
Trident Documentation

NetApp

Feb 21, 2020

1	What is Trident?	3
2	Trident for Kubernetes	5
2.1	Upgrading Trident	6
2.2	Deploying	10
2.3	Common tasks	16
2.4	Concepts	44
2.5	Known issues	52
2.6	Troubleshooting	52
3	Design and Architecture Guide	55
3.1	Introduction	55
3.2	Concepts and Definitions	56
3.3	NetApp Products and Integrations with Kubernetes	60
3.4	Kubernetes Cluster Architecture and Considerations	61
3.5	Storage for Kubernetes Infrastructure Services	64
3.6	Storage Configuration for Trident	70
3.7	Deploying Trident	76
3.8	Integrating Trident	79
3.9	Backup and Disaster Recovery	88
3.10	Security Recommendations	93
3.11	Frequently Asked Questions	93
4	Trident for Docker	103
4.1	Deploying	103
4.2	Host and storage configuration	104
4.3	Common tasks	121
4.4	Known issues	127
4.5	Troubleshooting	128
5	Requirements	129
5.1	Supported frontends (orchestrators)	129
5.2	Supported backends (storage)	129
5.3	Feature Gates	129
5.4	Supported host operating systems	130
5.5	Host configuration	130
5.6	Storage system configuration	130

6	Getting help	131
7	trident	133
7.1	Logging	133
7.2	Kubernetes	133
7.3	Docker	133
7.4	REST	134
8	tridentctl	135
8.1	create	135
8.2	delete	136
8.3	get	136
8.4	import volume	136
8.5	install	137
8.6	logs	137
8.7	uninstall	137
8.8	update	138
8.9	upgrade	138
8.10	version	138
9	REST API	139
10	Simple Kubernetes install	141
10.1	Prerequisites	141
10.2	Install Docker CE 17.03	141
10.3	Install the appropriate version of kubeadm, kubectl and kubelet	142
10.4	Configure the host	142
10.5	Create the cluster	142
10.6	Install the kubectl creds and untaint the cluster	142
10.7	Add an overlay network	142
10.8	Verify that all of the services started	142

Storage Orchestrator for Containers

CHAPTER 1

What is Trident?

Trident is a fully supported [open source project](#) maintained by [NetApp](#). It has been designed from the ground up to help you meet the sophisticated persistence demands of your containerized applications.

Through its support for popular container platforms like [Kubernetes](#) and [Docker](#), Trident understands the natural and evolving languages of those platforms, and translates requirements expressed or implied through them into an automated and orchestrated response from the infrastructure.

Today, that includes our [ONTAP \(AFF/FAS/Select/Cloud\)](#), [Element \(HCI/SolidFire\)](#), and [SANtricity \(E/EF-Series\)](#) data management software, plus the [Azure NetApp Files](#) service in [Azure](#) and our [Cloud Volumes Service](#) in [AWS](#). That list continues to grow.

Trident for Kubernetes

Trident integrates natively with [Kubernetes](#) and its [Persistent Volume framework](#) to seamlessly provision and manage volumes from systems running any combination of NetApp's [ONTAP \(AFF/FAS/Select/Cloud\)](#), [Element \(HCI/SolidFire\)](#), and [SANtricity \(E/EF-Series\)](#) data management platforms, plus our [Azure NetApp Files](#) service in Azure and [Cloud Volumes Service](#) in AWS.

Relative to other Kubernetes provisioners, Trident is novel in the following respects:

1. It is the first out-of-tree, out-of-process storage provisioner that works by watching events at the Kubernetes API Server, affording it levels of visibility and flexibility that cannot otherwise be achieved.
2. It is capable of orchestrating across multiple platforms at the same time through a unified interface. Rather than tying a request for a persistent volume to a particular system, Trident selects one from those it manages based on the higher-level qualities that the user is looking for in their volume.

Trident tightly integrates with Kubernetes to allow your users to request and manage persistent volumes using native Kubernetes interfaces and constructs. It's designed to work in such a way that your users can take advantage of the underlying capabilities of your storage infrastructure without having to know anything about it.

It automates the rest for you, the Kubernetes administrator, based on policies that you define.

A great way to get a feel for what we're trying to accomplish is to see Trident in action from the [perspective of an end user](#). This is a great demonstration of Kubernetes volume consumption when Trident is in the mix, through the lens of Red Hat's OpenShift platform, which is itself built on Kubernetes. All of the concepts in the video apply to any Kubernetes deployment.

While some details about Trident and NetApp storage systems are shown in the video to help you see what's going on behind-the-scenes, in standard deployments Trident and the rest of the infrastructure is completely hidden from the user.

Let's lift up the covers a bit to better understand Trident and what it is doing. This [introductory video](#) is a good place to start. While based on an earlier version of Trident, it explains the core concepts involved and is a great way to get started on the Trident journey.

2.1 Upgrading Trident

This section walks you through the upgrade process to move to the latest release of Trident.

2.1.1 Initiate the upgrade

Note: Before upgrading Trident, ensure that the required *feature gates* are enabled.

Warning: The 19.07.2 release of Trident removes support for **alpha Volume Snapshots**. Beginning with the 20.01 release, Trident only enables the creation of **beta Volume Snapshots**. When upgrading to 19.07.2 or above, all previous alpha snapshot CRs and CRDs (Volume Snapshot Classes, Volume Snapshots and Volume Snapshot Contents) must be removed before the upgrade is performed. Refer to [this blog](#) to understand the steps involved in migrating alpha snapshots to the beta spec.

The best way to upgrade to the latest version of Trident is to download the latest [installer bundle](#) and run:

```
./tridentctl uninstall -n <namespace>
./tridentctl install -n <namespace>
```

Upgrading Trident on Kubernetes 1.13

On Kubernetes 1.13, there are a couple of options when upgrading Trident:

- Install Trident in the desired namespace by executing the `tridentctl install` command with the `--csi` flag. This configures Trident to function as an enhanced CSI provisioner and is the preferred way to upgrade if using 1.13.
- If for some reason the *feature gates* required by Trident cannot be enabled, you can install Trident without the `--csi` flag. This will configure Trident to work in its traditional format without using the CSI specification.

Upgrading Trident on Kubernetes 1.14 and above

For Kubernetes 1.14 and greater, simply perform an uninstall followed by a reinstall to upgrade to the latest version of Trident.

2.1.2 What happens when you upgrade

By default the uninstall command will leave all of Trident's state intact by not deleting the PVC and PV used by the Trident deployment, allowing an uninstall followed by an install to act as an upgrade.

The 19.07 release of Trident will move away from etcd and use CRDs to maintain state. Upgrading Trident by running a `tridentctl install` will:

- initiate a one-time process that copies the metadata stored in the Trident PV into CRD objects.
- provide periodic updates throughout the upgrade with respect to the migration of Trident's metadata.

Warning: When upgrading Trident, do not interrupt the upgrade process. Ensure that the installer runs to completion.

PVs that have already been provisioned will remain available while Trident is offline, and Trident will provision volumes for any PVCs that are created in the interim once it is back online.

2.1.3 Next steps

When upgrading on Kubernetes versions 1.14 and above, the CSI provisioner will be used by default. For legacy PVs, all features made available by the previous Trident version will be supported. To make use of the rich set of features that will be provided in newer Trident releases, volumes can be upgraded using the `tridentctl upgrade` command.

The *Upgrading legacy volumes to CSI volumes* section explains how legacy non-CSI volumes can be upgraded to the CSI type.

If you encounter any issues, visit the *troubleshooting guide* for more advice.

2.1.4 Upgrading legacy volumes to CSI volumes

When upgrading Trident, there is a possibility of having legacy volumes that need to be ported to the CSI specification to make use of the complete set of features made available by Trident. A legacy PV that has been provisioned by Trident will still support the traditional set of features. For all additional features that Trident will provide, the Trident volume must be upgraded from a NFS/iSCSI type to the CSI type.

Before proceeding, you must determine if your Trident deployment is capable of upgrading legacy volumes. The version of Trident installed should be at least 19.07 and Trident should be configured as a CSI Provisioner. This can be confirmed by doing a `kubectl get pods -n <trident-namespace>`. The presence of a `trident-csi-<generated-id>` indicates that it is running as a CSI provisioner and supports upgrading legacy volumes.

```
$ tridentctl version
+-----+-----+
| SERVER VERSION | CLIENT VERSION |
+-----+-----+
| 19.07.2        | 19.07.2        |
+-----+-----+

$ kubectl get pods -n <trident-namespace>
NAME                                READY   STATUS    RESTARTS   AGE
trident-csi-426nx                   2/2     Running   0           20m
trident-csi-b5cf8fd7c-fnq24        4/4     Running   0           20m
```

Things to consider when upgrading volumes

When deciding to upgrade volumes to the CSI type, make sure to consider the following:

- It may not be required to upgrade all existing volumes. Previously created volumes will still continue to be accessible and function normally.
- A PV can be mounted as part of a Deployment/StatefulSet when upgrading. It is not required to bring down the Deployment/StatefulSet.

- A PV **cannot** be attached to a standalone pod when upgrading. You will have to shut down the pod before upgrading the volume.
- To upgrade a volume, it must be bound to a PVC. Volumes that are not bound to PVCs will need to be removed and imported before upgrading.

Example volume upgrade

Here is an example that shows how a volume upgrade is performed.

```
$ kubectl get pv
NAME                                CAPACITY    ACCESS MODES    RECLAIM POLICY    STATUS
↪CLAIM                              STORAGECLASS  REASON          AGE
default-pvc-1-a8475                 1073741824  RWO             Delete            Bound
↪default/pvc-1                      standard      19h
default-pvc-2-a8486                 1073741824  RWO             Delete            Bound
↪default/pvc-2                      standard      19h
default-pvc-3-a849e                 1073741824  RWO             Delete            Bound
↪default/pvc-3                      standard      19h
default-pvc-4-a84de                 1073741824  RWO             Delete            Bound
↪default/pvc-4                      standard      19h
trident                             2Gi         RWO             Retain            Bound
↪trident/trident                    19h
```

There are currently 4 PVs that have been created by Trident 19.04, using the `netapp.io/trident` provisioner.

```
$ kubectl describe pv default-pvc-2-a8486

Name:          default-pvc-2-a8486
Labels:        <none>
Annotations:   pv.kubernetes.io/provisioned-by: netapp.io/trident
               volume.beta.kubernetes.io/storage-class: standard
Finalizers:    [kubernetes.io/pv-protection]
StorageClass:  standard
Status:        Bound
Claim:         default/pvc-2
Reclaim Policy: Delete
Access Modes:  RWO
VolumeMode:   Filesystem
Capacity:     1073741824
Node Affinity: <none>
Message:
Source:
  Type:        NFS (an NFS mount that lasts the lifetime of a pod)
  Server:     10.xx.xx.xx
  Path:       /trid_1907_alpha_default_pvc_2_a8486
  ReadOnly:   false
```

The PV was created using the `netapp.io/trident` provisioner and is of type NFS. To support all new features provided by Trident, this PV will need to be upgraded to the CSI type.

To upgrade a legacy Trident volume to the CSI spec, you must execute the `tridentctl upgrade volume <name-of-trident-volume>` command.

```
$ ./tridentctl get volumes -n trident
+-----+-----+-----+-----+-----+
↪-----+-----+-----+
```

(continues on next page)

(continued from previous page)

```

|          NAME          | SIZE | STORAGE CLASS | PROTOCOL |          BACKEND UUID          |
|→          | STATE | MANAGED |          |          |
+-----+-----+-----+-----+-----+
|→-----+-----+-----+-----+-----+
| default-pvc-2-a8486 | 1.0 GiB | standard      | file     | c5a6f6a4-b052-423b-80d4-8fb491a14a22 |
|→-----+-----+-----+-----+-----+
| default-pvc-3-a849e | 1.0 GiB | standard      | file     | c5a6f6a4-b052-423b-80d4-8fb491a14a22 |
|→-----+-----+-----+-----+-----+
| default-pvc-1-a8475 | 1.0 GiB | standard      | file     | c5a6f6a4-b052-423b-80d4-8fb491a14a22 |
|→-----+-----+-----+-----+-----+
| default-pvc-4-a84de | 1.0 GiB | standard      | file     | c5a6f6a4-b052-423b-80d4-8fb491a14a22 |
|→-----+-----+-----+-----+-----+
|→-----+-----+-----+-----+-----+

$ ./tridentctl upgrade volume default-pvc-2-a8486 -n trident
+-----+-----+-----+-----+-----+
|→-----+-----+-----+-----+-----+
|          NAME          | SIZE | STORAGE CLASS | PROTOCOL |          BACKEND UUID          |
|→          | STATE | MANAGED |          |          |
+-----+-----+-----+-----+-----+
|→-----+-----+-----+-----+-----+
| default-pvc-2-a8486 | 1.0 GiB | standard      | file     | c5a6f6a4-b052-423b-80d4-8fb491a14a22 |
|→-----+-----+-----+-----+-----+
|→-----+-----+-----+-----+-----+

```

After upgrading the PV, performing a `kubectl describe pv` will show that the volume is a CSI volume.

```

$ kubectl describe pv default-pvc-2-a8486
Name:          default-pvc-2-a8486
Labels:        <none>
Annotations:   pv.kubernetes.io/provisioned-by: csi.trident.netapp.io
               volume.beta.kubernetes.io/storage-class: standard
Finalizers:    [kubernetes.io/pv-protection]
StorageClass:  standard
Status:        Bound
Claim:         default/pvc-2
Reclaim Policy: Delete
Access Modes:  RWO
VolumeMode:    Filesystem
Capacity:      1073741824
Node Affinity: <none>
Message:
Source:
  Type:        CSI (a Container Storage Interface (CSI) volume source)
  Driver:      csi.trident.netapp.io
  VolumeHandle: default-pvc-2-a8486
  ReadOnly:    false
  VolumeAttributes:
    backendUUID=c5a6f6a4-b052-423b-80d4-8fb491a14a22
    internalName=trid_1907_alpha_default_pvc_2_a8486
    name=default-pvc-2-a8486
    protocol=file
Events:        <none>

```

In this manner, volumes of the NFS/iSCSI type that were created by Trident can be upgraded to the CSI type, on a per-volume basis.

2.2 Deploying

This guide will take you through the process of deploying Trident and provisioning your first volume automatically. If you are a new user, this is the place to get started with using Trident.

If you are an existing user looking to upgrade, head on over to the *Upgrading Trident* section.

2.2.1 Before you begin

If you have not already familiarized yourself with the *basic concepts*, now is a great time to do that. Go ahead, we'll be here when you get back.

To deploy Trident you need:

Need Kubernetes?

If you do not already have a Kubernetes cluster, you can easily create one for demonstration purposes using our *simple Kubernetes install guide*.

- Full privileges to a *supported Kubernetes cluster*
- Access to a *supported NetApp storage system*
- *Volume mount capability* from all of the Kubernetes worker nodes
- A Linux host with `kubectl` (or `oc`, if you're using OpenShift) installed and configured to manage the Kubernetes cluster you want to use
- Enable the *Feature Gates* required by Trident
- If you are using Kubernetes with Docker EE 2.1, [follow their steps to enable CLI access](#).

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

2.2.2 1: Qualify your Kubernetes cluster

You made sure that you have everything in hand from the *previous section*, right? Right.

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

Note: With OpenShift, you will use `oc` instead of `kubectl` in all of the examples that follow, and you need to login as **system:admin** first by running `oc login -u system:admin`.

```
# Are you running a supported Kubernetes server version?
kubectl version

# Are you a Kubernetes cluster administrator?
kubectl auth can-i '*' '*' --all-namespaces

# Can you launch a pod that uses an image from Docker Hub and can reach your
# storage system over the pod network?
kubectl run -i --tty ping --image=busybox --restart=Never --rm -- \
  ping <management IP>
```

Identify your Kubernetes server version. You will be using it when you *Install Trident*.

2.2.3 2: Download & extract the installer

Note: Trident's installer is responsible for creating a Trident pod, configuring the CRD objects that are used to maintain its state and to initialize the CSI Sidecars that perform actions such as provisioning and attaching volumes to the cluster hosts.

Download the latest version of the [Trident installer bundle](#) from the *Downloads* section and extract it.

For example, if the latest version is 19.07.2:

```
wget https://github.com/NetApp/trident/releases/download/v19.07.2/trident-installer-
↪19.07.2.tar.gz
tar -xf trident-installer-19.07.2.tar.gz
cd trident-installer
```

2.2.4 3: Install Trident

Install Trident in the desired namespace by executing the `tridentctl install` command. The installation procedure slightly differs depending on the version of Kubernetes being used.

Installing Trident on Kubernetes 1.13

On Kubernetes 1.13, there are a couple of options when installing Trident:

- Install Trident in the desired namespace by executing the `tridentctl install` command with the `--csi` flag. This is the preferred method of installation and will support all features provided by Trident. The output observed will be similar to that shown *below*
- If for some reason the *feature gates* required by Trident cannot be enabled, you can install Trident without the `--csi` flag. This will configure Trident to work in its traditional format without using the CSI specification.

Installing Trident on Kubernetes 1.14 and above

Install 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.
```

(continues on next page)

(continued from previous page)

```
INFO Trident REST interface is up.                version=19.07.2
INFO Trident installation succeeded.
....
```

It will look like this when the installer is complete. Depending on the number of nodes in your Kubernetes cluster, you may observe more pods:

```
$ 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 |
+-----+-----+
| 19.07.2        | 19.07.2        |
+-----+-----+
```

If that's what you see, you're done with this step, but **Trident is not yet fully configured**. Go ahead and continue to the next step.

However, if the installer does not complete successfully or you don't see a **Running** `trident-csi-<generated id>`, then Trident had a problem and the platform was *not* installed.

To help figure out what went wrong, you could run the installer again using the `-d` argument, which will turn on debug mode and help you understand what the problem is:

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

After addressing the problem, you can clean up the installation and go back to the beginning of this step by first running:

```
./tridentctl uninstall -n trident
INFO Deleted Trident deployment.
INFO Deleted cluster role binding.
INFO Deleted cluster role.
INFO Deleted service account.
INFO Removed Trident user from security context constraint.
INFO Trident uninstallation succeeded.
```

If you continue to have trouble, visit the [troubleshooting guide](#) for more advice.

Customized Installation

Trident's installer allows 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`.

Users can also customize Trident's deployment files. Using the `--generate-custom-yaml` parameter will create 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

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

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

2.2.5 4: Create and Verify your first backend

You can now go ahead and create a backend that will be used by 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.

Visit the [backend configuration](#) of this guide for more details about how to craft the configuration file for your backend type.

Note: Many of the backends require some [basic preparation](#), so make sure that's been done before you try to use it. Also, we don't recommend an `ontap-nas-economy` backend or `ontap-nas-flexgroup` backend for this step as volumes of these types have specialized and limited capabilities relative to the volumes provisioned on other types of backends.

```
cp sample-input/<backend template>.json backend.json
# Fill out the template for your backend
vi backend.json
```

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

NAME	STORAGE DRIVER	UUID	STATE
nas-backend	ontap-nas	98e19b74-aec7-4a3d-8dcf-128e5033b214	online

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

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

After addressing the problem, simply go back to the beginning of this step and try again. If you continue to have trouble, visit the [troubleshooting guide](#) for more advice on how to determine what went wrong.

2.2.6 5: Add your first storage class

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

Basic too basic?

This is just a basic storage class to get you started. There's an art to *crafting differentiated storage classes* that you should explore further when you're looking at building them for production.

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 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.templ` 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.yaml

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

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

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

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

```
kubectl get sc basic
NAME          PROVISIONER          AGE
basic        csi.trident.netapp.io 15h

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

(continues on next page)

(continued from previous page)

```

        "aggr3",
        "aggr4"
    ]
}
}
]
}

```

2.2.7 6: Provision your first volume

Now you're ready to dynamically provision your first volume. How exciting! This is done by creating a Kubernetes [persistent volume claim \(PVC\)](#) object, and this is exactly how your users will do it too.

Create a persistent volume claim (PVC) for a volume that uses the storage class that you just created.

See `sample-input/pvc-basic.yaml` for an example. Make sure the storage class name matches the one that you created in 6.

```

kubect1 create -f sample-input/pvc-basic.yaml

kubect1 get pvc --watch
NAME          STATUS      VOLUME                                     CAPACITY   ACCESS_
↪MODES      STORAGECLASS  AGE
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

```

2.2.8 7: Mount the volume in a pod

Now that you have a volume, let's mount it. We'll 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

```

(continues on next page)

(continued from previous page)

```
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
Filesystem                                Size  Used Avail Use% Mounted on
↪Use% Mounted on
10.xx.xx.xx:/trid_1907_pvc_3acb0d1c_b1ae_11e9_8d9f_5254004dfdb7 1.0G 256K 1.0G 1% ↪ /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 could use it from another pod if you wanted to.

To delete the volume, simply delete the claim:

```
kubectl delete pvc basic
```

Check you out! You did it! Now you're dynamically provisioning Kubernetes volumes like a boss.

2.3 Common tasks

2.3.1 Managing Trident

Installing Trident

Follow the extensive *deployment* guide.

Upgrading Trident

The *Upgrade Guide* details the procedure for upgrading to the latest version of Trident.

Uninstalling Trident

The `uninstall` command in `tridentctl` will remove all of the resources associated with Trident except for the CRDs and related objects, making it easy to run the installer again to update to a more recent version.

```
./tridentctl uninstall -n <namespace>
```

To perform a complete removal of Trident, you will need to remove the finalizers for the CRDs created by Trident and delete the CRDs. Refer the *Troubleshooting Guide* for the steps to completely uninstall Trident.

Downgrading Trident

Downgrading to a previous release of Trident is **not recommended**. Refer the *Troubleshooting Guide* for downgrading Trident to a previous release.

2.3.2 Worker preparation

All of the worker nodes in the Kubernetes cluster need to be able to mount the volumes that users have provisioned for their pods.

If you are using the `ontap-nas`, `ontap-nas-economy`, `ontap-nas-flexgroup` driver for one of your back-ends, your workers will need the *NFS* tools. Otherwise they require the *iSCSI* tools.

Note: Recent versions of CoreOS have both installed by default.

Warning: You should always reboot your worker nodes after installing the NFS or iSCSI tools, or attaching volumes to containers may fail.

NFS

Install the following system packages:

RHEL / CentOS

```
sudo yum install -y nfs-utils
```

Ubuntu / Debian

```
sudo apt-get install -y nfs-common
```

iSCSI

RHEL / CentOS

1. Install the following system packages:

```
sudo yum install -y lsscsi iscsi-initiator-utils sg3_utils device-mapper-multipath
```

2. Enable multipathing:

```
sudo mpathconf --enable --with_multipathd y
```

3. Ensure that `iscsid` and `multipathd` are running:

```
sudo systemctl enable iscsid multipathd
sudo systemctl start iscsid multipathd
```

4. Start and enable `iscsi`:

```
sudo systemctl enable iscsi
sudo systemctl start iscsi
```

Ubuntu / Debian

1. Install the following system packages:

```
sudo apt-get install -y open-iscsi lsscsi sg3-utils multipath-tools scsitools
```

2. Enable multipathing:

```
sudo tee /etc/multipath.conf <<-'EOF'  
defaults {  
    user_friendly_names yes  
    find_multipaths yes  
}  
EOF  
  
sudo systemctl enable multipath-tools.service  
sudo service multipath-tools restart
```

3. Ensure that `open-iscsi` and `multipath-tools` are enabled and running:

```
sudo systemctl status multipath-tools  
sudo systemctl enable open-iscsi.service  
sudo service open-iscsi start  
sudo systemctl status open-iscsi
```

2.3.3 Backend configuration

A Trident backend defines the relationship between Trident and a storage system. It tells Trident how to communicate with that storage system and how Trident should provision volumes from it.

Trident will automatically offer up storage pools from backends that together match the requirements defined by a storage class.

To get started, choose the storage system type that you will be using as a backend:

Azure NetApp Files

Note: The Azure NetApp Files service does not support volumes less than 100 GB in size. To make it easier to deploy applications, Trident automatically creates 100 GB volumes if a smaller volume is requested.

Preparation

To configure and use an [Azure NetApp Files](#) backend, you will need:

- `subscriptionID` from an Azure subscription with Azure NetApp Files enabled
- `tenantID`, `clientID`, and `clientSecret` from an [App Registration](#) in Azure Active Directory with sufficient permissions to the Azure NetApp Files service
- Azure `location` that contains at least one [delegated subnet](#)

If you're using Azure NetApp Files for the first time or in a new location, some initial configuration is required that the [quickstart guide](#) will walk you through.

Backend configuration options

Parameter	Description	Default
version	Always 1	
storageDriver-Name	“azure-netapp-files”	
backendName	Custom name for the storage backend	Driver name + “_” + random characters
subscriptionID	The subscription ID from your Azure subscription	
tenantID	The tenant ID from an App Registration	
clientID	The client ID from an App Registration	
clientSecret	The client secret from an App Registration	
serviceLevel	One of “Standard”, “Premium” or “Ultra”	“” (random)
location	Name of the Azure location new volumes will be created in	“” (random)
virtualNetwork	Name of a virtual network with a delegated subnet	“” (random)
subnet	Name of a subnet delegated to <code>Microsoft.Netapp/volumes</code>	“” (random)
nfsMountOptions	Fine-grained control of NFS mount options	“-o nfsvers=3”
limitVolumeSize	Fail provisioning if requested volume size is above this value	“” (not enforced by default)

You can control how each volume is provisioned by default using these options in a special section of the configuration. For an example, see the configuration examples below.

Parameter	Description	Default
exportRule	The export rule(s) for new volumes	“0.0.0.0/0”
size	The default size of new volumes	“100G”

The `exportRule` value must be a comma-separated list of any combination of IPv4 addresses or IPv4 subnets in CIDR notation.

Example configurations

Example 1 - Minimal backend configuration for azure-netapp-files

This is the absolute minimum backend configuration. With this Trident will discover all of your NetApp accounts, capacity pools, and subnets delegated to ANF in every location worldwide, and place new volumes on one of them randomly.

This configuration is useful when you’re just getting started with ANF and trying things out, but in practice you’re going to want to provide additional scoping for the volumes you provision in order to make sure that they have the characteristics you want and end up on a network that’s close to the compute that’s using it. See the subsequent examples for more details.

```
{
  "version": 1,
  "storageDriverName": "azure-netapp-files",
  "subscriptionID": "9f87c765-4774-fake-ae98-a721add45451",
  "tenantID": "68e4f836-edc1-fake-bff9-b2d865ee56cf",
```

(continues on next page)

(continued from previous page)

```

"clientID": "dd043f63-bf8e-fake-8076-8de91e5713aa",
"clientSecret": "SECRET"
}

```

Example 2 - Single location and specific service level for azure-netapp-files

This backend configuration will place volumes in Azure’s “eastus” location in a “Premium” capacity pool. Trident automatically discovers all of the subnets delegated to ANF in that location and will place a new volume on one of them randomly.

```

{
  "version": 1,
  "storageDriverName": "azure-netapp-files",
  "subscriptionID": "9f87c765-4774-fake-ae98-a721add45451",
  "tenantID": "68e4f836-edc1-fake-bff9-b2d865ee56cf",
  "clientID": "dd043f63-bf8e-fake-8076-8de91e5713aa",
  "clientSecret": "SECRET",
  "location": "eastus",
  "serviceLevel": "Premium"
}

```

Example 3 - Advanced configuration for azure-netapp-files

This backend configuration further reduces the scope of volume placement to a single subnet, and also modifies some volume provisioning defaults.

```

{
  "version": 1,
  "storageDriverName": "azure-netapp-files",
  "subscriptionID": "9f87c765-4774-fake-ae98-a721add45451",
  "tenantID": "68e4f836-edc1-fake-bff9-b2d865ee56cf",
  "clientID": "dd043f63-bf8e-fake-8076-8de91e5713aa",
  "clientSecret": "SECRET",
  "location": "eastus",
  "serviceLevel": "Premium",
  "virtualNetwork": "my-virtual-network",
  "subnet": "my-subnet",
  "nfsMountOptions": "vers=3,proto=tcp,timeo=600",
  "limitVolumeSize": "500Gi",
  "defaults": {
    "exportRule": "10.0.0.0/24,10.0.1.0/24,10.0.2.100",
    "size": "200Gi"
  }
}

```

Example 4 - Virtual storage pools with azure-netapp-files

This backend configuration defines multiple *pools of storage* in a single file. This is useful when you have multiple capacity pools supporting different service levels and you want to create storage classes in Kubernetes that represent those.

This is just scratching the surface of the power of virtual storage pools and their labels.

```

{
  "version": 1,
  "storageDriverName": "azure-netapp-files",
  "subscriptionID": "9f87c765-4774-fake-ae98-a721add45451",

```

(continues on next page)

(continued from previous page)

```

"tenantID": "68e4f836-edc1-fake-bff9-b2d865ee56cf",
"clientID": "dd043f63-bf8e-fake-8076-8de91e5713aa",
"clientSecret": "SECRET",
"nfsMountOptions": "vers=3,proto=tcp,timeo=600",
"labels": {
  "cloud": "azure"
},
"location": "eastus",

"storage": [
  {
    "labels": {
      "performance": "gold"
    },
    "serviceLevel": "Ultra"
  },
  {
    "labels": {
      "performance": "silver"
    },
    "serviceLevel": "Premium"
  },
  {
    "labels": {
      "performance": "bronze"
    },
    "serviceLevel": "Standard",
  }
]
}

```

The following StorageClass definitions refer to the storage pools above. Using the `parameters.selector` field, each StorageClass calls out which pool may be used to host a volume. The volume will have the aspects defined in the chosen pool.

```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: gold
provisioner: csi.trident.netapp.io
parameters:
  selector: "performance=gold"
allowVolumeExpansion: true
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: silver
provisioner: csi.trident.netapp.io
parameters:
  selector: "performance=silver"
allowVolumeExpansion: true
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:

```

(continues on next page)

(continued from previous page)

```

name: bronze
provisioner: csi.trident.netapp.io
parameters:
  selector: "performance=bronze"
allowVolumeExpansion: true
    
```

Cloud Volumes Service for AWS

Note: The NetApp Cloud Volumes Service for AWS does not support volumes less than 100 GB in size. To make it easier to deploy applications, Trident automatically creates 100 GB volumes if a smaller volume is requested. Future releases of the Cloud Volumes Service may remove this restriction.

Preparation

To create and use a Cloud Volumes Service (CVS) for AWS backend, you will need:

- An [AWS account configured with NetApp CVS](#)
- API region, URL, and keys for your CVS account

Backend configuration options

Parameter	Description	Default
version	Always 1	
storageDriver-Name	“aws-cvs”	
backendName	Custom name for the storage backend	Driver name + “_” + part of API key
apiRegion	CVS account region	
apiURL	CVS account API URL	
apiKey	CVS account API key	
secretKey	CVS account secret key	
nfsMountOptions	Fine-grained control of NFS mount options	“-o nfsvers=3”
limitVolumeSize	Fail provisioning if requested volume size is above this value	“” (not enforced by default)
serviceLevel	The CVS service level for new volumes	“standard”

The required values `apiRegion`, `apiURL`, `apiKey`, and `secretKey` may be found in the CVS web portal in Account settings / API access.

Each backend provisions volumes in a single AWS region. To create volumes in other regions, you can define additional backends.

The `serviceLevel` values for CVS on AWS are `standard`, `premium`, and `extreme`.

You can control how each volume is provisioned by default using these options in a special section of the configuration. For an example, see the configuration examples below.

Parameter	Description	Default
exportRule	The export rule(s) for new volumes	"0.0.0.0/0"
snapshotReserve	Percentage of volume reserved for snapshots	"" (accept CVS default of 0)
size	The size of new volumes	"100G"

The `exportRule` value must be a comma-separated list of any combination of IPv4 addresses or IPv4 subnets in CIDR notation.

Example configurations

Example 1 - Minimal backend configuration for aws-cvs driver

```
{
  "version": 1,
  "storageDriverName": "aws-cvs",
  "apiRegion": "us-east-1",
  "apiURL": "https://cds-aws-bundles.netapp.com:8080/v1",
  "apiKey": "znHczZsrrtHisIsAbOguSaPIKeyAZNchRAGz1zZE",
  "secretKey": "rR0rUmWXfNioN1KhtHisISAnoTherboGuskey6pU"
}
```

Example 2 - Backend configuration for aws-cvs driver with single service level

This example shows a backend file that applies the same aspects to all Trident created storage in the AWS us-east-1 region.

```
{
  "version": 1,
  "storageDriverName": "aws-cvs",
  "backendName": "cvs-aws-us-east",
  "apiRegion": "us-east-1",
  "apiURL": "https://cds-aws-bundles.netapp.com:8080/v1",
  "apiKey": "znHczZsrrtHisIsAbOguSaPIKeyAZNchRAGz1zZE",
  "secretKey": "rR0rUmWXfNioN1KhtHisISAnoTherboGuskey6pU",
  "nfsMountOptions": "vers=3,proto=tcp,timeo=600",
  "limitVolumeSize": "50Gi",
  "serviceLevel": "premium",
  "defaults": {
    "snapshotReserve": "5",
    "exportRule": "10.0.0.0/24,10.0.1.0/24,10.0.2.100",
    "size": "200Gi"
  }
}
```

Example 3 - Backend and storage class configuration for aws-cvs driver with virtual storage pools

This example shows the backend definition file configured with *Virtual Storage Pools* along with *StorageClasses* that refer back to them.

In the sample backend definition file shown below, specific defaults are set for all storage pools, which set the `snapshotReserve` at 5% and the `exportRule` to 0.0.0.0/0. The virtual storage pools are defined in the storage section. In this example, each individual storage pool sets its own `serviceLevel`, and some pools overwrite the default values set above.

```
{
  "version": 1,
  "storageDriverName": "aws-cvs",
  "apiRegion": "us-east-1",
  "apiURL": "https://cds-aws-bundles.netapp.com:8080/v1",
  "apiKey": "EnterYourAPIKeyHere*****",
  "secretKey": "EnterYourSecretKeyHere*****",
  "nfsMountOptions": "vers=3,proto=tcp,timeo=600",

  "defaults": {
    "snapshotReserve": "5",
    "exportRule": "0.0.0.0/0"
  },

  "labels": {
    "cloud": "aws"
  },
  "region": "us-east-1",

  "storage": [
    {
      "labels": {
        "performance": "extreme",
        "protection": "extra"
      },
      "serviceLevel": "extreme",
      "defaults": {
        "snapshotReserve": "10",
        "exportRule": "10.0.0.0/24"
      }
    },
    {
      "labels": {
        "performance": "extreme",
        "protection": "standard"
      },
      "serviceLevel": "extreme"
    },
    {
      "labels": {
        "performance": "premium",
        "protection": "extra"
      },
      "serviceLevel": "premium",
      "defaults": {
        "snapshotReserve": "10"
      }
    },
    {
      "labels": {
        "performance": "premium",
        "protection": "standard"
      },
      "serviceLevel": "premium"
    }
  ],
}
```

(continues on next page)

(continued from previous page)

```

    {
      "labels": {
        "performance": "standard"
      },
      "serviceLevel": "standard"
    }
  ]
}

```

The following StorageClass definitions refer to the above Virtual Storage Pools. Using the `parameters.selector` field, each StorageClass calls out which virtual pool(s) may be used to host a volume. The volume will have the aspects defined in the chosen virtual pool.

The first StorageClass (`cvms-extreme-extra-protection`) will map to the first Virtual Storage Pool. This is the only pool offering extreme performance with a snapshot reserve of 10%. The last StorageClass (`cvms-extra-protection`) calls out any storage pool which provides a snapshot reserve of 10%. Trident will decide which Virtual Storage Pool is selected and will ensure the snapshot reserve requirement is met.

```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: cvms-extreme-extra-protection
provisioner: netapp.io/trident
parameters:
  selector: "performance=extreme; protection=extra"
allowVolumeExpansion: true
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: cvms-extreme-standard-protection
provisioner: netapp.io/trident
parameters:
  selector: "performance=premium; protection=standard"
allowVolumeExpansion: true
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: cvms-premium-extra-protection
provisioner: netapp.io/trident
parameters:
  selector: "performance=premium; protection=extra"
allowVolumeExpansion: true
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: cvms-premium
provisioner: netapp.io/trident
parameters:
  selector: "performance=premium; protection=standard"
allowVolumeExpansion: true
---
apiVersion: storage.k8s.io/v1
kind: StorageClass

```

(continues on next page)

```
metadata:
  name: cvs-standard
provisioner: netapp.io/trident
parameters:
  selector: "performance=standard"
allowVolumeExpansion: true
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: cvs-extra-protection
provisioner: netapp.io/trident
parameters:
  selector: "protection=extra"
allowVolumeExpansion: true
```

Element (HCI/SolidFire)

To create and use an Element backend, you will need:

- A *supported Element storage system*
- Complete *HCI/SolidFire backend preparation*
- Credentials to a HCI/SolidFire cluster admin or tenant user that can manage volumes

Preparation

All of your Kubernetes worker nodes must have the appropriate iSCSI tools installed. See the *worker configuration guide* for more details.

Note: Trident will use CHAP when functioning as an enhanced CSI Provisioner.

If you're using CHAP (which is the default for CSI), no further preparation is required. It is recommended to explicitly set the `UseCHAP` option to use CHAP with non-CSI Trident. Otherwise, see the *access groups guide* below.

Warning: Volume Access Groups are only supported by the conventional, non-CSI framework for Trident. When configured to work in CSI mode, Trident uses CHAP.

If neither `AccessGroups` or `UseCHAP` are set then one of the following rules applies: * If the default `trident` access group is detected then access groups are used. * If no access group is detected and Kubernetes version ≥ 1.7 then CHAP is used.

Backend configuration options

Parameter	Description	Default
version	Always 1	
storageDriver-Name	Always “solidfire-san”	
backendName	Custom name for the storage backend	“solidfire_” + storage (iSCSI) IP address
Endpoint	MVIP for the SolidFire cluster with tenant credentials	
SVIP	Storage (iSCSI) IP address and port	
TenantName	Tenant name to use (created if not found)	
InitiatorIFace	Restrict iSCSI traffic to a specific host interface	“default”
UseCHAP	Use CHAP to authenticate iSCSI	
AccessGroups	List of Access Group IDs to use	Finds the ID of an access group named “trident”
Types	QoS specifications (see below)	
limitVolume-Size	Fail provisioning if requested volume size is above this value	“” (not enforced by default)

Example configuration

Example 1 - Backend configuration for solidfire-san driver with three volume types

This example shows a backend file using CHAP authentication and modeling three volume types with specific QoS guarantees. Most likely you would then define storage classes to consume each of these using the IOPS storage class parameter.

```
{
  "version": 1,
  "storageDriverName": "solidfire-san",
  "Endpoint": "https://<user>:<password>@<mvip>/json-rpc/8.0",
  "SVIP": "<svip>:3260",
  "TenantName": "<tenant>",
  "UseCHAP": true,
  "Types": [
    {
      "Type": "Bronze",
      "Qos": {
        "minIOPS": 1000,
        "maxIOPS": 2000,
        "burstIOPS": 4000
      }
    },
    {
      "Type": "Silver",
      "Qos": {
        "minIOPS": 4000,
        "maxIOPS": 6000,
        "burstIOPS": 8000
      }
    },
    {
      "Type": "Gold",
      "Qos": {
        "minIOPS": 6000,
        "maxIOPS": 8000,
        "burstIOPS": 10000
      }
    }
  ]
}
```

Example 2 - Backend and Storage Class configuration for solidfire-san driver with Virtual Storage Pools

This example shows the backend definition file configured with *Virtual Storage Pools* along with StorageClasses that refer back to them.

In the sample backend definition file shown below, specific defaults are set for all storage pools, which set the type at Silver. The Virtual Storage Pools are defined in the storage section. In this example, some of the storage pool sets their own type, and some pools overwrite the default values set above.

```
{
  "version": 1,
  "storageDriverName": "solidfire-san",
```

(continues on next page)

(continued from previous page)

```

"Endpoint": "https://<user>:<password>@<mvip>/json-rpc/8.0",
"SVIP": "<svip>:3260",
"TenantName": "<tenant>",
"UseCHAP": true,
"Types": [{"Type": "Bronze", "Qos": {"minIOPS": 1000, "maxIOPS": 2000, "burstIOPS
↵": 4000}},
          {"Type": "Silver", "Qos": {"minIOPS": 4000, "maxIOPS": 6000, "burstIOPS
↵": 8000}},
          {"Type": "Gold", "Qos": {"minIOPS": 6000, "maxIOPS": 8000, "burstIOPS":
↵10000}}],

"defaults": {
  "type": "Silver"
},

"labels": {"store": "solidfire"},
"region": "us-east-1",

"storage": [
  {
    "labels": {"performance": "gold", "cost": "4"},
    "zone": "us-east-1a",
    "type": "Gold"
  },
  {
    "labels": {"performance": "silver", "cost": "3"},
    "zone": "us-east-1b",
    "type": "Silver"
  },
  {
    "labels": {"performance": "bronze", "cost": "2"},
    "zone": "us-east-1c",
    "type": "Bronze"
  },
  {
    "labels": {"performance": "silver", "cost": "1"},
    "zone": "us-east-1d"
  }
]
}

```

The following StorageClass definitions refer to the above Virtual Storage Pools. Using the `parameters.selector` field, each StorageClass calls out which virtual pool(s) may be used to host a volume. The volume will have the aspects defined in the chosen virtual pool.

The first StorageClass (`solidfire-gold-four`) will map to the first Virtual Storage Pool. This is the only pool offering gold performance with a Volume Type QoS of Gold. The last StorageClass (`solidfire-silver`) calls out any storage pool which offers a silver performance. Trident will decide which Virtual Storage Pool is selected and will ensure the storage requirement is met.

```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: solidfire-gold-four
provisioner: csi.trident.netapp.io
parameters:

```

(continues on next page)

(continued from previous page)

```

    selector: "performance=gold; cost=4"
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: solidfire-silver-three
provisioner: csi.trident.netapp.io
parameters:
  selector: "performance=silver; cost=3"
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: solidfire-bronze-two
provisioner: csi.trident.netapp.io
parameters:
  selector: "performance=bronze; cost=2"
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: solidfire-silver-one
provisioner: csi.trident.netapp.io
parameters:
  selector: "performance=silver; cost=1"
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: solidfire-silver
provisioner: csi.trident.netapp.io
parameters:
  selector: "performance=silver"

```

Using access groups

Note: Ignore this section if you are using CHAP, which we recommend to simplify management and avoid the scaling limit described below. In addition, if using Trident in CSI mode, you can safely ignore this section. Trident uses CHAP when installed as an enhanced CSI provisioner.

Trident can use volume access groups to control access to the volumes that it provisions. If CHAP is disabled it expects to find an access group called `trident` unless one or more access group IDs are specified in the configuration.

While Trident associates new volumes with the configured access group(s), it does not create or otherwise manage access groups themselves. The access group(s) must exist before the storage backend is added to Trident, and they need to contain the iSCSI IQNs from every node in the Kubernetes cluster that could potentially mount the volumes provisioned by that backend. In most installations that's every worker node in the cluster.

For Kubernetes clusters with more than 64 nodes, you will need to use multiple access groups. Each access group may contain up to 64 IQNs, and each volume can belong to 4 access groups. With the maximum 4 access groups configured, any node in a cluster up to 256 nodes in size will be able to access any volume.

If you're modifying the configuration from one that is using the default `trident` access group to one that uses others as well, include the ID for the `trident` access group in the list.

ONTAP (AFF/FAS/Select/Cloud)

To create and use an ONTAP backend, you will need:

- A *supported ONTAP storage system*
- Choose the *ONTAP storage driver* that you want to use
- Complete *ONTAP backend preparation* for the driver of your choice
- Credentials to an ONTAP SVM with *appropriate access*

Choosing a driver

Driver	Protocol
ontap-nas	NFS
ontap-nas-economy	NFS
ontap-nas-flexgroup	NFS
ontap-san	iSCSI

The `ontap-nas` and `ontap-san` drivers create an ONTAP FlexVol for each volume. ONTAP supports up to 1000 FlexVols per cluster node with a cluster maximum of 12,000 FlexVols. If your persistent volume requirements fit within that limitation, those drivers are the preferred solution due to the granular data management capabilities they afford.

If you need more persistent volumes than may be accommodated by the FlexVol limits, choose the `ontap-nas-economy` driver, which creates volumes as ONTAP Qtrees within a pool of automatically managed FlexVols. Qtrees offer far greater scaling, up to 100,000 per cluster node and 2,400,000 per cluster, at the expense of granular data management features.

Choose the `ontap-nas-flexgroup` driver to increase parallelism to a single volume that can grow into the petabyte range with billions of files. Some ideal use cases for FlexGroups include AI/ML/DL, big data and analytics, software builds, streaming, file repositories, etc. Trident uses all aggregates assigned to an SVM when provisioning a FlexGroup Volume. FlexGroup support in Trident also has the following considerations:

- Requires ONTAP version 9.2 or greater.
- As of this writing, FlexGroups only support NFSv3 (required to set `mountOptions: ["nfsvers=3"]` in the Kubernetes storage class).
- Recommended to enable the 64-bit NFSv3 identifiers for the SVM.
- The minimum recommended FlexGroup size is 100GB.
- Cloning is not supported for FlexGroup Volumes.

For information regarding FlexGroups and workloads that are appropriate for FlexGroups see the [NetApp FlexGroup Volume - Best Practices and Implementation Guide](#).

Remember that you can also run more than one driver, and create storage classes that point to one or the other. For example, you could configure a *Gold* class that uses the `ontap-nas` driver and a *Bronze* class that uses the `ontap-nas-economy` one.

Preparation

For all ONTAP backends, Trident requires at least one [aggregate assigned to the SVM](#).

ontap-nas, ontap-nas-economy, ontap-nas-flexgroups

All of your Kubernetes worker nodes must have the appropriate NFS tools installed. See the *worker configuration guide* for more details.

Trident uses NFS [export policies](#) to control access to the volumes that it provisions. It uses the `default` export policy unless a different export policy name is specified in the configuration.

While Trident associates new volumes (or qtrees) with the configured export policy, it does not create or otherwise manage export policies themselves. The export policy must exist before the storage backend is added to Trident, and it needs to be configured to allow access to every worker node in the Kubernetes cluster.

If the export policy is locked down to specific hosts, it will need to be updated when new nodes are added to the cluster, and that access should be removed when nodes are removed as well.

ontap-san

All of your Kubernetes worker nodes must have the appropriate iSCSI tools installed. See the *worker configuration guide* for more details.

Trident uses [igroups](#) to control access to the volumes (LUNs) that it provisions. It expects to find an igroup called `trident` unless a different igroup name is specified in the configuration.

While Trident associates new LUNs with the configured igroup, it does not create or otherwise manage igroups themselves. The igroup must exist before the storage backend is added to Trident, and it needs to contain the iSCSI IQNs from every worker node in the Kubernetes cluster.

The igroup needs to be updated when new nodes are added to the cluster, and they should be removed when nodes are removed as well.

Backend configuration options

Parameter	Description	Default
<code>version</code>	Always 1	
<code>storageDriver-Name</code>	“ontap-nas”, “ontap-nas-economy”, “ontap-nas-flexgroup”, or “ontap-san”	
<code>backendName</code>	Custom name for the storage backend	Driver name + “_” + <code>dataLIF</code>
<code>managementLIF</code>	IP address of a cluster or SVM management LIF	“10.0.0.1”
<code>dataLIF</code>	IP address of protocol LIF	Derived by the SVM unless specified
<code>svm</code>	Storage virtual machine to use	Derived if an SVM <code>managementLIF</code> is specified
<code>igroupName</code>	Name of the igroup for SAN volumes to use	“trident”
<code>username</code>	Username to connect to the cluster/SVM	
<code>password</code>	Password to connect to the cluster/SVM	
<code>storagePrefix</code>	Prefix used when provisioning new volumes in the SVM	“trident”
<code>limitAggregateUsage</code>	Fail provisioning if usage is above this percentage	“” (not enforced by default)
<code>limitVolume-Size</code>	Fail provisioning if requested volume size is above this value	“” (not enforced by default)
<code>nfsMountOptions</code>	Comma-separated list of NFS mount options (except <code>ontap-san</code>)	“”

A fully-qualified domain name (FQDN) can be specified for the managementLIF option. For the ontap-nas* drivers only, a FQDN may also be specified for the dataLIF option, in which case the FQDN will be used for the NFS mount operations. For the ontap-san driver, the default is to use all data LIF IPs from the SVM and to use iSCSI multipath. Specifying an IP address for the dataLIF for the ontap-san driver forces the driver to disable multipath and use only the specified address. For the ontap-nas-economy driver, the limitVolumeSize option will also restrict the maximum size of the volumes it manages for qtrees.

The nfsMountOptions parameter applies to all ONTAP drivers except ontap-san. The mount options for Kubernetes persistent volumes are normally specified in storage classes, but if no mount options are specified in a storage class, Trident will fall back to using the mount options specified in the storage backend's config file. If no mount options are specified in either the storage class or the config file, then Trident will not set any mount options on an associated persistent volume.

You can control how each volume is provisioned by default using these options in a special section of the configuration. For an example, see the configuration examples below.

Parameter	Description	Default
spaceAllocation	ontap-san* only: space-allocation for LUNs	“true”
spaceReserve	Space reservation mode; “none” (thin) or “volume” (thick)	“none”
snapshotPolicy	Snapshot policy to use	“none”
snapshotReserve	Percentage of volume reserved for snapshots	“0” if snapshotPolicy is “none”, else “”
splitOnClone	Split a clone from its parent upon creation	“false”
encryption	Enable NetApp volume encryption	“false”
unixPermissions	ontap-nas* only: mode for new volumes	“777”
snapshotDir	ontap-nas* only: access to the .snapshot directory	“false”
exportPolicy	ontap-nas* only: export policy to use	“default”
securityStyle	ontap-nas* only: security style for new volumes	“unix”

Example configuration

NFS Example for ontap-nas driver

```
{
  "version": 1,
  "storageDriverName": "ontap-nas",
  "managementLIF": "10.0.0.1",
  "dataLIF": "10.0.0.2",
  "svm": "svm_nfs",
  "username": "admin",
  "password": "secret",
  "nfsMountOptions": "nfsvers=4",
}
```

NFS Example for ontap-nas-flexgroup driver

```
{
  "version": 1,
  "storageDriverName": "ontap-nas-flexgroup",
  "managementLIF": "10.0.0.1",
  "dataLIF": "10.0.0.2",
}
```

(continues on next page)

(continued from previous page)

```
"svm": "svm_nfs",
"username": "vsadmin",
"password": "secret",
}
```

NFS Example for ontap-nas-economy driver

```
{
  "version": 1,
  "storageDriverName": "ontap-nas-economy",
  "managementLIF": "10.0.0.1",
  "dataLIF": "10.0.0.2",
  "svm": "svm_nfs",
  "username": "vsadmin",
  "password": "secret"
}
```

iSCSI Example for ontap-san driver

```
{
  "version": 1,
  "storageDriverName": "ontap-san",
  "managementLIF": "10.0.0.1",
  "dataLIF": "10.0.0.3",
  "svm": "svm_iscsi",
  "igroupName": "trident",
  "username": "vsadmin",
  "password": "secret"
}
```

User permissions

Trident expects to be run as either an ONTAP or SVM administrator, typically using the `admin` cluster user or a `vsadmin` SVM user, or a user with a different name that has the same role.

Note: If you use the “`limitAggregateUsage`” option, cluster admin permissions are required.

While it is possible to create a more restrictive role within ONTAP that a Trident driver can use, we don’t recommend it. Most new releases of Trident will call additional APIs that would have to be accounted for, making upgrades difficult and error-prone.

SANtricity (E-Series)

To create and use an E-Series backend, you will need:

- A *supported E-Series storage system*
- Complete *E-Series backend preparation*
- Credentials to the E-Series storage system

Preparation

All of your Kubernetes worker nodes must have the appropriate iSCSI tools installed. See the *worker configuration guide* for more details.

Trident uses host groups to control access to the volumes (LUNs) that it provisions. It expects to find a host group called `trident` unless a different host group name is specified in the configuration.

While Trident associates new volumes with the configured host group, it does not create or otherwise manage host groups themselves. The host group must exist before the storage backend is added to Trident, and it needs to contain a host definition with an iSCSI IQN for every worker node in the Kubernetes cluster.

Backend configuration options

Parameter	Description	Default
<code>version</code>	Always 1	
<code>storageDriverName</code>	Always “ <code>eseries-iscsi</code> ”	
<code>backendName</code>	Custom name for the storage backend	“ <code>eseries_</code> ” + <code>hostDataIP</code>
<code>webProxyHostname</code>	Hostname or IP address of the web services proxy	
<code>webProxyPort</code>	Port number of the web services proxy	80 for HTTP, 443 for HTTPS
<code>webProxyUseHTTP</code>	Use HTTP instead of HTTPS to communicate to the proxy	false
<code>webProxyVerifyTLS</code>	Verify certificate chain and hostname	false
<code>username</code>	Username for the web services proxy	
<code>password</code>	Password for the web services proxy	
<code>controllerA</code>	IP address for controller A	
<code>controllerB</code>	IP address for controller B	
<code>passwordArray</code>	Password for the storage array, if set	“”
<code>hostDataIP</code>	Host iSCSI IP address	
<code>poolNameSearchPattern</code>	Regular expression for matching available storage pools	“ <code>.+</code> ” (all)
<code>hostType</code>	E-Series Host types created by the driver	“ <code>linux_dm_mp</code> ”
<code>accessGroupName</code>	E-Series Host Group used by the driver	“ <code>trident</code> ”
<code>limitVolumeSize</code>	Fail provisioning if requested volume size is above this value	“” (not enforced by default)

Example configuration

Example 1 - Minimal backend configuration for `eseries-iscsi` driver

```
{
  "version": 1,
  "storageDriverName": "eseries-iscsi",
  "webProxyHostname": "localhost",
  "webProxyPort": "8443",
  "username": "rw",
  "password": "rw",
  "controllerA": "10.0.0.5",
  "controllerB": "10.0.0.6",
  "passwordArray": "",

```

(continues on next page)

(continued from previous page)

```

"hostDataIP": "10.0.0.101"
}

```

Example 2 - Backend and Storage Class configuration for eseries-iscsi driver with Virtual Storage Pools

This example shows the backend definition file configured with *Virtual Storage Pools* along with StorageClasses that refer back to them.

In the sample backend definition file shown below, the Virtual Storage Pools are defined in the `storage` section.

```

{
  "version": 1,
  "storageDriverName": "eseries-iscsi",
  "controllerA": "10.0.0.1",
  "controllerB": "10.0.0.2",
  "hostDataIP": "10.0.1.1",
  "username": "user",
  "password": "password",
  "passwordArray": "password",
  "webProxyHostname": "10.0.2.1",

  "labels": {"store": "eseries"},
  "region": "us-east",

  "storage": [
    {
      "labels": {"performance": "gold", "cost": "4"},
      "zone": "us-east-1a"
    },
    {
      "labels": {"performance": "silver", "cost": "3"},
      "zone": "us-east-1b"
    },
    {
      "labels": {"performance": "bronze", "cost": "2"},
      "zone": "us-east-1c"
    },
    {
      "labels": {"performance": "bronze", "cost": "1"},
      "zone": "us-east-1d"
    }
  ]
}

```

The following StorageClass definitions refer to the above virtual storage pools. Using the `parameters.selector` field, each StorageClass calls out which virtual pool(s) may be used to host a volume. The volume will have the aspects defined in the chosen virtual pool.

The first StorageClass (`eseries-gold-four`) will map to the first virtual storage pool. This is the only pool offering gold performance in zone `us-east-1a`. The last StorageClass (`eseries-bronze`) calls out any storage pool which offers a bronze performance. Trident will decide which virtual storage pool is selected and will ensure the storage requirement is met.

```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: eseries-gold-four

```

(continues on next page)

```
provisioner: netapp.io/trident
parameters:
  selector: "performance=gold; cost=4"
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: eseries-silver-three
provisioner: netapp.io/trident
parameters:
  selector: "performance=silver; cost=3"
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: eseries-bronze-two
provisioner: netapp.io/trident
parameters:
  selector: "performance=bronze; cost=2"
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: eseries-bronze-one
provisioner: netapp.io/trident
parameters:
  selector: "performance=bronze; cost=1"
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: eseries-bronze
provisioner: netapp.io/trident
parameters:
  selector: "performance=bronze"
```

2.3.4 Managing backends

Creating a backend configuration

We have an entire *backend configuration* guide to help you with this.

Creating a backend

Once you have a *backend configuration* file, run:

```
tridentctl create backend -f <backend-file>
```

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

```
tridentctl logs
```

Once you identify and correct the problem with the configuration file you can simply run the create command again.

Deleting a backend

Note: If Trident has provisioned volumes and snapshots from this backend that still exist, deleting the backend will prevent new volumes from being provisioned by it. The backend will continue to exist in a “Deleting” state and Trident will continue to manage those volumes and snapshots until they are deleted.

To delete a backend from Trident, run:

```
# Retrieve the backend name
tridentctl get backend

tridentctl delete backend <backend-name>
```

Viewing the existing backends

To view the backends that Trident knows about, run:

```
# Summary
tridentctl get backend

# Full details
tridentctl get backend -o json
```

Identifying the storage classes that will use a backend

This is an example of the kind of questions you can answer with the JSON that `tridentctl` outputs for Trident backend objects. This uses the `jq` utility, which you may need to install first.

```
tridentctl get backend -o json | jq '[.items[] | {backend: .name, storageClasses: [.
↪storage[].storageClasses]|unique}]'
```

Updating a backend

Once you have a new *backend configuration* file, run:

```
tridentctl update backend <backend-name> -f <backend-file>
```

If backend update fails, something was wrong with the backend configuration or you attempted an invalid update. You can view the logs to determine the cause by running:

```
tridentctl logs
```

Once you identify and correct the problem with the configuration file you can simply run the update command again.

2.3.5 Managing storage classes

Designing a storage class

The *StorageClass concept guide* will help you understand what they do and how you configure them.

Creating a storage class

Once you have a storage class file, run:

```
kubectl create -f <storage-class-file>
```

Deleting a storage class

To delete a storage class from Kubernetes, run:

```
kubectl delete storageclass <storage-class>
```

Any persistent volumes that were created through this storage class will remain untouched, and Trident will continue to manage them.

Viewing the existing storage classes

```
# Kubernetes storage classes
kubectl get storageclass

# Kubernetes storage class detail
kubectl get storageclass <storage-class> -o json

# Trident's synchronized storage classes
tridentctl get storageclass

# Trident's synchronized storage class detail
tridentctl get storageclass <storage-class> -o json
```

Setting a default storage class

Kubernetes v1.6 added the ability to set a default storage class. This is the storage class that will be used to provision a PV if a user does not specify one in a PVC.

You can define a default storage class by setting the annotation `storageclass.kubernetes.io/is-default-class` to `true` in the storage class definition. According to the specification, any other value or absence of the annotation is interpreted as `false`.

It is possible to configure an existing storage class to be the default storage class by using the following command:

```
kubectl patch storageclass <storage-class-name> -p '{"metadata": {"annotations":{"storageclass.kubernetes.io/is-default-class":"true"}}}'
```

Similarly, you can remove the default storage class annotation by using the following command:

```
kubectl patch storageclass <storage-class-name> -p '{"metadata": {"annotations":{"storageclass.kubernetes.io/is-default-class":"false"}}}'
```

There are also examples in the Trident installer bundle that include this annotation.

Note: You should only have one default storage class in your cluster at any given time. Kubernetes does not technically prevent you from having more than one, but it will behave as if there is no default storage class at all.

Identifying the Trident backends that a storage class will use

This is an example of the kind of questions you can answer with the JSON that `tridentctl` outputs for Trident backend objects. This uses the `jq` utility, which you may need to install first.

```
tridentctl get storageclass -o json | jq '[.items[] | {storageClass: .Config.name, ↵
↵backends: [.storage]|unique}]'
```

2.3.6 Managing volumes

On-Demand Volume Snapshots

Creating Kubernetes Volume Snapshots (using the beta specification) is available beginning with Trident 20.01. You can take a look at the documentation [here](#). You will need to make sure your Kubernetes release is 1.17 or above and this feature requires Trident 20.01 or later.

Expanding an NFS volume

Starting with v18.10, Trident supports volume resize for NFS PVs. More specifically, PVs provisioned on `ontap-nas`, `ontap-nas-economy`, `ontap-nas-flexgroup`, `aws-cvs` and `azure-netapp-files` backends can be expanded. Volume resize was introduced in Kubernetes v1.8 as an alpha feature and was promoted to beta in v1.11, which means this feature is enabled by default starting with Kubernetes v1.11.

To resize an NFS PV, the admin first needs to configure the storage class to allow volume expansion by setting the `allowVolumeExpansion` field to `true`:

```
$ cat storageclass-ontapnas.yaml
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: ontapnas
provisioner: csi.trident.netapp.io
parameters:
  backendType: ontap-nas
allowVolumeExpansion: true
```

If you have already created a storage class without this option, you can simply edit the existing storage class via `kubectrl edit storageclass` to allow volume expansion.

Next, we create a PVC using this storage class:

```
$ cat pvc-ontapnas.yaml
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: ontapnas20mb
spec:
  accessModes:
  - ReadWriteOnce
  resources:
    requests:
      storage: 20Mi
  storageClassName: ontapnas
```

Trident should create a 20MiB NFS PV for this PVC:

```
$ kubectl get pvc
NAME                                STATUS    VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS   AGE
↪ACCESS MODES    STORAGECLASS   AGE
ontapnas20mb    Bound     pvc-08f3d561-b199-11e9-8d9f-5254004dfdb7    20Mi       RWO            ontapnas       9s
↪
ontapnas       9s

$ kubectl get pv pvc-08f3d561-b199-11e9-8d9f-5254004dfdb7
NAME                                CAPACITY   ACCESS MODES   RECLAIM POLICY
↪STATUS    CLAIM                                STORAGECLASS   REASON    AGE
pvc-08f3d561-b199-11e9-8d9f-5254004dfdb7    20Mi       RWO            Delete
↪Bound     default/ontapnas20mb    ontapnas                                     2m42s
```

To resize the newly created 20MiB PV to 1GiB, we edit the PVC and set `spec.resources.requests.storage` to 1GB:

```
$ kubectl edit pvc ontapnas20mb
# Please edit the object below. Lines beginning with a '#' will be ignored,
# and an empty file will abort the edit. If an error occurs while saving this file,
↪will be
# reopened with the relevant failures.
#
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  annotations:
    pv.kubernetes.io/bind-completed: "yes"
    pv.kubernetes.io/bound-by-controller: "yes"
    volume.beta.kubernetes.io/storage-provisioner: csi.trident.netapp.io
  creationTimestamp: 2018-08-21T18:26:44Z
  finalizers:
  - kubernetes.io/pvc-protection
  name: ontapnas20mb
  namespace: default
  resourceVersion: "1958015"
  selfLink: /api/v1/namespaces/default/persistentvolumeclaims/ontapnas20mb
  uid: c1bd7fa5-a56f-11e8-b8d7-fa163e59eaab
spec:
  accessModes:
  - ReadWriteOnce
  resources:
    requests:
      storage: 1Gi
...
```

We can validate the resize has worked correctly by checking the size of the PVC, PV, and the Trident volume:

```
$ kubectl get pvc ontapnas20mb
NAME                                STATUS    VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS   AGE
↪MODES    STORAGECLASS   AGE
ontapnas20mb    Bound     pvc-08f3d561-b199-11e9-8d9f-5254004dfdb7    1Gi       RWO            ontapnas       4m44s
↪
ontapnas       4m44s

$ kubectl get pv pvc-08f3d561-b199-11e9-8d9f-5254004dfdb7
NAME                                CAPACITY   ACCESS MODES   RECLAIM POLICY
↪STATUS    CLAIM                                STORAGECLASS   REASON    AGE
pvc-08f3d561-b199-11e9-8d9f-5254004dfdb7    1Gi       RWO            Delete
↪Bound     default/ontapnas20mb    ontapnas                                     5m35s
```

(continues on next page)

(continued from previous page)

```
$ tridentctl get volume pvc-08f3d561-b199-11e9-8d9f-5254004dfdb7 -n trident
```

NAME	SIZE	STORAGE CLASS	PROTOCOL
pvc-08f3d561-b199-11e9-8d9f-5254004dfdb7	1.0 GiB	ontapnas	file
c5a6f6a4-b052-423b-80d4-8fb491a14a22	online	true	

Importing a volume

Trident version 19.04 and above allows importing an existing storage volume into Kubernetes with the `ontap-nas`, `ontap-nas-flexgroup`, `solidfire-san`, `azure-netapp-files` and `aws-cvs` drivers.

There are several use cases for importing a volume into Trident:

- Containerizing an application and reusing its existing data set
- Using a clone of a data set for an ephemeral application
- Rebuilding a failed Kubernetes cluster
- Migrating application data during disaster recovery

The `tridentctl` client is used to import an existing storage volume. Trident imports the volume by persisting volume metadata and creating the PVC and PV.

```
$ tridentctl import volume <backendName> <volumeName> -f <path-to-pvc-file>
```

To import an existing storage volume, specify the name of the Trident backend containing the volume, as well as the name that uniquely identifies the volume on the storage (i.e. ONTAP FlexVol, Element Volume, CVS Volume path'). The storage volume must allow read/write access and be accessible by the specified Trident backend.

The `-f` string argument is required and specifies the path to the YAML or JSON PVC file. The PVC file is used by the volume import process to create the PVC. At a minimum, the PVC file must include the name, namespace, accessModes, and storageClassName fields as shown in the following example.

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: my_claim
  namespace: my_namespace
spec:
  accessModes:
    - ReadWriteOnce
  storageClassName: my_storage_class
```

When Trident receives the import volume request the existing volume size is determined and set in the PVC. Once the volume is imported by the storage driver the PV is created with a ClaimRef to the PVC. The reclaim policy is initially set to `retain` in the PV. Once Kubernetes successfully binds the PVC and PV the reclaim policy is updated to match the reclaim policy of the Storage Class. If the reclaim policy of the Storage Class is `delete` then the storage volume will be deleted when the PV is deleted.

When a volume is imported with the `--no-manage` argument, Trident will not perform any additional operations on the PVC or PV for the lifecycle of the objects. Since Trident ignores PV and PVC events for `--no-manage` objects the storage volume is not deleted when the PV is deleted. Other operations such as volume clone and volume resize are also ignored. This option is provided for those that want to use Kubernetes for containerized workloads but otherwise want to manage the lifecycle of the storage volume outside of Kubernetes.

An annotation is added to the PVC and PV that serves a dual purpose of indicating that the volume was imported and if the PVC and PV are managed. This annotation should not be modified or removed.

Trident 19.07 handles the attachment of PVs and mounts the volume as part of importing it. For imports using earlier versions of Trident, there will not be any operations in the data path and the volume import will not verify if the volume can be mounted. If a mistake is made with volume import (e.g. the StorageClass is incorrect), you can recover by changing the reclaim policy on the PV to “Retain”, deleting the PVC and PV, and retrying the volume import command.

Note: The Element driver supports duplicate volume names. If there are duplicate volume names Trident’s volume import process will return an error. As a workaround, clone the volume and provide a unique volume name. Then import the cloned volume.

For example, to import a volume named `managed_volume` on a backend named `ontap_nas` use the following command:

```
$ tridentctl import volume ontap_nas managed_volume -f <path-to-pvc-file>
```

NAME	SIZE	STORAGE CLASS	PROTOCOL	
BACKEND UUID	STATE	MANAGED		
pvc-bf5ad463-afbb-11e9-8d9f-5254004dfdb7	1.0 GiB	standard	file	
c5a6f6a4-b052-423b-80d4-8fb491a14a22	online	true		

To import a volume named `unmanaged_volume` (on the `ontap_nas` backend) which Trident will not manage, use the following command:

```
$ tridentctl import volume nas_blog unmanaged_volume -f <path-to-pvc-file> --no-manage
```

NAME	SIZE	STORAGE CLASS	PROTOCOL	
BACKEND UUID	STATE	MANAGED		
pvc-df07d542-afbc-11e9-8d9f-5254004dfdb7	1.0 GiB	standard	file	
c5a6f6a4-b052-423b-80d4-8fb491a14a22	online	false		

When using the `--no-manage` flag, Trident renames the volume, but it does not validate if the volume was mounted. The import operation will fail if the volume was not mounted manually.

To import an `aws-cvs` volume on the backend called `awscvs_YEppr` with the volume path of `adroit-jolly-swift` use the following command:

```
$ tridentctl import volume awscvs_YEppr adroit-jolly-swift -f <path-to-pvc-file> -n
↳trident

+-----+-----+-----+-----+-----+
↳-----+-----+-----+-----+-----+
|          NAME          |  SIZE  | STORAGE CLASS | PROTOCOL |  |
↳          BACKEND UUID      | STATE  | MANAGED      |          |  |
+-----+-----+-----+-----+-----+
↳-----+-----+-----+-----+-----+
| pvc-a46ccab7-44aa-4433-94b1-e47fc8c0fa55 | 100 GiB | aws-storage   | file     |  |
↳e1a6e65b-299e-4568-ad05-4f0a105c888f | online | true          |          |  |
+-----+-----+-----+-----+-----+
↳-----+-----+-----+-----+-----+
```

Note: The AWS volume path is the portion of the volume's export path after the `./`. For example, if the export path is `10.0.0.1:/adroit-jolly-swift` then the volume path is `adroit-jolly-swift`.

To import an `azure-netapp-files` volume on the backend called `azurenetaappfiles_40517` with the volume path `importvol1`, you will use the following command:

```
$ tridentctl import volume azurenetaappfiles_40517 importvol1 -f <path-to-pvc-file> -n
↳trident

+-----+-----+-----+-----+-----+
↳-----+-----+-----+-----+-----+
|          NAME          |  SIZE  | STORAGE CLASS | PROTOCOL |  |
↳          BACKEND UUID      | STATE  | MANAGED      |          |  |
+-----+-----+-----+-----+-----+
↳-----+-----+-----+-----+-----+
| pvc-0ee95d60-fd5c-448d-b505-b72901b3a4ab | 100 GiB | anf-storage   | file     |  |
↳1c01274f-d94b-44a3-98a3-04c953c9a51e | online | true          |          |  |
+-----+-----+-----+-----+-----+
↳-----+-----+-----+-----+-----+
```

Note: The volume path for the ANF volume is present in the mount path after the `./`. For example, if the mount path is `10.0.0.2:/importvol1`, the volume path is `importvol1`.

Behavior of Drivers for Volume Import

- The `ontap-nas` and `ontap-nas-flexgroup` drivers do not allow duplicate volume names.
- To import a volume backed by the NetApp Cloud Volumes Service in AWS, identify the volume by its volume path instead of its name. An example is provided in the previous section.
- An ONTAP volume must be of type `rw` to be imported by Trident. If a volume is of type `dp` it is a SnapMirror destination volume; you must break the mirror relationship before importing the volume into Trident.

2.4 Concepts

2.4.1 Kubernetes and Trident objects

Both Kubernetes and Trident are designed to be interacted with through REST APIs by reading and writing resource objects.

Object overview

There are several different resource objects in play here, some that are managed through Kubernetes and others that are managed through Trident, that dictate the relationship between Kubernetes and Trident, Trident and storage, and Kubernetes and storage.

Perhaps the easiest way to understand these objects, what they are for and how they interact, is to follow a single request for storage from a Kubernetes user:

1. A user creates a *PersistentVolumeClaim* requesting a new *PersistentVolume* of a particular size from a *Kubernetes StorageClass* that was previously configured by the administrator.
2. The *Kubernetes StorageClass* identifies Trident as its provisioner and includes parameters that tell Trident how to provision a volume for the requested class.
3. Trident looks at its own *Trident StorageClass* with the same name that identifies the matching *Backends* and *StoragePools* that it can use to provision volumes for the class.
4. Trident provisions storage on a matching backend and creates two objects: a *PersistentVolume* in Kubernetes that tells Kubernetes how to find, mount and treat the volume, and a *Volume* in Trident that retains the relationship between the *PersistentVolume* and the actual storage.
5. Kubernetes binds the *PersistentVolumeClaim* to the new *PersistentVolume*. Pods that include the *PersistentVolumeClaim* will mount that *PersistentVolume* on any host that it runs on. and a *Volume*.

Throughout the rest of this guide, we will describe the different Trident objects and details about how Trident crafts the Kubernetes *PersistentVolume* object for storage that it provisions.

For further reading about the Kubernetes objects, we highly recommend that you read the [Persistent Volumes](#) section of the Kubernetes documentation.

Kubernetes *PersistentVolumeClaim* objects

A Kubernetes *PersistentVolumeClaim* object is a request for storage made by a Kubernetes cluster user.

In addition to the [standard specification](#), Trident allows users to specify the following volume-specific annotations if they want to override the defaults that you set in the backend configuration:

Annotation	Volume Option	Supported Drivers
trident.netapp.io/fileSystem	fileSystem	ontap-san, solidfire-san, eseries-iscsi
trident.netapp.io/cloneFromPVC	cloneSourceVolume	ontap-nas, ontap-san, solidfire-san, aws-cvs
trident.netapp.io/splitOnClone	splitOnClone	ontap-nas, ontap-san
trident.netapp.io/protocol	protocol	any
trident.netapp.io/exportPolicy	exportPolicy	ontap-nas, ontap-nas-economy, ontap-nas-flexgroup
trident.netapp.io/snapshotPolicy	snapshotPolicy	ontap-nas, ontap-nas-economy, ontap-nas-flexgroup, ontap-san
trident.netapp.io/snapshotReserve	snapshotReserve	ontap-nas, ontap-nas-flexgroup, ontap-san, aws-cvs
trident.netapp.io/snapshotDirectory	snapshotDirectory	ontap-nas, ontap-nas-economy, ontap-nas-flexgroup
trident.netapp.io/unixPermissions	unixPermissions	ontap-nas, ontap-nas-economy, ontap-nas-flexgroup
trident.netapp.io/blockSize	blockSize	solidfire-san

Note: The `trident.netapp.io/cloneFromPVC` and `trident.netapp.io/splitOnClone` annotations are deprecated. Users must create Volume Snapshots and use them to clone PVCs

If the created PV has the `Delete` reclaim policy, Trident will delete both the PV and the backing volume when the PV becomes released (i.e., when the user deletes the PVC). Should the delete action fail, Trident will mark the PV as such and periodically retry the operation until it succeeds or the PV is manually deleted. If the PV uses the `Retain` policy, Trident ignores it and assumes the administrator will clean it up from Kubernetes and the backend, allowing the volume to be backed up or inspected before its removal. Note that deleting the PV will not cause Trident to delete the backing volume; it must be removed manually via the REST API (i.e., `tridentctl`).

Trident supports the creation of Volume Snapshots using the CSI specification: you can create a Volume Snapshot and use it as a Data Source to clone existing PVCs. This way, point-in-time copies of PVs can be exposed to Kubernetes in the form of snapshots. The snapshots can then be used to create new PVs. Take a look at [On-Demand Volume Snapshots](#) to see how this would work.

See [Trident Volume objects](#) for a full description of the parameters and settings associated with Trident volumes.

Kubernetes PersistentVolume objects

A [Kubernetes PersistentVolume](#) object represents a piece of storage that's been made available to the Kubernetes cluster. They have a lifecycle that's independent of the pod that uses it.

Note: Trident creates `PersistentVolume` objects and registers them with the Kubernetes cluster automatically based on the volumes that it provisions. You are not expected to manage them yourself.

When a user creates a PVC that refers to a Trident-based `StorageClass`, Trident will provision a new volume using the corresponding storage class and register a new PV for that volume. In configuring the provisioned volume and corresponding PV, Trident follows the following rules:

- Trident generates a PV name for Kubernetes and an internal name that it uses to provision the storage. In both cases it is assuring that the names are unique in their scope.
- The size of the volume matches the requested size in the PVC as closely as possible, though it may be rounded up to the nearest allocatable quantity, depending on the platform.

Kubernetes StorageClass objects

Kubernetes `StorageClass` objects are specified by name in `PersistentVolumeClaims` to provision storage with a set of properties. The storage class itself identifies the provisioner that will be used and defines that set of properties in terms the provisioner understands.

It is one of two basic objects that need to be created and managed by you, the administrator. The other is the *Trident Backend object*.

A Kubernetes `StorageClass` object that uses Trident looks like this:

```
apiVersion: storage.k8s.io/v1beta1
kind: StorageClass
metadata:
  name: <Name>
provisioner: csi.trident.netapp.io
mountOptions: <Mount Options>
parameters:
  <Trident Parameters>
```

These parameters are Trident-specific and tell Trident how to provision volumes for the class.

The storage class parameters are:

Attribute	Type	Re-quired	Description
attributes	map[string]string	no	See the attributes section below
storagePools	map[string]StringList	no	Map of backend names to lists of storage pools within
additionalStorage-Pools	map[string]StringList	no	Map of backend names to lists of storage pools within
excludeStoragePools	map[string]StringList	no	Map of backend names to lists of storage pools within

Storage attributes and their possible values can be classified into two groups:

1. Storage pool selection attributes: These parameters determine which Trident-managed storage pools should be utilized to provision volumes of a given type.

Attribute	Type	Values	Offer	Request	Supported by
media ¹	string	hdd, hybrid, ssd	Pool contains media of this type; hybrid means both	Media type specified	ontap-nas, ontap-nas-economy, ontap-nas-flexgroup, ontap-san, solidfire-san
provisioningType	string	thin, thick	Pool supports this provisioning method	Provisioning method specified	thick: all but solidfire-san & aws-cvs, thin: all but eseries-iscsi
backendType	string	ontap-nas, ontap-nas-economy, ontap-nas-flexgroup, ontap-san, solidfire-san, eseries-iscsi, aws-cvs	Pool belongs to this type of backend	Backend specified	All drivers
snapshots	bool	true, false	Pool supports volumes with snapshots	Volume with snapshots enabled	ontap-nas, ontap-san, solidfire-san, aws-cvs
clones	bool	true, false	Pool supports cloning volumes	Volume with clones enabled	ontap-nas, ontap-san, solidfire-san, aws-cvs
encryption	bool	true, false	Pool supports encrypted volumes	Volume with encryption enabled	ontap-nas, ontap-nas-economy, ontap-nas-flexgroups, ontap-san
IOPS	int	positive integer	Pool is capable of guaranteeing IOPS in this range	Volume guaranteed these IOPS	solidfire-san

¹: Not supported by ONTAP Select systems

In most cases, the values requested will directly influence provisioning; for instance, requesting thick provisioning will result in a thickly provisioned volume. However, an Element storage pool will use its offered IOPS minimum and maximum to set QoS values, rather than the requested value. In this case, the requested value is used only to select the storage pool.

Ideally you will be able to use `attributes` alone to model the qualities of the storage you need to satisfy the needs of a particular class. Trident will automatically discover and select storage pools that match *all* of the `attributes` that you specify.

If you find yourself unable to use `attributes` to automatically select the right pools for a class, you can use the `storagePools` and `additionalStoragePools` parameters to further refine the pools or even to select a specific set of pools manually.

The `storagePools` parameter is used to further restrict the set of pools that match any specified `attributes`. In other words, Trident will use the intersection of pools identified by the `attributes` and `storagePools` parameters for provisioning. You can use either parameter alone or both together.

The `additionalStoragePools` parameter is used to extend the set of pools that Trident will use for provisioning, regardless of any pools selected by the `attributes` and `storagePools` parameters.

The `excludeStoragePools` parameter is used to filter the set of pools that Trident will use for provisioning and will remove any pools that match.

In the `storagePools` and `additionalStoragePools` parameters, each entry takes the form `<backend>:<storagePoolList>`, where `<storagePoolList>` is a comma-separated list of storage pools for the specified backend. For example, a value for `additionalStoragePools` might look like `ontapnas_192.168.1.100:aggr1,aggr2;solidfire_192.168.1.101:bronze`. These lists will accept regex values for both the backend and list values. You can use `tridentctl get backend` to get the list of backends and their pools.

2. **Kubernetes attributes:** These attributes have no impact on the selection of storage pools/backends by Trident during dynamic provisioning. Instead, these attributes simply supply parameters supported by Kubernetes Persistent Volumes.

Attribute	Type	Values	Description	Relevant Drivers	Kubernetes Version
<code>fsType</code>	string	<code>ext4</code> , <code>ext3</code> , <code>xf</code> s, etc.	The file system type for block volumes	<code>solidfire-san</code> , <code>ontap-san</code> , <code>eseries-iscsi</code>	All

The Trident installer bundle provides several example storage class definitions for use with Trident in `sample-input/storage-class-*.yaml`. Deleting a Kubernetes storage class will cause the corresponding Trident storage class to be deleted as well.

Kubernetes CustomResourceDefinition objects

Kubernetes Custom Resources are endpoints in the Kubernetes API that are defined by the administrator and are used to group similar objects. Kubernetes supports the creation of custom resources for storing a collection of objects. These resource definitions can be obtained by doing a `kubectl get crds`.

CRDs and their associated object metadata are stored by Kubernetes in its metadata store. This eliminates the need for a separate store for Trident.

Beginning with the 19.07 release, Trident uses a number of CustomResourceDefinitions (CRDs) to preserve the identity of Trident objects such as Trident backends, Trident Storage classes and Trident volumes. These objects are managed by Trident. In addition, the CSI volume snapshot framework introduces some CRDs that are required to define volume snapshots.

CRDs are a Kubernetes construct. Objects of the resources defined above are created by Trident. As a simple example, when a backend is created using `tridentctl`, a corresponding `tridentbackends` CRD object is created for consumption by Kubernetes.

Trident StorageClass objects

Note: With Kubernetes, these objects are created automatically when a Kubernetes StorageClass that uses Trident as a provisioner is registered.

Trident creates matching storage classes for Kubernetes StorageClass objects that specify `csi.trident.netapp.io/netapp.io/trident` in their provisioner field. The storage class's name will match that of the Kubernetes StorageClass object it represents.

Storage classes comprise a set of requirements for volumes. Trident matches these requirements with the attributes present in each storage pool; if they match, that storage pool is a valid target for provisioning volumes using that storage class.

One can create storage class configurations to directly define storage classes via the [REST API](#). However, for Kubernetes deployments, we expect them to be created as a side-effect of registering new [Kubernetes StorageClass objects](#).

Trident Backend objects

Backends represent the storage providers on top of which Trident provisions volumes; a single Trident instance can manage any number of backends.

This is one of the two object types that you will need to create and manage yourself. The other is the [Kubernetes StorageClass object](#) below.

For more information about how to construct these objects, visit the [backend configuration](#) guide.

Trident StoragePool objects

Storage pools represent the distinct locations available for provisioning on each backend. For ONTAP, these correspond to aggregates in SVMs. For HCI/SolidFire, these correspond to admin-specified QoS bands. For Cloud Volumes Service, these correspond to cloud provider regions. Each storage pool has a set of distinct storage attributes, which define its performance characteristics and data protection characteristics.

Unlike the other objects here, storage pool candidates are always discovered and managed automatically. [View your backends](#) to see the storage pools associated with them.

Trident Volume objects

Note: With Kubernetes, these objects are managed automatically and should not be manipulated by hand. You can view them to see what Trident provisioned, however.

Note: When deleting a Persistent Volume with associated snapshots, the corresponding Trident volume is updated to a "Deleting state". For the Trident volume to be deleted, the snapshots of the volume must be removed.

Volumes are the basic unit of provisioning, comprising backend endpoints such as NFS shares and iSCSI LUNs. In Kubernetes, these correspond directly to PersistentVolumes. Each volume must be created with a storage class, which determines where that volume can be provisioned, along with a size.

A volume configuration defines the properties that a provisioned volume should have.

Attribute	Type	Re-quired	Description
version	string	no	Version of the Trident API (“1”)
name	string	yes	Name of volume to create
storageClass	string	yes	Storage class to use when provisioning the volume
size	string	yes	Size of the volume to provision in bytes
protocol	string	no	Protocol type to use; “file” or “block”
internalName	string	no	Name of the object on the storage system; generated by Trident
snapshotPolicy	string	no	ontap-*: Snapshot policy to use
snapshotReserve	string	no	ontap-*: Percentage of volume reserved for snapshots
exportPolicy	string	no	ontap-nas*: Export policy to use
snapshotDirectory	bool	no	ontap-nas*: Whether the snapshot directory is visible
unixPermissions	string	no	ontap-nas*: Initial UNIX permissions
blockSize	string	no	solidfire-*: Block/sector size
fileSystem	string	no	File system type
cloneSourceVolume	string	no	ontap-{naslsan} & solidfire-* & aws-cvs*: Name of the volume to clone from
splitOnClone	string	no	ontap-{naslsan}: Split the clone from its parent

As mentioned, Trident generates `internalName` when creating the volume. This consists of two steps. First, it prepends the storage prefix – either the default, `trident`, or the prefix in the backend configurationd – to the volume name, resulting in a name of the form `<prefix>-<volume-name>`. It then proceeds to sanitize the name, replacing characters not permitted in the backend. For ONTAP backends, it replaces hyphens with underscores (thus, the internal name becomes `<prefix>_<volume-name>`), and for Element backends, it replaces underscores with hyphens. For E-Series, which imposes a 30-character limit on all object names, Trident generates a random string for the internal name of each volume. For CVS (AWS), which imposes a 16-to-36-character limit on the unique volume creation token, Trident generates a random string for the internal name of each volume.

One can use volume configurations to directly provision volumes via the *REST API*, but in Kubernetes deployments we expect most users to use the standard [Kubernetes PersistentVolumeClaim](#) method. Trident will create this volume object automatically as part of the provisioning process in that case.

2.4.2 How does provisioning work?

Provisioning in Trident has two primary phases. The first of these associates a storage class with the set of suitable backend storage pools and occurs as a necessary preparation before provisioning. The second encompasses the volume creation itself and requires choosing a storage pool from those associated with the pending volume’s storage class. This section explains both of these phases and the considerations involved in them, so that users can better understand how Trident handles their storage.

Associating backend storage pools with a storage class relies on both the storage class’s requested attributes and its `storagePools`, `additionalStoragePools`, and `excludeStoragePools` lists. When a user creates a storage class, Trident compares the attributes and pools offered by each of its backends to those requested by the storage class. If a storage pool’s attributes and name match all of the requested attributes and pool names, Trident adds that storage pool to the set of suitable storage pools for that storage class. In addition, Trident adds all storage pools listed in the `additionalStoragePools` list to that set, even if their attributes do not fulfill all or any of the storage class’s requested attributes. Use the `excludeStoragePools` list to override and remove storage pools from use for a storage class. Trident performs a similar process every time a user adds a new backend, checking whether its storage pools satisfy those of the existing storage classes and removing any that have been marked as excluded.

Trident then uses the associations between storage classes and storage pools to determine where to provision volumes. When a user creates a volume, Trident first gets the set of storage pools for that volume’s storage class, and, if the user specifies a protocol for the volume, it removes those storage pools that cannot provide the requested protocol (a

HCI/SolidFire backend cannot provide a file-based volume while an ONTAP NAS backend cannot provide a block-based volume, for instance). Trident randomizes the order of this resulting set, to facilitate an even distribution of volumes, and then iterates through it, attempting to provision the volume on each storage pool in turn. If it succeeds on one, it returns successfully, logging any failures encountered in the process. Trident returns a failure if and only if it fails to provision on **all** the storage pools available for the requested storage class and protocol.

2.4.3 Virtual Storage Pools

Virtual storage pools provide a layer of abstraction between Trident's storage backends and Kubernetes' StorageClasses. They allow an administrator to define aspects like location, performance, and protection for each backend in a common, backend-agnostic way without making a StorageClass specify which physical backend, backend pool, or backend type to use to meet desired criteria.

Today, Virtual Storage Pools can be created for Element and SANtricity backends, as well as the Cloud Volumes Service for AWS and Azure NetApp Files service.




Fig. 1: Virtual Storage Pools

The storage administrator defines the virtual pools and their aspects in a backend's JSON or YAML definition file. Any aspect specified outside the virtual pools list is global to the backend and will apply to all the virtual pools, while each virtual pool may specify one or more aspects individually (overriding any backend-global aspects).

Most aspects are specified in backend-specific terms. Crucially, the aspect values are not exposed outside the backend's driver and are not available for matching in StorageClasses. Instead, the administrator defines one or more labels for each virtual pool. Each label is a key:value pair, and labels may be common across unique backends. Like aspects, labels may be specified per-pool or global to the backend. Unlike aspects, which have predefined names and values, the administrator has full discretion to define label keys and values as needed.

A StorageClass identifies which virtual pool(s) to use by referencing the labels within a selector parameter. Virtual pool selectors support six operators:

Operator	Example	Description
=	performance=premium	A pool's label value must match
!=	performance!=extreme	A pool's label value must not match
in	location in (east, west)	A pool's label value must be in the set of values
notin	performance notin (silver, bronze)	A pool's label value must not be in the set of values
<key>	protection	A pool's label key must exist with any value
!<key>	!protection	A pool's label key must not exist

A selector may consist of multiple operators, delimited by semicolons; all operators must succeed to match a virtual pool.

2.5 Known issues

This page contains a list of known issues that may be observed when using Trident.

- Although we provide a deployment for Trident, it should never be scaled beyond a single replica. Similarly, only one instance of Trident should be run per Kubernetes cluster. Trident cannot communicate with other instances and cannot discover other volumes that they have created, which will lead to unexpected and incorrect behavior if more than one instance runs within a cluster.
- Volumes and storage classes created in the REST API will not have corresponding objects (PVCs or StorageClasses) created in Kubernetes; however, storage classes created via `tridentctl` or the REST API will be usable by PVCs created in Kubernetes.
- If Trident-based `StorageClass` objects are deleted from Kubernetes while Trident is offline, Trident will not remove the corresponding storage classes from its database when it comes back online. Any such storage classes must be deleted manually using `tridentctl` or the REST API.
- If a user deletes a PV provisioned by Trident before deleting the corresponding PVC, Trident will not automatically delete the backing volume. In this case, the user must remove the volume manually via `tridentctl` or the REST API.
- When using a backend across multiple Trident instances, it is recommended that each backend configuration file specify a different `storagePrefix` value for ONTAP backends or use a different `TenantName` for Element backends. Trident cannot detect volumes that other instances of Trident have created, and attempting to create an existing volume on either ONTAP or Element backends succeeds as Trident treats volume creation as an idempotent operation. Thus, if the `storagePrefix` or `TenantName` does not differ, there is a very slim chance to have name collisions for volumes created on the same backend.
- ONTAP cannot concurrently provision more than one FlexGroup at a time unless the set of aggregates are unique to each provisioning request.

2.6 Troubleshooting

- If there was a failure during install, run `tridentctl logs -l all -n trident` and look for problems in the logs for the `trident-main` and CSI containers (when using the CSI provisioner). Alternatively, you can use `kubectl logs` to retrieve the logs for the `trident-*****-****` pod.
- If the Trident pod fails to come up properly (e.g., when Trident pod is stuck in the `ContainerCreating` phase with fewer than 2 ready containers), running `kubectl -n trident describe deployment trident` and `kubectl -n trident describe pod trident-*****-****` can provide additional insights. Obtaining kubelet logs (e.g., via `journalctl -xeu kubelet`) can also be helpful.

- If there's not enough information in the Trident logs, you can try enabling the debug mode for Trident by passing the `-d` flag to the install parameter: `./tridentctl install -d -n trident`.
- If there are problems with mounting a PV to a container, ensure that `rpcbind` is installed and running. Use the required package manager for the host OS and check if `rpcbind` is running. You can check the status of the `rpcbind` service by running `systemctl status rpcbind` or its equivalent.
- If you encounter permission issues when installing Trident with Docker as the container runtime, attempt the installation of Trident with the `--in-cluster=false` flag. This will not use an installer pod and avoid permission troubles seen due to the `trident-installer` user.
- The *uninstall parameter* can help with cleaning up after a failed run. By default the script does not remove the CRDs that have been created by Trident, making it safe to uninstall and install again even in a running deployment.
- If you are looking to downgrade to an earlier version of Trident, first execute the `tridentctl uninstall` command to remove Trident. Download the desired [Trident version](#) and install using the `tridentctl install` command. Only consider a downgrade if there are no new PVs created and if no changes have been made to already existing PVs/backends/ storage classes. Since Trident now uses CRDs for maintaining state, all storage entities created (backends, storage classes, PVs) have *associated CRD objects* instead of data written into the PV that was used by the earlier installed version of Trident. **Newly created PVs will not be usable when moving back to an earlier version. Changes made to objects such as backends, PVs and storage classes (created/updated/deleted) will not be visible to Trident when downgraded.** The PV that was used by the earlier version of Trident installed will still be visible to Trident. Going back to an earlier version will not disrupt access for PVs that were already created using the older release, unless they have been upgraded.
- To completely remove Trident, execute the `tridentctl obliviante crd` command. This will remove all CRD objects and undefine the CRDs. Trident will no longer manage any PVs it had already provisioned. Remember that Trident will need to be reconfigured from scratch after this.
- After a successful install, if a PVC is stuck in the `Pending` phase, running `kubectl describe pvc` can provide additional information on why Trident failed to provision a PV for this PVC.
- If you require further assistance, please create a support bundle via `tridentctl logs -a -n trident` and send it to [NetApp Support](#).

3.1 Introduction

Containers have quickly become one of the most popular methods of packaging and deploying applications. The ecosystem surrounding the creation, deployment, and management of containerized applications has exploded, resulting in myriad solutions available to customers who simply want to deploy their applications with as little friction as possible.

Application teams love containers due to their ability to decouple the application from the underlying operating system. The ability to create a container on their own laptop, then deploy to a teammate's laptop, their on-premises data center, hyperscalars, and anywhere else means that they can focus their efforts on the application and its code, not on how the underlying operating system and infrastructure are configured.

At the same time, operations teams are only just now seeing the dramatic rise in popularity of containers. Containers are often approached first by developers for personal productivity purposes, which means the infrastructure teams are insulated from or unaware of their use. However, this is changing. Operations teams are now expected to deploy, maintain, and support infrastructures which host containerized applications. In addition, the rise of DevOps is pushing operations teams to understand not just the application, but the deployment method and platform at a much greater depth than ever before.

Fortunately there are robust platforms for hosting containerized applications. Arguably the most popular of those platforms is [Kubernetes](#), an open source [Cloud Native Computing Foundation \(CNCF\)](#) project, which orchestrates the deployment of containers, including connecting them to network and storage resources as needed.

Deploying an application using containers doesn't change its fundamental resource requirements. Reading, writing, accessing, and storing data doesn't change just because a container technology is now a part of the stack in addition to virtual and/or physical machines.

To facilitate the consumption of storage resources by containerized applications, [NetApp](#) created and released an open source project known as [Trident](#). Trident is a storage orchestrator which integrates with Docker and Kubernetes, as well as platforms built on those technologies, such as [NetApp Kubernetes Service](#), [Red Hat OpenShift](#), [Rancher](#), and [IBM Cloud Private](#). The goal of Trident is to make the provisioning, connection, and consumption of storage as transparent and frictionless for applications as possible; while operating within the constraints put forth by the storage administrator.

To achieve this goal, Trident automates the storage management tasks needed to consume storage for the storage administrator, the Kubernetes and Docker administrators, and the application consumers. Trident fills a critical role for storage administrators, who may be feeling pressure from application teams to provide storage resources in ways which have not previously been expected. Modern applications, and just as importantly modern development practices, have changed the storage consumption model, where resources are created, consumed, and destroyed quickly. According to [DataDog](#), containers have a median lifespan of just six days. This is dramatically different than storage resources for traditional applications, which commonly exist for years. Those which are deployed using container orchestrators have an even shorter lifespan of just a half day. Trident is the tool which storage administrators can rely on to safely, within the bounds given to it, provision the storage resources applications need, when they need them, and where they need them.

3.1.1 Target Audience

This document outlines the design and architecture considerations that should be evaluated when deploying containerized applications with persistence requirements within your organization. Additionally, you can find best practices for configuring Kubernetes and OpenShift with Trident.

It is assumed that you, the reader, have a basic understanding of containers, Kubernetes, and storage prior to reading this document. We will, however, explore and explain some of the concepts which are important to integrating Trident, and through it NetApp's storage platforms and services, with Kubernetes. Unless noted, Kubernetes and OpenShift can be used interchangeably in this document.

As with all best practices, these are suggestions based on the experience and knowledge of the NetApp team. Each should be considered according to your environment and targeted applications.

3.2 Concepts and Definitions

Kubernetes introduces several new concepts that storage, application, and platform administrators should understand. It is essential to evaluate the capability of each within the context of their use case. These concepts are discussed below and will be encountered throughout the document.

3.2.1 Kubernetes storage concepts

The Kubernetes storage paradigm includes several entities which are important to each stage of requesting, consuming, and managing storage for containerized applications. At a high level, Kubernetes uses the three types of objects for storage described below:

Persistent Volume Claim

Persistent Volume Claims (PVCs) are used by applications to request access to storage resources. At a minimum, this includes two key characteristics, Size and Access mode. Storage Class is optional:

- Size – The capacity desired by an application component
- Access mode – The rules for accessing the storage volume. The PVC can use one of three access modes:
 - Read Write Once (RWO) – Only one node is allowed to have read write access to the storage volume at a time.
 - Read Only Many (ROX) – Many nodes may access the storage volume in read-only mode
 - Read Write Many (RWX) – Many nodes may simultaneously read and write to the storage volume

- – Storage Class (optional) - Identifies which Storage Class to request for this PVC. See below for Storage Class information.

More information about PVCs can be found in the [Kubernetes](#) or [OpenShift](#) documentation.

Persistent Volume

Persistent Volumes (PVs) are Kubernetes objects that describe how to connect to a storage device. Kubernetes supports many different types of storage. However, this document covers only NFS and iSCSI devices since NetApp platforms and services support those protocols. At a minimum, the PV must contain these parameters:

- Capacity - The size of the volume represented by the object, e.g. “5Gi”
- Access mode - This is the same as for PVCs but can be dependent on the protocol used as shown below:
 - Read Write Once (RWO) - Supported by all PVs
 - Read Only Many (ROX) - Supported primarily by file and file-like protocols, e.g. NFS and CephFS. However, some block protocols are supported, such as iSCSI
 - Read Write Many (RWX) - Supported by file and file-like protocols only, such as NFS
- Protocol - Type of protocol (e.g. “iSCSI” or “NFS”) to use and additional information needed to access the storage. For example, an NFS PV will need the NFS server and a mount path.
- Reclaim policy - It describes the Kubernetes action when the PV is released. Three reclaim policy options are available:
 - Retain - Mark the volume as waiting for administrator action. The volume cannot be reissued to another PVC.
 - Recycle - Kubernetes will connect the volume to a temporary pod and issue a `rm -rf` command to clear the data after the volume is released. For our interests, this is only supported by NFS volumes.
 - Delete - Kubernetes will delete the PV when it is released. However, Kubernetes does not delete the storage which was referenced by the PV. A storage administrator will need to delete the volume on the backend.

More information about PVs can be found in the [Kubernetes](#) or [OpenShift](#) documentation.

Storage Class

Kubernetes uses the Storage Class object to describe storage with specific characteristics. An administrator may define several Storage Classes that each define different storage properties. The Storage Classes are used by the *PVC* to provision storage. A Storage Class may have a provisioner associated with it that will trigger Kubernetes to wait for the volume to be provisioned by the specified provider. In the case of Trident, the provisioner identifier used is `netapp.io/trident`. The value of the provisioner for Trident using its enhanced CSI provisioner is `csi.trident.netapp.io`.

A Storage Class object in Kubernetes has only two required fields:

- Name - It uniquely identifies the Storage Class from the PVC
- Provisioner - Identifies which plug-in to use to provision the volume. A full list of provisioners can be found in [the documentation](#)

The provisioner used may require additional attributes which will be specific to the provisioner used. Additionally, the Storage Class may have a reclaim policy and mount options specified which will be applied to all volumes created for that storage class.

More information about storage classes can be found in the [Kubernetes](#) or [OpenShift](#) documentation.

3.2.2 Kubernetes compute concepts

In addition to the basic Kubernetes storage concepts described above, it's important to understand the Kubernetes compute objects involved in the consumption of storage resources. This section discusses the Kubernetes compute objects.

Kubernetes, as a container orchestrator, dynamically assigns containerized workloads to cluster members according to any resource requirements that may have been expressed. For more information about what containers are, see the [Docker documentation](#).

Pods

A [pod](#) represents one or more containers which are related to each other and is the smallest (atomic) unit for Kubernetes. Containers which are members of the same pod are co-scheduled to the same node in the cluster. They typically share the same namespace, network, and storage resources. Though not every container in the pod may access the storage or be publicly accessible via the network.

Services

A Kubernetes [service](#) adds an abstraction layer over pods. It typically acts as an internal load balancer for replicated pods. The service enables the scaling of pods while maintaining a consistent service IP address. There are several types of services, a service can be reachable only within the cluster by declaring it of type ClusterIP, or may be exposed to the outside world by declaring it of type NodePort, LoadBalancer, or ExternalName.

Deployments

A Kubernetes [deployment](#) is one or more pods which are related to each other and often represent a “service” to a larger application being deployed. The application administrator uses deployments to declare the state of their application component and requests that Kubernetes ensure that the state is implemented at all times. This can include several options:

- Pods which should be deployed, including versions, storage, network, and other resource requests
- Number of replicas of each pod instance

The application administrator then uses the deployment as the interface for managing the application. For example, by increasing or decreasing the number of replicas desired the application can be horizontally scaled in or out. Updating the deployment with a new version of the application pod(s) will trigger Kubernetes to remove existing instances and redeploy using the new version. Conversely, rolling back to a previous version of the deployment will cause Kubernetes to revert the pods to the previously specified version and configuration.

StatefulSets

Deployments specify how to scale pods. When a webserver (which is managed as a Kubernetes deployment) is scaled up, Kubernetes will add more instances of that pod to reach the desired count. However, when a PVC is added to a deployment, the PVC is shared by all pod replicas. What if each pod needs unique persistent storage?

[StatefulSets](#) are a special type of deployment where separate persistent storage is requested along with each replica of the pod(s) so that each pod receives its own volume of storage. To accomplish this, the StatefulSet definition includes a template PVC which is used to request additional storage resources as the application is scaled out. This is generally used for stateful applications such as databases.

In order to accomplish the above, StatefulSets provide unique pod names and network identifiers that are persistent across pod restarts. They also allow ordered operations, including startup, scale-up, upgrades, and deletion.

As the number of pod replicas increase, the number of PVCs does as well. However, scaling down the application will not result in the PVCs being destroyed, as Kubernetes relies on the application administrator to clean up the PVCs in order to prevent inadvertent data loss.

3.2.3 Connecting containers to storage

When an application submits a PVC requesting storage, the Kubernetes engine will assign a PV which matches the requirement. If no PV exists which can meet the request expressed in the PVC, then it will wait until a provisioner creates a PV which matches the request before making the assignment. If no storage class was assigned, then the Kubernetes administrator would be expected to request a storage resource and introduce a PV.




Fig. 1: Kubernetes dynamic storage provisioning process

The storage is not connected to a Kubernetes node within a cluster until the pod has been scheduled. At that time, `kubelet`, the agent running on each node that is responsible for managing container instances, mounts the storage to the host according to the information in the PV. When the container(s) in the pod are instantiated on the host, `kubelet` mounts the storage devices into the container.

3.2.4 Destroying and creating pods

It's important to understand that Kubernetes creates and destroys pods (workloads), it does not "move" them like live VM migration performed by hypervisors. When Kubernetes scales down or needs to re-deploy a workload on a different host, the pod and the container(s) on the original host are stopped, destroyed, and the resources unmounted. The standard mount and instantiate process is then followed wherever in the cluster the same workload is re-deployed as a different pod with a different name, IP address, etc. (Note: Stateful sets are an exception and can re-deploy a pod with the same name). When an application being deployed relies on persistent storage, that storage must be accessible from any Kubernetes node deploying the workload within the cluster. Without a shared storage system available for persistence, the data would be abandoned, and usually deleted, on the source system when the workload is re-deployed elsewhere in the cluster.

To maintain a persistent pod that will always be deployed on the same node with the same name and characteristics, a `StatefulSet` must be used as described above.

3.2.5 Container Storage Interface

The Cloud Native Computing Foundation (CNCF) is actively working on a standardized Container Storage Interface (CSI). NetApp is active in the CSI Special Interest Group (SIG). CSI is meant to be a standard mechanism used by various container orchestrators to expose storage systems to containers. Trident v19.07 fully conforms with CSI 1.1

specifications and supports all volume operations. Trident's enhanced CSI support is production ready and currently supported on Kubernetes versions 1.13 and above.

3.3 NetApp Products and Integrations with Kubernetes

The NetApp portfolio of storage products integrates with many different aspects of a Kubernetes cluster, providing advanced data management capabilities which enhance the functionality, capability, performance, and availability of the Kubernetes deployment.

3.3.1 Trident

NetApp Trident is a dynamic storage provisioner for the containers ecosystem. It provides the ability to create storage volumes for containerized applications managed by Docker and Kubernetes. Trident is a fully supported, open source project hosted on [GitHub](#). Trident works with the portfolio of NetApp storage platforms to deliver storage on-demand to applications according to policies defined by the administrator. When used with Kubernetes, Trident is deployed using native paradigms and provides persistent storage to all namespaces in the cluster. For more information about Trident, visit [ThePub](#).

3.3.2 ONTAP

ONTAP is NetApp's multiprotocol, unified storage operating system that provides advanced data management capabilities for any application. ONTAP systems may have all-flash, hybrid, or all-HDD configurations and offer many different deployment models, including engineered hardware (FAS and AFF), white-box (ONTAP Select), and cloud-only (Cloud Volumes ONTAP). Trident supports all the above mentioned ONTAP deployment models.

Cloud Volumes ONTAP

[Cloud Volumes ONTAP](#) is a software-only storage appliance that runs the ONTAP data management software in the cloud. You can use Cloud Volumes ONTAP for production workloads, disaster recovery, DevOps, file shares, and database management. It extends enterprise storage to the cloud by offering storage efficiencies, high availability, data replication, data tiering and application consistency.

3.3.3 Element OS

Element OS enables the storage administrator to consolidate workloads by guaranteeing performance and enabling a simplified and streamlined storage footprint. Coupled with an API to enable automation of all aspects of storage management, Element OS enables storage administrators to do more with less effort.

Trident supports all Element OS clusters, more information can be found at [Element Software](#).

NetApp HCI

NetApp HCI simplifies the management and scale of the datacenter by automating routine tasks and enabling infrastructure administrators to focus on more important functions.

NetApp HCI is fully supported by Trident, it can provision and manage storage devices for containerized applications directly against the underlying HCI storage platform. For more information about NetApp HCI visit [NetApp HCI](#).

3.3.4 SANtricity

NetApp's E and EF Series storage platforms, using the SANtricity operating system, provides robust storage that is highly available, performant, and capable of delivering storage services for applications at any scale.

Trident can create and manage SANtricity volumes across the portfolio of products.

For more information about SANtricity and the storage platforms which use it, see [SANtricity Software](#).

3.3.5 Azure NetApp Files

[Azure NetApp Files](#) is an enterprise-grade Azure file share service, powered by NetApp. Run your most demanding file-based workloads in Azure natively, with the performance and rich data management you expect from NetApp.

3.3.6 Cloud Volumes Service for AWS

[NetApp Cloud Volumes Service](#) is a cloud native file service that provides NAS volumes over NFS and SMB with all-flash performance. This service enables any workload, including legacy applications, to run in the AWS cloud. It provides a fully managed service which offers consistent high performance, instant cloning, data protection and secure access to Elastic Container Service (ECS) instances.

3.4 Kubernetes Cluster Architecture and Considerations

Kubernetes is extremely flexible and is capable of being deployed in many different configurations. It supports clusters as small as a single node and as large as a [few thousand](#). It can be deployed using either physical or virtual machines on premises or in the cloud. However, single node deployments are mainly used for testing and are not suitable for production workloads. Also, hyperscalers such as AWS, Google Cloud and Azure abstract some of the initial and basic deployment tasks away. When deploying Kubernetes, there are a number of considerations and decisions to make which can affect the applications and how they consume storage resources.

3.4.1 Cluster concepts and components

A Kubernetes cluster typically consists of two types of nodes, each responsible for different aspects of functionality:

- Master nodes – These nodes host the control plane aspects of the cluster and are responsible for, among other things, the API endpoint which the users interact with and provides scheduling for pods across resources. Typically, these nodes are not used to schedule application workloads.
- Compute nodes – Nodes which are responsible for executing workloads for the cluster users.

The cluster has a number of Kubernetes intrinsic services which are deployed in the cluster. Depending on the service type, each service is deployed on only one type of node (master or compute) or on a mixture of node types. Some of these services, such as etcd and DNS, are mandatory for the cluster to be functional, while other services are optional. All of these services are deployed as pods within Kubernetes.

- etcd – It is a distributed key-value datastore. It is used heavily by Kubernetes to track the state and manage the resources associated with the cluster.
- DNS – Kubernetes maintains an internal DNS service to provide local resolution for the applications which have been deployed. This enables inter-pod communication to happen while referencing friendly names instead of internal IP addresses which can change as the container instances are scheduled.
- API Server - Kubernetes deploys the API server to allow interaction between kubernetes and the outside world. This is deployed on the master node(s).

- Dashboard – An optional component which provides a graphical interface to the cluster.
- Monitoring and logging – An optional components which can aid with resource reporting.

Note: We have not discussed Kubernetes container networking to allow pods to communicate with each other, or to outside the cluster. The choice of using an overlay network (e.g. Flannel) or a straight Layer-3 solution (e.g. Calico) is out of the scope of this document and does not affect access to storage resources by the pods.

3.4.2 Cluster architectures

There are three primary Kubernetes cluster architectures. These accommodate various methods of high availability and recoverability of the cluster, its services, and the applications running. Trident is installed the same no matter which Kubernetes architecture is chosen.

Master nodes are critical to the operation of the cluster. If no masters are running, or the master nodes are unable to reach a quorum, then the cluster is unable to schedule and execute applications. The master nodes are the control plane for the cluster and consequentially there should be special consideration given to their [sizing](#) and count.

Compute nodes are, generally speaking, much more disposable. However, extra resources must be built into the compute infrastructure to accommodate any workloads from failed nodes. Compute nodes can be added and removed from the cluster as needed quickly and easily to accommodate the scale of the applications which are being hosted. This makes it very easy to burst, and reclaim, resources based on real-time application workload.

Single master, compute

This architecture is the easiest to deploy but does not provide high availability of the core management services. In the event the master node is unavailable, no interaction can happen with the cluster until, at a minimum, the Kubernetes API server is returned to service.

This architecture can be useful for testing, qualification, proof-of-concept, and other non-production uses, however it should never be used for production deployments.

A single node used to host both the master service and the workloads is a variant of this architecture. Using a single node kubernetes cluster is useful when testing or experimenting with different concepts and capabilities. However, the limited scale and capacity make it unreasonable for more than very small tests. The Trident [:ref: quick start guide <Simple Kubernetes install>](#) outlines the process to instantiate a single node Kubernetes cluster with Trident that provides full functionality for testing and validation.

Multiple master, compute

Having multiple master nodes ensures that services remain available should master node(s) fail. In order to facilitate availability of master services, they should be deployed with odd numbers (e.g. 3,5,7,9 etc.) so that quorum (master node majority) can be maintained should one or more masters fail. In the HA scenario, Kubernetes will maintain a copy of the etcd databases on each master, but hold elections for the control plane function leaders *kube-controller-manager* and *kube-scheduler* to avoid conflicts. The worker nodes can communicate with any master's API server through a load balancer.

Deploying with multiple masters is the minimum recommended configuration for most production clusters.

Pros:

- Provides highly-available master services, ensuring that the loss of up to $(n/2) - 1$ master nodes will not affect cluster operations.




Fig. 2: Multiple master architecture

Cons:

- More complex initial setup.

Master, etcd, compute

This architecture isolates the etcd cluster from the other master server services. This removes workload from the master servers, enabling them to be sized smaller, and makes their scale out (or in) more simple. Deploying a Kubernetes cluster using this model adds a degree of complexity, however, it adds flexibility to the scale, support, and management of the etcd service used by Kubernetes, which may be desirable to some organizations.

Pros:

- Provides highly-available master services, ensuring that the loss of up to $(n/2) - 1$ master nodes will not affect cluster operations.
- Isolating etcd from the other master services reduces the workload for master servers.
- Decoupling etcd from the masters makes etcd administration and protection easier. Independent management allows for different protection and scaling schemes.

Cons:

- More complex initial setup.

Red Hat OpenShift infrastructure architecture

In addition to the architectures referenced above, Red Hat's OpenShift introduces the concept of [infrastructure nodes](#). These nodes host cluster services such as log aggregation, metrics collection and reporting, container registry services, and overlay network management and routing.

Red Hat recommends a minimum of three infrastructure nodes for production deployments. This ensures that the services have resources available and are able to migrate in the event of host maintenance or failure.




Fig. 3: Multiple master, etcd, compute architecture

This architecture enables the services which are critical to the cluster, i.e. registry, overlay network routing, and others to be hosted on dedicated nodes. These dedicated nodes may have additional redundancy, different CPU/RAM requirements, and other low-level differences from compute nodes. This also makes adding and removing compute nodes as needed easier, without needing to worry about core services being affected by a node being evacuated.

An additional option involves separating out the master and etcd roles into different servers in the same way as can be done in Kubernetes. This results in having master, etcd, infrastructure, and compute node roles. Further details, including examples of OpenShift node roles and potential deployment options, can be found in the [Red Hat documentation](#).

3.4.3 Choosing an architecture

Regardless of the architecture that you choose, it's important to understand the ramifications to high availability, scalability, and serviceability of the component services. Be sure to consider the effect on the applications being hosted by the Kubernetes or OpenShift cluster. The architecture of the storage infrastructure supporting the Kubernetes/OpenShift cluster and the hosted applications can also be affected by the chosen cluster architecture, such as where etcd is hosted.

3.5 Storage for Kubernetes Infrastructure Services

Trident focuses on providing persistence to Kubernetes applications, but before you can host those applications, you need to run a healthy, protected Kubernetes cluster. Those clusters are made up of a number of services with their own persistence requirements that need to be considered.

Node-local container storage, a.k.a. graph driver storage

One of the often overlooked components of a Kubernetes deployment is the storage which the container instances consume on the Kubernetes cluster nodes, usually referred to as [Graph Driver Storage](#). When a container is instantiated on a node it consumes capacity and IOPS to do many of its operations, as only data which is read from or written to a persistent volume is offloaded to the external storage. If the Kubernetes nodes are expected to host dozens, hundreds,




Fig. 4: OpenShift, Multiple master, infra, compute architecture

or more containers this may be a significant amount of temporary capacity and IOPS which are expected of the node-local storage.

Even if you don't have a requirement to keep the data, the containers still need enough performance and capacity to execute their application. The Kubernetes administrator and storage administrator should work closely together to determine the requirements for graph storage and ensure adequate performance and capacity is available.

The typical method of augmenting the graph driver storage is to use a block device mounted at the location where the container instance storage is located, e.g. `/var/lib/docker`. The host operating systems being used to underpin the Kubernetes deployment will each have different methods for how to replace the graph storage with something more robust than a simple directory on the node. Refer to the documentation from Red Hat, Ubuntu, SUSE, etc. for those instructions.

Note: Block protocol is specifically recommended for graph storage due to the nature of how the graph drivers work. In particular, they create thin clones, using a variety of methods depending on the driver, of the container image. NFS does not support this functionality and results in a full copy of the container image file system for each instance, resulting in significant performance and capacity implications.

If the Kubernetes nodes are virtualized, this could also be addressed by ensuring that the datastore they reside on meets the performance and capacity needs, however, the flexibility of having a separate device, even an additional virtual disk, should be carefully considered. Using a separate device gives the ability to independently control capacity, performance, and data protection to tailor the policies according to needs. Often the capacity and performance needed for graph storage can fluctuate dramatically, however, data protection is not necessary.

3.5.1 Persistent storage for cluster services

There are several critical services hosted on the kubernetes master nodes and other cluster critical services which may be hosted on other node types.

Using capacity provisioned from enterprise storage adds several benefits for each service as delineated below:

- Performance – Leveraging enterprise storage means being able to provide latency in-line with application needs and controlling performance using QoS policies. This can be used to limit performance for **bullies** or ensure performance for applications as needed.
- Resiliency – In the event of a node failure, being able to quickly recover the data and present it to a replacement ensures that the application is minimally affected.
- Data protection – Putting application data onto dedicated volumes allows the data to have its own snapshot, replication, and retention policies.
- Data security – Ensuring that data is encrypted and protected all the way to the disk level is important for meeting many compliance requirements.
- Scale – Using enterprise storage enables deploying fewer instances, with the instance count depending on compute resources, rather than having to increase the total count for resiliency and performance purposes. Data which is protected automatically, and provides adequate performance, means that horizontal scale-out doesn't need to compensate for limited performance.

There are some best practices which apply across all of the cluster services, or any application, which should be addressed as well.

- Determining the amount of acceptable data loss, i.e. the **Recovery Point Objective (RPO)**, is a critical part of deploying a production level system. Having an understanding of this will greatly simplify the decisions around other items described in this section.
- Cluster service volumes should have a snapshot policy which enables the rapid recovery of data according to your requirements, as defined by the RPO. This enables quick and easy restoration of a service to a point in time by reverting the volume or the ability to recover individual files if needed.
- Replication can provide both backup and disaster recovery capabilities, each service should be evaluated and have the recovery plan considered carefully. Using storage replication may provide rapid recovery with a higher RPO, or can provide a starting baseline which expedites restore operations without having to recover all data from other locations.

etcd considerations

You should carefully consider the high availability and data protection policy for the etcd instance used by the Kubernetes master(s). This service is arguably the most critical to the overall functionality, recoverability, and serviceability of the cluster as a whole, so it's imperative that its deployment meets your goals.

The most common method of providing high availability and resiliency to the etcd instance is to horizontally scale the application by having multiple instances across multiple nodes. A minimum of three nodes is recommended.

Kubernetes etcd will, by default, use local storage for its persistence. This holds true whether etcd is co-located with other master services or is hosted on dedicated nodes. To use enterprise storage for etcd a volume must be provisioned and mounted to the node at the correct location prior to deployment. This storage cannot be dynamically provisioned by Trident or any other storage class provisioner as it is needed prior to the Kubernetes cluster being operational.

Refer to your Kubernetes provider's documentation on where to mount the volume and/or customize the etcd configuration to use non-default storage.

logging considerations

Centralized logging for applications deployed to the Kubernetes cluster is an optional service. Using enterprise storage for logging has the same benefits as with etcd: performance, resiliency, protection, security, and scale.

Depending on how and when the service is deployed, the storage class may define how to dynamically provision storage for the service. Refer to your vendor's documentation on how to customize the storage for logging services.

metrics considerations

Many third-party metrics and analytics services are available for monitoring and reporting the status and health of every aspect of the cluster and application. Many of these tools require persistent storage with specific performance characteristics for it to function as intended.

Architecturally, many of these function similarly. An agent exists on each node which forwards that node's data to a centralized collector. The collector aggregates, analyzes, and displays the data. Similar to the logging service, using enterprise storage allows the aggregation service to move across nodes in the cluster seamlessly and ensures the data is protected in the event of node failure.

Each vendor has different recommendations and deployment methodology. Work with your vendor to identify requirements and, if needed, provision storage from an enterprise array to meet the requirements. This document will discuss the Red Hat OpenShift metrics service in a later chapter.

registry considerations

The users and applications interact most often with the registry service. Slow access to the registry service can have a dramatic effect on the Kubernetes cluster's perceived performance as a whole. A slow image push or pull operation can result in lengthy times for tasks which can directly affect the developer and application.

Fortunately, the registry is flexible with regard to storage protocol. Keep in mind different protocols have different implications.

- Object storage is the default recommendation and is the simplest to use for Kubernetes deployments which expect to have a significant scale or where the images need to be accessed across geographic regions.
- NFS is a good choice for many deployments as it allows a single repository for the container images while allowing many registry endpoints to front the capacity.
- Block protocols, such as iSCSI, can be used for registry storage, but they introduce a single point of failure. The block device can only be attached to a single registry node due to the single-writer limitation of the supported filesystems.

Each organization and each application has different priorities for protecting the images stored in the registry. Generally, registry images are either cached from upstream registries or have images pushed to them during the application build process. The Recovery Time Objective (RTO) influences the chosen protection scheme because it affects the recovery process. If RTO is not an issue, then the applications may be able to simply rebuild the container images and push them into a new instance of the registry. If faster RTO is desired, then use a replication policy which adheres to the desired recovery goal.

3.5.2 Design choices and guidelines when using ONTAP

A few design and implementation considerations should be addressed prior to deployment using Trident and ONTAP.

Storage Virtual Machines

Storage Virtual Machines (SVMs) are used for administrative delegation within ONTAP. By creating different SVMs, the storage administrator can isolate a particular user, group, or application to only specific resources. When Trident accesses the storage system via an SVM, it is prevented from doing many system level management tasks, which provides additional isolation of capabilities for storage provisioning and management tasks.

Trident can leverage SVMs with several different methods, each is explained below. Multiple Trident deployments, i.e. multiple Kubernetes clusters, does not change the below methods. When an SVM is shared with multiple Trident instances they simply need distinct prefixes defined in the backend configuration files.

SVM shared with non Trident-managed workloads

This configuration uses a single, or small number of, SVMs to host all of the workloads on the cluster and results in the containerized applications being hosted by the same SVM as other non-containerized workloads. The shared SVM model is common in organizations where there exists multiple network segments which are isolated and adding additional IP addresses is difficult or impossible.

There is nothing inherently wrong with this configuration, however, it is more challenging to apply policies which affect only the container workloads.

Dedicated SVM for Trident-managed workloads

By default, NetApp recommends creating an SVM used solely by Trident for provisioning and deprovisioning volumes. This enables the storage administrator to put controls in place to limit the amount of resources which Trident is able to consume.

As mentioned above, having multiple Kubernetes clusters connect to and consume storage from the same SVM is acceptable, the only change to the Trident configuration should be to *provide a different prefix*.

When creating different backends which connect to the same underlying SVM resources but have different features applied (e.g. snapshot policies) use different prefixes. The unique prefixes help identify the volumes to the storage administrator and ensure no confusion results.

Multiple SVMs dedicated to Trident-managed workloads

You may consider using multiple SVMs with Trident for many different reasons, including isolating applications and resource domains, strict control over resources, and to facilitate multitenancy. It's also worth considering using at least two SVMs with any Kubernetes cluster to isolate persistent storage for cluster services from application storage.

Using multiple SVMs with one dedicated to cluster services isolates and controls the workload. Since the persistent storage resources needed by the Kubernetes cluster must exist prior to Trident deployment, the Kubernetes cluster services SVM will not have dynamic provisioning in the same manner as the application SVM will have.

Kubernetes cluster services

Even for cluster services persistent volumes created by Trident, there should be serious consideration given to using per-volume QoS policies, including QoS minimums when possible, and customizing the volume options for the application. Below are the default recommendations for the cluster services, however you should evaluate your needs and adjust policies according to your data protection, performance, and availability requirements. Despite these recommendations, you will still want to evaluate and determine what works best for your Kubernetes cluster and applications.

etcd

- The default snapshot policy is often adequate for protecting against data corruption and loss, however, snapshots are not a backup strategy. Some consideration should be given to increasing the frequency, and decreasing the retention period for etcd volumes. For example, keeping 24 hourly snapshots or 48 snapshots taken every 30 minutes, but not retaining them for more than one or two days. Since any data loss for etcd can be problematic, having more frequent snapshots makes this scenario easier to recover from.
- If the disaster recovery plan is to recover the Kubernetes cluster as-is at the destination site, then these volumes should be replicated with SnapMirror or SnapVault.
- etcd does not have significant IOPS or throughput requirements, however, latency can play a critical role in the responsiveness of the Kubernetes API server. Whenever possible use the lowest latency storage or set your storage to the lowest latency possible.
- A QoS policy should be leveraged to provide a minimum amount of IOPS to the etcd volume(s). The minimum value will depend on the number of nodes and pods which are deployed to your Kubernetes cluster. Monitoring should be used to verify that the configured policy is adequate and adjusted over time as the Kubernetes cluster expands.

- The etcd volumes should have their export policy or iGroup limited to only the nodes which are hosting, or could potentially host, etcd instances.

logging

- Volumes which are providing storage capacity for aggregated logging services need to be protected, however, an average RPO is adequate in many instances since logging data is often not critical to recovery. If your application has strict compliance requirements, this may be different however.
- Using the default snapshot policy is generally adequate. Optionally, depending on the needs of the administrators, reducing the snapshot policy to one which keeps as few as seven daily snapshots may be acceptable.
- Logging volumes should be replicated to protect the historical data for use by the application and by administrators, however, recovery may be deprioritized for other services.
- Logging services have widely variable IOPS requirements and read/write patterns. It's important to consider the number of nodes, pods, and other objects in the cluster. Each of these will generate data which needs to be stored, indexed, analyzed, and presented, so a larger cluster may have substantially more data than expected.
- A QoS policy may be leveraged to provide both a minimum and maximum amount of throughput available. Note that the maximum may need to be adjusted as additional pods are deployed, so close monitoring of the performance should be used to verify that logging is not being adversely affected by storage performance.
- The volumes export policy or iGroup should be limited to nodes which host the logging service. This will depend on the particular solution used and the chosen configuration. For example, OpenShift's logging service is deployed to the infrastructure nodes.

metrics

- Kubernetes autoscale feature relies on metrics to provide data for when scale operations need to occur. Also, metrics data often plays a critical role in show-back and charge-back operations, so ensure that you are working to address the needs of the entire business with the RPO policy. Also, ensure that your RPO and RTO meet the needs of these functions.
- As the number of cluster nodes and deployed pods increases, so too does the amount of data which is collected and retained by the metrics service. It's important to understand the performance and capacity recommendations provided by the vendor for your metrics service as they can vary dramatically, particularly depending on the amount of time for which the data is retained and the number of metrics which are being monitored.
- A QoS policy can be used to limit the number of IOPS or throughput which the metrics services uses, however it is generally not necessary to use a minimum policy.
- It is recommended to limit the export policy or iGroup to the hosts which the metrics service is executed from. Note that it's important to understand the architecture of your metrics provider. Many have agents which run on all hosts in the cluster, however, those will report metrics to a centralised repository for storage and reporting. Only that group of nodes needs access.

registry

- Using a snapshot policy for the registry data may be valuable for recovering from data corruption or other issues, however it is not necessary. A basic snapshot policy is recommended, however, individual container images cannot be recovered (they are stored in a hashed manner), only a full volume revert can be used to recover data.
- The workload for the registry can vary widely, however, the general rule is that push operations happen infrequently, while pull operations happen frequently. If a CI/CD pipeline process is used to build, test, and deploy the application(s) this may result in a predictable workload. Alternatively, and even with a CI/CD system in use, the workload can vary based on application scaling requirements, build requirements, and even Kubernetes node add/remove operations. Close monitoring of the workload should be implemented to adjust as necessary.
- A QoS policy may be implemented to ensure that application instances are still able to pull and deploy new container images regardless of other workloads on the storage system. In the event of a disaster recovery

scenario, the registry may have a heavy read workload while applications are instantiated on the destination site. The configured QoS minimum policy will prevent other disaster recovery operations from slowing application deployment.

3.6 Storage Configuration for Trident

Each storage platform in NetApp's portfolio has unique capabilities that benefit applications, containerized or not. Trident works with each of the major platforms: ONTAP, Element, and E-Series. There is not one platform which is better suited for all applications and scenarios than another, however, the needs of the application and the team administering the device should be taken into account when choosing a platform.

The storage administrator(s), Kubernetes administrator(s), and the application team(s) should work with their NetApp team to ensure that the most appropriate platform is selected.

Regardless of which NetApp storage platform you will be connecting to your Kubernetes environment, you should follow the baseline best practices for the host operating system with the protocol that you will be leveraging. Optionally, you may want to consider incorporating application best practices, when available, with backend, storage class, and PVC settings to optimize storage for specific applications.

Some of the best practices guides are listed below, however, refer to the [NetApp Library](#) for the most current versions.

- ONTAP
 - [NFS Best Practice and Implementation Guide](#)
 - [SAN Administration Guide \(for iSCSI\)](#)
 - [iSCSI Express Configuration for RHEL](#)
- Element
 - [Configuring SolidFire for Linux](#)
 - [NetApp HCI Deployment Guide](#)
- E-Series
 - [Installing and Configuring for Linux Express Guide](#)

Some example application best practices guides:

- [ONTAP Best Practice Guidelines for MySQL](#)
- [MySQL Best Practices for NetApp SolidFire](#)
- [NetApp SolidFire and Cassandra](#)
- [Oracle Best Practices on NetApp SolidFire](#)
- [PostgreSQL Best Practices on NetApp SolidFire](#)

Not all applications will have specific guidelines, it's important to work with your NetApp team and to refer to the [library](#) for the most up-to-date recommendations and guides.

3.6.1 Best practices for configuring ONTAP and Cloud Volumes ONTAP

The following recommendations are guidelines for configuring ONTAP for containerized workloads which consume volumes that are dynamically provisioned by Trident. Each should be considered and evaluated for appropriateness in your environment.

Use SVM(s) which are dedicated to Trident

Storage Virtual Machines (SVMs) provide isolation and administrative separation between tenants on an ONTAP system. Dedicating an SVM to applications enables the delegation of privileges and enables applying best practices for limiting resource consumption.

There are several options available for the management of the SVM:

- Provide the cluster management interface in the backend configuration, along with appropriate credentials, and specify the SVM name.
- Create a dedicated management interface for the SVM using ONTAP System Manager or CLI.
- Share the management role with an NFS data interface.

In each case, the interface should be in DNS, and the DNS name should be used when configuring Trident. This helps to facilitate some DR scenarios, for example, SVM-DR without the use of network identity retention.

There is no preference between having a dedicated or shared management LIF for the SVM, however, you should ensure that your network security policies align with the approach you choose. Regardless, the management LIF should be accessible via DNS to facilitate maximum flexibility should SVM-DR be used in conjunction with Trident.

Limit the maximum volume count

ONTAP storage systems have a maximum volume count which varies based on the software version and hardware platform, see the [Hardware Universe](#) for your specific platform and ONTAP version to determine exact limits. When the volume count is exhausted, provisioning operations will fail not only for Trident but for all storage requests.

Trident's `ontap-nas` and `ontap-san` drivers provision a FlexVolume for each Kubernetes Persistent Volume (PV) which is created. The `ontap-nas-economy` driver creates approximately one FlexVolume for every 200 PVs. To prevent Trident from consuming all available volumes on the storage system, it is recommended that a limit be placed on the SVM. This can be done from the command line:

```
vserver modify -vserver <svm_name> -max-volumes <num_of_volumes>
```

The value for `max-volumes` will vary based on several criteria specific to your environment:

- The number of existing volumes in the ONTAP cluster
- The number of volumes you expect to provision outside of Trident for other applications
- The number of persistent volumes expected to be consumed by Kubernetes applications

The `max-volumes` value is total volumes provisioned across all nodes in the ONTAP cluster, not to an individual ONTAP node. As a result, some conditions may be encountered where an ONTAP cluster node may have far more, or less, Trident provisioned volumes than another.

For example, a 2-node ONTAP cluster has the ability to host a maximum of 2000 FlexVolumes. Having the maximum volume count set to 1250 appears very reasonable. However, if only [aggregates](#) from one node are assigned to the SVM, or the aggregates assigned from one node are unable to be provisioned against (e.g. due to capacity), then the other node will be the target for all Trident provisioned volumes. This means that the volume limit may be reached for that node before the `max-volumes` value is reached, resulting in impacting both Trident and other volume operations using that node. Avoid this situation by ensuring that aggregates from each node in the cluster are assigned to the SVM used by Trident in equal numbers.

In addition to controlling the volume count at the storage array, Kubernetes capabilities should also be leveraged as explained in the next chapter.

Limit the maximum size of volumes created by Trident

To configure the maximum size for volumes that can be created by Trident, use the `limitVolumeSize` parameter in your `backend.json` definition.

In addition to controlling the volume size at the storage array, Kubernetes capabilities should also be leveraged as explained in the next chapter.

Create and use an SVM QoS policy

Leveraging an ONTAP QoS policy, applied to the SVM, limits the number of IOPS consumable by the Trident provisioned volumes. This helps to [prevent a bully](#) or out-of-control container from affecting workloads outside of the Trident SVM.

Creating a QoS policy for the SVM can be done in a few steps. Refer to the documentation for your version of ONTAP for the most accurate information. The example below creates a QoS policy which limits the total IOPS available to the SVM to 5000.

```
# create the policy group for the SVM
qos policy-group create -policy-group <policy_name> -vserver <svm_name> -max-
↳throughput 5000iops

# assign the policy group to the SVM, note this will not work
# if volumes or files in the SVM have existing QoS policies
vserver modify -vserver <svm_name> -qos-policy-group <policy_name>
```

Additionally, if your version of ONTAP supports it, you may consider using a QoS minimum in order to guarantee an amount of throughput to containerized workloads. Adaptive QoS is not compatible with an SVM level policy.

The number of IOPS dedicated to the containerized workloads depends on many aspects. Among other things, these include:

- Other workloads using the storage array. If there are other workloads, not related to the Kubernetes deployment, utilizing the storage resources, then care should be taken to ensure that those workloads are not accidentally adversely impacted.
- Expected workloads running in containers. If workloads which have high IOPS requirements will be running in containers, then a low QoS policy will result in a bad experience.

It's important to remember that a QoS policy assigned at the SVM level will result in all volumes provisioned to the SVM sharing the same IOPS pool. If one, or a small number, of the containerized applications have a high IOPS requirement it could become a bully to the other containerized workloads. If this is the case, you may want to consider using external automation to assign per-volume QoS policies.

Limit storage resource access to Kubernetes cluster members

Limiting access to the NFS volumes and iSCSI LUNs created by Trident is a critical component of the security posture for your Kubernetes deployment. Doing so prevents hosts which are not a part of the Kubernetes cluster from accessing the volumes and potentially modifying data unexpectedly.

It's important to understand that namespaces are the logical boundary for resources in Kubernetes. The assumption is that resources in the same namespace are able to be shared, however, importantly, there is no cross-namespace capability. This means that even though PVs are global objects, when bound to a PVC they are only accessible by pods which are in the same namespace. It's critical to ensure that namespaces are used to provide separation when appropriate.

The primary concern for most organizations, with regard to data security in a Kubernetes context, is that a process in a container can access storage mounted to the host, but which is not intended for the container. Simply put, this is not possible, the underlying technology for containers i.e. `namespaces`, are designed to prevent this type of compromise. However, there is one exception: privileged containers.

A privileged container is one that is run with substantially more host-level permissions than normal. These are not denied by default, so disabling the capability using `pod security policies` is very important for preventing this accidental exposure.

For volumes where access is desired from both Kubernetes and external hosts, the storage should be managed in a traditional manner, with the PV introduced by the administrator and not managed by Trident. This ensures that the storage volume is destroyed only when both the Kubernetes and external hosts have disconnected and are no longer using the volume. Additionally, a custom export policy can be applied which enables access from the Kubernetes cluster nodes and targeted servers outside of the Kubernetes cluster.

For deployments which have dedicated infrastructure nodes (e.g. OpenShift), or other nodes which are not schedulable for user applications, separate export policies should be used to further limit access to storage resources. This includes creating an export policy for services which are deployed to those infrastructure nodes, such as, the OpenShift Metrics and Logging services, and standard applications which are deployed to non-infrastructure nodes.

Create an export policy

Create appropriate export policies for Storage Virtual Machines. Allow only Kubernetes nodes access to the NFS volumes.

Export policies contain one or more export rules that process each node access request. Use the `vserver export-policy create ONTAP CLI` to create the export policy. Add rules to the export policy using the `vserver export-policy rule create ONTAP CLI` command. Performing the above commands enables you to restrict which Kubernetes nodes have access to data.

Disable `showmount` for the application SVM

The `showmount` feature enables an NFS client to query the SVM for a list of available NFS exports. A pod deployed to the Kubernetes cluster could issue the `showmount -e` command against the data LIF and receive a list of available mounts, including those which it does not have access to. While this isn't, by itself, dangerous or a security compromise, it does provide unnecessary information potentially aiding an unauthorized user with connecting to an NFS export.

Disable `showmount` using SVM level ONTAP CLI command:

```
vserver nfs modify -vserver <svm_name> -showmount disabled
```

3.6.2 Best practices for configuring SolidFire

Solidfire Account

Create a SolidFire account. Each SolidFire account represents a unique volume owner and receives its own set of Challenge-Handshake Authentication Protocol (CHAP) credentials. You can access volumes assigned to an account either by using the account name and the relative CHAP credentials or through a volume access group. An account can have up to two-thousand volumes assigned to it, but a volume can belong to only one account.

SolidFire QoS

Use QoS policy if you would like to create and save a standardized quality of service setting that can be applied to many volumes.

Quality of Service parameters can be set on a per-volume basis. Performance for each volume can be assured by setting three configurable parameters that define the QoS: Min IOPS, Max IOPS, and Burst IOPS.

The following table shows the possible minimum, maximum, and Burst IOPS values for 4Kb block size.

IOPS Parameter	Definition	Min value	Default Value	Max Value(4Kb)
Min IOPS	The guaranteed level of performance for a volume.	50	50	15000
Max IOPS	The performance will not exceed this limit.	50	15000	200,000
Burst IOPS	Maximum IOPS allowed in a short burst scenario.	50	15000	200,000

Note: Although the Max IOPS and Burst IOPS can be set as high as 200,000, the real-world maximum performance of a volume is limited by cluster usage and per-node performance.

Block size and bandwidth have a direct influence on the number of IOPS. As block sizes increase, the system increases bandwidth to a level necessary to process the larger block sizes. As bandwidth increases the number of IOPS the system is able to attain decreases. For more information on QoS and performance, refer to the [NetApp SolidFire Quality of Service \(QoS\) Guide](#).

SolidFire authentication

SolidFire Element supports two methods for authentication: CHAP and Volume Access Groups (VAG). CHAP uses the CHAP protocol to authenticate the host to the backend. Volume Access Groups controls access to the volumes it provisions. NetApp recommends using CHAP for authentication as it's simpler and has no scaling limits.

Note: Trident as an enhanced CSI provisioner supports the use of CHAP authentication. VAGs should only be used in the traditional non-CSI mode of operation.

CHAP authentication (verification that the initiator is the intended volume user) is supported only with account-based access control. If you are using CHAP for authentication, 2 options are available: unidirectional CHAP and bidirectional CHAP. Unidirectional CHAP authenticates volume access by using the SolidFire account name and initiator secret. The bidirectional CHAP option provides the most secure way of authenticating the volume since the volume authenticates the host through the account name and the initiator secret, and then the host authenticates the volume through the account name and the target secret.

However, if CHAP is cannot be enabled and VAGs are required, create the access group and add the host initiators and volumes to the access group. Each IQN that you add to an access group can access each volume in the group with or without CHAP authentication. If the iSCSI initiator is configured to use CHAP authentication, account-based access control is used. If the iSCSI initiator is not configured to use CHAP authentication, then Volume Access Group access control is used.

Note: VAGs are only supported by Trident in non-CSI mode of operation.

For more information on how to setup Volume Access Groups and CHAP authentication, please refer the NetApp HCI Installation and setup guide.

3.6.3 Best practices for configuring E-Series

E-Series Disk Pools and Volume Groups

Create disk pools or volume groups based on your requirement and determine how the total storage capacity must be organized into volumes and shared among hosts. Both the disk pool and the volume group consist of a set of drives which are logically grouped to provide one or more volumes to an application host. All of the drives in a disk pool or volume group should be of the same media type.

E-Series Host Groups

Host groups are used by Trident to access the volumes (LUNs) that it provisions. By default, Trident uses the host group called “trident” unless a different host group name is specified in the configuration. Trident, by itself will not create or manage host groups. Host group has to be created before the E-Series storage backend is setup on Trident. Make sure that all the Kubernetes worker nodes iSCSI IQN names are updated in the host group.

E-Series Snapshot Schedule

Create a snapshot schedule and assign the volume created by Trident to a snapshot schedule so that volume backups can be taken at the required interval. Based on the snapshots taken as per the snapshot policy, rollback operations can be done on volumes by restoring a snapshot image to the base volume. Setting up Snapshot Consistency Groups are also ideal for applications that span multiple volumes. The purpose of a consistency group is to take simultaneous snapshot images of multiple volumes, thus ensuring consistent copies of a collection of volumes at a particular point in time. Snapshot schedule and Consistency group cannot be set through Trident. It has to be configured through SANtricity System Manager

3.6.4 Best practices for configuring Cloud Volumes Service on AWS

Create Export Policy

To make sure that only the authorized set of nodes have access to the volume provisioned through Cloud Volumes Service, set appropriate rules for the export policy while creating a Cloud Volumes Service. When provisioning volumes on Cloud Volume Services through Trident make sure to use `exportRule` parameter in the backend file to give access to the required Kubernetes nodes.

Create Snapshot Policy

Setting a snapshot policy for the volumes provisioned through Cloud Volume Service makes sure that snapshots are taken at required intervals. This guarantees data backup at regular intervals and data can be restored in the event of a data loss or data corruption. Snapshot Policy for volumes hosted by Cloud Volume Service can be set by selecting the appropriate schedule on the volumes details page.

Choosing the appropriate Service Level, Storage Capacity and Storage Bandwidth

AWS Cloud Volume Services offer different Service Levels like Standard, Premium and Extreme. These Service Levels cater to different storage capacity and storage bandwidth requirements. Make sure to select the appropriate Service Level based on your business needs.

The required size of allocated storage should be chosen during volume creation based on the specific needs of the application. There are two factors that need to be taken into consideration while deciding on the allocated storage. They are the storage requirements of the specific application and the bandwidth that you require at the peak or the edge.

The bandwidth depends on the combination of the service level and the allocated capacity that you have chosen. Therefore choose the right service level and allocated capacity keeping the required bandwidth in mind.

Limit the maximum size of volumes created by the Trident

It's possible to restrict the maximum size of volumes created by the Trident on AWS Cloud Volume Services using the `limitVolumeSize` parameter in the backend configuration file. Setting this parameter makes sure that provisioning fails if the requested volume size is above the set value.

3.7 Deploying Trident

The guidelines in this section provide recommendations for Trident installation with various Kubernetes configurations and considerations. As with all the other recommendations in this guide, each of these suggestions should be carefully considered to determine if it's appropriate and will provide benefit to your deployment.

3.7.1 Supported Kubernetes cluster architectures

Trident is supported with the following Kubernetes architectures.

Kubernetes Cluster Architectures	Supported	Default Install
Single master, compute	Yes	Yes
Multiple master, compute	Yes	Yes
Master, etcd, compute	Yes	Yes
Master, infrastructure, compute	Yes	Yes

3.7.2 Trident installation modes

Three ways to install Trident are discussed in this chapter.

Normal install mode

Normal installation involves running the `tridentctl install -n trident` command which deploys the Trident pod on the Kubernetes cluster. Trident installation is quite a straightforward process. For more information on installation and provisioning of volumes, refer to the [Deploying documentation](#).

Offline install mode

In many organizations, production and development environments do not have access to public repositories for pulling and posting images as these environments are completely secured and restricted. Such environments only allow pulling images from trusted private repositories. In such scenarios, make sure that a private registry instance is available. The trident image should be downloaded from a bastion host with internet access and pushed on to the private registry. To install Trident in offline mode, just issue the `tridentctl install -n trident` command with the `--trident-image` parameter set to the appropriate image name and location.

Remote install mode

Trident can be installed on a Kubernetes cluster from a remote machine. To do a remote install, install the appropriate version of `kubectl` on the remote machine from where you would be running the `tridentctl install` command. 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 you can connect to the required Kubernetes cluster. Complete the Trident deployment from the remote machine using the normal installation steps.

3.7.3 Trident Installation on Docker UCP 3.1

Docker EE Universal Control Plane (UCP) is the cluster management layer that sits on top of the Docker Enterprise Engine. Once deployed, administrators interact with their cluster via UCP instead of each node's individual Docker engines. Since the UCP supports both Kubernetes and Docker via a Web UI or CLI, administrators can use either a Kubernetes YAMLS or Docker Compose files to deploy images from the centralized location. It also provides cluster-wide monitoring of deployed containers, services, and pods.

Installing Trident for Kubernetes on UCP managed nodes is similar to installing Trident on Kubernetes. Refer to the following [documentation](#) for instructions on how to install Trident for Kubernetes.

Please note that starting with Docker EE 2.1 UCP and Trident 19.01, it's no longer required to specify the `--ucp-host` and `--ucp-bearer-token` parameters while installing and uninstalling Trident. Deploy the `tridentctl install -n <namespace>` command to start the installation on the UCP managed nodes.

3.7.4 Using Trident with the NetApp Kubernetes Service

The [NetApp Kubernetes Service](#) (NKS) is a universal control plane through which production grade Kubernetes clusters can be provisioned and run on the cloud provider of choice.

Refer to the following [documentation](#) for instructions on how to install Trident for NKS.

3.7.5 Deploying Trident as an enhanced CSI Provisioner

The Trident 19.07 release is built on a production-ready CSI 1.1 provisioner implementation. This allows Trident to absorb standardized features like snapshots, while still retaining its ability to innovate on the storage model.

To setup Trident as a CSI provisioner, refer to the [Deployment Guide](#). Ensure that the required [Feature Gates](#) are enabled. After deploying, you should consider [Upgrading existing PVs to CSI volumes](#).

3.7.6 CRDs for maintaining Trident's state

The 19.07 release of Trident introduces a set of [Custom Resource Definitions \(CRDs\)](#) for maintaining Trident's stateful information. CRDs are a Kubernetes construct used to group a set of similar objects together and classify them as user-defined resources. This translates to Trident no longer needing a dedicated etcd and a PV that it needs to use on the backend storage. All stateful objects used by Trident will be CRD objects that are present in the Kubernetes cluster's etcd.

Things to keep in mind about Trident's CRDs

1. When Trident is installed, a set of CRDs are created and can be used like any other resource type.
2. When [upgrading from a previous version of Trident](#) (one that used etcd to maintain state), the Trident installer will migrate data from the etcd key-value data store and create corresponding CRD objects.
3. [Downgrading](#) to a previous Trident version is not recommended.
4. When uninstalling Trident using the `tridentctl uninstall` command, Trident pods are deleted but the created CRDs will not be cleaned up. Refer to the [Uninstalling Guide](#) to understand how Trident can be completely removed and reconfigured from scratch.
5. Since the CRD objects that are used by Trident are stored in the Kubernetes cluster's etcd, [Trident disaster recovery workflows](#) will be different when compared to previous versions of Trident.

3.7.7 Trident Upgrade/Downgrade Process

Upgrading Trident

If you are looking to upgrade to the latest version of Trident, the [Upgrade section](#) provides a complete overview of the upgrade process.

Downgrading Trident

Downgrading to a previous release is not recommended. If you choose to downgrade, ensure that the PV used by the previous Trident installation is available.

Refer to the [Troubleshooting](#) section to understand what happens when a downgrade is attempted.

3.7.8 Recommendations for all deployments

Deploy Trident to a dedicated namespace

[Namespaces](#) provide administrative separation between different applications and are a barrier for resource sharing, for example, a PVC from one namespace cannot be consumed from another. Trident provides PV resources to all namespaces in the Kubernetes cluster and consequently leverages a service account which has elevated privileges.

Additionally, access to the Trident pod may enable a user to access storage system credentials and other sensitive information. It is important to ensure that application users and management applications do not have the ability to access the Trident object definitions or the pods themselves.

Use quotas and range limits to control storage consumption

Kubernetes has two features which, when combined, provide a powerful mechanism for limiting the resource consumption by applications. The [storage quota mechanism](#) allows the administrator to implement global, and storage class specific, capacity and object count consumption limits on a per-namespace basis. Further, using a [range limit](#) will ensure that the PVC requests must be within both a minimum and maximum value before the request is forwarded to the provisioner.

These values are defined on a per-namespace basis, which means that each namespace will need to have values defined which fall in line with their resource requirements. An example of [how to leverage quotas](#) can be found on [netapp.io](#).

3.7.9 Deploying Trident to OpenShift

OpenShift uses Kubernetes for the underlying container orchestrator. Consequently, the same recommendations will apply when using Trident with Kubernetes or OpenShift. However, there are some minor additions when using OpenShift which should be taken into consideration.

Deploy Trident to infrastructure nodes

Trident is a core service to the OpenShift cluster, provisioning and managing the volumes used across all projects. Consideration should be given to deploying Trident to the infrastructure nodes in order to provide the same level of care and concern.

To deploy Trident to the infrastructure nodes, the project for Trident must be created by an administrator using the `oc adm` command. This prevents the project from inheriting the default node selector, which forces the pod to execute on compute nodes.

```
# create the project which Trident will be deployed to using
# the non-default node selector
oc adm new-project <project_name> --node-selector="region=infra"

# deploy Trident using the project name
tridentctl install -n <project_name>
```

The result of the above command is that any pod deployed to the project will be scheduled to nodes which have the tag “region=infra”. This also removes the default node selector used by other projects which schedule pods to nodes which have the label “node-role.kubernetes.io/compute=true”.

3.8 Integrating Trident

3.8.1 Trident backend design

ONTAP

Choosing a backend driver for ONTAP

Four different backend drivers are available for ONTAP systems. These drivers are differentiated by the protocol being used and how the volumes are provisioned on the storage system. Therefore, give careful consideration regarding which driver to deploy.

At a higher level, if your application has components which need shared storage (multiple pods accessing the same PVC) NAS based drivers would be the default choice, while the block-based iSCSI driver meets the needs of non-shared storage. Choose the protocol based on the requirements of the application and the comfort level of the storage and infrastructure teams. Generally speaking, there is little difference between them for most applications, so often the decision is based upon whether or not shared storage (where more than one pod will need simultaneous access) is needed.

The four drivers for ONTAP backends are listed below:

- `ontap-nas` – Each PV provisioned is a full ONTAP FlexVolume
- `ontap-nas-economy` – Each PV provisioned is a qtree, with up to 200 qtrees per FlexVolume
- `ontap-nas-flexgroup` - Each PV provisioned as a full ONTAP FlexGroup, and all aggregates assigned to a SVM are used.
- `ontap-san` – Each PV provisioned is a LUN within its own FlexVolume

Choosing between the three NFS drivers has some ramifications to the features which are made available to the application.

Note that, in the tables below, not all of the capabilities are exposed through Trident. Some must be applied by the storage administrator after provisioning if that functionality is desired. The superscript footnotes distinguish the functionality per feature and driver.

Table 1: ONTAP NAS driver capabilities

ONTAP NFS Drivers	Snapshots	Clones	Multi-attach	QoS	Resize	Replication
<code>ontap-nas</code>	Yes ¹	Yes	Yes	Yes ¹	Yes	Yes ¹
<code>ontap-nas-economy</code>	Yes ¹²	Yes ¹²	Yes	Yes ¹²	Yes	Yes ¹²
<code>ontap-nas-flexgroup</code>	Yes ¹	No	Yes	Yes ¹	Yes	Yes ¹

The SAN driver capabilities are shown below.

Table 2: ONTAP SAN driver capabilities

ONTAP SAN Driver	Snapshots	Clones	Multi-attach	QoS	Resize	Replication
<code>ontap-san</code>	Yes ¹	Yes	Yes ³	Yes ¹	Yes	Yes ¹
<code>ontap-san-economy</code>	Yes ¹	Yes	Yes ³	Yes ¹²	Yes ¹	Yes ¹²

Footnote for the above tables:

Yes¹: Not Trident managed

Yes²: Trident managed, but not PV granular

Yes¹²: Not Trident managed and not PV granular

The features that are not PV granular are applied to the entire FlexVolume and all of the PVs (i.e. qtrees) will share a common schedule for each qtree.

As we can see in the above tables, much of the functionality between the `ontap-nas` and `ontap-nas-economy` is the same. However, since the `ontap-nas-economy` driver limits the ability to control the schedule at per-PV granularity, this may affect your disaster recovery and backup planning in particular. For development teams which desire to leverage PVC clone functionality on ONTAP storage, this is only possible when using the `ontap-nas` or `ontap-san` drivers (note, the `solidfire-san` driver is also capable of cloning PVCs).

Choosing a backend driver for Cloud Volumes ONTAP

Cloud Volumes ONTAP provides data control along with enterprise-class storage features for various use cases, including file shares and block-level storage serving NAS and SAN protocols (NFS, SMB / CIFS, and iSCSI). The compatible drivers for Cloud Volume ONTAP are `ontap-nas`, `ontap-nas-economy` or `ontap-san`. These are applicable for Cloud Volume ONTAP for AWS, Cloud Volume ONTAP for Azure, Cloud Volume ONTAP for GCP.

Element (HCI/SolidFire)

The `solidfire-san` driver used with the HCI/SolidFire platforms, helps the admin configure an Element backend for Trident on the basis of QoS limits. If you would like to design your backend to set the specific QoS limits on the volumes provisioned by Trident, use the `type` parameter in the backend file. The admin also can restrict the volume size that could be created on the storage using the `limitVolumeSize` parameter. Currently, Element OS storage features like volume resize and volume replication are not supported through the `solidfire-san` driver. These operations should be done manually through Element OS Web UI.

Table 3: SolidFire SAN driver capabilities

SolidFire Driver	Snapshots	Clones	Multi-attach	QoS	Resize	Replication
<code>solidfire-san</code>	Yes ¹	Yes	Yes ²	Yes	Yes	Yes ¹

Footnote:

Yes¹: Not Trident managed

SANtricity (E-Series)

To configure an E-Series backend for Trident, set the `storageDriverName` parameter to `eseries-iscsi` driver in the backend configuration. Once the E-Series backend has been configured, any requests to provision volume from the E-Series will be handled by Trident based on the host groups. Trident uses host groups to gain access to the LUNs that it provisions and by default, it looks for a host group named `trident` unless a different host group name is specified using the `accessGroupName` parameter in the backend configuration. The admin also can restrict the volume size that could be created on the storage using the `limitVolumeSize` parameter. Currently, E-Series storage features like volume resize and volume replication are not supported through the `eseries-iscsi` driver. These operations should be done manually through SANtricity System Manager.

Table 4: E-Series driver capabilities

E-Series Driver	Snapshots	Clones	Multi-attach	QoS	Resize	Replication
<code>eseries-iscsi</code>	Yes ¹	Yes ¹	No	No	Yes ¹	Yes ¹

Footnote:

Yes¹: Not Trident managed

Azure NetApp Files Backend Driver

Trident uses the `azure-netapp-files` driver to manage the [Azure NetApp Files](#) service.

More information about this driver and how to configure it can be found in Trident's [Azure NetApp Files backend documentation](#).

Table 5: Azure NetApp Files driver capabilities

Azure NetApp Files Driver	Snapshots	Clones	Multi-attach	QoS	Expand	Replication
<code>azure-netapp-files</code>	Yes ¹	Yes	Yes	Yes	Yes	Yes ¹

Footnote:

Yes¹: Not Trident managed

Cloud Volumes Service with AWS Backend Driver

Trident uses the `aws-cvs` driver to link with the Cloud Volumes Service on the AWS backend. To configure the AWS backend on Trident, you are required specify `apiRegion`, `apiURL`, `apiKey`, and the `secretKey` in the backend file. These values can be found in the CVS web portal in Account settings/API access. The supported service levels are aligned with CVS and include *standard*, *premium*, and *extreme*. More information on this driver may be found in the [Cloud Volumes Service for AWS Documentation](#). Currently, 100G is the minimum volume size that will be provisioned. Future releases of CVS may remove this restriction. Cloud Volume Service for AWS support fully-managed file services for both NFS, SMB, or dual protocol support.

Table 6: Cloud Volume Service driver capabilities

CVS for AWS Driver	Snapshots	Clones	Multi-attach	QoS	Expand	Replication
<code>aws-cvs</code>	Yes ¹	Yes	Yes	Yes	Yes	Yes ¹

Footnote:

Yes¹: Not Trident managed

The `aws-cvs` driver uses virtual storage pools. Virtual storage pools abstract the backend, letting Trident decide volume placement. The administrator defines the virtual storage pools in the `backend.json` file(s). Storage classes identify the virtual storage pools with the use of labels. More information on the virtual storage pools feature can be found in [Virtual Storage Pools Documentation](#).

3.8.2 Storage Class design

Individual Storage Classes need to be configured and applied to create a Kubernetes Storage Class object. This section discusses how to design a storage class for your application.

Storage Class design for specific backend utilization

Filtering can be used within a specific storage class object to determine which storage pool or set of pools are to be used with that specific storage class. Three sets of filters can be set in the Storage Class: *storagePools*, *additionalStoragePools*, and/or *excludeStoragePools*.

The *storagePools* parameter helps restrict storage to the set of pools that match any specified attributes. The *additionalStoragePools* parameter is used to extend the set of pools that Trident will use for provisioning along with the set of pools selected by the attributes and *storagePools* parameters. You can use either parameter alone or both together to make sure that the appropriate set of storage pools are selected.

The *excludeStoragePools* parameter is used to specifically exclude the listed set of pools that match the attributes.

Please refer to *Trident StorageClass Objects* on how these parameters are used.

Storage Class design to emulate QoS policies

If you would like to design Storage Classes to emulate Quality of Service policies, create a Storage Class with the *media* attribute as *hdd* or *ssd*. Based on the *media* attribute mentioned in the storage class, Trident will select the appropriate backend that serves *hdd* or *ssd* aggregates to match the media attribute and then direct the provisioning of the volumes on to the specific aggregate. Therefore we can create a storage class PREMIUM which would have *media* attribute set as *ssd* which could be classified as the PREMIUM QoS policy. We can create another storage class STANDARD which would have the media attribute set as 'hdd' which could be classified as the STANDARD QoS policy. We could also use the "IOPS" attribute in the storage class to redirect provisioning to an Element appliance which can be defined as a QoS Policy.

Please refer to *Trident StorageClass Objects* on how these parameters can be used.

Storage Class Design To utilize backend based on specific features

Storage Classes can be designed to direct volume provisioning on a specific backend where features such as thin and thick provisioning, snapshots, clones, and encryption are enabled. To specify which storage to use, create Storage Classes that specify the appropriate backend with the required feature enabled.

Please refer to *Trident StorageClass Objects* on how these parameters can be used.

Storage Class Design for Virtual Storage Pools

Virtual Storage Pools are available for Cloud Volumes Service for AWS, ANF, Element and E-Series backends.

Virtual Storage Pools allow an administrator to create a level of abstraction over backends which can be referenced through Storage Classes, for greater flexibility and efficient placement of volumes on backends. Different backends can be defined with the same class of service. Moreover, multiple Storage Pools can be created on the same backend but with different characteristics. When a Storage Class is configured with a selector with the specific labels, Trident chooses a backend which matches all the selector labels to place the volume. If the Storage Class selector labels matches multiple Storage Pools, Trident will choose one of them to provision the volume from.

Please refer to *Virtual Storage Pools* for more information and applicable parameters.

3.8.3 Virtual Storage Pool Design

While creating a backend, you can generally specify a set of parameters. It was impossible for the administrator to create another backend with the same storage credentials and with a different set of parameters. With the introduction of Virtual Storage Pools, this issue has been alleviated. Virtual Storage Pools is a level abstraction introduced between the backend and the Kubernetes Storage Class so that the administrator can define parameters along with labels which can be referenced through Kubernetes Storage Classes as a selector, in a backend-agnostic way. Currently Virtual Storage Pool is supported for Cloud Volumes Service for AWS, E-Series and SolidFire.

Design Virtual Storage Pools for emulating different Service Levels/QoS

It is possible to design Virtual Storage Pools for emulating service classes. Using the virtual pool implementation for Cloud Volume Service for AWS, let us examine how we can setup up different service classes. Configure the AWS-CVS backend with multiple labels, representing different performance levels. Set “servicelevel” aspect to the appropriate performance level and add other required aspects under each label. Now create different Kubernetes Storage Classes that would map to different virtual Storage Pools. Using the `parameters.selector` field, each StorageClass calls out which virtual pool(s) may be used to host a volume.

Design Virtual Pools for Assigning Specific Set of Aspects

Multiple Virtual Storage pools with a specific set of aspects can be designed from a single storage backend. For doing so, configure the backend with multiple labels and set the required aspects under each label. Now create different Kubernetes Storage Classes using the `parameters.selector` field that would map to different Virtual Storage Pools. The volumes that get provisioned on the backend will have the aspects defined in the chosen Virtual Storage Pool.

3.8.4 PVC characteristics which affect storage provisioning

Some parameters beyond the requested storage class may affect Trident’s provisioning decision process when creating a PVC.

Access mode

When requesting storage via a PVC, one of the mandatory fields is the access mode. The mode desired may affect the backend selected to host the storage request.

Trident will attempt to match the storage protocol used with the access method specified according to the following matrix. This is independent of the underlying storage platform.

Table 7: Protocols used by access modes

	ReadWriteOnce	ReadOnlyMany	ReadWriteMany
iSCSI	Yes	Yes	No
NFS	Yes	Yes	Yes

A request for a ReadWriteMany PVC submitted to a Trident deployment without an NFS backend configured will result in no volume being provisioned. For this reason, the requestor should use the access mode which is appropriate for their application.

3.8.5 Volume Operations

Modifying persistent volumes

Persistent volumes are, with two exceptions, immutable objects in Kubernetes. Once created, the reclaim policy and the size can be modified. However, this doesn't prevent some aspects of the volume from being modified outside of Kubernetes. This may be desirable in order to customize the volume for specific applications, to ensure that capacity is not accidentally consumed, or simply to move the volume to a different storage controller for any reason.

Note: Kubernetes in-tree provisioners do not support volume resize operations for NFS or iSCSI PVs at this time. Trident supports expanding NFS volumes. For a list of PV types which support volume resizing refer to the [Kubernetes documentation](#).

The connection details of the PV cannot be modified after creation.

Volume Move Operations

Storage administrators have the ability to move volumes between aggregates and controllers in the ONTAP cluster non-disruptively to the storage consumer. This operation does not affect Trident or the Kubernetes cluster, as long as the destination aggregate is one which the SVM Trident is using has access to. Importantly, if the aggregate has been newly added to the SVM, the backend will need to be “refreshed” by re-adding it to Trident. This will trigger Trident to reinventory the SVM so that the new aggregate is recognized.

However, moving volumes across backends is not supported automatically by Trident. This includes between SVMs in the same cluster, between clusters, or onto a different storage platform (even if that storage system is one which is connected to Trident).

If a volume is copied to another location, the *volume import feature* may be used to import current volumes into Trident.

Expanding volumes with ONTAP

To allow possible expansion later, set *allowVolumeExpansion* to *true* in your StorageClass associated with the volume. Whenever the Persistent Volume needs to be resized, edit the `spec.resources.requests.storage` annotation in the Persistent Volume Claim to the required volume size. Trident will automatically take care of resizing the volume on ONTAP.

Note:

1. Currently, Trident only supports NFS PV resize and not the iSCSI PV resize.
 2. Kubernetes, prior to version 1.12, does not support NFS PV resize as the admission controller may reject PVC size updates. The Trident team has changed Kubernetes to allow such changes starting with Kubernetes 1.12. While we recommend using Kubernetes 1.12, it is still possible to resize NFS PVs for earlier versions of Kubernetes that support resize. This is done by disabling the PersistentVolumeClaimResize admission plugin when the Kubernetes API server is started.
-

Import an existing volume into Kubernetes

Volume Import provides the ability to import an existing storage volume into a Kubernetes environment. This is currently supported by the `ontap-nas`, `ontap-nas-flexgroup`, `solidfire-san`, `azure-netapp-files`

and `aws-cvs` drivers. This feature is useful when porting an existing application into Kubernetes or during disaster recovery scenarios.

When using the ONTAP and `solidfire-san` drivers, use the command `tridentctl import volume <backend-name> <volume-name> -f /path/pvc.yaml` to import an existing volume into Kubernetes to be managed by Trident. The PVC YAML or JSON file used in the import volume command points to a storage class which identifies Trident as the provisioner. When using a HCI/SolidFire backend, ensure the volume names are unique. If the volume names are duplicated, clone the volume to a unique name so the volume import feature can distinguish between them.

If the `aws-cvs` driver is used, use the command `tridentctl import volume <backend-name> <volume path> -f /path/pvc.yaml` to import the volume into Kubernetes to be managed by Trident. This ensures a unique volume reference.

When the above command is executed, Trident will find the volume on the backend and read its size. It will automatically add (and overwrite if necessary) the configured PVC's volume size. Trident then creates the new PV and Kubernetes binds the PVC to the PV.

If a container was deployed such that it required the specific imported PVC, it would remain in a pending state until the PVC/PV pair are bound via the volume import process. After the PVC/PV pair are bound, the container should come up, provided there are no other issues.

For information, please see the [documentation](#).

3.8.6 Deploying OpenShift services using Trident

The OpenShift value-add cluster services provide important functionality to cluster administrators and the applications being hosted. The storage which these services use can be provisioned using the node-local resources, however, this often limits the capacity, performance, recoverability, and sustainability of the service. Leveraging an enterprise storage array to provide the capacity to these services can enable dramatically improved service, however, as with all applications, the OpenShift and storage administrators should work closely together to determine the best options for each. The Red Hat documentation should be leveraged heavily to determine the requirements and ensure that sizing and performance needs are met.

Registry service

Deploying and managing storage for the registry has been documented on [netapp.io](#) in [this blog post](#).

Logging service

Like other OpenShift services, the logging service is deployed using Ansible with configuration parameters supplied by the inventory file, a.k.a. hosts, provided to the playbook. There are two installation methods which will be covered: deploying logging during initial OpenShift install and deploying logging after OpenShift has been installed.

Warning: As of Red Hat OpenShift version 3.9, the official documentation recommends against NFS for the logging service due to concerns around data corruption. This is based on Red Hat testing of their products. ON-TAP's NFS server does not have these issues, and can easily back a logging deployment. Ultimately, the choice of protocol for the logging service is up to you, just know that both will work great when using NetApp platforms and there is no reason to avoid NFS if that is your preference.

If you choose to use NFS with the logging service, you will need to set the Ansible variable `openshift_enable_unsupported_configurations` to `true` to prevent the installer from failing.

Getting started

The logging service can, optionally, be deployed for both applications as well as for the core operations of the OpenShift cluster itself. If you choose to deploy operations logging, by specifying the variable `openshift_logging_use_ops` as `true`, two instances of the service will be created. The variables which control the logging instance for operations contain “ops” in them, whereas the instance for applications does not.

Configuring the Ansible variables according to the deployment method is important in order to ensure that the correct storage is utilized by the underlying services. Let’s look at the options for each of the deployment methods

Note: The tables below only contain the variables which are relevant for storage configuration as it relates to the logging service. There are many other options found in [the documentation](#) which should be reviewed, configured, and used according to your deployment.

The variables in the below table will result in the Ansible playbook creating a PV and PVC for the logging service using the details provided. This method is significantly less flexible than using the component installation playbook after OpenShift installation, however, if you have existing volumes available, it is an option.

Table 8: Logging variables when deploying at OpenShift install time

Variable	Details
<code>openshift_logging_storage</code>	Set to <code>true</code> to have the installer create an NFS PV for the logging service.
<code>openshift_logging_storage_hostname</code>	The hostname or IP address of the NFS host. This should be set to the data LIF for your virtual machine.
<code>openshift_logging_storage_mount_path</code>	The mount path for the NFS export. For example, if the volume is junctioned as <code>/openshift_logging</code> , you would use that path for this variable.
<code>openshift_logging_storage_name</code>	The name, e.g. <code>opense_logs</code> , of the PV to create.
<code>openshift_logging_storage_size</code>	The size of the NFS export, for example <code>100Gi</code> .

If your OpenShift cluster is already running, and therefore Trident has been deployed and configured, the installer can use dynamic provisioning to create the volumes. The following variables will need to be configured.

Table 9: Logging variables when deploying after OpenShift install

Variable	Details
<code>openshift_logging_es_pvc_dynamic</code>	Set to <code>true</code> to use dynamically provisioned volumes.
<code>openshift_logging_es_pvc_storage_class</code>	The name of the storage class which will be used in the PVC.
<code>openshift_logging_es_pvc_size</code>	The size of the volume requested in the PVC.
<code>openshift_logging_es_pvc_prefix</code>	A prefix for the PVCs used by the logging service.
<code>openshift_logging_es_ops_pvc_dynamic</code>	Set to <code>true</code> to use dynamically provisioned volumes for the ops logging instance.
<code>openshift_logging_es_ops_pvc_storage_class</code>	The name of the storage class for the ops logging instance.
<code>openshift_logging_es_ops_pvc_size</code>	The size of the volume request for the ops instance.
<code>openshift_logging_es_ops_pvc_prefix</code>	A prefix for the ops instance PVCs.

Note: A bug exists in OpenShift 3.9 which prevents a storage class from being used when the value for `openshift_logging_es_ops_pvc_dynamic` is set to `true`. However, this can be worked around by, counterintuitively, setting the variable to `false`, which will include the storage class in the PVC.

Deploy the logging stack

If you are deploying logging as a part of the initial OpenShift install process, then you only need to follow the standard deployment process. Ansible will configure and deploy the needed services and OpenShift objects so that the service

is available as soon as Ansible completes.

However, if you are deploying after the initial installation, the component playbook will need to be used by Ansible. This process may change slightly with different versions of OpenShift, so be sure to read and follow [the documentation](#) for your version.

Metrics service

The metrics service provides valuable information to the administrator regarding the status, resource utilization, and availability of the OpenShift cluster. It is also necessary for pod autoscale functionality and many organizations use data from the metrics service for their charge back and/or show back applications.

Like with the logging service, and OpenShift as a whole, Ansible is used to deploy the metrics service. Also, like the logging service, the metrics service can be deployed during an initial setup of the cluster or after it's operational using the component installation method. The following tables contain the variables which are important when configuring persistent storage for the metrics service.

Note: The tables below only contain the variables which are relevant for storage configuration as it relates to the metrics service. There are many other options found in the documentation which should be reviewed, configured, and used according to your deployment.

Table 10: Metrics variables when deploying at OpenShift install time

Variable	Details
openshift_metrics_storage	Set to <code>true</code> to have the installer create an NFS PV for the logging service.
openshift_metrics_storage	The hostname or IP address of the NFS host. This should be set to the data LIF for your SVM.
openshift_metrics_storage	The mount path for the NFS export. For example, if the volume is junctioned as <code>/openshift_metrics</code> , you would use that path for this variable.
openshift_metrics_storage	The name, <code>reg_name_metrics</code> , of the PV to create.
openshift_metrics_storage	The size of the NFS export, for example <code>100Gi</code> .

If your OpenShift cluster is already running, and therefore Trident has been deployed and configured, the installer can use dynamic provisioning to create the volumes. The following variables will need to be configured.

Table 11: Metrics variables when deploying after OpenShift install

Variable	Details
openshift_metrics_cassandra_pvc	A prefix to use for the metrics PVCs.
openshift_metrics_cassandra_pvc	The size of the volumes to request.
openshift_metrics_cassandra_pvc	The type of storage to use for metrics, this must be set to <code>dynamic</code> for Ansible to create PVCs with the appropriate storage class.
openshift_metrics_cassandra_pvc	The name of the storage class to use.

Deploying the metrics service

With the appropriate Ansible variables defined in your hosts/inventory file, deploy the service using Ansible. If you are deploying at OpenShift install time, then the PV will be created and used automatically. If you're deploying using the component playbooks, after OpenShift install, then Ansible will create any PVCs which are needed and, after Trident has provisioned storage for them, deploy the service.

The variables above, and the process for deploying, may change with each version of OpenShift. Ensure you review and follow [the deployment guide](#) for your version so that it is configured for your environment.

3.9 Backup and Disaster Recovery

Protecting application data is one of the fundamental purposes of any storage system. Whether an application is cloud native, 12 factor, a microservice, or any other architecture, the application data still needs to be protected.

NetApp's storage platforms provide data protection and recoverability options which vary based on recovery time and acceptable data loss requirements. Trident can provision volumes which can take advantage of some of these features, however, a full data protection and recovery strategy should be evaluated for each application with a persistence requirement.

3.9.1 Backing Up Kubernetes and Trident's state

Trident v19.07 and beyond will now utilize Kubernetes CRDs to store and manage its state. As a result, Trident will store its metadata in the Kubernetes cluster's etcd database. All Kubernetes objects are stored in the cluster's etcd. Periodically backing up the etcd cluster data is important to recover Kubernetes clusters under disaster scenarios.

This example provides a sample workflow to create etcd snapshots on a Kubernetes cluster using `etcdctl`.

etcdctl snapshot backup

The command `etcdctl snapshot save` enables us to take a point-in-time snapshot of the etcd cluster.

```
sudo docker run --rm -v /backup:/backup \
--network host \
-v /etc/kubernetes/pki/etcd:/etc/kubernetes/pki/etcd \
--env ETCDCCTL_API=3 \
k8s.gcr.io/etcd-amd64:3.2.18 \
etcdctl --endpoints=https://127.0.0.1:2379 \
--cacert=/etc/kubernetes/pki/etcd/ca.crt \
--cert=/etc/kubernetes/pki/etcd/healthcheck-client.crt \
--key=/etc/kubernetes/pki/etcd/healthcheck-client.key \
snapshot save /backup/etcd-snapshot.db
```

This command creates an etcd snapshot by spinning up an etcd container and saves it in the `/backup` directory.

etcdctl snapshot restore

In the event of a disaster, it is possible to spin up a Kubernetes cluster using the etcd snapshots. Use the `etcdctl snapshot restore` command to restore a specific snapshot taken to the `/var/lib/etcd` folder. After restoring, confirm if the `/var/lib/etcd` folder has been populated with the member folder. The following is an example of `etcdctl snapshot restore` command.

```
# etcdctl snapshot restore '/backup/etcd-snapshot-latest.db' ; mv /default.etcd/
↪member/ /var/lib/etcd/
```

Before you initialize the Kubernetes cluster, copy all the necessary certificates. Create the cluster with the `--ignore-preflight-errors=DirAvailable--var-lib-etcd` flag. After the cluster comes up make sure that the kube-system pods have started. Use the `kubectl get crd` command to verify if the custom resources created by Trident are present and retrieve Trident objects to make sure that all the data is available.

3.9.2 ONTAP Snapshots

The following sections talk in general about how ONTAP Snapshot technology can be used to take backups of the volume and how these snapshots can be restored. Snapshots play an important role by providing point-in-time recovery options for application data. However, snapshots are not backups by themselves, they will not protect against storage system failure or other catastrophes. But, they are a convenient, quick, and easy way to recover data in most scenarios.

Using ONTAP snapshots with containers

If the snapshot policy has not been defined in the backend, it defaults to using the `none` policy. This results in ONTAP taking no automatic snapshots. However, the storage administrator can take manual snapshots or change the snapshot policy via the ONTAP management interface. This will not affect Trident operation.

The snapshot directory is hidden by default. This helps facilitate maximum compatibility of volumes provisioned using the `ontap-nas` and `ontap-nas-economy` drivers.

Accessing the snapshot directory

Enable the `.snapshot` directory when using the `ontap-nas` and `ontap-nas-economy` drivers to allow applications to recover data from snapshots directly.

Restoring the snapshots

Restore a volume to a state recorded in a prior snapshot using the `volume snapshot restore` ONTAP CLI command. When you restore a Snapshot copy, the restore operation overwrites the existing volume configuration. Any changes made to the data in the volume after the Snapshot copy was created are lost.

```
cluster1::*> volume snapshot restore -vserver vs0 -volume vol3 -snapshot_
↳vol3_snap_archive
```

3.9.3 Data replication using ONTAP

Replicating data can play an important role in protecting against data loss due to storage array failure. Snapshots are a point-in-time recovery which provides a very quick and easy method of recovering data which has been corrupted or accidentally lost as a result of human or technological error. However, they cannot protect against catastrophic failure of the storage array itself.

ONTAP SnapMirror SVM Replication

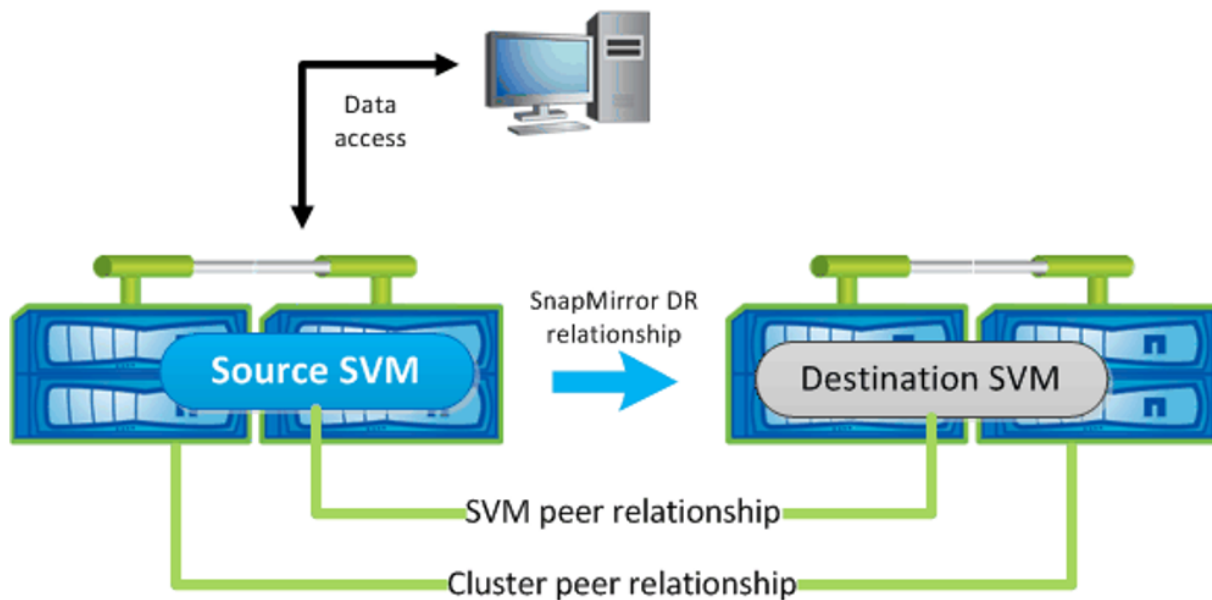
SnapMirror can be used to replicate a complete SVM which includes its configuration settings and its volumes. In the event of a disaster, SnapMirror destination SVM can be activated to start serving data and switch back to the primary when the systems are restored. Since Trident is unable to configure replication relationships itself, the storage administrator can use ONTAP's SnapMirror SVM Replication feature to automatically replicate volumes to a Disaster Recovery (DR) destination.

- A distinct backend should be created for each SVM which has SVM-DR enabled.
- Storage Classes should be crafted so as to not select the replicated backends except when desired. This is important to avoid having volumes which do not need the protection of a replication relationship to be provisioned onto the backend(s) that support SVM-DR.

- Application administrators should understand the additional cost and complexity associated with replicating the data and a plan for recovery should be determined before they leverage data replication.
- Before activating the SnapMirror destination SVM, stop all the scheduled SnapMirror transfers, abort all ongoing SnapMirror transfers, break the replication relationship, stop the source SVM, and then start the SnapMirror destination SVM.
- Trident does not automatically detect SVM failures. Therefore, upon a failure, the administrator needs to run the command `tridentctl backend update` to trigger Trident's failover to the new backend.

ONTAP SnapMirror SVM Replication Setup

- Set up peering between the Source and Destination Cluster and SVM.
- Setting up SnapMirror SVM replication involves creating the destination SVM by using the `-subtype dp-destination` option.
- Create a replication job schedule to make sure that replication happens in the required intervals.
- Create a SnapMirror replication from destination SVM to the source SVM using the `-identity-preserve true` option to make sure that source SVM configurations and source SVM interfaces are copied to the destination. From the destination SVM, initialize the SnapMirror SVM replication relationship.



SnapMirror SVM Replication Setup

SnapMirror SVM Disaster Recovery Workflow for Trident

The following steps describe how Trident and other containerized applications can resume functioning during a catastrophe using the SnapMirror SVM replication.

Disaster Recovery Workflow for Trident

Trident v19.07 and beyond will now utilize Kubernetes CRDs to store and manage its own state. It will use the Kubernetes cluster's etcd to store its metadata. Here we assume that the Kubernetes etcd data files and the certificates

are stored on NetApp FlexVolume. This FlexVolume resides in a SVM which has Snapmirror SVM DR relationship with a destination SVM at the secondary site. The following steps describe how we can recover a single master Kubernetes Cluster with Trident in the event of a disaster.

1. In the event of the source SVM failure, activate the SnapMirror destination SVM. Activating the destination SVM involves stopping scheduled SnapMirror transfers, aborting ongoing SnapMirror transfers, breaking the replication relationship, stopping the source SVM, and starting the destination SVM.
2. From the destination SVM, mount the volume which contains the Kubernetes etcd data files and certificates on to the host which will be setup as a master node.
3. Copy all the required certificates pertaining to the Kubernetes cluster under `/etc/kubernetes/pki` and the etcd member files under `/var/lib/etcd`.
4. Now create a Kubernetes cluster with the `kubeadm init` command along with the `--ignore-preflight-errors=DirAvailable--var-lib-etcd` flag. Please note that the hostnames used for the Kubernetes nodes must same as the source Kubernetes cluster.
5. Use the `kubectl get crd` command to verify if all the Trident custom resources have come up and retrieve Trident objects to make sure that all the data is available.
6. Update all the required backends to reflect the new destination SVM name using the `./tridentctl update backend <backend-name> -f <backend-json-file> -n <namespace>` command.

Disaster Recovery Workflow for Application Persistent Volumes

When the destination SVM is activated, all the volumes provisioned by Trident will start serving data. Once the Kubernetes cluster is setup on the destination side using the above mentioned procedure, all the deployments and pods are started and the containerized applications should run without any issues.

ONTAP SnapMirror Volume Replication

ONTAP SnapMirror Volume Replication is a disaster recovery feature which enables failover to destination storage from primary storage on a volume level. SnapMirror creates a volume replica or mirror of the primary storage on to the secondary storage by syncing snapshots.

ONTAP SnapMirror Volume Replication Setup


- The clusters in which the volumes reside and the SVMs that serve data from the volumes must be peered.
- Create a SnapMirror policy which controls the behavior of the relationship and specifies the configuration attributes for that relationship.
- Create a SnapMirror relationship between the destination volume and the source volume using the “snapmirror create” volume and assign the appropriate SnapMirror policy.
- After the SnapMirror relationship is created, initialize the relationship so that a baseline transfer from the source volume to the destination volume will be completed.

SnapMirror Volume Replication Setup

SnapMirror Volume Disaster Recovery Workflow for Trident

Disaster recovery using SnapMirror Volume Replication is not as seamless as the SnapMirror SVM Replication. The following steps describe how Trident and other applications can resume functioning during a catastrophe, from the secondary site .

Disaster Recovery Workflow for Trident



Trident v19.07 and beyond will now utilize Kubernetes CRDs to store and manage its own state. Trident will store its metadata in the Kubernetes cluster's etcd database. Here we assume that the Kubernetes etcd data files and the certificates are stored on NetApp FlexVolume which is SnapMirrored to the destination volume at the secondary site. The following steps describe how we can recover a single master Kubernetes Cluster with Trident.

1. In the event of a disaster, stop all the scheduled SnapMirror transfers and abort all ongoing SnapMirror transfers. Break the replication relationship between the destination and source volumes so that the destination volume becomes Read/Write.
2. From the destination SVM, mount the volume which contains the Kubernetes etcd data files and certificates on to the host which will be setup as a master node.
3. Copy all the required certificates pertaining to the Kubernetes cluster under `/etc/kubernetes/pki` and the etcd member files under `/var/lib/etcd`.
4. Now create a Kubernetes cluster with the `kubeadm init` command along with the `--ignore-preflight-errors=DirAvailable--var-lib-etcd` flag. Please note that the hostnames must same as the source Kubernetes cluster.
5. Use the `kubect1 get crd` command to verify if all the Trident custom resources have come up and retrieve Trident objects to make sure that all the data is available.
6. Clean up the previous backends and create new backends on Trident. Specify the new Management and Data LIF, new SVM name and password of the destination SVM.

Disaster Recovery Workflow for Application Persistent Volumes

In this section, let us examine how SnapMirror destination volumes can be made available for containerized workloads in the event of a disaster.

1. Stop all the scheduled SnapMirror transfers and abort all ongoing SnapMirror transfers. Break the replication relationship between the destination and source volume so that the destination volume becomes Read/Write. Clean up the deployments which were consuming PVC bound to volumes on the source SVM.
2. Once the Kubernetes cluster is setup on the destination side using the above mentioned procedure, clean up the deployments, PVCs and PV, from the Kubernetes cluster.

3. Create new backends on Trident by specifying the new Management and Data LIF, new SVM name and password of the destination SVM.
4. Now import the required volumes as a PV bound to a new PVC using the Trident import feature.
5. Re-deploy the application deployments with the newly created PVCs.

3.9.4 ElementOS snapshots

Backup data on an Element volume by setting a snapshot schedule for the volume, ensuring the snapshots are taken at the required intervals. Currently, it is not possible to set a snapshot schedule to a volume through the `solidfire-san` driver. Set it using the Element OS Web UI or Element OS APIs.

In the event of data corruption, we can choose a particular snapshot and rollback the volume to the snapshot manually using the Element OS Web UI or Element OS APIs. This reverts any changes made to the volume since the snapshot was created.

The *Creating Snapshots of Persistent Volumes* section details a complete workflow for creating Volume Snapshots and then using them to create PVCs.

3.10 Security Recommendations

3.10.1 Run Trident in its own namespace

It is important to prevent applications, application admins, users, and management applications from accessing Trident object definitions or the pods to ensure reliable storage and block potential malicious activity. To separate out the other applications and users from Trident, always install Trident in its own Kubernetes namespace. In our Installing Trident docs we call this namespace *trident*. Putting Trident in its own namespace assures that only the Kubernetes administrative personnel have access to the Trident pod and the artifacts (such as backend and CHAP secrets if applicable) stored in Trident's etcd datastore. Allow only administrators access to the trident namespace and thus access to *tridentctl* application.

3.10.2 CHAP authentication

NetApp recommends deploying bi-directional CHAP to ensure authentication between a host and the HCI/SolidFire backend. Trident uses a secret object that includes two CHAP passwords per tenant. Kubernetes manages the mapping of Kubernetes tenant to SF tenant. Upon volume creation time, Trident makes an API call to the HCI/SolidFire system to retrieve the secrets if the secret for that tenant doesn't already exist. Trident then passes the secrets on to Kubernetes. The kubelet located on each node accesses the secrets via the Kubernetes API and uses them to run/enable CHAP between each node accessing the volume and the HCI/SolidFire system where the volumes are located.

Since Trident v18.10, the `solidfire-san` driver defaults to use CHAP if the Kubernetes version is ≥ 1.7 and a Trident access group doesn't exist. Setting *AccessGroup* or *UseCHAP* in the backend configuration file overrides this behavior. CHAP is guaranteed by setting *UseCHAP* to `true` in the `backend.json` file.

3.11 Frequently Asked Questions

This section of the Design and Architecture Guide is divided into 3 areas and covers frequently asked questions for each:

1. *Trident for Kubernetes Installation*

2. *Trident Backend Configuration and Use*
3. *Trident Upgrade, Support, Licensing, and Troubleshooting*

3.11.1 Trident for Kubernetes Installation

This section covers Trident Installation on a Kubernetes cluster.

What are the supported versions of etcd?

Trident v19.07 does not require an etcd. It uses CRDs to maintain state.

Does Trident support an offline install from a private registry?

Yes, Trident can be installed offline.

Refer to the *Offline install* section for a step by step procedure.

Can Trident be installed remotely?

Trident v18.10 and above supports *remote install capability* from any machine that has kubectl access to the cluster. After kubectl access is verified (e.g. initiate a *kubectl get nodes* command from the remote machine to verify), follow the installation instructions.

Refer to *Deploying* for more information on how to install Trident.

Can we configure High Availability with Trident?

Since Trident is installed as a Kubernetes Deployment (ReplicaSet) with one instance, it has HA built in. Do not increase the number of replicas in the Trident deployment.

If the node where Trident is installed is lost or the pod is otherwise inaccessible, Kubernetes will automatically re-deploy the pod to a healthy node in your cluster.

Since Trident is control-plane only, currently mounted pods will not be affected if Trident is re-deployed.

Does Trident need access to kube-system namespace?

Trident reads from the Kubernetes API Server to determine when applications request new PVCs so it needs access to kube-system.

What are the roles and privileges used by Trident?

The Trident installer creates a Kubernetes ClusterRole which has specific access to the cluster's PersistentVolume, PersistentVolumeClaim, StorageClass and Secret resources of the Kubernetes cluster.

Refer to *Customized Installation* for more information.

Can we locally generate the exact manifest files Trident uses to install?

You can locally generate and modify the exact manifest files Trident uses to install if needed.

Refer to *Customized Installation* for instructions.

Can we share the same ONTAP backend SVM for two separate Trident instances for two separate Kubernetes clusters?

Although it is not advised, you can use the same backend SVM for 2 Trident instances. Specify a unique Trident volume name for each Trident instance during installation and/or specify a unique StoragePrefix parameter in the setup/backend.json file. This is to ensure the same FlexVol isn't used for both instances.

Refer to: *Customized Installation* for information on specifying a unique Trident volume name. Refer to: *Global Configuration* for information on creating a unique StoragePrefix.

Is it possible to install Trident under ContainerLinux (formerly CoreOS)?

Trident is simply a Kubernetes pod and can be installed wherever Kubernetes is running.

Refer to *Supported host operating systems* for more information.

Can we use Trident with NetApp Cloud Volumes ONTAP?

Yes, it is supported on AWS, Google Cloud and Azure.

Refer to *Supported backends* for more information.

Does Trident work with Cloud Volumes Services?

Yes, Trident supports the Azure NetApp Files service in Azure as well as the Cloud Volumes Service in AWS.

Refer to *Supported backends* for more information.

What versions of Kubernetes support Trident as an enhanced CSI Provisioner?

Kubernetes versions 1.13 and above support running Trident as a CSI Provisioner. Before installing Trident, ensure the required *feature gates* are enabled.

Refer to *Requirements* for a list of supported orchestrators.

Why should I install Trident to work as a CSI Provisioner?

With the 19.07 release, Trident now fully adheres to the latest CSI 1.1 specification and is production ready. This enables users to experience all features exposed by the current release, as well as future releases. Trident can continue to fix issues or add features without touching the Kubernetes core, while also absorbing any standardized future changes or features efficiently.

How do I install Trident to work as a CSI Provisioner?

The installation procedure is detailed under the *Deployment* section. Ensure that the *feature gates* are enabled.

How does Trident maintain state if it doesn't use etcd?

Trident uses *Custom Resource Definitions(CRDs)* to maintain its state. This eliminates the requirement for etcd and a Trident volume on the storage cluster. Trident no longer needs its separate PV; the information is stored as CRD objects that will be present in the Kubernetes cluster's etcd.

How do I uninstall Trident?

The *Uninstalling Trident* section explains how you can remove Trident.

3.11.2 Trident Backend Configuration and Use

This section covers Trident backend definition file configurations and use.

Do we need to define both Management and Data LIFs in an ONTAP backend definition file?

NetApp recommends having both in the backend definition file. However, the Management LIF is the only one that is absolutely mandatory.

Refer to *ONTAP (AFF/FAS/Select/Cloud)* for more information on backend definition files.

Can we specify a port in the DataLIF?

Trident 19.01 and later supports specifying a port in the DataLIF.

Configure it in the backend.json file as "*managementLIF*": *<ip address>*:*<port>*" For example, if the IP address of your management LIF is 192.0.2.1, and the port is 1000, configure "*managementLIF*": "192.0.2.1:1000",

Is it possible to update the Management LIF on the backend ?

Yes, it is possible to update the backend Management LIF using the `tridentctl update backend` command.

Refer to *Backend configuration* for more information on updating the backend.

Is it possible to update the Data LIF on the backend ?

No, it is not possible to update the Data LIF on the backend.

Can we create multiple backends in Trident for Kubernetes?

Trident can support many backends simultaneously, either with the same driver or different drivers.

Refer to *Backend configuration* for more information on creating backend definition files.

How does Trident store backend credentials?

Trident stores the backend credentials as Kubernetes Secrets.

How does Trident select a specific backend?

If the backend attributes cannot be used to automatically select the right pools for a class, the *storagePools* and *additionalStoragePools* parameters are used to select a specific set of pools.

Refer to *Storage Class design for specific backend utilization* in the Design and Architecture Guide for more information.

Can we make sure Trident will not provision from a specific backend?

The *excludeStoragePools* parameter is used to filter the set of pools that Trident will use for provisioning and will remove any pools that match.

Refer to *Kubernetes StorageClass Objects*

If there are multiple backends of the same kind, how does Trident select which backend to use?

If there are multiple backends configured of the same type, then Trident will select the appropriate backend based on the parameters present in the StorageClass and the PersistentVolumeClaim. For example, if there are multiple `ontap-nas` driver backends, then Trident will try to match parameters in the StorageClass and PersistentVolumeClaim combined and match a backend which can deliver the requirements listed in the StorageClass and PersistentVolumeClaim. If there are multiple backends that matches the request, then Trident will choose from one of them at random.

Does Trident support bi-directional CHAP with Element/SolidFire?

Bi-directional CHAP is supported with Element.

Refer to *CHAP authentication* in the Design and Architecture Guide for additional information.

How does Trident deploy Qtrees on an ONTAP volume? How many Qtrees can be deployed on a single volume through Trident?

The `ontap-nas-economy` driver will create up to 200 Qtrees in the same FlexVol, 100,000 Qtrees per cluster node, and 2.4M per cluster. When you enter a new PersistentVolumeClaim that is serviced by the economy driver, the driver looks to see if a FlexVol already exists that can service the new Qtree. If the FlexVol does not exist that can service the Qtree, a new FlexVol will be created.

Refer to *Choosing a driver* for more information.

How can we set Unix permissions for volumes provisioned on ONTAP NAS?

Unix Permissions can be set on the volume provisioned by Trident by setting a parameter in the backend definition file.

Refer to *ONTAP (AFF/FAS/Select/Cloud)* for more information.

How can we configure an explicit set of ONTAP NFS mount options while provisioning a volume?

By default, Trident does not set mount options to any value with Kubernetes.

To specify the mount options in the Kubernetes Storage Class, please follow the example given [here](#).

How do I set the provisioned volumes to a specific export policy?

To allow the appropriate hosts access to a volume, use the *exportPolicy* parameter configured in the backend definition file.

Refer to *ONTAP (AFF/FAS/Select/Cloud)* for more information.

How do I set volume encryption through Trident with ONTAP?

Encryption can be set on the volume provisioned by Trident by using the *encryption* parameter in the backend definition file.

Refer to *ONTAP (AFF/FAS/Select/Cloud)* for more information.

What is the best way to implement QoS for ONTAP through Trident?

Use *StorageClasses* to implement QoS for ONTAP.

Refer to *Storage Class design to emulate QoS policies* for more information.

How do we specify thin or thick provisioning through Trident?

The ONTAP drivers support either thin or thick provisioning. E-series only support thick provisioning. Element OS backends only support thin provisioning.

The ONTAP drivers default to thin provisioning. If thick provisioning is desired, you may configure either the backend definition file or the *StorageClass*. If both are configured, the *StorageClass* takes precedence. Configure the following for ONTAP:

- On the *StorageClass*, set the *provisioningType* attribute as *thick*.
- On the backend definition file, enable thick volumes by setting backend *spaceReserve* parameter as *volume*.

Refer to *ONTAP (AFF/FAS/Select/Cloud)* for more information.

How do I make sure that the volumes being used are not deleted even if I accidentally delete the PVC?

PVC protection is automatically enabled on Kubernetes starting from version 1.10.

Refer to *Storage Object in Use Protection* for additional information.

Can we use PVC resize functionality with NFS, Trident, and ONTAP?

PVC resize is supported with Trident. Note that *volume autogrow* is an ONTAP feature that is not applicable to Trident.

Refer to *Expanding Volumes* for more information.

If I have a volume that was created outside Trident can I import it into Trident?

Starting in Trident v19.04, you can use the volume import feature to bring volumes in to Kubernetes.

Refer to *Importing a volume* for more information.

Can I import a volume while it is in Snapmirror Data Protection (DP) or offline mode?

The volume import will fail if the external volume is in DP mode or offline. You will receive an error message.

```
Error: could not import volume: volume import failed to get size of volume: volume  
↪ <name> was not found (400 Bad Request) command terminated with exit code 1.
```

Make sure to remove the DP mode or put the volume online before importing the volume.

Refer to: *Behavior of Drivers for Volume Import* for additional information.

Can we use PVC resize functionality with iSCSI, Trident, and ONTAP?

PVC resize functionality with iSCSI is not supported with Trident.

How is resource quota translated to a NetApp cluster?

Kubernetes Storage Resource Quota should work as long as NetApp Storage has capacity. When the NetApp storage cannot honor the Kubernetes quota settings due to lack of capacity, Trident will try to provision but will error out.

Can you create Volume Snapshots using Trident?

The creation of Kubernetes Volume Snapshots is possible with the 20.01 release and above. This requires Kubernetes 1.17 and above.

Refer to the 20.01 Trident documentation: [On-Demand Volume Snapshots](#) for more information.

How do we take a snapshot backup of a volume provisioned by Trident with ONTAP?

This is available on `ontap-nas`, `ontap-san`, and `ontap-nas-flexgroup` drivers.

This is also available on the `ontap-nas-economy` drivers but on the FlexVol level granularity and not on the qtree level granularity.

To enable the ability to snapshot volumes provisioned by Trident, set the backend parameter option `snapshotPolicy` to the desired snapshot policy as defined on the ONTAP backend. Any snapshots taken by the storage controller will not be known by Trident.

Can we set a snapshot reserve percentage for a volume provisioned through Trident?

Yes, we can reserve a specific percentage of disk space for storing the snapshot copies through Trident by setting the `snapshotReserve` attribute in the backend definition file. If you have configured the `snapshotPolicy` and the `snapshotReserve` option in the backend definition file, then snapshot reserve percentage will be set according to the `snapshotReserve` percentage mentioned in the backend file. If the `snapshotReserve` percentage number is not mentioned, then ONTAP by default will take the snapshot reserve percentage as 5. In the case where the `snapshotPolicy` option is set to none, then the snapshot reserve percentage is set to 0.

Refer to: *ONTAP (AFF/FAS/Select/Cloud)* for more information.

Can we directly access the volume snapshot directory and copy files?

Yes, It is possible to access the snapshot directory on the volume provisioned by Trident by setting the *snapshotDir* parameter in the backend definition file.

Refer to: *ONTAP (AFF/FAS/Select/Cloud)* for more information.

Can we set up SnapMirror for Trident volumes through Trident?

Currently, SnapMirror has be set externally using ONTAP CLI or OnCommand System Manager.

How do I restore Persistent Volumes to a specific ONTAP snapshot?

To restore a volume to an ONTAP snapshot, follow the following steps:

- Quiesce the application pod which is using the Persistent volume .
- Revert to the required snapshot through ONTAP CLI or OnCommand System Manager.
- Restart the application pod.

How can I obtain complete Trident configuration details?

`tridentctl get` command provides more information on the Trident Configuration.

Refer to *tridentctl get* for more information on this command.

How can we separate out storage class usage for each customer/tenant?

Kubernetes does not allow storage classes in namespaces. However, we can use Kubernetes to limit usage of a specific storage class per namespace by using [Storage Resource Quotas](#) which are per namespace. To deny a specific namespace access to specific storage, set the resource quota to 0 for that storage class.

Does Trident provide insight into the capacity of the storage?

This is out of scope for Trident. NetApp offers [Cloud Insights](#) for monitoring and analysis.

Does the user experience change when using Trident as a CSI Provisioner?

No. From the user's point of view, there are no changes as far as the user experience and functionalities are concerned. The provisioner name used will be `csi.trident.netapp.io`. This method of installing Trident is recommended to use all new features provided by current and future releases.

How do I design a Disaster Workflow for Trident v19.07?

The *Data replication using ONTAP* section talks about backup and DR workflows using ONTAP.

3.11.3 Trident Upgrade, Support, Licensing, and Troubleshooting

This section covers upgrading Trident, Trident Support, Licensing and Troubleshooting.

How frequently is Trident released?

Trident is released every 3 months: January, April, July and October. This is one month after a Kubernetes release.

Does NetApp support Trident?

Although Trident is open source and provided for free, NetApp fully supports Trident provided your NetApp backend is supported.

How do I raise a support case for Trident?

To raise a support case, you could do the following

- Customers can reach their Support Account Manager and get help to raise a ticket.
- Raise a support case by contacting support at mysupport.netapp.com.

How do I generate a support log bundle using Trident?

You can create a support bundle by running `tridentctl logs -a`. In addition to the logs captured in the bundle, capture the kubelet log to diagnose the mount problems on the k8s side. The instructions to get the kubelet log varies based on how k8s is installed.

Refer to: *Troubleshooting*.

Does Trident support all the features that are released in a particular version of Kubernetes?

Trident usually doesn't support alpha features in Kubernetes. We may support beta features within the following two Trident releases after the Kubernetes beta release.

What do I do if I need to raise a request for a new feature on Trident?

If you would like to raise a request for a new feature, raise an issue at NetApp/Trident [Github](#) and make sure to mention "RFE" in the subject and description of the issue.

Where do I raise a defect for Trident?

If you would like to raise a defect against Trident, raise an issue at NetApp/Trident [Github](#). Make sure to include all the necessary information and logs pertaining to the issue.

What happens if I have quick question on Trident that I need clarification on? Is there a community or a forum for Trident?

If you have any questions, issues, or requests please reach out to us through our [Slack](#) team or [GitHub](#).

Does Trident have any dependencies on other NetApp products for its functioning?

Trident doesn't have any dependencies on other NetApp software products and it works as a standalone application. However, you must have a NetApp backend storage device.

Can I upgrade from an older version of Trident directly to a newer version (skipping a few versions)?

We support upgrading directly from a version up to one year back. For example, if you are currently on v18.04, v18.07, or v19.01, we will support directly upgrading to v19.04. We suggest testing upgrading in a lab prior to production deployment. Information on upgrading Trident can be obtained [here](#).

How can I upgrade to the most recent version of Trident?

Refer to *Upgrading Trident* for the steps involved in Upgrading Trident to the latest release.

Is it possible to downgrade Trident to a previous release?

Downgrading Trident is not recommended for the *following reasons*.

If the Trident pod is destroyed, will we lose the data?

No data will be lost if the Trident pod is destroyed. Trident's metadata will be stored in CRD objects. All PVs that have been provisioned by Trident will function normally.

My storage system's password has changed and Trident no longer works, how do I recover?

Update the backend's password with a `tridentctl update backend myBackend -f </path/to_new_backend.json> -n trident`. Replace *myBackend* in the example with your backend name, and */path/to_new_backend.json* with the path to the correct backend.json file.

Trident for Docker provides direct integration with the Docker ecosystem for NetApp's ONTAP, SolidFire, and E-Series storage platforms, as well as the Cloud Volumes Service in AWS. It supports the provisioning and management of storage resources from the storage platform to Docker hosts, with a robust framework for adding additional platforms in the future.

Multiple instances of Trident can run concurrently on the same host. This allows simultaneous connections to multiple storage systems and storage types, with the ability to customize the storage used for the Docker volume(s).

4.1 Deploying

1. Verify that your deployment meets all of the *requirements*.
2. Ensure that you have a supported version of Docker installed.

```
docker --version
```

If your version is out of date, [follow the instructions for your distribution](#) to install or update.

3. Verify that the protocol prerequisites are installed and configured on your host. See *Host Configuration*.
4. Create a configuration file. The default location is `/etc/netappdvp/config.json`. Be sure to use the correct options for your storage system.

```
# create a location for the config files
sudo mkdir -p /etc/netappdvp

# create the configuration file, see below for more configuration examples
cat << EOF > /etc/netappdvp/config.json
{
  "version": 1,
  "storageDriverName": "ontap-nas",
  "managementLIF": "10.0.0.1",
  "dataLIF": "10.0.0.2",
```

(continues on next page)

(continued from previous page)

```
"svm": "svm_nfs",
"username": "vsadmin",
"password": "secret",
"aggregate": "aggr1"
}
EOF
```

5. Start Trident using the managed plugin system.

```
docker plugin install netapp/trident-plugin:19.07 --alias netapp --grant-all-
↳permissions
```

6. Begin using Trident to consume storage from the configured system.

```
# create a volume named "firstVolume"
docker volume create -d netapp --name firstVolume

# create a default volume at container instantiation
docker run --rm -it --volume-driver netapp --volume secondVolume:/my_vol alpine_
↳ash

# remove the volume "firstVolume"
docker volume rm firstVolume
```

4.2 Host and storage configuration

4.2.1 Host Configuration

NFS

Install the following system packages:

- RHEL / CentOS

```
sudo yum install -y nfs-utils
```

- Ubuntu / Debian

```
sudo apt-get install -y nfs-common
```

iSCSI

- RHEL / CentOS

1. Install the following system packages:

```
sudo yum install -y lsscsi iscsi-initiator-utils sg3_utils device-mapper-
↳multipath
```

2. Start the multipathing daemon:

```
sudo mpathconf --enable --with_multipathd y
```

3. Ensure that *iscsid* and *multipathd* are enabled and running:

```
sudo systemctl enable iscsid multipathd
sudo systemctl start iscsid multipathd
```

4. Discover the iSCSI targets:

```
sudo iscsiadm -m discoverydb -t st -p <DATA_LIF_IP> --discover
```

5. Login to the discovered iSCSI targets:

```
sudo iscsiadm -m node -p <DATA_LIF_IP> --login
```

6. Start and enable *iscsi*:

```
sudo systemctl enable iscsi
sudo systemctl start iscsi
```

- Ubuntu / Debian

1. Install the following system packages:

```
sudo apt-get install -y open-iscsi lsscsi sg3-utils multipath-tools scsitools
```

2. Enable multipathing:

```
sudo tee /etc/multipath.conf <<-'EOF'
defaults {
    user_friendly_names yes
    find_multipaths yes
}
EOF
sudo service multipath-tools restart
```

3. Ensure that *iscsid* and *multipathd* are running:

```
sudo service open-iscsi start
sudo service multipath-tools start
```

4. Discover the iSCSI targets:

```
sudo iscsiadm -m discoverydb -t st -p <DATA_LIF_IP> --discover
```

5. Login to the discovered iSCSI targets:

```
sudo iscsiadm -m node -p <DATA_LIF_IP> --login
```

Traditional Install Method (Docker <= 1.12)

1. Ensure you have Docker version 1.10 or above

```
docker --version
```

If your version is out of date, update to the latest.

```
curl -fsSL https://get.docker.com/ | sh
```

Or, follow the instructions for your distribution.

2. After ensuring the correct version of Docker is installed, install and configure the NetApp Docker Volume Plugin. Note, you will need to ensure that NFS and/or iSCSI is configured for your system. See the installation instructions below for detailed information on how to do this.

```
# download and unpack the application
wget https://github.com/NetApp/trident/releases/download/v19.07.2/trident-
↳installer-19.07.2.tar.gz
tar xzf trident-installer-19.07.2.tar.gz

# move to a location in the bin path
sudo mv trident-installer/extras/bin/trident /usr/local/bin
sudo chown root:root /usr/local/bin/trident
sudo chmod 755 /usr/local/bin/trident

# create a location for the config files
sudo mkdir -p /etc/netappdvp

# create the configuration file, see below for more configuration examples
cat << EOF > /etc/netappdvp/ontap-nas.json
{
  "version": 1,
  "storageDriverName": "ontap-nas",
  "managementLIF": "10.0.0.1",
  "dataLIF": "10.0.0.2",
  "svm": "svm_nfs",
  "username": "vsadmin",
  "password": "secret",
  "aggregate": "aggr1"
}
EOF
```

3. After placing the binary and creating the configuration file(s), start the Trident daemon using the desired configuration file.

Note: Unless specified, the default name for the volume driver will be “netapp”.

```
sudo trident --config=/etc/netappdvp/ontap-nas.json
```

4. Once the daemon is started, create and manage volumes using the Docker CLI interface.

```
docker volume create -d netapp --name trident_1
```

Provision Docker volume when starting a container:

```
docker run --rm -it --volume-driver netapp --volume trident_2:/my_vol alpine ash
```

Destroy docker volume:

```
docker volume rm trident_1
docker volume rm trident_2
```

Starting Trident at System Startup

A sample unit file for systemd based systems can be found at `contrib/trident.service.example` in the git repo. To use the file, with CentOS/RHEL:

```
# copy the file to the correct location. you must use unique names for the
# unit files if you have more than one instance running
cp contrib/trident.service.example /usr/lib/systemd/system/trident.service

# edit the file, change the description (line 2) to match the driver name and the
# configuration file path (line 9) to reflect your environment.

# reload systemd for it to ingest changes
systemctl daemon-reload

# enable the service, note this name will change depending on what you named the
# file in the /usr/lib/systemd/system directory
systemctl enable trident

# start the service, see note above about service name
systemctl start trident

# view the status
systemctl status trident
```

Note that anytime the unit file is modified you will need to issue the command `systemctl daemon-reload` for it to be aware of the changes.

Docker Managed Plugin Method (Docker >= 1.13 / 17.03)

Note: If you have used Trident pre-1.13/17.03 in the traditional daemon method, please ensure that you stop the Trident process and restart your Docker daemon before using the managed plugin method.

```
# stop all running instances
pkill /usr/local/bin/netappdvp
pkill /usr/local/bin/trident

# restart docker
systemctl restart docker
```

Trident Specific Plugin Startup Options

- `config` - Specify the configuration file the plugin will use. Only the file name should be specified, e.g. `gold.json`, the location must be `/etc/netappdvp` on the host system. The default is `config.json`.
- `log-level` - Specify the logging level (`debug`, `info`, `warn`, `error`, `fatal`). The default is `info`.
- `debug` - Specify whether debug logging is enabled. Default is `false`. Overrides `log-level` if `true`.

Installing the Managed Plugin

1. Ensure you have Docker Engine 17.03 (nee 1.13) or above installed.

```
docker --version
```

If your version is out of date, follow the instructions for your distribution to install or update.

2. Create a configuration file. The config file must be located in the `/etc/netappdvp` directory. The default filename is `config.json`, however you can use any name you choose by specifying the `config` option with the file name. Be sure to use the correct options for your storage system.

```
# create a location for the config files
sudo mkdir -p /etc/netappdvp

# create the configuration file, see below for more configuration examples
cat << EOF > /etc/netappdvp/config.json
{
  "version": 1,
  "storageDriverName": "ontap-nas",
  "managementLIF": "10.0.0.1",
  "dataLIF": "10.0.0.2",
  "svm": "svm_nfs",
  "username": "vsadmin",
  "password": "secret",
  "aggregate": "aggr1"
}
EOF
```

3. Start Trident using the managed plugin system.

```
docker plugin install --grant-all-permissions --alias netapp netapp/trident-
↪plugin:19.07 config=myConfigFile.json
```

4. Begin using Trident to consume storage from the configured system.

```
# create a volume named "firstVolume"
docker volume create -d netapp --name firstVolume

# create a default volume at container instantiation
docker run --rm -it --volume-driver netapp --volume secondVolume:/my_vol alpine_
↪ash

# remove the volume "firstVolume"
docker volume rm firstVolume
```

4.2.2 Global Configuration

These configuration variables apply to all Trident configurations, regardless of the storage platform being used.

Option	Description	Example
version	Config file version number	1
storageDriverName	ontap-nas, ontap-san, ontap-nas-economy, ontap-nas-flexgroup, eseries-iscsi, solidfire-san, azure-netapp-files or aws-cvs.	ontap-nas
storagePrefix	Optional prefix for volume names. Default: "netappdvp_".	staging_
limitVolumeSize	Optional restriction on volume sizes. Default: "" (not enforced)	10g

Also, default option settings are available to avoid having to specify them on every volume create. The `size` option is available for all controller types. See the ONTAP config section for an example of how to set the default volume size.

Defaults Option	Description	Example
size	Optional default size for new volumes. Default: "1G"	10G

Storage Prefix

Do not use a `storagePrefix` (including the default) for Element backends. By default the `solidfire-san` driver will ignore this setting and not use a prefix. We recommend using either a specific `tenantID` for docker volume mapping or using the attribute data which is populated with the docker version, driver info and raw name from docker in cases where any name munging may have been used.

A note of caution: `docker volume rm` will *delete* these volumes just as it does volumes created by the plugin using the default prefix. Be very careful when using pre-existing volumes!

4.2.3 ONTAP Configuration

User Permissions

Trident does not need full permissions on the ONTAP cluster and should not be used with the cluster-level admin account. Below are the ONTAP CLI comands to create a dedicated user for Trident with specific permissions.

```
# create a new Trident role
security login role create -vserver [VSERVER] -role trident_role -cmddirname DEFAULT -
↳access none

# grant common Trident permissions
security login role create -vserver [VSERVER] -role trident_role -cmddirname "event_
↳generate-autosupport-log" -access all
security login role create -vserver [VSERVER] -role trident_role -cmddirname "network_
↳interface" -access readonly
security login role create -vserver [VSERVER] -role trident_role -cmddirname "version
↳" -access readonly
security login role create -vserver [VSERVER] -role trident_role -cmddirname "vserver
↳" -access readonly
```

(continues on next page)

(continued from previous page)

```

security login role create -vserver [VSERVER] -role trident_role -cmddirname "vserver_
↪nfs show" -access readonly
security login role create -vserver [VSERVER] -role trident_role -cmddirname "volume"
↪-access all
security login role create -vserver [VSERVER] -role trident_role -cmddirname
↪"snapmirror" -access all

# grant ontap-san Trident permissions
security login role create -vserver [VSERVER] -role trident_role -cmddirname "vserver_
↪iscsi show" -access readonly
security login role create -vserver [VSERVER] -role trident_role -cmddirname "lun" -
↪access all

# grant ontap-nas-economy Trident permissions
security login role create -vserver [VSERVER] -role trident_role -cmddirname "vserver_
↪export-policy create" -access all
security login role create -vserver [VSERVER] -role trident_role -cmddirname "vserver_
↪export-policy rule create" -access all

# create a new Trident user with Trident role
security login create -vserver [VSERVER] -username trident_user -role trident_role -
↪application ontapi -authmethod password
    
```

Configuration File Options

In addition to the global configuration values above, when using ONTAP these top level options are available.

Option	Description	Example
managementLIF	IP address of ONTAP management LIF	10.0.0.1
dataLIF	IP address of protocol LIF; will be derived if not specified	10.0.0.2
svm	Storage virtual machine to use (req. if management LIF is a cluster LIF)	svm_nfs
username	Username to connect to the storage device	vsadmin
password	Password to connect to the storage device	secret
aggregate	Aggregate for provisioning (optional; if set, must be assigned to the SVM)	aggr1
limitAggregateUsage	Fail provisioning if usage is above this percentage (optional)	75%

A fully-qualified domain name (FQDN) can be specified for the managementLIF option. For the ontap-nas* drivers only, a FQDN may also be specified for the dataLIF option, in which case the FQDN will be used for the NFS mount operations. For the ontap-san driver, the default is to use all data LIF IPs from the SVM and to use iSCSI multipath. Specifying an IP address for the dataLIF for the ontap-san driver forces the driver to disable multipath and use only the specified address.

For the ontap-nas and ontap-nas-economy drivers, an additional top level option is available. For NFS host configuration, see also: <http://www.netapp.com/us/media/tr-4067.pdf>

For the ontap-nas-flexgroup driver, the aggregate option in the configuration file is ignored. All aggregates assigned to the SVM are used to provision a FlexGroup Volume.

Option	Description	Example
nfsMountOptions	Fine grained control of NFS mount options; defaults to “-o nfsvers=3”	-o nfsvers=4

For the `ontap-san` driver, an additional top level option is available to specify an `igroup`.

Option	Description	Example
<code>igroupName</code>	The <code>igroup</code> used by the plugin; defaults to “ <code>netappdvp</code> ”	<code>myigroup</code>

For the `ontap-nas-economy` driver, the `limitVolumeSize` option will additionally limit the size of the FlexVols that it creates.

Option	Description	Example
<code>limitVolumeSize</code>	Maximum requestable volume size and <code>qtree</code> parent volume size	<code>300g</code>

Also, when using ONTAP, these default option settings are available to avoid having to specify them on every volume create.

Defaults Option	Description	Example
<code>spaceReserve</code>	Space reservation mode; “ <code>none</code> ” (thin provisioned) or “ <code>volume</code> ” (thick)	<code>none</code>
<code>snapshotPolicy</code>	Snapshot policy to use, default is “ <code>none</code> ”	<code>none</code>
<code>snapshotReserve</code>	Snapshot reserve percentage, default is “” to accept ONTAP’s default	<code>10</code>
<code>splitOnClone</code>	Split a clone from its parent upon creation, defaults to “ <code>false</code> ”	<code>false</code>
<code>encryption</code>	Enable NetApp Volume Encryption, defaults to “ <code>false</code> ”	<code>true</code>
<code>unixPermissions</code>	NAS option for provisioned NFS volumes, defaults to “ <code>777</code> ”	<code>777</code>
<code>snapshotDir</code>	NAS option for access to the <code>.snapshot</code> directory, defaults to “ <code>false</code> ”	<code>false</code>
<code>exportPolicy</code>	NAS option for the NFS export policy to use, defaults to “ <code>default</code> ”	<code>default</code>
<code>securityStyle</code>	NAS option for access to the provisioned NFS volume, defaults to “ <code>unix</code> ”	<code>mixed</code>
<code>fileSystemType</code>	SAN option to select the file system type, defaults to “ <code>ext4</code> ”	<code>xf</code>

Scaling Options

The `ontap-nas` and `ontap-san` drivers create an ONTAP FlexVol for each Docker volume. ONTAP supports up to 1000 FlexVols per cluster node with a cluster maximum of 12,000 FlexVols. If your Docker volume requirements fit within that limitation, the `ontap-nas` driver is the preferred NAS solution due to the additional features offered by FlexVols such as Docker-volume-granular snapshots and cloning.

If you need more Docker volumes than may be accommodated by the FlexVol limits, choose the `ontap-nas-economy` driver, which creates Docker volumes as ONTAP Qtrees within a pool of automatically managed FlexVols. Qtrees offer far greater scaling, up to 100,000 per cluster node and 2,400,000 per cluster, at the expense of some features. The `ontap-nas-economy` driver does not support Docker-volume-granular snapshots or cloning. The `ontap-nas-economy` driver is not currently supported in Docker Swarm, as Swarm does not orchestrate volume creation across multiple nodes.

Choose the `ontap-nas-flexgroup` driver to increase parallelism to a single volume that can grow into the petabyte range with billions of files. Some ideal use cases for FlexGroups include AI/ML/DL, big data and analytics, software builds, streaming, file repositories, etc. Trident uses all aggregates assigned to an SVM when provisioning a FlexGroup Volume. FlexGroup support in Trident also has the following considerations:

- Requires ONTAP version 9.2 or greater.
- As of this writing, FlexGroups only support NFS v3.
- Recommended to enable the 64-bit NFSv3 identifiers for the SVM.
- The minimum recommended FlexGroup size is 100GB.
- Cloning is not supported for FlexGroup Volumes.

For information regarding FlexGroups and workloads that are appropriate for FlexGroups see the [NetApp FlexGroup Volume - Best Practices and Implementation Guide](#).

To get advanced features and huge scale in the same environment, you can run multiple instances of the Docker Volume Plugin, with one using `ontap-nas` and another using `ontap-nas-economy`.

Example ONTAP Config Files

NFS Example for `ontap-nas` driver

```
{
  "version": 1,
  "storageDriverName": "ontap-nas",
  "managementLIF": "10.0.0.1",
  "dataLIF": "10.0.0.2",
  "svm": "svm_nfs",
  "username": "vsadmin",
  "password": "secret",
  "aggregate": "aggr1",
  "defaults": {
    "size": "10G",
    "spaceReserve": "none",
    "exportPolicy": "default"
  }
}
```

NFS Example for `ontap-nas-flexgroup` driver

```
{
  "version": 1,
  "storageDriverName": "ontap-nas-flexgroup",
  "managementLIF": "10.0.0.1",
  "dataLIF": "10.0.0.2",
  "svm": "svm_nfs",
  "username": "vsadmin",
  "password": "secret",
  "defaults": {
    "size": "100G",
    "spaceReserve": "none",
    "exportPolicy": "default"
  }
}
```

NFS Example for `ontap-nas-economy` driver

```
{
  "version": 1,
  "storageDriverName": "ontap-nas-economy",
  "managementLIF": "10.0.0.1",
  "dataLIF": "10.0.0.2",
  "svm": "svm_nfs",
  "username": "vsadmin",
  "password": "secret",
  "aggregate": "aggr1"
}
```

iSCSI Example for `ontap-san` driver

```

{
  "version": 1,
  "storageDriverName": "ontap-san",
  "managementLIF": "10.0.0.1",
  "dataLIF": "10.0.0.3",
  "svm": "svm_iscsi",
  "username": "vsadmin",
  "password": "secret",
  "aggregate": "aggr1",
  "igroupName": "myigroup"
}

```

4.2.4 Element OS/SolidFire Configuration

In addition to the global configuration values above, when using SolidFire, these options are available.

Option	Description	Example
Endpoint	Ex. <code>https://<login>:<password>@<mvip>/json-rpc/<element-version></code>	
SVIP	iSCSI IP address and port	10.0.0.7:3260
TenantName	SF Tenant to use (created if not found)	"docker"
InitiatorIFace	Specify interface when restricting iSCSI traffic to non-default interface	"default"
Types	QoS specifications	See below
LegacyNamePrefix	Prefix for upgraded Trident installs	"netappdvp-"

The SolidFire driver does not support Docker Swarm.

LegacyNamePrefix If you used a version of Trident prior to 1.3.2 and perform an upgrade with existing volumes, you'll need to set this value in order to access your old volumes that were mapped via the `volume-name` method.

Example Solidfire Config File

```

{
  "version": 1,
  "storageDriverName": "solidfire-san",
  "Endpoint": "https://admin:admin@192.168.160.3/json-rpc/8.0",
  "SVIP": "10.0.0.7:3260",
  "TenantName": "docker",
  "InitiatorIFace": "default",
  "Types": [
    {
      "Type": "Bronze",
      "Qos": {
        "minIOPS": 1000,
        "maxIOPS": 2000,
        "burstIOPS": 4000
      }
    },
    {
      "Type": "Silver",
      "Qos": {
        "minIOPS": 4000,

```

(continues on next page)

(continued from previous page)

```

        "maxIOPS": 6000,
        "burstIOPS": 8000
    },
    },
    {
        "Type": "Gold",
        "Qos": {
            "minIOPS": 6000,
            "maxIOPS": 8000,
            "burstIOPS": 10000
        }
    }
}
]
}

```

4.2.5 E-Series Configuration

In addition to the global configuration values above, when using E-Series, these options are available.

Option	Description	Example
webProxyHostname	Hostname or IP address of Web Services Proxy	localhost
webProxyPort	Port number of the Web Services Proxy (optional)	8443
webProxyUseHTTP	Use HTTP instead of HTTPS for Web Services Proxy (default = false)	true
webProxyVerifyTLS	Verify server's certificate chain and hostname (default = false)	true
username	Username for Web Services Proxy	rw
password	Password for Web Services Proxy	rw
controllerA	IP address of controller A	10.0.0.5
controllerB	IP address of controller B	10.0.0.6
passwordArray	Password for storage array if set	blank/empty
hostDataIP	Host iSCSI IP address (if multipathing just choose either one)	10.0.0.101
poolNameSearchPattern	Regular expression for matching storage pools available for Trident volumes (default = .+)	docker.*
hostType	Type of E-series Host created by Trident (default = linux_dm_mp)	linux_dm_mp
accessGroupName	Name of E-series Host Group to contain Hosts defined by Trident (default = netappdvp)	Docker-Hosts

Example E-Series Config File

Example for eseries-iscsi driver

```

{
  "version": 1,
  "storageDriverName": "eseries-iscsi",
  "webProxyHostname": "localhost",
  "webProxyPort": "8443",
  "webProxyUseHTTP": false,
  "webProxyVerifyTLS": true,
  "username": "rw",
  "password": "rw",
  "controllerA": "10.0.0.5",
  "controllerB": "10.0.0.6",

```

(continues on next page)

(continued from previous page)

```

"passwordArray": "",
"hostDataIP": "10.0.0.101"
}

```

E-Series Array Setup Notes

The E-Series Docker driver can provision Docker volumes in any storage pool on the array, including volume groups and DDP pools. To limit the Docker driver to a subset of the storage pools, set the `poolNameSearchPattern` in the configuration file to a regular expression that matches the desired pools.

When creating a docker volume you can specify the volume size as well as the disk media type using the `-o` option and the tags `size` and `mediaType`. Valid values for media type are `hdd` and `ssd`. Note that these are optional; if unspecified, the defaults will be a *1 GB* volume allocated from an *HDD pool*. An example of using these tags to create a 2 GiB volume from an SSD-based pool:

```
docker volume create -d netapp --name my_vol -o size=2G -o mediaType=ssd
```

The E-series Docker driver will detect and use any preexisting Host definitions without modification, and the driver will automatically define Host and Host Group objects as needed. The host type for hosts created by the driver defaults to `linux_dm_mp`, the native DM-MPIO multipath driver in Linux.

The current E-series Docker driver only supports iSCSI.

4.2.6 Cloud Volumes Service (CVS) on AWS Configuration

Note: The NetApp Cloud Volumes Service for AWS does not support volumes less than 100 GB in size. To make it easier to deploy applications, Trident automatically creates 100 GB volumes if a smaller volume is requested. Future releases of the Cloud Volumes Service may remove this restriction.

In addition to the global configuration values above, when using CVS on AWS, these options are available. The required values are all available in the CVS web user interface.

Option	Description	Example
<code>apiRegion</code>	CVS account region (required)	“us-east-1”
<code>apiURL</code>	CVS account API URL (required)	“https://cvs-aws-bundles.netapp.com:8080/v1”
<code>apiKey</code>	CVS account API key (required)	
<code>secretKey</code>	CVS account secret key (required)	
<code>nfsMountOptions</code>	NFS mount options; defaults to “-o nfsvers=3”	“vers=3,proto=tcp,timeo=600”
<code>serviceLevel</code>	Performance level (standard, premium, extreme), defaults to “standard”	“premium”

The required values `apiRegion`, `apiURL`, `apiKey`, and `secretKey` may be found in the CVS web portal in Account settings / API access.

Also, when using CVS on AWS, these default volume option settings are available.

Defaults Option	Description	Example
exportRule	NFS access list (addresses and/or CIDR subnets), defaults to "0.0.0.0/0"	"10.0.1.0/24,10.0.2.100"
snapshotReserve	Snapshot reserve percentage, default is "" to accept CVS default of 0	"10"
size	Volume size, defaults to "100GB"	"500G"

Example CVS on AWS Config File

```
{
  "version": 1,
  "storageDriverName": "aws-cvs",
  "apiRegion": "us-east-1",
  "apiURL": "https://cvs-aws-bundles.netapp.com:8080/v1",
  "apiKey": "znHczZsrrtHisIsAbOguSaPIKeyAZNchRAGzLzZE",
  "secretKey": "rR0rUmWXfNioN1KhtHisIsAnoTherboGuskey6pU",
  "region": "us-east-1",
  "serviceLevel": "premium",
  "storagePrefix": "cvs-",
  "limitVolumeSize": "200Gi",
  "defaults": {
    "snapshotReserve": "5",
    "exportRule": "10.0.0.0/24,10.0.1.0/24,10.0.2.100",
    "size": "100Gi"
  }
}
```

4.2.7 Azure NetApp Files Configuration

Note: The Azure NetApp Files service does not support volumes less than 100 GB in size. To make it easier to deploy applications, Trident automatically creates 100 GB volumes if a smaller volume is requested.

Preparation

To configure and use an [Azure NetApp Files](#) backend, you will need:

- `subscriptionID` from an Azure subscription with Azure NetApp Files enabled
- `tenantID`, `clientID`, and `clientSecret` from an [App Registration](#) in Azure Active Directory with sufficient permissions to the Azure NetApp Files service
- Azure `location` that contains at least one [delegated subnet](#)

If you're using Azure NetApp Files for the first time or in a new location, some initial configuration is required that the [quickstart guide](#) will walk you through.

Backend configuration options

Parameter	Description	Default
version	Always 1	
storageDriver-Name	“azure-netapp-files”	
backendName	Custom name for the storage backend	Driver name + “_” + random characters
subscriptionID	The subscription ID from your Azure subscription	
tenantID	The tenant ID from an App Registration	
clientID	The client ID from an App Registration	
clientSecret	The client secret from an App Registration	
serviceLevel	One of “Standard”, “Premium” or “Ultra”	“” (random)
location	Name of the Azure location new volumes will be created in	“” (random)
virtualNetwork	Name of a virtual network with a delegated subnet	“” (random)
subnet	Name of a subnet delegated to Microsoft .Netapp/ volumes	“” (random)
nfsMountOptions	Fine-grained control of NFS mount options	“-o nfsvers=3”
limitVolumeSize	Fail provisioning if requested volume size is above this value	“” (not enforced by default)

You can control how each volume is provisioned by default using these options in a special section of the configuration. For an example, see the configuration examples below.

Parameter	Description	Default
exportRule	The export rule(s) for new volumes	“0.0.0.0/0”
size	The default size of new volumes	“100G”

The `exportRule` value must be a comma-separated list of any combination of IPv4 addresses or IPv4 subnets in CIDR notation.

Example configurations

Example 1 - Minimal backend configuration for azure-netapp-files

This is the absolute minimum backend configuration. With this Trident will discover all of your NetApp accounts, capacity pools, and subnets delegated to ANF in every location worldwide, and place new volumes on one of them randomly.

This configuration is useful when you’re just getting started with ANF and trying things out, but in practice you’re going to want to provide additional scoping for the volumes you provision in order to make sure that they have the characteristics you want and end up on a network that’s close to the compute that’s using it. See the subsequent examples for more details.

```
{
  "version": 1,
  "storageDriverName": "azure-netapp-files",
  "subscriptionID": "9f87c765-4774-fake-ae98-a721add45451",
  "tenantID": "68e4f836-edc1-fake-bff9-b2d865ee56cf",
  "clientID": "dd043f63-bf8e-fake-8076-8de91e5713aa",
```

(continues on next page)

(continued from previous page)

```

"clientSecret": "SECRET"
}

```

Example 2 - Single location and specific service level for azure-netapp-files

This backend configuration will place volumes in Azure’s “eastus” location in a “Premium” capacity pool. Trident automatically discovers all of the subnets delegated to ANF in that location and will place a new volume on one of them randomly.

```

{
  "version": 1,
  "storageDriverName": "azure-netapp-files",
  "subscriptionID": "9f87c765-4774-fake-ae98-a721add45451",
  "tenantID": "68e4f836-edc1-fake-bff9-b2d865ee56cf",
  "clientID": "dd043f63-bf8e-fake-8076-8de91e5713aa",
  "clientSecret": "SECRET",
  "location": "eastus",
  "serviceLevel": "Premium"
}

```

Example 3 - Advanced configuration for azure-netapp-files

This backend configuration further reduces the scope of volume placement to a single subnet, and also modifies some volume provisioning defaults.

```

{
  "version": 1,
  "storageDriverName": "azure-netapp-files",
  "subscriptionID": "9f87c765-4774-fake-ae98-a721add45451",
  "tenantID": "68e4f836-edc1-fake-bff9-b2d865ee56cf",
  "clientID": "dd043f63-bf8e-fake-8076-8de91e5713aa",
  "clientSecret": "SECRET",
  "location": "eastus",
  "serviceLevel": "Premium",
  "virtualNetwork": "my-virtual-network",
  "subnet": "my-subnet",
  "nfsMountOptions": "vers=3,proto=tcp,timeo=600",
  "limitVolumeSize": "500Gi",
  "defaults": {
    "exportRule": "10.0.0.0/24,10.0.1.0/24,10.0.2.100",
    "size": "200Gi"
  }
}

```

Example 4 - Virtual storage pools with azure-netapp-files

This backend configuration defines multiple *pools of storage* in a single file. This is useful when you have multiple capacity pools supporting different service levels and you want to create storage classes in Kubernetes that represent those.

This is just scratching the surface of the power of virtual storage pools and their labels.

```

{
  "version": 1,
  "storageDriverName": "azure-netapp-files",
  "subscriptionID": "9f87c765-4774-fake-ae98-a721add45451",
  "tenantID": "68e4f836-edc1-fake-bff9-b2d865ee56cf",

```

(continues on next page)

(continued from previous page)

```

"clientID": "dd043f63-bf8e-fake-8076-8de91e5713aa",
"clientSecret": "SECRET",
"nfsMountOptions": "vers=3,proto=tcp,timeo=600",
"labels": {
  "cloud": "azure"
},
"location": "eastus",

"storage": [
  {
    "labels": {
      "performance": "gold"
    },
    "serviceLevel": "Ultra"
  },
  {
    "labels": {
      "performance": "silver"
    },
    "serviceLevel": "Premium"
  },
  {
    "labels": {
      "performance": "bronze"
    },
    "serviceLevel": "Standard",
  }
]
}

```

The following StorageClass definitions refer to the storage pools above. Using the `parameters.selector` field, each StorageClass calls out which pool may be used to host a volume. The volume will have the aspects defined in the chosen pool.

```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: gold
provisioner: csi.trident.netapp.io
parameters:
  selector: "performance=gold"
allowVolumeExpansion: true
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: silver
provisioner: csi.trident.netapp.io
parameters:
  selector: "performance=silver"
allowVolumeExpansion: true
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: bronze

```

(continues on next page)

```
provisioner: csi.trident.netapp.io
parameters:
  selector: "performance=bronze"
allowVolumeExpansion: true
```

4.2.8 Multiple Instances of Trident

Multiple instances of Trident are needed when you desire to have multiple storage configurations available simultaneously. The key to multiple instances is to give them different names using the `--alias` option with the containerized plugin, or `--volume-driver` option when instantiating Trident on the host.

Docker Managed Plugin (Docker >= 1.13 / 17.03)

1. Launch the first instance specifying an alias and configuration file

```
docker plugin install --grant-all-permissions --alias silver netapp/trident-
↪plugin:19.07 config=silver.json
```

2. Launch the second instance, specifying a different alias and configuration file

```
docker plugin install --grant-all-permissions --alias gold netapp/trident-
↪plugin:19.07 config=gold.json
```

3. Create volumes specifying the alias as the driver name

```
# gold volume
docker volume create -d gold --name ntapGold

# silver volume
docker volume create -d silver --name ntapSilver
```

Traditional (Docker <=1.12)

1. Launch the plugin with an NFS configuration using a custom driver ID:

```
sudo trident --volume-driver=netapp-nas --config=/path/to/config-nfs.json
```

2. Launch the plugin with an iSCSI configuration using a custom driver ID:

```
sudo trident --volume-driver=netapp-san --config=/path/to/config-iscsi.
↪json
```

3. Provision Docker volumes each driver instance:

- NFS

```
docker volume create -d netapp-nas --name my_nfs_vol
```

- iSCSI

```
docker volume create -d netapp-san --name my_iscsi_vol
```

4.3 Common tasks

4.3.1 Managing Trident

Installing Trident

Follow the extensive *deployment* guide.

Updating Trident

The plugin is not in the data path, therefore you can safely upgrade it without any impact to volumes that are in use. As with any plugin, during the upgrade process there will be a brief period where 'docker volume' commands directed at the plugin will not succeed, and applications will be unable to mount volumes until the plugin is running again. Under most circumstances, this is a matter of seconds.

1. List the existing volumes:

```
docker volume ls
DRIVER          VOLUME NAME
netapp:latest   my_volume
```

2. Disable the plugin:

```
docker plugin disable -f netapp:latest
docker plugin ls
ID                NAME          DESCRIPTION
↔ENABLED
7067f39a5df5     netapp:latest nDVP - NetApp Docker Volume Plugin  false
```

3. Upgrade the plugin:

```
docker plugin upgrade --skip-remote-check --grant-all-permissions netapp:latest
↔netapp/trident-plugin:19.07
```

Note: The 18.01 release of Trident replaces the nDVP. You should upgrade directly from the netapp/ndvp-plugin image to the netapp/trident-plugin image.

4. Enable the plugin:

```
docker plugin enable netapp:latest
```

5. Verify that the plugin is enabled:

```
docker plugin ls
ID                NAME          DESCRIPTION
↔ENABLED
7067f39a5df5     netapp:latest Trident - NetApp Docker Volume Plugin  true
```

6. Verify that the volumes are visible:

```
docker volume ls
DRIVER          VOLUME NAME
netapp:latest   my_volume
```

Uninstalling Trident

Note: Do not uninstall the plugin in order to upgrade it, as Docker may become confused as to which plugin owns any existing volumes; use the upgrade instructions in the previous section instead.

1. Remove any volumes that the plugin created.
2. Disable the plugin:

```
docker plugin disable netapp:latest
docker plugin ls
ID                NAME                DESCRIPTION                └
↔ENABLED
7067f39a5df5     netapp:latest       nDVP - NetApp Docker Volume Plugin  false
```

3. Remove the plugin:

```
docker plugin rm netapp:latest
```

4.3.2 Managing volumes

Creating and consuming storage from ONTAP, SolidFire, E-Series systems and the Cloud Volumes Service is easy with Trident. Simply use the standard `docker volume` commands with the Trident driver name specified when needed.

Create a Volume

```
# create a volume with a Trident driver using the default name
docker volume create -d netapp --name firstVolume

# create a volume with a specific Trident instance
docker volume create -d ntap_bronze --name bronzeVolume
```

If no options are specified, the defaults for the driver are used. The defaults are documented on the page for the storage driver you're using below.

The default volume size may be overridden per volume as follows:

```
# create a 20GiB volume with a Trident driver
docker volume create -d netapp --name my_vol --opt size=20G
```

Volume sizes are expressed as strings containing an integer value with optional units (e.g. “10G”, “20GB”, “3TiB”). If no units are specified, the default is ‘G’. Size units may be expressed either as powers of 2 (B, KiB, MiB, GiB, TiB) or powers of 10 (B, KB, MB, GB, TB). Shorthand units use powers of 2 (G = GiB, T = TiB, ...).

Volume Driver CLI Options

Each storage driver has a different set of options which can be provided at volume creation time to customize the outcome. Refer to the documentation below for your configured storage system to determine which options apply.

ONTAP Volume Options

Volume create options for both NFS and iSCSI:

- `size` - the size of the volume, defaults to 1 GiB
- `spaceReserve` - thin or thick provision the volume, defaults to `thin`. Valid values are `none` (thin provisioned) and `volume` (thick provisioned).
- `snapshotPolicy` - this will set the snapshot policy to the desired value. The default is `none`, meaning no snapshots will automatically be created for the volume. Unless modified by your storage administrator, a policy named “default” exists on all ONTAP systems which creates and retains six hourly, two daily, and two weekly snapshots. The data preserved in a snapshot can be recovered by browsing to the `.snapshot` directory in any directory in the volume.
- `snapshotReserve` - this will set the snapshot reserve to the desired percentage. The default is no value, meaning ONTAP will select the `snapshotReserve` (usually 5%) if you have selected a `snapshotPolicy`, or 0% if the `snapshotPolicy` is `none`. The default `snapshotReserve` value may be set in the config file for all ONTAP backends, and it may be used as a volume creation option for all ONTAP backends except `ontap-nas-economy`.
- `splitOnClone` - when cloning a volume, this will cause ONTAP to immediately split the clone from its parent. The default is `false`. Some use cases for cloning volumes are best served by splitting the clone from its parent immediately upon creation, since there is unlikely to be any opportunity for storage efficiencies. For example, cloning an empty database can offer large time savings but little storage savings, so it’s best to split the clone immediately.
- `encryption` - this will enable NetApp Volume Encryption (NVE) on the new volume, defaults to `false`. NVE must be licensed and enabled on the cluster to use this option.

NFS has additional options that aren’t relevant when using iSCSI:

- `unixPermissions` - this controls the permission set for the volume itself. By default the permissions will be set to `---rwxr-xr-x`, or in numerical notation `0755`, and root will be the owner. Either the text or numerical format will work.
- `snapshotDir` - setting this to `true` will make the `.snapshot` directory visible to clients accessing the volume. The default value is `false`, meaning that access to snapshot data is disabled by default. Some images, for example the official MySQL image, don’t function as expected when the `.snapshot` directory is visible.
- `exportPolicy` - sets the export policy to be used for the volume. The default is `default`.
- `securityStyle` - sets the security style to be used for access to the volume. The default is `unix`. Valid values are `unix` and `mixed`.

iSCSI has an additional option that isn’t relevant when using NFS:

- `fileSystemType` - sets the file system used to format iSCSI volumes. The default is `ext4`. Valid values are `ext3`, `ext4`, and `zfs`.
- `spaceAllocation` - setting this to `false` will turn off the LUN’s space-allocation feature. The default value is `true`, meaning ONTAP notifies the host when the volume has run out of space and the LUN in the volume cannot accept writes. This option also enables ONTAP to reclaim space automatically when your host deletes data.

Using these options during the docker volume create operation is super simple, just provide the option and the value using the `-o` operator during the CLI operation. These override any equivalent values from the JSON configuration file.

```
# create a 10GiB volume
docker volume create -d netapp --name demo -o size=10G -o encryption=true
```

(continues on next page)

(continued from previous page)

```
# create a 100GiB volume with snapshots
docker volume create -d netapp --name demo -o size=100G -o snapshotPolicy=default -o ↵
↳snapshotReserve=10

# create a volume which has the setUID bit enabled
docker volume create -d netapp --name demo -o unixPermissions=4755
```

The minimum volume size is 20MiB.

If the snapshot reserve is not specified and the snapshot policy is 'none', Trident will use a snapshot reserve of 0%.

```
# create a volume with no snapshot policy and no snapshot reserve
docker volume create -d netapp --name my_vol --opt snapshotPolicy=none

# create a volume with no snapshot policy and a custom snapshot reserve of 10%
docker volume create -d netapp --name my_vol --opt snapshotPolicy=none --opt ↵
↳snapshotReserve=10

# create a volume with a snapshot policy and a custom snapshot reserve of 10%
docker volume create -d netapp --name my_vol --opt snapshotPolicy=myPolicy --opt ↵
↳snapshotReserve=10

# create a volume with a snapshot policy, and accept ONTAP's default snapshot reserve ↵
↳ (usually 5%)
docker volume create -d netapp --name my_vol --opt snapshotPolicy=myPolicy
```

Element OS/SolidFire Volume Options

The SolidFire driver options expose the size and quality of service (QoS) policies associated with the volume. When the volume is created, the QoS policy associated with it is specified using the `-o type=service_level` nomenclature.

The first step to defining a QoS service level with the SolidFire driver is to create at least one type and specify the minimum, maximum, and burst IOPS associated with a name in the configuration file.

Example Configuration File with QoS Definitions

```
{
  "...": "...",
  "Types": [
    {
      "Type": "Bronze",
      "Qos": {
        "minIOPS": 1000,
        "maxIOPS": 2000,
        "burstIOPS": 4000
      }
    },
    {
      "Type": "Silver",
      "Qos": {
        "minIOPS": 4000,
        "maxIOPS": 6000,
        "burstIOPS": 8000
      }
    }
  ],
}
```

(continues on next page)

(continued from previous page)

```

    {
      "Type": "Gold",
      "Qos": {
        "minIOPS": 6000,
        "maxIOPS": 8000,
        "burstIOPS": 10000
      }
    }
  ]
}

```

In the above configuration we have three policy definitions: *Bronze*, *Silver*, and *Gold*. These names are well known and fairly common, but we could have just as easily chosen: *pig*, *horse*, and *cow*, the names are arbitrary.

```

# create a 10GiB Gold volume
docker volume create -d solidfire --name sfGold -o type=Gold -o size=10G

# create a 100GiB Bronze volume
docker volume create -d solidfire --name sfBronze -o type=Bronze -o size=100G

```

Other SolidFire Create Options

Volume create options for SolidFire:

- `size` - the size of the volume, defaults to 1GiB or config entry ... `"defaults": {"size": "5G"}`
- `blocksize` - use either 512 or 4096, defaults to 512 or config entry `DefaultBlockSize`

E-Series Volume Options

Media Type

The E-Series driver offers the ability to specify the type of disk which will be used to back the volume and, like the other drivers, the ability to set the size of the volume at creation time.

Currently only two values for `mediaType` are supported: `ssd` and `hdd`.

```

# create a 10GiB SSD backed volume
docker volume create -d eseries --name eseriesSsd -o mediaType=ssd -o size=10G

# create a 100GiB HDD backed volume
docker volume create -d eseries --name eseriesHdd -o mediaType=hdd -o size=100G

```

File System Type

The user can specify the file system type to use to format the volume. The default for `fileSystemType` is `ext4`. Valid values are `ext3`, `ext4`, and `xfs`.

```

# create a volume using xfs
docker volume create -d eseries --name xfsVolume -o fileSystemType=xfs

```

Cloud Volumes Service (CVS) on AWS Volume Options

Volume create options for the CVS on AWS driver:

- `size` - the size of the volume, defaults to 100 GB
- `serviceLevel` - the CVS service level of the volume, defaults to `standard`. Valid values are `standard`, `premium`, and `extreme`.
- `snapshotReserve` - this will set the snapshot reserve to the desired percentage. The default is no value, meaning CVS will select the snapshot reserve (usually 0%).

Using these options during the docker volume create operation is super simple, just provide the option and the value using the `-o` operator during the CLI operation. These override any equivalent values from the JSON configuration file.

```
# create a 200GiB volume
docker volume create -d netapp --name demo -o size=200G

# create a 500GiB premium volume
docker volume create -d netapp --name demo -o size=500G -o serviceLevel=premium
```

The minimum volume size is 100 GB.

Destroy a Volume

```
# destroy the volume just like any other Docker volume
docker volume rm firstVolume
```

Volume Cloning

When using the `ontap-nas`, `ontap-san`, `solidfire-san` and `aws-cvs` storage drivers, Trident can clone volumes. When using the `ontap-nas-flexgroup` or `ontap-nas-economy` drivers, cloning is not supported.

```
# inspect the volume to enumerate snapshots
docker volume inspect <volume_name>

# create a new volume from an existing volume. this will result in a new snapshot_
↳being created
docker volume create -d <driver_name> --name <new_name> -o from=<source_docker_volume>

# create a new volume from an existing snapshot on a volume. this will not create a_
↳new snapshot
docker volume create -d <driver_name> --name <new_name> -o from=<source_docker_volume>
↳ -o fromSnapshot=<source_snap_name>
```

Here is an example of that in action:

```
[me@host ~]$ docker volume inspect firstVolume

[
  {
    "Driver": "ontap-nas",
    "Labels": null,
    "Mountpoint": "/var/lib/docker-volumes/ontap-nas/netappdvp_firstVolume",
```

(continues on next page)

(continued from previous page)

```

    "Name": "firstVolume",
    "Options": {},
    "Scope": "global",
    "Status": {
      "Snapshots": [
        {
          "Created": "2017-02-10T19:05:00Z",
          "Name": "hourly.2017-02-10_1505"
        }
      ]
    }
  }
]

[me@host ~]$ docker volume create -d ontap-nas --name clonedVolume -o from=firstVolume
clonedVolume

[me@host ~]$ docker volume rm clonedVolume
[me@host ~]$ docker volume create -d ontap-nas --name volFromSnap -o from=firstVolume_
↪-o fromSnapshot=hourly.2017-02-10_1505
volFromSnap

[me@host ~]$ docker volume rm volFromSnap

```

Access Externally Created Volumes

Externally created block devices (or their clones) may be accessed by containers using Trident only if they have no partitions and if their filesystem is supported by Trident (example: an ext4-formatted /dev/sdc1 will not be accessible via Trident).

4.4 Known issues

1. Volume names must be a minimum of 2 characters in length

Note: This is a Docker client limitation. The client will interpret a single character name as being a Windows path. See [bug 25773](#).

2. Docker Swarm has certain behaviors that prevent us from supporting it with every storage and driver combination:
 - Docker Swarm presently makes use of volume name instead of volume ID as its unique volume identifier.
 - Volume requests are simultaneously sent to each node in a Swarm cluster.
 - Volume Plugins (including Trident) must run independently on each node in a Swarm cluster.

Due to the way ONTAP works and how the ontap-nas and ontap-san drivers function, they are the only ones that happen to be able to operate within these limitations.

The rest of the drivers are subject to issues like race conditions that can result in the creation of a large number of volumes for a single request without a clear “winner”; for example, Element has a feature that allows volumes to have the same name but different IDs.

NetApp has provided feedback to the Docker team, but does not have any indication of future recourse.

3. If a FlexGroup is in the process of being provisioned, ONTAP will not provision a second FlexGroup if the second FlexGroup has one or more aggregates in common with the FlexGroup being provisioned.

4.5 Troubleshooting

The most common problem new users run into is a misconfiguration that prevents the plugin from initializing. When this happens you will likely see a message like this when you try to install or enable the plugin:

```
Error response from daemon: dial unix /run/docker/plugins/<id>/netapp.sock:
↪connect: no such file or directory
```

This simply means that the plugin failed to start. Luckily, the plugin has been built with a comprehensive logging capability that should help you diagnose most of the issues you are likely to come across.

The method you use to access or tune those logs varies based on how you are running the plugin.

If there are problems with mounting a PV to a container, ensure that `rpcbind` is installed and running. Use the required package manager for the host OS and check if `rpcbind` is running. You can check the status of the `rpcbind` service by running `systemctl status rpcbind` or its equivalent.

4.5.1 Managed plugin method

If you are running Trident using the recommended managed plugin method (i.e., using `docker plugin` commands), the logs are passed through Docker itself, so they will be interleaved with Docker's own logging.

To view them, simply run:

```
# docker plugin ls
ID                NAME                DESCRIPTION
↪ENABLED
4fb97d2b956b     netapp:latest      nDVP - NetApp Docker Volume Plugin
↪false

# journalctl -u docker | grep 4fb97d2b956b
```

The standard logging level should allow you to diagnose most issues. If you find that's not enough, you can enable debug logging:

```
# install the plugin with debug logging enabled
docker plugin install netapp/trident-plugin:<version> --alias <alias>
↪debug=true

# or, enable debug logging when the plugin is already installed
docker plugin disable <plugin>
docker plugin set <plugin> debug=true
docker plugin enable <plugin>
```

4.5.2 Binary method

If you are not running as a managed plugin, you are running the binary itself on the host. The logs are available in the host's `/var/log/netappdvp` directory. If you need to enable debug logging, specify `-debug` when you run the plugin.

5.1 Supported frontends (orchestrators)

Trident supports multiple container engines and orchestrators, including:

- [NetApp Kubernetes Service](#)
- Kubernetes 1.11 to 1.15
- OpenShift 3.11
- Docker Enterprise 17.06 or later (latest: 3.0)

In addition, Trident should work with any distribution of Docker or Kubernetes that uses one of the supported versions as a base, such as Rancher or Tectonic.

5.2 Supported backends (storage)

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

- FAS/AFF/Select/Cloud ONTAP 8.3 or later
- HCI/SolidFire Element OS 8 or later
- E/EF-Series SANtricity
- Azure NetApp Files
- Cloud Volumes Service for AWS

5.3 Feature Gates

Trident requires some feature gates to be enabled for it to work. These feature gates are detailed below:

Feature Gate	Default	Stage	Since	Until
CSIDriverRegistry	False	Alpha	1.12	1.13
CSIDriverRegistry	True	Beta	1.14	•
CSINodeInfo	False	Alpha	1.12	1.13
CSINodeInfo	True	Beta	1.14	•

If using Kubernetes 1.13:

- Enable the `CSIDriverRegistry` and `CSINodeInfo` flags.

If using Kubernetes 1.14 and above:

- The feature gates are enabled by default.

5.4 Supported host operating systems

By default Trident itself runs in a container, therefore it will run on any Linux worker.

However, those workers do need to be able to mount the volumes that Trident provides using the standard NFS client or iSCSI initiator, depending on the backend(s) you're using.

These are the Linux distributions that are known to work:

- Debian 8 or later
- Ubuntu 16.04 or later
- CentOS 7.0 or later
- RHEL 7.0 or later
- CoreOS 1353.8.0 or later

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

5.5 Host configuration

Depending on the backend(s) in use, NFS and/or iSCSI utilities must be installed on all of the workers in the cluster. See the *worker preparation* guide for details.

5.6 Storage system configuration

Trident may require some changes to a storage system before a backend configuration can use it. See the *backend configuration* guide for details.

CHAPTER 6

Getting help

Trident is an officially supported NetApp project. That means you can reach out to NetApp using any of the [standard mechanisms](#) and get the enterprise grade support that you need.

There is also a vibrant public community of container users (including Trident developers) on the [#containers](#) channel in [NetApp's Slack team](#). This is a great place to ask general questions about the project and discuss related topics with like-minded peers.

Trident exposes several command-line options. Normally the defaults will suffice, but you may want to modify them in your deployment. They are:

7.1 Logging

- `-debug`: Optional; enables debugging output.
- `-loglevel <level>`: Optional; sets the logging level (debug, info, warn, error, fatal). Defaults to info.

7.2 Kubernetes

- `-k8s_pod`: Optional; however, either this or `-k8s_api_server` must be set to enable Kubernetes support. Setting this will cause Trident to use its containing pod's Kubernetes service account credentials to contact the API server. This only works when Trident runs as a pod in a Kubernetes cluster with service accounts enabled.
- `-k8s_api_server <insecure-address:insecure-port>`: Optional; however, either this or `-k8s_pod` must be used to enable Kubernetes support. When specified, Trident will connect to the Kubernetes API server using the provided insecure address and port. This allows Trident to be deployed outside of a pod; however, it only supports insecure connections to the API server. To connect securely, deploy Trident in a pod with the `-k8s_pod` option.
- `-k8s_config_path <file>`: Optional; path to a KubeConfig file.

7.3 Docker

- `-volume_driver <name>`: Optional; driver name used when registering the Docker plugin. Defaults to 'netapp'.
- `-driver_port <port-number>`: Optional; listen on this port rather than a UNIX domain socket.

- `-config <file>`: Path to a backend configuration file.

7.4 REST

- `-address <ip-or-host>`: Optional; specifies the address on which Trident's REST server should listen. Defaults to localhost. When listening on localhost and running inside a Kubernetes pod, the REST interface will not be directly accessible from outside the pod. Use `-address ""` to make the REST interface accessible from the pod IP address.
- `-port <port-number>`: Optional; specifies the port on which Trident's REST server should listen. Defaults to 8000.
- `-rest`: Optional; enable the REST interface. Defaults to true.

The [Trident installer bundle](#) includes a command-line utility, `tridentctl`, that provides simple access to Trident. It can be used to install Trident, as well as to interact with it directly by any Kubernetes users with sufficient privileges, to manage the namespace that contains the Trident pod.

For full usage information, run `tridentctl --help`. Here are the available commands and global options:

```
Usage:
  tridentctl [command]

Available Commands:
  create      Add a resource to Trident
  delete     Remove one or more resources from Trident
  get        Get one or more resources from Trident
  help       Help about any command
  install    Install Trident
  logs       Print the logs from Trident
  uninstall  Uninstall Trident
  update     Modify a resource in Trident
  version    Print the version of Trident

Flags:
  -d, --debug           Debug output
  -h, --help            help for tridentctl
  -n, --namespace string Namespace of Trident deployment
  -o, --output string   Output format. One of json|yaml|name|wide|ps (default)
  -s, --server string   Address/port of Trident REST interface
```

8.1 create

Add a resource to Trident

```
Usage:
  tridentctl create [command]

Available Commands:
  backend      Add a backend to Trident
```

8.2 delete

Remove one or more resources from Trident

```
Usage:
  tridentctl delete [command]

Available Commands:
  backend      Delete one or more storage backends from Trident
  snapshot     Delete one or more volume snapshots from Trident
  storageclass Delete one or more storage classes from Trident
  volume       Delete one or more storage volumes from Trident
```

8.3 get

Get one or more resources from Trident

```
Usage:
  tridentctl get [command]

Available Commands:
  backend      Get one or more storage backends from Trident
  snapshot     Get one or more snapshots from Trident
  storageclass Get one or more storage classes from Trident
  volume       Get one or more volumes from Trident
```

8.4 import volume

Import an existing volume to Trident

```
Usage:
  tridentctl import volume <backendName> <volumeName> [flags]

Aliases:
  volume, v

Flags:
  -f, --filename string  Path to YAML or JSON PVC file
  -h, --help             help for volume
  --no-manage            Create PV/PVC only, don't assume volume lifecycle management
```

8.5 install

Install Trident

```
Usage:
  tridentctl install [flags]

Flags:
  --csi                    Install CSI Trident (override for Kubernetes 1.13
↳only, requires feature gates).
  --dry-run                Run all the pre-checks, but don't install anything.
  --etcd-image string     The etcd image to install.
  --generate-custom-yaml Generate YAML files, but don't install anything.
  -h, --help              help for install
  --k8s-timeout duration  The number of seconds to wait before timing out on
↳Kubernetes operations. (default 3m0s)
  --pv string              The name of the PV used by Trident.
  --pvc string             The name of the PVC used by Trident.
  --silent                 Disable most output during installation.
  --trident-image string  The Trident image to install.
  --use-custom-yaml       Use any existing YAML files that exist in setup
↳directory.
  --volume-name string    The name of the storage volume used by Trident.
  --volume-size string    The size of the storage volume used by Trident.
↳(default "2Gi")
```

8.6 logs

Print the logs from Trident

```
Usage:
  tridentctl logs [flags]

Flags:
  -a, --archive           Create a support archive with all logs unless otherwise
↳specified.
  -h, --help              help for logs
  -l, --log string        Trident log to display. One of trident|etcd|auto|all (default
↳"auto")
  -p, --previous          Get the logs for the previous container instance if it exists.
```

8.7 uninstall

Uninstall Trident

```
Usage:
  tridentctl uninstall [flags]

Flags:
  -a, --all                Deletes almost all artifacts of Trident, including
↳the PVC and PV used by Trident;
                           however, it doesn't delete the volume used by
↳Trident from the storage backend. Use with caution!
```

(continues on next page)

(continued from previous page)

```
-h, --help          help for uninstall
--silent           Disable most output during uninstallation.
```

8.8 update

Modify a resource in Trident

```
Usage:
  tridentctl update [command]

Available Commands:
  backend      Update a backend in Trident
```

8.9 upgrade

Upgrade a resource in Trident

```
Usage:
  tridentctl upgrade [command]

Available Commands:
  volume      Upgrade one or more persistent volumes from NFS/iSCSI to CSI
```

8.10 version

Print the version of tridentctl and the running Trident service

```
Usage:
  tridentctl version
```

While *tridentctl* is the easiest way to interact with Trident's REST API, you can use the REST endpoint directly if you prefer.

This is particularly useful for advanced installations that are using Trident as a standalone binary in non-Kubernetes deployments.

For better security, Trident's *REST API* is restricted to localhost by default when running inside a pod. You will need to set Trident's `-address` argument in its pod configuration to change this behavior.

The API works as follows:

- GET `<trident-address>/trident/v1/<object-type>`: Lists all objects of that type.
- GET `<trident-address>/trident/v1/<object-type>/<object-name>`: Gets the details of the named object.
- POST `<trident-address>/trident/v1/<object-type>`: Creates an object of the specified type. Requires a JSON configuration for the object to be created; see the previous section for the specification of each object type. If the object already exists, behavior varies: backends update the existing object, while all other object types will fail the operation.
- DELETE `<trident-address>/trident/v1/<object-type>/<object-name>`: Deletes the named resource. Note that volumes associated with backends or storage classes will continue to exist; these must be deleted separately. See the section on backend deletion below.

To see an example of how these APIs are called, pass the debug (`-d`) flag to *tridentctl*.

CHAPTER 10

Simple Kubernetes install

Those that are interested in Trident and just getting started with Kubernetes frequently ask us for a simple way to install Kubernetes to try it out.

These instructions provide a bare-bones single node cluster that Trident will be able to integrate with for demonstration purposes.

Warning: The Kubernetes cluster these instructions build should never be used in production. Follow production deployment guides provided by your distribution for that.

This is a simplification of the [kubeadm install guide](#) provided by Kubernetes. If you're having trouble, your best bet is to revert to that guide.

10.1 Prerequisites

An Ubuntu 16.04 machine with at least 1 GB of RAM.

These instructions are very opinionated by design, and will not work with anything else. For more generic instructions, you will need to run through the entire [kubeadm install guide](#).

10.2 Install Docker CE 17.03

```
apt-get update && apt-get install -y curl apt-transport-https
curl -fsSL https://download.docker.com/linux/ubuntu/gpg | apt-key add -
cat <<EOF >/etc/apt/sources.list.d/docker.list
deb https://download.docker.com/linux/$(lsb_release -si | tr '[:upper:]' '[:lower:]')
  ↪$(lsb_release -cs) stable
```

(continues on next page)

(continued from previous page)

```
EOF
apt-get update && apt-get install -y docker-ce=$(apt-cache madison docker-ce | grep ↵
↵17.03 | head -1 | awk '{print $3}')

```

10.3 Install the appropriate version of kubeadm, kubectl and kubelet

```
curl -s https://packages.cloud.google.com/apt/doc/apt-key.gpg | apt-key add -
cat <<EOF >/etc/apt/sources.list.d/kubernetes.list
deb http://apt.kubernetes.io/ kubernetes-xenial main
EOF
apt-get update
apt-get install -y kubeadm=1.15.1 kubectl=1.15.1 kubelet=1.15.1 kubernetes-cni=0.7.5

```

10.4 Configure the host

```
swapoff -a
# Comment out swap line in fstab so that it remains disabled after reboot
vi /etc/fstab

```

10.5 Create the cluster

```
kubeadm init --kubernetes-version stable-1.15 --token-ttl 0 --pod-network-cidr=192.
↵168.0.0/16

```

10.6 Install the kubectl creds and untaint the cluster

```
mkdir -p $HOME/.kube
sudo cp -i /etc/kubernetes/admin.conf $HOME/.kube/config
sudo chown $(id -u):$(id -g) $HOME/.kube/config
kubectl taint nodes --all node-role.kubernetes.io/master-

```

10.7 Add an overlay network

```
kubectl apply -f https://docs.projectcalico.org/v3.7/manifests/calico.yaml

```

10.8 Verify that all of the services started

After completing those steps, you should see output similar to this within a few minutes:

```
# kubectl get po -n kube-system
```

RESTARTS	AGE	NAMESPACE	NAME	READY	STATUS	
↪		kube-system	calico-kube-controllers-658558ddf8-4hwf4	1/1	Running	0 ↪
↪	20m					
↪		kube-system	calico-node-75zdm	1/1	Running	0 ↪
↪	20m					
↪		kube-system	coredns-5c98db65d4-bgtlq	1/1	Running	0 ↪
↪	20m					
↪		kube-system	coredns-5c98db65d4-pr8rf	1/1	Running	0 ↪
↪	20m					
↪		kube-system	etcd-trident-1907	1/1	Running	0 ↪
↪	20m					
↪		kube-system	kube-apiserver-trident-1907	1/1	Running	0 ↪
↪	20m					
↪		kube-system	kube-controller-manager-trident-1907	1/1	Running	0 ↪
↪	20m					
↪		kube-system	kube-proxy-q4c75	1/1	Running	0 ↪
↪	20m					
↪		kube-system	kube-scheduler-trident-1907	1/1	Running	0 ↪
↪	20m					

Notice that all of the Kubernetes services are in a *Running* state. Congratulations! At this point your cluster is operational.

If this is the first time you're using Kubernetes, we highly recommend a [walkthrough](#) to familiarize yourself with the concepts and tools.