



Get started

Astra Trident

NetApp
September 04, 2024

Table of Contents

- Get started 1
 - Try it out 1
 - Requirements 1
 - Deployment overview 6
 - Deploy with Trident operator 9
 - Deploy with `tridentctl` 19
 - What's next? 22

Get started

Try it out

NetApp provides a ready-to-use lab image that you can request through [NetApp Test Drive](#).

Learn about the Test Drive

The Test Drive provides you with a sandbox environment that comes with a three-node Kubernetes cluster and Astra Trident installed and configured. It is a great way to familiarize yourself with Astra Trident and explore its features.

Another option is to see the [kubeadm Install Guide](#) provided by Kubernetes.



You should not use the Kubernetes cluster that you build using these instructions in production. Use the production deployment guides provided by your distribution for creating clusters that are production ready.

If this is the first time you're using Kubernetes, familiarize yourself with the concepts and tools [here](#).

Requirements

Get started by reviewing the supported frontends, backends, and host configuration.



To learn about the ports that Astra Trident uses, see [here](#).

Critical information about Astra Trident 22.10

You must read the following critical information before upgrading to Astra Trident 22.10.

Critical information about Astra Trident 22.10

- Kubernetes 1.25 is now supported in Trident. You must upgrade to Astra Trident 22.10 prior to upgrading to Kubernetes 1.25.
- Astra Trident now strictly enforces the use of multipathing configuration in SAN environments, with a recommended value of `find_multipaths: no` in `multipath.conf` file.



Use of non-multipathing configuration or use of `find_multipaths: yes` or `find_multipaths: smart` value in `multipath.conf` file will result in mount failures. Trident has recommended the use of `find_multipaths: no` since the 21.07 release.

Supported frontends (orchestrators)

Astra Trident supports multiple container engines and orchestrators, including the following:

- Anthos On-Prem (VMware) and Anthos on bare metal 1.9, 1.10, 1.11
- Kubernetes 1.20 - 1.25

- Mirantis Kubernetes Engine 3.5
- OpenShift 4.8, 4.9, 4.10, 4.11

The Trident operator is supported with these releases:

- Anthos On-Prem (VMware) and Anthos on bare metal 1.9, 1.10, 1.11
- Kubernetes 1.20 - 1.25
- OpenShift 4.8, 4.9, 4.10, 4.11

Astra Trident also works with a host of other fully-managed and self-managed Kubernetes offerings, including Google Kubernetes Engine (GKE), Amazon Elastic Kubernetes Services (EKS), Azure Kubernetes Service (AKS), Rancher, and VMWare Tanzu Portfolio.

Supported backends (storage)

To use Astra Trident, you need one or more of the following supported backends:

- Amazon FSx for NetApp ONTAP
- Azure NetApp Files
- Cloud Volumes ONTAP
- Cloud Volumes Service for GCP
- FAS/AFF/Select 9.3 or later
- NetApp All SAN Array (ASA)
- NetApp HCI/Element software 11 or above

Feature requirements

The table below summarizes the features available with this release of Astra Trident and the versions of Kubernetes it supports.

Feature	Kubernetes version	Feature gates required?
CSI Trident	1.20 - 1.25	No
Volume Snapshots	1.20 - 1.25	No
PVC from Volume Snapshots	1.20 - 1.25	No
iSCSI PV resize	1.20 - 1.25	No
ONTAP Bidirectional CHAP	1.20 - 1.25	No
Dynamic Export Policies	1.20 - 1.25	No
Trident Operator	1.20 - 1.25	No

Feature	Kubernetes version	Feature gates required?
Auto Worker Node Prep (beta)	1.20 - 1.25	No
CSI Topology	1.20 - 1.25	No

Tested host operating systems

Though Astra Trident does not officially "support" specific operating systems, the following are known to work:

- RedHat CoreOS (RHCOS) versions as supported by OpenShift Container Platform
- RHEL or CentOS 7
- Ubuntu 18.04 or later (latest 22.04)
- Windows Server 2019

By default, Astra Trident runs in a container and will, therefore, run on any Linux worker. However, those workers need to be able to mount the volumes that Astra Trident provides using the standard NFS client or iSCSI initiator, depending on the backends you are using.

The `tridentctl` utility also runs on any of these distributions of Linux.

Host configuration

Depending on the backend(s) in use, NFS and/or iSCSI utilities should be installed on all of the workers in the cluster. See [here](#) for more information.

Storage system configuration

Astra Trident might require some changes to a storage system before a backend configuration can use it. See [here](#) for details.

Container images and corresponding Kubernetes versions

For air-gapped installations, the following list is a reference of container images needed to install Astra Trident. Use the `tridentctl images` command to verify the list of needed container images.

Kubernetes version	Container image
v1.20.0	<ul style="list-style-type: none"> • docker.io/netapp/trident:22.10.0 • docker.io/netapp/trident-autosupport:22.10 • registry.k8s.io/sig-storage/csi-provisioner:v3.3.0 • registry.k8s.io/sig-storage/csi-attacher:v4.0.0 • registry.k8s.io/sig-storage/csi-resizer:v1.6.0 • registry.k8s.io/sig-storage/csi-snapshotter:v6.1.0 • registry.k8s.io/sig-storage/csi-node-driver-registrar:v2.5.1 • docker.io/netapp/trident-operator:22.10.0 (optional)
v1.21.0	<ul style="list-style-type: none"> • docker.io/netapp/trident:22.10.0 • docker.io/netapp/trident-autosupport:22.10 • registry.k8s.io/sig-storage/csi-provisioner:v3.3.0 • registry.k8s.io/sig-storage/csi-attacher:v4.0.0 • registry.k8s.io/sig-storage/csi-resizer:v1.6.0 • registry.k8s.io/sig-storage/csi-snapshotter:v6.1.0 • registry.k8s.io/sig-storage/csi-node-driver-registrar:v2.5.1 • docker.io/netapp/trident-operator:22.10.0 (optional)
v1.22.0	<ul style="list-style-type: none"> • docker.io/netapp/trident:22.10.0 • docker.io/netapp/trident-autosupport:22.10 • registry.k8s.io/sig-storage/csi-provisioner:v3.3.0 • registry.k8s.io/sig-storage/csi-attacher:v4.0.0 • registry.k8s.io/sig-storage/csi-resizer:v1.6.0 • registry.k8s.io/sig-storage/csi-snapshotter:v6.1.0 • registry.k8s.io/sig-storage/csi-node-driver-registrar:v2.5.1 • docker.io/netapp/trident-operator:22.10.0 (optional)

Kubernetes version	Container image
v1.23.0	<ul style="list-style-type: none"> • docker.io/netapp/trident:22.10.0 • docker.io/netapp/trident-autosupport:22.10 • registry.k8s.io/sig-storage/csi-provisioner:v3.3.0 • registry.k8s.io/sig-storage/csi-attacher:v4.0.0 • registry.k8s.io/sig-storage/csi-resizer:v1.6.0 • registry.k8s.io/sig-storage/csi-snapshotter:v6.1.0 • registry.k8s.io/sig-storage/csi-node-driver-registrar:v2.5.1 • docker.io/netapp/trident-operator:22.10.0 (optional)
v1.24.0	<ul style="list-style-type: none"> • docker.io/netapp/trident:22.10.0 • docker.io/netapp/trident-autosupport:22.10 • registry.k8s.io/sig-storage/csi-provisioner:v3.3.0 • registry.k8s.io/sig-storage/csi-attacher:v4.0.0 • registry.k8s.io/sig-storage/csi-resizer:v1.6.0 • registry.k8s.io/sig-storage/csi-snapshotter:v6.1.0 • registry.k8s.io/sig-storage/csi-node-driver-registrar:v2.5.1 • docker.io/netapp/trident-operator:22.10.0 (optional)
v1.25.0	<ul style="list-style-type: none"> • docker.io/netapp/trident:22.10.0 • docker.io/netapp/trident-autosupport:22.10 • registry.k8s.io/sig-storage/csi-provisioner:v3.3.0 • registry.k8s.io/sig-storage/csi-attacher:v4.0.0 • registry.k8s.io/sig-storage/csi-resizer:v1.6.0 • registry.k8s.io/sig-storage/csi-snapshotter:v6.1.0 • registry.k8s.io/sig-storage/csi-node-driver-registrar:v2.5.1 • docker.io/netapp/trident-operator:22.10.0 (optional)



On Kubernetes version 1.20 and above, use the validated `registry.k8s.gcr.io/sig-storage/csi-snapshotter:v6.x` image only if the `v1` version is serving the `volumesnapshots.snapshot.storage.k8s.gcr.io` CRD. If the `v1beta1` version is serving the CRD with/without the `v1` version, use the validated `registry.k8s.gcr.io/sig-storage/csi-snapshotter:v3.x` image.

Deployment overview

You can deploy Astra Trident using the Trident operator or with `tridentctl`.



Beginning with the 22.04 release, AES keys will no longer be regenerated every time Astra Trident is installed. With this release, Astra Trident will install a new secret object that persists across installations. This means, `tridentctl` in 22.04 can uninstall previous versions of Trident, but earlier versions cannot uninstall 22.04 installations.

Critical information about Astra Trident 22.10

You must read the following critical information before upgrading to Astra Trident 22.10.

Critical information about Astra Trident 22.10



- Kubernetes 1.25 is now supported in Trident. You must upgrade to Astra Trident 22.10 prior to upgrading to Kubernetes 1.25.
- Astra Trident now strictly enforces the use of multipathing configuration in SAN environments, with a recommended value of `find_multipaths: no` in `multipath.conf` file.

Use of non-multipathing configuration or use of `find_multipaths: yes` or `find_multipaths: smart` value in `multipath.conf` file will result in mount failures. Trident has recommended the use of `find_multipaths: no` since the 21.07 release.

Choose the deployment method

To determine which deployment method to use, consider the following:

When to use the Trident operator

The [Trident operator](#) is a great way to dynamically manage Astra Trident resources and automate the setup phase. There are some prerequisites that must be satisfied. See [the requirements](#).

The Trident operator provides several benefits as outlined below.

Self-healing capability

You can monitor an Astra Trident installation and actively take measures to address issues, such as when the deployment is deleted or if it is modified accidentally. When the operator is set up as a deployment, a `trident-operator-<generated-id>` pod is created. This pod associates a `TridentOrchestrator` CR with an Astra Trident installation and always ensures there is only one active `TridentOrchestrator`. In other words, the operator ensures that there is only one instance of Astra Trident in the cluster and controls its setup, making sure the installation is idempotent. When changes are made to the installation (such as, deleting the deployment or node daemonset), the operator identifies them and fixes them individually.

Easy updates to existing installations

You can easily update an existing deployment with the operator. You only need to edit the `TridentOrchestrator` CR to make updates to an installation. For example, consider a scenario where you need to enable Astra Trident to generate debug logs.

To do this, patch your `TridentOrchestrator` to set `spec.debug` to `true`:


```
kubectl patch torc <trident-orchestrator-name> -n trident --type=merge -p
'{"spec":{"debug":true}}'
```

After `TridentOrchestrator` is updated, the operator processes the updates and patches the existing installation. This might trigger the creation of new pods to modify the installation accordingly.

Automatically handles Kubernetes upgrades

When the Kubernetes version of the cluster is upgraded to a supported version, the operator updates an existing Astra Trident installation automatically and changes it to ensure that it meets the requirements of the Kubernetes version.



If the cluster is upgraded to an unsupported version, the operator prevents installing Astra Trident. If Astra Trident has already been installed with the operator, a warning is displayed to indicate that Astra Trident is installed on an unsupported Kubernetes version.

Manage your Kubernetes clusters using BlueXP (formerly Cloud Manager)

With [Astra Trident using BlueXP](#), you can upgrade to the latest version of Astra Trident, add and manage storage classes and connect them to Working Environments, and back up persistent volumes using Cloud Backup Service. BlueXP supports Astra Trident deployment using the Trident operator, either manually or using Helm.

When to use Helm

If you have other applications that you are managing using Helm, starting with Astra Trident 21.01, you can manage your deployment also using Helm.

When to use `tridentctl`

If you have an existing deployment that must be upgraded to or if you are looking to highly customize your deployment, you should consider using [tridentctl](#). This is the conventional method of deploying Astra Trident.

Considerations for moving between deployment methods

It is not hard to imagine a scenario where moving between deployment methods is desired. You should consider the following before attempting to move from a `tridentctl` deployment to an operator-based deployment, or vice-versa:

- Always use the same method for uninstalling Astra Trident. If you have deployed with `tridentctl`, you should use the appropriate version of the `tridentctl` binary to uninstall Astra Trident. Similarly, if you are deploying with the operator, you should edit the `TridentOrchestrator` CR and set `spec.uninstall=true` to uninstall Astra Trident.
- If you have an operator-based deployment that you want to remove and use `tridentctl` to deploy Astra Trident, you should first edit `TridentOrchestrator` and set `spec.uninstall=true` to uninstall Astra Trident. Then delete `TridentOrchestrator` and the operator deployment. You can then install using `tridentctl`.
- If you have a manual operator-based deployment, and you want to use Helm-based Trident operator deployment, you should manually uninstall the operator first, and then do the Helm install. This enables Helm to deploy the Trident operator with the required labels and annotations. If you do not do this, your

Helm-based Trident operator deployment will fail with label validation error and annotation validation error. If you have a `tridentctl`-based deployment, you can use Helm-based deployment without running into issues.

Understand the deployment modes

There are three ways to deploy Astra Trident.

Standard deployment

Deploying Trident on a Kubernetes cluster results in the Astra Trident installer doing two things:

- Fetching the container images over the Internet
- Creating a deployment and/or node daemonset, which spins up Astra Trident pods on all the eligible nodes in the Kubernetes cluster.

A standard deployment such as this can be performed in two different ways:

- Using `tridentctl install`
- Using the Trident operator. You can deploy Trident operator either manually or by using Helm.

This mode of installing is the easiest way to install Astra Trident and works for most environments that do not impose network restrictions.

Offline deployment

To perform an air-gapped deployment, you can use the `--image-registry` flag when invoking `tridentctl install` to point to a private image registry. If deploying with the Trident operator, you can alternatively specify `spec.imageRegistry` in your `TridentOrchestrator`. This registry should contain the [Trident image](#), the [Trident Autosupport image](#), and the CSI sidecar images as required by your Kubernetes version.

To customize your deployment, you can use `tridentctl` to generate the manifests for Trident's resources. This includes the deployment, daemonset, service account, and the cluster role that Astra Trident creates as part of its installation.

See these links for more information about customizing your deployment:

- [Customize your operator-based deployment](#)
- [Customize your `tridentctl`-based deployment](#)



If you are using a private image repository, you should add `/sig-storage` to the end of the private registry URL. When using a private registry for `tridentctl` deployment, you should use `--trident-image` and `--autosupport-image` in conjunction with `--image-registry`. If you are deploying Astra Trident by using the Trident operator, ensure that the orchestrator CR includes `tridentImage` and `autosupportImage` in the installation parameters.

Remote deployment

Here is a high-level overview of the remote deployment process:

- Deploy the appropriate version of `kubectl` on the remote machine from where you want to deploy Astra

Trident.

- Copy the configuration files from the Kubernetes cluster and set the `KUBECONFIG` environment variable on the remote machine.
- Initiate a `kubectl get nodes` command to verify that you can connect to the required Kubernetes cluster.
- Complete the deployment from the remote machine by using the standard installation steps.

Other known configuration options

When installing Astra Trident on VMWare Tanzu Portfolio products:

- The cluster must support privileged workloads.
- The `--kubelet-dir` flag should be set to the location of kubelet directory. By default, this is `/var/vcap/data/kubelet`.

Specifying the kubelet location using `--kubelet-dir` is known to work for Trident Operator, Helm, and `tridentctl` deployments.

Deploy with Trident operator

You can deploy Astra Trident using the Trident operator.

Critical information about Astra Trident 22.10

You must read the following critical information before upgrading to Astra Trident 22.10.

Critical information about Astra Trident 22.10

- Kubernetes 1.25 is now supported in Trident. You must upgrade to Astra Trident 22.10 prior to upgrading to Kubernetes 1.25.
- Astra Trident now strictly enforces the use of multipathing configuration in SAN environments, with a recommended value of `find_multipaths: no` in `multipath.conf` file.



Use of non-multipathing configuration or use of `find_multipaths: yes` or `find_multipaths: smart` value in `multipath.conf` file will result in mount failures. Trident has recommended the use of `find_multipaths: no` since the 21.07 release.

Trident operator deployment options

You can deploy the Trident operator in one of two ways:

- Using the Trident [Helm Chart](#): The Helm Chart deploys the Trident operator and installs Trident in one step.
- Manually: Trident provides a file that can be used to install the operator and create associated objects.
 - For clusters running Kubernetes 1.24 or lower, use [bundle_pre_1_25.yaml](#).
 - For clusters running Kubernetes 1.25 or higher, use [bundle_post_1_25.yaml](#).



If you have not already familiarized yourself with the [basic concepts](#), now is a great time to do that.

Verify prerequisites

To deploy Astra Trident, the following prerequisites should be met:

- You have full privileges to a supported Kubernetes cluster running a supported version of Kubernetes. Review the [requirements](#).
- You have access to a supported NetApp storage system.
- You have the capability to mount volumes from all of the Kubernetes worker nodes.
- You have a Linux host with `kubectl` (or `oc`, if you are using OpenShift) installed and configured to manage the Kubernetes cluster that you want to use.
- You have set the `KUBECONFIG` environment variable to point to your Kubernetes cluster configuration.
- You have enabled the [feature gates required by Astra Trident](#).
- If you are using Kubernetes with Docker Enterprise, [follow their steps to enable CLI access](#).

Got all that? Great! Let's get started.

Deploy the Trident operator and install Astra Trident using Helm

Perform the steps listed to deploy the Trident operator by using Helm.

What you'll need

In addition to the prerequisites listed above, to deploy Trident operator by using Helm, you need the following:

- A [supported Kubernetes version](#)
- Helm version 3

Steps

1. Add Trident's Helm repository:

```
helm repo add netapp-trident https://netapp.github.io/trident-helm-chart
```

2. Use the `helm install` command and specify a name for your deployment.
See the following example:

```
helm install <name> netapp-trident/trident-operator --version 22.10.0  
--create-namespace --namespace <trident-namespace>
```



If you already created a namespace for Trident, the `--create-namespace` parameter will not create an additional namespace.

There are two ways to pass configuration data during the install:

- `--values` (or `-f`): Specify a YAML file with overrides. This can be specified multiple times and the rightmost file will take precedence.
- `--set`: Specify overrides on the command line.

For example, to change the default value of `debug`, run the following `--set` command:

```
helm install <name> netapp-trident/trident-operator --version 22.10.0
--create-namespace --namespace --set tridentDebug=true
```

The `values.yaml` file, which is part of the Helm chart provides the list of keys and their default values.

`helm list` shows you details about the installation, such as name, namespace, chart, status, app version, revision number, and so on.

Deploy the Trident operator manually and install Trident

Perform the steps listed to manually deploy the Trident operator.

Step 1: Qualify your Kubernetes cluster

The first thing you need to do is log in to the Linux host and verify that it is managing a *working, supported Kubernetes cluster* that you have the necessary privileges to.



With OpenShift, use `oc` instead of `kubectl` in all of the examples that follow, and log in as **system:admin** first by running `oc login -u system:admin` or `oc login -u kube-admin`.

To verify your Kubernetes version, run the following command:

```
kubectl version
```

To see if you have Kubernetes cluster administrator privileges, run the following command:

```
kubectl auth can-i '*' '*' --all-namespaces
```

To verify if you can launch a pod that uses an image from Docker Hub and reach your storage system over the pod network, run the following command:

```
kubectl run -i --tty ping --image=busybox --restart=Never --rm -- \
ping <management IP>
```

Step 2: Download and set up the operator



Beginning with 21.01, the Trident operator is cluster scoped. Using the Trident operator to install Trident requires creating the `TridentOrchestrator` Custom Resource Definition (CRD) and defining other resources. You should perform these steps to set up the operator before you can install Astra Trident.

1. Download and extract the latest version of the Trident installer bundle from [the Assets section on GitHub](#).

```
wget
https://github.com/NetApp/trident/releases/download/v22.10.0/trident-
installer-22.10.0.tar.gz
tar -xf trident-installer-22.10.0.tar.gz
cd trident-installer
```

2. Use the appropriate CRD manifest to create the `TridentOrchestrator` CRD. You then create a `TridentOrchestrator` Custom Resource later on to instantiate an installation by the operator.

Run the following command:

```
kubectl create -f
deploy/crds/trident.netapp.io_tridentorchestrators_crd_post1.16.yaml
```

3. After the `TridentOrchestrator` CRD is created, create the following resources required for the operator deployment:

- A `ServiceAccount` for the operator
- A `ClusterRole` and `ClusterRoleBinding` to the `ServiceAccount`
- A dedicated `PodSecurityPolicy`
- The operator itself

The Trident installer contains manifests for defining these resources. By default, the operator is deployed in the `trident` namespace. If the `trident` namespace does not exist, use the following manifest to create one.

```
kubectl apply -f deploy/namespace.yaml
```

4. To deploy the operator in a namespace other than the default `trident` namespace, you should update the `serviceaccount.yaml`, `clusterrolebinding.yaml` and `operator.yaml` manifests and generate your `bundle.yaml`.

Run the following command to update the YAML manifests and generate your `bundle.yaml` using the `kustomization.yaml`:

```
kubectl kustomize deploy/ > deploy/bundle.yaml
```

Run the following command to create the resources and deploy the operator:

```
kubectl create -f deploy/bundle.yaml
```

5. To verify the status of the operator after you have deployed, do the following:

```
kubectl get deployment -n <operator-namespace>
```

NAME	READY	UP-TO-DATE	AVAILABLE	AGE
trident-operator	1/1	1	1	3m

```
kubectl get pods -n <operator-namespace>
```

NAME	READY	STATUS	RESTARTS
trident-operator-54cb664d-lnjxh	1/1	Running	0
3m			

The operator deployment successfully creates a pod running on one of the worker nodes in your cluster.



There should only be **one instance** of the operator in a Kubernetes cluster. Do not create multiple deployments of the Trident operator.

Step 3: Create `TridentOrchestrator` and install Trident

You are now ready to install Astra Trident using the operator! This will require creating `TridentOrchestrator`. The Trident installer comes with example definitions for creating `TridentOrchestrator`. This kicks off an installation in the `trident` namespace.

```

kubect1 create -f deploy/crds/tridentorchestrator_cr.yaml
tridentorchestrator.trident.netapp.io/trident created

kubect1 describe torc trident
Name:          trident
Namespace:
Labels:        <none>
Annotations:   <none>
API Version:   trident.netapp.io/v1
Kind:          TridentOrchestrator
...
Spec:
  Debug:       true
  Namespace:   trident
Status:
  Current Installation Params:
    IPv6:              false
    Autosupport Hostname:
    Autosupport Image: netapp/trident-autosupport:22.10
    Autosupport Proxy:
    Autosupport Serial Number:
    Debug:              true
    Image Pull Secrets:
    Image Registry:
    k8sTimeout:         30
    Kubelet Dir:        /var/lib/kubelet
    Log Format:          text
    Silence Autosupport: false
    Trident Image:      netapp/trident:21.04.0
  Message:            Trident installed Namespace:
trident
  Status:              Installed
  Version:              v21.04.0
Events:
  Type Reason Age From Message ---- -
  Installing 74s trident-operator.netapp.io Installing Trident Normal
  Installed 67s trident-operator.netapp.io Trident installed

```

The Trident operator enables you to customize the manner in which Astra Trident is installed by using the attributes in the `TridentOrchestrator` spec. See [Customize your Trident deployment](#).

The Status of `TridentOrchestrator` indicates if the installation was successful and displays the version of Trident installed.

Status	Description
Installing	The operator is installing Astra Trident using this <code>TridentOrchestrator</code> CR.
Installed	Astra Trident has successfully installed.
Uninstalling	The operator is uninstalling Astra Trident, because <code>spec.uninstall=true</code> .
Uninstalled	Astra Trident is uninstalled.
Failed	The operator could not install, patch, update or uninstall Astra Trident; the operator will automatically try to recover from this state. If this state persists you will require troubleshooting.
Updating	The operator is updating an existing installation.
Error	The <code>TridentOrchestrator</code> is not used. Another one already exists.

During the installation, the status of `TridentOrchestrator` changes from `Installing` to `Installed`. If you observe the `Failed` status and the operator is unable to recover by itself, you should check the logs of the operator. See the [troubleshooting](#) section.

You can confirm if the Astra Trident installation completed by taking a look at the pods that have been created:

```
kubectl get pod -n trident
```

NAME	READY	STATUS	RESTARTS	AGE
trident-csi-7d466bf5c7-v4cpw	5/5	Running	0	1m
trident-csi-mr6zc	2/2	Running	0	1m
trident-csi-xrp7w	2/2	Running	0	1m
trident-csi-zh2jt	2/2	Running	0	1m
trident-operator-766f7b8658-ldzsv	1/1	Running	0	3m

You can also use `tridentctl` to check the version of Astra Trident installed.

```
./tridentctl -n trident version
```

```
+-----+-----+
| SERVER VERSION | CLIENT VERSION |
+-----+-----+
| 21.04.0       | 21.04.0       |
+-----+-----+
```

Now you can go ahead and create a backend. See [post-deployment tasks](#).



For troubleshooting issues during deployment, see the [troubleshooting](#) section.

Customize Trident operator deployment

The Trident operator allows you to customize Astra Trident installation using the attributes in the `TridentOrchestrator` spec.

If you want to customize the installation beyond what `TridentOrchestrator` arguments allow, you should consider using `tridentctl` to generate custom YAML manifests that you can modify as needed.



`spec.namespace` is specified in `TridentOrchestrator` to signify which namespace where Astra Trident is installed. This parameter **cannot be updated after Astra Trident is installed**. Attempting to do so causes the `TridentOrchestrator` status to change to `Failed`. Astra Trident is not intended to be migrated across namespaces.

Configuration options

This table details `TridentOrchestrator` attributes:

Parameter	Description	Default
<code>namespace</code>	Namespace to install Astra Trident in	"default"
<code>debug</code>	Enable debugging for Astra Trident	false
<code>windows</code>	Setting to <code>true</code> enables installation on Windows worker nodes.	false
<code>IPv6</code>	Install Astra Trident over IPv6	false
<code>k8sTimeout</code>	Timeout for Kubernetes operations	30sec
<code>silenceAutosupport</code>	Don't send autosupport bundles to NetApp automatically	false
<code>enableNodePrep</code>	Manage worker node dependencies automatically (BETA)	false
<code>autosupportImage</code>	The container image for Autosupport Telemetry	"netapp/trident-autosupport:22.10.0"
<code>autosupportProxy</code>	The address/port of a proxy for sending Autosupport Telemetry	"http://proxy.example.com:8888"
<code>uninstall</code>	A flag used to uninstall Astra Trident	false
<code>logFormat</code>	Astra Trident logging format to be used [text,json]	"text"
<code>tridentImage</code>	Astra Trident image to install	"netapp/trident:21.04"

Parameter	Description	Default
imageRegistry	Path to internal registry, of the format <registry FQDN>[:port] [/subpath]	"k8s.gcr.io/sig-storage (k8s 1.19+) or quay.io/k8scsi"
kubeletDir	Path to the kubelet directory on the host	"/var/lib/kubelet"
wipeout	A list of resources to delete to perform a complete removal of Astra Trident	
imagePullSecrets	Secrets to pull images from an internal registry	
controllerPluginNodeSelector	Additional node selectors for pods running the Trident Controller CSI Plugin. Follows same format as pod.spec.nodeSelector.	No default; optional
controllerPluginTolerations	Overrides tolerations for pods running the Trident Controller CSI Plugin. Follows the same format as pod.spec.Tolerations.	No default; optional
nodePluginNodeSelector	Additional node selectors for pods running the Trident Node CSI Plugin. Follows same format as pod.spec.nodeSelector.	No default; optional
nodePluginTolerations	Overrides tolerations for pods running the Trident Node CSI Plugin. Follows the same format as pod.spec.Tolerations.	No default; optional



For more information on formatting pod parameters, see [Assigning Pods to Nodes](#).

Sample configurations

You can use the attributes mentioned above when defining `TridentOrchestrator` to customize your installation.

Example 1: Basic custom configuration

This is an example for a basic custom configuration.

```
cat deploy/crds/tridentorchestrator_cr_imagepullsecrets.yaml
apiVersion: trident.netapp.io/v1
kind: TridentOrchestrator
metadata:
  name: trident
spec:
  debug: true
  namespace: trident
  imagePullSecrets:
  - thisisasecret
```

Example 2: Deploy with node selectors

This example illustrates how Trident can be deployed with node selectors:

```
apiVersion: trident.netapp.io/v1
kind: TridentOrchestrator
metadata:
  name: trident
spec:
  debug: true
  namespace: trident
  controllerPluginNodeSelector:
    nodetype: master
  nodePluginNodeSelector:
    storage: netapp
```

Example 3: Deploy on Windows worker nodes

This example illustrates deployment on a Windows worker node.

```
$ cat deploy/crds/tridentorchestrator_cr.yaml
apiVersion: trident.netapp.io/v1
kind: TridentOrchestrator
metadata:
  name: trident
spec:
  debug: true
  namespace: trident
  windows: true
```

Deploy with `tridentctl`

You can deploy Astra Trident using `tridentctl`. It's a good idea to familiarize yourself with the [basic concepts](#). To customize your `tridentctl` deployment, refer to [Customize tridentctl deployment](#).

Critical information about Astra Trident 22.10

You must read the following critical information before upgrading to Astra Trident 22.10.

Critical information about Astra Trident 22.10

- Kubernetes 1.25 is now supported in Trident. You must upgrade to Astra Trident 22.10 prior to upgrading to Kubernetes 1.25.
- Astra Trident now strictly enforces the use of multipathing configuration in SAN environments, with a recommended value of `find_multipaths: no` in `multipath.conf` file.



Use of non-multipathing configuration or use of `find_multipaths: yes` or `find_multipaths: smart` value in `multipath.conf` file will result in mount failures. Trident has recommended the use of `find_multipaths: no` since the 21.07 release.

Verify prerequisites

To deploy Astra Trident, the following prerequisites should be met:

- Full privileges to a supported Kubernetes cluster.
- Access to a supported NetApp storage system.
- Capability to mount volumes from all of the Kubernetes worker nodes.
- A Linux host with `kubectl` (or `oc`, if you are using OpenShift) installed and configured to manage the Kubernetes cluster that you want to use.
- The `KUBECONFIG` environment variable points to your Kubernetes cluster configuration.

- The [feature gates required by Astra Trident](#) are enabled.
- If you are using Kubernetes with Docker Enterprise, [follow their steps to enable CLI access](#).

Step 1: Qualify your Kubernetes cluster

Log into the Linux host and verify it is managing a working, [supported Kubernetes cluster](#) and you have the necessary privileges.



With OpenShift, you use `oc` instead of `kubectl` in all of the examples that follow, and you should log in as **system:admin** first by running `oc login -u system:admin` or `oc login -u kube-admin`.

To check your Kubernetes version, run the following command:

```
kubectl version
```

To verify Kubernetes cluster administrator privileges, run the following command:

```
kubectl auth can-i '*' '*' --all-namespaces
```

To verify if you can launch a pod that uses an image from Docker Hub and reach your storage system over the pod network, run the following command:

```
kubectl run -i --tty ping --image=busybox --restart=Never --rm -- \
ping <management IP>
```

Identify your Kubernetes server version. You will use it when you install Astra Trident.

Step 2: Download and extract the installer



The Trident installer creates a Trident pod, configures the CRD objects that are used to maintain its state, and initializes the CSI sidecars that perform actions, such as provisioning and attaching volumes to the cluster hosts.

You can download and extract the latest version of the Trident installer bundle from [the Assets section on GitHub](#).

For example, if the latest version is 22.10.0:

```
wget https://github.com/NetApp/trident/releases/download/v22.10.0/trident-
installer-22.10.0.tar.gz
tar -xf trident-installer-22.10.0.tar.gz
cd trident-installer
```

Step 3: Install Astra Trident

Install Astra Trident in the desired namespace by executing the `tridentctl install` command.

```
./tridentctl install -n trident
....
INFO Starting Trident installation.                namespace=trident
INFO Created service account.
INFO Created cluster role.
INFO Created cluster role binding.
INFO Added finalizers to custom resource definitions.
INFO Created Trident service.
INFO Created Trident secret.
INFO Created Trident deployment.
INFO Created Trident daemonset.
INFO Waiting for Trident pod to start.
INFO Trident pod started.                        namespace=trident
pod=trident-csi-679648bd45-cv2mx
INFO Waiting for Trident REST interface.
INFO Trident REST interface is up.                version=22.10.0
INFO Trident installation succeeded.
....
```



To enable Astra Trident to run on Windows nodes, add the `--windows` flag to the install command: `$./tridentctl install --windows -n trident`.

Output similar to the following displays when the installer is complete. Depending on the number of nodes in your Kubernetes cluster, there might be more pods present:

```
kubectl get pod -n trident
NAME                                READY   STATUS    RESTARTS   AGE
trident-csi-679648bd45-cv2mx       4/4    Running   0           5m29s
trident-csi-vgc8n                   2/2    Running   0           5m29s

./tridentctl -n trident version
+-----+-----+
| SERVER VERSION | CLIENT VERSION |
+-----+-----+
| 22.10.0        | 22.10.0        |
+-----+-----+
```

To complete Astra Trident configuration, continue to [post-deployment tasks](#).

If the installer does not complete successfully or `trident-csi-<generated id>` does not have a **Running** status, the platform was not installed.



For troubleshooting issues during deployment, refer to [troubleshooting](#).

Customize `tridentctl` deployment

You can use the Astra Trident installer to customize deployment.

Learn about the installer

The Astra Trident installer enables you to customize attributes. For example, if you have copied the Trident image to a private repository, you can specify the image name by using `--trident-image`. If you have copied the Trident image as well as the needed CSI sidecar images to a private repository, it might be preferable to specify the location of that repository by using the `--image-registry` switch, which takes the form `<registry FQDN>[:port]`.

If you are using a distribution of Kubernetes, where `kubelet` keeps its data on a path other than the usual `/var/lib/kubelet`, you can specify the alternate path by using `--kubelet-dir`.

If you need to customize the installation beyond what the installer's arguments allow, you can also customize the deployment files. Using the `--generate-custom-yaml` parameter creates the following YAML files in the installer's setup directory:

- `trident-clusterrolebinding.yaml`
- `trident-deployment.yaml`
- `trident-crds.yaml`
- `trident-clusterrole.yaml`
- `trident-daemonset.yaml`
- `trident-service.yaml`
- `trident-namespace.yaml`
- `trident-serviceaccount.yaml`
- `trident-resourcequota.yaml`

After you have generated these files, you can modify them according to your needs and then use `--use-custom-yaml` to install your custom deployment.

```
./tridentctl install -n trident --use-custom-yaml
```

What's next?

After you deploy Astra Trident, you can proceed with creating a backend, creating a storage class, provisioning a volume, and mounting the volume in a pod.

Step 1: Create a backend

You can now go ahead and create a backend that will be used by Astra Trident to provision volumes. To do

this, create a `backend.json` file that contains the necessary parameters. Sample configuration files for different backend types can be found in the `sample-input` directory.

See [here](#) for more details about how to configure the file for your backend type.

```
cp sample-input/<backend template>.json backend.json
vi backend.json
```

```
./tridentctl -n trident create backend -f backend.json
+-----+-----+-----+
+-----+-----+
|   NAME   | STORAGE DRIVER |           UUID           |
STATE | VOLUMES |
+-----+-----+-----+
+-----+-----+
| nas-backend | ontap-nas      | 98e19b74-aec7-4a3d-8dcf-128e5033b214 |
online |         0 |
+-----+-----+-----+
+-----+-----+

```

If the creation fails, something was wrong with the backend configuration. You can view the logs to determine the cause by running the following command:

```
./tridentctl -n trident logs
```

After you address the problem, simply go back to the beginning of this step and try again. For more troubleshooting tips, see [the troubleshooting](#) section.

Step 2: Create a storage class

Kubernetes users provision volumes by using persistent volume claims (PVCs) that specify a [storage class](#) by name. The details are hidden from the users, but a storage class identifies the provisioner that is used for that class (in this case, Trident), and what that class means to the provisioner.

Create a storage class Kubernetes users will specify when they want a volume. The configuration of the class needs to model the backend that you created in the previous step, so that Astra Trident will use it to provision new volumes.

The simplest storage class to start with is one based on the `sample-input/storage-class-csi.yaml.template` file that comes with the installer, replacing `BACKEND_TYPE` with the storage driver name.

```

./tridentctl -n trident get backend
+-----+-----+-----+-----+
+-----+-----+
|   NAME   | STORAGE DRIVER |           UUID           |
STATE | VOLUMES |
+-----+-----+-----+-----+
+-----+-----+
| nas-backend | ontap-nas      | 98e19b74-aec7-4a3d-8dcf-128e5033b214 |
online |         0 |
+-----+-----+-----+-----+
+-----+-----+

cp sample-input/storage-class-csi.yaml.templ sample-input/storage-class-
basic-csi.yaml

# Modify __BACKEND_TYPE__ with the storage driver field above (e.g.,
ontap-nas)
vi sample-input/storage-class-basic-csi.yaml

```

This is a Kubernetes object, so you use `kubectl` to create it in Kubernetes.

```
kubectl create -f sample-input/storage-class-basic-csi.yaml
```

You should now see a **basic-csi** storage class in both Kubernetes and Astra Trident, and Astra Trident should have discovered the pools on the backend.

```

kubect1 get sc basic-csi
NAME          PROVISIONER          AGE
basic-csi     csi.trident.netapp.io 15h

./tridentctl -n trident get storageclass basic-csi -o json
{
  "items": [
    {
      "Config": {
        "version": "1",
        "name": "basic-csi",
        "attributes": {
          "backendType": "ontap-nas"
        },
        "storagePools": null,
        "additionalStoragePools": null
      },
      "storage": {
        "ontapnas_10.0.0.1": [
          "aggr1",
          "aggr2",
          "aggr3",
          "aggr4"
        ]
      }
    }
  ]
}

```

Step 3: Provision your first volume

Now you are ready to dynamically provision your first volume. This is done by creating a Kubernetes [persistent volume claim](#) (PVC) object.

Create a PVC for a volume that uses the storage class that you just created.

See `sample-input/pvc-basic-csi.yaml` for an example. Make sure the storage class name matches the one that you created.

```
kubectl create -f sample-input/pvc-basic-csi.yaml
```

```
kubectl get pvc --watch
```

NAME	STATUS	VOLUME	CAPACITY
basic	Pending		
basic	1s		
basic	Pending	pvc-3acb0d1c-b1ae-11e9-8d9f-5254004dfdb7	0
basic	5s		
basic	Bound	pvc-3acb0d1c-b1ae-11e9-8d9f-5254004dfdb7	1Gi
RWO	basic	7s	

Step 4: Mount the volumes in a pod

Now let us mount the volume. We will launch an nginx pod that mounts the PV under /usr/share/nginx/html.

```
cat << EOF > task-pv-pod.yaml
kind: Pod
apiVersion: v1
metadata:
  name: task-pv-pod
spec:
  volumes:
    - name: task-pv-storage
      persistentVolumeClaim:
        claimName: basic
  containers:
    - name: task-pv-container
      image: nginx
      ports:
        - containerPort: 80
          name: "http-server"
      volumeMounts:
        - mountPath: "/usr/share/nginx/html"
          name: task-pv-storage
EOF
kubectl create -f task-pv-pod.yaml
```

```
# Wait for the pod to start
kubectl get pod --watch

# Verify that the volume is mounted on /usr/share/nginx/html
kubectl exec -it task-pv-pod -- df -h /usr/share/nginx/html

# Delete the pod
kubectl delete pod task-pv-pod
```

At this point, the pod (application) no longer exists but the volume is still there. You can use it from another pod if you want to.

To delete the volume, delete the claim:

```
kubectl delete pvc basic
```

You can now do additional tasks, such as the following:

- [Configure additional backends.](#)
- [Create additional storage classes.](#)

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.