



Backup strategy for SAP HANA databases

SnapCenter Software 4.8

NetApp
February 12, 2024

This PDF was generated from https://docs.netapp.com/us-en/snapcenter-48/protect-hana/task_define_a_backup_strategy_for_sap_hana_databases.html on February 12, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Backup strategy for SAP HANA databases 1
 - Define a backup strategy for SAP HANA databases 1
 - Automatic discovery of resources on Linux host 1
 - Type of backups supported 2
 - How SnapCenter Plug-in for SAP HANA Database uses consistency group Snapshot copies..... 2
 - How SnapCenter manages housekeeping of log and data backups 3
 - Considerations for determining backup schedules for SAP HANA database 3
 - Number of backup jobs needed for SAP HANA databases 3
 - Backup naming conventions for Plug-in for SAP HANA databases 4

Backup strategy for SAP HANA databases

Define a backup strategy for SAP HANA databases

Defining a backup strategy before you create your backup jobs helps you to have the backups that you require to successfully restore or clone your resources. Your service-level agreement (SLA), recovery time objective (RTO), and recovery point objective (RPO) largely determine your backup strategy.

About this task

An SLA defines the level of service that is expected and addresses many service-related issues, including the availability and performance of the service. RTO is the time by which a business process must be restored after a disruption in service. RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. SLA, RTO, and RPO contribute to the data protection strategy.

Steps

1. Determine when you should back up your resources.
2. Decide how many backup jobs you require.
3. Decide how to name your backups.
4. Decide whether you want to create a Snapshot copy-based policy to back up application-consistent Snapshot copies of the database.
5. Decide whether you want to verify the integrity of the database.
6. Decide whether you want to use NetApp SnapMirror technology for replication or NetApp SnapVault technology for long-term retention.
7. Determine the retention period for the Snapshot copies on the source storage system and the SnapMirror destination.
8. Determine whether you want to run any commands before or after the backup operation and provide a prescript or postscript.

Automatic discovery of resources on Linux host

Resources are SAP HANA databases and Non-data Volume on the Linux host that are managed by SnapCenter. After installing the SnapCenter Plug-in for SAP HANA Database plug-in, the SAP HANA databases on that Linux host are automatically discovered and displayed in the Resources page.

Automatic discovery is supported for the following SAP HANA resources:

- Single containers

After installing or upgrading the plug-in, the single container resources located on a centralized host plug-in will continue as manually added resources.

After installing or upgrading the plug-in, the SAP HANA databases are automatically discovered only on the SAP HANA Linux hosts, which are directly registered into SnapCenter.

- Multitenant database container (MDC)

After installing or upgrading the plug-in, the MDC resources located on a centralized host plug-in will continue as manually added resource.

You must continue to manually add the MDC resources on the centralized host plug-in after upgrading to SnapCenter 4.3.

For SAP HANA Linux hosts directly registered in SnapCenter, installing or upgrading the plug-in will trigger an automatic discovery for resources on the host. After upgrading the plug-in, for every MDC resource that was located on the plug-in host, another MDC resource will be automatically discovered with a different GUID format and registered in SnapCenter. The new resource will be in locked state.

For example, in SnapCenter 4.2, if E90 MDC resource was located on the plug-in host and registered manually, after upgrading to SnapCenter 4.3, another E90 MDC resource with a different GUID will be discovered and registered in SnapCenter.

Automatic discovery is not supported for the following configurations:

- RDM and VMDK layouts



In case the above resources are discovered, the data protection operations are not supported on these resources.

- HANA multiple-host configuration
- Multiple instances on the same host
- Multitier scale out HANA System Replication
- Cascaded replication environment in System Replication mode

Type of backups supported

Backup type specifies the type of backup that you want to create. SnapCenter supports File-Based Backup and Snapshot copy-based backup types for SAP HANA databases.

File-Based Backup

File-Based Backups verify the integrity of the database. You can schedule the file-based backup operation to occur at specific intervals. Only active tenants are backed up. You cannot restore and clone File-Based backups from SnapCenter.

Snapshot copy based backup

Snapshot copy-based backups leverage NetApp Snapshot copy technology to create online, read-only copies of the volumes on which the SAP HANA databases reside.

How SnapCenter Plug-in for SAP HANA Database uses consistency group Snapshot copies

You can use the plug-in to create consistency group Snapshot copies for resource

groups. A consistency group is a container that can house multiple volumes so that you can manage them as one entity. A consistency group is simultaneous Snapshot copies of multiple volumes, providing consistent copies of a group of volumes.

You can also specify the wait time for the storage controller to consistently group Snapshot copies. The available wait time options are **Urgent**, **Medium**, and **Relaxed**. You can also enable or disable Write Anywhere File Layout (WAFL) sync during consistent group Snapshot copy operation. WAFL sync improves the performance of a consistency group Snapshot copy.

How SnapCenter manages housekeeping of log and data backups

SnapCenter manages the housekeeping of log and data backups on the storage system and file system levels, and within the SAP HANA backup catalog.

The Snapshot copies on the primary or secondary storage and their corresponding entries in the SAP HANA catalog are deleted based on the retention settings. The SAP HANA catalog entries are also deleted during backup and resource group deletion.

Considerations for determining backup schedules for SAP HANA database

The most critical factor in determining a backup schedule is the rate of change for the resource. You might back up a heavily used resource every hour, while you might back up a rarely used resource once a day. Other factors include the importance of the resource to your organization, your service-level agreement (SLA) and your recovery point objective (RPO).

Backup schedules have two parts, as follows:

- Backup frequency (how often backups are to be performed)

Backup frequency, also called schedule type for some plug-ins, is part of a policy configuration. For example, you might configure the backup frequency as hourly, daily, weekly, or monthly.

- Backup schedules (exactly when backups are to be performed)

Backup schedules are part of a resource or resource group configuration. For example, if you have a resource group that has a policy configured for weekly backups, you might configure the schedule to back up every Thursday at 10:00 p.m.

Number of backup jobs needed for SAP HANA databases

Factors that determine the number of backup jobs that you need include the size of the resource, the number of volumes used, the rate of change of the resource, and your Service Level Agreement (SLA).

Backup naming conventions for Plug-in for SAP HANA databases

You can either use the default Snapshot copy naming convention or use a customized naming convention. The default backup naming convention adds a timestamp to Snapshot copy names that helps you identify when the copies were created.

The Snapshot copy uses the following default naming convention:

```
resourcegroupname_hostname_timestamp
```

You should name your backup resource groups logically, as in the following example:

```
dts1_mach1x88_03-12-2015_23.17.26
```

In this example, the syntax elements have the following meanings:

- *dts1* is the resource group name.
- *mach1x88* is the host name.
- *03-12-2015_23.17.26* is the date and timestamp.

Alternatively, you can specify the Snapshot copy name format while protecting resources or resource groups by selecting **Use custom name format for Snapshot copy**. For example, `customtext_resourcegroup_policy_hostname` or `resourcegroup_hostname`. By default, the time stamp suffix is added to the Snapshot copy name.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.