IBM Spectrum Protect

Blueprint for IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP HANA

Document version 1.1

Thomas Prause
IBM Spectrum Protect for SAP development team

Dominic Müller-Wicke
IBM Spectrum Protect portfolio architect

CONTENTS

1 Introduction....1

2 Implementation requirements....1

2.1 Hardware and software prerequisites....2

2.1.1 IBM Spectrum Protect server compute node....2

2.1.2 Data Protection for SAP HANA compute nodes....2

3 Setting up and configuring the solution....2

3.1 Setting up and configuring IBM Spectrum Scale....3

3.2 Setting up and configuring the IBM Spectrum Protect server....5

3.3 Installing the IBM Spectrum Protect backup-archive client and application programming interface (API)....7

3.4 Installing the Data Protection for SAP HANA client....7

3.5 Configuring IBM Spectrum Protect server stanzas....7

3.6 Configuring the Data Protection for SAP HANA client....9

3.7 Optional tuning of the Data Protection for SAP HANA client....10

4 Verifying the system configuration....11

4.1 Verifying General Parallel File System (GPFS) settings....11

4.2 Verifying that the client scheduler is running....12

4.3 Verifying settings for the Data Protection for SAP HANA client....12

5 Operations....13

5.1 Running regular backup operations....13

5.2 Running backup operations with different retention times....13

5.3 Restoring and recovering the database 14

5.4 Restoring the database to an alternative system 14

5.5 Restoring long-term retention backups 16

5.6 Maintenance 17

5.7 Troubleshooting 17

5.8 Performance test results 18

5.8.1 Number of sessions 18

5.8.2 Buffer size 19

Notices 21

Trademarks 22

1 Introduction

You can build a large-scale data protection solution for an SAP HANA database management system by using the following IBM Spectrum® Protect software:

- IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP HANA
- IBM Spectrum Protect server

In addition, the described solution uses IBM Spectrum Scale as a shared file system to connect multiple nodes of an SAP HANA clustered environment.

By following the steps precisely, you can prepare your system to run backup and recovery operations for SAP HANA data by using IBM Spectrum Protect. The described settings and options are designed to yield optimal performance for medium-scale systems but can be modified for other systems.

This document covers the following main topics:

- Installation and configuration of IBM Spectrum Scale to provide shared storage for SAP HANA scale-out environments
- Configuration of the IBM Spectrum Protect server
- Installation and configuration of the IBM Spectrum Protect backup-archive client
- Installation and configuration of Data Protection for SAP HANA
- Configuration and scheduling of monthly backups of the SAP HANA database
- Restoring the SAP HANA database to an alternate system

2 Implementation requirements

The listed hardware and software requirements apply to a medium-sized SAP HANA database with up to 4 TB of data distributed across four compute nodes. Each compute node is equipped with 1 TB of memory that is dedicated to the database.

The business requirements for protecting the database were defined as follows:

- To support database recovery for points in time going back at least one month, a daily database backup is required, along with redo log backups every 15 minutes.
- To support long-term recovery of the database, an additional monthly backup is required, and the data is retained for one year.
- The amount of redo log data per day is 25% of the overall database capacity.

Based on the business requirements, the following backup data capacity can be calculated for the first year:

- Daily database backup with 30 days of retention: 30 × 4 TB = 120 TB
- Daily redo log backup with 30 days of retention: 30 × 1 TB = 30 TB
- Monthly database backup with 1 year of retention: 12 × 4 TB = 48 TB
- IBM Spectrum Protect server capacity: 120 TB + 30 TB + 48 TB = 198 TB

When you plan the production environment, you must provision space to accommodate the growth of the SAP HANA database and the growth of the IBM Spectrum Protect server database.

Planning tip: To help optimize the throughput of backup operations, the described solution uses an IBM Spectrum Protect server as documented in the Blueprints for large reference architectures on IBM AIX operating systems.

2.1 Hardware and software prerequisites

Before you start to configure the solution, review the hardware and software prerequisites.

2.1.1 IBM Spectrum Protect server compute node

For server hardware and software requirements, see the IBM Spectrum Protect Blueprints.

2.1.2 Data Protection for SAP HANA compute nodes

The described configuration uses four compute nodes to share the SAP HANA workload.

Hardware	Requirements	Blueprint component	Detailed description
Server and network	8 virtual processor cores, 1.7 GHz or faster 1 TB RAM 10 Gb Ethernet	VMware ESXi Version 6.7 or 7.0	• Virtual machine with a virtual hardware level of 13 or later • VMware tools • 8-core virtual CPU • 1 TB virtual RAM • Virtual network adapter of type E1000E
Disks for storage	Virtual disks can be assigned as raw device mapping (RDM) disks or as virtual disks. Virtual disks must be thick provisioned, and VMware snapshots should not be taken of the virtual disks.	When using virtual disks, create the virtual disks for the operating system and for IBM Spectrum Scale network shared disks (NSDs) in different VMware datastores.	Operating system disk: • Size: 100 GB • Quantity: 1 IBM Spectrum Scale NSD: • Size: 300 GB • Quantity: 4

The hardware that is used for this sample environment might not meet the latest SAP requirements for production use. To confirm that your configured system environment meets SAP requirements, follow the instructions in SAP technote 1943937 (Hardware Configuration Check Tool – Central Note).

3 Setting up and configuring the solution

When you set up and configure the solution correctly, you can take advantage of the advanced backup and restore functionality offered by IBM Spectrum Protect.

The setup and configuration process includes the following major steps:

1. Setting up and configuring IBM Spectrum Scale
2. Setting up and configuring the IBM Spectrum Protect server

3. Installing the IBM Spectrum Protect backup-archive client and application programming interface (API)
4. Installing the Data Protection for SAP HANA client
5. Configuring IBM Spectrum Protect server stanzas
6. Configuring the Data Protection for SAP HANA client

3.1 Setting up and configuring IBM Spectrum Scale

You can configure the IBM Spectrum Scale file system to share resources among the four SAP HANA compute nodes. An SAP HANA compute node is equivalent to an IBM Spectrum Scale cluster node.

Procedure

Complete the following steps:

1. On the SAP HANA compute node, configure TCP/IP settings as specified in the SAP HANA product documentation. If you use multiple 10 Gb or 40 Gb Ethernet network adapters, the preferred method is to use different adapters for communication between the IBM Spectrum Protect server and Data Protection for SAP HANA, and between the IBM Spectrum Scale compute nodes.

2. On the SAP HANA compute nodes, install IBM Spectrum Scale by taking the following actions:
   a. Download the IBM Spectrum Scale base software package from Passport Advantage.
   b. Download the latest IBM Spectrum Scale fix pack from Fix Central.
   c. Install the IBM Spectrum Scale base software. Follow the instructions in Installing IBM Spectrum Scale on Linux nodes and deploying protocols.

3. Add the IBM Spectrum Scale `bin` directory (`/usr/lpp/mmfs/bin`) to your PATH environment variable.

4. Resolve the dependencies. If you are using Red Hat Enterprise Linux in a minimal installation configuration, install the following packages:

   ```
   yum install ksh
   yum install m4
   yum install elfutils-libelf-devel
   yum groupinstall ‘Development Tools’
   ```

5. Ensure that the kernel is portable by issuing the following command:

   ```
   mmbuildgpl
   ```

6. Update the `/etc/hosts` file on all SAP HANA compute nodes with the Internet Protocol (IP) address, short name, and long name of all SAP HANA cluster nodes.

7. Configure a Secure Shell (SSH) automatic login procedure without a password between the SAP HANA compute nodes. Take one of the following actions:

   - If the `/root/.ssh/id_rsa.pub` file is not available on the SAP HANA compute nodes, generate an `id_rsa.pub` file. The file contains a public key. Issue the following commands from all SAP HANA compute nodes that are part of the cluster:

   ```
   ssh-keygen –t rsa
   cd /root/.ssh
   chmod 640 /root/.ssh/authorized_keys
   ```

- If the /root/.ssh/id_rsa.pub file is available on the SAP HANA compute nodes, append the contents of the id_rsa.pub file from each of the other systems in the cluster to the authorized_keys file on the SAP HANA compute nodes.

8. Verify that the login procedure is configured. Log in to the other computers in the cluster from the IBM Spectrum Protect server by running the ssh command without using a password.
9. If the operating system on the IBM Spectrum Protect server is running a firewall, open several ports for incoming network connections from other systems in the IBM Spectrum Scale cluster. For instructions, see Securing the IBM Spectrum Scale system using a firewall.
10. Create the IBM Spectrum Scale cluster by taking the following actions:

    a. Prepare a node input file that contains appropriate entries for all of your SAP HANA cluster nodes. For more information about the input file, see mmcrcluster command.

    b. Create an IBM Spectrum Scale cluster by running the mmcrcluster command:

    ```
    mmcrcluster -N nodefile
    ```

    c. Assign an IBM Spectrum Scale license to all SAP HANA compute nodes by issuing the command:

    ```
    mmchlicense server -N hana01,hana02,hana03,hana04
    ```

    d. Change the cluster configuration so that the cluster is started automatically when the system restarts. Issue the command:

    ```
    mmchconfig autoload=yes
    ```

    e. Start the IBM Spectrum Scale cluster on all cluster nodes by issuing the command:

    ```
    mmstartup -a
    ```

    f. Verify that the IBM Spectrum Scale cluster was started correctly by issuing the following command. The command output should display a status of active for all cluster nodes:

    ```
    mmgetstate -a
    ```

11. Create the IBM Spectrum Scale file system by taking the following actions:

    a. Prepare for the creation of IBM Spectrum Scale NSDs by creating an NSD input file for each SAP HANA compute node. For more information about the input file, see mmcrnsd command.

    b. Create the IBM Spectrum Scale NSDs by using the disks that are attached to the SAP HANA cluster nodes. Issue the command:

    ```
    mmcrnsd -F nsd.conf
    ```

    c. Create the IBM Spectrum Scale file system by reusing the previously created nsd.conf:

```
mmcrfs hana -F nsd.conf -D nfs4 -B 256k -A yes -L 128M -k nfs4 -
m 1 -M 2 -Q no -r 1 -R 2 -S relatime -T /hana -z no
```

d. Mount the newly created file system on all SAP HANA cluster nodes:

```
mmmount hana -a
```

e. Verify that the file system was mounted correctly by issuing the command:

```
mmlsmount hana01
```

12. Create directories for the SAP HANA instance.

    a. To create the SAP HANA data directory, issue the command:

    ```
    mkdir /hana/data
    ```

    b. To create the SAP Hana redo log directory, issue the command:

    ```
    mkdir /hana/log
    ```

    c. To create the directory that is used to store SAP software packages to be shared in the cluster, issue the command:

    ```
    mkdir /hana/sw
    ```

13. Install SAP HANA on the newly created file system by following the instructions in the SAP HANA documentation.

In the sample environment, SAP HANA 2.0 SPS 05 was installed. A database instance was created with an ID of XXL with a default tenant database of the same name.

3.2 Setting up and configuring the IBM Spectrum Protect server

For the documented solution, the IBM Spectrum Protect server must be set up in a proxy node relation. The proxy node relation supports the ability to send backup data in multiple streams from all IBM Spectrum Scale cluster nodes.

Before you begin

1. Review the hardware and software requirements for the server in the IBM Spectrum Protect Blueprints. Determine whether your solution will be a small, medium, or large configuration and provision the software and hardware accordingly.
2. Configure the storage device, storage pool, and policy domain (`hana`) as described in Configure the IBM Spectrum Protect server. However, instead of setting up a single node, you must set up four IBM Spectrum Scale cluster nodes, which will share the workload.

Procedure

To create the proxy node relation, complete the following steps:

1. Log in to the administrative console of the IBM Spectrum Protect server by using an administrative user ID.
2. Verify whether the IBM Spectrum Protect domain that was created as part of the server Blueprint matches the requirements. If not, adjust the settings or create another domain. Pay special attention to the retention time of the associated archive copy groups. This sample configuration uses the `p9b-aix1_hana` domain, which is created by the server Blueprint. The domain expires data after 90 days. Use the following command to check the archive copy groups:

```
query copygroup p9b-aix1_hana type=archive
```

3. Create IBM Spectrum Protect nodes on the server. These nodes will be used for file system backup. You can use hostnames for the node names, as shown in the following example:

   ```
   register node node1 password dom=p9b-aix1_hana
   register node node2 password dom=p9b-aix1_hana
   register node node3 password dom=p9b-aix1_hana
   register node node4 password dom=p9b-aix1_hana
   ```

4. Create IBM Spectrum Protect nodes for the SAP HANA database on the server. One node will serve as the target node (`hana_xxl`) and the other four will be used as agent nodes (`hana_node1-4`):

   ```
   register node hana_xxl password dom=p9b-aix1_hana
   register node hana_node1 password dom=p9b-aix1_hana
   register node hana_node2 password dom=p9b-aix1_hana
   register node hana_node3 password dom=p9b-aix1_hana
   register node hana_node4 password dom=p9b-aix1_hana
   ```

5. Grant proxy authority to the four server nodes over the node for the database:

   ```
   grant proxynode target=hana_xxl agent=hana_node1
   grant proxynode target=hana_xxl agent=hana_node2
   grant proxynode target=hana_xxl agent=hana_node3
   grant proxynode target=hana_xxl agent=hana_node4
   ```

6. Optional: Create IBM Spectrum Protect nodes for backup operations with different retention periods, for example:

   ```
   register node hana_xxl_monthly password dom=p9b-aix1_hana
   register node hana_xxl_quarterly password dom=p9b-aix1_hana
   ```

7. If you specified nodes with different retention periods, grant proxy authority to the four server nodes over the additional nodes, for example:

   ```
   grant proxynode target=hana_xxl_monthly agent=hana_node1
   grant proxynode target=hana_xxl_monthly agent=hana_node2
   grant proxynode target=hana_xxl_monthly agent=hana_node3
   grant proxynode target=hana_xxl_monthly agent=hana_node4
   grant proxynode target=hana_xxl_quarterly agent=hana_node1
   grant proxynode target=hana_xxl_quarterly agent=hana_node2
   grant proxynode target=hana_xxl_quarterly agent=hana_node3
   grant proxynode target=hana_xxl_quarterly agent=hana_node4
   ```

8. To verify that all of the required nodes are created, issue the commands:

   ```
   query proxy
   query node hana*
   ```

3.3 Installing the IBM Spectrum Protect backup-archive client and application programming interface (API)

The IBM Spectrum Protect backup-archive client can be used to protect user files and SAP HANA configuration files that are included in the Data Protection for SAP HANA client backup.

Requirement: The IBM Spectrum Protect backup-archive client must be installed on *all* nodes of the SAP HANA cluster. By copying the installation files to the shared folder `/hana/sw,` the installation files can be accessed from all nodes.

To install the IBM Spectrum Protect backup-archive client and API, follow the instructions in Install the UNIX and Linux backup-archive clients.

3.4 Installing the Data Protection for SAP HANA client

Install the Data Protection for SAP HANA client on only *one* node in the SAP HANA cluster. The software is automatically distributed to all other nodes in the cluster.

To install the Data Protection for SAP HANA client, follow the instructions in Installing.

3.5 Configuring IBM Spectrum Protect server stanzas

To ensure that the SAP HANA database and other files (for example, configuration files) can be backed up, you must configure IBM Spectrum Protect server stanzas.

Procedure

Complete the following steps:

1. Log in with the root user ID and edit the file `/opt/tivoli/tsm/client/ba/bin/dsm.sys`. Create a server stanza for backing up the local SAP HANA configuration. Ensure that the node names used in the server stanza match the node names that were previously specified for file system backup on the IBM Spectrum Protect server. For this environment, the file system backup will include only SAP HANA configuration files.

   The configuration is similar to the following example:

```
*
* The server stanza that is used for backup of SAP Hana
configuration files.
* This node can be used for additional file backup required
* on the SAP Hana cluster nodes.
*

SErvername            p9b-aix1
COMMMethod            TCPip
TCPPort               1500
TCPServeraddress      p9b-aix1.storage.tucson.ibm.com
passwordaccess        generate
nodename              node1
webports              1505,1506
managedservices       schedule
schedmode             prompted
schedlogname          /var/log/dsmsched.log
errorlogname          /var/log/dsmsched.err
exclude               /.../*
include               /hana/shared/X3L/global/hdb/custom/config/…/*
include            /usr/sap/XXL/home/…/*
exclude.dir        /usr/sap/XXL/home/sp_pwd
```

The preferred method is to exclude all files first and then include only selected folders. In this way, the changes in the SAP HANA configuration and the home directory of the database user are backed up. But the folder with the encrypted IBM Spectrum Protect passwords is not backed up.

Restriction: For configuration files, do not use the following path:
`/usr/sap/XXL/SYS/global/hdb/custom/config`

This path represents a symbolic link, which is not followed by default.

2. Create the following empty file:

   `/opt/tivoli/tsm/client/ba/bin/dsm.opt`

3. Edit the `/opt/tivoli/tsm/client/api/bin64/dsm.sys` file and create a server stanza to be used for the SAP HANA database backup. Ensure that the node names in the server stanzas match the node names that were previously configured for SAP HANA database backup operations on the IBM Spectrum Protect server. The server stanza is similar to the following example:

   ```
   *
   
   * The server stanza that is used for SAP Hana database
   backup.
   *
   SErvername  p9b-aix1-hana
                  COMMMethod         TCPip
                  TCPPort            1500
                  TCPServeraddress   p9b-
   aix1.storage.tucson.ibm.com
                  passwordaccess     generate
                  nodename           hana_node1
                  passworddir        /usr/sap/XXL/home/sp_pwd
   ```

   Restriction: Ensure that the password directory does not reside in a file system that is shared among the SAP HANA nodes.

4. Create the following empty file:

   `/opt/tivoli/tsm/client/api/bin64/dsm.opt`

Guidelines:

- If multiple server entries are required in the `/opt/tivoli/tsm/client/api/bin64/dsm.sys` file, the server stanza that is used for SAP HANA database backup operations should be the first stanza in the `dsm.sys` file. In this way, no server entry is required in the client options file, `dsm.opt`. The positive side effect is that a restart of the SAP HANA database to pick up changed environment variables like DSMI_CONFIG can be avoided.
- Do not add an entry for the ASNODE parameter in the server stanza for the SAP HANA database backup. The Data Protection for SAP HANA client will use the correct ASNODE parameter during the backup operation.

   - The preferred method is to use the PASSWORDACCESS GENERATE parameter in the server stanza. In a scale-out environment, this setting helps to meet the requirement that backup data must be restorable on any SAP HANA node. For both scale-out and single-node environments, this setting simplifies the configuration and handling of backup operations with different retention periods.

5. Verify the configuration by using the IBM Spectrum Protect backup-archive client command `dsmc`. The `dsmc` command-line client must be able to connect to the server without issues. Example:

   ```
   dsmc query opt
   ```

   When the command is run for the first time, you are prompted to provide the node name and password. During subsequent runs, the command displays node options without a prompt.

6. Start the IBM Spectrum Protect client acceptor:

   ```
   dsmcad
   ```

Setting up additional schedules

The client acceptor service enables the IBM Spectrum Protect server to schedule operations on this host. For SAP HANA scale-out environments, it is sufficient to run the scheduler on one host if backup operations cover data from the shared file system only. If operating system files also must be backed up, each host must have its own scheduler running. For example, to set up a schedule for backing up configuration files on HANA_NODE1, you would complete the following steps:

1. Start an IBM Spectrum Protect administrative client session and create a schedule to back up configuration files. For example, if the domain is HANA and the schedule name is HANA_CONFIG, you could run the following command, which will apply default values. Backup operations will be incremental, starting immediately:

   ```
   define schedule hana hana_config
   ```

2. Associate the node (for example, HANA_NODE1) with the newly defined schedule:

   ```
   define assoc hana hana_config hana_node1
   ```

3.6 Configuring the Data Protection for SAP HANA client

You can use a script to configure the Data Protection for SAP HANA client.

Before you begin

1. Log in as the root user on the SAP HANA cluster node that was previously used to install the Data Protection for SAP HANA client.

2. Verify access rights. The configuration script that comes with Data Protection for SAP HANA requires that the SAP HANA user ID that is used for data protection operations has the following privileges:

   - CATALOG READ privileges are required for front-end-capacity reporting and checking for obsolete backup generations (MAX_VERSIONS).

   - INIFILE ADMIN, SERVICE ADMIN, and DATABASE ADMIN privileges are required to make configuration changes in the SAP HANA environment during setup.

Tip: As part of the installation process, an entry is created in the SAP HANA `hdbuserstore` set. This entry is used by the Data Protection for SAP HANA client for data protection operations.

Procedure

1. To initiate the configuration, run the following script:
   `/opt/tivoli/tsm/tdp_hana/setup.sh`

After you enter the credentials to access the database, the software is distributed to all SAP HANA nodes that belong to the database instance.

2. Provide the mandatory parameters as prompted by the script. Your responses are similar to the following examples:

```
Do you want to use automatic password handling (passwordaccess
generate) [y/n] y
Please enter the IBM Spectrum Protect server name as defined in
dsm.sys: p9b-aix-hana
```

When prompted for the node name, enter the node that is designated as the target for database backup operations.

```
Enter the IBM Spectrum Protect node that will be used to store
the data from all SAP HANA nodes.
This will be used for parameter ASNODE in the IBM Spectrum
Protect for ERP profile: HANA_XXL
```

Finally, the SAP HANA configuration will be updated automatically to supply the appropriate profiles to the Data Protection for SAP HANA client for backup and restore operations. For SAP HANA 2.0 SPS 05, the threshold for multiple backup sessions will be set to 2 GB if more than one session was selected during configuration.

3.7 Optional tuning of the Data Protection for SAP HANA client

The previously described setup procedure creates a default configuration for the Data Protection for SAP HANA client. You can modify the configuration by editing the following profile:

`/usr/sap/SID/SYS/global/hdb/opt/hdbconfig/init`*`SID`*`.utl`

where *SID* is the system identifier.

Procedure

Optionally, update the profile based on the following guidelines:

- **Optimizing data deduplication.** If IBM Spectrum Protect disk container pools are used as a target for SAP HANA backup operations, you can increase the value of the `BUFFSIZE` option to 8 MB to help optimize data deduplication. In the configuration file, you can set the value as shown in the following example:

  ```
  BUFFSIZE   8388608
  ```

  However, the updated setting causes partially filled buffers to be stored during backup operations of small objects like the SAP HANA backup catalog or redo log files, and this can impact the performance of backup operations for log data and allocates more space on the IBM Spectrum Protect server than actually needed. To overcome the impact, you can use dedicated configuration files for database and redo log backup operations. To configure dedicated log backup operations, see technote 275755.

- **Setting retention policies.** The preferred method is to specify settings in the SAP HANA cockpit. The IBM Spectrum Protect setting of the archive copy group associated with the selected management class should be infinite.

  If the Data Protection for SAP HANA client is configured to delete obsolete backup generations (parameter `MAX_VERSIONS > 0`), the retention policy in SAP HANA should remove only backup entries for backups that are expired and have been deleted by the Data Protection for SAP HANA client. If expiration by the Data Protection for SAP HANA client is not enabled (parameter `MAX_VERSIONS = 0`), select the option to delete backups in the backup interface (Backint for SAP HANA database) when configuring the retention policy in the SAP HANA cockpit.

  Tip: This guideline refers to the retention and deletion of obsolete backup generations and should not be confused with retention sets, which are not supported in IBM Spectrum Protect for ERP. Keep in mind that Data Protection for SAP HANA stores all data as archive objects on the IBM Spectrum Protect server.

- **Specifying the target of backup operations.** The same processing parameters are used for incremental and differential backup operations and for redo log backup operations. As a result, the data is stored in the same management classes. Although the amount of data for incremental or differential backups is much less than the amount for a full database backup, the amount for incremental or differential backups can be considerably greater than the amount of data calculated for the storage of the redo logs. The configuration can be modified to direct the incremental and differential backups to the same management class as the full database backups. For instructions, see technote 533927.

4 Verifying the system configuration

Before you start running the new storage solution in a production environment, verify the system configuration.

4.1 Verifying General Parallel File System (GPFS) settings

To help ensure successful backup processing of the SAP HANA database, you must verify that the underlying IBM Spectrum Scale cluster is available and restarts automatically after a restart of one or more cluster nodes.

Procedure

Review the following settings:

1. Verify that the cluster is active and the file system is mounted by issuing the following commands:

```
# mmgetstate -a

 Node number  Node name        GPFS state
-------------------------------------------
       1      spppod3-vm6      active
       2      spppod3-vm8      active
       3      spppod3-vm9      active
       4      spppod3-vm10     active


# mmlsmount hana01 -L
```

In this example, the output shows that file system HANA01 is mounted on four nodes:

```
  192.168.64.70      spppod3-vm6
  192.178.64.72      spppod3-vm8
  192.168.64.74      spppod3-vm10
  192.168.64.73      spppod3-vm9
```

2. Verify that the cluster is enabled for automatic start after a node restart by issuing the command:

   ```
   # mmlsconfig autoload


   autoload yes
   ```

3. Verify that the file system is configured for automatic mount after a restart of the cluster or a cluster node by issuing the command:

   ```
   # mmlsfs hana -A


   flag                value                    description
   ------------------- ------------------------ -------------------
   ---
    -A                 yes                      Automatic mount
   option
   ```

4.2 Verifying that the client scheduler is running

To help ensure that backup operations occur as expected, verify that the client scheduler is running.

Procedure

1. Log in as the root user on the client node.
2. On all nodes where a client schedule is running, issue the command:

   ```
   dsmc query backup /hana/\*/custom/config/\*
   ```

   The command should display the backups of the SAP HANA configuration files without prompting for a node name or password, assuming that the schedule ran at least once. If no result is returned, check the include-exclude rules in the client system options file, `/opt/tivoli/tsm/client/ba/bin/dsm.sys.`
3. On all nodes where a client schedule is running, verify that the `dsmcad` process is running and check the scheduler error log:

   ```
   # ps -ef | grep dsmcad
   root     1424106       1  0 Apr16 ?        00:00:02 dsmcad
   # tail /var/log/dsmsched.err
   ```

4.3 Verifying settings for the Data Protection for SAP HANA client

You can verify settings for the Data Protection for SAP HANA client by running the following script at any time:

```
/opt/tivoli/tsm/tdp_hana/setup.sh
```

If a Data Protection for SAP HANA profile exists, the profile will not be changed. However, the script checks remote HANA nodes to ensure that the appropriate software is installed. In addition, the entries in the `hdbuserstore` are re-created and all modifications in the SAP HANA configuration are re-enabled. Optionally, you can enter the passwords for the IBM Spectrum Protect nodes in case the passwords have been changed on the server. The script automatically configures nodes that were added to the cluster after the initial installation of Data Protection for SAP HANA.

Procedure

1. As the database user, run the following command:

   ```
   /opt/tivoli/tsm/tdp_hana/backfm -p \
   /usr/sap/XXL/SYS/global/hdb/opt/hdbconfig/initXXL.utl
   ```

   All existing backups are displayed in a list.
2. To see the date when a selected object was backed up, press F6.
3. To verify that the expiration policy is working as expected, scroll down to the last element.

5 Operations

5.1 Running regular backup operations

Ensure that daily backup operations run on a predefined schedule. Database redo log files are backed up automatically by SAP HANA. You can schedule full database backups in the SAP HANA cockpit.

In the sample configuration in the IBM lab, one full backup operation was run daily. Depending on your business requirements, you can run additional full, incremental, or differential backups. For instructions about scheduling backups in the SAP HANA cockpit, see the *SAP HANA Administration Guide*.

5.2 Running backup operations with different retention times

To ensure that obsolete backup generations from daily backups are expired automatically, the backups with a different retention time must be separated from daily backups. Otherwise, backups with a different retention time would be counted and removed by the cleanup procedure for the daily backups. To prevent this, you must direct special backup operations to dedicated IBM Spectrum Protect nodes that are not used for daily backups.

For monthly and quarterly backup operations, special nodes were created as part of the IBM Spectrum Protect server configuration. To use these nodes, copy the Data Protection for SAP HANA profile that is used for the daily backups under a new name. For example, change the profile name to `initXXL_monthly.utl` and change the ASNODE parameter value to `hana_xxl_monthly`. Data Protection for SAP HANA is configured to use the nodes HANA_NODE1, and so on, for authentication. If you specify a different value for the ASNODE parameter, you do not have to set a password.

For every monthly backup operation, it would be necessary to change the SAP HANA configuration to use this special profile and revert the changes immediately after the backup operation. This process is cumbersome and should be automated as far as possible. This Blueprint is bundled with a `monthly_backup.sh` script that covers all the required steps. Copy this script as the SAP HANA database user to the file system of the SAP HANA node. If the script is located in the shared file system, the backup operations can be scheduled on any SAP HANA node. Otherwise, you must select the appropriate

SAP HANA node when scheduling the backup operations. In the sample configuration, the script was copied to the `/hana/tools` directory, which is in the file system that is shared among all SAP HANA nodes. The script was renamed to `monthly_backup_xxl.sh`.

Before using the script, the variables in the configuration section of the script must be reviewed and potentially modified for the current environment. With the default settings, a backup of the default tenant of the instance owned by the user ID that calls the script will be created.

To run the script monthly, you can set up a schedule in the IBM Spectrum Protect administrative console by using the command:

```
# define schedule hana monthly_backup action=command objects=”su –
xxladm -c /hana/tools/monthly_backup_xxl.sh” perunit=month
startdate=04/01/2021
```

The scheduler is operated by the root user. Therefore, a change of the user ID to the SAP HANA instance owner is required to run the script. Then, you can associate the node of the SAP HANA system where the script is located with this schedule:

```
# define assoc hana monthly_backup hana_node1
```

The script will then be invoked on the first day of each month by the IBM Spectrum Protect scheduler. Any other scheduler could be used as well.

If the backup is restored, the SAP HANA backup catalog will not contain an entry for the backup. As a result, the backup can be restored only by the tag. Therefore, the tag that is assigned to the backup with the BACKUP_TAG parameter in the `monthly_backup.sh` script must be unique. For example, if multiple “monthly” backups are created with the same tag, only the most recent backup can be restored by using this tag. The default setting would create a unique tag every day.

In the same way, backup operations that run quarterly, yearly, or at other intervals are possible. However, you must separate these backup operations from daily backups to prevent the special backups from being deleted by expiration processing.

5.3 Restoring and recovering the database

In the SAP HANA cockpit, you can initiate restore and recovery operations for the database. By using the recovery wizard, you can restore a specific backup or recover all data going back to a specific point in time. In the latter case, SAP HANA can automatically restore data from incremental or differential backups to optimize the recovery.

5.4 Restoring the database to an alternative system

The term *alternate restore* refers to restoring a database backup to a system or environment (target system) other than the system where the backup was created (source system). The restored database will be on a new system with a different system identifier (SID).

Procedure

1. Set up and configure the target system as described in the previous sections: Installing the IBM Spectrum Protect backup-archive client and application programming interface (API), Installing the Data Protection for SAP HANA client, Configuring IBM Spectrum Protect server stanzas, and Configuring the Data Protection for SAP HANA client. Verify that the system can back up logs and the database.
2. On each SAP HANA node, append a server stanza to the `/opt/tivoli/tsm/client/api/bin64/dsm.sys` file. The new stanza points to the server where the data from the source system is located. This stanza must *not* use

the passwordaccess generate option. In the following example, the new stanza includes the server name p9b-aix1-hana-rest. The new stanza lists the source server address as p9b-aix1.storage.tucson.ibm.com, which matches the previous stanza:

```
SErvername  p9b-aix1-hana
               COMMMethod         TCPip
               TCPPort            1500
               TCPServeraddress   p9b-aix1.storage.tucson.ibm.com
               passwordaccess     generate
               nodename           hana_node1
               passworddir        /usr/sap/XXL/home/sp_pwd


SErvername  p9b-aix1-hana-rest
               COMMMethod         TCPip
               TCPPort            1500
               TCPServeraddress   p9b-aix1.storage.tucson.ibm.com
```

3. Copy the Data Protection for SAP HANA profile and the configuration file (*.bki) in the folder and save them under a new name that matches the SID of the source system. For example, if you plan to restore the XXL system with data from the SRC system, you would specify the following profiles:

   ```
   /usr/sap/XXL/SYS/global/hdb/opt/hdbconfig/initXXL.utl and
   /usr/sap/XXL/SYS/global/hdb/opt/hdbconfig/initSRC.utl
   ```

   Accordingly, the configuration files would have the following names:

   ```
   /usr/sap/XXL/SYS/global/hdb/opt/hdbconfig/initXXL.bki and
   /usr/sap/XXL/SYS/global/hdb/opt/hdbconfig/initSRC.bki
   ```

4. Edit the CONFIG_FILE parameter in the initSRC.utl profile to point to initSRC.bki. The settings from the initSRC.utl profile will be used to retrieve the data. Immediately after the recovery is complete, the system would start backing up log files by using the settings from the initXXL.utl profile.
5. Edit the initSRC.utl profile to point to the appropriate IBM Spectrum Protect server and node where the backup data from the SRC source system is located.

   In the sample configuration, the following stanza was changed:

   ```
   SERVER         p9b-aix1-hana
     SESSIONS           7
     ASNODE             hana_xxl
     BRBACKUPMGTCLASS   mdb
     BRARCHIVEMGTCLASS  mlog1 mlog2
   ```

   The stanza was updated as shown:

   ```
   SERVER         p9b-aix1-hana-rest
     SESSIONS           7
     ASNODE             hana_xxl
     ADSMNODE           hana_src
     BRBACKUPMGTCLASS   mdb
     BRARCHIVEMGTCLASS  mlog1 mlog2
   ```

Note the changed value for the SERVER parameter. The ASNODE remains unchanged. Because the data will be restored from this node. The parameter ADSMNODE is added and denotes the name of the Spectrum Protect node that will be used to connect to the server. The node is no longer listed in the corresponding stanza in dsm.sys. The change of parameter PASSWORDACCESS from GENERATE to PROMPT is required because the data from node HANA_SRC was stored under a different operating system user. When using the PASSWORDACCESS GENERATE parameter, it is not possible to access the backup data from a different user. By using the default PASSWORDACCESS PROMPT parameter, the required access is granted.

6. To verify the configuration, issue the following command:

   ```
   /opt/tivoli/tsm/tdp_hana/hdbbackint -p \
   /usr/sap/XXL/SYS/global/hdb/opt/hdbconfig/initSRC.utl \
   -f inquire -u src
   ```

   The value for option `–u` is the SID of the source system and is not case sensitive. At the prompt, enter #NULL and press Ctrl + D. As result, all backup objects should be shown. If the output displays a `#NOTFOUND` message, review the configuration.

7. Start the restore operation in the SAP HANA cockpit:
   a. In the **Database Management** view of the SAP HANA instance, select the tenant database that is the target of the restore operation. (In the sample configuration, the target is XXL.)
   b. From the **Tenant Actions** menu, click **Copy Tenant Using Backup**.
   c. Follow the instructions in the wizard to complete the restore operation. In one of the windows, enter the name of the SAP HANA instance (SID) and the name of the database.

5.5 Restoring long-term retention backups

Backups that are retained for a longer period and are therefore no longer listed in the SAP HANA backup catalog require a special restore procedure.

Before you begin

1. Shut down the database.
2. Rename the Data Protection for SAP HANA profile, `init`*`SID`*`.utl`, to preserve it for future use.
3. Replace the profile with the profile that was used to create backups with a different retention time. For example:

   ```
   $ cd /usr/sap/XXL/SYS/global/hdb/opt/hdbconfig
   $ cp initXXL.utl initXXL.utl_orig
   $ cp initXXL_monthly.utl initXXL.utl
   ```

About this task

The starting point for the restore operation in the SAP HANA cockpit is the same as for the alternate restore.

Procedure

1. In the SAP HANA cockpit, select the **Copy Tenant Using Backup** option to restore a backup without using the backup catalog.

2. When prompted for the system identifier and database name, enter the values of the current database subject to restore (in the sample configuration, XXL would be entered for both values).
3. Follow the steps in the wizard.
4. When you are prompted to specify a backup prefix that uniquely identifies the backup to be restored, enter the prefix that is assigned to the backup with the option BACKUP_TAG in the `monthly_backup.sh` script.
5. Immediately after the restore operation is completed, revert the changes to the profiles:

   ```
   $ mv initXXL.utl_orig initXXL.utl
   ```
6. Run a full database backup operation.

5.6 Maintenance

To streamline maintenance operations, use the SAP HANA scheduler for the expiration of obsolete backup generations. The scheduler will remove both the backup objects from the IBM Spectrum Protect server and the entries from the SAP HANA backup catalog.

Alternatively, you can use the Data Protection for SAP HANA client for the same operation by specifying the parameter setting `MAX_VERSIONS > 0`. The Data Protection for SAP HANA client then deletes obsolete backup objects from IBM Spectrum Protect. The related entries in the SAP HANA backup catalog must then be removed manually.

5.7 Troubleshooting

Failed backup operations trigger alerts in the SAP HANA cockpit. In case of a failed backup operation, take one or more of the following actions to mitigate or solve the problem:

- Verify whether a subsequent backup operation was successful by opening the alert in the cockpit and scrolling down to the history. If you detect that a subsequent backup operation was successful, the alert indicated a temporary problem that no longer exists.

- Check entries in the SAP HANA log file `backint.log`. You can view the log file in the Database Explorer window in the SAP HANA cockpit. Complete the following steps:
  1. In the cockpit, click **Alert and Diagnostics > View trace and diagnostic files.**
  2. In the navigation panel, expand the database.
  3. For a scale-out configuration, navigate to the host where the master server runs and select the **other** entry under this host.
  4. View the `backup.log` file, which records the backup- and restore-related actions that were completed by SAP HANA.
  5. View the `backint.log` file, which records the input and output of the Data Protection for SAP HANA client. Depending on the type of the issue, you might have to search backward in the `backint.log` file to find information related to the alert. For more information, see the *IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP HANA Installation and User's Guide*.

  Tip: In a large scale-out cluster, it might be cumbersome to find the right host. If you can open a shell on one of the SAP HANA nodes, the following command shows all existing logs for the system database:

  ```
  > ls ${DIR_INSTANCE}/*/trace/ba*log
  /usr/sap/XXL/HDB00/node1/trace/backint.log
  /usr/sap/XXL/HDB00/node1/trace/backup.log
  ```

The following command shows the logs for all tenant databases:

```
> ls ${DIR_INSTANCE}/*/trace/DB*/ba*log
/usr/sap/XXL/HDB00/node1/trace/DB_XXL/backint.log
/usr/sap/XXL/HDB00/node1/trace/DB_XXL/backup.log
/usr/sap/XXL/HDB00/node1/trace/DB_X3L/backint.log
/usr/sap/XXL/HDB00/node1/trace/DB_X3L/backup.log
```

If the master name server was moved to another node, multiple log files might be available per database. In that case, use the timestamp to identify the current log file.

For more instructions about troubleshooting, see Troubleshooting.

5.8 Performance test results

This section describes the option settings that were used in the IBM lab for performance testing of the documented solution.

For general information about performance tuning, see Tuning performance.

For general information about setting up and tuning the IBM Spectrum Protect server, see the IBM Spectrum Protect Blueprints.

5.8.1 Number of sessions

SAP HANA can split the data from large services (usually the index server) into multiple streams. When you run the `setup.sh` script, the `init`*`SID`*`.utl` profile and the settings in SAP HANA are aligned to use matching values. When you increase the number of sessions, this alignment must be considered. To simplify this procedure, you can set the values in the `init`*`SID`*`.utl` profile to a high number. For instance, set both the parameter MAX_SESSIONS and SESSIONS in the server stanza to 16. The values are an upper limit that cannot be exceeded. The actual number of sessions per service depends on the SAP HANA configuration parameter, parallel_data_backup_backint_channels. By increasing the value of this parameter, the performance impact can be tested. Keep in mind that this value determines the number of sessions per SAP HANA service. In the sample environment, four index servers were running. When the parallel_data_backup_backint_channels parameter was set to 4, the result was 16 sessions on the IBM Spectrum Protect server.

Observations about the number of sessions

Testing in the IBM lab environment revealed that a single session targeted to an IBM Spectrum Protect disk container pool is limited to approximately 120 MB/s. Because four nodes are configured, the overall backup performance is approximately 465 MB/s.

Multiple parallel sessions yield comparable performance per session. Increasing the SAP HANA parameter parallel_data_backup_backint_channels to use multiple sessions per service improved performance as shown in the table.

Sessions from other IBM Spectrum Protect client workloads can impact performance. As shown in IBM lab tests, it is beneficial to complete multiple backup operations before modifying the number of sessions again.

Table 1: Performance results from IBM lab tests

Number of sessions per SAP HANA service	**Full database backup performance**
1	465 MB/s

Number of sessions per SAP HANA service	**Full database backup performance**
2	770 MB/s
3	1005 MB/s
4	1254 MB/s
5	1452 MB/s
6	1598 MB/s
7	1711 MB/s
8	1706 MB/s

5.8.2 Buffer size

The size of the transfer buffers in Data Protection for SAP HANA can affect performance. The buffers are controlled with the BUFFSIZE profile parameter. As mentioned previously, the BUFFSIZE value can have a positive effect on the data deduplication ratio of an IBM Spectrum Protect disk container pool. For more information, see Optional tuning of the Data Protection for SAP HANA client.

Observations about buffer size

For the transport of small objects like redo log files or the backup catalog file, a large BUFFSIZE value can increase processor usage and occupy more space on the IBM Spectrum Protect server than necessary.

For example, if the BUFFSIZE value is set to 8 MB, and the goal is to protect a backup catalog with a size of 2 MB, the result is 6 MB of unused capacity in the buffer. This additional and unused capacity will also be occupied on the IBM Spectrum Protect server.

When you set a larger buffer, you might be able to improve the performance of database backup operations. This advantage generally outweighs the disadvantage of having unused capacity in the buffer. For instructions about specifying a smaller BUFFSIZE value for redo logs and the backup catalog, see Optional tuning of the Data Protection for SAP HANA client

Document revisions

Version number	Description
1.0	Initial version
1.1	Revised instructions in chapter 5.4

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Trademarks