

Hitachi Ops Center Protector

7.4

Microsoft SQL Server Application Guide

This document is intended for database administrators who want to protect Microsoft SQL Server databases using Hitachi Ops Center Protector. It is assumed that the reader has a good working knowledge of Microsoft SQL Server, Hitachi Block Storage administration and network administration.

© 2016, 2022 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., or Hitachi Vantara LLC (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.

AI/AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Edge, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, 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 at <https://www.hitachivantara.com/en-us/company/legal.html>.

Contents

Preface.....	5
Software version.....	5
Intended audience.....	5
Related documents.....	6
Document conventions.....	6
Conventions for storage capacity values.....	7
Accessing product documentation.....	8
Getting help.....	9
Comments.....	9
Chapter 1: Before you begin.....	10
Supported configurations.....	10
Prerequisites.....	10
Application prerequisites.....	11
Hitachi Block prerequisites.....	11
Chapter 2: Microsoft SQL Server Backup workflows.....	15
About Microsoft SQL Server Policy Classifications.....	15
Block based workflows.....	15
How to create database restore points with block snapshots.....	16
Chapter 3: Microsoft SQL Server Restore workflows.....	19
About mounting Microsoft SQL Server block-based backups.....	19
About reverting Microsoft SQL Server block-based backups	20
How to mount a database from a block-based backup and attach it to an existing instance.	23
How to revert multiple databases from a block-based backup.	25
How to revert a user database online without affecting other databases in the same instance.	26
How to revert a user database and prepare it to apply transaction logs.....	27
Chapter 4: Reference.....	29
Nodes UI Reference.....	29
Microsoft SQL Server Node Wizard.....	29
Policies UI Reference.....	32
Microsoft SQL Server Classification Wizard.....	32

Microsoft SQL Server Database Selection Wizard.....	34
Microsoft SQL Server Instance Credentials Dialog.....	35
Microsoft SQL Server User Privileges.....	36
Restore UI Reference.....	36
SQL Server Mount Wizard.....	37
SQL Server Revert Wizard.....	41
Chapter 5: Troubleshooting.....	45
Troubleshooting Microsoft SQL Server.....	45
Cannot trigger a revert or revert fails without a force revert option.	45
Clustered SQL Server instances does not list all nodes.	45
Database(s) are not listed in Microsoft SQL Management Studio after offline revert.	46
Databases are not listed for when creating a SQL Server classification.....	46
Invalid Credentials; Unable to connect to instance.....	46
Issues with snapshots of mounted databases.	47
Mount Error - Database(s) already exists in instance.	47
No cluster nodes returned from cluster during node creation.....	47
Online Revert option is disabled.....	48
Unable to apply transaction log backups after successful online revert.....	48
Glossary.....	49

Preface

This guide describes how to backup and restore Microsoft SQL databases using Hitachi Ops Center Protector.

Ops Center Protector orchestrates the creation, retention and restoration of application-consistent and crash consistent snapshots and clones for Microsoft® SQL Server databases. Data protection policies are combined with data flow diagrams to automate local and remote snapshots and replications for end-to-end data protection and recovery solutions. These snapshots and clones then can be used to revert production databases to specific points in time and to create copies for repurposing scenarios.

Software version

This document revision applies to Ops Center Protector version 7.4. Please refer to the accompanying Release Notes for information on what's changed in this release.

Intended audience

This document is intended for database administrators who wants to protect Microsoft SQL Server databases using Hitachi Ops Center Protector. It is assumed that the reader has a good working knowledge of Microsoft SQL Server, Hitachi Block Storage administration and network administration.

If you are new to Ops Center Protector, we recommend that you start by referring to the *Hitachi Ops Center Protector User's Guide* so that you understand the basic concepts, workflows and user interface.

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	<ul style="list-style-type: none"> ▪ Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example: Click OK. ▪ Indicates emphasized words in list items.
<i>Italic</i>	<ul style="list-style-type: none"> ▪ 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: <pre>pairedisplay -g group</pre> (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: <code>pairedisplay -g oradb</code>

Convention	Description
< > angle brackets	<p>Indicates variables in the following scenarios:</p> <ul style="list-style-type: none"> Variables are not clearly separated from the surrounding text or from other variables. Example: <pre>Status-<report-name><file-version>.csv</pre> 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	<p>Indicates that you have a choice between two or more options or arguments. Examples:</p> <p>[a b] indicates that you can choose a, b, or nothing.</p> <p>{ a b } indicates that you must choose either a or b.</p>

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

Icon	Label	Description
	Note	Calls attention to important or additional information.
	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
	Caution	Warns the user of adverse conditions and/or consequences (for example, disruptive operations, data loss, or a system crash).
	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: <ul style="list-style-type: none"> ▪ 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 the Hitachi Vantara Support Website: <https://knowledge.hitachivantara.com/Documents>. 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](https://support.hitachivantara.com/en_us/contact-us.html) 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](https://community.hitachivantara.com) 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 us your comments on this document to doc.comments@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

Supported configurations

The following Microsoft SQL Server configurations and technologies are supported:

- Standalone Microsoft SQL Server with one or more instances and databases located on supported block storage.
- Microsoft SQL Server Instances using Windows Failover Cluster with databases located on supported block storage.

The following data protection technologies are supported:

- Block based snapshots.
- Local and remote replication of databases.

The following technologies are **not** supported:

- Always On availability groups.
- Transaction log backups using Protector. Transaction logs can be protected with the toolset provided by Microsoft or 3rd party tools.
- BLOBs which are stored outside of the database.
- Databases on Cluster Shared Volumes utilizing ReFS instead of NTFS.
- Databases on unsupported Block storage {see [Hitachi Block prerequisites \(on page 11\)](#)}

Prerequisites

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

To ensure that your hardware and software environment is fully supported, please refer to <https://www.hitachivantara.com/en-us/products/data-protection/ops-center-protector.html>.

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

Application prerequisites

To successfully protect a Microsoft SQL Server environment several prerequisites must be met:

- Protector Client software is installed on all nodes which are part of a Microsoft SQL Server setup
- For clustered setups, install the Protector client on each cluster node. No special installation steps are required. The procedure is the same as for standalone systems (For more details, please refer to “How to install/upgrade Protector on Windows and Linux” in the Protector User’s Guide)
- Clustered SQL Server instances must consist of a minimum of two nodes.
- A user account is provided for use by Protector, having the specified [Microsoft SQL Server User Privileges \(on page 36\)](#).
- In order to configure protection for clustered SQL Server setups, all cluster nodes must be running.



Note: This is a requirement for the setup. Once configured scheduled or RPO based backups will work even if only a subset of the cluster nodes is online.

- For block-based data protection, ensure that:
 - Databases are located on a single supported block storage.
 - Databases which should not be backed up together are on separate LUNs.



Note: It is recommended that each database uses a separate set of LUNs.



Note: To apply transaction logs after an online revert, multiple databases must not share LUNs and there must only be a single database in a backup. See [About reverting Microsoft SQL Server block-based backups \(on page 20\)](#) for more details.

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://www.hitachivantara.com/en-us/products/data-protection/ops-center-protector.html>.

- 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) and pre-configured RCU paths between arrays for remote replication technologies

Chapter 2: Microsoft SQL Server Backup workflows

The following topics describe the steps required to configure policies and data flows to implement a number of different data protection scenarios.

For a detailed introduction on how to work with the Protector user interface, please refer to *Hitachi Ops Center Protector User's Guide*.

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 \(on page 32\)](#). Once the source Microsoft SQL node has been selected, databases can be added using the [Microsoft SQL Server Database Selection Wizard \(on page 34\)](#).

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.

Block based workflows

This section addresses the workflows for block-based backups.

How to create database restore points with block snapshots

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.
- The Protector Client software has been installed on all nodes of the Microsoft SQL Server setup which should be protected.
- The Protector Client software has been installed on the node that will act as a proxy for the Hitachi storage device where the SQL Server data is located. Note that for a Thin Image snapshot, the source and destination LDEVs are located on the same device.
- The block storage device has been set up as per the Protector requirements and prerequisites. Refer to [Hitachi Block prerequisites \(on page 11\)](#).
- 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.
- A user has been created that provides the required privileges as detailed in [Microsoft SQL Server User Privileges \(on page 36\)](#). This user will be required when creating the SQL Server classification in the steps that follow.

This task describes the steps to follow when snapshotting databases that reside on a Hitachi Block storage device. The data flow and policy are as follows:



Figure 1 Microsoft SQL Server Node

Classification Type	Parameters	Value
Microsoft SQL Server Database	Node	Microsoft SQL Server application node hosting the instance.
	Instance	Instance hosting the databases
	User	User for backup
	Included Items	Refer to Microsoft SQL Server Database Selection Wizard (on page 34) on how to include database into the backup

Classification Type	Parameters	Value
	Backup Mode	Full Copy

Operation Type	Parameters	Value	Assigned Nodes
Snapshot	Mode	Hardware	SQL Server application node
	Hardware Type	Hitachi Block	
	Run Options	Run on RPO	
	RPO	2 hours	
	Retention	2 days	

Procedure

1. Locate the source *OS Host* nodes for your Microsoft SQL Server environment in the **Nodes Inventory** and check that the nodes are authorized and online.

These nodes represent the Protector Clients installed on your Microsoft SQL Server environment.

2. Create a new *SQL Server* node (unless a suitable one already exists) using the [Microsoft SQL Server Node Wizard](#) (on page 29).

The *SQL Server* node type is grouped under **Application** in the **Node Type Wizard**.

- a. Select one node which is part of your *Microsoft SQL Server* environment.
- b. Select which configuration you want the *SQL Server* 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 SQL Server database 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 **Block Storage Node Wizard** 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 SQL Server 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**, [Microsoft SQL Server Classification Wizard](#) (on page 32) and **Snapshot Operation Wizard**.

The *SQL Server* classification is grouped under **Application** in the **Policy Wizard**.


6. Draw a data flow as shown in the figure above, that shows only the *SQL Server* source node, using the Data Flow Wizard.

At this stage, the snapshot icon  is not shown.

7. Assign the *Snapshot* operation to the *SQL Server* source node. The *SQL Server-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 *SQL Server* database's data is located. Then select a **Snapshot Pool** from one of the available Thin Image 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 database 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 3: Microsoft SQL Server Restore workflows

About mounting Microsoft SQL Server block-based backups

Protector can mount a block-based backup to an existing SQL Server instance. This will make the disks comprising the backup available on the machines hosting the target instance and optionally attach the databases so the user can access them.



Note: We strongly recommend to always use the mount duplicate option in the mount dialog. While mounting the original is possible, it will modify the original backup data, which may prevent further mounts and reverts.

Mounting to clustered instances

When mounting to clustered instances, Protector will mount the block-based backup as traditional clustered disks and add them as a dependency for the selected target SQL Server instance. This ensures that the instance will still be able to failover, and the mounted databases remain online in case of a cluster failover.

Prerequisites

The following prerequisites must be met to mount a block-based Microsoft SQL Server database backup.

- A block-based backup of a Microsoft SQL Server is available.

For the target instance:

- A Microsoft SQL server node exists for the **target** SQL Server instance.
- Credentials for the SQL Server instance are available. Please refer to [Microsoft SQL Server User Privileges \(on page 36\)](#) for more details.
- The target instance supports the internal database version of the databases in the backup.
- All nodes of the target instance are up and running.
- When mounting to a clustered instance, the instance must consist of a minimum of two nodes. Mounting to single-node clusters is not supported.

The following table lists the supported configuration combinations for mount:

Source	Mount Target Standalone	Mount Target Cluster	Mount Target Source Node
Standalone	yes	yes	yes
Cluster with CSV	yes	yes	yes
Cluster with traditional clustered disks	yes	yes	no



Note: Mounted databases should not be used as a basis for further block-based backups. Including mounted databases in backups may cause problem during backup or when you try to unmount the database

About reverting Microsoft SQL Server block-based backups

Protector can revert blocked-based backups to restore databases captured at a previous point in time. When reverting, Protector will revert all disks which are part of the block-based backup. It is not possible to revert just a single database of a block-based backup.



Tip: If you want to restore a single database from a backup, mount the database and copy the files manually.

Protector offers two revert types. The following table summarizes the benefits and disadvantages of each option:

	Pro	Contra
Offline	<ul style="list-style-type: none"> ▪ Supports multiple databases. ▪ Supports system databases. 	<ul style="list-style-type: none"> ▪ Instance and all databases are temporarily offline during revert. ▪ It is not possible to apply transaction logs.
Online	<ul style="list-style-type: none"> ▪ Instances and other databases remain online. ▪ User has the option to apply transaction logs manually. 	<ul style="list-style-type: none"> ▪ Online revert does not support system databases.

Offline Revert

Offline Revert can revert both user and system databases, as well as multiple databases at the same time.

During the revert, the instance and databases on the instance will be offline, making them inaccessible for users. An offline revert will detach the user database(s) which are part of the backup and revert the data on disk. For backups containing only user databases, Protector can then attach selected databases.



Note:

Reverting System databases

System databases capture the configuration of the instance. Reverting them will change the configuration of the instance to what it was at the point in time of the backup. This includes user accounts, permissions, and the list of attached databases. As a result, user databases reverted together with system databases may be attached after the revert.

The following table illustrates the status of disk, instances, and databases after a successful revert.

Table 1 Status after offline revert

Backup content	Disk	Instance	Reverted database(s)	Other database(s)
User database(s)	reverted	online	Depends on user selection. Options are: <ul style="list-style-type: none"> ▪ Detached ▪ Attached and online 	Unchanged. Same as before the revert.
System database(s)	reverted	online	online	Depends on the state that is captured in the system database(s) at time of block-based backup.
Mix of system and user databases	reverted	online	System databases will be online. State of user databases depends on the state that is captured in the system databases at time of block-based backup.	Depends on the state that is captured in the system database(s) at time of block-based backup.

Online Revert

When performing an online revert, the instance as well as other databases will remain online and accessible to the user.

In addition, the user can choose if Microsoft SQL server should recover the database automatically to the point in time captured by the backup or if the user wants to apply transaction logs.

These benefits come with the additional restriction that the backup and the disks may only contain a single user database. When using online revert for multiple databases, each database must have its own set of disks. Furthermore, the databases must be backed up and reverted individually.

Reverting Clustered SQL Server instances

Protector supports reverting databases for instances using traditional clustered disks. It is **not** possible to revert Cluster Shared Volumes (CSV)

Prerequisites for reverting Microsoft SQL Server block-based backups

The following prerequisites must be met to revert a block-based Microsoft SQL Server database backup.

- A block-based backup of a Microsoft SQL Server is available.
- Credentials for the SQL Server instance are available. Please refer to [Microsoft SQL Server User Privileges \(on page 36\)](#) for more details.
- The SQL Server node which was used to create the backup is available and online.
- All nodes of the target instance are up and running.

Force Revert

Protector will not perform reverts in case it detects an enhanced data loss potential. If, for example, the snapshot contains multiple databases and only a single database was selected for backup, a regular revert will not be possible. Reverting the backup will revert the complete disk, which will not only affect the database selected for backup but the other databases sharing the same disks as well.

In these cases, Protector will not perform a revert unless the user specifies the force revert flag in the [SQL Server Revert Wizard \(on page 41\)](#). Once selected Protector will ignore all safety measures and attempt the revert.



WARNING: Forcing a revert is dangerous and has increased risk for data loss. Only use it when you understand why it is required and what the consequences will be. In most cases, it is a better option to mount the block-based backup and copy the data manually.

How to mount a database from a block-based backup and attach it to an existing instance.

Protector can mount block-based SQL Server backups to an existing SQL Server instance. This will make the disks comprising the backup available on the machines hosting the target instance and optionally attach the databases so the user can access them.

Before you begin

It is assumed that a SQL Server policy that created hardware snapshots or replications has been implemented and that at least one snapshot or replication has been created. See [How to create database restore points with block snapshots \(on page 16\)](#) for guidance on how to do this.

In addition, your target SQL Server instances must meet the requirements for mounting the instance. Please refer to “About mounting Microsoft SQL Server block-based backups” for details.

Note: Snapshots and replications cannot be used for mount operations if they are currently mounted elsewhere

This task describes the steps to follow when mounting a block-based SQL Server backup to an existing Microsoft SQL Server node

Procedure

1. Identify the destination where backup should be mounted. The destination Microsoft SQL Server environment must be represented by a Protector Microsoft SQL Server node. If the destination is not represented in Protector, then create a node using the [Microsoft SQL Server Node Wizard \(on page 29\)](#).
2. Ensure that the mount location is prepared by locating the node in the Nodes Inventory and checking it is authorized and online. In case of clustered SQL Server environments, ensure all cluster nodes are up and running.
3. Locate the block-based SQL Server backup by clicking the Restore link on the Navigation Sidebar to open the Restore Dashboard. Then use the search options to locate the backup you want to mount.
4. Select the **backup** and click *Mount* to open the [SQL Server Mount Wizard \(on page 37\)](#).
5. Select **SQL Mount** to perform an application-level mount. Click *Next*.
6. Select the target **Microsoft SQL Server Node** where the instance is located. Then select the instance and provide the credentials. Click *Next*.
7. Select which databases you want to attach to the instance. Click *Next*.
8. Review the selected databases and optionally specify a new name for each database. Click *Next*.



Note: Database names must be unique. If a database already exists on the target instance, the database will not be attached.

9. Choose which host groups the disks should be exposed to. Click *Next*.
10. Specify a mount location. Click *Next*.



Note: Which mount location options are available may differ depending on the type of SQL Server. Protector will disable options which are not applicable.

11. Select the mount mode and specify the Mount Pool if necessary, then click *Finish*.
The mount mode determines how the backup will be mounted. For replications, the mount mode is always set to Mount original.



Caution: Select the mount mode depending on the behavior required:

- **Mount Original** - Mounts the original snapshot. This option will expose your original backup to the Microsoft SQL Server. Any changes will persist even after the unmount. It may not be possible to mount or revert this backup in the future.
- **Mount duplicate** - Protector will create a cascaded duplicate of the original snapshot and mount the duplicate. The original snapshot is preserved. Use Cascade mode (the default setting) in the Block Snapshot Configuration Wizard when assigning the snapshot operation on the data flow to enable Mount duplicate.

A processing message will appear briefly, then the wizard will close, and the Jobs Inventory will be displayed. A new Job will appear at the top of the Jobs list, with the Progress entry initially indicating processing and finally showing successful completion.

12. Once the mount job completes successfully, the selected databases should be attached to the destination Microsoft SQL Server instance.

How to revert multiple databases from a block-based backup.

Protector can revert SQL Server environments to a state captured in a block-based SQL Server backup.

Before you begin

It is assumed that a SQL Server policy that created hardware snapshots or replications has been implemented and that at least one snapshot or replication has been created. See [How to create database restore points with block snapshots \(on page 16\)](#) for guidance on how to do this.

Procedure

1. Identify the SQL Node which was used to create the backup and locate it in the Nodes Inventory. Check the node is online and authorized.
2. Locate the block-based SQL Server backup by clicking the Restore link on the Navigation Sidebar to open the Restore Dashboard. Then use the search options to locate the backup you want to revert.
3. Select the **backup** and click *Revert* to open the [SQL Server Revert Wizard \(on page 41\)](#).
4. Select **offline** revert. Click *Next*.



Tip:

For more information regarding different revert types, please refer to [About reverting Microsoft SQL Server block-based backups \(on page 20\)](#).

5. Select which databases should be attached as part of the revert process. Click *Next*.
6. Check that there are no warnings. Click *Next*.



WARNING: Do **not** select the force revert option. Refer to [About reverting Microsoft SQL Server block-based backups \(on page 20\)](#) for more details about force revert.

7. To ensure the user does not accidentally perform a revert, you must type REVERT in uppercase, then click Finish.
A Processing message will appear briefly, then the wizard will close, and the Jobs Inventory will be displayed. A new Job will appear at the top of the Jobs list, with the Progress entry initially indicating processing and finally showing successful completion.
8. Once the revert job completes successfully, use the Microsoft SQL Server management tools to further use the reverted data.

How to revert a user database online without affecting other databases in the same instance.

Protector can revert SQL Server environments to a state captured in a block-based SQL Server backup.

Before you begin

It is assumed that a SQL Server policy that created hardware snapshots or replications has been implemented and that at least one snapshot or replication has been created. See [How to create database restore points with block snapshots \(on page 16\)](#) for guidance on how to do this.

In addition, your SQL Server instance must meet the requirements for an online revert. Please refer to [About reverting Microsoft SQL Server block-based backups \(on page 20\)](#) for details.

This task describes the steps to follow when performing an online revert of a single user database.

Procedure

1. Identify the SQL Node which was used to create the backup and locate it in the Nodes Inventory. Check the node is online and authorized.
2. Locate the block-based SQL Server backup by clicking the Restore link on the Navigation Sidebar to open the Restore Dashboard. Then use the search options to locate the backup you want to revert.
3. Select the **backup** and click *Revert* to open the [SQL Server Revert Wizard \(on page 41\)](#).
4. Select **online** revert. Then select auto-recover to bring the database online at the end the revert process. Click *Next*.



Tip:

For more information regarding different revert types, please refer to [About reverting Microsoft SQL Server block-based backups \(on page 20\)](#).

5. Check that there are no warnings. Click *Next*.



WARNING: Do **not** select the force revert option. Refer to [About reverting Microsoft SQL Server block-based backups \(on page 20\)](#) for more details about force revert.

6. To ensure the user does not accidentally perform a revert, you must type REVERT in uppercase, then click Finish.
A Processing message will appear briefly, then the wizard will close, and the Jobs Inventory will be displayed. A new Job will appear at the top of the Jobs list, with the Progress entry initially indicating processing and finally showing successful completion.
7. Once the revert job completes successfully, the user database will be attached and accessible.

How to revert a user database and prepare it to apply transaction logs.

Protector can revert SQL Server environments to a state captured in a block-based SQL Server backup.

Before you begin

It is assumed that a SQL Server policy that created hardware snapshots or replications has been implemented and that at least one snapshot or replication has been created. See [How to create database restore points with block snapshots \(on page 16\)](#) for guidance on how to do this.

In addition, your SQL Server instance must meet the requirements for an online revert. Please refer to [About reverting Microsoft SQL Server block-based backups \(on page 20\)](#) for details.

This task describes the steps to follow when performing an online revert of a single user database into a state which allows the user to apply transaction log backups manually.

Procedure

1. Identify the SQL Node which was used to create the backup and locate it in the Nodes Inventory. Check the node is online and authorized.
2. Locate the block-based SQL Server backup by clicking the Restore link on the Navigation Sidebar to open the Restore Dashboard. Then use the search options to locate the backup you want to revert.
3. Select the **backup** and click *Revert* to open the [SQL Server Revert Wizard \(on page 41\)](#).
4. Select **online** revert. Ensure auto-recover is **not** selected. Click *Next*.



Tip:

For more information regarding different revert types, please refer to [About reverting Microsoft SQL Server block-based backups \(on page 20\)](#).

5. Check that there are no warnings. Click *Next*.



WARNING: Do **not** select the force revert option. Refer to [About reverting Microsoft SQL Server block-based backups \(on page 20\)](#) for more details about force revert.

6. To ensure the user does not accidentally perform a revert, you must type REVERT in uppercase, then click Finish.
A Processing message will appear briefly, then the wizard will close, and the Jobs Inventory will be displayed. A new Job will appear at the top of the Jobs list, with the Progress entry initially indicating processing and finally showing successful completion.
7. Once the revert job completes successfully, the user database will be registered with the instance and await manual recovery. Use Microsoft's SQL management software or 3rd party tools to restore and apply transaction logs.

Chapter 4: Reference

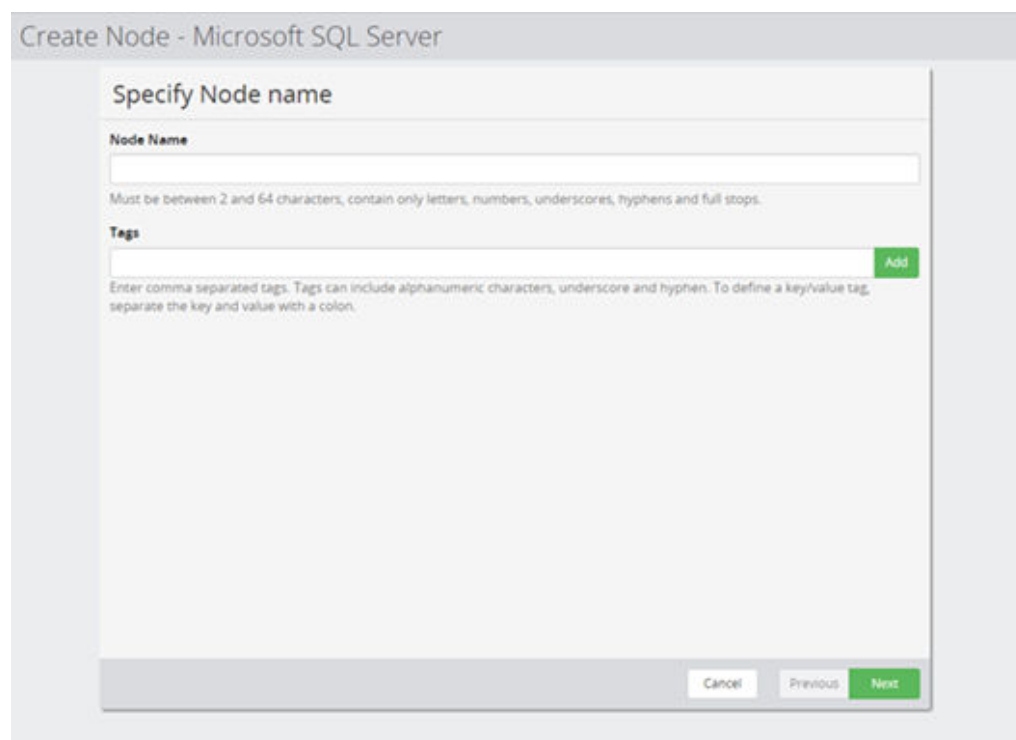
This section provides salient reference information that supports the workflows detailed in this guide.

Nodes UI Reference

This section describes the Nodes UI pertaining to the node types that are used to backup Microsoft SQL Server.

Microsoft SQL Server Node Wizard

Protector will launch this wizard when a new Microsoft SQL Server node is added to the Nodes Inventory.



The screenshot shows a wizard window titled "Create Node - Microsoft SQL Server". The current step is "Specify Node name". It features a text input field for "Node Name" with a placeholder "Node Name". Below the field is a note: "Must be between 2 and 64 characters, contain only letters, numbers, underscores, hyphens and full stops." There is also a "Tags" section with a text input field and a green "Add" button. Below the tags field is a note: "Enter comma separated tags. Tags can include alphanumeric characters, underscore and hyphen. To define a key/value tag, separate the key and value with a colon." At the bottom right of the wizard are three buttons: "Cancel", "Previous", and "Next".

Figure 2 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.

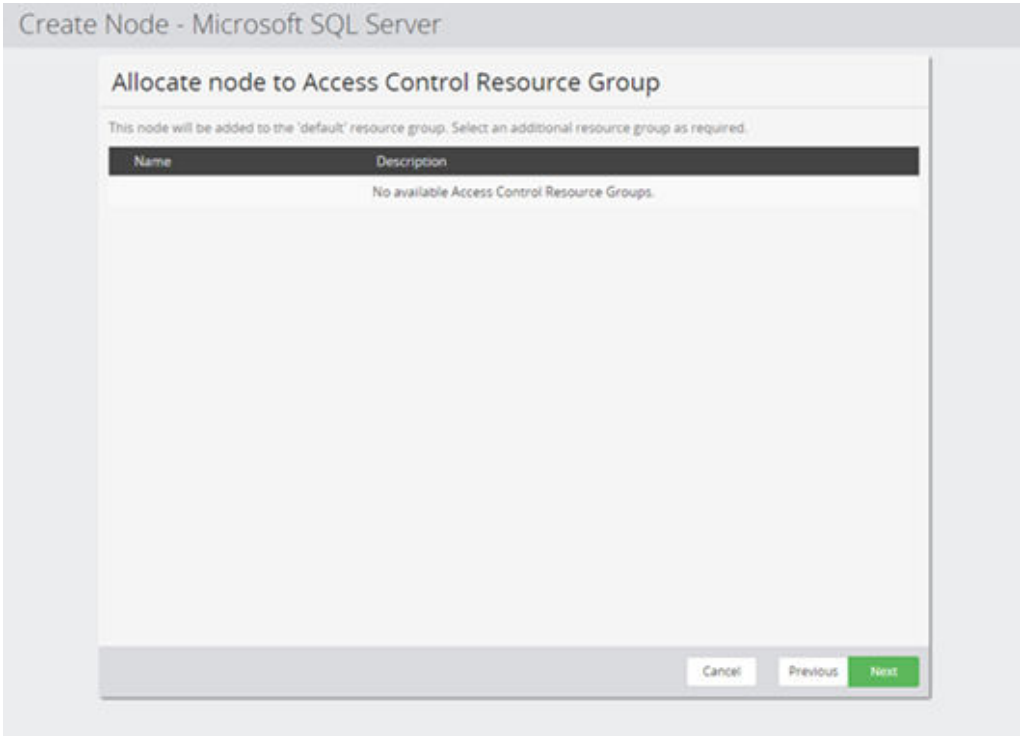


Figure 3 Microsoft SQL Node Wizard - Allocate node to Access Control Resource Group

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.

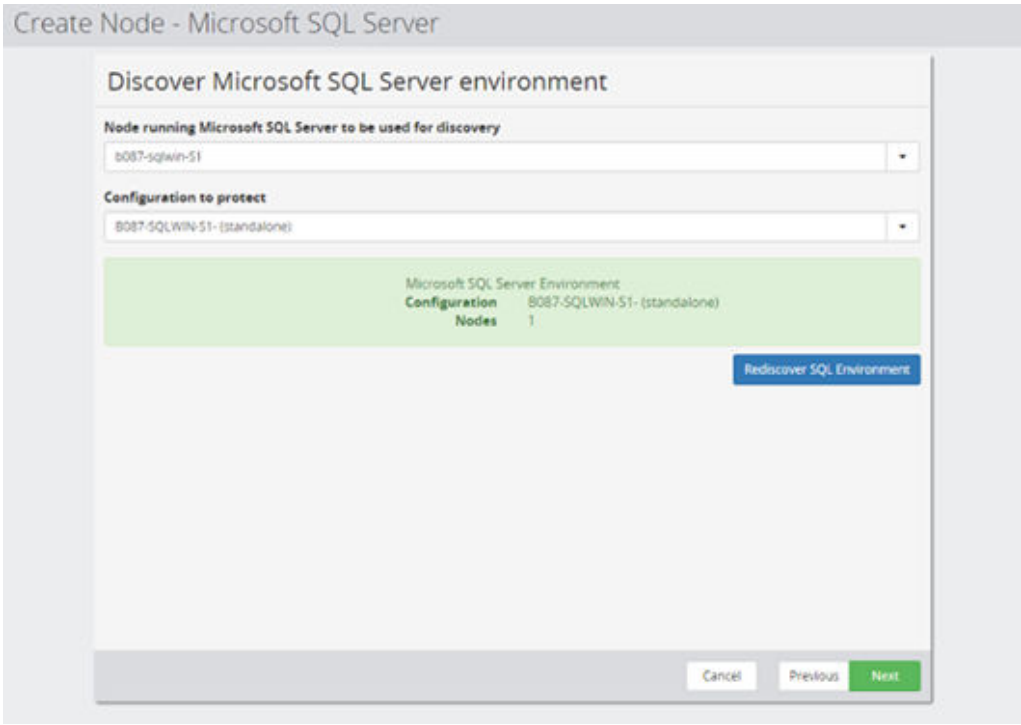


Figure 4 Microsoft SQL Node Wizard - Discover Microsoft SQL environment

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 5 Microsoft SQL Node Wizard - Summary

Control	Description
Summary	Summary of the settings and nodes selected in the node creation wizard.

Policies UI Reference

This section describes the Policies UI pertaining to the policies that are applied to backup Microsoft SQL Server.

Microsoft SQL Server Classification Wizard

This wizard is launched when a new MS SQL Server classification is added to a Policy.

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 \(on page 15\)](#) for details on how this classification works.

Figure 6 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.
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 35) SQL Server Instance Credentials Dialog.
Included Items	Lists the databases which will be included in the backup.
Add	Opens the Microsoft SQL Server Database Selection Wizard (on page 34) , which allows databases to be added to the Included Items list above.


Control	Description
Backup Mode	<p>Select which backup type is desired for the selected databases:</p> <ul style="list-style-type: none"> ▪ 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.

Microsoft SQL Server Database Selection Wizard

This wizard is displayed when the user chooses to include database in a Microsoft SQL Server classification.

Figure 7 Microsoft SQL Database Selection

Control	Description
Search	Enter a part of a database name and confirm to filter the list of known databases

Control	Description
Database List	<p>Select one or more databases</p> <div>  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. </div>

Microsoft SQL Server Instance Credentials Dialog

This dialog allows the user to specify which credentials should be used when protecting an instance

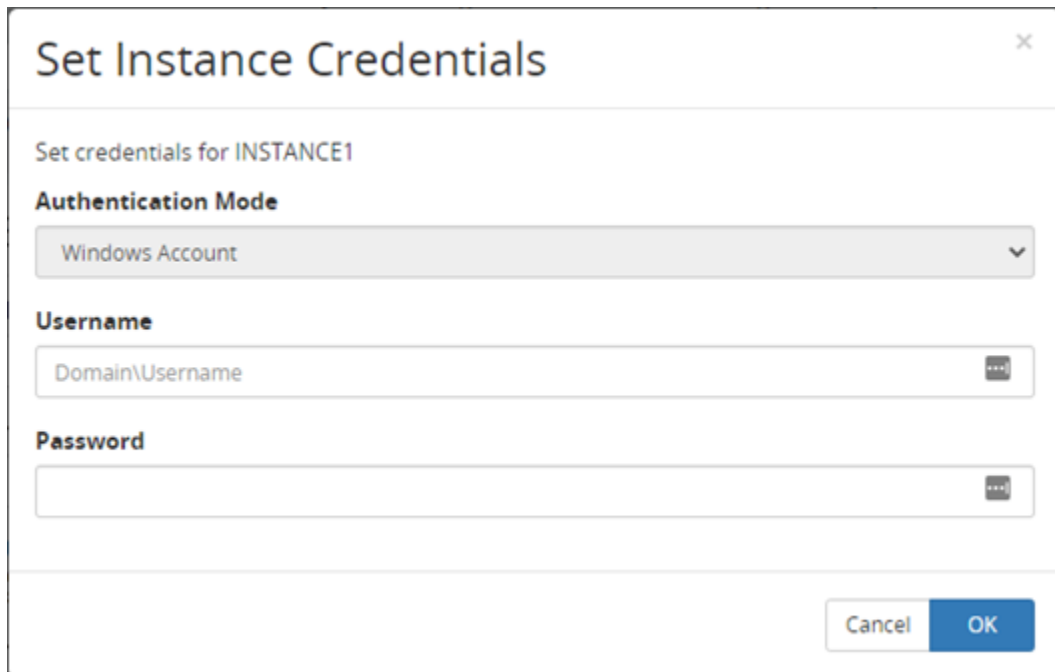


Figure 8 Microsoft SQL Server Instance Credentials Dialog

Control	Description
Authentication Mode	Lists the required account type.

Control	Description
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 36) 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

Restore UI Reference

This section describes the Restore UI pertaining to Microsoft SQL Server backups.

SQL Server Mount Wizard

When mounting block snapshots or replications created by a policy containing a Microsoft SQL Server classification, it is possible to choose between various levels of mount.

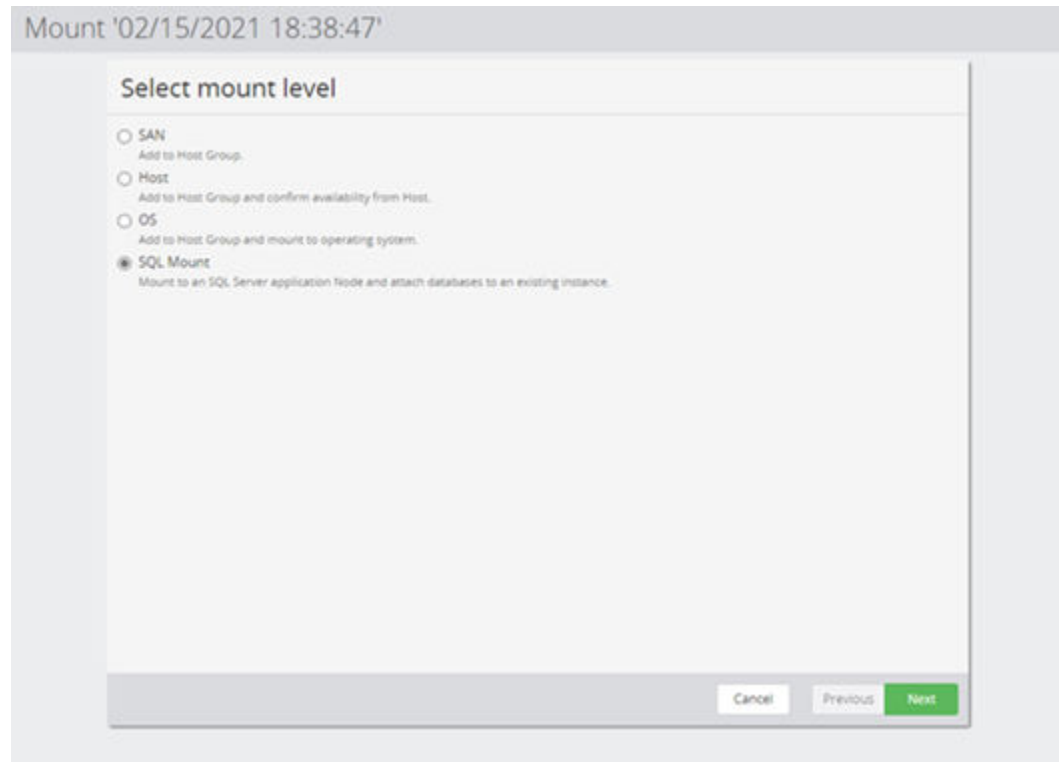


Figure 9 Mount Wizard - Select Mount Level

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.
SQL Mount	Adds the snapshot or replication to the Host Group(s) of a Microsoft SQL Server environment and attaches the database to an selected instance.

In case SQL Mount is selected, Protector will display the following additional wizard pages that allow application specific options to be configured:

Mount '02/15/2021 18:38:47'

Select Microsoft SQL target instance

Node

Instance

Windows Account for instance

Username

Password

Cancel Previous **Next**

Figure 10 Mount Wizard - Select target instance

Control	Description
Node	The SQL Server node representing the target Microsoft SQL Server environment.
Instance	The target instance the databases should be attached to.
Username	Username used to attach the database to the target instance. Refer to Microsoft SQL Server User Privileges (on page 36) for more details.
Password	Password for the user specified above.

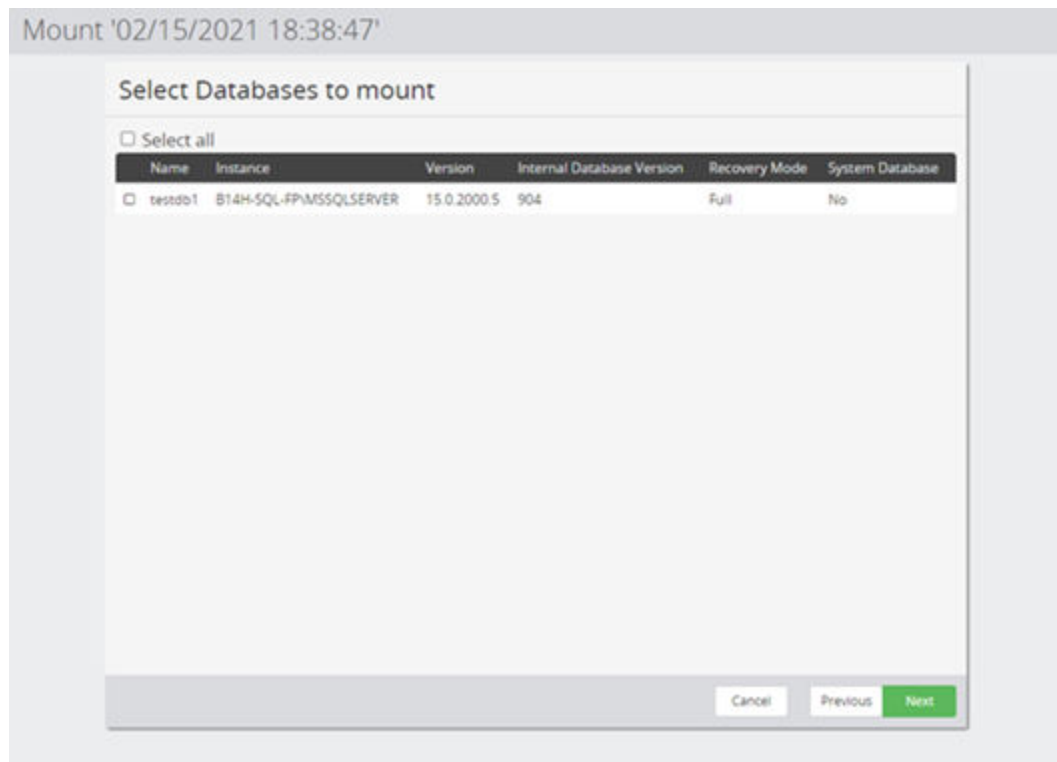


Figure 11 Mount Wizard - Select Database to Mount

Control	Description
Database list	Select which databases should be attached to the previously selected target instance. At least one database must be selected to proceed.

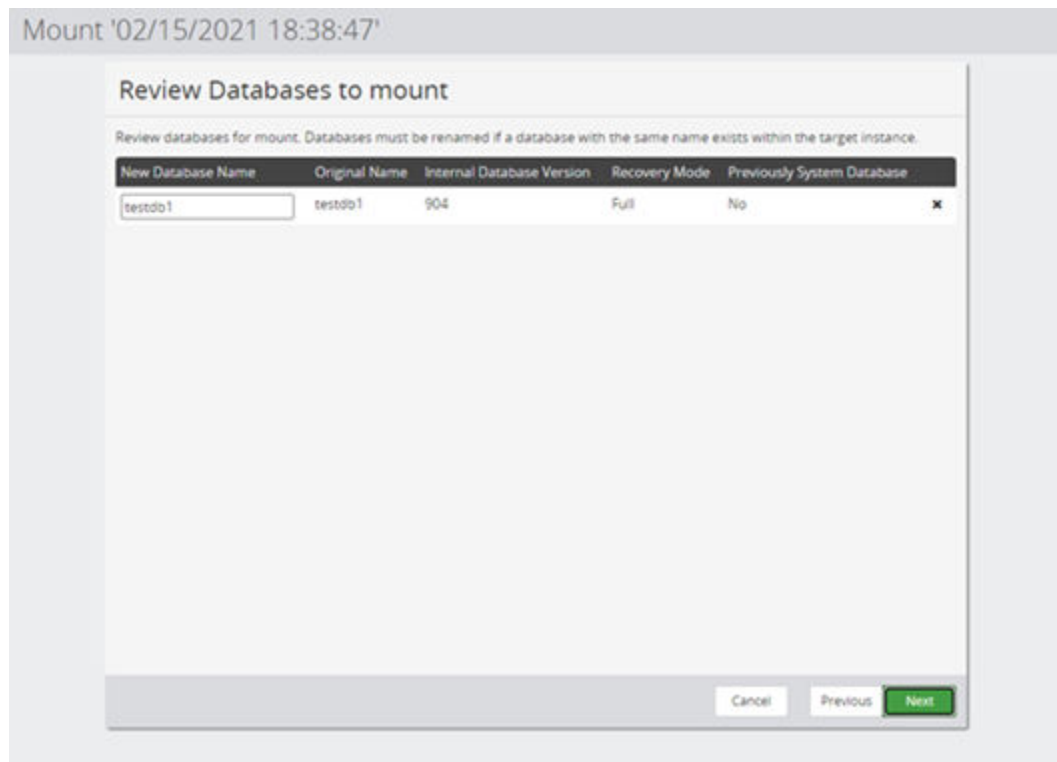



Figure 12 Mount Wizard - Review Database to Mount

Control	Description
Database list	Lists the previously selected databases for review.
New database name	<p>Defines the name of the mounted database on the target instance. The default value is the original name of the database, except for system database for which the Protector will suggest an alternate name.</p> <div> <p>Note: Database names must be unique on the target instance. Ensure you do not define a name which is already in use by the target instance.</p> </div>

SQL Server Revert Wizard

When reverting snapshots or local replications created by a policy containing a *Microsoft SQL Server* classification, Protector will first display the following wizard pages, allowing application specific options to be configured.

Figure 13 Microsoft SQL Revert Type

Control	Description
Revert Type	<p>Offline</p> <p>Performs an offline revert. This will take the instance offline and detach the reverted databases. All databases attached to this instance will be temporarily offline.</p> <p>It is possible to revert multiple system and/or user databases at the same time.</p> <p>Online</p> <p>Reverts user databases without affecting other databases attached to the target instance. To manually apply transaction logs, auto recovery must be skipped.</p> <div>  WARNING: Only possible if the block-based backup contains exactly a single user database. </div>

Control	Description
	Refer to About reverting Microsoft SQL Server block-based backups (on page 20) for more details on both options.
Perform auto-recover	<p>Enabled</p> <p>Selecting this will automatically recover the database(s) in the backup and bring them online. Once the database is recovered, it is not possible to apply transaction logs.</p> <p>Disabled</p> <p>De-Select if you want to manually roll forward to a point in time later than what is captured by the block-based backup. The database will not be accessible to users until transaction logs have been applied manually.</p>

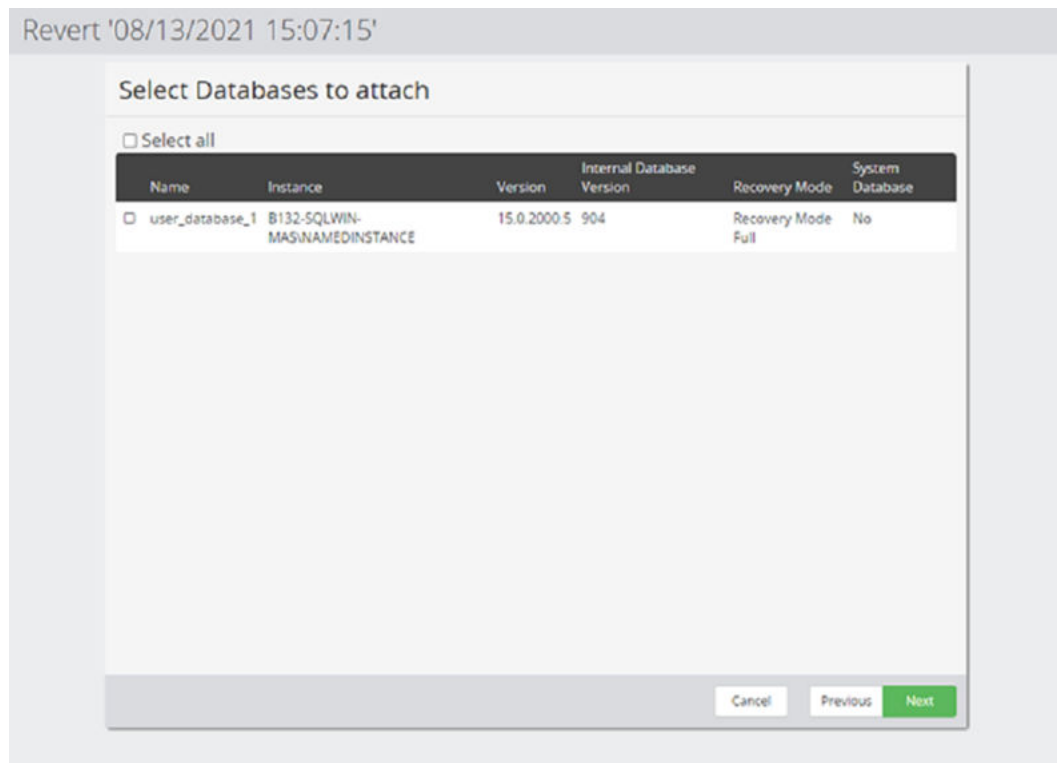


Figure 14 Microsoft SQL - Select Database to attach



Note: This step only applies when performing an offline revert of user databases.

Control	Description
Database list	Select which databases should be attached after revert completes.

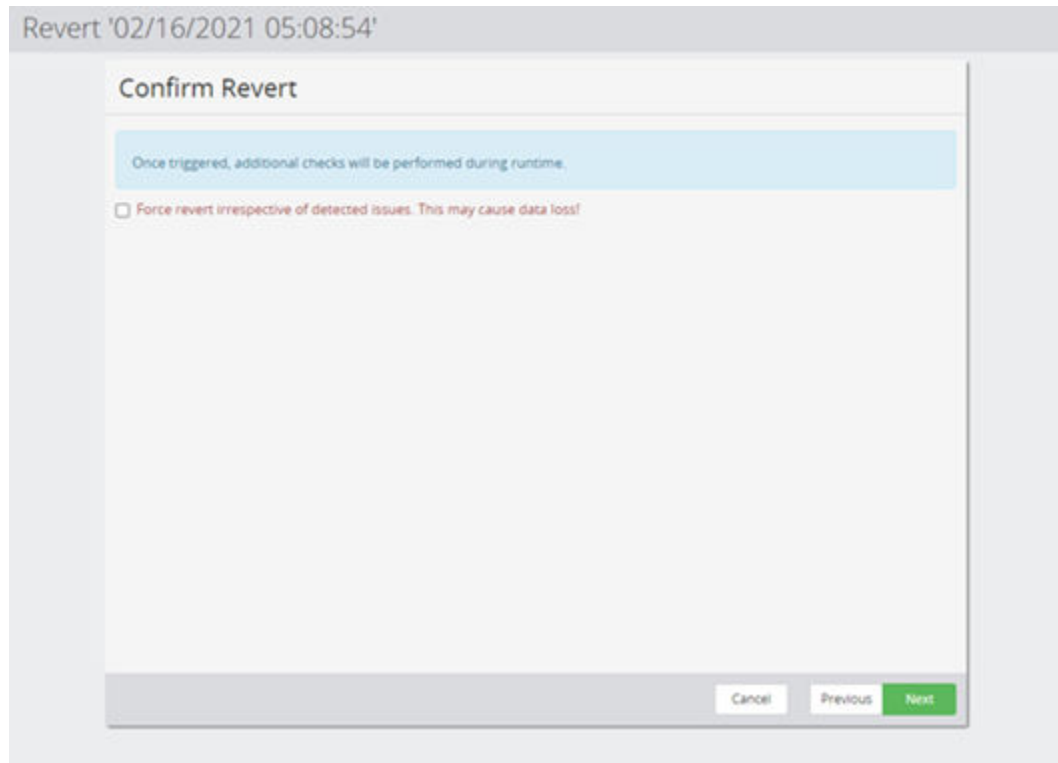



Figure 15 Microsoft SQL- Confirm Revert

Control	Description
Revert warnings	<p>This wizard page will list issues which will prevent Protector from performing a regular revert.</p> <p>In case any issues are listed here, it is recommended to mount the database instead of reverting it and restore the database manually from the mounted volume.</p>

Control	Description
	If any issues are found, it will only be possible to continue by selecting the force flag.
Force revert option	<p>This flag will instruct Protector to ignore problems preventing a regular revert.</p> <div> 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 database manually. Selecting this may cause data loss on the source system.</div>

Chapter 5: Troubleshooting

This chapter provides guidelines for how to troubleshoot issues that might occur when using Ops Center Protector.

Troubleshooting Microsoft SQL Server

This section provides guidelines for how to troubleshoot issues that might occur when using Microsoft SQL Server.

Cannot trigger a revert or revert fails without a force revert option.

Problem:

- It is not possible to complete the SQL Server Revert Wizard without selecting the “Force revert” option.
- A triggered revert fails with an error stating that a “Force revert” is required to continue.

Cause:

Your environment or block-based backup does not meet the requirements for normal revert.

Solution:

Consider a mount and manual copy as an alternative to the revert.

If you are sure about the consequences, you can trigger the revert with the force option selected. Depending on the situation, this may cause data loss!

If you are unsure on how to proceed, please open a support case.

Clustered SQL Server instances does not list all nodes.

Problem:

When creating a SQL Server node for a clustered SQL Server instance, the amount of nodes does not match cluster nodes in the setup.

Cause:

Protector will only consider nodes which are configured as a possible owner for the SQL Server instance

Database(s) are not listed in Microsoft SQL Management Studio after offline revert.

Solution:

Use Microsoft's Failover Cluster Manager to locate the cluster role for the instance and reconfigure the list possible owners.

Database(s) are not listed in Microsoft SQL Management Studio after offline revert.

Problem:

After a successful offline revert, the reverted databases are not listed in MSSQL Management Studio.

Cause:

The database was not selected to be attached when triggering the revert.

Solution:

Attach the reverted databases using Microsoft SQL Management Studio.

Databases are not listed for when creating a SQL Server classification.

Problem:

When adding databases to a classification, one or more recently created databases are not listed.

Cause:

The database information on the master might be out of date and not contain the databases, yet.

Solution:

Use the "Refresh Databases" button of the [Microsoft SQL Server Database Selection Wizard \(on page 34\)](#) to update the information.

Alternatively, the database and instance information can also be refreshed on the application node.

Invalid Credentials; Unable to connect to instance.

Problem:

During node or SQL Server classification creation or another operation Protector reports invalid credentials.

Cause:

The provided or stored credentials are invalid or out of date.

Solution for node creation, mount or revert configuration:

Provide valid credentials in the wizard.

Solution for backup:

When the issue is encountered during a backup, the credentials are out of date. Update the SQL Server node and redistribute the data flows utilizing this node.

Issues with snapshots of mounted databases.

Problem:

Databases which are mounted via protector are not included in the backup or require manual steps to unmount after a snapshot has been created.

Cause:

Including databases on block-based backups into a new block-based backup is not supported.

Solution:

Do not create block-based backups of block-based backups.

Mount Error - Database(s) already exists in instance.

Problem:

While mounting a database Protector reports that the database(s) already exist in the target instance.

Cause:

A database name must be unique within an instance. It is not possible to mount a database with a name that already exists.

Solution:

Review the attachment to the log message to determine which databases are affected. Retry the mount and define an alternate name for the database in the [SQL Server Mount Wizard \(on page 37\)](#).

No cluster nodes returned from cluster during node creation.

Problem:

When attempting to create a SQL Server Node for a clustered instance, the error “No cluster nodes returned from cluster” is displayed.

Cause:

Protector requires the cluster roles to follow the default naming schema of Microsoft SQL Server and Failover Cluster.

Solution:

Ensure the names of all cluster roles and resources are the default and none have been renamed.

Online Revert option is disabled.

Problem:

When triggering a revert, the online revert option is grayed-out.

Cause:

The block-based backup does not meet the requirements for an online revert.

Solution:

Either perform an offline revert or reconfigure your SQL Server environment so it meets the requirements for online revert and create a new backup. Refer to [About reverting Microsoft SQL Server block-based backups \(on page 20\)](#) for more details.

Unable to apply transaction log backups after successful online revert.

Problem:

After a successful online revert, it is not possible to apply transaction log backups.

Cause:

If a valid transaction log backup for the database which covers the required log sequence number (LSN) is available, this is usually caused by enabling auto-recover during the revert.

Solution:

Check that auto-recover is unchecked during revert, and the correct transaction logs are available.

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.

Continuous Data Protection (CDP)

A method of capturing the state of a file system in near real time. CDP shares much of the functionality of Live Backup, except that RPO is measured in minutes, data is retained for a much shorter period of time and is not indexed by the MDS. Typically, CDP and Live Backup are used in conjunction. CDP is only supported on source nodes running the Microsoft Windows operating system.

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.

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.

Live backup

A backup technique that avoids the need for bulk data transfers by continuously updating the repository with changes to the source file system. This is similar to CDP but with longer retention periods and RPOs being available. Live backups perform byte level change updates whereas batch backups perform block level change updates.

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.

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.

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.

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.

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.

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.

Hitachi Vantara



Corporate Headquarters

2535 Augustine Drive

Santa Clara, CA 95054 USA

HitachiVantara.com | community.HitachiVantara.com

Contact Information

USA: 1-800-446-0744

Global: 1-858-547-4526

HitachiVantara.com/contact