

Hitachi Data Instance Director

SAP HANA Application Guide

2018 Hitachi, Ltd. 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 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 Corporation 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 Corporation.

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 is a registered trademark of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, IBM, Lotus, MVS, OS/390, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, 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, the Microsoft Corporate 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.

Contents

Preface.....	6
Software version.....	6
Intended audience.....	6
Related documents.....	6
Document conventions.....	7
Conventions for storage capacity values.....	8
Accessing product documentation.....	9
Getting help.....	10
Comments.....	10
Chapter 1: Before you begin.....	11
Supported configurations.....	11
Prerequisites.....	11
Application software prerequisites.....	12
Hitachi Block prerequisites.....	12
Upgrading to Hitachi Data Instance Director 6.x.....	14
Differences between HDID 5.x and HDID 6.x.....	14
Chapter 2: Backup workflows.....	16
Block based workflows.....	16
How to create restore points with local snapshots.....	16
How to create restore points and full backups with snapshots of a local clone.....	20
How to create restore points and DR backups with snapshots of a remote clone.....	23
Chapter 3: Restore workflows.....	29
How to mount a snapshot or clone for repurposing.....	29
How to revert from a block snapshot or clone.....	30
Chapter 4: Reference.....	32
Nodes UI Reference.....	32
SAP HANA Application Node Wizard.....	32
Policies UI Reference.....	37
SAP HANA Database Classification Wizard.....	37

Chapter 5: Troubleshooting..... 40
 Troubleshooting SAP HANA..... 40

Glossary..... 41

Preface

This guide describes how to backup and restore SAP HANA databases using Hitachi Data Instance Director.

Data Instance Director orchestrates the creation, retention and restoration of application-consistent and crash consistent snapshots and clones for SAP HANA databases. Application data can be protected by creating snapshots or clones on Hitachi Block or NAS Storage. 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 Data Instance Director version 6.7. Please refer to the accompanying Release Notes (RN-93HDID018) for information on what's changed in this release.

Intended audience

This document is intended for database administrators who wants to protect SAP HANA databases using Hitachi Data Instance Director. It is assumed that the reader has a good working knowledge SAP HANA, Hitachi Block Storage administration and network administration.

If you are new to Data Instance Director, we recommend that you start by referring to the *Hitachi Data Instance Director User's Guide*, MK-93HDID014 so that you understand the basic concepts, workflows and user interface.

Related documents

Main product guides:

- *Hitachi Data Instance Director Software Release Notes*, RN-93HDID018
- *Hitachi Data Instance Director User's Guide*, MK-93HDID014
- *Hitachi Data Instance Director Quick Start Guide*, MK-93HDID015
- *Hitachi Data Instance Director Microsoft Exchange Server Application Guide*, MK-93HDID012
- *Hitachi Data Instance Director Microsoft SQL Server Application Guide*, MK-93HDID011
- *Hitachi Data Instance Director Oracle Application Guide*, MK-93HDID010
- *Hitachi Data Instance Director SAP HANA Application Guide*, MK-93HDID017
- *Hitachi Data Instance Director VMware Application Guide*, MK-93HDID022

Programming guides:

- *Hitachi Data Instance Director REST API User Guide*, MK-93HDID019
- *Hitachi Data Instance Director REST API Reference Guide*, MK-93HDID020
- *Hitachi Data Instance Director REST API Change Log*, MK-93HDID021





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: <code>pairedisplay -g group</code> <p>(For exceptions to this convention for variables, see the entry for angle brackets.)</p>
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 Hitachi Vantara Support Connect: <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

[Hitachi Vantara Support Connect](#) is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to Hitachi Vantara Support Connect 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 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 Corporation.

Thank you!

Chapter 1: Before you begin

Supported configurations

Data Instance Director supports the following SAP HANA configurations:

- SAP HANA 1.0 SPS 10 and newer.
- SAP HANA 2.0 SPS 00 and newer.
- Scale-up single node environments.
- MCOS (multiple components on one system).
- MDC (multi-tenant database containers) with only one database.

However, it does not currently support the following configurations:

- Scale-out cluster environments.
- System replication (application mirror).

Data Instance Director protects the following parts of an SAP HANA installation:

- All data LUNs for the specified SAP HANA instance.
- Configuration information listed in the *m_inifiles* table, including default and custom configuration files.

It does not protect the following:

- Applications running on top of SAP HANA.
- SAP HANA binaries.
- The log area (explicitly not supported by SAP - the data area is consistent without logs).
- The backup catalog (snapshot recovery is possible without the catalog).

Prerequisites

It is important that the following prerequisites are met before you attempt to implement any of the SAP HANA 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/data-instance-director.html#tech-specifications>.

For detailed information on installing the Data Instance Director Master, and Client components, please refer to the *Hitachi Data Instance Director User's Guide*, MK-93HDID014 .

Application software prerequisites

Before HDID can interact with SAP HANA to protect its data, ensure that:

- Hitachi block storage hosting the data LUNs has been configured.
- Each database is on its own set of disks.
- Each database has its data area on separate disks to its transaction logs.
- The database or default *system* user credentials are known and the following minimum roles are assigned:
 - BACKUP_OPERATOR
 - CATALOG_READ
- HDID Client components are installed on the SAP HANA server.

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 HDID support matrices at <https://www.hitachivantara.com/en-us/products/data-protection/data-instance-director.html#tech-specifications>:

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



Caution: ISM nodes and their associated CMDs used to control storage devices must not be shared with other applications.

- 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 UR journals must be set up.
- For GAD, the P-VOLs must be set up in a host group and quorum disks must be provided. Best practice is to provide one quorum disk per replication pair.
- For GAD, virtualized LDEVs are required (available on VSP G series only)
- Port security must be enabled on VSP G1000

- Primary volumes must be set up using Storage Navigator (or other tools) prior to selection in HDID
- For application consistent snapshots, the application must be installed and configured to use Hitachi Vantara block P-VOLs
- The password for authorizing a VSP node must contain only useable RAID Manager 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 Data Instance Director:
 - For standard mode (non-cascading) Thin Image (TI) the TI Pools must be set up
 - For cascade mode Thin Image (TI) the Dynamic Provisioning Pools must be set up
 - For ShadowImage (SI), TrueCopy (TC), Universal Replicator (UR) and Global-Active Device 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 HDID ISM node controlling the block storage device must have:
 - The correct version of Hitachi CCI installed.
 - Access to a dedicated Command Device (CMD) on the storage device, set up as follows:
 - Security disabled
 - User authentication enabled
 - Device group definition disabled
 - The CMD must be visible to the host OS where the HDID proxy resides
 - The CMD must be offline
 - Multiple active command devices may be visible to an HDID proxy as long as each one represents a different block storage device. Behavior is undefined if multiple active command devices represent the same block storage device, unless these are configured in the HDID 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 HDID must be created on the storage device with at least the following roles:
 - Storage Administrator (Provisioning)
 - Storage Administrator (Local Copy)
 - Storage Administrator (Remote Copy)

For Global-Active Device, the following role must be added:

 - Security Administrator (View & Modify).

Upgrading to Hitachi Data Instance Director 6.x

This section addresses the differences between SAP HANA application policies in HDID 6.x and legacy HDID versions:



Note: To use the new mount and revert options in HDID 6.x, existing data flows must be modified and reactivated.

Differences between HDID 5.x and HDID 6.x

The following changes have occurred between Hitachi Data Instance Director 5.x and Hitachi Data Instance Director 6.x:

- When mounting snapshots that were created in HDID 5.x, some of the wizard options that are available in HDID 6.x will not be displayed. However, the mount operation will complete successfully.



Note: After upgrading from HDID 5.x to 6.x, new data flows should be created in order to access the new mount options.

Chapter 2: 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 HDID user interface, please refer to *Hitachi Data Instance Director User's Guide*, MK-93HDID014.

Be aware that when performing backups on a SAP HANA database:

- It is only possible to back up a database when it is online. The database will remain online during the backup and after it has completed.
- As part of the backup process, HDID updates the backup catalog with the following information:
 - In progress snapshot backups.
 - Successful snapshot backups.
 - Failed snapshot backups, but only if the HDID backup fails after SAP HANA was quiesced by creating a save point.
 - Reference to the HDID backup ID.

The creation of an additional backup save point may cause a slight performance impact, but this is no different than that incurred by the default SAP HANA save point created every 5 minutes.

Block based workflows

This section addresses the workflows for block based backups.

How to create restore points with local snapshots

Before you begin

It is assumed that the following tasks have been performed:

- The SAP HANA application has been installed and any HDID prerequisites are met.
- The HDID Master software has been installed and licensed on a dedicated node.
- The HDID Client software has been installed on the source node where the SAP HANA application resides.
- The HDID 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 HDID requirements and prerequisites.
- Permissions have been granted to enable the HDID 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.

Data Instance Director enables you to create a point-in-time snapshot of the SAP HANA database by storing the changes instead of copying the whole database. The snapshot is created using Hitachi Thin Image technology. By creating a Thin Image snapshot, you can not only manage the storage space efficiently but also rapidly recover the database to a previous point in time.

Because Thin Image is differential, the primary volumes are required to reconstruct the entire data set, therefore if the primary data is lost then the snapshots are of no use. For this reason, Thin Image snapshots should not be relied upon for recovery from catastrophic primary data loss.

The data flow and policy are as follows:



Figure 1 Hardware Snapshot Data Flow

Table 1 SAP HANA Snapshot Policy

Classification Type	Parameters	Value
SAP HANA Database	Database Selection	TestDb (The selected databases must be located on the same Block device)

Operation Type	Parameters	Value	Assigned Nodes
Snapshot	Mode	Hardware	SAP HANA Database
	Hardware Type	Hitachi Block	
	RPO	8 Hours	
	Retention	1 Week	
	Run Options	Run on RPO	
	Source Options	Quiesce...	

Procedure

1. Locate the source *OS Host* node in the **Nodes Inventory** and check that it is authorized and online.
This node represents the HDID Client installed on the SAP HANA server.
2. Create a new *SAP HANA Database* node using the [SAP HANA Application Node Wizard \(on page 32\)](#) and check that the node it is authorized and online.
The *SAP HANA Database* node type is grouped under **Application** in the **Node Type Wizard**. This node will be used in the dataflow to represent the SAP HANA Database configuration to be protected.
 - a. Select the *OS Host* node identified above as the **Node running SAP HANA....**
 - b. Specify the **DB Credentials** for each the **Discovered Databases**, by clicking on the **Name** of each database in the table that are to be protected and that display the message `Password Required`.


3. Locate the node in the **Node Inventory** 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 HDID 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 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 the 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**, [SAP HANA Database Classification Wizard \(on page 37\)](#) and **Snapshot Operation Wizard**.
The *SAP HANA Database* classification is grouped under **Application** in the **Policy Wizard**.

6. Draw a data flow as shown in the figure above, that shows only the *SAP HANA Database* source node.

At this stage the snapshot icon  is not shown.

7. Assign the *Snapshot* operation to the *SAP HANA Database* source node. The *SAP HANA-Snapshot* policy will then be assigned automatically.
The **Hitachi Block Snapshot Operation Properties Dialog** is displayed.
8. Select the **Pool** by selecting the *Hitachi Block Device* node created in the steps above, followed by one of the available *Thin Image Pools*.
9. Leave the remaining parameters at their default settings, then click **OK**.
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 or warnings.
11. Locate the active data flow in the **Monitor Inventory** and open its **Monitor Details** 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.

12. Monitor the active data flow 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 snapshot.
13. Review the status of the *Hitachi Block Device* to ensure snapshots are being created. New snapshots will appear 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 create restore points and full backups with snapshots of a local clone

Before you begin

It is assumed that the following tasks have been performed:

- The SAP HANA Database application has been installed and any HDID prerequisites are met.
- The HDID Master software has been installed and licensed on a dedicated node.
- The HDID Client software has been installed on the source node where the SAP HANA Database application resides.
- The HDID Client software has been installed on the node that will act as a proxy for the Hitachi Block storage device. Note that for a Shadow Image replication, the source and destination LDEVs are located on the same device.
- The storage device has been set up as per the HDID requirements and prerequisites.
- Permissions have been granted to enable the HDID 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.

Snapshot of a local clone enables both rapid recovery to a point in time by using Thin Image snapshots, while providing an additional level of protection by creating a full clone of the database by using Hitachi ShadowImage technology. Taking snapshots of the clone adds the additional benefit of being able to roll back the backup copy to a given restore point.

Because Shadow Image is an in-system replication technology, it does not provide protection against a disaster at the local site, since both the primary and secondary volumes are co-located.

The data flow and policy are as follows:




Figure 2 Shadow Image Replication with Local Thin Image Snapshots Data Flow

Table 2 SAP HANA Replication/Snapshot Policy

Classification Type	Parameters	Value
SAP HANA Database	Database Selection	TestDb (All the selected databases must be located on the same Block device)


Operation Type	Parameter	Value	Assigned Nodes
Replicate	Run Options	Run on Schedule (see synch group schedule below)	Hitachi Block Device
	Source Options	Quiesce...	
Snapshot	Mode	Hardware	Hitachi Block Device
	Hardware Type	Hitachi Block	
	RPO	8 Hours	
	Retention	1 Week	
	Run Options	Run on Schedule (see synch group schedule below)	
	Source Options	Quiesce...	

Table 3 Synchronization Group Schedule

Schedule Item Type	Parameter	Value	Policy Operations
Trigger	N/A (this schedule defines a synchronization group name for local replications and snapshots. All parameters are ignored.)	N/A	Snapshot, Replication

Procedure

1. Locate the source *OS Host* node in the **Nodes Inventory** and check that it is authorized and online.
This node represents the HDID Client installed on the SAP HANA server.
2. Create a new *SAP HANA Database* node using the [SAP HANA Application Node Wizard \(on page 32\)](#) and check that the node is authorized and online.
The *SAP HANA Database* node type is grouped under **Application** in the **Node Type Wizard**. This node will be used in the dataflow to represent the SAP HANA Database configuration to be protected.
 - a. Select the *OS Host* node identified above as the **Node running SAP HANA....**
 - b. Specify the **DB Credentials** for each the **Discovered Databases**, by clicking on the **Name** of each database in the table that are to be protected and that display the message `Password Required`.
3. Locate the node in the **Nodes Inventory** that will control the Hitachi Block Devices via a CMD (Command Device) interface and check that it is authorized and online.
This node is used by HDID to orchestrate replication of the LDEV 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 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**. The Hitachi Block Device node appears in the replication data flow as the destination node.
5. Define a policy as shown in the table above using the **Policy Wizard**. This policy contains operations for the local replication and snapshot.
 - a. Define a *SAP HANA Database* classification using the [SAP HANA Database Classification Wizard \(on page 37\)](#).
The *SAP HANA Database* classification is grouped under **Application** in the **Policy Wizard**.
 - b. Define a *Replicate* operation using the **Replicate Operation Wizard**.
Shadow Image replication runs as a batch operation triggered by the RPO of the snapshot.
 - c. Define a local *Snapshot* operation using the **Snapshot Operation Wizard**.
Thin Image snapshots run based on the RPO. However we also want to synchronize the snapshot with the replication. This is done by defining a trigger schedule that is applied to both the snapshot and replication operations.
 - d. Define a *Trigger* schedule using the **Schedule Wizard**; accessed by clicking on **Manage Schedules** in the **Snapshot Operation Wizard** for the local snapshot.
Only the trigger schedule name is required; the parameters are not relevant here since the RPO of the snapshot dictates when the replication operation is triggered.
6. Draw a data flow as shown in the figure above, that shows the *SAP HANA Database* source node connected to the *Hitachi Block Device* via a *Batch* mover.
7. Assign the *SAP HANA-Replicate-Snapshot* policy to the *SAP HANA Database* source node.

8. Assign *SAP HANA-Replicate-Snapshot* policy's *Snapshot* operation to the *Hitachi Block Device* destination node.
The **Hitachi Block Snapshot Operation Properties Dialog** is displayed.
9. Select the **Pool** by selecting the local *Hitachi Block Device* node created in the steps above, followed by one of the available *Thin Image Pools*.
10. Leave the remaining snapshot parameters at their default settings, then click **OK**.
The snapshot icon  is now shown superimposed over the source node.
11. Assign the *Replicate* operation to the *Hitachi Block Device* node.
The **Hitachi Block Replication Operation Properties Dialog** is displayed.
12. 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**.
13. Compile and activate the data flow, checking carefully that there are no errors or warnings.
14. Locate the active data flow in the **Monitor Inventory** and open its **Monitor Details** page.
The policy will be invoked automatically to create and then maintain the replication according to the policy. Snapshot and replication operations will be triggered synchronously on the source node according to the RPO.
15. Monitor the active data flow to ensure the policy is operating as expected.
For a healthy data flow you will periodically see:
 - Replication and snapshot jobs appearing for the source node in the **Jobs** area triggered according to the RPO.
 - 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.
16. Review the status of the *Hitachi Block Device* to ensure snapshots and replications are being created.
New snapshots and a refreshed replication will appear 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 create restore points and DR backups with snapshots of a remote clone

Before you begin

It is assumed that the following tasks have been performed:

- The SAP HANA Database application has been installed and any HDID prerequisites are met.
- The HDID Master software has been installed and licensed on a dedicated node.
- The HDID Client software has been installed on the source node where the SAP HANA Database application resides.
- The HDID Client software has been installed on the nodes that will act as a proxy for both the primary and secondary Hitachi Block storage devices. Note that for a TrueCopy replication, the source and destination LDEVs are located on different devices.
- The storage devices have been set up as per the HDID requirements and prerequisites.
- Permissions have been granted to enable the HDID 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.

Snapshot of a remote clone enables both rapid recovery to a point in time by using Thin Image snapshots, while providing an additional level of protection by creating a full remote clone of the database using Hitachi TrueCopy technology. In synchronous replication, the storage system signals each write completion only once it is performed on the primary and secondary volume (copy on write).

This setup provides partial protection against a disaster at the local site and full protection at the remote site as the primary and secondary volumes are geographically separated. If necessary, production can be moved quickly to the remote site while the local site is being recovered. Taking snapshots of the remote clone adds the additional benefit of being able to roll back the backup copy to a given restore point from the remote site.

The data flow and policy are as follows:




Figure 3 TrueCopy Replication with Local and Remote Thin Image Snapshots Data Flow

Table 4 SAP HANA Replication/Snapshot Policy

Classification Type	Parameters	Value
SAP HANA Database	Database Selection	TestDb (All the selected databases must be located on the same Block device)

Operation Type	Parameter	Value	Assigned Nodes
Replicate	Run Options	N/A (TrueCopy is a continuous replication, so the Run option is ignored)	Secondary Hitachi Block Device
	Source Options	Quiesce...	
Snapshot (on local device)	Mode	Hardware	SAP HANA Database
	Hardware Type	Hitachi Block	
	RPO	8 Hours	
	Retention	1 Week	
	Run Options	Run on Schedule (see synch group schedule below)	
	Source Options	Quiesce...	
Snapshot (on remote device)	Mode	Hardware	Secondary Hitachi Block Device
	Hardware Type	Hitachi Block	
	RPO	8 hours (this must match the local snapshot)	
	Retention	1 Week (this can differ from the local snapshot)	
	Run Options	Run on Schedule (see synch group schedule below)	

Table 5 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 *OS Host* node in the **Nodes Inventory** and check that it is authorized and online.
This node represents the HDID Client installed on the SAP HANA server.
2. Create a new *SAP HANA Database* node using the [SAP HANA Application Node Wizard \(on page 32\)](#) and check that the node it is authorized and online.
The *SAP HANA Database* node type is grouped under **Application** in the **Node Type Wizard**. This node will be used in the dataflow to represent the SAP HANA Database configuration to be protected.
 - a. Select the *OS Host* node identified above as the **Node running SAP HANA....**
 - b. Specify the **DB Credentials** for each the **Discovered Databases**, by clicking on the **Name** of each database in the table that are to be protected and that display the message *Password Required*.
3. Locate the nodes in the **Nodes Inventory** 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 HDID 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 (Intelligent Storage Manager) nodes. The ISM nodes do not appear in the data flow.
4. Create new primary and secondary *Hitachi Block Device* nodes (unless ones already exists) using the **Hitachi Block Storage Node Wizard** and check that they are authorized and online.
The *Hitachi Block Device* node type is grouped under **Storage** in the **Node Type Wizard**. 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 *SAP HANA Database* node where the primary LDEV is mounted.
5. Define a policy as shown in the table above using the **Policy Wizard**. This policy contains operations for the replication, local and remote snapshots.

- a. Define an *SAP HANA Database* classification using the [SAP HANA Database Classification Wizard \(on page 37\)](#).
The *SAP HANA Database* classification is grouped under **Application** in the **Policy Wizard**.
- b. Define a *Replicate* operation using the **Replicate Operation Wizard**.
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**.
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**; accessed by clicking on **Manage Schedules** in the **Snapshot Operation Wizard** 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.
- e. Define a remote *Snapshot* operation using the **Snapshot Operation Wizard**.
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.

6. Draw a data flow as shown in the figure above, that shows the *SAP HANA Database* 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 (SVOL) volume is located.
7. Assign the *SAP HANA-Replicate-Snapshot-Snapshot* policy to the *SAP HANA Database* source node.
8. Assign the local *Snapshot* operation to the *SAP HANA Database* source node.
The **Hitachi Block Snapshot Operation Properties Dialog** is displayed.
9. Select the **Pool** by selecting the local *Hitachi Block Device* node created in the steps above, followed by one of the available *Thin Image Pools*.
10. Leave the remaining snapshot parameters at their default settings, then click **OK**.
The snapshot icon  is now shown superimposed over the source node.
11. Assign the remote *Snapshot* operation to the remote *Hitachi Block Device* node.
The **Hitachi Block Snapshot Operation Properties Dialog** is displayed.
12. Select the **Pool** by selecting the remote *Hitachi Block Device* node created in the steps above, followed by one of the available *Thin Image Pools*.
13. Leave the remaining snapshot parameters at their default settings, then click **OK**.
The snapshot icon  is now shown superimposed over the destination node.
14. Assign the *Replicate* operation to the remote *Hitachi Block Device* node.
The **Hitachi Block Replication Operation Properties Dialog** is displayed.

15. 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**.
16. Compile and activate the data flow, checking carefully that there are no errors or warnings.
17. Locate the active data flow in the **Monitor Inventory** and open its **Monitor Details** page.

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.
18. Monitor the active data flow to ensure the policy is operating as expected. For a healthy data flow you will periodically see:
 - Replication and snapshot jobs appearing for the source node in the **Jobs** area triggered according to the RPO.
 - 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.
19. Review the status of the local *Hitachi Block Device* to ensure snapshots are being created. Review the status of the remote *Hitachi Block Device* to ensure the replication is being performed and that snapshots are being created. New local and remote snapshots will appear 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: Restore workflows

The following topics describe the steps required to restore databases. These examples are performed from the **Restore Inventory**, however they can also be performed from the **Storage Inventory**. For a detailed introduction on how to work with the HDID user interface, please refer to *Hitachi Data Instance Director User's Guide*, MK-93HDID014.

In all of the following workflows, it is assumed that you have generated snapshots or clones of the database volumes using backup policies implemented in Data Instance Director. To prevent modifications of the existing configuration files and ensure they are available if required, the Data Instance Director will always restore configuration files to:

```
$HDID_HOME/saphana_backup/<SID>/<hdid_backup_id>
```

where:

\$HDID_HOME is the HDID installation dir,

<SID> is the SAP HANA SID,

<hdid_backup_id> is the snapshot ID shown in the Application column on the **Managed Storage** window.

How to mount a snapshot or clone for repurposing

A Block based clone or snapshot of a database can be used for repurposing by mounting it on a non-production server for development, test or analysis purposes. When using a Thin Image snapshots for repurposing, use a snapshot taken from a clone and not the production volume so that the performance of the production volume is not affected.

When mounting:

- It is only possible to mount a database on a node other than the source node.



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 mount operation fails.

- HDID Client software must be installed on the mount host.
- Only the backed up data files and a copy of the configuration files are included in the mounted snapshot or clone.
- HDID does not start or stop any application processes.

When unmounting:

- HDID removes the data area but does not check, stop, or start any application processes.

Procedure

1. Click **Restore** on the **Sidebar**.
2. Click the **Hitachi Block** button in the **Restore Dashboard**.
The **Hitachi Block Restore Inventory** is displayed, but no results are initially shown.
3. Ensure the search criteria are displayed by clicking **Show Search**.
4. Enter the required search criteria to find the desired snapshots and replications, then click **Search**.
All snapshots and replications meeting the search criteria will be displayed.
5. Select the snapshot or replication to be mounted.
6. Click **Mount**.



Note:

If you have only one SAP HANA Database server within your organization (i.e. you don't have a separate server for repurposing work), mount the snapshot to the ISM node. After the mount operation, manually copy the database to the SAP HANA Database server.

The **Mount Operation Wizard** opens.

7. Select the **Application** mount option, specify the **Host Group**, **Host** and **Mount Location**, then click **Finish**.
8. After mounting the snapshot or clone, it may be necessary to perform further manual recovery steps.

How to revert from a block snapshot or clone

A SAP HANA database can be reverted to an earlier state from a Block based snapshot or clone.

When reverting:

- It is possible to revert a database when SAP HANA is online or offline.
- The database will go offline during the revert process and remain offline after it has completed.
- Only the data area will be replaced but the log and shared areas will not be changed. The database administrator must use the SAP HANA Studio or DB Cockpit **Database Recovery Wizard** to complete the recovery process by:
 1. Bringing the database back online.
 2. Rolling forward to a specific point in time (the logs are not backed up, so they can be used for this purpose, if they have not been rotated out by other tools/mechanisms).



Note: Revert overwrites the original database and destroys all data in that database as a result.

Procedure

1. Click **Restore** on the **Sidebar**.
2. Click the **Hitachi Block** button in the **Restore Dashboard**.
The **Hitachi Block Restore Inventory** is displayed, but no results are initially shown.
3. Ensure the search criteria are displayed by clicking **Show Search**.
4. Enter the required search criteria to find the desired snapshots and replications, then click **Search**.
All snapshots and replications meeting the search criteria will be displayed.
5. Select the snapshot or replication to be used to revert from.
6. Click **Revert**.
The **Hitachi Block Snapshot Revert Wizard** is displayed.
7. To ensure the user does not accidentally perform a revert, the text `REVERT` must be typed in uppercase prior to clicking **Finish**.
The database is shutdown and the reversion process is performed.
8. Once reversion is complete, the database must be restarted manually by the database administrator who will need to choose how to recover the database (point-in-time, last known point, etc.).

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.

SAP HANA Application Node Wizard

This wizard is launched when a new SAP HANA Node is added to the Nodes Inventory.



Note: If you have a clustered SAP HANA environment and add or remove nodes to or from the cluster, the HDID SAP HANA application node must be updated so that the SAP HANA environment can be rediscovered. Any active data flows including that node must be reactivated to update the rules.

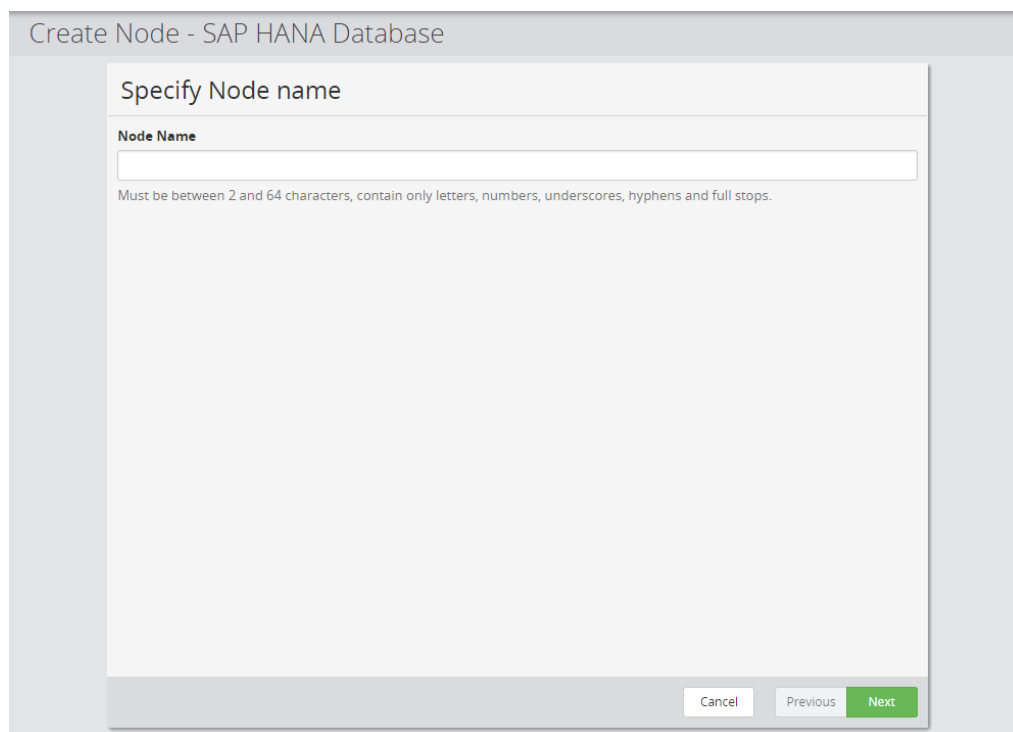
The image shows a screenshot of a web-based wizard titled "Create Node - SAP HANA Database". The main heading is "Specify Node name". Below this, there is a label "Node Name" followed by a text input field. A small note below the input field states: "Must be between 2 and 64 characters, contain only letters, numbers, underscores, hyphens and full stops." At the bottom right of the form, there are three buttons: "Cancel", "Previous", and "Next". The "Next" button is highlighted in green, indicating it is the active or recommended action.

Figure 4 SAP HANA Database Node Wizard - Specify Node Name

Control	Description
Node Name	Enter a name for the SAP HANA node.

Create Node - SAP HANA Database

Allocate node to Access Control Resource Group

This node will be added to the 'default' resource group. Select an additional resource group as required.

Name	Description
<input type="radio"/> myResGrp	

Cancel Previous Next

Figure 5 SAP HANA Database 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.

Create Node - SAP HANA Database

Discover SAP HANA Environment

Node running SAP HANA to be used for discovery

HANASERVER

SAP HANA Environment

Rediscover SAP HANA Environment

Cancel Previous Next

Figure 6 SAP HANA Database Node Wizard - Discover SAP HANA Environment

Control	Description
Node running SAP HANA to be used for discovery	Select the HDID client node that hosts the SAP HANA application. This node will then discover the SAP HANA environment.
Rediscover SAP HANA Environment	Click this button to refresh the cached details.

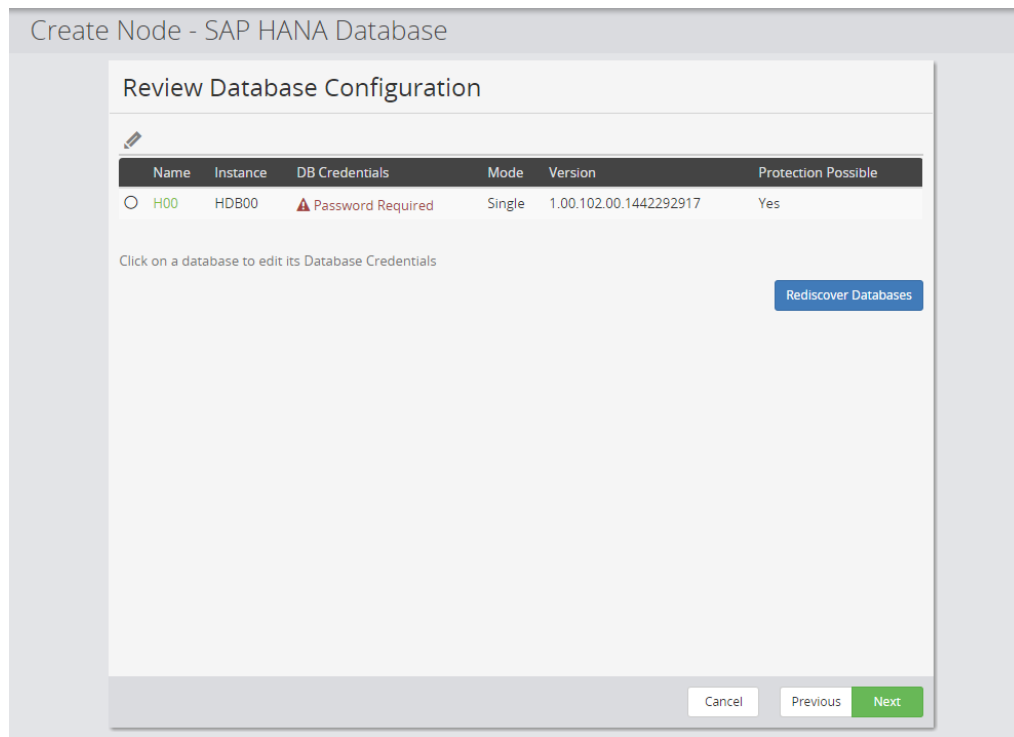


Figure 7 SAP HANA Database Node Wizard - Database Configuration

Control	Description
Name	Click on a database's Name to open the Specify Database credentials wizard (see below).
Rediscover Databases	Click this button to refresh the cached details.

Edit Node - SAP HANA Database 'mySapHanaServer'

Specify DB Credentials for database 'H00'

Username

Username which is used to run the operating system commands for this database

Password

Cancel Discard Previous Apply

Figure 8 SAP HANA Database Node Wizard - Specify Database Credentials

Control	Description
Username	Enter the username for the database.
Password	Enter the user's password for the database.

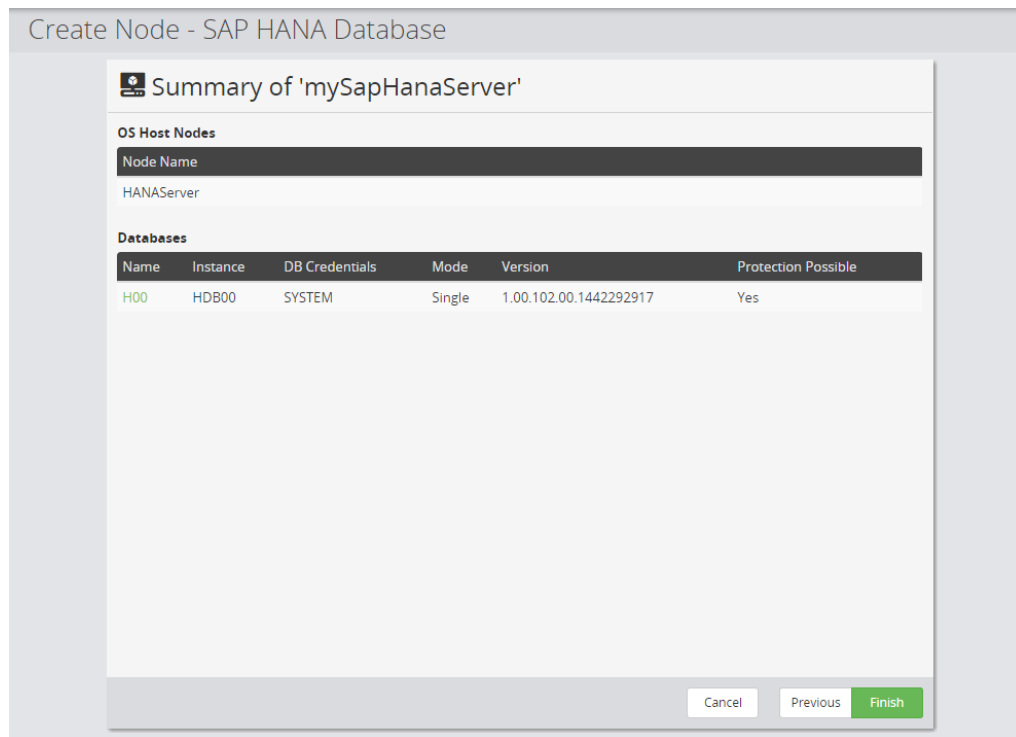


Figure 9 SAP HANA Database Node Wizard - Summary

Control	Description
Summary	Summary of the selected configuration.

Policies UI Reference

SAP HANA Database Classification Wizard

This wizard is launched when a new SAP HANA Database classification is added to a Policy.

The SAP HANA Database classification is used to define which databases are to be protected.



Note: When used in combination with a storage hardware backup operation, HDID will discover the underlying hardware paths at runtime. For storage hardware based backups, all the paths must exist on the same block hardware device.

Create Policy

SAP HANA Database Classification Attributes

Name	Instance	DB Credentials	Mode	Version
H00	HDB00	SYSTEM (edit)	Single	1.00.102.00.1442292917

[Select Database](#)

Cancel Discard Previous **Apply**

Figure 10 SAP HANA Database Wizard - Specify SAP HANA Database Classification Attributes

Control	Description
Databases	<p>Lists the currently selected databases.</p> <ul style="list-style-type: none"> 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.

Select Database

Select SAP HANA Node

mySapHanaServer

Databases

Filter database by name

Name	Instance	DB Credentials	Mode	Version
<input type="radio"/> H00	HDB00	SYSTEM	Single	1.00.102.00.1442292917

Refresh

Cancel OK

Figure 11 SAP HANA Database Wizard - Select SAP HANA Databases Dialog

Control	Description
Select SAP HANA Node	Select a node representing the SAP HAN 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.

Chapter 5: Troubleshooting

Troubleshooting SAP HANA

This section provides guidelines for how to troubleshoot issues that might occur when using SAP HANA.

There are currently no topics in this section.

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 Data Instance Director 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. See CDP below.

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 HDID 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 Data Instance Director Repository, Block or NAS device.

License key

A unique, alphanumeric code that is associated with the unique machine ID that is generated during the Data Instance Director 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 Data Instance Director 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 Data Instance Director Repository supports live backup, 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 Data Instance Director. 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.

Stubbing

The process by which the contents of a file, email body, or email attachment are replaced by a reference or link that points to the copy of the data that is stored on the HCP.

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

2845 Lafayette Street

Santa Clara, CA 95050-2639 USA

HitachiVantara.com | community.HitachiVantara.com

Contact Information

USA: 1-800-446-0744

Global: 1-858-547-4526

HitachiVantara.com/contact