



# Hitachi Data Instance Director v5.3 Backup for Scale-up Databases on SAP HANA

## Implementation Guide

By Mark Walker

October 2016

## Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to [SolutionLab@hds.com](mailto:SolutionLab@hds.com). To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

---

# Contents

<b>Solution Components .....</b>	<b>2</b>
Hardware Components .....	2
Software Components .....	4
<b>Solution Implementation .....</b>	<b>5</b>
Prerequisites to Install and Configure Hitachi Data Instance Director .....	6
Hitachi Data Instance Director Installation Procedures .....	8
Hitachi Data Instance Director Configuration .....	11
Database Restore Using Hitachi Data Instance Director Backups .....	16
System Copy Using Hitachi Data Instance Director Snapshots .....	19
SAP HANA Configuration File Handling in Hitachi Data Instance Director .....	21

# Hitachi Data Instance Director v5.3 Backup for Scale-up Databases on SAP HANA

## Implementation Guide

The implementation guide describes several scenarios for backing-up and restoring scale-up databases used for SAP HANA using Hitachi Data Instance Director (HDID) version 5.3. This includes information for the following:

- Setting-up the system
- Defining backup processes and using them to protect the SAP HANA database
- Copying the database to a second host

Protect your SAP HANA database using Hitachi Data Instance Director v5.3 for personalized recover point objective (RPO) and retention periods of SAP HANA backups.

Use Hitachi Data Instance Director to automate copy data management. Use this for backup and recovery of the data area and the configuration files for scale-up SAP HANA databases. It restores SAP HANA data to its original location. A restoration makes the configuration files available for manual integration into the running SAP HANA instance.

After you implement a resilient backup and restore system for your scale-up SAP HANA instance using Data Instance Director, make rapid snapshot-based backups of the HANA database at a frequency you define. These snapshots are stored and retained alongside the HANA configuration active at the time they were made for the period you need. Data Instance Director provides business-defined data protection.

If necessary, restore these snapshots with their configuration, to either of the following locations:

- The original host, in case of system failure
- To a new host, to facilitate a system copy of the HANA database

You can use the file backup functionality of Data Instance Director to make a backup of the log files used by SAP HANA. This enables a more fine-grained restore window.

Snapshot functionality using Data Instance Director adds very little overhead to the SAP HANA database during snapshot creation. It relies on the same savepoint mechanism used by HANA, which runs automatically every five minutes on the database.

Use Hitachi Data Instance Director v5.3 to protect your SAP HANA instance to implement a rapid, space-conscious backup policy. Use Data Instance Director with the native tools in SAP HANA to restore a snapshot of your database in case of catastrophic failure to get your database back up and running quickly.

---

**Note** — These procedures were developed in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow recommended practice by conducting proof-of-concept testing for acceptable results before implementing this solution in your production environment. Test the implementation in a non-production, isolated test environment that otherwise matches your production environment.

---

## Solution Components

You can use these hardware and software components with the procedures in this implementation guide for Hitachi Data Instance Director v5.3.

You can implement these Hitachi Data Instance Director v5.3 procedures on any SAP HANA appliance from Hitachi Data Systems using SAP HANA version SP10 or later. You can also use the procedures in this guide to back-up and restore a SAP HANA tailored datacenter integration (TDI) environment.

## Hardware Components

Your implementation of Hitachi Data Instance Director v5.3 will only use some of these hardware components. To use these procedures, you need:

- A SAP HANA appliance using a Hitachi Data Systems server platform and storage system
- A SAP HANA TDI system, with a storage system from Hitachi Data Systems

## Storage Systems

Hitachi Data Instance Director v5.3 works with any of these storage systems in a SAP HANA appliance.

### Hitachi Unified Storage VM

[Hitachi Unified Storage VM](#) is an entry-level enterprise storage platform. It combines storage virtualization services with unified block, file, and object data management. This versatile, scalable platform offers a storage virtualization system to provide central storage services to existing storage assets.

Unified management delivers end-to-end central storage management of all virtualized internal and external storage on Unified Storage VM. A unique, hardware-accelerated, object-based file system supports intelligent file tiering and migration, as well as virtual NAS functionality, without compromising performance or scalability.

The benefits of Unified Storage VM are the following:

- Enables the move to a new storage platform with less effort and cost when compared to the industry average
- Increases performance and lowers operating cost with automated data placement
- Supports scalable management for growing and complex storage environment while using fewer resources
- Achieves better power efficiency and with more storage capacity for more sustainable data centers
- Lowers operational risk and data loss exposure with data resilience solutions
- Consolidates management with end-to-end virtualization to prevent virtual server sprawl

### Hitachi Virtual Storage Platform Gx00 Models

[Hitachi Virtual Storage Platform Gx00 models](#) are based on industry-leading enterprise storage technology. With flash-optimized performance, these systems provide advanced capabilities previously available only in high-end storage arrays. With the Virtual Storage Platform Gx00 models, you can build a high performance, software-defined infrastructure to transform data into valuable information.

Hitachi Storage Virtualization Operating System provides storage virtualization, high availability, superior performance, and advanced data protection for all Virtual Storage Platform Gx00 models. This proven, mature software provides common features to consolidate assets, reclaim space, extend life, and reduce migration effort.

## Hitachi Virtual Storage Platform Fx00 Models

[Hitachi Virtual Storage Platform Fx00 models](#) deliver superior all-flash performance for business-critical applications, with continuous data availability. High-performance network attached storage with nondisruptive deduplication reduces the required storage capacity by up to 90% with the power to handle large, mixed-workload environments.

Hitachi Storage Virtualization Operating System provides storage virtualization, high availability, superior performance, and advanced data protection for all Virtual Storage Platform Fx00 models. This proven, mature software provides common features to consolidate assets, reclaim space, extend life, and reduce migration effort.

## Servers

The Intel CPU processors used in the server in this Hitachi Data Instance Director v5.3 solution can be any processor supported by the server blade and the SAP HANA database. These include versions 2, 3, and 4 of the certified Intel Xeon E5 processor and Xeon E7 processor. Check the support matrix for the server blade and the SAP HANA database supported CPU list for more information.

Hitachi Data Instance Director v5.3 works with any of these servers in a SAP HANA appliance.

### Hitachi Compute Blade 500

[Hitachi Compute Blade 500](#) combines the high-end features with the high compute density and adaptable architecture you need to lower costs and protect investment. Safely mix a wide variety of application workloads on a highly reliable, scalable, and flexible platform. Add server management and system monitoring at no cost with Hitachi Compute Systems Manager, which can seamlessly integrate with Hitachi Command Suite in IT environments using Hitachi storage.

### Hitachi Compute Blade 2000

[Hitachi Compute Blade 2000](#) is an enterprise-class blade server platform. It features the following:

- A balanced system architecture that eliminates bottlenecks in performance and throughput
- Configuration flexibility
- Sustainable power-saving capabilities
- Fast server failure recovery using a N+1 cold standby design that allows replacing failed servers within minutes

### Hitachi Compute Blade 2500

[Hitachi Compute Blade 2500](#) delivers enterprise computing power and performance with unprecedented scalability and configuration flexibility. Lower your costs and protect your investment.

Flexible I/O architecture and logical partitioning allow configurations to match application needs exactly with Hitachi Compute Blade 2500. Multiple applications easily and securely co-exist in the same chassis.

Add server management and system monitoring at no cost with Hitachi Compute Systems Manager. Seamlessly integrate with Hitachi Command Suite in Hitachi storage environments.

## Software Components

Your implementation of these procedures require using Hitachi Data Instance Director v5.3 with the SAP HANA appliance.

### Hitachi Data Instance Director

[Hitachi Data Instance Director](#) provides a modern, holistic approach to data protection, recovery, and retention.

Data Instance Director has a unique workflow-based policy engine, presented in a whiteboard-style user interface. It helps you to map copy data management processes to business priorities.

Data Instance Director includes a wide range of fully integrated storage-based and host-based incremental-forever data capture capabilities. These can be combined into complex workflows to automate and simplify copy data management.

## Solution Implementation

To implement an effective backup policy for a SAP HANA database using Hitachi Data Instance Director v5.3, consider your database usage patterns.

For example, if the SAP HANA instance supports SAP Business Warehouse, and data is created in the database overnight, then it is important to make a backup before data loading starts and when it finishes. Backing up during the day, when data changes are fewer and not as important, could do the following:

- Decrease database performance
- Use more storage space than necessary

Other usage patterns of the SAP HANA database require different backup strategies.

Also, consider the retention period of backups. Determine how long to retain the backups created by Data Instance Director. IT staff cannot solely provide this information. All users of the SAP database and supported applications must discuss this.

Once this information is available, you can create a backup strategy using Hitachi Data Instance Director.

The following are the three main use cases to use Hitachi Data Instance Director with SAP HANA:

- **Backup and restore using SAP HANA snapshots (snapshot backup)**

This procedure asks for the following to happen:

- (1) Hitachi Data Instance Director has a SAP HANA snapshot created.
- (2) Hitachi Data Instance Director stores the snapshot on the [storage system](#).

You can restore to the point-in-time when a particular snapshot was created. In case of total system failure with loss of SAP HANA data, only data created since the most recent snapshot would be lost.

See “Revert a Hitachi Data Instance Director Snapshot File” on page 17.

- **Backup and restore using SAP HANA snapshots plus SAP HANA log backup replay (point-in-time restore)**

This is very similar to the snapshot backup. This allows backup and restoration of the SAP HANA log area backups created automatically by the SAP HANA database every five minutes. You can perform this restoration to a more fine-grained point-in-time by doing the following:

- (3) Restore latest the snapshot.
- (4) Replay the data in the log backups since the latest snapshot.

At most, five minutes of data would be lost in case of total system failure with loss of the SAP HANA storage.

See “Revert SAP HANA Log Backup Files” on page 17.



### ■ Copy a system using Hitachi Data Instance Director snapshots

Use this procedure to make a copy of a SAP HANA instance and to transfer the instance to a second host. An example of using this is to create a pre-production database by copying the current production database.

See “System Copy Using Hitachi Data Instance Director Snapshots” on page 19.

Many of the steps necessary to perform these procedures are shared with other procedures.

Information concerning the handling of the SAP HANA configuration files is in SAP HANA Configuration File Handling in Hitachi Data Instance Director.

## Prerequisites to Install and Configure Hitachi Data Instance Director

The following are the prerequisites when using Hitachi Data Instance Director v5.3 with SAP HANA:

- 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
- The operating system ‘DB System Admin’ or ‘<SID>adm’ user credentials are known.

Remember the following constraints:

- These configurations are *supported*:
  - SAP HANA SPS 10 or later
  - Scale-up single node environment
  - MCOS (multiple components on one system)
  - MCOD (multiple components on one database)
- These configurations are *not supported*:
  - Scale-out cluster environments
  - MDC (multi-tenant database containers)
  - System replication (application mirror)

## Hitachi Data Instance Director Master Server Node

Hitachi Data Instance Director v5.3 requires a server to act as master node to do the following:

- Master orchestrator of backups
- Handler of job creation
- Manager and communicator with the other machines (clients and repository) in the backup topology

There must be only one active master node in the Data Instance Director topology.

The master node does not transfer a large amount of data to the other nodes in the system. This means that the speed of the network connection is not very important.

The Data Instance Director master node must operate on one of these versions of Microsoft® Windows Server®:

- Microsoft Windows Server 2008 (64-bit)
- Microsoft Windows Server 2008 R2 (64-bit)
- Microsoft Windows Server 2012 (64-bit)
- Microsoft Windows Server 2012 R2 (64-bit)

For full requirements for the master node, see the [Hitachi Data Instance Director web page](#). Detailed specifications are [in this datasheet](#) (PDF file).

## Hitachi Data Instance Director Repository Server Node

Hitachi Data Instance Director v5.3 requires a server to act as a repository node to handle the storage of backed-up data.

The repository node has strict network requirements. The recommendation is to have at least a 10 Gb Ethernet connection to the repository node if you will make file-based backups, such as a log back up.

The Disaster Instance Director repository node must operate on one of these versions of Microsoft Windows Server:

- Microsoft Windows Server 2008 (64-bit)
- Microsoft Windows Server 2008 R2 (64-bit)
- Microsoft Windows Server 2012 (64-bit)
- Microsoft Windows Server 2012 R2 (64-bit)

For full requirements for the repository node, see the [Hitachi Data Instance Director web page](#). Detailed specifications are [in this datasheet](#) (PDF file).

## Hitachi Data Instance Director Installation Procedures

The following are the installation procedures for the initial configuration and setup of Hitachi Data Instance Director v5.3 for use with SAP HANA:

1. “Install the Hitachi Data Instance Director Master Node” on page 8
2. “Install the Hitachi Data Instance Director Repository Node” on page 9
3. “Install the Hitachi Data Instance Director Software on the Linux Server Hosting the SAP HANA Database” on page 10
4. “Create the Block Storage Node” on page 10

### Install the Hitachi Data Instance Director Master Node

To install the master node for Hitachi Data Instance Director v5.3, do the following.

1. Identify the server node running Microsoft® Windows Server® to host the Hitachi Data Instance Director master server node.

This server node must have network connectivity to all other nodes in the Data Instance Directory topology:

- The repository node
- Any node to be backed-up (client nodes)

Verify that this server node meets the installation requirements in “Hitachi Data Instance Director Master Server Node” on page 7. This includes the node running a supported version of Microsoft Windows Server.

2. Install the Data Instance Director software on the master server node.
  - (1) Transfer the Data Instance Director installation file to the identified master server node,
  - (2) Start the installation by double-clicking the executable file (.exe file).
  - (3) Follow the prompts to install the software. The default installation directory is c:\Program Files\Hitachi\HDID.
  - (4) Verify that the **Master** option is clicked on the **Installation type** screen, to install the master node in the topology.

This installs the components required to make the current system the master server node.

There can be multiple master nodes on the same network. However, a Data Instance Director repository node can only be authorized and controlled by one master node.

You can install standby master nodes on a Windows Server failover cluster. See “Installation on a Windows cluster environment” in [Hitachi Data Instance Director User's Guide](#) (PDF).

This installs the main Data Instance Director user interface. This interface is only accessible on the master server node, where you set up and manage through the master node.

A master server node can also act as a client node and a repository node. This enables the master node to capture data on it. Also, the master server node is able to store data from source nodes in a repository on the master server node. Master nodes require installing Oracle Java.

For complete installation, refer to *Hitachi Data Instance Director User's Guide*.

## Install the Hitachi Data Instance Director Repository Node

To install the repository node for Hitachi Data Instance Director v5.3, do the following.

1. Identify the server node running Microsoft® Windows Server® to host the Hitachi Data Instance Director repository node.

This server node must have network connectivity to all other nodes in the Data Instance Directory topology:

- The master node
- Any node to be backed-up (client nodes)

Verify that this server node meets the installation requirements in “Hitachi Data Instance Director Repository Server Node,” starting on page 7. This includes the node running a supported version of Microsoft Windows Server.

The functionality of the repository server node can be performed by the master server node, if you want. The master server node fulfills the prerequisites of the master server node and the repository server node, including sufficient disk space to host the backed-up files.

---

**Note** — The repository server node does not store the SAP HANA database backups physically. The repository only stores the metadata and the SAP HANA configuration files.

---

2. Install the Data Instance Director software on the repository server node.
  - (1) Transfer the Data Instance Director installation file to the identified master repository node,
  - (2) Start the installation by double-clicking the executable file (.exe file).
  - (3) Follow the prompts to install the software. The default installation directory is c:\Program Files\Hitachi\HDID.
  - (4) Verify that the **Custom** option is clicked on the **Installation type** screen, to install the repository node in the topology.
  - (5) Verify that the **Repository** option is clicked on the **Custom install** screen.
  - (6) Select the directory to contain the repository data. This directory needs at least 200 GB of free space available.
  - (7) Give the node a name. The default name is the Microsoft® Windows Server® host name, but you can rename it.
  - (8) When asked, type one of the following:
    - The host name, if there is a DNS
    - IP address of the master server node
  - (9) Once installation completes, authorize the repository node from the Hitachi Data Instance Director graphical user interface on the master server node.
    - i. Select (highlight) the repository server node in the list of **Unauthorized nodes**.
    - ii. After selecting the repository server node, from the Node Manager, click the **Move** button (>>).

This installs the components required to make the current system the repository server node.

Any machine used as a hardware proxy node for storage-based protection must have the repository component installed because it provides the ISM capability. Repository/ISM nodes require installing Oracle Java.

For complete installation information, refer to *Hitachi Data Instance Director User's Guide*.

## Install the Hitachi Data Instance Director Software on the Linux Server Hosting the SAP HANA Database

To install the software for Hitachi Data Instance Director v5.3 on the Linux server hosting the SAP HANA database, do the following.

1. Log on to the Linux server hosting the HANA database as the root user.
2. Copy the Data Instance Director installation file to a temporary directory, such as /tmp.
3. Run the installation file from a terminal or shell by typing the following and pressing Enter:

```
. /HDI D-R<tab>
```

4. Accept the license and then click the **Client installation** option.
5. Give the node a name. The default name is the Linux host name of the server.

---

**Note** — The recommendation is to keep the node name the same as the Linux host name.

---

6. When asked, type one of the following:
  - The host name, if there is a DNS
  - IP address of the master server node
7. Once installation completes, authorize the repository node from the Hitachi Data Instance Director graphical user interface on the master server node.
  - (1) Select (highlight) the repository server node in the list of **Unauthorized nodes**.
  - (2) After selecting the repository server node, from the Node Manager, click the **Move** button (>>).

### Create the Block Storage Node

You create a block storage node on the storage array. This includes the following:

- Hitachi Unified Storage VM
- Hitachi Virtual Storage Platform Gx00
- Hitachi Virtual Storage Platform Fx00

To create a block storage node on the storage array, do the following.

1. Create a Hitachi Thin Image pool on the storage array.

On the storage array hosting the block storage, create a Hitachi Thin Image pool large enough to hold the SAP HANA backups. The pool size depends on the following:

- The size of the SAP HANA database
- The backup frequency (RPO) and retention period required.

Determine this information with the aid of all users of the SAP HANA database you need to back up.

There must be connectivity from the block storage to both of the following:

- The SAP HANA Linux node
- The Repository node

If you wish to use the procedure in “System Copy Using Hitachi Data Instance Director Snapshots” on page 19, then this storage requires Fibre Channel connectivity to the secondary host where you want to copy the SAP HANA database snapshot.

To act as a storage proxy, create a command device on the block storage to be made available to the master server node or repository server node.

2. Add the block storage node in the Data Instance Director graphic user interface.
  - (1) In the main Hitachi Data Instance Director graphical user interface, under **Agentless/Hardware Node**, click **Hitachi Block Storage** and click **Create**.
  - (2) Select (highlight) the storage proxy node on the Data Instance Director node that has access to the command device on the block storage.
  - (3) On the next screen, do the following:
    - i. Name the block server node.
    - ii. Enter the serial number of the block storage to use for the snapshot backups
    - iii. Enter the user name and password for the storage system.
    - iv. If necessary to make it easier to specify the pool to use for backups, enter the range of LDEV names, and/or Fibre Channel ports to use.
3. Once installation completes, authorize the block storage node from the Hitachi Data Instance Director graphical user interface on the master server node.
  - (1) Select (highlight) the block storage node in the list of **Unauthorized nodes**.
  - (2) After selecting the block storage node, from the Node Manager, click the **Move** button (>>).

## Hitachi Data Instance Director Configuration

To configure Hitachi Data Instance Director for use with SAP HANA in order to perform snapshot backup and log file backup, do the following procedures:

- “Add a SAP HANA Application Node in Hitachi Data Instance Director” on page 11
- “Create a Snapshot Backup Policy for SAP HANA” on page 12
- “Create a Hitachi Data Instance Director Data Flow to Call the Snapshot Backup Policy” on page 14
- “Create a Backup Policy to Save the SAP HANA Log Backup files” on page 15
- “Create a Hitachi Data Instance Director Data Flow to Call the SAP HANA Log Backup Policy” on page 16

## Add a SAP HANA Application Node in Hitachi Data Instance Director

Once adding the Linux node to Hitachi Data Instance Director v5.3 (“Install the Hitachi Data Instance Director Software on the Linux Server Hosting the SAP HANA Database” on page 10), add an application node for the SAP HANA database which resides on the Linux host to Data Instance Director.

To add the SAP HANA application node to Hitachi Data Instance Director, do the following.

1. In the main Hitachi Data Instance Director graphical user interface, under **Agentless/Hardware Node**, click **SAP HANA** and click **Create**.
2. Name the SAP HANA node, such as using the SAP HANA instance name. Select the Linux node that hosts the SAP HANA database.

---

**Note** — It is important to select the correct node. If the selected host name does not contain a SAP HANA database, then there will be an error when requesting a snapshot.

---

3. Under **SAP HANA Databases**, to open the **Enter DB information** dialog box, click **Add** and enter the following:
  - SID
  - System user password
  - <sid>adm Linux username and password
4. Once installation completes, authorize the SAP HANA node from the Hitachi Data Instance Director graphical user interface on the master server node.
  - (1) Select (highlight) the SAP HANA node in the list of **Unauthorized nodes**.
  - (2) After selecting the SAP HANA node, from the Node Manager, click the **Move** button (>>).

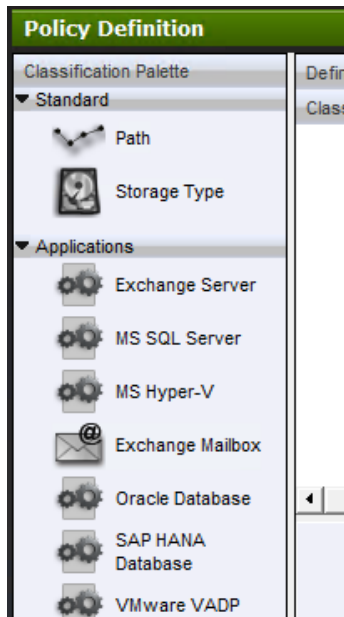
## Create a Snapshot Backup Policy for SAP HANA

In Hitachi Data Instance Director, a policy represents a business objective describing what needs to happen.

To create a snapshot backup policy for SAP HANA, do the following.

1. In the Data Instance Director graphic user interface, click the **Policy Definition** option. The main **Policy Definition** screen opens with a **Classification** (what) section and an **Operation** (how) section.

The left side of the screen is the **Classification Palette** area. It contains a list of different operations that Data Instance Directory performs. See Figure 1 on page 13.



**Figure 1**

2. Add the SAP HANA database to the policy.
  - (1) Under **Classification Palette**, expand the **Applications** area.
  - (2) Under **Applications**, drag **SAP HANA Database** symbol to **Classification** on the **Policy Definition** screen. A list of defined databases displays in the **Available Databases** list.
  - (3) Select the database to backup, and then add it to the **Included Databases** list by clicking **Include**.

---

**Note** — Only databases listed in the **Included Databases** list are backed-up.

---

3. Add the snapshot operation.
  - (1) Under **Operation Palette** on the right side of the screen, drag the **Snapshot** operation to the **Operation** section. The **Snapshot Attributes** show at the bottom of the screen.
  - (2) Verify these **Snapshot Attributes** settings:
    - **Mode** is set to **Hardware**
    - **Hardware Type** is set to **Hitachi Block**
  - (3) Set the **Retention Period**.

The **retention period** determines the amount of time that backups are kept on the block storage. This determines, in part, how much total disk space is required on the block storage for all backups that are made over the retention period.

Sufficient space to hold the SAP HANA configuration files is necessary also in the repository for the number of backups that will be made over the retention period.

The number of backups per day times number of days retention equals the total number of backups.



(4) Set the **Recovery Point Objective**.

The **recovery point objective** determines how often Hitachi Data Instance Director performs a snapshot.

(5) Select the **Quiesce configured applications before backup** check box.

The **Quiesce configured applications before backup** check box ensures making the snapshot in a coherent way. Select this check box to have SAP HANA perform database snapshots immediately before Data Instance Director makes the copy. This check box must be selected.

4. Save the policy for use in a data flow.

Save the policy with a meaningful name. This created SAP HANA snapshot policy can be used in a data flow.

### Create a Hitachi Data Instance Director Data Flow to Call the Snapshot Backup Policy

A *data flow* object in Hitachi Data Instance Director performs a backup operation using a policy object created in “Create a Snapshot Backup Policy for SAP HANA” on page 12 to backup the data.

To create a data flow object, do the following.

1. In the **Navigation** area of the main Hitachi Data Instance Director screen, click the **Data Flow** option.
2. Click the **Plus (+)** tab to create a new data flow.
3. From the **Source** area, drag the database to the **Data Flow** area. In the options area below the **Data Flow** area, a list of policies concerning the selected database displays.
4. Verify that the check boxes all back-up policies are selected. Specifically, verify that the newly created data flow has been selected.

As soon as the policy is selected, on the right of the screen, a list of check boxes for the snapshot display. The most important check box is the **Target Storage Pool**.

5. From the **Target Storage Pool** list, click the Hitachi Thin Image pool created in “Create the Block Storage Node” on page 10. This opens a dialog box to selection the storage node, and then the storage pool.
6. Save and activate the data flow. Then, distribute the data flow so that all nodes in the topology can see and use it.

### Create a Backup Policy to Save the SAP HANA Log Backup files

About every five minutes, SAP HANA makes a backup of the entries in its redo-log area. You can replay these redo log backups to recreate data in the SAP HANA database.

After SAP HANA backs up the redo log entries, Hitachi Data Instance Director can copy the logs to more secure storage. The live log area of SAP HANA cannot be backed-up while the database is running. However, SAP HANA makes a backup of the log area every five minutes or so. Data Instance Director can back up this log-backup.

---

**Note** — Backing up the redo log files of SAP HANA uses the file backup functionality in Data Instance Director. There is no difference between backing up the redo logs in SAP HANA and backing up any other file on the Linux file system.

---

Backups of the SAP HANA log files are made to the HDID repository, not to the Hitachi storage array.

Creating a backup of the active log area of the SAP HANA database is not feasible. Some files are always in use by SAP HANA. This means that Data Instance Director would give errors such as “files are in flux.”

To create a policy to create backups of the log backup files, do the following.

1. From the main Hitachi Data Instance Director user interface, under **Navigation**, click **Policy Definition** and then click **New Page**.
2. On the **Policy Definition** screen, under **Classification Palette**, click to expand **Standard**, then drag the **Path** symbol to the **Classification** area of the **Policy Definition** screen.
3. Select the symbol in the **Classification** area that you dropped, and, in the options at the bottom of the screen, provide the path where SAP HANA stores its backup logs in the **Enter paths** text box.

The storage of backup logs for SAP HANA is, by default, in this path:

```
/usr/sap/<HANA instance name>/HDB<HANA instance number>/backup/log
```

However, this path often is a link to some other directory. Data Instance Director cannot backup files from a directory containing symbolic links. You must find the real path where SAP HANA stores its backups. To find the real path, do the following in a terminal session:

- (1) Change to the log backup path.
- (2) Type this command: `readlink -f .`

The system replies with the real path of the current directory. Enter this path in the policy definition screen.

4. From the **Operation Palette**, drag a **Backup** operation to the policy.
5. As with the snapshot operation, set the **Recovery Point Object** frequency (backup frequency) and the **Retention Period**. Then, save the policy.

Since SAP HANA creates log backups every five minutes, a low value for the RPO is usually appropriate.

This defines a backup policy to copy the log backup files created by SAP HANA.

## Create a Hitachi Data Instance Director Data Flow to Call the SAP HANA Log Backup Policy

This creates the data flow to back up the log backups for SAP HANA.

To create a data flow to backup the SAP HANA log backups, do the following:

1. Create a new data flow.
  - (1) From the main Hitachi Data Instance Director user interface, click the **Data Flow** pane.
  - (2) To create a new data flow, click the plus tab (" + ").
2. Add in Linux node hosting the database.
  - (1) From the **Source** pane, drag the Linux node that hosts the SAP HANA database to the **Data Flow** pane. Do not drag the SAP HANA database node itself.
  - (2) Click the Linux node name in the data flow, and then select the check box next to the policies to back up. See [Create a Backup Policy to Save the SAP HANA Log Backup files](#).
3. Define the data flow.
  - (1) Under **Mover**, drag and drop **Batch** on top of the Linux node in the data flow.
  - (2) Under **Destination**, drag and drop the repository node on the **Batch** mover in the data flow.
4. Save and activate the data flow. Then, distribute the data flow so that all nodes in the topology can see and use it.

## Database Restore Using Hitachi Data Instance Director Backups

Two actions are necessary to restore the database from an HDID snapshot.

- “Revert a Hitachi Data Instance Director Snapshot File” on page 16
- “Revert SAP HANA Log Backup Files” on page 16

After restoring the data to SAP HANA, revert the database by following the steps in “Database Restore Using Hitachi Data Instance Director Backups” on page 16.

### Revert a Hitachi Data Instance Director Snapshot File

If it is necessary to restore the snapshot to the SAP HANA server, then Hitachi Data Instance Director can perform a revert operation on the snapshot. An example of when you would do this is if the database becomes corrupted

---

**Note** — Doing a revert process stops the SAP HANA database, if running, and destroys all available existing data available in the SAP HANA database. This keeps the information in the HANA log area, however.

---

To restore a Hitachi Data Instance Director snapshot file, do the following.

1. From the main Hitachi Data Instance Director graphical user interface, click the **Storage** option, and then click the **Managed Storage** tab.
2. Expand the snapshot list to view the snapshot to be restored.
3. Right-click the snapshot to restore, and then click **Revert** from the menu. A confirmation dialog box opens.
4. Confirm the restoration on the dialog box. This is necessary to prevent accidental overwrite of the database.

Check the progress of the revert operation in the **Log Manager**.

---

**Note** — The revert operation shuts down SAP HANA. Make sure that no users are connected to the SAP HANA system prior to doing a revert operation.

---

5. When finished reverting the snapshot, do one of the following:
  - Restore the database using the instructions in “Database Restore Using Hitachi Data Instance Director Backups” on page 16.
  - If you need a finer-grained restoration than available with the snapshot file, continue with “Revert SAP HANA Log Backup Files” on page 16.

Reverting the SAP HANA snapshot file to the HANA data area of the Linux server does not restore the database.

---

**Note** — Do not start the SAP HANA database at this point. **This is very important.** If you start the database without correctly restoring the snapshot, you cannot confirm the state of the database correctly. In addition, upon first startup, SAP HANA deletes any existing snapshot files from the file system.

---

## Revert SAP HANA Log Backup Files

Restoring the SAP HANA log backup files means being able to recover the database to a finer-grained point in time than the snapshot date.

---

**Note** — Only perform this procedure when the SAP HANA database is offline. Do this immediately after restoring the snapshot to the Linux file system using Hitachi Data Instance Director.

---

To restore a Hitachi Data Instance Director log backup file, do the following.

1. From the Hitachi Data Instance Director graphical user interface, click **Restore**, and then click **Search**. If the date of the log backup is before the date in the **Criteria** box, then change the date.

This shows a list of backup policies in the Data Instance Director repository.

2. Expand the policy tree for the log backup policy. A list of individual backup dates and times displays.
3. Right-click the backup date, and then click **Restore**.

---

**Note** — If possible, only restore only specific backup logs and not all the contents of the log backup directory.

---

For example, if you need to restore SAP HANA to 2:45 pm on Monday and the snapshot was taken (and restored) from 2:00 pm on Monday, restore log backups to the SAP HANA file system from 1:55 pm to 2:50 pm. This has the advantage of requiring much less disk space on the SAP HANA server. It also takes less time to restore.

Verify that the correct log backup files are present. If you do not verify this, the SAP HANA restoration can fail or can succeed with incorrect results. If in doubt, it is better to restore the whole directory and then delete the log backups after completing the restoration.

The **Restore Options** dialog box opens.

4. Click **Perform Restore** to start the restore process.

Leave the default options, unless there is not enough disk space. If there is not enough disk space, you can specify an alternate restore location.

A dialog box opens showing restoration progress, showing the names of the files being restored.

After finishing the restoration, the progress window closes.

The selected log files are now available in the SAP HANA file system, ready for use in “Database Restore Using Hitachi Data Instance Director Backups” on page 16.

## Database Restore Using SAP HANA Studio

To restore a database, you first need to do the following:

- “Revert a Hitachi Data Instance Director Snapshot File” on page 16 [all cases]
- “Revert SAP HANA Log Backup Files” on page 17 [when a finer-grained restoration is needed than the snapshot]

To complete the restoration operation of the SAP HANA database, perform a database restore using SAP HANA Studio.

To restore the SAP HANA database using SAP HANA studio, do the following.

1. Open SAP HANA Studio.
  - If the SAP HANA instance is already configured for use in SAP HANA Studio, the instance displays in the list of systems.
  - If the SAP HANA instance is not configured for use in SAP HANA Studio, you need to configure this instance before restoring the database. This is necessary even if the SAP HANA instance is offline, as will be when Data Instance Director reverts a snapshot.
2. To start the recovery wizard, right-click the SAP HANA system name, click **Backup and Recovery**, and then click **Recover System**.
3. To restore the SAP HANA system to the snapshot recovered by Data Instance Director, do one of the following:
  - Click the **Recover the database to a specific data backup or storage snapshot** option.
  - To recover SAP HANA to the state of the snapshot plus all available log backups, click the **Recover the database to its most recent state** option.
  - To restore to the snapshot plus a replay of some of the available log files, click the **Recover the database to the following point in time**.

When restoring to either the most recent state or a different point in time, restore the most recent snapshot taken before the destination time and have the necessary log backups available. If you restore an earlier snapshot and to replay more log backups, this adds considerably more time to restore the database.

4. On the **Specify Backup Location** window, do not change any settings and click **Next**. This restores the backup logs to their default location. SAP HANA examines this location automatically during the restoration.

Continue with the **Select a Backup** window. After a few seconds, this window lists all backups performed on the selected SAP HANA instance.

This listing of any backup currently present on the file system has a green symbol. The listing of other backups has a red symbol.

The snapshot reverted using Data Instance Director is the only snapshot in the list with a green symbol.

- If you are restoring the most recent snapshot, it is the first snapshot in the list.
- If you are not restoring the most recent snapshot, it is present elsewhere in the list. If you are reverting the database to a snapshot from an earlier point in time, it will not be the first one in the list.

5. Select the snapshot to restore, and then click **Next**. On the **Review Recovery Settings** screen, the wizard shows the action about to be performed.
6. To perform the restore, click **Finish**.

SAP HANA Studio performs the restoration of the system and then confirms success.

When finished, the SAP HANA database is running, reverted to the state of the snapshot reverted with Data Instance Director and, if used, the log backups to reach a specific date and time.

## System Copy Using Hitachi Data Instance Director Snapshots

Making a copy of a SAP HANA database using Hitachi Data Instance Director snapshots is very similar to the backup and restore scenario. To make a system copy, you complete the “restoration” on a second host.

This procedure assumes the following:

- The second host already has a SAP HANA instance running.
- The source and destination machines have a similar configuration to support the SAP HANA database size.

Perform the system copy process in two steps:

1. “Create a Snapshot for System Copy” on page 19
2. “Restore a System for System Copy” on page 19

### Create a Snapshot for System Copy

When copying a system, use the same procedures in the following:

- “Create a Snapshot Backup Policy for SAP HANA” on page 12
- “Create a Hitachi Data Instance Director Data Flow to Call the Snapshot Backup Policy” on page 14

### Restore a System for System Copy

---

**Note** — The second host must have Fibre Channel connectivity to the block storage for this procedure to work.

---

To restore a snapshot to a second host to make a copy of the initial SAP HANA database, do the following.

1. Verify installation of the client software for Hitachi Data Instance Director on the second host and that the second node is authorized in Data Instance Director. You can do this verification and authorization even if SAP HANA is not installed on the host.
2. Determine the mount location for the database snapshot on the destination host.

Data Instance Director proposes to mount the snapshot in its original location, but on the (second) destination host. In other words, Data Instance Director offers to mount the snapshot in /hana/data. This is probably not appropriate, as it will do the following:

- Overwrite any data current in this directory
- Only survive while the snapshot is mounted to the host, and will not persist across reboots

Normally it is more appropriate to mount the snapshot in a temporary location, and then move the data manually to its final destination.

3. Locate or create the directory to host the snapshot. An recommended directory to host the snapshot is the following:  
/mnt
4. On the **Managed Storage** tab of the Storage portion of the main Data Instance Director graphic user interface, expand the tree until you see the snapshot to restore to the destination in the list.
5. Mount the snapshot on the second host.
  - (1) Right-click the snapshot, and then click **Mount** from the menu.
  - (2) In the **Mount Snapshot Set** dialog box, set the following options:
  - (3) **Mount Host:** click the destination node from the list.
  - (4) **Mount Options:** Click **Mount** in directory, and select the directory to host the snapshot
  - (5) If the destination host has SAP HANA installed on it, stop SAP HANA. **This is important.**
  - (6) If there is a volume group called **vgdata**, temporarily rename this volume group to mount the snapshot, as the mount also creates a volume group called vgdata.

When the mount operation finishes, the following line appears in the log: Disks now visible to mount host

The selected mount point now contains the data contained in the snapshot. If using the recommended directory, the snapshot data is the following: /mnt/hana/data

If using the recommended mount directory, /hana/data still contains the original SAP HANA data area of the destination host. This data has not changed yet.

6. Prepare data for restoration as a system copy.
  - (1) If the /hana/data/ still contains the original SAP DATA data, delete the contents of the data area.
  - (2) If the snapshot data is in a temporary directory, such as /mnt/hana/data, copy the complete contents of the snapshot data area (/mnt/hana/data/\*) to the SAP HANA data area: /hana/data/
  - (3) After copying the snapshot data to the SAP HANA data area, unmount the snapshot.
  - (4) If you renamed the vgdata volume group, renamed this data group back to its original name.

The original snapshotted data is now available in the HANA data area, and the database can be restored.

7. Perform a snapshot restore using the command line tools in SAP HANA.
  - (1) Verify that the snapshot data and the other data files of the original SAP HANA system are in the destination node's HANA data area (/hana/data).
  - (2) Verify that the database is stopped.
  - (3) Perform a command-line database restore as the <sid>adm user.

The time the command will run depends on the size of the database.

When the restore has finished, there is a confirmation message, as shown in Figure 2 on page 21.

```

2016-07-04T16:08:19+00:00 P038413 155b6aae0c1 INFO RECOVERY state of se
rvic: nameserver, hdidtest:30001, volume: 1, RecoveryExecuteTopologyAndSSFSReco
veryFinished

2016-07-04T16:08:50+00:00 P038413 155b6aae0c1 INFO RECOVERY RECOVER DAT
A finished successfully

recoverSys finished successfully: 2016-07-04 16:08:51
[140585782675200, 40.681] 0
[140585782675200, 40.681] << ending recoverSys, rc = 0 (RC_TEST_OK), after 40.67
9 secs

```

**Figure 2**

The SAP HANA database is running on the destination node, restored to the state of the snapshot reverted using Data Instance Director.

## SAP HANA Configuration File Handling in Hitachi Data Instance Director

When a snapshot is made of a SAP HANA database instance using Hitachi Data Instance Director, there is also a backup made of the current configuration files of the instance.

Data Instance Director automatically makes these files available when restoring a snapshot. You must analyze these files manually. Then, you must merge them manually, if necessary, with the configuration of the running SAP HANA instance on snapshot restore.

The configuration files are placed in the following directory:

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

Where the following is true:

- **\$HDID\_HOME** is the Data Instance Director installation directory on the Linux host.
- **<SID>** is the SAP HANA SID.
- **<hdid\_backup\_id>** is the snapshot ID shown in the **Application** column on the **Managed Storage** screen.



## For More Information

Hitachi Data Systems Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the Hitachi Data Systems [Global Services](#) website.

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the Hitachi Data Systems Corporate [Resources](#) website. Click the **Product Demos** tab for a list of available recorded demonstrations.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the Hitachi Data Systems Services [Training and Certification](#) website.

For more information about Hitachi products and services, contact your sales representative or channel partner or visit the [Hitachi Data Systems](#) website.

---

 **Hitachi Data Systems**



Corporate Headquarters  
2845 Lafayette Street  
Santa Clara, CA 96050-2639 USA  
[www.HDS.com](http://www.HDS.com)    [community.HDS.com](http://community.HDS.com)

Regional Contact Information  
**Americas:** +1 408 970 1000 or [info@hds.com](mailto:info@hds.com)  
**Europe, Middle East and Africa:** +44 (0) 1753 618000 or [info.emea@hds.com](mailto:info.emea@hds.com)  
**Asia Pacific:** +852 3189 7900 or [hds.marketing.apac@hds.com](mailto:hds.marketing.apac@hds.com)

© Hitachi Data Systems Corporation 2016. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. VSP is a registered trademark of Hitachi Data Systems Corporation. Microsoft, SQL Server, and Windows Server are trademarks or registered trademarks of Microsoft Corporation. All other trademarks, service marks, and company names are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation.

AS-547-00. October 2016.