


# SAP Data Hub – Kubernetes

## Validation Guideline

SAP / Partner CONFIDENTIAL



Date	Alteration Reason	Version
07/13/2018	Initial draft	1.0
11/08/2018	Public Version for Partners	2.0
04/10/2019	Validation Checks extended	2.1

# TABLE OF CONTENTS

INTRODUCTION .....	3
Purpose of Kubernetes (K8s) validation for SAP Data Hub .....	3
VALIDATION PROCESS DESCRIPTION .....	4
Goal of Validation .....	4
Prerequisites .....	4
Key Deliverables .....	4
Roles and Responsibilities .....	4
<i>SAP responsibilities</i> .....	4
<i>Partner responsibilities</i> .....	5
Validation Outcome .....	5
Out of Scope .....	5
SAP DATA HUB TEST SETUP & CONDITIONS .....	6
General Approach.....	6
Scenarios.....	6
a) <i>SAP Data Hub Installer - ./install.sh --validate</i> .....	7
b) <i>Log into SAP Data Hub Launchpad</i> .....	7
c) <i>Execute a sample SAP Data Hub Data Pipeline</i> .....	8
d) <i>Check the SAP Data Hub Pipeline Modeler registry configuration</i> .....	8
e) <i>Create a connection to local SAP Vora in Connection Manager</i> .....	9
f) <i>Create an SAP Vora Streaming table in SAP Vora Tools</i> .....	10
g) <i>View/ Profile an SAP Vora table in Metadata Explorer</i> .....	11
SAP/Partner Support Guideline .....	12

[www.sap.com/contactsap](http://www.sap.com/contactsap)

© 2018 SAP SE or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company.

The information contained herein may be changed without prior notice. Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP or SAP affiliate company 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.

In particular, SAP SE or its affiliated companies have no obligation to pursue any course of business outlined in this document or any related presentation, or to develop or release any functionality mentioned therein. This document, or any related presentation, and SAP SE's or its affiliated companies' strategy and possible future developments, products, and/or platform directions and functionality are all subject to change and may be changed by SAP SE or its affiliated companies at any time for any reason without notice. The information in this document is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. All forward-looking statements are subject to various risks and uncertainties that could cause actual results to differ materially from expectations. Readers are cautioned not to place undue reliance on these forward-looking statements, and they should not be relied upon in making purchasing decisions.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies. See <http://www.sap.com/corporate-en/legal/copyright/index.epx> for additional trademark information and notices.



## INTRODUCTION

### **Purpose of Kubernetes (K8s) validation for SAP Data Hub**

The use of Kubernetes (K8s) is required to run SAP Data Hub. Kubernetes is an open source system for managing containerized applications across multiple hosts, providing basic mechanisms for deployment, maintenance, and scaling of applications. Originally developed by Google to manage very large applications clusters, Kubernetes (K8s) is adopted by many software vendors, who modify and enhance the original code set for their specific purposes and functions. As a result, several K8s “flavors” exist. Just as UNIX operating systems vary in implementation and kernel details, so do K8s versions from different software vendors. Only validated and tested K8s versions are approved to run productive SAP Data Hub scenarios. All supported Kubernetes versions are listed in the corresponding SAP note.

This document describes the validation process a partner has to follow to certify their K8s platform/version for SAP Data Hub compatibility.

The validation process covers the following aspects:

- Ensuring functional compatibility of the K8s software for SAP Data Hub
- Ensuring seamless supportability of SAP Data Hub with the corresponding K8s software
- Ensuring distribution and support model in sync with SAP Data Hub delivery model
- Ensuring legal compliance (e.g. regarding open source and 3<sup>rd</sup> party) and prevention of SAP license misuse

The validation process (sometimes also referred to as “certification”) requires a highly cooperative and goal-oriented communication between the corresponding Kubernetes vendor and SAP. Successful completion of the validation process will conclude in documenting the newly certified Kubernetes version in the corresponding SAP note. As a result of this validation, all constraints and all best practices for the corresponding technology shall be documented by the vendor. Updated versions of certified Kubernetes software require separate certification. This means, each new version must be certified.

This document shall formalize the major aspects of this cooperation between the partners and SAP, provides rules and guidelines that have to be followed by all participants in the project in order to ensure successful validation. It is a living document and will be maintained throughout its lifetime and should be available to all project participants for reference.

## VALIDATION PROCESS DESCRIPTION

### Goal of Validation

The aim of the SAP Data Hub Kubernetes validation is to test and publish a customer facing support statement for SAP Data Hub running on and being supported for K8s platforms, for use in production.

### Prerequisites

- A reasonable SAP business case which justifies the investment of human resources and infrastructure, required for validation and later support of the corresponding Kubernetes platform.
- List of minimum five customers who plan to run their SAP Data Hub systems on corresponding Kubernetes platform within their productive landscapes shall be identified to fortify the business case.
- Allocation of permanent and skilled resources, on-site SAP Headquarters in Walldorf/St. Leon Rot, to work with SAP Data Hub development on analyzing functional behavior during validation, resolving issues occurred during testing, and related customer support issues.
- Provisioning of reference systems (bare metal or virtualized).
- Existing SAP Technology Partner Status & DB/OS Platform certification.

### Key Deliverables

SAP expects the partner to provide / commit to the following deliverables as output of the validation activities:

- Publication of a “Kubernetes for SAP Data Hub” specific and customer facing best practice document, describing the exact configuration and learnings derived from validation activity.
- Availability of APIs to satisfy requirements towards prevention of SAP license misuse (e.g. SAP reporting tools validated with the corresponding Kubernetes version)
- Agreement that any non-SAP related issue will be routed to Kubernetes provider / vendor before it potentially gets passed on to the underlying cluster hardware vendor.

After successful validation, SAP will further update the following documents and resources to state official support for corresponding platform to its customers:

- Update of central [SAP Note 2464722](#) (SAP Data Hub 1.0) and [SAP Note 2686169](#) (SAP Data Hub 2), summarizing supported Kubernetes platforms.
- Update of SAP Data Hub documentation, if and where necessary.

### Roles and Responsibilities

Neither SAP nor the Kubernetes vendor alone can execute the validation activities on his or her own. A close collaboration is required between the two parties to successfully validate SAP Data Hub with the new Kubernetes software.

#### **SAP responsibilities**

- Held status meeting on a regular (e.g. bi-week schedule), to keep all stakeholders informed about the current status of validation activities and, if necessary, to take further actions to drive project success.
- Review functional test results.
- Provide validation parameters



### ***Partner responsibilities***

- Provide, setup and maintain test landscape (hardware & software)
- Execution and analysis of SAP Data Hub installation and of functional tests
- Make proposals on how to improve e.g. performance and apply necessary configuration
- Document any gained best practice for later publishing in corresponding best practice guide for SAP Data Hub on a particular platform
- Self-organized knowledge exchange with SAP Data Hub engineering
- Make sure to consult [SAP Note 2464722](#) (SAP Data Hub 1.0) and [SAP Note 2686169](#) (SAP Data Hub 2)  
- Prerequisites for installing SAP Data Hub

### **Validation Outcome**

While an unconstrained support statement through SAP is most desirable, the following levels of limited support statements may be considered during decision gates:

- Unconditioned support statement for production use cases
- Limited support statement to e.g. non-production workloads

### **Out of Scope**

The following constraints are of general nature and will remain even after successful completion of planned activities:

- Docker infrastructure and versions



## SAP DATA HUB TEST SETUP & CONDITIONS

### General Approach

The SAP Data Hub technology partner agrees to provide a reference landscape with bare metal or virtualized SAP Data Hub certified K8s systems to analyze and solve customer problems. In parallel, the SAP Data Hub technology partner is asked to take suitable measures to gain necessary knowledge how existing configuration of the K8s platform can be further improved over the course of test repetition. Outcome of this testing should be the publication of a K8s platform specific and customer facing best practice document, describing the exact configuration and learnings derived from validation activity.

### Scenarios

The following scenarios will be carried out and given KPIs need to be met.

- **Functional tests**
  - Verify e.g. log in to SAP Data Hub Launchpad or define a Data Pipeline

Once passed, the exact configuration, under which KPIs had been achieved, shall be documented in a customer facing best practice guide.

The K8s vendor and SAP may agree on a preliminary release statement, limited for non-productive use, if all other release conditions are fulfilled, and only certain SAP Data Hub validation checks (e.g. checkpoint store) are still pending.

### a) SAP Data Hub Installer - `./install.sh --validate`

After the successful installation of SAP Data Hub, the validate option of the install.sh script can be used to validate the installation.

`./install.sh --validate` provides the following checks:

- Verify if Vora tables can be created in related engines (data source, streaming table)
  - Execute SQL queries (create, insert, select, drop) against the SAP Vora disc engine
  - Execute SQL queries (create, insert, select, drop) against the SAP Vora relational engine
- Verify if Spark runs on Kubernetes cluster
- Verify if system management application is installed
- Checks that all diagnostics-\* pods are running
- Checks that a dummy HTTP GET request can be made to the auditlog pod
- Verify if internal SAP HANA instance has been initialized for SAP Data Hub and that user exists
- Runs smoke test inside the vora-textanalysis pod

<b>Description:</b>	E.g.: The V2Client executes simple SQL queries against SAP Vora transaction coordinator (for disc engine and relational engine).
<b>Scenarios:</b>	Mandatory for <input checked="" type="checkbox"/> functional tests
<b>Test execution:</b>	Via <code>./install.sh --validate</code> on the command line
<b>KPI:</b>	E.g.: Successful execution of SQL queries (create, insert, select, drop) against the SAP Vora disc engine and the SAP Vora relational engine.
<b>Output:</b>	KPI passed / not passed
<b>Documentation:</b>	<a href="https://help.sap.com/viewer/p/SAP_DATA_HUB">https://help.sap.com/viewer/p/SAP_DATA_HUB</a> -> Installation Guide -> References

### b) Log into SAP Data Hub Launchpad

After the installation of all components of SAP Data Hub, test the installation to ensure that it is working correctly. In this test scenario you log into the SAP Data Hub Launchpad which is a browser-based application that provides a single point of access to a range of user-facing applications such as SAP Vora Tools, SAP Data Hub Pipeline Modeler, SAP Data Hub Connection Management, System Management, and more. The applications that you can access depend on the user credentials of the active user.

<b>Description:</b>	In this test scenario you log into the SAP Data Hub Launchpad which is the entry point for all SAP Data Hub applications.
<b>Scenarios:</b>	Mandatory for <input checked="" type="checkbox"/> functional tests
<b>Test execution:</b>	Via launching of the SAP Data Hub Launchpad UI in a browser -> <code>http://&lt;node-ip&gt;:&lt;port&gt;</code> or <code>https://&lt;node-ip&gt;:&lt;port&gt;</code> and logging into it.
<b>KPI:</b>	Successful logon and display of the initial screen including tiles which can be opened.
<b>Output:</b>	KPI passed / not passed
<b>Documentation:</b>	<a href="https://help.sap.com/viewer/e66c399612e84a83a8abe97c0eeb443a/2.4.latest/en-US/8e12370aa28943bdb0952ad9431d2c2c.html">https://help.sap.com/viewer/e66c399612e84a83a8abe97c0eeb443a/2.4.latest/en-US/8e12370aa28943bdb0952ad9431d2c2c.html</a>

### c) Execute a sample SAP Data Hub Data Pipeline

After the installation of all components of SAP Data Hub, test the installation to ensure that it is working correctly. In this test scenario you use the SAP Data Hub Pipeline Modeler and execute a sample data pipeline.

<b>Description:</b>	In this test scenario you use the SAP Data Hub Pipeline Modeler and execute a sample data pipeline.
<b>Scenarios:</b>	Mandatory for <input checked="" type="checkbox"/> functional tests
<b>Test execution:</b>	Launch the SAP Data Hub Pipeline Modeler running as part of the Kubernetes deployment. Search and open the sample graph "com.sap.demo.datagenerator" and start/run it.
<b>KPI:</b>	Check that the status of the sample data pipeline changes to "Running". Open the Terminal user interface by right-clicking on the Terminal operator and check if the generated data is displayed.
<b>Output:</b>	KPI passed / not passed
<b>Documentation:</b>	<a href="https://help.sap.com/viewer/e66c399612e84a83a8abe97c0eeb443a/2.4.latest/en-US/c910ec74f2684688a91ebc44f3f1aa2f.html">https://help.sap.com/viewer/e66c399612e84a83a8abe97c0eeb443a/2.4.latest/en-US/c910ec74f2684688a91ebc44f3f1aa2f.html</a>

### d) Check the SAP Data Hub Pipeline Modeler registry configuration

After the installation of all components of SAP Data Hub, test the installation to ensure that it is working correctly. In this test scenario you check the pipeline modeler registry configuration:

- A Pipeline Modeler registry has been specified in System Management (-> Application Config)
  - Typically, this is the same as the registry as where the Data Hub images are stored
- Docker registry login credentials provided as a Kubernetes secret
  - See installation documentation for platform-specific instructions

<b>Description:</b>	In this test scenario you check the pipeline modeler registry configuration.
<b>Scenarios:</b>	Mandatory for <input checked="" type="checkbox"/> functional tests
<b>Test execution:</b>	<ul style="list-style-type: none"><li>• Click on the SAP Data Hub System Management tile</li><li>• Click on the Application configuration button</li><li>• Check the value of parameter 'Modeler: Docker registry for modeler images', e.g. myDataHubRegistry.azurecr.io</li><li>• Check the value of parameter 'Modeler: Name of the vSystem secret...', e.g. vlow-registry</li></ul>
<b>KPI:</b>	Successful configuration of the pipeline modeler registry.
<b>Output:</b>	KPI passed / not passed
<b>Documentation:</b>	n./a.



### e) Create a connection to local SAP Vora in Connection Manager

After the installation of all components of SAP Data Hub, test the installation to ensure that it is working correctly. In this test scenario you create a connection to the local SAP Vora in the connection manager. Please note, when running the validation for the first time a flowagent pod in Kubernetes will be deployed. The Flowagent is responsible for validating the parameters and the connectivity. Any validation errors during the validation will be logged in the corresponding flowagent pod.

<b>Description:</b>	In this test scenario you create a connection to the local SAP Vora via connection manager.
<b>Scenarios:</b>	Mandatory for <input checked="" type="checkbox"/> functional tests
<b>Test execution:</b>	<ul style="list-style-type: none"><li>• Click on the Connection Management tile</li><li>• Click on the create button</li><li>• Fill in following fields: id (VORA_LOCALHOST), Connection type (Vora), Host (vora-tx-coordinator), Port (10002), Tenant: (default), Use TLS (false), User ("yourUser"), Password ("yourPassword")</li><li>• Click on the Save Button</li><li>• Click on the Action button</li><li>• Click on the Check Status Button</li><li>• The opened window should display, that the connection is okay</li></ul>
<b>KPI:</b>	Successful creation of the connection to a local SAP Vora via the Connection Manager and successful status check.
<b>Output:</b>	KPI passed / not passed
<b>Documentation:</b>	<a href="https://help.sap.com/viewer/1246f58c6a74412580877af1f484f41a/2.4.latest/en-US/e259041c90734cb688e13a7931e7d721.html">https://help.sap.com/viewer/1246f58c6a74412580877af1f484f41a/2.4.latest/en-US/e259041c90734cb688e13a7931e7d721.html</a>

### f) Create an SAP Vora Streaming table in SAP Vora Tools

After the installation of all components of SAP Data Hub, test the installation to ensure that it is working correctly. In this test scenario you create a Vora Streaming table using Vora Tools.

<b>Description:</b>	In this test scenario you create a Vora Streaming table using Vora Tools.
<b>Scenarios:</b>	Mandatory for <input checked="" type="checkbox"/> functional tests
<b>Test execution:</b>	<ul style="list-style-type: none"><li>• Click on the Vora Tools tile</li><li>• Click on the SQL-Editor button</li><li>• Write the following code lines: <pre>CREATE TABLE helloWorld (a1 varchar (*), a2 int, a3 double) TYPE STREAMING STORE ON DISK; INSERT INTO helloWorld VALUES ('Hello', 2, 3.12); INSERT INTO helloWorld VALUES ('World', 3, 4.13); SELECT a1, a2, a3 FROM helloWorld;</pre></li><li>• Execute the statement via clicking on the run button or with the shortcut CTRL+Shift+R</li><li>• If everything was successful, the result set shows up and the execution Log doesn't have any errors</li></ul>
<b>KPI:</b>	Successful creation of a Vora Streaming table in Vora tools
<b>Output:</b>	KPI passed / not passed
<b>Documentation:</b>	<a href="https://help.sap.com/viewer/98fe4b43ad664cba874e0fc7ea8ecc76/2.4.latest/en-US">https://help.sap.com/viewer/98fe4b43ad664cba874e0fc7ea8ecc76/2.4.latest/en-US</a>

### g) View/ Profile an SAP Vora table in Metadata Explorer

After the installation of all components of SAP Data Hub, test the installation to ensure that it is working correctly. In this test scenario you View/Profile a Vora table in the Metadata Explorer. Consider when executing a profile job for the first time it can take up to 15 minutes due to building, pushing, pulling of required docker images.

<b>Description:</b>	In this test scenario you view/profile a Vora table in the Metadata Explorer.
<b>Scenarios:</b>	Mandatory for <input checked="" type="checkbox"/> functional tests
<b>Test execution:</b>	<ul style="list-style-type: none"><li>• Click on the Metadata Explorer tile</li><li>• After accessing the Metadata Explorer, click on the Browse Connection Button</li><li>• Search for the VORA_LOCALHOST Connection</li><li>• Click on the VORA_LOCALHOST Connection</li><li>• Browse for the schema, in which you stored your Vora table and click on it</li><li>• Select the helloWorld table and click on the More Action button</li><li>• Click on the Start Profiling Button to start the profiling of the table</li><li>• You can follow the progress of the profiling in the Monitoring component (Go back to the Metadata Explorer dashboard and click on the Monitoring quick link)</li><li>• After the profiling is finished, you can go back to the Browse Connections Menu and look at the fact sheet of the table</li></ul>
<b>KPI:</b>	After profiling a Vora table a fact sheet gets successfully created in Metadata Explorer
<b>Output:</b>	KPI passed / not passed
<b>Documentation:</b>	<a href="https://help.sap.com/viewer/bd9aff2447b24109bccbd7618cfd542e/2.4.latest/en-US/670607fd56914f8d9d673dc92475bebb.html">https://help.sap.com/viewer/bd9aff2447b24109bccbd7618cfd542e/2.4.latest/en-US/670607fd56914f8d9d673dc92475bebb.html</a>



## SAP/Partner Support Guideline

Please refer to [SAP Note 2698100](#) for details.

The following SAP Data Hub specific paragraph shall be added to the SAP/Partner Support Guide to ensure proper handling of potential support issues between involved parties.

### **Addendum Text**

If a customer reports an issue **while running SAP Data Hub on <partner><solution>**, customer opens an SAP support ticket via Service Marketplace. SAP Support will then route the ticket as follows:

1. If a problem is a known SAP Data Hub issue with a validated fix - SAP Support will recommend the appropriate fix to the customer. SAP Support will close the ticket if the problem disappears after the customer applies the recommended fix. However, if the problem still persists and the recommended fix does not work for SAP Data Hub, the ticket will be forwarded to <partner> for support
2. <partner> now owns the issue and, as needed, <partner> will involve SAP to troubleshoot the issue. If SAP and <partner> agree that the issue lies in the SAP Data Hub product then SAP Support will resume ownership of the support for the issue and forward the ticket to SAP Data Hub Development for further issue investigation and resolution.
3. SAP Data Hub DEV now owns the issue resolution. They may still involve <partner> or customer (depending on the situation) to help them troubleshoot the issue, but ultimately SAP is now responsible for driving the issue resolution, delivering a final fix to the customer, and providing supporting documentation to SAP Support for closing the fix.
4. If the issue does not reside in the SAP product, then <partner> Support will retain ownership of the issue and log a bug with <partner> Development for further issue investigation and resolution. <partner> owns the issue resolution and may still involve SAP Data Hub DEV and customer (depending on the situation) to help them troubleshoot the issue, but ultimately <partner> Development is now responsible for driving the issue resolution, delivering a final fix or workaround to the customer, and providing supporting documentation to <partner> Support for closing the fix.