



**SAP**<sup>®</sup> Partner  
Open Ecosystem

# ***Documentation***

---

## ***IBM Workload Scheduler integration with SAP HANA Database***

<b>Written by :</b> Miguel Sanders Uniforce	<b>Date :</b> June 30 2017
---	-------------------------------



## Table of Contents

<b>1. INTRODUCTION.....</b>	<b>4</b>
<b>2. INSTALLING AND CONFIGURING THE PLUG-IN FOR SAP HANA DATABASE....</b>	<b>5</b>
<b>3. DEFINING A SAP HANA DATABASE JOB .....</b>	<b>6</b>
<b>4. MONITORING A SAP HANA DATABASE JOB .....</b>	<b>9</b>
<b>5. APPENDIX A : USING THE SAP CLOUD CONNECTOR TO CONNECT TO A SAP HANA DATABASE .....</b>	<b>10</b>



## CHANGE HISTORY

Version	Date of change	Change detail
1.0	June 30 2017	Initial version by M. Sanders

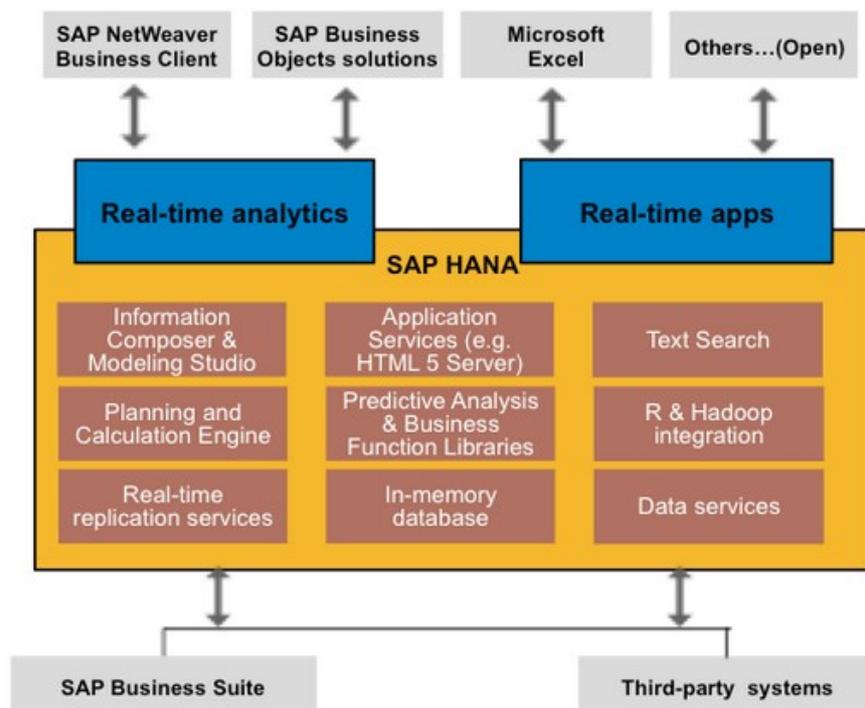
## 1. INTRODUCTION

SAP HANA is a column-oriented, in-memory relational DBMS. HANA is architected to enable applications to support both transactional and analytical processing on a single system with one copy of the data.

The DBMS, which runs on SUSE Linux and Red Hat Enterprise Linux, enables real-time analytics on transactional systems on a large scale and on a variety of data, including structured, unstructured, spatial, time series and streaming data. It provides features that support development for SAP and custom-built applications. SAP HANA combines database, advanced analytics (predictive, spatial, text analytics, sentiment analysis, search), enterprise information management (bulk load, real-time replication, transformation, cleansing) and application server capabilities all running in-memory, on one data copy and on a single platform.

SAP HANA supports multi-tenancy and data tiering, which enables petabyte-scale deployments for warm data (data that's less frequently accessed) to be stored on the disk, and offers a choice of deployment models and partners. The DBMS can be deployed on-premises, in the SAP Cloud Platform or as a hybrid of both.

By leveraging the IBM Workload Scheduler plug-in for SAP HANA Database, you will be able to schedule these mission critical jobs, both on-premise and on the SAP Cloud Platform, from a single interface .





## 2. INSTALLING AND CONFIGURING THE PLUG-IN FOR SAP HANA DATABASE

The following prerequisites must be met in order to use the IBM Workload Scheduler plug-in for SAP HANA Database.

- IBM Workload Scheduler 9.1 or later
- IBM Dynamic Workload Console 9.1 or later

To install and configure the IBM Workload Scheduler plug-in for SAP HANA Database, perform the following steps:

- Copy *com.ibm.scheduling.agent.saphdb\_<version>.jar* to the *<TWA\_HOME>/TWS/applicationJobPlugIn* folder on either the Master Domain Manager or a Dynamic Domain Manager
- Copy *com.ibm.scheduling.agent.saphdb\_<version>.jar* to the *<TWA\_HOME>/TWS/JavaExt/eclipse/plugins* folder on the Dynamic Agent that will run the SAP Data Services jobs
- Modify *config.ini* located in *<TWA\_HOME>/TWS/JavaExt/eclipse/configuration* on the Dynamic Agent that will run the SAP Data Services jobs. At the end of the line that starts with "osgi.bundles=", add the following:  
"*,com.ibm.scheduling.agent.saphdb@4:start*"
- Restart the WebSphere Application Server of either the Master Domain Manager or the Dynamic Domain Manager.
- Restart the Dynamic Agent
- Restart the Dynamic Workload Console

### 3. DEFINING A SAP HANA DATABASE JOB

From the Dynamic Workload Console, you can define a SAP HANA Database job as follows:

- In the Dynamic Workload Console navigation tree, expand *Administration* and select *Manage Workload Definitions*.
- Specify the name of the engine. Subsequently, the Workload Designer is displayed.
- In the Working List panel, select *New -> Job Definition -> Database and Integrations -> SAP HANA Database*.
- On the *SAP HANA Database* panel, fill in the job details
  - **Hostname** : Hostname of the SAP HANA Database or the SAP Cloud Connector if you want to connect to a SAP HANA Database hosted on the SAP Cloud Platform. Please refer to Appendix A for more information.
  - **Port** : Port number of the SAP HANA Database. The port number is 3<instance number>15. Please refer to SAP note 1592925 for more details.
  - **Database name** : The name of the database. SAP Multi-tenant Database Containers (MDC) are also supported. You can use the “Lookup...” button to retrieve the database names.
  - **Schema name** : The name of the schema. You can use the “Lookup...” button to retrieve the schema names.
  - **JDBC connection string** : Alternatively, you can enter the JDBC connection string.
  - **Username** : The username to log on to the SAP HANA Database.
  - **Password** : The password of the user that will be used to log on to the SAP HANA Database.
  - To validate the connection, click *Test Connection*.
- On the *SQL* panel, fill in the SQL details
  - **Standard SQL** : Fill in the DDL/DML statement.
  - **Stored Procedure** : Fill in the name of the stored procedure. You can use the “Lookup...” button to retrieve the stored procedure names. If the stored procedure requires input parameters, you can add them in the “Procedure Parameters” panel.



SAP HANA DATABASE (9.4.0.00) - DELTA#P\_HDB\_DEMO

Select an Action



General

Affinity

Recovery Options

SAP HANA Database

Connection Details

Predefined

\* Hostname

\* Port

Database name

Schema name

Custom

JDBC connection string

Credentials

\* Username

\* Password  ...

SAP HANA DATABASE (9.4.0.00) - DELTA#P\_HDB\_DEMO

Select an Action



General

Affinity

Recovery Options

SAP HANA Database

SQL

Versions

Standard SQL

Stored Procedure

\* Procedure Name

Filter

Name	Value
DISCOUNT	100



Alternatively, the *composer* command line can be used to define the job. Example:

```
DELTA#P_HDB_DEMO
TASK
<?xml version="1.0" encoding="UTF-8"?>
<jSDL:jobDefinition xmlns:jSDL="http://www.ibm.com/xmlns/prod/scheduling/1.0/jSDL"
xmlns:jSDLSAPHDB="http://www.ibm.com/xmlns/prod/scheduling/1.0/jSDLSAPHDB" name="SAPHDB">
<jSDL:application name="saphdb">
  <jSDLSAPHDB:saphdb>
    <jSDLSAPHDB:SAPHDBParameters>
      <jSDLSAPHDB:SAPHDBSQLParms>
        <jSDLSAPHDB:standardSQLOrStoredProcedure>
          <jSDLSAPHDB:storedProcedureGroup>
            <jSDLSAPHDB:storedProcedureName>SALES_REPORT</jSDLSAPHDB:storedProcedureName>
            <jSDLSAPHDB:storedProcedureParametersValues>
              <jSDLSAPHDB:storedProcedureParametersValue
                key="DISCOUNT">100</jSDLSAPHDB:storedProcedureParametersValue>
            </jSDLSAPHDB:storedProcedureParametersValues>
          </jSDLSAPHDB:storedProcedureGroup>
        </jSDLSAPHDB:standardSQLOrStoredProcedure>
      </jSDLSAPHDB:SAPHDBSQLParms>
      <jSDLSAPHDB:SAPHDBParms>
        <jSDLSAPHDB:connectionDetails>
          <jSDLSAPHDB:predefinedOrCustom>
            <jSDLSAPHDB:predefinedDatabase>
              <jSDLSAPHDB:hostname>sapcloudconnector</jSDLSAPHDB:hostname>
              <jSDLSAPHDB:port>30015</jSDLSAPHDB:port>
              <jSDLSAPHDB:databaseName>JSZ</jSDLSAPHDB:databaseName>
              <jSDLSAPHDB:schemaName>NEO_74WTJMFNT6QS0VXWLGWU9ADM8</jSDLSAPHDB:schemaName>
            </jSDLSAPHDB:predefinedDatabase>
          </jSDLSAPHDB:predefinedOrCustom>
        </jSDLSAPHDB:connectionDetails>
        <jSDLSAPHDB:creds>
          <jSDLSAPHDB:username>DEV_9EDGZ53LHH6VYFN435TLRX6DN</jSDLSAPHDB:username>
          <jSDLSAPHDB:password>${agent:session_password}</jSDLSAPHDB:password>
        </jSDLSAPHDB:creds>
      </jSDLSAPHDB:SAPHDBParms>
    </jSDLSAPHDB:SAPHDBParameters>
  </jSDLSAPHDB:saphdb>
</jSDL:application>
</jSDL:jobDefinition>
DESCRIPTION "Added by composer."
RECOVERY STOP
```



## 4. MONITORING A SAP HANA DATABASE JOB

Like regular jobs, you can monitor SAP HANA Database jobs by using either the Dynamic Workload Console or the *conman* command line.

Plan Name: Current Plan

@#@@HDB@DEMO@

Status	Internal Status	Job	Job Type	Workstation (Job)	Job Stream	Workstation (Job Stream)
<input checked="" type="checkbox"/>	Successful	SUCC P_HDB_DEMO	SAP HANA Database	DELTA	P_HDB_DEMO	MDM

```

= Job Number: 876296181
= Fri 06/30/2017 17:27:45 CEST
=====
SUM(_SYS_SS_VAR_VAR2.SALES_AMOUNT), SUM(_SYS_SS_VAR_VAR2.SALES_AMOUNT - _SYS_SS_VAR_VAR2.SALES_AMOUNT * TO_DOUBLE(?1) / 100), PRODUCT_NAME, REGION_NAME, SUB_REGION_NAME
100.0, 0.0, Shirts, Americas, North-America
90.0, 0.0, Jackets, Americas, North-America
85.0, 0.0, Purse, Americas, North-America
80.0, 0.0, Jackets, Americas, South-America
75.0, 0.0, Shirts, Americas, South-America
85.0, 0.0, Trousers, Asia, India
75.0, 0.0, Coats, Asia, Japan
65.0, 0.0, Shirts, Europe, Germany
65.0, 0.0, Jackets, Europe, Germany
=====
= Exit Status : 0
= Elapsed Time (hh:mm:ss) : 00:00:01
= Fri 06/30/2017 17:27:45 CEST
=====

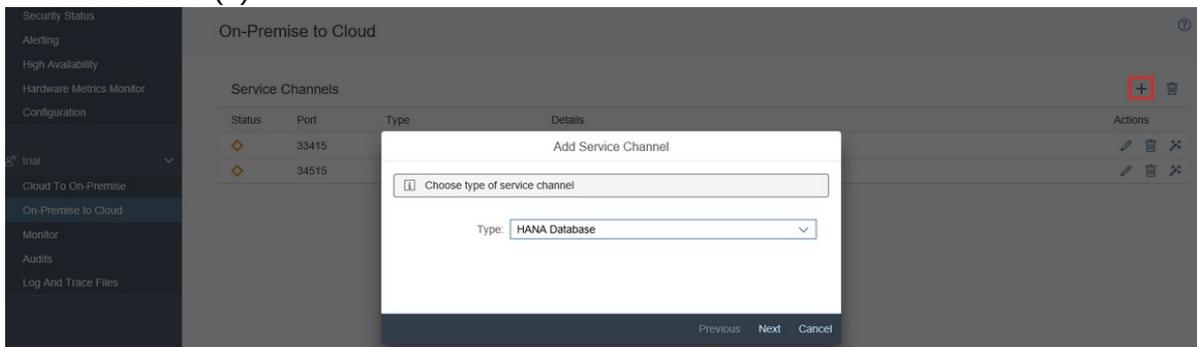
```

## 5. APPENDIX A : USING THE SAP CLOUD CONNECTOR TO CONNECT TO A SAP HANA DATABASE

Using service channels of the SAP Cloud Connector, you can connect to a HANA database on the SAP Cloud Platform. You can do this in section «On-Premise to Cloud» > «Service Channels» of the SAP Cloud Connector.

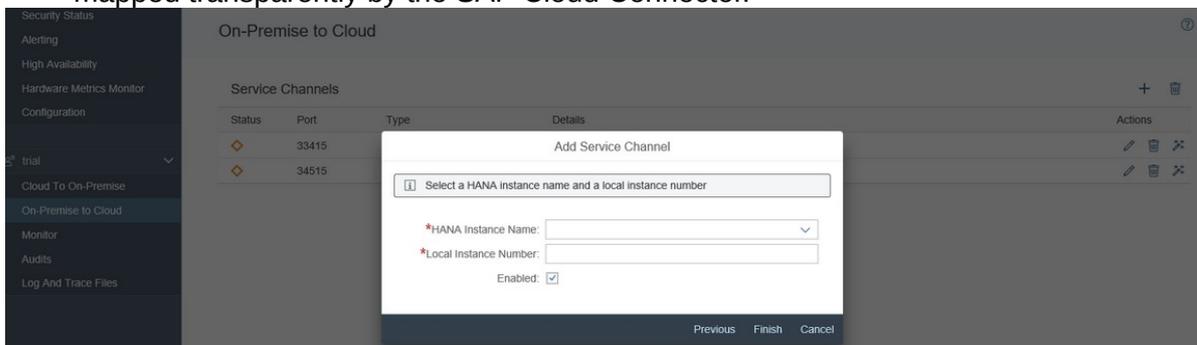
Follow the steps below to establish a Service Channel to a HANA instance of your subaccount.

- Logon to the SAP Cloud Connector.
- Choose *On-Premise To Cloud* from your subaccount menu.
- Choose Add (+)



- In the Add Service Channel dialog, leave the default value *HANA Database* in the <Type> field.
- Choose Next.
- Choose the HANA instance name from the drop down list of available HANA instances. If fetching the list failed, you need to specify the name yourself. It must match one of the names shown under SAP HANA/SAP ASE Databases & Schemas in the cockpit.
- Specify the local instance number. This is a double-digit number which computes the local port used to access the HANA instance in the cloud. The local port is derived from the local instance number as 3<instance number>15. For example, if the instance number is 22, then the local port will be 32215.

Note :The local port should not match the HANA port used in the cloud – they are mapped transparently by the SAP Cloud Connector.



- Leave the Enabled option selected to establish the channel immediately after clicking Finish, or deselect it if the channel should not be established yet.
- When you are done, choose Finish.