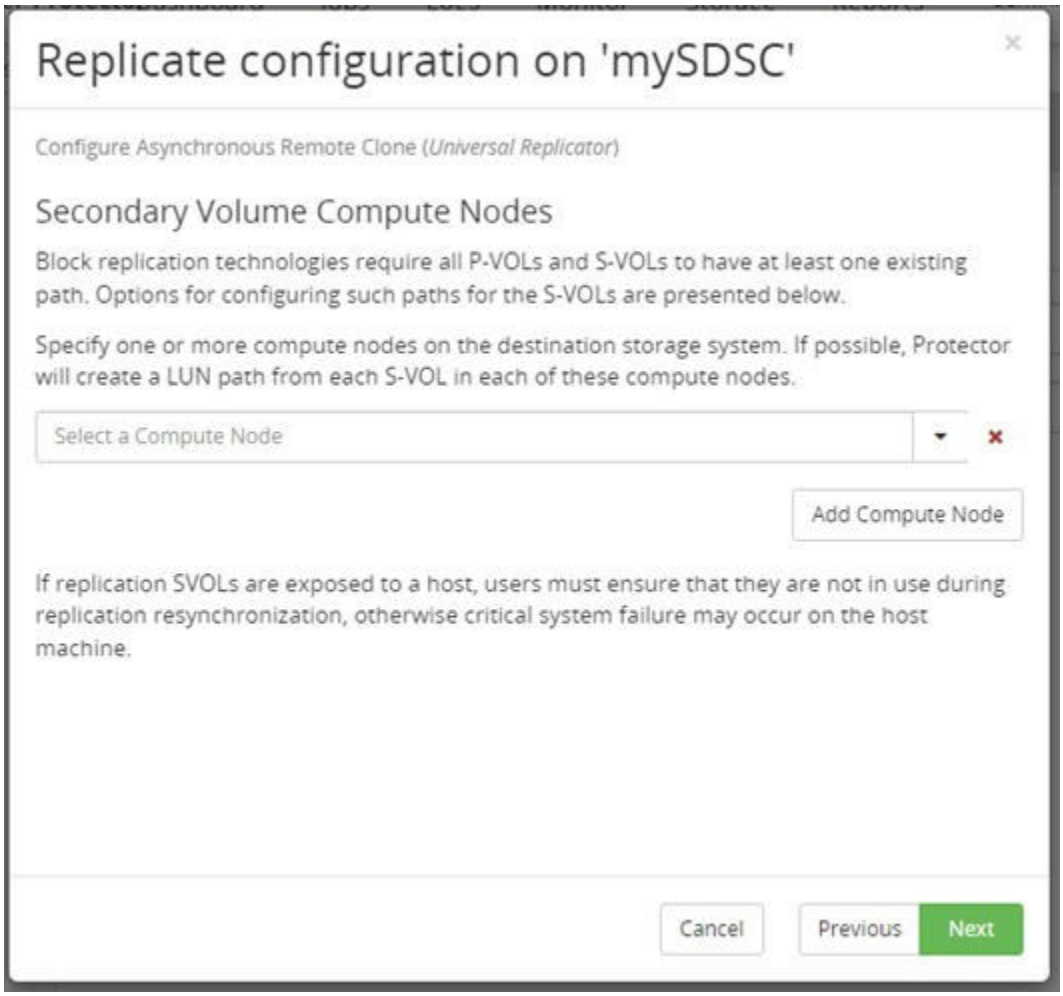


Figure 224 Software Block Replicate Configuration Wizard - Configure Asynchronous Remote Clone

Control	Description
Secondary Volume ComputeNodes	Select one or more compute nodes for the secondary volumes. Adding one node will automatically display the option to add more
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

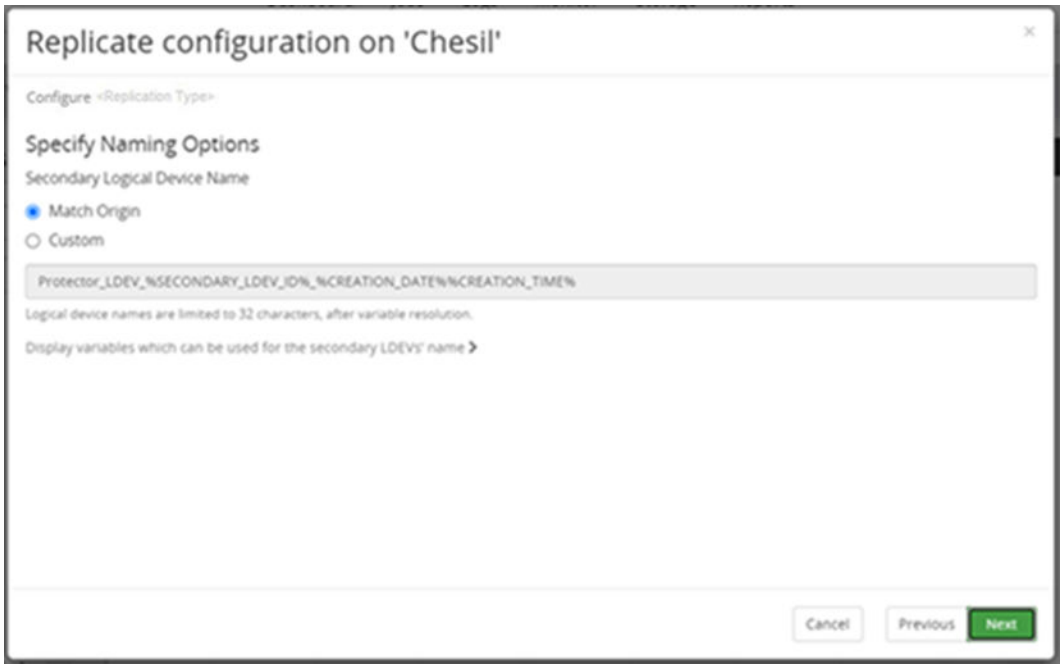


Figure 225 Replicate Configuration Wizard - Naming Options

Control	Description
Secondary Logical Device Name	Specifies how S-VOLs will be named: ▪ Match Origin - The S-VOL will be given the same name as that used for the origin P-VOL (i.e. the left-most volume in the data flow). ▪ Custom - The S-VOL will be named using the naming rule provided. The naming rule can consist of literal strings and/or one or more substitution variables listed. Click Display variables which can be used for the secondary LDEVs' name to view the available substitution variables: • `%ORIGIN_SERIAL%` - S/N of leftmost array in data flow. E.g. output string: `210613` • `%ORIGIN_LDEV_ID%` - ID of leftmost LDEV in data flow. E.g. output string: `00:3A:98` • `%ORIGIN_LDEV_NAME%` - name of leftmost LDEV in data flow. • `%PRIMARY_SERIAL%` - S/N of primary array in this operation. E.g. output string: `442302` • `%PRIMARY_LDEV_ID%` - ID of primary LDEV in this operation. E.g. output string: `00:4C:EB` • `%PRIMARY_LDEV_NAME%` - name of primary LDEV in this operation. • `%SECONDARY_SERIAL%` - S/N of secondary array in this operation. E.g. output string: `356323` • `%SECONDARY_LDEV_ID%` - ID of secondary LDEV in this operation. E.g. output string: `01:F4:35` • `%CREATION_DATE%` - date secondary LDEV was created by this operation. E.g. output string: `20180427` • `%CREATION_TIME%` - time secondary LDEV was created by this operation. 1130
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

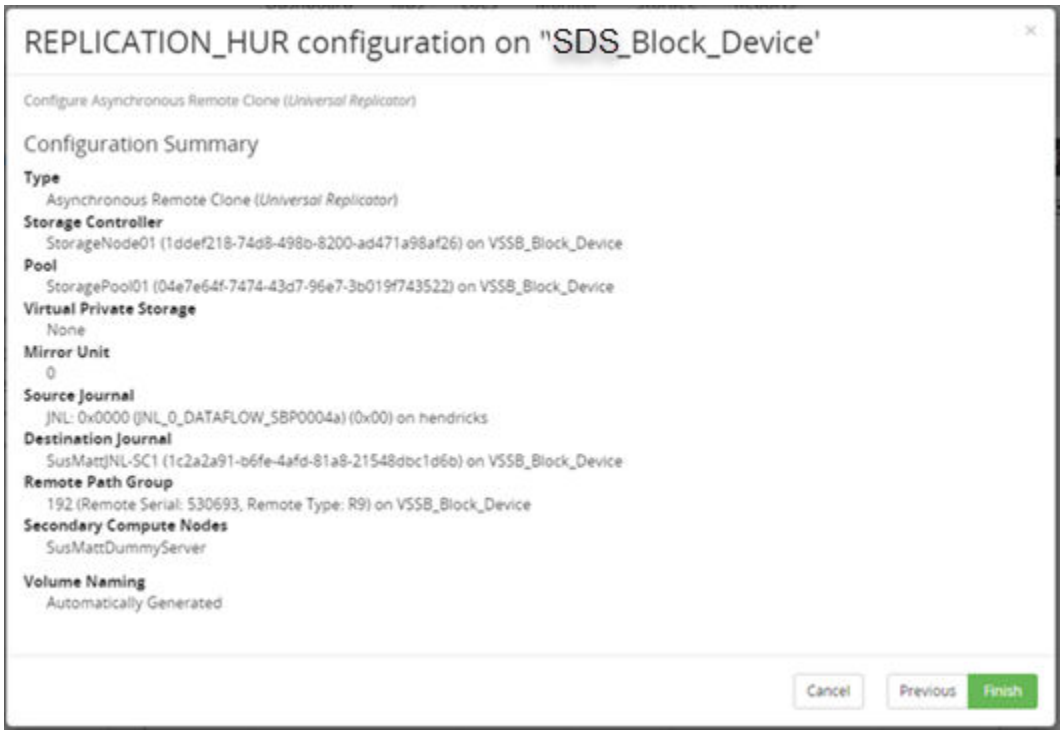


Figure 226 Software Block Replicate Configuration Wizard - Asynchronous Remote Clone (Universal Replicator) - Summary

Shows a summary of the replication configuration specified by the user.

Hitachi Block Mount Configuration Wizard

This wizard is displayed when you assign a mount operation to a Hitachi Block storage node on a data flow.

Note:

- When mounting a snapshot that contains a mounted subdirectory, the subdirectory will be mounted as expected. However, the volume referenced by the subdirectory will also be mounted as a separate drive. Unmount will unmount both the expected and additional mounts.
- The automated mount operation is not suitable for Oracle ASM. The disks are be presented to the OS but need to be manually mounted.

Note: Operating system specific behaviour:

OS	Note
Linux	When mounting a Linux snapshot to a different Linux machine; in order for the user and group names to be displayed correctly the users and groups must have the same ID's as the source.
SUSE Linux	SUSE Linux is not able to perform automated mount operations if hosted on VMware. (RHEL and OEL Linux work as expected).
AIX	The system command **`importvg`** is invoked by Protector to mount snapshots to the user specified location. **`importvg`** creates a directory for the user specified location plus an empty directory corresponding to the original mount point. Neither of these directories are removed by Protector when the snapshot is eventually unmounted, although neither will contain any data.

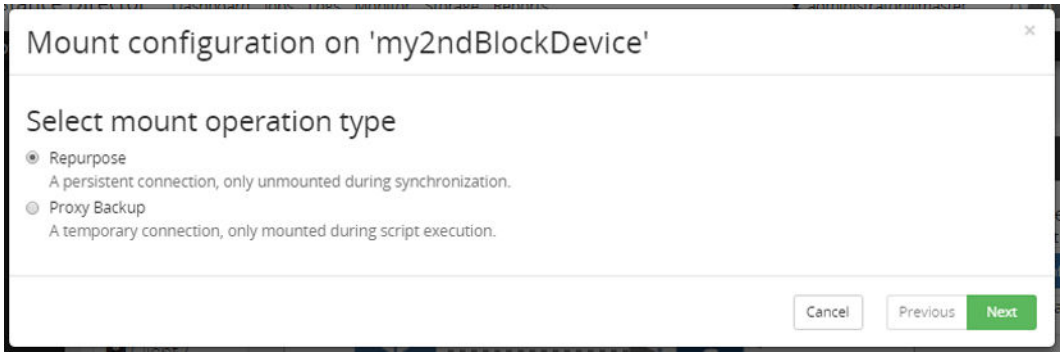


Figure 227 Mount Configuration Wizard - Mount Operation Type

Control	Description
Repurpose	Perform the mount sequence for the repurposing scenario (refer to About the repurposing mount sequence (on page 100) for details). **Note:** Repurpose is not valid for continuous replication; the rules compiler will issue an error if a continuous mover is used on the data flow in conjunction with this type of mount operation.

Control	Description
Proxy Backup	Perform the mount sequence for the proxy backup scenario (refer to About the proxy backup mount sequence (on page 101) for details). **Note:** A proxy backup of a live replication will generate a warning in the compiler to tell the user that their replication will be paused until the proxy backup is complete.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

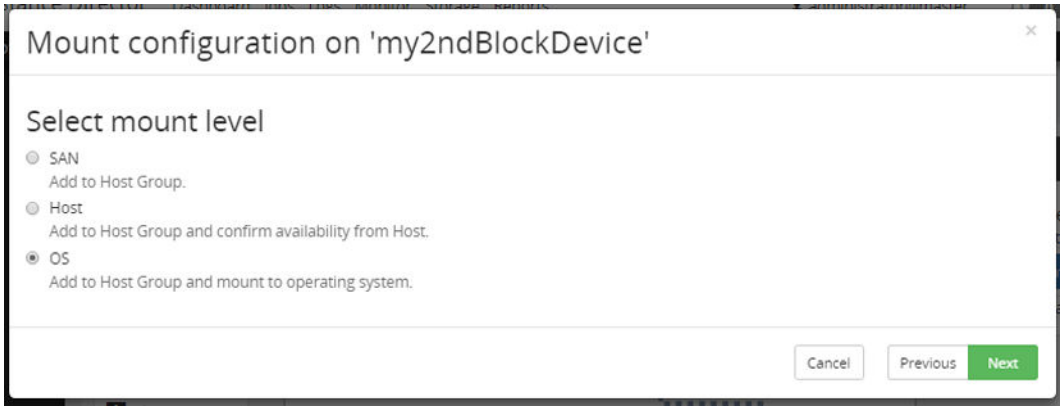


Figure 228 Mount Configuration Wizard - Mount Level

Control	Description
SAN	Adds the replication to a Host Group.
Host	Adds the replication to a Host Group and confirms that it is available from the specified Host.
OS	Adds the replication to a Host Group and mounts it on the specified Host's operating system.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

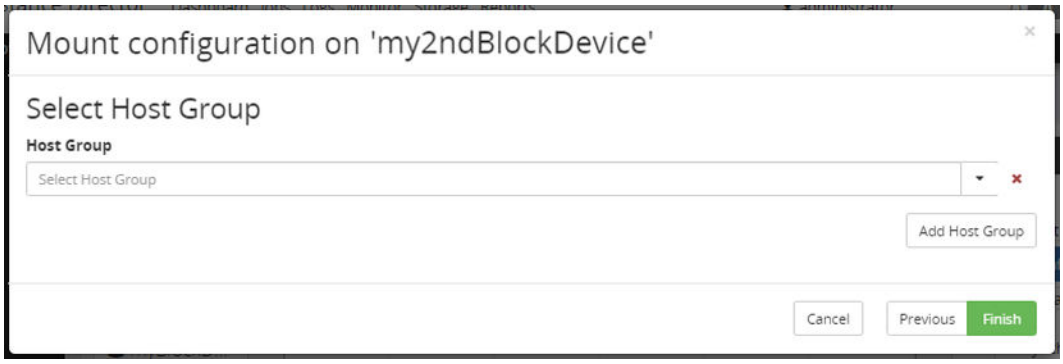


Figure 229 Mount Configuration Wizard - Select Host Group (SAN level mount only)

Control	Description
Host Group	Manually specify or select a host group to use to expose a snapshot or replication. **Note:** When exposing an LDEV, the host group specified must be in the same resource group as the secondary volumes.
Add Host Group	Click this button to add host groups when specifying a multi-path mount operation.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Finish	Commits the new changes. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

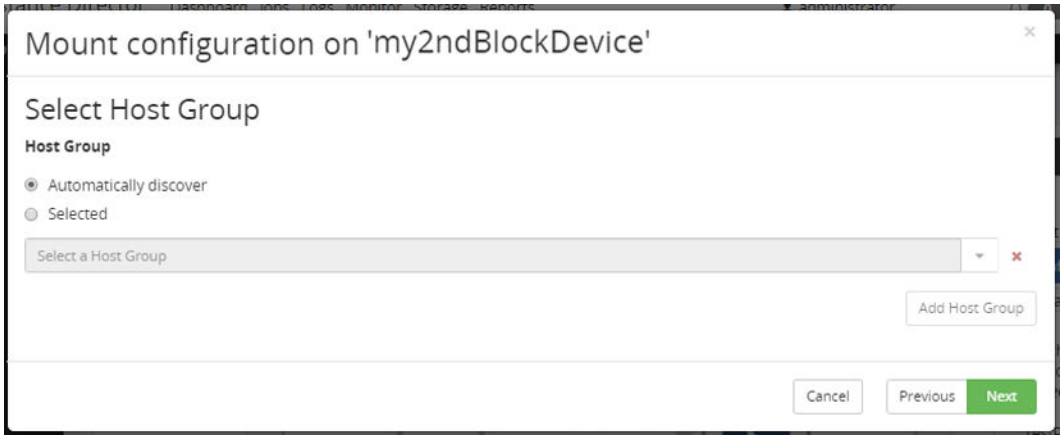


Figure 230 Mount Configuration Wizard - Host Group (Host and OS level mount only)

Control	Description
Automatically discover	Protector will automatically select a host group to use to expose the snapshot or replication.
Selected	The user must specify one or more host groups to use to expose the snapshot or replication.
Select a Host Group	Manually specify or select a host group to use to expose a snapshot or replication. **Note:** When exposing an LDEV, the host group specified must be in the same resource group as the secondary volumes.
Add Host Group	Click this button to add host groups when specifying a multi-path mount operation.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

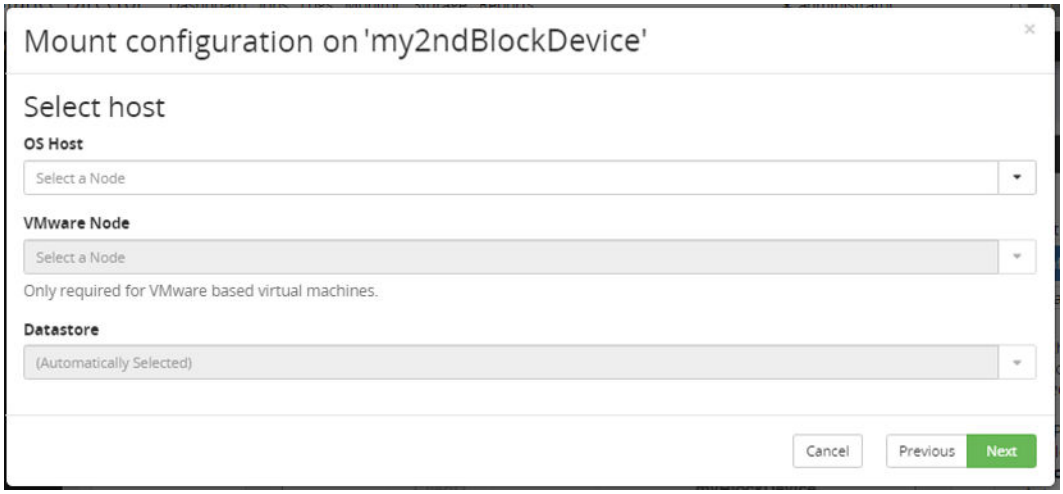


Figure 231 Mount Configuration Wizard - Select Host (Host and OS level mount only)

Control	Description
OS Host	Specify the machine to mount to or expose to. **Note:** Unless the user selects a host group, the machine where the volume is to be mounted must have an existing volume on the same storage device. If there is no connection between the mount host and the block storage device then Protector will fail the mount operation after a timeout of 30 minutes.
VMware host	Expose the volumes to the specified VMware host to enable them to be mounted to the VM. **Note:** Exposing using a VMware host requires that a VMware vCenter node be configured.
Datastore	Specifies a destination datastore when mounting to a VMware virtual machine which is part of a cluster, in which case the default datastore may not be a suitable place to save the RDM mount information. If the datastore field is left blank then mount information is saved alongside the VM.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Next	Takes the user to the next screen in the wizard.

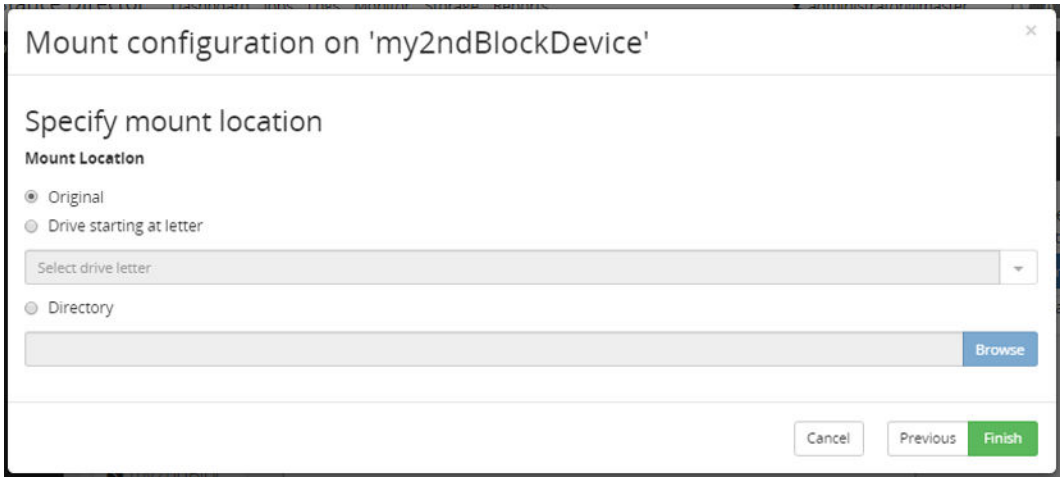


Figure 232 Mount Configuration Wizard - Specify mount location (OS level mount only)

Control	Description
Original	The replication is mounted at its original location. **Note:** Mounting at the original location will fail if there is already a volume mounted at that location.
Drive starting at letter	When mounting a replication that contains multiple volumes, the first volume will mount at the specified drive and subsequent drives are used for each additional volume.
Directory	When mounting a replication that contains multiple volumes, each volume will be assigned a separate subdirectory. Click Browse to view the drives and directories on the selected host. To create a new directory, type in the required path. **Note:** Protector does not check to make sure the directory selected as the mount point is empty. This means it is possible to mount a snapshot inside or even over the top of another mounted volume. This should be avoided.
Cancel	Discards all changes and reverts to the previous page.
Previous	Takes the user to the previous screen in the wizard.
Finish	Commits the new changes. Pages currently open in other tabs and windows will need to be reloaded before the changes are seen by the user.

Data Flow Details

This page displays the details of a Data Flow and enables you launch the wizard to edit, activate, deactivate and change access permissions.

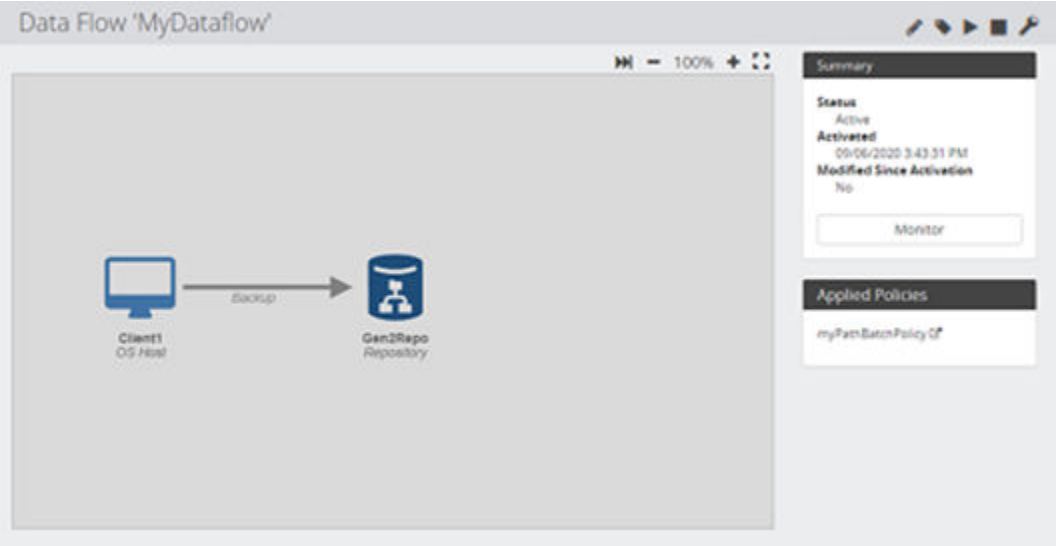


Figure 233 Data Flow Details

Control	Description
Edit	Launches the Data Flow Wizard (on page 361) to enable you to edit the data flow.
Tag	Modifies the tags of an existing object from the either the inventory screen or the details screen of the object.
Compile	Displays the Activate Data Flow Dialog (on page 357) and attempts to compile the rules for the data flow. If compilation is successful then the rules can be activated.
Deactivate	Deactivates the selected data flow and removes its rules. **Caution:** If the deactivated data flow contains storage hardware based operations, this will remove the pairing relationships.
Edit Permissions	Displays the Access Control Permissions Inventory (on page 349) to enable you to view and edit the data flow's permissions.
Data Flow Canvas	Shows the Data Flow diagram in read-only mode.
Applied Policies	The area to the right of the workspace lists all the policies that have been applied in the data flow. Click the node or mover of interest to view the applied policies and mover settings.

Control	Description
Summary	The area to the right of the workspace lists the Status of the data flow.
Status	The Status of the data flow can be Active or Inactive.
Activated	If the data flow is Active then Activated shows the date and time the data flow was activated.
Modified Since Activation	If the data flow is Active then Modified Since Activation shows if the data flow has been modified since it was activated, this can be either Yes or No. If the value is Yes then the currently distributed rules are different to those shown in the data flow.
Monitor	Displays the Monitor Details (on page 521) screen for this data flow

Jobs UI Reference

This section describes the Jobs UI, accessed via the Main Banner (on page 286).

For further information, refer to:

- Job Concepts (on page 109)
- Job Tasks (on page 257)

Jobs Inventory

The Jobs Inventory displays all currently running or recently completed jobs, for the purpose of operation confirmation and control.

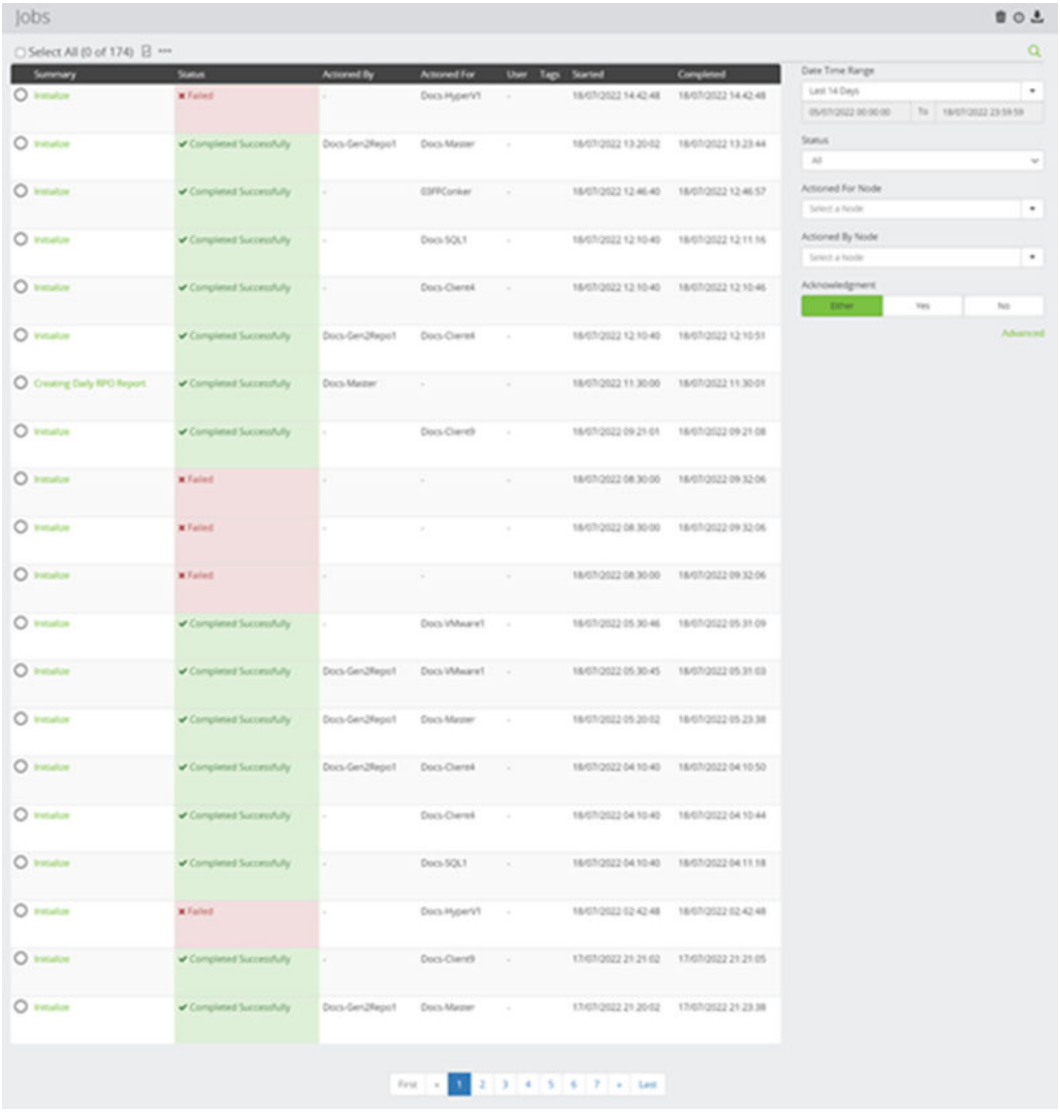


Figure 234 Jobs Inventory

Control	Description
Purge Jobs	Removes all jobs of a given age and older from the job manager database. The Purge Jobs Dialog (on page 496) is displayed prior to the jobs being purged.
Edit purge schedule	Allows the automatic job purge schedule to be configured. The Edit Job Purge Schedule (on page 497) is displayed.

Control	Description
 Export Jobs	Exports the job database as a file, respecting the currently specified filter terms. The Export Jobs Dialog (on page 498) is displayed prior to the jobs being exported.
 Acknowledge Selected Logs	Enabled only if one or more job entries are selected. Failed jobs will indicate that they have been acknowledged by including "Acknowledged" in their "Progress" column. The Acknowledge Jobs Dialog (on page 497) is displayed prior to the jobs being acknowledged.
 More Actions	Only enabled if the selected job is actionable. A menu appears with the available actions listed which can include: ▪ Cancel - *Restore*, *Resync* and *Software Update* jobs. ▪ Pause - *Replication* jobs. ▪ Resume- *Replication* jobs.
Summary	Indicates the type of action being performed.
Status	Indicates the status of a job as follows: ▪ In Progress ▪ Paused ▪ Cancelled ▪ Completed Successfully ▪ Failed
Node	The node for which the job is being actioned for. Note that not all job will have a node specified.
User	User who initiated the job if not a system initiated job.
Tags	List of tags associates with the job see Universal Tags (on page 39) for more detail.
Started/ Completed	Shows the data and time the job started/completed.
Filter on Date Time Range	Filters the job table so that only jobs from the specified Master date and time are displayed. Opens the Date Time Range Picker (on page 352).
Filter on Actioned By Node	Filters the job table so that only jobs relating to the specified node that performs an action are displayed.
Filter on Actioned For Node	Filters the job table so that only jobs relating to the specified node that requests an action are displayed.

Control	Description
Filter on Acknowledgement	Filters the job table so that only acknowledged or unacknowledged jobs are displayed.

Purge Jobs Dialog

The Purge Jobs dialog is displayed when the user attempts to remove jobs from the job database.

Caution: Once purged, the job cannot be recovered, so be careful not to purge important jobs. Consider exporting jobs before purging if in any doubt.

Note: Purging the jobs removes the job information from the database but does not reduce its overall size.

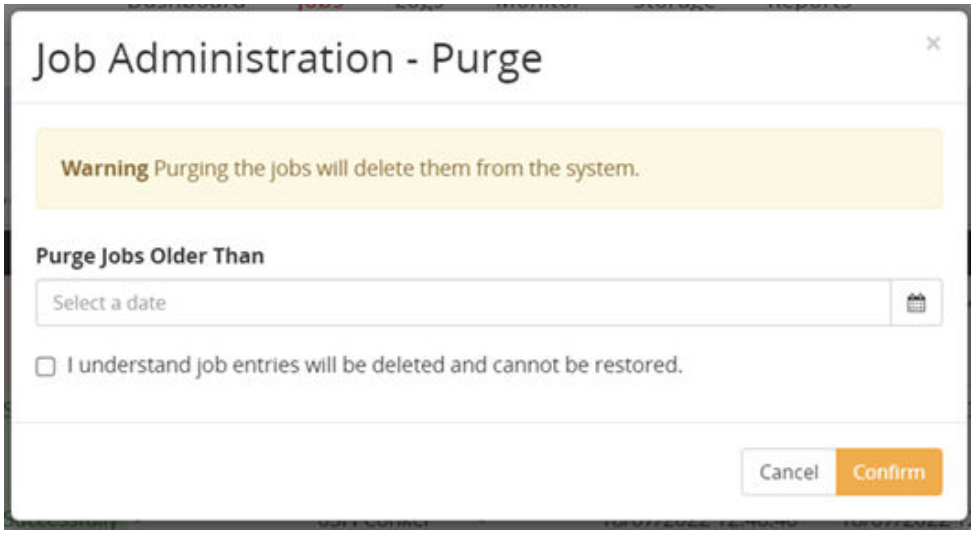


Figure 235 Purge Jobs Dialog

Control	Description
Purge Jobs Older Than	Specify the date and time, using the Date Time Picker (on page 351), on and before which jobs must originate for them to be deleted from the job database. Jobs newer than this date and time will not be deleted.
I understand job entries...	As an interlock, this checkbox must be ticked in order to proceed with the purge operation. The Confirm button will not be enabled unless the checkbox is ticked.

Edit Job Purge Schedule

The Job Purge Schedule dialog is displayed when the user clicks on the Edit Job Purge Schedule icon.

The Edit Job Purge Schedule dialogs allows a user to enable or disable the automatic purging of jobs. If enabled the user configures the maximum number of weeks jobs are kept for.

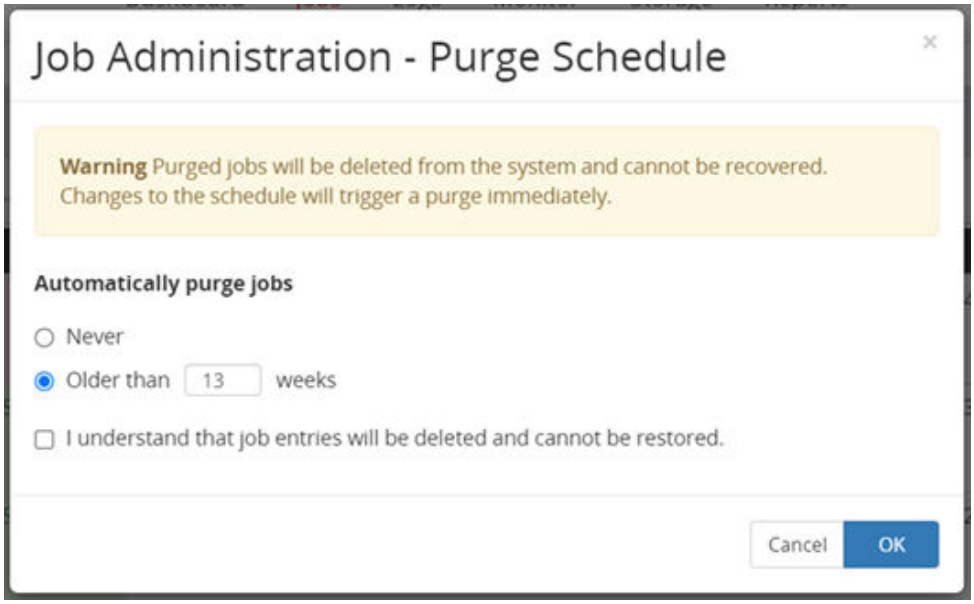


Figure 236 Job Purge Schedule Dialog

Control	Description
Never	Disable automatic job purging.
Older than *13* weeks	Enable automatic job purging causing job messages which are older than, or become older than, the specified number of weeks to be permanently deleted.
Acknowledgement checkbox	To enable the automatic purging feature, the user must acknowledge that purged jobs are not recoverable.

Acknowledge Jobs Dialog

The Acknowledge Jobs dialog is displayed when the user attempts to acknowledge a job in the job database.



Figure 237 Acknowledge Jobs Dialog

Control	Description
Acknowledgment	Optional. Enter a short message that describes the action taken or reason for acknowledging the job. This will be appended with the date and time of acknowledgment upon confirmation. Drag the handle at the bottom right of the edit box to expand it.

Export Jobs Dialog

The Job Export dialog is displayed when the user clicks on the export icon on the Jobs Inventory.

Note: Only jobs currently displayed are exported. Ensure that you have the correct filters applied to the jobs before exporting them.

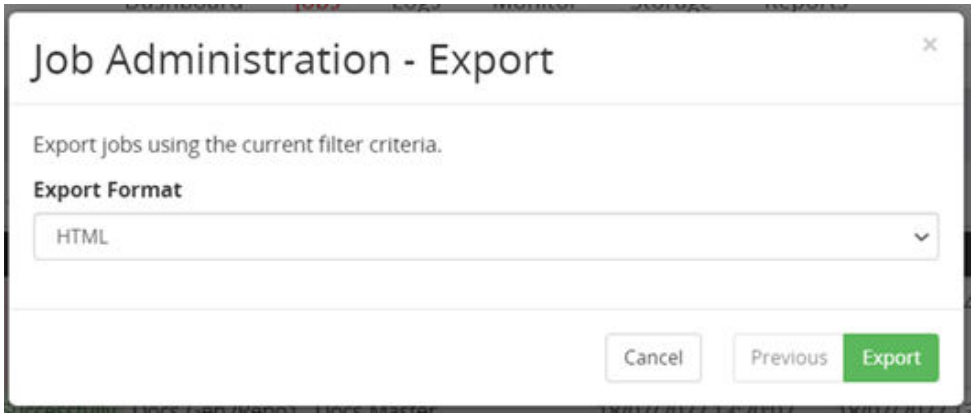


Figure 238 Export Jobs Dialog

Control	Description
Export Format	Select one of the following: ▪ HTML ▪ CSV
Export	Click to begin the export process. The dialog will display progress of the export process and the export file will automatically be downloaded when ready. The export will be automatically downloaded even if the dialog is closed.

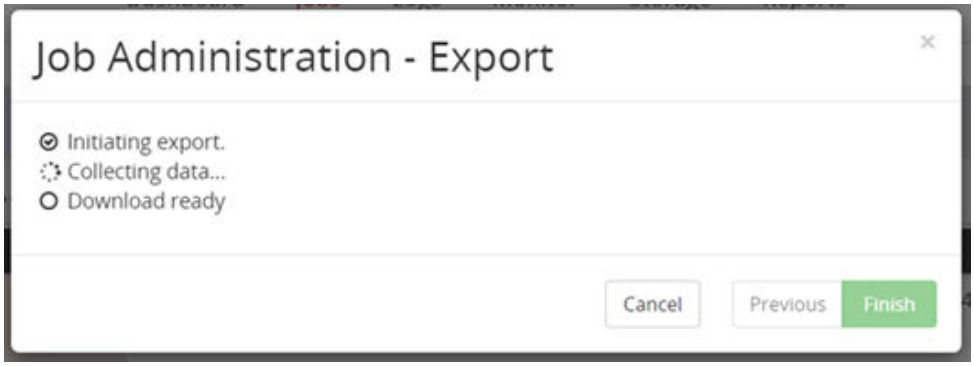


Figure 239 Export Jobs Dialog – Export in progress

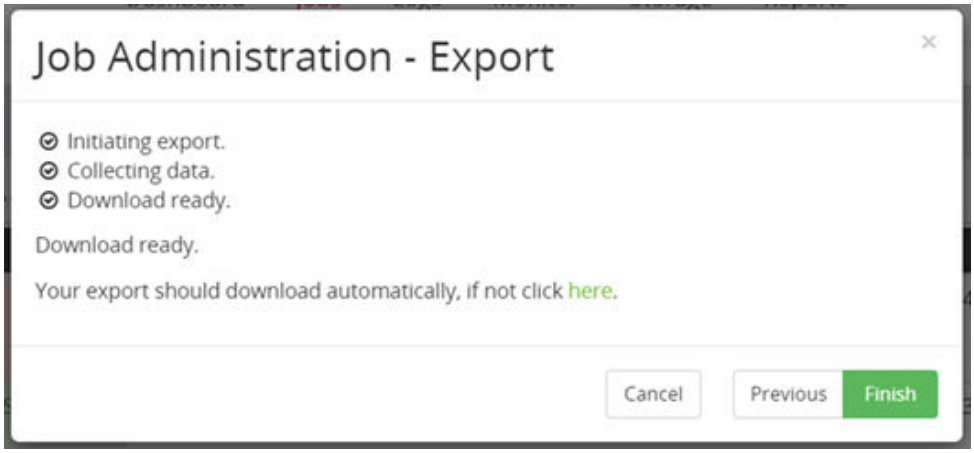


Figure 240 Export Jobs Dialog - Export completed

Job Details

This page displays the details of a Job.

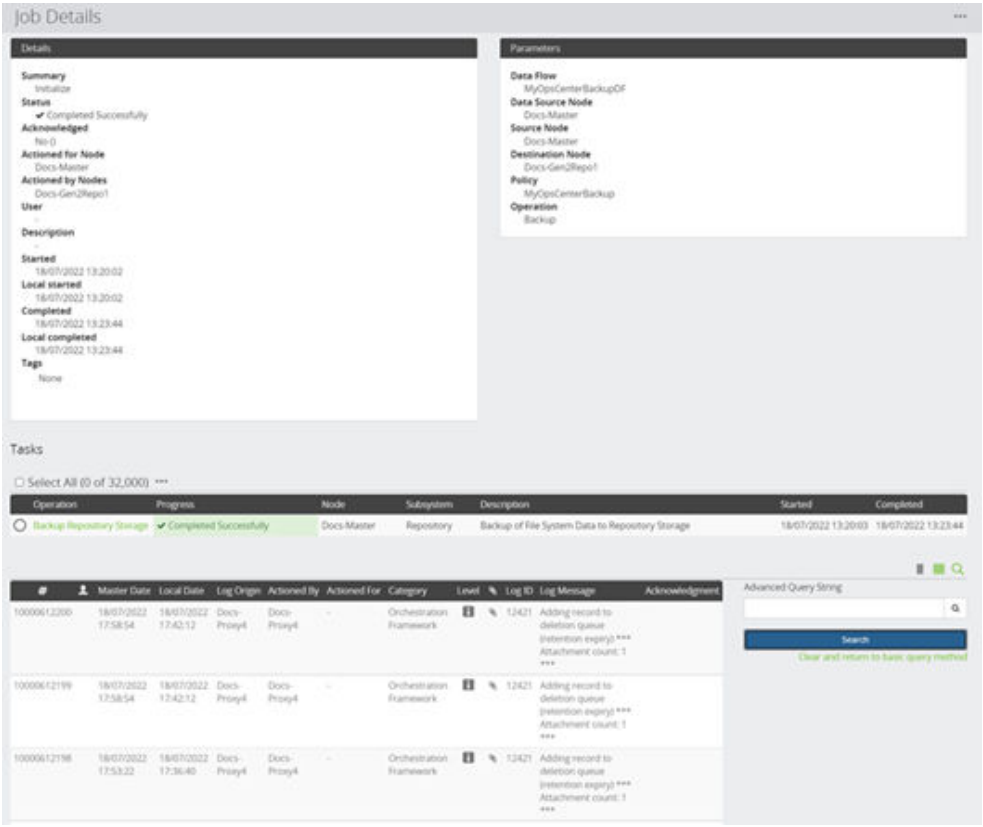


Figure 241 Job Details

Control	Description
●●● More Actions	Only enabled if the selected job is actionable. A menu appears with the available actions listed which can include: ▪ Cancel - *Restore*, *Resync* and *Software Update* jobs. ▪ Pause - *Replication* jobs. ▪ Resume- *Replication* jobs.
Details	Provides detailed information about the Job. For *Restore* and *Resync* jobs a *Progress* section is added that provides information about theprogress of the job. Refer to Monitor Details (on page 521). Once a job has completed the progress column will state Completes Successfully, Completed with Errors or Failed. For any status other than Completed Successfully refer to the logs section to determine the error.
Parameters	Identifies the Source, Destination, Policy and Data Flow associated with the Job.
Tasks	Identifies the Tasks associated with the job in action. For more details refer to Tasks Inventory (on page 501).

Control	Description
Logs	List log messages relating to the job. See Logs UI Reference (on page 509) for more details.
 Condensed View	Displays a reduced number of parameters in the log table.
 Extended View	Displays all available parameters in the log table.
 Show/ Hide Search	Toggles the display of the filtering options.
Filter on Log Level	Filters the displayed log results based on Log Level.

Note: If job details are viewed whilst the job is active additional progress is show. Whats displayed is dependent on the application and destination type.

Tasks Inventory

The Task Inventory displays all tasks in the system that belong to a job depending on the filter criteria.

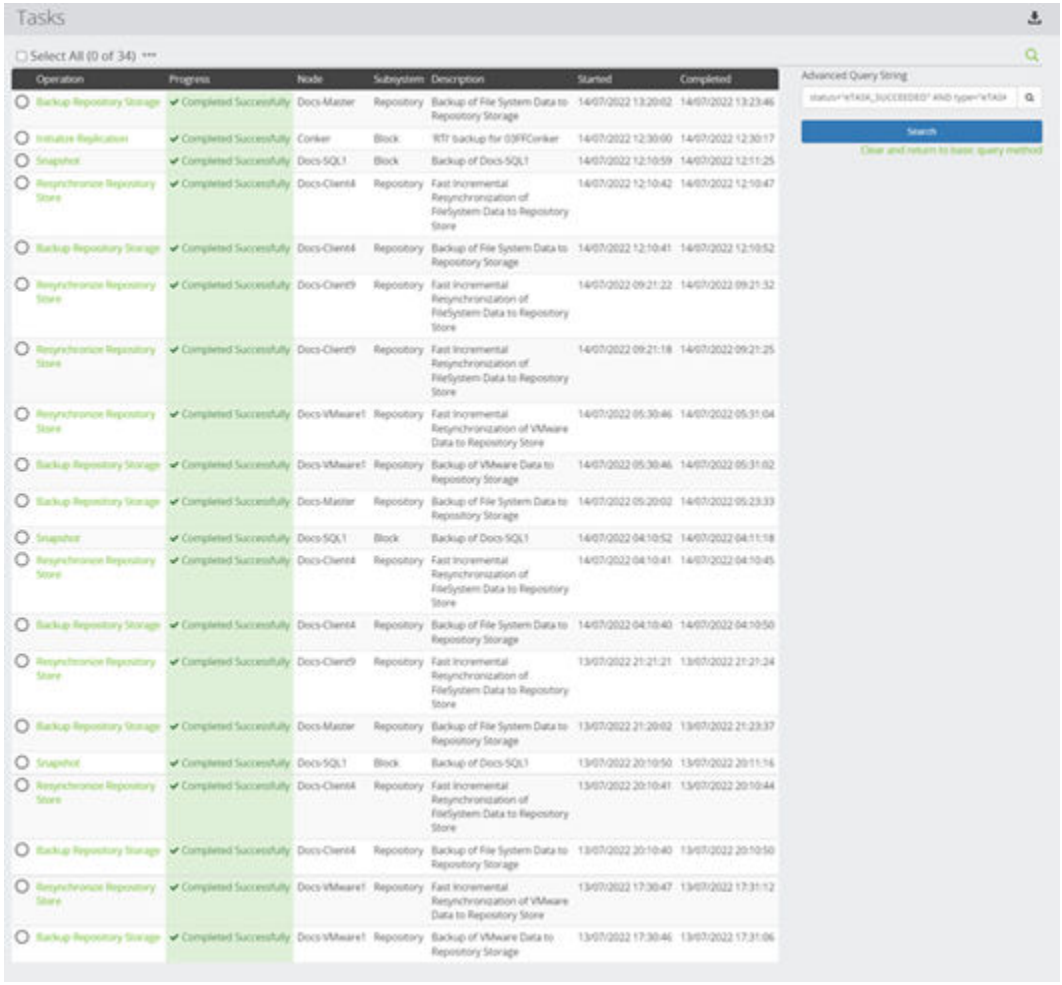


Figure 242 Tasks Inventory

Control	Description
●●● More Actions	Only enabled if the selected task is actionable. A menu appears with the available actions listed which can include: ▪ Cancel - *Restore*, *Resync* and *Software Update* tasks. ▪ Pause - *Replication* tasks. ▪ Resume- *Replication* tasks.
Operation	Click on the Operation link in the tasks list to open the Task Details (on page 504) for the selected task.
Progress	Indicates the status of a task as follows: ▪ In Progress ▪ Paused ▪ Cancelled

Control	Description
	▪ Succeeded ▪ Failed
Node	The node for which the task is being actioned for.
Subsystem	The part of the product responsible for the task being initiated.
Description	Description of the task being carried out.
Started/ Completed	Shows the data and time the task started/completed.
Filter on Date Time Range	Filters the task table so that only tasks from the specified Master date and time are displayed. Opens the Date Time Range Picker (on page 352).
Filter on Status	Filters the displayed results based on Status.
Filter on Node	Filters the displayed results based on Node name.

Export Tasks Dialog

TheTask Export dialog is displayed when the user clicks on the export icon on the Tasks Inventory page.

Note: Only tasks currently displayed are exported. Ensure that you have the correct filters applied to the tasks before exporting them.

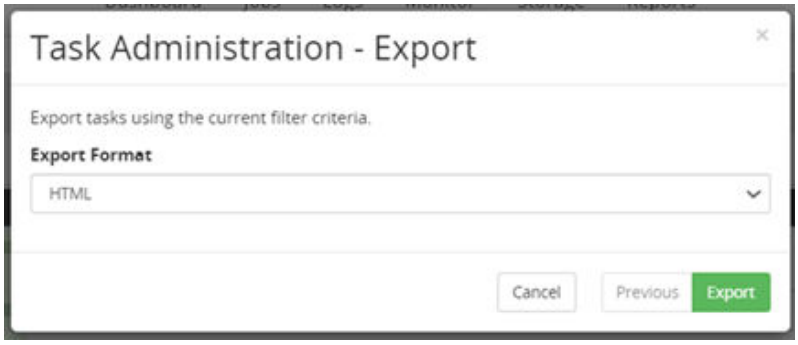


Figure 243 Tasks Export Dialog

Control	Description
Export Format	Select one of the following: ▪ HTML ▪ CSV

Control	Description
Export	Click to begin the export process. The dialog will display progress of the export process and the export file will automatically be downloaded when ready. The export will be automatically downloaded even if the dialog is closed.

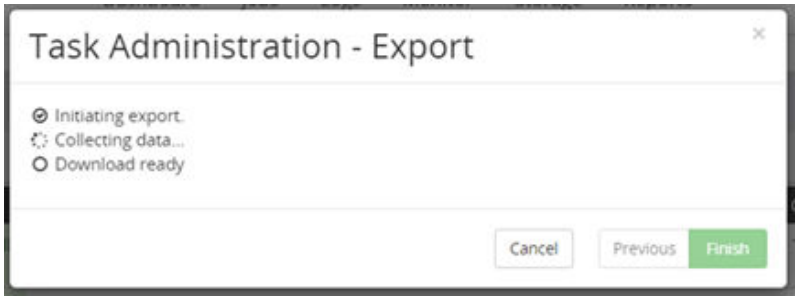


Figure 244 Tasks Export Dialog – Export in progress

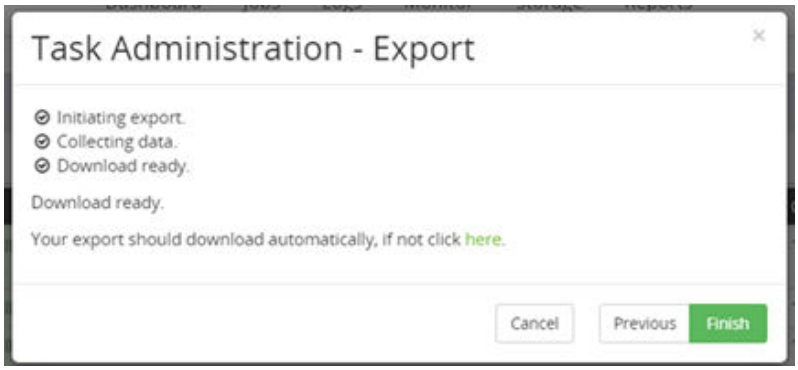


Figure 245 Tasks Export Dialog - Export completed

Task Details

This page displays the details of a Task.

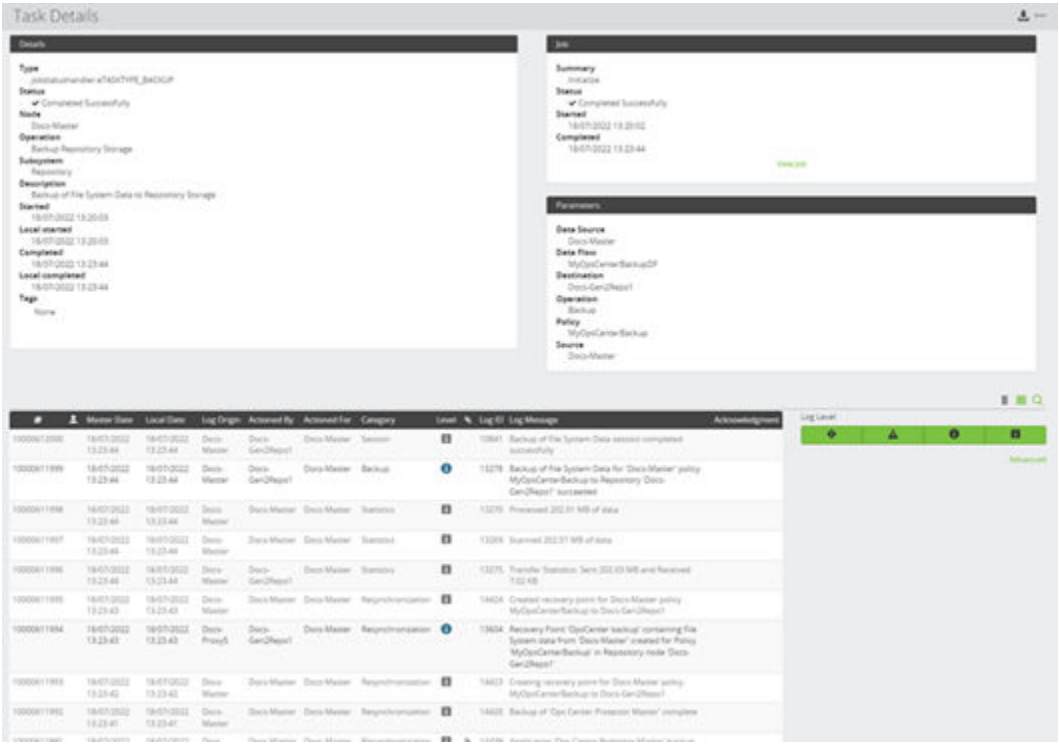


Figure 246 Task Details

Control	Description
Export Session Tasks	Exports the task database entries for the displayed session as a file, respectingthe currently specified filter terms. The Export Tasks Dialog (on page 503) is displayed prior to the tasks being exported.
More Actions	Only enabled if the selected task is actionable. A menu appears with the available actions listed which can include: ▪ Cancel - *Restore*, *Resync* and *Software Update* tasks. ▪ Pause - *Replication* tasks. ▪ Resume- *Replication* tasks.
Details	Provides detailed information about the Task. Once a task is completed, the progress column will state Completed Successfully,Completed with Errors or Failed.
Job	Identifies Summary, Status, Started/Completed data. Click on View Job to open the associated Jobs Details.
Parameters	Identifies the Source, Destination, Policy and Data Flow associated with the Task.
Logs	List log messages relating to the task. See Logs UI Reference (on page 509) for more details.

Control	Description
 Condensed View	Displays a reduced number of parameters in the log table.
 Extended View	Displays all available parameters in the log table.
 Show/ Hide Search	Toggles the display of the filtering options.
Filter on Log Level	Filters the displayed log results based on Log Level.

Note: If task details are viewed whilst the task is active additional progress is shown. What is displayed is dependent on the application and destination type.

Licenses UI Reference

This section describes the Licenses UI, accessed via the Navigation Sidebar (on page 291).

For further information, refer to:

- License Concepts (on page 128)
- License Tasks (on page 258)

Licenses Inventory

The License Inventory page displays the existing licenses and enables new licenses to be added.

Note: The Protector software installation includes a built-in evaluation key for sites that want to use the product on a trial-to-permanent basis. After the software is installed, it can be used for 5 weeks before a permanent license key is required.

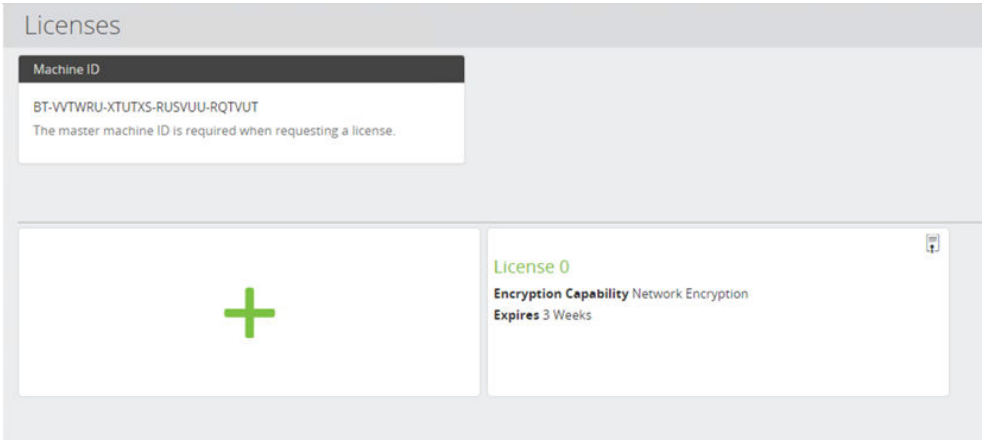


Figure 247 Licenses Inventory

Control	Description
Add	Click this tile to add a new license using the Activate License Wizard (on page 507)
Existing Licenses	Click these tiles to view currently active License Details (on page 509). The figure above shows both Gen2 and Gen3 license types.

Activate License Wizard

A Protector activation key can be entered using this wizard.

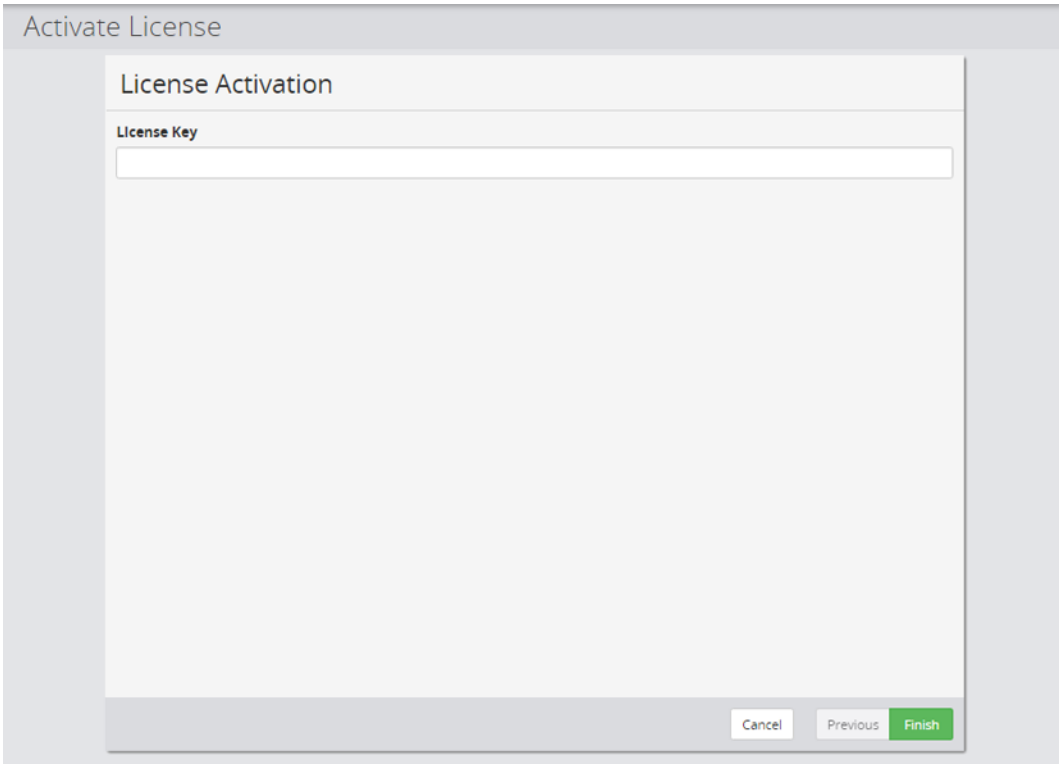


Figure 248 Activate License Wizard

Control	Description
License Key	Paste the new license key here. **Note:** ▪ In order to activate data flows, a valid Protector activation key must be provided. ▪ For cluster environments each machine in the cluster has a different machine ID, therefore a separate license key is required for each node.
Cancel	Cancels the activation process. The new license will not be activated
Finish	Completes the activation process. The new license will be displayed in the Licenses Inventory (on page 506) **Note:** Any node requiring encryption should have the Hub service restarted following the new license being entered. If the node is already authorized, two Hub service restarts may be required.

License Details

This page displays the details of the selected license.



Figure 249 License Details

Control	Description
License ID	A numeric identifier for the license
License Key	The license key provided by Hitachi Vantara
Expires	The time left before the license expires
Encryption capability	Indicates whether network (over the wire) and/or repository (in place) encryption has been enabled

Logs UI Reference

This section describes the Logs UI, accessed via the Main Banner (on page 286).

For further information, refer to:

- Log Concepts (on page 107)
- Log Tasks (on page 259)

Logs Inventory

The Logs Inventory displays system wide events for the purpose of operation confirmation, health monitoring, fault diagnosis and auditing.

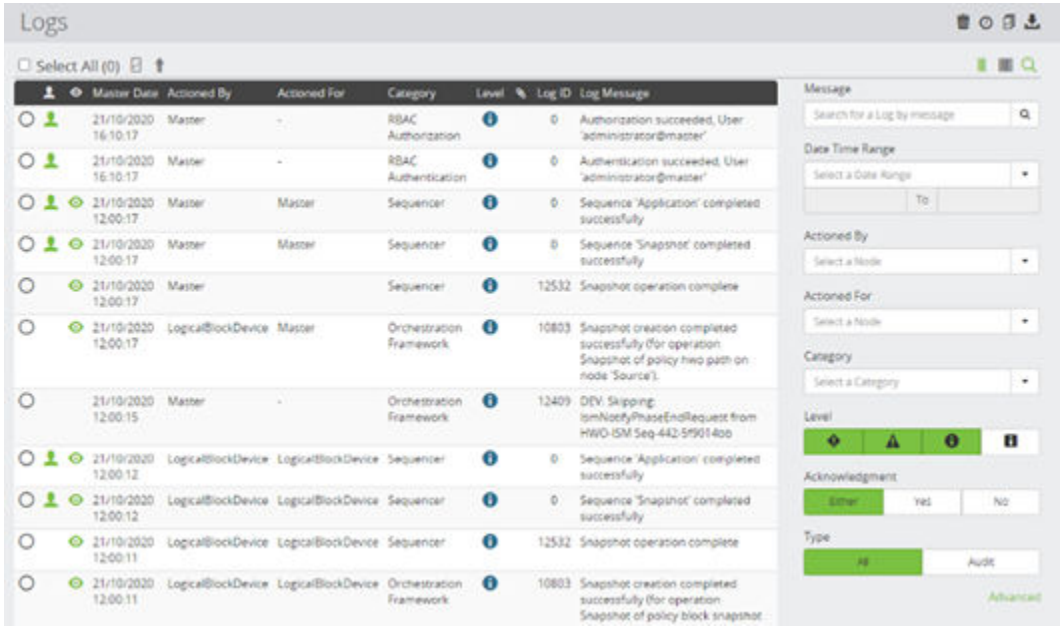


Figure 250 Logs Inventory

Control	Description
Purge Logs	Removes all logs of a given age and older from the log manager database. Logs marked as 'Audit' cannot be removed. The Purge Logs Dialog (on page 513) is displayed prior to the logs being purged.
Edit purge schedule	Allows the automatic log purge schedule to be configured. The Log Purge Schedule (on page 514) is displayed.
Acknowledge All Error Logs	Acknowledges all unacknowledged error logs in the log database. Unacknowledged error logs are displayed with a red background. Acknowledged logs are displayed in green and the acknowledgment message, date and time are displayed alongside the log entry. The Acknowledge Logs Dialog (on page 515) is displayed prior to the logs being acknowledged.
Export Logs	Exports the log database as a file, respecting the currently specified filter terms. The Export Logs Dialog (on page 516) is displayed prior to the logs being exported.
Acknowledge Selected Logs	Enabled only if one or more log entries are selected. Acknowledges all selected logs in the log database. Unacknowledged error logs are displayed with a red background. Acknowledged logs are displayed in green and the acknowledgment message, date and time are displayed alongside the log entry. The Acknowledge Logs Dialog (on page 515) is displayed prior to the logs being acknowledged.

Control	Description
 Create Log Notification	Enabled only if one log entry is selected. Creates a log notification condition for that log message type. The Notification Wizard (on page 662) is displayed to assist in creating the notification condition.
 Condensed View	Display a subset of the log fields in the log table.
 Extended View	Display all of the log fields in the log table.
Log Entry	Refer to the Log Fields table below for a full description of all fields.
Filter on Message	Filters the log table so that only logs containing the specified string in the Log Message field are displayed.
Filter on Date Time Range	Filters the log table so that only logs from the specified Master date and time are displayed. Opens the Date Time Range Picker (on page 352).
Filter on Actioned By	Filters the log table so that only logs relating to the specified node that performs an action are displayed.
Filter on Actioned For	Filters the log table so that only logs relating to the specified node that requests an action are displayed.
Filter on Category	Filters the log table so that only logs of the specified Category are displayed.
Filter on Log Level	Filters the log table so that only logs of the specified Level and above are displayed.
Filter on Acknowledgement	Filters the log table so that only acknowledged or unacknowledged logs are displayed.
Filter on Log Type	Filters the log table so that only audit logs are displayed.

Table 31 Log Fields

Field Name	Description
 Select Log Entry	Each log entry can be selected using the selection checkbox. Click one or more log entries to apply an action to that item.

Field Name	Description
# Sequence Number	All logs have a sequence number associated with them representing the order in which they arrived at the master node. Some logs that are used for statistics and reporting do not appear in the log inventory and consequently there will gaps in the displayed sequence numbers.
 Audit	Audit logs are marked with a icon. Click on the Audit icon if present to display the audit information for the corresponding log message. Audit messages cannot be removed from the log database.
 Session	Logs relating to a session are marked with a icon. Click on the Session icon if present to display all log messages that pertain to that same session. The Session Log Details (on page 517) is displayed.
Master Date	The date and time at which the log message was stored within the master node's log database. Each log entry appears in the log table in descending time order, with the latest log displayed first.
Local Date	The date and time at which the message was actually logged by the machine that generated the message. Due to network speed, connectivity issues or different time zones this can differ from the Master Date.
Actioned By	The machine from which the specific action is performed.
Action For	The machine that requested the specific action.
Category	The Protector component that generated the log. A special category named Session is used to indicate the start and end of a sequence of log entries relating to a specific transaction or job. Refer to Session Log Details (on page 517) for details.
Level	How important the log message is. This can be one of four values: ▪ Detail – These log messages can describe intricacies of a process, or non-critical information and do not require any action by the user. ▪ Information – These messages give a higher-level overview of a process, and usually indicate normal operation. These messages do not require any action by the user.

Field Name	Description
	▪ Warning – Messages logged at warning level can indicate that the user should investigate the subject of the message to see if it requires further action. Warning messages are not necessarily critical. ▪ Error – Error messages are only logged when the attention of the user is required. They appear with a red background and cause an error badge to appear on relevant nodes within the Monitor Details (on page 521). The error must be acknowledged to clear this badge (see Acknowledgement Selected Logs button).
Attachment	A icon is displayed if an attachment is included with the log entry. It provides further information that can aid understanding. Click on the Attachment icon if present to display the Log Attachments Dialog (on page 519) which shows additional supporting information that cannot be displayed within the log table.
Log ID	A unique identifier that enables that log message to be looked up in the knowledge base. Not all log entries have a Log ID assigned (i.e. the ID is set to 0). If you require assistance regarding an unassigned ID, please contact support.
Log Message	The actual text content of the log message. When searching this field, words that are entered in the search box must appear in the log message text. Log messages are displayed with alternating background shading to aid readability.
Acknowledgment	A note entered by the user who acknowledged the log message and the date & time of acknowledgement. (See Acknowledgement Selected Logs button).

Purge Logs Dialog

The Purge Logs dialog is displayed when the user attempts to remove logs from the log database.

Caution: Once purged, the log messages cannot be recovered, so be careful not to purge important log messages. Consider exporting logs before purging if in any doubt. Audit logs cannot be purged.

Note: Purging the logs removes the log information from the database but does not reduce its overall size.

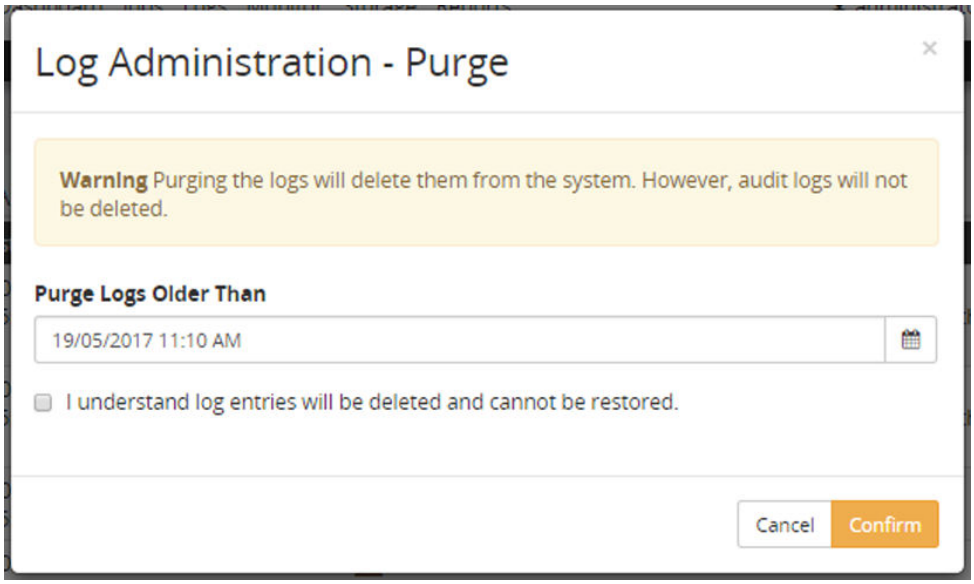


Figure 251 Purge Logs Dialog

Control	Description
Purge Logs Older Than	Specify the date and time, using the Date Time Picker (on page 351), on and before which logs must originate for them to be deleted from the log database. Logs newer than this date and time will not be deleted.
I understand log entries...	As an interlock, this checkbox must be ticked in order to proceed with the purge operation. The Confirm button will not be enabled unless the checkbox is ticked.

Log Purge Schedule

The Log Purge Schedule dialog is displayed when the user clicks on the Log Purge Schedule icon.

The Log Purge Schedule dialogs allows a user to enable or disable the automatic purging of logs. If enabled the user configures the maximum number of weeks logs are kept by for.

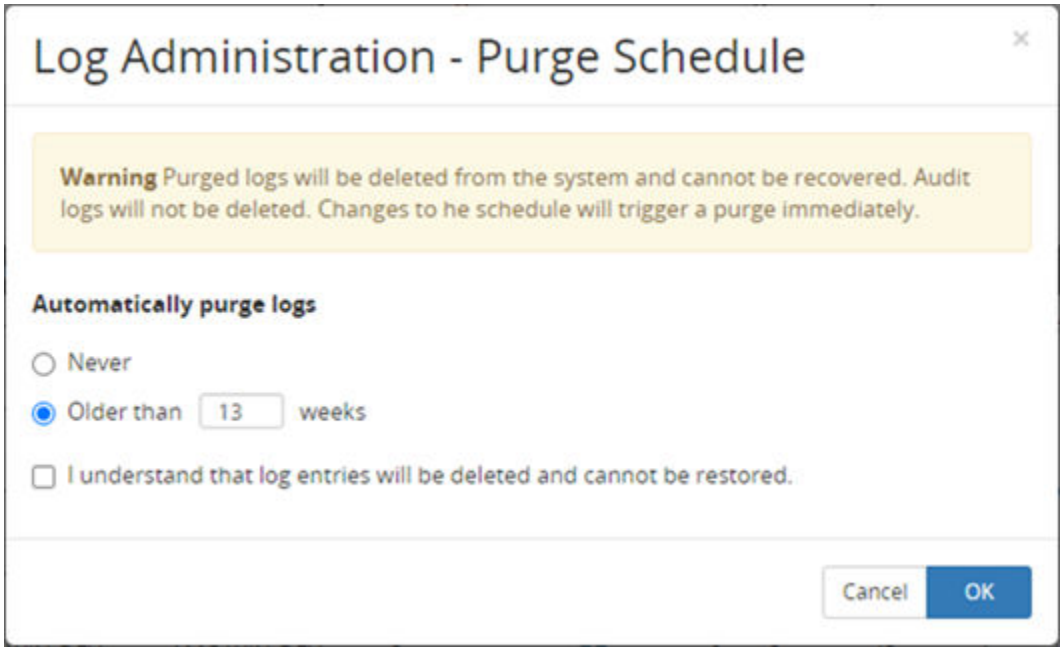


Figure 252 Log Purge Schedule Dialog

Control	Description
Never	Disable automatic log purging.
Older than *13* weeks	Enable automatic log purging causing log messages which are older than, or become older than, the specified number of weeks to be permanently deleted. **Note:** Audit logs are not deleted.
Acknowledgement checkbox	To enable the automatic purging feature, the user must acknowledge that purged logs are not recoverable.

Acknowledge Logs Dialog

The Acknowledge Logs dialog is displayed when the user attempts to acknowledge a log message in the log database.

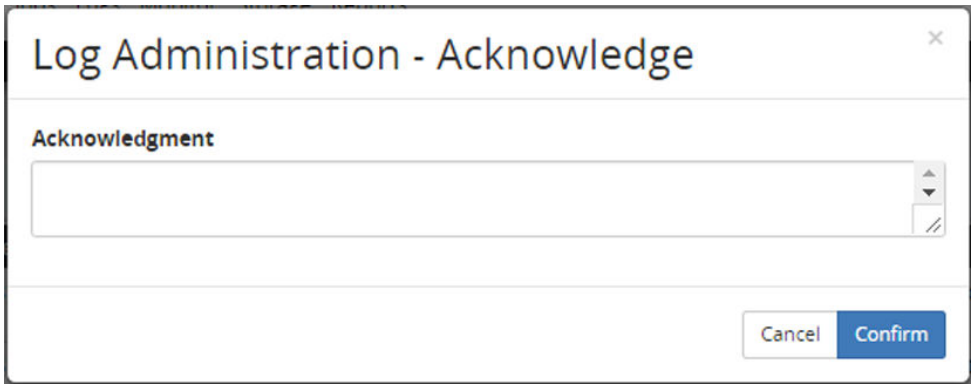


Figure 253 Acknowledge Logs Dialog

Control	Description
Acknowledgment	Optional. Enter a short message that describes the action taken or reason for acknowledging the log. This will be appended with the date and time of acknowledgment upon confirmation. Drag the handle at the bottom right of the edit box to expand it.

Export Logs Dialog

The Log Export dialog is displayed when the user clicks on the export icon on the Logs or Session page.

Note: Only logs currently displayed are exported. Ensure that you have the correct filters applied to the logs before exporting them.

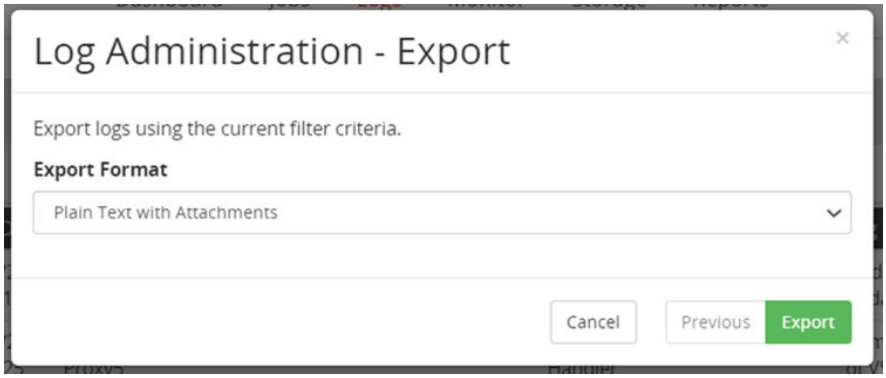


Figure 254 Export Logs Dialog

Control	Description
Export Format	Select one of the following: ▪ HTML ▪ CSV ▪ Plain Text ▪ Plain Text with Attachments
Export	Click to begin the export process. The dialog will display progress of the export process and the export file will automatically be downloaded when ready. The export will be automatically downloaded even if the dialog is closed.



Figure 255 Export Logs Dialog – Export in progress

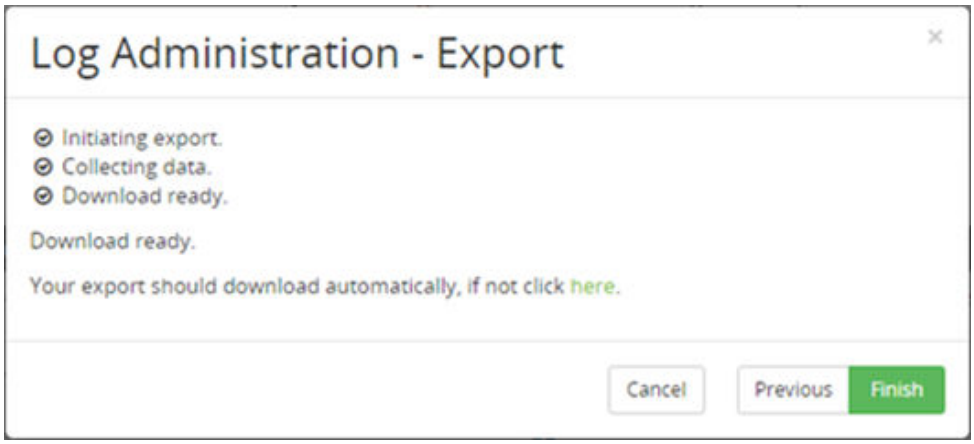


Figure 256 Export Logs Dialog - Export completed

Session Log Details

The Logs Session Details page is displayed when the user clicks on a Session icon in a log message.

A session contains all the logs related to a specific transaction or job. Log entries belonging to sessions are identified in the Logs Inventory (on page 509) by a Session icon. Several sessions can however be interleaved in the main Log, making it difficult to trace through a specific one, however all the logs from one specific session can be viewed on this page.

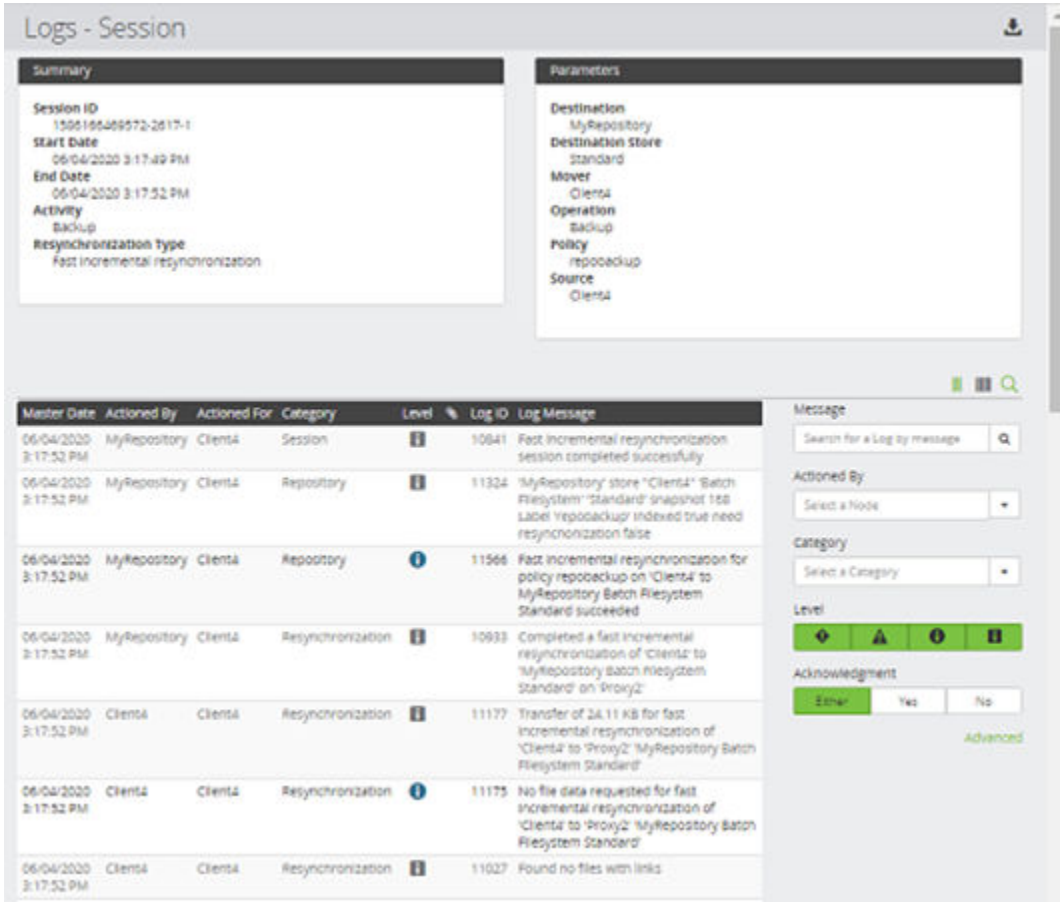


Figure 257 Session Details

Control	Description
Export Session Logs	Exports the log database entries for the displayed session as a file, respecting the currently specified filter terms. The Export Logs Dialog (on page 516) is displayed prior to the logs being exported.
Summary	Provides summary information about the session.
Parameters	Identifies additional details for the session, where possible including, Source, Destination, Policy and Operation.
Condensed View	Display a subset of the log fields in the session table. Green when selected, black when deselected.

Control	Description
Extended View	Display all of the log fields in the session table. Green when selected, black when deselected.
Log Entry	Each log entry for the selected session appears in the session log table in descending time order, with the latest log displayed first. The log table is described in detail in Logs Inventory (on page 509)
Filter on Message	Filters the log table so that only logs containing the specified string in the Log Message field are displayed.
Filter on Actioned By	Filters the log table so that only logs relating to the specified node that performed the action are displayed.
Filter on Category	Filters the log table so that only logs of the specified Category are displayed.
Filter on Log Level	Filters the log table so that only logs of the specified Level and above are displayed.
Filter on Acknowledgement	Filters the log table so that only acknowledged or unacknowledged logs are displayed.

Log Attachments Dialog

The Log Attachments dialog is displayed when the user clicks on the attachment icon for a log message.

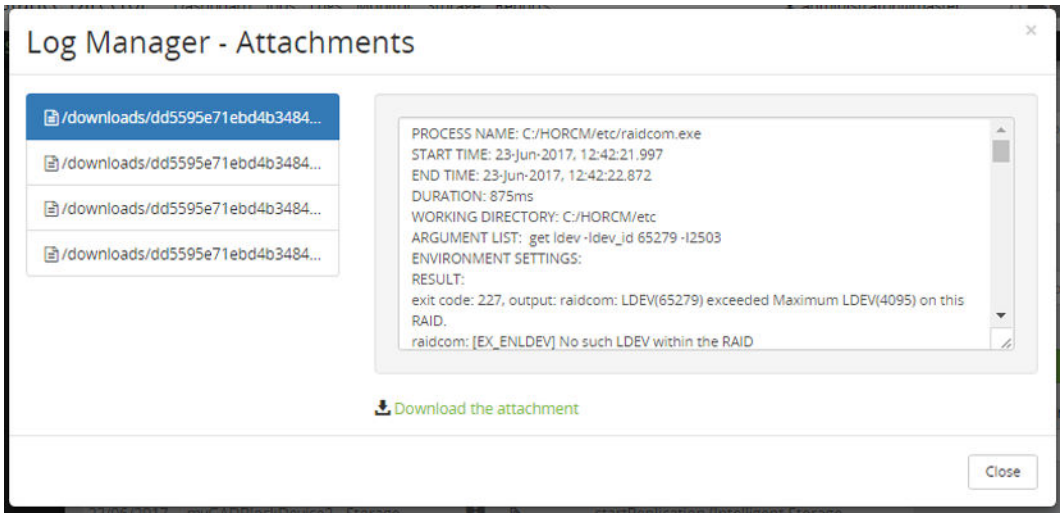


Figure 258 Log Manager - Attachments Dialog

Control	Description
Attachment File Name	Each attachment is selected by clicking the buttons down the left side of the dialog.
Attachment text	The attachment text for the selected attachment is displayed in a text box to the right.
Download the attachment	Click here to download the selected attachment as a `.txt` file. A tool tip is displayed indicating where the file will be downloaded. The download location may be configured via your browser. Your browser may prompt you when the file has downloaded and enable you to open it in a text editor.

Monitor UI Reference

This section describes the Monitor UI, accessed via the Main Banner (on page 286).

For further information, refer to:

- Monitor Concepts (on page 106)
- Monitor Tasks (on page 260)

Monitor Inventory

The Monitor Inventory lists all data flows which have been activated.

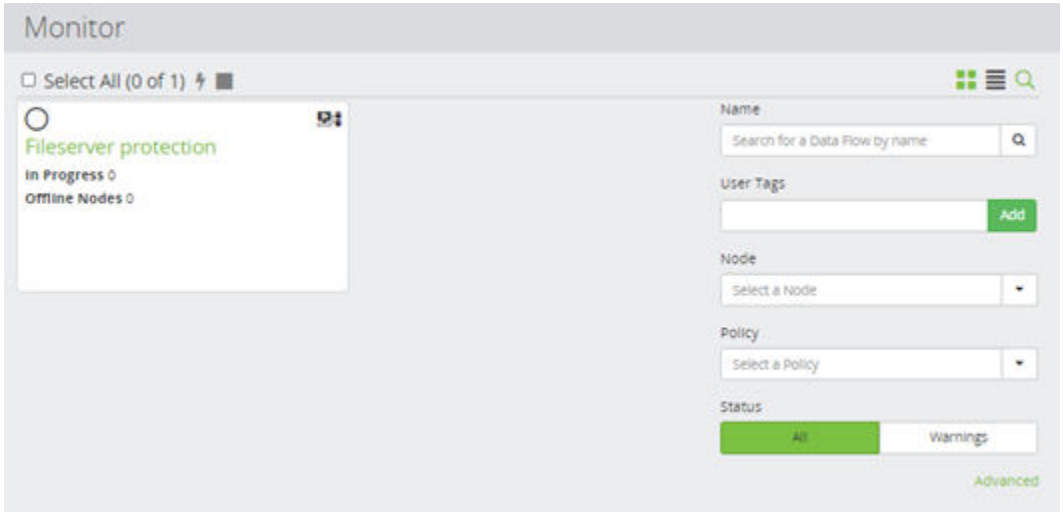


Figure 259 Monitor Inventory

Control	Description
Deactivate	Enabled only when one or more data flows are selected. Deactivates the selected data flows. **Caution:** If the data flow contains hardware storage based operations, this will remove the pairing relationships.
Data Flow	Click on the data flow name to open the Monitor Details (on page 521) for the selected data flow. If the data flow is in a warning state (e.g. a replication is paused or swapped, or the storage array is in an unexpected state), then a warning icon ( ) will be displayed here.
Filter on Data Flow Name	Filters the displayed data flows based on the data flow name.
Filter on User Tags	Filters the displayed results based on Tags contained in the data flow.
Filter on Node	Filters the displayed data flows based on the node name.
Filter on Policy	Filters the displayed data flows based on the policy name.
Status	Filter the displayed data flows based on the Status. It is possible to display All data flows or just the data flows with Warnings.

Monitor Details

This page displays the details of an active Data Flow and enables you to monitor activities.

The Monitor Details page is split into several areas:

- Active Data Flow and Status Badges
- Context sensitive details dependent on what is selected on the data flow including:
  - Applied Policies
  - Summary information
  - Node details
  - Network/Cache Utilization
  - Mover status
  - Linked Records for hardware operations
  - Jobs relating to the selected node(s)
  - Job Details
  - Logs relating to the selected node(s)

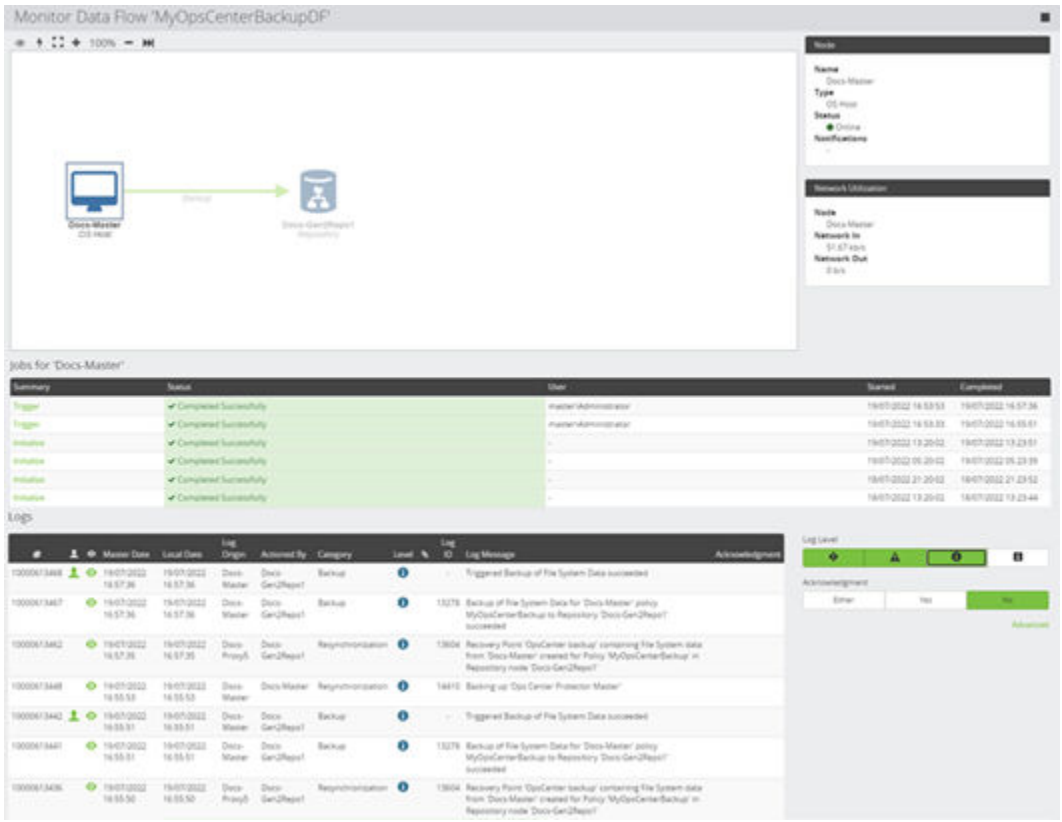


Figure 260 Monitor Data Flow Details Page (shown here with an OS Host node selected)

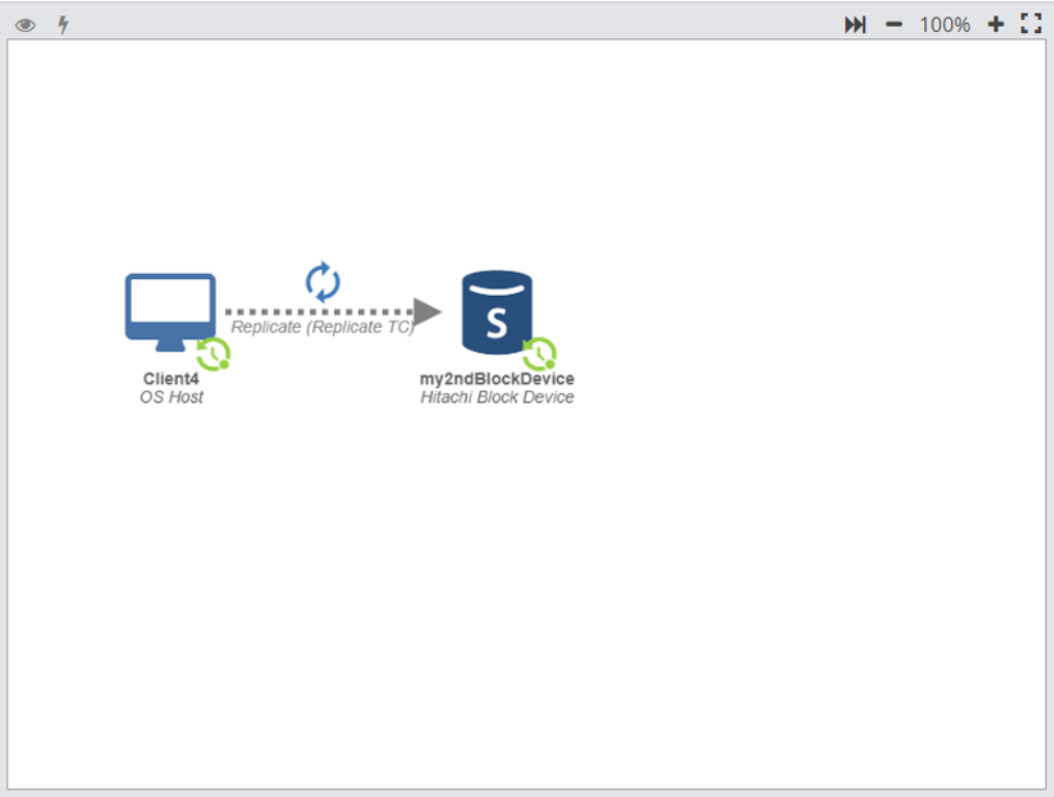
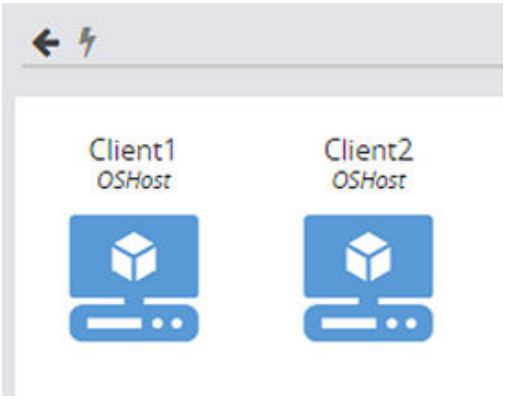


Figure 261 Active Data Flow Panel

Control	Description
Deactivate	Deactivates the data flow. The Monitor page is closed and the Monitor Inventory (on page 520) is displayed with the deactivated data flow removed from the inventory. To re-activate the data flow, locate it in the Data Flows Inventory (on page 355) and click Activate.
View Group Nodes	Displays the nodes within the selected node group and enables each node to be examined:  Click the ← Back button to return to the data flow view.

Control	Description
 Trigger Operation	Opens the Trigger Operation Dialog (on page 529) listing the available policy operations on the selected node. Selecting the required policy operation and clicking OK will instantly start the selected operation.
 Next Node	Click to move the focus to the next node on the data flow.
 Zoom In/Out	Click the buttons next to the workspace, press +/- on the keypad or hold down the **CTRL** key whilst using the mouse wheel to zoom in and out. The current zoom level is displayed between the zoom buttons. Click the zoom level to reset to 100% zoom.
 Fit to Screen	Click the button next to the workspace or press the **HOME** key to select a zoom level that allows the entire data flow to fit within the bounds of the workspace.
 Node	Select a node to view its details and current activity in the areas below and to the right. Source nodes, movers and destination nodes are represented by icons representing their type (see Node Type Icons (on page 657)).
Status Badge (E.g. )	A status badge may be superimposed on a node, or above a mover, to indicate that the node is in a particular state. Refer to Monitor Status Badges (on page 530) for a description of each badge. Nodes that cannot be contacted by the master are displayed in red (either because no network connection can be made or because the Protector services are not running on the node).
 Snapshot Badge	A snapshot badge may be superimposed on a node (in the bottom right corner) to indicate that the node has a policy containing a snapshot operation applied to it.
 Connector	Connectors between node indicate the direction and type of mover: ▪ Batch mover ▪ Continuous mover Select a connector to view its details in the area to the right. A badge may be superimposed above a connector to indicate a particular state. Refer to Monitor Status Badges (on page 530) for a description of each badge.

Control	Description
	Note: For some replication data flows it is possible to swap the direction. The arrowhead on the monitor data flow connector does not change direction when the replication is swapped.



Figure 262 Applied Policies Panel (displayed when nothing is selected on the data flow)

Control	Description
Policy Name	Lists the policies applied on the data flow. Click the button at the end of the policy name to open the Policy Details (on page 737) in a separated browser tab.

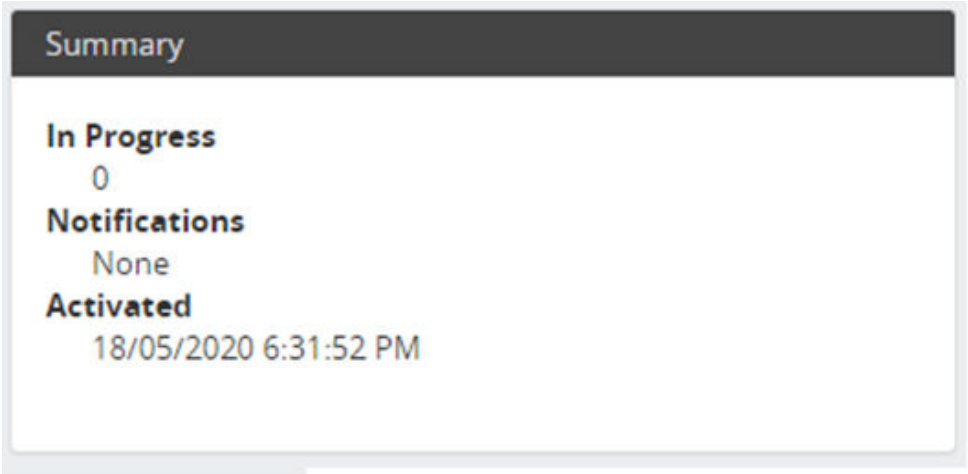


Figure 263 Summary Panel (displayed when nothing is selected on the data flow)

Control	Description
In Progress	Indicates the number of jobs in progress.
Notifications	A list of notifications relating to the data flow. If the data flow is in a warning state (e.g. a replication is paused or swapped, or the storage

Control	Description
	array is in an unexpected state), then a warning icon ( ) will be displayed alongside the notification text.
Activated	Shows the date and time the data flow was activated.



Figure 264 Node Panel (displayed when an OS Host or Repository node is selected)

Control	Description
Name	The node name.
Type	The node type.
Status	Whether the node is online or offline.
Notifications	A list of notifications relating to the node. If the node on the data flow displays a status badge it will only show the highest priority notification, whereas this field will indicate all status notification (see Monitor Status Badges (on page 530)).

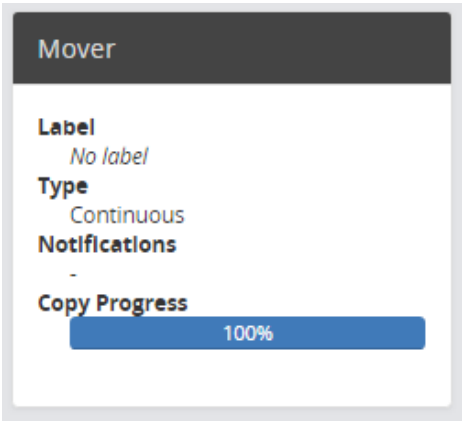


Figure 265 Mover (displayed when a mover is selected)

Control	Description
Label	The user defined mover label text.
Type	The mover type.
Notifications	A list of notifications relating to the mover. If the mover on the data flow displays a status badge it will only show the highest priority notification, whereas this field will indicate all status notification (see Monitor Status Badges (on page 530)).
Copy Progress	Displayed only for block hardware replications. Indicates the concordance between the source and destination.



Figure 266 Network/Cache Utilization Panel (displayed when an OS Host node is selected)

Control	Description
Node	Name of the node currently selected in the active data flow panel.
Network In	Shows the incoming network traffic for the node currently selected in the data flow.
Network Out	Shows the incoming and outgoing network traffic for the node currently selected in the data flow.

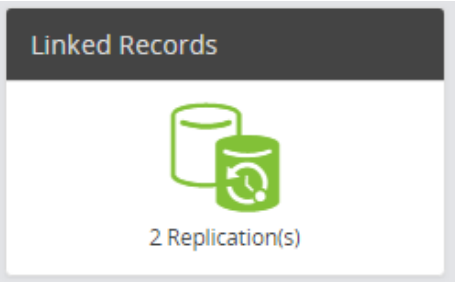


Figure 267 Linked Records Panel (displayed when an Hitachi Block Device destination node is selected)

Control	Description
 Replication(s)	Click this button to open the Hitachi Block Replications Inventory (on page 878) showing replications relating to the selected Hitachi Block Device node.

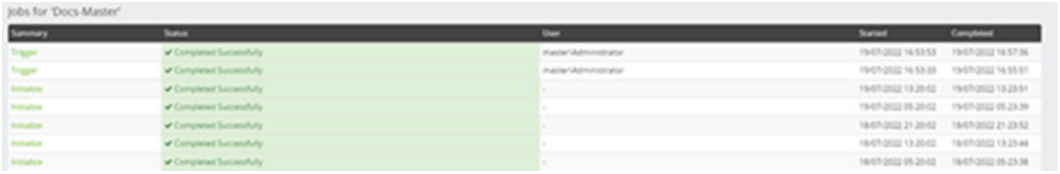


Figure 268 Jobs Panel (shows all jobs relating to the data flow or jobs relating to the selected node)

Control	Description
Jobs	Lists the jobs relating to the node currently selected in the data flow. Refer to Jobs Inventory (on page 493) for a description of the controls displayed in this area. Clicking on the jobs under Summary will take you to the tasks associated with the job.

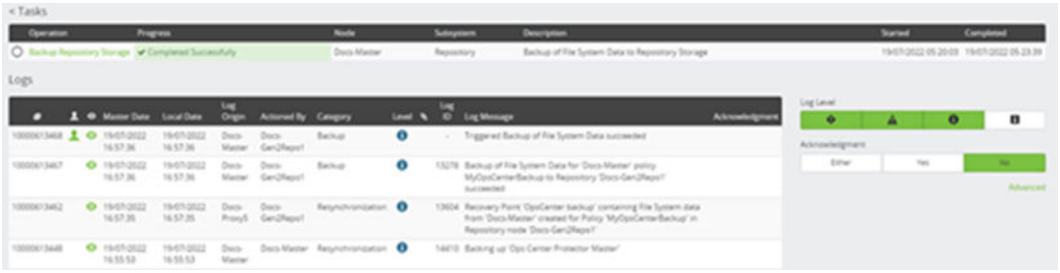


Figure 269 Tasks Details (shows all tasks relating to the jobs relating to the selected node)

Control	Description
Tasks	Lists the tasks relating to the node currently selected in the data flow. Refer to Tasks Inventory (on page 501) for description of the controls displayed in this area. Clicking on the tasks under Operation will display the information related to that task.



Figure 270 Logs Panel (shows all logs relating to the data flow or logs relating to the selected node)

Control	Description
Logs	Lists the logs relating to the node currently selected in the data flow. Refer to Logs Inventory (on page 509) for a description of the controls displayed in this area.
Filter on Log Level	Displays logs based on the selected Log Level. By default only Error level logs are displayed for more details select a lower log level option.
Filter on Acknowledgment	Displays logs based on the acknowledgment state.

Trigger Operation Dialog

This dialog is displayed when a policy operation is manually triggered.

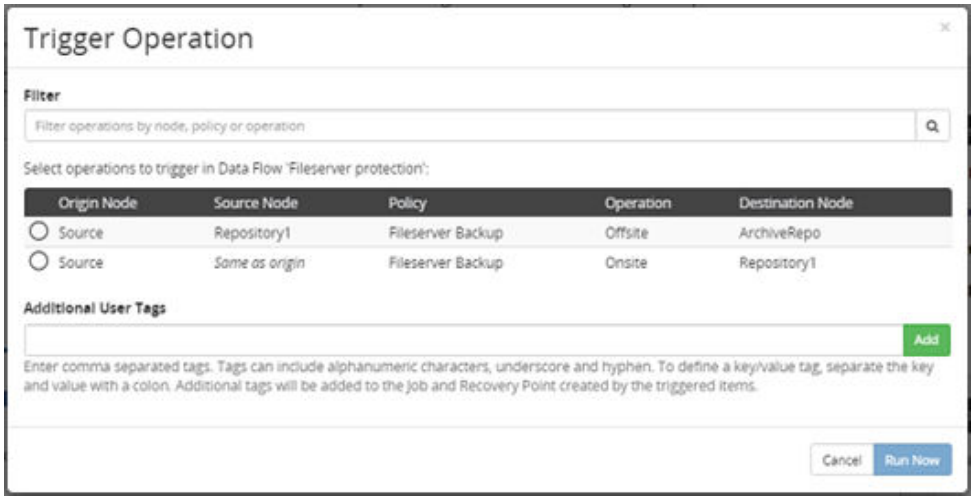


Figure 271 Trigger Operation Dialog

Control	Description
Policy Operation List	Select the name of the policy and operation that you want to trigger.
Filter	Allows the displayed triggers to be filtered by node, policy or operation.
Additional User Tags	Add the tags to be associated with the object being created.

Note: On triggering the policy operation, the task progress can be seen on the screen.

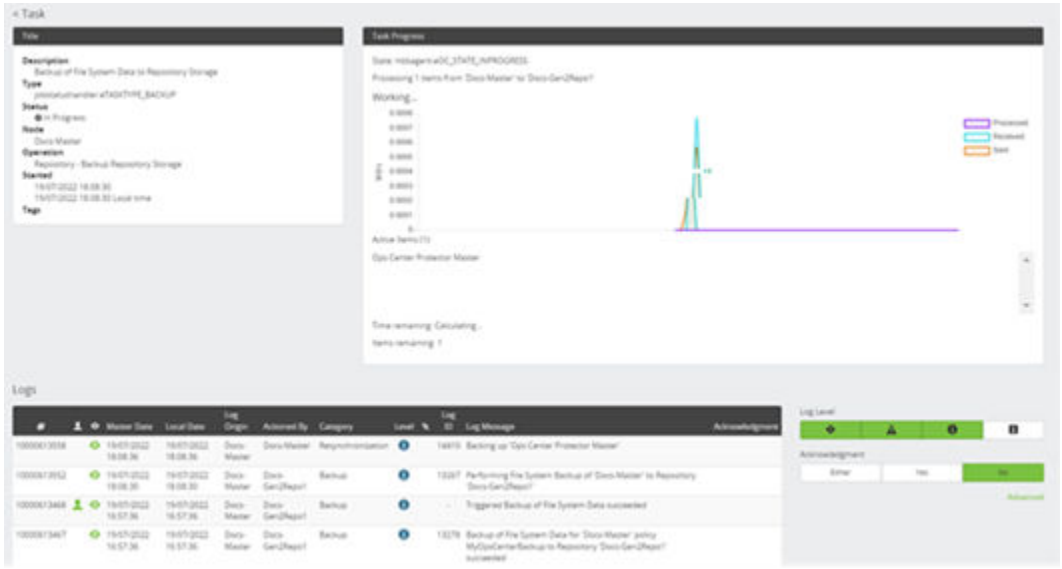


Figure 272 Task Trigger Operation Progress

Monitor Status Badges

Status badges are displayed on a node or mover to indicate an operating state, warning or error condition. Select the node or mover to view its status.

Note: Status badges reflect Protector's view of system state and may not concur with the underlying hardware state. For example, a batch replication will be paused by the storage device whenever it is within RPO and periodically unpaused while being re-evaluated due to the RPO being reached. These external pause/unpause states will not be indicated on the Monitor Details (on page 521) data flow. However, if the same replication is paused by the user from within Protector, then the paused state will be indicated.

Status badges are displayed in priority order on the Monitor Details (on page 521) data flow. When multiple status notifications exist simultaneously only the highest priority is shown. To see all notifications, select the node or mover in question and review the notifications in the associate node or mover panel adjacent to the data flow panel.

Table 32 Monitor Status Badges

Priority	Badge	State
N/A		Disconnected Node - Displayed when a node has gone offline.
1 (Highest)		Error - Displayed on a node or mover if it is in an error state.
2		Warning - Displayed on a node or mover if it is not in the expected state.
3		Paused - Displayed on a mover when a replication has been paused within Protector.
4		Swapped - Displayed on a mover when a replication has been swapped within Protector.
5		Resynchronizing - Displayed on a mover when (re)synchronization or replication re-evaluation is in progress.
=5		Restoring - Displayed on a node that is restoring data.

Node Groups UI Reference

This section describes the Node Groups UI, accessed via the Navigation Sidebar (on page 291).

For further information, refer to:

- Node Concepts (on page 47)
- Node Tasks (on page 261)
- Nodes UI Reference (on page 536)

Node Groups Inventory

This inventory lists all defined Node Groups.

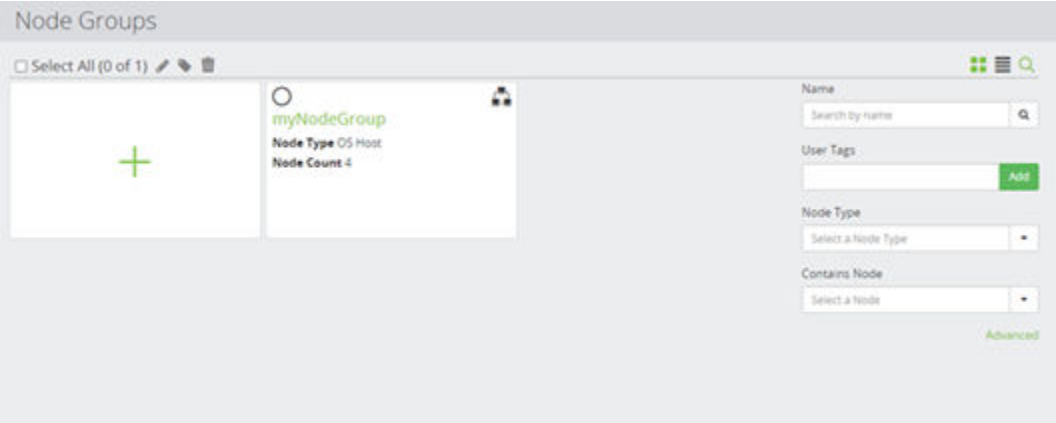


Figure 273 Node Groups Inventory

Control	Description
Edit Tags	Edits an existing node group in the inventory. The Node Group Wizard (on page 532) is launched to enable the group's attributes to be changed.
Delete Node	Enabled only when one or more Node Group is selected. Deletes the selected item from the inventory. The associated Nodes are not deleted.
Create a new item	Creates a new Node Group. The Node Group Wizard (on page 532) is launched to guide you through the process.
User defined Node Group(s)	Any number of user defined Node Groups can be created. These are displayed in the inventory. Click on the node group name to open the Node Group Details (on page 535) which enables the group to be viewed and edited.
Filter on Name	Filters the displayed results based on Node Name.
Filter on User Tags	Filters the displayed results based on Tags contained in the data flow.
Filter on Node Type	Filters the displayed results based on Node Type.
Filter on Contains Node	Filters the displayed results to groups that contain the node entered.

Node Group Wizard

This wizard is launched when a new Node Group is added to the Node Groups Inventory.

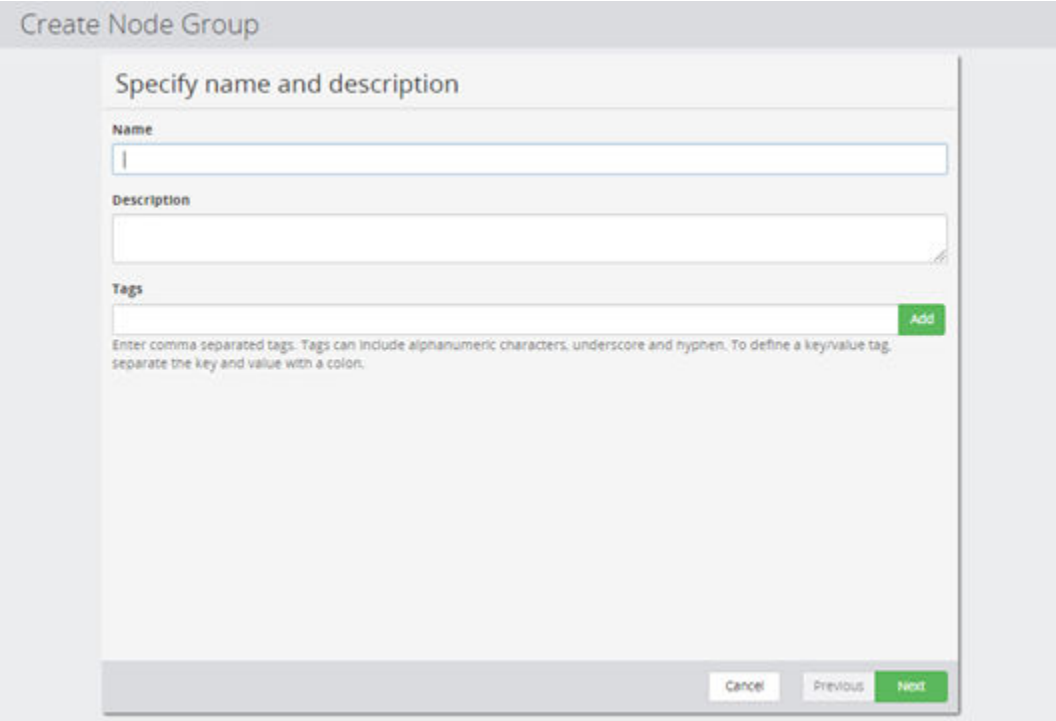


Figure 274 Node Group Wizard - Specify name and description

Control	Description
Name	Enter the name of the node group.
Description	Optional. Enter a description of the node group.
Tags	Add the tags to be associated with the object being created.

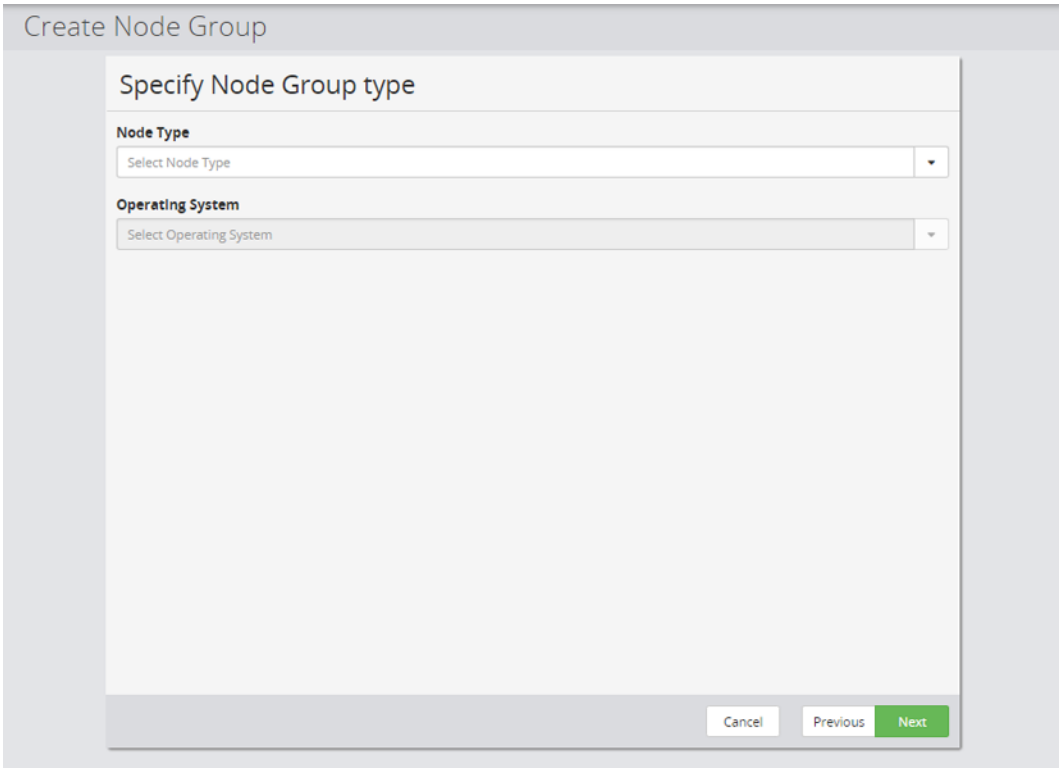


Figure 275 Node Group Wizard - Specify Node Group Type

Control	Description
Node Type	Specifies the node type that this group will contain. Only node types that can have a common policy assigned are available.
Operating System	If the *OS Host* node type is selected then this control is enabled so that a specific operating system can be selected.

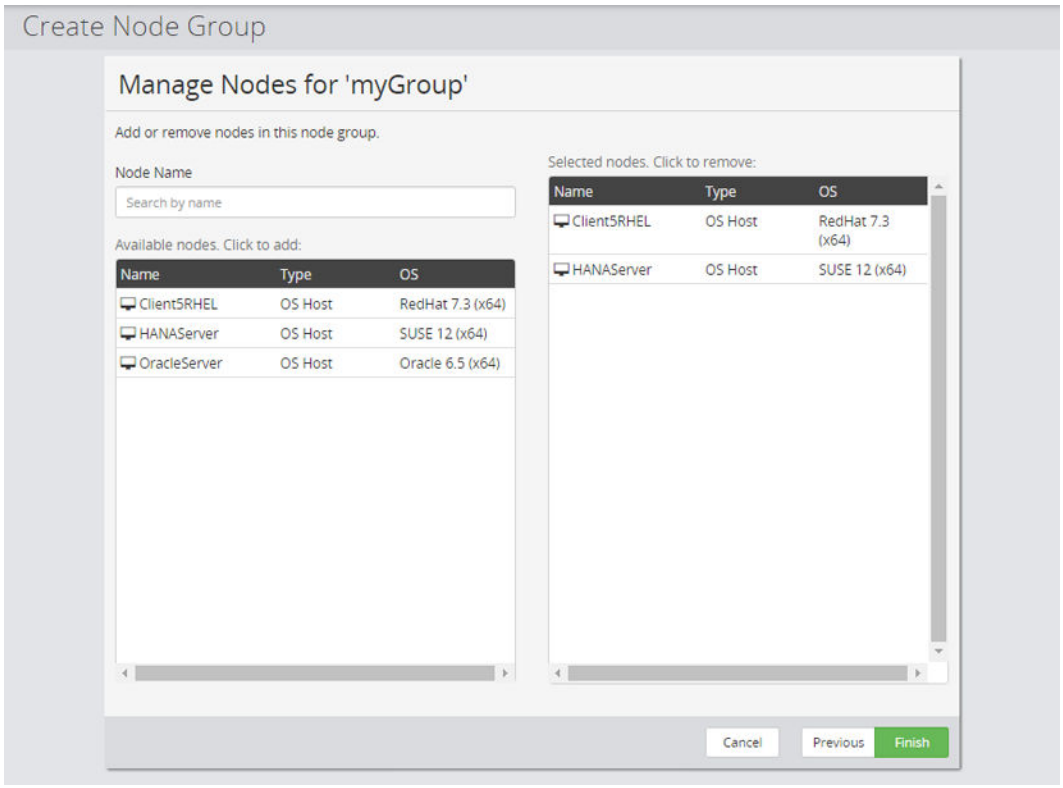


Figure 276 Node Group Wizard - Manage Nodes

Control	Description
Node Name	Filters the list of nodes displayed below by name.
Click to add	Lists all available nodes of the specified type. Click on the name of the node you want to add to this node group.
Click to remove	Lists all the nodes currently in this group. Click on the name of the node you want to remove from this node group.

Node Group Details

This page displays the details of a Node Group and enables you launch the wizard to edit them.

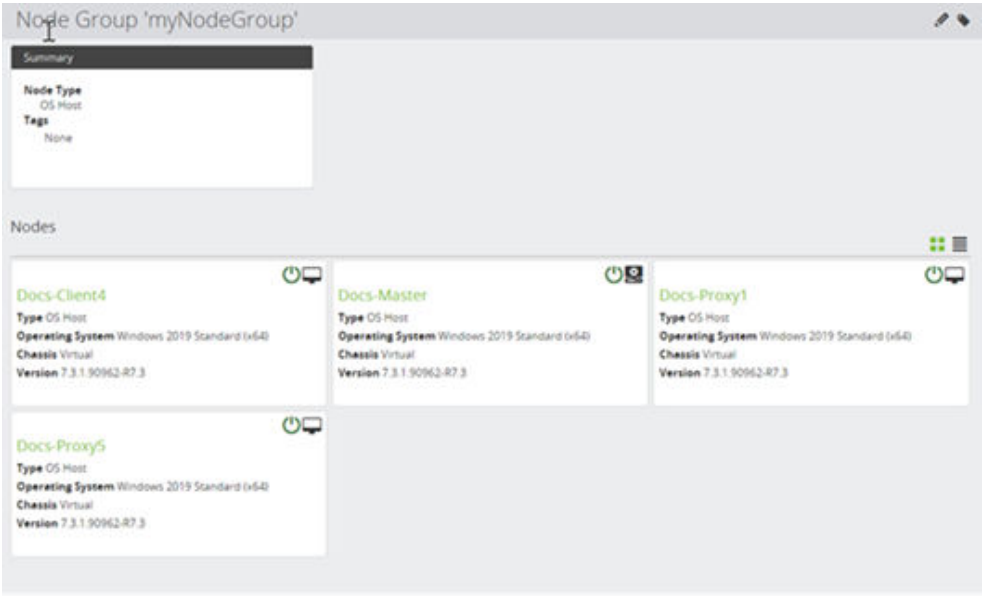


Figure 277 Node Group Details

Control	Description
Edit	Launches the Node Group Wizard (on page 532) to enable you to view and edit the node group.
Tag	Modifies the tags of an existing object from the either the inventory screen or the details screen of the object.
Nodes in group	Click the node name to open the Node Details (on page 650) to enable you to view the node's status and details.

Nodes UI Reference

This section describes the Nodes UI, accessed via the Navigation Sidebar (on page 291).

Nodes Inventory

This inventory lists all the nodes in a Protector network, be they online, offline, authorized or deauthorized.

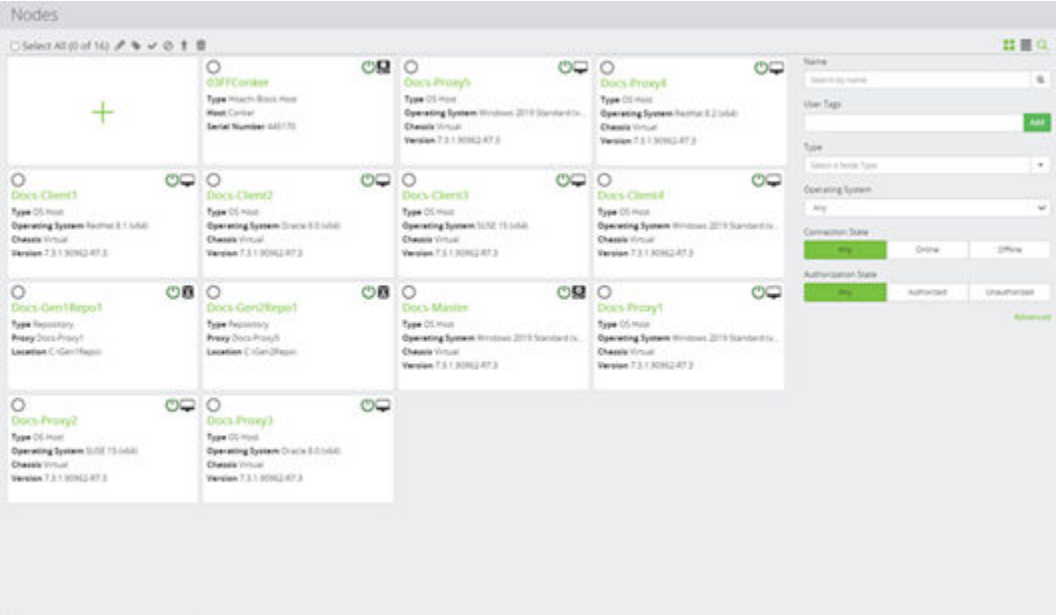


Figure 278 Nodes Inventory

Control	Description
Edit	Edits an existing node in the inventory. The Node Type Wizard (on page 539) is launched to enable the node's attributes to be changed.
Edit Tags	Modifies the tags of an existing object from the either the inventory screen or the details screen of the object.
Authorize	Enabled only if one or more unauthorized nodes is selected in the inventory. Attempts to authorize the selected nodes with the Master node. Only nodes that have been authorized by the Master node may perform Protector functions. **Note:** If an attempt is made to authorize an inactive or unknown node, or if the master node fails to communicate with the node, an error log is generated and the node remains unauthorized.
Deauthorize	Enabled only if one or more authorized nodes is selected in the inventory. Attempts to deauthorize the selected nodes. Nodes that have been deauthorized cannot perform Protector functions.

Control	Description
	Note: ▪ Before deauthorizing nodes that are no longer required, they must first be deleted from the current data flow definitions, then any modified active data flows must be reactivated. If any attempt is made to deauthorize a node that is implementing rules in the currently active data flows, an error message is generated for each such node and the node remains authorized. ▪ It is possible to de-authorize a node with active mounts. This should be avoided because once de-authorized it is not possible to unmount.
 Upgrade Clients	Enabled only if one or more authorized nodes is selected in the inventory. Attempts to remotely upgrade the Protector software installed on the selected nodes. The required upgrade installer and configuration files must be present in the `C:\Programs Files \Hitachi \Protector\runtime\updater` folder on the Master node. **Note:** ▪ Only *OS Host* nodes can be upgraded. It is recommended to upgrade nodes in batches of 20. ▪ It is recommended to manually upgrade the master node.
 Delete	Enabled only if one or more nodes is selected in the inventory. The node is deleted from the inventory. **Note:** If an *OS Host* node is still running the Protector hub process and is configured to use the current *Master* node, then the node will re-appear as an unauthorized node as it periodically reconnects to the *master* node. Protector should be uninstalled from the node to stop this periodic reconnection.
 Create a new item	Adds a new node to the inventory. The Node Type Wizard (on page 539) is launched to guide you through the process.
Existing Node(s)	Nodes on which Protector has been installed are automatically detected and listed here along side those that have been added by the user. The Node Details (on page 650) is displayed to enable the node's details to be viewed and edited.

Control	Description
	Note: ▪ DHCP renewal can cause temporary disconnection of a node. ▪ If the host of a virtual machine(s) creates a Windows restore point, then the virtual machine(s) can temporarily disconnect. ▪ If a node has been assigned to another master, it will not appear in the inventory.
Filter on Node Name	Filters the displayed results based on Node Name.
Filter on User Tags	Filters the displayed results based on Tags contained in the data flow.
Filter on Node Type	Filters the displayed results based on Node Type.
Filter on Operating System	Filters the displayed results based on Operating System.
Filter on Connection State	Filters the displayed results based on Connection State.
Filter on Authorization State	Filters the displayed results based on Authorization State.

Node Type Wizard

This wizard is launched when a new Node is added to or edited in the Nodes Inventory.

Nodes can be created to represent subsystems that interact with Protector. Data from these systems can then be integrated into the usual Protector workflow.

Authorization will fail if the supplied configuration is invalid. To alter the configuration, edit the node's settings. Any changes will come into effect the next time data flows using the node are reactivated.

Note: Node names are limited to a maximum of 64 characters.

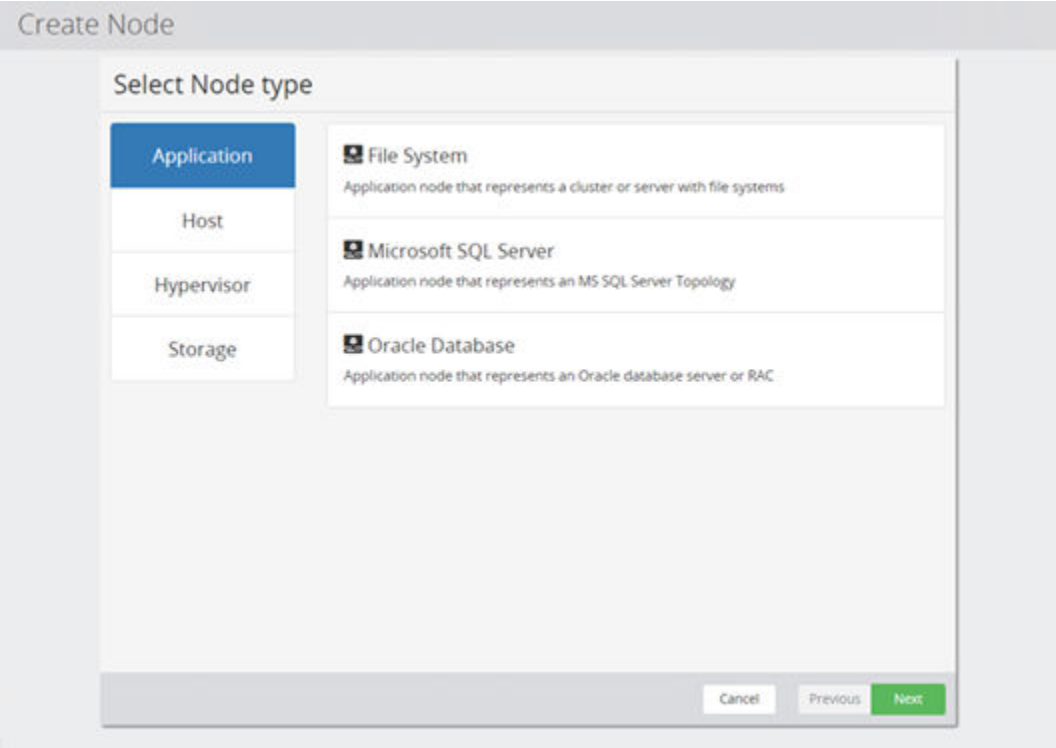


Figure 279 Node Type Wizard - Application

Control	Description
File System	Creates an application node representing a server or Windows Failover Cluster with file systems. The File System Node Wizard (on page 544) is launched to guide you through the process.
Microsoft SQL Server	Creates an application node representing Microsoft SQL Server . The Microsoft SQL Server Node Wizard (on page 547) is launched to guide you through the process.
Oracle Database	Creates an application node representing an Oracle database server or RAC. The Oracle Application Node Wizard (on page 551) is launched to guide you through the process.

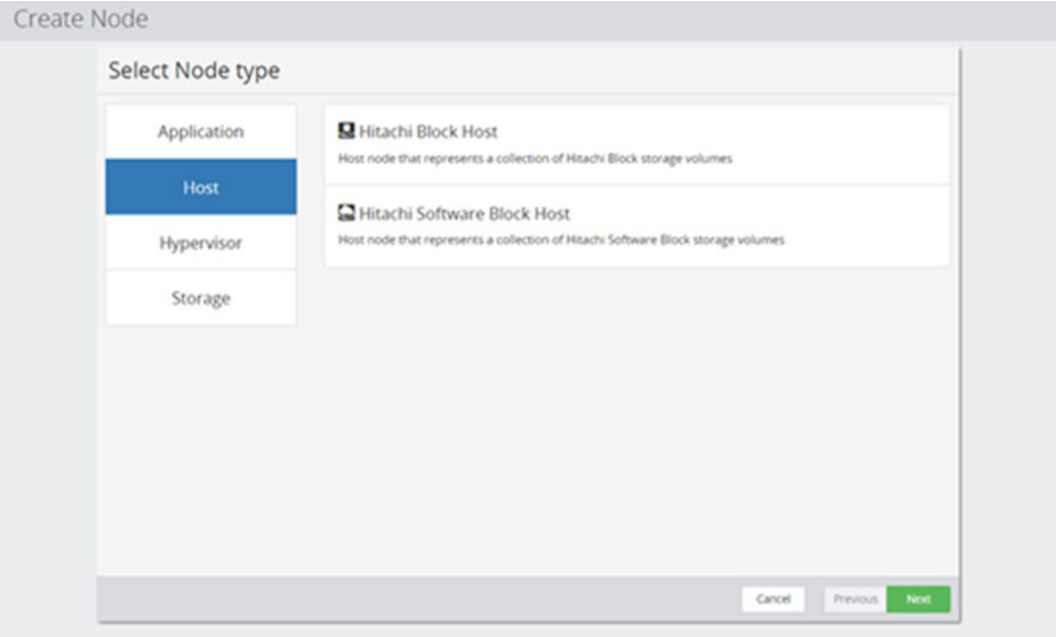


Figure 280 Select Node Type Wizard - Host

Control	Description
Hitachi Block Host	Creates a host node that represents a number of Hitachi Block Storage volumes. The Hitachi Block Host Node Wizard (on page 559) is launched to guide you through the process.
Hitachi Software Block Host	Creates a software host node that represents a number of Hitachi Software Block Storage volumes. The Hitachi Software Block Host Node Wizard (on page 564) is launched to guide you through the process.

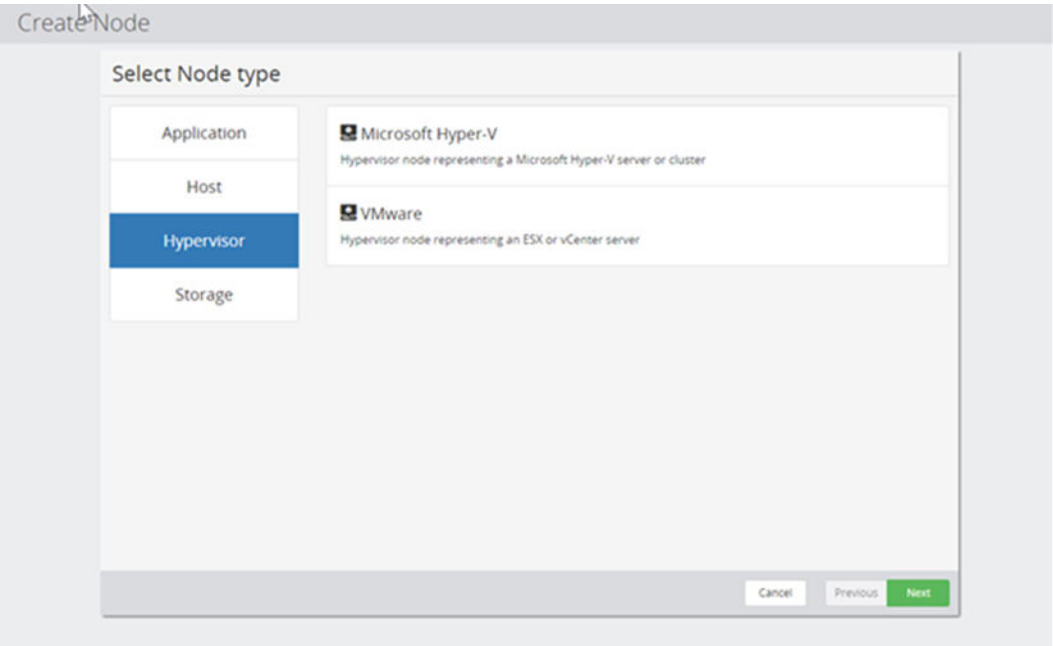


Figure 281 Select Node Type Wizard - Hypervisor

Control	Description
Microsoft Hyper-V	Creates a hypervisor node Microsoft Hyper-V. The Hyper-V Node Wizard (on page 568) is launched to guide you through the process.
VMware	Creates a hypervisor node representing VCenter. The VMware Node Wizard (on page 574) is launched to guide you through the process.

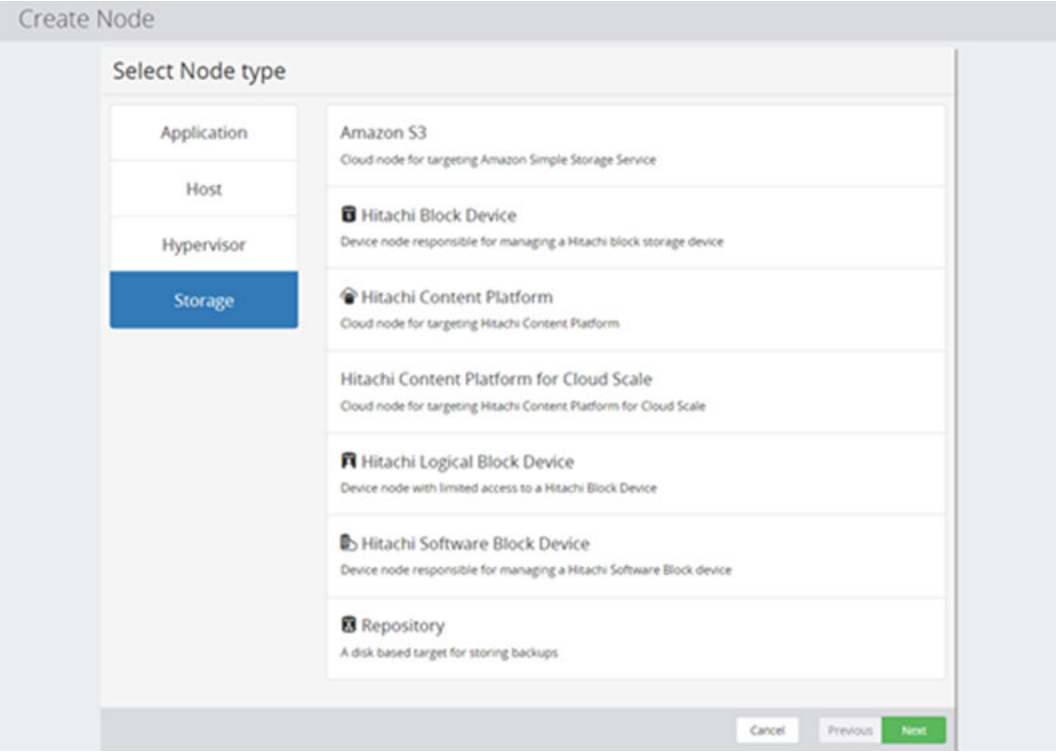


Figure 282 Select Node Type Wizard - Storage

Control	Description
Amazon S3	Creates a storage node responsible for managing an Amazon SimpleStorage Service backup. The Amazon S3 Storage Node Wizard (on page 641) is launched to guide you through the process.
Hitachi Block Device	Creates a device node responsible for managing a Hitachi Block device. The Hitachi Block Device Node Wizard (on page 583) is launched to guide you through the process.
Hitachi Content Platform	Creates a storage node for targeting Hitachi Content Platform. The Hitachi Content Platform Storage Node Wizard (on page 600) is launched to guide you through the process.
Hitachi Content Platform for cloud scale	Creates a storage node for targeting Hitachi Content Platform for cloud scale. The Amazon S3 Storage Node Wizard (on page 641) is launched to guide you through the process.

Control	Description
Hitachi Logical Block Device	Creates a device node with limited access to a Hitachi Block device. The Hitachi Logical Block Device Node Wizard (on page 609) is launched to guide you through the process.
Hitachi Software Block Device	Createsa device node responsible for managing a Hitachi Software Block device. The Hitachi Software Block Device Node Wizard (on page 626) is launched to guide you through the process.
Repository	Creates a disk based target for storing backups. The Repository Storage Node Wizard (on page 632) is launched to guide you through the process.

File System Node Wizard

Protector will launch this wizard when a new File System node is added to the Nodes Inventory.

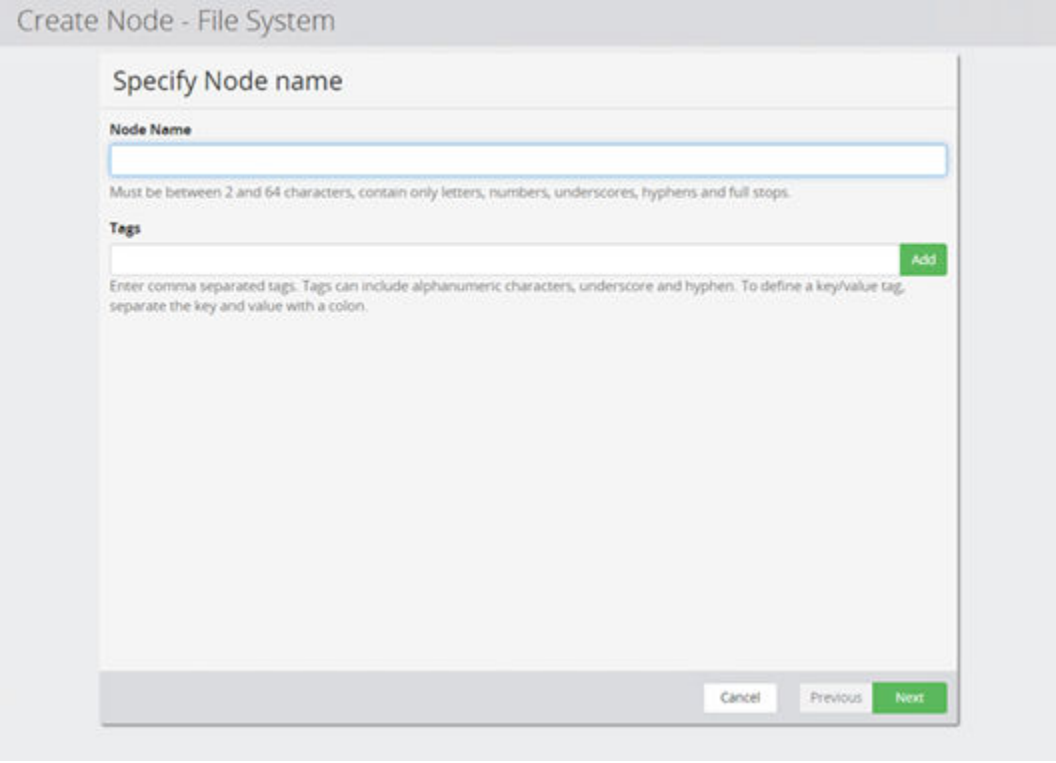


Figure 283 File System - Specify Node Name

Control	Description
Node Name	Enter a name for the File System node.
Tags	Add the tags to be associated with the object being created.

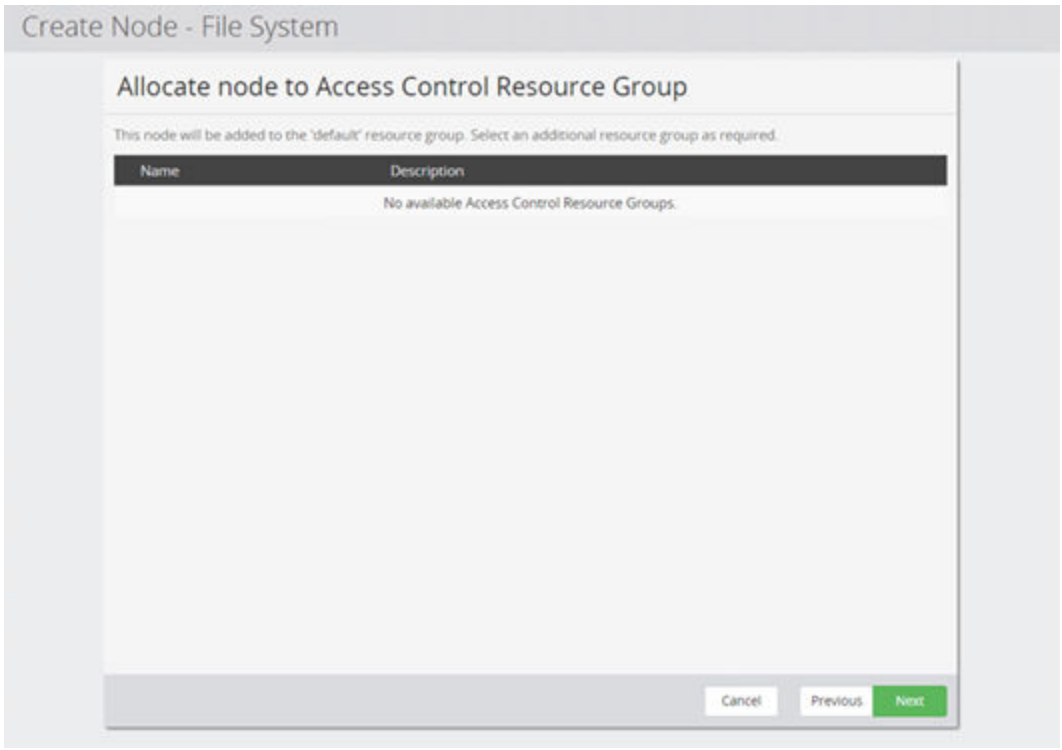


Figure 284 File System Node Wizard - Allocate node to Access Control Resource Group

Control	Description
Resource Groups	Select the resource group(s) to which this node will be allocated for the purposes of RBAC. All nodes are automatically allocated to the 'default' resource group.

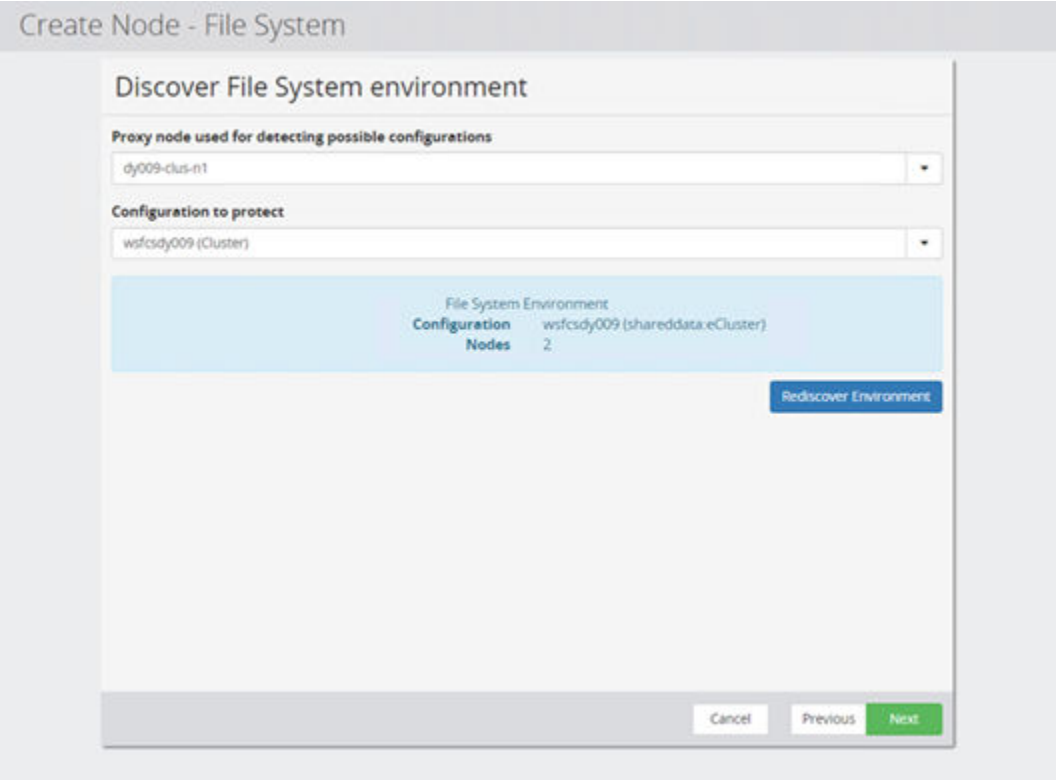


Figure 285 File System Node Wizard - Discover File System environment

Control	Description
Proxy Node used for detecting possible configurations	Select an OS Host node which represents the standalone server or which is part of the cluster.
Configuration to protect	Select the File System environment you want this node to represent. Choose the 'Standalone' option if you only have a standalone server. If you have a Windows Server Failover Cluster, you can select either 'Standalone' to represent only the OS Node or the cluster name to represent the entire cluster.
Rediscover Environment	Click in case you want to refresh the list of available File System configurations.

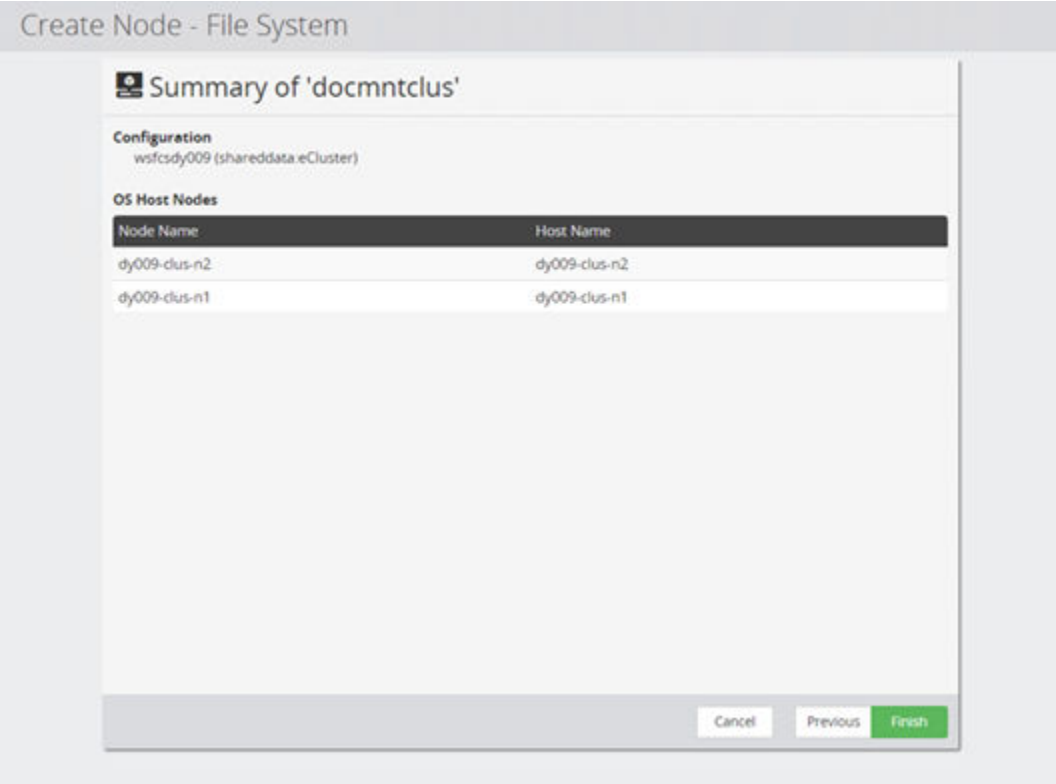


Figure 286 File System Node Wizard - Summary

Control	Description
Summary	Summary of the settings and nodes selected in the node creation wizard.

Microsoft SQL Server Node Wizard

Protector will launch this wizard when a new Microsoft SQL Server node is added to the Nodes Inventory.

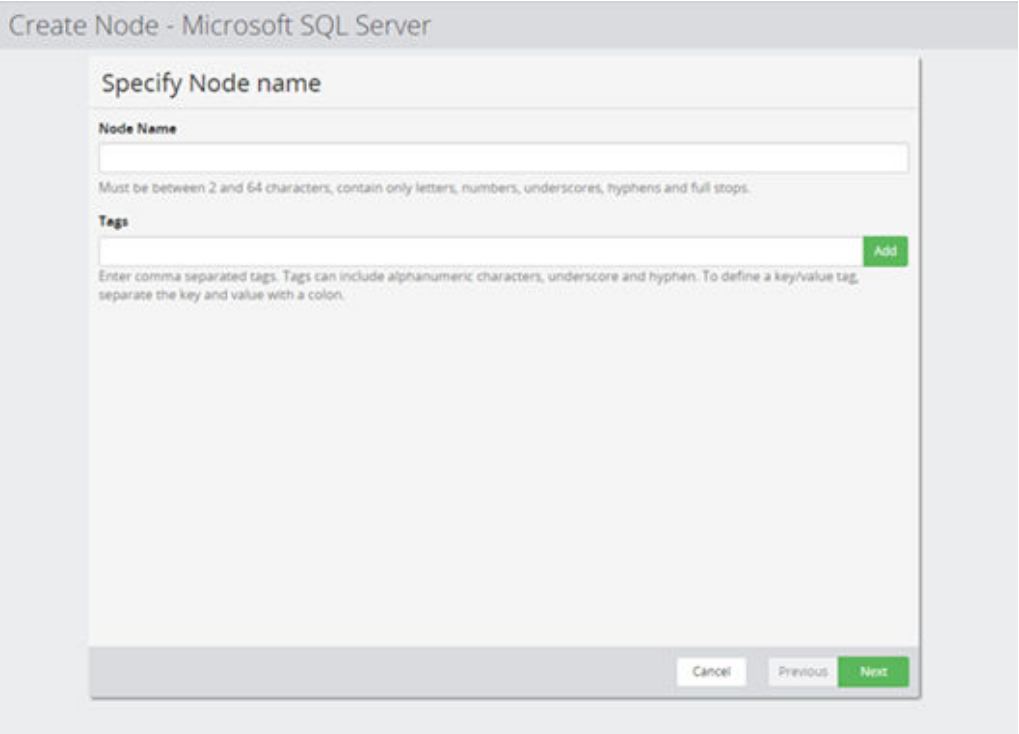


Figure 287 Microsoft SQL Node Wizard - Specify Node Name

Control	Description
Node Name	Enter a name for the Microsoft SQL Server node.
Tags	Add the tags to be associated with the object being created.

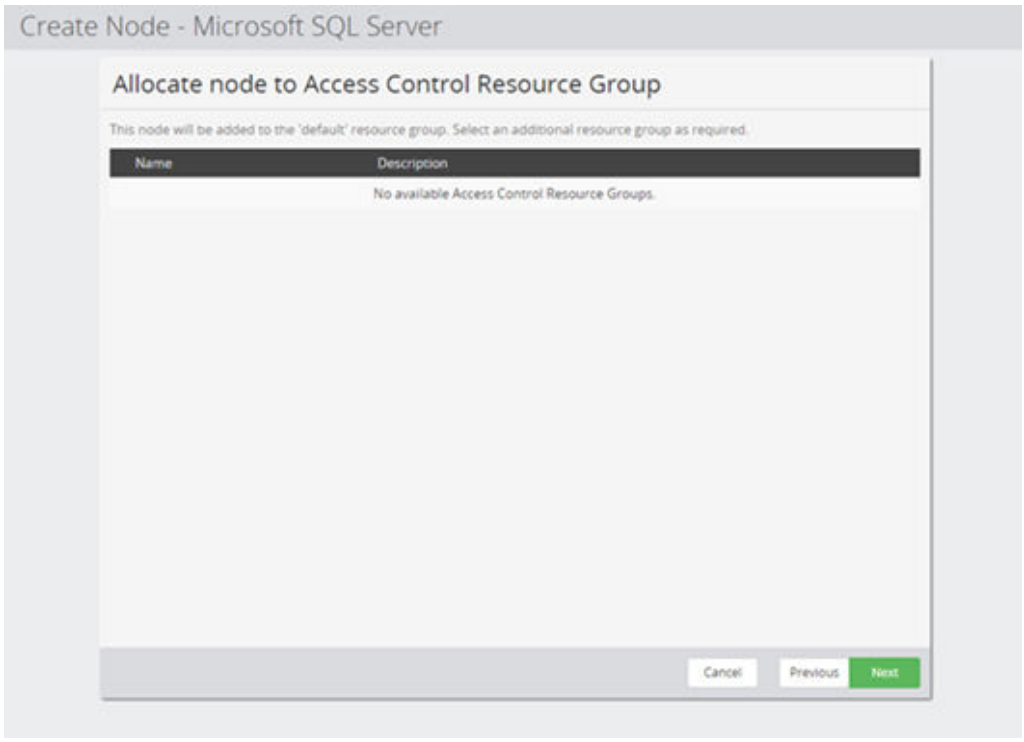


Figure 288 Microsoft SQL Node Wizard - Allocate node to Access Control Resource Group

Control	Description
Resource Groups	Select the resource group(s) to which this node will be allocated for the purposes of RBAC. All nodes are automatically allocated to the 'default' resource group.

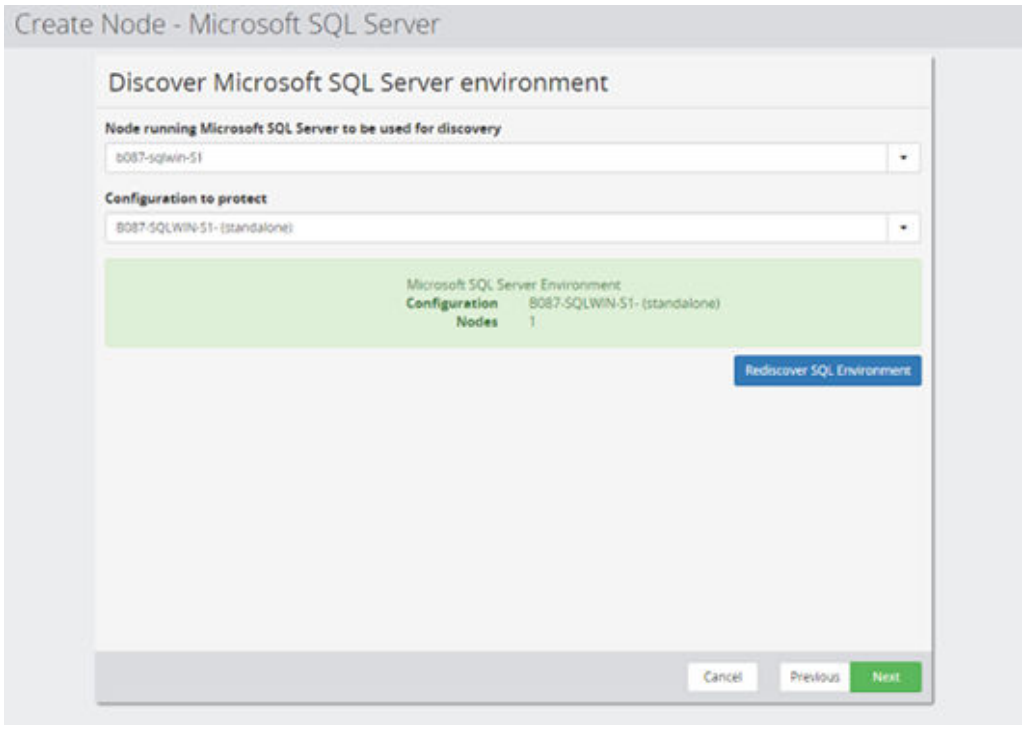


Figure 289 Microsoft SQL Node Wizard - Discover Microsoft SQL environment

Control	Description
Node running Microsoft SQL Server	Select an OS Host node which is part of the Microsoft SQL Server standalone or cluster environment
Configuration to protect	Select the Microsoft SQL Server environment you want this node to represent.
Rediscover SQL Server Environment	Click in case you want to refresh the list of available SQL Server configurations.

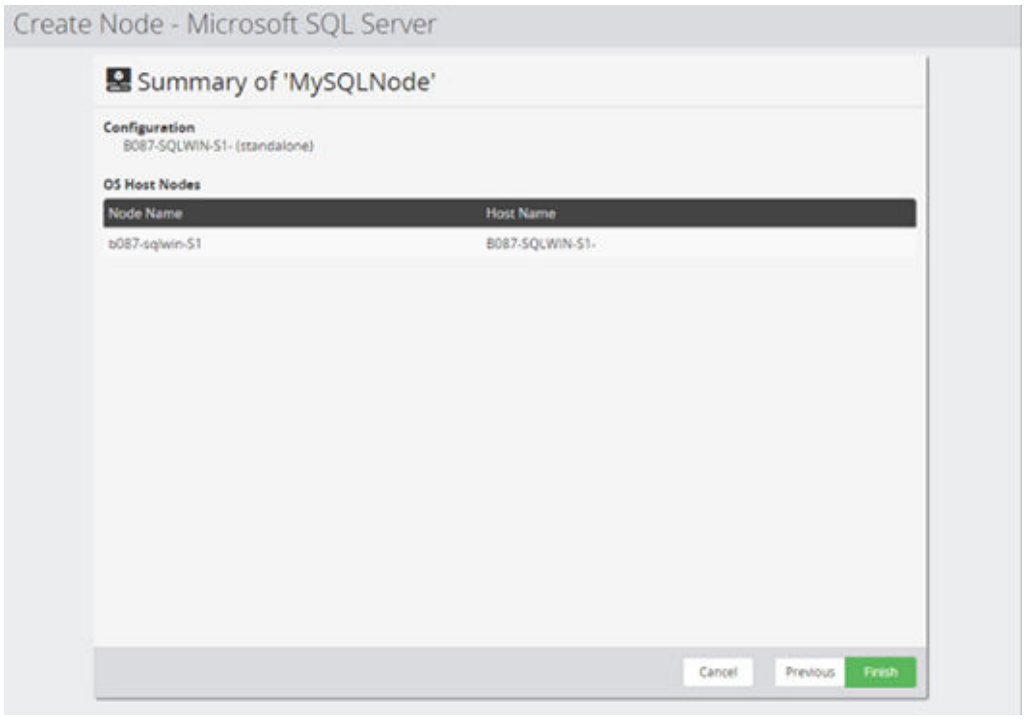


Figure 290 Microsoft SQL Node Wizard - Summary

Control	Description
Summary	Summary of the settings and nodes selected in the node creation wizard.

Oracle Application Node Wizard

This wizard is launched when a new Oracle Database Node is added to the Nodes Inventory.

Note: If you have a clustered Oracle environment and add or remove nodes to or from the cluster, the Protector Oracle application node must be updated so that the Oracle environment can be rediscovered. Any active data flows including that node must be reactivated to update the rules.

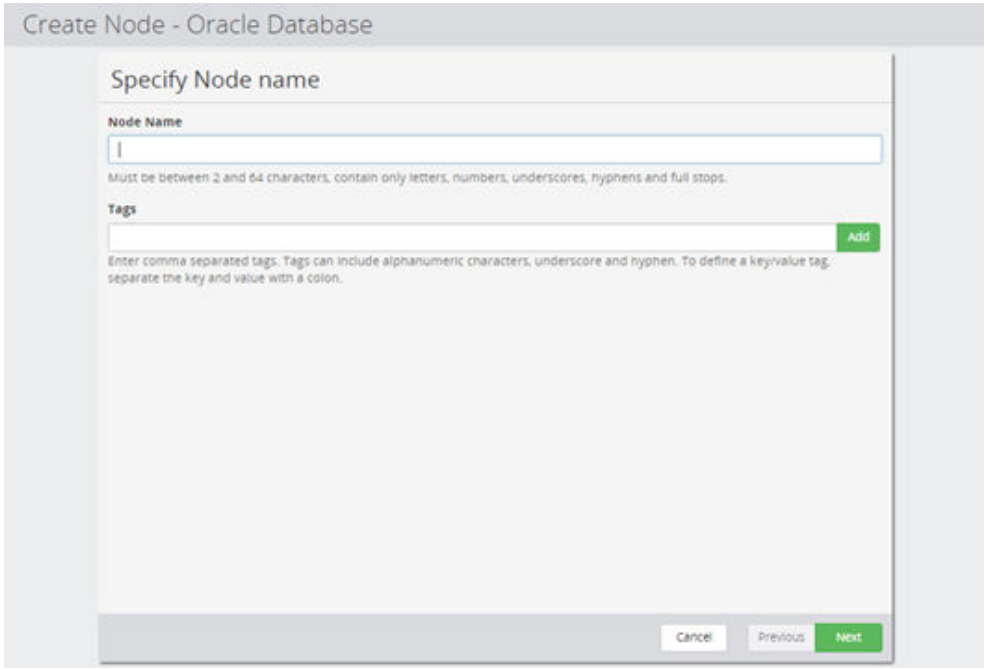


Figure 291 Oracle Database Node Wizard - Specify Node Name

Control	Description
Node Name	Enter a name for the Oracle node.
Tags	Modifies the tags of an existing object from the either the inventory screen or the details screen of the object.

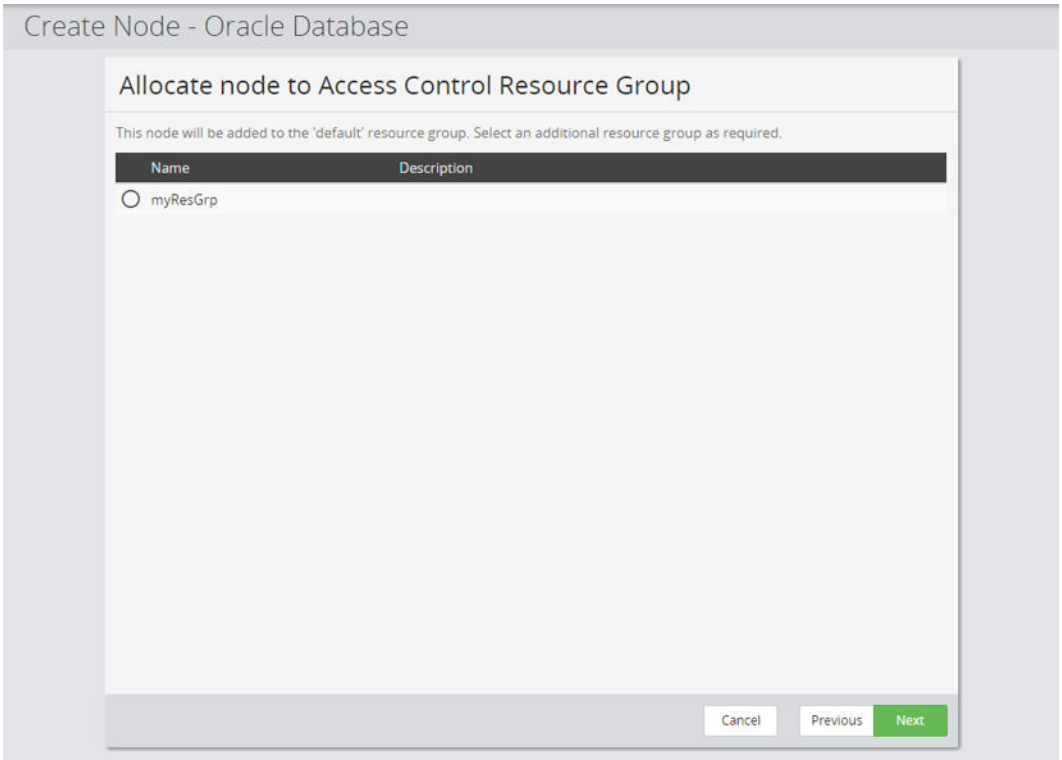


Figure 292 Oracle DB Node Wizard - Allocate node to Access Control Resource Group

Control	Description
Resource Groups	Select the resource group(s) to which this node will be allocated for the purposes of RBAC. All nodes are automatically allocated to the 'default' resource group.

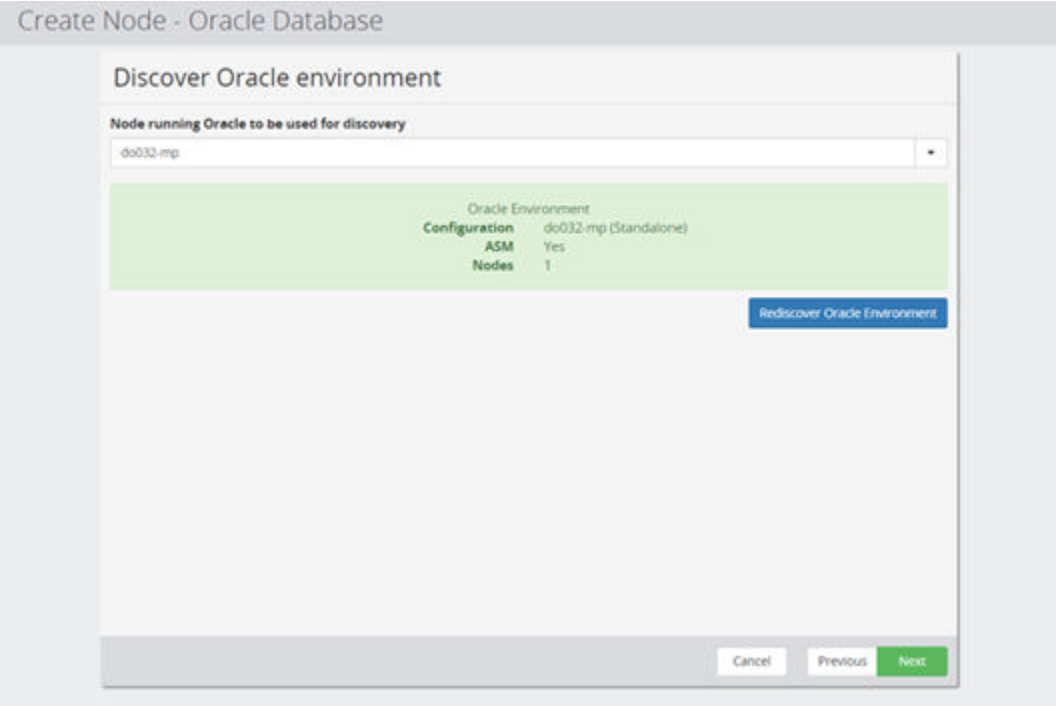


Figure 293 Oracle DB Node Wizard - Discover Oracle environment

Control	Description
Node running Oracle to be used for discovery	Select the Protector client node that communicates with the Oracle server. This node will then discover the Oracle environment.
Rediscover Oracle Environment	Click this button to refresh the cached details.

Note:

- If your Oracle setup uses ASM you can optionally specify which user is used to manage ASM.
- If your Oracle setup does not use ASM, this dialog will not be displayed.

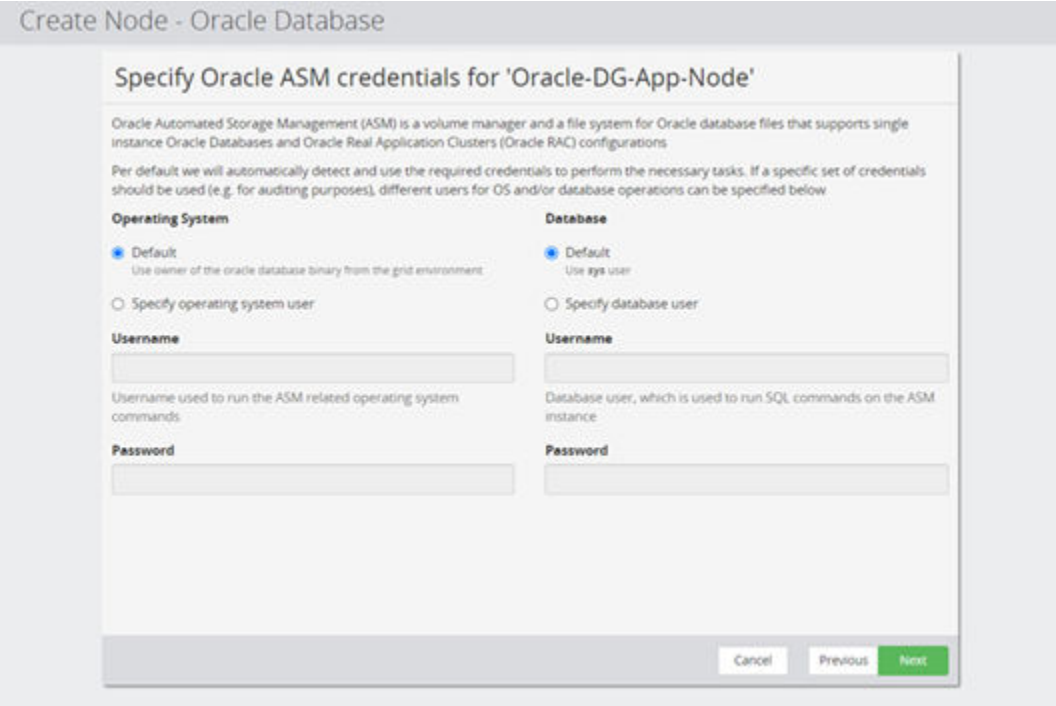


Figure 294 Oracle DB Node Wizard - Specify Oracle ASM credentials

Control	Description
Operating System	Select one of the following: ▪ Default - use the default owner of the Oracle database from the grid environment. ▪ Specify operating system user - specify the operating system user for the Oracle database.
Domain	For non-default operating system user only. If a Windows operating system is used, enter the domain name of the system to access Oracle ASM.
Username	For non-default operating system user only. Enter the Oracle Database username for the Oracle ASM.
Password	For non-default operating system user only. Enter the operating system username's password for the Oracle ASM.
Database	Select one of the following: ▪ Default - use the default database user to execute the SQL commands on the Oracle ASM instance ▪ Specify database user - specify the Oracle Database user who can execute the SQL commands on the Oracle ASM instance

Control	Description
Username	For non-default database user only. Enter the Oracle Database username of the user who can execute the SQL commands on the Oracle ASM instance.
Password	For non-default database user only. Enter the password of the user who can execute the SQL commands on the Oracle ASM instance.

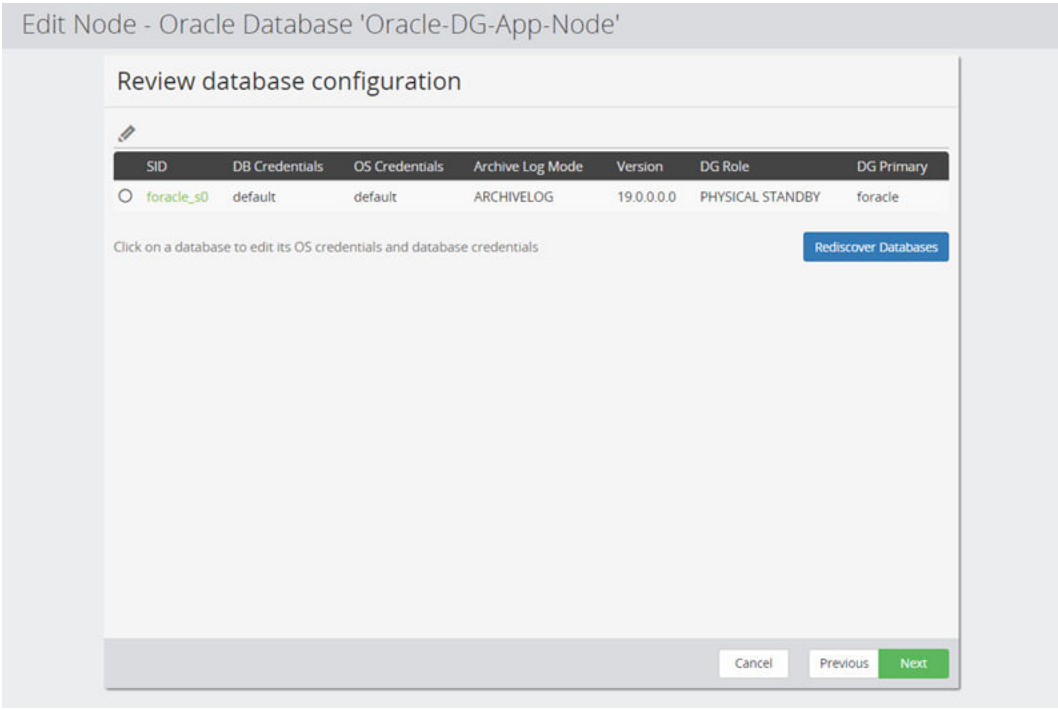


Figure 295 Oracle DB Node Wizard - Review database configuration

Control	Description
SID	Click on a database's SID to open the **Specify Database credentials** dialog (see below).
Rediscover Databases	Click this button to refresh the cached details.
DB Credentials	Displays the credentials based on the SID for database connection.
OS Credentials	Displays the OS credentials for system connection.
RMAN Catalog	Displays Oracle connection to the RMAN Catalog Database.
Archive Log Mode	Displays the archivelog information in archivelog mode.

Control	Description
Version	Displays the version number of Oracle Database.
DG Role	Displays the DataGuard roles Physical Standby, Logical Standby or Primary.
DG Primary	Displays the DataGuard Primary database name if the DG Role is Standby.

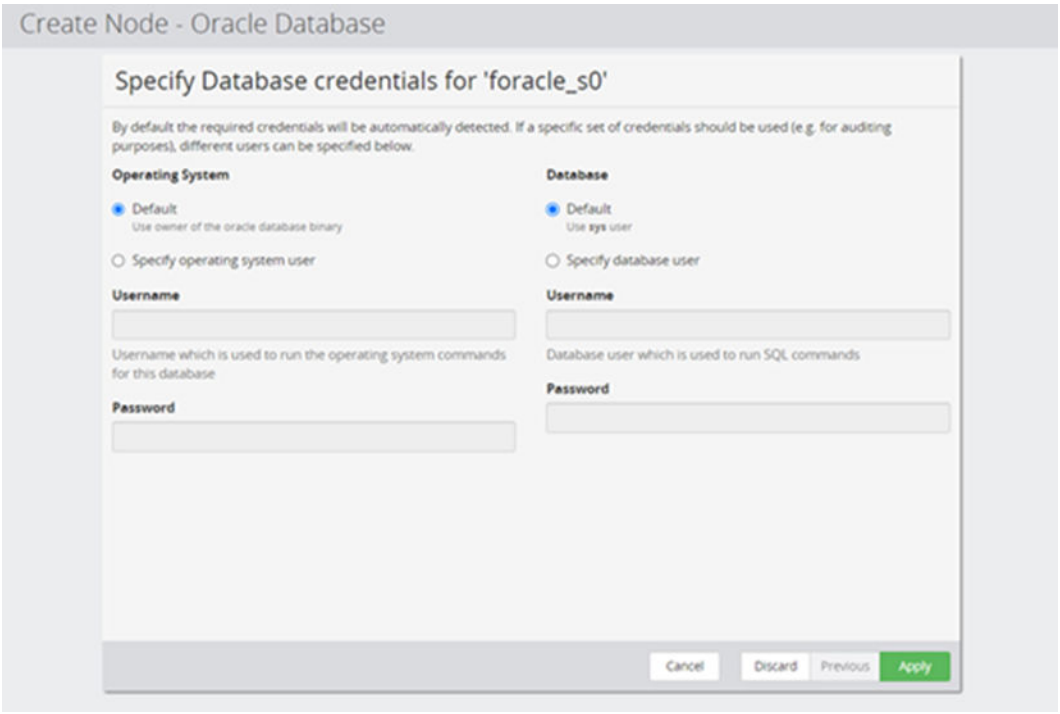


Figure 296 Oracle DB Node Wizard - Specify Database credentials

Control	Description
Operating System	Select one of the following: ▪ Default - use the owner of the Oracle database binary. ▪ Specify operating system user - specify the operating system user.
Domain	For non-default operating system user only. If a Windows operating system is used, enter the domain name of the system to access Oracle ASM.
Username	For non-default operating system user only. Enter the username which is used to run the operating system commands for this database.

Control	Description
Password	For non-default operating system user only. Enter the username's password.
Database	Select one of the following: ▪ Default - use the default database user to execute the SQL commands on the Oracle ASM instance ▪ Specify database user - specify the database user who can execute the SQL commands on the Oracle ASM instance
Username	For non-default database user only. Enter the username of the user who can execute the SQL commands.
Password	For non-default database user only. Enter the password of the user who can execute the SQL commands.

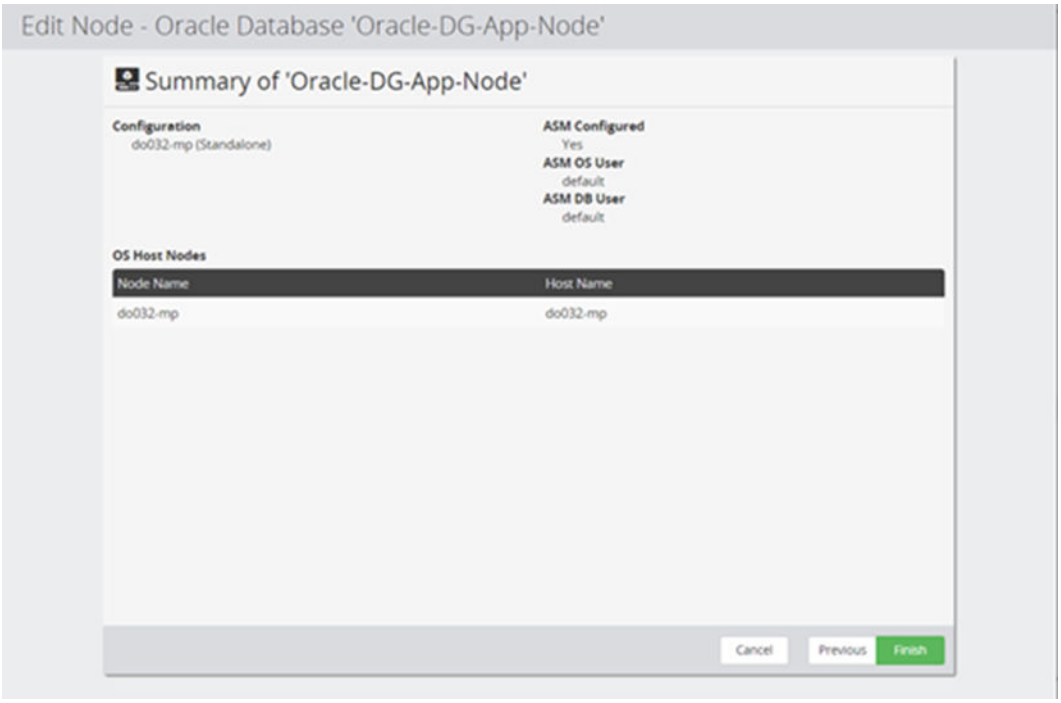


Figure 297 Oracle DB Node Wizard - Summary

Control	Description
Summary	Summary of the selected configuration.

Hitachi Block Host Node Wizard

A Hitachi Block Host node represents a group of Hitachi Block volumes attached to a host machine.

Note: A physical Block Device node must be created using the Hitachi Block Device Node Wizard (on page 583) before any Block Host nodes can be created on it.

Note: A Block Host node can also be based off of a Logical Block Device, this must also have been created before setting up the Block Host node using the Hitachi Logical Block Device Node Wizard (on page 609).

If a user is given access to a Block Host node by placing it in a resource group that they have access to, the associated physical Block Device node must also be placed in that resource group. It is not sufficient to include only a Logical Block Device node (based on the physical Block Device node) in that resource group.

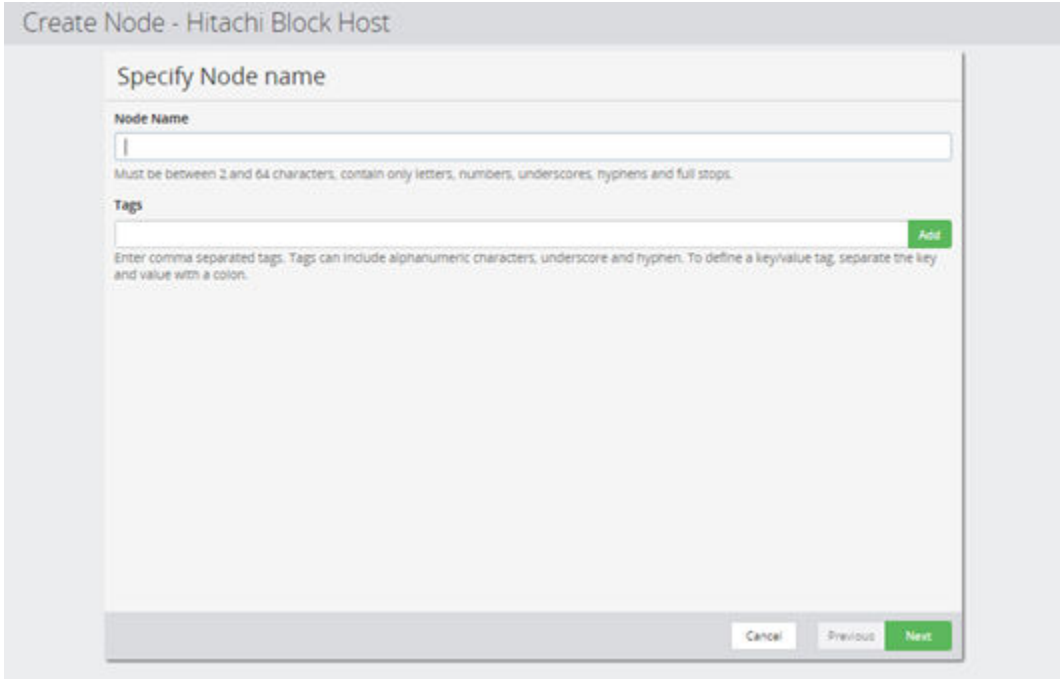


Figure 298 Hitachi Block Host Wizard - Specify Node name

Control	Description
Node Name	Enter a name for the Block Host node.
Tags	Add the tags to be associated with the object being created.

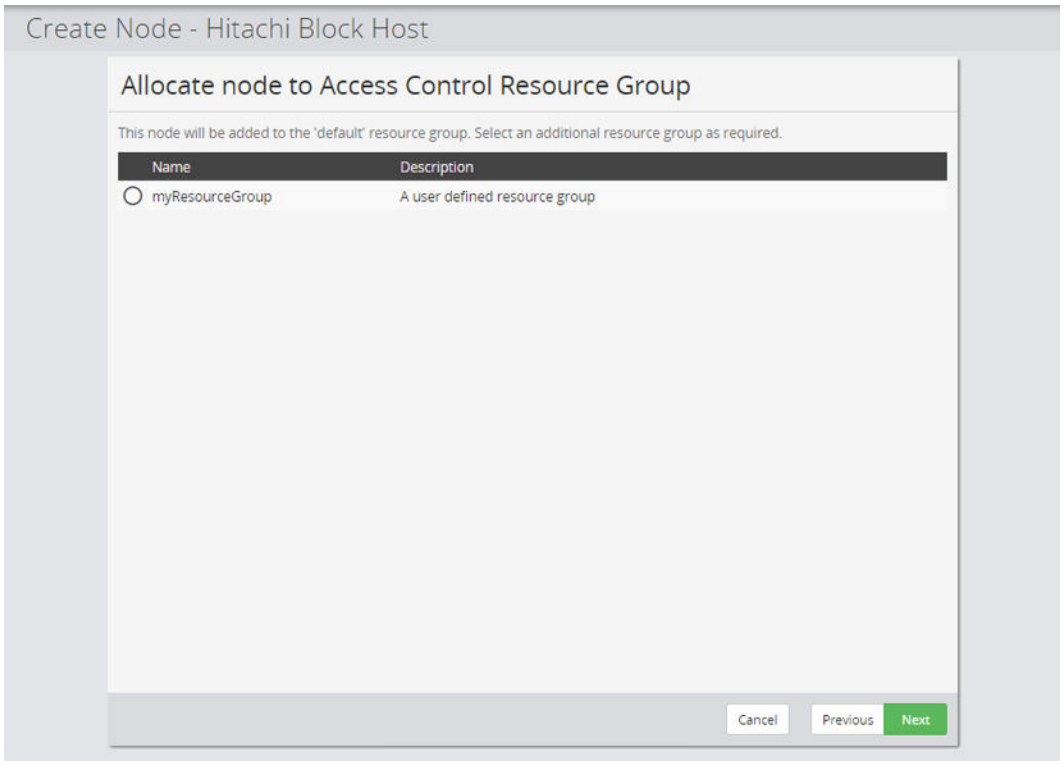


Figure 299 Hitachi Block Host Wizard - Allocate node to Access Control Resource Group

Control	Description
Resource Groups	Select the resource group(s) to which this node will be allocated for the purposes of RBAC. All nodes are automatically allocated to the 'default' resource group.

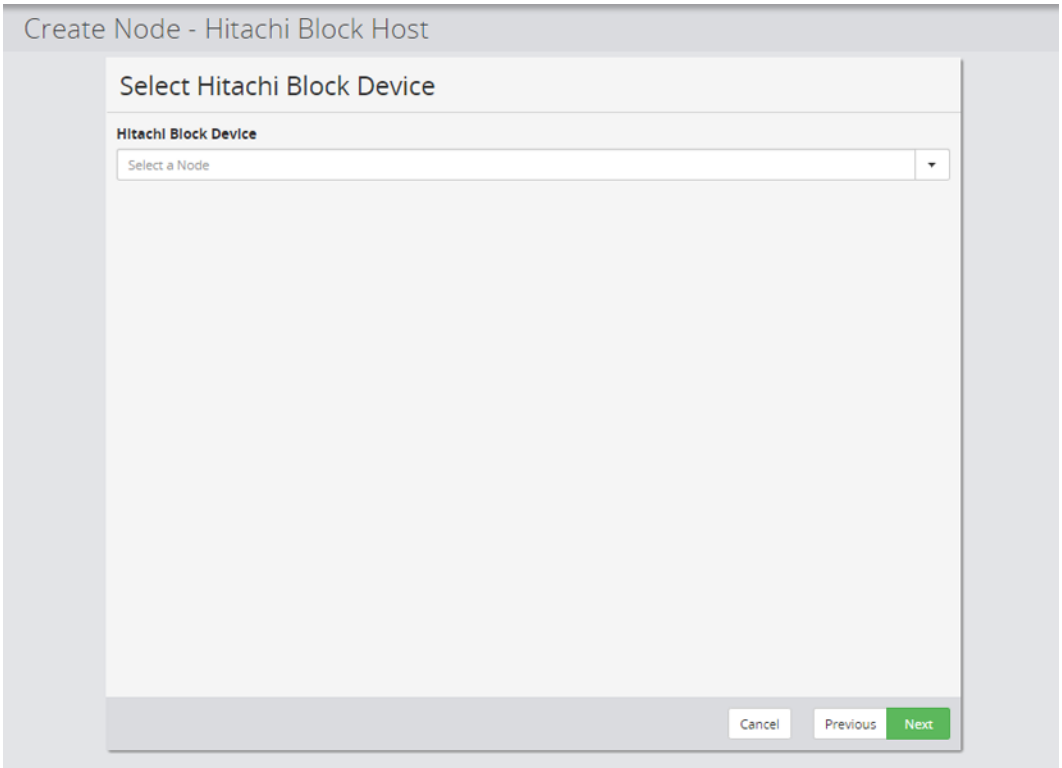


Figure 300 Hitachi Block Host Wizard - Select Hitachi Block Device

Control	Description
Hitachi Block Device	Enter or select a Hitachi Block node or Hitachi Logical Block node name from the dropdown list.

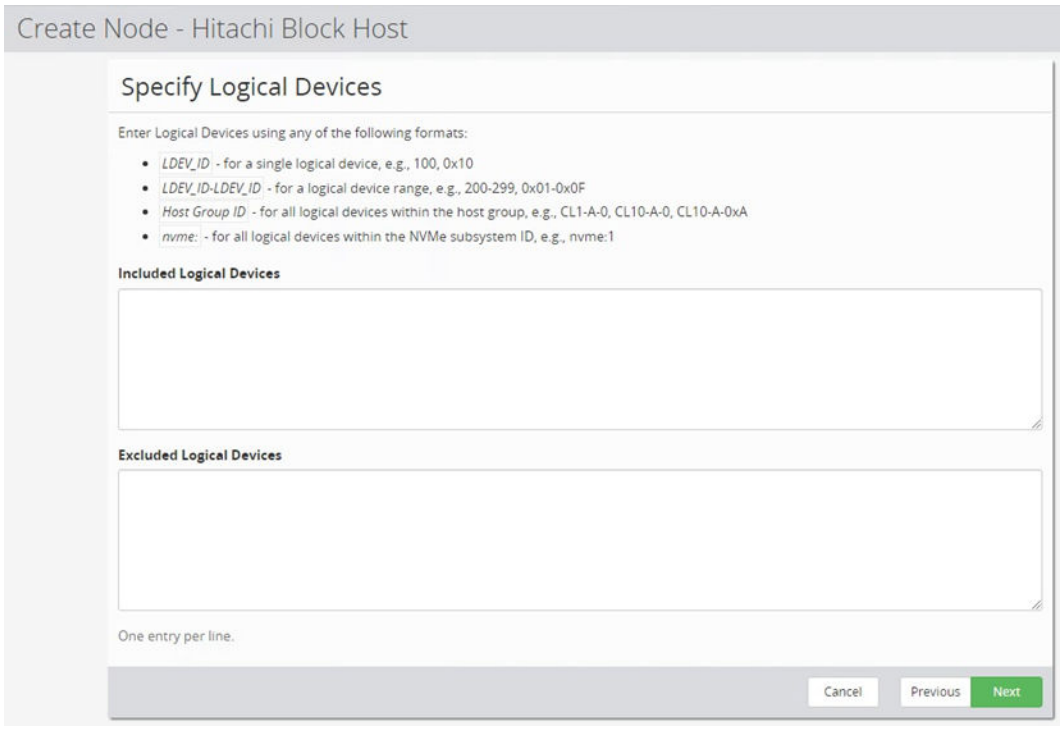


Figure 301 Hitachi Block Host Wizard - Specify Logical Devices

Control	Description
Included Logical Devices	Enter a list of LDEVs or Host Groups to be included in this Block Host (one logical device or host group specification per line) in the format described below.
Excluded Logical Devices	Enter a list of LDEVs or Host Groups to be excluded from this Block Host (one logical device or host group specification per line) in the format described below. **Tip:** This enables the inclusion of an entire Host Group but exclusion of individual LDEVs within that Host Group.

LDEVS are entered using the format:

identifier

Where *identifier* is one of:

- *ldev_id* a single LDEV ID in hex or decimal
- *ldev_id-ldev_id* a range of LDEV IDs
- *host_group_id* a host group identifier using the format:
  - `CL`*c-s-h* where:
    - `CL` is a literal string
    - *c* is the physical channel number in the range 1…n
    - *s* is the physical slot number in the range A…Z
    - *h* is the logical host group ID in the range 0…255

For example:

`0x00a0` for an LDEV specified in hexadecimal

`220-230` for a range of LDEVs specified in decimal

`CL2-A-12` for a host group

Note:

- Spaces and colons must not be present in the entries. If a space or colon is encountered on a line, the remaining text on that line from the space or colon will be disregarded. This allows entries that contain the LDEV name after a space or colon) to be present in the list.

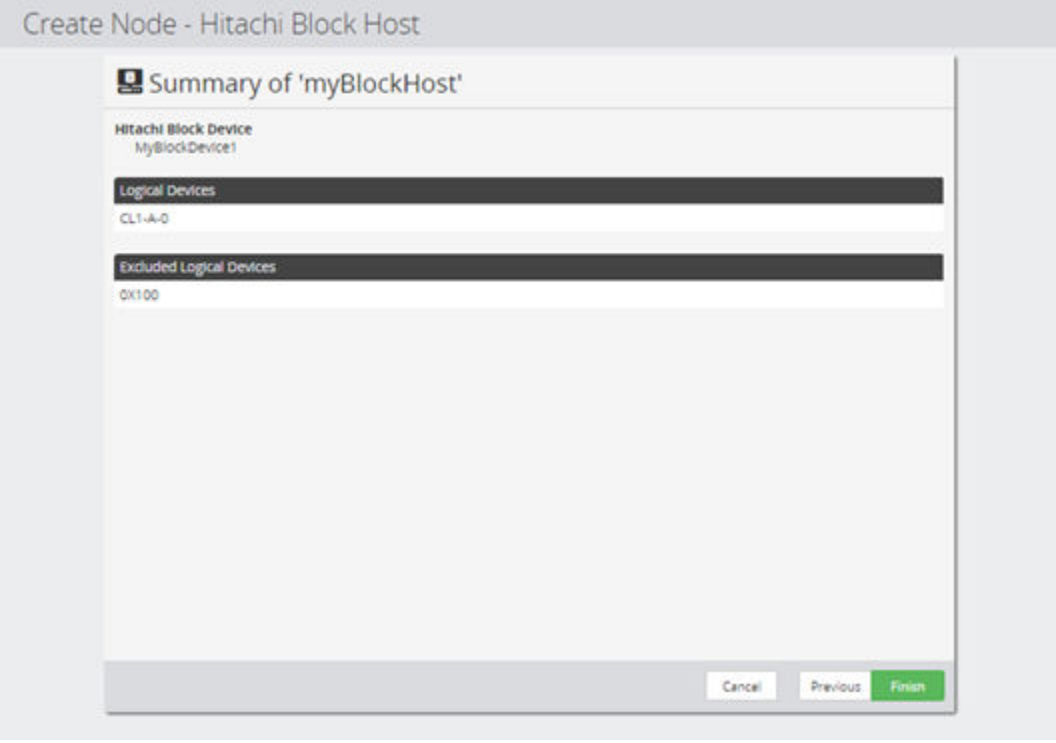


Figure 302 Hitachi Block Host Wizard - Summary

Control	Description
Summary	Summary of the settings entered.

Hitachi Software Block Host Node Wizard

A Hitachi Software Block Host node represents a group of Hitachi Block volumes attached to a host node.

Note: A physical Software Block Device node must be created using the Hitachi Software Block Device Node Wizard (on page 626) before and Software Block Host nodes can be created on it.

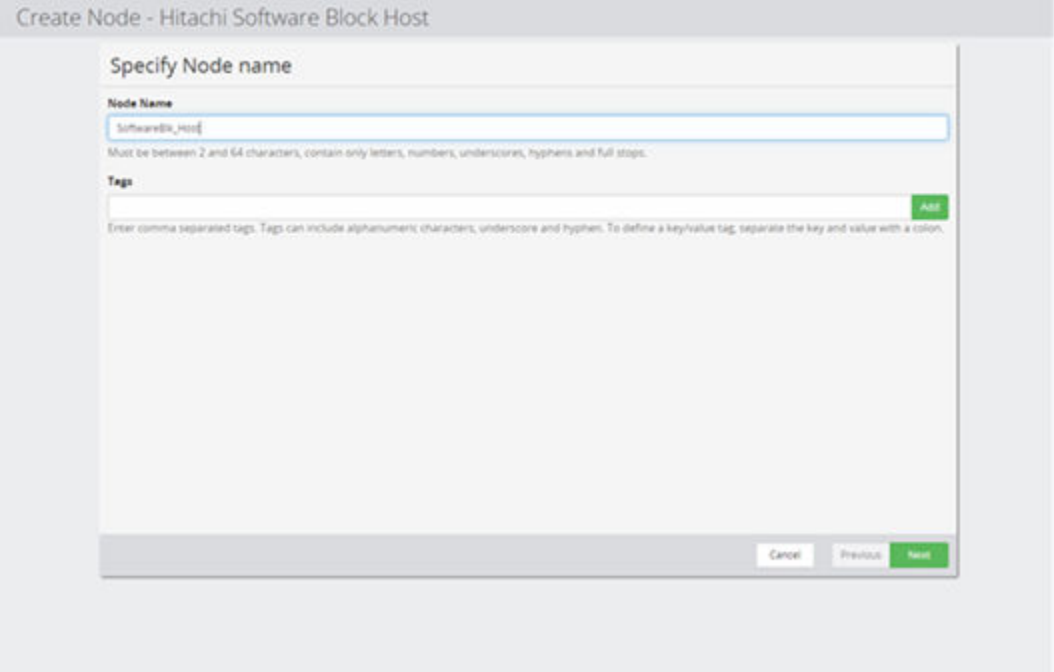


Figure 303 Hitachi Software Block Host Wizard - Specify Node name

Control	Description
Node Name	Enter a name for the Software Block Host node.
Tags	Add the tags to be associated with the object being created.

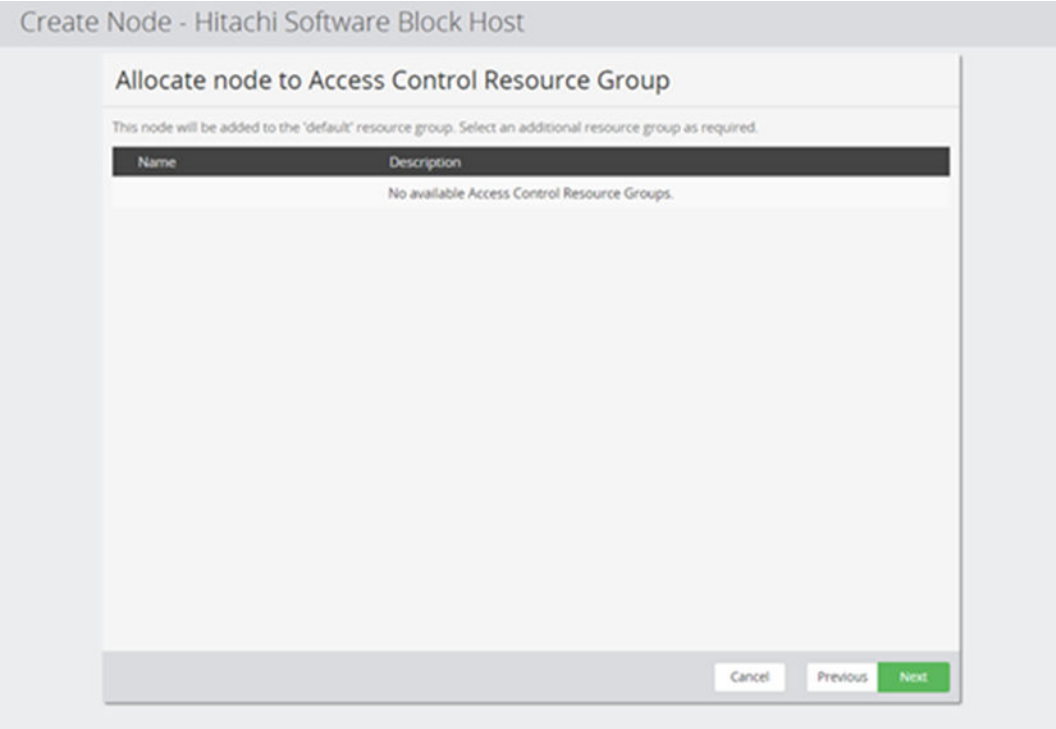


Figure 304 Hitachi Software Block Host Wizard - Allocate node to Access Control Resource Group

Control	Description
Resource Groups	Select the resource group(s) to which this node will be allocated for the purposes of RBAC. All nodes are automatically allocated to the 'default' resource group.

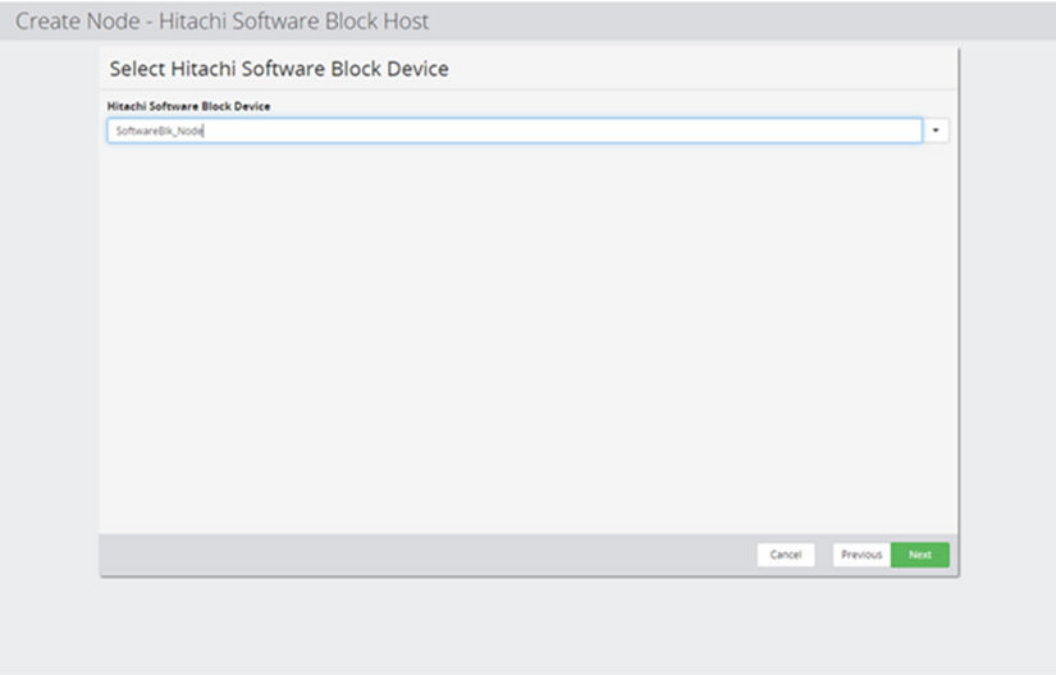


Figure 305 Hitachi Software Block Host Wizard - Select Hitachi Software Block Device

Control	Description
Hitachi Block Device	Enter or select a Hitachi Software Block node name from the dropdown list.

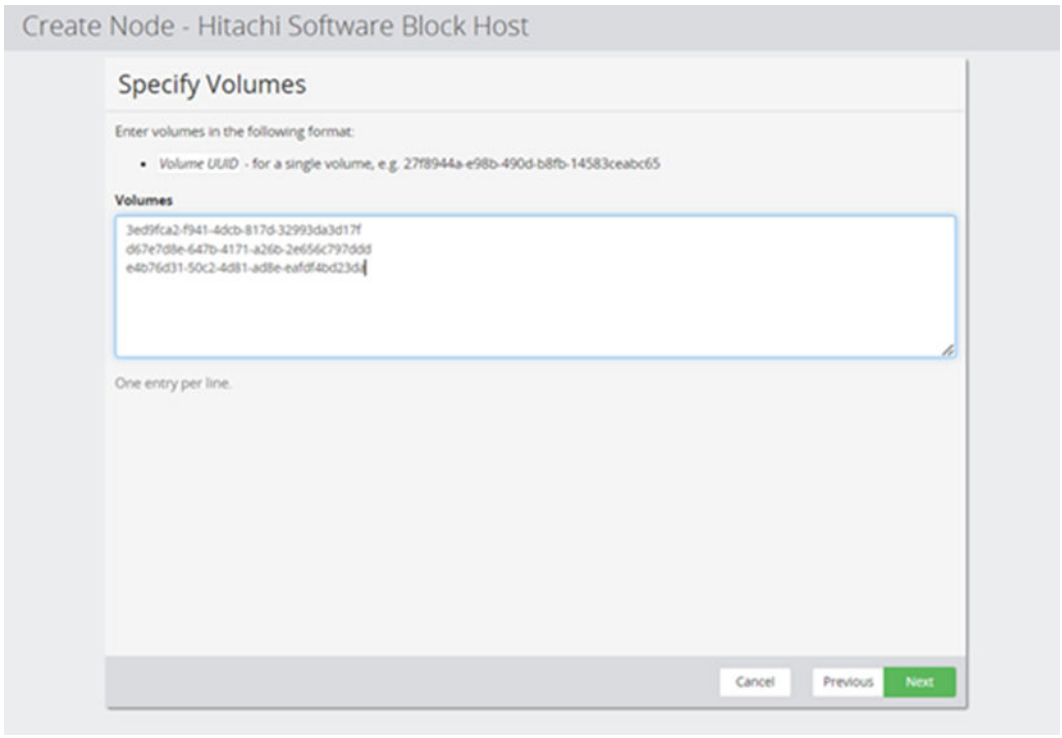


Figure 306 Hitachi Software Block Host Wizard - Specify Volumes

Control	Description
Volumes	Specify the volume ids associated with the Software Block device. **Note:** The volume IDs as those as displayed in the SDSB or SDSC storage UI.

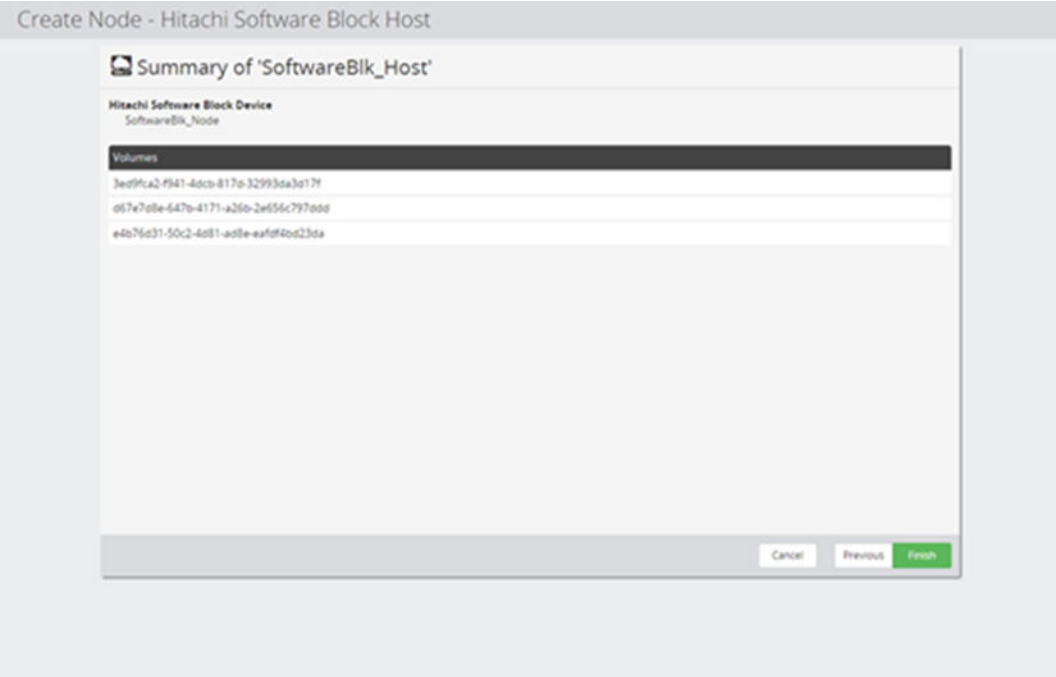


Figure 307 Hitachi Software Block Host Wizard - Summary

Control	Description
Summary	Summary of the settings entered.

Hyper-V Node Wizard

This wizard is launched when a new Hyper-V Node is added to the Nodes Inventory.

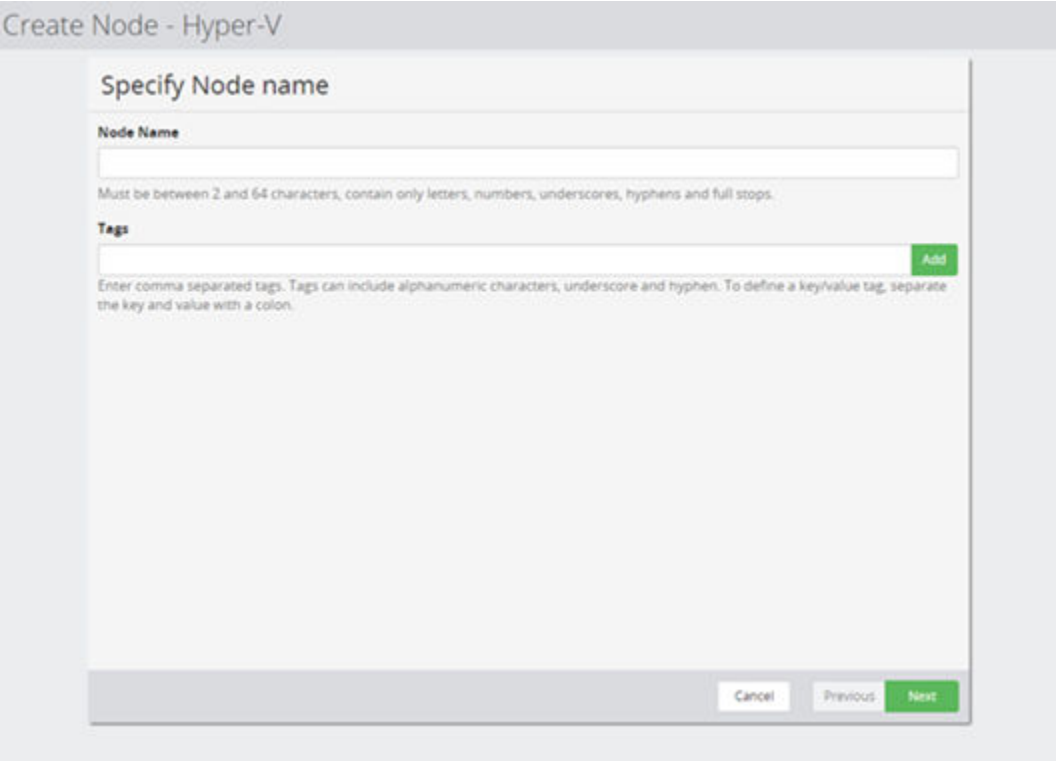


Figure 308 Hyper-V Node Wizard - Specify Node name

Control	Description
Node Name	Enter a name for the Hyper-V node.
Tags	Add the tags to be associated with the object being created

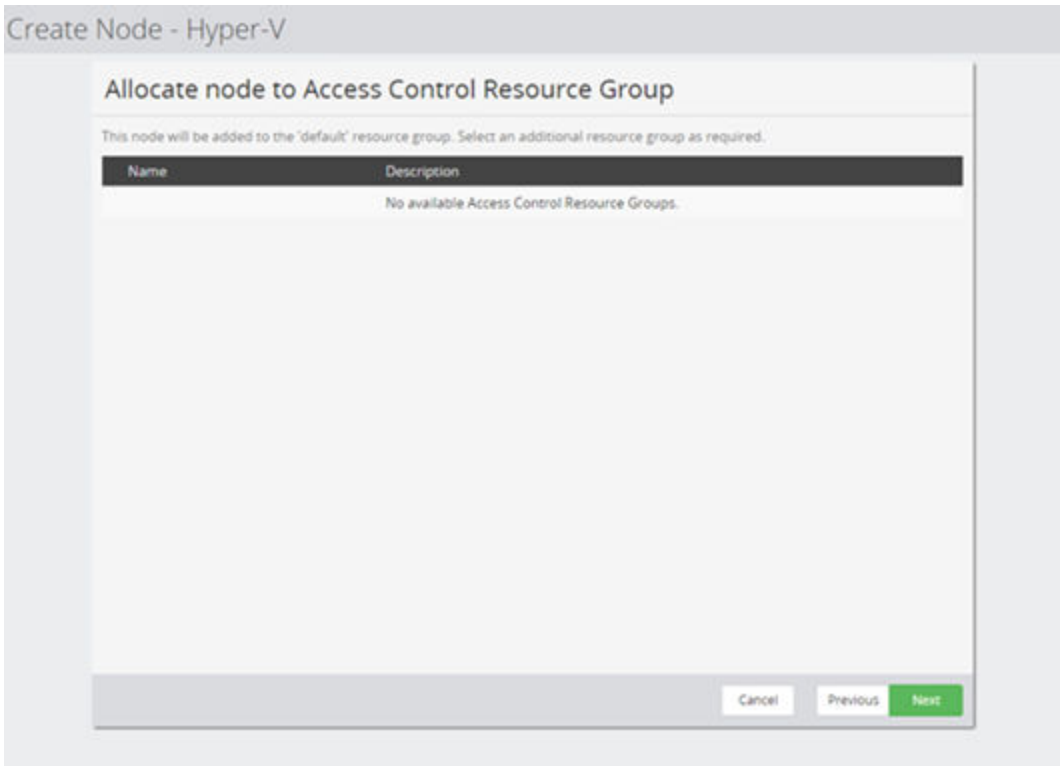


Figure 309 Hyper-V Node Wizard - Allocate node to Access Control Resource Group

Control	Description
Resource Groups	Select the resource group(s) to which this node will be allocated for the purposes of RBAC. All nodes are automatically allocated to the 'default' resource group.

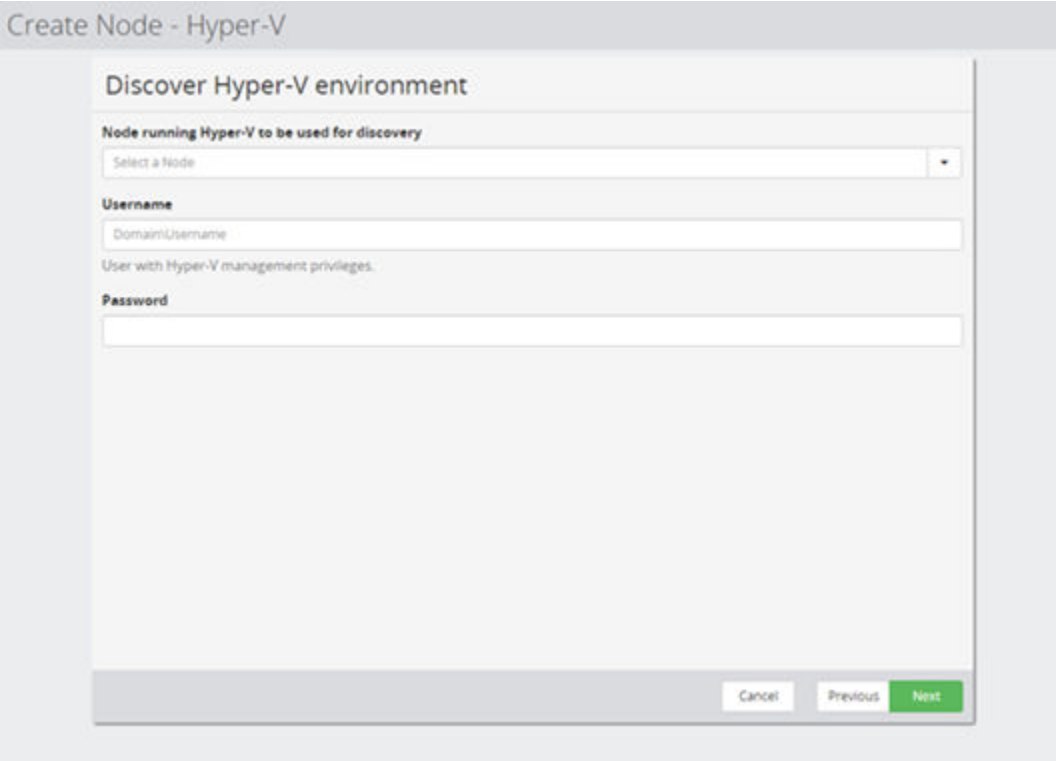


Figure 310 Hyper-V Node Wizard - Select node running Hyper-V server

Control	Description
Node running Hyper-V	Select an OS Host node which is part of the Microsoft Hyper-V standalone or cluster environment
Username	Enter the username that will be used to perform backups and restores on this Hyper-V environment. The user requires privileges as detailed in Microsoft Hyper-V user privileges (on page 573).
Password	Enter the password for the username provided above

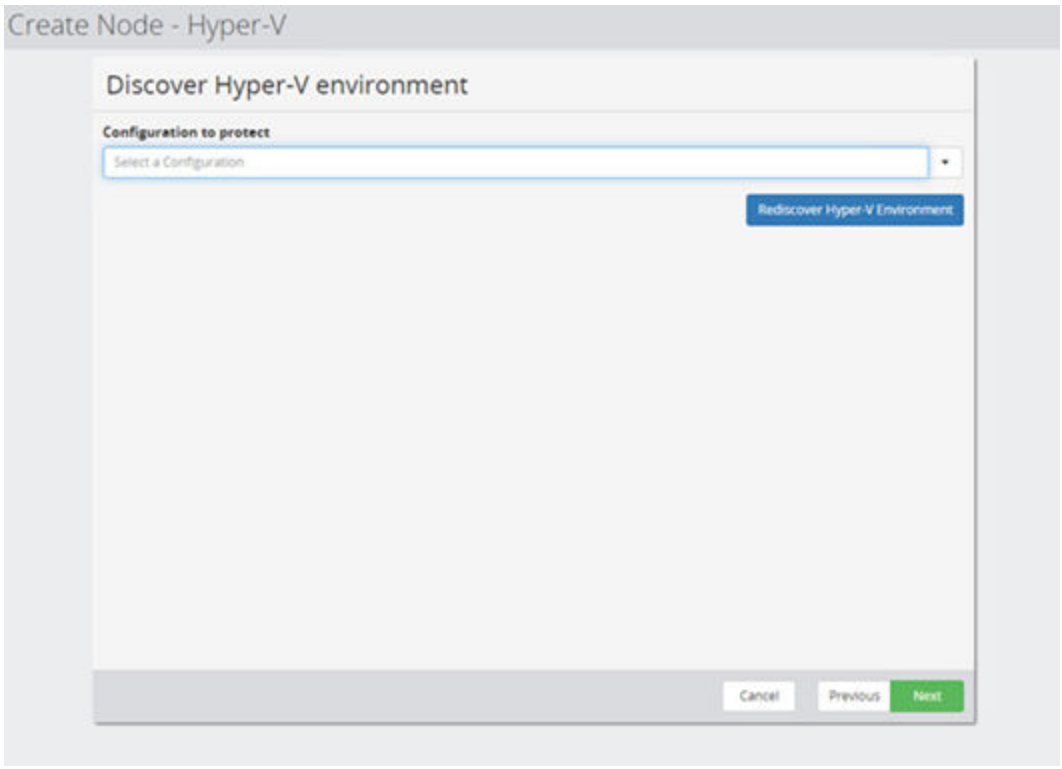


Figure 311 Hyper-V Node Wizard - Discover Hyper-V environment

Control	Description
Configuration to protect	Select the Hyper-V environment you want this node to represent
Rediscover Hyper-V Environment	Click in case you want to refresh the list of available Hyper-V configuration

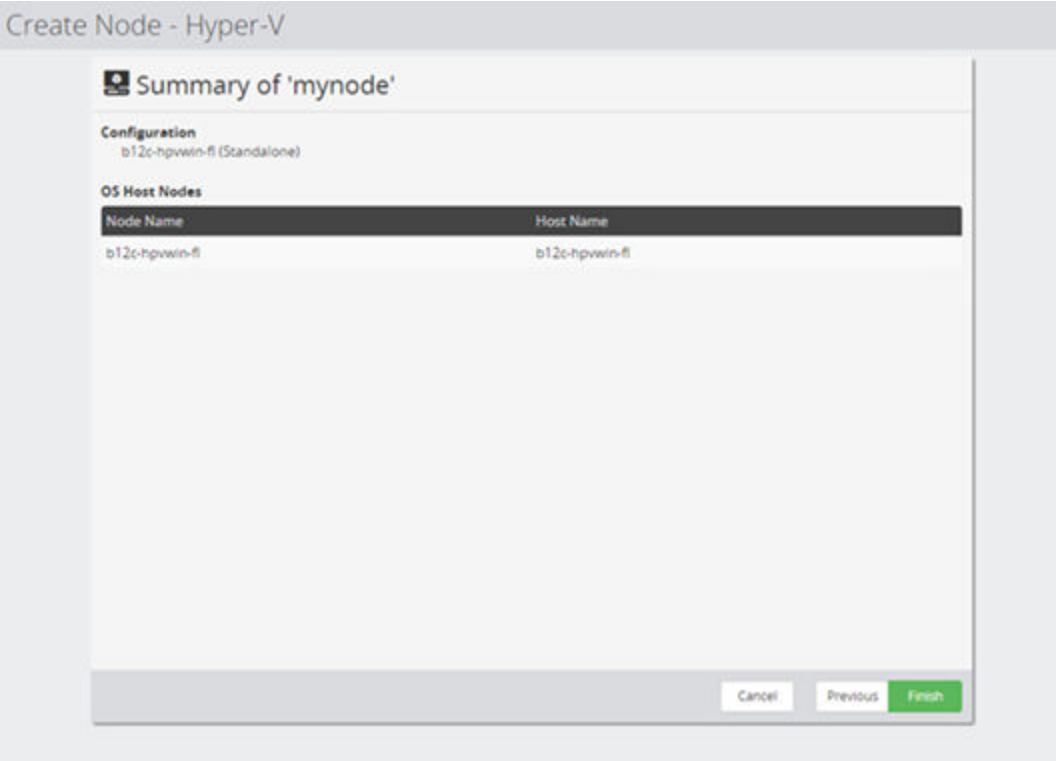


Figure 312 Hyper-V Node Wizard - Summary

Control	Description
Summary	Summary of the settings entered.

Microsoft Hyper-V user privileges

To lookup, protect and restore Hyper-V virtual machines Protector requires credentials that are valid for all machines comprising the Hyper-V node.

In case of a cluster, these credentials must be domain credentials. For a standalone configuration, a local user is sufficient.

The user needs to be a member of the following groups in the domain and all nodes of the Hyper-V setup:

- Users
- Domain Users
- Backup Operators
- Hyper-V Administrators
- Remote Management Users

In addition, for a cluster the user requires the permission to administrate the cluster. The following command needs to be executed on every cluster node:

```
Grant-ClusterAccess -User domain\username -Full
```

VMware Node Wizard

This wizard is launched when a new VMware Node is added to the Nodes Inventory.

Note: If you use vCenter to manage an ESX/ESXi host, then a proxy node cannot be created for that host. Create a proxy using the managing vCenter node instead.

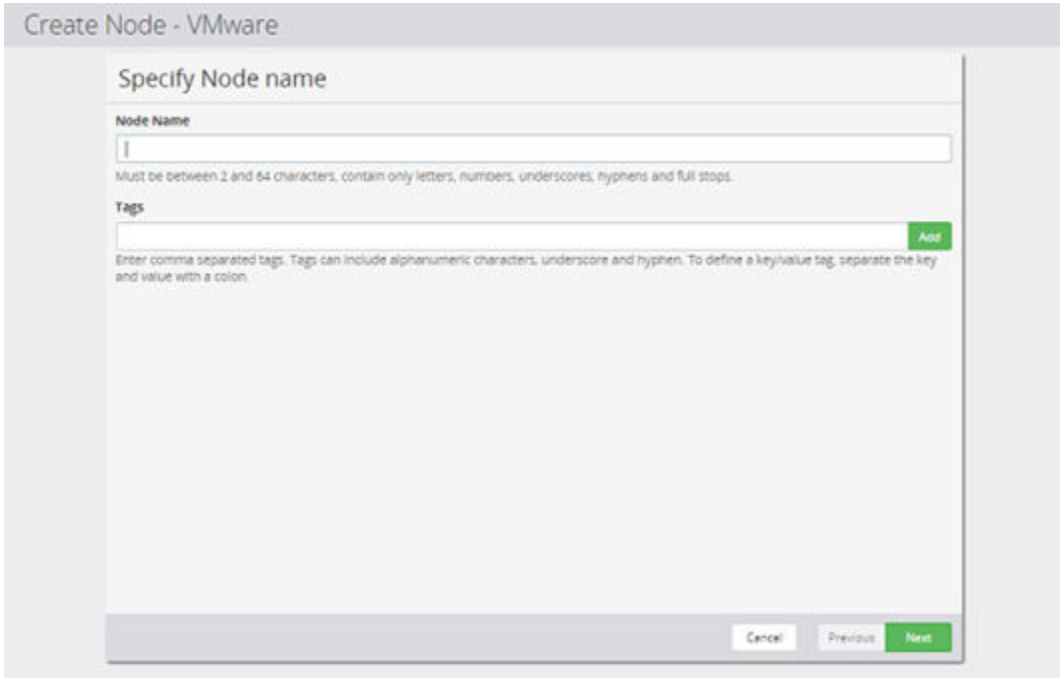


Figure 313 VMware Node Wizard - Specify Node name

Control	Description
Node Name	Enter a name for the VMware node.
Tags	Add the tags to be associated with the object being created.

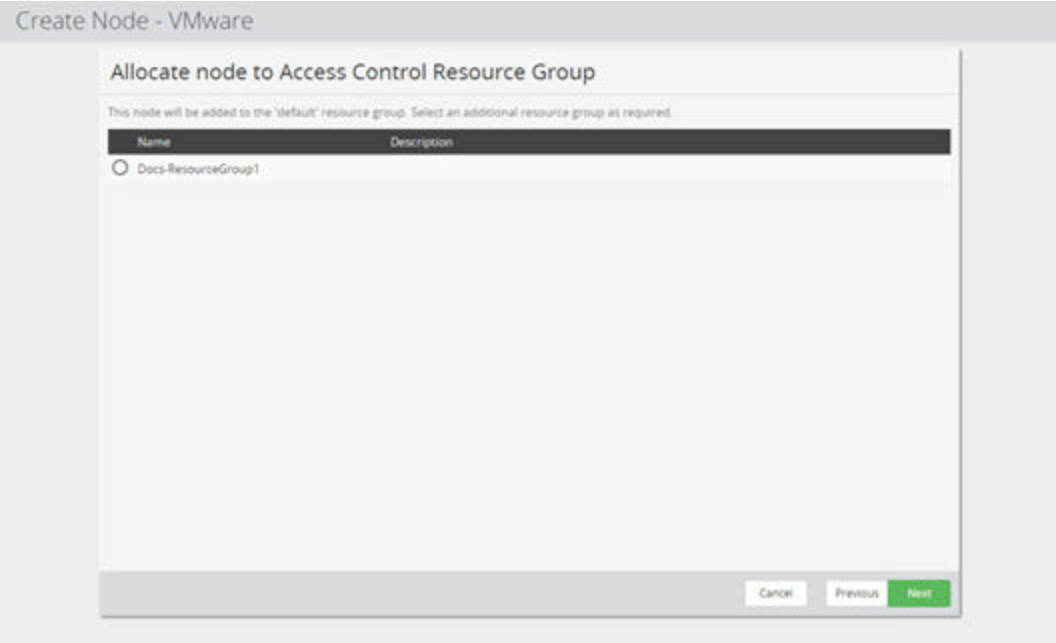


Figure 314 VMware Node Wizard - Allocate node to Access Control Resource Group

Control	Description
Resource Groups	Select the resource group(s) to which this node will be allocated for the purposes of RBAC. All nodes are automatically allocated to the 'default' resource group.

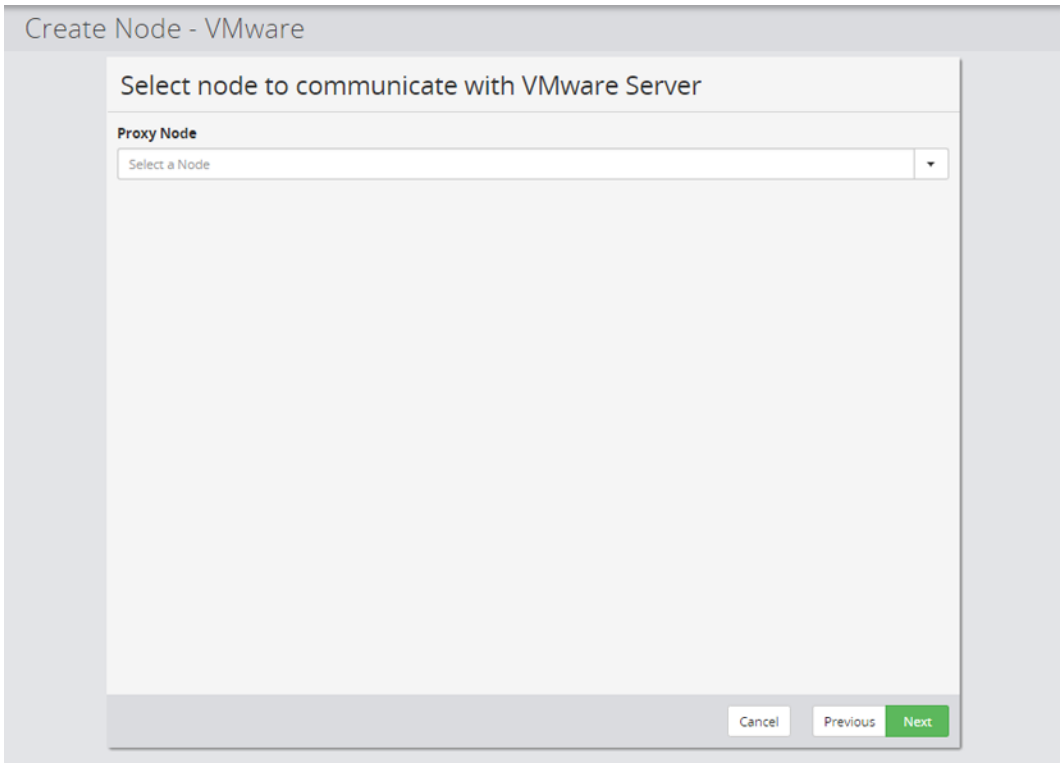


Figure 315 VMware Node Wizard - Select node to communicate with VMware Server

Control	Description
Proxy Node	Select a proxy node for the vCenter.

Control	Description
	Note: ▪ If using tags, VMware Power CLI 6.5.0 or later must be installed on the proxy node. You will need to restart the proxy node after completing the installation. ▪ If you change the proxy of a VMware node while the node is in an activate data flow, or if you change the vCenter credentials, then you must reactivate affected data flows in order for the changes to take effect. ▪ When performing host based backups, avoid excessive traffic across the network by selecting a proxy node that is as close as possible to the eventual destination of the backup data. ▪ If the proxy shares access over a SAN to disks used by the VMware datastores and this provides faster data transfer rates than the LAN (NBD), then the VMware *SAN Transport Mode* will be used during host based backup to transfer data directly from the datastores to the proxy. ▪ If the VMware node is a VM placed on the same Datastore and is a member of the same Datacenter as the VM(s) to be backed up, then this will provide faster data transfer rates than the LAN (NBD), then the VMware HotAdd Transport Mode will be used during host based backup to transfer data directly from the datastores to the proxy.

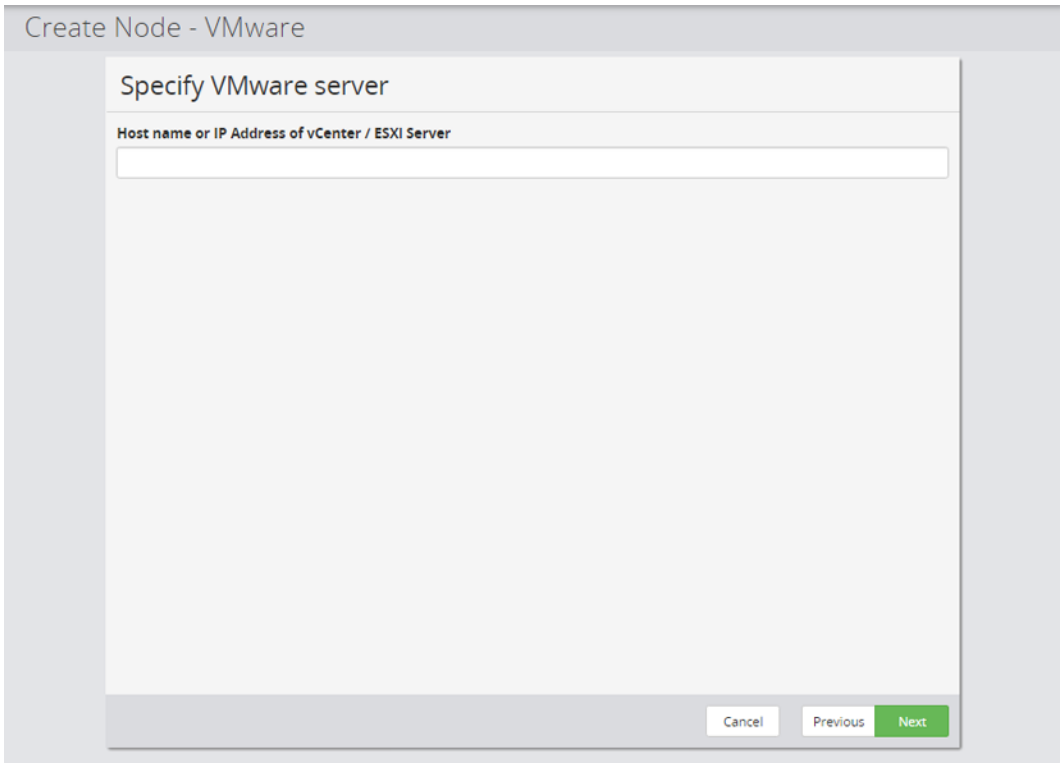


Figure 316 VMware Node Wizard - Specify VMware Server

The same VMware node can be used in both host based and block storage based data flows.

Control	Description
Host name or IP Address of vCenter Server	Specify the vCenter FQDN or IP address.

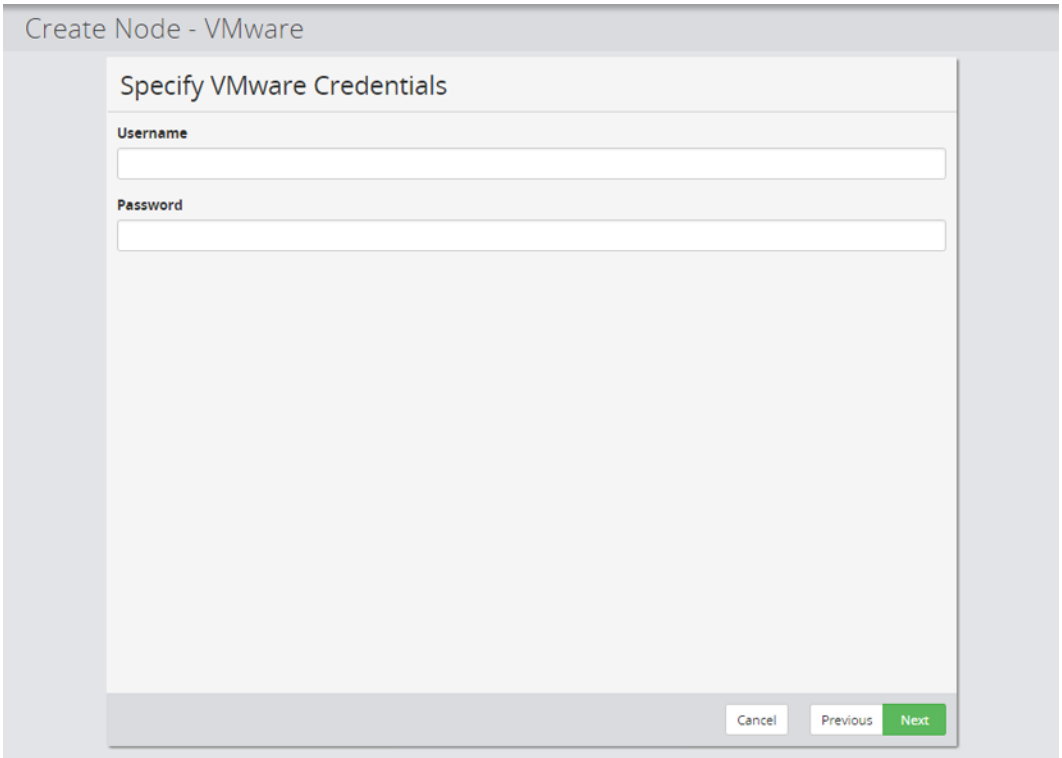


Figure 317 VMware Node Wizard - Specify VMware Credentials

Control	Description
Username	Enter the username for the vCenter. **Note:** The user specified here must have the specified VMware user privileges (on page 581).
Password	Enter the password for the vCenter.

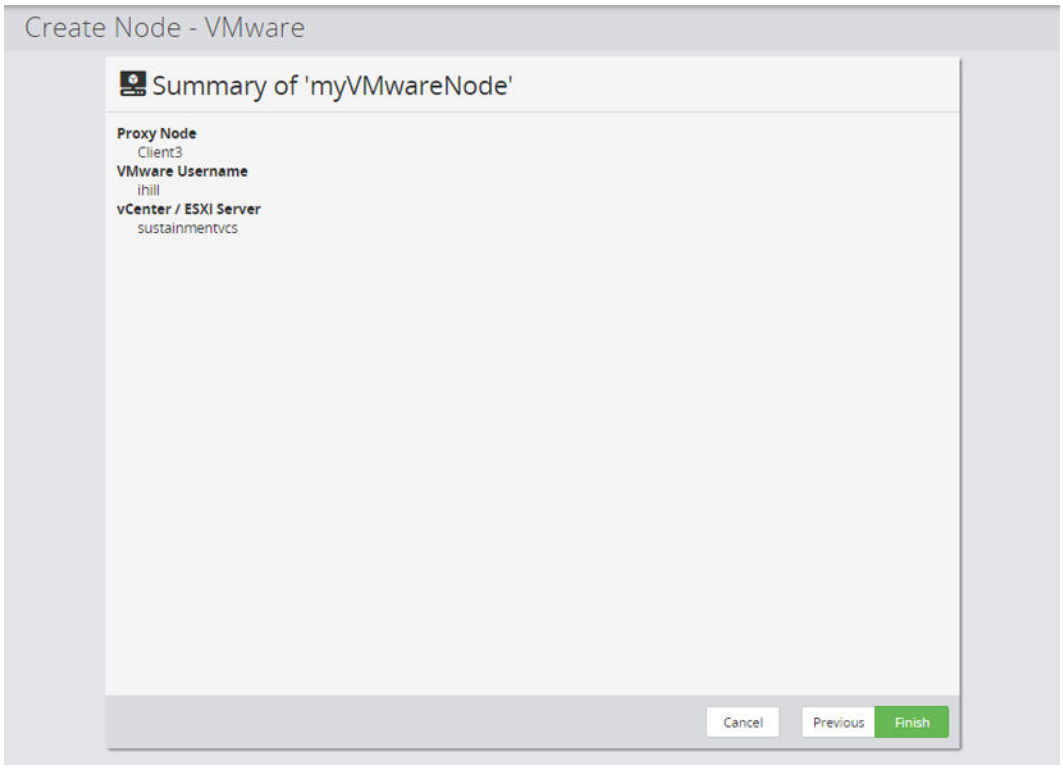


Figure 318 VMware Node Wizard - Summary

Control	Description
Summary	Summary of the settings entered.

VMware user privileges

The VMware user specified when interacting with Protector (i.e. in the context of a hypervisor proxy node or Site Recovery Manager SRA , Site Recovery Manager SRA or vRealize Orchestrator workflow) must have the following privileges assigned in vSphere:

Tip: Some privilege names have changed subtly between vSphere Client UI versions, so a little interpretation may be required. The names used here are consistent with those specified in https://techdocs.broadcom.com/us/en/vmware-cis/vsphere/vsphere/8-0/vsphere-security-8-0/defined-privileges.html

- Datastore:
  - Allocate space
  - Browse datastore
  - Low level file operations
  - Remove file
  - Rename datastore
  - Update virtual machine files
- Folder:
  - Create folder
- Global:
  - Disable methods
  - Enable methods
  - Licenses
  - Log event
  - Manage custom attributes
  - Set custom attribute
- Host:
  - Configuration:
    - Connection
    - Storage partition configuration
- Network:
  - Assign network
  - Configure
- Resource:
  - Assign virtual machine to resource pool
  - Migrate powered off virtual machine
  - Migrate powered on virtual machine
- Sessions:

- Validate session

- Virtual Machine:
  - Change Configuration:
    - Add existing disk
    - Add new disk
    - Add or remove device
    - Advanced
    - Acquire Disk lease
    - Change CPU count
    - Change resource
    - Change Memory
    - Change Settings
    - Change Swapfile placement
    - Configure Host USB device
    - Configure Raw device
    - Extend virtual disk
    - Modify device settings
    - Reload from path
    - Remove disk
    - Rename
    - Reset guest information
    - Set annotation
    - Toggle Disk change tracking
    - Upgrade virtual machine compatibility
  - Guest operations:
    - Guest operation modifications
    - Guest operation program execution
    - Guest operation queries
  - Interaction:
    - Answer question
    - Backup operation on virtual machine
    - Console interaction
    - Device connection
    - Guest operating system management by VIX API
    - Power off

        - Power on
    - Edit Inventory:
        - Create from existing
        - Create new
        - Register
        - Remove
        - Unregister
    - Provisioning:
        - Allow disk access
        - Allow read-only disk access
        - Allow virtual machine download
        - Allow virtual machine files upload
        - Mark as template
        - Mark as virtual machine
    - Snapshot management:
        - Create snapshot
        - Remove snapshot
        - Revert to snapshot
- dvPort group:
    - Create
    - Delete
- vApp:
    - Add virtual machine
    - Assign resource pool
    - Unregister
- vSphere Tagging:
    - Assign or Unassign vSphere Tag
    - Assign or Unassign vSphere Tag on Object

The System privileges (Anonymous, Read and View) are also required. These are automatically assigned to new and existing roles, but are not visible in the vSphere Client UI.

Hitachi Block Device Node Wizard

This wizard is launched when a new Hitachi Block Node is added to the Nodes Inventory.

Note:

A Hitachi Block Device Node can not share a proxy with a Hitachi Software Block Device Node.

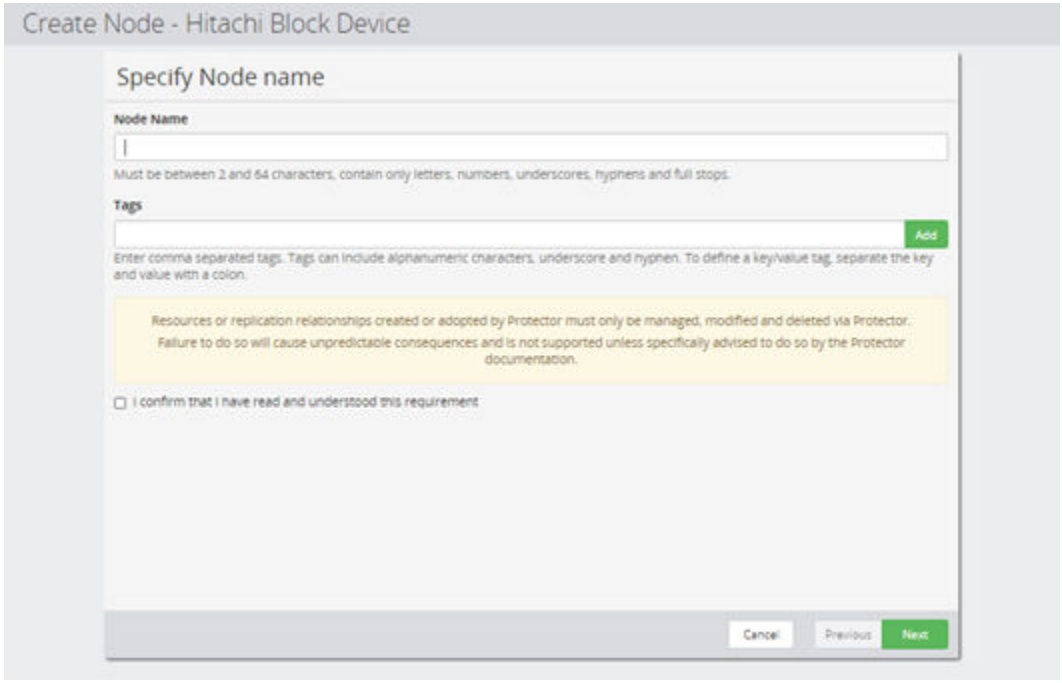


Figure 319 Hitachi Block Device Wizard - Specify Node Name

Control	Description
Node Name	Enter a name for the Hitachi Block storage node.
Tags	Add the tags to be associated with the object being created.
I confirm that ...	This checkbox must be checked to proceed with the node configuration. **Caution:** Resources or replication relationships created or adopted by Protector must only be managed, modified and deleted via Protector. Failure to do so will cause unpredictable consequences and is not supported unless specifically advised to do so by the Protector documentation.

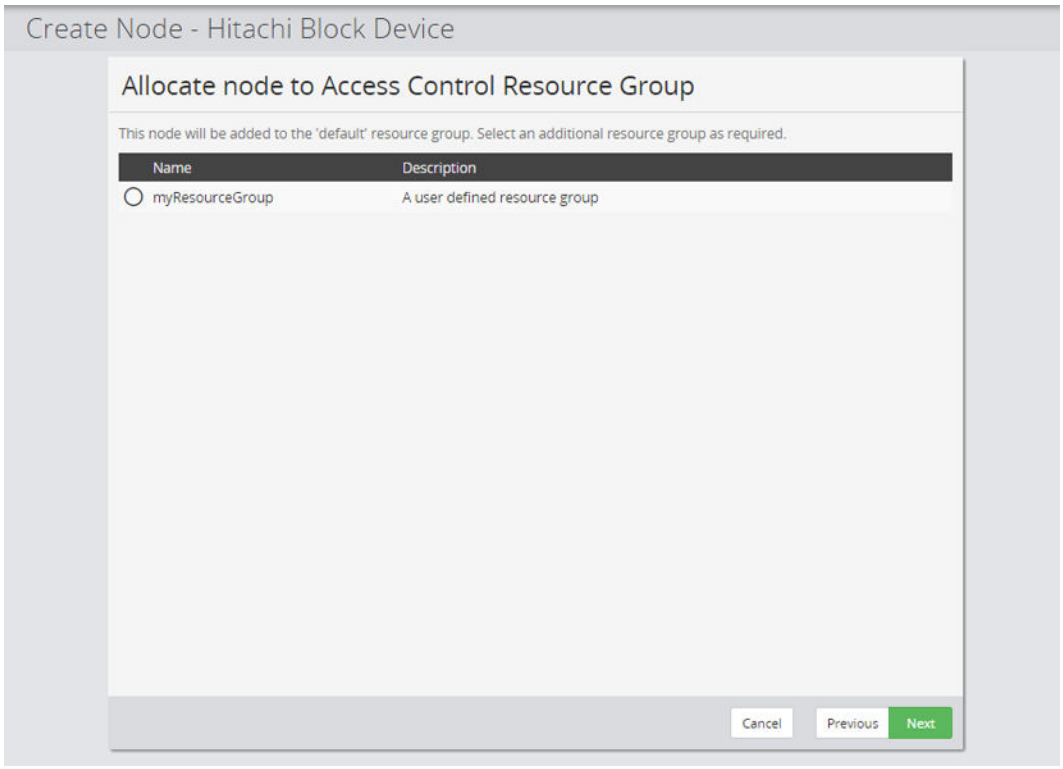


Figure 320 Hitachi Block Device Wizard - Allocate node to Access Control Resource Group

Control	Description
Resource Groups	Select the resource group(s) to which this node will be allocated for the purposes of RBAC. All nodes are automatically allocated to the 'default' resource group.

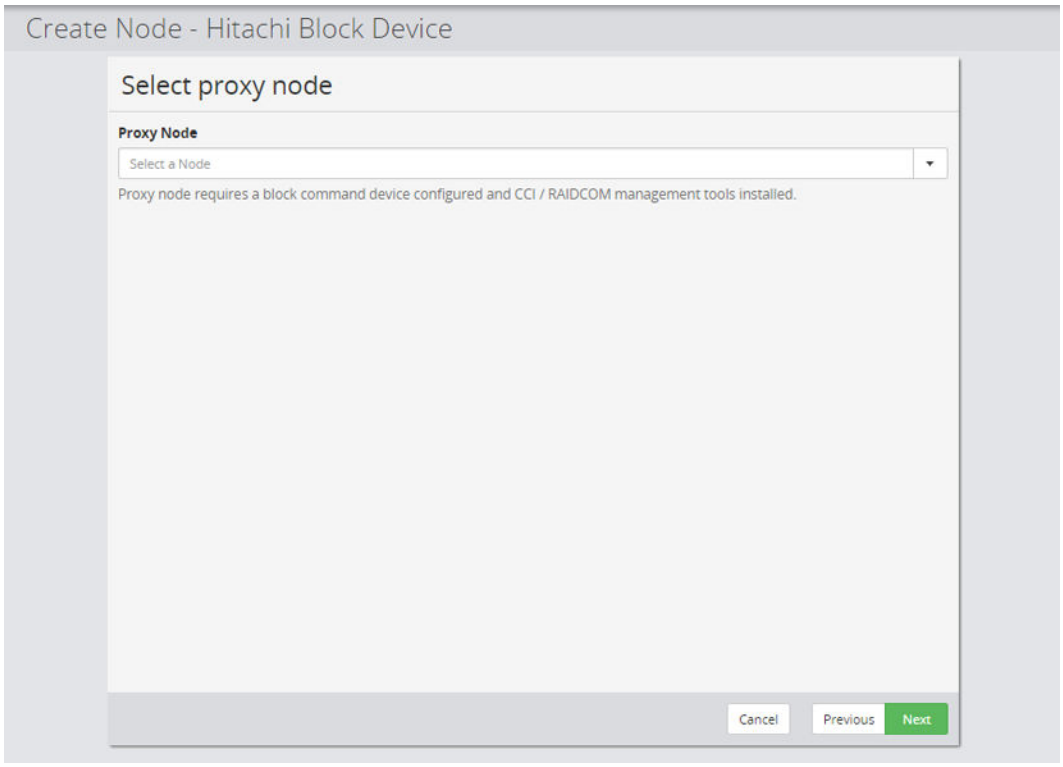


Figure 321 Hitachi Block Device Wizard - Select proxy node

Control	Description
Proxy Node	Select a Protector node to act as a proxy. **Caution:** ISM nodes and their associated CMDs used to control storage devices must not be shared with applications other than those forming part of the Ops Center suite. **Note:** The proxy node is responsible for interfacing with the Block storage device. It can be a Windows or Linux machine with the Protector Client software installed and must be connected via a command device to the Block storage device. The command device must *only* have user authentication enabled. The proxy node must have supported version of CCI installed. Refer to the Command Control Interface documentation available at https://knowledge.hitachivantara.com. **WARNING:** A Software Block Device Node can not share the same proxy node as a Block device node.

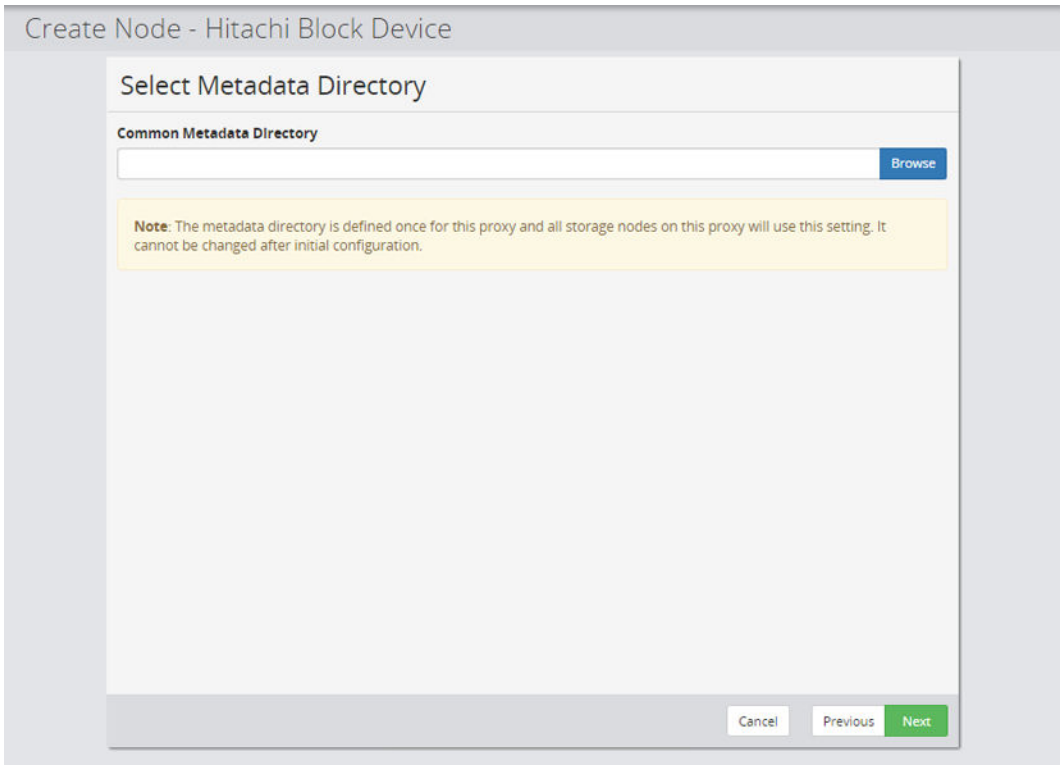


Figure 322 Hitachi Block Device Wizard - Select Metadata Directory

Control	Description
Common Metadata Directory	Enter a directory on the Proxy node where Protector can place metadata files related to Block snapshots and replications. Click the Browse button to open the Path Dialog (on page 350) if required. **Note:** The metadata directory is defined once for this proxy and all storage nodes on this proxy will use this setting. It cannot be changed after initial configuration.

Note: This screen is only displayed during node creation, it is not displayed in edit mode.

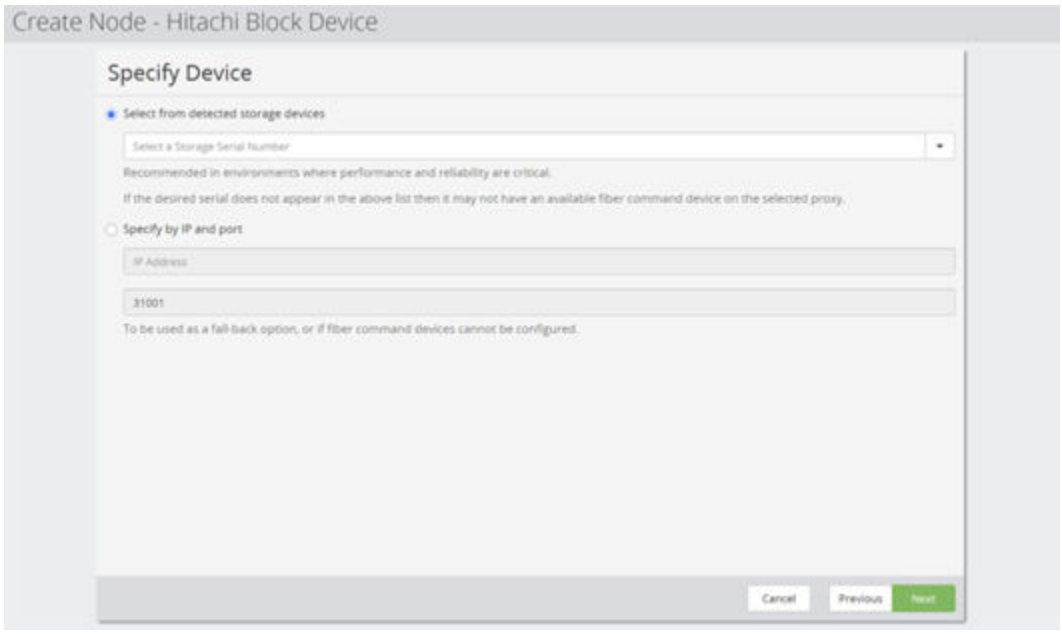


Figure 323 Hitachi Block Device Wizard - Specify Device

Control	Description
Select from detected storage devices	Select this option to specify the hardware storage device by serial number. A list of storage device serial numbers, available to the proxy node selected in the previous step, is displayed in the dropdown menu below.
Specify by IP or Hostname with a port	Select this option to specify the hardware storage device using the IP address or hostname and port number of an IP command device on the storage device. Additional In-band and Out-of-band command devices can be added at a step later in the wizard. **Note:** For VSP G1x00 storage devices, use the IP address of CTL1 or CTL2. DO NOT use the IP address of the SVP. For VSP One Bxx storage devices, use the IP address of CTL01 or CTL02. DO NOT use the ESM oe Service IP address. Later in the wizard you can add the second controller as an alternative address for failover purposes. When performance is important or where high workloads are expected, the use of In-band Channel command devices is highly recommended. Using Out-of-band command devices will result in slower performance.

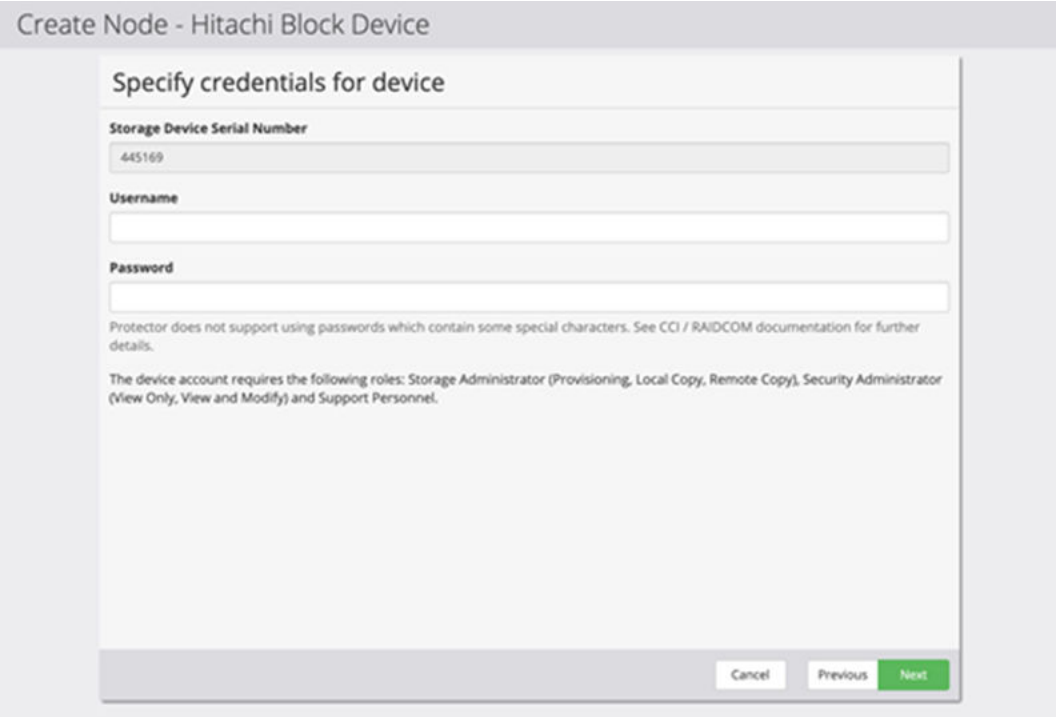


Figure 324 Hitachi Block Device Wizard - Specify credentials for device

Control	Description
Storage Device Serial Number	Displays the serial number of the Block device specified in the previous step. **Note:** When the serial number is displayed it is the serial number as provided by CCI, as such the serial number for the VSP 5000 series will have a “5” added to the beginning. Whereas the serial number for the VSP G1x00 and VSP F1500 will have a 3 added to the beginning.
Username	Enter the username for the Block device. **Note:** The username specified must be a member of the *Storage Administrator (Provisioning, Local Copy, Remote Copy)* and *Security Administrator (View Only, View and Modify)* on the Block device. In addition, the *Support Personnel* role is required for changing the Array SOM settings that allow LDEV resizing for replications. If the Block device cannot be accessed or its credentials are invalid then the node will fail authorization. The configuration wizard can be reopened to correct any errors.

Control	Description
Password	Enter the password for the Block device. **Note:** The password for authorizing a Block device must contain only useable CCI command characters: `A-Za-z0-9'-./:@\_`

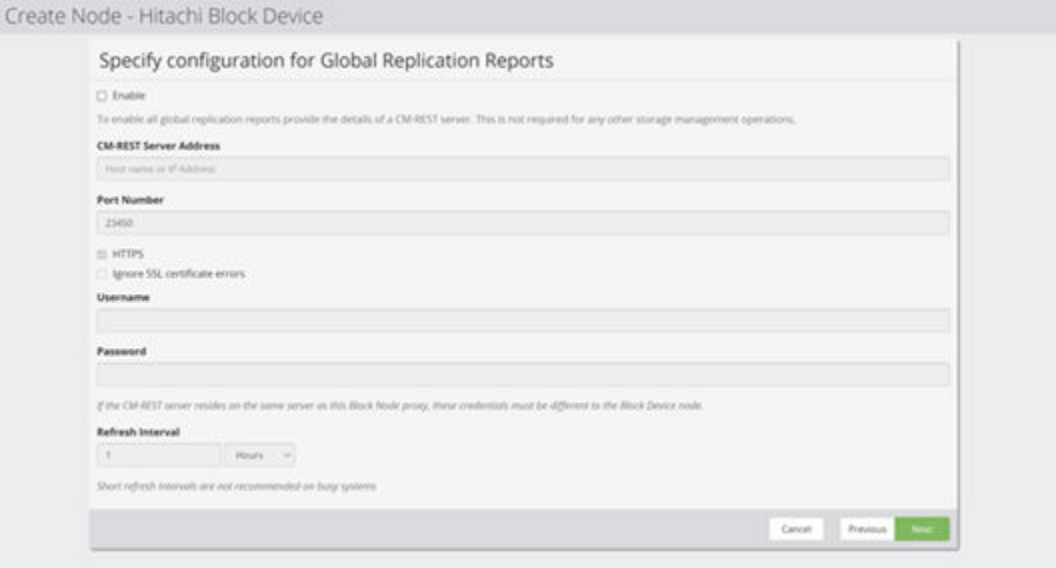


Figure 325 Hitachi Block Device Wizard - Specify Configuration for GRR

Control	Description
Enable	Allows the configuration of the Global Replication Roprts to be filled out. **Note:** GRR functionality uses certain CM-REST APIs which may require this array to be linked to an SVP. (Please see the CM-REST documentation for more information.) Full list of commands used can be found here Commands used to generate GRR (on page 968)
CM-REST Server Address	Enter either the Host Name or the IP Address of the CM-REST server.
Port Number	Enter the port number for the CM-REST server.
HTTPS	Select to enable secure communications to the CM-REST server.
Ignore SSL certificate errors	Select to ignore SSL certificate errors when communicating with the CM-REST server.

Control	Description
Username and Password	Enter the username and password for the storage array. **Note:** If CM-REST is on the same server as the Block Device Node then the credentials must be different.
Refresh Interval	Enter the required interval between refreshes of the GRR reports. **Note:** Populating the data in the GRR reports requires interaction with the storage array and therefore on busy systems it is not recommended to have short refresh times as this may affect performance of the system.

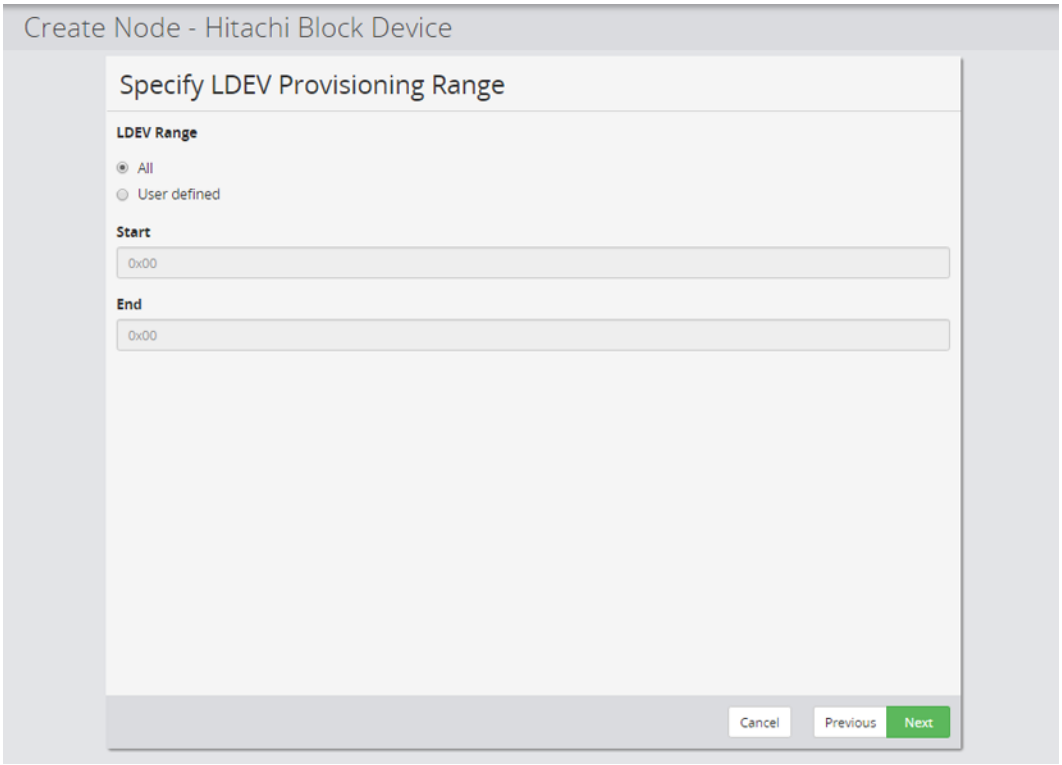


Figure 326 Hitachi Block Device Wizard - Specify LDEV Provisioning Range

Control	Description
All	Select this option if you want Protector to automatically detect the LDEV range from which snapshots and replications should be allocated.
User defined	Select this option if you want to manually specify the LDEV range from which snapshots and replications should be allocated.

Control	Description
	Note: All replication and snapshot S-VOLs must be created using free LDEV IDs that are mapped to the *meta_resource* group, and have virtual LDEV IDs matching their corresponding physical LDEV IDs. For fully provisioned snapshots and all replications, this applies to the operation that creates that snapshot or replication. For floating device snapshots and snapshots mounted using cascade mode, this applies to the mount or restore operation. For fully provisioned snapshots mounted using cascade mode, this applies both to the operation that creates that snapshot and to the mount or restore operation. If an operation tries to create one or more LDEVs, that operation will fail if there are not enough free LDEV IDs that meet the above conditions.
Start	Enabled only if User defined is selected. Enter the lower limit of the LDEV range to use for allocation.
End	Enabled only if User defined is selected. Enter the upper limit of the LDEV range to use for allocation.

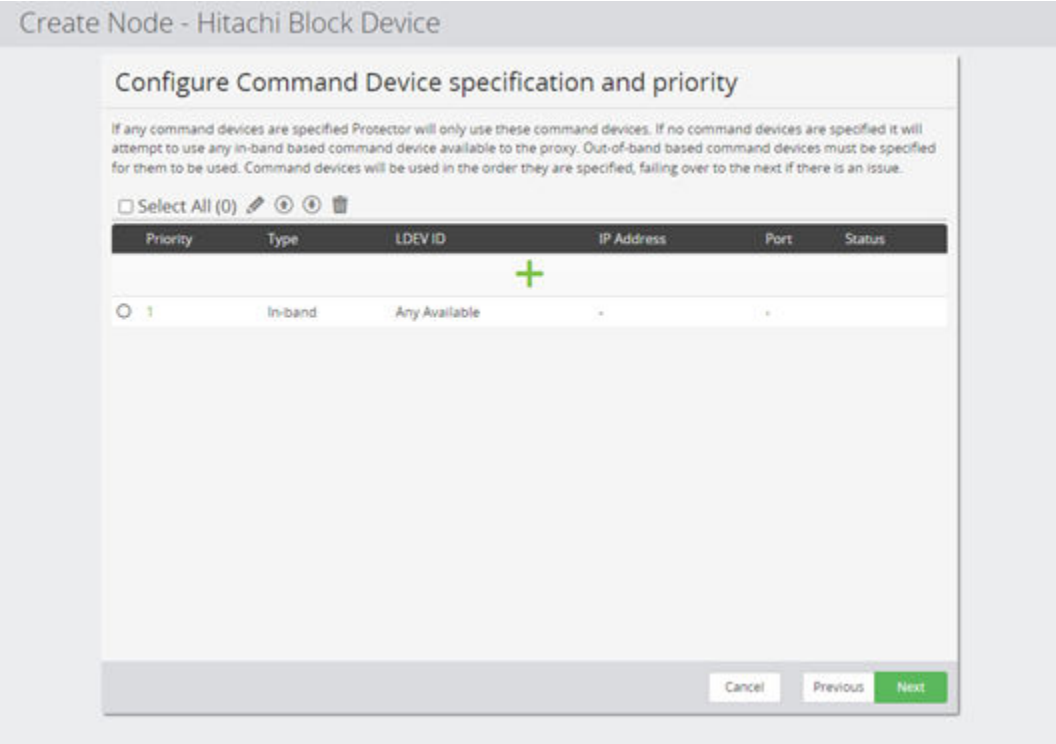


Figure 327 Hitachi Block Device Wizard - Command Device Specification and Priority

The user able to specify zero or more In-band or Out-of-band command devices in priority order.

If no command devices are specified, then Protector will attempt to control the hardware storage device via any connected command device, available to the Proxy Node specified, in an order specified by HORCM.

If one or more command devices are specified, then Protector will attempt to control the hardware storage device via a command device in the order specified by the user. If the first command device fails, Protector will progress to the next. If all specified command devices fail then the operation fails. Protector will not attempt to use any command devices that are not specified, even if they are available.

for example, it is possible to specify a specific In-band command device, followed by any In-band command device, followed by a specific Out-of-band command device.

Control	Description
Edit	Enabled when only one command device is selected. Launches the **Configure Command Device** wizard below to allow the settings to be edited.
Increase Priority	Enabled when only one command device is selected. Increases the priority of the selected command device.

Control	Description
Decrease Priority	Enabled when only one command device is selected. Decreases the priority of the selected command device.
Delete	Enabled when one or more command devices are selected. Deletes selected command device.
Add	Launches the **Configure Command Device** wizard below to guide you through setting up a In-band or Out-of-band command device.

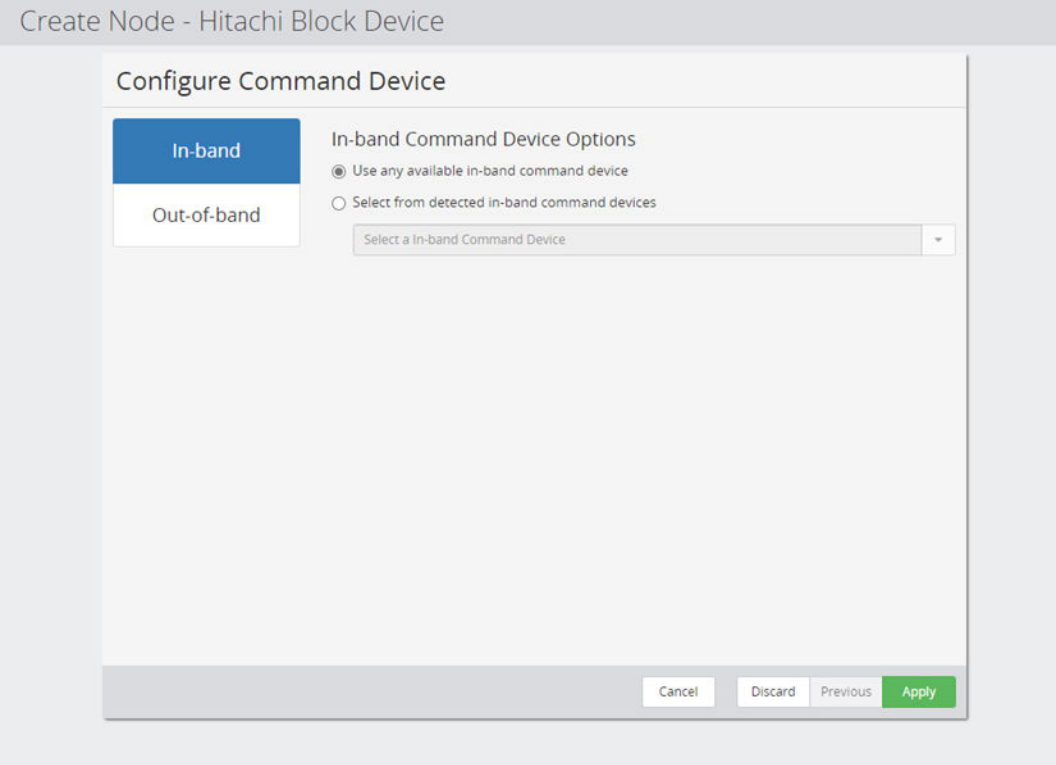


Figure 328 Hitachi Block Device Wizard - Configure Command Device - In-band

Control	Description
Use any available In-band command device	Select this option to insert an entry in the command device list that allows Protector to use any available In-band command device such as fibre or iSCSi.

Control	Description
Select from detected In-band command devices	Select this option to insert a specific In-band command device in the list. The detected In-band command devices are displayed in the dropdown menu below using their decimal LDEV ID.

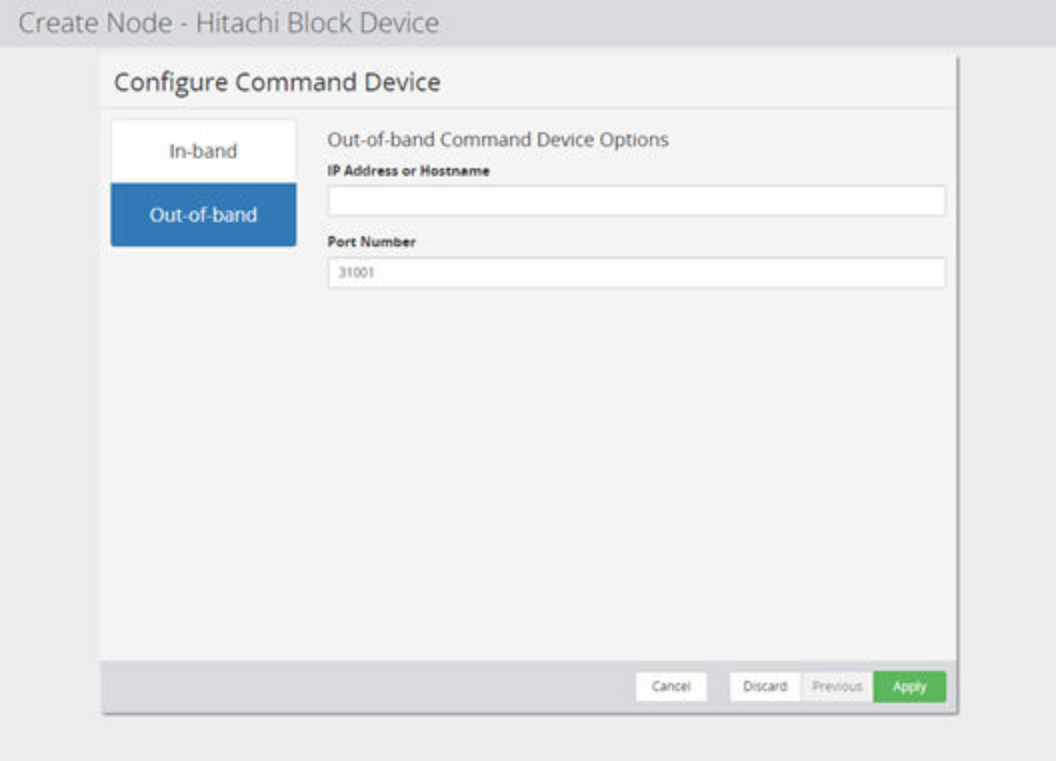


Figure 329 Hitachi Block Device Wizard - Configure Command Device - Out-of-band

Note: When configuring IP command devices for VSP G1x00 storage devices, we recommend adding one for CTL1 and one for CTL2, to maintain dual redundancy.

Control	Description
IP Address or Hostname	Enter the IP address or hostname of the command device to add to the list. **Note:** For HUS VM storage devices, use the IP address of the SVP. For VSP G1x00 storage devices, use the IP address of CTL1 or CTL2. DO NOT use the IP address of the SVP.

Control	Description
Port Number	Enter the port number of the command device.

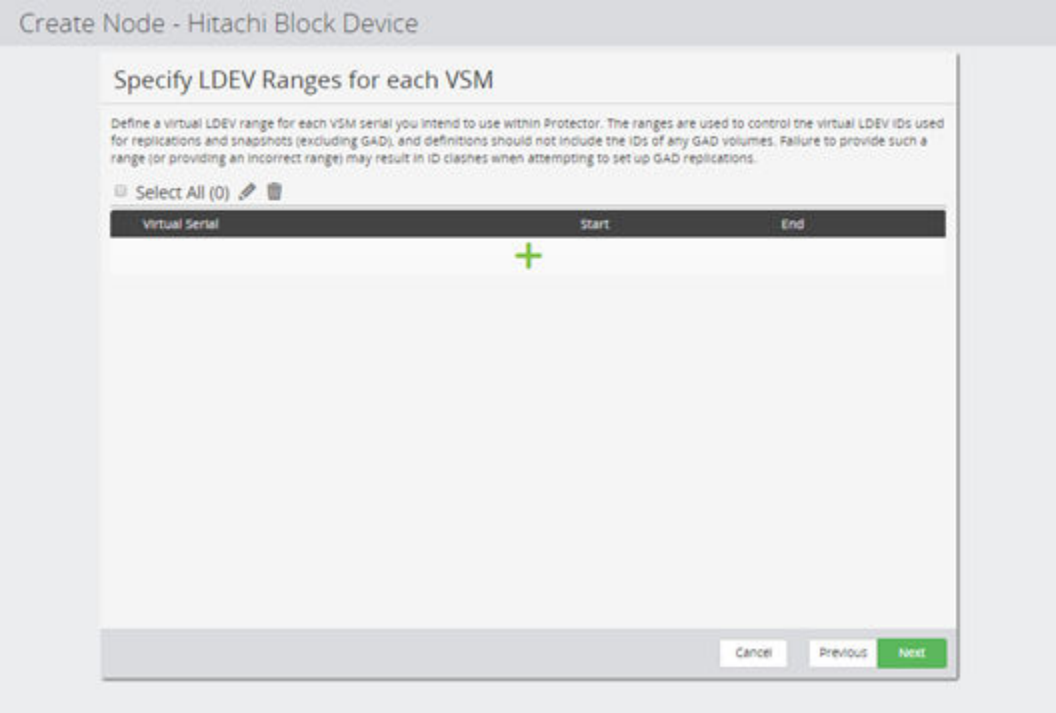


Figure 330 Hitachi Block Wizard - Specify LDEV Ranges for each VSM

Note: GAD replications require P-VOLs and S-VOLs to have matching virtual serial numbers and virtual LDEV IDs. To avoid virtual LDEV ID collisions between GAD volumes and non-GAD S-VOLs (created by Protector for other types of replications and snapshots), it is possible to define virtual LDEV ID ranges to be used by those non-GAD operations. Virtual LDEV ranges can be specified for each VSM (Virtual Storage Machine) to be used.

Control	Description
Edit	Enabled only if one Virtual LDEV range is selected. The **Configure Virtual LDEV Range** page of the wizard (see below) is displayed to enable a port to be added.
Delete	Enabled only if one or more Virtual LDEV ranges are selected. Deletes the Virtual LDEV range(s) from the list.
Add	Adds a new Virtual LDEV range to the list. The **Configure Virtual LDEV Range** page of the wizard (see below) is displayed to enable a Virtual LDEV Range to be added.

Control	Description
Virtual Serial(s)	The Virtual Serial(s) and associated LDEV Ranges that will be used.

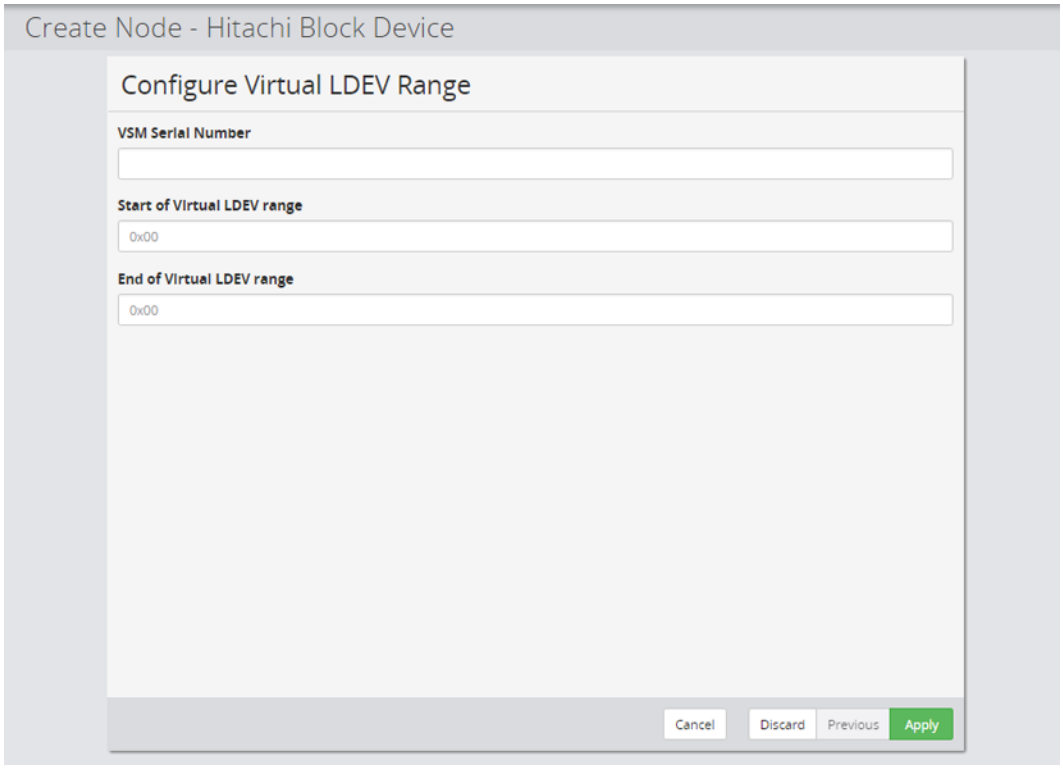


Figure 331 Hitachi Block Wizard - Configure Virtual LDEV Range

Caution: These ranges are used to control the virtual LDEV IDs used for non-GAD replications and snapshots. They must be defined to exclude the IDs of any GAD volumes. Failure to provide such a range (or providing an incorrect range) may result in ID clashes when attempting to set up GAD replications.

Control	Description
VSM Serial Number	Enter the serial number of the VSM you intend to use within Protector.
Start of LDEV range	Enter the lower limit of the LDEV range to use for allocation.
End of LDEV range	Enter the lower limit of the LDEV range to use for allocation.

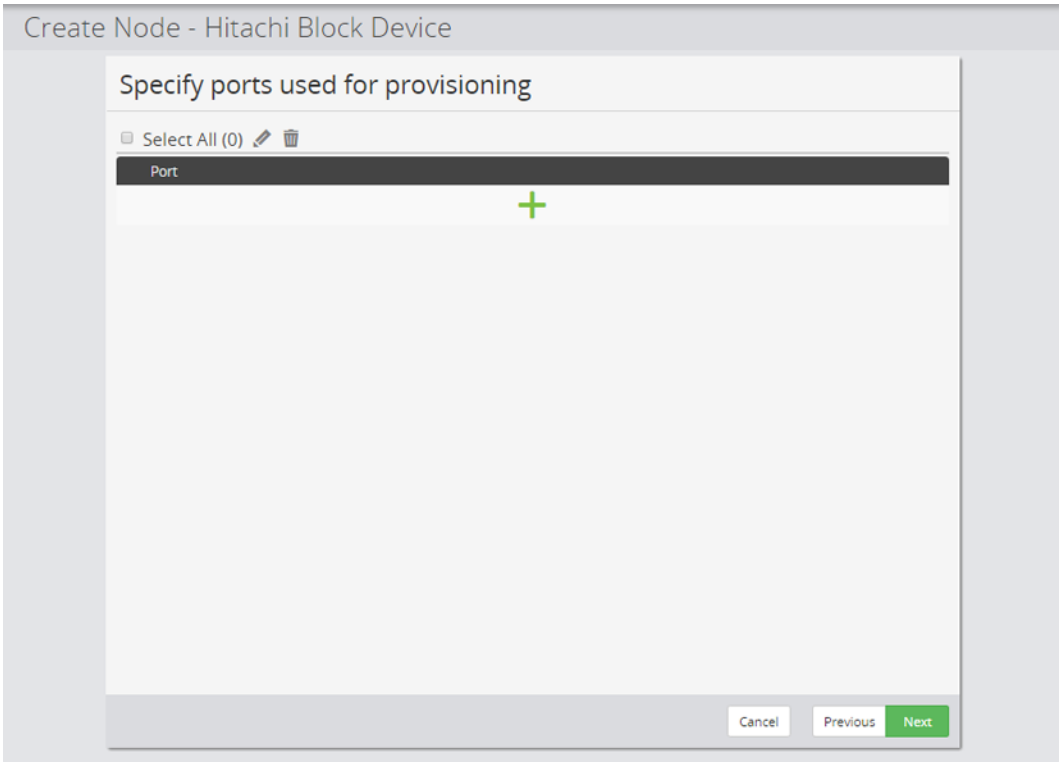


Figure 332 Hitachi Block Device Wizard - Specify ports used for provisioning

Control	Description
Edit	Enabled only if one provisioning port is selected. The **Specify Port** page of the wizard (see below) is displayed to enable a port to be added.
Delete	Enabled only if one or more provisioning ports are selected. Deletes the port(s) from the list.
Add	Adds a new provisioning port to the list. The **Specify Port** page of the wizard (see below) is displayed to enable a port to be added.
Provisioning Port(s)	Lists the ports that will be used for provisioning.

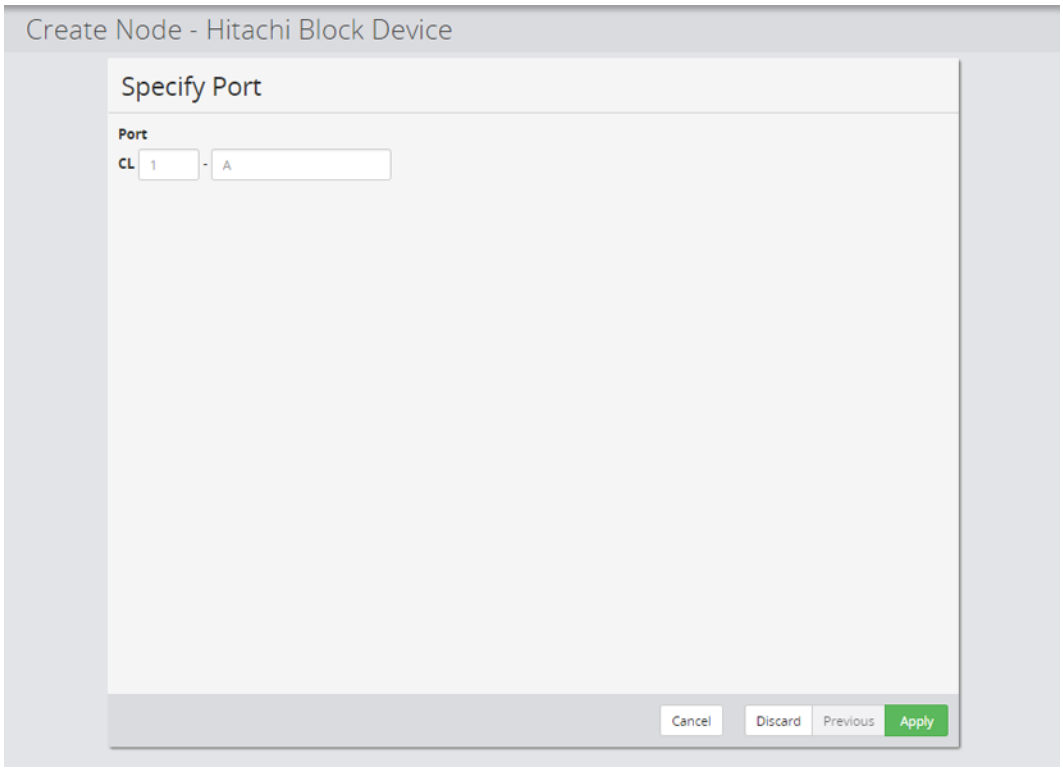


Figure 333 Hitachi Block Device Wizard - Specify Port

Note: If more than one provisioning port is selected, then the port with the least amount of LUNs will be used.

Control	Description
Port	Enter the port identifier in the following format: CL*c*-*s* where: ▪ *c* is the physical channel number in the range 1…n ▪ *s* is the physical slot number in the range A…Z

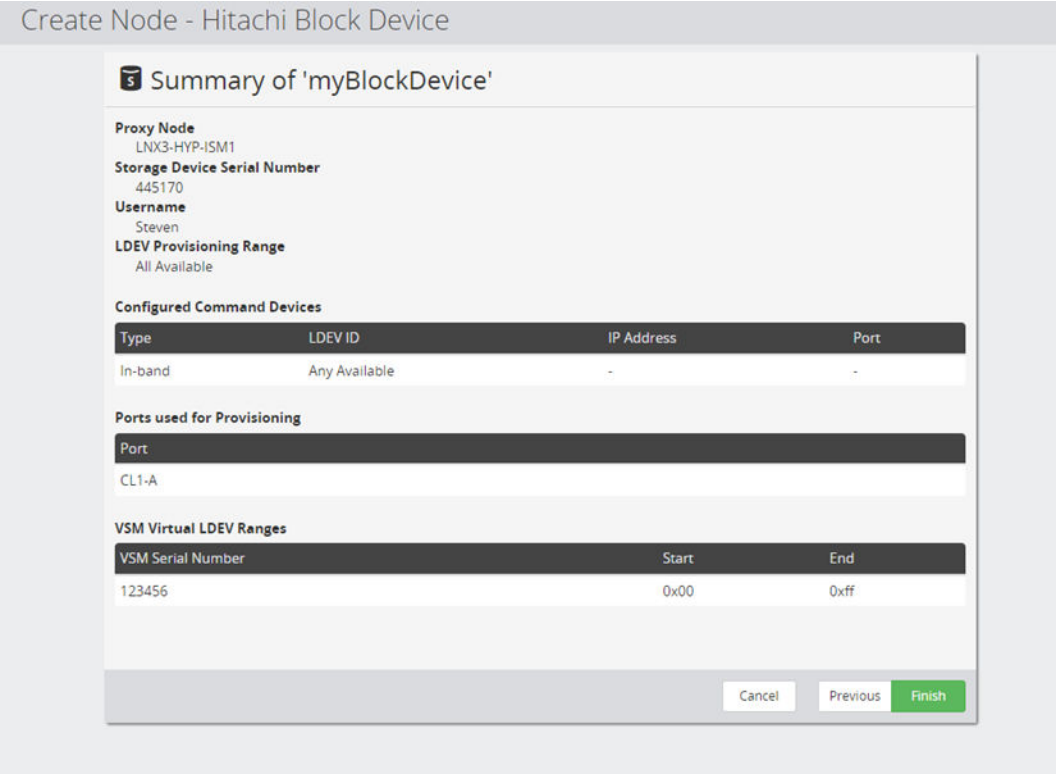


Figure 334 Hitachi Block Wizard - Summary

Control	Description
Summary	Summary of the settings entered.

Hitachi Content Platform Storage Node Wizard

This wizard is launched when a new Hitachi Content Platform Node is added to the Nodes Inventory.

Note: Ensure that the Proxy Node defined within the wizard and the Hitachi Content Platform node are time synced. If the time is different by over 5 minutes, then the node creation and future backup may fail. It is recommend that the system time on all nodes should have the correct date/time associated with their particular region and time zone.

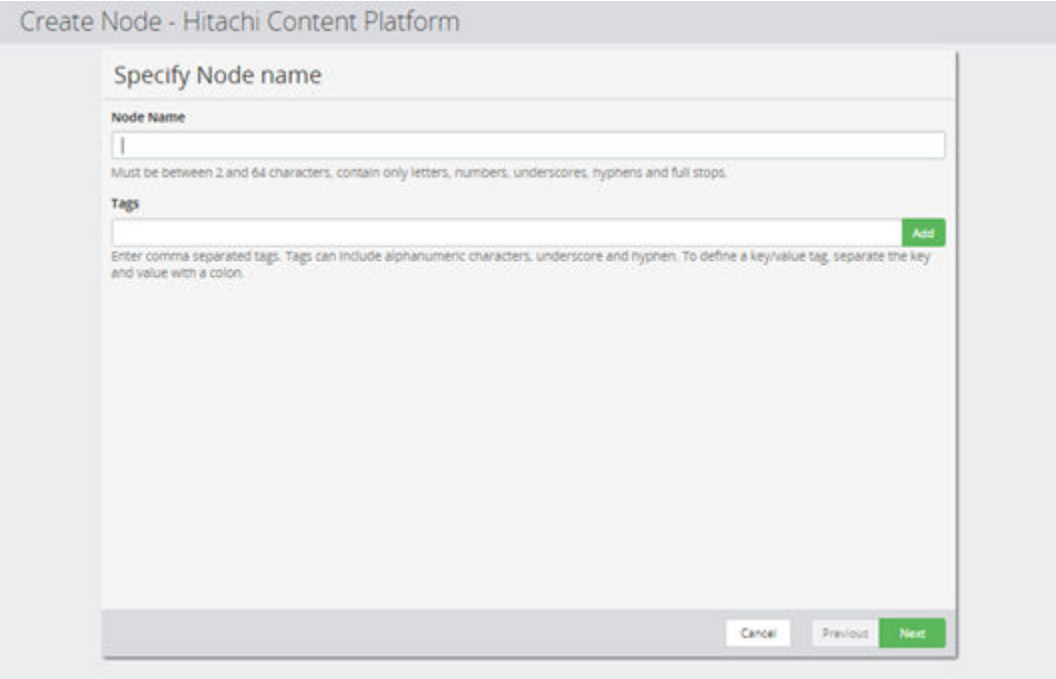


Figure 335 Hitachi Content Platform Wizard - Specify Node Name

Control	Description
Node Name	Enter a name for the HCP storage node.
Tags	Add the tags to be associated with the object being created.

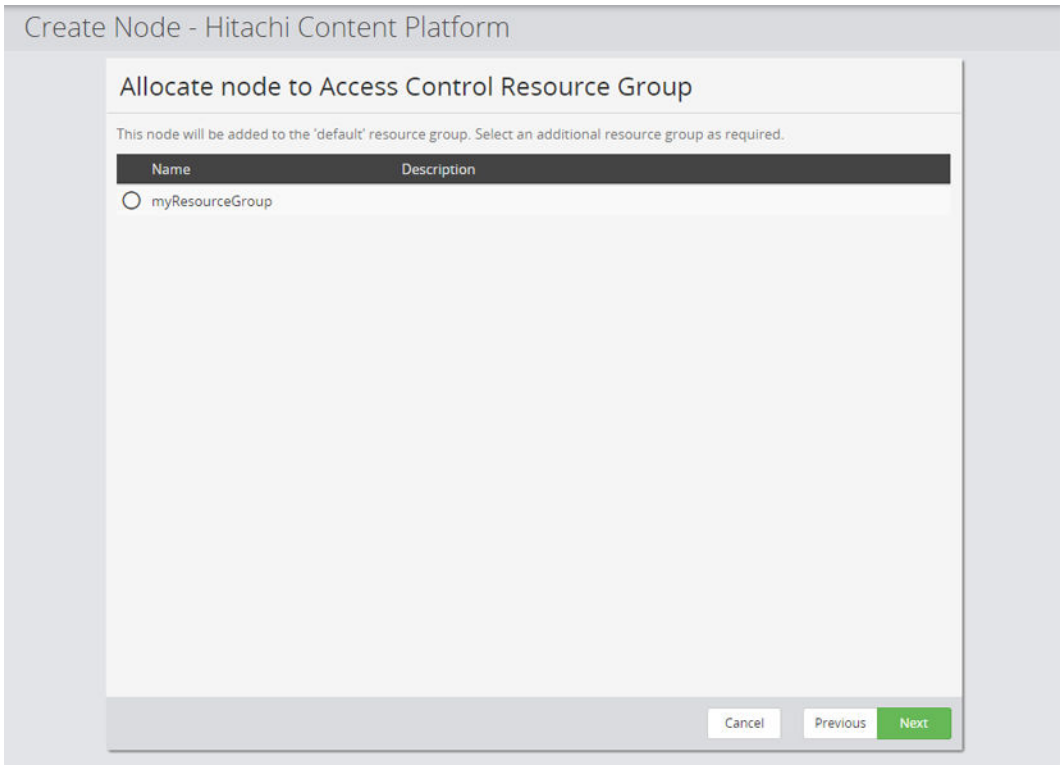


Figure 336 Hitachi Content Platform Wizard - Allocate node to Access Control Resource Group

Control	Description
Resource Groups	Select the resource group(s) to which this node will be allocated for the purposes of RBAC. All nodes are automatically allocated to the 'default' resource group.

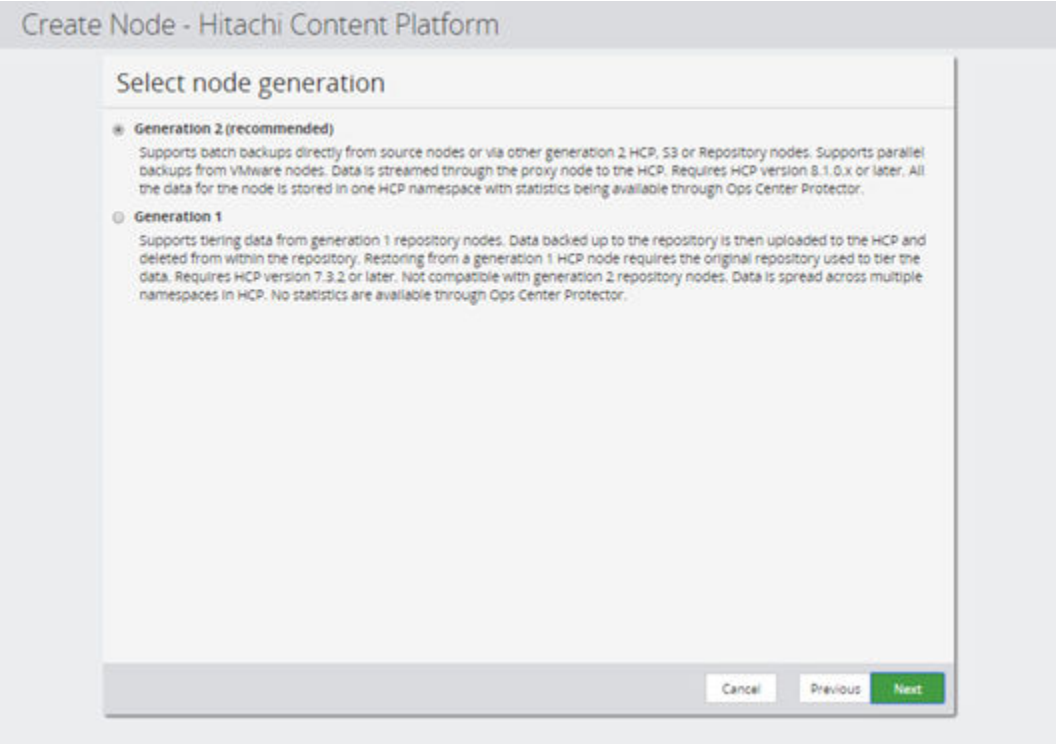


Figure 337 Hitachi Content Platform Wizard - Select Node Generation

Control	Description
Generation 2	Select this for a generation 2 storage node which is compliant with the new Universal Backup Infrastructure interface. This is the recommended option.
Generation 1	Select this for a generation 2 storage node which required data to be sent via a generation 1 repository. Clicking Next will take you to the Configure Tenant Page.

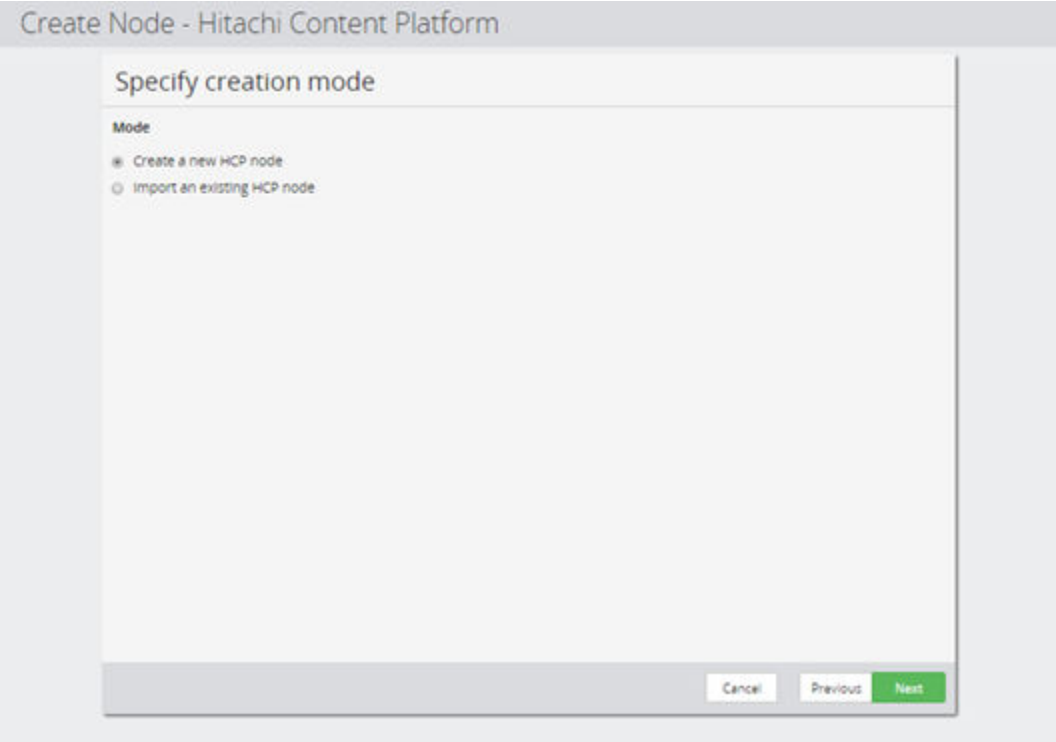


Figure 338 Hitachi Content Platform Wizard - Specify Node creation mode

Control	Description
Create a new HCP node	Select this to setup a new HCP node.
Import an existing HCP node	Select this to import an existing HCP node.

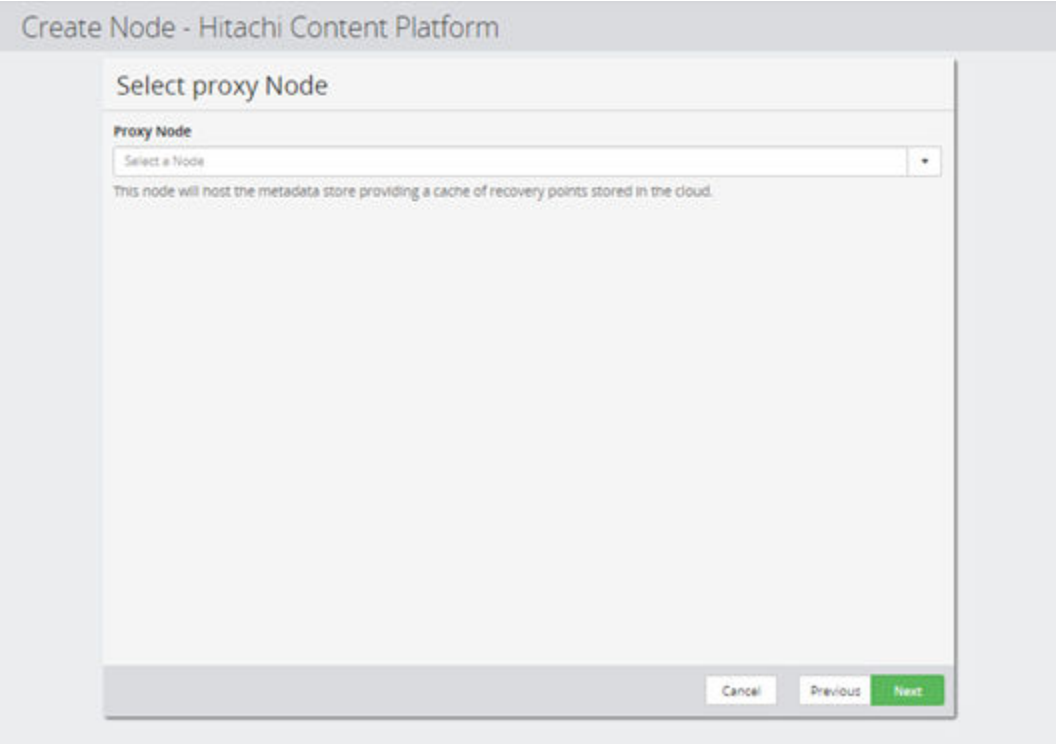


Figure 339 Hitachi Content Platform Wizard - Select Proxy Node

Control	Description
Proxy Node	Select a Protector node to act as a proxy.

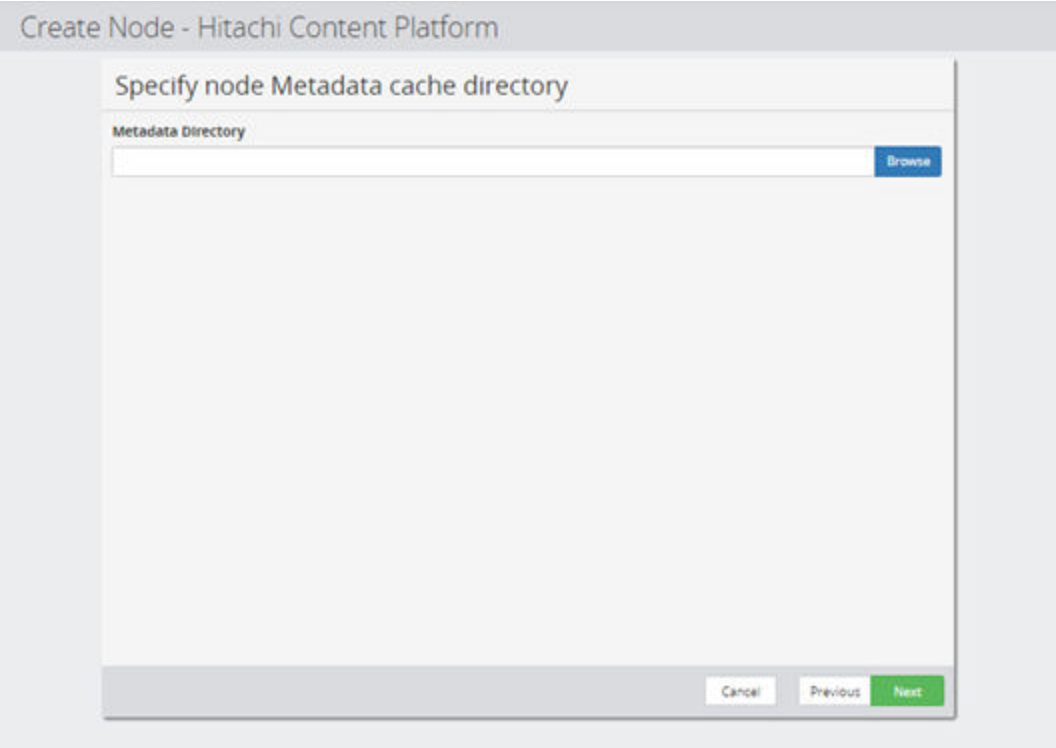


Figure 340 Hitachi Content Platform Wizard - Specify node Metadata cache directory

Control	Description
Metadata Directory	Enter a directory on the Proxy node where Protector can place metadata files related to HCP data. Click the Browse button to open the Path Dialog (on page 350) if required.

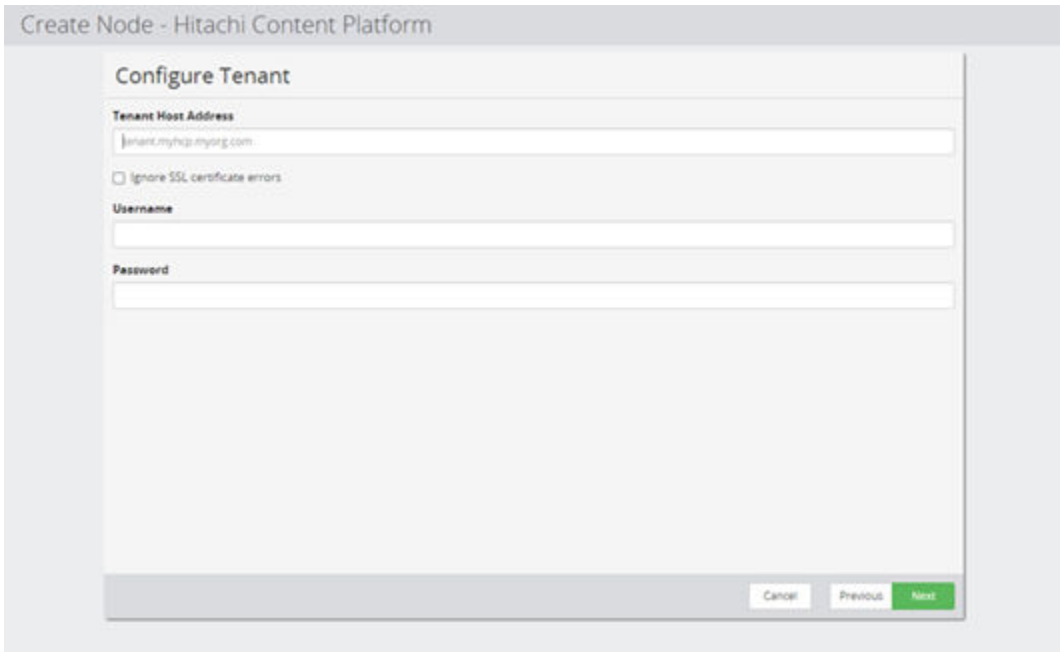


Figure 341 Hitachi Content Platform Wizard - Configure HCP Tenant

Control	Description
Tenant Host Address	Enter the host name for the tenant in the form `<tenant>.<hcp_name>.<domain>`.
Ignore SSL certificate errors	Allows the system to be used without properly configured SSL certificates.
Username	Enter a user name for the tenant specified above.
Password	Enter the password for the user.

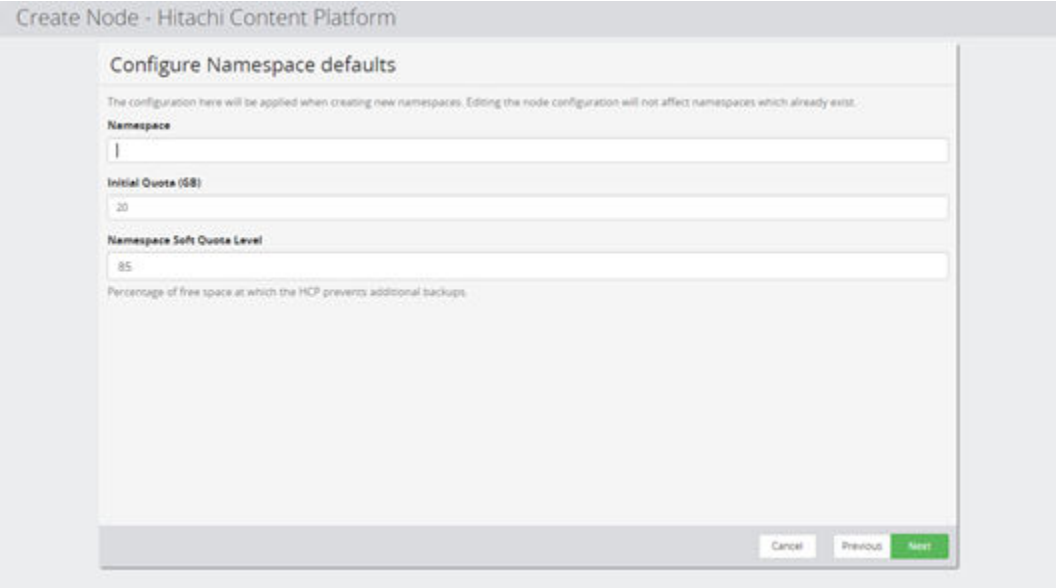


Figure 342 Hitachi Content Platform Wizard - Configure HCP Namespace defaults

Control	Description
Namespace	This is automatically generated for a generation 1 HCP node and is therefore greyed out. For a generation 2 HCP node enter the desired namespace
Initial Quota (GB)	Change the namespace quota from the default if required. It is important to estimate the quota required. Too small a value will result in tiering being stopped if the quota is exhausted. Too large a value will waste total tenant quota. This value is not displayed when importing an existing generation 2 HCP node. **Note:** Changing the quota for an existing HCP node will only affect newly created namespaces. Existing namespaces created by Protector on that node will not be affected.
Namespace Soft Quota Level	This is the percentage of free space at which additional backups are prevented

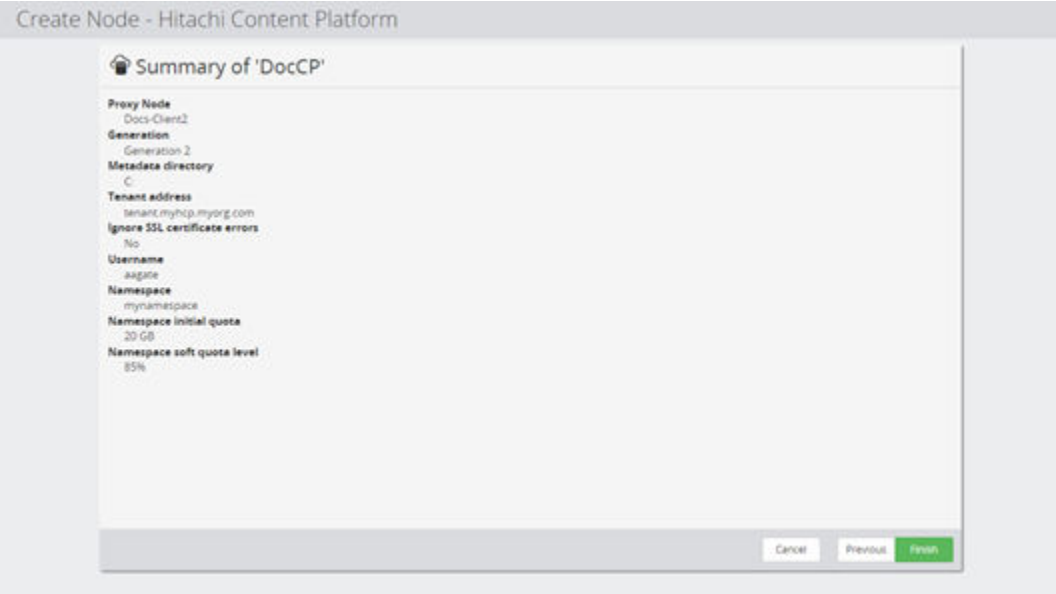


Figure 343 Hitachi Content Platform Wizard - Summary

Control	Description
Summary	Summary of the settings entered.

Hitachi Logical Block Device Node Wizard

Hitachi Logical Block Storage Nodes implement multi tenancy by limiting what resources each tenant has access to on a physical block storage device shared with other tenants.

The Hitachi Logical Block Storage Node Wizard is used by storage administrators to create logical views of a physical storage device, so that multiple tenants can use a common device without interfering with each other's resources. About multi-tenancy for Hitachi block storage (on page 50) describes how Protector implements this arrangement.

Note: A Physical Block node must be created using the Hitachi Block Device Node Wizard (on page 583) before any Logical Block nodes can be created on it.

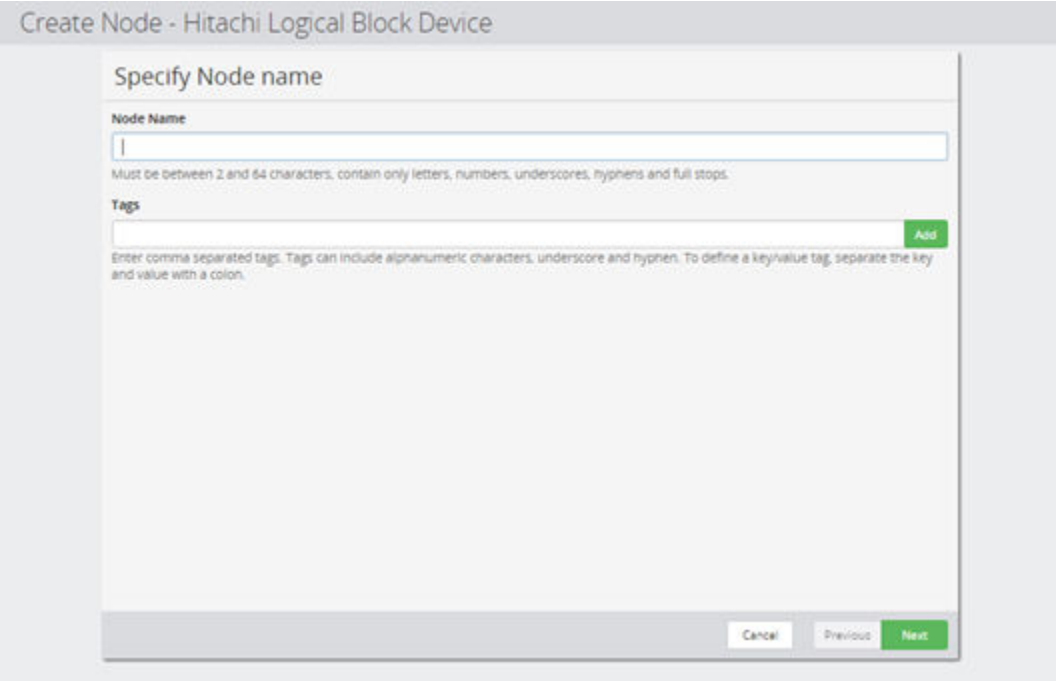


Figure 344 Hitachi Logical Block Wizard - Specify Node name

Control	Description
Node Name	Enter a name for the Logical Block node.
Tags	Add the tags to be associated with the object being created.
Add to access control resource group	Check this box if you want to add it to a different resource group.
Resource Group	Type the name of the resource group or select one from the drop down list.

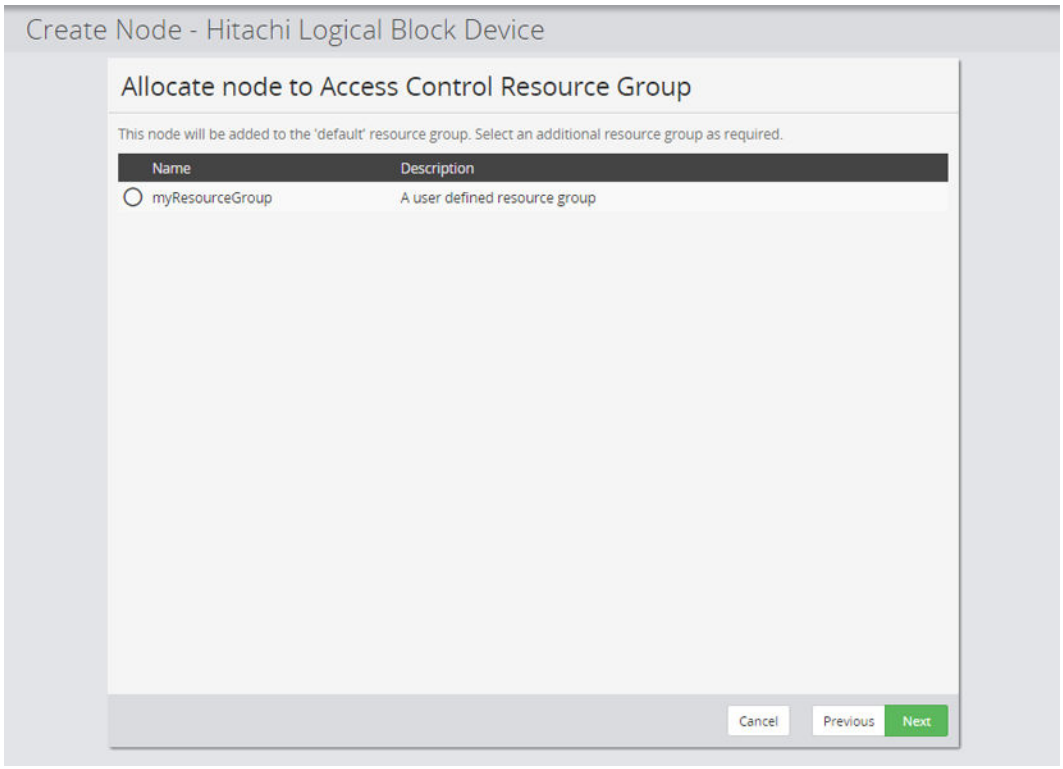


Figure 345 Hitachi Logical Block Wizard - Allocate node to Access Control Resource Group

Control	Description
Resource Groups	All nodes are automatically allocated to the 'default' RBAC resource group. Select the additional resource group(s), from those listed, to which this node will be allocated for the purposes of RBAC. Allocation to resource groups can also be done via the Access Control Resource Groups Inventory (on page 344).

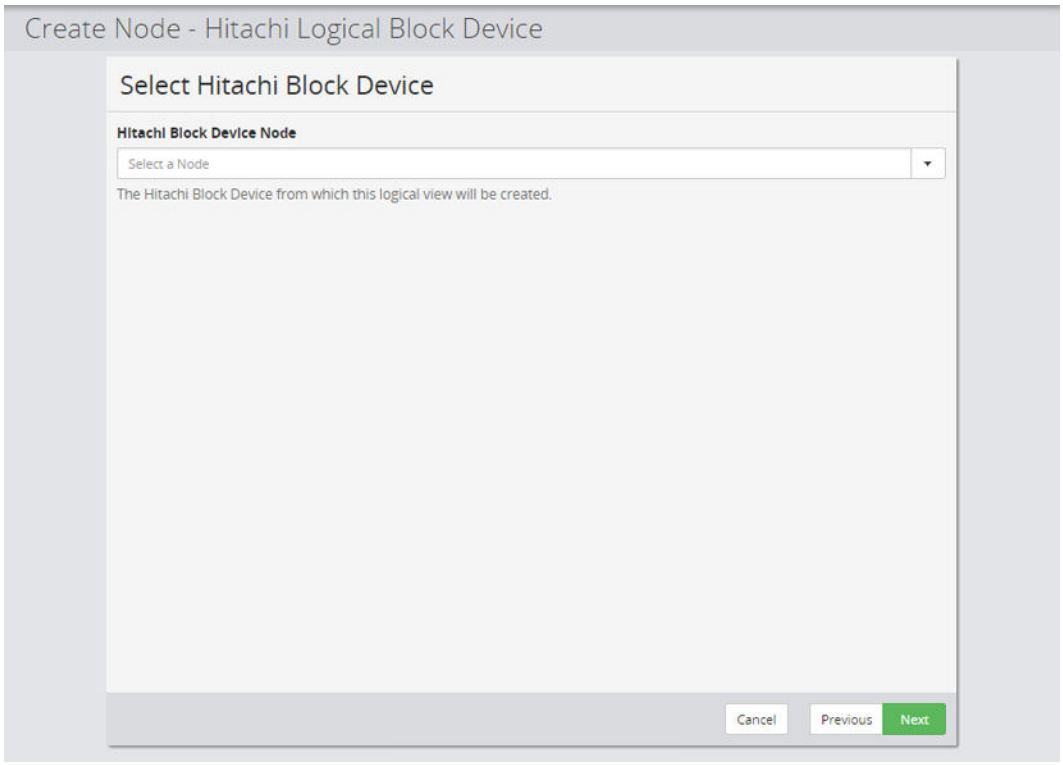


Figure 346 Hitachi Logical Block Wizard - Select Hitachi Block Device

Control	Description
Hitachi Block Device Node	Enter or select the name of the Physical Block node on which the Logical Block node will be based .

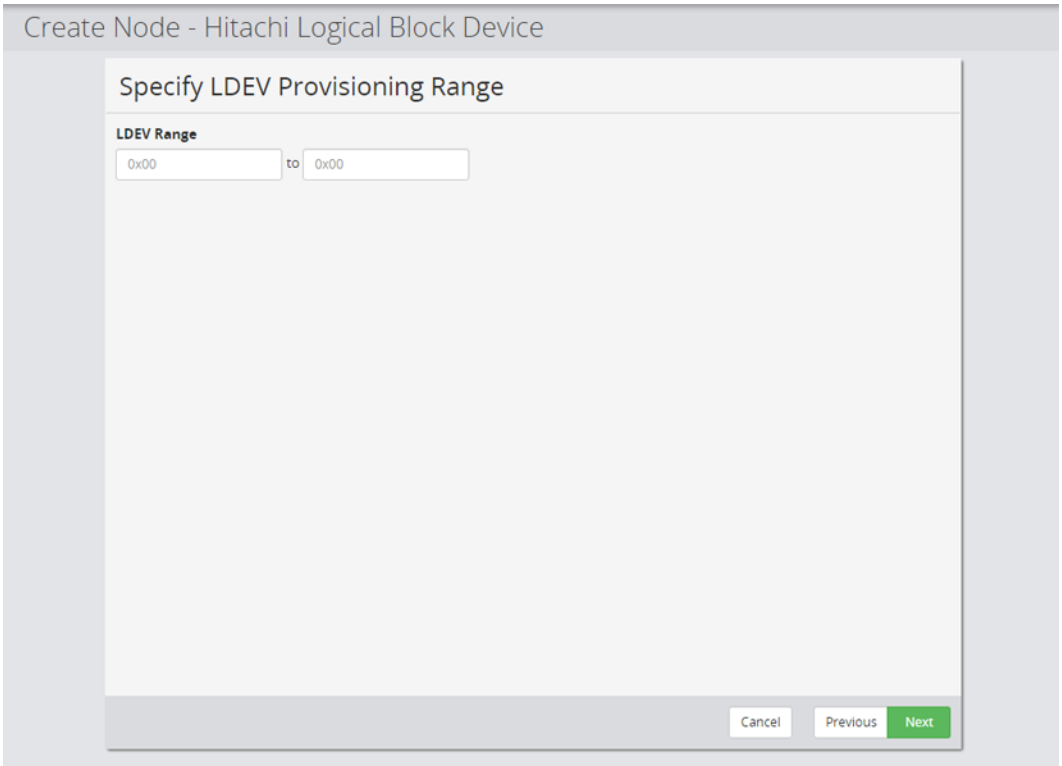


Figure 347 Hitachi Logical Block Wizard - Specify LDEV Provisioning Range

Control	Description
LDEV Range	Specify the LDEV range on the physical storage device that will be made available by this logical storage device for provisioning S-VOLs. **Note:** All replication and snapshot S-VOLs must be created using free LDEV IDs that are mapped to the *meta_resource* group, and have virtual LDEV IDs matching their corresponding physical LDEV IDs. For fully provisioned snapshots and all replications, this applies to the operation that creates that snapshot or replication. For floating device snapshots and snapshots mounted using cascade mode, this applies to the mount or restore operation. For fully provisioned snapshots mounted using cascade mode, this applies both to the operation that creates that snapshot and to the mount or restore operation. If an operation tries to create one or more LDEVs, that operation will fail if there are not enough free LDEV IDs that meet the above conditions.

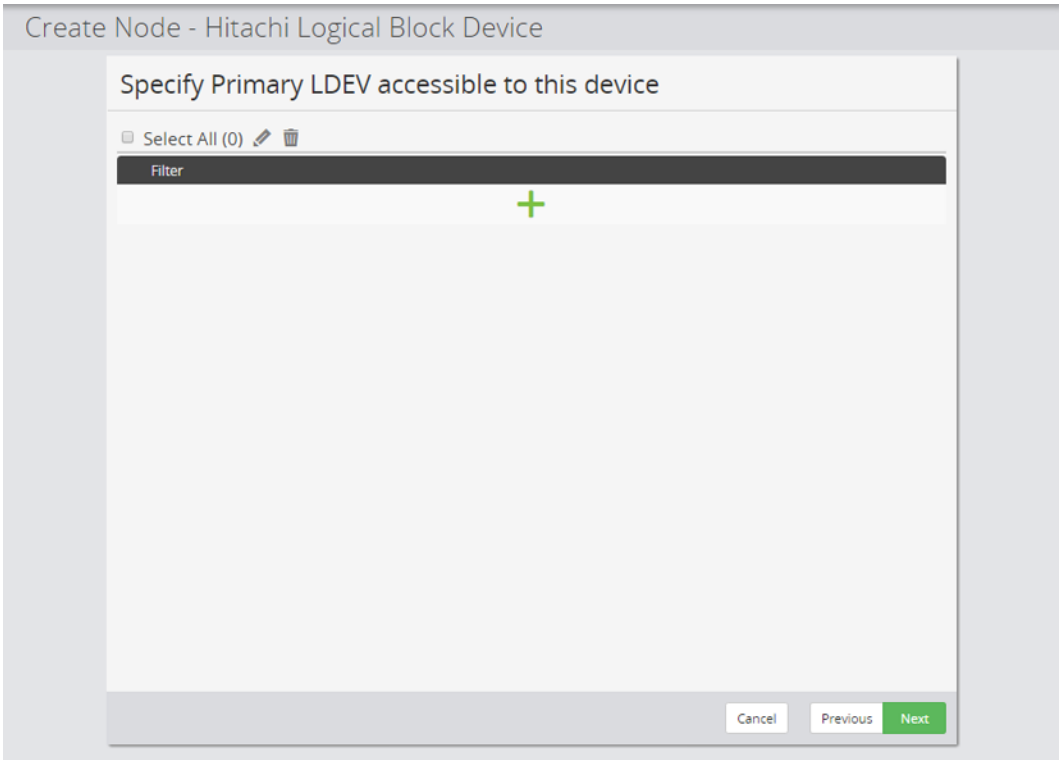


Figure 348 Hitachi Logical Block Wizard - Specify Primary LDEV accessible to this device

Specifies the P-VOL(s) on the physical block device that will be available to this logical block device for snapshotting or replication. LDEVs can be specified by name, range and host group.

Note: This information is optional. If nothing is specified here, then no P-VOLs will be available for snapshotting or replicating on this logical block device.

Control	Description
Edit	Enabled only if one Primary LDEV is selected. The **Specify Primary LDEVs filter** page of the wizard (see below) is displayed to enable an LDEV Filter to be added.
Delete	Enabled only if one or more LDEV Filters are selected. Deletes the LDEV Filter(s) from the list.
Add	Adds a new LDEV Filter to the list. The **Specify Primary LDEVs filter** page of the wizard (see below) is displayed to enable an LDEV Filter to be added.
Existing LDEV Filter(s)	Lists the LDEV Filter(s) that will be used.

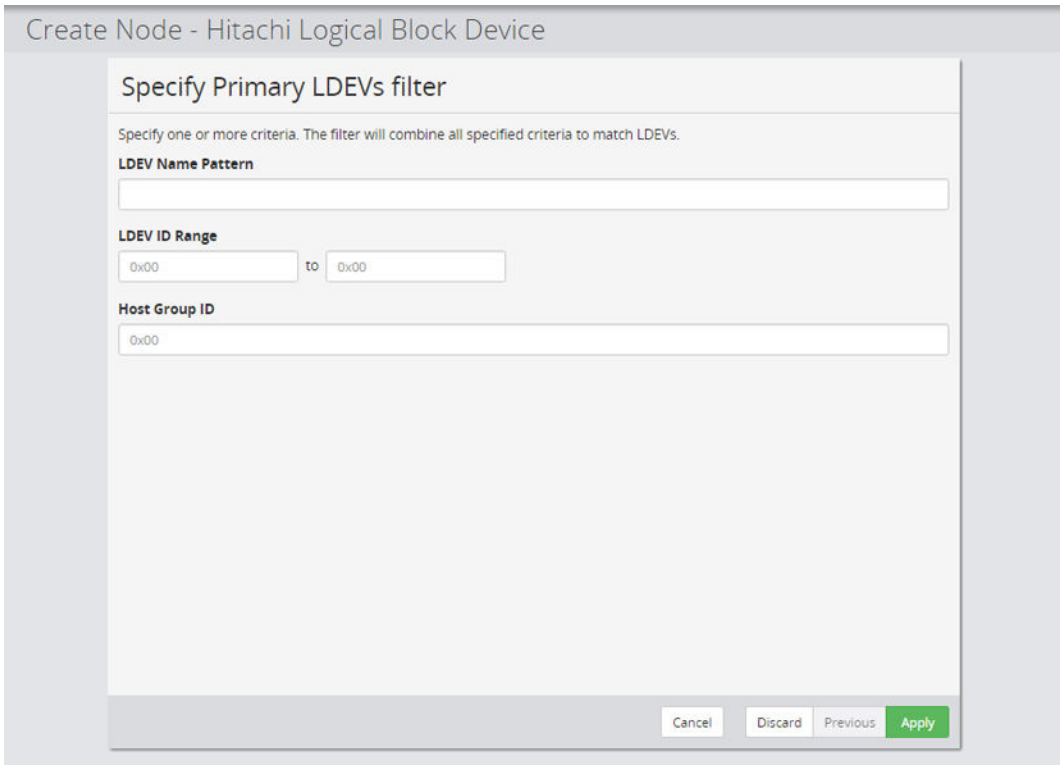


Figure 349 Hitachi Logical Block Wizard - Specify Primary LDEVs filter

Specify one or more of the following parameters to identify the Primary LDEV(s):

Note: The terms entered here are logically ANDed together, so an LDEV will only be selected if it matches the *Name Pattern* AND *ID Range* AND *Host Group ID*.

Control	Description
LDEV Name Pattern	Enter a regular expression (see Basic regular expressions (on page 624)) to identify matching Primary LDEV Names that will be made available on this logical block device.
LDEV ID Range	Enter the range of the Primary LDEV IDs that will be made available on this logical block device.
Host Group ID	Enter a Host Group identifier. **Tip:** Host Group ID is normally used alone rather than in combination with the other filter terms presented here.

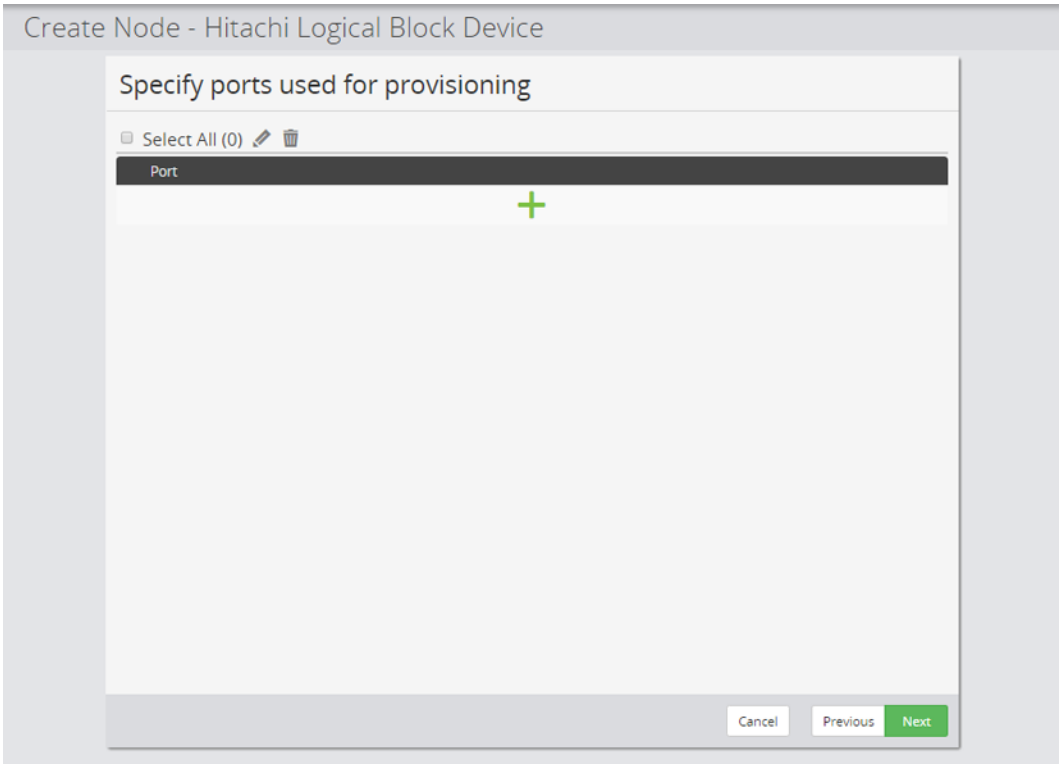


Figure 350 Hitachi Logical Block Wizard - Specify ports used for provisioning

Note: This information is optional. Protector creates dummy Host Groups for its own processing needs. If nothing is specified here, then Protector will use any available port when creating these dummy Host Groups.

Tip: Use this list to restrict Protector to using LUN allocations on designated ports only.

Control	Description
Edit	Enabled only if one provisioning port is selected. The **Specify Port** page of the wizard (see below) is displayed to enable a port to be edited.
Delete	Enabled only if one or more provisioning ports are selected. Deletes the port(s) from the list.
Add	Adds a new provisioning port for use by Protector to the list. The **Specify Port** page of the wizard (see below) is displayed to enable a port to be added.
Provisioning Port(s)	Lists the provisioning ports that will be used.

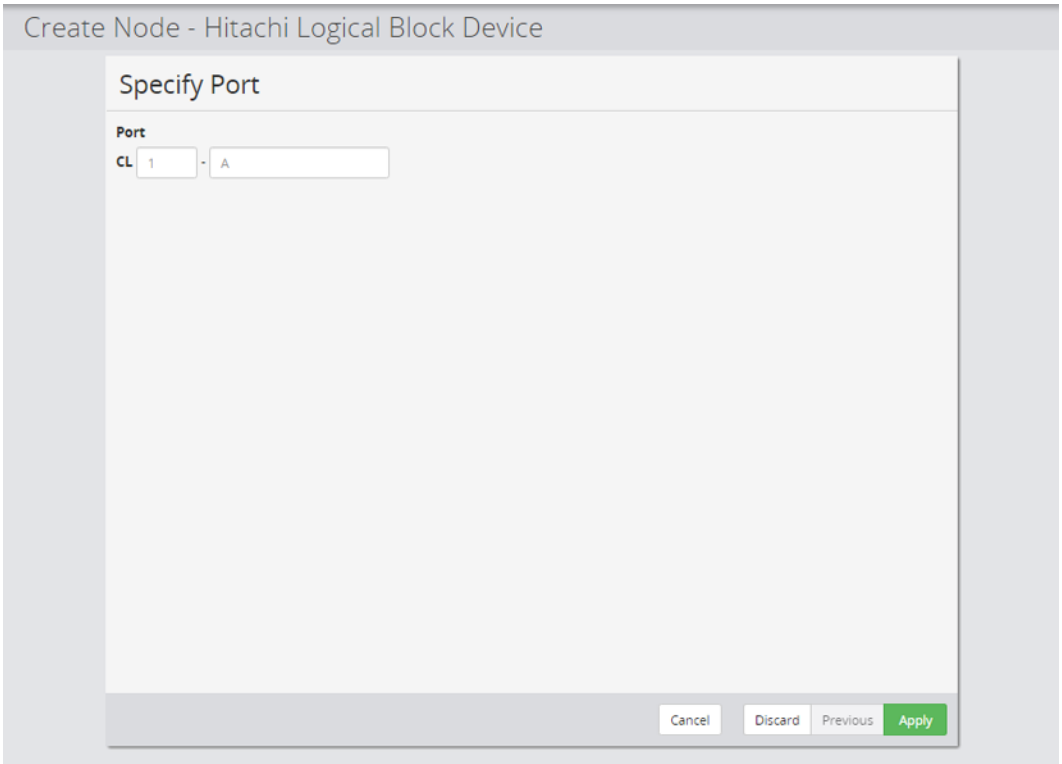


Figure 351 Hitachi Logical Block Wizard - Specify Port

Control	Description
Port	Specify a port that can be used by Protector for its own processing needs. Enter a port identifier in the following format: CL*c*-*s* where: ▪ *c* is the physical channel number in the range 1…n ▪ *s* is the physical slot number in the range A…Z



Figure 352 Hitachi Logical Block Wizard - Specify Journals accessible to this device

Note: This information is optional. If nothing is specified here, then no source or destination Journals will be available for Universal Replicator on this logical block device.

Control	Description
Edit	Enabled only if one Journal Range is selected. The **Specify Journal range** page of the wizard (see below) is displayed to enable a journal range to be added.
Delete	Enabled only if one or more Journal Ranges are selected. Deletes the Journal Range(s) from the list.
Add	Adds a new Journal Range to the list. The **Specify Journal range** page of the wizard (see below) is displayed to enable a Journal Range to be added.
Journal Ranges	The Journal Range to be used.

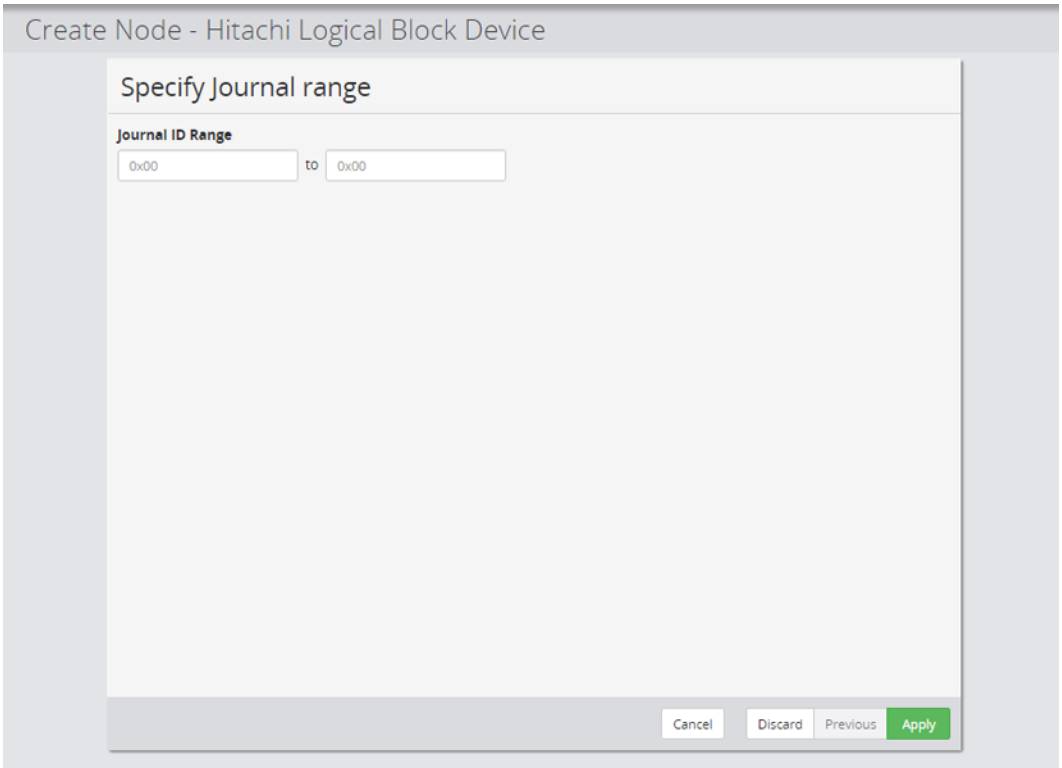


Figure 353 Hitachi Logical Block Wizard - Specify Journal range

Control	Description
Journal ID Range	Enter the Journal ID range that will be made available on this logical block device.

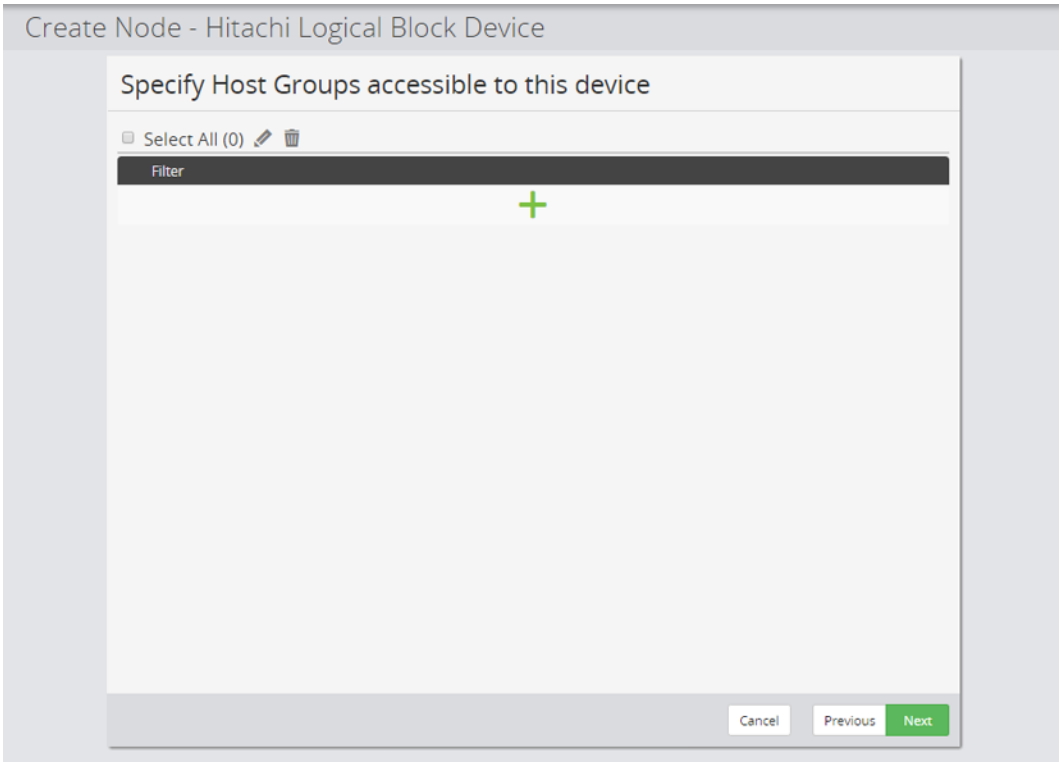


Figure 354 Hitachi Logical Block Wizard - Specify Host Groups accessible to this device

Note: This information is optional. If nothing is specified here, then no Host Groups will be available on this logical block device for:

- Mounting (see Hitachi Block Snapshot or Replication Mount Wizard (on page 772) and Hitachi Block Mount Configuration Wizard (on page 485)).
- Specifying GAD cross-path replication secondary Host Groups (see Hitachi Block Replication Configuration Wizard (on page 398)).

Control	Description
Edit	Enabled only if one Host Group is selected. The **Specify Host Group Filter** page of the wizard (see below) is displayed to enable a Host Group to be added.
Delete	Enabled only if one or more Host Group Filters are selected. Deletes the Host Group Filter(s) from the list.
Add	Adds a new Host Group Filter to the list. The **Specify Host Group Filter** page of the wizard (see below) is displayed to enable a Host Group Filter to be added.
Host Group Filter(s)	The Host Group Filter to be used.

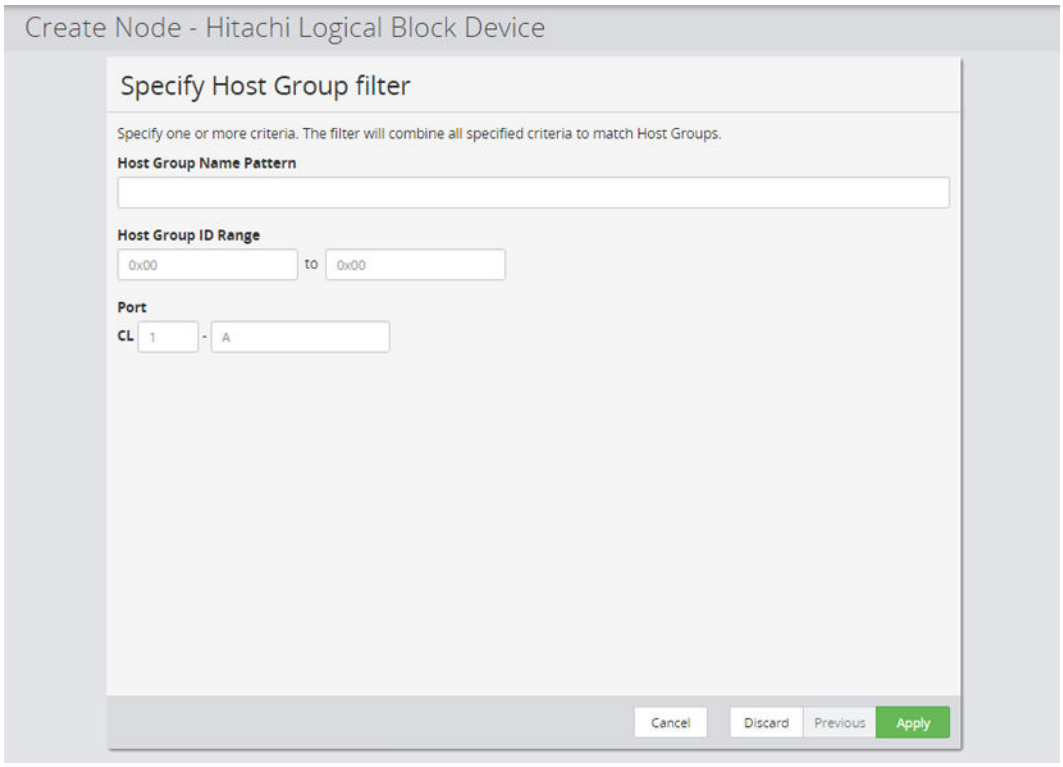


Figure 355 Hitachi Logical Block Wizard - Specify Host Group filter

Specify one or more of the following parameters to identify the Host Groups:

Note: The terms entered here are logically ANDed together, so a Host Group will only be selected if it matches the *Name Pattern* AND *ID Range* AND *Port ID*.

Control	Description
Host Group Name Pattern	Enter a regular expression (see Basic regular expressions (on page 624)) to identify matching Host Group Names.
Host Group ID range	Enter the lower and upper limit of the Host Group ID range to use.
Port ID	Enter the port identifier in the following format: CL*C*-*S* where: ▪ *C* is the physical channel number in the range 1…n ▪ *S* is the physical slot number in the range A…Z

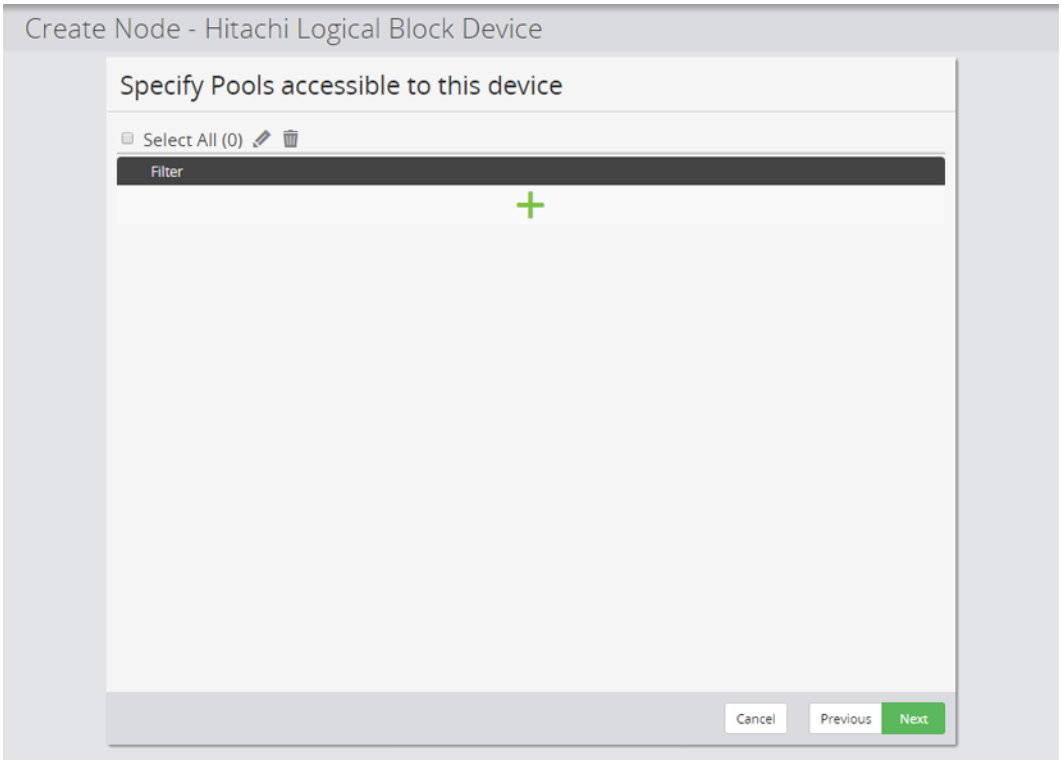


Figure 356 Hitachi Logical Block Wizard - Specify Pools accessible to this device

Note: This information is optional. If nothing is specified here, then no Pools will be available on this logical block device for creating snapshots or replications.

Control	Description
Edit	Enabled only if one Pool Filter is selected. The **Specify Pool filter** page of the wizard (see below) is displayed to enable a Pool Filter to be added.
Delete	Enabled only if one or more Pool Filters are selected. Deletes the Pool Filter(s) from the list.
Add	Adds a new Pool Filter to the list. The **Specify Pool filter** page of the wizard (see below) is displayed to enable a Pool Filter to be added.
Existing Pool Filter	The Pool Filter(s) to be used.

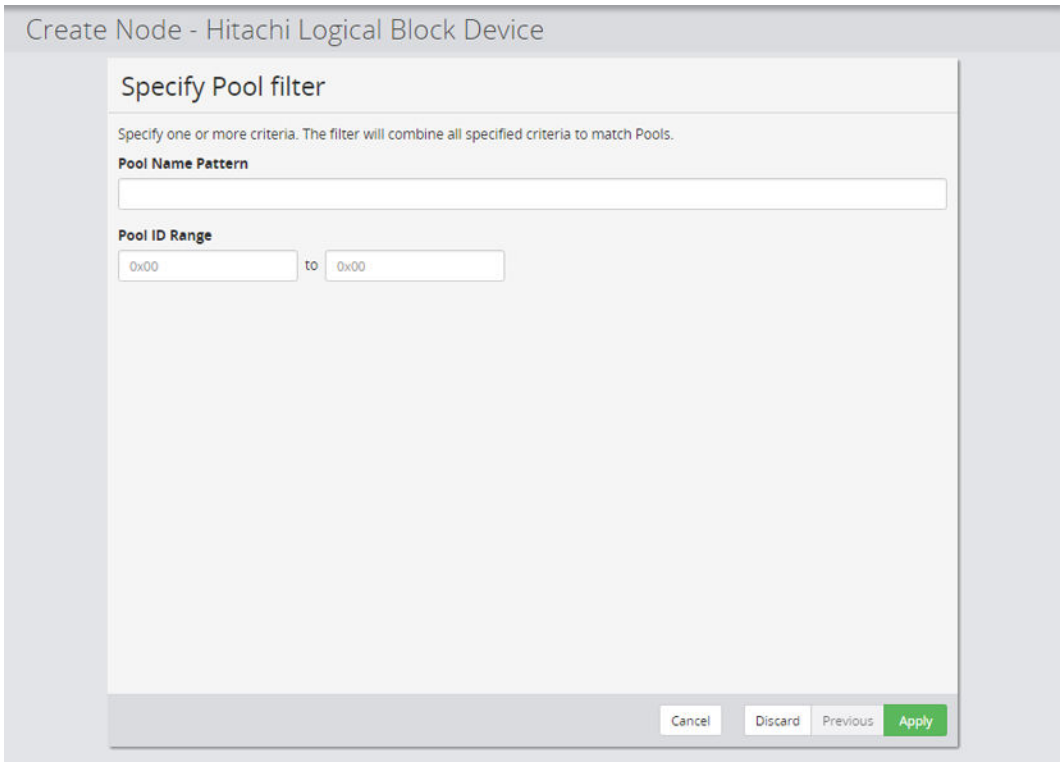


Figure 357 Hitachi Logical Block Wizard - Specify Pool filter

Specify one or more of the following parameters to identify the Pools:

Note: The terms entered here are logically ANDed together, so a Pool will only be selected if it matches the *Name Pattern* AND *ID Range*.

Control	Description
Pool Name Pattern	Enter a regular expression (see Basic regular expressions (on page 624)) to identify matching Pool Names.
Pool ID range	Enter the Pool ID range to use.

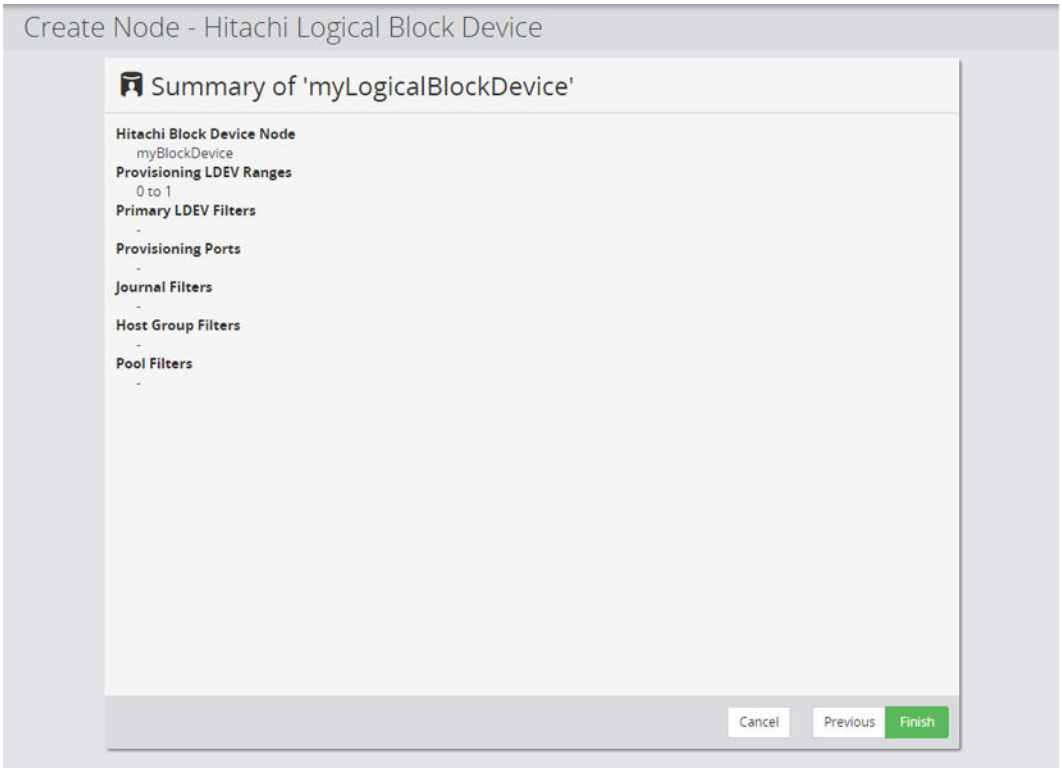


Figure 358 Hitachi Logical Block Wizard - Summary

Control	Description
Summary	Summary of the settings entered.

Basic regular expressions

Regular expressions (regex) are a powerful method of pattern matching in character strings. This topic describes a few basic regex terms. For a full explanation of regular expressions, please refer to one of the many books or online resources available.

Table 33 Commonly used regex matching terms

Type of Match	Regex Example	Action	Explanation
Contains the sub-string (case sensitive)	`this`	Matches any string containing the sub-string `this`, respecting the case e.g.: `Grab this thistle`	A sub-string on its own will be matched anywhere within the search string.
Contains the sub-string (case insensitive)	`(?i)this`	Matches any string containing the sub-string `this`, without regard to case e.g.: `Grab ThiS thIstle`	The `(?i)` *mode modifier* turns off case sensitive matching
Starts with the sub-string	`^this`	Matches any string starting with `this` e.g.: `this is a tree`	The `^` *anchor* means the sub-string must appear at the start of the string being searched
Ends with the sub-string	`that$`	Matches any string ending with `that` e.g.: `don't touch that`	The `$` *anchor* means the sub-string must appear at the end of the string being searched

Starts with any of the sub-strings	`^(this|that|other)`	Matches any string starting with `this` or `that` or `other`
		e.g.:
		`this time it matters`
		and:
		`that's life`
		and:
		`other types available`
			The `|` *alternation operator* indicates any sub-string in brackets will be matched

Type of Match	Regex Example	Action	Explanation

Contains any of the sub-strings	`(this|that)`	Matches any string containing `this` or `that`
		e.g.:
		`you can't touch this`
		and:
		`please don't do that`
			The `|` *alternation operator* indicates either sub-string in brackets will be matched

Starts with sub-string 1 and ends with sub-string 2	`^this.*that$`	Matches any string staring with `this` and ending with `that` e.g.: `this is the answer to that`	The `.` *metacharacter* matches any character, then the `*` *quantifier* indicates that zero or more of them will be matched

Hitachi Software Block Device Node Wizard

This wizard is launched when a new Hitachi Software Block Device Node is added to the Nodes Inventory.

Note: Only LINUX proxy nodes are supported for Software Block Device Nodes. Windows proxies are not supported.

Note: A Hitachi Software Block Device Node can not share a proxy with a Hitachi Block Device Node.

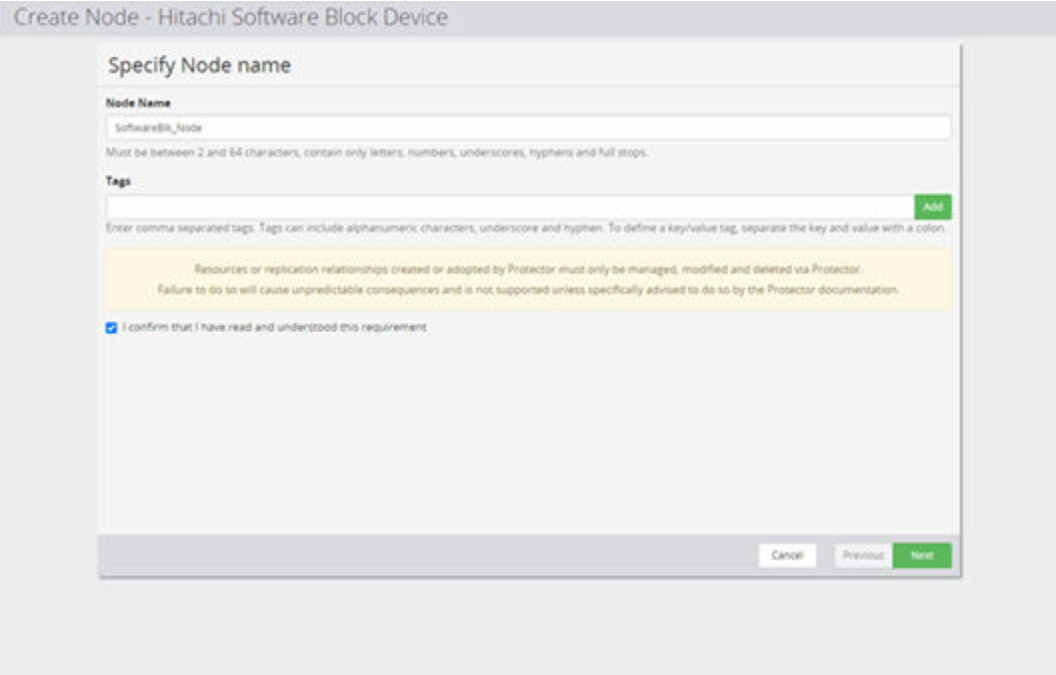


Figure 359 Hitachi Software Block Device - Specify Node Name

Control	Description
Node Name	Enter a name for the Hitachi Software Block storage node.
Tags	Add the tags to be associated with the object being created.
I confirm that ...	This checkbox must be checked to proceed with the node configuration. **Caution:** Resources or replication relationships created or adopted by Protector must only be managed, modified and deleted via Protector. Failure to do so will cause unpredictable consequences and is not supported unless specifically advised to do so by the Protector documentation.



Figure 360 Hitachi Software Block Device - Allocate node to Access Control Resource Group

Control	Description
Resource Groups	Select the resource group(s) to which this node will be allocated for the purposes of RBAC. All nodes are automatically allocated to the 'default' resource group.

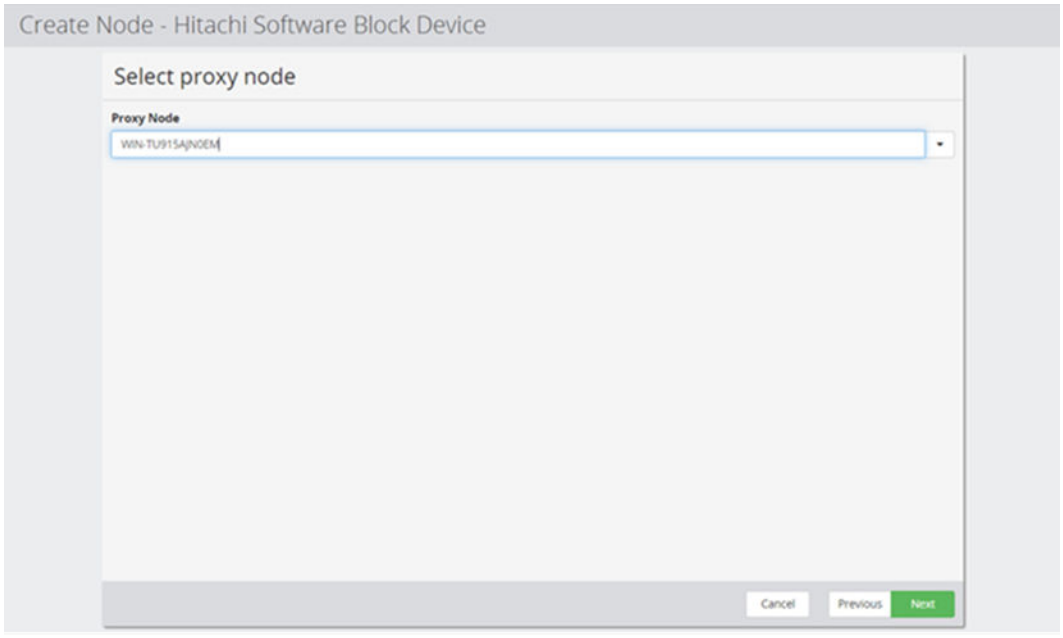


Figure 361 Hitachi Software Block Device - Select proxy node

Control	Description
Proxy Node	Select a Protector node to act as a proxy. **WARNING:** A Software Block Device Node can not share the same proxy node as a Block device node.

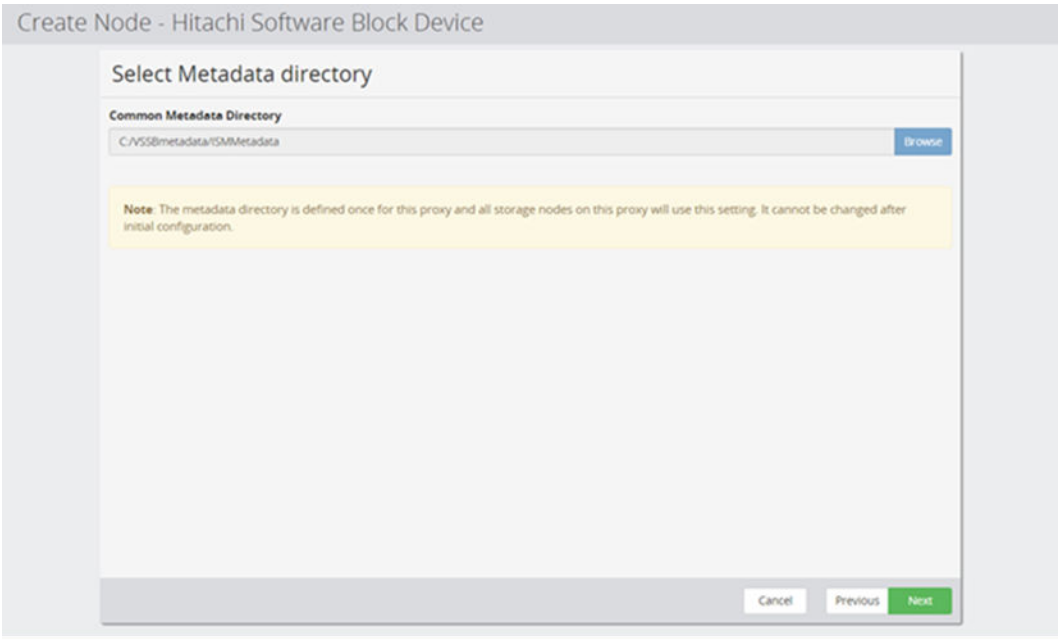


Figure 362 Hitachi Software Block Device - Select Metadata Directory

Control	Description
Common Metadata Directory	Enter a directory on the Proxy node where Protector can place metadata files related to Block snapshots and replications. Click the Browse button to open the Path Dialog (on page 350) if required. **Note:** The metadata directory is defined once for this proxy and all storage nodes on this proxy will use this setting. It cannot be changed after initial configuration.

Note: This screen is only displayed during node creation, it is not displayed in edit mode.

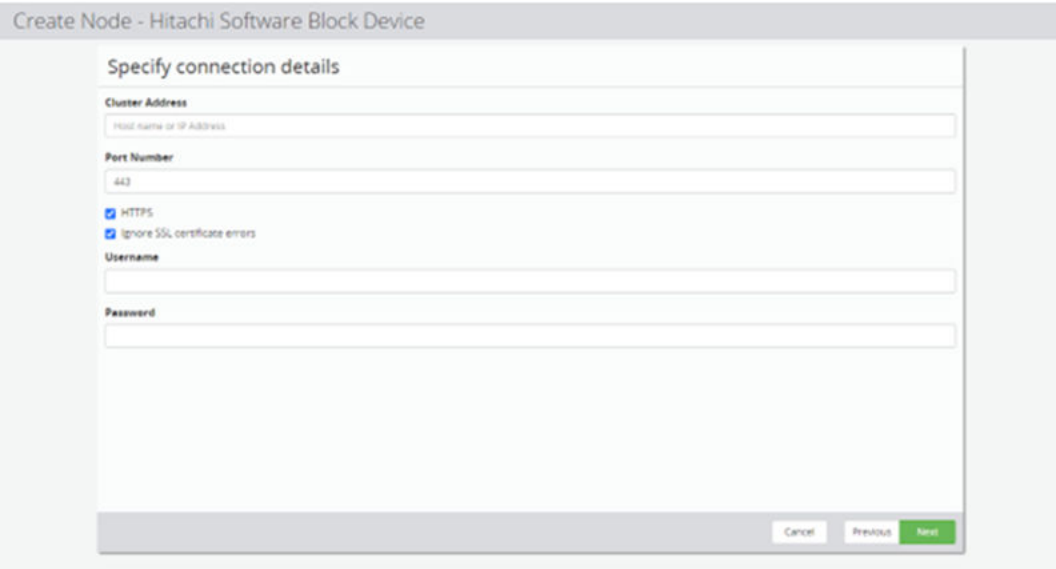


Figure 363 Hitachi Software Block Device - Specify Connection Details

Control	Description
Cluster Address	Enter either the Host Name or the IP Address of the Cluster server.
Port Number	Enter the port number for the Cluster server.
HTTPS	Select to enable secure communications to the Cluster server.
Ignore SSL certificate errors	Select to ignore SSL certificate errors when communicating with the Cluster server.
Username and Password	Enter the username and password for the storage array.

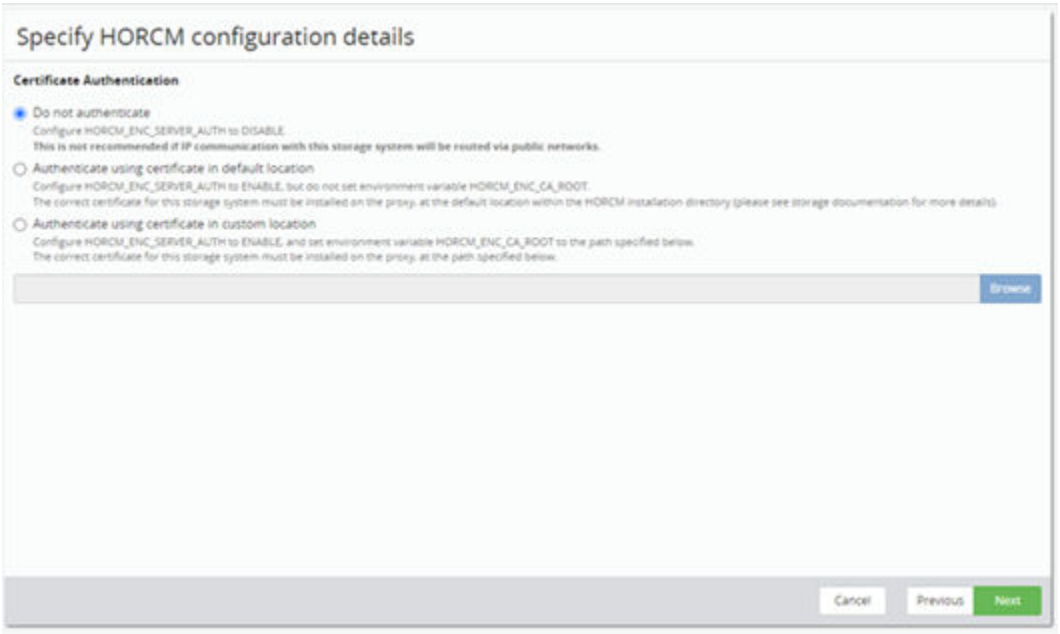


Figure 364 Hitachi Software Block Device - Certificate Authentication

Control	Description
Certificate Authentication	Select the authentication method required.

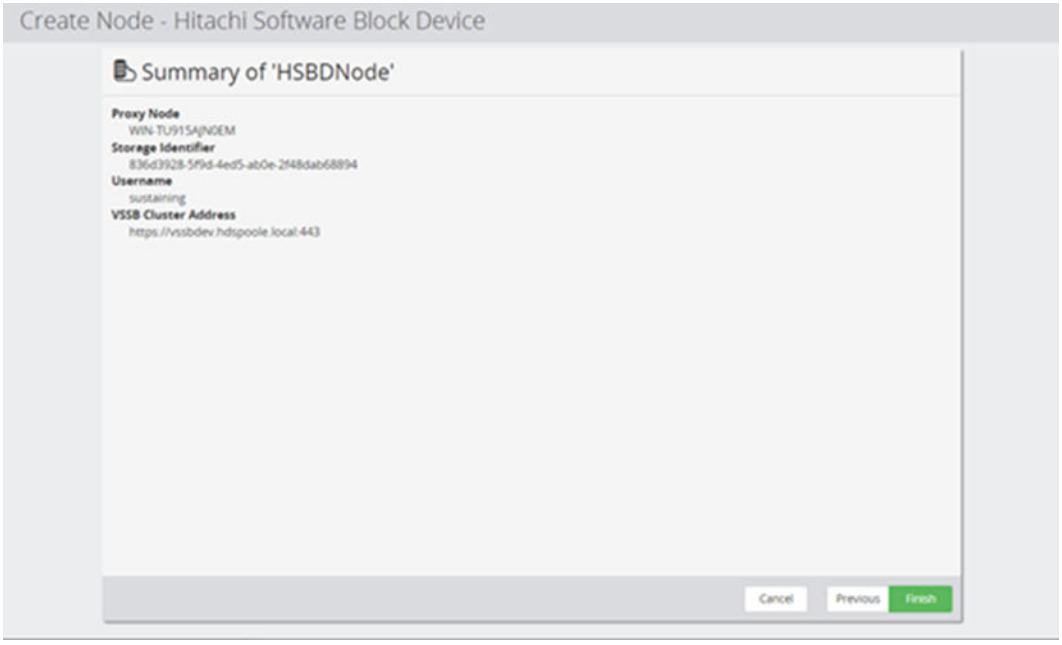


Figure 365 Hitachi Software Block Device - Summary

Control	Description
Summary	Summary of the settings entered.

Repository Storage Node Wizard

This wizard is launched when a new Repository Node is added to the Nodes Inventory.

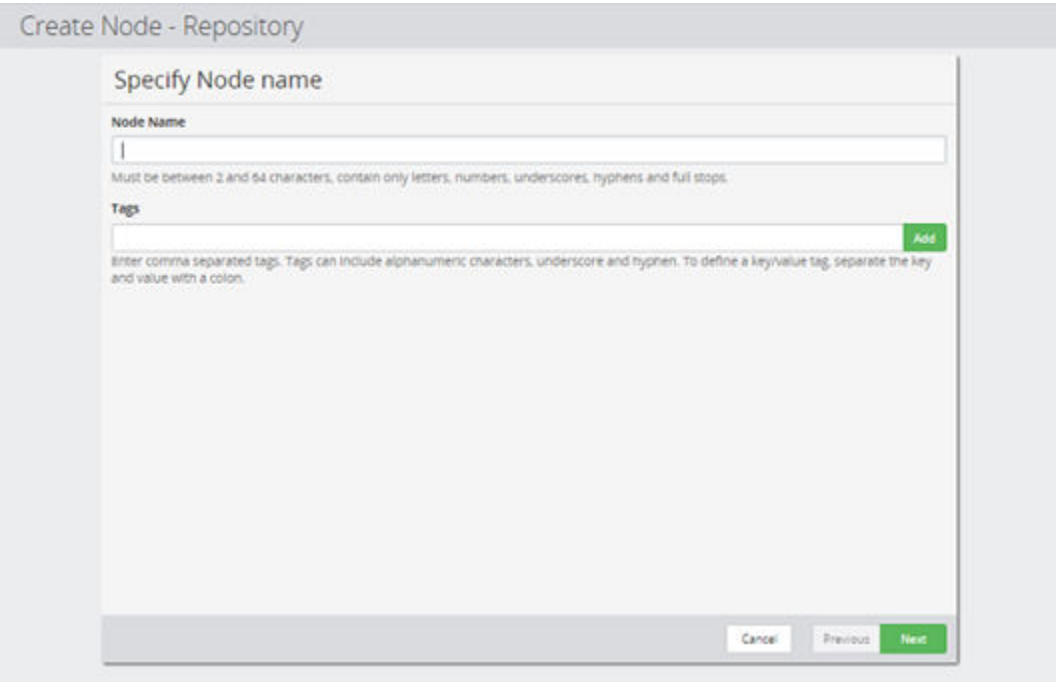


Figure 366 Repository Wizard - Specify Node name

Control	Description
Node Name	Enter a name for the Repository node. **Note:** Repository nodes that are hosted on the same Proxy Node must be uniquely named.
Tags	Add the tags to be associated with the object being created.

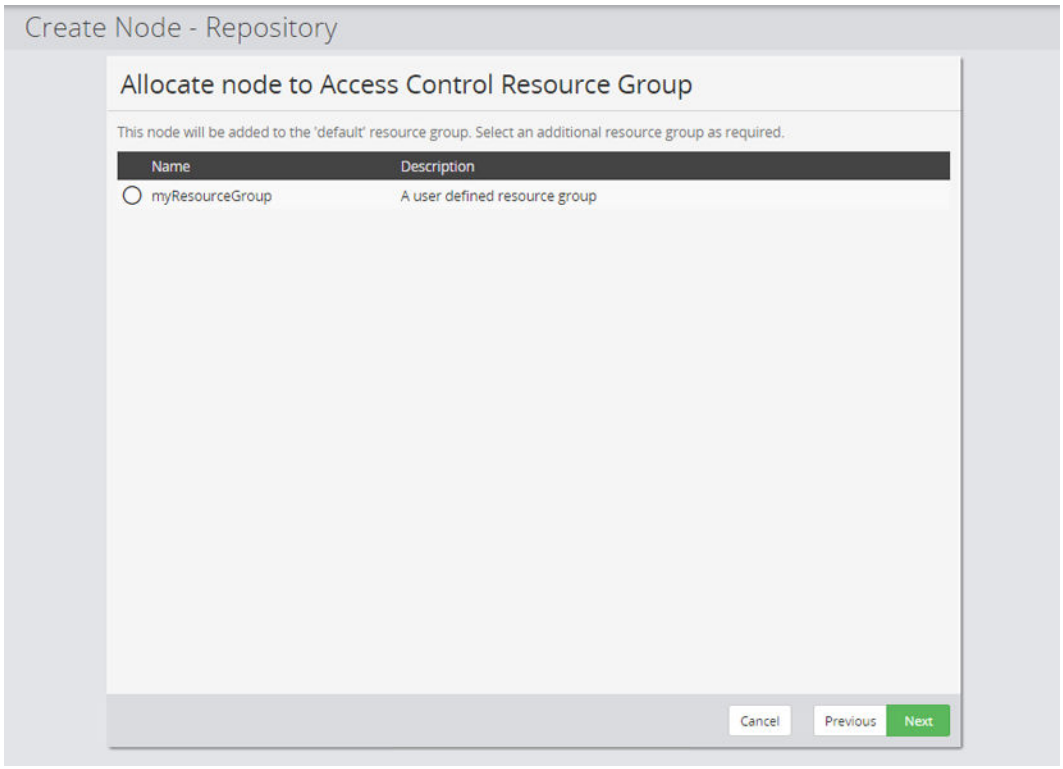


Figure 367 Repository Wizard - Allocate node to Access Control Resource Group

Control	Description
Resource Groups	Select the resource group(s) to which this node will be allocated for the purposes of RBAC. All nodes are automatically allocated to the 'default' resource group.

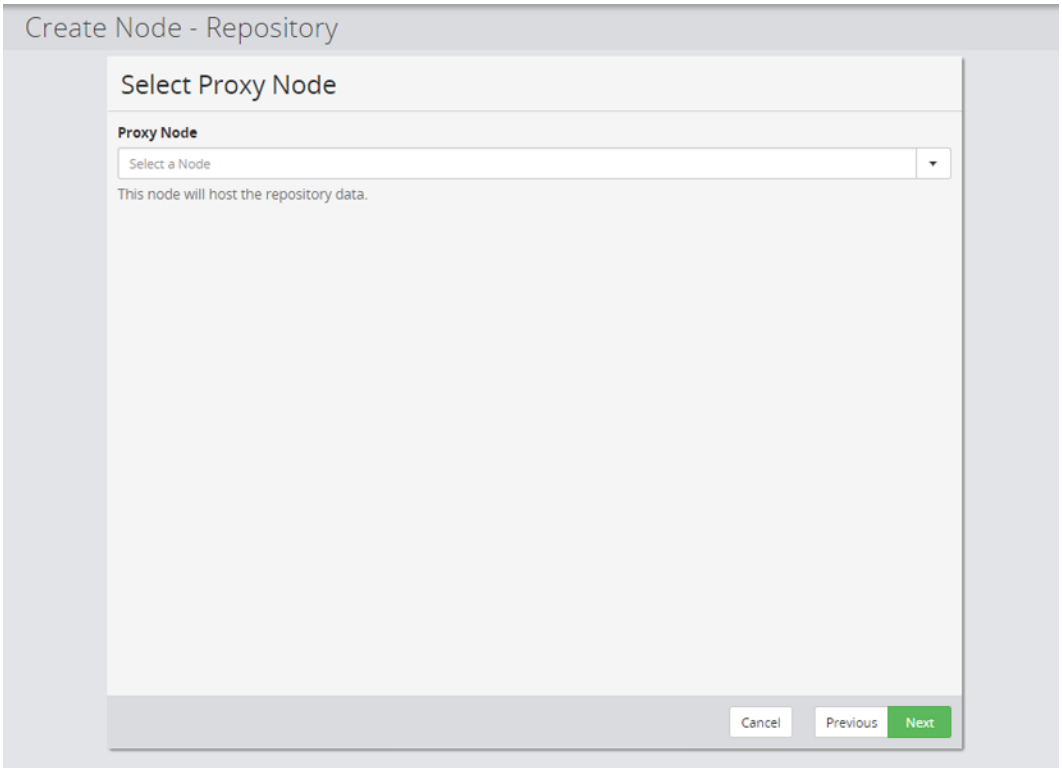


Figure 368 Repository Wizard - Select Proxy Node

Control	Description
Proxy Node	Enter or select the node that will host the Repository.

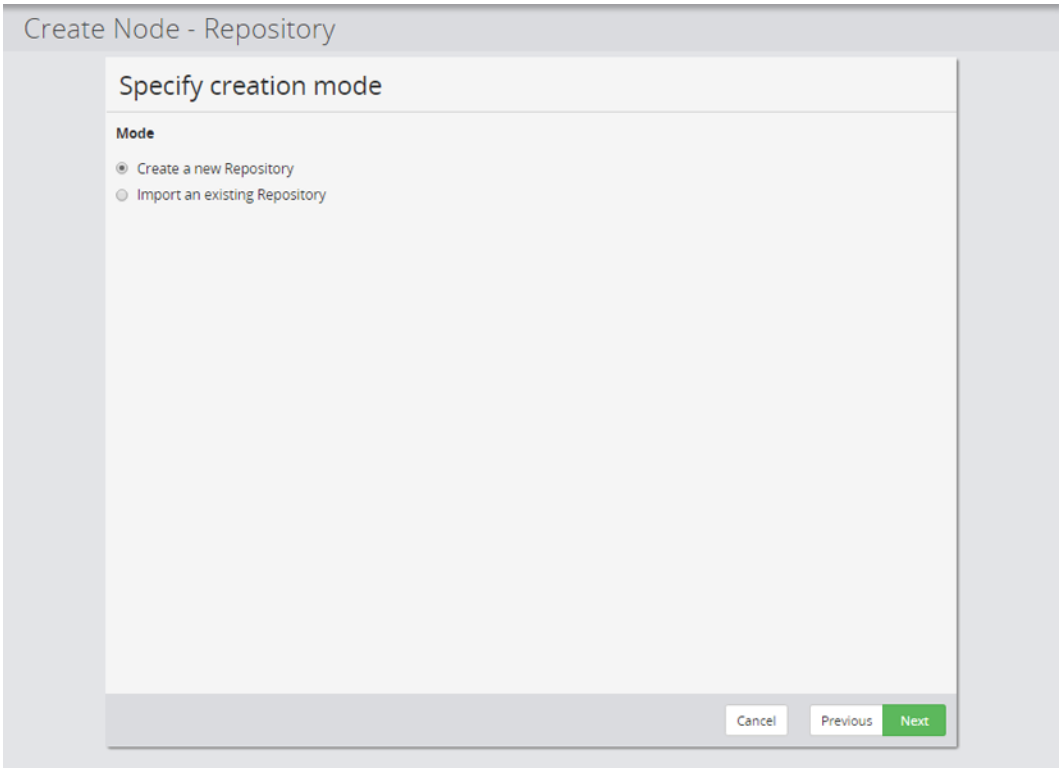


Figure 369 Repository Wizard - Specify creation mode

Control	Description
Mode	Select one of the following options: ▪ Create a new Repository ▪ Import an existing Repository If selection Import an existing Repository, you will be presented with the following wizard page.

Create New Repository

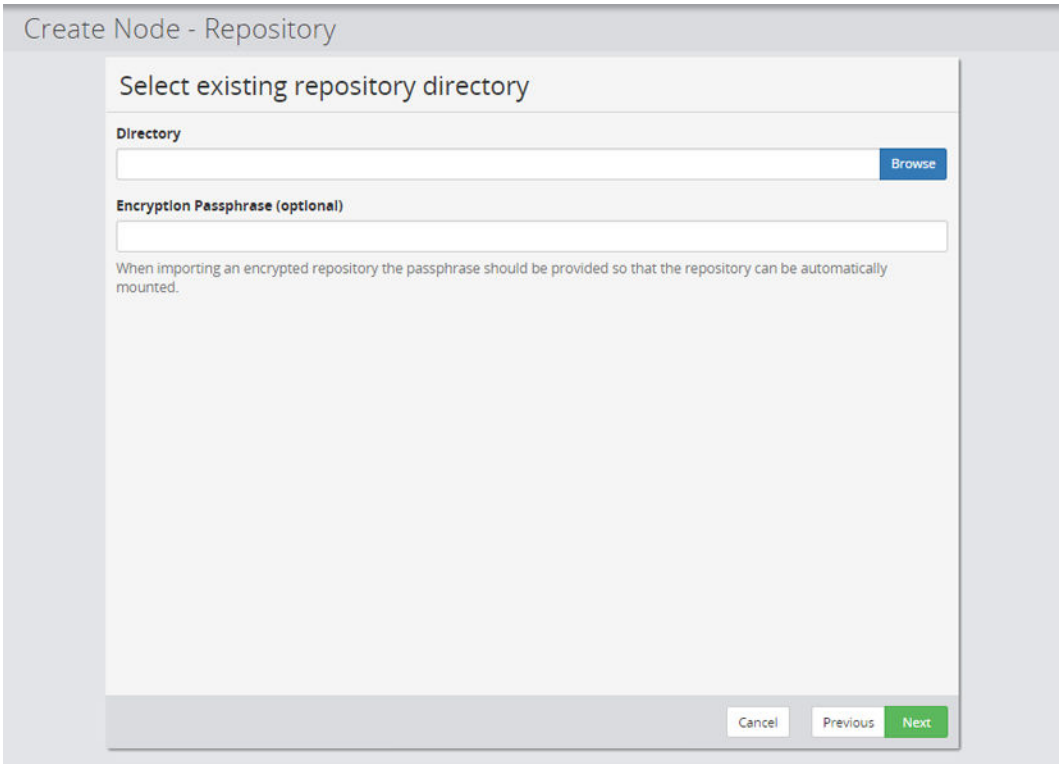


Figure 370 OR: Repository Wizard - Select existing repository directory

Control	Description
Directory	Enter the path of the directory where the existing repository resides or select the path by clicking the Browse button to open the Path Dialog (on page 350). Clicking next will take you to the Summary Page.
Encryption Passphrase (optional)	If the repository was encrypted, enter the password specified when it was created if you want to mount it immediately. The passphrase can be entered later if you intend to leave it unmounted for now.

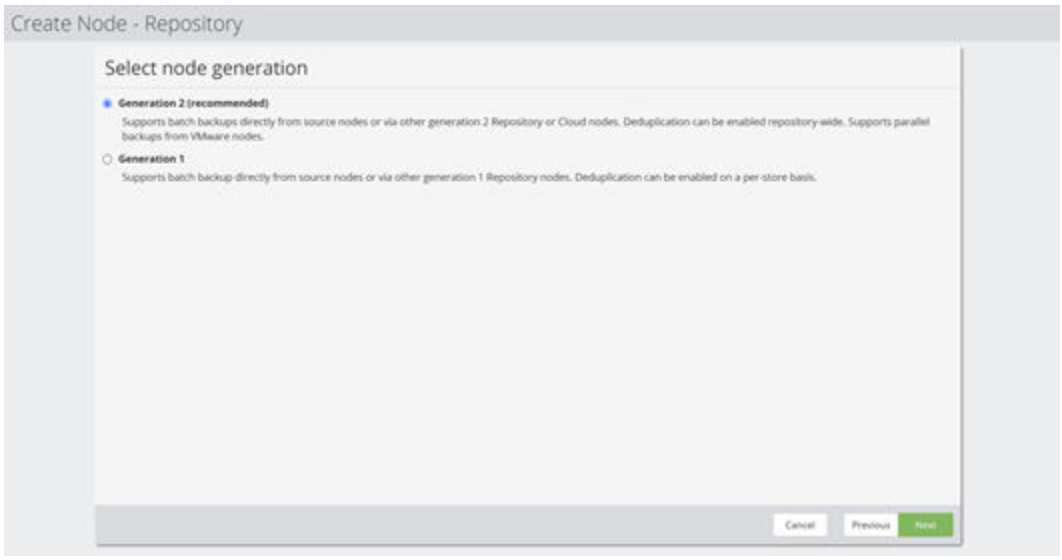


Figure 371 Repository Wizard – Select node generation

Control	Description
Generation 1	A generation 1 repository should only be selected if it is being used to receive data from another generation 1 repository. Generation 1 repositories are deprecated and wherever possible generation 2 repositories should be used.
Generation2	It is recommended to select a Generation 2 repository for all new repository setups. A generation 2 repository supports all new capabilities and utilizes the UBI framework.

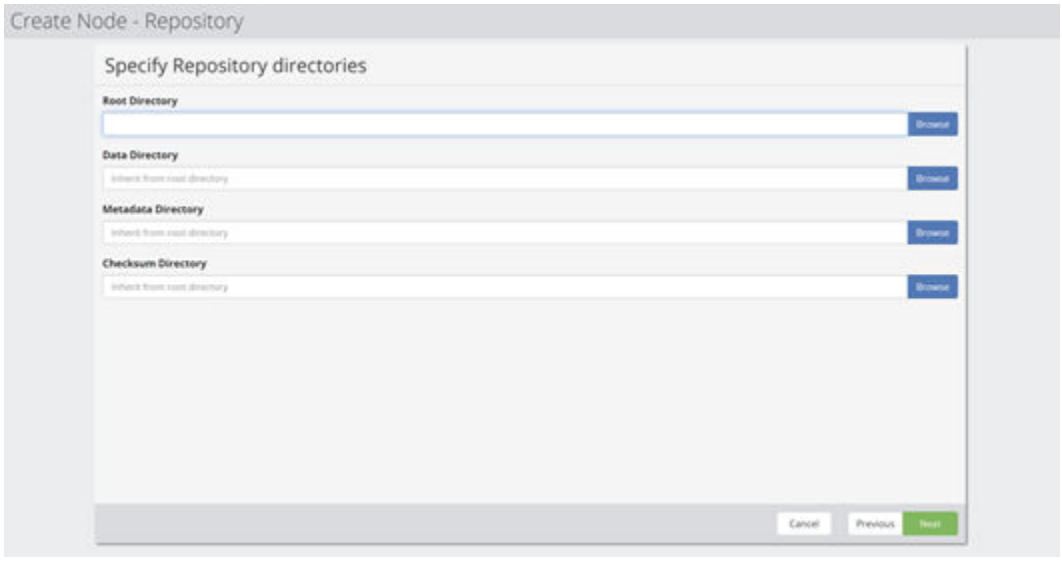


Figure 372 Repository Wizard – Specify Repository directories

Control	Description
Root Directory	The root directory is the location where the configuration data for the repository is held.
Data Directory	The data directory holds the actual backed up data contents of the repository.
Metadata Directory	The metadata directory holds the information that describes how the data in the data directory needs to be combined to make up a recovery point which can then be restored.
Checksum Directory	The checksum directory is used to store the checksums generated when Fine Change Detection is enabled.

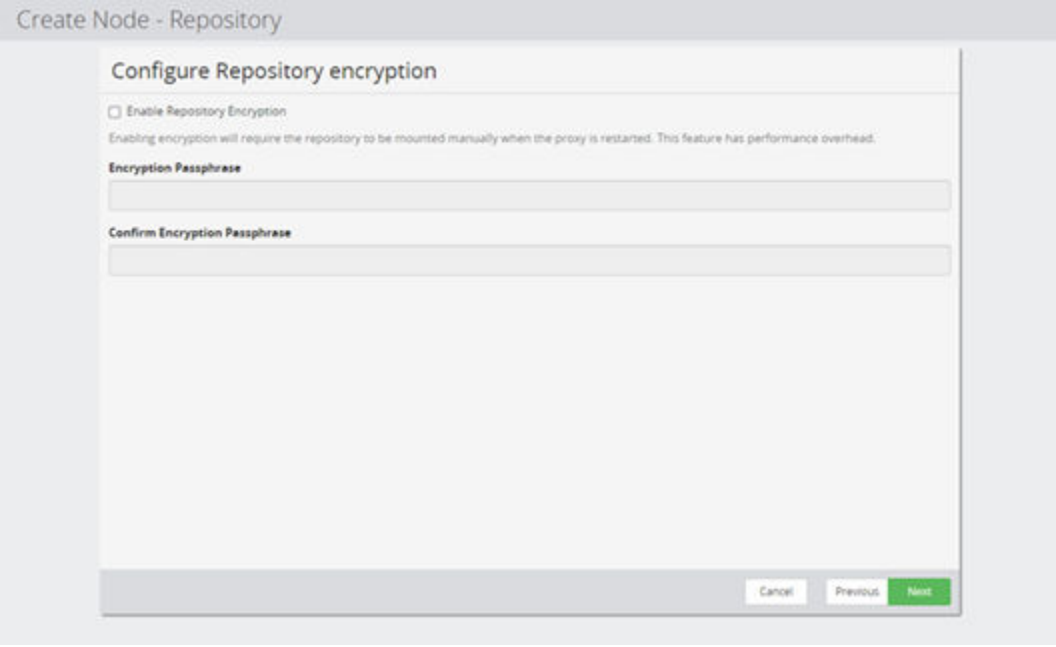


Figure 373 Repository Wizard – Configure Repository encryption

Control	Description
Enable Repository Encryption	You can optionally encrypt the contents of your repository to enhance security of the data stored in it. In-place encryption may not be allowed in some territories, this capability is controlled by license.
Encryption Passphrase	Enter an encryption password of your choosing. This will be required by anyone who accesses the repository. **Caution:** An encrypted repository cannot be recovered without the password. If the password is forgotten, then the data in that repository will be inaccessible.

Control	Description
Confirm Encryption Passphrase	Re-enter the encryption password to ensure it has been typed correctly.

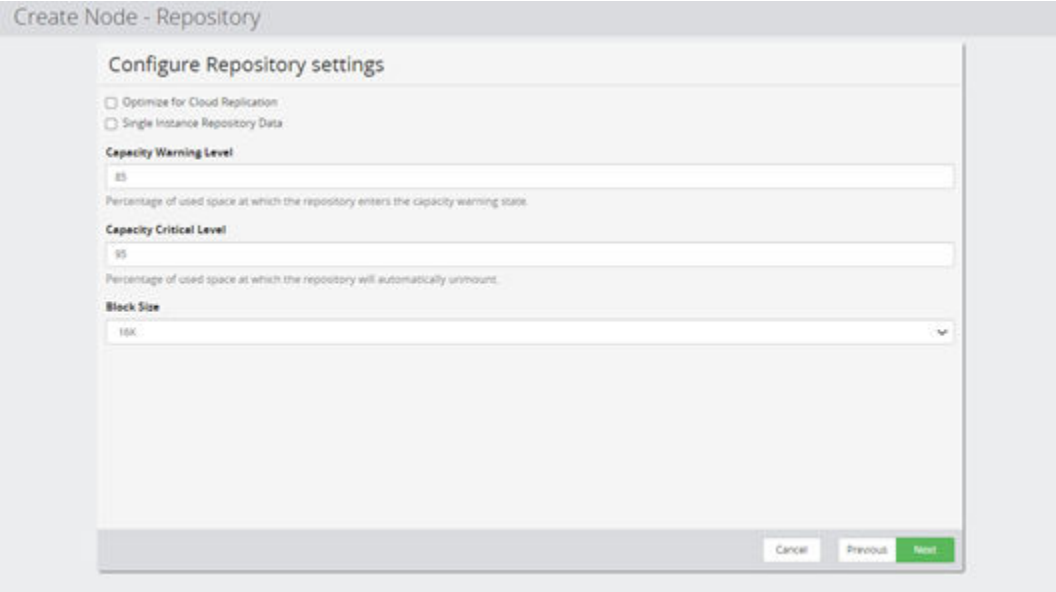


Figure 374 Repository Wizard – Configure Repository settings

Control	Description
Optimize for Cloud Replication	Aligns block boundaries between the repository and the cloud in order to make cloud uploads more efficient.
Single Instance Repository data	Enables the repository to perform file stream level single instancing. This reduces the space backups take up if you have multiple backups of the same file.
Capacity Warning Level	When the used space on the filesystem which hosts the repository reaches this percentage level warning logs will be generated advising of low disk space.
Capacity Critical Level	When the used space on the filesystem which hosts the repository reaches this percentage level an error log will be generated, and the repository will unmount.
Block Size	Sets the block size the repository uses. It is recommended that this matches the filesystem block size where possible.

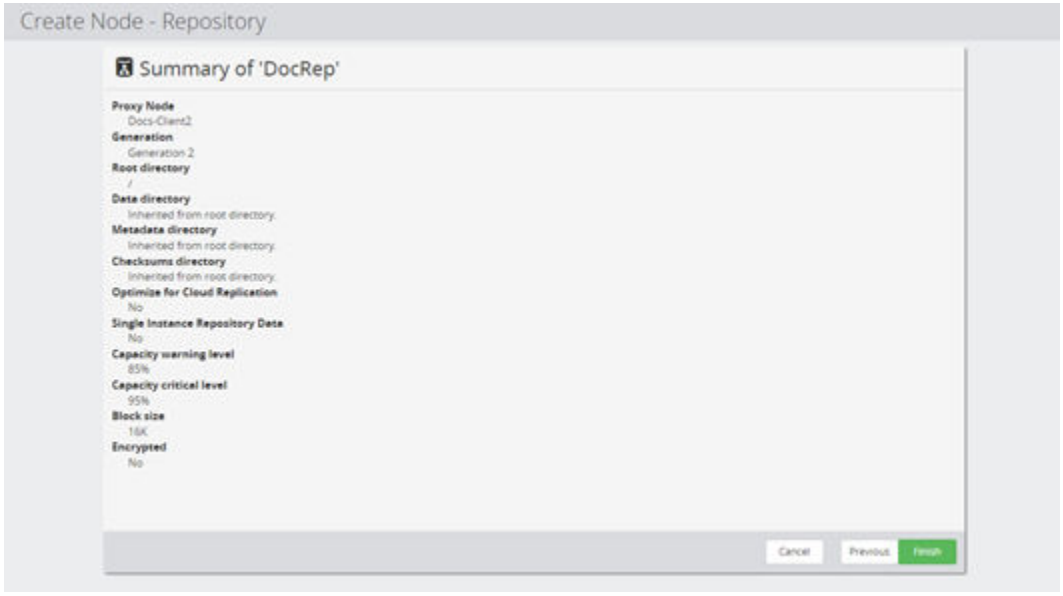


Figure 375 Repository Wizard - Summary

Control	Description
Summary	Summary of the parameters entered above.

Import an existing Repository

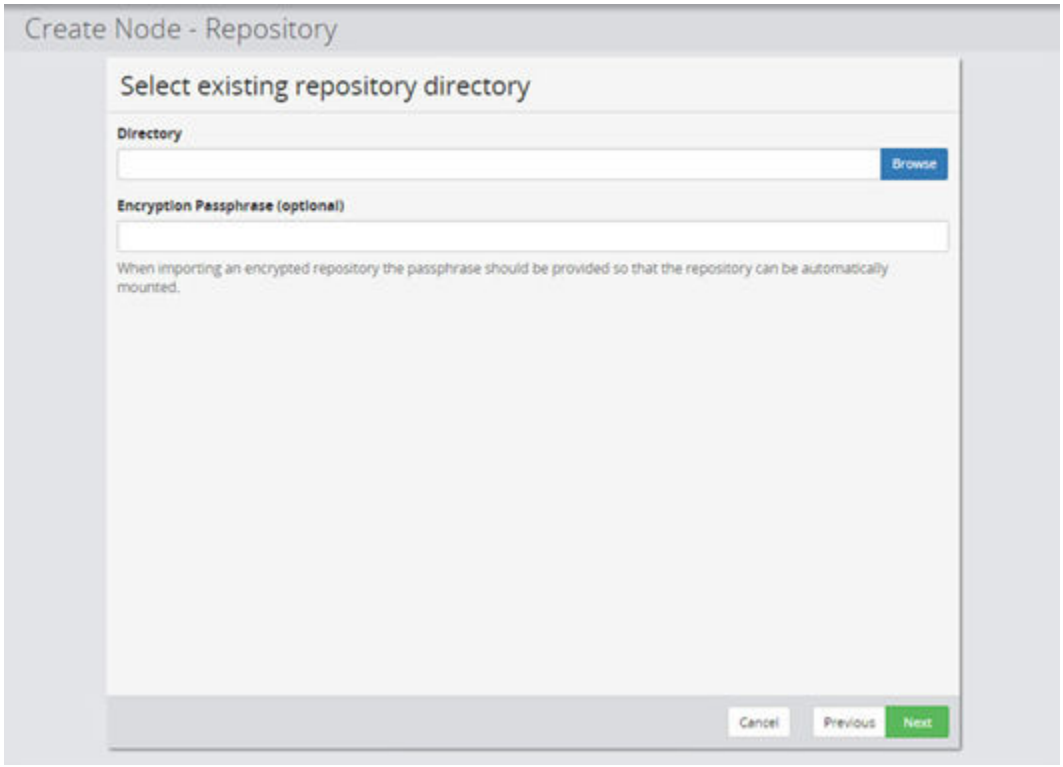


Figure 376 OR: Repository Wizard - Select existing repository directory for Import

Control	Description
Directory	Enter the path of the directory where the existing repository resides or select the path by clicking the Browse button to open the Path Dialog (on page 350). Clicking next will take you to the Summary Page.
Encryption Passphrase (optional)	If the repository was encrypted, enter the password specified when it was created if you want to mount it immediately. The passphrase can be entered later if you intend to leave it unmounted for now.

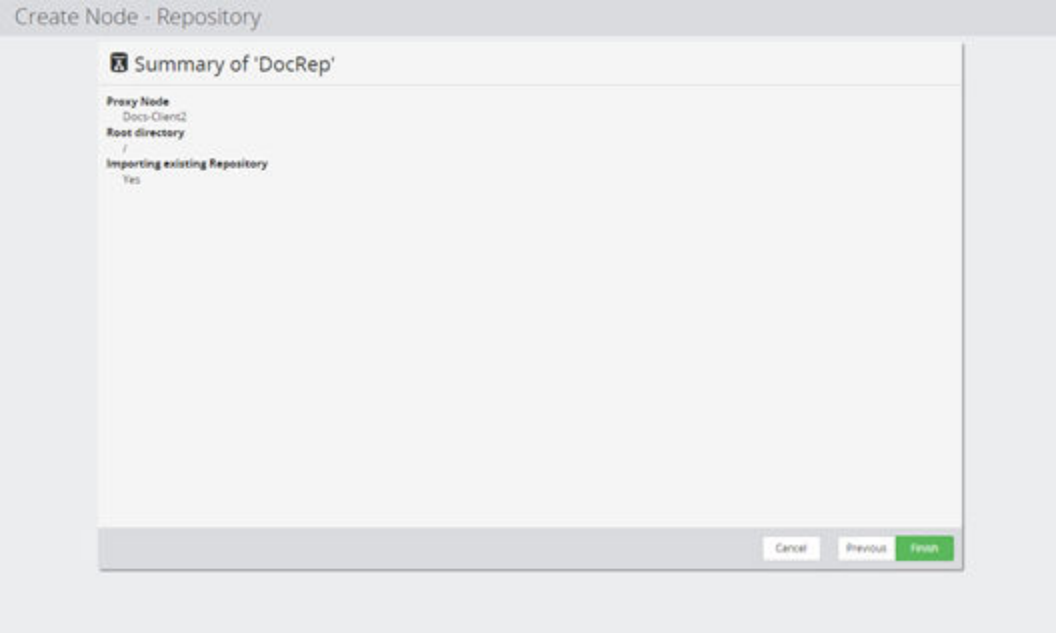


Figure 377 Repository Wizard – Summary of Import Repository Wizard

Control	Description
Summary	Summary of the parameters entered for Import Repository Wizard.

Amazon S3 Storage Node Wizard

This wizard is launched when a new Amazon S3 Node is added to the Nodes Inventory.

Note: Ensure that the Proxy Node defined within the wizard is time synced. If the time is different by over 5 minutes, then the node creation and future backup may fail. It is recommended that the system time on all nodes should have the correct date/time associated with their particular region and time zone.

Note: There are some differences between the Amazon S3 node creation and the HCP cloud scale node creation which will be noted in this section. ‘S3 Node’ will be used as a general term to represent an Amazon S3 node.

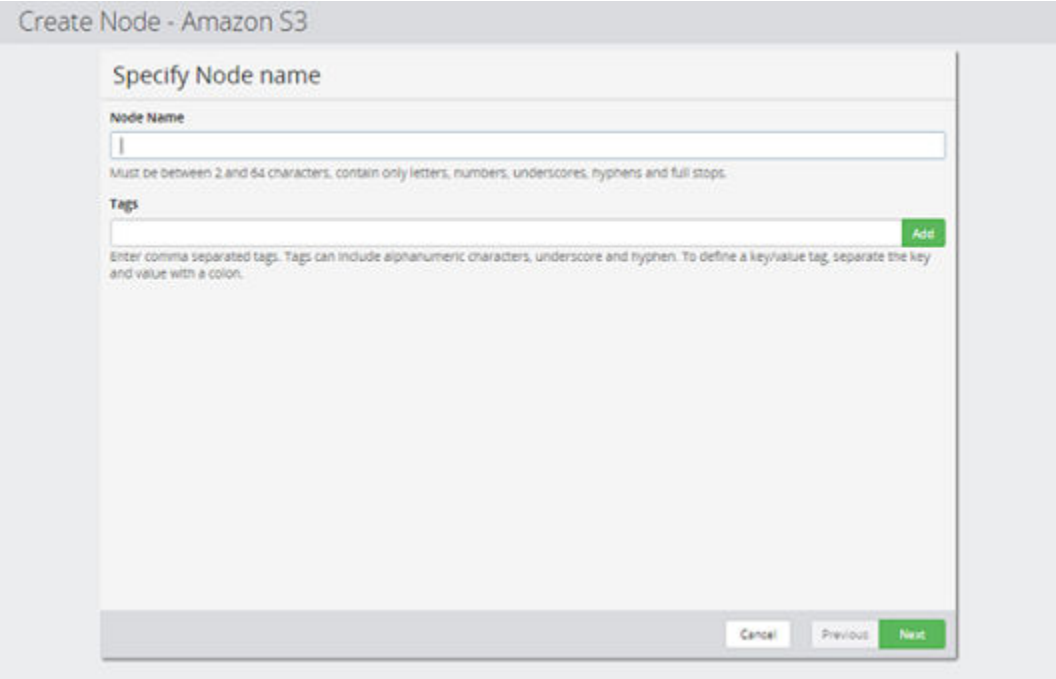


Figure 378 S3 Node Wizard - Specify Node Name

Control	Description
Node Name	Enter a name for the S3 node.
Tags	Add the tags to be associated with the object being created.

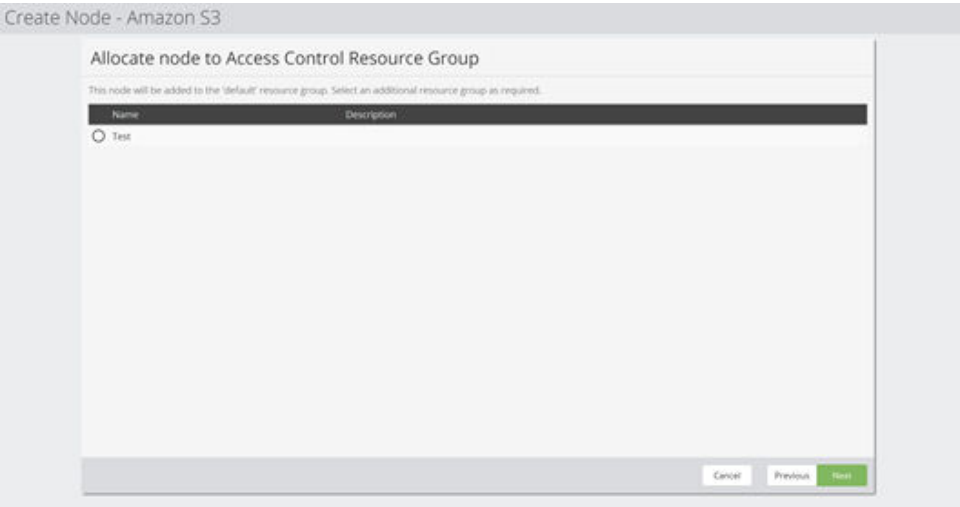


Figure 379 S3 Node Wizard - Allocate Node to Access Control Resource Group

Control	Description
Resource Groups	Select the resource group(s) to which this node will be allocated for the purposes of RBAC. All nodes are automatically allocated to the 'default' resource group.

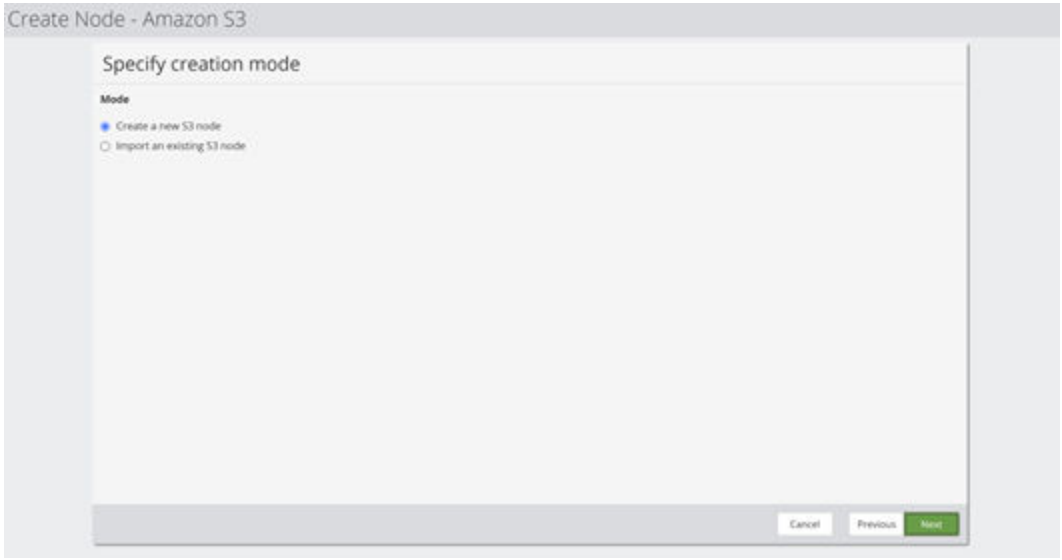


Figure 380 S3 Node Wizard - Specify Creation Mode

Control	Description
Mode	Select one of the following options: ▪ Create a new S3 node • This will start a new store in S3. ▪ Import an existing S3 Node ▪ • This will utilize a backup store that already exists. Subsequent screens are the same for either choices.

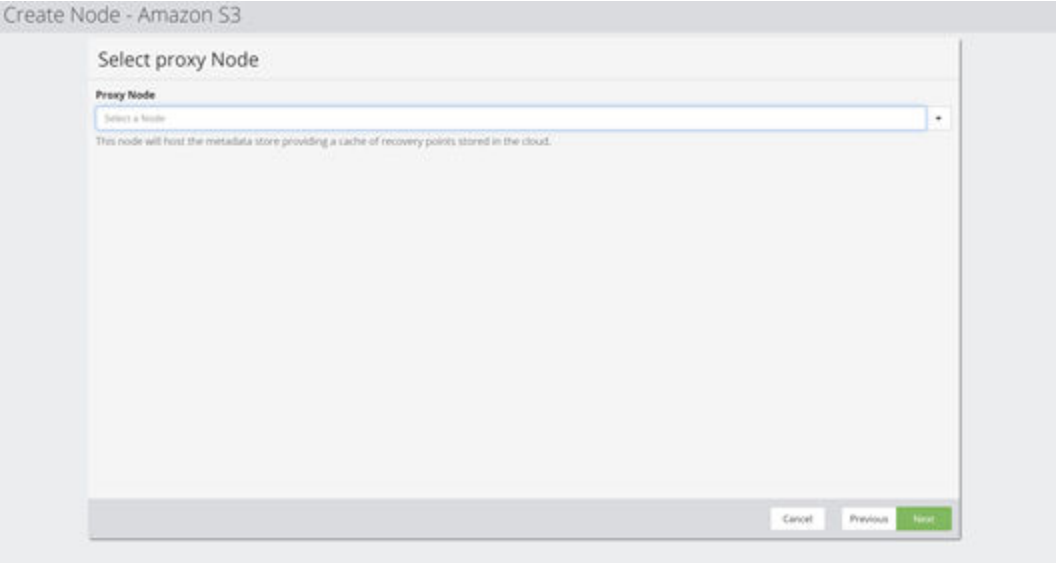


Figure 381 S3 Node Wizard - Select Proxy Node

Control	Description
Proxy Node	Enter or select the node that will host the S3 proxy.

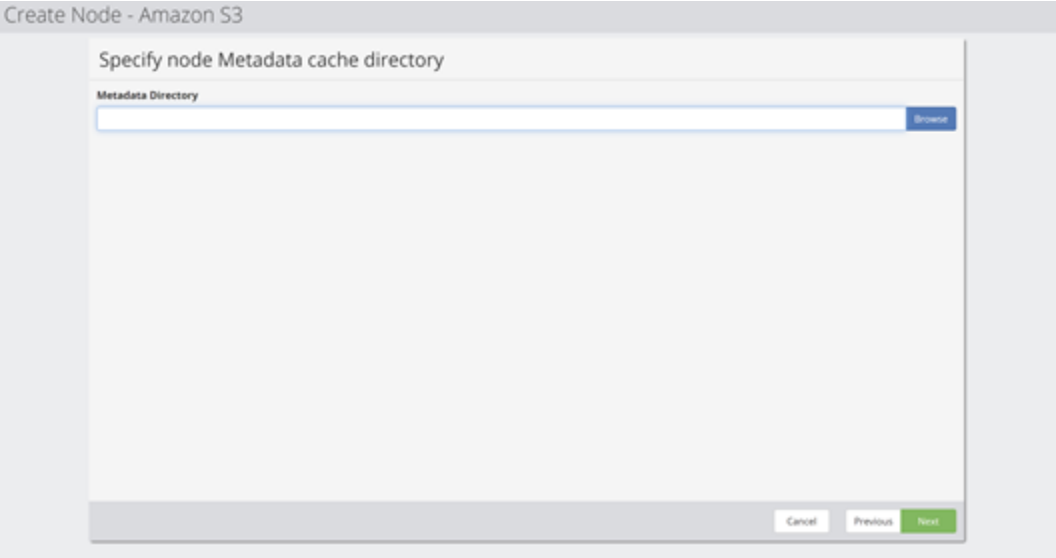


Figure 382 S3 Node Wizard - Specify Node Metadata Cache Directory

Control	Description
Directory	Enter the path of the directory where you want to create the new S3 metadata cache or select the path by clicking the Browse button to open the Path Dialog (on page 350). The metadata cache holds an index and metadata for backups that arestored in the S3 cloud. This is only a cache and can be rebuilt from the S3 cloud content if lost.

Amazon S3 specific configuration

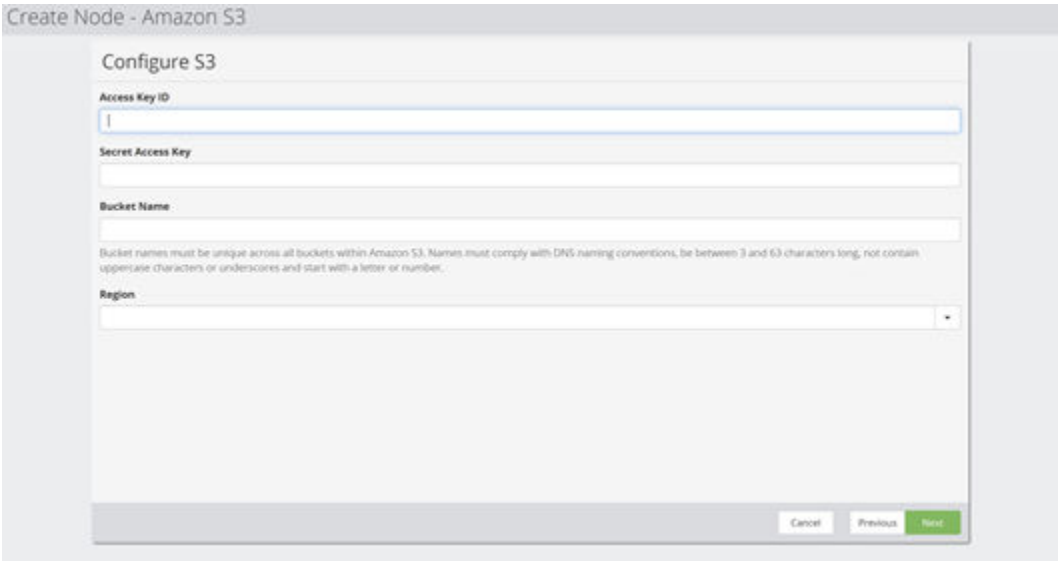


Figure 383 Amazon S3 Wizard - Configure S3

Control	Description
Access Key ID	Amazon S3 Access Key ID. This is the equivalent of a user . See Amazon S3 documentation for further information.
Secret Access Key	Amazon S3 Secret Access Key. This is the equivalent of a password. See Amazon S3 documentation for further information.
Bucket Name	This a name of where you will store data in the Amazon S3. It needs to be globally unique across all users of AWS across all regions. Consider the bucket name to be akin to a URL and it uses the DNS compliant naming conventions. See Amazon S3 documentation for further information.
Region	This is the region where the store will be created. Generally, you should select a store closest to your geographic region. See Amazon S3 documentation for further information.

HCP Cloud Scale Storage Node Wizard

This wizard is launched when a new HCP cloud scale node is added to the Nodes Inventory.

Note: Ensure that the Proxy Node defined within the wizard is time synced. If the time is different by over 5 minutes, then the node creation and future backup may fail. It is recommended that the system time on all nodes should have the correct date/time associated with their particular region and time zone.

Note: There are some differences between the Amazon S3 node creation and the HCP cloud scale node creation which will be noted in this section. ‘S3 Node’ will be used as a general term to represent HCP cloud scale node.

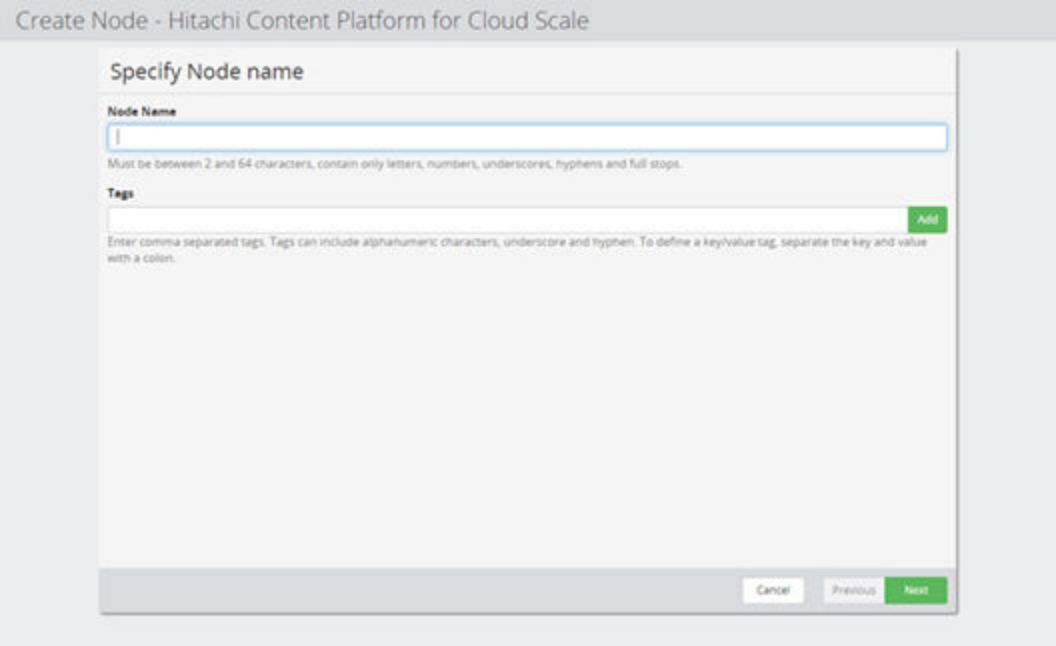


Figure 384 HCP Node Wizard - Specify Node Name

Control	Description
Node Name	Enter a name for the HCP node.
Tags	Add the tags to be associated with the object being created.

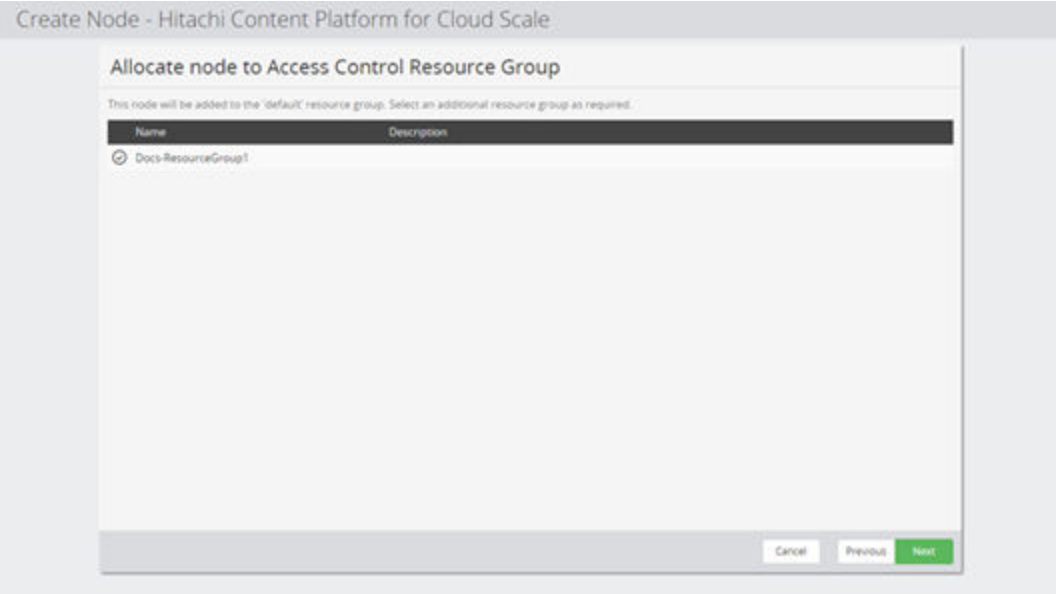


Figure 385 HCP Node Wizard - Allocate Node to Access Control Resource Group

Control	Description
Resource Groups	Select the resource group(s) to which this node will be allocated for the purposes of RBAC. All nodes are automatically allocated to the 'default' resource group.

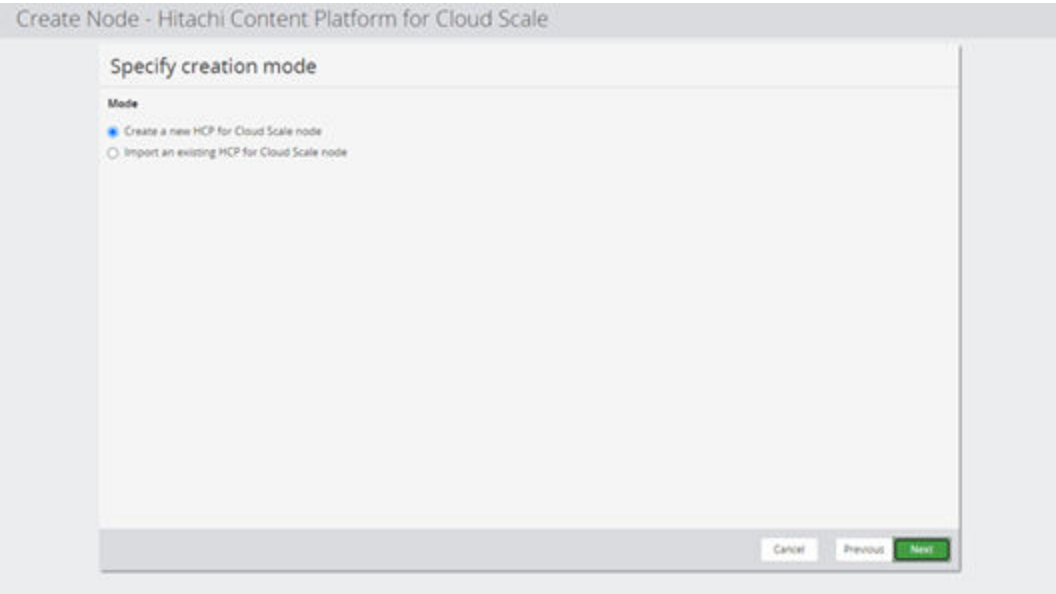


Figure 386 HCP Node Wizard - Specify Creation Mode

Control	Description
Mode	Select one of the following options: ▪ Create a new S3 node • This will start a new store in S3. ▪ Import an existing S3 Node ▪ • This will utilize a backup store that already exists. Subsequent screens are the same for either choices.

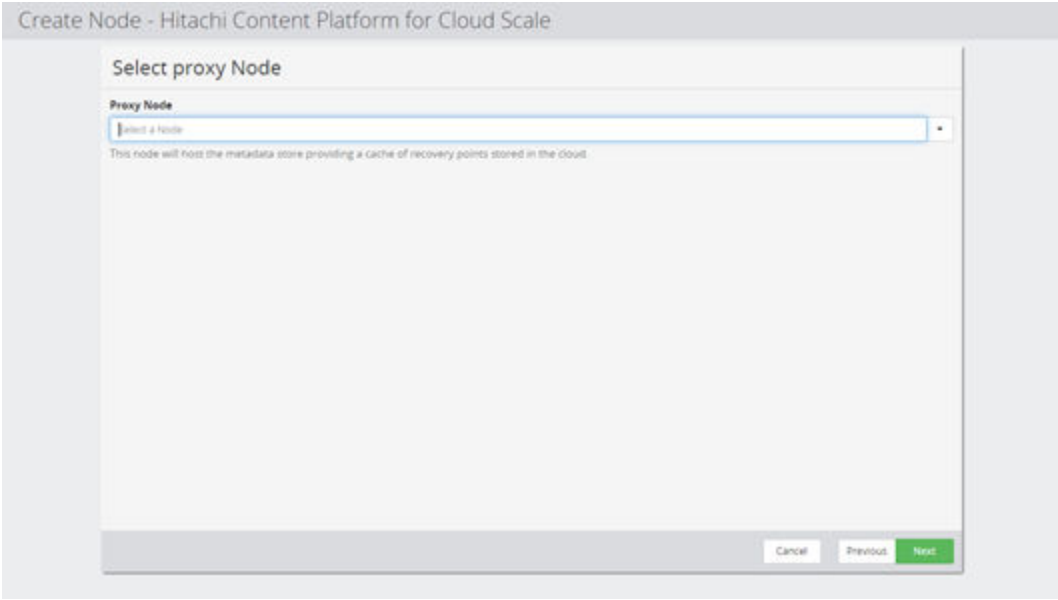


Figure 387 HCP Node Wizard - Select Proxy Node

Control	Description
Proxy Node	Enter or select the node that will host the S3 proxy.

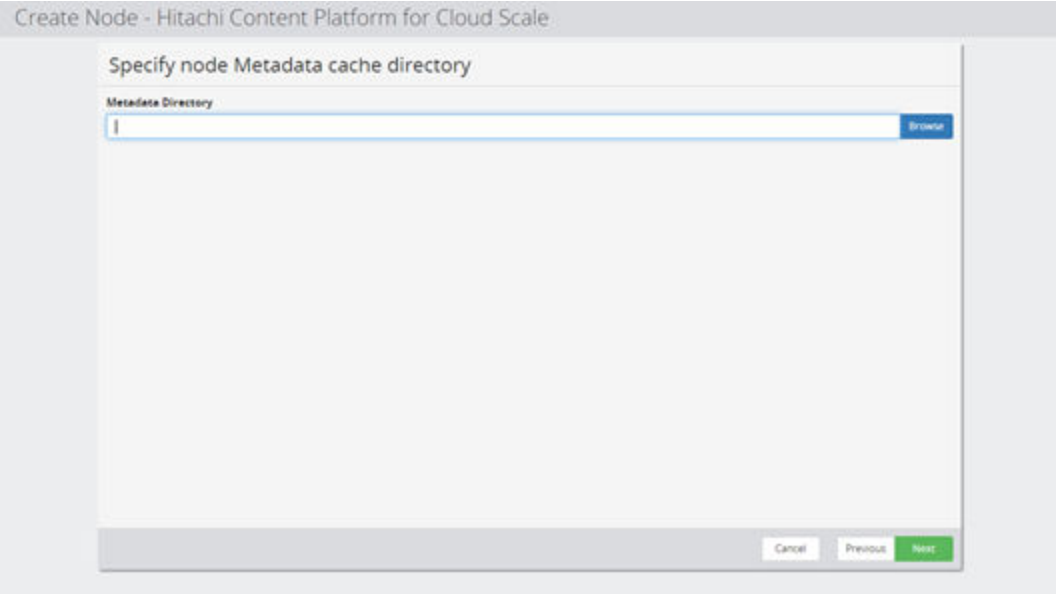


Figure 388 HCP Node Wizard - Specify Node Metadata Cache Directory

Control	Description
Directory	Enter the path of the directory where you want to create the new S3 metadata cache or select the path by clicking the Browse button to open the Path Dialog (on page 350). The metadata cache holds an index and metadata for backups that arestored in the S3 cloud. This is only a cache and can be rebuilt from the S3 cloud content if lost.

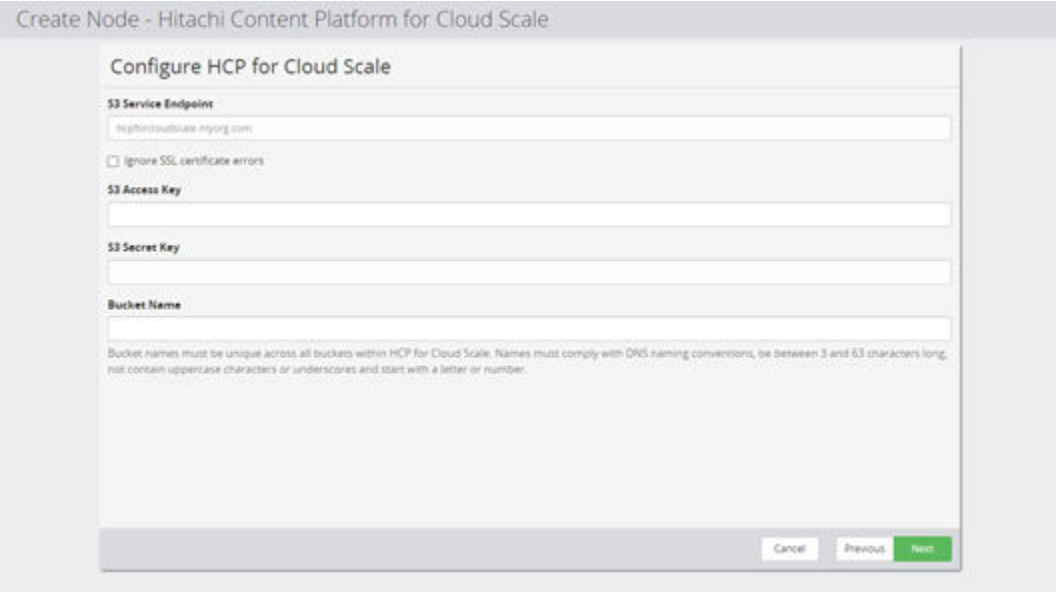


Figure 389 HCP cloud scale Wizard - Configure S3

Control	Description
S3 Service Endpoint	The HTTP endpoint to communicate with the HCP cloud scale S3 service.
Ignore SSL certificate errors	Ignore SSL errors when using self-signed certificates.
S3 Access Key	This is the equivalent of a user.
S3 Secret Key	This is the equivalent of a password.
Bucket Name	This a name of where you will store data on the HCP cloud scale device. Consider the bucket name to be akin to a URL and it uses the DNS compliant naming conventions.

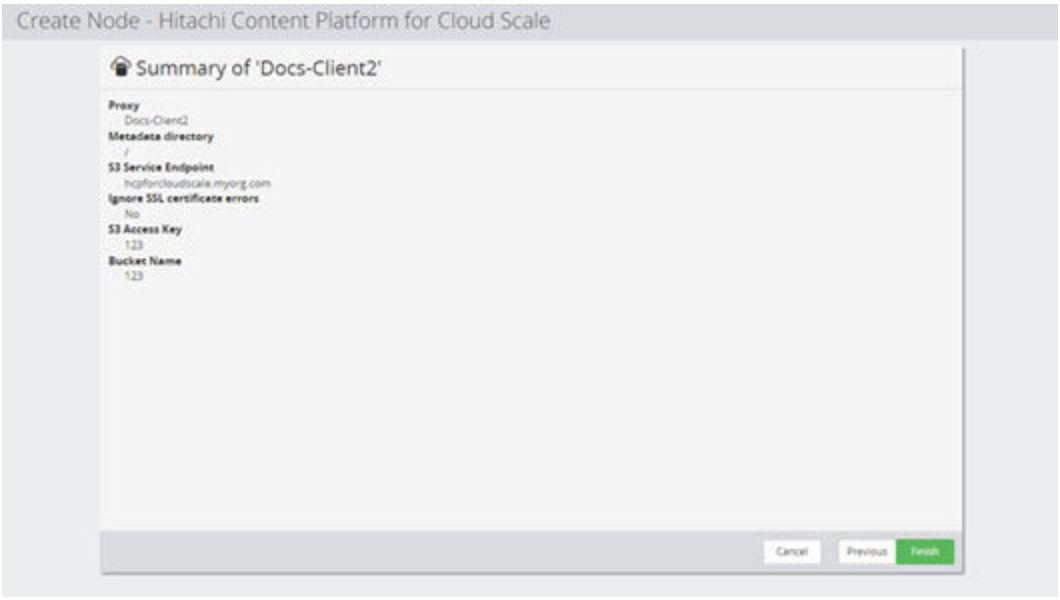


Figure 390 HCP Node Wizard - Summary

Control	Description
Summary	Summary of the parameters entered above.

Node Details

This page displays the details of an existing node and enables the wizard to be launched to edit them.

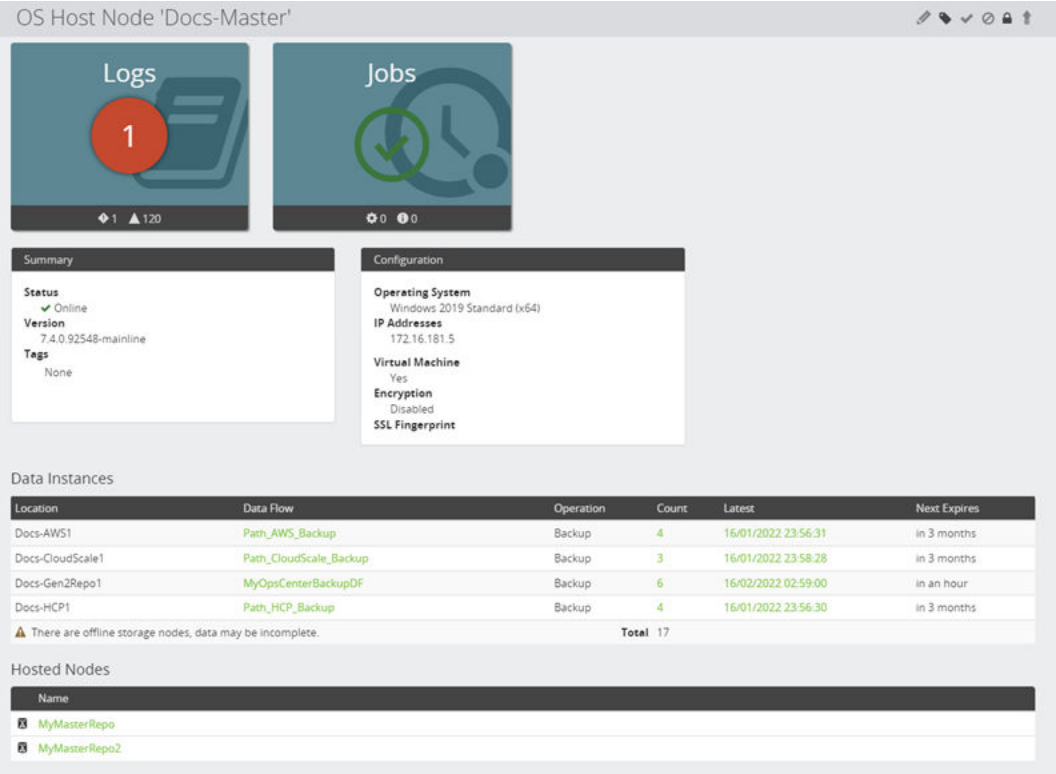


Figure 391 Node Details (OS Host)

Table 34 Common Controls (Displayed for all node types)

Control	Description
Logs	Indicates the number of unacknowledged error logs in red. If there are no errors then a green tick is displayed instead. Summarizes the number of unacknowledged errors and warnings. Opens the Logs Inventory (on page 509).
Jobs	Indicates the number of failed jobs in red. If there are no failed jobs then a green tick is displayed instead. Summarizes the number of failed, paused and in-progress jobs. Opens the Jobs Inventory (on page 493).
Summary	Indicates the following: ▪ Status (Online/Offline) - indicates if the node is contactable by the Master node. **Note:** OS Host have version. All other nodes have Proxy and Proxy Version • Proxy - Name of the OS Host node where this node reside. • Protector Version - The version number of the Protector software running on the proxy node. ▪ Tags

Control	Description
Data Instances	Indicates the following: ▪ Location – Name of the node where the data instance resides. ▪ Data Flow – Link to the Data Flow which describes the backup instance. ▪ Operation - Indicates operation type and name. ▪ Count – Indicates the number of recovery points for the storage system. ▪ Latest – Indicates the latest date and time when snapshot/backup is created. ▪  - Indicates that one or more node are offline and as such there may be instances residing on these nodes so there is a potential that the count is incomplete.

Table 35 OS Host Specific Controls

Control	Description
Configuration	Indicates the following: ▪ Operating System - the name of the operating system on the node. ▪ Virtual Machine - indicates the machine is running on a hypervisor. ▪ IP Addresses - all the IP addresses in use on the node.
Hosted Nodes	Lists all the nodes for which this node acts as the proxy.
 Enable \Disable Encryption	Enables disables SSL encyption on the node. For more details see How to enable or disable encryption on a node (on page 263).

Table 36 Oracle Specific Controls

Control	Description
 Edit	Launches the Oracle Application Node Wizard (on page 551) to enable the parameters to be edited.
Details	Displays the node's online status.
Configuration	Displays the node type and configuration.

Control	Description
Host Nodes	Lists all the nodes that act as the proxy for this node.

Table 37 VMware Specific Controls

Control	Description
 Edit	Launches the VMware Node Wizard (on page 574) to enable the parameters to be edited.
Configuration	Indicates the following: ▪ VMware Username. ▪ VMware Server.

Table 38 Hitachi Block Specific Controls

Control	Description
 Manage	Opens the storage Hitachi Block Device Details (on page 859) for this node.
 Edit	Launches the Hitachi Block Device Node Wizard (on page 583) to enable the parameters to be edited.
 Edit Tags	Add the tags to be associated with the object being created.
Configuration	Indicates the following: ▪ Storage Serial. ▪ Username. ▪ LDEV Range.
Port	List the available ports for this node.
Command Device Type	Lists the command devices and their type.
Virtual Serial	Lists the Virtual Serial devices.
Hosted Nodes	Lists all the nodes for which this node acts as the proxy.

Table 39 Hitachi Software Block Specific Controls

Control	Description
 Edit	Launches the Hitachi Block Device Node Wizard (on page 583) to enable the parameters to be edited.
 Edit Tags	Add the tags to be associated with the object being created.
Configuration	Indicates the following: ▪ Storage Identifier. ▪ Username. ▪ Cluster Address.

Table 40 Hitachi Block Host Specific Controls

Control	Description
 Edit	Launches the Hitachi Block Host Node Wizard (on page 559) to enable the parameters to be edited.
 Edit Tags	Add the tags to be associated with the object being created.
 Expand	Launches the Hitachi Block Host Resize Dialog (on page 354) to expand the one or more LDEVs associated with this Block Host node.
Summary	Displays the node's online status and tags.
Configuration	Indicates the Hitachi Block Device on which this node is based.
Included Logical Devices	List the LDEVs represented by this node.
Excluded Logical Devices	List the LDEVs excluded for this node.

Table 41 Hitachi Logical Block Specific Controls

Control	Description
 Manage	Opens the storage Hitachi Block Logical Device Details (on page 869) for this node.
 Edit	Launches the Hitachi Logical Block Device Node Wizard (on page 609) to enable the parameters to be edited.
 Edit Tags	Add the tags to be associated with the object being created.
Details	Displays the node's online status.
Configuration	Indicates the parent Hitachi Block Device Node.
LDEV Range	The LDEV provisioning ranges available for this node.

Table 42 Repository Specific Controls

Control	Description
 Manage	Opens the storage Generation 1 Repository Details (on page 921) for this node.
 Edit	Launches the Repository Storage Node Wizard (on page 632) to enable editing the capacity level.
 Edit Tags	Add the tags to be associated with the object being created.
Details	Displays the nodes proxy.
Configuration	Indicates the follow: ▪ Generation – The generation of the repository (affects supported options). ▪ Encryption – Repository encryption enabled or disabled. ▪ Cloud Optimized – (Generation 2) Data chunk size optimized for cloud storage. Enabled or disabled.

Control	Description
	▪ Single Instancing – (Generation 2) Store a single instance of a block. Enabled or disabled. ▪ Capacity Warning Level – Percentage of disk space at which the repository enters the warning state. ▪ Capacity Critical Level – Percentage of disk space at which the repository will automatically unmount. ▪ Block Size.
Paths	Indicates the following: ▪ Root Directory. ▪ Data Directory. ▪ Metadata Directory. ▪ Checksums Directory.

Table 43 Hitachi Content Platform Specific Controls

Control	Description
 Edit	Launches the Hitachi Content Platform Storage Node Wizard (on page 600) to enable the parameters to be edited.
 Edit Tags	Add the tags to be associated with the object being created.
Configuration	Indicates the following: ▪ Generation – The generation of the HCP node. ▪ Metadata directory. ▪ Username - The username used when connecting to the HCP Tenant Console. ▪ Tenant Address - The address used when connecting to the HCP Tenant Console. ▪ HTTPS - Use HTTP or HTTPS when transferring data to and from HCP.
Namespace	▪ Name. ▪ Created.

Control	Description
	▪ Hard Quota – The size of name space. ▪ Soft Quota – Percentage of hard quota.

Node Type Icons

The node icons and indicators used in the lists of unauthorized and authorized nodes.

Icon	Description
	Master Node - Indicates that this node is the master node.
	OS Host Node - A Workstation, Laptop or Server node with the Protector services installed.
	Repository Node - A node with the Protector Repository component installed. Used to store snapshots of data backed up from various types of source nodes.
	VMware vCenter Node - A node containing virtual machines, and accessed through a separate proxy node.
	Hitachi Content Platform Node - A node representing an Hitachi Content Platform.

Icon	Description
	Hitachi Block Node - A node representing a Hitachi Block Storage Device.
	Hitachi Block Host Node - A node that represents a number of Hitachi Block Storage volumes.
	Hitachi Logical Block Device Node - A logical view of Hitachi Block Storage Device node, having limited access to its resources.
	Oracle Node - Represents a cluster of servers running as an Oracle RAC that accesses databases on shared disks.
	Unknown Node - The master node is unable to determine any node details, and cannot connect to it.
	Inactive Node - Indicates that the node is not currently running the Protector hub process, or it is not currently attached to the network. The master node is currently unable to connect to the node.
	Node Group - Represents a group of nodes that can be treated identically when assigning policies or monitoring activities.

Icon	Description
	S3 Node - A node representing an S3 Bucket

Node Status Indicators

Small graphical indicators that provide visual cues as to a particular node's status.

Icon	Description
	Unauthorized Node - The node has Protector installed but has not been authorized by the Master node, so cannot interact with other Protector nodes.
	Authorized Online Node - The node has Protector installed and has been authorized by the Master node, so can interact with other Protector nodes.
	Offline - The node is not contactable from the Master so its status cannot be determined.
	Snapshotted Node - Indicates that the node has a policy containing a snapshot operation applied to it.

Notifications UI Reference

This section describes the Notifications UI, accessed via the Navigation Sidebar (on page 291).

For further information, refer to:

- Notification Concepts (on page 108)
- Notification Tasks (on page 264)

Notifications Inventory

This page lists all configured log notifications and enables you to launch the wizard to create and edit them.

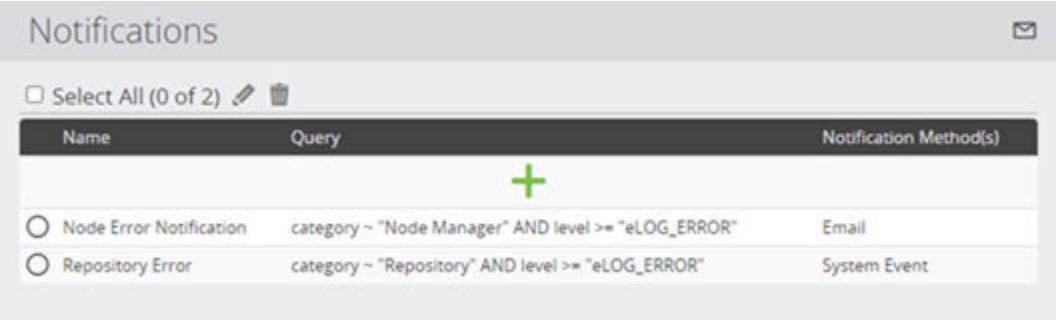


Figure 392 Notifications

Control	Description
Configure Email Settings	Launches the Email Notifications Settings Wizard (on page 660) to guide you through the process.
Edit	Enabled when only one Notification is selected. Launches the Notification Wizard (on page 662) to allow the settings to be edited.
Delete	Enabled when one or more Notifications are selected. Deletes selected Notifications.
Add	Launches the Notification Wizard (on page 662) to guide you through setting up a Notification.
Notification(s)	Any number of user defined Notifications can be created. These are displayed in the table in priority order. The table headings are as follows: ▪ Name – Name of the notification item. ▪ Query – The formatted query used to identify log messages which trigger the notification. ▪ Notification Method(s) - How the user will be notified.

Email Notifications Settings Wizard

This wizard provides the options for configuring the email account to send notifications from and to.

Note: Protector supports single step email account authentication. If the email account used to send notification is configured to use two step authentication, check that it is possible to allow applications to use less secure sign-in.

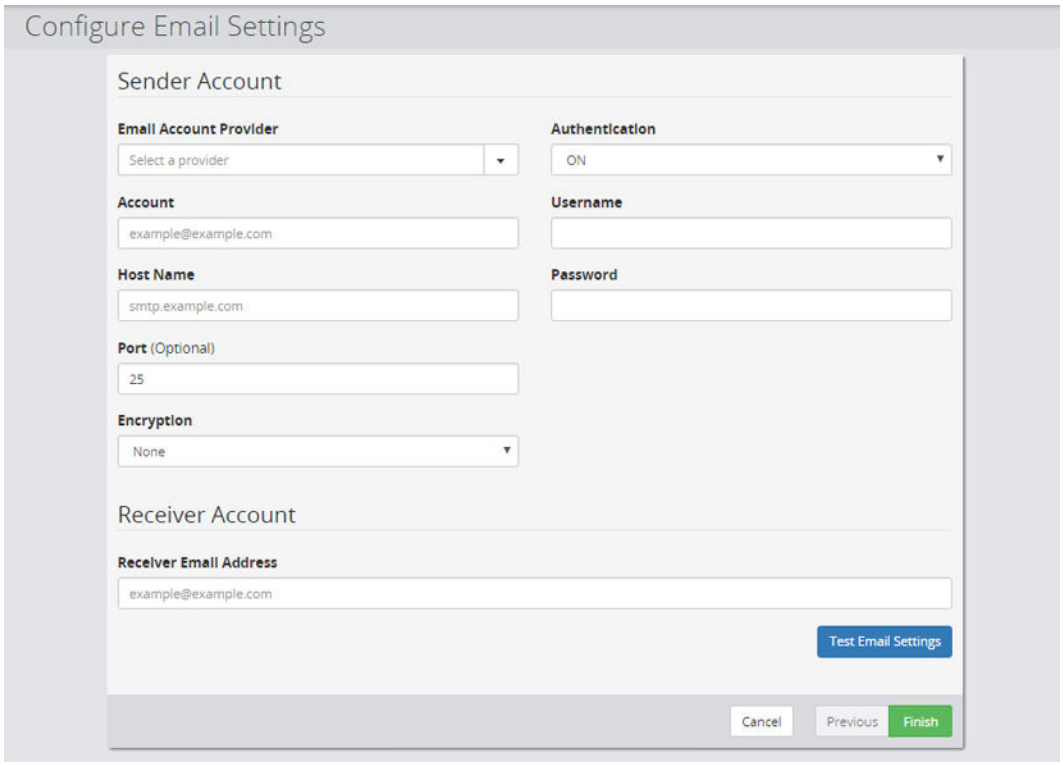


Figure 393 Configure Email Settings

Control	Description
Email Account Provider	Select the account provider. Can be one of: ▪ Google Mail ▪ Yahoo Mail ▪ Hotmail ▪ Microsoft Exchange ▪ Other
Account	Email address the notifications will be sent from.
Host Name	Host name for the email account that the notifications will be sent from.
Port (Optional)	The port number to use to contact the email server.
Encryption	Select the encryption method to be used. Can be one of: ▪ None ▪ START_TLS ▪ TLS
Authentication	Set Authentication either ON, OFF or NTLM.

Control	Description
Username	Only enabled when Authentication is set to On. Enter the username of the email account being used.
Password	Only enabled when Authentication is set to On. Enter the password of the email account being used.
Receiver Email Address	Enter the email address of the account(s) which should be the recipients of the notifications.
Test Email Settings	Clicking this button will send a test email based on the settings configured. If there are missing entries on the dialog then they will be highlighted.

Notification Wizard

This wizard provides the options for configuring notifications.

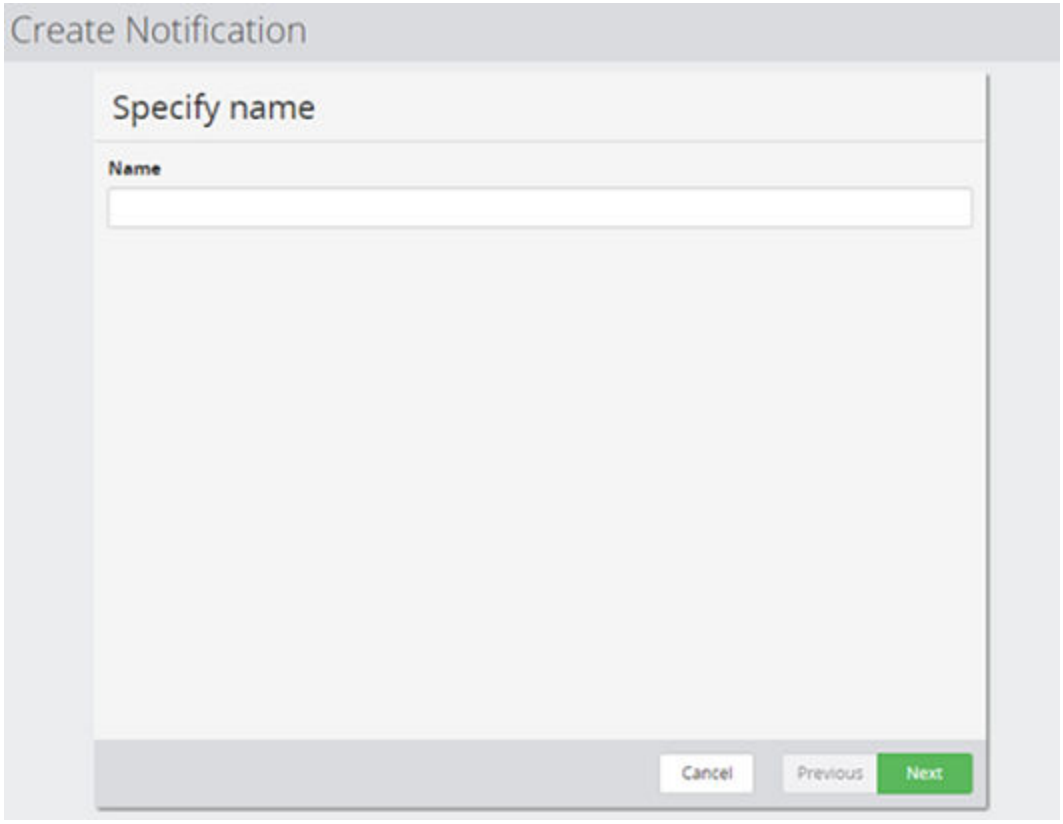


Figure 394 Log notification wizard – name

Control	Description
Name	Enter the name of the log notification.

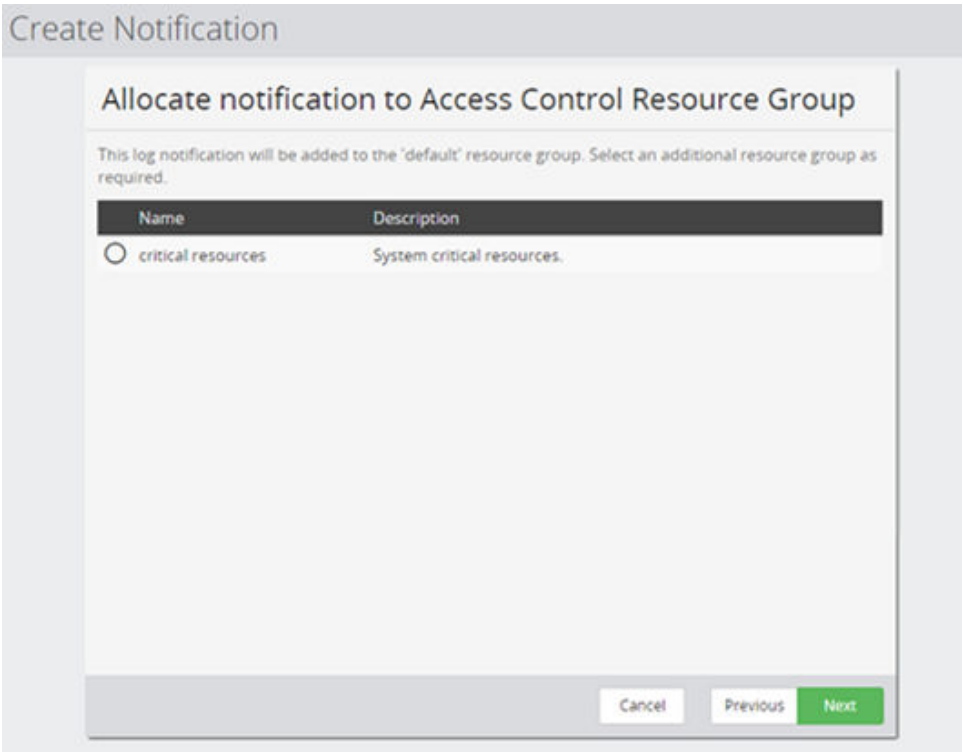


Figure 395 Log notification wizard – Allocate to Access Control Resource Group

Control	Description
Resource Groups	Select the resource groups to which this log notification will be allocated for the purposes of RBAC. All log notification are automatically allocated to the 'default' resource group.

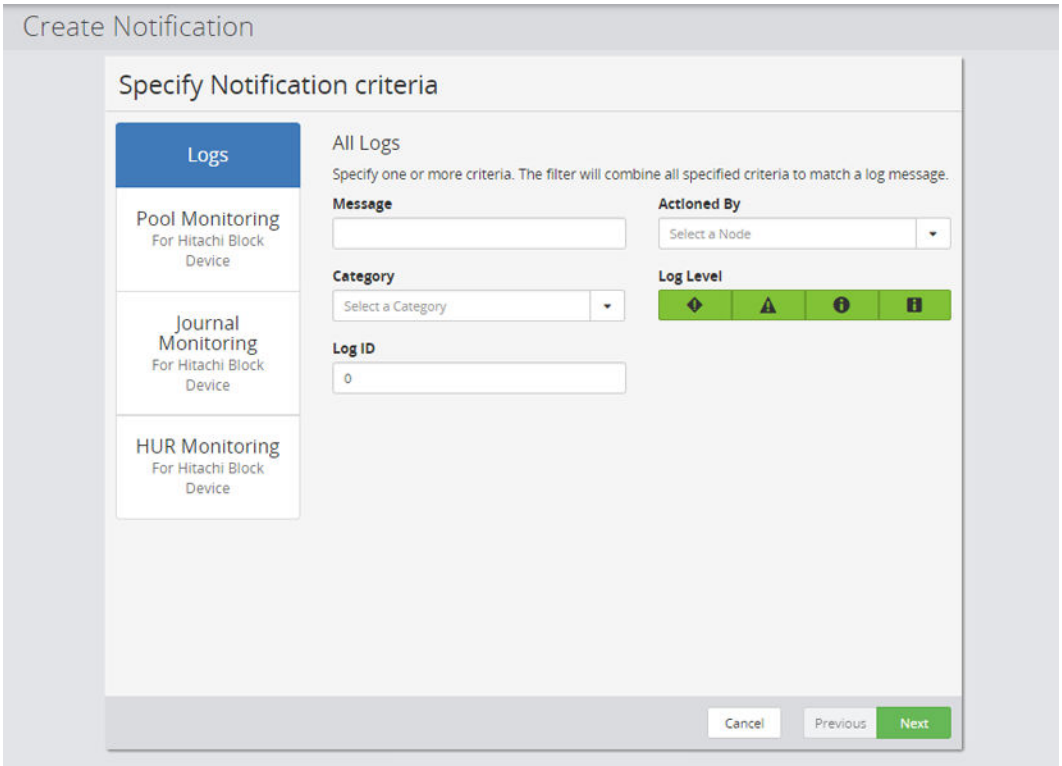


Figure 396 Notification Wizard - Logs

Control	Description
Message	Enter part or all of the log message that will trigger the notification if seen in the logs.
Category	Select the category of the log entry that will trigger the notification.
Log ID	If the log entry has a Log ID then it can be specified here.
Actioned By	Select the Actioned By node name from the log entry that will trigger the notification.
Log Level	Select the level of the log entry that will trigger the notification. Logs of the specified Level and above will be matched for notification. Refer to Logs Inventory (on page 509) for a description of Log Levels.

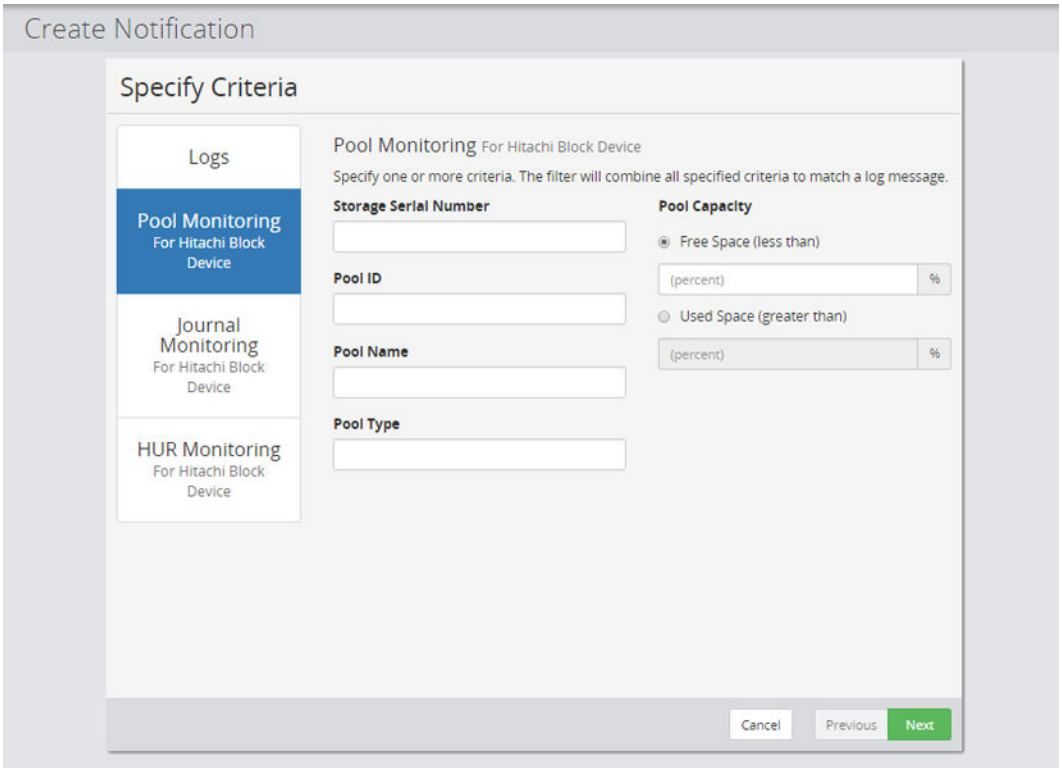


Figure 397 Notification Wizard - Hitachi Block Pool Monitoring

Control	Description
Storage Serial Number	Enter the Serial Number of the storage device for which the notification will be generated.
Pool ID	Enter the Pool ID for which the notification will be generated.
Pool Name	Enter the Pool Name for which the notification will be generated.
Pool Type	Enter the Pool Type for which the notification will be generated.
Pool Capacity	Select either of the following two options: ▪ Free Space (less than) ▪ Used Space (greater than) Then enter the percentage for which the notification will be generated. **Note:** Pool capacity is only checked for destination pools of active replications. Checking is performed every 10 minutes.

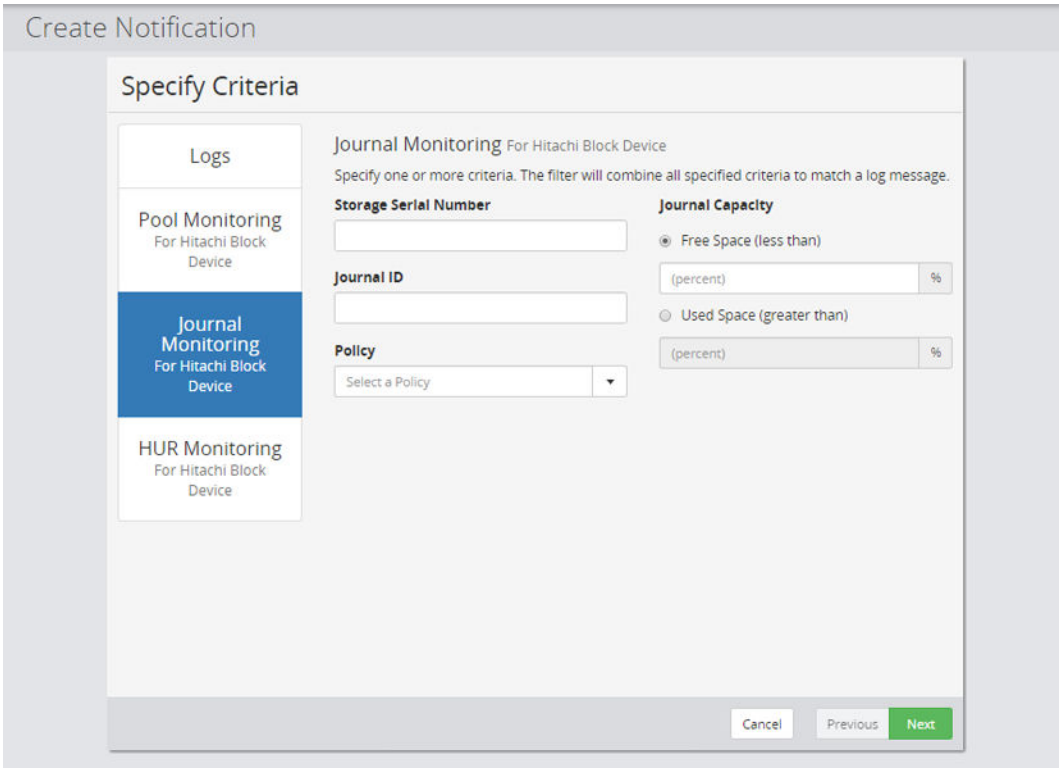


Figure 398 Notification Wizard - Hitachi Block Journal Monitoring

Control	Description
Storage Serial Number	Enter the Serial Number of the storage device for which the notification will be generated.
Journal ID	Enter the Journal ID for which the notification will be generated.
Policy	Select the Hitachi Block policy name for which the notification will be generated.
Journal Capacity	Select either of the following two options: ▪ Free Space (less than) ▪ Used Space (greater than) Then enter the percentage for which the notification will be generated.

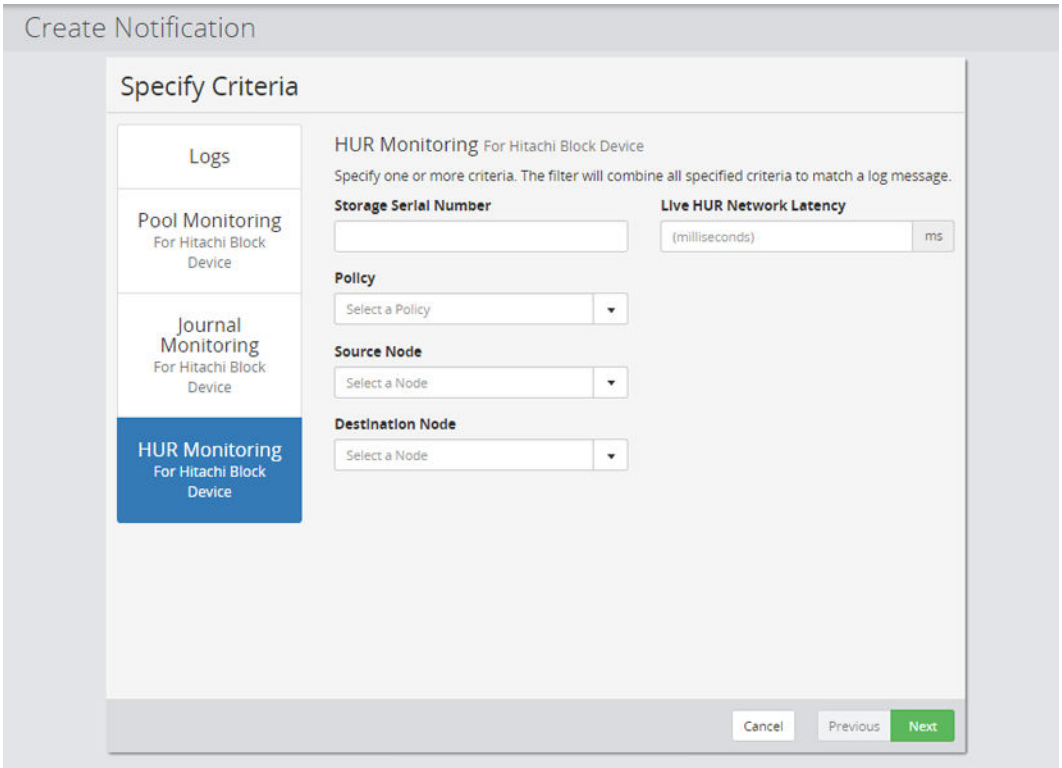


Figure 399 Notification Wizard - Hitachi Block UR Monitoring

Control	Description
Storage Serial Number	Enter the Serial Number of the storage device for which the notification will be generated.
Policy	Select the Block UR policy name for which the notification will be generated.
Source Node	Select the Source Node for which the notification will be generated.
Destination Node	Select the Destination Node for which the notification will be generated.
Live UR Network Latency	Enter the Network Latency above which the notification will be generated. A single notification is generated when this threshold is exceeded. The measurement must fall back below this threshold before another notification can be generated.

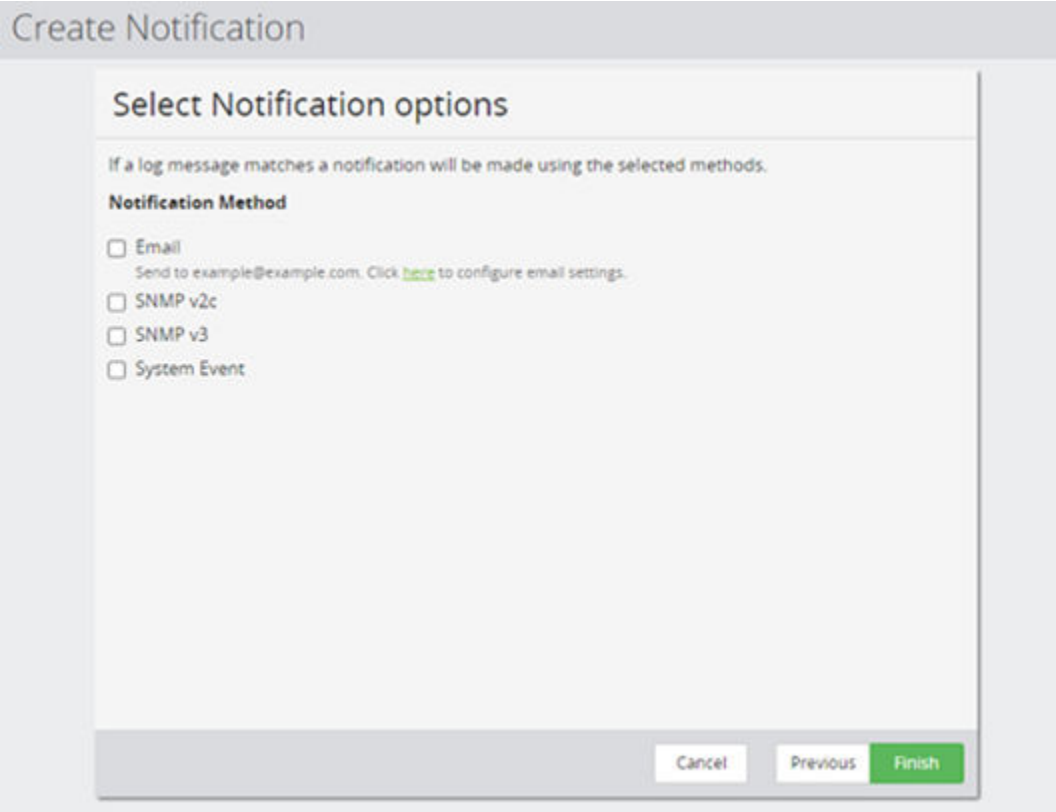


Figure 400 Log notification wizard – Select notification options

Control	Description
Notification Method	This defines how the event is notified to users: ▪ Email – An email will be sent to a configured email account whenever a log meets the defined criteria. If no email configuration is detected. Click the link to configure the email accounts from which notifications will be sent and by whom they will be received via the Email Notifications Settings Wizard (on page 660) ▪ SNMP v2c – An SNMP Version 2c message will be sent to the Management Station. Currently there is no user interface for setting up the SNMP parameters. These parameters must be configured by manually editing the configuration file `[INSTALL_DIR]\db\config\notification\SNMPv2c.cfg`. This file contains the command line for executing the notification event handler `nehsnmp.exe` where the parameter values are defined. Type `nehsnmp.exe -h` at a command prompt to obtain a detailed list of the parameters and allowed values. ▪ SNMP v3 – As for SNMP v2c above except a Version 3 message is sent. The parameters for the Version 3 message are different and therefore a separate configuration `[INSTALL_DIR]\db\config\notification\SNMPv3.cfg` must be edited. Please also refer to the note below for details on configuring the MIB file. ▪ System Event – This will create a system event entry on the master node when a matching log is found. The system event log can be found in the Application log section within the **Event Viewer** on a Windows system, and in `/var/log/messages` on a Linux system. If you create a System Event notification with a User condition then the user name must be prefixed with `//BUILTIN/`. ▪ Custom handlers – It is possible for advanced users to write their own custom handlers to add to the Notification Method list. Contact customer support for details.  **Note:** After changing configuration files, the hub service must be restarted for the changes to take effect. Open a command prompt in the **`\bin`** directory on the Master node type **`diagdata --stop hub`**, wait for the service to stop, then type **`diagdata --start`**

Note: After configuring `SNMPv2c.cfg` or `SNMPv3.cfg`, integrate the Management Information Base (MIB) file `HITACHI-Protector-MIB.txt` with the network management software. The MIB file describes the format of the SNMP Traps sent by Protector. The actual integration steps may vary depending on your network management software. For details, refer to the relevant network management software documentation. Navigate to the `[INSTALL DIR]\Hitachi\Protector\db\config\mibs` directory on the Master node. Locate and copy the MIB file to the MIB files directory (as defined in the relevant product documentation) on the host machine running the network management software.

Policies UI Reference

This section describes the Policies UI, accessed via the Navigation Sidebar (on page 291).

For further information, refer to:

- Policy Concepts (on page 90)
- Policy Tasks (on page 272)

Policies Inventory

This inventory lists all defined Policies. Each policy consists of a set of data classifications defining what to backup and a set of operations defining how to backup that data.

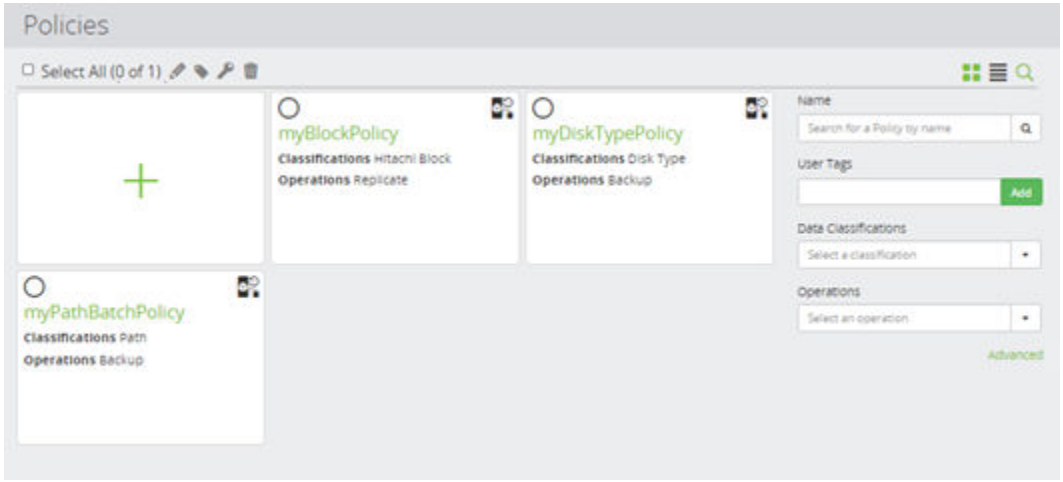


Figure 401 Policies Inventory

Control	Description
Edit	Edits an existing policy in the inventory. The Policy Wizard (on page 671) is launched to enable the policy's attributes to be changed.

Control	Description
 Tag	Modifies the tags of an existing object from the either the inventory screen or the details screen of the object.
 View Permissions	View an existing policy's access permissions. The Access Control Permissions Inventory (on page 349) is launched to enable the policy's access permissions to be read/view only.
 Delete	Enabled only when one or more Policies are selected. Deletes the selected item from the inventory. The associated Classifications and Operations are also deleted.
 Add	Creates a new Policy. The Policy Wizard (on page 671) is launched to guide you through the process.
 Policy	Any number of user defined Policies can be created. These are displayed in the inventory. The Policy Details (on page 737) is displayed to enable the policy to be viewed and edited.
Filter on Policy Name	Filters the displayed results based on the policy name.
Filter on Data Classification	Filters the displayed results based on classification type.
Filter on Operations	Filters the displayed results based on operation type

Policy Wizard

This wizard is launched when a new Policy is added to the Policies Inventory.

A Policy consists of the following parts:

- One or more Data Classification items that define what type of data is to be backed up.
  - Path and Disk Type classifications can also have additional Age, File Type and OS Type Filters applied to them.
- One or more Operation items that define what action is to be performed on that data.

Note:

- If the policy is in use in the current data flow definition, any modifications made to the policy only take effect following recompilation and reactivation of the affected data flows. If data operation items have been added or deleted, it may be necessary to first modify the data flow definition to take account of these changes before it is recompiled and reactivated.
- If a policy using the Replicate, CDP or Backup data operations has any part of its data classification changed, then all destination nodes for that policy need to be re-synchronized to include the change in scope of the policy.

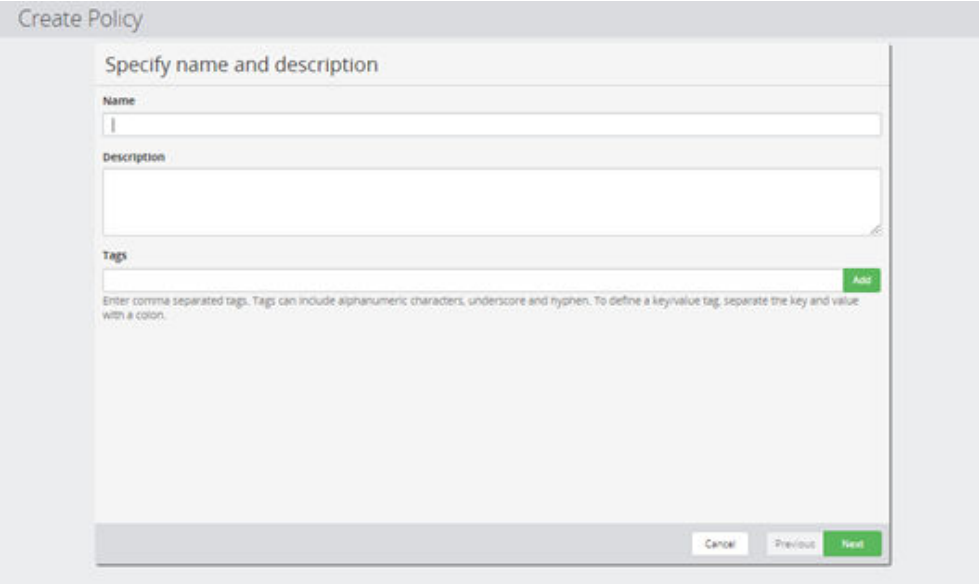


Figure 402 Policy Wizard - Specify name and description

Control	Description
Name	Enter a name for the policy.
Description	Optional. Enter a short description of the policy.
Tags	Add the tags to be associated with the object being created.

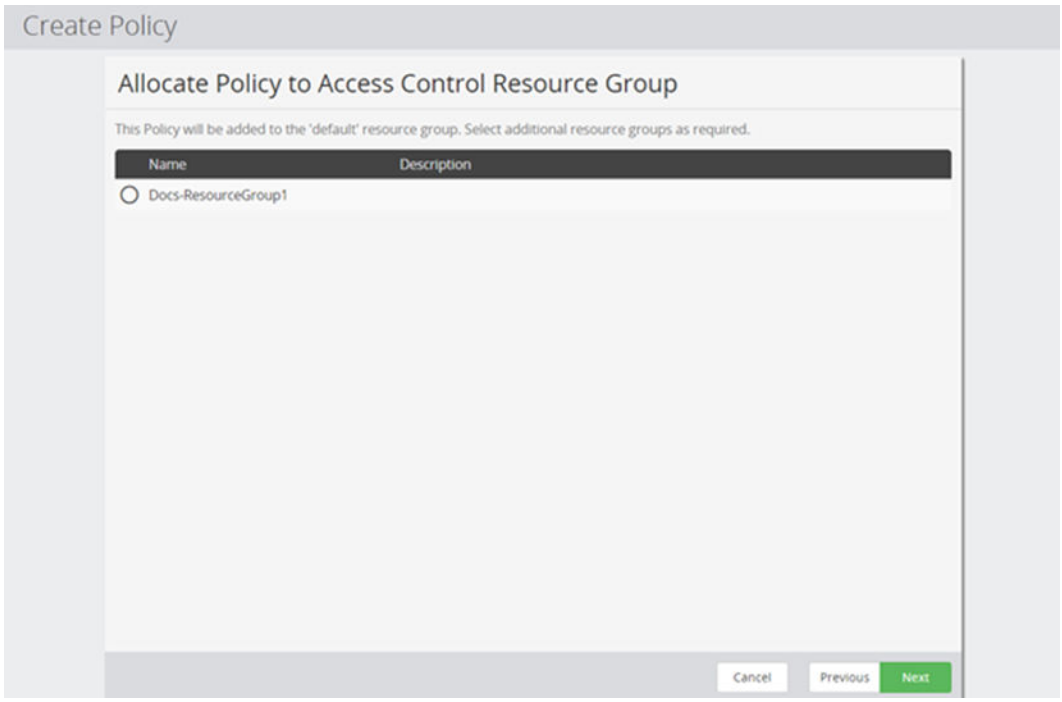


Figure 403 Policy Wizard - Allocate policy to Resource Group

Control	Description
Resource Groups	It allows the user to view the access permissions for those items granted to specific users and groups. **Note:** A single policy can be assigned to multiple resource groups.

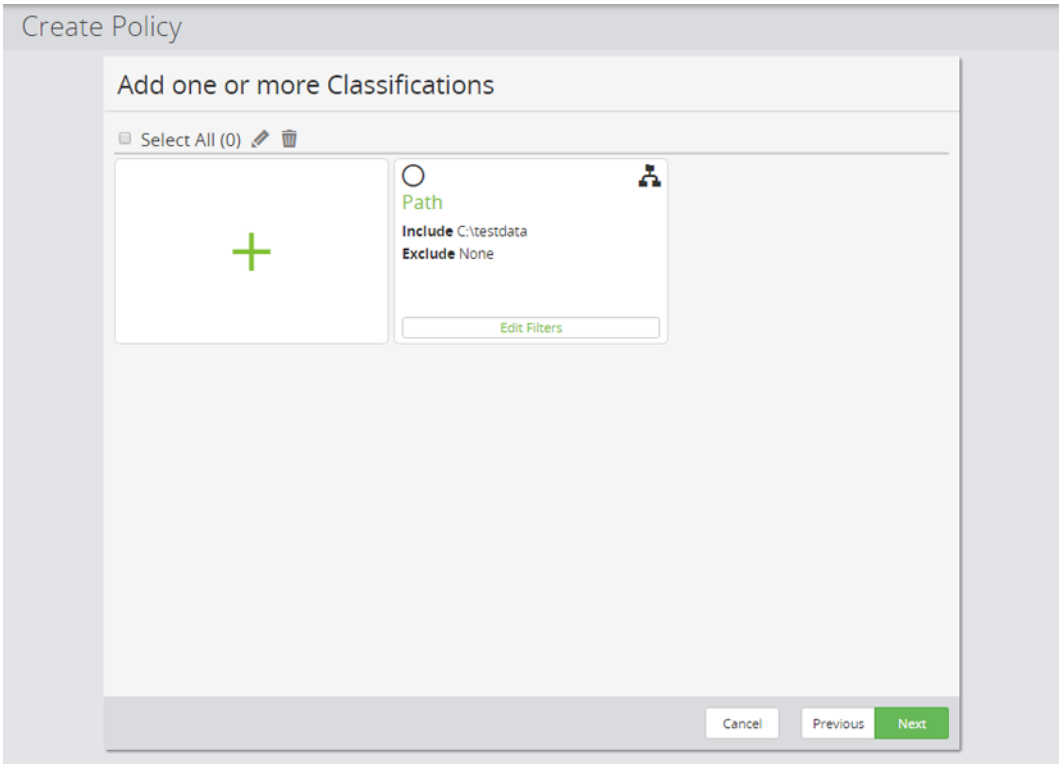


Figure 404 Policy Wizard - Classifications Inventory

Control	Description
Edit	Edits an existing classification in the inventory. The relevant Classification wizard is launched to enable the classification's attributes to be changed.
Delete	Deletes the selected item from the inventory.
Add	Creates a new Classification. The Classification Selection page (shown above) of this wizard is reopened to enable other classifications to be added to the policy.
Classification(s)	Any number of Classifications can be created. These are displayed in the inventory. The appropriate classification wizard is displayed to enable the classification to be viewed and edited.
Edit Filters	Launches the Classification Filters Wizard (on page 727) to enable you to add Filters to the classification.

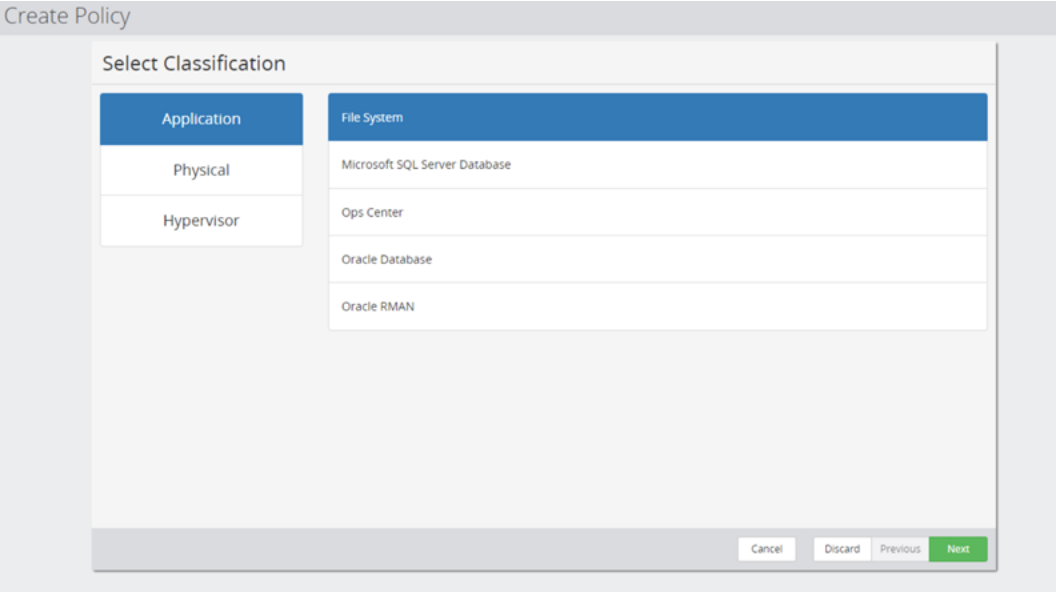


Figure 405 Policy Wizard - Select Classification (Application)

Each Application Classification Type defines a specific set of criteria for selecting data related to a specific application. For example, the Oracle type can be used to select data related to a specific Oracle database.

Control	Description
File System Application	Adds a File System Application classification to the policy. The File System Classification Wizard (on page 679) will be launched to guide you through the process.
Microsoft SQL Server Database	Adds a Microsoft SQL Server Database classification to the policy. The Microsoft SQL Server Classification Wizard (on page 681) will be launched to guide you through the process.
Ops Center Application	Adds an Ops Center Application classification to the policy. The Ops Center Classification Wizard (on page 685) will be launched to guide you through the process.
Oracle Database	Adds an Oracle Database classification to the policy. The Oracle Database Classification Wizard (on page 686) will be launched to guide you through the process.
Oracle RMAN Database	Adds an Oracle RMAN Database classification to the policy. The Oracle RMAN Classification Wizard (on page 688) will be launched to guide you through the process.

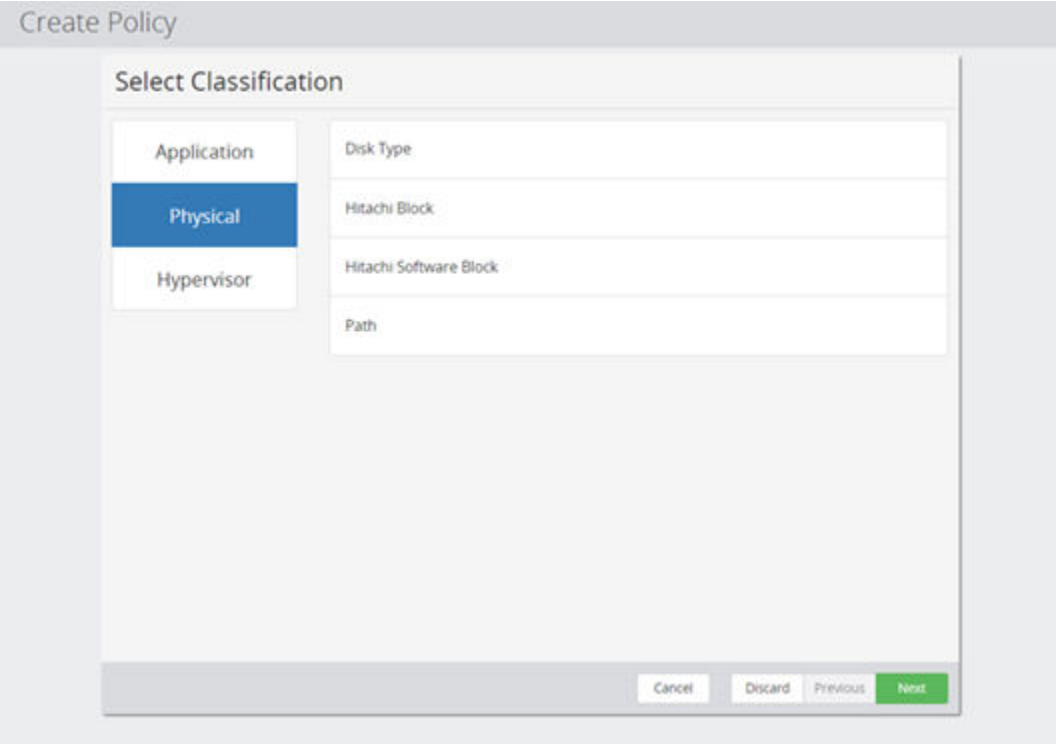


Figure 406 Policy Wizard - Select Classification (Physical)

Each Physical Classification type defines a specific set of criteria for selecting data on a node's file system. For example, the Path type can be used to select data by its folder location.

Control	Description
Disk Type	Adds a Storage Type classification to the policy. The Disk Type Classification Wizard (on page 708) will be launched to guide you through the process.
Hitachi Block	Adds an Hitachi Block classification to the policy. The Hitachi Block Classification Wizard (on page 709) will be launched to guide you through the process.
Hitachi Software Block	Adds an Hitachi Software Block classification to the policy. The Hitachi Software Block Classification Wizard (on page 712) will be launched to guide you through the process.
Path	Adds a Path classification to the policy. The Path Classification Wizard (on page 713) will be launched to guide you through the process.

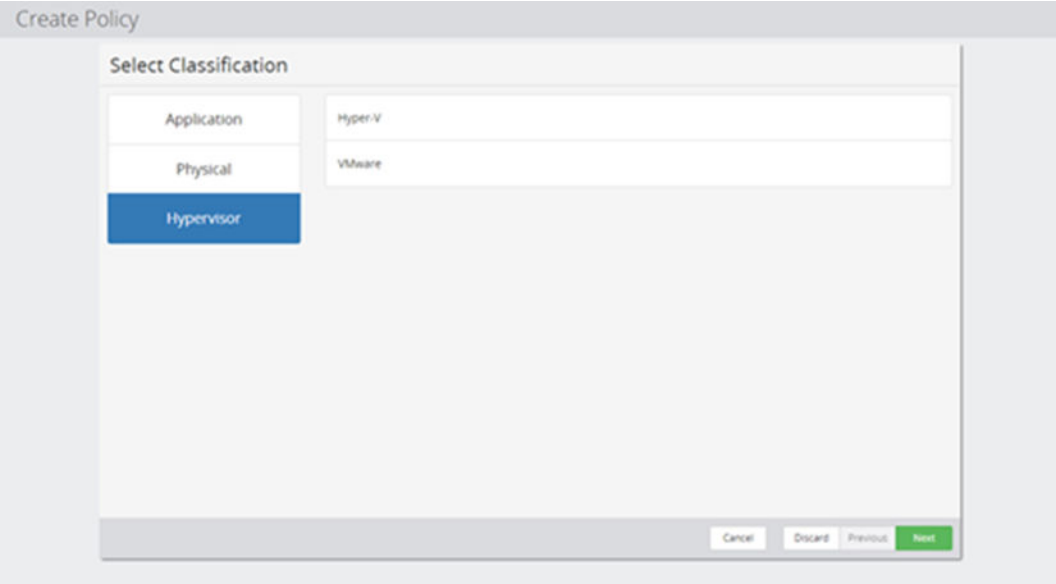


Figure 407 Policy Wizard - Select Classification (Hypervisor)

Each Hypervisor Classification Type defines a specific set of criteria for selecting data related to a specific hypervisor. For example, the VMware type can be used to select virtual machines on vCenter.

Control	Description
Hyper-V	Addsa Hyper-V classification to the policy. The Hyper-V Classification Wizard (on page 692) will be launched to guide you through the process.
VMware	Adds a VMware classification to the policy. The VMware Classification Wizard (on page 700) will be launched to guide you through the process.

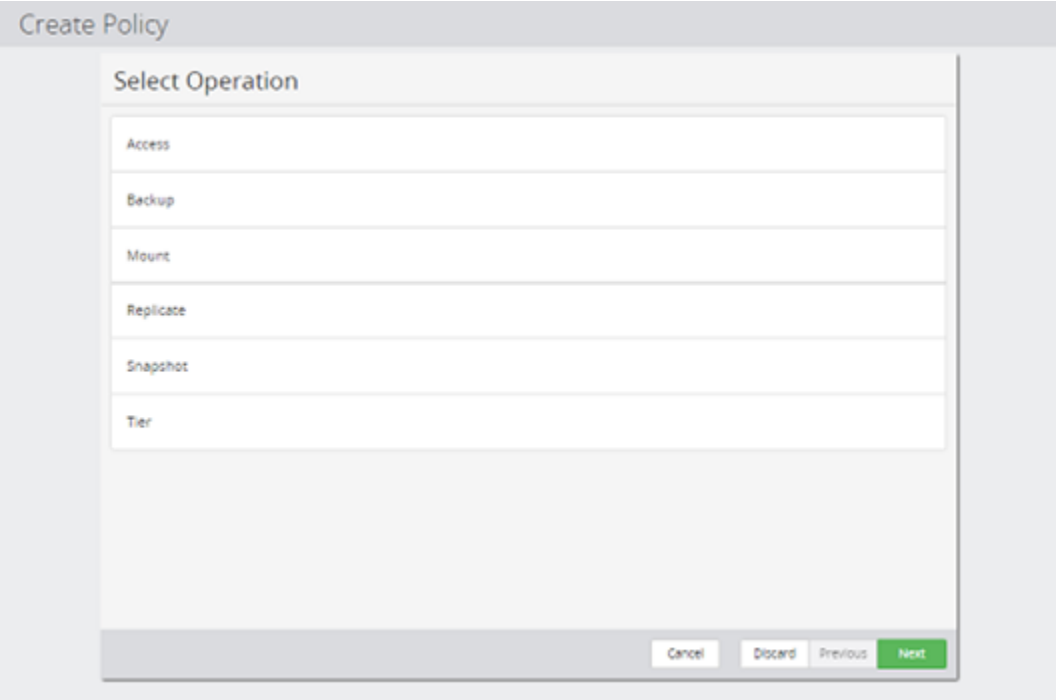


Figure 408 Policy Wizard - Select Operation

Each Operation Type represents an action that can be performed on data that matches the classification. For example, Snapshot can be used to create a point in time copy of the data specified by a Hitachi Block Classification.

Control	Description
Backup	Adds a backup operation to the policy. The Backup Operation Wizard (on page 718) will be launched to guide you through the process. **Note:** This operation cannot be applied to hardware storage devices.
Mount	Adds a mount operation to the policy. The Mount Operation Wizard (on page 721) will be launched to guide you through the process.
Replicate	Adds a replication operation to the policy. The Replicate Operation Wizard (on page 722) will be launched to guide you through the process.
Snapshot	Adds a snapshot operation to the policy. The Snapshot Operation Wizard (on page 724) will be launched to guide you through the process.
Tier	Adds a tier operation to the policy. The Tier Operation Wizard (on page 726) will be launched to guide you through the process.

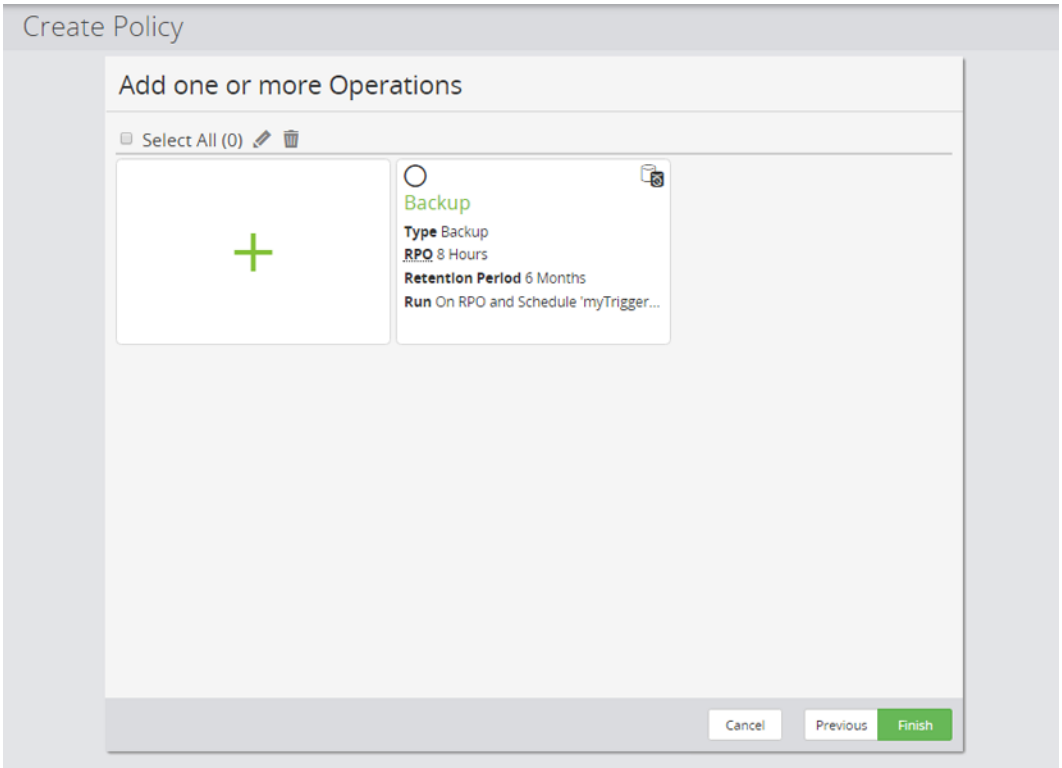


Figure 409 Policy Wizard - Add one or more Operations

Control	Description
Edit	Edits an existing operation in the inventory. The relevant Operation wizard is launched to enable the operation's attributes to be changed.
Delete	Enabled only when one or more Operations is selected. Deletes the selected item from the inventory.
Add	Creates a new Operation. The Operation Selection page of this wizard is reopened to enable other operations to be added to the policy.
Operation(s)	Launches the appropriate operations wizard to enable the operation to be viewed and edited.

File System Classification Wizard

This wizard is launched when a new File System Classification is added to the policy.

The File System Classification specifies which paths or drives of a a given File System node should be included in a backup.

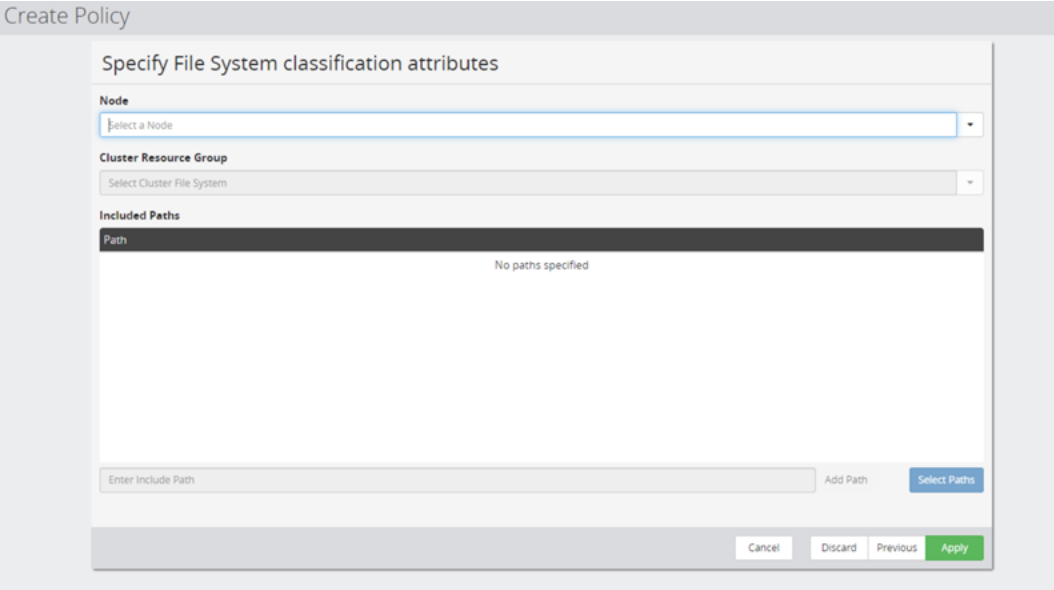


Figure 410 File System Wizard - Specify File System classification attributes

Control	Description
Node	Select the File System node you want to protect with this classification.
Cluster Resource Group	Select the Resource Group that holds the resources you want to protect. If you are using a standalone system, you will only have the "Standalone" option available. This option supports all the resources, paths, and disks in the filesystem node. In case of a Windows Server Failover Cluster, select the cluster resource group that contains the disks or paths you wish to protect. **Note:** A classification/policy can only protect a single resource group. To protect resources from multiple resource groups, use multiple policies and classifications.
Included Paths	Lists the paths and drives that are currently included in the classification. To remove a path from the list, simply click on the "X" next to it.
Add Path	Add the path provided in the “Enter Include Path” field to the list of included paths. **Note:** It is possible to add any path, but any paths that do not exist or are outside the resource group's scope will be skipped during the backup process.

Control	Description
Select Paths	Opens a Path Dialog (on page 350), which allows you to choose paths and drives that should be added to the list of included paths. The selection is limited to paths within the cluster resource group.

Microsoft SQL Server Classification Wizard

This wizard is launched when a new Microsoft SQL Server Database classification is added to the policy.

The Microsoft SQL Server Database classification conveniently specifies which should be included in a backup. Refer to About Microsoft SQL Server Policy Classifications (on page 97) for details on how this classification works.

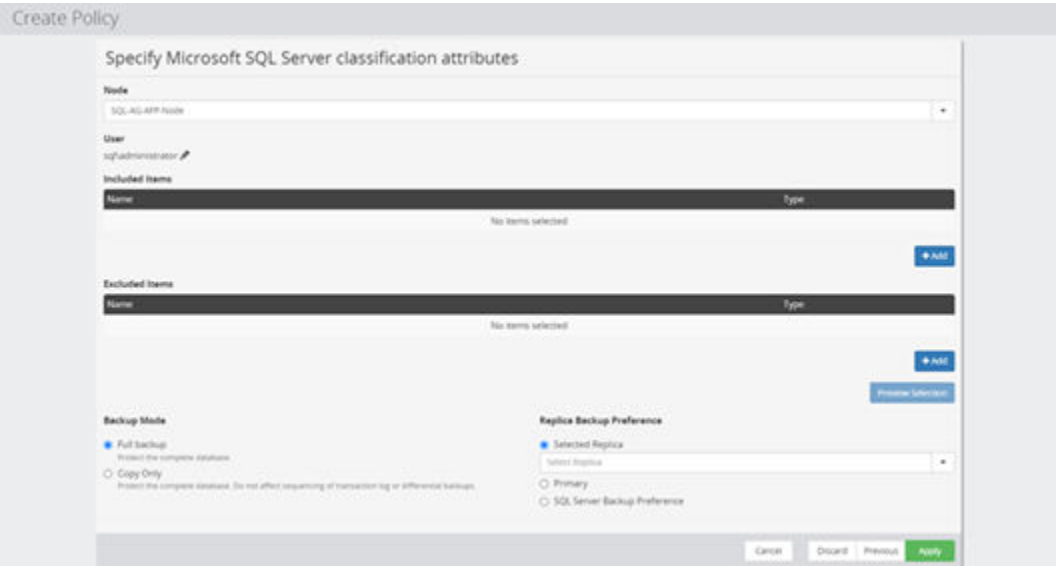


Figure 411 Microsoft SQL Server Wizard - Specify Microsoft SQL Server classification attributes

Control	Description
Node	Select the Microsoft SQL Server node you want to protect with this classification.
Instance	Select the SQL Server instance, hosting the databases you want to protect. This control is not available when an SQL Server Availability Group Node is selected. **Note:** Instance is visible only in case of Cluster or Standalone Nodes.

Control	Description
User	List the user which will be used to protect the instance. Editing the credentials will open the Microsoft SQL Server Instance Credentials Dialog (on page 684) SQL Server Instance Credentials Dialog.
Refresh Instances	Refreshes the instances and Databases in the Node. **Note:** Refresh Instance button is visible only in case of Cluster or Standalone Nodes.
Included Items	Lists the databases which will be included in the backup.
Add	Opens the Microsoft SQL Server Database Selection Wizard (on page 683), which allows databases to be added to the Included Items list above.
Excluded Items	Lists the databases that will be excluded from the backup policy.
Add	Opens the Microsoft SQL Server Database Selection Wizard (on page 683) which allows databases to be added to the Excluded Items list above.
Preview Selection	Opens the Microsoft SQL Server Classification Wizard (on page 681), that will preview which databases will be included if the classification is applied to a selected Microsoft SQL node.
Backup Mode	Select which backup type is desired for the selected databases: ▪ **Full backup** Create a point in time copy of the complete database. ▪ **Copy Only** Create a point in time copy of the complete database. Does not affect sequencing of differential backups or transaction logs.
Replica Backup Preference (Availability Group only)	**Selected Replica** Select a specific replica (instance) as a source for the backup, irrespective if it is currently acting as a primary or a secondary replica. This option ensures that backups are always created on the same storage. **Primary** The primary replica at the time of the backup will be used to protect the databases of the availability group. Ensure that all possible storage arrays are configured on the data flow, as the location of the primary replica may change over time. **SQL Server preference**

Control	Description
	The SQL Server's Availability Group preference will be honored and used as the source for the backup. Ensure that all possible storage arrays are configured on the data flow if the SQL Server backup preference allows for more than one replica to become the preferred backup source.

Microsoft SQL Server Database Selection Wizard

This wizard is displayed when the user chooses to include database in a Microsoft SQL Server classification.

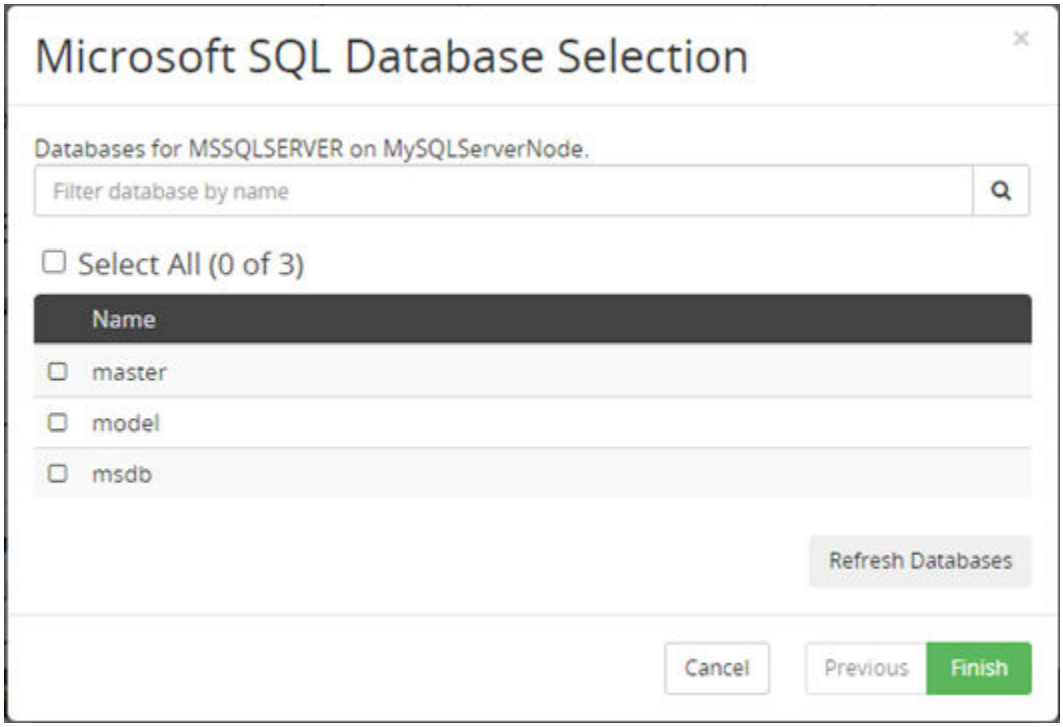


Figure 412 Microsoft SQL Database Selection

Control	Description
Search	Enter a part of a database name and confirm to filter the list of known databases
Database List	Select one or more databases **Note:** Do not select databases which are temporarily mounted by Protector. Including mounted databases in backups may cause problem during backup or when you try to unmount the database.

Microsoft SQL Server Instance Credentials Dialog

This dialog allows the user to specify which credentials should be used when protecting an instance

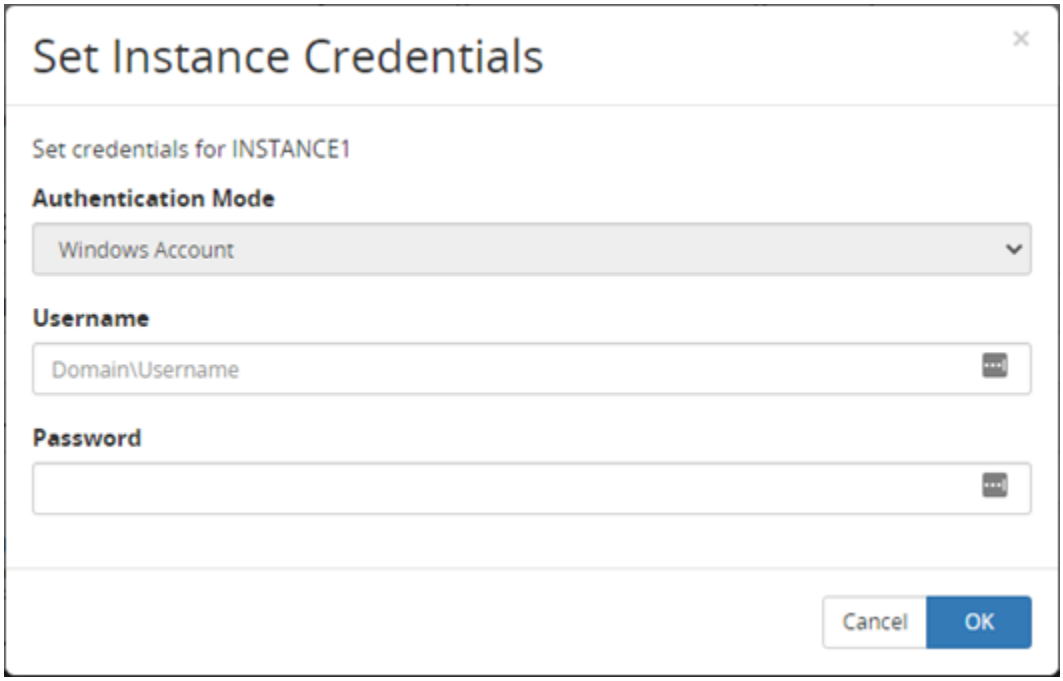


Figure 413 Microsoft SQL Server Instance Credentials Dialog

Control	Description
Authentication Mode	Lists the required account type.
Username	Specify the user which should be used to protect the databases of this instance. Use the format <NETBIOS DOMAIN NAME>\<username>. Refer to Microsoft SQL Server User Privileges (on page 684) for more details.
Password	Password for the user specified above

Microsoft SQL Server User Privileges

To create, mount or revert block-based Microsoft SQL Server backups Protector requires credentials with the necessary permissions.

The user account must meet the following requirements:

- Windows account
- Member of the sysadmin role for the SQL Server instance

Ops Center Classification Wizard

This wizard is launched when a new Ops Center Application classification is added to a Policy. The Ops Center classification is used to define which applications are to be protected.

Note: Ops Center Classification Wizard allows to backup only the master node at present

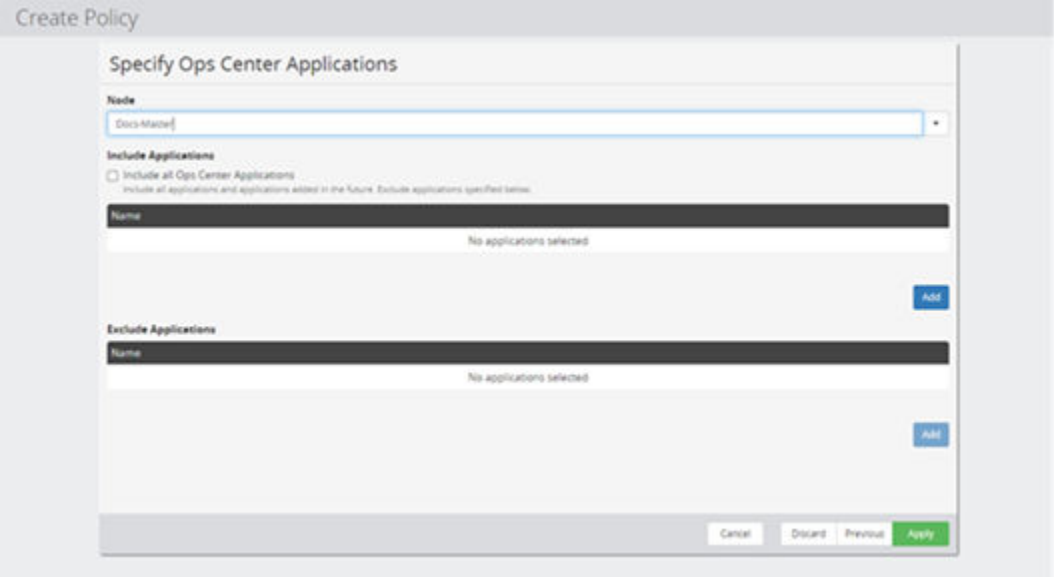


Figure 414 Ops Center Wizard - Specify Ops Center Applications attributes

Control	Description
Included Items	Lists the Ops Center resources that will be included in the backup policy.
Add	Opens the Ops Center Application Selection Wizard (on page 685) to enable Ops Center resources to be added to the include list above. **Note:** Only Ops Center Protector Master node can be backed up currently.
Excluded Items	Lists the Ops Center resources that will be excluded from the backup policy.
Remove	Each row has a remove button at the end of the row, the selected Ops Center resource is removed from the include/exclude list.
Add	Opens the Ops Center Application Selection Wizard (on page 685) to enable VMware resources to be added to the exclude list above.

Ops Center Application Selection Wizard

This wizard is displayed when the user includes or excludes Ops Center resources in a policy.



Figure 415 Ops Center Applications Selection for Inclusion/Exclusion Wizard

Control	Description
Available Applications	Select the Ops Center application you want to protect with this classification.

Oracle Database Classification Wizard

This wizard is launched when a new Oracle Database classification is added to a Policy.

The Oracle Database classification is used to define which databases are to be protected.

Note: When used in combination with a storage hardware backup operation, Protector will discover the underlying hardware paths at runtime. For Hitachi Block storage hardware based backups, all the paths must exist on the same block hardware device.

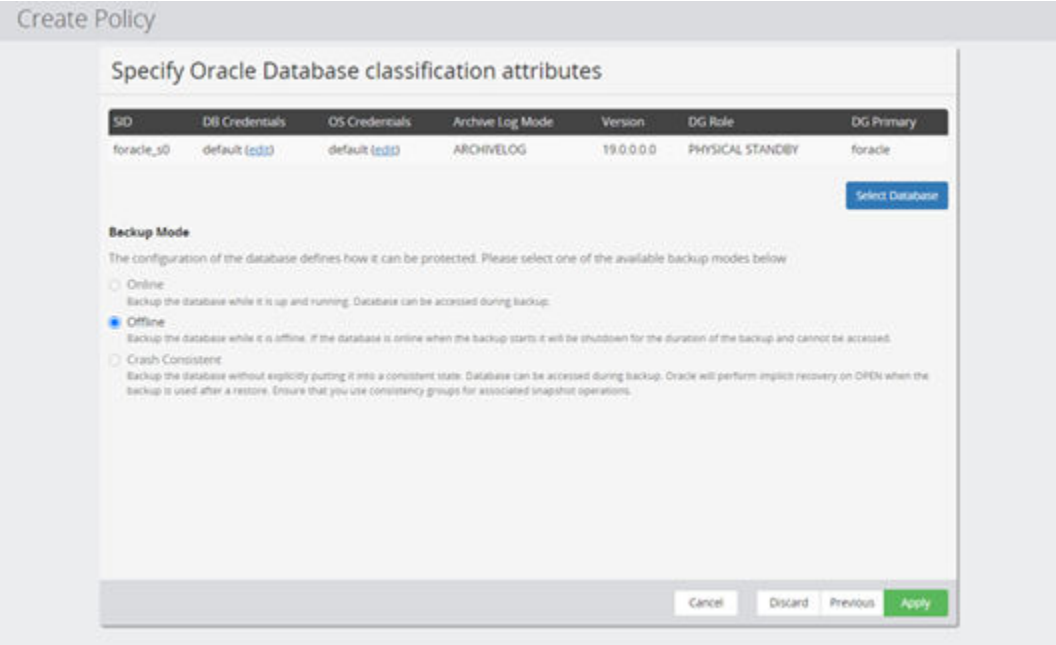


Figure 416 Oracle Database Wizard - Specify Oracle Database classification attributes

Control	Description
Databases	Lists the currently selected databases. Only the databases listed are backed up. ▪ For block based policies - Databases are discovered only once, when the policy is defined. ▪ For host based policies - This classification is not currently supported.
Select Database	Click to open the Select Database dialog shown below.
Backup Mode	Select one of the following: ▪ Online – Available only if Oracle Database is running in `ARCHIVELOG` mode. Oracle is briefly quiesced while an application consistent snapshot is made. ▪ Offline – Available always. Oracle is taken offline while an application consistent snapshot is made. ▪ Crash Consistent – Available only in Version 12c and if Oracle Database is running in `ARCHIVELOG` mode. Oracle remains running while a crash consistent snapshot is taken.

Control	Description
	Note: ▪ If the DataGuard Role is PHYSICAL STANDBY, backup mode is auto selected as Offline. ▪ If the Backup Mode is set to Online, ensure that Quiesce configured applications before backup is selected in the **Snapshot/Replicate Operation Attributes Wizard**. ▪ If the Backup Mode is set to Crash Consistent, ensure that the Use consistency group option is selected in the **Snapshot/Replicate Operation Properties Dialog** on the data flow.



Figure 417 Oracle Database Wizard - Select Oracle Databases Dialog

Control	Description
Oracle Node	Select a node representing the Oracle server hosting the database(s) to be selected for backup.
Filter database by name	Filters the databases list below to show only those entries that contain the filter string.
Databases	Select the database(s) to be backed up from the list.
Refresh	Click this button to refresh the cached details and clear the name filter.

Oracle RMAN Classification Wizard

This wizard is launched when a new Oracle RMAN classification is added to a Policy.

The Oracle RMAN classification allows to conveniently specify, which databases can or cannot access Oracle RMAN data using the access operation.

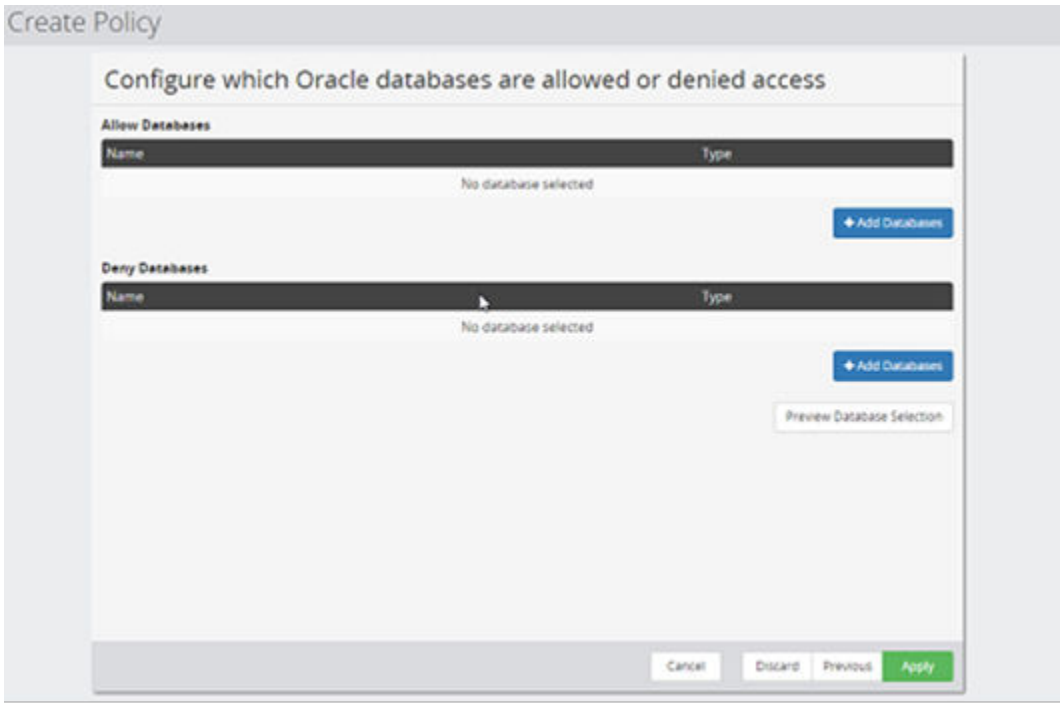


Figure 418 Oracle RMAN Wizard - Specify Oracle Databases (Allowed / Denied Access)

Control	Description
Allow Databases	Lists the databases that will be allowed access for Oracle RMAN access operations.
Add Databases	Opens the Oracle RMAN Database Selection Wizard (on page 690) to enable databases to be added to the **Allow Databases** list above.
Deny Databases	Lists the databases that will not be allowed access for Oracle RMAN access operations.
Add Databases	Opens the Oracle RMAN Database Selection Wizard (on page 690) to enable databases to be added to the "Deny Databases" list above.
Preview Database Selection	Click this button to preview the which databases are allowed access for an existing Oracle Database node.

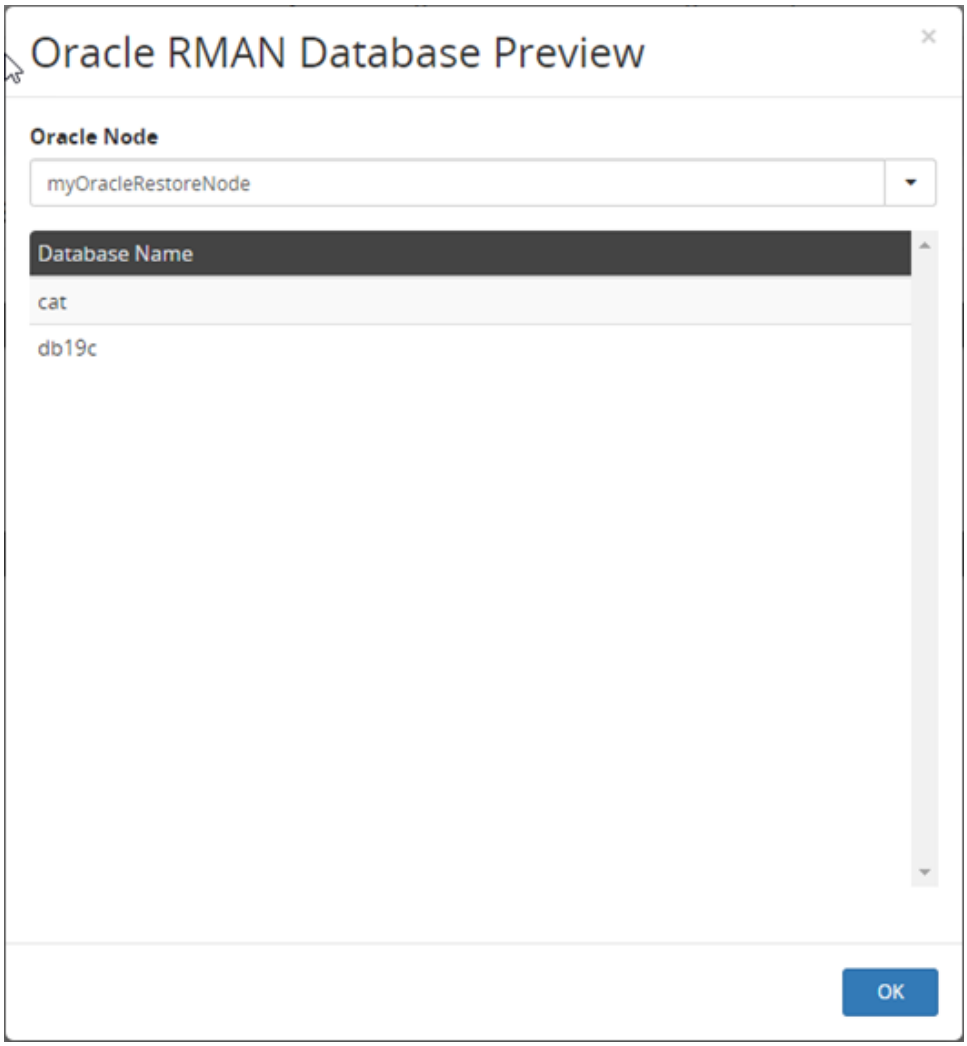


Figure 419 Oracle RMAN Database Preview

Control	Description
Oracle Node	Select a node representing the Oracle setup, which should be previewed.
Database List	List all databases on this node which would be allowed access with the defined classification.

Oracle RMAN Database Selection Wizard

This wizard is launched when a user adds entries to the list of allowed or denied databases in an Oracle RMAN classification.

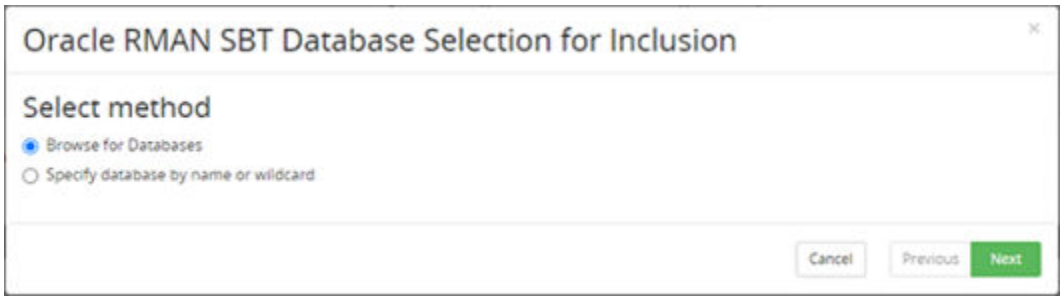


Figure 420 Oracle Database Selection – Select method

Control	Description
Browse for databases	Select this option to browse an existing Oracle node for databases. See **Oracle Database Selection – Browse by** below.
Specify databases by name or wildcard	Select this option specify a database by name pattern match. See **Oracle Database Selection – Specify name or wildcard** below.

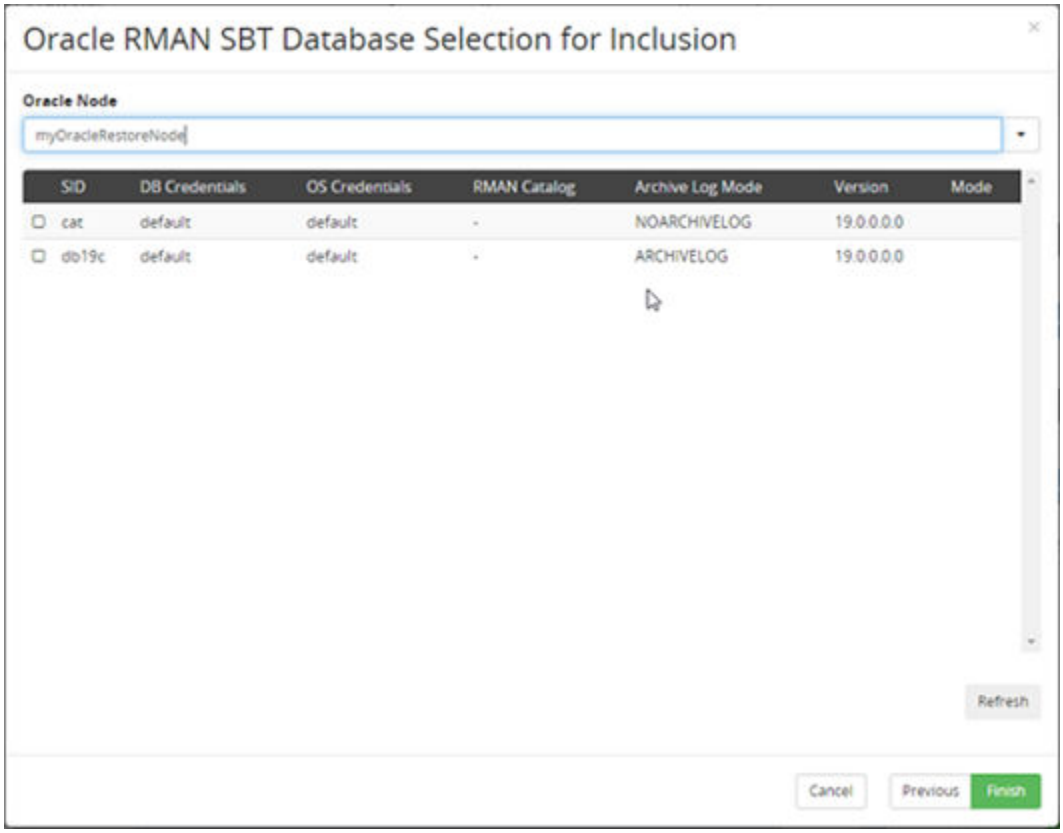


Figure 421 Oracle Database Selection – Browse by

Control	Description
Oracle Node	Select an Oracle database application node to browse for databases.
Database List	Lists the databases which exist on the selected node. You can select one or more databases.
Refresh	Refreshes the list of databases for the selected node. This operation may take a few minutes.

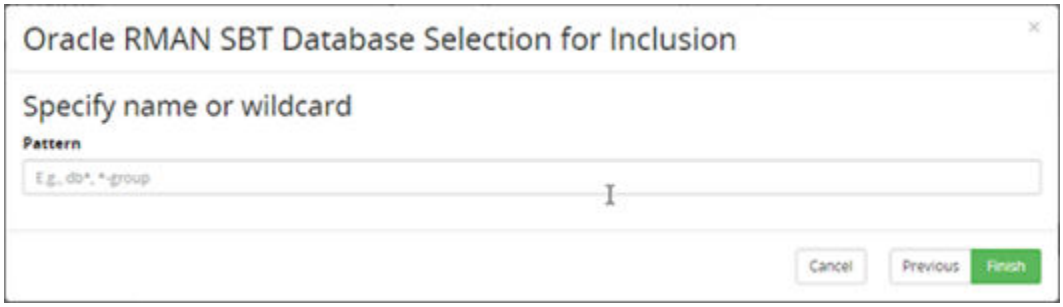


Figure 422 Oracle Database Selection – Specify name or wildcard

Control	Description
Pattern	Enter a case insensitive pattern that will be used to match the database by name. The '*' character can be used to match any sequence of characters. E.g.: IH_* would match any database type whose name begins with IH_. **Note:** Protector evaluates the pattern every time Oracle RMAN tries to access the data. If new databases are added later, they will still be allowed or denied access, depending if they match the pattern or not.

Hyper-V Classification Wizard

This wizard is launched when a new Hyper-V classification is added to policy.

The Hyper-V classification is used as a means to conveniently specify the Hyper-V resources which should be included in a backup. Refer to About Hyper-V policy classifications (on page 96) for details about how this classification works with host and block based operations.

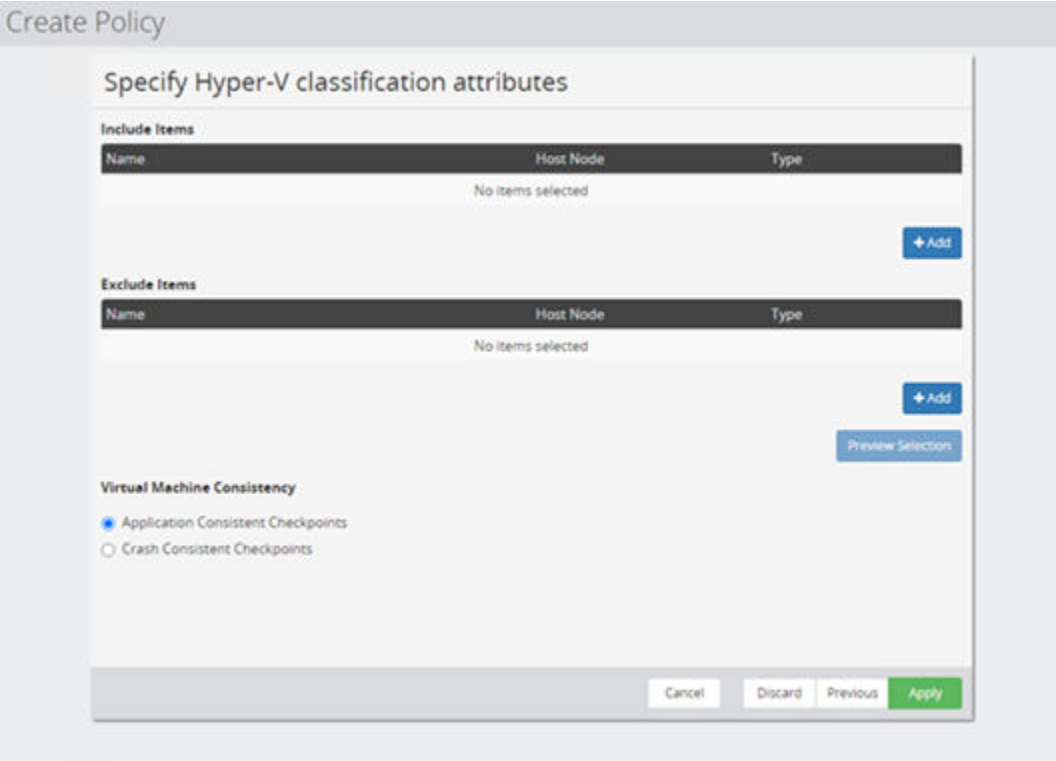


Figure 423 Hyper-V Wizard - Specify Hyper-V classification attributes

Control	Description
Included Items	Lists the Hyper-V resources that will be included in the backup.
Add	Opens the Hyper-V Resource Selection Wizard (on page 694) to enable Hyper-V Resources to be added to the include/exclude items list above.
Excluded Items	Lists the Hyper-V VMs that will be excluded from the backup policy.
Remove	Each row has a remove button at the end of the row. When clicked the selected Hyper-V resource is removed from the include/exclude list.
Preview Selection	Opens the Hyper-V Classification Preview Wizard (on page 694) Hyper-V Classification Preview Wizard, that will preview which VMs will be included if the classification is applied to a selected Hyper-V node.
Virtual Machine Consistency	Select which level of consistency is desired for the virtual machines: ▪ **Application consistent checkpoints** will use Hyper-V integration services to quiesce the data inside the VM. ▪ **Crash consistent checkpoints** will just use the data currently available on the virtual disks.

Hyper-V Classification Preview Wizard

This wizard is displayed when the user previews a Hyper-V policy classification.

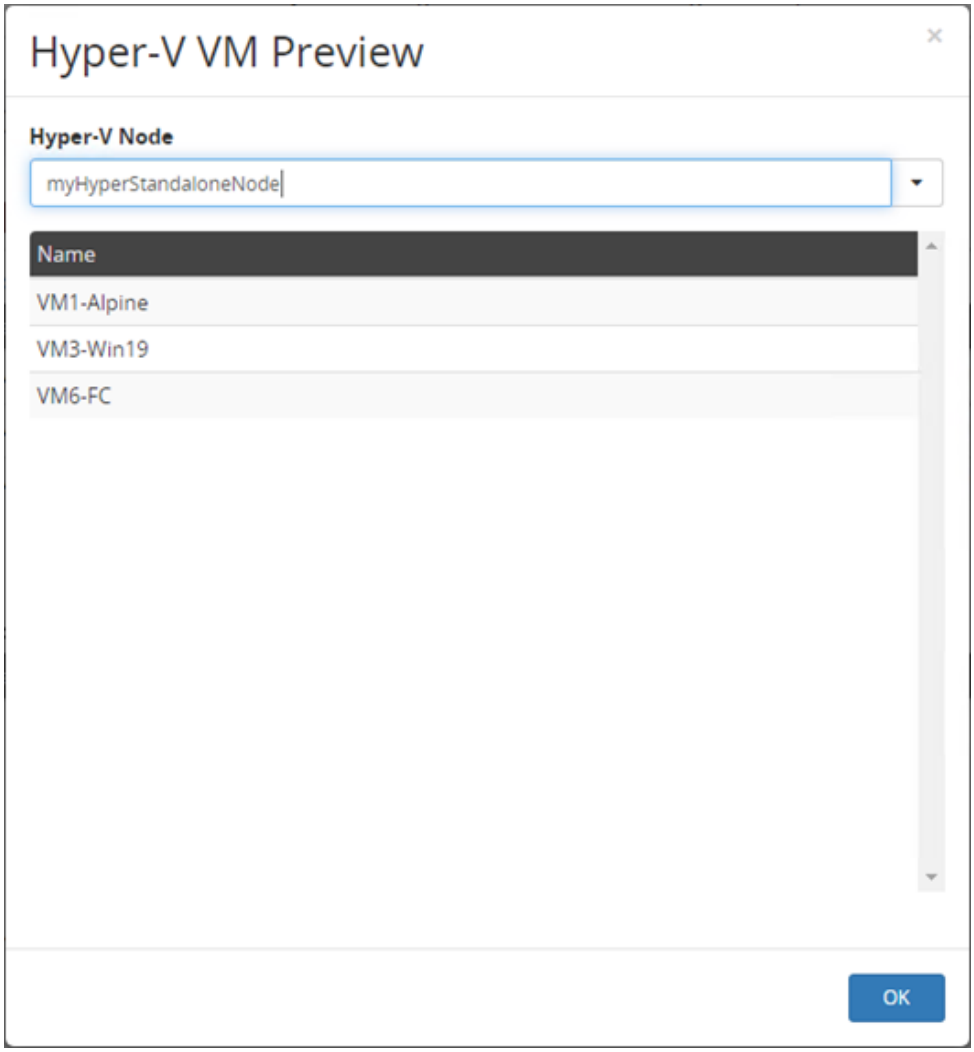


Figure 424 Hyper-V Classification Preview Wizard

Control	Description
Hyper-V Node	Select the Hyper-V app node you want to preview the classification for.
Virtual Machine List	Lists all virtual machines that would be considered for backup, based on the classification.

Hyper-V Resource Selection Wizard

This wizard is displayed when the user includes or excludes Hyper-V resources in a policy.

Caution: Protector tracks Hyper-V VMs via their unique ID. If the id of an explicitly selected VM is changed (e.g. by deleting and restoring the VM) it will not be included in the backup and an error will be logged.

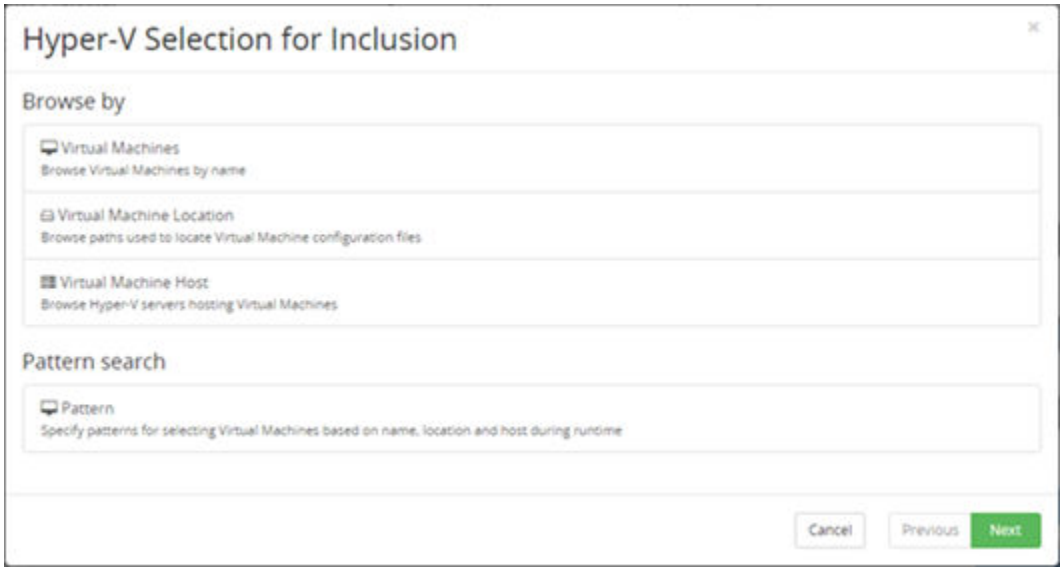


Figure 425 Hyper-V VM Selection for Inclusion

Control	Description
Virtual Machines	Displays flat list of the virtual machines configured on a Hyper-V node. One more VMs can be selected.. See Hyper-V Resource Selection Wizard – Browse Virtual Machines (on page 696) below.
Virtual Machine Locations	Displays a hierarchical view of a Hyper-V node's file system. See Hyper-V Resource Selection Wizard – Browse by Virtual Machine Locations (on page 697) below. **Note:** Protector will select all virtual machines, that have a configuration file under the select path. The list of VMs per path is re-evaluated at the beginning of each backup.
Virtual Machine Host	Displays a hierarchical view of Hyper-V nodes and VMs. It is possible to select one or more nodes of a cluster as well as individual VMs. See Hyper-V Resource Selection Wizard – Browse by Virtual Machine Hosts (on page 698) below. **Note:** If a host is selected Protector will select all virtual machines available on that host. The list of available virtual machines per host is re-evaluated at the beginning of each backup.

Control	Description
Pattern	Select if you want to specify a resource by type and name pattern match. See Hyper-V Resource Selection Wizard – Pattern search (on page 699) below.

Hyper-V Resource Selection Wizard – Browse Virtual Machines

This page of the wizard is displayed when the browse by virtual machines option is selected in the initial wizard page above.

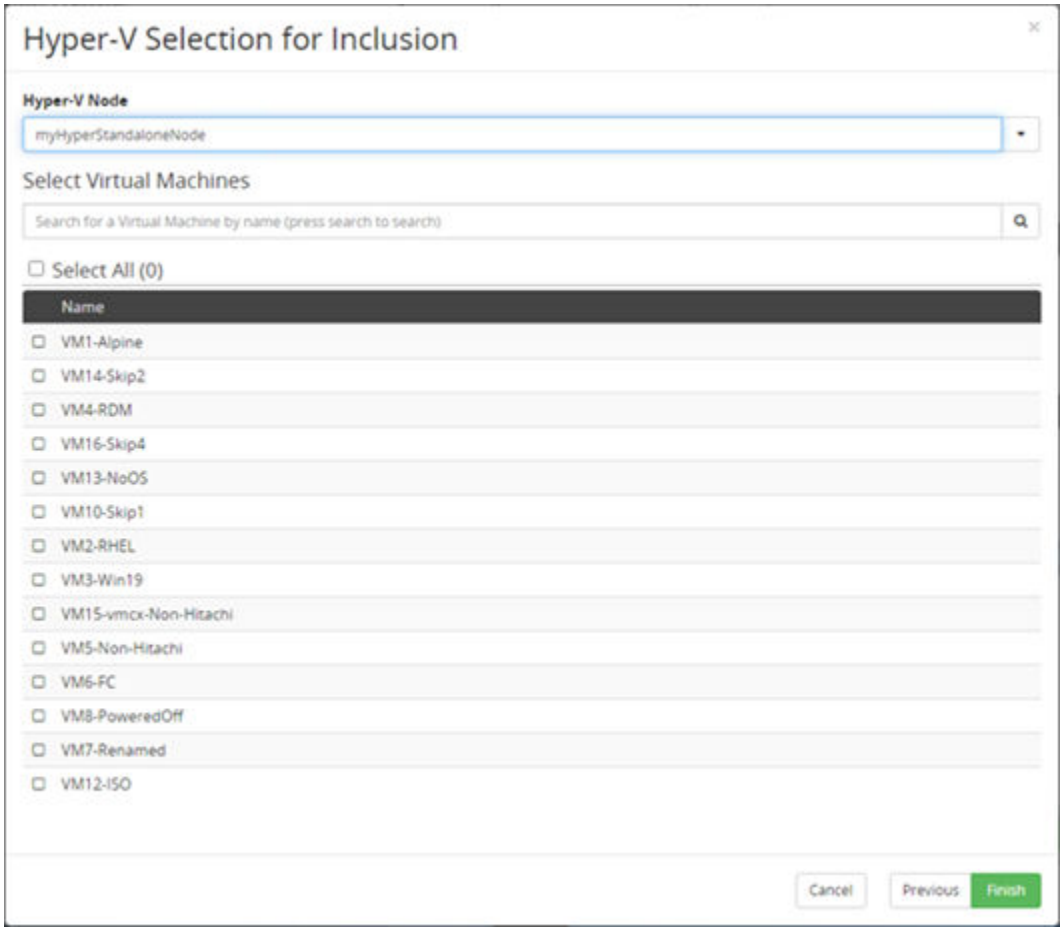


Figure 426 Hyper-V Policy Wizard - Classification – Browse Virtual Machines

Control	Description
Hyper-V Node	Select the Hyper-V app node you want to select VMs from.
Search	Enter a part of a virtual machine name and confirm to filter the list of virtual machines.

Control	Description
Virtual Machine List	Select one or more machines.

Hyper-V Resource Selection Wizard – Browse by Virtual Machine Locations

This page of the wizard is displayed when the browse by virtual machine paths option is selected in the initial wizard page above.

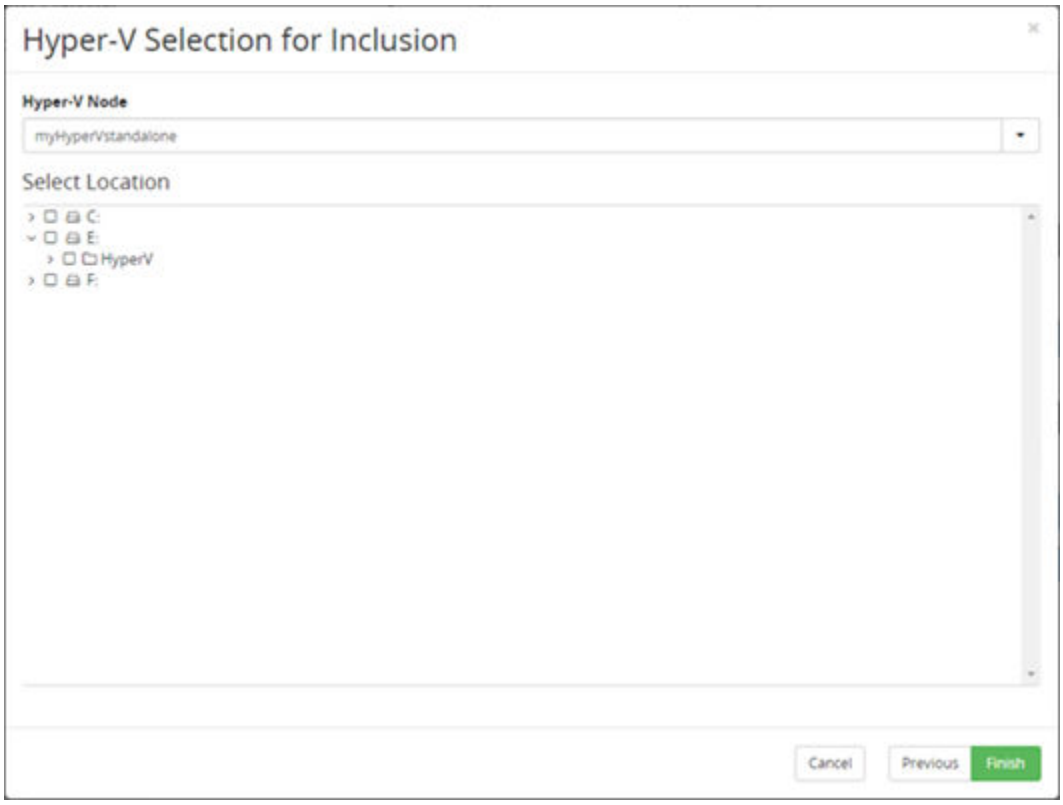


Figure 427 Hyper-VM - Browse by Virtual Machine Locations

Control	Description
Hyper-V Node	Select the Hyper-V app node you want to select paths from.
Location	Select one or more paths. Protector will search for virtual machines under the defined path at the beginning of each backup.

Control	Description
	Note: ▪ A virtual machine will only be selected for a backup if the virtual machine configuration is located under the specified path. If any included VMs utilize additional paths, these will be added as to the backup as well. This ensures that backed up virtual machines can be fully restored. ▪ For clustered Hyper-V nodes it is only possible to select paths which are located on cluster shared volumes, as only these paths are available on all nodes

Hyper-V Resource Selection Wizard – Browse by Virtual Machine Hosts

This page of the wizard is displayed when the browse by virtual machine host option is selected in the initial wizard page above.

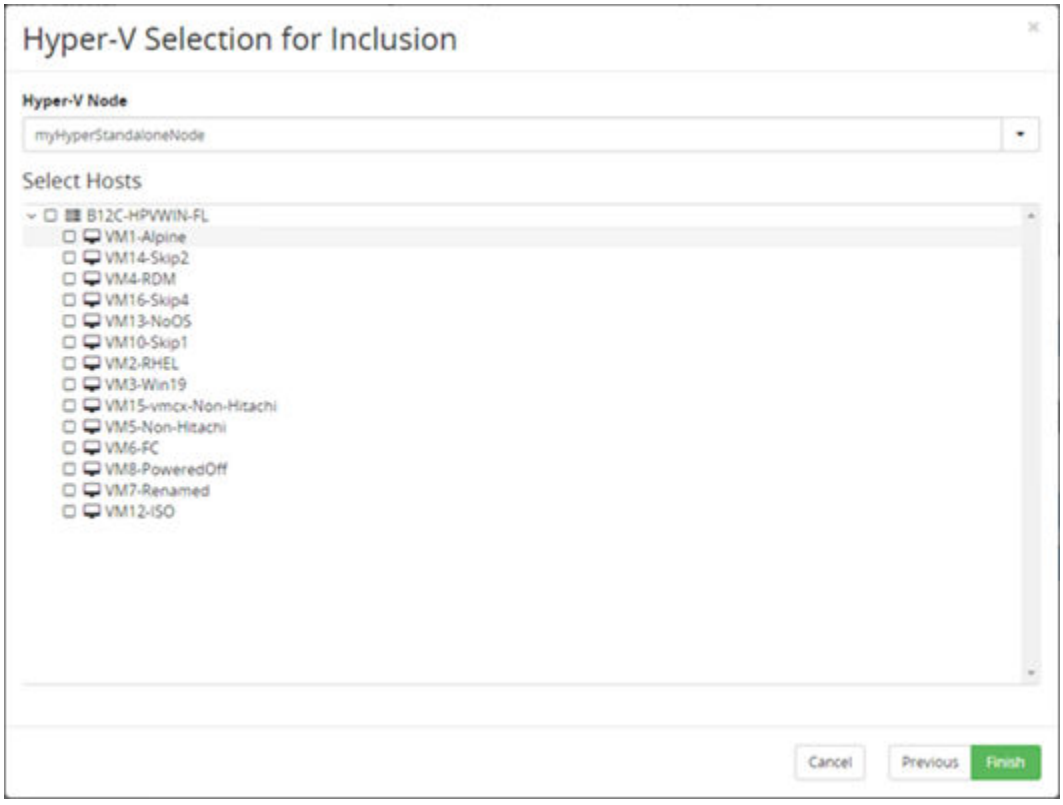


Figure 428 Hyper-V VM Selection - Browse by Virtual Machine Hosts

Control	Description
Hyper-V Node	Select the Hyper-V app node you want to select hosts or VMs from.
Hosts	Select either individual virtual machines or complete hosts. In case a host is selected, Protector will determine the list of VMs available on the host at the beginning of the backup.

Hyper-V Resource Selection Wizard – Pattern search

This page of the wizard is displayed when the pattern search option is selected in the initial wizard page above.

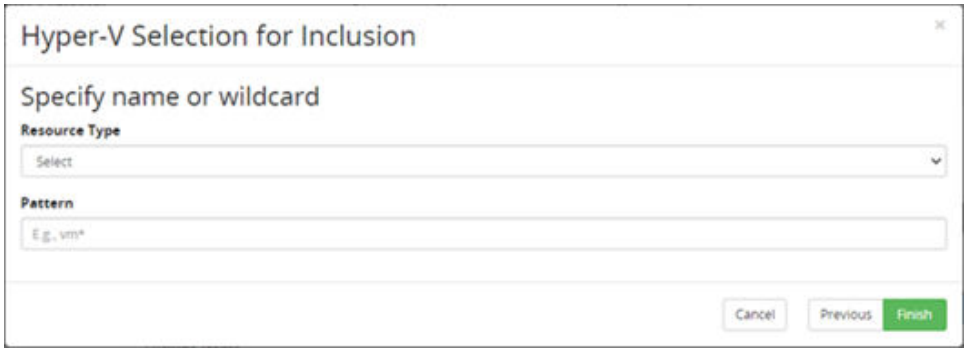


Figure 429 Hyper-V Resource Selection Wizard – Pattern search

Control	Description
Resource Type	Select a Hyper-V resource type, that will be matched by the provided name pattern. Available options: ▪ Virtual Machine ▪ Virtual Machine Location (Path) ▪ Virtual Machine Host
Pattern	Enter a case insensitive pattern that will be used to match the resource type by name. The '?' character matches any single character, while the '*' character can be used to match any sequence of characters. E.g.: **`IH_*`** would match any resource of the given type whose name begins **`IH_`**. **Note:** Resources are re-evaluated against the name pattern every time the policy is executed. New resources having a name that matches this pattern, added after the policy is activated, will be automatically included in the backup.

VMware Classification Wizard

This wizard is launched when a new VMware classification is added to policy.

The VMware classification is used as a means of conveniently specifying the VMware resources on a vCenter. Refer to About VMware policy classifications (on page 95) for details about how this classification works with host and block based operations.

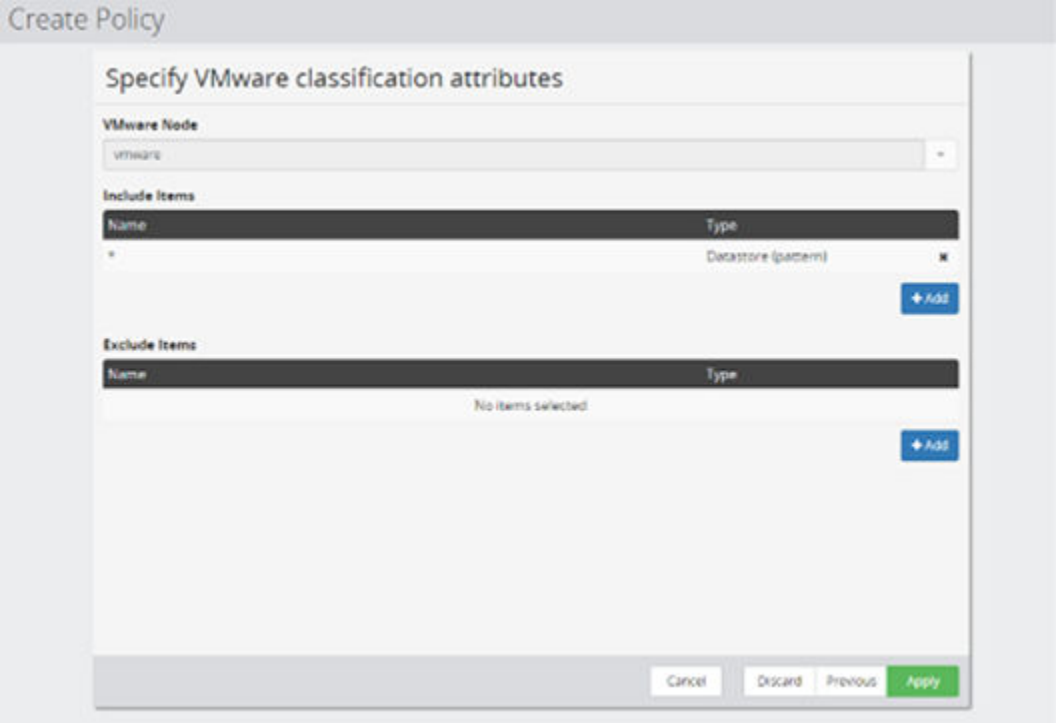


Figure 430 VMware Wizard - Specify VMware classification attributes

Note: If you attempt to edit a legacy VMware policy classification created prior to Protector 6.5, a message will be displayed asking you to reset the Include Items and Exclude Items lists.

Control	Description
Included Items	Lists the VMware resources that will be included in the backup policy. **Note:** Protector will not take a snapshot of a VM if that VM already has a Protector snapshot mounted to it.
Add	Opens the VMware Resource Selection Wizard (on page 701) to enable VMware resources to be added to the include list above.
Excluded Items	Lists the VMware resources that will be excluded from the backup policy.
Remove	Each row has a remove button at the end of the row, the selected VMware resource is removed from the include/exclude list.

Control	Description
Add	Opens the VMware Resource Selection Wizard (on page 701) to enable VMware resources to be added to the exclude list above.
VMware Node	Select the VMware node the policy is being created for.

VMware Resource Selection Wizard

This wizard is displayed when the user includes or excludes VMware resources in a policy.

Caution: Protector tracks VMware resources via their MoRef (Managed Object Reference). If a resource's MoRef is changed then it will not be included in the backup and a warning will be logged. Tracking resources via their MoRef means that they will be included in the backup even if vMotion moves them.



Figure 431 VMware Resource Selection for Inclusion/Exclusion Wizard - Select method

Control	Description
Browse for resources	Select this option to browse for VMware resources in similar ways to those provided in vSphere Client. See **VMware Resource Selection for Inclusion/Exclusion Wizard - Browse By** below.
Specify resource by name or wildcard	Select this option to specify a resource by type and name pattern match. See **VMware Resource Selection for Inclusion/Exclusion Wizard - Specify name or wildcard** below.

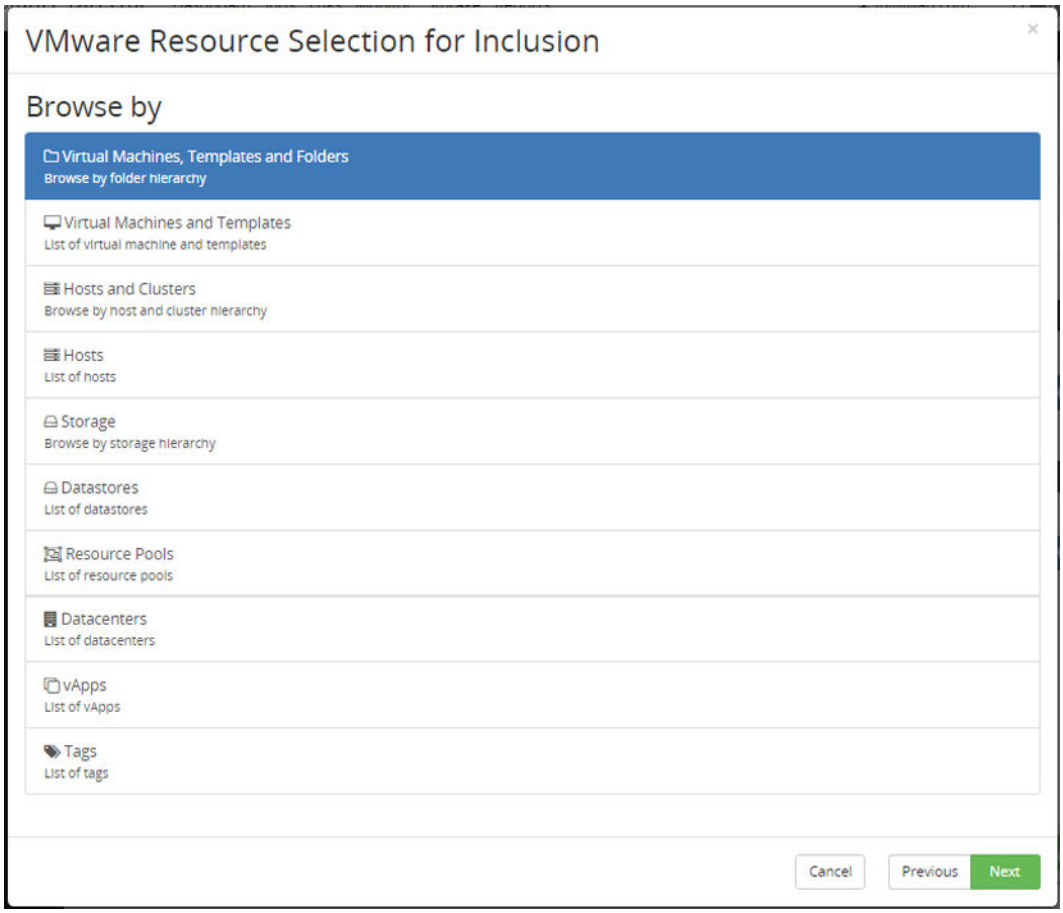


Figure 432 VMware Resource Selection for Inclusion/Exclusion Wizard - Browse by

This page of the wizard is displayed when the Browse for resources selection method is specified in the initial wizard page above.

Control	Description
Virtual Machines, Templates and Folders	Displays a hierarchical view ordered by datacenters, folders, and virtual machines and templates.
Virtual Machines and Templates	Displays a flat list of virtual machines and templates ordered alphabetically.
Hosts and Clusters	Displays a hierarchical view ordered by datacenters, hosts and virtual machines.
Hosts	Displays a flat list of hosts ordered alphabetically.
Storage	Displays a hierarchical view ordered by datacenters and datastores.
Datastores	Displays a flat list of datastores ordered alphabetically.
Resource Pools	Displays a flat list of resource pools ordered alphabetically.

Control	Description
Datacenters	Displays a flat list of datacenters ordered alphabetically.
vApps	Displays a flat list of vApps ordered alphabetically.
Tags	Displays a flat list of tags ordered alphabetically. **Note:** To browse by tags, the VMware proxy node must have PowerCLI installed. Refer to VMware Product Interoperability Matrices for vCenter Server/PowerCLI version compatibility.

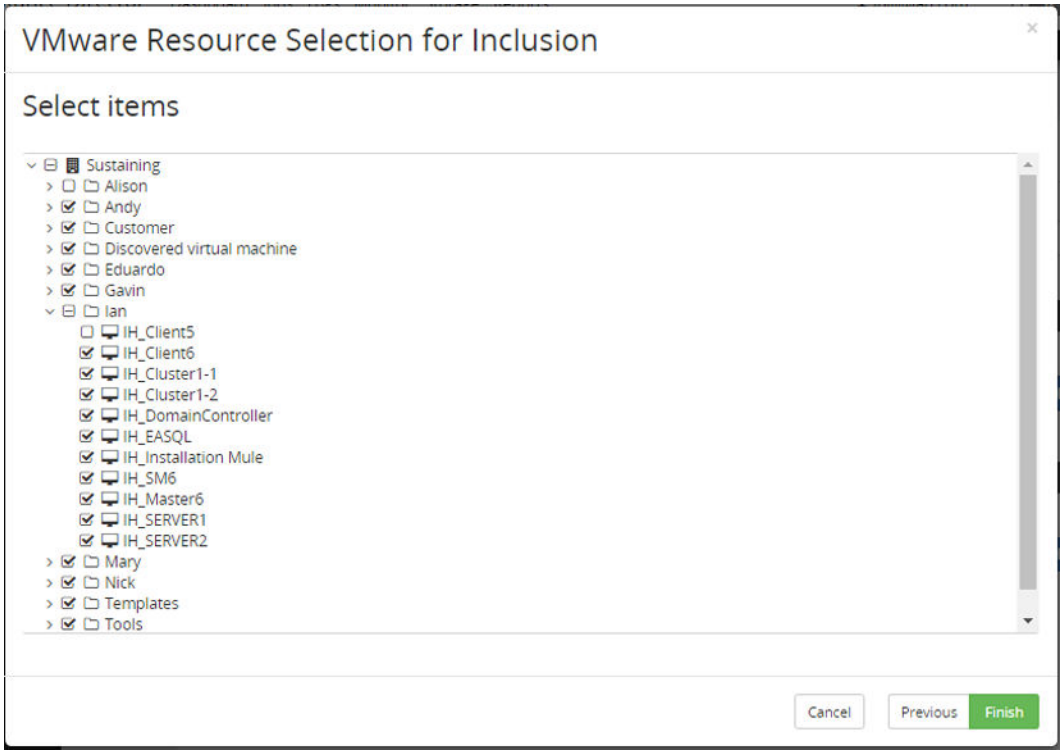


Figure 433 VMware Resource Selection for Inclusion/Exclusion Wizard - Virtual Machines, Templates and Folders Hierarchy

In a hierarchical view it is possible to select or deselect entire trees, sub-trees and individual nodes. For example, the screen-shot above shows the entire *Sustaining* datacenter selected, but with the *Alison* folder and the *IH_Client5* virtual machine deselected from a backup policy.

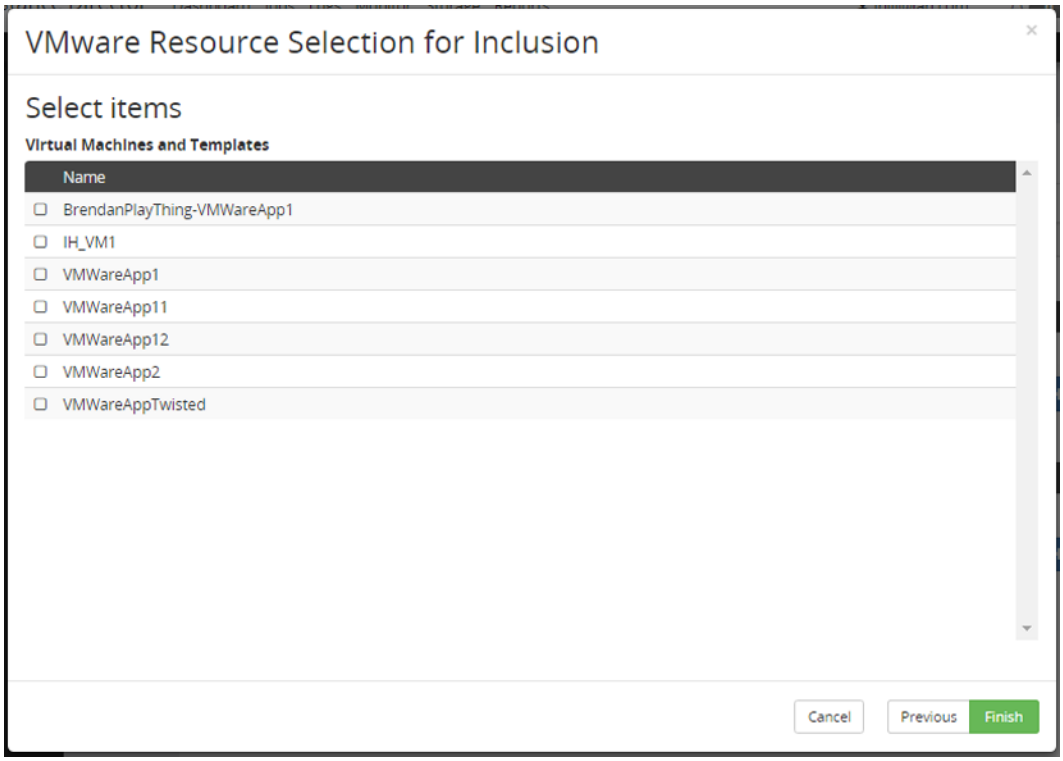


Figure 434 VMware Resource Selection for Inclusion/Exclusion Wizard - Virtual Machines and Templates List

In a flat list view it is possible to select or deselect multiple items of the same type.

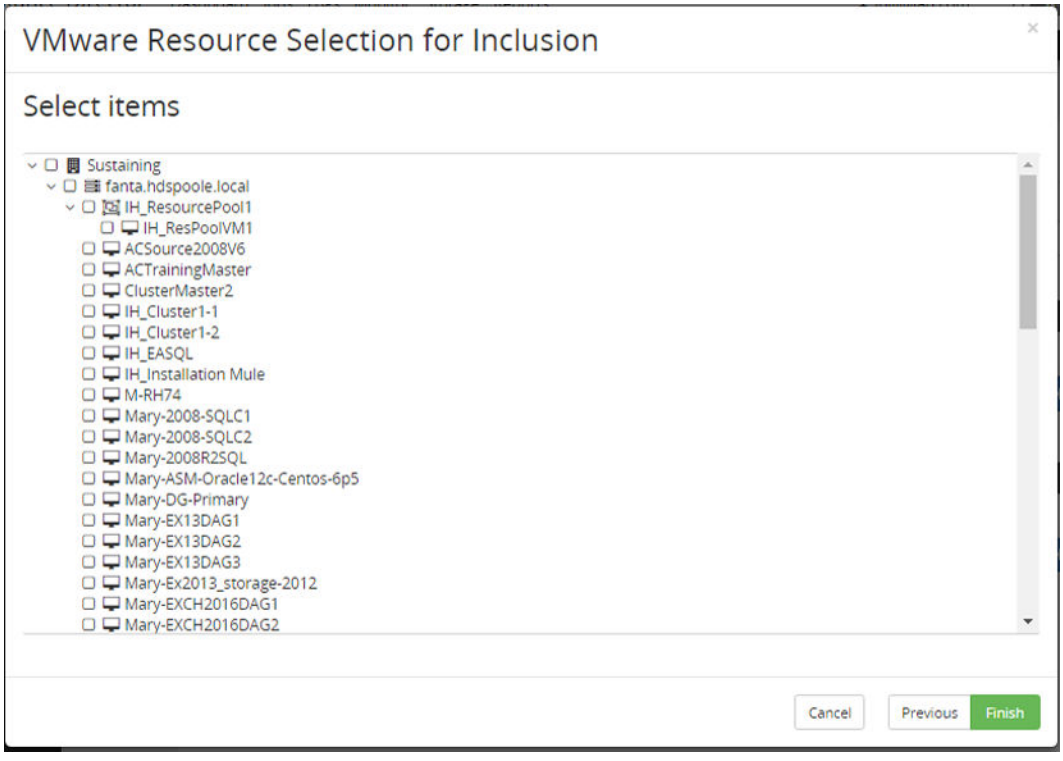


Figure 435 VMware Resource Selection for Inclusion/Exclusion Wizard - Hosts and Clusters Hierarchy

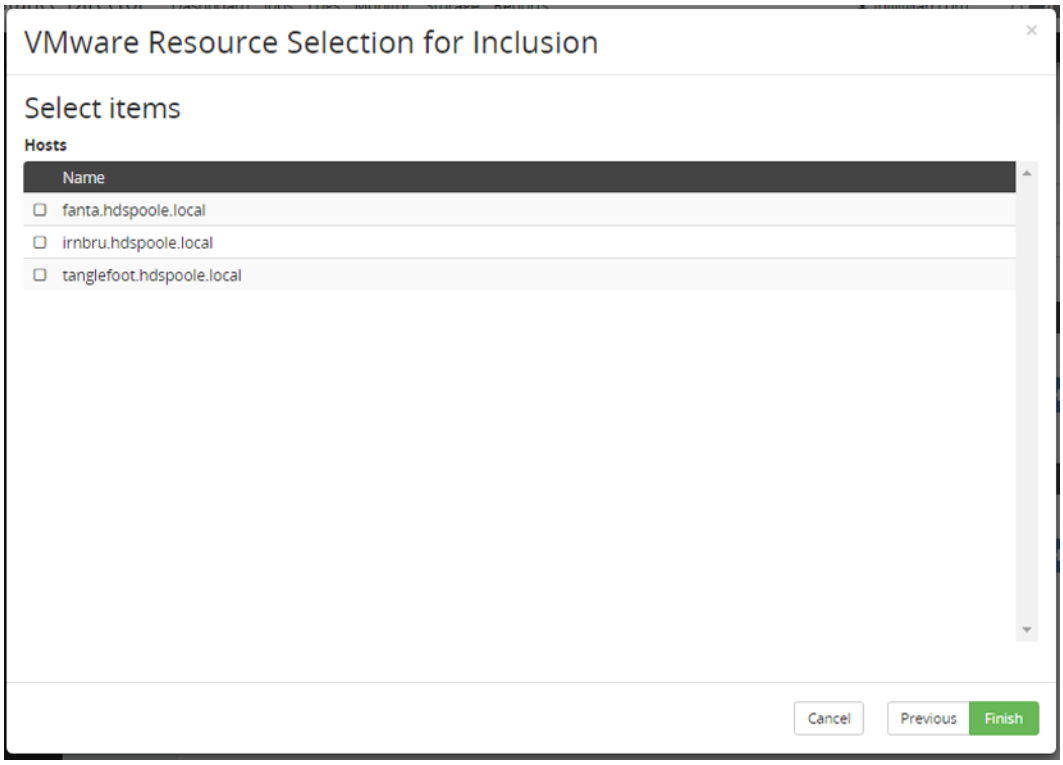


Figure 436 VMware Resource Selection for Inclusion/Exclusion Wizard - Hosts List

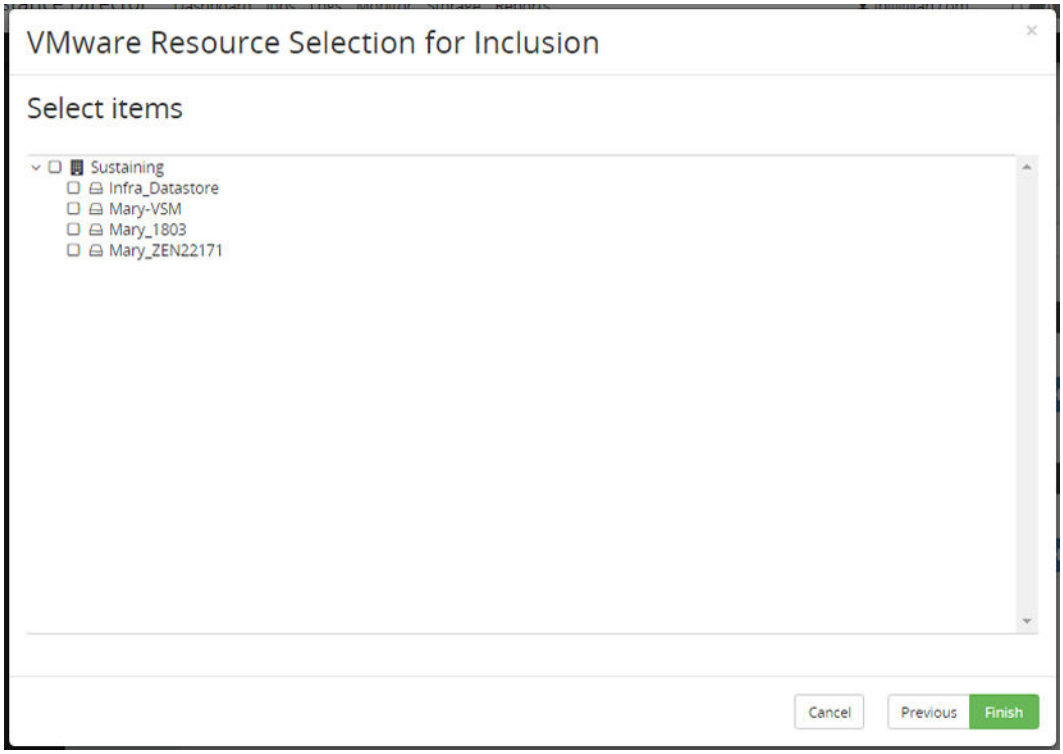


Figure 437 VMware Resource Selection for Inclusion/Exclusion Wizard - Storage Hierarchy

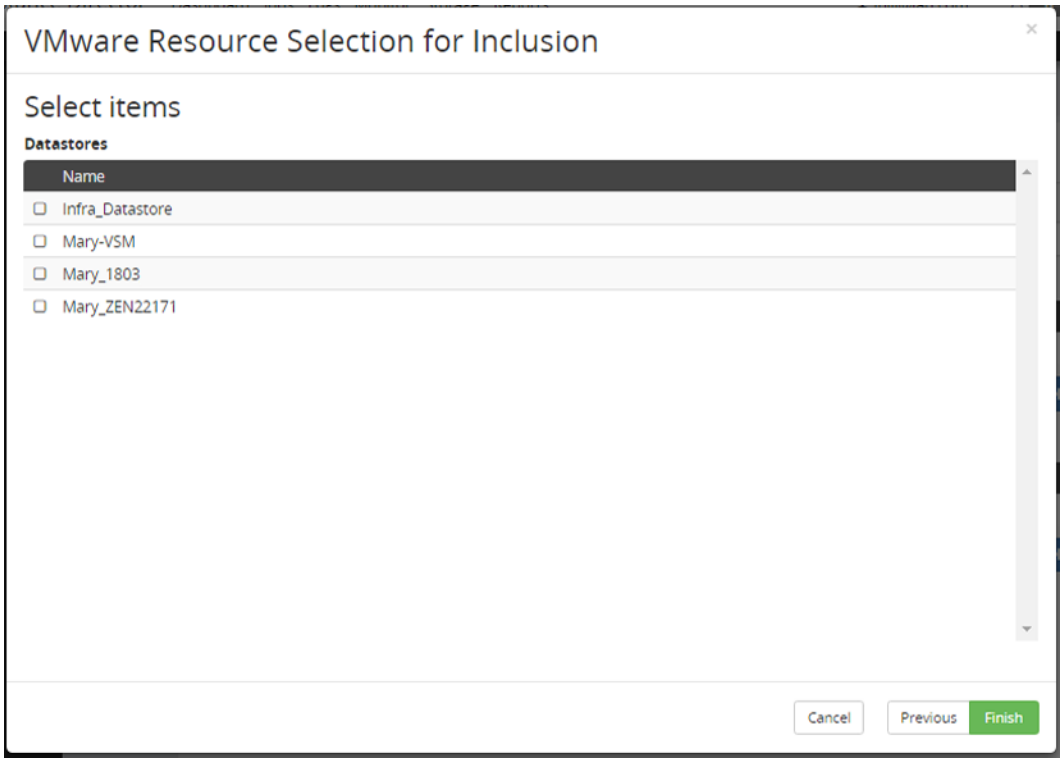


Figure 438 VMware Resource Selection for Inclusion/Exclusion Wizard - Datastores List

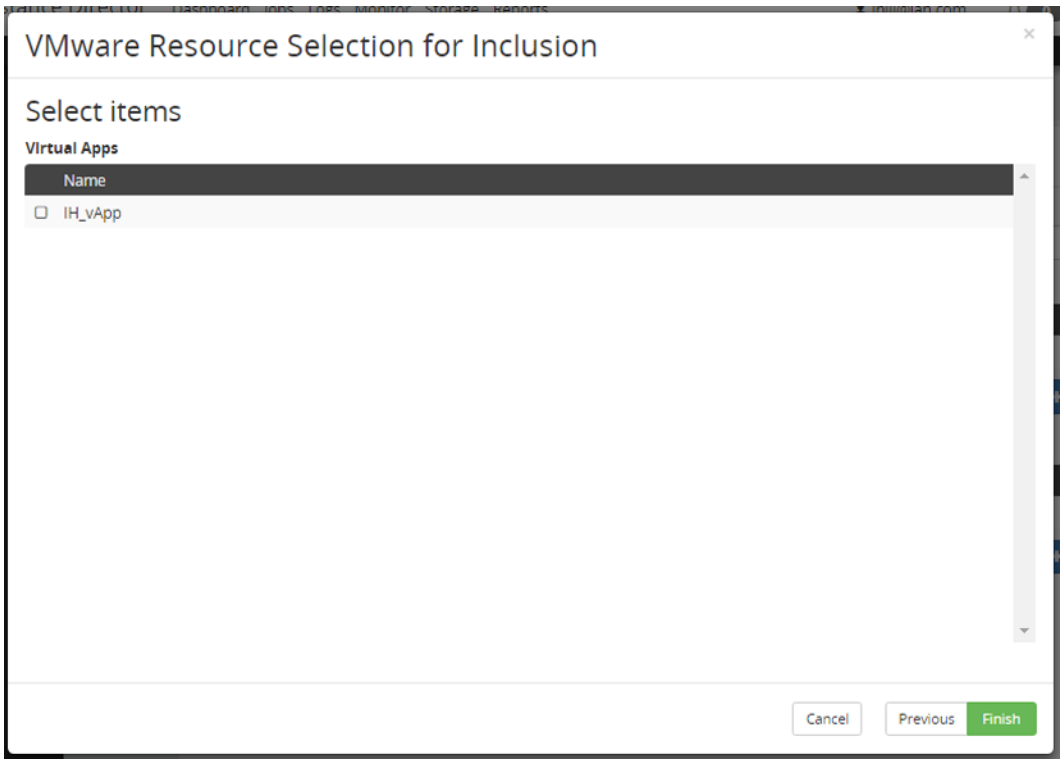


Figure 439 VMware Resource Selection for Inclusion/Exclusion Wizard - vApps List

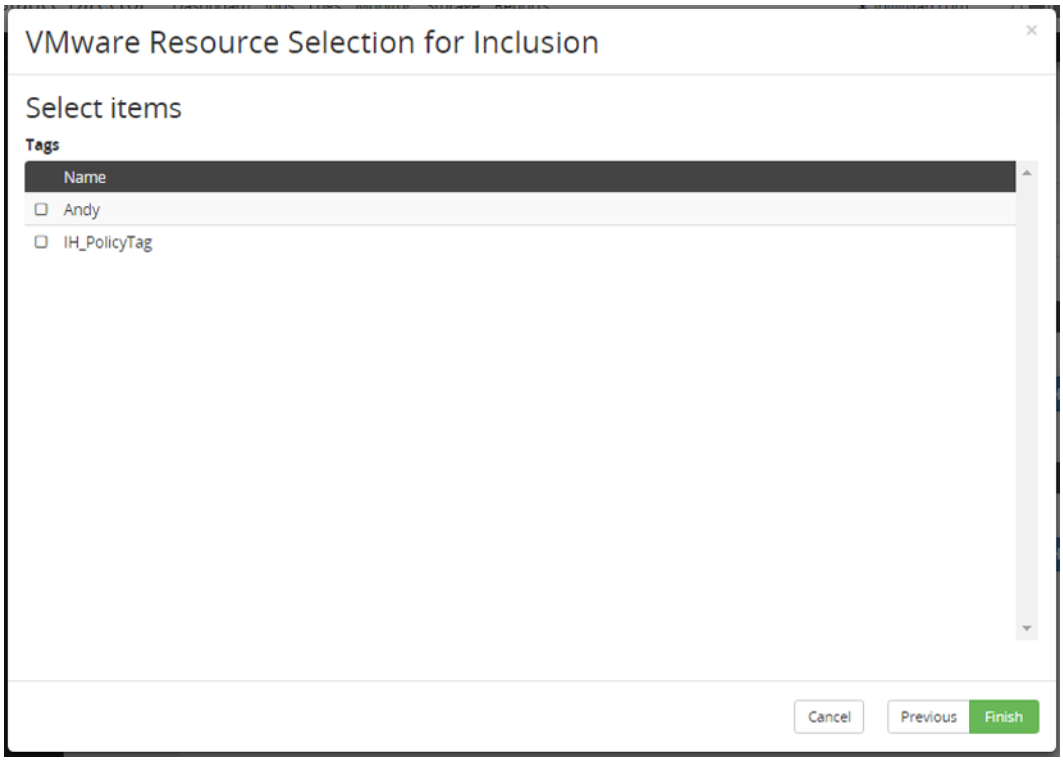


Figure 440 VMware Resource Selection for Inclusion/Exclusion Wizard - Tags List

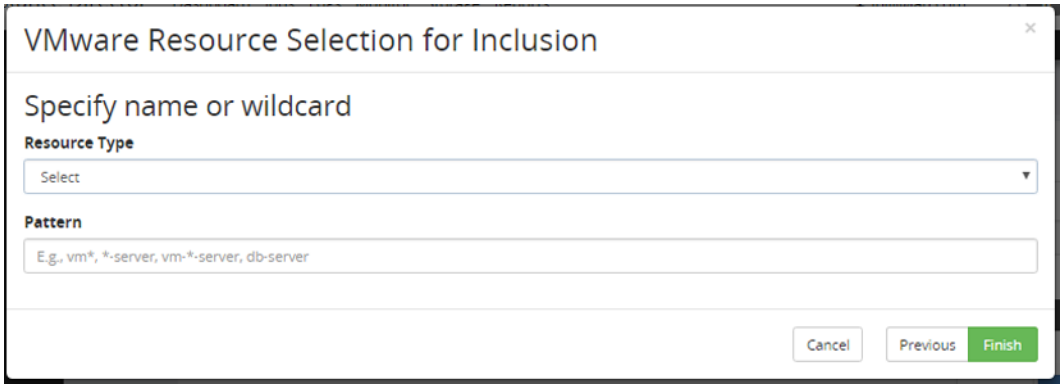


Figure 441 VMware Resource Selection for Inclusion/Exclusion Wizard - Specify name or wildcard

This page of the wizard is displayed when the Specify resource by name or wildcard selection method is specified in the initial wizard page above.

Control	Description
Resource Type	Select a VMware resource type that will be matched by the provided name pattern.

Control	Description
Pattern	Enter a case insensitive pattern that will be used to match the resource type by name. The '*' character can be used to match any sequence of characters. E.g.: `IH_*` would match any resource of the given type who's name begins `IH_`. **Note:** Resources are re-evaluated against the name pattern every time the policy is executed. New resources having a name that matches this pattern, added after the policy is activated, will be automatically included in the backup.

Disk Type Classification Wizard

This wizard is launched when a new Disk Type classification is added to a Policy.

The Disk Type classification is used to limit backups to files that are located on fixed devices and/or to files that are located on removable disks. To specify file types use the Classification Filters Wizard (on page 727) in conjunction with this Classification.

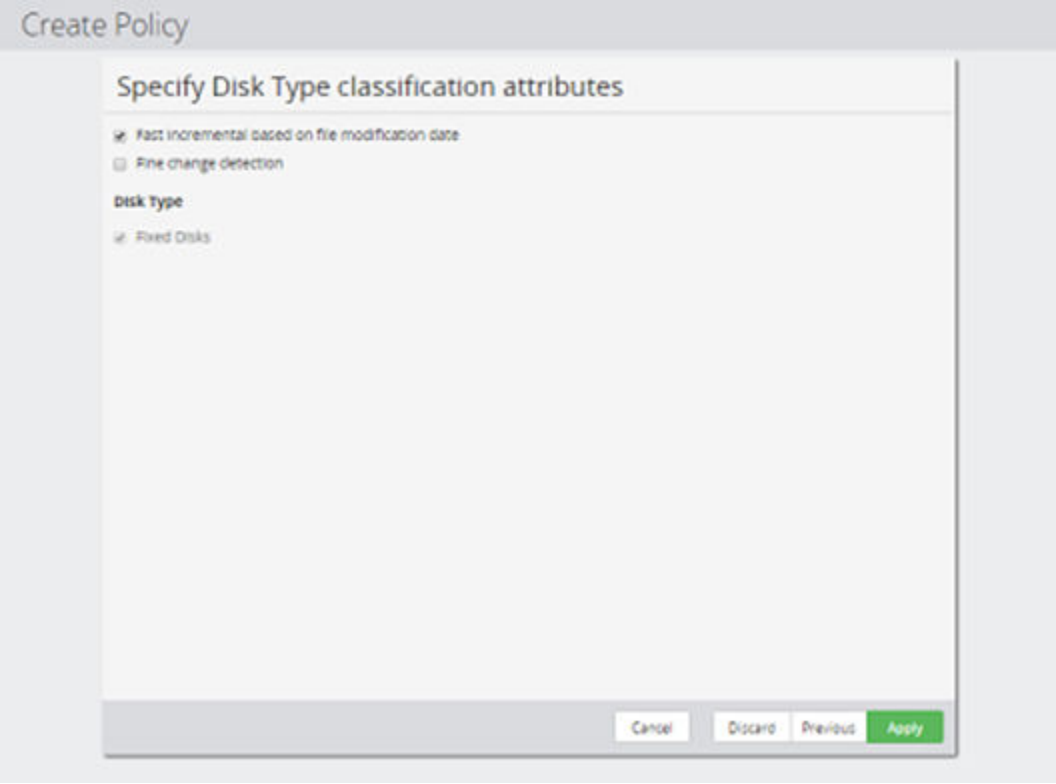


Figure 442 Physical Storage Type Wizard

Control	Description
Fixed Disks	Backs up all fixed disks, as this is the only disk type option it is permanently selected. All drives connected to SCSI, ATAPI, ATA, SAS, Fibre, RAID, and so on are considered fixed.
Fast incremental based on file modification date	This is only applicable for backups to Generation 2 storage nodes. If this is checked then Ops Center Protector decides what files need resynchronizing based on whether the modification date has changed. This reduces the time taken to resynchronize, but can be disabled if it is known that software is installed that will modify files without updating their size or modification date. NoteIf only file metadata changes between batch backups (e.g. file ownership or file permissions), then the changes are not captured. These changes are only captured when the file data changes.
Fine change detection	This is only applicable for backups to Generation 2 storage nodes. Reduces the amount of data transferred and stored during a resynchronization. An entire file is transferred if it has changed and is less than one block in size. (This option should be used sparingly as there is a processing overhead.)

Hitachi Block Classification Wizard

This wizard is launched when a new Hitachi Block classification is added to a Policy.

The Hitachi Block classification is used to define which LDEVS are to be protected.

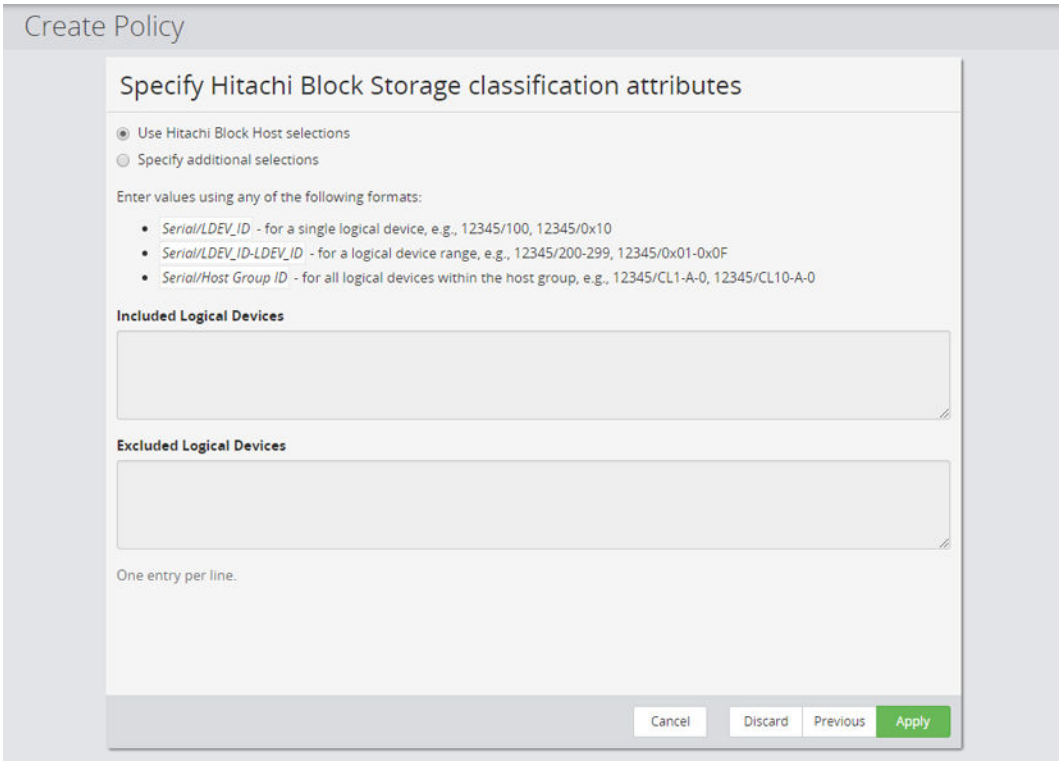


Figure 443 Hitachi Block Classification Wizard - Specify Hitachi Block Storage classification attributes

Control	Description
Use Hitachi Block Host selections	Select this option if the policy is used in conjunction with a Hitachi Block Host Node (see Hitachi Block Host Node Wizard (on page 559)). In this case the LDEVs were specified when the node was created.
Specify additional selections	Select this option to specify the LDEVs when used in conjunction with a: ▪ Hitachi Block Node (see Hitachi Block Device Node Wizard (on page 583)) or a ▪ Hitachi Logical Block Node (see Hitachi Logical Block Device Node Wizard (on page 609)) You can also select this option if the policy is used in conjunction with a Hitachi Block Host Node (see Hitachi Block Host Node Wizard (on page 559)). In this case the LDEVs specified here will be in addition to those specified when the node was created. **Note:** This option is not available for LDEVs using an NVMe subsystem.
Included Logical Devices	Enter a list of LDEVs or Host Groups to be included in this Block Host (one logical device or host group specification per line) in the format described below.

Control	Description
Excluded Logical Devices	Enter a list of LDEVs or Host Groups to be excluded from this Block Host (one logical device or host group specification per line) in the format described below. **Tip:** This enables the inclusion of an entire Host Group but exclusion of individual LDEVs within that Host Group.

LDEVs are entered in the following format:

serial_number/*identifier*

Where:

- *serial_number* is the serial number of the Block storage device
- *identifier* is one of:
  - *ldev_id* a single LDEV ID in hex or decimal
  - *ldev_id*-*ldev_id* a range of LDEV IDs
  - *host_group_id* a host group identifier using the format:
    - `CL`*c*-*s*-*h* where:
      - `CL` is a literal string
      - *c* is the physical channel number in the range 1…n
      - *s* is the physical slot number in the range A…Z
      - *h* is the logical host group ID in the range 0…255

For example:

`210613/0x00a0` for an LDEV specified in hexadecimal

`210613/220-230` for a range of LDEVs specified in decimal

`210613/CL2-A-12` for a host group

Note:

- Spaces and colons must not be present in the classification entries. If a space or colon is encountered on a line, the remaining text on that line from the space or colon will be disregarded. This allows existing entries (that contain the LDEV name after a space or colon) to be present in the list.
- When using pseudo multi-tenancy, if an LDEV or host group is manually specified that is not accessible then it will be excluded from the backup. An attachment is included in the log message that specifies what was included and excluded in the backup.
- If a selection of provisioned ports are specified, the port with the least amount of LUNs assigned to it will be used. The exception is when a TI Snapshot already exists and an update is made to the provisioned ports, in this case the port used will be taken from the existing TI snapshot set.
- When a replication already exists and a change is made to provisioned ports then the affected data flow must be reactivated. The existing replication will remain using the existing port after it has been resynchronized. New replications will use the updated settings.

Hitachi Software Block Classification Wizard

This wizard is launched when a new Hitachi Software Block classification is added to a Policy.

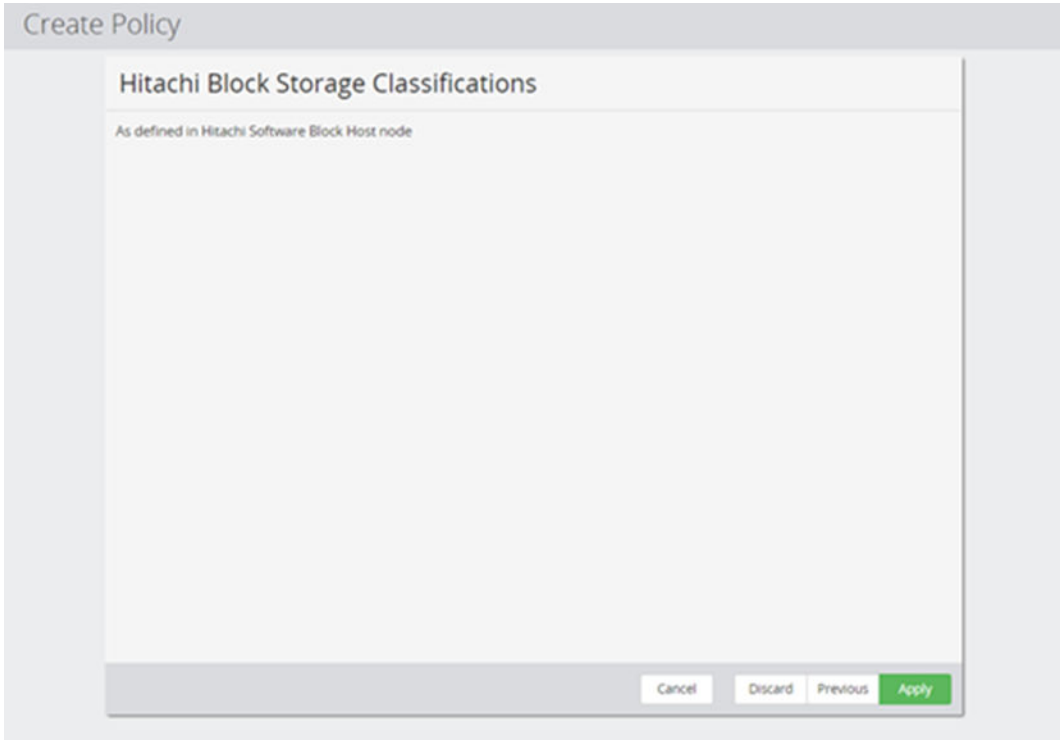


Figure 444 Hitachi Software Block Classification Wizard - Specify Hitachi Software Block Storage classification attributes

Control	Description
Hitachi Software Block Host Classifications	There are no configuration options for the software block host classification, the volumes are defined in the Software Block Host node.

Path Classification Wizard

This wizard is launched when a new Path classification is added to a Policy.

The Path classification allows you to define the path or paths for directories that contain files to be backed up. All files in each directory will be backed up. To specify file types use the Classification Filters Wizard (on page 727) in conjunction with this Classification.

Caution:

- Protector will not backup the files or directories referred to by a symbolic link. The symbolic link itself will be backed up, so upon restore the link will still work assuming the referenced files or directories exist in their original location. Essentially, Protector performs a shallow backup in respect of symbolic links.
- The following filesystem types are excluded from backups:
  - Windows: bad, cdrom, no root, ramdisk, unknown
  - Linux: autofs, binfmt_misc, configfs, debugfs, devpts, devtmpfs, fuse, group, hugetlbfs, mqueue, nfsd, proc, pstore, rootfs, rpc_pipefs, securityfs, selinuxfs, sysfs, tmpfs, vmblock
  - AIX: ahafs, autofs, cdrom, procfs

Note: When used in the context of storage hardware backups, Protector will discover the underlying hardware paths at runtime. For storage hardware based backups, all the paths must exist on the same block hardware device.

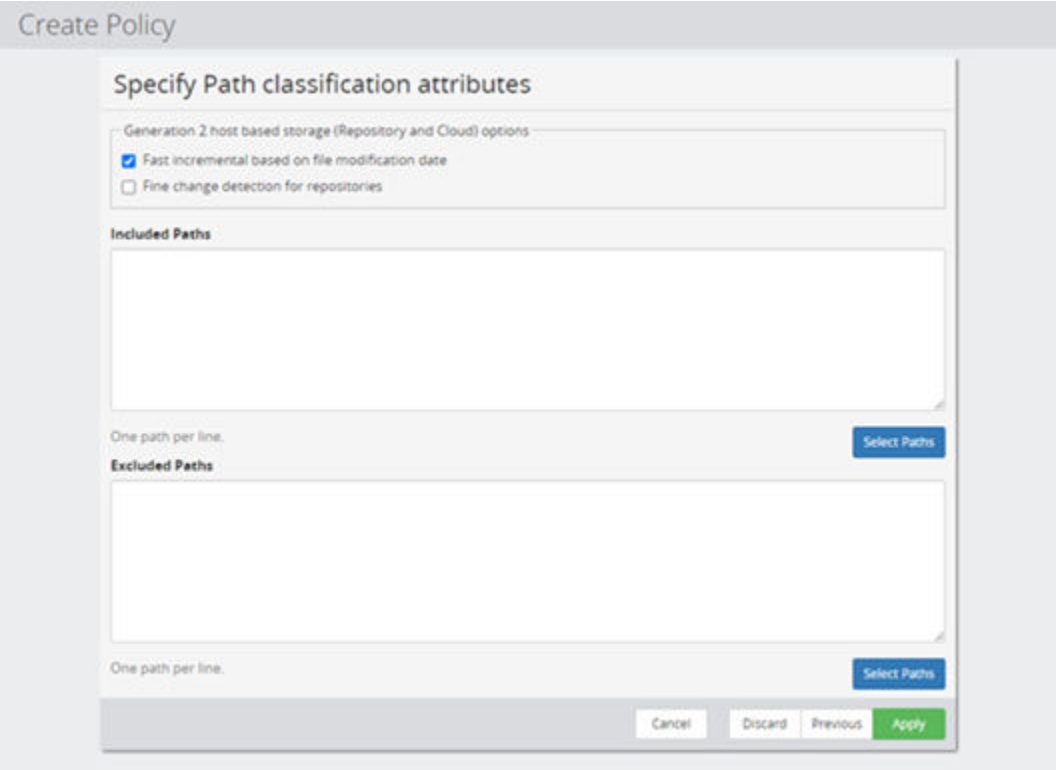


Figure 445 Physical Path Wizard - Specify Path classification attributes

Control	Description
Included Paths	Enter the paths to be included in the backup policy. The targeted directory paths are entered with one path specified per line. To specify file types add a Filter using the Classification Filters Wizard (on page 727). Refer also to Path selection rules (on page 716).
Excluded Paths	Enter the paths to be excluded from the backup policy. The targeted directory paths are entered with one path specified per line. Files themselves cannot be specified – only directories. If Included Paths to be backed up contains one or more sub-paths that are not required, then enter the paths of these sub-directories here. **Note:** If Excluded Paths contains paths that are not sub-paths of Included Paths, Protector will backup the entire file system. Refer also to Path selection rules (on page 716).
Select Paths	Opens the Select Classification Paths Dialog (on page 715) to enable you to select existing paths to be included or excluded. You can search for a path on a specific node and then apply that path broadly to all nodes or apply it to one node only.

Control	Description
Fast incremental based on file modification date	Only applicable to Generation 2 host based storage. If this is checked then Ops Center Protector decides what files need resynchronizing based on whether the modification date has changed. This reduces the time taken to resynchronize, but can be disabled if it is known that software is installed that will modify files without updating their size or modification date. Note: If only file metadata changes between batch backups (e.g. file ownership or file permissions), then the changes are not captured. These changes are only captured when the file data changes.
Fine change detection for repositories	Only applicable to Generation 2 repository backups. Reduces the amount of data transferred and stored during a resynchronization. An entire file is transferred if it has changed and is less than one block in size. (This option should be used sparingly as there is a processing overhead.)

Tip: It is possible to use wildcard characters (*) within each path entry. Only one wildcard character is permissible within each path element. That is: `C:\*\hello` is valid, whereas `C:\*document*\hello` is not valid.

It is possible to use variables within a Path classification. See About path macros (on page 98) for details.

Select Classification Paths Dialog

This dialog is displayed when selecting paths (e.g. in the Path Classification Wizard).

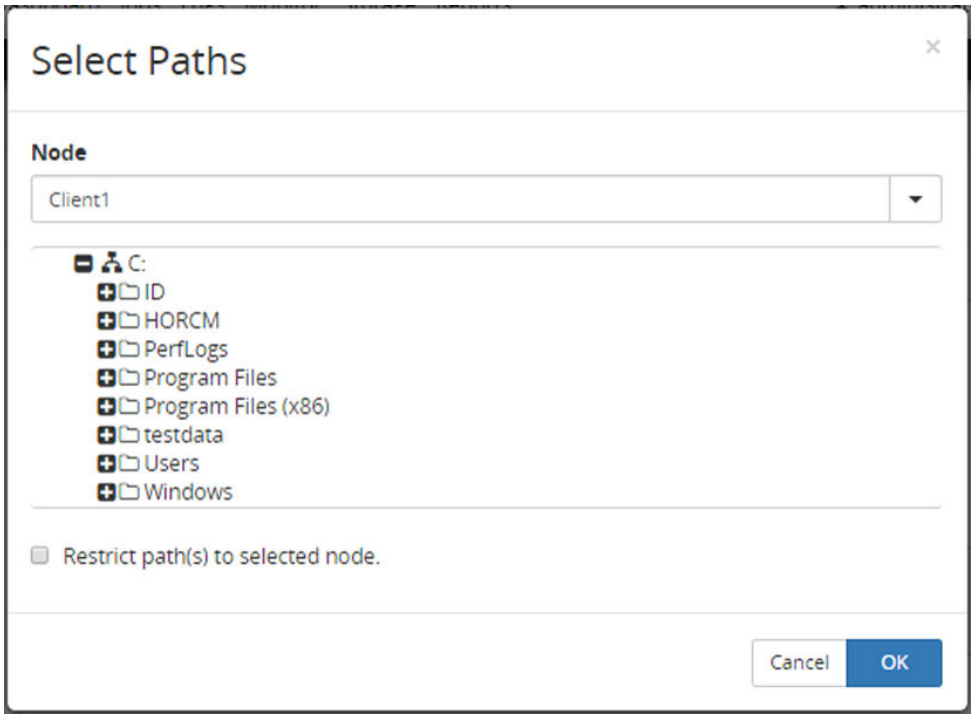


Figure 446 Select Paths Dialog

Control	Description
Node	Enter the name of a node or select one from the dropdown menu.
Directory Tree	Expand the directory tree for the node and select the required directory.
Restrict path(s) to selected node	If selected then the path will only be applied to the node selected above; the path will be qualified with the node name (i.e. `{node_name}path_name`) in the Path Classification Wizard (on page 713) path list. If deselected then the path will be applied to all nodes where the policy is assigned.

Path selection rules

A *Physical - Path* policy classification is a set of *Included Paths*, *Exclude Paths* and *Filters*. Paths must be specified following the rules listed below:

General path rules:

- Only directories can be specified by a path; files cannot.
- A path specification is only considered to refer to the leaf node of the path. For example:

  `/path_element_1/path_element_2`

  only refers to the `path_element_2` leaf node.
- `/home/` is the same as `/home`. The trailing delimiter is ignored.
- An *included path* must be defined to be a valid classification.
- An *excluded path* will never implicitly include other paths.
- An *excluded path* will be ignored if it is not a sub directory of an *included path*.
  - For example, if `/home` is set as both an *included path* and an *excluded path*, it will be included. This is because the *excluded path* it is not a sub-path of the *included path*.
- Only local files are backed up. Filesystem walking will not continue beyond a remote mount point.

Windows specific path rules:

- A windows path must start with `<drive letter>:\` (`<drive letter>` can be an actual drive letter or a `*` or `?`).
- Delimiters can be `/` or `\` and can be mixed within a path specification.

Posix specific path rules:

- A Posix path must start with `/`
- Delimiters can only be `/`

Wild card rules:

- The following characters that have special meaning in a path specification:
  - `*` matches zero or more characters.
  - `?` matches only a single character.
- If a path element contains a `*` it only matches directories at that level.
- Any number of `*` and `?` wildcards can be used in a path.

File filters:

- A file filter will only match files not directories.
- The same wildcard rules apply; directories will not be matched.

Path examples:

Example	Interpretation
`/home/**`	Selects all directories immediately beneath `/home`. The second `*` has no effect.
`/home/*d`	Selects all directories immediately beneath `/home` ending in `d`. The directory `/home` is not selected.
`/home/?`	Selects all directories immediately beneath `/home` that have a name that is exactly 1 character long. The directory `/home` is not selected.
`/home/*`	Selects all directories immediately beneath `/home`. The directory `/home` is not selected.
`?:\`	Selects all Windows drives, as will `*:\`
`/h*/*m*`	Selects all directories that contain an `m`, that are beneath all directories at the root level starting with `h`.
`/h*m*`	Selects all directories at the root level that start with `h` and also have `m` in the name. Note that a directory named `/hm` would be selected since `*` matches zero or more characters.
`/h*me`	Selects all directories at the root level that starts with `h` followed by any number of characters and ending with `me`.
`/h?me`	Selects all directories at the root level that starts with `h` followed by any single character and ending with `me`.

Backup Operation Wizard

This wizard is launched when a new Backup operation is added to a Policy.

The Backup operation makes complete or incremental copies of the source node's data (specified by the associated data classification). These backups are performed at times dictated by the Run Options.

Note: This operation cannot be applied to hardware storage devices.

Note: Operating System Specific Behaviour:

OS	Note
AIX	AIX host-based backups do not capture ACLs.

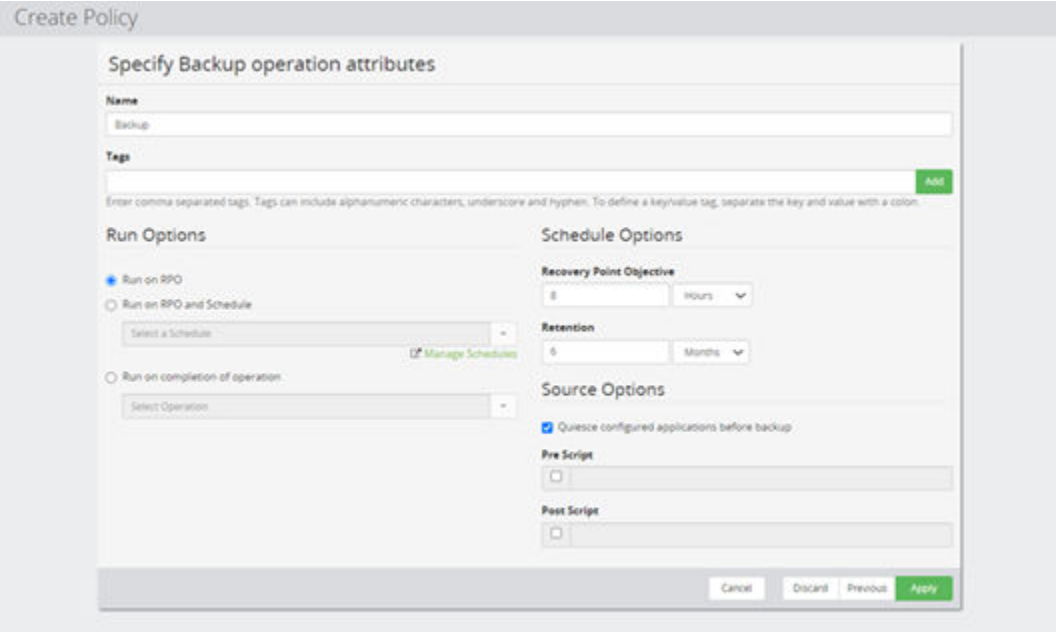


Figure 447 Backup Wizard - Specify Backup operation attributes

Control	Description
Name	The name of the operation. The default name is 'Backup'. This is used to identify the operation when the policy is used in a Data Flow diagram and when operations are to be triggered sequentially.
Tags	Add the tags to be associated with the object being created.
Run Options	Select one of the following: ▪ Run on RPO - the operation is triggered whenever the RPO is reached. ▪ Run on RPO and Schedule - the operation will trigger based on the RPO and the selected Schedule Name. ▪ Run on completion of operation - the operation will be triggered on completion of the operation entered in the Select Operation field. This is used when a policy contains more than one operation and it is necessary that the operations occur in a sequential order.
Select a Schedule Name	Enabled only when the Run Option is set to Scheduled. Enter an existing schedule name or select one from the dropdown menu.
Manage Schedules	Opens the Schedules Inventory (on page 849) in a new tab to enable you to add, edit, review or delete schedules. The times, dates and/or windows within which backups are written to a repository can be controlled by a named schedule that is then selected in the Schedule Name control. Once you have finished working in the Schedules tab, close it and return to the **Backup Wizard**.

Control	Description
Select Operation	Enabled only when the Run Option is set to Run on completion of operation. Enter an existing operation name or select one from the dropdown menu. This is used when a policy contains more than one operation and it is necessary that the operations occur in a sequential order.
Recovery Point Objective	This can be set to a specified period and is used in conjunction with the Run Options. If Run on RPO or Run on RPO and Schedule is selected, then a backup is created whenever the Recovery Point Objective time is reached, following the initial backup performed when the policy is activated.
Retention	The backup is retained for the specified period before being retired.
Quiesce configured applications before backup	This is used to ensure backup consistency. If this option is selected and any Application classifiers are included in the policy, then these applications will be temporarily quiesced into a consistent state prior to the snapshot occurring. See About application quiescing (on page 98).
Pre Script	Specify a script that will be run before the backup is performed. This is used to ensure consistent state for applications other than those with a predefined Application Classification Wizard. **Note:** ▪ The Pre script and Post script fields have no effect with VMware backups. If scripts need to be run either before or after a backup operation then those scripts need to be held locally on each virtual machine that must run them. See the following VMware articles for further information: https://pubs.vmware.com/vsphere-50 http://kb.vmware.com/selfservice ▪ In order to run a Python script as Pre and Post scripts is it necessary to precede the script name with the path and the name of the python binary as follows: `<script bin path>/<full script bin name> <script>` So for Windows this would look like this: `C:\Users\Administrator\AppData\Local\Programs\Python\Python38-32\python.exe E:\myScript.py` For Linux: `/usr/bin/python3 /opt/hitachi/Protector/scripts/myScript.py`
Post Script	Specify a script that will be run after the backup is performed. This is used to ensure consistent state for applications other than those with a predefined Application Classification Wizard.

Mount Operation Wizard

This wizard is launched when a new Mount operation is added to a Policy.

The Mount operation is used to automatically mount and unmount the destination of a Hitachi Block and File storage replication so that it can be used in Repurposing and Proxy Backup scenarios.

Note: The selected mount host must have a pre-existing LUN mounted from the corresponding storage device (this is required for the auto discover feature to work). If not then the automated mount operation will fail and the error message `"Cannot mount to machine that does not have path to storage"` will be logged.

Note:

- Only Replications can be auto-mounted using the Mount operation; Snapshots cannot.
- The storage paths used for Proxy Backup and Repurposing are not deleted immediately after use, therefore it is possible that they may be shown as offline in the OS's Disk Management UI.
- If a mount operation needs to mount multiple disks and one of the mount operations fail, the replication destination is shown as mounted rather than partially mounted. The logs will indicate the mount was only partially successful.

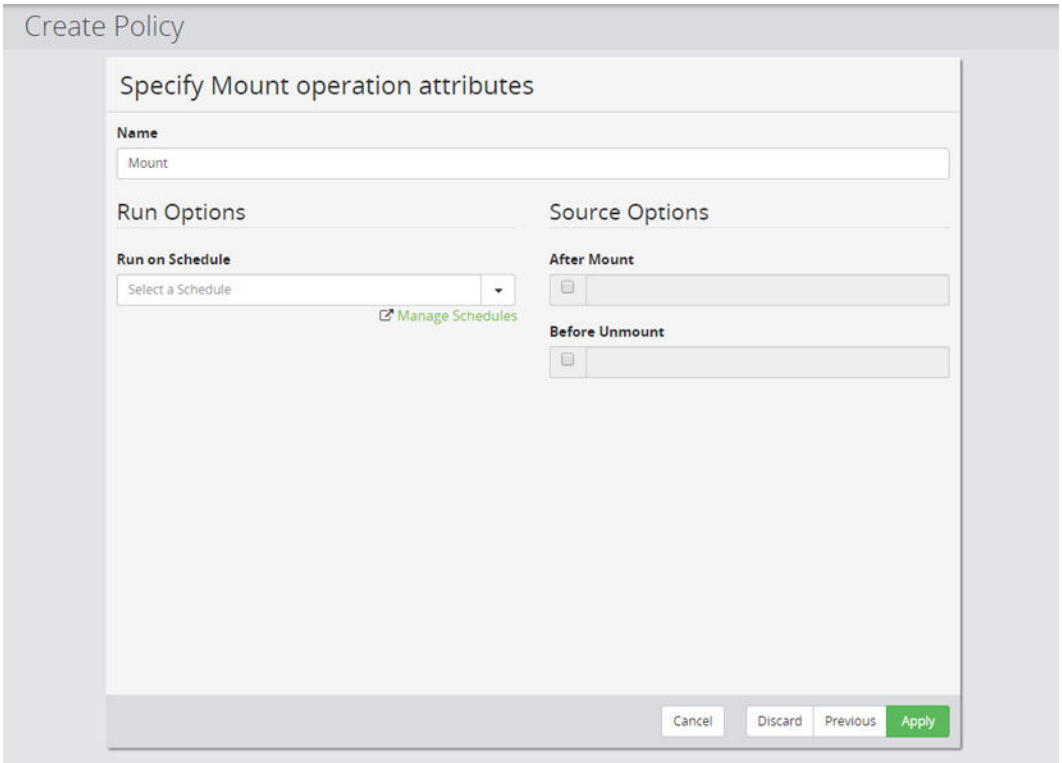


Figure 448 Mount Wizard - Specify Mount operation attributes

Control	Description
Name	The name of the operation. The default name is 'Mount'. This is used to identify the operation when the policy is used in a Data Flow diagram and when operations are to be triggered sequentially.
Run Options	Select one of the following options: ▪ Run on Schedule - specify a trigger schedule name in the Select a Schedule field to specify when the mount operation should occur.
Select a Schedule	The mount operation can have its own trigger schedule that can be used when applied to a continuous replication. The operation can also be part of a synchronization group so that it can participate in a batch replication sequence. Enter an existing schedule name or select one from the dropdown menu.
Manage Schedules	Opens the Schedules Inventory (on page 849) to enable you to add, edit, review or delete schedules.
After Mount	Specify a script that will be run on the mount host after the volume is mounted, allowing the script writer to invoke an action after the volumes have been mounted on the mount host, e.g.: schedule a host based backup. By default Protector will look for the script on the mount host in the `C:\Program Files\Hitachi\Protector\Scripts` folder.
Before Unmount	Specify a script that will be run on the mount host before the volume is unmounted, allowing the script writer to put the application into a safe state before the volumes are unmounted from the mount host. By default Protector will look for the script on the mount host in the `C:\Program Files\Hitachi\Protector\Scripts` folder.

Replicate Operation Wizard

This wizard is launched when a new Replicate operation is added to a Policy.

The Replicate operation maintains a complete copy of the source data (specified by the associated data classification) on the destination.

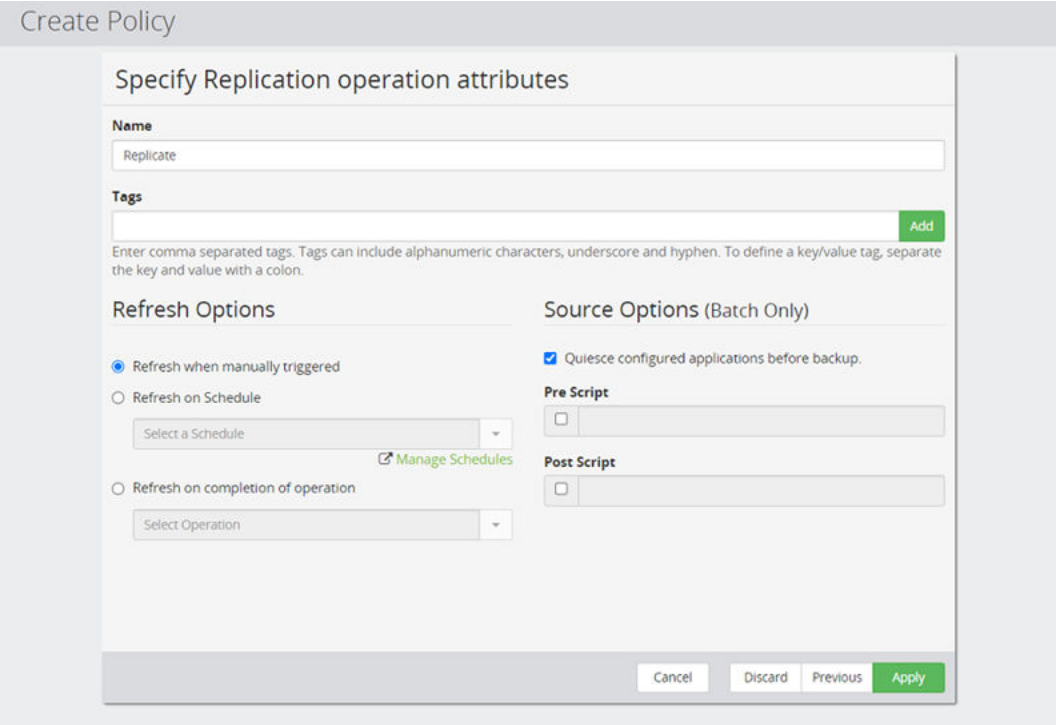


Figure 449 Replicate Wizard - Replication Operation Attributes

Control	Description
Name	The name of the operation. The default name is 'Replicate'. This is used to identify the operation when the policy is used in a Data Flow diagram and when operations are to be triggered sequentially.
Refresh Options	Replication can be performed in real time using a continuous mover or in batch mode using a batch mover. In the case of a continuous mover these options will cause the application to re-evaluate the LDEVs being used in the replication. ▪ Run when manually triggered - the operation is manually triggered. ▪ Run on Schedule - the operation will trigger based on the schedule specified in the Select a Schedule field. ▪ Run on completion of operation - the operation will be triggered on completion of the operation entered in the Select Operation field.
Select a Schedule	Enabled only when the Run Option is set to Scheduled. Enter an existing schedule name or select one from the dropdown menu.
Manage Schedules	Opens the Schedules Inventory (on page 849) in a new tab to enable you to add, edit, review or delete schedules. The times, dates and/or windows within which backups are written to a repository can be controlled by a named schedule that is then selected in the Schedule Name control. Once you have finished working in the Schedules tab, close it and return to the **Replicate Wizard**.

Control	Description
Run on completion of operation	Enabled only when the Run Option is set to Run on completion of operation. Enter an existing operation name or select one from the dropdown menu. This is used when a policy contains more than one operation and it is necessary that the operations occur in a sequential order.
Quiesce configured applications before backup	Only valid when performing a batch replication operation. This is used to ensure backup consistency. If this option is selected and any Application classifiers are included in the policy, then these applications will be temporarily quiesced into a consistent state prior to the snapshot occurring. See About application quiescing (on page 98). **Caution:** ▪ Oracle - Policies containing these application classifications must have the quiesce application option checked, otherwise backups may not be restorable. ▪ VMware - Policies containing this hypervisor classification require that VMware Tools is installed on the VM to perform quiesced backups, otherwise backups will fail.
Pre Script	Specify a script that will be run after replication COPY. This can be used to ensure consistent state for applications other than those with a predefined Application Classification Wizard.
Post Script	Specify a script that will be run after the replication is completed. This is used to ensure consistent state for applications other than those with a predefined Application Classification Wizard.

Snapshot Operation Wizard

This wizard is launched when a new Snapshot operation is added to a Policy.

The Snapshot operation is used to trigger snapshots of entire volumes. Multiple copies may be retained, enabling reversion to a specific point in time.

Caution:

- Snapshot operations always apply to the whole selected volume. Any associated data classifications are ignored.
- Snapshots will be automatically retired after the given retention period.

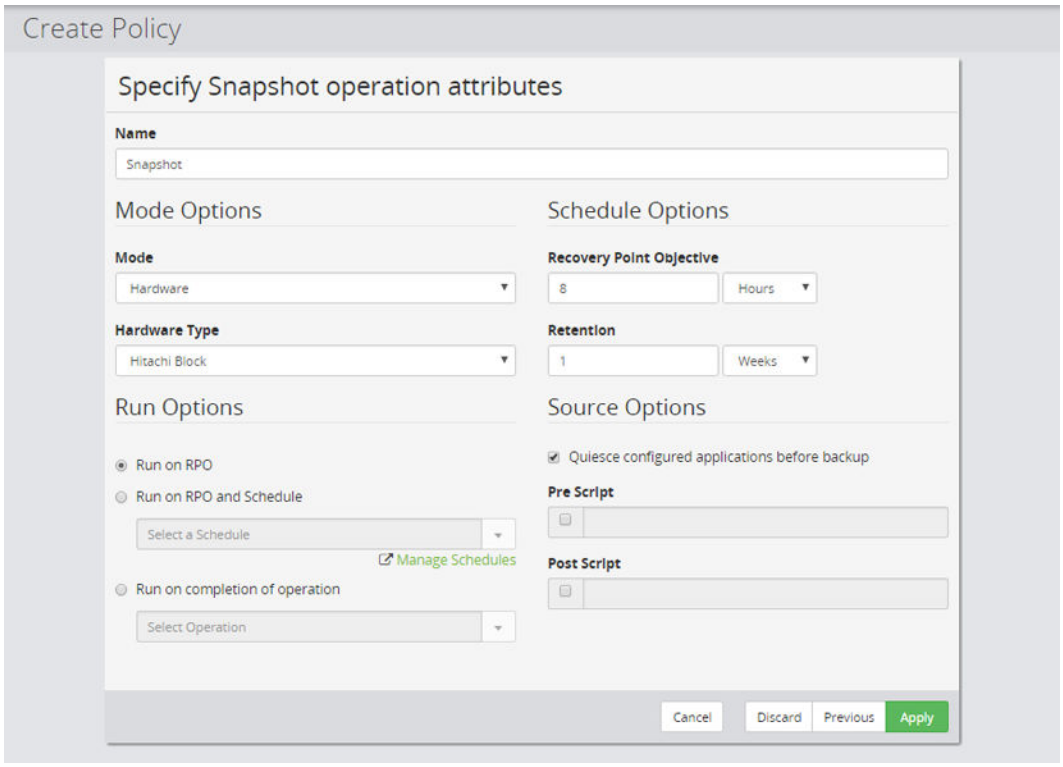


Figure 450 Snapshot Wizard

Control	Description
Name	The name of the operation. The default name is 'Snapshot'. This is used to identify the operation when the policy is used in a Data Flow diagram and when operations are to be triggered sequentially.
Mode	This can currently only be set to: ▪ Hardware - for snapshots of Hitachi Block.
Hardware Type	Only enabled if Mode is set to Hardware. This can be set to: ▪ Hitachi Block - for VSP and HUS VM block storage. ▪ Hitachi Software Block – for Software Defined Storage
Run Options	Select one of the following: ▪ Run on RPO - the operation is triggered whenever the RPO is reached. ▪ Run on RPO and Schedule - the operation will trigger based on the RPO and the schedule name entered in the Select a Schedule field. ▪ Run on completion of operation - the operation will be triggered on completion of the operation entered in the Select Operation field.

Control	Description
Select a Schedule	Enabled only when the Run Option is set to Scheduled. Enter an existing schedule name or select one from the dropdown menu.
Manage Schedules	Opens the Schedules Inventory (on page 849) in a new tab to enable you to add, edit, review or delete schedules. The times, dates and/or windows within which snapshots are created can be controlled by a named schedule that is then selected in the Schedule Name control. Once you have finished working in the Schedules tab, close it and return to the **Snapshot Wizard**.
Select Operation	Enabled only when the Run Option is set to Run on completion of operation. Enter an existing operation name or select one from the dropdown menu. This is used when a policy contains more than one operation and it is necessary that the operations occur in a sequential order.
Recovery Point Objective	This can be set to a specified period and is used in conjunction with the Run Options. If no schedules are defined (i.e. Run on RPO is selected), then a snapshot is created whenever the Recovery Point Objective time is reached.
Retention	The backup is retained for the specified period before being retired.
Quiesce configured applications before backup	This is used to ensure backup consistency. If this option is selected and any Application classifiers are included in the policy, these applications will be temporarily quiesced into a consistent state prior to the snapshot occurring. See About application quiescing (on page 98). **Caution:** ▪ Oracle - Policies containing these application classifications must have the quiesce application option checked, otherwise backups may not be restorable. ▪ VMware - Policies containing this hypervisor classification require that VMware Tools is installed on the VM to perform quiesced backups, otherwise backups will fail.
Pre Script	Specify a script that will be run before the snaphot is taken. This is used to ensure consistent state for applications other than those with a predefined Application Classification Wizard.
Post Script	Specify a script that will be run after the shapshot is taken. This is used to ensure consistent state for applications other than those with a predefined Application Classification Wizard.

Tier Operation Wizard

This wizard is launched when a new Tier operation is added to a Policy.

The Tier operation is used to trigger tiering of entire data from a repository store to a cloud storage platform.

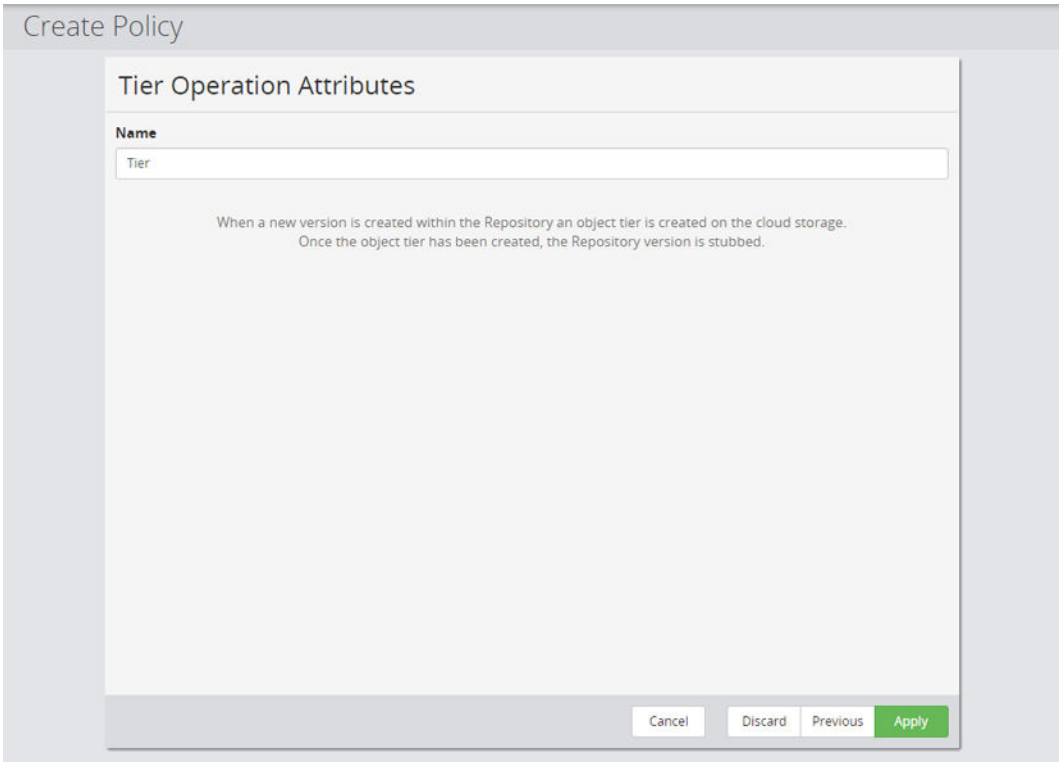


Figure 451 Tier Wizard

Control	Description
Name	The name of the operation. The default name is 'Tier'. This is used to identify the operation when the policy is used in a Data Flow diagram and when operations are to be triggered sequentially.

Classification Filters Wizard

This wizard is launched when Filters for a Path or Disk Type Classification are created or edited.

Classification Filters are applied to Path and Disk Type Classifications when used in conjunction with host based operations (Backup) to provide finer control over what is backed up.

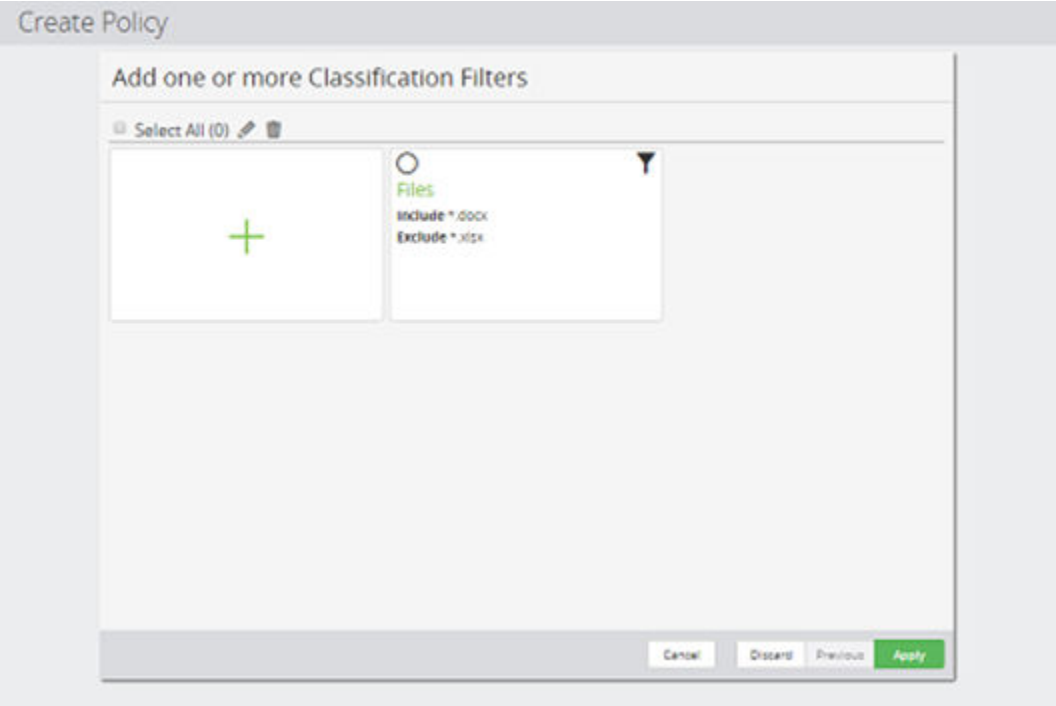


Figure 452 Classification Filter Wizard - Add one or more Classification Filters

Control	Description
Edit	Edits an existing filter in the inventory. The appropriate Filter wizard is launched to enable the filter's attributes to be changed.
Delete	Deletes the selected item from the inventory.
Add	Creates a new Classification Filter. The following wizard page is displayed allowing a Filter to be specified.
Existing Filter	Click on the filter name to Launch the appropriate Filter wizard to enable the filter to be viewed and edited.

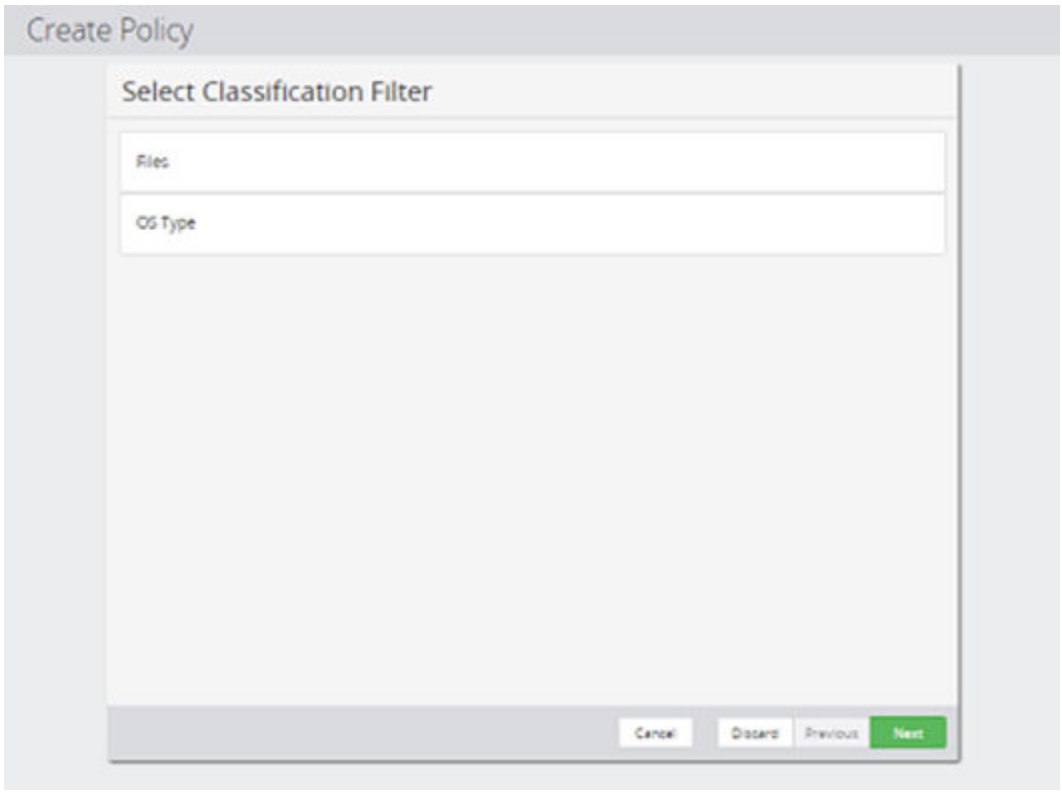


Figure 453 Filter Selection Wizard - Select Classification Filter

Control	Description
Files	Filters based on user defined file type.
OS Type	Filters based on OS Type.

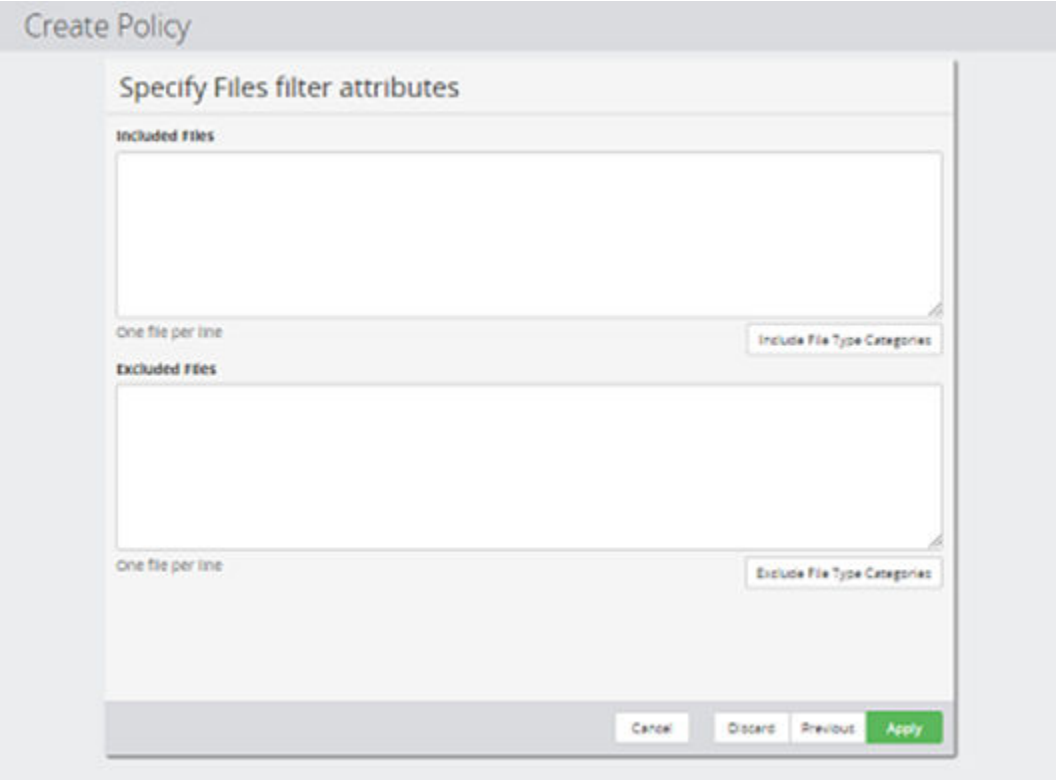


Figure 454 Filter Selection Wizard - Specify Files filter attributes

Control	Description
Include Files	The specified file types will be included in the classification. All files, regardless of path, whose name matches the given name will be included. Specify one file name and extension per line. To specify paths, create a path classifier using the Path Classification Wizard (on page 713) and attach this filter to it.
Include File Type Categories	Launches the Select File type categories dialog (on page 731).
Exclude Files	The specified file types will be excluded from the classification. Specify one file name and extension per line.
Exclude File Type Categories	Launches the Select File type categories dialog (on page 731).

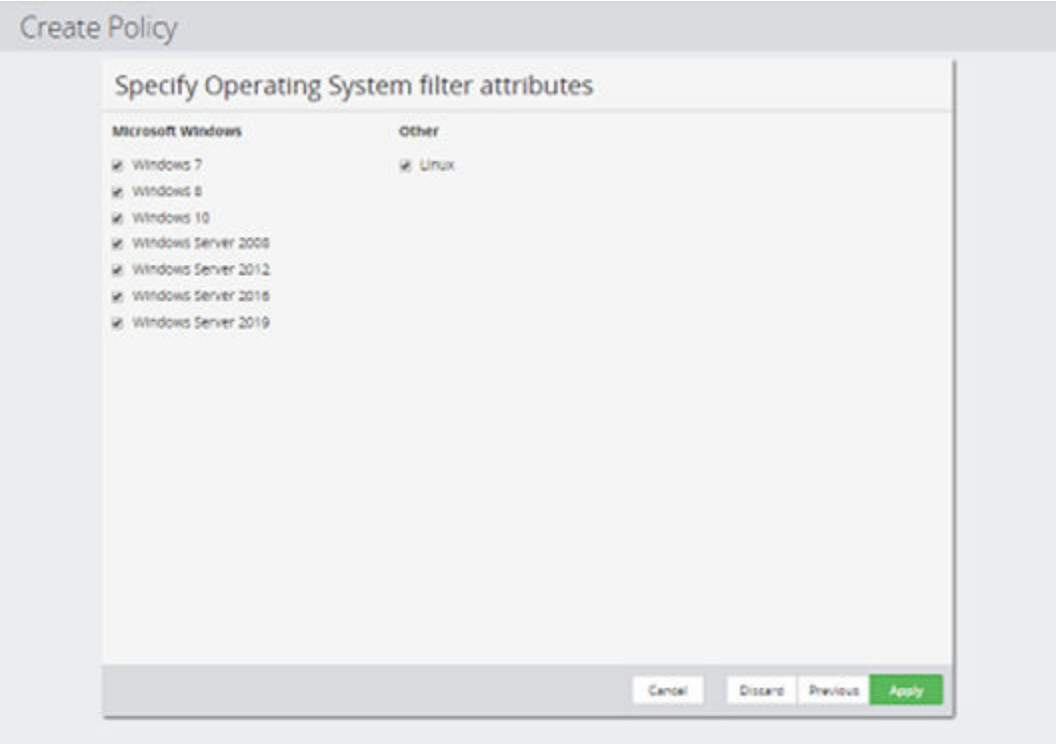


Figure 455 Filter Selection Wizard - Specify Operating System filter attributes

Control	Description
OS Type	Restricts the associated data classification to those machines running one or more of the selected operating system types.

Select File type categories dialog

This dialog is lauched when Include or Exclude file type catagories is selected during classification filter creation or editing.

It facilitates selection of groups of items such as ‘All Images’ or ‘All System’ files with a single selection. Once the selection has been made the list can still be edited in the Classification Filters Wizard (on page 727)

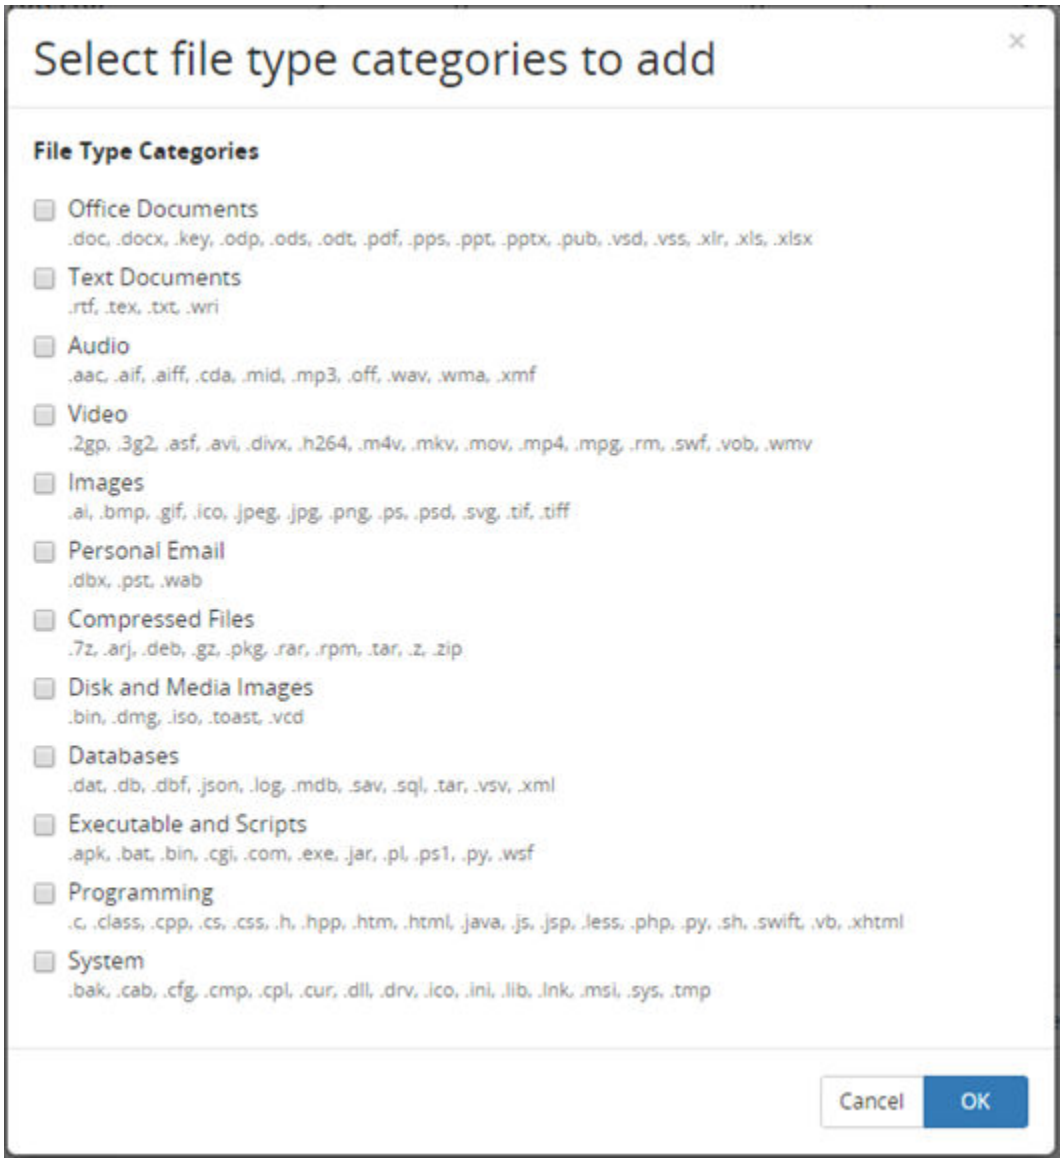


Figure 456 Select file type categories to add

Control	Description
File Type Categories	Select one or more categories to add those to the include or exclude list of the Classification Filters Wizard (on page 727).

Access Operation Wizard

This wizard is launched when a new Access operation is added to a Policy.

The access operation allows a node full or read-only access to Oracle RMAN data on a Gen2 repository or other UBI datastore.

Note: It is not possible to apply this operation to hardware storage devices or Gen1 repositories

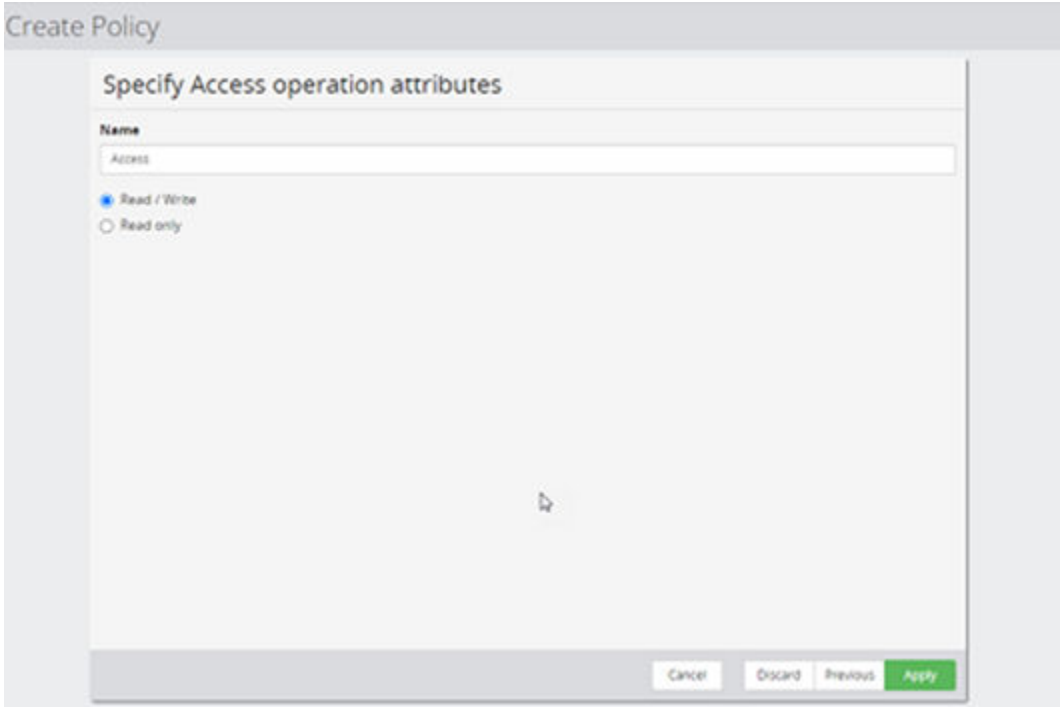


Figure 457 Access Operation Wizard

Control	Description
Name	The name of the operation. The default name is 'Access'. This is used to identify the operation when the policy is used in a Data Flow diagram and when displaying the triggers on the command line.
Read / Write	When selected, Nodes using this operation have full access to the data. It is possible to create, read, change and delete data.
Read only	When selected, Nodes using this operation will only be able to read data created previously or by other nodes. Creation of new data, modifications and deletes will be blocked.

Oracle RMAN Classification Wizard

This wizard is launched when a new Oracle RMAN classification is added to a Policy.

The Oracle RMAN classification allows to conveniently specify, which databases can or cannot access Oracle RMAN data using the access operation.

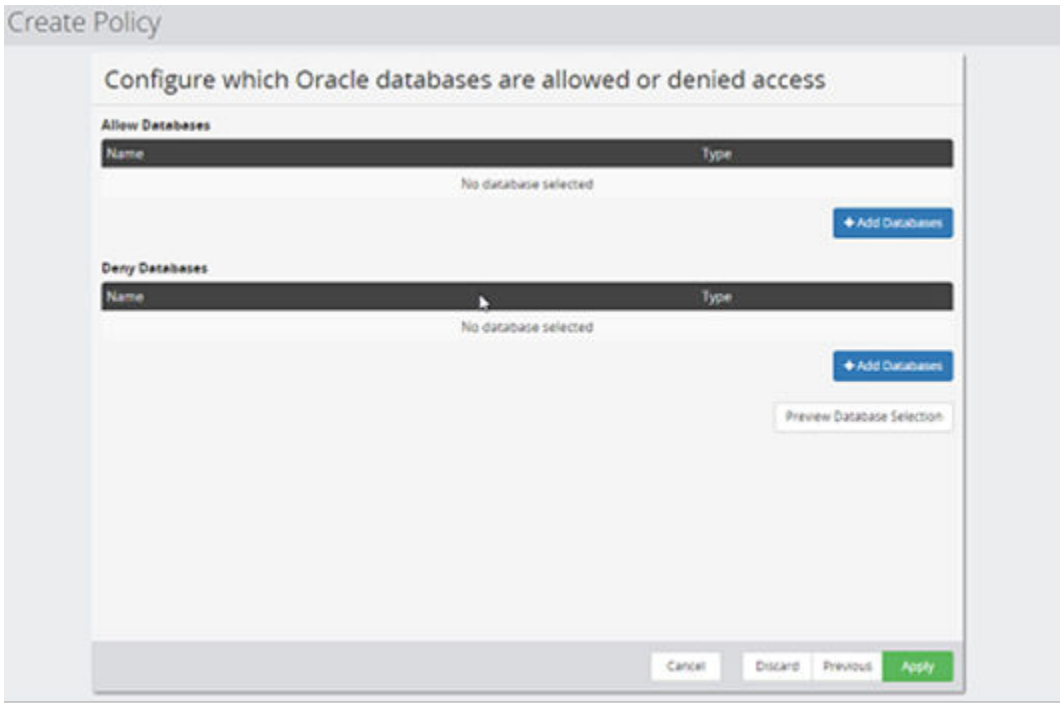


Figure 458 Oracle RMAN Wizard - Specify Oracle Databases (Allowed / Denied Access)

Control	Description
Allow Databases	Lists the databases that will be allowed access for Oracle RMAN access operations.
Add Databases	Opens the Oracle RMAN Database Selection Wizard (on page 690) to enable databases to be added to the **Allow Databases** list above.
Deny Databases	Lists the databases that will not be allowed access for Oracle RMAN access operations.
Add Databases	Opens the Oracle RMAN Database Selection Wizard (on page 690) to enable databases to be added to the "Deny Databases" list above.
Preview Database Selection	Click this button to preview the which databases are allowed access for an existing Oracle Database node.

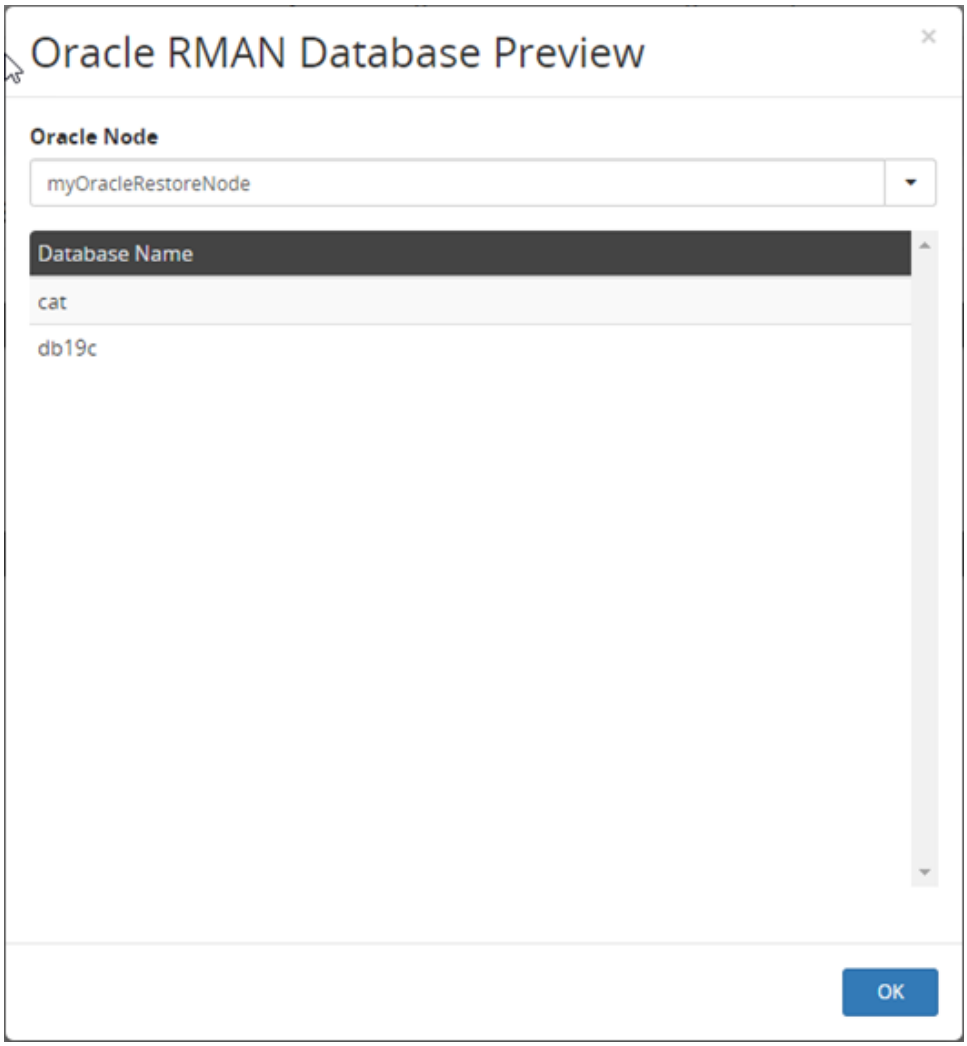


Figure 459 Oracle RMAN Database Preview

Control	Description
Oracle Node	Select a node representing the Oracle setup, which should be previewed.
Database List	List all databases on this node which would be allowed access with the defined classification.

Oracle RMAN Database Selection Wizard

This wizard is launched when a user adds entries to the list of allowed or denied databases in an Oracle RMAN classification.

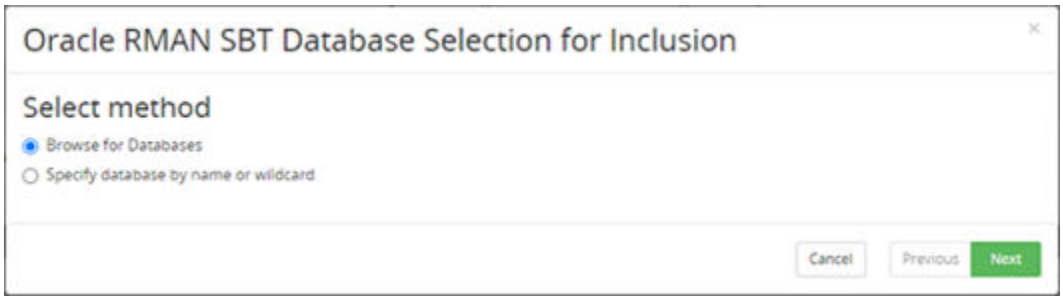


Figure 460 Oracle Database Selection – Select method

Control	Description
Browse for databases	Select this option to browse an existing Oracle node for databases. See **Oracle Database Selection – Browse by** below.
Specify databases by name or wildcard	Select this option specify a database by name pattern match. See **Oracle Database Selection – Specify name or wildcard** below.

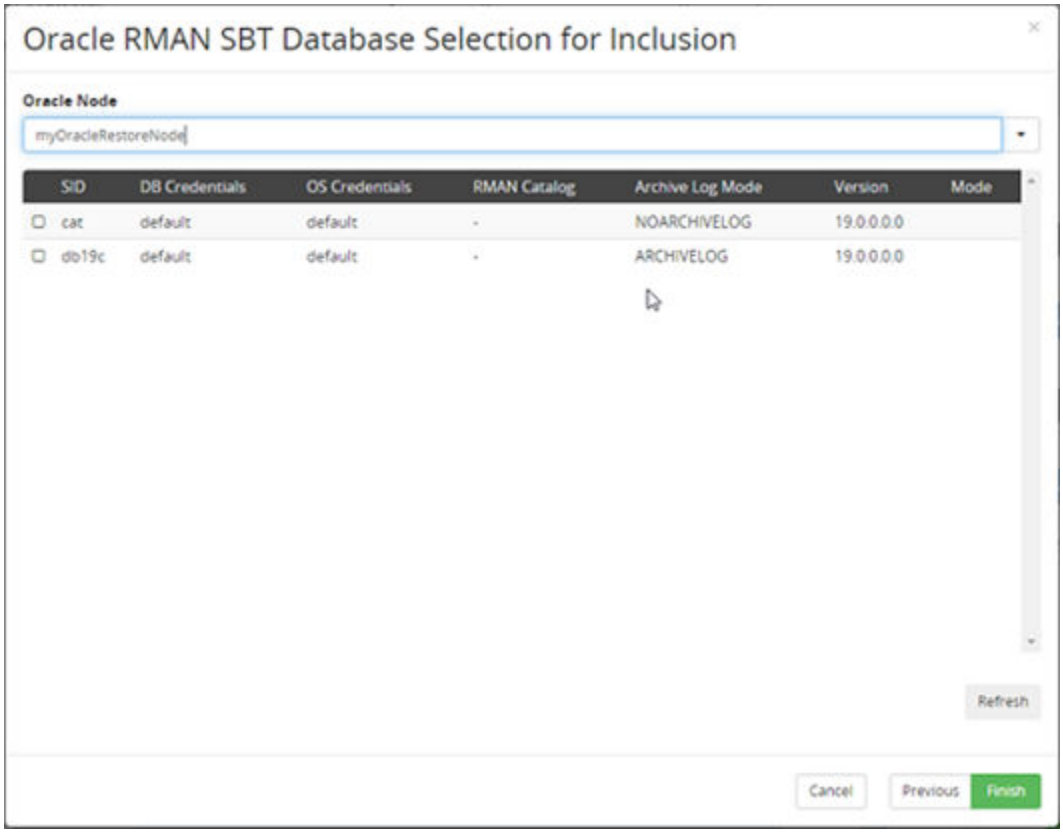


Figure 461 Oracle Database Selection – Browse by

Control	Description
Oracle Node	Select an Oracle database application node to browse for databases.
Database List	Lists the databases which exist on the selected node. You can select one or more databases.
Refresh	Refreshes the list of databases for the selected node. This operation may take a few minutes.

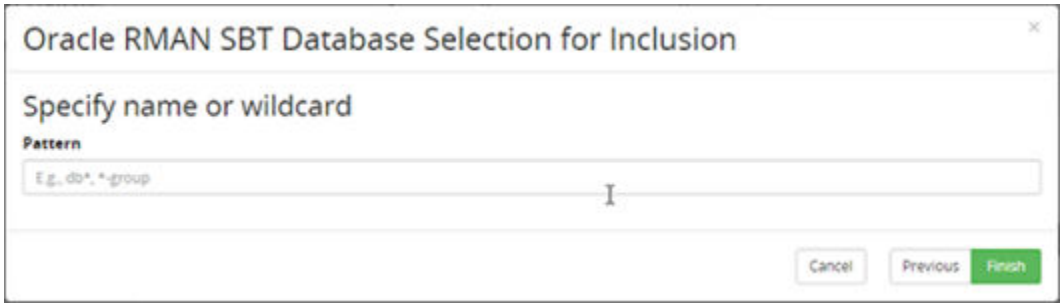


Figure 462 Oracle Database Selection – Specify name or wildcard

Control	Description
Pattern	Enter a case insensitive pattern that will be used to match the database by name. The '*' character can be used to match any sequence of characters. E.g.: IH_* would match any database type whose name begins with IH_. **Note:** Protector evaluates the pattern every time Oracle RMAN tries to access the data. If new databases are added later, they will still be allowed or denied access, depending if they match the pattern or not.

Policy Details

This page displays the details of a Policy and enables you to launch the wizard to edit them.

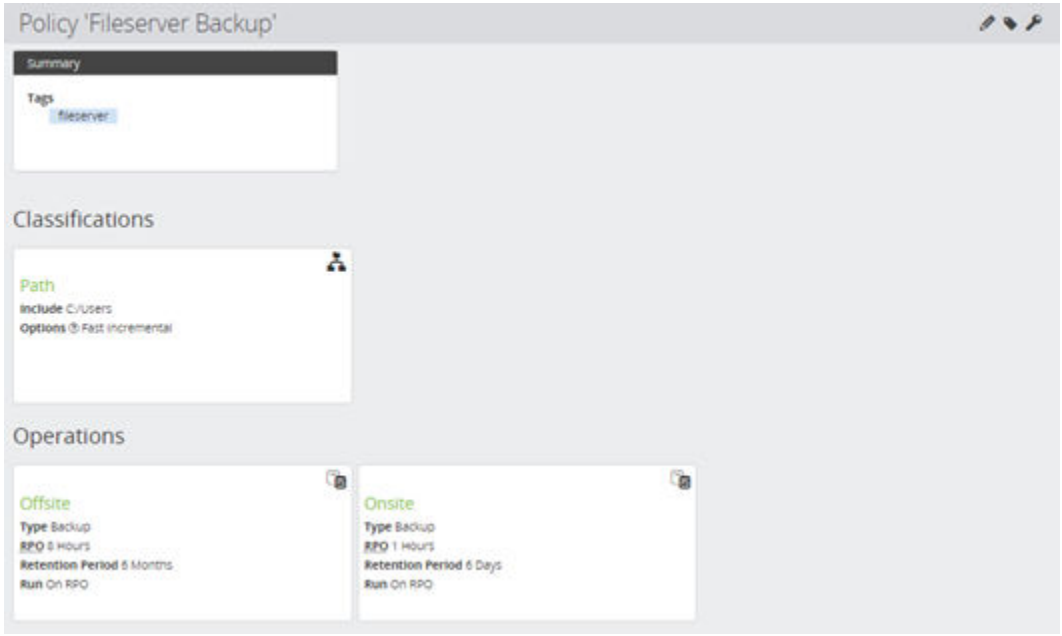


Figure 463 Policy Details

Control	Description
Edit	Launches the Policy Wizard (on page 671) at the Classifications page to enable the policy to be edited.
Tag	Modifies the tags of an existing object from the either the inventory screen or the details screen of the object.
Permissions	Displays the Access Control Permissions Inventory (on page 349) to enable you to view and edit the policy's permissions.
Classification(s)	Lists the classifications for this policy. Click on a classification name to open the Policy Classification Details (on page 738) to enable you to view the classification.
Operation(s)	Lists the operations for this policy. Click on an operation name to open the Policy Operation Details (on page 739) to enable you to view the operation.

Policy Classification Details

The Classification Details page is displayed when the user clicks on a classification in a policy.

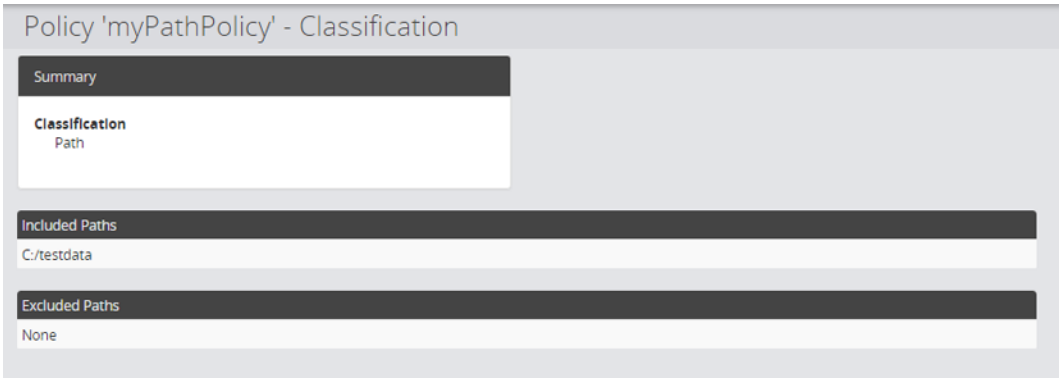


Figure 464 Classification Details - Path Policy

Policy Operation Details

The Operation Details page is displayed when the user clicks on an operation in a policy.

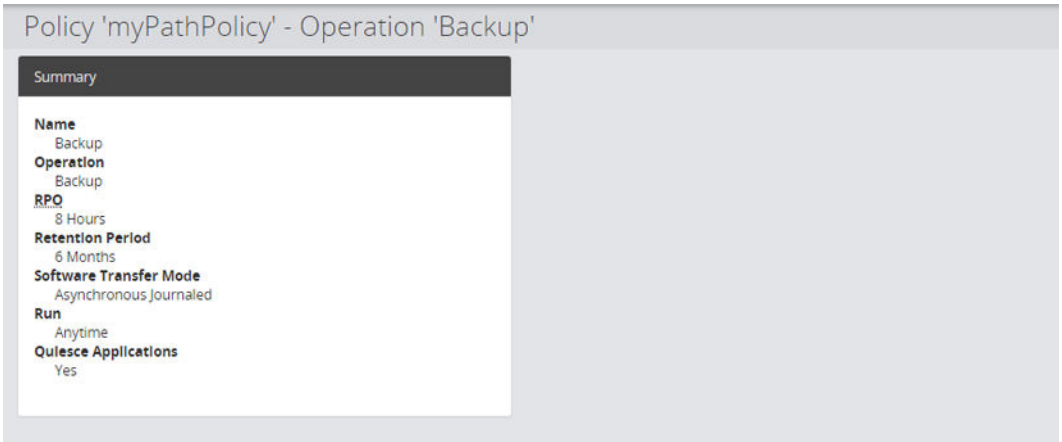


Figure 465 Operation Details - Backup Operation

Reports UI Reference

This section describes the Reports UI, accessed via the Main Banner (on page 286).

For further information, refer to:

- Report Concepts (on page 118)
- Report Tasks (on page 278)

Reports Dashboard

The Reports Dashboard displays the various report types generated by Protector.

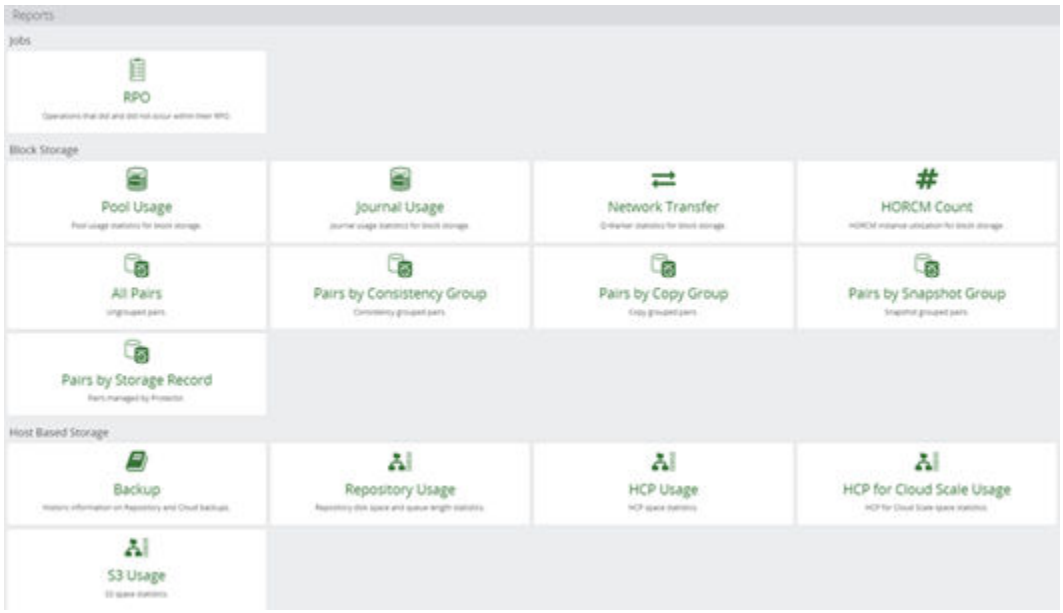


Figure 466 Reports Dashboard

Control	Description
RPO	Select to view the RPO Report (on page 741).
Repository Usage	Select to view the Repository Usage Report (on page 743).
Pool Usage	Select to view the Pool Usage Report (on page 747).
Journal Usage	Select to view the Journal Usage Report (on page 749).
Network Transfer	Select to view the Network Transfer Report (on page 750).
Backup	Select to view the Backup Report (on page 745).
HORCM Count	Select to view the Command Control Interface (CCI) Count Report (on page 752)

Control	Description
 HCP Usage	Select to view the HCP Usage Report (on page 753)
 HCP for cloud scale Usage	Select to view the HCP Cloud Scale Usage Report (on page 754)
 S3 Usage	Select to view the S3 Usage Report (on page 755)
 All Pairs	Select to view the Report Choices (on page 121)
 Pairs by Storage Record	Select to view the Report Choices (on page 121)
 Pairs by Consistency Group	Select to view the Report Choices (on page 121)
 Pairs by Copy Group	Select to view the Report Choices (on page 121)
 Pairs by Snapshot Group	Select to view the Report Choices (on page 121)

RPO Report

The RPO Report page shows the RPO report for all active operations. This can then be filtered using the filter controls.

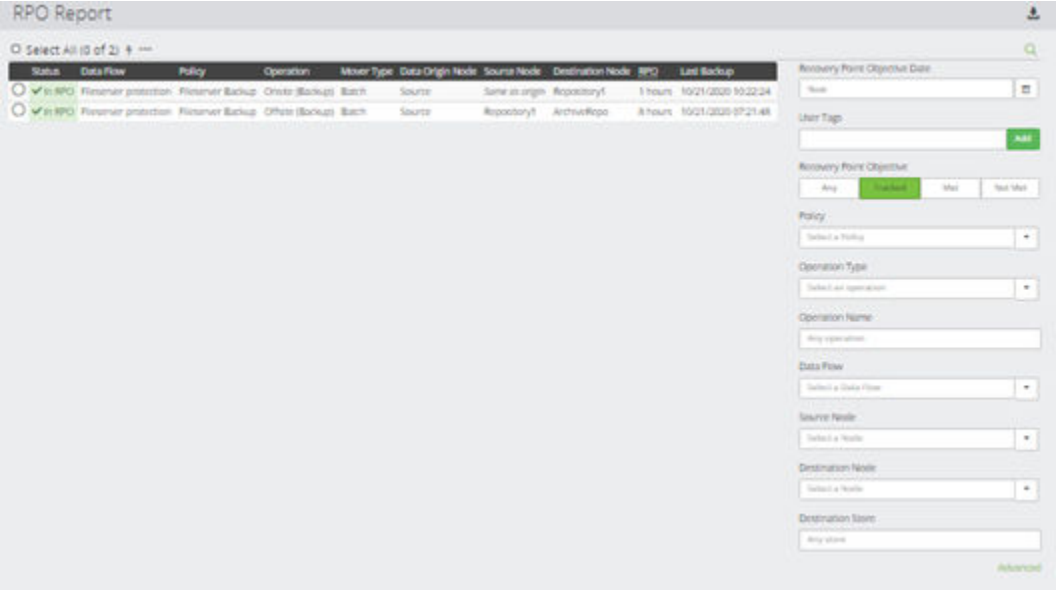


Figure 467 RPO Report

Control	Description
Export	Launches the Export Report Dialog (on page 761) to export the report in the chosen file format.
Trigger Operation	Triggers the selected operations from the table below.
More Actions	Displays a menu that provides access to UI pages relating to the selected RPO report entry.
Status	This table entry indicates when the policy operation is in or out of its RPO. Policy operations using continuous movers are show as not tracked.
Filter on Recovery Point Objective Date	Select a date and time for RPO. The Status column of the table will then indicate whether each policy operation meets that RPO date and time. Select Now for the current date and time.
Filter on User Tags	Filters the displayed results based on Tags.
Filter on Recovery Point Objective	Can be one of: ▪ Any - Shows operations whether they have met their RPO or not. ▪ Tracked - Shows only the operations which are tracked, i.e. filters out continuous hardware operations which do not have an RPO.

Control	Description
	▪ Met - Only shows operations which have met their RPO. ▪ Not Met - Only shows operations which have not met their RPO.
Filter on Policy	Filters the report so that only entries with the specified policy are displayed.
Filter on Operation Type	Filters the report so that only entries with the selected operations are displayed. Can be one or more of: ▪ Access ▪ Backup ▪ Mount ▪ Replicate ▪ Snapshot ▪ Tier
Filter on Operation Name	Filters the report so that only entries with the specified Operation Name are displayed.
Filter on Data Flow	Filters the report so that only entries with the specified data flow are displayed.
Filter on Source Node	Filters the report so that only entries with the specified Source Node name are displayed.
Filter on Destination Node	Filters the report so that only entries with the specified Destination Node name are displayed.
Filter on Destination Store	Filters the report so that only entries with the specified Destination Store are displayed.
Advanced	
Filter on Advanced Query String	Provides advanced search capabilities, allowing users to create filters based on combinations of API properties (see the API reference guide for more details).
Search	Returns the result on Advanced Query String.

Repository Usage Report

This page shows you the details of the Repository Usage.

Note: The repository's size increases by the number of deltas that are retained. For example, if your site completes one backup each day, and that backup is retained for two days with a change rate of 5% per day, then the size of the repository would be 100% plus 5% increase per day.

After a large set of data has been manually deleted, a large number of snapshots have expired, or a number of stores became inactive and were deleted from the repository, you will need to check the repository's status and remove the unused space. See Reducing repository size (on page 960) for more information.

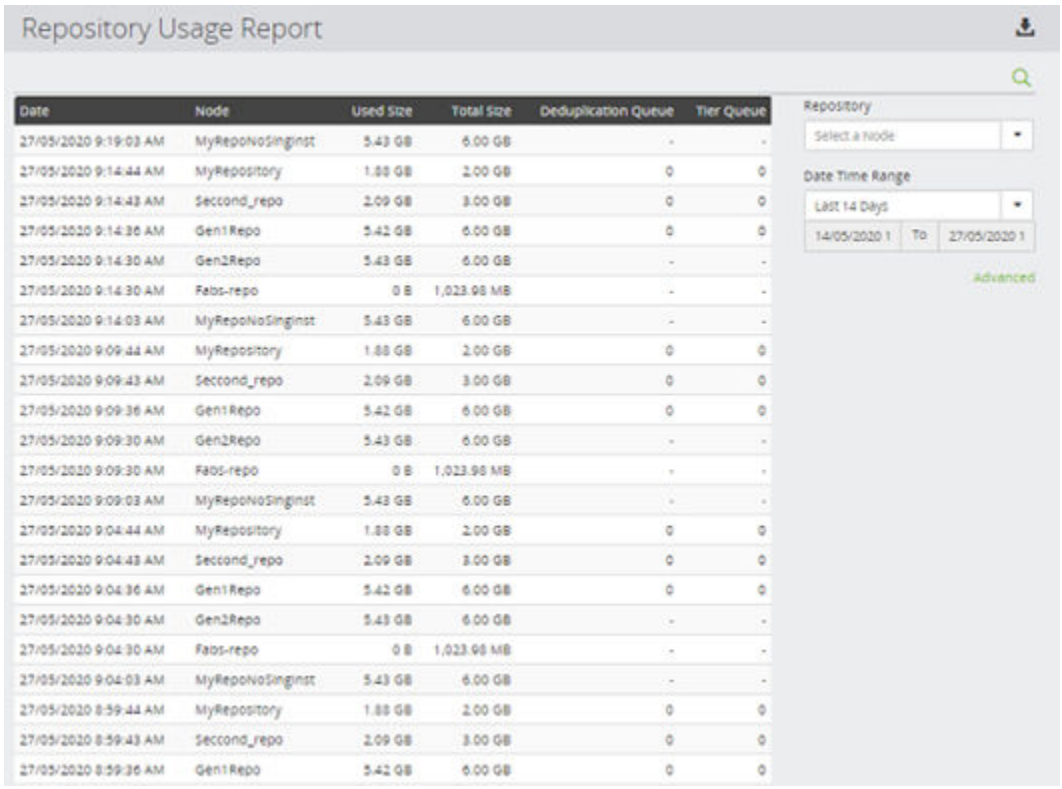


Figure 468 Repository Usage Report

Control	Description
Export	Launches the Export Report Dialog (on page 761) to export the report in the chosen file format.
Used Size	The actual logical size of the repository at any point in time. Typically this value will show rapid initial growth which rapidly reduces as the data change rate falls off.
Total Size	The allocated capacity of a repository. As this space is used thisvalue steps up by preset increments (100Gb by default).The total size will always be greater than or equal to the usedsize. When data is removed from a repository it will not shrink, new data will use the free space in the repoisitory before growing in size.

Control	Description
Deduplication Queue	Shows the frequency and duration of I/O activity for file-level deduplication within the repository.
Tier Queue	Shows the frequency and duration of I/O activity for file-level archiving within the repository.
Filter on Repository	Filters the report so that only entries for the specified repository are displayed.
Filter on Date Range	Filters the all report so that only entries within the specified date range are displayed. Opens the Date Time Range Picker (on page 352).
Advanced	
Filter on Advanced Query String	Provides advanced search capabilities, allowing users to create filters based on combinations of API properties (see the API reference guide for more details).
Search	Returns the result on Advanced Query String.

Backup Report

The Backup Report page shows the full session summary report for all repository, HCP, HCP cloud scale and Amazon S3 based backup sessions.

The Backup report shows a history of all repository, HCP , HCP cloud scale and Amazon S3 based resynchronizations that have taken place within Protector. This list of sessions can be filtered in a similar method to Log Manager to find a specific session or range of sessions. It is also possible to View Session details (as in the Logs Inventory (on page 509)) by selecting an entry in the report.

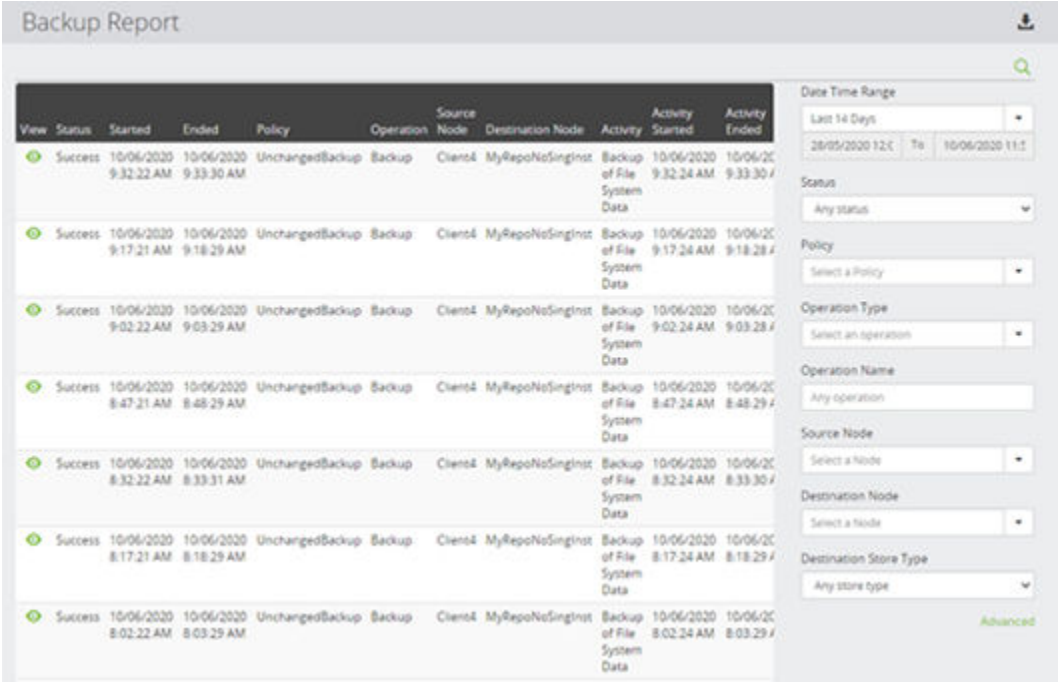


Figure 469 Repository Resynchronizations Report

Control	Description
Export	Launches the Export Report Dialog (on page 761) to export the report in the chosen file format.
View Session	Click on the Session icon to display all log messages that pertain to that session. The Session Log Details (on page 517) is displayed.
Filter on Date Range	Filters the report so that only entries within the specified date range are displayed. Opens the Date Time Range Picker (on page 352).
Filter on Status	Filter the report by the status of the operations listed.
Filter on Policy	Filter the report by the Policy.
Filter on Operation Type	Filters the report so that only entries with the selected Operation Type are displayed.
Filter on Operation Name	Filters the report so that only entries with the specified Operation Name are displayed.
Filter on Source Node	Filters the report so that only entries with the specified Source Node name are displayed.
Filter on Destination Node	Filters the report so that only entries with the specified Destination Node name are displayed.

Control	Description
Filter on Destination Store Type	Filters the report so that only entries with the specified Destination Store Type are displayed.
Advanced	
Filter on Advanced Query String	Provides advanced search capabilities, allowing users to create filters based on combinations of API properties (see the API reference guide for more details).
Search	Returns the result on Advanced Query String.

Pool Usage Report

This page shows you the details of the Pool Usage.

The Block Device Pool Usage report displays a list of Block pools currently containing records associated with Protector, their usage levels and configured warning and high watermark thresholds. Pools that have been allocated for use by Protector are only displayed and updated while snapshots or replications are held within them. A full list of pools is available in the Storage Navigator UI where the warning and high watermark (depletion) thresholds can be configured.

Entries are added to the report every 10 minutes (configurable) or whenever a snapshot or resynchronization event occurs. Used and free values may appear unchanged if the allocation has increased by less than the displayed resolution.

For dynamic pools, if the warning or high water threshold is reached then the entry is highlighted.

For snapshot (Thin Image) pools the warning threshold is not configurable. Snapshot pools will enter an error state when usage reaches a level of 5% below the high water threshold.

Caution: Should pool usage reach capacity, serious error conditions can arise.

For dynamic pools, the replications that use the pool are put into an error state and will refuse to accept additional data. Manual intervention is required before the pool can be used again.

For snapshot pools, the data in the pool will become corrupted, invalidating any and all snapshots stored within (hence a fixed 5% margin is defined to prevent this).

It is highly recommended that email notifications are set up to alert systems administrators when pool threshold events are logged. This can be done in the Hitachi Block Device Monitoring tab of the Notification Wizard (on page 662)

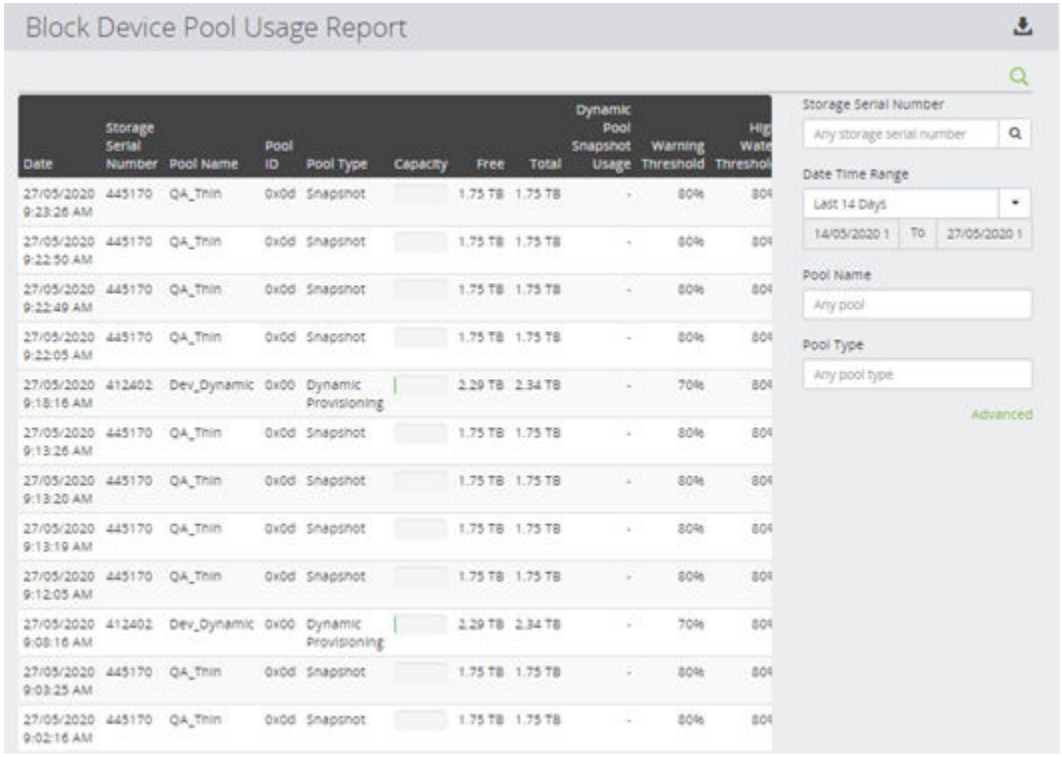


Figure 470 Pool Usage Report

Control	Description
Export	Launches the Export Report Dialog (on page 761) to export the report in the chosen file format.
Filter on Storage Serial Number	Filters the report so that only entries for Storage devices with the specified serial number are displayed.
Filter on Date Range	Filters the report so that only entries within the specified date range are displayed. Opens the Date Time Range Picker (on page 352).
Filter on Pool Name	Filters the report so that only entries for the specified Pool names are displayed.
Filter on Pool Type	Filters the report so that only entries for the specified Pool Types are displayed.
Advanced	
Filter on Advanced Query String	Provides advanced search capabilities, allowing users to create filters based on combinations of API properties (see the API reference guide for more details).
Search	Returns the result on Advanced Query String.

Journal Usage Report

The Block Device Journal Usage Report displays a list of all UR replication policies, showing the source and destination journal usage for both Hardware Block and Hitachi Software Block replications. The list is updated at regular intervals. Search terms can be constructed, based on the field names in the report, using the Filter controls.

Caution: Should journal usage reach capacity, serious error conditions can arise.

Note: You must manually refresh the Journal information from the Hitachi Block Device Details (on page 859) page for each storage array reflected within the report to ensure that the links within the report work for any newly created journals. This is not required for Hitachi Software Block devices

Email notifications may be set up to alert systems administrators when the journal usage rises above a predetermined threshold. This can be done using the Hitachi Block Device Monitoring tab of the Notification Wizard (on page 662).

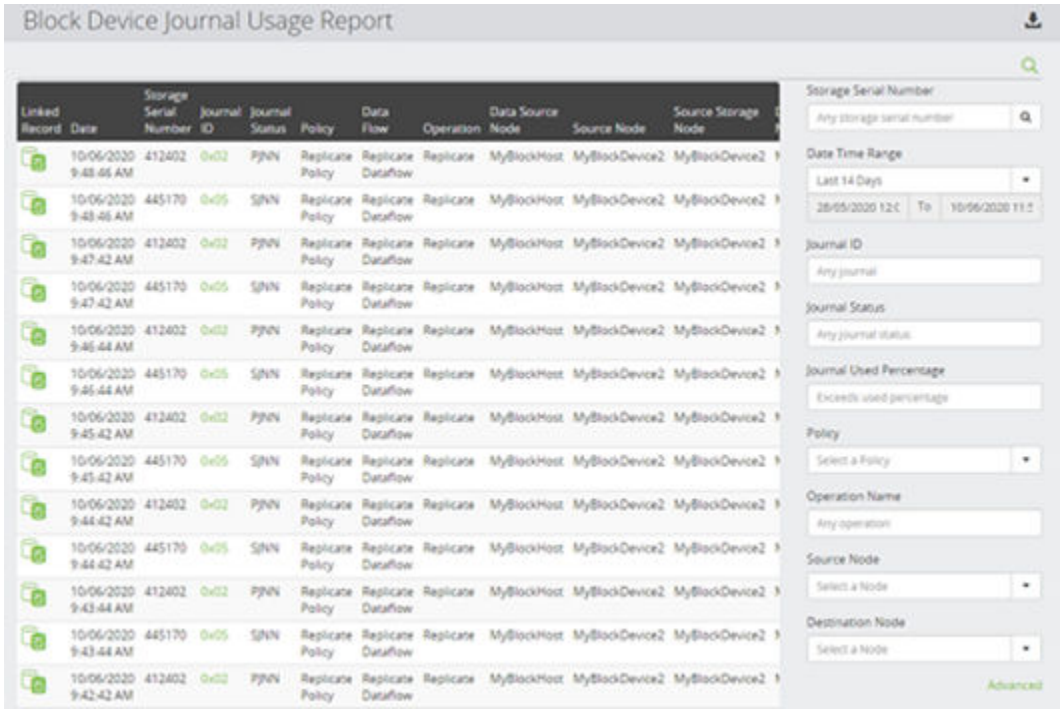


Figure 471 Journal Usage Report

Control	Description
Export	Launches the Export Report Dialog (on page 761) to export the report in the chosen file format.
Filter on Storage Serial Number	Filters the report so that only entries for Storage devices with the specified serial number are displayed.

Control	Description
Filter on Date Range	Filters the report so that only entries within the specified Date Range are displayed. Opens the Date Time Range Picker (on page 352).
Filter on Journal ID	Filters the report so that only entries for the specified Journal ID are displayed.
Filter on Journal Status	Filters the report so that only entries for the specified Journal Status are displayed.
Filter on Journal Used Percentage	Filters the report so that only entries for journals which have exceeded the specified Journal Usage Percentage are displayed.
Filter on Policy	Filters the report so that only entries for the specified Policy are displayed.
Filter on Operation Name	Filters the report so that only entries for the specified Operation Name are displayed.
Filter on Source Node	Filters the report so that only entries for the specified Source Node are displayed.
Filter on Source Storage Node	Filters the report so that only entries for the specified Source Storage Node are displayed.
Filter on Destination Node	Filters the report so that only entries for the specified Destination Node are displayed.
Advanced	
Filter on Advanced Query String	Provides advanced search capabilities, allowing users to create filters based on combinations of API properties (see the API reference guide for more details).
Search	Returns the result on Advanced Query String.

Network Transfer Report

The Network Transfer Time Report displays a list of all UR replication policies, showing the network transfer time (defined as the time taken for the data to be transferred from the source journal to the destination journal). Entries are added to the report every 1 minute. Search terms can be constructed, based on the field names in the report, using the Filter controls.

Email notifications may be set up to alert systems administrators when the network transfer time rises above a predetermined threshold. This can be done using the Hitachi Block Device Monitoring tab of the Notification Wizard (on page 662).

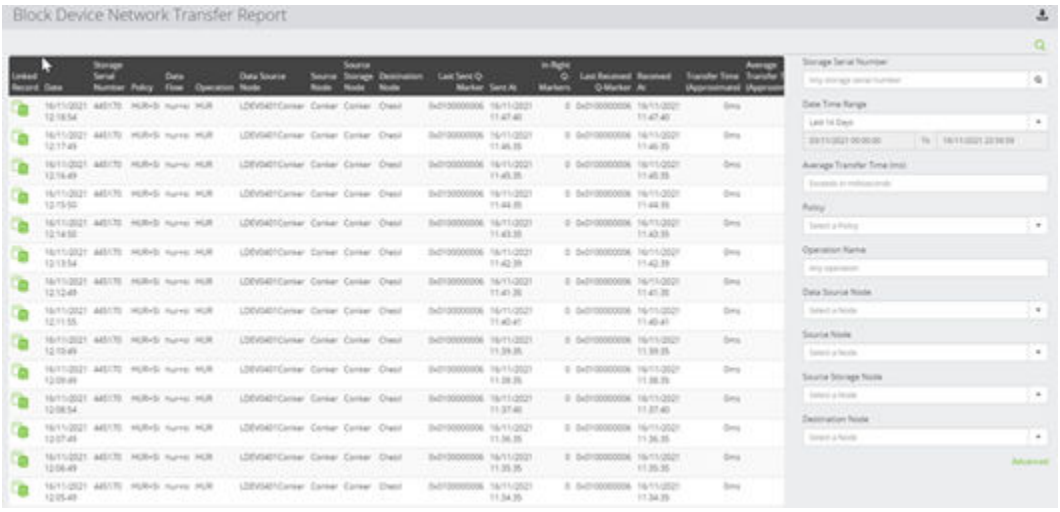


Figure 472 Block Device Network Transfer Report

Control	Description
Export	Launches the Export Report Dialog (on page 761) to export the report in the chosen file format.
Filter on Storage Serial Number	Filters the report so that only entries for Storage devices with the specified serial number are displayed.
Filter on Date Time Range	Filters the report so that only entries within the specified Date Range are displayed. Opens the Date Time Range Picker (on page 352).
Filter on Average Transfer Time	Filters the report so that only entries which have exceeded the specified Average Transfer Time number are displayed.
Filter on Policy	Filters the report so that only entries for the specified Policy are displayed.
Filter on Operation Name	Filters the report so that only entries for the specified Operation Name are displayed.
Filter on Data Source Node	Filters the report so that only entries for the specified Data Source Node are displayed.
Filter on Source Node	Filters the report so that only entries for the specified Source Node are displayed.
Filter on Source Storage Node	Filters the report so that only entries for the specified Source Storage Node are displayed.
Filter on Destination Node	Filters the report so that only entries for the specified Destination Node are displayed.
Advanced	

Control	Description
Filter on Advanced Query String	Provides advanced search capabilities, allowing users to create filters based on combinations of API properties (see the API reference guide for more details).
Search	Returns the result on Advanced Query String.

Command Control Interface (CCI) Count Report

The Command Control Interface (CCI) Count Report displays a count of Command Control Interface (CCI) instances (actual, temporary and current) for each in-band and out-of-band (IP) command device attached to each block storage device. The report is updated every hour (configurable). Search terms can be constructed, based on the field names in the report, using the Filter controls.

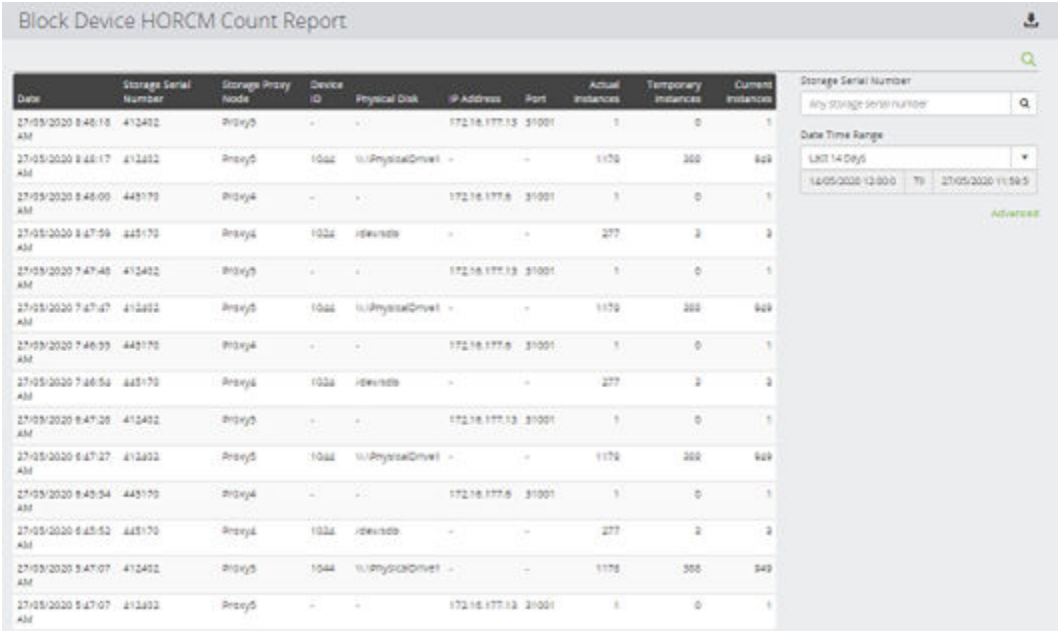


Figure 473 Block Device Command Control Interface (CCI) Count Report

Control	Description
Export	Launches the Export Report Dialog (on page 761) to export the report in the chosen file format.
Filter on Storage Serial Number	Filters the report so that only entries for Storage devices with the specified serial number are displayed.
Filter on Date Time Range	Filters the report so that only entries within the specified Date Range are displayed. Opens the Date Time Range Picker (on page 352).

Control	Description
Advanced	
Filter on Advanced Query String	Provides advanced search capabilities, allowing users to create filters based on combinations of API properties (see the API reference guide for more details).
Search	Returns the result on Advanced Query String.

HCP Usage Report

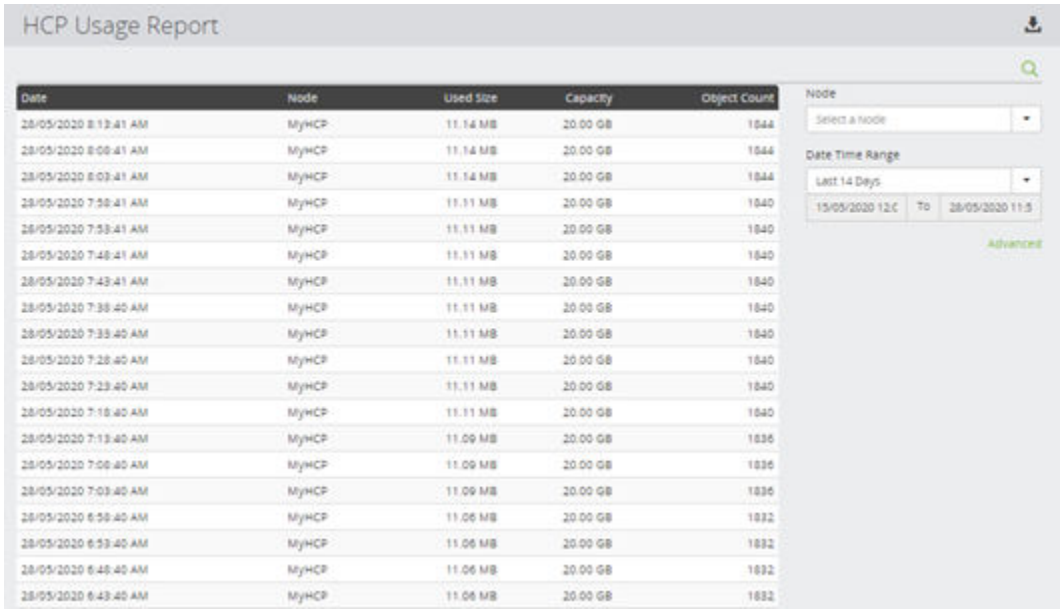


Figure 474 HCP Usage Report

Control	Description
Export	Launches the Export Report Dialog (on page 761) to export the report in the chosen file format.
Filter on Node	Filters the report so that only entries for the specified Node are displayed.
Filter on Date Time Range	Filters the report so that only entries within the specified date range are displayed. Opens the Date Time Picker (on page 351).
Advanced	

Control	Description
Filter on Advanced Query String	Provides advanced search capabilities, allowing users to create filters based on combinations of API properties (see the API reference guide for more details).
Search	Returns the result on Advanced Query String.

HCP Cloud Scale Usage Report

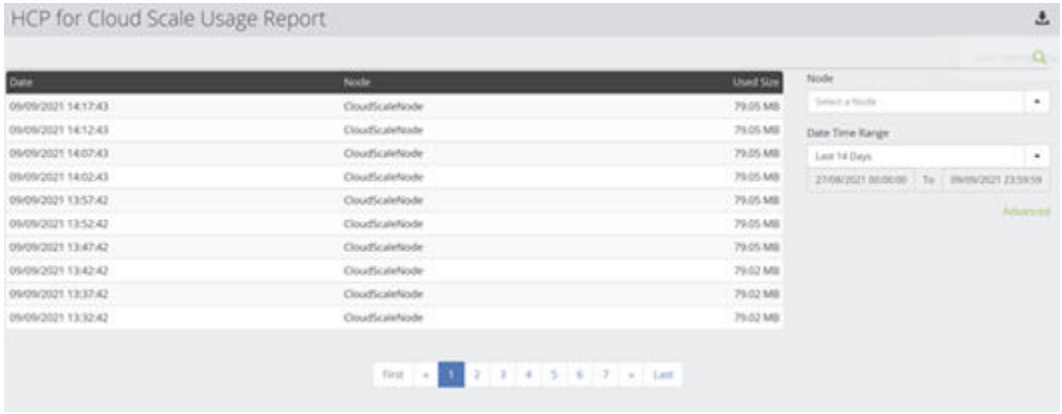


Figure 475 HCP cloud scale usage report

Control	Description
Export	Launches the Export Report Dialog (on page 761) to export the report in the chosen file format.
Filter on Node	Filters the report so that only entries for the specified Node are displayed.
Filter on Date Time Range	Filters the report so that only entries within the specified date range are displayed. Opens the Date Time Picker (on page 351).
Advanced	
Filter on Advanced Query String	Provides advanced search capabilities, allowing users to create filters based on combinations of API properties (see the API reference guide for more details).
Search	Returns the result on Advanced Query String.

S3 Usage Report

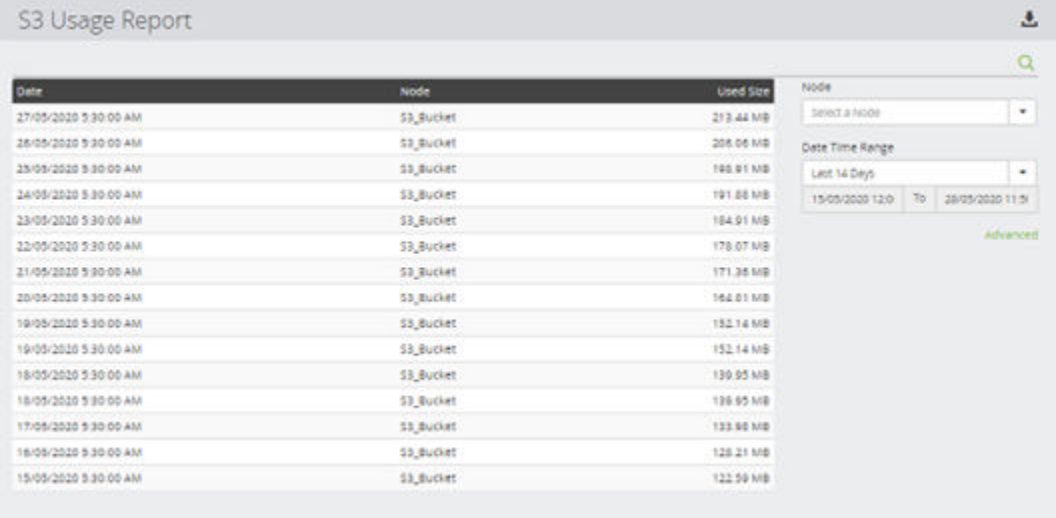


Figure 476 S3 Usage Report

Control	Description
Export	Launches the Export Report Dialog (on page 761) to export the report in the chosen file format.
Filter on Node	Filters the report so that only entries for the specified Node are displayed.
Filter on Date Time Range	Filtersthe report so that only entries within the specified date range are displayed. Opens the Date Time Picker (on page 351).
Advanced	
Filter on Advanced Query String	Provides advanced search capabilities, allowing users to create filters based on combinations of API properties (see the API reference guide for more details).
Search	Returns the result on Advanced Query String.

Reposotory Usage Report

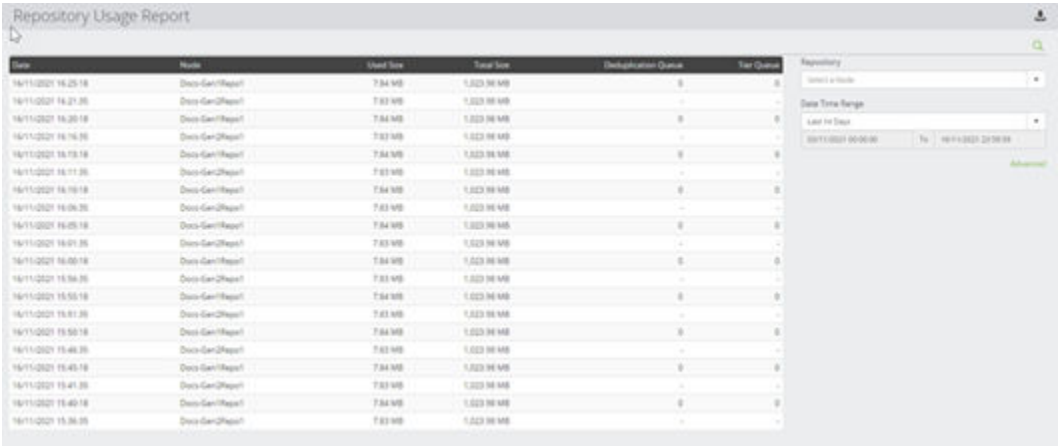


Figure 477 Repository Usage Report

Control	Description
Export	Launches the Export Report Dialog (on page 761) to export the report in the chosen file format.
Filter on Repository	Filters the report so that only entries for the specified Repository are displayed.
Filter on Date Time Range	Filtersthe report so that only entries within the specified date range are displayed. Opens the Date Time Picker (on page 351).
Advanced	
Filter on Advanced Query String	Provides advanced search capabilities, allowing users to create filters based on combinations of API properties (see the API reference guide for more details).
Search	Returns the result on Advanced Query String.

Global Replication Reports

The reports that make up the Global replication reports are all very similar in style only varying by the columns of data. . The Global Replication reports are made up of the following 5 Block Storage Reports



Figure 478 Global Replication Reports

Each Global Replication report provides similar flow as described below.

Each report has a condensed or expanded view. Expanded view will display all the fields for a replication / snapshot while condensed view will display the more commonly sought after information. Each report displays custom fields that are pertinent to that report.



Figure 479 Global Replication Report: View Report

In many columns for a given row in the replication reports, there are two values, one above the other. The first value represents the entry for the primary, the second represents the entry the secondary. So for instance, for the column "Array", the first entry is the array that the PVOLs resides, and the second entry is the array which hosts the SVOLs. Note that this is how the array interprets the source and destination volumes which may differ from Protector which always maintains the source and destination as the way it is drawn on the dataflow.



Figure 480 Global Replication Report

If the replication or the storage is one that is managed by Protector then additional actions are available. This is will be evident by the vertical expansion control at the end of the row. Selecting this control will provide the user with further options allow them to view with Protector:

- **Storage Record** - This will take you to the storage record for that replication / Snapshot. This is only available if the replication is defined in a Dataflow.
- **Primary Storage** - This will take you to the Storage Node details screen for that replication's primary storage. This is only available if that storage was configured in Protector.
- **Secondary Storage** -This will take you to the Storage Node details screen for that replication's secondary storage. This is only available if that storage was configured in Protector.
- **Policy** - This will take you to the Policy definition for that replication . This is only available if the replication is defined in a Dataflow.
- **Dataflow** - This will take you to the Dataflow configuration for that replication. This is only available if the replication is defined in a Dataflow.
- **Monitor** - This will take to the Monitor screen for that replication. This is only available if the replication is defined in a Dataflow.

Report Filters

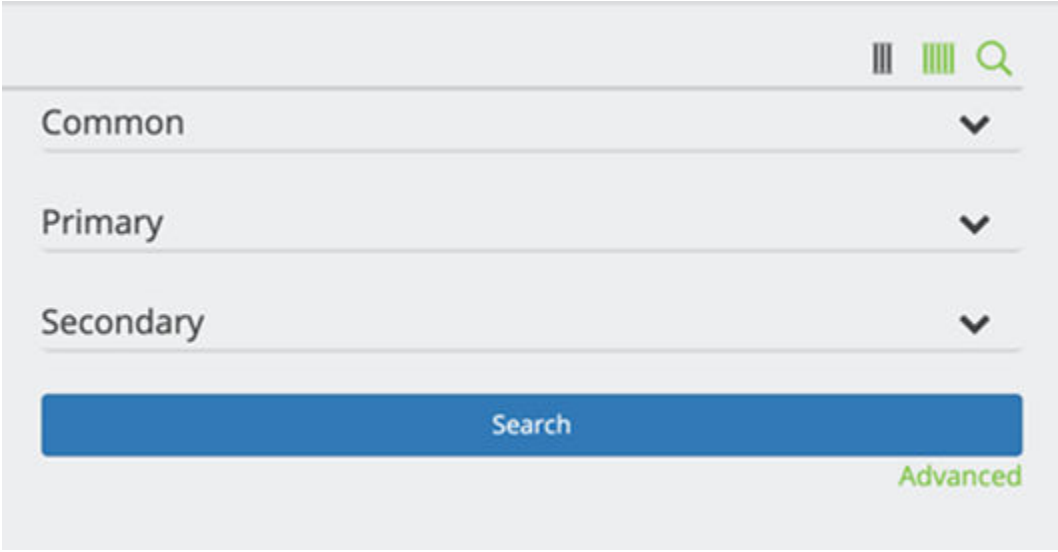


Figure 481 Report Filter

Each report comes with multiple filters that can help to manage the size of the report and be targeted to the users needs. The filters are grouped into three sections: Common, Primary, Secondary. These groups can be expanded to expose the filter attributes for that section . Common filter refers to attributes that are common to both the Primary and the Secondary side of a replication. Primary filter are attributes specific to the Primary side of a replication. Likewise, Secondary filter are attributes specific to the Secondary side of a replication. The filters can be used in conjunction to provide more targeted reports.

For Example:

- To view all replications going from one array to another, enter the primary array serial number in the Primary section and the secondary array serial number in the Secondary section.
- To view all replications coming into or going out of an array, enter the array serial number in the Common section.

Common filter attributes

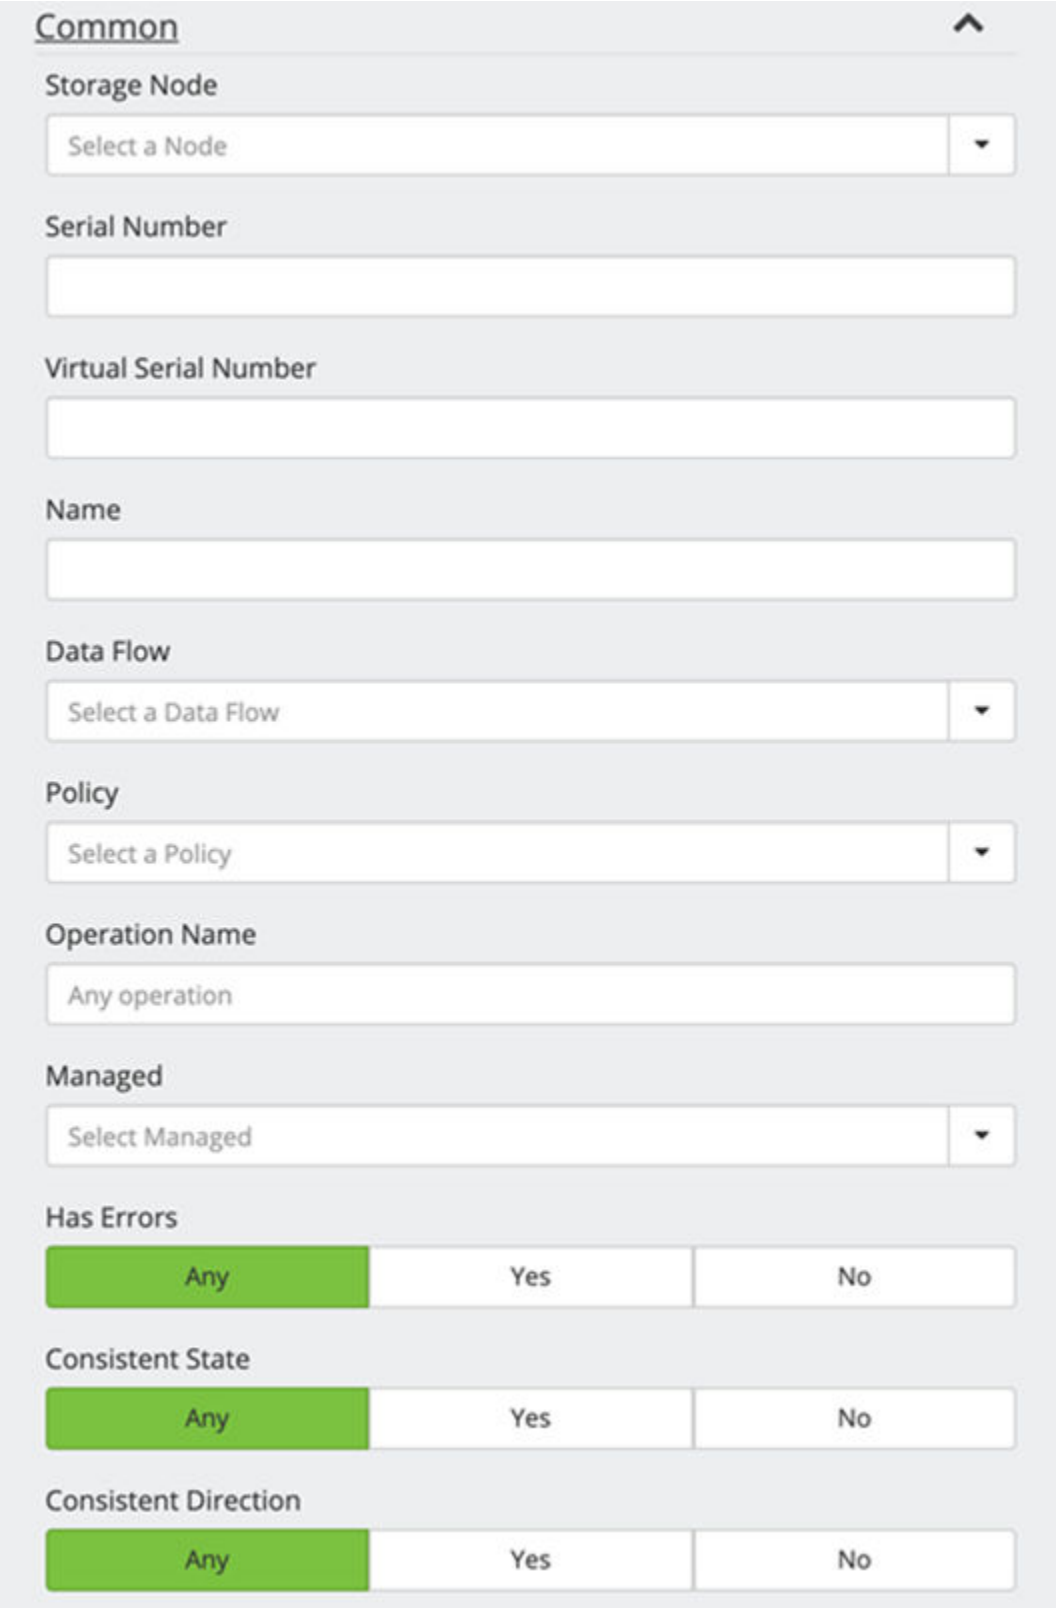


Figure 482 Common filter attributes

Primary and Secondary filter attributes

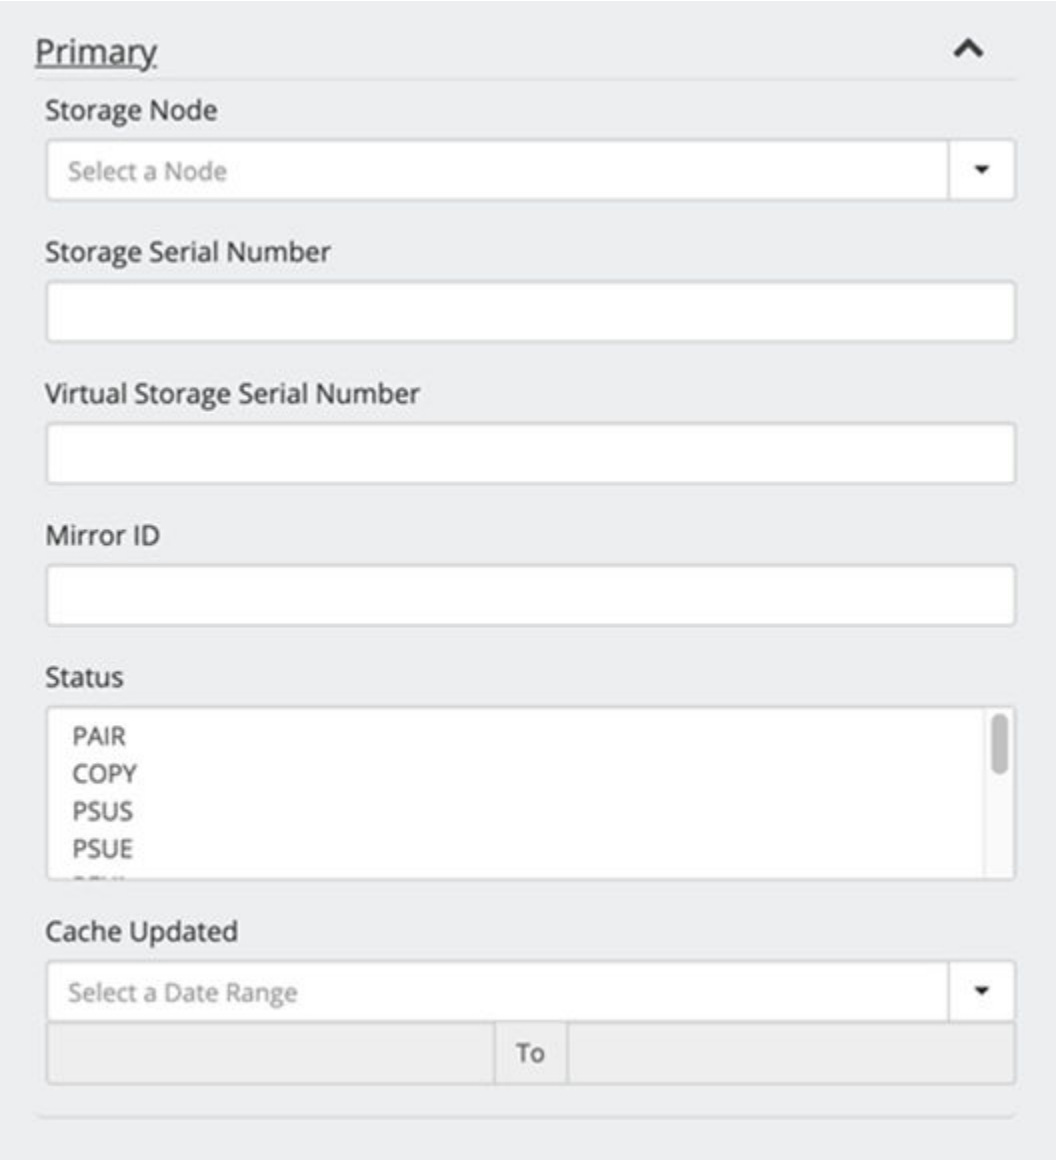


Figure 483 Primary and Secondary filter attributes

Export Report Dialog

This dialog is displayed when downloading a report.

Note:

- The report will be downloaded to a location determined by the web browser's settings.
- Any filtering applied when viewing the report will also be applied to the exported report.
- Exporting reports in HTML and PDF will have times exported in the same time zone as the Master node. Exporting in CSV or JSON will result in times being exported in UTC.



Figure 484 Export Report Dialog

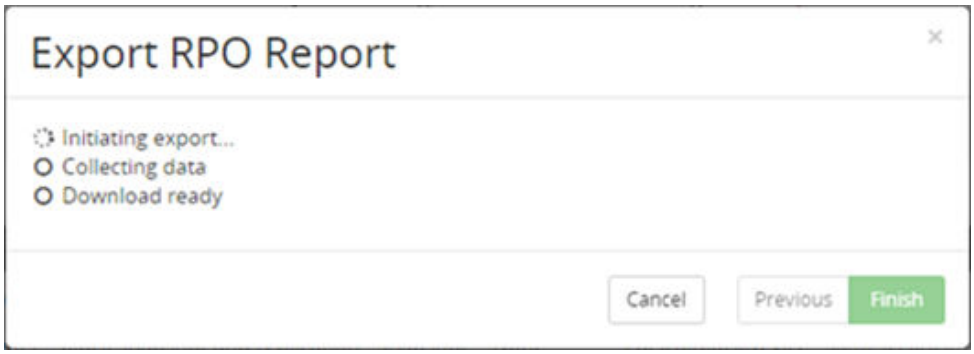


Figure 485 Export Report Dialog – Export in progress

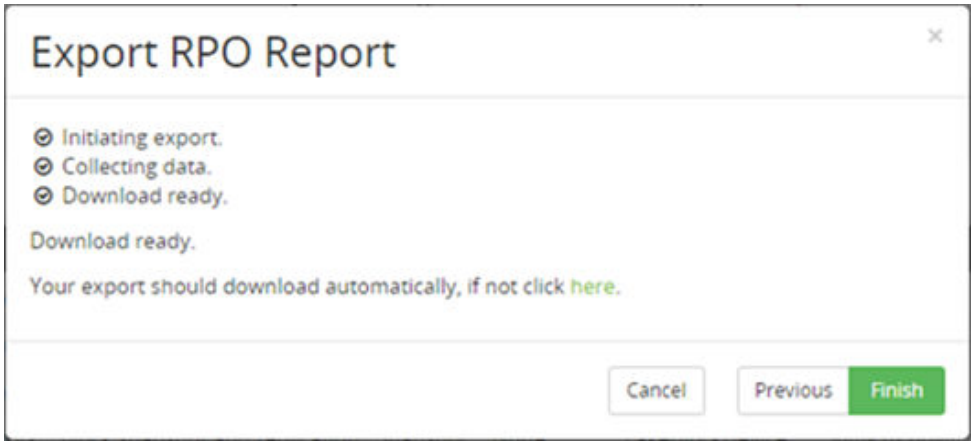


Figure 486 Export Report Dialog – Export completed

Control	Description
Export Format	Select the desired format to export the report in. Can be one of: ▪ HTML ▪ PDF

Control	Description
	▪ CSV ▪ JSON
Export	Click to begin the export process. The dialog will display progress of the export process and the export file will automatically be downloaded when ready. The export will be automatically downloaded even if the dialog is closed.

Restore UI Reference

This section describes the Restore UI, accessed via the Navigation Sidebar (on page 291).

For further information, refer to:

- Restore Concepts (on page 110)
- Restore Tasks (on page 279)

Restore Inventory

This inventory lists all the available backups from all storage types, allowing the user to select one or more for restoring.

Note: Nothing is displayed in the results area until a search has been performed.

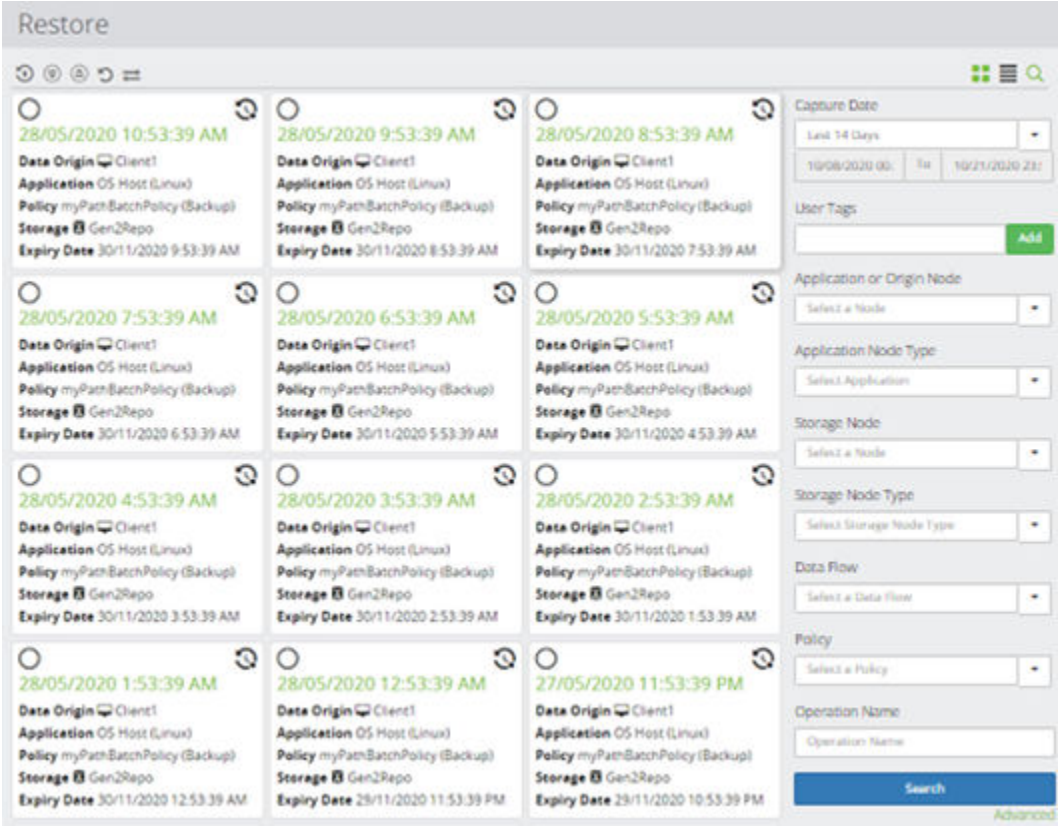


Figure 487 Restore All Inventory

Control	Description
Restore	Enabled when a Repository, HCP or S3 snapshot is selected. Refer to the Host Based Backup Restore Options (on page 826) section.
Mount	Enabled only when an Hitachi Block snapshot is selected. Refer to the Block Restore Options (on page 766) .
Unmount	Enabled only when an Hitachi Block snapshot is selected. Refer to the Block Restore Options (on page 766) section.
Revert	Enabled only when an Hitachi Block snapshot is selected. Refer to the Block Restore Options (on page 766) section.
Swap	Enabled only when an Hitachi Block replication is selected. Refer to the Block Restore Options (on page 766).
Repository Snapshot Tile	Refer to the Host Based Backup Restore Options (on page 826).

Control	Description
 Hitachi Block Snapshot Tile	Refer to the Block Restore Options (on page 766).
 Hitachi Block Replication Tile	Refer to the Block Restore Options (on page 766).
 HCP Snapshot Tile	Refer to the Host Based Backup Restore Options (on page 826) section
 S3 Snapshot Tile	Refer to the Host Based Backup Restore Options (on page 826) section
Filter on Capture Date	Filters the displayed results based on the Date Range within which the backup was created. Opens the Date Time Picker (on page 351)
Filter on User Tags	Filters the displayed results based on Tags.
Filter on Application or Origin Node	Filters the displayed results based on the Source node or Application node from which the backup originated.
Filter on Application Node Type	Filters the displayed results based on the Application node type from which the backup originated.
Filter on Storage Node	Filters the displayed results based on the Storage node where the backup is held.
Filter on Storage Node Type	Filters the displayed results based on the Storage node type where the backup is held.
Filter on Data Flow	Filters the displayed results based on the Data Flow that created the backup.
Filter on Policy	Filters the displayed results based on the Policy that created the backup.
Filter on Operation Name	Filters the displayed results based on the Policy that created the backup.
Search	Click this button to initiate the search based on the filter criteria entered above.

Block Restore Options

This section describes the Block Restore options, accessed via the Restore Inventory (on page 763). Refer to the following sections for details on how to restore the various block based protections operations.

Hitachi Block VMware Snapshot Restore Wizard

This wizard is displayed when you restore a VMware snapshot from a Hitachi Block device.

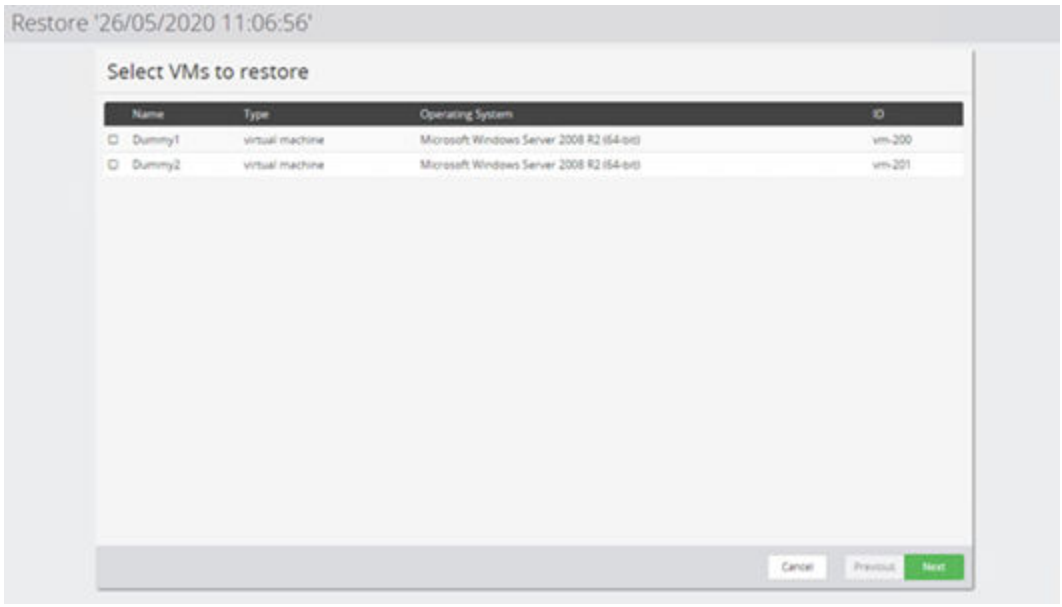


Figure 488 Restore VMware Snapshot Wizard - Select VMs to Restore

Control	Description
VMs in snapshot	Select the specific VMs within this snapshot that are to be restored.

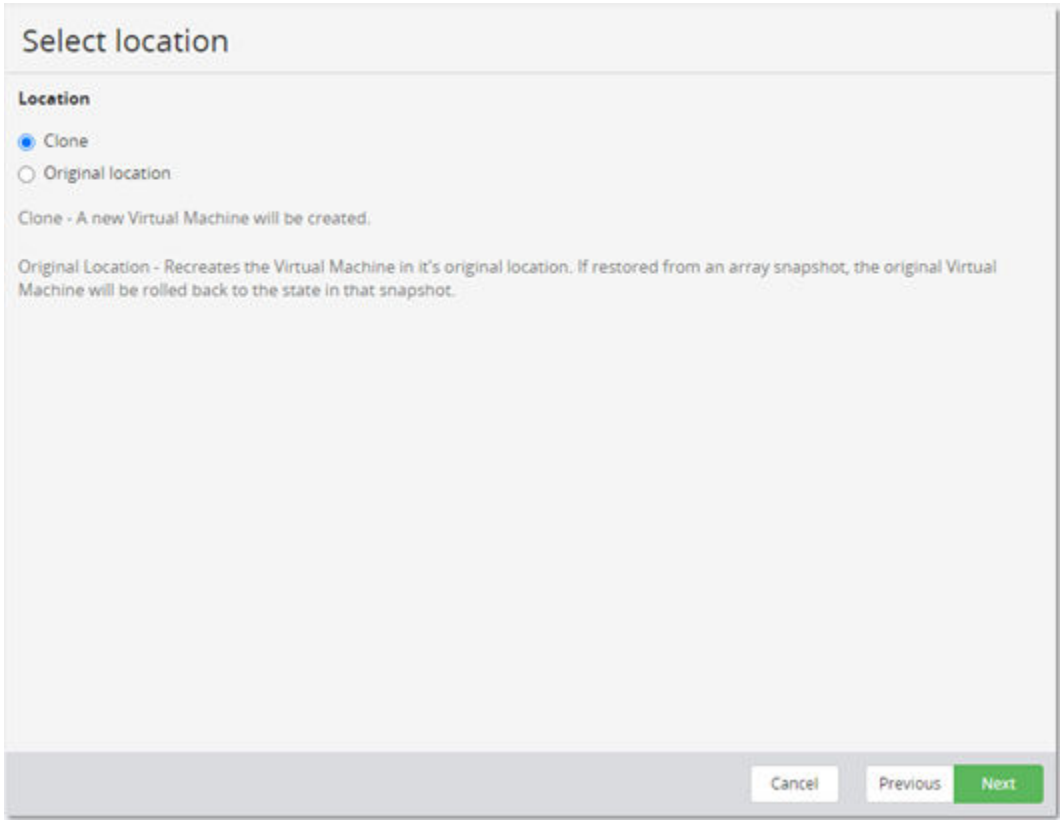


Figure 489 Restore VMware Snapshot Wizard - Select Location

Control	Description
Original location	If the Virtual Machine to be restored does not exist, it will be recreated. If the Virtual Machine to be restored does exist, it will first be backed up on the datastore and then rolled back to the state of the snapshot. **Note:** If you want to replace the existing VM with the restored one, then delete it before restoring.
Clone	The backup will be restored as a clone at the specified location. The wizard displays the **Set clone prefix and destination** page when Next is clicked. **Note:** A restored Virtual Machine will not overwrite any existing machines with the same name in the same location. If a VM of the same name exists at the restore location then the restore job will fail and log and error to that effect.

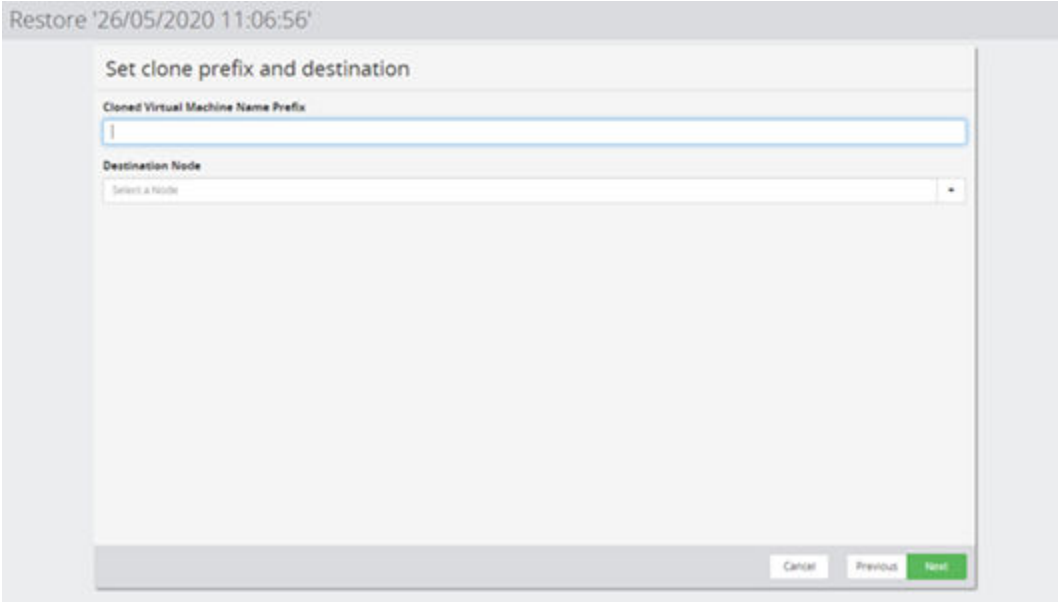


Figure 490 Restore VMware Snapshot Wizard - Clone Prefix and Destination

Control	Description
Cloned Virtual Machine Name Prefix	A prefix for the name(s) of the cloned VM(s) must be specified. If the resulting prefixed name is already used by an existing VM in the restore location then the restore will fail and an error will be logged.
Destination Node	Select the VMware Host or vCenter where the cloned VM(s) will be restored.

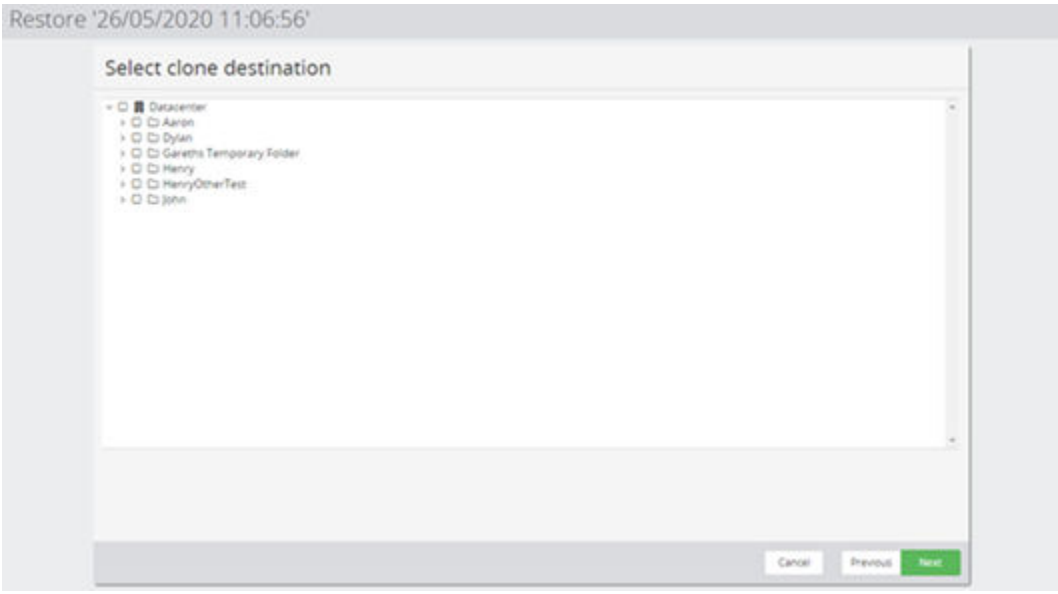


Figure 491 Restore VMware Snapshot Wizard - Select Clone Destination

Control	Description
Destination	Select the VMware Datacenter and sub-folder where the cloned VM(s) will be restored.

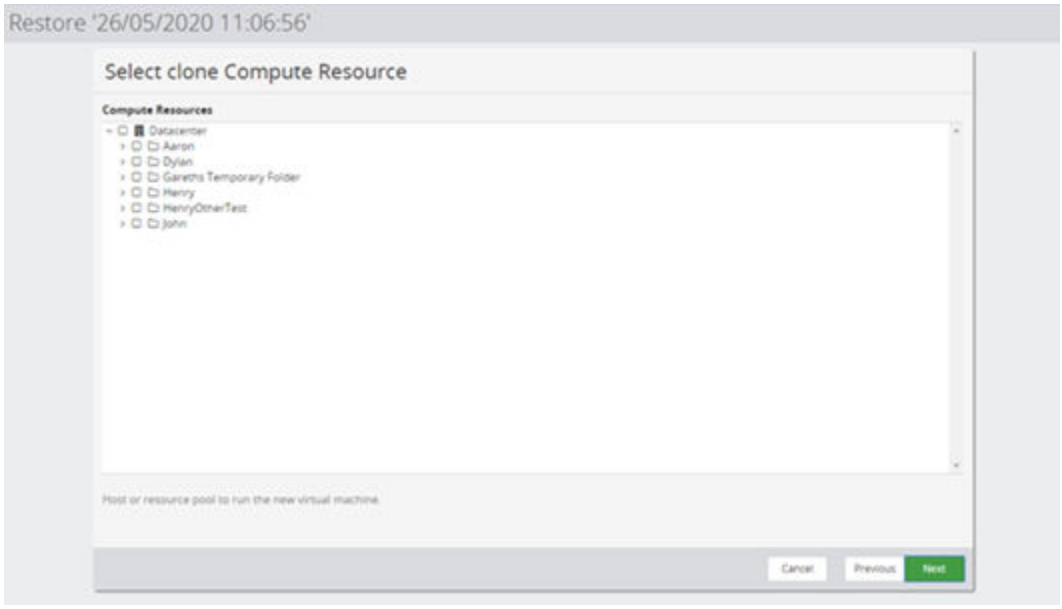


Figure 492 Restore VMware Snapshot Wizard - Select Clone Compute Resource

Control	Description
Compute Resources	Select the VMware Compute Resource where the cloned VM(s) will be run.

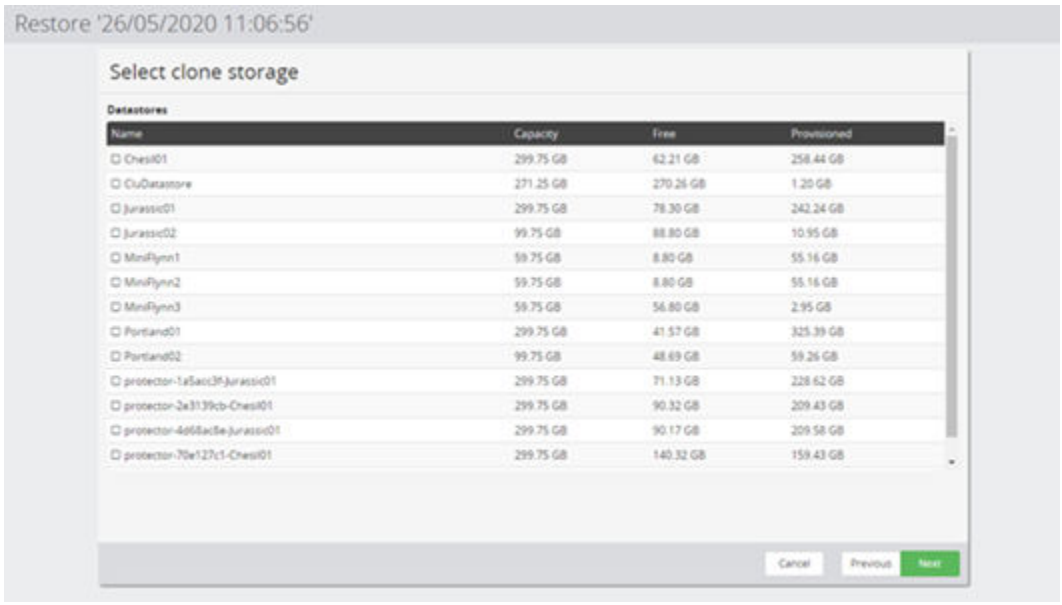


Figure 493 Restore VMware Snapshot Wizard - Select Clone Storage

Control	Description
Datastores	Select the VMware Datastore where the cloned VM(s) will stored.

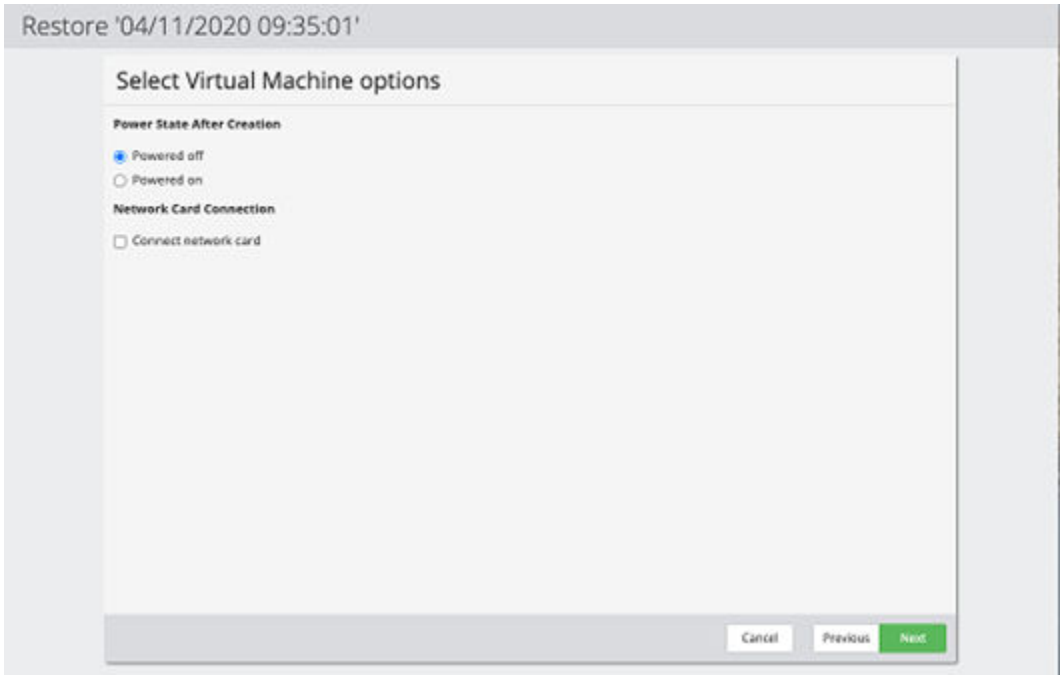


Figure 494 Restore VMware Snapshot Wizard - Select Virtual Machine Options

Control	Description
Powered off	Select this option to leave the restored VM(s) in the powered off state after they are restored.
Powered on	Select this option to place the restored VM(s) in the powered on state after they are restored.
Network Card State	Select this option to connect the restored VM network card on the VM(s) after they are restored.

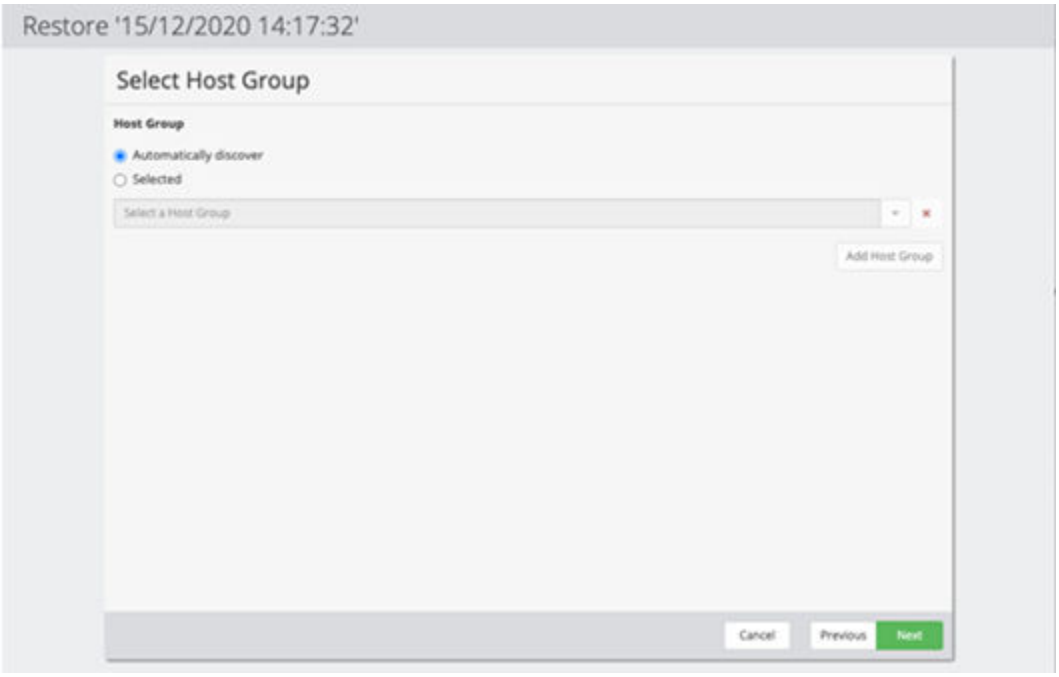


Figure 495 Restore VMware Snapshot Wizard – Select Host Groups

Control	Description
Automatically discover	Select this option to all the Host Groups to be automatically determined.
Selected	Use this option to specify the required Host Groups for exposing this restore point to the selected VMware system.

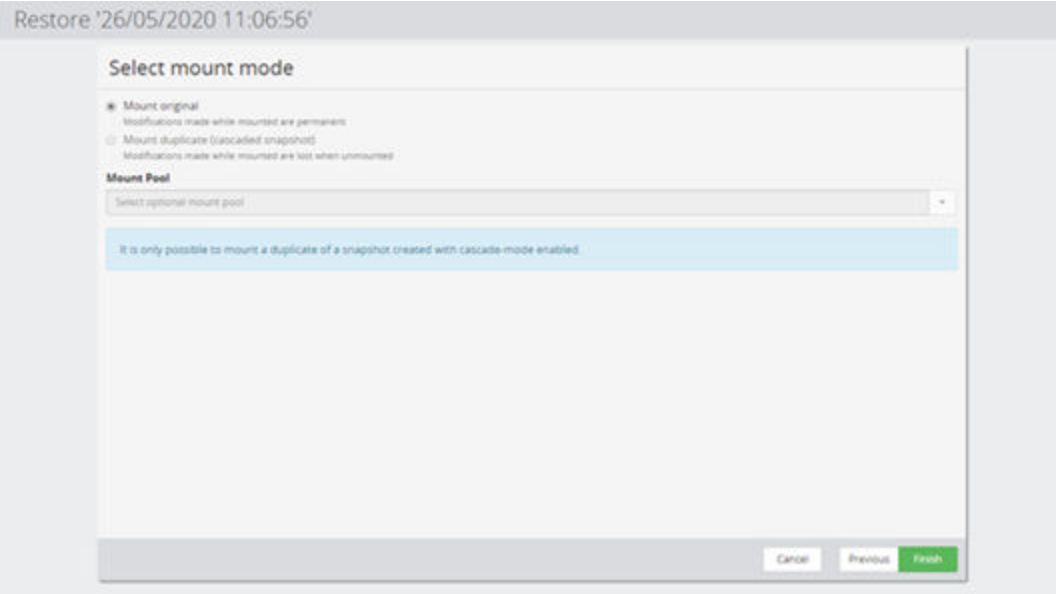


Figure 496 Restore VMware Snapshot Wizard - Select Mount Mode

Control	Description
Mount Original	Mounts the original (Level 1) snapshot and uses vMotion to move the restored VM(s) to the specified location. **Caution:** The process of restoring the VM(s) removes it from the snapshot. The metadata for the snapshot will be updated, and thus it will disappear from the snapshot.
Mount duplicate (cascaded snapshot)	Only enabled if the original (Level 1) snapshot was created in cascade mode. Mounts a copy of the original snapshot (i.e. a Level 2 snapshot). **Note:** The process of restoring the VM(s) removes it from the snapshot. However, because this is a copy, the original snapshot is preserved.
Mount Pool	Depending on the parameters specified for the snapshot operation on the data flow, a mount pool might be required. A message is displayed in a blue rectangle to explain if and why a mount pool is required.

Hitachi Block Snapshot or Replication Mount Wizard

For snapshots and replications using an NVME subsystem please refer to Hitachi Block NVMe Snapshot or Replication Mount Wizard (on page 781)

This wizard is displayed when you mount a snapshot or replication from a Hitachi Block device.

This wizard allows you to expose volumes or mount the selected snapshot or replication.

Caution:

- When a snapshot or replication is mounted, it is made writable. Any changes made to the mounted original snapshot or replication will persist even after it has been unmounted.
- Any attempt to mount two or more copies of the same volume(s) simultaneously on the same machine will fail. The base OS will experience problems if there is already an instance of that file system mounted. This is because snapshots and replications share the same disk partition IDs as the original volume. In extreme conditions it can cause corruption to both disks.

Note: The selected mount host must have a pre-existing LUN mounted from the corresponding storage device (this is required for the auto discover feature to work). If not then the automated mount operation will fail and the error message `"Cannot mount to machine that does not have path to storage"` will be logged.

Note:

- It is not possible to mount the SVOL of a GAD replication, paused or otherwise.
- For Oracle ASM the disks will be presented to the OS but will need to be manually mounted.
- When mounting a snapshot that contains a mounted sub directory, the subdirectory will be mounted as expected. However, the volume referenced by the subdirectory will also be mounted as a separate drive. Unmount will unmount both the expected and unexpected mounts.
- The mount operation can take several minutes to complete.
- If a mount operation needs to mount multiple disks and one of the mount operations fail, the snapshot/replication destination is shown as mounted rather than partially mounted. The logs will indicate the mount was only partially successful.
- It is not possible to perform an OS level mount of a backup that was created using a block host node as its source as there would be no way to know what filesystem(s) (if any) are present on the LDEVs to be mounted. If an OS level mount is desired, then the source should be either an application node (with the appropriate application classification) or an OS host node (with a file system path classification).

Note: Operating System Specific Behaviour

OS	Note
Linux	When mounting a Linux snapshot to a different Linux machine; in order for the user and group names to be displayed correctly the users and groups must have the same ID's as the source.
SUSE Linux	SUSE Linux is not able to perform automated mount operations if hosted on VMware. (RHEL and OEL Linux work as expected).
AIX	The system command **`importvg`** is invoked by Protector to mount snapshots to the user specified location. **`importvg`** creates a directory for the user specified location plus an empty directory corresponding to the original mount point. Neither of these directories are removed by Protector when the snapshot is eventually unmounted, although neither will contain any data.

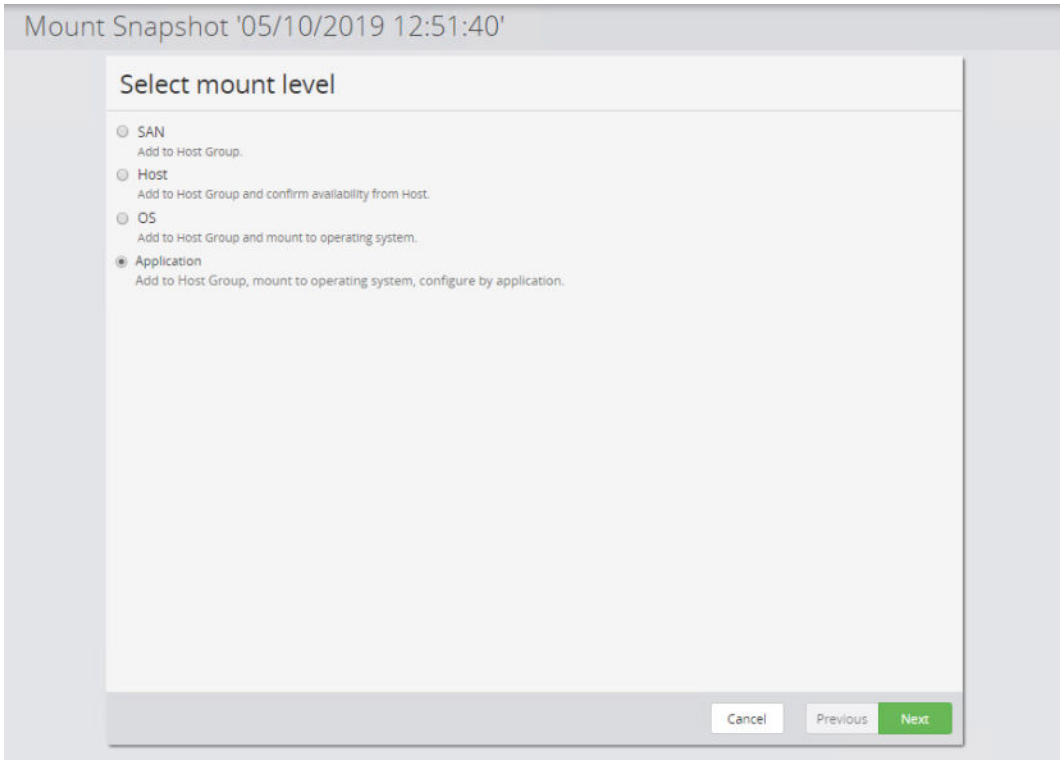


Figure 497 Mount Wizard - Select mount level

Control	Description
SAN	Adds the snapshot or replication to a Host Group.
Host	Adds the snapshot or replication to a Host Group and confirms that it is available from the specified Host.
OS	Adds the snapshot or replication to a Host Group and mounts it on the specified Host's operating system.
Application	Displayed only for application snapshots. Adds the snapshot or replication to a Host Group and mounts it on the specified Host's operating system. The final step in the wizard provides one of the following application specific mount options: ▪ Mount Wizard - Select Oracle Restore Options (on page 797)

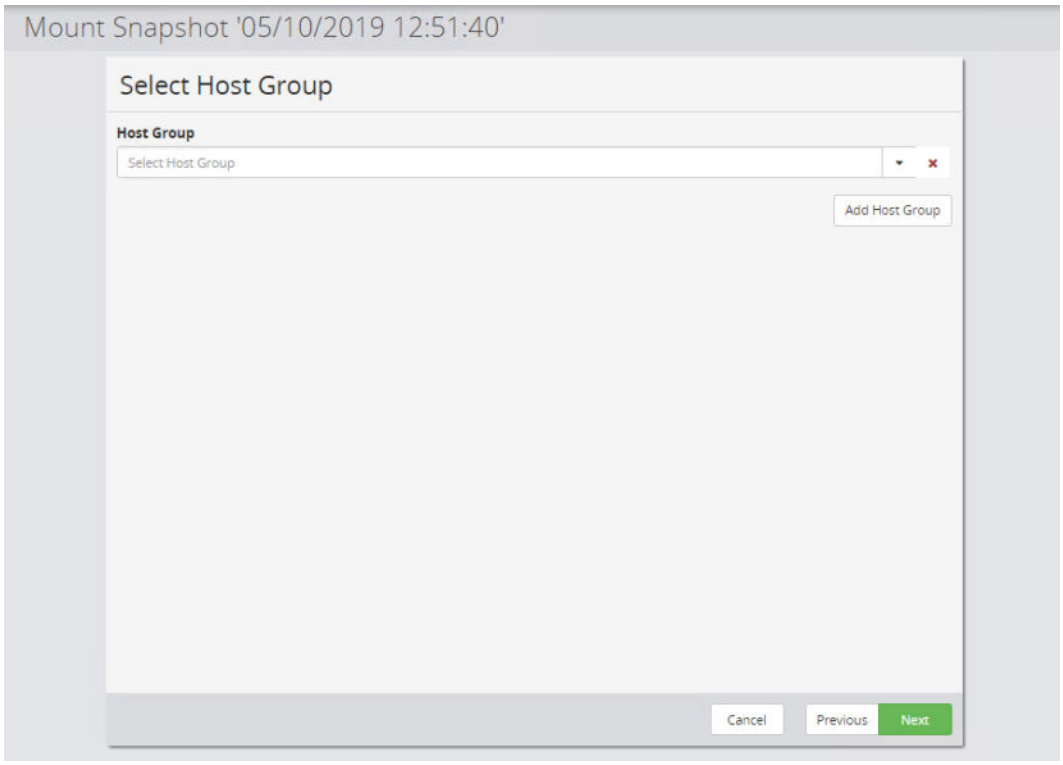


Figure 498 Mount Wizard - Select Host Group (SAN level mount only)

Control	Description
Host Group	Manually specify or select a host group to use to expose a snapshot or replication. **Note:** When exposing an LDEV, the host group specified must be in the same resource group as the secondary volumes.
Add Host Group	Click this button to add host groups when specifying a multi-path mount operation.

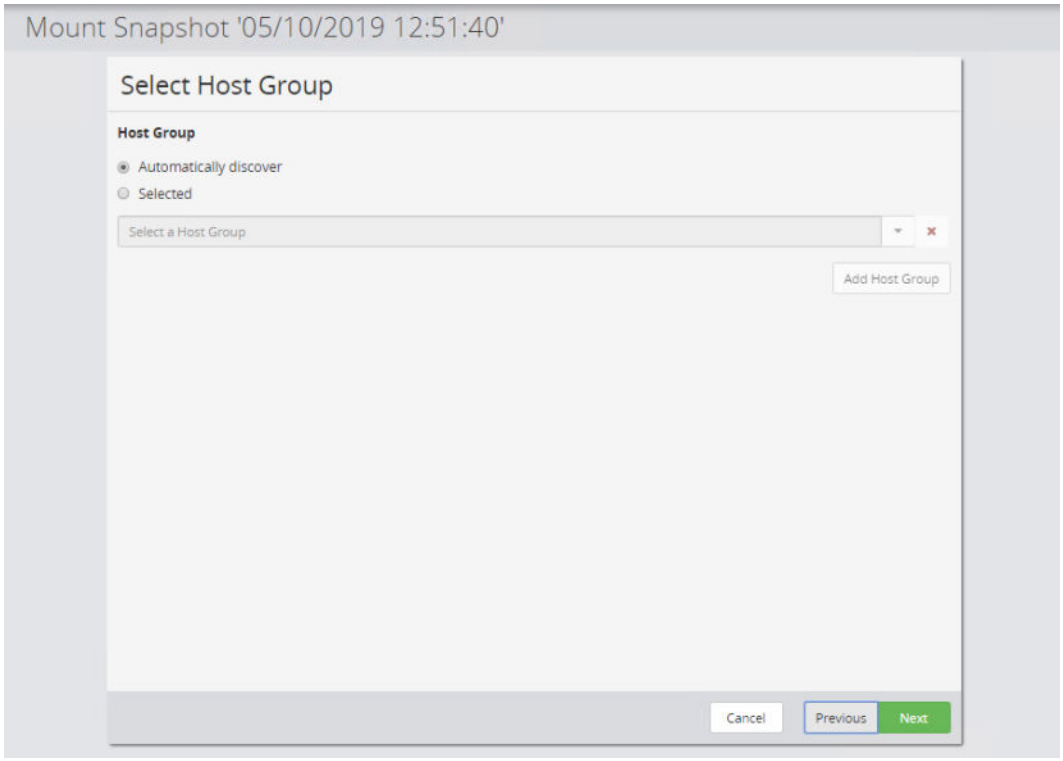


Figure 499 Mount Wizard - Select Host Group (Host and OS level mount only)

Control	Description
Automatically discover	Protector will automatically select a host group to use to expose the snapshot or replication.
Selected	The user must specify one or more host groups to use to expose the snapshot or replication.

Control	Description
Select a Host Group	Manually specify or select a host group to use to expose a snapshot or replication. **Note:** When exposing an LDEV, the host group specified must be in the same resource group as the secondary volumes.
Add Host Group	Click this button to add host groups when specifying a multi-path mount operation.

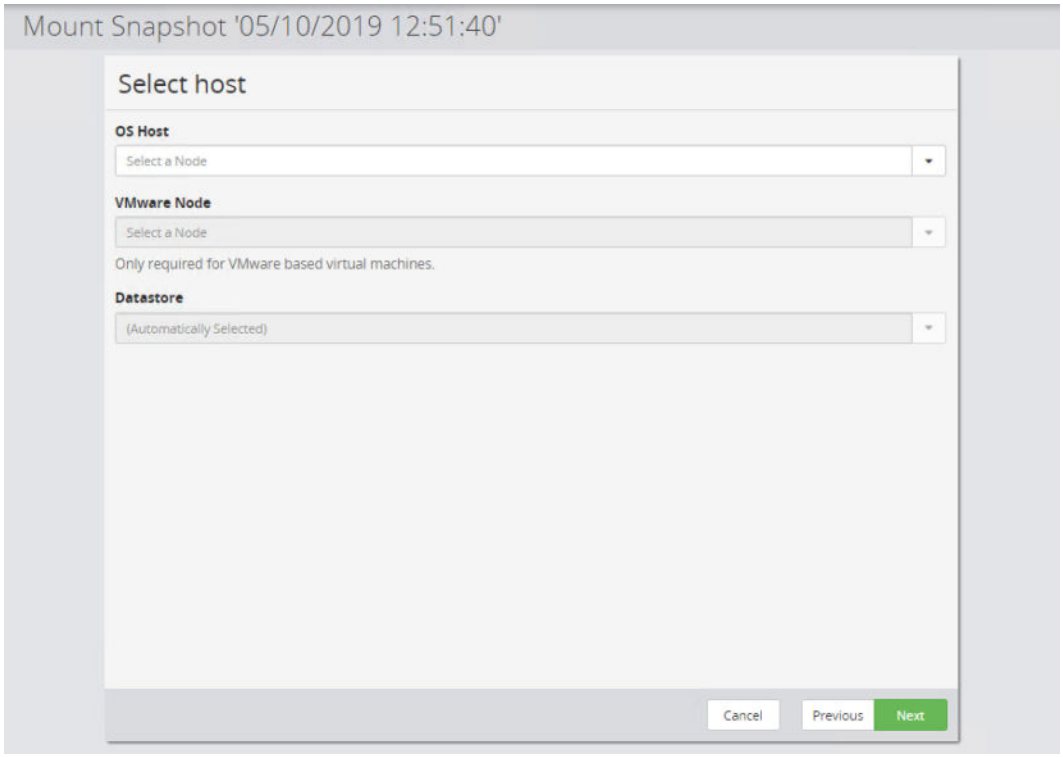


Figure 500 Mount Wizard - Select Host (Host and OS only level mount only)

Control	Description
OS Host	Specify a Protector *OS Node* to mount to or expose to. **Note:** Unless the user selects a host group, the machine where the volume is to be mounted must have an existing volume on the same storage device. If there is no connection between the mount host and the block storage device then Protector will fail the mount operation after a timeout of 30 minutes.

Control	Description
VMware Node	Expose the volumes to the specified VMware host and mount them to the VM as RDM disks. **Note:** Exposing using a VMware host requires that a VMware Server node be configured in Protector and that the mount target VM has Protector Client and VMware Tools installed.
Datastore	Specifies a destination datastore when mounting to a VMware VM which is part of a cluster, in which case the default datastore may not be a suitable place to save the RDM mount information. If the datastore field is left blank then mount information is saved alongside the VM.

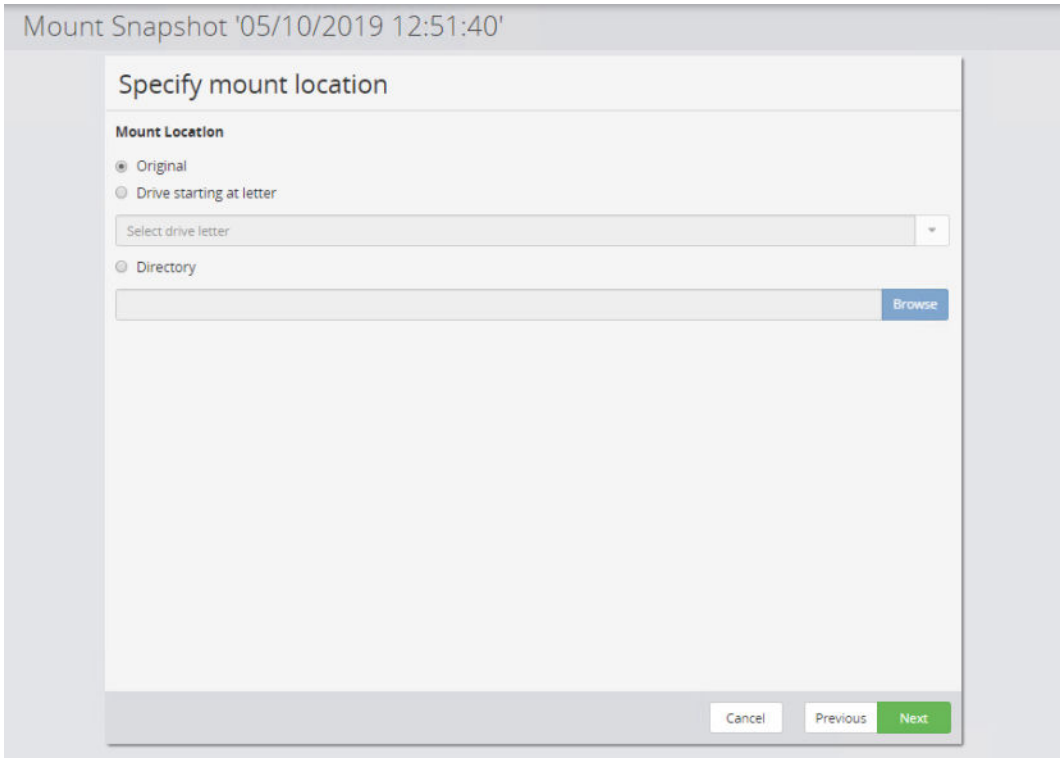


Figure 501 Mount Wizard - Specify mount location (OS level mount only)

Control	Description
Original	The snapshot or replication is mounted at its original location. **Note:** Mounting at the original location will fail if there is already a volume mounted at that location.

Control	Description
Drive starting at letter	When mounting a snapshot or replication that contains multiple volumes, the first volume will mount at the specified drive and subsequent drives are used for each additional volume.
Directory	When mounting a snapshot or replication that contains multiple volumes, each volume will be assigned a separate subdirectory. Click Browse to view the drives and directories on the selected host. To create a new directory, type in the required path. **Note:** Protector does not check to make sure the directory selected as the mount point is empty. This means it is possible to mount a snapshot inside or even over the top of another mounted volume. This should be avoided.

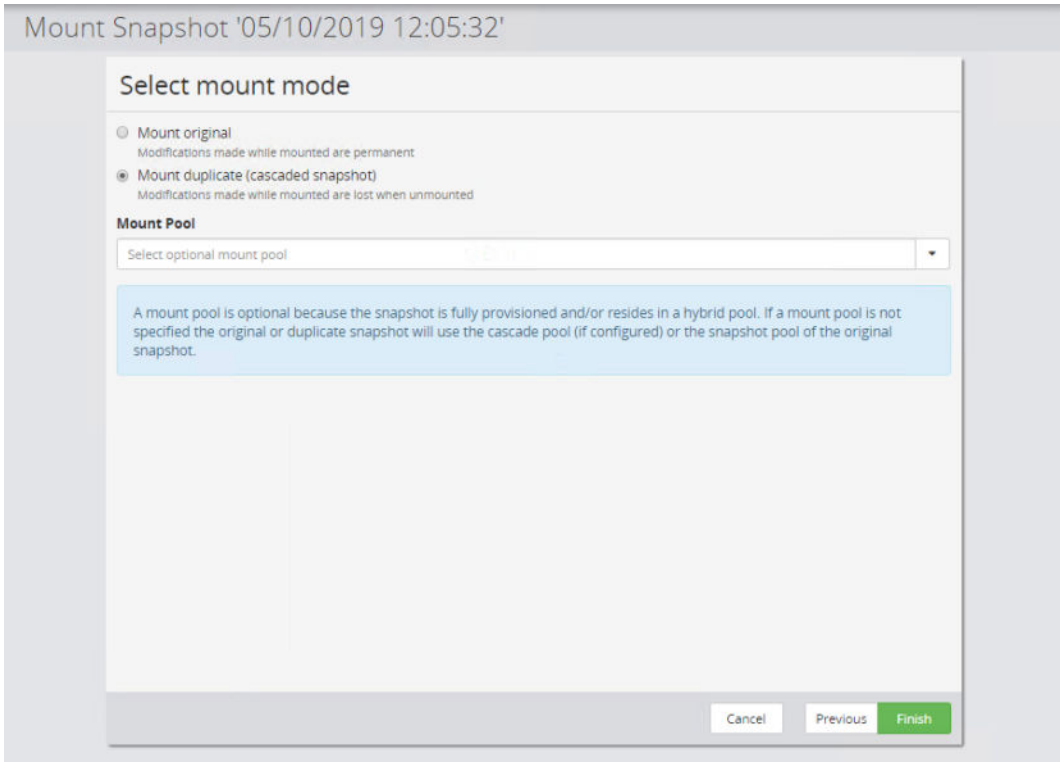


Figure 502 Mount Wizard - Select mount mode

Control	Description
Mount Original	Mounts the replication or the original (Level 1) snapshot. **Note:** If the mounted snapshot is modified then those changes will remain when the snapshot is unmounted.

<table>
<tr><th>Control</th><th>Description</th></tr>
<tr><td>Mount duplicate (cascaded snapshot)</td><td>Enabled only for snapshots and if the original (Level 1) snapshot was created in cascade mode. Mounts a copy of the original snapshot (i.e. a Level 2 snapshot).

Note: The mounted snapshot will be discarded when it is unmounted.</td></tr>
<tr><td>Mount Pool</td><td>Depending on the parameters specified for the snapshot operation on the data flow, a mount pool might or might not be required.

The following table lists all the scenarios:

<table>
<tr><th>Snapshot Pool Type (specified on data flow)</th><th>Provisioning Type (specified on data flow)</th><th>Mount Mode</th><th>Mount Pool</th></tr>
<tr><td>Thin Image</td><td>Floating Device</td><td>Original</td><td>Required(1)</td></tr>
<tr><td>Thin Image</td><td>Floating Device</td><td>Duplicate(4)</td><td>Required(1)</td></tr>
<tr><td>Thin Image</td><td>Fully Provisioned</td><td>Original</td><td>N/A(3)</td></tr>
<tr><td>Thin Image</td><td>Fully Provisioned</td><td>Duplicate(4)</td><td>Optional(2)</td></tr>
<tr><td>Hybrid</td><td>Floating Device</td><td>Original</td><td>Optional(2)</td></tr>
<tr><td>Hybrid</td><td>Floating Device</td><td>Duplicate(4)</td><td>Optional(2)</td></tr>
<tr><td>Hybrid</td><td>Fully Provisioned</td><td>Original</td><td>N/A(3)</td></tr>
<tr><td>Hybrid</td><td>Fully Provisioned</td><td>Duplicate(4)</td><td>Optional(2)</td></tr>
</table>
</td></tr>
</table>

Control	Description
	The following message is displayed in a blue rectangle to explain the options: ▪ (1) “A mount pool is required because the snapshot is not fully provisioned and resides in a Thin Image pool.” ▪ (2) “A mount pool is optional because the snapshot is fully provisioned and/or resides in a hybrid pool. If a mount pool is not specified the original or duplicate snapshot will use the cascade pool (if configured) or the snapshot pool of the original snapshot.” ▪ (3) “A mount pool is not required because the secondary LDEVs are already fully provisioned.” ▪ (4) “It is only possible to mount a duplicate of a snapshot created with cascade-mode enabled.”

Hitachi Block NVMe Snapshot or Replication Mount Wizard

This wizard is displayed when you mount a snapshot or replication from a Hitachi Block device using an NVMe subsystem.

This wizard allows you to expose volumes of the selected snapshot or replication at a SAN level.

Caution:

- When a snapshot or replication is mounted, it is made writable. Any changes made to the mounted original snapshot or replication will persist even after it has been unmounted.
- Any attempt to mount two or more copies of the same volume(s) simultaneously on the same machine will fail. The base OS will experience problems if there is already an instance of that file system mounted. This is because snapshots and replications share the same disk partition IDs as the original volume. In extreme conditions it can cause corruption to both disks.

Note:

The selected mount host must have a pre-existing LUN mounted from the corresponding storage device (this is required for the auto discover feature to work). If not then the automated mount operation will fail and the error message

`"Cannotmount to machine that does not have path to storage"`

will be logged.

Note:

- It is not possible to mount the SVOL of a GAD replication, paused or otherwise.
- For Oracle ASM the disks will be presented to the OS but will need to be manually mounted.
- When mounting a snapshot that contains a mounted sub directory, the subdirectorywill be mounted as expected. However, the volume referenced by the subdirectory will also be mounted as a separate drive. Unmount will unmount both the expected and unexpected mounts.
- The mount operation can take several minutes to complete.
- If a mount operation needs to mount multiple disks and one of the mount operations fail, the snapshot/replication destination is shown as mounted rather than partially mounted. The logs will indicate the mount was only partially successful.
- It is not possible to perform an OS level mount of a backup that was created using a block host node as its source as there would be no way to know what filesystem(s) (if any) are present on the LDEVs to be mounted. If an OS level mount is desired, then the source should be either an application node (with the appropriate application classification) or an OS host node (with a file system path classification).

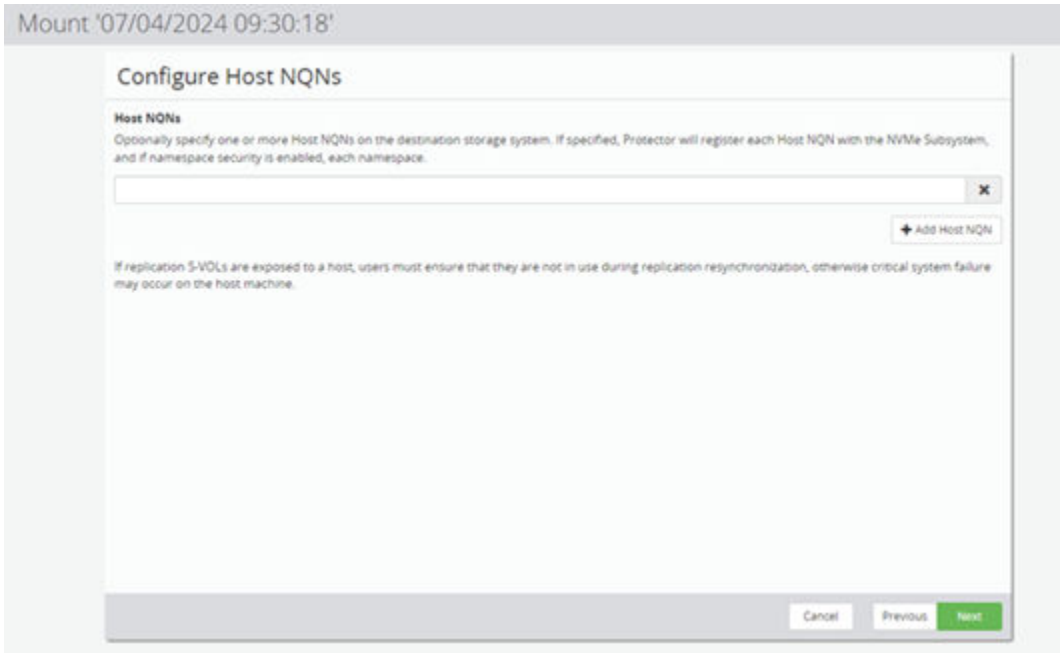


Figure 503 Mount Wizard - Select Host NQNs

Control	Description
Host NQNs	Manually specify or select a host group to use to expose a snapshot or replication. **Note:** When exposing an LDEV, the host group specified must be in the same resource group as the secondary volumes.
Add Host Group	Click this button to add host groups when specifying a multi-path mount operation.

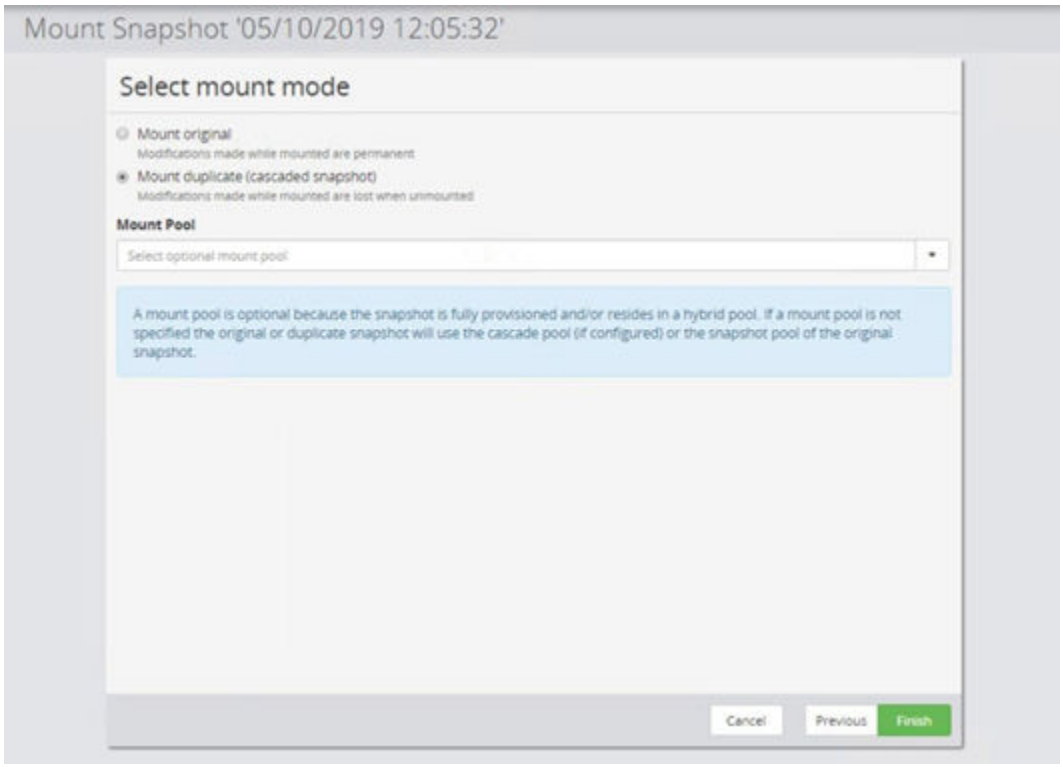


Figure 504 Mount Wizard - Select mount mode

Control	Description
Mount Original	Mounts the replication or the original (Level 1) snapshot. **Note:** If the mounted snapshot is modified then those changes will remain when the snapshot is unmounted.

Control	Description
Mount duplicate (cascaded snapshot)	Enabled only for snapshots and if the original (Level 1) snapshot was createdin cascade mode. Mounts a copy of the original snapshot (i.e. a Level 2 snapshot). **Note:** The mounted snapshot will be discarded when it is unmounted.
Mount Pool	Depending on the parameters specified for the snapshot operation on the data flow, a mount pool might or might not be required.

<table>
<tr><th>Control</th><th>Description</th></tr>
<tr><td></td><td>The following table lists all the scenarios:
<table>
<tr><th>Snapshot Pool Type (specified on data flow)</th><th>Provisioning Type (specified on data flow)</th><th>Mount Mode</th><th>Mount Pool</th><th></th></tr>
<tr><td>Thin Image</td><td>Floating Device</td><td>Original</td><td>Required(1)</td><td></td></tr>
<tr><td>Thin Image</td><td>Floating Device</td><td>Duplicate(4)</td><td>Required(1)</td><td></td></tr>
<tr><td>Thin Image</td><td>Fully Provisioned</td><td>Original</td><td>N/A(3)</td><td rowspan="7"></td></tr>
<tr><td>Thin Image</td><td>Fully Provisioned</td><td>Duplicate(4)</td><td>Optional(2)</td></tr>
<tr><td>Hybrid</td><td>Floating Device</td><td>Original</td><td>Optional(2)</td></tr>
<tr><td>Hybrid</td><td>Floating Device</td><td>Duplicate(4)</td><td>Optional(2)</td></tr>
<tr><td>Hybrid</td><td>Fully Provisioned</td><td>Original</td><td>N/A(3)</td></tr>
<tr><td>Hybrid</td><td>Fully Provisioned</td><td>Duplicate(4)</td><td>Optional(2)</td></tr>
</table></td></tr>
</table>

Control	Description
	The following message is displayed in a blue rectangle to explain the options: ▪ (1)“A mount pool is required because the snapshot is not fully provisioned and resides in a Thin Image pool.” ▪ (2) “A mount pool is optional because the snapshot is fully provisioned and/or resides in a hybrid pool. If a mount pool is not specifiedthe original or duplicate snapshot will use the cascade pool (if configured) or the snapshot pool of the original snapshot.” ▪ (3)“A mount pool is not required because the secondary LDEVs are already fully provisioned.” ▪ (4)“It is only possible to mount a duplicate of a snapshot created with cascade-mode enabled.”

When mounting block snapshots or replications created by a policy containing an *Oracle Database* classification, Protector will display the following additional wizard pages that allows application specific options to be configured:

Note: To perform the mount operation, a snapshot should be available.

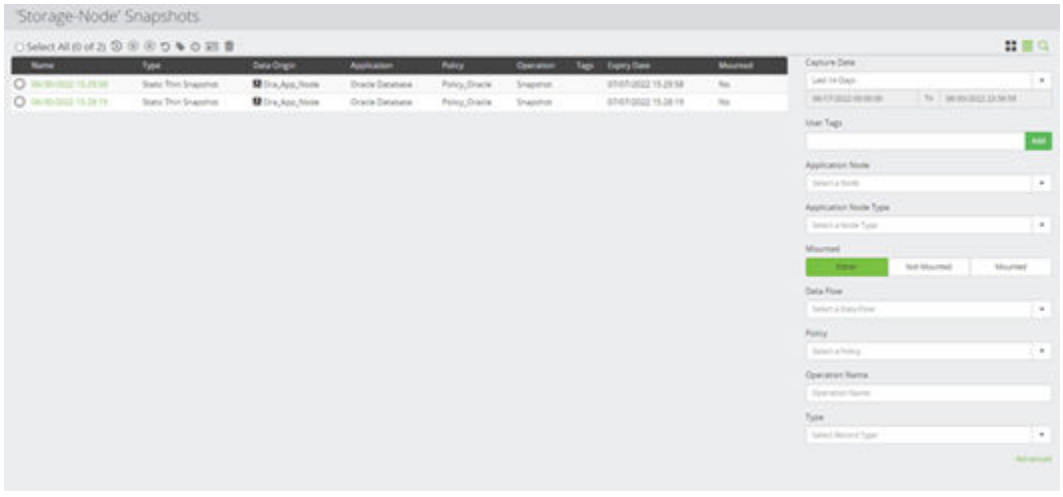


Figure 505 Snapshot to be mounted

Table 44 List of Snapshots to be mounted

Control	Description
Name	Displays the name of the snapshots to be mounted.
Type	Displays the type of the snapshot.
Data Origin	Displays the Application Node that was used to create the snapshot.
Application	Displays the application on which the snapshot is running.
Policy	Displays the policy associated with the snapshot.
Operation	Displays the type of the operation for the snapshot.
Tags	Displays the tags associated with that snapshot.
Expiry Date	Displays the expiry date for the snapshot.
Mounted	States the status of the snapshot if it is mounted or not mounted.
Filter on Capture Date	Filters the snapshot on the date range on which it was created.
Filter on User Tags	Filters the displayed results based on Tags.
Filter on Application Node	Filters the snapshot by the name of the node.
Filter on Application Node Type	Filters the snapshot on the type of the Application node.
Filter on Mounted	Filters the snapshot by either Mounted or Not Mounted.
Filter on Data Flow	Filters the snapshot based on the Data Flow associated with a snapshot.
Filter on Operation Name	Filters the snapshot so that only entries with the specified Operation Name are displayed.
Filter on Type	Filters the snapshot based on its type.
Filter on Advanced Query String	Provides advanced search capabilities, allowing users to create filters based on combinations of API properties (see the API reference guide for more details).

Control	Description
Search	Returns the result on Advanced Query String.

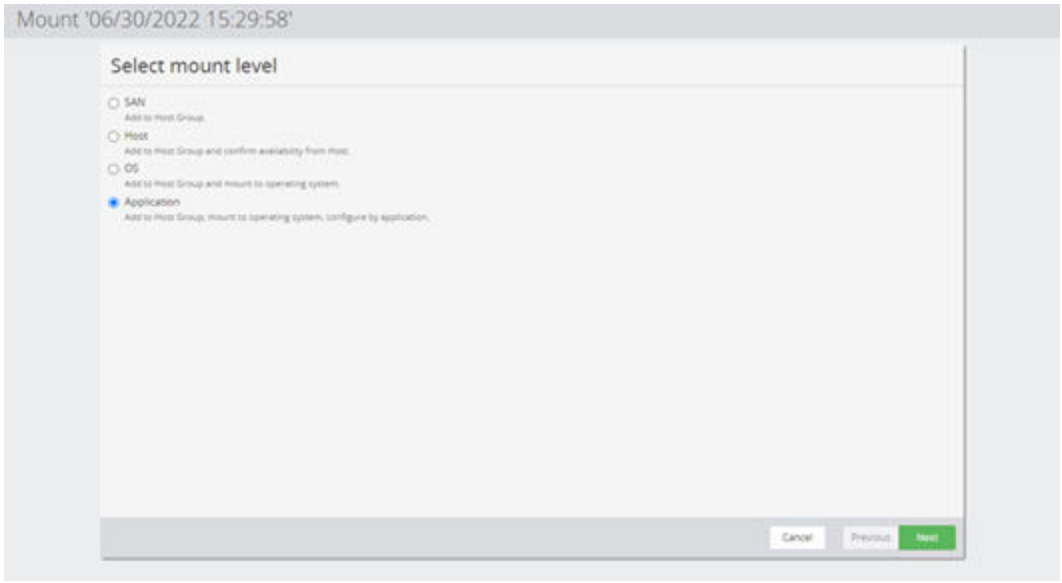


Figure 506 Select Mount Level

Table 45 Oracle - Select Mount Level

Control	Description
SAN	Add this record to a Host Group.
Host	Add this record to a Host Group and confirm availability from Host.
OS	Add this record to a Host Group and mount to operating system.
Application	Add this record to a Host Group, mount to operating system configure by application (This will include ASM operations).

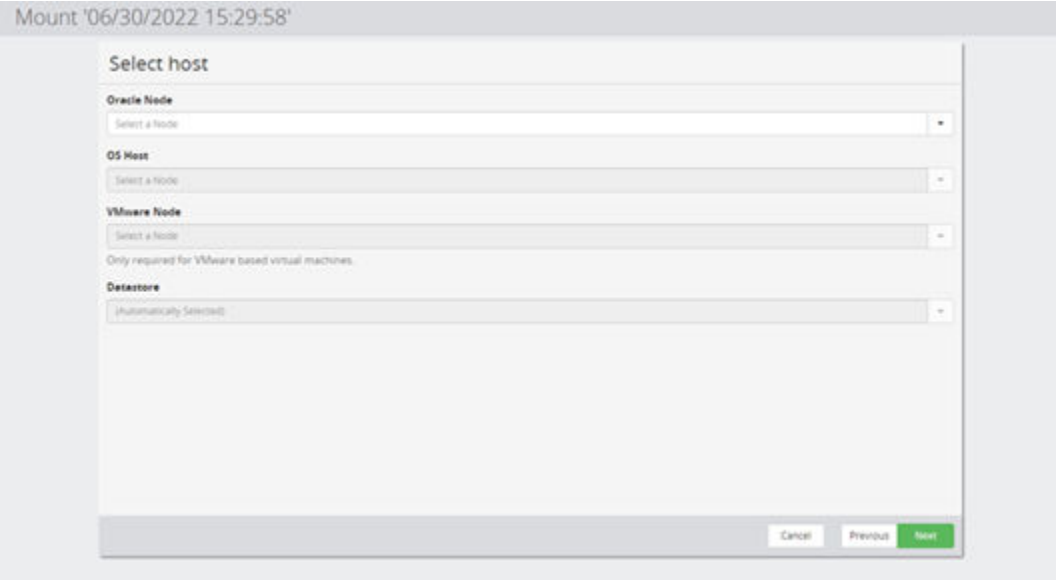


Figure 507 Oracle Mount - Select Host

Table 46 Oracle - Select Host

Control	Description
Oracle node	The oracle node where the user will mount the snapshot representing the target Oracle app server environment.
OS Host	The node which hosts the Oracle application node.
VMware Node	Select the VMware Host or vCenter where the VM's disks will be mounted.
Datastore	Select the Datastore where the disk will be mounted.

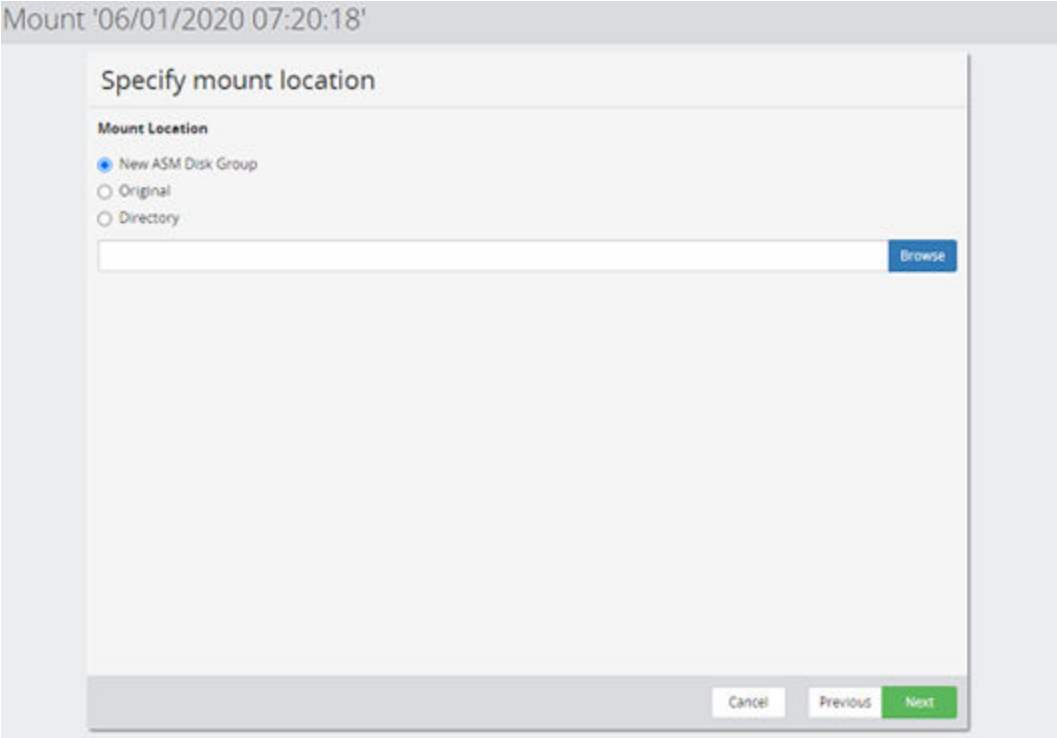


Figure 508 Mount Wizard - Specify Mount Location

Control	Description
New ASM Disk Group	Mount the ASM disk groups using a new generated name to avoid conflicts with the original database or existing ASM disk groups. The new disk group name is auto generated. It comprises the original name and a numeric suffix, allowing for multiple copies of the same database to be mounted to a single host. **Note:** This option is only valid for backups of ASM based Oracle databases. The option will be disabled automatically if Protector detects a filesystem based Oracle database in a backup created with Protector 7.1 or newer.
Original	Mount the database using the same ASM disk groups or filesystem paths as the original database the backup was created from. If the path or disk group name is in use by an existing database, the mount will fail.

Control	Description
Directory	Mount the database to the provided path. If the path in use by an existing database, the mount will fail. **Note:** This option is only valid for backups of filesystem based Oracle databases. The option will be deactivated automatically, if Protector detects an ASM based Oracle database in a backup created with Protector 7.1 or newer.

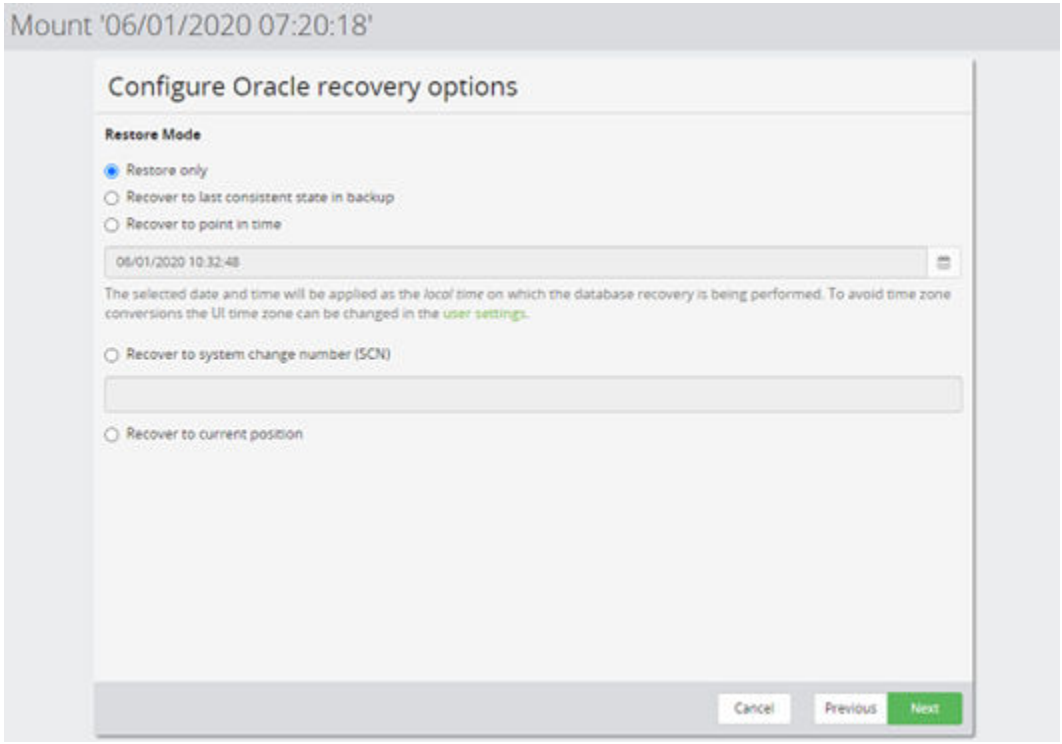


Figure 509 Mount Wizard - Select Oracle Restore Options

Control	Description	Logs Reset Post Mount	Requires RMAN catalog	Requires control/ spfile in RMAN backup
Restore only	The database is simply mounted. It is left to the database administrator to recover manually.	No	No	No

Control	Description	Logs Reset Post Mount	Requires RMAN catalog	Requires control/ spfile in RMAN backup
Recover to last consistent state in backup	The database is recovered to the consistent state which was captured by the backup. The database is brought online. This type of mount can be performed with the data in the backup alone.	Depends (see note 1)	No	No
Recover to point in time	A timestamp is entered which defines the point in time to recover. This requires a connection to the RMAN catalog and logs which are shared with/ available to the host. **Note:** This is not applicable for DataGuard configuration.	Depends (see note 1)	Yes	Yes (see note 2)
Recover to system change number (SCN)	A system change number is entered which defines the change point to recover. This requires a connection to the RMAN catalog and logs which are shared with/available to the host. **Note:** This is not applicable for DataGuard configuration.	Depends (see note 1)	Yes	Yes (see note 2)
Recover to current position	The database is recovered to the most current position possible. This requires a connection to the RMAN catalog and logs which are shared with/available to the host.	Depends (see note 1)	Yes	Yes (see note 2)

Control	Description	Logs Reset Post Mount	Requires RMAN catalog	Requires control/ spfile in RMAN backup
	Note: ▪ When mounting, the current position is the latest point in time which is provided by the archive logs referenced in RMAN catalog. It does not include any archive logs or redo logs on the source machine which have not been backed up via RMAN. ▪ This is not applicable for DataGuard configuration.			

Note: (1) In the table above, logs will only be reset when Open Database is selected in the **Post Recovery Options** page of the wizard (see below).

Caution: (2) In the table above, for some recovery scenarios the RMAN catalog needs to hold a control file. RMAN can be configured to add a control/spfile backup every time an archive log backup is performed.

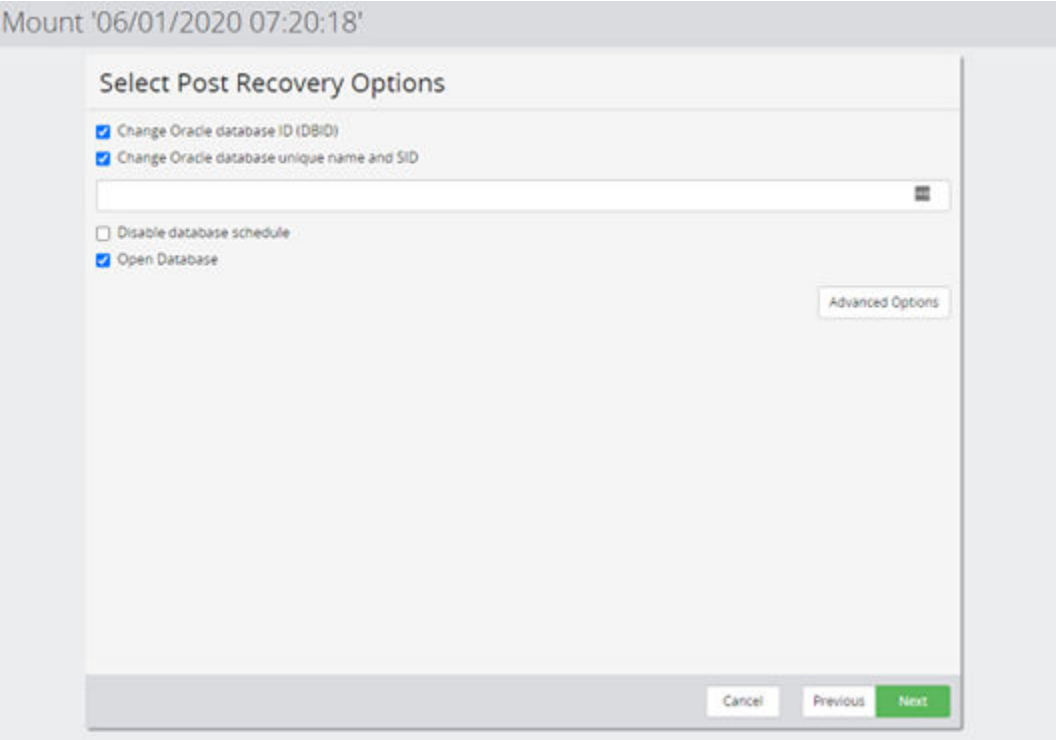


Figure 510 Mount Wizard - Select Post Recovery Options

Control	Description
Change Oracle database ID (DBID)	Creates a new DBID for the database. **Tip:** A DBID is a unique, Oracle generated number identifying each database. It is found in control files as well as datafile headers and is used to determine which database that file belongs to.
Change Oracle database unique name and SID	A new unique name and SID can be specified for the database. **Tip:** This changes the `unique_database_name` which is also used as the SID.
Disable database schedule	Disables database internal tasks scheduled for this database. **Tip:** The Oracle scheduler allows the administrator to schedule SQL commands as jobs. By selecting this option existing schedules will be disabled.
Open Database	If selected, then after recovery the database will placed in the *OPEN* state using the *RESETLOGS* or *NORESETLOGS* option, as per the requirements of the database. Otherwise the database will be left in the *MOUNT* state.
Advanced Options	Opens the **Advanced Options** page of the wizard.

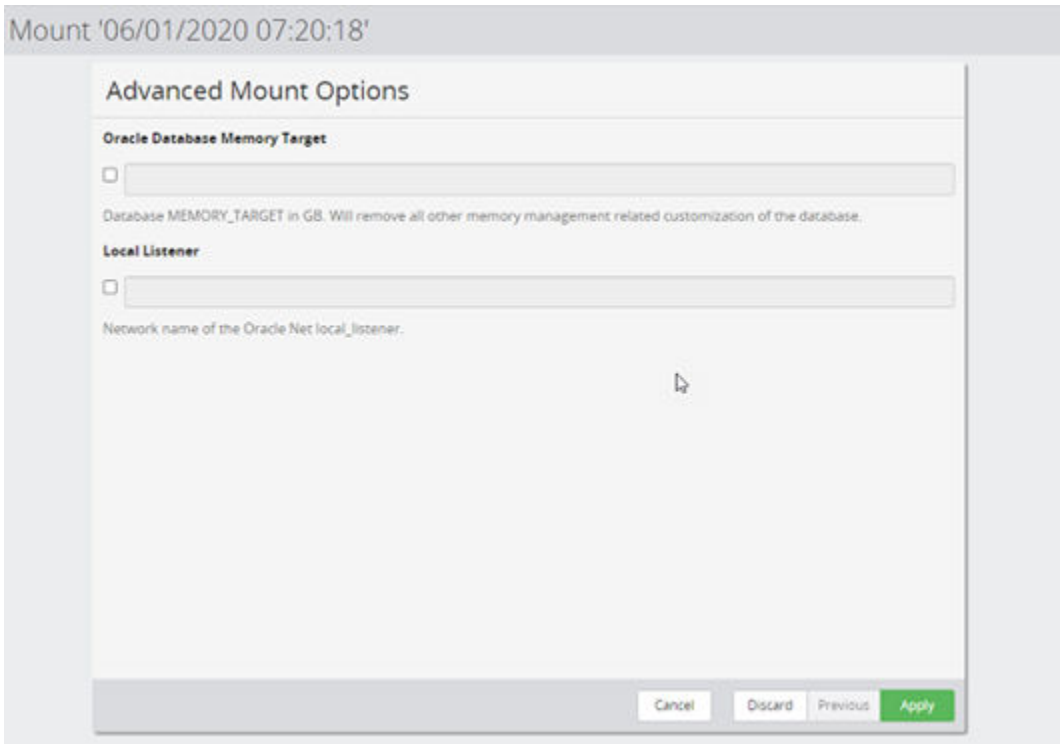


Figure 511 Mount Wizard - Advanced Mount Options

This page of the wizard is not be displayed if Restore only or Recover to last consistent state in backup is selected in a previous step.

Control	Description
Oracle database Memory Target	Sets the database `MEMORY_TARGET` in GB. Entering a value here will remove all other memory management related customization of the database. **Tip:** This allows Oracle databases from very powerful source systems to be deployed on less powerful systems. The PGA and SGA memory areas will be managed by Oracle within the given memory target.
Local Listener	Sets the network name of the Oracle Net `local_listener`.

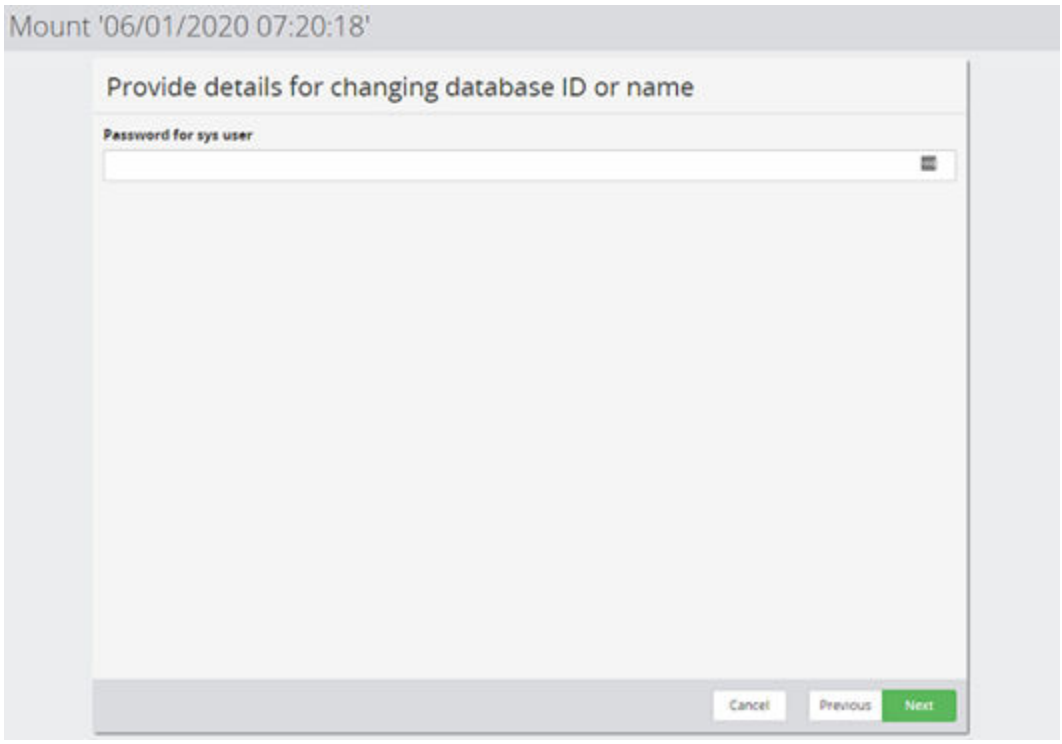


Figure 512 Mount Wizard - Provide details for changing database ID or name

This page of the wizard will only be displayed if either Change Oracle database ID (DBID) or Change Oracle database unique name and SID options are selected in a previous step.

Control	Description
Password for sys user	Depending on the Oracle version: ▪ Oracle 11g: The sys user password is required to change the Oracle database ID or database unique name. ▪ Oracle 12 or newer: This field can be left empty.

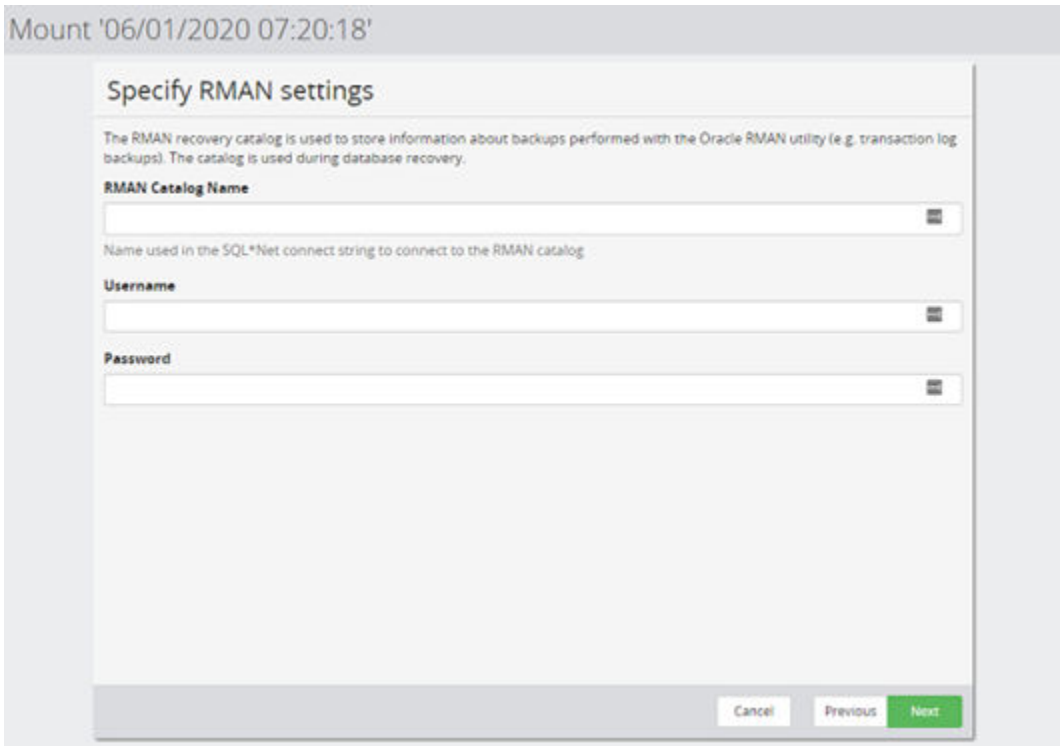


Figure 513 Mount Wizard - Specify RMAN credentials

This page of the wizard is not displayed if Restore only or Recover to last consistent state in backup is selected in a previous step.

Control	Description
RMAN Catalog Name	For RMAN only. Enter the RMAN Catalog Name as it is entered in the SQL*Net connect string to connect to the RMAN catalog.
Username	For RMAN only. Enter the username for the RMAN catalog.
Password	For RMAN only. Enter the password for the RMAN catalog.

Mount Wizard - Select Oracle Restore Options

When mounting block snapshots or replications created by a policy containing an *Oracle Database* classification, Protector will display the following additional wizard pages that allows application specific options to be configured:

Note: To perform the mount operation, a snapshot should be available.

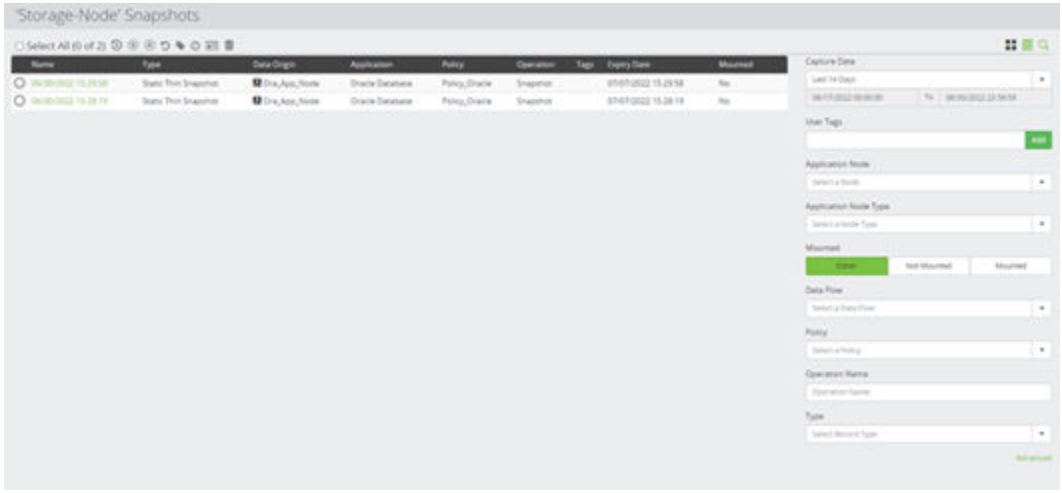


Figure 514 Snapshot to be mounted

Table 47 List of Snapshots to be mounted

Control	Description
Name	Displays the name of the snapshots to be mounted.
Type	Displays the type of the snapshot.
Data Origin	Displays the Application Node that was used to create the snapshot.
Application	Displays the application on which the snapshot is running.
Policy	Displays the policy associated with the snapshot.
Operation	Displays the type of the operation for the snapshot.
Tags	Displays the tags associated with that snapshot.
Expiry Date	Displays the expiry date for the snapshot.
Mounted	States the status of the snapshot if it is mounted or not mounted.
Filter on Capture Date	Filters the snapshot on the date range on which it was created.
Filter on User Tags	Filters the displayed results based on Tags.
Filter on Application Node	Filters the snapshot by the name of the node.

Control	Description
Filter on Application Node Type	Filters the snapshot on the type of the Application node.
Filter on Mounted	Filters the snapshot by either Mounted or Not Mounted.
Filter on Data Flow	Filters the snapshot based on the Data Flow associated with a snapshot.
Filter on Operation Name	Filters the snapshot so that only entries with the specified Operation Name are displayed.
Filter on Type	Filters the snapshot based on its type.
Filter on Advanced Query String	Provides advanced search capabilities, allowing users to create filters based on combinations of API properties (see the API reference guide for more details).
Search	Returns the result on Advanced Query String.

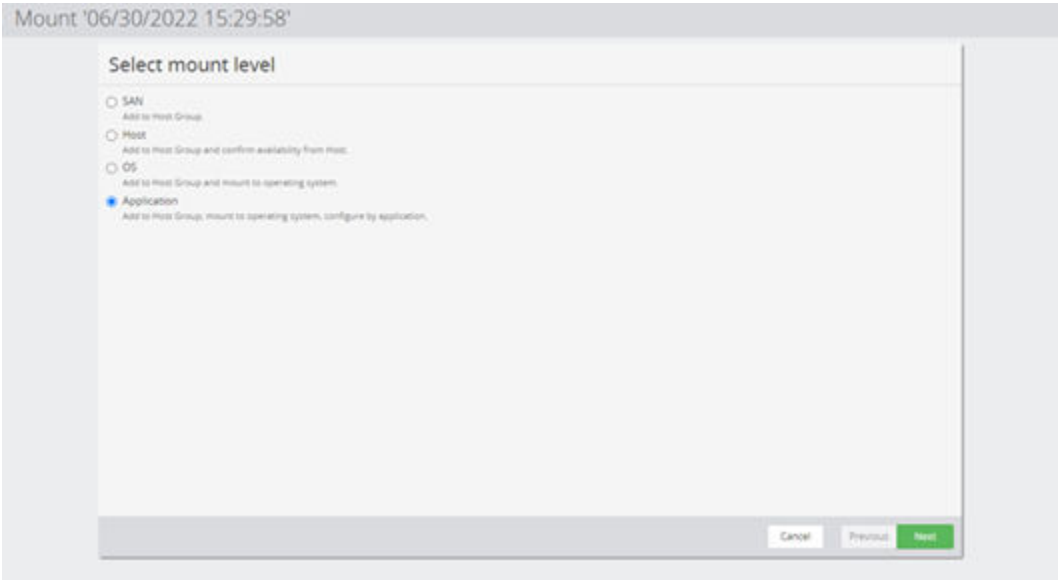


Figure 515 Select Mount Level

Table 48 Oracle - Select Mount Level

Control	Description
SAN	Add this record to a Host Group.

Control	Description
Host	Add this record to a Host Group and confirm availability from Host.
OS	Add this record to a Host Group and mount to operating system.
Application	Add this record to a Host Group, mount to operating system configure by application (This will include ASM operations).

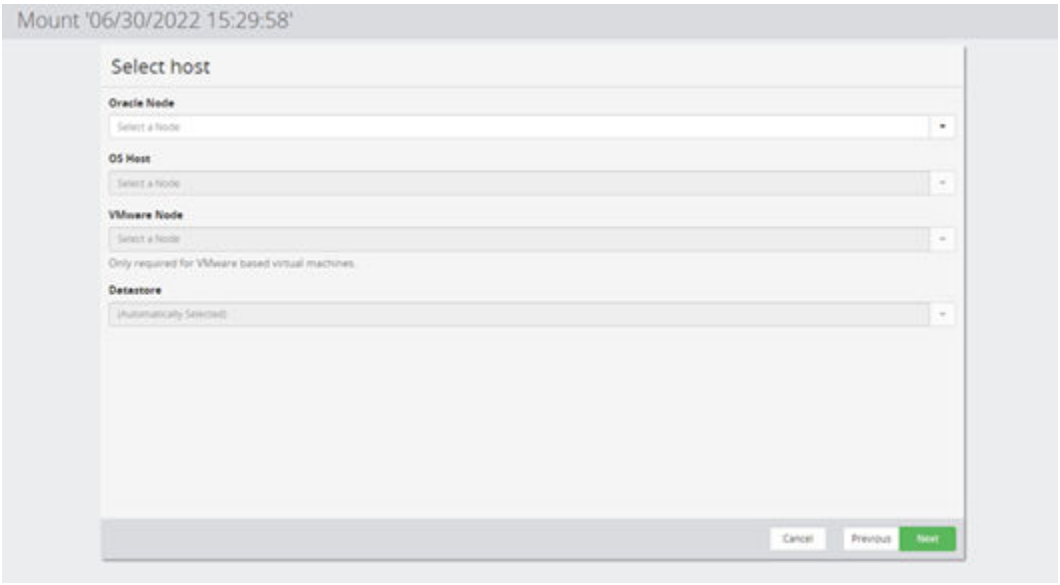


Figure 516 Oracle Mount - Select Host

Table 49 Oracle - Select Host

Control	Description
Oracle node	The oracle node where the user will mount the snapshot representing the target Oracle app server environment.
OS Host	The node which hosts the Oracle application node.
VMware Node	Select the VMware Host or vCenter where the VM's disks will be mounted.
Datastore	Select the Datastore where the disk will be mounted.

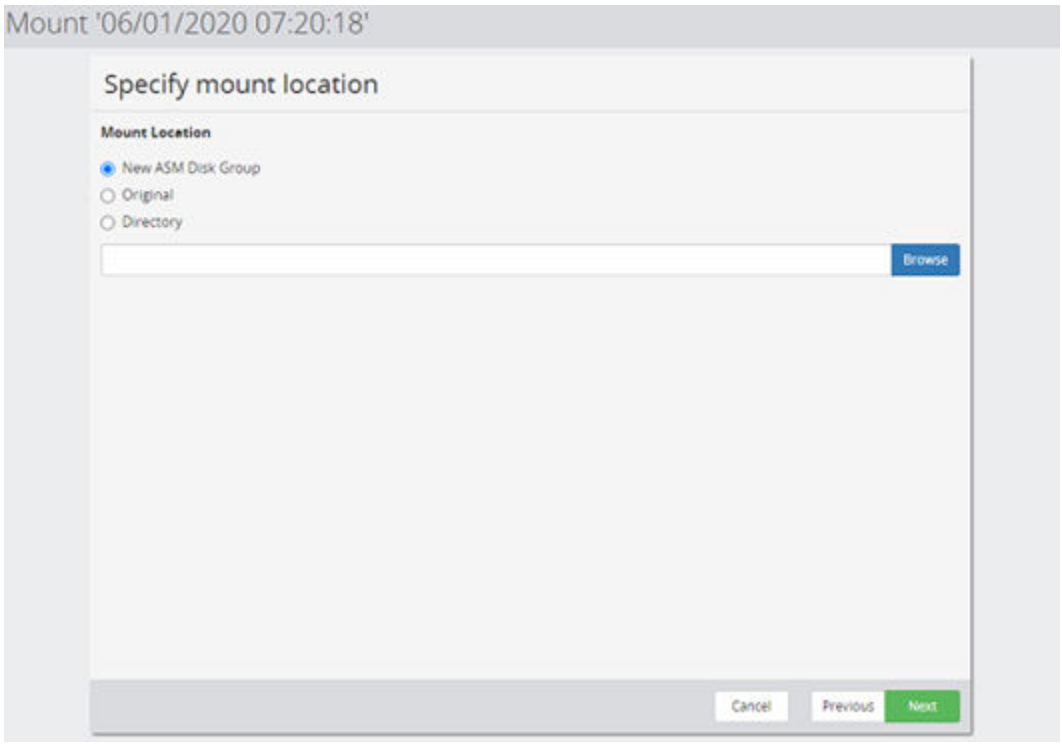


Figure 517 Mount Wizard - Specify Mount Location

Control	Description
New ASM Disk Group	Mount the ASM disk groups using a new generated name to avoid conflicts with the original database or existing ASM disk groups. The new disk group name is auto generated. It comprises the original name and a numeric suffix, allowing for multiple copies of the same database to be mounted to a single host. **Note:** This option is only valid for backups of ASM based Oracle databases. The option will be disabled automatically if Protector detects a filesystem based Oracle database in a backup created with Protector 7.1 or newer.
Original	Mount the database using the same ASM disk groups or filesystem paths as the original database the backup was created from. If the path or disk group name is in use by an existing database, the mount will fail.

Control	Description
Directory	Mount the database to the provided path. If the path in use by an existing database, the mount will fail. **Note:** This option is only valid for backups of filesystem based Oracle databases. The option will be deactivated automatically, if Protector detects an ASM based Oracle database in a backup created with Protector 7.1 or newer.

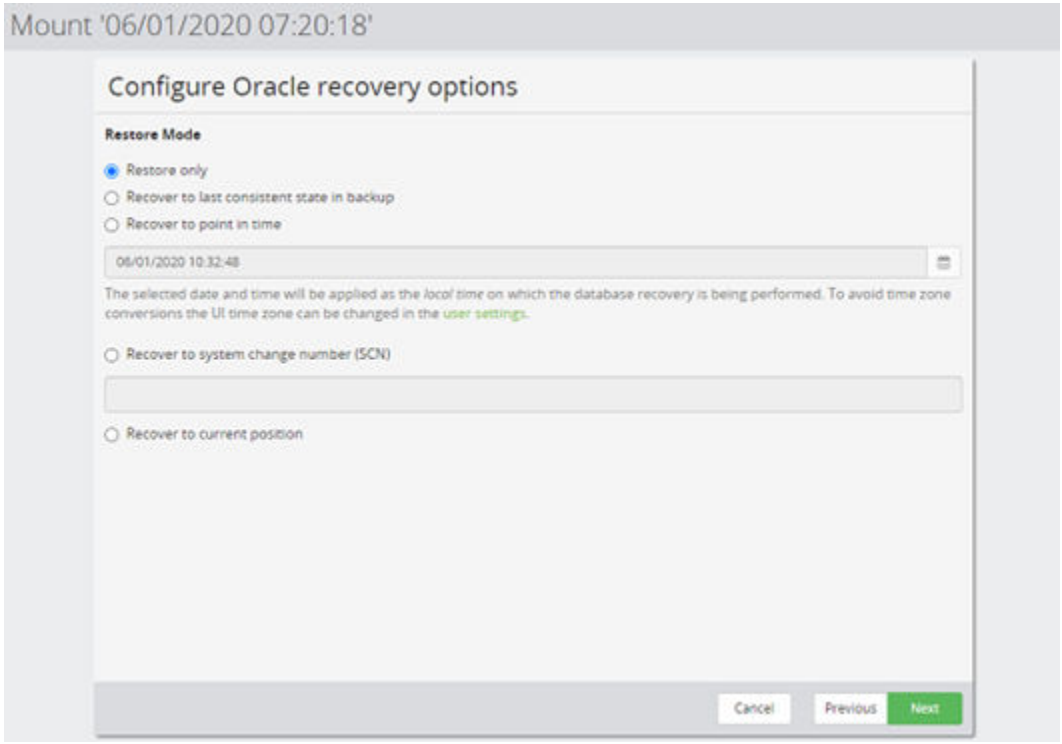


Figure 518 Mount Wizard - Select Oracle Restore Options

Control	Description	Logs Reset Post Mount	Requires RMAN catalog	Requires control/ spfile in RMAN backup
Restore only	The database is simply mounted. It is left to the database administrator to recover manually.	No	No	No

Control	Description	Logs Reset Post Mount	Requires RMAN catalog	Requires control/ spfile in RMAN backup
Recover to last consistent state in backup	The database is recovered to the consistent state which was captured by the backup. The database is brought online. This type of mount can be performed with the data in the backup alone.	Depends (see note 1)	No	No
Recover to point in time	A timestamp is entered which defines the point in time to recover. This requires a connection to the RMAN catalog and logs which are shared with/ available to the host. **Note:** This is not applicable for DataGuard configuration.	Depends (see note 1)	Yes	Yes (see note 2)
Recover to system change number (SCN)	A system change number is entered which defines the change point to recover. This requires a connection to the RMAN catalog and logs which are shared with/available to the host. **Note:** This is not applicable for DataGuard configuration.	Depends (see note 1)	Yes	Yes (see note 2)
Recover to current position	The database is recovered to the most current position possible. This requires a connection to the RMAN catalog and logs which are shared with/available to the host.	Depends (see note 1)	Yes	Yes (see note 2)

Control	Description	Logs Reset Post Mount	Requires RMAN catalog	Requires control/ spfile in RMAN backup
	Note: ▪ When mounting, the current position is the latest point in time which is provided by the archive logs referenced in RMAN catalog. It does not include any archive logs or redo logs on the source machine which have not been backed up via RMAN. ▪ This is not applicable for DataGuard configuration.			

Note: (1) In the table above, logs will only be reset when Open Database is selected in the **Post Recovery Options** page of the wizard (see below).

Caution: (2) In the table above, for some recovery scenarios the RMAN catalog needs to hold a control file. RMAN can be configured to add a control/spfile backup every time an archive log backup is performed.

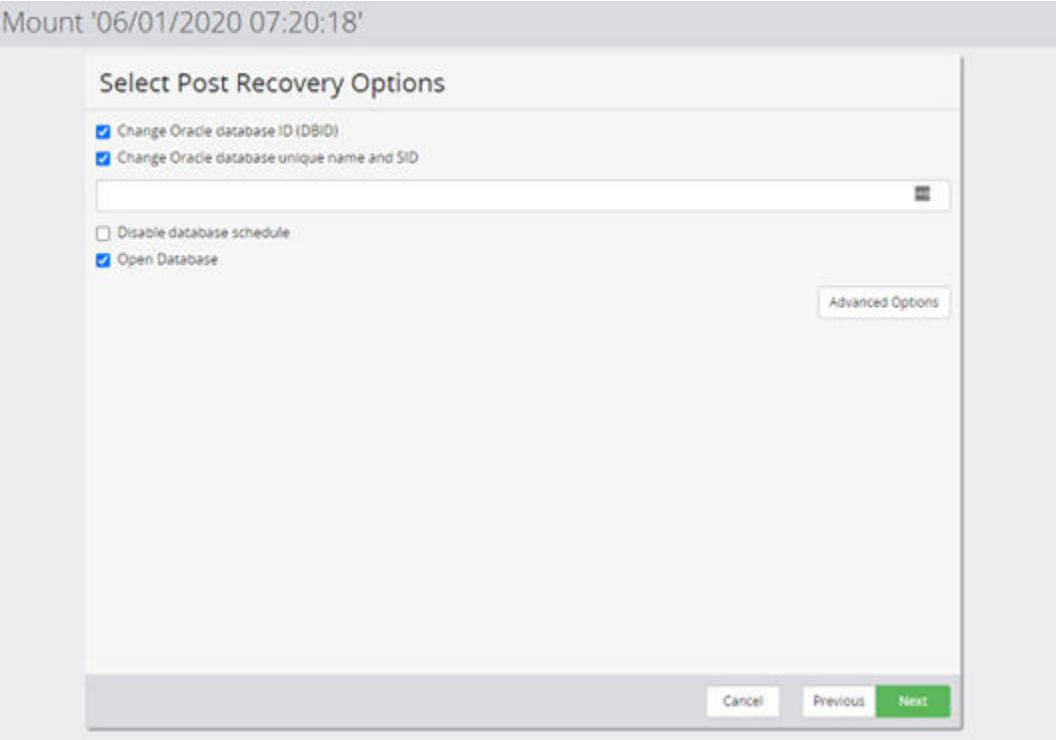


Figure 519 Mount Wizard - Select Post Recovery Options

Control	Description
Change Oracle database ID (DBID)	Creates a new DBID for the database. **Tip:** A DBID is a unique, Oracle generated number identifying each database. It is found in control files as well as datafile headers and is used to determine which database that file belongs to.
Change Oracle database unique name and SID	A new unique name and SID can be specified for the database. **Tip:** This changes the `unique_database_name` which is also used as the SID.
Disable database schedule	Disables database internal tasks scheduled for this database. **Tip:** The Oracle scheduler allows the administrator to schedule SQL commands as jobs. By selecting this option existing schedules will be disabled.
Open Database	If selected, then after recovery the database will placed in the *OPEN* state using the *RESETLOGS* or *NORESETLOGS* option, as per the requirements of the database. Otherwise the database will be left in the *MOUNT* state.
Advanced Options	Opens the **Advanced Options** page of the wizard.

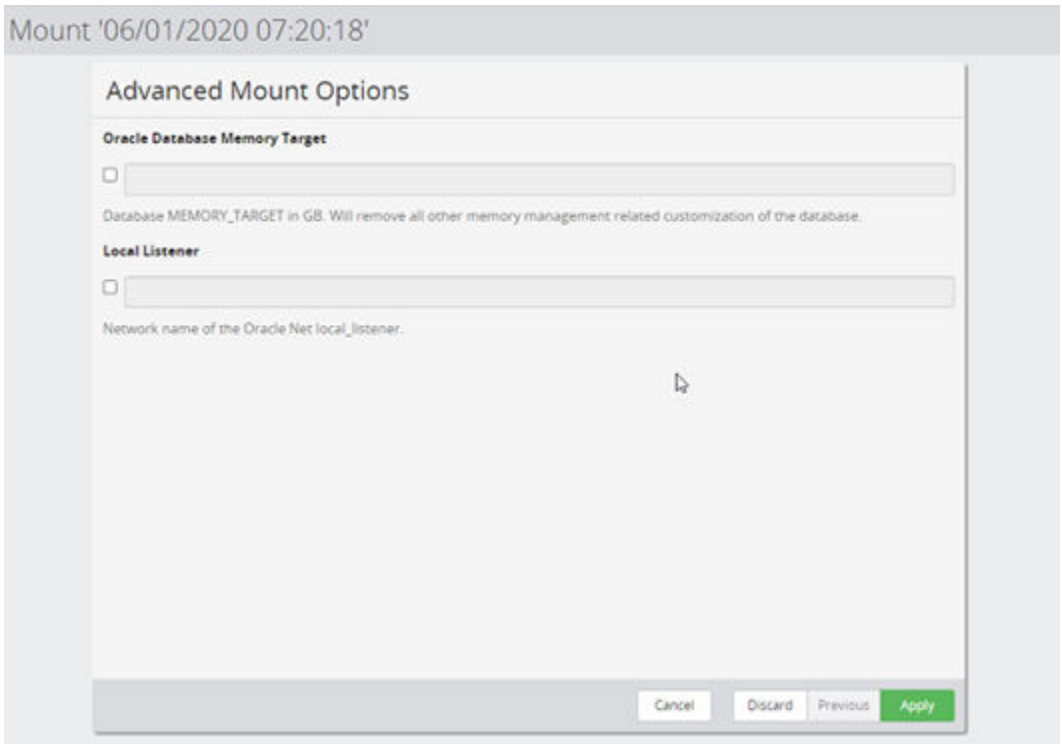


Figure 520 Mount Wizard - Advanced Mount Options

This page of the wizard is not be displayed if Restore only or Recover to last consistent state in backup is selected in a previous step.

Control	Description
Oracle database Memory Target	Sets the database `MEMORY_TARGET` in GB. Entering a value here will remove all other memory management related customization of the database. **Tip:** This allows Oracle databases from very powerful source systems to be deployed on less powerful systems. The PGA and SGA memory areas will be managed by Oracle within the given memory target.
Local Listener	Sets the network name of the Oracle Net `local_listener`.

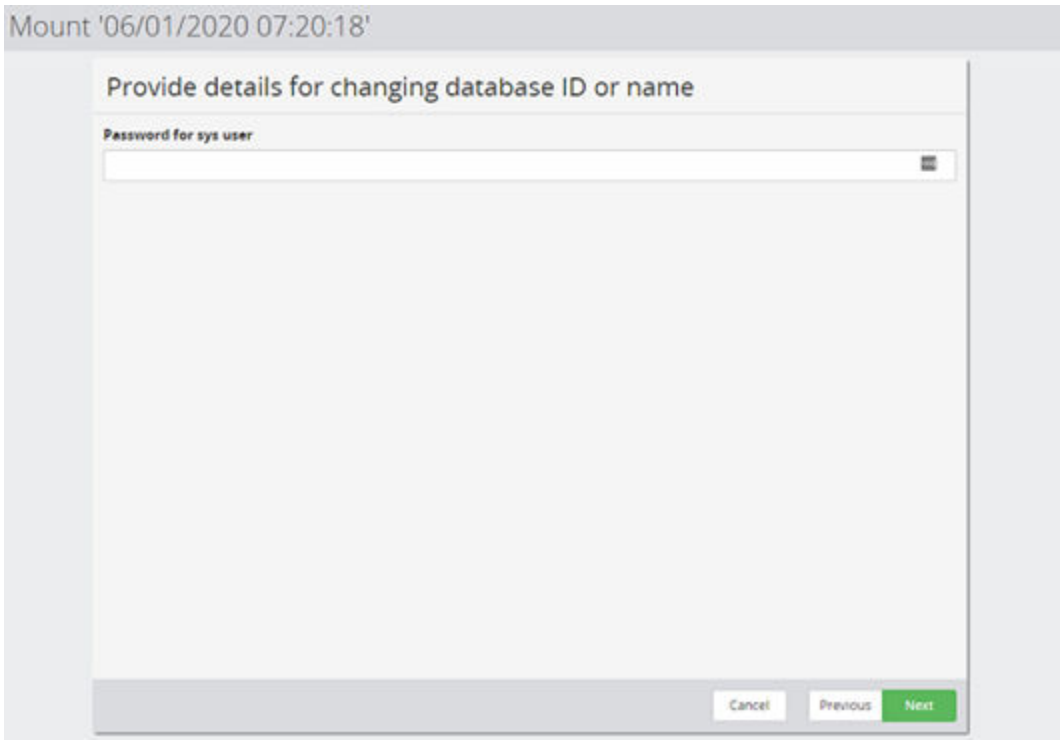


Figure 521 Mount Wizard - Provide details for changing database ID or name

This page of the wizard will only be displayed if either Change Oracle database ID (DBID) or Change Oracle database unique name and SID options are selected in a previous step.

Control	Description
Password for sys user	Depending on the Oracle version: ▪ Oracle 11g: The sys user password is required to change the Oracle database ID or database unique name. ▪ Oracle 12 or newer: This field can be left empty.

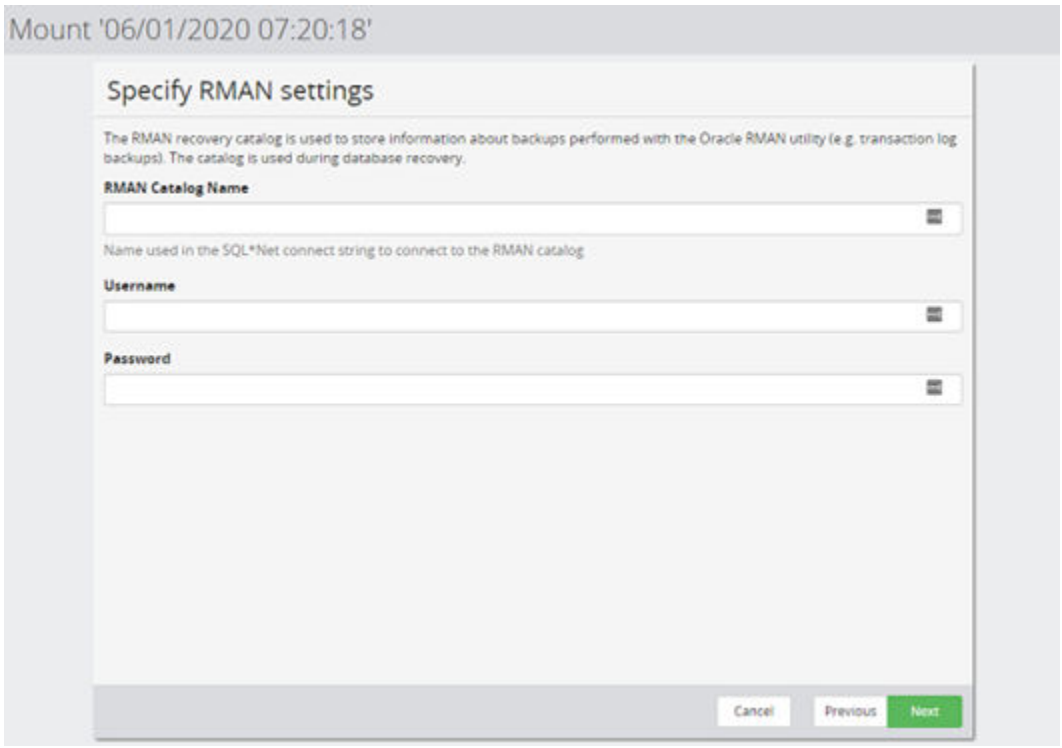


Figure 522 Mount Wizard - Specify RMAN credentials

This page of the wizard is not displayed if Restore only or Recover to last consistent state in backup is selected in a previous step.

Control	Description
RMAN Catalog Name	For RMAN only. Enter the RMAN Catalog Name as it is entered in the SQL*Net connect string to connect to the RMAN catalog.
Username	For RMAN only. Enter the username for the RMAN catalog.
Password	For RMAN only. Enter the password for the RMAN catalog.

Hitachi Block VMware Mount Wizard

This wizard is displayed when you mount a VMware snapshot or replication from a Hitachi Block device.

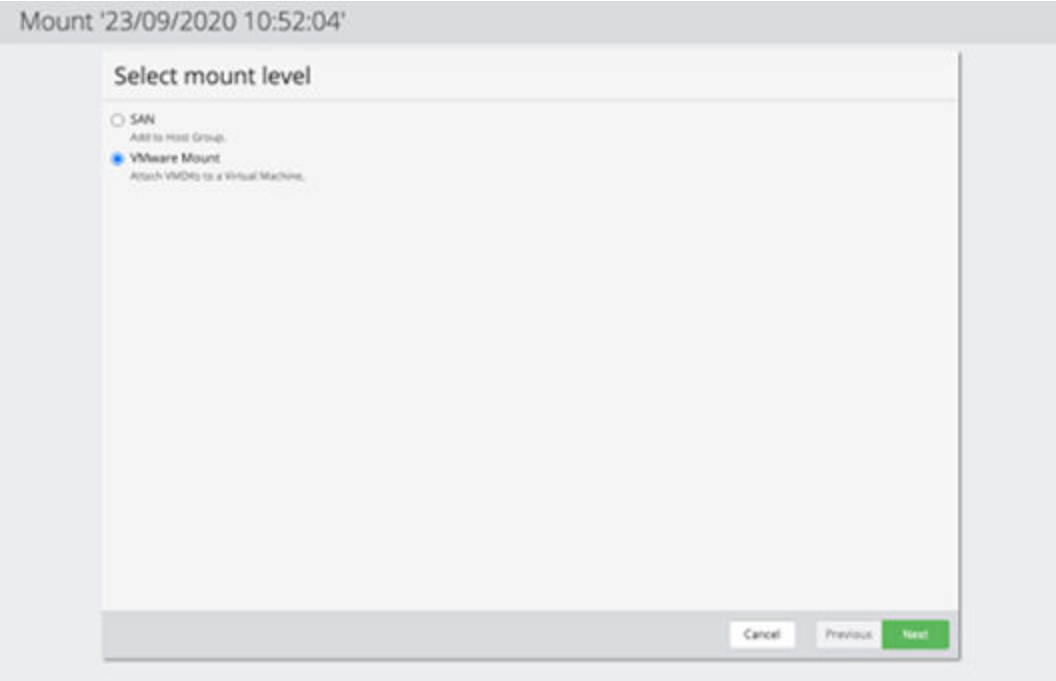


Figure 523 Mount VMware Wizard - Select Mount Level

Control	Description
SAN	Expose this record to a Host Group
VMware mount	Mount the disks of a VM from the record, to a target VM

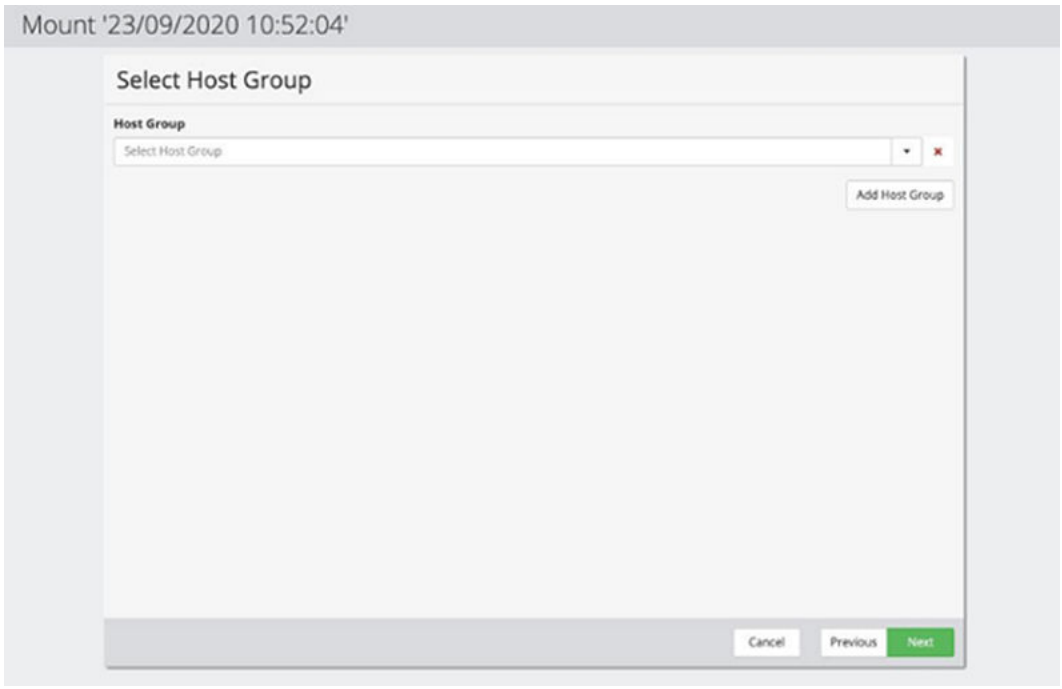


Figure 524 Mount VMware Wizard - Select Host Group

Control	Description
Host Group	Host Group to expose the record to. Multiple Host Groups can be added

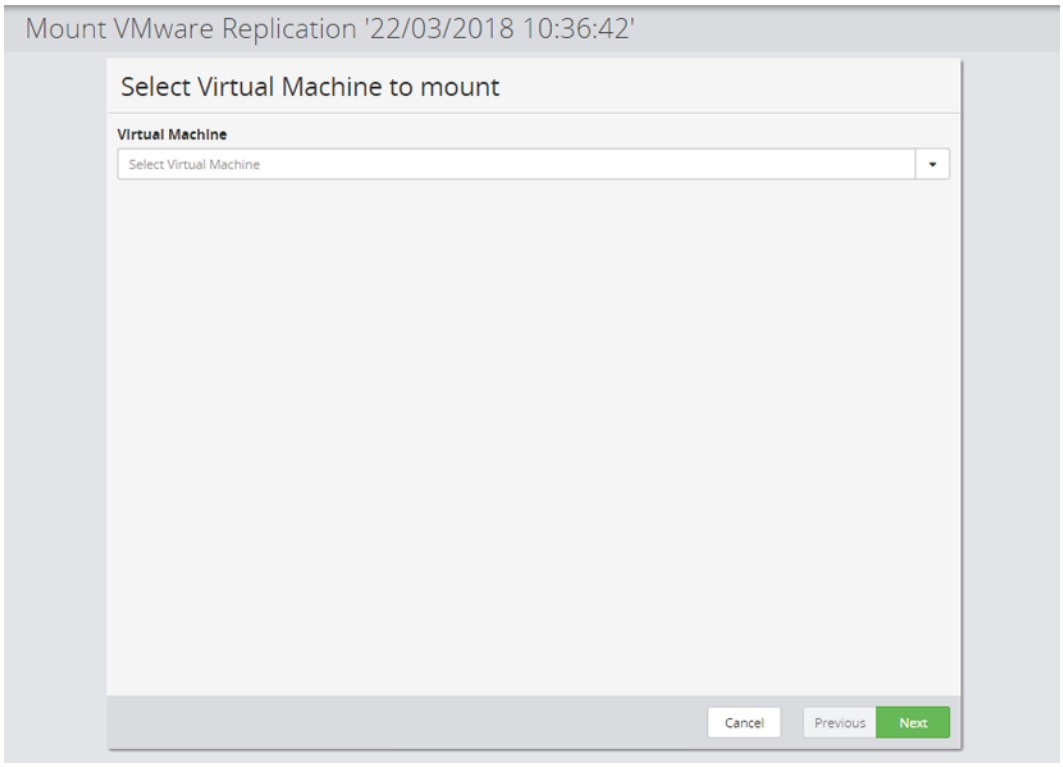


Figure 525 Mount VMware Wizard - Select Virtual Machine

Control	Description
Virtual Machine	Select the specific VM within this snapshot that is to have its disks mounted.

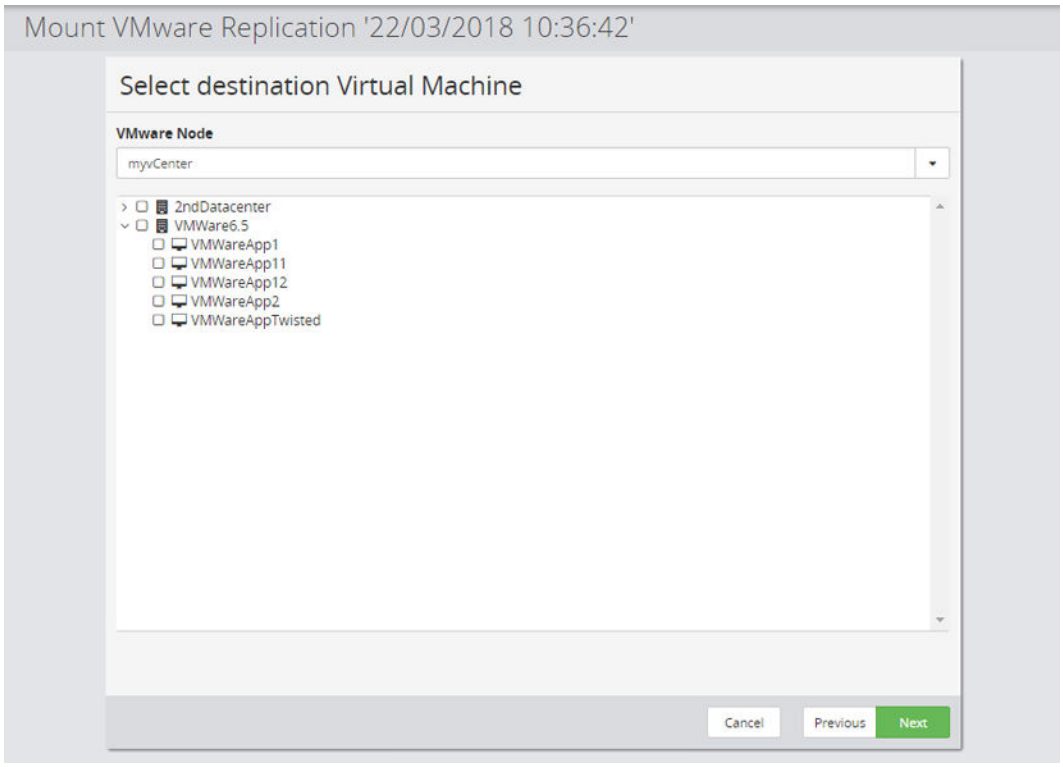


Figure 526 Mount VMware Wizard - Select Virtual Machine to Mount to

Control	Description
VMware Node	Select the VMware Host or vCenter where the VM's disks will be mounted.
Destination	Select the VMware Datacenter, sub-folder and VM where the VM's disks will be mounted.

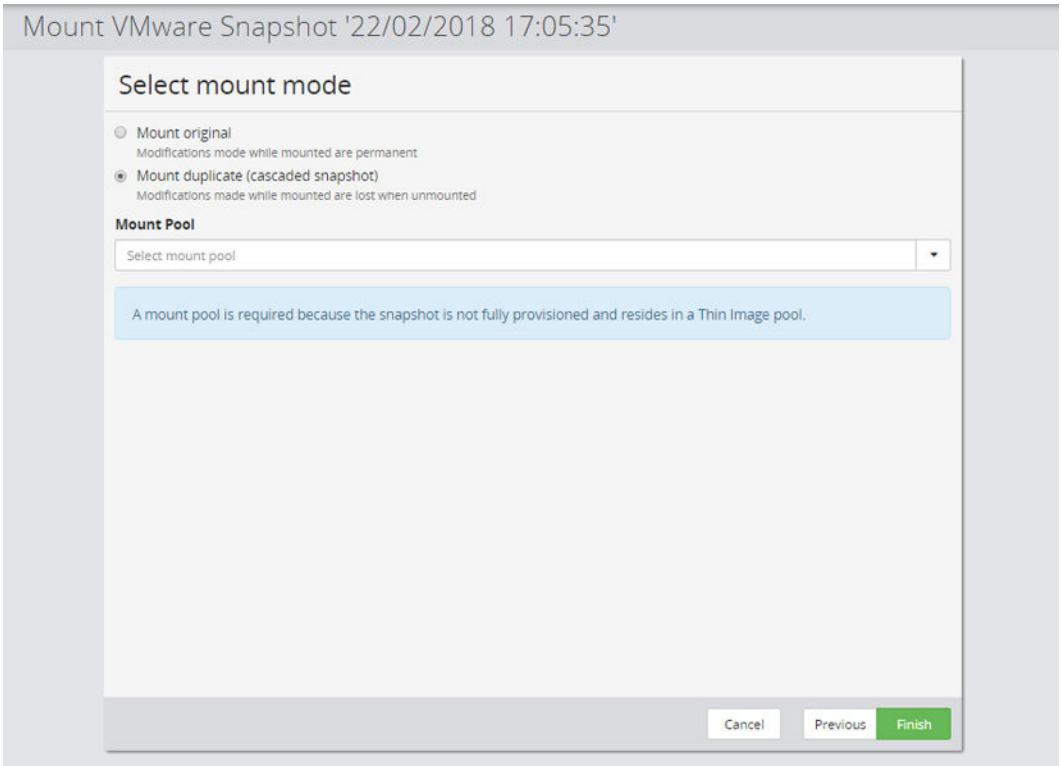


Figure 527 Mount VMware Wizard - Select Mount Mode

Control	Description
Mount Original	Mounts the VM's VMDKs using the replication or the original (Level 1) snapshot. **Caution:** Any changes made to the VMDKs will persist when they are unmounted.
Mount duplicate (cascaded snapshot)	Not available for replications. Only enabled if the original (Level 1) snapshot was created in cascade mode. Mounts the VM's VMDKs using a copy of the original snapshot (i.e. a Level 2 snapshot). **Caution:** Any changes made to the VMDKs will be lost when they are unmounted.
Mount Pool	Not available for replications. Depending on the parameters specified for the snapshot operation on the data flow, a mount pool might be required. A message is displayed in a blue rectangle to explain if and why a mount pool is required.

Hitachi Block Revert Wizard

This wizard allows a snapshot or local replication to be used to revert volume(s) to a prior state.

Caution: Reverting destroys any changes that were made to data on those volumes subsequent to the selected snapshot or local replication being taken.

Note:

- Before issuing a revert operation the user should ensure the volume being reverted is not having any other operation performed on it including being mounted, this can be confirmed in the Hitachi Block Replication Details (Storage) (on page 882). If it is mounted the revert will fail silently.
- For replications in a cascaded data flow, the immediate upstream and downstream volumes must be in an appropriate state (typically paused or suspended for swap). This is enforced by the storage array hardware.
- The revert operation can take some time to complete.

When reverting snapshots or local replications created by a policy containing an *Application Classification*, Protector will first display one of the following wizard pages that allow application specific options to be configured:

- Revert Wizard - Configure Oracle Recovery Options (on page 815)

Once any application specific options have been selected, the following wizard page will be displayed:

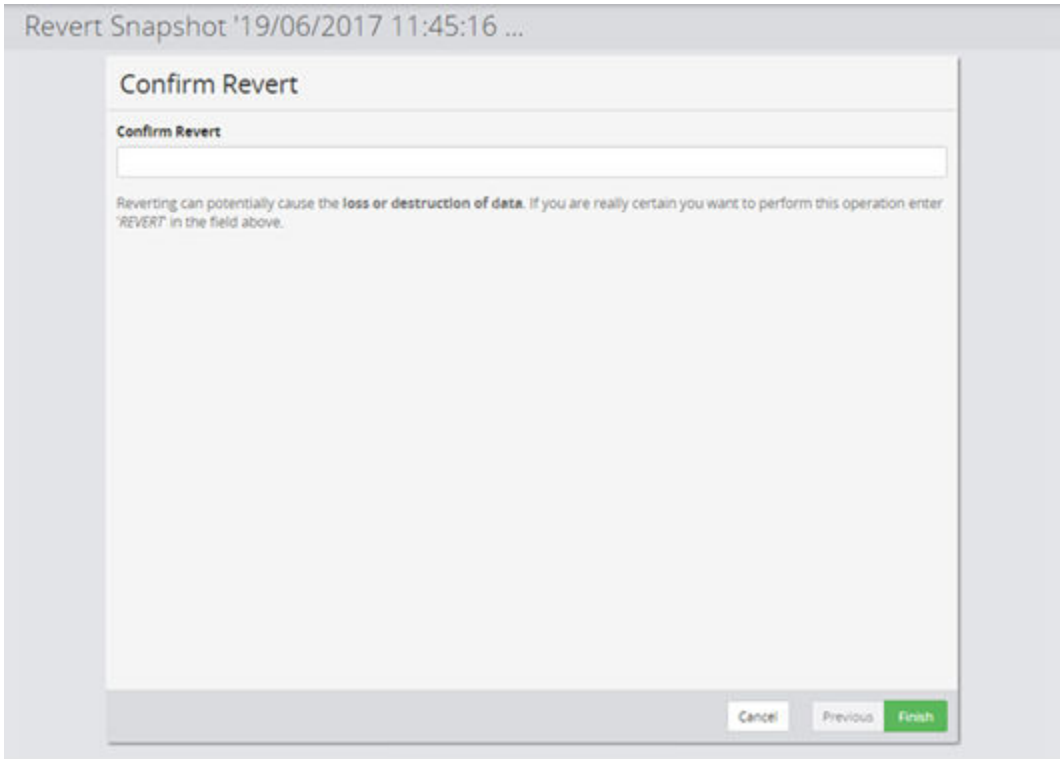


Figure 528 Revert Wizard - Confirm Revert

Control	Description
Confirm Revert	The word REVERT must be explicitly typed in to confirm the action.

Revert Wizard - Configure Oracle Recovery Options

When reverting snapshots or local replications created by a policy containing an *Oracle Database* classification, Protector will first display the following wizard page that allows application specific options to be configured.

Note: If an Oracle app node is of type DataGuard Standby then snapshot cannot be reverted.

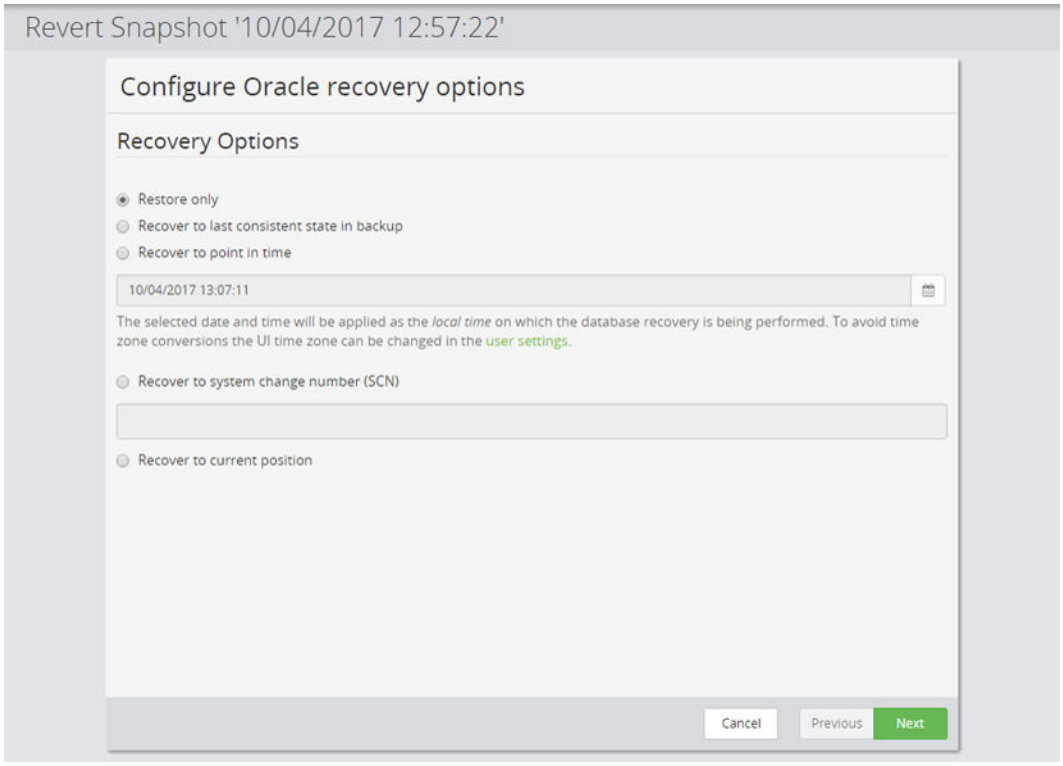


Figure 529 Revert Wizard - Configure Oracle Recovery Options

Control	Description	Logs Reset Post Recovery
Restore only	The database is simply reverted and it is left to the database administrator to recover manually.	No

Control	Description	Logs Reset Post Recovery
Recover to last consistent state in backup	The database is recovered to the consistent state which was captured by the backup. The database is brought online. This type of revert can be performed with the data in the backup alone; no RMAN catalogue is required.	Yes
Recover to point in time	A timestamp is entered (in 24 hour format: YYYY-MM-DD:HH:MM:SS), which defines the point in time to recover. The time entered must be after the time the snapshot was created and before the last available transaction. This option requires a connection to the RMAN catalog and logs which are shared with/available to the host.	Yes
Recover to system change number (SCN)	A system change number is entered which defines the change point to recover. The SCN entered must be after the snapshot was created and before the last available transaction. This option requires a connection to the RMAN catalog and logs which are shared with/available to the host.	Yes
Recover to current position	The database is recovered to the most current position possible. Because access to the latest redo logs is available, it is possible to recover to the last transaction. This requires a connection to the RMAN catalog and logs which are shared with/available to the host.	No

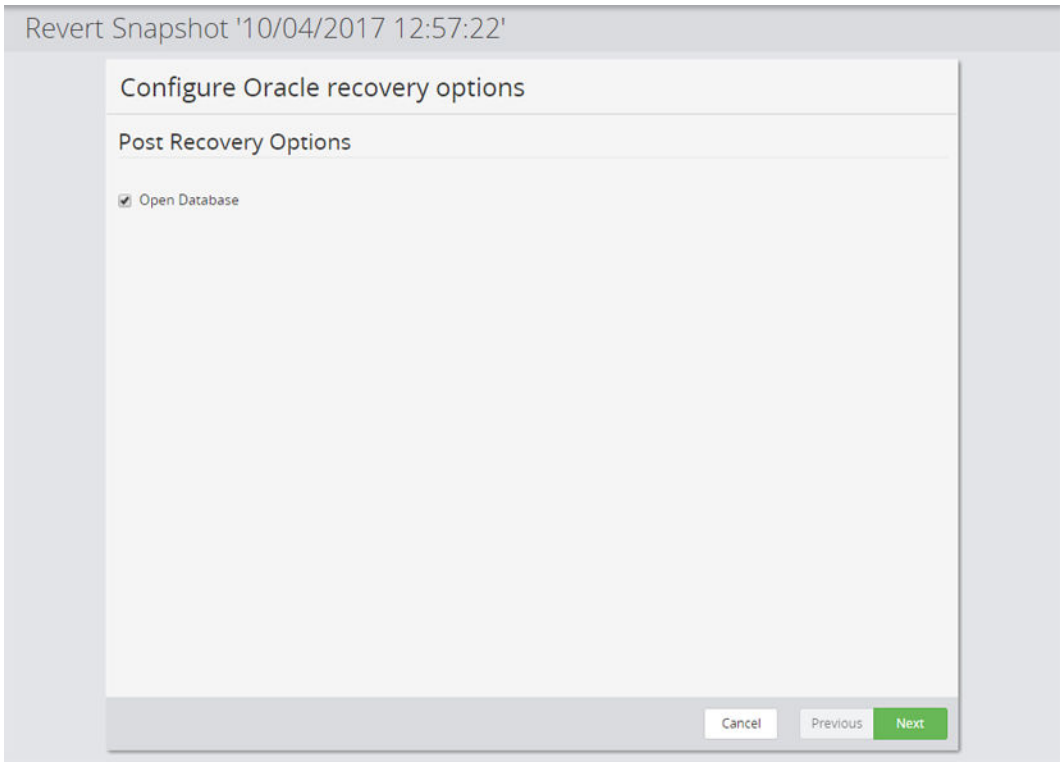


Figure 530 Revert Wizard - Post Recovery Options

Control	Description
Open Database	If selected, then after recovery the database will placed in the *OPEN* state using the *RESETLOGS* or *NORESETLOGS* option, as per the requirements of the database. Otherwise the database will be left in the *MOUNT* state.

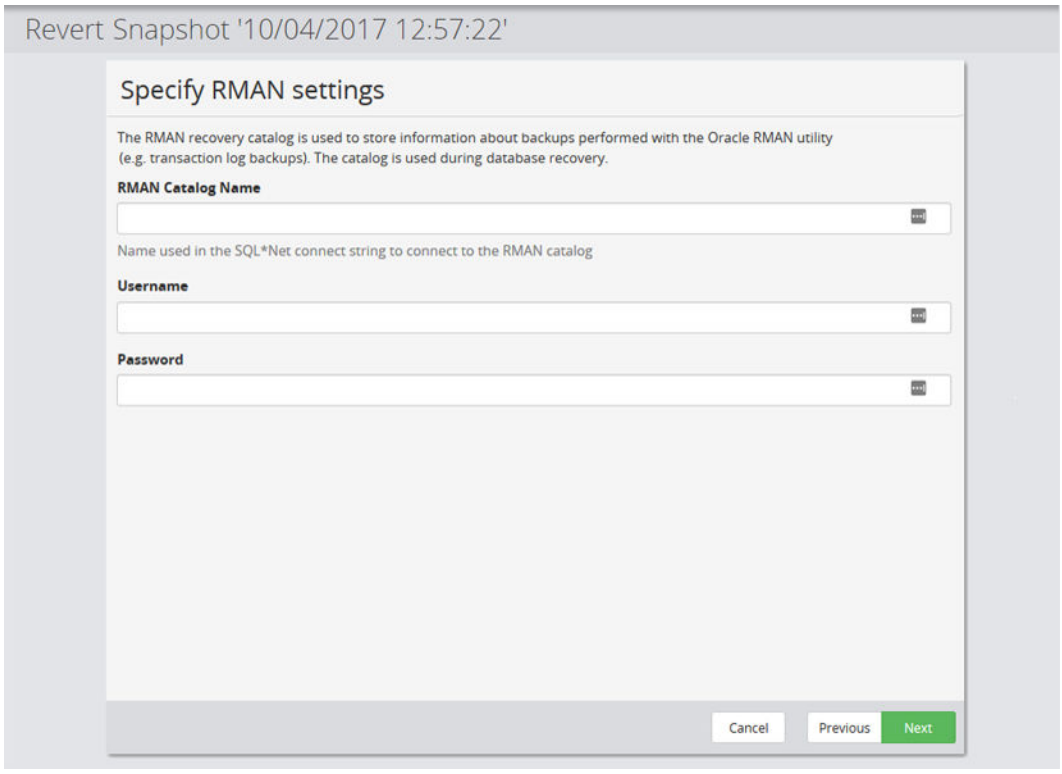


Figure 531 Revert Wizard - Specify RMAN Credentials

This page of the wizard is not displayed if Restore only or Recover to last consistent state in backup is selected in a previous step.

Control	Description
RMAN Catalog Name	For RMAN only. Enter the RMAN Catalog Name as it is entered in the SQL*Net connect string to connect to the RMAN catalog.
Username	For RMAN only. Enter the username for the RMAN catalog.
Password	For RMAN only. Enter the password for the RMAN catalog.

Hitachi Block Replication Swap Wizard

This wizard causes the direction of a replication to be reversed so that the secondary replicates to the primary.

Caution: When a replication is swapped, the secondary becomes writeable and the primary becomes read-only.

Note:

- If the swap cannot be completed then the replication pair enters the SSWS (suspended for swapping) state until the swap can be completed. The swap operation must be re-done by the user to complete a suspended swap.
- The flow direction of a replication pair should ONLY be determined by referring to the Storage Inventory (on page 858). Primary and secondary volume information shown in the Log Attachments Dialog (on page 519) should not be used to infer the flow direction following a swap.

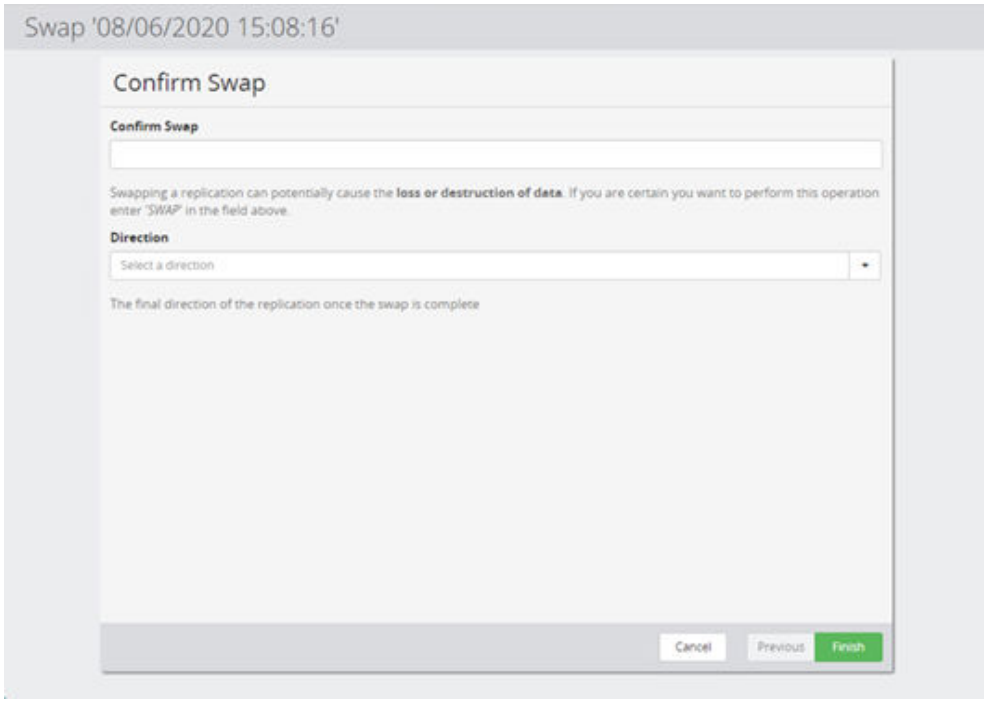


Figure 532 Replication Swap Wizard

Control	Description
Confirm Swap	The word SWAP must be explicitly typed in to confirm the action.
Direction	The final intended direction of the replication after the swap successfully completes.

Hitachi Block Snapshot Details (Restore)

This page shows the details of a Snapshot on a Hitachi Block storage device.

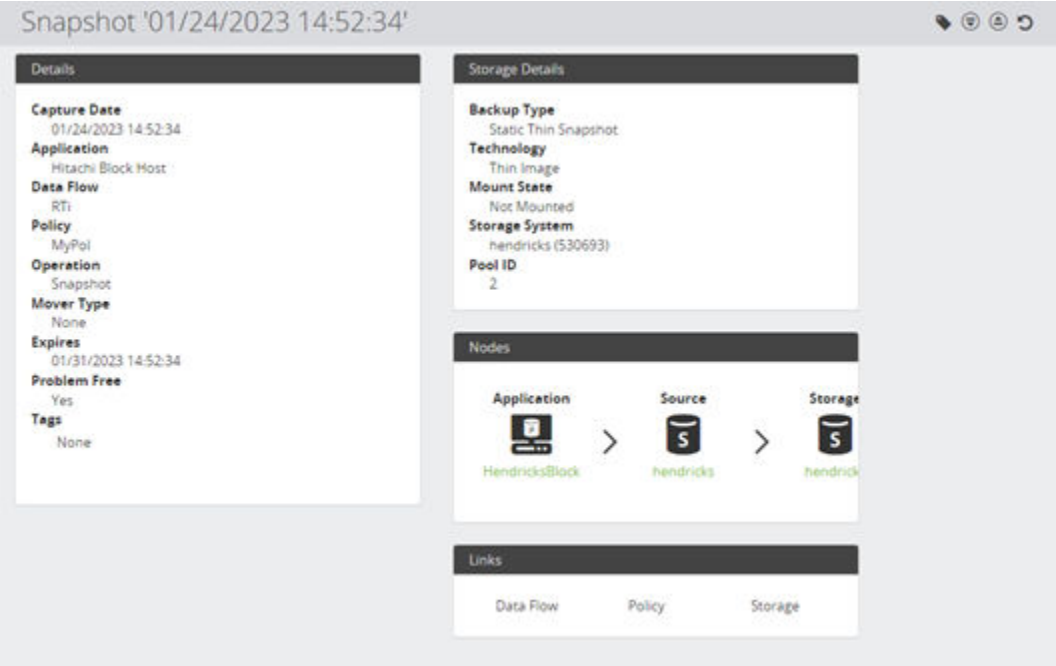


Figure 533 Snapshot Details (example shows VMware snapshot with included VMs listed)

Control	Description
Restore	Enabled only if a VMware snapshot is selected. Opens the Hitachi Block VMware Snapshot Restore Wizard (on page 766) to guide you through mounting the Snapshot.
Mount	Opens the Hitachi Block Snapshot or Replication Mount Wizard (on page 772) to guide you through mounting the Snapshot.
Unmount	Unmounts the selected Snapshots.
Revert	Enabled for non-VMware snapshots. Opens the Hitachi Block Revert Wizard (on page 813). to guide you through reverting the snapshot.
Details	Lists the Snapshot details. **Tip:** For Thin Image snapshots the *Type* field is set to *Static Thin Snapshot* indicating that the record is unchanging and requires the P-VOL to reconstruct an S-VOL.
Storage	Lists the following Storage details.

Control	Description
VMware Details for '<server>'	Only displayed for VMware snapshots. Lists the VMs included in the snapshot, the guest OS and whether the VM has been restored from the original snapshot and is thus no longer available to restore.
Nodes	Shows the nodes involved in the snapshot and opens the Node Details (on page 650) for node.
Links	Provides links to the following pages that provide additional information about the replication: ▪ Data Flow - opens the Data Flow Details (on page 491) where the snapshot operation is assigned. ▪ Policy - opens the Policy Details (on page 737) where the snapshot operation is defined. ▪ Storage - opens the Hitachi Block Snapshot Details (Storage) (on page 871) view for the snapshot.

Hitachi Block Replication Details (Restore)

This page shows the details of a Replication on a Hitachi Block storage device.

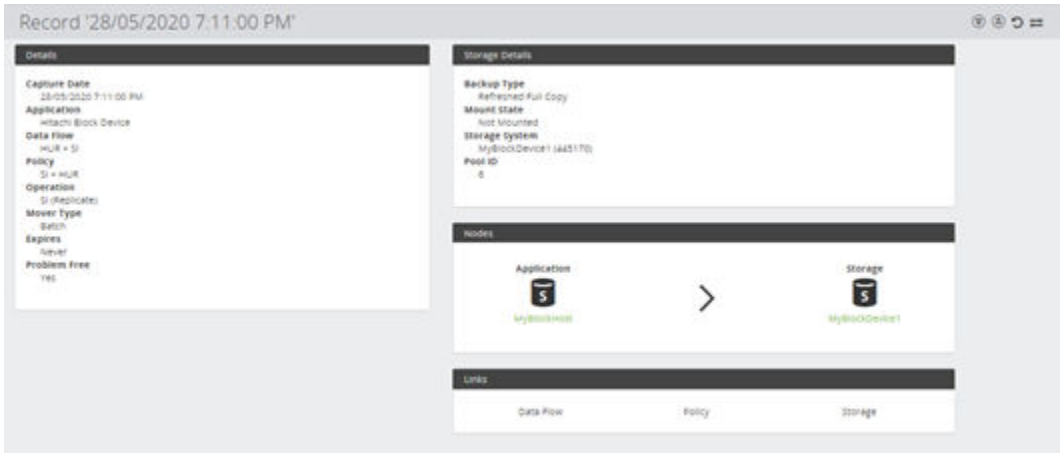


Figure 534 Replication Details

Control	Description
Mount	Opens the Hitachi Block Snapshot or Replication Mount Wizard (on page 772) to guide you through mounting the Replication.
Unmount	Unmounts the Replication.

Control	Description
Swap	Opens the Hitachi Block Replication Swap Wizard (on page 818) to guide you through swapping the Replication direction.
Nodes	Lists the nodes involved with the replication.
Storage	Lists the Storage details.
Links	Provides links to the following pages that provide additional information about the replication: ▪ Data Flow - opens the Data Flow Details (on page 491) where the replication operation is assigned. ▪ Policy - opens the Policy Details (on page 737) where the replication operation is defined. ▪ Storage - Opens the Hitachi Block Replication Details (Storage) (on page 882) for the replication, to help find and manually teardown deactivated replications.

File System Node Mount Wizard

When mounting block snapshots or replications created by a policy containing a File System classification, it is possible to choose between various mount levels

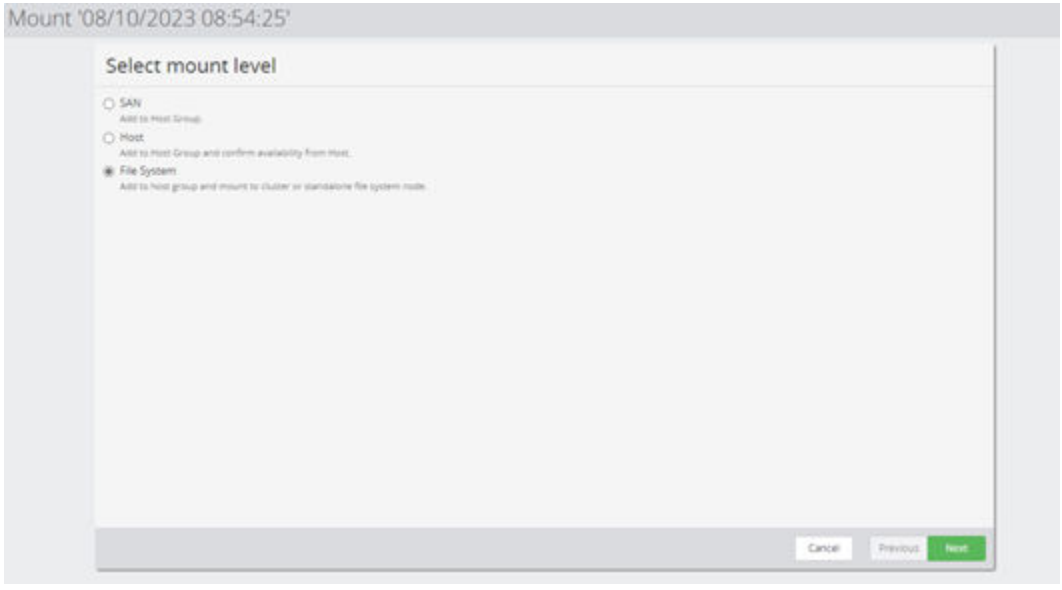


Figure 535 File System - Select Mount Level

Control	Description
SAN	Adds the snapshot or replication to a Host Group
Host	Adds the snapshot or replication to a Host Group and confirms that it is available on the specified Host.
File System	Adds the snapshot or replication to a Host Group and mounts the protected volumes to a File System node. In the case of Windows Server Failover Clusters, the snapshot can also be added as clustered disks

In case “File System” is selected, Protector will display the following additional wizard page that allows the application-specific options to be configured.

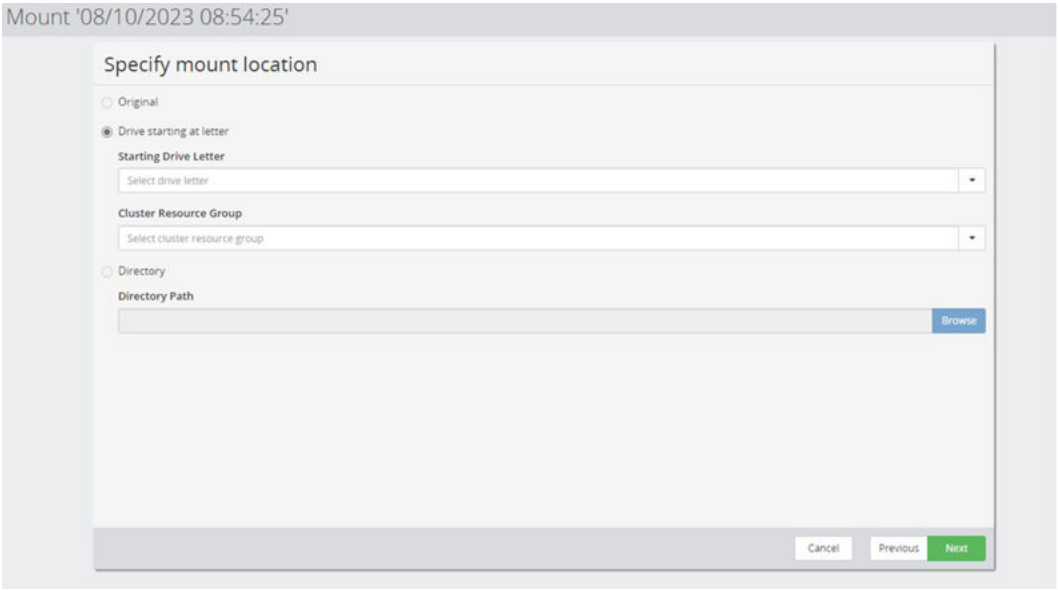


Figure 536 File System - Select Mount Location

Note: Not all options may be enabled depending on the type of snapshot or replication and the target File System node you want to mount to.

Control	Description
Original	Mount the snapshot or replication to the same mount point or drive letter used on the source at the time of backup.

Control	Description
Starting Drive Letter	Select which drive letter the snapshot or replication should be mounted to.
Cluster Resource Group	When mounting to a Windows Server Failover Cluster, select which resource group the snapshot or replication should be added to. **Note:** This option will not be offered for standalone File System nodes.
Directory	Select which path (mount point) the snapshot or replication should be mounted to.

File System Node Revert Wizard

When reverting snapshots or replications created by a policy containing a File System classification, Protector will first display the following wizard.

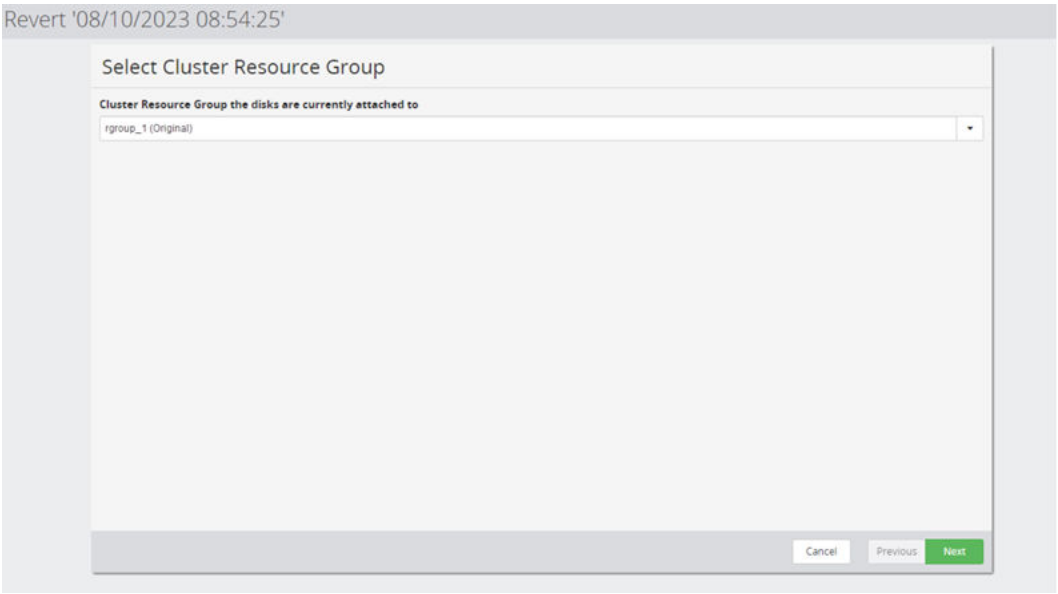


Figure 537 File System - Select Cluster Resource Group

Control	Description
Cluster Resource Group	Select the cluster resource group the disks are currently connected to. The cluster resource group that the disks originally belonged to will be automatically selected. However, if you have moved or recreated the resource group since the backup was created, you should choose the current resource group. For standalone File System nodes, the default selection will be "Standalone".

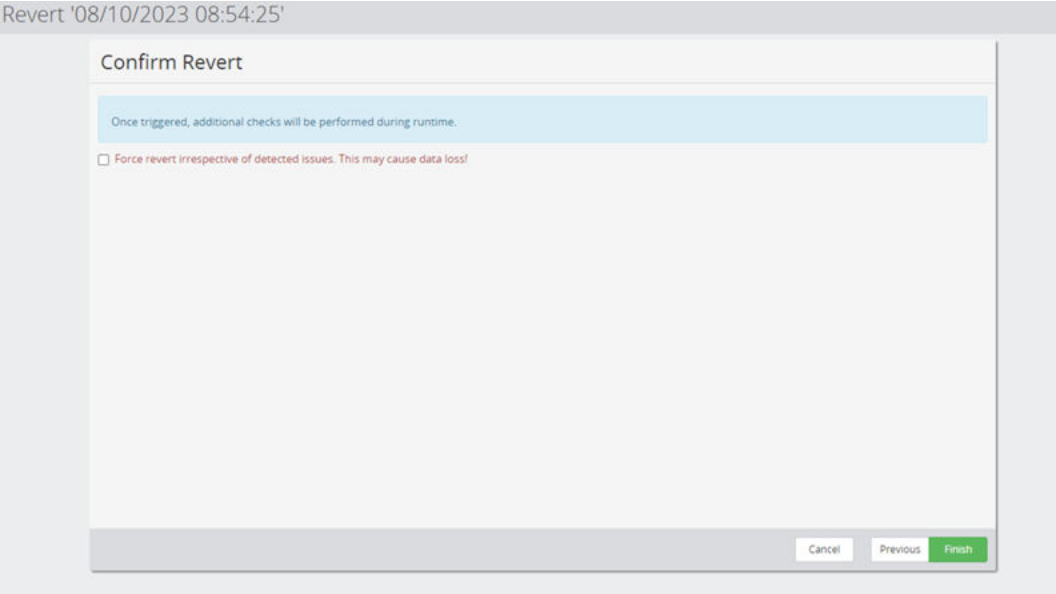


Figure 538 File System - Confirm Revert

Control	Description
Revert warnings	This wizard page will list issues which will prevent Protector from performing a regular revert. In case any issues are listed here, it is recommended to mount the data instead of reverting it and restore the data manually from the mounted volume. If any issues are found, it will only be possible to continue by selecting the force flag.

Control	Description
Fore revert option	This flag will instruct Protector to ignore problems preventing a regular revert. **WARNING:** Only select this option if you understand the implications. In most cases, it is better to mount the block-based backup and restore the data manually. **Selecting this may cause data loss on the source system.**

Host Based Backup Restore Options

This section describes the Host Based Restore options, accessed via the Restore Inventory (on page 763). Refer to the following sections for details on how to restore the various hast based protections operations.

Restore Repository Snapshot Wizard - File System

This wizard is displayed when you restore a file system snapshot from a repository store.

Caution: Be careful when restoring data to a location that is in the scope of one or more Ops Center Protector policies. The restored data will not be incorporated by Ops Center Protector into its data flows until after a resynchronization has occurred and synchronization errors for the restored data can be logged until that point. Perform a manual resynchronization to correct the inconsistent state.

Note: The following data restoration limitations apply:

- Because of issues with translating security streams, data originating from Linux machines cannot be restored to Windows machines and vice versa. Data can only be restored to compatible OS types.

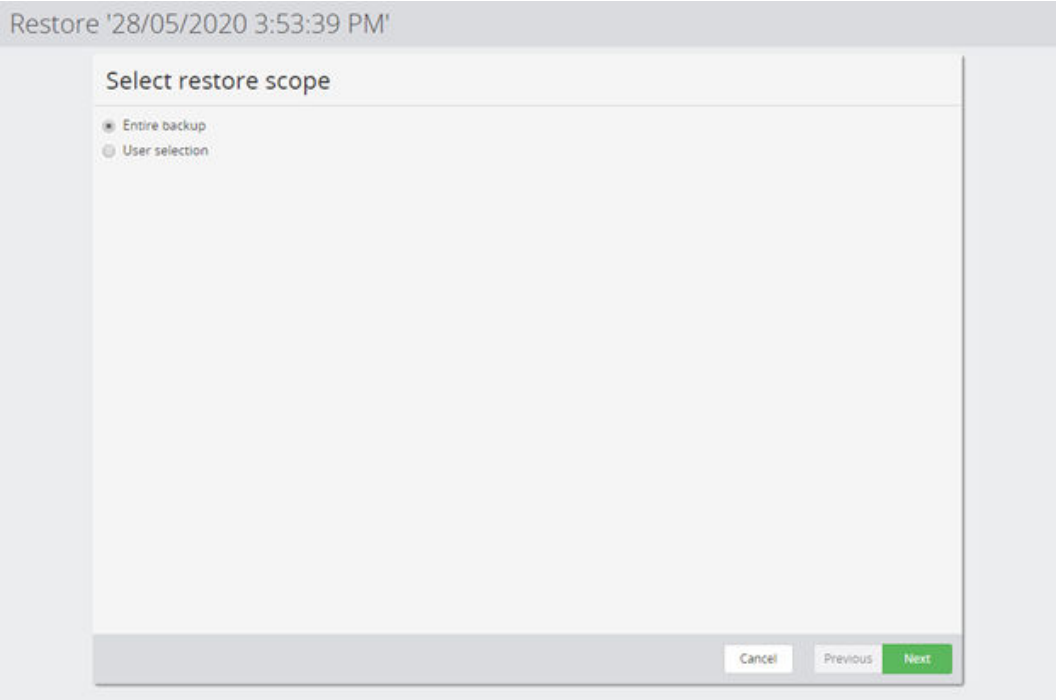


Figure 539 Restore File System Wizard - Select restore scope

Control	Description
Entire snapshot	The entire contents of the snapshot will be restored. The wizard displays the **Select restore options** page when Next is clicked.
User Selection	Only those files selected by the user will be restored. The wizard displays the **Select files and directories to restore** page when Next is clicked.

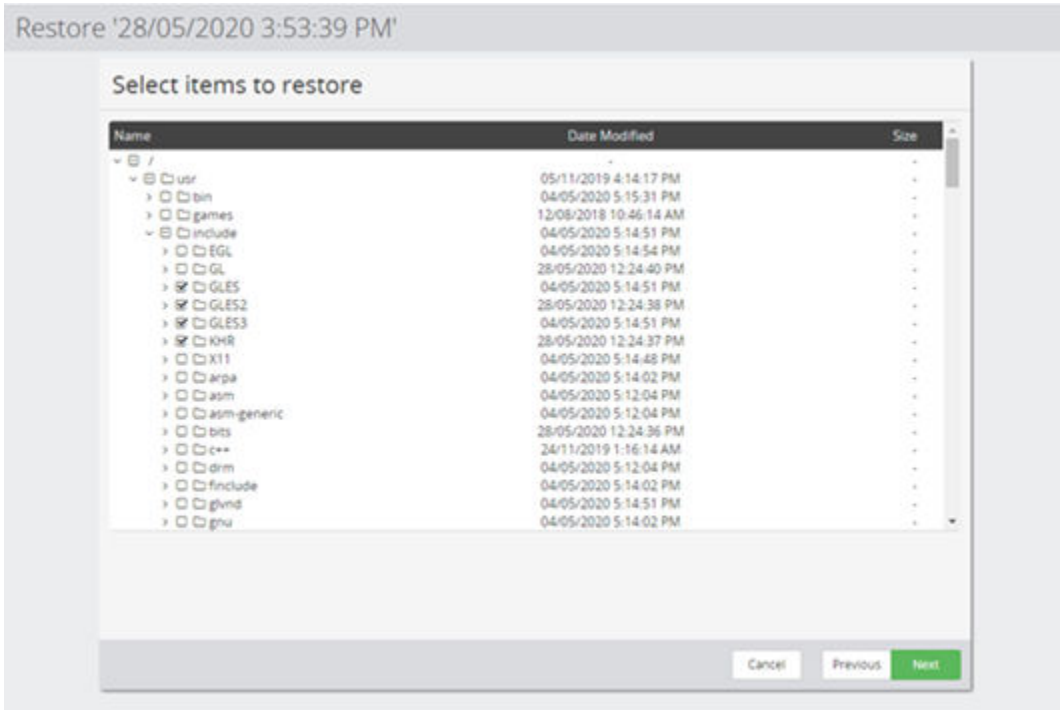


Figure 540 Restore File System Wizard - Select files and directories to restore

Control	Description
File System Contents	Select the directories and/or files you want to restore by clicking the checkbox to the left of its name (a tick mark is displayed next to every file that will be restored). To expand/collapse a folder, click the arrow symbol to the left of its name.

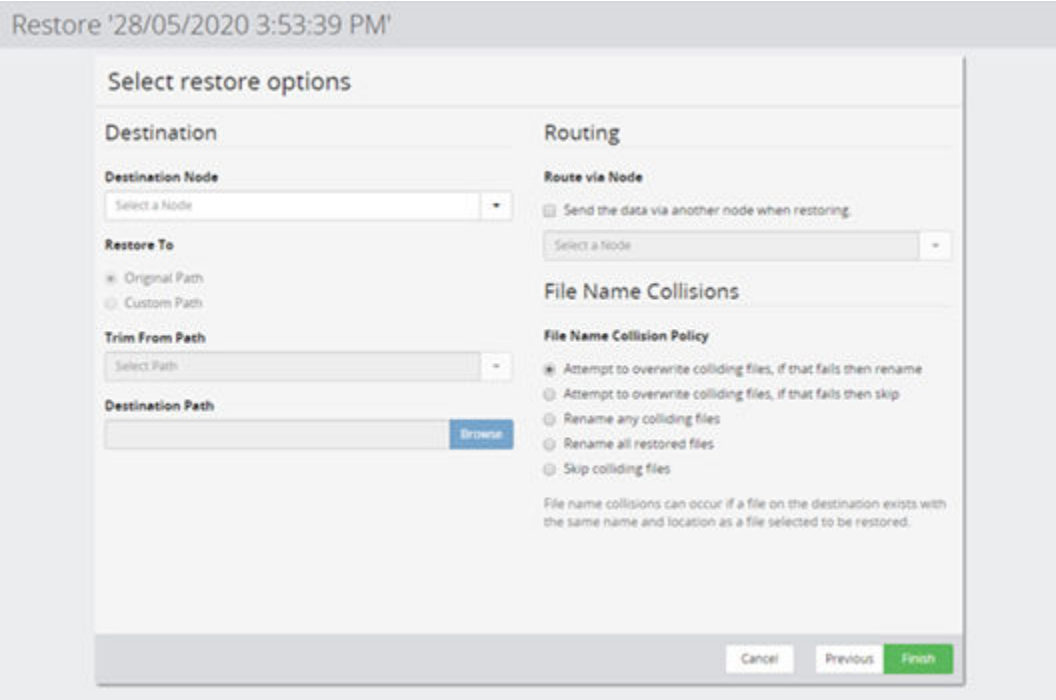


Figure 541 Restore File System Wizard - Select restore options

Control	Description
Destination Node	The target node for the restore.
Original Path	Use the original path names when restoring files.
Custom Path	Use a new path based on the Trim From Paths and Destination Path specification. Refer to the Trim From Path Examples in Destination Template Wizard (on page 379).
Trim From Path	Presents a list of path prefixes that can be removed from the files being restored, whilst ensuring files remain grouped together in their respective folders.
Destination Path	The path specification to be prepended once the paths are trimmed.
Route via Node	Select a node to route via. Primarily used for Internet Connected Nodes where the destination is not directly connected to the repository node, so a proxy must be specified.
Attempt to overwrite colliding files, if that fails then rename	If overwriting fails then the file will be restored by prepending the restore version (that is `(VersionNumber)_filename.ext`) to the file name.

Control	Description
Attempt to overwrite colliding files, if that fails then skip	If overwriting fails then the new file will NOT be restored.
Rename any colliding files	Differentiates all colliding files by prepending the restore version (that is `(VersionNumber)_filename.ext`) to the file name.
Rename all restored files	Differentiates all restored files by prepending the restore version (that is `(VersionNumber)_filename.ext`) to the file name.
Skip colliding files	Do not restore colliding files.

Restore Repository Snapshot Wizard – Ops Center

This wizard is displayed when you restore a Ops Center snapshot from a repository store.

Caution: Be careful when restoring data to a location that is in the scope of one or more Ops Center Protector policies. The restored data will not be incorporated by Ops Center Protector into its data flows until after a resynchronization has occurred and synchronization errors for the restored data can be logged until that point. Perform a manual resynchronization to correct the inconsistent state.

Note: The following data restoration limitations apply:

- Because of issues with translating security streams, data originating from Linux machines cannot be restored to Windows machines and vice versa. Data can only be restored to compatible OS types.

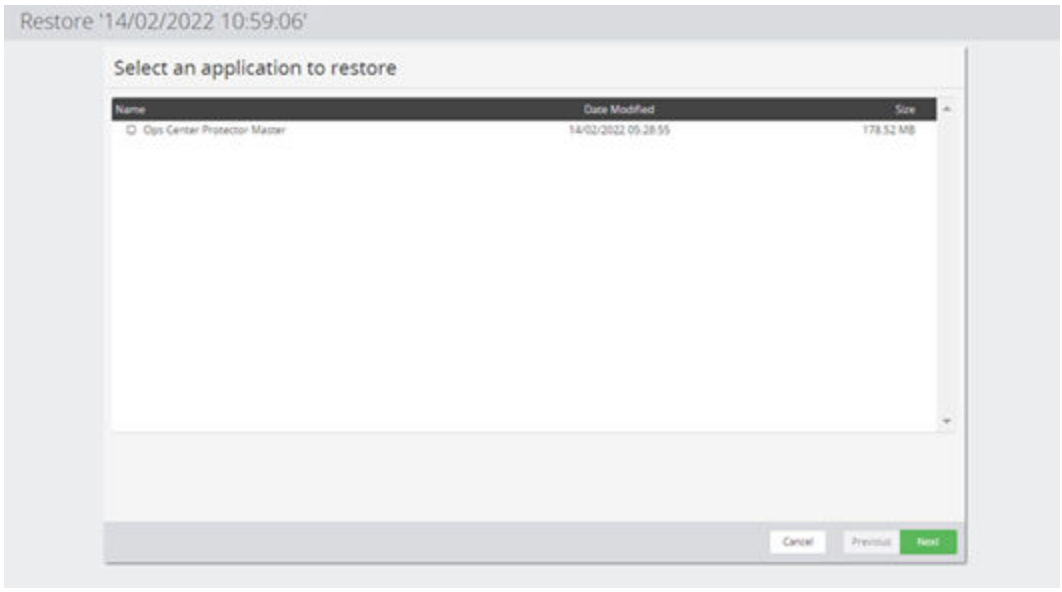


Figure 542 Restore Ops Center Application Wizard - Select application to restore

Control	Description
Applications to Restore	Select the applications you want to restore by clicking the checkbox to the left of the application name (a tick mark is displayed next to every file that will be restored). The wizard displays the **Select restore options** page when Next is clicked.

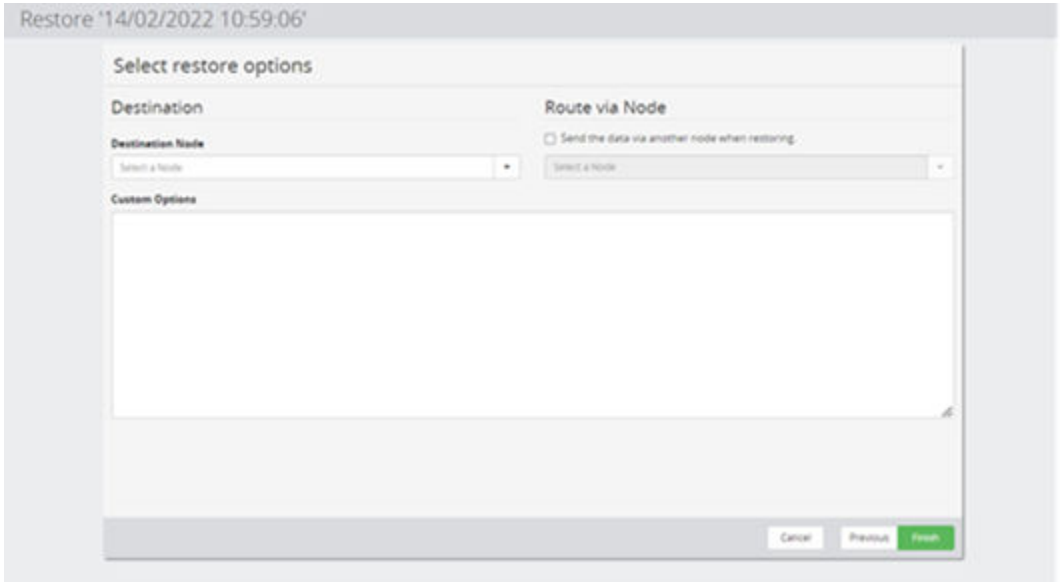


Figure 543 Restore Ops Center Application Wizard - Select restore options

Control	Description
Destination Node	The target node for the restore. This destination node must be set to the node hosting the application to be recovered.
Route via Node	Select a node to route via. Primarily used for Internet Connected Nodes where the destination is not directly connected to the repository node, so a proxy must be specified.
Custom Options	Custom options may be entered in the text box - these options will be saved to a file and that file given to the applications 'restore command/ script s parameter '%3'. These should only be entered if stated in the documentation for the application.

Restore HCP Snapshot Wizard - File System

This wizard is displayed when you restore a file system snapshot from an HCPstore.

Caution: Be careful when restoring data to a location that is in the scope of one or more Ops Center Protector policies. The restored data will not be incorporated by Ops Center Protector into its data flows until after a resynchronization has occurred and synchronization errors for the restored data can be logged until that point. Perform a manual resynchronization to correct the inconsistent state.

Note: The following data restoration limitations apply:

- Becauseof issues with translating security streams, data originating from Linux machines cannot be restored to Windows machines and vice versa. Data can only be restored to compatible OS types.

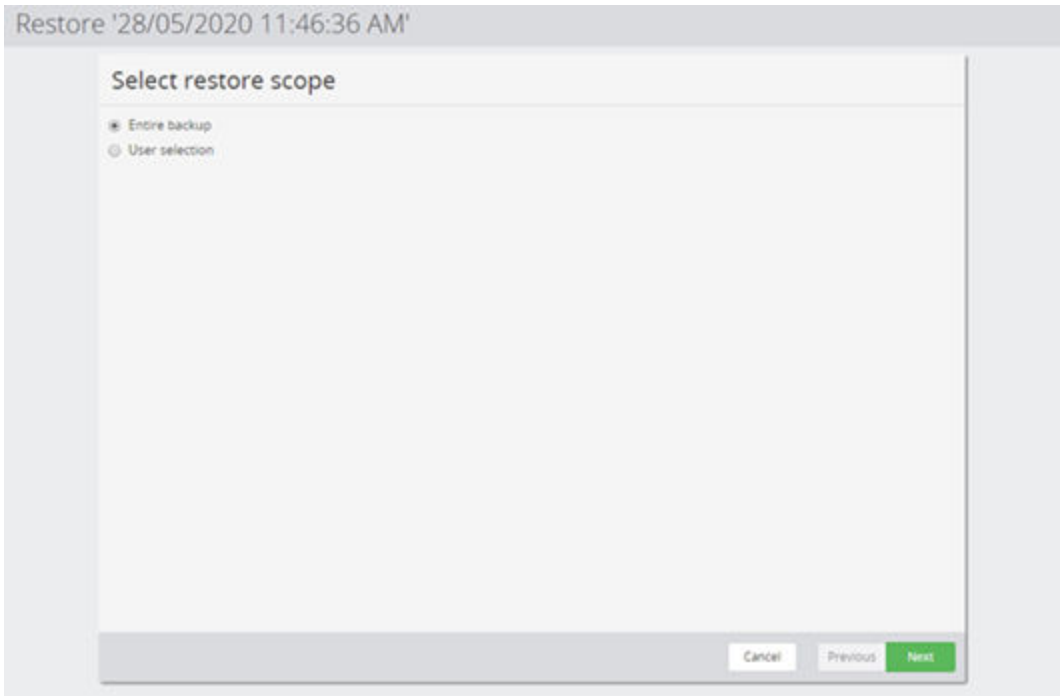


Figure 544 Restore HCP File System Wizard - Select restore scope

Control	Description
Entire snapshot	The entire contents of the snapshot will be restored. The wizard displays the **Select restore options** page when Next is clicked.
User Selection	Only those files selected by the user will be restored. The wizard displays the **Select files and directories to restore** page when Next is clicked.

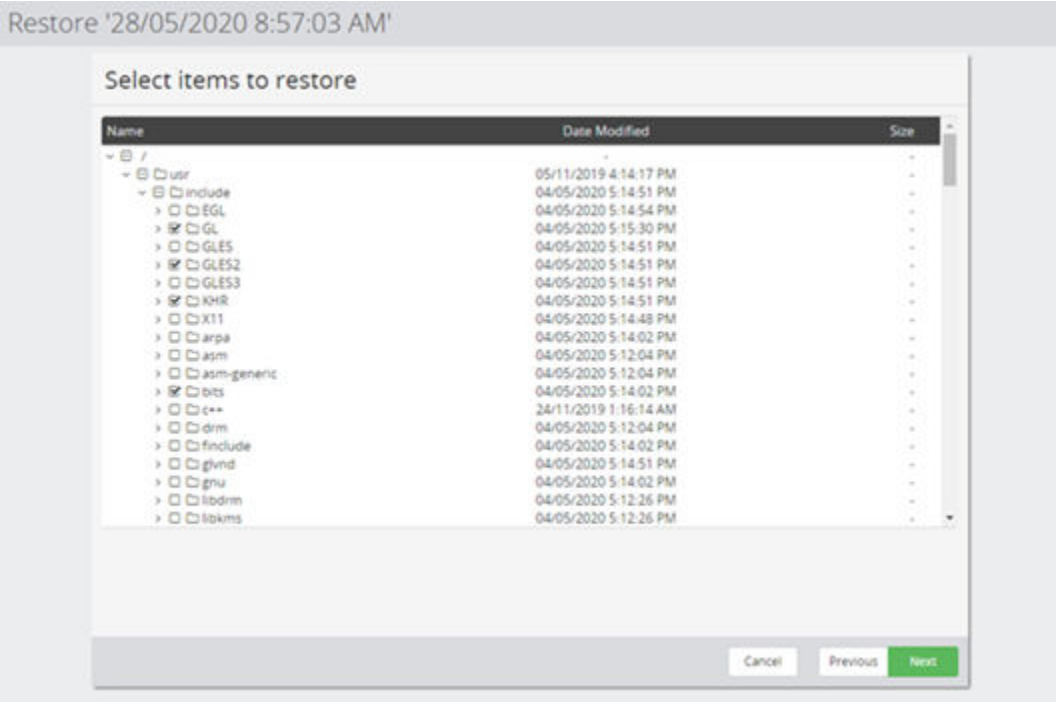


Figure 545 Restore HCP File System Wizard - Select items to restore

Control	Description
File System Contents	Select the directories and/or files you want to restore by clicking the checkbox to the left of its name (a tick mark is displayed next to every file that will be restored). To expand/collapse a folder, click the plus/minus symbol (+/-) to the left of its name.

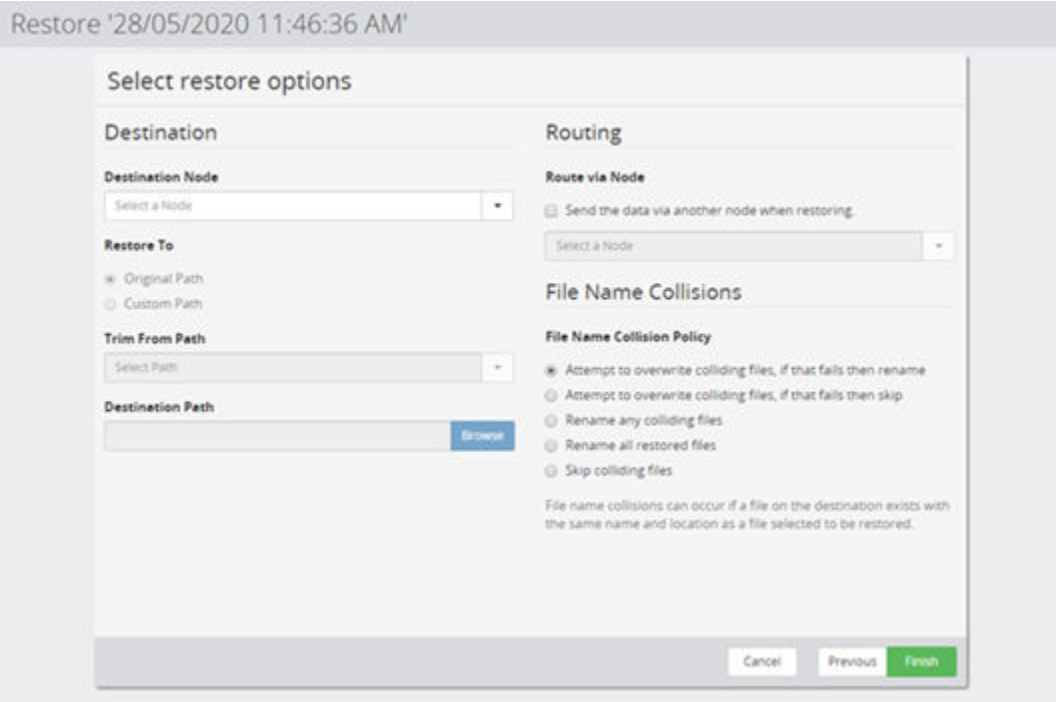


Figure 546 Restore HCP File System Wizard - Select restore options

Control	Description
Destination Node	The target node for the restore.
Original Path	Use the original path names when restoring files.
Custom Path	Use a new path based on the Trim From Paths and Destination Path specification. Refer to the Trim From Path Examples in Destination Template Wizard (on page 379).
Trim From Path	Presents a list of path prefixes that can be removed from the files being restored, whilst ensuring files remain grouped together in their respective folders.
Destination Path	The path specification to be prepended once the paths are trimmed.
Route via Node	Select a node to route via. Primarily used for Internet Connected Nodes where the destination is not directly connected to the repository node, so a proxy must be specified.
Attempt to overwrite colliding files, if that fails then rename	If overwriting fails then the file will be restored by prepending the restore version (that is `(VersionNumber)_filename.ext`) to the file name.

Control	Description
Attempt to overwrite colliding files, if that fails then skip	If overwriting fails then the new file will NOT be restored.
Rename any colliding files	Differentiates all colliding files by prepending the restore version (that is `(VersionNumber)_filename.ext`) to the file name.
Rename all restored files	Differentiates all restored files by prepending the restore version (that is `(VersionNumber)_filename.ext`) to the file name.
Skip colliding files	Do not restore colliding files.

Restore Amazon S3 / HCP Cloud Scale Snapshot Wizard - File System

This wizard is displayed when you restore a file system snapshot from an S3 store.

Caution: Be careful when restoring data to a location that is in the scope of one or more Ops Center Protector policies. The restored data will not be incorporated by Ops Center Protector into its data flows until after a resynchronization has occurred and synchronization errors for the restored data can be logged until that point. Perform a manual resynchronization to correct the inconsistent state.

Note: The following data restoration limitations apply:

- Because of issues with translating security streams, data originating from Linux machines cannot be restored to Windows machines and vice versa. Data can only be restored to compatible OS types.

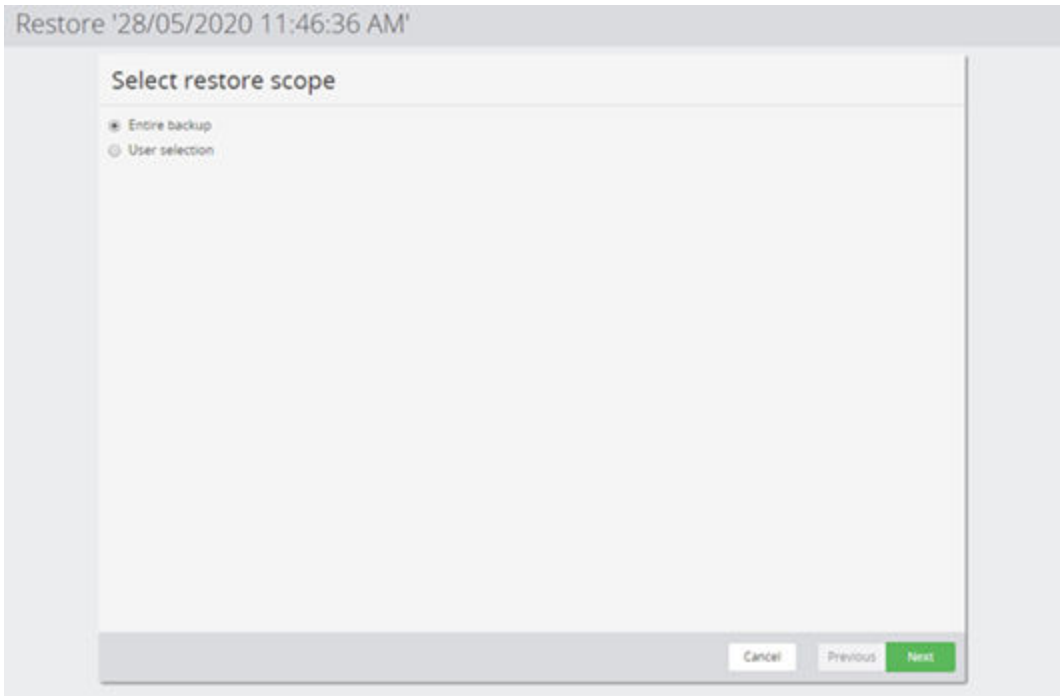


Figure 547 Restore S3 File System Wizard - Select restore scope

Control	Description
Entire snapshot	The entire contents of the snapshot will be restored. The wizard displays the **Select restore options** page when Next is clicked.
User Selection	Only those files selected by the user will be restored. The wizard displays the **Select files and directories to restore** page when Next is clicked.

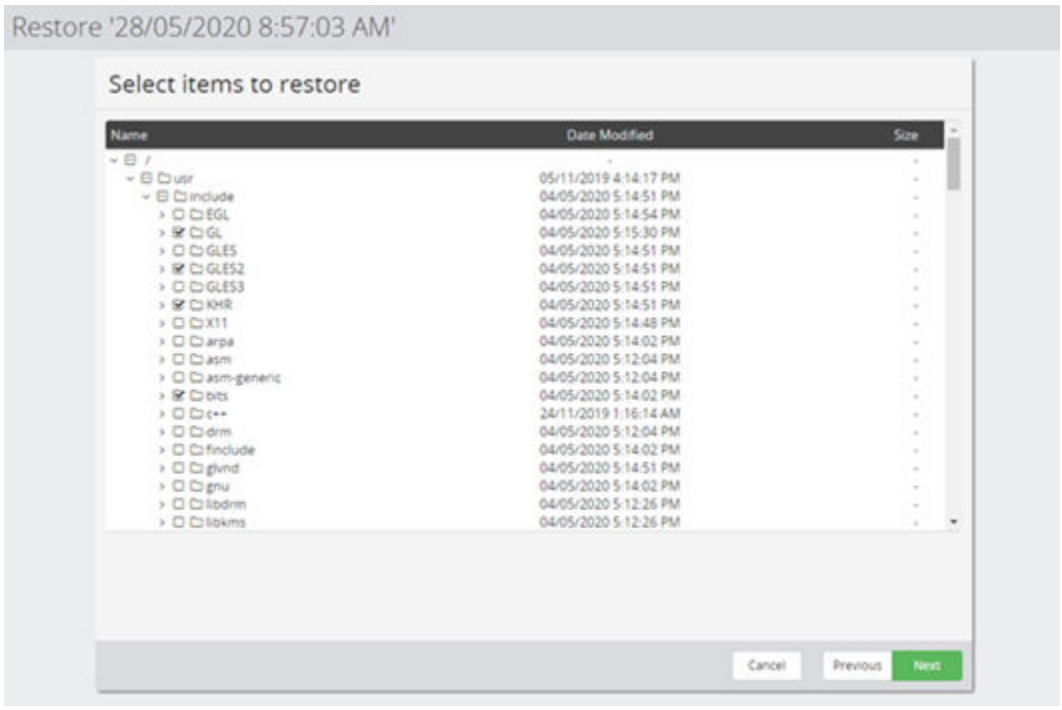


Figure 548 Restore S3 File System Wizard - Select items to restore

Control	Description
File System Contents	Select the directories and/or files you want to restore by clicking the checkbox to the left of its name (a tick mark is displayed next to every file that will be restored). To expand/collapse a folder, click the plus/minus symbol (+/-) to the left of its name.

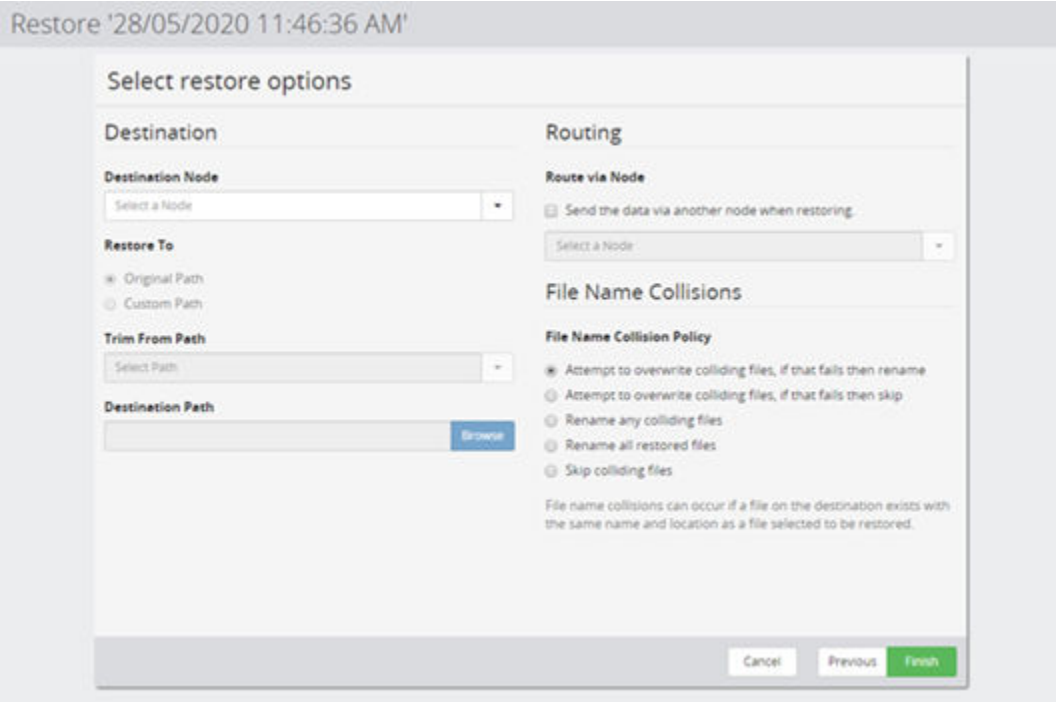


Figure 549 Restore S3 File System Wizard - Select restore options

Control	Description
Destination Node	The target node for the restore.
Original Path	Use the original path names when restoring files.
Custom Path	Use a new path based on the Trim From Paths and Destination Path specification. Refer to the Trim From Path Examples in Destination Template Wizard (on page 379).
Trim From Path	Presents a list of path prefixes that can be removed from the files being restored, whilst ensuring files remain grouped together in their respective folders.
Destination Path	The path specification to be prepended once the paths are trimmed.
Route via Node	Select a node to route via. Primarily used for Internet Connected Nodes where the destination is not directly connected to the repository node, so a proxy must be specified.
Attempt to overwrite colliding files, if that fails then rename	If overwriting fails then the file will be restored by prepending the restore version (that is `(VersionNumber)_filename.ext`) to the file name.

Control	Description
Attempt to overwrite colliding files, if that fails then skip	If overwriting fails then the new file will NOT be restored.
Rename any colliding files	Differentiates all colliding files by prepending the restore version (that is `(VersionNumber)_filename.ext`) to the file name.
Rename all restored files	Differentiates all restored files by prepending the restore version (that is `(VersionNumber)_filename.ext`) to the file name.
Skip colliding files	Do not restore colliding files.

Restore from host based backup Wizard - VMware

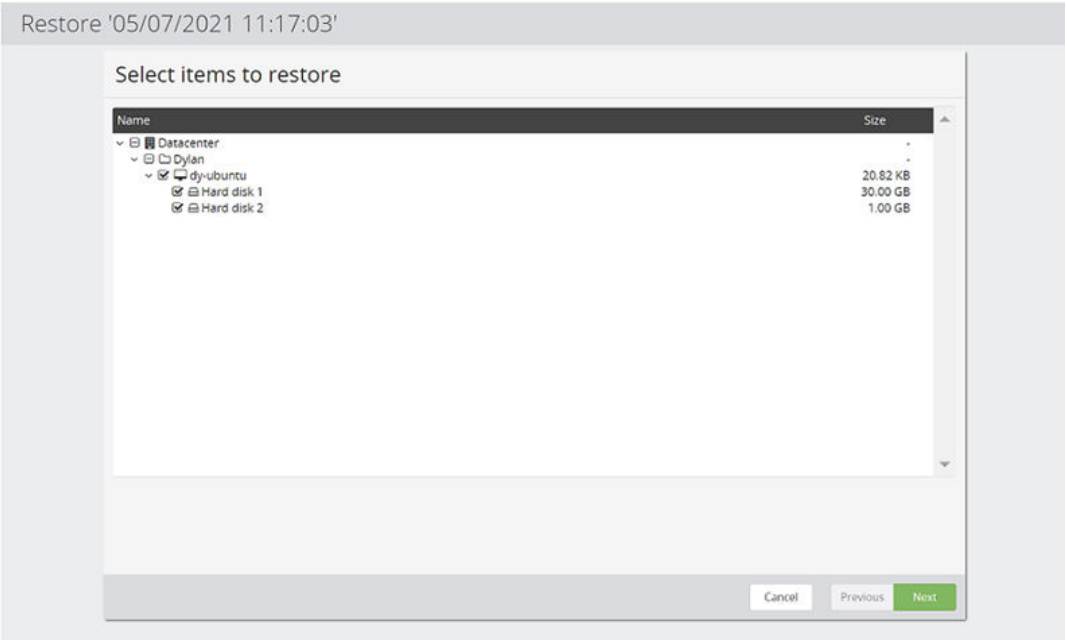


Figure 550 Restore VMware Host Based Backup Wizard - Select files and directories to restore

Control	Description
Virtual Machine List	Select the VMs and Folders to restore. To expand/collapse a folder, click the arrow symbol to the left.

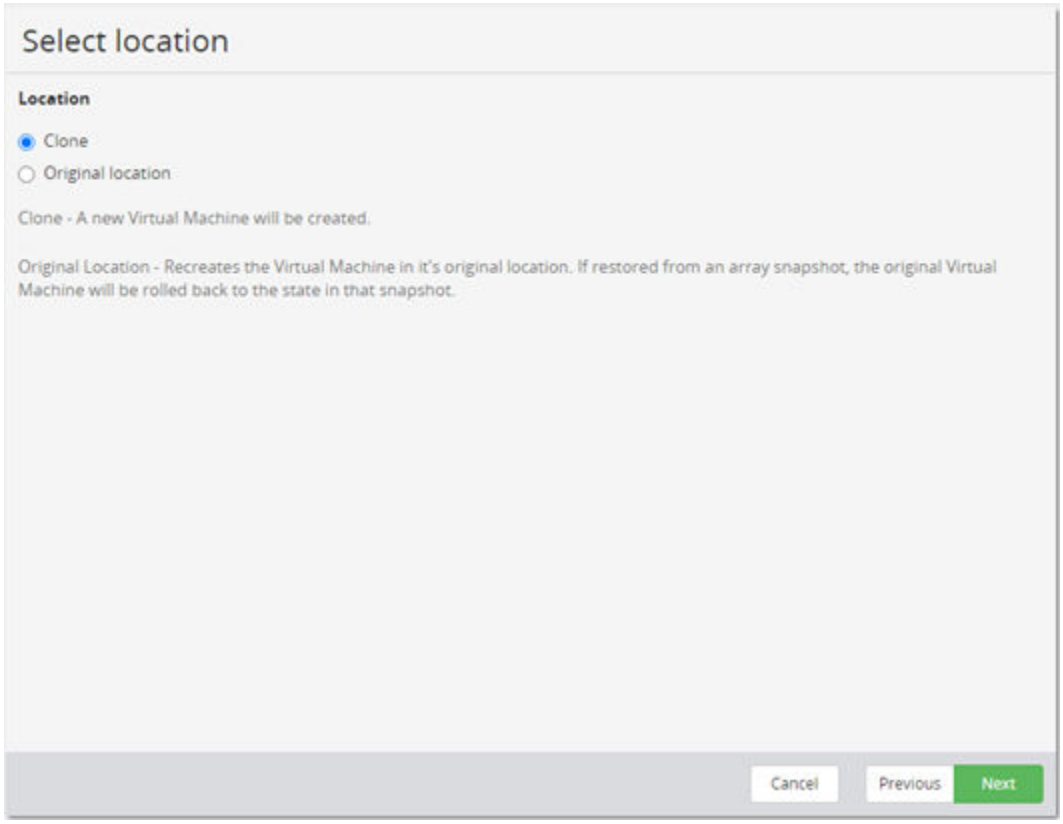


Figure 551 Restore VMware Host Based Backup Wizard - Select location

Control	Description
Original location	If the Virtual Machine to be restored does not exist, it will be recreated. If the Virtual Machine to be restored does exist, it will first be backed up on the datastore and then rolled back to the state of the snapshot. **Note:** If you want to replace the existing VM with the restored one, then delete it before restoring
Clone	The backup will be restored as a clone at the specified location. The wizard displays the **Set clone prefix and destination** page when Next is clicked. **Note:** A restored Virtual Machine will not overwrite any existing machines with the same name in the same location. If a VM of the same name exists at the restore location then the restore job will fail and log and error to that effect.

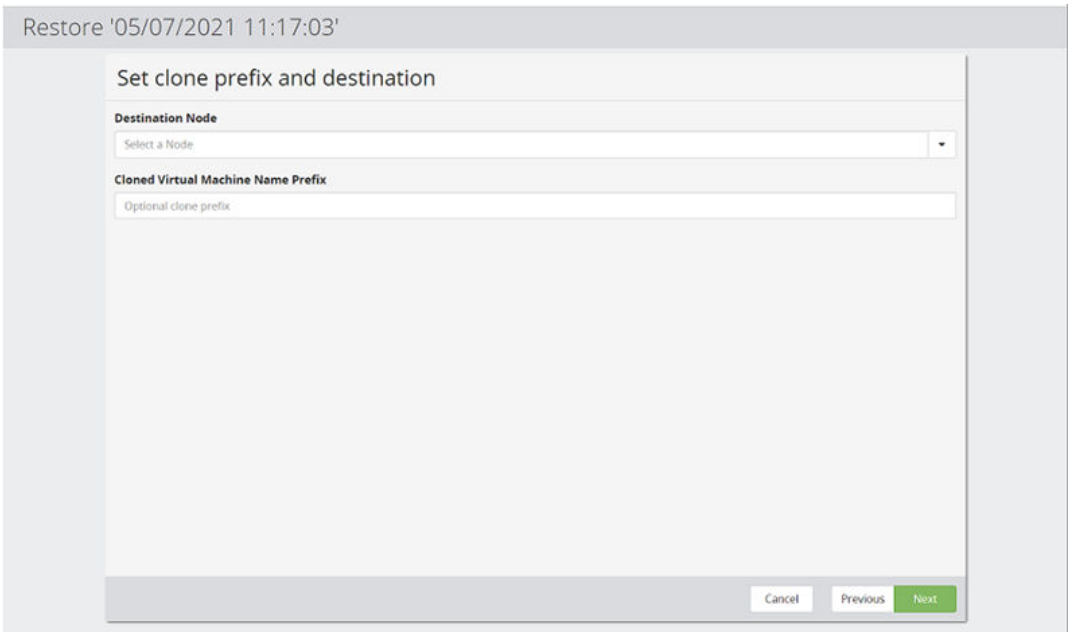


Figure 552 Restore VMware Host Based Backup Wizard - Set clone prefix and destination

Control	Description
Destination Node	Specifies which vCenter node the VMs are to be restored to.
Cloned Virtual Machine Name Prefix	Optional: Clones the original VMs with the prefix applied to the name **if** specified.

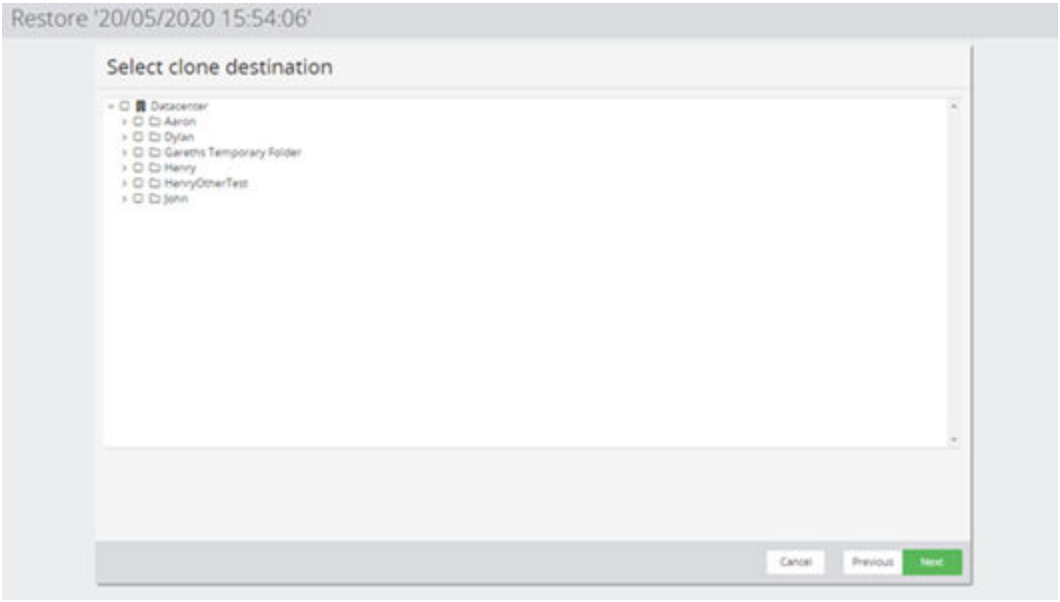


Figure 553 Restore VMware Host Based Backup Wizard - Select clone destination

Control	Description
Datacenters/ Folders	Select the datacenter or folder where the clones are to be located.

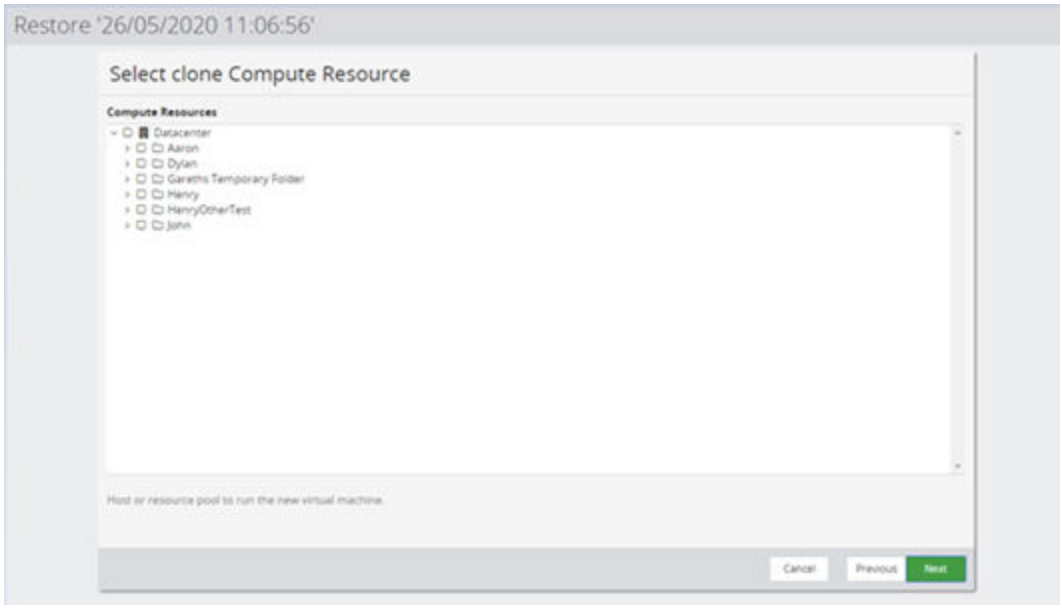


Figure 554 Restore VMware Host Based Backup Wizard - Select clone Compute Resource

Control	Description
Compute Resources	Select the host, vApp or resource pool where the clones are to be located.

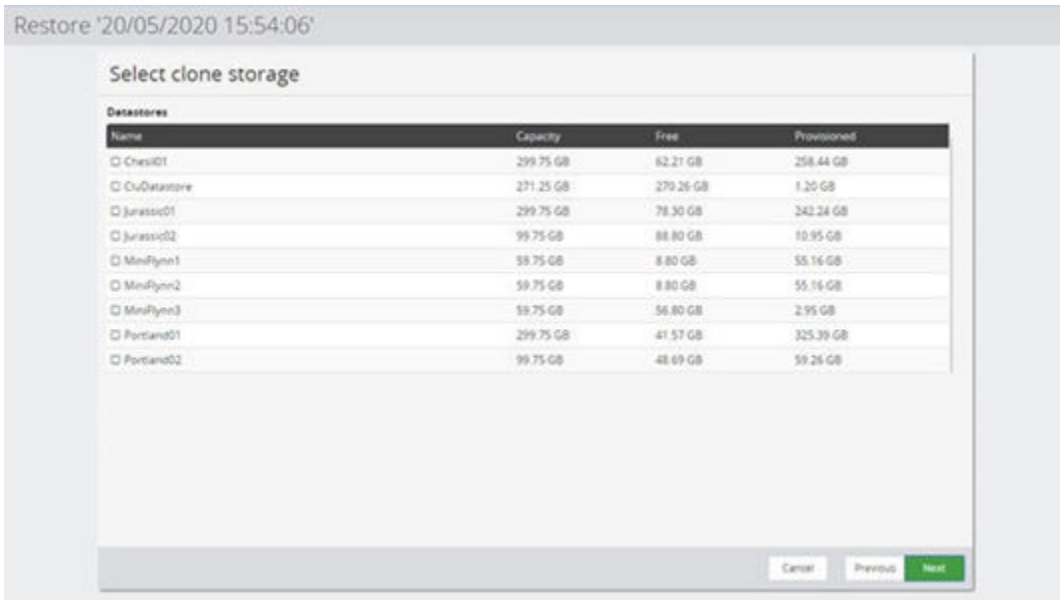


Figure 555 Restore VMware Host Based Backup Wizard - Select clone storage

Control	Description
Datastore	Select the datastore where the clones are to be located.

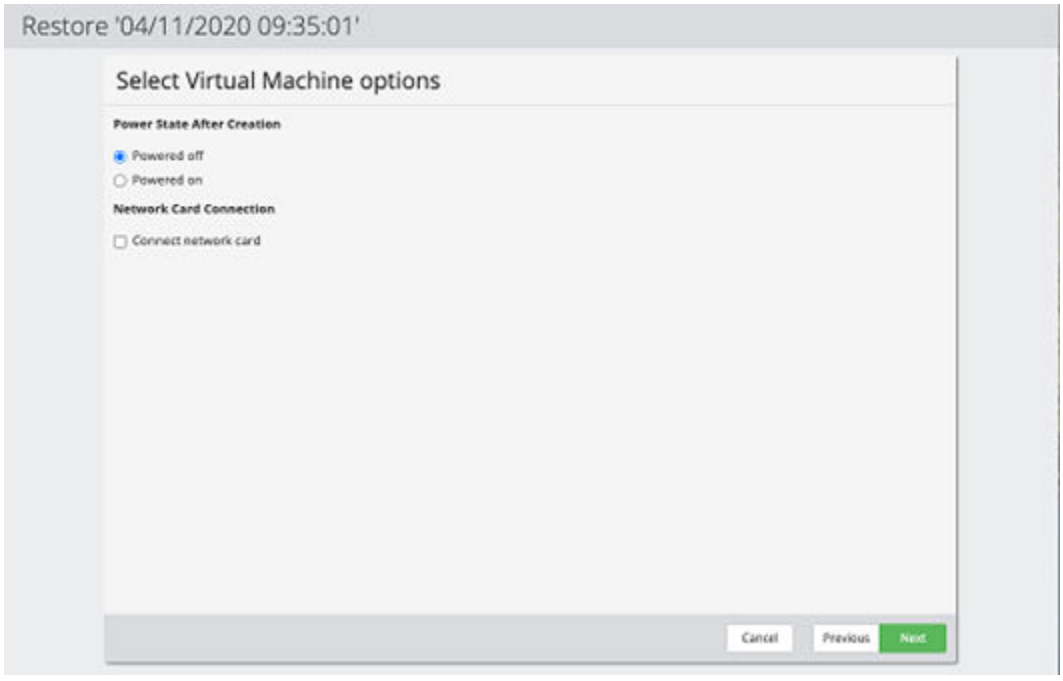


Figure 556 Restore VMware Host Based Backup Wizard - Select Virtual Machine options

Control	Description
Powered off	Select this option to leave the restored VM(s) in the powered off state after they are restored.
Powered on	Select this option to place the restored VM(s) in the powered on state after they are restored.
Network Card State	Select this option to connect the restored VM network card on the VM(s) after they are restored.

Repository Snapshot Details (Restore) - File System

This page provides details of a File System Snapshot within a Repository and contains an inventory of all the files within that snapshot.

Note: The Storage UI contains a similar page (Repository Snapshot Details (Storage) - File System (on page 917)) with different data and more options.

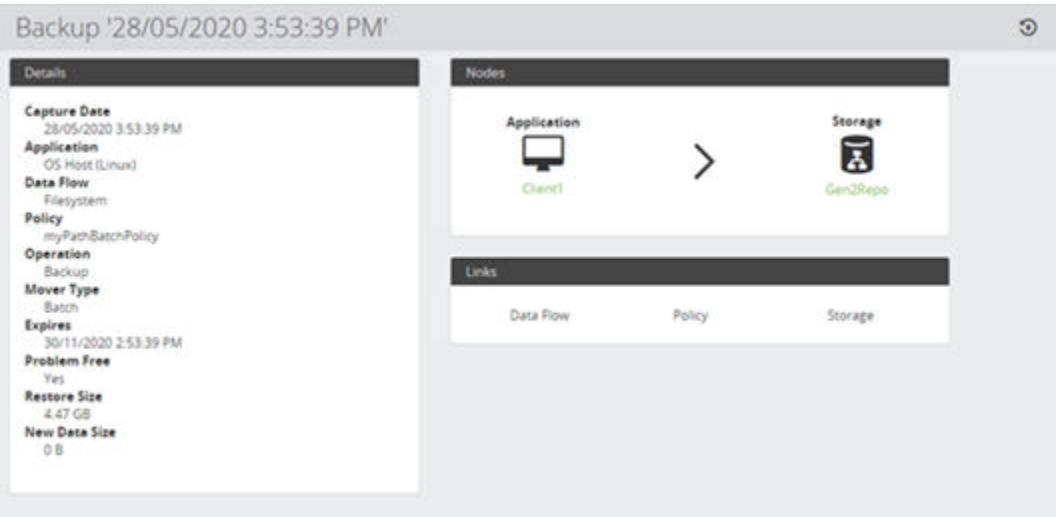


Figure 557 Repository Snapshot Details - File System

Control	Description
Restore	Restores the snapshot. Opens the Restore Repository Snapshot Wizard - File System (on page 826) to guide you through the process.
Details	Provides summary information about the backup.
Nodes	Provides summary information about the source and destination node for the backup, with links to the relevant node screen.

Control	Description
Links	Provides links to the Data Flow, Policy and Storage screen relevant to this backup.

Repository Snapshot Details (Restore) – Ops Center

This page provides details of an Ops Center Snapshot within a Repository and contains an inventory of all the files within that snapshot.

Note: The Storage UI contains a similar page (Repository Snapshot Details (Storage) - File System (on page 917)) with different data and more options.

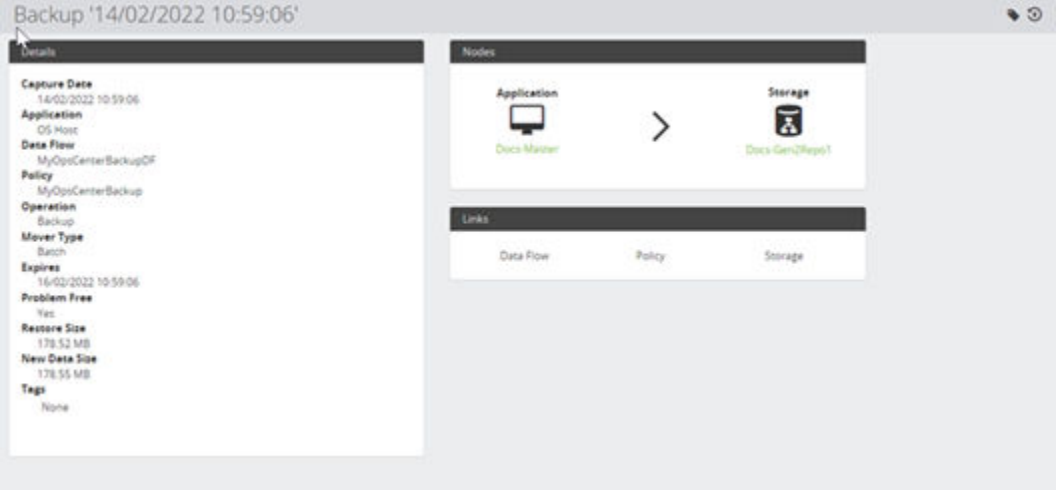


Figure 558 Repository Snapshot Details – Ops Center

Control	Description
Restore	Restores the snapshot. Opens the Restore Repository Snapshot Wizard – Ops Center (on page 830) to guide you through the process.
Details	Provides summary information about the backup.
Nodes	Provides summary information about the source and destination node for the backup, with links to the relevant node screen.
Links	Provides links to the Data Flow, Policy and Storage screen relevant to this backup.

Repository Snapshot Details (Restore) - VMware

This page provides details of a VMware Snapshot within a Repository and contains an inventory of all the VMware files within that snapshot.

Note: The Storage UI contains a similar page (Repository Snapshot Details (Storage) - VMware (on page 919)) with different data and more options.

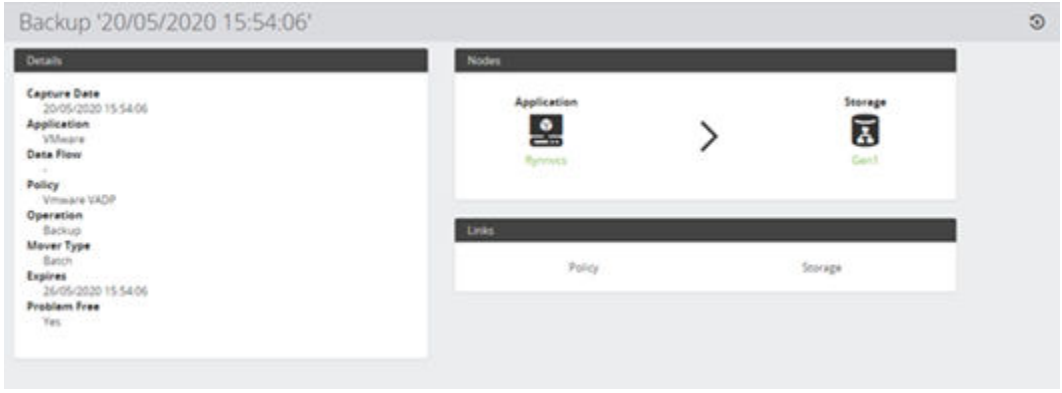


Figure 559 Repository Snapshot Details - VMware

Control	Description
Restore	Restores the snapshot. Opens the Restore from host based backup Wizard - VMware (on page 840) to guide you through the process.
Details	Provides summary information about the backup.
Nodes	Provides summary information about the source and destination node for the backup, with links to the relevant node screen.
Links	Provides links to the Data Flow, Policy and Storage screen relevant to this backup.

HCP Snapshot Details (Restore) - File System

This page provides details of a File System Snapshot within an HCP store.

Note: The Storage UI contains a similar page (Cloud Snapshot Details (Storage) - File System (on page 931)) with different data and more

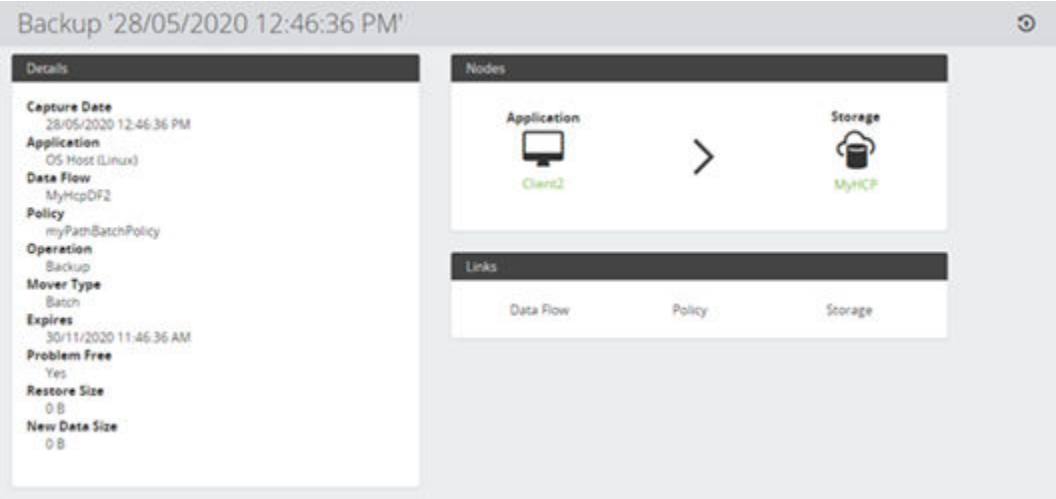


Figure 560 HCP Snapshot Details - File System

Control	Description
Restore	Restores the snapshot. Opens the Restore HCP Snapshot Wizard - File System (on page 832) to guide you through the process.
Details	Provides summary information about the backup.
Nodes	Providessummary information about the source and destination node for the backup, with links to the relevant node screen.
Links	Provides links to the Data Flow, Policy and Storage screen relevant to this backup.

Amazon S3 / HCP / HCP Cloud Scale Snapshot Details (Restore) - File System

This page provides details of a File System Snapshot within an S3 store.

Note: The Storage UI contains a similar page (Cloud Snapshot Details (Storage) - File System (on page 931)) with different data and more.

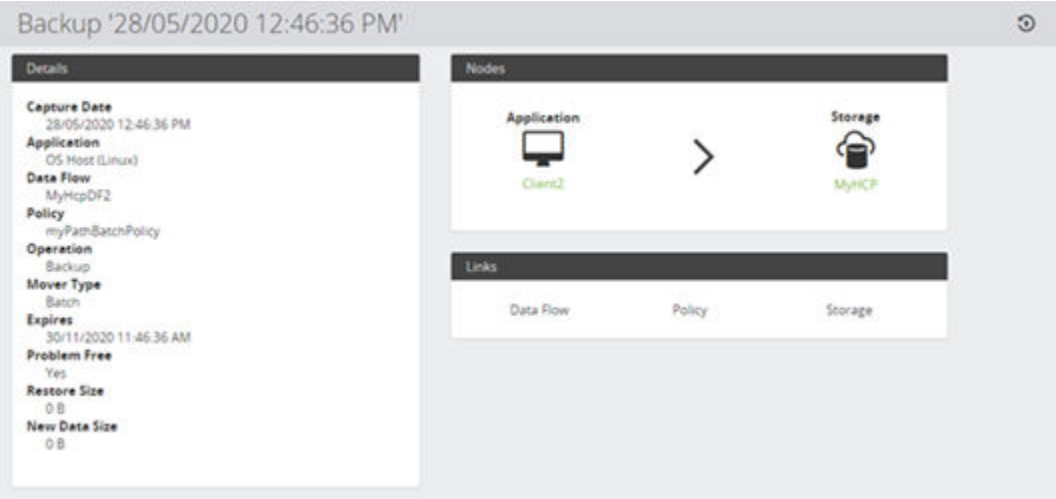


Figure 561 S3 Snapshot Details - File System

Control	Description
Restore	Restores the snapshot. Opens the Restore Amazon S3 / HCP Cloud Scale Snapshot Wizard - File System (on page 836) to guide you through the process.
Details	Provides summary information about the backup.
Nodes	Provides summary information about the source and destination node for the backup, with links to the relevant node screen.
Links	Provides links to the Data Flow, Policy and Storage screen relevant to this backup.

Schedules UI Reference

This section describes the Schedules UI, accessed via the Navigation Sidebar (on page 291).

For further information, refer to:

- Schedule Concepts (on page 101)
- Schedule Tasks (on page 279)

Schedules Inventory

This inventory lists all defined Schedules. Schedules are created to specify when backup operations should and should not occur.

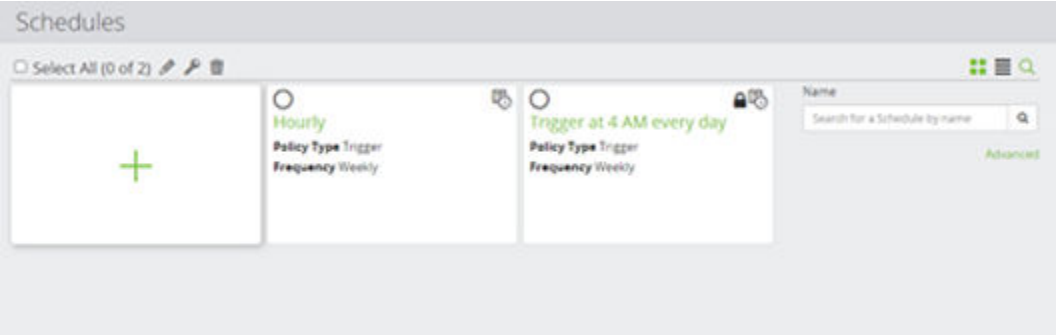


Figure 562 Schedules Inventory

Control	Description
Edit	Edits an existing schedule in the inventory. The Schedule Wizard (on page 850) is launched to enable the schedule's attributes to be changed.
Edit Permissions	Edits an existing schedule's access permissions. The Access Control Permissions Inventory (on page 349) is launched to enable the schedule's access permissions to be changed.
Delete	Enabled only when one or more Schedule is selected. Deletes the selected item(s) from the inventory.
Add	Creates a new Schedule. The Schedule Wizard (on page 850) is launched to guide you through the process.
Schedule(s)	Any number of user defined Schedules can be created. These are displayed in the inventory. Schedules should be defined to determine when operations should occur. The Schedule Details (on page 857) is displayed to enable the schedule to be viewed and edited. The schedule *Trigger at 4 AM every day* is built in and cannot be modified.
Filter on Schedule Name	Filters the displayed results based on the name.

Schedule Wizard

This wizard is launched when a new Schedule is added to the Schedule Inventory. It shows Schedule elements and allows them to be created and edited.

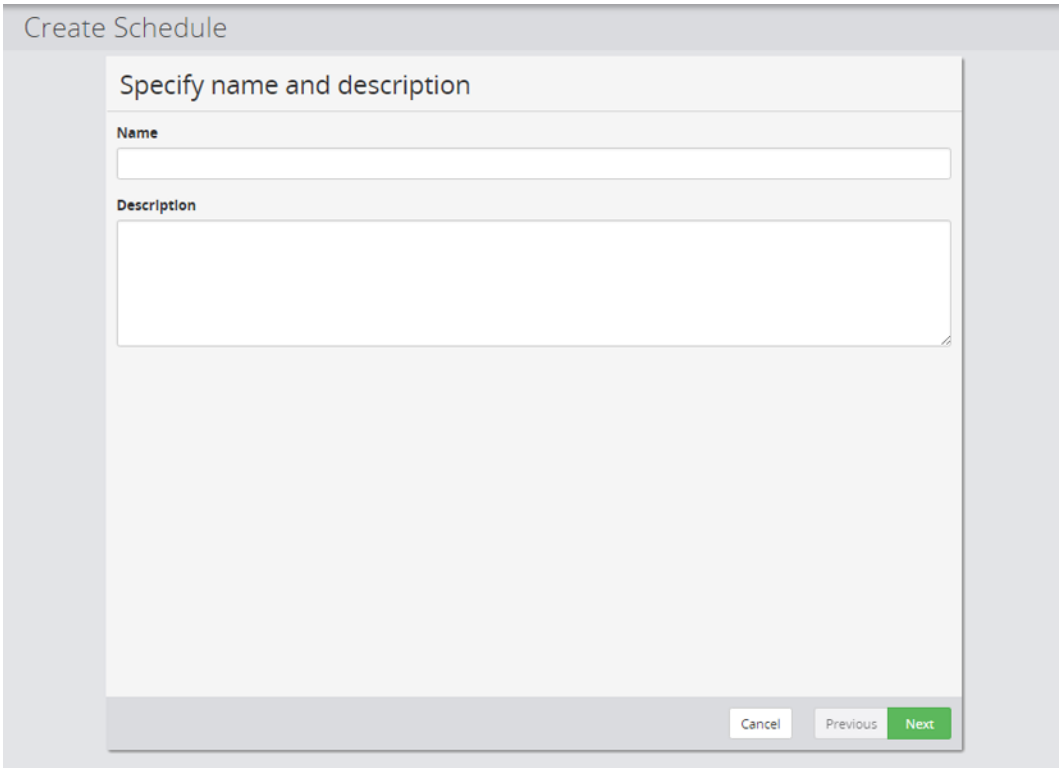


Figure 563 Schedule Wizard - Specify name and description

Control	Description
Name	Enter an identifying name for the schedule.
Description	Optional. Enter a short description of the schedule. This description can be viewed when selecting schedules for use in policies etc.

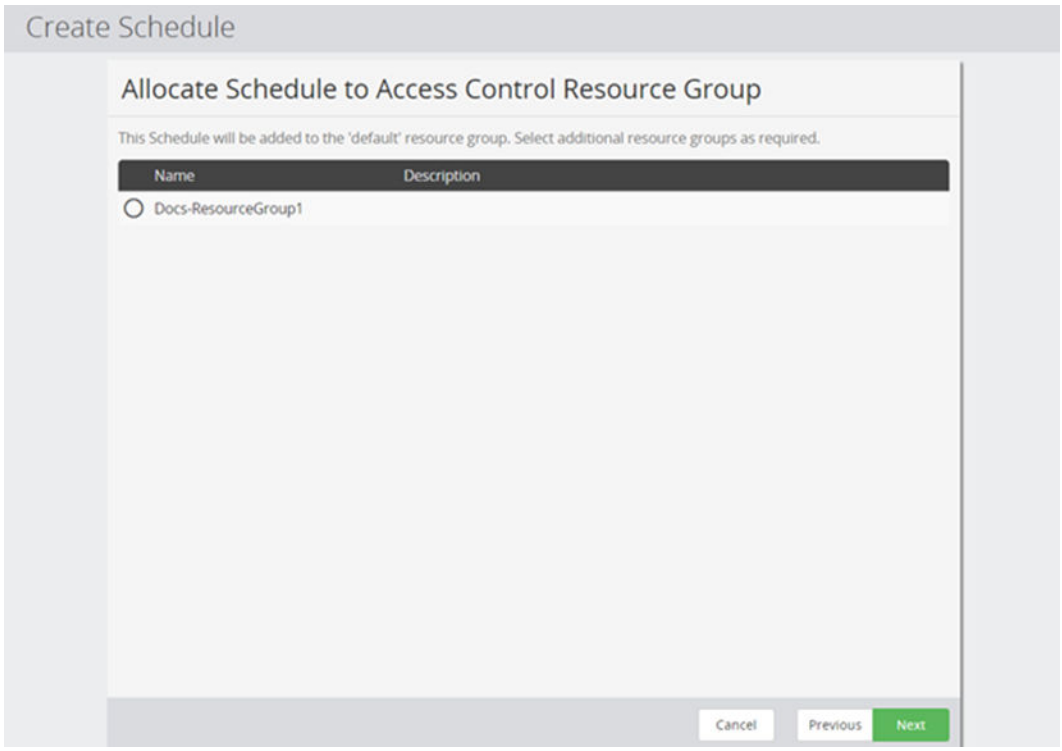


Figure 564 Schedule Wizard - Allocate Schedules to Resource Group

Control	Description
Resource Groups	It allows the user to view the access permissions for those items granted to specific users and groups. **Note:** A single Schedule can be assigned to multiple resource groups.

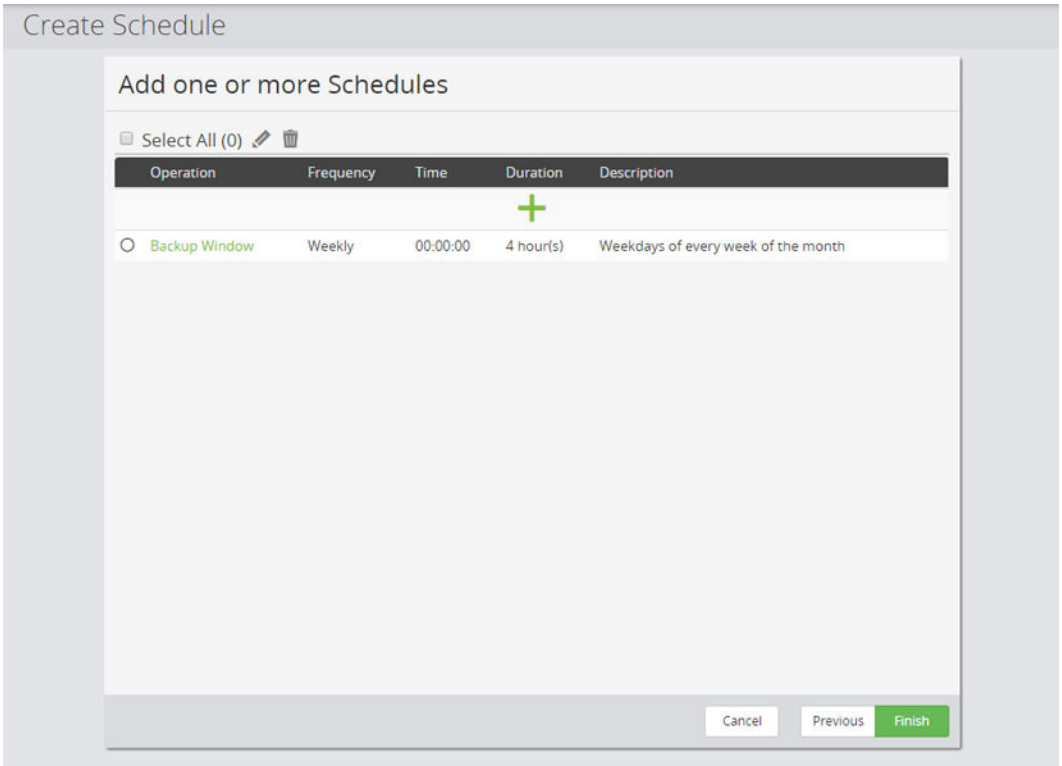


Figure 565 Schedule Wizard - Add one or more Schedule Items

Control	Description
Edit	Edits an existing schedule item in the list. The Schedule Item Wizard (on page 853) is launched to enable the schedule item's attributes to be changed.
Delete	Enabled only when one or more Schedule items are selected. Deletes the selected Schedule item(s) from the list.
Add	Creates a new Schedule item. A new schedule does not initially have any items so they must be created. The Schedule Item Wizard (on page 853) is launched to guide you through the process.
Schedule Element(s)	Any number of user defined Schedules items can be created. These are displayed in the table. Schedule items are combined to generate the required behaviour. The Schedule Item Wizard (on page 853) is displayed to enable the schedule item to be viewed and edited.

Schedule Item Wizard

This wizard is launched when a new Schedule item is added to a Schedule. It shows the Schedule item options and allows them to be configured.

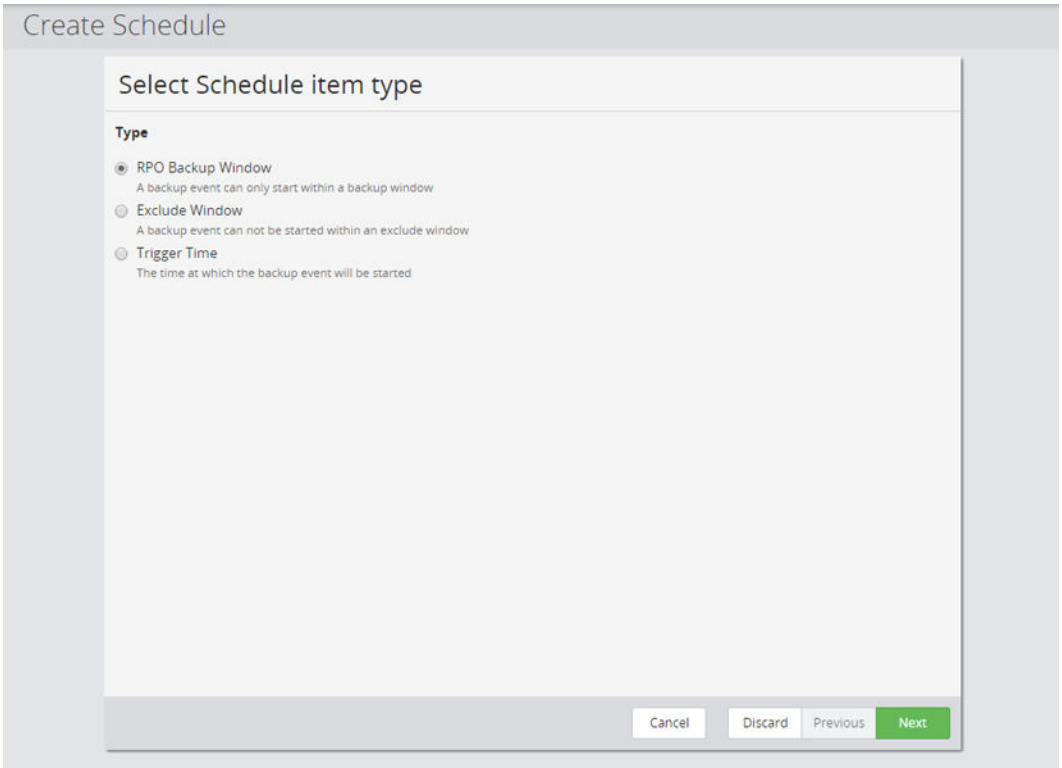


Figure 566 Schedule Item Wizard - Item Type

Table 50

Control	Description
RPO Backup Window	A backup event can only start during a backup window. By default, an implicit backup window is used that operates all day, every day. If a custom backup window element is created, then the default backup window is ignored.
Exclude Window	A backup event cannot start during an exclusion window.
Trigger Time	Similar to a backup window, except that it mandates that a backup event must occur at the specified time or date, unless it occurs during an exclude period.

The schedule item attributes are set by first selecting the corresponding operation type to display the corresponding wizard page:

- Weekly
- Monthly

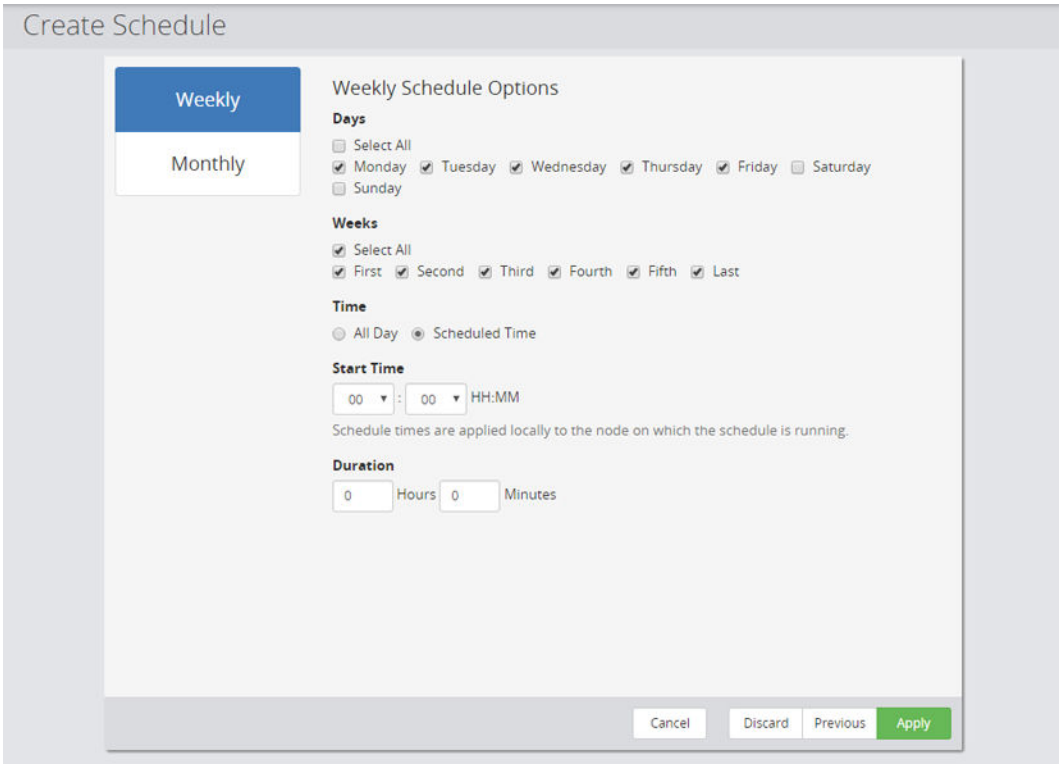


Figure 567 Schedule Item Wizard - Weekly

Weekly – As shown above, Weekly items are used to define schedules based on the days of the week. For example, you can create a schedule element that only occurs on weekends, or that never occurs on weekdays of the last week of the month. At least one day must be specified, while weeks are optional (every week is assumed if none are selected).

Note: The weeks listed do not refer to calendar weeks. The First Week of a month always includes the days from the 1st until the 7th, the Second Week includes the days from the 8th to the 14th, and so on. The Last week always refers to the 7 days prior to (and including) the last day of the month.

Control	Description
Days	Select one or more days of the week.
Weeks	Select one or more weeks of the month.
Time	Can be one of: ▪ All Day - operation can happen at any time of day. This option is disabled for the *Trigger* schedule item type. ▪ Scheduled Time - Set a start time and duration window within which the operation can or cannot occur.
Start Time	Only enabled once Scheduled Time has been selected. Select the time in HH:MM format for the operation to start.

Control	Description
Duration	Only enabled once Scheduled Time has been selected. Enter the duration in HH:MM format for the window in which the operation is allowed to or prevented from starting.

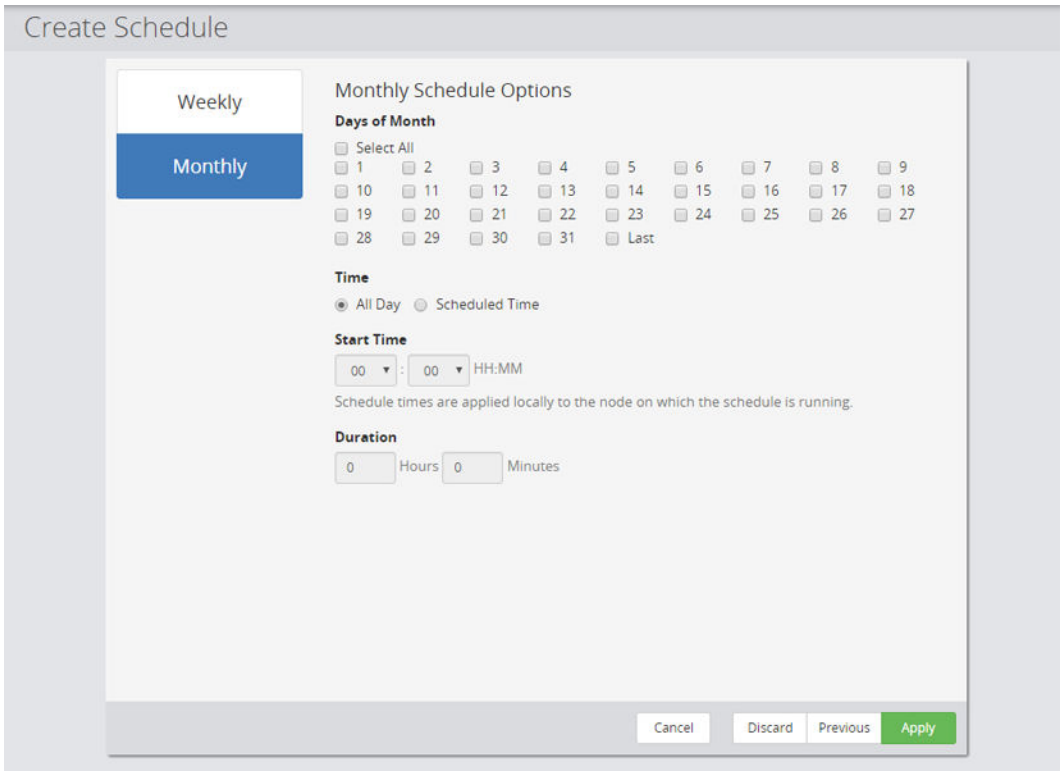


Figure 568 Schedule Item Wizard - Monthly

Monthly – As shown above, Monthly elements are used to define schedules based on days of the month. A monthly schedule element is applied on the same day or days each month. For instance, an element could be defined that triggers a backup on the last day of the month, or an element could be set to prevent backups from occurring on even numbered days. At least one day must be selected in order for a Monthly element to be valid. If the month in question does not contain the selected day (example: 29th Feb in a non-leap year) then the schedule element will not be applied.

Control	Description
Day of Month	Can be one of: ▪ Select All - Selects all date options ▪ 1 .. 31 - Select one or more individual days of the month. ▪ Last - Selects the last day of the month.

Control	Description
Time	Can be one of: ▪ All Day - operation can happen at any time of day. This option is disabled for the *Trigger* schedule item type. ▪ Scheduled Time - Set a start time and duration for the operation to occur.
Start Time	Only enabled once Scheduled Time has been selected. Select the time in HH:MM format for the operation to start.
Duration	Only enabled once Scheduled Time has been selected. Enter the duration in HH:MM format for the window in which the operation is allowed to start.

Schedule Details

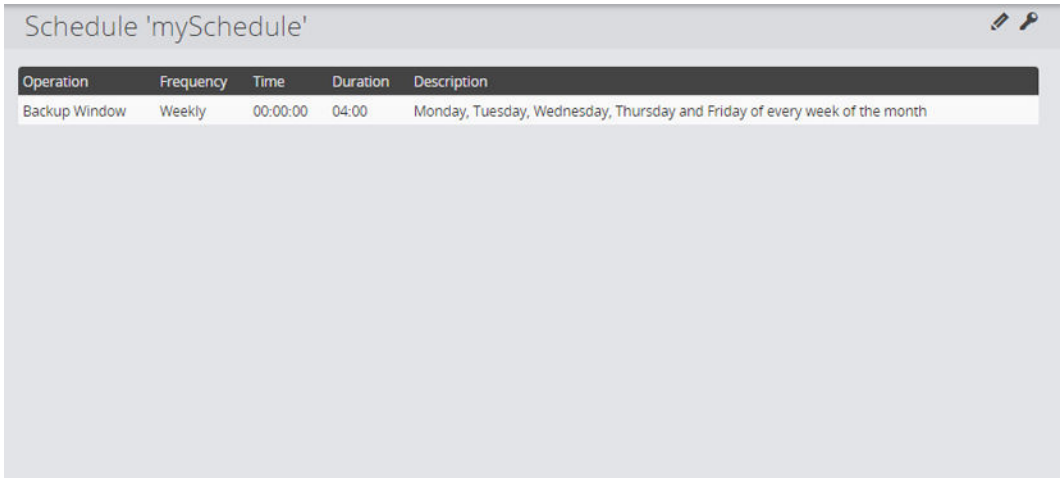


Figure 569 Schedule Details

Control	Description
Edit	Launches the Schedule Wizard (on page 850) to enable you to view and edit the schedule.
Permissions	Opens the Access Control Permissions Inventory (on page 349) to allow you to view and edit the permissions for this schedule.

Storage UI Reference

This section describes the Storage UI, accessed via the Main Banner (on page 286).

For further information, refer to:

- Storage Concepts (on page 114)
- Storage Tasks (on page 280)

Storage Inventory

This inventory contains storage information relating to all available storage nodes:

- Generation 1 Repository
- Generation 2 Repository
- Hitachi Block physical and logical storage devices
- HCP Storage Node
- Amazon S3 Storage Node
- HCP cloud scale storage node

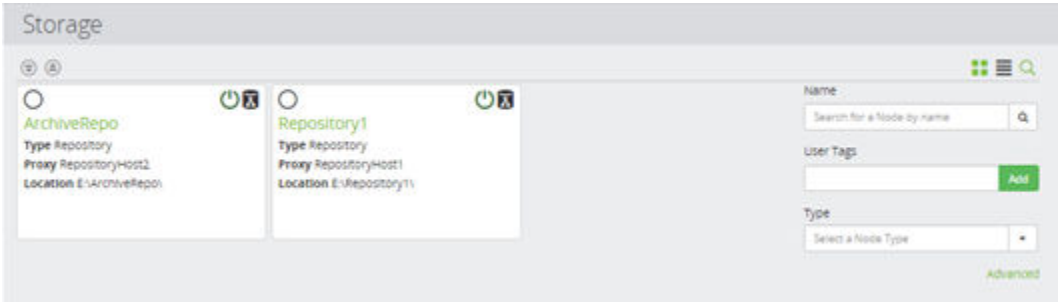


Figure 570 Storage Inventory

Control	Description
Storage Node(s)	Opens one of the following pages depending on the storage type: ▪ Hitachi Block Device Details (on page 859) ▪ Generation 1 Repository Details (on page 921) ▪ Generation 2 Repository Details (on page 922) ▪ Cloud Storage Details (on page 928)
Filter on Name	Filters the displayed results based on Node Name.
Filter on User Tags	Filters the displayed results based on Tags.

Control	Description
Filter on Type	Filters the displayed results based on Node Type: ▪ Amazon S3 ▪ HCP cloud scale ▪ Hitachi Block Device ▪ Hitachi Content Platform Generation 2 ▪ Hitachi Logical Block Device ▪ Hitachi Block Device ▪ Hitachi Software Block Device ▪ Repository Generation 1 ▪ Repository Generation 2

Hitachi Block Device Details

This page displays the details of a Block Storage device and enables monitoring of its activities, state, and the snapshots and replications stored within it.



Figure 571 Hitachi Block Device Details

Control	Description
Refresh	Refreshes the Host Groups, Journals, Pools and Logical Devices data. **Note:** A refresh must be performed to update information held on the storage array.

Control	Description
 Advanced Settings	Opens the Hitachi Block Device Advanced Settings Dialog (on page 893) to allow the advanced settings of the storage system to be modified.
Logs	Indicates the number of unacknowledged error logs in red. If there are no errors then a green tick is displayed instead. Numbers at the bottom show the unacknowledged errors and warnings. Opens the Logs Inventory (on page 509).
Jobs	Indicates the number of failed jobs in red. If there are no failed jobs then a green tick is displayed instead. Numbers at the bottom show active and paused jobs. Opens the Jobs Inventory (on page 493).
Details	Indicates the following: ▪ Status - online/offline ▪ Proxy - the node that acts as the ISM for the Block device. ▪ User Tags – Tags associated with the Block device.
 View Node Details	Takes you to the equivalent Node Management Node Details (on page 650).
Configuration	Shows the configuration of the Hitachi Block Device.
Object Cache	Shows last updated date and time for Pools, Journals, Resource Groups, Quora, Host Groups, Remote Path Groups, and Logical Devices of Hitachi Block Device. For details on the cache refreshing see Hitachi Block Device Scheduled Cache Refresh (on page 894)  **Note:** This screen in the UI does not automatically refresh. While the values are refreshed on a schedule it is necessary to reload the page in the browser to get an up to date view of the values.
 Refresh	Duplicates the Refresh button at the top of the page.
 Host Groups	Indicates the number of Host Groups on the Device visible to the CCI user identified when the Block storage node was configured. If Host Groups exist in a resource group to which the user does not have access then they will not be displayed here. Click on the icon to open the Hitachi Block Host Groups Inventory (on page 861).
 Journals	Indicates the number of Journals on the Device visible to the CCI user identified when the Block storage node was configured. Click on the icon to open the Hitachi Block Journals Inventory (on page 863)

Control	Description
 Pools	Indicates the number of Pools (Total, Snapshot and Dynamic) on the Device. Click on the icon to open the Hitachi Block Pools Inventory (on page 866)
 Logical Devices	Indicates the number of Logical Devices on the Device. If Logical Devices exist in a resource group to which the user does not have access then they will not be displayed here. Click on the icon to open the Hitachi Block Logical Devices Inventory (on page 868)
 Quorums	Indicates the number of Quorums on the device. Click on the icon to open the Hitachi Block Quorums Inventory (on page 895).
 Remote Path Groups	Indicates the number of Remote Path Groups on the device. Click on the icon to open the Hitachi Block Remote Path Groups Inventory (on page 892)
 Snapshots	Indicates the number of Snapshots on the Device that have been created by Protector. Click on the icon to open the Hitachi Block Snapshots Inventory (on page 869)
 Replications and Clones	Indicates the number of Replications on the Device that have been created or adopted by Protector. Click on the icon to open the Hitachi Block Replications Inventory (on page 878)
 Resource Groups	Indicates the number of Resource Groups on the device. Click on the icon to open the Hitachi Block Resource Groups Inventory (on page 896).
 NVMe Subsystems	Indicates the number of NVMe Subsystems on the device. Click on the icon to open the Hitachi Block NVMe Subsystem Details (on page 898)

Hitachi Block Host Groups Inventory

This inventory lists all Host Groups defined on a Hitachi Block storage device.

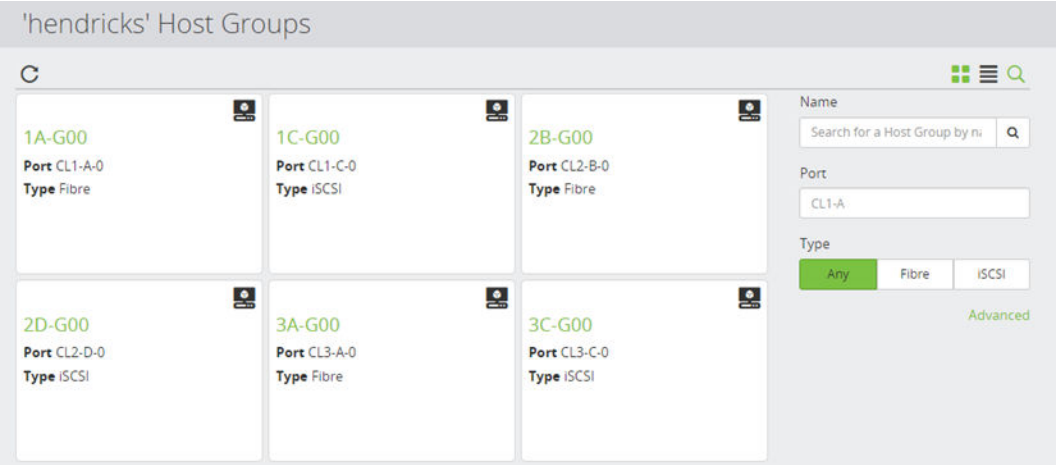


Figure 572 Host Groups Inventory

Control	Description
Refresh	Refreshes the displayed inventory. **Caution:** Host Group details are only refreshed when refreshed manually. To ensure the latest information is displayed, this page must be manually refreshed periodically.
Host Group(s)	Click on a Host Group to open the Hitachi Block Host Group Details (on page 862).
Filter on Name	Filters the displayed results based on the Host Group Name.
Filter on Port	Filters the displayed results based on the Port.

Hitachi Block Host Group Details

This page shows the details of a Host Group on a Hitachi Block storage device.

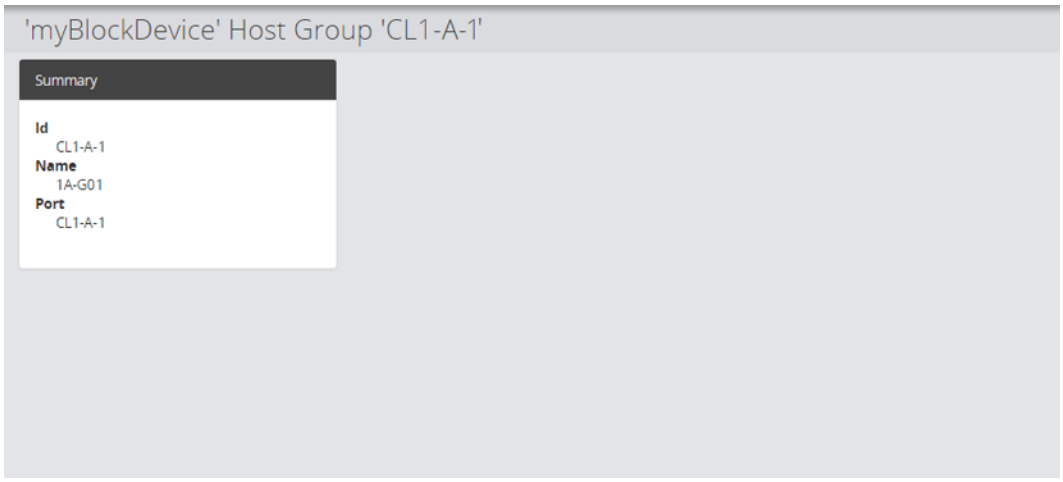


Figure 573 Host Group Details

Control	Description
Summary	Summarises Host Group details.

Hitachi Block Journals Inventory

This inventory lists all Journals on a Hitachi Block storage device.

Note: The Journal inventory is not automatically updated and should be periodically manually refreshed to ensure the latest information is available.

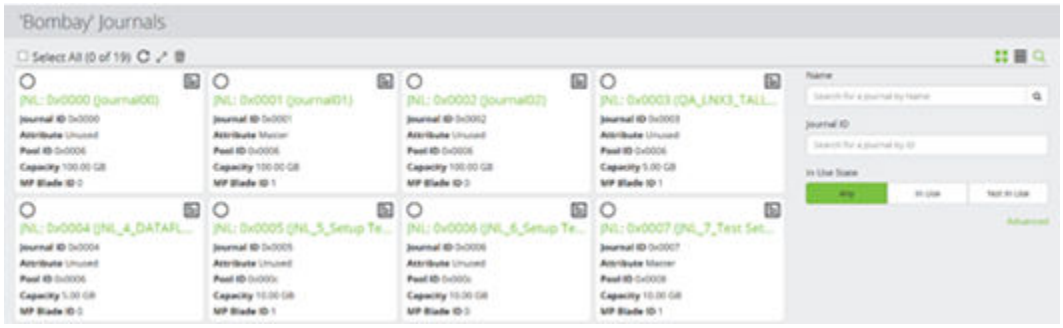


Figure 574 Journals Inventory

Control	Description
Refresh	Refreshes the displayed inventory.
Expand	Opens the Figure 575 Hitachi Block Journal Expansion Dialog (on page 864) to enable the capacity of the selected journal to be increased.

Control	Description
Delete	Deletes the selected journal(s) from the storage system. A confirmation dialog opens to confirm the action.
Journal(s)	Click on a Journal to open the Hitachi Block Journal Details (on page 864).
Filter Journal ID	Filters the displayed results based on the Journal ID.
Filter on Use State	Filters the displayed results based on the use state of the Journal



Figure 575 Hitachi Block Journal Expansion Dialog

Control	Description
New Journal Size	Enter a new size for the journal. The new size must be larger than the current size.

Hitachi Block Journal Details

This page shows the details of a Journal on a Hitachi Block storage device.

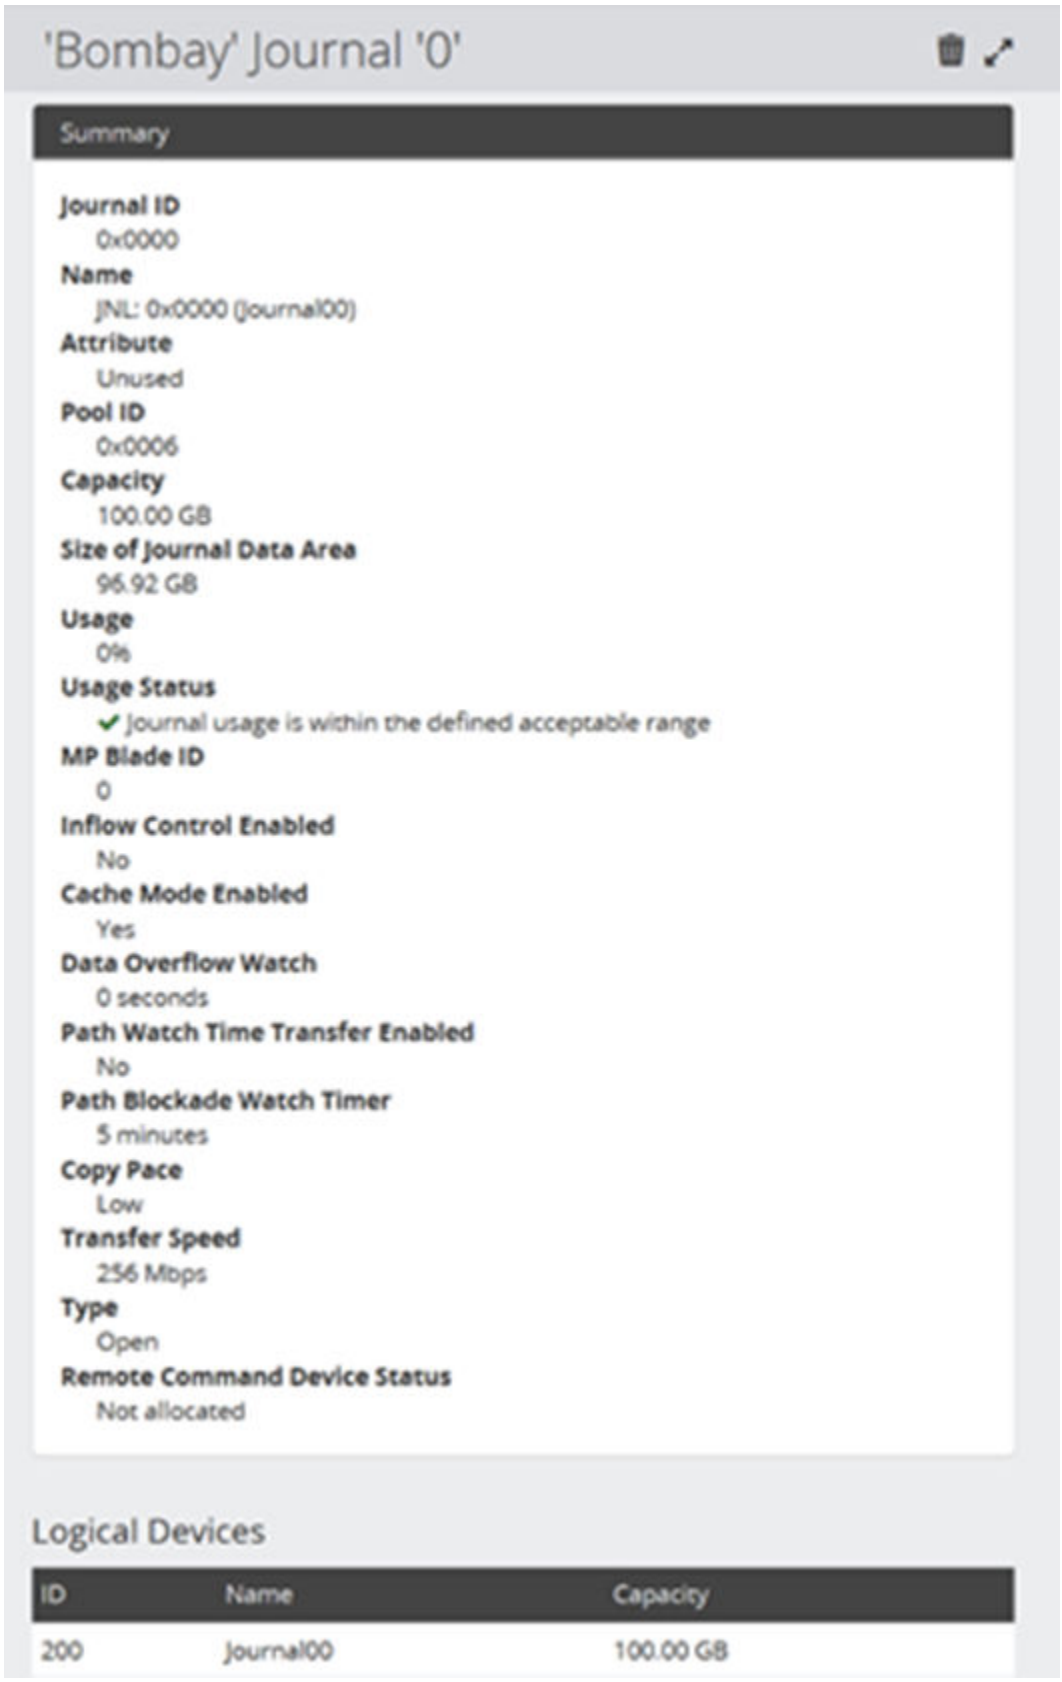


Figure 576 Journal Details

Control	Description
Expand	Opens the Figure 575 Hitachi Block Journal Expansion Dialog (on page 864) to enable the capacity of the selected journal to be increased.
Delete	Deletes the selected journal(s) from the storage system. A confirmation dialog opens to confirm the action.
Summary	Summarises the Journal details.
Logical Devices	Displays information on the Logical Devices for the Journal.

Hitachi Block Pools Inventory

This inventory lists all Pools on a Hitachi Block storage device.

Caution:

Filling a Thin Image pool to capacity will invalidate all snapshot data contained within that pool. All snapshots in the pool will have to be deleted before snapshotting can be resumed.

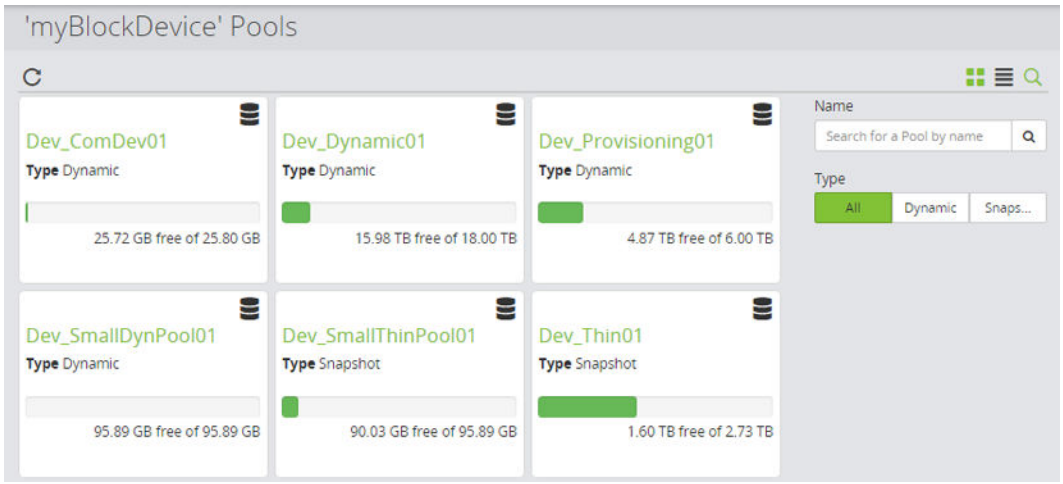


Figure 577 Pools Inventory

Control	Description
Refresh	Refreshes the displayed inventory. **Caution:** Pool usage is only refreshed when refreshed manually. If these values are being used to monitor pool usage they must be manually refreshed periodically.

Control	Description
 Pool Tile(s)	Click on a Pool name to open the Hitachi Block Pool Details (on page 867). **Note:** The pool capacity indicator bar changes colour when thresholds defined on and read from the storage array are crossed: ▪ Green - indicates pool usage is within the *normal* range. ▪ Orange - indicates pool usage is at or above the *warning* level. The *warning* threshold defaults to 70%. ▪ Red - indicates pool usage is at or above the *depletion* level. The *depletion* threshold defaults to 80%. (1) (1) Thin Image pools only have a single threshold defined on the array. Protector transitions straight from the normal state to the depletion state. For dynamic pools containing replication S-VOLs, the replication will transition to an error state when the pool is full.
Filter on Name	Filters the displayed results based on the Pool Name.
Filter on Type	Filters the displayed results based on the Pool Type: ▪ All ▪ Dynamic - Used to provision primary volumes, command devices and replications ▪ Snapshot - Used to provision snapshots

Hitachi Block Pool Details

This page shows the details of a Pool on a Hitachi Block storage device.

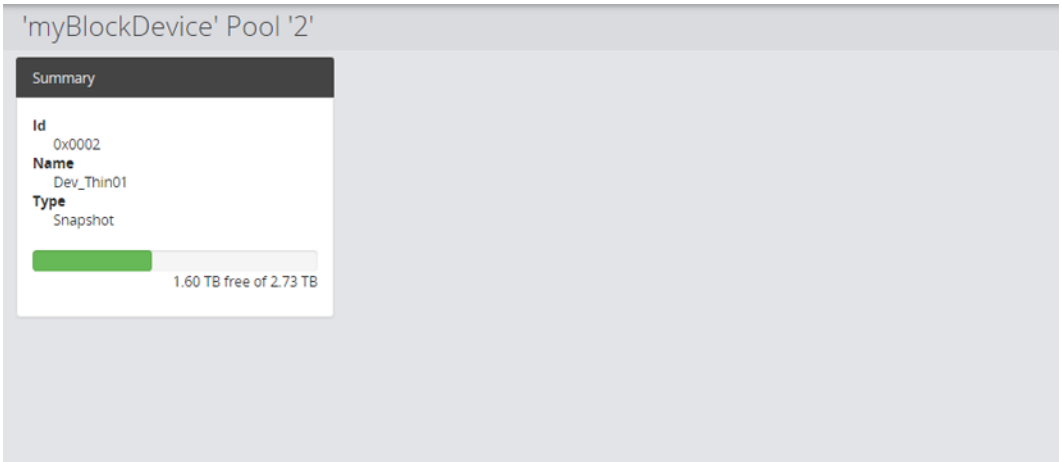


Figure 578 Pool Details

Control	Description
Summary	Summarises the Pool details.

Hitachi Block Logical Devices Inventory

This inventory lists all Logical Devices on a Hitachi Block storage device.

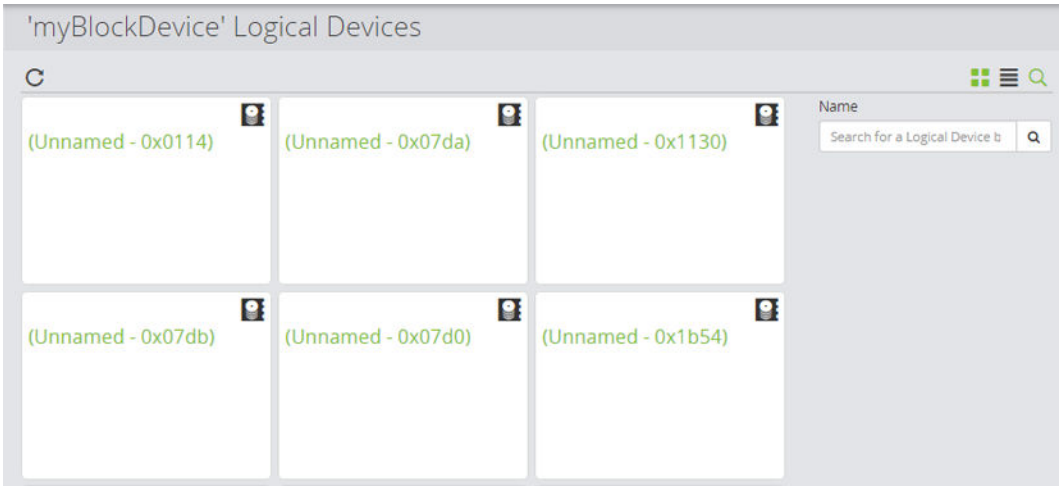


Figure 579 Logical Devices Inventory

Control	Description
Refresh	Refreshes the displayed inventory.

Control	Description
	Caution: Logical Device Inventory is only refreshed when refreshed manually. To ensure the latest information is displayed this page should be manually refreshed periodically.
Logical Device(s)	Click on a Logical Device to open the Hitachi Block Logical Device Details (on page 869).
Filter on Name	Filters the displayed results based on the Logical Device Name.

Hitachi Block Logical Device Details

This page shows the details of a Logical Device on a Hitachi Block storage device.

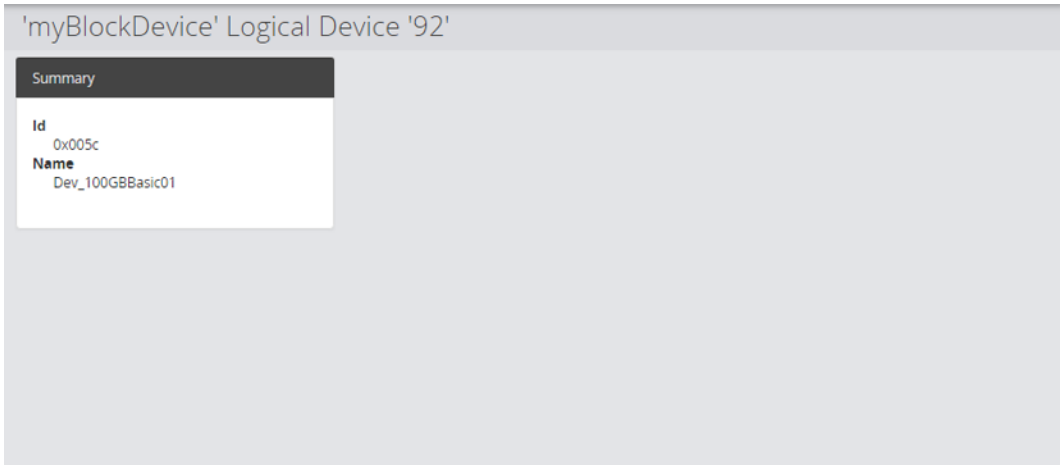


Figure 580 Logical Device Details

Control	Description
Summary	Summarises the Logical Device details.

Hitachi Block Snapshots Inventory

This page provides an inventory of all the Snapshots stored on a Block Storage Device.

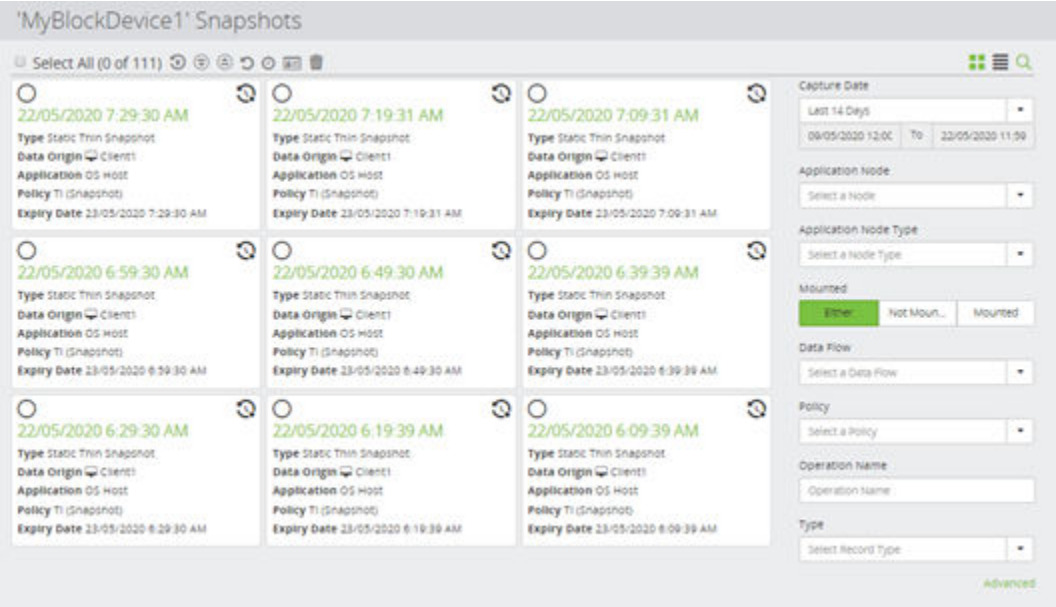


Figure 581 Hitachi Block Snapshots Inventory

Control	Description
Restore	Enabled only if one VMware Block snapshot is enabled. Opens the Hitachi Block VMware Snapshot Restore Wizard (on page 766) to guide you through restoring VM's from the snapshot Note: The restore operation may take several minutes to complete.
Mount	Enabled only if one or more Snapshots are selected. Opens the Hitachi Block Snapshot or Replication Mount Wizard (on page 772) to guide you through mounting the Snapshot. **Note:** The mount operation can take several minutes to complete.
Unmount	Enabled only if one mounted snapshot is selected. Unmounts the selected snapshot. **Note:** The unmount operation can take several minutes to complete.
Revert	Enabled only if one or more Snapshots are selected. Opens the Hitachi Block Revert Wizard (on page 813) to guide you through reverting the Snapshot.
Set Expiry Date	Opens the Hitachi Block Change Snapshot Expiry Date Dialog (on page 875) to enable the expiry date of the selected snapshot(s) to be modified.

Control	Description
 Delete	Enabled only if one or more Snapshots are selected. Deletes the snapshot from the pool. A confirmation dialog opens to confirm the action. If the snapshot record contains fully provisioned LDEVs then the user can view these details prior to deletion.
 Snapshot Date and Time	Click on the Snapshot's Date and Time to open the Hitachi Block Snapshot Details (Storage) (on page 871).
 DRU Protection	Indicates that the snapshot has DRU protection applied. The protection can be removed (subject to lock duration) via the Hitachi Block Record Details (Storage) (on page 875).
Filter on Capture Date	Filters the displayed results based on the Date Range specified in the Date Time Range Picker (on page 352).
Filter on Application Node	Filters the displayed results based on source application node.
Filter on Application Node Type	Filters the displayed results based on source application node type
Filter on Mount status	Filters the displayed results based on mount status of the snapshot
Filter on Dataflow	Filters the displayed results based on Dataflow
Filter on Policy	Filters the displayed results based on the Policy.
Filter on Operation Name	Filters the displayed results based on mount status of the operation name

Hitachi Block Snapshot Details (Storage)

This page shows the details of a Snapshot on a Hitachi Block storage device.

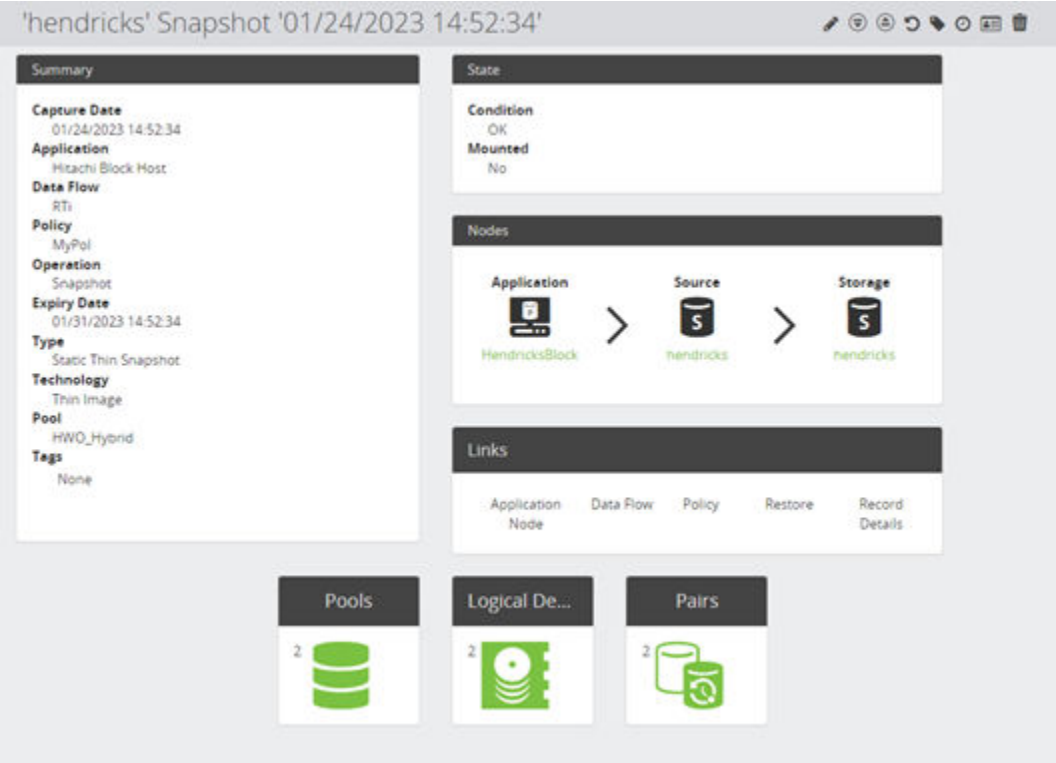


Figure 582 Snapshot Details

Control	Description
Rename Secondary LDEVs	Opens the Hitachi Block Rename Secondary LDEVs Dialog (on page 874) to enable renaming of LDEVs using a fixed string and/or substitution variables.
Mount	Opens the Hitachi Block Snapshot or Replication Mount Wizard (on page 772) to guide you through mounting the snapshot. **Note:** The mount operation can take several minutes to complete.
Unmount	Unmounts the snapshot. The export is deleted. **Note:** The unmount operation can take several minutes to complete.
Revert	Opens the Hitachi Block Revert Wizard (on page 813) to guide you through reverting the snapshot.
Set Expiry Date	Opens the Hitachi Block Change Snapshot Expiry Date Dialog (on page 875) to enable the expiry date of the snapshot to be modified.

Control	Description
 Remove DRU Protection	If the snapshot was created with DRU protection and the protection lock duration has expired, this will remove the protection.
 Delete	Deletes the snapshot from the pool. A confirmation dialog opens to confirm the action.
Summary	Identifies the hardware storage device, application, policy etc. Click on the policy name to open the Policy Details (on page 737).
State	Indicates the status of the snapshot.
Links	Provides links to the following pages that provide additional information about the snapshot: ▪ Application Node - opens the Node Details (on page 650) for the source node from which the snapshot data set originated. ▪ Data Flow - opens the Data Flow Details (on page 491) where the snapshot operation is assigned. ▪ Policy - opens the Policy Details (on page 737) where the snapshot operation is defined. ▪ Replication - displayed only if the snapshot is of a replication's destination. Opens the Hitachi Block Replication Details (Storage) (on page 882) ▪ Record Details - opens the Hitachi Block Record Details (Storage) (on page 875) for the snapshot. ▪ Restore - opens the Hitachi Block Snapshot Details (Restore) (on page 819) for the snapshot.
Nodes	Shows the source and destination nodes involved in the snapshot. You can click on either node to go to their Node Details (on page 650)
 Pools	Opens Hitachi Block Snapshot - Pools (on page 876), listing the pools used in this snapshot.
 Logical Devices	Opens Hitachi Block Snapshot - Logical Devices (on page 878), listing the LDEVs in this snapshot.
 Pairs	Opens Hitachi Block Snapshot - Pairs (on page 877), listing the pairs in this snapshot.

Hitachi Block Rename Secondary LDEVs Dialog

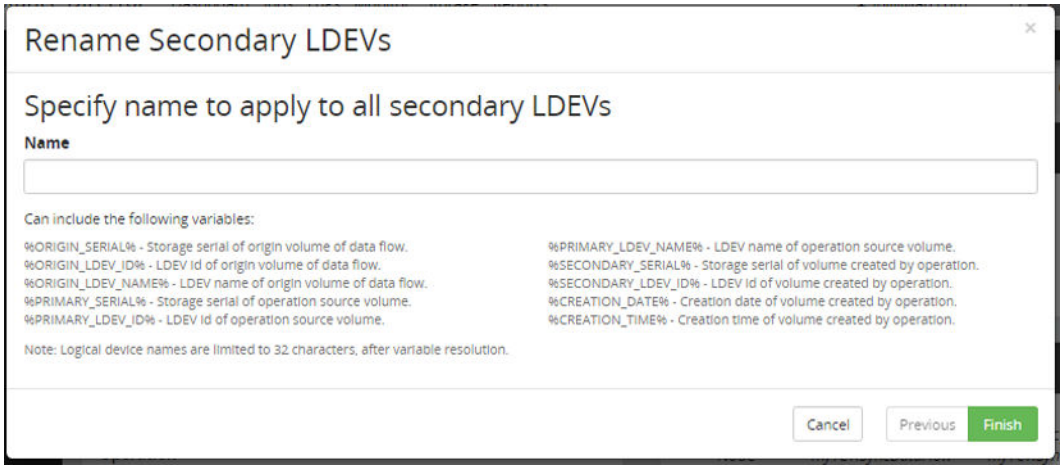


Figure 583 Rename Secondary LDEVs Dialog

Control	Description

Name	Enter the string to be used to rename all secondary LDEVs for the given snapshot or replication. The name can consist of literal strings and/or any of the variables listed below the edit box.
	▪ ORIGIN refers to the left-most item in a cascaded data flow.
	▪ PRIMARY refers to the item immediately to the left in a data flow.
	Tip: Use the cut and paste tools in your browser to insert the required variable (including the surrounding % delimiters) to ensure the correct format and spelling.
	For example:
	`Snap_%ORIGIN_SERIAL%_%ORIGIN_LDEV_ID%`
	might result in each secondary LDEV being renamed:
	`Snap_434525_00:3A:98`
	Note:
	▪ Logical device names are truncated to 32 characters, after variable resolution.
	▪ When renaming intermediate secondaries in a cascaded replication, the names of downstream secondaries will not update until those replications are retriggered.

Hitachi Block Change Snapshot Expiry Date Dialog



Figure 584 Change Snapshot Expiry Date Dialog

Control	Description
Expiry Date	Enter the date and time that the snapshot will expire and be deleted. Opens the Date Time Picker (on page 351).

Hitachi Block Record Details (Storage)

This page shows the details of a record on a Hitachi Block storage device.

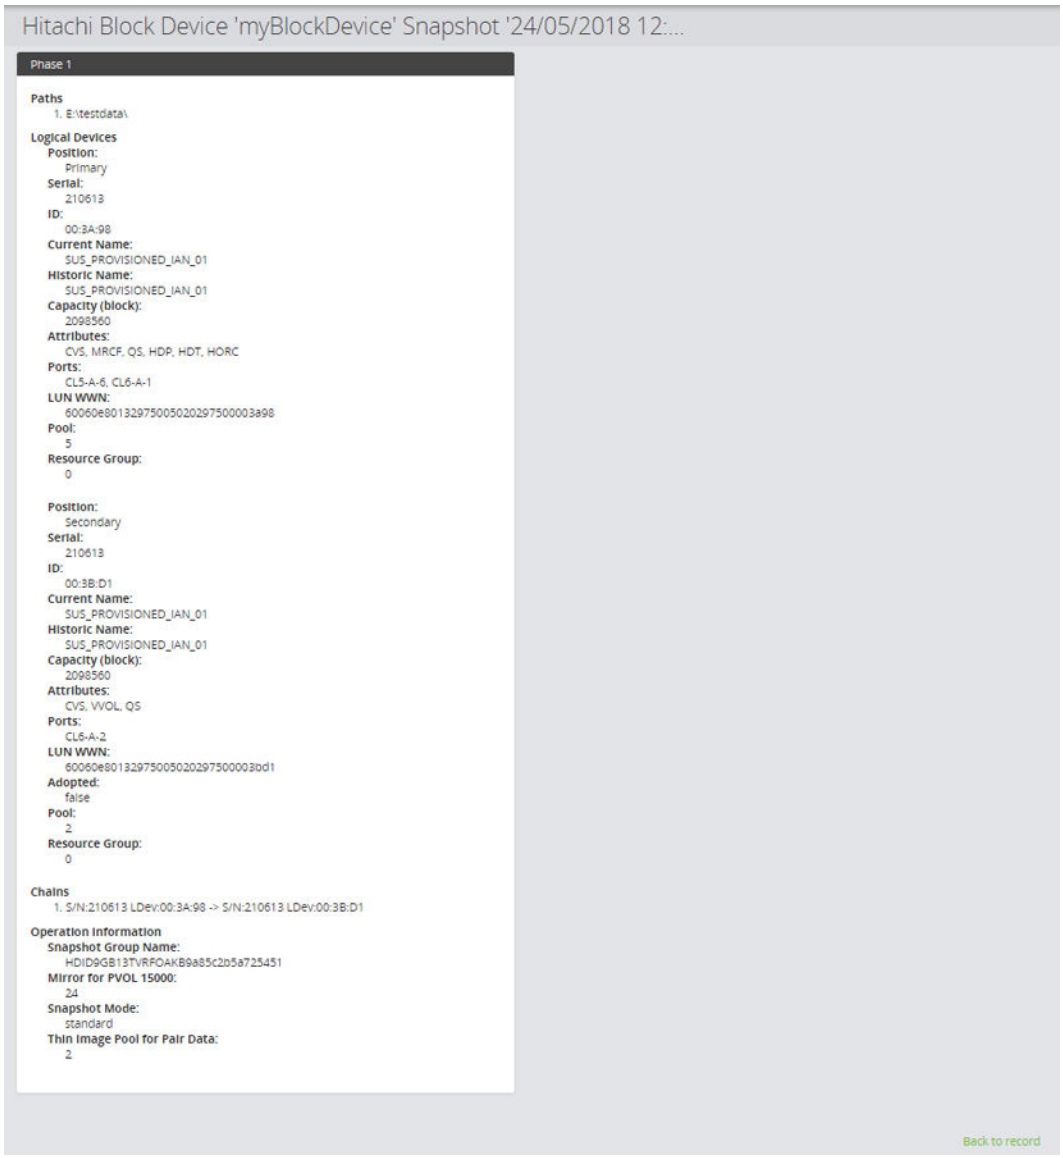


Figure 585 Record Details - Snapshot

Control	Description
Phase 1..n	Identifies the paths, logical devices, chains and operation information for the snapshot. For policies having multiple data classifications or an application classification, the operation is performed in multiple phases. The record for each phase is displayed separately.

Hitachi Block Snapshot - Pools

Lists the pools used in a particular snapshot.



Figure 586 Snapshot Pools

Control	Description
Snapshot Pools list	Each pool is listed, along with its attributes.

Hitachi Block Snapshot - Pairs

Lists the pairs in a particular snapshot. This data is collected from [raidcom get snapshot] – for more information about certain context-sensitive values, please consult the relevant CCI documentation.

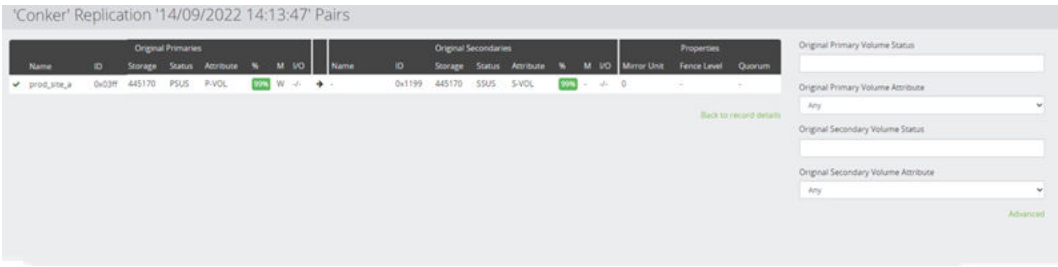


Figure 587 Snapshot Pairs

Control	Description
ID	Displays the ID of the LDEV.
Storage	Displays the serial number of the storage array.
Status	Not populated.
Attribute	Displays the status for this LDEV in this pair (eg COPY, PAIR, PSUS, SSUS). Not populated for Original Secondaries if the snapshot is floating and unmounted.
%	Not populated.
M	Not populated.
IO Mode	Not populated.
Direction (no column header)	Displays an arrow pointing from P-VOL to S-VOL.
Mirror Unit	Displays the mirror unit for the P-VOL of this pair.

Control	Description
Type	Displays the type of this pair (eg [TI]).
Fence Level	Not populated.
Quorum	Not populated.
Snapshot Pairs list	Each pair is listed, identifying the source LDEV's properties on the left and the destination LDEV's properties on the right, separated by an arrow indicating the direction of the snapshot.
Filter by Original Primary Volume Status	Filters the list by local volume status.
Filter by Original Primary Volume Attribute	Filters the list by local volume attribute.
Filter by Original Secondary Volume Status	Filters the list by remote volume status.
Filter by Original Secondary Volume Attribute	Filters the list by remote volume attribute.

Hitachi Block Snapshot - Logical Devices

Lists the LDEVs used in a particular snapshot.



Figure 588 Snapshot Logical Devices

Control	Description
Snapshot LDEVs list	Each LDEV is listed, along with its attributes.

Hitachi Block Replications Inventory

This page provides an inventory of all the Replications stored on a Block Storage node.

Caution: Deleting a replication record ( ) or removing it from a host group ( ), when it is allocated to a host, may cause the host or an application running on it to fail.

Note: After a replication operation that created or adopted S-VOLs has been removed from a data flow, or the data flow has been deactivated, the corresponding record (marked) should be torn down then deleted from the inventory once it is no longer required.

Records for replications in active data flows (marked *Refreshed Full Copy* or *Active Full Copy*) cannot be deleted.

If an S-VOL has been added to a host group, or has a LUN path allocated outside of Protector then it will not be deleted.

If an S-VOL has been renamed outside of Protector, then it will not be deleted because Protector will be unable to locate it.

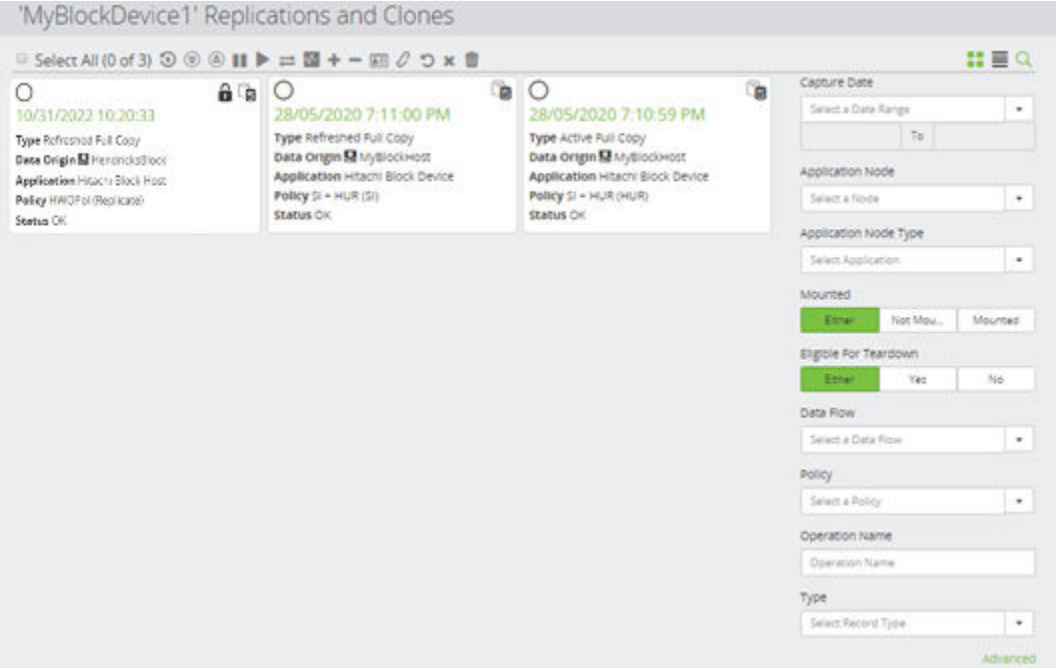


Figure 589 Hitachi Block Replications Inventory

Control	Description
Restore	Enabled only if a VMware Replication is selected. Opens the Hitachi Block VMware Snapshot Restore Wizard (on page 766).
Mount	Enabled only if one or more Replications are selected. Opens the Hitachi Block Snapshot or Replication Mount Wizard (on page 772) to guide you through mounting the Replication.

Control	Description
	Note: The mount operation can take several minutes to complete.
Unmount	Enabled only if one mounted replication is selected. Unmounts the selected replication. **Note:** The unmount operation can take several minutes to complete.
Pause	Enabled only if one or more Replications are selected. Pauses the Replication. Opens the following dialog:  Selecting this option will pause the replication and make the S-VOLs writable (instead of the P-VOL). This operation is carried out without needing to contact the primary array or storage proxy to allow continuation from a primary site failure. **Note:** This option is only relevant for remote replications – if selected for local replication it will be ignored.
Resume	Enabled only if one or more Replications are selected. Resumes a paused Replication.
Swap	Enabled only if one or more Replications are selected. Opens the Hitachi Block Replication Swap Wizard (on page 818) to guide you through swapping the Replication direction.
Unsuspend	If a Swap operation cannot be completed due to a P-VOL or data link fault between the primary and secondary device, then the replication will enter the SSWS state (suspended for swapping) indicating that the swap is not yet complete. Unsuspend enables the replication process to be re-established once the cause has been rectified.
Add to additional Host Groups	Enabled only if one or more replications are selected. Opens the Hitachi Block Add to Host Groups Dialog (on page 886) to enable LDEVs to be added to host groups in addition to the default `ProtectorProvisionedHostGroup` used by Protector.

Control	Description
	Note: A job is queued to perform the operation on the array and this may take some time to complete.
Remove from Host Groups	Enabled only if one replication is selected. Opens the Hitachi Block Remove from Host Groups Dialog (on page 887) to enable LDEVs to be removed from host groups, including the default `ProtectorProvisionedHostGroup` used by Protector. **Note:** A job is queued to perform the operation on the array and this may take some time to complete.
Dissociate	Enabled only if one or more Replications are selected. Dissociates a replication that is currently managed by Protector. Removes the selected replication(s) from Protector including state information such as direction and mount location. The replication remains active on the hardware device(s). The word 'DISSOCIATE' must be entered before the command is executed. **Caution:** Dissociating a replication removes all knowledge of the replication from Protector, including state information such as direction and mount location.
Teardown	Enabled only if one or more Replications are selected. Initiates the teardown of the selected replications. The word 'TEARDOWN' must be entered before the command is executed.
Delete	Enabled only if one or more Replications are selected. Deletes the replication record from Protector. The replication is also removed from the block storage device.
DRU Protection	Indicates that the replication has DRU protection applied. The protection can be removed (subject to lock duration) via the Hitachi Block Record Details (Storage) (on page 875).
Replication Date and Time	Click on the Replication's Date and Time to open the Hitachi Block Replication Details (Storage) (on page 882).
Eligible for Teardown	Displayed on a replication tile next to the replication icon. Indicates that a replication is eligible for teardown.
Filter on Capture Date Time	Filters the displayed results based on the replication Caputre Date Time range specified in the Date Time Picker (on page 351).

Control	Description
Filter on Application Node	Filters the displayed results based on the Source Application Node from which the replication is made.
Filter on Application Node Type	Filters the displayed results based on the Source Application Node Type.
Filter on Mounted status	Filters the displayed results based on the mount status
Filter on Eligible for Teardown Status	Filters the displayed results based on the Eligible for Teardown status
Filter on Data Flow	Filters the displayed results based on the Data Flow where the replication is defined
Filter on Policy	Filters the displayed results based on the Policy where the replication is defined.
Filter on Operation Name	Filters the displayed results based on the operation name defined within the policy.
Filter on Type	Filters the displayed results based on the replication type.

Hitachi Block Replication Details (Storage)

This page shows the details of a Replication on a Hitachi Block storage device.

Note: Replication records must be manually removed here after they have been removed from the data flow. Records for replications currently in an active data flow cannot be torn down or deleted.

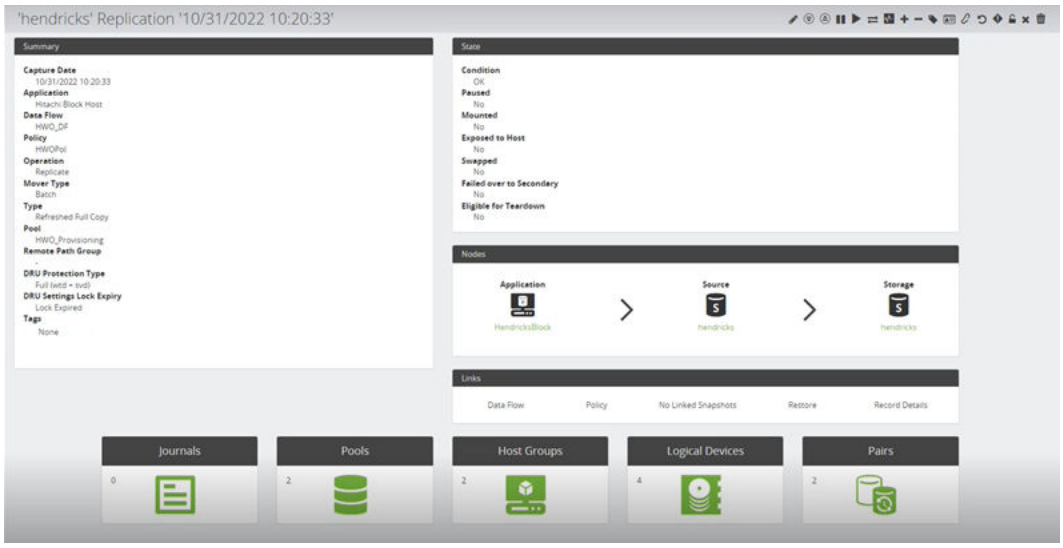


Figure 590 Replication Details

Control	Description
Rename Secondary LDEVs	Opens the Hitachi Block Rename Secondary LDEVs Dialog (on page 874) to enable renaming of LDEVs using a fixed string and/or substitution variables.
Mount	Opens the Hitachi Block Snapshot or Replication Mount Wizard (on page 772) to guide you through mounting the Replication. **Note:** The unmount operation can take several minutes to complete.
Unmount	Unmounts the selected replication. The export is deleted. **Note:** The unmount operation can take several minutes to complete.
Pause	Enabled only if one or more Replications are selected. Pauses the Replication. Opens the following dialog: 

Control	Description
	Selecting this option will pause the replication and make the S-VOLs writable without needing to contact the primary array or storage proxy to allow continuation from a primary site failure. **Note:** This option is only relevant for remote replications – if selected for local replication it will be ignored.
 Resume	Resumes a paused Replication.
 Swap	Opens the Hitachi Block Replication Swap Wizard (on page 818) to guide you through swapping the Replication direction.
 Unsuspend	If a Swap operation cannot be completed due to a P-VOL or data link fault between the primary and secondary device, then the replication will enter the SSWS state (suspended for swapping) indicating that the swap is not yet complete. Unsuspend enables the replication process to be re-established once the cause has been rectified.
 Add to additional Host Groups	Opens the Hitachi Block Add to Host Groups Dialog (on page 886) to enable LDEVs to be added to host groups in addition to the default `ProtectorProvisionedHostGroup` used by Protector. A job is queued to perform the operation on the array and this may take some time to complete.
 Remove from Host Groups	Opens the Hitachi Block Remove from Host Groups Dialog (on page 887) to enable LDEVs to be removed from host groups, including the default `ProtectorProvisionedHostGroup` used by Protector. A job is queued to perform the operation on the array and this may take some time to complete.
 Tag	Modifies the tags of an existing object from the either the inventory screen or the details screen of the object.
 Dissociate	▪ Dissociates a replication from Protector. ▪ Removes the selected replication(s) from Protector including state information such as direction and mount location. ▪ The replication remains active on the hardware device(s). ▪ The word 'DISSOCIATE' must be entered before the command is executed. **Caution:** Dissociating a replication removes all knowledge of the replication from Protector, including state information such as direction and mount location.

Control	Description
Force GAD IO Mode	Initiates the forcing of the GAD IO mode. The word 'FORCE' must be entered, and the storage array which you want to be the writable before the command is executed. The side selected will be made writable and if the Array and the Storage Proxy for the other side of the GAD is contactable it will be place into a Blocked state. **Note:** This feature can only be used on a paused or failed GAD replication. This can not be used on an Active Active GAD replication.
Remove DRU Protection	If the replication was created with DRU protection and the protection lock duration has expired, this will remove the protection.  A confirmation dialog opens to confirm the action of DRU Protection Removal. Clicking confirm will remove the DRU protection lock.
Teardown	Initiates the teardown of the replication. The word 'TEARDOWN' must be entered before the command is executed.
Delete	Deletes the replication record from Protector. The replication is also removed from the block storage device.
Summary	Identifies the hardware storage device, application, policy, DRU Protection Type, DRU Settings Lock Expiry, and general status.
State	Indicates the status of the replication.
Nodes	Indicates the Nodes involved in the replication. Opens the Node Details (on page 650) for the node selected.
Links	Provides links to the following pages that provide additional information about the replication: ▪ Data Flow - opens the Data Flow Details (on page 491) where the replication operation is assigned. ▪ Policy - opens the Policy Details (on page 737) where the replication operation is defined. ▪ Snapshots - displayed only if the replication's destination is snapshotted. Opens a filtered version of the Hitachi Block Snapshots Inventory (on page 869) that lists snapshot of the destination.

Control	Description
	▪ Restore - Opens the Hitachi Block Replication Details (Restore) (on page 821) ▪ Resource Details - Hitachi Block Record Details (Storage) (on page 875)
Journals	Opens Hitachi Block Replication - Journals (on page 888), listing the journals used in this replication. **Note:** You must manually refresh the journal information from the Hitachi Block Journals Inventory (on page 863) to ensure the current list of Journals is up to date
Pools	Opens Hitachi Block Replication - Pools (on page 889), listing the pools used in this replication. **Note:** : You must manually refresh the pool information from the Hitachi Block Pools Inventory (on page 866) to ensure the current list of Pools is up to date
Host Groups	Opens Hitachi Block Replication - Host Groups (on page 889), listing the host groups for this replication. **Note:** : You must manually refresh the Host Group information from the Hitachi Block Host Groups Inventory (on page 861) to ensure the current list of Host groups is up to date
Logical Devices	Opens Hitachi Block Replication - Logical Devices (on page 889), listing the LDEVs in this replication. **Note:** You must manually refresh the Logical Device information from the Hitachi Block Logical Devices Inventory (on page 868) to ensure the current list of Logical Devices is up to date
Pairs	Opens Hitachi Block Replication - Pairs (on page 890), listing the pairs in this replication.

Hitachi Block Add to Host Groups Dialog

This dialog is displayed when adding an existing replication S-VOL to another host group(s).

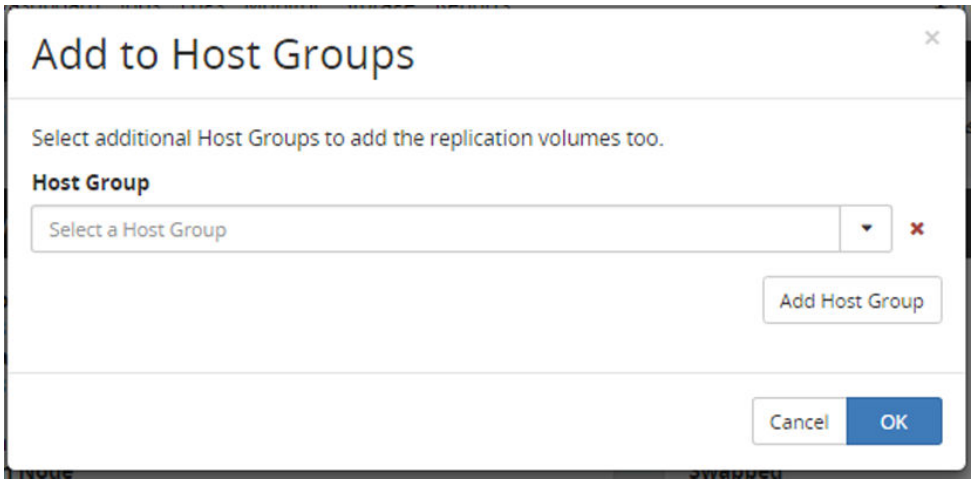


Figure 591 Add to Host Groups Dialog

Control	Description
Host Groups	Specify zero or more host groups that Protector will configure to provide access to the S-VOL. Click the Add Host Group button to insert another Host Group selection control. Click the Remove button next to a Host Group selection control to delete it. **Note:** If a LUN path to be created already exists, Protector will not attempt to add it again, or to change its ID. The specified host groups must be in the same resource group as the secondary volumes. For GAD replications, if the host group names and port IDs match between primary and secondary storage nodes, Protector will attempt to match the LUN IDs used for the S-VOLs with those of the respective P-VOLs. If this cannot be achieved then a warning will be logged and the next available LUN ID will be used.

Hitachi Block Remove from Host Groups Dialog

This dialog is displayed when removing an existing replication S-VOL from a host group(s).

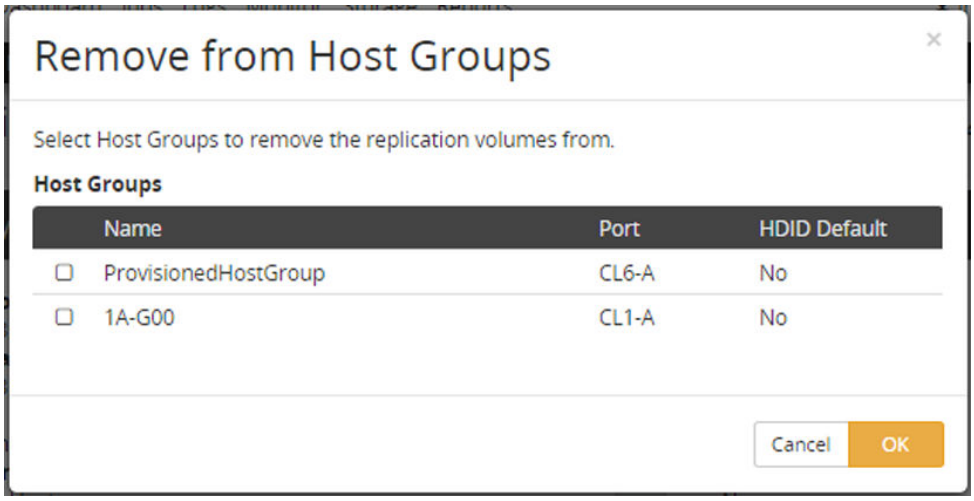


Figure 592 Remove from Host Groups Dialog

Control	Description
Host Groups	Specify one or more host groups to remove the S-VOL from. Click the check box next to the Host Group(s) from which the S-VOL is to be removed. **Note:** It is not possible to remove the last remaining Host Group if the S-VOL still has replication or snapshot pairings.

Hitachi Block Replication - Journals

Lists the journals used in a particular UR replication.

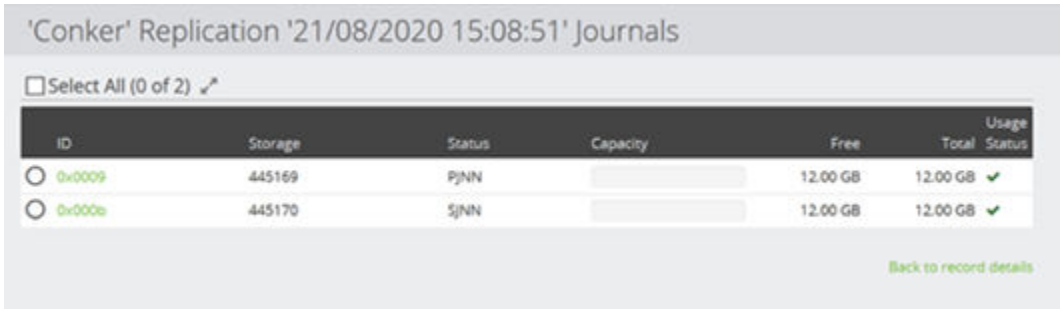


Figure 593 Replication Journals

Control	Description
Expand	Opens the Figure 575 Hitachi Block Journal Expansion Dialog (on page 864) to enable the capacity of the selected journal to be increased.

Control	Description
Replication Journals list	Each journal is listed, along with its attributes.

Hitachi Block Replication - Pools

Lists the pools used in a particular replication.



Figure 594 Replication Pools

Control	Description
Replication Pools list	Each pool is listed, along with its attributes.

Hitachi Block Replication - Host Groups

Lists the Host Groups used in a particular replication.

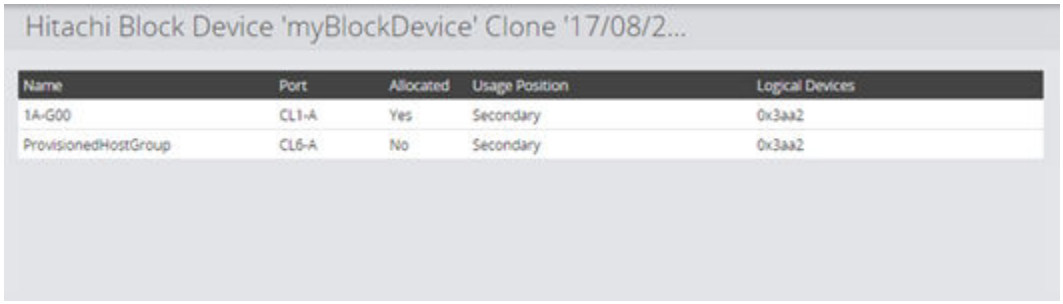


Figure 595 Replication Host Groups

Control	Description
Replication Host Groups list	Each host group is listed, along with its attributes.

Hitachi Block Replication - Logical Devices

Lists the LDEVs used in a particular replication.



Figure 596 Replication Logical Devices

Control	Description
Replication LDEVs list	Each LDEV is listed, along with its attributes.

Hitachi Block Replication - Pairs

Lists the pairs in a particular replication. Depending on pair type, this data is collected from [pairdisplay] or [raidcom get snapshot] – for more information about certain context-sensitive values, please consult the relevant CCI documentation.

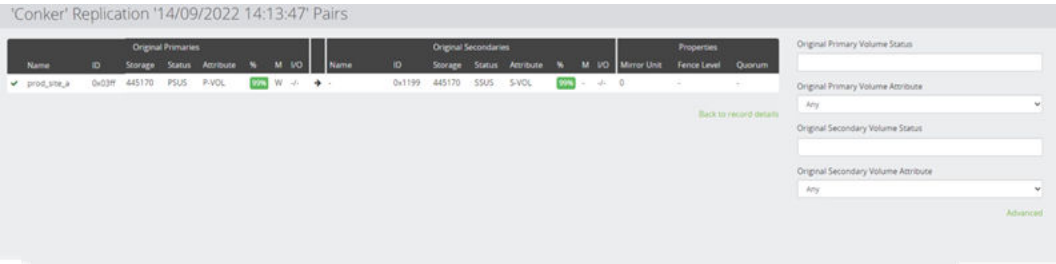


Figure 597 Replication Pairs

Control	Description
Name	Displays the Name of this LDEV.
ID	Displays the ID of this LDEV.
Storage	Displays the serial number of the storage array hosting this LDEV.
Status	Displays the status for this LDEV in this pair (eg COPY, PAIR, PSUS, SSUS, SSWS). **Note:** If pair status is not shown for the Pvol and/or Svol there might have been a change to the LU path allocation. This can happen when making changes outside of Protector such as using Administrator to provision High Availability pairs. To resolve it restart the Protector services on the Svol array ISM node.

Control	Description
Attribute	Displays the attribute for this LDEV in this pair (eg P-VOL, S-VOL, SMPL).
%	Displays the % value for this LDEV in this pair. The meaning of this value is context-sensitive, depending upon the pair status and pair type. This value is taken directly from the `pairdisplay` command. For guidance on how to interpret this value please consult the `pairdisplay` information in the CCI Command Reference.
M	Displays the M value for this LDEV in this pair. The meaning of this value is context-sensitive, depending upon the pair status and pair type. This value is taken directly from the `pairdisplay` command. For guidance on how to interpret this value please consult the `pairdisplay` information in the CCI Command Reference.
I/O Mode	Displays the I/O Mode for this LDEV (eg L/M, B/B). Only populated for pairs of type [Active-Active Remote Clone].
Direction (no column header)	Displays an arrow pointing from P-VOL to S-VOL.
Mirror Unit	Displays the mirror unit for the P-VOL of this pair.
Type	Displays the type of this pair (eg [SI], [TC]).
Fence Level	Displays the fence level for this pair. Only populated for pairs of remote replication types.
Quorum	Displays the quorum ID for this pair. Only populated for pairs of type [Active-Active Remote Clone].
Replication Pairs list	Each pair is listed, identifying the source LDEV's properties on the left and the destination LDEV's properties on the right, separated by an arrow indicating the direction of the replication.
Filter by Original Primary Volume Status	Filters the list by local volume status.
Filter by Original Primary Volume Attribute	Filters the list by local volume attribute.
Filter by Original Secondary Volume Status	Filters the list by remote volume status.
Filter by Original Secondary Volume Attribute	Filters the list by remote volume attribute.

Hitachi Block Remote Path Groups Inventory

This inventory lists all Remote Path Groups on a Hitachi Block storage device.

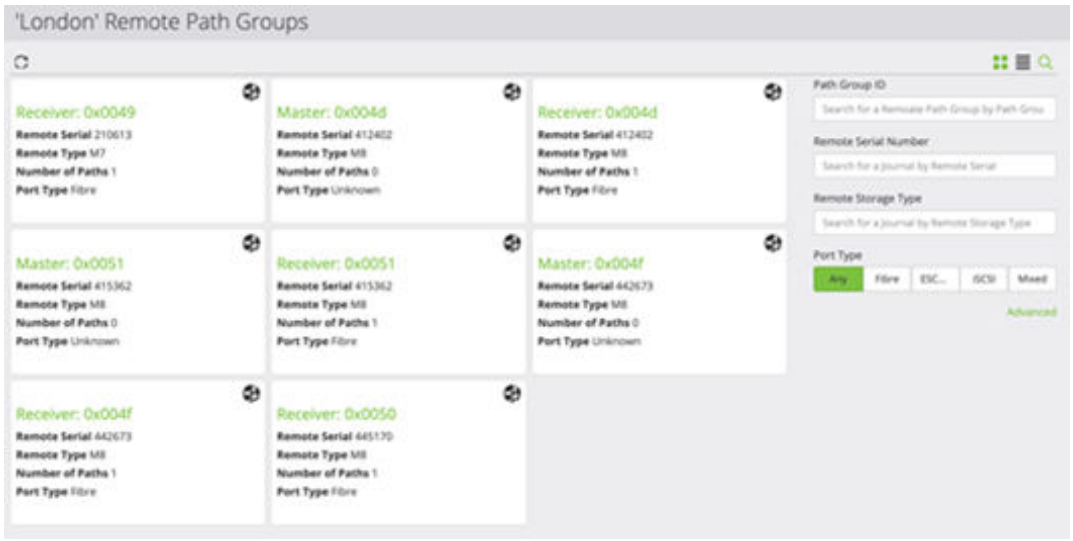


Figure 598 Remote Path Groups Inventory

Control	Description
Refresh	Refreshes the displayed inventory. **Caution:** Remote Path Group Inventory is only refreshed when refreshed manually. To ensure the latest information is displayed this page should be manually refreshed periodically.
Path Group	Click on a Path Group to open the Hitachi Block Remote Path Group Details (on page 892).
Filter on ID	Filters the displayed results based on the Path Group ID.
Filter on Serial	Filters the displayed results based on the serial number of the remote array.
Filter on Storage	Filters the displayed results based on the type of storage.
Filter on Port type	Filters the displayed results based on the type of port.

Hitachi Block Remote Path Group Details

This page shows the details of a Remote Path Group on a Hitachi Block storage device.



Figure 599 Remote Path Group Details

Control	Description
Summary	Summarises the Logical Device details.
Local	Details the remote path group.
Remote	Details the remote storage node.
Nodes	Provides links to both the local and remote storage screens.
Remote Paths	Shows the paths in the remote path group.
Outbound replications	Shows outbound (local) replications managed by Protector using the remote replication path.
Inbound replications	Shows inbound (remote) replications managed by Protector using the remote replication path.

Hitachi Block Device Advanced Settings Dialog

This dialog is displayed when the advanced settings for a block device node is being modified.

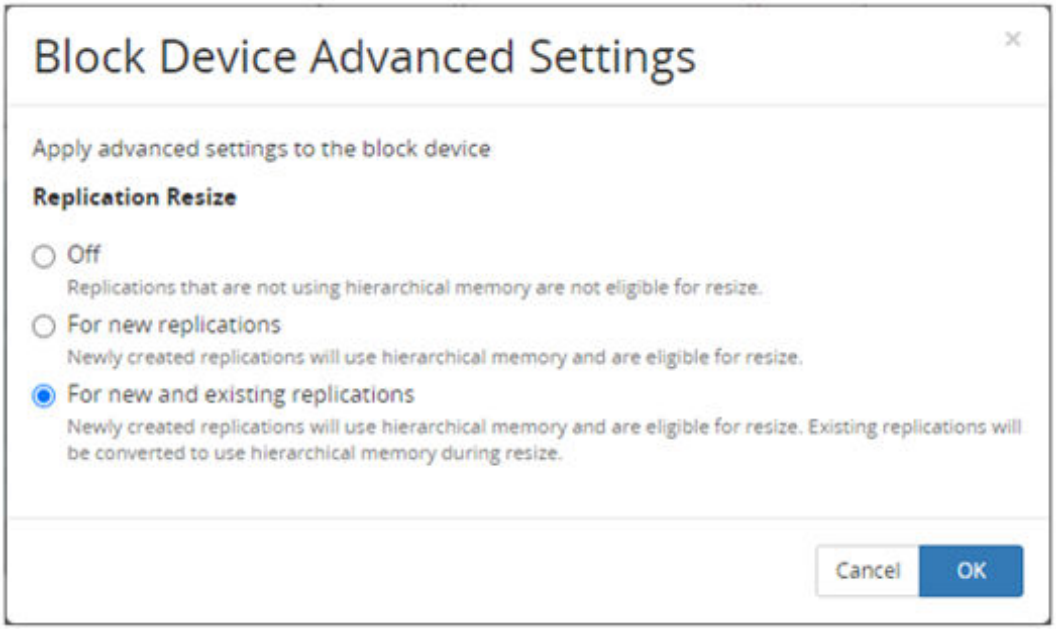


Figure 600 Block Device Advanced Settings Dialog

Control	Description
Replication Resize	Changes the replication resize option. This option determines if, or when replications will utilize Hierarchical Memory. A replication must be using hierarchical memory to support expansion of replicated Logical Devices. **Note:** The option "For new replications" may not always be available if it is not supported by the hardware.

Hitachi Block Device Scheduled Cache Refresh

The various caches for the Block Device are refreshed on a scheduled basis. The defaults for these refreshes are as detailed in the following table:

Item	Refresh Frequency
Pools	2 minutes
Journals	2 minutes
Resources	2 minutes
Quora	1 hour
Host Groups	1 hour
Remote Path Groups	1 hour

Item	Refresh Frequency
LDEVs	24 hours

These defaults are configurable on a per storage proxy basis so Block Devices sharing the same proxy node will have the same settings. If alternative configurations are required contact support for assistance.

WARNING: Excessive refreshing of these caches can put significant load on the storage arrays.

Hitachi Block Quorums Inventory

This page lists all the Quorums on a Hitachi Block Storage device.

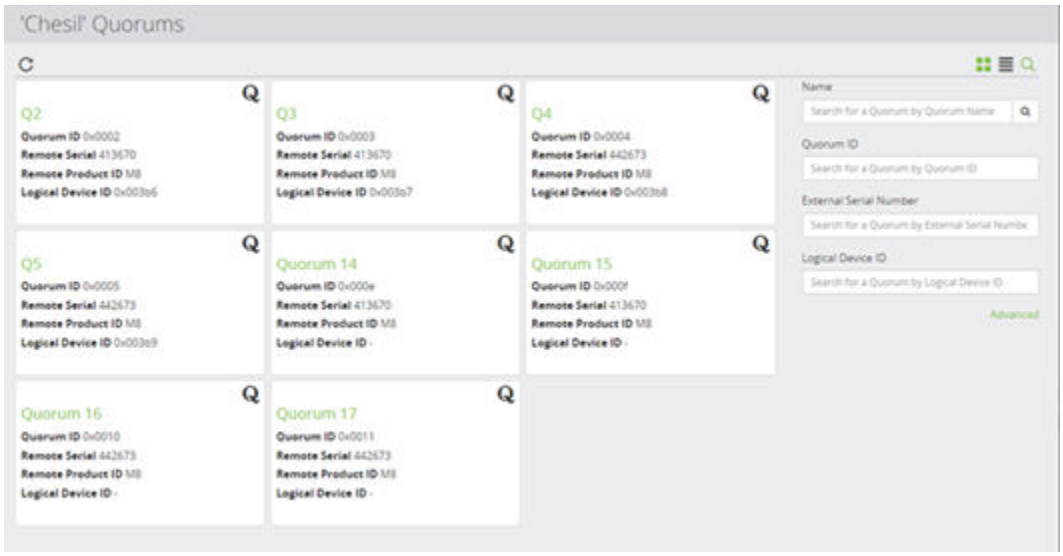


Figure 601 Quorums Inventory

Control	Description
Refresh	Refreshes the Cache and initiates the refreshing the Resources. **Note:** A refresh must be performed to update information held on the Quorum screen.
Filter on Name	Filters the displayed results based on the name of the Quorum.
Filter on Quorum ID.	Filters the displayed results based on the ID of the Quorum.

Control	Description
Filter on External Serial Number	Filters the displayed results based on the External Serial Number of the Quorum.
Filter on Logical Device ID.	Filters the displayed results based on the Logical Device ID. of the Quorum.

Hitachi Block Quorums Details

This page shows the details of the Quorums on a Hitachi Block Storage Device.

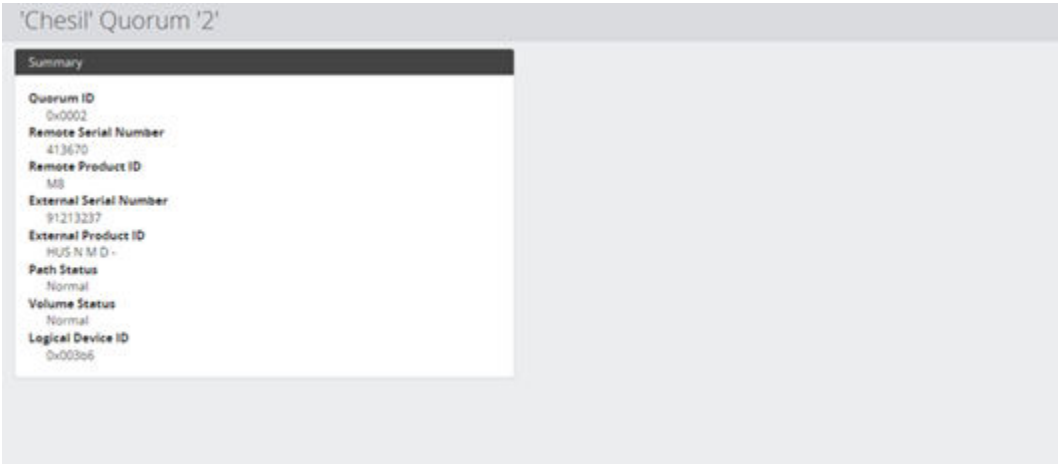


Figure 602 Quorums Details

Control	Description
Summary	Summarizes Quorum details.

Hitachi Block Resource Groups Inventory

This inventory lists all defined Resource Groups for this Block Device. Resource Group are created to define logical groups of computing resources in the context of Hitachi Block Storage Device.

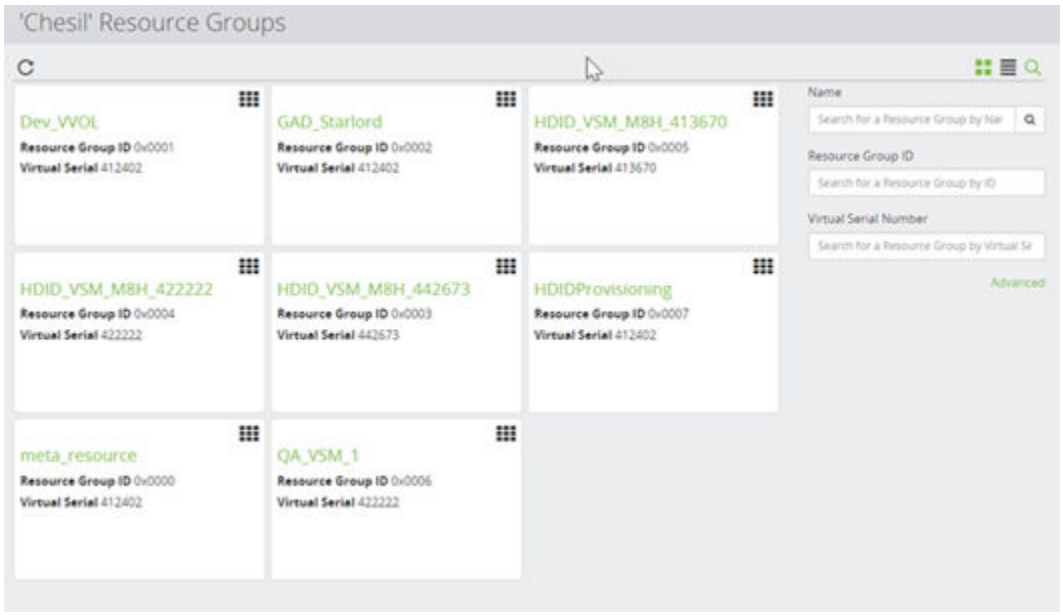


Figure 603 Hitachi Block Resource Groups Inventory

Control	Description
Refresh	Refreshes the Cache and initiates the refreshing the Resources. **Note:** A refresh must be performed to update information held on the Quorum screen.
Filter on Name	Filters the displayed results based on the name of the Resource Group .
Filter on Resource Group ID.	Filters the displayed results based on the ID of the Resource Group .
Filter on Virtual Serial Number	Filters the displayed results based on the Virtual Serial Number of the Resource Group .

Hitachi Block Resource Group Details

This page shows the details of the Resource Group on a Hitachi Block Storage Device.

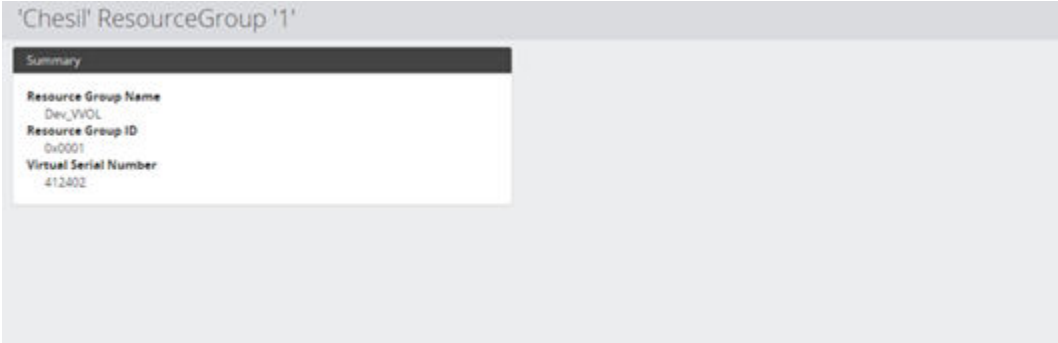


Figure 604 Resource Group Details

Control	Description
Summary	Summarizes Resource Group details.

Hitachi Block NVMe Subsystem Details

This page shows the details of the NVMe Subsystem on a Hitachi Block Storage Device.

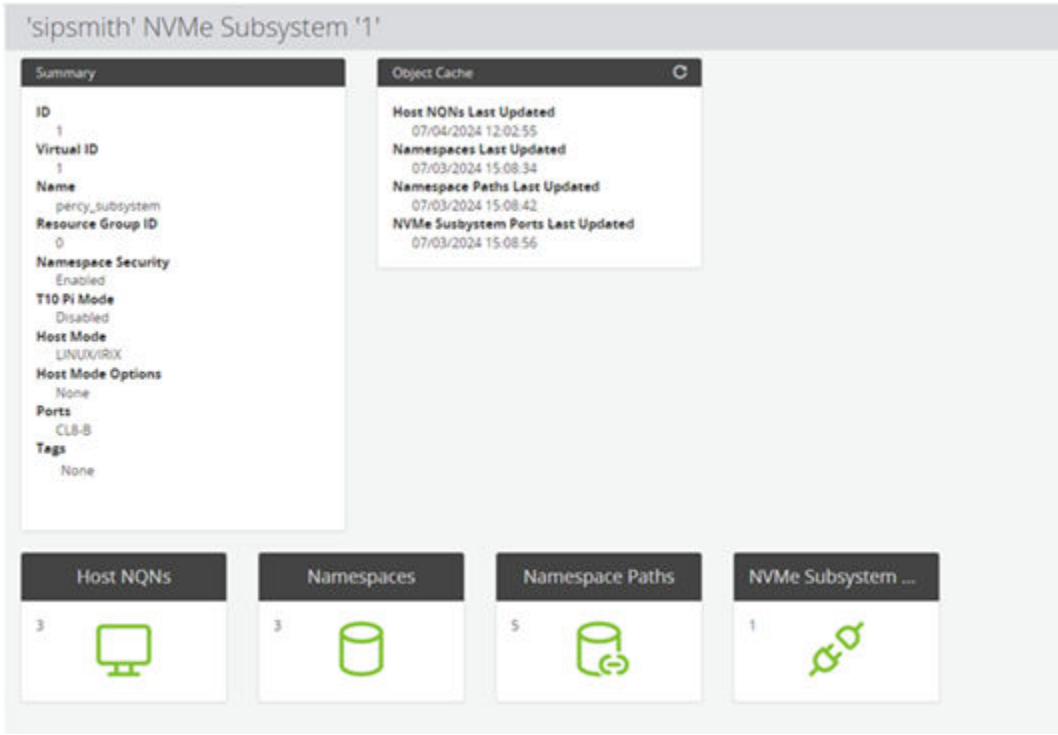


Figure 605 Hitachi NVMe Subsystem Details

Control	Description
Summary	Summarizes Resource Group details.

Control	Description
Object Cache	Displays the last update time for the various applicable caches.
 Host NQNs	Displays the configured Host NQNs, providing the following: ▪ Name ▪ Host NQN ▪ NVMe Subsystem ID ▪ NVMe Subsystem Name
 Namespaces	Displays the configured Namespaces, providing the following: ▪ Name ▪ ID ▪ NVMe Subsystem ID ▪ NVMe Subsystem Name ▪ LDev ID ▪ Capacity
 Namespace Paths	Displays the configured Namespace Paths, providing the following: ▪ NVMe Subsystem ID ▪ NVMe Subsystem Name ▪ Host NQN ▪ Namespace ID ▪ LDev ID
 NVMe Subsystems Ports	Displays the configured Namespace Ports, providing the following: ▪ NVMe Subsystem ID ▪ NVMe Subsystem Name ▪ Port

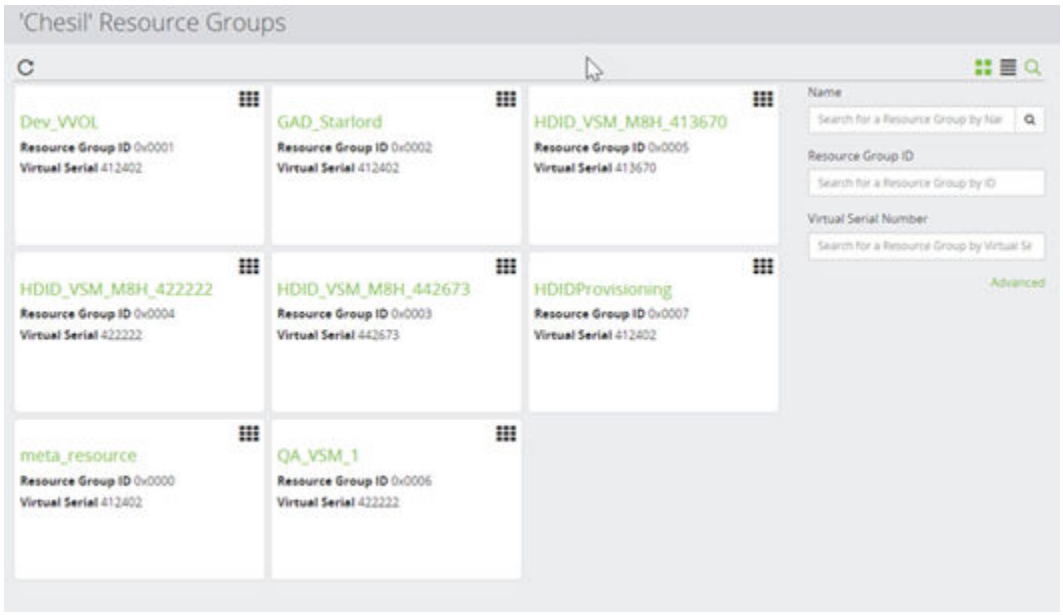


Figure 606 Hitachi Block Resource Groups Inventory

Control	Description
Refresh	Refreshes the Cache and initiates the refreshing the Resources. **Note:** A refresh must be performed to update information held on the Quorum screen.
Filter on Name	Filters the displayed results based on the name of the Resource Group .
Filter on Resource Group ID.	Filters the displayed results based on the ID of the Resource Group .
Filter on Virtual Serial Number	Filters the displayed results based on the Virtual Serial Number of the Resource Group .

Hitachi Block Resource Group Details

This page shows the details of the Resource Group on a Hitachi Block Storage Device.

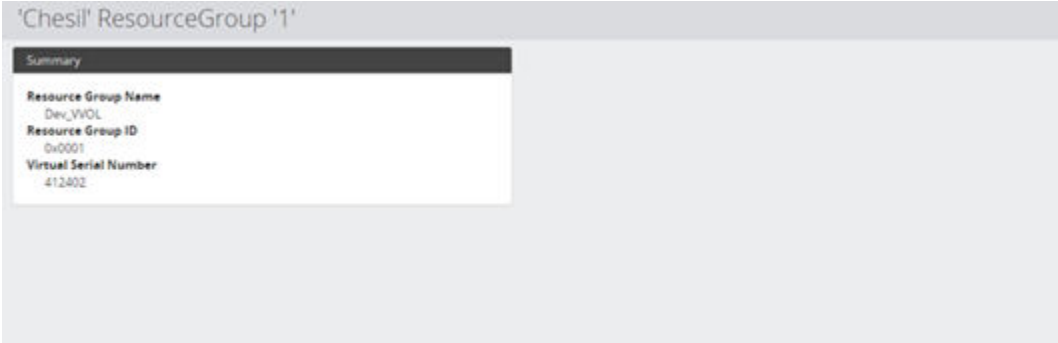


Figure 607 Hitachi Block Resource Group Details

Control	Description
Summary	Summarizes Resource Group details.

Hitachi Software Block Device Details

This page displays the details of a Software Block Storage device and enables monitoring of its activities, state, and the snapshots and replications stored within it. A Software Block device is used to represent both block and cloud software defined storage devices.

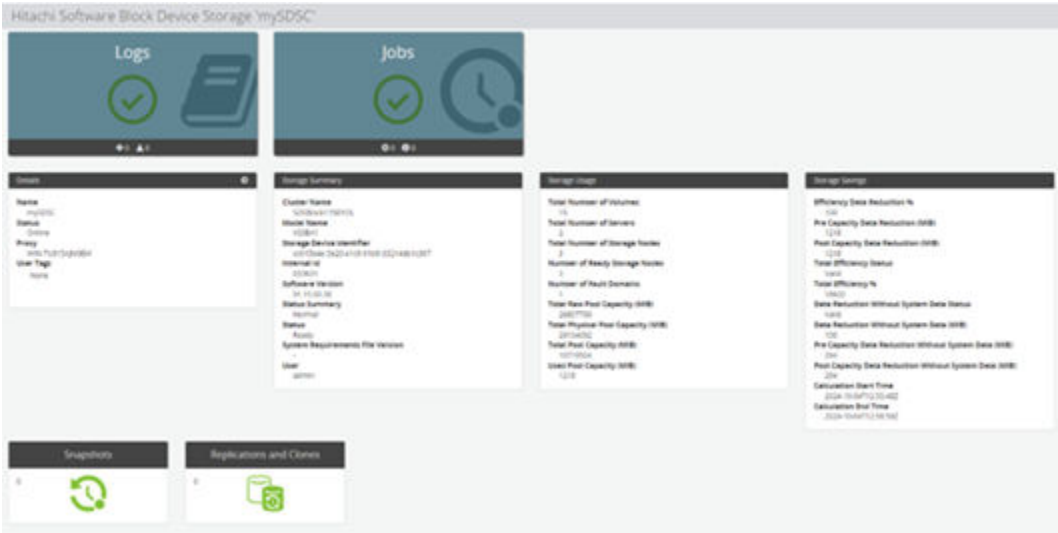


Figure 608 Hitachi Software Block Device Details

Control	Description
Logs	Indicates the number of unacknowledged error logs in red. If there are no errors then a green tick is displayed instead. Numbers at the bottom show the unacknowledged errors and warnings. Opens the Logs Inventory (on page 509).

Control	Description
Jobs	Indicates the number of failed jobs in red. If there are no failed jobs then a green tick is displayed instead. Numbers at the bottom show active and paused jobs. Opens the Jobs Inventory (on page 493).
Details	Indicates the following: ▪ Name - Name of the Software Block Device node ▪ Status - online/offline ▪ Proxy - the node that acts as the ISM for the Block device. ▪ User Tags – Tags associated with the Block device.
 View Node Details	Takes you to the equivalent Node Management Node Details (on page 650).
Storage Summary	Indicates the following details – Total Number of Volumes, Total Number of Servers, Total Number of Storage Nodes, Number of Ready Storage Nodes, Number of Fault Domains, Total Raw Pool Capacity (MiB), Total Physical Pool Capacity (MiB), Total Pool Capacity (MiB), Used Pool Capacity (MiB).
Configuration	Showsthe configuration of the Hitachi Block Device.
 Snapshots	Indicates the number of Snapshots on the Device that have been created by Protector. Click on the icon to open the Hitachi Software Block Snapshots Inventory (on page 902)
 Replications and Clones	Indicates the number of Replications on the Device that have been created or adopted by Protector. Click on the icon to open the Hitachi Software Block Replications Inventory (on page 905)

Hitachi Software Block Snapshots Inventory

This page provides an inventory of all the Snapshots stored on a Software Block Storage Device.

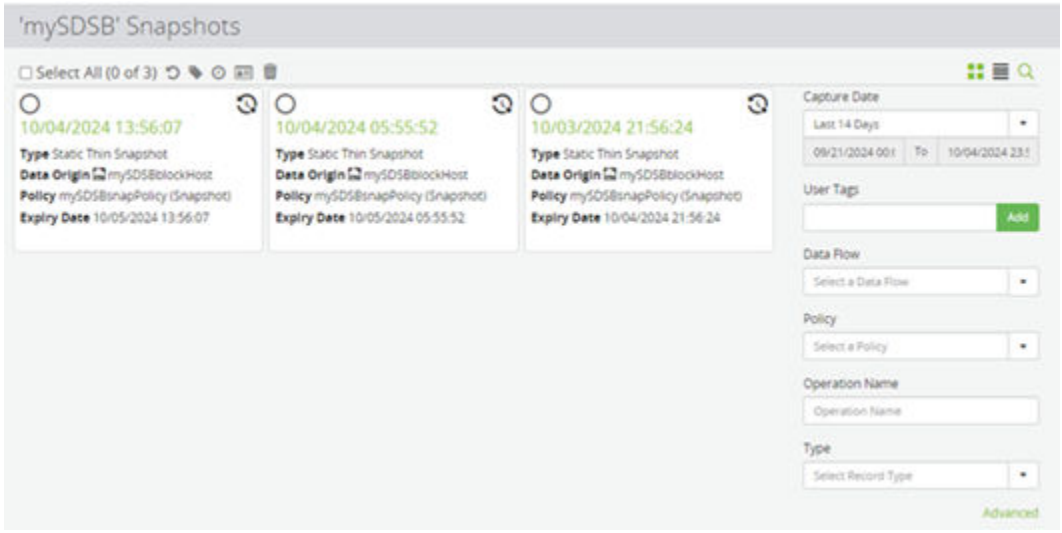


Figure 609 Hitachi Software Block Snapshots Inventory

Control	Description
Revert	Enabled only if one or more Snapshots are selected. Opens the Hitachi Block Revert Wizard (on page 813) to guide you through reverting the Snapshot.
Set Expiry Date	Opens the Hitachi Block Change Snapshot Expiry Date Dialog (on page 875) to enable the expiry date of the selected snapshot(s) to be modified.
Delete	Enabled only if one or more Snapshots are selected. Deletes the snapshot from the pool. A confirmation dialog opens to confirm the action. If the snapshot record contains fully provisioned LDEVs then the user can view these details prior to deletion.
Snapshot Date and Time	Click on the Snapshot's Date and Time to open the Hitachi Block Snapshot Details (Storage) (on page 871).
Filter on Capture Date	Filters the displayed results based on the Date Range specified in the Date Time Range Picker (on page 352).
Filter on Dataflow	Filters the displayed results based on Dataflow
Filter on Policy	Filters the displayed results based on the Policy.
Filter on Operation Name	Filters the displayed results based on mount status of the operation name

Control	Description
Type	Filters the displayed results based on type of the snapshot. Select from: ▪ Software Snapshot ▪ Static Full Snapshot ▪ Static Thin Snapshot

Hitachi Software Block Snapshot Details (Storage)

This page shows the details of a Snapshot on a Hitachi Software Block storage device.

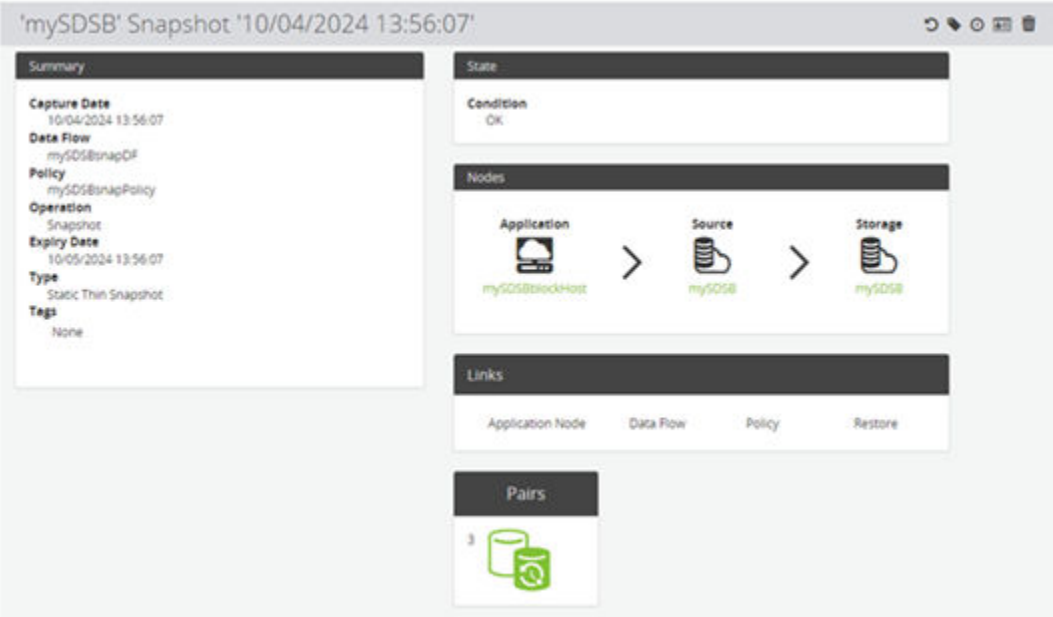


Figure 610 Storage Hitachi Software Block Snapshot Details

Control	Description
Revert	Opens the Hitachi Block Revert Wizard (on page 813) to guide you through reverting the snapshot.
Set Expiry Date	Opens the Hitachi Block Change Snapshot Expiry Date Dialog (on page 875) to enable the expiry date of the snapshot to be modified.
Delete	Deletes the snapshot from the pool. A confirmation dialog opens to confirm the action.
Summary	Identifies the hardware storage device, application, policy etc. Click on the policy name to open the Policy Details (on page 737).

Control	Description
State	Indicates the status of the snapshot.
Links	Provides links to the following pages that provide additional information about the snapshot: ▪ Application Node - opens the Node Details (on page 650) for the source node from which the snapshot data set originated. ▪ Data Flow - opens the Data Flow Details (on page 491) where the snapshot operation is assigned. ▪ Policy - opens the Policy Details (on page 737) where the snapshot operation is defined. ▪ Restore - opens the Hitachi Block Snapshot Details (Restore) (on page 819) for the snapshot.
Nodes	Shows the source and destination nodes involved in the snapshot. You can click on either node to go to their Node Details (on page 650)
Pairs	Opens Hitachi Block Snapshot - Pairs (on page 877), listing the pairs in this snapshot.

Hitachi Software Block Snapshot - Pairs

Lists the pairs in a particular snapshot. This data is collected from [raidcom get snapshot] – for more information about certain context-sensitive values, please consult the relevant CCI documentation.



Figure 611 Storage Hitachi Software Block - Snapshot Pairs

Control	Description
Volume UUID	Displays the UUID of the LDEV.
Direction (no columnheader)	Displays an arrow pointing from P-VOL to S-VOL.
Type	Displays the type of this pair (eg [TI]).

Hitachi Software Block Replications Inventory

This page provides an inventory of all the Replications stored on a Software Block Storage node.

Caution: Deleting a replication record ( ) or removing it from a host group ( ), when it is allocated to a host, may cause the host or an application running on it to fail.

Note: After a replication operation that created S-VOLs has been removed from a data flow, or the data flow has been deactivated, the corresponding record (marked) should be torn down then deleted from the inventory once it is no longer required.

Records for replications in active data flows (marked *RefreshedFull Copy* or *ActiveFull Copy*)cannot be deleted.



Figure 612 Hitachi Software Block Replications Inventory

Control	Description
Pause	Enabled only if one or more Replications are selected. Pauses the Replication. Opens the following dialog: 
Resume	Enabled only if one or more Replications are selected. Resumes a paused Replication.

Control	Description
 Swap	Enabled only if one or more Replications are selected. Opens the Hitachi Block Replication Swap Wizard (on page 818) to guide you through swapping the Replication direction.
 Unsuspend	If a Swap operation cannot be completed due to a P-VOL or data link fault between the primary and secondary device, then the replication will enter the SSWS state (suspended for swapping) indicating that the swap is not yet complete. Unsuspend enables the replication process to be re-established once the cause has been rectified.
	Enabled only if one or more Replications are selected. Dissociates a replication that is currently managed by Protector. Removes the selected replication(s) from Protector including state information such as direction and mount location. The replication remains active on the hardware device(s). The word 'DISSOCIATE' must be entered before the command is executed.  **Caution:** Dissociating a replication removes all knowledge ofthe replication from Protector, including state information such as direction and mount location.
 Teardown	Enabled only if one or more Replications are selected. Initiates the teardown of the selected replications. The word 'TEARDOWN' must be entered before the command is executed.
 Delete	Enabled only if one or more Replications are selected. Deletes the replication record from Protector. The replication is also removed from the block storage device.
 Eligible for Teardown	Displayed on a replication tile next to the replication icon. Indicates that a replication is eligible for teardown.
Filter on Capture Date Time	Filters the displayed results based on the replication Caputre Date Time range specified in the Date Time Picker (on page 351).
Filter on Application Node	Filters the displayed results based on the Source Application Node from which the replication is made.
Filter on Application Node Type	Filters the displayed results based on the Source Application Node Type.
Filter on Mounted status	Filters the displayed results based on the mount status

Control	Description
Filter on Eligible for Teardown Status	Filters the displayed results based on the Eligible for Teardown status
Filter on Data Flow	Filters the displayed results based on the Data Flow where the replication is defined
Filter on Policy	Filters the displayed results based on the Policy where the replication is defined.
Filter on Operation Name	Filters the displayed results based on the operation name defined within the policy.
Filter on Type	Filters the displayed results based on the replication type.

Hitachi Software Block Replication Details (Storage)

This page shows the details of a Replication on a Hitachi Software Block storage device.

Note: Replication records must be manually removed here after they have been removed from the data flow. Records for replications currently in an active data flow cannot be torn down or deleted.

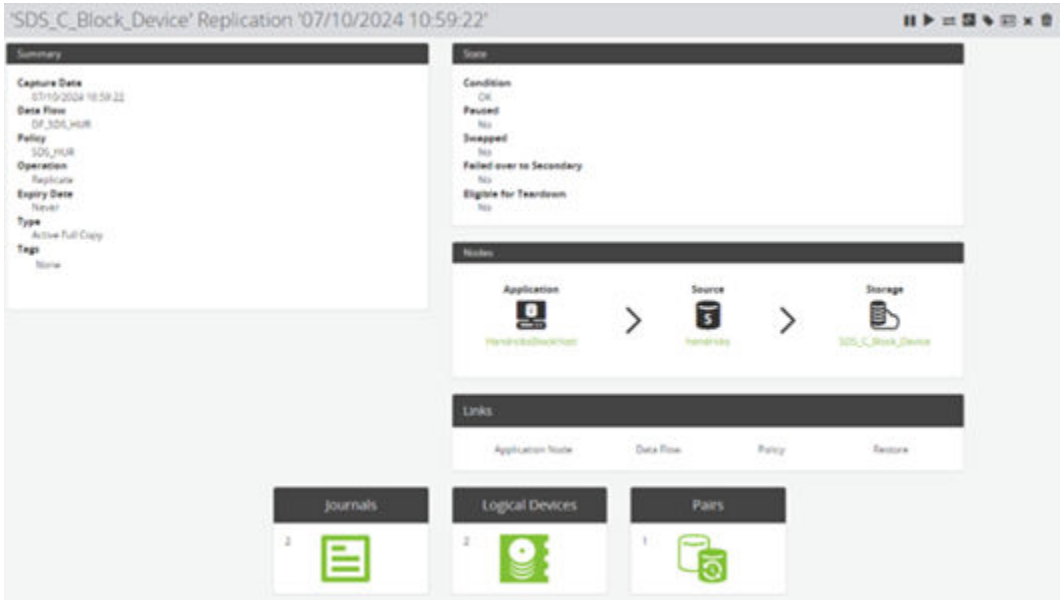


Figure 613 Hitachi Software Block Replication Details

Control	Description
Pause	Enabled only if one or more Replications are selected. Pauses the Replication. Opens the following dialog: 
Resume	Resumes a paused Replication.
Swap	Opens the Hitachi Block Replication Swap Wizard (on page 818) to guide you through swapping the Replication direction.
Unsuspend	If a Swap operation cannot be completed due to a P-VOL or data link fault between the primary and secondary device, then the replication will enter the SSWS state (suspended for swapping) indicating that the swap is not yet complete. Unsuspend enables the replication process to be re-established once the cause has been rectified.
Tag	Modifies the tags of an existing object from the either the inventory screen or the details screen of the object.
Teardown	Initiates the teardown of the replication. The word 'TEARDOWN' must be entered before the command is executed.
Delete	Deletes the replication record from Protector. The replication is also removed from the block storage device.
Summary	Identifies the software storage device, application, policy, and general status.
State	Indicates the status of the replication.
Nodes	Indicates the Nodes involved in the replication. Opens the Node Details (on page 650) for the node selected.

Control	Description
Links	Provides links to the following pages that provide additional information about the replication: ▪ Data Flow - opens the Data Flow Details (on page 491) where the replication operation is assigned. ▪ Policy - opens the Policy Details (on page 737) where the replication operation is defined. ▪ Restore - Opens the Hitachi Block Replication Details (Restore) (on page 821)
Journals	Opens Hitachi Software Block Replication - Journals (on page 910), listing the journals used in this replication.
Logical Devices	Opens Hitachi Software Block Replication Logical Devices (on page 910), listing the LDEVs in this replication.
Pairs	Opens Hitachi Software Block Replication - Pairs (on page 911), listing the pairs in this replication.

Hitachi Software Block Replication - Journals

Lists the journals used in a particular UR replication.

Figure 614 Replication Journals



Figure 615 Hitachi Software Block Replication Journals

Control	Description
Replication Journals list	Each journal is listed, along with its attributes.

Hitachi Software Block Replication Logical Devices

Lists the LDEVs used in a particular replication.



Figure 616 Hitachi Software Block Replication Logical Devices

Control	Description
Replication LDEVs list	Each LDEV is listed, along with its attributes.

Hitachi Software Block Replication - Pairs

Lists the pairs in a particular replication.



Figure 617 Hitachi Software Block Replicaiton - Pairs

Control	Description
Name	Displays the Name of this LDEV.
ID	Displays the ID of this LDEV.
Storage	Displays the serial number of the storage array hosting this LDEV.
Status	Displaysthe status for this LDEV in this pair (eg COPY, PAIR, PSUS, SSUS, SSWS). **Note:** If pair status is not shown for the Pvol and/or Svol there might have been a change to the LU path allocation. This can happen when making changes outside of Protector such as using Administrator to provision High Availability pairs. To resolve it restart the Protector services on the Svol array ISM node.
Attribute	Displays the attribute for this LDEV in this pair (eg P-VOL, S-VOL, SMPL).

Control	Description
%	Displays the % value for this LDEV in this pair. The meaning of this valueis context-sensitive, depending upon the pair status and pair type. This value is taken directly from the `pairdisplay`command. For guidance on how to interpret this value please consult the `pairdisplay`information in the CCI Command Reference.
M	Displays the M value for this LDEV in this pair. The meaning of this value is context-sensitive, depending upon the pair status and pair type. This value is taken directly from the `pairdisplay` command. For guidance on how to interpret this value please consult the `pairdisplay` information in the CCI Command Reference.
Direction (no column header)	Displays an arrow pointing from P-VOL to S-VOL.
Type	Displays the type of this pair (eg [SI], [TC]).
Filter by Original Primary Volume Status	Filters the list by local volume status.
Filter by Original Primary Volume Attribute	Filters the list by local volume attribute.
Filterby Original Secondary Volume Status	Filters the list by remote volume status.
Filter by Original Secondary Volume Attribute	Filters the list by remote volume attribute.

Gen1 Repository Stores Inventory

This inventory lists all the Stores within a Repository.

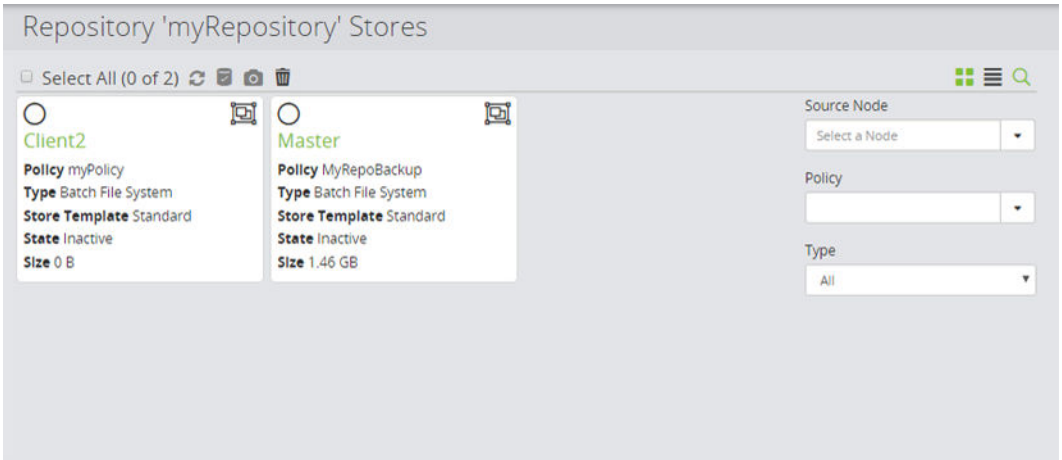


Figure 618 Repository Stores Inventory

Control	Description
Resynchronize	Enabled when one or more Stores are selected. Opens the Resynchronize Repository Store Dialog (on page 914) to resynchronize the stores with their sources.
Validate	Enabled when one or more Stores are selected. Validates the contents of the store against the respective source nodes. **Note:** Validation is done for checksums of data on the source machine against those in the repository stores. Checksums for data in repository stores that are tiered to the cloud are not validated against checksums held on the cloud storage platform.
User Snapshot	Enabled when only one Store is selected. Opens the Create Repository User Snapshot Dialog (on page 914) to create a single snapshot of the whole store – not just an individual policy. **Note:** User Snapshots are retained forever or until they are manually deleted.
Delete	Enabled when one or more Stores are selected. This option is only available if the store is inactive – that is, no active policies currently transfer data to a store with this name. Selecting it will delete this store and all the snapshots contained within it.
Existing Store(s)	Opens the Gen1 Repository Store Details (on page 915) enabling you to view and perform operations on the store.
Filter on Source Node	Filters the displayed results based on the Source node from which data is being ingested.

Control	Description
Filter on Policy	Filters the displayed results based on the Policy for which the store is used.
Filter on Type	Filters the displayed results based on the Store Type selected.

Resynchronize Repository Store Dialog

This dialog is displayed prior to a Repository Store being resynchronized by the user. It presents additional functions that can be performed during a resynchronization.

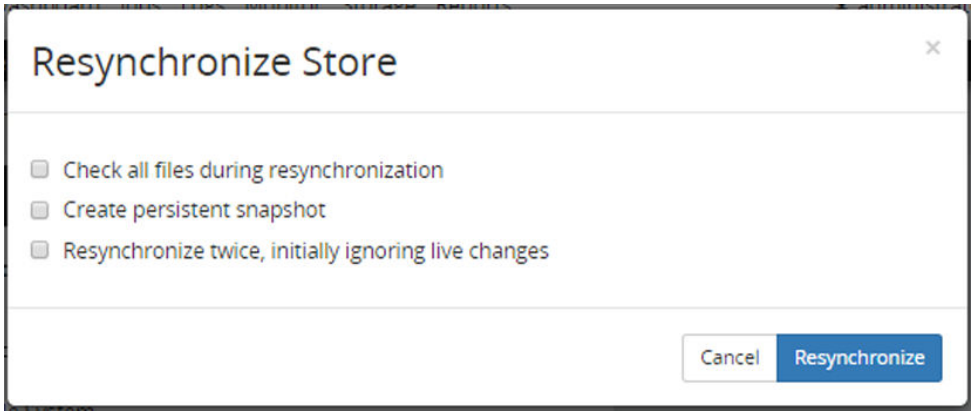


Figure 619 Resynchronize Store Dialog

Control	Description
Check all files during resynchronization	Ignores file modification date and checks every file for changes.
Create persistent snapshot	Once the resynchronization is complete, a persistent snapshot will be created. This snapshot will be retained indefinitely, but can be manually deleted.
Resynchronize twice, initially ignoring live changes	Two resynchronizations will be performed. The first will ignore any live changes that are made; the second will not ignore them. This option is only valid for a live policy and is ignored when using batch. It is used when there are problems with disk cache overflow. Using this option will reduce the disk usage.

Create Repository User Snapshot Dialog

This dialog is displayed when a User Snapshot is being created.

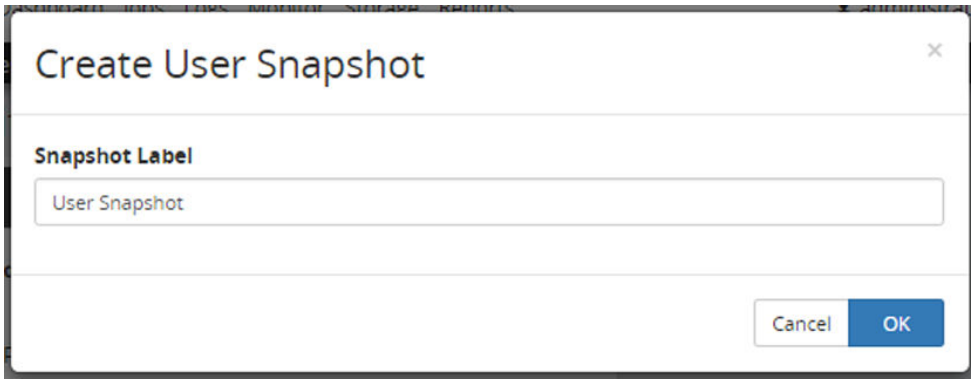


Figure 620 Create User Snapshot Dialog

Control	Description
Snapshot Label	Enter a name to identify the snapshot.

Gen1 Repository Store Details

This page provides details of a Store within a Repository and contains an inventory of all the Snapshots within that Store.

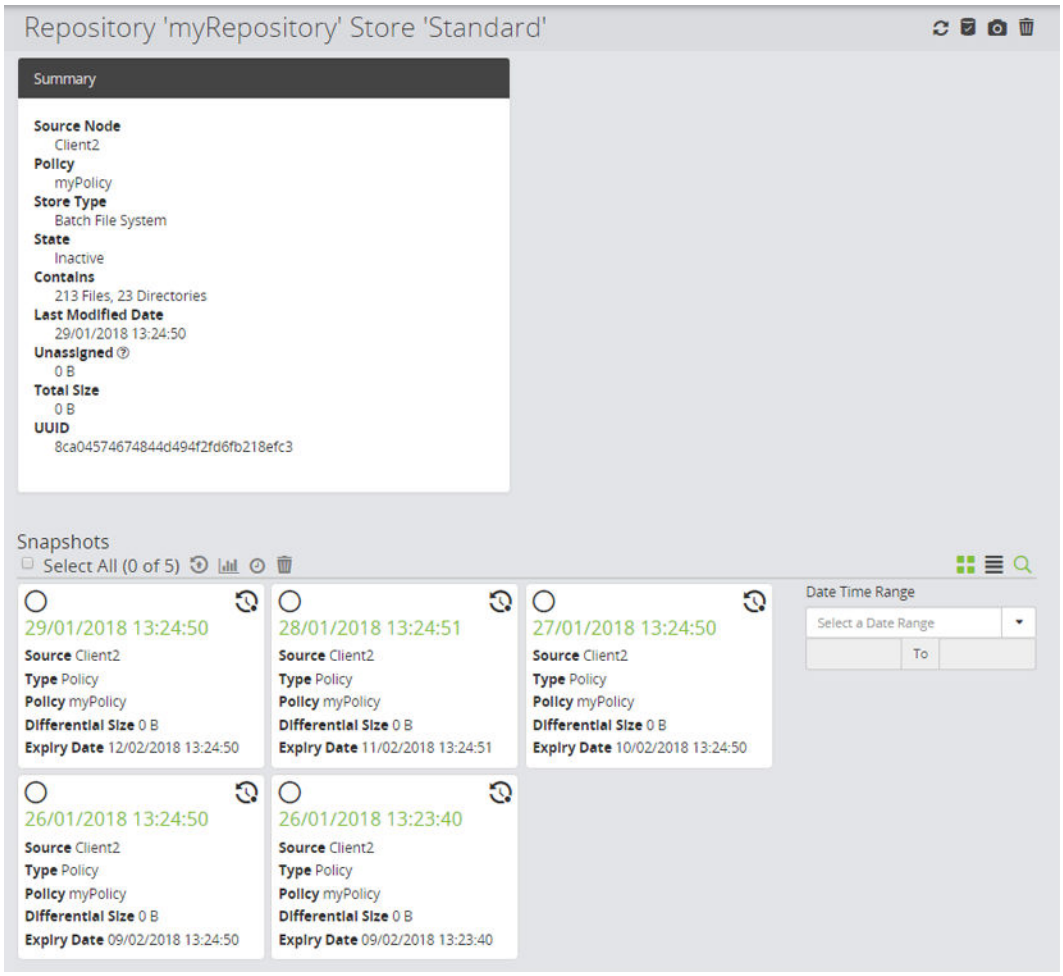


Figure 621 Repository Store Details

Control	Description
Resynchronize	Opens the Resynchronize Repository Store Dialog (on page 914) to resynchronize the stores with their sources.
Validate	Validates the contents of the store against the respective source nodes.
User Snapshot	Opens the Create Repository User Snapshot Dialog (on page 914) to create a single snapshot containing all the data in the store at that point in time. The snapshot is of the whole store – not just an individual policy. The data from all policies whose destination is that store will be contained within the same user snapshot. **Note:** User snapshots are retained forever or until they are manually deleted.

Control	Description
Delete	This option is only available if the store is inactive – that is, no active policies currently transfer data to a store with this name. Selecting it will delete this store and all the snapshots contained within it.
Restore	Enabled only when a single snapshot is selected. Restores the snapshot. Opens the Restore Repository Snapshot Wizard - File System (on page 826) to guide you through the process.
Analyze	Enabled only when a single snapshot is selected. Analyzes the snapshot to provide details about file count, size and changes. This is done on an individual snapshot and can take a few seconds to several minutes to complete depending on repository load and snapshot size.
Change Retention	Enabled only when one or more snapshots are selected. Opens the Change Repository Snapshot Retention Dialog (on page 920)
Delete	Enabled only when one or more snapshots are selected. Deletes the selected snapshots. **Note:** The latest snapshot cannot be deleted by the user. Also, if the snapshot is the only one remaining in the store then it will not be automatically retired. This provides a safeguard in the scenario where backups fail for a period of time, since at least one backup will be retained.
Snapshot(s)	Opens the Repository Snapshot Details (Storage) - File System (on page 917) to enable you to view the snapshot and perform operations on it.
Filter on Date Range	Filters the displayed results based on the dates entered in the Date Time Range Picker (on page 352).

Repository Snapshot Details (Storage) - File System

This page provides details of a Snapshot within a Store and contains an inventory of all the Files within that Snapshots.

Note: The Restore UI contains a similar page (Repository Snapshot Details (Restore) - File System (on page 845)) with different data and fewer options.

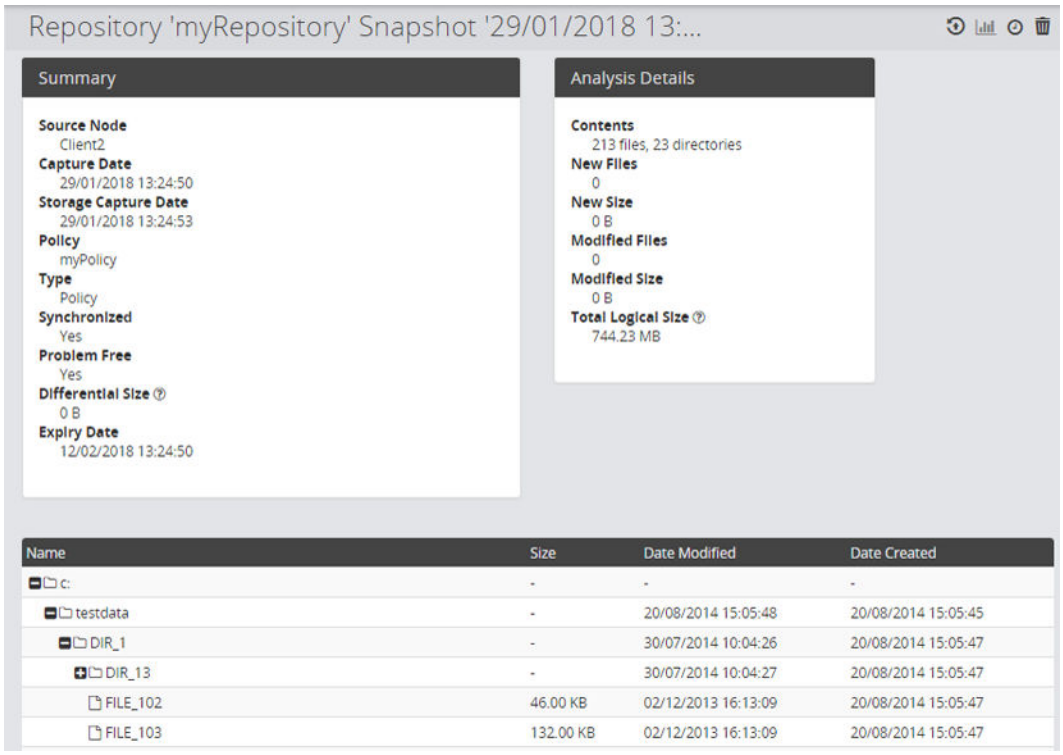


Figure 622 Repository Snapshot Details

Control	Description
Restore Snapshot	Restores the snapshot. Opens the Restore Repository Snapshot Wizard - File System (on page 826) to guide you through the process.
Analyze Snapshot	Analyzes the snapshot to provide details about file count, size and changes. This is done on an individual snapshot and can take a few seconds to several minutes to complete depending on repository load and snapshot size.
Change Retention	Enabled only when one or more snapshots are selected. Opens the Change Repository Snapshot Retention Dialog (on page 920)
Delete	Schedules the snapshot for deletion. The snapshot is not deleted immediately. No further operations can be performed on a snapshot pending deletion.
Summary	Provides summary information about the snapshot.
Analysis Details	Indicates details about file count, size and changes. It includes the logical size of all files contained within the snapshot.

Control	Description
Files System Contents	Displays the entire directory structure and files contained within this snapshot.

Repository Snapshot Details (Storage) - VMware

This page provides details of a VMware Snapshot within a Store and contains an inventory of all the files within that snapshot.

Note: The Restore UI contains a similar page (Repository Snapshot Details (Restore) - VMware (on page 846)) with different data and fewer options.

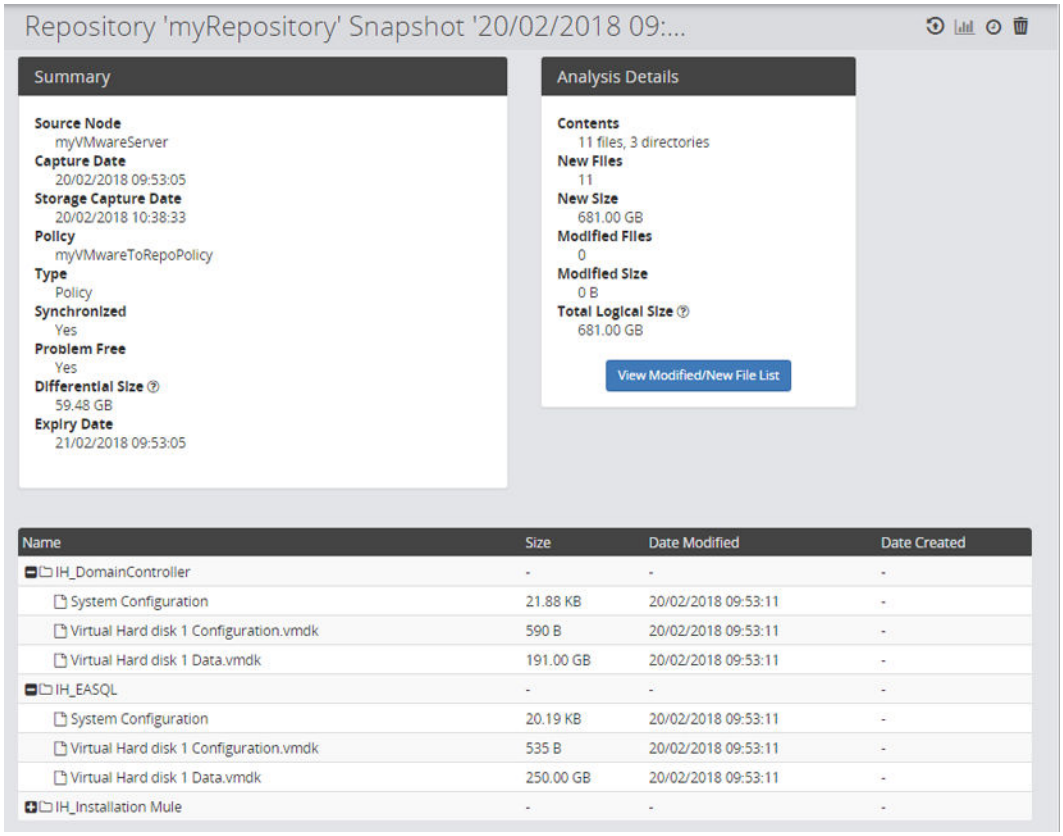


Figure 623 Repository Snapshot Details

Control	Description
Restore Snapshot	Restores the snapshot. Opens the Restore from host based backup Wizard - VMware (on page 840) to guide you through the process.

Control	Description
Analyze Snapshot	Analyzes the snapshot to provide details about file count, size and changes. This is done on an individual snapshot and can take a few seconds to several minutes to complete depending on repository load and snapshot size.
Change Retention	Enabled only when a one or more snapshots are selected. Opens the Change Repository Snapshot Retention Dialog (on page 920)
Delete	Schedules the snapshot for deletion. The snapshot is not deleted immediately. No further operations can be performed on a snapshot pending deletion.
Summary	Provides summary information about the snapshot.
Analysis Details	If the repository has never been analyzed then an Analyze button is displayed. Otherwise, it indicates details about file count, size and changes. It includes the logical size of all files contained within the snapshot. Click the View Modified/New File List to show details of the modified or new files in this snapshot (see below).
Snapshot Contents	Displays the entire contents within this snapshot.

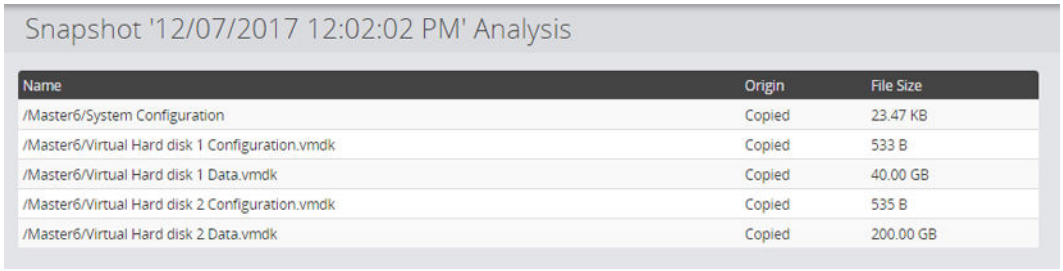


Figure 624 New/Modified Files

Change Repository Snapshot Retention Dialog

This dialog is displayed when the retention date of a snapshot is changed.

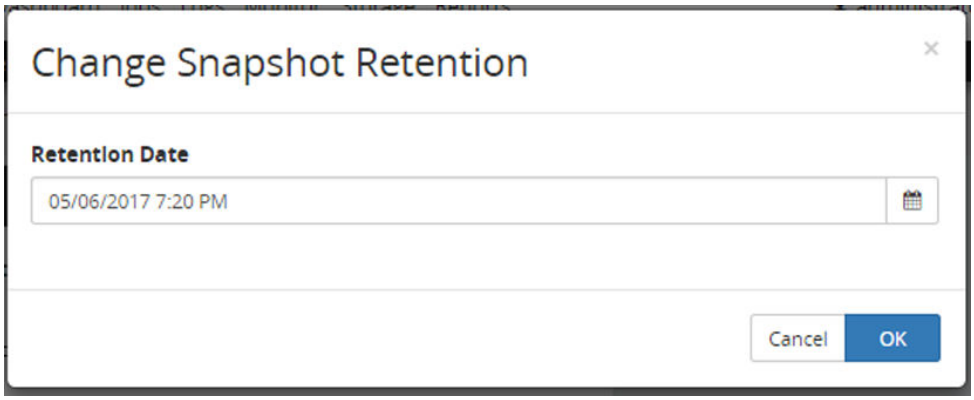


Figure 625 Change Snapshot Retention Dialog

Control	Description
Retention Date	Select the new date (using the Date Time Picker (on page 351)) after which the backup data will be retired

Generation 1 Repository Details

This page displays the details of a Gen1 Repository and enables you to monitor its activities, state and the snapshots stored within it.

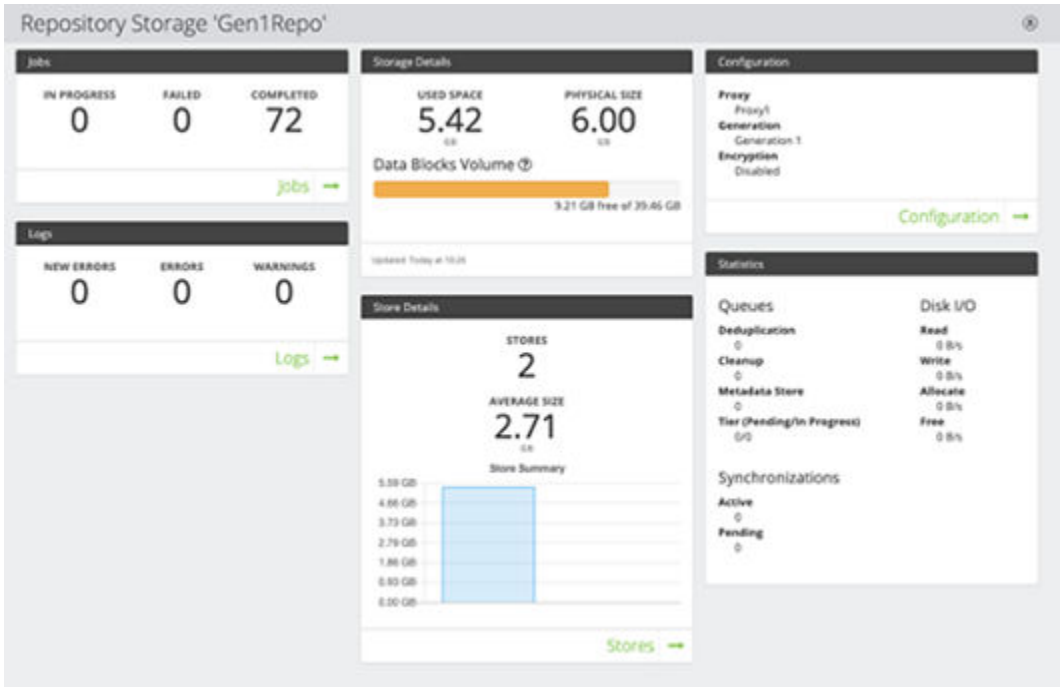


Figure 626 Repository Details

Control	Description
 Mount	Mounts the repository to either its original location or to an alternative location. If the repository is being mounted to an alternate location, a full resynchronization will be required.
 Unmount	Unmounts the repository.
Jobs	Indicates the number of In Progress jobs as well as the Failed and Completed jobs in the last 72 hours. The Jobs link opens the Jobs Inventory (on page 493).
Logs	Indicates the number of New Errors, Errors and Warnings in the last 72 hours. New errors are errors that have not been acknoleged. The Logs link opens the Logs Inventory (on page 509).
Storage Details	Indicates the following: ▪ Physical size of the repository as well as the used space within the physical size ▪ Disk space statistics - used, free and total
Stores Details	Indicates the number of Stores within the Repository and the size of each store. The Stores link opens the Gen1 Repository Stores Inventory (on page 912).
Configuration	Indicates the following: ▪ The proxy node – The node hosting the repository ▪ The generation of the repository (Generation 1 or 2) ▪ Encrypted -yes/no
Statistics	Indicates the following: ▪ Queues - indicates the background activities being performed. ▪ Synchronizations ▪ Capacity ▪ Disk I/O (Megabytes/second)

Generation 2 Repository Details

This page displays the details of a Gen2 Repository and enables you to monitor its activities, state and the snapshots stored within it.

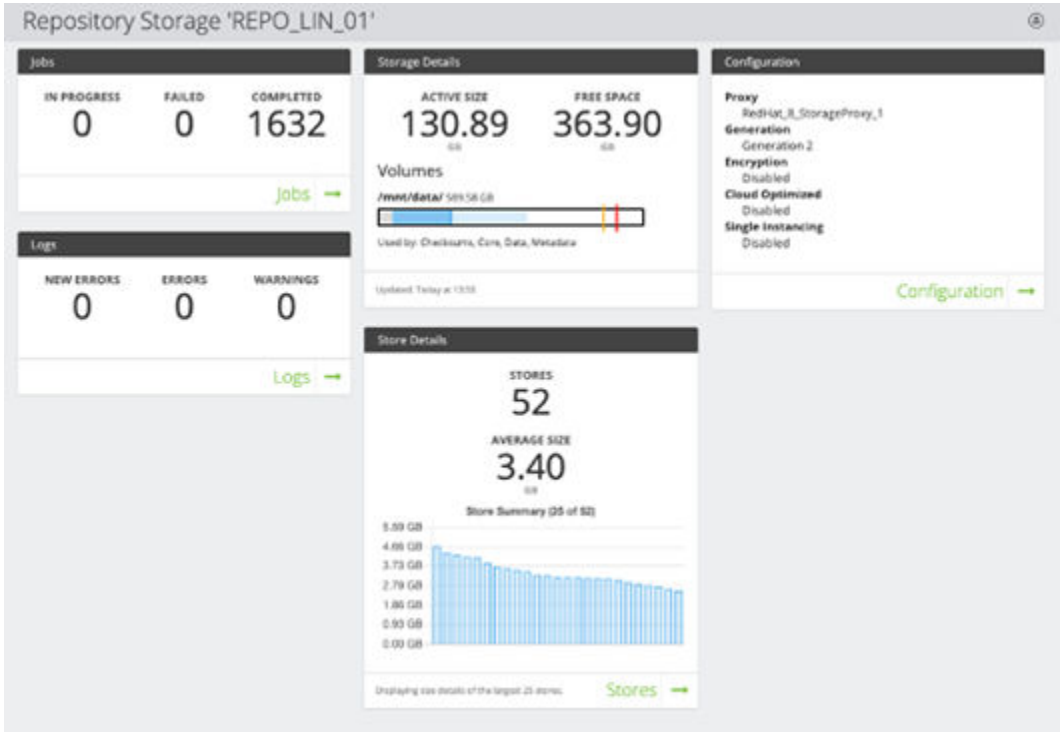


Figure 627 Repository Details

Control	Description
Mount	Mounts the repository to either its original location or to an alternative location. If the repository is being mounted to an alternate location, a full resynchronization will be required.
Unmount	Unmounts the repository.
Jobs	Indicates the number of In Progress jobs as well as the Failed and Completed jobs in the last 72 hours. The Jobs link opens the Jobs Inventory (on page 493).
Logs	Indicates the number of New Errors, Errors and Warnings in the last 72 hours. New errors are errors that have not been acknoleged. The Logs link opens the Logs Inventory (on page 509).

Control	Description
Storage Details	Indicates the following: ▪ Active Size – The size of the active data within the repository. This can be smaller than the physical size of the repository, in this case the active size can grow without the physical size growing. ▪ Free Space – The amount of free disk space available on the disk hosting the data part of the repository. ▪ Volumes – Shows disk space statistics for each volume the repository is hosted on. Typically repository metadata and data are stored on different volumes.
Stores Details	Indicates the number of Stores within the Repository and the size of each store. The Stores link opens the Gen2 Repository Stores Inventory (on page 924).
Configuration	Indicates the following: ▪ The proxy node – The node hosting the repository ▪ Generation - The generation of the ▪ Encrypted - Enabled / Disabled ▪ Cloud Optimized – Enables efficient transfer of data from the repository to cloud storage ▪ Single Instancing – Identical streams (files) are only stored once

Gen2 Repository Stores Inventory

This inventory lists all the Stores within a Repository.

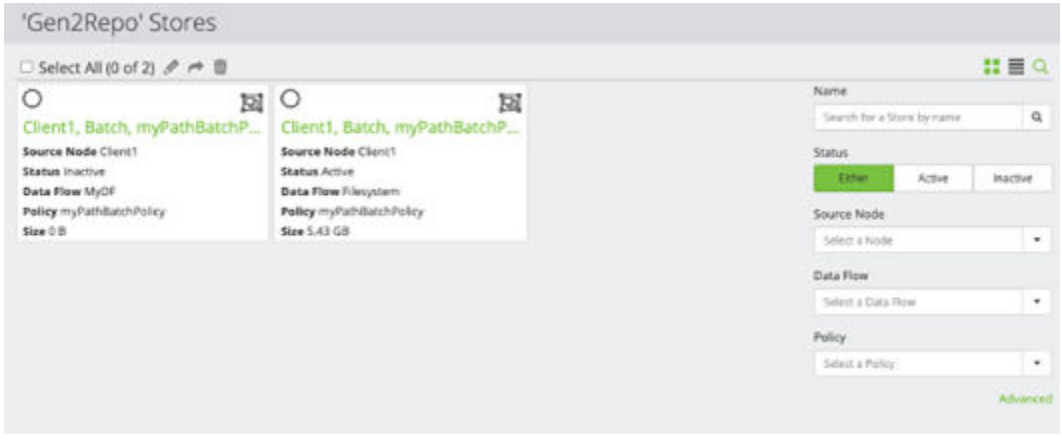


Figure 628 Gen2 Repository Stores Inventory

Control	Description
 Rename	Enabled when one Store is selected. This option renames the selected store and can be changed with affecting operations.
 Repurpose	Enabled when one store is selected. This option should only be used in very specific disaster recovery situations and under the guidance of customer support.
 Delete	Enabled when one or more Stores are selected. This option is only available if the store is inactive – that is, no active policies currently transfer data to a store with this name. Selecting it will delete this store and all the snapshots contained within it.
 Existing Store(s)	Opens the Generation 2 Repository Details (on page 922) enabling you to view and perform operations on the store.
Filter on Name	Filters the displayed results based on the name of the store.
Filter on Status	Filters the displayed results based on the status of the store.
Filter on Source Node	Filters the displayed results based on the Source node from which data is being ingested.
Filter on Dataflow	Filters the displayed results based on the Dataflow for which the store is used.
Filter on Policy	Filters the displayed results based on the Policy for which the store is used.

Gen2 Repository Store Details

This page provides details of a Store within a Repository and contains an inventory of all the Snapshots within that Store.

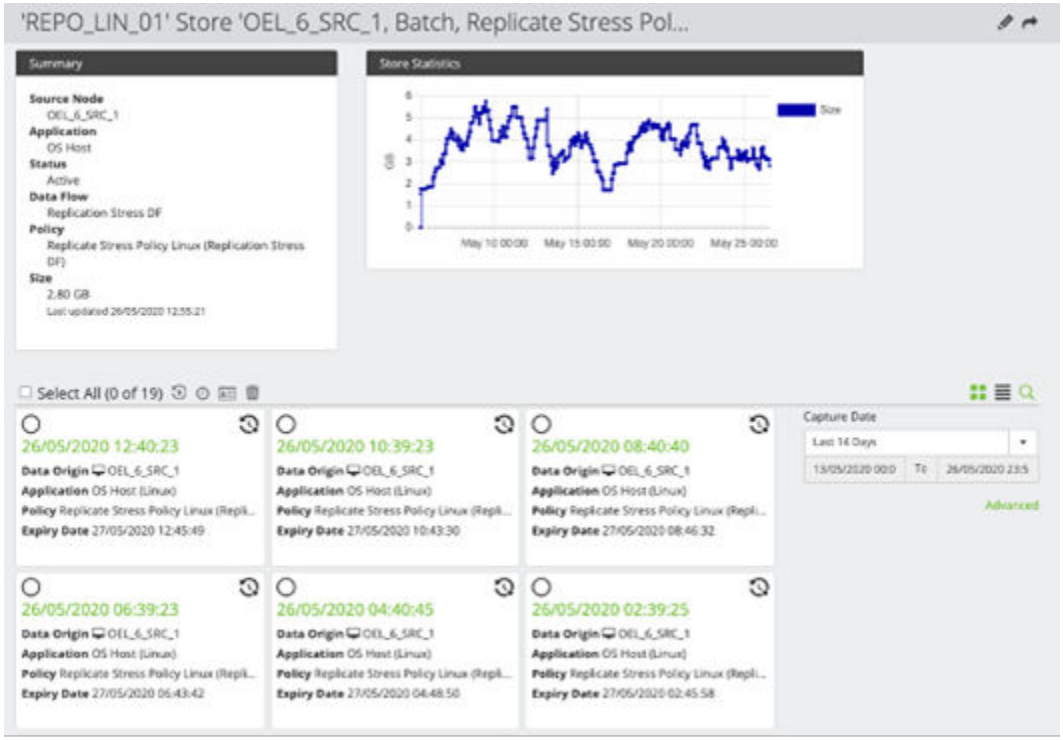


Figure 629 Gen2 Repository Store Details

Control	Description
Rename	Enabled when one Store is selected. This option renames the selected store and can be changed with affecting operations.
Repurpose	Enabled when one store is selected. This option should only be used in very specific disaster recovery situations and under the guidance of customer support.
Restore	Enabled only when a single snapshot is selected. Restores the snapshot. Opens the Restore Repository Snapshot Wizard - File System (on page 826) to guide you through the process..
Change Retention	Enabled only when one or more snapshots are selected. Opens the Change Repository Snapshot Retention Dialog (on page 920)
Transfer Ownership	Enabled only when one or more snapshots are selected. Allows the ownership of the recovery point to be changed. By changing the owning resource to a node that a user has access to enables the user to see and act upon the recovery point.
Delete	Enabled only when one or more snapshots are selected. Deletes the selected snapshots.

Control	Description
	Note: The latest snapshot cannot be deleted by the user. Also, if the snapshot is the only one remaining in the store then it will not be automatically retired. This provides a safeguard in the scenario where backups fail for a period of time, since at least one backup will be retained.
Filter on Date Range	Filters the displayed results based on the dates entered in the Date Time Range Picker (on page 352).

Repository Snapshot Details (Storage) - File System

Thispage provides details of a Recovery Point within a Store and contains an inventory of all the Files within that Recovery Point

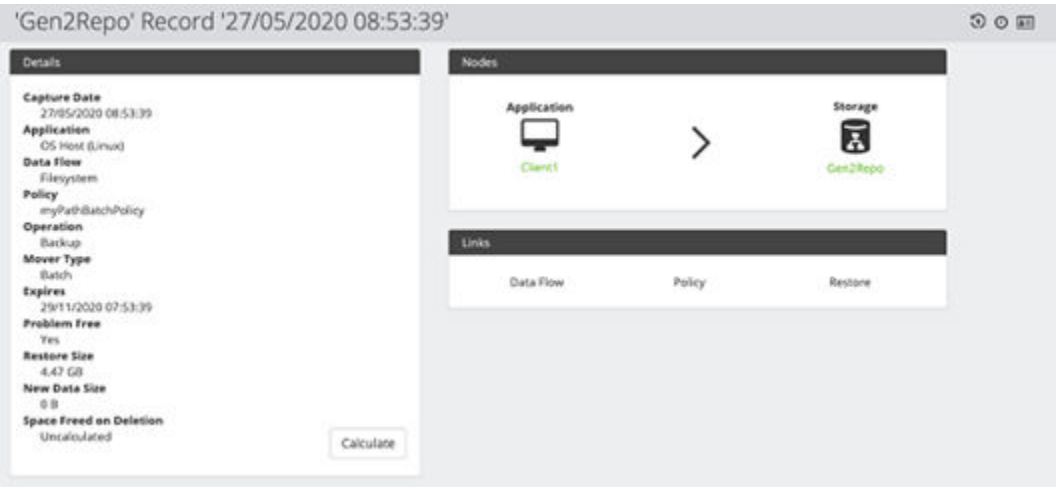


Figure 630 Gen2 Repository Recovery Point Details

Control	Description
Restore	Restores the snapshot. Opens the Restore Repository Snapshot Wizard - File System (on page 826) to guide you through the process.
Change Retention	Enabled only when one or more snapshots are selected. Opens the Change Repository Snapshot Retention Dialog (on page 920)
Transfer Ownership	Allows the ownership of the record to be changed. By changing the owning resource to a node that a user has access to enables the user to see and act upon the record.
Calculate	Calculates the amount of space that would be freed in the repository if this recovery point was deleted.

Control	Description
Files System Contents	Displays the entire directory structure and files contained within this snapshot.

Cloud Storage Details

This page displays the details of either an Amazon S3, HCP or HCP cloud scale container. It enables you to monitor its activities, state and the snapshots stored within it.

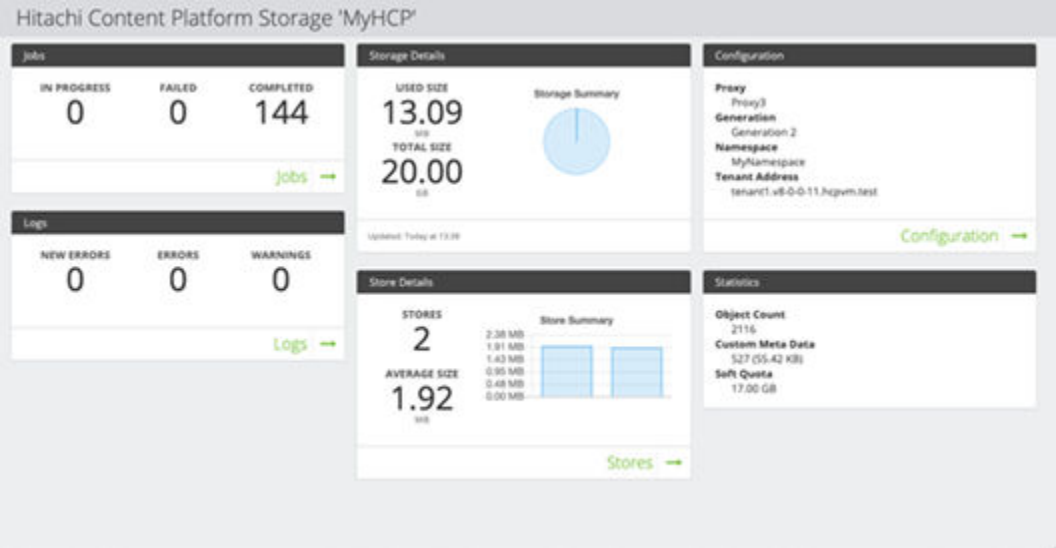


Figure 631 Cloud Details

Control	Description
Jobs	Indicates the number of In Progress jobs as well as the Failed and Completed jobs in the last 72 hours. The Jobs link opens the Jobs Inventory (on page 493).
Logs	Indicates the number of New Errors, Errors and Warnings in the last 72 hours. New errors are errors that have not been acknowledged. The Logs link opens the Logs Inventory (on page 509).
Storage Details	Indicates the following for HCP: ▪ The used size ▪ The total capacity Indicates the following for S3: ▪ The used size

Control	Description
Stores Details	Indicates the number of Stores within the Cloud and the size of each store. The Stores link opens the Cloud Stores Inventory (on page 929).
Configuration	Indicates the following for HCP: ▪ Proxy – The node hosting the HCP ▪ Generation - The generation of the HCP node ▪ Namespace - The HCP namespace used to store the data ▪ Tenant Address – The address of the HCP Tenant Indicates the following for S3 ▪ Proxy – The node hosting the S3 ▪ Bucket Name – The name of the bucket used to store the data ▪ Region – The Amazon region the bucket is located

Cloud Stores Inventory

This inventory lists all the Stores within a Cloud. This screen is the same for all cloud types.

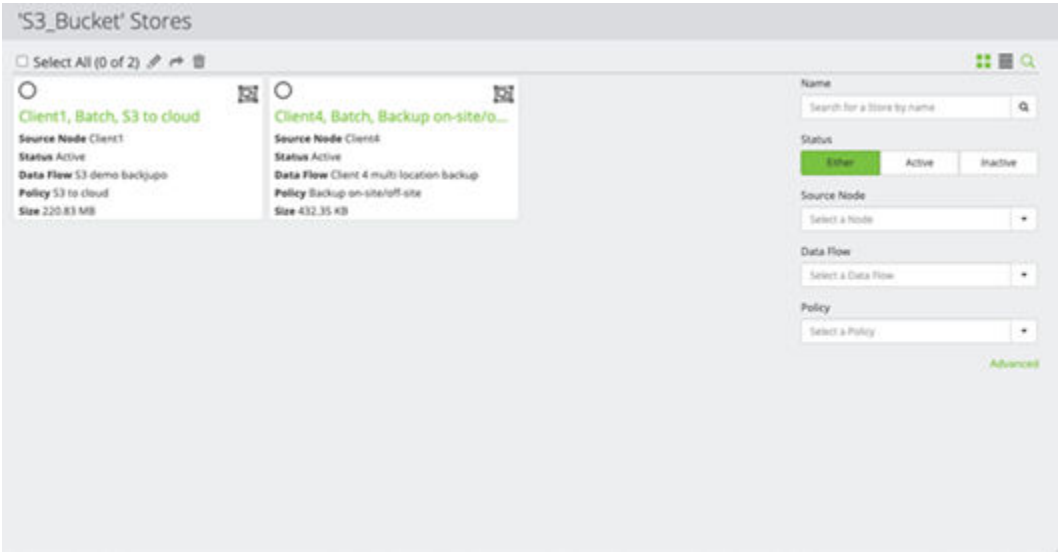


Figure 632 Cloud Stores Inventory

Control	Description
Rename	Enabled when one Store is selected. This option renames the selected store and can be changed with affecting operations.

Control	Description
Repurpose	Enabled when one store is selected. This option should only be used in very specific disaster recovery situations and under the guidance of customer support.
Delete	Enabled when one or more Stores are selected. This option is only available if the store is inactive – that is, no active policies currently transfer data to a store with this name. Selecting it will delete this store and all the snapshots contained within it.
Filter on Name	Filters the displayed results based on the name of the store.
Filter on Status	Filters the displayed results based on the status of the store.
Filter on Source Node	Filters the displayed results based on the Source node from which data is being ingested.
Filter on Dataflow	Filters the displayed results based on the Dataflow for which the store is used.
Filter on Policy	Filters the displayed results based on the Policy for which the store is used.

Cloud Store Details

This page provides details of a Store within a Cloud and contains an inventory of all the Snapshots within that Store. The screen is the same for all cloud types.

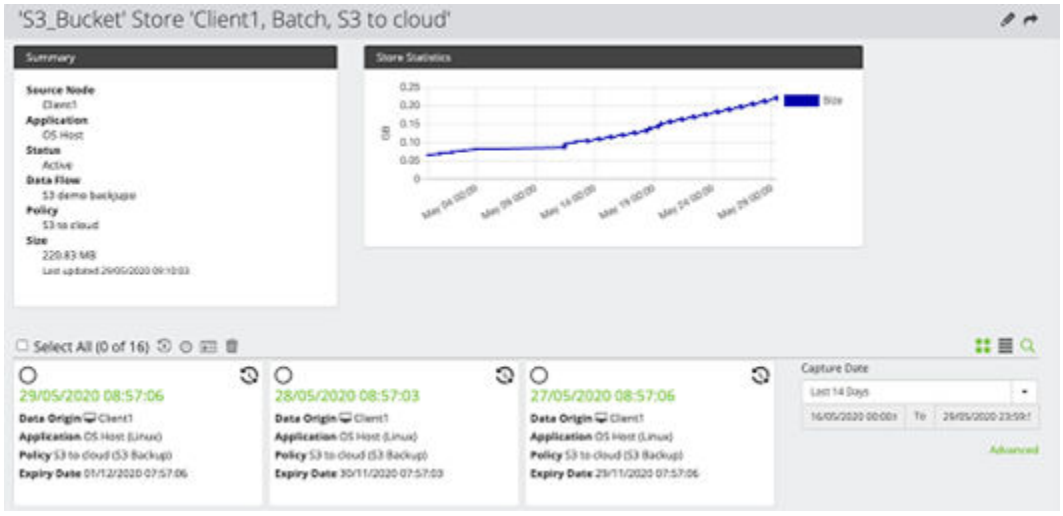


Figure 633 Cloud Store Details

Control	Description
Rename	Enabled when one Store is selected. This option renames the selected store and can be changed with affecting operations.
Repurpose	Enabled when one store is selected. This option should only be used in very specific disaster recovery situations and under the guidance of customer support.
Restore	Enabled only when a single snapshot is selected. Restores the snapshot. Opens the to guide you through the process.
Change Retention	Enabled only when one or more snapshots are selected. Opens the Change Repository Snapshot Retention Dialog (on page 920)
Transfer Ownership	Enabled only when one or more snapshots are selected. Allows the ownership of the recovery point to be changed. By changing the owning resource to a node that a user has access to enables the user to see and act upon the recovery point.
Set Expiry Date	Opens the Hitachi Block Change Snapshot Expiry Date Dialog (on page 875) to enable the expiry date of the selected snapshot(s) to be modified.
Transfer Owner	Enabled only when one or more snapshots are selected. Allows the ownership of the recovery point to be changed. By changing the owning resource to a node that a user has access to enables the user to see and act upon the recovery point.
Delete	Enabled only when one or more snapshots are selected. Deletes the selected snapshots. **Note:** The latest snapshot cannot be deleted by the user. Also, if the snapshot is the only one remaining in the store then it will not be automatically retired. This provides a safeguard in the scenario where backups fail for a period of time, since at least one backup will be retained.
Filter on Date Range	Filters the displayed results based on the dates entered in the Date Time Range Picker (on page 352).

Cloud Snapshot Details (Storage) - File System

This page provides details of a Recovery Point within a Store and contains an inventory of all the Files within that Recovery Point.

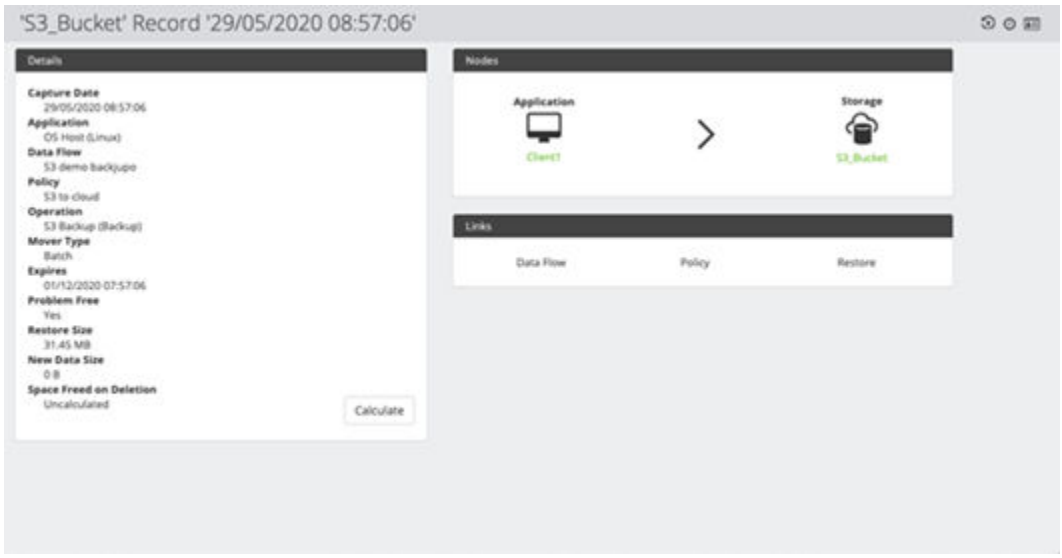


Figure 634 Cloud Recovery Point Details

Control	Description
Restore	Restores the snapshot. Opens the Restore Repository Snapshot Wizard - File System (on page 826) to guide you through the process..
Change Retention	Enabled only when one or more snapshots are selected. Opens the Change Repository Snapshot Retention Dialog (on page 920)
Transfer Ownership	Allows the ownership of the record to be changed. By changing the owning resource to a node that a user has access to enables the user to see and act upon the record.
Calculate	Amazon S3 only, Calculates the amount of space that would be freed in the repository if this recovery point was deleted.

Chapter 6: Command Line Interface Reference

This chapter describes the usage of command line tools installed with Ops Center Protector.

All commands are run from the `\bin` directory. By default these are located as follows.

For Windows:

```
C:\Program Files\Hitachi\Protector\bin
```

For Linux:

```
/opt/hitachi/protector/bin
```

We recommend that CLI based commands are run by experienced administrators only.

Note: Information about additional Ops Center Protector command line tools and their arguments is available from the Ops Center Protector Knowledge Base. These tools are only to be used with explicit guidance by customer support.

Gathering diagnostic information with `diagdata`

`diagdata` is a CLI tool that is used to aid diagnostic testing of Ops Center Protector. It provides the following functions:

- Stopping and starting of Ops Center Protector services
- Resetting the Ops Center Protector system state
- Raising and lowering of trace level for diagnostics
- Collection of Ops Center Protector state information to aid investigation of issues

Command line options

Table 51 General options

Abbreviated argument	Full argument	Description
`-h`	`--help`	Display help.

Table 52 Service options

Abbreviated argument	Full argument	Description
`-r` *arg*	`--restart` *arg*	Auto restart active services `yes | no`
`--start` *arg*		Starts the hub service. Options are `none`, `all`, `hub`. 'none' means no argument is required.
`--stop` *arg*		Stops the hub service. Options are `none`, `all`, `hub`. 'none' means no argument is required.

Table 53 Protector

Abbreviated argument	Full argument	Description
`-d` *arg*	`--trace-dir` *arg*	Set the trace directory.
`-T` *arg*	`--temp-dir` *arg*	Set the temp directory.
`--reset`		Resets the state of the Ops Center Protector node and flushes cache and logs. Can only be run if Ops Center Protector services are stopped (see `--stop`).

Table 54 Dump options

Abbreviated argument	Full argument	Description
`-f`	`--full`	Produce a full dump of the Ops Center Protector configuration on the system, including all trace files. Results are placed in a ZIP file named: `diag_`*`node name`*`_YYYYMMDDhhmm.zip`

Abbreviated argument	Full argument	Description
		in the directory: *Data-Instance-Director-install*\runtime This ZIP file can be requested by Hitachi Vantara support when investigating complex end user issues. Can be combined with `--spilt`
`-m`	`--mini`	Similar to `-f` but only saves a recent portion of all trace files. Can be combined with `--spilt`
`-c` *arg*	`--collect` *arg*	*arg* specifies which data must be collected. Use `all` to collect all components. For additional component names use `-c help`. Can be combined with `--spilt`
`-o` *arg*	`--output-dir` *arg*	*arg* specifies an optional output directory. Default is `Protector\runtime`
`-n` *arg*	`--tracename` *arg*	*arg* specifies a filter mask which uses wildcards to target only specific trace files e.g. `diagdata -f --tracename repos*`. Multiple filters can be used seperated by ';' e.g. `hub*;node*;repo*`. For Linux use `"hub*;node*;repo*"`).
`-a` *arg*	`--traceage` *arg*	*arg* specifies an age, in days, to target only trace files modified in that time frame.

Abbreviated argument	Full argument	Description
`-z`	`--nosnapshots`	Prevents `diagdata` from creating snapshots before a dump (this may result in some files not being added to the diagnostic dump).
`-Z`	`--withsnapshots`	Attempt diagdata snapshots regardless of configuration. The default mode of operation. For linux , snapshots will only be taken for LVM filesystems
`--no-md5`		Disable MD5 checksum on created zip files.

Table 55 Trace options

Abbreviated argument	Full argument	Description

`-t` *arg*	`--trace` *arg*	Sets the trace level of the node. Valid arguments are:
		`TraceAlways –`
		`TraceTesting`, `TraceNormal`, `TraceVerbose` `TraceDebug`
		Trace files are placed in the directory:
		`[Install-Dir]``\runtime\trace`

`-p`	`--deltrace`	Delete inactive trace files. Any trace files currently in use by the Protector system will not be deleted. Deletes the contents of *`[Install-Dir]`*`\runtime\trace`
`-x`	`--delalltrace`	Stop Protector services, delete all trace and then restart them.

Abbreviated argument	Full argument	Description
`-i`	`--info`	Get information about this node and current trace status.

Table 56 Windows crash dump file options

Abbreviated argument	Full argument	Description
`-l` *arg*	`--dumplevel` *arg*	Not available on Linux. Sets Protector crash dump level `full | default`.
`--test-dump`		Not available on Linux. Generates a Windows dump file for test purposes.

Table 57 Archive options

Abbreviated argument	Full argument	Description
`-S` *arg*	`--split-file` *arg*	Split an existing zip file (path specified by *arg*) into multiple chunks (configurable, 1GB default). Intended to be used if the file is too big to be uploaded to a file service.
`-C` *arg*	`--concatenate` *arg*	File name for the concatenated file. That must be followed by the list of the files to be concatenated. Example: --concatenate diag.zip part.000 part.001 part.002
`--split`		Split the zip in multiple chunks (configurable, 1GB default), this is used with the --full, --mini and --collect options. Intended to be used if the file is too big to be uploaded to a file service.

Usage Examples:

Stopping the Ops Center Protector services on any node:

```
diagdata --stop
```

Setting the trace level of the Ops Center Protector services on any node:

```
diagdata --trace TraceAlways
```

Collect a mini dianostic log from a node:

```
diagdata --mini
```

Collect a full diagnostic log from a node:

```
diagdata --full
```

Gathering diagnostic information with 'support manager cli'

Support manager cli is a powerful CLI tool that is based on **diagdata** and many of these commands should be familiar. The main advantage of the support manager CLI tool over **diagdata**, is that it is a remote tool and thus doesn't require specific access to any of the nodes vi remote desk top/ssh, which is always the case when running **diagdata** to control and collect information from nodes. Also, the support manager cli can run on the master as well, to collect data from that Master.

It provides the following functions from a command prompt on the Master node only and can also run commands that are targeted for the Master node:

- Remotely, stopping and starting of Ops Center Protector services.
- Remotely resetting the Ops Center Protector system state.
- Remotely raising and lowering of trace level for diagnostics.
- Remote collection of Ops Center Protector state information to aid investigation of issues.

It would be advantageous if the user also has experience of using **diagdata**.

This tool is invoked in the following way:

The executable is in `<Protector Installation>/bin`

1. Windows: From a cmd prompt or PowerShell prompt enter '`[./]supportmanagercli`'
2. Linux: From a terminal session `./supportmanagercli`

Like many of the CLI tools provided to support the Protector product, running the tool without any arguments, will print out the arguments/options that may be used. Many of these options, like **diagdata** can be concatenated together to create specific requirements.

Following the usage options below, there is a compressive set of notes that describes how the options can be used.

The table(s) below describe how each of the arguments would be used:

Abbreviated argument	Full argument	Description
	`--verbose`	More verbose output.
`-h`	`--help`	Produce this help message.

Table 58 Information options

Abbreviated argument	Full argument	Description
`-i`	`--info`	Get information about specified nodes.
	`--show-nodes`	Show list of specified nodes (or all nodes if none specified).

Table 59 Collection options

Abbreviated argument	Full argument	Description
`-m`	`--mini`	Perform mini dump on specified machines, and collect the resulting archives on the master.
`-f`	`--full`	Perform full dump on specified machines and collect the resulting archives on the master.
`-c <arg>`	`--collect <arg>`	Indicate which data must be collected from specified machines. Use 'all' to collect all components. For help use 'help'.
`-n <arg>`	`--tracename <arg>`	Optional filter mask, which uses wildcards to target only specific trace files e.g. --tracename repos* (multiple filters can be used separated by ';' e.g. hub*;node*;repo* for Linux use "hub*;node*;repo*")
`-a <arg>`	`--traceage <arg>`	Takes an integer which is the number of days to go back from today.

Abbreviated argument	Full argument	Description
`-z <arg>`	`--nosnapshots <arg>`	Optional filter mask, which uses age in days to target only trace files modified in that time frame.
	`--no-md5`	Disable MD5 checksum on created zip files
	`--split`	Split collected dump files into multiple chunks. Intended to be used if the file is too big to be uploaded to a file service. See the note 1 below
`-o <arg>`	`--output <arg>`	Optional output directory when trace is sent to the Master. The default value shown is correct for the installation location.

Note: The size of the 'chunks' used in the split option is configurable. This can be changed by editing the config file: `db/config/diagdata.cfg` and locating the following xml section on the associated node. The value is in byes with a default of 1GB – ie 1,073,741,824

```
<item name="File Part Chunk Size" argtype="single">
        <value type="uint32">1073741824</value>
</item>
```

Table 60 Management options

Abbreviated argument	Full argument	Description
`-t`***<arg>***	`--trace <arg>`	Set trace level on specified machines for new and running processes.
	`--trace-new <arg>`	Set trace level on specified machines for new processes only.
	`--trace-running <arg>`	Set trace level on specified machines for running processes only.

Abbreviated argument	Full argument	Description
`-p`	`--delalltrace`	Delete all trace files on specified machines. This option will stop the services, delete all trace, and start the services.

Table 61 Trace directory options

Abbreviated argument	Full argument	Description
`-d`*<arg>*	`--trace-dir`	Set the Protector trace directory on specified machine.

Table 62 Node selection options

Abbreviated argument	Full argument	Description
	`--all-nodes` *<arg>*	Specify that all nodes should be included.
	`--node-name` *<arg>*	Optional filter mask, which uses wildcards to include specific nodes by their name (multiple filters can be specified, separated by ',').
	`--node-type` *<arg>*	Optional filter mask, which uses wildcards to include specific nodes by their type (multiplefilters can be specified, separated by ',').
	`--test-node-filter`	Test a node filter only, without performing any operation on the nodes returned by that filter.

Table 63 Other shared options

Abbreviated argument	Full argument	Description
`-r<arg>`	`--restart`	Restart services on specified machines. Valid with management options and trace directory options.

How to build the options:

As mentioned above, this tool runs on the Master only, thus from a Windows or a Linux Master. It will run in a Windows CMD or Power Shell Window or in a Linux terminal window. It will not run from any other node type and it will not run from an AIX based node.

The 'binary' is located in the `<installation>/bin` area and is named 'supportmanagercli' (Note the binary name, as there are some other binary files with similar names in the same place)

Like many CLI tools, the options need to be 'strung' together. The following table(s), in no particular order will show how the options could be used:

Preceded any of the options with 'supportmanagercli' from a Windows CMD window and './ supportmanagercli' from Windows PowerShell Or Linux terminal window:

SYNOPSIS:

Table 64 supportmanagercli [OPTIONS] NodeSelection

Option	Description
`--show-nodes`	List all nodes as 3 lists 'OSHost Nodes involved', 'OSHost Nodes' and 'Agentless Nodes'
`--info --node-name <node name>`	Displays information about a specific node, nodes names using wild cards, a csv list (Note: no spaces between commas)
`--info --all-nodes`	Displays information about all nodes
`--info --node-type <node type>`	Displays information node types -ie OSHost, HyperV, HBBHCP, HitachiBlockDevice

`--test-node-filter`	This is used to test the filter:
	`--test-node-filter --node-name *<name>*`
	`--test-node-filter --node-type *<type>*`
	Or use just a wild card:
	`--test-node-filter --node-name *`

Option	Description
	`--test-node-filter --node-name *`
`--deltrace --node-name <node name>`	Delete inactive trace for a specific node
`--deltrace --all-nodes`	Delete inactive trace for all nodes
`--delalltrace --node-name <node name>`	Delete all trace files on specified node (Services will be restarted)
`--delalltrace --all-nodes`	Delete all trace files on all nodes (Services will be restarted)
`--trace-dir <dir> --node-name <node name>`	Set the trace directory on specified node (Requires the services to be restarted)
`--trace-dir <dir> --all-nodes`	Set all trace directories (Requires the services to be restarted)
`--trace <Trace Level> --node-name <node name>`	Set specified trace level on specific nodes (new and running processes)
`--trace <Trace Level> --all-nodes`	Set specified trace level on ALL nodes (new and running processes)
`--trace-new <Trace Level> --node-name <node name>`	Set specified trace level on specific nodes (new processes)
`--trace-new <Trace Level> --all-nodes`	Set specified trace level on ALL nodes (new processes)
`--trace-running <Trace Level> --node-name <node name>`	Set specified trace level on specific nodes (running processes)
`--trace-running <Trace Level> --all-nodes`	Set specified trace level on ALL nodes (running processes)
`--full --node-name <node name>`	Collect data from the node name (can use CSV - no space) and put it into the Master's <protector home>/runtime/remote trace collection/<time stamp directory>
`--full --all-nodes`	Collect data from all nodes and put it into the Master's `<protector home>/runtime/remote trace collection/<time stamp directory>`
`--mini --node-name <node name>`	Collect data from the node name (can use CSV - no space) and put it into the Master's `<protector home>/runtime/remote trace collection/<time stamp directory>`

Option	Description
`--mini --all-nodes`	Collect data from all nodes and put it into the Master's `<protector home>/runtime/remote trace collection/<time stamp directory>`
`--full --node-name <node name> --output-dir <dir>`	Specify a Master collection directory (<dir> will be created on the Master)
`--full --all-nodes --output-dir <dir>`	Specify a Master collection directory (<dir> will be created on the Master)
`--mini --node-name <node name> -output-dir <dir>`	Specify a Master collection directory (<dir> will be created on the Master)
`--mini --all-nodes -output-dir <dir>`	Specify a Master collection directory (<dir> will be created on the Master)
`--full --node-name <node name> --split --no-md5`	Collect data from the node name (can use CSV - no space) and put it into the Master's `<protector home>/runtime/remote trace collection/<time stamp directory>` and split the zip into chunks. Don't compute the MD5
`--full --all-nodes -split --no-md5`	Collect data from all nodes and put it into the Master's `<protector home>/runtime/remote trace collection/<time stamp directory>` and split the zip into chunks. Don't compute the MD5
`--mini --node-name <node name> --split --no-md5`	Collect data from the node name (can use CSV - no space) and put it into the Master's <`protector home>/runtime/remote trace collection/<time stamp directory`> and split the zip into chunks. Don't compute the MD5
`--mini --all-nodes -split --no-md5`	Collect data from all nodes and put it into the Master's `<protector home>/runtime/remote trace collection/<time stamp directory>` and split the zip into chunks. Don't compute the MD5

Option	Description
`--full --node-name <node name> --tracename`	Collect data from the node name (can use CSV - no space) and put it into the Master's `<protector home>/runtime/remote trace collection/<time stamp directory>` --tracename will get specific trace files and also works with wild cards. Case insensitive.
`--full --node-name <node name> --traceage`	Collect data from the node name (can use CSV - no space) and put it into the Master's <`protector home>/runtime/remote trace collection/<time stamp directory`> --traceage will get only trace up to n days old - ie --traceage 3 will collect that is between 0 days to 3 days old.
`--full --node-name <node name> --nosnapshots`	Collect data from the node name (can use CSV - no space) and put it into the Master's `<protector home>/runtime/remote trace collection/<time stamp directory`> --nosnapshots will disable snapshots when to collect data from nodes.
`--restart --node-name <node name>`	Send to restart nodes - Note using option '--all-nodes' will also restart the Master as well.

Finally, the --verbose option can be appended to any of the above. It will obviously create a lot of output. To work with extra output a redirect can be used to send all the data into a file. Also, the data can be filtered as well with the following windows or Linux commands:

To send data to a file, use the redirect '>':

```
customersupportcli
       --info --all-nodes > <file name>
```

To filter out a pattern, use the:

- Windows: `|findstr <pattern>`
  - `customersupportcli --info --all-nodes |findstr <pattern>`
- Linux: `|grep <pattern>`
  - `customersupportcli --info --all-nodes |grep <pattern>`
- PowerShell: `|select-string <pattern>`
  - `customersupportcli --info --all-nodes | select-string <pattern>`

Note: Always quote any arguments which have spaces, for example,

If setting a new trace directory on Windows to `C:\Program Files\Hitachi\trace` the command `supportmanagercli --trace-dir C:\Program Files\Hitachi\trace` will set the trace directory to `C:\Program`

Show all nodes:

supportmanagercli --show-nodes

OSHost nodes involved:

```
Docs-Client8

Docs-Client4

Docs-Client6

Docs-Proxy1

Docs-Client5

Docs-Proxy4

Docs-Client3

Docs-Proxy5

Docs-Proxy2

Docs-Client2

Docs-Proxy3

Docs-Master

Docs-Client1
```

OSHost Nodes:

```
Docs-Client2 (Type: Server, Virtual, Linux; Status: Authorized, Up;
Version: 7.4.0.92768-R7.4)

Docs-Client1 (Type: Server, Virtual, Linux; Status: Authorized, Up;
Version: 7.4.0.92768-R7.4)

Docs-Proxy3 (Type: Server, Virtual, Linux; Status: Authorized, Up;
Version: 7.4.0.92768-R7.4)

Docs-Client3 (Type: Server, Virtual, Linux; Status: Authorized, Up;
Version: 7.4.0.92768-R7.4)

Docs-Client4 (Type: Server, Virtual, Windows; Status: Authorized,
Up; Version: 7.4.0.92768-R7.4)

Docs-Proxy1 (Type: Server, Virtual, Windows; Status: Authorized, Up;
Version: 7.4.0.92768-R7.4)

Docs-Client5 (Type: Server, Virtual, Windows; Status: Authorized,
Up; Version: 7.4.0.92768-R7.4)

Docs-Client8 (Type: Server, Virtual, Windows; Status: Authorized,
Up; Version: 7.4.0.92768-R7.4)
```

```
Docs-Proxy5 (Type: Server, Virtual, Windows; Status: Authorized, Up;
Version: 7.4.0.92768-R7.4)

Docs-Client6 (Type: Server, Server/Workstation, AIX; Status:
Authorized, Up; Version: 7.4.0.92753-R7.4)

Docs-Proxy4 (Type: Server, Virtual, Linux; Status: Authorized, Up;
Version: 7.4.0.92768-R7.4)

Docs-Master (Type: Server, Virtual, Windows; Status: Authorized, Up,
Master; Version: 7.4.0.92768-R7.4)

Docs-Proxy2 (Type: Server, Virtual, Linux; Status: Authorized, Up;
Version: 7.4.0.92768-R7.4)
```

Agentless Nodes:

```
Chesil (Type: HitachiBlockDevice; Status: Authorized, Up; Proxy/Host
Node(s): Docs-Proxy5)

Docs-Gen1Repo1 (Type: Repository; Status: Authorized, Up; Proxy/Host
Node(s): Docs-Proxy1)

Gen1OnDoc-client8 (Type: Repository; Status: Authorized, Up; Proxy/
Host Node(s): Docs-Client8)

MyMasterRepo2 (Type: HBBRepository; Status: Authorized, Up; Proxy/
Host Node(s): Docs-Master)

Conker (Type: HitachiBlockDevice; Status: Authorized, Up; Proxy/Host
Node(s): Docs-Proxy4)

Docs-CloudScale1 (Type: HBBHCPCloudScale; Status: Authorized, Down;
Proxy/Host Node(s): Docs-Proxy5)

MyMasterRepo (Type: HBBRepository; Status: Authorized, Up; Proxy/
Host Node(s): Docs-Master)

Docs-Gen2Repo1 (Type: HBBRepository; Status: Authorized, Up; Proxy/
Host Node(s): Docs-Proxy5)

Docs-HCP1 (Type: HBBHCP; Status: Authorized, Up; Proxy/Host Node(s):
Docs-Proxy5)

Docs-HyperV1 (Type: HyperV; Status: Authorized, Up; Proxy/Host
Node(s): Docs-Client5, Docs-Client5)

Docs-AWS1 (Type: HBBS3; Status: Authorized, Down; Proxy/Host
Node(s): Docs-Proxy5)

Docs-SQL1 (Type: MSSQLSRV; Status: Authorized, Up; Proxy/Host
Node(s): Docs-Client4, Docs-Client4)

Docs-VMware1 (Type: VMware; Status: Authorized, Up; Proxy/Host
Node(s): Docs-Proxy1)
```

Retrieve the trace level from a single node:

This uses the windows 'findstr' to match the words "Hub Trace Level"

supportmanagercli --info --node-name Docs-Client4|findstr /c:"Hub Trace Level:"

[Docs-Client4]: Hub Trace Level: TraceAlways

Retrieve the trace level from a number of nodes: (use CSV no spaces)

supportmanagercli --info --node-name Docs-Client4,Docs-Proxy4 |findstr /c:"Hub Trace Level:"

[Docs-Proxy4]: Hub Trace Level: TraceAlways

[Docs-Client4]: Hub Trace Level: TraceAlways

Retrieve the trace level from all nodes:

supportmanagercli --info --all-nodes |findstr /c:"Hub Trace Level:"

[Docs-Proxy3]: Hub Trace Level: TraceAlways

[Docs-Client5]: Hub Trace Level: TraceAlways

[Docs-Proxy2]: Hub Trace Level: TraceAlways

[Docs-Client2]: Hub Trace Level: TraceAlways

[Docs-Client1]: Hub Trace Level: TraceAlways

[Docs-Master]: Hub Trace Level: TraceAlways

[Docs-Client6]: Hub Trace Level: TraceAlways

[Docs-Proxy5]: Hub Trace Level: TraceAlways

[Docs-Client3]: Hub Trace Level: TraceAlways

[Docs-Client8]: Hub Trace Level: TraceAlways

[Docs-Proxy1]: Hub Trace Level: TraceAlways

[Docs-Proxy4]: Hub Trace Level: TraceAlways

[Docs-Client4]: Hub Trace Level: TraceAlways

Set trace on a single node to TraceDebug

supportmanagercli --trace TraceDebug --node-name Docs-Proxy1

OSHost nodes involved:

Docs-Proxy1

[Docs-Proxy1]: Trace level for new processes set to TraceDebug

[Docs-Proxy1]: Trace level for running processes set to TraceDebug

[Docs-Proxy1]: Operation completed.

Operation completed on nodes:

Docs-Proxy1

Exiting...

Set trace on all nodes to TraceDebug

supportmanagercli --trace TraceDebug --all-nodes

OSHost nodes involved:

```
Docs-Proxy2
Docs-Client3
Docs-Proxy3
Docs-Master
Docs-Client8
Docs-Proxy4
Docs-Client1
Docs-Proxy5
Docs-Client2
Docs-Client5
Docs-Client6
Docs-Client4
Docs-Proxy1
```

[Docs-Proxy2]: Trace level for new processes set to TraceDebug

[Docs-Proxy2]: Trace level for running processes set to TraceDebug

[Docs-Proxy2]: Operation completed.

[Docs-Client3]: Trace level for new processes set to TraceDebug

[Docs-Client3]: Trace level for running processes set to TraceDebug

[Docs-Client3]: Operation completed.

[Docs-Proxy3]: Trace level for new processes set to TraceDebug

[Docs-Proxy3]: Trace level for running processes set to TraceDebug

[Docs-Proxy3]: Operation completed.

[Docs-Master]: Trace level for new processes set to TraceDebug

[Docs-Master]: Trace level for running processes set to TraceDebug

[Docs-Master]: Operation completed.

[Docs-Client8]: Trace level for new processes set to TraceDebug

[Docs-Client8]: Trace level for running processes set to TraceDebug

[Docs-Client8]: Operation completed.

[Docs-Proxy4]: Trace level for new processes set to TraceDebug

[Docs-Proxy4]: Trace level for running processes set to TraceDebug

[Docs-Proxy4]: Operation completed.

[Docs-Client1]: Trace level for new processes set to TraceDebug

[Docs-Client1]: Trace level for running processes set to TraceDebug

[Docs-Client1]: Operation completed.

[Docs-Proxy5]: Trace level for new processes set to TraceDebug

[Docs-Proxy5]: Trace level for running processes set to TraceDebug

[Docs-Proxy5]: Operation completed.

[Docs-Client2]: Trace level for new processes set to TraceDebug

[Docs-Client2]: Trace level for running processes set to TraceDebug

[Docs-Client2]: Operation completed.

[Docs-Client5]: Trace level for new processes set to TraceDebug

[Docs-Client5]: Trace level for running processes set to TraceDebug

[Docs-Client5]: Operation completed.

[Docs-Client6]: Trace level for new processes set to TraceDebug

[Docs-Client6]: Trace level for running processes set to TraceDebug

[Docs-Client6]: Operation completed.

[Docs-Client4]: Trace level for new processes set to TraceDebug

[Docs-Client4]: Trace level for running processes set to TraceDebug

[Docs-Client4]: Operation completed.

[Docs-Proxy1]: Trace level for new processes set to TraceDebug

[Docs-Proxy1]: Trace level for running processes set to TraceDebug

[Docs-Proxy1]: Operation completed.

Operation completed on nodes:

Docs-Master

Docs-Client8

Docs-Client5

Docs-Proxy3

Docs-Proxy1

Docs-Proxy2

Docs-Client2

Docs-Client4

Docs-Client3

Docs-Client1

Docs-Proxy4

Docs-Client6

Docs-Proxy5

Exiting...

And, to verify that all nodes are now set to TraceDebug:

supportmanagercli --info --all-nodes |findstr /c:"Hub Trace Level:"

[Docs-Client2]: Hub Trace Level: TraceDebug

[Docs-Client6]: Hub Trace Level: TraceDebug

[Docs-Client1]: Hub Trace Level: TraceDebug

[Docs-Client3]: Hub Trace Level: TraceDebug

[Docs-Proxy3]: Hub Trace Level: TraceDebug

[Docs-Client5]: Hub Trace Level: TraceDebug

[Docs-Proxy4]: Hub Trace Level: TraceDebug

[Docs-Proxy5]: Hub Trace Level: TraceDebug

[Docs-Proxy1]: Hub Trace Level: TraceDebug

[Docs-Client8]: Hub Trace Level: TraceDebug

[Docs-Proxy2]: Hub Trace Level: TraceDebug

[Docs-Client4]: Hub Trace Level: TraceDebug

[Docs-Master]: Hub Trace Level: TraceDebug

Triggering policies and operations with `hdidcmd`

Usage Example:

hdidcmd allows REST commands to be sent to the master. This topic demonstrates how to trigger a policy operation via the REST interface using this command line tool. Access to the REST API is controlled by RBAC, therefore the user:

1. Open a command prompt.
2. Login to Protector using the following command to begin a session:

   hdidcmd --activity "Login"

   --space *Authentication SpaceName*

   --userName *UserName*

   --password *Password*

   --ip *IpAddress*

   --port <*Port*>

Where:

- *AuthenticationSpaceName* is the *Authentication Space* where **Username** can be authenticated.
- *UserName* is the user name for the session. The specified *UserName* must have the *Trigger Operations* activity assigned to them and have visibility of the nodes, data flows and policies specified in the trigger command.
- *Password* is the password for *UserName*
- *IpAddress* is the IP address of the Master node
- *Port* is the port used to connect to the master server. This value is optional



Note: Sessions are maintained on a 'per command prompt' basis. If no activity is detected after 2 hours then the session will be terminated.

3. Trigger the required operation by specifying the *data flow*, *source node*, *policy* and *operation* using the following command:

hdidcmd --activity "Trigger Operation"

[--dataflowName *dataflowName* **| --dataflowId** *dataflowId***]**

[--sourceNodeName *sourceNodeName* **| --sourecNodeId** *sourceNodeId***]**

[--policyName *PolicyName* **| --policyId** *policyId***]**

[--operationName *operationName* **| --operationId** *operationId***]**

[--syncGroup *syncGroup* **]**

[--wait]

The following names (found in the UI) can be specified:

- *dataflowName* is the textual name of the data flow where *sourceNodeName* appears
- *sourceNodeName* is the textual name of the source node where the operation is assigned
- *PolicyName* is the textual name of the policy
- *operationName* is the textual name of the operation within *PolicyName*
- *synchGroup* is the textual name of the shedule that is used to synchronise a group of hardware based operations

Alternatively the following IDs can be specified:

- *dataflowId*
- *sourceNodeId*
- *policyId*
- *operationId*

These are determined using the following command:

hdidcmd --activity "View Triggers"

The **--wait** option causes the command to wait until all the jobs are either *completed*, *failed* or *paused*; i.e. in some state other than *in progress*.

4. Logout from Protector using the following command to end the session:

hdidcmd --activity "Logout"

```
hdidcmd --activity "Login" --space Master --userName "administrator" --password
3253GHrT3 --ip 192.168.45.149

Operation Successful!

hdidcmd --activity "View Triggers"
{
    "trigger": [
        {
            "id": "0",
            "policyId": "2deb212ef02445458a1da13230bcf8dc",
            "policyName": "myPathBackupPolicy",
            "operationId": 1,
            "operationName": "Backup",
            "dataflowName": "myOSHostToRepoDataFlow",
            "dataflowId": "7ef4f80e-81da-456d-b44b-07789f80b522",
            "sourceNodeId": "Client1@B2-289YQR-37XVXP-I2QLXH-B4IB2E[0-1-2]",
            "sourceNodeInstance": 19,
            "destinationNodeId": "myRepository@00-B36A96-A2E2B8-437AB6-0BEF11[0-1-
10]",
            "destinationNodeInstance": 20,
            "moverType": "eMOVER_BATCH",
            "syncGroup": "",
            "dataOriginNodeInstance": 19,
            "dataOriginNodeId": "Client1@B2-289YQR-37XVXP-I2QLXH-B4IB2E[0-1-2]"
        }
    ],
    "pageInfo": {
        "totalCount": 1,
        "end": true
    }
}
Operation Successful!

hdidcmd --activity "Trigger Operation" --dataflowName "myOSHostToRepoDataFlow" --
source
NodeName "Client1" --policyId "2deb212ef02445458a1da13230bcf8dc" --operationId 1
{
    "job": [
        {
            "id": "{6e127915-7553-464b-8511-0606b4bec3a4}"
        }
    ]
}
Operation Successful!

hdidcmd --activity "Logout"
```

```
Operation Successful!
```

Archiving master node setting with `mastersettings`

The master node configuration files can be zipped to an archive file using the **`mastersettings`** command. During the procedure, only the master node settings are zipped up; other information such as repositories and their metadata stores and log messages are excluded.

We recommend that you back up the master node configuration files once a week. You can set up a backup policy on a recurring schedule that runs the **`mastersettings`** CLI tool as a pre-script, then backs up the resulting zip file.

Note: A *host based front end capacity license* is required to perform backup of the master node using **`mastersettings`**.

Table 65 Command line options

Abbreviated Argument	Full Argument	Description
`-h`	`--help`	Displays command line parameter usage instructions.
`-e`	`--export` *`zipfile`*	Zips up all configuration files to `zipfile`. If no path is specified, the file will be created in the `$Protector_HOME\runtime` directory. If no `.zip` suffix is specified, then it will be appended. The file is created in zip format and can be read with commonly available tools such as WinZip.
`-c`	`--config` *`cfgfile`*	Used with `--export` and `--import`. Specifies a configuration file used to control what files are included and excluded in the archive. By default, the `cfgfile` name `mastersettings.cfg` is assumed. (See About the mastersettings.cfg file (on page 958).)

Abbreviated Argument	Full Argument	Description
`-i`	`--import` *zipfile*	Restores all master node configuration files from *zipfile* The CofioHub service must be stopped prior to importing.
`-p`	`--preclear True|False`	Used when importing a zip file. Indicates whether existing master node configuration files are to be deleted prior to restoring them from the archive.

Usage examples:

Archiving the master node configuration files:

```
mastersettings -e mastercfgfiles.zip
```

Configuring of mastersettings:

`mastersettings.cfg` is an XML file that contains information indicating which Ops Center Protector files are to be:

- included or ignored when creating the archive (exporting)
- deleted or ignored when restoring the archive (importing)

This file is located in the following directory:

```
$Protector_HOME$\db\config
```

The `mastersettings.cfg` file can contain information as described in the following table.

Table 66 Configuration file options:

Item Name	Description
`Path`	The path to be backed up. The path specified is relative to the `$Protector_HOME` directory.
`Depth`	The depth, defined as a value, within the path to be backed up: -1 (default) backs up all files in all subdirectories. 1 backs up only the files in the path. 2 backs up all files in the path and within subdirectories of the path (but not the subdirectories themselves).

Item Name	Description
`Export Mask`	Controls which files in the path are backed up. Can contain the following entries: (If no export mask is specified then all files will be matched and exported).
	`NameMask` A file name mask such as `*.cfg`.
	`MatchAction` Indicates how to process matching files: Either `CopyFile` or `IgnoreFile`
	NonMatchAction Indicates how to process non-matching files: Either `CopyFile` or `IgnoreFile`
`Import Mask`	Controls what is cleared from the path prior to a restore. Can contain the following entries: (If no import mask is specified then all files will be matched and deleted prior to the restore unless the option `--preclear false` is specified at the command line).
	`NameMask` A file name mask such as `*.cfg`.
	`MatchAction` Indicates how to process matching files: Either `DeleteFile` or `IgnoreFile`
	`NonMatchAction` Indicates how to process non-matching files: Either `DeleteFile` or `IgnoreFile`

Configuration file examples:

The following XML tells **`mastersettings.exe`** to export all files in and below the main Ops Center Protector configuration directory:

```
<item name="Entry" argtype="list" >
    <item name="Path" argtype="single" >
        <value type="string" >db/config</value>
    </item>
    <item name="Depth" argtype="single" >
        <value type="int32" >2</value>
    </item>
</item>
```

The following XML tells mastersettings.exe not to export any files from the media manager events directory, but to delete all events from that directory prior to a restore. This will ensure that all events that are logged prior to an import are deleted, because they will not be relevant after the master settings have been restored.

```
<item name="Entry" argtype="list" >
    <item name="Path" argtype="single" >
         <value type="string" >db/mediamanager/events</value>
    </item>
    <item name="Export Mask" argtype="list" >
        <item name="MatchAction" argtype="single" >
            <value type="string" >IgnoreFile</value>
        </item>
    </item>
    <item name="Import Mask" argtype="list" >
        <item name="MatchAction" argtype="single" >
            <value type="string" >DeleteFile</value>
        </item>
    </item>
</item>
```

About the mastersettings.cfg file

`mastersettings.cfg` is an XML file that contains information indicating which Ops Center Protector files are to be zipped up (exported) or deleted prior to a restore (imported). This file is located in the following directory:

`$Protector_HOME$\db\config`

The `mastersettings.cfg` file can contain information as described in the following table.

Table 67 Configuration settings

Item Name	Description
`Path`	The path to be backed up. The path specified is relative to the `$Protector_HOME` directory.
`Depth`	The depth, defined as a value, within the path to be backed up: -1 (default) backs up all files in all subdirectories. 1 backs up only the files in the path. 2 backs up all files in the path and within subdirectories of the path (but not the subdirectories themselves).
`Export Mask`	Controls which files in the path are backed up. Can contain the following entries: (If no export mask is specified then all files will be matched and exported).

Item Name	Description
	`NameMask` A file name mask such as `*.cfg`.
	`MatchAction` Indicates how to process matching files: Either `CopyFile` or `IgnoreFile`
	NonMatchAction Indicates how to process non-matching files: Either `CopyFile` or `IgnoreFile`
`Import Mask`	Controls what is cleared from the path prior to a restore. Can contain the following entries: (If no import mask is specified then all files will be matched and deleted prior to the restore unless the option `--preclear false` is specified at the command line).
	`NameMask` A file name mask such as `*.cfg`.
	`MatchAction` Indicates how to process matching files: Either `DeleteFile` or `IgnoreFile`
	`NonMatchAction` Indicates how to process non-matching files: Either `DeleteFile` or `IgnoreFile`

Examples

The following XML tells **`mastersettings`** to export all files in and below the main Ops Center Protector configuration directory:

```
<item name="Entry" argtype="list" >
    <item name="Path" argtype="single" >
        <value type="string" >db/config</value>
    </item>
    <item name="Depth" argtype="single" >
        <value type="int32" >2</value>
    </item>
</item>
```

The following XML tells **`mastersettings`** not to export any files from the intelligentstoragemanager events directory, but to delete all events from that directory prior to a restore. This will ensure that all events that are logged prior to an import are deleted, because they will not be relevant after the master settings have been restored.

```
<item name="Entry" argtype="list" >
    <item name="Path" argtype="single" >
```

```
         <value type="string" >db/intelligentstoragemanager/events</
value>
    </item>
    <item name="Export Mask" argtype="list" >
        <item name="MatchAction" argtype="single" >
            <value type="string" >IgnoreFile</value>
        </item>
    </item>
    <item name="Import Mask" argtype="list" >
        <item name="MatchAction" argtype="single" >
            <value type="string" >DeleteFile</value>
        </item>
    </item>
</item>
```

Reducing repository size

If there is a large amount of free space within a repository, due to removal of a large backup set, it is possible to reduce the overall size of the repository. However this should only be done if the free space is required for another purpose. The repository is designed to re-use available internal space before growing in size to accommodate more data. If it is necessary to free space within a repository raise a support ticket for assistance with the process.

Generating log messages with `sendlog`

Use the **`sendlog`** command to create Protector log entries from a script.

Table 68 Command line options

Abbreviated Argument	Full Argument	Description
`-h`	`--help`	Displays command line parameter usage instructions.
`-m` *arg*	`--message` *arg*	*Log Message* field in the Logs. Specify a string in quotes.
`-n` *arg*	`--node` *arg*	*Actioned By* field in the Logs. Specify a string in quotes.
`-l` *arg*	`--level` *arg*	*Level* field in the Logs. Specify one of the following strings: ▪ `detail` ▪ `information` ▪ `warning`

Abbreviated Argument	Full Argument	Description
		▪ `error` ▪ `reporting`
`-c` *arg*	`--category` *arg*	*Category* field in the Logs. Specify a string in quotes.
`-t` *arg*	`--tag` ***arg***	Log tag in the format: *key*`:String | NodeID | uint32 | int32| int64 | uint64 | ISODate | Bool:`*value*
`-a` *arg*	`--attachment` *arg*	*Attachments* field in the Logs. Specify a filename to attach.
`-u` *arg*	`--user` *arg*	*User* name field in the View Audit bubble. For use by Hitachi Vantara Support only.
`-z` *arg*	`--process` *arg*	Audit process. For use by Hitachi Vantara Support only.
`-d` *arg*	`--space` *arg*	*User* name (authentication space) field in the View Audit bubble. For use by Hitachi Vantara Support only.
`-i` *arg*	`--ipaddress` *arg*	*Address* field in the View Audit bubble. Specify an IP address. For use by Hitachi Vantara Support only.
`-p` *arg*	`--port` *arg*	*Address* field (port) in the View Audit bubble. Specify an IP port. For use by Hitachi Vantara Support only.
`-o` *arg*	`--action` *arg*	*Action* field in the View Audit bubble. Specify a string. For use by Hitachi Vantara Support only.
`-s` *arg*	`--success` *arg*	*Action Result* field in the View Audit bubble. For use by Hitachi Vantara Support only.

Usage Example:

Running the following command will create a log message in the category *Test*, at the level *Information* with the message *"Test log message"*:

```
sendlog -c "Test" -l information -m "Test log message"
```

Changing a node's profile with `setconfig`

If you need to change a client machine's Master node, you can do so by running the **`setconfig`** utility from the command line interface.

The following is information about the most commonly used options and command usage.

Table 69 Command line options

Abbreviated Argument	Full Argument	Description
`-h`	`--help`	Displays command line parameter usage instructions.
`-l`	`--list`	Lists the node's address settings.
`-c`	`--clear`	Clear the client node's master setting.
`-m` *address*	`--master` *address*	Set the IP address of the client node's master. This option is used with `-f`
`-f`	`--force`	Forcibly set the client node's master. This option is used with `-m` and `-c` if the node in question is already connected to an authorized master.
`-e` *name ipaddress*	`--resetmaster` *name ipaddress*	Reset name and IP address of master node e.g.: `--resetmaster myMaster 192.168.12.23`
`-n`	`--name` *name*	Set node name (excluding ID part)
`--full`		This option is used with `-n`. The name given is the full node name including ID
`-p`	`--prefer` *IPaddress*	Add preferred network(s)

Abbreviated Argument	Full Argument	Description
`-b`	`--bar` *`IPaddress`*	Bar network(s)
`-x`	`--fixed` *`IPaddress`*	Add fixed address(s). Sets the external-facing address of the Master when backing up over the Internet. This is the IP address that will be entered when installing Internet connected nodes.
`-r`	`--remove` *`IPaddress`*	Remove addresses from fixed/preferred/barred list(s)
`--firewalltrue`		Node is the other side of a firewall from the master. Provides the same behavior as selecting Internet connected node during the software installation.
`--deletesslcert` *`node-name`*		Delete the SSL certificate for the local node/specified node.
`--encrypted`		All data transfer to be encrypted
`--filterruleschеck`		Check the filter rules for the file system are up-to-date and triggers a resync if they needed updating.

Usage example:

Setting the Master node IP address:

```
setconfig --master 192.156.34.0
```

Installing and upgrading Protector from the command line

Use the **`-R<release_num>-<build_num>-<os>-<architecture>`** command to install or upgrade Protector from the command line. Notice that a separate installer is supplied on the ISO for each supported OS and architecture type.

Table 70 Command Line Options

Argument	Description
`--help`	Displays the list of valid options in a pop-up dialog.

Argument	Description
`--install-directory` *arg*	Specifies the installation directory. Windows Default is `C:\Program Files\Hitachi\Protector` Linux Default is `/opt/hitachi/protector`
`--node-type` *arg*	Specifies the installation type: ▪ `master` ▪ `client` (default)
`--mode` *arg*	Sets the installation mode: Windows options: ▪ `win32` - interactive wizard (default) ▪ `unattended` - no user interaction Linux options ▪ `gtk` - GUI using GTK libraries (default) ▪ `xwindow` - GUI using basic X11 ▪ `text` - Interactive text ▪ `unattended` - Non-interactive
`--unattendedmodeui` *arg*	Sets the UI displayed for unattended installation: ▪ `none` (default) ▪ `minimal` ▪ `minimalWithDialogs`
`--debugtrace` *arg*	Specifies a filename for debug trace output.
`--debuglevel` *arg*	Sets the level of verbosity for debug trace output: ▪ `0` - lowest ▪ `1` ▪ `2` (default) ▪ `3` ▪ `4` - highest
`--node-name` *arg*	Specifies the Protector node name. Default is the OS node name.

Argument	Description
`--user-account` *arg*	Used only when installing a master node. Specifies the account to use for initial login to Protector after installation of a master node.
`--ui-port` *arg*	Used only when installing a master node. Specifies the user interface port. Default is 443 (HTTPS).
`--master-name` *arg*	Used only when installing client nodes. Specifies the node name or IP address of the master. Clients use to contact the master.
`--internet-connected` *arg*	Used only when installing client nodes. Installs the node as an internet connected node: ▪ `0` - not internet connected (default) ▪ `1` - internet connected **Note:** This node will connect to the master over an unsecured network. Choosing this option will cause all data sent and received to be encrypted and will also enable this node to function behind a firewall. Over-the-wire encryption requires a license and may not be available in all territories.
`--installer-language` *arg*	Specifies the installation language: ▪ `en` - English (default) ▪ `ja` - Japanese ▪ `zh_CN` - Simplified Chinese
`--optionfile` *arg*	Specifies an installation option file. Instructs the installer to take the command line values from a file when `--mode unattended` is specified. The options are listed. line by line, as `key=value` pairs.
`--version`	Displays product version information in a pop-up dialog.

Running the following command will perform an unattended installation or upgrade of a Protector version 6.8 AIX PowerPC client node:

HDID-R6.8-6.8.0.59802-AIX-PPC-64.run --mode unattended --install-directory \opt\hitachi\hdid --node-type client --node-name myClient --master-name myMaster

Listing Oracle RMAN channel configurations with schedulershow

schedulershow is a CLI tool that is used to aid the Oracle database administrator to configure Oracle RMAN to save and restore data using a datastore managed by Ops Center Protector. It provides the following functionality.

- List the Oracle RMAN dataflows active on the current node
- Create a sample RMAN channel definition for a dataflow operation

In order for a node to store or access any data in a Protector managed datastore using Oracle RMAN, the node has to be part of a dataflow and granted access using the access operation. Once the dataflow is compiled and distributed, it can be listed with schedulershow and Oracle RMAN can access it using an SBT channel.

Table 71 Command line options

Abbreviated Argument	Full Argument	Description
`-h`	`--help`	Display help.
`-c`	`---config <trigger>`	Display RMAN channel configuration for a dataflow. The command provides a list of valid triggers when running the it without parameters..

Usage Example

List all Oracle RMAN dataflows active on this node

```
app/bin/schedulershow

Application DataFlow Source OperationName Storage Trigger

Oracle mysampledf1 myOracleNode Access myRepository
1f04eeec33b10190b79ea6d47a65f3943f5a9de0c1a5d1557afeb420ea604046

Oracle Demo2 myOracleNode Access myRepository
6e7b941d5e73ec5a3ce72b64e70c36c56e6d12f76a379d0982d130d1b505a895
```

Get the Oracle RMAN channel definition for a dataflow

```
app/bin/schedulershow -c
1f04eeec33b10190b79ea6d47a65f3943f5a9de0c1a5d1557afeb420ea604046

# Application = Oracle

# DataFlow=mysampledf1

# Source=myOracleNode

# OperationName=Access

# Storage=myRepository
```

```
#
Trigger=1f04eeec33b10190b79ea6d47a65f3943f5a9de0c1a5d1557afeb420ea60
4046

DEVICE TYPE 'SBT_TAPE' PARMS

'SBT_LIBRARY=/opt/hitachi/protector/app/lib/
libhsbt.so,SBT_PARMS=(TRIGGER=1f04eeec33b10190b79ea6d47a65f3943f5a9d
e0c1a5d1557afeb420ea604046)'
```

Query Storage Command Line Utility

This command line utility, included with Protector, allows users to query available consistency groups when replicating to software-defined storage. It differs from other Protector CLI tools in that it is a Python based tool and therefore has some additional requirements.

Prerequisites

- **`Operating System`**: Linux
- **`Python`**: Version 3.11 or newer
- **`pip`**: Python package installer
- **`CCI`**: Version 01-76-03/03 or newer
- **`FC Command Device`**: Recommended for better performance

Python pip usually ships with Python. If it is not available, it can be installed using:

```
python3 -m ensurepip –upgrade
```

While the tool might work on Windows, it is only supported on Linux as the SDS arrays support Linux exclusively. Installing Python on an existing Linux ISM is usually the easiest approach.

Installation

The tool is shipped with Protector as a Python package, which must be installed before use. To install the tool, run the following command:

```
python3 -m pip install --upgrade
      /opt/hitachi/protector/resources/cli/ocp_query_storage*.whl
```

Usage

Once installed, the ocp_query_storage command can be used to query the list of available consistency groups.

To list which consistency groups are available for new replications between a physical array and an SDS array, use the following command:

```
ocp-query-storage ctgs --storage-host <controller_ip> --storage-user <storage_user> --
sds-host
      <sds_hostname> --sds-username <sds_user>
```

The tool is self-documenting. For help and information about other command line options, use the --help `option`.

Note

Determining unused CTG IDs is a slow process and can take multiple minutes. We strongly recommend using Fibre Channel command devices, as the process can take more than 30 minutes when using IP command devices.

Commands used to generate GRR

The following commands are used to generate the data displayed in the GRR

Open Session `/v1/objects/storages/{storage_device_id}/sessions`

Close Session `/v1/objects/storages/{storage_device_id}/sessions/{session_id}`

Refresh Cache `/services/storage-cache-service/actions/refresh/invoke`

Get Host Group List `/v1/objects/storages/{storage_device_id}/host-groups`

Get Logical Devices List `/v1/objects/storages/{storage_device_id}/ldevs`

Get Local Clone Copy Group List `/v1/objects/storages/{storage_device_id}/local-clone-copygroups`

Get Local Replications List `/v1/objects/storages/{storage_device_id}/local-replications`

Get Ports List `/v1/objects/storages/{storage_device_id}/ports`

Get Remote Mirror Copy Group List `/v1/objects/storages/{storage_device_id}/remote-mirror-copygroups`

Get Remote Replications List `/v1/objects/storages/{storage_device_id}/remote-replications`

Get Remote Storages List `/v1/objects/storages/{storage_device_id}/remote-storages`

Get Snapshots List `/v1/objects/storages/{storage_device_id}/snapshot-replication`

Get Storage System `/v1/objects/storages/{storage_device_id}`

Get Storage Systems List `/v1/objects/storages`

Get Version `/configuration/version`

Chapter 7: Troubleshooting

This chapter provides guidelines for how to troubleshoot issues that might occur when using Ops Center Protector.

For troubleshooting activities carried out in conjunction with the support organization it may be necessary to collect diagnostics information from the machines involves in the dataflows displaying an issue. In order for the support organization to investigate the issue it is normally necessary to reproduce this issue with debug trace enabled and then collect diagdata from the machines involved.

To set a node in to debug trace run the following cli command from a command prompt in the <install_path>bin directory:

```
Diagdata -t TraceDebug
```

It will then ask if you want to set the trace level or the running processes – respond with Y to agree to that.

This need to be done on all nodes involved in the dataflow presenting an issue. Once this has been done recreate the issue on the system. Then on each node run the following command:

```
Diagdata -f
```

This will take a while to run but will generate a .zip file in <install_path>runtime send these zip files to the support organization. Once the zip files have been collected run the following command:

```
Diagdata -t TraceAlways
```

And again agree to setting the trace for running processes.

More detailed information on diagdata can be found here (Gathering diagnostic information with diagdata (on page 933)).

Troubleshooting General

This section provides guidelines for how to troubleshoot general issues that might occur when using Ops Center Protector.

Performance issues with the Log Database with very busy Protector instances

Problem:

On very busy systems there can be performance issues with the Log Database on the master node, this is often noticed with attempting to retire lige but can be on creating new log entries also.

Solution:

It is possible to modify the configuration of the PostgreSql database to allow for better performance. The amount of resources which can be assigned to the database will very per system, so it may be best to ramp up the resources in stages and monitor the overall stress on the server to make sure other performance issues are not introduced by allowing the database to use extra resources.

- shared_buffers
- effective_cache_size
- maintenance_work_mem
- work_mem

For example, if the system has 8GB of memory, you could set the following parameter/value pairs in postgresql.conf to:

- shared_buffers = 2GB
- effective_cache_size = 4GB
- maintenance_work_mem = 1GB
- work_mem = 26214KB

Additional tuning information can be found on the PostgreSQL website, however we have found that modifying these parameters have solved performance issues experienced so far.

Error copying file during Protector upgrade

Problem:

The following error is displayed while attempting to upgrade Protector:

```
Error copying file from C:/Program Files/Hitachi/Protector/xxxxxxxx
to C:/Program Files/Hitachi/Protector/xxxxxxxx
---------------------------
Abort Retry Ignore
---------------------------
```

Solution:

Please wait and retry, the binary is probably being slow to shutdown. If the problem persists seek assistance from Hitachi Vantara Support.

How to check if the link between nodes

Problem:

1. Testing using **`linktest`** verifies port-level connectivity between a pair of Protector nodes. For proper Protector operations the following must be satisfied
   a. Any and all Protector nodes must be able to connect to the master on port 30304
   b. The master must be able to connect to any other Protector node on port 30304.

c. Any repository/ISM/storage proxy node must be able to connect to another repository/ISM/storage proxy node on port 30304.

d. To begin the test, select a pair of Protector nodes to perform the below steps. One node will act as the **sender** and the other node will act as the **listener**. The steps will be repeated again where the roles between sender and listener are switched.

e. The **`diagdata`** command is located in the Protector default directory

```
/opt/hitachi/Protector/bin     [Protector 7.x - UNIX]
C:\Program Files\Hitachi\Protector\bin     [Protector 7.x - Windows]
```

f. The **`diagdata`** command will likely require that it be preceded by a "**`./`**" for the command to be found in UNIX environments.

g. All steps with **`diagdata`** are to be executed *on both nodes* before moving onto the next step.

2. Stop the Protector services**`diagdata --stop all`**

3. Double check Protector services are all stopped (except for the **`diagdata`** command itself):

   for versions below 7.0 use:

   ```
   diagdata --query
   ```

   for versions 7.0 and higher use:

   ```
   diagdata --info | more
   ```

4. Triple-check that port 30304 is closed. If the below command result is blank, that confirms the port is closed. If it is not closed (although unlikely), troubleshoot what OS process is holding the port open.

   ```
   netstat -a | grep 30304     [UNIX]
   netstat -a | findstr 30304     [Windows]
   ```

5. For ISM/storage-proxy nodes, check that no **HORCM** process(es) are running. If so, terminate them.

   ```
   ps -ef | grep HORC  and then  kill -9 {PID_of_HORCM_process}   [UNIX]
   Launch Task Manager , then right-click > "End Task" of selected HORCMGR process
   [Windows]
   ```

6. Then setup the listener on one node

   ```
   linktest --listen --port 30304
   ```

7. The listener should respond with

   ```
   Listening on socket…
   Waiting for client to connect…
   ```

8. Then setup the sender on the other node

```
linktest --send --address {IP_address} --port 30304 --ping
```

9. The sender should respond with

```
Connected to server
```

   followed by progress dots.

10. The listener should respond with

```
Connected to server
```

    followed by progress dots.

11. If the sender/listener response is anything but that shown in Steps 9 and 10 then there is a problem in reaching the listener from the sender.
12. Then `{CTRL+C}` on both nodes to exit the **`linktest`**.
13. Switch listener and sender roles between the two nodes by repeating the same Steps 6 through 12.
14. Identify a new node pair, being sure to identify those that will likely communicate across a firewall (e.g., master to/from remote ISM node, local ISM node to/from remote ISM node, etc) and repeat the steps above.
15. When the **`linktest`** is complete be sure to restart the HDID services on all the nodes tested

```
diagdata --start all
```

16. Note in which direction(s) the **`linktest`** failed between what node pairs and notify the team that manages the firewall on the **`linktest`** results.

Solution:

Environmental requirements for successful testing:

- Ensure that Protector is not blocked by a firewall.
- Enable port 30304 (Protector) so that it traverses the firewall in either direction, and port 443 (HTTPS) towards the master server.
- For suspected intermittent network issues, the **`--ping`** parameter may be used along with the **`--logfile`** parameter on both listener and sender nodes for an extended amount of time and left unattended; and when the test window has ended and the tool has been exited, the listener and sender logs may be inspected offline.

  Usage notes for **`linktest`**

```
C:\Program Files\Hitachi\HDID\bin>linktest
--listen or --send must be specified
Usage LINKTEST
        --listen  Listen mode, needs just the port specified
        --send    Send mode, needs port and target address specified
        --port
        --ping    Send a periodic ping
        --no-ack  Don't send/receive acks
        --flood   Size of packets to send
        --address IP address of the listener
        --logfile Name of file output log
```

  `C:\Program Files\Hitachi\HDID\bin>`

- The IP address is required with the `--address` parameter. Hostnames or FQDN should not be used.

Protector backups failing on some Linux distributions

Problem:

On some Linux distributions the **`avahi-daemon`** can interfere with Protector network communications. The may cause backups to fail.

Solution:

Disabling or un-installing the **`avahi-daemon`** resolves this issue.

Logging in with multiple tabs open causes circular log in request

Problem:

Protector allows a user to open multiple browser tabs to enable different UI pages to be viewed simultaneously. If the user is logged out, due to a session timeout for example, an **Unauthorized** dialog is displayed in the current tab. Logging back in to the current tab, then switching to another tab can result in another **Unauthorized** dialog being displayed, forcing the user to log in again. Under certain conditions, a circular log in request may be encountered, as the user switches back and forth between tabs.

Solution:

Log in to the current tab displaying the **Unauthorized** dialog, then if the **Unauthorized** dialog is seen in any other tabs, ignore it and click your browser's Page Reload button instead.

Moving system time backwards causes problems

Problem:

Moving system time backwards causes problems.

Solution:

Moving system time backwards is not recommended. If this is required then restart the Protector hub service after doing so.

Session ID and access failure messages using IE

Problem:

When attempting to access the Protector web UI using Microsoft Internet Explorer on Windows Server Editions, messages of the following type are displayed:

- Received non-login message with no session ID
- Failed to access <Page Name>

Solution:

This is most likely due to Internet Explorer's Enhanced Security Configuration (ESC), which is turned on by default with Server Editions of Windows, preventing cookies from being stored on the machine.

Turn off ESC when using Internet Explorer.

Unicode characters not displayed correctly in web UI

Problem:

When displaying Unicode characters, the Protector web UI may not render them correctly. This may happen when the text being displayed originated on a Protector client in one locale, but is viewed using a browser on a machine in a different locale.

Solution:

Ensure that the browser being used has the required fonts installed. Refer to your browser's documentation for details.

Troubleshooting Hitachi Block

This section provides guidelines for how to troubleshoot issues that might occur when using Hitachi Block storage.

SSB error code information

Problem:

Where can I find SSB error information generated by the storage arrays

Solution:

SSB error codes which are generated as a response to the commands sent to the storage arrays can be found in the attachments to the log messages. Here you will see all the commands and responses sent to/received from the storage arrays.

Cannot create any more snapshots

Problem:

When trying to create a snapshot, the following error message is logged:

```
Handler call failed: [...] failed to create snapshot pair [...]
```

Block storage has a limit of 1024 snapshots per LDEV. When this limit is reached, no more snapshots can be created.

Solution:

Review the RPO and Retention you have specified. If you set an overly aggressive RPO combined with long retention times then the snapshot limit can easily be breached.

Cannot create snapshots in the specified pool

Problem:

The following error message is logged when creating a snapshot:

```
Handler call failed: [] snapshots that are in a different pool
```

This indicates that the primary volume (P-VOL) in question already has a pool assigned for snapshots, and this is not the same pool as defined in the policy. Block storage will only allow one pool to be used for snapshots (S-VOLs) per P-VOL.

Solution:

Review the policy to ensure that the correct snapshot pool is specified.

Cannot make a mounted snapshot multipathed

Problem:

When mounting a snapshot there is always only one path to the disk.

Solution:

To make the mounted snapshot multipathed requires additional steps outside of Protector, dependent on the type of OS on mount host.

Cannot revert snapshot if source machine is unavailable

Problem:

If the source machine where the primary volume is mounted is unavailable, it is not possible to revert a snapshot using Protector. This is because Protector needs to unmount the primary volume prior to reverting it using the snapshot.

Solution:

In the case of a Hitachi Block hardware path classification, the source machine does not need to be available but it is important to unmount the volume before performing the revert.

Cannot tear down adopted replication having other pairings

Problem:

If a replication is adopted in Protector but other replications or snapshots are associated with the replication pair it will not be possible for Protector to remove the replication.

Solution:

Before removing the replication from Protector, first remove the associated snapshots and replications.

Clean-up actions not performed if replication setup fails

Problem:

If the set-up phase of a replication fails, Protector will not clean-up the changes made on the hardware.

Solution:

Detailed steps to perform the clean-up by hand can be found on the Hitachi Vantara knowledge base.

Error 10360 activating UR data flow with grouped source nodes

Problem:

When attempting to activate a data flow in which a node group is specified as the source of a UR replication, the rules compiler generates the following error, indicating that the journal is already assigned to one of the sources within the group:

`The Hitachi Block Storage node` *<Destination Name>* `has a target journal selected for data from` *<SourceName2 'OperationName'>* `operation 'UR' but that journal is already assigned to` *<SourceName1 'OperationName'>* `operation 'UR'`

Solution:

For UR, the selected journals must be unique to each operation and policy in the data flow. This means that a node group cannot be specified as a source when drawing a UR data flow.

Identify each source node explicitly on the data flow and assign unique journals for each replication operation.

Error when removing replications from Protector

Problem:

When removing a replication from Protector the following error message maybe logged:

`Replication pair status change *** Attachment count: 1 ***`

Solution:

In this circumstance (i.e. when removing a replication), this error can be safely ignored.

In all other circumstances this error should not be ignored and should be investigated.

ISM crashed, shutdown or rebooted whilst activating a replication or snapshot data flow

Problem:

If an ISM node crashes, shuts down or reboots whilst a replication or snapshot data flow is being activated, Protector will indicate that the job(s) associated with the activation has failed.

Solution:

Protector is resilient to this type of failure. Try retriggering the operations on the associated data flow(s) once the ISM has finished rebooting.

Policy returns 'Could not start HORCM instance [...]'

Problem:

This error message occurs when a policy is triggered and there is a problem with the command device (CMD) used by the HORCM instance to communicate with the storage device:

- The command device is being shared with another server.
- The command device has been initialized (and possibly put online) as a disk at the OS level.

Solution:

The array maintains a reference count for the command device instances. There can be a maximum of 4096 instances per storage system. If HORCM is not shutdown cleanly using **`horcmshutdown`** then the reference count will not be decremeted and the instance will be leaked. Over time this can cause an inability to start HORCM instances for that array.

Try detaching the command device from the Protector proxy node, clearing its command device attribute then setting it up again (this subtracts the number of instances for that command device from the total count). Alternatively, try unmapping the existing command device and mapping a new one.

Proxy node cannot access the LDEV's resource group

Problem:

When the user account does not have permission to access the LDEV the following error message is logged:

```
Handler call failed: refreshLogicalDevices Failed [...]
```

This is because the LDEV is not in a resource group that the user account can access.

This can be confirmed by navigating to `C:\HORCM\etc` on the proxy node and running the following command:

`raidcom.exe get ldev -ldev_id` *`LdevId`* **`-I0`**

The following message will be returned:

```
raidcom: [EX_EGPERM] Permission denied with the Resource Group
```

Solution:

Ensure that the user account specified for the block device proxy node has access to the resource group being used.

Renaming an LDEV fails due to unresolved variables

Problem:

When trying to name or rename an LDEV using a variable (e.g. `%ORIGIN_SERIAL%`), the following error message is logged:

```
Unable to generate custom secondary logical device name because one
or more variables could not be resolved.
```

If the referenced LDEV has no name then %XXXX_LDEV_NAME% variables will evaluate to an empty string and an error will be logged.

If upgrading from a version earlier than v6.6, the information required to resolve substitution variables may not be present in the metadata for exising LDEVs created prior to upgrade. For example the %ORIGIN_SERIAL% (serial number of the origin storage array for the LDEV) is not included in previous Protector versions.

- For replications, this metadata is added when the replication is next evaluated, following an upgrade to Protector v6.6.
- For snapshot, this metadata is never added since snapshots are not re-evaluated, and will eventually be retired.

Solution:

Review the naming string in the Secondary Logical Device Name field of the **Replication/ Snapshot Operation Properties** or **Rename Secondary LDEVs** dialog, and for:

- `%XXXX_LDEV_NAME%` variables - Ensure the existing LDEV has a name, then retry the naming operation.
- All other variables - For replications, ensure they have been re-evaluated, then retry the naming operation. For snapshots, remove the offending substitution variable(s).

Windows 2008 Server doesn't show changes to reverted disk

Problem:

On Windows 2008 Server, after successfully reverting a disk it is possible that changes are not reflected on the OS.

Solution:

Use the disk management console on the server to offline and then online the disk.

Snapshot fails with error "Cannot create snapshots: specified pool does not have a platform type of "OPEN""

Problem:

Attempt to create snapshot fails with the above error.

Solution:

One possible cause for this is that the destination pool has been deleted outside of Protector.

Replication fails with error "Replication policy specifies a pool that is not of type Thin Provisioning or Smart Tiers, cannot create secondary logical devices"

Problem:

Attempt to create or re-evaluate a replication fails with the above error.

Solution:

One possible cause for this is that the destination pool has been deleted outside of Protector.

Error received stating "Failed to start HORCM instance for command device"

Problem:

Many activities can result in the above error being generated. However within Hitachi there is no more information to help diagnose the issue.

Solution:

Locate the HORCM log file which is by default located in C:\HORCM for Windows or \HORCM for Linux. Investigate the errors produced in the log file, refer to the HORCM documentation for more details on this..

Troubleshooting HCP

This section provides guidelines for how to troubleshoot issues that might occur when using Hitachi Content Platform.

Changing the namespace size on HCP node has no effect

Problem:

Changing the namespace size once an HCP node has been created has no effect.

Solution:

To change the namespace size please use the HCP user interface.

Cannot connect to HCP or HTTP 503 Service Unavailable

Problem:

The messages:

```
Couldn't connect to server
```

or

```
HTTP status 503 Service Unavailable
```

might appear when attempting to authorizing an HCP node.

Solution:

Check that the Management API is enabled in the Tenant Console on HCP for that tenant.

Couldn’t resolve HCP host name

Problem:

The message:

```
Couldn't resolve host name
```

might appear when attempting to authorize an HCP node, or might be logged when tiering.

Solution:

Check that the HCP tenant can be contacted; for example, by pinging the tenant using the host name to ensure DNS is working correctly.

If system level credentials were provided then ensure that the Allow HCP system-level users to manage me and search across my space check box is checked on the HCP’s **Tenant’s Overview** window.

HCP connection returns HTTP status 403 Forbidden

Problem:

The message:

```
HTTP status 403 Forbidden
```

might appear when authorizing an HCP node or might be logged when tiering.

Solution:

Check all the HCP node details are correct including Username, Password and Tenant.

HCP returns Request Entity Too Large error

Problem:

The message:

```
Request Entity Too Large
```

might be logged when tiering to HCP.

Solution:

Check on HCP that the namespace has not exceeded its quota and that the tenant has not exceeded its quota.

HCP restore finished with minor issues

Problem:

The message:

`The restore finished with minor issues.`

might be logged when restoring from HCP.

Solution:

This may occur if files have been deleted from HCP by an application other than Protector. The log attachment lists the paths that the missing files were originally backed up from by Protector.

HCP Node stays offline after node creation

Problem:

After creating a HCP node, it remains offline:

Solution:

Ensure that the Proxy Node defined within the wizard and the Hitachi Content Platform node are time synced. If the time is different by over 5 minutes, then the node creation and future backup may fail. It is recommend that the system time on all nodes should have the correct date/time associated with their particular region and time zone.

Troubleshooting Oracle Database

This section provides guidelines for how to troubleshoot issues that might occur when using Oracle Databases.

An online backup or mount fails

Problem:

In the event of a failure, Protector attempts to revert the operations performed up to the point of the failure, but success cannot be guaranteed.

Solution:

If an Oracle online backup fails, then check that the operations described in the attachments to the log entry `backupOnline Worksteps` are reverted. It is especially important that the step `endOnlineBackup` is performed. If not then execute the steps described in the attachment to the log entry `endOnlineBackup`.

If an Oracle offline backup fails, check that the database is started.

If a mount fails then check if the status of the snapshot is mounted. If so, then perform an un-mount. In addition to this, check that an eventual ASM Diskgroup (**`asmcmd lsdg`**) is un-mounted and removed from the resource group (**`asmcmd umount`** *ASM_DISKGROUP*; **`srvctl remove diskgroup -g`** *ASM_DISKGROUP*). Do not try to issue a mount operation to another server, prior a successful un-mount operation.

Failed to discover Oracle environment when creating application node

Problem:

The following error message is seen when attempting to discover an Oracle Environment in the 'Create Node – Oracle Database' section:

```
Failed to discover Oracle Environment
```

Solution:

1. Click **Rediscover Oracle Environment**

   Note: If click fails, continue with the below steps.

2. Ensure that Oracle is running. If not, start Oracle and repeat the discover process.
3. Edit the file `/etc/nsswitch.conf`
   a. Using a suitable editor, open the file `/etc/nsswitch.conf`
   b. Locate the line that starts with `hosts` and comment it out with a `#`
   c. Copy the line and re-order it so that `dns` is first and `files` is last. E.g.:

   ```
   #hosts: files dns myhostname
   hosts: dns myhostname files
   ```

   d. Save the file.
   e. Attempt to discover the Oracle Environment again. The edited file should be picked up straight away and thus it should not be necessary to restart or reboot.

Oracle database snapshot fails to mount

Problem:

The following error message is logged if you attempt to mount an Oracle database snapshot to a location where one is already mounted:

```
Handler call failed: [...] ASM Diskgroups [...] mounted
```

Solution:

You cannot mount a second snapshot of the database to a location where one is already mounted.

RAC lock on database failed

The following error message is logged, during online snapshot operations, for all but one node in a RAC environment:

```
Handler call failed: [...] RAC Lock Database failed, check if [...]
```

Only one RAC node will succeed to lock the database prior to online snapshotting, all other nodes in the RAC will fail.

Solution:

Check that the operation has succeeded on exactly one of the RAC nodes.

Cannot find Oracle database metadata files on mount

Problem:

The metadata files relating to the mounted Oracle database files cannot be located.

Solution:

The following information message is logged when an Oracle database is mounted:

```
OracleHandler [...]: mount for 'Oracle DB: [...]' finished [...]
```

The attachment lists the destination paths where the Oracle metadata files have been placed.

Error when reverting on ASM in normal/high redundancy mode

Problem:

If you attempt to revert a database running on ASM in normal or high redundancy mode, the following warning message is logged during backup operations, followed by the error message:

```
Handler call failed: [...] normal or high redundancy used; revert
not possible with one mirror
```

Solution:

ASM normal/high redundancy modes are currently not supported.

Warning during backup if data/redo files in the same directory

Problem:

The following warning message is logged during backup operations when the data and redo files of the Oracle database are located in the same directory:

```
OracleHandler [...]: Oracle Database Files and Redo in the [...]
```

Solution:

It is not recommended to have redo logs and data on the same physical disks, as this may cause unusable backups for anything other than crash consistent backups. Please refer to **Oracle application software prerequisites** for more information.

Oracle RAC RPO based policy creates more snapshots than expected

Problem:

When using the 'All Day' schedule option for an Oracle RAC policy and relying on the RPO setting alone, more snapshots than expected may be created. This can happen if rules are activated while one node in the RAC is busy or restarting, causing the activation time to be offset for that node.

Solution:

Try using the ‘Scheduled Time’ option rather than an RPO. Also check that the RAC nodes’ times are synchronized.

Listing Oracle RMAN channel configurations with schedulershow

schedulershow is a CLI tool that is used to aid the Oracle database administrator to configure Oracle RMAN to save and restore data using a datastore managed by Ops Center Protector. It provides the following functionality.

- List the Oracle RMAN dataflows active on the current node
- Create a sample RMAN channel definition for a dataflow operation

In order for a node to store or access any data in a Protector managed datastore using Oracle RMAN, the node has to be part of a dataflow and granted access using the access operation. Once the dataflow is compiled and distributed, it can be listed with schedulershow and Oracle RMAN can access it using an SBT channel.

Table 72 Command line options

Abbreviated Argument	Full Argument	Description
`-h`	`--help`	Display help.
`-c`	`---config <trigger>`	Display RMAN channel configuration for a dataflow. The command provides a list of valid triggers when running the it without parameters..

Usage Example

List all Oracle RMAN dataflows active on this node

```
app/bin/schedulershow

Application DataFlow Source OperationName Storage Trigger

Oracle mysampledf1 myOracleNode Access myRepository
1f04eeec33b10190b79ea6d47a65f3943f5a9de0c1a5d1557afeb420ea604046

Oracle Demo2 myOracleNode Access myRepository
6e7b941d5e73ec5a3ce72b64e70c36c56e6d12f76a379d0982d130d1b505a895
```

Get the Oracle RMAN channel definition for a dataflow

```
app/bin/schedulershow -c
1f04eeec33b10190b79ea6d47a65f3943f5a9de0c1a5d1557afeb420ea604046
```

```
# Application = Oracle
# DataFlow=mysampledf1
# Source=myOracleNode
# OperationName=Access
# Storage=myRepository
#
Trigger=1f04eeec33b10190b79ea6d47a65f3943f5a9de0c1a5d1557afeb420ea60
4046
DEVICE TYPE 'SBT_TAPE' PARMS
'SBT_LIBRARY=/opt/hitachi/protector/app/lib/
libhsbt.so,SBT_PARMS=(TRIGGER=1f04eeec33b10190b79ea6d47a65f3943f5a9d
e0c1a5d1557afeb420ea604046)'
```

Troubleshooting VMware

This section provides guidelines for how to troubleshoot issues that might occur when using VMware.

VM MAC Conflict alarm when restoring cloned VM

Problem:

When restoring a cloned VM with the original VM present, vSphere Client may display the following critical alarm:

```
VM MAC Conflict
```

Solution:

The alarm can be ignored.

When the VM is restored, vSphere may detect a transient MAC conflict between the original and cloned VM before a new MAC address is automatically assigned to the clone.

Restoring VMs to original location fails with 'Restore failed to recover all the required VMs'

Problem:

The following messages are displayed in the logs when attempting to restore a VM to its original location:

```
Handler 'VMwareESX' call failed: Restore failed to recover all the
required VMs
Restore failed to recover all the required VMs *** Attachment
count: 1 ***
```

The attachment identifies the VMDK file associated with the VM that failed to restore.

Cause:

If a VM only resides on one datastore, Protector will not consolidate that VMs snapshots when it is restored (thus all its intermediate snapshots are preserved). This can cause a restore failure under certain conditions.

Solution:

Try selecting Clone instead of Original location when specifying the restore location in Protector. This will cause Protector to consolidate the VM's snapshots.

SAN transport message logged for non-SAN datastore

Problem:

The following message is logged when performing an incremental backup of a datastore that is not accessible using SAN Transport Mode:

```
Disk 'Virtual Hard disk <n> Data.vmdk' snapshot opened with 'san'
transport mode
```

Solution:

This log message may be generated if you are using a virtual machine as the proxy node. For some versions of vCenter, the incorrect transfer mode is reported to Protector. Either ignore the log message or use a physical proxy node to prevent the message.

Host Based SAN Recovery fails with error writing to Virtual Disk

Problem:

Log messages like that found in Figure 635 Log messages from a failed HBB VMware restore attempting to use SAN transport (on page 986) are logged when performing Host based restore that attempts to use the SAN Transport Mode:

14:46:21						
24/01/2023 14:46:21	Repo	GrootVCS	Restore		13277	Restore of VMware Data to 'GrootVCS' from Repository 'Repo' failed. Error: Recovery failed
24/01/2023 14:46:21	Repo	GrootVCS	Statistics		13275	Transfer Statistics: Sent 7.79 KB and Received 21.04 MB
24/01/2023 14:46:15	GrootVCS	GrootVCS	Restore		13557	Recovery failed to recover virtual machine(s) to VMware host 'grootvcs.hdspoole.local'
24/01/2023 14:46:13	GrootVCS	GrootVCS	Restore		13520	Error writing virtual disk data for Hard disk 1 on ge-SanStore1ThickDisk *** Attachment count: 1 ***
24/01/2023 14:46:13	Gaymer	-	VMware		10150	VDDK error: Failed to write virtual disk data to host system Code: 16000 *** Attachment count: 1 ***
24/01/2023 14:45:07	GrootVCS	GrootVCS	Restore		14546	Disk 'Hard disk 1' for SanStore1ThickDisk opened with 'san' transport mode
24/01/2023 14:44:55	GrootVCS	GrootVCS	Restore		13506	Created virtual machine ge-SanStore1ThickDisk

Figure 635 Log messages from a failed HBB VMware restore attempting to use SAN transport

The attachment from log message 10150:

```
VixDiskLib: Detected DiskLib error 4096062 (One of the
parameters supplied is invalid). VixDiskLib: VixDiskLib_Write:
Write 2048 sectors to disk at 0 failed. Error 16000 (One of the
parameters supplied is invalid) (DiskLib error 4096062: One of
the parameters supplied is invalid) at 8019.
```

The attachment from log message 13520:

```
Virtual Disk:Hard disk 1 Virtual Machine:<VM_NAME> Host:
<VMWARE_NODE> vddk error: Failed to write virtual disk data to
host system.
```

Solution:

The above errors are seen because VMware has selected the SAN transport method as the proxy and target datastore are both on the same SAN. However, the Proxy OS does not have write access to the VMFS volume. Ensure that the VMFS volume is exposed but not mounted, and has both read and write access.

In Windows this can be achieved by the diskpart CLI tools as described below:

`list disk` will display all disks (as shown by the disk management app.

`select disk` is used to select a disk for further inspection/modification

`detail disk` is used to display information about the currently selected disk

`attribute disk clear readonly` is used to clear the disk's 'read-only' attribute

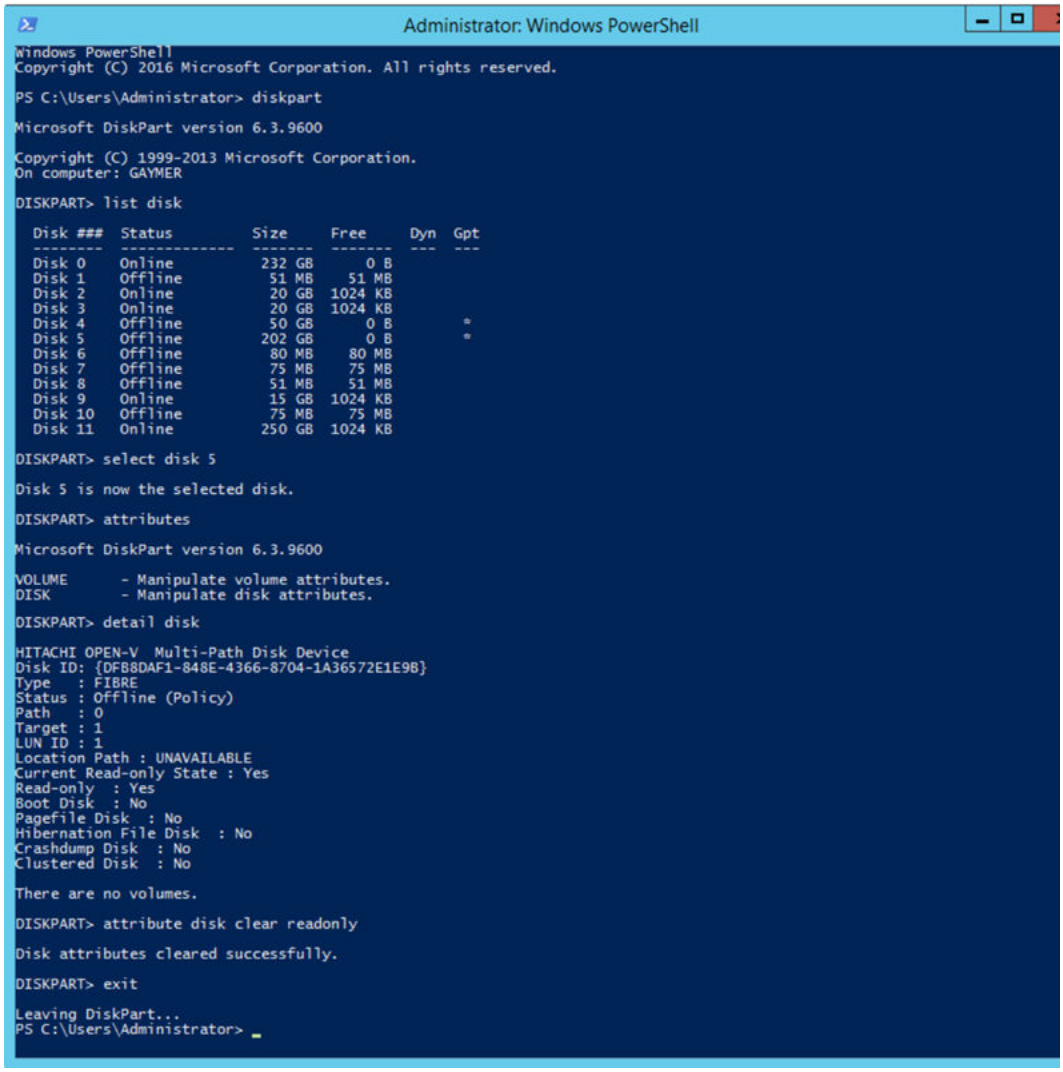


Figure 636 Diskpart CLI tool

In linux the readwrite options can be checked via `lsblk`, e.g.

```
[root@LinuxHost ~]# lsblk -o NAME,FSTYPE,RO
NAME                    FSTYPE                    RO
sdb                                                0
└─sdb1                  VMFS_volume_member         0

[root@LinuxHost ~]# lsblk -o NAME,FSTYPE,MODE
NAME                    FSTYPE             MODE
sdb                                        brw-rw----
└─sdb1                  VMFS_volume_member brw-rw----
```

SRM recovery fails with 'Cannot process consistency group [...] expected [...] role target'

Problem:

The following message is displayed by SRM when performing a test or real fail-over or fail-back recovery operation:

```
Failed to sync data on replica devices.

Cause:

Cannot process consistency group '{<CTG ID>}' with role
'promotedTarget' when expected consistency group with role 'target'
```

Cause:

SRM checks that replications are in the expected state before performing the recovery operation. This message may be generated if the continuous TrueCopy replication between the production and recovery site has been paused or swapped outside of SRM.

SRM replications must not be paused or swapped outside of SRM.

Solution:

In SRM, perform Discover Devices, then check the Status of the datastores. If the status is *Failover complete*, check if the corresponding TC replication in Protector is in the paused state and un-pause it if required.

Datastores status should be either *Outgoing Replication* or *Incoming Replication* before starting a failover.

vRO Ad Hoc Backup fails with '[...] Tag 'HDID/Protector Ad Hoc' already in use [...]'

Problem:

The following message is logged in vRO when running the 'Ad Hoc Backup' workflow:

```
(com.hitachivantara.protector.backup/performAdHocBackupOf) Error
in           (Dynamic Script Module name : performAdHocBackupOf#17)

HdidPluginException[Cannot perform Ad Hoc Backup, Tag 'HDID/
Protector           Ad Hoc' already in use. Other VM(s) may be
backed up because the following objects are           tagged:
[Name: , Type: VirtualMachine, Id: vm-4398]

Please try again later or contact your system administrator if the
failure persists.
```

Cause:

If you manually assign the 'HDID/Protector Ad Hoc' tag to a VM and subsequently delete that VM, then run the 'Ad Hoc Backup' workflow on any other VM, the operation will fail with the above error.

Solution:

Delete the 'HDID/Protector Ad Hoc' tag from vSphere then recreate it.

The 'Ad Hoc' tag should never be manually assigned. It should only be automatically assigned by Protector vRO workflow scripts.

vRO Ad Hoc Backup fails with 'Cause: VMwareException[Tagging cardinality violation]'

Problem:

The following message is logged in vRO when running the 'Ad Hoc Backup' workflow:

```
(com.hitachivantara.protector.backup/performAdHocBackupOf) Error in (Dynamic Script
Module
        name : performAdHocBackupOf#9) MasterException[failed to associate Tag [name:
Protector Ad
        Hoc, Id: urn:vmomi:InventoryServiceTag:...] to VM [vm-...],
 Cause:
        VMwareException[Tagging cardinality violation] ]
```

Cause:

For the vRO Ad Hoc Backup Workflow to work, it needs to tag the desired VM with the ‘HDID/Protector Ad Hoc’ tag.

The error *‘Tagging cardinality* violation’ indicates that the VM already has a tag that belongs to the same tag category as the 'HDID/Protector Ad Hoc' tag. And that this tag category is restricted to “Tags Per Object: One tag”. As such only one tag from the category can be assigned to the desired VM.

Solution:

Move the 'HDID/Protector Ad Hoc' tag into its own category.

If you cannot create a new category, then alter the category tag cardinality such that “Tags Per Object” is “Many tags”. (See screenshot below)

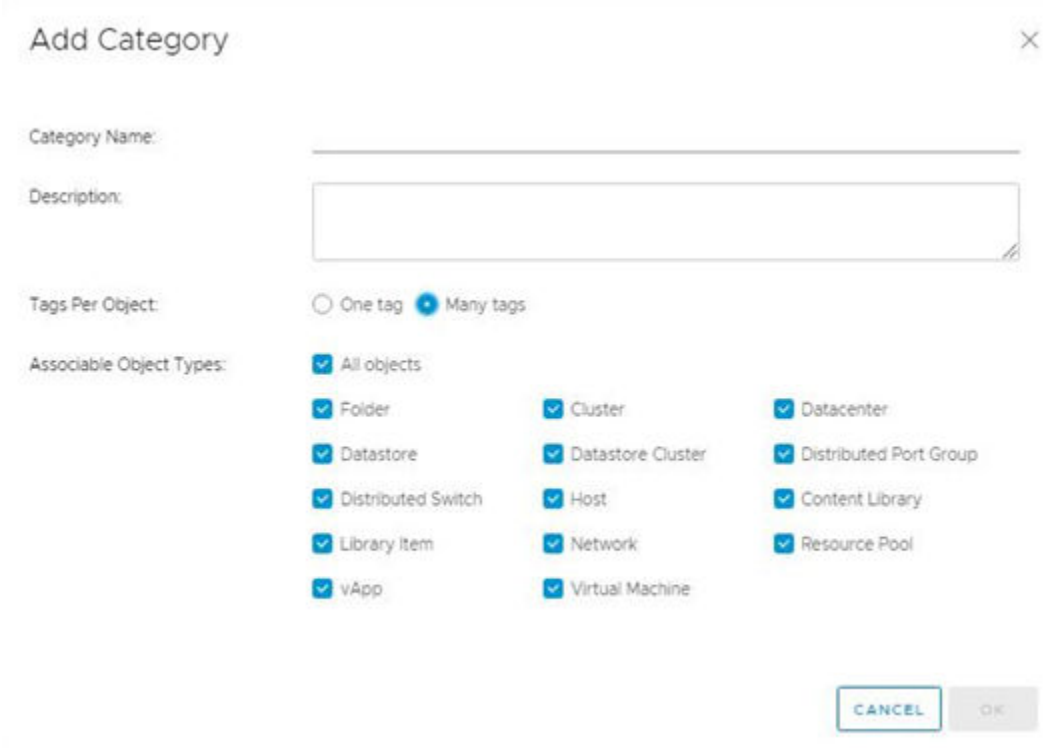


Figure 637 Creating a new Category for vRO tags

Troubleshooting Amazon S3

This section provides guidelines for how to troubleshoot issues that might occur when using Amazon S3.

Amazon S3 Node stays offline after node creation

Problem:

After creating an Amazon S3 node, it remains offline:

Solution:

Ensure that the Proxy Node defined within the wizard and the Amazon S3 Proxy node are time synced. If the time is different by over 5 minutes, then the node creation and future backup may fail. It is recommend that the system time on all nodes should have the correct date/time associated with their particular region and time zone.

Troubleshooting Oracle RMAN

This section provides guidelines for how to troubleshoot issues that might occur when using Oracle RMAN.

ORA-27211: Failed to load Media Management Library

Oracle RMAN error message:

```
ORA-19554: error allocating device, device type: SBT_TAPE, device
name:

ORA-27211: Failed to load Media Management Library

Additional information: 4156

The Media Management Library can not be loaded.
```

Cause:

Oracle RMAN cannot locate Protectors Oracle RMAN integration library

Solution:

Check the RMAN channel configuration and ensure the SBT_LIBARY path section is correct and the filesystem permissions allow the oracle user to read the file.

The schedulershow command will generate channel configurations with the correct path, however you will have to check the filesystem permission manually.

ORA-19511: ... sbtinit2: initialize communication failed

Oracle RMAN error message:

```
ORA-19554: error allocating device, device type: SBT_TAPE, device
name:

ORA-27023: skgfqsbi: media manager protocol error

ORA-19511: non RMAN, but media manager or vendor specific failure,
error text:

sbtinit2: initialize communication failed
```

Cause:

There are two possible causes for this error message

- The trigger / channel is no longer valid
- The UBI datastore (e.g Gen2 Repository, Hitachi Content Platform Gen2, Amazon S3) is not available

Solution:

- Check the target storage (e.g. Gen2 Repository, Hitachi Content Platform Gen2, Amazon S3) is available and listed as online
- Check the Oracle Application node is part of the dataflow and the dataflow is activated
- Compare the channel definition (specifically the trigger id) with the definitions listed by the schedulershow command

ORA-19511: ... sbtbackup: write to store not allowed

Oracle RMAN error message:

```
ORA-27028: skgfqcre: sbtbackup returned error

ORA-19511: non RMAN, but media manager or vendor specific failure,
error text:

sbtbackup: write to store not allowed
```

Cause

The policy does not allow this node to write into the store.

Solution:

Ensure you use an “Access“ operation with read/write permissions on the dataflow.

ORA-19511: ... Data Folder not found

Oracle RMAN error message:

```
ORA-27027: sbtremove2 returned error

ORA-19511: non RMAN, but media manager or vendor specific failure,
error text:

sbtremove: internal - Data Folder not found
```

Cause

The server lost the connection to the target storage during the operation

Solution:

- Ensure the datastore used in this dataflow is online and available
- Restart RMAN
- Use RMAN to perform a consistency check (e.g. crosscheck backup / crosscheck backupset)
- Retry the operation

If this does not solve your problem, please contact support

ORA-19511: ... no Root Object

Oracle RMAN error message:

```
ORA-27028: skgfqcre: sbtbackup returned error

ORA-19511: non RMAN, but media manager or vendor specific failure,
error text:

internal - no Root Object
```

Cause

The RMAN integration cannot establish the connection to the storage target.

Solution:

- Ensure the datastore used in this dataflow is online and available
- Restart RMAN

- Use RMAN to perform a consistency check (e.g. crosscheck backup / crosscheck backupset)
- Retry the operation

If this does not solve your problem, please contact support

ORA-19511: ... database <dbname> not allowed

Oracle RMAN error message:

```
ORA-27028: skgfqcre: sbtbackup returned error

ORA-19511: non RMAN, but media manager or vendor specific failure,
error text:

sbtbackup: processing database myoradb not allowed
```

Cause

The database is not allowed access in the **Oracle RMAN classification**

Solution:

- Change the Oracle RMAN classification of the policy to allow the database access
- Activate the dataflow
- Restart RMAN and retry the operation

ORA-19511: ... Store exists but cannot be accessed

Oracle RMAN error message:

```
RMAN-03002: failure of backup command at 06/03/2020 12:33:44

ORA-27023: skgfqsbi: media manager protocol error

ORA-19511: non RMAN, but media manager or vendor specific failure,
error text:

sbtinit2: init store failed -- initializeStore: Store exists but
cannot be accessed by this node. Store permissions may have to be
adjusted
```

Cause

Store permissions are not correct and need to be updated. The owner node must be a node on the dataflow.

Solution:

Change the store ownership, by performing the following steps:

- Identify the following information
  - Name of the dataflow
  - Name of the policy
  - Name of a Oracle application node using the policy on this dataflow
  - Name of the storage target node (e.g. gen2 repository used to store the data)
- List the stores inventory of the storage target node
- Select the relevant store for the dataflow from the list and select the "Repurpose" action
  - Ensure you pick the correct store by comparing the dataflow and policy names. The Application should be "Oracle Database"
- Fill in the following information and repurpose

Control	Value
Source Node	Select the Oracle application node
Policy	Leave blank
Data Flow	Leave blank
Automatically rename store	Ensure this is not selected

- Restart RMAN and retry the operation

Troubleshooting Hyper-V

This section provides guidelines for how to troubleshoot issues that might occur when using Hyper-V.

A virtual machine is not included in a backup.

Problem:

When performing a backup, it does not include a virtual machine. Protector does not report any errors for the virtual machine during the backup job.

Solution:

The Hyper-V classification does most likely not include this virtual machine.

Review and update the Hyper-V classification of the policy. Use the preview functionality for a list of VMs included by the policy.

A virtual machine is no longer included in backups after it has been restored.

Problem:

A virtual machine has been restored, now it is no longer included in backups.

Cause:

When selecting a virtual machine explicitly (via browse) in a Hyper-V classification, Protector will refer to this VM via its ID. Restoring a virtual machine assigns a new ID to the restored VM. As the classification still refers to the original ID, the restored VM will not be included in the backup.

Solution:

Update the Hyper-V classification to include the restored VM and re-activate the data flow.

Restore to original fails stating the node does longer exists

Problem:

When you try to restore to the original, an error is displayed, indicating that the node no longer exists.

Solution:

Perform a clone restore instead

Restore to Hyper-V cluster fails as virtual machine already exist

Problem:

Restoring a virtual machine to a Hyper-V cluster fails, stating that the virtual machine already exists, however, Hyper-V Manager does not list the VM.

Cause:

This is usually caused by a partially deleted VM. In case of a clustered Hyper-V there is the virtual machine and an associated cluster role. When the virtual machine is deleted via the Microsoft's Failover Cluster Manager or SCVMM both the VM And the role are deleted. However, if the virtual machine is deleted via PowerShell or Hyper-V Manager the cluster role may remain.

Solution:

Ensure both the virtual machine and the associated cluster role are deleted, before you re-try the restore.

Verification of Hyper-V credential fails

Problem:

When creating or editing a Hyper-V node the provided credentials are rejected.

Solution:

double check you provided the correct domain, user, and password

verify Hyper-V and WMI services are running on all nodes, which are part of the Hyper-V setup

In case of a clustered Hyper-V environment, verify the cluster service is running on all nodes

Virtual FC adapters, pass through disks or mapped ISOs are missing after a restore

Problem:

After restoring the virtual machine, it does no longer contain virtual fibre channel adapters, mapping to pass through disks or mappings to ISOs.

Solution:

It is currently not possible to protect or restore these connections to external resources. Use your Hyper-V management tool to re-add them.

VM is skipped during backup because config version is too small

Problem:

During a backup job, a virtual machine is skipped. The associated error message shows that the configuration version is too low.

Solution:

Use the Microsoft's Hyper-V Manager or PowerShell to upgrade the virtual machine version. See https://docs.microsoft.com/en-us/windows-server/virtualization/hyper-v/deploy/upgrade-virtual-machine-version-in-hyper-v-on-windows-or-windows-server for additional guidance.

Glossary

Archive
A copy that is created for long-term retention.

Asynchronous journalling
Transactions are written to disk and also placed in a journal log file, to protect against data loss in the event of a system failure. Transactions from the log file are sent to the destination machine.

Asynchronous replication
Transactions are held in memory before being sent over the network. If the network is unavailable then transactions are written to disk and sent to the destination machine when the connection is re-established. Asynchronous replication is optimal for connections with sporadic efficiency.

Backup
A copy that is created for operational and disaster recovery.

Bandwidth throttling
Used to control when and what proportion of available network bandwidth is used by Ops Center Protector for replication.

Batch backup
A process by which the repository is updated periodically using scheduled resynchronizations. This method involves a scan of the source machine's file system, but only the changed bytes are transferred and stored. This method is useful for data that does not change often, such as data contained on the operating system disk. Linux based source nodes are only able to perform batch backups.

Clone
An operation where a copy of the database is created in another storage location in a local or remote site.

COPY
A hardware orchestration related status code that indicates that a volume pair is being created. An initial copy or resynchronization is being performed.

Data flow
Identifies the data sources, movers and destinations participating in a backup, along with interconnection paths between them. Policies are assigned to each node to determine what type of data is backed up.

Data source
A machine hosting a file system or application where the Protector client software is installed.

Deduplication
A method of reducing the amount of storage space that your organization requires, to archive data, by replacing multiple instances of identical data with references to a single instance of that data.

Destination node
A machine that is capable of receiving data for the purposes of archiving. This machine might be the Ops Center Protector Repository or Block device.

Dynamic Provisioning Virtual Volume (DP-VOL)
Dynamic Provisioning Virtual Volume. A virtual volume that has no memory space. Used in Dynamic Provisioning.

Hitachi Open Remote Copy Manager (HORCM)

HORCM is a daemon process on the CCI server that communicates with the storage system and remote servers.

Intelligent Storage Manager (ISM)

A Protector Client node that acts as a proxy to Block storage devices.

ISM may also refer to the Intelligent Storage Manager process that runs within the Protector Client software.

License key

A unique, alphanumeric code that is associated with the unique machine ID that is generated during the Ops Center Protector installation. The license key must be activated in order to use the software.

Logical Device (LDEV)

An individual drive (or multiple drives in a RAID configuration) in the storage system. An LDEV might or might not contain any data and might or might not be assigned to any hosts. Each LDEV has a unique identifier, or address, within the storage system. The identifier is composed of the logical disk controller (LDKC) number, control unit (CU) number, and LDEV number.

The LDEV IDs within a storage system do not change. An LDEV formatted for use by open-system hosts is called a logical unit (LU).

Master node

The machine that controls the actions of other nodes within the Ops Center Protector network.

Metadata Store (MDS)

Records metadata that describes items that are held in repositories. The MDS supports indexing of stored data, thus enabling fast searches when locating data for restoration.

Mover

Defines the type of data movement operation to be performed between source and destination nodes, during the creation of a data flow. Batch movers perform block level data transfers on a scheduled basis, whereas continuous movers perform byte level data transfers on a near-continuous basis.

Node Group

Multiple machines of the same type can be assigned to one or more node groups. Within the Data Flow page, you can assign policies to nodes within node groups en-mass.

PAIR

A hardware orchestration related status code that indicates that a volume pair is now created. The initial copy has finished, and the paired volumes are synchronized.

PFUL

A hardware orchestration related status code that indicates that the amount of the data in the journal volume exceeds the threshold. The volume pair is not split and data continues to be copied.

PFUS

A hardware orchestration related status code that indicates that the amount of the data in the journal volume has reached 100% and the volume pair is split and data is no longer being copied.

Policy

A configurable data protection objective that is mapped to machines or groups, and to the data management agents that implement the policy. Multiple policies can be assigned to a single node.

Primary Volume (P-VOL)

The volume in a copy pair that contains the original data to be replicated. The data on the P-VOL is duplicated synchronously or asynchronously to the secondary volume (S-VOL).

PSUE

A hardware orchestration related status code that indicates that a volume pair is split due to an error.

PSUS

A hardware orchestration related status code that indicates that a volume pair is split by operation, or deleted from the storage system on the secondary site. This value is output for the P-VOL.

Raw Device Mapping (RDM)

Raw Device Mapping enables a LUN from a SAN to be directly connected to a VMware VM.

Recovery Point Objective (RPO)

The frequency at which a backup will occur. This governs the point in time to which data can be recovered should a restore be needed.

Refreshed Thin Image (RTI)

A local replication technique based on Thin Image snapshot. The S-VOL is refreshed based on a schedule or on demand. The S-VOL is a thin copy of the P-VOL and therefore needs the P-VOL to be available. Because the S-VOL is refreshed, its ID remains unchanged whenever its contents are updated.

Replication

An operation where a copy of the data is created in another local or remote location automatically.

Repository

A destination node that stores data from one or more source nodes. The Ops Center Protector Repository supports batch backup, archiving, and versioning policies.

Secondary Volume (S-VOL)

The volume in a copy pair that is the copy of the original data on the primary volume (P-VOL). See also primary volume.

SMPL

A hardware orchestration related status code that indicates that a volume is un-paired.

Snapshot (Thin Image)

A point in time copy of the data that is based on references to the original data.

Source node

Any node (server, workstation or virtual machine) that hosts data to be protected by Ops Center Protector. The source node has an Active Data Change Agent, which is responsible for monitoring the host file system and performing the relevant actions defined by the policies. Nodes need to be configured as a source node if they need to transfer locally stored data to a destination node, or implement data tracking, blocking and auditing functions. A node can be both a source and destination simultaneously.

SSUS

A hardware orchestration related status code that indicates that a volume pair is split by operation, or deleted from the storage system on the secondary site. This value is output for the S-VOL.

SSWS

A hardware orchestration related status code that indicates that the P-VOL and S-VOL are switched. The S-VOL is writable.

Synchronous replication

Transactions are transferred to the remote storage device immediately and the write operation is signaled as completed only once data is confirmed as written to both primary and secondary volumes. Synchronous replication is optimal for connections with high efficiency.

Virtual Storage Machine (VSM)

Virtual Storage Machine. A virtualised block storage device that exists within a physical storage array.

Conventions for storage capacity values

Physical storage capacity values (for example, disk drive capacity) are calculated based on the following values:

Physical capacity unit	Value
1 kilobyte (KB)	1,000 (10^3) bytes
1 megabyte (MB)	1,000 KB or $1{,}000^2$ bytes
1 gigabyte (GB)	1,000 MB or $1{,}000^3$ bytes
1 terabyte (TB)	1,000 GB or $1{,}000^4$ bytes
1 petabyte (PB)	1,000 TB or $1{,}000^5$ bytes
1 exabyte (EB)	1,000 PB or $1{,}000^6$ bytes

Logical capacity values (for example, logical device capacity, cache memory capacity) are calculated based on the following values:

Logical capacity unit	Value
1 block	512 bytes
1 cylinder	Mainframe: 870 KB Open-systems: ▪ OPEN-V: 960 KB ▪ Others: 720 KB
1 KB	1,024 (2^{10}) bytes
1 MB	1,024 KB or $1{,}024^2$ bytes
1 GB	1,024 MB or $1{,}024^3$ bytes
1 TB	1,024 GB or $1{,}024^4$ bytes
1 PB	1,024 TB or $1{,}024^5$ bytes
1 EB	1,024 PB or $1{,}024^6$ bytes

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact