

Business Data Restore and Recovery: Enterprise Database SAP HANA prospective

Rajesh Azmeera

SAP HANA Cloud Architect, NJ, USA.

¹Corresponding Author: rajsapbasis007@gmail.com

Abstract: Making backups of collected data is critically important in any industry. Backups protect against human errors, hardware failure, virus attacks, power failure, and natural disasters. Backups can help save time and money if these failures occur. No matter what the scenario or business need, learn how Metallic Database Backup can protect your multi-cloud environments. Database Backup customers can expect: Single control pane for managing your cloud and hybrid cloud environments, Enterprise-ready, hybrid-cloud protection, Agile deployment, simplified backup, automatic updates, Zero-trust access model, with early threat detection, Manage multiple database services from a single interface, Unlimited retention. It's been challenging to backup and restore data from business continuity process management point of view.

Keywords: SAP Enterprise Database, Backup, Recovery, HANA Database, RTO (Recovery Time Objective), RPO (Recovery Point of Objective)

Date of Submission: 11-10-2023

Date of acceptance: 25-10-2023

I. Introduction:

Every data or database should be backed up to restore back in critical times. Different Database backup types available to achieve strong RTO[1]. Strong RTO and RPO depends on two factors Database backups and restore[2]. The quicker the restore will reduce business downtime and the length of data will reduce loss of business data.

1.1 Full Database Backup

A full database backup contains all used data pages and the log files written during the backup. The transaction log is not truncated when you perform a full database backup.

You always need a full database backup to restore the database. You cannot restore a differential database backup or a transaction log backup without a full database backup.

Full database backups may result in a significant amount of disk I/O and should therefore be performed at a time when the workload is low.

The data of the SAP, master, and msdb database must be backed up at regular intervals. These backups are essential to enable restore operations, if the database is damaged. In general, the SAP database <SAPSID> should be backed up once a day. It is sufficient to backup the master and msdb databases once a week or month, depending on the amount of database activity in your system. For administrative reasons, it is however easier to backup the master and msdb databases whenever you backup the SAP data. This hardly has an impact on system throughput as these backups do not require much additional space or time.

1.2 Differential Database Backup

A differential database backup contains only those data pages that have been modified since the last full database backup. Apart from this, a differential database backup is identical to a full database backup.

The size of a differential database backup is much smaller than the size of a full database backup. Differential database backups are optional. The backup strategy will become more complicated if it includes differential database backups. An example of integrating differential database backups in a backup strategy is to perform a full database backup only once a week and to perform differential database backups daily.

1.3 Transaction Log Backup

The transaction log of the SAP database records all changes made to the database. It may never be deleted and must be backed up separately.

Transaction log backups save the log files. They are mandatory when you use the Full or Bulk-Logged Recovery Model since they are needed to truncate the log. Two subsequent transaction log backups do not contain redundant information. Transaction log backups have almost no impact on performance and can therefore be performed when the workload is high.

SAP recommends that you schedule a transaction log backup every 30 to 60 minutes on a normal workday. If your SAP system(s) are used productively at weekends, schedule additional backups for weekends.

1.4 Log Shipping

You can use log shipping to feed transaction logs from one database to another on a constant basis. Continually backing up the transaction logs from a source database and then copying and restoring the logs to a destination database keeps the destination database synchronized with the source database.

You can use the Database Maintenance Plan Wizard to configure log shipping as a maintenance task that can be scheduled to run automatically at defined intervals.

Log shipping enables you to run a backup server and provides high availability as required by large enterprise systems. Log shipping requires extra administration and well-defined procedures to monitor the status of all log-shipping servers.

II. Discussion: SAP HANA Backup solutioning:

2.1 HANA DB Backup Management

To ensure the recoverability of the database for up to 60 days (Production) and 30 days (non prod) by facilitating regular database backups. We need to ensure recoverability to any point in time within the backup retention period and from any longer retention backup without roll forward. Success criterion is the Backup success rate to ensure backups are taken so the RPO and RTO can be achieved.

2.2 HANA DB Restore Management

Have the Strategy and Resources in place such that in the event of a Failure Scenario, the HANA DB is returned to service within the RTO (Recovery Time Objective) and RPO (Recovery Point Objective) set by your organization. Backups are verified by means of redirected restore once per month.

Success criteria are: Restore testing is performed once per month to ensure the validity of backups – all production backups and some non-productive backups are tested once in the backup retention cycle. In addition in a Failure Scenario service is returned to the business within the RPO and RTO.

2.3 What is RTO and RPO?

Backup and recovery apply to an entire HANA database. It is not possible to back up and/or recover individual database objects in isolation (or individual nodes). In addition to the database backup, external configuration files must always be saved to guarantee full recoverability of the database.

HANA database backups are full backups, differential or incremental are possible.

In addition HANA, writes all changes to log files. Log switches are triggered when the log is full or after a defined (configured) time.

2.4 RTO (Recovery Time Objective)

RTO is the time that it takes to recover the system.

Consists of:

The time taken to restore the data files (using a full backup and possibly incremental or differential backups)

The time needed to recover the database to the required point in time (Recovery Time)

Recovery time is influenced by the number of logs that need to be applied. It depends on the time difference between the full data backup to be restored and the point in time that needs to be reached by applying the database changes that are stored in the archive log files. The total time needed depends on:

The time needed to restore the archive log files from backup media

The time needed to apply the logs to the database

The restore time can be affected by any bottlenecks or problems in any of the components involved in the restore: TSM, drives, SAN, server and storage.

The recovery time is primarily dependent on the number of logs that are written.

To minimize the Recovery Time Objective (RTO), some systems are utilizing HANA system replication. Due to the time needed to restore and recover a system using TSM, where a system is using HANA system replication the preferred option to restore service is to perform a takeover to the secondary node if possible.

2.5 Recovery Point Objective (RPO)

The RPO (Recovery Point Objective) is the maximum time that is possible to be lost when restoring data. An RPO of 10 minutes means in the worst case scenario, 10 minutes of data loss is possible when the system needs to be restored and recovered.

To ensure the agreed Recovery Point Objective can be met, the log file switch frequency needs to be set accordingly with parameter `log_backup_timeout_s`.

III. Methodology:

How to perform backup and restore on HANA Database. You can use studio, sql commands and cockpit.

In this article, we will demonstrate how to perform backup and restore using cockpit on AWS storage S3

You need third party agent backint to trigger backups to S3

Deploy backint agent software and parameters

```

./sapmedia/HANA2SP5Rev56/awsBackint # more template.rs
[DEFAULT]
hana_systemdb_global_ini_file =
install_awscli = n
s3_bucket_name =
s3_bucket_owner_account_id =
awscli_version_match = True
download_using_http = False
hana_opt_dir = /usr/sap/<sid>/SYS/global/hdb/opt
modify_global_ini = none
s3_bucket_region = ''
instance_partition = aws
perform_s3_validations = True
download_sign_file = False
download_sign_using_s3 = False
input_signature_filepath =
instance_region = us-west-2
sp_license_file_name = THIRD_PARTY_LICENSES.txt
s3_sse_kms_arn =
s3_bucket_folder =
agent_tar_file = aws-backint-agent.tar.gz
signature_file_name = aws-backint-agent.sig
hana_sid =
verify_signature = n

```

3.1 Backup config parameters:

[backup]

`catalog_backup_parameter_file = /usr/sap/<SID>/SYS/global/hdb/opt/hdbconfig/aws-backint-agent-config.yaml`

`data_backup_parameter_file = /usr/sap/<SID>/SYS/global/hdb/opt/hdbconfig/aws-backint-agent-config.yaml`

`log_backup_parameter_file = /usr/sap/<SID>/SYS/global/hdb/opt/hdbconfig/aws-backint-agent-config.yaml`

`catalog_backup_using_backint = true`

`log_backup_using_backint = true`

`parallel_data_backup_backint_channels = 8`

`data_backup_buffer_size = 4096`

`max_recovery_backint_channels = 1`

[communication]

`tcp_backlog = 2048`

[persistence]

`enable_auto_log_backup = yes`

`verify_signature = yes`

`input_signature_filepath = https://s3.amazonaws.com/awssap-backint-agent/binary/latest/aws-backint-agent.sig`

Trigger install.sh script which will launch AWS Installer and provide all the parameters AWS Agent Installer

```
# more aws-backint-agent-install.rsp
[DEFAULT]
hana_systemdb_global_ini_file =
install_awscli = n
s3_bucket_name =
s3_bucket_owner_account_id = 720964816748
awscli_version_match = True
download_using_http = False
hana_opt_dir = /usr/sap/<sid>/SYS/global/hdb/opt
modify_global_ini = none
s3_bucket_region = us-west-2
instance_partition = aws
perform_s3_validations = True
download_sign_file = False
download_sign_using_s3 = False
input_signature_filepath =
instance_region = us-west-2
3p_license_file_name = THIRD_PARTY_LICENSES.txt
s3_sse_kms_arn =
s3_bucket_folder =
agent_tar_file = aws-backint-agent.tar.gz
signature_file_name = aws-backint-agent.sig
hana_sid =
verify_signature = n
hana_sidadm_user =
confirm_install = n
download_using_s3 = False
installation_directory = /hana/shared/
```

Require all S3 parameters to point backup storage

```
absolute_path backint_agent.....: /
nt/aws-backint-agent
hana_systemdb_global_ini_file.....: 1
install_awscli.....: 1
awscli_exists.....: 1
s3_bucket_name.....: 1
s3_bucket_owner_account_id.....: 1
awscli_version_match.....: 1
download_using_http.....: 1
hana_opt_dir.....: /
opt
modify_global_ini.....: 1
s3_bucket_region.....: 1
backint_agent_home_directory.....: /
nt
instance_partition.....: 1
perform_s3_validations.....: 1
binary_name.....: 1
download_sign_file.....: 1
download_sign_using_s3.....: 1
input_signature_filepath.....: 1
instance_region.....: 1
3p_license_file_name.....: 1
s3_sse_kms_arn.....: 1
s3_bucket_folder.....: 1
agent_tar_file.....: 1
signature_file_name.....: 1
create_response_file.....: 1
hana_sid.....: 1
verify_signature.....: 1
hana_sidadm_user.....: 1
confirm_install.....: 1
download_using_s3.....: 1
installation_directory.....: /

parameter modify_global_ini set to none. The following parameters
modified
[communication][tcp_backlog].....: 2048
[backup][max_recovery_backint_channels].....: 1
```

Installation is completed successfully.

```
summary of installation
*****

checking and creating /hana/shared/aws-backint-agent directory...
downloading backint agent from http .....
creating aws backint agent launch script.....
creating aws backint agent configuration file.....
creating aws backint agent log file.....
creating symbolic links in hana opt directory.....


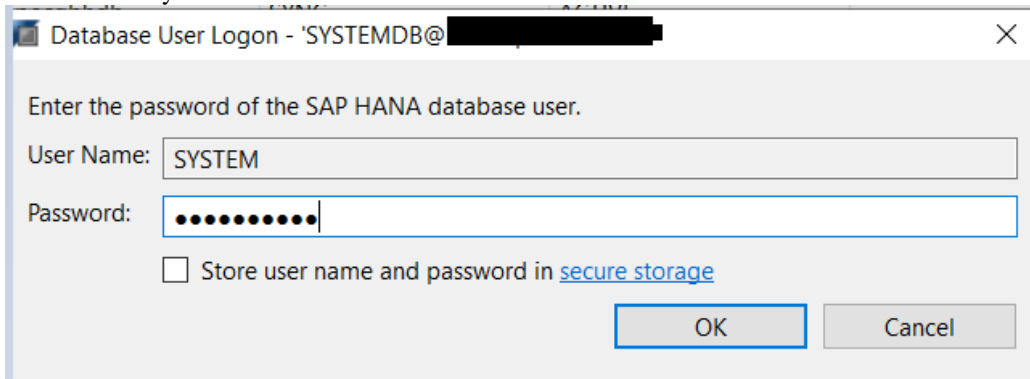
Installation completed

check the install log /sapmedia/HANA2SP5Rev56/awsBackint/aws-backint
install-20211011154627.log
```

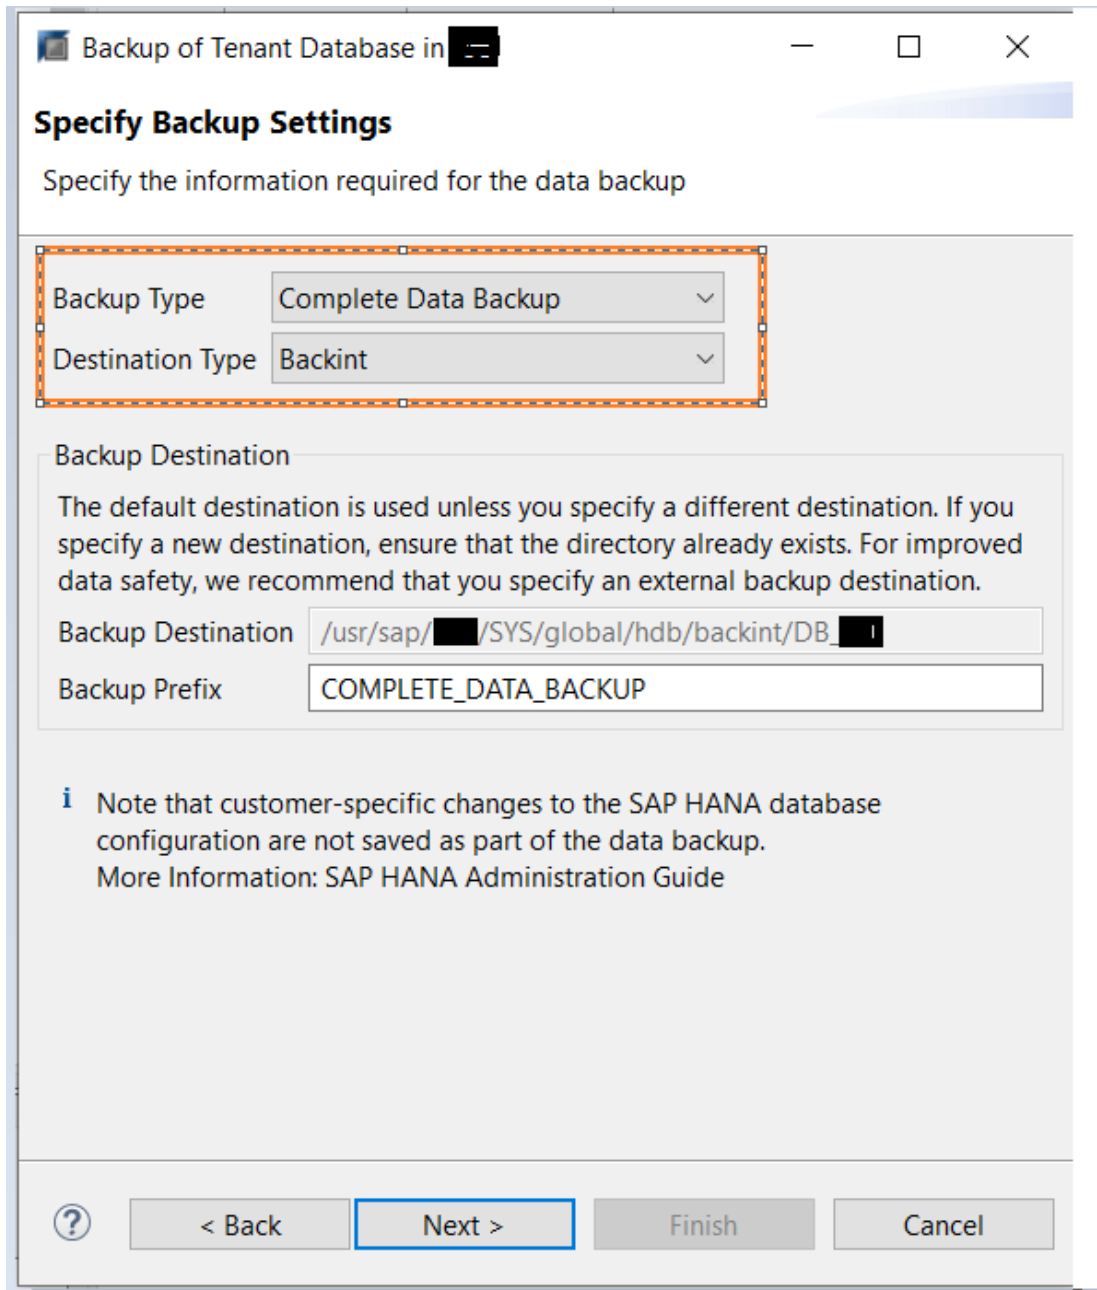
Installer log is successful. In case of errors it would describe what parameters are missing.

3.2 Backup to AWS S3 from HANA Database.

Login to Production System Database and enter credentials



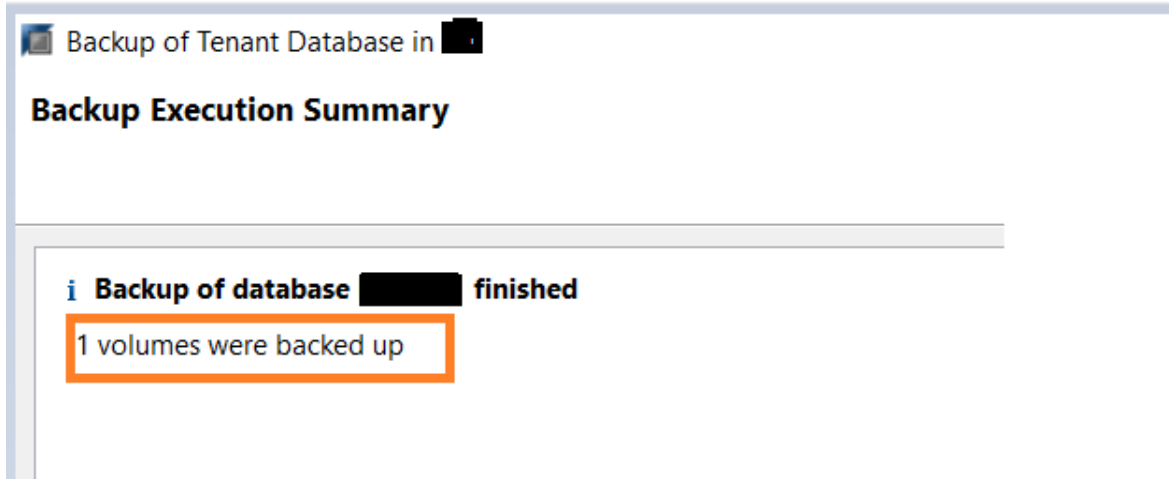
Backup Tenant Database



Complete Data Backup



Backup is in progress with percentage of completion.



Backup is completed successfully

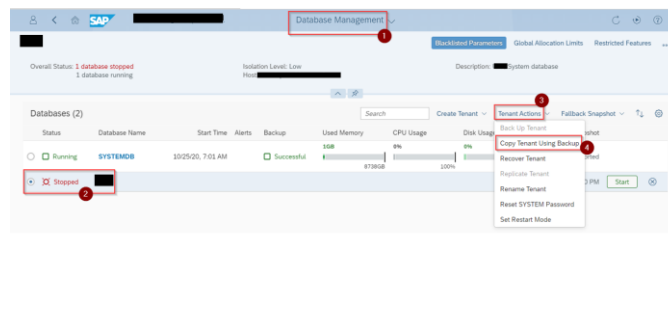
IV. Restore Database

source Database was a production which was taken in above steps. Login to target Database and replace target database yaml file with source.

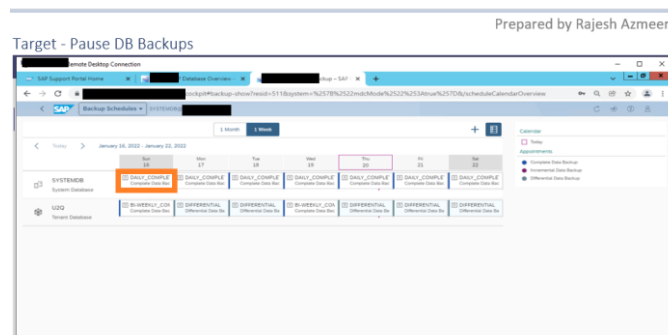
```
Mv aws-backint-agent-config.yaml aws-backint-agent-config.yaml_source
mv aws-backint-agent-config.yaml_Targ aws-backint-agent-config.yaml
```

4.1 Target – Copy Tenant using backup

Login to cockpit -> choose the SYSTEMDB source -> database management Chosen -> Copy Tenant using Backup



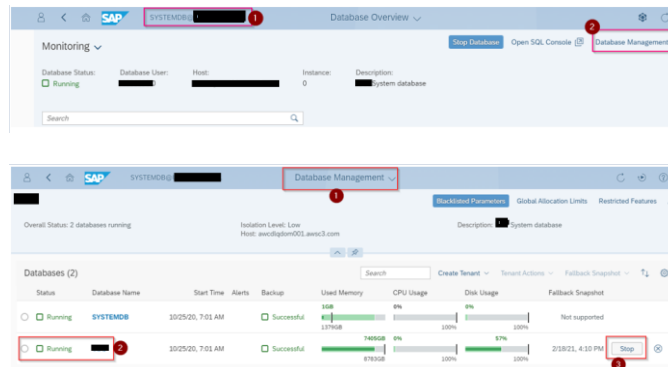
Pause/Stop Target Database backups



4.2 Target - Stop tenant DB

Using – Cockpit

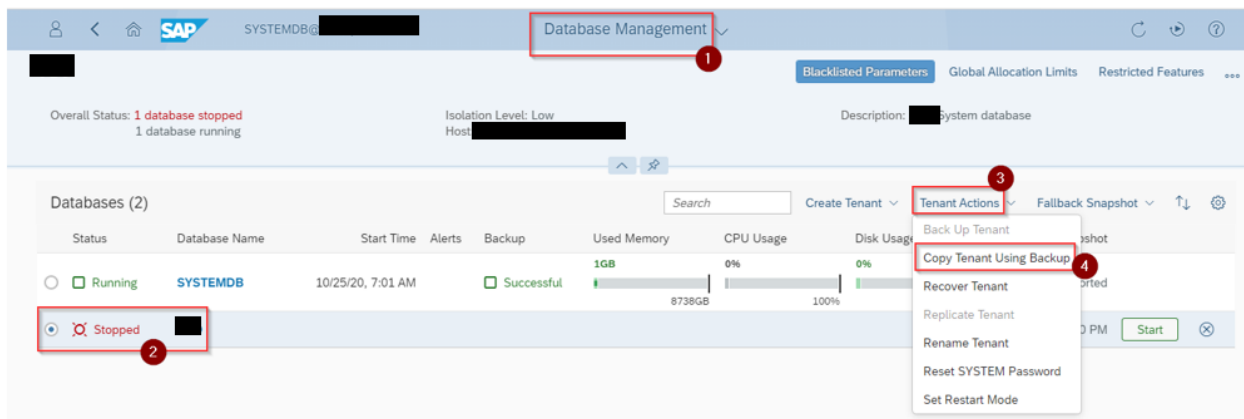
Login to cockpit -> choose the systemdb (Target) -> databasemanagement



Target System – Copy Tenant using backup

Login to cockpit -> choose the SYSTEMDB -> database management

Chosen -> Copy Tenant using Backup



Prepared by Rajesh Azmeera

1 Database Copy Type 2 Backup Catalog 3 Backup Catalog Location 4 Source System Type 5 Source Database

1. Database Copy Type

Copy from:

- Full data backup only
- Data and log backups

2. Backup Catalog

Select the data backup from a backup catalog:

- No
- Yes

3. Backup Catalog Location

- Backint
- File system

4. Source System Type

- Multi container
- Single container

Copy Database [redacted]

1 Database Copy Type 2 Backup Catalog 3 Backup Catalog Location 4 Source System Type 5 Source Database 6 Backup to be Used

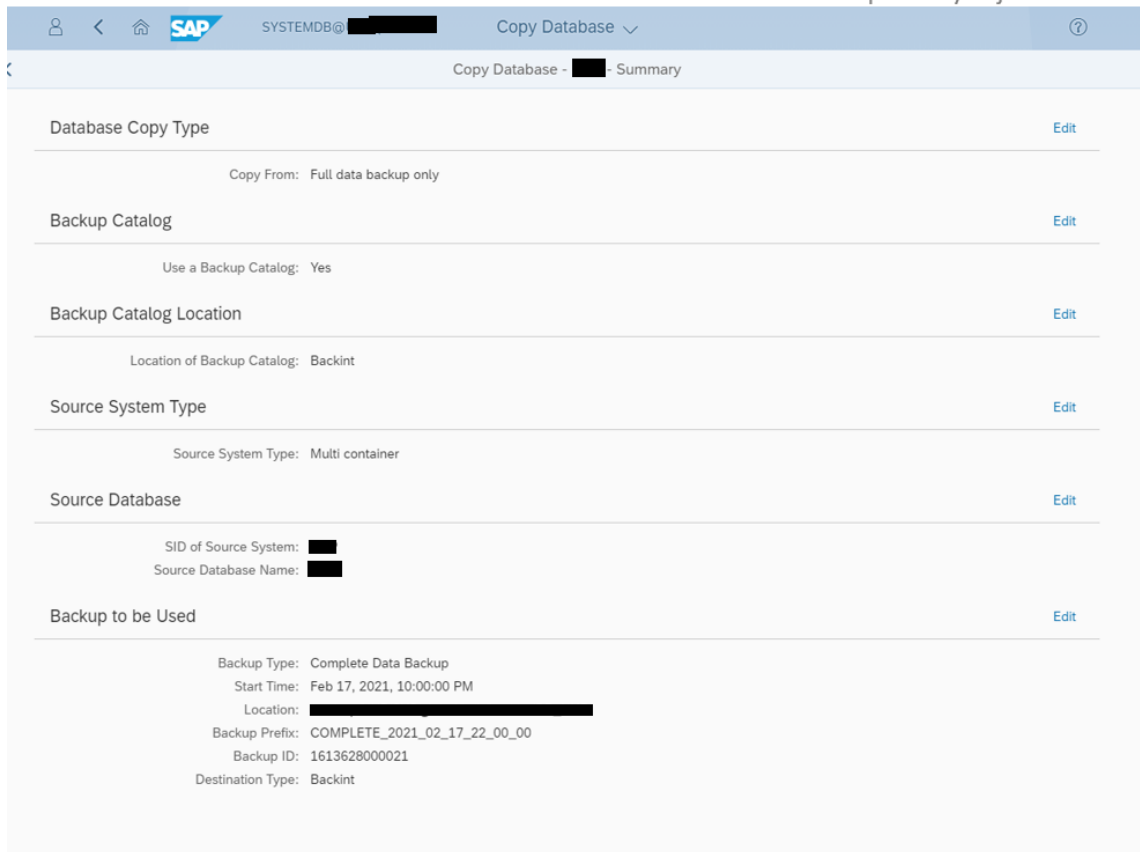
6. Backup to be Used

Filtered by: Destination Type (Backint)

Start Time	Backup Type	Status	Location	Backup Prefix
2/17/21, 10:00 PM	Complete Data Backup	<input checked="" type="checkbox"/>	/usr/sas/[redacted]SYS/global/hdb/backint/DB_UHP/	COMPLETE_2021_02_17_22_00_00
2/13/21, 10:00 PM	Complete Data Backup	<input type="checkbox"/>	[redacted]	COMPLETE_2021_02_13_22_00_00
2/10/21, 10:00 PM	Complete Data Backup	<input type="checkbox"/>	[redacted]	COMPLETE_2021_02_10_22_00_00
2/6/21, 10:00 PM	Complete Data Backup	<input type="checkbox"/>	[redacted]	COMPLETE_2021_02_06_22_00_00
2/3/21, 10:00 PM	Complete Data Backup	<input type="checkbox"/>	[redacted]	COMPLETE_2021_02_03_22_00_00
1/30/21, 10:00 PM	Complete Data Backup	<input type="checkbox"/>	[redacted]	COMPLETE_2021_01_30_22_00_00
1/27/21, 10:00 PM	Complete Data Backup	<input type="checkbox"/>	[redacted]	COMPLETE_2021_01_27_22_00_00
1/23/21, 10:00 PM	Complete Data Backup	<input type="checkbox"/>	[redacted]	COMPLETE_2021_01_23_22_00_00
1/21/21, 12:17 AM	Complete Data Backup	<input type="checkbox"/>	[redacted]	COMPLETE_2021_01_21_00_17_00
1/16/21, 10:00 PM	Complete Data Backup	<input type="checkbox"/>	[redacted]	COMPLETE_2021_01_16_22_00_00

[Review](#)

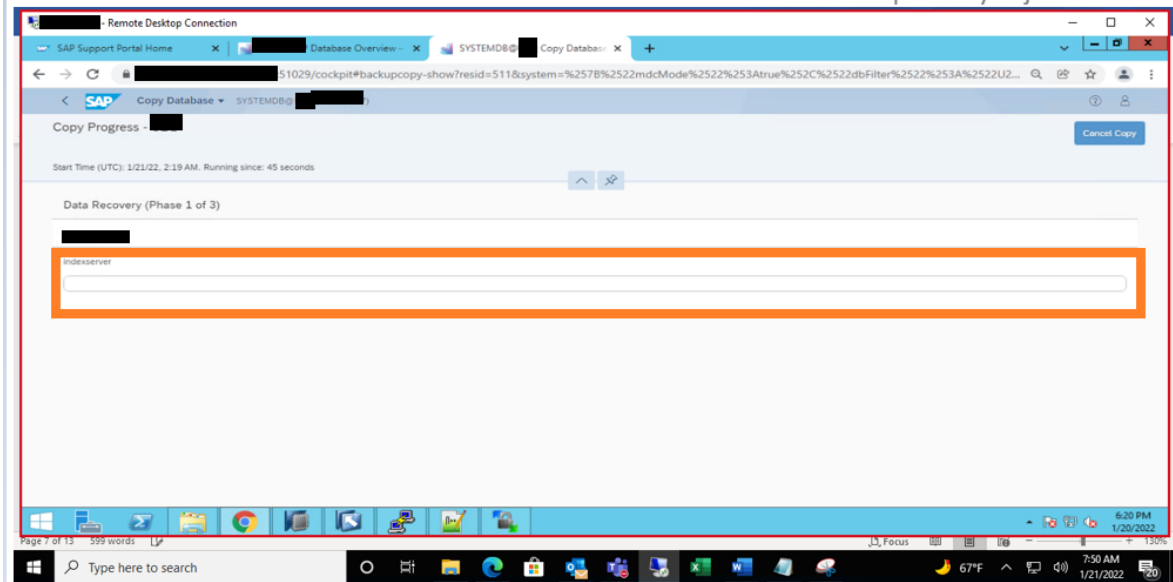
Prepared by Rajesh Azmeera



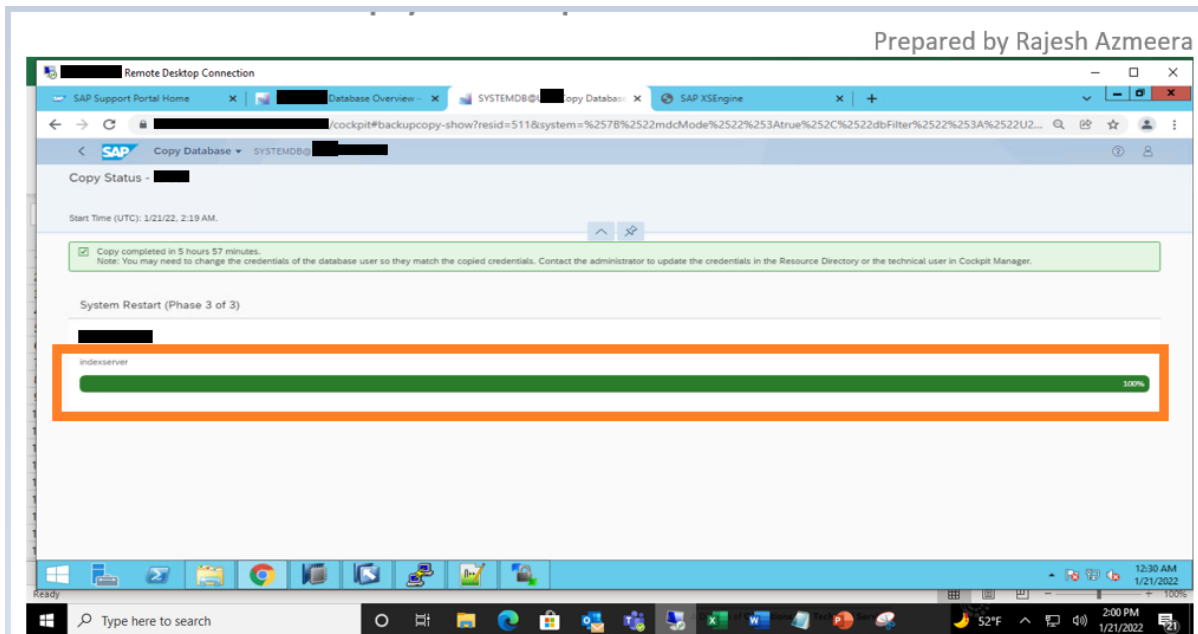
The screenshot shows the SAP 'Copy Database' summary page. The interface includes a top navigation bar with the SAP logo and user information. The main content area is titled 'Copy Database - Summary' and contains several sections, each with an 'Edit' link:

- Database Copy Type:** Copy From: Full data backup only
- Backup Catalog:** Use a Backup Catalog: Yes
- Backup Catalog Location:** Location of Backup Catalog: Backint
- Source System Type:** Source System Type: Multi container
- Source Database:** SID of Source System: [redacted], Source Database Name: [redacted]
- Backup to be Used:** Backup Type: Complete Data Backup, Start Time: Feb 17, 2021, 10:00:00 PM, Location: [redacted], Backup Prefix: COMPLETE_2021_02_17_22_00_00, Backup ID: 1613628000021, Destination Type: Backint

Prepared by Rajesh Azmeera



The screenshot shows a 'Remote Desktop Connection' window displaying the SAP 'Copy Progress' interface. The browser address bar shows a URL with parameters for the copy operation. The main content area is titled 'Copy Progress' and includes a 'Cancel Copy' button. Below this, it shows 'Start Time (UTC): 1/21/22, 2:19 AM. Running since: 45 seconds'. The 'Data Recovery (Phase 1 of 3)' section contains a table with one row labeled 'indexserver' and a corresponding empty input field, which is highlighted with an orange border. The Windows taskbar at the bottom shows the time as 7:50 AM on 1/21/2022.



Restore is completed and Database will come up and available for business users.

V. Conclusion:

With unlimited storage S3, HANA Database backups should be taken periodically to restore back when ever needed. SAP with AWS using backint tool using multiple channels runs backup faster and can be restore in no time. Understanding and implementing SAP HANA backup is crucial for every business using SAP HANA. With careful planning, regular testing, and vigilant monitoring, you can safeguard your business's valuable data against unexpected losses. Regular DR test using these backups can be tested to make sure that systems will be available and data won't be lost.

AWS Backup for SAP HANA makes it easy to perform backup and restore operations for your SAP HANA databases on AWS. AWS customers can now centrally manage and automate data protection activities, including backup, restore and system copy. Customers will benefit from a native AWS experience which can be scaled to simplify the administration across multiple AWS resources and accounts.

VI. Disclaimer:

The information presented in this paper is for general informational purposes only and does not constitute legal, financial, or professional advice. The authors and publishers do not assume any liability or responsibility for any errors or omissions in the content or for any actions taken based on the information provided. The use of any product, service, or solution mentioned in this paper should be based on the reader's due diligence and discretion.

Declarations:

Ethical Approval: Not Applicable

Funding: Not Applicable

Availability of data and materials: Not Applicable

References:

- [1]. USGS science for a changing world, Backup-Secure. [Online]. Available at: <https://www.usgs.gov/data-management/backup-secure#:~:text=Making%20backups%20of%20collected%20data,money%20if%20these%20failures%20occur>.
- [2]. Anjani Singh, Migrating SAP HANA-based systems using AWS Migration Hub Orchestrator. AWS Blogs. [Online]. Available at: <https://aws.amazon.com/blogs/awsfor/sap-on-hana-migration-using-aws-migration-hub-orchestrator/>
- [3]. Introduction to the Universe of SAP S/4HANA Cloud, SAP Help. [Online]. Available at: https://help.sap.com/docs/SAP_S4HANA_CLOUD?version=2308.501
- [4]. Rezza Prayogi, SAP Book, SAP BASIS from Zero to Hero
- [5]. Empower, On Premise Cloud. Cloud blog [Online]. Available at: <https://www.empowersuite.com/en/blog/on-premise-vs-cloud>
- [6]. Raja Gupta, Sandip Jha. DevOps with SAP. SAP PRESS. [Online]. Available at: https://www.sap-press.com/devops-with-sap_5694/
- [7]. Icloudhb. Migrating SAP to the Cloud. Website. [Online]. Available at: <https://www.1cloudhub.com/migrating-sap-cloud/>

- [8]. SAP HANA Database Backup and Recovery, Help.sap. [Online]. Available at: https://help.sap.com/docs/SAP_HANA_PLATFORM/6b94445c94ae495c83a19646e7c3fd56/c4663eabbb571014923a90c70cec566c.html?version=2.0.03&locale=en-us
- [9]. AWS Backup. AWS Docs. [Online]. Available at: <https://docs.aws.amazon.com/sap/latest/sap-hana/aws-backint-agent-backup.html>
- [10]. Install and configure AWS Backint agents. [Online]. Available at: <https://docs.aws.amazon.com/sap/latest/sap-hana/aws-backint-agent-s3-installing-configuring.html>
- [11]. Denys van Kempen. SAP HANA 2.0 Introduction. SAP PRESS Book.
- [12]. Rich Dennis, HANA Advance concepts.