Header Data

Released On	09.04.2015 10:04:52
Release Status	Released for Customer
Component	HAN-DB SAP HANA Database
Other Components	HAN-LM-INS SAP HANA Installation
Priority	Recommendations / Additional Info
Category	Consulting

Symptom

This SAP Note describes problems and provides tips during the system copy or migration of an SAP system with or to SAP HANA with the current SWPM 1.0 from SAP Service Marketplace.

Other Terms

Migration, system copy, SAP HANA, r3load, migmon, migmonctrl, migration monitor controller

Reason and Prerequisites

The current version of the SWPM is located here: http://service.sap.com/sltoolset
-> System Provisioning
 -> Software Provisioning Manager 1.0
 -> Download

Solution

SAP HANA SP 9, multitenant database containers

As of SAP HANA SP9, the following new feature is available: multitenant database containers. This feature enables you to run several completely separate instances, known as tenants, within a database.

The *software provisioning manager* also supports this feature. The first version of SWPM SP7 supports the following features:

- Using an existing tenant to install and deinstall SAP systems
- Using an existing tenant to perform a system copy or to migrate systems by means of R3load

With SWPM 1.0 SP7 PL3, additional features are now supported:

- Creating and dropping tenants as part of the installation/deinstallation process
- A system copy by means of a backup/restore in a tenant
- Renaming an SAP system whose data lies in a tenant

For more information about this SWPM, please refer to the SCN at: http://scn.sap.com/community/it-management/alm/software-logistics/blog/2013/10/04/migration-to-sap-hana-latest-news-about-software-provisioning-manager

Known problems

Exporting the source system

October 24, 2014:

Using the new table splitting tool SAPuptool in SUM with an Oracle source database.

As of SUM Support Package 11 Patch 6, select count(*) will receive the Oracle hint "/* parallel */".
With this hint, Oracle parallelizes the SQL statement, leading to a better performance.
The number of parallel threads is defined via the Oracle parameter "parallel_max_servers".
If the number is too high, this may lead to terminations if too many SAPuptool processes are running in parallel.

Setting up the target system

- We recommend that you implement the latest revision of the SAP HANA database.
- You must pay attention to SAP Notes 1921023, 1990894, and 1991638 with regard to this.

Known problems:

4/9/2015

If you want to use virtual host names, see SAP Note 2141442.

February 06, 2015

Error during execution of HdbCmdClazz 'EXECUTE_DATABASE_REORG' : Fri Feb 06 14:31:10 CET 2015 | com.sap.hdb.core.main.cmd.HdbCmdMain | For input string: "undefined"

The error occurs only on scale-out systems if the file HDB_TABLE_CLASSIFICATION.TXT does not exist in the directory <exportDVD>/ABAP/DB. As a workaround, create a blank file with this name in the directory. Then choose "Retry" in the *software provisioning manager*.
The error is corrected as of version SWPM 1.0 SP7 PL 5.

February 05, 2015

SAP HANA Multitenant database containers and SAP system on one server.

The current version of the *Software Provisioning Manager* does not support so-called *one host installation* in which the SAP HANA server and the SAP system are installed on one server.
As a workaround, the only option is to install the SAP system on another server.
The scenario is supported as of version SWPM 1.0 SP7 PL4.

January 13, 2015:

Error during an installation/system copy involving a scale-out system whereby the tenant is not located on the master host.
Error message:
System (Last error reported by the step:
Process call '/usr/sap/ERP/SYS/exe/uc/linuxx86_64/R3load -testconnect' exits with error code 2. For details see log file(s) R3load.exe.log.).

<sid>adm 1> hdbuserstore list
KEY DEFAULT
  ENV : ld9968.wdf.sap.corp:30044;ld2432.wdf.sap.corp:30044;ld9970.wdf.sap.corp:30044
  USER: SAPERP

In this example, the tenant is located on the node ld2431.
The *software provisioning manager* enters the potential master hosts of the scale-out system in hdbuserstore in descending order.
If the tenant's master service is located on a node other than the first node in hdbuserstore, and if this first node can be accessed, the client tries to connect to this host. This action fails because the tenant is running on a different node.
The first node in hdbuserstore must be the node on which the tenant's master service is running. To determine the relevant node, you can execute the following SQL statement on the system database:

select * **from** SYS_DATABASES.M_SERVICES **where** database_name = '<TENANT>' **and** COORDINATOR_TYPE = 'MASTER'

The user store is then corrected using this command:

hdbuserstore set DEFAULT "ld2431.wdf.sap.corp:30044;ld9968.wdf.sap.corp:30044;ld2432.wdf.sap.corp:30044;ld9970.wdf.sap.corp:30044" SAPERP <password>

The client should then be able to successfully connect to the host. As soon as the node ld2431 fails, all services are moved to the next free potential master host. The client automatically connects to this node.

December 22, 2014:

Installation and system copy. The ABAP roles and AFL roles are not created.
SWPM 1.0 SP7 PL3 corrects this.

Software Provisioning Manager 1.0 SP7 no longer creates the roles AFL_* and ABAP_*.
This is an error in the HdbSlLib. To subsequently create the roles after the installation, you must execute the following command in the installation directory:

./sapjvm/sapjvm_6/bin/java -classpath com.sap.hdb.sl.lib.jar:com.sap.hdb.core.jar:ngdbc.jar com.sap.hdb.core.main.cmd.HdbCmdMain
HDB_CMD_CLAZZ_NAME=UPDATE_REPAIR_ROLES_AND_PRIVILEGES JDBC_PATH=. DATABASE_SID=**<DBSID>**
DATABASE_HOSTNAME=**<hostName>** INSTANCE_NUMBER=**<Number>** DATABASE_SYSTEM_USER_PASSWORD=**<Password>**
DATABASE_SQL_USER_NAME=**<Name>** DATABASE_SQL_USER_PASSWORD=**<Password>**

November 14, 2014:

SCM on SAP HANA: Error in the dialog phase of the *software provisioning manager*.
This is corrected with SWPM 1.0 SP6 PL6.

HdbCmdOut.log des com.sap.hdb.sl.lib.jar in the installation directory:
com.sap.hdb.core.main.cmd.HdbCmdMain | Error during execution of HdbCmdClazz section!
com.sap.hdb.core.main.cmd.HdbCmdMain | No enum const class com.sap.hdb.sl.lib.instance.Instance$ServiceName.SCRIPTSERVER

This error occurs in the archive com.sap.hdb.sl.lib.jar.

July 15, 2014:

If you create a system copy of an SAP system based on SAP NetWeaver 7.40 SP5 or higher, in which the ABAP source search is active (switch: SRIS_SWITCH_SOURCE_SEARCH), see also SAP Note 2031638.

July 11, 2014:

For potential problems during the installation of scale-out systems, see SAP Note: 2041412 - Error during initial landscape redistribution during

migration via SWPM

May 26, 2014:

If the output of the ABAP report SMIGR_CREATE_DDL is very large, read the following SAP Note: 2022386

March 19, 2014:

Additional patches for the R3load problem.

With SWPM 1.0 SP4 PL8 and SWPM 1.0 SP5 PL0, the option "-loadprocedure dbsl" is appended to the R3load calls for every migration to SAP HANA.
This safely deactivates the mass loader function of R3load.
During the loading of data into the SAP HANA database, if you use a R3load that contains the fix (see SAP Note 1990894), you can replace this R3load in the archive SAPEXEDB.SAR. Proceed as described in SAP Note 1826587.
To reactivate the mass loader function, set the environment variable HDB_MASSIMPORT=YES before the start of the SWPM. This also overrides the above option "-loadprocedure dbsl".

March 17, 2014:

Fix and ABAP check report for R3load problem

There is now a fix for the R3load problem described on March 12, 2014.
SAP Note 1990894 contains a table listing the fixed versions.
SAP Note 1991638 contains an ABAP report that can be implemented in the target system and checks the tables that might be affected by the error.
This SAP Note also contains a PDF document describing how to bypass the error in the software provisioning manger by deactivating the mass loader; the setting of the environment variable HDB_MASS_IMPORT=NO is not sufficient.

We recommend that you use the fixed version of the R3load in all cases. With SWPM 1.0 SP4 PL8, the mass loader in R3load is deactivated until further action is taken.
This takes place using the R3load parameter "-loadprocedure dbsl". If you use the fixed version of the R3load, you can activate the mass loader again.
You do this by setting the environment variable HDB_MASSIMPORT=YES before executing the software provisioning manager.

March 12, 2014:

Due to a problem in R3load, you must use the mass loader function, which speeds up loading of the LOBS.

You should not use any R3loads for the import that you have compiled after November 2013.
You can check the version or date using the command "R3load -version" ("compile time").
For older R3loads, you should continue to set the following environment variable: "HDB_MASSIMPORT=NO"
Set this variable before you execute the SWPM.

February 2, 2014:

> For migrations on the AIX platform, see SAP Note: 1976256 - AIX: TCP/IP connections hang

> February 4, 2014:

> The installation terminates with the following message in HdbCmdOut.log:
> '| com.sap.hdb.core.main.cmd.HdbCmdMain | 0 Unexpected result size. Expected 1 for sql query: SELECT ACTIVE_STATUS ...'

> This is corrected with SWPM 1.0 Support Package 4 Patch Level 5.

> The error occurs if the "new" statistics server is active. This is part of the index server. As a result, there is no longer a separate service for the statistics server. This is the reason why the query fails. Until the fix is available, we recommend that you activate the statistics server only after installation.

> For more information about the new statistics server, see: http://www.saphana.com/docs/DOC-4403

> October 21, 2013:

> There is an error during the import. R3load reports:
> "(RFF) ERROR: no entry for table , aborting the task."

> The error occurs during a homogeneous system copy from SAP HANA to SAP HANA using R3load. The cause of this is in the R3load.
> For a detailed description of the problem, see SAP Note 1928043.

Installation of the current SAP HANA client

> To install the current SAP HANA client directly, follow the instructions in SAP Note 1825053.

> Kernel archive

> Before the migration, check whether there are updated kernel archives for your kernel on SAP Service Marketplace (SAPEXE.SAR and SAPEXEDB.SAR). If there are, download the latest archives. During the migration, specify the path for the archives in the dialog "Unpack Archives".

If there are no kernel archives that contain the latest patches of the R3 tool (R3load, R3ldctl, DBSL), proceed as described in SAP Note 1826587.

Export dump DVD

Exchange the file /ABAP/DB/HDB/rowstorelist.txt from SAP Note 1659383.

Check whether the files created by the ABAP report SMIGR_CREATE_DDL are in the directory ABAP/DB/HDB and if they can be read by everyone on Unix or Linux. The same applies to the file SQLFiles.LST in the directory ABAP/DB. If required, set the necessary authorizations (for example, chmod 755 /ABAP/DB/HDB/* chmod 755 /ABAP/DB/SQLFiles.LST)

SAP HANA parameters that must be set during the import

With SWPM 1. 0 SP 1, there is an SAPinst dialog with the name "SAP HANA Import Parameters". This is displayed in the "Custom" installation mode. In the "Typical" mode, you can return to the dialog using the option "Detail -> Revise" on the summary screen. You can use the "Only during import" option to delete them again after the import or reset them to the original value. The dialog is already filled with the most important parameters. Currently, it is not necessary to set additional parameters.

SAP HANA mass import

If the SAP HANA target database is used in revision 52 or higher and kernel 7.21 or higher, the import with R3load can be accelerated by activating the mass importer. To do so, you set the following environment variable before the start of the SWPM:

HDB_MASSIMPORT=YES

Caution: If you activate the mass importer for an SAP HANA revision lower than 52, there may be a loss of data (see SAP Note 1806935).

You can now start the setup of the target system.

Analysis of the R3load export and import times
In most cases, a few long-running packages prescribe the entire migration runtime. The runtimes can be reduced simply by splitting the packages in question or extracting long-running tables from the packages.
You can use the archive MIGTIME.SAR to analyze the runtimes of the individual packages. It is stored on the SWPM in the directory /COMMON/INSTALL/MIGTIME.SAR. It is unpacked by means of SAPCAR to the installation directory (the directory that also contains sapinst_dev.log). It contains documentation in addition to the tools.

If you have a very high number of files (> 10,000) in the installation directory, the analysis using migtime.jar might take a long time (> 1h). To reduce this time, you can proceed as follows:

- Create a directory as follows: "mkdir timeAnalyzer"

- Copy all log files to this new directory: "cp *.log timeAnalyzer"

- Navigate to the directory and delete the following files: "cd timeAnalyzer", "rm *__D*.log"

- Start the TimeAnalyzer from this directory.

Validity

This document is not restricted to a software component or software component version

References

This document refers to:

SAP Notes
2141442   SAP HANA client software: hdbuserstore not found when using virtual host names
2041412   Error during initial landscape redistribution during migration via SWPM
2031638   R3load terminates during creation of an index for a field of the type LOB
2022386   Migration to SAP HANA: Software provisioning manager terminates when executing HdbSlLib: java.lang.OutOfMemoryError: Java heap space
1991638   Supplementary SAP Note to SAP Note 1990894
1990894   Data corruption in tables with data type DECFLOAT during SAP HANA migration
1958346   BW on SAP HANA: Landscape redistribution check procedures
1930853   HdbUserStore contains too many SAP HANA nodes
1921023   SMIGR_CREATE_DDL: Corrections and enhancements for SAP HANA
1900822   Using SAP HANA landscape reorg from SWPM
1900445   HANA R3load import: (DB) ERROR: DDL statement failed
1873729   createoderby.jar to optimize the export during systemcopy.
1826587   SAP NetWeaver installation using SWPM 1.0:exchange R3load

1825053 Installation of the current HANA client with SWPM
1819123 BW on SAP HANA SP5: landscape redistribution
1806935 SAP HANA: Corrupt database after R3load import
1783927 Prerequisites for table splitting with target SAP HANA database
1706931 Inst. SAP Sys. Based on NW 7.3 and higher: SAP HANA DB, Win
1706930 Inst. SAP Sys. Based on NW 7.3 and higher: SAP HANA DB, UNIX
1659383 RowStore list for SAP Netweaver 7.30/7.31 on SAP HANA database
1600929 SAP BW powered by SAP HANA DB: Information

This document is referenced by:

SAP Notes (16)

1900445 HANA R3load import: (DB) ERROR: DDL statement failed
1667731 Rapid database migration of SAP BW to SAP HANA
1659383 RowStore list for SAP Netweaver 7.30/7.31 on SAP HANA database
1826587 SAP NetWeaver installation using SWPM 1.0:exchange R3load
1825053 Installation of the current HANA client with SWPM
1821999 Rapid database migration of SAP Business Suite to SAP HANA
1600929 SAP BW powered by SAP HANA DB: Information
1819123 BW on SAP HANA SP5: landscape redistribution
1806935 SAP HANA: Corrupt database after R3load import
2096175 Software Provisioning Manager older then SWPM 1.0 SP7 will not be able to connect to a Tenant Database
1930853 HdbUserStore contains too many SAP HANA nodes
1873729 createoderby.jar to optimize the export during systemcopy.
1783927 Prerequisites for table splitting with target SAP HANA database
1783238 SMIGR_CREATE_DDL: Estimated row count for fact tables
1706931 Inst. SAP Sys. Based on NW 7.3 and higher: SAP HANA DB, Win
1706930 Inst. SAP Sys. Based on NW 7.3 and higher: SAP HANA DB, UNIX

Attachments

File Name	File Size (KB)	Mime Type
MIGTIME.SAR	349	application/octet-stream
MIGMON.SAR	1525	application/octet-stream