



SAP HANA Installation Guide

- SAP In-Memory Appliance (SAP HANA) 1.0

2011-09-02

Copyright

© 2011 SAP AG. All rights reserved. SAP, R/3, SAP NetWeaver, Duet, PartnerEdge, ByDesign, SAP BusinessObjects Explorer, StreamWork, and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and other countries. Business Objects and the Business Objects logo, BusinessObjects, Crystal Reports, Crystal Decisions, Web Intelligence, Xcelsius, and other Business Objects products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of Business Objects Software Ltd. Business Objects is an SAP company. Sybase and Adaptive Server, iAnywhere, Sybase 365, SQL Anywhere, and other Sybase products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of Sybase, Inc. Sybase is an SAP company. All other product and service names mentioned are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary. These materials are subject to change without notice. These materials are provided by SAP AG and its affiliated companies ("SAP Group") for informational purposes only, without representation or warranty of any kind, and SAP Group shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP Group products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

2011-09-02

Contents

Chapter 1	Installing SAP HANA Overview.....	5
1.1	Software Components.....	5
1.2	Software Download.....	6
1.3	Hardware and Software Requirements.....	7
1.4	SAP HANA Guides	7
Chapter 2	Installing the SAP In-Memory Database Components.....	9
Chapter 3	Installing Trigger-Based Replication.....	11
Chapter 4	Installing ETL-Based Replication.....	13
Chapter 5	Installing Log-Based Replication.....	15
5.1	Installing the SAP HANA Load Controller.....	15
5.1.1	Use and Prerequisites.....	15
5.1.2	Deploying the SAP HANA Load Controller and Related Components.....	17
5.1.3	Manually Configuring SAP HANA Load Controller on SAP HANA System	18
5.1.4	Using the SAP HANA Load Controller.....	21
5.2	Installing the SAP Host Agent.....	22
5.3	Installing the Sybase Replication Server.....	22
5.3.1	Installation and Configuration of the Sybase Replication Server and ECDA	26
5.3.2	Installation and Configuration of the Sybase Replication Agent	28

Installing SAP HANA Overview

The SAP HANA Installation Guide is structured according to the available replication scenarios and the components that are needed for these scenarios:

- SAP In-Memory Database Components (required for all scenarios)
- Trigger-Based Replication
- ETL-Based Replication
- Log-Based Replication

1.1 Software Components

SAP HANA appliance software is available in different editions:

- SAP HANA appliance software platform edition 1.0
- SAP HANA appliance software enterprise edition 1.0
- SAP HANA appliance software enterprise extended edition 1.0

Components	Platform	Enterprise	Extended
SAP HANA database 1.00	x	x	x
SAP HANA studio 1.00	x	x	x
SAP HANA client 1.00	x	x	x
SAP Host Agent 7.20	x	x	x
LT Replication AddOn		x	x
LT Replication Server		x	x
SAP BusinessObjects Data Services 4.0		x	x
Sybase Adaptive Server Enterprise (ASE) 15			x
Sybase Replication Server 15 (incl. ECDA)			x
Sybase Replication Server Agent 15			x
SAP HANA Load Controller 1.00			x

The **SAP HANA appliance software platform edition 1.0** is intended for customers who want to use the ETL-Based Replication and already have a license for SAP BusinessObjects Data Services. It comprises the following components:

- SAP HANA database 1.00 (formerly known as SAP In-Memory Database - Server)
- SAP HANA studio 1.00 (formerly known as SAP In-Memory Database - Studio)

- SAP HANA client 1.00 (formerly known as SAP In-Memory Database - Client)
- SAP HANA Host Agent 7.20

The **SAP HANA appliance software enterprise edition 1.0** is intended for customers who want to use either Trigger-Based Replication or ETL-Based Replication and do not already have all necessary licenses for SAP BusinessObjects Data Services. It comprises the following components:

- SAP HANA database 1.00 (formerly known as SAP In-Memory Database - Server)
- SAP HANA studio 1.00 (formerly known as SAP In-Memory Database - Studio)
- SAP HANA client 1.00 (formerly known as SAP In-Memory Database - Client)
- SAP HANA Host Agent 7.20
- LT Replication AddOn
- LT Replication Server
- SAP BusinessObjects Data Services 4.0

The **SAP HANA appliance software enterprise extended edition 1.0** is intended for customers who want to use the full potential of all available replication scenarios including the Log-Based Replication. It comprises the following components:

- SAP HANA database 1.00 (formerly known as SAP In-Memory Database - Server)
- SAP HANA studio 1.00 (formerly known as SAP In-Memory Database - Studio)
- SAP HANA client 1.00 (formerly known as SAP In-Memory Database - Client)
- SAP HANA Host Agent 7.20
- LT Replication AddOn
- LT Replication Server
- SAP BusinessObjects Data Services 4.0
- Sybase Adaptive Server Enterprise (ASE) 15
- Sybase Replication Server 15 (incl. ECDA)
- Sybase Replication Server Agent 15
- SAP HANA Load Controller 1.00

1.2 Software Download

The components of SAP HANA and of the SAP In-Memory Database can only be installed by certified hardware partners on validated hardware running a specific operating system. Any other system or content developed with such systems is not supported by SAP. For further information please refer to the information page of the product version. Support Package Stacks (SPS) can be downloaded and applied to appliances according to agreements with the respective hardware partner.

1.3 Hardware and Software Requirements

Note:

You find the complete list of all SAP HANA components and the respective SAP HANA hardware and software requirements in the Product Availability Matrix (PAM) on SAP Service Marketplace <https://service.sap.com/pam>.

Software Requirements

Operating System for SAP HANA 1.0

SUSE Linux Enterprise Server (SLES) 11 SP1 (see [SAP Note 1310037](#) for information about installing SLES 11 SP1 in an SAP environment)

SAPCAR 7.10

(see [SAP Note 212876 - The new archiving tool SAPCAR](#))

1.4 SAP HANA Guides

For more information about SAP HANA landscape, security, installation, and administration, see the resources listed in the table below:

Topic	Guide/Tool	Quick Link
SAP HANA Landscape Deployment & Installation	SAP HANA Knowledge Center on SAP Service Marketplace	https://service.sap.com/hana <ul style="list-style-type: none"> • SAP HANA 1.0 Master Guide • SAP HANA 1.0 Installation Guide
SAP HANA Administration & Security	SAP HANA Knowledge Center on SAP Help Portal	http://help.sap.com/hana <ul style="list-style-type: none"> • SAP HANA 1.0 Technical Operations Manual • SAP HANA 1.0 Security Guide

Installing the SAP In-Memory Database Components

The SAP In-Memory Database Components are:

- SAP In-Memory Database 1.0 - Server
- SAP In-Memory Database 1.0 – Studio (incl. IMDB JDBC 1.0)
- SAP In-Memory Database 1.0 - Clients
 - IMDB ODBO 1.0
 - IMDB JDBC 1.0
 - IMDB ODBC 1.0
 - IMDB SQLDBC 1.0

For information about SAP HANA installation and upgrade, see the SAP HANA Installation & Implementation Knowledge Center on SAP Service Marketplace (SMP) <https://service.sap.com/hana>
-> *Installation & Upgrade*.

Installing Trigger-Based Replication

For information about SAP HANA installation and upgrade for the Trigger-Based Replication, see the SAP HANA Installation & Implementation Knowledge Center on SAP Service Marketplace (SMP) <https://service.sap.com/hana> -> *Installation & Upgrade*.


Installing ETL-Based Replication

In the Log-Based Replication scenario the ERP system is the data provider for SAP HANA and those data can be analyzed by means of MS EXCEL as end user client on top.

In addition to those SAP HANA components the customer can add existing SAP BusinessObjects products to the SAP HANA landscape. The information below provides a list of the available documentation to integrate Data Services to your SAP HANA landscape.

In addition to those SAP HANA components the customer can add existing SAP BusinessObjects products to the SAP HANA landscape to use the SAP BusinessObjects end user clients like SAP BusinessObjects Explorer or SAP BusinessObjects Xcelsius on top to analyze those data.

SAP HANA 1.0 Landscape Including External SAP BusinessObjects Servers



If you are not using BusinessObjects Enterprise with Data Services you need Information platform services installed together with Data Services. You will find the necessary installation and upgrade information in the following guides:


- [Information platform services Installation Guide for UNIX](#)
- [Information platform services Installation Guide for Windows](#)
- [Data Services Installation Guide for UNIX](#)
- [Data Services Installation Guide for Windows](#)
- [Data Services Upgrade Guide](#)
- [Sybase Adaptive Enterprise 15.5](#)

Installing Log-Based Replication

5.1 Installing the SAP HANA Load Controller

5.1.1 Use and Prerequisites

The figure below shows how the Sybase Replication Server components and the SAP HANA Load Controller enable continuous data replication from the ERP system that is the data provider for SAP HANA. The role of the SAP HANA Load Controller in connection with the information modeler is also shown:



The SAP HANA Load Controller coordinates the SAP HANA data replication as follows:

- The information modeler starts the SAP HANA Load Controller.
- The Load Controller starts the initial load of the SAP ERP data (source system) to the SAP In-Memory Database on the SAP HANA system (target system). This is done using the R3load component on the SAP ERP system (source system) by exporting the data from the source system and importing the data to the in-memory database on the SAP HANA system (target system).
- In parallel, the Sybase Replication Agent on the SAP ERP system is started and performs log mining on the ERP source database. It relays all relevant information to the Replication Server on the SAP HANA system.
- As the main component, the Sybase Replication Server coordinates both the Replication Agent and the ECDA component.
- The Sybase Enterprise Connect Data Access (ECDA) connects to SAP HANA as the target database using an ODBC driver and mirrors the changes from the source database of the ERP system that are relayed through the Replication Server.
- After the initial load has been successfully executed and the changes in the source database during the initial load have been successfully replicated, the Sybase Replication Server and the related components perform the ongoing continuous data replication from the SAP ERP system to the SAP HANA box.

Prerequisites

- The SAP HANA system (target system) and its components have to fulfill the requirements (see section *Hardware and Software Requirements* in the SAP HANA Master Guide).
- You have installed the Sybase Replication Server, Sybase ECDA and Sybase Replication Agent as described here: [Installation and Configuration of the Sybase Replication Server and ECDA](#) .
- You have downloaded the SAP HANA Load Controller from the SAP Service Marketplace (see [Software Download](#)).
- The SAP ERP system (source system) has to fulfill the following requirements:
 - R3load: 700 Patch 134 or above, 701 Patch 37 or above, 720 Patch 65 or above
 - R3ldctl: 700 Patch 1 or above, 701 Patch 1 or above, 720 Patch 2 or above
 - SAP Host Agent Patch 46 or above

Deployment and Configuration of the SAP HANA Load Controller and Related Components

To enable the SAP HANA Load Controller, you have to deploy and partially configure the following components:

- Deploy the following files on the SAP ERP source system/data provider:


SAP Host Agent 7.20 (including the file hostctrl). For more information, see [Installing the SAP Host Agent](#).

A SAP Host Agent is normally part of an SAP ERP system and does not need to be installed.
- Deploy or install the following files or components on the SAP HANA system as target system:
 - HANALDCTR.SAR: contains the SAP HANA Load Controller itself

- SAP Host Agent 7.20 (incl. the file hostctrl). For more information, [Installing the SAP Host Agent](#).

Note that you only need the SAPHOSTAGENT_linuxx86_64.SAR file, since SAP HANA is running on the operating system SUSE Linux Enterprise Server (SLES) 11 SP1. For more information, see Hardware and Software Requirements in the SAP Master Guide.

Components needed for SAP HANA Load Controller



Extract the SAR archive files to obtain the following contents:

- HANALDCTR.SAR (=SAP HANA Load Controller), creates the following folders: bin, config, exe, export, import, log, READ.me, sql, wrk
- Patches for R3load and R3ldctl are available on SAP Service Marketplace (see [Software Download](#)).

5.1.2 Deploying the SAP HANA Load Controller and Related Components

On SAP ERP System (source system)

Prepare the file system

1. Login as the <sid>adm user; <sid> is the instance name of the ERP system (3 letters)
2. `cd /usr/sap/<sid>`
3. `mkdir HANA`
4. `mkdir HANA/export`
5. `mkdir HANA/export/DATA`
6. `mkdir HANA/export/wrk`
7. `mkdir HANA/export/DATA/ALL`

Mount this directory on target system as <initload>/export/DATA/ALL

Deploy the patches for R3load and R3ldctl

```
cp <PATH>/R3load<dbms_type> /sapmnt/<SID>/exe
```

On the SAP HANA system (target system)

Deploy the HANALDCTR.SAR file with the SAP HANA Load Controller

1. Login as the <sid>adm user; <sid> is the instance name of the SAP In-Memory Database system (3 letters)
2. cd /usr/sap/<sid>
3. mkdir HANA
4. mkdir HANA/LoadControl
5. cd HANA/LoadControl
6. SAPCAR -xvf HANALDCTR.SAR (creates the following folders: bin, config, exe, export, import, log, READ.me, sql, wrk)

Install SAP Host Agent 7.20 (including the file hostctrl)

For more information, see "Installing SAP Host Agent" . Note that you only need the SAPHOSTAGENT_linuxx86_64.SAR file, since SAP HANA is running on the operating system SUSE Linux Enterprise Server (SLES) 11 SP1. See *Hardware and Software Requirements* in the SAP HANA Master Guide.

5.1.3 Manually Configuring SAP HANA Load Controller on SAP HANA System

Configure the secure connections for the source system and the SAP In-Memory Database

1. Manage the password for the connection to the SAP Host Agent using the following program: <LoadControllerDir>/exe/LdCtlHASarter. Use the following command on the source system to provide the <sid>adm password to the Load Controller: <LoadControllerDir>/exe/LdCtlHASarter -loadControllerDir <LoadControllerDir> --updatePassword <password of sidadm-user>
2. Manage the secure database connection using the hdbuserstore tool. The connection to the database on the SAP HANA system uses the userkey "HANAREP". Provide the password of the database user to the Load Controller using the following command: hdbuserstore SET HANAREP <newdbhost>:3<instance number>15 <SAP HANA user> <Password>

Configure the configuration file reupload.ini for the SAP HANA Load Controller

1. Navigate to /usr/sap/<sid>/HANA/LoadControl/config/reupload.ini
2. Change the existing INI file reupload.ini.new to reupload.ini by configuring the following settings and parameters (Note that all parameters listed here need to be maintained. For example settings, see the end of the list):

portno:	Port number for R3load socket communication. Ensure that this port number is not in use on the source system or target system by other components. If you are configuring multiple data transfer channels, each LoadController needs a separate port number.
exportDir:	Location of the export directory on the source system . Default: /usr/sap/<sourceSID>/HANA/export .
newDBInstanceNO:	Instance number of the target system (SAP HANA instance).
newDBpasswd:	SAP HANA password. The default value is MANAGER
dbshdb_port:	Do not change
export dbshdb_port:	Do not change
schemaname:	Schema name in SAP In-Memory Database (default is SYSTEM)
#source system	
sourceSID:	SAPSYSTEMNAME (SID) of the target system
sourceHost:	Server name of an application server of the ABAP source system on which R3load is running and you have installed the SAP Host Agent configuration files.
upper_source_dbms_type:	dbms_type of the source system (in uppercase). For example: DB6
osuser:	OS user <sidadm> of the source system.
ospasswd:	Password of the OS user <sidadm> of the source
SourceDBUser:	Database user or schema of the source system.
migrationkey:	Migration key for R3load. To obtain the migration key, contact SAP Support.
SYBASE:	Location of Sybase installation
export SYBASE	Do not change
# SYBASE_OCS (location of OCS libraries)	
SYBASE_OCS:	Location of Sybase OCS libraries
export SYBASE_OCS:	Do not change
# SYBASE_REP (location of replication server)	
SYBASE_REP:	Location of Sybase Replication Server Software
export SYBASE_REP:	Do not change

repScriptPath:	Do not change
#isql (isql binary to use)	
isql:	Location and name of the Sybase program isql (<Path>/isql)
# Replication Agent and Replication Server names (entries in interfaces file)	
prs01:	Name of Sybase Replication Server
rax01:	Name of Sybase Replication Agent
# Scripts use default ID 'sa' and passwords below	
rs_pwd:	
ra_pwd:	
# set primary (ERP) connection name attributes (DS.DB)	
pds01:	Sybase connection name of primary dataserer
pdb01:	Sybase connection name of primary database
# set replicate (SAP HANA) connection name attributes (DS.DB)	
rds01:	Sybase connection name of dataserer on SAP HANA
rdb01:	Sybase connection name of database on SAP HANA

An example for the complete configuration of all settings and parameters in the configuration file `reload.ini` for the SAP HANA Load Controller:

```
# socket Portnumber
portno=5701

# Path to export Directory on Host of Sourcesystem
exportDir=/usr/sap/RSS/HANA/export
newDBInstanceNO=00
dbs_hdb_dbhost=lu0904.wdf.sap.corp
dbs_hdb_port=3${newDBInstanceNO}15
export dbs_hdb_port
dbs_hdb_userkey=HANAREP
export dbs_hdb_userkey

#<schemaname in newDB>
schemaname=SYSTEM
#source system
sourceSID=RSS
sourceHost=lu01161.wdf.sap.corp
upper_source_dbms_type=XYZ
# XYZ DB2
osuser=rssadm
ospasswd=schoko01
SourceDBUser=Sapsr3
```

```

#migrationkey=a
migrationkey=1G5fdEM50DqSq3egt6h]pGhK
SYBASE=/usr/sap/sybase
export SYBASE
# SYBASE_OCS (location of OCS libraries)
SYBASE_OCS=OCS-15_0
export SYBASE_OCS
# SYBASE_REP (location of replication server)
SYBASE_REP=REP-15_5
export SYBASE_REP
repScriptPath=/usr/sap/sybase/${SYBASE_REP}/scripts
#isql (isql binary to use)
isql=/usr/sap/sybase/${SYBASE_OCS}/bin/isql
# Replication Server name (entries in interfaces file)
prs01=HANARS1
# Scripts use default ID 'sa' and passwords below
rs_pwd=
ra_pwd=
if [ "${upper_source_dbms_type}" == "XYZ" ]; then
# Replication Agent name for orcale (entries in interfaces file)
rax01=HANARAO1
# set primary (ERP) connection name attributes (DS.DB)
pds01=HANARAO1
pdb01=HANARAO1
# set replicate (NewDB) connection name attributes (DS.DB)
rdb01=XYZ2NEWDB
rds01=HANADC1SVC
elif [ "${upper_source_dbms_type}" == "DB6" ]; then
# Replication Agent name for DB6 (entries in interfaces file)
rax01=HANARAU1
# set primary (ERP) connection name attributes (DS.DB)
pds01=HANARAU1
pdb01=HANARAU1
# set replicate (NewDB) connection name attributes (DS.DB)
rdb01=UDB2NEWDB
rds01=HANADC2SVC
else
echo "ERROR: unknown dbms_type!"
exit -40
fi

```

Verify the successful deployment and configuration of the SAP HANA Load Controller

- After configuring the SAP HANA Load Controller, test the connectivity to the source database, SAP In-Memory Database, Sybase Replication Agent, and Sybase Replication Server by executing the following command: `bin/LoadController.SH -testDBconnects` as os user `<sid>adm`. This call creates a number status tables in the SAP HANA system.
- If this test call has been successfully executed the following output appears:

```

check remote Shell Calls and exportDir
check saphost agent connection to source system and database
Connection to source system via saphostagent is possible.
Source Database is running.
check target database
Target Database is running.
Repagent and repserver are running.

```

5.1.4 Using the SAP HANA Load Controller

1. Navigate to the folder `/usr/sap/<sid>/HANA/LoadControl/exe`

2. Execute the command LoadController.SH

```
Usage: LoadController.SH  
-loadTablelist <tablename 1> <tablename 2> <tablename 3> ... <tablename n>  
-loadTablelistWithoutReplication <tablename 1> <tablename 2> <tablename 3> ... <tablename n>  
-onlyCreate <tablename 1> <tablename 2> <tablename 3> ... <tablename n>  
-allDBobjectsCreate  
-bruteForceLoad <tablename>  
-testDBconnects  
-getStatus  
-h  
-help
```

5.2 Installing the SAP Host Agent

Installation

For information about installing and using the SAP Host Agent, see


- Installation: [SAP Note 1031096 - Installing Package SAPHOSTAGENT](#)
- Documentation: [The Documentation attached to the above SAP note](#)
- How to use SAP Host Agent: [How to use the SAP Host Agent](#)

Troubleshooting

- Check the SAP Host Agent log file on the source system. In some cases the process appears to start, however the log shows errors mean that the agent cannot start. In this case re-install the SAP Host Agent.
- Check the network configuration of SLES 11 SP1 as documented in [SAP Note 1310037](#) (search for 127.0.0.2).
- If the source system is AIX, follow SAP Note 1140980 (configure pam.d).
- If the error is "invalid credentials" first check the user password and that the user is not locked on the source system.

5.3 Installing the Sybase Replication Server

The figure below shows how the Sybase Replication Server components are used to enable continuous data replication from the ERP system that is the data provider for the SAP In-Memory Database.



Summary of How the SAP HANA Load Controller Coordinates the SAP HANA Data Replication

- The information modeler starts the SAP HANA Load Controller.
- The Load Controller starts the initial load of the SAP ERP data (source system) to the SAP In-Memory Database on the SAP HANA box (target system). This is done using the SAP R3load component on the SAP ERP system (source system) by exporting the data from the source system and importing the data into the SAP In-Memory Database on the SAP HANA box (target system).
- In parallel, the Sybase Replication Agent on the SAP ERP system is started and performs log mining on the ERP source database. It relays all relevant information to the Replication Server on the SAP HANA box.
- As the main component, the Sybase Replication Server coordinates both the Replication Agent and the ECDA component.
- The Sybase Enterprise Connect Data Access (ECDA) connects to SAP HANA as a target database using an ODBC driver and mirrors the changes from the source database of the ERP system that are relayed through the Replication Server.
- After the initial load has been successfully executed and the changes in the source database during the initial load have been successfully replicated, the Sybase Replication Server and the related components perform the ongoing continuous data replication from the SAP ERP system to the SAP HANA box.

Important Links

- [Sybase Replication Server Reference Manual](#)
- [Introduction to Sybase Replication Command Language](#)
- [Replication Server Commands](#)

- [Sybase Replication Agent Reference Manual](#)
- [Replication Agent Comand Reference](#)

Components to Be Installed

The Sybase Replication software includes three main processes:

1. Sybase **Replication Agent** (rep agent): component does the log mining on the source database and relays all relevant information to the Replication Server
2. Sybase 'Enterprise Connect Data Access' (**ECDA**): component that connects to the target database (SAP In-Memory Database) through ODBC
3. Sybase **Replication Server** (rep server): main component that accepts data from the Replication Agent distributes and applies that data to the target database using ECDA/ODBC for connectivity

Installation

To set up the Sybase-based replication system, perform the following steps:

1. Install the Sybase Replication Server and the ECDA on the SAP HANA box

[Installation and Configuration of the Sybase Replication Server and ECDA](#)

2. Install the Sybase Replication Agent on the SAP ERP system database

[Installation and Configuration of the Sybase Replication Agent](#)

Configuration

1. Configure the components and connect them to the source SAP ERP system and the SAP In-Memory Database target system as detailed in the installation documents.
2. Select which tables to replicate.

For information about configuring data replication, see [SAP HANA 1.0 Modeling Guide](#).

APPENDIX

Default Server/Service names and Port numbers used in TAR images

Component	Default Server Name	Default Service Name	Default Port	Comments
Replication Server	HANARS1	N/A	2121	
Replication Server Embedded RSSD	HA-NARS1_ERSSD	N/A	2122	Internally set to RS port + 1
ECDA	HANADC1	HANADC2SVC	2131	
RepAgent for DB2	HANARAU1	N/A	2141	
RepAgent DB2 Embedded RASD			2142	Internally set to RAU port + 1

Default Login Information Used in TAR Images

Component	Default User	Default Password	Comment
Replication Server	sa	Null	% isql --Usa --P --SHA-NARS1
Replication Server Embedded RSSD	HA-NARS1_RSSD_prim	HANARS1_RSSD_prim_ps	%isql --U<username> -P<password> -Sservername
ECDA	system (IMDB default user)	manager (IMDB default user password)	%isql --Usystem --Pmanager --SHANADC2SVC
RepAgent	sa	Null	%isql --Usa -P

Performing an Initial Test of Replication

If you want to do a manual test of replication before continuing with the SAP HANA Load Controller setup you can perform the following steps:

1. Create a demo table in both the source and target systems: Create Table Test(i int). We will use this table to test replication

2. Mark the table for replication; the table's owner in the target system is explicitly set to system

```

pdb_setreptable Test, system.test,
mark
go

```

3. Start replication

```

resume
go

```

4. Insert a value into your test table in the source database and check whether it is replicated to the target database

5. Verify that the content is replicated into your target database; this might take a minute

6. Stop replication for your test table

```

pdb_setreptable Test, disable
go
pdb_setreptable Test, unmark
go

```

5.3.1 Installation and Configuration of the Sybase Replication Server and ECDA

This task details the steps required to install and configure the Sybase Replication Server and Enterprise Connect Data Access (ECDA).

Installation files

The installation file for the replication components for the SAP HANA system is supplied in a single TAR image. The TAR image contains the Replication Server and ECDA images that are already preconfigured for SAP HANA usage. Apart from unpacking the tar image only a few manual steps are required for installation.

Important: Do not deviate from the default values for Replication Server/Agent names. This could break the setup scripts.

Use the most current SYBASE software version (drop<number>) from the SAP Service Marketplace.

There is one TAR image that accompany this package: - hana_drop<number>_rsdc_linux.tgz (Note: pre-GA version would include drop<number> in the tar filename). This package is for the SAP HANA system.

Installation and configuration sequence

Installation and configuration consists of three distinct stages

1. Decompressing and un-tarring the TAR image
2. Providing configuration settings unique to the host environment
3. Running an installation script to apply the configuration options to the servers

Installation prerequisites

- The TAR images expect that they are installed in the following location `/usr/sap/sybase`
Either create a directory under `/usr/sap/sybase` or a soft link to `/usr/sap/sybase`
Ensure that sufficient space (100+GB) is available for `/usr/sap/sybase`
- The TAR image `hana_drop<number>_rsdc_linux.tgz` must be installed on the SAP HANA system.
This image contains the pre-configured installations for Replication Server (RS) and Enterprise Connect Data Access (ECDA).
- The user installing images must have write permission to `/usr/sap/sybase`
- Prior to installing the TAR images for Sybase Replication for SAP HANA, SAP In-Memory Database and SAP In-Memory Computing studio should be installed.
- Install target replicate database (RDB).
In the case of SAP HANA target database is SAP In-Memory Database.
- Install SAP In-Memory Computing studio on Windows machine.

Ensure that SAP In-Memory Computing studio is able to connect to SAP In-Memory Database replicate database.

Installation of TAR images

On the SAP HANA system (where SAP In-Memory Database runs) perform the following:

1. % cd /usr/sap/sybase
2. % tar xvfz hana_drop<number>_rsdc_linux.tgz

Note:

Some operating systems require the compressed file to be uncompressed separately – the tar command may not perform this operation automatically. Use GNU unzip or similar utility.

3. % chown -R <sid>adm:sapsys /usr/sap/sybase/*

Note:

For the Sybase Replication Server SAP recommends using the same <sid>adm user as is used for SAP In-Memory Database.

Startup and Connectivity establishment

This section deals with basic post installation activities including starting various replication components and testing basic connectivity to individual components.

SAP HANA system

On the SAP HANA system (where SAP In-Memory Database runs) perform the following as <sid>adm:

1. Source /usr/sap/sybase/SYBASE.sh:
% . /usr/sap/sybase/SYBASE.sh
2. Go to the installation directory:
% cd \$SYBASE/postinstall/scripts
3. Make sure you are using the UTF-8 locale; otherwise errors might occur when trying to run the Replication Server:
export LANG=en_US.UTF-8
4. Fill in the important configuration details in postinstallcfg.res file. The following parameters need to be changed: RSHOST, NEWDBHOST, NEWDBNAME, NEWDBINST, NEWDBUSER, RAHOST

Note:

This is a critical step. Attempting skip or doing an incomplete work in this regard would cost more time later on.

RSHOST – The machine name of the HANA system where this Replication Server product is installed. This value may be an IP address of the machine name, or logical host name.

NEWDBHOST – The machine name of the HANA system where SAP In-Memory Database is installed (should be the same as RSHOST)

NEWDBNAME – The name of the SAP In-Memory Database instance.

NEWDBINST – The instance number for the SAP In-Memory Database instance

RAHOST – The machine name where the SAP/ERP system runs and the Replication Agent will be installed. This value may be an IP address of the machine name, or logical host name.

5. Run the configuration script:

```
% . postinstallcfg.sh
```

When prompted Enter the requested passwords.

For "Enter SAP In-Memory Database admins users password" enter the SAP In-Memory Database SYSTEM user password.

For a new installation press 'Enter' for the RSSD admin user password (the default password is blank, it can be changed after the installation).

For an upgrade: if the 'sa' password was changed in the previous version, enter the correct password.

This script will perform the following tasks, among other things:

- Modify the Sybase interface file entries which identify TCP/IP host and port entries
- Modify the ECDA configuration file entries for connectivity to SAP In-Memory Database via ODBC
- Boot ECDA
- Boot Replication Server
- Does a test connection to Replication Server, and ECDA

6. The script reports output to system.out as it processes, and some errors may be reported. Note that these errors occur as expected - the script attempts to stop any prior processes and if the processes are not running (as expected on first install) an error from the shut down request may be reported.

7. At the end of processing, if the install failed, the following message will be in the postinstall log:

```
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Postinstall failed.
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
```

5.3.2 Installation and Configuration of the Sybase Replication Agent

This task describes the installation and configuration of the Sybase Replication Agent that exists and executes on the SAP/ERP system.

Installation files

The installation of the Replication Agent is delivered in a single TAR image. The Replication Agent is already preconfigured for SAP HANA usage. Apart from unpacking the tar image only a few manual steps are required for installation.

Important: Do not deviate from the default values for Replication Server/Agent names. This might break the setup scripts.

Use the most current SYBASE software version (drop<number>) from the SAP Service Marketplace; for details see Software download.

There is one TAR image that accompany this package: - [hana_drop<number>_rax_<platform>.tgz (Note: pre-GA version would include drop<number> in the tar filename)

The <platform> represents the operating system the Replication Agent supports. This is meant for running on the ERP database system.

Installation and configuration sequence

Installation and configuration consists of three distinct stages

1. Decompressing and un-tarring the TAR image
2. Providing configuration settings unique to the host environment
3. Running an installation script to apply the configuration options to the servers

Installation prerequisites

- The TAR images expect that they are installed in the following location /usr/sap/sybase
Either create a directory under /usr/sap/sybase or a soft link to /usr/sap/sybase
Ensure that sufficient space (5+GB) is available for usr/sap/sybase
- The TAR image hana_drop<number>_rax_<platform>.tar is a platform specific one, and it should be installed on the machine where SAP ERP database runs.
This image contains the pre-configured installation for Replication Agent
- The user installing images must have write permission to /usr/sap/sybase
- Installed source primary database (PDB), in the case of SAP HANA source database is DB2 UDB.
- Successfully Installation of Replication Server on the SAP HANA system ([Installation and Configuration of the Sybase Replication Server and ECDA](#)).

Installation of TAR images

On the machine where SAP ERP database runs, perform the following:

1. % cd /usr/sap/sybase
2. % tar xvfz hana_drop<number>_rax_<platform>

Note:

Some operating systems require the compressed file to be uncompressed separately – the tar command may not perform this operation automatically. Use GNU unzip or similar utility.

3. % chown -R <sid>adm:sapsys /usr/sap/sybase/*

Note:

For the Sybase Replication Agent SAP recommends to use the same <sid>adm user as is used for the ERP source system.

Startup and Connectivity establishment

This section deals with basic post installation activities including starting various replication components and testing basic connectivity to individual components.

SAP ERP database system

After installing the tar image of the Replication Agent binary on the ERP database system, please do the following post install activity. On the machine where SAP ERP database runs, perform the following as user <sid>adm of the source system:

1. Source /usr/sap/sybase/SYBASE.sh:

```
% . /usr/sap/sybase/SYBASE.sh
```

2. Go to the postinstall directory:

```
% cd $SYBASE/postinstall
```

3. Prepare your DB2 system for replication by setting up a DB2 as primary data source:

Note: in a customer installation scenario the DB2 steps outlined below can be performed by the DB2 DBA's well in advance of the actual SAP HANA installation. This can save time during the SAP HANA installation activity.

A. Creating a DB2 user for Replication Agent

The Replication Agent will use this user to communicate with DB2. The user needs to have the privileges of the DBADM group in order to set "DATA CAPTURE CHANGES" attribute on tables to be replicated.

- a. Create user repagent on operating system level (Note: do NOT use upper case username, DB2 login will fail).
- b. Log in as db2 administrator user, e.g. for system D3D as user DB2D3D.
- c. Start db2 command prompt:

```
lsi155:db2d3d 5> db2
```

- d. Create dbadm privileges to your new db2 user, repagent => grant dbadm on database to repagent.
- e. Enable archive log writing (it should already be enabled on a productive system)

```
connect
```

```
UPDATE DATABASE CONFIGURATION using LOGARCHMETH1 LOGRETAIN
```

- f. You may have to restart the database for this setting to take effect:

```
stop force
```

```
start
```

- g. You may also need to perform a backup before you can connect to the database again

```
backup db d3d
```

- h. If you are on a test system and don't really need the backup, you can backup to {{/dev/null}} to speed up the process (depending on the database size this can run for a while)

```
backup db d3d to /dev/null
```

- i. You can monitor the progress of your backup in another db2 session with the list utilities show detail command

B. Determining the port number of the DB2 instance

- a. The port number is dependent on your SAP system name and can be found in /etc/services, e.g. to determine the port of your DB2 database for SAP instance D3D:

```
db2d3d@lsi155:~> grep sapdb2 /etc/services sapdb2D3D 5912/tcp # SAP DB2 Communication Port
```

4. Edit /usr/sap/sybase/postinstall/HANAUDBSourceEnv.txt. Important parameters to change are:

udb_instance_home	The home directory for DB2
pds_host_name	The machine name where the SAP/ERP system runs and the Replication Agent is being installed. This value may be an IP address of the machine name, or logical host name.
pds_port_number	The client port number for DB2, as determined under B. Determining the port number of the DB2 instance above.
pds_username	The DB2 user created under A. Creating a DB2 user for Replication Agent, above. The user name from the example is repagent. Typically it is not the ERP database user name.
pdb_archive_path	The path to the UDB archive log files. Archive files will be read from this directory path and no other. Set pdb_archive_path to point to the location specified by either LOGARCHMETH1 or LOGARCHMETH2.
hana_host_name	The machine name of the SAP HANA system where the Sybase Replication Server and SAP In-Memory Database are installed. This value may be an IP address of the machine name, or logical host name.

This is an important step. Please take time to ensure that information you provide in the configuration file is correct. Time spent in ensuring that the information provided here is correct, will save significant time in preventing setup and configuration issues.

5. Run /usr/sap/sybase/postinstall/hana_udb_validate_env.sh. This step will validate basic settings and file access permissions to prevent avoidable errors. Only move to the next step if you have a successful validation return from this script .
6. Run /usr/sap/sybase/postinstall/hana_udb_repagent_setup.sh. You will be prompted to supply the password for the 'repagent' user. Hit <return> to accept the default 'null' password. This script will perform the following tasks, among other things:
 - Start the Replication Agent
 - Configure the Replication Agent with host and port information for connecting to the SAP/ERP/DB2 system as well as the Replication Server on the SAP HANA box
 - Test connectivity to both SAP/ERP/DB2 system as well as the Replication Server on the SAP HANA box
 - Initialize the Replication Agent to read the DB2 transaction log

- Set the Replication Agent into “replicating” state.
7. At the end of processing, the `hana_udb_repagent_setup.sh` script will report success or failure of installation. If there is any failure, the script may be re-run once a problem has been determined and corrected.