Dell EMC PowerStore

Protecting Your Data

Version 2.x

June 2021
Rev. A01

Notes, cautions, and warnings

NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Contents

Additional Resources

As part of an improvement effort, revisions of the software and hardware are periodically released. Some functions that are described in this document are not supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information about product features. Contact your service provider if a product does not function properly or does not function as described in this document.

Where to get help

Support, product, and licensing information can be obtained as follows:

- **Product information**

  For product and feature documentation or release notes, go to the PowerStore Documentation page at https://www.dell.com/powerstoredocs.

- **Troubleshooting**

  For information about products, software updates, licensing, and service, go to https://www.dell.com/support and locate the appropriate product support page.

- **Technical support**

  For technical support and service requests, go to https://www.dell.com/support and locate the **Service Requests** page. To open a service request, you must have a valid support agreement. Contact your Sales Representative for details about obtaining a valid support agreement or to answer any questions about your account.

1

Introduction

This chapter contains the following information:

Topics:

Data protection

PowerStore provides both local and remote data protection. Using the PowerStore Manager, you can protect your data locally by creating snapshots (point-in-time copies) of volumes, volume groups, virtual machines, or file systems. You can also apply remote protection by replicating your data to a remote system for redundancy in the event of a disaster.

PowerStore enables you to create custom protection policies, which are sets of rules for snapshot creation, replication, or both, and assign them to storage resources . Protection policies apply the defined rules on the storage resource, providing it with local and/or remote protection.

Snapshots

Snapshots are read-only, point-in-time copies of data of a volume, volume group, virtual machine, or file system. Creating a snapshot saves the state of the storage resource at that particular point-in-time. Using snapshots, you can easily protect your data locally and restore a storage resource to a previous state.

You can manually create snapshots at any time. It is also possible to configure snapshot rules as part of a protection policy and assign them to the relevant storage resources. The system automatically creates snapshots of the relevant resource according to the schedule specified in the protection policy.

If data corruption occurs or data is accidentally deleted, you can recover the data from the snapshots or restore the volume or volume group to the point in time when the snapshot was created.

For file systems, you can create two access types of read-only file snapshot: protocol and .snapshot. The default access type is protocol, which can be exported as an SMB share, NFS export, or both. You can share and mount the snapshot on a client like any other file system. For .snapshot access types, you can access the files within the snapshot from the production file system in the `.snapshot` subdirectory of each directory.

You can also create write-order consistent and application consistent snapshots of volumes:

- Write-order consistent snapshots - PowerStore holds all writes on the volume group members to provide a uniform point-in-time copy, and therefore ensure consistent protection across all member volumes. You can generate write-order consistent snapshots from the PowerStore Manager.
- Application consistent snapshots - You can create application consistent snapshots of a volume or a volume group using AppSync. When creating an application consistent snapshot, all incoming I/O for a given application is quiesced while the snapshot is created.

To verify whether a snapshot is write-order consistent or application consistent, look at the **Write-Order Consistent** and **Application Consistent** columns in the snapshot tables for a volume or volume group in PowerStore Manager.

NOTE: If you cannot see these columns, you can add them, using the **Add Filters** option.

Mapping snapshots to hosts is not supported in PowerStore. To allow a connected host to access a snapshot, you can create a thin clone - a writable, space efficient copy of the snapshot - and map it to a host. You can update the thin clone from different snapshots using the refresh operation.

For details on the possible snapshot-related operations you can perform, using the PowerStore Manager, refer to the Snapshots section.

Replication

Data replication is a process in which storage data is duplicated to a remote system, which provides an enhanced level of redundancy in case the main production system fails. Replication minimizes the downtime-associated costs of a system failure and simplifies recovery following a natural disaster or human error.

PowerStore supports asynchronous remote replication for volumes and volume groups. Replication is not supported on file systems or for virtual volumes.

To configure replication:

1. Create a remote connection between the source and destination systems
2. Configure a protection policy with a replication rule that best meets your business needs
3. Assign the protection policy to a storage resource

PowerStore enables you to failover control to the remote system and reverse the direction of a remote protection session. Failover may be required in the following cases:

- If you want to migrate data to a new system and then switch to working from it without losing data. In this case, failover can be performed with no data loss.
- When there is no access to the data in the source system, you can switch to the remote system and continue to work, using the latest point-in-time remote protection copy. In such a case there may be some data loss, since the latest copy in the remote system does not include data changes made between the time this copy was created and the time the data in the system became inaccessible.
- When the data in the source system is accessible but its integrity may be compromised. In such a case, you should revert to the latest point-in-time protection copy created before the data was compromised.

You can perform a failover test on the destination storage resource to test the system disaster recovery readiness.

For detailed information on replication-related procedures you can perform, refer to the Replication section.

Protection policies

A protection policy consists of snapshot rules, replication rules, or both, that you can create to establish consistent data protection across storage resources. After configuring a protection policy, you can assign it to new or existing storage resources.

Each protection policy can only include one replication rule, and up to four snapshot rules. A replication/snapshot rule can be included in multiple policies.

Protection policies automatically manage snapshots or replication operations, based on the rules included in them. You can create policies with various rules that provide different levels of protection to meet your local and remote protection needs, and assign a policy to multiple storage resources to provide identical protection to those resources.

Based on your user privileges, you can create or modify relevant rules and policies.

If you want to create a new snapshot or replication rule, ensure that you review the parameters and your business requirements with an administrator before proceeding. This helps achieve and maintain consistent policies across the system.

For detailed information on protection policies-related procedures you can perform, refer to the Protection Policies section.

2

Snapshots

This chapter contains the following information:

Topics:

Create a snapshot

Creating a snapshot saves the state of the storage resource and all files and data within it at a particular point in time. You can use snapshots to restore the entire storage resource to a previous state. You can Create a snapshot of a volume, volume group, file system, or virtual machine.

Before creating a snapshots, consider the following:

- Snapshots are not full copies of the original data. Do not rely on snapshots for mirrors, disaster recovery, or high-availability tools. Because snapshots are partially derived from the real-time data of the storage resources, they can become inaccessible if the storage resource becomes inaccessible.
- Although snapshots are space efficient, they consume overall system storage capacity . Ensure that the system has enough capacity to accommodate snapshots.
- When configuring snapshots, review the snapshot retention policy that is associated with the storage resource. You may want to change the retention policy in the associated rules or manually set a different retention policy, depending on the purpose of the snapshot.
- Manual snapshots that are created with PowerStore Manager are retained for one week after creation (unless configured otherwise).
- If the maximum number of snapshots is reached, no more can be created. In this case, to enable creation of new snapshots, you are required to delete existing snapshots.

To manually create a snapshot of a volume:

NOTE: You can also create a snapshot for a volume group, file system or a virtual machine.

1. To open the Volumes window, select **Storage** > **Volumes**.
2. Click the check box next to the relevant volume to select it and then select **Protect** > **Create Snapshot**.
3. In the **Create Snapshot of Volume** slide-out panel, enter a unique name for the snapshot, and set the **Local Retention Policy**.
   NOTE: Retention period is set to one week by default. You can set a different retention period or select the **No Automatic Deletion** for indefinite retention.
4. Click **Create Snapshot**.

Create a thin clone

Thin clones are writable copies of a snapshot, volume, volume group, or file system that can be accessed by a host. Unlike a full clone, a thin clone is a space efficient copy that shares data blocks with its parent object and not a full backup of the original resource. A thin clone can be created directly as a copy of the parent object or using one of its snapshots.

Thin clones retain full read access to the original resource. You can modify the data within the thin clone while preserving the original snapshot.

Using thin clones, you can establish hierarchical points in time to preserve data over different stages of data modifications. If the parent resource is deleted, migrated, or replicated, the thin clone is unaffected.

You can perform the following actions to thin clones:

- Map thin clones to different hosts.
- Refresh the thin clone (not applicable for file systems).
- Restore the thin clone from a backup.
- Apply protection policies to thin clones.

To create a thin clone of a volume or volume group:

1. Select **Storage** > **Volumes** or **Storage** > **Volume Groups** to open the relevant resource window.
2. Click the check box next to the relevant volume or volume group and then select **Repurpose** > **Create Thin Clone using Volume** (or **Create Thin Clone** for volume group).
3. In the **Create Thin Clone** slide-out window perform the following:
   - Enter thin clone name.
   - Enter description (optional).
   - Set performance policy (only for thin clones created from volumes) .
   - Set host connectivity (only for thin clones created from volumes).
   - Set protection policy.
4. Click **Clone**.

To create a thin clone from a snapshot:

1. Open the relevant storage resource window.
2. Click a resource to open its Overview window.
3. Click the **Protection** tab.
4. Click **Snapshots**.
5. Select a snapshot from the table and then select **More actions** > **Create Thin Clone using Snapshot**.

Using clones to access read-only snapshots from hosts

Mapping and unmapping block snapshots to hosts is not supported in PowerStore. To allow a connected host to access a snapshot, create a thin clone of the snapshot and map it to a host. After creating the thin clone, you can use the refresh operation to update the thin clone from different snapshots. For more information, see Refresh a storage resource.

Similar to block snapshots, file snapshots can be mounted on hosts either directly (to allow read-only access) or by creating a thin clone (to allow read-write access). To mount the file system directly, the snapshots can be exported as NFS export or SMB share.

You can export snapshots using one of the following access types:

- Protocol - The snapshot is exported with a new share name.
- .snapshot - You can see the snapshot on Unix/Linux under the .snapshot directory of the file system, and on Windows, by right-clicking the file system and selecting the **Previous Version** option.

Refresh a storage resource

The refresh operation is used to replace the contents of a storage resource with contents from a related resource (a clone or an indirect child snapshot). You can create a duplicate of the production environment to be used for various purposes (such as test and dev, reporting etc.). To keep the duplicated environment up-to-date, it should be updated with a storage resource that includes the recent changes.

You can use the refresh operation in the following scenarios:

- Refresh a thin clone from the base volume.
- Refresh a storage resource or thin clone from another thin clone in the family.
- Refresh a storage resource or thin clone from a snapshot of a related thin clone or base volume.

NOTE: For file systems, you can refresh a snapshot of a file system with its direct parent file system.

If you refresh the thin clone of a snapshot that has derivative snapshots, the derivative snapshots remain unchanged and the family hierarchy stays intact. If you refresh a volume group, the point-in-time image on all member volumes is also refreshed.

When refreshing a resource from a snapshot that was replicated from a remote system, check the creation time value to ensure that you are using the correct snapshot. The **Source Data Time** value of replicated snapshots reflects the original source data time, and the **Creation Time** value is updated to the time of replication.

NOTE: Because the refresh operation replaces the contents of a storage resource, it is recommended to take a snapshot of the resource before refreshing it. Creating a backup allows you to revert to a previous point in time.

Before refreshing a snapshot, it is mandatory to shut down the application and unmount the file system that is running on the production host, and then flush the host cache to prevent data corruption during the refresh operation.

To refresh a storage resource from a snapshot:

1. Open the relevant storage resource list window.
2. Select the storage resource from which the snapshot was taken to open its Overview window.
3. Click the **Protection** tab, and then click **Snapshots**.
4. From the snapshots list, select the snapshot you want to use for the refresh operation.
5. Click **More Actions** > **Refresh using Snapshot**.
6. In the **Refresh using Snapshot** slide-out panel, select the volume or clone you want to refresh from the **Volume being refreshed** drop-down list.
7. Select whether to create a backup snapshot for the refreshed volume (the option is selected by default).
8. Click **Refresh**

To refresh a storage resource using a related volume:

1. Open the relevant storage resource list window.
2. Select the checkbox next to the storage resource and then select **Repurpose** > **Refresh Using Related Volume**.
3. In the **Refresh using Related Volume** slide-out panel, click the **Select volume to refresh from** and select the source volume.
4. Select whether to create a backup snapshot for the refreshed volume (the option is selected by default).
5. Click **Refresh**.

Restore a storage resource from a snapshot

The restore operation is used to reconstruct an environment following and event that may have compromised its data. You can use the restore operation to replace the contents of a parent storage resource with data from a directly associated snapshot. Restoring resets the data in the parent storage resource to the point in time at which the snapshot was taken.

Before restoring a snapshot, it is mandatory to shut down the application and unmount the file system that is running on the production host, and then flush the host cache to prevent data corruption during the restore operation.

If you restore a volume group, all member volumes are restored to the point in time associated with the source snapshot.

When restoring a file system you can only use a direct child snapshot as the source for restore.

When restoring a resource from a snapshot that was replicated from a remote system, check the source data time value to ensure that you are using the correct snapshot.

To restore a storage resource:

1. Check the check box next to the storage resource you wish to restore.
2. Select **Protect** > **Restore from Snapshot**.
3. In the **Restore Volume from Snapshot** slide-out panel, select the snapshot to use for the restore operation.
4. Select whether to create a backup snapshot of the restored object (the option is selected by default).
   NOTE: Because the restore operation replaces the contents of a storage resource, it is recommended to create a snapshot prior to restoring. Creating a backup allows you to revert to the original data.
5. Click **Restore**.

NOTE: You can also restore the storage resource by selecting the resource snapshot from the **Snapshots** view of the resource **Protection** tab, and then clicking **More Actions** > **Restore from Snapshot**.

3

Protection Policies

This chapter contains the following information:

Topics:

Create snapshot rules

You can create snapshot rules to control parameters such as the frequency of snapshot creation, and snapshots retention period. Snapshot rules, combined with replication rules, enable you to configure and apply consistent data protection policies to storage resources based on the data protection requirements.

If you want to create a new snapshot rule in addition to the existing rules, it is recommended to review the business requirements with an administrator before proceeding. This can help in achieving and maintaining consistent policies across the system.

To create a new snapshot rule:

1. Select **Protection** > **Protection Policies**.
2. In the **Protection Policies** window, click **Snapshot Rules** on the **Protection** bar .
3. In the **Snapshot Rules** window, click **Create**.
4. In the **Create Snapshot Rule** slide-out panel, enter a name for the new rule.
5. Set the following:
   - Select the days on which a snapshot will be created.
   - Set the frequency/start time:
     - For a snapshot to be taken at a fixed interval, select this option and set the number of hours after which a snapshot will be created.
     - For a snapshot to be taken at a particular time of the selected days, select the **Time of day** option and set the time and time zone.
   - Set the retention period.
   - For file snapshots , select the file snapshot access type.
6. Click **Create**.

Create replication rules

A replication rule is a set of parameters the system uses to synchronize data in a replication session. The parameters include selecting a replication destination and setting a recovery point objective (RPO).

After you have configured a replication rule, you can choose to use it in a new or existing protection policy, which then automatically changes or applies the replication session parameters for any storage resource that uses the protection policy.

You cannot change a protection policy to use a different replication rule with a different remote system. To change a protection policy with a replication rule using a different remote system, remove the old policy before assigning a new one.

NOTE: Changing a remote system requires a full synchronization.

If you want to create a new replication rule in addition to the existing rules, it is recommended to review the parameters and your business requirements with an administrator before proceeding. This can help in achieving and maintaining consistent policies across the system.

To create a new replication rule:

1. Select **Protection** > **Protection Policies**.
2. In the **Protection Policies** window, click **Replication Rules** on the **Protection** bar .
3. In the **Replication Rules** window, click **Create**.
4. In the **Create Replication Rule** slide-out panel, enter a name for the new rule.
5. Set the following:
   - Select an existing replication destination or configure a new destination.
   - Set the RPO.
   - Set the alert threshold.
6. Click **Create**.

Recovery point objective

Recovery point objective (RPO) indicates the acceptable amount of data, measured in units of time, that may be lost in case a failure occurs. When you set up a replication rule, you can configure automatic synchronization based on the RPO. Possible RPO values range from 5 minutes to 24 hours. The default RPO value is one hour.

NOTE: A smaller RPO interval provides more protection and consumes less space. However, it has a higher performance impact, resulting in more network traffic. A higher RPO interval may result in more space consumption, which can affect snapshot schedules and space thresholds.

Alert threshold

When you configure a replication rule, you can specify an alert threshold, which is the amount of time the system will wait before generating a compliance alert when a replication session does not meet the RPO. Setting the alert threshold to zero means that alerts will be generated if the actual synchronization time exceeds the RPO.

Create a protection policy

About this task

Create a protection policy to provide local and/or remote protection for your storage resources. Each protection policy can include one replication rule and up to four snapshot rules. A rule can be included in multiple policies.

Steps

1. Select **Protection** > **Protection Policies**
2. In the **Protection Policies** window, click **Create**.
3. In the **Create Protection Policy** slide-out panel, set the new policy's name.
4. Select the snapshot rules you want to include in the policy or create a new snapshot rule (refer to Create Snapshot Rules).
5. Select the replication rules you want to include in the policy or create a new replication rule (refer to Create Replication Rules).
6. Click **Create**.

Results

When you create a protection policy that includes a replication rule, the policy is automatically replicated to the remote system and assigned to destination resources created by the policy. The replicated policy and associated rules names are identical to the policy and rules on the source system with the name of the remote system appended at the end. Changes made to the original policy or included rules, are replicated to the remote system to maintain synchronization. After a replication failover, the replicated policy becomes active on the destination system.

The replicated policies and rules are managed by the system and are not displayed in the destination system policy and rules tables. However, you can see the rules details in the **Protection** tab of the replicating volumes or volume groups, by hovering over the replicated policy name.

Modify a protection policy

You can modify a protection policy by adding and removing snapshot and replication rules.

About this task

NOTE: Changing the settings of a protection policy applies the new settings to all objects to which the protection policy is assigned. If you need to change the protection policy for one resource, it is recommended to create a new protection policy, and assign it to that resource instead.

You cannot change the replication destination on a replication rule used in protection policies which are assigned to one or more storage resources. To reconfigure replication to a different remote system, unassign the protection policy and assign a new one with a different replication rule. Unassigning a protection policy with a replication rule will delete the associated replication session and assigning a new protection policy will create a new one, which requires a full synchronization to the new destination.

Steps

1. Select **Protection** > **Protection Policies**.
2. Select the check box next to the relevant policy and click **Modify**.
3. In the **Properties** slide-out panel, you can modify the following parameters:
   - Policy name
   - Selected snapshot rules
   - Selected replication rules
4. Click **Apply**.

Assign a protection policy to a storage resource

Assign a protection policy to one or more storage resources to apply the snapshot and replication rules included in the policy to the storage resource. The protection policy automatically performs snapshot operations and replication based on the specified parameters.

If a protection policy that meets your data protection requirements is available, you can assign it to a storage resource at anytime. When you assign a new protection policy that contains a replication rule to the storage resource, a complete initial synchronization is required.

You can assign protection policy to a storage resource during the resource creation or at a later stage.

NOTE: It is not possible to assign a protection policy containing replication rules to a storage resource that does not support replication.

To assign a protection policy to an existing storage resource:

1. Select the check box of the storage resource to which you want to assign a protection policy.
2. Select **Protect** > **Assign Protection Policy**.
3. From the **Assign Protection Policy** slide-out panel, select the protection policy.
4. Click **Apply**.

To assign a different protection policy with an existing storage resource:

1. Select the relevant storage resource to open its **Overview** window.
2. Click the **Protection** tab.
3. Next to the assigned protection policy name, click **Change**.
4. In the **Change Protection Policy** slide-out panel, select a different protection policy.
5. Click **Apply**.

NOTE: You can only change assignment of protection policies that do not have a replication rule, or if the remote system specified the new policy is the same as the one specified in the old policy. To change assignment of a protection policy with a replication rule using a different remote system, remove the old policy before assigning a new one.

Unassign a protection policy

Prerequisites

Removing the protection policy from a storage resource results in the following:

- Scheduled snapshots and replication based on the rules associated with the policy stop.
- Existing snapshots remain, and are retained in the system, based on the snapshot rule settings when they were created.
- The destination storage resource stays in read-only mode. You can clone the destination storage resource to get a read/write copy or change the **replication destination** attribute in the **Properties** page of the storage resource.

NOTE: You cannot unassign a protection policy while importing is in progress.

Steps

1. Select the check box of the storage resource to which you want to assign a protection policy.
2. Select **Protect** > **Unassign Protection Policy**.
3. Click **Unassign** to confirm.

4

Replication

This chapter contains the following information:

Topics:

Remote systems

Configuring a remote system connection between the source and destination systems enables remote replication. In PowerStore, the remote system connection is associated with the replication rule. You can create a remote system connection ahead of time, or while creating a replication rule.

NOTE: It is possible to create a remote connection between systems running different versions (1.x, 2.x).

Before creating a remote system connection, ensure that you have obtained the following remote system details:

- System IP address
- User authentication credentials for connecting to the system

To add a remote system connection:

1. Select **Protection** > **Remote Systems**.
2. In the **Remote Systems** window, click **Add**.
3. In the **Add Remote System** slide-out panel, configure the following:
   - Management IP address
   - Description (optional)
   - Network latency
   - Username and password
4. Click **Add**.

In the Remote Systems table you can:

- View remote systems connection status.
- Click a remote system to modify its attributes. You can change the management cluster IP address, description and network latency of a remote system connection.
- Select a remote system and click **Delete** to remove it. You cannot delete a remote system in the following instances:
  - If there are active replication sessions.
  - If there are remote protection policies active in the system associated with the remote system.
  - If there is a replication rule associated with the remote system.
- Select a remote system and click **More Actions** > **Verify and Update** to verify and update the connection to the remote system. Verify and update detects changes in the local and remote systems and reestablishes data connections, while also taking the Challenge Handshake Authentication Protocol (CHAP) settings into account.
- Monitor the management and data connection status for troubleshooting purposes.

Synchronization

PowerStore enables you to asynchronously update the destination resource with changes (such as changes in content, size, and membership) that occurred on the source resource since the last synchronization cycle.

Synchronization can occur either automatically - according to a set schedule - or manually. Snapshots are synchronized from the source system to the destination system, and maintain block sharing efficiency.

NOTE: When you add volumes to a volume group or change the size of the volume group during an asynchronous replication session, the changes do not immediately appear on the destination. You can either perform a manual synchronization or wait until the synchronization occurs based on the RPO.

You can synchronize a replication session when it is in the following states:

- Operating normally
- System paused

While a replication session is synchronizing, you can take the following actions:

- Planned failover from the source system
- Fail over from the destination system
- Pause replication sessions from the source or destination system
- Delete a replication session by removing a protection policy

If synchronization fails, the replication session is placed in a system paused state. When the system recovers, the replication session continues from the same point as when the system was paused.

Failover

Failing over a replication session includes switching roles between the source and destination systems and reversing the direction of the replication session.

There are two types of failovers:

- Planned failover - User initiated, includes synchronization between source and destination to prevent data loss.
- Unplanned failover - Initiated by the destination system in response to source system failure.

During a replication session failover, the system performs the following actions:

- Stop I/Os on the source object.
- Synchronize the source and destination storage objects (occurs only in a planned failover).
- Stop the replication session.
- Reverse roles between source and destination systems.
- Promote the latest object version on the new source.
- Resume I/Os on the new source (initiated by the user).

After a failover, you can access applications on the new source system to recover data.

Performing a failover test

After you set up a replication session, you can test the connection to ensure that your sites are correctly configured and prepared for disaster recovery.

During a failover test, the system performs a failover and production access is provided to the destination site using replicated data or a point-in-time snapshot. The destination storage resource is available in read/write mode, and production access is enabled for hosts and applications. You can verify your disaster recovery configuration while replication continues to run in the background.

When you wish to stop the failover test, select one of the following actions:

- Failover to the current test data - If you made changes to the data during the failover test, you can use the updated test data. This will stop the test and preserve the test data. Any data replicated from the source during the test will be discarded and the destination system will become the source.

  NOTE: You must acknowledge these changes before failing over to the test data.
- Stop the failover test - When you stop the test, production access to the destination will be disabled for hosts and applications and the destination storage resource will be updated with the latest data synched from the source system. You can create a backup snapshot of the test data before stopping the failover test.

Restrictions

A failover test can only be performed under the following conditions:

- The PowerStore system version on both the source and destination system is 2.x or later.
- The replication session state is not Initializing, Failing Over, Failed Over, Paused for NDU/Migration, or Failover Test in Progress.

During the failover test, you cannot execute the following actions on the destination system:

- Change volume group membership
- Increase volume group size
- Change volume group name
- Start migration
- Remove a protection policy

NOTE: You can still perform these actions from the source system.

You cannot perform a planned failover while a failover test is in progress. Stop the failover test to perform a planned failover. However, unplanned failovers may still occur uninterrupted in response to a disaster. If possible, it is recommended to stop the failover test before an unplanned failover, because any data replicated to the destination after the failover test started will be lost.

You can also pause and resume replication sessions during a failover test. If you delete a replication session during a failover test, the test will be cancelled.

Starting a failover test

You can start a failover test from the current destination data, or from any snapshot.

There are two ways to start a failover test:

- From **Protection** > **Replication**, select the replication session you want to test, then select **Start Failover Test**.
- From the **Protection** tab of the resource, select **Replication**, then select **Start Failover Test**.

After the failover test starts, an alert is raised on the replication session. The alert is cleared after the test is stopped.

Stopping a failover test

Before you stop the failover test, it is recommended that you unmount file systems and stop any running applications on the destination resource to avoid data corruption.

There are two ways to stop a failover test:

- From **Protection** > **Replication**, select the replication session that has a test in progress, then select **Stop Failover Test**.
- From the **Protection** tab of the resource with a test in progress, select **Replication**, then select **Stop Failover Test**.

You can also choose to create a snapshot to save the test data that was created during the failover test.

Planned Failover

When you perform a planned failover, the replication session is manually failed over from the source system to the destination system. Prior to the failover, the destination system is synchronized with the source system, to prevent any data loss.

Before performing a planned failover, make sure that you stop I/O operations for any applications and hosts. You cannot pause a replication session that is undergoing a planned failover.

During a planned failover, you can take the following actions:

- Perform an unplanned failover.
- Delete the replication session by removing the protection policy on the storage resource.

You cannot initiate a planned failover when a failover test is in progress.

You can initiate a planned failover test from the current source data, or from any snapshot.

There are two ways to initiate a planned failover:

- From **Protection** > **Replication**, select the relevant replication session, and then select **Planned Failover**.
- From the **Protection** tab of the resource, select **Replication**, and then select **Planned Failover**.

After a planned failover, the replication session is inactive. To synchronize the destination storage resource and resume the replication session use the **Reprotect** action. You can also select the auto-reprotect option before failing over, which automatically initiates the synchronization in the opposite direction (at the next RPO) after the failover is complete, and returns the source and the target system to a normal state.

Unplanned Failover

Unplanned failover occurs following events such as source system failure, or events on the source system that leads to downtime for production access. Unplanned failover is initiated from the destination system, and provides production access to the original destination resource from a synchronized point-in-time snapshot.

When the connection to the source system is re-established, the original source resource is placed into destination mode. After an unplanned failover, you can restore the system from the latest data or any point-in-time snapshot. Reprotect the replication session to synchronize the destination storage resource, and then resume the replication session.

A

Use cases

This chapter contains the following information:

Topics:

Snapshot and thin clone use cases

You can use snapshots and thin clones to restore corrupted volumes and create test environments.

Snapshots are read-only copies that can be used to save the current state of an object. You can use snapshots to quickly recover data if there is corruption or user error. Snapshots cannot be directly accessed by a host.

Thin clones are writable copies of a snapshot, volume, or volume group that can be accessed by a host. Thin clones can be created directly as a copy of the parent object or using one of its snapshots. Both snapshots and thin clones are space efficient copies that share data blocks with their parent object.

Using snapshots and thin clones for partial recovery of a volume

You can use snapshots and thin clones to recover part of a volume, such as individual files or database records, from a previous point in time. First, create a thin clone using the snapshot that contains the data you need to recover. Then, provide host access to the clone, and recover data from the host.

Using snapshots to restore a volume or volume group

You can use snapshots to roll back a volume to a previous point of time, if there is corruption. To revert a volume or volume group to a previous point in time, use the volume restore operation and supply a snapshot from before the corruption occurred. The restore operation is instantaneous. You can also create a backup snapshot to save the state of the volume or volume group before you use the restore operation.

Using thin clones to test a patch before applying it to the production volume

Before installing a patch or software update of a critical application on a volume, you can take a thin clone of the volume, then apply the update to the thin clone. After you have installed the update and verified that the update is safe for your environment, you can install the update on the other volumes.

Create thin clones for development use

Instead of provisioning volumes or volume groups for each individual developer, you can create thin clones. Creating thin clones of the volume or volume group enables you to distribute the same data and configuration to each developer. The thin clones also take up less space than if you had created a full clone of the volume, or provisioned individual volumes or volume groups. You can also take snapshots of thin clones and replicate them.

Replication use cases

You can use replication for planned downtime, such as during inter-cluster migration, the installation of a major software update, and disaster recovery.

Intercluster migration

If you need to migrate a storage object to another PowerStore cluster, you can set up a one-time replication between the two clusters, followed by a planned fail over to the new cluster to complete the migration. After the migration, dismantle the source object to reclaim space on the original cluster.

Using replication for planned downtime

Planned downtime is a situation where you take the source system offline for maintenance or testing, while operating off the destination system. Before the planned downtime, both the source and destination are running with an active replication session. There is no data loss in planned downtime.

In this scenario, the source system, Boston, is taken offline for maintenance, and the destination system, New York, is used as the production system during the maintenance period. After maintenance is over, return production to the Boston system.

To start planned downtime, select **Planned Failover** on the Boston source system. The New York destination system is fully synchronized with the source to ensure that there is no data loss. The session remains paused, while the Boston source system becomes read-only and the destination becomes read/write. The New York destination storage resource can provide access to the host. On the New York destination storage resource, select **Reprotect** to resume replication in the reverse direction.

To resume operations on the Boston system after maintenance, select **Planned Failover** on the New York system. After the failover is complete, **Reprotect** on the Boston system.

NOTE: To replicate data from the destination to the source with the reprotect operation, ensure that there is a replication policy on the destination system that has a replication rule pointing to the source system. For example, if the regular replication session is from a site in Boston to a site in New York, the replication policy on the destination storage resource in New York must point to Boston.

Using replication for disaster recovery

In this disaster recovery scenario, the source system, Boston, is unavailable due to a natural or human-caused disaster. A destination system, New York, was created, which contains a full copy, or replica, of the production data. Data access can be restored by failing over to New York because a replication session was configured between the Boston and New York systems.

Using replicas for disaster recovery minimizes potential data loss. The replica is up-to-date with the last time that the destination synchronized with the source, as specified in the associated replication rule. The amount of potential data loss is based on the recovery point objective (RPO) setting in the associated replication rule. The replication session can be failed over to the New York destination system, using the latest data that was replicated from Boston.

After the session is failed over to the New York system, it becomes read/write. When originally establishing a replication session between the source and destination systems, the storage resource was given the correct access permissions to the host and share. Creating the correct host access on the destination system ahead of time reduces downtime in an event of a disaster.

To resume operations on the Boston system, when it is available:

1. From the New York system, select the **Reprotect** option, which resumes the replication session in the reverse direction.
2. After the systems are synchronized, select the **Planned Failover** option on the New York system.
3. Select the checkbox to auto-reprotect the system after failing over. Or, after the failover is complete, on the Boston system, select **Reprotect**.

NOTE: To replicate data from the destination to the source with the reprotect operation, ensure that there is a replication policy on the destination system that has a replication rule pointing to the source system. For example, if the replication session is from a site in Boston to a site in New York, the replication policy on the target storage resource in New York must point to Boston.