



Manage and protect applications

Trident

NetApp
October 30, 2026

Table of Contents

Manage and protect applications	1
Use Trident Protect AppVault objects to manage buckets	1
Configure AppVault authentication and passwords	1
AppVault creation examples	5
View AppVault information	10
Remove an AppVault	11
Define an application for management with Trident Protect	12
Create an AppVault CR	12
Define an application	12
Protect applications using Trident Protect	17
Create an on-demand snapshot	17
Create an on-demand backup	18
Create a data protection schedule	20
Delete a snapshot	24
Delete a backup	25
Check the status of a backup operation	25
Enable backup and restore for azure-netapp-files (ANF) operations	25
Restore applications	26
Restore applications using Trident Protect	26
Use advanced Trident Protect restore settings	40
Modify resources as you restore them using custom resources	42
Replicate applications using NetApp SnapMirror and Trident Protect	47
Namespace annotations and labels during restore and failover operations	48
Execution hooks during failover and reverse operations	49
Set up a replication relationship	49
Reverse application replication direction	58
Migrate applications using Trident Protect	61
Backup and restore operations	61
Migrate applications from one storage class to another storage class	62
Protect KubeVirt virtual machines with Trident Protect	64
KubeVirt protection modes	64
Resources collected during backup and snapshot	65
Protect VMs using namespace-based application definitions	67
Dynamic VM protection with label selectors	68
Protect specific VMs with includedVirtualMachines	69
Restore specific KubeVirt VMs	70
Create a VM-scoped application using the CLI	75
Related information	75
Manage Trident Protect execution hooks	75
Types of execution hooks	75
Important notes about custom execution hooks	76
Execution hook filters	77
Execution hook examples	77

Create an execution hook	77
Manually run an execution hook	80

Manage and protect applications

Use Trident Protect AppVault objects to manage buckets

The bucket custom resource (CR) for Trident Protect is known as an AppVault. AppVault objects are the declarative Kubernetes workflow representation of a storage bucket. An AppVault CR contains the configurations necessary for a bucket to be used in protection operations, such as backups, snapshots, restore operations, and SnapMirror replication. Only administrators can create AppVaults.

You need to create an AppVault CR manually or from the command line when you perform data protection operations on an application. The AppVault CR is specific to your environment, and you can use the examples on this page as a guide when creating AppVault CRs.



Ensure the AppVault CR is on the cluster where Trident Protect is installed. If the AppVault CR does not exist or you cannot access it, the command line shows an error.

Configure AppVault authentication and passwords

Before you create an AppVault CR, ensure the AppVault and the data mover you choose can authenticate with the provider and any related resources.

Data mover repository passwords

When you create AppVault objects using CRs or the Trident Protect CLI plugin, you can specify a Kubernetes secret with custom passwords for Restic and Kopia encryption. If you don't specify a secret, Trident Protect uses a default password.

- When manually creating AppVault CRs, use the **spec.dataMoverPasswordSecretRef** field to specify the secret.
- When creating AppVault objects using the Trident Protect CLI, use the `--data-mover-password -secret-ref` argument to specify the secret.

Create a data mover repository password secret

Use the following examples to create the password secret. When you create AppVault objects, you can instruct Trident Protect to use this secret to authenticate with the data mover repository.



- Depending on which data mover you are using, you only need to include the corresponding password for that data mover. For example, if you are using Restic and do not plan to use Kopia in the future, you can include only the Restic password when you create the secret.
- Keep the password in a safe place. You will need it to restore data on the same cluster or a different one. If the cluster or the `trident-protect` namespace is deleted, you will not be able to restore your backups or snapshots without the password.

Use a CR

```

apiVersion: v1
data:
  KOPIA_PASSWORD: <base64-encoded-password>
  RESTIC_PASSWORD: <base64-encoded-password>
kind: Secret
metadata:
  name: my-optional-data-mover-secret
  namespace: trident-protect
type: Opaque

```

Use the CLI

```

kubectl create secret generic my-optional-data-mover-secret \
--from-literal=KOPIA_PASSWORD=<plain-text-password> \
--from-literal=RESTIC_PASSWORD=<plain-text-password> \
-n trident-protect

```

S3-compatible storage IAM permissions

When you access S3-compatible storage such as Amazon S3, Generic S3, [StorageGrid S3](#), or [ONTAP S3](#) using Trident Protect, you need to ensure that the user credentials you provide have the necessary permissions to access the bucket. The following is an example of a policy that grants the minimum required permissions for access with Trident Protect. You can apply this policy to the user that manages S3-compatible bucket policies.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
        "s3:GetObject",
        "s3:ListBucket",
        "s3:DeleteObject"
      ],
      "Resource": "*"
    }
  ]
}

```

For more information about Amazon S3 policies, refer to the examples in the [Amazon S3 documentation](#).

EKS Pod Identity for Amazon S3 (AWS) authentication

Trident Protect supports EKS Pod Identity for Kopia data mover operations. This feature enables secure access to S3 buckets without storing AWS credentials in Kubernetes secrets.

Requirements for EKS Pod Identity with Trident Protect

Before using EKS Pod Identity with Trident Protect, ensure the following:

- Your EKS cluster has Pod Identity enabled.
- You have created an IAM role with the necessary S3 bucket permissions. To learn more, refer to [S3-compatible storage IAM permissions](#).
- The IAM role is associated with the following Trident Protect service accounts:
 - <trident-protect>-controller-manager
 - <trident-protect>-resource-backup
 - <trident-protect>-resource-restore
 - <trident-protect>-resource-delete

For detailed instructions on enabling Pod Identity and associating IAM roles with service accounts, refer to the [AWS EKS Pod Identity documentation](#).

AppVault Configuration

When using EKS Pod Identity, configure your AppVault CR with the `useIAM: true` flag instead of explicit credentials:

```
apiVersion: protect.trident.netapp.io/v1
kind: AppVault
metadata:
  name: eks-protect-vault
  namespace: trident-protect
spec:
  providerType: AWS
  providerConfig:
    s3:
      bucketName: trident-protect-aws
      endpoint: s3.example.com
      useIAM: true
```

AppVault key generation examples for cloud providers

When defining an AppVault CR, you need to include credentials to access the resources hosted by the provider, unless you are using IAM authentication. How you generate the keys for the credentials will differ depending on the provider. The following are command line key generation examples for several providers. You can use the following examples to create keys for the credentials of each cloud provider.

Google Cloud

```
kubectl create secret generic <secret-name> \  
--from-file=credentials=<mycreds-file.json> \  
-n trident-protect
```

Amazon S3 (AWS)

```
kubectl create secret generic <secret-name> \  
--from-literal=accessKeyID=<objectstorage-accesskey> \  
--from-literal=secretAccessKey=<amazon-s3-trident-protect-src-bucket  
-secret> \  
-n trident-protect
```

Microsoft Azure

```
kubectl create secret generic <secret-name> \  
--from-literal=accountKey=<secret-name> \  
-n trident-protect
```

Generic S3

```
kubectl create secret generic <secret-name> \  
--from-literal=accessKeyID=<objectstorage-accesskey> \  
--from-literal=secretAccessKey=<generic-s3-trident-protect-src-bucket  
-secret> \  
-n trident-protect
```

ONTAP S3

```
kubectl create secret generic <secret-name> \  
--from-literal=accessKeyID=<objectstorage-accesskey> \  
--from-literal=secretAccessKey=<ontap-s3-trident-protect-src-bucket  
-secret> \  
-n trident-protect
```

StorageGrid S3

```
kubectl create secret generic <secret-name> \  
--from-literal=accessKeyID=<objectstorage-accesskey> \  
--from-literal=secretAccessKey=<storagegrid-s3-trident-protect-src  
-bucket-secret> \  
-n trident-protect
```

AppVault creation examples

The following are example AppVault definitions for each provider.

AppVault CR examples

You can use the following CR examples to create AppVault objects for each cloud provider.



- You can optionally specify a Kubernetes secret that contains custom passwords for the Restic and Kopia repository encryption. Refer to [Data mover repository passwords](#) for more information.
- For Amazon S3 (AWS) AppVault objects, you can optionally specify a sessionToken, which is useful if you are using single sign-on (SSO) for authentication. This token is created when you generate keys for the provider in [AppVault key generation examples for cloud providers](#).
- For S3 AppVault objects, you can optionally specify an egress proxy URL for outbound S3 traffic using the `spec.providerConfig.S3.proxyURL` key.

Google Cloud

```
apiVersion: protect.trident.netapp.io/v1
kind: AppVault
metadata:
  name: gcp-trident-protect-src-bucket
  namespace: trident-protect
spec:
  dataMoverPasswordSecretRef: my-optional-data-mover-secret
  providerType: GCP
  providerConfig:
    gcp:
      bucketName: trident-protect-src-bucket
      projectID: project-id
  providerCredentials:
    credentials:
      valueFromSecret:
        key: credentials
        name: gcp-trident-protect-src-bucket-secret
```

Amazon S3 (AWS)

```
---
apiVersion: protect.trident.netapp.io/v1
kind: AppVault
metadata:
  name: amazon-s3-trident-protect-src-bucket
  namespace: trident-protect
spec:
  dataMoverPasswordSecretRef: my-optional-data-mover-secret
```

```

providerType: AWS
providerConfig:
  s3:
    bucketName: trident-protect-src-bucket
    endpoint: s3.example.com
    proxyURL: http://10.1.1.1:3128
providerCredentials:
  accessKeyID:
    valueFromSecret:
      key: accessKeyID
      name: s3-secret
  secretAccessKey:
    valueFromSecret:
      key: secretAccessKey
      name: s3-secret
  sessionToken:
    valueFromSecret:
      key: sessionToken
      name: s3-secret

```



For EKS environments using Pod Identity with Kopia data mover, you can remove the `providerCredentials` section and add `useIAM: true` under the `s3` configuration instead.

Microsoft Azure

```

apiVersion: protect.trident.netapp.io/v1
kind: AppVault
metadata:
  name: azure-trident-protect-src-bucket
  namespace: trident-protect
spec:
  dataMoverPasswordSecretRef: my-optional-data-mover-secret
  providerType: Azure
  providerConfig:
    azure:
      accountName: account-name
      bucketName: trident-protect-src-bucket
  providerCredentials:
    accountKey:
      valueFromSecret:
        key: accountKey
        name: azure-trident-protect-src-bucket-secret

```

Generic S3

```
apiVersion: protect.trident.netapp.io/v1
kind: AppVault
metadata:
  name: generic-s3-trident-protect-src-bucket
  namespace: trident-protect
spec:
  dataMoverPasswordSecretRef: my-optional-data-mover-secret
  providerType: GenericS3
  providerConfig:
    s3:
      bucketName: trident-protect-src-bucket
      endpoint: s3.example.com
      proxyURL: http://10.1.1.1:3128
  providerCredentials:
    accessKeyID:
      valueFromSecret:
        key: accessKeyID
        name: s3-secret
    secretAccessKey:
      valueFromSecret:
        key: secretAccessKey
        name: s3-secret
```

ONTAP S3

```
apiVersion: protect.trident.netapp.io/v1
kind: AppVault
metadata:
  name: ontap-s3-trident-protect-src-bucket
  namespace: trident-protect
spec:
  dataMoverPasswordSecretRef: my-optional-data-mover-secret
  providerType: OntapS3
  providerConfig:
    s3:
      bucketName: trident-protect-src-bucket
      endpoint: s3.example.com
      proxyURL: http://10.1.1.1:3128
  providerCredentials:
    accessKeyID:
      valueFromSecret:
        key: accessKeyID
        name: s3-secret
    secretAccessKey:
```

```
valueFromSecret:
  key: secretAccessKey
  name: s3-secret
```

StorageGrid S3

```
apiVersion: protect.trident.netapp.io/v1
kind: AppVault
metadata:
  name: storagegrid-s3-trident-protect-src-bucket
  namespace: trident-protect
spec:
  dataMoverPasswordSecretRef: my-optional-data-mover-secret
  providerType: StorageGridS3
  providerConfig:
    s3:
      bucketName: trident-protect-src-bucket
      endpoint: s3.example.com
      proxyURL: http://10.1.1.1:3128
  providerCredentials:
    accessKeyID:
      valueFromSecret:
        key: accessKeyID
        name: s3-secret
    secretAccessKey:
      valueFromSecret:
        key: secretAccessKey
        name: s3-secret
```

AppVault creation examples using the Trident Protect CLI

You can use the following CLI command examples to create AppVault CRs for each provider.



- You can optionally specify a Kubernetes secret that contains custom passwords for the Restic and Kopia repository encryption. Refer to [Data mover repository passwords](#) for more information.
- For S3 AppVault objects, you can optionally specify an egress proxy URL for outbound S3 traffic using the `--proxy-url <ip_address:port>` argument.

Google Cloud

```
tridentctl-protect create vault GCP <vault-name> \
--bucket <mybucket> \
--project <my-gcp-project> \
```

```
--secret <secret-name>/credentials \
--data-mover-password-secret-ref <my-optional-data-mover-secret> \
-n trident-protect
```

Amazon S3 (AWS)

```
tridentctl-protect create vault AWS <vault-name> \
--bucket <bucket-name> \
--secret <secret-name> \
--endpoint <s3-endpoint> \
--data-mover-password-secret-ref <my-optional-data-mover-secret> \
-n trident-protect
```

Microsoft Azure

```
tridentctl-protect create vault Azure <vault-name> \
--account <account-name> \
--bucket <bucket-name> \
--secret <secret-name> \
--data-mover-password-secret-ref <my-optional-data-mover-secret> \
-n trident-protect
```

Generic S3

```
tridentctl-protect create vault GenericS3 <vault-name> \
--bucket <bucket-name> \
--secret <secret-name> \
--endpoint <s3-endpoint> \
--data-mover-password-secret-ref <my-optional-data-mover-secret> \
-n trident-protect
```

ONTAP S3

```
tridentctl-protect create vault OntapS3 <vault-name> \
--bucket <bucket-name> \
--secret <secret-name> \
--endpoint <s3-endpoint> \
--data-mover-password-secret-ref <my-optional-data-mover-secret> \
-n trident-protect
```

StorageGrid S3

```
tridentctl-protect create vault StorageGridS3 <vault-name> \
--bucket <bucket-name> \
```

```
--secret <secret-name> \
--endpoint <s3-endpoint> \
--data-mover-password-secret-ref <my-optional-data-mover-secret> \
-n trident-protect
```

Supported `providerConfig.s3` configuration options

See the following table for the S3 provider configuration options:

Parameter	Description	Default	Example
<code>providerConfig.s3.skipCertValidation</code>	Disable SSL/TLS certificate verification.	false	"true", "false"
<code>providerConfig.s3.secure</code>	Enable secure HTTPS communication with the S3 endpoint.	true	"true", "false"
<code>providerConfig.s3.proxyURL</code>	Specify the URL of the proxy server used to connect to S3.	None	http://proxy.example.com:8080
<code>providerConfig.s3.rootCA</code>	Provide a custom root CA certificate for SSL/TLS verification.	None	"CN=MyCustomCA"
<code>providerConfig.s3.useIAM</code>	Enable IAM authentication for accessing S3 buckets. Applicable for EKS Pod Identity.	false	true, false

View AppVault information

You can use the Trident Protect CLI plugin to view information about AppVault objects that you have created on the cluster.

Steps

1. View the contents of an AppVault object:

```
tridentctl-protect get appvaultcontent gcp-vault \
--show-resources all \
-n trident-protect
```

Example output:

```
+-----+-----+-----+-----+
+-----+
| CLUSTER | APP | TYPE | NAME |
```

```

TIMESTAMP |
+-----+-----+-----+-----+
+-----+
|          | mysql | snapshot | mysnap          | 2024-
08-09 21:02:11 (UTC) |
| production1 | mysql | snapshot | hourly-e7db6-20240815180300 | 2024-
08-15 18:03:06 (UTC) |
| production1 | mysql | snapshot | hourly-e7db6-20240815190300 | 2024-
08-15 19:03:06 (UTC) |
| production1 | mysql | snapshot | hourly-e7db6-20240815200300 | 2024-
08-15 20:03:06 (UTC) |
| production1 | mysql | backup   | hourly-e7db6-20240815180300 | 2024-
08-15 18:04:25 (UTC) |
| production1 | mysql | backup   | hourly-e7db6-20240815190300 | 2024-
08-15 19:03:30 (UTC) |
| production1 | mysql | backup   | hourly-e7db6-20240815200300 | 2024-
08-15 20:04:21 (UTC) |
| production1 | mysql | backup   | mybackup5          | 2024-
08-09 22:25:13 (UTC) |
|          | mysql | backup   | mybackup          | 2024-
08-09 21:02:52 (UTC) |
+-----+-----+-----+-----+
+-----+

```

2. Optionally, to see the AppVaultPath for each resource, use the flag `--show-paths`.

The cluster name in the first column of the table is only available if a cluster name was specified in the Trident Protect helm installation. For example: `--set clusterName=production1`.

Remove an AppVault

You can remove an AppVault object at any time.



Do not remove the `finalizers` key in the AppVault CR before deleting the AppVault object. If you do so, it can result in residual data in the AppVault bucket and orphaned resources in the cluster.

Before you begin

Ensure that you have deleted all snapshot and backup CRs being used by the AppVault you want to delete.

Remove an AppVault using the Kubernetes CLI

1. Remove the AppVault object, replacing `appvault-name` with the name of the AppVault object to remove:

```
kubectl delete appvault <appvault-name> \
```

```
-n trident-protect
```

Remove an AppVault using the Trident Protect CLI

1. Remove the AppVault object, replacing `appvault-name` with the name of the AppVault object to remove:

```
tridentctl-protect delete appvault <appvault-name> \  
-n trident-protect
```

Define an application for management with Trident Protect

You can define an application that you want to manage with Trident Protect by creating an application CR and an associated AppVault CR.

Create an AppVault CR

You need to create an AppVault CR that will be used when performing data protection operations on the application, and the AppVault CR needs to reside on the cluster where Trident Protect is installed. The AppVault CR is specific to your environment; for examples of AppVault CRs, refer to [AppVault custom resources](#).

Define an application

You need to define each application that you want to manage with Trident Protect. You can define an application for management by either manually creating an application CR or by using the Trident Protect CLI.



To define a VM-scoped Application that protects specific KubeVirt virtual machines (VMs) by name instead of entire namespaces, use the `spec.includedVirtualMachines` field in the Application CR or the `--virtual-machines` flag in the CLI. For details, field reference, and YAML examples, refer to [Protect KubeVirt virtual machines with Trident Protect](#).

Add an application using a CR

Steps

1. Create the destination application CR file:
 - a. Create the custom resource (CR) file and name it (for example, `maria-app.yaml`).
 - b. Configure the following attributes:
 - **metadata.name:** (*Required*) The name of the application custom resource. Note the name you choose because other CR files needed for protection operations refer to this value.
 - **spec.includedNamespaces:** (*Required*) Use namespace and label selector to specify the namespaces and resources that the application uses. The application namespace must be part of this list. The label selector is optional and can be used to filter resources within each specified namespace.
 - **spec.includedClusterScopedResources:** (*Optional*) Use this attribute to specify cluster-scoped resources to be included in the application definition. This attribute allows you to

select these resources based on their group, version, kind, and labels.

- **groupVersionKind:** (*Required*) Specifies the API group, version, and kind of the cluster-scoped resource.
- **labelSelector:** (*Optional*) Filters the cluster-scoped resources based on their labels.
- **metadata.annotations.protect.trident.netapp.io/skip-vm-freeze:** (*Optional*) This annotation is only applicable to applications defined from virtual machines, such as in KubeVirt environments, where filesystem freezes occur before snapshots. Specify whether this application can write to the filesystem during a snapshot. If set to true, the application ignores the global setting and can write to the filesystem during a snapshot. If set to false, the application ignores the global setting and the filesystem is frozen during a snapshot. If specified but the application has no virtual machines in the application definition, the annotation is ignored. If not specified, the application follows the [global Trident Protect freeze setting](#).

If you need to apply this annotation after an application has already been created, you can use the following command:



```
kubectl annotate application -n <application CR namespace> <application CR name>
protect.trident.netapp.io/skip-vm-freeze="true"
```

Example YAML:

```
apiVersion: protect.trident.netapp.io/v1
kind: Application
metadata:
  annotations:
    protect.trident.netapp.io/skip-vm-freeze: "false"
  name: my-app-name
  namespace: my-app-namespace
spec:
  includedNamespaces:
    - namespace: namespace-1
      labelSelector:
        matchLabels:
          app: example-app
    - namespace: namespace-2
      labelSelector:
        matchLabels:
          app: another-example-app
  includedClusterScopedResources:
    - groupVersionKind:
        group: rbac.authorization.k8s.io
        kind: ClusterRole
        version: v1
```

```
labelSelector:
  matchLabels:
    mylabel: test
```

2. (Optional) If needed, you can add resource filtering to the same CR to include or exclude specific resources:

- **Generic filter example:**

- **resourceFilter.resourceSelectionCriteria:** (Required for filtering) Use `Include` or `Exclude` to include or exclude a resource defined in `resourceMatchers`. Add the following `resourceMatchers` parameters to define the resources to be included or excluded:
 - **resourceFilter.resourceMatchers:** An array of `resourceMatcher` objects. If you define multiple elements in this array, they match as an OR operation, and the fields inside each element (group, kind, version) match as an AND operation.
 - **resourceMatchers[].group:** (Optional) Group of the resource to be filtered.
 - **resourceMatchers[].kind:** (Optional) Kind of the resource to be filtered.
 - **resourceMatchers[].version:** (Optional) Version of the resource to be filtered.
 - **resourceMatchers[].names:** (Optional) Names in the Kubernetes metadata.name field of the resource to be filtered.
 - **resourceMatchers[].namespaces:** (Optional) Namespaces in the Kubernetes metadata.name field of the resource to be filtered.
 - **resourceMatchers[].labelSelectors:** (Optional) Label selector string in the Kubernetes metadata.name field of the resource as defined in the [Kubernetes documentation](#). For example: "trident.netapp.io/os=linux".



When both `resourceFilter` and `labelSelector` are used, `resourceFilter` runs first, and then `labelSelector` is applied to the resulting resources.

For example:

```
spec:
  resourceFilter:
    resourceSelectionCriteria: "Include"
    resourceMatchers:
      - group: my-resource-group-1
        kind: my-resource-kind-1
        version: my-resource-version-1
        names: ["my-resource-names"]
        namespaces: ["my-resource-namespaces"]
        labelSelectors: ["trident.netapp.io/os=linux"]
      - group: my-resource-group-2
        kind: my-resource-kind-2
        version: my-resource-version-2
        names: ["my-resource-names"]
        namespaces: ["my-resource-namespaces"]
```

```
labelSelectors: ["trident.netapp.io/os=linux"]
```

- **PVC-only filter example:**

To define a PVC-only application, you must also include `PersistentVolume` and `VolumeSnapshotClass` in the resource filter. Snapshot and backup operations depend on `PersistentVolume` (the cluster-scoped volume bound to each PVC) and `VolumeSnapshotClass` (the snapshot driver), and will fail without them. For example:

```
apiVersion: protect.trident.netapp.io/v1
kind: Application
metadata:
  name: my-pvc-app
  namespace: my-app-namespace
spec:
  includedNamespaces:
  - namespace: my-app-namespace
  resourceFilter:
    resourceMatchers:
    - kind: PersistentVolumeClaim
      version: v1
    - kind: PersistentVolume
      version: v1
    - kind: VolumeSnapshotClass
      version: v1
    resourceSelectionCriteria: Include
```

3. After you create the application CR to match your environment, apply the CR. For example:

```
kubectl apply -f maria-app.yaml
```

Add an application using the CLI

Steps

1. Create and apply the application definition using one of the following examples, replacing values in brackets with information from your environment. You can include namespaces and resources in the application definition using comma-separated lists with the arguments shown in the examples.

You can optionally use an annotation when you create an app to specify whether the application can write to the filesystem during a snapshot. This is only applicable to applications defined from virtual machines, such as in KubeVirt environments, where filesystem freezes occur before snapshots. If you set the annotation to `true`, the application ignores the global setting and can write to the filesystem during a snapshot. If you set it to `false`, the application ignores the global setting and the filesystem is frozen during a snapshot. If you use the annotation but the application has no virtual machines in the application definition, the annotation is ignored. If you don't use the annotation, the application follows the [global Trident Protect freeze setting](#).

+

To specify the annotation when you use the CLI to create an application, you can use the `--annotation` flag.

+

* Create the application and use the global setting for filesystem freeze behavior:

+

```
tridentctl-protect create application <my_new_app_cr_name> --namespaces
<namespaces_to_include> --csr <cluster_scoped_resources_to_include>
--namespace <my-app-namespace>
```

+

* Create the application and configure the local application setting for filesystem freeze behavior:

+

```
tridentctl-protect create application <my_new_app_cr_name> --namespaces
<namespaces_to_include> --csr <cluster_scoped_resources_to_include>
--namespace <my-app-namespace> --annotation
protect.trident.netapp.io/skip-vm-freeze=<"true"|"false">
```

+

* You can use `--resource-filter-include` and `--resource-filter-exclude` flags to include or exclude resources based on `resourceSelectionCriteria` such as group, kind, version, labels, names, and namespaces, as shown in the following example:

+

```
tridentctl-protect create application <my_new_app_cr_name> --namespaces
<namespaces_to_include> --csr <cluster_scoped_resources_to_include>
--namespace <my-app-namespace> --resource-filter-include
'[{ "Group": "apps", "Kind": "Deployment", "Version": "v1", "Names": [ "my-
deployment" ], "Namespaces": [ "my-namespace" ], "LabelSelectors": [ "app=my-
app" ] } ] '
```

+

* To define a PVC-only application, you must also include `PersistentVolume` and `VolumeSnapshotClass`` in the resource filter. Snapshot and backup operations depend on `PersistentVolume` (the cluster-scoped volume bound to each PVC) and `VolumeSnapshotClass` (the snapshot driver), and will fail without them. For example:

+

```
tridentctl-protect create app my-pvc-app --namespaces <my-app-
```

```
namespace> --resource-filter-include
' [{"Kind": "PersistentVolumeClaim", "Version": "v1"}, {"Kind": "PersistentVo
lume", "Version": "v1"}, {"Kind": "VolumeSnapshotClass", "Version": "v1"} ] '
-n <my-app-namespace>
```

Protect applications using Trident Protect

You can protect all apps managed by Trident Protect by taking snapshots and backups using an automated protection policy or on an ad-hoc basis.



You can configure Trident Protect to freeze and unfreeze filesystems during data protection operations. [Learn more about configuring filesystem freezing with Trident Protect.](#)

Create an on-demand snapshot

You can create an on-demand snapshot at any time.



Cluster-scoped resources are included in a backup, snapshot, or clone if they are explicitly referenced in the application definition or if they have references to any of the application namespaces.

Create a snapshot using a CR

Steps

1. Create the custom resource (CR) file and name it `trident-protect-snapshot-cr.yaml`.
2. In the file you created, configure the following attributes:
 - **metadata.name:** (*Required*) The name of this custom resource; choose a unique and sensible name for your environment.
 - **spec.applicationRef:** The Kubernetes name of the application to snapshot.
 - **spec.appVaultRef:** (*Required*) The name of the AppVault where the snapshot contents (metadata) should be stored.
 - **spec.reclaimPolicy:** (*Optional*) Defines what happens to the AppArchive of a snapshot when the snapshot CR is deleted. This means that even when set to `Retain`, the snapshot will be deleted. Valid options:
 - `Retain` (default)
 - `Delete`

```
apiVersion: protect.trident.netapp.io/v1
kind: Snapshot
metadata:
  namespace: my-app-namespace
  name: my-cr-name
spec:
```

```
applicationRef: my-application
appVaultRef: appvault-name
reclaimPolicy: Delete
```

3. After you populate the `trident-protect-snapshot-cr.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-snapshot-cr.yaml
```

Create a snapshot using the CLI

Steps

1. Create the snapshot, replacing values in brackets with information from your environment. For example:

```
tridentctl-protect create snapshot <my_snapshot_name> --appvault
<my_appvault_name> --app <name_of_app_to_snapshot> -n
<application_namespace>
```

Create an on-demand backup

You can back up an app at any time.



Cluster-scoped resources are included in a backup, snapshot, or clone if they are explicitly referenced in the application definition or if they have references to any of the application namespaces.

Before you begin

Ensure that the AWS session token expiration is sufficient for any long-running s3 backup operations. If the token expires during the backup operation, the operation can fail.

- Refer to the [AWS API documentation](#) for more information about checking the current session token expiration.
- Refer to the [AWS IAM documentation](#) for more information about credentials with AWS resources.

Create a backup using a CR

Steps

1. Create the custom resource (CR) file and name it `trident-protect-backup-cr.yaml`.
2. In the file you created, configure the following attributes:
 - **metadata.name:** *(Required)* The name of this custom resource; choose a unique and sensible name for your environment.
 - **spec.applicationRef:** *(Required)* The Kubernetes name of the application to back up.
 - **spec.appVaultRef:** *(Required)* The name of the AppVault where the backup contents should be stored.

- **spec.dataMover:** (*Optional*) A string indicating which backup tool to use for the backup operation. Possible values (case sensitive):
 - Restic
 - Kopia (default)
- **spec.reclaimPolicy:** (*Optional*) Defines what happens to a backup when released from its claim. Possible values:
 - Delete
 - Retain (default)
- **spec.snapshotRef:** (*Optional*): Name of the snapshot to use as the source of the backup. If not provided, a temporary snapshot will be created and backed up.

Example YAML:

```
---
apiVersion: protect.trident.netapp.io/v1
kind: Backup
metadata:
  namespace: my-app-namespace
  name: my-cr-name
spec:
  applicationRef: my-application
  appVaultRef: appvault-name
  dataMover: Kopia
```

3. After you populate the `trident-protect-backup-cr.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-backup-cr.yaml
```

Create a backup using the CLI

Steps

1. Create the backup, replacing values in brackets with information from your environment. For example:

```
tridentctl-protect create backup <my_backup_name> --appvault <my-vault-name> --app <name_of_app_to_back_up> --data-mover <Kopia_or_Restic> -n <application_namespace>
```

You can optionally use the `--full-backup` flag to specify whether a backup should be non-incremental. By default, all backups are incremental. When this flag is used, the backup becomes non-incremental. It is best practice to perform a full backup periodically and then perform incremental backups in between full backups to minimize the risk associated with restores.

Supported backup annotations

The following table describes the annotations you can use when creating a backup CR:

Annotation	Type	Description	Default value
protect.trident.netapp.io/full-backup	string	Specifies whether a backup should be non-incremental. Set to <code>true</code> to create a non-incremental backup. It is best practice to perform a full backup periodically and then perform incremental backups in between full backups to minimize the risk associated with restores.	"false"
protect.trident.netapp.io/snapshots-hot-completion-timeout	string	The maximum time allowed for the overall snapshot operation to complete.	"60m"
protect.trident.netapp.io/volume-snapshots-ready-to-use-timeout	string	The maximum time allowed for volume snapshots to reach the ready-to-use state.	"30m"
protect.trident.netapp.io/volume-snapshots-created-timeout	string	The maximum time allowed for volume snapshots to be created.	"5m"
protect.trident.netapp.io/pvc-bind-timeout-sec	string	Maximum time (in seconds) to wait for any newly created PersistentVolumeClaims (PVCs) to reach the <code>Bound</code> phase before the operations fails.	"1200" (20 minutes)

Create a data protection schedule

A protection policy protects an app by creating snapshots, backups, or both at a defined schedule. You can choose to create snapshots and backups hourly, daily, weekly, and monthly, and you can specify the number of copies to retain. You can schedule a non-incremental full backup by using the `full-backup-rule` annotation. By default, all backups are incremental. Performing a full backup periodically, along with incremental backups in between, helps reduce the risk associated with restores.



- You can create schedules for snapshots only by setting `backupRetention` to zero and `snapshotRetention` to a value greater than zero. Setting `snapshotRetention` to zero means any scheduled backups will still create snapshots, but those are temporary and get deleted immediately after the backup is completed.
- Cluster-scoped resources are included in a backup, snapshot, or clone if they are explicitly referenced in the application definition or if they have references to any of the application namespaces.

Create a schedule using a CR

Steps

1. Create the custom resource (CR) file and name it `trident-protect-schedule-cr.yaml`.
2. In the file you created, configure the following attributes:
 - **metadata.name:** (*Required*) The name of this custom resource; choose a unique and sensible name for your environment.
 - **spec.dataMover:** (*Optional*) A string indicating which backup tool to use for the backup operation.

Possible values (case sensitive):

- Restic
- Kopia (default)
- **spec.applicationRef:** The Kubernetes name of the application to back up.
- **spec.appVaultRef:** (*Required*) The name of the AppVault where the backup contents should be stored.
- **spec.backupRetention:** (*Required*) The number of backups to retain. Zero indicates that no backups should be created (snapshots only).
- **spec.backupReclaimPolicy:** (*Optional*) Determines what happens to a backup if the backup CR is deleted during its retention period. After the retention period, backups are always deleted.
Possible values (case sensitive):
 - Retain (default)
 - Delete
- **spec.snapshotRetention:** (*Required*) The number of snapshots to retain. Zero indicates that no snapshots should be created.
- **spec.snapshotReclaimPolicy:** (*Optional*) Determines what happens to a snapshot if the snapshot CR is deleted during its retention period. After the retention period, snapshots are always deleted. Possible values (case sensitive):
 - Retain
 - Delete (default)
- **spec.granularity:** The frequency at which the schedule should run. Possible values, along with required associated fields:
 - Hourly (requires that you specify `spec.minute`)
 - Daily (requires that you specify `spec.minute` and `spec.hour`)
 - Weekly (requires that you specify `spec.minute`, `spec.hour`, and `spec.dayOfWeek`)
 - Monthly (requires that you specify `spec.minute`, `spec.hour`, and `spec.dayOfMonth`)
 - Custom
- **spec.dayOfMonth:** (*Optional*) The day of the month (1 - 31) that the schedule should run. This field is required if the granularity is set to `Monthly`. The value must be provided as a string.
- **spec.dayOfWeek:** (*Optional*) The day of the week (0 - 7) that the schedule should run. Values of 0 or 7 indicate Sunday. This field is required if the granularity is set to `Weekly`. The value must be provided as a string.
- **spec.hour:** (*Optional*) The hour of the day (0 - 23) that the schedule should run. This field is required if the granularity is set to `Daily`, `Weekly`, or `Monthly`. The value must be provided as a string.
- **spec.minute:** (*Optional*) The minute of the hour (0 - 59) that the schedule should run. This field is required if the granularity is set to `Hourly`, `Daily`, `Weekly`, or `Monthly`. The value must be provided as a string.
- **spec.runImmediately:** (*Optional*) Set to `true` to trigger a one-time, immediate baseline run (backup and/or snapshot per retention settings) when the schedule is created. Defaults to `false`. This does not modify subsequent recurrence.

Example YAML for backup and snapshot schedule:

```
---
apiVersion: protect.trident.netapp.io/v1
kind: Schedule
metadata:
  namespace: my-app-namespace
  name: my-cr-name
spec:
  dataMover: Kopia
  applicationRef: my-application
  appVaultRef: appvault-name
  backupRetention: "15"
  snapshotRetention: "15"
  granularity: Daily
  hour: "0"
  minute: "0"
```

Example YAML for snapshot-only schedule:

```
---
apiVersion: protect.trident.netapp.io/v1
kind: Schedule
metadata:
  namespace: my-app-namespace
  name: my-snapshot-schedule
spec:
  applicationRef: my-application
  appVaultRef: appvault-name
  backupRetention: "0"
  snapshotRetention: "15"
  granularity: Daily
  hour: "2"
  minute: "0"
```

Example YAML for schedule with immediate run:

```
---
apiVersion: protect.trident.netapp.io/v1
kind: Schedule
metadata:
  namespace: my-app-namespace
  name: my-daily-schedule-run-immediately
spec:
```

```

applicationRef: my-application
appVaultRef: appvault-name
backupRetention: "7"
snapshotRetention: "7"
granularity: Daily
hour: "3"
minute: "0"
runImmediately: true

```

3. After you populate the `trident-protect-schedule-cr.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-schedule-cr.yaml
```

Create a schedule using the CLI

Steps

1. Create the protection schedule, replacing values in brackets with information from your environment. For example:



You can use `tridentctl-protect create schedule --help` to view detailed help information for this command.

```

tridentctl-protect create schedule <my_schedule_name> \
  --appvault <my_appvault_name> \
  --app <name_of_app_to_snapshot> \
  --backup-retention <how_many_backups_to_retain> \
  --backup-reclaim-policy <Retain|Delete (default Retain)> \
  --data-mover <Kopia_or_Restic> \
  --day-of-month <day_of_month_to_run_schedule> \
  --day-of-week <day_of_week_to_run_schedule> \
  --granularity <frequency_to_run> \
  --hour <hour_of_day_to_run> \
  --minute <minute_of_hour_to_run> \
  --recurrence-rule <recurrence> \
  --snapshot-retention <how_many_snapshots_to_retain> \
  --snapshot-reclaim-policy <Retain|Delete (default Delete)> \
  --full-backup-rule <string> \
  --run-immediately <true|false> \
  -n <application_namespace>

```

The following flags provide additional control over your schedule:

- **Full backup scheduling:** Use the `--full-backup-rule` flag to schedule non-incremental full backups. This flag only works with `--granularity Daily`. Possible values:

- **Always:** Create a full backup every day.
- **Specific weekdays:** Specify one or more days separated by commas (for example, "Monday, Thursday"). Valid values: Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, Sunday.



The `--full-backup-rule` flag does not work with Hourly, Weekly, or Monthly granularity.

- **Immediate baseline protection:** Use `--run-immediately true` to create an initial backup or snapshot immediately when the schedule is created, rather than waiting for the first scheduled run time. Default is `false`.
- **Snapshot-only schedules:** Set `--backup-retention 0` and specify a value greater than zero for `--snapshot-retention`.

Supported schedule annotations

The following table describes the annotations you can use when creating a schedule CR:

Annotation	Type	Description	Default value
<code>protect.trident.netapp.io/full-backup-rule</code>	string	Specifies the rule for scheduling full backups. You can set it to <code>Always</code> for constant full backup or customize it based on your requirements. For example, if you choose daily granularity, you can specify the weekdays on which full backup should occur (for example, "Monday, Thursday"). Valid weekday values are: Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, Sunday. Note that this annotation can only be used with schedules that have <code>granularity</code> set to <code>Daily</code> .	Not set (all backups are incremental)
<code>protect.trident.netapp.io/snapshots-hot-completion-timeout</code>	string	The maximum time allowed for the overall snapshot operation to complete.	"60m"
<code>protect.trident.netapp.io/volume-snapshots-ready-to-use-timeout</code>	string	The maximum time allowed for volume snapshots to reach the ready-to-use state.	"30m"
<code>protect.trident.netapp.io/volume-snapshots-created-timeout</code>	string	The maximum time allowed for volume snapshots to be created.	"5m"
<code>protect.trident.netapp.io/pvc-bind-timeout-sec</code>	string	Maximum time (in seconds) to wait for any newly created PersistentVolumeClaims (PVCs) to reach the <code>Bound</code> phase before the operations fails.	"1200" (20 minutes)

Delete a snapshot

Delete the scheduled or on-demand snapshots that you no longer need.

Steps

1. Remove the snapshot CR associated with the snapshot:

```
kubectl delete snapshot <snapshot_name> -n my-app-namespace
```

Delete a backup

Delete the scheduled or on-demand backups that you no longer need.



Ensure the reclaim policy is set to `Delete` to remove all backup data from object storage. The default setting of the policy is `Retain` to avoid accidental data loss. If the policy is not changed to `Delete`, the backup data will remain in object storage and will require manual deletion.

Steps

1. Remove the backup CR associated with the backup:

```
kubectl delete backup <backup_name> -n my-app-namespace
```

Check the status of a backup operation

You can use the command line to check the status of a backup operation that is in progress, has completed, or has failed.

Steps

1. Use the following command to retrieve status of the backup operation, replacing values in brackets with information from your environment:

```
kubectl get backup -n <namespace_name> <my_backup_cr_name> -o jsonpath='{.status}'
```

Enable backup and restore for azure-netapp-files (ANF) operations

If you have installed Trident Protect, you can enable space-efficient backup and restore functionality for storage backends that use the `azure-netapp-files` storage class and were created prior to Trident 24.06. This functionality works with NFSv4 volumes and does not consume additional space from the capacity pool.

Before you begin

Ensure the following:

- You have installed Trident Protect.
- You have defined an application in Trident Protect. This application will have limited protection functionality until you complete this procedure.
- You have `azure-netapp-files` selected as the default storage class for your storage backend.

▼ Expand for configuration steps

1. Do the following in Trident if the ANF volume was created prior to upgrading to Trident 24.10:

- a. Enable the snapshot directory for each PV that is azure-netapp-files based and associated with the application:

```
tridentctl update volume <pv name> --snapshot-dir=true -n trident
```

- b. Confirm that the snapshot directory has been enabled for each associated PV:

```
tridentctl get volume <pv name> -n trident -o yaml | grep  
snapshotDir
```

Response:

```
snapshotDirectory: "true"
```

When the snapshot directory is not enabled, Trident Protect chooses the regular backup functionality, which temporarily consumes space in the capacity pool during the backup process. In this case, ensure that sufficient space is available in the capacity pool to create a temporary volume of the size of the volume being backed up.

Result

The application is ready for backup and restore using Trident Protect. Each PVC is also available to be used by other applications for backups and restores.

Restore applications

Restore applications using Trident Protect

You can use Trident Protect to restore your application from a snapshot or backup. Restoring from an existing snapshot will be faster when restoring the application to the same cluster.



- When you restore an application, all execution hooks configured for the application are restored with the app. If a post-restore execution hook is present, it runs automatically as part of the restore operation.
- Restoring from a backup to a different namespace or to the original namespace is supported for qtree volumes. However, restoring from a snapshot to a different namespace or to the original namespace is not supported for qtree volumes.
- You can use advanced settings to customize restore operations. To learn more, refer to [Use advanced Trident Protect restore settings](#).
- To restore specific KubeVirt virtual machines (VMs) from an archive without restoring all namespace resources, use the `includedVirtualMachines` field on a restore CR. For supported CRs, field reference, and YAML examples, refer to [Protect KubeVirt virtual machines with Trident Protect](#).

Restore from a backup to a different namespace

When you restore a backup to a different namespace using a BackupRestore CR, Trident Protect restores the application in a new namespace and creates an application CR for the restored application. To protect the restored application, create on-demand backups or snapshots, or establish a protection schedule.



- Restoring a backup to a different namespace with existing resources will not alter any resources that share names with those in the backup. To restore all resources in the backup, either delete and re-create the target namespace, or restore the backup to a new namespace.
- When using a CR to restore to a new namespace, you must manually create the destination namespace before applying the CR. Trident Protect automatically creates namespaces only when using the CLI.

Before you begin

Ensure that the AWS session token expiration is sufficient for any long-running s3 restore operations. If the token expires during the restore operation, the operation can fail.

- Refer to the [AWS API documentation](#) for more information about checking the current session token expiration.
- Refer to the [AWS IAM documentation](#) for more information about credentials with AWS resources.



When you restore backups using Kopia as the data mover, you can optionally specify annotations in the CR or using the CLI to control the behavior of the temporary storage used by Kopia. Refer to the [Kopia documentation](#) for more information about the options you can configure. Use the `tridentctl-protect create --help` command for more information about specifying annotations with the Trident Protect CLI.

Use a CR

Steps

1. Create the custom resource (CR) file and name it `trident-protect-backup-restore-cr.yaml`.
2. In the file you created, configure the following attributes:
 - **metadata.name:** (*Required*) The name of this custom resource; choose a unique and sensible name for your environment.
 - **spec.appArchivePath:** The path inside AppVault where the backup contents are stored. You can use the following command to find this path:

```
kubectl get backups <BACKUP_NAME> -n my-app-namespace -o  
jsonpath='{.status.appArchivePath}'
```

- **spec.appVaultRef:** (*Required*) The name of the AppVault where the backup contents are stored.
- **spec.destinationApplicationName:** (*Optional*) The name for the restored application. If provided, the restored application uses this name. If not provided, the restored application uses the source application name.
- **spec.namespaceMapping:** The mapping of the source namespace of the restore operation to the destination namespace. Replace `my-source-namespace` and `my-destination-namespace`

with information from your environment.

```
---
apiVersion: protect.trident.netapp.io/v1
kind: BackupRestore
metadata:
  name: my-cr-name
  namespace: my-destination-namespace
spec:
  appArchivePath: my-backup-path
  appVaultRef: appvault-name
  destinationApplicationName: my-new-app-name
  namespaceMapping: [{"source": "my-source-namespace",
"destination": "my-destination-namespace"}]
```

3. (Optional) If you need to select only certain resources of the application to restore, add filtering that includes or excludes resources marked with particular labels:



Trident Protect selects some resources automatically because of their relationship with resources that you select. For example, if you select a persistent volume claim resource and it has an associated pod, Trident Protect will also restore the associated pod.

- **resourceFilter.resourceSelectionCriteria:** (Required for filtering) Use `Include` or `Exclude` to include or exclude a resource defined in `resourceMatchers`. Add the following `resourceMatchers` parameters to define the resources to be included or excluded:
 - **resourceFilter.resourceMatchers:** An array of `resourceMatcher` objects. If you define multiple elements in this array, they match as an OR operation, and the fields inside each element (group, kind, version) match as an AND operation.
 - **resourceMatchers[].group:** (Optional) Group of the resource to be filtered.
 - **resourceMatchers[].kind:** (Optional) Kind of the resource to be filtered.
 - **resourceMatchers[].version:** (Optional) Version of the resource to be filtered.
 - **resourceMatchers[].names:** (Optional) Names in the Kubernetes metadata.name field of the resource to be filtered.
 - **resourceMatchers[].namespaces:** (Optional) Namespaces in the Kubernetes metadata.name field of the resource to be filtered.
 - **resourceMatchers[].labelSelectors:** (Optional) Label selector string in the Kubernetes metadata.name field of the resource as defined in the [Kubernetes documentation](#). For example: `"trident.netapp.io/os=linux"`.

For example:

```
spec:
  resourceFilter:
    resourceSelectionCriteria: "Include"
```

```
resourceMatchers:
- group: my-resource-group-1
  kind: my-resource-kind-1
  version: my-resource-version-1
  names: ["my-resource-names"]
  namespaces: ["my-resource-namespaces"]
  labelSelectors: ["trident.netapp.io/os=linux"]
- group: my-resource-group-2
  kind: my-resource-kind-2
  version: my-resource-version-2
  names: ["my-resource-names"]
  namespaces: ["my-resource-namespaces"]
  labelSelectors: ["trident.netapp.io/os=linux"]
```

4. After you populate the `trident-protect-backup-restore-cr.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-backup-restore-cr.yaml
```

Use the CLI

Steps

1. Restore the backup to a different namespace, replacing values in brackets with information from your environment. The `namespace-mapping` argument uses colon-separated namespaces to map source namespaces to the correct destination namespaces in the format `source1:dest1,source2:dest2`. For example:

```
tridentctl-protect create backuprestore <my_restore_name> \
--backup <backup_namespace>/<backup_to_restore> \
--namespace-mapping <source_to_destination_namespace_mapping> \
--destination-app-name<custom_app_name>\
-n <application_namespace>
```

Restore from a backup to the original namespace

You can restore a backup to the original namespace at any time. When you perform an in-place restore, Trident Protect automatically manages protection schedules and in-progress operations to prevent invalid recovery points:

- All enabled protection schedules for the application are disabled before the restore begins. This prevents scheduled backups or snapshots from running while the application resources are being restored.
- After the restore completes successfully, only the schedules that were enabled before the restore are re-enabled. Schedules that were already disabled remain disabled.
- Any in-progress backup or snapshot operations are cancelled before the restore begins. If an operation does not cancel within 5 minutes, the restore proceeds and logs a warning in the restore CR status.

Before you begin

Ensure that the AWS session token expiration is sufficient for any long-running s3 restore operations. If the token expires during the restore operation, the operation can fail.

- Refer to the [AWS API documentation](#) for more information about checking the current session token expiration.
- Refer to the [AWS IAM documentation](#) for more information about credentials with AWS resources.



When you restore backups using Kopia as the data mover, you can optionally specify annotations in the CR or using the CLI to control the behavior of the temporary storage used by Kopia. Refer to the [Kopia documentation](#) for more information about the options you can configure. Use the `tridentctl-protect create --help` command for more information about specifying annotations with the Trident Protect CLI.

Use a CR

Steps

1. Create the custom resource (CR) file and name it `trident-protect-backup-ipr-cr.yaml`.
2. In the file you created, configure the following attributes:
 - **metadata.name:** (*Required*) The name of this custom resource; choose a unique and sensible name for your environment.
 - **spec.appArchivePath:** The path inside AppVault where the backup contents are stored. You can use the following command to find this path:

```
kubectl get backups <BACKUP_NAME> -n my-app-namespace -o  
jsonpath='{.status.appArchivePath}'
```

- **spec.appVaultRef:** (*Required*) The name of the AppVault where the backup contents are stored.

For example:

```
---  
apiVersion: protect.trident.netapp.io/v1  
kind: BackupInplaceRestore  
metadata:  
  name: my-cr-name  
  namespace: my-app-namespace  
spec:  
  appArchivePath: my-backup-path  
  appVaultRef: appvault-name
```

3. (*Optional*) If you need to select only certain resources of the application to restore, add filtering that includes or excludes resources marked with particular labels:



Trident Protect selects some resources automatically because of their relationship with resources that you select. For example, if you select a persistent volume claim

resource and it has an associated pod, Trident Protect will also restore the associated pod.

- **resourceFilter.resourceSelectionCriteria:** (Required for filtering) Use `Include` or `Exclude` to include or exclude a resource defined in `resourceMatchers`. Add the following `resourceMatchers` parameters to define the resources to be included or excluded:
 - **resourceFilter.resourceMatchers:** An array of `resourceMatcher` objects. If you define multiple elements in this array, they match as an OR operation, and the fields inside each element (group, kind, version) match as an AND operation.
 - **resourceMatchers[].group:** (*Optional*) Group of the resource to be filtered.
 - **resourceMatchers[].kind:** (*Optional*) Kind of the resource to be filtered.
 - **resourceMatchers[].version:** (*Optional*) Version of the resource to be filtered.
 - **resourceMatchers[].names:** (*Optional*) Names in the Kubernetes metadata.name field of the resource to be filtered.
 - **resourceMatchers[].namespaces:** (*Optional*) Namespaces in the Kubernetes metadata.name field of the resource to be filtered.
 - **resourceMatchers[].labelSelectors:** (*Optional*) Label selector string in the Kubernetes metadata.name field of the resource as defined in the [Kubernetes documentation](#). For example: `"trident.netapp.io/os=linux"`.

For example:

```
spec:
  resourceFilter:
    resourceSelectionCriteria: "Include"
    resourceMatchers:
      - group: my-resource-group-1
        kind: my-resource-kind-1
        version: my-resource-version-1
        names: ["my-resource-names"]
        namespaces: ["my-resource-namespaces"]
        labelSelectors: ["trident.netapp.io/os=linux"]
      - group: my-resource-group-2
        kind: my-resource-kind-2
        version: my-resource-version-2
        names: ["my-resource-names"]
        namespaces: ["my-resource-namespaces"]
        labelSelectors: ["trident.netapp.io/os=linux"]
```

4. After you populate the `trident-protect-backup-ipr-cr.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-backup-ipr-cr.yaml
```

Use the CLI

Steps

1. Restore the backup to the original namespace, replacing values in brackets with information from your environment. The `backup` argument uses a namespace and backup name in the format `<namespace>/<name>`. For example:

```
tridentctl-protect create backupinplacerestore <my_restore_name> \
--backup <namespace/backup_to_restore> \
-n <application_namespace>
```

Restore from a backup to a different cluster

You can restore a backup to a different cluster if there is an issue with the original cluster.



- When you restore backups using Kopia as the data mover, you can optionally specify annotations in the CR or using the CLI to control the behavior of the temporary storage used by Kopia. Refer to the [Kopia documentation](#) for more information about the options you can configure. Use the `tridentctl-protect create --help` command for more information about specifying annotations with the Trident Protect CLI.
- When using a CR to restore to a new namespace, you must manually create the destination namespace before applying the CR. Trident Protect automatically creates namespaces only when using the CLI.

Before you begin

Ensure the following prerequisites are met:

- The destination cluster has Trident Protect installed.
- The destination cluster has access to the bucket path of the same AppVault as the source cluster, where the backup is stored.
- Ensure that your local environment can connect to the object storage bucket defined in the AppVault CR when running the `tridentctl-protect get appvaultcontent` command. If network restrictions prevent access, run the Trident Protect CLI from within a pod on the destination cluster instead.
- Ensure that the AWS session token expiration is sufficient for any long-running restore operations. If the token expires during the restore operation, the operation can fail.
 - Refer to the [AWS API documentation](#) for more information about checking the current session token expiration.
 - Refer to the [AWS documentation](#) for more information about credentials with AWS resources.

Steps

1. Verify that the AppVault CR exists on the destination cluster using the Trident Protect CLI plugin:

```
tridentctl-protect get appvault --context <destination_cluster_name>
```



If the AppVault CR does not exist on the destination cluster, create it by following the steps

in [Use Trident Protect AppVault objects to manage buckets](#).

2. View the backup contents of the available AppVault on the destination cluster, and note `appArchivePath` of the backup you want to restore:

```
tridentctl-protect get appvaultcontent <appvault_name> \
--show-resources backup \
--show-paths \
--context <destination_cluster_name>
```

Running this command displays the available backups in the AppVault, including their originating clusters, corresponding application names, timestamps, and archive paths.

Example output:

```
+-----+-----+-----+-----+
+-----+-----+-----+-----+
|  CLUSTER  |  APP  |  TYPE  |  NAME          |  TIMESTAMP
|  PATH     |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
| production1 | wordpress | backup | wordpress-bkup-1 | 2024-10-30
08:37:40 (UTC) | backuppath1 |
| production1 | wordpress | backup | wordpress-bkup-2 | 2024-10-30
08:37:40 (UTC) | backuppath2 |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
```

3. Restore the application to the destination cluster using the AppVault name and archive path:



When using a CR, ensure that the namespace intended for the application restore exists on the destination cluster.

Use a CR

4. Create the custom resource (CR) file and name it `trident-protect-backup-restore-cr.yaml`.
5. In the file you created, configure the following attributes:
 - **metadata.name:** *(Required)* The name of this custom resource; choose a unique and sensible name for your environment.
 - **spec.appVaultRef:** *(Required)* The name of the AppVault where the backup contents are stored.
 - **spec.appArchivePath:** *(Required)* The path inside AppVault where the backup contents are stored. Use the command from step 2 to view the backup contents and find `appArchivePath` for the backup you want to restore.
 - **spec.destinationApplicationName:** *(Optional)* The name for the restored application. If provided,

the restored application uses this name. If not provided, the restored application uses the source application name.

- **spec.namespaceMapping:** The mapping of the source namespace of the restore operation to the destination namespace. Replace `my-source-namespace` and `my-destination-namespace` with information from your environment.

For example:

```
apiVersion: protect.trident.netapp.io/v1
kind: BackupRestore
metadata:
  name: my-cr-name
  namespace: my-destination-namespace
spec:
  appVaultRef: appvault-name
  appArchivePath: my-backup-path
  destinationApplicationName: my-new-app-name
  namespaceMapping: [{"source": "my-source-namespace",
"destination": "my-destination-namespace"}]
```

6. After you populate the `trident-protect-backup-restore-cr.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-backup-restore-cr.yaml
```

Use the CLI

4. Use the following command to restore the application, replacing values in brackets with information from your environment. The namespace-mapping argument uses colon-separated namespaces to map source namespaces to the correct destination namespaces in the format `source1:dest1,source2:dest2`. For example:

```
tridentctl-protect create backuprestore <restore_name> \
--namespace-mapping <source_to_destination_namespace_mapping> \
--appvault <appvault_name> \
--path <backup_path> \
--destination-app-name <custom_app_name> \
--context <destination_cluster_name> \
-n <application_namespace>
```

Restore from a snapshot to a different namespace

You can restore data from a snapshot using a custom resource (CR) file either to a different namespace or the original source namespace. When you restore a snapshot to a different namespace using a `SnapshotRestore` CR, Trident Protect restores the application in a new namespace and creates an application CR for the

restored application. To protect the restored application, create on-demand backups or snapshots, or establish a protection schedule.



- SnapshotRestore supports the `spec.storageClassMapping` attribute, but only when the source and destination storage classes use the same storage backend. If you attempt to restore to a `StorageClass` that uses a different storage backend, the restore operation will fail.
- When using a CR to restore to a new namespace, you must manually create the destination namespace before applying the CR. Trident Protect automatically creates namespaces only when using the CLI.

Before you begin

Ensure that the AWS session token expiration is sufficient for any long-running s3 restore operations. If the token expires during the restore operation, the operation can fail.

- Refer to the [AWS API documentation](#) for more information about checking the current session token expiration.
- Refer to the [AWS IAM documentation](#) for more information about credentials with AWS resources.

Use a CR

Steps

1. Create the custom resource (CR) file and name it `trident-protect-snapshot-restore-cr.yaml`.
2. In the file you created, configure the following attributes:
 - **metadata.name:** (*Required*) The name of this custom resource; choose a unique and sensible name for your environment.
 - **spec.appVaultRef:** (*Required*) The name of the AppVault where the snapshot contents are stored.
 - **spec.appArchivePath:** The path inside AppVault where the snapshot contents are stored. You can use the following command to find this path:

```
kubectl get snapshots <SNAPSHOT_NAME> -n my-app-namespace -o  
jsonpath='{.status.appArchivePath}'
```

- **spec.destinationApplicationName:** (*Optional*) The name for the restored application. If provided, the restored application uses this name. If not provided, the restored application uses the source application name.
- **spec.namespaceMapping:** The mapping of the source namespace of the restore operation to the destination namespace. Replace `my-source-namespace` and `my-destination-namespace` with information from your environment.

```
---  
apiVersion: protect.trident.netapp.io/v1  
kind: SnapshotRestore  
metadata:
```

```

name: my-cr-name
namespace: my-destination-namespace
spec:
  appVaultRef: appvault-name
  appArchivePath: my-snapshot-path
  namespaceMapping: [{"source": "my-source-namespace",
"destination": "my-destination-namespace"}]

```

3. (Optional) If you need to select only certain resources of the application to restore, add filtering that includes or excludes resources marked with particular labels:



Trident Protect selects some resources automatically because of their relationship with resources that you select. For example, if you select a persistent volume claim resource and it has an associated pod, Trident Protect will also restore the associated pod.

- **resourceFilter.resourceSelectionCriteria:** (Required for filtering) Use `Include` or `Exclude` to include or exclude a resource defined in `resourceMatchers`. Add the following `resourceMatchers` parameters to define the resources to be included or excluded:
 - **resourceFilter.resourceMatchers:** An array of `resourceMatcher` objects. If you define multiple elements in this array, they match as an OR operation, and the fields inside each element (`group`, `kind`, `version`) match as an AND operation.
 - **resourceMatchers[].group:** (Optional) Group of the resource to be filtered.
 - **resourceMatchers[].kind:** (Optional) Kind of the resource to be filtered.
 - **resourceMatchers[].version:** (Optional) Version of the resource to be filtered.
 - **resourceMatchers[].names:** (Optional) Names in the Kubernetes `metadata.name` field of the resource to be filtered.
 - **resourceMatchers[].namespaces:** (Optional) Namespaces in the Kubernetes `metadata.name` field of the resource to be filtered.
 - **resourceMatchers[].labelSelectors:** (Optional) Label selector string in the Kubernetes `metadata.name` field of the resource as defined in the [Kubernetes documentation](#). For example: `"trident.netapp.io/os=linux"`.

For example:

```

spec:
  resourceFilter:
    resourceSelectionCriteria: "Include"
    resourceMatchers:
      - group: my-resource-group-1
        kind: my-resource-kind-1
        version: my-resource-version-1
        names: ["my-resource-names"]
        namespaces: ["my-resource-namespaces"]
        labelSelectors: ["trident.netapp.io/os=linux"]
      - group: my-resource-group-2

```

```
kind: my-resource-kind-2
version: my-resource-version-2
names: ["my-resource-names"]
namespaces: ["my-resource-namespaces"]
labelSelectors: ["trident.netapp.io/os=linux"]
```

4. After you populate the `trident-protect-snapshot-restore-cr.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-snapshot-restore-cr.yaml
```

Use the CLI

Steps

1. Restore the snapshot to a different namespace, replacing values in brackets with information from your environment.
 - The `snapshot` argument uses a namespace and snapshot name in the format `<namespace>/<name>`.
 - The `namespace-mapping` argument uses colon-separated namespaces to map source namespaces to the correct destination namespaces in the format `source1:dest1,source2:dest2`.

For example:

```
tridentctl-protect create snapshotrestore <my_restore_name> \
--snapshot <namespace/snapshot_to_restore> \
--namespace-mapping <source_to_destination_namespace_mapping> \
--destination-app-name <custom_app_name> \
-n <application_namespace>
```

Restore from a snapshot to the original namespace

You can restore a snapshot to the original namespace at any time. When you perform an in-place restore, Trident Protect automatically manages protection schedules and in-progress operations to prevent invalid recovery points:

- All enabled protection schedules for the application are disabled before the restore begins. This prevents scheduled backups or snapshots from running while the application resources are being restored.
- After the restore completes successfully, only the schedules that were enabled before the restore are re-enabled. Schedules that were already disabled remain disabled.
- Any in-progress backup or snapshot operations are cancelled before the restore begins. If an operation does not cancel within 5 minutes, the restore proceeds and logs a warning in the restore CR status.

Before you begin

Ensure that the AWS session token expiration is sufficient for any long-running s3 restore operations. If the token expires during the restore operation, the operation can fail.

- Refer to the [AWS API documentation](#) for more information about checking the current session token expiration.
- Refer to the [AWS IAM documentation](#) for more information about credentials with AWS resources.

Use a CR

Steps

1. Create the custom resource (CR) file and name it `trident-protect-snapshot-ipr-cr.yaml`.
2. In the file you created, configure the following attributes:
 - **metadata.name:** (*Required*) The name of this custom resource; choose a unique and sensible name for your environment.
 - **spec.appVaultRef:** (*Required*) The name of the AppVault where the snapshot contents are stored.
 - **spec.appArchivePath:** The path inside AppVault where the snapshot contents are stored. You can use the following command to find this path:

```
kubectl get snapshots <SNAPSHOT_NAME> -n my-app-namespace -o
jsonpath='{.status.appArchivePath}'
```

+

```
---
apiVersion: protect.trident.netapp.io/v1
kind: SnapshotInplaceRestore
metadata:
  name: my-cr-name
  namespace: my-app-namespace
spec:
  appVaultRef: appvault-name
  appArchivePath: my-snapshot-path
```

+

. (*Optional*) If you need to select only certain resources of the application to restore, add filtering that includes or excludes resources marked with particular labels:

+

NOTE: Trident Protect selects some resources automatically because of their relationship with resources that you select. For example, if you select a persistent volume claim resource and it has an associated pod, Trident Protect will also restore the associated pod.

+

* **resourceFilter.resourceSelectionCriteria:** (*Required for filtering*) Use `Include` or `Exclude` to include or exclude a resource defined in `resourceMatchers`. Add the following `resourceMatchers` parameters to define the resources to be included or excluded:

resourceFilter.resourceMatchers: An array of `resourceMatcher` objects. If you define multiple elements in this array, they match as an OR operation, and the fields inside each element (group,

kind, version) match as an AND operation.

* **resourceMatchers[].group**: (*Optional*) Group of the resource to be filtered.

* **resourceMatchers[].kind**: (*Optional*) Kind of the resource to be filtered.

* **resourceMatchers[].version**: (*Optional*) Version of the resource to be filtered.

* **resourceMatchers[].names**: (*Optional*) Names in the Kubernetes metadata.name field of the resource to be filtered.

* **resourceMatchers[].namespaces**: (*Optional*) Namespaces in the Kubernetes metadata.name field of the resource to be filtered.

* **resourceMatchers[].labelSelectors**: (*Optional*) Label selector string in the Kubernetes metadata.name field of the resource as defined in the [Kubernetes documentation](#). For example:

"trident.netapp.io/os=linux".

+

For example:

+

```
spec:
  resourceFilter:
    resourceSelectionCriteria: "Include"
    resourceMatchers:
      - group: my-resource-group-1
        kind: my-resource-kind-1
        version: my-resource-version-1
        names: ["my-resource-names"]
        namespaces: ["my-resource-namespaces"]
        labelSelectors: ["trident.netapp.io/os=linux"]
      - group: my-resource-group-2
        kind: my-resource-kind-2
        version: my-resource-version-2
        names: ["my-resource-names"]
        namespaces: ["my-resource-namespaces"]
        labelSelectors: ["trident.netapp.io/os=linux"]
```

+

. After you populate the trident-protect-snapshot-ipr-cr.yaml file with the correct values, apply the CR:

+

```
kubectl apply -f trident-protect-snapshot-ipr-cr.yaml
```

Use the CLI

Steps

1. Restore the snapshot to the original namespace, replacing values in brackets with information from your environment. For example:

```
tridentctl-protect create snapshotinplacerestore <my_restore_name> \  
--snapshot <namespace/snapshot_to_restore> \  
-n <application_namespace>
```

Check the status of a restore operation

You can use the command line to check the status of a restore operation that is in progress, has completed, or has failed.

Steps

1. Use the following command to retrieve status of the restore operation, replacing values in brackets with information from your environment:

```
kubectl get backuprestore -n <namespace_name> <my_restore_cr_name> -o  
jsonpath='{.status}'
```

Use advanced Trident Protect restore settings

You can customize restore operations using advanced settings such as annotations, namespace settings, and storage options to meet your specific requirements.

Namespace annotations and labels during restore and failover operations

During restore and failover operations, labels and annotations in the destination namespace are made to match the labels and annotations in the source namespace. Labels or annotations from the source namespace that don't exist in the destination namespace are added, and any labels or annotations that already exist are overwritten to match the value from the source namespace. Labels or annotations that exist only on the destination namespace remain unchanged.



If you use Red Hat OpenShift, it's important to note the critical role of namespace annotations in OpenShift environments. Namespace annotations ensure that restored pods adhere to the appropriate permissions and security configurations defined by OpenShift security context constraints (SCCs) and can access volumes without permission issues. For more information, refer to the [OpenShift security context constraints documentation](#).

You can prevent specific annotations in the destination namespace from being overwritten by setting the Kubernetes environment variable `RESTORE_SKIP_NAMESPACE_ANNOTATIONS` before you perform the restore or failover operation. For example:

```
helm upgrade trident-protect -n trident-protect netapp-trident-  
protect/trident-protect \  
--set-string  
restoreSkipNamespaceAnnotations="{<annotation_key_to_skip_1>,<annotation_k  
ey_to_skip_2>}" \  
--reuse-values
```



When performing restore or failover operation, any namespace annotations and labels specified in `restoreSkipNamespaceAnnotations` and `restoreSkipNamespaceLabels` are excluded from the restore or failover operation. Ensure these settings are configured during the initial Helm installation. To learn more, refer to [Configure additional Trident Protect helm chart settings](#).

If you installed the source application using Helm with the `--create-namespace` flag, special treatment is given to the `name` label key. During the restore or failover process, Trident Protect copies this label to the destination namespace, but updates the value to the destination namespace value if the value from source matches the source namespace. If this value doesn't match the source namespace it is copied to the destination namespace with no changes.

Example

The following example presents a source and destination namespace, each with different annotations and labels. You can see the state of the destination namespace before and after the operation, and how the annotations and labels are combined or overwritten in the destination namespace.

Before the restore or failover operation

The following table illustrates the state of the example source and destination namespaces before the restore or failover operation:

Namespace	Annotations	Labels
Namespace ns-1 (source)	<ul style="list-style-type: none">• <code>annotation.one/key: "updatedvalue"</code>• <code>annotation.two/key: "true"</code>	<ul style="list-style-type: none">• <code>environment=production</code>• <code>compliance=hipaa</code>• <code>name=ns-1</code>
Namespace ns-2 (destination)	<ul style="list-style-type: none">• <code>annotation.one/key: "true"</code>• <code>annotation.three/key: "false"</code>	<ul style="list-style-type: none">• <code>role=database</code>

After the restore operation

The following table illustrates the state of the example destination namespace after the restore or failover operation. Some keys have been added, some have been overwritten, and the `name` label has been updated to match the destination namespace:

Namespace	Annotations	Labels
Namespace ns-2 (destination)	<ul style="list-style-type: none">• <code>annotation.one/key: "updatedvalue"</code>• <code>annotation.two/key: "true"</code>• <code>annotation.three/key: "false"</code>	<ul style="list-style-type: none">• <code>name=ns-2</code>• <code>compliance=hipaa</code>• <code>environment=production</code>• <code>role=database</code>

Supported fields

This section describes additional fields available for restore operations.

Storage class mapping

The `spec.storageClassMapping` attribute defines a mapping from a storage class present in the source application to a new storage class on the target cluster. You can use this when migrating applications between clusters with different storage classes or when changing the storage backend for BackupRestore operations.

Example:

```
storageClassMapping:
- destination: "destinationStorageClass1"
  source: "sourceStorageClass1"
- destination: "destinationStorageClass2"
  source: "sourceStorageClass2"
```

Supported annotations

This section lists the supported annotations for configuring various behaviors in the system. If an annotation is not explicitly set by the user, the system will use the default value.

Annotation	Type	Description	Default value
protect.trident.netapp.io/data-mover-timeout-sec	string	The maximum time (in seconds) allowed for data mover operation to be stalled.	"300"
protect.trident.netapp.io/kopia-content-cache-size-limit-mb	string	The maximum size limit (in megabytes) for the Kopia content cache.	"1000"
protect.trident.netapp.io/pvc-bind-timeout-sec	string	Maximum time (in seconds) to wait for any newly created PersistentVolumeClaims (PVCs) to reach the Bound phase before the operations fails. Applies to all restore CR types (BackupRestore, BackupInplaceRestore, SnapshotRestore, SnapshotInplaceRestore). Use a higher value if your storage backend or cluster often requires more time.	"1200" (20 minutes)

Modify resources as you restore them using custom resources

Resource transformations enable you to modify a resource as it is being restored. This is useful when the restored version needs to differ from the original, for example, updating a hostname in a ConfigMap when restoring an application to a different cluster.

How resource modification works

The `transformations` field in `SnapshotRestore`, `BackupRestore`, `AppMirrorRelationship`, and other restore resources enables you to modify Kubernetes resources during the restore process. This is useful when the restored application needs to be configured differently from the original, for example, changing hostnames, registry URLs, resource limits, or environment variables.

Resource transformations use [RFC 6902 JSON Patch](#) operations and [RFC 6901 JSON Pointer](#) paths to target and modify specific fields within Kubernetes resources.

The following shows the basic structure of a restore object that contains resource transformations:

```
apiVersion: protect.trident.netapp.io/v1
kind: SnapshotRestore
metadata:
  name: my-restore
  namespace: target-namespace
spec:
  appVaultRef: my-vault
  appArchivePath: /path/to/snapshot
  namespaceMapping:
    - source: source-ns
      destination: target-ns
  transformations:
    - resource:
        kind: ConfigMap           # Required: resource kind
        group: ""                # Optional: API group (empty for core
resources)
        version: ""              # Optional: API version
        name: ""                 # Optional: specific resource name
      operations:
        - op: replace             # Operation type
          path: "/data/key"       # JSON Pointer path
          value: "new-value"      # New value (for add/replace/test)
```

Supported resources

You can use resource transformations with resources that match the following criteria:

- **kind** (required): The Kubernetes resource kind (for example, ConfigMap, Deployment, Pod)
- **group** (optional): The API group (for example, apps, route.openshift.io) — omit for core resources
- **version** (optional): The API version (for example, v1, v1beta1)
- **name** (optional): Apply only to a specific resource by name



Modifying PersistentVolumeClaims and Namespaces is currently not supported.

Supported operations

You can use the following operations to modify resources: add, copy, move, remove, replace, and test.

Add a value to a resource

Use the `add` operation to add a new field or value at the specified path. You can add data to objects or arrays.

Use a CR

The following example adds a node selector to a Deployment resource:

```
transformations:
- resource:
  kind: Deployment
  operations:
  - op: add
    path: "/spec/template/spec/nodeSelector"
    value:
      "topology.kubernetes.io/zone": "us-east-1a"
      disktype: "ssd"
```

Use the CLI

Use the following command to run this transformation using the command line:

```
tridentctl-protect create snapshotrestore <restore_name> \
--snapshot <namespace/snapshot-name> \
--namespace-mapping <source-ns>:<dest-ns> \
--transformation
'apps,v1,Deployment:add{"path":"/spec/template/spec/nodeSelector","value":{"topology.kubernetes.io/zone":"us-east-1a","disktype":"ssd"}}'
```

Copy a value within a resource

Use the `copy` operation to copy a value from one path to another within the same resource. The source is not changed.

Use a CR

The following example duplicates a data key for a ConfigMap object:

```
transformations:
- resource:
  kind: ConfigMap
  operations:
  - op: copy
    from: "/data/source-key"
    path: "/data/backup-key"
```

Use the CLI

Use the following command to run this transformation using the command line:

```
tridentctl-protect create snapshotrestore <restore_name> \
  --snapshot <namespace/snapshot-name> \
  --namespace-mapping <source-ns>:<dest-ns> \
  --transformation ',v1,ConfigMap:copy{"from":"/data/source-
key","path":"/data/backup-key"}'
```

Move a value within a resource

Use the `move` operation to move a value from one path to another within the same resource. The source is removed and the value is placed at the destination.

Use a CR

The following example renames a data key for a ConfigMap object:

```
transformations:
- resource:
  kind: ConfigMap
  operations:
  - op: move
    from: "/data/OLD_KEY"
    path: "/data/NEW_KEY"
```

Use the CLI

Use the following command to run this transformation using the command line:

```
tridentctl-protect create snapshotrestore <restore_name> \
  --snapshot <namespace/snapshot-name> \
  --namespace-mapping <source-ns>:<dest-ns> \
  --transformation
',v1,ConfigMap:move{"from":"/data/OLD_KEY","path":"/data/NEW_KEY"}'
```

Remove a value from a resource

Use the `remove` operation to remove a field or value at the specified path.

Use a CR

The following example removes an annotation from a ConfigMap resource:

```
transformations:
- resource:
```

```

    kind: ConfigMap
  operations:
    - op: remove
      path: "/metadata/annotations/kubectl.kubernetes.io~1last-
        applied-configuration"

```



In the path for the above example, ~1 is the JSON Pointer escape sequence for /.

Use the CLI

Use the following command to run this transformation using the command line:

```

tridentctl-protect create snapshotrestore <restore_name> \
  --snapshot <namespace/snapshot-name> \
  --namespace-mapping <source-ns>:<dest-ns> \
  --transformation
  ',v1,ConfigMap:remove{"path":"/metadata/annotations/kubectl.kubernetes.
  io~1last-applied-configuration"}'

```

Replace a value within a resource

Use the `replace` operation to replace an existing value within a resource at the specified path. The JSON path must already exist.

Use a CR

The following example changes a hostname for a `Route` object:

```

transformations:
  - resource:
      kind: Route
      group: route.openshift.io
    operations:
      - op: replace
        path: "/spec/host"
        value: "prod.example.com"

```

Use the CLI

Use the following command to run this transformation using the command line:

```

tridentctl-protect create snapshotrestore <restore_name> \
  --snapshot <namespace/snapshot-name> \
  --namespace-mapping <source-ns>:<dest-ns> \
  --transformation

```

```
'route.openshift.io,v1,Route:replace{"path":"/spec/host","value":"prod.example.com"}'
```

Test the resource modification

Use the `test` operation to verify that the value at a path matches an expected value before the other operations are applied. If the test fails, the entire modification is rolled back.

Use a CR

The following example updates `database-host` only if environment is `set` to `staging`:

```
transformations:
- resource:
  kind: ConfigMap
  operations:
  - op: test
    path: "/data/environment"
    value: "staging"
  - op: replace
    path: "/data/database-host"
    value: "prod-db.example.com"
```

Use the CLI

Use the following command to run this transformation using the command line:

```
tridentctl-protect create snapshotrestore <restore_name> \
--snapshot <namespace/snapshot-name> \
--namespace-mapping <source-ns>:<dest-ns> \
--transformation
',v1,ConfigMap:test{"path":"/data/environment","value":"staging"},repla
ce{"path":"/data/database-host","value":"prod-db.example.com"}'
```

Replicate applications using NetApp SnapMirror and Trident Protect

Using Trident Protect, you can use the asynchronous replication capabilities of NetApp SnapMirror technology to replicate data and application changes from one storage backend to another, on the same cluster or between different clusters.

Namespace annotations and labels during restore and failover operations

During restore and failover operations, labels and annotations in the destination namespace are made to match the labels and annotations in the source namespace. Labels or annotations from the source namespace that don't exist in the destination namespace are added, and any labels or annotations that already exist are overwritten to match the value from the source namespace. Labels or annotations that exist only on the destination namespace remain unchanged.



If you use Red Hat OpenShift, it's important to note the critical role of namespace annotations in OpenShift environments. Namespace annotations ensure that restored pods adhere to the appropriate permissions and security configurations defined by OpenShift security context constraints (SCCs) and can access volumes without permission issues. For more information, refer to the [OpenShift security context constraints documentation](#).

You can prevent specific annotations in the destination namespace from being overwritten by setting the Kubernetes environment variable `RESTORE_SKIP_NAMESPACE_ANNOTATIONS` before you perform the restore or failover operation. For example:

```
helm upgrade trident-protect -n trident-protect netapp-trident-protect/trident-protect \
  --set-string
  restoreSkipNamespaceAnnotations="{<annotation_key_to_skip_1>,<annotation_key_to_skip_2>}" \
  --reuse-values
```



When performing restore or failover operation, any namespace annotations and labels specified in `restoreSkipNamespaceAnnotations` and `restoreSkipNamespaceLabels` are excluded from the restore or failover operation. Ensure these settings are configured during the initial Helm installation. To learn more, refer to [Configure additional Trident Protect helm chart settings](#).

If you installed the source application using Helm with the `--create-namespace` flag, special treatment is given to the `name` label key. During the restore or failover process, Trident Protect copies this label to the destination namespace, but updates the value to the destination namespace value if the value from source matches the source namespace. If this value doesn't match the source namespace it is copied to the destination namespace with no changes.

Example

The following example presents a source and destination namespace, each with different annotations and labels. You can see the state of the destination namespace before and after the operation, and how the annotations and labels are combined or overwritten in the destination namespace.

Before the restore or failover operation

The following table illustrates the state of the example source and destination namespaces before the restore or failover operation:

Namespace	Annotations	Labels
Namespace ns-1 (source)	<ul style="list-style-type: none"> • annotation.one/key: "updatedvalue" • annotation.two/key: "true" 	<ul style="list-style-type: none"> • environment=production • compliance=hipaa • name=ns-1
Namespace ns-2 (destination)	<ul style="list-style-type: none"> • annotation.one/key: "true" • annotation.three/key: "false" 	<ul style="list-style-type: none"> • role=database

After the restore operation

The following table illustrates the state of the example destination namespace after the restore or failover operation. Some keys have been added, some have been overwritten, and the `name` label has been updated to match the destination namespace:

Namespace	Annotations	Labels
Namespace ns-2 (destination)	<ul style="list-style-type: none"> • annotation.one/key: "updatedvalue" • annotation.two/key: "true" • annotation.three/key: "false" 	<ul style="list-style-type: none"> • name=ns-2 • compliance=hipaa • environment=production • role=database



You can configure Trident Protect to freeze and unfreeze filesystems during data protection operations. [Learn more about configuring filesystem freezing with Trident Protect.](#)

Execution hooks during failover and reverse operations

When using AppMirror relationship to protect your application, there are specific behaviors related to execution hooks that you should be aware of during failover and reverse operations.

- During failover, the execution hooks are automatically copied from the source cluster to the destination cluster. You do not need to manually recreate them. After failover, execution hooks are present on the application and will execute during any relevant actions.
- During reverse or reverse resync, any existing execution hooks on the application are removed. When the source application becomes the destination application, these execution hooks are not valid and are deleted to prevent their execution.

To learn more about execution hooks, refer to [Manage Trident Protect execution hooks](#).

Set up a replication relationship

Setting up a replication relationship involves the following:

- Choosing how frequently you want Trident Protect to take an app snapshot (which includes the app's Kubernetes resources as well as the volume snapshots for each of the app's volumes)
- Choosing the replication schedule (includes Kubernetes resources as well as persistent volume data)
- Setting the time for the snapshot to be taken

Steps

1. On the source cluster, create an AppVault for the source application. Depending on your storage provider, modify an example in [AppVault custom resources](#) to fit your environment:

Create an AppVault using a CR

- a. Create the custom resource (CR) file and name it (for example, `trident-protect-appvault-primary-source.yaml`).
- b. Configure the following attributes:
 - **metadata.name:** (*Required*) The name of the AppVault custom resource. Make note of the name you choose, because other CR files needed for a replication relationship refer to this value.
 - **spec.providerConfig:** (*Required*) Stores the configuration necessary to access the AppVault using the specified provider. Choose a `bucketName` and any other necessary details for your provider. Make note of the values you choose, because other CR files needed for a replication relationship refer to these values. Refer to [AppVault custom resources](#) for examples of AppVault CRs with other providers.
 - **spec.providerCredentials:** (*Required*) Stores references to any credential required to access the AppVault using the specified provider.
 - **spec.providerCredentials.valueFromSecret:** (*Required*) Indicates that the credential value should come from a secret.
 - **key:** (*Required*) The valid key of the secret to select from.
 - **name:** (*Required*) Name of the secret containing the value for this field. Must be in the same namespace.
 - **spec.providerCredentials.secretAccessKey:** (*Required*) The access key used to access the provider. The **name** should match **spec.providerCredentials.valueFromSecret.name**.
 - **spec.providerType:** (*Required*) Determines what provides the backup; for example, NetApp ONTAP S3, generic S3, Google Cloud, or Microsoft Azure. Possible values:
 - `aws`
 - `azure`
 - `gcp`
 - `generic-s3`
 - `ontap-s3`
 - `storagegrid-s3`
- c. After you populate the `trident-protect-appvault-primary-source.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-appvault-primary-source.yaml -n
trident-protect
```

Create an AppVault using the CLI

- a. Create the AppVault, replacing values in brackets with information from your environment:

```
tridentctl-protect create vault Azure <vault-name> --account
<account-name> --bucket <bucket-name> --secret <secret-name> -n
trident-protect
```

2. On the source cluster, create the source application CR:

Create the source application using a CR

- a. Create the custom resource (CR) file and name it (for example, `trident-protect-app-source.yaml`).
- b. Configure the following attributes:
 - **metadata.name:** *(Required)* The name of the application custom resource. Make note of the name you choose, because other CR files needed for a replication relationship refer to this value.
 - **spec.includedNamespaces:** *(Required)* An array of namespaces and associated labels. Use namespace names and optionally narrow the scope of the namespaces with labels to specify resources that exist in the namespaces listed here. The application namespace must be part of this array.

Example YAML:

```
---
apiVersion: protect.trident.netapp.io/v1
kind: Application
metadata:
  name: my-app-name
  namespace: my-app-namespace
spec:
  includedNamespaces:
    - namespace: my-app-namespace
      labelSelector: {}
```

- c. After you populate the `trident-protect-app-source.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-app-source.yaml -n my-app-
namespace
```

Create the source application using the CLI

- a. Create the source application. For example:

```
tridentctl-protect create app <my-app-name> --namespaces
```

```
<namespaces-to-be-included> -n <my-app-namespace>
```

3. Optionally, on the source cluster, take a snapshot of the source application. This snapshot is used as the basis for the application on the destination cluster. If you skip this step, you'll need to wait for the next scheduled snapshot to run so that you have a recent snapshot. To create an on-demand snapshot, refer to [Create an on-demand snapshot](#).
4. On the source cluster, create the replication schedule CR:

Alongside the schedule provided below, it is recommended to create a separate daily snapshot schedule with a retention period of 7 days to maintain a common snapshot between peered ONTAP clusters. This ensures that snapshots are available for up to 7 days, but the retention period can be customized based on user requirements.



If a failover happens, the system can use these snapshots for up to 7 days for reverse operations. This approach makes the reverse process faster and more efficient because only the changes made since the last snapshot will be transferred, not all the data.

If an existing schedule for the application already meets the desired retention requirements, no additional schedules are required.

Create the replication schedule using a CR

a. Create a replication schedule for the source application:

- i. Create the custom resource (CR) file and name it (for example, `trident-protect-schedule.yaml`).
- ii. Configure the following attributes:
 - **metadata.name:** (*Required*) The name of the schedule custom resource.
 - **spec.appVaultRef:** (*Required*) This value must match the `metadata.name` field of the AppVault for the source application.
 - **spec.applicationRef:** (*Required*) This value must match the `metadata.name` field of the source application CR.
 - **spec.backupRetention:** (*Required*) This field is required, and the value must be set to 0.
 - **spec.enabled:** Must be set to `true`.
 - **spec.granularity:** Must be set to `Custom`.
 - **spec.recurrenceRule:** Define a start date in UTC time and a recurrence interval.
 - **spec.snapshotRetention:** Must be set to 2.

Example YAML:

```
---
apiVersion: protect.trident.netapp.io/v1
kind: Schedule
metadata:
  name: appmirror-schedule
  namespace: my-app-namespace
```

```
spec:
  appVaultRef: my-appvault-name
  applicationRef: my-app-name
  backupRetention: "0"
  enabled: true
  granularity: Custom
  recurrenceRule: |-
    DTSTART:20220101T000200Z
    RRULE:FREQ=MINUTELY;INTERVAL=5
  snapshotRetention: "2"
```

- iii. After you populate the `trident-protect-schedule.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-schedule.yaml -n my-app-namespace
```

Create the replication schedule using the CLI

- a. Create the replication schedule, replacing values in brackets with information from your environment:

```
tridentctl-protect create schedule --name appmirror-schedule
--app <my_app_name> --appvault <my_app_vault> --granularity
Custom --recurrence-rule <rule> --snapshot-retention
<snapshot_retention_count> -n <my_app_namespace>
```

Example:

```
tridentctl-protect create schedule --name appmirror-schedule
--app <my_app_name> --appvault <my_app_vault> --granularity
Custom --recurrence-rule "DTSTART:20220101T000200Z
\nRRULE:FREQ=MINUTELY;INTERVAL=5" --snapshot-retention 2 -n
<my_app_namespace>
```

5. On the destination cluster, create a source application AppVault CR that is identical to the AppVault CR you applied on the source cluster and name it (for example, `trident-protect-appvault-primary-destination.yaml`).
6. Apply the CR:

```
kubectl apply -f trident-protect-appvault-primary-destination.yaml -n
trident-protect
```

7. Create a destination AppVault CR for the destination application on the destination cluster. Depending on your storage provider, modify an example in [AppVault custom resources](#) to fit your environment:
 - a. Create the custom resource (CR) file and name it (for example, `trident-protect-appvault-secondary-destination.yaml`).
 - b. Configure the following attributes:
 - **metadata.name:** (*Required*) The name of the AppVault custom resource. Make note of the name you choose, because other CR files needed for a replication relationship refer to this value.
 - **spec.providerConfig:** (*Required*) Stores the configuration necessary to access the AppVault using the specified provider. Choose a `bucketName` and any other necessary details for your provider. Make note of the values you choose, because other CR files needed for a replication relationship refer to these values. Refer to [AppVault custom resources](#) for examples of AppVault CRs with other providers.
 - **spec.providerCredentials:** (*Required*) Stores references to any credential required to access the AppVault using the specified provider.
 - **spec.providerCredentials.valueFromSecret:** (*Required*) Indicates that the credential value should come from a secret.
 - **key:** (*Required*) The valid key of the secret to select from.
 - **name:** (*Required*) Name of the secret containing the value for this field. Must be in the same namespace.
 - **spec.providerCredentials.secretAccessKey:** (*Required*) The access key used to access the provider. The **name** should match **spec.providerCredentials.valueFromSecret.name**.
 - **spec.providerType:** (*Required*) Determines what provides the backup; for example, NetApp ONTAP S3, generic S3, Google Cloud, or Microsoft Azure. Possible values:
 - `aws`
 - `azure`
 - `gcp`
 - `generic-s3`
 - `ontap-s3`
 - `storagegrid-s3`
 - c. After you populate the `trident-protect-appvault-secondary-destination.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-appvault-secondary-destination.yaml
-n trident-protect
```

8. On the destination cluster, create an AppMirrorRelationship CR file.



When using a CR, manually create the destination namespace before applying the CR. Trident Protect automatically creates namespaces only when using the CLI.

Create an AppMirrorRelationship using a CR

- a. Create the custom resource (CR) file and name it (for example, `trident-protect-`

relationship.yaml).

b. Configure the following attributes:

- **metadata.name:** (Required) The name of the AppMirrorRelationship custom resource.
- **spec.destinationAppVaultRef:** (Required) This value must match the name of the AppVault for the destination application on the destination cluster.
- **spec.namespaceMapping:** (Required) The destination and source namespaces must match the application namespace defined in the respective application CR.
- **spec.sourceAppVaultRef:** (Required) This value must match the name of the AppVault for the source application.
- **spec.sourceApplicationName:** (Required) This value must match the name of the source application you defined in the source application CR.
- **spec.sourceApplicationUID:** (Required) This value must match the UID of the source application you defined in the source application CR.
- **spec.storageClassName:** (Optional) Choose the name of a valid storage class on the cluster. The storage class must be linked to an ONTAP storage VM that is peered with the source environment. If the storage class is not provided, the default storage class on the cluster will be used by default.
- **spec.storageClassMapping:** (Optional) Maps source storage classes to destination storage classes for individual PVCs within the same AMR. Each entry requires a `source` and `destination` storage class name. When set, `storageClassMapping` overrides `storageClassName` entirely; PVCs without a matching source entry use the cluster default storage class, not `storageClassName`.

When using `storageClassMapping` in an AppMirrorRelationship CR, be aware of the following restrictions:

- `storageClassMapping` overrides `storageClassName` entirely. When both fields are set, `storageClassName` is ignored.
 - PVCs with source storage classes that have no matching entry in `storageClassMapping` use the cluster default storage class, not `storageClassName`.
 - Each entry requires both a `source` and `destination` value.
 - Changing the destination storage class can force deletion and recreation of existing destination PVCs.
- **spec.recurrenceRule:** Define a start date in UTC time and a recurrence interval.

Example YAML:

```
---
apiVersion: protect.trident.netapp.io/v1
kind: AppMirrorRelationship
metadata:
  name: amr-16061e80-1b05-4e80-9d26-d326dc1953d8
  namespace: my-app-namespace
spec:
  desiredState: Established
```

```

destinationAppVaultRef: generic-s3-trident-protect-dst-
bucket-8fe0b902-f369-4317-93d1-ad7f2edc02b5
namespaceMapping:
  - destination: my-app-namespace
    source: my-app-namespace
recurrenceRule: |-
  DTSTART:20220101T000200Z
  RRULE:FREQ=MINUTELY;INTERVAL=5
sourceAppVaultRef: generic-s3-trident-protect-src-bucket-
b643cc50-0429-4ad5-971f-ac4a83621922
sourceApplicationName: my-app-name
sourceApplicationUID: 7498d32c-328e-4ddd-9029-122540866aeb
storageClassName: sc-vsrm-2

```

- c. After you populate the `trident-protect-relationship.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-relationship.yaml -n my-app-namespace
```

Create an AppMirrorRelationship using the CLI

- a. Create and apply the AppMirrorRelationship object, replacing values in brackets with information from your environment:

```

tridentctl-protect create appmirrorrelationship
<name_of_appmirrorrelationship> --destination-app-vault
<my_vault_name> --source-app-vault <my_vault_name> --recurrence
-rule <rule> --namespace-mapping <ns_mapping> --source-app-id
<source_app_UID> --source-app <my_source_app_name> --storage
-class <storage_class_name> --storageclass-mapping
<src_class:dst_class> -n <application_namespace>

```



Use `--storageclass-mapping` to map individual source storage classes to destination storage classes. Specify one or more `source:destination` pairs separated by commas (for example, `"class-a:class-b,class-c:class-d"`). Use `*:destination-class` as a wildcard to map all unmapped PVCs to a single destination storage class. When `--storageclass-mapping` is set, `--storage-class` is ignored.

Example using a single destination storage class (`--storage-class`):

```

tridentctl-protect create appmirrorrelationship my-amr
--destination-app-vault appvault2 --source-app-vault appvault1
--recurrence-rule

```

```
"DTSTART:20220101T000200Z\nRRULE:FREQ=MINUTELY;INTERVAL=5"
--source-app my-app --namespace-mapping "my-source-ns1:my-dest-
ns1,my-source-ns2:my-dest-ns2" --source-app-id 373f24c1-5769-
404c-93c3-5538af6ccc36 --storage-class my-storage-class -n my-
dest-ns1
```

9. (Optional) On the destination cluster, check the state and status of the replication relationship:

```
kubectl get amr -n my-app-namespace <relationship name> -o=jsonpath
='{.status}' | jq
```

Fail over to destination cluster

Using Trident Protect, you can fail over replicated applications to a destination cluster. This procedure stops the replication relationship and brings the app online on the destination cluster. Trident Protect does not stop the app on the source cluster if it was operational.

Steps

1. On the destination cluster, edit the AppMirrorRelationship CR file (for example, `trident-protect-relationship.yaml`) and change the value of **spec.desiredState** to `Promoted`.
2. Save the CR file.
3. Apply the CR:

```
kubectl apply -f trident-protect-relationship.yaml -n my-app-namespace
```

4. (Optional) Create any protection schedules that you need on the failed over application.
5. (Optional) Check the state and status of the replication relationship:

```
kubectl get amr -n my-app-namespace <relationship name> -o=jsonpath
='{.status}' | jq
```

Resync a failed over replication relationship

The resync operation re-establishes the replication relationship. After you perform a resync operation, the original source application becomes the running application, and any changes made to the running application on the destination cluster are discarded.

The process stops the app on the destination cluster before re-establishing replication.



Any data written to the destination application during failover will be lost.

Steps

1. Optional: On the source cluster, create a snapshot of the source application. This ensures that the latest changes from the source cluster are captured.

2. On the destination cluster, edit the AppMirrorRelationship CR file (for example, `trident-protect-relationship.yaml`) and change the value of `spec.desiredState` to `Established`.
3. Save the CR file.
4. Apply the CR:

```
kubectl apply -f trident-protect-relationship.yaml -n my-app-namespace
```

5. If you created any protection schedules on the destination cluster to protect the failed over application, remove them. Any schedules that remain cause volume snapshot failures.

Reverse resync a failed over replication relationship

When you reverse resync a failed over replication relationship, the destination application becomes the source application, and the source becomes the destination. Changes made to the destination application during failover are kept.

Steps

1. On the original destination cluster, delete the AppMirrorRelationship CR. This causes the destination to become the source. If there are any protection schedules remaining on the new destination cluster, remove them.
2. Set up a replication relationship by applying the CR files you originally used to set up the relationship to the opposite clusters.
3. Ensure the new destination (original source cluster) is configured with both AppVault CRs.
4. Set up a replication relationship on the opposite cluster, configuring values for the reverse direction.

Reverse application replication direction

When you reverse replication direction, Trident Protect moves the application to the destination storage backend while continuing to replicate back to the original source storage backend. Trident Protect stops the source application and replicates the data to the destination before failing over to the destination app.

In this situation, you are swapping the source and destination.

Steps

1. On the source cluster, create a shutdown snapshot:

Create a shutdown snapshot using a CR

- a. Disable the protection policy schedules for the source application.
- b. Create a ShutdownSnapshot CR file:
 - i. Create the custom resource (CR) file and name it (for example, `trident-protect-shutdownsnapshot.yaml`).
 - ii. Configure the following attributes:
 - **metadata.name:** *(Required)* The name of the custom resource.
 - **spec.AppVaultRef:** *(Required)* This value must match the `metadata.name` field of the AppVault for the source application.

- **spec.ApplicationRef:** *(Required)* This value must match the metadata.name field of the source application CR file.

Example YAML:

```
---
apiVersion: protect.trident.netapp.io/v1
kind: ShutdownSnapshot
metadata:
  name: replication-shutdown-snapshot-afc4c564-e700-4b72-86c3-c08a5dbe844e
  namespace: my-app-namespace
spec:
  appVaultRef: generic-s3-trident-protect-src-bucket-04b6b4ec-46a3-420a-b351-45795e1b5e34
  applicationRef: my-app-name
```

- After you populate the `trident-protect-shutdownsnapshot.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-shutdownsnapshot.yaml -n my-app-namespace
```

Create a shutdown snapshot using the CLI

- Create the shutdown snapshot, replacing values in brackets with information from your environment. For example:

```
tridentctl-protect create shutdownsnapshot <my_shutdown_snapshot>
--appvault <my_vault> --app <app_to_snapshot> -n
<application_namespace>
```

- On the source cluster, after the shutdown snapshot completes, get the status of the shutdown snapshot:

```
kubectl get shutdownsnapshot -n my-app-namespace
<shutdown_snapshot_name> -o yaml
```

- On the source cluster, find the value of **shutdownsnapshot.status.appArchivePath** using the following command, and record the last part of the file path (also called the **basename**; this will be everything after the last slash):

```
k get shutdownsnapshot -n my-app-namespace <shutdown_snapshot_name> -o
jsonpath='{.status.appArchivePath}'
```

4. Perform a fail over from the new destination cluster to the new source cluster, with the following change:



In step 2 of the fail over procedure, include the `spec.promotedSnapshot` field in the AppMirrorRelationship CR file, and set its value to the basename you recorded in step 3 above.

5. Perform the reverse resync steps in [Reverse resync a failed over replication relationship](#).

6. Enable protection schedules on the new source cluster.

Result

The following actions occur because of the reverse replication:

- A snapshot is taken of the original source app's Kubernetes resources.
- The original source app's pods are gracefully stopped by deleting the app's Kubernetes resources (leaving PVCs and PVs in place).
- After the pods are shut down, snapshots of the app's volumes are taken and replicated.
- The SnapMirror relationships are broken, making the destination volumes ready for read/write.
- The app's Kubernetes resources are restored from the pre-shutdown snapshot, using the volume data replicated after the original source app was shut down.
- Replication is re-established in the reverse direction.

Fail back applications to the original source cluster

Using Trident Protect, you can achieve "fail back" after a failover operation by using the following sequence of operations. In this workflow to restore the original replication direction, Trident Protect replicates (resyncs) any application changes back to the original source application before reversing the replication direction.

This process starts from a relationship that has completed a failover to a destination and involves the following steps:

- Start with a failed over state.
- Reverse resync the replication relationship.



Do not perform a normal resync operation, as this will discard data written to the destination cluster during the fail over procedure.

- Reverse the replication direction.

Steps

1. Perform the [Reverse resync a failed over replication relationship](#) steps.
2. Perform the [Reverse application replication direction](#) steps.

Delete a replication relationship

You can delete a replication relationship at any time. When you delete the application replication relationship, it results in two separate applications with no relationship between them.

Steps

1. On the current destination cluster, delete the AppMirrorRelationship CR:

```
kubectl delete -f trident-protect-relationship.yaml -n my-app-namespace
```

Migrate applications using Trident Protect

You can migrate your applications between clusters or to different storage classes by restoring backup data.



When you migrate an application, all execution hooks configured for the application are migrated with the app. If a post-restore execution hook is present, it runs automatically as part of the restore operation.

Backup and restore operations

To perform backup and restore operations for the following scenarios, you can automate specific backup and restore tasks.

Clone to same cluster

To clone an application to the same cluster, create a snapshot or backup and restore the data to the same cluster.

Steps

1. Do one of the following:
 - a. [Create a snapshot](#).
 - b. [Create a backup](#).
2. On the same cluster, do one of the following, depending on if you created a snapshot or a backup:
 - a. [Restore your data from the snapshot](#).
 - b. [Restore your data from the backup](#).

Clone to different cluster

To clone an application to a different cluster (perform a cross-cluster clone), create a backup on the source cluster, and then restore the backup to a different cluster. Make sure that Trident Protect is installed on the destination cluster.



You can replicate an application between different clusters using [SnapMirror replication](#).

Steps

1. [Create a backup](#).
2. Ensure that the AppVault CR for the object storage bucket that contains the backup has been configured on the destination cluster.
3. On the destination cluster, [restore your data from the backup](#).

Migrate applications from one storage class to another storage class

You can migrate applications from one storage class to a different storage class by restoring a backup to the destination storage class.

For example (excluding the secrets from the restore CR):

```
apiVersion: protect.trident.netapp.io/v1
kind: SnapshotRestore
metadata:
  name: "${snapshotRestoreCRName}"
spec:
  appArchivePath: "${snapshotArchivePath}"
  appVaultRef: "${appVaultCRName}"
  namespaceMapping:
    - destination: "${destinationNamespace}"
      source: "${sourceNamespace}"
  storageClassMapping:
    - destination: "${destinationStorageClass}"
      source: "${sourceStorageClass}"
  resourceFilter:
    resourceMatchers:
      kind: Secret
      version: v1
    resourceSelectionCriteria: exclude
```

Restore the snapshot using a CR

Steps

1. Create the custom resource (CR) file and name it `trident-protect-snapshot-restore-cr.yaml`.
2. In the file you created, configure the following attributes:
 - **metadata.name:** *(Required)* The name of this custom resource; choose a unique and sensible name for your environment.
 - **spec.appArchivePath:** The path inside AppVault where the snapshot contents are stored. You can use the following command to find this path:

```
kubectl get snapshots <my-snapshot-name> -n trident-protect -o
jsonpath='{.status.appArchivePath}'
```

- **spec.appVaultRef:** *(Required)* The name of the AppVault where the snapshot contents are stored.
- **spec.namespaceMapping:** The mapping of the source namespace of the restore operation to the destination namespace. Replace `my-source-namespace` and `my-destination-namespace` with information from your environment.

```

---
apiVersion: protect.trident.netapp.io/v1
kind: SnapshotRestore
metadata:
  name: my-cr-name
  namespace: trident-protect
spec:
  appArchivePath: my-snapshot-path
  appVaultRef: appvault-name
  namespaceMapping: [{"source": "my-source-namespace",
"destination": "my-destination-namespace"}]

```

3. Optionally, if you need to select only certain resources of the application to restore, add filtering that includes or excludes resources marked with particular labels:

- **resourceFilter.resourceSelectionCriteria:** (Required for filtering) Use `include` or `exclude` to include or exclude a resource defined in `resourceMatchers`. Add the following `resourceMatchers` parameters to define the resources to be included or excluded:
 - **resourceFilter.resourceMatchers:** An array of `resourceMatcher` objects. If you define multiple elements in this array, they match as an OR operation, and the fields inside each element (group, kind, version) match as an AND operation.
 - **resourceMatchers[].group:** (*Optional*) Group of the resource to be filtered.
 - **resourceMatchers[].kind:** (*Optional*) Kind of the resource to be filtered.
 - **resourceMatchers[].version:** (*Optional*) Version of the resource to be filtered.
 - **resourceMatchers[].names:** (*Optional*) Names in the Kubernetes metadata.name field of the resource to be filtered.
 - **resourceMatchers[].namespaces:** (*Optional*) Namespaces in the Kubernetes metadata.name field of the resource to be filtered.
 - **resourceMatchers[].labelSelectors:** (*Optional*) Label selector string in the Kubernetes metadata.name field of the resource as defined in the [Kubernetes documentation](#). For example: `"trident.netapp.io/os=linux"`.

For example:

```

spec:
  resourceFilter:
    resourceSelectionCriteria: "include"
    resourceMatchers:
      - group: my-resource-group-1
        kind: my-resource-kind-1
        version: my-resource-version-1
        names: ["my-resource-names"]
        namespaces: ["my-resource-namespaces"]
        labelSelectors: ["trident.netapp.io/os=linux"]
      - group: my-resource-group-2

```

```
kind: my-resource-kind-2
version: my-resource-version-2
names: ["my-resource-names"]
namespaces: ["my-resource-namespaces"]
labelSelectors: ["trident.netapp.io/os=linux"]
```

4. After you populate the `trident-protect-snapshot-restore-cr.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-snapshot-restore-cr.yaml
```

Restore the snapshot using the CLI

Steps

1. Restore the snapshot to a different namespace, replacing values in brackets with information from your environment.
 - The `snapshot` argument uses a namespace and snapshot name in the format `<namespace>/<name>`.
 - The `namespace-mapping` argument uses colon-separated namespaces to map source namespaces to the correct destination namespaces in the format `source1:dest1,source2:dest2`.

For example:

```
tridentctl-protect create snapshotrestore <my_restore_name>
--snapshot <namespace/snapshot_to_restore> --namespace-mapping
<source_to_destination_namespace_mapping>
```

Protect KubeVirt virtual machines with Trident Protect

Trident Protect enables you to back up and restore KubeVirt virtual machines (VMs) running on OpenShift Virtualization. You can protect all VMs across one or more namespaces at once, or use the `includedVirtualMachines` field to target specific VMs by name. Trident Protect automatically collects all dependent VM resources during backup, and during restore, you can recover the entire archive or selectively restore individual VMs.

KubeVirt protection modes

Trident Protect supports two ways to define which KubeVirt VMs to protect. Choose one mode when you define an application:

Mode	Field	Description
Namespace-based	<code>includedNamespaces</code>	Protect all matching resources in one or more namespaces. Use a <code>labelSelector</code> on the namespace entry to control which VMs are included dynamically.
VM-scoped	<code>includedVirtualMachines</code>	Protect specific KubeVirt VMs by namespace and name. Use this mode when VMs are spread across namespaces or when you do not want to protect all namespace resources.

You must specify exactly one field (`includedNamespaces` or `includedVirtualMachines`) in the `ApplicationSpec`. You cannot use both in the same application.

Both modes support additional filtering:

- Use `resourceFilter` to include or exclude specific resources.
- Use `includedClusterScopedResources` to include compatible cluster-scoped resources, such as `StorageClass` objects.



Trident Protect can freeze and unfreeze KubeVirt VM filesystems during data protection operations to ensure consistency. For configuration details, see [Protecting data with KubeVirt VMs](#).

Resources collected during backup and snapshot

When you create a backup or snapshot for an application that includes KubeVirt VMs, Trident Protect collects the selected VM resources and automatically discovers their dependent resources.

The collection behavior applies in two cases:

- VM-scoped applications that use `includedVirtualMachines`.
- Namespace-based applications that contain KubeVirt `VirtualMachine` resources.

VM-scoped collection

For VM-scoped applications, Trident Protect starts with each `VirtualMachine` listed in `includedVirtualMachines` and discovers the dependent resources needed to restore that VM.

Resource type	Scope	How it is discovered
<code>VirtualMachine</code>	Namespaced	Direct match from the <code>includedVirtualMachines</code> list.
<code>VirtualMachineInstance</code>	Namespaced	Same namespace and name as the VM, if the VM is running.
<code>VirtualMachineInstanceType</code>	Namespaced	Referenced by <code>specinstancetype</code> when <code>kind: VirtualMachineInstanceType</code> .
<code>VirtualMachineClusterInstanceType</code>	Cluster	Referenced by <code>specinstancetype</code> when <code>kind</code> is unset or <code>kind: VirtualMachineClusterInstanceType</code> .

Resource type	Scope	How it is discovered
VirtualMachinePreference	Namespaced	Referenced by <code>spec.preference</code> when <code>kind: VirtualMachinePreference</code> .
VirtualMachineClusterPreference	Cluster	Referenced by <code>spec.preference</code> when <code>kind</code> is unset or <code>kind: VirtualMachineClusterPreference</code> .
DataVolume	Namespaced	Referenced by <code>spec.template.spec.volumes[].dataVolume.name</code> .
PersistentVolumeClaim	Namespaced	Discovered from DataVolume-backing PVCs, which use the same name as the DataVolume; direct PVC references in <code>persistentVolumeClaim.claimName</code> ; memory dump PVCs in <code>memoryDump.claimName</code> ; and ephemeral PVCs in <code>ephemeral.persistentVolumeClaim.claimName</code> .
PersistentVolume	Cluster	Discovered from the backing PV for each collected PVC by using <code>spec.volumeName</code> .
Secret	Namespaced	Discovered from volume secrets in <code>secret.secretName</code> ; containerDisk image pull secrets; CloudInitNoCloud and CloudInitConfigDrive user-data and network-data secret references; Sysprep secrets; KernelBoot container image pull secrets; and AccessCredentials SSH or password secrets.
ConfigMap	Namespaced	Discovered from volume ConfigMaps in <code>configMap.name</code> and Sysprep ConfigMaps.
ServiceAccount	Namespaced	Discovered from <code>serviceAccount.serviceAccountName</code> volume references.
NetworkAttachmentDefinition	Namespaced	Discovered from Multus networks in <code>spec.template.spec.networks[].multus.networkName</code> when the network name does not contain <code>/</code> . Only same-namespace NetworkAttachmentDefinition references are collected. Cross-namespace references are skipped.
Pod	Namespaced	Discovered for <code>virt-launcher</code> Pods by using the label selector <code>kubevirt.io=virt-launcher,vm.kubevirt.io/name=<vmName></code> .
Storage class, Trident volumes, and Trident backends	Cluster	Collected the same way as namespace-based applications, based on the PVCs discovered for the VM.



After all VM-scoped namespaced resources are collected, Trident Protect applies the standard `resourceFilter`, if one is configured. Cluster-scoped resource collection then proceeds as normal.

Namespace-based collection

For namespace-based applications, Trident Protect scans the selected namespaces and collects all matching resources. If any KubeVirt `VirtualMachine` objects are found, Trident Protect automatically collects their dependent resources, including instance types, preferences, and `virt-launcher` Pods, the same way it does for VM-scoped applications.

This means VM-specific resources are included in your backup or snapshot even when the application is defined with `includedNamespaces`.

Protect VMs using namespace-based application definitions

When you create a namespace-based application, Trident Protect scans the selected namespaces and collects all resources, including any KubeVirt VMs. For each virtual machine found, Trident Protect also collects the dependent resources needed to back up and restore that VM.

Example namespace-based application CR:

```
apiVersion: protect.trident.netapp.io/v1
kind: Application
metadata:
  name: <application_name>
  namespace: <application_namespace>
spec:
  includedNamespaces:
    - namespace: <vm_namespace>
```

After you create the application, use the standard Trident Protect workflows for backups and snapshots:

- For snapshot YAML examples, see [Create an on-demand snapshot](#).
- For backup YAML examples, see [Create an on-demand backup](#).

Restore behavior for namespace-based VM applications

Because namespace-based applications collect resources from the selected namespaces, the backup or snapshot can include KubeVirt runtime objects and KubeVirt-managed metadata. Some of these objects should not be restored directly because they represent runtime state from the source cluster.

During restore, Trident Protect applies KubeVirt-aware transforms to remove source-cluster metadata and skip transient KubeVirt resources. This helps KubeVirt recreate the VM cleanly in the restored environment.

Resource	Transform
VirtualMachine	Removes labels and annotations that start with <code>kubevirt.io/</code> .
VirtualMachineInstance	Removes labels and annotations that start with <code>kubevirt.io/</code> .
VirtualMachineInstanceMigration	Ignored and not restored.
VirtualMachineSnapshot	Ignored and not restored.
VirtualMachineSnapshotContent	Ignored and not restored.
Pod (virt-launcher)	Skipped when the pod has the <code>kubevirt.io=virt-launcher</code> label.



Trident Protect restores KubeVirt resources in a defined order to avoid dependency issues. CDI DataVolumes (`cdi.kubevirt.io`) are restored first so VM disk data is available before the VM objects are created. After that, VirtualMachines (`kubevirt.io`) are restored before VirtualMachineInstances.

Dynamic VM protection with label selectors

If you want to protect VMs dynamically without listing each VM name, use namespace-based application definitions with label selectors.

Only the VirtualMachine CR needs labels. Dependent resources such as DataVolumes, PVCs, Secrets, and ConfigMaps are collected automatically.

Steps

1. Label virtual machine CRs with a shared label:

```
kubectl label virtualmachine <vm_name_1> app=<label_value> -n
<vm_namespace>
kubectl label virtualmachine <vm_name_2> app=<label_value> -n
<vm_namespace>
```

2. Create an application CR with includedNamespaces and a labelSelector:

```
apiVersion: protect.trident.netapp.io/v1
kind: Application
metadata:
  name: <application_name>
  namespace: <application_namespace>
spec:
  includedNamespaces:
    - namespace: <vm_namespace>
    labelSelector:
      matchLabels:
        app: <label_value>
```

3. Create backups and restores using the standard Trident Protect workflow.

Additional details:

- You can add or remove VMs from protection at any time by changing labels:

```
kubectl label virtualmachine <vm_name_to_add> app=<label_value> -n
<vm_namespace>
kubectl label virtualmachine <vm_name_to_remove> app
=<different_label_value> -n <vm_namespace>
```

- You can also use label selectors across multiple namespaces:

```
spec:
  includedNamespaces:
```

```

- namespace: <vm_namespace_1>
  labelSelector:
    matchLabels:
      <label_key>: <label_value>
- namespace: <vm_namespace_2>
  labelSelector:
    matchLabels:
      <label_key>: <label_value>

```

Protect specific VMs with includedVirtualMachines

Use `includedVirtualMachines` when you want to protect a named set of `VirtualMachines` across one or more namespaces and selectively restore specific VMs from an archive.

Both the application and restore CRs use the same `includedVirtualMachines` selector format.

Use this selector in different contexts:

- In an application CR, it defines which VMs are protected.
- In a restore CR, it defines which VMs are restored from the backup or snapshot archive.



On restore CRs, `includedVirtualMachines` is optional. If you do not include this field, the entire archive is restored. When specified, only the listed VMs and their dependent resources are restored.

The YAML below shows the shared selector structure:

```

includedVirtualMachines:
- namespace: <vm_namespace_1>
  names:
    - <vm_name_1>
    - <vm_name_2>
- namespace: <vm_namespace_2>
  names:
    - <vm_name_3>

```



Multiple entries in `includedVirtualMachines` can reference different namespaces, which enables multi-namespace VM-aware application definitions and restores.

Table 1. includedVirtualMachines selector fields

Field	Type	Description
namespace	string	Kubernetes namespace that contains the target virtual machine objects.
names	list of strings	One or more virtual machine names in that namespace.

Example VM-based application:

```
apiVersion: protect.trident.netapp.io/v1
kind: Application
metadata:
  name: <application_name>
  namespace: <application_namespace>
spec:
  includedVirtualMachines:
    - namespace: <vm_namespace_1>
      names:
        - <vm_name_1>
        - <vm_name_2>
    - namespace: <vm_namespace_2>
      names:
        - <vm_name_3>
  includedClusterScopedResources:
    - group: storage.k8s.io
      kind: StorageClass
```

Restore specific KubeVirt VMs

You can use `includedVirtualMachines` on restore CRs to restore only selected KubeVirt VMs from a backup, snapshot, or replicated snapshot archive.

When `includedVirtualMachines` is specified on a restore CR, Trident Protect restores only the listed `VirtualMachines` and their dependent resources, such as `VirtualMachineInstance`, `DataVolume`, `PersistentVolumeClaim`, `PersistentVolume`, `Secret`, `ConfigMap`, `ServiceAccount`, and other VM-related objects.

If you do not specify `includedVirtualMachines` on the restore CR, Trident Protect restores the full archive according to the restore CR configuration.

Supported restore CRs and CLI commands

The following table lists the restore CRs that support `includedVirtualMachines` and their corresponding CLI commands.

Restore CR	CLI command	Description
BackupRestore	<code>tridentctl-protect create backuprestore</code>	Restores from a backup to new or different namespace mappings. Supports <code>namespaceMapping</code> and <code>storageClassMapping</code> .
BackupInplaceRestore	<code>tridentctl-protect create backupinplacerestore</code>	Restores from a backup into the original namespace or namespaces.

Restore CR	CLI command	Description
SnapshotRestore	<code>tridentctl-protect create snapshotrestore</code>	Restores from a snapshot to new or different namespace mappings. Supports <code>namespaceMapping</code> and <code>storageClassMapping</code> .
SnapshotInplaceRestore	<code>tridentctl-protect create snapshotinplace restore</code>	Restores from a snapshot into the original namespace or namespaces.
ReplicateSnapshotRestore	YAML only, or use the supported CLI command if available in your environment	Restores from a replicated snapshot. Supports <code>namespaceMapping</code> , <code>storageClassMapping</code> , and in-place restore by using the <code>inPlaceRestore</code> flag.



The `tridentctl-protect` CLI does not have a flag for `includedVirtualMachines` on restore commands. To restore specific VMs, you must apply a restore CR YAML file directly using `kubectl apply`. You can write the YAML from scratch, or use the CLI with `--dry-run --output yaml` to generate a starting file, add the `includedVirtualMachines` field under `spec`, and then apply it.

VM-scoped restore behavior

The `includedVirtualMachines` field on a restore CR uses the same selector format as the application CR:

```
includedVirtualMachines:
- namespace: <source_vm_namespace>
  names:
  - <vm_name_1>
  - <vm_name_2>
```

For restore CRs, the namespace values in `includedVirtualMachines` refer to the source namespaces as they exist in the backup, snapshot, or replicated snapshot archive.

If you use `namespaceMapping`, Trident Protect first selects the requested VMs from the source archive and then applies namespace mappings during restore. If you use `storageClassMapping`, Trident Protect applies the storage class mappings during the restore transform phase.

Generate restore YAML with the CLI

You can use `tridentctl-protect` to generate a base restore CR and then edit the YAML to add `includedVirtualMachines`.

```
tridentctl-protect create <restore_command> <restore_name> \
  <restore_options> \
  --dry-run --output yaml > <restore_name>.yaml
```

Edit the generated YAML and add the `includedVirtualMachines` field under `spec`, then apply the file:

```
kubectl apply -f <restore_name>.yaml
```

Example: Define a VM-based application

The following application protects three specific VMs across two namespaces.

```
apiVersion: protect.trident.netapp.io/v1
kind: Application
metadata:
  name: production-vms
  namespace: prod
spec:
  includedVirtualMachines:
    - namespace: prod
      names:
        - app-server
        - database-server
    - namespace: monitoring
      names:
        - prometheus-vm
  includedClusterScopedResources:
    - group: storage.k8s.io
      kind: StorageClass
```

Example: Create a backup

Creating a backup for a VM-based application uses the standard backup CR. No VM-specific fields are required on the backup CR.

```
apiVersion: protect.trident.netapp.io/v1
kind: Backup
metadata:
  name: daily-backup
  namespace: prod
spec:
  applicationRef: production-vms
  appVaultRef: s3-vault
```

Example: Restore a single VM from a backup

The following example restores only `database-server` from the backup archive. The other VMs in the archive, such as `app-server` and `prometheus-vm`, are skipped.

```

apiVersion: protect.trident.netapp.io/v1
kind: BackupRestore
metadata:
  name: restore-db-only
  namespace: prod-vms-restore
spec:
  appVaultRef: s3-vault
  appArchivePath: backups/production-vms/2026-03-10T00-00-00Z
  namespaceMapping:
    - source: prod
      destination: prod-dr
  includedVirtualMachines:
    - namespace: prod
      names:
        - database-server

```

```
kubectl apply -f restore-db-only.yaml
```

Only `database-server` and its dependent resources, such as `DataVolumes`, `PVCs`, `PVs`, `Secrets`, `ConfigMaps`, and related `KubeVirt` resources, are restored.

Example: Restore a single VM in place from a backup

The following example restores only `database-server` into its original namespace from the backup archive.

```

apiVersion: protect.trident.netapp.io/v1
kind: BackupInplaceRestore
metadata:
  name: ipr-db-only
  namespace: prod
spec:
  appVaultRef: s3-vault
  appArchivePath: backups/production-vms/2026-03-10T00-00-00Z
  includedVirtualMachines:
    - namespace: prod
      names:
        - database-server

```

```
kubectl apply -f ipr-db-only.yaml
```

Example: Restore a single VM from a snapshot

The following example restores only `ubuntu-vm-blue-a627be38` from a snapshot archive.

```
apiVersion: protect.trident.netapp.io/v1
kind: SnapshotRestore
metadata:
  name: blue-vm-snap-restore
  namespace: blue
spec:
  appVaultRef: my-appvault
  appArchivePath: snapshots/protectctl-blue-vm/snap-1
  includedVirtualMachines:
    - namespace: blue
      names:
        - ubuntu-vm-blue-a627be38
```

```
kubectl apply -f blue-vm-snap-restore.yaml
```

Combine VM restore with other filters

You can use `includedVirtualMachines` together with `resourceFilter`, `namespaceMapping`, and `storageClassMapping` in the same restore CR to fine-tune what gets restored.

When you add a `resourceFilter`, Trident Protect first selects the VMs listed in `includedVirtualMachines` and their dependent resources. It then applies the resource filter to further narrow down what is included or excluded from those results.

```
spec:
  includedVirtualMachines:
    - namespace: <vm_namespace>
      names:
        - <vm_name>
  resourceFilter:
    resourceSelectionCriteria: Include
    resourceMatchers:
      - kinds: ["PersistentVolumeClaim"]
        names: ["<pvc_name>"]
```

For `BackupRestore`, `SnapshotRestore`, and `ReplicateSnapshotRestore`, `namespaceMapping` and `storageClassMapping` work alongside `includedVirtualMachines`. Namespace references in `includedVirtualMachines` correspond to source namespaces as they appear in the archive. Namespace and storage class mappings are applied during the restore transform phase after filtering is complete.

For restore behavior details, refer to [Restore applications using Trident Protect](#).

Create a VM-scoped application using the CLI

Use the `--virtual-machines` flag to scope an application to specific KubeVirt VMs instead of whole namespaces. The `--virtual-machines` flag is mutually exclusive with `--namespaces`.

Steps

1. Create the application by running one of the following commands, replacing values in brackets with information from your environment. The format for the `--virtual-machines` flag is `<namespace> (<vm_name>)`. For multiple VMs, use comma-separated names within the parentheses:
 - Single VM in a namespace:

```
tridentctl-protect create application <application_name> \  
-n <application_namespace> \  
--virtual-machines "<vm_namespace>(<vm_name>)"
```

- Multiple VMs across namespaces:

```
tridentctl-protect create application <application_name> \  
-n <application_namespace> \  
--virtual-machines "<vm_namespace_1>(<vm_name_1>,<vm_name_2>  
>,<vm_namespace_2>(<vm_name_3>)"
```

Related information

- [Define an application for management](#)
- [Protect applications](#)
- [Restore applications using Trident Protect](#)
- [Use advanced restore settings](#)
- [Install the Trident Protect CLI plugin](#)
- [Use AppVault objects to manage buckets](#)

Manage Trident Protect execution hooks

An execution hook is a custom action that you can configure to run in conjunction with a data protection operation of a managed app. For example, if you have a database app, you can use an execution hook to pause all database transactions before a snapshot, and resume transactions after the snapshot is complete. This ensures application-consistent snapshots.

Types of execution hooks

Trident Protect supports the following types of execution hooks, based on when they can be run:

- Pre-snapshot

- Post-snapshot
- Pre-backup
- Post-backup
- Post-restore
- Post-failover

Order of execution

When a data protection operation is run, execution hook events take place in the following order:

1. Any applicable custom pre-operation execution hooks are run on the appropriate containers. You can create and run as many custom pre-operation hooks as you need, but the order of execution of these hooks before the operation is neither guaranteed nor configurable.
2. Filesystem freezes occur, if applicable. [Learn more about configuring filesystem freezing with Trident Protect.](#)
3. The data protection operation is performed.
4. Frozen filesystems are unfrozen, if applicable.
5. Any applicable custom post-operation execution hooks are run on the appropriate containers. You can create and run as many custom post-operation hooks as you need, but the order of execution of these hooks after the operation is neither guaranteed nor configurable.

If you create multiple execution hooks of the same type (for example, pre-snapshot), the order of execution of those hooks is not guaranteed. However, the order of execution of hooks of different types is guaranteed. For example, the following is the order of execution of a configuration that has all of the different types of hooks:

1. Pre-snapshot hooks executed
2. Post-snapshot hooks executed
3. Pre-backup hooks executed
4. Post-backup hooks executed



The preceding order example only applies when you run a backup that does not use an existing snapshot.



You should always test your execution hook scripts before enabling them in a production environment. You can use the 'kubectl exec' command to conveniently test the scripts. After you enable the execution hooks in a production environment, test the resulting snapshots and backups to ensure they are consistent. You can do this by cloning the app to a temporary namespace, restoring the snapshot or backup, and then testing the app.



If a pre-snapshot execution hook adds, changes, or removes Kubernetes resources, those changes are included in the snapshot or backup and in any subsequent restore operation.

Important notes about custom execution hooks

Consider the following when planning execution hooks for your apps.

- An execution hook must use a script to perform actions. Many execution hooks can reference the same script.

- Trident Protect requires the scripts that execution hooks use to be written in the format of executable shell scripts.
- Script size is limited to 96KB.
- Trident Protect uses execution hook settings and any matching criteria to determine which hooks are applicable to a snapshot, backup, or restore operation.



Because execution hooks often reduce or completely disable the functionality of the application they are running against, you should always try to minimize the time your custom execution hooks take to run. If you start a backup or snapshot operation with associated execution hooks but then cancel it, the hooks are still allowed to run if the backup or snapshot operation has already begun. This means that the logic used in a post-backup execution hook cannot assume that the backup was completed.

Execution hook filters

When you add or edit an execution hook for an application, you can add filters to the execution hook to manage which containers the hook will match. Filters are useful for applications that use the same container image on all containers, but might use each image for a different purpose (such as Elasticsearch). Filters allow you to create scenarios where execution hooks run on some but not necessarily all identical containers. If you create multiple filters for a single execution hook, they are combined with a logical AND operator. You can have up to 10 active filters per execution hook.

Each filter you add to an execution hook uses a regular expression to match containers in your cluster. When a hook matches a container, the hook will run its associated script on that container. Regular expressions for filters use the Regular Expression 2 (RE2) syntax, which does not support creating a filter that excludes containers from the list of matches. For information on the syntax that Trident Protect supports for regular expressions in execution hook filters, see [Regular Expression 2 \(RE2\) syntax support](#).



If you add a namespace filter to an execution hook that runs after a restore or clone operation and the restore or clone source and destination are in different namespaces, the namespace filter is only applied to the destination namespace.

Execution hook examples

Visit the [NetApp Verda GitHub project](#) to download real execution hooks for popular apps such as Apache Cassandra and Elasticsearch. You can also see examples and get ideas for structuring your own custom execution hooks.



The scripts in the Verda repository are provided as-is and are not officially supported by NetApp. NetApp support covers only the execution hooks framework within Trident Protect.

Create an execution hook

You can create a custom execution hook for an app using Trident Protect. You need to have Owner, Admin, or Member permissions to create execution hooks.

Use a CR

Steps

1. Create the custom resource (CR) file and name it `trident-protect-hook.yaml`.

2. Configure the following attributes to match your Trident Protect environment and cluster configuration:

- **metadata.name:** (*Required*) The name of this custom resource; choose a unique and sensible name for your environment.
- **spec.applicationRef:** (*Required*) The Kubernetes name of the application for which to run the execution hook.
- **spec.stage:** (*Required*) A string indicating which stage during the action that the execution hook should run. Possible values:
 - Pre
 - Post
- **spec.action:** (*Required*) A string indicating which action the execution hook will take, assuming any execution hook filters specified are matched. Possible values:
 - Snapshot
 - Backup
 - Restore
 - Failover
- **spec.enabled:** (*Optional*) Indicates whether this execution hook is enabled or disabled. If not specified, the default value is true.
- **spec.hookSource:** (*Required*) A string containing the base64-encoded hook script.
- **spec.timeout:** (*Optional*) A number defining how long in minutes that the execution hook is allowed to run. The minimum value is 1 minute, and the default value is 25 minutes if not specified.
- **spec.arguments:** (*Optional*) A YAML list of arguments that you can specify for the execution hook.
- **spec.matchingCriteria:** (*Optional*) An optional list of criteria key value pairs, each pair making up an execution hook filter. You can add up to 10 filters per execution hook.
- **spec.matchingCriteria.type:** (*Optional*) A string identifying the execution hook filter type. Possible values:
 - ContainerImage
 - ContainerName
 - PodName
 - PodLabel
 - NamespaceName
- **spec.matchingCriteria.value:** (*Optional*) A string or regular expression identifying the execution hook filter value.

Example YAML:

```
apiVersion: protect.trident.netapp.io/v1
kind: ExecHook
metadata:
  name: example-hook-cr
  namespace: my-app-namespace
```

```

annotations:
  astra.netapp.io/astra-control-hook-source-id:
/account/test/hookSource/id
spec:
  applicationRef: my-app-name
  stage: Pre
  action: Snapshot
  enabled: true
  hookSource: IyEvYmluL2Jhc2gKZWNoYAiZXhhbXBsZSBzY3JpcHQiCg==
  timeout: 10
  arguments:
    - FirstExampleArg
    - SecondExampleArg
  matchingCriteria:
    - type: containerName
      value: mysql
    - type: containerImage
      value: bitnami/mysql
    - type: podName
      value: mysql
    - type: namespaceName
      value: mysql-a
    - type: podLabel
      value: app.kubernetes.io/component=primary
    - type: podLabel
      value: helm.sh/chart=mysql-10.1.0
    - type: podLabel
      value: deployment-type=production

```

3. After you populate the CR file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-hook.yaml
```

Use the CLI

Steps

1. Create the execution hook, replacing values in brackets with information from your environment. For example:

```

tridentctl-protect create exehook <my_exec_hook_name> --action
<action_type> --app <app_to_use_hook> --stage <pre_or_post_stage>
--source-file <script-file> -n <application_namespace>

```

Manually run an execution hook

You can manually run an execution hook for testing purposes or if you need to re-run the hook manually after a failure. You need to have Owner, Admin, or Member permissions to manually run execution hooks.

Manually running an execution hook consists of two basic steps:

1. Create a resource backup, which collects resources and creates a backup of them, determining where the hook will run
2. Run the execution hook against the backup

▼ Step 1: Create a resource backup

Use a CR

Steps

1. Create the custom resource (CR) file and name it `trident-protect-resource-backup.yaml`.
2. Configure the following attributes to match your Trident Protect environment and cluster configuration:
 - **metadata.name:** *(Required)* The name of this custom resource; choose a unique and sensible name for your environment.
 - **spec.applicationRef:** *(Required)* The Kubernetes name of the application for which to create the resource backup.
 - **spec.appVaultRef:** *(Required)* The name of the AppVault where the backup contents are stored.
 - **spec.appArchivePath:** The path inside AppVault where the backup contents are stored. You can use the following command to find this path:

```
kubectl get backups <BACKUP_NAME> -n my-app-namespace -o  
jsonpath='{.status.appArchivePath}'
```

Example YAML:

```
---  
apiVersion: protect.trident.netapp.io/v1  
kind: ResourceBackup  
metadata:  
  name: example-resource-backup  
spec:  
  applicationRef: my-app-name  
  appVaultRef: my-appvault-name  
  appArchivePath: example-resource-backup
```

3. After you populate the CR file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-resource-backup.yaml
```

Use the CLI

Steps

1. Create the backup, replacing values in brackets with information from your environment. For example:

```
tridentctl protect create resourcebackup <my_backup_name> --app  
<my_app_name> --appvault <my_appvault_name> -n <my_app_namespace>  
--app-archive-path <app_archive_path>
```

2. View the status of the backup. You can use this example command repeatedly until the operation is complete:

```
tridentctl protect get resourcebackup -n <my_app_namespace>  
<my_backup_name>
```

3. Verify that the backup was successful:

```
kubectl describe resourcebackup <my_backup_name>
```

▼ Step 2: Run the execution hook

Use a CR

Steps

1. Create the custom resource (CR) file and name it `trident-protect-hook-run.yaml`.
2. Configure the following attributes to match your Trident Protect environment and cluster configuration:
 - **metadata.name:** *(Required)* The name of this custom resource; choose a unique and sensible name for your environment.
 - **spec.applicationRef:** *(Required)* Ensure this value matches the application name from the ResourceBackup CR you created in step 1.
 - **spec.appVaultRef:** *(Required)* Ensure this value matches the appVaultRef from the ResourceBackup CR you created in step 1.
 - **spec.appArchivePath:** Ensure this value matches the appArchivePath from the ResourceBackup CR you created in step 1.

```
kubectl get backups <BACKUP_NAME> -n my-app-namespace -o  
jsonpath='{.status.appArchivePath}'
```

- **spec.action:** (*Required*) A string indicating which action the execution hook will take, assuming any execution hook filters specified are matched. Possible values:
 - Snapshot
 - Backup
 - Restore
 - Failover
- **spec.stage:** (*Required*) A string indicating which stage during the action that the execution hook should run. This hook run will not run hooks in any other stage. Possible values:
 - Pre
 - Post

Example YAML:

```
---
apiVersion: protect.trident.netapp.io/v1
kind: ExecHooksRun
metadata:
  name: example-hook-run
spec:
  applicationRef: my-app-name
  appVaultRef: my-appvault-name
  appArchivePath: example-resource-backup
  stage: Post
  action: Failover
```

3. After you populate the CR file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-hook-run.yaml
```

Use the CLI

Steps

1. Create the manual execution hook run request:

```
tridentctl protect create exehooksruntime <my_exec_hook_run_name> -n
<my_app_namespace> --action snapshot --stage <pre_or_post> --app
<my_app_name> --appvault <my_appvault_name> --path
<my_backup_name>
```

2. Check the status of the execution hook run. You can run this command repeatedly until the operation is complete:

```
tridentctl protect get exehooksruntime -n <my_app_namespace>
```

```
<my_exec_hook_run_name>
```

3. Describe the exehooksruntime object to see the final details and status:

```
kubectl -n <my_app_namespace> describe exehooksruntime  
<my_exec_hook_run_name>
```

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