IBM Technical Brief

SAP® HANA® DB Migration from x86 to POWER® via Backup/Recover

Mark Gordon

IBM POWER and SAP Solutions

Version: 1.0
Date: December 4, 2017

1. Disclaimers

This paper demonstrates a way to minimize downtime and migrate a HANA® database from HANA 1.0 SPS12 to HANA 2.0 SPS02 using HANA backup/restore. It is not a best practices guide. This process has been tested on our lab systems in the configuration described.

2. Trademarks

SUSE ® is a registered trademark of SUSE LLC.
SAP ®, SAP NetWeaver ®, and SAP HANA ® are registered trademarks of SAP SE.
Intel® is a registered trademark of Intel Corporation.
IBM® and POWER® are registered trademarks of IBM Corporation.

3. Version Changes

Version 1.0: September 6, 2017 – initial version
Version 1.1: December 4, 2017 – add scenario for backup from Single-DB source

4. Acknowledgements

Thank you to Lou Lamprinakos and Damir Rubic for reading the draft and offering suggestions for improvements.

5. Feedback

Please send comments or suggestions for changes to gordonmr@us.ibm.com.

6. Introduction

This paper demonstrates a process to migrate a scale-out HANA V1.0 DB on Intel to a single node HANA 2.0 DB on IBM POWER LE (Little Endian) using HANA backup and recover. When backup/recover can be used, migration from x86 to POWER LE can be faster and simpler than classical migration or SUM DMO.

HANA 2.0 on POWER runs on the Little-Endian SLES12 Operating System. HANA on Intel is also LE. *As described in SAP note 1642148 items 44 & 45, a DB backup from x86 HANA 1.0 SPS10 (or later) can be restored to HANA on POWER 2.0.*

In addition, per SAP "HANA Recovery" online doc at https://help.sap.com/doc/6b94445c94ae495c83a19646e7c3fd56/2.0.00/en-US/c3c66b63bb571014b3e5ad8618cda1ad.html the source HANA system may be Single-DB or Multi-Tenant.

The source and target DBs must have similar topology - both must also have the same number of indexservers. This does not mean that both source and target must be scale-out systems. In this example, we will consolidate a three node HANA cluster onto a single LPAR with three indexservers. This fulfills the requirement that the topology be the same for source and target.

7. DB Migration Steps

In this example, we will migrate a Netweaver 7.5 DB server from x86 to POWER. The application server is kernel 7.49 on SUSE SLES11 SP4 x86. The source DB server is a three-node scale-out system on HANA 1.0 SPS 12 on SUSE SLES11 SP4 x86.

The target DB will be a single LPAR POWER system on SLES12 SP2 running HANA 2.0 SPS02. The application server will remain on x86 SLES11 SP4.

- Application server – atssg140
- Source DB servers – atssg141, atssg142, atssg143
- Target (POWER) DB server – atssg86

In this paper, we will:
- Verify that the SAP system is supported on HANA 2.0 - section 7.1
- Backup Scenario 1 – Source is Single Database System - section 7.2
- Backup Scenario 2 – Source is Multi-Tenant - section 7.3
- Install target DB and prepare for restore - section 7.4.
- Recover HN1 tenant onto target DB - section 7.5
- Update SAP application server and restart SAP instance - section 7.6
- (Optional) Reduce number of index servers – section 7.7

7.1. Verify that the SAP system is supported on HANA 2.0

Review the SAP PAM (https://support.sap.com/en/release-upgrade-maintenance/product-availability-matrix.html) to confirm that the SAP version (in our case Netweaver 7.5) is supported on HANA 2.0. The PAM also shows the kernel versions that are compatible with both HANA 1.0 and 2.0

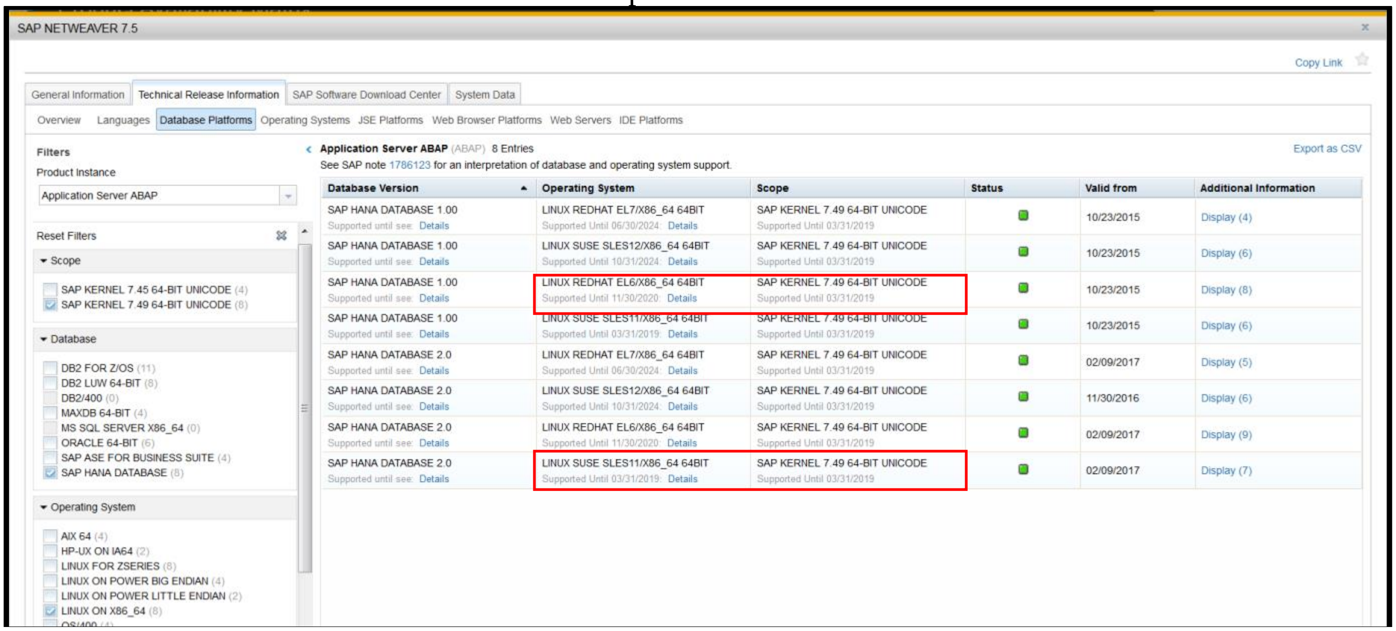

Figure 1: SAP PAM for Netweaver 7.5

Figure 1 shows the SAP kernel 7.49 on SLES11 SP4 is supported with HANA 1.0 and HANA 2.0 for Netweaver 7.5.

Confirm in SAP note 2218464 that HANA 2.0 on POWER supports Netweaver 7.5. HANA on POWER support has some limitations that are not contained in the PAM.

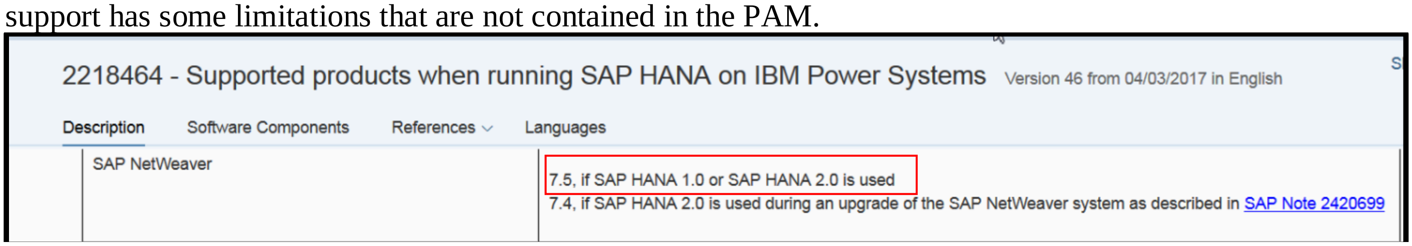

Figure 2: SAP note 2218464

Read SAP note 2399995 to review hardware requirements for HANA 2.0 – POWER 8 is required.

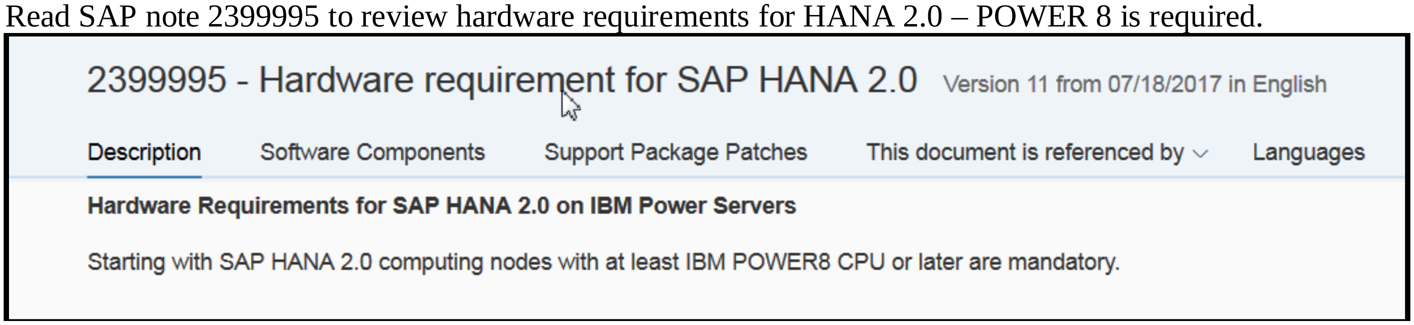

Figure 3: SAP note 2399995

Review SAP note 2235581 to check Linux requirements for HANA 2.0 on POWER. We will install SLES for SAP Applications 12 SP2.

2235581 - SAP HANA: Supported Operating Systems Version 20 from 08/07/2017 in English

Description | Software Components | References | Attachments | Attributes | Languages

Operating System for SAP HANA 2.0 on IBM Power Servers

For an SAP HANA system on IBM Power servers the following operating system is available for SAP HANA 2.0:

- SUSE Linux Enterprise Server (SLES) for SAP Applications 12 SP1
  - SAP Note 2205917 - SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12
  - SAP Note 1984787 - SUSE LINUX Enterprise Server 12: Installation notes
  - SAP Note 2055470 - HANA on POWER Planning and Installation Specifics - Central Note
- SUSE Linux Enterprise Server (SLES) for SAP Applications 12 SP2
  - SAP Note - 2205917 SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12
  - SAP Note - 1984787 SUSE LINUX Enterprise Server 12: Installation notes
  - SAP Note 2055470 - HANA on POWER Planning and Installation Specifics - Central Note

Figure 4: SAP note 2235581

7.2. Backup Scenario 1 – Source is Single Database System

The source DB (HN1) is a three-node scale-out system, without a standby node. There are three indexservers, one xsengine and one nameserver.

HN1 (SYSTEM) atssg141.svl.ibm.com 00

Overview | Landscape | Alerts | Performance | Volumes | Configuration | System Information | Diagnosis Files | Trace Configura

Show: Service Host: <All>

Service/Volume	Service	Total Volume Size (MB)	Data Volume Size (MB)	Log Volume Size (MB)
atssg141:30001	nameserver	455	326	129
atssg141:30003	indexserver	6,679	4,629	2,049
atssg141:30007	xsengine	341	324	17
atssg142:30003	indexserver	9,101	7,052	2,049
atssg143:30003	indexserver	6,145	4,096	2,049

Figure 5: Source DB topology

Stop the SAP application servers before backup of the HANA DB.

Backup the entire HANA 1.0 system using the userid SYSTEM. This will back up indexservers, nameserver, and xsengine persistence. The restore will use only the indexserver and xs components, when restoring into the HN1 client of the HANA 2.0 system.

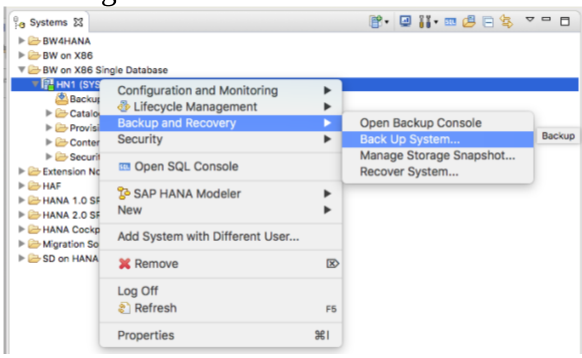

Figure 6: Single DB Backup

The Backup prefix 'SINGLE_DB_SOURCE" is used to name the files.

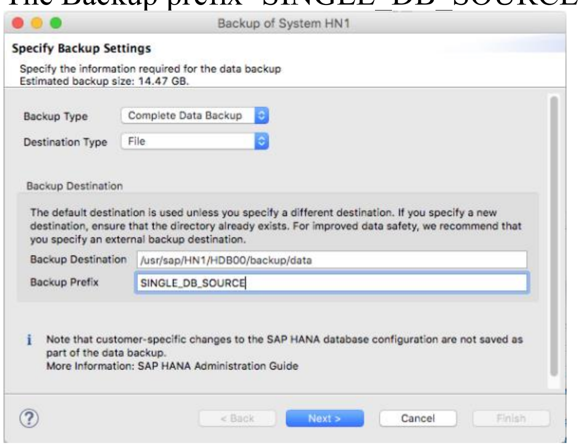

Figure 7: Single DB Backup Settings

Five volumes are backed up – three indexservers, nameserver, and xs. (See Figure 5: Source DB topology)

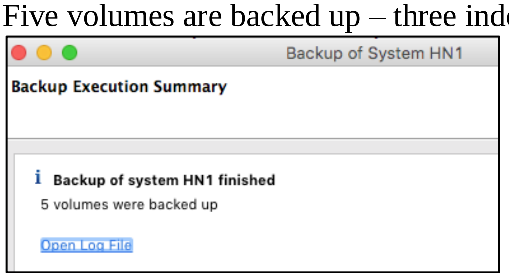

Figure 8: Single DB Backup Summary

After the backup is complete, note the files (which will be copied to the target system).

```
atssg141:/hana/shared/HN1/HDB00/backup/data> ls -l
total 21137596
-rw-r----- 1 hn1adm sapsys     163840 Dec  4 17:18 SINGLE_DB_SOURCE_databackup_0_1
-rw-r----- 1 hn1adm sapsys   83894272 Dec  4 17:18 SINGLE_DB_SOURCE_databackup_1_1
-rw-r----- 1 hn1adm sapsys   83894272 Dec  4 17:18 SINGLE_DB_SOURCE_databackup_2_1
-rw-r----- 1 hn1adm sapsys 4747960320 Dec  4 17:20 SINGLE_DB_SOURCE_databackup_3_1
-rw-r----- 1 hn1adm sapsys 3741327360 Dec  4 17:20 SINGLE_DB_SOURCE_databackup_4_1
-rw-r----- 1 hn1adm sapsys 6996107264 Dec  4 17:21 SINGLE_DB_SOURCE_databackup_5_1
atssg141:/hana/shared/HN1/HDB00/backup/data>
```

Figure 9: Single DB Backup Files

Since the HANA backup does not copy customization of the ini files, review the contents of the ini files to determine which changes need to be applied on the target system.

```
atssg141:/hana/shared/HN1/HDB00/backup/data> find /hana -name '*ini'
/hana/global.ini
/hana/shared/HN1/global/hdb/custom/config/global.ini
/hana/shared/HN1/global/hdb/custom/config/nameserver.ini
/hana/shared/HN1/global/hdb/custom/config/webdispatcher.ini
/hana/shared/HN1/global/hdb/custom/config/statisticsserver.ini
/hana/shared/HN1/global/hdb/custom/config/indexserver.ini
```

Figure 10: Single DB ini files

7.3. Backup Scenario 2 – Source is Multi-Tenant Database

We will backup the HN1 tenant, which will subsequently be recovered onto HANA 2.0.

The HN1 tenant has persistence for indexservers and xs. The nameserver is part of the SystemDB. Note here that the HN1 SYSTEM userid can display only the HN1 tenant resources.

HN1@HN1 (SYSTEM) atssg141.svl.ibm.com 00

Overview | Landscape | Alerts | Performance | Volumes | Configuration | System Information | Diagnosis Files | Trace Configuration

Show: Service   Host: <All>

Service/Volume	Service	Total Volume Size (MB)	Data Volume Size (MB)	Log Volume Size (MB)	Path
atssg141:30003	indexserver	7,358	5,309	2,049	
atssg141:30007	xsengine	341	324	17	
atssg142:30003	indexserver	10,977	8,928	2,049	
atssg143:30003	indexserver	6,513	4,464	2,049	

Figure 11: HN1 tenant DB volumes

Using the SYSTEM ID in SYSTEMDB, backup the tenant.

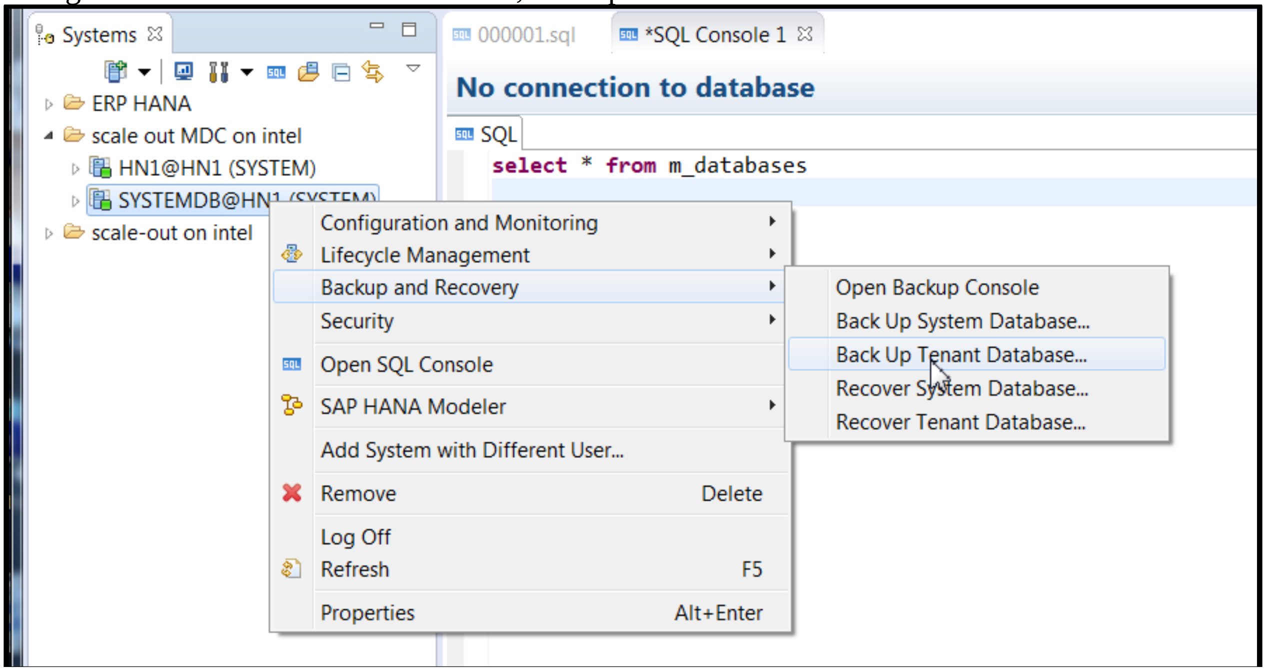

Figure 12: Backup tenant DB

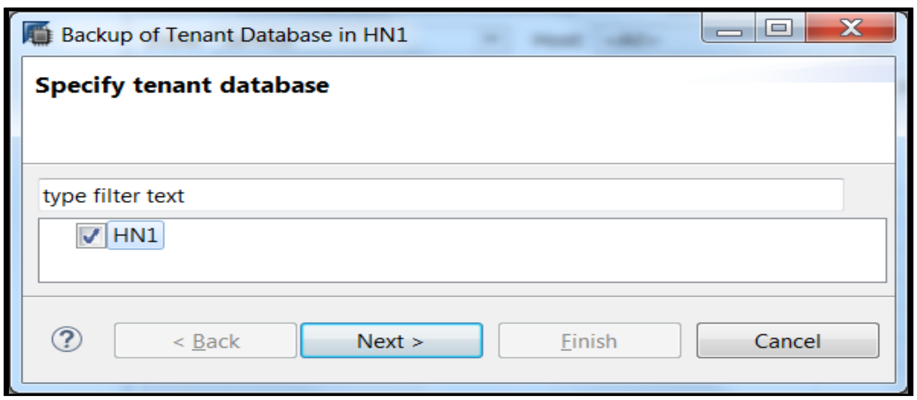

Figure 13: Backup select tenant

In Figure 14, note the ‘Backup Prefix’, which will be used for the restore.

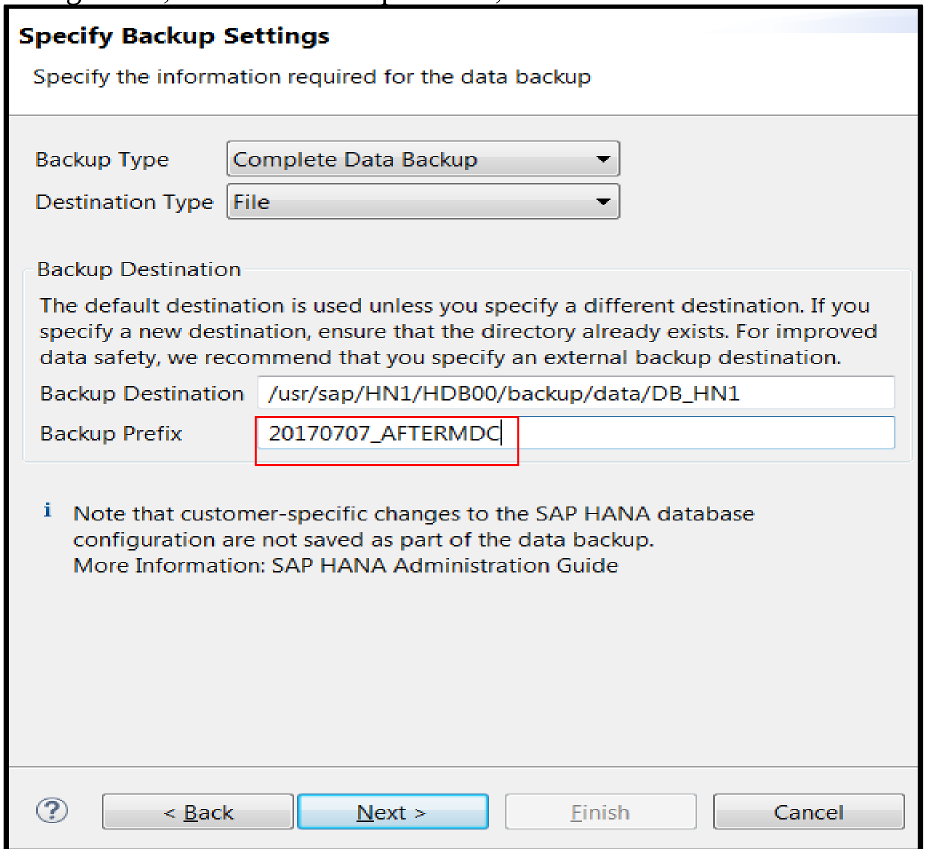

Figure 14: Backup destination

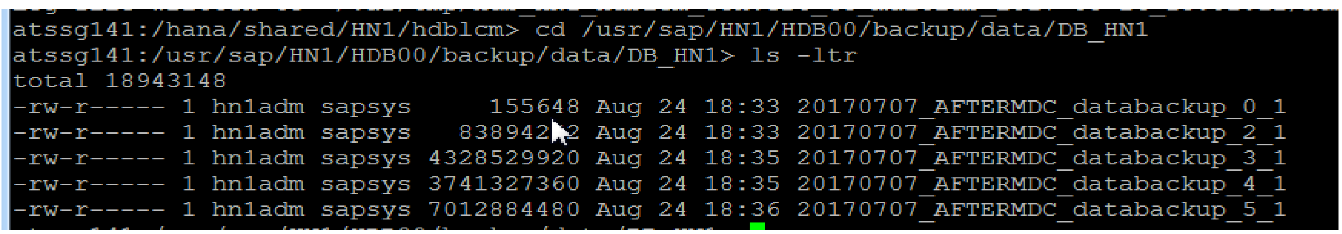

```
atssg141:/hana/shared/HN1/hdblcm> cd /usr/sap/HN1/HDB00/backup/data/DB_HN1
atssg141:/usr/sap/HN1/HDB00/backup/data/DB_HN1> ls -ltr
total 18943148
-rw-r----- 1 hn1adm sapsys     155648 Aug 24 18:33 20170707_AFTERMDC_databackup_0_1
-rw-r----- 1 hn1adm sapsys   83894272 Aug 24 18:33 20170707_AFTERMDC_databackup_2_1
-rw-r----- 1 hn1adm sapsys 4328529920 Aug 24 18:35 20170707_AFTERMDC_databackup_3_1
-rw-r----- 1 hn1adm sapsys 3741327360 Aug 24 18:35 20170707_AFTERMDC_databackup_4_1
-rw-r----- 1 hn1adm sapsys 7012884480 Aug 24 18:36 20170707_AFTERMDC_databackup_5_1
```

Figure 15: Backup files

Review the configuration parameters for the system and tenant, to determine which need to be re-applied on the target DB.

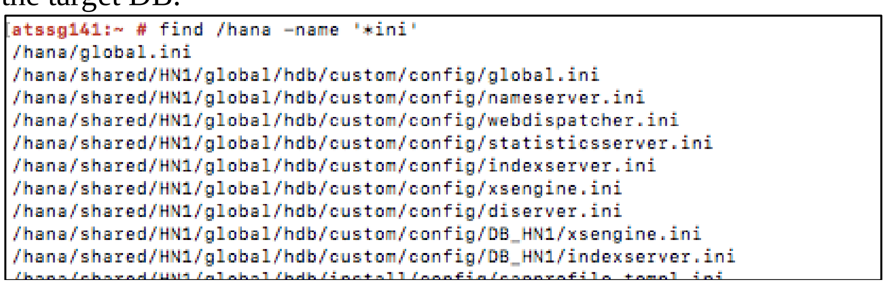

```
atssg141:~ # find /hana -name '*ini'
/hana/global.ini
/hana/shared/HN1/global/hdb/custom/config/global.ini
/hana/shared/HN1/global/hdb/custom/config/nameserver.ini
/hana/shared/HN1/global/hdb/custom/config/webdispatcher.ini
/hana/shared/HN1/global/hdb/custom/config/statisticsserver.ini
/hana/shared/HN1/global/hdb/custom/config/indexserver.ini
/hana/shared/HN1/global/hdb/custom/config/xsengine.ini
/hana/shared/HN1/global/hdb/custom/config/diserver.ini
/hana/shared/HN1/global/hdb/custom/config/DB_HN1/xsengine.ini
/hana/shared/HN1/global/hdb/custom/config/DB_HN1/indexserver.ini
```

Figure 16: Multi-tenant source ini files

7.4. Install target DB and prepare for restore

Install the HANA 2.0 SPS2 HN1 database on the POWER system.

```
SAP HANA Database System Installation
   Installation Parameters
      Remote Execution: ssh
      Database Isolation: low
      Installation Path: /hana/shared
      Local Host Name: atssg86
      SAP HANA System ID: HN1
      Instance Number: 00
      Local Host Worker Group: default
      System Usage: custom
      Location of Data Volumes: /hana/data/HN1
      Location of Log Volumes: /hana/log/HN1
      Certificate Host Names: atssg86 -> atssg86
      System Administrator Home Directory: /usr/sap/HN1/home
      System Administrator Login Shell: /bin/sh
      System Administrator User ID: 1001
      ID of User Group (sapsys): 79
   Software Components
      SAP HANA Database
         Install version 2.00.020.00.1500920972
         Location: /repo/Power_HANA/v20_SPS02_install/20170804/51052326/DATA_UNITS/HDB_SERVER_LINUX_PPC64LE/server
      SAP HANA AFL (incl.PAL,BFL,OFL,HIE)
         Do not install
      SAP HANA EML AFL
         Do not install
      SAP HANA EPM-MDS
         Do not install
      SAP HANA Database Client
         Do not install
      SAP HANA Smart Data Access
         Do not install
      SAP HANA XS Advanced Runtime
         Do not install

Do you want to continue? (y/n):
```

Figure 17: hdblcm

The m_databases view on the target LPAR shows we have HN1 tenant and the system DB installed.

SYSTEMDB@HN1 (SYSTEM) atssg86.svl.ibm.com 00

SQL | Result

```
select * from m_databases
```

	DATABASE_NAME	DESCRIPTION	ACTIVE_STATUS	ACTIVE_STATUS_DETAILS	OS_USER	OS_GROUP	RESTART_MODE
1	SYSTEMDB	SystemDB-HN1-00	YES				DEFAULT
2	HN1	HN1-00	YES				DEFAULT

Figure 18: m_databases on target DB

Check configuration of HN1 tenant. Initially, there is persistence for only one indexserver.

HN1@HN1 (SYSTEM) atssg86.svl.ibm.com 00 La

Overview | Landscape | Alerts | Performance | Volumes | Configuration | System Information | Diagnosis Files | Trace Configuration

Show: Service Host: <All>

Service/Volume	Service	Total Volume Size (MB)	Data Volume Size (MB)	Log Volume Size (MB)	Path
atssg86:30003	indexserver	4,045	1,996	2,049	
atssg86:30007	xsengine	337	320	17	

Figure 19: Volumes on target DB after initial install

The initial indexserver is on SQL port 30015. We note this for later, when we remove added indexservers.

HN1@HN1 (SYSTEM) atssg86.svl.ibm.com 00

SQL | Result

```sql
select * from m_services
```

	HOST	PORT	SERVICE_NAME	PROCESS_ID	DETAIL	ACTIVE_STATUS	SQL_PORT	COORDINATOR_TYPE
1	atssg86	30,000	daemon	18,226		YES	0	NONE
2	atssg86	30,001	nameserver	18,242	master	YES	0	MASTER
3	atssg86	30,002	preprocessor	18,548		YES	0	NONE
4	atssg86	30,006	webdispatcher	19,770		YES	0	NONE
5	atssg86	30,010	compileserver	18,546		YES	0	NONE
6	atssg86	30,003	indexserver	18,591	master	YES	30,015	MASTER
7	atssg86	30,007	xsengine	18,593		YES	0	NONE

Next, we need to create two more indexservers on the target DB server, so that the topology will match the source – one XS volume, and three indexserver volumes.

SYSTEMDB@HN1 (SYSTEM) atssg86.svl.ibm.com 00

SQL

```sql
alter database HN1 add 'indexserver' ;
alter database HN1 add 'indexserver' ;
```

Figure 20: add indexservers

Now, there are three indexservers on HN1 on the POWER system.

HN1@HN1 (SYSTEM) atssg86.svl.ibm.com 00

SQL | Result

```sql
select * from m_services
```

	HOST	PORT	SERVICE_NAME	PROCESS_ID	DETAIL	ACTIVE_STATUS	SQL_PORT	COORDINATOR_TYPE
1	atssg86	30,000	daemon	18,226		YES	0	NONE
2	atssg86	30,001	nameserver	18,242	master	YES	0	MASTER
3	atssg86	30,002	preprocessor	18,548		YES	0	NONE
4	atssg86	30,006	webdispatcher	19,770		YES	0	NONE
5	atssg86	30,010	compileserver	18,546		YES	0	NONE
6	atssg86	30,003	indexserver	18,591	master	YES	30,015	MASTER
7	atssg86	30,007	xsengine	18,593		YES	0	NONE
8	atssg86	30,040	indexserver	33,961		YES	30,041	SLAVE
9	atssg86	30,043	indexserver	34,418		YES	30,044	SLAVE

Figure 21: m_services after adding indexservers

And likewise, there are now volumes on the target DB for three indexservers, and one XS server. The target HN1 tenant topology now matches the source HN1.

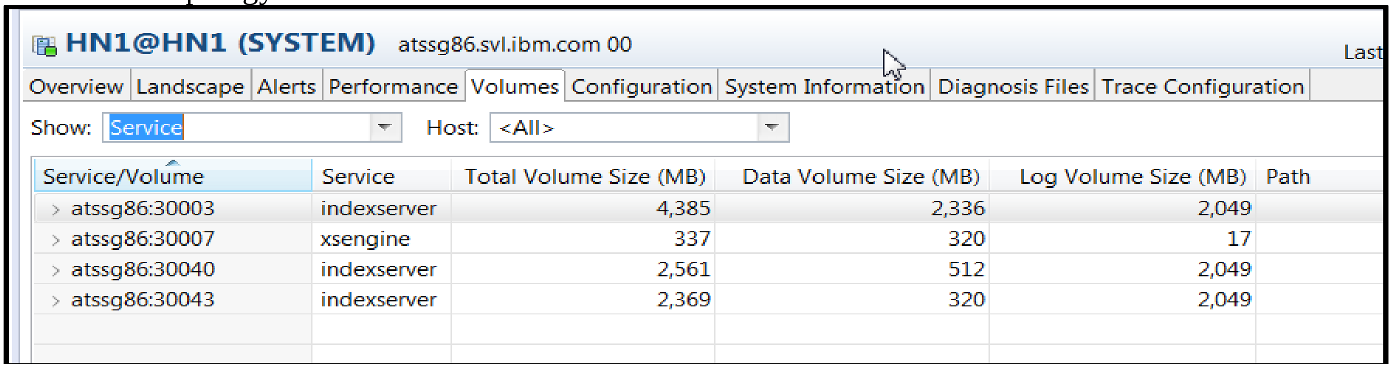

Figure 22: Target DB volumes after add indexservers

The target DB is now ready to restore the backup made x86. The restore will move the DB from x86 to POWER, and consolidate three x86 nodes to one LPAR.

7.5. Recover HN1 tenant onto target DB

The recover process is the same, whether the source system is Single DB or Multi-Tenant.

First, backup the HN1 tenant on the target DB. It will be replaced by the recovery operation.

Second, copy the HN1 backup files from the source DB to the target.

Figure 23: Copy files to target DB

Now, using the SYSTEM userid in the SYSTEMDB, recover the tenant DB.

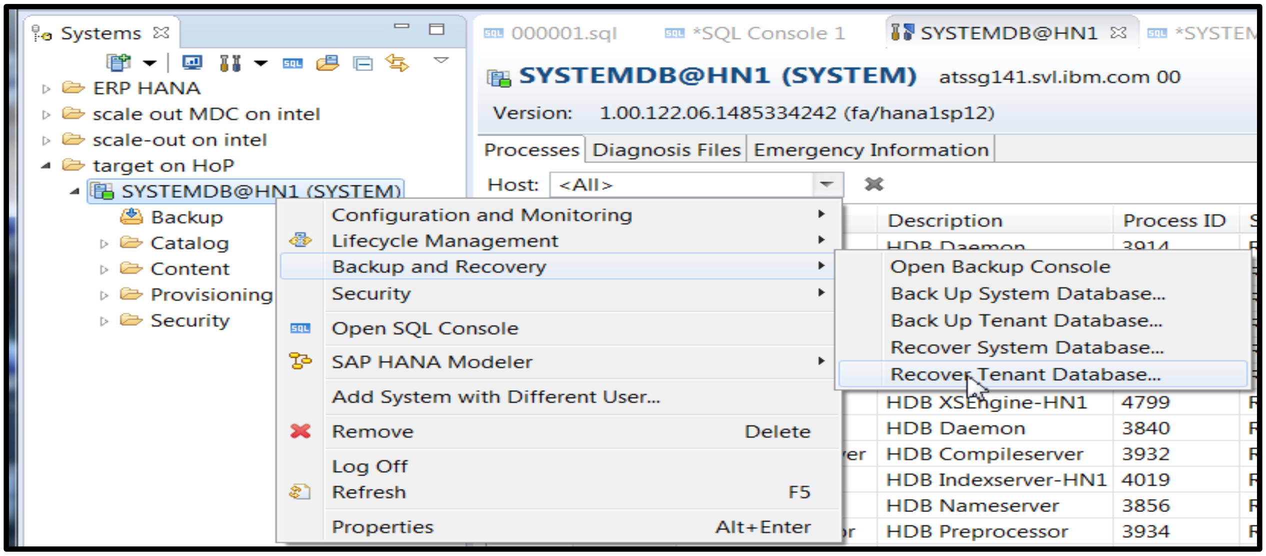

Figure 24: Recover DB

Choose the tenant to recover.

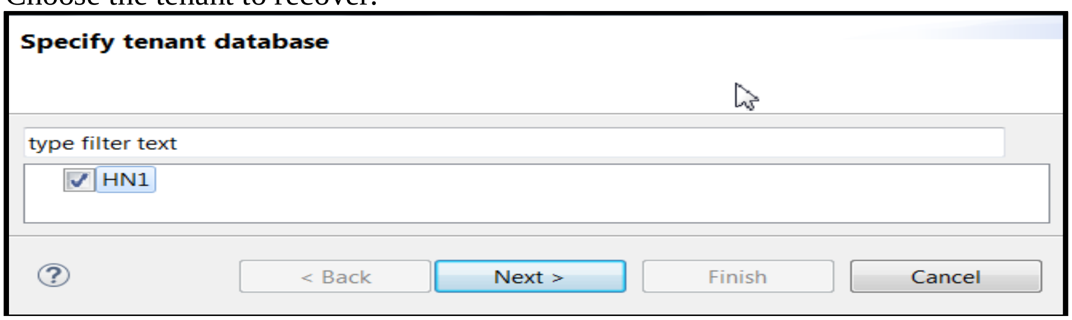

Figure 25: Recover DB tenant

Since we copied the files to the new DB server, we cannot use the catalog for recover. We will specify the location of the files, and which backup set to recover.

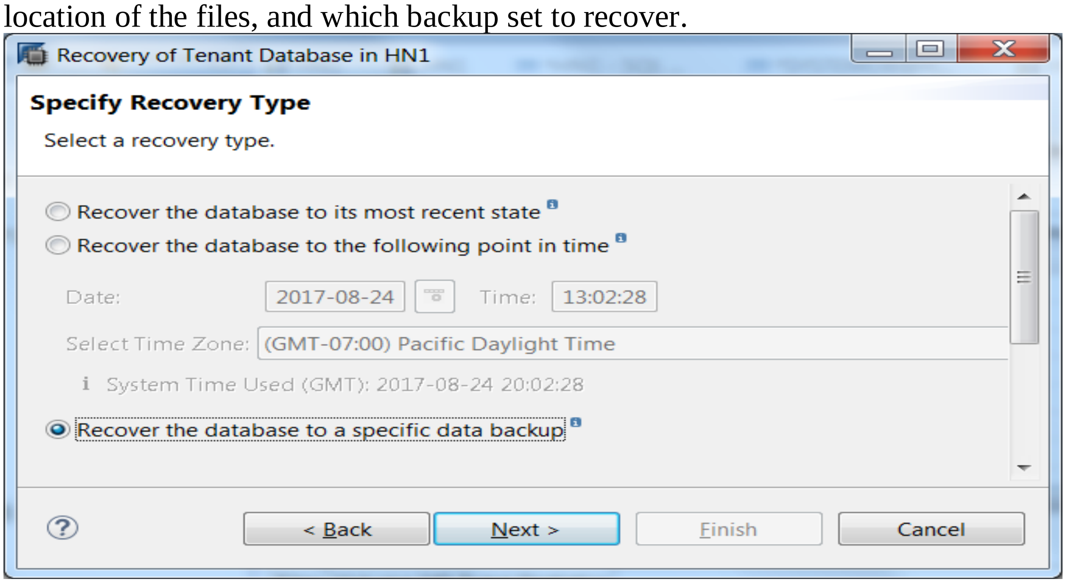

Figure 26: Recover DB type

Select “Specify backup without catalog”.

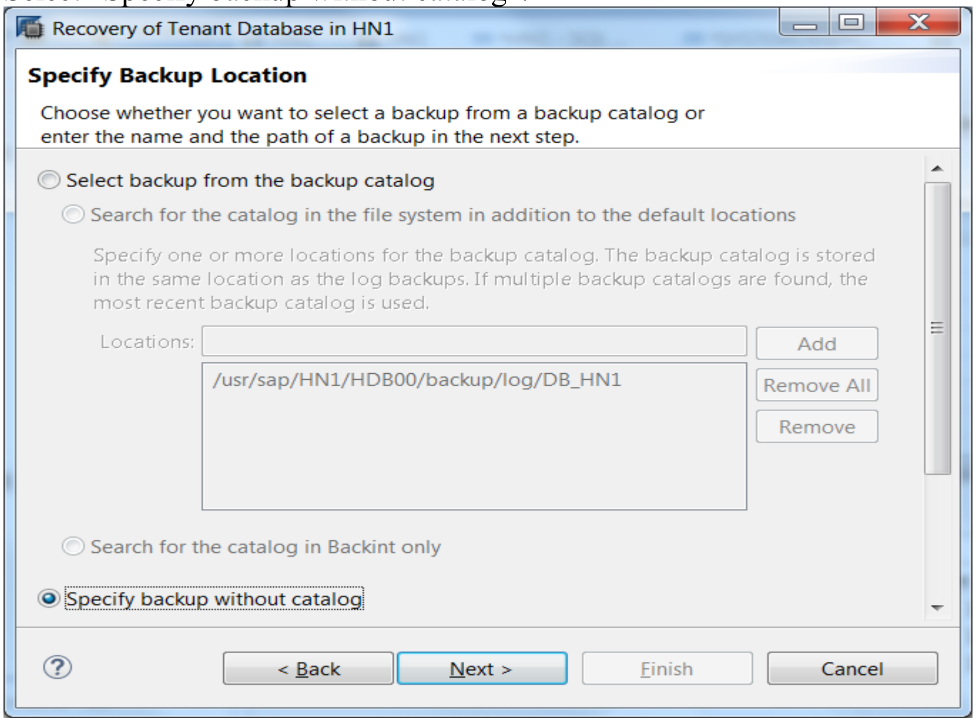

Figure 27: Recover DB backup location

Specify location where files were copied, and use the same prefix that was used above in Figure 14 when the backup was created.

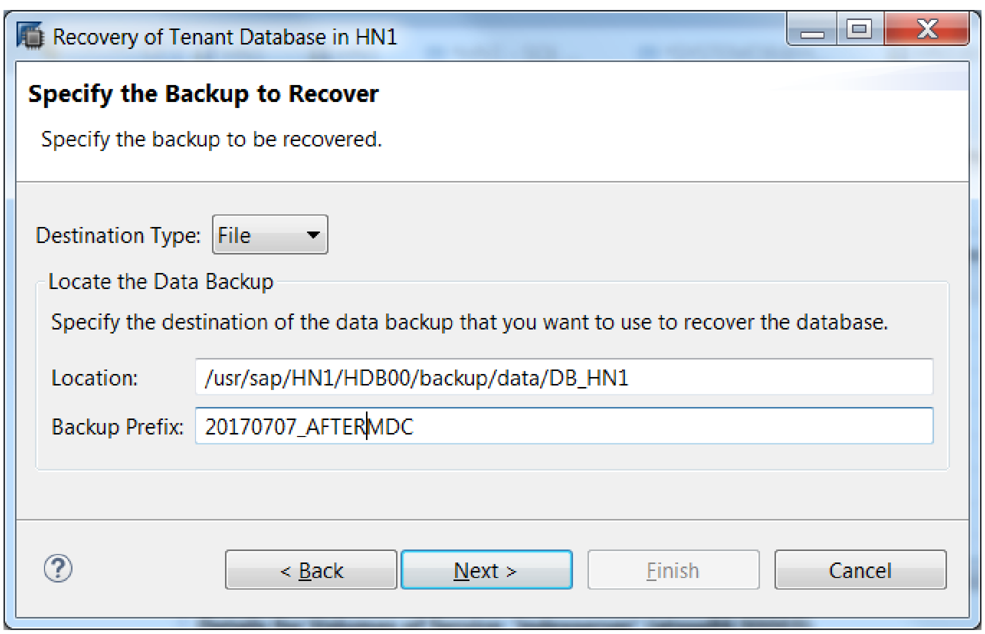

Figure 28: Recover DB choose files

Keep the defaults on this screen.

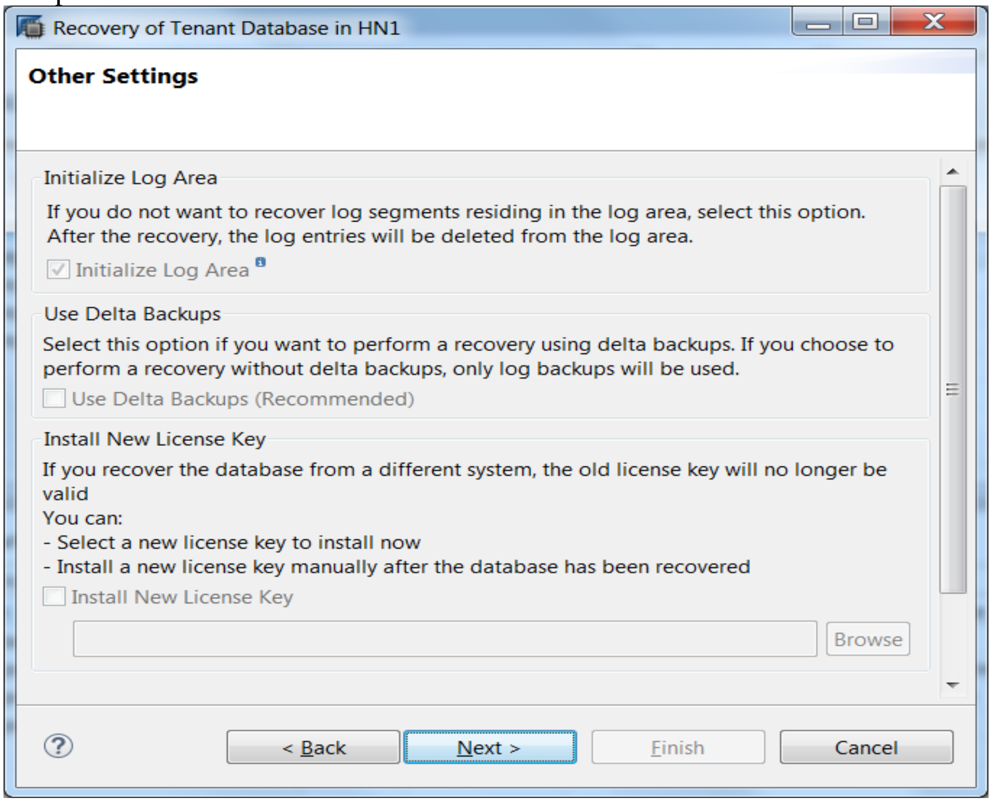

Figure 29: Recover DB other settings

HANA Studio displays what it is about to do.

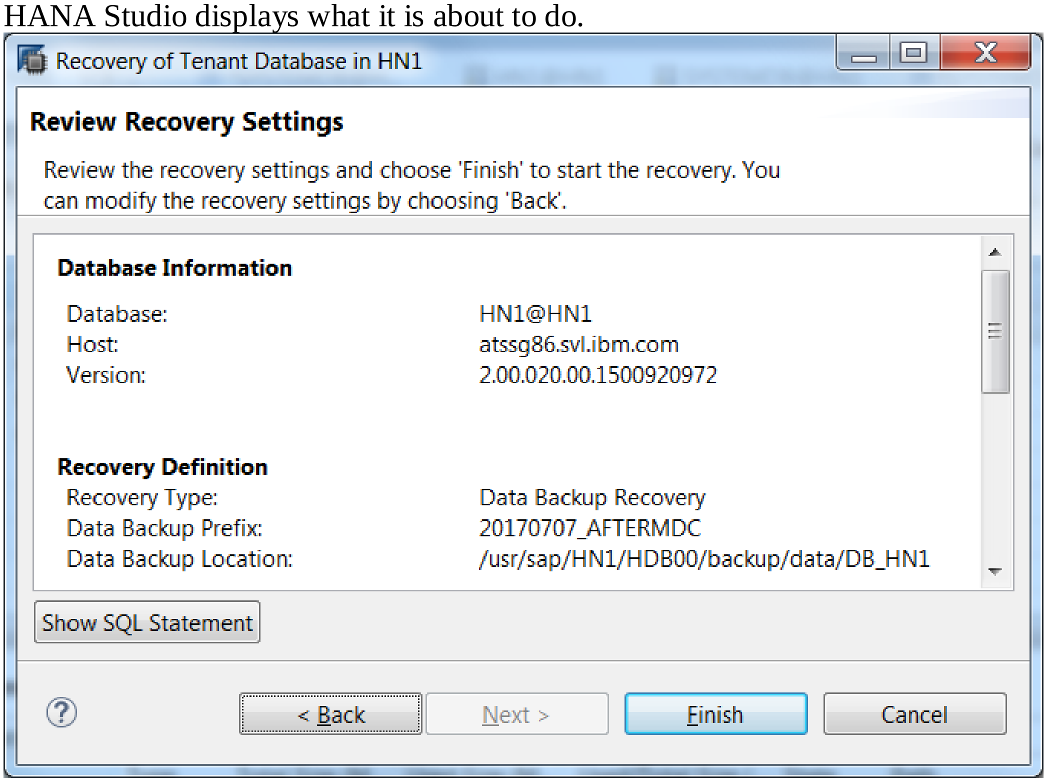

Figure 30: Recover DB settings

Press OK to stop the DB and go on.

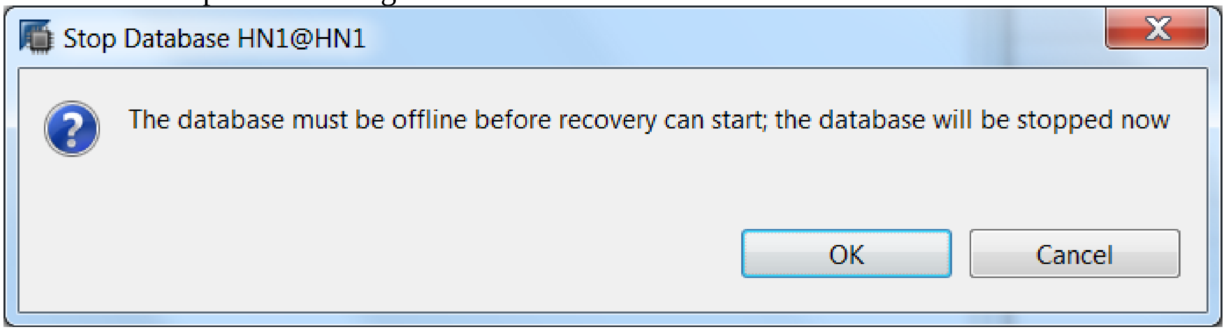

Figure 31: Pop-up warning

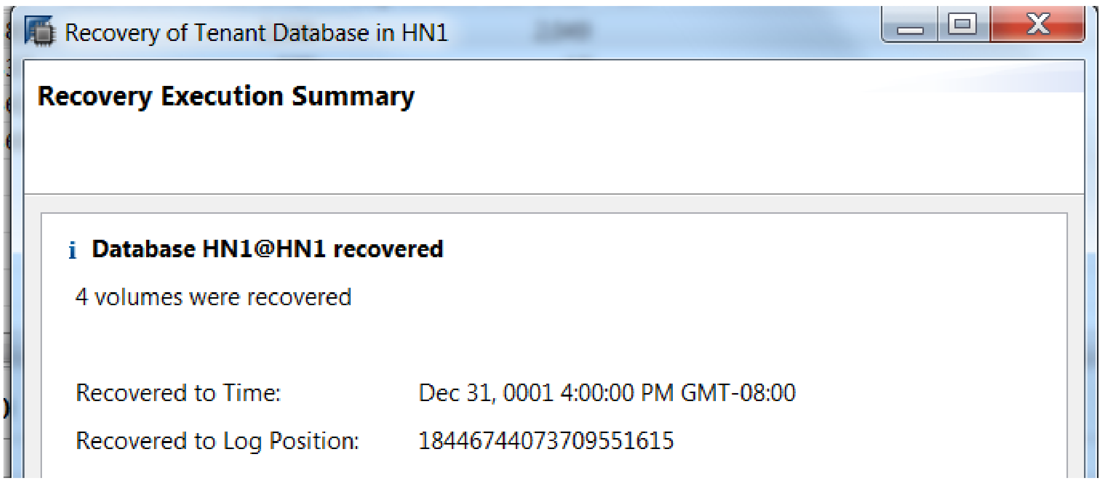

Figure 32: Recover complete

The HN1 tenant DB has been copied from x86 to POWER.

Shutdown, backup and restart the HANA DB on the target system.

7.6. Update SAP application server and restart SAP instance

Update the HANA client using the HANA 2.0 SPS02 install DVD, so that the application server HANA client matches the DB version.

```
atssg140:/tmp #  /repo/Power*/v20_SPS02_install/20170804/51*/DA*/HDB_CLIENT_LINUX_X86_64/hdbinst
SAP HANA Database Client installation kit detected.


SAP HANA Lifecycle Management - Client Installation 2.2.23.1499440855
*********************************************************************


Select a SAP HANA Database Client installation:

No  | Installation Path                     | Version                   | Mode
-----------------------------------------------------------------------------
[1] | /usr/sap/B7Z/hdbclient                | 1.00.120.050.1484020229 | 64bit
    |                                       |                           |
 2  | Install new SAP HANA Database Client | 2.2.23.1499440855         | 64bit

Enter number [1]: 1
```

Figure 33: Update HANA client on application server

The SAP ABAP application server uses the "hdbuserstore" configuration to locate the DB server. This is still pointing to the old scale-out x86 DB server.

```
atssg140:/tmp # su - b7zadm
atssg140:b7zadm 53> hdbuserstore LIST
DATA FILE       : /home/b7zadm/.hdb/atssg140/SSFS_HDB.DAT
KEY FILE        : /home/b7zadm/.hdb/atssg140/SSFS_HDB.KEY

KEY DEFAULT
  ENV : atssg141.svl.ibm.com:30015;atssg143.svl.ibm.com:30015;atssg142.svl.ibm.com:30015
  USER: SAPABAP1
```

Figure 34: Original hdbuserstore LIST configuration

We will update hdbuserstore on the application server to point to the new DB server.

```
  USER: SAPABAP1
atssg140:b7zadm 54> hdbuserstore SET DEFAULT atssg86:30015 SAPABAP1 password
```

Start the application server and login to SAP, and we're done. The DB is running on HANA on POWER.

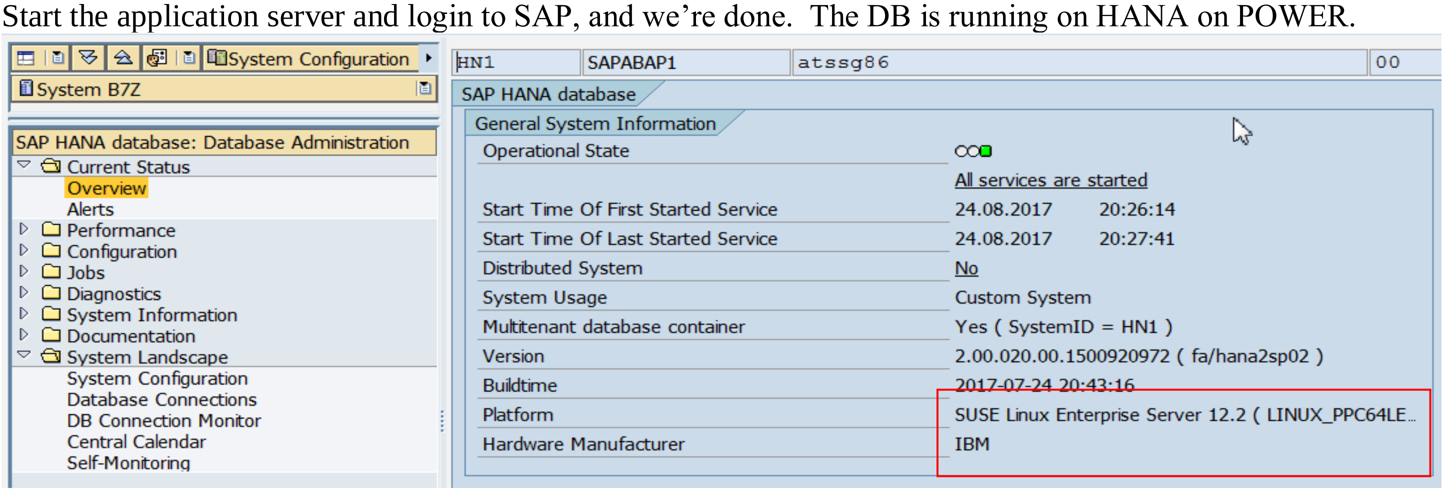

Figure 35: DBACOCKPIT transaction display DB

7.7. (Optional) Reduce number of index servers

Additional changes, such changing table partitioning or changing the number of indexservers, can be implemented now or later. SAP notes 1986612 and 2447887 have information on removing indexservers.

To get the full benefit of changing from a scale-out to scale-up DB server, it can be beneficial to remove extra indexservers after migration so that there is one indexserver, as described in SAP note 2103956. See SAP note 1986612 for the steps to remove an indexserver.

Here, tables and partitions are distributed across all three indexservers.

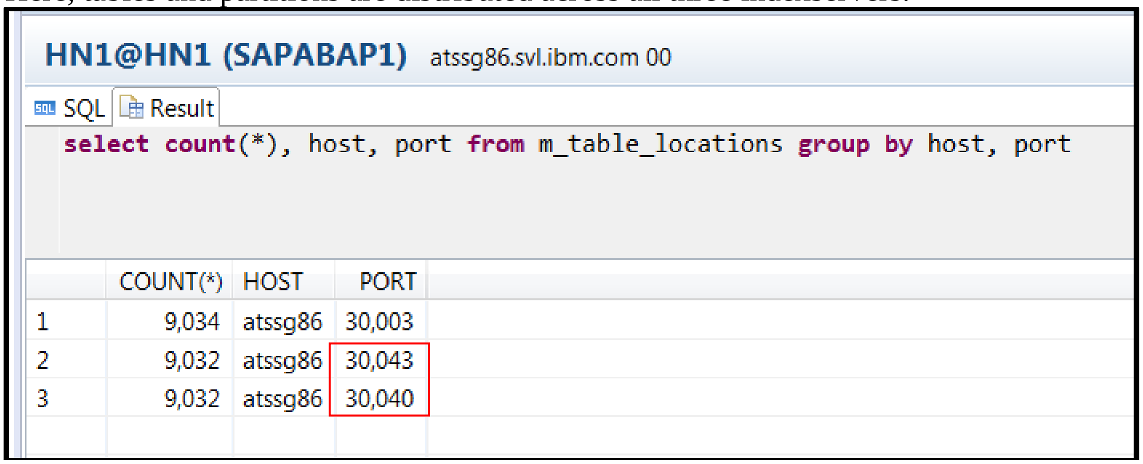

Figure 36: m_table_locations with three indexservers

As described in SAP note 1986612, set the two added indexservers to be inactive, so that tables and partitions will be moved off during reorg. Port 30003 in Figure 36 is the default indexserver.

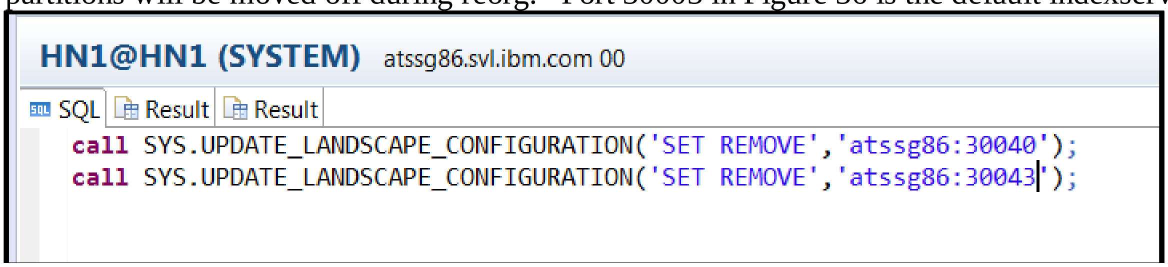

Figure 37: SET REMOVE

Create the reorg job.

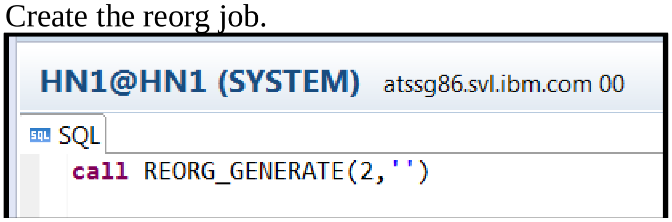

Figure 38: Create reorg job

Then start it.

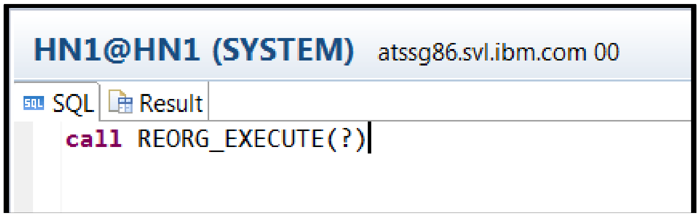

Figure 39: Execute reorg job

Track status of reorg job in HANA Studio, or using the queries in note 1986612.

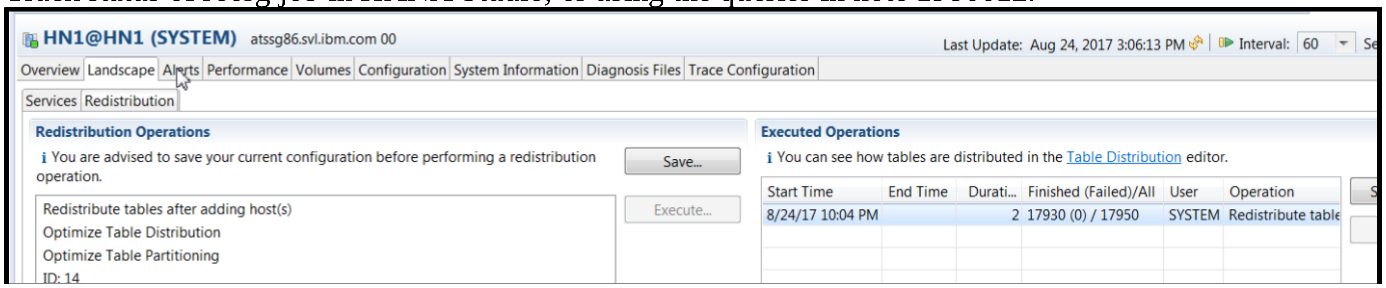

Figure 40: REORG job status

Check status of the reorg. Here, it is finished.

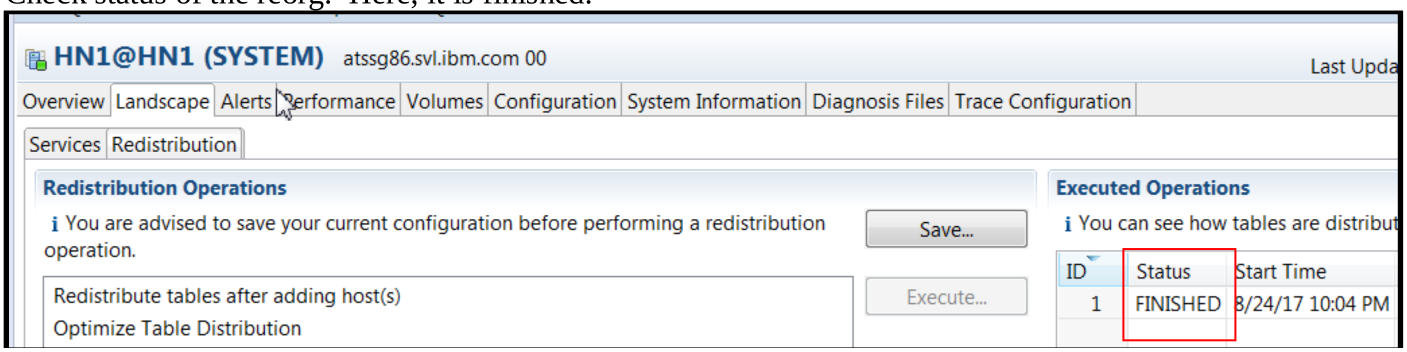

Figure 41: Reorg is finished

After the reorg has finished, the volume status shows the tables have been removed.

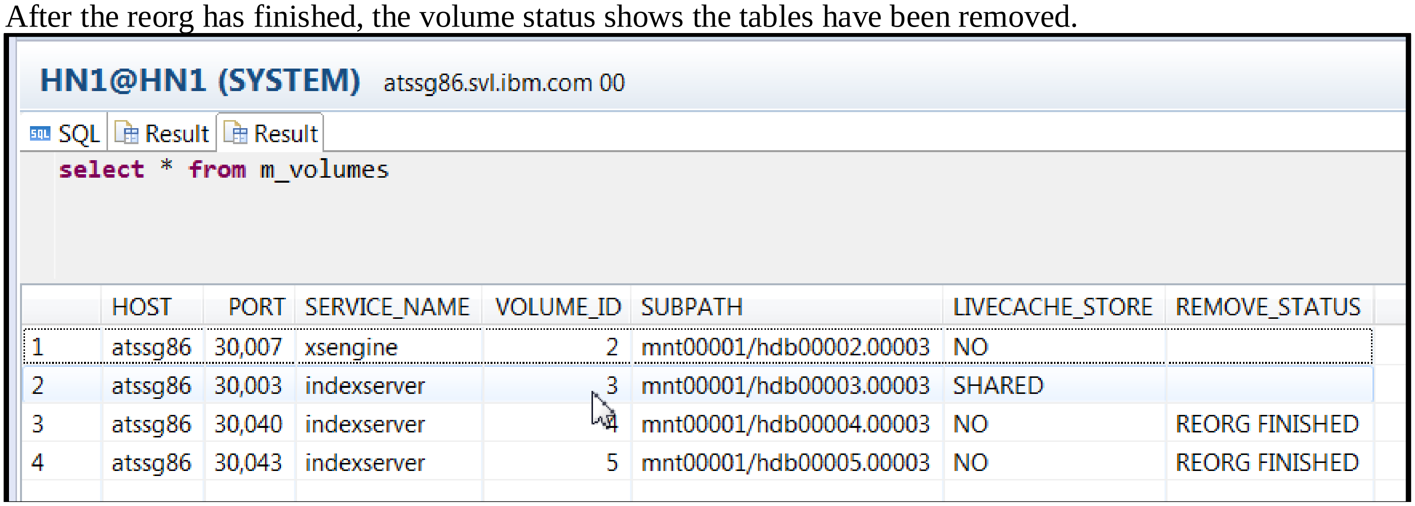

Figure 42: check volume status

Now we can remove the indexservers, which will also remove their volumes.

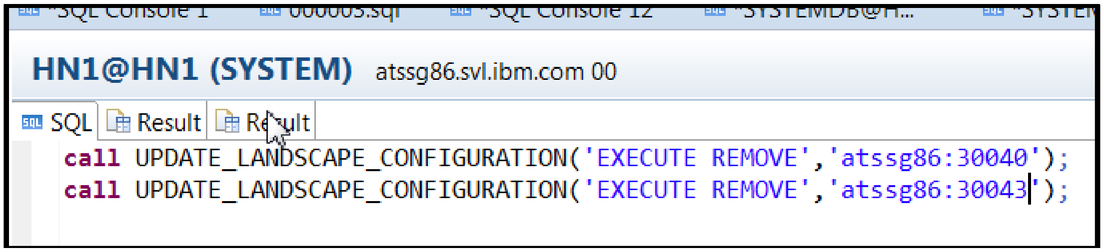

Figure 43: Remove the indexservers

Display the volumes for the HN1 tenant. There is now only one indexserver for HN1.

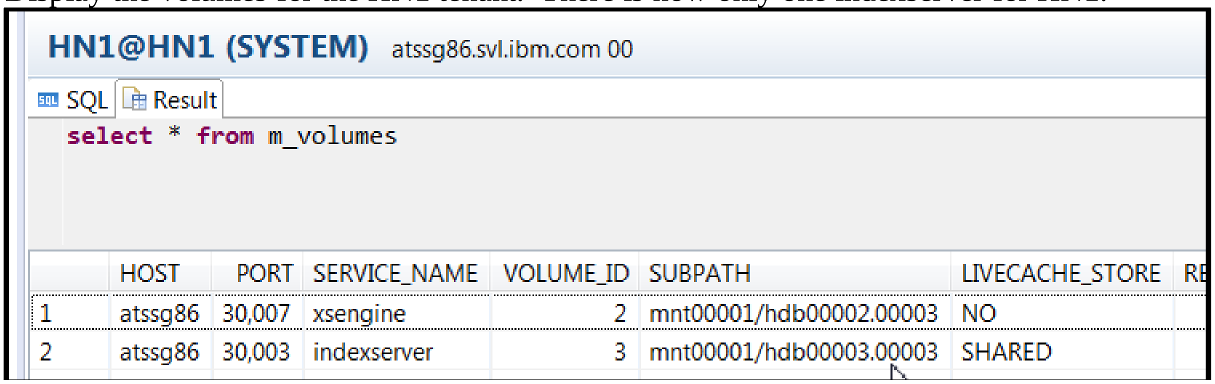

Figure 44: volume status after extra indexservers removed

Rather than using HANA reorg to move all tables off an indexserver, one can move individual tables and partitions using SQL command line.

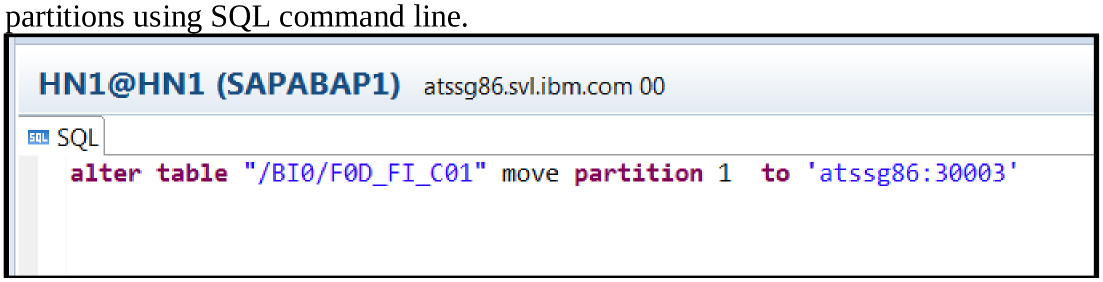

Figure 45: move partition with SQL

Note that reducing the number of indexservers as done above does not affect the partitioning of tables. Tables that were partitioned on the scale-out configuration will still be partitioned. Partitioning a large table in a single indexserver may improve performance, so it is best to test the performance impact before converting a table from partitioned to unpartitioned.

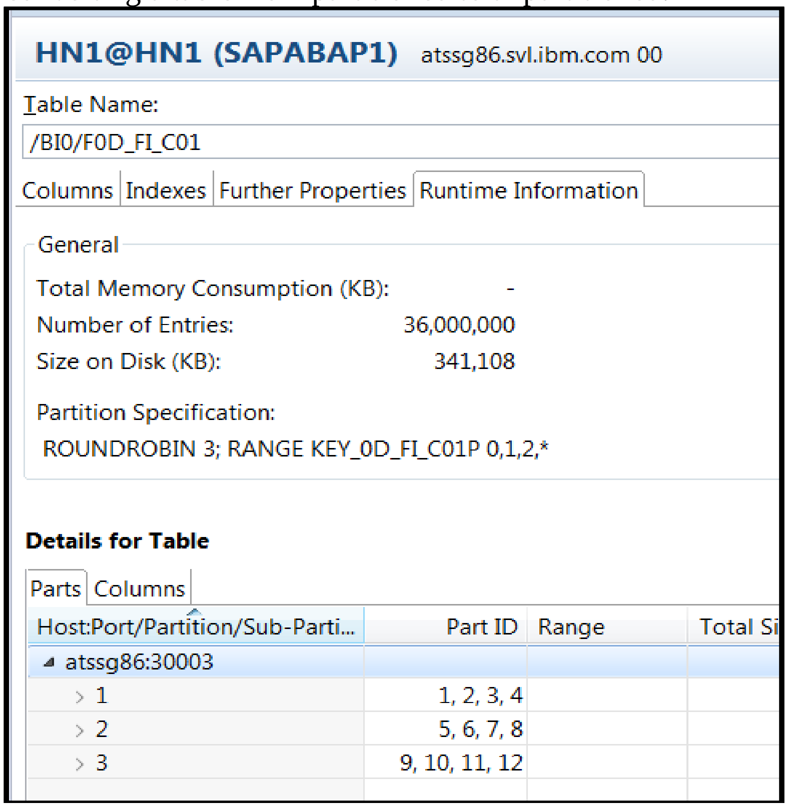

Figure 46: Partitioned table on single indexserver

8. Summary

Our goal for this paper was to demonstrate the recently available backup-based migration process to HANA on POWER. For systems where the SAP software version supports HANA 1.0 on x86 and HANA 2.0 on POWER, this can offer a fast and simple DB migration path to HANA on POWER.

The key steps were:

- convert the source DB to multi-tenant
- backup tenant on source x86 HANA DB
- install HANA on POWER
- add HANA server processes (e.g. indexservers) on POWER to match x86 topology
- restore tenant into HANA DB on POWER
- (optional) remove HANA on POWER server processes added to match source system topology.

