



Creating, monitoring, and troubleshooting protection relationships

Active IQ Unified Manager 9.7

NetApp
April 17, 2024

This PDF was generated from <https://docs.netapp.com/us-en/active-iq-unified-manager-97/data-protection/concept-types-of-snapmirror-protection.html> on April 17, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Creating, monitoring, and troubleshooting protection relationships 1
 - Types of SnapMirror protection 1
 - Setting up protection relationships in Unified Manager 2
 - Performing a protection relationship failover and failback 9
 - Resolving a protection job failure 13
 - Resolving lag issues 16

Creating, monitoring, and troubleshooting protection relationships

Unified Manager enables you to create protection relationships, to monitor and troubleshoot mirror protection and backup vault protection of data stored on managed clusters, and to restore data when it is overwritten or lost.

Types of SnapMirror protection

Depending on the deployment of your data storage topology, Unified Manager enables you to configure multiple types of SnapMirror protection relationships. All variations of SnapMirror protection offer failover disaster recovery protection, but offer differing capabilities in performance, version flexibility, and multiple backup copy protection.

Traditional SnapMirror Asynchronous protection relationships

Traditional SnapMirror Asynchronous protection provides block replication mirror protection between source and destination volumes.

In traditional SnapMirror relationships, mirror operations execute faster than they would in alternative SnapMirror relationships because the mirror operation is based on block replication. However, traditional SnapMirror protection requires that the destination volume run under the same or later minor version of ONTAP software as the source volume within the same major release (for example, version 8.x to 8.x, or 9.x to 9.x).

SnapMirror Asynchronous protection with version-flexible replication

SnapMirror Asynchronous protection with version-flexible replication provides logical replication mirror protection between source and destination volumes, even if those volumes are running under different versions of ONTAP 8.3 or later software (for example, version 8.3 to 8.3.1, or 8.3 to 9.1, or 9.2.2 to 9.2).

In SnapMirror relationships with version-flexible replication, mirror operations do not execute as quickly as they would in traditional SnapMirror relationships.

Because of slower execution, SnapMirror with version-flexible replication protection is not suitable to implement in either of the following circumstances:

- The source object contains more than 10 million files to protect.
- The recovery point objective for the protected data is two hours or less. (That is, the destination must always contain mirrored, recoverable data that is no more than two hours older than data at the source.)

In either of the listed circumstances, the faster block-replication based execution of default SnapMirror protection is required.

SnapMirror Asynchronous protection with version-flexible replication and backup option

SnapMirror Asynchronous protection with version-flexible replication and backup option provides mirror protection between source and destination volumes and the capability to store multiple copies of the mirrored data at the destination.

The storage administrator can specify which Snapshot copies are mirrored from source to destination and can also specify how long to retain those copies at the destination, even if they are deleted at the source.

In SnapMirror relationships with version-flexible replication and backup option, mirror operations do not execute as quickly as they would in traditional SnapMirror relationships.

SnapMirror Unified Replication (mirror and vault)

SnapMirror unified replication allows you to configure disaster recovery and archiving on the same destination volume. As with SnapMirror, unified data protection performs a baseline transfer the first time you invoke it. A baseline transfer under the default unified data protection policy “MirrorAndVault” makes a Snapshot copy of the source volume, then transfers that copy and the data blocks it references to the destination volume. Like SnapVault, unified data protection does not include older Snapshot copies in the baseline.

SnapMirror Synchronous protection with strict synchronization

SnapMirror Synchronous protection with “strict” synchronization ensures that the primary and secondary volumes are always a true copy of each other. If a replication failure occurs when attempting to write data to the secondary volume, then the client I/O to the primary volume is disrupted.

SnapMirror Synchronous protection with regular synchronization

SnapMirror Synchronous protection with “regular” synchronization does not require that the primary and secondary volume are always a true copy of each other; thereby ensuring availability of the primary volume. If a replication failure occurs when attempting to write data to the secondary volume, the primary and secondary volumes fall out of sync and client I/O will continue to the primary volume.



The Restore button and the Relationship operation buttons are not available when monitoring synchronous protection relationships from the Health: All Volumes view or the Volume / Health details page.

Setting up protection relationships in Unified Manager

There are several steps that you must perform to use Unified Manager and OnCommand Workflow Automation to set up SnapMirror and SnapVault relationships to protect your data.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have established peer relationships between two clusters or two storage virtual machines (SVMs).
- OnCommand Workflow Automation must be integrated with Unified Manager:
 - [Set up OnCommand Workflow Automation](#)
 - [Verifying Unified Manager data source caching in Workflow Automation](#)

Steps

1. Depending on the type of protection relationship you want to create, do one of the following:

- [Create a SnapMirror protection relationship.](#)
 - [Create a SnapVault protection relationship.](#)
2. If you want to create a policy for the relationship, depending on the relationship type you are creating, do one of the following:
 - [Create a SnapVault policy.](#)
 - [Create a SnapMirror policy.](#)
 3. [Create a SnapMirror or SnapVault schedule.](#)

Configuring a connection between Workflow Automation and Unified Manager

You can configure a secure connection between OnCommand Workflow Automation (WFA) and Unified Manager. Connecting to Workflow Automation enables you to use protection features such as SnapMirror and SnapVault configuration workflows, as well as commands for managing SnapMirror relationships.

Before you begin

- The installed version of Workflow Automation must be 5.1 or greater.



The “WFA pack for managing Clustered Data ONTAP” is included in WFA 5.1 so there is no need to download this pack from the NetAppStorage Automation Store and install it separately onto your WFA server as was required in the past. [WFA pack for managing ONTAP](#)

- You must have the name of the database user that you created in Unified Manager to support WFA and Unified Manager connections.

This database user must have been assigned the Integration Schema user role.

- You must be assigned either the Administrator role or the Architect role in Workflow Automation.
- You must have the host address, port number 443, user name, and password for the Workflow Automation setup.
- You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **General > Workflow Automation**.
2. In the **Database User** area of the **Workflow Automation** page, select the name, and enter the password for the database user that you created to support Unified Manager and Workflow Automation connections.
3. In the **Workflow Automation Credentials** area of the page, enter the host name or IP address (IPv4 or IPv6), and the user name and password for the Workflow Automation setup.

You must use the Unified Manager server port (port 443).

4. Click **Save**.
5. If you use a self-signed certificate, click **Yes** to authorize the security certificate.

The Workflow Automation page displays.

6. Click **Yes** to reload the web UI, and add the Workflow Automation features.

Related information

[NetApp Documentation: OnCommand Workflow Automation \(current releases\)](#)

Verifying Unified Manager data source caching in Workflow Automation

You can determine whether Unified Manager data source caching is working correctly by checking if data source acquisition is successful in Workflow Automation. You might do this when you integrate Workflow Automation with Unified Manager to ensure that Workflow Automation functionality is available after the integration.

Before you begin

You must be assigned either the Administrator role or the Architect role in Workflow Automation to perform this task.

Steps

1. From the Workflow Automation UI, select **Execution > Data Sources**.
2. Right-click the name of the Unified Manager data source, and then select **Acquire Now**.
3. Verify that the acquisition succeeds without errors.

Acquisition errors must be resolved for Workflow Automation integration with Unified Manager to succeed.

What happens when OnCommand Workflow Automation is reinstalled or upgraded

Before reinstalling or upgrading OnCommand Workflow Automation, you must first remove the connection between OnCommand Workflow Automation and Unified Manager, and ensure that all OnCommand Workflow Automation currently running or scheduled jobs are stopped.

You must also manually delete Unified Manager from OnCommand Workflow Automation.

After you reinstall or upgrade OnCommand Workflow Automation, you must set up the connection with Unified Manager again.

Removing OnCommand Workflow Automation setup from Unified Manager

You can remove the OnCommand Workflow Automation setup from Unified Manager when you no longer want to use Workflow Automation.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **General > Workflow Automation** in the left Setup menu.

2. In the **Workflow Automation** page, click **Remove Setup**.

Creating a SnapMirror protection relationship from the Volume / Health details page

You can use the Volume / Health details page to create a SnapMirror relationship so that data replication is enabled for protection purposes. SnapMirror replication enables you to restore data from the destination volume in the event of data loss on the source.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

About this task

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have an intercluster relationship and the destination cluster has not yet been discovered

You can perform up to 10 protection jobs simultaneously with no performance impact. You might experience some performance impact when you run between 11 and 30 jobs simultaneously. Running more than 30 jobs simultaneously is not recommended.

Steps

1. In the **Protection** tab of the **Volume / Health** details page, right-click in the topology view the name of a volume that you want to protect.
2. Select **Protect > SnapMirror** from the menu.

The Configure Protection dialog box is displayed.

3. Click **SnapMirror** to view the **SnapMirror** tab and to configure the destination information.
4. Click **Advanced** to set the space guarantee, as needed, and then click **Apply**.
5. Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.
6. Click **Apply**.

You are returned to the Volume / Health details page.

7. Click the protection configuration job link at the top of the **Volume / Health** details page.

The job's tasks and details are displayed in the Job details page.

8. In the **Job** details page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
9. When the job tasks are complete, click **Back** on your browser to return to the **Volume / Health** details page.

The new relationship is displayed in the Volume / Health details page topology view.

Results

Depending on the destination SVM you specified during configuration or on the options you enabled in your Advanced settings, the resulting SnapMirror relationship might be one of several possible variations:

- If you specified a destination SVM that runs under the same or a newer version of ONTAP compared to that of the source volume, a block-replication-based SnapMirror relationship is the default result.
- If you specified a destination SVM that runs under the same or a newer version of ONTAP (version 8.3 or higher) compared to that of the source volume, but you enabled version-flexible replication in the Advanced settings, a SnapMirror relationship with version-flexible replication is the result.
- If you specified a destination SVM that runs under an earlier version of ONTAP 8.3, or a version that is higher than that of the source volume and the earlier version supports version-flexible replication, a SnapMirror relationship with version-flexible replication is the automatic result.

Creating a SnapVault protection relationship from the Volume / Health details page

You can create a SnapVault relationship using the Volume / Health details page so that data backups are enabled for protection purposes on volumes.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation to perform this task.

About this task

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have an intercluster relationship and the destination cluster has not yet been discovered

Steps

1. In the **Protection** tab of the **Volume / Health** details page, right-click a volume in the topology view that you want to protect.
2. Select **Protect > SnapVault** from the menu.

The Configure Protection dialog box is launched.

3. Click **SnapVault** to view the **SnapVault** tab and to configure the secondary resource information.
4. Click **Advanced** to set deduplication, compression, autogrow, and space guarantee as needed, and then click **Apply**.
5. Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.
6. Click **Apply**.

You are returned to the Volume / Health details page.

7. Click the protection configuration job link at the top of the **Volume / Health** details page.

The Job details page is displayed.

8. Click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.

When the job tasks are complete, the new relationships are displayed in the Volume / Health details page topology view.

Creating a SnapVault policy to maximize transfer efficiency

You can create a new SnapVault policy to set the priority for a SnapVault transfer. You use policies to maximize the efficiency of transfers from the primary to the secondary in a protection relationship.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.
- You must have already completed Destination Information area in the Configure Protection dialog box.

Steps

1. From the **SnapVault** tab of the **Configure Protection** dialog box, click the **Create Policy** link in the **Relationship Settings** area.

The SnapVault tab is displayed.

2. In the **Policy Name** field, type the name that you want to give the policy.
3. In the **Transfer Priority** field, select the transfer priority that you want to assign to the policy.
4. In the **Comment** field, enter a comment for the policy.
5. In the **Replication Label** area, add or edit a replication label, as necessary.
6. Click **Create**.

The new policy is displayed in the Create Policy drop-down list.

Creating a SnapMirror policy to maximize transfer efficiency

You can create a SnapMirror policy to specify the SnapMirror transfer priority for protection relationships. SnapMirror policies enable you to maximize transfer efficiency from the source to the destination by assigning priorities so that lower-priority transfers are scheduled to run after normal-priority transfers.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.
- This task assumes that you have already completed the Destination Information area in the Configure Protection dialog box.

Steps

1. From the **SnapMirror** tab of the **Configure Protection** dialog box, click the **Create Policy** link in the **Relationship Settings** area.

The Create SnapMirror Policy dialog box is displayed.

2. In the **Policy Name** field, type a name you want to give the policy.
3. In the **Transfer Priority** field, select the transfer priority you want to assign to the policy.
4. In the **Comment** field, enter an optional comment for the policy.
5. Click **Create**.

The new policy is displayed in the SnapMirror Policy drop-down list.

Creating SnapMirror and SnapVault schedules

You can create basic or advanced SnapMirror and SnapVault schedules to enable automatic data protection transfers on a source or primary volume so that transfers take place more frequently or less frequently, depending on how often the data changes on your volumes.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have already completed the Destination Information area in the Configure Protection dialog box.
- You must have set up Workflow Automation to perform this task.

Steps

1. From the **SnapMirror** tab or **SnapVault** tab of the **Configure Protection** dialog box, click the **Create Schedule** link in the **Relationship Settings** area.

The Create Schedule dialog box is displayed.

2. In the **Schedule Name** field, type the name you want to give to the schedule.
3. Select one of the following:
 - **Basic**
Select if you want to create a basic interval-style schedule.
 - **Advanced**
Select if you want to create a cron-style schedule.
4. Click **Create**.

The new schedule is displayed in the SnapMirror Schedule or SnapVault Schedule drop-down list.

Performing a protection relationship failover and failback

When a source volume in your protection relationship is disabled because of a hardware failure or a disaster, you can use the protection relationship features in Unified Manager to make the protection destination read/write accessible and fail over to that volume until the source is online again; then, you can fail back to the original source when it is available to serve data.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up OnCommand Workflow Automation to perform this operation.

Steps


1. Break the SnapMirror relationship.

You must break the relationship before you can convert the destination from a data protection volume to a read/write volume, and before you can reverse the relationship.

2. Reverse the protection relationship.

When the original source volume is available again, you might decide to reestablish the original protection relationship by restoring the source volume. Before you can restore the source, you must synchronize it with the data written to the former destination. You use the reverse resync operation to create a new protection relationship by reversing the roles of the original relationship and synchronizing the source volume with the former destination. A new baseline Snapshot copy is created for the new relationship.

The reversed relationship looks similar to a cascaded relationship:



3. Break the reversed SnapMirror relationship.

When the original source volume is resynchronized and can again serve data, use the break operation to break the reversed relationship.

4. Remove the relationship.

When the reversed relationship is no longer required, you should remove that relationship before

reestablishing the original relationship.

5. [Resynchronize the relationship](#).

Use the resynchronize operation to synchronize data from the source to the destination and to reestablish the original relationship.

Breaking a SnapMirror relationship from the Volume / Health details page

You can break a protection relationship from the Volume / Health details page and stop data transfers between a source and destination volume in a SnapMirror relationship. You might break a relationship when you want to migrate data, for disaster recovery, or for application testing. The destination volume is changed to a read-write volume. You cannot break a SnapVault relationship.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

Steps

1. In the **Protection** tab of the **Volume / Health** details page, select from the topology the SnapMirror relationship you want to break.
2. Right-click the destination and select **Break** from the menu.

The Break Relationship dialog box is displayed.

3. Click **Continue** to break the relationship.
4. In the topology, verify that the relationship is broken.

Reversing protection relationships from the Volume / Health details page

When a disaster disables the source volume in your protection relationship, you can use the destination volume to serve data by converting it to read/write while you repair or replace the source. When the source is again available to receive data, you can use the reverse resynchronization operation to establish the relationship in the reverse direction, synchronizing the data on the source with the data on the read/write destination.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.
- The relationship must not be a SnapVault relationship.
- A protection relationship must already exist.
- The protection relationship must be broken.
- Both the source and destination must be online.

- The source must not be the destination of another data protection volume.

About this task

- When you perform this task, data on the source that is newer than the data on the common Snapshot copy is deleted.
- Policies and schedules created on the reverse resynchronization relationship are the same as those on the original protection relationship.

If policies and schedules do not exist, they are created.

Steps

1. From the **Protection** tab of the **Volume / Health** details page, locate in the topology the SnapMirror relationship on which you want to reverse the source and destination, and right-click it.

2. Select **Reverse Resync** from the menu.

The Reverse Resync dialog box is displayed.

3. Verify that the relationship displayed in the **Reverse Resync** dialog box is the one for which you want to perform the reverse resynchronization operation, and then click **Submit**.

The Reverse Resync dialog box is closed and a job link is displayed at the top of the Volume / Health details page.

4. Click **View Jobs** on the **Volume / Health** details page to track the status of each reverse resynchronization job.

A filtered list of jobs is displayed.

5. Click the Back arrow on your browser to return to the **Volume / Health** details page.

The reverse resynchronization operation is finished when all job tasks are completed successfully.

Removing a protection relationship from the Volume / Health details page

You can remove a protection relationship to permanently delete an existing relationship between the selected source and destination: for example, when you want to create a relationship using a different destination. This operation removes all metadata and cannot be undone.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

Steps

1. In the **Protection** tab of the **Volume / Health** details page, select from the topology the SnapMirror relationship you want to remove.

2. Right-click the name of the destination and select **Remove** from the menu.

The Remove Relationship dialog box is displayed.

3. Click **Continue** to remove the relationship.

The relationship is removed from the Volume / Health details page.

Resynchronizing protection relationships from the Volume / Health details page

You can resynchronize data on a SnapMirror or SnapVault relationship that was broken and then the destination was made read/write so that data on the source matches the data on the destination. You might also resynchronize when a required common Snapshot copy on the source volume is deleted causing SnapMirror or SnapVault updates to fail.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up OnCommand Workflow Automation.

Steps

1. From the **Protection** tab of the **Volume / Health** details page, locate in the topology the protection relationship that you want to resynchronize and right-click it.
2. Select **Resynchronize** from the menu.

Alternatively, from the **Actions** menu, select **Relationship > Resynchronize** to resynchronize the relationship for which you are currently viewing the details.

The Resynchronize dialog box is displayed.

3. In the **Resynchronization Options** tab, select a transfer priority and the maximum transfer rate.
4. Click **Source Snapshot Copies**; then, in the **Snapshot Copy** column, click **Default**.

The Select Source Snapshot Copy dialog box is displayed.

5. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
6. Click **Submit**.

You are returned to the Resynchronize dialog box.

7. If you selected more than one source to resynchronize, click **Default** for the next source for which you want to specify an existing Snapshot copy.
8. Click **Submit** to begin the resynchronization job.

The resynchronization job is started, you are returned to the Volume / Health details page and a jobs link is displayed at the top of the page.

9. Click **View Jobs** on the **Volume / Health** details page to track the status of each resynchronization job.

A filtered list of jobs is displayed.

10. Click the Back arrow on your browser to return to the **Volume / Health** details page.

The resynchronization job is finished when all job tasks successfully complete.

Resolving a protection job failure

This workflow provides an example of how you might identify and resolve a protection job failure from the Unified Manager dashboard.

Before you begin

Because some tasks in this workflow require that you log in using the Administrator role, you must be familiar with the roles required to use various functionality.

About this task

In this scenario, you access the Dashboard page to see if there are any issues with your protection jobs. In the Protection Incident area, you notice that there is a Job Terminated incident, showing a Protection Job Failed error on a volume. You investigate this error to determine the possible cause and potential resolution.

Steps

1. In the **Protection Incidents** panel of the Dashboard **Unresolved Incidents and Risks** area, you click the **Protection job failed** event.



The linked text for the event is written in the form `object_name:/object_name - Error Name`, such as `cluster2_src_svm:/cluster2_src_vol2 - Protection Job Failed`.

The Event details page for the failed protection job displays.

2. Review the error message in the Cause field of the **Summary** area to determine the problem and evaluate potential corrective actions.

See [Identifying the problem and performing corrective actions for a failed protection job](#).

Identifying the problem and performing corrective actions for a failed protection job

You review the job failure error message in the Cause field on the Event details page and determine that the job failed because of a Snapshot copy error. You then proceed to the Volume / Health details page to gather more information.

Before you begin

You must have the Application Administrator role.

About this task

The error message provided in the Cause field on the Event details page contains the following text about the

failed job:

```
Protection Job Failed. Reason: (Transfer operation for
relationship 'cluster2_src_svm:cluster2_src_vol2->cluster3_dst_svm:
managed_svc2_vol3' ended unsuccessfully. Last error reported by
Data ONTAP: Failed to create Snapshot copy 0426cluster2_src_vol2snap
on volume cluster2_src_svm:cluster2_src_vol2. (CSM: An operation
failed due to an ONC RPC failure..))
*Job Details*
```

This message provides the following information:

- A backup or mirror job did not complete successfully.

The job involved a protection relationship between the source volume `cluster2_src_vol2` on the virtual server `cluster2_src_svm` and the destination volume `managed_svc2_vol3` on the virtual server named `cluster3_dst_svm`.

- A Snapshot copy job failed for `0426cluster2_src_vol2snap` on the source volume `cluster2_src_svm:/cluster2_src_vol2`.

In this scenario, you can identify the cause and potential corrective actions of the job failure. However, resolving the failure requires that you access either the System Manager web UI or the ONTAP CLI commands.

Steps

1. You review the error message and determine that a Snapshot copy job failed on the source volume, indicating that there is probably a problem with your source volume.

Optionally, you could click the **Job Details** link at the end of the error message, but for the purposes of this scenario, you choose not to do that.

2. You decide that you want to try to resolve the event, so you do the following:
 - a. Click the **Assign To** button and select **Me** from the menu.
 - b. Click the **Acknowledge** button so that you do not continue to receive repeat alert notifications, if alerts were set for the event.
 - c. Optionally, you can also add notes about the event.
3. Click the **Source** field in the **Summary** pane to see details about the source volume.


The **Source** field contains the name of the source object: in this case, the volume on which the Snapshot copy job was scheduled.

The Volume / Health details page displays for `cluster2_src_vol2`, showing the content of the Protection tab.

4. Looking at the protection topology graph, you see an error icon associated with the first volume in the topology, which is the source volume for the SnapMirror relationship.

You also see the horizontal bars in the source volume icon, indicating the warning and error thresholds set

for that volume.



5. You place your cursor over the error icon to see the pop-up dialog box that displays the threshold settings and see that the volume has exceeded the error threshold, indicating a capacity issue.
6. Click the **Capacity** tab.

Capacity information about volume `cluster2_src_vol2` displays.

7. In the **Capacity** panel, you see that there is an error icon in the bar graph, again indicating that the volume capacity has surpassed the threshold level set for the volume.
8. Below the capacity graph, you see that volume autogrow has been disabled and that a volume space guarantee has been set.

You could decide to enable autogrow, but for the purposes of this scenario, you decide to investigate further before making a decision about how to resolve the capacity problem.

9. You scroll down to the **Events** list and see that Protection Job Failed, Volume Days Until Full, and Volume Space Full events were generated.
10. In the **Events** list, you click the **Volume Space Full** event to get more information, having decided that this event seems most relevant to your capacity issue.

The Event details page displays the Volume Space Full event for the source volume.

11. In the **Summary** area, you read the Cause field for the event: The full threshold set at 90% is breached. 45.38 MB (95.54%) of 47.50 MB is used.
12. Below the **Summary** area, you see Suggested Corrective Actions.



The Suggested Corrective Actions display only for some events, so you do not see this area for all types of events.

You click through the list of suggested actions that you might perform to resolve the Volume Space Full event:

- Enable autogrow on this volume.

- Resize the volume.
- Enable and run deduplication on this volume.
- Enable and run compression on this volume.

13. You decide to enable autogrow on the volume, but to do so, you must determine the available free space on the parent aggregate and the current volume growth rate:

- Look at the parent aggregate, `cluster2_src_aggr1`, in the **Related Devices** pane.



You can click the name of the aggregate to get further details about the aggregate.

You determine that the aggregate has sufficient space to enable volume autogrow.

- At the top of the page, look at the icon indicating a critical incident and review the text below the icon.

You determine that "Days to Full: Less than a day | Daily Growth Rate: 5.4%".

14. Go to System Manager or access the ONTAP CLI to enable the `volume autogrow` option.



Make note of the names of the volume and aggregate so you have them available when enabling autogrow.

15. After resolving the capacity issue, return to the Unified Manager **Event** details page and mark the event as resolved.

Resolving lag issues

This workflow provides an example of how you might resolve a lag issue. In this scenario, you are an administrator or operator accessing the Unified Manager **Dashboard** page to see if there are any problems with your protection relationships and, if they exist, to find solutions.

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

In the **Dashboard** page, you look at the Unresolved Incidents and Risks area and see a SnapMirror Lag error in the Protection pane under Protection Risks.

Steps

- In the **Protection** pane on the **Dashboard** page, locate the SnapMirror relationship lag error and click it.

The Event details page for the lag error event is displayed.


- From the **Event** details page you can perform one or more of the following tasks:
 - Review the error message in the Cause field of the Summary area to determine if there is any suggested corrective action.

- Click the object name, in this case a volume, in the Source field of the Summary area to get details about the volume.
 - Look for notes that might have been added about this event.
 - Add a note to the event.
 - Assign the event to a specific user.
 - Acknowledge or resolve the event.
3. In this scenario, you click the object name (in this case, a volume) in the Source field of the **Summary** area to get details about the volume.

The Protection tab of the Volume / Health details page is displayed.

4. In the **Protection** tab, you look at the topology diagram.

You note that the volume with the lag error is the last volume in a three-volume SnapMirror cascade. The volume you selected is outlined in dark gray, and a double orange line from the source volume indicates a SnapMirror relationship error.



5. Click each of the volumes in the SnapMirror cascade.

As you select each volume, the protection information in the Summary, Topology, History, Events, Related Devices, and Related Alerts areas changes to display details relevant to the selected volume.

6. You look at the **Summary** area and position your cursor over the information icon in the **Update Schedule** field for each volume.

In this scenario, you note that the SnapMirror policy is DPDefault, and the SnapMirror schedule updates hourly at five minutes after the hour. You realize that all of the volumes in the relationship are attempting to complete a SnapMirror transfer at the same time.

7. To resolve the lag issue, you modify the schedules for two of the cascaded volumes so that each destination begins a SnapMirror transfer after its source has completed a transfer.

Copyright information

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

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

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

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

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

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

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

Trademark information

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