

SUSE Virtualization

Integrating SUSE® Virtualization with NetApp® Storage

Enabling Persistent Storage for Virtual Machines Using NetApp ON-
TAP and Trident

SUSE Virtualization
NetApp ONTAP
NetApp Trident

Gopala Krishnan A, Partner Solution Architect (SUSE)
Suresh S, Partner Solution Architect (SUSE)
Terry Smith, Ecosystem Solution Innovations Director (SUSE)

Integrating SUSE® Virtualization with NetApp® Storage

Enabling Persistent Storage for Virtual Machines Using NetApp ONTAP and Trident

Date: 2025-03-24

Summary

This guide covers how to integrate NetApp® external storage with SUSE® Virtualization to provide persistent storage for virtual machines. It provides detailed steps for provisioning LUNs, configuring iSCSI and multipath, installing and configuring Trident™ Container Storage Interface (CSI), and using a NetApp-backed StorageClass to create and attach volumes to virtual machines.

Disclaimer

Documents published as part of the series SUSE Technical Reference Documentation 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 Architecture 5
- 3 Deployment 7
- 4 Validation 15
- 5 Summary 17
- 6 Frequently Asked Questions (FAQ) 18
- 7 Legal notice 20
- 8 GNU Free Documentation License 21

1 Introduction

SUSE® Virtualization is a cloud-native virtualization platform that unifies virtual machine (VM) and container management, streamlining operations, and boosting efficiency with a single, cohesive platform for hybrid cloud infrastructure. NetApp® ONTAP® delivers unified data management, allowing you to manage, protect, and move data across your hybrid environment. By leveraging NetApp® Trident™, you can deliver high-performance storage for virtual machines (VMs) hosted in your SUSE® Virtualization environment.

1.1 Scope

This guide covers the detailed steps required for setting up a three-node SUSE Virtualization cluster to consume NetApp storage. It includes steps for LUN provisioning, multipath configuration, Trident installation, StorageClass definition, and basic validation.

1.2 Audience

This guide is intended for IT professionals responsible for virtualization and infrastructure management, such as infrastructure architects, systems administrators, or platform engineers.

To successfully follow this guide, you should have:

- Familiarity with SUSE Virtualization deployment and configuration
- Experience with NetApp storage provisioning and management
- Understanding of iSCSI concepts and configuration
- Working knowledge of Kubernetes

1.3 Prerequisites

You will need the following, minimum resources to implement the described integration:

- **SUSE Virtualization cluster** with at least three nodes for high availability

SUSE Virtualization (<https://www.suse.com/products/harvester>) is a modern, open, interoperable, hyperconverged infrastructure (HCI) solution built on Kubernetes. It is an open source alternative designed for operators seeking a cloud-native HCI solution. SUSE Virtualization runs on bare metal servers and provides integrated virtualization and distributed storage capabilities.

See [SUSE Virtualization: Hardware and Network Requirements](https://documentation.suse.com/cloudnative/virtualization/latest/en/installation-setup/requirements.html) and [Trident: Requirements](https://docs.netapp.com/us-en/trident/trident-get-started/requirements.html) for more details.

This guide applies to SUSE Virtualization v1.6.0 and newer.

- **NetApp ONTAP storage**

NetApp ONTAP (<https://www.netapp.com/ontap-data-management-software/>) intelligent data infrastructure must be available and accessible by the SUSE Virtualization cluster nodes over a network.

This guide applies to ONTAP 9 and newer.

- **NetApp Trident CSI driver**

Trident (<https://docs.netapp.com/us-en/trident>) is an open source storage provisioner and orchestrator maintained by NetApp that enables you to create storage volumes for containerized applications in Kubernetes.

This guide applies to Trident 25.10.0 and newer.

- **Administrative access:** Ensure that you have administrative access to the SUSE Virtualization UI and the underlying Kubernetes cluster to apply configurations and manage resources.

See [SUSE Virtualization: Troubleshooting FAQ](https://documentation.suse.com/cloudnative/virtualization/latest/en/troubleshooting/faq.html).

2 Architecture

The diagram below illustrates the basic architecture for SUSE Virtualization integrated with NetApp ONTAP storage.

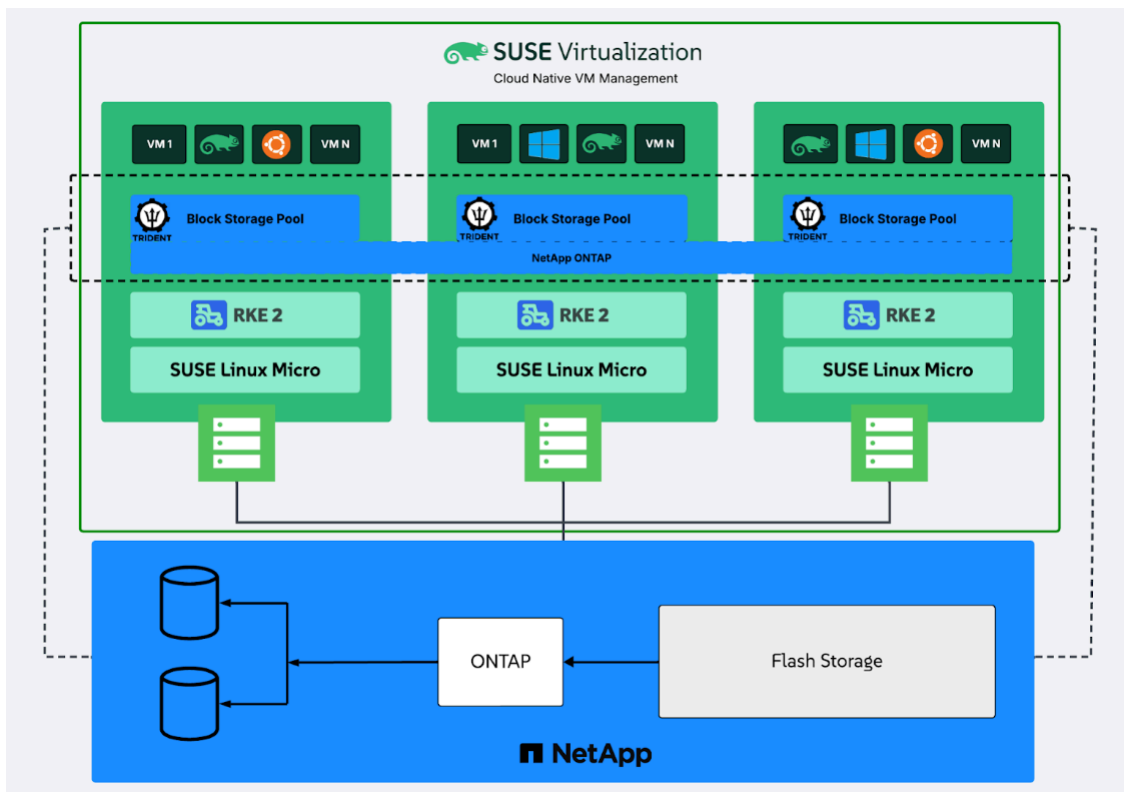


FIGURE 1: ARCHITECTURE: SUSE VIRTUALIZATION WITH NETAPP ONTAP AND TRIDENT

The key components of the solution are:

- **SUSE Virtualization** is the cloud native solution that provides unified and streamlined orchestration for VMs and containers. Built on Kubernetes, SUSE Virtualization delivers a flexible, scalable, and resilient infrastructure platform for all of your workloads.
- **NetApp ONTAP** data management software powers ONTAP storage systems. With a global namespace that supports up to 24 hardware nodes of all types, ONTAP software supports unified storage (block and file), enabling significant workload consolidation that you can manage using the ONTAP System Manager GUI. In this solution, ONTAP is used to provision iSCSI LUNs on the ONTAP storage back-end.
- **NetApp Trident** is a Container Storage Interface (CSI) that delivers dynamic storage orchestration and natively integrates with Kubernetes. Trident enables consumption and management of storage resources across all popular NetApp storage platforms, in the public cloud or on-premises. Trident runs as a single controller pod and a node pod on each worker node in the cluster.

When deployed together in the illustrated configuration, SUSE Virtualization nodes connect over the iSCSI protocol to NetApp ONTAP storage using the Trident CSI. High availability storage access is available when the cluster environment is configured with multipath support.

3 Deployment

This section provides practical guidance for deploying a functional, proof-of-concept implementation of the solution with additional considerations for scaling to production.

For deployment purposes, the solution can be illustrated as a layered stack. The deployment process consists of installing and configuring each layer of this stack, typically from bottom (compute platform) to top (software applications). Guidance for deploying each layer of the solution are provided in the following sections.

3.1 Preparing the compute platform

1. Ensure each node meets minimum hardware requirements (<https://documentation.suse.com/cloudnative/virtualization/latest/en/installation-setup/requirements.html>) .






Note

Be sure hardware virtualization (VT-x/AMD-V) is enabled on each node.

2. Deploy the nodes, storage, and network as illustrated in the architecture diagram.

3.2 Deploying SUSE Virtualization

1. Install a SUSE Virtualization high-availability cluster onto the compute platform. See [official documentation \(https://documentation.suse.com/cloudnative/virtualization/latest/en/introduction/overview.html\)](https://documentation.suse.com/cloudnative/virtualization/latest/en/introduction/overview.html)  for detailed guidance.
2. After deployment, confirm you have access to the cluster through the SUSE Virtualization UI and CLI. For tips, see [SUSE Virtualization: Authentication \(https://documentation.suse.com/cloudnative/virtualization/latest/en/installation-setup/authentication.html\)](https://documentation.suse.com/cloudnative/virtualization/latest/en/installation-setup/authentication.html)  and [SUSE Virtualization: Troubleshooting - FAQ \(https://documentation.suse.com/cloudnative/virtualization/latest/en/troubleshooting/faq.html\)](https://documentation.suse.com/cloudnative/virtualization/latest/en/troubleshooting/faq.html) .

3.3 Configuring iSCSI

Use CloudInit to create a persistent configuration of the iSCSI initiator.

1. Define the configuration of the iSCSI initiator.
 - a. Create the `01_iscsi_ip.yaml` CloudInit resource file listed below in the `/oem` directory on each cluster node.

```
apiVersion: node.harvesterhci.io/v1beta1
kind: CloudInit
metadata:
  name: iscsi-ip-config
spec:
  matchSelector: {}
  kubernetes.io/hostname: harvester-node-1 # Replace with actual hostname
  filename: 01_iscsi_ip.yaml
  contents: |
    stages:
      initramfs:
        - name: "Configure iSCSI network interfaces"
          files:
            - path: /etc/NetworkManager/system-connections/iscsi-
path1.nmconnection
              owner: 0
              group: 0
              permissions: 0600
              content: |
                [connection]
                id=iscsi-path1
                type=ethernet
                interface-name=eno3np0
                autoconnect=true
                [ipv4]
                method=manual
                addresses=192.168.100.87/24 # Replace the IP address
                never-default=true
            - path: /etc/NetworkManager/system-connections/iscsi-
path2.nmconnection
              owner: 0
              group: 0
              permissions: 0600
              content: |
                [connection]
                id=iscsi-path2
                type=ethernet
                interface-name=eno4np1
```

```
autoconnect=true
[ipv4]
method=manual
addresses=192.168.200.87/24 # Replace the IP address
never-default=true
```

- b. Apply the configuration on each node.

```
kubectl apply -f 01_iscsi_ip.yaml
```

- c. Verify the iSCSI initiator is deployed.

```
kubectl get cloudinit
NAME                AGE
iscsi-ip-config    20s
```

2. Enable the iSCSI service daemon.

- a. Create 99_iscsid.yaml in the /oem directory any node.

```
apiVersion: node.harvesterhci.io/v1beta1
kind: CloudInit
metadata:
  name: iscsid-start
spec:
  matchSelector: {}
  filename: 99_iscsid.yaml
  contents: |
    stages:
      network:
        - name: "Enable and start iscsid"
          systemctl:
            enable:
              - iscsid
            start:
              - iscsid
```

- b. Apply the configuration.

```
kubectl apply -f 99_iscsid.yaml
```

- c. Verify the status of iscsid on each node.

```
systemctl status iscsid
```

```
● iscsid.service - Open-iSCSI
```

```

Loaded: loaded (/usr/lib/systemd/system/iscsid.service; disabled;)
Active: active (running) since Fri 2026-03-27 20:59:47 IST; 8s ago
TriggeredBy: ● iscsid.socket
  Docs: man:iscsid(8)
       man:iscsiuio(8)
       man:iscsiadm(8)
Main PID: 27651 (iscsid)
Status: "Ready to process requests"
Tasks: 1
  CPU: 3ms
CGroup: /system.slice/iscsid.service
        └─27651 /sbin/iscsid -f

Mar 27 20:59:47 docserver systemd[1]: Starting Open-iSCSI...
Mar 27 20:59:47 docserver systemd[1]: Started Open-iSCSI.

```

3.4 Enabling Multipath

1. Create a CloudInit resource to auto-start the multipath daemon (multipathd) as shown below.

```

apiVersion: node.harvesterhci.io/v1beta1
kind: CloudInit
metadata:
  name: multipathd-config
spec:
  matchSelector: {}
  filename: 99_multipathd.yaml
  contents: |
    stages:
      network:
        - name: "Configure multipath.conf"
          files:
            - path: /etc/multipath.conf
              permissions: 0644
              owner: 0
              group: 0
              content: |
                defaults {
                  find_multipaths no
                  user_friendly_names yes
                }
                blacklist {
                  device {
                    vendor ".*"

```

```

        product ".*"
    }
}
blacklist_exceptions {
    device {
        vendor "NETAPP"
        product "LUN"
    }
}
- name: "Enable and start multipathd"
systemctl:
    enable:
        - multipathd
    start:
        - multipathd

```

2. Apply the configuration on any node.

```
kubectl apply -f 99_multipathd.yaml
```

3. Verify the status of `multipathd`.

```
systemctl status multipathd
```

```

● multipathd.service - Device-Mapper Multipath Device Controller
   Loaded: loaded (/usr/lib/systemd/system/multipathd.service; disabled;)
   Active: active (running) since Fri 2026-03-27 21:01:08 IST; 11s ago
 TriggeredBy: ○ multipathd.socket
   Main PID: 28854 (multipathd)
    Status: "up"
     Tasks: 7
        CPU: 18ms
    CGroup: /system.slice/multipathd.service
           └─28854 /sbin/multipathd -d -s

Mar 27 21:01:08 docserver systemd[1]: Starting Device-Mapper Multipath Device
Controller...
Mar 27 21:01:08 docserver multipathd[28854]: multipathd v0.10.2+122+suse.51e02cc:
start up

```

3.5 Deploying the Trident CSI

1. Prepare the worker nodes.

Before installing NetApp Trident, ensure that each SUSE Virtualization node meets the prerequisites for Kubernetes storage integration. Follow the official NetApp documentation for [Worker Node Preparation for Trident CSI \(https://docs.netapp.com/us-en/trident/trident-use/worker-node-prep.html#iscsi-self-healing-capabilities\)](https://docs.netapp.com/us-en/trident/trident-use/worker-node-prep.html#iscsi-self-healing-capabilities).

2. Create a Kubernetes secret for your NetApp ONTAP storage controller credentials.

```
kubectl create secret generic ontap-secret -n trident \
  --from-literal=username=admin \
  --from-literal=password='YourPassword'
```

3. Install NetApp Trident.

The preferred installation method uses `tridentctl`, which is demonstrated in the following steps. Be sure to check the [NetApp Trident Installation Guide \(https://docs.netapp.com/us-en/trident/trident-get-started/kubernetes-deploy-tridentctl.html#install-trident-using-tridentctl\)](https://docs.netapp.com/us-en/trident/trident-get-started/kubernetes-deploy-tridentctl.html#install-trident-using-tridentctl) for the latest details about obtaining and using `tridentctl`.

a. Download the `trident-install` package.

```
wget https://github.com/NetApp/trident/releases/download/v25.10.0/trident-installer-25.10.0.tar.gz
```

b. Extract the `trident-install` and associated files.

```
tar -xzf trident-installer-25.10.0.tar.gz
```

c. Change to the extracted directory.

```
cd trident-installer
```

d. Initiate the installation.

```
./tridentctl install -n trident
```

e. Verify that all Trident pods are in the `Running` state.

```
kubectl get pods -n trident
```

NAME	READY	STATUS	RESTARTS	AGE
trident-controller-5dd5f9df49-xkp7t	6/6	Running	0	2m37s
trident-node-linux-7qmcw	2/2	Running	0	2m37s
trident-node-linux-9k4xr	2/2	Running	0	2m37s

4. Configure the Trident back-end.

- a. Create a back-end configuration file, `ontap-backend.json` as below.

```
{
  "version": 1,
  "storageDriverName": "ontap-san",
  "managementLIF": "10.10.10.50",
  "dataLIF": "192.168.100.10",
  "svm": "harvester_svm",
  "username": "admin",
  "password": "YourPassword",
  "igroupName": "harvester_igroup"
}
```

- b. Apply the back-end.

```
./tridentctl create backend -f ontap-backend.json -n trident
```

- c. Validate the back-end.

```
./tridentctl get backend -n trident
```

NAME	STORAGE DRIVER	UUID	STATE	USER-STATE	VOLUMES
ontapsan_XXXXX	ontap-san	xxxxxxxx	online	normal	0

3.6 Defining the NetApp StorageClass and SnapshotClass

1. Create the NetApp StorageClass.

- a. Create a file, named `netapp-storageclass.yaml`, as shown below.

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: netapp-storageclass
  annotations:
    storageclass.kubernetes.io/is-default-class: "true"
provisioner: csi.trident.netapp.io
```

```
parameters:
  fsType: ext4
  backendType: "ontap-san"
reclaimPolicy: Delete
volumeBindingMode: Immediate
allowVolumeExpansion: true
```

b. Apply the StorageClass.

```
kubectl apply -f netapp-storageclass.yaml
```

c. Verify that the StorageClass is available.

```
kubectl get storageclass
```

NAME	PROVISIONER	VOLUMEBINDINGMODE	ALLOWVOLUMEEXPANSION
netapp-storage	csi.trident.netapp.io	Immediate	true

2. Create the NetApp SnapshotClass.

a. Create a file, named `netapp-snapshotclass.yaml`, with the content below.

```
apiVersion: snapshot.storage.k8s.io/v1
kind: VolumeSnapshotClass
metadata:
  name: netapp-snapclass
driver: csi.trident.netapp.io
deletionPolicy: Delete
```

b. Apply the SnapshotClass.

```
kubectl apply -f netapp-snapshotclass.yaml
```

c. Verify the that the SnapshotClass is deployed.

```
kubectl get volumesnapshotclass
```

NAME	DRIVER	DELETIONPOLICY	AGE
netapp-snapclass	csi.trident.netapp.io	Delete	6s

3.7 Registering the NetApp CSI driver in SUSE Virtualization

1. Log in to the SUSE Virtualization UI.
2. Go to *Advanced > Settings > csi-driver-config*.
3. Click *Edit Setting*.
4. Set the provisioner to `csi.trident.netapp.io`.
5. Click *Save*.
6. Configure the NetApp CSI driver as the default StorageClass.
 - a. Unset Longhorn as the default StorageClass.

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

- b. Set the NetApp CSI driver as the default StorageClass.

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

4 Validation

Verify that you can create a virtual machine configured with the NetApp StorageClass and which is backed by NetApp ONTAP iSCSI storage.

4.1 Creating a VM

1. Log in to the SUSE Virtualization UI.
2. Navigate to *Virtual Machines* and click *Create*.
3. Select a namespace for your VM.



Note

Only the `harvester-public` namespace is visible to all users by default.

4. Provide a name for your VM.
5. (Optional) Select a VM template (for example, iso-image, raw-image, windows-iso-image) to speed up provisioning.
6. On the *Basics* tab, configure the following settings:
 - *CPU*: 2
 - *Memory*: 4 GB
 - *SSHKey*: Select an existing key or upload a new one (required if you need SSH access).
 - (Optional) *Overcommit configuration*: Set CPU, memory, and storage overcommit ratios to allow scheduling of additional VMs even when physical resources are fully used.
 - (Optional) *Other advanced settings*: Run strategy, OS type, or custom cloud-init data.
7. On the *Volumes* tab, select or add disks.



Note

One writable root disk is created by default.

- a. Click *Add Disk* and select cd-rom as the disk type.
 - b. Select the ISO image you uploaded and set its boot order to 1.
 - c. Choose the StorageClass named netapp-storageclass, which is provisioned by Trident.
 - d. For the root disk, select the netapp-storageclass StorageClass and specify the desired size.
8. The management network is added by default.
You can remove it if using VLAN networks. Optionally, you can add additional interfaces (VLAN, bridge, masquerade).
 9. Review all settings and click *Create* to deploy the VM.

4.2 Verifying the VM Storage

You can use the SUSE Virtualization UI or CLI to verify that the VM is backed by NetApp CSI and the ONTAP iSCSI back-end.

4.2.1 Verifying with the SUSE Virtualization UI

1. Open *Virtual Machines* > <Your VM> > *Volumes*.
2. Confirm that the attached disk is using the `netapp-storageclass` StorageClass.

4.2.2 Verifying with the SUSE Virtualization CLI

1. Log in to the SUSE Virtualization management node CLI as `root`.
2. List the PersistentVolumeClaims (PVCs) in the VM's namespace.

```
kubectl get pvc -n <namespace>
```

3. Confirm that the **STORAGECLASS** column shows `netapp-storageclass`.

5 Summary

Organizations are accelerating their digital transformation by modernizing applications and infrastructure, transitioning from traditional, virtual machine environments to modern, cloud-native platforms powered by Kubernetes. SUSE Virtualization delivers a unified and flexible platform that allows organizations to seamlessly run both virtual machine (VM) and container workloads side-by-side on the same infrastructure.

NetApp Trident extends this capability by providing enterprise-grade, dynamic storage provisioning for Kubernetes and SUSE Virtualization clusters. Through its integration with NetApp ONTAP, Trident enables resilient, scalable, and efficient iSCSI-based storage for mission-critical workloads, ensuring consistent performance, automation, and data protection across hybrid environments.

This guide details the steps to configure and validate the integration of SUSE Virtualization with NetApp Trident and NetApp ONTAP, showcasing how enterprises can build a consistent, automated, and production-ready virtualization environment.

Dive into the product documentation to continue your learning journey:

- [SUSE Virtualization documentation \(https://documentation.suse.com/cloudnative/virtualization/\)](https://documentation.suse.com/cloudnative/virtualization/) ↗
- [NetApp Trident documentation \(https://docs.netapp.com/us-en/trident/\)](https://docs.netapp.com/us-en/trident/) ↗
- [NetApp ONTAP documentation \(https://docs.netapp.com/us-en/ontap/\)](https://docs.netapp.com/us-en/ontap/) ↗

For further guidance on how to modernize and optimize your IT infrastructure, contact your SUSE or NetApp representative or e-mail SUSE at isv-cosell@suse.com (<mailto:isv-cosell@suse.com>) ↗.

6 Frequently Asked Questions (FAQ)

1. What is the role of Trident CSI in the SUSE Virtualization environment?

NetApp Trident acts as a Container Storage Interface (CSI) driver that integrates NetApp storage with Kubernetes-based platforms such as SUSE Virtualization.

It enables dynamic provisioning of persistent storage volumes for virtual machines and containers using NetApp ONTAP storage systems.

2. Why is multipath configuration required for NetApp iSCSI storage?

Multipath provides high availability and redundancy for storage connections.

When multiple network paths exist between SUSE Virtualization nodes and the NetApp storage system, multipath ensures that:

- Storage traffic can fail over automatically if a path becomes unavailable.
- Performance is improved through path load balancing.
- Storage access remains uninterrupted during network or interface failures.

3. Why must the `netapp-storageclass` be configured as the default StorageClass?

Setting `netapp-storageclass` as the default ensures that newly created volumes automatically use NetApp-backed storage unless another StorageClass is explicitly selected.

This simplifies VM provisioning and guarantees that persistent volumes are dynamically created through Trident using the NetApp ONTAP back-end.

4. Can you deploy the Trident CSI using methods other than `tridentctl`?

Yes. While using `tridentctl` is the recommended installation method, Trident can also be deployed using Helm charts.

Helm-based deployment is useful when you manage Kubernetes applications through a package management workflow.

5. Can SUSE Virtualization run both virtual machines and container workloads?

Yes. SUSE Virtualization supports running virtual machines and Kubernetes workloads on the same infrastructure.

This unified platform simplifies operations and allows organizations to modernize applications while maintaining existing VM-based workloads.

7 Legal notice

Copyright © 2006–2026 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 series SUSE Technical Reference Documentation 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.

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.