



VMware Tanzu

VMware Tanzu Container Platform  
**SAP Data Intelligence on  
Tanzu - Deployment Guide**

---

Tanzu Solutions Engineering

September 2022

[Overview](#)

[Prerequisites](#)

[Pre-Installation Tips for SAP Data Intelligence on Tanzu Kubernetes Operations](#)

[PKI Setup](#)

[Private Registry \(Harbor\)](#)

[Configure Projects](#)

[S3 Datastore](#)

[TKG Configuration](#)

[AKODeploymentConfig](#)

[TKG Documentation](#)

[Service Engine Sharing](#)

[Post SAP DI Install Procedures](#)

[NSX-T ALB L7](#)

[Create Ingress](#)

[NSX-T ALB L4 + Contour L7](#)

[Create Ingress](#)

[Create HostRule](#)

[SAP DI Configuration](#)

[Configure Modeling Tenants to Use Private Registry](#)

[Configure Connection to SDL to DI\\_DATA\\_LAKE](#)

[Configure ML Scenario Connection to S3 Data Store](#)

[Using Self-Signed Certificates](#)

[TKG](#)

[Trust Custom CA Certificates in New Clusters](#)

[SAP DI](#)

[Installation for Modeling Test](#)

[Installation for Backup/Restore Target](#)

[Observability](#)

[Backup/Restore](#)

[Useful Tools](#)

[slcb](#)

[vctl](#)



## Overview

This guide is to assist application operators and infrastructure engineering teams to deliver a successful, manageable install of SAP Data Intelligence on VMware Tanzu.

VMware Tanzu simplifies operation of Kubernetes for multi-cloud deployment by centralizing management and governance for clusters and teams across on-premises, public clouds and edge.

It delivers an open source aligned Kubernetes distribution with consistent operations and management to support infrastructure and app modernization.

This document provides added instructions for installing and configuring SAP Data Intelligence Tanzu for Kubernetes Operations on a vSphere environment backed by a Virtual Distributed Switch (VDS) based on the reference architecture defined [here](#) and does not cover any deployment procedures for the underlying SDDC components; and does not replace any documentation provided by SAP or VMware.

## Prerequisites

Consult the [Tanzu for Kubernetes Operations on vSphere](#) - Deployment Guide for specific details on configuring your VMware SDDC install to support Tanzu and the install of SAP Data Intelligence.

In addition, consult the following SAP materials for installation of SAP Data Intelligence on a Kubernetes cluster.

---

SAP Data	<a href="https://help.sap.com/docs/SAP_DATA_INTE">https://help.sap.com/docs/SAP_DATA_INTE</a>
Intelligence	<a href="#">LLIGENCE_ON-PREMISE</a>
Documentation	

---

SAP Data  
Intelligence  
Blogs

<https://blogs.sap.com/>

(click on SAP Data Intelligence tag)

SAP Press Book: [https://www.sap-press.com/sap-data-intelligence\\_5369/](https://www.sap-press.com/sap-data-intelligence_5369/)

SAP Data  
Intelligence

The  
Comprehensive  
Guide

---

This guide assumes that the Application Operations team has a working vSphere and TKG setup to install SAP Data Intelligence on.

## Pre-Installation Tips for SAP Data Intelligence on Tanzu Kubernetes Operations

### PKI Setup

Certificates will be needed for TLS Termination for the following:

Private Registry
------------------

Private S3 Datastore
NSX-T ALB

These certs should have proper SAN set and be signed by a trusted certificate authority.

If using self-signed certificates, see the [section](#) below for additional information and configurations needed.

### Private Registry (Harbor)

#### Configure Projects

Create two new projects in your Harbor registry. One is for the SAP DI installation container images and the other one is for the modeling images built by Kaniko running on Kubernetes.

The names of both projects will be entered in during installation of SAP DI.

Harbor docs for project creation	<a href="https://goharbor.io/docs/2.5.0/working-with-projects/create-projects/">https://goharbor.io/docs/2.5.0/working-with-projects/create-projects/</a>
----------------------------------	---

### S3 Datastore

You will configure multiple buckets to support SAP DI backup and the SAP DI\_DATA\_LAKE connection for file storage.

This datastore may also be used for backups with Velero.

## TKG Configuration

### AKODeploymentConfig

The SAP DI application runs in ClusterIP mode, so in using the NSX ALB, you will have to configure ingress to use this mode.

### TKG Documentation

NSX-T ALB L4 + Contour L7	<a href="https://docs.vmware.com/en/VMware-Tanzu-Kubernetes-Grid/1.5/vmware-tanzu-kubernetes-grid-15/GUID-packages-ingress-contour.html">https://docs.vmware.com/en/VMware-Tanzu-Kubernetes-Grid/1.5/vmware-tanzu-kubernetes-grid-15/GUID-packages-ingress-contour.html</a>
---------------------------	---

## NSX-T LB (Avi)

### Service Engine Sharing

There is a limitation in the NSX-T LB on sharing Service Engines (SE's) with an application running in Kubernetes **ClusterIP** mode. If you have multiple installations of SAP DI running on separate TKG Workload Clusters (say Dev/Test and separately Production) but sharing the same Avi Controller, you will need to assign each cluster their own Service Engine Group and specify separate `AKOConfigs` and cluster labels.

## Post SAP DI Install Procedures

### NSX-T ALB L4 + Contour L7

#### Create Ingress

The following example is for creating an HTTPProxy to the `vsystem` service on port 8797 that serves as the entry point to the SAP DI application.

```
# httpproxy.yaml
apiVersion: projectcontour.io/v1
kind: HTTPProxy
```

```
metadata:
  name: sapdi
  namespace: sap
spec:
  virtualhost:
    fqdn: di-contour.tanzu.test
    tls:
      passthrough: true
  tcpproxy:
    services:
      - name: vsystem
        port: 8797
```

- Configure with your FQDN and with any other customizations
- Deploy this file with `kubectl apply -f`

See docs for more information:

NSX-T ALB Documentation	<a href="https://avinetworks.com/docs/ako/1.6/custom-resource-definitions/#hostrule">https://avinetworks.com/docs/ako/1.6/custom-resource-definitions/#hostrule</a>
AKO Documentation	<a href="https://github.com/vmware/load-balancer-and-ingress-services-for-kubernetes/blob/release-1.7.2/docs/crds/hostrule.md">https://github.com/vmware/load-balancer-and-ingress-services-for-kubernetes/blob/release-1.7.2/docs/crds/hostrule.md</a>

## SAP DI Configuration



## Configure Modeling Tenants to Use Private Registry

For each SAP DI tenant created, you will need to upload the registry password secret. Consult the SAP Documentation for the process.

[https://help.sap.com/doc/52d0f61d1a734dfc97ad55599f9e23c8/3.2.8/en-US/onprem\\_loio52d0f61d1a734dfc97ad55599f9e23c8.pdf](https://help.sap.com/doc/52d0f61d1a734dfc97ad55599f9e23c8/3.2.8/en-US/onprem_loio52d0f61d1a734dfc97ad55599f9e23c8.pdf)

### 5.4.3.2 Provide Access Credentials for a Password Protected Container Registry

#### Configure Connection to SDL to DI\_DATA\_LAKE

Installation Guide: Configure the Connection to DI_DATA_LAKE	<a href="https://help.sap.com/docs/SAP_DATA_INTELLIGENCE_ON-PREMISE/a8d90a56d61a49718ebcb5f65014bbe7/c20c0251b3654103818e35ee28a8493c.html">https://help.sap.com/docs/SAP_DATA_INTELLIGENCE_ON-PREMISE/a8d90a56d61a49718ebcb5f65014bbe7/c20c0251b3654103818e35ee28a8493c.html</a>
Administration Guide: SDL	<a href="https://help.sap.com/docs/SAP_DATA_INTELLIGENCE_ON-PREMISE/b13b5722c8ff4bf9bb097251310031d0/a6b555f56d8c4641bd1a248231202050.html">https://help.sap.com/docs/SAP_DATA_INTELLIGENCE_ON-PREMISE/b13b5722c8ff4bf9bb097251310031d0/a6b555f56d8c4641bd1a248231202050.html</a>

## Configure ML Scenario Connection to S3 Data Store

Go to Connection Management

Create a storage connection to an S3 compatible store.

## Using Self-Signed Certificates

It's highly recommended that you use a PKI management system to create and distribute TLS certificates. In the event that you use self-signed certificates, you will

need to create ones with specific SAN and then you will need to update the following.

## TKG

### Trust Custom CA Certificates in New Clusters

TKG Product Documentation	<a href="https://docs.vmware.com/en/VMware-Tanzu-Kubernetes-Grid/1.5/vmware-tanzu-kubernetes-grid-15/GUID-cluster-lifecycle-secrets.html#trust-custom-ca-certificates-in-new-clusters-6">https://docs.vmware.com/en/VMware-Tanzu-Kubernetes-Grid/1.5/vmware-tanzu-kubernetes-grid-15/GUID-cluster-lifecycle-secrets.html#trust-custom-ca-certificates-in-new-clusters-6</a>
---------------------------	---

## SAP DI

### Installation for Modeling Test

During the SAP DI installation process, the installer will check the connection to the Private Registry and deploy a test modeling container. If the certificate is self-signed and you receive the following error,

<a href="https://me.sap.com/notes/3046261">https://me.sap.com/notes/3046261</a>	Requires SAP S Account access

you will need to follow the steps described in

[https://help.sap.com/docs/SAP\\_DATA\\_INTELLIGENCE\\_ON-PREMISE/a8d90a56d61a49718ebcb5f65014bbe7/39e8e391d5984e919725e601f089db74.html](https://help.sap.com/docs/SAP_DATA_INTELLIGENCE_ON-PREMISE/a8d90a56d61a49718ebcb5f65014bbe7/39e8e391d5984e919725e601f089db74.html)

Make sure to follow the Note: File name must be "cert".

### Installation for Backup/Restore Target

During installation, you will also need to update the self-signed certificate for the s3 datastore. See the procedure used for Installation for Modeling Test.

**Important: Must contain both the registry cert and the CA cert chained together.**

**If the certs are not done correctly prior to installation, fixing the certs, then retrying won't work. You'll have to completely uninstall and reinstall DI.**

## Observability

There are several options for observability on this platform.

SAP DI has its own built-in Kibana/Elasticsearch install for logs and Grafana/Prometheus install for monitoring data. You can then integrate third-party solutions with this Prometheus install through its REST API endpoints or to other Prometheus instances through federation.

Administration Guide: Integrating SAP Data Intelligence Diagnostics with an Application Performance Management Solution	<a href="https://help.sap.com/docs/SAP_DATA_INTELLIGENCE_ON-PREMISE/b13b5722c8ff4bf9bb097251310031d0/38d318e4f5de4a26991fa5e276b6c20a.html">https://help.sap.com/docs/SAP_DATA_INTELLIGENCE_ON-PREMISE/b13b5722c8ff4bf9bb097251310031d0/38d318e4f5de4a26991fa5e276b6c20a.html</a>
--	---

You may also use the TKG packages to install both Prometheus and Grafana into a separate shared services cluster.

TKG Documentation	<a href="https://docs.vmware.com/en/VMware-Tanzu-Kubernetes-Grid/1.5/vmware-tanzu-kubernetes-grid-15/GUID-packages-monitoring.html">https://docs.vmware.com/en/VMware-Tanzu-Kubernetes-Grid/1.5/vmware-tanzu-kubernetes-grid-15/GUID-packages-monitoring.html</a>
-------------------	---

There are also options to use Tanzu Observability and if using Tanzu Mission Control, integrate the two.

VMware Tanzu Observability	<a href="https://tanzu.vmware.com/observability">https://tanzu.vmware.com/observability</a>
----------------------------	---

VMware Tanzu Mission Control blog post on Observability integration	<a href="https://tanzu.vmware.com/content/tanzu-mission-control/vmware-tanzu-mission-control-introduces-integration-with-vmware-tanzu-observability-by-wavefront">https://tanzu.vmware.com/content/tanzu-mission-control/vmware-tanzu-mission-control-introduces-integration-with-vmware-tanzu-observability-by-wavefront</a>
Tanzu Observability docs on integration with TMC	<a href="https://docs.wavefront.com/tmc.html">https://docs.wavefront.com/tmc.html</a>

## Backup/Restore

SAP DI has built-in backup and requires a S3 Datastore to save checkpoints.

SAP DI Documentation	<a href="https://help.sap.com/docs/SAP_DATA_INTELLIGENCE_ON-PREMISE/a8d90a56d61a49718ebcb5f65014bbe7/e8d4c33e6cd648b0af9fd674dbf6e76c.html">https://help.sap.com/docs/SAP_DATA_INTELLIGENCE_ON-PREMISE/a8d90a56d61a49718ebcb5f65014bbe7/e8d4c33e6cd648b0af9fd674dbf6e76c.html</a>
SAP DI Backup and restore note	<a href="https://launchpad.support.sap.com/#/notes/2918288">https://launchpad.support.sap.com/#/notes/2918288</a>

## Useful Tools

SAP provides several tools for installation and operation of DI that an application operations team should explore in depth.

slcb

Installation Guide: Install SAP Data Intelligence with SLC Bridge in	<a href="https://help.sap.com/docs/SAP_DATA_INTELLIGENCE_ON-">https://help.sap.com/docs/SAP_DATA_INTELLIGENCE_ON-</a>
--	---

Unattended Mode	<a href="https://help.sap.com/docs/SAP_DATA_INTELLIGENCE_ON-PREMISE/a8d90a56d61a49718ebcb5f65014bbe7/40d59a4d9c4c4a1ab064ee93496547de.html">PREMISE/a8d90a56d61a49718ebcb5f65014bbe7/40d59a4d9c4c4a1ab064ee93496547de.html</a>
-----------------	--

vctl

<p>Product Documentation</p> <p>System Management Command-Line Client Reference for SAP Data Intelligence</p>	<a href="https://help.sap.com/docs/SAP_DATA_INTELLIGENCE_ON-PREMISE/0b99871a1d994d2ea89598fe59d16cf9/b13e0b4fe2574b3881e3c7eflad2a640.html?locale=en-US&amp;state=PRODUCTION&amp;version=3.2.8">https://help.sap.com/docs/SAP_DATA_INTELLIGENCE_ON-PREMISE/0b99871a1d994d2ea89598fe59d16cf9/b13e0b4fe2574b3881e3c7eflad2a640.html?locale=en-US&amp;state=PRODUCTION&amp;version=3.2.8</a>
<p>Blog Posts by SAP</p>	<p><a href="https://blogs.sap.com/2019/12/18/zen-and-the-art-of-sap-data-intelligence.-episode-3-vctl-the-hidden-pearl-you-must-know/">https://blogs.sap.com/2019/12/18/zen-and-the-art-of-sap-data-intelligence.-episode-3-vctl-the-hidden-pearl-you-must-know/</a></p> <p><a href="https://blogs.sap.com/2020/05/14/create-additional-users-in-sap-data-intelligence-trial-with-vctl/">https://blogs.sap.com/2020/05/14/create-additional-users-in-sap-data-intelligence-trial-with-vctl/</a></p>