

SAP Data Intelligence 3 on CaaS Platform 4.2

Installation Guide

SUSE CaaS Platform 4.2
SUSE Linux Enterprise Server for SAP applications 15
SAP Data Intelligence 3

Dr. Ulrich Schairer, SAP Solutions Architect (SUSE)
Kevin Klinger, SAP Solutions Architect (SUSE)

SAP Data Intelligence 3 on CaaS Platform 4.2

Installation Guide

Date: 2025-08-28

SAP Data Intelligence 3 is the tool set to govern big amount of data. SUSE CaaS Platform 4 is the Kubernetes base that makes deploying SAP Data Intelligence 3 easy. This document describes the installation and configuration of SUSE CaaS Platform 4 and SAP Data Intelligence 3.

Disclaimer: Documents published as part of the SUSE Best Practices series have been contributed voluntarily by SUSE employees and third parties. They are meant to serve as examples of how particular actions can be performed. They have been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. SUSE cannot verify that actions described in these documents do what is claimed or whether actions described have unintended consequences. SUSE LLC, its affiliates, the authors, and the translators may not be held liable for possible errors or the consequences thereof.

Contents


- 1 Introduction 4
- 2 Prerequisites 4
- 3 Installing SUSE CaaSP 4.2 5
- 4 Adding secure private Docker Registry for container images 9
- 5 SUSE Enterprise Storage 23
- 6 Installing SAP Data Intelligence 3 on top of SUSE CaaSP 4.2 58
- 7 Legal notice 86
- 8 GNU Free Documentation License 87

1 Introduction

This guide describes the on-premises installation of SAP Data Intelligence 3 on SUSE CaaS Platform 4.2.

2 Prerequisites

2.1 Hardware

For **sizing** information, see in addition the SAP documentation: [Sizing Guide for SAP Data Intelligence \(https://help.sap.com/viewer/835f1e8d0dde4954ba0f451a9d4b5f10/3.0.latest/en-US\)](https://help.sap.com/viewer/835f1e8d0dde4954ba0f451a9d4b5f10/3.0.latest/en-US) 

At least eight nodes are needed for a Kubernetes cluster for production use.

- Minimal requirements:
 - three master nodes
 - four worker nodes
 - one or two loadbalancers (these can be virtual machines)

For the installation of SUSE CaaSP 4.2, additional hosts are needed:

- Management host
- Registry for storing container images

These can be virtual machines.


2.2 Software

- SUSE Linux Enterprise 15 SP1
- SUSE CaaSP 4.2

3 Installing SUSE CaaSP 4.2

3.1 Documentation

SUSE CaaS Platform 4.2 is documented here:

- SUSE CaaS Platform product documentation (<https://documentation.suse.com/suse-caasp/4.2/>) 

3.2 Preparations

On all the nodes, install SUSE Linux Enterprise 15 SP1 or higher, as per the documentation for CaaS Platform 4.2.

On each respective node, the following modules or products are required.

- Management host:
 - SUSE Linux Enterprise 15 SP1
 - SUSE Linux Enterprise 15 SP1 Containers Modules
 - SUSE Linux Enterprise 15 SP1 Public Cloud
 - SUSE CaaSP 4
- Kubernetes master nodes:
 - SUSE Linux Enterprise 15 SP1
 - SUSE Linux Enterprise 15 SP1 Public Cloud
 - SUSE CaaSP 4

- Kubernetes worker nodes:
 - SUSE Linux Enterprise 15 SP1
 - SUSE Linux Enterprise 15 SP1 Public Cloud
 - SUSE CaaSP 4
- Loadbalancer host:
 - SUSE Linux Enterprise Server for SAP applications 15 SP1
or
 - SUSE Linux Enterprise 15 SP1 plus High Availability Extension

3.3 Installing the SUSE CaaSP 4 cluster nodes

- Install SUSE Linux Enterprise 15 SP1.



Note

Use the "Expert Partitioner" to disable and remove any automatically-configured swap partitions on the Kubernetes nodes.

See the relevant product documentation:

- SUSE Linux Enterprise Server 15 SP1 Deployment Guide (<https://documentation.suse.com/sles/15-SP1/single-html/SLES-deployment/#book-sle-deployment>) ↗
- SUSE CaaS Platform Deployment Guide (<https://documentation.suse.com/suse-caasp/4.2/>) ↗

3.4 Installing the loadbalancer for the Kubernetes cluster

- Install SUSE Linux Enterprise 15 SP1
- Install ha-proxy or nginx
- Configure the loadbalancer

See the relevant product documentation:

- SUSE CaaS Platform Deployment Guide (<https://documentation.suse.com/suse-caasp/4.2/single-html/caasp-deployment/>) 

3.5 Installing the management workstation

- Install SUSE Linux Enterprise 15 SP1
- Add the necessary SUSE Linux Enterprise 15 SP1 modules: value

```
$ sudo SUSEConnect -r CAASP_REGISTRATION_CODE
$ sudo SUSEConnect -p sle-module-containers/15.1/x86_64
$ sudo SUSEConnect -p caasp/4.0/x86_64 -r CAASP_REGISTRATION_CODE
$ sudo SUSEConnect -p sle-module-python2/15.1/x86_64
```

3.6 Bootstrapping the SUSE CaaS 4 cluster

- Run the skuba tool for initialization of the cluster.
- Make sure that ssh is working between all nodes without using passwords, and configure ssh-agent.

```
$ eval `ssh-agent`
$ ssh-add <path to key>
```

```
$ skuba cluster init --control-plane <LB IP/FQDN> my-cluster
```

- Bootstrap the cluster:

```
$ cd my-cluster
$ skuba node bootstrap --target <IP/FQDN> <NODE NAME>
```

- Add additional master nodes:

```
$ cd my-cluster
$ skuba node join --role master --target <IP/FQDN> <NODE NAME>
```

- Repeat this for all the master nodes.

- Add the worker nodes:

```
$ cd my-cluster  
$ skuba node join --role worker --target <IP/FQDN> <NODE NAME>
```

- Repeat this for all worker nodes.
- Finally, check the cluster status.


```
$ cd my-cluster  
$ skuba cluster status  
$ cp -av ~/my-cluster/admin.conf ~/.kube/config  
$ kubectl get nodes -o wide
```


4 Adding secure private Docker Registry for container images



Tip

This step is optional if you already have a private secure Docker Registry. If you skip this chapter, follow the instructions in [Section 5, “SUSE Enterprise Storage”](#).

To satisfy the requirements for SAP Data Intelligence 3, you also need a Docker Registry. The easiest way to build and manage one is using the [Harbor project \(https://goharbor.io/\)](https://goharbor.io/) .

To this end, you need to create a dedicated server for your Docker registry and the Harbor stack.




Warning

As Docker only allows characters within the range [a-z],[A-Z],[0-9] and '-' for domain names, make sure that your FQDN does not contain any other characters.

In our example, the server will be connected to a local bridge which provides common services (DNS, SMT, Docker-registry) for the SAP Data Intelligence stack. The FQDN of this server will be harbor-registry.example.com.


4.1 Prerequisites

Find the prerequisites for Harbor here: [Harbor Installation Prerequisites \(https://goharbor.io/docs/2.1.0/install-config/installation-prereqs/\)](https://goharbor.io/docs/2.1.0/install-config/installation-prereqs/) .

Before you can set up Harbor, you need to install Docker and Docker Compose.

- To install Docker, run:

```
# zypper in -y docker
```

- To install Docker Compose, you must download the executable from its [GitHub repository \(https://github.com/docker/compose\)](https://github.com/docker/compose)  and save it into a directory within your \$PATH.

For example, run:

```
# curl -L "https://github.com/docker/compose/releases/download/1.25.5/docker-compose-$(uname -s)-$(uname -m)" -o /usr/local/bin/docker-compose
```

```
# chmod +x /usr/local/bin/docker-compose

# ln -s /usr/local/bin/docker-compose /usr/bin/docker-compose
```

The next steps will generate the certificates used to make Harbor secure. These can also be found in the document [Configure HTTPS Access to Harbor \(https://goharbor.io/docs/2.1.0/install-config/configure-https/\)](https://goharbor.io/docs/2.1.0/install-config/configure-https/).

1. First, generate a CA certificate private key:

```
# openssl genrsa -out ca.key 4096
```

2. Then, generate a certificate with the given key for your domain.



Note

For all further steps, replace `<FQDN>` with your fully qualified domain name. In our example, this would be `harbor-registry.example.com`.

```
# openssl req -x509 -new -nodes -sha512 -days 3650 \
-subj "/C=DE/ST=BW/O=SUSE/CN=<FQDN>" \
-key ca.key \
-out ca.crt
```

Your CA certificate is now ready for use.

3. Next, you must generate a server certificate as follows:

```
# openssl genrsa -out <FQDN>.key 4096
```

4. Generate a certificate signing request (CSR):

```
# openssl req -sha512 -new \
-subj "/C=DE/ST=BW/O=SUSE/CN=<FQDN>" \
-key <FQDN>.key \
-out <FQDN>.csr
```

5. Create an `x509 v3` extension file with the following content:

```
authorityKeyIdentifier=keyid,issuer
basicConstraints=CA:FALSE
keyUsage = digitalSignature, nonRepudiation, keyEncipherment, dataEncipherment
```

```
extendedKeyUsage = serverAuth
subjectAltName = @alt_names

[alt_names]
DNS.1=<FQDN>
DNS.2=<hostname>
```

6. Use the extension file to generate a certificate:

```
# openssl x509 -req -sha512 -days 3650 \
  -extfile v3.ext \
  -CA ca.crt -CAkey ca.key -CAcreateserial \
  -in $fqdn.csr \
  -out $fqdn.crt
```

7. Copy the .crt and .key files to the system's certificate directory:

```
# cp <FQDN>.crt /etc/pki/trust/anchors/
# cp <FQDN>.key /etc/pki/trust/anchors/
```

8. As Docker interprets .crt files as CA certificates and .cert files as clients, you must convert your .crt file as follows:

```
# openssl x509 -inform PEM -in <FQDN>.crt -out <FQDN>.cert
```

9. You can now copy the newly created certificates to your Docker certificate directory. If the directory /etc/docker/certs.d does not exist, create it.

```
# mkdir /etc/docker/certs.d/<FQDN>
# cp <FQDN>.cert /etc/docker/certs.d/<FQDN>/
# cp <FQDN>.key /etc/docker/certs.d/<FQDN>/
# cp ca.crt /etc/docker/certs.d/<FQDN>/
```



Note

If you want to expose your registry on a port other than 443, you must create a directory /etc/docker/certs.d/<FQDN>:<Port> and copy the certificates to this directory instead.

10. To introduce the certificates to Docker, restart the Docker daemon:

```
# systemctl restart docker
```

4.2 Setting up Harbor

1. Fetch the Harbor installer and extract its contents:

```
# wget https://github.com/goharbor/harbor/releases/download/v2.1.1/harbor-online-  
installer-v2.1.1.tgz  
# tar xvf harbor-online-installer-v2.1.1.tgz
```

2. Enter the extracted directory:

```
# cd harbor
```

3. Within this directory, you should find a file called `harbor.yml.tpl`. It contains the configuration for the Harbor Registry and must be adjusted.
 - a. First, edit the `hostname` field and enter your FQDN.
 - b. Next, update the **HTTPS** configuration. The subentries `certificate` and `private_key` must be adjusted so they point to the `.crt` and `.key` files you created in [Section 4.1, "Prerequisites"](#).

This should look like the following:

If you want to expose registry on a port other than 443, you can change the port sub-entry to match your desired port.

- c. You should also change the administrator password for Harbor, which by default is defined in the field harbor_admin_password: Harbor12345.
 - d. The last field to mention is data_volume: /data, which defines where all Harbor data will be stored. If you want Harbor to store the data somewhere else, enter the path to the desired directory here.
4. When done, save your changes, and rename the file by removing the .tmpl suffix:

```
# mv harbor.yml.tmpl harbor.yml
```

5. As Harbor uses nginx as a reverse proxy for all services, you must run the prepare script to configure it correctly:

```
# ./prepare
```

6. You can now start the needed containers using Docker Compose:

```
# docker-compose up -d
```

Harbor should now be up and running.

7. Distribute the CA certificate to all Kubernetes nodes, so they can access the registry. Run the following commands on all the nodes:

```
# scp <FQDN>:/etc/docker/certs.d/<FQDN>/ca.crt /etc/docker/certs.d/<FQDN>/<FQDN>.crt
# systemctl restart docker
```

For the example <FQDN> of harbor-registry.example.com, this looks like the following:

```
# scp harbor-registry.example.com:/etc/docker/certs.d/harbor-registry.example.com/
ca.crt /etc/docker/certs.d/harbor-registry.example.com/harbor-
registry.example.com.crt
# systemctl restart docker
```

4.3 Verifying configuration and setting up Harbor projects

To verify Harbor is running, you may access its Web front-end by visiting `https://<FQDN>` from your browser.

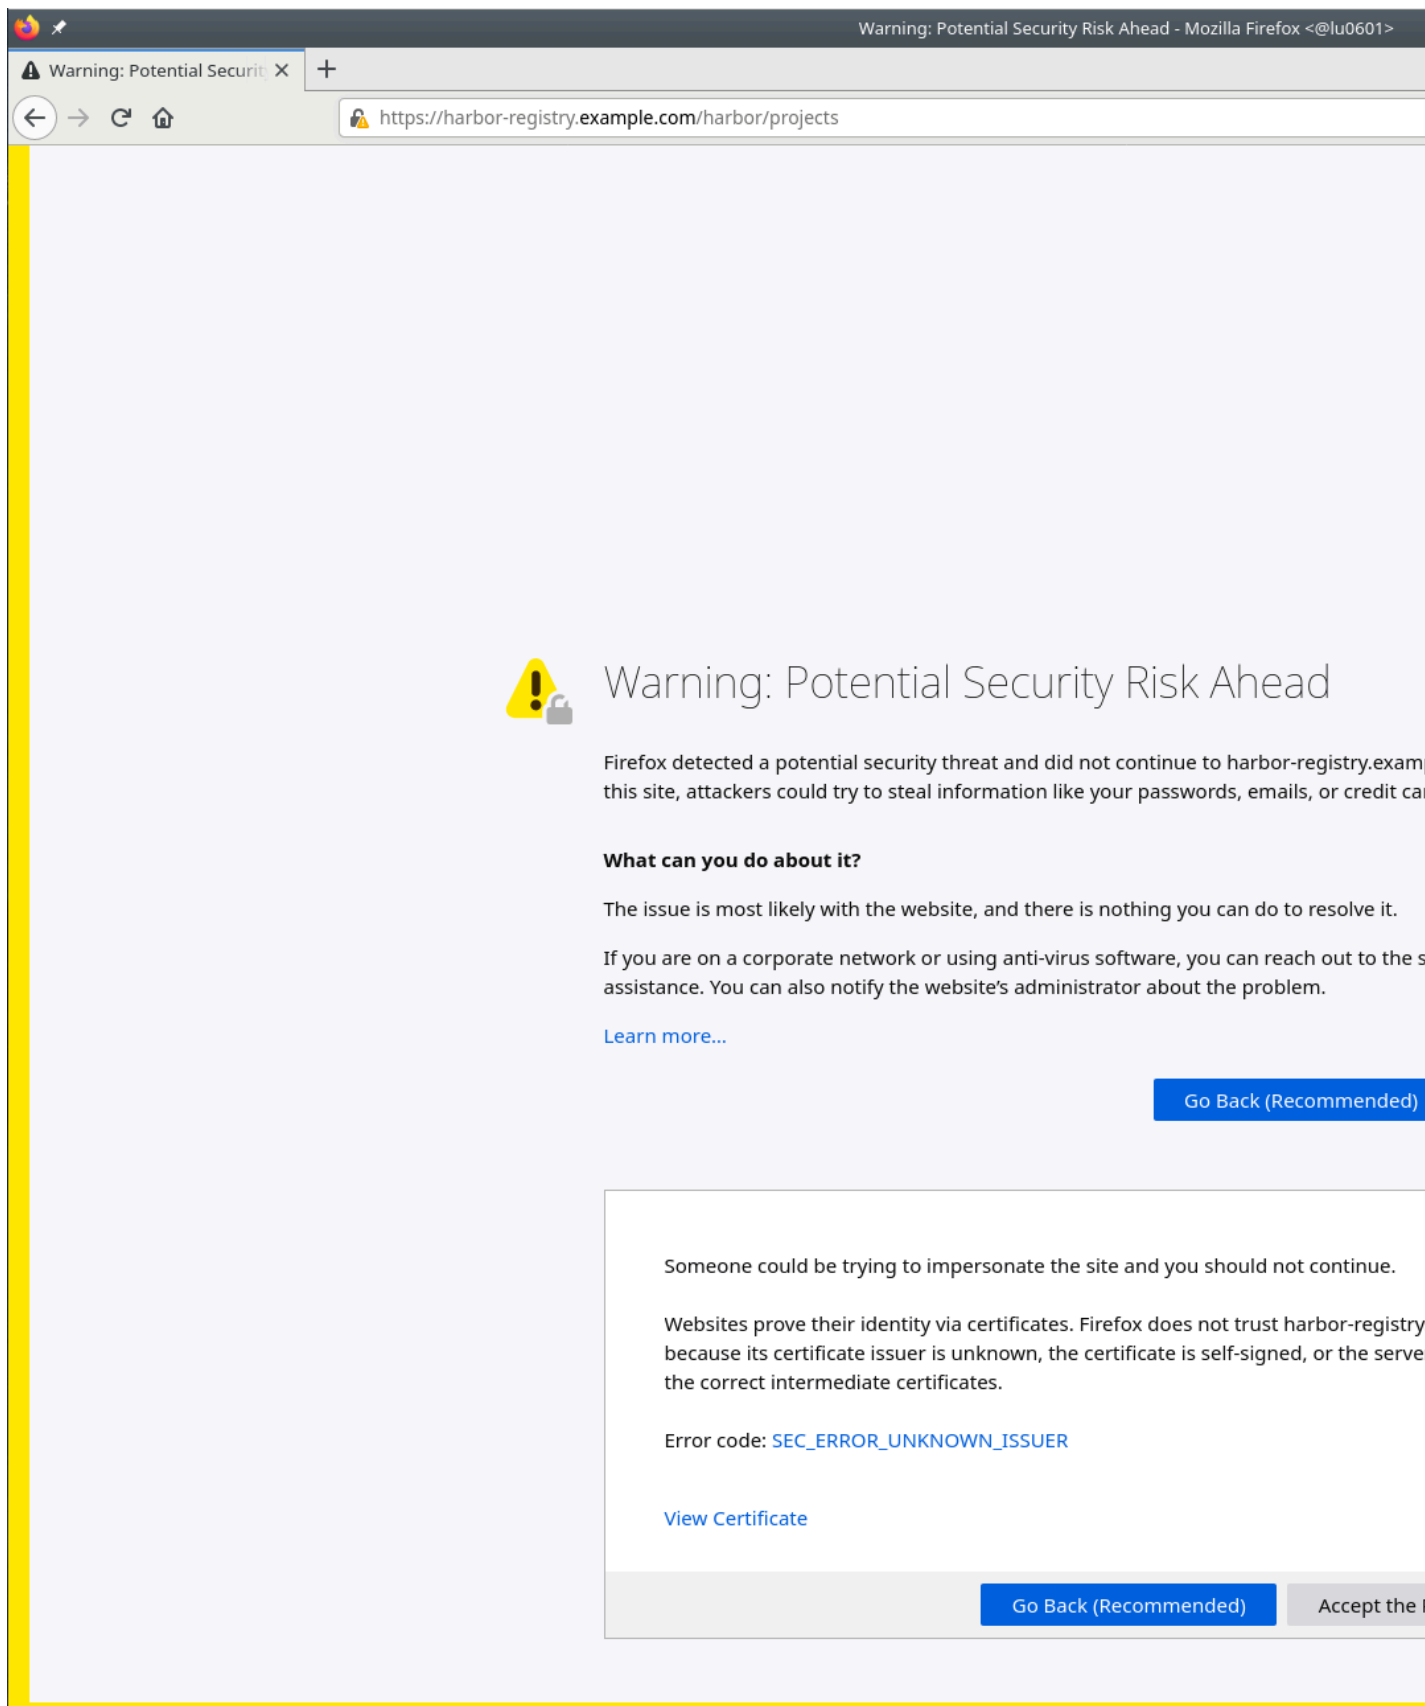


FIGURE 2: HARBOR CERTIFICATE WARNING

You can toggle the lower box by clicking Advanced... Then click Accept the Risk and Continue. You will be redirected to the login page of your Harbor registry.

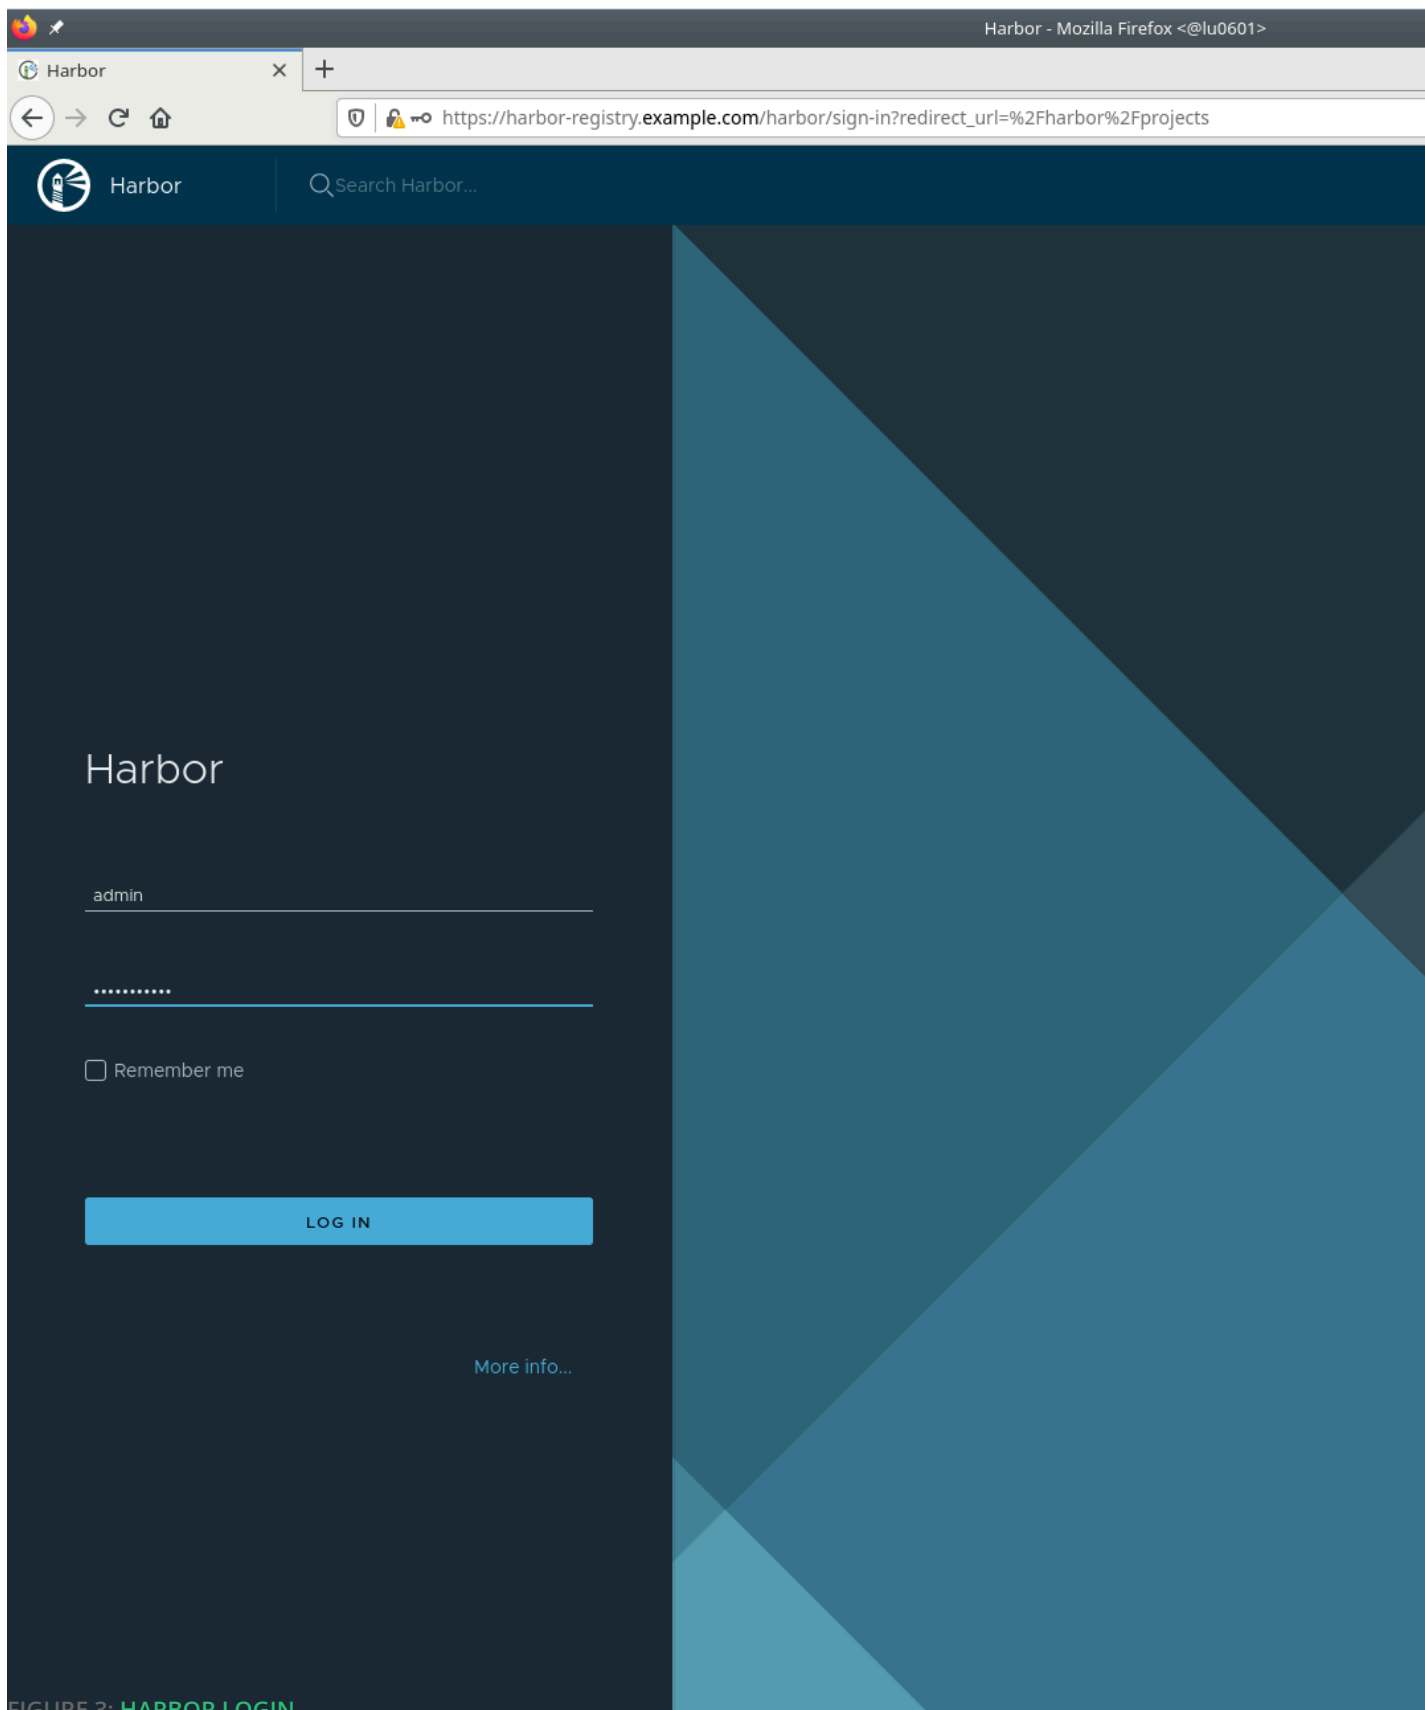


FIGURE 3: HARBOR LOGIN

Enter admin as the user name and enter the password specified in the harbor.yml file. By default, this is "Harbor12345".

By default, you will be redirected to the project page which holds the "libraries" project.

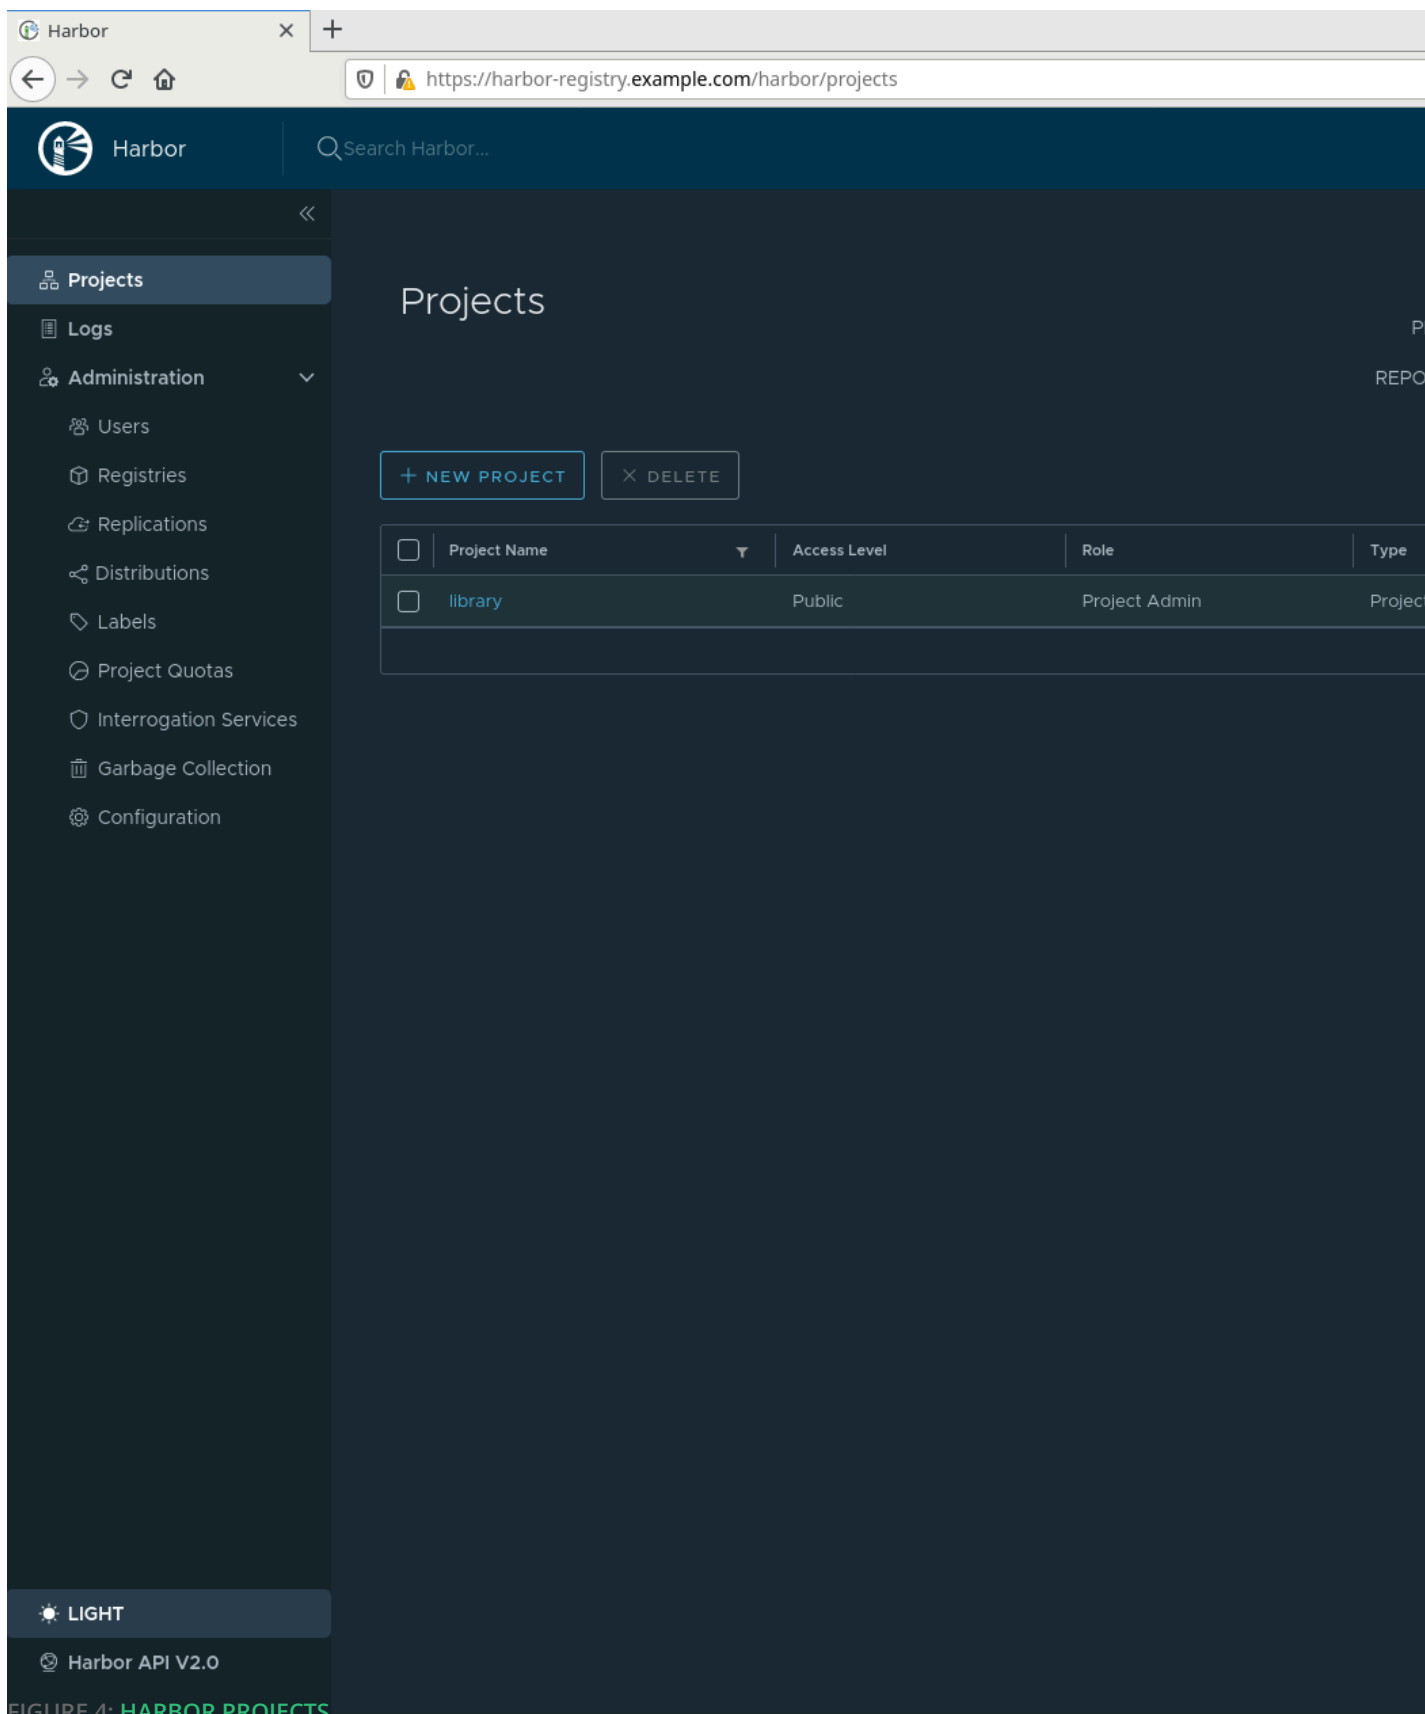


FIGURE 4: HARBOR PROJECTS

You should also check if the Docker clients on your Kubernetes nodes can access the registry. To do so, run:

```
# docker login <FQDN>
```

You will be prompted to enter a user name and password. Use admin for the user name, and the password you set in harbor.yml (default: "Harbor12345").

If Docker can access the registry, you will see a message displayed that states "Login Succeeded" or similar.

If your machines cannot resolve the FQDN of your registry, edit your /etc/hosts file and add a line with the following information:

```
<IP> <FQDN> <Hostname>
```

In our example, this will look as follows:

```
192.168.180.100 harbor-registry.example.com harbor-registry
```

Create the namespaces on your registry that are needed for SAP Data Intelligence 3:

- com.sap.hana.container
- com.sap.datahub.linuxx86_64
- com.sap.datahub.linuxx86_64.gcc6
- consul
- elasticsearch
- fabric8
- google_containers
- grafana
- kibana
- prom
- vora
- kaniko-project
- com.sap.bds.docker


5 SUSE Enterprise Storage



Tip


This step is optional if you already have a storage that provides RBD volumes and/or S3 buckets. If you skip this chapter, follow the instructions in [Section 6.4, “Installing SAP Data Intelligence 3”](#).

An on-premises installation of SAP Data Intelligence 3 requires SUSE Enterprise Storage 5.5 or higher.

If you plan to use SUSE Enterprise Storage not only for your Kubernetes dynamic storage class but also for your Kubernetes control plane (virtualized or not), you should reserve enough resources to address the `etcd` requirements specified in the [etcd Hardware recommendations](https://etcd.io/docs/current/op-guide/hardware/) (<https://etcd.io/docs/current/op-guide/hardware/>) .

The following steps will deploy a minimalist, virtualized, test-oriented instance of SUSE Enterprise Storage 6. In our example, we will build a four-nodes (1 Admin + 3 OSD) Ceph cluster.

5.1 Before starting

- Obtain registration codes for SUSE Linux Enterprise Server 15 SP1 and SUSE Enterprise Storage from <https://scc.suse.com> , or have SMT/RMT properly set up and already mirroring these products.

- SCC (<https://scc.suse.com>) 

SUSE Customer Center

suse.com Help English

INTERNAL TOOLS

Products

MY ORGANIZATIONS (2)

Connect to an organization

Manage my organizations

MY TOOLS

Support

Activate subscriptions

Packages

Patches

Container Images

Available Architectures

Release schedule

aarch64

ppc64le

s390x

x86_64

beta

Dec 14 2018

released

Jun 21 2019

Type

Product Class

Requires Regcode

EULA

Description

Still supported?

Trial product code

Download information

Installation media

Base

Module

Extension

7261

Yes

https://updates.suse.com/SUSE/Products/SLE-Product-SLES/15-SP1/x86_64/product.license/

SUSE Linux Enterprise offers a comprehensive suite of products built on a single code base. The platform addresses business needs from the smallest thin-client devices to the world's most powerful high-performance computing and mainframe servers. SUSE Linux Enterprise offers common management tools and technology certifications across the platform, and each product is enterprise-class.

Show technical details

Yes

No trial offered

Show download information

Show installation media

Extensions

SUSE CaaS Platform 4.0

SUSE Enterprise Storage 6

SUSE Linux Enterprise High Availability Extension 15 SP1

SUSE Linux Enterprise Live Patching 15 SP1

SUSE Linux Enterprise Server LTSS 15 SP1

SUSE Linux Enterprise Workstation Extension 15 SP1

SUSE Package Hub 15 SP1

Modules

Basesystem Module 15 SP1

Containers Module 15 SP1

Desktop Applications Module 15 SP1

Development Tools Module 15 SP1

Legacy Module 15 SP1

Public Cloud Module 15 SP1

Python 2 Module 15 SP1

SUSE Cloud Application Platform Tools Module 15 SP1

Server Applications Module 15 SP1

Transactional Server Module 15 SP1

Web and Scripting Module 15 SP1

Migratable from

openSUSE Leap 15.1


SUSE Linux Enterprise Server 15

© 2021 SUSE, All Rights Reserved

Privacy and Cookie Policy

FIGURE 5: SUSE CUSTOMER CENTER - SUSE LINUX ENTERPRISE SERVER

- SMT (<https://documentation.suse.com/smt/11.3/html/SLE-smt/index.html>) 



The screenshot displays the SUSE Customer Center interface. The top navigation bar includes the SUSE logo, a hamburger menu, and links for suse.com, Help, and language selection (English). The left sidebar contains sections for 'INTERNAL TOOLS' (Products), 'MY ORGANIZATIONS (2)', and 'MY TOOLS' (Support, Activate subscriptions, Packages, Patches, Container Images). The main content area shows the product page for 'SUSE Linux Enterprise Server 15 SP1 x86_64', which is marked as 'released'. The page includes filters for available architectures (aarch64, ppc64le, s390x, x86_64) and release schedules (beta, released). The product details section lists the type (Base), product class (7261), and requires regcode (Yes). The description states that SUSE Linux Enterprise offers a comprehensive suite of products built on a single code base. The page also lists extensions (SUSE CaaS Platform 4.0, SUSE Enterprise Storage 6, etc.) and modules (Basesystem Module 15 SP1, Containers Module 15 SP1, etc.). The bottom of the page shows the migratable from section, listing openSUSE Leap 15.1 and SUSE Linux Enterprise Server 15.

SUSE Customer Center

INTERNAL TOOLS

Products

MY ORGANIZATIONS (2)

MY TOOLS

Support

Activate subscriptions

Packages

Patches

Container Images

Products > SUSE Linux Enterprise Server 15 SP1 x86_64 **released**

Available Architectures: aarch64 ppc64le s390x **x86_64**

Release schedule: **beta** Dec 14 2018 **released** Jun 21 2019

Type: **Base** | Module | Extension

Product Class: 7261

Requires Regcode: Yes

EULA: https://updates.suse.com/SUSE/Products/SLE-Product-SLES/15-SP1/x86_64/product.license/

Description: SUSE Linux Enterprise offers a comprehensive suite of products built on a single code base. The platform addresses business needs from the smallest thin-client devices to the world's most powerful high-performance computing and mainframe servers. SUSE Linux Enterprise offers common management tools and technology certifications across the platform, and each product is enterprise-class.

Still supported? Yes

Trial product code: No trial offered

Download information: [Show download information](#)

Installation media: [Show installation media](#)

Extensions:

- SUSE CaaS Platform 4.0
- SUSE Enterprise Storage 6
- SUSE Linux Enterprise High Availability Extension 15 SP1
- SUSE Linux Enterprise Live Patching 15 SP1
- SUSE Linux Enterprise Server LTSS 15 SP1
- SUSE Linux Enterprise Workstation Extension 15 SP1
- SUSE Package Hub 15 SP1

Modules:

- Basesystem Module 15 SP1 **recommended**
- Containers Module 15 SP1
- Desktop Applications Module 15 SP1 **added in migration**
- Development Tools Module 15 SP1 **added in migration**
- Legacy Module 15 SP1 **added in migration**
- Public Cloud Module 15 SP1
- Python 2 Module 15 SP1 **added in migration**
- SUSE Cloud Application Platform Tools Module 15 SP1
- Server Applications Module 15 SP1 **recommended**
- Transactional Server Module 15 SP1
- Web and Scripting Module 15 SP1 **added in migration**

Migratable from:

- openSUSE Leap 15.1 **online**
- SUSE Linux Enterprise Server 15 **online**

© 2021 SUSE, All Rights Reserved

[Privacy and Cookie Policy](#)

FIGURE 6: "SUSE CUSTOMER CENTER - SUSE ENTERPRISE STORAGE

5.2

FIGURE 7: SLES SETUP INITIAL PAGE



3. Select your preferred language and keyboard layout. Then select the option "SUSE Linux Enterprise Server 15 SP1" and click the "Next" button.
4. You need to agree to the license agreement. Click "Next" again.
5. The registration screen should be displayed.



Registration

SUSE Linux

Please select

☒ Register System

E-mail

Your

Register

Enter

☐ Register System

Local

http

☐ Skip Registration

[Help](#)

6. Now, either enter your e-mail and the registration code you collected from the SCC as described in the previous chapter, or select "Register System via local SMT Server" and enter the URL of your SMT server.
7. If required, you can use the **Network Configuration** button at the top right to configure your network settings.

Add On Product

☐ I would like to install a

- ☐ Scan Using SLP...
- ☒ Specify URL...
- ☐ FTP...
- ☐ HTTP...
- ☐ HTTPS...
- ☐ SMB/CIFS
- ☐ NFS...
- ☐ CD...
- ☐ DVD...
- ☐ Hard Disk...
- ☐ USB Mass Storage
- ☐ Local Directory...
- ☐ Local ISO Image..

☒ Download repository

[Help](#)

8. When the system is registered, you will see the "Add On Product" page. You can skip it by clicking the "Next" button again.
9. The "Suggested Partitioning" page will be displayed. You may edit the proposed partitioning if needed. For our example, we accept the proposal and continue by clicking the "Next" button.

Suggested Partitioning

Initial layout proposed with the default Guid

Changes to partitioning:


- Create GPT on /dev/vda
- Create partition /dev/vda1 (8.00 MiB)
- Create partition /dev/vda2 (16.40 GiB)
- Create partition /dev/vda3 (7.93 GiB)
- Create partition /dev/vda4 (15.66 GiB)
- 9 subvolume actions ([see details](#))

[Help](#)

[Release Notes...](#)

10. On the "Clock and Time Zone" page, you should select your preferred Region and Time Zone.

Clock and Time Zone



Region

USA

Time Zone

Eastern (New York)

☒ Hardware Clock Set to UTC

Date and Time: 2020-09-0

[Help](#)

[Release Notes...](#)

Click "Other Settings" to open the "Change Date and Time" page. Select the "Synchronize with NTP-Server" option. Then select an NTP server and click the "Accept" button.

Change Date and Time

☐ Manually

Current Time

10:57:47

Current Date

2020-09-03

☒ Change the Time Now

☒ Synchronize with NTP Server

NTP Server Address

us.pool.ntp.org

Synchronize

☒ Run NTP as daemon

☐ Save NTP Configuration

[Help](#)

[Release Notes...](#)



Important

All machines in the cluster must synchronize with the same NTP Server!

11. Next, you can create a user with name and password of your choice. When done, confirm and click "Next".

Local User

☒ Create New

User's

Username

Password

Confirm

☐ Use

☐ Auto

☐ Skip User C

[Help](#)

[Release Notes...](#)

12. Now the "Installation Settings" page is displayed. Disable the firewall. To do so, click the "disable" button located behind "Firewall will be enabled".

Installation Settings

Click a headline to make changes.

Software

- Product: SUSE Linux Enterprise Server
- Product: Basesystem Module
- Product: Server Applications Module
- Product: SUSE Enterprise Storage 6
- Patterns:
 - Minimal Base System
- Size of Packages to Install: 1.3 GiB
- Downloading from Remote Repositories

Booting

- Boot Loader Type: GRUB2
- Enable Trusted Boot: no
- Status Location: /dev/vda (MBR)
- Change Location:
 - Install bootcode into MBR ([do not](#))
 - Do not install boot code into a partition
- Order of Hard Disks: /dev/vda, /dev/vdb

Security

- CPU Mitigations: [Auto](#)
- Firewall will be disabled ([enable](#))
- SSH service will be enabled ([disable](#))

Network Configuration

- Configured with DHCP: eth0
- Hostname: Set by DHCP

Kdump

- Kdump status: enabled
- Value(s) of crashkernel option: 166M,high
- Dump format: lzo
- Target of dumps: /var/crash
- Number of dumps: 5

- To finalize the installation, go to "Extension and Module Selection" and select "SUSE Linux Enterprise Storage 6 x86_64".

FIGURE 14: SLES SETUP INSTALLATION SETTINGS PAGE

[Help](#)

[Release Notes...](#)

Extension and Module Selection

Available Extensions and Modules

- ☐ SUSE CaaS Platform 4.0 x86_64
- ☒ SUSE Enterprise Storage 6 x86_64
- ☐ SUSE Linux Enterprise High Availability Extension 15 SP1 x86_64
- ☐ SUSE Linux Enterprise Live Patching 15 SP1 x86_64
- ☐ SUSE Linux Enterprise Workstation Extension 15 SP1 x86_64
- ☐ SUSE Package Hub 15 SP1 x86_64
- ☒ Basesystem Module 15 SP1 x86_64
- ☐ Containers Module 15 SP1 x86_64
- ☐ Desktop Applications Module 15 SP1 x86_64
- ☐ Development Tools Module 15 SP1 x86_64
- ☐ Legacy Module 15 SP1 x86_64
- ☐ Public Cloud Module 15 SP1 x86_64
- ☐ Python 2 Module 15 SP1 x86_64
- ☐ SUSE Cloud Application Platform Tools Module 15 SP1 x86_64

Details (English only)

SUSE Enterprise Storage 6 for SUSE Linux Enterprise Server 15 SP1 x86_64

[Help](#)

Set up at least four machines the same way.

5.3 Setting up SUSE Enterprise Storage

Check that your machines have the correct host names:

```
# hostname -f
```

This should output the FQDN of the machine. If this is not the case, you can use the following command to set the name:

```
# hostnamectl --set-hostname <FQDN>
```

For the document at hand, we use the following machine names:

- admin.example.com
- mon1.example.com
- mon2.example.com
- mon3.example.com

To ensure that DNS resolution for these names works properly, you should edit the /etc/hosts file on the nodes as follows:

```
127.0.0.1    localhost
<Admin-IP>  admin.example.com admin
<Mon1-IP>   mon1.example.com mon1
<Mon2-IP>   mon2.example.com mon2
<Mon3-IP>   mon3.example.com mon3
```

If you do not know the IP address of a machine, run:

```
$ ip a
```

```
1: lo: <LOOPBACK
    link/loopback
    inet 127.0.0.1
        valid 1
0: eth0@if9: <
    link/ether
    inet 172.17.0.1
        valid 1
```

FIGURE 16: FETCH IP

5.3.1 Preparing the nodes

1. On the OSD Nodes, install salt-minion and deepsea:

```
# zypper in -y salt-minion deepsea
```

2. On the Admin Node, install salt-minion, salt-master, and deepsea:

```
# zypper in -y salt-minion salt-master deepsea
```

3. Set the Salt master on all machines by editing the associated line:

```
# echo "master: admin.example.com" > /etc/salt/minion
```

4. On the Admin node, enable the salt-master service:

```
# systemctl enable salt-master --now
```

5. Enable the salt-minion service on all nodes:

```
# systemctl enable salt-minion --now
```

6. Now accept the Salt keys on the admin node:

```
# salt-key --accept-all -y
```



Tip

To check the keys, use the command salt-key -L. This should output something like:

```
Admin:~ # salt-key -L
Accepted Keys:
admin.example.com
mon1.example.com
mon2.example.com
mon3.example.com
Denied Keys:
Unaccepted Keys:
Rejected Keys:
```

FIGURE 17: SALT KEYS

7. Make all nodes DeepSea minions:

```
# echo "deepsea_minions: '*'" > /srv/pillar/ceph/deepsea_minions.sls
```

8. Synchronize your Salt minions:

```
# salt '*' saltutil.sync_all
```

9. Make sure your desired disks are all cleared. To this end, you can use:

```
# wipefs -a /dev/vdb
```



Note

Be aware that you might need to change the device (vdb) here. Ensure you clear all disks on all nodes.

10. Apply the cleared disks by using:

```
# salt '*' state.apply ceph.subvolume
```

5.3.2 Deploying the cluster



Tip

You can watch the progress of the Ceph stages in a separate terminal with the command deepsea monitor.

Now you can run the first stage of deploying the Ceph cluster.

1. First, prepare the cluster:

```
# salt-run state.orch ceph.stage.prep
```

2. The result for the preparation stage should look similar to this:

```
Parsing orchestration ceph.stage.prep steps... []
Parsing ceph.stage.prep steps... ✓

[1/17] ceph.salt.crc.master on
admin.example.com..... ✓ (0.5s)

[2/17] ceph.sync on
admin.example.com..... ✓ (0.8s)

[3/17] ceph.salt-api on
admin.example.com..... ✓ (18s)

[4/17] ceph.repo on
admin.example.com..... ✓ (0.3s)

[5/17] ceph.metapackage on
admin.example.com..... ✓ (0.5s)

[6/17] ceph.updates on
admin.example.com..... ✓ (10s)

[7/17] filequeue.remove(item=lock)..... ✓ (0.0s)

[8/17] filequeue.enqueue(item=complete)..... ✓ (0.0s)

[9/17] minions.ready(timeout=300)..... ✓ (0.3s)

[10/17] ceph.salt.crc.minion on
admin.example.com..... ✓ (1s)
mon1.example.com..... ✓ (1s)
mon3.example.com..... ✓ (1s)
mon2.example.com..... ✓ (2s)

[11/17] ceph.apparmor on
admin.example.com..... ✓ (14s)
mon1.example.com..... ✓ (12s)
mon3.example.com..... ✓ (13s)
mon2.example.com..... ✓ (13s)

[12/17] ceph.repo on
admin.example.com..... ✓ (0.4s)
mon1.example.com..... ✓ (0.5s)
mon3.example.com..... ✓ (0.4s)
mon2.example.com..... ✓ (0.4s)

[13/17] ceph.metapackage on
admin.example.com..... ✓ (0.5s)
mon1.example.com..... ✓ (0.6s)
mon3.example.com..... ✓ (0.6s)
mon2.example.com..... ✓ (0.6s)

[14/17] ceph.packages.common on
admin.example.com..... ✓ (28s)
mon1.example.com..... ✓ (65s)
mon3.example.com..... ✓ (69s)
mon2.example.com..... ✓ (69s)

[15/17] ceph.sync on
admin.example.com..... ✓ (2s)
mon1.example.com..... ✓ (3s)
mon3.example.com..... ✓ (2s)
mon2.example.com..... ✓ (2s)

[16/17] ceph.mines on
admin.example.com..... ✓ (0.7s)
mon1.example.com..... ✓ (0.8s)
mon3.example.com..... ✓ (0.8s)
mon2.example.com..... ✓ (0.7s)


[17/17] ceph.updates on
admin.example.com..... ✓ (10s)
mon1.example.com..... ✓ (13s)
mon3.example.com..... ✓ (13s)
mon2.example.com..... ✓ (13s)

Ended stage: ceph.stage.prep succeeded=17/17 time=199.7s
```

FIGURE 18: CLUSTER DEPLOYMENT - PREPARATION

3. The next stage is collecting information about the nodes:

```
# salt-run state.orch ceph.stage.discovery
```



```
Parsing orchestration ceph.stage.discovery steps... []
Parsing ceph.stage.discovery steps... ✓

[1/3] minions.ready(timeout=300)..... ✓ (0.4s)
[2/3] ceph.refresh on
      admin.example.com..... ✓ (0.5s)
[3/3] populate.proposals..... ✓ (5s)
Ended stage: ceph.stage.discovery succeeded=3/3 time=12.7s
```

FIGURE 19: CLUSTER DEPLOYMENT - STAGE 1

4. Before you can run the last three stages, you must provide a role configuration for the nodes. This will be stored in /srv/pillar/ceph/proposals/policy.cfg. This example uses the following configuration:

```
#General config
config/stack/default/global.yml
config/stack/default/ceph/cluster.yml

#CEPH Cluster members
cluster-ceph/cluster/admin.example.com.sls
cluster-ceph/cluster/mon1.example.com.sls
cluster-ceph/cluster/mon2.example.com.sls
cluster-ceph/cluster/mon3.example.com.sls

#CEPH Admin nodes
role-admin/cluster/mon1.example.com.sls
role-admin/cluster/mon2.example.com.sls
role-admin/cluster/mon3.example.com.sls

#CEPH Master node
role-master/cluster/admin.example.com.sls

#CEPH Manager nodes
role-mgr/cluster/mon1.example.com.sls
role-mgr/cluster/mon2.example.com.sls
role-mgr/cluster/mon3.example.com.sls

#CEPH Monitor nodes
role-mon/cluster/mon1.example.com.sls
role-mon/cluster/mon2.example.com.sls
role-mon/cluster/mon3.example.com.sls
```

```
#CEPH RGW nodes
role-rgw/cluster/admin.example.com.sls
role-rgw/cluster/mon1.example.com.sls
role-rgw/cluster/mon2.example.com.sls
role-rgw/cluster/mon3.example.com.sls

#CEPH Storage nodes
role-storage/cluster/admin.example.com.sls
role-storage/cluster/mon1.example.com.sls
role-storage/cluster/mon2.example.com.sls
role-storage/cluster/mon3.example.com.sls
```

5. You can now safely deploy your configuration:

```
# salt-run state.orch ceph.stage.configure
```



```

Parsing orchestration ceph.stage.configure steps... []
Parsing ceph.stage.configure steps... ✓

Stage initialization output:
deepsea_minions      : valid
yaml_syntax          : valid

[1/15] push.proposal..... ✓ (0.0s)
[2/15] ceph.refresh on
      mon1.example.com..... ✓ (2s)
      admin.example.com..... ✓ (1s)
      mon2.example.com..... ✓ (1s)
      mon3.example.com..... ✓ (2s)
[3/15] advise.networks..... ✓ (0.6s)
[4/15] ceph.packages on
      mon1.example.com..... ✓ (53s)
      admin.example.com..... ✓ (34s)
      mon2.example.com..... ✓ (58s)
      mon3.example.com..... ✓ (57s)
[5/15] ceph.admin.key on
      admin.example.com..... ✓ (0.6s)
[6/15] ceph.osd.key on
      admin.example.com..... ✓ (0.4s)
[7/15] ceph.mon.key on
      admin.example.com..... ✓ (0.4s)
[8/15] ceph.mgr.key on
      admin.example.com..... ✓ (2s)
[9/15] ceph.igw.key on
      admin.example.com..... ✓ (0.9s)
[10/15] ceph.mds.key on
      admin.example.com..... ✓ (0.8s)
[11/15] ceph.rgw.key on
      admin.example.com..... ✓ (2s)
[12/15] ceph.ganesha.key on
      admin.example.com..... ✓ (0.8s)
[13/15] ceph.monitoring.prometheus.exporters.node_exporter on
      mon1.example.com..... ✓ (68s)
      admin.example.com..... ✓ (41s)
      mon2.example.com..... ✓ (69s)
      mon3.example.com..... ✓ (64s)
[14/15] ceph.ssl on
      admin.example.com..... ✓ (0.7s)
[15/15] ceph.ssl.distribute_ca on
      admin.example.com..... ✓ (1s)

Ended stage: ceph.stage.configure succeeded=15/15 time=146.0s

```

FIGURE 20: CLUSTER DEPLOYMENT - STAGE 2

6. Deploy your configuration to the cluster:

```
# salt-run state.orch ceph.stage.deploy
```

```

Parsing orchestration ceph.stage.deploy steps... 
Parsing ceph.stage.deploy steps... ✓

Stage initialization output:
firewall      : not installed
apparmor      : disabled
fsid          : valid
public_network : valid
public_interface : valid
cluster_network : valid
cluster_interface : valid
ip_version    : valid
monitors      : valid
subvolume     : valid
mgrs          : valid
storage       : valid
storage_role  : valid
rgw           : valid
ganesha       : valid
master_role   : valid
time_server   : valid
fqdn          : valid

[1/51] ceph.time on
mon3.example.com..... ✓ (43s)
admin.example.com..... ✓ (0.5s)
mon2.example.com..... ✓ (47s)
mon1.example.com..... ✓ (47s)

[2/51] ceph.configuration.check on
admin.example.com..... ✓ (0.3s)

[3/51] ceph.configuration.create on
admin.example.com..... ✓ (5s)

[4/51] ceph.configuration on
mon3.example.com..... ✓ (0.6s)
admin.example.com..... ✓ (0.5s)
mon2.example.com..... ✓ (0.5s)
mon1.example.com..... ✓ (0.6s)

[50/51] ceph.monitoring.prometheus.exporters.mgr_exporter on
admin.example.com..... ✓ (0.9s)

[51/51] ceph.monitoring.prometheus.exporters.rbd_exporter on
admin.example.com..... ✓ (18s)

Ended stage: ceph.stage.deploy succeeded=51/51 time=302.0s

```

FIGURE 21: CLUSTER DEPLOYMENT - STAGE 3

7. When the deployment stage has been successfully passed, check the cluster health to insure that all is running properly.

```
# ceph -s
```

```
ses55-admin:~ # ceph -s
cluster:
  id:      520a906e-e6a3-324f-ba18-856c2e155395
  health: HEALTH_OK

services:
  mon: 3 daemons, quorum ses55-osd0,ses55-osd1,ses55-osd2
  mgr: ses55-osd0(active), standbys: ses55-osd1, ses55-osd2
  osd: 16 osds: 16 up, 16 in

data:
  pools:   0 pools, 0 pgs
  objects: 0 objects, 0B
  usage:   16.1GiB used, 302GiB / 318GiB avail
  pgs:
```

FIGURE 22: CHECKING CLUSTER HEALTH

8. The last stage to run is deploying the service roles, which were specified in the `policy.cfg` file:

```
# salt-run state.orch ceph.stage.services
```

```

Parsing orchestration ceph.stage.services steps... []
Parsing ceph.stage.services steps... ✓

Stage initialization output:
No minions matched the target. No command was sent, no jid was assigned.
kernel_module      : valid
openattic-disabled  : valid

[1/17] ceph.rgw.auth on
      admin.example.com..... ✓ (4s)

[2/17] ceph.rgw.users on
      admin.example.com..... ✓ (34s)

[3/17] ceph.rgw.dashboard on
      admin.example.com..... ✓ (5s)

[4/17] ceph.rgw on
      mon3.example.com..... ✓ (79s)
      admin.example.com..... ✓ (23s)
      mon2.example.com..... ✓ (78s)
      mon1.example.com..... ✓ (69s)

[5/17] ceph.monitoring.prometheus.exporters.ceph_rgw_exporter on
      admin.example.com..... ✓ (31s)

[6/17] ceph.wait on
      admin.example.com..... ✓ (7s)

[7/17] ceph.processes.rgw on
      mon3.example.com..... ✓ (0.5s)
      admin.example.com..... ✓ (0.5s)
      mon2.example.com..... ✓ (0.5s)
      mon1.example.com..... ✓ (0.4s)

[8/17] ceph.rgw.restart on
      admin.example.com..... ✓ (3s)

[9/17] ceph.wait on
      admin.example.com..... ✓ (6s)

[10/17] ceph.processes.rgw on
      mon3.example.com..... ✓ (0.4s)
      admin.example.com..... ✓ (0.4s)
      mon2.example.com..... ✓ (0.5s)
      mon1.example.com..... ✓ (0.4s)

[11/17] ceph.rgw.restart on
      mon1.example.com..... ✓ (3s)

[12/17] ceph.wait on
      admin.example.com..... ✓ (7s)

[13/17] ceph.processes.rgw on
      mon3.example.com..... ✓ (0.4s)
      admin.example.com..... ✓ (0.5s)
      mon2.example.com..... ✓ (0.4s)
      mon1.example.com..... ✓ (0.5s)

[14/17] ceph.rgw.restart on
      mon2.example.com..... ✓ (3s)

[15/17] ceph.wait on
      admin.example.com..... ✓ (7s)

[16/17] ceph.processes.rgw on
      mon3.example.com..... ✓ (0.8s)
      admin.example.com..... ✓ (0.6s)
      mon2.example.com..... ✓ (0.7s)
      mon1.example.com..... ✓ (0.7s)

[17/17] ceph.rgw.restart on
      mon3.example.com..... ✓ (3s)

Ended stage: ceph.stage.services succeeded=17/17 time=206.8s

```

FIGURE 23: CLUSTER DEPLOYMENT - STAGE 4

5.4 Access the dashboard and create a new pool

After the Ceph cluster is up and running, you must create a pool for SAP Data Intelligence 3.

In our example, we will use the dashboard for this purpose.

1. The dashboard is published by any of the monitor nodes. To access it, use a Web browser:

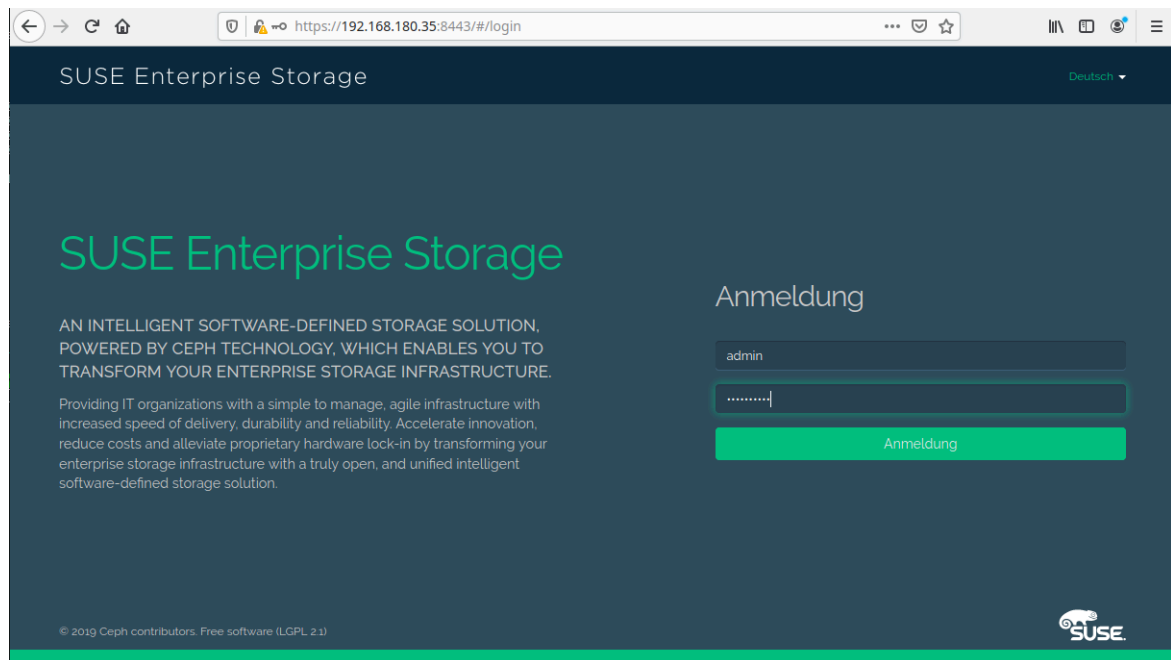


FIGURE 24: CLUSTER DASHBOARD

Get the credentials to log in by running the following command on the Admin Node:

```
# salt-call grains.get dashboard_creds
```

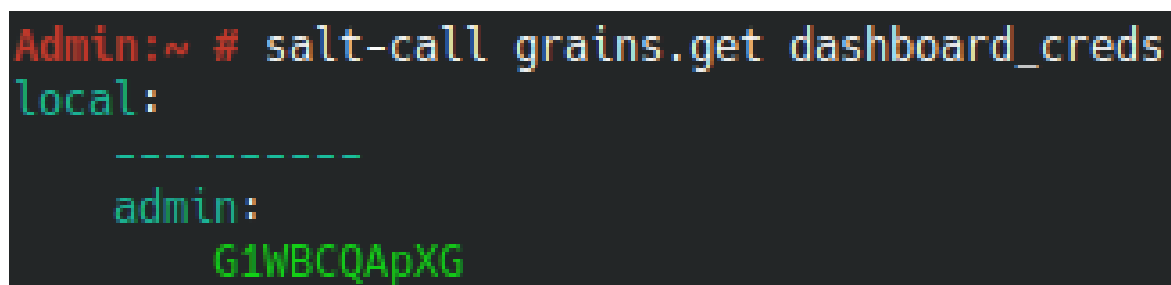


FIGURE 25: CLUSTER DASHBOARD LOGIN

After logging in, the landing page should appear.

2. Select the "Pools" tab at the top of the page:

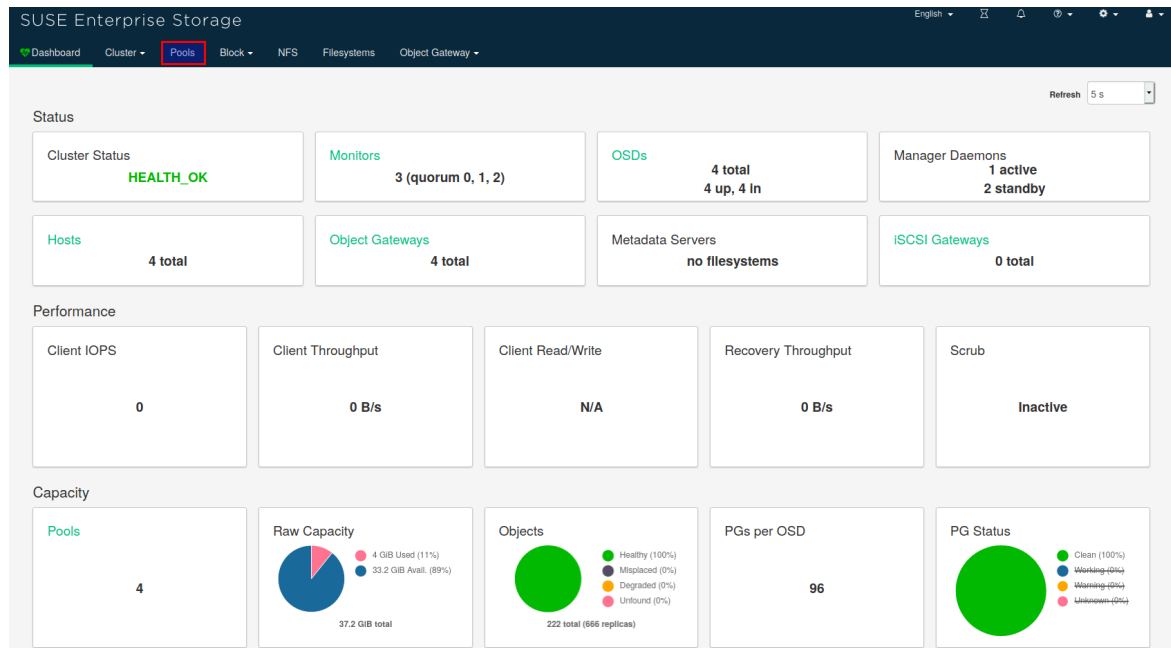


FIGURE 26: CLUSTER DASHBOARD LANDING PAGE

The Pools page gives an overview of the currently defined pools.

3. Click the "Create" button at the top of the table to create a new pool:

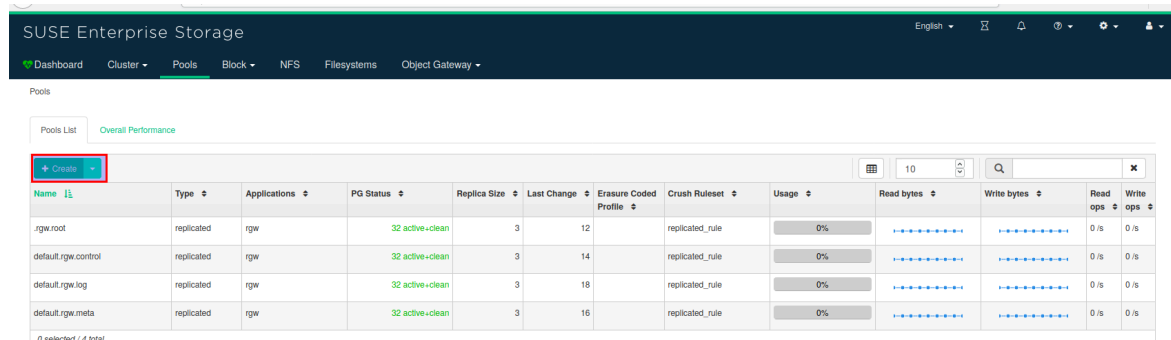


FIGURE 27: CLUSTER DASHBOARD POOLS

4. Enter the name of the pool. As "Pool type", select replicated.
5. On the left side of "Applications", click the pencil symbol and select rbid.
6. Confirm the creation of the pool by clicking the "CreatePool" button at the bottom.

SUSE Enterprise Storage

Dashboard
Cluster
Pools
Block
NFS
Filesystems
Object Gateway

Pools » Create

CreatePool

Name *

datahub

Pool type *

replicated

Placement groups *

64

Calculation help

Crush ruleset

replicated rule

Replicated size *

3

Applications

rd

bd

Compression

Mode

none

RBD Configuration

Quality of Service +

CreatePool

Cancel

FIGURE 28: CLUSTER DASHBOARD CREATING POOLS

! Important

Make note of the name of the pool. It will be needed in [Section 6.4.1.2, “Creating Storage Class”](#).

- After this, the pools page is displayed again, and the newly created pool is shown in the table of pools.

Name	Type	Applications	PG Status	Replica Size	Last Change	Erasure Coded Profile	Crush Ruleset	Usage	Read bytes	Write bytes	Read ops	Write ops
rgw.root	replicated	rgw	32 active+clean	3	12		replicated_rule	0%			0/s	0/s
datahub	replicated	rbd	64 active+clean	3	32		replicated_rule	0%			0/s	0/s
default.rgw.control	replicated	rgw	32 active+clean	3	14		replicated_rule	0%			0/s	0/s
default.rgw.log	replicated	rgw	32 active+clean	3	18		replicated_rule	0%			0/s	0/s
default.rgw.meta	replicated	rgw	32 active+clean	3	16		replicated_rule	0%			0/s	0/s

FIGURE 29: CLUSTER DASHBOARD NEW POOL

- Now provide access to this pool through an RBD device. Go to the RDB overview page by selecting "Block→Images"




FIGURE 30: CLUSTER DASHBOARD ACCESSING POOL

An overview of the configured RDBs is displayed.

- Click the **Create** button.



FIGURE 31: CLUSTER DASHBOARD RDB OVERVIEW

- Enter the name, and if it is not already selected, select the previously created pool. Select the size of the RBD and confirm the creation by clicking the "CreateRBD" button.

CreateRBD

Name * datahubRBD

Pool * datahub

☐ Use a dedicated data pool

Size * 10 GiB

Features

- ☐ Deep flatten
- ☒ Layering
- ☐ Exclusive lock
- ☐ Object map (requires exclusive-lock)
- ☐ Journaling (requires exclusive-lock)
- ☐ Fast diff (interlocked with object-map)

[Advanced...](#)

CreateRBD **Cancel**

FIGURE 32: CLUSTER DASHBOARD CREATING NEW RDB

The overview page of the RBDs is shown again. It now contains your newly created RBD.

SUSE Enterprise Storage

Dashboard Cluster Pools **Block** NFS Filesystems Object Gateway

Block > Images

Images Trash Overall Performance

+ Create -

Name	Pool	Size	Objects	Object size	Provisioned	Total provisioned	Parent
datahubRBD	datahub	10 GiB	2.6 k	4 MiB	0 B	0 B	-

0 selected / 1 total

FIGURE 33: CLUSTER DASHBOARD RDB UPDATED OVERVIEW

At this point, the SUSE Enterprise Storage cluster is ready for usage with SAP Data Intelligence.

6 Installing SAP Data Intelligence 3 on top of SUSE CaaSP 4.2

6.1 Documentation

- SAP Notes:
 - Release Notes for SAP Data Intelligence 3 (<https://launchpad.support.sap.com/#/notes/2871970>) ↗
 - Installation Guide for SAP Data Intelligence 3 (<https://help.sap.com/viewer/a8d90a56d61a49718ebcb5f65014bbe7/3.0.latest/en-US>) ↗

6.2 Planning the installation with the SAP Maintenance Planner

For the installation of SAP Data Intelligence, you should start here: [SAP Maintenance Planner \(https://support.sap.com/en/alm/solution-manager/processes-72/maintenance-planner.html\)](https://support.sap.com/en/alm/solution-manager/processes-72/maintenance-planner.html) ↗



Note

You need to have your SAP S-User available.

- The landing page of the SAP Maintenance Planner looks as follows:

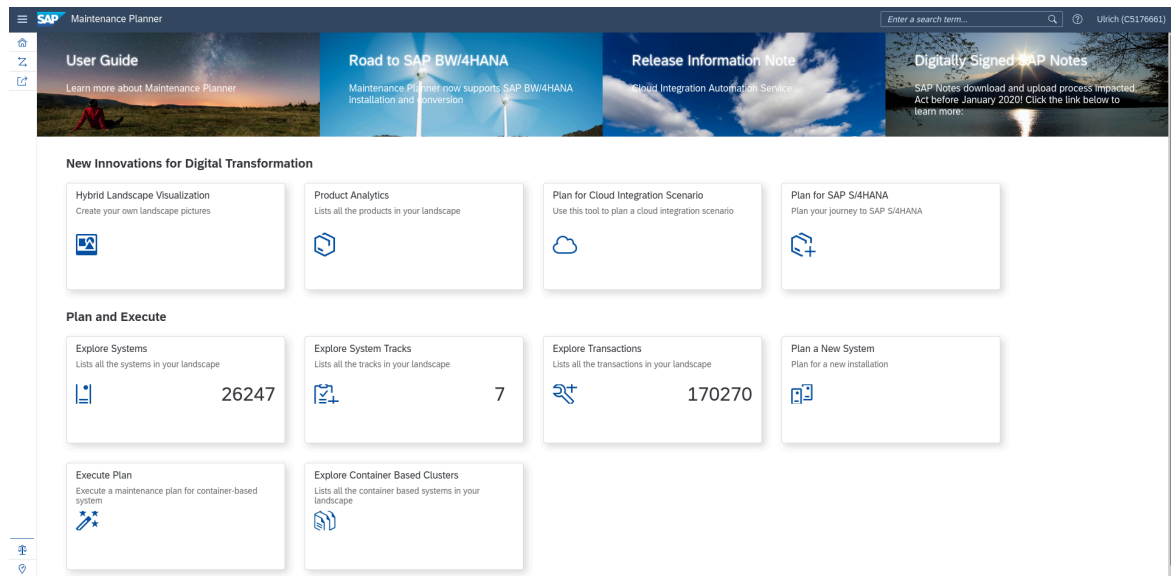


FIGURE 34: SAP MAINTENANCE PLANNER START PAGE

Click "Plan a New System" on the right.

- The next page displayed looks as follows:

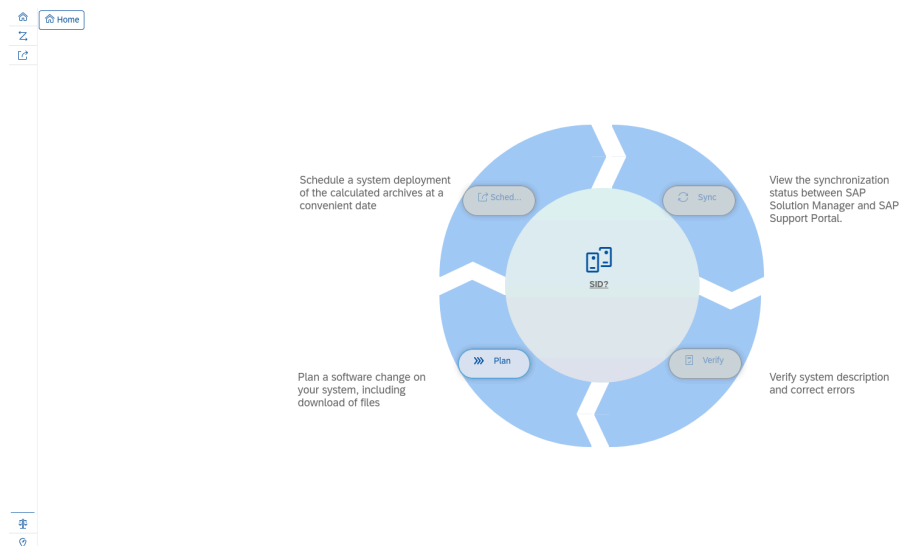


FIGURE 35: SAP MAINTENANCE PLANNER: SELECT PLAN

You will see a circle where all options except the option "Plan" are greyed out. Click the "Plan" option.

- The next page shows the "Define Change" step of your planning:

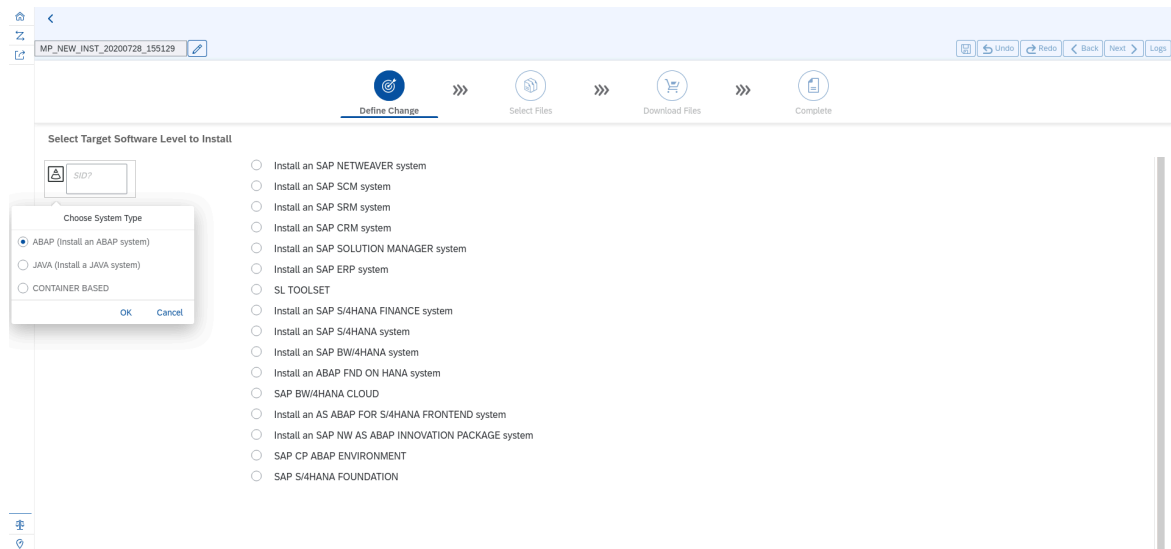


FIGURE 36: SAP MAINTENANCE PLANNER: SELECT CONTAINER BASED PRODUCT

- On the left, there is a window with three toggle buttons. Select "CONTAINER BASED" and click "OK":

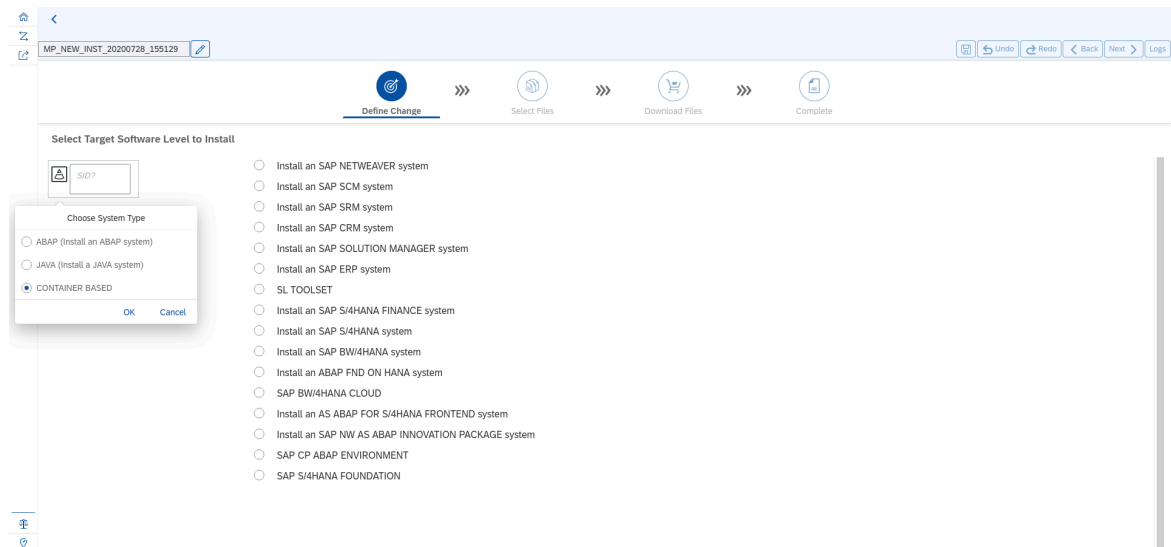


FIGURE 37: SAP MAINTENANCE PLANNER: SELECT CONTAINER BASED PRODUCT #2

- To the right, the option "SAP DATA INTELLIGENCE" should appear. When you select it, a sub-selection should open with choices for "SAP DATA INTELLIGENCE 3" and "SAP DATA HUB 2". Select "SAP DATA INTELLIGENCE 3":




FIGURE 38: SAP MAINTENANCE PLANNER: SELECT DATA INTELLIGENCE 3

- A pop-up window will appear to inform you about the related SAP note. Click "Continue" to proceed:




FIGURE 39: SAP MAINTENANCE PLANNER: SELECT CONTINUE

- On the right hand side, a drop-down box is shown with "Select Support Package Stack". Click this, and select from the available patch levels as needed:

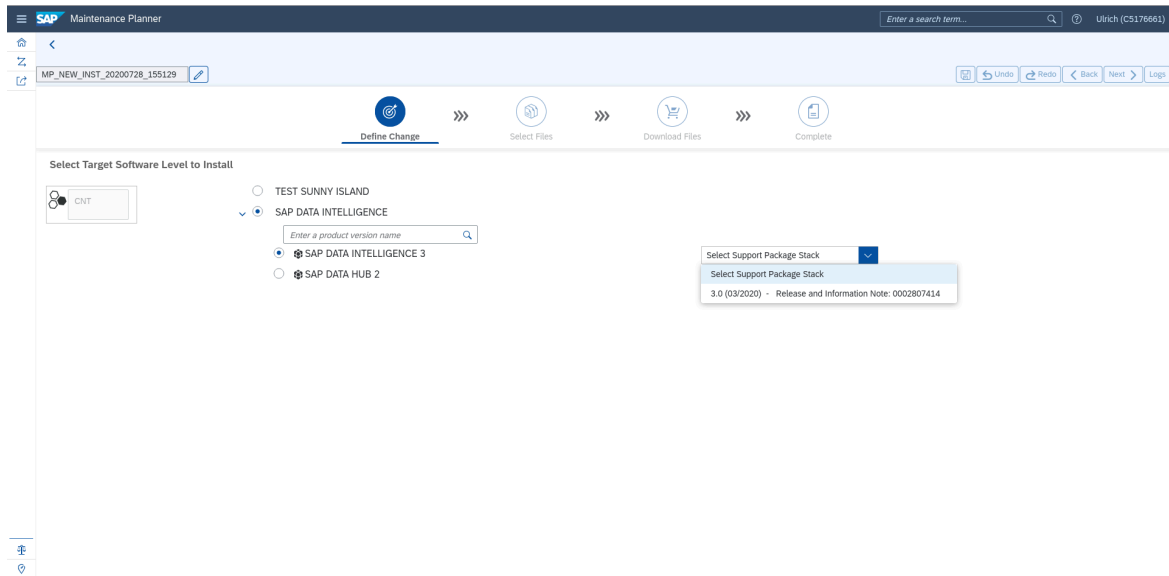


FIGURE 40: SAP MAINTENANCE PLANNER: SELECT FROM AVAILABLE PATCH LEVELS AS YOU NEED.

- To the left, the sub-selection for "SAP DATA INTELLIGENCE 3" will need changes. Select what you need and click the "Confirm Selection" at the very right:

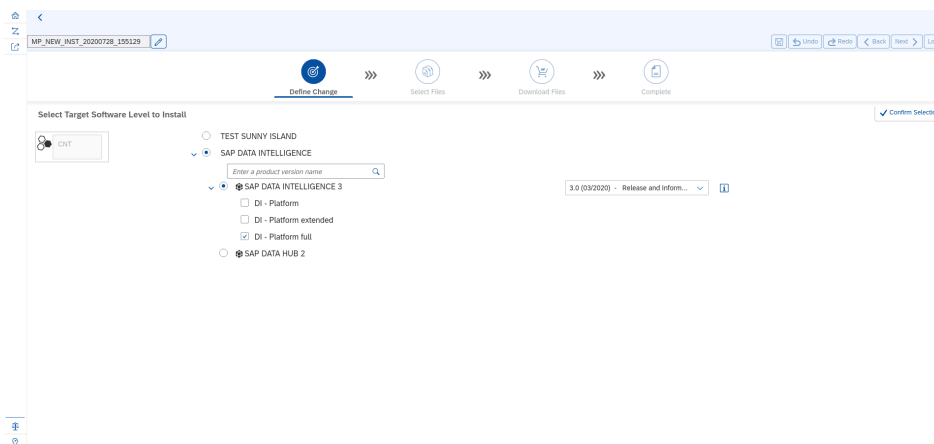


FIGURE 41: SAP MAINTENANCE PLANNER: SELECT ACCORDING YOUR NEEDS AND CONFIRM.

- An overview of your selection is shown. If this fits your needs, click the "Next" button at the upper right corner:

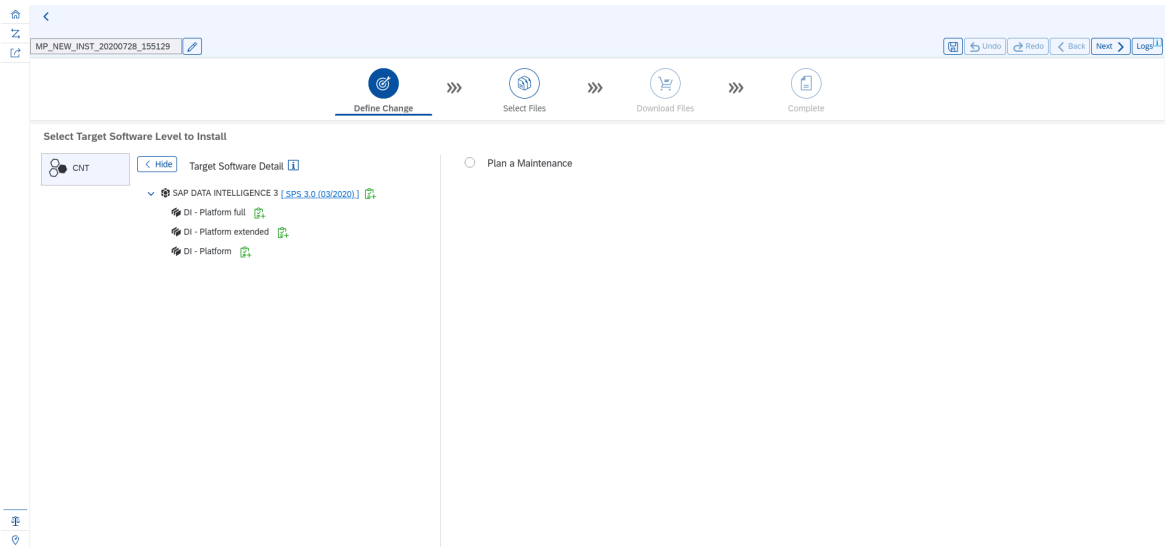


FIGURE 42: SAP MAINTENANCE PLANNER: SELECT NEXT IF SATISFIED

- Next, select the operating system upon which SAP Data Intelligence will be installed. Select "Linux on x86_64 64bit" and click "Confirm Selection":




FIGURE 43: SAP MAINTENANCE PLANNER: SELECT LINUX AND CONFIRM

- The next page shows the preselected files to use and download. Again, click the "Next" button on the upper right:

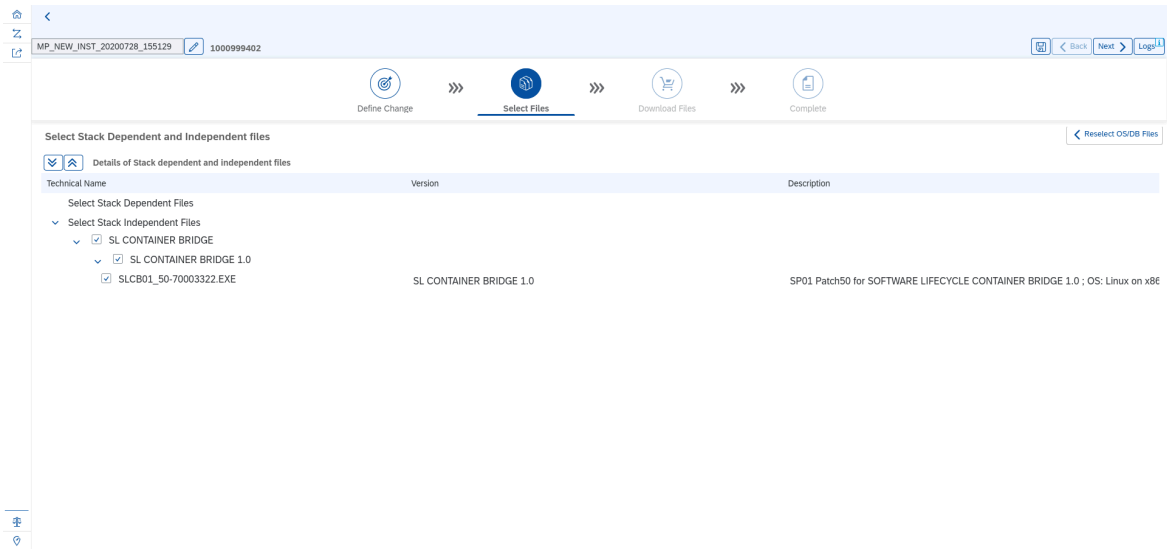


FIGURE 44: SAP MAINTENANCE PLANNER: CONFIRM

- You should now arrive at the "Download Files" page. The required **SLC bridge** is already preselected. Click "Next" to proceed:




FIGURE 45: SAP MAINTENANCE PLANNER: EXECUTE PLAN

- Your maintenance plan is shown as PDF. Confirm everything by clicking the "Execute Plan" button on the upper right side:

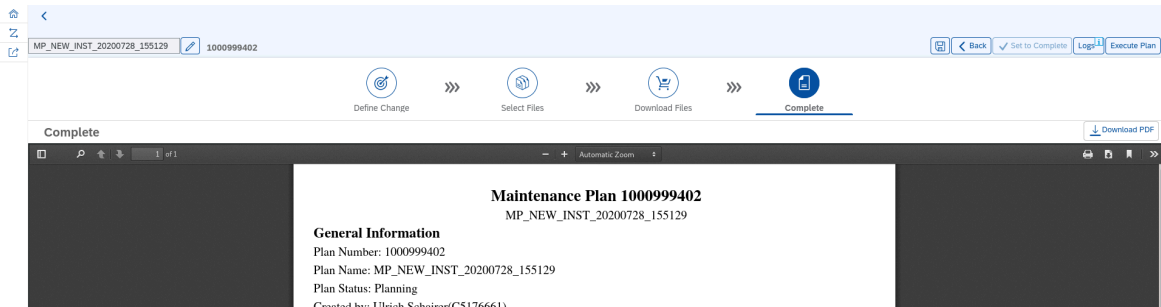


FIGURE 46: SAP MAINTENANCE PLANNER: DOWNLOAD STACK.XML AND SLC BRIDGE INSTALLER

- Download the **SLC Bridge Installer** and copy it to your management workstation. You will need this file in *Section 6.3, "Installing the SAP SLC Bridge"*.

- You will be prompted to enter the FQDN and the port of the machine your SLC bridge will run on:

Home MP_EP_20200728_160349 Next > Logs

Overview Deploy Complete

Execute plan overview

Enter Transaction ID
1000999402 X

Please enter the details of the host where the container bridge is already installed, or is to be installed:

Fully Qualified Domain Name
Port
9000

If not already downloaded, please download the below files:
Click [here](#) to download Stack XML
Click [here](#) to download SL CONTAINER BRIDGE 1.0.Linux on x86_64 64bit Database independent [SLCB01_50-70003322.EXE ,76422 KB]

FIGURE 47: SAP MAINTENANCE PLANNER: ENTER FQDN OF HOST WHERE THE SLC BRIDGE WILL RUN

- Fill in the values. An example looks as follows:

Home MP_EP_20200728_160349 Next > Logs

Overview Deploy Complete

Execute plan overview

Enter Transaction ID
1000999402 X

Please enter the details of the host where the container bridge is already installed, or is to be installed:

Fully Qualified Domain Name
youmode.example.com
Port
30778

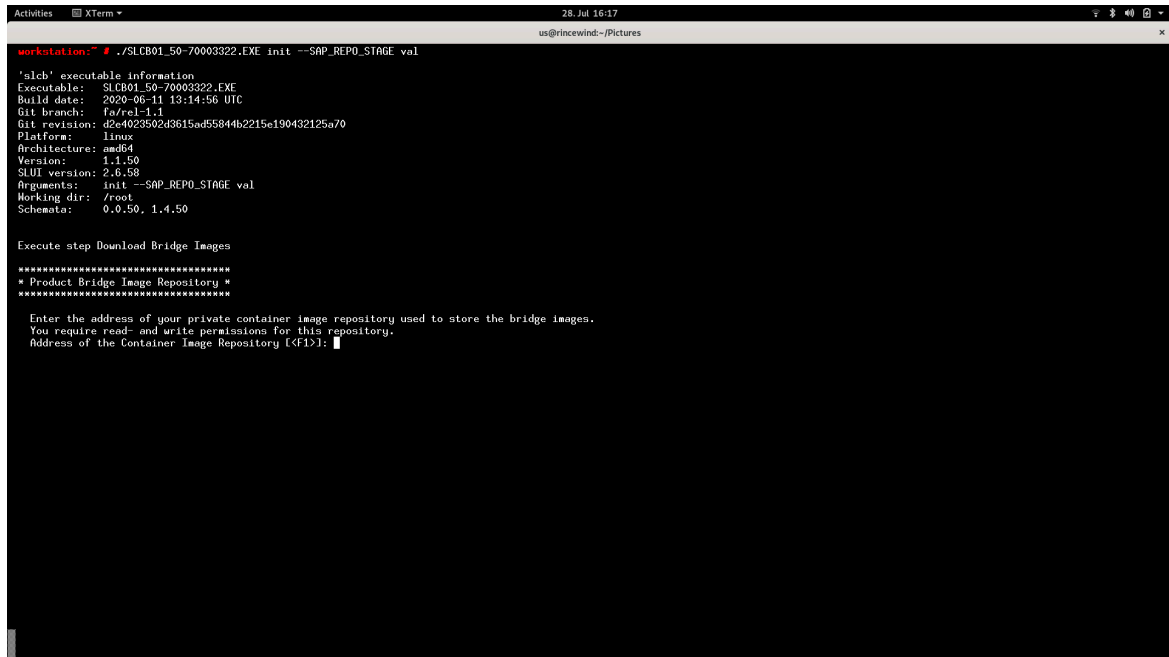
If not already downloaded, please download the below files:
Click [here](#) to download Stack XML
Click [here](#) to download SL CONTAINER BRIDGE 1.0.Linux on x86_64 64bit Database independent [SLCB01_50-70003322.EXE ,76422 KB]

FIGURE 48: SAP MAINTENANCE PLANNER: EXAMPLE FOR HOST AND PORT

6.3 Installing the SAP SLC Bridge

- Download the file containing the SLC Bridge Installer.
- If you have not already done so, copy this file to the management workstation.
- Run the **SLC Bridge Installer** on the management workstation.

```
$ ./SLCB01_<YOUR DOWNLOADED VERSION>.EXE init
```

A screenshot of a terminal window titled 'workstation:~' showing the execution of the SLC Bridge Installer. The command entered is './SLCB01_50-70003322.EXE init --SAP_REPO_STAGE val'. The output displays executable information, including build date (2020-06-11), Git branch (f6/rel-1.1), and platform (linux). It then prompts the user to enter the address of a private container image repository. The terminal output is as follows:

```
workstation:~ # ./SLCB01_50-70003322.EXE init --SAP_REPO_STAGE val

'slcb' executable information
Executable: SLCB01_50-70003322.EXE
Build date: 2020-06-11 13:14:56 UTC
Git branch: f6/rel-1.1
Git revision: d2e4023502d3615ad55844b2215e190432125a70
Platform: linux
Architecture: amd64
Version: 1.1.50
SLUI version: 2.6.58
Arguments: init --SAP_REPO_STAGE val
Working dir: /root
Schema: 0.0.50, 1.4.50

Execute step Download Bridge Images
*****
* Product Bridge Image Repository *
*****

Enter the address of your private container image repository used to store the bridge images.
You require read- and write permissions for this repository.
Address of the Container Image Repository [<F1>]:
```

FIGURE 49: SAP SLC BRIDGE INSTALLER

- This interactive script gathers all the necessary information to run the SAP SLC Bridge, and at the end, deploys it into the CaaSP cluster.

```

Activities  XTerm  28. Jul 16:25
us@rincewind:~/Pictures

Execute step Wait for Kubernetes Object SLCBridgeRoleBindingDeployment
Execute step Wait for Kubernetes Object DefaultsMap
Execute step Wait for Kubernetes Object ProductHistory
Execute step Wait for Kubernetes Object MasterSecret
Execute step Wait for Kubernetes Object NginxConf
Execute step Wait for Kubernetes Object NginxSecret
Execute step Wait for Kubernetes Object SLCBridgePod
Execute step Show availability of SLC Bridge Base deployment and ip address of SLC Bridge Base service
Execute step SL Container Bridge
*****
*      Message      *
*****

Deployment "slcbridgebase" has 1 available replicas in namespace "sap-slcbridge"
Service slcbridgebase-service is listening on any of the kubernetes nodes on "https://node:30616/docs/index.html"

Choose action Next In[<F1>]: n
*****
* Provide feedback to SAP SE *
*****

Dear user, please help us improve our software by providing your feedback (press <F1> for more information).

> 1. Fill out questionnaire
> 2. Send analytics data only
> 3. No feedback
possible values {1,2,3} [<F1>]: 3

Execute step Show availability of SLC Bridge Base deployment and ip address of SLC Bridge Base service
*****
*      Information    *
*****

Execution finished successfully

Choose action Exit [e]: e

```

FIGURE 50: SAP SLC BRIDGE INFORMATION

- Identify the service port for the SLC Bridge:

```

# kubectl -n sap-slcbridge get pods
NAME                                READY   STATUS    RESTARTS   AGE
di-platform-full-product-bridge    2/2     Running   0           3d23h
slcbridgebase-858f895bd6-74gps      2/2     Running   1           10d
# kubectl -n sap-slcbridge get svc
NAME                                TYPE           CLUSTER-IP   EXTERNAL-IP
PORT(S)          AGE
di-platform-full-product-bridge-service  ClusterIP      10.101.2.224  <none>
9000/TCP         3d23h
slcbridgebase-service                    NodePort        10.97.79.26   <none>
9000:30778/TCP   10d

```

In our example, the port number on which the SLC Bridge listens is 30778.

- Make a note of this information. It is needed for the installation process via SAP Maintenance Planner.

```

Activities  XTerm  28. Jul 16:27
us@rincewind:~/Pictures

Execute step Wait for Kubernetes Object MasterSecret
Execute step Wait for Kubernetes Object NginxConf
Execute step Wait for Kubernetes Object NginxSecret
Execute step Wait for Kubernetes Object SLCBridgePod
Execute step Show availability of SLC Bridge Base deployment and ip address of SLC Bridge Base service
Execute step SL Container Bridge
*****
*      Message      *
*****

Deployment "slcbridgebase" has 1 available replicas in namespace "sap-slcbridge"
Service slcbridgebase-service is listening on any of the kubernetes nodes on "https://node:30616/docs/index.html"

Choose action Next [n/<F1>]: n
*****
* Provide feedback to SAP SE *
*****

Dear user, please help us improve our software by providing your feedback (press <F1> for more information).

> 1. Fill out questionnaire
> 2. Send analytics data only
> 3. No feedback
possible values {1,2,3} [<F1>]: 3

Execute step Show availability of SLC Bridge Base deployment and ip address of SLC Bridge Base service
*****
*      Information      *
*****

Execution finished successfully

Choose action Exit [e]: e
workstation: # kubectl get pod -n sap-slcbridge
NAME                                READY  STATUS   RESTARTS  AGE
slcbridgebase-858f895bd6-nj56b      2/2    Running  0          2m3s
workstation: # kubectl get svc -n sap-slcbridge
NAME                                TYPE        CLUSTER-IP    EXTERNAL-IP  PORT(S)    AGE
slcbridgebase-service               NodePort    10.100.207.95 <none>       9000:30616/TCP  2m43s
workstation: #

```

FIGURE 51: SAP SLC BRIDGE

6.4 Installing SAP Data Intelligence 3

This section describes the installation of SAP Data Intelligence 3 on top of SUSE CaaSP 4.2.

6.4.1 Preparations

Before the installation of SAP DI 3 can start, some preparation work must be done:

- Create a namespace on the Kubernetes cluster.
- Define a (default) storage class on the Kubernetes cluster.
- Adapt PodSecurityPolicies.
- Create the necessary ClusterRoleBindings.
- If you use self-signed certificates for the private registry, a special secret must be created.



Note

Unless otherwise specified, all these tasks are run from the management workstation.

6.4.1.1 Creating namespace for SAP Data Intelligence 3 on Kubernetes

Define the namespace into which SAP Data Intelligence 3 will be installed:

```
$ kubectl create namespace <YOUR NAMESPACE>
```

6.4.1.2 Creating Storage Class

- Create the storage class to provide volumes for SAP Data Intelligence 3 on SUSE Enterprise Storage.
- Make sure that you have:
 - the connection data for your SUSE Enterprise Storage at hand
 - the IP addresses and port number (default: 6789) of the monitor nodes of your SES cluster
 - created a data pool on your SES cluster for use with SAP Data Intelligence 3
 - the name of this pool (datahub in this example) available
- Edit the example below to fit your environment.

```
$ cat > storageClass.yaml <<EOF
apiVersion: storage.kubernetes.io/v1
kind: StorageClass
metadata:
  annotations:
    storageclass.kubernetes.io/is-default-class: "true"
  name: datahub
  namespace: default
parameters:
  adminId: admin
  adminSecretName: ceph-admin-secret
  adminSecretNamespace: default
  imageFeatures: layering
  imageFormat: "2"
  monitors: <IP ADDRESS OF MONITOR 1>:6789, <IP ADDRESS OF MONITOR 2>:6789, <IP
ADDRESS OF MONITOR 3 >:6789
  pool: datahub
  userId: admin
  userSecretName: ceph-user-secret
provisioner: kubernetes.io/rbd
reclaimPolicy: Delete
```

```
volumeBindingMode: Immediate
EOF

$ kubectl create -f storageClass.yaml
```

- Create secrets for the StorageClass.
 - Create the secrets needed to access the storage.
 - Obtain the keys from your SES cluster. These are located in ceph.admin.keyring and ceph.user.keyring.
You must encode the keys with base64.

```
$ echo <YOUR KEY HERE> | base64
```

```
$ cat > ceph-admin-secret.yaml <<EOF
apiVersion: v1
kind: Secret
metadata:
  name: ceph-admin-secret
type: "kubernetes.io/rbd"
data:
  key: <YOUR BASE64 ENCODED KEY HERE>
EOF
```

```
$ cat > ceph-user-secret.yaml <<EOF
apiVersion: v1
kind: Secret
metadata:
  name: ceph-user-secret
type: "kubernetes.io/rbd"
data:
  key: <YOUR BASE64 ENCODED KEY HERE>
EOF

$ kubectl apply -f ceph-admin-secret.yaml
$ kubectl apply -f ceph-user-secret.yaml
```

- Create the credentials for accessing the StorageClass from the namespace where DI 3 will be installed into.

```
$ kubectl -n <YOUR NAMESPACE FOR DI 3> create -f ceph-admin-secret.yaml
$ kubectl -n <YOUR NAMESPACE FOR DI 3> create -f ceph-user-secret.yaml
```

6.4.1.3 Creating PodSecurityPolicies and ClusterRoleBindings

- PodSecurityPolicies

```
$ kubectl edit psp suse.caasp.psp.privileged
```

Change the `pathPrefix` in `allowedHostPaths` to `/`

- ClusterRoleBindings

```
$ cat > clusterrolebinding.yaml << EOF
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: suse:caasp:psp:privileged:default
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: suse:caasp:psp:privileged
subjects:
- kind: ServiceAccount
  name: default
  namespace: XXX
- kind: ServiceAccount
  name: vora-vsystem-XXX
  namespace: XXX
- kind: ServiceAccount
  name: vora-vsystem-XXX-vrep
  namespace: XXX
- kind: ServiceAccount
  name: XXX-elasticsearch
  namespace: XXX
- kind: ServiceAccount
  name: XXX-fluentd
  namespace: XXX
- kind: ServiceAccount
  name: XXX-nodeexporter
  namespace: XXX
- kind: ServiceAccount
  name: vora-vflow-server
  namespace: XXX
- kind: ServiceAccount
  name: mlf-deployment-api
  namespace: XXX
EOF
$ sed -i s/XXX/<your-di-namespace>/g clusterrolebinding.yaml
$ kubectl apply -f clusterrolebinding.yaml
```


- Additional changes

```
$ kubectl edit clusterrolebinding system:node
```

Insert the following at the end of the file:

```
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: Group
  name: system:nodes
```

- If you use self-signed SSL certificates for the secure private registry, create a secret for accessing this registry.



Note

The certificate chain should be saved in pem format into a single file called cert.

```
export NAMESPACE=<your namespace>
mv cert cert_with_carriage_return
tr -d '\r' < cert_with_carriage_return > cert
kubectl create secret generic cmcertificates --from-file=cert -n $NAMESPACE
```

6.4.2 Installing SAP Data Intelligence 3

After you successfully finished the preparation stages, proceed with the installation of SAP Data Intelligence 3. To do so, several steps must be executed.

6.4.2.1 Connecting to the SLC Bridge

- Point your browser to the SLC Bridge service:
 - <https://<yournode>:<yourport>/docs/index.html>
- Fill in your credentials you created during the deployment of the SLC Bridge.

6.4.2.2 Installing the workflow

- Connect and authenticate to the SLC Bridge service created above. Use the credentials created during the setup of the SLC Bridge:




FIGURE 52: CONNECT AND AUTHENTICATE TO THE SLCB SERVICE

- Select "Planned software Changes":




FIGURE 53: INSTALLING PLANNED SOFTWARE CHANGES

- Select the SAP Data Intelligence deployment required by your needs. Click "Next":

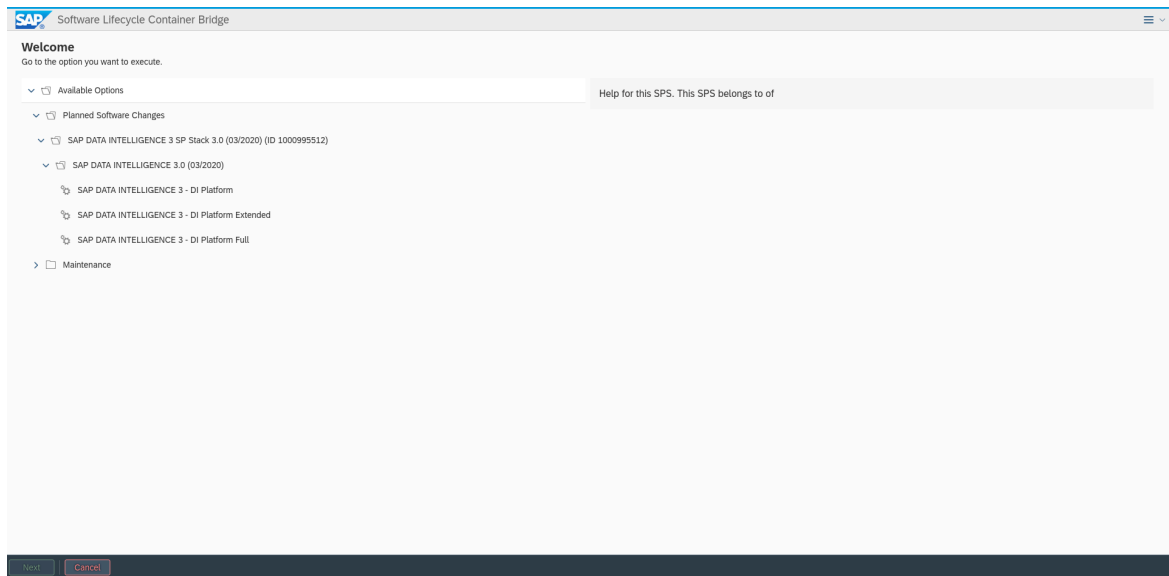


FIGURE 54: INSTALLING THE REQUIRED SAP DI DEPLOYMENT

- Enter the Kubernetes namespace created beforehand, for example di310. Click "Next":

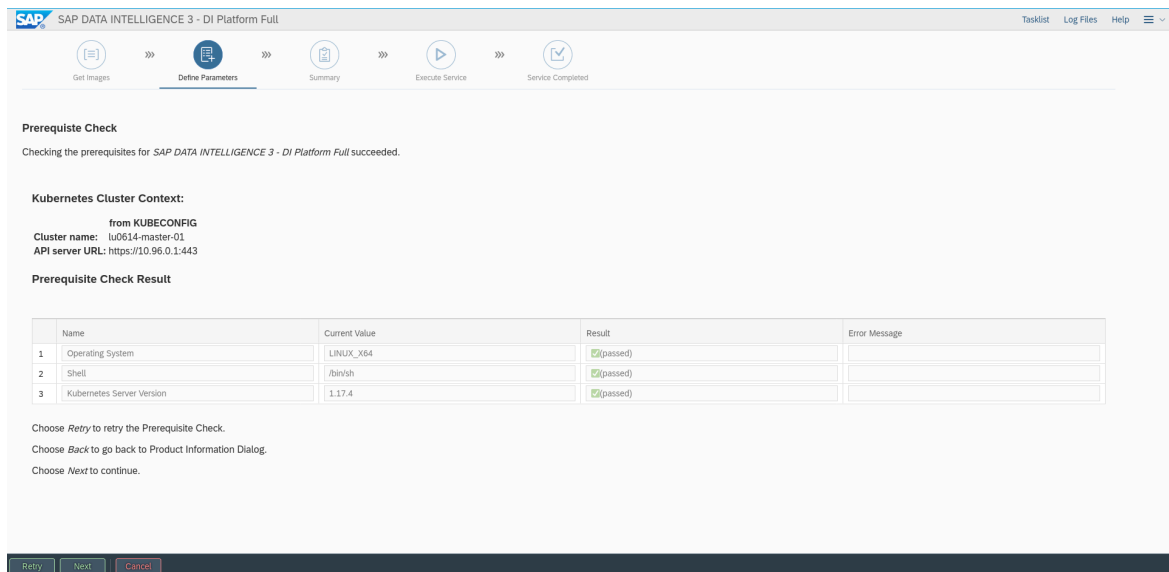


FIGURE 55: INSTALLING PRE-REQUISITES CHECK

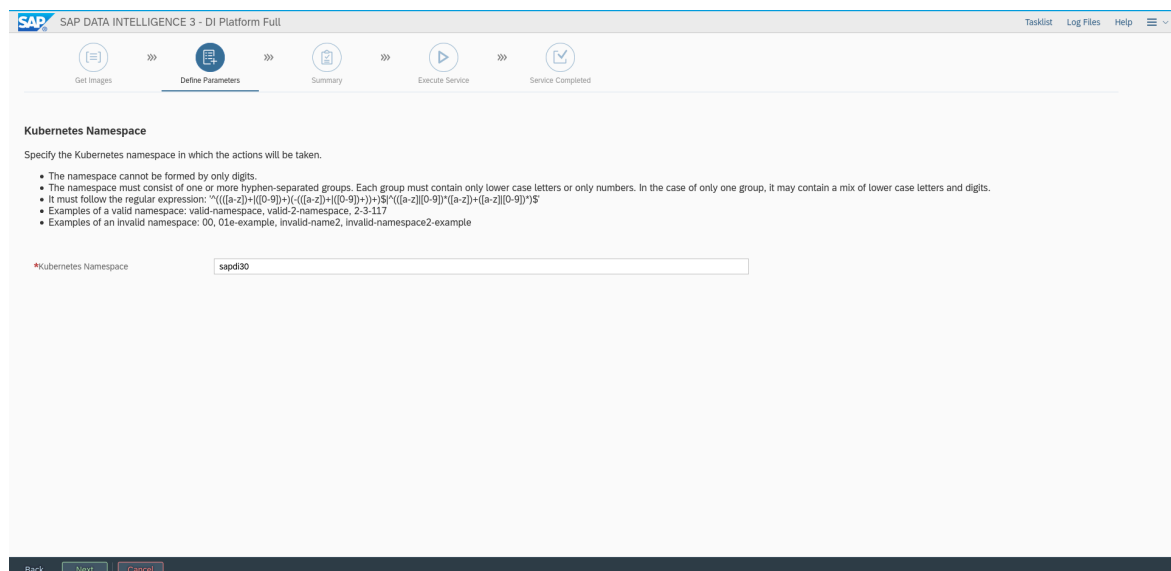


FIGURE 56: INSTALLATION ENTER KUBERNETES NAMESPACE

- Select "Advanced Installation":

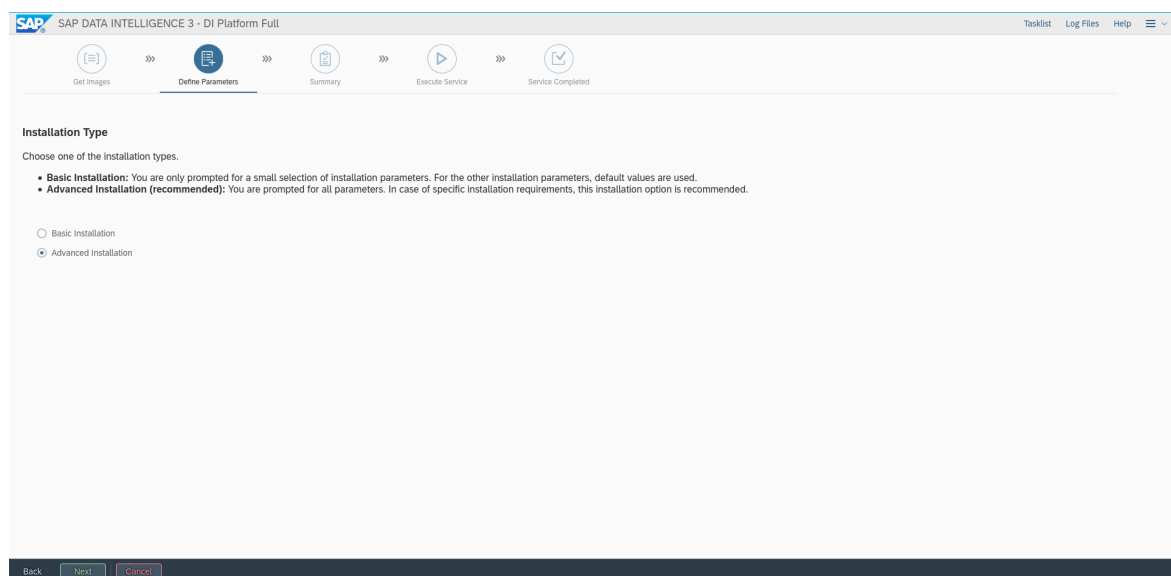


FIGURE 57: INSTALLING ADVANCED INSTALLATION

- Enter the URI of your Private secure registry. Click "Next":

The screenshot shows the 'Define Parameters' step of the SAP Data Intelligence 3 - DI Platform Full installation wizard. The progress bar at the top indicates the current step. The main content area is titled 'Container Image Repository' and contains the following text: 'Specify the container image repository to push the SAP Data Intelligence images. This container image repository will be used by Kubernetes and by SAP Data Intelligence Modeler. The container image repository must be accessible from the installation host including the necessary authentication.' Below this text, there are examples: '012345678910.dkr.ecr.us-east-1.amazonaws.com, eu.gcr.io/my-project-name, myregistry.azurecr.io, myhost:5000'. A text input field labeled 'Container Image Repository' contains the value 'lu0591.wdf.sap.corp:5000'. At the bottom, there are 'Back', 'Next', and 'Cancel' buttons.

FIGURE 58: INSTALLING PRIVATE CONTAINER REGISTRY

- Enter a password for the system tenant in SAP DI 3.0:

The screenshot shows the 'Define Parameters' step of the SAP Data Intelligence 3 - DI Platform Full installation wizard. The progress bar at the top indicates the current step. The main content area is titled 'SAP Data Intelligence System Tenant Administrator Password' and contains the following text: 'Specify a password for the "system" user of "system" tenant.' Below this text, there are requirements: 'It must contain between 8 and 255 characters.', 'It must contain at least one lower case, one upper case, one numerical and one special character.', 'The allowed special characters are . @ # \$ % * + _ ? ! .', and 'It cannot contain spaces.' There are two text input fields: one labeled '*Password' and one labeled '*Confirm'. At the bottom, there are 'Back', 'Next', and 'Cancel' buttons.

FIGURE 59: ENTER SYSTEM TENANT PASSWORD

- Assign a name to the initially created tenant, for example "default":

The screenshot shows the SAP Data Intelligence 3 - DI Platform Full interface. The top navigation bar includes the SAP logo, the title 'SAP DATA INTELLIGENCE 3 - DI Platform Full', and links for 'Tasklist', 'Log Files', and 'Help'. Below the navigation bar is a progress bar with five steps: 'Get Images', 'Define Parameters' (current step), 'Summary', 'Execute Service', and 'Service Completed'. The main content area is titled 'SAP Data Intelligence Initial Tenant Name'. It contains the instruction: 'Specify a name for the SAP Data Intelligence initial tenant that is going to be created automatically.' followed by a list of requirements:

- It must contain between 4 and 60 characters.
- It must consist of lower case letters, digits or hyphens.
- It must not begin or end with hyphens and must not contain multiple consecutive hyphens.
- It must follow the regular expression: "[a-z0-9]+(-[a-z0-9]+)*\$"

 Below the requirements is a text input field labeled 'Tenant Name' with the value 'default' entered. At the bottom of the form are three buttons: 'Back', 'Next', and 'Cancel'.

FIGURE 60: ASSIGN A NAME TO CREATED DEFAULT TENANT

- Create the administrator user for the default tenant in SAP DI 3.0:

The screenshot shows the SAP Data Intelligence 3 - DI Platform Full interface. The top navigation bar includes the SAP logo, the title 'SAP DATA INTELLIGENCE 3 - DI Platform Full', and links for 'Tasklist', 'Log Files', and 'Help'. Below the navigation bar is a progress bar with five steps: 'Get Images', 'Define Parameters' (current step), 'Summary', 'Execute Service', and 'Service Completed'. The main content area is titled 'SAP Data Intelligence Initial Tenant Administrator Username'. It contains the instruction: 'Specify a name for administrator user of "default" tenant.' followed by a list of requirements:

- It must contain between 4 and 60 characters.
- It must consist of lower case letters, digits or hyphens.
- It must not begin or end with hyphens and must not contain multiple consecutive hyphens.
- It must follow the regular expression: "[a-z0-9]+(-[a-z0-9]+)*\$"

 Below the requirements is a text input field labeled 'Username' which is currently empty. At the bottom of the form are three buttons: 'Back', 'Next', and 'Cancel'.

FIGURE 61: CREATE ADMIN USER NAME OF DEFAULT TENANT

- Set the password for the administrator user of the default tenant:

SAP DATA INTELLIGENCE 3 - DI Platform Full

Tasklist Log Files Help

Get Images Define Parameters Summary Execute Service Service Completed

SAP Data Intelligence Initial Tenant Administrator Password

Specify a password for "admin" user of "default" tenant.

- It must contain between 8 and 255 characters.
- It must contain at least one lower case, one upper case, one numerical and one special character.
- The allowed special characters are . @ # \$ % ^ + _ ? !.
- It cannot contain spaces.

*Password

*Confirm

Back Next Cancel

FIGURE 62: SET ADMIN USER PASSWORD

- If you need a proxy to connect to the Internet, set the proxy settings accordingly:

SAP DATA INTELLIGENCE 3 - DI Platform Full

Tasklist Log Files Help

Get Images Define Parameters Summary Execute Service Service Completed

Cluster Proxy Settings

Choose if you want to configure proxy settings. It is necessary when the Kubernetes cluster is running behind a proxy.

☐ Choose if you want to configure proxy settings on the cluster

Back Next Cancel

FIGURE 63: SET PROXY SETTINGS

- Select if you want to use a checkpoint storage:

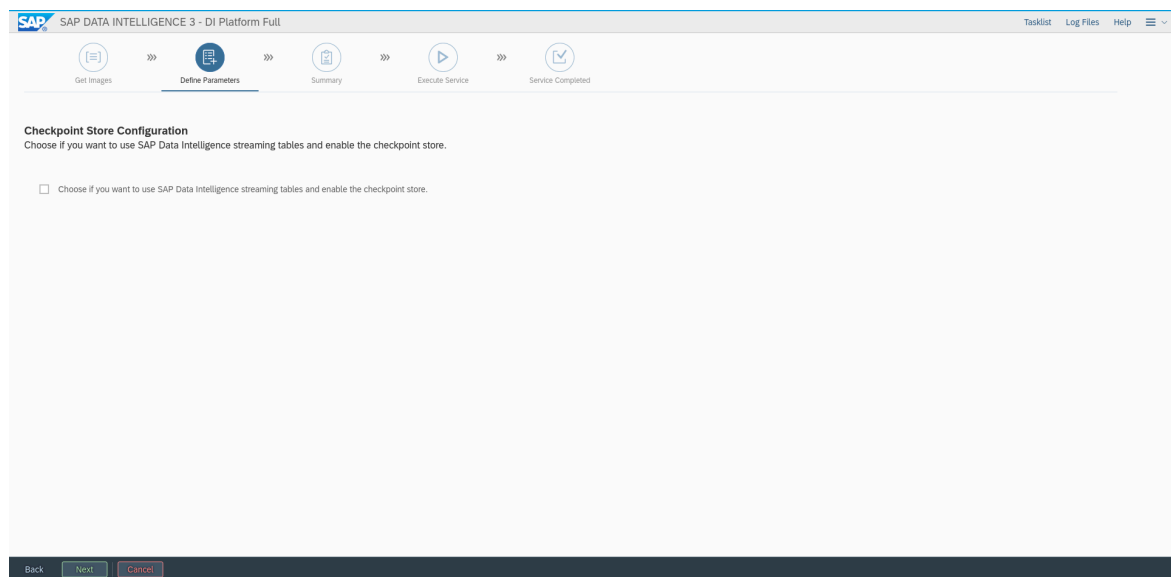


FIGURE 64: CONFIGURE CHECKPOINT STORAGE

- Define the storage class that should be used by SAP DI 3.0. Enter the name of the storage class you created previously:

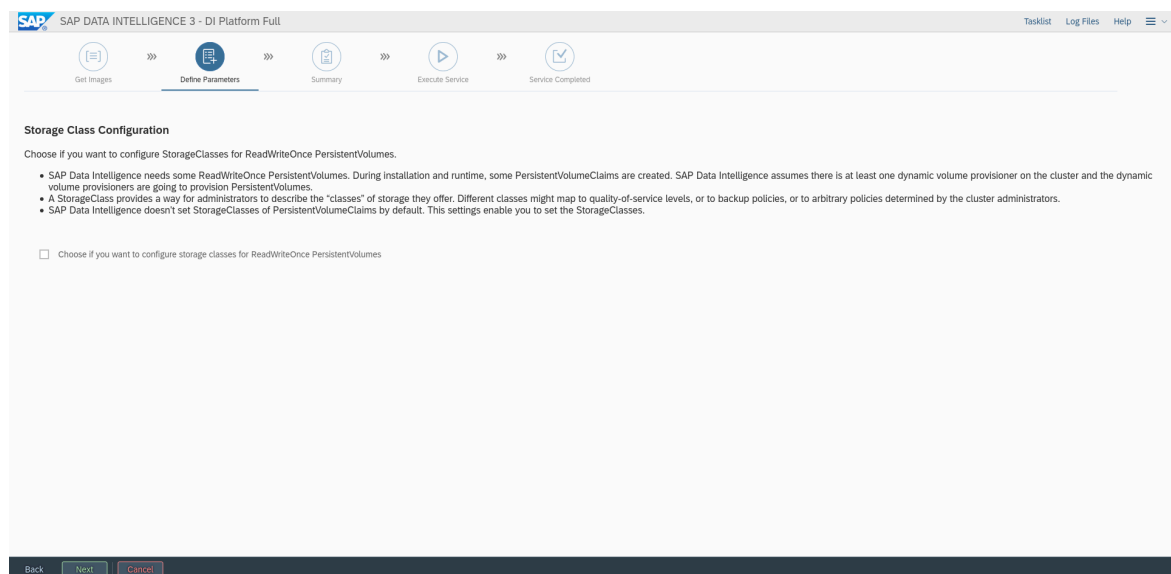


FIGURE 65: DEFINE STORAGE CLASS TO BE USED

- For the SAP DI 3 installation on SUSE CaaSP 4.2 the docker log path needs to be adapted. Check the check box and click "Next":




FIGURE 66: ADAPT DOCKER LOG PATH

- Enter the docker log path: `"/var/log/containers"`

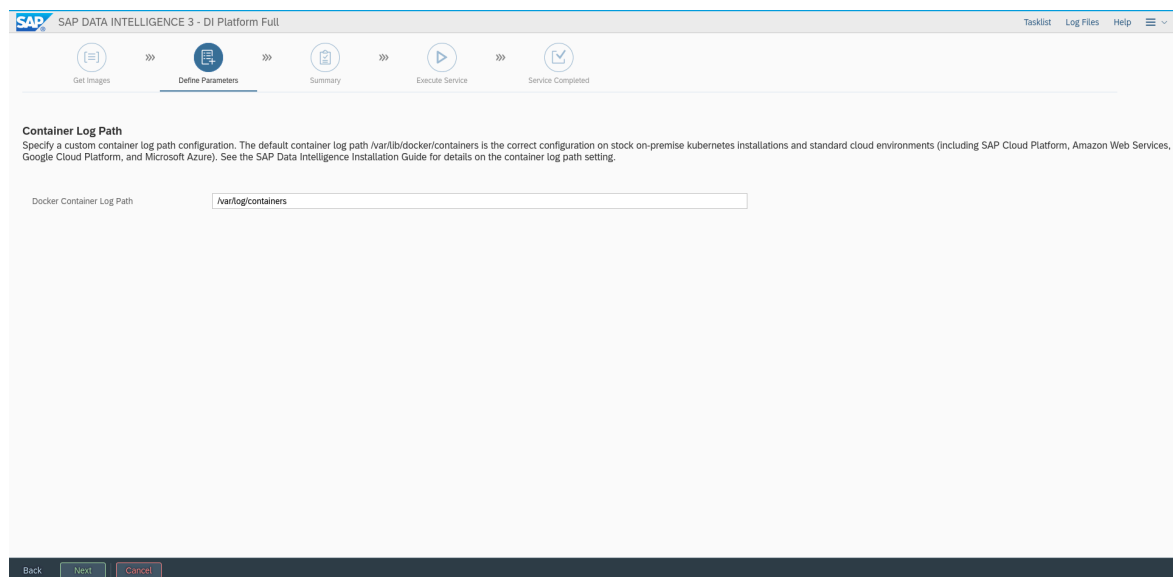


FIGURE 67: ENTER DOCKER LOG PATH

- Enable Kaniko:

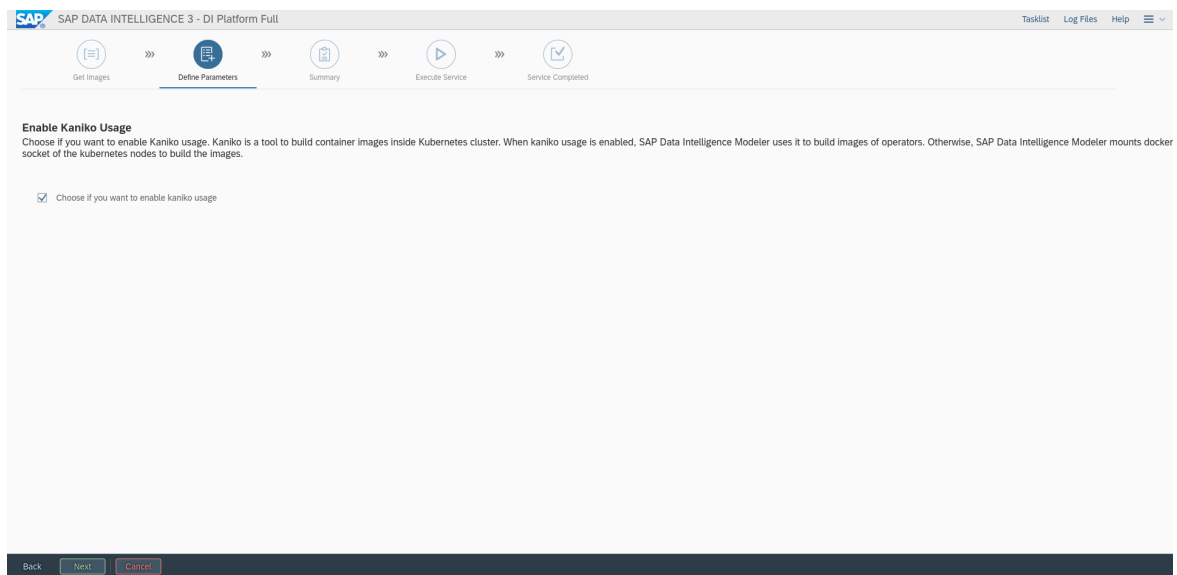


FIGURE 68: ENABLE KANIKO

- Here, a different private registry can be configured if needed. To proceed, click "Next":




FIGURE 69: CONFIGURE DOCKER REGISTRY FOR SAP DI MODELER IMAGES

- Enable the loading of NFS kernel modules. This ensures that the NFS kernel modules are loaded on all Kubernetes nodes. Check the check box and click "Next":

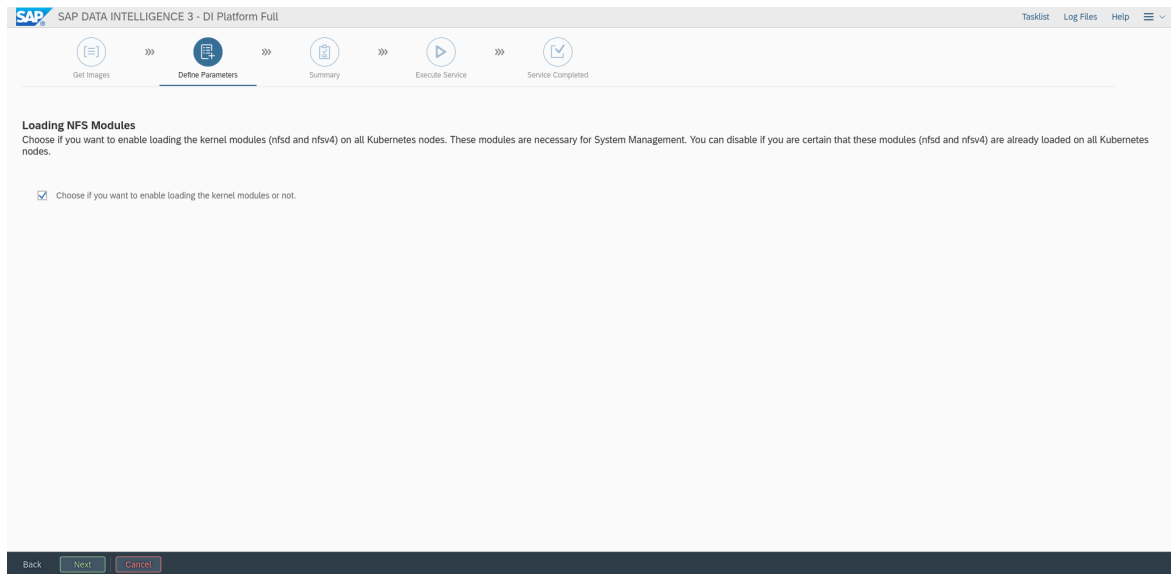


FIGURE 70: LOADING NFS KERNEL MODULES

- If needed, enable Network policies. Click "Next":




FIGURE 71: ENABLE NETWORK POLICIES

- Configure timeout during installation, leave the default and click "Next":

SAP DATA INTELLIGENCE 3 - DI Platform Full

Tasklist Log Files Help

Get Images Define Parameters Summary Execute Service Service Completed

Timeout
Timeout in seconds for DataHub installation. The default duration is enough for most of the environments. Increasing the value is necessary when the network or volume provisioner is so slow that deployments of components fail because of timeouts without any other issues.

*Timeout in seconds 3600

Back Next Cancel

FIGURE 72: CONFIGURING TIMEOUT DURING INSTALLATION OF SAP DI 3

- In the field "Additional Installation Parameters", enter:
`-e diagnostic.fluentd.logDriverFormat=regex -e vsystem.vRep.exportsMask=true`

SAP DATA INTELLIGENCE 3 - DI Platform Full

Tasklist Log Files Help

Get Images Define Parameters Summary Execute Service Service Completed

Additional Installation Parameters
You can specify additional installation parameters. The parameters are documented in the section "Configuration Parameters for Kubernetes Deployment" in the official SAP Data Intelligence documentation. Use -e flag for each additional parameter that you give and use spaces between them.

• Example: -e hana.memoryRequest=7Gi -e storageGateway.replicas=1

Additional Installation Parameters -e diagnostic.fluentd.logDriverFormat=regex -e vsystem.vRep.exportsMask=true into

Back Next Cancel

FIGURE 73: ADDITIONAL INSTALLATION PARAMETERS 1

SAP DATA INTELLIGENCE 3 - DI Platform Full

Tasklist Log Files Help

Get Images Define Parameters Summary Execute Service Service Completed

Enter Logon Information for SAP Support Portal Access

Enter S-User credentials to log on to the SAP Support Portal and upload system identity information.

If there is no Internet connectivity to the SAP Support Portal, leave the Username empty. In this case you can download the *systemidentity.xml* file from the *Log Files* menu and upload it manually to the Maintenance Planner.

User Name

Password

Back Next Cancel

FIGURE 74: ADDITIONAL INSTALLATION PARAMETERS 2

- Check the Summary page. Check if the settings are correct:

SAP DATA INTELLIGENCE 3 - DI Platform Full

Tasklist Log Files Help

Get Images Define Parameters Summary Execute Service Service Completed

Parameter Summary

Choose Next to start the deployment with the displayed parameter values or choose Back to revise the parameters.

Kubernetes Namespace

Kubernetes Namespace

Installation Type

☐ Basic Installation

☒ Advanced Installation

Container Image Repository

Container Image Repository

SAP Data Intelligence System Tenant Administrator Password

Password

Confirm

Back Next Cancel

FIGURE 75: SUMMARY OF INSTALLATION PARAMETERS


- Click "Next" to start the deployment of SAP DI 3.0.

Your installation should now be finished.

7 Legal notice

Copyright © 2006–2025 SUSE LLC and contributors. All rights reserved.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or (at your option) version 1.3; with the Invariant Section being this copyright notice and license. A copy of the license version 1.2 is included in the section entitled "GNU Free Documentation License".

SUSE, the SUSE logo and YaST are registered trademarks of SUSE LLC in the United States and other countries. For SUSE trademarks, see <https://www.suse.com/company/legal/> .

Linux is a registered trademark of Linus Torvalds. All other names or trademarks mentioned in this document may be trademarks or registered trademarks of their respective owners.

Documents published as part of the SUSE Best Practices series have been contributed voluntarily by SUSE employees and third parties. They are meant to serve as examples of how particular actions can be performed. They have been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. SUSE cannot verify that actions described in these documents do what is claimed or whether actions described have unintended consequences. SUSE LLC, its affiliates, the authors, and the translators may not be held liable for possible errors or the consequences thereof.

Below we draw your attention to the license under which the articles are published.

8 GNU Free Documentation License

Copyright © 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition. The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.

- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all

Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

Copyright (c) YEAR YOUR NAME.

Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.2

```
or any later version published by the Free Software Foundation;  
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.  
A copy of the license is included in the section entitled "GNU  
Free Documentation License".
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “ with... Texts.” line with this:

```
with the Invariant Sections being LIST THEIR TITLES, with the  
Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.