

SAP Data Intelligence 3.1 on Rancher Kubernetes Engine 2

Installation Guide

SUSE Linux Enterprise Server for SAP Applications 15 SP2
Rancher Kubernetes Engine 2
SAP Data Intelligence 3.1

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

SAP Data Intelligence 3.1 on Rancher Kubernetes Engine 2

Installation Guide

Date: 2025-03-13

SAP Data Intelligence 3 is the tool set to govern big amount of data. Rancher Kubernetes Engine (RKE) 2 is the Kubernetes base that makes deploying SAP Data Intelligence 3 easy. This document describes the installation and configuration of RKE 2 from SUSE 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 Preparations 6
- 4 Installing Rancher Kubernetes Engine 2 cluster 6
- 5 Installing SAP Data Intelligence 3.1 26
- 6 Maintenance tasks 44
- 7 Legal notice 48
- 8 GNU Free Documentation License 49

1 Introduction

This guide describes the on-premises installation of SAP Data Intelligence 3.1 (SAP DI 3.1) on top of Rancher Kubernetes Engine (RKE) 2. In a nutshell, the installation of SAP DI 3.1 consists of the following steps:

- Installing SUSE Linux Enterprise Server 15 SP2
- Installing RKE 2 Kubernetes cluster on the dedicated nodes
- Deploying SAP DI 3.1 on RKE 2 Kubernetes cluster
- Performing post-installation steps for SAP DI 3.1
- Testing the installation of SAP DI 3.1

2 Prerequisites

2.1 Hardware requirements

This chapter describes the hardware requirements for installing SAP DI 3.1 on RKE 2 on top of SUSE Linux Enterprise Server 15 SP2. Only the AMD64/Intel 64 architecture is applicable for our use case.

2.1.1 Hardware Sizing

Correct hardware sizing is very important for setting up SAP DI 3.1 on RKE 2.

- Minimal hardware requirements for a generic SAP DI 3 deployment:
 - At least 7 nodes are needed for the Kubernetes cluster.
 - Minimum sizing of the nodes needs to be as shown below:

Server Role	Count	RAM	CPU	Disk space
Management Workstation	1	16 GiB	4	> 100 GiB
Master Node	3	16 GiB	4	> 120 GiB
Worker Node	4	32 GiB	8	> 120 GiB

- Minimal hardware requirements for an SAP DI 3 deployment for production use:
 - At least seven nodes are needed for the Kubernetes cluster.
 - Minimum sizing of the nodes needs to be as shown below:

Server Role	Count	RAM	CPU	Disk space
Management Workstation	1	16 GiB	4	> 100 GiB
Master Node	3	16 GiB	4	> 120 GiB
Worker Node	4	64 GiB	16	> 120 GiB

- For more information about the requirements for RKE, read the documentation at:
 - <https://docs.rke2.io/>
- For more detailed sizing information about SAP DI 3, read the "Sizing Guide for SAP Data Intelligence" at:
 - <https://help.sap.com/viewer/835f1e8d0dde4954ba0f451a9d4b5f10/3.1.latest/en-US>

2.2 Software requirements

The following list contains the software components needed to install SAP DI 3.1 on RKE:

- SUSE Linux Enterprise Server 15 SP2
- Rancher Kubernetes Engine 2
- SAP Software Lifecycle Bridge
- SAP Data Intelligence 3.1

- Secure private registry for container images, for example <https://documentation.suse.com/sbp/all/single-html/SBP-Private-Registry/index.html> ↗
- Access to a storage solution providing dynamically physical volumes
- If it is planned to use Vora's streaming tables checkpoint store, an S3 bucket like object store is needed
- If it is planned to enable backup of SAP DI 3.1 during installation access to an S3-compatible object store is needed

3 Preparations

- Get a SUSE Linux Enterprise Server subscription.
- Download the installer for SUSE Linux Enterprise Server 15 SP2.
- Download the installer for RKE 2.
- Check the storage requirements.
- Create a or get access to a private container registry.
- Get an SAP S-user to access software and documentation by SAP.
- Read the relevant SAP documentation:
 - Release Note for SAP DI 3 (<https://launchpad.support.sap.com/#/notes/2871970>) ↗
 - Release Note for SAP SLC Bridge (<https://launchpad.support.sap.com/#/notes/2589449>) ↗
 - Installation Guide at help.sap.com (<https://help.sap.com/viewer/a8d90a56d61a49718e-bcb5f65014bbe7/3.1.latest/en-US>) ↗

4 Installing Rancher Kubernetes Engine 2 cluster

The installation of Rancher Kubernetes Engine 2 cluster is straight forward. After the installation and basic configuration of the operating system the Kubernetes cluster configuration is created on the management host. Subsequently, the Kubernetes cluster will be deployed.

The following sections describe the installation steps in more detail.

4.1 Preparing management host and Kubernetes cluster nodes

All servers in our scenario use SUSE Linux Enterprise 15 SP2 (SLES 15 SP2) on the AMD64/Intel 64 architecture. The documentation for SUSE Linux Enterprise Server can be found at:

- <https://documentation.suse.com/sles/15-SP2/> 

4.1.1 Installing SUSE Linux Enterprise Server 15 SP2

On each server in your environment for SAP Data Intelligence 3.1, install SUSE Linux Enterprise Server 15 SP2 as the operating system. This chapter describes all recommended steps for the installation.



Tip

If you have already set up all machines and the operating system, skip this chapter and follow the instructions in [Section 4.2, "Configuring the Kubernetes nodes"](#).

- It is recommended to use a static network configuration. During the installation setup, the first time to adjust this is when the registration page is displayed. In the upper right corner, click the button "Network Configuration ...":

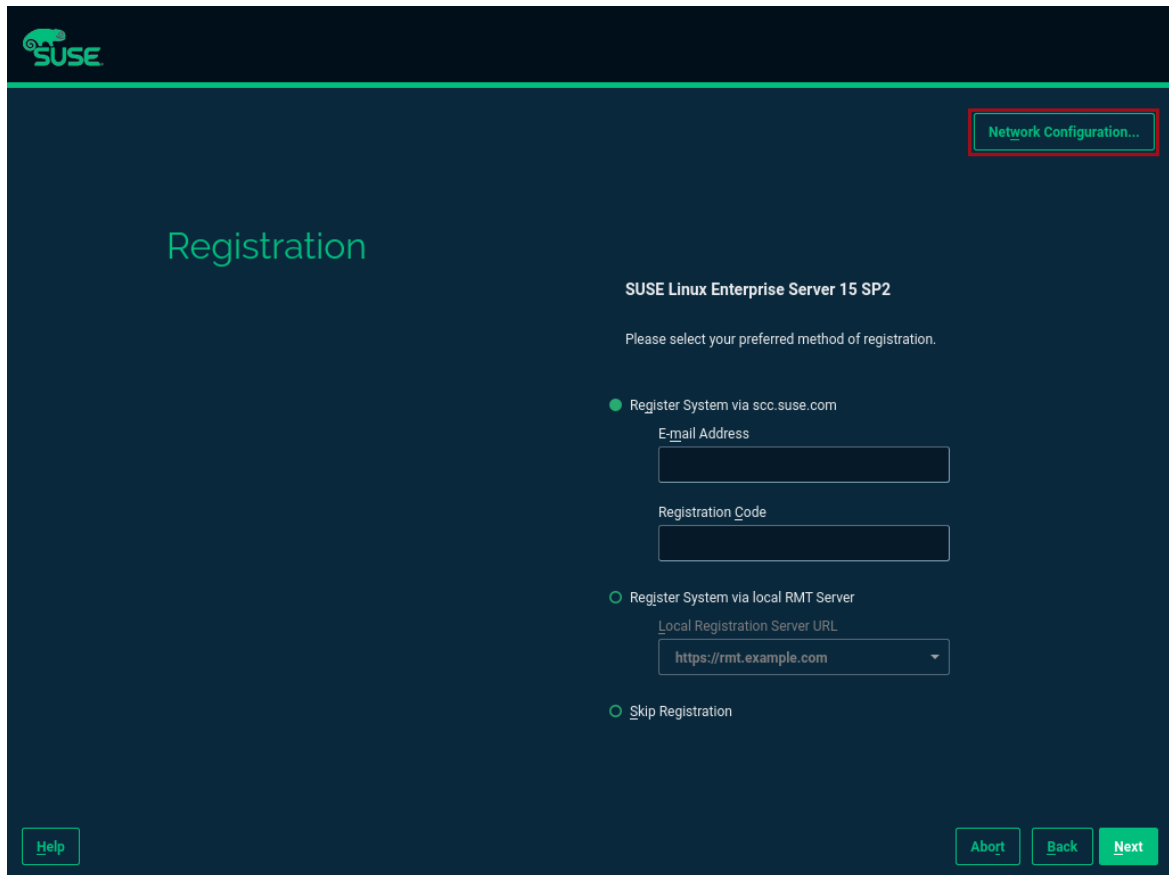


FIGURE 1: SLES SETUP REGISTRATION PAGE

- The **Network Settings** page is displayed. By default, the network adapter is configured to use DHCP. To change this, click the Button "Edit".

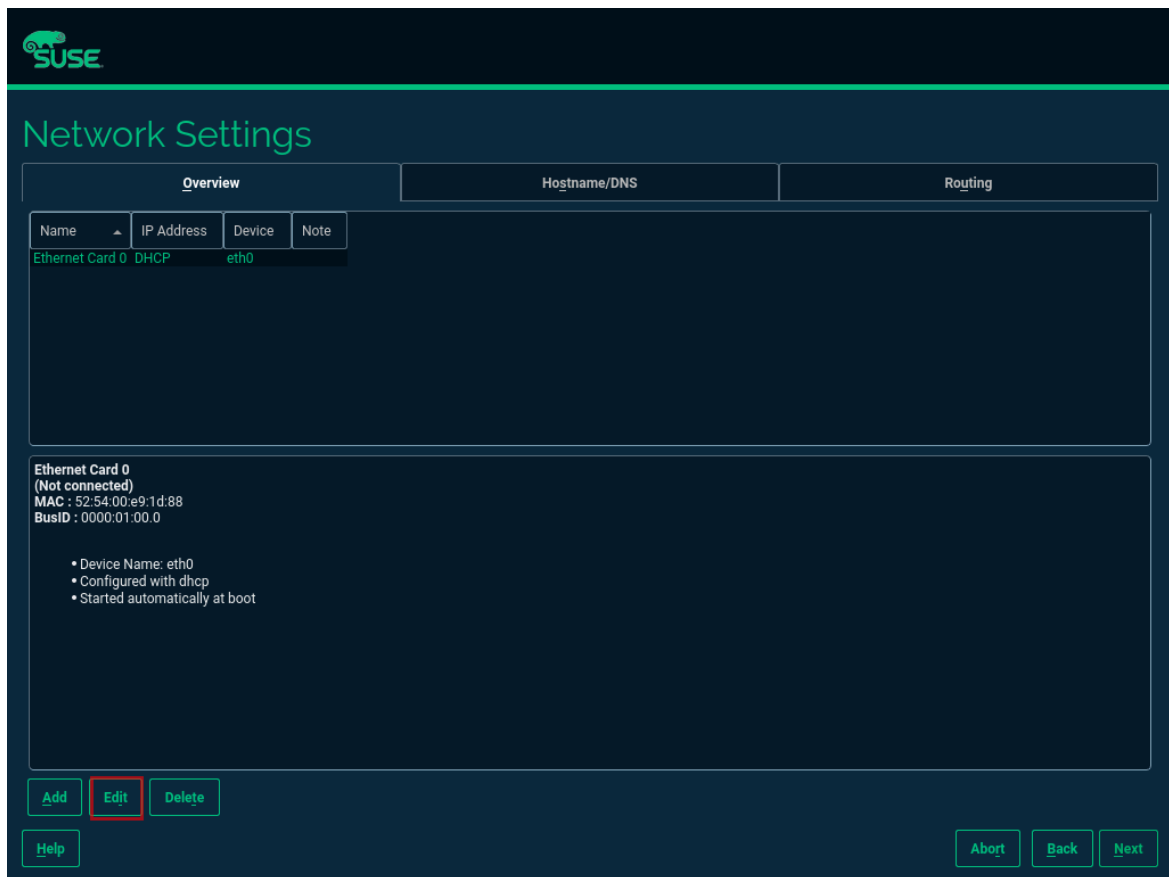


FIGURE 2: SLES SETUP NETWORK SETTINGS

- On the **Network Card Setup** page, select "Statically Assigned IP Address" and fill in the fields "IP Address", "Subnet Mask" and "Hostname".




FIGURE 3: SLES SETUP NETWORK CARD

- During the installation, you also need to adjust the extensions that need to be installed. The **Container Module** is needed to operate RKE 2.

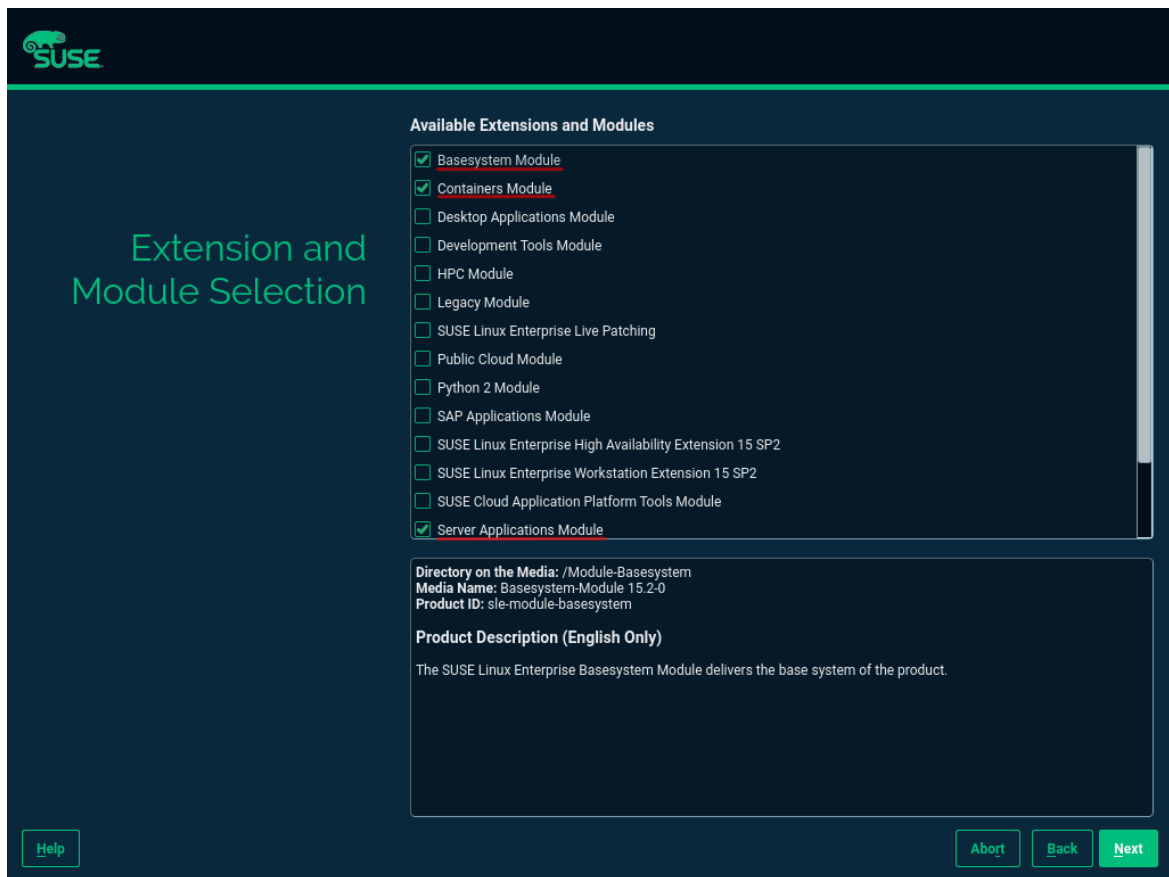


FIGURE 4: SLES SETUP EXTENSIONS

- As no graphical interface is needed, it is recommended to install a text-based server.

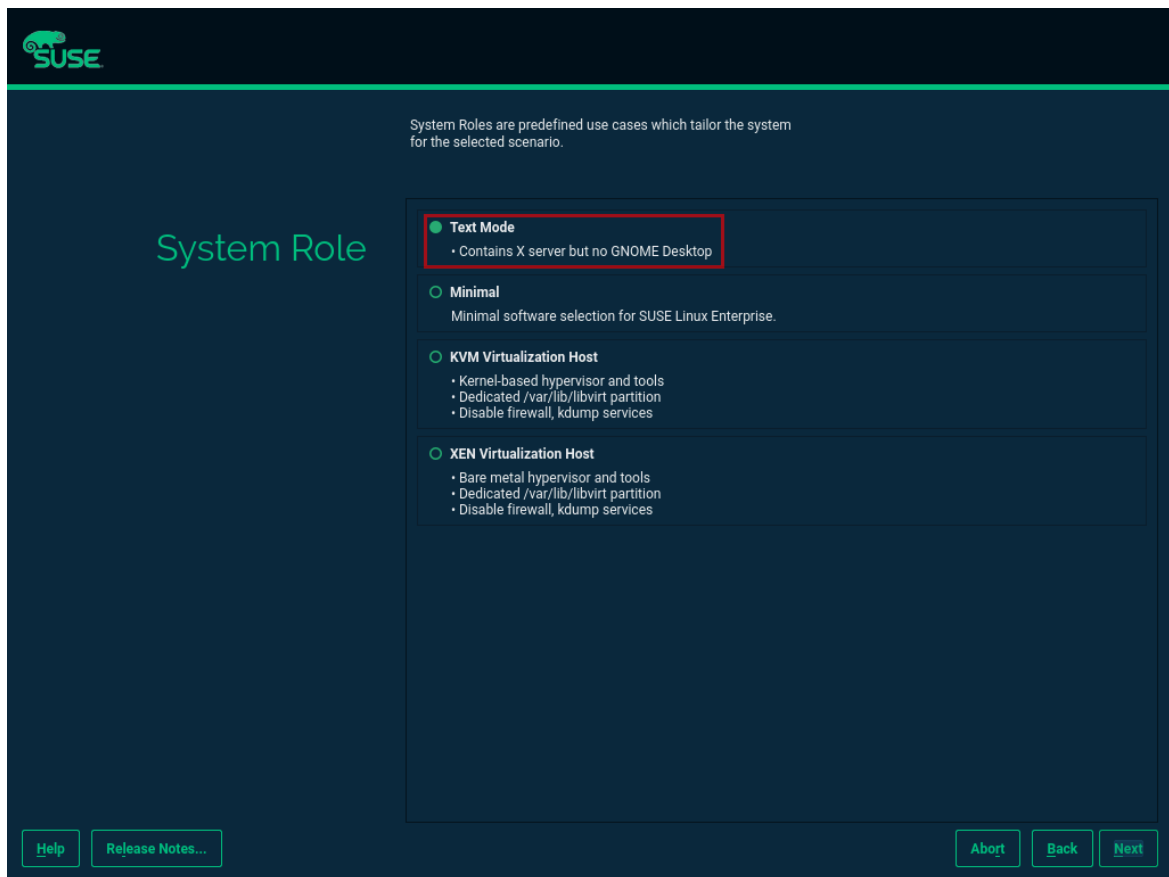


FIGURE 5: SLES SETUP SYSTEM ROLE

- To run Kubernetes the swap partition needs to be disabled. To do so, the partition proposal during installation can be adjusted.

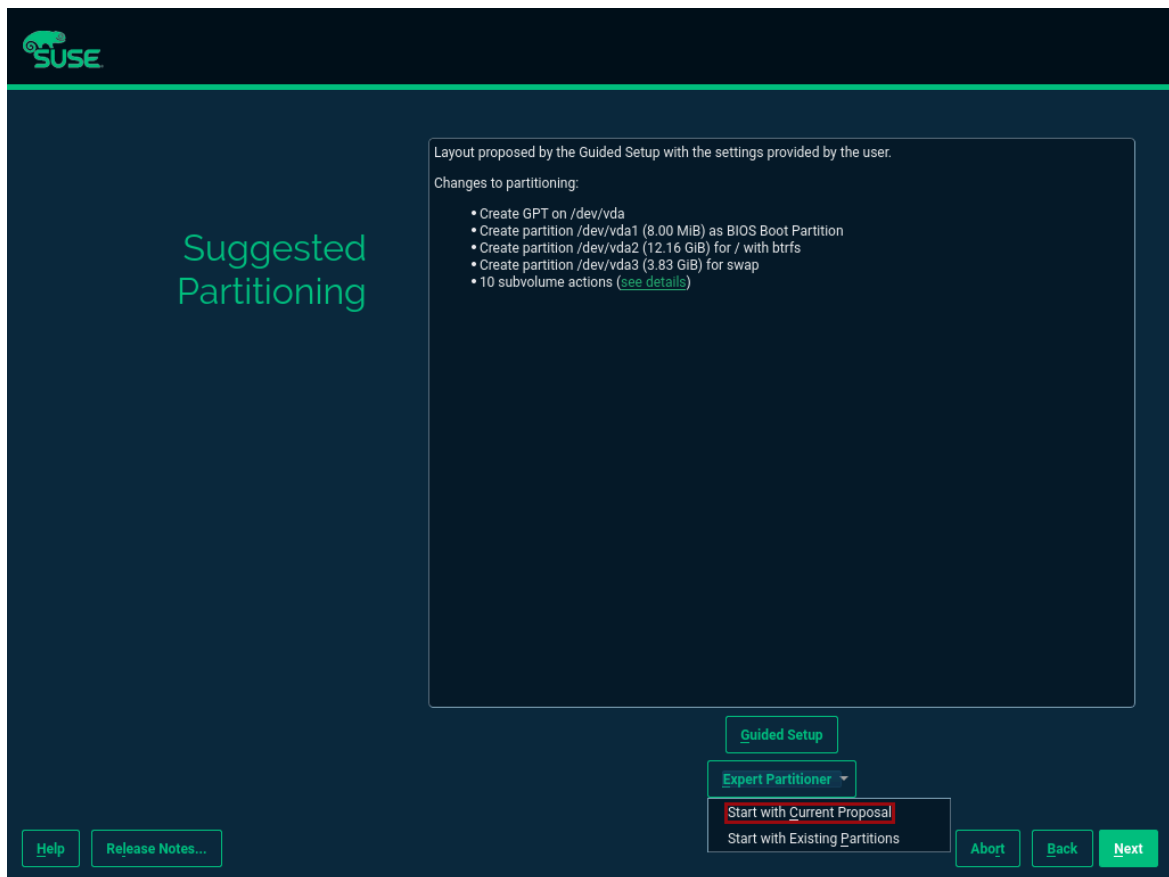


FIGURE 6: SLES SETUP PARTITIONING

- When opening the **Expert Partitioner**, the Swap partition needs to be selected to delete it.

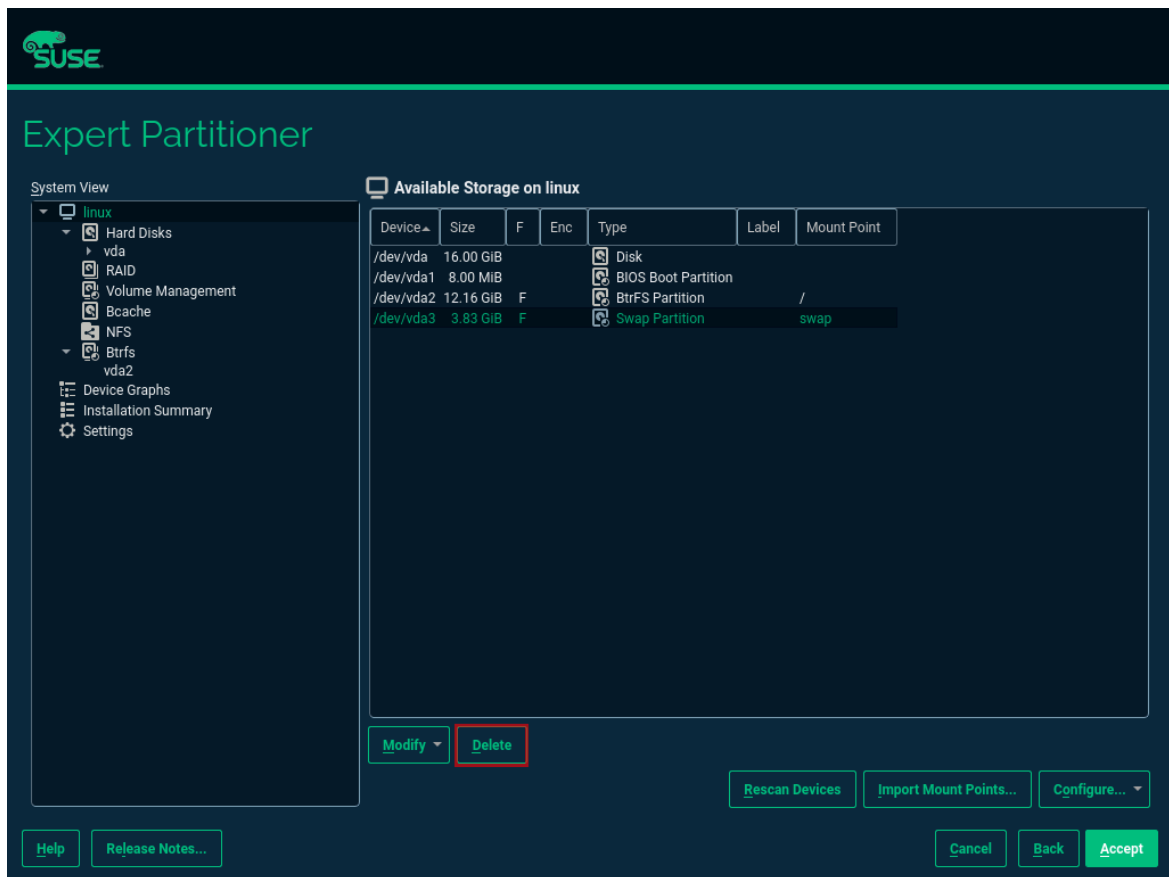


FIGURE 7: SLES SETUP EXPERT PARTITIONER SWAP

- After deleting the swap partition, there will be some space left that can be used to enlarge the main partition. To do so, the resize page can be called.

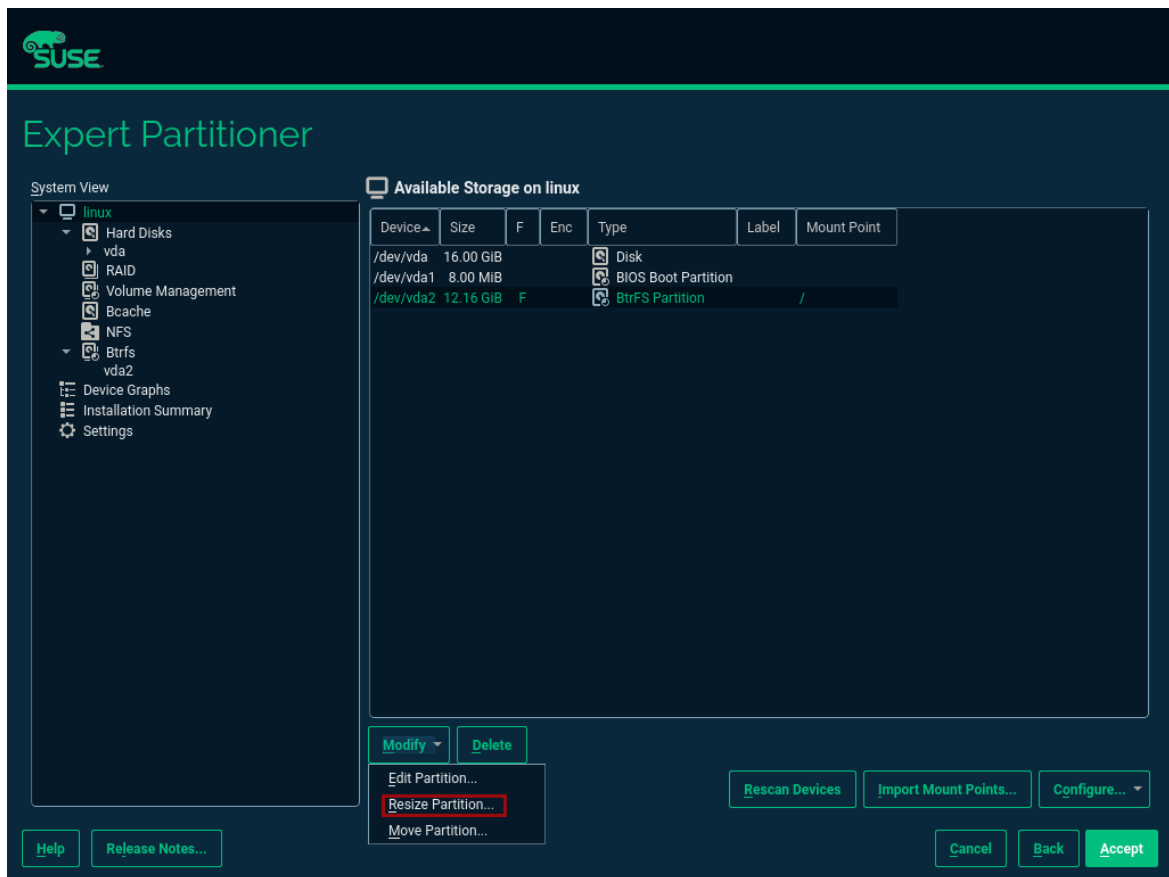


FIGURE 8: SLES SETUP EXPERT PARTITIONER RESIZE

- The easiest way to use all the unused space is to select the "Maximum Size" option here.

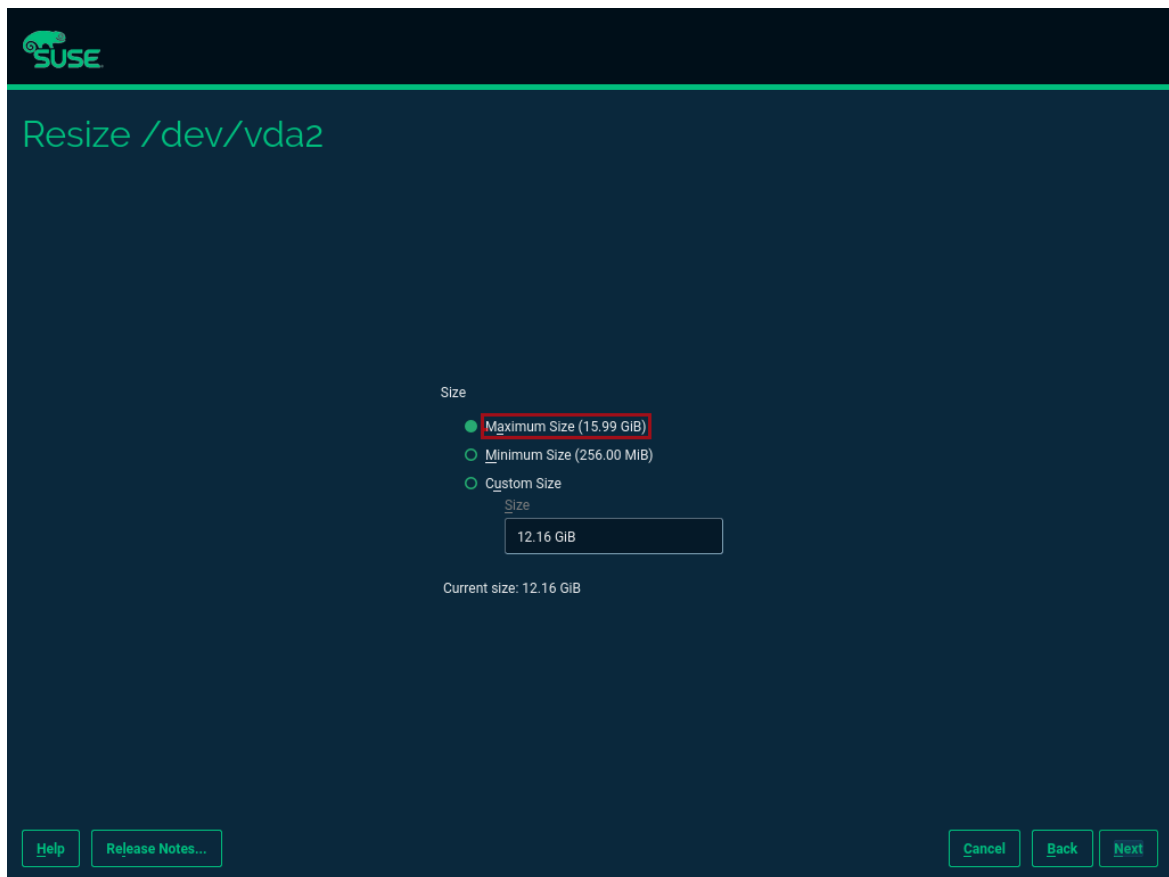


FIGURE 9: SLES SETUP RESIZE DISK

- Next, enable the NTP time synchronization. This can be done when the **Clock and Time Zone** page shows up during installation. To enable NTP, click the "Other Settings ..." button.

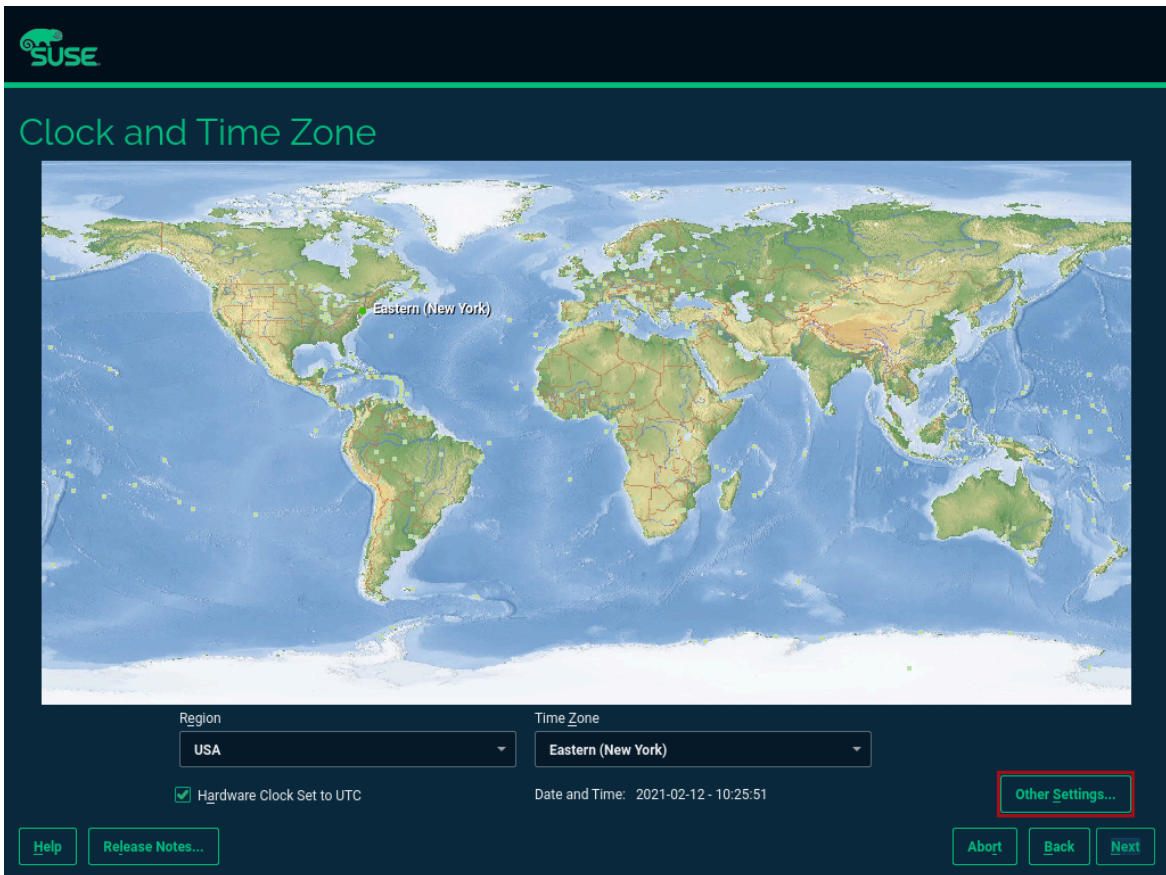


FIGURE 10: SLES SETUP TIMEZONE

- Select the "Synchronize with NTP Server" option. A custom NTP server address can be added if desired. Ensure to mark the check boxes for "Run NTP as daemon" and "Save NTP Configuration".

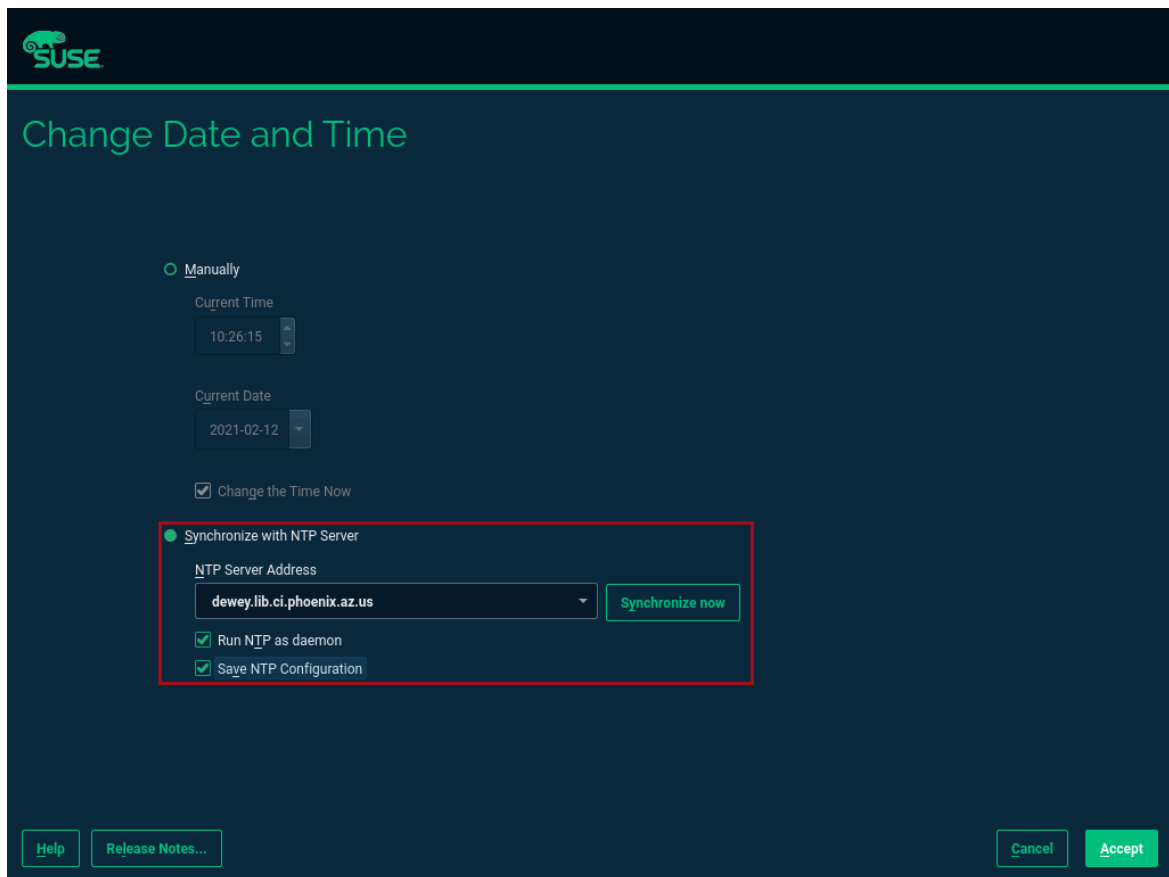


FIGURE 11: SLES SETUP NTP

- When the **Installation Settings** page is displayed, make sure that:
 - The firewall will be disabled
 - The SSH service will be enabled
 - Kdump status is disabled

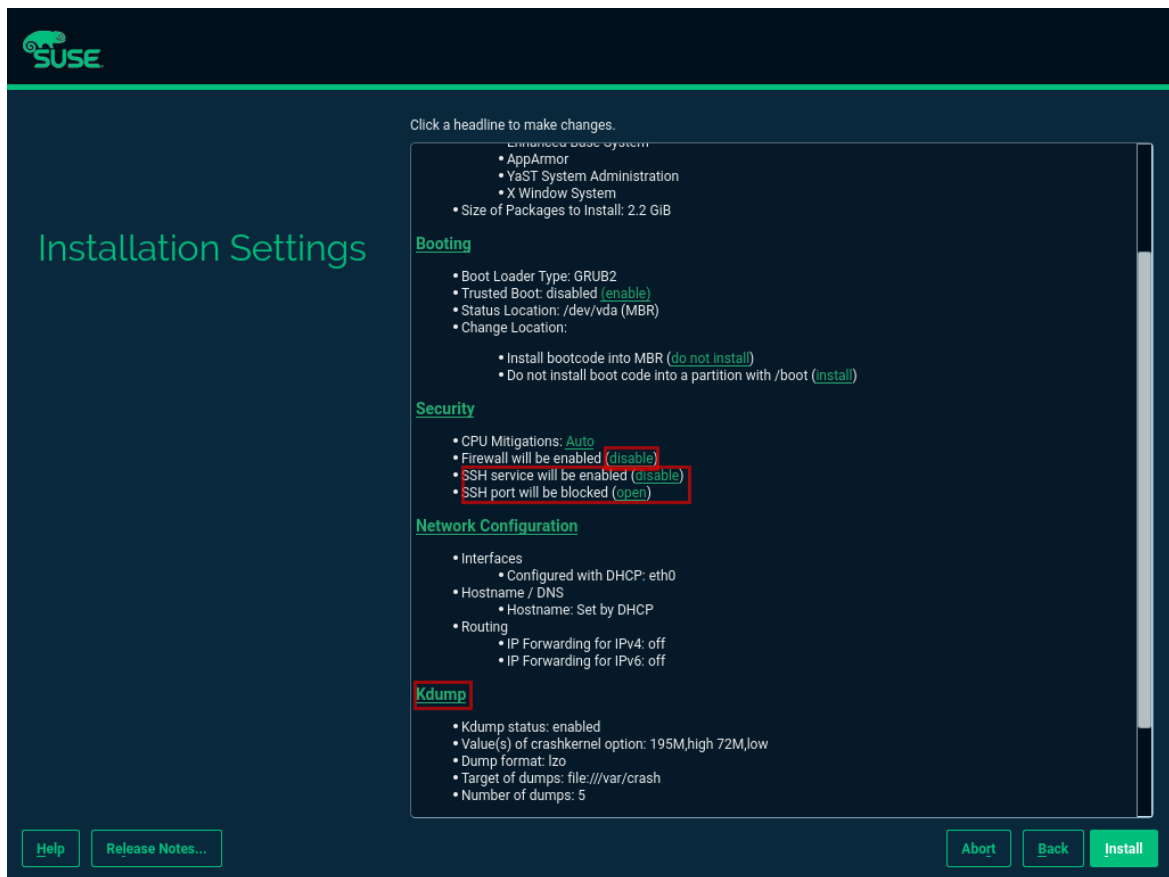


FIGURE 12: SLES SETUP SUMMARY

- To disable Kdump, click its label. This opens the **Kdump Start-Up** page. On that page, make sure "Disable Kdump" is selected.

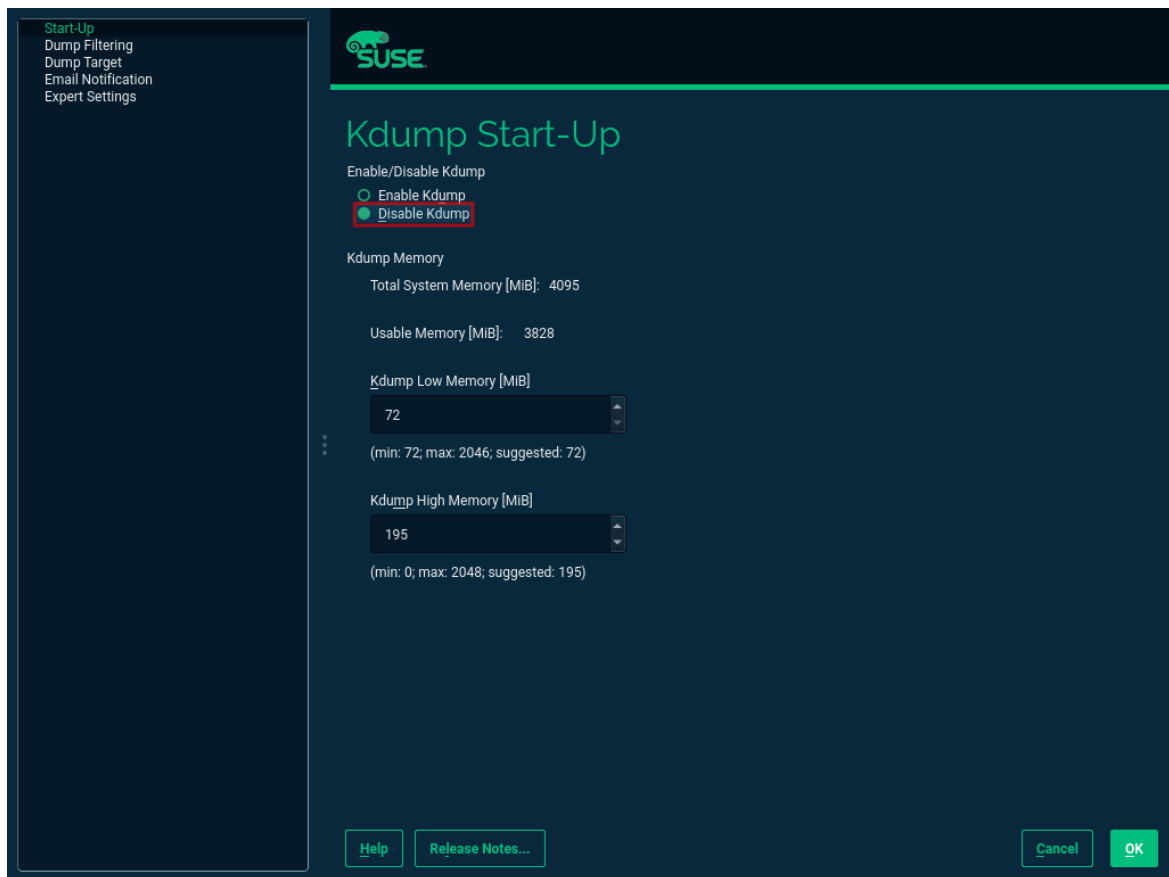


FIGURE 13: SLES SETUP KDUMP

Finish the installation and proceed to the next chapter.

4.2 Configuring the Kubernetes nodes

For the purpose of this guide, the workstation will be used to orchestrate all other machines via Salt.

4.2.1 Installing and configuring Salt minions

- First, register all systems to the [SUSE Customer Center \(https://scc.suse.com\)](https://scc.suse.com) or an SMT or RMT server to obtain updates during installation and afterward.

- When using an SMT or RMT server, the address must be specified:

```
$ sudo SUSEConnect --url "https://<SMT/RMT-address>"
```

- When registering via SUSE Customer Center, use your subscription and e-mail address:

```
$ sudo SUSEConnect -r <SubscriptionCode> -e <EmailAddress>
```

- The base system is required by all other modules. To start the installation, run:

```
$ sudo SUSEConnect -p sle-module-basesystem/15.2/x86_64
```

- Before you can use the workstation for orchestration, install and configure Salt on all Kubernetes nodes:

```
$ sudo zypper in -y salt-minion  
$ sudo echo "master: <WorkstationIP>" > /etc/salt/minion  
$ sudo systemctl enable salt-minion --now
```

4.3 Configuring the management workstation

The management workstation is used to deploy and maintain the Kubernetes cluster and workloads running on it.

4.3.1 Installing and configuring Salt masters

It is recommended to use Salt to orchestrate all Kubernetes nodes. You can skip this activity, but this means that every node must be configured manually afterward.

- To install Salt, run:

```
$ sudo zypper in -y salt-master  
$ sudo systemctl enable salt-master --now
```

- Make sure all Kubernetes nodes show up when running:

```
$ salt-key -L
```

- Accept and verify all minion keys:

```
$ salt-key -A -y
$ salt-key -L
```

- Since the RKE deployment needs SSH, an ssh key is needed. To generate a new one, run:

```
$ ssh-keygen -t rsa -b 4096
```

- Distribute the generated key to all other nodes with the command:

```
$ ssh-copy-id -i <path to your sshkey> root@<nodeIP>
```

4.3.2 Configuring Kubernetes nodes

- Check the status of the firewall and disable it if this is not yet done:

```
$ sudo salt '*' cmd.run 'systemctl status firewalld'
$ sudo salt '*' cmd.run 'systemctl disable firewalld --now'
```

- Check the status of Kdump and disable it if this is not yet done:

```
$ sudo salt '*' cmd.run 'systemctl status kdump'
$ sudo salt '*' cmd.run 'systemctl disable kdump --now'
```

- Make sure swap is disabled and disable if this is not yet done:

```
$ sudo salt '*' cmd.run 'cat /proc/swaps'
$ sudo salt '*' cmd.run 'swapoff -a'
```

- Check the NTP time synchronization and enable it if this is not yet done:

```
$ sudo salt '*' cmd.run 'systemctl status chronyd'
$ sudo salt '*' cmd.run 'systemctl enable chronyd --now'
$ sudo salt '*' cmd.run 'chronyc sources'
```

- Make sure the SSH server is running:

```
$ sudo salt '*' cmd.run 'systemctl status sshd'
$ sudo salt '*' cmd.run 'systemctl enable sshd --now'
```

- Activate the needed SUSE modules:

```
$ sudo salt '*' cmd.run 'SUSEConnect -p sle-module-containers/15.2/x86_64'
```

- Install the packages required to run SAP Data Intelligence:

```
$ sudo salt '*' cmd.run 'zypper in -y nfs-client nfs-kernel-server xfsprogs ceph-common open-iscsi'
```

- Enable `open-iscsid`:

```
$ sudo salt '*' cmd.run 'systemctl status iscsid'
$ sudo salt '*' cmd.run 'systemctl enable iscsid --now'
```

4.4 Installing Rancher Kubernetes Engine 2

To install RKE 2 on the cluster nodes, download the RKE 2 install script and copy it to each of the Kubernetes cluster nodes. The single steps are described in the following sections.

For more detailed information, read the [RKE 2 Quick Start guide \(https://docs.rke2.io/install/quick-start/\)](https://docs.rke2.io/install/quick-start/).

4.4.1 Downloading the RKE 2 install script

To download the RKE 2 install script, run the following command:

```
$ curl -sfL https://get.rke2.io --output install.sh
$ chmod 0700 install.sh
```

4.4.2 Deploying RKE 2

Now deploy the Kubernetes cluster:

- In a first step, deploy the Kubernetes master nodes.
- The second step is to deploy the worker nodes of the Kubernetes cluster.
- Finally, configure and test the access to the RKE 2 cluster from the management workstation.

4.4.2.1 Deploying RKE 2 master nodes

This section describes the deployment of the RKE 2 master nodes.

- Copy the `install.sh` script you downloaded before to all of your Kubernetes nodes (masters and workers).

```
$ export INSTALL_RKE2_TYPE="server"
$ export INSTALL_RKE2_VERSION=v1.19.8+rke2r1
$ ./install.sh
```

The command above downloads a TAR archive and extracts it to the local machine.

- Create a first configuration file for the RKE 2 deployment:

```
$ sudo mkdir -p /etc/rancher/rke2
$ sudo cat <<EOF > /etc/rancher/rke2/config.yaml
disable: rke2-ingress-nginx
EOF
```

- With the following command, start the actual deployment:

```
$ sudo systemctl enable --now rke2-server.service
```

- For the remaining master nodes, proceed with the command:

```
$ sudo mkdir -p /etc/rancher/rke2/
```

- Copy the authentication token from the first master node found at `/var/lib/rancher/server/token`. Save this token for later usage.

- Create the file `/etc/rancher/rke2/config.yaml` on the other nodes of the RKE 2 cluster.

```
$ sudo cat <<EOF > /etc/rancher/rke2/config.yaml
server: https://<ip of first master node>:9345
token: <add token gained from first master node>
disable: rke2-nginx-ingress
EOF
```

- Distribute this file to the remaining master and worker nodes.

4.4.2.2 Deploying RKE 2 worker nodes

This section describes the deployment of the RKE 2 worker nodes.

- If not already done, copy the install script to the worker nodes.
- Create the file `/etc/rancher/rke2/config.yaml` for the worker nodes.

- Set the environment variables to install the RKE 2 worker nodes and execute the install script.

```
$ export INSTALL_RKE2_VERSION=v1.19.8+rke2r1
$ export INSTALL_RKE2_TYPE="agent"
$ sudo ./install.sh
$ sudo systemctl enable --now rke2-agent.service
```

- You can monitor the installation progress via the systemd journal.

```
$ sudo journalctl -f -u rke2-agent
```

4.4.2.3 Checking the installation

- Download a matching `kubectl` version to the management workstation. Below find an example for `kubectl` version 1.19.8:

```
$ curl -LO https://storage.googleapis.com/kubernetes-release/release/v1.19.8/bin/
linux/amd64/kubectl
$ chmod a+x kubectl
$ sudo cp -av kubectl /usr/bin/kubectl
```

- Get the KUBECONFIG file from the first master node and copy it to the management workstation:

```
$ scp <first master node>:/etc/rancher/rke2/rke2.yaml <management workstation>:/
path/where/kubeconfig/should/be/placed
```

- Replace "127.0.0.1" in `rke2.yaml` with the IP address of the first master node:

```
$ sed -e -i 's/127.0.0.1/<ip of first master node>/' rke2.yaml
```

- Verify it by running the following command:

```
$ export KUBECONFIG=<PATH to your kubeconfig>
$ kubectl version
$ kubectl get nodes
```

Your RKE 2 cluster should now be ready to use.

5 Installing SAP Data Intelligence 3.1

This section describes the installation of SAP DI 3.1 on an RKE 2-powered Kubernetes cluster.

5.1 Preparations

The following steps need to be executed before the deployment of SAP DI 3.1 can start:

- Create a namespace for SAP DI 3.1.
- Create an access to a secure private registry.
- Create a default storage class.
- Download and install SAP SLCBridge.
- Download the `stack.xml` file for provisioning the DI 3.1 installation.
- Check if the `nfsd` and `nfsv4` kernel modules are loaded and/or loadable on the Kubernetes nodes.

5.1.1 Creating namespace for SAP DI 3.1 in the Kubernetes cluster

Log in your management workstation and create the namespace in the Kubernetes cluster where DI 3.1 will be deployed.

```
$ kubectl create ns <NAMESPACE for DI 31>
$ kubectl get ns
```

5.1.2 Creating cert file to access the secure private registry

Create a file named `cert` that contains the SSL certificate chain for the secure private registry. This imports the certificates into SAP DI 3.1.

```
$ cat CA.pem > cert
$ kubectl -n <NAMESPACE for DI 31> create secret generic cmcertificates --from-file=cert
```

5.2 Creating default storage class

To install SAP DI 3.1, a default storage class is needed to provision the installation with physical volumes (PV). Below find an example for a ceph/rbd based storage class that uses the CSI.

- Create the yaml files for the storage class; contact your storage admin to get the required information.
- Create config-map:

```
$ cat << EOF > csi-config-map.yaml
---
apiVersion: v1
kind: ConfigMap
data:
  config.json: |-
    [
      {
        "clusterID": "<ID of your ceph cluster>",
        "monitors": [
          "<IP of Monitor 1>:6789",
          "<IP of Monitor 2>:6789",
          "<IP of Monitor 3>:6789"
        ]
      }
    ]
metadata:
  name: ceph-csi-config
EOF
```

- Create a secret to access the storage:

```
$ cat << EOF > csi-rbd-secret.yaml
---
apiVersion: v1
kind: Secret
metadata:
  name: csi-rbd-secret
  namespace: default
stringData:
  userID: admin
  userKey: AQCR7htglvJzBxAAtPN0YUeSiDzyTeQe0lveDQ==
EOF
```

- Download the file:

```
$ curl -LO https://raw.githubusercontent.com/ceph/ceph-csi/master/deploy/rbd/kubernetes/csi-rbdplugin-provisioner.yaml
```

- Download the file:

```
$ curl -LO https://raw.githubusercontent.com/ceph/ceph-csi/master/deploy/rbd/kubernetes/csi-rbdplugin.yaml
```

- Create a pool on the Ceph storage where the PVs will be created, and insert the poolname and the Ceph cluster ID:

```
$ cat << EOF > csi-rbd-sc.yaml
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: csi-rbd-sc
provisioner: rbd.csi.ceph.com
parameters:
  clusterID: <your ceph cluster id>
  pool: <your pool>
  csi.storage.k8s.io/provisioner-secret-name: csi-rbd-secret
  csi.storage.k8s.io/provisioner-secret-namespace: default
  csi.storage.k8s.io/node-stage-secret-name: csi-rbd-secret
  csi.storage.k8s.io/node-stage-secret-namespace: default
reclaimPolicy: Delete
mountOptions:
  - discard
EOF
```

- Create config for encryption. This is needed, else the deployment of the CSI driver for ceph/rbd will fail.

```
$ cat << EOF > kms-config.yaml
---
apiVersion: v1
kind: ConfigMap
data:
  config.json: |-
    {
      },
      "vault-tokens-test": {
        "encryptionKMSType": "vaulttokens",
        "vaultAddress": "http://vault.default.svc.cluster.local:8200",
```

```

    "vaultBackendPath": "secret/",
    "vaultTLSServerName": "vault.default.svc.cluster.local",
    "vaultCAVerify": "false",
    "tenantConfigName": "ceph-csi-kms-config",
    "tenantTokenName": "ceph-csi-kms-token",
    "tenants": {
      "my-app": {
        "vaultAddress": "https://vault.example.com",
        "vaultCAVerify": "true"
      },
      "an-other-app": {
        "tenantTokenName": "storage-encryption-token"
      }
    }
  }
}
}
}
}
metadata:
  name: ceph-csi-encryption-kms-config
EOF

```

- Deploy the ceph/rbd CSI and storage class:

```

$ kubectl apply -f csi-config-map.yaml
$ kubectl apply -f csi-rbd-secret.yaml
$ kubectl apply -f \
  https://raw.githubusercontent.com/ceph/ceph-csi/master/deploy/rbd/kubernetes/csi-
provisioner-rbac.yaml
$ kubectl apply -f \
  https://raw.githubusercontent.com/ceph/ceph-csi/master/deploy/rbd/kubernetes/csi-
nodeplugin-rbac.yaml
$ kubectl apply -f csi-rbdplugin-provisioner.yaml
$ kubectl apply -f csi-rbdplugin.yaml
$ kubectl apply -f csi-rbd-sc.yaml
$ kubectl apply -f kms-config.yaml
$ kubectl patch storageclass csi-rbd-sc \
  -p '{"metadata": {"annotations":{"storageclass.kubernetes.io/is-default-
class":"true"}}}'

```

- Check your storage class:

```

$ kubectl get sc

```

NAME	PROVISIONER	RECLAIMPOLICY	VOLUMEBINDINGMODE
csi-rbd-sc (default)	rbd.csi.ceph.com	Delete	Immediate
ALLOWVOLUMEEXPANSION	AGE		false
	103m		

5.3 Using Longhorn for physical volumes

A possible valid alternative is to deploy Longhorn storage for serving the PVs of SAP DI 3. For more information, visit <https://longhorn.io>.

Longhorn uses the CSI for accessing the storage.

5.3.1 Prerequisites

Each node in the Kubernetes cluster where Longhorn is installed must fulfill the following requirements:

- A matching Kubernetes version (this is because we are installing SAP DI 3)
- open-iscsi
- Support for the XFS file system
- `nfsv4` client must be installed
- `curl`, `lsblk`, `blkid`, `findmnt`, `grep`, `awk` must be installed
- Mount propagations must be enabled on Kubernetes cluster

A check script provided by Longhorn project can be installed on the management workstation.

```
$ curl -sSfL https://raw.githubusercontent.com/longhorn/longhorn/v1.1.0/scripts/environment_check.sh | bash
```

On the Kubernetes worker nodes that should act as storage nodes, add sufficient disk drives. Create mount points for these disks, then create the XFS file system on top, and mount them. Longhorn will be configured to use these disks for storing data. For detailed information about disk sizes, see the [SAP Sizing Guide for SAP DI 3](https://help.sap.com/viewer/835f1e8d0dde4954ba0f451a9d4b5f10/3.1.latest/en-US) (<https://help.sap.com/viewer/835f1e8d0dde4954ba0f451a9d4b5f10/3.1.latest/en-US>).

Make sure as well that `iscsid` is started on the Longhorn nodes:

```
$ sudo systemctl enable --now iscsid
```

5.3.2 Installing Longhorn

The installation of Longhorn is straight forward. This guide follows the documentation of Longhorn which can be found at <https://longhorn.io/docs/1.1.0/>.

- Start the deployment with the command:

```
$ kubectl apply -f https://raw.githubusercontent.com/longhorn/longhorn/v1.1.0/
deploy/longhorn.yaml
```

- Monitor the deployment progress with the following command:

```
$ kubectl get pods \
  --namespace longhorn-system \
  --watch
```

5.3.3 Configuring Longhorn

The Longhorn storage administration is done via a built-in UI dashboard. To access this UI, an Ingress needs to be configured.

5.3.3.1 Creating an Ingress with basic authentication

- Create a basic auth file:

```
$ USER=<USERNAME_HERE>; \
  PASSWORD=<PASSWORD_HERE>; \
  echo "${USER}:${(openssl passwd -stdin -apr1 <<< ${PASSWORD})}" >> auth
```

- Create a secret from the file auth:

```
$ kubectl -n longhorn-system create secret generic basic-auth --from-file=auth
```

- Create the Ingress with basic authentication:

```
$ cat <<EOF > longhorn-ingress.yaml
apiVersion: networking.k8s.io/v1beta1
kind: Ingress
metadata:
  name: longhorn-ingress
  namespace: longhorn-system
  annotations:
    # type of authentication
    nginx.ingress.kubernetes.io/auth-type: basic
    # prevent the controller from redirecting (308) to HTTPS
    nginx.ingress.kubernetes.io/ssl-redirect: 'false'
    # name of the secret that contains the user/password definitions
```

```
    nginx.ingress.kubernetes.io/auth-secret: basic-auth
    # message to display with an appropriate context why the authentication is
    required
    nginx.ingress.kubernetes.io/auth-realm: 'Authentication Required '
spec:
  rules:
  - http:
      paths:
      - path: /
        backend:
          serviceName: longhorn-frontend
          servicePort: 80
EOF

$ kubectl -n longhorn-system apply -f longhorn-ingress.yaml
```

5.3.3.2 Adding disk space for Longhorn

This section describes how to add disk space to the Longhorn implementation.

- Prepare the disks:
 - Create a mount point for the disks.
 - Create a partition and file system on the disk.
 - Mount the file system of the disk to the created mount point.
 - Add an entry for this file system to `fstab`.
 - Test this setup (for example: `umount` file system, run `mount -a`, check with `lsblk` if file system is mounted properly)
- Configure additional disks using the Longhorn UI:
 - Access the UI of Longhorn through the URL configured in the Ingress (for example "http://node:").
 - Authenticate with the user and password set in the previous chapter.

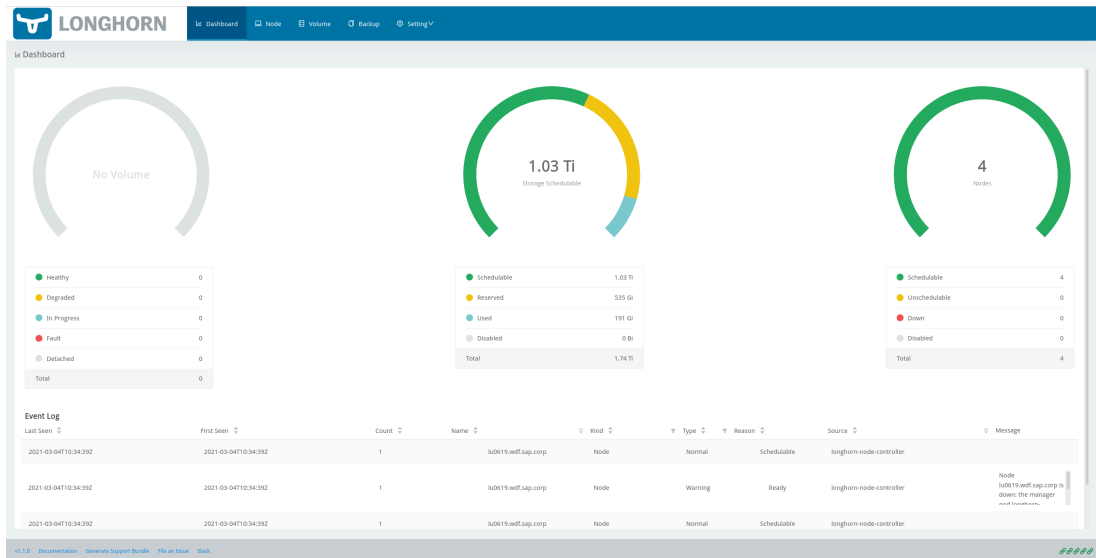


FIGURE 14: LONGHORN UI OVERVIEW

On this Overview page, click the nodes tab.

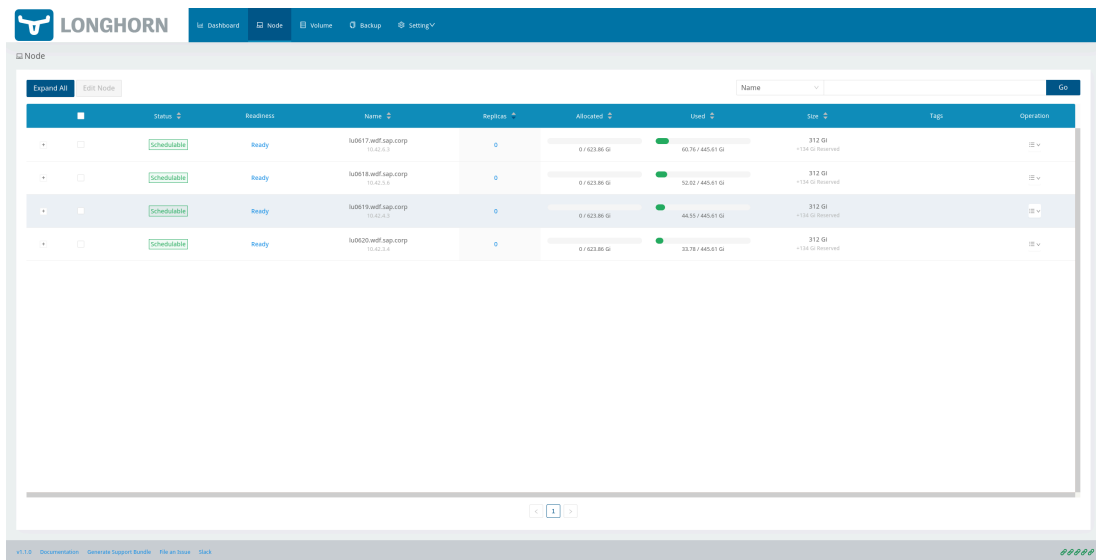


FIGURE 15: LONGHORN UI NODES

- Hover over the settings icon on the right side.

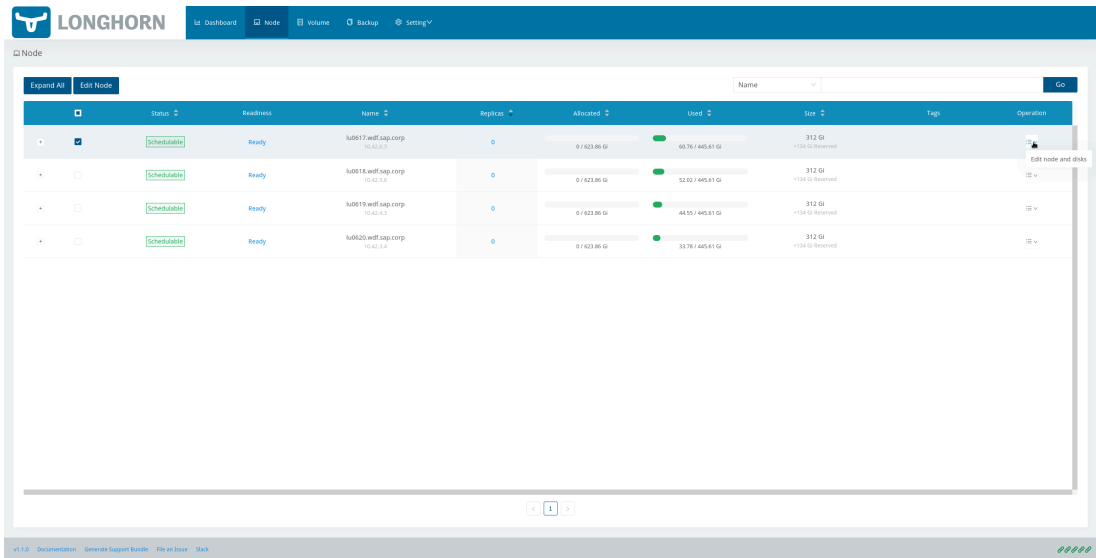


FIGURE 16: LONGHORN UI EDIT NODE

- Click "Edit Node and Disks".

Edit Node and Disks
✕

Node Scheduling

Enable Disable

Eviction Requested

True False

Node Tags

+ New Node Tag

Conditions

MountPropagation Ready Schedulable

+ New Disk Tag

Conditions

Ready Schedulable

Storage Available	Storage Scheduled	Storage Max
384.86Gi	0Gi	445.61Gi

Name

Path

Storage Reserved

Gi

Scheduling

Enable Disable

Eviction Requested

True False 🗑️

Add Disk

Cancel
Save

FIGURE 17: LONGHORN UI ADD DISK

- Click the "Add Disks" button.

Conditions

MountPropagation Ready Schedulable

+ New Disk Tag

Conditions

Ready Schedulable

Storage Available	Storage Scheduled	Storage Max
384.86Gi	0Gi	445.61Gi

Name

default-disk-d538223c04d0d9c0

Path

/var/lib/longhorn/

Storage Reserved

133.68

Gi

Scheduling

Enable Disable

Eviction Requested

True False



+ New Disk Tag

Storage Available	Storage Scheduled	Storage Max
0Gi	0Gi	0Gi

Name

nvme-1

Path

/longhorn-data/

Storage Reserved

100

Gi

Scheduling

Enable Disable

Eviction Requested

True False



Add Disk

Cancel

Save

- Fill in the mount point and mark it as **schedulable**.
- Click "Save".
- Repeat this for other disks on the other nodes.
- Check the status in the UI of Longhorn:
 - Point the browser to the URL defined in the Ingress.
 - Authenticate with the user and password created above.
- The UI displays an overview of the Longhorn storage.

For more details, see the [Longhorn documentation \(https://longhorn.io/docs/1.1.0/\)](https://longhorn.io/docs/1.1.0/) ↗

5.3.4 Creating a Storage Class on top of Longhorn

The following command creates a storage class named `longhorn` for the use of SAP DI 3.1.

```
$ kubectl create -f https://raw.githubusercontent.com/longhorn/longhorn/v1.1.0/examples/storageclass.yaml
```

Annotate this storage class as default:

```
$ kubectl patch storageclass longhorn \
  -p '{"metadata": {"annotations":{"storageclass.kubernetes.io/is-default-class":"true"}}}'
```

5.3.5 Longhorn documentation

For more details, see the [Longhorn documentation \(https://longhorn.io/docs/1.1.0/\)](https://longhorn.io/docs/1.1.0/) ↗.

5.4 Using NetApp Trident and NetApp StorageGRID as block and object store

5.4.1 Installing NetApp Trident

Install NetApp Trident according to the documentation at <https://netapp-trident.readthedocs.io/en/stable-v21.01/>:

- Create a storage class on top of NetApp Trident and make it the default storage class.
- Create an **igroup** and add the initiator names from the RKE 2 worker nodes:
 - Log in to the RKE2 worker nodes and query the initiator names from `/etc/iscsi/initiator.iscsi`.
 - Make sure the initiator names are unique! If they are not, replace them with new ones, generated with `iscsi-iname`.
 - Create an **igroup** inside your **vserver** and add all initiator names.

5.4.2 Installing NetApp StorageGRID

If an SAP Data Intelligence checkpoint store must be used, or AI/ML scenarios should be used, an object store is required. To use StorageGRID as an object store for SAP Data Intelligence, complete the following installation (if StorageGRID is not used as an appliance) and configuration steps:

- Prepare NetApp StorageGRID for Object Store:
 - Use the instructions from the publicly available documentation [Deploying StorageGRID in a Kubernetes Cluster \(https://netapp.io/2019/01/15/deploying-storagegrid-in-a-kubernetes-cluster/\)](https://netapp.io/2019/01/15/deploying-storagegrid-in-a-kubernetes-cluster/) to install StorageGRID in the cluster.
 - Select **Manage Certificates**, and make sure that the certificates being used are issued from an official Trust Center: NetApp- StorageGRID/SSL-Certificate-Configuration - as outlined in the document [SSL Certificate Configuration for StorageGRID \(https://github.com/NetApp-StorageGRID/SSL-Certificate-Configuration\)](https://github.com/NetApp-StorageGRID/SSL-Certificate-Configuration).

- Following the instructions in section [Configuring tenant accounts and connections](https://docs.netapp.com/sgws-114/topic/com.netapp.doc.sg-swift/GUID-F0D7793E-678D-42C1-91DA-29697823B492.html) (<https://docs.netapp.com/sgws-114/topic/com.netapp.doc.sg-swift/GUID-F0D7793E-678D-42C1-91DA-29697823B492.html>), configure a tenant.
- Following the instructions in section [Creating an S3 bucket](https://docs.netapp.com/sgws-114/index.jsp?topic=%2Fcom.netapp.doc.sg-tenant-admin%2FGUID-AC8EC92E-1E79-41FA-8A8B-960484C5699D.html&resultof=%22%63%72%65%61%74%65%22%20%22%63%72%65%61%74%22%20%22%62%75%63%6b%65%74%22%20) (<https://docs.netapp.com/sgws-114/index.jsp?topic=%2Fcom.netapp.doc.sg-tenant-admin%2FGUID-AC8EC92E-1E79-41FA-8A8B-960484C5699D.html&resultof=%22%63%72%65%61%74%65%22%20%22%63%72%65%61%74%22%20%22%62%75%63%6b%65%74%22%20>), log in to the created tenant and create a bucket (such as DataHub).
- Use the instructions in section [Creating another user's S3 access keys](https://docs.netapp.com/sgws-114/topic/com.netapp.doc.sg-tenant-admin/GUID-042C8C01-E479-4696-AFA2-A9212C10E723.html) (<https://docs.netapp.com/sgws-114/topic/com.netapp.doc.sg-tenant-admin/GUID-042C8C01-E479-4696-AFA2-A9212C10E723.html>) to define an access key; then save the generated access key and the corresponding secret.

5.5 Downloading the SLC Bridge

The SLC Bridge can be obtained:

- from the SAP software center https://support.sap.com/en/tools/software-logistics-tools.html#section_622087154. Choose **Download SLC Bridge**.
- via the information in the release notes of the SLC Bridge at <https://launchpad.support.sap.com/#/notes/2589449>.
- via <https://help.sap.com/viewer/a8d90a56d61a49718ebcb5f65014bbe7/3.1.latest/en-US/8ae38791d71046fab1f25ee0f682dc4c.html>.

Download the SLC Bridge software to the management workstation.

5.6 Installing the SLC Bridge

Rename the SLC Bridge binary to `slcb` and make it executable. Deploy the SLC Bridge to the Kubernetes cluster.

```
$ mv SLCB01_XX-70003322.EXE slcb
$ chmod 0700 slcb
```

```
$ export KUBECONFIG=<<KUBE_CONFIG>
$ ./slcb init
```

During the interactive installation, the following information is needed:

- URL of secure private registry
- Choose **expert mode**
- Choose **NodePort** for the service

Take a note of the service port of the SLC Bridge. It is needed for the installation of SAP DI 3.1 or for the reconfiguration of DI 3.1, for example to enable backup. If you forgot to note it down, the following command will list the service port:

```
$ kubectl -n sap-slcbridge get svc
```

5.7 Creating and downloading Stack XML for the SAP DI installation

Follow the steps described in the chapter [Install SAP Data Intelligence with SLC Bridge in a Cluster with Internet Access \(https://help.sap.com/viewer/a8d90a56d61a49718ebcb5f65014bbe7/3.1.latest/en-US/7e4847e241c340b3a3c50a5db11b46e2.html\)](https://help.sap.com/viewer/a8d90a56d61a49718ebcb5f65014bbe7/3.1.latest/en-US/7e4847e241c340b3a3c50a5db11b46e2.html) of the SAP DI 3.1 Installation Guide.

5.7.1 Creating Stack XML

You can create the Stack XML via the SAP Maintenance Planner. Access the tool via <https://support.sap.com/en/alm/solution-manager/processes-72/maintenance-planner.html>. Go to the Maintenance Planner at <https://apps.support.sap.com/sap/support/mp> published on the SAP Web site and generate a Stack XML file with the container image definitions of the SAP Data Intelligence release that you want to install. Download the Stack XML file to a local directory. Copy stack.xml to the management workstation.

5.8 Running the installation of SAP DI

The installation of SAP DI 3.1 is invoked by:

```
$ export KUBECONFIG=<path to kubeconfig>
$ ./slcb execute --useStackXML MP_Stack_XXXXXXXXXX_XXXXXXXXX.xml --url https://
<node>:<service port>/docs/index.html
```


This starts an interactive process for configuring and deploying SAP DI 3.1.

The table below lists some parameters available for an SAP DI 3.1 installation:

Parameter	Condition	Recommendation
Kubernetes Namespace	Always	set to namespace created beforehand
Installation Type	installation or update	either
Container Registry	Always	add the uri for the secure private registry
Checkpoint Store Configuration	installation	whether to enable Checkpoint Store
Checkpoint Store Type	if Checkpoint Store is enabled	use S3 object store from SES
Checkpoint Store Validation	if Checkpoint is enabled	Object store access will be verified
Container Registry Settings for Pipeline Modeler	optional	used if a second container registry is used
StorageClass Configuration	optional, needed if a different StorageClass is used for some components	leave the default
Default StorageClass	detected by SAP DI installer	The Kubernetes cluster shall have a storage class annotated as default SC
Enable Kaniko Usage	optional if running on Docker	enable
Container Image Repository Settings for SAP Data Intelligence Modeler	mandatory	

Parameter	Condition	Recommendation
Container Registry for Pipeline Modeler	optional	Needed if a different container registry is used for the pipeline modeler images
Loading NFS Modules	optional	Make sure that nfsd and nfsv4 kernel modules are loaded on all worker nodes
Additional Installer Parameters	optional	

For more details about input parameters for an SAP DI 3.1 installation, visit the section [Required Input Parameters \(https://help.sap.com/viewer/a8d90a56d61a49718e-bcb5f65014bbe7/3.1.latest/en-US/abfa9c73f7704de2907ea7ff65e7a20a.html\)](https://help.sap.com/viewer/a8d90a56d61a49718e-bcb5f65014bbe7/3.1.latest/en-US/abfa9c73f7704de2907ea7ff65e7a20a.html) of the SAP Data Intelligence Installation Guide.

5.9 Post-installation tasks

After the installation workflow is successfully finished, you need to carry out some more tasks:

- Obtain or create an SSL certificate to securely access the SAP DI installation:

- Create a certificate request using `openssl`, for example:

```
$ openssl req -newkey rsa:2048 -keyout <hostname>.key -out <hostname>.csr
```

- Decrypt the key:

```
$ openssl rsa -in <hostname>.key -out decrypted-<hostname>.key
```

- Let a CA sign the `<hostname>.csr` You will receive a `<hostname>.cert`.
- Create a secret from the certificate and the key in the SAP DI 3 namespace:

```
$ export NAMESPACE=<SAP DI 3 namespace>
$ kubectl -n $NAMESPACE create secret tls vsystem-tls-certs --key decrypted-<hostname>.key --cert <hostname>.cert
```

- Deploy an `nginx-ingress` controller

- For more information, see <https://kubernetes.github.io/ingress-nginx/deploy/#bare-metal>
- Create the `nginx-ingress` controller as a **nodePort** service according to the [Ingress nginx documentation](#):

```
$ kubectl apply -f https://raw.githubusercontent.com/kubernetes/ingress-nginx/controller-v0.46.0/deploy/static/provider/baremetal/deploy.yaml
```

- Determine the port the `nginx` controller is redirecting HTTPS to:

```
$ kubectl -n ingress-nginx get svc ingress-nginx-controller
```

The output should be similar to the below:

```
kubectl -n ingress-nginx get svc ingress-nginx-controller
NAME                                TYPE           CLUSTER-IP   EXTERNAL-IP   PORT(S)
ingress-nginx-controller            NodePort       10.43.86.90  <none>        80:31963/
TCP,443:31106/TCP                   53d
```

In our example here, the TLS port is be 31106. Note the port IP down as you will need it to access the SAP DI installation from the outside.

- Create an Ingress to access the SAP DI installation:

```
$ cat <<EOF > ingress.yaml
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  annotations:
    kubernetes.io/ingress.class: nginx
    nginx.ingress.kubernetes.io/force-ssl-redirect: "true"
    nginx.ingress.kubernetes.io/secure-backends: "true"
    nginx.ingress.kubernetes.io/backend-protocol: HTTPS
    nginx.ingress.kubernetes.io/proxy-body-size: "0"
    nginx.ingress.kubernetes.io/proxy-buffer-size: 16k
    nginx.ingress.kubernetes.io/proxy-connect-timeout: "30"
    nginx.ingress.kubernetes.io/proxy-read-timeout: "1800"
    nginx.ingress.kubernetes.io/proxy-send-timeout: "1800"
  name: vsystem
spec:
  rules:
  - host: "<hostname FQDN must match SSL certificate"
    http:
```

```
paths:
  - backend:
      serviceName: vsystem
      servicePort: 8797
      path: /
tls:
  - hosts:
      - "<hostname FQDN must match SSL certificate>"
      secretName: vsystem-tls-certs
EOF
$ kubectl apply -f ingress.yaml
```

- Connecting to `https://hostname: <ingress service port>` brings up the SAP DI login dialog.

5.10 Testing the SAP Data Intelligence 3 installation

Finally, the SAP DI 3 installation should be verified with some very basic tests:

- Log in to SAP DI's launchpad
- Create example pipeline
- Create ML Scenario
- Test machine learning
- Download [vctl](#)

For details, see the [SAP DI 3 Installation Guide \(https://help.sap.com/viewer/a8d90a56d61a49718e-bcb5f65014bbe7/3.1.latest/en-US/1551785f3d7e4d37af7fe99185f7acb6.html\)](https://help.sap.com/viewer/a8d90a56d61a49718e-bcb5f65014bbe7/3.1.latest/en-US/1551785f3d7e4d37af7fe99185f7acb6.html) ↗

6 Maintenance tasks

This section provides some tips about what should and could be done to maintain the Kubernetes cluster, the operating system and the SAP DI 3 deployment.

6.1 Backup

It is good practice to keep backups of all relevant data to be able to restore the environment in case of a failure.

To perform regular backups, follow the instructions as outlined in the respective documentation below:

- For RKE 2, consult section [Backups and Disaster Recovery \(https://rancher.com/docs/rke/latest/en/etcd-snapshots/\)](https://rancher.com/docs/rke/latest/en/etcd-snapshots/) ↗
- SAP Data Intelligence 3 can be configured to create regular backups. For more information, visit [help.sap.com https://help.sap.com/viewer/a8d90a56d61a49718e-bcb5f65014bbe7/3.1.latest/en-US/e8d4c33e6cd648b0af9fd674dbf6e76c.html](https://help.sap.com/viewer/a8d90a56d61a49718e-bcb5f65014bbe7/3.1.latest/en-US/e8d4c33e6cd648b0af9fd674dbf6e76c.html) ↗.

6.2 Upgrade or update

This section explains how you can keep your installation of SAP DI, RKE 2 and SUSE Linux Enterprise Server up-to-date.

6.2.1 Updating the operating system

To obtain updates for SUSE Linux Enterprise Server 15 SP2, the installation must be registered either to SUSE Customer Center, an SMT or RMT server, or SUSE Manager with a valid subscription.

- SUSE Linux Enterprise Server 15 SP2 can be updated on the command line using [zypper](#):

```
$ sudo zypper ref -s
$ sudo zypper lu
$ sudo zypper patch
```

- Other methods for updating SUSE Linux Enterprise Server 15 SP2 are described in the [product documentation \(https://documentation.suse.com/sles\)](https://documentation.suse.com/sles) ↗.

If an update requires a reboot of the server, make sure that this can be done safely.

- For example, block access to SAP DI, and drain and cordon the Kubernetes node before rebooting:

```
$ kubectl edit ingress <put in some dummy port>
$ kubectl drain <node>
```

- Check the status of the node:

```
$ kubectl get node <node>
```

The node should be marked as **not schedulable**.

- On RKE 2 master nodes, run the command:

```
$ sudo systemctl stop rke2-server
```

- On RKE 2 worker nodes, run the command:

```
$ sudo systemctl stop rke2-agent
```

- Update SUSE Linux Enterprise Server 15 SP2:

```
$ ssh node  
$ sudo zypper patch
```

- Reboot the nodes if necessary or start the appropriate RKE 2 service.

- On master nodes, run the command:

```
$ sudo systemctl start rke2-server
```

- On worker nodes, run the command:

```
$ sudo systemctl start rke2-agent
```

- Check if the respective nodes are back and uncordon them.


```
$ kubectl get nodes  
$ kubectl uncordon <node>
```

6.2.2 Updating RKE 2

For an update of RKE 2, first read


- the [SAP DI 3.1 documentation about upgrading Kubernetes](https://help.sap.com/viewer/a8d90a56d61a49718ebcb5f65014bbe7/3.1.latest/en-US/1ca2ac1d9c5a4bd98c5aaf57e53a81bf.html) (https://help.sap.com/viewer/a8d90a56d61a49718ebcb5f65014bbe7/3.1.latest/en-US/1ca2ac1d9c5a4bd98c5aaf57e53a81bf.html) [↗](#) and
- [section Upgrade Basics](https://docs.rke2.io/upgrade/basic_upgrade/) (https://docs.rke2.io/upgrade/basic_upgrade/) [↗](#) of the RKE 2 documentation.

Then perform the following actions:

- Create a backup of everything.
- Block access to SAP DI.
- Run the update of RKE 2 according to the [RKE 2 documentation, section Upgrade Basics \(https://docs.rke2.io/upgrade/basic_upgrade/\)](https://docs.rke2.io/upgrade/basic_upgrade/) .

6.2.3 Updating SAP Data Intelligence

For an update of SAP DI 3.1., follow SAP's update guide and notes:

- [Upgrading SAP Data Intelligence \(https://help.sap.com/viewer/a8d90a56d61a49718e-bcb5f65014bbe7/3.1.latest/en-US/b87299d2e8bc436baadfa020abb59892.html\)](https://help.sap.com/viewer/a8d90a56d61a49718e-bcb5f65014bbe7/3.1.latest/en-US/b87299d2e8bc436baadfa020abb59892.html) 

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.