



Configure and manage backends

Trident

NetApp
October 30, 2026

Table of Contents

- Configure and manage backends 1
 - Configure Trident backends 1
 - Azure NetApp Files 1
 - Configure an Azure NetApp Files backend 1
 - Prepare to configure an Azure NetApp Files backend 4
 - Azure NetApp Files backend configuration options and examples 7
 - Google Cloud NetApp Volumes 18
 - Configure Google Cloud NetApp Volumes 18
 - Configure Google Cloud NetApp Volumes for SAN workloads 22
 - Prepare to configure a Google Cloud NetApp Volumes backend 28
 - Google Cloud NetApp Volumes backend configuration options and examples 28
 - Configure auto-tiering for Google Cloud NetApp Volumes 39
 - Configure Google Cloud NetApp Volumes ONTAP-mode storage pools 42
 - Patch value of deletionPolicy to retain 184

Configure and manage backends

Configure Trident backends

A 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 automatically offers up storage pools from backends that match the requirements defined by a storage class. Learn how to configure the backend for your storage system.

- [Configure an Azure NetApp Files backend](#)
- [Configure a Google Cloud NetApp Volumes backend](#)
- [Configure a NetApp HCI or SolidFire backend](#)
- [Configure a backend with ONTAP or Cloud Volumes ONTAP NAS drivers](#)
- [Configure a backend with ONTAP or Cloud Volumes ONTAP SAN drivers](#)
- [Use Trident with Amazon FSx for NetApp ONTAP](#)

Azure NetApp Files

Configure an Azure NetApp Files backend

Use Azure NetApp Files as a backend for Trident.

This backend supports NFS and SMB volumes.

Trident supports managed identities and workload identity for Azure Kubernetes Service (AKS) clusters.

Supported Azure cloud environments

Trident supports Azure NetApp Files backends in multiple Azure cloud environments.

Supported Azure clouds include:

- Azure Commercial
- Azure Government (Azure Government / MAG)
- Azure China

When you deploy Trident or configure an Azure NetApp Files backend, ensure that Azure Resource Manager and authentication endpoints match your Azure cloud environment.

Review Azure NetApp Files driver support

Trident provides the following Azure NetApp Files storage driver.

Supported access modes include *ReadWriteOnce* (RWO), *ReadOnlyMany* (ROX), *ReadWriteMany* (RWX), and *ReadWriteOncePod* (RWOP).

Driver	Protocol	volumeMode	Access modes supported	File systems supported
azure-netapp-files	NFS SMB	Filesystem	RWO, ROX, RWX, RWOP	nfs, smb

Review considerations

- Azure NetApp Files does not support volumes smaller than 50 GiB. Trident creates a 50-GiB volume when a smaller volume is requested.
- Trident supports SMB volumes mounted to pods running on Windows nodes only.
- Azure NetApp Files deployments in non-Commercial Azure clouds require cloud-specific Azure Resource Manager and authentication endpoints. Ensure that Trident and any backend configuration use the endpoints appropriate for your Azure cloud environment.

Use managed identities for AKS

Trident supports [managed identities](#) for AKS clusters.

If you use `tridentctl` to create or manage Azure NetApp Files backends, ensure that it is configured for the correct Azure cloud environment.

To use managed identities, you must have:

- A Kubernetes cluster deployed using AKS
- Managed identities configured on the AKS Kubernetes cluster
- Trident installed with `cloudProvider` set to "Azure"

Trident operator

Edit `tridentorchestrator_cr.yaml` and set `cloudProvider` to "Azure".

```
apiVersion: trident.netapp.io/v1
kind: TridentOrchestrator
metadata:
  name: trident
spec:
  debug: true
  namespace: trident
  imagePullPolicy: IfNotPresent
  cloudProvider: "Azure"
```

Helm

The following example installs Trident and sets `cloudProvider` using the environment variable `$CP`:

```
helm install trident trident-operator-100.2606.0.tgz --create-namespace
```

```
--namespace <trident-namespace> --set cloudProvider=$CP
```

tridentctl

The following example installs Trident and sets the `cloud-provider` flag to Azure:

```
tridentctl install --cloud-provider="Azure" -n trident
```

Use workload identity for AKS

Workload identity enables Kubernetes pods to access Azure resources by authenticating as a workload identity.

If you use `tridentctl` to create or manage Azure NetApp Files backends, ensure that it is configured for the correct Azure cloud environment.

To use workload identity, you must have:

- A Kubernetes cluster deployed using AKS
- Workload identity and `oidc-issuer` configured on the AKS Kubernetes cluster
- Trident installed with `cloudProvider` set to "Azure" and `cloudIdentity` set to the workload identity value

Trident operator

Edit `tridentorchestrator_cr.yaml` and set `cloudProvider` to "Azure".

Set `cloudIdentity` to `azure.workload.identity/client-id: xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxx`.

```
apiVersion: trident.netapp.io/v1
kind: TridentOrchestrator
metadata:
  name: trident
spec:
  debug: true
  namespace: trident
  imagePullPolicy: IfNotPresent
  cloudProvider: "Azure"
  cloudIdentity: 'azure.workload.identity/client-id: xxxxxxxx-xxxx-
xxxx-xxxx-xxxxxxxxxx' # Edit
```

Helm

Set the values for the **cloud-provider (CP)** and **cloud-identity (CI)** flags using the following environment variables:

```
export CP="Azure"
```




- `subscriptionID`, `tenantID`, `clientID`, `location`, and `clientSecret` are optional when using managed identities on an AKS cluster.
- `tenantID`, `clientID`, and `clientSecret` are optional when using a cloud identity on an AKS cluster.
- Azure NetApp Files deployments in non-Commercial Azure clouds require cloud-specific Azure Resource Manager and authentication endpoints. Ensure that Trident and any backend configuration use the endpoints appropriate for your Azure cloud environment.

- A capacity pool. Refer to [Microsoft: Create a capacity pool for Azure NetApp Files](#).
- A subnet delegated to Azure NetApp Files. Refer to [Microsoft: Delegate a subnet to Azure NetApp Files](#).
- `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. The App Registration should use either:
 - The Owner or Contributor role [predefined by Azure](#).
 - A [custom Contributor role](#) at the subscription level (`assignableScopes`) with the following permissions that are limited to only what Trident requires. After creating the custom role, [assign the role using the Azure portal](#).

▼ Custom contributor role

```
{
  "id": "/subscriptions/<subscription-id>/providers/Microsoft.Authorization/roleDefinitions/<role-definition-id>",
  "properties": {
    "roleName": "custom-role-with-limited-perms",
    "description": "custom role providing limited permissions",
    "assignableScopes": [
      "/subscriptions/<subscription-id>"
    ],
    "permissions": [
      {
        "actions": [
          "Microsoft.NetApp/netAppAccounts/capacityPools/read",
          "Microsoft.NetApp/netAppAccounts/capacityPools/write",

          "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/read",

          "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/write",

          "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/delete",

          "Microsoft.NetApp/netAppAccounts/capacityPools/volumes/snapshots/read",

```

```

"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/snapshots/write",

"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/snapshots/delete",

"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/MountTargets/read",
    "Microsoft.Network/virtualNetworks/read",
    "Microsoft.Network/virtualNetworks/subnets/read",

"Microsoft.Features/featureProviders/subscriptionFeatureRegistrations/read",

"Microsoft.Features/featureProviders/subscriptionFeatureRegistrations/write",

"Microsoft.Features/featureProviders/subscriptionFeatureRegistrations/delete",
    "Microsoft.Features/features/read",
    "Microsoft.Features/operations/read",
    "Microsoft.Features/providers/features/read",
    "Microsoft.Features/providers/features/register/action",

"Microsoft.Features/providers/features/unregister/action",

"Microsoft.Features/subscriptionFeatureRegistrations/read"
    ],
    "notActions": [],
    "dataActions": [],
    "notDataActions": []
  }
]
}
}

```

- The Azure location that contains at least one [delegated subnet](#). As of Trident 22.01, the location parameter is a required field at the top level of the backend configuration file. Location values specified in virtual pools are ignored.
- To use Cloud Identity, get the client ID from a [user-assigned managed identity](#) and specify that ID in `azure.workload.identity/client-id: xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx`.

Additional requirements for SMB volumes

To create an SMB volume, you must have:

- Active Directory configured and connected to Azure NetApp Files. Refer to [Microsoft: Create and manage Active Directory connections for Azure NetApp Files](#).
- A Kubernetes cluster with a Linux controller node and at least one Windows worker node running Windows Server 2022. Trident supports SMB volumes mounted to pods running on Windows nodes only.
- At least one Trident secret containing your Active Directory credentials so Azure NetApp Files can authenticate to Active Directory. To generate secret `smbcreds`:

```
kubectl create secret generic smbcreds --from-literal username=user
--from-literal password='password'
```

- A CSI proxy configured as a Windows service. To configure a `csi-proxy`, refer to [GitHub: CSI Proxy](#) or [GitHub: CSI Proxy for Windows](#) for Kubernetes nodes running on Windows.

Azure NetApp Files backend configuration options and examples

Learn about NFS and SMB backend configuration options for Azure NetApp Files and review configuration examples.

Backend configuration options

Trident uses your backend configuration (subnet, virtual network, service level, and location) to create Azure NetApp Files volumes on capacity pools that are available in the requested location and match the requested service level and subnet.

Azure NetApp Files backends provide these configuration options.

Parameter	Description	Default
<code>version</code>	Backend configuration version.	Always 1
<code>storageDriverName</code>	Name of the storage driver	"azure-netapp-files"
<code>backendName</code>	Custom name for the storage backend	Driver name + "_" + random characters
<code>subscriptionID</code>	The subscription ID from your Azure subscription Optional when managed identities is enabled on an AKS cluster.	
<code>tenantID</code>	The tenant ID from an App Registration Optional when managed identities or cloud identity is used on an AKS cluster.	

Parameter	Description	Default
<code>clientID</code>	The client ID from an App Registration Optional when managed identities or cloud identity is used on an AKS cluster.	
<code>clientSecret</code>	The client secret from an App Registration Optional when managed identities or cloud identity is used on an AKS cluster.	
<code>serviceLevel</code>	One of Standard, Premium, or Ultra	"" (random)
<code>location</code>	Name of the Azure location where the new volumes will be created Optional when managed identities is enabled on an AKS cluster.	
<code>cloudConfiguration.cloudName</code>	Identifies the Azure cloud environment for the backend. Set to <code>AzureGovernment</code> when using Azure Government.	
<code>resourceGroups</code>	List of resource groups for filtering discovered resources	[] (no filter)
<code>netappAccounts</code>	List of NetApp accounts for filtering discovered resources	[] (no filter)
<code>capacityPools</code>	List of capacity pools for filtering discovered resources	[] (no filter, random)
<code>virtualNetwork</code>	Name of a virtual network with a delegated subnet	""
<code>subnet</code>	Name of a subnet delegated to <code>Microsoft.Netapp/volumes</code>	""
<code>networkFeatures</code>	Set of VNet features for a volume, may be <code>Basic</code> or <code>Standard</code> . Network Features is not available in all regions and might have to be enabled in a subscription. Specifying <code>networkFeatures</code> when the functionality is not enabled causes volume provisioning to fail.	""

Parameter	Description	Default
nfsMountOptions	<p>Fine-grained control of NFS mount options.</p> <p>Ignored for SMB volumes.</p> <p>To mount volumes using NFS version 4.1, include <code>nfsvers=4</code> in the comma-delimited mount options list to choose NFS v4.1.</p> <p>Mount options set in a storage class definition override mount options set in backend configuration.</p>	"nfsvers=3"
limitVolumeSize	Fail provisioning if the requested volume size is above this value	"" (not enforced by default)
debugTraceFlags	<p>Debug flags to use when troubleshooting. Example, <code>\{"api": false, "method": true, "discovery": true\}</code>. Do not use this unless you are troubleshooting and require a detailed log dump.</p>	null
nasType	<p>Configure NFS or SMB volumes creation.</p> <p>Options are <code>nfs</code>, <code>smb</code> or <code>null</code>. Setting to <code>null</code> defaults to NFS volumes.</p>	nfs
supportedTopologies	<p>Represents a list of regions and zones that are supported by this backend.</p> <p>For more information, refer to Use CSI Topology.</p>	
qosType	Represents the QoS type: Auto or Manual.	Auto
maxThroughput	<p>Sets the maximum throughput allowed in MiB/sec.</p> <p>Supported only for manual QoS capacity pools.</p>	4 MiB/sec



For more information on Network Features, refer to [Configure network features for an Azure NetApp Files volume](#).

Consider Azure cloud environments (26.06)

Starting with the 26.06 release, Trident supports creating and managing Azure NetApp Files backends in multiple Azure cloud environments.

Supported Azure clouds include:

- Azure Commercial
- Azure Government (Azure Government / MAG)
- Azure China

When you deploy Trident or create an Azure NetApp Files backend, ensure that Azure Resource Manager and authentication endpoints match your Azure cloud environment.

If the endpoints do not match, `tridentctl` cannot authenticate and backend creation fails.

For an Azure Government backend, set `cloudConfiguration.cloudName` in the backend configuration:

```
cloudConfiguration:
  cloudName: AzureGovernment
```

Required permissions and resources

If you receive a "No capacity pools found" error when creating a PVC, it is likely your app registration doesn't have the required permissions and resources (subnet, virtual network, capacity pool) associated.

If debug is enabled, Trident logs the Azure resources discovered when the backend is created.

Verify an appropriate role is being used.

The values for `resourceGroups`, `netappAccounts`, `capacityPools`, `virtualNetwork`, and `subnet` can be specified using short or fully-qualified names.

Fully-qualified names are recommended in most situations as short names can match multiple resources with the same name.



If the vNet is located in a different resource group from the Azure NetApp Files (ANF) storage account, specify the resource group for the virtual network while configuring the `resourceGroups` list for the backend.

The `resourceGroups`, `netappAccounts`, and `capacityPools` values are filters that restrict the set of discovered resources to those available to this storage backend and may be specified in any combination.

Fully-qualified names follow this format:

Type	Format
Resource group	<resource group>
NetApp account	<resource group>/<netapp account>
Capacity pool	<resource group>/<netapp account>/<capacity pool>
Virtual network	<resource group>/<virtual network>
Subnet	<resource group>/<virtual network>/<subnet>

Volume provisioning

You can control default volume provisioning by specifying the following options in a special section of the configuration file.

Refer to [Example configurations](#) for details.

Parameter	Description	Default
exportRule	Export rules for new volumes. exportRule must be a comma-separated list of any combination of IPv4 addresses or IPv4 subnets in CIDR notation. Ignored for SMB volumes.	"0.0.0.0/0"
snapshotDir	Access to the .snapshot directory	true, false (Set explicitly).
size	The default size of new volumes	"100G"
unixPermissions	The unix permissions of new volumes (4 octal digits). Ignored for SMB volumes.	"" (preview feature, requires whitelisting in subscription)

Example configurations

The following examples show basic configurations that leave most parameters to default. This is the easiest way to define a backend.

▼ Minimal configuration

This is the absolute minimum backend configuration.

With this configuration, Trident discovers all of your NetApp accounts, capacity pools, and subnets delegated to Azure NetApp Files in the configured location, and places new volumes on one of those pools and subnets randomly.

Because `nasType` is omitted, the `nfs` default applies and the backend will provision for NFS volumes.

This configuration is ideal when you are just getting started with Azure NetApp Files and trying things out, but in practice you are going to want to provide additional scoping for the volumes you provision.

```
---
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: backend-tbc-anf-1
  namespace: trident
spec:
  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
```

▼ Managed identities for AKS

This backend configuration omits `subscriptionID`, `tenantID`, `clientID`, and `clientSecret`, which are optional when using managed identities.

```
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: backend-tbc-anf-1
  namespace: trident
spec:
  version: 1
  storageDriverName: azure-netapp-files
  capacityPools:
    - resource-group-1/netapp-account-1/ultra-pool
  resourceGroups:
    - resource-group-1
  netappAccounts:
    - resource-group-1/netapp-account-1
  virtualNetwork: resource-group-1/eastus-prod-vnet
  subnet: resource-group-1/eastus-prod-vnet/eastus-anf-subnet
```

▼ Cloud identity for AKS

This backend configuration omits `tenantID`, `clientID`, and `clientSecret`, which are optional when using a cloud identity.

```
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: backend-tbc-anf-1
  namespace: trident
spec:
  version: 1
  storageDriverName: azure-netapp-files
  capacityPools:
    - ultra-pool
  resourceGroups:
    - aks-ami-eastus-rg
  netappAccounts:
    - smb-na
  virtualNetwork: eastus-prod-vnet
```

```
subnet: eastus-anf-subnet
location: eastus
subscriptionID: 9f87c765-4774-fake-ae98-a721add45451
```

▼ Specific service level configuration with capacity pool filters

This backend configuration places volumes in Azure's `eastus` location in an `Ultra` capacity pool. Trident automatically discovers all of the subnets delegated to Azure NetApp Files in that location and places 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: Ultra
capacityPools:
  - application-group-1/account-1/ultra-1
  - application-group-1/account-1/ultra-2
```

▼ Backend example with manual QoS capacity pools

This backend configuration places volumes in Azure's `eastus` location with manual QoS capacity pools.

```
---
version: 1
storageDriverName: azure-netapp-files
backendName: anf1
location: eastus
labels:
  clusterName: test-cluster-1
  cloud: anf
  nasType: nfs
defaults:
  qosType: Manual
storage:
  - serviceLevel: Ultra
    labels:
      performance: gold
    defaults:
      maxThroughput: 10
  - serviceLevel: Premium
    labels:
```

```

    performance: silver
  defaults:
    maxThroughput: 5
- serviceLevel: Standard
  labels:
    performance: bronze
  defaults:
    maxThroughput: 3

```

▼ Advanced configuration

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: Ultra
capacityPools:
- application-group-1/account-1/ultra-1
- application-group-1/account-1/ultra-2
virtualNetwork: application-group-1/eastus-prod-vnet
subnet: application-group-1/eastus-prod-vnet/my-subnet
networkFeatures: Standard
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
  snapshotDir: "true"
  size: 200Gi
  unixPermissions: "0777"

```

▼ Virtual pool configuration

This backend configuration defines multiple storage pools 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.

Virtual pool labels were used to differentiate the pools based on performance.

```

---
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
resourceGroups:
  - application-group-1
networkFeatures: Basic
nfsMountOptions: vers=3,proto=tcp,timeo=600
labels:
  cloud: azure
storage:
  - labels:
      performance: gold
      serviceLevel: Ultra
      capacityPools:
        - application-group-1/netapp-account-1/ultra-1
        - application-group-1/netapp-account-1/ultra-2
      networkFeatures: Standard
  - labels:
      performance: silver
      serviceLevel: Premium
      capacityPools:
        - application-group-1/netapp-account-1/premium-1
  - labels:
      performance: bronze
      serviceLevel: Standard
      capacityPools:
        - application-group-1/netapp-account-1/standard-1
        - application-group-1/netapp-account-1/standard-2

```

▼ Supported topologies configuration

Trident facilitates provisioning of volumes for workloads based on regions and availability zones.

The `supportedTopologies` block in this backend configuration is used to provide a list of regions and zones per backend.

The region and zone values specified here must match the region and zone values from the labels on each Kubernetes cluster node.

These regions and zones represent the list of permissible values that can be provided in a storage class. For storage classes that contain a subset of the regions and zones provided in a backend, Trident creates volumes in the mentioned region and zone.

For more information, refer to [Use CSI Topology](#).

```

---
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: Ultra
capacityPools:
  - application-group-1/account-1/ultra-1
  - application-group-1/account-1/ultra-2
supportedTopologies:
  - topology.kubernetes.io/region: eastus
    topology.kubernetes.io/zone: eastus-1
  - topology.kubernetes.io/region: eastus
    topology.kubernetes.io/zone: eastus-2

```

Storage class definitions

The following StorageClass definitions refer to the storage pools above.

Example definitions using `parameter.selector` field

Using `parameter.selector` you can specify for each StorageClass the virtual pool that is 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
provisioner: csi.trident.netapp.io
parameters:
  selector: performance=bronze
allowVolumeExpansion: true

```

Example definitions for SMB volumes

Using `nasType`, `node-stage-secret-name`, and `node-stage-secret-namespace`, you can specify an SMB volume and provide the required Active Directory credentials.

▼ Basic configuration on default namespace

```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: anf-sc-smb
provisioner: csi.trident.netapp.io
parameters:
  backendType: "azure-netapp-files"
  trident.netapp.io/nasType: "smb"
  csi.storage.k8s.io/node-stage-secret-name: "smbcreds"
  csi.storage.k8s.io/node-stage-secret-namespace: "default"

```

▼ Using different secrets per namespace

```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: anf-sc-smb
provisioner: csi.trident.netapp.io
parameters:
  backendType: "azure-netapp-files"
  trident.netapp.io/nasType: "smb"
  csi.storage.k8s.io/node-stage-secret-name: "smbcreds"
  csi.storage.k8s.io/node-stage-secret-namespace: ${pvc.namespace}

```

▼ Using different secrets per volume

```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: anf-sc-smb
provisioner: csi.trident.netapp.io

```

```
parameters:
  backendType: "azure-netapp-files"
  trident.netapp.io/nasType: "smb"
  csi.storage.k8s.io/node-stage-secret-name: ${pvc.name}
  csi.storage.k8s.io/node-stage-secret-namespace: ${pvc.namespace}
```



nasType: smb filters for pools which support SMB volumes.
nasType: nfs or nasType: null filters for NFS pools.

Create the backend

After you create the backend configuration file, run the following command:

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

If you use a non-Commercial Azure cloud, ensure that `tridentctl` is configured to use the Azure Resource Manager and authentication endpoints for your Azure cloud environment.

If the backend creation fails, check your backend configuration and view the logs to determine the cause:

```
tridentctl logs
```

After you identify and correct the problem with the configuration file, you can run the create command again.

Google Cloud NetApp Volumes

Configure Google Cloud NetApp Volumes

You can configure Google Cloud NetApp Volumes as a backend for Trident to provision storage for Kubernetes workloads.

Overview

Trident supports Google Cloud NetApp Volumes for both NAS (NFS and SMB) and block (iSCSI) workloads.

- NAS workloads use the `google-cloud-netapp-volumes` backend
- Block (iSCSI) workloads use the `google-cloud-netapp-volumes-san` backend

NAS volumes provide file-based storage and are accessed using NFS or SMB protocols. These volumes support shared access across multiple pods or nodes.

Block volumes provide raw block storage and are accessed as iSCSI devices attached to Kubernetes nodes. These volumes are used when applications require block-level access.

This applies to the following environments:

- Trident 26.02 and later

- Google Kubernetes Engine (GKE) or Red Hat OpenShift
- Google Cloud NetApp Volumes storage pools

To configure block (iSCSI) storage, see [Configure block storage \(iSCSI\)](#).

Prepare to configure

Cloud identity enables Kubernetes workloads to access Google Cloud resources by authenticating as a workload identity instead of using static credentials.

To use cloud identity with Google Cloud NetApp Volumes, you must have:

- A Kubernetes cluster deployed using Google Kubernetes Engine (GKE)
- Workload identity enabled on the GKE cluster and the metadata server enabled on the node pools
- A Google Cloud service account with the Google Cloud NetApp Volumes Admin role (roles/netapp.admin) or an equivalent custom role
- Trident installed with the cloud provider set to GCP and the cloud identity annotation configured

Trident operator

To install Trident using the Trident operator, edit `tridentorchestrator_cr.yaml`:

```
apiVersion: trident.netapp.io/v1
kind: TridentOrchestrator
metadata:
  name: trident
spec:
  namespace: trident
  cloudProvider: "GCP"
  cloudIdentity: "iam.gke.io/gcp-service-account: cloudvolumes-admin-sa@mygcpproject.iam.gserviceaccount.com"
```

Helm

Set the cloud provider and cloud identity when installing Trident with Helm:

```
helm install trident trident-operator-100.6.0.tgz \
  --set cloudProvider=GCP \
  --set cloudIdentity="iam.gke.io/gcp-service-account: cloudvolumes-admin-sa@mygcpproject.iam.gserviceaccount.com"
```

tridentctl

Install Trident by specifying the cloud provider and cloud identity:

```
tridentctl install \
```

```
--cloud-provider=GCP \
--cloud-identity="iam.gke.io/gcp-service-account: cloudvolumes-admin-
sa@mygcpproject.iam.gserviceaccount.com" \
-n trident
```

Configure NAS storage



For Google Cloud NetApp Volumes UNIFIED storage pools, Trident applies UNIFIED-specific naming and validation rules during volume operations.

When locating a volume, Trident can evaluate multiple compatible volume name variants (for example, hyphen and underscore formats) to improve import and discovery reliability.

Driver details

Trident provides the `google-cloud-netapp-volumes` driver to provision NAS storage from Google Cloud NetApp Volumes.

The driver supports the following access modes:

- ReadWriteOnce (RWO)
- ReadOnlyMany (ROX)
- ReadWriteMany (RWX)
- ReadWriteOncePod (RWOP)

Driver	Protocol	volumeMode	Access modes supported	File systems supported
google-cloud-netapp-volumes	NFS SMB	Filesystem	RWO, ROX, RWX, RWOP	nfs, smb

Configure a Trident NAS backend

```
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: gcnv-nas
  namespace: trident
spec:
  version: 1
  storageDriverName: google-cloud-netapp-volumes
  projectNumber: "<project-number>"
  location: "<region>"
  sdkTimeout: "600"
  storage:
    - labels:
        cloud: gcp
```

```
network: "<vpc-network>"
```

Provision NAS volumes

NAS volumes are provisioned using the `google-cloud-netapp-volumes` backend and support NFS and SMB protocols.

StorageClass for NFS volumes

To provision NFS volumes, set `nasType` to `nfs`.

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: gcnv-nfs
provisioner: csi.trident.netapp.io
parameters:
  backendType: "google-cloud-netapp-volumes"
  trident.netapp.io/nasType: "nfs"
allowVolumeExpansion: true
```

StorageClass for SMB volumes

To provision SMB volumes, set `nasType` to `smb` and provide credentials.

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: gcnv-smb
provisioner: csi.trident.netapp.io
parameters:
  backendType: "google-cloud-netapp-volumes"
  trident.netapp.io/nasType: "smb"
  csi.storage.k8s.io/node-stage-secret-name: "smbcreds"
  csi.storage.k8s.io/node-stage-secret-namespace: "default"
allowVolumeExpansion: true
```

PersistentVolumeClaim example (RWX)

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: gcnv-nas-rwx
spec:
```

```

accessModes:
  - ReadWriteMany
resources:
  requests:
    storage: 100Gi
storageClassName: gcnv-nfs

```

PersistentVolumeClaim example (RWO)

```

apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: gcnv-nas-rwo
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 100Gi
  storageClassName: gcnv-nfs

```



NAS volumes use `volumeMode: Filesystem`.

Configure Google Cloud NetApp Volumes for SAN workloads

You can configure Trident to provision block storage volumes using the iSCSI protocol from Google Cloud NetApp Volumes. SAN volumes are provisioned from Flex Unified storage pools by using the `google-cloud-netapp-volumes-san` storage driver.



This driver is dedicated to block workloads and does not support NAS protocols.



The `google-cloud-netapp-volumes-san` backend is required to provision iSCSI block volumes. The `google-cloud-netapp-volumes` backend supports NAS protocols only and cannot be used for SAN workloads.

Overview

Trident supports Google Cloud NetApp Volumes SAN (iSCSI) workloads using the `google-cloud-netapp-volumes-san` driver.

SAN volumes are provisioned from Flex Unified storage pools and presented to Kubernetes nodes as iSCSI block devices.

This applies to the following environments:

- Trident 26.02 and later
- Google Kubernetes Engine (GKE) or Red Hat OpenShift
- Google Cloud NetApp Volumes Flex Unified storage pools
- iSCSI-based workloads

Flex Unified storage pools

Flex Unified storage pools provide block storage using the iSCSI protocol and are required for SAN provisioning:

- Flex Unified REGIONAL pools are supported.
- Flex Unified ZONAL pools are supported starting with Trident 26.02.1.
- Only the **Flex** service level is supported for SAN workloads.

Configure a Trident SAN backend

```
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: gcnv-san
  namespace: trident
spec:
  version: 1
  storageDriverName: google-cloud-netapp-volumes-san
  projectNumber: "<project-number>"
  location: "<region>"
  sdkTimeout: "600"
  storage:
    - labels:
        cloud: gcp
        performance: flex
        network: "<vpc-network>"
        serviceLevel: Flex
```

Create a StorageClass

After configuring the SAN backend, create a StorageClass that references the google-cloud-netapp-volumes-san driver.

The filesystem type is defined in the StorageClass, not in the backend.

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: gcnv-san
provisioner: csi.trident.netapp.io
```

```
parameters:
  backendType: "google-cloud-netapp-volumes-san"
  fsType: "ext4"
allowVolumeExpansion: true
```

Supported filesystem types:

- ext4 (default)
- ext3
- xfs



The SAN driver supports only the Flex service level and does not use NAS-specific backend parameters such as `exportRule`, `unixPermissions`, `nasType`, `snapshotDir`, `nfsMountOptions`, or tiering-related settings.

Provision block volumes

ReadWriteOnce (RWO)

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: gcnv-san-rwo
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 100Gi
  storageClassName: gcnv-san
```

ReadWriteOncePod (RWOP)

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: gcnv-san-rwop
spec:
  accessModes:
    - ReadWriteOncePod
  resources:
    requests:
      storage: 100Gi
  storageClassName: gcnv-san
```

ReadOnlyMany (ROX)

A common pattern for ROX is to clone an existing ReadWriteOnce volume and mount the clone as read-only.

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: gcnv-san-rox
spec:
  accessModes:
    - ReadOnlyMany
  resources:
    requests:
      storage: 100Gi
  storageClassName: gcnv-san
  dataSource:
    kind: PersistentVolumeClaim
    name: gcnv-san-rwo
```

ReadWriteMany (RWX) — raw block only

ReadWriteMany is supported only when `volumeMode: Block`.

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: gcnv-san-raw-rwx
spec:
  accessModes:
    - ReadWriteMany
  volumeMode: Block
  resources:
    requests:
      storage: 100Gi
  storageClassName: gcnv-san
```

Block volume behavior

Block volumes are provisioned as iSCSI LUNs and presented to Kubernetes nodes as block devices.

Block volumes:

- Use the iSCSI protocol
- Support filesystem and raw block presentation
- Are attached and managed by Trident

- Support multiple Kubernetes access modes

Access modes

Block volumes provisioned by Trident support the following access modes:

- `ReadWriteOnce` (RWO)
- `ReadOnlyMany` (ROX)
- `ReadWriteOncePod` (RWOP)
- `ReadWriteMany` (RWX), supported only when `volumeMode: Block`

volumeMode behavior

The `volumeMode` field controls how a block volume is exposed:

- `Filesystem`
Trident formats and mounts the volume.
- `Block`
Trident attaches the device and exposes it as a raw block device.

Supported operations

Block volumes provisioned using the `google-cloud-netapp-volumes-san` driver support:

- Create
- Delete
- Clone
- Snapshot
- Resize
- Import

Extra GiB overprovisioning behavior

Google Cloud NetApp Volumes block volumes include internal metadata overhead. This overhead reduces the kernel-visible device size compared to the provisioned capacity.

Testing shows:

- Approximately 300 KiB overhead on initial creation
- Up to approximately 107 MiB overhead after a resize

Because Google Cloud NetApp Volumes accepts only whole-GiB allocations, Trident ensures that the usable device size always meets or exceeds the PVC request by:

- Rounding the requested size up to the next whole GiB
- Adding an additional 1 GiB buffer

Example:

- PVC request: 100 GiB
- Provisioned size in Google Cloud NetApp Volumes: 101 GiB
- Usable space visible to the application: at least 100 GiB

Pod examples

Filesystem-mounted block volume (RWO)

```
apiVersion: v1
kind: Pod
metadata:
  name: app-rwo
spec:
  containers:
  - name: app
    image: ubuntu:22.04
    command: ["sleep", "infinity"]
    volumeMounts:
    - name: data
      mountPath: /mnt/data
  volumes:
  - name: data
    persistentVolumeClaim:
      claimName: gcnv-san-rwo
```

Raw block device (RWX)

```
apiVersion: v1
kind: Pod
metadata:
  name: app-raw-rwx
spec:
  containers:
  - name: app
    image: ubuntu:22.04
    command: ["sleep", "infinity"]
    volumeDevices:
    - name: data
      devicePath: /dev/xda
  volumes:
  - name: data
    persistentVolumeClaim:
      claimName: gcnv-san-raw-rwx
```

Attach and mount behavior

For SAN volumes provisioned from Google Cloud NetApp Volumes:

- Trident creates a Logical Unit Number (LUN) in a Flex Unified storage pool.
- During publish, Trident maps the LUN to a per-node host group.
- During node staging, Trident:
 - Logs in to the iSCSI target
 - Discovers the LUN
 - Configures multipath
- If `volumeMode: Filesystem`, Trident formats the device if required and mounts it.
- If `volumeMode: Block`, Trident attaches the device and exposes it directly to the pod without formatting or mounting.



SAN block volumes do not provide distributed locking or write coordination. When a block volume is accessed by multiple nodes (ReadWriteMany with `volumeMode: Block`), the application or filesystem must manage concurrency.

Prepare to configure a Google Cloud NetApp Volumes backend

Before you can configure your Google Cloud NetApp Volumes backend, you need to ensure the following requirements are met.

Prerequisites for NFS or SMB volumes

If you are using Google Cloud NetApp Volumes for the first time or in a new location, some initial configuration is required to set up Google Cloud NetApp Volumes and create an NFS or SMB volume. Refer to [Before you begin](#).

Ensure that you have the following before configuring Google Cloud NetApp Volumes backend:

- A Google Cloud account configured with Google Cloud NetApp Volumes service. Refer to [Google Cloud NetApp Volumes](#).
- Project number of your Google Cloud account. Refer to [Identifying projects](#).
- A Google Cloud service account with the NetApp Volumes Admin (`roles/netapp.admin`) role. Refer to [Identity and Access Management roles and permissions](#).
- API key file for your GCNV account. Refer to [Create a service account key](#)
- A storage pool. Refer to [Storage pools overview](#).

For more information about how to set up access to Google Cloud NetApp Volumes, refer to [Set up access to Google Cloud NetApp Volumes](#).

Google Cloud NetApp Volumes backend configuration options and examples

Learn about backend configuration options for Google Cloud NetApp Volumes and review configuration examples.

Backend configuration options

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

Parameter	Description	Default
version		Always 1
storageDriverName	Name of the storage driver	The value of <code>storageDriverName</code> must be specified as "google-cloud-netapp-volumes".
backendName	(Optional) Custom name of the storage backend	Driver name + "_" + part of API key
storagePools	Optional parameter used to specify storage pools for volume creation.	
projectNumber	Google Cloud account project number. The value is found on the Google Cloud portal home page.	
location	<p>The Google Cloud location where Trident creates GCNV volumes. When creating cross-region Kubernetes clusters, volumes created in a <code>location</code> can be used in workloads scheduled on nodes across multiple Google Cloud regions.</p> <p>Cross-region traffic incurs an additional cost.</p>	
apiKey	<p>API key for the Google Cloud service account with the <code>netapp.admin</code> role.</p> <p>It includes the JSON-formatted contents of a Google Cloud service account's private key file (copied verbatim into the backend configuration file).</p> <p>The <code>apiKey</code> must include key-value pairs for the following keys: <code>type</code>, <code>project_id</code>, <code>client_email</code>, <code>client_id</code>, <code>auth_uri</code>, <code>token_uri</code>, <code>auth_provider_x509_cert_url</code>, and <code>client_x509_cert_url</code>.</p>	
nfsMountOptions	Fine-grained control of NFS mount options.	"nfsvers=3"
limitVolumeSize	Fail provisioning if the requested volume size is above this value.	"" (not enforced by default)
serviceLevel	The service level of a storage pool and its volumes. The values are <code>flex</code> , <code>standard</code> , <code>premium</code> , or <code>extreme</code> .	
labels	Set of arbitrary JSON-formatted labels to apply on volumes	""
network	Google Cloud network used for GCNV volumes.	

Parameter	Description	Default
<code>debugTraceFlags</code>	<p>Debug flags to use when troubleshooting. Example, <code>{"api":false, "method":true}</code>.</p> <p>Do not use this unless you are troubleshooting and require a detailed log dump.</p>	<code>null</code>
<code>nasType</code>	<p>Configure NFS or SMB volumes creation.</p> <p>Options are <code>nfs</code>, <code>smb</code> or <code>null</code>. Setting to <code>null</code> defaults to NFS volumes.</p>	<code>nfs</code>
<code>supportedTopologies</code>	<p>Represents a list of regions and zones that are supported by this backend.</p> <p>For more information, refer to Use CSI Topology. For example:</p> <pre>supportedTopologies: - topology.kubernetes.io/region: asia-east1 topology.kubernetes.io/zone: asia-east1-a</pre>	

Volume provisioning options

You can control default volume provisioning in the `defaults` section of the configuration file.

Parameter	Description	Default
<code>exportRule</code>	The export rules for new volumes. Must be a comma-separated list of any combination of IPv4 addresses.	<code>"0.0.0.0/0"</code>
<code>snapshotDir</code>	Access to the <code>.snapshot</code> directory	<code>true</code> , <code>false</code> (default behavior might vary. Set explicitly) <code>"false"</code> for NFSv3
<code>snapshotReserve</code>	Percentage of volume reserved for snapshots	<code>""</code> (accept default of 0)
<code>unixPermissions</code>	The unix permissions of new volumes (4 octal digits).	<code>""</code>

Example configurations

The following examples show basic configurations that leave most parameters to default. This is the easiest way to define a backend.

▼ Minimal configuration

This is the absolute minimum backend configuration. With this configuration, Trident discovers all of your storage pools delegated to Google Cloud NetApp Volumes in the configured location, and places new volumes on one of those pools randomly. Because `nasType` is omitted, the `nfs` default applies and the backend will provision for NFS volumes.

This configuration is ideal when you are just getting started with Google Cloud NetApp Volumes and trying things out, but in practice you might need to provide additional scoping for the volumes you provision.



Replace <id_value> and <key_value> with your service account credentials.

```
---
apiVersion: v1
kind: Secret
metadata:
  name: backend-tbc-gcnv-secret
type: Opaque
stringData:
  private_key_id: "<id_value>"
  private_key: |
    -----BEGIN PRIVATE KEY-----
    <key_value>
    -----END PRIVATE KEY-----
---
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: backend-tbc-gcnv
spec:
  version: 1
  storageDriverName: google-cloud-netapp-volumes
  projectNumber: "123455380079"
  location: europe-west6
  serviceLevel: premium
  apiKey:
    type: service_account
    project_id: my-gcnv-project
    client_email: myproject-prod@my-gcnv-project.iam.gserviceaccount.com
    client_id: "103346282737811234567"
    auth_uri: https://accounts.google.com/o/oauth2/auth
    token_uri: https://oauth2.googleapis.com/token
    auth_provider_x509_cert_url:
https://www.googleapis.com/oauth2/v1/certs
    client_x509_cert_url:
https://www.googleapis.com/robot/v1/metadata/x509/myproject-prod%40my-
gcnv-project.iam.gserviceaccount.com
  credentials:
    name: backend-tbc-gcnv-secret
```

▼ Configuration for SMB volumes

```

apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: backend-tbc-gcnv1
  namespace: trident
spec:
  version: 1
  storageDriverName: google-cloud-netapp-volumes
  projectNumber: "123456789"
  location: asia-east1
  serviceLevel: flex
  nasType: smb
  apiKey:
    type: service_account
    project_id: cloud-native-data
    client_email: trident-sample@cloud-native-
data.iam.gserviceaccount.com
    client_id: "123456789737813416734"
    auth_uri: https://accounts.google.com/o/oauth2/auth
    token_uri: https://oauth2.googleapis.com/token
    auth_provider_x509_cert_url:
https://www.googleapis.com/oauth2/v1/certs
    client_x509_cert_url:
https://www.googleapis.com/robot/v1/metadata/x509/trident-
sample%40cloud-native-data.iam.gserviceaccount.com
  credentials:
    name: backend-tbc-gcnv-secret

```

▼ Configuration with StoragePools filter

```

---
apiVersion: v1
kind: Secret
metadata:
  name: backend-tbc-gcnv-secret
type: Opaque
stringData:
  private_key_id: "<id_value>"
  private_key: |
    -----BEGIN PRIVATE KEY-----
    <key_value>
    -----END PRIVATE KEY-----
---

```

```

apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: backend-tbc-gcnv
spec:
  version: 1
  storageDriverName: google-cloud-netapp-volumes
  projectNumber: "123455380079"
  location: europe-west6
  serviceLevel: premium
  storagePools:
    - premium-pool1-europe-west6
    - premium-pool2-europe-west6
  apiKey:
    type: service_account
    project_id: my-gcnv-project
    client_email: myproject-prod@my-gcnv-project.iam.gserviceaccount.com
    client_id: "103346282737811234567"
    auth_uri: https://accounts.google.com/o/oauth2/auth
    token_uri: https://oauth2.googleapis.com/token
    auth_provider_x509_cert_url:
https://www.googleapis.com/oauth2/v1/certs
    client_x509_cert_url:
https://www.googleapis.com/robot/v1/metadata/x509/myproject-prod%40my-
gcnv-project.iam.gserviceaccount.com
  credentials:
    name: backend-tbc-gcnv-secret

```

▼ Virtual pool configuration

This backend configuration defines multiple virtual pools in a single file. Virtual pools are defined in the storage section. They are useful when you have multiple storage pools supporting different service levels and you want to create storage classes in Kubernetes that represent those. Virtual pool labels are used to differentiate the pools. For instance, in the example below performance label and serviceLevel type is used to differentiate virtual pools.

You can also set some default values to be applicable to all virtual pools, and overwrite the default values for individual virtual pools. In the following example, snapshotReserve and exportRule serve as defaults for all virtual pools.

For more information, refer to [Virtual pools](#).

```

---
apiVersion: v1
kind: Secret
metadata:
  name: backend-tbc-gcnv-secret
type: Opaque

```

```

stringData:
  private_key_id: "<id_value>"
  private_key: |
    -----BEGIN PRIVATE KEY-----
    <key_value>
    -----END PRIVATE KEY-----

---
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: backend-tbc-gcnv
spec:
  version: 1
  storageDriverName: google-cloud-netapp-volumes
  projectNumber: "123455380079"
  location: europe-west6
  apiKey:
    type: service_account
    project_id: my-gcnv-project
    client_email: myproject-prod@my-gcnv-project.iam.gserviceaccount.com
    client_id: "103346282737811234567"
    auth_uri: https://accounts.google.com/o/oauth2/auth
    token_uri: https://oauth2.googleapis.com/token
    auth_provider_x509_cert_url:
https://www.googleapis.com/oauth2/v1/certs
    client_x509_cert_url:
https://www.googleapis.com/robot/v1/metadata/x509/myproject-prod%40my-
gcnv-project.iam.gserviceaccount.com
  credentials:
    name: backend-tbc-gcnv-secret
  defaults:
    snapshotReserve: "10"
    exportRule: 10.0.0.0/24
  storage:
    - labels:
        performance: extreme
        serviceLevel: extreme
      defaults:
        snapshotReserve: "5"
        exportRule: 0.0.0.0/0
    - labels:
        performance: premium
        serviceLevel: premium
    - labels:
        performance: standard

```

```
serviceLevel: standard
```

▼ Cloud identity for GKE

When you install Trident with a GKE cloud identity, backends authenticate as a workload identity and omit the `apiKey` block. Cloud identity is configured the same way for both Google Cloud NetApp Volumes drivers. The only difference is the `storageDriverName` value: use `google-cloud-netapp-volumes` for NAS (NFS or SMB) volumes and `google-cloud-netapp-volumes-san` for SAN (iSCSI block) volumes.



Cloud identity support for the `google-cloud-netapp-volumes-san` driver requires Trident 26.06 or later. SAN volumes are provisioned from Flex Unified storage pools, so set `serviceLevel` to `flex`.

NAS driver (google-cloud-netapp-volumes)

```
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: backend-tbc-gcp-gcnv
spec:
  version: 1
  storageDriverName: google-cloud-netapp-volumes
  projectNumber: '012345678901'
  network: gcnv-network
  location: us-west2
  serviceLevel: Premium
  storagePool: pool-premium1
```

SAN driver (google-cloud-netapp-volumes-san)

```
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: backend-tbc-gcp-gcnv-san
spec:
  version: 1
  storageDriverName: google-cloud-netapp-volumes-san
  projectNumber: '012345678901'
  network: gcnv-network
  location: us-west2
  serviceLevel: flex
  storagePool: pool-flex1
```

▼ Supported topologies configuration

Trident facilitates provisioning of volumes for workloads based on regions and availability zones. The `supportedTopologies` block in this backend configuration is used to provide a list of regions and zones

per backend. The region and zone values specified here must match the region and zone values from the labels on each Kubernetes cluster node. These regions and zones represent the list of permissible values that can be provided in a storage class. For storage classes that contain a subset of the regions and zones provided in a backend, Trident creates volumes in the mentioned region and zone. For more information, refer to [Use CSI Topology](#).

```
---
version: 1
storageDriverName: google-cloud-netapp-volumes
subscriptionID: 9f87c765-4774-fake-ae98-a721add45451
tenantID: 68e4f836-edc1-fake-bff9-b2d865ee56cf
clientID: dd043f63-bf8e-fake-8076-8de91e5713aa
clientSecret: SECRET
location: asia-east1
serviceLevel: flex
supportedTopologies:
  - topology.kubernetes.io/region: asia-east1
    topology.kubernetes.io/zone: asia-east1-a
  - topology.kubernetes.io/region: asia-east1
    topology.kubernetes.io/zone: asia-east1-b
```

What's next?

After you create the backend configuration file, run the following command:

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

To verify that the backend is successfully created, run the following command:

```
kubectl get tridentbackendconfig
```

NAME	BACKEND NAME	BACKEND UUID
PHASE STATUS		
backend-tbc-gcnv	backend-tbc-gcnv	b2fd1ff9-b234-477e-88fd-713913294f65
Bound Success		

If the backend creation fails, something is wrong with the backend configuration. You can describe the backend using the `kubectl get tridentbackendconfig <backend-name>` command or view the logs to determine the cause by running the following command:

```
tridentctl logs
```

After you identify and correct the problem with the configuration file, you can delete the backend and run the create command again.

Storage class definitions

The following is a basic `StorageClass` definition that refers to the backend above.

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: gcnv-nfs-sc
provisioner: csi.trident.netapp.io
parameters:
  backendType: "google-cloud-netapp-volumes"
```

Example definitions using the `parameter.selector` field:

Using `parameter.selector` you can specify for each `StorageClass` the [virtual pool](#) that is 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: extreme-sc
provisioner: csi.trident.netapp.io
parameters:
  selector: performance=extreme
  backendType: google-cloud-netapp-volumes

---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: premium-sc
provisioner: csi.trident.netapp.io
parameters:
  selector: performance=premium
  backendType: google-cloud-netapp-volumes

---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: standard-sc
provisioner: csi.trident.netapp.io
parameters:
  selector: performance=standard
  backendType: google-cloud-netapp-volumes
```

For more details on storage classes, refer to [Create a storage class](#).

Example definitions for SMB volumes

Using `nasType`, `node-stage-secret-name`, and `node-stage-secret-namespace`, you can specify an SMB volume and provide the required Active Directory credentials. Any Active Directory user/password with any/no permissions can be used for the node stage secret.

▼ Basic configuration on default namespace

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: gcnv-sc-smb
provisioner: csi.trident.netapp.io
parameters:
  backendType: "google-cloud-netapp-volumes"
  trident.netapp.io/nasType: "smb"
  csi.storage.k8s.io/node-stage-secret-name: "smbcreds"
  csi.storage.k8s.io/node-stage-secret-namespace: "default"
```

▼ Using different secrets per namespace

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: gcnv-sc-smb
provisioner: csi.trident.netapp.io
parameters:
  backendType: "google-cloud-netapp-volumes"
  trident.netapp.io/nasType: "smb"
  csi.storage.k8s.io/node-stage-secret-name: "smbcreds"
  csi.storage.k8s.io/node-stage-secret-namespace: ${pvc.namespace}
```

▼ Using different secrets per volume

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: gcnv-sc-smb
provisioner: csi.trident.netapp.io
parameters:
  backendType: "google-cloud-netapp-volumes"
  trident.netapp.io/nasType: "smb"
  csi.storage.k8s.io/node-stage-secret-name: ${pvc.name}
  csi.storage.k8s.io/node-stage-secret-namespace: ${pvc.namespace}
```



nasType: smb filters for pools which support SMB volumes. nasType: nfs or nasType: null filters for NFS pools.

PVC definition example

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: gcnv-nfs-pvc
spec:
  accessModes:
    - ReadWriteMany
  resources:
    requests:
      storage: 100Gi
  storageClassName: gcnv-nfs-sc
```

To verify if the PVC is bound, run the following command:

```
kubectl get pvc gcnv-nfs-pvc
```

NAME	STATUS	VOLUME	CAPACITY
ACCESS MODES	STORAGECLASS	AGE	
gcnv-nfs-pvc	Bound	pvc-b00f2414-e229-40e6-9b16-ee03eb79a213	100Gi
RWX	gcnv-nfs-sc	1m	

Configure auto-tiering for Google Cloud NetApp Volumes

Auto-tiering is configured through Trident backend parameters and PersistentVolumeClaim annotations during volume provisioning. You can configure auto-tiering for Google Cloud NetApp Volumes using Trident.

Overview

Auto-tiering allows Trident to provision volumes that automatically move inactive data from a performance tier to a capacity tier.

This reduces storage cost while preserving performance for frequently accessed data.

Trident applies auto-tiering settings only at volume creation time. Post-provisioning changes are not supported in Trident 26.06.

Concepts

Auto-tiering

Auto-tiering moves infrequently accessed data from a performance tier to a capacity tier based on access

patterns.

Data movement occurs asynchronously and is not immediate.

Tiering policy

The tiering policy determines whether auto-tiering is enabled for a volume.

The following policies are supported:

- * `auto`: Enables automatic tiering based on access patterns
- * `none`: Disables auto-tiering

Cooling days

Cooling days specify the minimum number of days a block of data must remain inactive before it becomes eligible for tiering.

Cooling days apply only when the tiering policy is set to `auto`.

Configuration model

Configuration scopes

Auto-tiering can be configured at multiple scopes:

- **Storage pool scope**
Applies to all volumes provisioned from the pool.
- **Volume scope**
Applies to a single volume through `PersistentVolumeClaim` annotations.

Trident determines the effective configuration based on where each setting is defined.

Configuration precedence

When the same setting is defined at multiple scopes, Trident applies the following precedence order:

1. `PersistentVolumeClaim` annotations
2. Trident backend configuration
3. Storage pool defaults

Settings defined at a higher precedence override lower-level values.

Supported functionality in Trident 26.06

Trident 26.06 supports the following auto-tiering capabilities for Google Cloud NetApp Volumes:

- Enabling or disabling auto-tiering during volume provisioning
- Defining a tiering policy in the Trident backend configuration
- Overriding the tiering policy and cooling days per volume using PVC annotations
- Configuring cooling days for volumes with auto-tiering enabled

Unsupported functionality in Trident 26.06

The following operations are not supported:

- Modifying auto-tiering settings after volume creation
- Changing tiering policies on existing volumes using Kubernetes updates
- Applying auto-tiering settings outside of Trident-managed provisioning workflows

Backend configuration parameters

The following parameters control auto-tiering behavior when defined in the Trident backend configuration:

Parameter	Required	Description
tieringPolicy	No	Tiering policy for volumes (<code>auto</code> or <code>none</code>)
tieringMinimumCoolingDays	No	Number of inactive days before data is tiered (range: 2–183, default: 31)

Volume-level overrides using PersistentVolumeClaim annotations

Supported annotations

PersistentVolumeClaim annotations allow per-volume overrides of auto-tiering settings.

Annotation	Description
<code>trident.netapp.io/tieringPolicy</code>	Overrides the tiering policy for the volume
<code>trident.netapp.io/tieringMinimumCoolingDays</code>	Overrides the cooling days value for the volume

Example: PersistentVolumeClaim with auto-tiering overrides

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: auto-tiering-pvc
  annotations:
    trident.netapp.io/tieringPolicy: auto
    trident.netapp.io/tieringMinimumCoolingDays: "45"
spec:
  accessModes:
    - ReadWriteOnce
  storageClassName: google-cloud-netapp-volumes-auto-tiering
  resources:
    requests:
      storage: 500Gi
```

Behavior and limitations

Provisioning behavior

- Auto-tiering settings are evaluated and applied only at volume creation time.
- Trident does not reconcile tiering configuration after provisioning.
- Cooling days are ignored when the tiering policy is set to `none`.

Platform limitations

- Auto-tiering is supported only for NAS volumes (NFS and SMB).
- Block volumes (iSCSI) do not support auto-tiering.
- The Google Cloud NetApp Volumes storage pool must have auto-tiering enabled in Google Cloud.

Supported values

- Valid range for `tieringMinimumCoolingDays`: 2 to 183
- Default value: 31

Configure Google Cloud NetApp Volumes ONTAP-mode storage pools

Beginning with Trident 26.06, you can configure existing `ontap-san` and `ontap-nas` backends to provision storage on Google Cloud NetApp Volumes (GCNV) ONTAP-mode storage pools. This feature is a Tech Preview.



This feature is a Tech Preview in Trident 26.06. Tech Preview features are not supported for production use. The functionality, configuration fields, and supported matrix can change before general availability (GA).

How it works

Trident reuses the existing `ontap-san` and `ontap-nas` driver logic and routes ONTAP REST calls through a GCNV proxy endpoint. This approach lets you provision and manage storage on GCNV-hosted ONTAP clusters while keeping the same operational workflows that you use for a direct ONTAP backend. ONTAP-mode does not introduce a new storage driver. You enable it per backend by adding a `gcnv` configuration block to an `ontap-san` or `ontap-nas` backend.

The provisioning path is as follows:

PersistentVolumeClaim → Trident (`ontap-san` or `ontap-nas`) → ONTAP REST client → GCNV proxy → ONTAP cluster in the GCNV pool

Supported and unsupported scope

This Tech Preview supports the following:

- Drivers: `ontap-san` (iSCSI) and `ontap-nas` (NFS or SMB).
- Lifecycle operations through the proxy: backend initialization, volume creation and deletion, and related ONTAP REST flows.

- Authentication models: Workload Identity Pool (WIP), service account key, and Application Default Credentials (ADC) fallback.

The following are out of scope for this Tech Preview:

- The `ontap-san-economy`, `ontap-nas-economy`, and `ontap-nas-flexgroup` drivers.
- The ASA r2 personality path.
- ZAPI fallback. ONTAP-mode uses ONTAP REST only.

Prerequisites

Before you configure an ONTAP-mode backend, ensure that you have the following:

- Trident 26.06 or later.
- A GCNV ONTAP-mode storage pool in your target project and location.
- GCP IAM permissions for proxy operations, scoped appropriately for your environment.
- A backend configuration that includes `proxyURL`, `projectNumber`, `location`, and `poolID`.
- For iSCSI SAN workflows, the node-side iSCSI and multipath prerequisites for your platform. Refer to [Prepare the worker node](#).

Configure the backend

A `TridentBackendConfig` requires credentials. For ONTAP-mode, that secret holds GCP proxy credentials—not ONTAP managementLIF / username / password. `svm` is optional: each GCNV ONTAP-mode pool currently has one SVM, and Trident derives it from the pool when omitted.

Use the same service account pattern as native GCNV: non-sensitive fields in `gcnv.apiKey`, `private_key` and `private_key_id` in the secret. See [Google Cloud NetApp Volumes configuration examples](#).

ontap-san backend

```
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: gcnv-ontap-san
  namespace: trident
spec:
  version: 1
  storageDriverName: ontap-san
  backendName: gcnv-ontap-san
  credentials:
    name: gcnv-sa-secret
    type: secret
  gcnv:
    proxyURL: "https://netapp.googleapis.com"
    projectNumber: "<project-number>"
    location: "<region-or-zone>"
    poolID: "<pool-id>"
```

```
svm: "<svm-name>"
```

ontap-nas backend

```
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: gcnv-ontap-nas
  namespace: trident
spec:
  version: 1
  storageDriverName: ontap-nas
  backendName: gcnv-ontap-nas
  credentials:
    name: gcnv-sa-secret
    type: secret
  gcnv:
    proxyURL: "https://netapp.googleapis.com"
    projectNumber: "<project-number>"
    location: "<region-or-zone>"
    poolID: "<pool-id>"
    apiKey:
      type: service_account
      project_id: "<project-id>"
      client_email: "<service-account-email>"
      client_id: "<client-id>"
      auth_uri: "https://accounts.google.com/o/oauth2/auth"
      token_uri: "https://oauth2.googleapis.com/token"
      auth_provider_x509_cert_url:
"https://www.googleapis.com/oauth2/v1/certs"
      client_x509_cert_url: "<client-x509-cert-url>"
```

Secret

```
apiVersion: v1
kind: Secret
metadata:
  name: gcnv-sa-secret
  namespace: trident
type: Opaque
stringData:
  private_key_id: "<private-key-id>"
  private_key: |
    -----BEGIN PRIVATE KEY-----
    <private-key>
```

-----END PRIVATE KEY-----



Do not put `private_key` or `private_key_id` in the `TridentBackendConfig` spec—the CRD rejects them. Use `ontap-san` with the same `gcnv` and credential pattern for block storage.

```
svm: "<svm-name>"
```

== Authentication

Trident resolves credentials for proxy access in the following order:

- . Workload Identity Pool (WIP)
- . Service account key
- . Application Default Credentials (ADC)

For this Tech Preview, use the following practices:

- * Prefer Workload Identity Pool where it is available.
- * Do not embed raw private keys in version-controlled backend files. Store credentials in a Kubernetes secret.
- * Apply least-privilege IAM scoping to the service account.

== Validation and initialization behavior

Trident validates an ONTAP-mode backend during initialization and fails fast rather than partially initializing.

Backend initialization fails if a required ``gcnv`` field is missing, the driver is unsupported for ONTAP-mode, or credential resolution fails.

== Known limitations

- * This feature is a Tech Preview. Behavior and the supported matrix can change before GA.
- * ZAPI fallback is not used in ONTAP-mode.
- * Protocol and node-readiness requirements still apply. For example, iSCSI workflows require the node-side prerequisites described in `xref:{relative_path}../trident-use/worker-node-prep.html` [Prepare the worker node].
- * Existing ONTAP and GCNV environment constraints still apply.

== Upgrade and compatibility

ONTAP-mode is opt-in per backend through the `gcnv` configuration block:

- * Existing non-GCNV ONTAP backends are unaffected.
- * Mixed deployments that combine direct ONTAP backends and ONTAP-mode proxy backends are supported through backend-level configuration. Each backend is configured independently.
- * To stop using ONTAP-mode, remove or replace the affected backends. No global switch is required.

== What's next?

Apply secret before TBC; use `kubectl apply` and `-n trident` to match examples.

```
:leveloffset: -1
```

```
[[ID81190148e4d4e430d9663c5cb39ced5e]]
= Configure a NetApp HCI or SolidFire backend
:hardbreaks:
:icons: font
:relative_path: ./trident-use/
:imagesdir: {root_path}{relative_path}../media/
```

[.lead]

Learn how to create and use an Element backend with your Trident installation.

== Element driver details

Trident provides the `solidfire-san` storage driver to communicate with the cluster. Supported access modes are: `_ReadWriteOnce_` (RWO), `_ReadOnlyMany_` (ROX), `_ReadWriteMany_` (RWX), `_ReadWriteOncePod_` (RWOP).

The `solidfire-san` storage driver supports `_file_` and `_block_` volume modes. For the `Filesystem` volumeMode, Trident creates a volume and creates a filesystem. The filesystem type is specified by the `StorageClass`.

```
[cols=5,options="header"]
|===
|Driver
|Protocol
|VolumeMode
|Access modes supported
```

|File systems supported

|`solidfire-san`

a|iSCSI

a|Block

a|RWO, ROX, RWX, RWOP

a|No Filesystem. Raw block device.

|`solidfire-san`

a|iSCSI

a|Filesystem

a|RWO, RWOP

a|`xfs`, `ext3`, `ext4`

|===

== Before you begin

You'll need the following before creating an Element backend.

- * A supported storage system that runs Element software.
- * Credentials to a NetApp HCI/SolidFire cluster admin or tenant user that can manage volumes.
- * All of your Kubernetes worker nodes should have the appropriate iSCSI tools installed. Refer to [xref:{relative_path}../trident-use/worker-node-prep.html](#) [worker node preparation information].

== Backend configuration options

See the following table for the backend configuration options:

[cols=3,options="header"]

|===

|Parameter |Description |Default

|`version` | |Always 1

|`storageDriverName` | Name of the storage driver |Always "solidfire-san"

|`backendName` |Custom name or the storage backend |"solidfire_" +
storage (iSCSI) IP address

|`Endpoint` |MVIP for the SolidFire cluster with tenant credentials |

|`SVIP` |Storage (iSCSI) IP address and port |

|`labels` |Set of arbitrary JSON-formatted labels to apply on volumes. |""

```
|`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. Trident uses CHAP. |true

|`AccessGroups` |List of Access Group IDs to use |Finds the ID of an
access group named "trident"

| `Types` |QoS specifications |

|`limitVolumeSize` |Fail provisioning if requested volume size is above
this value |"" (not enforced by default)

| `debugTraceFlags` |Debug flags to use when troubleshooting. Example,
{"api":false, "method":true} |null

|===

WARNING: Do not use `debugTraceFlags` unless you are troubleshooting and
require a detailed log dump.

== 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.

[source,yaml]
```

```
version: 1
storageDriverName: solidfire-san
Endpoint: https://<user>:<password>@<mvip>/json-rpc/8.0
SVIP: <svip>:3260
TenantName: <tenant>
labels:
k8scluster: dev1
backend: dev1-element-cluster
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 pools

This example shows the backend definition file configured with virtual pools along with StorageClasses that refer back to them.

Trident copies labels present on a storage pool to the backend storage LUN at provisioning. For convenience, storage administrators can define labels per virtual pool and group volumes by label.

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

[source,yaml]

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
type: Silver
labels:
store: solidfire
k8scluster: dev-1-cluster
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 pools. Using the `parameters.selector` field, each StorageClass calls out which virtual pool(s) can 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 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 pool is selected and ensures the storage requirement is met.

[source,yaml]

```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
name: solidfire-gold-four
provisioner: csi.trident.netapp.io
parameters:
selector: performance=gold; cost=4
fsType: ext4

```

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: solidfire-silver-three
  provisioner: csi.trident.netapp.io
  parameters:
    selector: performance=silver; cost=3
    fsType: ext4
```

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: solidfire-bronze-two
  provisioner: csi.trident.netapp.io
  parameters:
    selector: performance=bronze; cost=2
    fsType: ext4
```

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: solidfire-silver-one
  provisioner: csi.trident.netapp.io
  parameters:
    selector: performance=silver; cost=1
    fsType: ext4
```

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: solidfire-silver
  provisioner: csi.trident.netapp.io
  parameters:
    selector: performance=silver
    fsType: ext4
```

```
== Find more information
```

```
* xref:{relative_path}../trident-concepts/vol-access-groups.html[Volume
access groups^]
```

```
= ONTAP SAN drivers
```

```
:leveloffset: +1
```

```
[[IDa54d6b2c3978aa8048489900af7dd15b]]
```

```
= ONTAP SAN driver overview
```

```
:hardbreaks:
```

```
:icons: font
```

```
:relative_path: ./trident-use/
```

```
:imagesdir: {root_path}{relative_path}../media/
```

```
[.lead]
```

Learn about configuring an ONTAP backend with ONTAP and Cloud Volumes ONTAP SAN drivers.

```
== ONTAP SAN driver details
```

Trident provides the following SAN storage drivers to communicate with the ONTAP cluster. Supported access modes are: `_ReadWriteOnce_` (RWO), `_ReadOnlyMany_` (ROX), `_ReadWriteMany_` (RWX), `_ReadWriteOncePod_` (RWOP).

//IMPORTANT: If you are using Astra Control for protection, recovery, and mobility, read <<Astra Control driver compatibility>>.

```
[cols="2, 1, 1, 2, 2", options="header"]
```

```
|==
```

```
|Driver
```

```
|Protocol
```

```
|volumeMode
```

```
|Access modes supported
```

```
|File systems supported
```

```
|`ontap-san`
```

```
a|iSCSI
```

```
SCSI over FC
```

```
a|Block
```

```
a|RWO, ROX, RWX, RWOP
```

```
a|No filesystem; raw block device
```

```
|`ontap-san`
```

```
a|iSCSI
```

```
SCSI over FC
```

```
a|Filesystem
```

```
a|RWO, RWOP
```

ROX and RWX are not available in Filesystem volume mode.

```
a|`xfs`, `ext3`, `ext4`
```

```
|`ontap-san`
```

```
a|NVMe/TCP
```

Refer to <<Additional considerations for NVMe/TCP>>.

```
a|Block
```

```
a|RWO, ROX, RWX, RWOP
```

```
a|No filesystem; raw block device
```

```
|`ontap-san`
```

```
a|NVMe/TCP
```

Refer to <<Additional considerations for NVMe/TCP>>.

```
a|Filesystem
```

```
a|RWO, RWOP
```

ROX and RWX are not available in Filesystem volume mode.

```
a|`xfs`, `ext3`, `ext4`
```

```
|`ontap-san-economy`
```

```
a|iSCSI
```

```
a|Block
```

```
a|RWO, ROX, RWX, RWOP
```

```
a|No filesystem; raw block device
```

```
|`ontap-san-economy`
```

```
a|iSCSI
```

```
a|Filesystem
```

```
a|RWO, RWOP
```

ROX and RWX are not available in Filesystem volume mode.

```
a|`xfs`, `ext3`, `ext4`
```

```
|===
```

```
//=== Astra Control driver compatibility
```

```
//Astra Control provides seamless protection, disaster recovery, and  
mobility (moving volumes between Kubernetes clusters) for volumes created  
with the `ontap-nas`, `ontap-nas-flexgroup`, and `ontap-san` drivers.  
Refer to link:https://docs.netapp.com/us-en/astra-control-  
center/use/replicate\_snapmirror.html#replication-prerequisites[Astra  
Control replication prerequisites^] for details.
```

```
[WARNING]
```

```
=====
```

```
* Use `ontap-san-economy` only if persistent volume usage count is
```

expected to be higher than [link:https://docs.netapp.com/us-en/ontap/volumes/storage-limits-reference.html](https://docs.netapp.com/us-en/ontap/volumes/storage-limits-reference.html)[supported ONTAP volume limits^].

* Use `ontap-nas-economy` only if persistent volume usage count is expected to be higher than [link:https://docs.netapp.com/us-en/ontap/volumes/storage-limits-reference.html](https://docs.netapp.com/us-en/ontap/volumes/storage-limits-reference.html)[supported ONTAP volume limits^] and the `ontap-san-economy` driver cannot be used.

* Do not use use `ontap-nas-economy` if you anticipate the need for data protection, disaster recovery, or mobility.

* NetApp does not recommend using Flexvol autogrow in all ONTAP drivers, except ontap-san. As a workaround, Trident supports the use of snapshot reserve and scales Flexvol volumes accordingly.

====

== 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. For Amazon FSx for NetApp ONTAP deployments, Trident expects to be run as either an ONTAP or SVM administrator, using the cluster `fsxadmin` user or a `vsadmin` SVM user, or a user with a different name that has the same role. The `fsxadmin` user is a limited replacement for the cluster admin user.

NOTE: If you use the `limitAggregateUsage` parameter, cluster admin permissions are required. When using Amazon FSx for NetApp ONTAP with Trident, the `limitAggregateUsage` parameter will not work with the `vsadmin` and `fsxadmin` user accounts. The configuration operation will fail if you specify this parameter.

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.

== Additional considerations for NVMe/TCP

Trident supports the non-volatile memory express (NVMe) protocol using the `ontap-san` driver including:

- * IPv6
- * Snapshots and clones of NVMe volumes
- * Resizing an NVMe volume
- * Importing an NVMe volume that was created outside of Trident so that its lifecycle can be managed by Trident
- * NVMe-native multipathing

* Graceful or ungraceful shutdown of the K8s nodes (24.06)

Trident does not support:

- * DH-HMAC-CHAP that is natively supported by NVMe
- * Device mapper (DM) multipathing
- * LUKS encryption

NOTE: NVMe is supported only with ONTAP REST APIs and not supported with ONTAPI (ZAPI).

```
[[ID6c1b5041023d13177621d1519be66cba]]  
= Prepare to configure backend with ONTAP SAN drivers  
:hardbreaks:  
:icons: font  
:relative_path: ./trident-use/  
:imagesdir: {root_path}{relative_path}../media/
```

[.lead]

Understand the requirements and authentication options for configuring an ONTAP backend with ONTAP SAN drivers.

== Requirements

For all ONTAP backends, Trident requires at least one aggregate be assigned to the SVM.

NOTE: [link:https://docs.netapp.com/us-en/asa-r2/get-started/learn-about.html](https://docs.netapp.com/us-en/asa-r2/get-started/learn-about.html) [ASA r2 systems^] differ from other ONTAP systems (ASA, AFF, and FAS) in the implementation of their storage layer. In ASA r2 systems, storage availability zones are used instead of aggregates. Refer to [link:https://kb.netapp.com/on-prem/ASAr2/ASAr2_KBs/su_create_by_SVM_admin_using_CLI_fails_with_error_No_candidate_aggregates_are_available_for_storage_services](https://kb.netapp.com/on-prem/ASAr2/ASAr2_KBs/su_create_by_SVM_admin_using_CLI_fails_with_error_No_candidate_aggregates_are_available_for_storage_services) [this^] Knowledge Base article on how to assign aggregates to SVMs in ASA r2 systems.

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 `san-dev` class that uses the `ontap-san` driver and a `san-default` class that uses the `ontap-san-economy` one.

All your Kubernetes worker nodes must have the appropriate iSCSI tools installed. Refer to [xref:{relative_path}worker-node-prep.html](#) [Prepare the worker node] for details.

== Authenticate the ONTAP backend

Trident offers two modes of authenticating an ONTAP backend.

- * **Credential-based:** The username and password to an ONTAP user with the required permissions. It is recommended to use a pre-defined security login role, such as ``admin`` or ``vsadmin`` to ensure maximum compatibility with ONTAP versions.

- * **Certificate-based:** Trident can also communicate with an ONTAP cluster using a certificate installed on the backend. Here, the backend definition must contain Base64-encoded values of the client certificate, key, and the trusted CA certificate if used (recommended).

You can update existing backends to move between credential-based and certificate-based methods. However, only one authentication method is supported at a time. To switch to a different authentication method, you must remove the existing method from the backend configuration.

WARNING: If you attempt to provide *both credentials and certificates*, backend creation will fail with an error that more than one authentication method was provided in the configuration file.

=== Enable credential-based authentication

Trident requires the credentials to an SVM-scoped/cluster-scoped admin to communicate with the ONTAP backend. It is recommended to make use of standard, pre-defined roles such as ``admin`` or ``vsadmin``. This ensures forward compatibility with future ONTAP releases that might expose feature APIs to be used by future Trident releases. A custom security login role can be created and used with Trident, but is not recommended.

A sample backend definition will look like this:

```
// start tabbed area
```

```
[role="tabbed-block"]
```

```
=====
```

```
.YAML
```

```
--
```

```
[source,yaml]
```

```
version: 1
backendName: ExampleBackend
storageDriverName: ontap-san
```

managementLIF: 10.0.0.1
svm: svm_nfs
username: vsadmin
password: password

```
--  
  
.JSON  
--  
[source,json]
```

```
{  
"version": 1,  
"backendName": "ExampleBackend",  
"storageDriverName": "ontap-san",  
"managementLIF": "10.0.0.1",  
"svm": "svm_nfs",  
"username": "vsadmin",  
"password": "password"  
}
```

```
--  
  
====  
  
// end tabbed area
```

Keep in mind that the backend definition is the only place the credentials are stored in plain text. After the backend is created, usernames/passwords are encoded with Base64 and stored as Kubernetes secrets. The creation or update of a backend is the only step that requires knowledge of the credentials. As such, it is an admin-only operation, to be performed by the Kubernetes/storage administrator.

=== Enable certificate-based authentication

New and existing backends can use a certificate and communicate with the ONTAP backend. Three parameters are required in the backend definition.

- * `clientCertificate`: Base64-encoded value of client certificate.
- * `clientPrivateKey`: Base64-encoded value of associated private key.
- * `trustedCACertificate`: Base64-encoded value of trusted CA certificate. If using a trusted CA, this parameter must be provided. This can be ignored if no trusted CA is used.

A typical workflow involves the following steps.

.Steps

```
. Generate a client certificate and key. When generating, set Common Name
(CN) to the ONTAP user to authenticate as.
+
```

```
openssl req -x509 -nodes -days 1095 -newkey rsa:2048 -keyout k8senv.key -out k8senv.pem -subj
"/C=US/ST=NC/L=RTP/O=NetApp/CN=admin"
```

```
. Add trusted CA certificate to the ONTAP cluster. This might be already
handled by the storage administrator. Ignore if no trusted CA is used.
+
```

```
security certificate install -type server -cert-name <trusted-ca-cert-name> -vserver <vserver-name>
ssl modify -vserver <vserver-name> -server-enabled true -client-enabled true -common-name <common-
name> -serial <SN-from-trusted-CA-cert> -ca <cert-authority>
```

```
. Install the client certificate and key (from step 1) on the ONTAP
cluster.
+
```

```
security certificate install -type client-ca -cert-name <certificate-name> -vserver <vserver-name>
security ssl modify -vserver <vserver-name> -client-enabled true
```

```
+
NOTE: After running this command, ONTAP prompts for certificate input.
Paste the contents of the `k8senv.pem` file generated in step 1, then
enter `END` to complete the installation.
. Confirm the ONTAP security login role supports `cert` authentication
method.
+
```

```
security login create -user-or-group-name admin -application ontapi -authentication-method cert
security login create -user-or-group-name admin -application http -authentication-method cert
```

```
. Test authentication using certificate generated. Replace <ONTAP
Management LIF> and <vserver name> with Management LIF IP and SVM name.
+
```

```
curl -X POST -Lk https://<ONTAP-Management-LIF>/servlets/netapp.servlets.admin.XMLrequest\_filer --key
k8senv.key --cert ~/k8senv.pem -d '<?xml version="1.0" encoding="UTF-8"?><netapp
xmlns="http://www.netapp.com/filer/admin" version="1.21" vfiler="<vserver-name>"><vserver-get></vserver-
```

```
get></netapp>'
```

```
. Encode certificate, key and trusted CA certificate with Base64.  
+
```

```
base64 -w 0 k8senv.pem >> cert_base64  
base64 -w 0 k8senv.key >> key_base64  
base64 -w 0 trustedca.pem >> trustedca_base64
```

```
. Create backend using the values obtained from the previous step.  
+
```

```
cat cert-backend.json  
{  
  "version": 1,  
  "storageDriverName": "ontap-san",  
  "backendName": "SanBackend",  
  "managementLIF": "1.2.3.4",  
  "svm": "vserver_test",  
  "clientCertificate": "Faaaakkkkeeee...Vaaallluuuueeee",  
  "clientPrivateKey": "LS0tFaKE...0VaLuES0tLS0K",  
  "trustedCACertificate": "QNFinfO...SiqOyN",  
  "storagePrefix": "myPrefix_"  
}
```

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

```
-----  
| NAME | STORAGE DRIVER | UUID | STATE | VOLUMES |  
-----  
| SanBackend | ontap-san | 586b1cd5-8cf8-428d-a76c-2872713612c1 | online | 0 |  
-----
```

=== Update authentication methods or rotate credentials

You can update an existing backend to use a different authentication method or to rotate their credentials. This works both ways: backends that make use of username/password can be updated to use certificates; backends that utilize certificates can be updated to username/password based. To do this, you must remove the existing authentication method and add the new authentication method. Then use the updated backend.json file containing the required parameters to execute `tridentctl backend update`.

```
cat cert-backend-updated.json  
{  
  "version": 1,  
  "storageDriverName": "ontap-san",  
  "backendName": "SanBackend",
```

```
"managementLIF": "1.2.3.4",
"svm": "vserver_test",
"username": "vsadmin",
"password": "password",
"storagePrefix": "myPrefix_"
}
```

```
#Update backend with tridentctl
tridentctl update backend SanBackend -f cert-backend-updated.json -n trident
```

```
-----
| NAME | STORAGE DRIVER | UUID | STATE | VOLUMES |
-----
| SanBackend | ontap-san | 586b1cd5-8cf8-428d-a76c-2872713612c1 | online | 9 |
-----
```

NOTE: When rotating passwords, the storage administrator must first update the password for the user on ONTAP. This is followed by a backend update. When rotating certificates, multiple certificates can be added to the user. The backend is then updated to use the new certificate, following which the old certificate can be deleted from the ONTAP cluster.

Updating a backend does not disrupt access to volumes that have already been created, nor impact volume connections made after. A successful backend update indicates that Trident can communicate with the ONTAP backend and handle future volume operations.

=== Create custom ONTAP role for Trident

You can create an ONTAP cluster role with minimum privileges so that you do not have to use the ONTAP admin role to perform operations in Trident. When you include the username in a Trident backend configuration, Trident uses the ONTAP cluster role you created to perform the operations.

Refer to

link:https://github.com/NetApp/trident/tree/master/contrib/ontap/trident_role[Trident custom-role generator] for more information about creating Trident custom roles.

```
[role="tabbed-block"]
=====
```

.Using ONTAP CLI

--

. Create a new role using the following command:

+

```
`security login role create <role_name> -cmddirname "command" -access all
-vserver <svm_name>`
```

. Create a username for the Trident user:

```

+
`security login create -username <user_name\> -application ontapi
-authmethod <password\> -role <name_of_role_in_step_1\> -vserver
<svm_name\> -comment "user_description"`
. Map the role to the user:
+
`security login modify username <user_name\> -vserver <svm_name\> -role
<role_name\> -application ontapi -application console -authmethod
<password\>`

```

--

.Using System Manager

--

Perform the following steps in ONTAP System Manager:

```

. *Create a custom role*:
.. To create a custom role at the cluster-level, select *Cluster >
Settings*.
+
(Or) To create a custom role at the SVM level, select *Storage > Storage
VMs > `required SVM`> Settings > Users and Roles*.
.. Select the arrow icon (*->*) next to *Users and Roles*.
.. Select *+Add* under *Roles*.
.. Define the rules for the role and click *Save*.
. *Map the role to the Trident user*:
+ Perform the following steps on the *Users and Roles* page:
.. Select Add icon *++* under *Users*.
.. Select the required username, and select a role in the drop-down menu
for *Role*.
.. Click *Save*.

```

--

====

Refer to the following pages for more information:

```

* link:https://kb.netapp.com/on-prem/ontap/Ontap\_OS/OS-KBs/FAQ\_\_Custom\_roles\_for\_administration\_of\_ONTAP[Custom roles for
administration of ONTAP^] or link:https://docs.netapp.com/us-en/ontap/authentication/define-custom-roles-task.html[Define custom
roles^]
* link:https://docs.netapp.com/us-en/ontap-automation/rest/rbac\_roles\_users.html#rest-api[Work with roles and users^]

```

```
== Authenticate connections with bidirectional CHAP
```

Trident can authenticate iSCSI sessions with bidirectional CHAP for the `ontap-san` and `ontap-san-economy` drivers. This requires enabling the `useCHAP` option in your backend definition. When set to `true`, Trident configures the SVM's default initiator security to bidirectional CHAP and set the username and secrets from the backend file. NetApp recommends using bidirectional CHAP to authenticate connections. See the following sample configuration:

```
[source,yaml]
```

```
version: 1
storageDriverName: ontap-san
backendName: ontap_san_chap
managementLIF: 192.168.0.135
svm: ontap_iscsi_svm
useCHAP: true
username: vsadmin
password: password
chapInitiatorSecret: cl9qxlm36DKyawxy
chapTargetInitiatorSecret: rxigXgkeslpwxyz
chapTargetUsername: iJF4heBRT0TCwxyz
chapUsername: uh2aNCLSD6cNwxyz
```

WARNING: The `useCHAP` parameter is a Boolean option that can be configured only once. It is set to false by default. After you set it to true, you cannot set it to false.

In addition to `useCHAP=true`, the `chapInitiatorSecret`, `chapTargetInitiatorSecret`, `chapTargetUsername`, and `chapUsername` fields must be included in the backend definition. The secrets can be changed after a backend is created by running `tridentctl update`.

```
=== How it works
```

By setting `useCHAP` to true, the storage administrator instructs Trident to configure CHAP on the storage backend. This includes the following:

- * Setting up CHAP on the SVM:

- ** If the SVM's default initiator security type is none (set by default) *and* there are no pre-existing LUNs already present in the volume, Trident will set the default security type to `CHAP` and proceed to configuring the CHAP initiator and target username and secrets.

- ** If the SVM contains LUNs, Trident will not enable CHAP on the SVM. This ensures that access to LUNs that are already present on the SVM isn't

restricted.

* Configuring the CHAP initiator and target username and secrets; these options must be specified in the backend configuration (as shown above).

After the backend is created, Trident creates a corresponding `tridentbackend` CRD and stores the CHAP secrets and usernames as Kubernetes secrets. All PVs that are created by Trident on this backend will be mounted and attached over CHAP.

=== Rotate credentials and update backends

You can update the CHAP credentials by updating the CHAP parameters in the `backend.json` file. This will require updating the CHAP secrets and using the `tridentctl update` command to reflect these changes.

WARNING: When updating the CHAP secrets for a backend, you must use `tridentctl` to update the backend. Do not update the credentials on the storage cluster using the ONTAP CLI or ONTAP System Manager as Trident will not be able to pick up these changes.

```
cat backend-san.json
```

```
{
"version": 1,
"storageDriverName": "ontap-san",
"backendName": "ontap_san_chap",
"managementLIF": "192.168.0.135",
"svm": "ontap_iscsi_svm",
"useCHAP": true,
"username": "vsadmin",
"password": "password",
"chapInitiatorSecret": "cl9qxUpDaTeD",
"chapTargetInitiatorSecret": "rqxigXgkeUpDaTeD",
"chapTargetUsername": "iJF4heBRT0TCwxyz",
"chapUsername": "uh2aNCLSD6cNwxyz",
}
```

```
/tridentctl update backend ontap_san_chap -f backend-san.json -n trident
```

```
-----
| NAME | STORAGE DRIVER | UUID | STATE | VOLUMES |
-----
| ontap_san_chap | ontap-san | aa458f3b-ad2d-4378-8a33-1a472ffbeb5c | online | 7 |
-----
```

Existing connections will remain unaffected; they will continue to remain active if the credentials are updated by Trident on the SVM. New connections use the updated credentials and existing connections continue to remain active. Disconnecting and reconnecting old PVs will result in them using the updated credentials.

```
[[ID22ad79984b37ef03f51ece399251e9f2]]
```

= ONTAP SAN configuration options and examples

:hardbreaks:

:icons: font

:relative_path: ./trident-use/

:imagesdir: {root_path}{relative_path}../media/

[.lead]

Create and use ONTAP SAN drivers with your Trident installation. This section provides backend configuration examples and details for mapping backends to StorageClasses. [link:https://docs.netapp.com/us-en/asa-r2/get-started/learn-about.html](https://docs.netapp.com/us-en/asa-r2/get-started/learn-about.html) [ASA r2 systems^] differ from other ONTAP systems (ASA, AFF, and FAS) in the implementation of their storage layer. These variations impact the usage of certain parameters as notated.

[link:https://docs.netapp.com/us-en/asa-r2/learn-more/hardware-comparison.html](https://docs.netapp.com/us-en/asa-r2/learn-more/hardware-comparison.html) [Learn more about the differences between ASA r2 systems and other ONTAP systems^]. In the Trident backend configuration, you need not specify that your system is ASA r2. When you select `ontap-san` as the `storageDriverName`, Trident detects automatically the ASA r2 or other ONTAP systems. Some backend configuration parameters are not applicable to ASA r2 systems as noted in the table below.

NOTE: Only the `ontap-san` driver (with iSCSI, NVMe/TCP, and FC protocols) is supported for ASA r2 systems.

== Backend configuration options

Beginning with the 26.06 release, you can enable ONTAP balanced placement so that ONTAP selects the aggregate for new volumes. Refer to [xref:{relative_path}ontap-balanced-placement.html](#) [Use balanced placement with ONTAP backends] for requirements and examples.

See the following table for the backend configuration options:

```
[cols="1,3,2",options="header"]
```

```
|===
```

```
|Parameter |Description |Default
```

```
|`version` | |Always 1
```

```
|`storageDriverName` | Name of the storage driver.
```

Cannot be modified after creation. To update this parameter, you must create a new backend. |`ontap-san` or `ontap-san-economy`

|`backendName` |Custom name or the storage backend |Driver name + "_" + dataLIF

|`managementLIF` a|IP address of a cluster or SVM management LIF.

A fully-qualified domain name (FQDN) can be specified.

Can be set to use IPv6 addresses if Trident was installed using the IPv6 flag. IPv6 addresses must be defined in square brackets, such as `[28e8:d9fb:a825:b7bf:69a8:d02f:9e7b:3555]`.

For seamless MetroCluster switchover, see the <<mcc-best>>.

NOTE: If you are using "vsadmin" credentials, `managementLIF` must be that of the SVM; if using "admin" credentials, `managementLIF` must be that of the cluster.

|"10.0.0.1", "[2001:1234:abcd::fefe]"

|`dataLIF` |IP address of protocol LIF.

Can be set to use IPv6 addresses if Trident was installed using the IPv6 flag. IPv6 addresses must be defined in square brackets, such as `[28e8:d9fb:a825:b7bf:69a8:d02f:9e7b:3555]`.

Do not specify for iSCSI. Trident uses link:<https://docs.netapp.com/us-en/ontap/san-admin/selective-lun-map-concept.html>[ONTAP Selective LUN Map^] to discover the iSCSI LIFs needed to establish a multi path session. A warning is generated if `dataLIF` is explicitly defined.

Omit for Metrocluster. See the <<mcc-best>>.

|Derived by the SVM

|`svm` |Storage virtual machine to use

Omit for Metrocluster. See the <<mcc-best>>.

|Derived if an SVM `managementLIF` is specified

|`useCHAP` |Use CHAP to authenticate iSCSI for ONTAP SAN drivers [Boolean].

Set to `true` for Trident to configure and use bidirectional CHAP as the default authentication for the SVM given in the backend. Refer to [xref:{relative_path}ontap-san-prep.html](#)[Prepare to configure backend with ONTAP SAN drivers] for details.

Not supported for FCP or NVMe/TCP.

|`false`

|`chapInitiatorSecret` |CHAP initiator secret. Required if `useCHAP=true`
|""

|`labels` |Set of arbitrary JSON-formatted labels to apply on volumes |""

|`chapTargetInitiatorSecret` |CHAP target initiator secret. Required if
`useCHAP=true` |""

|`chapUsername` |Inbound username. Required if `useCHAP=true` |""

|`chapTargetUsername` |Target username. Required if `useCHAP=true` |""

|`clientCertificate` |Base64-encoded value of client certificate. Used for
certificate-based auth |""

|`clientPrivateKey` |Base64-encoded value of client private key. Used for
certificate-based auth |""

|`trustedCACertificate` |Base64-encoded value of trusted CA certificate.
Optional. Used for certificate-based authentication. |""

|`username` |Username needed to communicate with the ONTAP cluster. Used
for credential-based authentication.

For Active Directory authentication, see [xref:{relative_path}../trident-use/ontap-san-examples.html#authenticate-trident-to-a-backend-svm-using-active-directory-credentials](#)[Authenticate Trident to a backend SVM using Active Directory credentials]. |""

|`password` |Password needed to communicate with the ONTAP cluster. Used
for credential-based authentication.

For Active Directory authentication, see [xref:{relative_path}../trident-use/ontap-san-examples.html#authenticate-trident-to-a-backend-svm-using-active-directory-credentials](#)[Authenticate Trident to a backend SVM using Active Directory credentials].
|""

|`svm` |Storage virtual machine to use |Derived if an SVM `managementLIF`
is specified

|`storagePrefix` |Prefix used when provisioning new volumes in the SVM.

Cannot be modified later. To update this parameter, you need to create a

new backend.|`trident`

|`aggregate`

a|Aggregate for provisioning (optional; if set, must be assigned to the SVM). For the `ontap-nas-flexgroup` driver, this option is ignored. If not assigned, any of the available aggregates can be used to provision a FlexGroup volume.

NOTE: When the aggregate is updated in SVM, it is updated in Trident automatically by polling SVM without having to restart the Trident Controller. When you have configured a specific aggregate in Trident to provision volumes, if the aggregate is renamed or moved out of the SVM, the backend moves to failed state in Trident while polling the SVM aggregate. You must either change the aggregate to one that is present on the SVM or remove it altogether to bring the backend back online.

Do not specify for ASA r2 systems.

Do not specify when `useBalancedPlacement` is set to `true`.

a|""

|`useBalancedPlacement`

a|Delegate aggregate selection for new volumes to ONTAP [Boolean]. Requires ONTAP 9.19.1 or later and `useREST` set to `true`. Cannot be combined with `aggregate`.

Refer to [xref:{relative_path}ontap-balanced-placement.html](#) [Use balanced placement with ONTAP backends] for requirements and examples.

|`false`

|`limitAggregateUsage` |Fail provisioning if usage is above this percentage.

If you are using an Amazon FSx for NetApp ONTAP backend, do not specify `limitAggregateUsage`. The provided `fsxadmin` and `vsadmin` do not contain the permissions required to retrieve aggregate usage and limit it using Trident.

Do not specify for ASA r2 systems.|"" (not enforced by default)

|`limitVolumeSize` |Fail provisioning if requested volume size is above this value.

Also restricts the maximum size of the volumes it manages for LUNs.

|"" (not enforced by default)

```
|`lunsPerFlexvol` |Maximum LUNs per Flexvol, must be in range [50, 200]
|`100`
```

```
|`debugTraceFlags` |Debug flags to use when troubleshooting. Example,
{"api":false, "method":true}
```

Do not use unless you are troubleshooting and require a detailed log dump. |`null`

```
|`useREST` a|Boolean parameter to use ONTAP REST APIs.
```

`useREST` When set to `true`, Trident uses ONTAP REST APIs to communicate with the backend; when set to `false`, Trident uses ONTAPI (ZAPI) calls to communicate with the backend. This feature requires ONTAP 9.11.1 and later. In addition, the ONTAP login role used must have access to the `ontapi` application. This is satisfied by the pre-defined `vsadmin` and `cluster-admin` roles. Beginning with the Trident 24.06 release and ONTAP 9.15.1 or later, `useREST` is set to `true` by default; change `useREST` to `false` to use ONTAPI (ZAPI) calls.

`useREST` is fully qualified for NVMe/TCP.

NOTE: NVMe is supported only with ONTAP REST APIs and not supported with ONTAPI (ZAPI).

If specified, always set to `true` for ASA r2 systems. |`true` for ONTAP 9.15.1 or later, otherwise `false`.

```
a|`sanType` |
```

Use to select `iscsi` for iSCSI, `nvme` for NVMe/TCP or `fcp` for SCSI over Fibre Channel (FC).

```
| `iscsi` if blank
```

```
|`formatOptions`
```

a| Use `formatOptions` to specify command line arguments for the `mkfs` command, which is applied whenever a volume is formatted. This allows you to format the volume according to your preferences. Make sure to specify the formatOptions similar to that of the mkfs command options, excluding the device path.

Example: "-E nodiscard"

Supported for `ontap-san` and `ontap-san-economy` drivers with iSCSI protocol.

Additionally, supported for ASA r2 systems when using iSCSI and NVMe/TCP protocols.

a|

|`limitVolumePoolSize` |Maximum requestable FlexVol size when using LUNs in ontap-san-economy backend.

| "" (not enforced by default)

|`denyNewVolumePools` |Restricts `ontap-san-economy` backends from creating new FlexVol volumes to contain their LUNs. Only preexisting Flexvols are used for provisioning new PVs.

|

|===

=== Recommendations for using formatOptions

Trident recommends the following options to expedite the formatting process:

* *-E nodiscard (ext3, ext4):* Do not attempt to discard blocks at mkfs time (discarding blocks initially is useful on solid state devices and sparse / thin-provisioned storage). This replaces the deprecated option "-K" and it is applicable to ext3, ext4 file systems.

* *-K (xfs):* Do not attempt to discard blocks at mkfs time. This option is applicable to xfs file system.

=== Authenticate Trident to a backend SVM using Active Directory credentials

You can configure Trident to authenticate to a backend SVM using Active Directory (AD) credentials. Before an AD account can access the SVM, you must configure AD domain controller access to the cluster or SVM. For cluster administration with an AD account, you must create domain tunnel. Refer to link:<https://docs.netapp.com/us-en/ontap/authentication/enable-ad-users-groups-access-cluster-svm-task.html>[Configure Active Directory domain controller access in ONTAP^] for details.

.steps

. Configure Domain Name System (DNS) settings for a backend SVM:

+

```
`vserver services dns create -vserver <svm_name> -dns-servers  
<dns_server_ip1>,<dns_server_ip2>`
```

```
. Run the following command to create a computer account for the SVM in
Active Directory:
+
`vserver active-directory create -vserver DataSVM -account-name ADSERVER1
-domain demo.netapp.com`
. Use this command to create an AD user or group to manage the cluster or
SVM
+
`security login create -vserver <svm_name> -user-or-group-name
<ad_user_or_group> -application <application> -authentication-method
domain -role vsadmin`
. In the Trident backend configuration file, set the `username` and
`password` parameters to the AD user or group name and password,
respectively.
```

== Backend configuration options for provisioning volumes

You can control default provisioning using these options in the `defaults` section of the configuration. For an example, see the configuration examples below.

```
[cols="1,3,2",options="header"]
|===
|Parameter |Description |Default
|`spaceAllocation` |Space-allocation for LUNs |"true"
*If specified, set to `true` for ASA r2 systems*.

|`spaceReserve` |Space reservation mode; "none" (thin) or "volume"
(thick).

*Set to `none` for ASA r2* systems.|"none"

|`snapshotPolicy` |Snapshot policy to use.

*Set to `none` for ASA r2 systems*.|"none"

|`qosPolicy` |QoS policy group to assign for volumes created. Choose one
of qosPolicy or adaptiveQosPolicy per storage pool/backend.

Using QoS policy groups with Trident requires ONTAP 9.8 or later. You
should use a non-shared QoS policy group and ensuring the policy group is
applied to each constituent individually. A shared QoS policy group
enforces the ceiling for the total throughput of all workloads. |""

|`adaptiveQosPolicy` |Adaptive QoS policy group to assign for volumes
```

created. Choose one of qosPolicy or adaptiveQosPolicy per storage pool/backend |""

|`snapshotReserve` |Percentage of volume reserved for snapshots.

Do not specify for ASA r2 systems. |"0" if `snapshotPolicy` is "none", otherwise ""

|`splitOnClone` |Split a clone from its parent upon creation |"false"

|`encryption` |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.

If NAE is enabled on the backend, any volume provisioned in Trident is NAE enabled.

For more information, refer to: xref:{relative_path}../trident-reco/security-reco.html[How Trident works with NVE and NAE]. |"false"

If specified, set to `true` for ASA r2 systems.

|`luksEncryption` | Enable LUKS encryption. Refer to xref:{relative_path}../trident-reco/security-luks.html[Use Linux Unified Key Setup (LUKS)].

| ""

Set to `false` for ASA r2 systems.

|`tieringPolicy` |Tiering policy to use "none"

Do not specify for ASA r2 systems .|

|`nameTemplate` | Template to create custom volume names.

|""

|===

=== Volume provisioning examples

Here's an example with defaults defined:

[source,yaml]

```
version: 1
storageDriverName: ontap-san
managementLIF: 10.0.0.1
svm: trident_svm
username: admin
password: <password>
labels:
k8scluster: dev2
backend: dev2-sanbackend
storagePrefix: alternate-trident
debugTraceFlags:
api: false
method: true
defaults:
spaceReserve: volume
qosPolicy: standard
spaceAllocation: 'false'
snapshotPolicy: default
snapshotReserve: '10'
```

NOTE: For all volumes created using the `ontap-san` driver, Trident adds an extra 10 percent capacity to the FlexVol to accommodate the LUN metadata. The LUN is provisioned with the exact size that the user requests in the PVC. Trident adds 10 percent to the FlexVol (shows as Available size in ONTAP). Users now get the amount of usable capacity they requested. This change also prevents LUNs from becoming read-only unless the available space is fully utilized. This does not apply to ontap-san-economy.

For backends that define `snapshotReserve`, Trident calculates the size of volumes as follows:

Total volume size = [(PVC requested size) / (1 - (snapshotReserve percentage) / 100)] * 1.1

The 1.1 is the extra 10 percent Trident adds to the FlexVol to accommodate the LUN metadata. For `snapshotReserve` = 5%, and PVC request = 5 GiB, the total volume size is 5.79 GiB and the available size is 5.5 GiB. The `volume show` command should show results similar to this example:

image:../media/vol-show-san.png[Shows the output of the volume show command.]

Currently, resizing is the only way to use the new calculation for an existing volume.

== Minimal configuration examples

The following examples show basic configurations that leave most parameters to default. This is the easiest way to define a backend.

NOTE: If you are using Amazon FSx on NetApp ONTAP with Trident, NetApp recommends that you specify DNS names for LIFs instead of IP addresses.

.ONTAP SAN example

[%collapsible%closed]

====

This is a basic configuration using the `ontap-san` driver.

[source,yaml]

```
version: 1
storageDriverName: ontap-san
managementLIF: 10.0.0.1
svm: svm_iscsi
labels:
k8scluster: test-cluster-1
backend: testcluster1-sanbackend
username: vsadmin
password: <password>
```

====

[#mcc-best]

.MetroCluster example

[%collapsible%closed]

====

You can configure the backend to avoid having to manually update the backend definition after switchover and switchback during [xref:{relative_path}../trident-reco/backup.html#svm-replication-and-recovery](#)[SVM replication and recovery].

For seamless switchover and switchback, specify the SVM using `managementLIF` and omit the `svm` parameters. For example:

[source,yaml]

```
version: 1
storageDriverName: ontap-san
managementLIF: 192.168.1.66
username: vsadmin
password: password
```

```
====
```

```
.ONTAP SAN economy example  
[%collapsible%closed]
```

```
====
```

```
[source,yaml]
```

```
version: 1  
storageDriverName: ontap-san-economy  
managementLIF: 10.0.0.1  
svm: svm_iscsi_eco  
username: vsadmin  
password: <password>
```

```
====
```

```
.Certificate-based authentication example  
[%collapsible%closed]
```

```
====
```

```
In this basic configuration example `clientCertificate`,  
`clientPrivateKey`, and `trustedCACertificate` (optional, if using trusted  
CA) are populated in `backend.json` and take the base64-encoded values of  
the client certificate, private key, and trusted CA certificate,  
respectively.  
[source,yaml]
```

```
version: 1  
storageDriverName: ontap-san  
backendName: DefaultSANBackend  
managementLIF: 10.0.0.1  
svm: svm_iscsi  
useCHAP: true  
chapInitiatorSecret: cl9qxlM36DKyawxy  
chapTargetInitiatorSecret: rqxigXgkeslpwxyz  
chapTargetUsername: iJF4heBRT0TCwxyz  
chapUsername: uh2aNCLSD6cNwxyz  
clientCertificate: ZXR0ZXJwYXB...ICMgJ3BhcGVyc2  
clientPrivateKey: vciwKlyAgZG...0cnksIGRlc2NyaX  
trustedCACertificate: zcyBbaG...b3lgb3dulGNsYXNz
```

```
====
```

```
.Bidirectional CHAP examples  
[%collapsible%closed]
```

====

These examples create a backend with `useCHAP` set to `true`.

```
.ONTAP SAN CHAP example
[source,yaml]
```

```
version: 1
storageDriverName: ontap-san
managementLIF: 10.0.0.1
svm: svm_iscsi
labels:
k8scluster: test-cluster-1
backend: testcluster1-sanbackend
useCHAP: true
chapInitiatorSecret: cl9qxlm36DKyawxy
chapTargetInitiatorSecret: rqxigXgkeslpwxyz
chapTargetUsername: iJF4heBRT0TCwxyz
chapUsername: uh2aNCLSD6cNwxyz
username: vsadmin
password: <password>
```

```
.ONTAP SAN economy CHAP example
[source,yaml]
```

```
version: 1
storageDriverName: ontap-san-economy
managementLIF: 10.0.0.1
svm: svm_iscsi_eco
useCHAP: true
chapInitiatorSecret: cl9qxlm36DKyawxy
chapTargetInitiatorSecret: rqxigXgkeslpwxyz
chapTargetUsername: iJF4heBRT0TCwxyz
chapUsername: uh2aNCLSD6cNwxyz
username: vsadmin
password: <password>
```

====

```
.Balanced placement example
[%collapsible%closed]
```

====

This example enables ONTAP balanced placement so that ONTAP selects the aggregate for new volumes. Balanced placement requires ONTAP 9.19.1 or later. Refer to [xref:{relative_path}ontap-balanced-placement.html](#) [Use

```
balanced placement with ONTAP backends] for details.  
[source,yaml]
```

```
version: 1  
storageDriverName: ontap-san  
backendName: san-balanced-placement  
managementLIF: 10.0.0.1  
svm: svm_iscsi  
username: vsadmin  
password: <password>  
useBalancedPlacement: true
```

```
====
```

```
.NVMe/TCP example  
[%collapsible%closed]
```

```
====
```

You must have an SVM configured with NVMe on your ONTAP backend. This is a basic backend configuration for NVMe/TCP.

```
[source,yaml]
```

```
version: 1  
backendName: NVMeBackend  
storageDriverName: ontap-san  
managementLIF: 10.0.0.1  
svm: svm_nvme  
username: vsadmin  
password: password  
sanType: nvme  
useREST: true
```

```
====
```

```
.SCSI over FC (FCP) example  
[%collapsible%closed]
```

```
====
```

You must have an SVM configured with FC on your ONTAP backend. This is a basic backend configuration for FC.

```
[source,yaml]
```

```
version: 1
```

```
backendName: fcp-backend
storageDriverName: ontap-san
managementLIF: 10.0.0.1
svm: svm_fc
username: vsadmin
password: password
sanType: fcp
useREST: true
```

```
====
```

```
.Backend configuration example with nameTemplate
[%collapsible%closed]
```

```
====
```

```
[source,yaml]
```

```
version: 1
storageDriverName: ontap-san
backendName: ontap-san-backend
managementLIF: <ip address>
svm: svm0
username: <admin>
password: <password>
defaults:
nameTemplate: "{{.volume.Name}}_{{.labels.cluster}}_{{.volume.Namespace}}_{{.vo\
lume.RequestName}}"
labels:
cluster: ClusterA
PVC: "{{.volume.Namespace}}_{{.volume.RequestName}}"
```

```
====
```

```
.formatOptions example for ontap-san-economy driver
[%collapsible%closed]
```

```
====
```

```
[source,yaml]
```

```
version: 1
storageDriverName: ontap-san-economy
managementLIF: ""
svm: svm1
username: ""
password: "!"
storagePrefix: whelk_
```

debugTraceFlags:
method: true
api: true
defaults:
formatOptions: -E nodiscard

====

== Examples of backends with virtual pools

In these sample backend definition files, specific defaults are set for all storage pools, such as `spaceReserve` at none, `spaceAllocation` at false, and `encryption` at false. The virtual pools are defined in the storage section.

Trident sets provisioning labels in the "Comments" field. Comments are set on the FlexVol volume Trident copies all labels present on a virtual pool to the storage volume at provisioning. For convenience, storage administrators can define labels per virtual pool and group volumes by label.

In these examples, some of the storage pools set their own `spaceReserve`, `spaceAllocation`, and `encryption` values, and some pools override the default values.

.ONTAP SAN example
[%collapsible%closed]
====
[source,yaml]

version: 1
storageDriverName: ontap-san
managementLIF: 10.0.0.1
svm: svm_iscsi
useCHAP: true
chapInitiatorSecret: cl9qxl36DKyawxy
chapTargetInitiatorSecret: rxigXgkeslpwxyz
chapTargetUsername: iJF4heBRT0TCwxyz
chapUsername: uh2aNCLSD6cNwxyz
username: vsadmin
password: <password>
defaults:
spaceAllocation: "false"
encryption: "false"
qosPolicy: standard
labels:
store: san_store

```

kubernetes-cluster: prod-cluster-1
region: us_east_1
storage:
- labels:
  protection: gold
  creditpoints: "40000"
  zone: us_east_1a
  defaults:
  spaceAllocation: "true"
  encryption: "true"
  adaptiveQosPolicy: adaptive-extreme
- labels:
  protection: silver
  creditpoints: "20000"
  zone: us_east_1b
  defaults:
  spaceAllocation: "false"
  encryption: "true"
  qosPolicy: premium
- labels:
  protection: bronze
  creditpoints: "5000"
  zone: us_east_1c
  defaults:
  spaceAllocation: "true"
  encryption: "false"

```

```
====
```

```

.ONTAP SAN economy example
[%collapsible%closed]
====
[source,yaml]

```

```

version: 1
storageDriverName: ontap-san-economy
managementLIF: 10.0.0.1
svm: svm_iscsi_eco
useCHAP: true
chapInitiatorSecret: cl9qxlm36DKyawxy
chapTargetInitiatorSecret: rqxigXgkeslpwxyz
chapTargetUsername: iJF4heBRT0TCwxyz
chapUsername: uh2aNCLSD6cNwxyz
username: vsadmin
password: <password>
defaults:
spaceAllocation: "false"
encryption: "false"
labels:
store: san_economy_store

```

```
region: us_east_1
storage:
- labels:
app: oracledb
cost: "30"
zone: us_east_1a
defaults:
spaceAllocation: "true"
encryption: "true"
- labels:
app: postgresdb
cost: "20"
zone: us_east_1b
defaults:
spaceAllocation: "false"
encryption: "true"
- labels:
app: mysqldb
cost: "10"
zone: us_east_1c
defaults:
spaceAllocation: "true"
encryption: "false"
- labels:
department: legal
creditpoints: "5000"
zone: us_east_1c
defaults:
spaceAllocation: "true"
encryption: "false"
```

```
====

.NVMe/TCP example
[%collapsible%closed]
====

[source,yaml]
```

```
version: 1
storageDriverName: ontap-san
sanType: nvme
managementLIF: 10.0.0.1
svm: nvme_svm
username: vsadmin
password: <password>
useREST: true
defaults:
spaceAllocation: "false"
encryption: "true"
storage:
```

```
- labels:
app: testApp
cost: "20"
defaults:
spaceAllocation: "false"
encryption: "false"
```

```
====
```

```
== Map backends to StorageClasses
```

The following StorageClass definitions refer to the <<Examples of backends with virtual pools>>. Using the `parameters.selector` field, each StorageClass calls out which virtual pools can be used to host a volume. The volume now has the aspects defined in the chosen virtual pool.

* The `protection-gold` StorageClass maps to the first virtual pool in the `ontap-san` backend. This is the only pool offering gold-level protection.

+

```
[source,yaml]
```

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
name: protection-gold
provisioner: csi.trident.netapp.io
parameters:
selector: "protection=gold"
fsType: "ext4"
```

* The `protection-not-gold` StorageClass maps to the second and third virtual pool in `ontap-san` backend. These are the only pools offering a protection level other than gold.

+

```
[source,yaml]
```

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
name: protection-not-gold
provisioner: csi.trident.netapp.io
parameters:
selector: "protection!=gold"
fsType: "ext4"
```

* The `app-mysqldb` StorageClass maps to the third virtual pool in `ontap-san-economy` backend. This is the only pool offering storage pool

```
configuration for the mysqlldb type app.  
+  
[source,yaml]
```

```
apiVersion: storage.k8s.io/v1  
kind: StorageClass  
metadata:  
name: app-mysqldb  
provisioner: csi.trident.netapp.io  
parameters:  
selector: "app=mysqlldb"  
fsType: "ext4"
```

```
* The `protection-silver-creditpoints-20k` StorageClass maps to the second  
virtual pool in `ontap-san` backend. This is the only pool offering  
silver-level protection and 20000 creditpoints.  
+  
[source,yaml]
```

```
apiVersion: storage.k8s.io/v1  
kind: StorageClass  
metadata:  
name: protection-silver-creditpoints-20k  
provisioner: csi.trident.netapp.io  
parameters:  
selector: "protection=silver; creditpoints=20000"  
fsType: "ext4"
```

```
* The `creditpoints-5k` StorageClass maps to the third virtual pool in  
`ontap-san` backend and the fourth virtual pool in the `ontap-san-economy`  
backend. These are the only pool offerings with 5000 creditpoints.  
+  
[source,yaml]
```

```
apiVersion: storage.k8s.io/v1  
kind: StorageClass  
metadata:  
name: creditpoints-5k  
provisioner: csi.trident.netapp.io  
parameters:  
selector: "creditpoints=5000"  
fsType: "ext4"
```

```
* The `my-test-app-sc` StorageClass maps to the `testAPP` virtual pool in  
the `ontap-san` driver with `sanType: nvme`. This is the only pool  
offering `testApp`.
```

```
+  
[source,yaml]
```

```
apiVersion: storage.k8s.io/v1  
kind: StorageClass  
metadata:  
name: my-test-app-sc  
provisioner: csi.trident.netapp.io  
parameters:  
selector: "app=testApp"  
fsType: "ext4"
```

Trident decides which virtual pool is selected and ensures the storage requirement is met.

```
:leveloffset: -1
```

= ONTAP NAS drivers

```
:leveloffset: +1
```

```
[[IDc18456bedb2b3f98a3ee30d89c002992]]
```

= ONTAP NAS driver overview

```
:hardbreaks:
```

```
:icons: font
```

```
:relative_path: ./trident-use/
```

```
:imagesdir: {root_path}{relative_path}../media/
```

```
[.lead]
```

Learn about configuring an ONTAP backend with ONTAP and Cloud Volumes ONTAP NAS drivers.

== ONTAP NAS driver details

Trident provides the following NAS storage drivers to communicate with the ONTAP cluster. Supported access modes are: `_ReadWriteOnce_ (RWO)`, `_ReadOnlyMany_ (ROX)`, `_ReadWriteMany_ (RWX)`, `_ReadWriteOncePod_ (RWOP)`.

//IMPORTANT: If you are using Astra Control for protection, recovery, and mobility, read <<Astra Control driver compatibility>>.

```
[cols="2, 1, 1, 2, 2", options="header"]
```

```
|===  
|Driver  
|Protocol  
|volumeMode  
|Access modes supported  
|File systems supported
```

```
|`ontap-nas`
```

```
a|NFS
```

```
SMB
```

```
a|Filesystem
```

```
a|RWO, ROX, RWX, RWOP
```

```
a|", `nfs`, `smb`
```

```
|`ontap-nas-economy`
```

```
a|NFS
```

```
SMB
```

```
a|Filesystem
```

```
a|RWO, ROX, RWX, RWOP
```

```
a|", `nfs`, `smb`
```

```
|`ontap-nas-flexgroup`
```

```
a|NFS
```

```
SMB
```

```
a|Filesystem
```

```
a|RWO, ROX, RWX, RWOP
```

```
a|", `nfs`, `smb`
```

```
|===
```

```
//=== Astra Control driver compatibility
```

//Astra Control provides seamless protection, disaster recovery, and mobility (moving volumes between Kubernetes clusters) for volumes created with the `ontap-nas`, `ontap-nas-flexgroup`, and `ontap-san` drivers. Refer to link:https://docs.netapp.com/us-en/astra-control-center/use/replicate_snapmirror.html#replication-prerequisites[Astra Control replication prerequisites^] for details.

```
[WARNING]
```

```
=====
```

* Use `ontap-san-economy` only if persistent volume usage count is expected to be higher than link:<https://docs.netapp.com/us-en/ontap/volumes/storage-limits-reference.html>[supported ONTAP volume limits^].

* Use `ontap-nas-economy` only if persistent volume usage count is

expected to be higher than [link:https://docs.netapp.com/us-en/ontap/volumes/storage-limits-reference.html](https://docs.netapp.com/us-en/ontap/volumes/storage-limits-reference.html)[supported ONTAP volume limits^] and the `ontap-san-economy` driver cannot be used.

* Do not use use `ontap-nas-economy` if you anticipate the need for data protection, disaster recovery, or mobility.

* NetApp does not recommend using Flexvol autogrow in all ONTAP drivers, except ontap-san. As a workaround, Trident supports the use of snapshot reserve and scales Flexvol volumes accordingly.

====

== 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.

For Amazon FSx for NetApp ONTAP deployments, Trident expects to be run as either an ONTAP or SVM administrator, using the cluster `fsxadmin` user or a `vsadmin` SVM user, or a user with a different name that has the same role. The `fsxadmin` user is a limited replacement for the cluster admin user.

NOTE: If you use the `limitAggregateUsage` parameter, cluster admin permissions are required. When using Amazon FSx for NetApp ONTAP with Trident, the `limitAggregateUsage` parameter will not work with the `vsadmin` and `fsxadmin` user accounts. The configuration operation will fail if you specify this parameter.

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.

```
[[ID6080d27a027c38f0355d857104f2ebd1]]
```

```
= Prepare to configure a backend with ONTAP NAS drivers
```

```
:hardbreaks:
```

```
:icons: font
```

```
:relative_path: ./trident-use/
```

```
:imagesdir: {root_path}{relative_path}../media/
```

```
[.lead]
```

Understand the requirements, authentication options, and export policies for configuring an ONTAP backend with ONTAP NAS drivers. Beginning with the 25.10 release, NetApp Trident supports

[link:https://docs.netapp.com/us-en/ontap-afx/index.html](https://docs.netapp.com/us-en/ontap-afx/index.html)[NetApp AFX storage

system^]. NetApp AFX storage systems differ from other ONTAP systems (ASA, AFF, and FAS) in the implementation of their storage layer. In the Trident backend configuration, you need not specify that your system is AFX. When you select ``ontap-nas`` as the ``storageDriverName``, Trident detects automatically the AFX systems.

NOTE: Only the ``ontap-nas`` driver (with NFS protocol) is supported for AFX systems; SMB protocol is not supported.

== Requirements

- * For all ONTAP backends, Trident requires at least one aggregate be assigned to the SVM.
- * You can 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.
- * All your Kubernetes worker nodes must have the appropriate NFS tools installed. Refer to [xref:{relative_path}worker-node-prep.html\[here\]](#) for more details.
- * Trident supports SMB volumes mounted to pods running on Windows nodes only. Refer to [<<Prepare to provision SMB volumes>>](#) for details.

== Authenticate the ONTAP backend

Trident offers two modes of authenticating an ONTAP backend.

- * **Credential-based:** This mode requires sufficient permissions to the ONTAP backend. It is recommended to use an account associated with a pre-defined security login role, such as ``admin`` or ``vsadmin`` to ensure maximum compatibility with ONTAP versions.
- * **Certificate-based:** This mode requires a certificate installed on the backend for Trident to communicate with an ONTAP cluster. Here, the backend definition must contain Base64-encoded values of the client certificate, key, and the trusted CA certificate if used (recommended).

You can update existing backends to move between credential-based and certificate-based methods. However, only one authentication method is supported at a time. To switch to a different authentication method, you must remove the existing method from the backend configuration.

WARNING: If you attempt to provide *both credentials and certificates*, backend creation will fail with an error that more than one authentication method was provided in the configuration file.

=== Enable credential-based authentication

Trident requires the credentials to an SVM-scoped/cluster-scoped admin to communicate with the ONTAP backend. It is recommended to make use of standard, pre-defined roles such as `admin` or `vsadmin`. This ensures forward compatibility with future ONTAP releases that might expose feature APIs to be used by future Trident releases. A custom security login role can be created and used with Trident, but is not recommended.

A sample backend definition will look like this:

```
// start tabbed area

[role="tabbed-block"]
====

.YAML
--
[source,yaml]
```

```
version: 1
backendName: ExampleBackend
storageDriverName: ontap-nas
managementLIF: 10.0.0.1
dataLIF: 10.0.0.2
svm: svm_nfs
credentials:
name: secret-backend-creds
```

```
--

.JSON
--
[source,json]
```

```
{
"version": 1,
"backendName": "ExampleBackend",
"storageDriverName": "ontap-nas",
"managementLIF": "10.0.0.1",
"dataLIF": "10.0.0.2",
"svm": "svm_nfs",
"credentials": {
"name": "secret-backend-creds"
}
}
```

```
--  
====
```

```
// end tabbed area
```

Keep in mind that the backend definition is the only place the credentials are stored in plain text. After the backend is created, usernames/passwords are encoded with Base64 and stored as Kubernetes secrets. The creation/updating of a backend is the only step that requires knowledge of the credentials. As such, it is an admin-only operation, to be performed by the Kubernetes/storage administrator.

=== Enable certificate-based Authentication

New and existing backends can use a certificate and communicate with the ONTAP backend. Three parameters are required in the backend definition.

- * `clientCertificate`: Base64-encoded value of client certificate.
- * `clientPrivateKey`: Base64-encoded value of associated private key.
- * `trustedCACertificate`: Base64-encoded value of trusted CA certificate. If using a trusted CA, this parameter must be provided. This can be ignored if no trusted CA is used.

A typical workflow involves the following steps.

.Steps

- . Generate a client certificate and key. When generating, set Common Name (CN) to the ONTAP user to authenticate as.

+

```
openssl req -x509 -nodes -days 1095 -newkey rsa:2048 -keyout k8senv.key -out k8senv.pem -subj  
"/C=US/ST=NC/L=RTP/O=NetApp/CN=vsadmin"
```

- . Add trusted CA certificate to the ONTAP cluster. This might be already handled by the storage administrator. Ignore if no trusted CA is used.

+

```
security certificate install -type server -cert-name <trusted-ca-cert-name> -vserver <vserver-name>  
ssl modify -vserver <vserver-name> -server-enabled true -client-enabled true -common-name <common-  
name> -serial <SN-from-trusted-CA-cert> -ca <cert-authority>
```

- . Install the client certificate and key (from step 1) on the ONTAP

```
cluster.  
+
```

```
security certificate install -type client-ca -cert-name <certificate-name> -vserver <vserver-name>  
security ssl modify -vserver <vserver-name> -client-enabled true
```

```
. Confirm the ONTAP security login role supports `cert` authentication  
method.  
+
```

```
security login create -user-or-group-name vsadmin -application ontapi -authentication-method cert -vserver  
<vserver-name>  
security login create -user-or-group-name vsadmin -application http -authentication-method cert -vserver  
<vserver-name>
```

```
. Test authentication using certificate generated. Replace <ONTAP  
Management LIF> and <vserver name> with Management LIF IP and SVM name.  
You must ensure the LIF has its service policy set to `default-data-  
management`.  
+
```

```
curl -X POST -Lk https://<ONTAP-Management-LIF>/servlets/netapp.servlets.admin.XMLrequest\_filer --key  
k8senv.key --cert ~/k8senv.pem -d '<?xml version="1.0" encoding="UTF-8"?><netapp  
xmlns="http://www.netapp.com/filer/admin" version="1.21" vfiler="<vserver-name>"><vserver-get></vserver-  
get></netapp>'
```

```
. Encode certificate, key and trusted CA certificate with Base64.  
+
```

```
base64 -w 0 k8senv.pem >> cert_base64  
base64 -w 0 k8senv.key >> key_base64  
base64 -w 0 trustedca.pem >> trustedca_base64
```

```
. Create backend using the values obtained from the previous step.  
+
```

```
cat cert-backend-updated.json  
{  
  "version": 1,  
  "storageDriverName": "ontap-nas",  
  "backendName": "NasBackend",  
  "managementLIF": "1.2.3.4",  
  "dataLIF": "1.2.3.8",  
  "svm": "vserver_test",
```

```
"clientCertificate": "Faaaakkkkeeee...Vaaalllluuuuueeee",
"clientPrivateKey": "LS0tFaKE...0VaLuES0tLS0K",
"storagePrefix": "myPrefix_"
}
```

```
#Update backend with tridentctl
tridentctl update backend NasBackend -f cert-backend-updated.json -n trident
```

```
-----
| NAME | STORAGE DRIVER | UUID | STATE | VOLUMES |
-----
| NasBackend | ontap-nas | 98e19b74-aec7-4a3d-8dcf-128e5033b214 | online | 9 |
-----
```

=== Update authentication methods or rotate credentials

You can update an existing backend to use a different authentication method or to rotate their credentials. This works both ways: backends that make use of username/password can be updated to use certificates; backends that utilize certificates can be updated to username/password based. To do this, you must remove the existing authentication method and add the new authentication method. Then use the updated backend.json file containing the required parameters to execute `tridentctl update backend`.

```
cat cert-backend-updated.json
```

```
[source,json]
```

```
{
"version": 1,
"storageDriverName": "ontap-nas",
"backendName": "NasBackend",
"managementLIF": "1.2.3.4",
"dataLIF": "1.2.3.8",
"svm": "vserver_test",
"username": "vsadmin",
"password": "password",
"storagePrefix": "myPrefix_"
}
```

```
#Update backend with tridentctl
tridentctl update backend NasBackend -f cert-backend-updated.json -n trident
```

```
-----
| NAME | STORAGE DRIVER | UUID | STATE | VOLUMES |
-----
| NasBackend | ontap-nas | 98e19b74-aec7-4a3d-8dcf-128e5033b214 | online | 9 |
-----
```

NOTE: When rotating passwords, the storage administrator must first update the password for the user on ONTAP. This is followed by a backend update. When rotating certificates, multiple certificates can be added to the user. The backend is then updated to use the new certificate, following which the old certificate can be deleted from the ONTAP cluster.

Updating a backend does not disrupt access to volumes that have already been created, nor impact volume connections made after. A successful backend update indicates that Trident can communicate with the ONTAP backend and handle future volume operations.

=== Create custom ONTAP role for Trident

You can create an ONTAP cluster role with minimum privileges so that you do not have to use the ONTAP admin role to perform operations in Trident. When you include the username in a Trident backend configuration, Trident uses the ONTAP cluster role you created to perform the operations.

Refer to

link:https://github.com/NetApp/trident/tree/master/contrib/ontap/trident_role[Trident custom-role generator] for more information about creating Trident custom roles.

[role="tabbed-block"]

=====

.Using ONTAP CLI

--

. Create a new role using the following command:

+

```
`security login role create <role_name> -cmddirname "command" -access all  
-vserver <svm_name>`
```

. Create a username for the Trident user:

+

```
`security login create -username <user_name> -application ontapi  
-authmethod <password> -role <name_of_role_in_step_1> -vserver  
<svm_name> -comment "user_description"`
```

. Map the role to the user:

+

```
`security login modify username <user_name> -vserver <svm_name> -role  
<role_name> -application ontapi -application console -authmethod  
<password>`
```

--

.Using System Manager

--

Perform the following steps in ONTAP System Manager:

```
. *Create a custom role*:
.. To create a custom role at the cluster-level, select *Cluster >
Settings*.
+
(Or) To create a custom role at the SVM level, select *Storage > Storage
VMs > `required SVM`> Settings > Users and Roles*.
.. Select the arrow icon (*->*) next to *Users and Roles*.
.. Select *+Add* under *Roles*.
.. Define the rules for the role and click *Save*.
. *Map the role to the Trident user*:
+ Perform the following steps on the *Users and Roles* page:
.. Select Add icon *++* under *Users*.
.. Select the required username, and select a role in the drop-down menu
for *Role*.
.. Click *Save*.
```

--

====

Refer to the following pages for more information:

```
* link:https://kb.netapp.com/on-prem/ontap/Ontap\_OS/OS-
KBs/FAQ\_\_Custom\_roles\_for\_administration\_of\_ONTAP[Custom roles for
administration of ONTAP^] or link:https://docs.netapp.com/us-
en/ontap/authentication/define-custom-roles-task.html[Define custom
roles^]
* link:https://docs.netapp.com/us-en/ontap-
automation/rest/rbac\_roles\_users.html#rest-api[Work with roles and users^]
```

== Manage NFS export policies

Trident uses NFS export policies to control access to the volumes that it provisions.

Trident provides two options when working with export policies:

```
* Trident can dynamically manage the export policy itself; in this mode of
operation, the storage administrator specifies a list of CIDR blocks that
represent admissible IP addresses. Trident adds applicable node IPs that
fall in these ranges to the export policy automatically at publish time.
Alternatively, when no CIDRs are specified, all global-scoped unicast IPs
```

found on the node that the volume being published to will be added to the export policy.

* Storage administrators can create an export policy and add rules manually. Trident uses the default export policy unless a different export policy name is specified in the configuration.

=== Dynamically manage export policies

Trident provides the ability to dynamically manage export policies for ONTAP backends. This provides the storage administrator the ability to specify a permissible address space for worker node IPs, rather than defining explicit rules manually. It greatly simplifies export policy management; modifications to the export policy no longer require manual intervention on the storage cluster. Moreover, this helps restrict access to the storage cluster only to worker nodes that are mounting volumes and have IPs in the range specified, supporting a fine-grained and automated management.

NOTE: Do not use Network Address Translation (NAT) when using dynamic export policies. With NAT, the storage controller sees the frontend NAT address and not the actual IP host address, so access will be denied when no match is found in the export rules.

==== Example

There are two configuration options that must be used. Here's an example backend definition:

```
[source,yaml]
```

```
version: 1
storageDriverName: ontap-nas-economy
backendName: ontap_nas_auto_export
managementLIF: 192.168.0.135
svm: svm1
username: vsadmin
password: password
autoExportCIDRs:
- 192.168.0.0/24
autoExportPolicy: true
```

NOTE: When using this feature, you must ensure that the root junction in your SVM has a previously created export policy with an export rule that permits the node CIDR block (such as the default export policy). Always follow NetApp recommended best practice to dedicate an SVM for Trident.

Here is an explanation of how this feature works using the example above:

* ``autoExportPolicy`` is set to ``true``. This indicates that Trident creates an export policy for each volume provisioned with this backend for the ``svml`` SVM and handle the addition and deletion of rules using ``autoexportCIDRs`` address blocks. Until a volume is attached to a node, the volume uses an empty export policy with no rules to prevent unwanted access to that volume. When a volume is published to a node Trident creates an export policy with the same name as the underlying qtree containing the node IP within the specified CIDR block. These IPs will also be added to the export policy used by the parent FlexVol volume

** For example:

```
*** backend UUID 403b5326-8482-40db-96d0-d83fb3f4daec
*** `autoExportPolicy` set to `true`
*** storage prefix `trident`
*** PVC UUID a79bcf5f-7b6d-4a40-9876-e2551f159c1c
*** qtree named trident_pvc_a79bcf5f_7b6d_4a40_9876_e2551f159c1c creates
an export policy for the FlexVol named `trident-403b5326-8482-40db96d0-
d83fb3f4daec`, an export policy for the qtree named
`trident_pvc_a79bcf5f_7b6d_4a40_9876_e2551f159c1c`, and an empty export
policy named `trident_empty` on the SVM. The rules for the FlexVol export
policy will be a superset of any rules contained in the qtree export
policies. The empty export policy will be reused by any volumes that are
not attached.
```

* ``autoExportCIDRs`` contains a list of address blocks. This field is optional and it defaults to `["0.0.0.0/0", ":::/0"]`. If not defined, Trident adds all globally-scoped unicast addresses found on the worker nodes with publications.

In this example, the ``192.168.0.0/24`` address space is provided. This indicates that Kubernetes node IPs that fall within this address range with publications will be added to the export policy that Trident creates. When Trident registers a node it runs on, it retrieves the IP addresses of the node and checks them against the address blocks provided in ``autoExportCIDRs``. At publish time, after filtering the IPs, Trident creates the export policy rules for the client IPs for the node it is publishing to.

You can update ``autoExportPolicy`` and ``autoExportCIDRs`` for backends after you create them. You can append new CIDRs for a backend that is automatically managed or delete existing CIDRs. Exercise care when deleting CIDRs to ensure that existing connections are not dropped. You can also choose to disable ``autoExportPolicy`` for a backend and fall back to a manually created export policy. This will require setting the ``exportPolicy`` parameter in your backend config.

After Trident creates or updates a backend, you can check the backend using ``tridentctl`` or the corresponding ``tridentbackend`` CRD:

`/tridentctl get backends ontap_nas_auto_export -n trident -o yaml`

items:

- backendUUID: 403b5326-8482-40db-96d0-d83fb3f4daec

config:

aggregate: ""

autoExportCIDRs:

- 192.168.0.0/24

autoExportPolicy: true

backendName: ontap_nas_auto_export

chapInitiatorSecret: ""

chapTargetInitiatorSecret: ""

chapTargetUsername: ""

chapUsername: ""

dataLIF: 192.168.0.135

debug: false

debugTraceFlags: null

defaults:

encryption: "false"

exportPolicy: <automatic>

fileSystemType: ext4

When a node is removed, Trident checks all export policies to remove the access rules corresponding to the node. By removing this node IP from the export policies of managed backends, Trident prevents rogue mounts, unless this IP is reused by a new node in the cluster.

For previously existing backends, updating the backend with ``tridentctl update backend`` ensures that Trident manages the export policies automatically. This creates two new export policies named after the backend's UUID and qtree name when they are needed. Volumes that are present on the backend will use the newly created export policies after they are unmounted and mounted again.

NOTE: Deleting a backend with auto-managed export policies will delete the dynamically created export policy. If the backend is re-created, it is treated as a new backend and will result in the creation of a new export policy.

If the IP address of a live node is updated, you must restart the Trident pod on the node. Trident will then update the export policy for backends it manages to reflect this IP change.

== Prepare to provision SMB volumes

With a little additional preparation, you can provision SMB volumes using

`ontap-nas` drivers.

WARNING: You must configure both NFS and SMB/CIFS protocols on the SVM to create an `ontap-nas-economy` SMB volume for ONTAP on-premises clusters. Failure to configure either of these protocols will cause SMB volume creation to fail.

NOTE: `autoExportPolicy` is not supported for SMB volumes.

.Before you begin

Before you can provision SMB volumes, you must have the following.

- * A Kubernetes cluster with a Linux controller node and at least one Windows worker node running Windows Server 2022. Trident supports SMB volumes mounted to pods running on Windows nodes only.

- * At least one Trident secret containing your Active Directory credentials. To generate secret `smbcreds`:

+

```
kubectl create secret generic smbcreds --from-literal username=user --from-literal password='password'
```

- * A CSI proxy configured as a Windows service. To configure a `csi-proxy`, refer to link:<https://github.com/kubernetes-csi/csi-proxy>[GitHub: CSI Proxy^] or link:<https://github.com/Azure/aks-engine/blob/master/docs/topics/csi-proxy-windows.md>[GitHub: CSI Proxy for Windows^] for Kubernetes nodes running on Windows.

.Steps

- . For on-premises ONTAP, you can optionally create an SMB share or Trident can create one for you.

+

NOTE: SMB shares are required for Amazon FSx for ONTAP.

+

You can create the SMB admin shares in one of two ways either using the link:<https://learn.microsoft.com/en-us/troubleshoot/windows-server/system-management-components/what-is-microsoft-management-console>[Microsoft Management Console^] Shared Folders snap-in or using the ONTAP CLI. To create the SMB shares using the ONTAP CLI:

+

- .. If necessary, create the directory path structure for the share.

+

The `vserver cifs share create` command checks the path specified in the

-path option during share creation. If the specified path does not exist, the command fails.

.. Create an SMB share associated with the specified SVM:
+

`vserver cifs share create -vserver vserver_name -share-name share_name -path path [-share-properties share_properties,...] [other_attributes] [-comment text]`

.. Verify that the share was created:
+

`vserver cifs share show -share-name share_name`

+
NOTE: Refer to link:<https://docs.netapp.com/us-en/ontap/smb-config/create-share-task.html>[Create an SMB share^] for full details.

. When creating the backend, you must configure the following to specify SMB volumes. For all FSx for ONTAP backend configuration options, refer to [xref:{relative_path}trident-fsx-examples.html](#)[FSx for ONTAP configuration options and examples].

+
[cols="1,2,1",options="header"]
|===
|Parameter |Description | Example

|`smbShare`
| You can specify one of the following: the name of an SMB share created using the Microsoft Management Console or ONTAP CLI; a name to allow Trident to create the SMB share; or you can leave the parameter blank to prevent common share access to volumes.

This parameter is optional for on-premises ONTAP.

This parameter is required for Amazon FSx for ONTAP backends and cannot be blank.

| `smb-share`

|`nasType`
| *Must set to `smb`. * If null, defaults to `nfs`.
|`smb`

|`securityStyle`

```
|Security style for new volumes.
```

```
*Must be set to `ntfs` or `mixed` for SMB volumes.*
```

```
|`ntfs` or `mixed` for SMB volumes
```

```
|`unixPermissions`
```

```
|Mode for new volumes. *Must be left empty for SMB volumes.*
```

```
|""
```

```
|===
```

```
// Secure SMB - Added for 25.06 (June 2025), Writer: Arul Deepa A
```

```
=== Enable secure SMB
```

Beginning with the 25.06 release, NetApp Trident supports secure provisioning of SMB volumes created using `ontap-nas` and `ontap-nas-economy` backends. When secure SMB is enabled, you can provide controlled access to SMB the shares for Active Directory (AD) users and user groups using Access Control Lists (ACLs).

.Points to remember

- * Importing `ontap-nas-economy` volumes is not supported.

- * Only read-only clones are supported for `ontap-nas-economy` volumes.

- * If Secure SMB is enabled, Trident will ignore the SMB share mentioned in the backend.

- * Updating the PVC annotation, storage class annotation, and backend field does not update the SMB share ACL.

- * The SMB share ACL specified in the annotation of the clone PVC will take precedence over those in the source PVC.

- * Ensure that you provide valid AD users while enabling secure SMB.

Invalid users will not be added to the ACL.

- * If you provide the same AD user in the backend, storage class, and PVC with different permissions, the permission priority will be: PVC, storage class, and then backend.

- * Secure SMB is supported for `ontap-nas` managed volume imports and not applicable to unmanaged volume imports.

.Steps

1. Specify adAdminUser in TridentBackendConfig as shown in the following example:

```
+
```

```
[source,yaml]
```

apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig

```
metadata:
name: backend-tbc-ontap
namespace: trident
spec:
version: 1
storageDriverName: ontap-nas
managementLIF: 10.193.176.x
svm: svm0
useREST: true
defaults:
adAdminUser: tridentADtest
credentials:
name: backend-tbc-ontap-invest-secret
```

```
2. Add the annotation in the storage class.
+
Add the `trident.netapp.io/smbShareAdUser` annotation to the storage class
to enable secure SMB without fail.
The user value specified for the annotation
`trident.netapp.io/smbShareAdUser` should be the same as the username
specified in the `smbcreds` secret.
You can choose one of the following for `smbShareAdUserPermission`:
`full_control`, `change`, or `read`. The default permission is
`full_control`.

[source,yaml]
```

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
name: ontap-smb-sc
annotations:
trident.netapp.io/smbShareAdUserPermission: change
trident.netapp.io/smbShareAdUser: tridentADuser
parameters:
backendType: ontap-nas
csi.storage.k8s.io/node-stage-secret-name: smbcreds
csi.storage.k8s.io/node-stage-secret-namespace: trident
trident.netapp.io/nasType: smb
provisioner: csi.trident.netapp.io
reclaimPolicy: Delete
volumeBindingMode: Immediate
```

```
[start=3]
3. Create a PVC.
+
The following example creates a PVC:
[source,yaml]
```

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
name: my-pvc4
namespace: trident
annotations:
trident.netapp.io/snapshotDirectory: "true"
trident.netapp.io/smbShareAccessControl: |
read:
- tridentADtest
spec:
accessModes:
- ReadWriteOnce
resources:
requests:
storage: 1Gi
storageClassName: ontap-smb-sc
```

```
[[ID9f3b65a2a29fe942b1e333ff0043be30]]
= ONTAP NAS configuration options and examples
:hardbreaks:
:icons: font
:relative_path: ./trident-use/
:imagesdir: {root_path}{relative_path}../media/
```

[.lead]

Create and use ONTAP NAS drivers with your Trident installation. This section provides backend configuration examples and details for mapping backends to StorageClasses. Beginning with the 25.10 release, NetApp Trident supports [link:https://docs.netapp.com/us-en/ontap-afx/index.html](https://docs.netapp.com/us-en/ontap-afx/index.html)[NetApp AFX storage systems^]. NetApp AFX storage systems differ from other ONTAP-based systems (ASA, AFF, and FAS) in the implementation of their storage layer.

NOTE: Only the `ontap-nas` driver (with NFS protocol) is supported for NetApp AFX systems; SMB protocol is not supported.

== Backend configuration options

In the Trident backend configuration, you need not specify that your system is an NetApp AFX storage system. When you select `ontap-nas` as the `storageDriverName`, Trident detects automatically the AFX storage system. Some backend configuration parameters are not applicable to AFX storage systems.

Beginning with the 26.06 release, you can enable ONTAP balanced placement so that ONTAP selects the aggregate for new volumes. Balanced placement is required for AFX storage systems on ONTAP 9.19.1 or later. Refer to

xref:{relative_path}ontap-balanced-placement.html[Use balanced placement with ONTAP backends] for requirements and examples.

The following table displays backend configuration options:

```
[cols="1,3,2",options="header"]
|==
|Parameter |Description |Default
|`version` | |Always 1
```

```
|`storageDriverName` a| Name of the storage driver.
```

Cannot be modified after creation. To update this parameter, you must create a new backend.

NOTE: For NetApp AFX systems, only `ontap-nas` is supported. |`ontap-nas`, `ontap-nas-economy`, or `ontap-nas-flexgroup`

```
|`backendName` |Custom name or the storage backend |Driver name + "_" +
dataLIF
```

```
|`managementLIF` |IP address of a cluster or SVM management LIF
```

A fully-qualified domain name (FQDN) can be specified.

Can be set to use IPv6 addresses if Trident was installed using the IPv6 flag. IPv6 addresses must be defined in square brackets, such as
`[28e8:d9fb:a825:b7bf:69a8:d02f:9e7b:3555]`.

For seamless MetroCluster switchover, see the <<mcc-best>>.

```
|"10.0.0.1", "[2001:1234:abcd::fefe]"
```

```
|`dataLIF` |IP address of protocol LIF.
```

NetApp recommends specifying `dataLIF`. If not provided, Trident fetches dataLIFs from the SVM. You can specify a fully-qualified domain name (FQDN) to be used for the NFS mount operations, allowing you to create a round-robin DNS to load-balance across multiple dataLIFs.

Can be changed after initial setting. Refer to <<Update `dataLIF` after initial configuration>>.

Can be set to use IPv6 addresses if Trident was installed using the IPv6 flag. IPv6 addresses must be defined in square brackets, such as
`[28e8:d9fb:a825:b7bf:69a8:d02f:9e7b:3555]`.

Omit for Metrocluster. See the <<mcc-best>>.

|Specified address or derived from SVM, if not specified (not recommended)

|`svm` |Storage virtual machine to use

Omit for Metrocluster. See the <<mcc-best>>.

|Derived if an SVM `managementLIF` is specified

|`autoExportPolicy` |Enable automatic export policy creation and updating
[Boolean].

Using the `autoExportPolicy` and `autoExportCIDRs` options, Trident can
manage export policies automatically. |false

|`autoExportCIDRs` |List of CIDRs to filter Kubernetes' node IPs against
when `autoExportPolicy` is enabled.

Using the `autoExportPolicy` and `autoExportCIDRs` options, Trident can
manage export policies automatically. |["0.0.0.0/0", "::/0"]`

|`labels` |Set of arbitrary JSON-formatted labels to apply on volumes |""

|`clientCertificate` |Base64-encoded value of client certificate. Used
for certificate-based auth |""

|`clientPrivateKey` |Base64-encoded value of client private key. Used for
certificate-based auth |""

|`trustedCACertificate` |Base64-encoded value of trusted CA certificate.
Optional. Used for certificate-based auth |""

|`username` |Username to connect to the cluster/SVM. Used for credential-
based auth.

For Active Directory authentication, see [xref:{relative_path}../trident-use/ontap-san-examples.html#authenticate-trident-to-a-backend-svm-using-active-directory-credentials](#)[Authenticate Trident to a backend SVM using Active Directory credentials]. |

|`password` |Password to connect to the cluster/SVM. Used for credential-
based auth.

For Active Directory authentication, see [xref:{relative_path}../trident-use/ontap-san-examples.html#authenticate-trident-to-a-backend-svm-using-active-directory-credentials](#)[Authenticate Trident to a backend SVM using Active Directory credentials]. |

|`storagePrefix` a|Prefix used when provisioning new volumes in the SVM. Cannot be updated after you set it

NOTE: When using ontap-nas-economy and a storagePrefix that is 24 or more characters, the qtrees do not have the storage prefix embedded, though it is in the volume name. |"trident"

|`aggregate`

a|Aggregate for provisioning (optional; if set, must be assigned to the SVM). For the `ontap-nas-flexgroup` driver, this option is ignored. If not assigned, any of the available aggregates can be used to provision a FlexGroup volume.

NOTE: When the aggregate is updated in SVM, it is updated in Trident automatically by polling SVM without having to restart the Trident Controller. When you have configured a specific aggregate in Trident to provision volumes, if the aggregate is renamed or moved out of the SVM, the backend is moved to failed state in Trident while polling the SVM aggregate. You must either change the aggregate to one that is present on the SVM or remove it altogether to bring the backend back online.

Do not specify for AFX storage systems.

Do not specify when `useBalancedPlacement` is set to `true`.

a|""

|`useBalancedPlacement`

a|Delegate aggregate selection for new volumes to ONTAP [Boolean]. Requires ONTAP 9.19.1 or later and `useREST` set to `true`. Cannot be combined with `aggregate`.

Refer to xref:{relative_path}ontap-balanced-placement.html[Use balanced placement with ONTAP backends] for requirements and examples.

If specified, always set to `true` for AFX storage systems on ONTAP 9.19.1 or later.

|`false`; `true` for AFX storage systems on ONTAP 9.19.1 or later

|`limitAggregateUsage` |Fail provisioning if usage is above this percentage.

Does not apply to Amazon FSx for ONTAP.

Do not specify for AFX storage systems.|"" (not enforced by default)

|flexgroupAggregateList a| List of aggregates for provisioning (optional;

if set, must be assigned to the SVM). All aggregates assigned to the SVM are used to provision a FlexGroup volume. Supported for the `*ontap-nas-flexgroup*` storage driver.

NOTE: When the aggregate list is updated in SVM, the list is updated in Trident automatically by polling SVM without having to restart the Trident Controller. When you have configured a specific aggregate list in Trident to provision volumes, if the aggregate list is renamed or moved out of SVM, the backend is moved to failed state in Trident while polling the SVM aggregate. You must either change the aggregate list to one that is present on the SVM or remove it altogether to bring the backend back online.

| ""

|`limitVolumeSize` |Fail provisioning if requested volume size is above this value. |"" (not enforced by default)

|`debugTraceFlags` |Debug flags to use when troubleshooting. Example, {"api":false, "method":true}

Do not use `debugTraceFlags` unless you are troubleshooting and require a detailed log dump.|null

|`nasType` | Configure NFS or SMB volumes creation.

Options are `nfs`, `smb` or null. Setting to null defaults to NFS volumes. *If specified, always set to `nfs` for AFX storage systems*.|`nfs`

|`nfsMountOptions` |Comma-separated list of NFS mount options.

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 falls back to using the mount options specified in the storage backend's configuration file.

If no mount options are specified in the storage class or the configuration file, Trident does not set any mount options on an associated persistent volume. |""

|`qtreesPerFlexvol` |Maximum Qtrees per FlexVol, must be in range [50, 300] |"200"

|`smbShare`

|You can specify one of the following: the name of an SMB share created using the Microsoft Management Console or ONTAP CLI; a name to allow Trident to create the SMB share; or you can leave the parameter blank to

prevent common share access to volumes.

This parameter is optional for on-premises ONTAP.

This parameter is required for Amazon FSx for ONTAP backends and cannot be blank.

|`smb-share`

|`useREST` |Boolean parameter to use ONTAP REST APIs.

`useREST` When set to `true`, Trident uses ONTAP REST APIs to communicate with the backend; when set to `false`, Trident uses ONTAPI (ZAPI) calls to communicate with the backend. This feature requires ONTAP 9.11.1 and later. In addition, the ONTAP login role used must have access to the `ontapi` application. This is satisfied by the pre-defined `vsadmin` and `cluster-admin` roles.

Beginning with the Trident 24.06 release and ONTAP 9.15.1 or later, `useREST` is set to `true` by default; change `useREST` to `false` to use ONTAPI (ZAPI) calls.

If specified, always set to `true` for AFX storage systems.

|`true` for ONTAP 9.15.1 or later, otherwise `false`.

|`limitVolumePoolSize` |Maximum requestable FlexVol size when using Qtrees in ontap-nas-economy backend.

| "" (not enforced by default)

|`denyNewVolumePools` |Restricts `ontap-nas-economy` backends from creating new FlexVol volumes to contain their Qtrees. Only preexisting Flexvols are used for provisioning new PVs.

|

|`adAdminUser`

| Active Directory admin user or user group with full access to SMB shares. Use this parameter to provide admin rights to the SMB share with full control. |

|===

== Backend configuration options for provisioning volumes

You can control default provisioning using these options in the `defaults`

section of the configuration. For an example, see the configuration examples below.

```
[cols="1,3,2",options="header"]
|===
|Parameter |Description |Default
|`spaceAllocation` |Space-allocation for Qtrees |"true"

|`spaceReserve` |Space reservation mode; "none" (thin) or "volume" (thick)
|"none"

|`snapshotPolicy` |Snapshot policy to use |"none"

|`qosPolicy` |QoS policy group to assign for volumes created. Choose one
of qosPolicy or adaptiveQosPolicy per storage pool/backend |""

|`adaptiveQosPolicy` |Adaptive QoS policy group to assign for volumes
created. Choose one of qosPolicy or adaptiveQosPolicy per storage
pool/backend.

Not supported by ontap-nas-economy. |""

|`snapshotReserve` |Percentage of volume reserved for snapshots |"0" if
`snapshotPolicy` is "none", otherwise ""

|`splitOnClone` |Split a clone from its parent upon creation |"false"

|`encryption` |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.

If NAE is enabled on the backend, any volume provisioned in Trident is NAE
enabled.

For more information, refer to: xref:{relative_path}../trident-
reco/security-reco.html[How Trident works with NVE and NAE]. |"false"

|`tieringPolicy` |Tiering policy to use "none" |

|`unixPermissions` |Mode for new volumes |"777" for NFS volumes; empty
(not applicable) for SMB volumes

|`snapshotDir` |Controls access to the `.snapshot` directory |`true`,
`false` (Set explicitly).

|`exportPolicy` |Export policy to use |"default"
```

```
|`securityStyle` |Security style for new volumes.
```

NFS supports `mixed` and `unix` security styles.

SMB supports `mixed` and `ntfs` security styles.

```
|NFS default is `unix`.
```

SMB default is `ntfs`.

```
|`nameTemplate` | Template to create custom volume names.
```

```
|""
```

```
|===
```

NOTE: Using QoS policy groups with Trident requires ONTAP 9.8 or later. You should use a non-shared QoS policy group and ensure the policy group is applied to each constituent individually. A shared QoS policy group enforces the ceiling for the total throughput of all workloads.

=== Volume provisioning examples

Here's an example with defaults defined:

```
[source,yaml]
```

```
version: 1
storageDriverName: ontap-nas
backendName: customBackendName
managementLIF: 10.0.0.1
dataLIF: 10.0.0.2
labels:
k8scluster: dev1
backend: dev1-nasbackend
svm: trident_svm
username: cluster-admin
password: <password>
limitAggregateUsage: 80%
limitVolumeSize: 50Gi
nfsMountOptions: nfsvers=4
debugTraceFlags:
api: false
method: true
defaults:
spaceReserve: volume
qosPolicy: premium
exportPolicy: myk8scluster
```

snapshotPolicy: default
snapshotReserve: "10"

For ``ontap-nas`` and ``ontap-nas-flexgroups``, Trident now uses a new calculation to ensure that the FlexVol is sized correctly with the snapshotReserve percentage and PVC. When the user requests a PVC, Trident creates the original FlexVol with more space by using the new calculation. This calculation ensures that the user receives the writable space they requested for in the PVC, and not less space than what they requested. Before v21.07, when the user requests a PVC (for example, 5 GiB), with the snapshotReserve to 50 percent, they get only 2.5 GiB of writeable space. This is because what the user requested for is the whole volume and ``snapshotReserve`` is a percentage of that. With Trident 21.07, what the user requests for is the writeable space and Trident defines the ``snapshotReserve`` number as the percentage of the whole volume. This does not apply to ``ontap-nas-economy``. See the following example to see how this works:

The calculation is as follows:

Total volume size = <PVC requested size> / (1 - (<snapshotReserve percentage> / 100))

For snapshotReserve = 50%, and PVC request = 5 GiB, the total volume size is $5 / .5 = 10$ GiB and the available size is 5 GiB, which is what the user requested in the PVC request. The ``volume show`` command should show results similar to this example:

image:.../media/volume-show-nas.png[Shows the output of the volume show command.]

Existing backends from previous installs, provision volumes as explained above when upgrading Trident. For volumes that you created before upgrading, you should resize their volumes for the change to be observed. For example, a 2 GiB PVC with ``snapshotReserve=50`` earlier resulted in a volume that provides 1 GiB of writable space. Resizing the volume to 3 GiB, for example, provides the application with 3 GiB of writable space on a 6 GiB volume.

== Minimal configuration examples

The following examples show basic configurations that leave most parameters to default. This is the easiest way to define a backend.

NOTE: If you are using Amazon FSx on NetApp ONTAP with Trident, the recommendation is to specify DNS names for LIFs instead of IP addresses.

```
.ONTAP NAS economy example
[%collapsible%closed]
====
[source,yaml]
```

```
version: 1
storageDriverName: ontap-nas-economy
managementLIF: 10.0.0.1
dataLIF: 10.0.0.2
svm: svm_nfs
username: vsadmin
password: password
```

```
====

.ONTAP NAS Flexgroup example
[%collapsible%closed]
====
[source,yaml]
```

```
version: 1
storageDriverName: ontap-nas-flexgroup
managementLIF: 10.0.0.1
dataLIF: 10.0.0.2
svm: svm_nfs
username: vsadmin
password: password
```

```
====

[#mcc-best]
.MetroCluster example
[%collapsible%closed]
====

You can configure the backend to avoid having to manually update the
backend definition after switchover and switchback during
xref:{relative_path}../trident-reco/backup.html#svm-replication-and-
recovery[SVM replication and recovery].

For seamless switchover and switchback, specify the SVM using
`managementLIF` and omit the `dataLIF` and `svm` parameters. For example:
```

```
[source,yaml]
```

```
version: 1
storageDriverName: ontap-nas
managementLIF: 192.168.1.66
username: vsadmin
password: password
```

```
====
```

```
.SMB volumes example
[%collapsible%closed]
====
[source,yaml]
```

```
version: 1
backendName: ExampleBackend
storageDriverName: ontap-nas
managementLIF: 10.0.0.1
nasType: smb
securityStyle: ntfs
unixPermissions: ""
dataLIF: 10.0.0.2
svm: svm_nfs
username: vsadmin
password: password
```

```
====
```

```
.Certificate-based authentication example
[%collapsible%closed]
```

```
====
```

```
This is a minimal backend configuration example. `clientCertificate`,
`clientPrivateKey`, and `trustedCACertificate` (optional, if using trusted
CA) are populated in `backend.json` and take the base64-encoded values of
the client certificate, private key, and trusted CA certificate,
respectively.
[source,yaml]
```

```
version: 1
backendName: DefaultNASBackend
```

storageDriverName: ontap-nas
managementLIF: 10.0.0.1
dataLIF: 10.0.0.15
svm: nfs_svm
clientCertificate: ZXR0ZXJwYXB...lCMgJ3BhcGVyc2
clientPrivateKey: vciwKlyAgZG...0cnkslGRlc2NyaX
trustedCACertificate: zcyBbaG...b3lgb3dulGNsYXNz
storagePrefix: myPrefix_

====

.Auto export policy example
[%collapsible%closed]

====

This example shows you how you can instruct Trident to use dynamic export policies to create and manage the export policy automatically. This works the same for the `ontap-nas-economy` and `ontap-nas-flexgroup` drivers.
[source,yaml]

version: 1
storageDriverName: ontap-nas
managementLIF: 10.0.0.1
dataLIF: 10.0.0.2
svm: svm_nfs
labels:
k8scluster: test-cluster-east-1a
backend: test1-nasbackend
autoExportPolicy: true
autoExportCIDRs:
- 10.0.0.0/24
username: admin
password: password
nfsMountOptions: nfsvers=4

====

.Balanced placement example
[%collapsible%closed]

====

This example enables ONTAP balanced placement so that ONTAP selects the aggregate for new volumes. Balanced placement requires ONTAP 9.19.1 or later. Refer to [xref:{relative_path}ontap-balanced-placement.html](#)[Use balanced placement with ONTAP backends] for details.
[source,yaml]

version: 1

```
storageDriverName: ontap-nas
backendName: nas-balanced-placement
managementLIF: 10.0.0.1
dataLIF: 10.0.0.2
svm: svm_nfs
username: vsadmin
password: password
useBalancedPlacement: true
```

```
====
.IPv6 addresses example
[%collapsible%closed]
====
This example shows `managementLIF` using an IPv6 address.
[source,yaml]
```

```
version: 1
storageDriverName: ontap-nas
backendName: nas_ipv6_backend
managementLIF: "[5c5d:5edf:8f:7657:bef8:109b:1b41:d491]"
labels:
k8scluster: test-cluster-east-1a
backend: test1-ontap-ipv6
svm: nas_ipv6_svm
username: vsadmin
password: password
```

```
====
.Amazon FSx for ONTAP using SMB volumes example
[%collapsible%closed]
====
The `smbShare` parameter is required for FSx for ONTAP using SMB volumes.
[source,yaml]
```

```
version: 1
backendName: SMBBackend
storageDriverName: ontap-nas
managementLIF: example.mgmt.fqdn.aws.com
nasType: smb
dataLIF: 10.0.0.15
svm: nfs_svm
smbShare: smb-share
clientCertificate: ZXR0ZXJwYXB...ICMgJ3BhcGVyc2
clientPrivateKey: vciwKlyAgZG...0cnkslGRlc2NyaX
trustedCACertificate: zcyBbaG...b3lgb3dulGNsYXNz
```

storagePrefix: myPrefix_

```
====
```

```
.Backend configuration example with nameTemplate  
[%collapsible%closed]
```

```
====
```

```
[source,yaml]
```

```
version: 1  
storageDriverName: ontap-nas  
backendName: ontap-nas-backend  
managementLIF: <ip address>  
svm: svm0  
username: <admin>  
password: <password>  
defaults:  
nameTemplate: "{{.volume.Name}}_{{.labels.cluster}}_{{.volume.Namespace}}_{{.volume.RequestName}}"  
labels:  
cluster: ClusterA  
PVC: "{{.volume.Namespace}}_{{.volume.RequestName}}"
```

```
====
```

```
== Examples of backends with virtual pools
```

In the sample backend definition files shown below, specific defaults are set for all storage pools, such as `spaceReserve` at none, `spaceAllocation` at false, and `encryption` at false. The virtual pools are defined in the storage section.

Trident sets provisioning labels in the "Comments" field. Comments are set on FlexVol for `ontap-nas` or FlexGroup for `ontap-nas-flexgroup`. Trident copies all labels present on a virtual pool to the storage volume at provisioning. For convenience, storage administrators can define labels per virtual pool and group volumes by label.

In these examples, some of the storage pools set their own `spaceReserve`, `spaceAllocation`, and `encryption` values, and some pools override the default values.

```
.ONTAP NAS example  
[%collapsible%open]
```

```
====
```

```
version: 1
storageDriverName: ontap-nas
managementLIF: 10.0.0.1
svm: svm_nfs
username: admin
password: <password>
nfsMountOptions: nfsvers=4
defaults:
spaceReserve: none
encryption: "false"
qosPolicy: standard
labels:
store: nas_store
k8scluster: prod-cluster-1
region: us_east_1
storage:
- labels:
app: msoffice
cost: "100"
zone: us_east_1a
defaults:
spaceReserve: volume
encryption: "true"
unixPermissions: "0755"
adaptiveQosPolicy: adaptive-premium
- labels:
app: slack
cost: "75"
zone: us_east_1b
defaults:
spaceReserve: none
encryption: "true"
unixPermissions: "0755"
- labels:
department: legal
creditpoints: "5000"
zone: us_east_1b
defaults:
spaceReserve: none
encryption: "true"
unixPermissions: "0755"
- labels:
app: wordpress
cost: "50"
zone: us_east_1c
defaults:
spaceReserve: none
encryption: "true"
unixPermissions: "0775"
```

- labels:
app: mysqldb
cost: "25"
zone: us_east_1d
defaults:
spaceReserve: volume
encryption: "false"
unixPermissions: "0775"

====

.ONTAP NAS FlexGroup example
[%collapsible%open]

====

[source,yaml]

version: 1
storageDriverName: ontap-nas-flexgroup
managementLIF: 10.0.0.1
svm: svm_nfs
username: vsadmin
password: <password>
defaults:
spaceReserve: none
encryption: "false"
labels:
store: flexgroup_store
k8scluster: prod-cluster-1
region: us_east_1
storage:
- labels:
protection: gold
creditpoints: "50000"
zone: us_east_1a
defaults:
spaceReserve: volume
encryption: "true"
unixPermissions: "0755"
- labels:
protection: gold
creditpoints: "30000"
zone: us_east_1b
defaults:
spaceReserve: none
encryption: "true"
unixPermissions: "0755"
- labels:
protection: silver
creditpoints: "20000"
zone: us_east_1c

```
defaults:
spaceReserve: none
encryption: "true"
unixPermissions: "0775"
- labels:
protection: bronze
creditpoints: "10000"
zone: us_east_1d
defaults:
spaceReserve: volume
encryption: "false"
unixPermissions: "0775"
```

```
====

.ONTAP NAS economy example
[%collapsible%open]
====

[source,yaml]
```

```
version: 1
storageDriverName: ontap-nas-economy
managementLIF: 10.0.0.1
svm: svm_nfs
username: vsadmin
password: <password>
defaults:
spaceReserve: none
encryption: "false"
labels:
store: nas_economy_store
region: us_east_1
storage:
- labels:
department: finance
creditpoints: "6000"
zone: us_east_1a
defaults:
spaceReserve: volume
encryption: "true"
unixPermissions: "0755"
- labels:
protection: bronze
creditpoints: "5000"
zone: us_east_1b
defaults:
spaceReserve: none
encryption: "true"
unixPermissions: "0755"
- labels:
```

```

department: engineering
creditpoints: "3000"
zone: us_east_1c
defaults:
spaceReserve: none
encryption: "true"
unixPermissions: "0775"
- labels:
department: humanresource
creditpoints: "2000"
zone: us_east_1d
defaults:
spaceReserve: volume
encryption: "false"
unixPermissions: "0775"

```

```
====
```

```
== Map backends to StorageClasses
```

The following StorageClass definitions refer to <<Examples of backends with virtual pools>>. Using the `parameters.selector` field, each StorageClass calls out which virtual pools can be used to host a volume. The volume has the aspects defined in the chosen virtual pool.

* The `protection-gold` StorageClass maps to the first and second virtual pool in the `ontap-nas-flexgroup` backend. These are the only pools offering gold level protection.

```
+
```

```
[source,yaml]
```

```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
name: protection-gold
provisioner: csi.trident.netapp.io
parameters:
selector: "protection=gold"
fsType: "ext4"

```

* The `protection-not-gold` StorageClass maps to the third and fourth virtual pool in the `ontap-nas-flexgroup` backend. These are the only pools offering protection level other than gold.

```
+
```

```
[source,yaml]
```

```

apiVersion: storage.k8s.io/v1
kind: StorageClass

```

metadata:
name: protection-not-gold
provisioner: csi.trident.netapp.io
parameters:
selector: "protection!=gold"
fsType: "ext4"

```
* The `app-mysqldb` StorageClass maps to the fourth virtual pool in the
`ontap-nas` backend. This is the only pool offering storage pool
configuration for mysqldb type app.
+
[source,yaml]
```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
name: app-mysqldb
provisioner: csi.trident.netapp.io
parameters:
selector: "app=mysqldb"
fsType: "ext4"

```
* The `protection-silver-creditpoints-20k` StorageClass maps to the third
virtual pool in the `ontap-nas-flexgroup` backend. This is the only pool
offering silver-level protection and 20000 creditpoints.
+
[source,yaml]
```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
name: protection-silver-creditpoints-20k
provisioner: csi.trident.netapp.io
parameters:
selector: "protection=silver; creditpoints=20000"
fsType: "ext4"

```
* The `creditpoints-5k` StorageClass maps to the third virtual pool in the
`ontap-nas` backend and the second virtual pool in the `ontap-nas-economy`
backend. These are the only pool offerings with 5000 creditpoints.
+
[source,yaml]
```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
name: creditpoints-5k

```
provisioner: csi.trident.netapp.io
parameters:
selector: "creditpoints=5000"
fsType: "ext4"
```

Trident decides which virtual pool is selected and ensures the storage requirement is met.

== Update `dataLIF` after initial configuration

You can change the dataLIF after initial configuration by running the following command to provide the new backend JSON file with updated dataLIF.

```
tridentctl update backend <backend-name> -f <path-to-backend-json-file-with-updated-dataLIF>
```

NOTE: If PVCs are attached to one or multiple pods, you must bring down all corresponding pods and then bring them back up in order for the new dataLIF to take effect.

== Secure SMB examples

=== Backend configuration with ontap-nas driver

[source,yaml]

```
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
name: backend-tbc-ontap-nas
namespace: trident
spec:
version: 1
storageDriverName: ontap-nas
managementLIF: 10.0.0.1
svm: svm2
nasType: smb
defaults:
adAdminUser: tridentADtest
credentials:
name: backend-tbc-ontap-invest-secret
```

=== Backend configuration with ontap-nas-economy driver

[source,yaml]

```
apiVersion: trident.netapp.io/v1
```

kind: TridentBackendConfig
metadata:
name: backend-tbc-ontap-nas
namespace: trident
spec:
version: 1
storageDriverName: ontap-nas-economy
managementLIF: 10.0.0.1
svm: svm2
nasType: smb
defaults:
adAdminUser: tridentADtest
credentials:
name: backend-tbc-ontap-invest-secret

```
=== Backend configuration with storage pool
```

```
[source,yaml]
```

apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
name: backend-tbc-ontap-nas
namespace: trident
spec:
version: 1
storageDriverName: ontap-nas
managementLIF: 10.0.0.1
svm: svm0
useREST: false
storage:
- labels:
app: msoffice
defaults:
adAdminUser: tridentADuser
nasType: smb
credentials:
name: backend-tbc-ontap-invest-secret

```
=== Storage class example with ontap-nas driver
```

```
[source,yaml]
```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
name: ontap-smb-sc
annotations:
trident.netapp.io/smbShareAdUserPermission: change
trident.netapp.io/smbShareAdUser: tridentADtest

parameters:
backendType: ontap-nas
csi.storage.k8s.io/node-stage-secret-name: smbcreds
csi.storage.k8s.io/node-stage-secret-namespace: trident
trident.netapp.io/nasType: smb
provisioner: csi.trident.netapp.io
reclaimPolicy: Delete
volumeBindingMode: Immediate

NOTE: Ensure that you add `annotations` to enable secure SMB. Secure SMB does not work without the annotations, irrespective of configurations set in the Backend or PVC.

=== Storage class example with ontap-nas-economy driver

[source,yaml]

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
name: ontap-smb-sc
annotations:
trident.netapp.io/smbShareAdUserPermission: change
trident.netapp.io/smbShareAdUser: tridentADuser3
parameters:
backendType: ontap-nas-economy
csi.storage.k8s.io/node-stage-secret-name: smbcreds
csi.storage.k8s.io/node-stage-secret-namespace: trident
trident.netapp.io/nasType: smb
provisioner: csi.trident.netapp.io
reclaimPolicy: Delete
volumeBindingMode: Immediate

=== PVC example with a single AD user

[source,yaml]

apiVersion: v1
kind: PersistentVolumeClaim
metadata:
name: my-pvc4
namespace: trident
annotations:
trident.netapp.io/smbShareAccessControl: |
change:
- tridentADtest
read:
- tridentADuser
spec:

```
accessModes:
- ReadWriteOnce
resources:
requests:
storage: 1Gi
storageClassName: ontap-smb-sc
```

```
=== PVC example with multiple AD users
```

```
[source,yaml]
```

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
name: my-test-pvc
annotations:
trident.netapp.io/smbShareAccessControl: |
full_control:
- tridentTestuser
- tridentuser
- tridentTestuser1
- tridentuser1
change:
- tridentADuser
- tridentADuser1
- tridentADuser4
- tridentTestuser2
read:
- tridentTestuser2
- tridentTestuser3
- tridentADuser2
- tridentADuser3
spec:
accessModes:
- ReadWriteOnce
resources:
requests:
storage: 1Gi
```

```
:leveloffset: -1

[[ID69228bf6fc125ae66c1fe432a75d408e]]
= Use balanced placement with ONTAP backends
:hardbreaks:
:icons: font
:linkattrs:
:relative_path: ./
:relative_path: ./trident-use/
```

```
:imagesdir: {root_path}{relative_path}../media/
```

```
[.lead]
```

You can enable ONTAP balanced placement on ONTAP backends. With balanced placement, ONTAP selects the optimal aggregate for each new volume, so you do not need to manage aggregate assignment in Trident.

== How balanced placement works

By default, Trident discovers the aggregates assigned to an SVM, treats each aggregate as a storage pool, and selects an aggregate when it provisions a volume.

When you enable balanced placement:

- * Trident reports the entire SVM as a single storage pool instead of reporting individual aggregates.
- * ONTAP selects the aggregate for each new volume and balances volumes across the available aggregates.
- * Trident creates volumes using the ONTAP `/application/containers` REST API.`

Balanced placement is available for unified ONTAP systems (AFF, FAS, and ASA) and NetApp AFX storage systems.

== Requirements

- * The backend must use ONTAP 9.19.1 or later.
- * Trident must communicate with ONTAP using REST. Balanced placement is not available with ZAPI.
- * The backend configuration must not specify aggregates. You cannot combine balanced placement with explicit aggregate selection.

== Default values

The default value of `useBalancedPlacement` depends on the storage system:`

```
[cols="1,1,2", options="header"]
```

```
|==
```

```
|Storage system |Default |Notes
```

```
|Unified ONTAP (AFF, FAS, ASA)
```

```
|`false`
```

```
|Trident selects aggregates unless you enable balanced placement.
```

```
|NetApp AFX (ONTAP 9.19.1 or later)
```

```
|`true`
```

```
|Balanced placement is required. Setting `useBalancedPlacement` to `false` causes backend validation to fail.
```

```
|NetApp AFX (earlier than ONTAP 9.19.1)
|`false`
|Balanced placement is not available.
|===
```

== Configuration option

```
[cols="1,2,1", options="header"]
|===
|Parameter |Description |Default
```

```
|`useBalancedPlacement`
|Delegates aggregate selection for new volumes to ONTAP. Set this option at the root of the backend configuration or on individual virtual pools.
|`false`; `true` for AFX on ONTAP 9.19.1 or later
|===
```

== Balanced placement with virtual pools

Balanced placement works with or without virtual pools:

- * Without virtual pools, Trident reports the SVM as a single storage pool and creates all new volumes using balanced placement.
- * With virtual pools, Trident reports the virtual pools and matches volumes to them as usual. If the matching virtual pool has balanced placement enabled, ONTAP selects the aggregate for the volume.

Set `useBalancedPlacement` at the root of the backend configuration to apply it to all virtual pools, or set it on individual virtual pools to enable it selectively.

== Examples

The following example enables balanced placement on a unified ONTAP backend:

```
version: 1
storageDriverName: ontap-nas
backendName: nas-balanced-placement
managementLIF: 10.0.0.1
dataLIF: 10.0.0.2
svm: trident_svm
```

```
username: cluster-admin
password: <password>
useBalancedPlacement: true
```

The following example enables balanced placement using
`TridentBackendConfig`:

```
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: backend-tbc-ontap-nas
spec:
  version: 1
  storageDriverName: ontap-nas
  backendName: nas-balanced-placement
  managementLIF: 10.0.0.1
  svm: trident_svm
  useBalancedPlacement: true
  credentials:
    name: backend-tbc-ontap-nas-secret
```

The following example enables balanced placement on the `gold` virtual pool only. Trident continues to select aggregates for volumes that match the `bronze` virtual pool:

```
version: 1
storageDriverName: ontap-nas
backendName: nas-virtual-pools
managementLIF: 10.0.0.1
svm: trident_svm
username: cluster-admin
password: <password>
storage:
  - labels:
    performance: gold
  useBalancedPlacement: true
  - labels:
    performance: bronze
```

== Backend validation errors

Backend creation or update fails in the following cases:

```
[cols="2,2", options="header"]
|==
|Configuration |Result
```

```
|`useBalancedPlacement: true` on ONTAP earlier than 9.19.1
|Validation fails. Upgrade to ONTAP 9.19.1 or later.

|`useBalancedPlacement: true` on a backend using ZAPI
|Validation fails. Configure the backend to use REST.

|`useBalancedPlacement: true` combined with specific aggregates
|Validation fails. Remove the aggregate configuration or disable balanced
placement.

|`useBalancedPlacement: false` on AFX with ONTAP 9.19.1 or later
|Validation fails. Balanced placement is required for AFX on ONTAP 9.19.1
or later.
|===
```

= Amazon FSx for NetApp ONTAP

```
:leveloffset: +1
```

```
[[ID0c2db12fe24e0c9049260e65f5f817a4]]
```

= Use Trident with Amazon FSx for NetApp ONTAP

```
:hardbreaks:
:icons: font
:relative_path: ./trident-use/
:imagesdir: {root_path}{relative_path}../media/
```

```
[.lead]
```

<https://docs.aws.amazon.com/fsx/latest/ONTAPGuide/what-is-fsx-ontap.html>[Amazon FSx for NetApp ONTAP^] is a fully managed AWS service that runs file systems powered by the NetApp ONTAP storage operating system. It provides ONTAP features, performance, and administration with the scalability and operational simplicity of AWS. A file system is the primary resource in Amazon FSx and is analogous to an on-premises ONTAP cluster. Each file system contains one or more storage virtual machines (SVMs), and each SVM contains one or more volumes that store files and directories. This integration enables Kubernetes clusters running in Amazon Elastic Kubernetes Service (EKS) to provision ONTAP-backed persistent volumes for block and file workloads.

== Requirements

In addition to `xref:{relative_path}../trident-get-started/requirements.html`[Trident requirements], to integrate FSx for

ONTAP with Trident, you need:

- * An existing Amazon EKS cluster or self-managed Kubernetes cluster with `kubectl` installed.
- * An existing Amazon FSx for NetApp ONTAP file system and storage virtual machine (SVM) that is reachable from your cluster's worker nodes.
- * Worker nodes that are prepared for `xref:{relative_path}worker-node-prep.html`[NFS or iSCSI].

+

NOTE: Ensure you follow the node preparation steps required for Amazon Linux and Ubuntu

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIs.html>[Amazon Machine Images^] (AMIs) depending on your EKS AMI type.

== Considerations

* SMB volumes:

- ** SMB volumes are supported using the `ontap-nas` driver only.
- ** SMB volumes are not supported with Trident EKS add-on.
- ** Trident supports SMB volumes mounted to pods running on Windows nodes only. Refer to `xref:{relative_path}worker-node-prep.html#prepare-to-provision-smb-volumes`[Prepare to provision SMB volumes] for details.
- * Prior to Trident 24.02, volumes created on Amazon FSx file systems that have automatic backups enabled, could not be deleted by Trident. To prevent this issue in Trident 24.02 or later, specify the `fsxFileSystemID`, AWS `apiRegion`, AWS `apiKey`, and AWS `secretKey` in the backend configuration file for AWS FSx for ONTAP.

+

NOTE: If you are specifying an IAM role to Trident, then you can omit specifying the `apiRegion`, `apiKey`, and `secretKey` fields to Trident explicitly. For more information, refer to `xref:{relative_path}../trident-use/trident-fsx-examples.html`[FSx for ONTAP configuration options and examples].

=== Simultaneous usage of Trident SAN/iSCSI and EBS-CSI driver

If you plan to use `ontap-san` drivers (e.g., iSCSI) with AWS (EKS, ROSA, EC2, or any other instance), the multipath configuration required on the nodes might conflict with the Amazon Elastic Block Store (EBS) CSI driver. To ensure that multipathing functions without interfering with EBS disks on the same node, you need to exclude EBS in your multipathing setup. This example shows a `multipath.conf` file that includes the required Trident settings while excluding EBS disks from multipathing:

```
defaults {  
find_multipaths no  
}
```

```

blacklist {
device {
vendor "NVME"
product "Amazon Elastic Block Store"
}
}

```

== Authentication

Trident offers two modes of authentication.

* **Credential-based(Recommended):** Stores credentials securely in AWS Secrets Manager. You can use the ``fsxadmin`` user for your file system or the ``vsadmin`` user configured for your SVM.

+

WARNING: Trident expects to be run as a ``vsadmin`` SVM user or as a user with a different name that has the same role. Amazon FSx for NetApp ONTAP has an ``fsxadmin`` user that is a limited replacement of the ONTAP ``admin`` cluster user. We strongly recommend using ``vsadmin`` with Trident.

* **Certificate-based:** Trident will communicate with the SVM on your FSx file system using a certificate installed on your SVM.

For details on enabling authentication, refer to the authentication for your driver type:

* `xref:{relative_path}ontap-nas-prep.html` [ONTAP NAS authentication]
 * `xref:{relative_path}ontap-san-prep.html` [ONTAP SAN authentication]

== Tested Amazon Machine Images (AMIs)

EKS cluster supports various operating systems, but AWS has optimized certain Amazon Machine Images (AMIs) for containers and EKS. The following AMIs have been tested with NetApp Trident 25.02.

```

|===
| AMI | NAS | NAS-economy | iSCSI | iSCSI-economy
| AL2023_x86_64_STANDARD | Yes | Yes | Yes | Yes
| AL2_x86_64 | Yes | Yes | Yes* | Yes*
| BOTTLEROCKET_x86_64 | Yes** | Yes | N/A | N/A
| AL2023_ARM_64_STANDARD | Yes | Yes | Yes | Yes
| AL2_ARM_64 | Yes | Yes | Yes* | Yes*
| BOTTLEROCKET_ARM_64 | Yes** | Yes | N/A | N/A
|===

```

* * Unable to delete the PV without restarting the node

* ** Doesn't work with NFSv3 with Trident version 25.02.

NOTE: If your desired AMI is not listed here, it does not mean that it is

not supported; it simply means it has not been tested. This list serves as a guide for AMIs are known to work.

Tests performed with:

- * EKS version: 1.32
- * Installation Method: Helm 25.06 and as an AWS add-On 25.06
- * For NAS both NFSv3 and NFSv4.1 were tested.
- * For SAN only iSCSI was tested, not NVMe-oF.

Tests performed:

- * Create: Storage Class, pvc, pod
- * Delete: pod, pvc (regular, qtree/lun - economy, NAS with AWS backup)

== Find more information

- * <https://docs.aws.amazon.com/fsx/latest/ONTAPGuide/what-is-fsx-ontap.html>[Amazon FSx for NetApp ONTAP documentation^]
- * <https://www.netapp.com/blog/amazon-fsx-for-netapp-ontap/>[Blog post on Amazon FSx for NetApp ONTAP^]

```
[[ID13d686507798e3a130195377bd71d154]]
```

= Create an IAM role and AWS Secret

:hardbreaks:

:icons: font

:relative_path: ./trident-use/

:imagesdir: {root_path}{relative_path}../media/

[.lead]

You can configure Kubernetes pods to access AWS resources by authenticating as an AWS IAM role instead of by providing explicit AWS credentials.

[NOTE]

To authenticate using an AWS IAM role, you must have a Kubernetes cluster deployed using EKS.

== Create AWS Secrets Manager secret

Since Trident will be issuing APIs against an FSx vserver to manage the storage for you, it will need credentials to do so. The secure way to pass those credentials is through an AWS Secrets Manager secret. Therefore, if you don't already have one, you'll need to create an AWS Secrets Manager secret that contains the credentials for the vsadmin account.

This example creates an AWS Secrets Manager secret to store Trident CSI credentials:

[source,console]

```
aws secretsmanager create-secret --name trident-secret --description "Trident CSI credentials" \
--secret-string "{\"username\":\"vsadmin\", \"password\":\"<svmpassword>\"}"
```

== Create IAM Policy

Trident also needs AWS permissions to run correctly. Therefore, you need to create a policy that gives Trident the permissions it needs.

The following examples creates an IAM policy using the AWS CLI:

[source,console]

```
aws iam create-policy --policy-name AmazonFSxNCSIDriverPolicy --policy-document file://policy.json \
--description "This policy grants access to Trident CSI to FSxN and Secrets manager"
```

Policy JSON example:

[source,json]

```
{
  "Statement": [
    {
      "Action": [
        "fsx:DescribeFileSystems",
        "fsx:DescribeVolumes",
        "fsx:CreateVolume",
        "fsx:RestoreVolumeFromSnapshot",
        "fsx:DescribeStorageVirtualMachines",
        "fsx:UntagResource",
        "fsx:UpdateVolume",
        "fsx:TagResource",
        "fsx>DeleteVolume"
      ],
      "Effect": "Allow",
      "Resource": ""
    },
    {
      "Action": "secretsmanager:GetSecretValue",
      "Effect": "Allow",
      "Resource": "arn:aws:secretsmanager:<aws-region>:<aws-account-id>:secret:<aws-secret-manager-name>"
    }
  ],
}
```

```
"Version": "2012-10-17"
}
```

```
////
```

To enable automatic backend configuration for Amazon FSx, add the following actions to the `policy.json` file while creating an IAM policy:

```
* `fsx:CreateStorageVirtualMachine`
* `fsx:DescribeStorageVirtualMachines`
* `secretsmanager:CreateSecret`
* `secretsmanager>DeleteSecret`
* `secretsmanager:TagResource`
```

Policy JSON file example for automatic backend configuration:

```
{
  "Statement": [
    {
      "Action": [
        "fsx:CreateStorageVirtualMachine",
        "fsx:DescribeFileSystems",
        "fsx:DescribeStorageVirtualMachines",
        "fsx:DescribeVolumes",
        "fsx:CreateVolume",
        "fsx:RestoreVolumeFromSnapshot",
        "fsx:DescribeStorageVirtualMachines",
        "fsx:UntagResource",
        "fsx:UpdateVolume",
        "fsx:TagResource",
        "fsx>DeleteVolume"
      ],
      "Effect": "Allow",
      "Resource": ""
    },
    {
      "Action": [
        "secretsmanager:GetSecretValue",
        "secretsmanager:CreateSecret",
        "secretsmanager>DeleteSecret",
        "secretsmanager:TagResource"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:secretsmanager:<aws-region>:<aws-account-id>:secret:"
    }
  ],
  "Version": "2012-10-17"
}
```

```
////
```

```
=== Create Pod Identity or IAM role for Service account association (IRSA)
```

You can configure a Kubernetes service account to assume an AWS Identity and Access Management (IAM) role with EKS Pod Identity or IAM role for Service account association (IRSA). Any Pods that are configured to use the service account can then access any AWS service that the role has permissions to access.

```
[role="tabbed-block"]
```

```
=====
```

```
.Pod Identity
```

```
--
```

Amazon EKS Pod Identity associations provide the ability to manage credentials for your applications, similar to the way that Amazon EC2 instance profiles provide credentials to Amazon EC2 instances.

***Install Pod Identity on your EKS cluster*:**

You can create Pod identity via the AWS console or using the following AWS CLI command:

```
aws eks create-addon --cluster-name <EKS_CLUSTER_NAME> --addon-name eks-pod-identity-agent
```

For more information refer to
link:<https://docs.aws.amazon.com/eks/latest/userguide/pod-id-agent-setup.html>[Set up the Amazon EKS Pod Identity Agent].

***Create trust-relationship.json*:**

Create trust-relationship.json to enable EKS Service Principal to assume this role for Pod Identity. Then create a role with this trust policy:

```
aws iam create-role \
--role-name fsxn-csi-role --assume-role-policy-document file://trust-relationship.json \
--description "fsxn csi pod identity role"
```

***trust-relationship.json file*:**

```
[source,JSON]
```

```
{
"Version": "2012-10-17",
"Statement": [
```

```
{
"Effect": "Allow",
"Principal": {
"Service": "pods.eks.amazonaws.com"
},
"Action": [
"sts:AssumeRole",
"sts:TagSession"
]
}
]
```

***Attach the role policy to the IAM role*:**

Attach the role policy from the previous step to the IAM role that was created:

```
aws iam attach-role-policy \
--policy-arn arn:aws:iam::aws:111122223333:policy/fsxn-csi-policy \
--role-name fsxn-csi-role
```

***Create a pod identity association*:**

Create a pod identity association between IAM role and the Trident service account (trident-controller)

```
aws eks create-pod-identity-association \
--cluster-name <EKS_CLUSTER_NAME> \
--role-arn arn:aws:iam::111122223333:role/fsxn-csi-role \
--namespace trident --service-account trident-controller
```

--

.IAM role for Service account association (IRSA)

--

***Using the AWS CLI*:**

```
aws iam create-role --role-name AmazonEKS_FSxN_CSI_DriverRole \
--assume-role-policy-document file:///trust-relationship.json
```

***trust-relationship.json file*:**

```
[source,JSON]
```

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Federated": "arn:aws:iam::<account_id>:oidc-provider/<oidc_provider>"
      },
      "Action": "sts:AssumeRoleWithWebIdentity",
      "Condition": {
        "StringEquals": {
          "<oidc_provider>:aud": "sts.amazonaws.com",
          "<oidc_provider>:sub": "system:serviceaccount:trident:trident-controller"
        }
      }
    }
  ]
}
```

Update the following values in the `trust-relationship.json` file:

```
* *<account_id>* - Your AWS account ID
* *<oidc_provider>* - The OIDC of your EKS cluster. You can obtain the
oidc_provider by running:
+
[source,console]
```

```
aws eks describe-cluster --name my-cluster --query "cluster.identity.oidc.issuer" \
--output text | sed -e "s/^https:\/\///"
```

***Attach the IAM role with the IAM policy*:**

Once the role has been created, attach the policy (that was created in the step above) to the role using this command:

```
[source,console]
```

```
aws iam attach-role-policy --role-name my-role --policy-arn <IAM policy ARN>
```

***Verify OICD provider is associated*:**

Verify that your OIDC provider is associated with your cluster. You can verify it using this command:

```
[source,console]
```

```
aws iam list-open-id-connect-providers | grep $oidc_id | cut -d "/" -f4
```

If the output is empty, use the following command to associate IAM OIDC to your cluster:

```
[source,console]
```

```
eksctl utils associate-iam-oidc-provider --cluster $cluster_name --approve
```

If you are using eksctl, use the following example to create an IAM role for service account in EKS:

```
[source,console]
```

```
eksctl create iamserviceaccount --name trident-controller --namespace trident \
--cluster <my-cluster> --role-name AmazonEKS_FSxN_CSI_DriverRole --role-only \
--attach-policy-arn <IAM-Policy ARN> --approve
```

```
--
```

```
====
```

```
[[ID148595af6461fb2b1c8b947aa890a87f]]
= Install Trident
:hardbreaks:
:icons: font
:relative_path: ./trident-use/
:imagesdir: {root_path}{relative_path}../media/
```

```
[.lead]
```

Trident streamlines Amazon FSx for NetApp ONTAP storage management in Kubernetes to enable your developers and administrators focus on application deployment. You can install Trident using one of the following methods:

- * Helm
- * EKS add-on

If you want to make use of the snapshot functionality, install the CSI snapshot controller add-on. Refer to

```
link:https://docs.aws.amazon.com/eks/latest/userguide/csi-snapshot-controller.html[Enable snapshot functionality for CSI volumes^] for more information.
```

```
== Install Trident via helm
```

```
[role="tabbed-block"]
```

```
=====
```

```
.Pod Identity
```

```
--
```

```
. Add the Trident Helm repository:
```

```
+
```

```
[source,console]
```

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

```
+
```

```
. Install Trident using the following example:
```

```
+
```

```
[source,console]
```

```
helm install trident-operator netapp-trident/trident-operator --version 100.2502.1 --namespace trident --create-namespace
```

```
You can use the `helm list` command to review installation details such as name, namespace, chart, status, app version, and revision number.
```

```
+
```

```
[source,console]
```

```
helm list -n trident
```

```
+
```

```
NAME NAMESPACE REVISION UPDATED STATUS CHART APP VERSION
```

```
trident-operator trident 1 2024-10-14 14:31:22.463122 +0300 IDT deployed trident-operator-100.2502.0 25.02.0
```

```
--
```

```
.Service account association (IRSA)
```

```
--  
  
. Add the Trident Helm repository:  
+  
[source,console]
```

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

```
. Set the values for *cloud provider* and *cloud identity*:  
+  
[source,console]
```

```
helm install trident-operator netapp-trident/trident-operator --version 100.2502.1 \  
--set cloudProvider="AWS" \  
--set cloudIdentity=""eks.amazonaws.com/role-arn:  
arn:aws:iam::<accountID>:role/<AmazonEKS_FSxN_CSI_DriverRole>" \  
--namespace trident \  
--create-namespace
```

You can use the `helm list` command to review installation details such as name, namespace, chart, status, app version, and revision number.

```
+  
[source,console]
```

helm list -n trident

```
+
```

NAME	NAMESPACE	REVISION	UPDATED	STATUS	CHART	APP VERSION
------	-----------	----------	---------	--------	-------	-------------

trident-operator	trident	1	2024-10-14 14:31:22.463122 +0300	IDT deployed	trident-operator-100.2510.0	25.10.0
------------------	---------	---	----------------------------------	--------------	-----------------------------	---------

```
////  
+  
* Beginning with 25.02 release, Trident supports automatic backend  
configuration. Trident creates the backend and storage classes seamlessly  
after Trident installation. To enable automatic backend configuration, add  
the `ontapConfigurator` parameter and specify `authType`, `fsxnID`,  
`protocols` during the installation.  
+  
[source,console]
```

```
helm install trident trident-operator-100.2502.0.tgz --set cloudProvider=$CP --set cloudIdentity=$CI
--namespace trident \
--set ontapConfigurator.enabled=true \
--set ontapConfigurator.svms[0].fsxnID="fs-0dfeaa884a68b1cab" \
--set ontapConfigurator.svms[0].protocols[0]=iscsi \
--set ontapConfigurator.svms[0].protocols[1]=nfs \
--set ontapConfigurator.svms[0].authType="awsarn"
```

NOTE: To disable automatic backend configuration, upgrade the Trident release version and set the `*ontapConfigurator*` to `*false*`.

////

--
====

[NOTE]

====

If you're planning to use iSCSI, make sure iSCSI is enabled on your client machine. If you're using AL2023 Worker node OS, you can automate the installation of the iSCSI client by adding the node prep parameter in the helm installation:

[source,console]

```
helm install trident-operator netapp-trident/trident-operator --version 100.2502.1 --namespace trident --create
--namespace --set nodePrep={iscsi}
```

====

== Install Trident via the EKS add-on

The Trident EKS add-on includes the latest security patches, bug fixes, and is validated by AWS to work with Amazon EKS. The EKS add-on enables you to consistently ensure that your Amazon EKS clusters are secure and stable and reduce the amount of work that you need to do in order to install, configure, and update add-ons.

=== Prerequisites

Ensure that you have the following before configuring the Trident add-on for AWS EKS:

- * An Amazon EKS cluster account with add-on subscription
- * AWS permissions to the AWS marketplace:
 - `"aws-marketplace:ViewSubscriptions",
 - "aws-marketplace:Subscribe",
 - "aws-marketplace:Unsubscribe`"
- * AMI type: Amazon Linux 2 (AL2_x86_64) or Amazon Linux 2 Arm(AL2_ARM_64)

```

* Node type: AMD or ARM
* An existing Amazon FSx for NetApp ONTAP file system

=== Enable the Trident add-on for AWS

[role="tabbed-block"]
=====

.Management console

--

. Open the Amazon EKS console at
https://console.aws.amazon.com/eks/home#/clusters.
. On the left navigation pane, select *Clusters*.
. Select the name of the cluster that you want to configure the NetApp
Trident CSI add-on for.
. Select *Add-ons* and then select *Get more add-ons*.
. Follow these steps to select the add-on:
.. Scroll down to the *AWS Marketplace add-ons* section and type
*"Trident"* in the search box.
.. Select the check box at the top right corner of the Trident by NetApp
box.
.. Select *Next*.
. On the *Configure selected add-ons* settings page, do the following:
+
NOTE: *Skip these steps if you are using Pod Identity association.*
+
.. Select the *Version* you would like to use.
.. If you're using IRSA authentication, make sure to set configuration
values available in the Optional configuration settings:
+
* Select the *Version* you would like to use.
* Follow the *Add-on configuration schema* and set the
*configurationValues* parameter on the *Configuration values* section to
the role-arn you created on the previous step ( Value should be in the
following format):

+
[source,JSON]

```

```
{
```

```

"cloudIdentity": "'eks.amazonaws.com/role-arn: <role ARN>'",
"cloudProvider": "AWS"

```

```
}
```

If you select Override for the Conflict resolution method, one or more of the settings for the existing add-on can be overwritten with the Amazon EKS add-on settings. If you don't enable this option and there's a conflict with your existing settings, the operation fails. You can use the resulting error message to troubleshoot the conflict. Before selecting this option, make sure that the Amazon EKS add-on doesn't manage settings that you need to self-manage.

```
+
```

```
. Choose *Next*.
```

```
. On the *Review and add* page, choose *Create*.
```

```
+
```

After the add-on installation is complete, you see your installed add-on.

```
--
```

```
.AWS CLI
```

```
--
```

```
*1. Create the `add-on.json` file*:
```

```
*For Pod Identity, use the following format*:
```

```
NOTE: Use the  
[source,json]
```

```
{  
  "clusterName": "<eks-cluster>",  
  "addonName": "netapp_trident-operator",  
  "addonVersion": "v25.6.0-eksbuild.1",  
}
```

```
*For IRSA authentication, use the following format*:  
[source,json]
```

```
{  
  "clusterName": "<eks-cluster>",  
  "addonName": "netapp_trident-operator",  
  "addonVersion": "v25.6.0-eksbuild.1",  
  "serviceAccountRoleArn": "<role ARN>",  
  "configurationValues": {  
    "cloudIdentity": "eks.amazonaws.com/role-arn: <role ARN>",  
    "cloudProvider": "AWS"  
  }  
}
```

```
}
```

NOTE: Replace ``<role ARN>`` with the ARN of the role that was created in the previous step.

```
////
```

* Beginning with 25.02 release, Trident supports automatic backend configuration. Trident creates the backend and storage classes seamlessly after Trident installation. To enable automatic backend configuration, add the `ontapConfigurator` parameter and specify `authType`, `fsxnID`, `protocols` during the installation.

```
+
```

```
[source,json]
```

```
{
"clusterName": "<eks-cluster>",
"addonName": "netapp_trident-operator",
"addonVersion": "v25.6.0-eksbuild.1",
"serviceAccountRoleArn": "arn:aws:iam::123456:role/astratrident-role",
"configurationValues": {
"cloudIdentity": "eks.amazonaws.com/role-arn:
arn:aws:iam::464262061435:role/AmazonEKS_FSXN_CSI_DriverRole",
"ontapConfigurator": {
"enabled": true,
"svms": [
{
"authType": "awsarn",
"fsxnID": "fs-0dfeaa884a68b1cab",
"protocols": [
"nfs",
"iscsi"
]
}
]
}
}
}
```

```
+
```

NOTE: To disable automatic backend configuration, upgrade the Trident release version and set the `*ontapConfigurator*` to `*false*`.

```
////
```

2. Install the Trident EKS add-on.

```
[source,console]
```

```
aws eks create-addon --cli-input-json file://add-on.json
```

```
--
```

```
.eksctl
```

```
--
```

The following example command installs the Trident EKS add-on:

```
[source,console]
```

```
eksctl create addon --name netapp_trident-operator --cluster <cluster_name> --force
```

```
--
```

```
====
```

```
=== Update the Trident EKS add-on
```

```
[role="tabbed-block"]
```

```
====
```

```
.Management console
```

```
--
```

```
. Open the Amazon EKS console
```

```
https://console.aws.amazon.com/eks/home#/clusters.
```

```
. On the left navigation pane, select *Clusters*.
```

```
. Select the name of the cluster that you want to update the NetApp  
Trident CSI add-on for.
```

```
. Select the *Add-ons* tab.
```

```
. Select *Trident by NetApp* and then select *Edit*.
```

```
. On the *Configure Trident by NetApp* page, do the following:
```

```
.. Select the *Version* you would like to use.
```

```
.. Expand the *Optional configuration settings* and modify as needed.
```

```
.. Select *Save changes*.
```

```
--
```

```
.AWS CLI
```

```
--
```

The following example updates the EKS add-on:

```
[source,console]
```

```
aws eks update-addon --cluster-name <eks_cluster_name> --addon-name netapp_trident-operator --addon
```

```
-version v25.6.0-eksbuild.1 \
--service-account-role-arn <role-ARN> --resolve-conflict preserve \
--configuration-values "{\"cloudIdentity\" : \"eks.amazonaws.com/role-arn: <role ARN>\"}"
```

```
--
.eksctl

--

* Check the current version of your FSxN Trident CSI add-on. Replace `my-
cluster` with your cluster name.
+
[source,console]
```

```
eksctl get addon --name netapp_trident-operator --cluster my-cluster
```

```
*Example output:*
```

```
NAME VERSION STATUS ISSUES IAMROLE UPDATE AVAILABLE CONFIGURATION VALUES
netapp_trident-operator v25.6.0-eksbuild.1 ACTIVE 0 {"cloudIdentity":"eks.amazonaws.com/role-arn:
arn:aws:iam::139763910815:role/AmazonEKS_FSXN_CSI_DriverRole"}
```

```
* Update the add-on to the version returned under UPDATE AVAILABLE in the
output of the previous step.
+
[source,console]
```

```
eksctl update addon --name netapp_trident-operator --version v25.6.0-eksbuild.1 --cluster my-cluster --force
```

If you remove the `--force` option and any of the Amazon EKS add-on settings conflict with your existing settings, then updating the Amazon EKS add-on fails; you receive an error message to help you resolve the conflict. Before specifying this option, make sure that the Amazon EKS add-on does not manage settings that you need to manage, because those settings are overwritten with this option.

For more information about other options for this setting, see [link:https://eksctl.io/usage/addons/\[Addons\]](https://eksctl.io/usage/addons/[Addons]).

For more information about Amazon EKS Kubernetes field management, see [link:https://docs.aws.amazon.com/eks/latest/userguide/kubernetes-field-management.html](https://docs.aws.amazon.com/eks/latest/userguide/kubernetes-field-management.html) [Kubernetes field management].

```
--
====
```

=== Uninstall/remove the Trident EKS add-on

You have two options for removing an Amazon EKS add-on:

* ***Preserve add-on software on your cluster*** - This option removes Amazon EKS management of any settings. It also removes the ability for Amazon EKS to notify you of updates and automatically update the Amazon EKS add-on after you initiate an update. However, it preserves the add-on software on your cluster. This option makes the add-on a self-managed installation, rather than an Amazon EKS add-on. With this option, there's no downtime for the add-on. Retain the `--preserve` option in the command to preserve the add-on.

* ***Remove add-on software entirely from your cluster*** - NetApp recommends that you remove the Amazon EKS add-on from your cluster only if there are no resources on your cluster that are dependent on it. Remove the `--preserve` option from the `delete` command to remove the add-on.

[NOTE]

If the add-on has an IAM account associated with it, the IAM account is not removed.

[role="tabbed-block"]

=====

.Management console

--

. Open the Amazon EKS console at
<https://console.aws.amazon.com/eks/home#/clusters>.

. In the left navigation pane, select ***Clusters***.

. Select the name of the cluster that you want to remove the NetApp Trident CSI add-on for.

. Select the ***Add-ons*** tab and then select ***Trident by NetApp***.

. Select ***Remove***.

. In the ***Remove netapp_trident-operator confirmation*** dialog, do the following:

.. If you want Amazon EKS to stop managing settings for the add-on, select ***Preserve on cluster***. Do this if you want to retain the add-on software on your cluster so that you can manage all of the settings of the add-on on your own.

.. Enter ***netapp_trident-operator***.

.. Select ***Remove***.

+

--

.AWS CLI

--

Replace `my-cluster` with the name of your cluster, and then run the following command.

[source,console]

```
aws eks delete-addon --cluster-name my-cluster --addon-name netapp_trident-operator --preserve
```

--

.eksctl

--

The following command uninstalls the Trident EKS add-on:

[source,console]

```
eksctl delete addon --cluster K8s-arm --name netapp_trident-operator
```

--

====

[[IDfe83bbdf2cb1e31fdef92d79a0256142]]

= Configure a storage class

:hardbreaks:

:icons: font

:relative_path: ./trident-use/

:imagesdir: {root_path}{relative_path}../media/

[.lead]

The <https://kubernetes.io/docs/concepts/storage/storage-classes/> [Kubernetes StorageClass object^] identifies a provisioner and instructs the provisioner how to provision volumes. This section shows you how to configure a Kubernetes StorageClass object that specifies Trident as the provisioner.

== Create a StorageClass Object

When you create a StorageClass for FSx for ONTAP, Trident will automatically create the backend configuration.

NOTE: If you'd like to manually configure the storage backend, please refer to the <<create-a-kubernetes-storageclass-without-automatic-backend-configuration>> section for how to create the Trident backend and storage class separately.

=== Specify required StorageClass parameters

The following three parameters need to be defined when creating a StorageClass:

```
[cols="2,1,1,3", options="header"]
|===
|Parameter|Required|Type|Description
|`fsxFilesystemID`|Yes|string|FSx for NetApp ONTAP filesystem ID
|`storageDriverName`|Yes|string|Trident storage driver (for example,
`ontap-nas` or `ontap-san`)
|`credentialsName`|Yes|string|Name of the Kubernetes Secret that
contains FSx for ONTAP credentials
|===
```

==== Specify optional parameters

You can pass optional backend parameters through the StorageClass. Define all optional values as strings in the StorageClass `parameters` section. For a complete list of backend parameters, see:

[xref:{relative_path}trident-fsx-storage-backend.html#backend-advanced-configuration-and-examples\[FSx for NetApp ONTAP backend configuration\]](#).

==== Example StorageClass configuration files.

The following example shows a StorageClass that triggers automatic backend configuration.

```
[role="tabbed-block"]
=====

.YAML
--

[source,yaml]
```

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: ontap-fsx-demo
annotations:
  description: "Demo StorageClass for FSx for NetApp ONTAP"
```

```
provisioner: csi.trident.netapp.io
parameters:
fsxFilesystemID: "fs-0abc123"
storageDriverName: "ontap-nas"
credentialsName: trident-fsx-credentials
allowVolumeExpansion: true
reclaimPolicy: Delete
volumeBindingMode: Immediate
```

```
--

.JSON
--

[source,json]
```

```
{
"apiVersion": "storage.k8s.io/v1",
"kind": "StorageClass",
"metadata": {
"name": "ontap-fsx-demo",
"annotations": {
"description": "Demo StorageClass for FSx for NetApp ONTAP"
}
},
"provisioner": "csi.trident.netapp.io",
"parameters": {
"fsxFilesystemID": "fs-0abc123",
"storageDriverName": "ontap-nas",
"credentialsName": "trident-fsx-credentials"
},
"allowVolumeExpansion": true,
"reclaimPolicy": "Delete",
"volumeBindingMode": "Immediate"
}
```

```
--
====

==== Create the StorageClass

Once you have created your configuration file, run the following command to
create the storage class.

[source,console]
```

```
kubectl create -f storage-class-ontapnas.yaml
```

You should now see a `*basic-csi*` storage class in both Kubernetes and

```
Trident, and Trident should have discovered the pools on the backend.
```

```
[source,console]
```

```
kubectl get sc basic-csi
```

NAME PROVISIONER AGE

```
basic-csi csi.trident.netapp.io 15h
```

After you apply the StorageClass, Trident creates the backend automatically. You can then create PersistentVolumeClaims that reference this StorageClass.

==== Verify backend configuration status

Trident records the result of backend creation in StorageClass annotations.

```
[cols="2,2", options="header"]
```

```
|===
```

```
|Annotation |Description
```

```
|`trident.netapp.io/configuratorStatus` |Configuration result (`Success`  
or `Failure`)
```

```
|`trident.netapp.io/configuratorMessage` |Detailed status or error message  
|`trident.netapp.io/configuratorName` |Name of the internal configurator  
resource
```

```
|`trident.netapp.io/managed` |Indicates the StorageClass is managed by  
Trident
```

```
|`trident.netapp.io/additionalStoragePools` |Storage pools created for  
this backend
```

```
|===
```

To verify status, run:

```
[source,bash]
```

```
kubectl get storageclass ontap-fsx-demo -o yaml
```

Confirm that `trident.netapp.io/configuratorStatus` is set to `Success`. If the value is `Failure`, review `trident.netapp.io/configuratorMessage` for the error.

==== Add additional FSxN file systems

If you need additional storage capacity while continuing to use the same

StorageClass, add additional FSxN file system IDs.

Edit the StorageClass and add the following annotation:

```
[source,yaml]
```

metadata:

annotations:

trident.netapp.io/additionalFsxNFileSystemID: ["fs-xxxxxxxxxxxxxxxxxx"]

After you apply the change, Trident updates the backend configuration and updates the StorageClass annotations.

==== Operational considerations and limitations

* Deleting a StorageClass that has the automatic backend configuration usually deletes the associated Trident backend. This can disrupt storage connectivity and break running workloads. Validate the impact before you delete a managed StorageClass.

* Automatic backend configuration is supported only for AWS FSx for NetApp ONTAP.

=== Create a Kubernetes StorageClass without automatic backend configuration

If you want to create the Trident backend and StorageClass separately then follow these steps.

==== Understand how automatic backend configuration works

Trident derives backend configuration from the StorageClass definition. When you apply the StorageClass, Trident validates the required parameters, creates the backend, and annotates the StorageClass with status.

Trident creates the VolumeSnapshotClass only once. Trident reuses the same VolumeSnapshotClass for subsequent StorageClasses.

=== Create the Trident backend

To create a Trident backend, you need to create a configuration file in either JSON or YAML format. The file needs to specify the type of storage you want (NAS or SAN), the file system, and SVM to get it from and how to authenticate with it. The following example shows how to define NAS-based storage and using an AWS secret to store the credentials to the SVM you want to use:

```
[role="tabbed-block"]
====

.YAML
--

[source,YAML]
```

```
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: backend-tbc-ontap-nas
  namespace: trident
spec:
  version: 1
  storageDriverName: ontap-nas
  backendName: tbc-ontap-nas
  svm: svm-name
  aws:
    fsxFileSystemID: fs-xxxxxxxxxx
  credentials:
    name: "arn:aws:secretsmanager:us-west-2:xxxxxxx:secret:secret-name"
    type: awsarn
```

```
--

.JSON
--

[source,JSON]
```

```
{
  "apiVersion": "trident.netapp.io/v1",
  "kind": "TridentBackendConfig",
  "metadata": {
    "name": "backend-tbc-ontap-nas"
    "namespace": "trident"
  },
  "spec": {
    "version": 1,
    "storageDriverName": "ontap-nas",
    "backendName": "tbc-ontap-nas",
    "svm": "svm-name",
    "aws": {
      "fsxFileSystemID": "fs-xxxxxxxxxx"
    },
    "managementLIF": null,
    "credentials": {
      "name": "arn:aws:secretsmanager:us-west-2:xxxxxxx:secret:secret-name",
      "type": "awsarn"
    }
  }
}
```

```
}
```

```
--  
====
```

```
==== FSx for ONTAP driver details
```

You can integrate Trident with Amazon FSx for NetApp ONTAP using the following drivers:

```
[cols="1,4", options="header"]  
|===  
|Driver Name|Description  
|`ontap-san` | Each PV provisioned is a LUN within its own Amazon FSx for  
NetApp ONTAP volume. Recommended for block storage.  
|`ontap-nas` | Each PV provisioned is a full Amazon FSx for NetApp ONTAP  
volume. Recommended for NFS and SMB.  
|`ontap-san-economy` | Each PV provisioned is a LUN with a configurable  
number of LUNs per Amazon FSx for NetApp ONTAP volume.  
|`ontap-nas-economy` | Each PV provisioned is a qtree, with a configurable  
number of qtrees per Amazon FSx for NetApp ONTAP volume.  
|`ontap-nas-flexgroup` | Each PV provisioned is a full Amazon FSx for  
NetApp ONTAP FlexGroup volume.  
|===
```

For driver details, refer to `xref:{relative_path}../trident-use/ontap-nas.html` [NAS drivers] and `xref:{relative_path}../trident-use/ontap-san.html` [SAN drivers].

```
==== Create the backend
```

After creating the configuration file run the following commands to create and validate the Trident Backend Configuration (TBC):

* Create trident backend configuration (TBC) from yaml file and run the following command:

```
+  
[source,console]
```

```
kubectl create -f backendconfig.yaml -n trident
```

```
+
```

tridentbackendconfig.trident.netapp.io/backend-tbc-ontap-nas created

```
* Validate the trident backend configuration (TBC) was created
successfully:
+
[source,console]
```

Kubecttl get tbc -n trident

```
+
```

NAME BACKEND NAME BACKEND UUID PHASE STATUS

backend-tbc-ontap-nas tbc-ontap-nas 933e0071-66ce-4324-b9ff-f96d916ac5e9 Bound Success

For more information on other configuration options, see the <<first-link,Backend-advanced-configuration-and-examples>> section below.

=== Configure a Storage Class **without** automatic backend configuration

The following are examples of Storage Class configurations for use with Trident and FSx for ONTAP.

==== Storage Class for NFS

You can use this example to setup StorageClass for volumes using NFS (Refer to Trident Attribute section below for the full list of attributes):

[source,YAML]

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
name: ontap-gold
provisioner: csi.trident.netapp.io
parameters:
backendType: "ontap-nas"
provisioningType: "thin"
snapshots: "true"
```

==== Storage Class for iSCSI

Use this example to setup StorageClass for volumes using iSCSI:

[source,YAML]

```
apiVersion: storage.k8s.io/v1
```

```
kind: StorageClass
metadata:
name: ontap-gold
provisioner: csi.trident.netapp.io
parameters:
backendType: "ontap-san"
provisioningType: "thin"
snapshots: "true"
```

```
==== Storage Class using NFSv3 and AWS Bottlerocket
To provision NFSv3 volumes on AWS Bottlerocket, add the required
`mountOptions` to the storage class:
```

```
[source,YAML]
```

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
name: ontap-gold
provisioner: csi.trident.netapp.io
parameters:
backendType: "ontap-nas"
media: "ssd"
provisioningType: "thin"
snapshots: "true"
mountOptions:
- nfsvers=3
- nolock
```

```
==== Trident StorageClass attributes
```

These parameters determine which Trident-managed storage pools should be utilized to provision volumes of a given type.

```
[cols=",,,,",options="header",]
```

```
|===
```

```
|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 ontap; thin: all ontap &
solidfire-san
```

```
|backendType |string a| ontap-nas, ontap-nas-economy, ontap-nas-flexgroup,
ontap-san, solidfire-san, azure-netapp-files, ontap-san-economy |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

|clones |bool |true, false |Pool supports cloning volumes |Volume with
clones enabled |ontap-nas, ontap-san, solidfire-san

|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 or FSx for ONTAP systems

```

Refer to `xref:{relative_path}../trident-reference/objects.html` [Kubernetes and Trident objects] for details on how storage classes interact with the `PersistentVolumeClaim` and parameters for controlling how Trident provisions volumes.`

==== Create the storage class

Once you have configured the `StorageClass`, you can create it in Kubernetes.

.Steps

. This is a Kubernetes object, so use `kubect1` to create it in Kubernetes.`

+

[source,console]

`kubect1 create -f storage-class-ontapnas.yaml`

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

+

[source,console]

`kubect1 get sc basic-csi`

== Provision SMB volumes

You can provision SMB volumes using the `ontap-nas` driver.

However, to do so you must complete these steps:

xref:{relative_path}worker-node-prep.html#prepare-to-provision-smb-volumes[Prepare to provision SMB volumes].

== [[first-link]]Backend advanced configuration and examples

See the following table for the backend configuration options:

```
[cols=3,options="header"]
```

```
|==
```

```
|Parameter |Description |Example
```

```
|`version`
```

```
|
```

```
|Always 1
```

```
|`storageDriverName`
```

```
| Name of the storage driver
```

```
|`ontap-nas`, `ontap-nas-economy`, `ontap-nas-flexgroup`, `ontap-san`,  
`ontap-san-economy`
```

```
|`backendName`
```

```
|Custom name or the storage backend
```

```
| Driver name + "_" + dataLIF
```

```
|`managementLIF`
```

```
|IP address of a cluster or SVM management LIF
```

A fully-qualified domain name (FQDN) can be specified.

Can be set to use IPv6 addresses if Trident was installed using the IPv6 flag. IPv6 addresses must be defined in square brackets, such as [28e8:d9fb:a825:b7bf:69a8:d02f:9e7b:3555].

If you provide the `fsxFilesystemID` under the `aws` field, you need not to provide the `managementLIF` because Trident retrieves the SVM `managementLIF` information from AWS. So, you must provide credentials for a user under the SVM (For example: vsadmin) and the user must have the `vsadmin` role.

```
|"10.0.0.1", "[2001:1234:abcd::fefe]"
```

```
|`dataLIF`
|IP address of protocol LIF.

*ONTAP NAS drivers*: NetApp recommends specifying dataLIF. If not
provided, Trident fetches dataLIFs from the SVM. You can specify a fully-
qualified domain name (FQDN) to be used for the NFS mount operations,
allowing you to create a round-robin DNS to load-balance across multiple
dataLIFs. Can be changed after initial setting.

*ONTAP SAN drivers*: Do not specify for iSCSI. Trident uses ONTAP
Selective LUN Map to discover the iSCSI LIFs needed to establish a multi
path session. A warning is generated if dataLIF is explicitly defined.

Can be set to use IPv6 addresses if Trident was installed using the IPv6
flag. IPv6 addresses must be defined in square brackets, such as
[28e8:d9fb:a825:b7bf:69a8:d02f:9e7b:3555].

|

|`autoExportPolicy`
|Enable automatic export policy creation and updating [Boolean].

Using the `autoExportPolicy` and `autoExportCIDRs` options, Trident can
manage export policies automatically.
|`false`

|`autoExportCIDRs`
|List of CIDRs to filter Kubernetes' node IPs against when
`autoExportPolicy` is enabled.

Using the `autoExportPolicy` and `autoExportCIDRs` options, Trident can
manage export policies automatically.
|["0.0.0.0/0", "::/0"]

|`labels`
|Set of arbitrary JSON-formatted labels to apply on volumes
|""

|`clientCertificate`
|Base64-encoded value of client certificate. Used for certificate-based
auth
|""

|`clientPrivateKey`
|Base64-encoded value of client private key. Used for certificate-based
auth
```

```

|""

|`trustedCACertificate`
|Base64-encoded value of trusted CA certificate. Optional. Used for
certificate-based authentication.
|""

|`username`
|Username to connect to the cluster or SVM. Used for credential-based
authentication. For example, vsadmin.
|

|`password`
|Password to connect to the cluster or SVM. Used for credential-based
authentication.
|

|`svm`
|Storage virtual machine to use
|Derived if an SVM managementLIF is specified.

|`storagePrefix`
|Prefix used when provisioning new volumes in the SVM.

Cannot be modified after creation. To update this parameter, you will need
to create a new backend.
|`trident`

|`limitAggregateUsage` |*Do not specify for Amazon FSx for NetApp ONTAP.*

The provided `fsxadmin` and `vsadmin` do not contain the permissions
required to retrieve aggregate usage and limit it using Trident.

|Do not use.

|`limitVolumeSize`
|Fail provisioning if requested volume size is above this value.

Also restricts the maximum size of the volumes it manages for qtrees and
LUNs, and the `qtreesPerFlexvol` option allows customizing the maximum
number of qtrees per FlexVol volume

|"" (not enforced by default)

|`lunsPerFlexvol`
|Maximum LUNs per Flexvol volume, must be in range [50, 200].

```

SAN only.

| ``100``

| `debugTraceFlags`

| Debug flags to use when troubleshooting. Example, {"api":false,
"method":true}

Do not use `debugTraceFlags` unless you are troubleshooting and require a detailed log dump.

| null

| `nfsMountOptions`

| Comma-separated list of NFS mount options.

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 configuration file.

If no mount options are specified in the storage class or the configuration file, Trident will not set any mount options on an associated persistent volume.

| ""

| `nasType` | Configure NFS or SMB volumes creation.

Options are `nfs`, `smb`, or null.

Must set to `smb` for SMB volumes. Setting to null defaults to NFS volumes.

| `nfs`

| `qtreesPerFlexvol`

| Maximum Qtrees per FlexVol volume, must be in range [50, 300]

| ``200``

| `smbShare`

| You can specify one of the following: the name of an SMB share created using the Microsoft Management Console or ONTAP CLI or a name to allow Trident to create the SMB share.

This parameter is required for Amazon FSx for ONTAP backends.

| `smb-share`

| `useREST`

```
|Boolean parameter to use ONTAP REST APIs.
```

When set to `true`, Trident will use ONTAP REST APIs to communicate with the backend.

This feature requires ONTAP 9.11.1 and later. In addition, the ONTAP login role used must have access to the `ontap` application. This is satisfied by the pre-defined `vsadmin` and `cluster-admin` roles.

```
|`false`
```

```
|`aws`| You can specify the following in the configuration file for AWS FSx for ONTAP:
```

- `fsxFileSystemID`: Specify the ID of the AWS FSx file system.
- `apiRegion`: AWS API region name.
- `apikey`: AWS API key.
- `secretKey`: AWS secret key.

```
| ``  
``
```

```
````  
````  
````
```

```
|`credentials`|Specify the FSx SVM credentials to store in AWS Secrets Manager.
```

- `name`: Amazon Resource Name (ARN) of the secret, which contains the credentials of SVM.
- `type`: Set to `awsarn`.

Refer to

[link:https://docs.aws.amazon.com/secretsmanager/latest/userguide/create\\_secret.html](https://docs.aws.amazon.com/secretsmanager/latest/userguide/create_secret.html)[Create an AWS Secrets Manager secret^] for more information.

```
|===
```

== Backend configuration options for provisioning volumes

You can control default provisioning using these options in the `defaults` section of the configuration. For an example, see the configuration examples below.

```
[cols=3,options="header"]
```

```
|===
```

```
|Parameter |Description |Default
```

```
|`spaceAllocation` |Space-allocation for LUNs |`true`
```

```
|`spaceReserve` |Space reservation mode; "none" (thin) or "volume" (thick)
|`none`
```

```
|`snapshotPolicy` |Snapshot policy to use |`none`
```

```
|`qosPolicy` |QoS policy group to assign for volumes created. Choose one
of qosPolicy or adaptiveQosPolicy per storage pool or backend.
```

Using QoS policy groups with Trident requires ONTAP 9.8 or later.

You should use a non-shared QoS policy group and ensuring the policy group is applied to each constituent individually. A shared QoS policy group enforces the ceiling for the total throughput of all workloads.

```
|""
```

```
|`adaptiveQosPolicy` |Adaptive QoS policy group to assign for volumes
created. Choose one of qosPolicy or adaptiveQosPolicy per storage pool or
backend.
```

Not supported by ontap-nas-economy.

```
|""
```

```
|`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 (NVE) on the new volume;
defaults to `false`. NVE must be licensed and enabled on the cluster to
use this option.
```

If NAE is enabled on the backend, any volume provisioned in Trident will be NAE enabled.

For more information, refer to: `xref:{relative_path}../trident-reco/security-reco.html`[How Trident works with NVE and NAE]. |`false`

```
|`luksEncryption`
```

```
| Enable LUKS encryption. Refer to xref:{relative_path}../trident-
reco/security-luks.html[Use Linux Unified Key Setup (LUKS)].
```

SAN only.

```
| ""
```

```
|`tieringPolicy`
```

```
|Tiering policy to use `none`
|

|`unixPermissions` |Mode for new volumes.

Leave empty for SMB volumes.
|""

|`securityStyle` |Security style for new volumes.

NFS supports `mixed` and `unix` security styles.

SMB supports `mixed` and `ntfs` security styles.

|NFS default is `unix`.

SMB default is `ntfs`.

|===

[[ID4eeff0cce9e87af0419e6cf58a587a4c]]
= Configure a PVC
:hardbreaks:
:icons: font
:relative_path: ./trident-use/
:imagesdir: {root_path}{relative_path}../media/

[.lead]
This sections includes instructions on how to create a
PersistentVolumeClaim (PVC) that uses the configured Kubernetes
StorageClass to request a PV.
If successful, you can then mount the PV to a pod.

== Create the PVC
A https://kubernetes.io/docs/concepts/storage/persistent-
volumes\[_PersistentVolumeClaim_\] (PVC) is a request for access to the
PersistentVolume on the cluster. The PVC can be configured to request
storage of a certain size or access mode. Using the associated
StorageClass, the cluster administrator can control more than
PersistentVolume size and access mode--such as performance or service
level.

After you create the Trident backend and StorageClass you can create a
PVC. After the PVC is created, you can mount the volume in a pod.
```

```
=== Sample manifests
```

The following examples show basic PVC configuration options.

```
.PVC with RWX access
```

This example shows a basic PVC with RWX access that is associated with a StorageClass named `basic-csi`.

```
[source,YAML]
```

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
name: pvc-storage
spec:
accessModes:
- ReadWriteMany
resources:
requests:
storage: 1Gi
storageClassName: ontap-gold
```

```
.PVC using iSCSI example
```

This example shows a basic PVC for iSCSI with RWO access that is associated with a StorageClass named `protection-gold`.

```
[source,YAML]
```

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
name: pvc-san
spec:
accessModes:
- ReadWriteOnce
resources:
requests:
storage: 1Gi
storageClassName: protection-gold
```

```
=== Create PVC
```

```
.Steps
```

```
. Create the PVC.
```

```
+
```

```
[source,console]
```

```
kubectrl create -f pvc.yaml
```

```
. Verify the PVC status.
+
[source,console]
```

kubectl get pvc

```
NAME STATUS VOLUME CAPACITY ACCESS MODES STORAGECLASS AGE
pvc-storage Bound pv-name 2Gi RWO 5m
```

Refer to [xref:{relative\\_path}../trident-reference/objects.html](#) [Kubernetes and Trident objects] for details on how storage classes interact with the `PersistentVolumeClaim` and parameters for controlling how Trident provisions volumes.

```
[[ID1f24b3c8114100df6f6735e9f650ed3e]]
= Deploy an application
:hardbreaks:
:icons: font
:relative_path: ../trident-use/
:imagesdir: {root_path}{relative_path}../media/
```

[.lead]

When the storage class and PVC are created, you can mount the PV to a pod. This section lists the example command and configuration to attach the PV to a pod.

= Deploy a sample application

.Steps

```
. Mount the volume in a pod.
+
[source,console]
```

kubectl create -f pv-pod.yaml

These examples show basic configurations to attach the PVC to a pod:

\*Basic configuration\*:

+

```
[source,yaml]
```

```
kind: Pod
apiVersion: v1
metadata:
 name: pv-pod
spec:
 volumes:
 - name: pv-storage
 persistentVolumeClaim:
 claimName: basic
 containers:
 - name: pv-container
 image: nginx
 ports:
 - containerPort: 80
 name: "http-server"
 volumeMounts:
 - mountPath: "/my/mount/path"
 name: pv-storage
```

NOTE: You can monitor the progress using `kubectl get pod --watch`.

. Verify that the volume is mounted on `/my/mount/path`.

+

```
[source,console]
```

```
kubectl exec -it pv-pod -- df -h /my/mount/path
```

```
Filesystem Size Used Avail Use% Mounted on
192.168.188.78:/trident_pvc_ae45ed05_3ace_4e7c_9080_d2a83ae03d06 1.1G 320K 1.0G 1%
/my/mount/path
```

You can now delete the Pod. The Pod application will no longer exist, but the volume will remain.

```
[source,console]
```

```
kubectl delete pod pv-pod
```

```
[[IDe0b765829fe12f9eb530fedf547fdff9]]
= Configure the Trident EKS add-on on an EKS cluster
:hardbreaks:
:icons: font
```

```
:relative_path: ./trident-use/
:imagesdir: {root_path}{relative_path}../media/
```

```
[.lead]
```

NetApp Trident streamlines Amazon FSx for NetApp ONTAP storage management in Kubernetes to enable your developers and administrators focus on application deployment. The NetApp Trident EKS add-on includes the latest security patches, bug fixes, and is validated by AWS to work with Amazon EKS. The EKS add-on enables you to consistently ensure that your Amazon EKS clusters are secure and stable and reduce the amount of work that you need to do in order to install, configure, and update add-ons.

#### == Prerequisites

Ensure that you have the following before configuring the Trident add-on for AWS EKS:

- \* An Amazon EKS cluster account with permissions to work with add-ons. Refer to link:<https://docs.aws.amazon.com/eks/latest/userguide/eks-add-ons.html>[Amazon EKS add-ons^].

- \* AWS permissions to the AWS marketplace:

```
`"aws-marketplace:ViewSubscriptions",
 "aws-marketplace:Subscribe",
 "aws-marketplace:Unsubscribe`
```

- \* AMI type: Amazon Linux 2 (AL2\_x86\_64) or Amazon Linux 2 Arm(AL2\_ARM\_64)

- \* Node type: AMD or ARM

- \* An existing Amazon FSx for NetApp ONTAP file system

#### == Steps

- . Make sure to create IAM role and AWS secret to enable EKS pods to access AWS resources. For instructions, see [xref:{relative\\_path}../trident-use/trident-fsx-iam-role.html](#)[Create an IAM role and AWS Secret^].

- . On your EKS Kubernetes cluster, navigate to the *Add-ons* tab.

+

image:.../media/aws-eks-01.png[EKS cluster Add-ons tab]

+

- . Go to *AWS Marketplace add-ons* and choose the *\_storage\_* category.

+

image:.../media/aws-eks-02.png[AWS Marketplace add-ons storage category]

+

- . Locate *NetApp Trident* and select the checkbox for the Trident add-on, and click *Next*.

- . Choose the desired version of the add-on.

+

image:.../media/aws-eks-03.png[Trident add-on version selection]

+

```

. Configure the required add-on settings.
+
image:.../media/aws-eks-04.png[Trident add-on configuration settings]
+
. If you are using IRSA (IAM roles for service account), refer to the
additional configuration steps link:https://docs.netapp.com/us-en/trident/trident-use/trident-fsx-install-trident.html#enable-the-trident-add-on-for-aws[here].
+
/////
* Beginning with 25.02 release, Trident supports automatic backend
configuration. Trident creates the backend and storage classes seamlessly
after Trident installation. To enable automatic backend configuration, add
the `ontapConfigurator` parameter and specify `authType`, `fsxnID`, and
`protocols` in the Add-on configuration schema along with `cloudIdentity`
during the installation.
+
[source,json]

```

```

{
"ontapConfigurator": {
"enabled": true,
"svms": [
{
"authType": "awsarn",
"fsxnID": "fs-0dfeaa884a68b1cab",
"protocols": [
"nfs",
"iscsi"
]
}
]
}
}
}

```

```

+
NOTE: To disable automatic backend configuration, upgrade the Trident
release version and set the *ontapConfigurator* to *false*.
/////
. Select *Create*.
. Verify that the status of the add-on is _Active_.
+
image:.../media/aws-eks-05.png[Trident add-on active status]
+
. Run the following command to verify that Trident is properly installed
on the cluster:

```

+

kubectl get pods -n trident

. Continue the setup and configure the storage backend. For information, see [xref:{relative\\_path}trident-fsx-storage-backend.html](#)[Configure the Storage Backend].

== Install/uninstall the Trident EKS add-on using CLI

.Install the NetApp Trident EKS add-on using CLI:

The following example command installs the Trident EKS add-on:

```
`eksctl create addon --cluster clusterName --name netapp_trident-operator --version v25.6.0-eksbuild.1` (with a dedicated version)
```

The following example command installs the Trident EKS add-on version 25.6.1:

```
`eksctl create addon --cluster clusterName --name netapp_trident-operator --version v25.6.1-eksbuild.1` (with a dedicated version)
```

The following example command installs the Trident EKS add-on version 25.6.2:

```
`eksctl create addon --cluster clusterName --name netapp_trident-operator --version v25.6.2-eksbuild.1` (with a dedicated version)
```

.Uninstall the NetApp Trident EKS add-on using CLI:

The following command uninstalls the Trident EKS add-on:

eksctl delete addon --cluster K8s-arm --name netapp\_trident-operator

```
:leveloffset: -1
```

```
[[ID9ab0f76f0a264cc932440a9504778b90]]
```

```
= Create backends with kubectl
```

```
:hardbreaks:
```

```
:icons: font
```

```
:relative_path: ./trident-use/
```

```
:imagesdir: {root_path}{relative_path}../media/
```

```
[.lead]
```

A 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. After Trident is installed, the next step is to create a backend. The `TridentBackendConfig` Custom`

Resource Definition (CRD) enables you to create and manage Trident backends directly through the Kubernetes interface. You can do this by using `kubectl` or the equivalent CLI tool for your Kubernetes distribution.

== `TridentBackendConfig`

`TridentBackendConfig` (`tbc`, `tbconfig`, `tbackendconfig`) is a frontend, namespaced CRD that enables you to manage Trident backends using `kubectl`. Kubernetes and storage admins can now create and manage backends directly through the Kubernetes CLI without requiring a dedicated command-line utility (`tridentctl`).

Upon the creation of a `TridentBackendConfig` object, the following happens:

- \* A backend is created automatically by Trident based on the configuration you provide. This is represented internally as a `TridentBackend` (`tbe`, `tridentbackend`) CR.
- \* The `TridentBackendConfig` is uniquely bound to a `TridentBackend` that was created by Trident.

Each `TridentBackendConfig` maintains a one-to-one mapping with a `TridentBackend`. The former is the interface provided to the user to design and configure backends; the latter is how Trident represents the actual backend object.

WARNING: `TridentBackend` CRs are created automatically by Trident. You *should not* modify them. If you want to make updates to backends, do this by modifying the `TridentBackendConfig` object.

See the following example for the format of the `TridentBackendConfig` CR:  
[source,yaml]

```
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
 name: backend-tbc-ontap-san
spec:
 version: 1
 backendName: ontap-san-backend
 storageDriverName: ontap-san
 managementLIF: 10.0.0.1
 dataLIF: 10.0.0.2
 svm: trident_svm
 credentials:
 name: backend-tbc-ontap-san-secret
```

You can also take a look at the examples in the [https://github.com/NetApp/trident/tree/stable/v21.07/trident-installer/sample-input/backends-samples\[trident-installer^\]](https://github.com/NetApp/trident/tree/stable/v21.07/trident-installer/sample-input/backends-samples[trident-installer^]) directory for sample configurations for the desired storage platform/service.

The ``spec`` takes backend-specific configuration parameters. In this example, the backend uses the ``ontap-san`` storage driver and uses the configuration parameters that are tabulated here. For the list of configuration options for your desired storage driver, refer to the `xref:{relative_path}backends.html[backend configuration information for your storage driver^]`.

The ``spec`` section also includes ``credentials`` and ``deletionPolicy`` fields, which are newly introduced in the ``TridentBackendConfig`` CR:

- \* ``credentials``: This parameter is a required field and contains the credentials used to authenticate with the storage system/service. This is set to a user-created Kubernetes Secret. The credentials cannot be passed in plain text and will result in an error.
- \* ``deletionPolicy``: This field defines what should happen when the ``TridentBackendConfig`` is deleted. It can take one of two possible values:
  - \*\* ``delete``: This results in the deletion of both ``TridentBackendConfig`` CR and the associated backend. This is the default value.
  - \*\* ``retain``: When a ``TridentBackendConfig`` CR is deleted, the backend definition will still be present and can be managed with ``tridentctl``. Setting the deletion policy to ``retain`` lets users downgrade to an earlier release (pre-21.04) and retain the created backends. The value for this field can be updated after a ``TridentBackendConfig`` is created.

NOTE: The name of a backend is set using ``spec.backendName``. If unspecified, the name of the backend is set to the name of the ``TridentBackendConfig`` object (metadata.name). It is recommended to explicitly set backend names using ``spec.backendName``.

TIP: Backends that were created with ``tridentctl`` do not have an associated ``TridentBackendConfig`` object. You can choose to manage such backends with ``kubectl`` by creating a ``TridentBackendConfig`` CR. Care must be taken to specify identical config parameters (such as ``spec.backendName``, ``spec.storagePrefix``, ``spec.storageDriverName``, and so on). Trident will automatically bind the newly-created ``TridentBackendConfig`` with the pre-existing backend.

== Steps overview

To create a new backend by using ``kubectl``, you should do the following:

. Create a <https://kubernetes.io/docs/concepts/configuration/secret/> [Kubernetes Secret<sup>^</sup>]. The secret contains the credentials Trident needs to communicate with the storage cluster/service.

. Create a `TridentBackendConfig` object. This contains specifics about the storage cluster/service and references the secret created in the previous step.

After you create a backend, you can observe its status by using `kubectl get tbc <tbc-name> -n <trident-namespace>` and gather additional details.

#### == Step 1: Create a Kubernetes Secret

Create a Secret that contains the access credentials for the backend. This is unique to each storage service/platform. Here's an example:

```
kubectl -n trident create -f backend-tbc-ontap-san-secret.yaml
```

```
[source,yaml]
```

```
apiVersion: v1
kind: Secret
metadata:
name: backend-tbc-ontap-san-secret
type: Opaque
stringData:
username: cluster-admin
password: password
```

This table summarizes the fields that must be included in the Secret for each storage platform:

```
[cols=3,options="header"]
|===
|Storage platform Secret Fields description
|Secret
|Fields description

|Azure NetApp Files
a|clientID
a|The client ID from an app registration

|Element (NetApp HCI/SolidFire)
a|Endpoint
a|MVIP for the SolidFire cluster with tenant credentials
```

```

|ONTAP
a|username
a|Username to connect to the cluster/SVM. Used for credential-based
authentication

|ONTAP
a|password
a|Password to connect to the cluster/SVM. Used for credential-based
authentication

|ONTAP
a|clientPrivateKey
a|Base64-encoded value of client private key. Used for certificate-based
authentication

|ONTAP
a|chapUsername
a|Inbound username. Required if useCHAP=true. For `ontap-san` and `ontap-
san-economy`

|ONTAP
a|chapInitiatorSecret
a|CHAP initiator secret. Required if useCHAP=true. For `ontap-san` and
`ontap-san-economy`

|ONTAP
a|chapTargetUsername
a|Target username. Required if useCHAP=true. For `ontap-san` and `ontap-
san-economy`

|ONTAP
a|chapTargetInitiatorSecret
a|CHAP target initiator secret. Required if useCHAP=true. For `ontap-san`
and `ontap-san-economy`

|===

```

The Secret created in this step will be referenced in the `spec.credentials` field of the `TridentBackendConfig` object that is created in the next step.

== Step 2: Create the `TridentBackendConfig` CR

You are now ready to create your `TridentBackendConfig` CR. In this example, a backend that uses the `ontap-san` driver is created by using the `TridentBackendConfig` object shown below:

```
kubectl -n trident create -f backend-tbc-ontap-san.yaml
```

```
[source,yaml]
```

```
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
 name: backend-tbc-ontap-san
spec:
 version: 1
 backendName: ontap-san-backend
 storageDriverName: ontap-san
 managementLIF: 10.0.0.1
 dataLIF: 10.0.0.2
 svm: trident_svm
 credentials:
 name: backend-tbc-ontap-san-secret
```

== Step 3: Verify the status of the `TridentBackendConfig` CR

Now that you created the `TridentBackendConfig` CR, you can verify the status. See the following example:

```
kubectl -n trident get tbc backend-tbc-ontap-san
NAME BACKEND NAME BACKEND UUID PHASE STATUS
backend-tbc-ontap-san ontap-san-backend 8d24fce7-6f60-4d4a-8ef6-bab2699e6ab8 Bound Success
```

A backend was successfully created and bound to the `TridentBackendConfig` CR.

Phase can take one of the following values:

- \* `Bound`: The `TridentBackendConfig` CR is associated with a backend, and that backend contains `configRef` set to the `TridentBackendConfig` CR's uid.
- \* `Unbound`: Represented using `""`. The `TridentBackendConfig` object is not bound to a backend. All newly created `TridentBackendConfig` CRs are in this phase by default. After the phase changes, it cannot revert to Unbound again.
- \* `Deleting`: The `TridentBackendConfig` CR's `deletionPolicy` was set to delete. When the `TridentBackendConfig` CR is deleted, it transitions to the Deleting state.
- \*\* If no persistent volume claims (PVCs) exist on the backend, deleting the `TridentBackendConfig` will result in Trident deleting the backend as well as the `TridentBackendConfig` CR.
- \*\* If one or more PVCs are present on the backend, it goes to a deleting

state. The `TridentBackendConfig` CR subsequently also enters deleting phase. The backend and `TridentBackendConfig` are deleted only after all PVCs are deleted.

\* `Lost`: The backend associated with the `TridentBackendConfig` CR was accidentally or deliberately deleted and the `TridentBackendConfig` CR still has a reference to the deleted backend. The `TridentBackendConfig` CR can still be deleted irrespective of the `deletionPolicy` value.

\* `Unknown`: Trident is unable to determine the state or existence of the backend associated with the `TridentBackendConfig` CR. For example, if the API server is not responding or if the `tridentbackends.trident.netapp.io` CRD is missing. This might require intervention.

At this stage, a backend is successfully created! There are several operations that can additionally be handled, such as `xref:{relative_path}backend_ops_kubectl.html[backend updates and backend deletions^]`.

== (Optional) Step 4: Get more details

You can run the following command to get more information about your backend:

```
kubectl -n trident get tbc backend-tbc-ontap-san -o wide
```

```
NAME BACKEND NAME BACKEND UUID PHASE STATUS STORAGE DRIVER DELETION POLICY
backend-tbc-ontap-san ontap-san-backend 8d24fce7-6f60-4d4a-8ef6-bab2699e6ab8 Bound Success ontap-san delete
```

In addition, you can also obtain a YAML/JSON dump of `TridentBackendConfig`.

```
kubectl -n trident get tbc backend-tbc-ontap-san -o yaml
```

```
[source,yaml]
```

```
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
 creationTimestamp: 2021-04-21T20:45:11Z
finalizers:
- trident.netapp.io
generation: 1
name: backend-tbc-ontap-san
```

```
namespace: trident
resourceVersion: "947143"
uid: 35b9d777-109f-43d5-8077-c74a4559d09c
spec:
 backendName: ontap-san-backend
 credentials:
 name: backend-tbc-ontap-san-secret
 managementLIF: 10.0.0.1
 dataLIF: 10.0.0.2
 storageDriverName: ontap-san
 svm: trident_svm
 version: 1
 status:
 backendInfo:
 backendName: ontap-san-backend
 backendUUID: 8d24fce7-6f60-4d4a-8ef6-bab2699e6ab8
 deletionPolicy: delete
 lastOperationStatus: Success
 message: Backend 'ontap-san-backend' created
 phase: Bound
```

``backendInfo`` contains the ``backendName`` and the ``backendUUID`` of the backend that got created in response to the ``TridentBackendConfig`` CR. The ``lastOperationStatus`` field represents the status of the last operation of the ``TridentBackendConfig`` CR, which can be user-triggered (for example, user changed something in ``spec``) or triggered by Trident (for example, during Trident restarts). It can either be Success or Failed. ``phase`` represents the status of the relation between the ``TridentBackendConfig`` CR and the backend. In the example above, ``phase`` has the value Bound, which means that the ``TridentBackendConfig`` CR is associated with the backend.

You can run the ``kubectl -n trident describe tbc <tbc-cr-name>`` command to get details of the event logs.

WARNING: You cannot update or delete a backend which contains an associated ``TridentBackendConfig`` object using ``tridentctl``. To understand the steps involved in switching between ``tridentctl`` and ``TridentBackendConfig``, [xref:{relative\\_path}backend\\_options.html\[see here^\]](#).

= Manage backends

:leveloffset: +1

[[IDe9ae4e209cab058314439133fb08ffa7]]

```
= Manage Trident backends with kubectl
:hardbreaks:
:icons: font
:relative_path: ./trident-use/
:imagesdir: {root_path}{relative_path}../media/
```

[.lead]

Learn about how to perform backend management operations by using  
`kubectl`.

== Delete a backend

By deleting a `TridentBackendConfig`, you instruct Trident to delete/retain backends (based on `deletionPolicy`). To delete a backend, ensure that `deletionPolicy` is set to delete. To delete just the `TridentBackendConfig`, ensure that `deletionPolicy` is set to retain. This ensures the backend is still present and can be managed by using `tridentctl`.

Run the following command:

```
kubectl delete tbc <tbc-name> -n trident
```

Trident does not delete the Kubernetes Secrets that were in use by `TridentBackendConfig`. The Kubernetes user is responsible for cleaning up secrets. Care must be taken when deleting secrets. You should delete secrets only if they are not in use by the backends.

== View the existing backends

Run the following command:

```
kubectl get tbc -n trident
```

You can also run `tridentctl get backend -n trident` or `tridentctl get backend -o yaml -n trident` to obtain a list of all backends that exist. This list will also include backends that were created with `tridentctl`.

== Update a backend

There can be multiple reasons to update a backend:

- \* Credentials to the storage system have changed. To update credentials, the Kubernetes Secret that is used in the `TridentBackendConfig` object must be updated. Trident will automatically update the backend with the

latest credentials provided. Run the following command to update the Kubernetes Secret:

+

```
kubectl apply -f <updated-secret-file.yaml> -n trident
```

\* Parameters (such as the name of the ONTAP SVM being used) need to be updated.

\*\* You can update `TridentBackendConfig` objects directly through Kubernetes using the following command:

+

```
kubectl apply -f <updated-backend-file.yaml>
```

\*\* Alternatively, you can make changes to the existing `TridentBackendConfig` CR using the following command:

+

```
kubectl edit tbc <tbc-name> -n trident
```

[NOTE]

====

\* If a backend update fails, the backend continues to remain in its last known configuration. You can view the logs to determine the cause by running `kubectl get tbc <tbc-name> -o yaml -n trident` or `kubectl describe tbc <tbc-name> -n trident`.

\* After you identify and correct the problem with the configuration file, you can re-run the update command.

====

```
[[IDc68a5701a9cd87410dca286cdf93743d]]
```

```
= Manage Trident backends with tridentctl
```

```
:hardbreaks:
```

```
:icons: font
```

```
:relative_path: ./trident-use/
```

```
:imagesdir: {root_path}{relative_path}../media/
```

```
[.lead]
```

```
Learn about how to perform backend management operations by using
`tridentctl`.
```

```
== Create a backend
```

After you create a `xref:{relative_path}backends.html[backend configuration file^]`, run the following command:

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

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

```
tridentctl logs -n trident
```

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

== Delete a backend

To delete a backend from Trident, do the following:

- . Retrieve the backend name:  
+

```
tridentctl get backend -n trident
```

- . Delete the backend:  
+

```
tridentctl delete backend <backend-name> -n trident
```

NOTE: If Trident has provisioned volumes and snapshots from this backend that still exist, deleting the backend prevents new volumes from being provisioned by it. The backend will continue to exist in a "Deleting" state.

== View the existing backends

To view the backends that Trident knows about, do the following:

- \* To get a summary, run the following command:  
+

```
tridentctl get backend -n trident
```

```
* To get all the details, run the following command:
+
```

```
tridentctl get backend -o json -n trident
```

```
== Update a backend
```

After you create a new backend configuration file, run the following command:

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

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 the following command:

```
tridentctl logs -n trident
```

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

```
== Identify the storage classes that use a backend
```

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

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

This also applies for backends that were created by using `TridentBackendConfig`.

```
[[ID82917feee67c733dd086172de0defcf0]]
= Choose a Trident backend management method
:hardbreaks:
:icons: font
:relative_path: ./trident-use/
:imagesdir: {root_path}{relative_path}../media/
```

[.lead]

Learn about the different ways of managing backends in Trident.

#### == Options for managing backends

With the introduction of `TridentBackendConfig`, administrators now have two unique ways of managing backends. This poses the following questions:

- \* Can backends created using `tridentctl` be managed with `TridentBackendConfig`?
- \* Can backends created using `TridentBackendConfig` be managed using `tridentctl`?

#### == Manage `tridentctl` backends using `TridentBackendConfig`

This section covers the steps required to manage backends that were created using `tridentctl` directly through the Kubernetes interface by creating `TridentBackendConfig` objects.

This will apply to the following scenarios:

- \* Pre-existing backends, that don't have a `TridentBackendConfig` because they were created with `tridentctl`.
- \* New backends that were created with `tridentctl`, while other `TridentBackendConfig` objects exist.

In both scenarios, backends will continue to be present, with Trident scheduling volumes and operating on them. Administrators have one of two choices here:

- \* Continue using `tridentctl` to manage backends that were created using it.
- \* Bind backends created using `tridentctl` to a new `TridentBackendConfig` object. Doing so would mean the backends will be managed using `kubectrl` and not `tridentctl`.

To manage a pre-existing backend using `kubectrl`, you will need to create a `TridentBackendConfig` that binds to the existing backend. Here is an overview of how that works:

- . Create a Kubernetes Secret. The secret contains the credentials Trident needs to communicate with the storage cluster/service.
- . Create a `TridentBackendConfig` object. This contains specifics about the storage cluster/service and references the secret created in the previous step. Care must be taken to specify identical config parameters (such as `spec.backendName`, `spec.storagePrefix`, `spec.storageDriverName`, and so on). `spec.backendName` must be set to

the name of the existing backend.

=== Step 0: Identify the backend

To create a `TridentBackendConfig` that binds to an existing backend, you will need to obtain the backend configuration. In this example, let us assume a backend was created using the following JSON definition:

```
tridentctl get backend ontap-nas-backend -n trident
```

```

| NAME | STORAGE DRIVER | UUID | STATE | VOLUMES |

| ontap-nas-backend | ontap-nas | 52f2eb10-e4c6-4160-99fc-96b3be5ab5d7 | online | 25 |

```

```
cat ontap-nas-backend.json
```

```
[source,json]
```

```
{
 "version": 1,
 "storageDriverName": "ontap-nas",
 "managementLIF": "10.10.10.1",
 "dataLIF": "10.10.10.2",
 "backendName": "ontap-nas-backend",
 "svm": "trident_svm",
 "username": "cluster-admin",
 "password": "admin-password",
 "defaults": {
 "spaceReserve": "none",
 "encryption": "false"
 },
 "labels": {
 "store": "nas_store"
 },
 "region": "us_east_1",
 "storage": [
 {
 "labels": {
 "app": "msoffice",
 "cost": "100"
 },
 "zone": "us_east_1a",
 "defaults": {
 "spaceReserve": "volume",
 "encryption": "true",
 "unixPermissions": "0755"
 }
 }
]
}
```

```

},
{
 "labels": {
 "app": "mysqldb",
 "cost": "25"
 },
 "zone": "us_east_1d",
 "defaults": {
 "spaceReserve": "volume",
 "encryption": "false",
 "unixPermissions": "0775"
 }
}
]
}

```

=== Step 1: Create a Kubernetes Secret

Create a Secret that contains the credentials for the backend, as shown in this example:

cat tbc-ontap-nas-backend-secret.yaml

```
[source,yaml]
```

```

apiVersion: v1
kind: Secret
metadata:
 name: ontap-nas-backend-secret
type: Opaque
stringData:
 username: cluster-admin
 password: admin-password

```

```

kubectl create -f tbc-ontap-nas-backend-secret.yaml -n trident
secret/backend-tbc-ontap-san-secret created

```

=== Step 2: Create a `TridentBackendConfig` CR

The next step is to create a `TridentBackendConfig` CR that will automatically bind to the pre-existing `ontap-nas-backend` (as in this example). Ensure the following requirements are met:

- \* The same backend name is defined in `spec.backendName`.
- \* Configuration parameters are identical to the original backend.

```
* Virtual pools (if present) must retain the same order as in the original
backend.
* Credentials are provided through a Kubernetes Secret and not in plain
text.
```

In this case, the `TridentBackendConfig` will look like this:

```
cat backend-tbc-ontap-nas.yaml
```

```
[source,yaml]
```

```
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
 name: tbc-ontap-nas-backend
spec:
 version: 1
 storageDriverName: ontap-nas
 managementLIF: 10.10.10.1
 dataLIF: 10.10.10.2
 backendName: ontap-nas-backend
 svm: trident_svm
 credentials:
 name: mysecret
 defaults:
 spaceReserve: none
 encryption: 'false'
 labels:
 store: nas_store
 region: us_east_1
 storage:
 - labels:
 app: msoffice
 cost: '100'
 zone: us_east_1a
 defaults:
 spaceReserve: volume
 encryption: 'true'
 unixPermissions: '0755'
 - labels:
 app: mysqldb
 cost: '25'
 zone: us_east_1d
 defaults:
 spaceReserve: volume
 encryption: 'false'
 unixPermissions: '0775'
```

```
kubectl create -f backend-tbc-ontap-nas.yaml -n trident
tridentbackendconfig.trident.netapp.io/tbc-ontap-nas-backend created
```

=== Step 3: Verify the status of the `TridentBackendConfig` CR

After the `TridentBackendConfig` has been created, its phase must be `Bound`. It should also reflect the same backend name and UUID as that of the existing backend.

```
kubectl get tbc tbc-ontap-nas-backend -n trident
NAME BACKEND NAME BACKEND UUID PHASE STATUS
tbc-ontap-nas-backend ontap-nas-backend 52f2eb10-e4c6-4160-99fc-96b3be5ab5d7 Bound Success
```

#confirm that no new backends were created (i.e., TridentBackendConfig did not end up creating a new backend)

```
tridentctl get backend -n trident
```

```

| NAME | STORAGE DRIVER | UUID | STATE | VOLUMES |

| ontap-nas-backend | ontap-nas | 52f2eb10-e4c6-4160-99fc-96b3be5ab5d7 | online | 25 |

```

The backend will now be completely managed using the `tbc-ontap-nas-backend` `TridentBackendConfig` object.

== Manage `TridentBackendConfig` backends using `tridentctl`

`tridentctl` can be used to list backends that were created using `TridentBackendConfig`. In addition, administrators can also choose to completely manage such backends through `tridentctl` by deleting `TridentBackendConfig` and making sure `spec.deletionPolicy` is set to `retain`.

=== Step 0: Identify the backend

For example, let us assume the following backend was created using `TridentBackendConfig`:

```
kubectl get tbc backend-tbc-ontap-san -n trident -o wide
NAME BACKEND NAME BACKEND UUID PHASE STATUS STORAGE DRIVER DELETION POLICY
backend-tbc-ontap-san ontap-san-backend 81abcb27-ea63-49bb-b606-0a5315ac5f82 Bound Success ontap-san delete
```

```
tridentctl get backend ontap-san-backend -n trident
```

```

| NAME | STORAGE DRIVER | UUID | STATE | VOLUMES |

| ontap-san-backend | ontap-san | 81abcb27-ea63-49bb-b606-0a5315ac5f82 | online | 33 |

```

---

From the output, it is seen that `TridentBackendConfig` was created successfully and is bound to a backend [observe the backend's UUID].

=== Step 1: Confirm `deletionPolicy` is set to `retain`

Let us take a look at the value of `deletionPolicy`. This needs to be set to `retain`. This ensures that when a `TridentBackendConfig` CR is deleted, the backend definition will still be present and can be managed with `tridentctl`.

```
kubectl get tbc backend-tbc-ontap-san -n trident -o wide
NAME BACKEND NAME BACKEND UUID PHASE STATUS STORAGE DRIVER DELETION POLICY
backend-tbc-ontap-san ontap-san-backend 81abcb27-ea63-49bb-b606-0a5315ac5f82 Bound Success ontap-
san delete
```

## Patch value of deletionPolicy to retain

```
kubectl patch tbc backend-tbc-ontap-san --type=merge -p '{"spec":{"deletionPolicy":"retain"}}' -n trident
tridentbackendconfig.trident.netapp.io/backend-tbc-ontap-san patched
```

#Confirm the value of deletionPolicy

```
kubectl get tbc backend-tbc-ontap-san -n trident -o wide
NAME BACKEND NAME BACKEND UUID PHASE STATUS STORAGE DRIVER DELETION POLICY
backend-tbc-ontap-san ontap-san-backend 81abcb27-ea63-49bb-b606-0a5315ac5f82 Bound Success ontap-
san retain
```

NOTE: Do not proceed to the next step unless `deletionPolicy` is set to `retain`.

=== Step 2: Delete the `TridentBackendConfig` CR

The final step is to delete the `TridentBackendConfig` CR. After confirming the `deletionPolicy` is set to `retain`, you can go ahead with the deletion:

```
kubectl delete tbc backend-tbc-ontap-san -n trident
tridentbackendconfig.trident.netapp.io "backend-tbc-ontap-san" deleted
```

```
tridentctl get backend ontap-san-backend -n trident
```

---

| NAME | STORAGE DRIVER | UUID | STATE | VOLUMES |
|------|----------------|------|-------|---------|
|------|----------------|------|-------|---------|

---

|                   |           |                                      |        |    |
|-------------------|-----------|--------------------------------------|--------|----|
| ontap-san-backend | ontap-san | 81abcb27-ea63-49bb-b606-0a5315ac5f82 | online | 33 |
|-------------------|-----------|--------------------------------------|--------|----|

---

Upon the deletion of the `TridentBackendConfig` object, Trident simply removes it without actually deleting the backend itself.

```
:leveloffset: -1
```

```
:leveloffset: -1
```

```
:leveloffset: -1
```

```
<<<
```

```
Copyright information
```

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

\*Trademark information\*

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