

Data protection and disaster recovery ONTAP 9

NetApp August 25, 2025

This PDF was generated from https://docs.netapp.com/us-en/ontap/peering/index.html on August 25, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Data protection and disaster recovery	1
Cluster and SVM peering	1
Learn about ONTAP cluster and SVM peering	1
Prepare for cluster and SVM peering	1
Configure intercluster LIFs	5
Configure peer relationships	18
Enable ONTAP cluster peering encryption on peer relationships	27
Remove ONTAP cluster peering encryption from peer relationships	27
Manage local snapshots	29
Learn about managing local ONTAP snapshots	29
Configure custom snapshot policies	29
Manage snapshots manually	33
Manage the snapshot reserve	36
Restore files from snapshots	40
SnapMirror volume replication	46
Learn about SnapMirror volume replication	46
Configure SnapMirror volume replication	72
Manage SnapMirror volume replication	92
Manage SnapMirror SVM replication	122
Learn about ONTAP SnapMirror SVM replication	122
Replicate SVM configurations	130
Serve data from a SnapMirror SVM DR destination	142
Reactivate the SnapMirror source SVM	147
Convert an ONTAP SnapMirror volume DR relationship to an SVM DR relationship	159
Delete an ONTAP SnapMirror SVM replication relationship	160
Manage SnapMirror root volume replication	161
Learn about ONTAP SnapMirror root volume replication	161
Create and initialize ONTAP load-sharing mirror relationships	162
Update an ONTAP load-sharing mirror relationship	163
Promote an ONTAP load-sharing mirror	164
Back up to the cloud.	165
Install an ONTAP SnapMirror cloud license	165
Back up data to the cloud using ONTAP SnapMirror	166
Back up data using BlueXP backup and recovery	169
Archive and compliance using SnapLock technology	172
Learn about ONTAP SnapLock	172
Configure SnapLock.	177
Manage WORM files	192
Move an ONTAP SnapLock volume	207
Lock an ONTAP snapshot for protection against ransomware attacks	208
Consistency groups	215
Learn about ONTAP consistency groups	215
Learn about ONTAP consistency group limits	220

Configure a single ONTAP consistency group	221
Configure a hierarchical ONTAP consistency group	225
Protect ONTAP consistency groups	229
Modify member volumes in an ONTAP consistency group	237
Modify ONTAP consistency group geometry	242
Modify ONTAP consistency group application and component tags	247
Clone an ONTAP consistency group	248
Delete an ONTAP consistency group	250
SnapMirror active sync.	251
Introduction	251
Plan	262
Configure	270
Manage SnapMirror active sync and protect data	310
Troubleshoot	327
ONTAP Mediator for MetroCluster and SnapMirror active sync	336
Learn about ONTAP Mediator	336
What's new in ONTAP Mediator	337
Install or upgrade	342
Manage ONTAP Mediator	388
Maintain the host OS for ONTAP Mediator	417
Learn about MetroCluster IP site management with ONTAP System Manager	418
Data protection using tape backup	419
Learn about tape backup of FlexVol volumes with ONTAP	419
Tape backup and restore workflow in ONTAP	419
Use cases for choosing a tape backup engine	420
Manage tape drives	421
About tape drives	426
Transfer data between storage systems	434
NDMP for FlexVol volumes	438
About NDMP for FlexGroup volumes	459
About NDMP with SnapLock volumes	459
Manage node-scoped NDMP mode for FlexVol volumes	459
Manage SVM-scoped NDMP mode for FlexVol volumes	461
About dump engine for FlexVol volumes	468
About SMTape engine for FlexVol volumes	479
Monitor tape backup and restore operations for FlexVol volumes	484
Error messages for tape backup and restore of FlexVol volumes	488
NDMP configuration	507
Learn about ONTAP NDMP configuration	507
Learn about ONTAP NDMP configuration workflow.	508
Prepare ONTAP NDMP configurations	509
Verify ONTAP NDMP tape device connections	512
Enable tape reservations for ONTAP NDMP backup operations	513
Configure SVM-scoped NDMP	514
Configure node-scoped NDMP	523

Configure backup applications for ONTAP NDMP configuration	. 528
Replication between NetApp Element software and ONTAP overview	. 528

Data protection and disaster recovery

Cluster and SVM peering

Learn about ONTAP cluster and SVM peering

You can create peer relationships between source and destination clusters and between source and destination storage virtual machines (SVMs). You must create peer relationships between these entities before you can replicate snapshots using SnapMirror.

ONTAP 9.3 offers enhancements that simplify the way you configure peer relationships between clusters and SVMs. The cluster and SVMs peering procedures are available for all ONTAP 9 versions. You should use the appropriate procedure for your version of ONTAP.

You perform the procedures using the command-line interface (CLI), not System Manager or an automated scripting tool.

Prepare for cluster and SVM peering

ONTAP peering basics

You must create *peer relationships* between source and destination clusters and between source and destination SVMs before you can replicate snapshots using SnapMirror. A peer relationship defines network connections that enable clusters and SVMs to exchange data securely.

Clusters and SVMs in peer relationships communicate over the intercluster network using *intercluster logical interfaces (LIFs)*. An intercluster LIF is a LIF that supports the "intercluster-core" network interface service and is typically created using the "default-intercluster" network interface service policy. You must create intercluster LIFs on every node in the clusters being peered.

Intercluster LIFs use routes that belong to the system SVM to which they are assigned. ONTAP automatically creates a system SVM for cluster-level communications within an IPspace.

Fan-out and cascade topologies are both supported. In a cascade topology, you need only create intercluster networks between the primary and secondary clusters and between the secondary and tertiary clusters. You need not create an intercluster network between the primary and the tertiary cluster.

It is possible (but not advisable) for an administrator to remove the intercluster-core service from the default-intercluster service policy. If this occurs, LIFs created using "default-intercluster" will not actually be intercluster LIFs. To confirm that the default-intercluster service policy contains the intercluster-core service, use the following command:

(

network interface service-policy show -policy default-intercluster

Learn more about network interface service-policy show in the ONTAP command reference.

ONTAP peering prerequisites

Before you set up cluster peering, you should confirm that the connectivity, port, IP address, subnet, firewall, and cluster-naming requirements are met.

Beginning with ONTAP 9.6, Cluster Peering provides TLS 1.2 AES-256 GCM encryption support for data replication by default. The default security ciphers ("PSK-AES256-GCM-SHA384") are required for Cluster Peering to work even if encryption is disabled.



Beginning with ONTAP 9.11.1, DHE-PSK security ciphers are available by default.

Beginning with ONTAP 9.15.1, Cluster Peering provides TLS 1.3 encryption support for data replication by default.

Connectivity requirements

Every intercluster LIF on the local cluster must be able to communicate with every intercluster LIF on the remote cluster.

Although it is not required, it is typically simpler to configure the IP addresses used for intercluster LIFs in the same subnet. The IP addresses can reside in the same subnet as data LIFs, or in a different subnet. The subnet used in each cluster must meet the following requirements:

- The subnet must belong to the broadcast domain that contains the ports that are used for intercluster communication.
- The subnet must have enough IP addresses available to allocate to one intercluster LIF per node.

For example, in a four-node cluster, the subnet used for intercluster communication must have four available IP addresses.

Each node must have an intercluster LIF with an IP address on the intercluster network.

Intercluster LIFs can have an IPv4 address or an IPv6 address.



ONTAP enables you to migrate your peering networks from IPv4 to IPv6 by optionally allowing both protocols to be present simultaneously on the intercluster LIFs. In earlier releases, all intercluster relationships for an entire cluster were either IPv4 or IPv6. This meant that changing protocols was a potentially disruptive event.

Port requirements

You can use dedicated ports for intercluster communication, or share ports used by the data network. Ports must meet the following requirements:

• All ports that are used to communicate with a given remote cluster must be in the same IPspace.

You can use multiple IPspaces to peer with multiple clusters. Pair-wise full-mesh connectivity is required only within an IPspace.

• The broadcast domain that is used for intercluster communication must include at least two ports per node so that intercluster communication can fail over from one port to another port.

Ports added to a broadcast domain can be physical network ports, VLANs, or interface groups (ifgrps).

- All ports must be cabled.
- All ports must be in a healthy state.
- The MTU settings of the ports must be consistent.

Firewall requirements

()

Beginning with ONTAP 9.10.1, firewall policies are deprecated and wholly replaced with LIF service policies. For more information, see Configure firewall policies for LIFs.

Firewalls and the intercluster firewall policy must allow the following protocols:

- Bidirectional ICMP traffic
- Bidirectional initiated TCP traffic to the IP addresses of all the intercluster LIFs over ports 11104 and 11105
- Bidirectional HTTPS between the intercluster LIFs

Although HTTPS is not required when you set up cluster peering using the CLI, HTTPS is required later if you use System Manager to configure data protection.

The default intercluster firewall policy allows access through the HTTPS protocol and from all IP addresses (0.0.0.0/0). You can modify or replace the policy if necessary.

Cluster requirement

Clusters must meet the following requirement:

• A cluster cannot be in a peer relationship with more than 255 clusters.

Use shared or dedicated ONTAP ports

You can use dedicated ports for intercluster communication, or share ports used by the data network. In deciding whether to share ports, you need to consider network bandwidth, the replication interval, and port availability.



You can share ports on one peered cluster while using dedicated ports on the other.

Network bandwidth

If you have a high-speed network, such as 10 GbE, you might have enough local LAN bandwidth to perform replication using the same 10 GbE ports used for data access.

Even then, you should compare your available WAN bandwidth to your LAN bandwidth. If the available WAN bandwidth is significantly less than 10 GbE, you might need to use dedicated ports.



The one exception to this rule might be when all or many nodes in the cluster replicate data, in which case bandwidth utilization is typically spread across nodes.

If you are not using dedicated ports, the maximum transmission unit (MTU) size of the replication network should typically be the same as the MTU size of the data network.

Replication interval

If replication takes place in off-peak hours, you should be able to use data ports for replication even without a 10-GbE LAN connection.

If replication takes place during normal business hours, you need to consider the amount of data that will be replicated and whether it requires so much bandwidth that it could cause contention with data protocols. If network utilization by data protocols (SMB, NFS, iSCSI) is above 50%, you should use dedicated ports for intercluster communication, to allow for non-degraded performance if node failover occurs.

Port availability

If you determine that replication traffic is interfering with data traffic, you can migrate intercluster LIFs to any other intercluster-capable shared port on the same node.

You can also dedicate VLAN ports for replication. The bandwidth of the port is shared between all VLANs and the base port.

Use custom ONTAP IPspaces to isolate replication traffic

You can use custom IPspaces to separate the interactions that a cluster has with its peers. Called *designated intercluster connectivity*, this configuration allows service providers to isolate replication traffic in multitenant environments.

Suppose, for example, that you want replication traffic between Cluster A and Cluster B to be separated from replication traffic between Cluster A and Cluster C. To accomplish this, you can create two IPspaces on Cluster A.

One IPspace contains the intercluster LIFs that you use to communicate with Cluster B. The other contains the intercluster LIFs that you use to communicate with Cluster C, as shown in the following illustration.



Related information

• Learn about ONTAP IPspace configuration

Configure intercluster LIFs

Configure ONTAP intercluster LIFs on shared data ports

You can configure intercluster LIFs on ports shared with the data network. Doing so reduces the number of ports you need for intercluster networking.

Steps

1. List the ports in the cluster:

network port show

Learn more about network port show in the ONTAP command reference.

The following example shows the network ports in cluster01:

cluster01::> network port show						
						Speed
(Mbps)						
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper
cluste	r01-01					
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	eOc	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
cluste	r01-02					
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	eOc	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000

2. Create intercluster LIFs on either an admin SVM (Default IPspace) or a system SVM (custom IPspace):

Option	Description
In ONTAP 9.6 and later:	<pre>network interface create -vserver system_SVM -lif LIF_name -service -policy default-intercluster -home -node node -home-port port -address port_IP -netmask netmask</pre>

Option	Description
In ONTAP 9.5 and earlier:	<pre>network interface create -vserver system_SVM -lif LIF_name -role intercluster -home-node node -home -port port -address port_IP -netmask netmask</pre>

Learn more about network interface create in the ONTAP command reference.

The following example creates intercluster LIFs cluster01_icl01 and cluster01_icl02:

```
cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl01 -service-
policy default-intercluster -home-node cluster01-01 -home-port e0c
-address 192.168.1.201
-netmask 255.255.255.0
cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl02 -service-
policy default-intercluster -home-node cluster01-02 -home-port e0c
-address 192.168.1.202
-netmask 255.255.255.0
```

3. Verify that the intercluster LIFs were created:

Option	Description
In ONTAP 9.6 and later:	network interface show -service-policy default-intercluster
In ONTAP 9.5 and earlier:	network interface show -role intercluster

Learn more about network interface show in the ONTAP command reference.

cluster01::> network interface show -service-policy default-intercluster Logical Status Network Current Current Is Interface Admin/Oper Address/Mask Node Port Vserver Home _____ ____ _____ ____ cluster01 cluster01 icl01 up/up 192.168.1.201/24 cluster01-01 e0c true cluster01 icl02 up/up 192.168.1.202/24 cluster01-02 e0c true

4. Verify that the intercluster LIFs are redundant:

Option	Description
In ONTAP 9.6 and later:	network interface show -service-policy default-intercluster -failover
In ONTAP 9.5 and earlier:	network interface show -role intercluster -failover

Learn more about network interface show in the ONTAP command reference.

The following example shows that the intercluster LIFs <code>cluster01_icl01</code> and <code>cluster01_icl02</code> on the <code>e0c</code> port will fail over to the <code>e0d</code> port.

<pre>cluster01::> network interface show -service-policy default-intercluster -failover</pre>				
	Logical	Home	Failover	Failover
Vserver	Interface	Node:Port	Policy	Group
cluster03	1			
	cluster01_icl01	cluster01-01:e0c lo	ocal-only	
192.168.3	1.201/24			
		Failover Targets:	cluster01-01:e0c	,
			cluster01-01:e0d	
192.168.3	cluster01_icl02 1.201/24	cluster01-02:e0c lo	ocal-only	
		Failover Targets:	cluster01-02:e0c	,
			cluster01-02:e0d	

Configure ONTAP intercluster LIFs on dedicated ports

You can configure intercluster LIFs on dedicated ports. Doing so typically increases the available bandwidth for replication traffic.

Steps

1. List the ports in the cluster:

network port show

Learn more about network port show in the ONTAP command reference.

The following example shows the network ports in cluster01:

cluster01::> network port show						
						Speed
(Mbps)						
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper
cluste	r01-01					
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	eOc	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	eOf	Default	Default	up	1500	auto/1000
cluste	r01-02					
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	eOc	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	eOf	Default	Default	up	1500	auto/1000

2. Determine which ports are available to dedicate to intercluster communication:

network interface show -fields home-port, curr-port

Learn more about network interface show in the ONTAP command reference.

The following example shows that ports elle and ell have not been assigned LIFs:

cluster01::> network interfa	ce show -f	ields home-port,curr-port
vserver lif	home-port	curr-port
Cluster cluster01-01_clus1	e0a	e0a
Cluster cluster01-01_clus2	e0b	e0b
Cluster cluster01-02_clus1	e0a	e0a
Cluster cluster01-02_clus2	e0b	e0b
cluster01		
cluster_mgmt	eOc	eOc
cluster01		
cluster01-01_mgmt1	eOc	eOc
cluster01		
cluster01-02_mgmt1	e0c	e0c

3. Create a failover group for the dedicated ports:

```
network interface failover-groups create -vserver system_SVM -failover-group
failover_group -targets physical _or_logical_ports
```

The following example assigns ports ele and elf to the failover group interclusterld on the system SVM clusterll:

```
cluster01::> network interface failover-groups create -vserver cluster01
-failover-group
intercluster01 -targets
cluster01-01:e0e,cluster01-01:e0f,cluster01-02:e0e,cluster01-02:e0f
```

4. Verify that the failover group was created:

network interface failover-groups show

Learn more about network interface failover-groups show in the ONTAP command reference.

cluster01::> network interface failover-groups show Failover Group Vserver Targets _____ _____ _____ Cluster Cluster cluster01-01:e0a, cluster01-01:e0b, cluster01-02:e0a, cluster01-02:e0b cluster01 Default cluster01-01:e0c, cluster01-01:e0d, cluster01-02:e0c, cluster01-02:e0d, cluster01-01:e0e, cluster01-01:e0f cluster01-02:e0e, cluster01-02:e0f intercluster01 cluster01-01:e0e, cluster01-01:e0f cluster01-02:e0e, cluster01-02:e0f

5. Create intercluster LIFs on the system SVM and assign them to the failover group.

Option	Description
In ONTAP 9.6 and later:	<pre>network interface create -vserver system_SVM -lif LIF_name -service -policy default-intercluster -home -node node -home- port port -address port_IP -netmask netmask -failover -group failover_group</pre>
In ONTAP 9.5 and earlier:	<pre>network interface create -vserver system_SVM -lif LIF_name -role intercluster -home-node node -home -port port -address port_IP -netmask netmask -failover-group failover_group</pre>

Learn more about network interface create in the ONTAP command reference.

The following example creates intercluster LIFs <code>cluster01_icl01</code> and <code>cluster01_icl02</code> in the failover group intercluster01:

```
cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl01 -service-
policy default-intercluster -home-node cluster01-01 -home-port e0e
-address 192.168.1.201
-netmask 255.255.255.0 -failover-group intercluster01
cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl02 -service-
policy default-intercluster -home-node cluster01-02 -home-port e0e
-address 192.168.1.202
-netmask 255.255.255.0 -failover-group intercluster01
```

6. Verify that the intercluster LIFs were created:

Option	Description
In ONTAP 9.6 and later:	network interface show -service-policy default-intercluster
In ONTAP 9.5 and earlier:	network interface show -role intercluster

Learn more about network interface show in the ONTAP command reference.

```
cluster01::> network interface show -service-policy default-intercluster
        Logical Status Network
                                       Current
Current Is
Vserver Interface Admin/Oper Address/Mask Node
                                                Port
Home
_____ ___
cluster01
        cluster01 icl01
                 up/up 192.168.1.201/24 cluster01-01 e0e
true
        cluster01 icl02
                 up/up 192.168.1.202/24 cluster01-02 eOf
true
```

7. Verify that the intercluster LIFs are redundant:

Option	Description
In ONTAP 9.6 and later:	network interface show -service-policy default-intercluster -failover
In ONTAP 9.5 and earlier:	network interface show -role intercluster -failover

Learn more about network interface show in the ONTAP command reference.

The following example shows that the intercluster LIFs <code>cluster01_icl01</code> and <code>cluster01_icl02</code> on the SVMe0e port will fail over to the <code>e0f</code> port.

```
cluster01::> network interface show -service-policy default-intercluster
-failover
       Logical
                     Home
                                        Failover
                                                      Failover
Vserver Interface
                     Node:Port
                                        Policy
                                                       Group
        cluster01
        cluster01 icl01 cluster01-01:e0e local-only
intercluster01
                        Failover Targets: cluster01-01:e0e,
                                         cluster01-01:e0f
        cluster01 icl02 cluster01-02:e0e local-only
intercluster01
                        Failover Targets: cluster01-02:e0e,
                                         cluster01-02:e0f
```

Configure ONTAP intercluster LIFs in custom IPspaces

You can configure intercluster LIFs in custom IPspaces. Doing so allows you to isolate replication traffic in multitenant environments.

When you create a custom IPspace, the system creates a system storage virtual machine (SVM) to serve as a container for the system objects in that IPspace. You can use the new SVM as the container for any intercluster LIFs in the new IPspace. The new SVM has the same name as the custom IPspace.

Steps

1. List the ports in the cluster:

network port show

Learn more about network port show in the ONTAP command reference.

The following example shows the network ports in cluster01:

cluster01::> network port show						
						Speed
(Mbps)						
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper
cluste	r01-01					
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	eOc	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	eOf	Default	Default	up	1500	auto/1000
cluste	r01-02					
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	eOc	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	eOf	Default	Default	up	1500	auto/1000

2. Create custom IPspaces on the cluster:

network ipspace create -ipspace ipspace

The following example creates the custom IPspace ipspace-IC1:

cluster01::> network ipspace create -ipspace ipspace-IC1

Learn more about network ipspace create in the ONTAP command reference.

3. Determine which ports are available to dedicate to intercluster communication:

network interface show -fields home-port, curr-port

Learn more about network interface show in the ONTAP command reference.

The following example shows that ports elle and ell have not been assigned LIFs:

cluster01::> network interface show -fields home-port, curr-port			
vserver lif	home-p	port curr-port	
Cluster cluster01_clus1	e0a	e0a	
Cluster cluster01_clus2	e0b	e0b	
Cluster cluster02_clus1	e0a	e0a	
Cluster cluster02_clus2	e0b	e0b	
cluster01			
cluster_mgmt	eOc	eOc	
cluster01			
cluster01-01_mgmt	e0c	eOc	
cluster01			
cluster01-02_mgmt	e0c	eOc	

4. Remove the available ports from the default broadcast domain:

```
network port broadcast-domain remove-ports -broadcast-domain Default -ports
ports
```

A port cannot be in more than one broadcast domain at a time. Learn more about network port broadcast-domain remove-ports in the ONTAP command reference.

The following example removes ports elle and ellf from the default broadcast domain:

```
cluster01::> network port broadcast-domain remove-ports -broadcast
-domain Default -ports
cluster01-01:e0e,cluster01-01:e0f,cluster01-02:e0e,cluster01-02:e0f
```

5. Verify that the ports have been removed from the default broadcast domain:

network port show

Learn more about network port show in the ONTAP command reference.

The following example shows that ports e0e and e0f have been removed from the default broadcast domain:

cluster01::> network port show Speed (Mbps) Node Broadcast Domain Link Admin/Oper Port IPspace MTU _____ ____ _____ cluster01-01 e0a Cluster Cluster up 9000 auto/1000 e0b Cluster Cluster up 9000 auto/1000 e0c Default Default 1500 auto/1000 up e0d Default Default 1500 auto/1000 up e0e Default _ 1500 auto/1000 up eOf Default _ 1500 auto/1000 up Default auto/1000 eOg Default 1500 up cluster01-02 e0a Cluster Cluster 9000 auto/1000 up e0b Cluster 9000 auto/1000 Cluster up 1500 auto/1000 e0c Default Default up 1500 auto/1000 e0d Default Default up e0e Default 1500 auto/1000 _ up e0f Default _ up 1500 auto/1000 Default 1500 auto/1000 eOg Default up

6. Create a broadcast domain in the custom IPspace:

network port broadcast-domain create -ipspace ipspace -broadcast-domain broadcast domain -mtu MTU -ports ports

The following example creates the broadcast domain <code>ipspace-IC1-bd</code> in the IPspace <code>ipspace-IC1:</code>

```
cluster01::> network port broadcast-domain create -ipspace ipspace-IC1
-broadcast-domain
ipspace-IC1-bd -mtu 1500 -ports cluster01-01:e0e,cluster01-01:e0f,
cluster01-02:e0e,cluster01-02:e0f
```

7. Verify that the broadcast domain was created:

network port broadcast-domain show

Learn more about network port broadcast-domain show in the ONTAP command reference.

cluster01::> network port broadcast-domain show				
IPspace	Broadcast			Update
Name	Domain Name	MTU	Port List	Status Details
Cluster	Cluster	9000		
			cluster01-01:e0a	complete
			cluster01-01:e0b	complete
			cluster01-02:e0a	complete
			cluster01-02:e0b	complete
Default	Default	1500		
			cluster01-01:e0c	complete
			cluster01-01:e0d	complete
			cluster01-01:e0f	complete
			cluster01-01:e0g	complete
			cluster01-02:e0c	complete
			cluster01-02:e0d	complete
			cluster01-02:e0f	complete
			cluster01-02:e0g	complete
ipspace	-IC1			
	ipspace-IC1-b	d		
		1500		
			cluster01-01:e0e	complete
			cluster01-01:e0f	complete
			cluster01-02:e0e	complete
			cluster01-02:e0f	complete

8. Create intercluster LIFs on the system SVM and assign them to the broadcast domain:

Option	Description
In ONTAP 9.6 and later:	<pre>network interface create -vserver system_SVM -lif LIF_name -service -policy default-intercluster -home -node node -home-port port -address port_IP -netmask netmask</pre>
In ONTAP 9.5 and earlier:	<pre>network interface create -vserver system_SVM -lif LIF_name -role intercluster -home-node node -home -port port -address port_IP -netmask netmask</pre>

The LIF is created in the broadcast domain that the home port is assigned to. The broadcast domain has a default failover group with the same name as the broadcast domain. Learn more about network interface create in the ONTAP command reference.

The following example creates intercluster LIFs <code>cluster01_icl01</code> and <code>cluster01_icl02</code> in the broadcast domain <code>ipspace-IC1-bd</code>:

```
cluster01::> network interface create -vserver ipspace-IC1 -lif
cluster01_icl01 -service-
policy default-intercluster -home-node cluster01-01 -home-port e0e
-address 192.168.1.201
-netmask 255.255.255.0
cluster01::> network interface create -vserver ipspace-IC1 -lif
cluster01_icl02 -service-
policy default-intercluster -home-node cluster01-02 -home-port e0e
-address 192.168.1.202
-netmask 255.255.255.0
```

9. Verify that the intercluster LIFs were created:

Option	Description
In ONTAP 9.6 and later:	network interface show -service-policy default-intercluster
In ONTAP 9.5 and earlier:	network interface show -role intercluster

Learn more about network interface show in the ONTAP command reference.

```
cluster01::> network interface show -service-policy default-intercluster
         Logical Status Network
                                        Current
Current Is
Vserver Interface Admin/Oper Address/Mask Node
                                                 Port
Home
_____ ___
ipspace-IC1
         cluster01 icl01
                 up/up 192.168.1.201/24 cluster01-01 e0e
true
         cluster01 icl02
                 up/up
                        192.168.1.202/24 cluster01-02 eOf
true
```

10. Verify that the intercluster LIFs are redundant:

Option	Description
In ONTAP 9.6 and later:	network interface show -service-policy default-intercluster -failover
In ONTAP 9.5 and earlier:	network interface show -role intercluster -failover

Learn more about network interface show in the ONTAP command reference.

The following example shows that the intercluster LIFs <code>cluster01_icl01</code> and <code>cluster01_icl02</code> on the SVM <code>e0e</code> port fail over to the `eOf` port:

```
cluster01::> network interface show -service-policy default-intercluster
-failover
                    Home
Node:Port
                                        Failover
      Logical
                                                      Failover
                                        Policy
Vserver Interface
                                                       Group
  _____
        ------
ipspace-IC1
        cluster01 icl01 cluster01-01:e0e local-only
intercluster01
                        Failover Targets: cluster01-01:e0e,
                                        cluster01-01:e0f
        cluster01 icl02 cluster01-02:e0e local-only
intercluster01
                        Failover Targets: cluster01-02:e0e,
                                         cluster01-02:e0f
```

Configure peer relationships

Create ONTAP cluster peer relationships

Before you can protect your data by replicating it to a remote cluster for data backup and disaster recovery purposes, you should create a cluster peer relationship between the local and remote cluster.

About this task

This procedure applies to FAS, AFF, and ASA systems. If you have an ASA r2 system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, ASA A20, or ASA C30), follow these steps to create set up snapshot replication. ASA r2 systems provide a simplified ONTAP experience specific to SAN-only customers.

Several default protection policies are available. You must have created your protection policies if you want to use custom policies.

Before you begin

If you are using the ONTAP CLI, you must have created intercluster LIFs on every node in the clusters being

peered using one of the following methods:

- Configure intercluster LIFs on shared data ports
- Configure intercluster LIFs on dedicated data ports
- Configure intercluster LIFs in custom IPspaces

Steps

Perform this task using ONTAP System Manager or the ONTAP CLI.

System Manager

- 1. In the local cluster, click **Cluster > Settings**.
- 2. In the **Intercluster Settings** section, click **Add Network Interfaces** and enter the IP address and subnet mask to add intercluster network interfaces for the cluster.

Repeat this step on the remote cluster.

- 3. In the remote cluster, click **Cluster > Settings**.
- 4. Click in the **Cluster Peers** section and select **Generate Passphrase**.
- 5. Select the remote ONTAP cluster version.
- 6. Copy the generated passphrase.
- 7. In the local cluster, under **Cluster peers**, click and select **Peer cluster**.
- 8. In the Peer cluster window, paste the passphrase and click Initiate cluster peering.

CLI

1. On the destination cluster, create a peer relationship with the source cluster:

```
cluster peer create -generate-passphrase -offer-expiration
<MM/DD/YYYY HH:MM:SS|1...7days|1...168hours> -peer-addrs
<peer_LIF_IPs> -initial-allowed-vserver-peers <svm_name|*> -ipspace
<ipspace>
```

If you specify both -generate-passphrase and -peer-addrs, only the cluster whose intercluster LIFs are specified in -peer-addrs can use the generated password.

You can ignore the -ipspace option if you are not using a custom IPspace. Learn more about cluster peer create in the ONTAP command reference.

If you are creating the peering relationship in ONTAP 9.6 or later and you do not want cross-cluster peering communications to be encrypted, you must use the -encryption-protocol-proposed none option to disable encryption.

The following example creates a cluster peer relationship with an unspecified remote cluster, and preauthorizes peer relationships with SVMs vs1 and vs2 on the local cluster:

The following example creates a cluster peer relationship with the remote cluster at intercluster LIF IP addresses 192.140.112.103 and 192.140.112.104, and pre-authorizes a peer relationship with any SVM on the local cluster:

The following example creates a cluster peer relationship with an unspecified remote cluster, and preauthorizes peer relationships with SVMsvs1 and vs2 on the local cluster:

again.

2. On source cluster, authenticate the source cluster to the destination cluster:

cluster peer create -peer-addrs <peer LIF IPs> -ipspace <ipspace>

Learn more about cluster peer create in the ONTAP command reference.

The following example authenticates the local cluster to the remote cluster at intercluster LIF IP addresses 192.140.112.101 and 192.140.112.102:

```
cluster01::> cluster peer create -peer-addrs
192.140.112.101,192.140.112.102
Notice: Use a generated passphrase or choose a passphrase of 8 or
more characters.
        To ensure the authenticity of the peering relationship, use
a phrase or sequence of characters that would be hard to guess.
Enter the passphrase:
Confirm the passphrase:
Clusters cluster02 and cluster01 are peered.
```

Enter the passphrase for the peer relationship when prompted.

3. Verify that the cluster peer relationship was created:

cluster peer show -instance

4. Check the connectivity and status of the nodes in the peer relationship:

```
cluster peer health show
```

```
cluster01::> cluster peer health show
Node cluster-Name
                                 Node-Name
          Ping-Status
                                RDB-Health Cluster-Health
Avail...
_____ _ ____
_____
cluster01-01
         cluster02
                                cluster02-01
           Data: interface reachable
           ICMP: interface reachable true
                                          true
true
                                 cluster02-02
           Data: interface reachable
           ICMP: interface reachable true
                                         true
true
cluster01-02
         cluster02
                                 cluster02-01
           Data: interface reachable
           ICMP: interface reachable true
                                         true
true
                                 cluster02-02
           Data: interface reachable
           ICMP: interface reachable true true
true
```

Other ways to do this in ONTAP

To perform these tasks with	See this content
System Manager Classic (available with ONTAP 9.7 and earlier)	Volume disaster recovery preparation overview

Create ONTAP intercluster SVM peer relationships

You can use the vserver peer create command to create a peer relationship between SVMs on local and remote clusters.

Before you begin

- The source and destination clusters must be peered.
- You must have "pre-authorized" peer relationships for the SVMs on the remote cluster.

For more information, see Creating a cluster peer relationship.

About this task

You can "pre-authorize" peer relationships for multiple SVMs by listing the SVMs in the -initial-allowed -vserver option when you create a cluster peer relationship. For more information, see Creating a cluster peer relationship.

Steps

1. On the data protection destination cluster, display the SVMs that are pre-authorized for peering:

```
vserver peer permission show
```

On the data protection source cluster, create a peer relationship to a pre-authorized SVM on the data protection destination cluster:

```
vserver peer create -vserver local SVM -peer-vserver remote SVM
```

Learn more about vserver peer create in the ONTAP command reference.

The following example creates a peer relationship between the local SVM pvs1 and the pre-authorized remote SVM vs1:

cluster01::> vserver peer create -vserver pvs1 -peer-vserver vs1

3. Verify the SVM peer relationship:

```
vserver peer show
```

cluster01::> vserver peer show				
	Peer	Peer		Peering
Remote				
Vserver	Vserver	State	Peer Cluster	Applications
Vserver				
		,		
pvsl	vsl	peered	cluster02	snapmırror
VSL				

Add ONTAP intercluster SVM peer relationships

If you create an SVM after configuring a cluster peer relationship, you will need to add a peer relationship for the SVM manually. You can use the vserver peer create command to create a peer relationship between SVMs. After the peer relationship has been created, you can run vserver peer accept on the remote cluster to authorize the peer relationship.

Before you begin

The source and destination clusters must be peered.

About this task

You can create a peer relationships between SVMs in the same cluster for local data backup. Learn more about vserver peer create in the ONTAP command reference.

Administrators occasionally use the vserver peer reject command to reject a proposed SVM peer relationship. If the relationship between SVMs is in the rejected state, you must delete the relationship before you can create a new one. Learn more about vserver peer reject in the ONTAP command reference.

Steps

1. On the data protection source cluster, create a peer relationship with an SVM on the data protection destination cluster:

```
vserver peer create -vserver local_SVM -peer-vserver remote_SVM -applications
snapmirror|file-copy|lun-copy -peer-cluster remote cluster
```

The following example creates a peer relationship between the local SVMpvs1 and the remote SVMvs1

```
cluster01::> vserver peer create -vserver pvs1 -peer-vserver vs1
-applications snapmirror -peer-cluster cluster02
```

If the local and remote SVMs have the same names, you must use a *local name* to create the SVM peer relationship:

```
cluster01::> vserver peer create -vserver vs1 -peer-vserver
vs1 -applications snapmirror -peer-cluster cluster01
-local-name cluster1vs1LocallyUniqueName
```

2. On the data protection source cluster, verify that the peer relationship has been initiated:

vserver peer show-all

Learn more about vserver peer show-all in the ONTAP command reference.

The following example shows that the peer relationship between SVM $_{\tt pvs1}$ and SVM $_{\tt vs1}$ has been initiated:

```
cluster01::> vserver peer show-all
                                      Peering
         Peer
                 Peer
                                      Applications
Vserver
        Vserver State Peer Cluster
                                      _____
         _____
                  _____
                           _____
_____
pvs1
         vs1
                 initiated Cluster02
                                       snapmirror
```

3. On the data protection destination cluster, display the pending SVM peer relationship:

vserver peer show

Learn more about vserver peer show in the ONTAP command reference.

The following example lists the pending peer relationships for cluster02:

```
 vserver peer showPeerPeerVserverVserverVserverpvs1pvs1
```

4. On the data protection destination cluster, authorize the pending peer relationship:

vserver peer accept -vserver local SVM -peer-vserver remote SVM

Learn more about vserver peer accept in the ONTAP command reference.

The following example authorizes the peer relationship between the local SVM $_{\tt vs1}$ and the remote SVM $_{\tt pvs1:}$

cluster02::> vserver peer accept -vserver vs1 -peer-vserver pvs1

5. Verify the SVM peer relationship:

vserver peer show

Enable ONTAP cluster peering encryption on peer relationships

Beginning with ONTAP 9.6, cluster peering encryption is enabled by default on all newly created cluster peering relationships. Cluster peering encryption uses a pre-shared key (PSK) and the Transport Security Layer (TLS) to secure cross-cluster peering communications. This adds an additional layer of security between the peered clusters.

About this task

If you are upgrading peered clusters to ONTAP 9.6 or later, and the peering relationship was created in ONTAP 9.5 or earlier, cluster peering encryption must be enabled manually after upgrading. Both clusters in the peering relationship must be running ONTAP 9.6 or later in order to enable cluster peering encryption.

Steps

1. On the destination cluster, enable encryption for communications with the source cluster:

```
cluster peer modify source_cluster -auth-status-admin use-authentication
-encryption-protocol-proposed tls-psk
```

- 2. When prompted enter a passphrase.
- 3. On the data protection source cluster, enable encryption for communication with the data protection destination cluster:

```
cluster peer modify data_protection_destination_cluster -auth-status-admin
use-authentication -encryption-protocol-proposed tls-psk
```

4. When prompted, enter the same passphrase entered on the destination cluster.

Learn more about cluster peer modify in the ONTAP command reference.

Remove ONTAP cluster peering encryption from peer relationships

By default, cluster peering encryption is enabled on all peer relationships created in ONTAP 9.6 or later. If you do not want to use encryption for cross-cluster peering communications, you can disable it.

Steps

1. On the destination cluster, modify communications with the source cluster to discontinue use of cluster

peering encryption:

• To remove encryption, but maintain authentication enter:

```
cluster peer modify <source_cluster> -auth-status-admin use-
authentication -encryption-protocol-proposed none
```

- To remove encryption and authentication:
 - a. Modify the cluster peering policy to allow unauthenticated access:

```
cluster peer policy modify -is-unauthenticated-access-permitted true
```

b. Modify encryption and authentication access:

```
cluster peer modify <source_cluster> -auth-status no-
authentication
```

- 2. When prompted enter the passphrase.
- 3. Confirm the passphrase by re-entering it.
- 4. On the source cluster, disable encryption for communication with the destination cluster:
 - To remove encryption, but maintain authentication enter:

```
cluster peer modify <destination_cluster> -auth-status-admin use-
authentication -encryption-protocol-proposed none
```

- To remove encryption and authentication:
 - a. Modify the cluster peering policy to allow unauthenticated access:

```
cluster peer policy modify -is-unauthenticated-access-permitted true
```

b. Modify encryption and authentication access:

```
cluster peer modify <destination_cluster> -auth-status no-
authentication
```

5. When prompted, enter and re-enter the same passphrase you used on the destination cluster.

Manage local snapshots

Learn about managing local ONTAP snapshots

A *snapshot* is a read-only, point-in-time image of a volume. The image consumes minimal storage space and incurs negligible performance overhead because it records only changes to files since the last snapshot.

You can use a snapshot to restore the entire contents of a volume, or to recover individual files or LUNs. snapshots are stored in the directory .snapshot on the volume.

In ONTAP 9.4 and later, a FlexVol volume can contain up to 1023 snapshots. In ONTAP 9.3 and earlier, a volume can contain up to 255 snapshots.



Beginning with ONTAP 9.8, FlexGroup volumes can contain 1023 snapshots. For more information, see Protect FlexGroup volumes using snapshots.

Configure custom snapshot policies

Learn about configuring custom ONTAP snapshot policies

A *snapshot policy* defines how the system creates snapshots. The policy specifies when to create snapshots, how many copies to retain, and how to name them. For example, a system might create one snapshot every day at 12:10 a.m., retain the two most recent copies, and name the copies "daily.*timestamp*."

The default policy for a volume automatically creates snapshots on the following schedule, with the oldest snapshots deleted to make room for newer copies:

- A maximum of six hourly snapshots taken five minutes past the hour.
- A maximum of two daily snapshots taken Monday through Saturday at 10 minutes after midnight.
- A maximum of two weekly snapshots taken every Sunday at 15 minutes after midnight.

Unless you specify a snapshot policy when you create a volume, the volume inherits the snapshot policy associated with its containing storage virtual machine (SVM).

When to configure a custom ONTAP snapshot policy

If the default snapshot policy is not appropriate for a volume, you can configure a custom policy that modifies the frequency, retention, and name of snapshots. The schedule will be dictated mainly by the rate of change of the active file system.

You might back up a heavily used file system like a database every hour, while you back up rarely used files once a day. Even for a database, you will typically run a full backup once or twice a day, while backing up transaction logs every hour.

Other factors are the importance of the files to your organization, your Service Level Agreement (SLA), your Recovery Point Objective (RPO), and your Recovery Time Objective (RTO). Generally speaking, you should retain only as many snapshots as necessary.

Create an ONTAP snapshot job schedule

A snapshot policy requires at least one snapshot job schedule. You can use System Manager or the job schedule cron create command to create a job schedule. Learn more about job schedule cron create in the ONTAP command reference.

About this task

This procedure applies to FAS, AFF, and ASA systems. If you have an ASA r2 system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, ASA A20, or ASA C30), follow these steps to create a snapshot job schedule. ASA r2 systems provide a simplified ONTAP experience specific to SAN-only customers.

By default, ONTAP forms the names of snapshots by appending a timestamp to the job schedule name.

If you specify values for both day of the month and day of the week, the values are considered independently. For example, a cron schedule with the day specification Friday and the day of the month specification 13 runs every Friday and on the 13th day of each month, not just on every Friday the 13th.

System Manager

- 1. Navigate to Protection > Overview and expand Local policy settings.
- 2. In the **Schedules** pane, click \rightarrow .
- 3. In the Schedules window, click + Add.
- 4. In the Add schedule window, enter the schedule name, and choose the context and schedule type.
- 5. Click Save.

CLI

1. Create a job schedule:

```
job schedule cron create -name <job_name> -month <month> -dayofweek
<day of week> -day <day of month> -hour <hour> -minute <minute>
```

For -month, -dayofweek, and -hour, you can specify all to run the job every month, day of the week, and hour, respectively.

Beginning with ONTAP 9.10.1, you can include the Vserver for your job schedule:

```
job schedule cron create -name <job_name> -vserver <Vserver_name>
-month <month> -dayofweek <day_of_week> -day <day_of_month> -hour
<hour> -minute <minute>
```

The following example creates a job schedule named myweekly that runs on Saturdays at 3:00 a.m.:

```
cluster1::> job schedule cron create -name myweekly -dayofweek
"Saturday" -hour 3 -minute 0
```

The following example creates a schedule named myweeklymulti that specifies multiple days, hours and minutes:

```
job schedule cron create -name myweeklymulti -dayofweek
"Monday,Wednesday,Sunday" -hour 3,9,12 -minute 0,20,50
```

Create an ONTAP snapshot policy

A snapshot policy specifies when to create snapshots, how many copies to retain, and how to name them. For example, a system might create one snapshot every day at 12:10 a.m., retain the two most recent copies, and name them "daily.timestamp." A snapshot policy can contain up to five job schedules.

About this task

This procedure applies to FAS, AFF, and ASA systems. If you have an ASA r2 system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, ASA A20, or ASA C30), follow these steps to create a snapshot policy. ASA r2 systems provide a simplified ONTAP experience specific to SAN-only customers.

By default, ONTAP forms the names of snapshots by appending a timestamp to the job schedule name:

daily.2017-05-14_0013/	hourly.2017-05-15_1106/
daily.2017-05-15_0012/	hourly.2017-05-15_1206/
hourly.2017-05-15_1006/	hourly.2017-05-15_1306/

You can substitute a prefix for the job schedule name if you prefer.

The snapmirror-label option is for SnapMirror replication. For more information, see Defining a rule for a policy.

Steps

You can create a snapshot policy using System Manager or the ONTAP CLI. The procedure creates a snapshot policy on the local cluster only.
System Manager

- 1. Navigate to Protection > Overview and expand Local policy settings.
- 2. In the Snapshot policies pane, click ->.
- 3. In the Snapshot policies tab, click + Add.
- 4. In the Add snapshot policy window, enter the policy name, and choose the scope.
- 5. Click 🕂 Add .
- 6. To select a schedule click the currently displayed schedule name, click V, and choose a different schedule.
- 7. Enter the maximum snapshots to retain, and, if needed, enter the SnapMirror label and the SnapLock retention period.
- 8. Click Save.

CLI

1. Create a snapshot policy:

```
volume snapshot policy create -vserver <SVM> -policy <policy_name>
-enabled true|false -schedule1 <schedule1_name> -count1
<copies_to_retain> -prefix1 <snapshot_prefix> -snapmirror-label1
<snapshot_label> ... -schedule5 <schedule5_name> -count5
<copies_to_retain> -prefix5 <snapshot_prefix> -snapmirror-label5
<snapshot_label>
```

The following example creates a snapshot policy named snap_policy_daily that runs on a daily schedule. The policy has a maximum of five snapshots, each with the name daily.timestamp and the SnapMirror label daily:

```
cluster1::> volume snapshot policy create -vserver vs0 -policy
snap_policy_daily -schedule1 daily -count1 5 -snapmirror-label1
daily
```

Manage snapshots manually

Create and delete snapshots manually

You can create snapshots manually when you can't wait for a scheduled snapshot to be created, and you can delete snapshots when they are no longer needed.

About this task

This procedure applies to FAS, AFF, and ASA systems. If you have an ASA r2 system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, ASA A20, or ASA C30), follow these steps to create an on-demand snapshot. ASA r2 systems provide a simplified ONTAP experience specific to SAN-only customers.

Create a snapshot manually

You can manually create a snapshot using System Manager or the ONTAP CLI.

```
System Manager
Steps

1. Navigate to Storage > Volumes and select the Snapshot copies tab.

2. Click + Add.

3. In the Add a snapshot window, accept the default snapshot name or edit it if desired.

4. Optional: Add a SnapMirror label.

5. Click Add.

CLI

1. Create a snapshot:

volume snapshot create -vserver <SVM> -volume <volume> -snapshot
```

Delete snapshots manually

<snapshot name>

You can manually delete a snapshot using System Manager or the ONTAP CLI.

System Manager

Steps

- 1. Navigate to **Storage > Volumes** and select the **Snapshot copies** tab.
- 2. Locate the snapshot you want to delete, click , and select **Delete**.
- 3. In the **Delete snapshot** window, select **Delete snapshot**.
- 4. Click **Delete**.

CLI

1. Use the volume snapshot show command to verify which snapshots you want to delete.

volume snapshot show -vserver <SVM> -volume <volume>

In this example, the command shows the snapshots on the volume vol3 in the SVM vs3.

cluster::> volume snapshot show -vserver vs3 -volume vol3						
				Blc	ocks	
Vserver	Volume	Snapshot	Size	Total &	Used%	
vs3	vol3					
		snap1.2013-05-01_0015	100KB	0%	38%	
		snap1.2013-05-08_0015	76KB	0%	32%	
		snap2.2013-05-09_0010	76KB	0%	32%	
		snap2.2013-05-10_0010	76KB	0%	32%	
		snap3.2013-05-10_1005	72KB	0%	31%	
		snap3.2013-05-10_1105	72KB	0%	31%	
		snap3.2013-05-10_1205	72KB	0%	31%	
		snap3.2013-05-10_1305	72KB	0%	31%	
		snap3.2013-05-10_1405	72KB	0%	31%	
		snap3.2013-05-10_1505	72KB	0%	31%	
10 entrie	es were (displayed.				

2. Delete a snapshot:

If you want to	Enter this command
Delete a single snapshot	<pre>volume snapshot delete -vserver _svm_namevolume _vol_namesnapshot _snapshot_name_</pre>

If you want to	Enter this command
Delete multiple snapshots	<pre>volume snapshot delete -vserver _svm_namevolume _vol_name_ -snapshot _snapshot_name1_[,_snapshot_nam e2_,]</pre>
Delete all snapshots	<pre>volume snapshot delete -vserver _svm_namevolume _vol_namesnapshot *</pre>

Calculate reclaimable space before deleting snapshots

Beginning with ONTAP 9.10.1, you can use System Manager to select snapshots you want to delete and calculate the reclaimable space before you delete them.

Steps

- 1. Click **Storage > Volumes**.
- 2. Select the volume from which you want to delete snapshots.
- 3. Click Snapshot Copies.
- 4. Select one or more snapshots.
- 5. Click Calculate Reclaimable Space.

Manage the snapshot reserve

Learn about managing the ONTAP snapshot reserve

The *snapshot reserve* sets aside a percentage of disk space for snapshots, five percent by default. Because snapshots use space in the active file system when the snapshot reserve is exhausted, you might want to increase the snapshot reserve as needed. Alternatively, you can autodelete snapshots when the reserve is full.

When to increase the snapshot reserve

In deciding whether to increase the snapshot reserve, it's important to remember that a snapshot records only changes to files since the last snapshot was made. It consumes disk space only when blocks in the active file system are modified or deleted.

This means that the rate of change of the file system is the key factor in determining the amount of disk space used by snapshots. No matter how many snapshots you create, they will not consume disk space if the active file system has not changed.

A FlexVol volume containing database transaction logs, for example, might have a snapshot reserve as large as 20% to account for its greater rate of change. Not only will you want to create more snapshots to capture the more frequent updates to the database, you will also want to have a larger snapshot reserve to handle the additional disk space the snapshots consume.



A snapshot consists of pointers to blocks rather than copies of blocks. You can think of a pointer as a "claim" on a block: ONTAP "holds" the block until the snapshot is deleted.



A Snapshot copy consumes disk space only when blocks in the active file system are modified or deleted.

How deleting protected files can lead to less file space than expected

A snapshot points to a block even after you delete the file that used the block. This explains why an exhausted snapshot reserve might lead to the counter-intuitive result in which deleting an entire file system results in less space being available than the file system occupied.

Consider the following example. Before deleting any files, the df command output is as follows:

Filesystem	kbytes	used	avail	capacity
/vol/vol0/	3000000	3000000	0	100%
/vol/vol0/.snapshot	1000000	500000	500000	50%

After deleting the entire file system and making a snapshot of the volume, the df command generates the following output:

 Filesystem
 kbytes
 used
 avail
 capacity

 /vol/vol0/
 300000
 2500000
 500000
 83%

 /vol/vol0/.snapshot
 100000
 3500000
 0
 350%

As the output shows, the entire 3 GB formerly used by the active file system is now being used by snapshots, in addition to the 0.5 GB used before the deletion.

Because the disk space used by the snapshots now exceeds the snapshot reserve, the overflow of 2.5 GB "spills" into the space reserved for active files, leaving you with 0.5 GB free space for files where you might reasonably have expected 3 GB.

Learn more about the commands described in this procedure in the ONTAP command reference.

Monitor ONTAP snapshot disk consumption

You can monitor snapshot disk consumption using the df command. The command displays the amount of free space in the active file system and the snapshot reserve.

Step

1. Display snapshot disk consumption: df

The following example shows snapshot disk consumption:

cluster1::> df Filesystem kbytes used avail capacity /vol/vol0/ 300000 300000 0 100% /vol/vol0/.snapshot 100000 500000 50000 50%

Learn more about the commands described in this procedure in the ONTAP command reference.

Check available ONTAP snapshot reserve on a volume

You might want to check how much snapshot reserve is available on a volume by using the snapshot-reserve-available parameter with the volume show command. Learn more about volume show in the ONTAP command reference.

Step

1. Check the snapshot reserve available on a volume:

```
vol show -vserver SVM -volume volume -fields snapshot-reserve-available
```

The following example displays the available snapshot reserve for vol1:

```
cluster1::> vol show -vserver vs0 -volume vol1 -fields snapshot-reserve-
available
vserver volume snapshot-reserve-available
------
vs0 vol1 4.84GB
```

Modify the ONTAP snapshot reserve

You might want to configure a larger snapshot reserve to prevent snapshots from using space reserved for the active file system. You can decrease the snapshot reserve when you no longer need as much space for snapshots.

Step

1. Modify the snapshot reserve:

```
volume modify -vserver SVM -volume volume -percent-snapshot-space snap reserve
```

Learn more about volume modify in the ONTAP command reference.

The following example sets the snapshot reserve for vol1 to 10 percent:

```
cluster1::> volume modify -vserver vs0 -volume vol1 -percent-snapshot
-space 10
```

Autodelete ONTAP snapshots

You can use the volume snapshot autodelete modify command to trigger automatic deletion of snapshots when the Snapshot reserve is exceeded. By default, the oldest snapshots are deleted first. Learn more about volume snapshot autodelete modify in the ONTAP command reference.

About this task

LUN and file clones are deleted when there are no more snapshots to be deleted.

Step

1. Autodelete snapshots:

```
volume snapshot autodelete modify -vserver SVM -volume volume -enabled
true|false -trigger volume|snap reserve
```

The following example autodeletes snapshots for vol1 when the snapshot reserve is exhausted:

```
cluster1::> volume snapshot autodelete modify -vserver vs0 -volume vol1
-enabled true -trigger snap_reserve
```

Restore files from snapshots

Restore a file from an ONTAP snapshot on an NFS or SMB client

A user on an NFS or SMB client can restore a file directly from a snapshot without the intervention of a storage system administrator.

Every directory in the file system contains a subdirectory named .snapshot accessible to NFS and SMB users. The .snapshot subdirectory contains subdirectories corresponding to the snapshots of the volume:

\$ ls .snapshot daily.2017-05-14_0013/ hourly.2017-05-15_1106/ daily.2017-05-15_0012/ hourly.2017-05-15_1206/ hourly.2017-05-15_1006/ hourly.2017-05-15_1306/

Each subdirectory contains the files referenced by the snapshot. If users accidentally delete or overwrite a file, they can restore the file to the parent read-write directory by copying the file from the snapshot subdirectory to the read-write directory:

```
$ ls my.txt
ls: my.txt: No such file or directory
$ ls .snapshot
daily.2017-05-14_0013/ hourly.2017-05-15_1106/
daily.2017-05-15_0012/ hourly.2017-05-15_1206/
hourly.2017-05-15_1006/ hourly.2017-05-15_1306/my.txt
my.txt
$ cp .snapshot/hourly.2017-05-15_1306/my.txt .
$ ls my.txt
my.txt
```

Enable and disable NFS and SMB client access to ONTAP snapshot directory

You can enable and disable access to the snapshot directory using the ONTAP CLI -snapdir-access option of the volume modify command, and beginning with ONTAP 9.10.1, you can use System Manager to enable or disable client systems to access to a snapshot directory on a volume. Enabling access makes the snapshot directory visible to clients and allows Windows clients to map a drive to the snapshot directory to view and access its contents. NFS and SMB clients can then restore a file or LUN from a snapshot.

You can enable or disable access to a volume's snapshot directory by editing the volume settings or by editing the volume's share settings.

Enable or disable client access to snapshot directory by editing a volume

Steps

You can enable and disable client snapshot directory access by using ONTAP System Manager or the ONTAP CLI. The snapshot directory on a volume is accessible to clients by default.

System Manager

- 1. Click **Storage > Volumes**.
- 2. Select the volume containing the snapshots directory you want to either show or hide.
- 3. Click and select Edit.
- 4. In the Snapshot Copies (Local) Settings section, select or deselect Show the Snapshot copies directory to clients.
- 5. Click Save.

CLI

1. Check the snapshot directory access status:

```
volume show -vserver <SVM_name> -volume <vol_name> -fields snapdir-
access
```

Example:

Learn more about volume show in the ONTAP command reference.

2. Enable or disable the snapshot directory access:

```
volume modify -vserver <SVM_name> -volume <vol_name> -snapdir-access
<true|false>
```

The following example enables snapshot directory access on vol1:

```
clus1::> volume modify -vserver vs0 -volume vol1 -snapdir-access
true
Volume modify successful on volume vol1 of Vserver vs0.
```

Learn more about volume modify in the ONTAP command reference.

Enable or disable client access to snapshot directory by editing a share

The snapshot directory on a volume is accessible to clients by default.

Steps

- 1. Click Storage > Shares.
- 2. Select the volume containing the snapshots directory you want to either show or hide.
- 3. Click i and select Edit.
- 4. In the Share Properties section, select or deselect Allow clients to access snapshots directory.
- 5. Click Save.

Restore a single file from an ONTAP snapshot

You can use the volume snapshot restore-file command to restore a single file or LUN from a snapshot. You can restore the file to a different location in the parent readwrite volume if you do not want to replace an existing file.

About this task

If you are restoring an existing LUN, a LUN clone is created and backed up in the form of a snapshot. During the restore operation, you can read from and write to the LUN.

Files with streams are restored by default.

Steps

1. List the snapshots in a volume:

volume snapshot show -vserver SVM -volume volume

Learn more about volume snapshot show in the ONTAP command reference.

The following example shows the snapshots in vol1:

clus1::> volume snapshot show -vserver vs1 -volume vol1						
Vserver	Volume	Snapshot	State	Size	Total%	Used%
vs1	voll	hourly.2013-01-25_0005	valid	224KB	 0%	0%
		daily.2013-01-25_0010	valid	92KB	0%	0%
		hourly.2013-01-25_0105	valid	228KB	0%	08
		hourly.2013-01-25_0205	valid	236KB	0%	08
		hourly.2013-01-25_0305	valid	244KB	0%	0%
		hourly.2013-01-25_0405	valid	244KB	0%	0%
		hourly.2013-01-25_0505	valid	244KB	0%	0%
7 entrie	es were	displayed.				

2. Restore a file from a snapshot:

```
volume snapshot restore-file -vserver SVM -volume volume -snapshot snapshot
-path file_path -restore-path destination_path
```

Learn more about volume snapshot restore-file in the ONTAP command reference.

The following example restores the file myfile.txt:

```
cluster1::> volume snapshot restore-file -vserver vs0 -volume vol1
-snapshot daily.2013-01-25 0010 -path /myfile.txt
```

Restore part of a file from an ONTAP snapshot

You can use the volume snapshot partial-restore-file command to restore a range of data from a snapshot to a LUN or to an NFS or SMB container file, assuming you know the starting byte offset of the data and the byte count. You might use this command to restore one of the databases on a host that stores multiple databases in the same LUN.

Beginning with ONTAP 9.12.1, partial restore is available for volumes using SnapMirror active sync.

Steps

1. List the snapshots in a volume:

volume snapshot show -vserver SVM -volume volume

Learn more about volume snapshot show in the ONTAP command reference.

The following example shows the snapshots in vol1:

clus1::> volume snapshot show -vserver vs1 -volume vol1						
Vserver	Volume	Snapshot	State	Size	Total%	Used%
vs1	voll	hourly.2013-01-25_0005	valid	224KB	 0%	 0%
		daily.2013-01-25_0010	valid	92KB	0%	0 %
		hourly.2013-01-25_0105	valid	228KB	0%	08
		hourly.2013-01-25_0205	valid	236KB	0%	0%
		hourly.2013-01-25_0305	valid	244KB	0%	08
		hourly.2013-01-25_0405	valid	244KB	0%	08
		hourly.2013-01-25_0505	valid	244KB	0%	0%
7 entrie	es were	displayed.				

2. Restore part of a file from a snapshot:

volume snapshot partial-restore-file -vserver SVM -volume volume -snapshot snapshot -path file path -start-byte starting byte -byte-count byte count

The starting byte offset and byte count must be multiples of 4,096.

The following example restores the first 4,096 bytes of the file myfile.txt:

cluster1::> volume snapshot partial-restore-file -vserver vs0 -volume vol1 -snapshot daily.2013-01-25_0010 -path /myfile.txt -start-byte 0 -byte-count 4096

Restore the contents of a volume from an ONTAP snapshot

You can recover a volume to an earlier point in time by restoring from a snapshot. You can use System Manager or the volume snapshot restore command to restore the contents of a volume from a snapshot. Learn more about volume snapshot restore in the ONTAP command reference.

About this task

If the volume has SnapMirror relationships, manually replicate all mirror copies of the volume immediately after you restore from a snapshot. Not doing so can result in unusable mirror copies that must be deleted and recreated.

Steps

You can use System Manager or the ONTAP CLI to restore from an earlier snapshot.

System Manager

- 1. Click **Storage** and select a volume.
- 2. Under **Snapshot copies**, click inext to the snapshot you want to restore, and select **Restore**.

CLI

1. List the snapshots in a volume:

```
volume snapshot show -vserver <SVM> -volume <volume>
```

The following example shows the snapshot in vol1:

```
clus1::> volume snapshot show -vserver vs1 -volume vol1
Vserver Volume Snapshot
                                    Size Total% Used%
                              State
_____ _____
vs1 vol1 hourly.2013-01-25 0005 valid 224KB
           hourly.2013-01-25_0000 valid 92KB 0%
                                           08
                                                 0%
                                                08
                                                08
           hourly.2013-01-25 0205 valid 236KB
                                                08
                                           0 응
           hourly.2013-01-25_0305 valid 244KB 0%
                                                 0%
           hourly.2013-01-25 0405 valid 244KB
                                           0%
                                                 0%
           hourly.2013-01-25_0505 valid 244KB 0%
                                                 08
7 entries were displayed.
```

2. Restore the contents of a volume from a snapshot:

```
volume snapshot restore -vserver <SVM> -volume <volume> -snapshot
<snapshot>
```

The following example restores the contents of vol1:

```
cluster1::> volume snapshot restore -vserver vs0 -volume vol1
-snapshot daily.2013-01-25 0010
```

SnapMirror volume replication

Learn about SnapMirror volume replication

Learn about ONTAP SnapMirror asynchronous disaster recovery

SnapMirror is disaster recovery technology, designed for failover from primary storage to secondary storage at a geographically remote site. As its name implies, SnapMirror creates a replica, or *mirror,* of your working data in secondary storage from which you can continue to serve data in the event of a catastrophe at the primary site.

If the primary site is still available to serve data, you can simply transfer any needed data back to it, and not serve clients from the mirror at all. As the failover use case implies, the controllers on the secondary system should be equivalent or nearly equivalent to the controllers on the primary system to serve data efficiently from mirrored storage.

Data protection relationships

Data is mirrored at the volume level. The relationship between the source volume in primary storage and the destination volume in secondary storage is called a *data protection relationship*. The clusters in which the volumes reside and the SVMs that serve data from the volumes must be peered. A peer relationship enables clusters and SVMs to exchange data securely.



This figure illustrates SnapMirror data protection relationships:

A SnapMirror data protection relationship typically mirrors the Snapshot copies available on the source volume.

Scope of data protection relationships

You can create a data protection relationship directly between volumes or between the SVMs that own the volumes. In an *SVM data protection relationship,* all or part of the SVM configuration, from NFS exports and SMB shares to RBAC, is replicated, as well as the data in the volumes that the SVM owns.

You can also use SnapMirror for special data protection applications:

- A *load-sharing mirror* copy of the SVM root volume ensures that data remains accessible in the event of a node outage or failover.
- A data protection relationship between SnapLock volumes lets you replicate WORM files to secondary

storage.

Archive and compliance using SnapLock technology

• Beginning with ONTAP 9.13.1, you can use SnapMirror asynchronous to protect consistency groups. Beginning with ONTAP 9.14.1, you can use SnapMirror asynchronous to replicate volume-granular snapshots to the destination cluster using the consistency group relationship. For more information, see Configure SnapMirror asynchronous protection.

How SnapMirror data protection relationships are initialized

The first time you invoke SnapMirror, it performs a *baseline transfer* from the source volume to the destination volume. The *SnapMirror policy* for the relationship defines the contents of the baseline and any updates.

A baseline transfer under the default SnapMirror policy MirrorAllSnapshots involves the following steps:

- Make a snapshot of the source volume.
- Transfer the snapshot and all the data blocks it references to the destination volume.
- Transfer the remaining, less recent snapshots on the source volume to the destination volume for use in case the "active" mirror is corrupted.

How SnapMirror data protection relationships are updated

Updates are asynchronous, following the schedule you configure. Retention mirrors the snapshot policy on the source.

At each update under the MirrorAllSnapshots policy, SnapMirror creates a snapshot of the source volume and transfers that snapshot and any snapshots that have been made since the last update. In the following output from the snapmirror policy show command for the MirrorAllSnapshots policy, note the following:

- Create Snapshot is "true", indicating that MirrorAllSnapshots creates a snapshot when SnapMirror updates the relationship.
- MirrorAllSnapshots has rules "sm_created" and "all_source_snapshots", indicating that both the snapshot created by SnapMirror and any snapshots that have been made since the last update are transferred when SnapMirror updates the relationship.

cluster dst::> snapmirror policy show -policy MirrorAllSnapshots -instance Vserver: vs0 SnapMirror Policy Name: MirrorAllSnapshots SnapMirror Policy Type: async-mirror Policy Owner: cluster-admin Tries Limit: 8 Transfer Priority: normal Ignore accesstime Enabled: false Transfer Restartability: always Network Compression Enabled: false Create Snapshot: true Comment: SnapMirror asynchronous policy for mirroring all snapshots and the latest active file system. Total Number of Rules: 2 Total Keep: 2 Rules: SnapMirror Label Keep Preserve Warn Schedule Prefix ____ _____ _____ ___ _____ ___ sm created 1 false 0 all_source_snapshots 1 false 0 -

MirrorLatest policy

The preconfigured MirrorLatest policy works exactly the same way as MirrorAllSnapshots, except that only the snapshot created by SnapMirror is transferred at initialization and update.

Schedule Prefix	Rules:	SnapMirror Label	Кеер	Preserve Wa:	rn
		sm_created	1	false	0 -

Related information

snapmirror policy show

Learn about ONTAP SnapMirror synchronous disaster recovery

Beginning with ONTAP 9.5, SnapMirror synchronous (SM-S) technology is supported on all FAS and AFF platforms that have at least 16 GB of memory and on all ONTAP Select

platforms. SnapMirror synchronous technology is a per-node, licensed feature that provides synchronous data replication at the volume level.

This functionality addresses the regulatory and national mandates for synchronous replication in financial, healthcare, and other regulated industries where zero data loss is required.

SnapMirror synchronous operations allowed

The limit on the number of SnapMirror synchronous replication operations per HA pair depends on the controller model.

The following table lists the number of SnapMirror synchronous operations that are allowed per HA pair according to platform type and ONTAP release.

Platform	Releases earlier than ONTAP 9.9.1	ONTAP 9.9.1	ONTAP 9.10.1	ONTAP 9.11.1 through ONTAP 9.14.1
AFF	80	160	200	400
ASA	80	160	200	400
FAS	40	80	80	80
ONTAP Select	20	40	40	40

Supported features

The following table indicates the features supported with SnapMirror synchronous and the ONTAP releases in which support is available.

Feature	Release first supported	Additional information
Antivirus on the primary volume of the SnapMirror synchronous relationship	ONTAP 9.6	
Application-created snapshot replication	ONTAP 9.7	If a snapshot is tagged with the appropriate label at the time of the snapshot create operation, using the CLI or the ONTAP API, SnapMirror synchronous replicates the snapshots, both user created or those created with external scripts, after quiescing the applications. Scheduled snapshots created using a snapshot policy are not replicated. For more information about replicating application-created snapshots, see the Knowledge Base article: How to replicate application created snapshots with SnapMirror synchronous.
Clone auto delete	ONTAP 9.6	

FabricPool aggregates with tiering policy of None, Snapshot, or Auto are supported with SnapMirror synchronous source and destination.	ONTAP 9.5	The destination volume in a FabricPool aggregate cannot be set to All tiering policy.
FC	ONTAP 9.5	Over all networks for which latency does not exceed 10ms
FC-NVMe	ONTAP 9.7	
File clones	ONTAP 9.7	
FPolicy on the primary volume of the SnapMirror synchronous relationship	ONTAP 9.6	
Hard and soft quotas on the primary volume of the SnapMirror synchronous relationship	ONTAP 9.6	The quota rules are not replicated to the destination; therefore, the quota database is not replicated to the destination.
Intra-cluster synchronous relationships	ONTAP 9.14.1	High availability is provided when source and destination volumes are placed on different HA pairs. If the entire cluster goes down, access to volumes will not be possible until the cluster is recovered. Intra-cluster SnapMirror synchronous relationships will contribute to the overall limit of simultaneous relationships per HA pair.
iSCSI	ONTAP 9.5	
LUN clones and NVMe namespace clones	ONTAP 9.7	
LUN clones backed by application- created snapshots	ONTAP 9.7	
Mixed protocol access (NFS v3 and SMB)	ONTAP 9.6	
NDMP/NDMP restore	ONTAP 9.13.1	Both the source and destination cluster must be running ONTAP 9.13.1 or later to use NDMP with SnapMirror Synchronous. For more information, see Transfer data using ndmp copy.
Non-disruptive SnapMirror synchronous operations (NDO) on AFF/ASA platforms, only.	ONTAP 9.12.1	Support for non-disruptive operations enables you to perform many common maintenance tasks without scheduling down time. Operations supported include takeover and giveback, and volume move, provided that a single node is surviving among each of the two clusters.
NFS v4.2	ONTAP 9.10.1	
NFS v4.0	ONTAP 9.6	
NFS v4.1	ONTAP 9.6	
NVMe/TCP	9.10.1	
Removal of high metadata operation frequency limitation	ONTAP 9.6	

Security for sensitive data in-transit using TLS 1.2 encryption	ONTAP 9.6	
Single file and partial file restore	ONTAP 9.13.1	
SMB 2.0 or later	ONTAP 9.6	
SnapMirror synchronous mirror-mirror cascade	ONTAP 9.6	The relationship from the destination volume of the SnapMirror synchronous relationship must be an SnapMirror asynchronous relationship.
SVM disaster recovery	ONTAP 9.6	 * A SnapMirror synchronous source can also be a SVM disaster recovery source, for example, a fan-out configuration with SnapMirror synchronous as one leg and SVM disaster recovery as the other. * A SnapMirror synchronous source cannot be an SVM disaster recovery destination because SnapMirror synchronous does not support cascading a data protection source. You must release the synchronous relationship before performing an SVM disaster recovery flip resync in the destination cluster. * A SnapMirror synchronous destination cannot be an SVM disaster recovery source because SVM disaster recovery does not support replication of DP volumes. A flip resync of the synchronous source would result in the SVM disaster recovery excluding the DP volume in the destination cluster.
Tape-based restore to the source volume	ONTAP 9.13.1	
Timestamp parity between source and destination volumes for NAS	ONTAP 9.6	If you have upgraded from ONTAP 9.5 to ONTAP 9.6, the timestamp is replicated only for any new and modified files in the source volume. The timestamp of existing files in the source volume is not synchronized.

Unsupported features

The following features are not supported with SnapMirror synchronous relationships:

- Consistency groups
- DP_Optimized (DPO) systems
- FlexGroup volumes
- FlexCache volumes
- Global throttling
- In a fan-out configuration, only one relationship can be a SnapMirror synchronous relationship; all the other relationships from the source volume must be SnapMirror asynchronous relationships.
- LUN move
- MetroCluster configurations

- Mixed SAN and NVMe access LUNs and NVMe namespaces are not supported on the same volume or SVM.
- SnapCenter
- SnapLock volumes
- Tamperproof snapshots
- Tape backup or restore using dump and SMTape on the destination volume
- Throughput floor (QoS Min) for source volumes
- Volume SnapRestore
- VVol

Modes of operation

SnapMirror synchronous has two modes of operation based on the type of the SnapMirror policy used:

• Sync mode

In Sync mode, application I/O operations are sent in parallel to the primary and secondary storage systems. If the write to the secondary storage is not completed for any reason, the application is allowed to continue writing to the primary storage. When the error condition is corrected, SnapMirror synchronous technology automatically resynchronizes with the secondary storage and resumes replicating from primary storage to secondary storage in synchronous mode.

In Sync mode, RPO=0 and RTO is very low until a secondary replication failure occurs at which time RPO and RTO become indeterminate, but equal the time to repair the issue that caused secondary replication to fail and for the resync to complete.

StrictSync mode

SnapMirror synchronous can optionally operate in StrictSync mode. If the write to the secondary storage is not completed for any reason, the application I/O fails, thereby ensuring that the primary and secondary storage are identical. Application I/O to the primary resumes only after the SnapMirror relationship returns to the InSync status. If the primary storage fails, application I/O can be resumed on the secondary storage, after failover, with no loss of data.

In StrictSync mode RPO is always zero, and RTO is very low.

Relationship status

The status of a SnapMirror synchronous relationship is always in the InSync status during normal operation. If the SnapMirror transfer fails for any reason, the destination is not in sync with the source and can go to the OutofSync status.

For SnapMirror synchronous relationships, the system automatically checks the relationship status (InSync or OutofSync) at a fixed interval. If the relationship status is OutofSync, ONTAP automatically triggers the auto resync process to bring back the relationship to the InSync status. Auto resync is triggered only if the transfer fails due to any operation, such as unplanned storage failover at source or destination or a network outage. User-initiated operations such as snapmirror quiesce and snapmirror break do not trigger auto resync.

If the relationship status becomes OutofSync for a SnapMirror synchronous relationship in the StrictSync mode, all I/O operations to the primary volume are stopped. The OutofSync state for SnapMirror synchronous relationship in the Sync mode is not disruptive to the primary and I/O operations are allowed on the primary volume.

Related information

- NetApp Technical Report 4733: SnapMirror synchronous configuration and best practices
- snapmirror break
- snapmirror quiesce

Default ONTAP data protection policies

ONTAP includes several default protection policies you can use for your data protection relationships. The policy you use depends on the protection relationship type.

If the default policies don't meet your data protection relationships needs, you can create a custom policy.

List of default protection policies and descriptions

Default protection policies and their associated policy types are described below.

Name	Description	Policy type
Asynchronous	A unified SnapMirror asynchronous and vault policy for mirroring the latest active file system and daily and weekly snapshots with an hourly transfer schedule.	Asynchrono us
AutomatedFailOver	Policy for SnapMirror synchronous with zero RTO guarantee where client I/O will not be disrupted on replication failure.	Synchronou s
AutomatedFailOverDuplex	Policy for SnapMirror synchronous with zero RTO guarantee and bi-directional sync replication.	Synchronou s
CloudBackupDefault	Vault policy with daily rule.	Asynchrono us
Continuous	Policy for S3 bucket mirroring.	Continuous
DailyBackup	Vault policy with a daily rule and a daily transfer schedule.	Asynchrono us
DPDefault	SnapMirror asynchronous policy for mirroring all snapshots and the latest active file system.	Asynchrono us
MirrorAllSnapshots	SnapMirror asynchronous policy for mirroring all snapshots and the latest active file system.	Asynchrono us
MirrorAllSnapshotsDiscardNetwork	SnapMirror asynchronous policy for mirroring all snapshots and the latest active file system excluding the network configurations.	Asynchrono us
MirrorAndVault	A unified SnapMirror asynchronous and vault policy for mirroring the latest active file system and daily and weekly snapshots.	Asynchrono us
MirrorAndVaultDiscardNetwork	A unified SnapMirror asynchronous and vault policy for mirroring the latest active file system and daily and weekly snapshots excluding the network configurations.	Asynchrono us
MirrorLatest	SnapMirror asynchronous policy for mirroring the latest active file system.	Asynchrono us

Name	Description	Policy type
SnapCenterSync	Policy for SnapMirror synchronous for SnapCenter with Application Created Snapshot configuration.	Synchronou s
StrictSync	Policy for SnapMirror synchronous where client access will be disrupted on replication failure.	Synchronou s
Synchronous	Policy for SnapMirror synchronous where client access will not be disrupted on replication failure.	Synchronou s
Unified7year	Unified SnapMirror policy with 7-year retention.	Asynchrono us
XDPDefault	Vault policy with daily and weekly rules.	Asynchrono us

Learn about workloads supported by ONTAP StrictSync and Sync policies

StrictSync and Sync policies support all LUN-based applications with FC, iSCSI, and FC-NVMe protocols, as well as NFSv3 and NFSv4 protocols for enterprise applications such as databases, VMWare, quota, SMB, and so on. Beginning with ONTAP 9.6, SnapMirror synchronous can be used for enterprise file services such as electronic design automation (EDA), home directories, and software build workloads.

In ONTAP 9.5, for a Sync policy, you need to consider a few important aspects while selecting the NFSv3 or NFSv4 workloads. The amount of data read or write operations by workloads is not a consideration, as Sync policy can handle high read or write IO workloads. In ONTAP 9.5, workloads that have excessive file creation, directory creation, file permission changes, or directory permission changes may not be suitable (these are referred to as high-metadata workloads). A typical example of a high-metadata workload is a DevOps workload in which you create multiple test files, run automation, and delete the files. Another example is parallel build workload that generate multiple temporary files during compilation. The impact of a high rate of write metadata activity is that it can cause synchronization between mirrors to temporarily break which stalls the read and write IOs from the client.

Beginning with ONTAP 9.6, these limitations are removed and SnapMirror synchronous can be used for enterprise file services workloads that include multiuser environments, such as home directories and software build workloads.

Related information

SnapMirror synchronous Configuration and Best Practices

Learn about vault archiving using ONTAP SnapMirror technology

SnapMirror vault policies replace SnapVault technology in ONTAP 9.3 and later. You use a SnapMirror vault policy for disk-to-disk snapshot replication for standards compliance and other governance-related purposes. In contrast to a SnapMirror relationship, in which the destination usually contains only the snapshots currently in the source volume, a vault destination typically retains point-in-time snapshots created over a much longer period.

You might want to keep monthly snapshots of your data over a 20-year span, for example, to comply with government accounting regulations for your business. Since there is no requirement to serve data from vault storage, you can use slower, less expensive disks on the destination system.

The figure below illustrates SnapMirror vault data protection relationships.



A SnapVault data protection relationship typically retains point-in-time Snapshot copies created over a longer period than the Snapshot copies on the source volume.

How vault data protection relationships are initialized

The SnapMirror policy for the relationship defines the contents of the baseline and any updates.

A baseline transfer under the default vault policy XDPDefault makes a snapshot of the source volume, then transfers that copy and the data blocks it references to the destination volume. Unlike SnapMirror relationships, a vault backup does not include older snapshots in the baseline.

How vault data protection relationships are updated

Updates are asynchronous, following the schedule you configure. The rules you define in the policy for the relationship identify which new snapshots to include in updates and how many copies to retain. The labels defined in the policy ("monthly," for example) must match one or more labels defined in the snapshot policy on the source. Otherwise, replication fails.

At each update under the XDPDefault policy, SnapMirror transfers snapshots that have been made since the last update, provided they have labels matching the labels defined in the policy rules. In the following output from the snapmirror policy show command for the XDPDefault policy, note the following:

- Create Snapshot is "false", indicating that XDPDefault does not create a snapshot when SnapMirror updates the relationship.
- XDPDefault has rules "daily" and "weekly", indicating that all snapshots with matching labels on the source are transferred when SnapMirror updates the relationship.

cluster dst::> snapmirror policy show -policy XDPDefault -instance Vserver: vs0 SnapMirror Policy Name: XDPDefault SnapMirror Policy Type: vault Policy Owner: cluster-admin Tries Limit: 8 Transfer Priority: normal Ignore accesstime Enabled: false Transfer Restartability: always Network Compression Enabled: false Create Snapshot: false Comment: Default policy for XDP relationships with daily and weekly rules. Total Number of Rules: 2 Total Keep: 59 Rules: SnapMirror Label Keep Preserve Warn Schedule Prefix _____ _____ ___ 7 false 0 daily weekly 52 false 0 -

Related information

• snapmirror policy show

Learn about ONTAP SnapMirror unified replication

SnapMirror *unified replication* allows you to configure disaster recovery and archiving on the same destination volume. When unified replication is appropriate, it offers benefits in reducing the amount of secondary storage you need, limiting the number of baseline transfers, and decreasing network traffic.

How unified data protection relationships are initialized

As with SnapMirror, unified data protection performs a baseline transfer the first time you invoke it. The SnapMirror policy for the relationship defines the contents of the baseline and any updates.

A baseline transfer under the default unified data protection policy MirrorAndVault makes a snapshot of the source volume, then transfers that copy and the data blocks it references to the destination volume. Like vault archiving, unified data protection does not include older snapshots in the baseline.

How unified data protection relationships are updated

At each update under the MirrorAndVault policy, SnapMirror creates a snapshot of the source volume and transfers that snapshot and any snapshots that have been made since the last update, provided they have labels matching the labels defined in the snapshot policy rules. In the following output from the snapmirror policy show command for the MirrorAndVault policy, note the following:

- Create Snapshot is "true", indicating that MirrorAndVault creates a snapshot when SnapMirror updates the relationship.
- MirrorAndVault has rules "sm_created", "daily", and "weekly", indicating that both the snapshot created by SnapMirror and the snapshots with matching labels on the source are transferred when SnapMirror updates the relationship.

```
cluster dst::> snapmirror policy show -policy MirrorAndVault -instance
                   Vserver: vs0
     SnapMirror Policy Name: MirrorAndVault
     SnapMirror Policy Type: mirror-vault
              Policy Owner: cluster-admin
               Tries Limit: 8
          Transfer Priority: normal
  Ignore accesstime Enabled: false
    Transfer Restartability: always
Network Compression Enabled: false
           Create Snapshot: true
                   Comment: A unified SnapMirror synchronous and
SnapVault policy for
                          mirroring the latest file system and daily
and weekly snapshots.
      Total Number of Rules: 3
                Total Keep: 59
                    Rules: SnapMirror Label Keep Preserve Warn
Schedule Prefix
                            _____
                                               ----- -----
----- -----
                                               1 false 0 -
                           sm created
                                                 7 false 0-
                           daily
                                                 52 false 0 -
                           weekly
```

Unified7year policy

The preconfigured Unified7year policy works exactly the same way as MirrorAndVault, except that a fourth rule transfers monthly snapshots and retains them for seven years.

Schodulo Drofiu	Rules:	SnapMirror Label	Кеер	Preserve Warn		
Schedule Fiellx						
		sm_created	1	false	0 —	
_		daily	7	false	0 –	
-		weekly	52	false	0 -	
-		monthly	84	false	0 —	
-						

Protect against possible data corruption

Unified replication limits the contents of the baseline transfer to the snapshot created by SnapMirror at initialization. At each update, SnapMirror creates another snapshot of the source and transfers that snapshot and any new snapshots that have labels matching the labels defined in the snapshot policy rules.

You can protect against the possibility that an updated snapshot is corrupted by creating a copy of the last transferred snapshot on the destination. This "local copy" is retained regardless of the retention rules on the source, so that even if the snapshot originally transferred by SnapMirror is no longer available on the source, a copy of it will be available on the destination.

When to use unified data replication

You need to weigh the benefit of maintaining a full mirror against the advantages that unified replication offers in reducing the amount of secondary storage, limiting the number of baseline transfers, and decreasing network traffic.

The key factor in determining the appropriateness of unified replication is the rate of change of the active file system. A traditional mirror might be better suited to a volume holding hourly snapshots of database transaction logs, for example.

Related information

• snapmirror policy show

When an ONTAP data protection destination volume grows automatically

During a data protection mirror transfer, the destination volume grows automatically in size if the source volume has grown, provided there is available space in the aggregate that contains the volume.

This behavior occurs irrespective of any automatic growth setting on the destination. You cannot limit the volume's growth or prevent ONTAP from growing it.

By default, data protection volumes are set to the grow_shrink autosize mode, which enables the volume to grow or shrink in response to the amount of used space. The max-autosize for data protection volumes is equal to the maximum FlexVol size and is platform dependent. For example:

• FAS8200, default DP volume max-autosize = 100TB

For more information, see NetApp Hardware Universe.

Learn about ONTAP data protection fan-out and cascade deployments

You can use a *fan-out* deployment to extend data protection to multiple secondary systems. You can use a *cascade* deployment to extend data protection to tertiary systems.

Both fan-out and cascade deployments support any combination of SnapMirror DR, SnapVault, or unified replication. Beginning with ONTAP 9.5, SnapMirror synchronous relationships support fan-out deployments with one or more SnapMirror asynchronous relationships. Only one relationship in the fan-out configuration can be a SnapMirror synchronous relationship, all the other relationships from the source volume must be SnapMirror asynchronous relationships. SnapMirror synchronous relationships also support cascade deployments (beginning with ONTAP 9.6); however, the relationship from the destination volume of the SnapMirror synchronous relationship must be a SnapMirror asynchronous relationship must be a SnapMirror synchronous relationship. SnapMirror asynchronous relationship from the destination volume of the SnapMirror synchronous relationship must be a SnapMirror asynchronous relationship. SnapMirror asynchronous relationship must be a SnapMirror asynchronous relationship.



You can use a *fan-in* deployment to create data protection relationships between multiple primary systems and a single secondary system. Each relationship must use a different volume on the secondary system.

You should be aware that volumes that are part of a fan-out or cascade configuration can take longer to

resynchronize. It is not uncommon to see the SnapMirror relationship reporting the status "preparing" for an extended time period.

How fan-out deployments work

SnapMirror supports *multiple-mirrors* and *mirror-vault* fan-out deployments.

A multiple-mirrors fan-out deployment consists of a source volume that has a mirror relationship to multiple secondary volumes.



A mirror-vault fan-out deployment consists of a source volume that has a mirror relationship to a secondary

volume and a SnapVault relationship to a different secondary volume.



Beginning with ONTAP 9.5, you can have fan-out deployments with SnapMirror synchronous relationships; however, only one relationship in the fan-out configuration can be a SnapMirror synchronous relationship, all the other relationships from the source volume must be SnapMirror asynchronous relationships.



How cascade deployments work

SnapMirror supports mirror-mirror, mirror-vault, vault-mirror, and vault-vault cascade deployments.

A mirror-mirror cascade deployment consists of a chain of relationships in which a source volume is mirrored to a secondary volume, and the secondary volume is mirrored to a tertiary volume. If the secondary volume becomes unavailable, you can synchronize the relationship between the primary and tertiary volumes without

performing a new baseline transfer.

Beginning with ONTAP 9.6, SnapMirror synchronous relationships are supported in a mirror-mirror cascade deployment. Only the primary and secondary volumes can be in a SnapMirror synchronous relationship. The relationship between the secondary volumes and tertiary volumes must be asynchronous.



A mirror-vault cascade deployment consists of a chain of relationships in which a source volume is mirrored to a secondary volume, and the secondary volume is vaulted to a tertiary volume.



Vault-mirror and vault-vault cascade deployments are also supported:

- A vault-mirror cascade deployment consists of a chain of relationships in which a source volume is vaulted to a secondary volume, and the secondary volume is mirrored to a tertiary volume.
- A vault-vault cascade deployment consists of a chain of relationships in which a source volume is vaulted to a secondary volume, and the secondary volume is vaulted to a tertiary volume.

Related information

• Resume protection in a fan-out configuration with SnapMirror active sync

Learn about ONTAP SnapMirror licensing

Beginning with ONTAP 9.3, licensing has been simplified for replicating between ONTAP instances. In ONTAP 9 releases, the SnapMirror license supports both vault and mirror relationships. You can use a SnapMirror license to support ONTAP replication for both backup and disaster recovery use cases.

Prior to the ONTAP 9.3 release, a separate SnapVault license was needed to configure *vault* relationships between ONTAP instances, where the DP instance could retain a higher number of snapshots to support backup use cases with longer retention times, and a SnapMirror license was needed to configure *mirror* relationships between ONTAP instances, where each ONTAP instance would maintain the same number of snapshots (that is, a *mirror* image) to support disaster recovery use cases to make cluster failovers possible. Both SnapMirror and SnapVault licenses continue to be used and supported for ONTAP 8.x and 9.x releases.

While SnapVault licenses continue to function and are supported for both ONTAP 8.x and 9.x releases, the SnapMirror license can be used in place of a SnapVault license and can be used for both mirror and vault configurations.

For ONTAP asynchronous replication, beginning with ONTAP 9.3 a single unified replication engine is used to configure extended data protection mode (XDP) policies, where the SnapMirror license can be configured for a mirror policy, a vault policy, or a mirror-vault policy. A SnapMirror license is required on both the source and destination clusters. A SnapVault license is not required if a SnapMirror license is already installed. The SnapMirror asynchronous perpetual license is included in the ONTAP One software suite that's installed on new AFF and FAS systems.

Data protection configuration limits are determined using several factors, including your ONTAP version, hardware platform, and the licenses installed. For more information, see Hardware Universe.

SnapMirror synchronous license

Beginning with ONTAP 9.5, SnapMirror synchronous relationships are supported. You require the following licenses for creating a SnapMirror synchronous relationship:

• The SnapMirror synchronous license is required on both the source cluster and the destination cluster.

The SnapMirror synchronous license is part of the ONTAP One license suite.

If your system was purchased before June 2019 with a Premium or Flash Bundle, you can download a NetApp master key to get the required SnapMirror synchronous license from the NetApp Support Site: Master License Keys.

• The SnapMirror license is required on both the source cluster and the destination cluster.

SnapMirror cloud license

Beginning with ONTAP 9.8, the SnapMirror cloud license provides asynchronous replication of snapshots from ONTAP instances to object storage endpoints. Replication targets can be configured using both on-premises object stores as well as S3 and S3-compatible public cloud object storage services. SnapMirror cloud relationships are supported from ONTAP systems to pre-qualified object storage targets.

SnapMirror cloud is not available as a standalone license. Only one license is needed per ONTAP cluster. In addition to a SnapMirror cloud license, the SnapMirror asynchronous license is also required.

You require the following licenses for creating a SnapMirror cloud relationship:

- Both a SnapMirror license and a SnapMirror cloud license for replicating directly to the object store endpoint.
- When configuring a multi-policy replication workflow (for example, Disk-to-Disk-to-Cloud), a SnapMirror license is required on all ONTAP instances, while the SnapMirror cloud license is only required for the source cluster which is replicating directly to the object storage endpoint.

Beginning with ONTAP 9.9.1, you can use System Manager for SnapMirror cloud replication.

A list of authorized SnapMirror cloud third-party applications is published on the NetApp web site.

Data Protection Optimized license

Data Protection Optimized (DPO) licenses are no longer being sold, and DPO is not supported on current platforms; however, if you have a DPO license installed on a supported platform, NetApp continues to provide

support until the end of availability of that platform.

DPO is not included with the ONTAP One license bundle, and you cannot upgrade to the ONTAP One license bundle if the DPO license is installed on a system.

For information about supported platforms, see Hardware Universe.

ONTAP DPO systems feature enhancements

Beginning with ONTAP 9.6, the maximum number of FlexVol volumes supported increases when the DP_Optimized (DPO) license is installed. Beginning with ONTAP 9.4, systems with the DPO license support SnapMirror backoff, cross-volume background deduplication, use of snapshot blocks as donors, and compaction.

Beginning with ONTAP 9.6, the maximum supported number of FlexVol volumes on secondary or data protection systems has increased, enabling you to scale up to 2,500 FlexVol volumes per node, or up to 5,000 in failover mode. The increase in FlexVol volumes is enabled with the DP_Optimized (DPO) license. A SnapMirror license is still required on both the source and destination nodes.

Beginning with ONTAP 9.4, the following feature enhancements are made to DPO systems:

• SnapMirror backoff: In DPO systems, replication traffic is given the same priority that client workloads are given.

SnapMirror backoff is disabled by default on DPO systems.

• Volume background deduplication and cross-volume background deduplication: Volume background deduplication and cross-volume background deduplication are enabled in DPO systems.

You can run the storage aggregate efficiency cross-volume-dedupe start -aggregate aggregate_name -scan-old-data true command to deduplicate the existing data. The best practice is to run the command during off-peak hours to reduce the impact on performance.

Learn more about storage aggregate efficiency cross-volume-dedupe start in the ONTAP command reference.

• Increased savings by using snapshot blocks as donors: The data blocks that are not available in the active file system but are trapped in snapshots are used as donors for volume deduplication.

The new data can be deduplicated with the data that was trapped in snapshots, effectively sharing the snapshot blocks as well. The increased donor space provides more savings, especially when the volume has a large number of snapshots.

• Compaction: Data compaction is enabled by default on DPO volumes.

Learn about path name pattern matching in ONTAP SnapMirror commands

You can use pattern matching to specify the source and destination paths in snapmirror commands.

snapmirror commands use fully qualified path names in the following format: vserver:volume. You can
abbreviate the path name by not entering the SVM name. If you do this, the snapmirror command assumes
the local SVM context of the user.

Assuming that the SVM is called "vserver1" and the volume is called "vol1", the fully qualified path name is vserver1:vol1.

You can use the asterisk (*) in paths as a wildcard to select matching, fully qualified path names. The following table provides examples of using the wildcard to select a range of volumes.

*	Matches all paths.
vs*	Matches all SVMs and volumes with SVM names beginning with $\ensuremath{\mathtt{vs}}$.
:*src	Matches all SVMs with volume names containing the src text.
:vol	Matches all SVMs with volume names beginning with vol.

<pre>vs1::> snapmirror show -destination-path *:*dest*</pre>											
Progress											
Source		Destination	Mirror	Relationship	Total						
Last											
Path	Туре	Path	State	Status	Progress						
Healthy Updat	ed										
	·										
vsl:sm_src2											
	DP	vs2:sm_dest1									
			Snapmirrored	Idle	-						
true -											

Learn more about snapmirror show in the ONTAP command reference.

Learn about extended queries for ONTAP SnapMirror relationship operations

You can use *extended queries* to perform SnapMirror operations on many SnapMirror relationships at one time. For example, you might have multiple uninitialized SnapMirror relationships that you want to initialize using one command.

About this task

You can apply extended queries to the following SnapMirror operations:

- Initializing uninitialized relationships
- Resuming quiesced relationships
- Resynchronizing broken relationships

- · Updating idle relationships
- Aborting relationship data transfers

Step

1. Perform a SnapMirror operation on many relationships:

```
snapmirror command {-state state } *
```

The following command initializes SnapMirror relationships that are in an Uninitialized state:

vs1::> snapmirror initialize {-state Uninitialized} *

Learn more about snapmirror initialize in the ONTAP command reference.

Compatible ONTAP versions for SnapMirror relationships

The source and destination volumes must be running compatible ONTAP versions before creating a SnapMirror data protection relationship. Before you upgrade ONTAP, you should verify that your current ONTAP version is compatible with your target ONTAP version for SnapMirror relationships.

Unified replication relationships

For SnapMirror relationships of type "XDP", using on premises or Cloud Volumes ONTAP releases:

Beginning with ONTAP 9.9.0:

• ONTAP 9.x.0 releases are cloud-only releases and support Cloud Volumes ONTAP systems. The asterisk (*) after the release version indicates a cloud-only release.



ONTAP 9.16.0 is an exception to the cloud-only rule because it provides support for ASA r2 systems. The plus sign (+) after the release version indicates an ASA r2 supported release. ASA r2 systems support SnapMirror relationships only to other ASA r2 systems.

• ONTAP 9.x.1 releases are general releases and support both on-premises and Cloud Volumes ONTAP systems.



When advanced capacity balancing is enabled on volumes in clusters running ONTAP 9.16.1 or later, SnapMirror transfers are not supported to clusters running ONTAP versions earlier than ONTAP 9.16.1.



Interoperability is bidirectional.

Interoperability for ONTAP version 9.4 and later

TA P

ver

S	io	

n																						
	9.1 7.1	9.1 6.1	9.1 6.0 +	9.1 5.1	9.1 5.0 *	9.1 4.1	9.1 4.0 *	9.1 3.1	9.1 3.0 *	9.1 2.1	9.1 2.0 *	9.1 1.1	9.1 1.0 *	9.1 0.1	9.1 0.0 *	9.9 .1	9.9 .0*	9.8	9.7	9.6	9.5	9.4
9.1 7.1	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	No	No	No	No	No	No	No	No	No	No	No	No
9.1 6.1	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	No	No	No	No	No	No	No	No	No	No
9.1 6.0 +	Ye s	Ye s	Ye s	Ye s	No	No	No	No	No	No	No	No										
9.1 5.1	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	No	No	No	No	No	No
9.1 5.0 *	Ye s	Ye s	No	Ye s	Ye s	Ye s	No	Ye s	No	No	No	No	No	No								
9.1 4.1	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	No	No	No	No	No
9.1 4.0 *	Ye s	Ye s	No	Ye s	No	Ye s	Ye s	Ye s	No	Ye s	No	Ye s	No	Ye s	No	Ye s	No	No	No	No	No	No
9.1 3.1	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	No	No	No	No
9.1 3.0 *	Ye s	Ye s	No	Ye s	No	Ye s	No	Ye s	Ye s	Ye s	No	Ye s	No	Ye s	No	Ye s	No	Ye s	No	No	No	No
9.1 2.1	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	No	No	No
9.1 2.0 *	No	Ye s	No	Ye s	No	Ye s	No	Ye s	No	Ye s	Ye s	Ye s	No	Ye s	No	Ye s	No	Ye s	Ye s	No	No	No
9.1 1.1	No	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	No	No												
9.1 1.0 *	No	No	No	Ye s	Ye s	Ye s	No	Ye s	No	Ye s	Ye s	Ye s	No	No								
9.1 0.1	No	No	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	No										

9.1 0.0 *	No	No	No	Ye s	No	Ye s	No	Ye s	No	Ye s	No	Ye s	No	Ye s	Ye s	Ye s	No	Ye s	Ye s	Ye s	Ye s	No
9.9 .1	No	No	No	Ye s	No																	
9.9 .0*	No	No	No	No	No	Ye s	Ye s	Ye s	Ye s	Ye s	Ye s	No										
9.8	No	No	No	No	No	No	No	Ye s	No													
9.7	No	No	No	No	No	No	No	No	No	Ye s	No											
9.6	No	No	No	No	No	No	No	No	No	No	No	Ye s	No									
9.5	No	No	No	No	No	No	No	No	No	No	No	No	No	Ye s								
9.4	No	No	No	No	No	No	No	No	No	No	No	No	No	No	No	No	No	No	No	No	Ye s	Ye s

SnapMirror synchronous relationships

 (\mathbf{i})

SnapMirror synchronous is not supported for ONTAP cloud instances.

ONTA P versio n	Interoperates with these previous ONTAP versions														
	9.17.1	9.16.1	9.15.1	9.14.1	9.13.1	9.12.1	9.11.1	9.10.1	9.9.1	9.8	9.7	9.6	9.5		
9.17.1	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No		
9.16.1	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No		
9.15.1	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No		
9.14.1	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No		
9.13.1	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No		
9.12.1	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No		
9.11.1	No	Yes	Yes	No	No	No	No								
9.10.1	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No		
9.9.1	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No		
9.8	No	No	No	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	No		
9.7	No	No	No	No	Yes	Yes	No	No	Yes	Yes	Yes	Yes	Yes		
9.6	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes	Yes		
9.5	No	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes		
SnapMirror SVM disaster recovery relationships

For SVM disaster recovery data and SVM protection:

SVM disaster recovery is supported only between clusters running the same version of ONTAP. **Version-independence is not supported for SVM replication**.

For SVM disaster recovery for SVM migration:

- Replication is supported in a single direction from an earlier version of ONTAP on the source to the same or later version of ONTAP on the destination.
- The ONTAP version on the target cluster must be no more than two major on-premises versions newer or two major cloud versions newer (beginning with ONTAP 9.9.0), as shown in the table below.
 - $\,\circ\,$ Replication is not supported for long-term data protection use cases.

The asterisk (*) after the release version indicates a cloud-only release.

To determine support, locate the source version in the left table column, and then locate the destination version on the top row (DR/Migration for like versions and Migration only for newer versions).

So urc e	Des	tinat	ion																			
	9.4	9.5	9.6	9.7	9.8	9.9 .0*	9.9 .1	9.1 0.0 *	9.1 0.1	9.1 1.0 *	9.1 1.1	9.1 2.0 *	9.1 2.1	9.1 3.0 *	9.1 3.1	9.1 4.0 *	9.1 4.1	9.1 5.0 *	9.1 5.1	9.1 6.0	9.1 6.1	9.1 7.1
9.4	DR /Mi gra tion	Mig rati on	Mig rati on																			
9.5		DR /Mi gra tion	Mig rati on	Mig rati on																		
9.6			DR /Mi gra tion	Mig rati on	Mig rati on																	
9.7				DR /Mi gra tion	Mig rati on	Mig rati on																
9.8					DR /Mi gra tion	Mig rati on	Mig rati on															
9.9 .0*						DR /Mi gra tion	Mig rati on	Mig rati on	Mig rati on	Mig rati on	Mig rati on											

9.9 .1			DR /Mi gra tion	Mig rati on	Mig rati on	Mig rati on	Mig rati on											
9.1 0.0 *				DR /Mi gra tion	Mig rati on	Mig rati on	Mig rati on	Mig rati on	Mig rati on									
9.1 0.1					DR /Mi gra tion	Mig rati on	Mig rati on	Mig rati on	Mig rati on									
9.1 1.0 *						DR /Mi gra tion	Mig rati on	Mig rati on	Mig rati on	Mig rati on	Mig rati on							
9.1 1.1							DR /Mi gra tion	Mig rati on	Mig rati on	Mig rati on	Mig rati on							
9.1 2.0 *								DR /Mi gra tion	Mig rati on	Mig rati on	Mig rati on	Mig rati on	Mig rati on					
9.1 2.1									DR /Mi gra tion	Mig rati on	Mig rati on	Mig rati on	Mig rati on					
9.1 3.0 *										DR /Mi gra tion	Mig rati on	Mig rati on	Mig rati on	Mig rati on	Mig rati on			
9.1 3.1											DR /Mi gra tion	Mig rati on	Mig rati on	Mig rati on	Mig rati on			
9.1 4.0 *												DR /Mi gra tion	Mig rati on	Mig rati on	Mig rati on	Mig rati on	Mig rati on	
9.1 4.1													DR /Mi gra tion	Mig rati on	Mig rati on	Mig rati on	Mig rati on	

9.1 5.0 *									DR /Mi gra tion	Mig rati on	Mig rati on	Mig rati on	Mig rati on
9.1 5.1										DR /Mi gra tion	Mig rati on	Mig rati on	Mig rati on
9.1 6.0											DR /Mi gra tion	Mig rati on	Mig rati on
9.1 6.1												DR /Mi gra tion	Mig rati on
9.1 7.1													DR /Mi gra tion

SnapMirror disaster recovery relationships

For SnapMirror relationships of type "DP" and policy type "async-mirror":

()

DP-type mirrors cannot be initialized beginning with ONTAP 9.11.1 and are completely deprecated in ONTAP 9.12.1. For more information, see Deprecation of data protection SnapMirror relationships.

 (\mathbf{i})

In the following table, the column on the left indicates the ONTAP version on the source volume, and the top row indicates the ONTAP versions you can have on your destination volume.

Source	Destinatio	on							
	9.11.1	9.10.1	9.9.1	9.8	9.7	9.6	9.5	9.4	9.3
9.11.1	Yes	No	No	No	No	No	No	No	No
9.10.1	Yes	Yes	No	No	No	No	No	No	No
9.9.1	Yes	Yes	Yes	No	No	No	No	No	No
9.8	No	Yes	Yes	Yes	No	No	No	No	No
9.7	No	No	Yes	Yes	Yes	No	No	No	No
9.6	No	No	No	Yes	Yes	Yes	No	No	No
9.5	No	No	No	No	Yes	Yes	Yes	No	No
9.4	No	No	No	No	No	Yes	Yes	Yes	No
9.3	No	No	No	No	No	No	Yes	Yes	Yes



Interoperability is not bidirectional.

Learn about ONTAP SnapMirror limitations

You should be aware of basic SnapMirror limitations before creating a data protection relationship.

• A destination volume can have only one source volume.



A source volume can have multiple destination volumes. The destination volume can be the source volume for any type of SnapMirror replication relationship.

- Depending on the array model, you can fan out a maximum of eight or sixteen destination volumes from a single source volume. See the Hardware Universe to learn details for your specific configuration.
- You cannot restore files to the destination of a SnapMirror DR relationship.
- · Source or destination SnapVault volumes cannot be 32-bit.
- The source volume for a SnapVault relationship should not be a FlexClone volume.



The relationship will work, but the efficiency offered by FlexClone volumes will not be preserved.

Configure SnapMirror volume replication

ONTAP SnapMirror replication workflow

SnapMirror offers three types of data protection relationship: SnapMirror DR, archive (previously known as SnapVault), and unified replication. You can follow the same basic workflow to configure each type of relationship.

Beginning with general availability in ONTAP 9.9.1, SnapMirror active sync provides Zero Recovery Time Objective (Zero RTO) or Transparent Application Failover (TAF) to enable automatic failover of business-critical applications in SAN environments.

For each type of SnapMirror data protection relationship, the workflow is the same: create a destination volume, create a job schedule, specify a policy, create and initialize the relationship.

Beginning with ONTAP 9.3, you can use the snapmirror protect command to configure a data protection relationship in a single step. Even if you use snapmirror protect, you need to understand each step in the workflow.



Related information

• snapmirror protect

Configure an ONTAP SnapMirror replication relationship in one step

Beginning with ONTAP 9.3, you can use the snapmirror protect command to configure a data protection relationship in a single step. You specify a list of volumes to be replicated, an SVM on the destination cluster, a job schedule, and a SnapMirror policy. snapmirror protect does the rest.

Before you begin

• The source and destination clusters and SVMs must be peered.

Cluster and SVM peering

• The language on the destination volume must be the same as the language on the source volume.

About this task

The snapmirror protect command chooses an aggregate associated with the specified SVM. If no aggregate is associated with the SVM, it chooses from all the aggregates in the cluster. The choice of aggregate is based on the amount of free space and the number of volumes on the aggregate.

The snapmirror protect command then performs the following steps:

- Creates a destination volume with an appropriate type and amount of reserved space for each volume in the list of volumes to be replicated.
- Configures a replication relationship appropriate for the policy you specify.
- Initializes the relationship.

The name of the destination volume is of the form *source_volume_name_dst*. In case of a conflict with an existing name, the command appends a number to the volume name. You can specify a prefix and/or suffix in the command options. The suffix replaces the system-supplied dst suffix.

In ONTAP 9.4 and later, a destination volume can contain up to 1019 snapshots. In ONTAP 9.3 and earlier, a destination volume can contain up to 251 snapshots.



Initialization can be time-consuming. snapmirror protect does not wait for initialization to complete before the job finishes. For this reason, you should use the snapmirror show command rather than the job show command to determine when initialization is complete.

Beginning with ONTAP 9.5, SnapMirror synchronous relationships can be created by using the snapmirror protect command.

Learn more about snapmirror protect in the ONTAP command reference.

Step

1. Create and initialize a replication relationship in one step:

You must replace the variables in angle brackets with the required values before running this command.

```
snapmirror protect -path-list <SVM:volume> -destination-vserver
<destination_SVM> -policy <policy> -schedule <schedule> -auto-initialize
<true|false> -destination-volume-prefix <prefix> -destination-volume
-suffix <suffix>
```



You must run this command from the destination SVM or the destination cluster. The -auto -initialize option defaults to "true".

The following example creates and initializes a SnapMirror DR relationship using the default MirrorAllSnapshots policy:

```
cluster_dst::> snapmirror protect -path-list svm1:volA, svm1:volB
-destination-vserver svm_backup -policy MirrorAllSnapshots -schedule
replication_daily
```



You can use a custom policy if you prefer. For more information, see Creating a custom replication policy.

The following example creates and initializes a SnapVault relationship using the default XDPDefault policy:

```
cluster_dst::> snapmirror protect -path-list svm1:volA, svm1:volB
-destination-vserver svm_backup -policy XDPDefault -schedule
replication_daily
```

The following example creates and initializes a unified replication relationship using the default MirrorAndVault policy:

```
cluster_dst::> snapmirror protect -path-list svm1:volA, svm1:volB
-destination-vserver svm backup -policy MirrorAndVault
```

The following example creates and initializes a SnapMirror synchronous relationship using the default Sync policy:

```
cluster_dst::> snapmirror protect -path-list svm1:volA, svm1:volB
-destination-vserver svm_sync -policy Sync
```



For SnapVault and unified replication policies, you might find it useful to define a schedule for creating a copy of the last transferred snapshot on the destination. For more information, see Defining a schedule for creating a local copy on the destination.

After you finish

Use the snapmirror show command to verify that the SnapMirror relationship was created.

Learn more about snapmirror show in the ONTAP command reference.

Related information

job show

Configure a replication relationship one step at a time

Create an ONTAP SnapMirror destination volume

You can use the volume create command on the destination to create a destination

volume. The destination volume should be the same or greater in size than the source volume. Learn more about volume create in the ONTAP command reference.

Step

1. Create a destination volume:

```
volume create -vserver SVM -volume volume -aggregate aggregate -type DP -size
size
```

The following example creates a 2-GB destination volume named volA dst:

```
cluster_dst::> volume create -vserver SVM_backup -volume volA_dst
-aggregate node01_aggr -type DP -size 2GB
```

Create an ONTAP SnapMirror replication job schedule

The job schedule determines when SnapMirror automatically updates the data protection relationship to which the schedule is assigned. You can use System Manager or the job schedule cron create command to create a replication job schedule. Learn more about job schedule cron create in the ONTAP command reference.

About this task

You assign a job schedule when you create a data protection relationship. If you do not assign a job schedule, you must update the relationship manually.

Steps

You can create a replication job schedule using System Manager or the ONTAP CLI.

System Manager

- 1. Navigate to **Protection > Overview** and and expand **Local policy settings**.
- 2. In the **Schedules** pane, click \rightarrow .
- 3. In the Schedules window, click + Add.
- 4. In the Add schedule window, enter the schedule name, and choose the context and schedule type.
- 5. Click Save.

CLI

1. Create a job schedule:

job schedule cron create -name <job_name> -month <month> -dayofweek
<day of week> -day <day of month> -hour <hour> -minute <minute>

For -month, -dayofweek, and -hour, you can specify all to run the job every month, day of the week, and hour, respectively.

Beginning with ONTAP 9.10.1, you can include the Vserver for your job schedule:

```
job schedule cron create -name <job_name> -vserver <Vserver_name>
-month <month> -dayofweek <day_of_week> -day <day_of_month> -hour
<hour> -minute <minute>
```



The minimum supported schedule (RPO) for FlexVol volumes in a volume SnapMirror relationship is 5 minutes. The minimum supported schedule (RPO) for FlexGroup volumes in a volume SnapMirror relationship is 30 minutes.

The following example creates a job schedule named my_weekly that runs on Saturdays at 3:00 a.m.:

cluster_dst::> job schedule cron create -name my_weekly -dayofweek
"Saturday" -hour 3 -minute 0

Customize a SnapMirror replication policy

Create a custom ONTAP SnapMirror replication policy

You can create a custom replication policy if the default policy for a relationship is not suitable. You might want to compress data in a network transfer, for example, or modify the number of attempts SnapMirror makes to transfer snapshots.

You can use a default or custom policy when you create a replication relationship. For a custom archive (formerly SnapVault) or unified replication policy, you must define one or more *rules* that determine which

snapshots are transferred during initialization and update. You might also want to define a schedule for creating local snapshots on the destination.

The *policy type* of the replication policy determines the type of relationship it supports. The table below shows the available policy types.

Policy type	Relationship type
async-mirror	SnapMirror DR
vault	SnapVault
mirror-vault	Unified replication
strict-sync-mirror	SnapMirror synchronous in the StrictSync mode (supported beginning with ONTAP 9.5)
sync-mirror	SnapMirror synchronous in the Sync mode (supported beginning with ONTAP 9.5)



When you create a custom replication policy, it is a good idea to model the policy after a default policy.

Steps

You can create custom data protection policies with System Manager or the ONTAP CLI. Beginning with ONTAP 9.11.1, you can use System Manager to create custom mirror and vault policies, and to display and select legacy policies. This capability is also available in ONTAP 9.8P12 and later patches of ONTAP 9.8.

Create custom protection policies on both the source and destination cluster.

System Manager

- 1. Click Protection > Overview > Local Policy Settings.
- 2. Under **Protection Policies**, click \rightarrow .
- 3. In the **Protection Policies** pane, click + Add.
- 4. Enter the new policy name, and select the policy scope.
- 5. Choose a policy type. To add a vault-only or mirror-only policy, choose **Asynchronous**, and click **Use a legacy policy type**.
- 6. Complete the required fields.
- 7. Click Save.
- 8. Repeat these steps on the other cluster.

CLI

1. Create a custom replication policy:

```
snapmirror policy create -vserver <SVM> -policy _policy_ -type
<async-mirror|vault|mirror-vault|strict-sync-mirror|sync-mirror>
-comment <comment> -tries <transfer_tries> -transfer-priority
<low|normal> -is-network-compression-enabled <true|false>
```

Beginning with ONTAP 9.5, you can specify the schedule for creating a common snapshot schedule for SnapMirror synchronous relationships by using the <code>-common-snapshot-schedule</code> parameter. By default, the common snapshot schedule for SnapMirror synchronous relationships is one hour. You can specify a value from 30 minutes to two hours for the snapshot schedule for SnapMirror synchronous relationships.

The following example creates a custom replication policy for SnapMirror DR that enables network compression for data transfers:

cluster_dst::> snapmirror policy create -vserver svm1 -policy DR_compressed -type async-mirror -comment "DR with network compression enabled" -is-network-compression-enabled true

The following example creates a custom replication policy for SnapVault:

cluster_dst::> snapmirror policy create -vserver svm1 -policy
my snapvault -type vault

The following example creates a custom replication policy for unified replication:

```
cluster_dst::> snapmirror policy create -vserver svm1 -policy
my unified -type mirror-vault
```

The following example creates a custom replication policy for SnapMirror synchronous relationship in the StrictSync mode:

```
cluster_dst::> snapmirror policy create -vserver svm1 -policy
my_strictsync -type strict-sync-mirror -common-snapshot-schedule
my_sync_schedule
```

Learn more about snapmirror policy create in the ONTAP command reference.

After you finish

For "vault" and "mirror-vault" policy types, you must define rules that determine which snapshots are transferred during initialization and update.

Use the snapmirror policy show command to verify that the SnapMirror policy was created.

```
Learn more about snapmirror policy show in the ONTAP command reference.
```

Define a rule for an ONTAP SnapMirror policy

For custom policies with the vault or mirror-vault policy type, you must define at least one rule that determines which snapshots are transferred during initialization and update. You can also define rules for default policies with the vault or mirror-vault policy type.

About this task

Every policy with the vault or mirror-vault policy type must have a rule that specifies which snapshots to replicate. The rule bi-monthly, for example, indicates that only snapshots assigned the SnapMirror label bi-monthly should be replicated. You specify the SnapMirror label when you configure the snapshot policy on the source.

Each policy type is associated with one or more system-defined rules. These rules are automatically assigned to a policy when you specify its policy type. The table below shows the system-defined rules.

System-defined rule	Used in policy types	Result
sm_created	async-mirror, mirror-vault, Sync, StrictSync	A snapshot created by SnapMirror is transferred on initialization and update.
all_source_snapshots	async-mirror	New snapshots on the source are transferred on initialization and update.
daily	vault,mirror-vault	New snapshots on the source with the SnapMirror label daily are transferred on initialization and update.

weekly	vault,mirror-vault	New snapshots on the source with the SnapMirror label weekly are transferred on initialization and update.
monthly	mirror-vault	New snapshots on the source with the SnapMirror label monthly are transferred on initialization and update.
app_consistent	Sync, StrictSync	Snapshots with the SnapMirror label app_consistent on source are synchronously replicated to the destination. Supported beginning with ONTAP 9.7.

Except for the "async-mirror" policy type, you can specify additional rules as needed, for default or custom policies. For example:

- For the default MirrorAndVault policy, you might create a rule called bi-monthly to match snapshots on the source with the bi-monthly SnapMirror label.
- For a custom policy with the mirror-vault policy type, you might create a rule called bi-weekly to match snapshots on the source with the bi-weekly SnapMirror label.

Step

1. Define a rule for a policy:

```
snapmirror policy add-rule -vserver SVM -policy policy_for_rule -snapmirror
-label snapmirror-label -keep retention_count
```

The following example adds a rule with the SnapMirror label bi-monthly to the default MirrorAndVault policy:

cluster_dst::> snapmirror policy add-rule -vserver svm1 -policy MirrorAndVault -snapmirror-label bi-monthly -keep 6

The following example adds a rule with the SnapMirror label bi-weekly to the custom my_snapvault policy:

```
cluster_dst::> snapmirror policy add-rule -vserver svm1 -policy
my snapvault -snapmirror-label bi-weekly -keep 26
```

The following example adds a rule with the SnapMirror label app consistent to the custom Sync policy:

```
cluster_dst::> snapmirror policy add-rule -vserver svm1 -policy Sync
-snapmirror-label app_consistent -keep 1
```

Learn more about snapmirror policy add-rule in the ONTAP command reference.

You can then replicate snapshots from the source cluster that match this SnapMirror label:

```
cluster_src::> snapshot create -vserver vs1 -volume vol1 -snapshot
snapshot1 -snapmirror-label app consistent
```

Define an ONTAP SnapMirror schedule to create a local copy on the destination

For SnapVault and unified replication relationships, you can protect against the possibility that an updated snapshot is corrupted by creating a copy of the last transferred snapshot on the destination. This "local copy" is retained regardless of the retention rules on the source, so that even if the snapshot originally transferred by SnapMirror is no longer available on the source, a copy of it will be available on the destination.

About this task

You specify the schedule for creating a local copy in the -schedule option of the snapmirror policy add-rule command.

Step

1. Define a schedule for creating a local copy on the destination:

```
snapmirror policy add-rule -vserver SVM -policy policy_for_rule -snapmirror
-label snapmirror-label -schedule
```

For an example of how to create a job schedule, see Creating a replication job schedule.

The following example adds a schedule for creating a local copy to the default MirrorAndVault policy:

```
cluster_dst::> snapmirror policy add-rule -vserver svm1 -policy
MirrorAndVault -snapmirror-label my_monthly -schedule my_monthly
```

The following example adds a schedule for creating a local copy to the custom my_unified policy:

```
cluster_dst::> snapmirror policy add-rule -vserver svm1 -policy
my_unified -snapmirror-label my_monthly -schedule my_monthly
```

Learn more about snapmirror policy add-rule in the ONTAP command reference.

The relationship between the source volume in primary storage and the destination volume in secondary storage is called a *data protection relationship*. You can use the snapmirror create command to create SnapMirror DR, SnapVault, or unified replication data protection relationships.



This procedure applies to FAS, AFF, and ASA systems. If you have an ASA r2 system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, ASA A20, or ASA C30), follow these steps to create a replication relationship. ASA r2 systems provide a simplified ONTAP experience specific to SAN-only customers.

Beginning with ONTAP 9.11.1, you can use System Manager to select pre-created and custom mirror and vault policies, to display and select legacy policies, and to override the transfer schedules defined in a protection policy when protecting volumes and storage VMs. This capability is also available in ONTAP 9.8P12 and later patches of ONTAP 9.8.



If you are using ONTAP 9.8P12 or later ONTAP 9.8 patch release and you configured SnapMirror using System Manager, you should use ONTAP 9.9.1P13 or later and ONTAP 9.10.1P10 or later patch releases if you plan to upgrade to ONTAP 9.9.1 or ONTAP 9.10.1 releases.

Before you begin

• The source and destination clusters and SVMs must be peered.

Cluster and SVM peering

• The language on the destination volume must be the same as the language on the source volume.

About this task

Until ONTAP 9.3, SnapMirror invoked in DP mode and SnapMirror invoked in XDP mode used different replication engines, with different approaches to version-dependence:

• SnapMirror invoked in DP mode used a *version-dependent* replication engine in which the ONTAP version was required to be the same on primary and secondary storage:

```
cluster_dst::> snapmirror create -type DP -source-path ... -destination
-path ...
```

• SnapMirror invoked in XDP mode used a *version-flexible* replication engine that supported different ONTAP versions on primary and secondary storage:

```
cluster_dst::> snapmirror create -type XDP -source-path ...
-destination-path ...
```

With improvements in performance, the significant benefits of version-flexible SnapMirror outweigh the slight advantage in replication throughput obtained with version-dependent mode. For this reason, beginning with ONTAP 9.3, XDP mode has been made the new default, and any invocations of DP mode on the command

line or in new or existing scripts are automatically converted to XDP mode.

Existing relationships are not affected. If a relationship is already of type DP, it will continue to be of type DP. The table below shows the behavior you can expect.

If you specify	The type is…	The default policy (if you do not specify a policy) is…
DP	XDP	MirrorAllSnapshots (SnapMirror DR)
Nothing	XDP	MirrorAllSnapshots (SnapMirror DR)
XDP	XDP	XDPDefault (SnapVault)

See also the examples in the procedure below.

The only exceptions to conversion are as follows:

• SVM data protection relationships continue to default to DP mode.

Specify XDP explicitly to obtain XDP mode with the default MirrorAllSnapshots policy.

- Load-sharing data protection relationships continue to default to DP mode.
- SnapLock data protection relationships continue to default to DP mode.
- Explicit invocations of DP continue to default to DP mode if you set the following cluster-wide option:

options replication.create_data_protection_rels.enable on

This option is ignored if you do not explicitly invoke DP.

Beginning with ONTAP 9.14.1, the -backoff-level option is added to the snapmirror create, snapmirror modify, and snapmirror restore commands to enable you to specify the backoff level per relationship. The option is supported only with FlexVol SnapMirror relationships. The optional command specifies the SnapMirror backoff level due to client ops. Backoff values can be high, medium or none. The default value is high.

Beginning with ONTAP 9.5, SnapMirror synchronous relationships are supported.

In ONTAP 9.4 and later, a destination volume can contain up to 1019 snapshots. In ONTAP 9.3 and earlier, a destination volume can contain up to 251 snapshots.

Steps

You can use System Manager or the ONTAP CLI to create a replication relationship.

System Manager

- 1. Select the volume or LUN to protect: click **Storage > Volumes** or **Storage > LUNs**, and then click the desired volume or LUN name.
- 2. Click 🔵 Protect
- 3. Select the destination cluster and storage VM.
- 4. The asynchronous policy is selected by default. To select a synchronous policy, click More Options.
- 5. Click Protect.
- 6. Click the **SnapMirror (Local or Remote)** tab for the selected volume or LUN to verify that protection is set up correctly.

CLI

1. From the destination cluster, create a replication relationship:

You must replace the variables in angle brackets with the required values before running this command.

```
snapmirror create -source-path <SVM:volume> -destination-path
<SVM:volume> -type <DP|XDP> -schedule <schedule> -policy <policy>
```



The schedule parameter is not applicable when creating SnapMirror synchronous relationships.

The following example creates a SnapMirror DR relationship using the default MirrorLatest policy:

```
cluster_dst::> snapmirror create -source-path svm1:volA -destination
-path svm_backup:volA_dst -type XDP -schedule my_daily -policy
MirrorLatest
```

The following example creates a SnapVault relationship using the default XDPDefault policy:

cluster_dst::> snapmirror create -source-path svm1:volA -destination
-path svm_backup:volA_dst -type XDP -schedule my_daily -policy
XDPDefault

The following example creates a unified replication relationship using the default MirrorAndVault policy:

cluster_dst:> snapmirror create -source-path svm1:volA -destination
-path svm_backup:volA_dst -type XDP -schedule my_daily -policy
MirrorAndVault

The following example creates a unified replication relationship using the custom my unified policy:

cluster_dst::> snapmirror create -source-path svm1:volA -destination
-path svm_backup:volA_dst -type XDP -schedule my_daily -policy
my unified

The following example creates a SnapMirror synchronous relationship using the default Sync policy:

```
cluster_dst::> snapmirror create -source-path svm1:volA -destination
-path svm backup:volA dst -type XDP -policy Sync
```

The following example creates a SnapMirror synchronous relationship using the default StrictSync policy:

```
cluster_dst::> snapmirror create -source-path svm1:volA -destination
-path svm backup:volA dst -type XDP -policy StrictSync
```

The following example creates a SnapMirror DR relationship. With the DP type automatically converted to XDP and with no policy specified, the policy defaults to the MirrorAllSnapshots policy:

```
cluster_dst::> snapmirror create -source-path svm1:volA -destination
-path svm backup:volA dst -type DP -schedule my daily
```

The following example creates a SnapMirror DR relationship. With no type or policy specified, the policy defaults to the MirrorAllSnapshots policy:

cluster_dst::> snapmirror create -source-path svm1:volA -destination
-path svm backup:volA dst -schedule my_daily

The following example creates a SnapMirror DR relationship. With no policy specified, the policy defaults to the XDPDefault policy:

```
cluster_dst::> snapmirror create -source-path svm1:volA -destination
-path svm backup:volA dst -type XDP -schedule my daily
```

The following example creates a SnapMirror synchronous relationship with the predefined policy SnapCenterSync:

cluster_dst::> snapmirror create -source-path svm1:volA -destination
-path svm backup:volA dst -type XDP -policy SnapCenterSync



The predefined policy SnapCenterSync is of type Sync. This policy replicates any snapshot that is created with the snapmirror-label of "app_consistent".

After you finish

Use the snapmirror show command to verify that the SnapMirror relationship was created.

Learn more about snapmirror show in the ONTAP command reference.

Related information

• Create and delete SnapMirror failover test volumes.

Other ways to do this in ONTAP

To perform these tasks with	See this content
System Manager Classic (available with ONTAP 9.7 and earlier)	Volume backup using SnapVault overview

Related information

• snapmirror create

Initialize an ONTAP SnapMirror replication relationship

For all relationship types, initialization performs a *baseline transfer*: it makes a snapshot of the source volume, then transfers that copy and all the data blocks it references to the destination volume. Otherwise, the contents of the transfer depend on the policy.

Before you begin

The source and destination clusters and SVMs must be peered.

Cluster and SVM peering

About this task

Initialization can be time-consuming. You might want to run the baseline transfer in off-peak hours.

Beginning with ONTAP 9.5, SnapMirror synchronous relationships are supported.

You should be aware that if a filesystem is rebooted for any reason, such as a node reboot, takeover/giveback, or panic, then initialization will not automatically resume and must be restarted manually.

Step

1. Initialize a replication relationship:

```
snapmirror initialize -source-path <SVM:volume>|<cluster://SVM/volume>, ...
-destination-path <SVM:volume>|<cluster://SVM/volume>, ...
```



You must run this command from the destination SVM or the destination cluster.

The following example initializes the relationship between the source volume volA on svm1 and the destination volume volA_dst on svm_backup:

cluster_dst::> snapmirror initialize -source-path svm1:volA -destination
-path svm_backup:volA_dst

Learn more about snapmirror initialize in the ONTAP command reference.

Ensure a common snapshot in an ONTAP mirror-vault deployment

You can use the snapmirror snapshot-owner create command to preserve a labeled snapshot on the secondary in a mirror-vault deployment. Doing so ensures that a common snapshot exists for the update of the vault relationship.

About this task

If you use a combination mirror-vault fan-out or cascade deployment, you should keep in mind that updates will fail if a common snapshot does not exist on the source and destination volumes.

This is never an issue for the mirror relationship in a mirror-vault fan-out or cascade deployment, since SnapMirror always creates a snapshot of the source volume before it performs the update.

It might be an issue for the vault relationship, however, because SnapMirror does not create a snapshot of the source volume when it updates a vault relationship. You need to use the snapmirror snapshot-owner create to ensure that there is at least one common snapshot on both the source and destination of the vault relationship. Learn more about data protection fan-out and cascade deployments.

Steps

1. On the source volume, assign an owner to the labeled snapshot you want to preserve:

```
snapmirror snapshot-owner create -vserver <SVM> -volume <volume> -snapshot
<snapshot> -owner <owner>
```

The following example assigns ApplicationA as the owner of the snap1 snapshot:

```
clust1::> snapmirror snapshot-owner create -vserver vs1 -volume vol1
-snapshot snap1 -owner ApplicationA
```

Learn more about snapmirror snapshot-owner create in the ONTAP command reference.

2. Update the mirror relationship, as described in Updating a replication relationship manually.

Alternatively, you can wait for the scheduled update of the mirror relationship.

3. Transfer the labeled snapshot to the vault destination:

snapmirror update -source-path <SVM:volume>|<cluster://SVM/volume>, ...

-destination-path <SVM:volume>|<cluster://SVM/volume>, ... -source-snapshot snapshot

The following example transfers the snap1 snapshot

```
clust1::> snapmirror update -vserver vs1 -volume vol1
-source-snapshot snap1
```

The labeled snapshot will be preserved when the vault relationship is updated.

Learn more about snapmirror update in the ONTAP command reference.

4. On the source volume, remove the owner from the labeled snapshot:

```
snapmirror snapshot-owner delete -vserver SVM -volume volume -snapshot
snapshot -owner
```

The following examples removes ApplicationA as the owner of the snap1 snapshot:

clust1::> snapmirror snapshot-owner delete -vserver vs1 -volume vol1
-snapshot snap1 -owner ApplicationA

Learn more about snapmirror snapshot-owner delete in the ONTAP command reference.

Example: Configure an ONTAP SnapMirror vault-vault cascade

An example will show in concrete terms how you can configure replication relationships one step at a time. You can use the vault-vault cascade deployment configured in the example to retain more than 251 snapshots labeled my-weekly.

Before you begin

The source and destination clusters and SVMs must be peered.

About this task

The example assumes the following:

- You have configured snapshots on the source cluster with the SnapMirror labels my-daily, my-weekly, and my-monthly.
- You have configured destination volumes named volA on the secondary and tertiary destination clusters.
- You have configured replication job schedules named my_snapvault on the secondary and tertiary destination clusters.

The example shows how to create replication relationships based on two custom policies:

- The snapvault_secondary policy retains 7 daily, 52 weekly, and 180 monthly snapshots on the secondary destination cluster.
- The snapvault tertiary policy retains 250 weekly snapshots on the tertiary destination cluster.

Steps

1. On the secondary destination cluster, create the snapvault secondary policy:

cluster_secondary::> snapmirror policy create -policy snapvault_secondary
-type vault -comment "Policy on secondary for vault to vault cascade" -vserver
svm secondary

2. On the secondary destination cluster, define the my-daily rule for the policy:

```
cluster_secondary::> snapmirror policy add-rule -policy snapvault_secondary
-snapmirror-label my-daily -keep 7 -vserver svm secondary
```

3. On the secondary destination cluster, define the my-weekly rule for the policy:

cluster_secondary::> snapmirror policy add-rule -policy snapvault_secondary -snapmirror-label my-weekly -keep 52 -vserver svm secondary

4. On the secondary destination cluster, define the my-monthly rule for the policy:

cluster_secondary::> snapmirror policy add-rule -policy snapvault_secondary -snapmirror-label my-monthly -keep 180 -vserver svm secondary

5. On the secondary destination cluster, verify the policy:

cluster_secondary::> snapmirror policy show snapvault_secondary -instance

Vserver: svm secondary SnapMirror Policy Name: snapvault secondary SnapMirror Policy Type: vault Policy Owner: cluster-admin Tries Limit: 8 Transfer Priority: normal Ignore accesstime Enabled: false Transfer Restartability: always Network Compression Enabled: false Create Snapshot: false Comment: Policy on secondary for vault to vault cascade Total Number of Rules: 3 Total Keep: 239 Rules: SnapMirror Label Keep Preserve Warn Schedule Prefix _____ ____ ___ _____ ___ 7 false 0 my-daily my-weekly 52 false 0 my-monthly 180 false 0 -

6. On the secondary destination cluster, create the relationship with the source cluster:

cluster_secondary::> snapmirror create -source-path svm_primary:volA
-destination-path svm_secondary:volA -type XDP -schedule my_snapvault -policy
snapvault secondary

7. On the secondary destination cluster, initialize the relationship with the source cluster:

```
cluster_secondary::> snapmirror initialize -source-path svm_primary:volA
-destination-path svm secondary:volA
```

8. On the tertiary destination cluster, create the snapvault tertiary policy:

```
cluster_tertiary::> snapmirror policy create -policy snapvault_tertiary -type
vault -comment "Policy on tertiary for vault to vault cascade" -vserver
svm tertiary
```

9. On the tertiary destination cluster, define the my-weekly rule for the policy:

```
cluster_tertiary::> snapmirror policy add-rule -policy snapvault_tertiary
-snapmirror-label my-weekly -keep 250 -vserver svm tertiary
```

10. On the tertiary destination cluster, verify the policy:

cluster_tertiary::> snapmirror policy show snapvault_tertiary -instance

Vserver: svm tertiary SnapMirror Policy Name: snapvault tertiary SnapMirror Policy Type: vault Policy Owner: cluster-admin Tries Limit: 8 Transfer Priority: normal Ignore accesstime Enabled: false Transfer Restartability: always Network Compression Enabled: false Create Snapshot: false Comment: Policy on tertiary for vault to vault cascade Total Number of Rules: 1 Total Keep: 250 Rules: SnapMirror Label Keep Preserve Warn Schedule Prefix _____ ____ _____ ___ _____ ___ 250 false 0my-weekly

11. On the tertiary destination cluster, create the relationship with the secondary cluster:

```
cluster_tertiary::> snapmirror create -source-path svm_secondary:volA
-destination-path svm_tertiary:volA -type XDP -schedule my_snapvault -policy
snapvault_tertiary
```

12. On the tertiary destination cluster, initialize the relationship with the secondary cluster:

```
cluster_tertiary::> snapmirror initialize -source-path svm_secondary:volA
-destination-path svm_tertiary:volA
```

Related information

- snapmirror create
- snapmirror initialize
- snapmirror policy add-rule
- snapmirror policy create
- snapmirror policy show

Manage SnapMirror volume replication

Convert an existing ONTAP SnapMirror DP-type relationship to XDP

If you are upgrading to ONTAP 9.12.1 or later, you must convert DP-type relationships to

XDP before upgrading. ONTAP 9.12.1 and later does not support DP-type relationships. You can easily convert an existing DP-type relationship to XDP to take advantage of version-flexible SnapMirror.

Before upgrading to ONTAP 9.12.1, you must convert existing DP-type relationships to XDP before you can upgrade to ONTAP 9.12.1 and later releases.

About this task

- SnapMirror does not automatically convert existing DP-type relationships to XDP. To convert the relationship, you need to break and delete the existing relationship, create a new XDP relationship, and resync the relationship.
- When planning your conversion, you should be aware that background preparation and the data warehousing phase of an XDP SnapMirror relationship can take a long time. It is not uncommon to see the SnapMirror relationship reporting the status "preparing" for an extended time period.



After you convert a SnapMirror relationship type from DP to XDP, space-related settings, such as autosize and space guarantee are no longer replicated to the destination.

Steps

1. From the destination cluster, ensure that the SnapMirror relationship is type DP, that the mirror state is SnapMirrored, the relationship status is Idle, and the relationship is healthy:

snapmirror show -destination-path <SVM:volume>

The following example shows the output from the snapmirror show command:

cluster dst::>snapmirror show -destination-path svm backup:volA dst Source Path: svml:volA Destination Path: svm backup:volA dst Relationship Type: DP SnapMirror Schedule: -Tries Limit: -Throttle (KB/sec): unlimited Mirror State: Snapmirrored Relationship Status: Idle Transfer Snapshot: -Snapshot Progress: -Total Progress: -Snapshot Checkpoint: -Newest Snapshot: snapmirror.10af643c-32d1-11e3-954b-123478563412 2147484682.2014-06-27 100026 Newest Snapshot Timestamp: 06/27 10:00:55 Exported Snapshot: snapmirror.10af643c-32d1-11e3-954b-123478563412 2147484682.2014-06-27 100026 Exported Snapshot Timestamp: 06/27 10:00:55 Healthy: true



You might find it helpful to retain a copy of the snapmirror show command output to keep track existing of the relationship settings. Learn more about snapmirror show in the ONTAP command reference.

2. From the source and the destination volumes, ensure that both volumes have a common snapshot:

volume snapshot show -vserver <SVM> -volume <volume>

The following example shows the volume snapshot show output for the source and the destination volumes:

```
cluster src:> volume snapshot show -vserver vsm1 -volume volA
---Blocks---
Vserver Volume Snapshot State Size Total% Used%
_____ _ _____ _____
_____ ____
svml volA
weekly.2014-06-09 0736 valid 76KB 0% 28%
weekly.2014-06-16 1305 valid 80KB 0% 29%
daily.2014-06-26 0842 valid 76KB 0% 28%
hourly.2014-06-26 1205 valid 72KB 0% 27%
hourly.2014-06-26 1305 valid 72KB 0% 27%
hourly.2014-06-26 1405 valid 76KB 0% 28%
hourly.2014-06-26 1505 valid 72KB 0% 27%
hourly.2014-06-26 1605 valid 72KB 0% 27%
daily.2014-06-27 0921 valid 60KB 0% 24%
hourly.2014-06-27 0921 valid 76KB 0% 28%
snapmirror.10af643c-32d1-11e3-954b-123478563412 2147484682.2014-06-
27 100026
valid 44KB 0% 19%
11 entries were displayed.
cluster dest:> volume snapshot show -vserver svm backup -volume volA dst
---Blocks---
Vserver Volume Snapshot State Size Total% Used%
_____ __ ____
_____ ____
svm backup volA dst
weekly.2014-06-09 0736 valid 76KB 0% 30%
weekly.2014-06-16 1305 valid 80KB 0% 31%
daily.2014-06-26 0842 valid 76KB 0% 30%
hourly.2014-06-26 1205 valid 72KB 0% 29%
hourly.2014-06-26 1305 valid 72KB 0% 29%
hourly.2014-06-26 1405 valid 76KB 0% 30%
hourly.2014-06-26 1505 valid 72KB 0% 29%
hourly.2014-06-26 1605 valid 72KB 0% 29%
daily.2014-06-27 0921 valid 60KB 0% 25%
hourly.2014-06-27 0921 valid 76KB 0% 30%
snapmirror.10af643c-32d1-11e3-954b-123478563412 2147484682.2014-06-
27 100026
```

3. To ensure scheduled updates will not run during the conversion, quiesce the existing DP-type relationship:

```
snapmirror quiesce -source-path <SVM:volume> -destination-path
<SVM:volume>
```



You must run this command from the destination SVM or the destination cluster.

The following example quiesces the relationship between the source volume volA on svm1 and the destination volume volA_dst on svm_backup:

cluster dst::> snapmirror quiesce -destination-path svm backup:volA dst

Learn more about snapmirror quiesce in the ONTAP command reference.

4. Break the existing DP-type relationship:

snapmirror break -destination-path <SVM:volume>



You must run this command from the destination SVM or the destination cluster.

The following example breaks the relationship between the source volume volA on svm1 and the destination volume volA dst on svm backup:

cluster_dst::> snapmirror break -destination-path svm_backup:volA_dst

Learn more about snapmirror break in the ONTAP command reference.

5. If automatic deletion of snapshots is enabled on the destination volume, disable it:

```
volume snapshot autodelete modify -vserver _SVM_ -volume _volume_
-enabled false
```

The following example disables snapshot autodelete on the destination volume vola dst:

```
cluster_dst::> volume snapshot autodelete modify -vserver svm_backup
-volume volA dst -enabled false
```

6. Delete the existing DP-type relationship:

```
snapmirror delete -destination-path <SVM:volume>
```

Learn more about snapmirror-delete in the ONTAP command reference.



You must run this command from the destination SVM or the destination cluster.

The following example deletes the relationship between the source volume volA on svm1 and the destination volume volA dst on svm backup:

cluster_dst::> snapmirror delete -destination-path svm_backup:volA_dst

7. Release the origin SVM disaster recovery relationship on the source:

```
snapmirror release -destination-path <SVM:volume> -relationship-info
-only true
```

The following example releases the SVM disaster recovery relationship:

```
cluster_src::> snapmirror release -destination-path svm_backup:volA_dst
-relationship-info-only true
```

Learn more about snapmirror release in the ONTAP command reference.

8. You can use the output you retained from the snapmirror show command to create the new XDP-type relationship:

snapmirror create -source-path <SVM:volume> -destination-path
<SVM:volume> -type XDP -schedule <schedule> -policy <policy>

The new relationship must use the same source and destination volume. Learn more about the commands described in this procedure in the ONTAP command reference.



You must run this command from the destination SVM or the destination cluster.

The following example creates a SnapMirror disaster recovery relationship between the source volume volA on svm1 and the destination volume volA_dst on svm_backup using the default MirrorAllSnapshots policy:

```
cluster_dst::> snapmirror create -source-path svm1:volA -destination
-path svm_backup:volA_dst
-type XDP -schedule my_daily -policy MirrorAllSnapshots
```

9. Resync the source and destination volumes:

```
snapmirror resync -source-path <SVM:volume> -destination-path
<SVM:volume>
```

To improve resync time, you can use the -quick-resync option, but you should be aware that storage efficiency savings can be lost.



You must run this command from the destination SVM or the destination cluster. Although resync does not require a baseline transfer, it can be time-consuming. You might want to run the resync in off-peak hours.

The following example resyncs the relationship between the source volume volA on svm1 and the destination volume volA_dst on svm_backup:

cluster_dst::> snapmirror resync -source-path svm1:volA -destination
-path svm backup:volA dst

Learn more about snapmirror resync in the ONTAP command reference.

10. If you disabled automatic deletion of snapshots, reenable it:

```
volume snapshot autodelete modify -vserver <SVM> -volume <volume>
-enabled true
```

After you finish

1. Use the snapmirror show command to verify that the SnapMirror relationship was created.

Learn more about snapmirror show in the ONTAP command reference.

2. Once the SnapMirror XDP destination volume begins updating snapshots as defined by the SnapMirror policy, use the output of snapmirror list-destinations command from the source cluster to display the new SnapMirror XDP relationship.

Additional information about DP-type relationships

Beginning with ONTAP 9.3, XDP mode is the default, and any invocations of DP mode on the command line or in new or existing scripts are automatically converted to XDP mode.

Existing relationships are not affected. If a relationship is already of type DP, it will continue to be of type DP. Beginning with ONTAP 9.5, MirrorAndVault is the default policy when no data protection mode is specified or when XDP mode is specified as the relationship type. The table below shows the expected behavior.

If you specify	The type is	The default policy (if you do not specify a policy) is…
DP	XDP	MirrorAllSnapshots (SnapMirror DR)

Nothing	XDP	MirrorAndVault (unified replication)
XDP	XDP	MirrorAndVault (unified replication)

As the table shows, the default policies assigned to XDP in different circumstances ensure that the conversion maintains the functional equivalence of the previous types. Of course, you can use different policies as needed, including policies for unified replication:

If you specify	And the policy is	The result is
DP	MirrorAllSnapshots	SnapMirror DR
	XDPDefault	SnapVault
	MirrorAndVault	Unified replication
XDP	MirrorAllSnapshots	SnapMirror DR
	XDPDefault	SnapVault
	MirrorAndVault	Unified replication

The only exceptions to conversion are as follows:

• SVM data protection relationships continue to default to DP mode in ONTAP 9.3 and earlier.

Beginning with ONTAP 9.4, SVM data protection relationships default to XDP mode.

- Root volume load-sharing data protection relationships continue to default to DP mode.
- SnapLock data protection relationships continue to default to DP mode in ONTAP 9.4 and earlier.

Beginning with ONTAP 9.5, SnapLock data protection relationships default to XDP mode.

• Explicit invocations of DP continue to default to DP mode if you set the following cluster-wide option:

options replication.create_data_protection_rels.enable on

This option is ignored if you do not explicitly invoke DP.

Related information

- snapmirror create
- snapmirror delete
- snapmirror quiesce
- snapmirror release
- snapmirror resync

Convert the type of an ONTAP SnapMirror relationship

Beginning with ONTAP 9.5, SnapMirror synchronous is supported. You can convert an SnapMirror asynchronous relationship to a SnapMirror synchronous relationship or vice versa without performing a baseline transfer.

About this task

You cannot convert an SnapMirror asynchronous relationship to a SnapMirror synchronous relationship or vice versa by changing the SnapMirror policy.

Steps

- Converting an SnapMirror asynchronous relationship to a SnapMirror synchronous relationship
 - a. From the destination cluster, delete the SnapMirror asynchronous relationship:

```
snapmirror delete -destination-path <SVM:volume>
```

```
cluster2::>snapmirror delete -destination-path vs1_dr:vol1
```

b. From the source cluster, release the SnapMirror relationship without deleting the common snapshots:

```
snapmirror release -relationship-info-only true -destination-path
<destination SVM>:<destination volume>
```

```
cluster1::>snapmirror release -relationship-info-only true
-destination-path vs1 dr:vol1
```

c. From the destination cluster, create a SnapMirror synchronous relationship:

```
snapmirror create -source-path src_SVM:src_volume -destination-path
<destination SVM>:<destination volume> -policy sync-mirror
```

cluster2::>snapmirror create -source-path vs1:vol1 -destination-path vs1_dr:vol1 -policy sync

d. Resynchronize the SnapMirror synchronous relationship:

snapmirror resync -destination-path <destination_SVM:destination_volume>

cluster2::>snapmirror resync -destination-path vs1_dr:vol1

- Converting a SnapMirror synchronous relationship to an SnapMirror asynchronous relationship
 - a. From the destination cluster, quiesce the existing SnapMirror synchronous relationship:

snapmirror quiesce -destination-path <destination_SVM>:<destination_volume>

cluster2::> snapmirror quiesce -destination-path vs1 dr:vol1

b. From the destination cluster, delete the SnapMirror asynchronous relationship:

snapmirror delete -destination-path <SVM:volume>

cluster2::>snapmirror delete -destination-path vs1 dr:vol1

c. From the source cluster, release the SnapMirror relationship without deleting the common snapshots:

```
snapmirror release -relationship-info-only true -destination-path
<destination SVM:destination volume>
```

cluster1::>snapmirror release -relationship-info-only true
-destination-path vs1 dr:vol1

d. From the destination cluster, create an SnapMirror asynchronous relationship:

```
snapmirror create -source-path src_SVM:src_volume -destination-path
<destination SVM:destination volume> -policy MirrorAllSnapshots
```

```
cluster2::>snapmirror create -source-path vs1:vol1 -destination-path
vs1 dr:vol1 -policy sync
```

e. Resynchronize the SnapMirror synchronous relationship:

```
snapmirror resync -destination-path <destination_SVM:destination_volume>
```

cluster2::>snapmirror resync -destination-path vs1_dr:vol1

Related information

- snapmirror create
- snapmirror delete
- snapmirror quiesce
- snapmirror release
- snapmirror resync

Convert the mode of an ONTAP SnapMirror synchronous relationship

Beginning with ONTAP 9.5, SnapMirror synchronous relationships are supported. You can convert the mode of a SnapMirror synchronous relationship from StrictSync to Sync or vice versa.

About this task

You cannot modify the policy of a SnapMirror synchronous relationship to convert its mode.

Steps

1. From the destination cluster, quiesce the existing SnapMirror synchronous relationship:

snapmirror quiesce -destination-path <destination_SVM>:<destination_volume>

cluster2::> snapmirror quiesce -destination-path vs1 dr:vol1

2. From the destination cluster, delete the existing SnapMirror synchronous relationship:

snapmirror delete -destination-path <destination SVM>:<destination volume>

cluster2::> snapmirror delete -destination-path vs1 dr:vol1

3. From the source cluster, release the SnapMirror relationship without deleting the common snapshots:

snapmirror release -relationship-info-only true -destination-path
<destination SVM>:<destination volume>

```
cluster1::> snapmirror release -relationship-info-only true -destination
-path vs1 dr:vol1
```

4. From the destination cluster, create a SnapMirror synchronous relationship by specifying the mode to which you want to convert the SnapMirror synchronous relationship:

```
snapmirror create -source-path vs1:vol1 -destination-path
<destination_SVM>:<destination_volume> -policy Sync|StrictSync
```

cluster2::> snapmirror create -source-path vs1:vol1 -destination-path vs1_dr:vol1 -policy Sync

5. From the destination cluster, resynchronize the SnapMirror relationship:

snapmirror resync -destination-path <destination_SVM>:<destination_volume>

cluster2::> snapmirror resync -destination-path vs1_dr:vol1

Related information

- snapmirror create
- snapmirror delete

- snapmirror quiesce
- snapmirror release
- snapmirror resync

Create and delete ONTAP SnapMirror failover test volumes

Beginning with ONTAP 9.14.1, you can use System Manager to create a volume clone to test SnapMirror failover and disaster recovery without disrupting the active SnapMirror relationship. When you finish testing, you can clean up the associated data and delete the test volume.

Create a SnapMirror failover test volume

About this task

- You can perform failover tests on synchronous and SnapMirror asynchronous relationships.
- · A volume clone is created to perform the disaster recovery test.
- The clone volume is created on the same storage VM as the SnapMirror destination.
- You can use FlexVol and FlexGroup SnapMirror relationships.
- If a test clone already exists for the selected relationship, you cannot create another clone for that relationship.
- SnapLock vault relationships are not supported.

Before you begin

- You must be a cluster administrator.
- The SnapMirror license must be installed on the source and destination cluster.

Steps

- 1. On the destination cluster, select **Protection > Relationships**.
- 2. Select inext to the relationship source and choose Test Failover.
- 3. In the Test Failover window, select Test Failover.
- 4. Select **Storage > Volumes**, and verify that the test failover volume is listed.
- 5. Select Storage > Shares.
- 6. Select + Add and choose Share.
- 7. In the Add share window, type a name for the share in the Share Name field.
- 8. In the Folder field, select Browse, select the test clone volume, and Save.
- 9. At the bottom of the Add share window, choose Save.
- 10. In the **Storage > Shares** pane, locate the share you created and select v to view the share information.
- 11. Under SMB/CIFS Access, copy or make note of the access path for the share; for example, \\123.456.7.890\failover_test.
- 12. Use the SMB access path to open the share on the client and verify that the test volume has read and write capabilities.

Clean up failover data and delete the test volume

After you have completed failover testing, you can clean up all data associated with the test volume and delete it.

Steps

- 1. On the destination cluster, select **Protection > Relationships**.
- 2. Select inext to the relationship source and choose **Clean Up Test Failover**.
- 3. In the Clean Up Test Failover window, select Clean Up.
- 4. Select **Storage > Volumes** and verify that the test volume was deleted.

Serve data from a SnapMirror DR destination volume

Make the ONTAP SnapMirror destination volume writeable

You need to make the destination volume writeable before you can serve data from the volume to clients. To serve data from a mirror destination when a source becomes unavailable, stop scheduled transfers to the destination, and then break the SnapMirror relationship to make the destination writable.

About this task

You must perform this task from the destination SVM or the destination cluster.

Steps

You can use System Manager or the ONTAP CLI to make a destination volume writable.
System Manager

- 1. Select the protection relationship: click **Protection > Relationships**, and then click the desired volume name.
- 2. Click :
- 3. Stop scheduled transfers : click **Pause**.
- 4. Make the destination writable: click **Break**.
- 5. Go to the main **Relationships** page to verify that the relationship state displays as "broken off".

Next steps

You need to reverse resynchronize the replication relationship after you make a destination volume writeable.

When the disabled source volume is available again, you should reverse resynchronize the relationship again to copy the current data to the original source volume.

CLI

1. Stop scheduled transfers to the destination:

snapmirror quiesce -source-path <SVM:volume|cluster://SVM/volume>
-destination-path <SVM:volume|cluster://SVM/volume>

The following example stops scheduled transfers between the source volume volA on svm1 and the destination volume volA dst on svm backup:

```
cluster_dst::> snapmirror quiesce -source-path svm1:volA
-destination-path svm backup:volA dst
```

Learn more about snapmirror quiesce in the ONTAP command reference.

2. Stop ongoing transfers to the destination:

```
snapmirror abort -source-path <SVM:volume|cluster://SVM/volume>
-destination-path <SVM:volume|cluster://SVM/volume>
```



This step is not required for SnapMirror synchronous relationships (supported beginning with ONTAP 9.5).

The following example stops ongoing transfers between the source volume volA on svm1 and the destination volume volA_dst on svm_backup:

cluster_dst::> snapmirror abort -source-path svm1:volA -destination
-path svm backup:volA dst

Learn more about snapmirror abort in the ONTAP command reference.

3. Break the SnapMirror DR relationship:

```
snapmirror break -source-path <SVM:volume|cluster://SVM/volume>
-destination-path <SVM:volume|cluster://SVM/volume>
```

The following example breaks the relationship between the source volume volA on svm1 and the destination volume volA dst on svm backup:

cluster_dst::> snapmirror break -source-path svm1:volA -destination
-path svm backup:volA dst

Learn more about snapmirror break in the ONTAP command reference.

Next steps

You need to resynchronize the replication relationship after you make a destination volume writeable.

Other ways to do this in ONTAP

To perform these tasks with	See this content
System Manager Classic (available with ONTAP 9.7 and earlier)	Volume disaster recovery overview

Configure the ONTAP SnapMirror destination volume for data access

After making the destination volume writeable, you must configure the volume for data access. NAS clients, NVMe subsystem, and SAN hosts can access the data from the destination volume until the source volume is reactivated.

NAS environment:

- 1. Mount the NAS volume to the namespace using the same junction path that the source volume was mounted to in the source SVM.
- 2. Apply the appropriate ACLs to the SMB shares at the destination volume.
- 3. Assign the NFS export policies to the destination volume.
- 4. Apply the quota rules to the destination volume.
- 5. Redirect clients to the destination volume.
- 6. Remount the NFS and SMB shares on the clients.

SAN environment:

- 1. Map the LUNs in the volume to the appropriate initiator group.
- 2. For iSCSI, create iSCSI sessions from the SAN host initiators to the SAN LIFs.

3. On the SAN client, perform a storage re-scan to detect the connected LUNs.

For information about NVMe environment, see SAN administration.

Reactivate the original ONTAP SnapMirror source volume

You can reestablish the original data protection relationship between the source and destination volumes when you no longer need to serve data from the destination.

About this task

- The procedure below assumes that the baseline in the original source volume is intact. If the baseline is not intact, you must create and initialize the relationship between the volume you are serving data from and the original source volume before performing the procedure.
- Background preparation and the data warehousing phase of an XDP SnapMirror relationship can take a long time. It is not uncommon to see the SnapMirror relationship reporting the status "preparing" for an extended time period.

Steps

1. Reverse the original data protection relationship:

```
snapmirror resync -source-path SVM:volume -destination-path SVM:volume
```

Learn more about snapmirror resync in the ONTAP command reference.



You must run this command from the original source SVM or the original source cluster. Although resync does not require a baseline transfer, it can be time-consuming. You might want to run the resync in off-peak hours. The command fails if a common snapshot does not exist on the source and destination. Use snapmirror initialize to re-initialize the relationship. Learn more about snapmirror initialize in the ONTAP command reference.

The following example reverses the relationship between the original source volume, volA on svm1, and the volume you are serving data from, volA_dst on svm_backup:

cluster_src::> snapmirror resync -source-path svm_backup:volA_dst
-destination-path svm1:volA

2. When you are ready to reestablish data access to the original source, stop access to the original destination volume. One way to do this is to stop the original destination SVM:

```
vserver stop -vserver SVM
```



You must run this command from the original destination SVM or the original destination cluster. This command stops user access to the entire original destination SVM. You may want to stop access to the original destination volume using other methods.

The following example stops the original destination SVM:

cluster dst::> vserver stop svm backup

Learn more about vserver stop in the ONTAP command reference.

3. Update the reversed relationship:

snapmirror update -source-path SVM:volume -destination-path SVM:volume



You must run this command from the original source SVM or the original source cluster.

The following example updates the relationship between the volume you are serving data from, volA_dst on svm backup, and the original source volume, volA on svm1:

cluster_src::> snapmirror update -source-path svm_backup:volA_dst
-destination-path svm1:volA

Learn more about snapmirror update in the ONTAP command reference.

4. From the original source SVM or the original source cluster, stop scheduled transfers for the reversed relationship:

snapmirror quiesce -source-path SVM:volume -destination-path SVM:volume



You must run this command from the original source SVM or the original source cluster.

The following example stops scheduled transfers between the original destination volume, volA_dst on svm backup, and the original source volume, volA on svm1:

```
cluster_src::> snapmirror quiesce -source-path svm_backup:volA_dst
-destination-path svm1:volA
```

Learn more about snapmirror quiesce in the ONTAP command reference.

5. When the final update is complete and the relationship indicates "Quiesced" for the relationship status, run the following command from the original source SVM or the original source cluster to break the reversed relationship::

snapmirror break -source-path SVM:volume -destination-path SVM:volume



You must run this command from the original source SVM or the source cluster.

The following example breaks the relationship between the original destination volume, volA_dst on svm backup, and the original source volume, volA on svm1:

```
cluster_scr::> snapmirror break -source-path svm_backup:volA_dst
-destination-path svm1:volA
```

Learn more about snapmirror break in the ONTAP command reference.

From the original source SVM or the original source cluster, delete the reversed data protection relationship:

snapmirror delete -source-path SVM:volume -destination-path SVM:volume



You must run this command from the original source SVM or the original source cluster.

The following example deletes the reversed relationship between the original source volume, volA on svm1, and the volume you are serving data from, volA dst on svm backup:

```
cluster_src::> snapmirror delete -source-path svm_backup:volA_dst
-destination-path svm1:volA
```

Learn more about snapmirror delete in the ONTAP command reference.

7. Release the reversed relationship from the original destination SVM or the original destination cluster.

snapmirror release -source-path SVM:volume -destination-path SVM:volume



You must run this command from the original destination SVM or the original destination cluster.

The following example releases the reversed relationship between the original destination volume, volA dst on svm backup, and the original source volume, volA on svm1:

cluster_dst::> snapmirror release -source-path svm_backup:volA_dst -destination-path svm1:volA

Learn more about snapmirror release in the ONTAP command reference.

8. Reestablish the original data protection relationship from the original destination:

snapmirror resync -source-path SVM:volume -destination-path SVM:volume

The following example reestablishes the relationship between the original source volume, volA on svm1, and the original destination volume, volA_dst on svm_backup:

```
cluster_dst::> snapmirror resync -source-path svm1:volA -destination
-path svm backup:volA dst
```

Learn more about snapmirror resync in the ONTAP command reference.

9. If needed, start the original destination SVM:

vserver start -vserver SVM

The following example starts the original destination SVM:

cluster dst::> vserver start svm backup

Learn more about vserver start in the ONTAP command reference.

After you finish

Use the snapmirror show command to verify that the SnapMirror relationship was created.

Learn more about snapmirror show in the ONTAP command reference.

Restore files from a SnapMirror destination volume

Restore a file, LUN, or NVMe namespace from an ONTAP SnapMirror destination

You can restore a single file, LUN, a set of files or LUNs from a snapshot, or an NVMe namespace from a SnapMirror destination volume. Beginning with ONTAP 9.7, you can also restore NVMe namespaces from a SnapMirror synchronous destination. You can restore files to the original source volume or to a different volume.

Before you begin

To restore a file or LUN from a SnapMirror synchronous destination (supported beginning with ONTAP 9.5), you must first delete and release the relationship.

About this task

The volume to which you are restoring files or LUNs (the destination volume) must be a read-write volume:

- SnapMirror performs an *incremental restore* if the source and destination volumes have a common snapshot (as is typically the case when you are restoring to the original source volume).
- Otherwise, SnapMirror performs a *baseline restore*, in which the specified snapshot and all the data blocks it references are transferred to the destination volume.

Steps

1. List the snapshots in the destination volume:

volume snapshot show -vserver <SVM> -volume volume

Learn more about volume snapshot show in the ONTAP command reference.

The following example shows the snapshots on the vserverB:secondary1 destination:

<pre>cluster_dst::> volume snapshot show -vserver vserverB -volume secondary1</pre>						
Vserver Used%	Volume	Snapshot	State	Size	Total%	
vserverB 0%	secondary1	hourly.2013-01-25_0005	valid	224KB	0%	
0%		daily.2013-01-25_0010	valid	92KB	0%	
0%		hourly.2013-01-25_0105	valid	228KB	0%	
0%		hourly.2013-01-25_0205	valid	236KB	0%	
0%		hourly.2013-01-25_0305	valid	244KB	0%	
0%		hourly.2013-01-25_0405	valid	244KB	0%	
0%		hourly.2013-01-25_0505	valid	244KB	0%	
7 ontrios v	oro digolaria	d				
/ entries W	ете атвртауе	u.				

2. Restore a single file or LUN or a set of files or LUNs from a snapshot in a SnapMirror destination volume:

```
snapmirror restore -source-path <SVM:volume>|<cluster://SVM/volume>, ...
-destination-path <SVM:volume>|<cluster://SVM/volume>, ... -source-snapshot
snapshot -file-list <source_file_path,@destination_file_path>
```

 (\mathbf{i})

You must run this command from the destination SVM or the destination cluster.

The following command restores the files file1 and file2 from the snapshot daily.2013-01-25_0010 in the original destination volume secondary1, to the same location in the active file system of the original source volume primary1:

```
cluster_dst::> snapmirror restore -source-path vserverB:secondary1
-destination-path vserverA:primary1 -source-snapshot daily.2013-01-
25_0010 -file-list /dir1/file1,/dir2/file2
[Job 3479] Job is queued: snapmirror restore for the relationship with
destination vserverA:primary1
```

The following command restores the files file1 and file2 from the snapshot daily.2013-01-25_0010 in the original destination volume secondary1, to a different location in the active file system of the original source volume primary1. The destination file path begins with the @ symbol followed by the path of the file from the root of the original source volume. In this example, file1 is restored to /dir1/file1.new and file2 is restored to /dir2.new/file2 on primary1:

```
cluster_dst::> snapmirror restore -source-path vserverB:secondary1
-destination-path vserverA:primary1 -source-snapshot daily.2013-01-
25_0010 -file-list
/dir/file1,@/dir1/file1.new,/dir2/file2,@/dir2.new/file2
[Job 3479] Job is queued: snapmirror restore for the relationship with
destination vserverA:primary1
```

The following command restores the files file1 and file3 from the snapshot daily.2013-01-25_0010 in the original destination volume secondary1, to different locations in the active file system of the original source volume primary1, and restores file2 from snap1 to the same location in the active file system of primary1.

In this example, the file file1 is restored to /dir1/file1.new and file3 is restored to /dir3.new/file3:

cluster_dst::> snapmirror restore -source-path vserverB:secondary1 -destination-path vserverA:primary1 -source-snapshot daily.2013-01-25_0010 -file-list /dir/file1,@/dir1/file1.new,/dir2/file2,/dir3/file3,@/dir3.new/file3 [Job 3479] Job is queued: snapmirror restore for the relationship with destination vserverA:primary1

Related information

snapmirror restore

Restore volume contents from an ONTAP SnapMirror destination

You can restore the contents of an entire volume from a snapshot in a SnapMirror destination volume. You can restore the volume's contents to the original source volume or to a different volume.

About this task

This procedure applies to FAS, AFF, and ASA systems. If you have an ASA r2 system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, ASA A20, or ASA C30), follow these steps to restore data. ASA r2 systems provide a simplified ONTAP experience specific to SAN-only customers.

The destination volume for the restore operation must be one of the following:

• A read-write volume, in which case SnapMirror performs an *incremental restore*, provided that the source and destination volumes have a common snapshot (as is typically the case when you are restoring to the original source volume).



The command fails if there is not a common snapshot. You cannot restore the contents of a volume to an empty read-write volume.

• An empty data protection volume, in which case SnapMirror performs a *baseline restore*, in which the specified snapshot and all the data blocks it references are transferred to the source volume.

Restoring the contents of a volume is a disruptive operation. SMB traffic must not be running on the SnapVault primary volume when a restore operation is running.

If the destination volume for the restore operation has compression enabled, and the source volume does not have compression enabled, disable compression on the destination volume. You need to re-enable compression after the restore operation is complete.

Any quota rules defined for the destination volume are deactivated before the restore is performed. You can use the volume guota modify command to reactivate quota rules after the restore operation is complete.

When data in a volume is lost or corrupted, you can roll back your data by restoring from an earlier snapshot.

This procedure replaces the current data on the source volume with data from an earlier snapshot version. You should perform this task on the destination cluster.

Steps

You can restore a volume's contents using System Manager or the ONTAP CLI.

System Manager

- 1. Click **Protection > Relationships**, and then click the source volume name.
- 2. Click and then select **Restore**.
- 3. Under **Source**, the source volume is selected by default. Click **Other Volume** if you want to choose a volume other than the source.
- 4. Under **Destination**, choose the snapshot you want to restore.
- 5. If your source and destination are located on different clusters, on the remote cluster, click Protection
 > Relationships to monitor the restore progress.

CLI

1. List the snapshots in the destination volume:

```
volume snapshot show -vserver <SVM> -volume <volume>
```

The following example shows the snapshots on the vserverB:secondary1 destination:

```
cluster dst::> volume snapshot show -vserver vserverB -volume
secondary1
Vserver Volume
                      Snapshot
                                           State Size
Total% Used%
_____
           _____
                       _____ _
                                                     ____
_____ ___
vserverB secondary1 hourly.2013-01-25 0005 valid
                                                    224KB
                                                              0%
0%
                      daily.2013-01-25 0010 valid
                                                    92KB
                                                              0%
0%
                      hourly.2013-01-25 0105 valid
                                                    228KB
                                                              0%
0%
                      hourly.2013-01-25 0205 valid
                                                              0%
                                                    236KB
0 %
                      hourly.2013-01-25 0305 valid
                                                    244KB
                                                              0%
0%
                      hourly.2013-01-25 0405 valid
                                                    244KB
                                                              0%
0 %
                      hourly.2013-01-25 0505 valid
                                                    244KB
                                                              0%
0%
7 entries were displayed.
```

2. Restore the contents of a volume from a snapshot in a SnapMirror destination volume:

snapmirror restore -source-path <SVM:volume>|<cluster://SVM/volume> -destination-path <SVM:volume>|<cluster://SVM/volume> -source-snapshot <snapshot>

You must run this command from the original source SVM or the original source cluster.

The following command restores the contents of the original source volume primary1 from the snapshot daily.2013-01-25 0010 in the original destination volume secondary1:

```
cluster_src::> snapmirror restore -source-path vserverB:secondary1
-destination-path vserverA:primary1 -source-snapshot daily.2013-01-
25_0010
Warning: All data newer than snapshot daily.2013-01-25_0010 on
volume vserverA:primary1 will be deleted.
Do you want to continue? {y|n}: y
[Job 34] Job is queued: snapmirror restore from source
vserverB:secondary1 for the snapshot daily.2013-01-25_0010.
```

3. Remount the restored volume and restart all applications that use the volume.

Other ways to do this in ONTAP

To perform these tasks with	See this content
System Manager Classic (available with ONTAP 9.7 and earlier)	Volume restore using SnapVault overview

Related information

- snapmirror restore
- volume snapshot show

Update an ONTAP SnapMirror replication relationship manually

You might need to update a replication relationship manually if an update fails because the source volume has been moved.

About this task

SnapMirror aborts any transfers from a moved source volume until you update the replication relationship manually.

Beginning with ONTAP 9.5, SnapMirror synchronous relationships are supported. Although the source and destination volumes are in sync at all times in these relationships, the view from the secondary cluster is synchronized with the primary only on an hourly basis. If you want to view the point-in-time data at the destination, you should perform a manual update by running the snapmirror update command.

Step

1. Update a replication relationship manually:

```
snapmirror update -source-path <SVM:volume>|<cluster://SVM/volume>, ...
-destination-path <SVM:volume>|<cluster://SVM/volume>, ...
```



You must run this command from the destination SVM or the destination cluster. The command fails if a common snapshot does not exist on the source and destination. Use snapmirror initialize to re-initialize the relationship. Learn more about snapmirror initialize in the ONTAP command reference.

The following example updates the relationship between the source volume volA on svml and the destination volume volA_dst on svm_backup:

```
cluster_src::> snapmirror update -source-path svm1:volA -destination
-path svm_backup:volA_dst
```

Learn more about snapmirror update in the ONTAP command reference.

Resynchronize an ONTAP SnapMirror replication relationship

You need to resynchronize a replication relationship after you make a destination volume writeable, after an update fails because a common snapshot does not exist on the source and destination volumes, or if you want to change the replication policy for the relationship.

Beginning with ONTAP 9.8, you can use System Manager to perform a reverse resynchronization operation to delete an existing protection relationship and reverse the functions of the source and destination volumes. Then you use the destination volume to serve data while you repair or replace the source, update the source, and reestablish the original configuration of the systems.

About this task

- Although resync does not require a baseline transfer, it can be time-consuming. You might want to run the resync in off-peak hours.
- Volumes that are part of a fan-out or cascade configuration can take longer to resynchronize. It is not uncommon to see the SnapMirror relationship reporting the status "preparing" for an extended time period.



System Manager does not support reverse resynchronization with intracluster relationships. You can use the ONTAP CLI to perform reverse resync operations with intracluster relationships.

Steps

You can use System Manager or the ONTAP CLI to perform this task. If you use the ONTAP CLI, the procedure is the same regardless of whether you are making a destination volume writable or you are updating the replication relationship.

System Manager reverse resync

After you break a relationship to make a destination writable, reverse resynchronize the relationship:

- 1. On the destination cluster, click **Protection > Relationships**.
- 2. Hover over the broken off relationship you want to reverse, click **‡**, and select **Reverse Resync**.
- 3. In the Reverse resync relationship window, click Reverse resync.
- 4. Under **Relationships**, monitor the reverse resynchronization progress by viewing **Transfer Status** for the relationship.

Next steps

When the original source is available again, you can reestablish the original relationship by breaking the reversed relationship and performing another reverse resync operation. The reverse resync process will copy any changes from the site that is serving data to the original source and make the original source read-writable again.

System Manager resync

- 1. Click **Protection > Relationships**.
- 2. Hover over the relationship you want to resynchronize, and click and then select **Break**.
- 3. When the relationship state displays "Broken off," click and then select **Resync**.
- 4. Under **Relationships**, monitor the resynchronization progress by checking the relationship state. The state changes to "Mirrored" when resynchronization is complete.

CLI

1. Resync the source and destination volumes:

```
snapmirror resync -source-path <SVM:volume|cluster://SVM/volume>
-destination-path <SVM:volume|cluster://SVM/volume> -type DP|XDP
-policy <policy>
```



You must run this command from the destination SVM or the destination cluster.

The following example resynchronizes the relationship between the source volume volA on svm1 and the destination volume volA dst on svm backup:

cluster_dst::> snapmirror resync -source-path svm1:volA -destination
-path svm backup:volA dst

Learn more about snapmirror resync in the ONTAP command reference.

Delete an ONTAP SnapMirror volume replication relationship

You can use the snapmirror delete and snapmirror release commands to delete a volume replication relationship. You can then delete unneeded destination

volumes manually.

About this task

The snapmirror release command deletes any SnapMirror-created snapshots from the source. You can use the -relationship-info-only option to preserve the snapshots.

Steps

1. Quiesce the replication relationship:

```
snapmirror quiesce -destination-path <SVM:volume>|<cluster://SVM/volume>
```

```
cluster dst::> snapmirror quiesce -destination-path svm backup:volA dst
```

Learn more about snapmirror quiesce in the ONTAP command reference.

 (Optional) Break the replication relationship if you require the destination volume to be a read/write volume. You can skip this step if you plan to delete the destination volume or if you don't need the volume to be read/write:

```
snapmirror break -source-path <SVM:volume>|<cluster://SVM/volume>, ...
-destination-path <SVM:volume>|<cluster://SVM/volume>, ...
```

```
cluster_dst::> snapmirror break -source-path svm1:volA -destination-path
svm backup:volA dst
```

Learn more about snapmirror break in the ONTAP command reference.

3. Delete the replication relationship:

```
snapmirror delete -source-path <SVM:volume>|<cluster://SVM/volume>, ...
-destination-path <SVM:volume>|<cluster://SVM/volume>, ...
```



You must run this command from the destination cluster or destination SVM.

The following example deletes the relationship between the source volume volA on svm1 and the destination volume volA_dst on svm_backup:

```
cluster_dst::> snapmirror delete -source-path svm1:volA -destination
-path svm backup:volA dst
```

Learn more about snapmirror delete in the ONTAP command reference.

4. Release replication relationship information from the source SVM:

```
snapmirror release -source-path <SVM:volume>|<cluster://SVM/volume>, ...
-destination-path <SVM:volume>|<cluster://SVM/volume>, ...
```



You must run this command from the source cluster or source SVM.

The following example releases information for the specified replication relationship from the source SVM svm1:

cluster_src::> snapmirror release -source-path svm1:volA -destination
-path svm backup:volA dst

Learn more about snapmirror release in the ONTAP command reference.

Manage storage efficiency on ONTAP SnapMirror volumes

SnapMirror preserves storage efficiency on the source and destination volumes except when postprocess data compression is enabled on the destination volume. In that case, all storage efficiency is lost on the destination volume. To correct this issue, you need to disable postprocess compression on the destination volume, update the relationship manually, and re-enable storage efficiency.

About this task

You can use the volume efficiency show command to determine whether efficiency is enabled on a volume.

Learn more about volume efficiency show in the ONTAP command reference.

You can check if SnapMirror is maintaining storage efficiency by viewing the SnapMirror audit logs and locating the transfer description. If the transfer description displays transfer_desc=Logical Transfer with Storage Efficiency, SnapMirror is maintaining storage efficiency. If the transfer description displays transfer_desc=Logical Transfer, SnapMirror is not maintaining storage efficiency. For example:

```
Fri May 22 02:13:02 CDT 2020 ScheduledUpdate[May 22 02:12:00]:cc0fbc29-
b665-11e5-a626-00a09860c273 Operation-Uuid=39fbcf48-550a-4282-a906-
df35632c73a1 Group=none Operation-Cookie=0 action=End source=<sourcepath>
destination=<destpath> status=Success bytes_transferred=117080571
network_compression_ratio=1.0:1 transfer_desc=Logical Transfer - Optimized
Directory Mode
```

Before you begin

• The source and destination clusters and SVMs must be peered.

Cluster and SVM peering

- · You must disable postprocess compression on the destination volume.
- Logical Transfer with storage: Beginning with ONTAP 9.3, manual update is no longer required to re-enable storage efficiency. If SnapMirror detects that postprocess compression has been disabled, it automatically re-enables storage efficiency at the next scheduled update. Both the source and the destination must be running ONTAP 9.3.
- Beginning with ONTAP 9.3, AFF systems manage storage efficiency settings differently from FAS systems

after a destination volume is made writeable:

• After you make a destination volume writeable using the snapmirror break command, the caching policy on the volume is automatically set to auto (the default).



This behavior is applicable to FlexVol volumes, only, and it does not apply to FlexGroup volumes.

Learn more about snapmirror break in the ONTAP command reference.

 On resync, the caching policy is automatically set to none, and deduplication and inline compression are automatically disabled, regardless of your original settings. You must modify the settings manually as needed.



Manual updates with storage efficiency enabled can be time-consuming. You might want to run the operation in off-peak hours.

Steps

1. Update a replication relationship and re-enable storage efficiency:

```
snapmirror update -source-path <SVM:volume>|<cluster://SVM/volume>, ...
-destination-path <SVM:volume>|<cluster://SVM/volume>, ... -enable
-storage-efficiency true
```



You must run this command from the destination SVM or the destination cluster. The command fails if a common snapshot does not exist on the source and destination. Use snapmirror initialize to re-initialize the relationship. Learn more about snapmirror initialize in the ONTAP command reference.

The following example updates the relationship between the source volume volA on svm1 and the destination volume volA_dst on svm_backup, and re-enables storage efficiency:

```
cluster_dst::> snapmirror update -source-path svm1:volA -destination
-path svm backup:volA dst -enable-storage-efficiency true
```

Learn more about snapmirror update in the ONTAP command reference.

Use ONTAP SnapMirror global throttling

Global network throttling is available for all SnapMirror and SnapVault transfers at a pernode level.

About this task

SnapMirror global throttling restricts the bandwidth used by incoming and/or outgoing SnapMirror and SnapVault transfers. The restriction is enforced cluster wide on all nodes in the cluster.

For example, if the outgoing throttle is set to 100 MBps, each node in the cluster will have the outgoing

bandwidth set to 100 MBps. If global throttling is disabled, it is disabled on all nodes.

Although data transfer rates are often expressed in bits per second (bps), the throttle values must be entered in kilobytes per second (KBps).



In ONTAP 9.9.1 and earlier releases, the throttle has no effect on volume move transfers or load-sharing mirror transfers. Beginning with ONTAP 9.10.0, you can specify an option to throttle a volume move operation. For details, see How to throttle volume move in ONTAP 9.10 and later.

Global throttling works with the per-relationship throttle feature for SnapMirror and SnapVault transfers. The per-relationship throttle is enforced until the combined bandwidth of per-relationship transfers exceeds the value of the global throttle, after which the global throttle is enforced. A throttle value 0 implies that global throttling is disabled.



SnapMirror global throttling has no effect on SnapMirror synchronous relationships when they are In-Sync. However, the throttle does effect SnapMirror synchronous relationships when they perform an asynchronous transfer phase such as an initialization operation or after an Out Of Sync event. For this reason, enabling global throttling with SnapMirror synchronous relationships is not recommended.

Steps

1. Enable global throttling:

```
options -option-name replication.throttle.enable on | off
```

The following example shows how to enable SnapMirror global throttling on cluster_dst:

cluster dst::> options -option-name replication.throttle.enable on

2. Specify the maximum total bandwidth used by incoming transfers on the destination cluster:

options -option-name replication.throttle.incoming.max kbs <KBps>

The recommended minimum throttle bandwidth is 4 kilobytes per second (KBps) and the maximum is up to 2 terabytes per second (TBps). The default value for this option is unlimited, which means there is no limit on total bandwidth used.

The following example shows how to set the maximum total bandwidth used by incoming transfers to 100 megabits per second (Mbps):

```
cluster_dst::> options -option-name
replication.throttle.incoming.max_kbs 12500
```



100 megabits per second (Mbps) = 12500 kilobytes per second (KBps)

3. Specify the maximum total bandwidth used by outgoing transfers on the source cluster:

options -option-name replication.throttle.outgoing.max kbs <KBps>

The recommended minimum throttle bandwidth is 4 KBps and the maximum is up to 2 TBps. The default value for this option is unlimited, which means there is no limit on total bandwidth used. Parameter values are in kilobytes per second (KBps).

The following example shows how to set the maximum total bandwidth used by outgoing transfers to 100 Mbps:

cluster_src::> options -option-name
replication.throttle.outgoing.max_kbs 12500

Manage SnapMirror SVM replication

Learn about ONTAP SnapMirror SVM replication

You can use SnapMirror to create a data protection relationship between SVMs. In this type of data protection relationship, all or part of the SVM's configuration, from NFS exports and SMB shares to RBAC, is replicated, as well as the data in the volumes that the SVM owns.

Supported relationship types

Only data-serving SVMs can be replicated. The following data protection relationship types are supported:

• SnapMirror DR, in which the destination typically contains only the snapshots currently on the source.

Beginning with ONTAP 9.9.1, this behavior changes when you are using the mirror-vault policy. Beginning with ONTAP 9.9.1, you can create different snapshot policies on the source and destination, and the snapshots on the destination are not overwritten by snapshots on the source:

- They are not overwritten from the source to the destination during normal scheduled operations, updates and resync
- $\circ\,$ They are not deleted during break operations.
- They are not deleted during flip-resync operations.
 When you configure an SVM disaster relationship using the mirror-vault policy using ONTAP 9.9.1 and later, the policy behaves as follows:
- User-defined snapshot policies at the source are not copied to the destination.
- System-defined snapshot policies are not copied to the destination.
- Volume association with user and system defined snapshot policies are not copied to the destination.

SVM.

• SnapMirror unified replication, in which the destination is configured for both DR and long-term retention.

For more information about SnapMirror unified replication, see SnapMirror unified replication basics.

The *policy type* of the replication policy determines the type of relationship it supports. The following table shows the available policy types.

Policy type	Relationship type
async-mirror	SnapMirror DR
mirror-vault	Unified replication

XDP replaces DP as the SVM replication default in ONTAP 9.4

Beginning with ONTAP 9.4, SVM data protection relationships default to XDP mode. SVM data protection relationships continue to default to DP mode in ONTAP 9.3 and earlier.

Existing relationships are not affected by the XDP default. If a relationship is already of type DP, it will continue to be of type DP. The following table shows the behavior you can expect.

If you specify	The type is…	The default policy (if you do not specify a policy) is…
DP	XDP	MirrorAllSnapshots (SnapMirror DR)
Nothing	XDP	MirrorAllSnapshots (SnapMirror DR)
XDP	XDP	MirrorAndVault (Unified replication)

You can find information about converting DP relationships to XDP relationships and other details here: Convert an existing ONTAP DP-type relationship to XDP.



Version-independence is not supported for SVM replication. In an SVM disaster recovery configuration, the destination SVM must be on a cluster running the same ONTAP version as the source SVM cluster to support failover and fail back operations.

Compatible ONTAP versions for SnapMirror relationships

How SVM configurations are replicated

The content of an SVM replication relationship is determined by the interaction of the following fields:

• The -identity-preserve true option of the snapmirror create command replicates the entire SVM configuration.

The -identity-preserve false option replicates only the volumes and authentication and authorization configurations of the SVM, and the protocol and name service settings listed in Configurations replicated in SVM disaster recovery relationships.

- The -discard-configs network option of the snapmirror policy create command excludes LIFs and related network settings from SVM replication, for use in cases where the source and destination SVMs are in different subnets.
- The -vserver-dr-protection unprotected option of the volume modify command excludes the specified volume from SVM replication.

Otherwise, SVM replication is almost identical to volume replication. You can use virtually the same workflow for SVM replication as you use for volume replication.

Support details

The following table shows support details for SnapMirror SVM replication.

Resource or feature	Support details
Deployment types	 Single source to single destination Beginning with ONTAP 9.4, fan-out. You can fan- out to two destinations only. By default, only one -identity-preserve true relationship is allowed per source SVM.
Relationship types	SnapMirror disaster recoverySnapMirror unified replication
Replication scope	Intercluster only. You cannot replicate SVMs in the same cluster.
Autonomous Ransomware Protection	 Supported beginning with ONTAP 9.12.1. For more information, see Autonomous Ransomware Protection.
Consistency groups asynchronous support	Beginning with ONTAP 9.14.1, a maximum of 32 SVM disaster recovery relationships are supported when consistency groups exist. See Protect a consistency group and Consistency group limits for more information.
FabricPool	Beginning with ONTAP 9.6, SnapMirror SVM replication is supported with FabricPool. When in an SVM DR relationship, source and destination volumes do not need to use FabricPool aggregates, but they must use the same tiering policy.
	Beginning with ONTAP 9.12.1, SnapMirror SVM replication is supported with FabricPool and FlexGroup volumes working in conjunction. Prior to 9.12.1, any two of these features worked together, but not all three together.

MetroCluster	 Beginning with ONTAP 9.11.1, both sides of a SVM disaster recovery relationship within a MetroCluster configuration can act as a source for additional SVM disaster recovery configurations. Beginning with ONTAP 9.5, SnapMirror SVM replication is supported on MetroCluster configurations. In releases earlier than ONTAP 9.10.X, a MetroCluster configuration cannot be the destination of an SVM disaster recovery relationship. In ONTAP 9.10.1 and later releases, a MetroCluster configuration can be the destination of an SVM disaster recovery relationship for migration purposes only, and it must meet all necessary requirements described in TR-4966: Migrating a SVM within a MetroCluster configuration can be the source of an SVM disaster recovery relationship. Only an active SVM within a MetroCluster configuration can be the source of an SVM disaster recovery relationship. A source can be a sync-source SVM before switchover or a sync-destination SVM cannot be the source of an SVM disaster recovery relationship. When a MetroCluster configuration is in a steady state, the MetroCluster sync-destination SVM cannot be the source of an SVM disaster recovery relationship, since the volumes are not online. When the sync-source SVM is the source of an SVM disaster recovery relationship, the source SVM disaster recovery relationship information is replicated to the MetroCluster partner. During the switchover and switchback processes, replication to the SVM disaster recovery destination might fail. However, after the switchover or switchback processes, recovery scheduled updates will succeed.
Consistency group	Supported beginning with ONTAP 9.14.1. For more information, see Protect a consistency group
ONTAP S3	Not supported with SVM disaster recovery.
SponMirror Symphremetry	Not ourported with CV/A4 disaster reservery.
Snapiviirror Synchronous	Not supported with SVM disaster recovery.

Version-independence	Not supported.
Volume encryption	• Encrypted volumes on the source are encrypted on the destination.
	 Onboard Key Manager or KMIP servers must be configured on the destination.
	 New encryption keys are generated at the destination.
	 If the destination does not contain a node that supports volume .encryption, replication succeeds, but the destination volumes are not encrypted.

Configurations replicated in SVM disaster recovery relationships

The following table shows the interaction of the snapmirror create -identity-preserve option and the snapmirror policy create -discard-configs network option:

Configuration replicated		-identity-preserve true		-identity-preser ve false
		Policy without -discard -configs network Set	Policy with -discard -configs network Set	
Network	NAS LIFs	Yes	No	No
	LIF Kerberos configuration	Yes	No	No
	SAN LIFs	No	No	No
	Firewall policies	Yes	Yes	No
	Service policies	Yes	Yes	No
	Routes	Yes	No	No
	Broadcast domain	No	No	No
	Subnet	No	No	No
	IPspace	No	No	No

SMB	SMB server	Yes	Yes	No
	Local groups and local user	Yes	Yes	Yes
	Privilege	Yes	Yes	Yes
	Shadow copy	Yes	Yes	Yes
	BranchCache	Yes	Yes	Yes
	Server options	Yes	Yes	Yes
	Server security	Yes	Yes	No
	Home directory, share	Yes	Yes	Yes
	Symlink	Yes	Yes	Yes
	Fpolicy policy, Fsecurity policy, and Fsecurity NTFS	Yes	Yes	Yes
	Name mapping and group mapping	Yes	Yes	Yes
	Audit information	Yes	Yes	Yes
NFS	Export policies	Yes	Yes	No
	Export policy rules	Yes	Yes	No
	NFS server	Yes	Yes	No
RBAC	Security certificates	Yes	Yes	No
	Login user, public key, role, and role configuration	Yes	Yes	Yes
	SSL	Yes	Yes	No

Name services	DNS and DNS hosts	Yes	Yes	No
	UNIX user and UNIX group	Yes	Yes	Yes
	Kerberos realm and Kerberos keyblocks	Yes	Yes	No
	LDAP and LDAP client	Yes	Yes	No
	Netgroup	Yes	Yes	No
	NIS	Yes	Yes	No
	Web and web access	Yes	Yes	No
Volume	Object	Yes	Yes	Yes
	Snapshots and snapshot policy	Yes	Yes	Yes
	Autodelete policy	No	No	No
	Efficiency policy	Yes	Yes	Yes
	Quota policy and quota policy rule	Yes	Yes	Yes
	Recovery queue	Yes	Yes	Yes

Root volume	Namespace	Yes	Yes	Yes
	User data	No	No	No
	Qtrees	No	No	No
	Quotas	No	No	No
	File-level QoS	No	No	No
	Attributes: state of the root volume, space guarantee, size, autosize, and total number of files	No	No	No
Storage QoS	QoS policy group	Yes	Yes	Yes
Fibre Channel (FC)		No	No	No
iSCSI		No	No	No
LUNs	Object	Yes	Yes	Yes
	igroups	No	No	No
	portsets	No	No	No
	Serial numbers	No	No	No
SNMP	v3 users	Yes	Yes	No

SVM disaster recovery storage limits

The following table shows the recommended maximum number of volumes and SVM disaster recovery relationships supported per storage object. You should be aware that limits are often platform dependent. Refer to the Hardware Universe to learn the limits for your specific configuration.

Storage object	Limit
SVM	300 Flexible volumes
HA pair	1,000 Flexible Volumes
Cluster	128 SVM disaster relationships

Related information

- snapmirror create
- snapmirror policy create

Replicate SVM configurations

ONTAP SnapMirror SVM replication workflow

SnapMirror SVM replication involves creating the destination SVM, creating a replication job schedule, and creating and initializing a SnapMirror relationship.

You should determine which replication workflow best suits your needs:

- Replicate an entire SVM configuration
- Exclude LIFs and related network settings from SVM replication
- Exlude network, name service, and other settings from SVM configuration

Criteria for placing volumes on ONTAP SnapMirror destination SVMs

When replicating volumes from the source SVM to the destination SVM, it's important to know the criteria for selecting aggregates.

Aggregates are selected based on the following criteria:

- Volumes are always placed on non-root aggregates.
- Non-root aggregates are selected based on the available free space and the number of volumes already hosted on the aggregate.

Aggregates with more free space and fewer volumes are given priority. The aggregate with the highest priority is selected.

- Source volumes on FabricPool aggregates are placed on FabricPool aggregates on the destination with the same tiering-policy.
- If a volume on the source SVM is located on a Flash Pool aggregate, then the volume is placed on a Flash Pool aggregate on the destination SVM, if such an aggregate exists and has enough free space.
- If the -space-guarantee option of the volume that is replicated is set to volume, only aggregates with free space greater than the volume size are considered.
- The volume size grows automatically on the destination SVM during replication, based on the source volume size.

If you want to pre-reserve the size on the destination SVM, you must resize the volume. The volume size does not shrink automatically on the destination SVM based on the source SVM.

If you want to move a volume from one aggregate to another, you can use the volume move command on the destination SVM.

Replicate an entire ONTAP SVM configuration

You can create an SVM disaster recovery (SVM DR) relationship to replicate one SVM

configuration to another. In the event of a disaster at the primary site, you can quickly activate the destination SVM.

Before you begin

The source and destination clusters and SVMs must be peered. For more information, see Create a cluster peer relationship and Create an SVM intercluster peer relationship.

Learn more about the commands described in this procedure in the ONTAP command reference.

About this task

This workflow assumes that you are already using a default policy or a custom replication policy.

Beginning with ONTAP 9.9.1, when you use the mirror-vault policy, you can create different snapshot policies on the source and destination SVM, and the snapshots on the destination are not overwritten by snapshots on the source. For more information, see Understanding SnapMirror SVM replication.

Complete this procedure from the destination. If you need to create a new protection policy, for instance, when your source storage VM has SMB configured, you should create the policy and use the **Identity preserve** option.

For details see Create custom data protection policies.

Steps

You can perform this task from System Manager or the ONTAP CLI.

System Manager

- 1. On the destination cluster, click **Protection > Relationships**.
- 2. Under Relationships, click Protect and choose Storage VMs (DR).
- Select a protection policy. If you created a custom protection policy, select it, then choose the source cluster and storage VM you want to replicate. You can also create a new destination storage VM by entering a new storage VM name.
- 4. If desired, change the destination settings to override identity preserve and to include or exclude network interfaces and protocols.

5. Click Save.

CLI

1. Create a destination SVM:

vserver create -vserver <SVM name> -subtype dp-destination

The SVM name must be unique across the source and destination clusters.

The following example creates a destination SVM named svm backup:

```
cluster_dst:> vserver create -vserver svm_backup -subtype dp-
destination
```

Learn more about vserver create in the ONTAP command reference.

2. From the destination cluster, create an SVM peer relationship using the vserver peer create command.

For more information, see Create an SVM intercluster peer relationship.

Learn more about vserver peer create in the ONTAP command reference.

3. Create a replication job schedule:

job schedule cron create -name <job_name> -month <month> -dayofweek
<day_of_week> -day <day_of_month> -hour <hour> -minute <minute>

For -month, -dayofweek, and -hour, you can specify all to run the job every month, day of the week, and hour, respectively.



The minimum supported schedule (RPO) for FlexVol volumes in an SVM SnapMirror relationship is 15 minutes. The minimum supported schedule (RPO) for FlexGroup volumes in an SVM SnapMirror relationship is 30 minutes.

The following example creates a job schedule named my_weekly that runs on Saturdays at 3:00 a.m.:

```
cluster_dst::> job schedule cron create -name my_weekly -dayofweek
saturday -hour 3 -minute 0
```

Learn more about job schedule cron create in the ONTAP command reference.

4. From the destination SVM or the destination cluster, create a replication relationship:

```
snapmirror create -source-path <SVM_name>: -destination-path
<SVM_name>: -type <DP|XDP> -schedule <schedule> -policy <policy>
-identity-preserve true
```



You must enter a colon (:) after the SVM name in the -source-path and -destination-path options.

The following example creates a SnapMirror DR relationship using the default MirrorAllSnapshots policy:

```
cluster_dst::> snapmirror create -source-path svm1: -destination
-path svm_backup: -type XDP -schedule my_daily -policy
MirrorAllSnapshots -identity-preserve true
```

The following example creates a unified replication relationship using the default MirrorAndVault policy:

cluster_dst:> snapmirror create -source-path svml: -destination-path svm_backup: -type XDP -schedule my_daily -policy MirrorAndVault -identity-preserve true

Assuming you have created a custom policy with the policy type <code>async-mirror</code>, the following example creates a SnapMirror DR relationship:

```
cluster_dst::> snapmirror create -source-path svm1: -destination
-path svm_backup: -type XDP -schedule my_daily -policy my_mirrored
-identity-preserve true
```

Assuming you have created a custom policy with the policy type mirror-vault, the following example creates a unified replication relationship:

```
cluster_dst::> snapmirror create -source-path svm1: -destination
-path svm_backup: -type XDP -schedule my_daily -policy my_unified
-identity-preserve true
```

Learn more about snapmirror create in the ONTAP command reference.

5. Stop the destination SVM:

vserver stop -vserver <SVM name>

The following example stops a destination SVM named svm_backup:

cluster_dst::> vserver stop -vserver svm_backup

Learn more about vserver stop in the ONTAP command reference.

6. From the destination SVM or the destination cluster, initialize the SVM replication relationship:

```
snapmirror initialize -source-path <SVM_name>: -destination-path
<SVM name>:
```



You must enter a colon (:) after the SVM name in the -source-path and -destination-path options.

The following example initializes the relationship between the source SVM, svml, and the destination SVM, svm_{backup} :

```
cluster_dst::> snapmirror initialize -source-path svm1: -destination
-path svm backup:
```

Learn more about snapmirror initialize in the ONTAP command reference.

Exclude LIFs and related network settings from ONTAP SnapMirror SVM replication

If the source and destination SVMs are in different subnets, you can use the -discard -configs network option of the snapmirror policy create command to exclude LIFs and related network settings from SVM replication.

Before you begin

The source and destination clusters and SVMs must be peered.

For more information, see Create a cluster peer relationship and Create an SVM intercluster peer relationship.

About this task

The -identity-preserve option of the snapmirror create command must be set to true when you create the SVM replication relationship.

Steps

1. Create a destination SVM:

vserver create -vserver SVM -subtype dp-destination

The SVM name must be unique across the source and destination clusters.

The following example creates a destination SVM named svm backup:

cluster dst:> vserver create -vserver svm backup -subtype dp-destination

2. From the destination cluster, create an SVM peer relationship using the vserver peer create command.

For more information, see Create an SVM intercluster peer relationship.

Learn more about vserver peer create in the ONTAP command reference.

3. Create a job schedule:

job schedule cron create -name job_name -month month -dayofweek day_of_week -day day of month -hour hour -minute minute

For -month, -dayofweek, and -hour, you can specify all to run the job every month, day of the week, and hour, respectively.



The minimum supported schedule (RPO) for FlexVol volumes in an SVM SnapMirror relationship is 15 minutes. The minimum supported schedule (RPO) for FlexGroup volumes in an SVM SnapMirror relationship is 30 minutes.

The following example creates a job schedule named my weekly that runs on Saturdays at 3:00 a.m.:

cluster_dst::> job schedule cron create -name my_weekly -dayofweek
"Saturday" -hour 3 -minute 0

4. Create a custom replication policy:

```
snapmirror policy create -vserver SVM -policy policy -type async-
mirror|vault|mirror-vault -comment comment -tries transfer_tries -transfer
-priority low|normal -is-network-compression-enabled true|false -discard
-configs network
```

The following example creates a custom replication policy for SnapMirror DR that excludes LIFs:

cluster_dst::> snapmirror policy create -vserver svm1 -policy
DR exclude LIFs -type async-mirror -discard-configs network

The following example creates a custom replication policy for unified replication that excludes LIFs:

```
cluster_dst::> snapmirror policy create -vserver svm1 -policy
unified_exclude_LIFs -type mirror-vault -discard-configs network
```



Consider creating the same custom SnapMirror policy on the source cluster for future failover and failback scenarios.

Learn more about snapmirror policy create in the ONTAP command reference.

5. From the destination SVM or the destination cluster, run the following command to create a replication relationship:

```
snapmirror create -source-path SVM: -destination-path SVM: -type DP|XDP
-schedule schedule -policy policy -identity-preserve true|false -discard
-configs true|false
```



You must enter a colon (:) after the SVM name in the -source-path and -destination -path options. See the examples below.

The following example creates a SnapMirror DR relationship that excludes LIFs:

```
cluster_dst::> snapmirror create -source-path svm1: -destination-path
svm_backup: -type XDP -schedule my_weekly -policy DR_exclude_LIFs
-identity-preserve true
```

The following example creates a SnapMirror unified replication relationship that excludes LIFs:

```
cluster_dst::> snapmirror create -source-path svm1: -destination-path
svm_backup: -type XDP -schedule my_weekly -policy unified_exclude_LIFs
-identity-preserve true -discard-configs true
```

Learn more about snapmirror create in the ONTAP command reference.

6. Stop the destination SVM:

vserver stop

SVM name

The following example stops the destination SVM named svm_backup:

cluster dst::> vserver stop -vserver svm backup

7. From the destination SVM or the destination cluster, initialize a replication relationship:

snapmirror initialize -source-path SVM: -destination-path SVM:

The following example initializes the relationship between the source, svml and the destination, svm_backup:

cluster_dst::> snapmirror initialize -source-path svml: -destination
-path svm_backup:

Learn more about snapmirror initialize in the ONTAP command reference.

After you finish

You must configure the network and protocols on the destination SVM for data access in the event a disaster occurs.

Related information

- snapmirror create
- snapmirror initialize
- snapmirror policy create

Exclude network, name service, and other settings from SVM replication with ONTAP

You might want to exclude network, name service, and other settings from an SVM replication relationship to avoid conflicts or configuration differences with the destination SVM.

You can use the -identity-preserve false option of the snapmirror create command to replicate only the volumes and security configurations of an SVM. Some protocol and name service settings are also preserved.

About this task

For a list of preserved protocol and name service settings, see Configurations replicated in SVM DR relationships.

Before you begin

The source and destination clusters and SVMs must be peered.

For more information, see Create a cluster peer relationship and Create an SVM intercluster peer relationship.

Steps

1. Create a destination SVM:

vserver create -vserver SVM -subtype dp-destination

The SVM name must be unique across the source and destination clusters.

The following example creates a destination SVM named svm backup:

cluster dst:> vserver create -vserver svm backup -subtype dp-destination

2. From the destination cluster, create an SVM peer relationship using the vserver peer create command.

For more information, see Create an SVM intercluster peer relationship.

Learn more about vserver peer create in the ONTAP command reference.

3. Create a replication job schedule:

job schedule cron create -name job_name -month month -dayofweek day_of_week
-day day of month -hour hour -minute minute

For -month, -dayofweek, and -hour, you can specify all to run the job every month, day of the week, and hour, respectively.



The minimum supported schedule (RPO) for FlexVol volumes in an SVM SnapMirror relationship is 15 minutes. The minimum supported schedule (RPO) for FlexGroup volumes in an SVM SnapMirror relationship is 30 minutes.

The following example creates a job schedule named my weekly that runs on Saturdays at 3:00 a.m.:

```
cluster_dst::> job schedule cron create -name my_weekly -dayofweek
"Saturday" -hour 3 -minute 0
```

4. Create a replication relationship that excludes network, name service, and other configuration settings:

```
snapmirror create -source-path SVM: -destination-path SVM: -type DP|XDP
-schedule schedule -policy policy -identity-preserve false
```



You must enter a colon (:) after the SVM name in the <code>-source-path</code> and <code>-destination</code> <code>-path</code> options. See the examples below. You must run this command from the destination SVM or the destination cluster.

The following example creates a SnapMirror DR relationship using the default MirrorAllSnapshots policy. The relationship excludes network, name service, and other configuration settings from SVM replication:

cluster_dst::> snapmirror create -source-path svm1: -destination-path svm_backup: -type XDP -schedule my_daily -policy MirrorAllSnapshots -identity-preserve false

The following example creates a unified replication relationship using the default MirrorAndVault policy. The relationship excludes network, name service, and other configuration settings:

cluster_dst:> snapmirror create svml: -destination-path svm_backup: -type XDP -schedule my_daily -policy MirrorAndVault -identity-preserve false

Assuming you have created a custom policy with the policy type <code>async-mirror</code>, the following example creates a SnapMirror DR relationship. The relationship excludes network, name service, and other configuration settings from SVM replication:

cluster_dst::> snapmirror create -source-path svm1: -destination-path svm_backup: -type XDP -schedule my_daily -policy my_mirrored -identity -preserve false

Assuming you have created a custom policy with the policy type mirror-vault, the following example creates a unified replication relationship. The relationship excludes network, name service, and other configuration settings from SVM replication:

cluster_dst::> snapmirror create -source-path svm1: -destination-path svm_backup: -type XDP -schedule my_daily -policy my_unified -identity -preserve false

Learn more about snapmirror create in the ONTAP command reference.

5. Stop the destination SVM:

vserver stop

SVM name

The following example stops a destination SVM named dvs1:

destination_cluster::> vserver stop -vserver dvs1

6. If you are using SMB, you must also configure an SMB server.

See SMB only: Creating an SMB server.

7. From the destination SVM or the destination cluster, initialize the SVM replication relationship:

snapmirror initialize -source-path SVM_name: -destination-path SVM_name:

Learn more about snapmirror initialize in the ONTAP command reference.

After you finish

You must configure the network and protocols on the destination SVM for data access in the event a disaster occurs.

Specify local tiers to use for ONTAP SnapMirror SVM DR relationships

After a disaster recovery SVM is created, you can use the aggr-list option with vserver modify command to limit which local tiers are used to host SVM DR destination volumes.

Steps

1. Create a destination SVM:

vserver create -vserver SVM -subtype dp-destination

2. Modify the disaster recovery SVM's aggr-list to limit the local tiers that are used to host the disaster recovery SVM's volume:

cluster dest::> vserver modify -vserver SVM -aggr-list <comma-separated-list>

Create an SMB server for an ONTAP SnapMirror destination SVM in a DR relationship

If the source SVM has an SMB configuration, and you chose to set identitypreserve to false, you must create an SMB server for the destination SVM. An SMB server is required for some SMB configurations, such as shares during initialization of the SnapMirror relationship.

Steps

1. Start the destination SVM by using the vserver start command.

```
destination_cluster::> vserver start -vserver dvs1
[Job 30] Job succeeded: DONE
```

Learn more about vserver start in the ONTAP command reference.

2. Verify that the destination SVM is in the running state and subtype is dp-destination by using the vserver show command.

Learn more about vserver show in the ONTAP command reference.

3. Create a LIF by using the network interface create command.
```
destination_cluster::>network interface create -vserver dvs1 -lif NAS1
-role data -data-protocol cifs -home-node destination_cluster-01 -home
-port a0a-101 -address 192.0.2.128 -netmask 255.255.255.128
```

Learn more about network interface create in the ONTAP command reference.

4. Create a route by using the network route create command.

```
destination_cluster::>network route create -vserver dvs1 -destination
0.0.0/0
-gateway 192.0.2.1
```

Network management

Learn more about network route create in the ONTAP command reference.

5. Configure DNS by using the vserver services dns create command.

```
destination_cluster::>vserver services dns create -domains
mydomain.example.com -vserver
dvs1 -name-servers 192.0.2.128 -state enabled
```

Learn more about vserver services dns create in the ONTAP command reference.

6. Add the preferred domain controller by using the vserver cifs domain preferred-dc add command.

```
destination_cluster::>vserver cifs domain preferred-dc add -vserver dvs1
-preferred-dc
192.0.2.128 -domain mydomain.example.com
```

Learn more about vserver cifs domain preferred-dc add in the ONTAP command reference.

7. Create the SMB server by using the vserver cifs create command.

```
destination_cluster::>vserver cifs create -vserver dvs1 -domain
mydomain.example.com
-cifs-server CIFS1
```

Learn more about vserver cifs create in the ONTAP command reference.

8. Stop the destination SVM by using the vserver stop command.

```
destination_cluster::> vserver stop -vserver dvs1
[Job 46] Job succeeded: DONE
```

Learn more about vserver stop in the ONTAP command reference.

Exclude volumes from an ONTAP SnapMirror SVM DR relationship

By default, all RW data volumes of the source SVM are replicated. If you do not want to protect all the volumes on the source SVM, you can use the -vserver-dr -protection unprotected option of the volume modify command to exclude volumes from SVM replication.

Steps

1. Exclude a volume from SVM replication:

```
volume modify -vserver SVM -volume volume -vserver-dr-protection unprotected
```

Learn more about volume modify in the ONTAP command reference.

The following example excludes the volume volA_src from SVM replication:

```
cluster_src::> volume modify -vserver SVM1 -volume volA_src -vserver-dr
-protection unprotected
```

If you later want to include a volume in the SVM replication that you originally excluded, run the following command:

volume modify -vserver SVM -volume volume -vserver-dr-protection protected

The following example includes the volume volA src in the SVM replication:

```
cluster_src::> volume modify -vserver SVM1 -volume volA_src -vserver-dr
-protection protected
```

2. Create and initialize the SVM replication relationship as described in Replicating an entire SVM configuration.

Serve data from a SnapMirror SVM DR destination

ONTAP SnapMirror SVM disaster recovery workflow

To recover from a disaster and serve data from the destination SVM, you must activate the destination SVM. Activating the destination SVM involves stopping scheduled SnapMirror transfers, aborting ongoing SnapMirror transfers, breaking the replication relationship, stopping the source SVM, and starting the destination SVM.



Configure the ONTAP SnapMirror SVM destination volume as writable

You need to make SVM destination volumes writeable before you can serve data to clients.

The procedure is largely identical to the procedure for volume replication, with one exception. If you set -identity-preserve true when you created the SVM replication relationship, you must stop the source SVM before activating the destination SVM.

About this task

Learn more about the commands described in this procedure in the ONTAP command reference.



In a disaster recovery scenario, you cannot perform a SnapMirror update from the source SVM to the disaster recovery destination SVM because your source SVM and its data will be inaccessible, and because updates since the last resync might be bad or corrupt.

Beginning with ONTAP 9.8, you can use System Manager to activate a destination storage VM after a disaster. Activating the destination storage VM makes the SVM destination volumes writable and enables you to serve data to clients.

Steps

You can perform this task from System Manager or the ONTAP CLI.

System Manager

- 1. If the source cluster is accessible, verify that the SVM is stopped: navigate to **Storage > Storage VMs** and check the **State** column for the SVM.
- 2. If the source SVM state is "Running", stop it: select i and choose **Stop**.
- 3. On the destination cluster, locate the desired protection relationship: navigate to **Protection > Relationships**.
- 4. Hover over the desired source storage VM name, click ; , and choose **Activate destination Storage VM**.
- 5. In the Activate destination storage VM window, select Activate the destination storage VM and break the relationship.
- 6. Click Activate.

CLI

1. From the destination SVM or the destination cluster, quiesce the SVM to stop scheduled transfers to the destination:

snapmirror quiesce -source-path <SVM>: -destination-path <SVM>:



You must enter a colon (:) after the SVM name in the -source-path and -destination-path options. See the example below.

The following example stops scheduled transfers between the source SVM ${\tt svm1}$ and the destination SVM ${\tt svm_backup}$:

```
cluster_dst::> snapmirror quiesce -source-path svm1: -destination
-path svm backup:
```

Learn more about snapmirror quiesce in the ONTAP command reference.

2. From the destination SVM or the destination cluster, stop ongoing transfers to the destination:

snapmirror abort -source-path <SVM>: -destination-path <SVM>:



You must enter a colon (:) after the SVM name in the -source-path and -destination-path options. See the example below.

The following example stops ongoing transfers between the source SVM ${\tt svm1}$ and the destination SVM ${\tt svm_backup}$:

cluster_dst::> snapmirror abort -source-path svm1: -destination-path
svm backup:

Learn more about snapmirror abort in the ONTAP command reference.

3. From the destination SVM or the destination cluster, break the replication relationship:

snapmirror break -source-path <SVM>: -destination-path <SVM>:



You must enter a colon (:) after the SVM name in the -source-path and -destination-path options. See the example below.

The following example breaks the relationship between the source SVM svm1 and the destination SVM svm backup:

```
cluster_dst::> snapmirror break -source-path svm1: -destination-path
svm backup:
```

Learn more about snapmirror break in the ONTAP command reference.

4. If you set -identity-preserve true when you created the SVM replication relationship, stop the source SVM:

vserver stop -vserver <SVM>

The following example stops the source SVM svm1:

cluster src::> vserver stop svm1

5. Start the destination SVM:

vserver start -vserver <SVM>

The following example starts the destination SVM svm backup:

cluster dst::> vserver start svm backup

After you finish

Configure SVM destination volumes for data access, as described in Configuring the destination volume for data access.

Reactivate the SnapMirror source SVM

ONTAP SnapMirror source SVM reactivation workflow

If the source SVM exists after a disaster, you can reactivate it and protect it by recreating the SVM disaster recovery relationship.



Reactivate the original ONTAP SnapMirror source SVM

You can reestablish the original data protection relationship between the source and destination SVM when you no longer need to serve data from the destination. The procedure is largely identical to the procedure for volume replication, with one exception. You must stop the destination SVM before reactivating the source SVM.

Before you begin

If you have increased the size of destination volume while serving data from it, before you reactivate the source volume, you should manually increase max-autosize on the original source volume to ensure it can grow sufficiently.

When a destination volume grows automatically

About this task

Beginning with ONTAP 9.11.1, you can reduce resynchronization time during a disaster recovery rehearsal by using the CLI -quick-resync true option of the snapmirror resync command while performing a reverse resync of an SVM DR relationship. A quick resync can reduce the time it takes to return to production

by bypassing the data warehouse rebuild and restore operations. Learn more about snapmirror resync in the ONTAP command reference.



Quick resync does not preserve the storage efficiency of the destination volumes. Enabling quick resync might increase the volume space used by the destination volumes.

This procedure assumes that the baseline in the original source volume is intact. If the baseline is not intact, you must create and initialize the relationship between the volume you are serving data from and the original source volume before performing the procedure.

Beginning with ONTAP 9.8, you can use System Manager to reactivate a source storage VM after a disaster. Reactivating the source storage VM stops the destination storage VM, and it reenables replication from the source to the destination.

When you use System Manager to reactivate the source storage VM, System Manager performs the following operations in the background:

- Creates a reverse SVM DR relationship from the original destination to original source using SnapMirror resync
- · Stops the destination SVM
- · Updates the SnapMirror relationship
- Breaks the SnapMirror relationship
- · Restarts the original SVM
- Issues a SnapMirror resync of the original source back to the original destination
- Cleans up the SnapMirror relationships

Steps

You can perform this task from System Manager or the ONTAP CLI.

System Manager

- 1. From the destination cluster, click **Protection > Relationships**, and locate the desired protection relationship.
- 2. Hover over the source relationship name, click ; and select **Reactivate Source Storage VM**.
- 3. In the Reactivate source storage VM window, click Reactivate.
- 4. Under **Relationships**, monitor the source reactivation progress by viewing **Transfer Status** for the protection relationship. When reactivation is complete, the relationship state should return to "Mirrored".

CLI

1. From the original source SVM or the original source cluster, create a reverse SVM DR relationship using the same configuration, policy, and identity-preserve setting as the original SVM DR relationship:

```
snapmirror create -source-path <SVM>: -destination-path <SVM>:
```



You must enter a colon (:) after the SVM name in the -source-path and -destination-path options. See the example below.

The following example creates a relationship between the SVM from which you are serving data, svm_backup, and the original source SVM, svm1:

```
cluster_src::> snapmirror create -source-path svm_backup:
  -destination-path svm1:
```

Learn more about snapmirror create in the ONTAP command reference.

2. From the original source SVM or the original source cluster, run the following command to reverse the data protection relationship:

snapmirror resync -source-path <SVM>: -destination-path <SVM>:



You must enter a colon (:) after the SVM name in the -source-path and -destination-path options. See the example below.

Although resync does not require a baseline transfer, it can be time-consuming. You might want to run the resync in off-peak hours.



The command fails if a common snapshot does not exist on the source and destination. Use snapmirror initialize to reinitialize the relationship.

The following example reverses the relationship between the original source SVM, svm1, and the SVM from which you are serving data, svm backup:

```
cluster_src::> snapmirror resync -source-path svm_backup:
-destination-path svm1:
```

Example using -quick-resync option:

```
cluster_src::> snapmirror resync -source-path svm_backup:
-destination-path svm1: -quick-resync true
```

3. When you are ready to reestablish data access to the original source SVM, stop the original destination SVM to disconnect any clients currently connected to the original destination SVM.

```
vserver stop -vserver <SVM>
```

The following example stops the original destination SVM which is currently serving data:

cluster_dst::> vserver stop svm_backup

4. Verify that the original destination SVM is in the stopped state by using the vserver show command.

5. From the original source SVM or the original source cluster, run the following command to perform the final update of the reversed relationship to transfer all changes from the original destination SVM to the original source SVM:

snapmirror update -source-path <SVM>: -destination-path <SVM>:



You must enter a colon (:) after the SVM name in the -source-path and -destination-path options. See the example below.

The following example updates the relationship between the original destination SVM from which you are serving data, svm backup, and the original source SVM, svm1:

```
cluster_src::> snapmirror update -source-path svm_backup:
-destination-path svm1:
```

Learn more about snapmirror update in the ONTAP command reference.

6. From the original source SVM or the original source cluster, run the following command to stop scheduled transfers for the reversed relationship:

snapmirror quiesce -source-path <SVM>: -destination-path <SVM>:



You must enter a colon (:) after the SVM name in the -source-path and -destination-path options. See the example below.

The following example stops scheduled transfers between the SVM you are serving data from, svm backup, and the original SVM, svm1:

```
cluster_src::> snapmirror quiesce -source-path svm_backup:
-destination-path svm1:
```

7. When the final update is complete and the relationship indicates "Quiesced" for the relationship status, run the following command from the original source SVM or the original source cluster to break the reversed relationship:

snapmirror break -source-path <SVM>: -destination-path <SVM>:



You must enter a colon (:) after the SVM name in the -source-path and -destination-path options. See the example below.

The following example breaks the relationship between the original destination SVM from which you were serving data, svm backup, and the original source SVM, svm1:

```
cluster_src::> snapmirror break -source-path svm_backup:
  -destination-path svm1:
```

Learn more about snapmirror break in the ONTAP command reference.

8. If the original source SVM was previously stopped, from the original source cluster, start the original source SVM:

vserver start -vserver <SVM>

The following example starts the original source SVM:

```
cluster src::> vserver start svm1
```

9. From the original destination SVM or the original destination cluster, reestablish the original data protection relationship:

snapmirror resync -source-path <SVM>: -destination-path <SVM>:



You must enter a colon (:) after the SVM name in the -source-path and -destination-path options. See the example below.

The following example reestablishes the relationship between the original source SVM, svm1, and the original destination SVM, svm backup:

```
cluster_dst::> snapmirror resync -source-path svm1: -destination
-path svm backup:
```

10. From the original source SVM or the original source cluster, run the following command to delete the reversed data protection relationship:

snapmirror delete -source-path <SVM>: -destination-path <SVM>:



You must enter a colon (:) after the SVM name in the -source-path and -destination-path options. See the example below.

The following example deletes the reversed relationship between the original destination SVM, svm_backup, and the original source SVM, svm1:

```
cluster_src::> snapmirror delete -source-path svm_backup:
-destination-path svm1:
```

11. From the original destination SVM or the original destination cluster, release the reversed data protection relationship:

snapmirror release -source-path <SVM>: -destination-path <SVM>:



You must enter a colon (:) after the SVM name in the -source-path and -destination-path options. See the example below.

The following example releases the reversed relationship between the original destination SVM,

svm_backup, and the original source SVM, svm1

cluster_dst::> snapmirror release -source-path svm_backup: -destination-path svm1:

After you finish

Use the snapmirror show command to verify that the SnapMirror relationship was created.

Learn more about snapmirror show in the ONTAP command reference.

Related information

- snapmirror create
- snapmirror delete
- snapmirror initialize
- snapmirror quiesce
- snapmirror release
- snapmirror resync

Reactivate the original ONTAP SnapMirror source SVM for FlexGroup volumes

You can reestablish the original data protection relationship between the source and destination SVM when you no longer need to serve data from the destination. To reactivate the original source SVM when you are using FlexGroup volumes, you need to perform some additional steps, including deleting the original SVM DR relationship and releasing the original relationship before you reverse the relationship. You also need to release the reversed relationship and recreate the original relationship before stopping scheduled transfers.

Steps

1. From the original destination SVM or the original destination cluster, delete the original SVM DR relationship:

snapmirror delete -source-path SVM: -destination-path SVM:



You must enter a colon (:) after the SVM name in the -source-path and -destination -path options. See the example below.

The following example deletes the original relationship between the original source SVM, svm1, and the original destination SVM, svm_backup:

cluster_dst::> snapmirror delete -source-path svm1: -destination-path
svm backup:

2. From the original source SVM or the original source cluster, release the original relationship while keeping the snapshots intact:

```
snapmirror release -source-path SVM: -destination-path SVM: -relationship-info
-only true
```



You must enter a colon (:) after the SVM name in the -source-path and -destination -path options. See the example below.

The following example releases the original relationship between the original source SVM, svm1, and the original destination SVM, svm backup.

```
cluster_src::> snapmirror release -source-path svm1: -destination-path
svm_backup: -relationship-info-only true
```

3. From the original source SVM or the original source cluster, create a reverse SVM DR relationship using the same configuration, policy, and identity-preserve setting as the original SVM DR relationship:

snapmirror create -source-path SVM: -destination-path SVM:



You must enter a colon (:) after the SVM name in the -source-path and -destination -path options. See the example below.

The following example creates a relationship between the SVM from which you are serving data, svm backup, and the original source SVM, svm1:

```
cluster_src::> snapmirror create -source-path svm_backup: -destination
-path svm1:
```

4. From the original source SVM or the original source cluster, run the following command to reverse the data protection relationship:

snapmirror resync -source-path SVM: -destination-path SVM:



You must enter a colon (:) after the SVM name in the -source-path and -destination -path options. See the example below.

Although resync does not require a baseline transfer, it can be time-consuming. You might want to run the resync in off-peak hours.



The command fails if a common snapshot does not exist on the source and destination. Use snapmirror initialize to reinitialize the relationship.

The following example reverses the relationship between the original source SVM, svm1, and the SVM from which you are serving data, svm_backup:

```
cluster_src::> snapmirror resync -source-path svm_backup: -destination
-path svm1:
```

5. When you are ready to reestablish data access to the original source SVM, stop the original destination SVM to disconnect any clients currently connected to the original destination SVM.

vserver stop -vserver SVM

The following example stops the original destination SVM which is currently serving data:

cluster dst::> vserver stop svm backup

6. Verify that the original destination SVM is in the stopped state by using the vserver show command.

cluster_dst::>	vserver	show			
			Admin	Operational	Root
Vserver	Туре	Subtype	State	State	Volume
Aggregate					
svm backup	data	default	stopped	stopped	rv
aggr1			o copposi		

7. From the original source SVM or the original source cluster, run the following command to perform the final update of the reversed relationship to transfer all changes from the original destination SVM to the original source SVM:

snapmirror update -source-path SVM: -destination-path SVM:



You must enter a colon (:) after the SVM name in the -source-path and -destination -path options. See the example below.

The following example updates the relationship between the original destination SVM from which you are serving data,svm backup, and the original source SVM, svm1:

```
cluster_src::> snapmirror update -source-path svm_backup: -destination
-path svm1:
```

Learn more about snapmirror update in the ONTAP command reference.

8. From the original source SVM or the original source cluster, run the following command to stop scheduled transfers for the reversed relationship:

snapmirror quiesce -source-path SVM: -destination-path SVM:



You must enter a colon (:) after the SVM name in the -source-path and -destination -path options. See the example below.

The following example stops scheduled transfers between the SVM you are serving data from, svm backup, and the original SVM, svm1:

cluster_src::> snapmirror quiesce -source-path svm_backup: -destination
-path svm1:

Learn more about snapmirror quiesce in the ONTAP command reference.

9. When the final update is complete and the relationship indicates "Quiesced" for the relationship status, run the following command from the original source SVM or the original source cluster to break the reversed relationship:

snapmirror break -source-path SVM: -destination-path SVM:



You must enter a colon (:) after the SVM name in the -source-path and -destination -path options. See the example below.

The following example breaks the relationship between the original destination SVM from which you were serving data, svm backup, and the original source SVM, svm1:

```
cluster_src::> snapmirror break -source-path svm_backup: -destination
-path svm1:
```

Learn more about snapmirror break in the ONTAP command reference.

10. If the original source SVM was previously stopped, from the original source cluster, start the original source SVM:

vserver start -vserver SVM

The following example starts the original source SVM:

cluster_src::> vserver start svm1

11. From the original source SVM or the original source cluster, delete the reversed SVM DR relationship:

snapmirror delete -source-path SVM: -destination-path SVM:



You must enter a colon (:) after the SVM name in the -source-path and -destination -path options. See the example below.

The following example deletes the reversed relationship between the original destination SVM, svm_backup, and the original source SVM, svm1:

```
cluster_src::> snapmirror delete -source-path svm_backup: -destination
-path svm1:
```

12. From the original destination SVM or the original destination cluster, release the reversed relationship while keeping the snapshots intact:

```
snapmirror release -source-path SVM: -destination-path SVM: -relationship-info
-only true
```



You must enter a colon (:) after the SVM name in the -source-path and -destination -path options. See the example below.

The following example releases the reversed relationship between the original destination SVM, svm_backup, and the original source SVM, svm1:

```
cluster_dst::> snapmirror release -source-path svm_backup: -destination
-path svm1: -relationship-info-only true
```

13. From the original destination SVM or the original destination cluster, recreate the original relationship. Use the same configuration, policy, and identity-preserve setting as the original SVM DR relationship:

snapmirror create -source-path SVM: -destination-path SVM:



You must enter a colon (:) after the SVM name in the -source-path and -destination -path options. See the example below.

The following example creates a relationship between the original source SVM, svm1, and the original destination SVM, svm backup:

```
cluster_dst::> snapmirror create -source-path svm1: -destination-path
svm backup:
```

14. From the original destination SVM or the original destination cluster, reestablish the original data protection relationship:

```
snapmirror resync -source-path SVM: -destination-path SVM:
```



You must enter a colon (:) after the SVM name in the -source-path and -destination -path options. See the example below.

The following example reestablishes the relationship between the original source SVM, svm1, and the original destination SVM, svm_backup:

```
cluster_dst::> snapmirror resync -source-path svm1: -destination-path
svm_backup:
```

Related information

- snapmirror create
- snapmirror delete
- snapmirror initialize
- snapmirror quiesce
- snapmirror release
- snapmirror resync

Resynchronize the data on an ONTAP SnapMirror destination SVM

ONTAP 9.11.1 introduces an option to bypass a full data warehouse rebuild when you perform a disaster recovery rehearsal, enabling you to return to production faster.

Beginning with ONTAP 9.8, you can use System Manager to resynchronize the data and configuration details from the source storage VM to the destination storage VM in a broken protection relationship and reestablish the relationship.

You perform the resync operation only from the destination of the original relationship. The resync deletes any data in the destination storage VM that is newer than the data in the source storage VM.

Steps

You can use System Manager or the ONTAP CLI to perform this task.

System Manager

- 1. From the destination, select the desired protection relationship: click **Protection > Relationships**.
- 2. Optionally, select **Perform a quick resync** to bypass a full data warehouse rebuild during a disaster recovery rehearsal.
- 3. Click and click **Resync**.
- 4. Under **Relationships**, monitor the resynchronization progress by viewing **Transfer Status** for the relationship.

CLI

1. From the destination cluster, resynchronize the relationship:

```
snapmirror resync -source-path <svm>: -destination-path <svm>:
-quick-resync true|false
```

Related information

snapmirror resync

Convert an ONTAP SnapMirror volume DR relationship to an SVM DR relationship

You can convert replication relationships between volumes to a replication relationship between the storage virtual machines (SVMs) that own the volumes, provided that each volume on the source (except the root volume) is being replicated, and each volume on the source (including the root volume) has the same name as the volume on the destination.

About this task

Use the volume rename command when the SnapMirror relationship is idle to rename destination volumes if necessary. Learn more about volume rename in the ONTAP command reference.

Steps

1. From the destination SVM or the destination cluster, run the following command to resync the source and destination volumes:

```
snapmirror resync -source-path <SVM:volume> -destination-path <SVM:volume>
-type DP|XDP -policy <policy>
```

Learn more about snapmirror resync in the ONTAP command reference.



Although resync does not require a baseline transfer, it can be time-consuming. You might want to run the resync in off-peak hours.

The following example resyncs the relationship between the source volume volA on svm1 and the destination volume volA on svm backup:

```
cluster_dst::> snapmirror resync -source-path svm1:volA -destination
-path svm backup:volA
```

2. Create an SVM replication relationship between the source and destination SVMs, as described in Replicating SVM configurations.

You must use the -identity-preserve true option of the snapmirror create command when you create your replication relationship.

Learn more about snapmirror create in the ONTAP command reference.

3. Stop the destination SVM:

vserver stop -vserver SVM

Learn more about vserver stop in the ONTAP command reference.

The following example stops the destination SVM svm backup:

cluster_dst::> vserver stop svm_backup

4. From the destination SVM or the destination cluster, run the following command to resync the source and destination SVMs:

```
snapmirror resync -source-path <SVM>: -destination-path <SVM>: -type DP|XDP
-policy <policy>
```



You must enter a colon (:) after the SVM name in the -source-path and -destination -path options. See the example below.

Although resync does not require a baseline transfer, it can be time-consuming. You might want to run the resync in off-peak hours.

The following example resyncs the relationship between the source SVM svm1 and the destination SVM svm_backup:

```
cluster_dst::> snapmirror resync -source-path svm1: -destination-path
svm_backup:
```

Related information

- snapmirror create
- snapmirror resync

Delete an ONTAP SnapMirror SVM replication relationship

You can use the snapmirror delete and snapmirror release commands to delete an SVM replication relationship. You can then delete unneeded destination volumes manually. Learn more about the commands described in this procedure in the ONTAP command reference.

About this task

The snapmirror release command deletes any SnapMirror-created snapshots from the source. You can use the -relationship-info-only option to preserve the snapshots.

Steps

1. Run the following command from the destination SVM or the destination cluster to break the replication relationship:

snapmirror break -source-path SVM: -destination-path SVM:



You must enter a colon (:) after the SVM name in the -source-path and -destination -path options. See the example below.

The following example breaks the relationship between the source SVM svm1 and the destination SVM svm_backup:

```
cluster_dst::> snapmirror break -source-path svm1: -destination-path
svm_backup:
```

Learn more about snapmirror break in the ONTAP command reference.

2. Run the following command from the destination SVM or the destination cluster to delete the replication relationship:

snapmirror delete -source-path SVM: -destination-path SVM:



You must enter a colon (:) after the SVM name in the -source-path and -destination -path options. See the example below.

The following example deletes the relationship between the source SVM svm1 and the destination SVM svm_backup:

```
cluster_dst::> snapmirror delete -source-path svm1: -destination-path
svm backup:
```

Learn more about snapmirror delete in the ONTAP command reference.

3. Run the following command from the source cluster or source SVM to release the replication relationship information from the source SVM:

snapmirror release -source-path SVM: -destination-path SVM:



You must enter a colon (:) after the SVM name in the -source-path and -destination -path options. See the example below.

The following example releases information for the specified replication relationship from the source SVM svm1:

cluster_src::> snapmirror release -source-path svm1: -destination-path
svm backup:

Learn more about snapmirror release in the ONTAP command reference.

Manage SnapMirror root volume replication

Learn about ONTAP SnapMirror root volume replication

Every SVM in a NAS environment has a unique namespace. The SVM *root volume,* containing operating system and related information, is the entry point to the namespace hierarchy. To ensure that data remains accessible to clients in the event of a node outage

or failover, you should create a load-sharing mirror copy of the SVM root volume.

The main purpose of load-sharing mirrors for SVM root volumes is no longer for load sharing; instead, their purpose is for disaster recovery.

- If the root volume is temporarily unavailable, the load-sharing mirror automatically provides read-only access to root volume data.
- If the root volume is permanently unavailable, you can promote one of the load-sharing volumes to provide write access to root volume data.

Create and initialize ONTAP load-sharing mirror relationships

You should create a load-sharing mirror (LSM) for each SVM root volume that serves NAS data in the cluster. For clusters consisting of two or more HA pairs, you should consider load-sharing mirrors of SVM root volumes to ensure the namespace remains accessible to clients in the event that both nodes of an HA pair fail. Load-sharing mirrors are not suitable for clusters consisting of a single HA pair.

Before you begin

Beginning with ONTAP 9.16.1, when you create a load-sharing mirror relationship, the destination SVM cannot have a storage limit enabled.

About this task

If you create an LSM on the same node, and the node is unavailable, you have a single point of failure, and you do not have a second copy to ensure the data remains accessible to clients. But when you create the LSM on a node other than the one containing the root volume, or on a different HA pair, your data is still accessible in the event of an outage.

For example, in a four-node cluster with a root volume on three nodes:

- For the root volume on HA1 node 1, create the LSM on HA2 node 1 or HA2 node 2.
- For the root volume on HA 1 node 2, create the LSM on HA 2 node 1 or HA 2 node 2.
- For the root volume on HA 2 node 1, create the LSM on HA 1 node 1 or HA 1 node 2.

Steps

1. Create a destination volume for the LSM:

```
volume create -vserver <SVM> -volume <volume> -aggregate <aggregate>
-type DP -size <size>
```

The destination volume should be the same or greater in size than the root volume.

It is a best practice to name the root and destination volume with suffixes, such as root and m1.

Learn more about volume create in the ONTAP command reference.

The following example creates a load-sharing mirror volume for the root volume svm1_root in cluster_src:

```
cluster_src:> volume create -vserver svm1 -volume svm1_m1 -aggregate
aggr_1 -size 1gb -state online -type DP
```

- 2. Create a replications job schedule.
- Create a load-sharing mirror relationship between the SVM root volume and the destination volume for the LSM:

```
snapmirror create -source-path <SVM:volume> -destination-path
<SVM:volume> -type LS -schedule <schedule>
```

The following example creates a load-sharing mirror relationship between the root volume svm1_root and the load-sharing mirror volume svm1_m1:

```
cluster_src::> snapmirror create -source-path svml:svml_root
-destination-path svml:svml_m1 -type LS -schedule hourly
```

The type attribute of the load-sharing mirror changes from DP to LS.

Learn more about snapmirror create in the ONTAP command reference.

4. Initialize the load-sharing mirror:

snapmirror initialize-ls-set -source-path <SVM:volume>

The following example initializes the load-sharing mirror for the root volume svm1 root:

cluster src::> snapmirror initialize-ls-set -source-path svm1:svm1 root

Learn more about snapmirror initialize in the ONTAP command reference.

Update an ONTAP load-sharing mirror relationship

Load-sharing mirror (LSM) relationships are updated automatically for SVM root volumes after a volume in the SVM is mounted or unmounted, and during volume create operations that include the junction-path option. You can manually update a LSM relationship if you want it updated before the next scheduled update.

Load-sharing mirror relationships update automatically in the following circumstances:

- · It's time for a scheduled update
- · A mount or unmount operation is performed on a volume in the SVM root volume

• A volume create command is issued that includes the <code>junction-path</code> option

Learn more about volume create in the ONTAP command reference.

Step

1. Update a load-sharing mirror relationship manually:

You must replace the variables in angle brackets with the required values before running this command.

snapmirror update-ls-set -source-path <SVM:volume>

The following example updates the load-sharing mirror relationship for the root volume svm1 root:

cluster src::> snapmirror update-ls-set -source-path svm1:svm1 root

Learn more about snapmirror update in the ONTAP command reference.

Promote an ONTAP load-sharing mirror

If a root volume is permanently unavailable, you can promote the load-sharing mirror (LSM) volume to provide write access to root volume data.

Before you begin

You must use advanced privilege level commands for this task.

Steps

1. Change to advanced privilege level:

set -privilege advanced

2. Promote an LSM volume:

You must replace the variables in angle brackets with the required values before running this command.

```
snapmirror promote -destination-path <SVM:volume>
```

The following example promotes the volume svm1 m2 as the new SVM root volume:

Enter y. ONTAP makes the LSM volume a read/write volume, and deletes the original root volume if it is accessible.



The promoted root volume might not have all of the data that was in the original root volume if the last update did not occur recently.

Learn more about snapmirror promote in the ONTAP command reference.

3. Return to admin privilege level:

set -privilege admin

4. Rename the promoted volume following the naming convention you used for the root volume:

You must replace the variables in angle brackets with the required values before running this command.

```
volume rename -vserver <SVM> -volume <volume> -newname <new_name>
```

The following example renames the promoted volume svm1 m2 with the name svm1 root:

```
cluster_src::> volume rename -vserver svm11 -volume svm1_m2 -newname
svm1_root
```

5. Protect the renamed root volume, as described in step 3 through step 4 in Creating and initializing loadsharing mirror relationships.

Back up to the cloud

Install an ONTAP SnapMirror cloud license

SnapMirror cloud relationships can be orchestrated using pre-qualified third-party backup applications. Beginning with ONTAP 9.9.1, you can also use System Manager to orchestrate SnapMirror cloud replication. Both SnapMirror and SnapMirror cloud capacity licenses are required when using System Manager to orchestrate on-premises ONTAP to

object storage backups. You will also need to request and install the SnapMirror cloud API license.

About this task

The SnapMirror cloud and SnapMirror S3 licenses are cluster licenses, not node licenses, so they are *not* delivered with the ONTAP One license bundle. These licenses are included in the separate ONTAP One Compatibility bundle. If you want to enable SnapMirror cloud, you need to request this bundle.

Additionally, System Manager orchestration of SnapMirror cloud backups to object storage requires a SnapMirror cloud API key. This API license is a single-instance cluster-wide license, meaning it does not need to be installed on every node in the cluster.

Steps

You need to request and download the ONTAP One Compatibility bundle and the SnapMirror cloud API license and then install them using System Manager.

1. Locate and record the cluster UUID for the cluster you want to license.

The cluster UUID is required when you submit your request to order the ONTAP One Compatibility bundle for your cluster.

- 2. Contact your NetApp sales team and request the ONTAP One Compatibility bundle.
- Request the SnapMirror cloud API license by following the instructions provided on the NetApp Support Site.

Request SnapMirror cloud API license key

- 4. When you've received and downloaded the license files, use System Manager to upload the ONTAP Cloud Compatibility NLF and the SnapMirror cloud API NLF to the cluster:
 - a. Click Cluster > Settings.
 - b. In the Settings window, click Licenses.
 - c. In the Licenses window, click + Add .
 - d. In the **Add License** dialog box, click **Browse** to select the NLF you downloaded, and then click **Add** to upload the file to the cluster.

Related information

Back up data to the cloud using SnapMirror

NetApp Software License Search

Back up data to the cloud using ONTAP SnapMirror

Beginning with ONTAP 9.9.1, you can back up your data to the cloud and to restore your data from cloud storage to a different volume by using System Manager. You can use either StorageGRID or ONTAP S3 as your cloud object store.

Beginning with ONTAP 9.16.1:

• SnapMirror cloud backup supports fan-out relationships. This means that SnapMirror backups can be created simultaneously on two different object stores. With ONTAP 9.16.1, SnapMirror cloud supports two fan-out relationships. Fan-outs can be to two object stores and to one or two buckets in two different object

stores. Attempts to create more than two fan-out relationships will fail.

• SnapMirror cloud supports backups of volumes migrated to the cloud using a more efficient synchronization process using existing ONTAP REST APIs. The functionality supports SnapMirror cloud backups from a migrated volume in the cloud to the same destination object store endpoint without the need for performing a re-baseline operation. Both FlexVol and FlexGroup volumes are supported.

Before using the SnapMirror cloud feature, you should request a SnapMirror cloud API license key from the NetApp Support Site: Request SnapMirror cloud API license key.

Following the instructions, you should provide a simple description of your business opportunity and request the API key by sending an email to the provided email address. You should receive an email response within 24 hours with further instructions on how to acquire the API key.

Add a cloud object store

Before you configure SnapMirror cloud backups, you need to add a StorageGRID or ONTAP S3 cloud object store.

Steps

- 1. Click Protection > Overview > Cloud Object Stores.
- 2. Click + Add .

Back up using the default policy

You can quickly configure a SnapMirror cloud backup for an existing volume using the default cloud protection policy, DailyBackup.

Steps

- 1. Click Protection > Overview and select Back Up Volumes to Cloud.
- 2. If this is your first time backing up to the cloud, enter your SnapMirror cloud API license key in the license field as indicated.
- 3. Click Authenticate and Continue.
- 4. Select a source volume.
- 5. Select a cloud object store.
- 6. Click Save.

Create a custom cloud backup policy

If you do not want to use the default DailyBackup cloud policy for your SnapMirror cloud backups, you can create your own policy.

Steps

- 1. Click Protection > Overview > Local Policy Settings and select Protection Policies.
- 2. Click Add and enter the new policy details.
- 3. In the **Policy Type** section, select **Back up to Cloud** to indicate that you are creating a cloud policy.
- 4. Click Save.

Create a backup from the Volumes page

You can use the System Manager Volumes page when you want to select and create cloud backups for

multiple volumes at one time or when you want to use a custom protection policy.

Steps

- 1. Click **Storage > Volumes**.
- 2. Select the volumes you want to back up to the cloud, and click Protect.
- 3. In the Protect Volume window, click More Options.
- 4. Select a policy.

You can select the default policy, DailyBackup, or a custom cloud policy you created.

- 5. Select a cloud object store.
- 6. Click Save.

Restore from the cloud

You can use System Manager to restore backed up data from cloud storage to a different volume on the source cluster.



If you are using ONTAP 9.16.1 or later and you are performing a SnapMirror cloud single file restore to a FlexGroup volume, you should only restore files to a new directory in the FlexGroup volume, and granular data must be set to advanced on the destination FlexGroup volume. For more information about setting the -granular-data advanced option, see Balance ONTAP FlexGroup volumes by redistributing file data.

Steps

- 1. From the source Cluster of a SnapMirror-to-Cloud relationship, click **Storage > Volumes**.
- 2. Select the volume you want to restore.
- 3. Select the Back Up to Cloud tab.
- 4. Click i next to the source volume you want to restore to display the menu, and select **Restore**.
- 5. Under **Source**, select a storage VM and then enter the name of the volume to which you want the data restored.
- 6. Under **Destination**, select the snapshot you want to restore.
- 7. Click Save.

Delete a SnapMirror cloud relationship

You can use System Manager to delete a cloud relationship.

Steps

- 1. Click **Storage > Volumes** and select the volume you want to delete.
- 2. Click inext to the source volume and select Delete.
- 3. Select **Delete the cloud object store endpoint (optional)** if you want to delete the cloud object store endpoint.
- 4. Click Delete.

Remove a cloud object store

You can use System Manager to remove a cloud object store if it is not part of a cloud backup relationship. When a cloud object store is part of a cloud backup relationship, it cannot be deleted.

Steps

- 1. Click Protection > Overview > Cloud Object Stores.
- 2. Select the object store you want to delete, click and select **Delete**.

Back up data using BlueXP backup and recovery

Beginning with ONTAP 9.9.1, you can use System Manager to back up data in the cloud using BlueXP backup and recovery (formerly Cloud Backup service).

BlueXP backup and recovery supports FlexVol read-write volumes and data-protection (DP) volumes. Beginning with ONTAP 9.12.1, BlueXP backup and recovery supports FlexGroup volumes and SnapLock volumes.

Learn more about BlueXP backup and recovery.

Before you begin

You should perform the following procedures to establish an account in BlueXP. For the service account, you need to create the role as "Account Admin". (Other service account roles do not have the required privileges needed to establish a connection from System Manager.)

- 1. Create an account in BlueXP.
- 2. Create a connector in BlueXP with one of the following cloud providers:
 - Microsoft Azure
 - Amazon Web Services (AWS)
 - Google Cloud Platform (GCP)
 - StorageGRID (ONTAP 9.10.1)



Beginning with ONTAP 9.10.1, you can select StorageGRID as a cloud backup provider, but only if BlueXP is deployed on premises. The BlueXP Connector must be installed on premises and available through the BlueXP software-as-a-service (SaaS) application.

- 3. Subscribe to BlueXP backup and recovery in BlueXP (requires the appropriate license).
- 4. Generate an access key and a secret key using BlueXP.

Register the cluster with BlueXP

You can register the cluster with BlueXP by using either BlueXP or System Manager.

Steps

- 1. In System Manager, go to Protection Overview.
- 2. Under BlueXP backup and recovery, provide the following details:
 - Client ID
 - Client secret key

3. Select Register and Continue.

Enable BlueXP backup and recovery

After the cluster is registered with BlueXP, you need to enable the BlueXP backup and recovery and initiate the first backup to the cloud.

Steps

- 1. In System Manager, select Protection > Overview, then scroll to the Cloud Backup Service section.
- 2. Enter the **Client ID** and **Client Secret**.



Beginning with ONTAP 9.10.1, you can learn about the cost of using the cloud by selecting **Learn more about the cost of using the cloud**.

- 3. Select Connect and Enable Cloud Backup Service.
- 4. On the **Enable BlueXP backup and recovery** page, provide the following details, depending on the provider you selected.

For this cloud provider	Enter the following data
Azure	 Azure Subscription ID Region Resource group name (existing or new)
AWS	 AWS Account ID Access key Secret key Region
Google Cloud Project (GCP)	 Google Cloud Project name Google Cloud Access key Google Cloud Secret key Region
StorageGRID (ONTAP 9.10.1 and later, and only for on-premises deployment of BlueXP)	 Server SG Access Key SG Secret Key

5. Select a **Protection policy**:

- Existing policy: Choose an existing policy.
- **New Policy**: Specify a name and set up a transfer schedule.



Beginning with ONTAP 9.10.1, you can specify whether you want to enable archiving with Azure or AWS.



If you enable archiving for a volume with Azure or AWS, you cannot disable the archiving.

If you enable archiving for Azure or AWS, specify the following:

- The number of days after which the volume is archived.
- The number of backups to retain in the archive. Specify "0" (zero) to archive up to the latest backup.
- For AWS, select the archive storage class.
- 6. Select the volumes you want to back up.
- 7. Select Save.

Edit the protection policy used for BlueXP backup and recovery

You can change which protection policy is used with BlueXP backup and recovery.

Steps

- 1. In System Manager, select Protection > Overview, then scroll to the Cloud Backup Service section.
- 2. Select ; then Edit.
- 3. Select a Protection policy:
 - Existing policy: Choose an existing policy.
 - New Policy: Specify a name and set up a transfer schedule.



Beginning with ONTAP 9.10.1, you can specify whether you want to enable archiving with Azure or AWS.



If you enable archiving for a volume with Azure or AWS, you cannot disable the archiving.

If you enable archiving for Azure or AWS, specify the following:

- The number of days after which the volume is archived.
- The number of backups to retain in the archive. Specify "0" (zero) to archive up to the latest backup.
- For AWS, select the archive storage class.
- 4. Select Save.

Protect new volumes or LUNs on the cloud

When you create a new volume or LUN, you can establish a SnapMirror protection relationship that enables backing up to the cloud for the volume or LUN.

Before you begin

- You should have a SnapMirror license.
- Intercluster LIFs should be configured.
- NTP should be configured.

• Cluster must be running ONTAP 9.9.1 or later.

About this task

You cannot protect new volumes or LUNs on the cloud for the following cluster configurations:

- The cluster cannot be in a MetroCluster environment.
- SVM-DR is not supported.
- FlexGroup volumes cannot be backed up using BlueXP backup and recovery.

Steps

- 1. When provisioning a volume or LUN, on the **Protection** page in System Manager, select the checkbox labeled **Enable SnapMirror (Local or Remote)**.
- 2. Select the BlueXP backup and recovery policy type.
- 3. If the BlueXP backup and recovery is not enabled, select **Enable Backup using BlueXP backup and recovery**.

Protect existing volumes or LUNs on the cloud

You can establish a SnapMirror protection relationship for existing volumes and LUNs.

Steps

- 1. Select an existing volume or LUN, and select **Protect**.
- 2. On the **Protect Volumes** page, specify **Backup using BlueXP backup and recovery** for the protection policy.
- 3. Select Protect.
- 4. On the **Protection** page, select the checkbox labeled **Enable SnapMirror (Local or Remote)**.
- 5. Select Connect and enable BlueXP backup and recovery.

Restore data from backup files

You can perform backup management operations, such as restoring data, updating relationships, and deleting relationships, only when using the BlueXP interface. Refer to Restoring data from backup files for more information.

Archive and compliance using SnapLock technology

Learn about ONTAP SnapLock

SnapLock is a high-performance compliance solution for organizations that use WORM storage to retain files in unmodified form for regulatory and governance purposes.

SnapLock helps to prevent deletion, change, or renaming of data to meet regulations such as SEC 17a-4(f), HIPAA, FINRA, CFTC, and GDPR. With SnapLock, you can create special-purpose volumes in which files can be stored and committed to a non-erasable, non-writable state either for a designated retention period or indefinitely. SnapLock allows this retention to be performed at the file level through standard open file protocols such as CIFS and NFS. The supported open file protocols for SnapLock are NFS (versions 2, 3, and 4) and CIFS (SMB 1.0, 2.0, and 3.0).

Using SnapLock, you commit files and snapshots to WORM storage, and set retention periods for WORM-

protected data. SnapLock WORM storage uses NetApp snapshot technology and can leverage SnapMirror replication, and SnapVault backups as the base technology for providing backup recovery protection for data. Learn more about WORM storage: Compliant WORM storage using NetApp SnapLock - TR-4526.

You can use an application to commit files to WORM over NFS or CIFS, or use the SnapLock autocommit feature to commit files to WORM automatically. You can use a *WORM appendable file* to retain data that is written incrementally, like log information. For more information see Use volume append mode to create WORM appendable files.

SnapLock supports data protection methods that should satisfy most compliance requirements:

- You can use SnapLock for SnapVault to WORM-protect snapshots on secondary storage. See Commit snapshots to WORM.
- You can use SnapMirror to replicate WORM files to another geographic location for disaster recovery. See Mirror WORM files.

SnapLock is a license-based feature of NetApp ONTAP. A single license entitles you to use SnapLock in strict Compliance mode, to satisfy external mandates like SEC Rule 17a-4(f), and a looser Enterprise mode, to meet internally mandated regulations for the protection of digital assets. SnapLock licenses are part of the ONTAP One software suite.

SnapLock is supported on all AFF and FAS systems as well as ONTAP Select. SnapLock is not a softwareonly solution; it is an integrated hardware and software solution. This distinction is important for strict WORM regulations such as SEC 17a-4(f), which requires an integrated hardware and software solution. For more information, refer to SEC Guidance to Broker-Dealers on the Use of Electronic Storage Media.

What you can do with SnapLock

After you configure SnapLock, you can complete the following tasks:

- Commit files to WORM
- Commit snapshots to WORM for secondary storage
- · Mirror WORM files for disaster recovery
- Retain WORM files during litigation using Legal Hold
- Delete WORM files using the privileged delete feature
- Set the file retention period
- Move a SnapLock volume
- · Lock a snapshot for protection against ransomware attacks
- Review snapLock use with the Audit Log
- Use SnapLock APIs

SnapLock Compliance and Enterprise modes

SnapLock Compliance and Enterprise modes differ mainly in the level at which each mode protects WORM files:

SnapLock mode	Protection level	WORM file deleting during
		retention

Compliance mode	At the disk level	Cannot be deleted
Enterprise mode	At the file level	Can be deleted by the compliance administrator using an audited "privileged delete" procedure

After the retention period has elapsed, you are responsible for deleting any files you no longer need. Once a file has been committed to WORM, whether under Compliance or Enterprise mode, it cannot be modified, even after the retention period has expired.

You cannot move a WORM file during or after the retention period. You can copy a WORM file, but the copy will not retain its WORM characteristics.

The following table shows the differences in capabilities supported by SnapLock Compliance and Enterprise modes:

Capability	SnapLock Compliance	SnapLock Enterprise
Enable and delete files using privileged delete	No	Yes
Reinitialize disks	No	Yes
Destroy SnapLock aggregates and volumes during retention period	No	Yes, with the exception of the SnapLock audit log volume
Rename aggregates or volumes	No	Yes
Use non-NetApp disks	No	No
Use the SnapLock volume for audit logging	Yes	Yes, beginning with ONTAP 9.5

Supported and unsupported features with SnapLock

The following table shows the features that are supported with SnapLock Compliance mode, SnapLock Enterprise mode, or both:

Feature	Supported with SnapLock Compliance	Supported with SnapLock Enterprise
Consistency Groups	No	No
Encrypted volumes	Yes, learn more about Encryption and SnapLock.	Yes, learn more about Encryption and SnapLock.

FabricPools on SnapLock aggregates	No	Yes, beginning with ONTAP 9.8. Learn more about FabricPool on SnapLock Enterprise aggregates.
Flash Pool aggregates	Yes.	Yes.
FlexClone	You can clone SnapLock volumes, but you cannot clone files on a SnapLock volume.	You can clone SnapLock volumes, but you cannot clone files on a SnapLock volume.
FlexGroup volumes	Yes, beginning with ONTAP 9.11.1. Learn more about FlexGroup volumes.	Yes, beginning with ONTAP 9.11.1. Learn more about FlexGroup volumes.
LUNs	No. Learn more about LUN support with SnapLock.	No. Learn more about LUN support with SnapLock.
MetroCluster configurations	Yes, beginning with ONTAP 9.3. Learn more about MetroCluster support.	Yes, beginning with ONTAP 9.3. Learn more about MetroCluster support.
Multi-admin verification (MAV)	Yes, beginning with ONTAP 9.13.1. Learn more about MAV support.	Yes, beginning with ONTAP 9.13.1. Learn more about MAV support.
SAN	No	No
Single-file SnapRestore	No	Yes
SnapMirror active sync	No	No
SnapRestore	No	Yes
SMTape	No	No
SnapMirror Synchronous	No	No
SSDs	Yes.	Yes.
Storage efficiency features	Yes, beginning with ONTAP 9.9.1. Learn more about storage efficiency support.	Yes, beginning with ONTAP 9.9.1. Learn more about storage efficiency support.

FabricPool on SnapLock Enterprise aggregates

FabricPools are supported on SnapLock Enterprise aggregates beginning with ONTAP 9.8. However, your account team needs to open a product variance request documenting that you understand that FabricPool data tiered to a public or private cloud is no longer protected by SnapLock because a cloud admin can delete that

data.



Any data that FabricPool tiers to a public or private cloud is no longer protected by SnapLock because that data can be deleted by a cloud administrator.

FlexGroup volumes

SnapLock supports FlexGroup volumes beginning with ONTAP 9.11.1; however, the following features are not supported:

- Legal-hold
- Event-based retention
- SnapLock for SnapVault (supported beginning with ONTAP 9.12.1)

You should also be aware of the following behaviors:

- The volume compliance clock (VCC) of a FlexGroup volume is determined by the VCC of the root constituent. All non-root constituents will have their VCC closely synced to the root VCC.
- SnapLock configuration properties are set only on the FlexGroup as a whole. Individual constituents cannot have different configuration properties, such as default retention time and autocommit period.

LUN support

LUNs are supported in SnapLock volumes only in scenarios where snapshots created on a non-SnapLock volume are transferred to a SnapLock volume for protection as part of SnapLock vault relationship. LUNs are not supported in read/write SnapLock volumes. Tamperproof snapshots however are supported on both SnapMirror source volumes and destination volumes that contain LUNs.

MetroCluster support

SnapLock support in MetroCluster configurations differs between SnapLock Compliance mode and SnapLock Enterprise mode.

SnapLock Compliance

- Beginning with ONTAP 9.3, SnapLock Compliance is supported on unmirrored MetroCluster aggregates.
- Beginning with ONTAP 9.3, SnapLock Compliance is supported on mirrored aggregates, but only if the aggregate is used to host SnapLock audit log volumes.
- SVM-specific SnapLock configurations can be replicated to primary and secondary sites using MetroCluster.

SnapLock Enterprise

- SnapLock Enterprise aggregates are supported.
- Beginning with ONTAP 9.3, SnapLock Enterprise aggregates with privileged delete are supported.
- SVM-specific SnapLock configurations can be replicated to both sites using MetroCluster.

MetroCluster configurations and compliance clocks

MetroCluster configurations use two compliance clock mechanisms, the Volume Compliance Clock (VCC) and the System Compliance Clock (SCC). The VCC and SCC are available to all SnapLock configurations. When you create a new volume on a node, its VCC is initialized with the current value of the SCC on that node. After the volume is created, the volume and file retention time is always tracked with the VCC.
When a volume is replicated to another site, its VCC is also replicated. When a volume switchover occurs, from Site A to Site B, for example, the VCC continues to be updated on Site B while the SCC on Site A halts when Site A goes offline.

When Site A is brought back online and the volume switchback is performed, the Site A SCC clock restarts while the VCC of the volume continues to be updated. Because the VCC is continuously updated, regardless of switchover and switchback operations, the file retention times do not depend on SCC clocks and do not stretch.

Multi-admin verification (MAV) support

Beginning with ONTAP 9.13.1, a cluster administrator can explicitly enable multi-admin verification on a cluster to require quorum approval before some SnapLock operations are executed. When MAV is enabled, SnapLock volume properties such as default-retention-time, minimum-retention-time, maximum-retention-time, volume-append-mode, autocommit-period and privileged-delete will require quorum approval. Learn more about MAV.

Storage efficiency

Beginning with ONTAP 9.9.1, SnapLock supports storage efficiency features, such as data compaction, crossvolume-deduplication, and adaptive compression for SnapLock volumes and aggregates. For more information about storage efficiency, see ONTAP storage efficiency overview.

Encryption

ONTAP offers both software- and hardware-based encryption technologies for ensuring that data at rest cannot be read if the storage medium is repurposed, returned, misplaced, or stolen.

Disclaimer: NetApp cannot guarantee that SnapLock-protected WORM files on self-encrypting drives or volumes will be retrievable if the authentication key is lost or if the number of failed authentication attempts exceeds the specified limit and results in the drive being permanently locked. You are responsible for ensuring against authentication failures.



Encrypted volumes are supported on SnapLock aggregates.

7-Mode Transition

You can migrate SnapLock volumes from 7-Mode to ONTAP by using the Copy-Based Transition (CBT) feature of the 7-Mode Transition Tool. The SnapLock mode of the destination volume, Compliance or Enterprise, must match the SnapLock mode of the source volume. You cannot use Copy-Free Transition (CFT) to migrate SnapLock volumes.

Configure SnapLock

Learn about configuring ONTAP SnapLock

Before you use SnapLock, you need to configure SnapLock by completing various tasks such as install the SnapLock license for each node that hosts an aggregate with a SnapLock volume, initialize the Compliance Clock, create a SnapLock aggregate for clusters running ONTAP releases earlier than ONTAP 9.10.1, create and mount a SnapLock volume, and more.

Initialize the ONTAP Compliance Clock

SnapLock uses the *volume Compliance Clock* to ensure against tampering that might alter the retention period for WORM files. You must first initialize the *system ComplianceClock* on each node that hosts a SnapLock aggregate.

Beginning with ONTAP 9.14.1, you can initialize or reinitialize the system Compliance Clock when there are no SnapLock volumes or no volumes with snapshot locking enabled. The ability to reinitialize enables system administrators to reset the system Compliance Clock in instances where it might have been incorrectly initialized or to correct clock drift on the system. In ONTAP 9.13.1 and earlier releases, once you initialize the Compliance Clock on a node, you cannot initialize it again.

Before you begin

To reinitialize the Compliance Clock:

- All nodes in the cluster must be in the healthy state.
- All volumes must be online.
- No volumes can be present the the recovery queue.
- No SnapLock volumes can be present.
- No volumes with snapshot locking enabled can be present.

General requirements for initializing the Compliance Clock:

- You must be a cluster administrator to perform this task.
- The SnapLock license must be installed on the node.

About this task

The time on the system Compliance Clock is inherited by the *volume Compliance Clock*, the latter of which controls the retention period for WORM files on the volume. The volume Compliance Clock is initialized automatically when you create a new SnapLock volume.



The initial setting of the system Compliance Clock is based on the current hardware system clock. For that reason, you should verify that the system time and time zone are correct before initializing the system Compliance Clock on each node. Once you initialize the system Compliance Clock on a node, you cannot initialize it again when SnapLock volumes or volumes with locking enabled are present.

Steps

You can use the ONTAP CLI to initialize the Compliance Clock or, beginning with ONTAP 9.12.1, you can use System Manager to initialize the Compliance Clock.

System Manager

- 1. Navigate to **Cluster > Overview**.
- 2. In the Nodes section, click Initialize SnapLock Compliance Clock.
- 3. To display the **Compliance Clock** column and to verify that the Compliance Clock is initialized, in the **Cluster > Overview > Nodes** section, click **Show/Hide** and select **SnapLock Compliance Clock**.

CLI

1. Initialize the system Compliance Clock:

```
snaplock compliance-clock initialize -node node name
```

The following command initializes the system Compliance Clock on node1:

cluster1::> snaplock compliance-clock initialize -node node1

Learn more about snaplock compliance-clock initialize in the ONTAP command reference.

2. When prompted, confirm that the system clock is correct and that you want to initialize the Compliance Clock:

Warning: You are about to initialize the secure ComplianceClock of the node "nodel" to the current value of the node's system clock. This procedure can be performed only once on a given node, so you should ensure that the system time is set correctly before proceeding.

The current node's system clock is: Mon Apr 25 06:04:10 GMT 2016

Do you want to continue? (y|n): y

3. Repeat this procedure for each node that hosts a SnapLock aggregate.

Enable Compliance Clock resynchronization for an NTP-configured system

You can enable the SnapLock Compliance Clock synchronization feature when an NTP server is configured.

Before you begin

- This feature is available only at the advanced privilege level.
- You must be a cluster administrator to perform this task.
- The SnapLock license must be installed on the node.
- This feature is available only for Cloud Volumes ONTAP, ONTAP Select, and VSIM platforms.

About this task

When the SnapLock secure clock daemon detects a skew beyond the threshold, ONTAP uses the system time to reset both the system and volume Compliance Clocks. A period of 24 hours is set as the skew threshold. This means that the system Compliance Clock is synchronized to the system clock only if the skew is more than a day old.

The SnapLock secure clock daemon detects a skew and changes the Compliance Clock to the system time. Any attempt at modifying the system time to force the Compliance Clock to synchronize to the system time fails, since the Compliance Clock synchronizes to the system time only if the system time is synchronized with the NTP time.

Steps

1. Enable the SnapLock Compliance Clock synchronization feature when an NTP server is configured:

```
snaplock compliance-clock ntp
```

The following command enables the system Compliance Clock synchronization feature:

```
cluster1::*> snaplock compliance-clock ntp modify -is-sync-enabled true
```

Learn more about snaplock compliance-clock ntp modify in the ONTAP command reference.

- 2. When prompted, confirm that the configured NTP servers are trusted and that the communications channel is secure to enable the feature:
- 3. Check that the feature is enabled:

snaplock compliance-clock ntp show

The following command checks that the system Compliance Clock synchronization feature is enabled:

```
cluster1::*> snaplock compliance-clock ntp show
```

Enable clock sync to NTP system time: true

Learn more about snaplock compliance-clock ntp show in the ONTAP command reference.

Create an ONTAP SnapLock aggregate

You use the volume <code>-snaplock-type</code> option to specify a Compliance or Enterprise SnapLock volume type. For releases earlier than ONTAP 9.10.1, you must create a separate SnapLock aggregate. Beginning with ONTAP 9.10.1, SnapLock and non-SnapLock volumes can exist on the same aggregate; therefore, you are no longer required to create a separate SnapLock aggregate if you are using ONTAP 9.10.1.

Before you begin

- You must be a cluster administrator to perform this task.
- The SnapLock license must be installed on the node. This license in included in ONTAP One.
- The Compliance Clock on the node must be initialized.

• If you have partitioned the disks as "root", "data1", and "data2", you must ensure that spare disks are available.

Upgrade considerations

When upgrading to ONTAP 9.10.1, existing SnapLock and non-SnapLock aggregates are upgraded to support the existence of both SnapLock and non-SnapLock volumes; however, the existing SnapLock volume attributes are not automatically updated. For example, data-compaction, cross-volume-dedupe, and cross-volume-background-dedupe fields remain unchanged. New SnapLock volumes created on existing aggregates have the same default values as non-SnapLock volumes, and the default values for new volumes and aggregates are platform dependent.

Revert considerations

If you need to revert to an ONTAP version earlier than 9.10.1, you must move all SnapLock Compliance, SnapLock Enterprise, and SnapLock volumes to their own SnapLock aggregates.

About this task

- You cannot create Compliance aggregates with the SyncMirror option.
- You can create mirrored Compliance aggregates in a MetroCluster configuration only if the aggregate is used to host SnapLock audit log volumes.



In a MetroCluster configuration, SnapLock Enterprise is supported on mirrored and unmirrored aggregates. SnapLock Compliance is supported only on unmirrored aggregates.

Steps

1. Create a SnapLock aggregate:

```
storage aggregate create -aggregate <aggregate_name> -node <node_name>
-diskcount <number_of_disks> -snaplock-type <compliance|enterprise>
```

The following command creates a SnapLock Compliance aggregate named aggr1 with three disks on node1:

cluster1::> storage aggregate create -aggregate aggr1 -node node1
-diskcount 3 -snaplock-type compliance

Learn more about storage aggregate create in the ONTAP command reference.

Create and mount ONTAP SnapLock volumes

You must create a SnapLock volume for the files or snapshots that you want to commit to the WORM state. Beginning with ONTAP 9.10.1, any volume you create, regardless of the aggregate type, is created by default as a non-SnapLock volume. You must use the <code>-snaplock-type</code> option to explicitly create a SnapLock volume by specifying either Compliance or Enterprise as the SnapLock type. By default, the SnapLock type is set to non-snaplock.

Before you begin

- The SnapLock aggregate must be online.
- You should verify that a SnapLock license is installed. If a SnapLock license is not installed on the node, you must install it. This license is included with ONTAP One. Prior to ONTAP One, the SnapLock license was included in the Security and Compliance bundle. The Security and Compliance bundle is no longer offered but is still valid. Although not currently required, existing customers can choose to upgrade to ONTAP One.
- The Compliance Clock on the node must be initialized.

About this task

With the proper SnapLock permissions, you can destroy or rename an Enterprise volume at any time. You cannot destroy a Compliance volume until the retention period has elapsed. You can never rename a Compliance volume.

You can clone SnapLock volumes, but you cannot clone files on a SnapLock volume. The clone volume will be of the same SnapLock type as the parent volume.



LUNs are not supported in SnapLock volumes. LUNs are supported in SnapLock volumes only in scenarios where snapshots created on a non-SnapLock volume are transferred to a SnapLock volume for protection as part of SnapLock vault relationship. LUNs are not supported in read/write SnapLock volumes. Tamperproof snapshots however are supported on both SnapMirror source volumes and destination volumes that contain LUNs.

Perform this task using ONTAP System Manager or the ONTAP CLI.

System Manager

Beginning with ONTAP 9.12.1, you can use System Manager to create a SnapLock volume.

Steps

- 1. Navigate to Storage > Volumes and click Add.
- 2. In the Add Volume window, click More Options.
- 3. Enter the new volume information, including the name and size of the volume.
- 4. Select Enable SnapLock and choose the SnapLock type, either Compliance or Enterprise.
- 5. In the Auto-Commit Files section, select Modified and enter the amount of time a file should remain unchanged before it is automatically committed. The minimum value is 5 minutes and the maximum value is 10 years.
- 6. In the Data Retention section, select the minimum and maximum retention period.
- 7. Select the default retention period.
- 8. Click Save.
- 9. Select the new volume in the **Volumes** page to verify the SnapLock settings.

CLI

1. Create a SnapLock volume:

```
volume create -vserver <SVM_name> -volume <volume_name> -aggregate
<aggregate name> -snaplock-type <compliance|enterprise>
```

Learn more about volume create in the ONTAP command reference. The following options are not available for SnapLock volumes: -nvfail, -atime-update, -is -autobalance-eligible, -space-mgmt-try-first, and vmalign.

The following command creates a SnapLock Compliance volume named vol1 on aggr1 on vs1:

```
cluster1::> volume create -vserver vs1 -volume vol1 -aggregate aggr1
-snaplock-type compliance
```

Mount a SnapLock volume

You can mount a SnapLock volume to a junction path in the SVM namespace for NAS client access.

Before you begin

The SnapLock volume must be online.

About this task

- You can mount a SnapLock volume only under the root of the SVM.
- You cannot mount a regular volume under a SnapLock volume.

Steps

1. Mount a SnapLock volume:

volume mount -vserver SVM_name -volume volume_name -junction-path path

Learn more about volume mount in the ONTAP command reference.

The following command mounts a SnapLock volume named <code>vol1</code> to the junction path /sales in the <code>vs1</code> namespace:

cluster1::> volume mount -vserver vs1 -volume vol1 -junction-path /sales

Set the ONTAP SnapLock retention time

You can set the retention time for a file explicitly, or you can use the default retention period for the volume to derive the retention time. Unless you set the retention time explicitly, SnapLock uses the default retention period to calculate the retention time. You can also set file retention after an event.

About retention period and retention time

The *retention period* for a WORM file specifies the length of time the file must be retained after it is committed to the WORM state. The *retention time* for a WORM file is the time after which the file no longer needs to be retained. A retention period of 20 years for a file committed to the WORM state on 10 November 2020 6:00 a.m., for example, would yield a retention time of 10 November 2040 6:00 a.m.



Beginning with ONTAP 9.10.1, you can set a retention time up to October 26, 3058 and a retention period up to 100 years. When you extend retention dates, older policies are converted automatically. In ONTAP 9.9.1 and earlier releases, unless you set the default retention period to infinite, the maximum supported retention time is January 19 2071 (GMT).

Important replication considerations

When establishing a SnapMirror relationship with a SnapLock source volume using a retention date later than January 19th 2071 (GMT), the destination cluster must be running ONTAP 9.10.1 or later or the SnapMirror transfer will fail.

Important revert considerations

ONTAP prevents you from reverting a cluster from ONTAP 9.10.1 to an earlier ONTAP version when there are any files with a retention period later than "January 19, 2071 8:44:07 AM".

Understanding the retention periods

A SnapLock Compliance or Enterprise volume has four retention periods:

• Minimum retention period (min), with a default of 0

 (\mathbf{i})

- Maximum retention period (max), with a default of 30 years
- Default retention period, with a default equal to min for both Compliance mode and Enterprise mode beginning with ONTAP 9.10.1. In ONTAP releases earlier than ONTAP 9.10.1, the default retention period depends on the mode:
 - For Compliance mode, the default is equal to max.
 - For Enterprise mode, the default is equal to min.
- Unspecified retention period.

Beginning with ONTAP 9.8, you can set the retention period on files in a volume to unspecified, to enable the file to be retained until you set an absolute retention time. You can set a file with absolute retention time to unspecified retention and back to absolute retention as long as the new absolute retention time is later than the absolute time you previously set.

Beginning with ONTAP 9.12.1, WORM files with the retention period set to unspecified are guaranteed to have a retention period set to the minimum retention period configured for the SnapLock volume. When you change the file retention period from unspecified to an absolute retention time, the new retention time specified must be greater than the minimum retention time already set on the file.

So, if you do not set the retention time explicitly before committing a Compliance-mode file to the WORM state, and you do not modify the defaults, the file will be retained for 30 years. Similarly, if you do not set the retention time explicitly before committing an Enterprise-mode file to the WORM state, and you do not modify the defaults, the file will be retained for 0 years, or, effectively, not at all.

Set the default retention period

You can use the volume snaplock modify command to set the default retention period for files on a SnapLock volume.

Before you begin

The SnapLock volume must be online.

About this task

The following table shows the possible values for the default retention period option:



The default retention period must be greater than or equal to (>=) the minimum retention period and less than or equal to (<=) the maximum retention period.

Value	Unit	Notes
0 - 65535	seconds	
0 - 24	hours	
0 - 365	days	
0 - 12	months	

Value	Unit	Notes
0 - 100	years	Beginning with ONTAP 9.10.1. For earlier ONTAP releases, the value is 0 - 70.
max	-	Use the maximum retention period.
min	-	Use the minimum retention period.
infinite	-	Retain the files forever.
unspecified	-	Retain the files until an absolute retention period is set.

The values and ranges for the maximum and minimum retention periods are identical, except for max and min, which are not applicable. For more information about this task, see Set the retention time overview.

You can use the volume snaplock show command to view the retention period settings for the volume. Learn more about volume snaplock show in the ONTAP command reference.



After a file has been committed to the WORM state, you can extend but not shorten the retention period.

Steps

1. Set the default retention period for files on a SnapLock volume:

```
volume snaplock modify -vserver SVM_name -volume volume_name -default
-retention-period default_retention_period -minimum-retention-period
min_retention_period -maximum-retention-period max_retention_period
```

Learn more about volume snaplock modify in the ONTAP command reference.



The following examples assume that the minimum and maximum retention periods have not been modified previously.

The following command sets the default retention period for a Compliance or Enterprise volume to 20 days:

```
cluster1::> volume snaplock modify -vserver vs1 -volume vol1 -default
-retention-period 20days
```

The following command sets the default retention period for a Compliance volume to 70 years:

```
cluster1::> volume snaplock modify -vserver vs1 -volume vol1 -maximum
-retention-period 70years
```

The following command sets the default retention period for an Enterprise volume to 10 years:

cluster1::> volume snaplock modify -vserver vs1 -volume vol1 -default
-retention-period max -maximum-retention-period 10years

The following commands set the default retention period for an Enterprise volume to 10 days:

```
cluster1::> volume snaplock modify -vserver vs1 -volume vol1 -minimum
-retention-period 10days
cluster1::> volume snaplock modify -vserver vs1 -volume vol1 -default
-retention-period min
```

The following command sets the default retention period for a Compliance volume to infinite:

```
cluster1::> volume snaplock modify -vserver vs1 -volume vol1 -default
-retention-period infinite -maximum-retention-period infinite
```

Set the retention time for a file explicitly

You can set the retention time for a file explicitly by modifying its last access time. You can use any suitable command or program over NFS or CIFS to modify the last access time.

About this task

After a file has been committed to WORM, you can extend but not shorten the retention time. The retention time is stored in the atime field for the file.



You cannot explicitly set the retention time of a file to infinite. That value is only available when you use the default retention period to calculate the retention time.

Steps

1. Use a suitable command or program to modify the last access time for the file whose retention time you want to set.

In a UNIX shell, use the following command to set a retention time of 21 November 2020 6:00 a.m. on a file named document.txt:

touch -a -t 202011210600 document.txt



You can use any suitable command or program to modify the last access time in Windows.

Set the file retention period after an event

Beginning with ONTAP 9.3, you can define how long a file is retained after an event occurs by using the SnapLock *Event Based Retention (EBR)* feature.

Before you begin

• You must be a SnapLock administrator to perform this task.

Create a SnapLock administrator account

• You must have logged in on a secure connection (SSH, console, or ZAPI).

About this task

The *event retention policy* defines the retention period for the file after the event occurs. The policy can be applied to a single file or all the files in a directory.

- If a file is not a WORM file, it will be committed to the WORM state for the retention period defined in the policy.
- If a file is a WORM file or a WORM appendable file, its retention period will be extended by the retention period defined in the policy.

You can use a Compliance-mode or Enterprise-mode volume.



EBR policies cannot be applied to files under a Legal Hold.

For advanced usage, see Compliant WORM Storage Using NetApp SnapLock.

Using EBR to extend the retention period of already existing WORM files

EBR is convenient when you want to extend the retention period of already existing WORM files. For example, it might be your firm's policy to retain employee W-4 records in unmodified form for three years after the employee changes a withholding election. Another company policy might require that W-4 records be retained for five years after the employee is terminated.

In this situation, you could create an EBR policy with a five-year retention period. After the employee is terminated (the "event"), you would apply the EBR policy to the employee's W-4 record, causing its retention period to be extended. That will usually be easier than extending the retention period manually, particularly when a large number of files is involved.

Steps

1. Create an EBR policy:

```
snaplock event-retention policy create -vserver SVM_name -name policy_name
-retention-period retention period
```

The following command creates the EBR policy employee_exit on vs1 with a retention period of ten years:

```
cluster1::>snaplock event-retention policy create -vserver vs1 -name
employee exit -retention-period 10years
```

2. Apply an EBR policy:

snaplock event-retention apply -vserver SVM_name -name policy_name -volume
volume name -path path name

The following command applies the EBR policy employee_exit on vs1 to all the files in the directory d1:

cluster1::>snaplock event-retention apply -vserver vs1 -name
employee_exit -volume vol1 -path /d1

Related information

- snaplock event-retention policy create
- snaplock event-retention apply

Create an ONTAP SnapLock-protected audit log

If you are using ONTAP 9.9.1 or earlier, you must first create a SnapLock aggregate and then you must create a SnapLock-protected audit log before performing a privileged delete or SnapLock volume move. The audit log records the creation and deletion of SnapLock administrator accounts, modifications to the log volume, whether privileged delete is enabled, privileged delete operations, and SnapLock volume move operations.

Beginning with ONTAP 9.10.1, you no longer create a SnapLock aggregate. You must use the -snaplock-type option to explicitly create a SnapLock volume by specifying either Compliance or Enterprise as the SnapLock type.

Before you begin

If you are using ONTAP 9.9.1 or earlier, you must be a cluster administrator to create a SnapLock aggregate.

About this task

You cannot delete an audit log until the log file retention period has elapsed. You cannot modify an audit log even after the retention period has elapsed. This is true for both SnapLock Compliance and Enterprise modes.



In ONTAP 9.4 and earlier, you cannot use a SnapLock Enterprise volume for audit logging. You must use a SnapLock Compliance volume. In ONTAP 9.5 and later, you can use either a SnapLock Enterprise volume or a SnapLock Compliance volume for audit logging. In all cases, the audit log volume must be mounted at the junction path /snaplock_audit_log. No other volume can use this junction path.

You can find the SnapLock audit logs in the /snaplock_log directory under the root of the audit log volume, in subdirectories named privdel_log (privileged delete operations) and system_log (everything else). Audit log file names contain the timestamp of the first logged operation, making it easy to search for records by the approximate time that operations were executed.

- You can use the snaplock log file show command to view the log files on the audit log volume.
- You can use the *snaplock log file archive* command to archive the current log file and create a new one, which is useful in cases where you need to record audit log information in a separate file.

Learn more about snaplock log file show and snaplock log file archive in the ONTAP command reference.



A data protection volume cannot be used as a SnapLock audit log volume.

Steps

1. Create a SnapLock aggregate.

Create a SnapLock aggregate

2. On the SVM that you want to configure for audit logging, create a SnapLock volume.

Create a SnapLock volume

3. Configure the SVM for audit logging:

```
snaplock log create -vserver SVM_name -volume snaplock_volume_name -max-log
-size size -retention-period default_retention_period
```



The minimum default retention period for audit log files is six months. If the retention period of an affected file is longer than the retention period of the audit log, the retention period of the log inherits the retention period of the file. So, if the retention period for a file deleted using privileged delete is 10 months, and the retention period of the audit log is 8 months, the retention period of the log is extended to 10 months. For more information about retention time and default retention period, see Set the retention time.

The following command configures SVM1 for audit logging using the SnapLock volume logVol. The audit log has a maximum size of 20 GB and is retained for eight months.

```
SVM1::> snaplock log create -vserver SVM1 -volume logVol -max-log-size
20GB -retention-period 8months
```

Learn more about snaplock log create in the ONTAP command reference.

4. On the SVM that you configured for audit logging, mount the SnapLock volume at the junction path /snaplock audit log.

Mount a SnapLock volume

Verify ONTAP SnapLock settings

You can use the volume file fingerprint start and volume file fingerprint dump commands to view key information about files and volumes, including the file type (regular, WORM, or WORM appendable), the volume expiration date, and so forth.

Steps

1. Generate a file fingerprint:

```
volume file fingerprint start -vserver <SVM_name> -file <file_path>
```

```
svm1::> volume file fingerprint start -vserver svm1 -file
/vol/sle/vol/f1
File fingerprint operation is queued. Run "volume file fingerprint show
-session-id 16842791" to view the fingerprint session status.
```

The command generates a session ID that you can use as input to the volume file fingerprint dump command.



You can use the volume file fingerprint show command with the session ID to monitor the progress of the fingerprint operation. Make sure that the operation has completed before attempting to display the fingerprint.

2. Display the fingerprint for the file:

volume file fingerprint dump -session-id <session ID>

```
svm1::> volume file fingerprint dump -session-id 33619976
        Vserver:svm1
        Session-ID:33619976
        Volume:slc vol
        Path:/vol/slc vol/f1
        Data
Fingerprint:MOFJVevxNSJm3C/4Bn5oEEYH51CrudOzZYK4r5Cfy1g=Metadata
Fingerprint:8iMjqJXiNcqgXT5XuRhLiEwIrJEihDmwS0hrexnjgmc=Fingerprint
Algorithm:SHA256
        Fingerprint Scope:data-and-metadata
        Fingerprint Start Time:1460612586
        Formatted Fingerprint Start Time: Thu Apr 14 05:43:06 GMT 2016
        Fingerprint Version:3
        **SnapLock License:available**
        Vserver UUID:acf7ae64-00d6-11e6-a027-0050569c55ae
        Volume MSID:2152884007
        Volume DSID:1028
        Hostname:my host
        Filer ID:5f18eda2-00b0-11e6-914e-6fb45e537b8d
        Volume Containing Aggregate:slc aggr1
        Aggregate ID:c84634aa-c757-4b98-8f07-eefe32565f67
        **SnapLock System ComplianceClock:1460610635
        Formatted SnapLock System ComplianceClock: Thu Apr 14 05:10:35
GMT 2016
        Volume SnapLock Type:compliance
        Volume ComplianceClock:1460610635
        Formatted Volume ComplianceClock: Thu Apr 14 05:10:35 GMT 2016
        Volume Expiry Date:1465880998**
```

Is Volume Expiry Date Wraparound: false Formatted Volume Expiry Date: Tue Jun 14 05:09:58 GMT 2016 Filesystem ID:1028 File ID:96 File Type:worm File Size:1048576 Creation Time:1460612515 Formatted Creation Time: Thu Apr 14 05:41:55 GMT 2016 Modification Time:1460612515 Formatted Modification Time: Thu Apr 14 05:41:55 GMT 2016 Changed Time:1460610598 Is Changed Time Wraparound: false Formatted Changed Time: Thu Apr 14 05:09:58 GMT 2016 Retention Time:1465880998 Is Retention Time Wraparound:false Formatted Retention Time: Tue Jun 14 05:09:58 GMT 2016 Access Time:-Formatted Access Time:-Owner ID:0 Group ID:0 Owner SID:-Fingerprint End Time: 1460612586 Formatted Fingerprint End Time: Thu Apr 14 05:43:06 GMT 2016

Manage WORM files

Manage WORM files with ONTAP SnapLock

You can manage WORM files in the following ways:

- Commit files to WORM
- Commit snapshots to WORM on a vault destination
- Mirror WORM files for disaster recovery
- Retain WORM files during litigation
- Delete WORM files

Commit files to WORM using ONTAP SnapLock

You can commit files to WORM (write once, read many) either manually or by committing them automatically. You can also create WORM appendable files.

Commit files to WORM manually

You commit a file to WORM manually by making the file read-only. You can use any suitable command or program over NFS or CIFS to change the read-write attribute of a file to read-only. You might choose to manually commit files if you want to ensure an application has finished writing to a file so that the file isn't

committed prematurely or if there are scaling issues for the autocommit scanner because of a high number of volumes.

Before you begin

- The file you want to commit must reside on a SnapLock volume.
- The file must be writable.

About this task

The volume ComplianceClock time is written to the ctime field of the file when the command or program is executed. The ComplianceClock time determines when the retention time for the file has been reached.

Steps

1. Use a suitable command or program to change the read-write attribute of a file to read-only.

In a UNIX shell, use the following command to make a file named document.txt read-only:

chmod -w document.txt

In a Windows shell, use the following command to make a file named document.txt read-only:

```
attrib +r document.txt
```

Commit files to WORM automatically

The SnapLock autocommit feature enables you to commit files to WORM automatically. The autocommit feature commits a file to WORM state on a SnapLock volume if the file did not change for the autocommitperiod

duration. The autocommit feature is disabled by default.

Before you begin

- The files you want to autocommit must reside on a SnapLock volume.
- The SnapLock volume must be online.
- The SnapLock volume must be a read-write volume.



The SnapLock autocommit feature scans through all of the files in the volume and commits a file if it meets the autocommit requirement. There might be a time interval between when the file is ready for autocommit and when it is actually committed by the SnapLock autocommit scanner. However, the file is still protected from modifications and deletion by the file system as soon as it is eligible for autocommit.

About this task

The *autocommit period* specifies the amount of time that files must remain unchanged before they are autocommitted. Changing a file before the autocommit period has elapsed restarts the autocommit period for the file.

The following table shows the possible values for the autocommit period:

Value	Unit	Notes
none	-	The default.
5 - 5256000	minutes	-
1 - 87600	hours	-
1 - 3650	days	-
1 - 120	months	-
1 - 10	years	-



The minimum value is 5 minutes and the maximum value is 10 years.

Steps

1. Autocommit files on a SnapLock volume to WORM:

volume snaplock modify -vserver SVM_name -volume volume_name -autocommit -period autocommit period

Learn more about volume snaplock modify in the ONTAP command reference.

The following command autocommits the files on volume vol1 of SVM vs1, as long as the files remain unchanged for 5 hours:

```
cluster1::>volume snaplock modify -vserver vs1 -volume vol1 -autocommit
-period 5hours
```

Create a WORM appendable file

A WORM appendable file retains data written incrementally, like log entries. You can use any suitable command or program to create a WORM appendable file, or you can use the SnapLock *volume append mode* feature to create WORM appendable files by default.

Use a command or program to create a WORM appendable file

You can use any suitable command or program over NFS or CIFS to create a WORM appendable file. A WORM appendable file retains data written incrementally, like log entries. Data is appended to the file in 256 KB chunks. As each chunk is written, the previous chunk becomes WORM-protected. You cannot delete the file until the retention period has elapsed.

Before you begin

The WORM appendable file must reside on a SnapLock volume.

About this task

Data does not have to be written sequentially to the active 256 KB chunk. When data is written to byte

n×256KB+1 of the file, the previous 256 KB segment becomes WORM-protected.

Any unordered writes beyond the current active 256 KB chunk will result in the active 256KB chunk being reset to the latest offset and will cause writes to older offsets to fail with a 'Read Only File System (ROFS)' error. The write offsets are dependent on the client application. A client that does not conform to the WORM append file write semantics can cause incorrect termination of the write contents. Therefore, it is recommended to either ensure that the client follows the offset restrictions for unordered writes, or to ensure synchronous writes by mounting the file system in synchronous mode.

Steps

1. Use a suitable command or program to create a zero-length file with the desired retention time.

In a UNIX shell, use the following command to set a retention time of 21 November 2020 6:00 a.m. on a zero-length file named document.txt:

touch -a -t 202011210600 document.txt

2. Use a suitable command or program to change the read-write attribute of the file to read-only.

In a UNIX shell, use the following command to make a file named document.txt read-only:

chmod 444 document.txt

3. Use a suitable command or program to change the read-write attribute of the file back to writable.



This step is not deemed a compliance risk because there is no data in the file.

In a UNIX shell, use the following command to make a file named document.txt writable:

chmod 777 document.txt

4. Use a suitable command or program to start writing data to the file.

In a UNIX shell, use the following command to write data to document.txt:

echo test data >> document.txt



Change the file permissions back to read-only when you no longer need to append data to the file.

Use volume append mode to create WORM appendable files

Beginning with ONTAP 9.3, you can use the SnapLock *volume append mode* (VAM) feature to create WORM appendable files by default. A WORM appendable file retains data written incrementally, like log entries. Data is appended to the file in 256 KB chunks. As each chunk is written, the previous chunk becomes WORM-protected. You cannot delete the file until the retention period has elapsed.

Before you begin

- The WORM appendable file must reside on a SnapLock volume.
- The SnapLock volume must be unmounted and empty of snapshots and user-created files.

About this task

Data does not have to be written sequentially to the active 256 KB chunk. When data is written to byte n×256KB+1 of the file, the previous 256 KB segment becomes WORM-protected.

If you specify an autocommit period for the volume, WORM appendable files that are not modified for a period greater than the autocommit period are committed to WORM.



VAM is not supported on SnapLock audit log volumes.

Steps

1. Enable VAM:

```
volume snaplock modify -vserver SVM_name -volume volume_name -is-volume-append
-mode-enabled true false
```

Learn more about volume snaplock modify in the ONTAP command reference.

The following command enables VAM on volume vol1 of SVMvs1:

```
cluster1::>volume snaplock modify -vserver vs1 -volume vol1 -is-volume
-append-mode-enabled true
```

2. Use a suitable command or program to create files with write permissions.

The files are WORM-appendable by default.

Commit snapshots to WORM on an ONTAP vault destination

You can use SnapLock for SnapVault to WORM-protect snapshots on secondary storage. You perform all of the basic SnapLock tasks on the vault destination. The destination volume is automatically mounted read-only, so there is no need to explicitly commit the snapshots to WORM.

Before you begin

- If you want to use System Manager to configure the relationship, both the source and the destination clusters must be running ONTAP 9.15.1 or later.
- On the destination cluster:
 - Install the SnapLock license.
 - Initialize the Compliance Clock.
 - If you are using the CLI with an ONTAP release earlier than 9.10.1, create a SnapLock aggregate.
- · The protection policy must be of type "vault".
- The source and destination aggregates must be 64-bit.

- The source volume cannot be a SnapLock volume.
- If you are using the ONTAP CLI, the source and destination volumes must be created in peered clusters and SVMs.

About this task

The source volume can use NetApp or non-NetApp storage.



You cannot rename a snapshot that is committed to the WORM state.

You can clone SnapLock volumes, but you cannot clone files on a SnapLock volume.



LUNs are not supported in SnapLock volumes. LUNs are supported in SnapLock volumes only in scenarios where snapshots created on a non-SnapLock volume are transferred to a SnapLock volume for protection as part of SnapLock vault relationship. LUNs are not supported in read/write SnapLock volumes. Tamperproof snapshots, however, are supported on both SnapMirror source volumes and destination volumes that contain LUNs.

Beginning with ONTAP 9.10.1, SnapLock and non-SnapLock volumes can exist on the same aggregate; therefore, you are no longer required to create a separate SnapLock aggregate if you are using ONTAP 9.10.1. You use the volume '-snaplock-type' option to specify a Compliance or Enterprise SnapLock volume type. In ONTAP releases earlier than ONTAP 9.10.1, the SnapLock mode, Compliance or Enterprise, is inherited from the aggregate. Version-flexible destination volumes are not supported. The language setting of the destination volume must match the language setting of the source volume.

A SnapLock volume that is a vault destination has a default retention period assigned to it. The value for this period is initially set to a minimum of 0 years for SnapLock Enterprise volumes and a maximum of 30 years for SnapLock Compliance volumes. Each NetApp snapshot is committed with this default retention period at first. The retention period can be extended later, if needed. For more information, see <u>Set retention time overview</u>.

Beginning with ONTAP 9.14.1, you can specify retention periods for specific SnapMirror labels in the SnapMirror policy of the SnapMirror relationship so that the replicated snapshots from the source to the destination volume are retained for the retention-period specified in the rule. If no retention period is specified, the default-retention-period of the destination volume is used.

Beginning with ONTAP 9.13.1, you can instantaneously restore a locked snapshot on the destination SnapLock volume of a SnapLock vault relationship by creating a FlexClone with the snaplock-type option set to non-snaplock and specifying the snapshot as the "parent-snapshot" when executing the volume clone creation operation. Learn more about creating a FlexClone volume with a SnapLock type.

For MetroCluster configurations, you should be aware of the following:

- You can create a SnapVault relationship only between sync-source SVMs, not between a sync-source SVM and a sync-destination SVM.
- You can create a SnapVault relationship from a volume on a sync-source SVM to a data-serving SVM.
- You can create a SnapVault relationship from a volume on a data-serving SVM to a DP volume on a syncsource SVM.

The following illustration shows the procedure for initializing a SnapLock vault relationship:

Steps

You can use the ONTAP CLI to create a SnapLock vault relationship or, beginning with ONTAP 9.15.1, you can use System Manager to create a SnapLock vault relationship.

System Manager

- 1. If the volume doesn't already exist, on the source cluster, navigate to **Storage > Volumes** and select **Add**.
- 2. In the Add Volume window, choose More Options.
- 3. Enter the volume name, size, export policy and share name.
- 4. Save your changes.
- 5. On the destination cluster, navigate to **Protection > Relationships**.
- 6. Above the **Source** column, select **Protect** and choose **Volumes** from the menu.
- 7. In the Protect volumes window, choose Vault as the protection policy.
- 8. In the Source section, select the cluster, storage VM, and volume you want to protect.
- 9. In the **Destination** section, under **Configuration details**, select **Lock destination snapshots**, and then choose **SnapLock for SnapVault** as the locking method. **Locking method** is not displayed if the policy type selected is not of type vault, if the SnapLock license is not installed, or if the Compliance Clock is not initialized.
- 10. If it is not already enabled, select Initialize SnapLock Compliance Clock.
- 11. Save your changes.

CLI

1. On the destination cluster, create a SnapLock destination volume of type DP that is either the same or greater in size than the source volume:

```
volume create -vserver <SVM_name> -volume <volume_name> -aggregate
<aggregate_name> -snaplock-type <compliance|enterprise> -type DP
-size <size>
```

The following command creates a 2GB SnapLock Compliance volume named dstvolB in SVM2 on the aggregate node01 aggr:

cluster2::> volume create -vserver SVM2 -volume dstvolB -aggregate
node01_aggr -snaplock-type compliance -type DP -size 2GB

- 2. On the destination cluster, set the default retention period.
- 3. Create a new replication relationship between the non-SnapLock source and the new SnapLock destination you created.

This example creates a new SnapMirror relationship with destination SnapLock volume dstvolB using a policy of XDPDefault to vault snapshots labeled daily and weekly on an hourly schedule:

cluster2::> snapmirror create -source-path SVM1:srcvolA -destination
-path SVM2:dstvolB -vserver SVM2 -policy XDPDefault -schedule hourly



Create a custom replication policy or a custom schedule if the available defaults are not suitable.

4. On the destination SVM, initialize the SnapVault relationship created:

snapmirror initialize -destination-path <destination path>

The following command initializes the relationship between the source volume srcvolA on SVM1 and the destination volume dstvolB on SVM2:

cluster2::> snapmirror initialize -destination-path SVM2:dstvolB

5. After the relationship is initialized and idle, use the snapshot show command on the destination to verify the SnapLock expiry time applied to the replicated snapshots.

This example lists the snapshots on volume dstvolB that have the SnapMirror label and the SnapLock expiration date:

```
cluster2::> snapshot show -vserver SVM2 -volume dstvolB -fields
snapmirror-label, snaplock-expiry-time
```

Related information

- Cluster and SVM peering
- Volume backup using SnapVault
- snapmirror initialize

Mirror WORM files with ONTAP SnapMirror for disaster recovery

You can use SnapMirror to replicate WORM files to another geographic location for disaster recovery and other purposes. Both the source volume and destination volume must be configured for SnapLock, and both volumes must have the same SnapLock mode, Compliance or Enterprise. All key SnapLock properties of the volume and files are replicated.

Prerequisites

The source and destination volumes must be created in peered clusters with peered SVMs. For more information, see Cluster and SVM peering.

About this task

 Beginning with ONTAP 9.5, you can replicate WORM files with the XDP (extended data protection) type SnapMirror relationship rather than the DP (data protection) type relationship. XDP mode is ONTAP version-independent, and is able to differentiate files stored in the same block, making it much easier to resync replicated Compliance-mode volumes. For information on how to convert an existing DP-type relationship to an XDP-type relationship, see Data Protection.

- A resync operation on a DP type SnapMirror relationship fails for a Compliance-mode volume if SnapLock determines that it will result in a loss of data. If a resync operation fails, you can use the volume clone create command to make a clone of the destination volume. You can then resync the source volume with the clone.
- A SnapMirror relationship of a SnapLock volume only supports the MirrorAllSnapshots policy of type async-mirror. The retention period of a SnapLock volume is determined by the maximum retention period among all the WORM files that it holds. Because the destination is a DR copy of the source, the retention period of the destination SnapLock volume will be same as the source.
- A SnapMirror relationship of type XDP between SnapLock compliant volumes supports a resync after a break even if data on the destination has diverged from the source post the break.

On a resync, when data divergence is detected between the source the destination beyond the common snapshot, a new snapshot is cut on the destination to capture this divergence. The new snapshot and the common snapshot are both locked with a retention time as follows:

- The volume expiry time of the destination
- If the volume expiry time is in the past or has not been set, then the snapshot is locked for a period of 30 days
- If the destination has legal-holds, the actual volume expiry period is masked and shows up as 'indefinite'; however, the snapshot is locked for the duration of the actual volume expiry period.

If the destination volume has an expiry period that is later than the source, the destination expiry period is retained and will not be overwritten by the expiry period of the source volume post the resync.

If the destination has legal-holds placed on it that differ from the source, a resync is not allowed. The source and destination must have identical legal-holds or all legal-holds on the destination must be released before a resync is attempted.

A locked snapshot on the destination volume created to capture the divergent data can be copied to the source using the CLI by running the snapmirror update -s snapshot command. The snapshot once copied will continue to be locked at the source as well.

- SVM data protection relationships are not supported.
- Load-sharing data protection relationships are not supported.

The following illustration shows the procedure for initializing a SnapMirror relationship:

System Manager

Beginning with ONTAP 9.12.1, you can use System Manager to set up SnapMirror replication of WORM files.

Steps

- 1. Navigate to **Storage > Volumes**.
- 2. Click Show/Hide and select SnapLock Type to display the column in the Volumes window.
- 3. Locate a SnapLock volume.
- 4. Click and select **Protect**.
- 5. Choose the destination cluster and the destination storage VM.
- 6. Click More Options.
- 7. Select Show legacy policies and select DPDefault (legacy).
- 8. In the **Destination Configuration details** section, select **Override transfer schedule** and select **hourly**.
- 9. Click Save.
- 10. To the left of the source volume name, click the arrow to expand the volume details, and on the right side of the page, review the remote SnapMirror protection details.
- 11. On the remote cluster, navigate to **Protection Relationships**.
- 12. Locate the relationship and click the destination volume name to view the relationship details.
- 13. Verify that the destination volume SnapLock type and other SnapLock information.

CLI

- 1. Identify the destination cluster.
- 2. On the destination cluster, install the SnapLock license, initialize the Compliance Clock, and, if you are using an ONTAP release earlier than 9.10.1, create a SnapLock aggregate.
- 3. On the destination cluster, create a SnapLock destination volume of type DP that is either the same size as or greater in size than the source volume:

volume create -vserver SVM_name -volume volume_name -aggregate aggregate_name -snaplock-type compliance|enterprise -type DP -size size



Beginning with ONTAP 9.10.1, SnapLock and non-SnapLock volumes can exist on the same aggregate; therefore, you are no longer required to create a separate SnapLock aggregate if you are using ONTAP 9.10.1. You use the volume -snaplock-type option to specify a Compliance or Enterprise SnapLock volume type. In ONTAP releases earlier than ONTAP 9.10.1, the SnapLock mode—Compliance or Enterprise—is inherited from the aggregate. The language setting of the destination volume must match the language setting of the source volume.

The following command creates a 2 GB SnapLock Compliance volume named dstvolB in SVM2 on the aggregate node01_aggr:

```
cluster2::> volume create -vserver SVM2 -volume dstvolB -aggregate
node01_aggr -snaplock-type compliance -type DP -size 2GB
```

4. On the destination SVM, create a SnapMirror policy:

snapmirror policy create -vserver SVM name -policy policy name

The following command creates the SVM-wide policy SVM1-mirror:

SVM2::> snapmirror policy create -vserver SVM2 -policy SVM1-mirror

5. On the destination SVM, create a SnapMirror schedule:

job schedule cron create -name *schedule_name* -dayofweek *day_of_week* -hour *hour* -minute *minute*

The following command creates a SnapMirror schedule named weekendcron:

SVM2::> job schedule cron create -name weekendcron -dayofweek
"Saturday, Sunday" -hour 3 -minute 0

6. On the destination SVM, create a SnapMirror relationship:

```
snapmirror create -source-path source_path -destination-path
destination path -type XDP|DP -policy policy name -schedule schedule name
```

The following command creates a SnapMirror relationship between the source volume srcvolA on SVM1 and the destination volume dstvolB on SVM2, and assigns the policy SVM1-mirror and the schedule weekendcron:

```
SVM2::> snapmirror create -source-path SVM1:srcvolA -destination
-path SVM2:dstvolB -type XDP -policy SVM1-mirror -schedule
weekendcron
```



The XDP type is available in ONTAP 9.5 and later. You must use the DP type in ONTAP 9.4 and earlier.

7. On the destination SVM, initialize the SnapMirror relationship:

```
snapmirror initialize -destination-path destination path
```

The initialization process performs a *baseline transfer* to the destination volume. SnapMirror makes a snapshot of the source volume, then transfers the copy and all the data blocks that it references to the destination volume. It also transfers any other snapshots on the source volume to the destination volume.

The following command initializes the relationship between the source volume srcvolA on SVM1 and the destination volume dstvolB on SVM2:

SVM2::> snapmirror initialize -destination-path SVM2:dstvolB

Related information

- Cluster and SVM peering
- Volume disaster recovery preparation
- Data protection
- snapmirror create
- snapmirror initialize
- snapmirror policy create

Retain WORM files during litigation using ONTAP SnapLock Legal Hold

Beginning with ONTAP 9.3, you can retain Compliance-mode WORM files for the duration of a litigation by using the *Legal Hold* feature.

Before you begin

• You must be a SnapLock administrator to perform this task.

Create a SnapLock administrator account

• You must have logged in on a secure connection (SSH, console, or ZAPI).

About this task

A file under a Legal Hold behaves like a WORM file with an indefinite retention period. It is your responsibility to specify when the Legal Hold period ends.

The number of files you can place under a Legal Hold depends on the space available on the volume.

Steps

1. Start a Legal Hold:

```
snaplock legal-hold begin -litigation-name <litigation_name> -volume
<volume name> -path <path name>
```

The following command starts a Legal Hold for all the files in vol1:

```
cluster1::>snaplock legal-hold begin -litigation-name litigation1
-volume vol1 -path /
```

2. End a Legal Hold:

snaplock legal-hold end -litigation-name <litigation_name> -volume

```
<volume_name> -path <path_name>
```

The following command ends a Legal Hold for all the files in vol1:

```
cluster1::>snaplock legal-hold end -litigation-name litigation1 -volume
vol1 -path /
```

Related information

- snaplock legal-hold begin
- snaplock legal-hold end

Delete WORM files with ONTAP SnapLock

You can delete Enterprise-mode WORM files during the retention period using the privileged delete feature.

Before you can use this feature, you must create a SnapLock administrator account and then using the account, enable the feature.

Create a SnapLock administrator account

You must have SnapLock administrator privileges to perform a privileged delete. These privileges are defined in the vsadmin-snaplock role. If you have not already been assigned that role, you can ask your cluster administrator to create an SVM administrator account with the SnapLock administrator role.

Before you begin

- You must be a cluster administrator to perform this task.
- You must have logged in on a secure connection (SSH, console, or ZAPI).

Steps

1. Create an SVM administrator account with the SnapLock administrator role:

security login create -vserver SVM_name -user-or-group-name user_or_group_name -application application -authmethod authentication_method -role role -comment comment

The following command enables the SVM administrator account SnapLockAdmin with the predefined vsadmin-snaplock role to access SVM1 using a password:

```
cluster1::> security login create -vserver SVM1 -user-or-group-name
SnapLockAdmin -application ssh -authmethod password -role vsadmin-
snaplock
```

Learn more about security login create in the ONTAP command reference.

Enable the privileged delete feature

You must explicitly enable the privileged delete feature on the Enterprise volume that contains the WORM files you want to delete.

About this task

The value of the -privileged-delete option determines whether privileged delete is enabled. Possible values are enabled, disabled, and permanently-disabled.



permanently-disabled is the terminal state. You cannot enable privileged delete on the volume after you set the state to permanently-disabled.

Steps

1. Enable privileged delete for a SnapLock Enterprise volume:

volume snaplock modify -vserver SVM_name -volume volume_name -privileged -delete disabled|enabled|permanently-disabled

The following command enables the privileged delete feature for the Enterprise volume dataVol on SVM1:

```
SVM1::> volume snaplock modify -vserver SVM1 -volume dataVol -privileged
-delete enabled
```

Delete Enterprise-mode WORM files

You can use the privileged delete feature to delete Enterprise-mode WORM files during the retention period.

Before you begin

- You must be a SnapLock administrator to perform this task.
- You must have created a SnapLock audit log and enabled the privileged delete feature on the Enterprise volume.

About this task

You cannot use a privileged delete operation to delete an expired WORM file. You can use the volume file retention show command to view the retention time of the WORM file that you want to delete. Learn more about volume file retention show in the ONTAP command reference.

Step

1. Delete a WORM file on an Enterprise volume:

volume file privileged-delete -vserver SVM_name -file file_path

The following command deletes the file /vol/dataVol/f1 on the SVMsVM1:

SVM1::> volume file privileged-delete -file /vol/dataVol/f1

Move an ONTAP SnapLock volume

Beginning with ONTAP 9.8, you can move a SnapLock volume to a destination aggregate of the same type, either Enterprise to Enterprise, or Compliance to Compliance. You must be assigned the SnapLock security role to move a SnapLock volume.

Create a SnapLock security administrator account

You must have SnapLock security administrator privileges to perform a SnapLock volume move. This privilege is granted to you with the *snaplock* role, introduced in ONTAP 9.8. If you have not already been assigned that role, you can ask your cluster administrator to create a SnapLock security user with this SnapLock security role.

Before you begin

- You must be a cluster administrator to perform this task.
- · You must have logged in on a secure connection (SSH, console, or ZAPI).

About this task

The snaplock role is associated with the admin SVM, unlike the vsadmin-snaplock role, which is associated with the data SVM.

Step

1. Create an SVM administrator account with the SnapLock administrator role:

security login create -vserver SVM_name -user-or-group-name user_or_group_name -application application -authmethod authentication_method -role role -comment comment

The following command enables the SVM administrator account SnapLockAdmin with the predefined snaplock role to access admin SVM cluster1 using a password:

```
cluster1::> security login create -vserver cluster1 -user-or-group-name
SnapLockAdmin -application ssh -authmethod password -role snaplock
```

Learn more about security login create in the ONTAP command reference.

Move a SnapLock volume

You can use the volume move command to move a SnapLock volume to a destination aggregate.

Before you begin

• You must have created a SnapLock-protected audit log before performing SnapLock volume move.

Create an audit log.

• If you are using a version of ONTAP earlier than ONTAP 9.10.1, the destination aggregate must be the same SnapLock type as the SnapLock volume you want to move; either Compliance to Compliance or Enterprise to Enterprise. Beginning with ONTAP 9.10.1, this restriction is removed and an aggregate can include both Compliance and Enterprise SnapLock volumes, as well as non-SnapLock volumes.

• You must be a user with the SnapLock security role.

Steps

1. Using a secure connection, log in to the ONTAP cluster management LIF:

ssh snaplock_user@cluster_mgmt_ip

2. Move a SnapLock volume:

volume move start -vserver SVM_name -volume SnapLock_volume_name -destination
-aggregate destination_aggregate_name

3. Check the status of the volume move operation:

volume move show -volume SnapLock_volume_name -vserver SVM_name -fields
volume,phase,vserver

Lock an ONTAP snapshot for protection against ransomware attacks

Beginning with ONTAP 9.12.1, you can lock a snapshot on a non-SnapLock volume to provide protection from ransomware attacks. Locking snapshots ensures that they can't be deleted accidentally or maliciously.

You use the SnapLock compliance clock feature to lock snapshots for a specified period so that they cannot be deleted until the expiration time is reached. Locking snapshots makes them tamperproof, protecting them from ransomware threats. You can use locked snapshots to recover data if a volume is compromised by a ransomware attack.

Beginning with ONTAP 9.14.1, snapshot locking supports long-term retention snapshots on SnapLock vault destinations and on non-SnapLock SnapMirror destination volumes. Snapshot locking is enabled by setting the retention period using SnapMirror policy rules associated with an existing policy label. The rule overrides the default retention period set on the volume. If there is no retention period associated with the SnapMirror label, the default retention period of the volume is used.

Tamperproof snapshot requirements and considerations

- If you are using the ONTAP CLI, all nodes in the cluster must be running ONTAP 9.12.1 or later. If you are using System Manager, all nodes must be running ONTAP 9.13.1 or later.
- The SnapLock license must be installed on the cluster. This license is included in ONTAP One.
- The compliance clock on the cluster must be initialized.
- When snapshot locking is enabled on a volume, you can upgrade the clusters to a version of ONTAP later than ONTAP 9.12.1; however, you cannot revert to an earlier version of ONTAP until all locked snapshots have reached their expiration date and are deleted and snapshot locking is disabled.
- When a snapshot is locked, the volume expiry time is set to the expiry time of the snapshot. If more than one snapshot is locked, the volume expiry time reflects the largest expiry time among all snapshots.
- The retention period for locked snapshots takes precedence over the snapshot keep count, which means the keep count limit is not honored if the snapshot retention period for locked snapshots has not expired.
- In a SnapMirror relationship, you can set a retention period on a mirror-vault policy rule, and the retention period is applied for snapshots replicated to the destination if the destination volume has snapshot locking enabled. The retention period takes precedence over keep count; for example, snapshots that have not passed their expiry will be retained even if the keep count is exceeded.

- You can rename a snapshot on a non-SnapLock volume. Snapshot rename operations on the primary volume of a SnapMirror relationship are reflected on the secondary volume only if the policy is MirrorAllSnapshots. For other policy types, the renamed snapshot is not propagated during updates.
- If you are using the ONTAP CLI, you can restore a locked snapshot with the volume snapshot restore command only if the locked snapshot is the most recent. If there are any unexpired snapshots later than the one being restored, the snapshot restore operation fails.

Features supported with tamperproof snapshots

- Cloud Volumes ONTAP
- FlexGroup volumes

Snapshot locking is supported on FlexGroup volumes. Snapshot locking occurs only on the root constituent snapshot. Deleting the FlexGroup volume is allowed only if the root constituent expiration time has passed.

• FlexVol to FlexGroup conversion

You can convert a FlexVol volume with locked snapshots to a FlexGroup volume. Snapshots remain locked after the conversion.

SnapMirror asynchronous

The compliance clock must be initialized on both the source and destination.

• SVM data mobility (used for migrating or relocating an SVM from a source cluster to a destination cluster)

Supported beginning with ONTAP 9.14.1.

- SnapMirror policy rules using the -schedule parameter
- SVM DR

The compliance clock must be initialized on both the source and destination.

• Volume clone and file clone

You can create volume clones and file clones from a locked snapshot.

Unsupported features

The following features currently are not supported with tamperproof snapshots:

- Consistency groups
- FabricPool

Tamperproof snapshots provide immutable protections that cannot be deleted. Because FabricPool requires the ability to delete data, FabricPool and snapshot locks cannot be enabled on the same volume.

- FlexCache volumes
- SMtape
- SnapMirror active sync
- SnapMirror synchronous

Enable snapshot locking when creating a volume

Beginning with ONTAP 9.12.1, you can enable snapshot locking when you create a new volume or when you modify an existing volume by using the <code>-snapshot-locking-enabled</code> option with the volume <code>create</code> and <code>volume modify</code> commands in the CLI. Beginning with ONTAP 9.13.1, you can use System Manager to enable snapshot locking.

System Manager

- 1. Navigate to **Storage > Volumes** and select **Add**.
- 2. In the Add Volume window, choose More Options.
- 3. Enter the volume name, size, export policy and share name.
- 4. Select **Enable Snapshot locking**. This selection is not displayed if the SnapLock license is not installed.
- 5. If it is not already enabled, select **Initialize SnapLock Compliance Clock**.
- 6. Save your changes.
- 7. In the Volumes window, select the volume you updated and choose Overview.
- 8. Verify that SnapLock Snapshot Locking displays as Enabled.

CLI

1. To create a new volume and enable snapshot locking, enter the following command:

```
volume create -vserver <vserver_name> -volume <volume_name> -snapshot
-locking-enabled true
```

The following command enables snapshot locking on a new volume named vol1:

```
> volume create -volume voll -aggregate aggr1 -size 100m -snapshot
-locking-enabled true
Warning: snapshot locking is being enabled on volume "voll" in
Vserver "vs1". It cannot be disabled until all locked snapshots are
past their expiry time. A volume with unexpired locked snapshots
cannot be deleted.
Do you want to continue: {yes|no}: y
[Job 32] Job succeeded: Successful
```

Enable snapshot locking on an existing volume

Beginning with ONTAP 9.12.1, you can enable snapshot locking on an existing volume using the ONTAP CLI. Beginning with ONTAP 9.13.1, you can use System Manager to enable snapshot locking on an existing volume.

System Manager

- 1. Navigate to **Storage > Volumes**.
- 2. Select and choose **Edit > Volume**.
- 3. In the **Edit Volume** window, locate the Snapshots (Local) Settings section and select **Enable snapshot locking**.

This selection is not displayed if the SnapLock license is not installed.

- 4. If it is not already enabled, select **Initialize SnapLock Compliance Clock**.
- 5. Save your changes.
- 6. In the **Volumes** window, select the volume you updated and choose **Overview**.
- 7. Verify that **SnapLock snapshot locking** displays as **Enabled**.

CLI

1. To modify an existing volume to enable snapshot locking, enter the following command:

```
volume modify -vserver <vserver_name> -volume <volume_name> -snapshot
-locking-enabled true
```

Create a locked snapshot policy and apply retention

Beginning with ONTAP 9.12.1, you can create snapshot policies to apply a snapshot retention period and apply the policy to a volume to lock snapshots for the specified period. You can also lock a snapshot by manually setting a retention period. Beginning with ONTAP 9.13.1, you can use System Manager to create snapshot locking policies and apply them to a volume.

Create a snapshot locking policy

System Manager

- 1. Navigate to Storage > Storage VMs and select a storage VM.
- 2. Select Settings.
- 3. Locate Snapshot Policies and select ->.
- 4. In the Add Snapshot Policy window, enter the policy name.
- 5. Select + Add.
- 6. Provide the snapshot schedule details, including the schedule name, maximum snapshots to keep, and SnapLock retention period.
- 7. In the **SnapLock Retention Period** column, enter the number of hours, days, months or years to retain the snapshots. For example, a snapshot policy with a retention period of 5 days locks a snapshot for 5 days from the time it is created, and it cannot be deleted during that time. The following retention period ranges are supported:
 - · Years: 0 100
 - Months: 0 1200
 - Days: 0 36500
 - Hours: 0 24
- 8. Save your changes.

CLI

1. To create a snapshot policy, enter the following command:

```
volume snapshot policy create -policy <policy_name> -enabled true
-schedule1 <schedule1_name> -count1 <maximum snapshots> -retention-period1
<retention period>
```

The following command creates a snapshot locking policy:

cluster1> volume snapshot policy create -policy lock_policy -enabled true -schedule1 hourly -count1 24 -retention-period1 "1 days"

A snapshot is not replaced if it is under active retention; that is, the retention count will not be honored if there are locked snapshots that have not yet expired.

Apply a locking policy to a volume
- 1. Navigate to **Storage > Volumes**.
- 2. Select and choose **Edit > Volume**.
- 3. In the Edit Volume window, select Schedule snapshots.
- 4. Select the locking snapshot policy from the list.
- 5. If snapshot locking is not already enabled, select **Enable snapshot locking**.
- 6. Save your changes.

CLI

1. To apply a snapshot locking policy to an existing volume, enter the following command:

```
volume modify -volume <volume_name> -vserver <vserver_name> -snapshot
-policy <policy name>
```

Apply retention period during manual snapshot creation

You can apply a snapshot retention period when you manually create a snapshot. Snapshot locking must be enabled on the volume; otherwise, the retention period setting is ignored.

System Manager

- 1. Navigate to **Storage > Volumes** and select a volume.
- 2. In the volume details page, select the Snapshots tab.
- 3. Select + Add.
- 4. Enter the snapshot name and the SnapLock expiration time. You can select the calendar to choose the retention expiration date and time.
- 5. Save your changes.
- In the Volumes > Snapshots page, select Show/Hide and choose SnapLock Expiration Time to display the SnapLock Expiration Time column and verify that the retention time is set.

CLI

1. To create a snapshot manually and apply a locking retention period, enter the following command:

```
volume snapshot create -volume <volume_name> -snapshot <snapshot name>
-snaplock-expiry-time <expiration date time>
```

The following command creates a new snapshot and sets the retention period:

cluster1> volume snapshot create -vserver vs1 -volume vol1 -snapshot snap1 -snaplock-expiry-time "11/10/2022 09:00:00"

- 1. Navigate to **Storage > Volumes** and select a volume.
- 2. In the volume details page, select the **Snapshots** tab.
- 3. Select the snapshot, select ; and choose **Modify SnapLock Expiration Time**. You can select the calendar to choose the retention expiration date and time.
- 4. Save your changes.
- 5. In the Volumes > Snapshots page, select Show/Hide and choose SnapLock Expiration Time to display the SnapLock Expiration Time column and verify that the retention time is set.

CLI

1. To manually apply a retention period to an existing snapshot, enter the following command:

```
volume snapshot modify-snaplock-expiry-time -volume <volume_name> -snapshot
<snapshot name> -snaplock-expiry-time <expiration date time>
```

The following example applies a retention period to an existing snapshot:

```
cluster1> volume snapshot modify-snaplock-expiry-time -volume vol1
-snapshot snap2 -snaplock-expiry-time "11/10/2022 09:00:00"
```

Modify an existing policy to apply long-term retention

In a SnapMirror relationship, you can set a retention period on a mirror-vault policy rule, and the retention period is applied for snapshots replicated to the destination if the destination volume has snapshot locking enabled. The retention period takes precedence over keep count; for example, snapshots that have not passed their expiry will be retained even if the keep count is exceeded.

Beginning with ONTAP 9.14.1, you can modify an existing SnapMirror policy by adding a rule to set long-term retention of snapshots. The rule is used to override the default volume retention period on SnapLock vault destinations and on non-SnapLock SnapMirror destination volumes.

1. Add a rule to an existing SnapMirror policy:

```
snapmirror policy add-rule -vserver <SVM name> -policy <policy name>
-snapmirror-label <label name> -keep <number of snapshots> -retention-period
[<integer> days|months|years]
```

The following example creates a rule that applies a retention period of 6 months to the existing policy called "lockvault":

```
snapmirror policy add-rule -vserver vs1 -policy lockvault -snapmirror
-label test1 -keep 10 -retention-period "6 months"
```

Learn more about snapmirror policy add-rule in the ONTAP command reference.

Consistency groups

Learn about ONTAP consistency groups

A consistency group is a collection of volumes that are managed as a single unit. In ONTAP, consistency groups provide easy management and a protection guarantee for an application workload spanning multiple volumes.

You can use consistency groups to simplify your storage management. Imagine you have an important database spanning twenty LUNs. You could manage the LUNs on an individual basis or treat the LUNs as a solitary dataset, organizing them into a single consistency group.

Consistency groups facilitate application workload management, providing easily configured local and remote protection policies and simultaneous crash-consistent or application-consistent snapshots of a collection of volumes at a point in time. Snapshots of a consistency groups enable an entire application workload to be restored.

Learn about consistency groups

Consistency groups support any FlexVol volume regardless of protocol (NAS, SAN, or NVMe) and can be managed through the ONTAP REST API or in System Manager under the **Storage > Consistency Groups** menu item. Beginning with ONTAP 9.14.1, consistency groups can be managed with the ONTAP CLI.

Consistency groups can exist as individual entities—as a collection of volumes—or in a hierarchical relationship, which consists of other consistency groups. Individual volumes can have their own volume-granular snapshot policy. In addition, there can be a consistency group-wide snapshot policy. The consistency group can only have one SnapMirror active sync relationship and shared SnapMirror policy, which can be used to recover the entire consistency group.

The following diagram illustrates how you might use an individual consistency group. The data for an application hosted on SVM1 spans two volumes: vol1 and vol2. A snapshot policy on the consistency group captures snapshots of the data every 15 minutes.



Larger application workloads might require multiple consistency groups. In these situations, you can create

hierarchical consistency groups, where a single consistency group becomes the child components of a parent consistency group. The parent consistency group can include up to five child consistency groups. Like in individual consistency groups, a remote SnapMirror active sync protection policy can be applied to the entire configuration of consistency groups (parent and children) to recover the application workload.

In the following example, an application is hosted on SVM1. The administrator has created a parent consistency group, SVM1_app, which includes two child consistency groups: SVM1appDataCG for the data and SVM1app_logCG for the logs. Each child consistency group has its own snapshot policy. Snapshots of the volumes in SVM1appDataCG are taken every 15 minutes. Snapshots of SVM1app_logCG are taken hourly. The parent consistency group SVM1_app has an SnapMirror active sync policy which replicates the data to ensure continued service in the event of a disaster.



Beginning with ONTAP 9.12.1, consistency groups support cloning and modifying the members of the consistency by adding or removing volumes in both System Manager and the ONTAP REST API. Beginning with ONTAP 9.12.1, the ONTAP REST API also supports:

- Creating consistency groups with new NFS or SMB volumes or NVMe namespaces.
- Adding new or existing NFS or SMB volumes or NVMe namespaces to existing consistency groups.

For more information about the ONTAP REST API, refer to ONTAP REST API reference documentation.

Monitor consistency groups

Beginning with ONTAP 9.13.1, consistency groups offer real-time and historical capacity and performance monitoring, offering insights about the performance of applications and individual consistency groups.

Monitoring data is refreshed every five minutes and is maintained for up to one year. You can track metrics for:

- Performance: IOPS, latency, and throughput
- Capacity: Size, logical used, available

You can view monitoring data in the **Overview** tab of the consistency group menu in System Manager or by requesting it in the REST API. Beginning with ONTAP 9.14.1, you can view consistency group metrics with the CLI using the consistency-group metrics show command. Learn more about consistency-group metrics show in the ONTAP command reference.



In ONTAP 9.13.1, you can only retrieve historical metrics using the REST API. Beginning with ONTAP 9.14.1, historical metrics are also available in System Manager.

Protect consistency groups

Consistency groups offer application-consistent protection, ensuring consistency of your data across multiple volumes or LIFs. When creating a snapshot of a consistency group, a "fence" is established on the consistency group. The fence initiates a queue for I/O until after the snapshot operation completes, ensuring point-in-time consistency of data across all entities in the consistency group. The fence can cause a transient spike in latency during snapshot creation operations, such as a scheduled snapshot policy or creating a snapshot with System Manager. Learn more about the context of REST API and CLI in the ONTAP REST API documentation and ONTAP command reference.

Consistency groups offer protection through:

- Snapshot policies
- SnapMirror active sync
- MetroCluster support (beginning with ONTAP 9.11.1)
- SnapMirror asynchronous (beginning with ONTAP 9.13.1)
- SVM disaster recovery (beginning with ONTAP 9.14.1)

Creating a consistency group does not automatically enable protection. Local and remote protection policies can be set when creating or after creating a consistency group.

To configure protection on a consistency group, see Protect a consistency group.

In order to use remote protection, you must meet the requirements for SnapMirror active sync.



SnapMirror active sync relationships cannot be established on volumes mounted for NAS access.

Multi-admin verification support for consistency groups

Beginning with ONTAP 9.16.1, you can use multi-admin verification (MAV) with consistency groups to ensure that certain operations, such as creating, modifying, or deleting consistency groups, can be executed only after approvals from designated administrators. This prevents compromised, malicious, or inexperienced administrators from making undesirable changes to existing configurations.

Learn more

Consistency groups in MetroCluster configurations

Beginning with ONTAP 9.11.1, you can provision consistency groups with new volumes on a cluster within a

MetroCluster configuration. These volumes are provisioned on mirrored aggregates.

After they are provisioned, you can move volumes associated with consistency groups between mirrored and unmirrored aggregates. Therefore, volumes associated with consistency groups can be located on mirrored aggregates, unmirrored aggregates, or both. You can modify mirrored aggregates containing volumes associated with consistency groups to become unmirrored. Similarly, you can modify unmirrored aggregates containing volumes associated with consistency groups to enable mirroring.

Volumes and snapshots associated with consistency groups placed on mirrored aggregates are replicated to the remote site (site B). The contents of the volumes on site B provide a write-order guarantee for the consistency group, allowing you to recover from site B in the event of a disaster. You can access consistency group snapshots using consistency group with the REST API and System Manager on clusters running ONTAP 9.11.1 or later. Beginning with ONTAP 9.14.1, you can also access snapshots with the ONTAP CLI.

If some or all the volumes associated with a consistency group are located on unmirrored aggregates that are not currently accessible, GET or DELETE operations on the consistency group behave as if the local volumes or hosting aggregates are offline.

Consistency group configurations for replication

If site B is running ONTAP 9.10.1 or earlier, only the volumes associated with the consistency groups located on mirrored aggregates are replicated to site B. The consistency group configurations are only replicated to site B, if both sites are running ONTAP 9.11.1 or later. After site B is upgraded to ONTAP 9.11.1, data for consistency groups on site A that have all their associated volumes placed on mirrored aggregates are replicated to site B.



It's recommended you maintain at least 20% free space for mirrored aggregates for optimal storage performance and availability. Although the recommendation is 10% for non-mirrored aggregates, the additional 10% of space may be used by the filesystem to absorb incremental changes. Incremental changes increase space utilization for mirrored aggregates due to ONTAP's copy-on-write snapshot-based architecture. Failure to adhere to these best practices may have a negative impact on performance.

Upgrade considerations

When upgrading to ONTAP 9.10.1 or later, consistency groups created with SnapMirror active sync (previously known as SnapMirror Business Continuity) in ONTAP 9.8 and 9.9.1 are automatically upgraded and become manageable under **Storage > Consistency Groups** in System Manager or the ONTAP REST API For more information about upgrading from ONTAP 9.8 or 9.9.1, see SnapMirror active sync upgrade and revert considerations.

Consistency group snapshots created in the REST API can be managed through System Manager's Consistency Group interface and through consistency group REST API endpoints. Beginning with ONTAP 9.14.1, consistency group snapshots can also be managed with the ONTAP CLI.



Snapshots created with the ONTAPI commands cg-start and cg-commit are not recognized as consistency group snapshots and thus cannot be managed through System Manager's consistency group interface or the consistency group endpoints in the ONTAP REST API. Beginning with ONTAP 9.14.1, these snapshots can be mirrored to the destination volume if you are using a SnapMirror asynchronous policy. For more information, see Configure SnapMirror asynchronous.

Supported features by release

	ONTAP 9.16.1	ONTAP 9.15.1	ONTAP 9.14.1	ONTAP 9.13.1	ONTAP 9.12.1	ONTAP 9.11.1	ONTAP 9.10.1
Hierarchical consistency groups	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
Local protection with snapshots	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
SnapMirror active sync	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
MetroCluster support	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	
Two-phase commits (REST API only)	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	
Application and component tags	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark		
Clone consistency groups	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark		
Add and remove volumes	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark		
Create CGs with new NAS volumes	\checkmark	\checkmark	\checkmark	\checkmark	REST API only		
Create CGs with new NVMe Namespaces	\checkmark	\checkmark	\checkmark	\checkmark	REST API only		
Move volumes between child consistency groups	\checkmark	\checkmark	\checkmark	\checkmark			
Modify consistency group geometry	\checkmark	\checkmark	\checkmark	\checkmark			
Monitoring	\checkmark	\checkmark	\checkmark	\checkmark			
Multi-admin verification	\checkmark						
SnapMirror asynchronous (single consistency groups only)	\checkmark	\checkmark	\checkmark	\checkmark			
SVM disaster recovery (single consistency groups only)	\checkmark	\checkmark	\checkmark				
CLI support	\checkmark	\checkmark	\checkmark				

Learn more about consistency groups

Consistency Groups for Application Management & Protection

With NetApp ONTAP 9.10.1 + System Manager





© 2022 NetApp, Inc. All rights reserved.

Related information

- ONTAP automation documentation
- SnapMirror active sync
- SnapMirror asynchronous disaster recovery basics
- MetroCluster documentation
- Multi-admin verification
- ONTAP command reference

Learn about ONTAP consistency group limits

When planning and managing your consistency groups, account for object limits at the scope of both the cluster and the parent or child consistency group.

Enforced limits

The following table captures limits for consistency groups. Separate limits apply for consistency groups using SnapMirror active sync. For more information, see SnapMirror active sync limits.

Limit	Scope	Minimum	Maximum
Number of consistency groups	Cluster	0	Same as maximum volume count in cluster*
Number of parent consistency groups	Cluster	0	Same as maximum volume count in cluster
Number of individual and parent consistency groups	Cluster	0	Same as maximum volume count in cluster

Number of volumes in a consistency group	Single consistency group	1 volume	80 volumes	
Number of volumes in a consistency group with SnapMirror asynchronous	Single consistency group	1 volume	 In ONTAP 9.15.1 and later: 80 volumes In ONTAP 9.13.1 and 9.14.1: 16 volumes 	
Number of volumes in the child of a parent consistency group	Parent consistency group	1 volume	80 volumes	
Number of volumes in a child consistency group	Child consistency group	1 volume	80 volumes	
Number of child consistency groups in a parent consistency group	Parent consistency group	1 consistency group	5 consistency groups	
Number of SVM disaster recovery relationships where a consistency group exists (available beginning ONTAP 9.14.1)	ber of SVM disaster overy relationships re a consistency up exists (available inning ONTAP 9.14.1)		32	

* A maximum of 50 consistency groups enabled with SnapMirror asynchronous can be hosted on a cluster.

Unenforced limits

The minimum supported snapshot schedule for consistency groups is 30 minutes. This is based on testing for FlexGroup volumes, which share the same Snapshot infrastructure as consistency groups.

Configure a single ONTAP consistency group

Consistency groups can be created with existing volumes or new LUNs or volumes (depending on the version of ONTAP). A volume or LUN can only be associated with one consistency group at a time.

About this task

• In ONTAP 9.10.1 through 9.11.1, modifying the member volumes of a consistency group after it is created is not supported.

Beginning with ONTAP 9.12.1, you can modify the member volumes of a consistency group. For more information on this process, refer to Modify a consistency group.

• Beginning with ONTAP 9.17.1, you can select the NVMe protocol to map a host to an NVMe subsystem for VMware workloads in a SnapMirror active sync configuration.

Create a consistency group with new LUNs or volumes

In ONTAP 9.10.1 through 9.12.1, you can create a consistency group using new LUNs. Beginning with ONTAP 9.13.1, System Manager also supports creating a consistency group with new NVMe namespaces or new NAS volumes. (This is also supported in the ONTAP REST API beginning with ONTAP 9.12.1.)

System Manager (ONTAP 9.16.1 and earlier)

Steps

- 1. Select Storage > Consistency groups.
- 2. Select +Add then select the protocol for your storage object.

In ONTAP 9.10.1 through 9.12.1, the only option for a new storage object is **Using new LUNs**. Beginning with ONTAP 9.13.1, System Manager supports creating consistency groups with new NVMe namespaces and new NAS volumes.

- 3. Name the consistency group. Designate the number of volumes or LUNs and the capacity per volume or LUN.
 - a. **Application Type**: If you are using ONTAP 9.12.1 or later, select an application type. If no value is selected, the consistency group will be assigned the type of **Other** by default. Learn more about tagging consistency in Application and component tags. If you plan to create a consistency group with a remote protection policy, you must use **Other**.
 - b. For **New LUNs**: Select the host operating system and LUN format. Enter the host initiator information.
 - c. For **New NAS volumes**: choose the appropriate export option (NFS or SMB/CIFS) based on the NAS configuration of your SVM.
 - d. For **New NVMe namespaces**: Select the host operating system and NVMe subsystem.
- 4. To configure protection policies, add a child consistency group, or access permissions, select **More options**.
- 5. Select Save.
- 6. Confirm your consistency group has been created by returning to the main consistency group menu where it will appear once the job completes. If you set a protection policy, you will know it has been applied when you see a green shield under look under the appropriate policy, remote or local.

System Manager (ONTAP 9.17.1 and later)

Steps

- 1. Select **Protection > Consistency groups**.
- 2. Select +Add then select the protocol for your storage object.
- 3. Name the consistency group. Designate the number of volumes or LUNs and the capacity per volume or LUN.

Application Type: Select an application type. If no value is selected, the consistency group will be assigned the type of **Other** by default. Learn more about tagging consistency in Application and component tags. If you plan to create a consistency group with a remote protection policy, you must use **Other**.

- a. For **New LUNs**: Select the host operating system and LUN format. Enter the host initiator information.
- b. For **New NAS volumes**: choose the appropriate export option (NFS or SMB/CIFS) based on the NAS configuration of your SVM.
- c. For New NVMe namespaces: Select the host operating system and NVMe subsystem.
- 4. To configure protection policies, add a child consistency group, or access permissions, select **More options**.
- 5. Select Save.

6. Confirm your consistency group has been created by returning to the main consistency group menu where it will appear once the job completes. If you set a protection policy, you will know it has been applied when you see a green shield under look under the appropriate policy, remote or local.

CLI

Beginning with ONTAP 9.14.1, you can create a new consistency group with new volumes using the ONTAP CLI. The specific parameters depends on whether the volumes are SAN, NVMe, or NFS.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Create a consistency group with NFS volumes

1. Create the consistency group:

```
consistency-group create -vserver <SVM_name> -consistency-group
<consistency-group-name> -volume-prefix <prefix_for_new_volume_names>
-volume-count <number> -size <size> -export-policy <policy name>
```

Create a consistency group with SAN volumes

1. Create the consistency group:

```
consistency-group create -vserver <SVM_name> -consistency-group
<consistency-group-name> -lun <lun_name> -size <size> -lun-count <number>
-lun-os-type <LUN_operating_system_format> -igroup <igroup_name>
```

Create a consistency group with NVMe namespaces

1. Create the consistency group:

```
consistency-group create -vserver <SVM_name> -consistency-group
<consistency_group_name> -namespace <namespace_name> -volume-count <number>
-namespace-count <number> -size <size> -subsystem <subsystem_name>
```

Learn more about consistency-group create in the ONTAP command reference.

After you're done

1. Confirm your consistency group has been created using the consistency-group show command.

Learn more about consistency-group show in the ONTAP command reference.

Create a consistency group with existing volumes

You can use existing volumes to create a consistency group.

System Manager (ONTAP 9.16.1 and earlier)

Steps

- 1. Select Storage > Consistency groups.
- 2. Select +Add then Using existing volumes.
- 3. Name the consistency group and select the storage VM.
 - a. **Application Type**: If you are using ONTAP 9.12.1 or later, select an application type. If no value is selected, the consistency group will be assigned the type of **Other** by default. Learn more about tagging consistency in Application and component tags. If the consistency group has a SnapMirror active sync relationship, you must use **Other**.



In versions of ONTAP earlier than ONTAP 9.15.1, SnapMirror active sync is referred to as SnapMirror Business Continuity.

4. Select the existing volumes to include. Only volumes that are not already part of a consistency group will be available for selection.



If creating a consistency group with existing volumes, the consistency group supports FlexVol volumes. Volumes with or SnapMirror synchronous or SnapMirror asynchronous relationships can be added to consistency groups, but they are not consistency group-aware. Consistency groups do not support S3 buckets or storage VMs with SVMDR relationships.

- 5. Select Save.
- 6. Confirm your consistency group has been created by returning to the main consistency group menu where it appears once the ONTAP job completes. If you have chosen a protection policy, confirm it was properly set by selecting your consistency group from the menu. If you set a protection policy, you know it has been applied when you see a green shield under look under the appropriate policy, remote or local.

CLI

Beginning with ONTAP 9.14.1, you can create a consistency group with existing volumes using the ONTAP CLI.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Steps

1. Issue the consistency-group create command. The -volumes parameter accepts a commaseparated list of volume names.

```
consistency-group create -vserver <SVM_name> -consistency-group
<consistency-group-name> -volume <volumes>
```

Learn more about consistency-group create in the ONTAP command reference.

2. View your consistency group using the consistency-group show command.

Learn more about consistency-group show in the ONTAP command reference.

Next steps

- Protect a consistency group
- Modify a consistency group
- Clone a consistency group

Configure a hierarchical ONTAP consistency group

Hierarchical consistency groups enable you to manage large workloads spanning multiple volumes, creating a parent consistency group that serves as an umbrella for child consistency groups.

Hierarchical consistency groups have a parent that can include up to five individual consistency groups. Hierarchical consistency groups can support different local snapshot policies across consistency groups or individual volumes. If you use a remote protection policy, that will apply for the entire hierarchical consistency group (parent and children).

Beginning with ONTAP 9.13.1, you can modify the geometry of your consistency groups and move volumes between child consistency groups.

For object limits on consistency groups, see Object limits for consistency groups.

Create a hierarchical consistency group with new LUNs or volumes

When creating a hierarchical consistency group, you can populate it with new LUNs. Beginning with ONTAP 9.13.1, you can also use new NVMe namespaces and NAS volumes.

Steps

- 1. Select Storage > Consistency groups.
- 2. Select +Add then select the protocol for your storage object.

In ONTAP 9.10.1 through 9.12.1, the only option for a new storage object is **Using new LUNs**. Beginning with ONTAP 9.13.1, System Manager supports creating consistency groups with new NVMe namespaces and new NAS volumes.

- 3. Name the consistency group. Designate the number of volumes or LUNs and the capacity per volume or LUN.
 - a. **Application Type**: If you are using ONTAP 9.12.1 or later, select an application type. If no value is selected, the consistency group will be assigned the type of **Other** by default. Learn more about tagging consistency in Application and component tags. If you plan to use a remote protection policy, you must choose **Other**.
- 4. Select the host operating system and LUN format. Enter the host initiator information.
 - a. For **New LUNs**: Select the host operating system and LUN format. Enter the host initiator information.
 - b. For **New NAS volumes**: choose the appropriate export option (NFS or SMB/CIFS) based on the NAS configuration of your SVM.
 - c. For New NVMe namespaces: Select the host operating system and NVMe subsystem.
- 5. To add a child consistency group, select **More options** then **+Add child consistency group**.
- 6. Select the performance level, the number of LUNs or volumes, and capacity per LUN or volume. Designate the appropriate export configurations or operating system information based on the protocol you are using.
- 7. Optionally, select a local snapshot policy and set the access permissions.
- 8. Repeat for up to five child consistency groups.
- 9. Select Save.
- 10. Confirm your consistency group has been created by returning to the main consistency group menu where it will appear once the ONTAP job completes. If you set a protection policy, look under the appropriate policy, remote or local, which should display a green shield with a checkmark in it.

CLI

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

When creating a hierarchical consistency group in the CLI with new volumes, you must create each child consistency group individually.

Step

1. Create the new consistency group using the consistency-group create command.

consistency-group create -vserver <SVM_name> -consistency-group <consistency_group_name> -parent-consistency-group <parent_consistency_group_name> -volume-prefix <volume_prefix> -volume
-count <number_of_volumes> -size <size>

- 2. When prompted by the CLI, confirm you want to create the new parent consistency group. Enter y.
- 3. Optionally, repeat step 1 to create more child consistency groups.

Learn more about consistency-group create in the ONTAP command reference.

Create a hierarchical consistency group with existing volumes

You can organize existing volumes into a hierarchical consistency group.

Steps

- 1. Select Storage > Consistency groups.
- 2. Select +Add then Using existing volumes.
- 3. Select the storage VM.
- 4. Select the existing volumes to include. Only volumes that are not already part of a consistency group will be available for selection.
- 5. To add a child consistency group, select **+Add Child Consistency Group**. Create the necessary consistency groups, which will be named automatically.
 - a. Component Type: If you are using ONTAP 9.12.1 or later, select a component type of "data", "logs", or "other". If no value is selected, the consistency group will be assigned the type of Other by default. Learn more about tagging consistency in Application and component tags. If you plan to use a remote protection policy, you must use Other.
- 6. Assign existing volumes to each consistency group.
- 7. Optionally, select a local snapshot policy.
- 8. Repeat for up to five child consistency groups.
- 9. Select Save.
- 10. Confirm your consistency group has been created by returning to the main consistency group menu where it will appear once the ONTAP job completes. If you have chosen a protection policy, confirm it was properly set by selecting your consistency group from the menu; under the appropriate policy type, you will see a green shield with a checkmark inside of it.

CLI

Beginning with ONTAP 9.14.1, you can create an hierarchical consistency group using the CLI.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Steps

1. Provision a new parent consistency group and assign volumes to a new child consistency group:

```
consistency-group create -vserver <svm_name> -consistency-group
<child_consistency_group_name> -parent-consistency-group
<parent consistency group name> -volumes <volume names>
```

2. Enter y to confirm you want to create a new parent and child consistency group.

Learn more about consistency-group create in the ONTAP command reference.

Next steps

- Modify the geometry of a consistency groups
- Modify a consistency group

• Protect a consistency group

Protect ONTAP consistency groups

Consistency groups offer easily managed local and remote protection for SAN, NAS, and NVMe applications that span multiple volumes.

Creating a consistency group does not automatically enable protection. Protection policies can be set at the time of creation or after creating your consistency group. You can protect consistency groups using:

- Local snapshots
- SnapMirror active sync (referred to as SnapMirror Business Continuity in versions of ONTAP before 9.15.1)
- MetroCluster (beginning 9.11.1)
- SnapMirror asynchronous (beginning 9.13.1)
- Asynchronous SVM disaster recovery (beginning 9.14.1)

If you are utilizing nested consistency groups, you can set different protection policies for the parent and child consistency groups.

Beginning with ONTAP 9.11.1, consistency groups offer two-phase consistency group snapshot creation. The two-phase snapshot operation executes a pre-check, ensuring the snapshot is captured successfully.

Recovery can occur for an entire consistency group, a single consistency group in a hierarchical configuration, or for individual volumes within the consistency group. Recovery can be achieved by selecting the consistency group you want to recover from, selecting the snapshot type, and then identifying the snapshot to base the restoration on. For more information about this process, see Restore a volume from an earlier snapshot.

Configure a local snapshot policy

Setting a local snapshot protection policy allows you to create a policy spanning all volumes in a consistency group.

About this task

The minimum supported snapshot schedule for consistency groups is 30 minutes. This is based on testing for FlexGroup volumes, which share the same Snapshot infrastructure as consistency groups.

Steps

- 1. Select Storage > Consistency groups.
- 2. Select the consistency group you have created from the Consistency group menu.
- 3. At the top right of the overview page for the consistency group, select Edit.
- 4. Check the box next to Schedule Snapshot copies (local).
- 5. Select a snapshot policy. To configure a new, custom policy, refer to Create a custom data protection policy.
- 6. Select Save.
- 7. Return to the consistency group overview menu. In the left column under **Snapshots (Local)**, the status will say protected next to $\sqrt[4]{}$.

CLI

Beginning with ONTAP 9.14.1, you can modify the protection policy of a consistency group using the CLI.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Step

1. Issue the following command to set or modify the protection policy:

If you are modifying the protection policy of a child consistency, you must identify the parent consistency group using the -parent-consistency-group parent_consistency_group_name parameter.

```
consistency-group modify -vserver svm_name -consistency-group
consistency_group_name -snapshot-policy policy_name
```

Create an on-demand snapshot

If you need to create a snapshot of your consistency group outside of a normally scheduled policy, you can create one on-demand.

Steps

- 1. Navigate to Storage > Consistency groups.
- 2. Select the consistency group for which you want to create an on-demand snapshot.
- 3. Switch to the Snapshot copies tab then select +Add.
- 4. Provide a **Name** and a **SnapMirror Label**. In the dropdown menu for **Consistency**, select **Application consistent** or **Crash consistent**.
- 5. Select Save.

CLI

Beginning with ONTAP 9.14.1, you can create an on-demand snapshot of a consistency group using the CLI.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Step

1. Create the snapshot:

By default, the snapshot type is crash-consistent. You can modify the snapshot type with the optional -type parameter.

```
consistency-group snapshot create -vserver svm_name -consistency-group
consistency_group_name -snapshot snapshot_name
```

Create two-phase consistency group snapshots

Beginning with ONTAP 9.11.1, consistency groups support two-phase commits for consistency group (CG) snapshot creation, which execute a precheck before committing the snapshot. This feature is only available with the ONTAP REST API.

Two-phase CG snapshot creation is only available for snapshot creation, not provisioning consistency groups or restoring consistency groups.

A two-phase CG snapshot breaks the snapshot creation process into two phases:

- 1. In the first phase, the API executes prechecks and triggers snapshot creation. The first phase includes includes a timeout parameter, designating the amount of time for the snapshot to commit successfully.
- 2. If the request in phase one completes successfully, you can invoke the second phase within the designated interval from the first phase, committing the snapshot to the appropriate endpoint.

Before you begin

- To use two-phase CG snapshot creation, all nodes in the cluster must be running ONTAP 9.11.1 or later.
- Only one active invocation of a consistency group snapshot operation is supported on a consistency group instance at a time, whether it be a one-phase or two-phase. Attempting to invoke a snapshot operation while another one is in progress results in a failure.

• When you invoke the snapshot creation, you can set an optional timeout value of between 5 and 120 seconds. If no timeout value is provided, the operation times out at the default of 7 seconds. In the API, set the timeout value with the action timeout parameter. In the CLI, use the -timeout flag.

Steps

You can complete a two-phase snapshot with the REST API or, beginning with ONTAP 9.14.1, the ONTAP CLI. This operation is not supported in System Manager.



If you invoke the snapshot creation with the API, you must commit the snapshot with the API. If you invoke the snapshot creation with the CLI, you must commit the snapshot with the CLI. Mixing methods is not supported.

CLI

Beginning with ONTAP 9.14.1, you can create a two-phase snapshot using the CLI.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Steps

1. Initiate the snapshot:

```
consistency-group snapshot start -vserver svm_name -consistency-group
consistency_group_name -snapshot snapshot_name [-timeout time_in_seconds
-write-fence {true|false}]
```

2. Verify the snapshot was taken:

consistency-group snapshot show

3. Commit the snapshot:

```
consistency-group snapshot commit svm_name -consistency-group
consistency_group_name -snapshot snapshot_name
```

API

1. Invoke the snapshot creation. Send a POST request to the consistency group endpoint using the action=start parameter.

```
curl -k -X POST 'https://<IP_address>/application/consistency-
groups/<cg-uuid>/snapshots?action=start&action_timeout=7' -H
"accept: application/hal+json" -H "content-type: application/json"
-d '
{
    "name": "<snapshot_name>",
    "consistency_type": "crash",
    "comment": "<comment>",
    "snapmirror_label": "<SnapMirror_label>"
}'
```

2. If the POST request succeeds, the output includes a snapshot uuid. Using that uuid, submit a PATCH request to commit the snapshot.

```
curl -k -X PATCH 'https://<IP_address>/application/consistency-
groups/<cg_uuid>/snapshots/<snapshot_id>?action=commit' -H "accept:
application/hal+json" -H "content-type: application/json"
For more information about the ONTAP REST API, see
link:https://docs.netapp.com/us-en/ontap-
automation/reference/api_reference.html[API reference^] or the
link:https://devnet.netapp.com/restapi.php[ONTAP REST API page^] at
the NetApp Developer Network for a complete list of API endpoints.
```

Set remote protection for a consistency group

Consistency groups offer remote protection through SnapMirror active sync and, beginning with ONTAP 9.13.1, SnapMirror asynchronous.

Configure protection with SnapMirror active sync

You can use SnapMirror active sync to ensure snapshots of consistency groups created on your consistency group are copied to the destination. To learn more about SnapMirror active sync or how to configure SnapMirror active sync using the CLI, see Configure protection for business continuity.

Before you begin

- SnapMirror active sync relationships cannot be established on volumes mounted for NAS access.
- The policy labels in the source and destination cluster must match.
- SnapMirror active sync will not replicate snapshots by default unless a rule with a SnapMirror label is added to the predefined AutomatedFailOver policy and the snapshots are created with that label.

To learn more about this process, refer to Protect with SnapMirror active sync.

- Cascade deployments are not supported with SnapMirror active sync.
- Beginning with ONTAP 9.13.1, you can non-disruptively add volumes to a consistency group with an active SnapMirror active sync relationship. Any other changes to a consistency group require you to break the SnapMirror active sync relationship, modify the consistency group, then reestablish and resynchronize the relationship.



To configure SnapMirror active sync with the CLI, see Protect with SnapMirror active sync.

Steps for System Manager

- 1. Ensure you have met the prerequisites for using SnapMirror active sync.
- 2. Select Storage > Consistency groups.
- 3. Select the consistency group you have created from the Consistency group menu.
- 4. At the top right of the overview page, select More then Protect.
- 5. System Manager auto-fills source-side information. Select the appropriate cluster and storage VM for the destination. Select a protection policy. Ensure that **Initialize relationship** is checked.
- 6. Select Save.

7. The consistency group needs to initialize and synchronize. Confirm synchronization has completed successfully by returning to the **Consistency group** menu. The **SnapMirror (Remote)** status displays Protected next to .

Configure SnapMirror asynchronous

Beginning with ONTAP 9.13.1, you can configure SnapMirror asynchronous protection for a single consistency group. Beginning with ONTAP 9.14.1, you can use SnapMirror asynchronous to replicate volume-granular snapshots to the destination cluster using the consistency group relationship.

About this task

To replicate volume-granular snapshots, you must be running ONTAP 9.14.1 or later. For MirrorAndVault and Vault policies, the volume-granular snapshot policy's SnapMirror label must match the consistency group's SnapMirror policy rule. Volume-granular snapshots abide by the keep value of the consistency group's SnapMirror policy, which is calculated independently of the consistency group snapshots. For example, if you have a policy to keep two snapshots on the destination, you can have two volume-granular snapshots and two consistency group snapshots.

When resynchronizing the SnapMirror relationship with volume-granular snapshots, you can preserve volumegranular snapshots with the -preserve flag. Volume-granular snapshots newer than consistency group snapshots are preserved. If there is not a consistency group snapshot, no volume-granular snapshots can be transferred in the resync operation.

Before you begin

- SnapMirror asynchronous protection is only available for a single consistency group. It is not supported for hierarchical consistency groups. To convert a hierarchical consistency group into a single consistency group, see modify consistency group architecture.
- The policy labels in the source and destination cluster must match.
- You can non-disruptively add volumes to a consistency group with an active SnapMirror asynchronous relationship. Any other changes to a consistency group require you to break the SnapMirror relationship, modify the consistency group, then reestablish and resynchronize the relationship.
- Consistency groups enabled for protection with SnapMirror asynchronous have different limits. For more information, see Consistency group limits.
- If you have configured an SnapMirror asynchronous protection relationship for multiple individual volumes, you can convert those volumes into a consistency group while retaining the existing snapshots. To convert volumes successfully:
 - There must be a common snapshot of the volumes.
 - You must break the existing SnapMirror relationship, add the volumes to a single consistency group, then resynchronize the relationship using the following workflow.

Steps

- 1. From the destination cluster, select **Storage > Consistency groups**.
- 2. Select the consistency group you have created from the Consistency group menu.
- 3. At the top right of the overview page, select More then Protect.
- 4. System Manager auto-fills source-side information. Select the appropriate cluster and storage VM for the destination. Select a protection policy. Ensure that **Initialize relationship** is checked.

When selecting an asynchronous policy, you have the option to Override Transfer Schedule.



The minimum supported schedule (recovery point objective, or RPO) for consistency groups with SnapMirror asynchronous is 30 minutes.

- 5. Select Save.
- 6. The consistency group needs to initialize and synchronize. Confirm synchronization has completed successfully by returning to the **Consistency group** menu. The **SnapMirror (Remote)** status displays Protected next to \heartsuit .

Configure SVM disaster recovery

Beginning with ONTAP 9.14.1, SVM disaster recovery supports consistency groups, enabling you to mirror consistency group information from the source to the destination cluster.

If you are enabling SVM disaster recovery on an SVM that already contains a consistency group, following the SVM configuration workflows for System Manager or the ONTAP CLI.

If you are adding a consistency group to an SVM that is in an active and healthy SVM disaster recovery relationship, you must update the SVM disaster recovery relationship from the destination cluster. For more information, see Update a replication relationship manually. You must update the relationship any time you expand the consistency group.

Limitations

- SVM disaster recovery does not support hierarchical consistency groups.
- SVM disaster recovery does not support consistency groups protected with SnapMirror asynchronous. You must break the SnapMirror relationship before configuring SVM disaster recovery.
- Both clusters must be running ONTAP 9.14.1 or later.
- Fan-out relationships are not supported for SVM disaster recovery configurations that contain consistency groups.
- For other limits, see consistency group limits.

Visualize relationships

System Manager visualizes LUN maps under the **Protection > Relationships** menu. When you select a source relationship, System Manager displays a visualization of the source relationships. By selecting a volume, you can delve deeper into these relationships to see a list of the contained LUNs and the initiator group relationships. This information can be downloaded as an Excel workbook from the individual volume view; the download operation runs in the background.

Related information

- · Clone a consistency group
- Configure snapshots
- · Create custom data protection policies
- Recover from snapshots
- · Restore a volume from an earlier snapshot
- SnapMirror active sync overview
- ONTAP automation documentation
- SnapMirror asynchronous disaster recovery basics

Modify member volumes in an ONTAP consistency group

Beginning with ONTAP 9.12.1, you can modify a consistency group by removing volumes or adding volumes (expanding the consistency group). Beginning with ONTAP 9.13.1, you can move volumes between child consistency groups if they share a common parent.

Add volumes to a consistency group

Beginning with ONTAP 9.12.1, you can non-disruptively add volumes to a consistency group.

About this task

- You cannot add volumes associated with another consistency group.
- Consistency groups support NAS, SAN, and NVMe protocols.
- You can add up to 16 volumes at a time to a consistency group if the adjustments are within the overall consistency group limits.
- Beginning with ONTAP 9.13.1, you can non-disruptively add volumes to a consistency group with an active SnapMirror active sync or SnapMirror asynchronous protection policy.
- When you add volumes to a consistency group protected by SnapMirror active sync, the status of the SnapMirror active sync relationship status changes to "Expanding" until mirroring and protection are configured for the new volume. If a disaster occurs on the primary cluster before this process completes, the consistency group reverts back to its original composition as part of the failover operation.
- In ONTAP 9.12.1 and earlier, you *cannot* add volumes to a consistency group in an SnapMirror active sync relationship. You must first delete the SnapMirror active sync relationship, modify the consistency group, then restore protection with SnapMirror active sync.
- Beginning with ONTAP 9.12.1, the ONTAP REST API supports adding *new* or existing volumes to a consistency group. For more information about the ONTAP REST API, refer to ONTAP REST API reference documentation.

Beginning with ONTAP 9.13.1, this functionality is supported in System Manager.

- When expanding a consistency group, snapshots of the consistency group captured before the modification will be considered partial. Any restore operation based on that snapshot will reflect the consistency group at the point-in-time of the snapshot.
- If you are using ONTAP 9.10.1 through 9.11.1, you cannot modify a consistency group. To change the configuration of a consistency group in ONTAP 9.10.1 or 9.11.1, you must delete the consistency group, then create a new consistency group with the volumes you want to include.
- Beginning with ONTAP 9.14.1, you can replicate volume-granular snapshots to the destination cluster when using SnapMirror asynchronous. When expanding a consistency group using SnapMirror asynchronous, volume-granular snapshots are only replicated after expanding the consistency group when the SnapMirror policy is MirrorAll or MirrorAndVault. Only volume-granular snapshots newer than the baseline consistency group snapshot are replicated.
- If you add volumes to a consistency group in an SVM disaster recovery relationship (supported beginning with ONTAP 9.14.1), you must update the SVM disaster recovery relationship from the destination cluster after expanding the consistency group. For more information, see Update a replication relationship manually.
- If you are using NVMe with ONTAP 9.17.1, you cannot modify a consistency group.

Beginning with ONTAP 9.12.1, you can perform this operation with System Manager.

- 1. Select Storage > Consistency groups.
- 2. Select the consistency group that you want to modify.
- 3. If you are modifying a single consistency group, at the top of the **Volumes** menu, select **More** and then **Expand** to add a volume.

If you are modifying a child consistency group, identify the parent consistency group you want to modify. Select the > button to view the child consistency groups, then select in next to the name of the child consistency group you want to modify. From that menu, select **Expand**.

- 4. Select up to 16 volumes to add to the consistency group.
- 5. Select **Save**. When the operation completes, view the newly added volumes in the consistency group's **Volumes** menu.

CLI

Beginning with ONTAP 9.14.1, you can add volumes to a consistency group using the ONTAP CLI.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Add existing volumes

1. Issue the following command. The -volumes parameter accepts a comma-separated list of volumes.



Only include the *-parent-consistency-group* parameter if the consistency group is in an hierarchical relationship.

```
consistency-group volume add -vserver svm_name -consistency-group
consistency_group_name -parent-consistency-group parent_consistency_group
-volume volumes
```

Add new volumes

The procedure to add new volumes depends on the protocol you are using.



Only include the -parent-consistency-group parameter if the consistency group is in a hierarchical relationship.

· To add new volumes without exporting them:

```
consistency-group volume create -vserver SVM_name -consistency-group
child_consistency_group -parent-consistency-group existingParentCg -volume
volume_name -size size
```

• To add new NFS volumes:

consistency-group volume create -vserver SVM_name -consistency-group consistency-group-name -volume volume-prefix -volume-count number -size size -export-policy policy_name

• To add new SAN volumes:

consistency-group volume create -vserver SVM_name -consistency-group
consistency-group-name -lun lun_name -size size -lun-count number -igroup
igroup name

• To add new NVMe namespaces:

```
consistency-group volume create -vserver SVM_name -consistency-group
consistency_group_name -namespace namespace_name -volume-count number
-namespace-count number -size size -subsystem subsystem name
```

Remove volumes from a consistency group

Volumes removed from a consistency group are not deleted. They remain active in the cluster.

About this task

- You cannot remove volumes from a consistency group in a SnapMirror active sync or SVM disaster recovery relationship. You must first delete the SnapMirror active sync relationship to modify the consistency group and then reestablish the relationship.
- If a consistency group has no volumes in it following the remove operation, the consistency group will be deleted.
- When a volume is removed from a consistency group, existing snapshots of the consistency group remain but are considered invalid. The existing snapshots cannot be used to restore the contents of the consistency group. Volume-granular snapshots remain valid.
- If you delete a volume from the cluster, it is automatically removed from the consistency group.
- To change the configuration of a consistency group in ONTAP 9.10.1 or 9.11.1, you must delete the consistency group then create a new consistency group with the desired member volumes.
- Deleting a volume from the cluster will automatically remove it from the consistency group.

Beginning with ONTAP 9.12.1, you can perform this operation with System Manager.

Steps

- 1. Select Storage > Consistency groups.
- 2. Select the single or child consistency group that you want to modify.
- 3. In the **Volumes** menu, select the checkboxes next to the individual volumes you want to remove from the consistency group.
- 4. Select Remove volumes from the consistency group.
- 5. Confirm that you understand removing the volumes will cause all snapshots of the consistency group to become invalid and select **Remove**.

CLI

Beginning with ONTAP 9.14.1, you can remove volumes from a consistency group using the CLI.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Step

1. Remove the volumes. The -volumes parameter accepts a comma-separated list of volumes.

Only include the <code>-parent-consistency-group</code> parameter if the consistency group is in an hierarchical relationship.

```
consistency-group volume remove -vserver SVM_name -consistency-group
consistency_group_name -parent-consistency-group
parent consistency group name -volume volumes
```

Move volumes between consistency groups

Beginning with ONTAP 9.13.1, you can move volumes between child consistency groups that share a parent.

About this task

- You can only move volumes between consistency groups nested under the same parent consistency group.
- Existing consistency group snapshots become invalid and no longer accessible as consistency group snapshots. Individual volume snapshots remain valid.
- · Snapshots of the parent consistency group remain valid.
- If you move all volumes out of a child consistency group, that consistency group will be deleted.
- Modifications to a consistency group must abide by consistency group limits.

Beginning with ONTAP 9.12.1, you can perform this operation with System Manager.

Steps

1. Select Storage > Consistency groups.

- 2. Select the parent consistency group that contains the volumes you want to move. Find the child consistency group and then expand the **Volumes** menu. Select the volumes you want to move.
- 3. Select Move.
- 4. Choose whether you want to move the volumes to a new consistency group or an existing group.
 - a. To move to an existing consistency group, select **Existing child consistency group** then choose the consistency group's name from the dropdown menu.
 - b. To move to a new consistency group, select **New child consistency group**. Enter a name for the new child consistency group and select a component type.
- 5. Select Move.

CLI

Beginning with ONTAP 9.14.1, you can move volumes between consistency groups using the ONTAP CLI.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Move volumes to a new child consistency group

1. The following command creates a new child consistency group that contains the designated volumes.

When you create the new consistency group, you can designate new snapshot, QoS, and tiering policies.

```
consistency-group volume reassign -vserver SVM_name -consistency-group
source_child_consistency_group -parent-consistency-group
parent_consistency_group -volume volumes -new-consistency-group
consistency_group_name [-snapshot-policy policy -qos-policy policy -tiering
-policy policy]
```

Move volumes to an existing child consistency group

1. Reassign the volumes. The -volumes parameter accepts a comma-separated list of volume names.

```
consistency-group volume reassign -vserver SVM_name -consistency-group
source_child_consistency_group -parent-consistency-group
parent_consistency_group -volume volumes -to-consistency-group
target_consistency_group
```

Related information

Consistency group limits

Clone a consistency group

Modify ONTAP consistency group geometry

Beginning with ONTAP 9.13.1, you can modify the geometry of a consistency group. Modifying the geometry of a consistency group enables you to alter the configuration of child or parent consistency groups without disruption to ongoing IO operations.

Modifying consistency group geometry has an impact on existing snapshots of the consistency group. For details, refer to the specific modification to geometry you want to perform.



You cannot modify the geometry of a consistency group that is configured with a remote protection policy. You must first break the protection relationship, modify the geometry, then restore remote protection.

Add a new child consistency group

Beginning with ONTAP 9.13.1, you can add a new child consistency group to an existing parent consistency group.

About this task

- A parent consistency group can contain a maximum of five child consistency groups. See consistency group limits for other limits.
- You cannot add a child consistency group to a single consistency group. You must first promote the consistency group, then you can add a child consistency group.
- Existing snapshots of the consistency group captured before the expand operation will be considered partial. Any restore operation based on that snapshot will reflect the consistency group at the point-in-time of the snapshot.

Beginning with ONTAP 9.13.1, you can perform this operation with System Manager.

Add a new child consistency group

- 1. Select Storage > Consistency groups.
- 2. Select the parent consistency group you want to which you want to add a child consistency group.
- 3. Next to the parent consistency group's name, select More then Add new child consistency group.
- 4. Enter a name for your consistency group.
- 5. Choose whether you would like to add new or existing volumes.
 - a. If you are adding existing volumes, select **Existing volumes** then choose the volumes from the dropdown menu.
 - b. If you are adding new volumes, select **New volumes** then designate the number of volumes and their size.
- 6. Select Add.

CLI

Beginning with ONTAP 9.14.1, you can add a child consistency group using the ONTAP CLI.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Add a child consistency group with new volumes

1. Create the new consistency group. Provide values for the consistency group name, volume prefix, number of volumes, volume size, storage service, and export policy name:

```
consistency-group create -vserver SVM_name -consistency-group
consistency_group -parent-consistency-group parent_consistency_group
-volume-prefix prefix -volume-count number -size size -storage-service
service -export-policy policy_name
```

Add a child consistency group with existing volumes

1. Create the new consistency group. The volumes parameter accepts a comma-separated list of volume names.

```
consistency-group create -vserver SVM_name -consistency-group
new_consistency_group -parent-consistency-group parent_consistency_group
-volumes volume
```

Detach a child consistency group

Beginning with ONTAP 9.13.1, you can remove a child consistency group from its parent, converting it into an individual consistency group.

About this task

- Detaching a child consistency group causes the parent consistency group's snapshots to become invalid and inaccessible. Volume granular snapshots remain valid.
- Existing snapshots of the individual consistency group remain valid.
- This operation will fail if there is an existing single consistency group that has the same name as the child consistency group you intend to detach. If you encounter this scenario, you must rename the consistency group when you detach it.

Example 4. Steps

System Manager

Beginning with ONTAP 9.13.1, you can perform this operation with System Manager.

Detach a child consistency group

- 1. Select Storage > Consistency groups.
- 2. Select the parent consistency group that contains the child you want to detach.
- 3. Next to the child consistency group you want to detach, select More then Detach from parent.
- 4. Optionally, rename the consistency group and select an application type.

5. Select Detach.

CLI

Beginning with ONTAP 9.14.1, you can detach a child consistency group using the ONTAP CLI.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Detach a child consistency group

1. Detach the consistency group. Optionally, rename the detached consistency group with the -new -name parameter.

```
consistency-group detach -vserver SVM_name -consistency-group
child_consistency_group -parent-consistency-group parent_consistency_group
[-new-name new_name]
```

Move an existing single consistency group under a parent consistency group

Beginning with ONTAP 9.13.1, you can convert an existing single consistency group to a child consistency group. You can either move the consistency group under an existing parent consistency group or create a new parent consistency group during the move operation.

About this task

- The parent consistency group must have four or fewer children. A parent consistency group can contain a maximum of five child consistency groups. See consistency group limits for other limits.
- Existing snapshots of the *parent* consistency group captured before this operation are considered partial. Any restore operation based on one of those snapshots reflects the consistency group at the point-in-time

of the snapshot.

• Existing consistency group snapshots of the single consistency group remain valid.

Example 5. Steps

System Manager

Beginning with ONTAP 9.13.1, you can perform this operation with System Manager.

Move an existing single consistency group under a parent consistency group

- 1. Select Storage > Consistency groups.
- 2. Select the consistency group you want to convert.
- 3. Select More then Move under different consistency group.
- 4. Optionally, enter a new name for the consistency group and select a component type. By default, the component type will be Other.
- 5. Choose if you want to migrate to an existing parent consistency group or create a new parent consistency group:
 - a. To migrate to an existing parent consistency group, select **Existing consistency group** then choose the consistency group from the dropdown menu.
 - b. To create a new parent consistency group, select **New consistency group** then provide a name for the new consistency group.
- 6. Select Move.

CLI

Beginning with ONTAP 9.14.1, you can move a single consistency group under a parent consistency group using the ONTAP CLI.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Move a consistency group under a new parent consistency group

1. Create the new parent consistency group. The -consistency-groups parameter will migrate any existing consistency groups to the new parent.

consistency-group attach -vserver svm_name -consistency-group
parent_consistency_group -consistency-groups child_consistency_group

Move a consistency group under an existing consistency group

1. Move the consistency group:

```
consistency-group add -vserver SVM_name -consistency-group
consistency_group -parent-consistency-group parent_consistency_group
```

Promote a child consistency group

Beginning with ONTAP 9.13.1, you can promote a single consistency group to a parent consistency group. When you promote the single consistency group to a parent, you also create a new child consistency group that inherits all of the volumes in the original, single consistency group.

About this task

- If you want to convert a child consistency group to a parent consistency group, you must first detach the child consistency group then follow this procedure.
- Existing snapshots of the consistency group remain valid after you promote the consistency group.

System Manager

Beginning with ONTAP 9.13.1, you can perform this operation with System Manager.

Promote a child consistency group

- 1. Select Storage > Consistency groups.
- 2. Select the consistency group you want to promote.
- 3. Select More then Promote to parent consistency group.
- 4. Enter a Name and select a Component type for the child consistency group.
- 5. Select Promote.

CLI

Beginning with ONTAP 9.14.1, you can move a single consistency group under a parent consistency group using the ONTAP CLI.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Promote a child consistency group

1. Promote the consistency group. This command will create one parent and one child consistency group.

```
consistency-group promote -vserver SVM_name -consistency-group
existing consistency group -new-name new child consistency group
```

Demote a parent to a single consistency group

Beginning with ONTAP 9.13.1, you can demote a parent consistency group to a single consistency group. Demoting the parent flattens the hierarchy of the consistency group, removing all associated child consistency groups. All volumes in the consistency group will remain under the new, single consistency group.

About this task

• Existing snapshots of the *parent* consistency group remain valid after you demote it to a single consistency. Existing snapshots of any of the associated *child* consistency groups of that parent become invalid upon demotion. The individual volume snapshots within the child consistency group continue to be accessible as volume-granular snapshots.

Example 6. Steps

System Manager

Beginning with ONTAP 9.13.1, you can perform this operation with System Manager.

Demote a consistency group

- 1. Select Storage > Consistency groups.
- 2. Select the parent consistency group you want to demote.
- 3. Select More then Demote to single consistency group.
- 4. A warning will advise you that all associated child consistency groups will be deleted and their volumes will be moved under the new single consistency group. Select **Demote** to confirm you understand the impact.

CLI

Beginning with ONTAP 9.14.1, you can demote a consistency group using the ONTAP CLI.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Demote a consistency group

1. Demote the consistency group. Use the optional -new-name parameter to rename the consistency group.

consistency-group demote -vserver SVM_name -consistency-group
parent_consistency_group [-new-name new_consistency_group_name]

Modify ONTAP consistency group application and component tags

Beginning with ONTAP 9.12.1, consistency groups support component and application tagging. Application and component tags are a management tool, enabling you to filter and identify different workloads in your consistency groups.

About this task

Consistency groups offer two types of tags:

- **Application tags**: these apply to individual and parent consistency groups. Application tags provide labeling for workloads such as MongoDB, Oracle, or SQL Server. The default application tag for consistency groups is Other.
- **Component tags**: Children in hierarchal consistency groups have component tags instead of application tags. The options for component tags are "data", "logs", or "other". The default value is Other.

You can apply tags when creating consistency groups or after the consistency groups have been created.



If the consistency group has a SnapMirror active sync relationship, you must use **Other** as the application or component tag.

Steps

Beginning with ONTAP 9.12.1, you can modify application and component tags using System Manager. Beginning with ONTAP 9.14.1, you can modify the application and component tags using the ONTAP CLI.

System Manager

- 1. Select Storage > Consistency groups.
- 2. Select the consistency group whose tag you want to modify. Select the **i** next to the consistency group's name then **Edit**.
- 3. In the dropdown menu, choose the appropriate application or component tag.
- 4. Select Save.

CLI

Beginning with ONTAP 9.14.1, you can modify the application or component tag of an existing consistency group using the ONTAP CLI.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Modify the application tag

1. Application tags accept a limited number of preset strings. To see the accepted list of strings, run the following command:

consistency-group modify -vserver svm_name -consistency-group
consistency_group -application-type ?

2. Choose the appropriate string from the output, the modify the consistency group: consistency-group modify -vserver *svm_name* -consistency-group *consistency group* -application-type application type

Modify the component tag

1. Modify the component type. The component type can be data, logs, or other. If you are using SnapMirror active sync, it must be "other." consistency-group modify -vserver svm -consistency-group child_consistency_group -parent-consistency-group parent_consistency_group -application-component-type [data|logs|other]

Clone an ONTAP consistency group

Beginning with ONTAP 9.12.1, you can clone a consistency group to create a copy of a consistency group and its contents. Cloning a consistency group creates a copy of the consistency group configuration, its metadata such as application type, and all the volumes and its contents such as files, directories, LUNs or NVMe namespaces.

About this task

When cloning a consistency group, you can clone it with its current configuration, but with volume contents as they are or based on an existing consistency group snapshot.
Cloning a consistency group is supported only for the entire consistency group. You cannot clone an individual child consistency group in a hierarchical relationship: only the complete consistency group configuration can be cloned.

When you clone a consistency group, the following components are not cloned:

- iGroups
- LUN maps
- NVMe subsystems
- NVMe namespace subsystem maps

Before you begin

- When you clone a consistency group, ONTAP will not create SMB shares for the cloned volumes if a share name is not specified. * Cloned consistency groups are not mounted if a junction path is not specified.
- If you attempt to clone a consistency group based on a snapshot that does not reflect the consistency group's current constituent volumes, the operation will fail.
- After you clone a consistency group, you need to perform the appropriate mapping operation.

Refer to Map igroups to multiple LUNs or Map an NVMe namespace to a subsystem for more information.

• Cloning a consistency group is not supported for a consistency group in a SnapMirror active sync relationship or with any associated DP volumes.

System Manager

Steps

- 1. Select Storage > Consistency groups.
- 2. Select the consistency group you want to clone from the Consistency Group menu.
- 3. At the top right of the overview page for the consistency group, select **Clone**.
- 4. Enter a name for the new, cloned consistency group or accept the default name.
 - a. Choose if you want to enable Thin Provisioning.
 - b. Choose **Split Clone** if you want to dissociate the consistency group from its source and allocate additional disk space for the cloned consistency group.
- 5. To clone the consistency group in its current state, choose Add a new Snapshot copy.

To clone the consistency group based on a snapshot, choose **Use an existing snapshot**. Selecting this option will open a new sub-menu. Choose the snapshot that you want to use as the basis for the clone operation.

- 6. Select Clone.
- 7. Return to the **Consistency Group** menu to confirm your consistency group has been cloned.

CLI

Beginning with ONTAP 9.14.1, you can clone a consistency group using the CLI with cluster admin credentials.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Clone a consistency group

1. The consistency-group clone create command clones the consistency group at its current point-in-time status. To base the clone operation on a snapshot, include the -source-snapshot parameter.

consistency-group clone create -vserver svm_name -consistency-group clone_name -source-consistency-group consistency_group_name [-sourcesnapshot snapshot_name]

Learn more about consistency-group clone create in the ONTAP command reference.

Next steps

- Map igroups to multiple LUNs
- Map an NVMe namespace to a subsystem

Delete an ONTAP consistency group

If you decide that you no longer need a consistency group, you can delete it.

About this task

- Deleting a consistency group deletes the instance of the consistency group and does *not* impact the constituent volumes or LUNs. Deleting a consistency group does not result in deletion of the snapshots present on each volume, but they will no longer be accessible as consistency group snapshots. The snapshots can, however, continue to be managed as ordinary volume granular snapshots.
- ONTAP automatically deletes a consistency group if all of the volumes in the consistency group are deleted.
- Deleting a parent consistency group results in the deletion of all associated child consistency groups.
- If you are using a version of ONTAP between 9.10.1 to 9.12.0, volumes can only be removed from a consistency group if the volume itself is deleted, in which case the volume is automatically removed from the consistency group. Beginning with ONTAP 9.12.1, you can remove volumes from a consistency group without deleting the consistency group. For more information on this process, refer to Modify a consistency group.

Example 7. Steps

System Manager

- 1. Select Storage > Consistency groups.
- 2. Select the consistency group you would like to delete.
- 3. Next to the name of the consistency group, select **‡** then **Delete**.

CLI

Beginning with ONTAP 9.14.1, you can delete a consistency group using the CLI.

Before you begin

- You must be at the admin privilege level to perform this task.
- Beginning with ONTAP 9.15.1, any user at the admin privilege level can perform this task. In ONTAP 9.14.1, you must be a cluster or SVM administrator to perform this task.

Delete a consistency group

1. Delete the consistency group:

```
consistency-group delete -vserver svm_name -consistency-group
consistency_group_name
```

SnapMirror active sync

Introduction

Learn about ONTAP SnapMirror active sync

SnapMirror active sync (also referred to as SnapMirror Business Continuity *[SM-BC]*), enables business services to continue operating even through a complete site failure, supporting applications to fail over transparently using a secondary copy. There is no manual intervention or custom scripting required to trigger a failover with SnapMirror active sync.

Support for SnapMirror active sync varies depending on your version of ONTAP:

ONTAP version	Supported clusters	Supported protocols	Supported configurations
9.17.1 and later	 AFF ASA C-Series ASA r2 	 iSCSI FC NVMe 	 Asymmetric active/active Asymmetric active/activ e does not support ASA r2 and NVMe For more information about NVMe support, see NVMe configuratio n, support, and limitations. Symmetric active/active
9.15.1 and later	 AFF ASA C-Series 	• iSCSI • FC	 Asymmetric active/active Symmetric active/active Symmetric active/active configurations support 2-node clusters in ONTAP 9.15.1. 4- node clusters are supported in ONTAP 9.16.1 and later.
9.9.1 and later	 AFF ASA C-Series * 	• iSCSI • FC	Asymmetric active/active

Primary and secondary clusters must be of the same type: either ASA, ASA r2, or AFF.



Beginning July 2024, content from technical reports previously published as PDFs has been integrated with ONTAP product documentation. The ONTAP SnapMirror active sync documentation now includes content from *TR-4878: SnapMirror active sync*.

Benefits

SnapMirror active sync provides the following benefits:

- Continuous availability for business-critical applications.
- Ability to host critical applications alternately from primary and secondary sites.
- Simplified application management using consistency groups for dependent write-order consistency.
- The ability to test failover for each application.
- Instantaneous creation of mirror clones without impacting application availability.
- The ability to deploy protected and non-protected workloads in the same ONTAP cluster.
- LUN, NVMe namespace, NVMe subsystem, or storage unit identity remains the same, so the application sees them as a shared virtual device.
- The ability to reuse secondary clusters with flexibility to create instantaneous clones for application usage for dev-test, UAT or reporting purposes without impacting application performance or availability.

SnapMirror active sync allows you to protect your data LUNs or NVMe namespaces, which enables applications to fail over transparently for the purpose of business continuity in the event of a disaster. For more information, see Use cases.

Key concepts

SnapMirror active sync uses consistency groups and either the ONTAP Mediator or, beginning with ONTAP 9.17.1, the ONTAP Cloud Mediator to ensure your data is replicated and served even in the event of a disaster scenario. When planning your SnapMirror active sync deployment, it is important to understand the essential concepts in SnapMirror active sync and its architecture.

Asymmetry and symmetry

In symmetric active/active configurations, both sites can access local storage for active I/O. Symmetric active/active is optimized for clustered applications including VMware vMSC, Windows Failover Cluster with SQL, and Oracle RAC.

In asymmetric active/active configurations data on the secondary site is proxied to a LUN, namespace or storage unit.

For more information, see SnapMirror active sync architecture.

Consistency group

For AFF and ASA systems a consistency group is a collection of FlexVol volumes that provide a consistency guarantee for the application workload that must be protected for business continuity. In ASA r2 systems, a consistency group is a collection of storage units.

The purpose of a consistency group is to take simultaneous snapshot images of a collection of volumes or storage units, thus ensuring crash-consistent copies of the collection at a point in time. A consistency group ensures all volumes of a dataset are quiesced and then snapped at precisely the same point in time. This provides a data-consistent restore point across volumes or storage units supporting the dataset. A consistency group thereby maintains dependent write-order consistency. If you decide to protect applications for business continuity, the group of volumes or storage units corresponding to this application must be added to a

consistency group so a data protection relationship is established between a source and a destination consistency group. The source and destination consistency must contain the same number and type of volumes.

Constituent

An individual volume, LUN, or NVMe namespace (beginning with ONTAP 9.17.1) that is part of the consistency group protected in the SnapMirror active sync relationship.

ONTAP Mediator

The ONTAP Mediator receives health information about peered ONTAP clusters and nodes, orchestrating between the two and determining if each node/cluster is healthy and running. ONTAP Mediator provides health information about:

- Peer ONTAP clusters
- Peer ONTAP cluster nodes
- Consistency groups (which define the failover units in a SnapMirror active sync relationship); for each consistency group, the following information is provided:
 - Replication state: Uninitialized, In Sync, or Out of Sync
 - · Which cluster hosts the primary copy
 - Operation context (used for planned failover)

With this ONTAP Mediator health information, clusters can differentiate between distinct types of failures and determine whether to perform an automated failover. ONTAP Mediator is one of the three parties in the SnapMirror active sync quorum along with both ONTAP clusters (primary and secondary). To reach consensus, at least two parties in the quorum must agree to a certain operation.



Beginning with ONTAP 9.15.1, System Manager displays the status of your SnapMirror active sync relationship from either cluster. You can also monitor the ONTAP Mediator's status from either cluster in System Manager. In earlier releases of ONTAP, System Manager displays the status of SnapMirror active sync relationships from the source cluster.

ONTAP Cloud Mediator

ONTAP Cloud Mediator is available beginning with ONTAP 9.17.1. ONTAP Cloud Mediator provides the same services as ONTAP Mediator, except that it is hosted in the cloud using BlueXP.

Planned failover

A manual operation to change the roles of copies in a SnapMirror active sync relationship. The primary sites becomes the secondary, and the secondary becomes the primary.

Primary-first and primary bias

SnapMirror active sync uses a primary-first principle that gives preference to the primary copy to serve I/O in case of a network partition.

Primary-bias is a special quorum implementation that improves availability of a SnapMirror active sync protected dataset. If the primary copy is available, primary-bias comes into effect when the ONTAP Mediator is not reachable from both clusters.

Primary-first and primary bias are supported in SnapMirror active sync beginning with ONTAP 9.15.1. Primary copies are designated in System Manager and output with the REST API and CLI.

Automatic unplanned failover (AUFO)

An automatic operation to perform a failover to the mirror copy. The operation requires assistance from the ONTAP Mediator to detect that the primary copy is unavailable.

Out of Sync (OOS)

When the application I/O is not replicating to the secondary storage system, it will be reported as **out of sync**. An out of sync status means the secondary volumes are not synchronized with the primary (source) and that SnapMirror replication is not occurring.

If the mirror state is Snapmirrored, this indicates a transfer failure or failure due to an unsupported operation.

SnapMirror active sync supports automatic resync, enabling copies to return to an InSync state.

Beginning with ONTAP 9.15.1, SnapMirror active sync supports automatic reconfiguration in fan-out configurations.

Uniform and non-uniform configuration

- **Uniform host access** means that hosts from both sites are connected to all paths to storage clusters on both sites. Cross-site paths are stretched across distances.
- **Non-uniform host access** means hosts in each site are connected only to the cluster in the same site. Cross-site paths and stretched paths aren't connected.



Uniform host access is supported for any SnapMirror active sync deployment; non-uniform host access is only supported for symmetric active/active deployments.

Zero RPO

RPO stands for recovery point objective, which is the amount of data loss deemed acceptable during a given time period. Zero RPO signifies that no data loss is acceptable.

Zero RTO

RTO stands for recovery time objective, which is the amount of time that is deemed acceptable for an application to return to normal operations non-disruptively following an outage, failure, or other data loss event. Zero RTO signifies that no amount of downtime is acceptable.

ONTAP SnapMirror active sync architecture

The SnapMirror active sync architecture enables active workloads on both clusters, where primary workloads can be served simultaneously from both clusters. Regulations for financial institutions in some countries require businesses to be periodically serviceable from their secondary data centers as well, called "Tick-Tock" deployments, which SnapMirror active sync enables.

The data protection relationship to protect for business continuity is created between the source storage system and destination storage system, by adding the application specific LUNs or NVMe namespaces from different volumes within a storage virtual machine (SVM) to the consistency group. Under normal operations, the enterprise application writes to the primary consistency group, which synchronously replicates this I/O to the mirror consistency group.



Even though two separate copies of the data exist in the data protection relationship, because SnapMirror active sync maintains the same LUN or NVMe namespace identity, the application host sees this as a shared virtual device with multiple paths while only one LUN or NVMe namespace copy is being written to at a time. When a failure renders the primary storage system offline, ONTAP detects this failure and uses the Mediator for re-confirmation; if neither ONTAP nor the Mediator are able to ping the primary site, ONTAP performs the automatic failover operation. This process results in failing over only a specific application without the need for the manual intervention or scripting which was previously required for the purpose of failover.

Other points to consider:

- Unmirrored volumes which exist outside of protection for business continuity are supported.
- Only one other SnapMirror asynchronous relationship is supported for volumes being protected for business continuity.
- · Cascade topologies are not supported with protection for business continuity.

The role of mediators

SnapMirror active sync uses a mediator to act as a passive witness to SnapMirror active sync copies. In the event of a network partition or unavailability of one copy, SnapMirror active sync uses the mediator to determine which copy continues to serve I/O, while discontinuing I/O on the other copy. In addition to the on-

premises ONTAP Mediator, beginning with ONTAP 9.17.1, you can install ONTAP Cloud Mediator to provide the same functionality in a cloud deployment. You can use ONTAP Mediator or ONTAP Cloud Mediator, but you cannot use both at the same time.

The Mediator plays a crucial role in SnapMirror active sync configurations as a passive quorum witness, ensuring quorum maintenance and facilitating data access during failures. It acts as a ping proxy for controllers to determine liveliness of peer controllers. Although the Mediator does not actively trigger switchover operations, it provides a vital function by allowing the surviving node to check its partner's status during network communication issues. In its role as a quorum witness, the ONTAP Mediator provides an alternate path (effectively serving as a proxy) to the peer cluster.

Furthermore, it allows clusters to get this information as part of the quorum process. It uses the node management LIF and cluster management LIF for communication purposes. It establishes redundant connections through multiple paths to differentiate between site failure and InterSwitch Link (ISL) failure. When a cluster loses connection with the Mediator software and all its nodes due to an event, it is considered not reachable. This triggers an alert and enables automated failover to the mirror consistency group in the secondary site, ensuring uninterrupted I/O for the client. The replication data path relies on a heartbeat mechanism, and if a network glitch or event persists beyond a certain period, it can result in heartbeat failures, causing the relationship to go out-of-sync. However, the presence of redundant paths, such as LIF failover to another port, can sustain the heartbeat and prevent such disruptions.

ONTAP Mediator

ONTAP Mediator is installed in a third failure domain, distinct from the two ONTAP clusters it monitors. There are three key components in this setup:

- Primary ONTAP cluster hosting the SnapMirror active sync primary consistency group
- · Secondary ONTAP cluster hosting the mirror consistency group
- ONTAP Mediator

ONTAP Mediator is used for the following purposes:

- Establish a quorum
- · Continuous availability via automatic failover (AUFO)
- Planned failovers (PFO)



÷.

ONTAP Mediator 1.7 can manage ten cluster pairs for the purpose of business continuity.

When the ONTAP Mediator is not available, you cannot perform planned or automated failovers. The application data continues to synchronously replicate without any interruption to for zero data loss.

ONTAP Cloud Mediator

Beginning with ONTAP 9.17.1, ONTAP Cloud Mediator is available as a cloud-based service in BlueXP for use with SnapMirror active sync. Similar to ONTAP Mediator, ONTAP Cloud Mediator provides the following functionality in a SnapMirror active sync relationship:

- Provides a persistent and fenced store for HA or SnapMirror active sync metadata.
- · Serves as ping proxy for controller liveliness.
- Provides synchronous node health query functionality to aid in quorum determination.

The ONTAP Cloud Mediator helps simplify SnapMirror active sync deployment by using the BlueXP cloud

service as a third site that you do not need to manage. The ONTAP Cloud Mediator service provides the same functionality as the on-premises ONTAP Mediator; however, ONTAP Cloud Mediator reduces the operational complexity of maintaining a third site. In contrast, ONTAP Mediator is available as a package and must be installed on a Linux host running at a third site with independent power and network infrastructure for its operations.

SnapMirror active sync operation workflow

The following figure illustrates the design of SnapMirror active sync at a high level.



The diagram shows an enterprise application that is hosted on an storage VM (SVM) at the primary data center. The SVM contains five volumes, three of which are part of a consistency group. The three volumes in the consistency group are mirrored to a secondary data center. In normal circumstances, all write operations are performed to the primary data center; in effect, this data center serves as the source for I/O operations, while the secondary data center serves as a destination.

In the event of a disaster scenario at the primary data center, ONTAP directs the secondary data center to act as the primary, serving all I/O operations. Only the volumes that are mirrored in the consistency group are served. Any operations pertaining to the other two volumes on the SVM is be affected by the disaster event.

Symmetric active/active

SnapMirror active sync offers asymmetric and symmetric solutions.

In *asymmetric configurations*, the primary storage copy exposes an active-optimized path and actively serves client I/O. The secondary site uses a remote path for I/O. The storage paths for the secondary site are considered active-non-optimized. Access to the write LUN is proxied from the secondary site. NVMe protocol is not supported in asymmetric configurations.

In *symmetric active/active configurations*, active-optimized paths are exposed on both sites, are host specific, and are configurable, meaning hosts on either side are able to access local storage for active I/O. Beginning with ONTAP 9.16.1, symmetric active/active is supported on clusters with up to four nodes. Beginning with ONTAP 9.17.1, symmetric active/active configurations support NVMe protocol on two node clusters.



Symmetric active/active is targeted for clustered applications including VMware Metro Storage Cluster, Oracle RAC, and Windows Failover Clustering with SQL.

Use cases for ONTAP SnapMirror active sync

The demands of a globally connected business environment demand rapid recovery of business-critical application data with zero data loss in the event of a disruption such as a cyber attack, power outage, or natural disaster. These demands are heightened in arenas such as finance and those adhering to regulatory mandates such as the General Data Protection Regulation (GDPR).

SnapMirror active sync provides the following use cases:

Application deployment for zero recovery time objective (RTO)

In a SnapMirror active sync deployment, you have a primary and secondary cluster. A LUN in the primary cluster (1LP) has a mirror (L1s) on the secondary; both LUNs share the same serial ID and are reported as read-write LUNs to the host. In asymmetric configurations read and write operations, however, are only serviced to the primary LUN, 1LP. Any writes to the mirror L1s are served by proxy.

Application deployment for zero RTO or transparent application failover (TAF)

TAF is based on host MPIO software-based path failover to achieve non-disruptive access to the storage. Both LUN copies—for example, primary (L1P) and mirror copy (L1S)--have the same identity (serial number) and are reported as read-writable to the host. In asymmetric configurations however, reads and writes are serviced

only by the primary volume. I/Os issued to the mirror copy are proxied to the primary copy. The host's preferred path to L1 is VS1:N1 based on asymmetric logical unit access (ALUA) access state Active Optimized (A/O). ONTAP Mediator is required as part of the deployment, primarily to perform failover (planned or unplanned) in the event of a storage outage on the primary.

TAF operates in two modes: Automated Failover and Automated Failover Duplex. With Automated Failover, reads and writes are serviced only by the primary volume, therefore, IOs issued to the mirror copy (which cannot service writes on its own) are proxied to the primary copy. With Automated Failover Duplex, both the primary and secondary copies can service IOs so no proxy is necessary.

If you are using NVMe for host access with ONTAP 9.17.1, only Automated Failover Duplex is supported.

SnapMirror active sync uses ALUA, a mechanism that allows an application host multipathing software with paths advertised with priorities and access availability for the application host communication with the storage array. ALUA marks active optimized paths to the controllers owning the LUN and others as active non-optimized paths, used only if the primary path fails.

SnapMirror active sync with NVMe protocol uses ANA, which enables application hosts to discover optimized and non-optimized paths to NVMe namespaces that are being protected. The ONTAP NVMe target publishes the appropriate path states to enable application hosts to use the optimal path for a protected NVMe namespace.

Clustered applications

Clustered applications including VMware Metro Storage Cluster, Oracle RAC, and Windows Failover Clustering with SQL require simultaneous access so the VMs can be failed over to other site without any performance overhead. SnapMirror active sync symmetric active/active serves IO locally with bidirectional replication to meet the requirements of clustered applications. Beginning with ONTAP 9.16.1, symmetric active/active is supported in a configuration in four-node clusters, expanding from the two-node cluster limit in ONTAP 9.15.1.

Disaster scenario

Synchronously replicate multiple volumes for an application between sites at geographically dispersed locations. You can automatically failover to the secondary copy in case of disruption of the primary, thus enabling business continuity for tier one applications. When the site hosting the primary cluster experiences a disaster, the host multipathing software marks all paths through the cluster as down and uses paths from the secondary cluster. The result is a non-disruptive failover enabled by ONTAP Mediator to the mirror copy.

Windows failover

SnapMirror active sync provides flexibility with easy-to-use application-level granularity and automatic failover. SnapMirror active sync uses proven SnapMirror synchronous replication over IP network to replicate data at high speeds over LAN or WAN, to achieve high data availability and fast data replication for your businesscritical applications such as Oracle, Microsoft SQL Server, and so on, in both virtual and physical environments.

SnapMirror active sync enables mission-critical business services to continue operating even through a complete site failure, with TAF to the secondary copy. No manual intervention or no additional scripting are required to trigger this failover.

Deployment strategy and best practices for ONTAP SnapMirror active sync

It is important that your data protection strategy clearly identifies the workloads threats need to be protected for business continuity. The most critical step in your data protection strategy is to have clarity in your enterprise application data layout so that you can decide how you are distributing the volumes and protecting business continuity. Because failover occurs at the consistency group level on a per-application basis, make sure to add the necessary data volumes to the consistency group.

SVM configuration

The diagram captures a recommended storage VM (SVM) configuration for SnapMirror active sync.



- For data volumes:
 - Random read workloads are isolated from sequential writes; therefore, depending on the database size, the data and log files are typically placed on separate volumes.
 - For large critical databases, the single data file is on FlexVol 1 and its corresponding log file is on FlexVol 2.
 - For better consolidation, small-to-medium-size noncritical databases are grouped such that all the data files are on FlexVol 1 and their corresponding log files are on FlexVol 2. However, you will lose application-level granularity through this grouping.
 - Another variant is to have all the files within the same FlexVol 3, with data files in LUN1 and its log files in LUN 2.
- If your environment is virtualized, you would have all the VMs for various enterprise applications shared in a datastore. Typically, the VMs and application binaries are asynchronously replicated using SnapMirror.

Plan

Prerequisites for ONTAP SnapMirror active sync

When planning your SnapMirror active sync deployment, ensure you have met the various hardware, software, and system configuration requirements.

Hardware

The following table outlines the supported NetApp cluster configurations.

Cluster type	Supported models	Supported features	Maximum supported cluster nodes
AFF	A-Series, C-Series	Automated Failover Duplex (Symmetric Active/Active), Automated Failover (Asymmetric Active/Active)	 2 (ONTAP 9.9.1 or later) 4 (ONTAP 9.16.1 with Symmetric Active/Active configurations)
ASA	A-Series, C-Series	Automated Failover Duplex (Symmetric Active/Active), Automated Failover (Asymmetric Active/Active)	 2 (ONTAP 9.9.1 or later 4 (ONTAP 9.16.1 with Symmetric Active/Active configurations)
ASA r2	All	Automated Failover Duplex (Symmetric Active/Active)	2

The table below outlines the capability for replication between cluster types.

Cluster type 1	Cluster type 2	Replication supported?
AFF A-Series	AFF C-Series	Yes
ASA r2 A-Series	ASA r2 C-Series	Yes
AFF	ASA	No
ASA	ASA r2	No
ASA r2	ASA r2	Yes

Software

ONTAP 9.9.1 or later

- ONTAP Mediator 1.2 or later
- A Linux server or virtual machine for ONTAP Mediator running one of the following:

ONTAP Mediator version	Supported Linux versions
1.10	Red Hat Enterprise Linux
	 Compatible: 9.5¹
	 Recommended: 10, 9.6, 9.4, and 8.10
	• Rocky Linux 10, 9.6, and 8.10
1.9.1	Red Hat Enterprise Linux
	 Compatible: 9.3, 9.1, 8.9, 8.7, 8.6, 8.5, and 8.4¹
	 Recommended: 9.5, 9.4, 9.2, 9.0, 8.10, and 8.8
	• Rocky Linux 9.5 and 8.10
1.9	Red Hat Enterprise Linux
	 Compatible: 9.3, 9.1, 8.9, 8.7, 8.6, 8.5, and 8.4⁻¹
	 Recommended: 9.5, 9.4, 9.2, 9.0, 8.10, and 8.8
	• Rocky Linux 9.5 and 8.10
1.8	Red Hat Enterprise Linux:
	 Compatible: 8.7, 8.6, 8.5, and 8.4⁻¹
	 Recommended: 9.4, 9.3, 9.2, 9.1, 9.0, 8.10, 8.9, and 8.8
	• Rocky Linux 9.4 and 8.10
1.7	Red Hat Enterprise Linux:
	 Compatible: 8.7, 8.6, 8.5, and 8.4⁻¹
	 Recommended: 9.3, 9.2, 9.1, 9.0, 8.9, and 8.8
	• Rocky Linux 9.3 and 8.9
1.6	Red Hat Enterprise Linux:
	 Compatible: 8.7, 8.6, 8.5, and 8.4⁻¹
	 Recommended: 9.2, 9.1, 9.0, and 8.8
	• Rocky Linux 9.2 and 8.8
1.5	• Red Hat Enterprise Linux: 8.5, 8.4, 8.3, 8.2, 8.1, 8.0, 7.9, 7.8, 7.7, and 7.6
	• CentOS: 7.9, 7.8, 7.7, and 7.6

1.4	 Red Hat Enterprise Linux: 8.5, 8.4, 8.3, 8.2, 8.1, 8.0, 7.9, 7.8, 7.7, and 7.6 CentOS: 7.9, 7.8, 7.7, and 7.6
1.3	 Red Hat Enterprise Linux: 8.3, 8.2, 8.1, 8.0, 7.9, 7.8, 7.7, and 7.6 CentOS: 7.9, 7.8, 7.7, and 7.6
1.2	Red Hat Enterprise Linux: 8.1, 8.0, 7.9, 7.8, 7.7, and 7.6CentOS: 7.9, 7.8, 7.7, and 7.6

1. Compatible means that Red Hat no longer supports these RHEL versions, but ONTAP Mediator can still be installed on them.

Licensing

- SnapMirror synchronous license must be applied on both clusters.
- SnapMirror license must be applied on both clusters.



If your ONTAP storage systems were purchased before June 2019, see NetApp ONTAP Master License Keys to get the required SnapMirror synchronous license.

Networking environment

- Inter-cluster latency round trip time (RTT) must be less than 10 milliseconds.
- Beginning with ONTAP 9.14.1, SCSI-3 persistent reservations are supported with SnapMirror active sync.

Supported protocols

SnapMirror active sync supports SAN protocols.

- The FC and iSCSI protocols are supported beginning with ONTAP 9.9.1.
- The NVMe protocol is supported with VMware workloads beginning with ONTAP 9.17.1.



NVMe/TCP with VMware depends on the resolution of VMware Bug ID: TR1049746.

SnapMirror active sync does not support the following with the NVMe protocol:

- · 4-node symmetric active/active configurations
- · Changes in consistency group size

You cannot expand or shrink a consistency group when using the NVMe protocol with SnapMirror active sync.

· Coexistence of LUNs and namespaces in the same consistency group is not supported.

IPspace

The default IPspace is required by SnapMirror active sync for cluster peer relationships. Custom IPspace is not

supported.

NTFS Security Style

NTFS security style is not supported on SnapMirror active sync volumes.

ONTAP Mediator

- ONTAP Mediator must be provisioned externally and attached to ONTAP for transparent application failover.
- To be fully functional and to enable automatic unplanned failover, the external ONTAP Mediator should be provisioned and configured with ONTAP clusters.
- ONTAP Mediator must be installed in a third failure domain, separate from the two ONTAP clusters.
- When installing ONTAP Mediator, you should replace the self-signed certificate with a valid certificate signed by a mainstream reliable CA.
- For more information about ONTAP Mediator, see Prepare to install ONTAP Mediator.

Other prerequisites

- SnapMirror active sync relationships are not supported on read-write destination volumes. Before you can use a read-write volume, you must convert it to a DP volume by creating a volume-level SnapMirror relationship and then deleting the relationship. For details, see Convert an existing SnapMirror relationships to SnapMirror active sync.
- Storage VMs using SnapMirror active sync cannot be joined to Active Directory as a client computed.

Further information

- Hardware Universe
- ONTAP Mediator overview

ONTAP SnapMirror active sync interoperability

SnapMirror active sync is compatible with numerous operating systems, application hosts, and other features in ONTAP.



For specific supportability and interoperability details not covered here, consult the Interoperability Matrix Tool (IMT).

Application hosts

SnapMirror active sync support applications hosts including Hyper-V, Red Hat Enterprise Linux (RHEL), VMware, VMware vSphere Metro Storage Cluster (vMSC), Windows Server, and, beginning with ONTAP 9.14.1, Windows Server Failover Cluster.

Operating systems

SnapMirror active sync is supported with numerous operating systems, including:

- AIX via PVR (beginning ONTAP 9.11.1)
- HP-UX (beginning ONTAP 9.10.1)

- Solaris 11.4 (beginning ONTAP 9.10.1)
- NVMe support with ESXi (beginning ONTAP 9.17.1)

ΑΙΧ

Beginning with ONTAP 9.11.1, AIX is supported with SnapMirror active sync via standard engineering Feature Policy Variance Request (FPVR) with the agreement that the following stipulations are understood:

- SnapMirror active sync can provide zero RPO data protection, but the failover process with AIX requires additional steps to recognize the path change. LUNs that are not part of a root volume group will experience an I/O pause until a cfgmgr command is run. This can be automated, and most applications will resume operations without further disruption.
- LUNs that are part of a root volume group should generally not be protected with SnapMirror active sync. It's not possible to run the cfgmgr command after a failover, meaning that a reboot is required to recognize the changes in SAN paths. You can still achieve zero RPO data protection of the root volume group, but failover will be disruptive.

Consult your NetApp account team for further information about SnapMirror active sync with AIX.

HP-UX

Beginning with ONTAP 9.10.1, SnapMirror active sync for HP-UX is supported.

Automatic unplanned failover with HP-UX

An automatic unplanned failover (AUFO) event on the isolated master cluster may be caused by dual event failure when the connection between the primary and the secondary cluster is lost and the connection between the primary cluster and the mediator is also lost. This is considered a rare event, unlike other AUFO events.

- In this scenario, it might take more than 120 seconds for I/O to resume on the HP-UX host. Depending on the applications that are running, this might not lead to any I/O disruption or error messages.
- To remediate, you must restart applications on the HP-UX host that have a disruption tolerance of less than 120 seconds.

Solaris

Beginning with ONTAP 9.10.1, SnapMirror active sync supports Solaris 11.4.

To ensure the Solaris client applications are non-disruptive when an unplanned site failover switchover occurs in an SnapMirror active sync environment, modify the default Solaris OS settings. To configure Solaris with the recommended settings, see the Knowledge Base article Solaris Host support recommended settings in SnapMirror active sync.

ONTAP interoperability

SnapMirror active sync integrates with components of ONTAP to extends its data protection capabilities.

FabricPool

SnapMirror active sync supports source and destination volumes on FabricPool aggregates with tiering policies of None, Snapshot or Auto. SnapMirror active sync does not support FabricPool aggregates using a tiering policy of All.

Fan-out configurations

In fan-out configurations, your source volume can be mirrored to a SnapMirror active sync destination endpoint and to one or more SnapMirror asynchronous relationships.



SnapMirror active sync supports fan-out configurations with the MirrorAllSnapshots policy and, beginning with ONTAP 9.11.1, the MirrorAndVault policy. Fan-out configurations are not supported in SnapMirror active sync with the XDPDefault policy.

Beginning with ONTAP 9.15.1, SnapMirror active sync supports automatic reconfiguration in the fan-out leg after a failover event. If the failover from the primary to the secondary site has succeeded, the tertiary site is automatically reconfigured to treat the secondary site as the source. The async fan-out leg can be a consistency group relationship or an independent volume relationship. The reconfiguration will work for either of the cases. Reconfiguration is triggered by either a planned or unplanned failover. Reconfiguration also occurs upon failback to the primary site.

For information about managing your fan-out configuration in earlier releases of ONTAP, see resume protection in the fan-out configuration.

NDMP restore

Beginning with ONTAP 9.13.1, you can use NDMP to copy and restore data with SnapMirror active sync. Using NDMP allows you to move data onto the SnapMirror active sync source to complete a restore without pausing protection. This is particularly useful in fan-out configurations.

SnapCenter

SnapMirror active sync is supported with SnapCenter beginning with SnapCenter 5.0. SnapCenter enables the creation of snapshots that can be used to protect and recover applications and virtual machines, enabling always available storage solutions with application-level granularity.

SnapRestore

SnapMirror active sync supports partial and single file SnapRestore.

Single file SnapRestore

Beginning with ONTAP 9.11.1, single-file SnapRestore is supported for SnapMirror active sync volumes. You can restore a single file from a snapshot replicated from the SnapMirror active sync source to the destination. Because volumes can contain one or more LUNs, this feature helps you implement a less disruptive restore operation, granularly restoring a single LUN without disrupting the other LUNs. Single File SnapRestore has two options: in-place and out-of-place.

Partial file SnapRestore

Beginning in ONTAP 9.12.1, partial LUN restore is supported for SnapMirror active sync volumes. You can restore a data from application-created snapshots that have been replicated between the SnapMirror active sync source (volume) and the destination (snapshot) volumes. Partial LUN or file restore may be necessary if you need to restore a database on a host that stores multiple databases on the same LUN. Using this functionality requires you to know the starting byte offset of the data and byte count.

Large LUNs and large volumes

Support for large LUNs and large volumes (greater than 100 TB) depends on the version of ONTAP you are using and your platform.

ONTAP 9.12.1P2 and later

 For ONTAP 9.12.1 P2 and later, SnapMirror active sync supports Large LUNs and large volumes greater than 100 TB on ASA and AFF (A-Series and C-Series). Primary and secondary clusters must be of the same type: either ASA or AFF. Replication from AFF A-Series to AFF C-Series and vice versa is supported.



For ONTAP Releases 9.12.1P2 and later, you must ensure that both the primary and secondary clusters are either All-Flash SAN Arrays (ASA) or All Flash Array (AFF), and that they both have ONTAP 9.12.1 P2 or later installed. If the secondary cluster is running a version earlier than ONTAP 9.12.1P2 or if the array type is not the same as primary cluster, the synchronous relationship can go out of sync if the primary volume grows larger than 100 TB.

ONTAP 9.9.1 - 9.12.1P1

• For ONTAP releases between ONTAP 9.9.1 and 9.12.1 P1 (inclusive), Large LUNs and large volumes greater than 100TB are supported only on All-Flash SAN Arrays. Replication from AFF A-Series to AFF C-Series and vice versa is supported.



For ONTAP releases between ONTAP 9.9.1 and 9.12.1 P2, you must ensure that both the primary and secondary clusters are All-Flash SAN Arrays, and that they both have ONTAP 9.9.1 or later installed. If the secondary cluster is running a version earlier than ONTAP 9.9.1 or if it is not an All-Flash SAN Array, the synchronous relationship can go out of sync if the primary volume grows larger than 100 TB.

More information

• How to configure an AIX host for SnapMirror active sync

Object limits for ONTAP SnapMirror active sync

When preparing to use SnapMirror active sync, be aware of the following object limits.

Consistency groups in a cluster

Consistency group limits for a cluster with SnapMirror active sync are calculated based on relationships and depend on the version of ONTAP used. Limits are platform-independent.

ONTAP version	Maximum number of relationships
ONTAP 9.11.1 and later	50*
ONTAP 9.10.1	20
ONTAP 9.9.1	5

* Beginning with ONTAP 9.16.1, SnapMirror active sync supports four-node clusters in symmetric active/active configurations. In a four-node cluster, 100 consistency groups are supported.

Volumes per consistency group

The maximum number of volumes per consistency group with SnapMirror active sync is platform independent.

ONTAP version	Maximum number of volumes supported in a consistency group relationship
ONTAP 9.15.1 and later	80
ONTAP 9.10.1-9.14.1	16
ONTAP 9.9.1	12

Volumes

Volume limits in SnapMirror active sync are calculated based on the number of endpoints, not the number of relationships. A consistency group with 12 volumes contributes 12 endpoints on both the primary and secondary cluster. Both SnapMirror active sync and SnapMirror synchronous relationships contribute to the total number of endpoints.



These limits apply to FAS, AFF, and ASA systems. If you have an ASA r2 system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, or ASA A20), see ASA r2 documentation.

The maximum endpoints per platform are included in the following table.

Platform	Endpoints per sync	HA for SnapM	irror active	Overall sync a endpoints per	nd SnapMirror HA	active sync
	ONTAP 9.11.1 and later	ONTAP 9.10.1	ONTAP 9.9.1	ONTAP 9.11.1 and later	ONTAP 9.10.1	ONTAP 9.9.1
AFF	400*	200	60	400	200	80
ASA	400*	200	60	400	200	80

* Beginning with ONTAP 9.16.1, SnapMirror active sync supports four-node clusters in symmetric active/active configurations. The total limit for a four-node cluster is 800 endpoints.

SAN object limits

SAN object limits are included in the following table. The limits apply regardless of the platform.

Object in a SnapMirror active sync relationship	Count
LUNs per volume	 256 (ONTAP 9.9.1 - ONTAP 9.15.0) 512 (ONTAP 9.15.1 and later)
Number of unique LUNs, namespaces, or storage units per 2 x 2 SnapMirror active sync solution	4,096
Number of unique LUNs, namespaces, or storage units per 4 x 4 SnapMirror active-sync solution (available beginning with ONTAP 9.16.1)	6,144
LIFs per SVM (with at least one volume in a SnapMirror active sync relationship)	256
Inter-cluster LIFs per node	4
Inter-cluster LIFs per cluster	8

NVMe object limits

Beginning with ONTAP 9.17.1, SnapMirror active sync supports the NVMe protocol. NVMe object limits are included in the following table.

Maximum objects in a SnapMirror active sync relationship	Count
Number of namespace maps per node	4K
Cluster size	2 nodes
Number of consistency groups per HA pair	50
Number of volumes in a single NVMe SnapMirror active sync consistency group	80
Number of volumes in an HA pair	400
NVMe subsystems per consistency group	16
Namespace maps per consistency group	256

Related information

- Hardware Universe
- Consistency group limits

Configure

Configure ONTAP clusters for SnapMirror active sync

SnapMirror active sync uses peered clusters to protect your data in the event of a failover scenario. Before you configure ONTAP Mediator or ONTAP Cloud Mediator for SnapMirror active sync, you must first ensure the cluster is configured correctly.

Before you begin

Before you configure ONTAP Mediator or ONTAP Cloud Mediator, you should confirm the following:

1. A cluster peering relationship exists between the clusters.



The default IPspace is required by SnapMirror active sync for cluster peer relationships. A custom IPspace isn't supported.

Creating a cluster peer relationship

2. The SVMs are created on each cluster.

Creating an SVM

3. A peer relationship exists between the SVMs on each cluster.

Creating an SVM peering relationship

4. The volumes exist for your LUNs.

Creating a volume

5. At least one SAN LIF is created on each node in the cluster.

Considerations for LIFs in a cluster SAN environment

Creating a LIF

6. The necessary LUNs are created and mapped to an igroup, which is used to map LUNs to the initiator on the application host.

Create LUNs and map igroups

7. The application host is re-scanned to discover any new LUNs.

Configure the ONTAP Mediator for SnapMirror active sync

SnapMirror active sync uses peered clusters to protect your data in the event of a failover scenario. ONTAP Mediator is a key resource that enables business continuity by monitoring the health of each cluster. To configure SnapMirror active sync, you must first install ONTAP Mediator and verify that your primary and secondary clusters are configured properly.

Once you have installed ONTAP Mediator and configured your clusters, initialize ONTAP Mediator for SnapMirror active sync using self-signed certificates. You must then create, initialize, and map the consistency group for SnapMirror active sync.

ONTAP Mediator

ONTAP Mediator provides a persistent and fenced store for high availability (HA) metadata used by the ONTAP clusters in a SnapMirror active sync relationship. Additionally, ONTAP Mediator provides a synchronous node health query functionality to aid in quorum determination and serves as a ping proxy for controller liveliness detection.

Each cluster peer relationship can only be associated with a single ONTAP Mediator instance. HA Mediator instances aren't supported. When a cluster is in several peer relationships with other clusters, the following ONTAP Mediator options are available:

- If SnapMirror active sync is configured on each relationship, each cluster peer relationship can have its own unique ONTAP Mediator instance.
- The cluster can use the same ONTAP Mediator instance for all peer relationships.

For example, if cluster B has a peer relationship with cluster A, cluster C, and cluster D, all three cluster peer relationships can have a unique associated ONTAP Mediator instance when SnapMirror active sync is configured on each relationship. Alternatively, cluster B can use the same ONTAP Mediator instance for all three peer relationships. In this scenario, the same instance of ONTAP Mediator is listed three times for the cluster.

Beginning with ONTAP 9.17.1, you can configure ONTAP Cloud Mediator to monitor the health of your cluster in a SnapMirror active sync configuration, however, you cannot use both Mediators at the same time.

Prerequisites for ONTAP Mediator

• ONTAP Mediator includes its own set of prerequisites. You must meet these prerequisites before installing ONTAP Mediator.

For more information, see Prepare to install the ONTAP Mediator service.

• By default, the ONTAP Mediator provides service through TCP port 31784. You should make sure that port 31784 is open and available between the ONTAP clusters and the ONTAP Mediator.

Install ONTAP Mediator and confirm cluster configuration

Perform each of the following steps to install ONTAP Mediator and verify the cluster configuration. For each step, you should confirm that the specific configuration has been performed. Each step includes a link to the specific procedure that you need to follow.

Steps

1. Install ONTAP Mediator before verifying that your source and destination clusters are configured correctly.

Prepare to install or upgrade ONTAP Mediator

2. Confirm that a cluster peering relationship exists between the clusters.



The default IPspace is required by SnapMirror active sync for cluster peer relationships. A custom IPspace isn't supported.

Configure ONTAP clusters for SnapMirror active sync

Initialize ONTAP Mediator for SnapMirror active sync using self-signed certificates

Once you have installed ONTAP Mediator and confirmed you cluster configuration, you must initialize ONTAP Mediator for cluster monitoring. You can initialize ONTAP Mediator using System Manager or the ONTAP CLI.

System Manager

With System Manager, you can configure ONTAP Mediator for automated failover. You can also replace the self-signed SSL and CA with the third party validated SSL Certificate and CA if you have not already done so.



From ONTAP 9.14.1 through 9.8, SnapMirror active sync is referred to as SnapMirror Business Continuity (SM-BC).

ONTAP Mediator 1.9 and later

- 1. Navigate to **Protection > Overview > Mediator > Configure**.
- 2. Select **Add**, and enter the following ONTAP Mediator information:
 - IPv4 address
 - Username
 - Password
 - · Certificate
- 3. You can provide the Certificate input in two ways:
 - Option (a): Select Import to navigate to the intermediate.crt file and import it.
 - **Option (b)**: Copy the content of the intermediate.crt file and paste it in the **Certificate** field.

When all details are entered correctly, the provided certificate is installed on all the peer clusters.

Overview								
< Intercluster settings Network interfaces		Protected Volume pr Snapshots (l data rotection (local)			Alfine		→
0 1.2.3.5	Configure me	diator					×	
	IP address	User name	Password	Port	Cluster peers	Certificate		cted.
Cluster peers	Mediator IP	MediatorAccount Username		31784	Cluster Peer Name	BEGIN CERTIFICATE MIGOTCCBCGgAwlBAglUc4 SKmEzb+aawrTrURp.Jemid5b BdwDQYJKoZlhvcNAQEL BQAwgaMxCzAJBgNVBAYTAI VTMRNwEQYDVQQIDApDY	Import	ed up to cloud.
Mediator (?)								Protect for business continuity
Not configured.							Cancel	s you protect a consistency group with a zero recovery e objective.
configure							Close	e NetApp SnapCenter for application-consistent protection.
Storage VM peers Prefered storage vws e	1	Bucket pro SnapMirror (1) The chart Back up to c (1) The chart	Ditection (local or remote) t wasn't generated be cloud t wasn't generated be	acause no bucket acause no bucket	s were found. s were found.			

When the certificate addition is complete, ONTAP Mediator is added to the ONTAP cluster.

The following image demonstrates a successful ONTAP Mediator configuration:

PEEREI	CLUSTER NAME	
\odot	Peer Cluster Name	
Mec	liator 🕐	ţộ
Mec	liator ⑦	Ę

ONTAP Mediator 1.8 and earlier

- 1. Navigate to **Protection > Overview > Mediator > Configure**.
- 2. Select Add, and enter the following ONTAP Mediator information:
 - IPv4 address
 - Username
 - Password
 - · Certificate
- 3. You can provide the Certificate input in two ways:
 - **Option (a)**: Select **Import** to navigate to the ca.crt file and import it.
 - Option (b): Copy the content of the ca.crt file and paste it in the Certificate field.

When all details are entered correctly, the provided certificate is installed on all the peer clusters.

Overview							
< Intercluster settings Network interfaces		Protected data Volume protection Snapshots (local)					→
0 1.2.3.5	Configure mediator					X	
	IP address	User name	Password	Port	Cluster peers	Certificate	cted.
Cluster peers	Mediator IP	MediatorAccount Username		31784	Cluster Peer Name	BEGIN CERTIFICATE MIIGOTCCBCGgAwlBAglUc4 sKmEzb+aaw/TrURpJemid5b B4wDQYJKoZIhvcNAQEL BQAwgaMxCZAJBgNVBAYTAI VTMRMwEQYDQQIDApDY	ed up to cloud.
Mediator ③							Protect for business continuity
Not configured.						C	s you protect a consistency group with a zero recovery e objective.
Configure						Clo	se NetApp SnapCenter for application-consistent protection.
Storage VM peers PEERED STORADE VMS 0		Bucket pro SnapMirror (1) The chart Back up to c (1) The chart	otection (local or remote) t wasn't generated be cloud t wasn't generated be	ecause no bucket ecause no bucket	s were found. s were found.		

When the certificate addition is complete, ONTAP Mediator is added to the ONTAP cluster.

The following image demonstrates a successful ONTAP Mediator configuration:

	<u>*</u>
PEERED CLUSTER NAME	
Peer Cluster Name	
Mediator (?)	63
Mediator	
Peer Cluster Name	

CLI

You can initialize ONTAP Mediator from either the primary or secondary cluster using the ONTAP CLI. When you issue the mediator add command on one cluster, ONTAP Mediator is automatically added on the other cluster.

When using ONTAP Mediator to monitor a SnapMirror active sync relationship, ONTAP Mediator cannot be initialized in ONTAP without a valid self-signed or certificate authority (CA) certificate. You add a valid certificate to the certificate store for peered clusters. When using ONTAP Mediator to monitor MetroCluster IP systems, HTTPS isn't used after the initial configuration; therefore, certificates aren't required.

ONTAP Mediator 1.9 and later

- 1. Find the ONTAP Mediator CA certificate from the ONTAP Mediator Linux VM/host software installation location cd /opt/netapp/lib/ontap mediator/ontap mediator/server config.
- 2. Add a valid certificate authority to the certificate store on the peered cluster.

Example:

```
[root@ontap-mediator_config]# cat intermediate.crt
----BEGIN CERTIFICATE----
<certificate_value>
-----END CERTIFICATE-----
```

3. Add the ONTAP Mediator CA certificate to an ONTAP cluster. When prompted, insert the CA certificate obtained from ONTAP Mediator. Repeat the steps on all of the peer clusters:

security certificate install -type server-ca -vserver <vserver name>

```
[root@ontap-mediator ~]# cd
/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config
[root@ontap-mediator_config]# cat intermediate.crt
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----
```

```
C1_test_cluster::*> security certificate install -type server-ca
-vserver C1_test_cluster
Please enter Certificate: Press when done
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----
You should keep a copy of the CA-signed digital certificate for
future reference.
The installed certificate's CA and serial number for reference:
CA: ONTAP Mediator CA
serial: D86D8E4E87142XXX
The certificate's generated name for reference: ONTAPMediatorCA
C1 test cluster::*>
```

4. View the self-signed CA certificate installed using the generated name of the certificate:

security certificate show -common-name <common name>

Example:

Initialize ONTAP Mediator on one of the clusters. ONTAP Mediator is automatically added for the other cluster:

snapmirror mediator add -mediator-address <ip_address> -peer-cluster
<peer cluster name> -username user name

```
C1_test_cluster::*> snapmirror mediator add -mediator-address
1.2.3.4 -peer-cluster C2_test_cluster -username mediatoradmin
Notice: Enter the mediator password.
Enter the password: *****
Enter the password again: *****
```

6. Optionally, check the job ID status job show -id to verify if the SnapMirror Mediator add command is successful.

```
C1 test cluster::*> snapmirror mediator show
  This table is currently empty.
  C1 test cluster::*> snapmirror mediator add -peer-cluster
  C2 test cluster -type on-prem -mediator-address 1.2.3.4 -username
  mediatoradmin
  Notice: Enter the mediator password.
  Enter the password:
  Enter the password again:
  Info: [Job: 87] 'mediator add' job queued
  C1_test_cluster::*> job show -id 87
                          Owning
                                   Node State
  Job ID Name
                         Vserver
   _____
           ----- ------
                               _____
   _____
  87 mediator add C1 test cluster C2_test Running
  Description: Creating a mediator entry
  C1 test cluster::*> job show -id 87
                          Owning
  Job ID Name
                          Vserver
                                   Node
                                                    State
   _____ _ ____
   _____
  87 mediator add C1 test cluster C2 test Success
  Description: Creating a mediator entry
  C1 test cluster::*> snapmirror mediator show
  Mediator Address Peer Cluster Connection Status Quorum Status
  Type
   _____
  1.2.3.4 C2_test_cluster connected true
  on-prem
  C1 test cluster::*>
7. Check the status of the ONTAP Mediator configuration:
```

```
snapmirror mediator show
```

Mediator Address	Peer Cluster	Connection Status	Quorum Status
1.2.3.4	C2_test_cluster	connected	true

Quorum Status indicates whether the SnapMirror consistency group relationships are synchronized with ONTAP Mediator; a status of true indicates successful synchronization.

ONTAP Mediator 1.8 and earlier

- 1. Find the ONTAP Mediator CA certificate from the ONTAP Mediator Linux VM/host software installation location cd /opt/netapp/lib/ontap_mediator/ontap_mediator/server_config.
- 2. Add a valid certificate authority to the certificate store on the peered cluster.

Example:

```
[root@ontap-mediator_config]# cat ca.crt
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----
```

3. Add the ONTAP Mediator CA certificate to an ONTAP cluster. When prompted, insert the CA certificate obtained from the ONTAP Mediator. Repeat the steps on all of the peer clusters:

security certificate install -type server-ca -vserver <vserver name>

```
[root@ontap-mediator ~]# cd
/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config
```

```
[root@ontap-mediator_config]# cat ca.crt
----BEGIN CERTIFICATE----
<certificate_value>
-----END CERTIFICATE-----
```

```
C1_test_cluster::*> security certificate install -type server-ca
-vserver C1_test_cluster
Please enter Certificate: Press when done
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----
You should keep a copy of the CA-signed digital certificate for
future reference.
The installed certificate's CA and serial number for reference:
CA: ONTAP Mediator CA
serial: D86D8E4E87142XXX
The certificate's generated name for reference: ONTAPMediatorCA
C1 test cluster::*>
```

4. View the self-signed CA certificate installed using the generated name of the certificate:

security certificate show -common-name <common name>

Example:

Initialize ONTAP Mediator on one of the clusters. ONTAP Mediator is automatically added for the other cluster:

snapmirror mediator add -mediator-address <ip_address> -peer-cluster
<peer cluster name> -username user name

```
C1_test_cluster::*> snapmirror mediator add -mediator-address
1.2.3.4 -peer-cluster C2_test_cluster -username mediatoradmin
Notice: Enter the mediator password.
Enter the password: *****
Enter the password again: *****
```

6. Optionally, check the job ID status job show -id to verify if the SnapMirror Mediator add command is successful.

```
C1 test cluster::*> snapmirror mediator show
  This table is currently empty.
  C1 test cluster::*> snapmirror mediator add -peer-cluster
  C2 test cluster -type on-prem -mediator-address 1.2.3.4 -username
  mediatoradmin
  Notice: Enter the mediator password.
  Enter the password:
  Enter the password again:
  Info: [Job: 87] 'mediator add' job queued
  C1_test_cluster::*> job show -id 87
                         Owning
                                   Node State
  Job ID Name
                         Vserver
   _____ __
            _____ ____
                               _____
   _____
  87 mediator add C1 test cluster C2_test Running
  Description: Creating a mediator entry
  C1 test cluster::*> job show -id 87
                         Owning
  Job ID Name
                                   Node
                                                    State
                         Vserver
   _____ ____
   _____
  87 mediator add C1 test cluster C2 test Success
  Description: Creating a mediator entry
  C1 test cluster::*> snapmirror mediator show
  Mediator Address Peer Cluster Connection Status Quorum Status
  Type
   _____
  1.2.3.4 C2_test_cluster connected true
  on-prem
  C1 test cluster::*>
7. Check the status of the ONTAP Mediator configuration:
```

```
snapmirror mediator show
```

Mediator Address	Peer Cluster	Connection Status	Quorum Status
1.2.3.4	C2_test_cluster	connected	true

Quorum Status indicates whether the SnapMirror consistency group relationships are synchronized with ONTAP Mediator; a status of true indicates successful synchronization.

Re-initialize ONTAP Mediator with third-party certificates

You might need to re-initialize ONTAP Mediator. There might be situations that require the re-initialization of ONTAP Mediator such as a change in the ONTAP Mediator IP address, certificate expiration, and so on.

The following procedure illustrates the re-initialization of ONTAP Mediator for a specific case when a selfsigned certificate needs to be replaced by a third-party certificate.

About this task

You need to replace the SnapMirror active sync cluster's self-signed certificates with third-party certificates, remove the ONTAP Mediator configuration from ONTAP, and then add ONTAP Mediator.

System Manager

With System Manager, you need to remove the ONTAP Mediator version configured with the old self-signed certificate from the ONTAP cluster and re-configure the ONTAP cluster with the new third-party certificate.

Steps

1. Select the menu options icon and select **Remove** to remove ONTAP Mediator.



This step does not remove the self-signed server-ca from the ONTAP cluster. NetApp recommends navigating to the **Certificate** tab and removing it manually before performing the next step below to add a third-party certificate:

	User name	Password	Port	Cluster peers	Certificate	
Mediator IP	1		31784	Peer Cluster Name	I	
nove						

2. Add ONTAP Mediator again with the correct certificate.
ONTAP Mediator is now configured with the new third-party self-signed certificate.

Overview								
< Intercluster settings Network interfaces	Protected data Volume protection Snapshots (local)							.→
⊘ 1.2.3.5	Configure mediator						×	
	IP address	User name	Password	Port	Cluster peers	Certificate		cted.
Cluster peers	Mediator IP	MediatorAccount Username		31784	Cluster Peer Name	BGIN CERTIFICATE MIGOTCCBCGgAwiBAgiUc4 sKmEzb+aawrTrURpJemid5b BdwDQYJKoZihveNAQEL BQAwgaNkcZJBgNVBAYTAI VTMRMwEQYDVQQIDApDY	Import	rd up to cloud.
Mediator (?)								Protect for business continuity
Not configured.							Cancel	s you protect a consistency group with a zero recovery e objective.
							Close	e NetApp SnapCenter for application-consistent protection,
Storage VM peers	:	Bucket pro SnapMirror I (1) The chart Back up to c (1) The chart	otection (local or remote) wasn't generated bec loud wasn't generated bec	ause no bucket: ause no bucket:	s were found. s were found.			

CLI

You can re-initialize ONTAP Mediator from either the primary or secondary cluster by using the ONTAP CLI to replace the self-signed certificate with the third-party certificate.

ONTAP Mediator 1.9 and later

1. Remove the self-signed intermediate.crt installed earlier when you used self-signed certificates for all clusters. In the example below, there are two clusters:

Example:

```
C1_test_cluster::*> security certificate delete -vserver
C1_test_cluster -common-name ONTAPMediatorCA
2 entries were deleted.
C2_test_cluster::*> security certificate delete -vserver
C2_test_cluster -common-name ONTAPMediatorCA *
2 entries were deleted.
```

2. Remove the previously configured ONTAP Mediator from the SnapMirror active sync cluster using -force true:

Example:

```
C1 test cluster::*> snapmirror mediator show
Mediator Address Peer Cluster Connection Status Quorum Status
                   -----
---- ---
1.2.3.4
              C2 test cluster connected
                                                 true
C1 test cluster::*> snapmirror mediator remove -mediator-address
1.2.3.4 -peer-cluster C2 test cluster -force true
Warning: You are trying to remove the ONTAP Mediator configuration
with force. If this configuration exists on the peer cluster, it
could lead to failure of a SnapMirror failover operation. Check if
this configuration
        exists on the peer cluster C2 test cluster and remove it as
well.
Do you want to continue? \{y|n\}: y
Info: [Job 136] 'mediator remove' job queued
C1 test cluster::*> snapmirror mediator show
This table is currently empty.
```

3. Refer to the steps described in Replace self-signed certificates with trusted third-party certificates for instructions on how to obtain certificates from a subordinate CA, referred to as intermediate.crt. Replace self-signed certificates with trusted third-party certificates



The intermediate.crt has certain properties that it derives from the request that need to be sent to the PKI authority, defined in the file /opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/open ssl ca.cnf

4. Add the new third-party ONTAP Mediator CA certificate intermediate.crt from the ONTAP Mediator Linux VM/host software installation location:

Example:

```
[root@ontap-mediator ~]# cd
/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config
[root@ontap-mediator_config]# cat intermediate.crt
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----
```

5. Add the intermediate.crt file to the peered cluster. Repeat this step for all peer clusters:

Example:

```
C1_test_cluster::*> security certificate install -type server-ca
-vserver C1_test_cluster
Please enter Certificate: Press when done
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----
You should keep a copy of the CA-signed digital certificate for
future reference.
The installed certificate's CA and serial number for reference:
CA: ONTAP Mediator CA
serial: D86D8E4E87142XXX
The certificate's generated name for reference: ONTAPMediatorCA
C1_test_cluster::*>
```

6. Remove the previously configured ONTAP Mediator from the SnapMirror active sync cluster:

Example:

C1_test_cluster::*> snapmirror mediator show Mediator Address Peer Cluster Connection Status Quorum Status 1.2.3.4 C2_test_cluster connected true C1_test_cluster::*> snapmirror mediator remove -mediator-address 1.2.3.4 -peer-cluster C2_test_cluster Info: [Job 86] 'mediator remove' job queued C1_test_cluster::*> snapmirror mediator show This table is currently empty.

7. Add ONTAP Mediator again:

Example:

Quorum Status indicates whether the SnapMirror consistency group relationships are synchronized with the mediator; a status of true indicates successful synchronization.

ONTAP Mediator 1.8 and earlier

1. Remove the self-signed ca.crt installed earlier when you used self-signed certificates for all clusters. In the example below, there are two clusters:

Example:

```
C1_test_cluster::*> security certificate delete -vserver
C1_test_cluster -common-name ONTAPMediatorCA
2 entries were deleted.
C2_test_cluster::*> security certificate delete -vserver
C2_test_cluster -common-name ONTAPMediatorCA *
2 entries were deleted.
```

2. Remove the previously configured ONTAP Mediator from the SnapMirror active sync cluster using -force true:

Example:

```
C1 test cluster::*> snapmirror mediator show
Mediator Address Peer Cluster Connection Status Quorum Status
_____
1.2.3.4
               C2 test cluster connected
                                                true
C1 test cluster::*> snapmirror mediator remove -mediator-address
1.2.3.4 -peer-cluster C2 test cluster -force true
Warning: You are trying to remove the ONTAP Mediator configuration
with force. If this configuration exists on the peer cluster, it
could lead to failure of a SnapMirror failover operation. Check if
this configuration
        exists on the peer cluster C2 test cluster and remove it as
well.
Do you want to continue? \{y|n\}: y
Info: [Job 136] 'mediator remove' job queued
C1 test cluster::*> snapmirror mediator show
This table is currently empty.
```

3. Refer to the steps described in Replace self-signed certificates with trusted third-party certificates for instructions on how to obtain certificates from a subordinate CA, referred to as ca.crt. Replace self-signed certificates with trusted third-party certificates



The ca.crt has certain properties that it derives from the request that need to be sent to the PKI authority, defined in the file /opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/open ssl_ca.cnf

4. Add the new third-party ONTAP Mediator CA certificate ca.crt from the ONTAP Mediator Linux VM/host software installation location:

Example:

```
[root@ontap-mediator ~]# cd
/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config
[root@ontap-mediator_config]# cat ca.crt
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----
```

5. Add the intermediate.crt file to the peered cluster. Repeat this step for all peer clusters:

Example:

```
Cl_test_cluster::*> security certificate install -type server-ca
-vserver Cl_test_cluster
Please enter Certificate: Press when done
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----
You should keep a copy of the CA-signed digital certificate for
future reference.
The installed certificate's CA and serial number for reference:
CA: ONTAP Mediator CA
serial: D86D8E4E87142XXX
The certificate's generated name for reference: ONTAPMediatorCA
Cl_test_cluster::*>
```

6. Remove the previously configured ONTAP Mediator from the SnapMirror active sync cluster:

Example:

```
C1_test_cluster::*> snapmirror mediator show

Mediator Address Peer Cluster Connection Status Quorum Status

1.2.3.4 C2_test_cluster connected true

C1_test_cluster::*> snapmirror mediator remove -mediator-address

1.2.3.4 -peer-cluster C2_test_cluster

Info: [Job 86] 'mediator remove' job queued

C1_test_cluster::*> snapmirror mediator show

This table is currently empty.
```

7. Add ONTAP Mediator again:

Example:

Quorum Status indicates whether the SnapMirror consistency group relationships are synchronized with the mediator; a status of true indicates successful synchronization.

Related information

- job show
- · security certificate delete
- security certificate install
- security certificate show
- snapmirror mediator add
- snapmirror mediator remove
- · snapmirror mediator show

Prepare to configure ONTAP Cloud Mediator

Before you configure ONTAP Cloud Mediator, you must ensure that the prerequisites are met.

Firewall requirements

The firewall setting on the domain controller must allow HTTPS traffic to api.bluexp.netapp.com from both clusters.

Proxy server requirements

If you use proxy servers for SnapMirror active sync, ensure the proxy servers are created and you have the following proxy server information:

- HTTPS proxy IP
- Port
- Username
- Password

Latency

The recommended ping latency between the BlueXP cloud server and SnapMirror active sync cluster peers is less than 200 ms.

Root CA certificates

Check the cluster for certificates

ONTAP comes with well-known root CA certificates pre-installed so in most cases you do not need to install the BlueXP server's root CA certificate. Before you begin the ONTAP Cloud Mediator configuration, you can check the cluster to verify that the certificates exist:

Example:

```
C1_cluster% openssl s_client -showcerts -connect
api.bluexp.netapp.com:443|egrep "s:|i:"
depth=2 C = US, O = DigiCert Inc, OU = www.digicert.com, CN = DigiCert
Global Root G2
verify return:1
depth=1 C = US, O = Microsoft Corporation, CN = Microsoft Azure RSA TLS
Issuing CA 04
verify return:1
depth=0 C = US, ST = WA, L = Redmond, O = Microsoft Corporation, CN =
*.azureedge.net
verify return:1
0 s:/C=US/ST=WA/L=Redmond/O=Microsoft Corporation/CN=*.azureedge.net
i:/C=US/O=Microsoft Corporation/CN=Microsoft Azure RSA TLS Issuing CA
04
1 s:/C=US/O=Microsoft Corporation/CN=Microsoft Azure RSA TLS Issuing CA
```

```
04
   i:/C=US/O=DigiCert Inc/OU=www.digicert.com/CN=DigiCert Global Root G2
 2 s:/C=US/O=DigiCert Inc/OU=www.digicert.com/CN=DigiCert Global Root G2
   i:/C=US/O=DigiCert Inc/OU=www.digicert.com/CN=DigiCert Global Root G2
<====
C1 cluster::> security certificate show -common-name DigiCert*
          Serial Number
                         Certificate Name
Vserver
                                                                 Type
_____ _
                              _____
_____
C1 cluster 0CE7E0EXXXXX46FE8FE560FC1BFXXXXX DigiCertAssuredIDRootCA
server-ca
    Certificate Authority: DigiCert Assured ID Root CA
         Expiration Date: Mon Nov 10 05:30:00 2031
C1 cluster 0B931C3XXXXX67EA6723BFC3AF9XXXXX DigiCertAssuredIDRootG2
server-ca
    Certificate Authority: DigiCert Assured ID Root G2
         Expiration Date: Fri Jan 15 17:30:00 2038
C1 cluster 0BA15AFXXXXXA0B54944AFCD24AXXXXX DigiCertAssuredIDRootG3
server-ca
    Certificate Authority: DigiCert Assured ID Root G3
         Expiration Date: Fri Jan 15 17:30:00 2038
C1 cluster 083BE05XXXXX46B1A1756AC9599XXXXX DigiCertGlobalRootCA server-ca
    Certificate Authority: DigiCert Global Root CA
         Expiration Date: Mon Nov 10 05:30:00 2031
C1 cluster 033AF1EXXXXA9A0BB2864B11D0XXXXX DigiCertGlobalRootG2 server-ca
    Certificate Authority: DigiCert Global Root G2
         Expiration Date: Fri Jan 15 17:30:00 2038
C1 cluster 055556BXXXXXA43535C3A40FD5AXXXXX DigiCertGlobalRootG3 server-ca
    Certificate Authority: DigiCert Global Root G3
         Expiration Date: Fri Jan 15 17:30:00 2038
C1 cluster 02AC5C2XXXXX409B8F0B79F2AE4XXXXX DigiCertHighAssuranceEVRootCA
server-ca
    Certificate Authority: DigiCert High Assurance EV Root CA
         Expiration Date: Mon Nov 10 05:30:00 2031
C1 cluster 059B1B5XXXXX2132E23907BDA77XXXXX DigiCertTrustedRootG4 server-
са
    Certificate Authority: DigiCert Trusted Root G4
         Expiration Date: Fri Jan 15 17:30:00 2038
```

Check proxy server for installed certificates

If you are using a proxy to connect to the ONTAP Cloud Mediator service in BlueXP, ensure that the proxy server's root CA certificates are installed in ONTAP:

Example:

```
C1_cluster% openssl s_client -showcerts -proxy <ip:port> -connect
api.bluexp.netapp.com:443 |egrep "s:|i:"
```

Download the CA certificate:

If necessary, you can download the root-CA certificates fom the certificate authority's website and install them on the clusters.

Example:

```
C1_cluster::> security certificate install -type server-ca -vserver
C1_cluster
C2_cluster::> security certificate install -type server-ca -vserver
C2_cluster
```

Configure the ONTAP Cloud Mediator for SnapMirror active sync

Beginning with ONTAP 9.17.1, you can use ONTAP Cloud Mediator to enable business continuity by monitoring the health of each cluster. ONTAP Cloud Mediator is a cloud-based service. When you use ONTAP Cloud Mediator with SnapMirror active sync, you must first confirm that BlueXP services and client information are configured and ensure proper cluster peering.

As with ONTAP Mediator, ONTAP Cloud Mediator provides a persistent and fenced store for high availability (HA) metadata used by the ONTAP clusters in a SnapMirror active sync relationship. ONTAP Cloud Mediator provides a synchronous node health query functionality to aid in quorum determination and serves as a ping proxy for controller liveliness detection.

Before you begin

Before you configure ONTAP Cloud Mediator, you should confirm the following information:

• The cluster is configured.

Configure ONTAP clusters for SnapMirror active sync

• You have a valid BlueXP subscription.

Subscribe to BlueXP data services (standard mode)

• You have copied your BlueXP organization ID from the BlueXP console and created a BlueXP member service account to use when you configure ONTAP Cloud Mediator. When you create the service account, the organization must be set to the subscription where you configured the ONTAP Cloud Mediator. The category must be set to **Application**, and the role type must be **ONTAP Mediator Setup Role**. You must

save the client ID and client secret when you create the role.

Add BlueXP members and service accounts

Steps

You can add ONTAP Cloud Mediator using System Manager or the ONTAP CLI.

System Manager

- 1. Navigate to **Protection > Overview > Mediator** and select **Add**.
- 2. In the Add a mediator window, select Cloud as the mediator type and enter the following information:
 - BlueXP organization ID
 - BlueXP client ID
 - BlueXP client secret
- 3. Select the cluster peer.
- 4. If you are using an HTTP proxy and it's not already configured, enter the HTTP proxy information for the local and remote hosts.

It's recommended that you use a different proxy server for each cluster peer.

- 5. Optional: If a root CA certificate needs to be installed in ONTAP, especially when using a proxy server, paste the certificate in the text box provided.
- 6. Select Add.
- 7. Navigate to **Protection > Overview** and check the status of the relationship between the SnapMirror active sync clusters and ONTAP Cloud Mediator.

CLI

1. Configure ONTAP Cloud Mediator:

```
snapmirror mediator add -peer-cluster <peerClusterName> -type cloud -bluexp
-org-id <BlueXP Organization ID> -service-account-client-id <Service
Account Client ID> -use-http-proxy-local <true|false> -use-http-proxy
-remote <true|false>
```

2. Check ONTAP Cloud Mediator status:

snapmirror mediator show

Example:

```
C1_cluster::> snapmirror mediator show
Mediator Address Peer Cluster Connection Status Quorum Status
Type
-------
0.0.0.0 C2_cluster connected true
cloud
```

Protect with ONTAP SnapMirror active sync

SnapMirror active sync offers asymmetric protection and, beginning with ONTAP 9.15.1, symmetric active/active protection.

Configure asymmetric protection

Configuring asymmetric protection using SnapMirror active sync involves selecting LUNs on the ONTAP source cluster and adding them to a consistency group.

Before you begin

- You must have a SnapMirror synchronous license.
- You must be a cluster or storage VM administrator.
- All constituent volumes in a consistency group must be in a single storage VM (SVM).
 - LUNs can reside on different volumes.
- The source and destination cluster cannot be the same.
- You cannot establish SnapMirror active sync consistency group relationships across ASA clusters and non-ASA clusters.
- The default IPspace is required by SnapMirror active sync for cluster peer relationships. Custom IPspace is not supported.
- The name of the consistency group must be unique.
- The volumes on the secondary (destination) cluster must be type DP.
- The primary and the secondary SVMs must be in a peered relationship.

Steps

You can configure a consistency group using the ONTAP CLI or System Manager.

Beginning with ONTAP 9.10.1, ONTAP offers a consistency group endpoint and menu in System Manager, offering additional management utilities. If you are using ONTAP 9.10.1 or later, see Configure a consistency group then configure protection to create a SnapMirror active sync relationship.



From ONTAP 9.14.1 through 9.8, SnapMirror active sync is referred to as SnapMirror Business Continuity (SM-BC).

System Manager

- 1. On the primary cluster, navigate to **Protection > Overview > Protect for Business Continuity > Protect LUNs**.
- 2. Select the LUNs you want to protect and add them to a protection group.
- 3. Select the destination cluster and SVM.
- 4. Initialize relationship is selected by default. Click Save to begin protection.
- 5. Go to **Dashboard > Performance** to verify IOPS activity for the LUNs.
- 6. On the destination cluster, use System Manager to verify that the protection for business continuity relationship is in sync: **Protection > Relationships**.

CLI

1. Create a consistency group relationship from the destination cluster. destination::> snapmirror create -source-path source-path -destination-path destination-path -cg-item-mappings volume-paths -policy policy-name

You can map up to 12 constituent volumes using the cg-item-mappings parameter on the snapmirror create command.

The following example creates two consistency groups: cg_src_ on the source with `vol1 and vol2 and a mirrored destination consistency group, cg_dst.

```
destination::> snapmirror create -source-path vs1_src:/cg/cg_src
-destination-path vs1_dst:/cg/cg_dst -cg-item-mappings
vol src1:@vol dst1,vol src2:@vol dst2 -policy AutomatedFailOver
```

2. From the destination cluster, initialize the consistency group.

```
destination::>snapmirror initialize -destination-path destination-
consistency-group
```

3. Confirm that the initialization operation completed successfully. The status should be InSync.

snapmirror show

4. On each cluster, create an igroup so you can map LUNs to the initiator on the application host. lun igroup create -igroup name -protocol fcp|iscsi -ostype os -initiator initiator_name

Learn more about lun igroup create in the ONTAP command reference.

5. On each cluster, map LUNs to the igroup:

lun map -path path_name -igroup igroup_name

6. Verify the LUN mapping completed successfully with the lun map command. Then, you can discover the new LUNs on the application host.

Configure symmetric active/active protection

You can establish symmetric protection using System Manager or the ONTAP CLI. In both interfaces, there are different steps for uniform and non-uniform configurations.

Before you begin

- Both clusters must be running ONTAP 9.15.1 or later.
- Symmetric active/active configurations require the AutomatedFailoverDuplex protection policy. Alternately, you can create a custom SnapMirror policy provided the -type is automated-failoverduplex.
- In ONTAP 9.15.1, symmetric active/active is only supported on 2-node clusters.
- Beginning with ONTAP 9.16.1 GA, SnapMirror active sync supports symmetric active/active configurations on four-node clusters.
 - To use SnapMirror active sync on a four-node cluster, you must be running ONTAP 9.16.1 GA or later.
 - Before deploying a four-node configuration, you must create a cluster peer relationship.
 - Review the limits for four-node clusters.
 - If you revert to a two-node cluster, you must remove the SnapMirror active sync relationships from the cluster before reverting.
 - You can use the four-node configuration to upgrade storage and controllers. This process is nondisruptive and expands the cluster while moving volumes into the new nodes. For more information, see refresh a cluster.
- Beginning with ONTAP 9.17.1, you can configure symmetric active/active protection on NVMe namespaces only when both clusters are running ONTAP 9.17.1 or later.

Configure symmetric active/active protection using a SCSI SnapMirror active sync configuration

Steps

You can use System Manager or the ONTAP CLI to configure symmetric active/active protection using SCSI protocol host mappings.

System Manager

Steps for a uniform configuration

- 1. On the primary site, create a consistency group using new LUNs.
 - a. When creating the consistency group, specify host initiators to create igroups.
 - b. Select the checkbox to Enable SnapMirror then choose the AutomatedFailoverDuplex policy.
 - c. In the dialog box that appears, select the **Replicate initiator groups** checkbox to replicate igroups. In **Edit proximity settings**, set proximal SVMs for your hosts.
 - d. Select Save.

Steps for a non-uniform configuration

- 1. On the primary site, create a consistency group using new LUNs.
 - a. When creating the consistency group, specify host initiators to create igroups.
 - b. Select the checkbox to Enable SnapMirror then choose the AutomatedFailoverDuplex policy.
 - c. Select **Save** to create the LUNs, consistency group, igroup, SnapMirror relationship, and igroup mapping.
- 2. On the secondary site, create an igroup and map the LUNs.
 - a. Navigate to Hosts > SAN Initiator Groups.
 - b. Select +Add to create a new igroup.
 - c. Provide a Name, select the Host Operating System, then choose Initiator Group Members.
 - d. Select **Save** to initialize the relationship.
- 3. Map the new igroup to the destination LUNs.
 - a. Navigate to Storage > LUNs.
 - b. Select all the LUNs to map to the igroup.
 - c. Select More then Map to Initiator Groups.

CLI

Steps for a uniform configuration

1. Create a new SnapMirror relationship grouping all the volumes in the application. Ensure you designate the AutomatedFailOverDuplex policy to establish bidirectional sync replication.

```
snapmirror create -source-path <source_path> -destination-path
<destination_path> -cg-item-mappings <source_volume:@destination_volume>
-policy AutomatedFailOverDuplex
```

Example:

The following example creates two consistency groups: cg_src on the source with vol1 and vol2, and a mirrored consistency group on the destination, cg_dst.

```
destination::> snapmirror create -source-path vs1_src:/cg/cg_src
-destination-path vs1_dst:/cg/cg_dst -cg-item-mappings
vol src1:@vol dst1,vol src2:@vol dst2 -policy AutomatedFailOver
```

- 2. Initialize the SnapMirror relationship: snapmirror initialize -destination-path <destination-consistency-group>
- 3. Confirm the operation has succeeded by waiting for the Mirrored State to show as SnapMirrored and the Relationship Status as Insync.

```
snapmirror show -destination-path <destination path>
```

- 4. On your host, configure host connectivity with access to each cluster according to your needs.
- 5. Establish the igroup configuration. Set the preferred paths for initiators on the local cluster. Specify the option to replicate the configuration to the peer cluster for inverse affinity.

```
SiteA::> igroup create -vserver <svm_name> -ostype <os_type> -igroup
<igroup name> -replication-peer <peer svm name> -initiator <host>
```



Beginning with ONTAP 9.16.1, use the -proximal-vserver local parameter in this command.

```
SiteA::> igroup add -vserver <svm_name> -igroup <igroup_name> -ostype
<os type> -initiator <host>
```



Beginning with ONTAP 9.16.1, use the -proximal-vserver peer parameter in this command.

- 6. From the host, discover the paths and verify the hosts have an active/optimized path to the storage LUN from the preferred cluster.
- 7. Deploy the application and distribute the VM workloads across clusters to achieve the required load balancing.

Steps for a non-uniform configuration

1. Create a new SnapMirror relationship grouping all the volumes in the application. Ensure you designate the AutomatedFailOverDuplex policy to establish bidirectional sync replication.

```
snapmirror create -source-path <source_path> -destination-path
<destination_path> -cg-item-mappings <source_volume:@destination_volume>
-policy AutomatedFailOverDuplex
```

Example:

The following example creates two consistency groups: cg_src on the source with vol1 and vol2, and a mirrored consistency group on the destination, cg_dst.

```
destination::> snapmirror create -source-path vs1_src:/cg/cg_src
-destination-path vs1_dst:/cg/cg_dst -cg-item-mappings
vol src1:@vol dst1,vol src2:@vol dst2 -policy AutomatedFailOver
```

- 2. Initialize the SnapMirror relationship: snapmirror initialize -destination-path <destination-consistency-group>
- 3. Confirm the operation has succeeded by waiting for the Mirrored State to show as SnapMirrored and the Relationship Status as Insync.

```
snapmirror show -destination-path <destination_path>
```

- 4. On your host, configure host connectivity with access to each cluster according to your needs.
- 5. Establish the igroup configurations on both the source and destination clusters.

```
# primary site
SiteA::> igroup create -vserver <svm_name> -igroup <igroup_name> -initiator
<host_1_name_>
# secondary site
```

```
SiteB::> igroup create -vserver <svm_name> -igroup <igroup_name> -initiator
<host_2_name>
```

- 6. From the host, discover the paths and verify the hosts have an active/optimized path to the storage LUN from the preferred cluster.
- 7. Deploy the application and distribute the VM workloads across clusters to achieve the required load balancing.

Configure symmetric active/active protection using an NVMe SnapMirror active sync configuration

Before you begin

In addition to the requirements for configuring symmetric active/active protection, you should be aware of the supported and unsupported configurations when using the NVMe protocol.

- · Consistency groups can have one or more subsystem.
- Volumes within the consistency group can have namespace maps from multiple subsystems.
- Subsystems cannot have namespace maps that belong to more than one consistency group.
- Subsystems cannot have some namespace maps that belong to a consistency group and some namespace maps that do not belong to a consistency group.
- Subsystems must have namespace maps that are part of the same consistency group.

Steps

Beginning with ONTAP 9.17.1, you can use System Manager or the ONTAP CLI to create a consistency group and configure symmetric active/active protection using NVMe protocol host mappings.

System Manager

- 1. On the primary site, create a consistency group using new volumes or NVMe namespaces.
- 2. select +Add and choose Using new NVMe namespaces.
- 3. Enter the consistency group name.
- 4. Select More.
- 5. In the **Protection** section, select **Enable SnapMirror** then choose the AutomatedFailoverDuplex policy.
- 6. In the Host mapping section, choose either Existing NVMe subsystem or New NVMe subsystem.
- 7. Select In proximity to to change the proximal SVM. The source SVM is selected by default.
- 8. If necessary, add another NVMe subsystem.

CLI

1. Create a new SnapMirror relationship grouping all the volumes containing all NVMe namespaces used by the application. Ensure you designate the AutomatedFailOverDuplex policy to establish bidirectional sync replication.

```
snapmirror create -source-path <source_path> -destination-path
<destination_path> -cg-item-mappings <source_volume:@destination_volume>
-policy AutomatedFailOverDuplex
```

Example:

```
DST::> snapmirror create -source-path vs_src:/cg/cg_src_1
-destination-path vs_dst:/cg/cg_dst_1 -cg-item-mappings
vs_src_vol1:@vs_dst_vol1,vs_src_vol2:@vs_dst_vol2 -policy
AutomatedFailOverDuplex
```

2. Initialize the SnapMirror relationship:

snapmirror initialize -destination-path <destination-consistency-group>

Example:

DST::> snapmirror initialize -destination-path vs1:/cg/cg_dst_1

3. Confirm the operation has succeeded by waiting for the Mirrored State to show as SnapMirrored and the Relationship Status as Insync.

snapmirror show -destination-path <destination path>

The NVMe subsystems associated with the NVMe namespaces in the primary volumes are automatically replicated to the secondary cluster.

- 4. On your host, configure host connectivity with access to each cluster according to your needs.
- 5. Specify the SVM that is proximal to each of your hosts. This enables host access to the NVMe namespace using a path from the preferred cluster. This might be the SVM in the primary cluster *or*

the SVM in DR cluster.

The following command indicates that SVM VS_A is proximal to host H1 and set VS_A as the proximal SVM:

SiteA::> vserver nvme subsystem host add -subsystem ss1 -host-nqn <H1_NQN>
-proximal-vservers <VS_A>

The following command indicates that SVM VS_B is proximal to host H2 and sets VS_B as the proximal SVM:

```
SiteB::> vserver nvme subsystem host add -subsystem ss1 -host-nqn <H2_NQN>
-proximal-vservers <VS B>
```

- 6. From the host, discover the paths and verify the hosts have an active/optimized path to the storage from the preferred cluster.
- 7. Deploy the application and distribute the VM workloads across clusters to achieve the required load balancing.

Related information

- snapmirror create
- snapmirror initialize
- snapmirror show

Convert an existing ONTAP SnapMirror relationship to SnapMirror active sync relationship

If you've configured SnapMirror protection, you can convert the relationship to SnapMirror active sync. Beginning with ONTAP 9.15.1, you can convert the relationship to use symmetric active/active protection.

Convert an existing iSCSI or FC SnapMirror relationship to an asymmetric SnapMirror active sync relationship

If you have an existing iSCSI or FC SnapMirror synchronous relationship between a source and destination cluster, you can convert it to an asymmetric SnapMirror active sync relationship. This allows you to associate the mirrored volumes with a consistency group, ensuring zero RPO across a multi-volume workload. Additionally, you can retain existing SnapMirror snapshots if you need to revert to a point in time prior to establishing the SnapMirror active sync relationship.

About this task

- You must be a cluster and SVM administrator on the primary and secondary clusters.
- You cannot convert zero RPO to zero RTO sync by changing the SnapMirror policy.
- You must ensure the LUNs are unmapped before issuing the snapmirror create command.

If existing LUNs on the secondary volume are mapped and the AutomatedFailover policy is configured, the snapmirror create command triggers an error.

Before you begin

- A zero RPO SnapMirror synchrnous relationship must exist between the primary and secondary cluster.
- All LUNs on the destination volume must be unmapped before the zero RTO SnapMirror relationship can

be created.

 SnapMirror active sync only supports SAN protocols (not NFS/CIFS). Ensure no constituent of the consistency group is mounted for NAS access.

Steps

1. From the secondary cluster, perform a SnapMirror update on the existing relationship:

SiteB::>snapmirror update -destination-path vs1 dst:vol1

2. Verify that the SnapMirror update completed successfully:

SiteB::>snapmirror show

3. Pause each of the zero RPO synchronous relationships:

SiteB::>snapmirror quiesce -destination-path vs1_dst:vol1

SiteB::>snapmirror quiesce -destination-path vs1 dst:vol2

4. Delete each of the zero RPO synchronous relationships:

SiteB::>snapmirror delete -destination-path vs1 dst:vol1

SiteB::>snapmirror delete -destination-path vs1 dst:vol2

5. Release the source SnapMirror relationship but retain the common snapshots:

```
SiteA::>snapmirror release -relationship-info-only true -destination-path vs1 dst:vol1
```

SiteA::>snapmirror release -relationship-info-only true -destination-path vs1 dst:vol2

6. Create a zero RTO SnapMirror synchronous relationship:

SiteB::> snapmirror create -source-path vs1_src:/cg/cg_src -destination-path
vs1_dst:/cg/cg_dst -cg-item-mappings vol1:@vol1,vol2:@vol2 -policy
AutomatedFailover

7. Resynchronize the consistency group:

SiteB::> snapmirror resync -destination-path vs1 dst:/cg/cg dst

8. Rescan host LUN I/O paths to restore all paths to the LUNs.

Convert an existing iSCSI or FC SnapMirror relationship to symmetric active/active

Beginning with ONTAP 9.15.1, you can convert an existing iSCSI or FC SnapMirror relationship to a SnapMirror active sync symmetric active/active relationship.

Before you begin

- You must be running ONTAP 9.15.1 or later.
- A zero RPO SnapMirror synchrnous relationship must exist between the primary and secondary cluster.

- All LUNs on the destination volume must be unmapped before the zero RTO SnapMirror relationship can be created.
- SnapMirror active sync only supports SAN protocols (not NFS/CIFS). Ensure no constituent of the consistency group is mounted for NAS access.

Steps

1. From the secondary cluster, perform a SnapMirror update on the existing relationship:

SiteB::>snapmirror update -destination-path vs1 dst:vol1

2. Verify that the SnapMirror update completed successfully:

SiteB::>snapmirror show

3. Pause each of the zero RPO synchronous relationships:

SiteB::>snapmirror quiesce -destination-path vs1_dst:vol1

SiteB::>snapmirror quiesce -destination-path vs1_dst:vol2

4. Delete each of the zero RPO synchronous relationships:

SiteB::>snapmirror delete -destination-path vs1 dst:vol1

SiteB::>snapmirror delete -destination-path vs1 dst:vol2

5. Release the source SnapMirror relationship but retain the common snapshots:

```
SiteA::>snapmirror release -relationship-info-only true -destination-path vs1 dst:vol1
```

SiteA::>snapmirror release -relationship-info-only true -destination-path vs1 dst:vol2

6. Create a zero RTO SnapMirror synchronous relationship with the AutomatedFailoverDuplex policy:

SiteB::> snapmirror create -source-path vs1_src:/cg/cg_src -destination-path vs1_dst:/cg/cg_dst -cg-item-mappings vol1:@vol1,vol2:@vol2 -policy AutomatedFailoverDuplex

- 7. If the existing hosts are local the primary cluster, add the host to the secondary cluster and establish connectivity with respective access to each cluster.
- 8. On the secondary site, delete the LUN maps on the igroups associated with remote hosts.



Ensure the igroup does not contain maps for non-replicated LUNs.

SiteB::> lun mapping delete -vserver <svm name> -igroup <igroup> -path <>

9. On the primary site, modify the initiator configuration for existing hosts to set the proximal path for initiators on the local cluster.

SiteA::> igroup initiator add-proximal-vserver -vserver <svm name> -initiator

<host> -proximal-vserver <server>

10. Add a new igroup and initiator for the new hosts and set the host proximity for host affinity to its local site. Ennable igroup replication to replicate the configuration and invert the host locality on the remote cluster.

```
SiteA::> igroup modify -vserver vsA -igroup ig1 -replication-peer vsB
SiteA::> igroup initiator add-proximal-vserver -vserver vsA -initiator host2
-proximal-vserver vsB
```

- 11. Discover the paths on the hosts and verify the hosts have an Active/Optimized path to the storage LUN from the preferred cluster
- 12. Deploy the application and distribute the VM workloads across clusters.
- 13. Resynchronize the consistency group:

SiteB::> snapmirror resync -destination-path vs1_dst:/cg/cg_dst

14. Rescan host LUN I/O paths to restore all paths to the LUNs.

Related information

- snapmirror create
- snapmirror delete
- snapmirror quiesce
- snapmirror release
- snapmirror resync
- snapmirror show

Convert ONTAP SnapMirror active sync relationship type

Beginning with ONTAP 9.15.1, you can convert between types of SnapMirror active sync protection: from asymmetric to symmetric active/active and vice versa.

Convert to a symmetric active/active relationship

You can convert a iSCSI or FC SnapMirror active sync relationship with asymmetric protection to use symmetric active/active.

Before you begin

- Both clusters must be running ONTAP 9.15.1 or later.
- Symmetric active/active configurations require the AutomatedFailoverDuplex protection policy. Alternately, you can create a custom SnapMirror policy provided the -type is automated-failoverduplex.

System Manager

Steps for a uniform configuration

- 1. Remove the destination igroup:
 - a. On the destination cluster, navigate to **Hosts > SAN Initiator Groups**.
 - b. Select the igroup with the SnapMirror relationship, then **Delete**.
 - c. In the dialog box, select the Unmap the associated LUNs box then Delete.
- 2. Edit the SnapMirror active sync relationship.
 - a. Navigate to **Protection > Relationships**.
 - b. Select the kabob menu next to the relationship you want to modify then Edit.
 - c. Modify the Protection Policy to AutomatedFailoverDuplex.
 - d. Selecting AutoMatedFailoverDuplex prompts a dialog box to modify host proximity settings. For the initiators, select the appropriate option for **Initiator proximal to** then **Save**.
 - e. Select Save.
- 3. In the **Protection** menu, confirm the operation succeeded when the relationship displays as InSync.

Steps for a non-uniform configuration

- 1. Remove the destination igroup:
 - a. On the secondary site, navigate to **Hosts > SAN Initiator Groups**.
 - b. Select the igroup with the SnapMirror relationship, then Delete.
 - c. In the dialog box, select the **Unmap the associated LUNs** box then **Delete**.
- 2. Create a new igroup:
 - a. In the SAN Initiator Groups menu on the destination site, select Add.
 - b. Provide a Name, select the Host Operating System, then choose Initiator Group Members.
 - c. Select Save.
- 3. Map the new igroup to the destination LUNs.
 - a. Navigate to **Storage** > **LUNs**.
 - b. Select all the LUNs to map to the igroup.
 - c. Select More then Map to Initiator Groups.
- 4. Edit the SnapMirror active sync relationship.
 - a. Navigate to **Protection > Relationships**.
 - b. Select the kabob menu next to the relationship you want to modify then Edit.
 - c. Modify the Protection Policy to AutomatedFailoverDuplex.
 - d. Selecting AutoMatedFailoverDuplex initiates the option to modify host proximity settings. For the initiators, select the appropriate option for **Initiator proximal to** then **Save**.
 - e. Select Save.
- 5. In the **Protection** menu, confirm the operation succeeded when the relationship displays as InSync.

CLI

Steps for a uniform configuration

1. Modify the SnapMirror policy from AutomatedFailover to AutomatedFailoverDuplex:

```
snapmirror modify -destination-path <destination_path> -policy
AutomatedFailoverDuplex
```

2. Modifying the policy triggers a resync. Wait for the resync to complete and confirm the relationship is Insync:

snapmirror show -destination-path <destination path>

- 3. If the existing hosts are local the primary cluster, add the host to the second cluster and establish connectivity with respective access to each cluster.
- 4. On the secondary site, delete the LUN maps on the igroups associated with remote hosts.



Ensure the igroup does not contain maps for non-replicated LUNs.

SiteB::> lun mapping delete -vserver <svm_name> -igroup <igroup> -path <>

5. On the primary site, set the privilege level to advanced:

SiteA::> set -privilege advanced

6. Modify the initiator configuration for existing hosts to set the proximal path for initiators on the local cluster.

```
SiteA::*> igroup initiator add-proximal-vserver -vserver <svm_name>
-initiator <host> -proximal-vserver <server>
```



You can set the privilege level back to admin after you complete this step.

7. Add a new igroup and initiator for the new hosts and set the host proximity for host affinity to its local site. Ennable igroup replication to replicate the configuration and invert the host locality on the remote cluster.

```
SiteA::> igroup modify -vserver vsA -igroup ig1 -replication-peer vsB
SiteA::> igroup initiator add-proximal-vserver -vserver vsA -initiator
host2 -proximal-vserver vsB
```

- 8. Discover the paths on the hosts and verify the hosts have an Active/Optimized path to the storage LUN from the preferred cluster
- 9. Deploy the application and distribute the VM workloads across clusters.

Steps for a non-uniform configuration

1. Modify the SnapMirror policy from AutomatedFailover to AutomatedFailoverDuplex:

snapmirror modify -destination-path <destination_path> -policy
AutomatedFailoverDuplex

2. Modifying the policy triggers a resync. Wait for the resync to complete and confirm the relationship is

```
Insync:
```

snapmirror show -destination-path <destination_path>

- 3. If the existing hosts are local to the primary cluster, add the host to the second cluster and establish connectivity with respective access to each cluster.
- 4. On the secondary site, add a new igroup and initiator for the new hosts and set the host proximity for host affinity to its local site. Map the LUNs to the igroup.

```
SiteB::> igroup create -vserver <svm_name> -igroup <igroup>
SiteB::> igroup add -vserver <svm_name> -igroup <igroup> -initiator
<host_name>
SiteB::> lun mapping create -igroup <igroup> -path <path name>
```

- 5. Discover the paths on the hosts and verify the hosts have an Active/Optimized path to the storage LUN from the preferred cluster
- 6. Deploy the application and distribute the VM workloads across clusters.

Convert from symmetric active/active to an asymmetric iSCSI or FC relationship

If you've configured symmetric active/active protection using iSCSI or FC, you can convert the relationship to asymmetric protection using the ONTAP CLI.

Steps

- 1. Move all the VM workloads to the host local to the source cluster.
- 2. Remove the igroup configuration for the hosts not managing the VM instances then modify the igroup configuration to terminate igroup replication.

igroup modify -vserver <svm name> -igroup <igroup> -replication-peer -

3. On the secondary site, unmap the LUNs.

SiteB::> lun mapping delete -vserver <svm_name> -igroup <igroup> -path <>

4. On the secondary site, delete the symmetric active/active relationship.

SiteB::> snapmirror delete -destination-path <destination_path>

- 5. On the primary site, release the symmetric active/active relationship. SiteA::> snapmirror release -destination-path <destination_path> -relationship -info-only true
- From the secondary site, create a relationship to the same set of volumes with the AutomatedFailover policy to resynchronize the relationship.

```
SiteB::> snapmirror create -source-path <source_path> -destination-path
<destination_path> -cg-item-mappings <source:@destination> -policy
AutomatedFailover
SiteB::> snapmirror resync -destination-path vs1:/cg/cg1 dst -policy
```



The consistency group on the secondary site needs to be deleted before recreating the relationship. The destination volumes must be converted to type DP. To convert the volumes to DP, perform the snapmirror resync command with a non-AutomatedFailover policy: MirrorAndVault, MirrorAllSnapshots, or Sync.

7. Confirm the relationship Mirror State is Snapmirrored the Relationship Status is Insync.

snapmirror show -destination-path destination path

8. Re-discover the paths from the host.

Related information

- snapmirror delete
- · snapmirror modify
- snapmirror release
- snapmirror resync
- snapmirror show

Manage SnapMirror active sync and protect data

Create a common snapshot between ONTAP consistency groups

In addition to the regularly scheduled snapshot operations, you can manually create a common snapshot between the volumes in the primary SnapMirror consistency group and the volumes in the secondary SnapMirror consistency group.

About this task

The scheduled snapshot creation interval is 12 hours.

Before you begin

• The SnapMirror group relationship must be in sync.

Steps

1. Create a common snapshot:

destination::>snapmirror update -destination-path vs1 dst:/cg/cg dst

2. Monitor the progress of the update:

destination::>snapmirror show -fields newest-snapshot

Related information

snapmirror show

Perform a planned failover of ONTAP clusters in a SnapMirror active sync relationship

In a planned failover of ONTAP clusters in a SnapMirror active sync relationship, you switch the roles of the primary and secondary clusters, so that the secondary cluster takes over from the primary cluster. During a failover, what is normally the secondary cluster processes input and output requests locally without disrupting client operations.

You may want to perform a planned failover to test the health of your disaster recovery configuration or to perform maintenance on the primary cluster.

About this task

A planned failover is initiated by the administrator of the secondary cluster. The operation requires switching the primary and secondary roles so that the secondary cluster takes over from the primary. The new primary cluster can then begin processing input and output requests locally without disrupting client operations.

Before you begin

- The SnapMirror active sync relationship must be in sync.
- You cannot initiate a planned failover when a nondisruptive operation is in process. Nondisruptive operations include volume moves, aggregate relocations, and storage failovers.
- The ONTAP Mediator must be configured, connected, and in quorum.

Steps

You can perform a planned failover using the ONTAP CLI or System Manager.

System Manager



From ONTAP 9.14.1 through 9.8, SnapMirror active sync is referred to as SnapMirror Business Continuity (SM-BC).

- 1. In System Manager, select Protection > Overview > Relationships.
- 2. Identify the SnapMirror active sync relationship you want to failover. Next to its name, select the ... next to the relationship's name, then select **Failover**.
- 3. To monitor the status of the failover, use the snapmirror failover show in the ONTAP CLI.

CLI

1. From the destination cluster, initiate the failover operation:

```
destination::>snapmirror failover start -destination-path
vs1_dst:/cg/cg_dst
```

2. Monitor the progress of the failover:

destination::>snapmirror failover show

3. When the failover operation is complete, you can monitor the SnapMirror synchronous protection relationship status from the destination:

destination::>snapmirror show

Related information

- snapmirror failover show
- snapmirror failover start
- snapmirror show

Recover from automatic unplanned ONTAP cluster failover operations

An automatic unplanned failover (AUFO) operation occurs when the primary cluster is down or isolated. The ONTAP Mediator detects when a failover occurs and, and executes an automatic unplanned failover to the the secondary cluster. The secondary cluster is converted to the primary and begins serving clients. This operation is performed only with assistance from the ONTAP Mediator.



After the automatic unplanned failover, it is important to rescan the host LUN I/O paths so that there is no loss of I/O paths.

Reestablish the protection relationship after an unplanned failover

You can reestablish the protection relationship using System Manager or the ONTAP CLI.

System Manager

i.

Steps

From ONTAP 9.14.1 through 9.8, SnapMirror active sync is referred to as SnapMirror Business Continuity (SM-BC).

- 1. Navigate to Protection > Relationships and wait for the relationship state to show "InSync."
- 2. To resume operations on the original source cluster, click and select **Failover**.

CLI

You can monitor the status of the automatic unplanned failover using the snapmirror failover show command.

For example:

```
ClusterB::> snapmirror failover show -instance
Start Time: 9/23/2020 22:03:29
Source Path: vs1:/cg/scg3
Destination Path: vs3:/cg/dcg3
Failover Status: completed
Error Reason:
End Time: 9/23/2020 22:03:30
Primary Data Cluster: cluster-2
Last Progress Update: -
Failover Type: unplanned
Error Reason codes: -
```

Refer to the EMS reference to learn about event messages and about corrective actions.

Resume protection in a fan-out configuration after failover

Beginning with ONTAP 9.15.1, SnapMirror active sync supports automatic reconfiguration in the fan-out leg after a failover event. The async fan-out leg can be a consistency group relationship or an independent volume relationship. For more information, see fan-out configurations.

If you're using ONTAP 9.14.1 or earlier and you experience a failover on the secondary cluster in the SnapMirror active sync relationship, the SnapMirror asynchronous destination becomes unhealthy. You must manually restore protection by deleting and recreating the relationship with the SnapMirror asynchronous endpoint.

Steps

- 1. Verify the failover has completed successfully: snapmirror failover show
- 2. On the SnapMirror asynchronous endpoint, delete the fan-out endpoint: snapmirror delete -destination_path destination_path
- 3. On the third site, create a SnapMirror asynchronous relationships between the new SnapMirror active sync primary volume and the async fan-out destination volume:

snapmirror create -source-path source_path -destination-path destination_path
-policy MirrorAllSnapshots -schedule schedule

- 4. Resynchronize the relationship: snapmirror resync -destination-path destination_path
- 5. Verify the relationship status and heath: snapmirror show

Related information

- snapmirror create
- snapmirror delete
- snapmirror failover show
- snapmirror resync
- snapmirror show

Monitor ONTAP SnapMirror active sync operations

You can monitor the following SnapMirror active sync operations to ensure the health of your SnapMirror active sync configuration:

- ONTAP Mediator
- · Planned failover operations
- · Automatic unplanned failover operations
- · SnapMirror active sync availability



Beginning with ONTAP 9.15.1, System Manager displays the status of your SnapMirror active sync relationship from either cluster. You can also monitor the ONTAP Mediator's status from either cluster in System Manager.

ONTAP Mediator

During normal operations, the ONTAP Mediator state should be connected. If it's in any other state, this might indicate an error condition. You can review the Event Management System (EMS) messages to determine the error and appropriate corrective actions.

Planned failover operations

You can monitor status and progress of a planned failover operation using the snapmirror failover show command. For example:

ClusterB::> snapmirror failover start -destination-path vs1:/cg/dcg1

Once the failover operation is complete, you can monitor the SnapMirror protection status from the new destination cluster. For example:

```
ClusterA::> snapmirror show
```

Refer to the EMS reference to learn about event messages and corrective actions.

Automatic unplanned failover operations

During an unplanned automatic failover, you can monitor the status of the operation using the snapmirror failover show command.

```
ClusterB::> snapmirror failover show -instance
Start Time: 9/23/2020 22:03:29
Source Path: vs1:/cg/scg3
Destination Path: vs3:/cg/dcg3
Failover Status: completed
Error Reason:
End Time: 9/23/2020 22:03:30
Primary Data Cluster: cluster-2
Last Progress Update: -
Failover Type: unplanned
Error Reason codes: -
```

Refer to the EMS reference to learn about event messages and about corrective actions.

SnapMirror active sync availability

You can check the availability of the SnapMirror active sync relationship using a series of commands, either on the primary cluster, the secondary cluster, or both.

Commands you use include the snapmirror mediator show command on both the primary and secondary cluster to check the connection and quorum status, the snapmirror show command, and the volume show command. For example:

SMBC A::*> snapmirror mediator show Mediator Address Peer Cluster Connection Status Quorum Status _____ ____ 10.236.172.86 SMBC B connected true SMBC B::*> snapmirror mediator show Mediator Address Peer Cluster Connection Status Quorum Status _____ 10.236.172.86 SMBC A connected true SMBC B::*> snapmirror show -expand Progress Destination Mirror Relationship Total Source Last Path Type Path State Status Progress Healthy Updated _____ ____ ____ ____ _____ _____ _____ vs0:/cg/cg1 XDP vs1:/cg/cg1_dp Snapmirrored Insync - true vs0:vol1 XDP vs1:vol1 dp Snapmirrored Insync true 2 entries were displayed. SMBC A::*> volume show -fields is-smbc-master, smbc-consensus, is-smbcfailover-capable -volume vol1 vserver volume is-smbc-master is-smbc-failover-capable smbc-consensus _____ _____ vs0 vol1 true false Consensus SMBC_B::*> volume show -fields is-smbc-master, smbc-consensus, is-smbcfailover-capable -volume vol1 dp vserver volume is-smbc-master is-smbc-failover-capable smbc-consensus _____ _ _____ vs1 vol1 dp false true No-consensus

Related information

- snapmirror failover show
- snapmirror failover start
- · snapmirror mediator show

Add or remove volumes to an ONTAP consistency group

As your application workload requirements change, you may need to add or remove volumes from a consistency group to ensure business continuity. The process of adding and removing volumes in an an active SnapMirror active sync relationship depends on

the version of ONTAP you are using.

In most instances, this is a disruptive process requiring you to delete the SnapMirror relationship, modify the consistency group, then resume protection. Beginning with ONTAP 9.13.1, adding volumes to a consistency group with an active SnapMirror relationship is a non-disruptive operation.

About this task

- In ONTAP 9.9.1, you can add or remove volumes to a consistency group using the ONTAP CLI.
- Beginning with ONTAP 9.10.1, it is recommended that you manage consistency groups through System Manager or with the ONTAP REST API.

If you want to change the composition of the consistency group by adding or removing a volume, you must first delete the original relationship and then create the consistency group again with the new composition.

• Beginning with ONTAP 9.13.1, you can non-disruptively add volumes to a consistency group with an active SnapMirror relationship from the source or destination. This action is not supported with the NVMe protocol.

Removing volumes is a disruptive operation. You must delete the SnapMirror relationship before removing volumes.

ONTAP 9.9.1-9.13.0

Before you begin

- You cannot begin to modify the consistency group while it is in the InSync state.
- The destination volume should be of type DP.
- The new volume you add to expand the consistency group must have a pair of common snapshots between the source and destination volumes.

Steps

The examples shown in two volume mappings: $vol_src1 \leftrightarrow vol_dst1$ and $vol_src2 \leftrightarrow vol_dst2$, in a consistency group relationship between the end points $vsl_src:/cg/cg_src$ and $vsl_dst:/cg/cg_dst$.

1. On the source and destination clusters, verify there is a common snapshot between the source and destination clusters with the command snapshot show -vserver svm_name -volume volume name -snapshot snapmirror

source::>snapshot show -vserver vs1_src -volume vol_src3 -snapshot
snapmirror*

destination::>snapshot show -vserver vs1_dst -volume vol_dst3 -snapshot snapmirror*

2. If no common snapshot exists, create and initialize a FlexVol SnapMirror relationship:

```
destination::>snapmirror initialize -source-path vs1_src:vol_src3
-destination-path vs1_dst:vol_dst3
```

3. Delete the consistency group relationship:

destination::>snapmirror delete -destination-path vs1_dst:/cg/cg_dst

4. Release the source SnapMirror relationship and retain the common snapshots:

```
source::>snapmirror release -relationship-info-only true -destination-path
vs1 dst:vol dst3
```

5. Unmap the LUNs and delete the existing consistency group relationship:

```
destination::>lun mapping delete -vserver vs1_dst -path <lun_path> -igroup
<igroup_name>
```



The destination LUNs are unmapped, while the LUNs on the primary copy continue to serve the host I/O.

destination::>snapmirror delete -destination-path vs1_dst:/cg/cg_dst

source::>snapmirror release -destination-path vs1_dst:/cg/cg_dst
-relationship-info-only true

6. If you are using ONTAP 9.10.1 through 9.13.0, delete and recreate and the consistency group on

the source with the correct composition. Follow the steps in Delete a consistency group and then Configure a single consistency group. In ONTAP 9.10.1 and later, you must perform the delete and create operations in System Manager or with the ONTAP REST API; there is no CLI procedure.

If you are using ONTAP 9.9.1, skip to the next step.

7. Create the new consistency group on the destination with the new composition:

```
destination::>snapmirror create -source-path vs1_src:/cg/cg_src
-destination-path vs1_dst:/cg/cg_dst -cg-item-mappings vol_src1:@vol_dst1,
vol_src2:@vol_dst2, vol_src3:@vol_dst3
```

8. Resynchronize the zero RTO consistency group relationship to ensure it is in sync:

```
destination::>snapmirror resync -destination-path vs1_dst:/cg/cg_dst
```

9. Remap the LUNs that you unmapped in Step 5:

```
destination::> lun map -vserver vs1 dst -path lun path -igroup igroup name
```

10. Rescan host LUN I/O paths to restore all paths to the LUNs.

ONTAP 9.13.1 and later

Beginning with ONTAP 9.13.1, you can non-disruptively add volumes to a consistency group with an active SnapMirror active sync relationship. SnapMirror active sync supports adding volumes from both the source or destination.



From ONTAP 9.14.1 through 9.8, SnapMirror active sync is referred to as SnapMirror Business Continuity (SM-BC).

For details on adding volumes from the source consistency group, see Modify a consistency group.

Add a volume from the destination cluster

- 1. On the destination cluster, select **Protection > Relationships**.
- 2. Find the SnapMirror configuration you want to add volumes to. Select then **Expand**.
- 3. Select the volume relationships whose volumes are to be added to consistency group
- 4. Select **Expand**.

Related information

- snapmirror delete
- snapmirror initialize
- snapmirror release
- snapmirror resync

Upgrade and revert with ONTAP SnapMirror active sync

SnapMirror active sync is supported beginning with ONTAP 9.9.1. Upgrading and reverting your ONTAP cluster or controllers has implications on your SnapMirror active sync relationships depending on the ONTAP version to which you are upgrading or

reverting.

Refresh a cluster

Beginning with ONTAP 9.16.1, SnapMirror active sync supports four-node clusters in symmetric active/active configurations. You can use the four-node cluster to upgrade controllers and storage.

Before you begin

- Review the requirements for four-node clusters.
- You can create asymmetrical configurations during the tech refresh process; however, you should return to a symmetrical configuration after completing the refresh.
- These instructions apply to an existing four-node configuration with 50 or fewer consistency groups and 400 or fewer volume endpoints.

Steps

- 1. Move all the SnapMirror active sync volumes onto a *single* high-availability (HA) pair.
- 2. Remove the unused nodes from the cluster.
- 3. Add the new nodes to the cluster.
- 4. Move all the volumes into the new nodes.
- 5. Remove the unused nodes from the cluster then replace them with the new nodes.

Upgrade ONTAP with SnapMirror active sync

To use SnapMirror active sync, all nodes on the source and destination clusters must be running ONTAP 9.9.1 or later.

When upgrading ONTAP with active SnapMirror active sync relationships, you should use automated nondisruptive upgrade (ANDU). Using ANDU ensures your SnapMirror active sync relationships are in sync and healthy during the upgrade process.

There are no configuration steps to prepare SnapMirror active sync deployments for ONTAP upgrades. However, it is recommended that before and after the upgrade, you should check that:

- SnapMirror active sync relationships are in sync.
- There are no errors related to SnapMirror in the event log.
- The Mediator is online and healthy from both clusters.
- All hosts can see all paths properly to protect LUNs.



When you upgrade clusters from ONTAP 9.9.1 or 9.9.1 to ONTAP 9.10.1 and later, ONTAP creates new consistency groups on both source and destination clusters for SnapMirror active sync relationships that can be configured using System Manager.



The snapmirror quiesce and snampirror resume commands are not supported with SnapMirror active sync.

Revert to ONTAP 9.9.1 from ONTAP 9.10.1

To revert relationships from 9.10.1 to 9.9.1, SnapMirror active sync relationships must be deleted, followed by the 9.10.1 consistency group instance. Consistency groups with an active SnapMirror active sync relationship
cannot be deleted. Any FlexVol volumes that were upgraded to 9.10.1 previously associated with another Smart Container or Enterprise App in 9.9.1 or earlier will no longer be associated on revert. Deleting consistency groups does not delete the constituent volumes or volume granular snapshots. Refer to Delete a consistency group for more information on this task in ONTAP 9.10.1 and later.

Revert from ONTAP 9.9.1



SnapMirror active sync is not supported with mixed ONTAP clusters than include releases earlier than ONTAP 9.9.1.

When you revert from ONTAP 9.9.1 to an earlier release of ONTAP, you must be aware of the following:

- If the cluster hosts an SnapMirror active sync destination, reverting to ONTAP 9.8 or earlier is not allowed until the relationship is broken and deleted.
- If the cluster hosts an SnapMirror active sync source, reverting to ONTAP 9.8 or earlier is not allowed until the relationship is released.
- All user-created custom SnapMirror active sync policies must be deleted before reverting to ONTAP 9.8 or earlier.

To meet these requirements, see Remove a SnapMirror active sync configuration.

Steps

1. Confirm your readiness to revert, entering the following command from one of the clusters in the SnapMirror active sync relationship:

cluster::> system node revert-to -version 9.7 -check-only

The following sample output shows a cluster that is not ready to revert with instructions for clean up.

```
cluster::> system node revert-to -version 9.7 -check-only
Error: command failed: The revert check phase failed. The following
issues must be resolved before revert can be completed. Bring the data
LIFs down on running vservers. Command to list the running vservers:
vserver show -admin-state running Command to list the data LIFs that are
up: network interface show -role data -status-admin up Command to bring
all data LIFs down: network interface modify {-role data} -status-admin
down
Disable snapshot policies.
    Command to list snapshot policies: "snapshot policy show".
    Command to disable snapshot policies: "snapshot policy modify
-vserver
   * -enabled false"
   Break off the initialized online data-protection (DP) volumes and
delete
   Uninitialized online data-protection (DP) volumes present on the
local
   node.
    Command to list all online data-protection volumes on the local
```

```
node:
   volume show -type DP -state online -node <local-node-name>
    Before breaking off the initialized online data-protection volumes,
   quiesce and abort transfers on associated SnapMirror relationships
and
   wait for the Relationship Status to be Quiesced.
    Command to quiesce a SnapMirror relationship: snapmirror quiesce
    Command to abort transfers on a SnapMirror relationship: snapmirror
   abort
    Command to see if the Relationship Status of a SnapMirror
relationship
   is Quiesced: snapmirror show
    Command to break off a data-protection volume: snapmirror break
    Command to break off a data-protection volume which is the
destination
   of a SnapMirror relationship with a policy of type "vault":
snapmirror
   break -delete-snapshots
    Uninitialized data-protection volumes are reported by the
"snapmirror
   break" command when applied on a DP volume.
    Command to delete volume: volume delete
   Delete current version snapshots in advanced privilege level.
    Command to list snapshots: "snapshot show -fs-version 9.9.1"
    Command to delete snapshots: "snapshot prepare-for-revert -node
   <nodename>"
   Delete all user-created policies of the type active-strict-sync-
mirror
   and active-sync-mirror.
   The command to see all active-strict-sync-mirror and active-sync-
mirror
   type policies is:
   snapmirror policy show -type
   active-strict-sync-mirror, active-sync-mirror
   The command to delete a policy is :
    snapmirror policy delete -vserver <SVM-name> -policy <policy-name>
```

2. Once you've satisfied the requirements of the revert check, see Revert ONTAP.

Related information

- network interface
- snapmirror break
- snapmirror policy delete

- snapmirror policy show
- snapmirror quiesce
- snapmirror show

Remove an ONTAP SnapMirror active sync configuration

If you no longer require zero RTO SnapMirror synchronous protection, you can delete your SnapMirror active sync relationship.

Remove an asymmetric configuration

- Before you delete the SnapMirror active sync relationship, all LUNs in the destination cluster must be unmapped.
- After the LUNs are unmapped and the host is rescanned, the SCSI target notifies the hosts that the LUN inventory has changed. The existing LUNs on the zero RTO secondary volumes change to reflect a new identity after the zero RTO relationship is deleted. Hosts discover the secondary volume LUNs as new LUNs that have no relationship to the source volume LUNs.
- The secondary volumes remain DP volumes after the relationship is deleted. You can issue the snapmirror break command to convert them to read/write.
- Deleting the relationship is not allowed in the failed-over state when the relationship is not reversed.

Steps

1. From the secondary cluster, remove the SnapMirror active sync consistency group relationship between the source endpoint and destination endpoint:

destination::>snapmirror delete -destination-path vs1_dst:/cg/cg_dst

2. From the primary cluster, release the consistency group relationship and the snapshots created for the relationship:

```
source::>snapmirror release -destination-path vs1 dst:/cg/cg dst
```

- 3. Perform a host rescan to update the LUN inventory.
- Beginning with ONTAP 9.10.1, deleting the SnapMirror relationship does not delete the consistency group. If you want to delete the consistency group, you must use System Manager or the ONTAP REST API. See Delete a consistency group for more information.

Remove iSCSI or FC symmetric active/active configuration

You can remove a symmetric configuration using System Manager or the ONTAP CLI. In both interfaces, there are different steps for uniform and non-uniform configurations.

System Manager

Steps for a uniform configuration

- 1. On the primary site, remove the remote hosts from the igroup and terminate replication.
 - a. Navigate to **Hosts > SAN Initiator Groups**.
 - b. Select the igroup you want to modify then Edit.
 - c. Remove the remote initiator and terminate igroup replication. Select **Save**.
- 2. On the secondary site, delete the replicated relationship by unmapping the LUNs.
 - a. Navigate to Hosts > SAN Initiator Groups.
 - b. Select the igroup with the SnapMirror relationship, then Delete.
 - c. In the dialog box, select the **Unmap the associated LUNs** box then **Delete**.
 - d. Navigate to **Protection > Relationships**.
 - e. Select the SnapMirror active sync relationship then **Release** to delete the relationships.

Steps for a non-uniform configuration

- 1. On the primary site, remove the remote hosts from the igroup and terminate replication.
 - a. Navigate to **Hosts > SAN Initiator Groups**.
 - b. Select the igroup you want to modify then Edit.
 - c. Remove the remote initiator and terminate igroup replication. Select Save.
- 2. On the secondary site, remove the SnapMirror active sync relationship.
 - a. Navigate to Protection > Relationships.
 - b. Select the SnapMirror active sync relationship then **Release** to delete the relationships.

CLI

Steps for a uniform configuration

- 1. Move all the VM workloads to the host local to source cluster of SnapMirror active sync.
- 2. On the source cluster, remove the initiators from the igroup and modify the igroup configuration to terminate igroup replication.

SiteA::> igroup remove -vserver <svm_name> -igroup <igroup_name> -os-type
<os_type> -initiator <host2>
SiteA::> igroup modify -vserver <svm_name> -igroup <igroup_name> -os-type
<os type> -replication-peer "-"

3. On the secondary site, delete the LUN mapping and remove the igroup configuration:

SiteB::> lun mapping delete -vserver <svm_name> -igroup <igroup_name> -path
<>
SiteB::> igroup delete -vserver <svm name> -igroup <igroup name>

4. On the secondary site, delete the SnapMirror active sync relationship.

SiteB::> snapmirror delete -destination-path destination path

5. On the primary site, release the SnapMirror active sync relationship from primary site.

SiteA::> snapmirror release -destination-path <destination path>

6. Rediscover the paths to verify that only the local path is available to the host.

Steps for a non-uniform configuration

- 1. Move all the VM workloads to the host local to source cluster of SnapMirror active sync.
- 2. On the source cluster, remove the initiators from the igroup.

```
SiteA::> igroup remove -vserver <svm_name> -igroup <igroup_name> -initiator
<host2>
```

3. On the secondary site, delete the LUN mapping and remove the igroup configuration:

```
SiteB::> lun mapping delete -vserver <svm_name> -igroup <igroup_name> -path
<>
SiteB::> igroup delete -vserver <svm name> -igroup <igroup_name>
```

4. On the secondary site, delete the SnapMirror active sync relationship.

SiteB::> snapmirror delete -destination-path <destination path>

5. On the primary site, release the SnapMirror active sync relationship from primary site.

SiteA::> snapmirror release -destination-path <destination path>

6. Rediscover the paths to verify that only the local path is available to the host.

Remove an NVMe symmetric active/active configuration

System Manager

Steps

- 1. On the source cluster, navigate to **Protection > Replication**.
- 2. Locate the relationship you want to remove, select **and choose Delete**.

CLI

1. From the destination cluster, delete the SnapMirror active sync relationship.

```
snapmirror delete -destination-path <destination_path> -unmap-namespace
true
```

Example:

```
DST::> snapmirror delete -destination-path vs1:/cg/cg_dst_1 -force
true
```

The subsystem and its namespaces are removed from the secondary cluster.

2. From the source cluster, release the SnapMirror active sync relationship from primary site.

snapmirror release -destination-path <destination path>

Example:

SRC::> snapmirror release -destination-path vs1:/cg/cg dst 1

3. Rediscover the paths to verify that only the local path is available to the host.

Related information

- snapmirror break
- snapmirror delete
- snapmirror release

Remove ONTAP Mediator or ONTAP Cloud Mediator

If you want to remove an existing ONTAP Mediator or ONTAP Cloud Mediator configuration from your ONTAP clusters, you can do so by using the snapmirror mediator remove command. For example, you can use only one type of Mediator at a time, so you must remove one instance before you install the other.

Steps

You can remove ONTAP Mediator or ONTAP Cloud Mediator by completing one of the following steps.

ONTAP Mediator

1. Remove ONTAP Mediator:

snapmirror mediator remove -mediator-address <address> -peer-cluster
<peerClusterName>

Example:

snapmirror mediator remove -mediator-address 12.345.678.90 -peer -cluster cluster xyz

ONTAP Cloud Mediator

1. Remove ONTAP Cloud Mediator:

```
snapmirror mediator remove -peer-cluster <peerClusterName> -type cloud
```

Example:

snapmirror mediator remove -peer-cluster cluster_xyz -type cloud

Related information

• snapmirror mediator remove

Troubleshoot

ONTAP SnapMirror delete operation fails in takover state

Use the following information if the snapmirror delete command fails when a SnapMirror active sync consistency group relationship is in takeover state.

Issue:

When ONTAP 9.9.1 is installed on a cluster, executing the snapmirror delete command fails when a SnapMirror active sync consistency group relationship is in takeover state.

Example:

```
C2_cluster::> snapmirror delete vs1:/cg/dd
Error: command failed: RPC: Couldn't make connection
```

Solution

When the nodes in a SnapMirror active sync relationship are in takeover state, perform the SnapMirror delete and release operation with the "-force" option set to true.

Example:

```
C2_cluster::> snapmirror delete vs1:/cg/dd -force true
Warning: The relationship between source "vs0:/cg/ss" and destination
          "vs1:/cg/dd" will be deleted, however the items of the
destination
          Consistency Group might not be made writable, deletable, or
modifiable
          after the operation. Manual recovery might be required.
Do you want to continue? {y|n}: y
Operation succeeded: snapmirror delete for the relationship with
destination "vs1:/cg/dd".
```

Related information

• snapmirror delete

Failure creating an ONTAP SnapMirror relationship and initializing consistency group

Use the following information if the creation of a SnapMirror relationship and consistency group initialization fails.

Issue:

Creation of SnapMirror relationship and consistency group initialization fails.

Solution:

Ensure that you have not exceeded the limit of consistency groups per cluster. Consistency group limits in SnapMirror active sync are platform independent and differ based on the version of ONTAP. See Object limits for guidance specific to your ONTAP version.

Error:

If the consistency group is stuck initializing, check the status of your consistency group initializations with the ONTAP REST API, System Manager or the command sn show -expand.



From ONTAP 9.14.1 through 9.8, SnapMirror active sync is referred to as SnapMirror Business Continuity (SM-BC).

Solution:

If consistency groups fail to initialize, remove the SnapMirror active sync relationship, delete the consistency group, then recreate the relationship and initialize it. This workflow differs depending on the version of ONTAP you are using.

If you are using ONTAP 9.9.1

If you are using ONTAP 9.10.1 or later

 Remove the SnapMirror active sync configuration Create a consistency group relationship then Initialize the consistency group relationship 	 Under Protection > Relationships, find the SnapMirror active sync relationship on the consistency group. Select , then Delete to remove the SnapMirror active sync relationship.
	2. Delete the consistency group
	3. Configure the consistency group

Planned ONTAP cluster failover unsuccessful

Use the following information if the planned failover operation is unsuccessful.

Issue:

After executing the snapmirror failover start command, the output for the snapmirror failover show command displays a message indicates that a nondisruptive operation is in progress.

Example:

08:35:04

Cause:

A planned failover cannot begin when a nondisruptive operation is in progress, including volume move, aggregate relocation, and storage failover.

Solution:

Wait for the nondisruptive operation to complete and try the failover operation again.

Related information

- snapmirror failover show
- snapmirror failover start

ONTAP Mediator or ONTAP Cloud Mediator not reachable or Mediator quorum status is false

Use the following information if the ONTAP Mediator or ONTAP Cloud Mediator is not reachable or the Mediator quorum status is false.

Issue:

After executing the snapmirror failover start command, the output for the snapmirror failover show command displays a message indicating that either the ONTAP Mediator or ONTAP Cloud Mediator is not configured.

See Configure the ONTAP Mediator and clusters for SnapMirror active sync or Configure the ONTAP Cloud Mediator for SnapMirror active sync.

Example:

Cause:

Mediator is not configured or there are network connectivity issues.

Solution:

If the ONTAP Mediator is not configured, you must configure the ONTAP Mediator before you can establish a SnapMirror active sync relationship. Fix any network connectivity issues. Make sure Mediator is connected and quorum status is true on both the source and destination site using the snapmirror mediator show command. For more information, see Configure the ONTAP Mediator.

Example:

```
cluster::> snapmirror mediator show
Mediator Address Peer Cluster Connection Status Quorum Status
10.234.10.143 cluster2 connected true
```

Related information

- snapmirror failover show
- snapmirror failover start
- · snapmirror mediator show

ONTAP Cloud Mediator is reachable but responding slowly

Use the following information if the ONTAP Cloud Meditor fails with an error that says the ping latency is higher than the recommended latency.

Issue:

System Manager: The Cloud Mediator service is reachable but it's responding slowly.

CLI:

The mediator add command fails with the error:

Error: command failed: The ping latency of the BlueXP cloud server is $\langle x \rangle$ ms which is higher than twice the recommended latency of 200 ms.

Cause:

The clusters might not be located in proximity to the BlueXP cloud or there are network path bottlenecks.

Solution:

- Check the geographical location and proximity to the BlueXP cloud (US East).
- Optimize network path or address bottlenecks.
- Measure round trip time (RTT) using network tools, and reduce latency to within recommended limits.
- Use an HTTP proxy to improve performance.

See Configure the ONTAP Cloud Mediator and clusters for SnapMirror active sync.

Automatic unplanned failover not triggered on Site B

Use the following information if a failure on Site A does not trigger an unplanned failover on Site B.

Issue:

A failure on Site A does not trigger an unplanned failover on Site B.

Possible cause #1:

The ONTAP Mediator or the ONTAP Cloud Mediator is not configured. To determine if this is the cause, issue the snapmirror mediator show command on the Site B cluster.

Example:

```
Cluster2::> snapmirror mediator show
This table is currently empty.
```

This example indicates that the Mediator is not configured on Site B.

Solution:

Ensure that Mediator is configured on both clusters, that the status is connected, and quorum is set to True.

Possible cause #2:

SnapMirror consistency group is out of sync. To determine if this is the cause, view the event log to view if the consistency group was in sync during the time at which the Site A failure occurred.

Example:

Solution:

Complete the following steps to perform a forced failover on Site B.

- 1. Unmap all LUNs belonging to the consistency group from Site B.
- 2. Delete the SnapMirror consistency group relationship using the force option.
- 3. Enter the snapmirror break command on the consistency group constituent volumes to convert volumes from DP to R/W, to enable I/O from Site B.
- 4. Boot up the Site A nodes to create a zero RTO relationship from Site B to Site A.
- 5. Release the consistency group with relationship-info-only on Site A to retain common snapshot and unmap the LUNs belonging to the consistency group.
- 6. Convert volumes on Site A from R/W to DP by setting up a volume level relationship using either the Sync policy or Async policy.
- 7. Issue the snapmirror resync to synchronize the relationships.
- 8. Delete the SnapMirror relationships with the Sync policy on Site A.
- Release the SnapMirror relationships with Sync policy using relationship-info-only true on Site B.
- 10. Create a consistency group relationship from Site B to Site A.
- 11. Perform a consistency group resync from Site A, and then verify that the consistency group is in sync.
- 12. Rescan host LUN I/O paths to restore all paths to the LUNs.

Related information

- snapmirror break
- snapmirror mediator show
- snapmirror resync

Link between Site B and ONTAP Mediator down and Site A down

To check on the connection of the ONTAP Mediator or the ONTAP Cloud Mediator, use the snapmirror mediator show command. If the connection status is unreachable and Site B is unable to reach Site A, you will have an output similar to the one below. Follow the steps in the solution to restore connection

Example:

Using ONTAP Cloud Mediator output `snapmirror mediator show`command:

```
cluster::> snapmirror mediator showMediator Address Peer ClusterConnection Status Quorum Status Type0.0.0.0C1_clusterunreachabletruecloud
```

Using ONTAP Mediator output `snapmirror mediator show`command:

```
cluster::> snapmirror mediator show
Mediator Address Peer Cluster Connection Status Quorum Status
_____ ____
10.237.86.17 C1 cluster unreachable
                                        true
SnapMirror consistency group relationship status is out of sync.
C2 cluster::> snapmirror show -expand
              Destination Mirror Relationship
Source
                                          Total
Last
Path
      Type Path State Status
                                          Progress Healthy
Updated
_____
vs0:/cg/src cg 1 XDP vs1:/cg/dst cg 1 Snapmirrored OutOfSync - false -
vs0:zrto cg 655724 188a RW1 XDP vs1:zrto cg 655755 188c DP1 Snapmirrored
OutOfSync - false -
vs0:zrto cg 655733 188a RW2 XDP vs1:zrto cg 655762 188c DP2 Snapmirrored
OutOfSync - false -
vs0:zrto_cg_655739_188b_RW1 XDP vs1:zrto_cg_655768_188d_DP1 Snapmirrored
OutOfSync - false -
vs0:zrto cg 655748 188b RW2 XDP vs1:zrto cg 655776 188d DP2 Snapmirrored
OutOfSync - false -
5 entries were displayed.
Site B cluster is unable to reach Site A.
C2 cluster::> cluster peer show
Peer Cluster Name Cluster Serial Number Availability
Authentication
_____
_____
            1-80-000011 Unavailable ok
C1 cluster
```

Solution

Force a failover to enable I/O from Site B and then establish a zero RTO relationship from Site B to Site A. Complete the following steps to perform a forced failover on Site B.

- 1. Unmap all LUNs belonging to the consistency group from Site B.
- 2. Delete the SnapMirror consistency group relationship using the force option.
- 3. Enter the SnapMirror break command (snapmirror break -destination_path *svm*:_volume_) on the consistency group constituent volumes to convert volumes from DP to RW, to enable I/O from Site B.

You must issue the SnapMirror break command for each relationship in the consistency group. For example, if there are three volumes in the consistency group, you will issue the command for each volume.

- 4. Boot up the Site A nodes to create a zero RTO relationship from Site B to Site A.
- 5. Release the consistency group with relationship-info-only on Site A to retain common snapshot and unmap the LUNs belonging to the consistency group.
- 6. Convert volumes on Site A from RW to DP by setting up a volume level relationship using either Sync policy or Async policy.
- 7. Issue the snapmirror resync command to synchronize the relationships.
- 8. Delete the SnapMirror relationships with Sync policy on Site A.
- 9. Release the SnapMirror relationships with Sync policy using relationship-info-only true on Site B.
- 10. Create a consistency group relationship between Site B to Site A.
- 11. From the source cluster, resynchronize the consistency group. Verify the consistency group state is in sync.
- 12. Rescan the host LUN I/O paths to restore all paths to the LUNs.

Related information

- snapmirror break
- snapmirror mediator show
- snapmirror resync
- snapmirror show

Link between Site A and ONTAP Mediator down and Site B down

When using SnapMirror active sync, you may lose connectivity between the ONTAP Mediator or your peered clusters. You can diagnose the issue by checking the connection, availability, and consensus status of the different parts of the SnapMirror active sync relationship then forcefully resuming connection.

Table 1. Determining the cause

What to check	CLI command	Indicator
Mediator from Site A	snapmirror mediator show	The connection status displays as unreachable
Site B connectivity	cluster peer show	Availability displays as unavailable
Consensus status of the SnapMirror active sync volume	volume show <i>volume_name</i> -fields smbc-consensus	The sm-bc consensus field displays Awaiting-consensus

For additional information about diagnosing and resolving this issue, refer to the Knowledge Base article Link between Site A and Mediator down and Site B down when using SnapMirror active sync.

Related information

- cluster peer show
- snapmirror mediator show

ONTAP SnapMirror delete operation fails when fence is set on destination volume

Use the following information if the SnapMirror delete operation fails when any of the destination volumes have redirection fence set.

Issue:

SnapMirror delete operation fails when any of the destination volumes have redirection fence set.

Solution

Performing the following operations to retry the redirection and remove the fence from the destination volume.

- SnapMirror resync
- SnapMirror update

Volume move operation stuck when ONTAP primary is down

Use the following information if a volume move operation is stuck indefinitely in cutover deferred state when the primary site is down in a SnapMirror active sync relationship.

Issue:

A volume move operation is stuck indefinitely in cutover deferred state when the primary site is down in a SnapMirror active sync relationship.

When the primary site is down, the secondary site performs an automatic unplanned failover (AUFO). When a volume move operation is in progress when the AUFO is triggered the volume move becomes stuck.

Solution:

Abort the volume move instance that is stuck and restart the volume move operation.

ONTAP SnapMirror release fails when unable to delete snapshot

Use the following information if the SnapMirror release operation fails when the snapshot cannot be deleted.

Issue:

The SnapMirror release operation fails when the snapshot cannot be deleted.

Solution:

The snapshot contains a transient tag. Use the snapshot delete command with the -ignore-owners option to remove the transient snapshot.

snapshot delete -volume <volume_name> -snapshot <snapshot_name> -ignore-owners
true -force true

Retry the snapmirror release command.

Related information

snapmirror release

Volume move reference snapshot shows as the newest for ONTAP SnapMirror relationship

Use the following information if the volume move reference snapshot shows as the newest for the SnapMirror relationship after a volume move operation.

Issue:

After performing a volume move operation on a consistency group volume, the volume move reference snapshot might incorrectly display as the newest for the SnapMirror relationship.

You can view the newest snapshot with the following command:

snapmirror show -fields newest-snapshot status -expand

Solution:

Manually perform a snapmirror resync or wait for the next automatic resync operation after the volume move operation completes.

Related information

- snapmirror resync
- snapmirror show

ONTAP Mediator for MetroCluster and SnapMirror active sync

Learn about ONTAP Mediator

This documentation refers to the on-premise version of ONTAP Mediator. For information about ONTAP Cloud Mediator, available beginning with ONTAP 9.17.1, see the SnapMirror active sync documentation.

ONTAP Mediator provides several functions for ONTAP features:

- Provides a persistent and fenced store for HA metadata.
- Serves as a ping proxy for controller liveliness.
- Provides synchronous node health query functionality to aid in quorum determination.

ONTAP Mediator provides two additional systemctl services:

• ontap_mediator.service

Maintains the REST API server for managing the ONTAP relationships.

* mediator-scst.service

Controls the startup and shutdown of the iSCSI module (SCST).

Tools provided for the system administrator

Tools provided for the system administrator:

* /usr/local/bin/mediator_change_password

Sets a new API password when the current API username and password are provided.

* /usr/local/bin/mediator_change_user

Sets a new API username when the current API username and password are provided.

* /usr/local/bin/mediator_generate_support_bundle

Generates a local tgz file containing all useful support information needed for communication with NetApp customer support. This includes application configuration, logs, and some system information. The bundles are generated on the local disk and can be transferred manually, as needed. Storage location: /opt/netapp/data/support_bundles/

* /usr/local/bin/uninstall_ontap_mediator

Removes the ONTAP Mediator package and the SCST kernel module. This includes all configuration, logs, and mailbox data.

* /usr/local/bin/mediator_unlock_user

Releases a lock-out on the API user account if the authentication retry limit was reached. This feature is used to prevent brute force password derivation. It prompts the user for the correct username and password.

* /usr/local/bin/mediator_add_user

(Support only) Used to add the API user upon installation.

Special Notes

ONTAP Mediator relies on SCST to provide iSCSI (See http://scst.sourceforge.net/index.html). This package is a kernel module that is compiled during installation specifically for the kernel. Any updates to the kernel might require SCST to be re-installed. Alternatively, uninstall then re-install ONTAP Mediator, then reconfigure the ONTAP relationship.



Any updates to the server OS kernel should be coordinated with a maintenance window in ONTAP.

What's new in ONTAP Mediator

New enhancements to ONTAP Mediator are provided with each release. Here's what's new.

Enhancements

For SCST version information, see the SCST support matrix.

ONTAP Mediator version Enhancements

1.10	 Support for RHEL: Compatible: 9.5. Recommended: 10.0, 9.6, 9.4, and 8.10. Support for Rocky Linux 10.0, 9.6, and 8.10. Upgraded the base Python version from Python 3.9 to Python 3.12.
1.9.1	 Support for RHEL: Compatible: 9.3, 9.1, 8.9, 8.7, 8.6, 8.5, and 8.4. Recommended: 9.5, 9.4, 9.2, 9.0, 8.10, and 8.8. Support for Rocky Linux 9.5 and 8.10. Added new certificates for code signature verification. Added support for skipping code signature checks using the -skip-code -signature-check flag. Includes installer warnings when expired code signature certificates are detected.
1.9	 Support for RHEL: Compatible: 9.3, 9.1, 8.9, 8.7, 8.6, 8.5, and 8.4. Recommended: 9.5, 9.4, 9.2, 9.0, 8.10, and 8.8. Support for Rocky Linux 9.5 and 8.10. FIPS support for RHEL and Rocky Linux. Added performance enhancements for larger scalability. Improved filenames to simplify the setup of PKI-signed certificates.
1.8	 Support for RHEL: Compatible: 8.7, 8.6, 8.5, and 8.4. Recommended: 9.4, 9.3, 9.2, 9.1, 9.0, 8.10, 8.9, and 8.8. Support for Rocky Linux 9.4 and 8.10.
1.7	 Support for RHEL: Compatible: 8.7, 8.6, 8.5, and 8.4. Recommended: 9.3, 9.2, 9.1, 9.0, 8.9, and 8.8. Support for Rocky Linux 9.3 and 8.9. Support for SAN (Subject Alternative Name) data in self-signed certificates and third-party signed certificates.

1.6	 Python 3.9 updates. Support for RHEL: Compatible: 8.7, 8.6, 8.5, and 8.4. Recommended: 9.2, 9.1, 9.0, and 8.8. Support for Rocky Linux 9.2 and 8.8. Discontinued support for RHEL 7.x / CentOS all releases.
1.5	 Support for RHEL 8.5, 8.4, 8.3, 8.2, 8.1, 8.0, 7.9, 7.8, 7.7, and 7.6. Support for CentOS 7.9, 7.8, 7.7, and 7.6. Includes deprecation warnings for RHEL 7.x / CentOS 7.x. Optimizes speed for larger scale SnapMirror active sync systems. Cryptographic code-signature added to the installer.
1.4	 Support for RHEL 8.5, 8.4, 8.3, 8.2, 8.1, 8.0, 7.9, 7.8, 7.7, and 7.6. Support for CentOS 7.9, 7.8, 7.7, and 7.6. Added support for UFEI-based firmware's Secure Boot (SB).
1.3	 Support for RHEL 8.3, 8.2, 8.1, 8.0, 7.9, 7.8, 7.7, and 7.6. Support for CentOS 7.9, 7.8, 7.7, and 7.6.
1.2	 Support for RHEL 8.1, 8.0, 7.9, 7.8, 7.7, and 7.6. Support for CentOS 7.9, 7.8, 7.7, and 7.6. Support for HTTPs mailboxes. For use with ONTAP 9.8+ MCC-IP AUSO and SnapMirror active sync ZRTO.
1.1	 Support for RHEL 8.0 and 7.6. Support for CentOS 7.6. Eliminates Perl dependencies.
1.0	 Support for iSCSI mailboxes. For use with ONTAP 9.7+ MCC-IP AUSO. Support for RHEL/CentOS 7.6.

OS support matrix

OS for ONTA	1.10	1.9.1	1.9	1.8	1.7	1.6	1.5	1.4	1.3	1.2	1.1	1.0
Ρ												
Mediat												
or												

RHEL 10.0	Yes	Yes	No	No	No	No	No	No	No	No	No	No
RHEL 9.6	Yes	Yes	No	No	No	No	No	No	No	No	No	No
RHEL 9.5	Comp atible	Yes	Yes	No	No	No	No	No	No	No	No	No
RHEL 9.4	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No
RHEL 9.3	No	Comp atible	Comp atible	Yes	Yes	No	No	No	No	No	No	No
RHEL 9.2	No	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No
RHEL 9.1	No	Comp atible	Comp atible	Yes	Yes	Yes	No	No	No	No	No	No
RHEL 9.0	No	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No
RHEL 8.10	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No
RHEL 8.9	No	Comp atible	Comp atible	Yes	Yes	No	No	No	No	No	No	No
RHEL 8.8	No	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No
RHEL 8.7	No	Comp atible	Comp atible	Yes	Yes	Yes	No	No	No	No	No	No
RHEL 8.6	No	Comp atible	Comp atible	Yes	Yes	Yes	No	No	No	No	No	No
RHEL 8.5	No	Comp atible	Comp atible	Yes	Yes	Yes	Yes	Yes	No	No	No	No
RHEL 8.4	No	Comp atible	Comp atible	Yes	Yes	Yes	Yes	Yes	No	No	No	No
RHEL 8.3	Obsole te	Obsole te	Obsole te	Obsole te	Obsole te	Obsole te	Yes	Yes	Yes	No	No	No

RHEL 8.2	Obsole te	Obsole te	Obsole te	Obsole te	Obsole te	Obsole te	Yes	Yes	Yes	No	No	No
RHEL 8.1	Obsole te	Obsole te	Obsole te	Obsole te	Obsole te	Obsole te	Yes	Yes	Yes	Yes	No	No
RHEL 8.0	Obsole te	Obsole te	Obsole te	Obsole te	Obsole te	Obsole te	Yes	Yes	Yes	Yes	Yes	No
RHEL and CentO S 7.9	Obsole te	Obsole te	Obsole te	Obsole te	Obsole te	Obsole te	Yes	Yes	Yes	Comp atible	No	No
RHEL and CentO S 7.8	Obsole te	Obsole te	Obsole te	Obsole te	Obsole te	Obsole te	Yes	Yes	Yes	Yes	No	No
RHEL and CentO S 7.7	Obsole te	Obsole te	Obsole te	Obsole te	Obsole te	Obsole te	Yes	Yes	Yes	Yes	No	No
RHEL and CentO S 7.6	Obsole te	Obsole te	Obsole te	Obsole te	Obsole te	Obsole te	Yes	Yes	Yes	Yes	Yes	Yes (RHEL only)
CentO S 8 and stream	No	No	No	No	No	No	No	No	No	N/A	N/A	N/A
Rocky Linux 10.0	Yes	No	No	No	No	No	No	No	No	No	No	No
Rocky Linux 9	Yes	Yes	Yes	Yes	Yes	Yes	N/A	N/A	N/A	N/A	N/A	N/A
Rocky Linux 8	Yes	Yes	Yes	Yes	Yes	Yes	N/A	N/A	N/A	N/A	N/A	N/A
Oracle Linux 10	No	No	No	No	No	No	No	No	No	No	No	No

Oracle Linux	No											
9												

- "Yes" means that the OS is recommended for ONTAP Mediator installation and is fully compatible and supported.
- "No" means that the OS and ONTAP Mediator are not compatible.
- "Compatible" means that Red Hat no longer supports these RHEL versions, but ONTAP Mediator can still be installed on them.
- ONTAP Mediator 1.6 adds support for Rocky Linux 9 and 8.
- ONTAP Mediator 1.5 was the last supported release for RHEL 7.x branch operating systems.
- Centos 8 was removed for all releases due to its rebranching. Centos Stream was deemed as not a suitable production target OS. No support is planned.

SCST support matrix

The following table shows the supported SCST version for each version of ONTAP Mediator.

ONTAP Mediator version	Supported SCST version
ONTAP Mediator 1.10	scst-3.9.tar.gz
ONTAP Mediator 1.9.1	scst-3.8.0.tar.bz2
ONTAP Mediator 1.9	scst-3.8.0.tar.bz2
ONTAP Mediator 1.8	scst-3.8.0.tar.bz2
ONTAP Mediator 1.7	scst-3.7.0.tar.bz2
ONTAP Mediator 1.6	scst-3.7.0.tar.bz2
ONTAP Mediator 1.5	scst-3.6.0.tar.bz2
ONTAP Mediator 1.4	scst-3.6.0.tar.bz2
ONTAP Mediator 1.3	scst-3.5.0.tar.bz2
ONTAP Mediator 1.2	scst-3.4.0.tar.bz2
ONTAP Mediator 1.1	scst-3.4.0.tar.bz2
ONTAP Mediator 1.0	scst-3.3.0.tar.bz2

Install or upgrade

ONTAP Mediator installation workflow summary

Installing ONTAP Mediator includes preparing for the installation, providing access to repositories, downloading the installation package, verifying the code signature, installing the ONTAP Mediator package, and performing post-installation configuration tasks.

To install or upgrade ONTAP Mediator, you must ensure all prerequisites are met.



Upgrade host OS and Mediator

If you are upgrading and existing version of ONTAP Mediator, you must first uninstall the previous version, and then install the new version. If you are installing ONTAP Mediator for the first time, you can skip this step.

Provide repository access

You should enable access to repositories so that ONTAP Mediator can access the required packages during the installation process.



Download the ONTAP Mediator installation package

Download the ONTAP Mediator installation package from the ONTAP Mediator download page.

5

Verify the code signature of the ONTAP Mediator installation package

NetApp recommends verifying the ONTAP Mediator code signature before installing the ONTAP Mediator installation package.



Install ONTAP Mediator

To install ONTAP Mediator, you must get the installation package and run the installer on the host.



Verify the ONTAP Mediator installation

After you install ONTAP Mediator, verify that it is running successfully.



Perform post-installation configuration tasks

After ONTAP Mediator is installed and running, additional configuration tasks must be performed to use the ONTAP Mediator features.

Prepare to install or upgrade ONTAP Mediator

To install ONTAP Mediator, you must ensure all prerequisites are met, fetch the installation package, and run the installer on the host. This procedure is used for an installation or an upgrade of an existing installation.

- Beginning with ONTAP 9.8, you can use any version of ONTAP Mediator to monitor an SnapMirror active sync relationship.
- Beginning with ONTAP 9.7, you can use any version of ONTAP Mediator to monitor a MetroCluster IP configuration.

Installation and upgrade considerations

Review the following considerations before you upgrade or install ONTAP Mediator.



ONTAP Mediator 1.8 and earlier is not compatible with Red Hat Enterprise Linux (RHEL) FIPS mode and will prevent it from installing successfully. You can check if FIPS mode is enabled using the fips-mode-setup --check command. You can disable FIPS mode using the fips-modesetup --disable command. Reboot after disabling FIPS mode to successfully install ONTAP Mediator 1.8 or earlier.

- You should upgrade ONTAP Mediator to the latest version that is available. Previous versions of ONTAP Mediator remain backwards compatible with all ONTAP versions but recent versions include security patches for all third-party elements.
- When you upgrade to a new ONTAP Mediator version, the installer automatically upgrades to the recommended SCST version unless a higher version is available. For instructions on manually installing a higher SCST version, see Manage ONTAP Mediator. For supported versions, see the SCST support matrix.
 - If an installation failure occurs, you might need to upgrade to a later version of ONTAP Mediator.
 - From June 15, 2025, you can't install or upgrade ONTAP Mediator 1.9 and 1.8 because their code signing certificates have expired. If the installation or upgrade fails, use the ONTAP Mediator 1.9.1 patch version instead.
- If you install the yum-utils package, you can use the needs-restarting command.
- ONTAP Mediator 1.10 and earlier versions do not support IPv6.

Host requirements

Follow these requirements when installing RHEL or Rocky Linux and configuring the associated repositories.



i

If you modify the installation or configuration process, you might need to perform additional steps.

Linux distribution requirements

- Install RHEL or Rocky Linux according to Red Hat's best practices. Since CentOS 8.x has reached end-oflife, compatible versions of CentOS 8.x are not recommended.
- When installing ONTAP Mediator, ensure the system has access to the required repository so the installation program can retrieve and install all required software dependencies.
- To enable the yum installer to find dependent software in the RHEL repositories, register the system during installation or afterwards using a valid Red Hat subscription.



See the Red Hat Subscription Manager documentation for further information.

Networking requirements

Ensure that the following ports are available and unused for ONTAP Mediator:

Port/services	Source	Direction	Destination	Purpose

22/tcp	Management host	Inbound	ONTAP Mediator	(Optional) SSH / ONTAP Mediator management
31784/tcp	Cluster management LIFs	Inbound	ONTAP Mediator web server	(Required) REST API (HTTPS)
3260/tcp 1	Node Data LIFs or Node Management LIFs	Bidirectional	ONTAP Mediator iSCSI targets	(Required for MCCIP) iSCSI data connection for mailboxes

- 1. For SMBC customers, ONTAP doesn't require port 3260 to be enabled or connected.
 - If using a third-party firewall, refer to Firewall requirements for ONTAP Mediator.
 - For Linux hosts without internet access, make sure the required packages are available in a local repository.

If you are using Link Aggregation Control Protocol (LACP) in a Linux environment, configure the kernel and set the sysctl net.ipv4.conf.all.arp_ignore to 2.

OS requirements

Your OS must meet the following requirements:

- · 64-bit physical installation or virtual machine
- 8 GB RAM
- 1 GB disk space (used for applications installation, server logs, and the database)
- User: Root access

The following table shows the supported OSs for each version of ONTAP Mediator.

ONTAP Mediator version	Supported Linux versions
1.10	Red Hat Enterprise Linux
	 Compatible: 9.5¹
	 Recommended: 10, 9.6, 9.4, and 8.10
	• Rocky Linux 10, 9.6, and 8.10
1.9.1	Red Hat Enterprise Linux
	 Compatible: 9.3, 9.1, 8.9, 8.7, 8.6, 8.5, and 8.4⁻¹
	 Recommended: 9.5, 9.4, 9.2, 9.0, 8.10, and 8.8
	• Rocky Linux 9.5 and 8.10

1.9	 Red Hat Enterprise Linux Compatible: 9.3, 9.1, 8.9, 8.7, 8.6, 8.5, and 8.4¹ Recommended: 9.5, 9.4, 9.2, 9.0, 8.10, and 8.8 Rocky Linux 9.5 and 8.10
1.8	 Red Hat Enterprise Linux: Compatible: 8.7, 8.6, 8.5, and 8.4⁻¹ Recommended: 9.4, 9.3, 9.2, 9.1, 9.0, 8.10, 8.9, and 8.8 Rocky Linux 9.4 and 8.10
1.7	 Red Hat Enterprise Linux: Compatible: 8.7, 8.6, 8.5, and 8.4¹ Recommended: 9.3, 9.2, 9.1, 9.0, 8.9, and 8.8 Rocky Linux 9.3 and 8.9
1.6	 Red Hat Enterprise Linux: Compatible: 8.7, 8.6, 8.5, and 8.4¹ Recommended: 9.2, 9.1, 9.0, and 8.8 Rocky Linux 9.2 and 8.8
1.5	 Red Hat Enterprise Linux: 8.5, 8.4, 8.3, 8.2, 8.1, 8.0, 7.9, 7.8, 7.7, and 7.6 CentOS: 7.9, 7.8, 7.7, and 7.6
1.4	 Red Hat Enterprise Linux: 8.5, 8.4, 8.3, 8.2, 8.1, 8.0, 7.9, 7.8, 7.7, and 7.6 CentOS: 7.9, 7.8, 7.7, and 7.6
1.3	 Red Hat Enterprise Linux: 8.3, 8.2, 8.1, 8.0, 7.9, 7.8, 7.7, and 7.6 CentOS: 7.9, 7.8, 7.7, and 7.6
1.2	 Red Hat Enterprise Linux: 8.1, 8.0, 7.9, 7.8, 7.7, and 7.6 CentOS: 7.9, 7.8, 7.7, and 7.6

1. Compatible means that Red Hat no longer supports these RHEL versions, but ONTAP Mediator can still be installed on them.

OS required packages

The following packages are required by ONTAP Mediator:



The packages are either pre-installed or automatically installed by the ONTAP Mediator installer.

All RHEL/CentOS versions	Additional packages for RHEL 10.x / Rocky Linux 10	Additional packages for RHEL 9.x / Rocky Linux 9	Additional packages for RHEL 8.x / Rocky Linux 8
 openssl 	• python3.12	elfutils-libelf-devel	elfutils-libelf-devel
 openssl-devel 	• python3.12-devel	 policycoreutils-python- 	 policycoreutils-python-
 kernel-devel-\$ 		utils	utils
(uname -r)		 python3 	 redhat-lsb-core
• gcc		 python3-devel 	• python39
• make			 python39-devel
 libselinux-utils 			
• patch			
• bzip2			
 perl-Data-Dumper 			
 perl-ExtUtils- MakeMaker 			
 efibootmgr 			
• mokutil			

The Mediator installation package is a self-extracting compressed tar file that includes:

- An RPM file containing all dependencies that cannot be obtained from the supported release's repository.
- An install script.

A valid SSL certification is recommended.

OS upgrade considerations and kernel compatibility

- All library packages, except the kernel, can safely be updated but might require a reboot to apply the changes within the ONTAP Mediator application. A service window is recommended when a reboot is required.
- You should keep the OS kernel up to date. The kernel core can be upgraded to a version listed as supported in the ONTAP Mediator version matrix. A reboot is mandatory, so you should plan a maintenance window for the outage.
 - You must uninstall the SCST kernel module before rebooting and then re-install it after.
 - You must have a supported version of the SCST ready to reinstall before starting the kernel OS upgrade.



- The kernel version must match the operating system version.
- Upgrading to a kernel beyond the supported OS release for the specific ONTAP Mediator release is not supported. (This likely indicates that the tested SCST module won't compile).

Install ONTAP Mediator when UEFI Secure Boot is enabled

ONTAP Mediator can be installed on a system with or without UEFI Secure Boot enabled.

About this task

You can choose to disable UEFI Secure Boot before installing ONTAP Mediator if it is not needed or if you are troubleshooting ONTAP Mediator installation issues. Disable the UEFI Secure Boot option from your machine settings.



For detailed instructions on disabling UEFI Secure Boot, refer to the documentation for your host OS.

To install ONTAP Mediator with UEFI Secure Boot enabled, you must register a security key before the service can start. The key is generated during the SCST installation's compile step and saved as a private-public key pair on your machine. Use the mokutil utility to add the public key as a Machine Owner Key (MOK) to your UEFI firmware, enabling the system to trust and load the signed module. Save the mokutil passphrase in a secure location as this is required when rebooting your system to activate the MOK.

Steps

1. Check if UEFI Secure Boot is enabled on your system:

```
mokutil --sb-state
```

The results indicate whether UEFI Secure Boot is enabled on this system.

lf	Go to
UEFI secure boot is enabled	The step where you run the mokutil utility
UEFI secure boot is disabled	Upgrade the host operating system and then ONTAP Mediator



 You are prompted to create a passphrase that you must store in a secure location. You'll need this passphrase to enable the key in the UEFI Boot Manager.

- ONTAP Mediator 1.2.0 and earlier versions do not support this mode.
- 2. If the mokutil utility is not installed, run the following command:

yum install mokutil

3. Add the public key to the MOK list:

```
mokutil --import
/opt/netapp/lib/ontap_mediator/ontap_mediator/SCST_mod_keys/scst_module_key.de
r
```



You can leave the private key in its default location or move it to a secure location. However, the public key must be maintained in its existing location for use by the Boot Manager. For further information, see the following README.module-signing file:

```
[root@hostname ~]# ls
/opt/netapp/lib/ontap_mediator/ontap_mediator/SCST_mod_keys/
README.module-signing scst_module_key.der scst_module_key.priv
```

4. Reboot the host and use your device's UEFI Boot Manager to approve the new MOK. You'll need the

passphrase provided for the mokutil utility in the step where you check if UEFI Secure Boot is enabled on your system.

Upgrade the host OS and ONTAP Mediator

To upgrade the host OS for ONTAP Mediator to a later version, you must first uninstall ONTAP Mediator.

About this task

When you upgrade the host OS for ONTAP Mediator to a later major version (for example, from 7.x to 8.x) using the leapp-upgrade tool, you must uninstall ONTAP Mediator because the tool tries to detect new versions of any RPMs that are installed in the repositories that are registered with the system.

Because an .rpm file was installed as part of the ONTAP Mediator installer, it is included in that search. However, because that .rpm file was unpacked as part of the installer and not downloaded from a registered repository, an upgrade cannot be found. In this case, the leapp-upgrade tool uninstalls the package.

In order to preserve the log files, which will be used to triage support cases, you should back up the files before you upgrade the OS and restore them after a reinstall of the ONTAP Mediator package. ONTAP clusters that are connected to it will need to be reconnected after the ONTAP Mediator installation.



The following steps should be performed in order. Immediately after you reinstall ONTAP Mediator, you should stop ontap_mediator, replace the log files, and restart it. This is to ensure that logs are not lost.

Steps

1. Back up the log files.

```
[rootmediator-host ~]# tar -czf ontap_mediator_file_backup.tgz -C
/opt/netapp/lib/ontap_mediator ./log
./ontap_mediator/server_config/ontap_mediator.user_config.yaml
[rootmediator-host ~]# tar -tf ontap_mediator_file_backup.tgz
./log/
./log/ontap_mediator.log
./log/ontap_mediator_stdout.log
./log/ontap_mediator_stdout.log
./log/install_20230419134611.log
./log/scst.log
./log/ontap_mediator_syslog.log
./log/ontap_mediator_syslog.log
./ontap_mediator/server_config/ontap_mediator.user_config.yaml
[rootmediator-host ~]#
```

2. Perform upgrade with leapp-upgrade tool.

```
[rootmediator-host ~]# leapp preupgrade --target 8.4
..<snip upgrade checks>..
..<fix issues found>..
[rootmediator-host ~]# leapp upgrade --target 8.4
..<snip upgrade>..
[rootmediator-host ~]# cat /etc/os-release | head -2
NAME="Red Hat Enterprise Linux"
VERSION="8.4 (Ootpa)"
[rootmediator-host ~]#
```

3. Reinstall ONTAP Mediator.



Perform the rest of the steps immediately after reinstalling ONTAP Mediator to prevent a loss of log files.

```
[rootmediator-host ~]# ontap-mediator-1.10/ontap-mediator-1.10
ONTAP Mediator: Self Extracting Installer
   ..<snip installation>..
[rootmediator-host ~]#
```

4. Stop ontap_mediator.

```
[rootmediator-host ~]# systemctl stop ontap_mediator
[rootmediator-host ~]#
```

5. Replace the log files.

```
[rootmediator-host ~]# tar -xf ontap_mediator_log_backup.tgz -C
/opt/netapp/lib/ontap_mediator
[rootmediator-host ~]#
```

6. Start ontap_mediator.

```
[rootmediator-host ~]# systemctl start ontap_mediator
[rootmediator-host ~]#
```

7. Reconnect all ONTAP clusters to the upgraded ONTAP Mediator

siteA::> metrocluster configuration-settings mediator show Mediator IP Port Node Configuration Connection Status Status _____ 172.31.40.122 31784 siteA-node2 false true siteA-node1 false true siteB-node2 false true siteB-node2 false true siteA::> metrocluster configuration-settings mediator remove Removing the mediator and disabling Automatic Unplanned Switchover. It may take a few minutes to complete. Please enter the username for the mediator: mediatoradmin Please enter the password for the mediator: Confirm the mediator password: Automatic Unplanned Switchover is disabled for all nodes... Removing mediator mailboxes... Successfully removed the mediator. siteA::> metrocluster configuration-settings mediator add -mediator -address 172.31.40.122 Adding the mediator and enabling Automatic Unplanned Switchover. It may take a few minutes to complete. Please enter the username for the mediator: mediatoradmin Please enter the password for the mediator: Confirm the mediator password: Successfully added the mediator. siteA::> metrocluster configuration-settings mediator show Mediator IP Port Node Configuration Connection Status Status _____ _____ 172.31.40.122 31784 siteA-node2 true true siteA-node1 true true siteB-node2 true true siteB-node2 true true siteA::>

For SnapMirror active sync, if you installed your TLS certificate outside of the /opt/netapp directory, then you will not need to reinstall it. If you were using the default generated self-signed certificate or put your custom certificate in the /opt/netapp directory, then you should back it up and restore it.

```
peer1::> snapmirror mediator show
Mediator Address Peer Cluster Connection Status Quorum Status
_____ ____
172.31.49.237 peer2
                            unreachable
                                           true
peer1::> snapmirror mediator remove -mediator-address 172.31.49.237
-peer-cluster peer2
Info: [Job 39] 'mediator remove' job queued
peer1::> job show -id 39
                      Owning
Job ID Name
                      Vserver
                               Node
                                            State
_____ ____
39
  mediator remove peer1 peer1-node1
                                            Success
    Description: Removing entry in mediator
peer1::> security certificate show -common-name ONTAPMediatorCA
Vserver Serial Number Certificate Name
Type
_____ ____
_____
peer1
       4A790360081F41145E14C5D7CE721DC6C210007F
                    ONTAPMediatorCA
server-ca
   Certificate Authority: ONTAP Mediator CA
      Expiration Date: Mon Apr 17 10:27:54 2073
peer1::> security certificate delete -common-name ONTAPMediatorCA *
1 entry was deleted.
peer1::> security certificate install -type server-ca -vserver
peer1
Please enter Certificate: Press <Enter> when done
 .. < snip ONTAP Mediator CA public key>..
You should keep a copy of the CA-signed digital certificate for
future reference.
```

The installed certificate's CA and serial number for reference: CA: ONTAP Mediator CA serial: 44786524464C5113D5EC966779D3002135EA4254 The certificate's generated name for reference: ONTAPMediatorCA peer2::> security certificate delete -common-name ONTAPMediatorCA * 1 entry was deleted. peer2::> security certificate install -type server-ca -vserver peer2 Please enter Certificate: Press <Enter> when done .. < snip ONTAP Mediator CA public key>.. You should keep a copy of the CA-signed digital certificate for future reference. The installed certificate's CA and serial number for reference: CA: ONTAP Mediator CA serial: 44786524464C5113D5EC966779D3002135EA4254 The certificate's generated name for reference: ONTAPMediatorCA peer1::> snapmirror mediator add -mediator-address 172.31.49.237 -peer-cluster peer2 -username mediatoradmin Notice: Enter the mediator password. Enter the password: Enter the password again: Info: [Job: 43] 'mediator add' job queued peer1::> job show -id 43 Owning Job ID Name Vserver Node State _____ _____ 43 mediator add peerl peerl-node2 Success Description: Creating a mediator entry peer1::> snapmirror mediator show Mediator Address Peer Cluster Connection Status Quorum Status _____ ____ 172.31.49.237 peer2 connected true peer1::>

Related information

- security certificate delete
- security certificate install
- security certificate show
- snapmirror mediator add
- snapmirror mediator remove

Provide repository access for ONTAP Mediator installation

You should enable access to repositories so that ONTAP Mediator can access the required packages during the installation process.

Steps

1. Determine which repositories must be accessed, as shown in the following table:

If your operating system is	You must provide access to these repositories
RHEL 10.x	 rhel-10-for-x86_64-baseos-rpms
	 rhel-10-tor-x86_64-appstream-rpms
RHEL 9.x	 rhel-9-for-x86_64-baseos-rpms
	 rhel-9-for-x86_64-appstream-rpms
RHEL 8.x	 rhel-8-for-x86_64-baseos-rpms
	 rhel-8-for-x86_64-appstream-rpms
RHEL 7.x	 rhel-7-server-optional-rpms
CentOS 7.x	C7.6.1810 - Base repository
Rocky Linux 10	• appstream
	• baseos
Rocky Linux 9	• appstream
	• baseos
Rocky Linux 8	• appstream
	• baseos

2. Use one of the following procedures to enable access to the repositories listed above so ONTAP Mediator can access the required packages during the installation process.



If ONTAP Mediator has dependencies on Python modules present in the "extras" and "optional" repositories, it might need to access the rhel-x-for-x86_64-extras-rpms and rhel-x-for-x86_64-optional-rpms files.

Procedure for RHEL 10.x operating system

Use this procedure if your operating system is **RHEL 10.x** to enable access to repositories:

Steps

1. Subscribe to the required repository:

```
subscription-manager repos --enable rhel-10-for-x86 64-baseos-rpms
```

```
subscription-manager repos --enable rhel-10-for-x86_64-appstream-
rpms
```

The following example shows the execution of this command:

```
[root@localhost ~]# subscription-manager repos --enable rhel-10-for-
x86_64-baseos-rpms
Repository 'rhel-10-for-x86_64-baseos-rpms' is enabled for this
system.
[root@localhost ~]# subscription-manager repos --enable rhel-10-for-
x86_64-appstream-rpms
Repository 'rhel-10-for-x86_64-appstream-rpms' is enabled for this
system.
```

2. Run the yum repolist command.

The newly subscribed repositories should appear in the list.

Use this procedure if your operating system is **RHEL 9.x** to enable access to repositories:

Steps

1. Subscribe to the required repository:

```
subscription-manager repos --enable rhel-9-for-x86 64-baseos-rpms
```

```
subscription-manager repos --enable rhel-9-for-x86 64-appstream-rpms
```

The following example shows the execution of this command:

```
[root@localhost ~]# subscription-manager repos --enable rhel-9-for-
x86_64-baseos-rpms
Repository 'rhel-9-for-x86_64-baseos-rpms' is enabled for this
system.
[root@localhost ~]# subscription-manager repos --enable rhel-9-for-
x86_64-appstream-rpms
Repository 'rhel-9-for-x86_64-appstream-rpms' is enabled for this
system.
```

2. Run the yum repolist command.

The newly subscribed repositories should appear in the list.
Use this procedure if your operating system is **RHEL 8.x** to enable access to repositories:

Steps

1. Subscribe to the required repository:

subscription-manager repos --enable rhel-8-for-x86 64-baseos-rpms

```
subscription-manager repos --enable rhel-8-for-x86 64-appstream-rpms
```

The following example shows the execution of this command:

```
[root@localhost ~]# subscription-manager repos --enable rhel-8-for-
x86_64-baseos-rpms
Repository 'rhel-8-for-x86_64-baseos-rpms' is enabled for this
system.
[root@localhost ~]# subscription-manager repos --enable rhel-8-for-
x86_64-appstream-rpms
Repository 'rhel-8-for-x86_64-appstream-rpms' is enabled for this
system.
```

2. Run the yum repolist command.

The newly subscribed repositories should appear in the list.

Use this procedure if your operating system is **RHEL 7.x** to enable access to repositories:

Steps

1. Subscribe to the required repository:

```
subscription-manager repos --enable rhel-7-server-optional-rpms
```

The following example shows the execution of this command:

```
[root@localhost ~]# subscription-manager repos --enable rhel-7-
server-optional-rpms
Repository 'rhel-7-server-optional-rpms' is enabled for this system.
```

2. Run the yum repolist command.

The following example shows the execution of this command. The "rhel-7-server-optional-rpms" repository should appear in the list.

```
[root@localhost ~]# yum repolist
Loaded plugins: product-id, search-disabled-repos, subscription-
manager
rhel-7-server-optional-rpms | 3.2 kB 00:00:00
rhel-7-server-rpms | 3.5 kB 00:00:00
(1/3): rhel-7-server-optional-rpms/7Server/x86 64/group
| 26 kB 00:00:00
(2/3): rhel-7-server-optional-rpms/7Server/x86 64/updateinfo
| 2.5 MB 00:00:00
(3/3): rhel-7-server-optional-rpms/7Server/x86 64/primary db
| 8.3 MB 00:00:01
repo id
                                             repo name
status
rhel-7-server-optional-rpms/7Server/x86 64
                                             Red Hat Enterprise
Linux 7 Server - Optional (RPMs)
                                   19,447
rhel-7-server-rpms/7Server/x86 64
                                             Red Hat Enterprise
Linux 7 Server (RPMs)
                                   26,758
repolist: 46,205
[root@localhost ~]#
```

Use this procedure if your operating system is **CentOS 7.x** to enable access to repositories:



The following examples are showing a repository for CentOS 7.6 and might not work for other CentOS versions. Use the base repository for your version of CentOS.

Steps

- 1. Add the C7.6.1810 Base repository. The C7.6.1810 Base vault repository contains the "kerneldevel" package needed for ONTAP Mediator.
- 2. Add the following lines to /etc/yum.repos.d/CentOS-Vault.repo.

```
[C7.6.1810-base]
name=CentOS-7.6.1810 - Base
baseurl=http://vault.centos.org/7.6.1810/os/$basearch/
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-7
enabled=1
```

3. Run the yum repolist command.

The following example shows the execution of this command. The CentOS-7.6.1810 - Base repository should appear in the list.

```
Loaded plugins: fastestmirror
Loading mirror speeds from cached hostfile
 * base: distro.ibiblio.org
 * extras: distro.ibiblio.org
 * updates: ewr.edge.kernel.org
C7.6.1810-base
                                              | 3.6 kB 00:00:00
(1/2): C7.6.1810-base/x86 64/group gz
                                             | 166 kB 00:00:00
(2/2): C7.6.1810-base/x86 64/primary db
                                             | 6.0 MB 00:00:04
repo id
                            repo name
                                                    status
C7.6.1810-base/x86 64
                            CentOS-7.6.1810 - Base 10,019
base/7/x86 64
                            CentOS-7 - Base
                                                    10,097
extras/7/x86 64
                            CentOS-7 - Extras
                                                    307
updates/7/x86 64
                            CentOS-7 - Updates
                                                    1,010
repolist: 21,433
[root@localhost ~]#
```

Use this procedure if your operating system is **Rocky Linux 10**, **Rocky Linux 9**, or **Rocky Linux 8** to enable access to repositories:

Steps

1. Subscribe to the required repositories:

dnf config-manager --set-enabled baseos

dnf config-manager --set-enabled appstream

2. Perform a clean operation:

dnf clean all

3. Verify the list of repositories:

dnf repolist

Example for Rocky Linux 10

```
[root@localhost ~]# dnf config-manager --set-enabled baseos
[root@localhost ~]# dnf clean all
[root@localhost ~]# dnf repolist
repo id repo name
appstream Rocky Linux 10 - AppStream
baseos Rocky Linux 10 - BaseOS
[root@localhost ~]#
```

Example for Rocky Linux 9

```
[root@localhost ~]# dnf config-manager --set-enabled baseos
[root@localhost ~]# dnf config-manager --set-enabled appstream
[root@localhost ~]# dnf clean all
[root@localhost ~]# dnf repolist
repo id repo name
appstream Rocky Linux 9 - AppStream
baseos Rocky Linux 9 - BaseOS
[root@localhost ~]#
```

```
Example for Rocky Linux 8
```

```
[root@localhost ~]# dnf config-manager --set-enabled baseos
[root@localhost ~]# dnf config-manager --set-enabled appstream
[root@localhost ~]# dnf clean all
[root@localhost ~]# dnf repolist
repo id repo name
appstream Rocky Linux 8 - AppStream
baseos Rocky Linux 8 - BaseOS
[root@localhost ~]#
```

Download the ONTAP Mediator installation package

Download the ONTAP Mediator installation package as part of the installation process.

Steps

1. Download the ONTAP Mediator installation package from the ONTAP Mediator download page.

ONTAP Mediator download page

2. Confirm that the Mediator installation package is in the current working directory:

[root@sdot-r730-0003a-d6 ~] # ls ontap-mediator-1.10.tgz

ontap-mediator-1.10.tgz



For ONTAP Mediator versions 1.4 and earlier, the installer is named ontap-mediator.

If you are at a location without access to the internet, you must ensure that the installer has access to the required packages.

- 3. If necessary, move the Mediator installation package from the download directory to the installation directory on the Linux Mediator host.
- 4. Unzip the installer package:

```
tar xvfz ontap-mediator-1.10.tgz
```

```
ontap-mediator-1.10/
ontap-mediator-1.10/csc-prod-ONTAP-Mediator.pem
ontap-mediator-1.10/csc-prod-chain-ONTAP-Mediator.pem
ontap-mediator-1.10/tsa-prod-ONTAP-Mediator.pem
ontap-mediator-1.10/tsa-prod-chain-ONTAP-Mediator.pem
ontap-mediator-1.10/ONTAP-Mediator-production.pub
ontap-mediator-1.10/ontap-mediator-1.10
ontap-mediator-1.10/ontap-mediator-1.10.sig.tsr
ontap-mediator-1.10/ontap-mediator-1.10.tsr
ontap-mediator-1.10/ontap-mediator-1.10.sig
```

Verify the ONTAP Mediator code signature

NetApp recommends verifying the ONTAP Mediator code signature before installing the ONTAP Mediator installation package; however, this process is optional.

Before you begin

Before verifying the ONTAP Mediator code signature, your system must meet the following requirements.



- From June 15th, 2025, you can't install or upgrade to ONTAP Mediator 1.9 and 1.8 because the code signature verification certificates have expired. Instead, install or upgrade to ONTAP Mediator 1.10.
- If the system doesn't meet the below requirements, the verification process isn't required and you can go directly to Install the ONTAP Mediator installation package.
- openssl versions 1.0.2 to 3.0 for basic verification
- openssl version 1.1.0 or later for Time Stamping Authority (TSA) operations
- · Public internet access for OCSP verification

The following files are included in the download package:

File	Description
ONTAP-Mediator-production.pub	The public key used to verify the signature
csc-prod-chain-ONTAP-Mediator.pem	The public certification CA chain of trust
csc-prod-ONTAP-Mediator.pem	The certificate used to generate the key
ontap-mediator-1.10	The product installation executable for version 1.10
ontap-mediator-1.10.sig	The SHA-256 hashed, then RSA-signed using the csc-prod key, signature for the installer

ontap-mediator-1.10.sig.tsr	The revocation request for use by OCSCP for the installer's signature
ontap-mediator-1.10.tsr	The timestamp signing request file
tsa-prod-ONTAP-Mediator.pem	The public certificate for the TSR
tsa-prod-chain-ONTAP-Mediator.pem	The public certificate CA Chain for the TSR

Steps

- 1. Perform the revocation check on csc-prod-ONTAP-Mediator.pem by using Online Certificate Status Protocol (OCSP).
 - a. Find the OCSP URL used to register the certificate because developer certificates might not provide a uri.

```
openssl x509 -noout -ocsp uri -in csc-prod-chain-ONTAP-Mediator.pem
```

b. Generate an OCSP request for the certificate.

```
openssl ocsp -issuer csc-prod-chain-ONTAP-Mediator.pem -CAfile csc-
prod-chain-ONTAP-Mediator.pem -cert csc-prod-ONTAP-Mediator.pem
-reqout req.der
```

c. Connect to the OCSP Manager to send the OCSP request:

```
openssl ocsp -issuer csc-prod-chain-ONTAP-Mediator.pem -CAfile csc-
prod-chain-ONTAP-Mediator.pem -cert csc-prod-ONTAP-Mediator.pem -url
${ocsp_uri} -resp_text -respont resp.der -verify_other csc-prod-
chain-ONTAP-Mediator.pem
```

2. Verify the trust chain of the CSC and expiration dates against the local host:

```
openssl verify
```



The opensel version from the PATH must have a valid cert.pem (not self-signed).

openssl verify -untrusted csc-prod-chain-ONTAP-Mediator.pem -CApath \${OPENSSLDIR} csc-prod-ONTAP-Mediator.pem # Failure action: The Code-Signature-Check certificate has expired or is invalid. Download a newer version of the ONTAP Mediator. openssl verify -untrusted tsa-prod-chain-ONTAP-Mediator.pem -CApath \${OPENSSLDIR} tsa-prod-ONTAP-Mediator.pem # Failure action: The Time-Stamp certificate has expired or is invalid. Download a newer version of the ONTAP Mediator.

3. Verify the ontap-mediator-1.10.sig.tsr and ontap-mediator-1.10.tsr files using the associated certificates:

openssl ts -verify



.tsr files contain the time stamp response associated with the installer and the code signature. Processing confirms that the time stamp has a valid signature from TSA and that your input file has not changed.

The verification is performed locally on your machine. Independently, there is no need to access TSA servers.

```
openssl ts -verify -data ontap-mediator-1.10.sig -in ontap-mediator-
1.10.sig.tsr -CAfile tsa-prod-chain-ONTAP-Mediator.pem -untrusted tsa-
prod-ONTAP-Mediator.pem
openssl ts -verify -data ontap-mediator-1.10 -in ontap-mediator-1.10.tsr
-CAfile tsa-prod-chain-ONTAP-Mediator.pem -untrusted tsa-prod-ONTAP-
Mediator.pem
```

4. Verify signatures against the key:

```
openssl -dgst -verify
```

openssl dgst -sha256 -verify ONTAP-Mediator-production.pub -signature ontap-mediator-1.10.sig ontap-mediator-1.10

```
[root@scspa2695423001 ontap-mediator-1.10]# pwd
/root/ontap-mediator-1.10
[root@scspa2695423001 ontap-mediator-1.10]# ls -1
total 63660
-r--r-- 1 root root 8582 Feb 19 15:02 csc-prod-chain-ONTAP-
Mediator.pem
-r--r-- 1 root root 2373 Feb 19 15:02 csc-prod-ONTAP-
Mediator.pem
-r-xr-xr-- 1 root root 65132818 Feb 20 15:17 ontap-mediator-1.10
-rw-r--r-- 1 root root 384 Feb 20 15:17 ontap-mediator-1.10.sig
-rw-r--r-- 1 root root 5437 Feb 20 15:17 ontap-mediator-
1.10.sig.tsr
-rw-r--r-- 1 root root 5436 Feb 20 15:17 ontap-mediator-1.10.tsr
-r--r-- 1 root root
                         625 Feb 19 15:02 ONTAP-Mediator-
production.pub
-r--r-- 1 root root 3323 Feb 19 15:02 tsa-prod-chain-ONTAP-
Mediator.pem
-r--r-- 1 root root 1740 Feb 19 15:02 tsa-prod-ONTAP-
Mediator.pem
[root@scspa2695423001 ontap-mediator-1.10]#
[root@scspa2695423001 ontap-mediator-1.10]#
/root/verify ontap mediator signatures.sh
++ openssl version -d
++ cut -d '"' -f2
+ OPENSSLDIR=/etc/pki/tls
+ openssl version
OpenSSL 1.1.1k FIPS 25 Mar 2021
++ openssl x509 -noout -ocsp uri -in csc-prod-chain-ONTAP-Mediator.pem
+ ocsp uri=http://ocsp.entrust.net
+ echo http://ocsp.entrust.net
http://ocsp.entrust.net
+ openssl ocsp -issuer csc-prod-chain-ONTAP-Mediator.pem -CAfile csc-
prod-chain-ONTAP-Mediator.pem -cert csc-prod-ONTAP-Mediator.pem -reqout
req.der
+ openssl ocsp -issuer csc-prod-chain-ONTAP-Mediator.pem -CAfile csc-
prod-chain-ONTAP-Mediator.pem -cert csc-prod-ONTAP-Mediator.pem -url
http://ocsp.entrust.net -resp text -respout resp.der -verify other csc-
prod-chain-ONTAP-Mediator.pem
OCSP Response Data:
    OCSP Response Status: successful (0x0)
    Response Type: Basic OCSP Response
    Version: 1 (0x0)
    Responder Id: C = US, O = "Entrust, Inc.", CN = Entrust Extended
Validation Code Signing CA - EVCS2
```

```
Produced At: Feb 28 05:01:00 2023 GMT
    Responses:
    Certificate ID:
      Hash Algorithm: shal
      Issuer Name Hash: 69FA640329AB84E27220FE0927647B8194B91F2A
      Issuer Key Hash: CE894F8251AA15A28462CA312361D261FBF8FE78
      Serial Number: 511A542B57522AEB7295A640DC6200E5
    Cert Status: good
    This Update: Feb 28 05:00:00 2023 GMT
   Next Update: Mar 4 04:59:59 2023 GMT
    Signature Algorithm: sha512WithRSAEncryption
         3c:1d:49:b0:93:62:37:3e:c7:38:e3:9f:9f:62:82:73:ed:f4:
         ea:00:6b:f1:01:cd:79:57:92:f1:9d:5d:85:9b:60:59:f8:6c:
         e6:f4:50:51:f3:4c:8a:51:dd:50:68:16:8f:20:24:7e:39:b0:
         44:94:8d:b0:61:da:b9:08:36:74:2d:44:55:62:fb:92:be:4a:
         e7:6c:8c:49:dd:0c:fd:d8:ce:20:08:0d:0f:5a:29:a3:19:03:
         9f:d3:df:41:f4:89:0f:73:18:3f:ac:bb:a7:a3:96:7d:c5:70:
         4c:57:cd:17:17:c6:8a:60:d1:37:c9:2d:81:07:2a:d7:a6:02:
         ee:ce:88:16:22:db:e3:43:64:1e:9b:0d:4d:31:66:fa:ab:a5:
         52:99:94:4a:4a:d0:52:c5:34:f5:18:c7:15:5b:ce:74:c2:fc:
         61:ea:55:aa:f1:2f:82:a3:6a:95:8d:7e:2b:38:49:4f:bf:b1:
         68:7b:1b:24:8b:1f:4d:c5:77:f0:71:af:9c:34:c8:7a:82:50:
         09:a2:19:6e:c6:30:4f:da:a2:79:08:f9:d0:ff:85:d9:2a:84:
         cf:0c:aa:75:8f:72:c9:a7:a2:83:e8:8b:cf:ed:0c:69:75:b6:
         2a:7b:6b:58:99:01:d8:34:ad:e1:89:25:27:1b:fa:d9:6d:32:
         97:3a:0b:0a:8e:a3:9e:e3:f4:e0:d6:1a:c9:b5:14:8c:3e:54:
         3b:37:17:1a:93:44:84:8b:4a:87:97:1e:76:43:3e:d3:ec:8b:
         7e:56:4a:3f:01:31:c0:e5:58:fb:50:ce:6f:b1:e7:35:f9:b7:
         a3:ef:6b:3b:21:95:37:a6:5b:8f:f0:15:18:36:65:89:a1:9c:
         9b:69:00:b4:b1:65:6a:bc:11:2d:d4:9b:b4:97:cc:cb:7a:0c:
         16:11:c1:75:58:7e:13:ab:56:3c:3f:93:5b:95:24:c6:54:52:
         1f:86:a9:16:ce:d9:ea:8b:3a:f3:4f:c4:8f:ad:de:e8:3e:3c:
         d2:51:51:ad:33:7f:d8:c5:33:24:26:f1:2d:9d:0e:9f:55:d0:
         68:bf:af:bd:68:4a:40:08:bc:92:a0:62:54:7d:16:7b:36:29:
         15:b1:cd:58:8e:fb:4a:f2:3e:94:8b:fe:56:95:cc:24:32:af:
         5f:71:99:18:ed:0c:64:94:f7:54:48:87:48:d0:6d:b3:42:04:
         96:03:73:a2:8e:8a:6a:b2:af:ee:56:19:a1:c6:35:12:59:ad:
         19:6a:fe:e0:f1:27:cc:96:4e:f0:4f:fb:6a:bd:ce:05:2c:aa:
         79:7c:df:02:5c:ca:53:7d:60:12:88:7c:ce:15:c7:d4:02:27:
         c1:ab:cf:71:30:1e:14:ba
WARNING: no nonce in response
Response verify OK
csc-prod-ONTAP-Mediator.pem: good
        This Update: Feb 28 05:00:00 2023 GMT
        Next Update: Mar 4 04:59:59 2023 GMT
```

```
+ openssl verify -untrusted csc-prod-chain-ONTAP-Mediator.pem -CApath
/etc/pki/tls csc-prod-ONTAP-Mediator.pem
csc-prod-ONTAP-Mediator.pem: OK
+ openssl verify -untrusted tsa-prod-chain-ONTAP-Mediator.pem -CApath
/etc/pki/tls tsa-prod-ONTAP-Mediator.pem
tsa-prod-ONTAP-Mediator.pem: OK
+ openssl ts -verify -data ontap-mediator-1.10.sig -in ontap-mediator-
1.10.sig.tsr -CAfile tsa-prod-chain-ONTAP-Mediator.pem -untrusted tsa-
prod-ONTAP-Mediator.pem
Using configuration from /etc/pki/tls/openssl.cnf
Verification: OK
+ openssl ts -verify -data ontap-mediator-1.10 -in ontap-mediator-
1.10.tsr -CAfile tsa-prod-chain-ONTAP-Mediator.pem -untrusted tsa-prod-
ONTAP-Mediator.pem
Using configuration from /etc/pki/tls/openssl.cnf
Verification: OK
+ openssl dgst -sha256 -verify ONTAP-Mediator-production.pub -signature
ontap-mediator-1.10.sig ontap-mediator-1.10
Verified OK
[root@scspa2695423001 ontap-mediator-1.10]#
```

Install the ONTAP Mediator installation package

To install ONTAP Mediator, you must get the installation package and run the installer on the host.

Steps

1. Run the installer and respond to the prompts as required:

```
./ontap-mediator-1.10/ontap-mediator-1.10 -y
```

[root@scs000099753 ~]# ./ontap-mediator-1.10/ontap-mediator-1.10 -y



To skip the automatic signature check during installation, use the following command: ./ontap-mediator-1.10/ontap-mediator-1.10 -y --skip-code-signature -check

The installation process proceeds to create the required accounts and install required packages. If you have a previous version of Mediator installed on the host, you will be prompted to confirm that you want to upgrade.

2. Beginning with ONTAP Mediator 1.4, the Secure Boot mechanism is enabled on UEFI systems. When Secure Boot is enabled, you must take additional steps to register the security key after installation:

• Follow instructions in the README file to sign the SCST kernel module .:

/opt/netapp/lib/ontap_mediator/ontap_mediator/SCST_mod_keys/README.modulesigning

• Locate the required keys:

/opt/netapp/lib/ontap_mediator/ontap_mediator/SCST_mod_keys



After installation, the README files and key location are also provided in the system output.

```
[root@mediator host ~]# tar -zxvf ontap-mediator-1.10.tgz
ontap-mediator-1.10/
ontap-mediator-1.10/csc-prod-chain-ONTAP-Mediator.pem
ontap-mediator-1.10/csc-prod-ONTAP-Mediator.pem
ontap-mediator-1.10/tsa-prod-ONTAP-Mediator.pem
ontap-mediator-1.10/tsa-prod-chain-ONTAP-Mediator.pem
ontap-mediator-1.10/ONTAP-Mediator-production.pub
ontap-mediator-1.10/ontap-mediator-1.10
ontap-mediator-1.10/ontap-mediator-1.10.sig.tsr
ontap-mediator-1.10/ontap-mediator-1.10.tsr
ontap-mediator-1.10/ontap-mediator-1.10.sig
[root@mediator host ~]#./ontap-mediator-1.10.0/ontap-mediator-1.10.0
ONTAP Mediator: Self Extracting Installer
+ Extracting the ONTAP Mediator installation/upgrade archive
+ Performing the ONTAP Mediator run-time code signature check
   Using openssl from the path: /usr/bin/openssl configured for
CApath:/etc/pki/tls
Error querying OCSP responder
80BBA032607F0000:error:1E800080:HTTP
routines:OSSL HTTP REQ CTX nbio:failed reading
data:crypto/http/http client.c:549:
80BBA032607F0000:error:1E800067:HTTP
routines:OSSL HTTP REQ CTX exchange:error
receiving:crypto/http/http client.c:901:server=http://ocsp.entrust.net:
80
   WARNING: The OCSP check failed while attempting to test the Code-
Signature-Check certificate
   Continue without code signature checking (only recommended if
integrity has been established manually)? y(es)/N(o): yes
 SKIPPING: Code signature check, manual override due to lack of OCSP
response
+ Unpacking the ONTAP Mediator installer
ONTAP Mediator requires two user accounts. One for the service
(netapp), and one for use by ONTAP to the mediator API (mediatoradmin).
Using default account names: netapp + mediatoradmin
Enter ONTAP Mediator user account (mediatoradmin) password:
Re-Enter ONTAP Mediator user account (mediatoradmin) password:
```

```
+ Checking if SELinux is in enforcing mode
The installer will change the SELinux context type of
/opt/netapp/lib/ontap mediator/pyenv/bin/uwsgi from type 'lib t' to
'bin t'.
+ Checking for default Linux firewall
+ Installing required packages.
Updating Subscription Management repositories.
Unable to read consumer identity
This system is not registered with an entitlement server. You can use
"rhc" or "subscription-manager" to register.
Last metadata expiration check: 5 days, 14:34:13 ago on Thu 10 Jul 2025
01:28:32 AM EDT.
Package openssl-1:3.2.2-16.el10.x86 64 is already installed.
Package libselinux-utils-3.8-1.el10.x86 64 is already installed.
Package perl-Data-Dumper-2.189-512.el10.x86 64 is already installed.
Package bzip2-1.0.8-25.el10.x86 64 is already installed.
Package efibootmgr-18-8.el10.x86 64 is already installed.
Package mokutil-2:0.6.0-11.el10.x86 64 is already installed.
Package policycoreutils-python-utils-3.8-1.el10.noarch is already
installed.
Package python3-3.12.9-1.el10.x86 64 is already installed.
Dependencies resolved.
______
______
_____
Package
Architecture
                                    Version
Repository
                                       Size
_______
______
_____
Installing:
elfutils-libelf-devel
x86 64
                                    0.192-5.el10
                                       50 k
AppStream
```

gcc x86 64 AppStream kernel-devel x86 64 AppStream make x86 64 BaseOS openssl-devel x86 64 AppStream patch x86 64 AppStream perl-ExtUtils-MakeMaker noarch AppStream python3-devel x86 64 AppStream 334 k python3-pip noarch AppStream Installing dependencies: annobin-docs noarch AppStream annobin-plugin-gcc x86 64 AppStream bison x86 64 AppStream cmake-filesystem x86 64 AppStream срр x86 64 AppStream dwz x86 64 AppStream efi-srpm-macros noarch

14.2.1-7.el10 37 M 6.12.0-55.9.1.el10_0 22 M 1:4.4.1-9.el10 591 k 1:3.2.2-16.el10 3.9 M 2.7.6-26.el10 134 k 2:7.70-513.el10 297 k 3.12.9-1.el10 23.3.2-7.el10 3.2 M 12.92-1.el10 94 k 12.92-1.el10 985 k 3.8.2-9.el10 1.0 M 3.30.5-2.el10 29 k 14.2.1-7.el10 12 M 0.15-7.el10 139 k 6-6.el10

AppStream flex x86 64 AppStream fonts-srpm-macros noarch AppStream forge-srpm-macros noarch AppStream gcc-plugin-annobin x86 64 AppStream glibc-devel x86 64 AppStream go-srpm-macros noarch AppStream kernel-headers x86 64 AppStream kernel-srpm-macros noarch AppStream libxcrypt-devel x86 64 AppStream libzstd-devel x86 64 AppStream 53 k lua-srpm-macros noarch AppStream m4 x86 64 AppStream ocaml-srpm-macros noarch AppStream openblas-srpm-macros noarch AppStream package-notes-srpm-macros noarch

25 k 2.6.4-19.el10 303 k 1:2.0.5-18.el10 29 k 0.4.0-6.el10 23 k 14.2.1-7.el10 62 k 2.39-37.el10 641 k 3.6.0-4.el10 29 k 6.12.0-55.9.1.el10 0 2.3 M 1.0-25.el10 11 k 4.4.36-10.el10 33 k 1.5.5-9.el10 1-15.el10 10 k 1.4.19-11.el10 309 k 10-4.el10 10 k 2-19.el10 9.0 k

AppStream perl-AutoSplit noarch AppStream perl-Benchmark noarch AppStream perl-CPAN-Meta-Requirements noarch AppStream perl-CPAN-Meta-YAML noarch AppStream perl-Devel-PPPort x86 64 AppStream perl-ExtUtils-Command noarch AppStream perl-ExtUtils-Constant noarch AppStream perl-ExtUtils-Install noarch AppStream perl-ExtUtils-Manifest noarch AppStream perl-ExtUtils-ParseXS noarch AppStream perl-File-Compare noarch AppStream perl-File-Copy noarch AppStream perl-I18N-Langinfo x86 64 AppStream perl-JSON-PP noarch AppStream perl-Test-Harness noarch AppStream

11 k 5.74-512.el10 23 k 1.25-512.el10 28 k 2.143-11.el10 39 k 0.018-512.el10 29 k 3.72-512.el10 223 k 2:7.70-513.el10 16 k 0.25-512.el10 47 k 2.22-511.el10 47 k 1:1.75-511.el10 37 k 1:3.51-512.el10 190 k 1.100.800-512.el10 15 k 2.41-512.el10 22 k 0.24-512.el10 28 k 1:4.16-512.el10 69 k 1:3.48-512.el10 288 k

perl-lib x86 64 AppStream perl-srpm-macros noarch AppStream perl-version x86 64 AppStream pyproject-srpm-macros noarch AppStream python-srpm-macros noarch AppStream python3-pyparsing noarch BaseOS qt6-srpm-macros noarch AppStream 11 k redhat-rpm-config noarch AppStream rust-toolset-srpm-macros noarch AppStream systemtap-sdt-devel x86 64 AppStream systemtap-sdt-dtrace x86 64 AppStream zlib-ng-compat-devel x86 64 AppStream Installing weak dependencies: perl-CPAN-Meta noarch AppStream perl-Encode-Locale noarch AppStream perl-Time-HiRes x86 64

0.65-512.el10 16 k 1-57.el10 9.7 k 8:0.99.32-4.el10 68 k 1.16.2-1.el10 16 k 3.12-9.1.el10 26 k 3.1.1-7.el10 273 k 6.8.1-3.el10 288-1.el10 83 k 1.84.1-1.el10 13 k 5.2-2.el10 78 k 5.2-2.el10 72 k 2.2.3-1.el10 41 k 2.150010-511.el10 202 k 1.05-31.el10 21 k

62 k AppStream perl-devel 4:5.40.1-512.el10 x86 64 772 k AppStream perl-doc noarch 5.40.1-512.el10 4.9 M AppStream Transaction Summary ______ _____ Install 63 Packages Total size: 94 M Installed size: 282 M Downloading Packages: BaseOS Packages Red Hat Enterprise Linux 10 439 kB/s | 3.7 kB 00:00 Importing GPG key 0xFD431D51: Userid : "Red Hat, Inc. (release key 2) <security@redhat.com>" Fingerprint: 567E 347A D004 4ADE 55BA 8A5F 199E 2F91 FD43 1D51 : /etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release From Key imported successfully Importing GPG key 0x5A6340B3: Userid : "Red Hat, Inc. (auxiliary key 3) <security@redhat.com>" Fingerprint: 7E46 2425 8C40 6535 D56D 6F13 5054 E4A4 5A63 40B3 : /etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release From Key imported successfully Running transaction check Transaction check succeeded. Running transaction test Transaction test succeeded. Running transaction Preparing : 1/1Installing : perl-version-8:0.99.32-4.el10.x86 64 1/63 Installing : perl-File-Copy-2.41-512.el10.noarch 2/63 Installing : perl-CPAN-Meta-Requirements-2.143-11.el10.noarch 3/63 Installing : perl-Time-HiRes-4:1.9777-511.el10.x86 64 4/63 Installing : perl-JSON-PP-1:4.16-512.el10.noarch

5/63	
Installing	: perl-File-Compare-1.100.800-512.el10.noarch
0/03	\cdot porl-Extiltilg-DarceVC-1.2 51-512 ell0 pearch
THSCALLING	: peri-Excourts-ParsexS-1:5.51-512.erro.moarch
Installing	• m4-1 4 19-11 e110 x86 64
8/63	. M4 1.4.19 11.0110.800_04
Installing	• make-1•4 4.1-9.el10.x86 64
9/63	
Installing	: bison-3.8.2-9.el10.x86 64
10/63	_
Installing	: flex-2.6.4-19.el10.x86_64
11/63	
Installing	: perl-ExtUtils-Command-2:7.70-513.el10.noarch
12/63	
Installing	: perl-ExtUtils-Manifest-1:1.75-511.el10.noarch
13/63	
Installing	: systemtap-sdt-devel-5.2-2.el10.x86_64
14/63	
Installing	: rust-toolset-srpm-macros-1.84.1-1.el10.noarch
15/63	
Installing	: qt6-srpm-macros-6.8.1-3.el10.noarch
16/63	
Installing	: python3-pip-23.3.2-7.ell0.noarch
I//03	\cdot purproduct array magness 1 16 2-1 ollo possed
18/63	. pyproject-sipm-macros-1.10.2-1.e110.noarch
Installing	• perl-srpm-macros-1-57 ell0 poarch
19/63	· port bipm macrob 1 of orto modelon
Installing	: perl-lib-0.65-512.el10.x86 64
20/63	
Installing	: perl-doc-5.40.1-512.el10.noarch
21/63	
Installing	: perl-I18N-Langinfo-0.24-512.el10.x86_64
22/63	
Installing	: perl-Encode-Locale-1.05-31.el10.noarch
23/63	
Installing	: perl-ExtUtils-Constant-0.25-512.el10.noarch
24/63	
Installing	: perl-Devel-PPPort-3.72-512.el10.x86_64
25/63	
Installing	: perl-CPAN-Meta-YAML-U.U18-512.ellU.noarch
Z0/03	• $perl=CDAN=Meta=2$ 150010=511 0110 percent
27/63	. perr-GrAN-meta-2.130010-311.0110.0041C0
Installing	: perl-Benchmark-1.25-512.el10 poarch

28/63	
Installing	: perl-Test-Harness-1:3.48-512.el10.noarch
29/63	
Installing	: perl-AutoSplit-5.74-512.el10.noarch
30/63	
Installing	: package-notes-srpm-macros-0.5-13.el10.noarch
31/63	
Installing	\cdot openesi-devel-1.3 2 2-16 ello x86 64
22/62	. openssi dever 1.3.2.2 10.0110.x00_04
52/05	
Installing	: openblas-srpm-macros-2-19.ell0.noarch
33/63	
Installing	: ocaml-srpm-macros-10-4.el10.noarch
34/63	
Installing	: lua-srpm-macros-1-15.el10.noarch
35/63	
Installing	: libzstd-devel-1.5.5-9.el10.x86_64
36/63	
Installing	: kernel-srpm-macros-1.0-25.el10.noarch
37/63	-
Installing	: kernel-headers-6.12.0-55.9.1.el10 0.x86 64
38/63	
Tratalling	\cdot libuarumt_double4 4 26-10 pl10 \times 26 64
	: IIDxCrypt-devei-4.4.56-10.e110.x06_64
39/63	
Installing	: glibc-devel-2.39-37.ell0.x86_64
40/63	
Installing	: efi-srpm-macros-6-6.el10.noarch
41/63	
Installing	: dwz-0.15-7.el10.x86_64
42/63	
Installing	: cpp-14.2.1-7.el10.x86_64
43/63	
Installing	: gcc-14.2.1-7.el10.x86 64
44/63	_
Installing	: gcc-plugin-annobin-14.2.1-7.el10.x86 64
45/63	• 900 Fragm annorm 110101 / 00100000_01
Installing	· cmaka-filesystem-3 30 5-2 allo v86 64
	. Chiake IIIesystem 5.50.5 2.0110.800_04
40/03	
Installing	: 211D-ng-compat-dever-2.2.3-1.e110.x86_64
47/63	
Installing	: elfutils-libelf-devel-0.192-5.el10.x86_64
48/63	
Installing	: annobin-docs-12.92-1.el10.noarch
49/63	
Installing	: annobin-plugin-gcc-12.92-1.el10.x86_64
50/63	
Installing	: fonts-srpm-macros-1:2.0.5-18.el10.noarch

51/63 Installing : forge-srpm-macros-0.4.0-6.el10.noarch 52/63 Installing : go-srpm-macros-3.6.0-4.el10.noarch 53/63 Installing : python-srpm-macros-3.12-9.1.el10.noarch 54/63 Installing : redhat-rpm-config-288-1.el10.noarch 55/63 Running scriptlet: redhat-rpm-config-288-1.el10.noarch 55/63 Installing : python3-pyparsing-3.1.1-7.el10.noarch 56/63 Installing : systemtap-sdt-dtrace-5.2-2.el10.x86 64 57/63 Installing : perl-devel-4:5.40.1-512.el10.x86 64 58/63 Installing : perl-ExtUtils-Install-2.22-511.el10.noarch 59/63 Installing : perl-ExtUtils-MakeMaker-2:7.70-513.el10.noarch 60/63 Installing : kernel-devel-6.12.0-55.9.1.el10_0.x86_64 61/63 Running scriptlet: kernel-devel-6.12.0-55.9.1.el10 0.x86 64 61/63 Installing : python3-devel-3.12.9-1.el10.x86 64 62/63 Installing : patch-2.7.6-26.el10.x86 64 63/63 Running scriptlet: patch-2.7.6-26.el10.x86 64 63/63 Installed products updated. Installed: annobin-docs-12.92-1.el10.noarch annobin-plugin-gcc-12.92-1.el10.x86 64 bison-3.8.2-9.el10.x86 64 cmake-filesystem-3.30.5-2.el10.x86 64 cpp-14.2.1-7.el10.x86 64 dwz-0.15-7.el10.x86 64 efi-srpm-macros-6-6.el10.noarch elfutils-libelf-devel-0.192-5.el10.x86 64 flex-2.6.4-19.el10.x86 64 fontssrpm-macros-1:2.0.5-18.el10.noarch forge-srpm-macros-0.4.0-6.el10.noarch gcc-14.2.1-7.el10.x86 64 gcc-plugin-annobin-14.2.1-7.el10.x86_64 glibc-devel-2.39-37.el10.x86_64 gosrpm-macros-3.6.0-4.el10.noarch

```
kernel-devel-6.12.0-55.9.1.el10 0.x86 64 kernel-headers-6.12.0-
55.9.1.el10 0.x86 64 kernel-srpm-macros-1.0-25.el10.noarch
libxcrypt-devel-4.4.36-10.el10.x86 64
                                             libzstd-devel-1.5.5-
9.el10.x86 64
                                             m4-1.4.19-
 lua-srpm-macros-1-15.el10.noarch
11.el10.x86 64
                                      make-1:4.4.1-9.el10.x86 64
ocaml-srpm-macros-10-4.el10.noarch
                                              openblas-srpm-macros-2-
19.el10.noarch
  openssl-devel-1:3.2.2-16.el10.x86 64
                                             package-notes-srpm-
macros-0.5-13.el10.noarch patch-2.7.6-26.el10.x86_64
perl-AutoSplit-5.74-512.el10.noarch
                                            perl-Benchmark-1.25-
512.el10.noarch
perl-CPAN-Meta-2.150010-511.el10.noarch perl-CPAN-Meta-
Requirements-2.143-11.el10.noarch perl-CPAN-Meta-YAML-0.018-
512.el10.noarch perl-Devel-PPPort-3.72-512.el10.x86 64
                                                               perl-
Encode-Locale-1.05-31.el10.noarch
 perl-ExtUtils-Command-2:7.70-513.el10.noarch perl-ExtUtils-Constant-
0.25-512.el10.noarch perl-ExtUtils-Install-2.22-511.el10.noarch
perl-ExtUtils-MakeMaker-2:7.70-513.el10.noarch perl-ExtUtils-Manifest-
1:1.75-511.el10.noarch
perl-ExtUtils-ParseXS-1:3.51-512.el10.noarch perl-File-Compare-
1.100.800-512.el10.noarch perl-File-Copy-2.41-512.el10.noarch
perl-I18N-Langinfo-0.24-512.el10.x86 64
                                        perl-JSON-PP-1:4.16-
512.el10.noarch
  perl-Test-Harness-1:3.48-512.el10.noarch perl-Time-HiRes-
4:1.9777-511.el10.x86 64 perl-devel-4:5.40.1-512.el10.x86 64
perl-doc-5.40.1-512.el10.noarch
                                             perl-lib-0.65-
512.el10.x86 64
perl-srpm-macros-1-57.el10.noarch
                                             perl-version-8:0.99.32-
4.el10.x86 64
                        pyproject-srpm-macros-1.16.2-1.el10.noarch
python-srpm-macros-3.12-9.1.el10.noarch
                                            python3-devel-3.12.9-
1.el10.x86 64
python3-pip-23.3.2-7.el10.noarch
                                             python3-pyparsing-
3.1.1-7.el10.noarch
                              qt6-srpm-macros-6.8.1-3.el10.noarch
redhat-rpm-config-288-1.el10.noarch
                                            rust-toolset-srpm-
macros-1.84.1-1.el10.noarch
 systemtap-sdt-devel-5.2-2.el10.x86 64 systemtap-sdt-dtrace-
5.2-2.el10.x86 64
                     zlib-ng-compat-devel-2.2.3-1.el10.x86 64
Complete!
OS package installations finished
+ Installing ONTAP Mediator. (Log: /root/ontap_mediator.vdizgQ/ontap-
mediator-1.10.0/ontap-mediator-1.10.0/install 20250715160240.log)
    This step will take several minutes. Use the log file to view
progress.
    Sudoer config verified
```

ONTAP Mediator rsyslog and logging rotation enabled + Install successful. (Moving log to /opt/netapp/lib/ontap mediator/log/install 20250715160240.log) + WARNING: This system supports UEFI Secure Boot (SB) is currently disabled on this system. If SB is enabled in the future, SCST will not work unless the following action is taken: Using the keys in /opt/netapp/lib/ontap mediator/ontap mediator/SCST mod keys follow instructions in /opt/netapp/lib/ontap mediator/ontap mediator/SCST mod keys/README.modu le-signing to sign the SCST kernel module. Note that reboot will be needed. SCST will not start automatically when Secure Boot is enabled and not configured properly. + Note: ONTAP Mediator generated a self-signed server certificate for temporary use on this host. If the DNS name or IP address for the host is changed, the certificate will no longer be valid. The default certificates should be replaced with secure trusted certificates signed by a known certificate authority prior to use for production. For more information, see /opt/netapp/lib/ontap mediator/README + Note: ONTAP Mediator uses a kernel module compiled specifically for the current OS. Using 'yum update' to upgrade the kernel might cause service interruption. For more information, see /opt/netapp/lib/ontap mediator/README root@mediator host:~# systemctl status ontap mediator • ontap mediator.service - ONTAP Mediator Loaded: loaded (/etc/systemd/system/ontap mediator.service; enabled; preset: disabled) Active: active (running) since Tue 2025-07-15 16:07:29 EDT; 4min 9s ago Invocation: 395e9479487e4e308be2ae030c800c7f Process: 28745 ExecStartPre=/opt/netapp/lib/ontap mediator/tools/otm logs fs.sh (code=exited, status=0/SUCCESS) Main PID: 28759 (python) Tasks: 1 (limit: 22990) Memory: 66.8M (peak: 68.8M) CPU: 2.865s

```
CGroup: /system.slice/ontap mediator.service
             -28759 /opt/netapp/lib/ontap mediator/pyenv/bin/python
/opt/netapp/lib/ontap mediator/ontap mediator/server
Jul 15 16:07:29 mediator host systemd[1]: Starting
ontap mediator.service - ONTAP Mediator...
Jul 15 16:07:29 mediator host systemd[1]: Started
ontap mediator.service - ONTAP Mediator.
root@mediator host:~# systemctl status mediator-scst
• mediator-scst.service
     Loaded: loaded (/etc/systemd/system/mediator-scst.service;
enabled; preset: disabled)
    Active: active (running) since Tue 2025-07-15 16:07:29 EDT; 4min
15s ago
 Invocation: f1d3be6ca1f9492b943e61872676f384
    Process: 28653 ExecStart=/etc/init.d/scst start (code=exited,
status=0/SUCCESS)
    Process: 28738 ExecStartPost=/usr/sbin/modprobe scst vdisk
(code=exited, status=0/SUCCESS)
   Main PID: 28696 (iscsi-scstd)
     Tasks: 1 (limit: 22990)
    Memory: 5.2M (peak: 35.2M)
        CPU: 547ms
     CGroup: /system.slice/mediator-scst.service
             -28696 /usr/local/sbin/iscsi-scstd
Jul 15 16:07:28 mediator host systemd[1]: Starting mediator-
scst.service...
Jul 15 16:07:29 mediator host iscsi-scstd[28694]: max data seg len
1048576, max queued cmds 2048
Jul 15 16:07:29 mediator host scst[28653]: Loading and configuring SCST
Jul 15 16:07:29 mediator host systemd[1]: Started mediator-
scst.service.
root@mediator host:~#
```

Verify the ONTAP Mediator installation status

After you install ONTAP Mediator, verify that it is running successfully.

Steps

1. View the status of ONTAP Mediator:

```
a. systemctl status ontap_mediator
```

```
[root@scspr1915530002 ~]# systemctl status ontap mediator
 ontap mediator.service - ONTAP Mediator
Loaded: loaded (/etc/systemd/system/ontap mediator.service; enabled;
vendor preset: disabled)
Active: active (running) since Mon 2022-04-18 10:41:49 EDT; 1 weeks 0
days ago
Process: 286710 ExecStop=/bin/kill -s INT $MAINPID (code=exited,
status=0/SUCCESS)
Main PID: 286712 (uwsgi)
Status: "uWSGI is ready"
Tasks: 3 (limit: 49473)
Memory: 139.2M
CGroup: /system.slice/ontap mediator.service
      -286712 /opt/netapp/lib/ontap mediator/pyenv/bin/uwsgi --ini
/opt/netapp/lib/ontap mediator/uwsgi/ontap mediator.ini
      -286716 /opt/netapp/lib/ontap mediator/pyenv/bin/uwsgi --ini
/opt/netapp/lib/ontap mediator/uwsgi/ontap mediator.ini
      └─286717 /opt/netapp/lib/ontap mediator/pyenv/bin/uwsgi --ini
/opt/netapp/lib/ontap mediator/uwsgi/ontap mediator.ini
[root@scspr1915530002 ~]#
```

b. systemctl status mediator-scst

```
[root@scspr1915530002 ~]# systemctl status mediator-scst
Loaded: loaded (/etc/systemd/system/mediator-scst.service;
enabled; vendor preset: disabled)
Active: active (running) since Mon 2022-04-18 10:41:47 EDT; 1
weeks 0 days ago
Process: 286595 ExecStart=/etc/init.d/scst start (code=exited,
status=0/SUCCESS)
Main PID: 286662 (iscsi-scstd)
Tasks: 1 (limit: 49473)
Memory: 1.2M
CGroup: /system.slice/mediator-scst.service
____286662 /usr/local/sbin/iscsi-scstd
[root@scspr1915530002 ~]#
```

2. Confirm the ports that are used by ONTAP Mediator:

netstat

[root@scspr1905507001 ~]# netstat -anlt grep -E '3260 31784'				
tcp	0	0 0.0.0.31784	0.0.0.0:*	LISTEN
tcp	0	0 0.0.0.3260	0.0.0.0:*	LISTEN
tcp6	0	0 :::3260	:::*	LISTEN

Post-installation ONTAP Mediator configuration

After ONTAP Mediator is installed and running, additional configuration tasks must be performed in the ONTAP storage system to use the ONTAP Mediator features:

- To use ONTAP Mediator in a MetroCluster IP configuration, see Configure ONTAP Mediator from a MetroCluster IP configuration.
- To use SnapMirror active sync, see Install ONTAP Mediator and confirm the ONTAP cluster configuration.

Configure ONTAP Mediator security policies

ONTAP Mediator supports several configurable security settings. The default values for all settings are provided in a low_space_threshold_mib: 10 read-only file:

/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediator.user_c
onfig.yaml

All values that are placed in the <code>ontap_mediator.user_config.yaml</code> will override the default values and be maintained across all ONTAP Mediator upgrades.

After you modify ontap mediator.user config.yaml, restart ONTAP Mediator:

```
systemctl restart ontap mediator
```

Modify ONTAP Mediator attributes

The ONTAP Mediator attributes described in this section can be modified if required.



Other default values in the ontap_mediator.config.yaml should not be changed because modified values are not maintained during ONTAP Mediator upgrades.

You modify ONTAP Mediator attributes by copying the required variables to the ontap_mediator.user_config.yaml file to override the default settings.

Install third-party SSL certificates

If you need to replace the default self-signed certificates with third-party SSL certificates, modify certain attributes in the following files:

 /opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediator.con fig.yaml • /opt/netapp/lib/ontap_mediator/uwsgi/ontap_mediator.ini

The variables in these files are used to control the certificate files used by ONTAP Mediator.

ONTAP Mediator 1.9 and later

The default variables listed in the following table are included in the

/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediator.con
fig.yaml file.

Variable	Path
cert_path	<pre>/opt/netapp/lib/ontap_mediator/ontap_m ediator/server_config/ontap_mediator_s erver.crt</pre>
key_path	<pre>/opt/netapp/lib/ontap_mediator/ontap_m ediator/server_config/ontap_mediator_s erver.key</pre>
ca_cert_path	<pre>/opt/netapp/lib/ontap_mediator/ontap_m ediator/server_config/intermediate.crt</pre>
ca_key_path	<pre>/opt/netapp/lib/ontap_mediator/ontap_m ediator/server_config/intermediate.key</pre>
ca_serial_path	<pre>/opt/netapp/lib/ontap_mediator/ontap_m ediator/server_config/intermediate.srl</pre>
cert_valid_days	1095
x509_passin_pwd	pass:ontap

- cert_valid_days is used to set the expiration of client certificates. The maximum value is three years (1095 days).
- x509 passin pwd is the passphrase for the signed client certificate.

The default variables listed in the following table are included in the /opt/netapp/lib/ontap_mediator/uwsgi/ontap_mediator.ini file.

Variable	Path
mediator_cert	<pre>/opt/netapp/lib/ontap_mediator/ontap_m ediator/server_config/ontap_mediator_s erver.crt</pre>
mediator_key	<pre>/opt/netapp/lib/ontap_mediator/ontap_m ediator/server_config/ontap_mediator_s erver.key</pre>
ca_cert_path	<pre>/opt/netapp/lib/ontap_mediator/ontap_m ediator/server_config/intermediate.crt</pre>

ONTAP Mediator 1.8 and earlier

The default variables listed in the following table are included in the

/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediator.con
fig.yaml file.

Variable	Path
cert_path	<pre>/opt/netapp/lib/ontap_mediator/ontap_m ediator/server_config/ontap_mediator_s erver.crt</pre>
key_path	<pre>/opt/netapp/lib/ontap_mediator/ontap_m ediator/server_config/ontap_mediator_s erver.key</pre>
ca_cert_path	<pre>/opt/netapp/lib/ontap_mediator/ontap_m ediator/server_config/ca.crt</pre>
ca_key_path	<pre>/opt/netapp/lib/ontap_mediator/ontap_m ediator/server_config/ca.key</pre>
ca_serial_path	<pre>/opt/netapp/lib/ontap_mediator/ontap_m ediator/server_config/ca.srl</pre>
cert_valid_days	1095
x509_passin_pwd	pass:ontap

- cert_valid_days is used to set the expiration of client certificates. The maximum value is three years (1095 days).
- x509 passin pwd is the passphrase for the signed client certificate.

The default variables listed in the following table are included in the /opt/netapp/lib/ontap_mediator/uwsgi/ontap_mediator.ini file.

Variable	Path
mediator_cert	<pre>/opt/netapp/lib/ontap_mediator/ontap_m ediator/server_config/ontap_mediator_s erver.crt</pre>
mediator_key	<pre>/opt/netapp/lib/ontap_mediator/ontap_m ediator/server_config/ontap_mediator_s erver.key</pre>
ca_cert_path	<pre>/opt/netapp/lib/ontap_mediator/ontap_m ediator/server_config/ca.crt</pre>

If you modify these attributes, restart ONTAP Mediator to apply the changes. For detailed instructions on how to replace default certificates with third-party certificates, refer to Replace self-signed certificates with trusted third party certificates.

Password attack protection

The following settings provide protection against brute-force password guessing attacks.

To enable the feature, set a value for the window seconds and the retry limit.

Examples:

• Provide a 5-minute window for guesses, and then reset the count to zero failures:

authentication lock window seconds: 300

• Lock the account if five failures occur within the window timeframe:

authentication retry limit: 5

• Reduce the impact of brute-force password guessing attacks by setting a delay that occurs prior to rejecting each attempt, which slows the attacks.

authentication_failure_delay_seconds: 5

```
authentication_failure_delay_seconds: 0 # seconds (float) to delay
failed auth attempts prior to response, 0 = no delay
authentication_lock_window_seconds: null # seconds (int) since the
oldest failure before resetting the retry counter, null = no window
authentication_retry_limit: null # number of retries to allow
before locking API access, null = unlimited
```

Password complexity rules

The following fields control the password complexity rules of the ONTAP Mediator API user account.

```
password_min_length: 8
password_max_length: 64
password_uppercase_chars: 0  # min. uppercase characters
password_lowercase_chars: 1  # min. lowercase character
password_special_chars: 1  # min. non-letter, non-digit
password_nonletter_chars: 2  # min. non-letter characters (digits, specials, anything)
```

Control of free space

There are settings that control the required free space on the /opt/netapp/lib/ontap mediator disk.

If the space is lower than the set threshold, the service will issue a warning event.

```
low space threshold mib: 10
```

Control of reserve log space

The RESERVE_LOG_SPACE is controlled by specific settings. By default, the ONTAP Mediator installation creates a separate disk space for the logs. The installer creates a new fixed-size file with a total of 700MB of disk space to be used explicitly for ONTAP Mediator logging.

To disable this feature and use the default disk space, perform the following steps:

1. Change the value of RESERVE_LOG_SPACE from 1 to 0 in the following file:

/opt/netapp/lib/ontap_mediator/tools/mediator_env

- 2. Restart the Mediator:
 - a. cat /opt/netapp/lib/ontap_mediator/tools/mediator_env | grep
 "RESERVE LOG SPACE"

RESERVE LOG SPACE=0

```
b. systemctl restart ontap mediator
```

To re-enable the feature, change the value from 0 to 1 and restart the Mediator.



Toggling between disk spaces does not purge existing logs. All previous logs are backed up and then moved to the current disk space after toggling and restarting the Mediator.

Manage ONTAP Mediator

Manage ONTAP Mediator, including changing user credentials, stopping and re-enabling the service, verifying its health, and installing or uninstalling SCST for host maintenance. You can also manage certificates, such as regenerating self-signed certificates, replacing them with trusted third-party certificates, and troubleshooting certificate-related issues.

Change the username

You can change the username using the following procedure.

About this task

Perform this task on the Linux host where you installed ONTAP Mediator.

If you are unable to reach this command, you might need to run the command using the full path as shown in the following example:

/usr/local/bin/mediator_username

Steps

Change the username by choosing one of the following options:

• **Option (a)**: Run the command mediator_change_user and respond to the prompts as shown in the following example:

• Option (b): Run the following command:

MEDIATOR_USERNAME=mediator MEDIATOR_PASSWORD=mediator2 MEDIATOR NEW USERNAME=mediatoradmin mediator change user

```
[root@mediator-host ~] # MEDIATOR_USERNAME=mediator
MEDIATOR_PASSWORD='mediator2' MEDIATOR_NEW_USERNAME=mediatoradmin
mediator_change_user
The account username has been modified successfully.
[root@mediator-host ~]#
```

Change the password

You can change the password using the following procedure.

About this task

Perform this task on the Linux host where you installed ONTAP Mediator.

If you are unable to reach this command, you might need to run the command using the full path as shown in the following example:

/usr/local/bin/mediator change password

Steps

Change the password by choosing one of the following options:

• Option (a): Run the mediator_change_password command and respond to the prompts as shown in the following example:

• Option (b): Run the following command:

```
MEDIATOR_USERNAME=mediatoradmin MEDIATOR_PASSWORD=mediator1
MEDIATOR_NEW_PASSWORD=mediator2 mediator_change_password
```

The example shows that the password is changed from "mediator1" to "mediator2".

```
[root@mediator-host ~]# MEDIATOR_USERNAME=mediatoradmin
MEDIATOR_PASSWORD=mediator1 MEDIATOR_NEW_PASSWORD=mediator2
mediator_change_password
The password has been updated successfully.
[root@mediator-host ~]#
```

Stop ONTAP Mediator

To stop ONTAP Mediator, perform the following steps:

Steps

1. Stop ONTAP Mediator:

systemctl stop ontap mediator

2. Stop SCST:

systemctl stop mediator-scst

3. Disable ONTAP Mediator and SCST:

systemctl diable ontap mediator mediator-scst

Re-enable ONTAP Mediator

To re-enable ONTAP Mediator, perform the following steps:

Steps

1. Enable ONTAP Mediator and SCST:

systemctl enable ontap_mediator mediator-scst

2. Start SCST:

systemctl start mediator-scst

3. Start ONTAP Mediator:

systemctl start ontap_mediator

Verify ONTAP Mediator is healthy

After you install ONTAP Mediator, verify that it is running successfully.

Steps

- 1. View the status of ONTAP Mediator:
 - a. systemctl status ontap_mediator

```
[root@scspr1915530002 ~]# systemctl status ontap mediator
ontap mediator.service - ONTAP Mediator
Loaded: loaded (/etc/systemd/system/ontap mediator.service; enabled;
vendor preset: disabled)
Active: active (running) since Mon 2022-04-18 10:41:49 EDT; 1 weeks 0
days ago
Process: 286710 ExecStop=/bin/kill -s INT $MAINPID (code=exited,
status=0/SUCCESS)
Main PID: 286712 (uwsgi)
Status: "uWSGI is ready"
Tasks: 3 (limit: 49473)
Memory: 139.2M
CGroup: /system.slice/ontap mediator.service
      -286712 /opt/netapp/lib/ontap mediator/pyenv/bin/uwsgi --ini
/opt/netapp/lib/ontap mediator/uwsgi/ontap mediator.ini
      -286716 /opt/netapp/lib/ontap mediator/pyenv/bin/uwsgi --ini
/opt/netapp/lib/ontap mediator/uwsgi/ontap mediator.ini
      -286717 /opt/netapp/lib/ontap mediator/pyenv/bin/uwsgi --ini
/opt/netapp/lib/ontap mediator/uwsgi/ontap mediator.ini
[root@scspr1915530002 ~]#
```

b. systemctl status mediator-scst

```
[root@scspr1915530002 ~]# systemctl status mediator-scst
Loaded: loaded (/etc/systemd/system/mediator-scst.service;
enabled; vendor preset: disabled)
Active: active (running) since Mon 2022-04-18 10:41:47 EDT; 1
weeks 0 days ago
Process: 286595 ExecStart=/etc/init.d/scst start (code=exited,
status=0/SUCCESS)
Main PID: 286662 (iscsi-scstd)
Tasks: 1 (limit: 49473)
Memory: 1.2M
CGroup: /system.slice/mediator-scst.service
____286662 /usr/local/sbin/iscsi-scstd
[root@scspr1915530002 ~]#
```

2. Confirm the ports that are used by ONTAP Mediator:

netstat

[root@scspr1905507001 ~]# netstat -anlt | grep -E '3260|31784' tcp 0 0 0.0.0.0:31784 0.0.0.0:* LISTEN tcp 0 0 0.0.0.0:3260 0.0.0.0:* LISTEN tcp6 0 0 :::3260 :::* LISTEN

Perform host maintenance

If the VM running ONTAP Mediator needs to be upgraded with a newer kernel, this can cause compatibility issues with the SCST kernel modules used by ONTAP Mediator. In this scenario, NetApp recommends that you manually uninstall and reinstall SCST.

Step 1: Manually uninstall SCST

To uninstall SCST, you need the SCST tar bundle that is used for the installed version of ONTAP Mediator.

Steps

1. Download the appropriate SCST bundle (as shown in the following table) and extract it.

For this version	Use this tar bundle…
ONTAP Mediator 1.10	scst-3.9.tar.gz
ONTAP Mediator 1.9.1	scst-3.8.0.tar.bz2
ONTAP Mediator 1.9	scst-3.8.0.tar.bz2
--------------------	--------------------
ONTAP Mediator 1.8	scst-3.8.0.tar.bz2
ONTAP Mediator 1.7	scst-3.7.0.tar.bz2
ONTAP Mediator 1.6	scst-3.7.0.tar.bz2
ONTAP Mediator 1.5	scst-3.6.0.tar.bz2
ONTAP Mediator 1.4	scst-3.6.0.tar.bz2
ONTAP Mediator 1.3	scst-3.5.0.tar.bz2
ONTAP Mediator 1.1	scst-3.4.0.tar.bz2
ONTAP Mediator 1.0	scst-3.3.0.tar.bz2

a. Access the open source package from the SCST sourceforge downloads.

- b. Select **Download released versions**.
- c. Extract the bundle to your VM.

2. Run the following uninstall commands in the scst directory:

- a. systemctl stop mediator-scst
- b. make scstadm_uninstall
- C. make iscsi_uninstall
- d. make usr_uninstall
- e. make scst_uninstall
- f. depmod

Step 2: Manually install SCST

To manually install SCST, you need the SCST tar bundle that is used for the installed version of ONTAP Mediator (see the SCST table).



Perform this step before you install the ONTAP Mediator. If the SCST version you're using is newer than the version bundled with the ONTAP Mediator installer, the installer skips this step.

1. Run the following install commands in the scst directory:

```
a. make 2releaseb. make scst_installc. make usr install
```

```
d. make iscsi install
```

```
e. make scstadm_install
```

f. depmod

If you're performing a first-time installation and want to pre-install ONTAP Mediator, run the following command before continuing with the next step:

mkdir -p
/opt/netapp/lib/ontap mediator/ontap mediator/SCST mod keys

9. cp scst/src/certs/scst_module_key.der /opt/netapp/lib/ontap_mediator/ontap_mediator/SCST_mod_keys/

```
h. patch /etc/init.d/scst < /opt/netapp/lib/ontap mediator/systemd/scst.patch</pre>
```



If you're performing a first-time installation and want to pre-install SCST before installing ONTAP Mediator, you can skip this step. The ONTAP Mediator installer applies any relevant patches to the SCST components during installation.

- 2. Optionally, if Secure Boot is enabled, before you reboot, perform the following steps:
 - a. Determine each file name for the scst_vdisk, scst, and iscsi_scst modules:

```
[root@localhost ~]# modinfo -n scst_vdisk
[root@localhost ~]# modinfo -n scst
[root@localhost ~]# modinfo -n iscsi_scst
```

b. Determine the kernel release:

```
[root@localhost ~]# uname -r
```

c. Sign each module file with the kernel:

```
[root@localhost ~]# /usr/src/kernels/<KERNEL-RELEASE>/scripts/sign-
file \sha256 \
/opt/netapp/lib/ontap_mediator/ontap_mediator/SCST_mod_keys/scst_modu
le_key.priv \
/opt/netapp/lib/ontap_mediator/ontap_mediator/SCST_mod_keys/scst_modu
le_key.der \
_module-filename_
```

d. Install the UEFI key with the firmware.

Instructions for installing the UEFI key are located at:

/opt/netapp/lib/ontap_mediator/ontap_mediator/SCST_mod_keys/README.module-

signing

The generated UEFI key is located at:

```
/opt/netapp/lib/ontap_mediator/ontap_mediator/SCST_mod_keys/scst_module_key.
der
```

3. Reboot the system:

reboot

Uninstall ONTAP Mediator

If necessary, you can remove ONTAP Mediator.

Before you begin

You must disconnect ONTAP Mediator from ONTAP before removing it.

About this task

Perform this task on the Linux host where you installed ONTAP Mediator.

If you are unable to reach this command, you might need to run the command using the full path as shown in the following example:

/usr/local/bin/uninstall_ontap_mediator

Step

1. Uninstall ONTAP Mediator:

uninstall_ontap_mediator

```
[root@mediator-host ~]# uninstall_ontap_mediator
ONTAP Mediator: Self Extracting Uninstaller
+ Removing ONTAP Mediator. (Log:
/tmp/ontap_mediator.GmRGdA/uninstall_ontap_mediator/remove.log)
+ Remove successful.
[root@mediator-host ~]#
```

Regenerate a temporary self-signed certificate

Beginning with ONTAP Mediator 1.7, you can regenerate a temporary self-signed certificate using the following procedure.



This procedure is only supported on systems running ONTAP Mediator 1.7 or later.

About this task

• Perform this task on the Linux host where you installed ONTAP Mediator.

- You can perform this task only if the generated self-signed certificates have become obsolete due to changes to the hostname or IP address of the host after installing ONTAP Mediator.
- After the temporary self-signed certificate has been replaced by a trusted third-party certificate, you do *not* use this task to regenerate a certificate. The absence of a self-signed certificate will cause this procedure to fail.

Step

To regenerate a new temporary self-signed certificate for the current host, perform the following step:

1. Restart ONTAP Mediator:

```
./make self signed certs.sh overwrite
```

```
[root@xyz000123456 ~]# cd
/opt/netapp/lib/ontap mediator/ontap mediator/server config
[root@xyz000123456 server config]# ./make self signed certs.sh overwrite
Adding Subject Alternative Names to the self-signed server certificate
#
# OpenSSL example configuration file.
Generating self-signed certificates
Generating RSA private key, 4096 bit long modulus (2 primes)
.....++++
e is 65537 (0x010001)
Generating a RSA private key
+
writing new private key to 'ontap mediator server.key'
____
Signature ok
subject=C = US, ST = California, L = San Jose, O = "NetApp, Inc.", OU =
ONTAP Core Software, CN = ONTAP Mediator, emailAddress =
support@netapp.com
Getting CA Private Key
```

Replace self-signed certificates with trusted third-party certificates

If supported, you can replace self-signed certificates with trusted third-party certificates.

()

- Third-party certificates are only supported beginning with ONTAP 9.16.1 and in some earlier ONTAP patch releases. See NetApp Bugs Online Bug ID CONTAP-243278.
- Third-party certificates are only supported on systems running ONTAP Mediator 1.7 or later.

About this task

- Perform this task on the Linux host where you installed ONTAP Mediator.
- You can perform this task if the generated self-signed certificates need to be replaced by certificates obtained from a trusted subordinate certificate authority (CA). To accomplish this, you should have access to a trusted public-key infrastructure (PKI) authority.
- The following image shows the purposes of each ONTAP Mediator certificate.



• The following image shows configuration for the web server setup and ONTAP Mediator setup.

ONTAP Mediator Certificates Root Cert diate Cert erver Gert •root ca.crt intermediate.crt ontap mediator server.crt intermediate.key •ontap_mediator_server.key intermediate.srl ontap_mediator_server.ort ntermediate.crt ontap_mediator_server chain.crt uwsgi/ontap_mediator.ini WebServer Setup set-placeholder = mediator_cert=/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediator_server_chain.crt • set-placeholder = mediator_key=/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediator_server.key • set-placeholder = ca_certificate=/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/root_ca.crt ontap_mediator.user_config.yaml ONTAP Mediator Server Setup cert_path: '/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediator_server.crt' • key_path: '/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediator_server.key' ca_cert_path: '/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/intermediate.crt' · ca_key_path: '/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/intermediate.key' ca_serial_path: '/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/intermediate.srl'

Step 1: Obtain a certificate from a third-party issuing a CA certificate

You can obtain a certificate from a PKI authority using the following procedure.

The following example demonstrates replacing the self-signed certificate actors with the third-party certificate actors located at /opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/.



The example illustrates the necessary criteria for the certificates required for ONTAP Mediator. You can obtain the certificates from a PKI authority in a way that might be different to this procedure. Adjust the procedure according to your business need.

- 1. Create a private key intermediate.key and a configuration file <code>openssl_ca.cnf</code> that will be consumed by the PKI authority to generate a certificate.
 - a. Generate the private key intermediate.key:

Example

openssl genrsa -aes256 -out intermediate.key 4096

b. The configuration file openssl_ca.cnf (located at

/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/openssl_ca.c
nf) defines the properties that the generated certificate must have.

2. Use the private key and configuration file to create a certificate signing request intermediate.csr:

Example:

```
openssl req -key <private_key_name>.key -new -out
<certificate csr name>.csr -config <config file name>.cnf
```

```
[root@scs000216655 server_config]# openssl req -key intermediate.key
-new -config openssl_ca.cnf -out intermediate.csr
Enter pass phrase for intermediate.key:
[root@scs000216655 server_config]# cat intermediate.csr
-----BEGIN CERTIFICATE REQUEST-----
<certificate_value>
-----END CERTIFICATE REQUEST-----
```

3. Send the certificate signing request intermediate.csr to a PKI authority for their signature.

The PKI authority verifies the request and signs the .csr, generating the certificate intermediate.crt. Additionally, you need to obtain the root_intermediate.crt certificate that signed the intermediate.crt certificate from the PKI authority.



For SnapMirror Business Continuity (SM-BC) clusters, you must add the intermediate.crt and root_intermediate.crt certificates to an ONTAP cluster. See Configure ONTAP Mediator and clusters for SnapMirror active sync.

ONTAP Mediator 1.8 and earlier

- 1. Create a private key ca.key and a configuration file openssl_ca.cnf that will be consumed by the PKI authority to generate a certificate.
 - a. Generate the private key ca.key:

Example

openssl genrsa -aes256 -out ca.key 4096

- b. The configuration file openssl_ca.cnf (located at /opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/openssl_ca.c nf) defines the properties that the generated certificate must have.
- 2. Use the private key and configuration file to create a certificate signing request ca.csr:

Example:

```
openssl req -key <private_key_name>.key -new -out
<certificate csr name>.csr -config <config file name>.cnf
```

```
[root@scs000216655 server_config]# openssl req -key ca.key -new
-config openssl_ca.cnf -out ca.csr
Enter pass phrase for ca.key:
[root@scs000216655 server_config]# cat ca.csr
-----BEGIN CERTIFICATE REQUEST-----
<certificate_value>
-----END CERTIFICATE REQUEST-----
```

3. Send the certificate signing request ca.csr to a PKI authority for their signature.

The PKI authority verifies the request and signs the .csr, generating the certificate ca.crt. Additionally, you need to obtain the root_ca.crt that signed the `ca.crt certificate from the PKI authority.



For SnapMirror Business Continuity (SM-BC) clusters, you must add the ca.crt and root_ca.crt certificates to an ONTAP cluster. See Configure ONTAP Mediator and clusters for SnapMirror active sync.

Step 2: Generate a server certificate by signing with a third-party CA certification

A server certificate must be signed by the private key intermediate.key and the third-party certificate intermediate.crt. Additionally, the configuration file

/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/openssl_server.cnf
contains certain attributes that specify the properties required for server certificates issued by OpenSSL.

The following commands can generate a server certificate.

Steps

1. To generate a server certificate signing request (CSR), run the following command from the /opt/netapp/lib/ontap_mediator/ontap_mediator/server_config folder:

```
openssl req -config openssl_server.cnf -extensions v3_req -nodes -newkey
rsa:4096 -sha512 -keyout ontap_mediator_server.key -out
ontap_mediator_server.csr
```

2. To generate a server certificate from the CSR, run the following command from the /opt/netapp/lib/ontap_mediator/ontap_mediator/server_config folder:



These files were obtained from a PKI authority. If you are using a different certificate name, replace intermediate.crt and intermediate.key with the relevant file names.

```
openssl x509 -extfile openssl_server.cnf -extensions v3_req -CA
intermediate.crt -CAkey intermediate.key -CAcreateserial -sha512 -days 1095
-req -in ontap mediator server.csr -out ontap mediator server.crt
```

• The -CAcreateserial option is used to generate the intermediate.srl files.

ONTAP Mediator 1.8 and earlier

A server certificate must be signed by the private key ca.key and the third-party certificate ca.crt. Additionally, the configuration file

/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/openssl_server.cnf
contains certain attributes that specify the properties required for server certificates issued by OpenSSL.

The following commands can generate a server certificate.

Steps

1. To generate a server certificate signing request (CSR), run the following command from the /opt/netapp/lib/ontap_mediator/ontap_mediator/server_config folder:

```
openssl req -config openssl_server.cnf -extensions v3_req -nodes -newkey
rsa:4096 -sha512 -keyout ontap_mediator_server.key -out
ontap_mediator_server.csr
```

2. To generate a server certificate from the CSR, run the following command from the /opt/netapp/lib/ontap_mediator/ontap_mediator/server_config folder:



These files were obtained from a PKI authority. If you are using a different certificate name, replace ca.crt and ca.key with the relevant file names.

```
openssl x509 -extfile openssl_server.cnf -extensions v3_req -CA ca.crt
-CAkey ca.key -CAcreateserial -sha512 -days 1095 -req -in
ontap_mediator_server.csr -out ontap_mediator_server.crt
```

 $^\circ$ The -CAcreateserial option is used to generate the <code>ca.srl</code> files.

Step 3: Replace new third-party CA certificate and server certificate in ONTAP Mediator configuration

The certificate configuration is supplied to ONTAP Mediator in the configuration file located at /opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediator.con fig.yaml. The file includes the following attributes:

```
cert_path:
'/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_medi
ator_server.crt'
key_path:
'/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_medi
ator_server.key'
ca_cert_path:
'/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/intermedia
te.crt'
ca_key_path:
'/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/intermedia
te.key'
ca_serial_path:
'/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/intermedia
te.srl'
```

- cert path and key path are server certificate variables.
- ca_cert_path, ca_key_path, and ca_serial_path are CA certificate variables.

Steps

- 1. Replace all intermediate.* files with the third-party certificates.
- Create a certificate chain from the intermediate.crt and ontap_mediator_server.crt certificates:

```
cat ontap_mediator_server.crt intermediate.crt >
ontap_mediator_server_chain.crt
```

3. Update the /opt/netapp/lib/ontap_mediator/uwsgi/ontap_mediator.ini file.

Update the values of mediator_cert, mediator_key, and ca_certificate:

```
set-placeholder = mediator_cert =
/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediator_
server_chain.crt
set-placeholder = mediator_key =
/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediator_
server.key
set-placeholder = ca_certificate =
/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/root_intermedia
```

- The mediator cert value is the path of the ontap mediator server chain.crt file.
- The mediator_key value is the key path in the ontap_mediator_server.crt file, which is ontap_mediator_server.key.
- ° The ca certificate value is the path of the root intermediate.crt file.
- 4. Verify that the following attributes of the newly generated certificates are set correctly:
 - Linux Group Owner: netapp:netapp
 - Linux permissions: 600
- 5. Restart ONTAP Mediator:

systemctl restart ontap_mediator

ONTAP Mediator 1.8 and earlier

The certificate configuration is supplied to ONTAP Mediator in the configuration file located at /opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediator.con fig.yaml. The file includes the following attributes:

```
cert_path:
'/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_medi
ator_server.crt'
key_path:
'/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_medi
ator_server.key'
ca_cert_path:
'/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ca.crt'
ca_key_path:
'/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ca.key'
ca_serial_path:
'/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ca.srl'
```

• cert path and key path are server certificate variables.

• ca cert path, ca key path, and ca serial path are CA certificate variables.

Steps

- 1. Replace all ca. * files with the third-party certificates.
- 2. Create a certificate chain from the ca.crt and ontap_mediator_server.crt certificates:

cat ontap_mediator_server.crt ca.crt > ontap_mediator_server_chain.crt

3. Update the /opt/netapp/lib/ontap_mediator/uwsgi/ontap_mediator.ini file.

Update the values of mediator_cert, mediator_key, and ca_certificate:

- The mediator_key value is the key path in the ontap_mediator_server.crt file, which is ontap_mediator_server.key.
- ° The ca_certificate value is the path of the root_ca.crt file.
- 4. Verify that the following attributes of the newly generated certificates are set correctly:
 - Linux Group Owner: netapp:netapp
 - ° Linux permissions: 600
- 5. Restart ONTAP Mediator:

systemctl restart ontap_mediator

Step 4: Optionally, use a different path or name for your third-party certificates

You can use third-party certificates with a different name other than intermediate.* or store the third-party certificates in a different location.

Steps

1. Configure the

/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediator. user_config.yaml file to override the default variable values in the ontap_mediator.config.yaml file.

If you obtained intermediate.crt from a PKI authority and you store its private key intermediate.key at the location

/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config, the
ontap mediator.user config.yaml file should look like the following example:



If you used intermediate.crt to sign the ontap_mediator_server.crt certificate, the intermediate.srl file is generated. See Step 2: Generate a server certificate by signing with a third-party CA certification for more information.

```
[root@scs000216655 server config]# cat
ontap mediator.user config.yaml
# This config file can be used to override the default settings in
ontap mediator.config.yaml
# To override a setting, copy the property key from
ontap mediator.config.yaml to this file and
# set the property to the desired value. e.g.,
# The default value for 'default mailboxes per target' is 4 in
ontap mediator.config.yaml
#
# To override this value with 6 mailboxes per target, add the
following key/value pair
# below this comment:
# 'default mailboxes per target': 6
cert path:
'/opt/netapp/lib/ontap mediator/ontap mediator/server config/ontap m
ediator server.crt'
key path:
'/opt/netapp/lib/ontap mediator/ontap mediator/server config/ontap m
ediator server.key'
ca cert path:
'/opt/netapp/lib/ontap mediator/ontap mediator/server config/interme
diate.crt'
ca key path:
'/opt/netapp/lib/ontap mediator/ontap mediator/server config/interme
diate.key'
ca serial path:
'/opt/netapp/lib/ontap mediator/ontap mediator/server config/interme
diate.srl'
```

a. If you are using a certificate structure where the root_intermediate.crt certificate provides an intermediate.crt certificate that signs the ontap_mediator_server.crt certificate, create a certificate chain from the intermediate.crt and ontap_mediator_server.crt certificates:



You should have obtained the intermediate.crt and ontap_mediator_server.crt certificates from a PKI authority earlier in the procedure.

```
cat ontap_mediator_server.crt intermediate.crt >
ontap_mediator_server_chain.crt
```

b. Update the /opt/netapp/lib/ontap mediator/uwsgi/ontap mediator.ini file.

Update the values of mediator_cert, mediator_key, and ca_certificate:

```
set-placeholder = mediator_cert =
/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediat
or_server_chain.crt
```

```
set-placeholder = mediator_key =
/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediat
or server.key
```

```
set-placeholder = ca_certificate =
/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/root_interme
diate.crt
```

- The mediator cert value is the path of the ontap_mediator_server_chain.crt file.
- The mediator_key value is the key path in the ontap_mediator_server.crt file, which is ontap mediator server.key.
- The ca certificate value is the path of the root intermediate.crt file.

i

For SnapMirror Business Continuity (SM-BC) clusters, you must add the intermediate.crt and root_intermediate.crt certificates to an ONTAP cluster. See Configure ONTAP Mediator and clusters for SnapMirror active sync.

- c. Verify that the following attributes of the newly generated certificates are set correctly:
 - Linux Group Owner: netapp:netapp
 - Linux permissions: 600
- 2. Restart ONTAP Mediator when the certificates are updated in the configuration file:

systemctl restart ontap_mediator

ONTAP Mediator 1.8 and earlier

You can use third-party certificates with a different name other than ca.* or store the third-party certificates in a different location.

Steps

1. Configure the

```
/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediator.
user_config.yaml file to override the default variable values in the
ontap_mediator.config.yaml file.
```

If you obtained ca.crt from a PKI authority and you store its private key ca.key at the location /opt/netapp/lib/ontap_mediator/ontap_mediator/server_config, the ontap_mediator.user_config.yaml file should look like the following example:

i

If you used ca.crt to sign the ontap_mediator_server.crt certificate, the ca.srl file is generated. See Step 2: Generate a server certificate by signing with a third-party CA certification for more information.

```
[root@scs000216655 server config]# cat
ontap mediator.user config.yaml
# This config file can be used to override the default settings in
ontap mediator.config.yaml
# To override a setting, copy the property key from
ontap mediator.config.yaml to this file and
# set the property to the desired value. e.g.,
#
# The default value for 'default mailboxes per target' is 4 in
ontap mediator.config.yaml
#
# To override this value with 6 mailboxes per target, add the
following key/value pair
# below this comment:
#
# 'default mailboxes per target': 6
cert path:
'/opt/netapp/lib/ontap mediator/ontap mediator/server config/ontap m
ediator server.crt'
key path:
'/opt/netapp/lib/ontap mediator/ontap mediator/server config/ontap m
ediator server.key'
ca cert path:
'/opt/netapp/lib/ontap mediator/ontap mediator/server config/ca.crt'
ca key path:
'/opt/netapp/lib/ontap mediator/ontap mediator/server config/ca.key'
ca serial path:
'/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ca.srl'
```

a. If you are using a certificate structure where the root_ca.crt certificate provides an ca.crt certificate that signs the ontap_mediator_server.crt certificate, create a certificate chain from the ca.crt and ontap_mediator_server.crt certificates:



You should have obtained the ca.crt and ontap_mediator_server.crt certificates from a PKI authority earlier in the procedure.

cat ontap_mediator_server.crt ca.crt > ontap_mediator_server_chain.crt

b. Update the /opt/netapp/lib/ontap_mediator/uwsgi/ontap_mediator.ini file.

Update the values of mediator_cert, mediator_key, and ca_certificate:

```
set-placeholder = mediator_cert =
/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediat
or_server_chain.crt
set-placeholder = mediator_key =
/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/ontap_mediat
or_server.key
set-placeholder = ca_certificate =
```

```
/opt/netapp/lib/ontap_mediator/ontap_mediator/server_config/root_ca.crt
```

- The mediator cert value is the path of the ontap mediator server chain.crt file.
- The mediator_key value is the key path in the ontap_mediator_server.crt file, which is ontap_mediator_server.key.
- The ca_certificate value is the path of the root_ca.crt file.



For SnapMirror Business Continuity (SM-BC) clusters, you must add the ca.crt and root_ca.crt certificates to an ONTAP cluster. See Configure ONTAP Mediator and clusters for SnapMirror active sync.

- c. Verify that the following attributes of the newly generated certificates are set correctly:
 - Linux Group Owner: netapp:netapp
 - Linux permissions: 600
- 2. Restart ONTAP Mediator when the certificates are updated in the configuration file:

```
systemctl restart ontap_mediator
```

Troubleshoot certificate-related issues

You can check certain properties of the certificates.

Verify certificate expiration

Use the following command to identify the certificate validity range.

```
[root@mediator_host server_config]# openssl x509 -in intermediate.crt
-text -noout
Certificate:
    Data:
...
    Validity
    Not Before: Feb 22 19:57:25 2024 GMT
    Not After : Feb 15 19:57:25 2029 GMT
```

ONTAP Mediator 1.8 and earlier

```
[root@mediator_host server_config]# openssl x509 -in ca.crt -text
-noout
Certificate:
    Data:
...
    Validity
    Not Before: Feb 22 19:57:25 2024 GMT
    Not After : Feb 15 19:57:25 2029 GMT
```

Verify X509v3 extensions in CA certification

Use the following command to verify the X509v3 extensions in the CA certification.

The properties defined within v3_ca in openssl_ca.cnf are displayed as X509v3 extensions in intermediate.crt.

```
[root@mediator host server config]# pwd
/opt/netapp/lib/ontap mediator/ontap mediator/server config
[root@mediator host server config]# cat openssl ca.cnf
. . .
[ v3 ca ]
subjectKeyIdentifier = hash
authorityKeyIdentifier = keyid:always,issuer
basicConstraints = critical, CA:true
keyUsage = critical, cRLSign, digitalSignature, keyCertSign
[root@mediator host server config]# openssl x509 -in intermediate.crt
-text -noout
Certificate:
    Data:
. . .
        X509v3 extensions:
            X509v3 Subject Key Identifier:
9F:06:FA:47:00:67:BA:B2:D4:82:70:38:B8:48:55:B5:24:DB:FC:27
            X509v3 Authority Key Identifier:
keyid:9F:06:FA:47:00:67:BA:B2:D4:82:70:38:B8:48:55:B5:24:DB:FC:27
            X509v3 Basic Constraints: critical
                CA:TRUE
            X509v3 Key Usage: critical
                Digital Signature, Certificate Sign, CRL Sign
```

ONTAP Mediator 1.8 and earlier

The properties defined within v3_ca in openssl_ca.cnf are displayed as X509v3 extensions in ca.crt.

```
[root@mediator host server config]# pwd
/opt/netapp/lib/ontap mediator/ontap mediator/server config
[root@mediator host server config]# cat openssl ca.cnf
. . .
[ v3 ca ]
subjectKeyIdentifier = hash
authorityKeyIdentifier = keyid:always,issuer
basicConstraints = critical, CA:true
keyUsage = critical, cRLSign, digitalSignature, keyCertSign
[root@mediator host server config]# openssl x509 -in ca.crt -text
-noout
Certificate:
    Data:
. . .
        X509v3 extensions:
            X509v3 Subject Key Identifier:
9F:06:FA:47:00:67:BA:B2:D4:82:70:38:B8:48:55:B5:24:DB:FC:27
            X509v3 Authority Key Identifier:
keyid:9F:06:FA:47:00:67:BA:B2:D4:82:70:38:B8:48:55:B5:24:DB:FC:27
            X509v3 Basic Constraints: critical
                CA:TRUE
            X509v3 Key Usage: critical
                Digital Signature, Certificate Sign, CRL Sign
```

Verify X509v3 extensions in server certificate and subject Alt Names

The v3_req properties defined in the openssl_server.cnf configuration file are displayed as X509v3 extensions in the certificate.

In the following example, you can obtain the variables in the alt_names sections by running the commands hostname -I on the Linux VM on which ONTAP Mediator is installed.

Check with your network administrator for the correct values of the variables.

```
[root@mediator host server config]# pwd
/opt/netapp/lib/ontap mediator/ontap mediator/server config
[root@mediator host server config]# cat openssl server.cnf
. . .
[ v3 req ]
basicConstraints = CA:false
extendedKeyUsage
                     = serverAuth
                     = keyEncipherment, dataEncipherment
keyUsaqe
                     = @alt names
subjectAltName
[ alt names ]
DNS.1 = abc.company.com
DNS.2 = abc-v6.company.com
IP.1 = 1.2.3.4
IP.2 = abcd:abcd:abcd:abcd:abcd
[root@mediator host server config]# openssl x509 -in intermediate.crt
-text -noout
Certificate:
   Data:
. . .
        X509v3 extensions:
           X509v3 Basic Constraints:
               CA:FALSE
           X509v3 Extended Key Usage:
                TLS Web Server Authentication
           X509v3 Key Usage:
                Key Encipherment, Data Encipherment
           X509v3 Subject Alternative Name:
                DNS:abc.company.com, DNS:abc-v6.company.com, IP
Address:1.2.3.4, IP Address:abcd:abcd:abcd:abcd:abcd
```

```
ONTAP Mediator 1.8 and earlier
```

```
[root@mediator host server config]# pwd
/opt/netapp/lib/ontap mediator/ontap mediator/server config
[root@mediator host server config]# cat openssl server.cnf
. . .
[ v3 req ]
basicConstraints
                  = CA:false
extendedKeyUsage = serverAuth
keyUsage
                     = keyEncipherment, dataEncipherment
subjectAltName = @alt names
[ alt names ]
DNS.1 = abc.company.com
DNS.2 = abc-v6.company.com
IP.1 = 1.2.3.4
IP.2 = abcd:abcd:abcd:abcd:abcd
[root@mediator host server config]# openssl x509 -in ca.crt -text
-noout
Certificate:
   Data:
. . .
       X509v3 extensions:
           X509v3 Basic Constraints:
               CA:FALSE
           X509v3 Extended Key Usage:
               TLS Web Server Authentication
           X509v3 Key Usage:
               Key Encipherment, Data Encipherment
           X509v3 Subject Alternative Name:
                DNS:abc.company.com, DNS:abc-v6.company.com, IP
Address:1.2.3.4, IP Address:abcd:abcd:abcd:abcd:abcd:abcd
```

Verify that a private key matches with a certificate

You can verify whether a particular private key matches with a certificate.

Use the following OpenSSL commands on the key and certificate respectively.

```
[root@mediator_host server_config]# openssl rsa -noout -modulus -in
intermediate.key | openssl md5
Enter pass phrase for intermediate.key:
(stdin) = 14c6b98b0c7c59012b1de89eee4a9dbc
[root@mediator_host server_config]# openssl x509 -noout -modulus -in
intermediate.crt | openssl md5
(stdin) = 14c6b98b0c7c59012b1de89eee4a9dbc
```

ONTAP Mediator 1.8 and earlier

```
[root@mediator_host server_config]# openssl rsa -noout -modulus -in
ca.key | openssl md5
Enter pass phrase for ca.key:
(stdin) = 14c6b98b0c7c59012b1de89eee4a9dbc
[root@mediator_host server_config]# openssl x509 -noout -modulus -in
ca.crt | openssl md5
(stdin) = 14c6b98b0c7c59012b1de89eee4a9dbc
```

If the -modulus attribute for both match, it indicates that the private key and certificate pair are compatible and can work with each other.

Verify that a server certificate is created from a particular CA certificate

You can use the following command to verify that the server certificate is created from a particular CA certificate.

ONTAP Mediator 1.9 and later

```
[root@mediator_host server_config]# openssl verify -CAfile root_ca.crt
--untrusted intermediate.crt ontap_mediator_server.crt
ontap_mediator_server.crt: OK
[root@mediator_host server_config]#
```

ONTAP Mediator 1.8 and earlier

```
[root@mediator_host server_config]# openssl verify -CAfile ca.crt
ontap_mediator_server.crt
ontap_mediator_server.crt: OK
```

If the Online Certificate Status Protocol (OCSP) validation is being used, use the command openssl-verify.

Maintain the host OS for ONTAP Mediator

For optimal performance, you should maintain the host OS for ONTAP Mediator on a regular basis.

Reboot the host

Reboot the host when the clusters are healthy. While ONTAP Mediator is offline, the clusters are at risk of not being able to react properly to failures. A service window is recommended if a reboot is required.

ONTAP Mediator will automatically resume during a reboot and will re-enter the relationships that were previously configured with ONTAP clusters.

Host package updates

Any library or yum packages (except the kernel) can be safely updated but might require a reboot to take effect. A service window is recommended if a reboot is required.

If you install the yum-utils package, use the needs-restarting command to detect if any package changes require a reboot.

You should reboot if any of the ONTAP Mediator dependencies are updated because they will not take immediate effect on running processes.

Host OS minor kernel upgrades

SCST must be compiled for the kernel that is being used. To update the OS, a maintenance window is required.

Steps

Perform the following steps to upgrade the host OS kernel.

- 1. Stop ONTAP Mediator.
- 2. Uninstall the SCST package. (SCST doesn't provide an upgrade mechanism.)
- 3. Upgrade the OS, and reboot.
- 4. Re-install the SCST package.
- 5. Re-enable ONTAP Mediator.

Host changes to the hostname or IP

About this task

- Perform this task on the Linux host where you installed ONTAP Mediator.
- You can perform this task only if the generated self-signed certificates have become obsolete due to changes to the hostname or IP address of the host after installing ONTAP Mediator.
- After the temporary self-signed certificate has been replaced by a trusted third-party certificate, you do *not* use this task to regenerate a certificate. The absence of a self-signed certificate will cause this procedure to fail.

Step

To regenerate a new temporary self-signed certificate for the current host, perform the following step:

1. Restart ONTAP Mediator:

./make_self_signed_certs.sh overwrite

```
[root@xyz000123456 ~]# cd
/opt/netapp/lib/ontap mediator/ontap mediator/server config
[root@xyz000123456 server_config]# ./make_self_signed_certs.sh overwrite
Adding Subject Alternative Names to the self-signed server certificate
#
# OpenSSL example configuration file.
Generating self-signed certificates
Generating RSA private key, 4096 bit long modulus (2 primes)
.....++++
e is 65537 (0x010001)
Generating a RSA private key
writing new private key to 'ontap mediator server.key'
____
Signature ok
subject=C = US, ST = California, L = San Jose, O = "NetApp, Inc.", OU =
ONTAP Core Software, CN = ONTAP Mediator, emailAddress =
support@netapp.com
Getting CA Private Key
[root@xyz000123456 server config]# systemctl restart ontap mediator
```

Learn about MetroCluster IP site management with ONTAP System Manager

MetroCluster configurations synchronously mirror data and configuration between two ONTAP clusters in separate locations. Beginning with ONTAP 9.8, you can use System Manager as a simplified interface for managing a MetroCluster IP configuration.



You can only perform MetroCluster operations using System Manager in a MetroCluster IP configuration. In a MetroCluster FC configuration, you can still use System Manager to manage each node in your MetroCluster configuration, but you can't perform any MetroCluster-specific operations.

Typically, you set up and configure clusters in a MetroCluster configuration in two separate geographical sites. You then set up peering between the clusters so that they synchronize and share data. The two clusters in the peered network provide bidirectional disaster recovery (DR), where each cluster can be the source and backup of the other cluster. In eight-node or four-node MetroCluster IP configurations, each site consists of storage controllers configured as one or two high availability (HA) pairs.

You can install ONTAP Mediator in a third location to monitor the state of the nodes and their DR partners. ONTAP Mediator can implement a Mediator-assisted unplanned switchover (MAUSO) in the case of a disaster.

You can also perform a negotiated switchover to bring down one of the clusters for planned maintenance. The partner cluster handles all data I/O operations for both clusters until you bring up the cluster that you performed maintenance on and perform a switchback operation.

You can find procedures for setting up and managing a MetroCluster IP configuration using System Manager in the MetroCluster documentation.

Data protection using tape backup

Learn about tape backup of FlexVol volumes with ONTAP

ONTAP supports tape backup and restore through Network Data Management Protocol (NDMP). NDMP allows you to back up data in storage systems directly to tape, resulting in efficient use of network bandwidth. ONTAP supports both dump and SMTape engines for tape backup.

You can perform a dump or SMTape backup or restore by using NDMP-compliant backup applications. Only NDMP version 4 is supported.

Tape backup using dump

Dump is a snapshot-based backup in which your file system data is backed up to tape. The ONTAP dump engine backs up files, directories, and the applicable access control list (ACL) information to tape. You can back up an entire volume, an entire qtree, or a subtree that is not an entire volume or an entire qtree. Dump supports baseline, differential, and incremental backups.

Tape backup using SMTape

SMTape is a snapshot-based disaster recovery solution from ONTAP that backs up blocks of data to tape. You can use SMTape to perform volume backups to tapes. However, you cannot perform a backup at the qtree or subtree level. SMTape supports baseline, differential, and incremental backups.

Beginning with ONTAP 9.13.1, Tape backup using SMTape is supporting with SnapMirror active sync.

Tape backup and restore workflow in ONTAP

You can perform tape backup and restore operations by using an NDMP-enabled backup application.

About this task

The tape backup and restore workflow provides an overview of the tasks that are involved in performing tape backup and restore operations. For detailed information about performing a backup and restore operation, see the backup application documentation.

Steps

- 1. Set up a tape library configuration by choosing an NDMP-supported tape topology.
- 2. Enable NDMP services on your storage system.

You can enable the NDMP services either at the node level or at the storage virtual machine (SVM) level. This depends on the NDMP mode in which you choose to perform the tape backup and restore operation.

3. Use NDMP options to manage NDMP on your storage system.

You can use NDMP options either at the node level or at the SVM level. This depends on the NDMP mode in which you choose to perform the tape backup and restore operation.

You can modify the NDMP options at the node level by using the system services ndmp modify command and at the SVM level by using the vserver services ndmp modify command. Learn more about system services ndmp modify and vserver services ndmp modify in the ONTAP command reference.

4. Perform a tape backup or restore operation by using an NDMP-enabled backup application.

ONTAP supports both dump and SMTape engines for tape backup and restore.

For more information about using the backup application (also called *Data Management Applications* or *DMAs*) to perform backup or restore operations, see your backup application documentation.

Related information

Common NDMP tape backup topologies

Understanding dump engine for FlexVol volumes

Use cases for choosing a tape backup engine

ONTAP supports two backup engines: SMTape and dump. You should be aware of the use cases for the SMTape and dump backup engines to help you choose the backup engine to perform tape backup and restore operations.

Dump can be used in the following cases:

- · Direct Access Recovery (DAR) of files and directories
- Backup of a subset of subdirectories or files in a specific path
- · Excluding specific files and directories during backups
- Preserving backup for long durations

SMTape can be used in the following cases:

- Disaster recovery solution
- Preserving deduplication savings and deduplication settings on the backed up data during a restore operation
- Backup of large volumes

Manage tape drives

Manage tape drives overview

You can verify tape library connections and view tape drive information before performing a tape backup or restore operation. You can use a nonqualified tape drive by emulating this to a qualified tape drive. You can also assign and remove tape aliases in addition to viewing existing aliases.

When you back up data to tape, the data is stored in tape files. File marks separate the tape files, and the files have no names. You specify a tape file by its position on the tape. You write a tape file by using a tape device. When you read the tape file, you must specify a device that has the same compression type that you used to write that tape file.

Commands for managing tape drives, media changers, and tape drive operations in ONTAP

There are commands for viewing information about tape drives and media changers in a cluster, bringing a tape drive online and taking it offline, modifying the tape drive cartridge position, setting and clearing tape drive alias name, and resetting a tape drive. You can also view and reset tape drive statistics.

If you want to	Use this command
Bring a tape drive online	storage tape online
Clear an alias name for tape drive or media changer	storage tape alias clear
Enable or disable a tape trace operation for a tape drive	storage tape trace
Modify the tape drive cartridge position	storage tape position
Reset a tape drive	storagetaperesetImage: Constraint of the storageThis command is available only at the advanced privilege level.
Set an alias name for tape drive or media changer	storage tape alias set
Take a tape drive offline	storage tape offline
View information about all tape drives and media changers	storage tape show
View information about tape drives attached to the cluster	storage tape show-tape-drivesystem node hardware tape drive show

If you want to	Use this command
View information about media changers attached to the cluster	storage tape show-media-changer
View error information about tape drives attached to the cluster	storage tape show-errors
View all ONTAP qualified and supported tape drives attached to each node in the cluster	storage tape show-supported-status
View aliases of all tape drives and media changers attached to each node in the cluster	storage tape alias show
Reset the statistics reading of a tape drive to zero	storage stats tape zero tape_name You must use this command at the nodeshell.
View tape drives supported by ONTAP	storage show tape supported $[-v]$ You must use this command at the nodeshell. You can use the $-v$ option to view more details about each tape drive.
View tape device statistics to understand tape performance and check usage pattern	storage stats tape tape_name You must use this command at the nodeshell.

Learn more about storage tape in the ONTAP command reference.

Use a nonqualified tape drive

You can use a nonqualified tape drive on a storage system if it can emulate a qualified tape drive. It is then treated like a qualified tape drive. To use a nonqualified tape drive, you must first determine whether it emulates any of the qualified tape drives.

About this task

A nonqualified tape drive is one that is attached to the storage system, but not supported or recognized by ONTAP.

Steps

1. View the nonqualified tape drives attached to a storage system by using the storage tape showsupported-status command.

The following command displays tape drives attached to the storage system and the support and qualification status of each tape drive. The nonqualified tape drives are also listed. tape_drive_vendor_name is a nonqualified tape drive attached to the storage system, but not supported by ONTAP.

cluster1::> storage tape show-supported-status -node Node1 Node: Node1 Is Tape Drive Supported Support Status _____ _____ _____ "tape drive vendor name" false Nongualified tape drive Hewlett-Packard C1533A true Qualified Qualified Hewlett-Packard C1553A true Hewlett-Packard Ultrium 1 true Qualified Sony SDX-300C Oualified true Sony SDX-500C Oualified true StorageTek T9840C Dynamically Qualified true StorageTek T9840D Dynamically Qualified true Tandberg LTO-2 HH Dynamically Qualified true

2. Emulate the qualified tape drive.

NetApp Downloads: Tape Device Configuration Files

Related information

What qualified tape drives are

Assign tape aliases in ONTAP

For easy device identification, you can assign tape aliases to a tape drive or medium changer. Aliases provide a correspondence between the logical names of backup devices and a name permanently assigned to the tape drive or medium changer.

Steps

1. Assign an alias to a tape drive or medium changer by using the storage tape alias set command.

Learn more about storage tape alias set in the ONTAP command reference.

You can view the serial number (SN) information about the tape drives by using the system node hardware tape drive show command and about tape libraries by using the system node hardware tape library show commands.

The following command sets an alias name to a tape drive with serial number SN[123456]L4 attached to the node, cluster1-01:

cluster-01::> storage tape alias set -node cluster-01 -name st3
-mapping SN[123456]L4

The following command sets an alias name to a media changer with serial number SN[65432] attached to the node, cluster1-01:

```
cluster-01::> storage tape alias set -node cluster-01 -name mc1
-mapping SN[65432]
```

Related information

What tape aliasing is

Removing tape aliases

Remove tape aliases in ONTAP

You can remove aliases by using the storage tape alias clear command when persistent aliases are no longer required for a tape drive or medium changer.

Steps

1. Remove an alias from a tape drive or medium changer by using the storage tape alias clear command.

Learn more about storage tape alias clear in the ONTAP command reference.

The following command removes the aliases of all tape drives by specifying the scope of the alias clear operation to tape:

cluster-01::>storage tape alias clear -node cluster-01 -clear-scope tape

After you finish

If you are performing a tape backup or restore operation using NDMP, then after you remove an alias from a tape drive or medium changer, you must assign a new alias name to the tape drive or medium changer to continue access to the tape device.

Related information

What tape aliasing is

Assigning tape aliases

Enabling or disabling tape reservations

You can control how ONTAP manages tape device reservations by using the tape.reservations option. By default, tape reservation is turned off.

About this task

Enabling the tape reservations option can cause problems if tape drives, medium changers, bridges, or libraries do not work properly. If tape commands report that the device is reserved when no other storage systems are using the device, this option should be disabled.

Steps

1. To use either the SCSI Reserve/Release mechanism or SCSI Persistent Reservationsor to disable tape reservations, enter the following commandat the clustershell:

scsi selects the SCSI Reserve/Release mechanism.

persistent selects SCSI Persistent Reservations.

off disables tape reservations.

Related information

What tape reservations are

Commands for verifying tape library connections in ONTAP

You can view information about the connection path between a storage system and a tape library configuration attached to the storage system. You can use this information to verify the connection path to the tape library configuration or for troubleshooting issues related to the connection paths.

You can view the following tape library details to verify the tape library connections after adding or creating a new tape library, or after restoring a failed path in a single-path or multipath access to a tape library. You can also use this information while troubleshooting path-related errors or if access to a tape library fails.

- · Node to which the tape library is attached
- Device ID
- NDMP path
- Tape library name
- Target port and initiator port IDs
- · Single-path or multipath access to a tape library for every target or FC initiator port
- Path-related data integrity details, such as "Path Errors" and "Path Qual"
- LUN groups and LUN counts

If you want to	Use this command
View information about a tape library in a cluster	system node hardware tape library show
View path information for a tape library	storage tape library path show
View path information for a tape library for every initiator port	storage tape library path show-by- initiator
View connectivity information between a storage tape library and cluster	storage tape library config show

Learn more about storage tape library in the ONTAP command reference.

About tape drives

Qualified tape drives overview

You must use a qualified tape drive that has been tested and found to work properly on a storage system. You can follow tape aliasing and also enable tape reservations to ensure that only one storage system accesses a tape drive at any particular time.

A qualified tape drive is a tape drive that has been tested and found to work properly on storage systems. You can qualify tape drives for existing ONTAP releases by using the tape configuration file.

Format of the tape configuration file

The tape configuration file format consists of fields such as vendor ID, product ID, and details of compression types for a tape drive. This file also consists of optional fields for enabling the autoload feature of a tape drive and changing the command timeout values of a tape drive.

Item	Size	Description
vendor_id (string)	up to 8 bytes	The vendor ID as reported by the SCSI Inquiry command.
product_id(string)	up to 16 bytes	The product ID as reported by the SCSI Inquiry command.
id_match_size(number)		The number of bytes of the product ID to be used for matching to detect the tape drive to be identified, beginning with the first character of the product ID in the Inquiry data.
vendor_pretty (string)	up to 16 bytes	If this parameter is present, it is specified by the string displayed by the command, storage tape show -device-names; otherwise, INQ_VENDOR_ID is displayed.
product_pretty(string)	up to 16 bytes	If this parameter is present, it is specified by the string displayed by the command, storage tape show -device-names; otherwise, INQ_PRODUCT_ID is displayed.

The following table displays the format of the tape configuration file:



The <code>vendor_pretty</code> and <code>product_pretty</code> fields are optional, but if one of these fields has a value, the other must also have a value.

The following table explains the description, density code, and compression algorithm for the various compression types, such as 1, m, h, and a:

Item	Size	Description
{l m h a}_description=(string)	up to 24 bytes	The string to print for the nodeshell command, sysconfig -t, that describes characteristics of the particular density setting.
<pre>{l m h a}_density=(hex codes)</pre>		The density code to be set in the SCSI mode page block descriptor corresponding to the desired density code for I, m, h, or a.
<pre>{l m h a}_algorithm=(hex codes)</pre>		The compression algorithm to be set in the SCSI Compression Mode Page corresponding to the density code and the desired density characteristic.

The following table describes the optional fields available in the tape configuration file:

Field	Description
autoload=(Boolean yes/no)	This field is set to yes if the tape drive has an automatic loading feature; that is, after tape cartridge is inserted, the tape drive becomes ready without the need to execute a SCSI load (start/stop unit) command. The default for this field is no.
cmd_timeout_0x	Individual timeout value. You must use this field only if you want to specify a different timeout value from the one being used as a default by the tape driver. The sample file lists the default SCSI command timeout values used by the tape drive. The timeout value can be expressed in minutes (m), seconds (s), or milliseconds (ms).(i)You should not change this field.

You can download and view the tape configuration file from the NetApp Support Site.

Example of a tape configuration file format

The tape configuration file format for the HP LTO5 ULTRIUM tape drive is as follows:

vendor_id="HP"

product_id="Ultrium 5-SCSI"

id_match_size=9

vendor pretty="Hewlett-Packard"

product pretty="LTO-5"

- l description="LTO-3(ro)/4 4/800GB"
- l_density=0x00
- l_algorithm=0x00
- m description="LTO-3(ro)/4 8/1600GB cmp"

m_density=0x00

m_algorithm=0x01

h_description="LTO-5 1600GB"

h_density=0x58

h algorithm=0x00

- a description="LTO-5 3200GB cmp"
- a density=0x58
- a_algorithm=0x01

autoload="yes"

Related information

NetApp Tools: Tape Device Configuration Files

How the storage system qualifies a new tape drive dynamically

The storage system qualifies a tape drive dynamically by matching its vendor ID and product ID with the information contained in the tape qualification table.

When you connect a tape drive to the storage system, it looks for a vendor ID and product ID match between the information obtained during tape discovery and the information in the internal tape qualification table. If the storage system discovers a match, it marks the tape drive as qualified and can access the tape drive. If the storage system cannot find a match, the tape drive remains in the unqualified state and is not accessed.

Tape devices overview

Tape devices overview

A tape device is a representation of a tape drive. It is a specific combination of rewind
type and compression capability of a tape drive.

A tape device is created for each combination of rewind type and compression capability. Therefore, a tape drive or tape library can have several tape devices associated with it. You must specify a tape device to move, write, or read tapes.

When you install a tape drive or tape library on a storage system, ONTAP creates tape devices associated with the tape drive or tape library.

ONTAP detects tape drives and tape libraries and assigns logical numbers and tape devices to them. ONTAP detects the Fibre Channel, SAS, and parallel SCSI tape drives and libraries when they are connected to the interface ports. ONTAP detects these drives when their interfaces are enabled.

Tape device name format

Each tape device has an associated name that appears in a defined format. The format includes information about the type of device, rewind type, alias, and compression type.

The format of a tape device name is as follows:

rewind_type st alias_number compression_type

rewind_type is the rewind type.

The following list describes the various rewind type values:

• r

ONTAP rewinds the tape after it finishes writing the tape file.

• nr

ONTAP does not rewind the tape after it finishes writing the tape file. You must use this rewind type when you want to write multiple tape files on the same tape.

• ur

This is the unload/reload rewind type. When you use this rewind type, the tape library unloads the tape when it reaches the end of a tape file, and then loads the next tape, if there is one.

You must use this rewind type only under the following circumstances:

- The tape drive associated with this device is in a tape library or is in a medium changer that is in the library mode.
- The tape drive associated with this device is attached to a storage system.
- Sufficient tapes for the operation that you are performing are available in the library tape sequence defined for this tape drive.



If you record a tape using a no-rewind device, you must rewind the tape before you read it.

st is the standard designation for a tape drive.

alias_number is the alias that ONTAP assigns to the tape drive. When ONTAP detects a new tape drive,

ONTAP assigns an alias to the tape drive.

compression type is a drive-specific code for the density of data on the tape and the type of compression.

The following list describes the various values for compression_type:

• a

Highest compression

• h

High compression

• m

Medium compression

• |

Low compression

Examples

nrst0a specifies a no-rewind device on tape drive 0 using the highest compression.

Example of a listing of tape devices

The following example shows the tape devices associated with HP Ultrium 2-SCSI:

```
Tape drive (fc202_6:2.126L1) HP Ultrium 2-SCSI

rst01 - rewind device, format is: HP (200GB)

nrst01 - no rewind device, format is: HP (200GB)

urst01 - unload/reload device, format is: HP (200GB)

rst0m - rewind device, format is: HP (200GB)

urst0m - no rewind device, format is: HP (200GB)

urst0m - unload/reload device, format is: HP (200GB)

rst0h - rewind device, format is: HP (200GB)

nrst0h - no rewind device, format is: HP (200GB)

urst0h - no rewind device, format is: HP (200GB)

urst0h - no rewind device, format is: HP (200GB)

urst0h - no rewind device, format is: HP (200GB)

rst0a - no rewind device, format is: HP (200GB)

urst0a - no rewind device, format is: HP (400GB w/comp)

urst0a - unload/reload device, format is: HP (400GB w/comp)

urst0a - unload/reload device, format is: HP (400GB w/comp)
```

The following list describes the abbreviations in the preceding example:

- GB—Gigabytes; this is the capacity of the tape.
- w/comp—With compression; this shows the tape capacity with compression.

ONTAP supports a maximum of 64 simultaneous tape drive connections, 16 medium changers, and 16 bridge or router devices for each storage system (per node) in any mix of Fibre Channel, SCSI, or SAS attachments.

Tape drives or medium changers can be devices in physical or virtual tape libraries or stand-alone devices.



Although a storage system can detect 64 tape drive connections, the maximum number of backup and restore sessions that can be performed simultaneously depends upon the scalability limits of the backup engine.

Related information

Scalability limits for dump backup and restore sessions

Tape aliasing

Tape aliasing overview

Aliasing simplifies the process of device identification. Aliasing binds a physical path name (PPN) or a serial number (SN) of a tape or a medium changer to a persistent, but modifiable alias name.

The following table describes how tape aliasing enables you to ensure that a tape drive (or tape library or medium changer) is always associated with a single alias name:

Scenario	Reassigning of the alias
When the system reboots	The tape drive is automatically reassigned its previous alias.
When a tape device moves to another port	The alias can be adjusted to point to the new address.
When more than one system uses a particular tape device	The user can set the alias to be the same for all the systems.



When you upgrade from Data ONTAP 8.1.x to Data ONTAP 8.2.x, the tape alias feature of Data ONTAP 8.2.x modifies the existing tape alias names. In such a case you might have to update the tape alias names in the backup application.

Assigning tape aliases provides a correspondence between the logical names of backup devices (for example, st0 or mc1) and a name permanently assigned to a port, a tape drive, or a medium changer.



st0 and st00 are different logical names.



Logical names and serial numbers are used only to access a device. After the device is accessed, it returns all error messages by using the physical path name.

There are two types of names available for aliasing: physical path name and serial number.

What physical path names are

Physical path names (PPNs) are the numerical address sequences that ONTAP assigns to tape drives and tape libraries based on the SCSI-2/3 adapter or switch (specific location) they are connected to the storage system. PPNs are also known as electrical names.

PPNs of direct-attached devices use the following format: host_adapter.device_id_lun



The LUN value is displayed only for tape and medium changer devices whose LUN values are not zero; that is, if the LUN value is zero the lun part of the PPN is not displayed.

For example, the PPN 8.6 indicates that the host adapter number is 8, the device ID is 6, and the logical unit number (LUN) is 0.

SAS tape devices are also direct-attached devices. For example, the PPN 5c.4 indicates that in a storage system, the SAS HBA is connected in slot 5, SAS tape is connected to port C of the SAS HBA, and the device ID is 4.

PPNs of Fibre Channel switch-attached devices use the following format: switch:port_id. device_id_lun

For example, the PPN MY_SWITCH:5.3L2 indicates that the tape drive connected to port 5 of a switch called MY_SWITCH is set with device ID 3 and has the LUN 2.

The LUN (logical unit number) is determined by the drive. Fibre Channel, SCSI tape drives and libraries, and disks have PPNs.

PPNs of tape drives and libraries do not change unless the name of the switch changes, the tape drive or library moves, or the tape drive or library is reconfigured. PPNs remain unchanged after reboot. For example, if a tape drive named MY_SWITCH:5.3L2 is removed and a new tape drive with the same device ID and LUN is connected to port 5 of the switch MY_SWITCH, the new tape drive would be accessible by using MY_SWITCH:5.3L2.

What serial numbers are

A serial number (SN) is a unique identifier for a tape drive or a medium changer. ONTAP generates aliases based on SN instead of the WWN.

Since the SN is a unique identifier for a tape drive or a medium changer, the alias remains the same regardless of the multiple connection paths to the tape drive or medium changer. This helps storage systems to track the same tape drive or medium changer in a tape library configuration.

The SN of a tape drive or a medium changer does not change even if you rename the Fibre Channel switch to which the tape drive or medium changer is connected. However, in a tape library if you replace an existing tape drive with a new one, then ONTAP generates new aliases because the SN of the tape drive changes. Also, if you move an existing tape drive to a new slot in a tape library or remap the tape drive's LUN, ONTAP generates a new alias for that tape drive.



You must update the backup applications with the newly generated aliases.

The SN of a tape device uses the following format: SN [XXXXXXXX] L [X]

x is an alphanumeric character and Lx is the LUN of the tape device. If the LUN is 0, the Lx part of the string is not displayed.

Each SN consists of up to 32 characters; the format for the SN is not case-sensitive.

Considerations when configuring multipath tape access in ONTAP

You can configure two paths from the storage system to access the tape drives in a tape library. If one path fails, the storage system can use the other paths to access the tape drives without having to immediately repair the failed path. This ensures that tape operations can be restarted.

You must consider the following when configuring multipath tape access from your storage system:

• In tape libraries that support LUN mapping, for multipath access to a LUN group, LUN mapping must be symmetrical on each path.

Tape drives and media changers are assigned to LUN groups (set of LUNs that share the same initiator path set) in a tape library. All tape drives of a LUN group must be available for backup and restore operations on all multiple paths.

- A maximum of two paths can be configured from the storage system to access the tape drives in a tape library.
- Multipath tape access supports load balancing. Load balancing is disabled by default.

In the following example, the storage system accesses LUN group 0 through two initiator paths: 0b and 0d. In both these paths, the LUN group has the same LUN number, 0, and LUN count, 5. The storage system accesses LUN group 1 through only one initiator path, 3d.

STSW-3070-2_cluster::> storage tape library config show				
Node Target Port Initiato	LUN Group r	LUN Count	Library Name Library	
STSW-3070-2_cluster-0 510a09800000412d	1 0 0b	5	IBM 3573-TL_1	
0d	1	2	TBM 3573-TI 2	
50050763124b4d6f	3d	-		
3 entries were displa	yed			

Learn more about storage tape library config show in the ONTAP command reference.

How you add tape drives and libraries to storage systems

You can add tape drives and libraries to storage system dynamically (without taking the

storage system offline).

When you add a new medium changer, the storage system detects its presence and adds it to the configuration. If the medium changer is already referenced in the alias information, no new logical names are created. If the library is not referenced, the storage system creates a new alias for the medium changer.

In a tape library configuration, you must configure a tape drive or medium changer on LUN 0 of a target port for ONTAP to discover all medium changers and tape drives on that target port.

What tape reservations are

Multiple storage systems can share access to tape drives, medium changers, bridges, or tape libraries. Tape reservations ensure that only one storage system accesses a device at any particular time by enabling either the SCSI Reserve/Release mechanism or SCSI Persistent Reservations for all tape drives, medium changers, bridges, and tape libraries.



All the systems that share devices in a library, whether switches are involved or not, must use the same reservation method.

The SCSI Reserve/Release mechanism for reserving devices works well under normal conditions. However, during interface error recovery procedures, reservations can be lost. If this occurs, initiators other than the reserved owner can access the device.

Reservations made with SCSI Persistent Reservations are not affected by error recovery mechanisms, such as loop reset or target reset; however, not all devices implement SCSI Persistent Reservations correctly.

Transfer data between storage systems

Transfer data using ndmpcopy

The ndmpcopy nodeshell command transfers data between storage systems that support NDMP v4. You can perform both full and incremental data transfers. You can transfer full or partial volumes, gtrees, directories, or individual files.

About this task

Using ONTAP 8.x and earlier releases, incremental transfers are limited to a maximum of two levels (one full and up to two incremental backups).

Beginning with ONTAP 9.0 and later releases, incremental transfers are limited to a maximum of nine levels (one full and up to nine incremental backups).

You can run ndmpcopy at the nodeshell command line of the source and destination storage systems, or a storage system that is neither the source nor the destination of the data transfer. You can also run ndmpcopy on a single storage system that is both the source and the destination of the data transfer.

You can use IPv4 or IPv6 addresses of the source and destination storage systems in the ndmpcopy command. The path format is /vserver_name/volume_name \[path\].

Learn more about the options you can use with the ndmpcopy command.

Steps

1. Enable NDMP service on the source and destination storage systems:

If you are performing data transfer at the source or destination in	Use the following command		
SVM-scoped NDMP mode	vserver	services ndmp on For NDMP authentication in the admin SVM, the user account is admin and the user role is admin or backup. In the data SVM, the user account is vsadmin and the user role is vsadmin or vsadmin- backup role.	
Node-scoped NDMP mode	system services ndmp on		

2. Transfer data within a storage system or between storage systems using the ndmpcopy command at the nodeshell:

```
::> system node run -node <node_name> < ndmpcopy [options]
source_IP:source_path destination_IP:destination_path [-mcs {inet|inet6}] [-
mcd {inet|inet6}] [-md {inet|inet6}]</pre>
```



DNS names are not supported in ndmpcopy. You must provide the IP address of the source and the destination. The loopback address (127.0.0.1) is not supported for the source IP address or the destination IP address.

- The ndmpcopy command determines the address mode for control connections as follows:
 - The address mode for control connection corresponds to the IP address provided.
 - You can override these rules by using the -mcs and -mcd options.
- If the source or the destination is the ONTAP system, then depending on the NDMP mode (nodescoped or SVM-scoped), use an IP address that allows access to the target volume.
- source_path and destination_path are the absolute path names till the granular level of volume, qtree, directory or file.
- -mcs specifies the preferred addressing mode for the control connection to the source storage system.

inet indicates an IPv4 address mode and inet6 indicates an IPv6 address mode.

 -mcd specifies the preferred addressing mode for the control connection to the destination storage system.

inet indicates an IPv4 address mode and inet6 indicates an IPv6 address mode.

 -md specifies the preferred addressing mode for data transfers between the source and the destination storage systems.

inet indicates an IPv4 address mode and inet6 indicates an IPv6 address mode.

If you do not use the -md option in the ndmpcopy command, the addressing mode for the data connection is determined as follows:

- If either of the addresses specified for the control connections is an IPv6 address, the address mode for the data connection is IPv6.
- If both the addresses specified for the control connections are IPv4 addresses, the ndmpcopy command first attempts an IPv6 address mode for the data connection.

If that fails, the command uses an IPv4 address mode.



An IPv6 address, if specified, must be enclosed within square brackets.

This sample command migrates data from a source path (source_path) to a destination path (destination path).

```
> ndmpcopy -sa admin:<ndmp_password> -da admin:<ndmp_password>
-st md5 -dt md5 192.0.2.129:/<src_svm>/<src_vol>
192.0.2.131:/<dst_svm>/<dst_vol>
```

This sample command explicitly sets the control connections and the data connection to use IPv6 address mode:

```
> ndmpcopy -sa admin:<ndmp_password> -da admin:<ndmp_password> -st
md5 -dt md5 -mcs inet6 -mcd inet6 -md
inet6 [2001:0db8:1:1:209:6bff:feae:6d67]:/<src_svm>/<src_vol>
[2001:0ec9:1:1:200:7cgg:gfdf:7e78]:/<dst_svm>/<dst_vol>
```

Learn more about the commands described in this procedure in the ONTAP command reference.

Options for the ndmpcopy command

You should understand the options available for the ndmpcopy nodeshell command to successfully transfer data.

The following table lists the available options.

Option	Description
-sa username:[password]	This option sets the source authentication user name and password for connecting to the source storage system. This is a mandatory option.For a user without admin privilege, you must specify the user's system-generated NDMP-specific password. The system-generated password is mandatory for both admin and non-admin users.

Option	Description		
-da username:[password]	This option sets the destination authentication user name and password for connecting to the destination storage system. This is a mandatory option.		
-st{md5 text}	This option sets the source authentication type to be used when connecting to the source storage system.This is a mandatory option and therefore the user should provide either the text or md5 option.		
-dt{md5 text}	This option sets the destination authentication type to be used when connecting to the destination storage system.		
-1	This option sets the dump level used for the transfer to the specified value of level.Valid values are 0, 1, to 9, where 0 indicates a full transfer and 1 to 9 specifies an incremental transfer. The default is 0.		
-d	This option enables generation of ndmpcopy debug log messages. The ndmpcopy debug log files are located in the /mroot/etc/log root volume. The ndmpcopy debug log file names are in the ndmpcopy.yyyymmdd format.		
-f	This option enables the forced mode. This mode enables system files to be overwritten in the /etc directory on the root of the 7-Mode volume.		
-h	This option prints the help message.		
-p	This option prompts you to enter the password for source and destination authorization. This password overrides the password specified for -sa and -da options.(i)You can use this option only when the command is running in an interactive console		
-exclude	This option excludes specified files or directories from the path specified for data transfer. The value can be a comma-separated list of directory or file names such as .pst or .txt. The maximum number of exclude patterns supported is 32 and the maximum number of characters supported is 255.		

NDMP for FlexVol volumes

About NDMP for FlexVol volumes

The Network Data Management Protocol (NDMP) is a standardized protocol for controlling backup, recovery, and other types of data transfer between primary and secondary storage devices, such as storage systems and tape libraries.

By enabling NDMP support on a storage system, you enable that storage system to communicate with NDMPenabled network-attached backup applications (also called *Data Management Applications* or *DMAs*), data servers, and tape servers participating in backup or recovery operations. All network communications occur over TCPIP or TCP/IPv6 network. NDMP also provides low-level control of tape drives and medium changers.

You can perform tape backup and restore operations in either node-scoped NDMP mode or storage virtual machine (SVM) scoped NDMP mode.

You must be aware of the considerations that you have to take into account while using NDMP, list of environment variables, and supported NDMP tape backup topologies. You can also enable or disable the enhanced DAR functionality. The two authentication methods supported by ONTAP for authenticating NDMP access to a storage system are: plaintext and challenge.

Related information

Environment variables supported by ONTAP

About NDMP modes of operation

About NDMP modes of operation

You can choose to perform tape backup and restore operations either at the node level or at the storage virtual machine (SVM) level. To perform these operations successfully at the SVM level, NDMP service must be enabled on the SVM.

If you upgrade from Data ONTAP 8.2 to Data ONTAP 8.3, the NDMP mode of operation used in 8.2 will continue to be retained post the upgrade from 8.2 to 8.3.

If you install a new cluster with Data ONTAP 8.2 or later, NDMP is in the SVM-scoped NDMP mode by default. To perform tape backup and restore operations in the node-scoped NDMP mode, you must explicitly enable the node-scoped NDMP mode.

Related information

Commands for managing node-scoped NDMP mode

Managing node-scoped NDMP mode for FlexVol volumes

Managing SVM-scoped NDMP mode for FlexVol volumes

What node-scoped NDMP mode is

In the node-scoped NDMP mode, you can perform tape backup and restore operations at the node level. The NDMP mode of operation used in Data ONTAP 8.2 will continue to be retained post the upgrade from 8.2 to 8.3.

In the node-scoped NDMP mode, you can perform tape backup and restore operations on a node that owns

the volume. To perform these operations, you must establish NDMP control connections on a LIF hosted on the node that owns the volume or tape devices.



This mode is deprecated and will be removed in a future major release.

Related information

Managing node-scoped NDMP mode for FlexVol volumes

What SVM-scoped NDMP mode is

You can perform tape backup and restore operations at the storage virtual machine (SVM) level successfully if the NDMP service is enabled on the SVM. You can back up and restore all volumes hosted across different nodes in the SVM of a cluster if the backup application supports the CAB extension.

An NDMP control connection can be established on different LIF types. In the SVM-scoped NDMP mode, these LIFs belong to either the data SVM or admin SVM. The connection can be established on a LIF only if the NDMP service is enabled on the SVM that owns this LIF.

A data LIF belongs to the data SVM and the intercluster LIF, node-management LIF, and cluster-management LIF belong to the admin SVM.

In the SVM-scoped NDMP mode, the availability of volumes and tape devices for backup and restore operations depends on the LIF type on which the NDMP control connection is established and the status of the CAB extension. If your backup application supports the CAB extension and a volume and the tape device share the same affinity, then the backup application can perform a local backup or restore operation, instead of a three-way backup or restore operation.

Related information

Managing SVM-scoped NDMP mode for FlexVol volumes

Considerations when using NDMP

You must take into account a number of considerations when starting the NDMP service on your storage system.

- Each node supports a maximum of 16 concurrent backups, restores, or combination of the two using connected tape drives.
- NDMP services can generate file history data at the request of NDMP backup applications.

File history is used by backup applications to enable optimized recovery of selected subsets of data from a backup image. File history generation and processing might be time-consuming and CPU-intensive for both the storage system and the backup application.



SMTape does not support file history.

If your data protection is configured for disaster recovery—where the entire backup image will be recovered—you can disable file history generation to reduce backup time. See your backup application documentation to determine whether it is possible to disable NDMP file history generation.

• Firewall policy for NDMP is enabled by default on all LIF types.

• In node-scoped NDMP mode, backing up a FlexVol volume requires that you use the backup application to initiate a backup on a node that owns the volume.

However, you cannot back up a node root volume.

• You can perform NDMP backup from any LIF as permitted by the firewall policies.

If you use a data LIF, you must select a LIF that is not configured for failover. If a data LIF fails over during an NDMP operation, the NDMP operation fails and must be run again.

- In node-scoped NDMP mode and storage virtual machine (SVM) scoped NDMP mode with no CAB extension support, the NDMP data connection uses the same LIF as the NDMP control connection.
- During LIF migration, ongoing backup and restore operations are disrupted.

You must initiate the backup and restore operations after the LIF migration.

• The NDMP backup path is of the format /vserver_name/volume_name/path_name.

path name is optional, and specifies the path of the directory, file, or snapshot.

• When a SnapMirror destination is backed up to tape by using the dump engine, only the data in the volume is backed up.

However, if a SnapMirror destination is backed up to tape using SMTape, then the metadata is also backed up. The SnapMirror relationships and the associated metadata are not backed up to tape. Therefore, during restore, only the data on that volume is restored, but the associated SnapMirror relationships are not restored.

Related information

What Cluster Aware Backup extension does

System administration

Environment variable

Environment variables overview

Environment variables are used to communicate information about a backup or restore operation between an NDMP-enabled backup application and a storage system.

For example, if a user specifies that a backup application should back up /vserver1/vol1/dir1, the backup application sets the FILESYSTEM environment variable to /vserver1/vol1/dir1. Similarly, if a user specifies that a backup should be a level 1 backup, the backup application sets the LEVEL environment variable to 1 (one).



The setting and examining of environment variables are typically transparent to backup administrators; that is, the backup application sets them automatically.

A backup administrator rarely specifies environment variables; however, you might want to change the value of an environment variable from that set by the backup application to characterize or work around a functional or performance problem. For example, an administrator might want to temporarily disable file history generation to determine if the backup application's processing of file history information is contributing to performance issues or functional problems. Many backup applications provide a means to override or modify environment variables or to specify additional environment variables. For information, see your backup application documentation.

Environment variables supported by ONTAP

Environment variables are used to communicate information about a backup or restore operation between an NDMP-enabled backup application and a storage system. ONTAP supports environment variables, which have an associated default value. However, you can manually modify these default values.

If you manually modify the values set by the backup application, the application might behave unpredictably. This is because the backup or restore operations might not be doing what the backup application expected them to do. But in some cases, judicious modification might help in identifying or working around problems.

The following tables list the environment variables whose behavior is common to dump and SMTape and those variables that are supported only for dump and SMTape. These tables also contain descriptions of how the environment variables that are supported by ONTAP work if they are used:



In most cases, variables that have the value, ${\tt Y}$ also accept ${\tt T}$ and ${\tt N}$ also accept ${\tt F}.$

Environment variables supported for dump and SMTape

Environment variable	Valid values	Default	Description
DEBUG	Y or N	Ν	Specifies that debugging information is printed.
FILESYSTEM	string	none	Specifies the path name of the root of the data that is being backed up.
NDMP_VERSION	return_only	none	You should not modify the NDMP_VERSION variable. Created by the backup operation, the NDMP_VERSION variable returns the NDMP version. ONTAP sets the NDMP_VERSION variable during a backup for internal use and to pass to a backup application for informational purposes. The NDMP version of an NDMP session is not set with this variable.

Environment variable	Valid values	Default	Description
PATHNAME_SEPARATO R	return_value	none	Specifies the path name separator character. This character depends on the file system being backed up. For ONTAP, the character "/" is assigned to this variable. The NDMP server sets this variable before starting a tape backup operation.
TYPE	dump or smtape	dump	Specifies the type of backup supported to perform tape backup and restore operations.
VERBOSE	Y or N	Ν	Increases the log messages while performing a tape backup or restore operation.

Environment variables supported for dump

Environment variable	Valid values	Default	Description
ACL_START	return_only	none	Created by the backup operation, the ACL_START variable is an offset value used by a direct access restore or restartable NDMP backup operation. The offset value is the byte offset in the dump file where the ACL data (Pass V) begins and is returned at the end of a backup. For a direct access restore operation to correctly restore backed- up data, the ACL_START value must be passed to the restore operation when it begins. An NDMP restartable backup operation uses the ACL_START value to communicate to the backup application where the nonrestartable portion of the backup stream begins.
BASE_DATE	0, -1, or DUMP_DATE value	-1	Specifies the start date for incremental backups. When set to -1, the BASE_DATE incremental specifier is disabled. When set to 0 on a level 0 backup, incremental backups are enabled. After the initial backup, the value of the DUMP_DATE variable from the previous incremental backup is assigned to the BASE_DATE variable. These variables are an alternative to the LEVEL/UPDATE based incremental backups.

Environment variable	Valid values	Default	Description
DIRECT	Y OR N	Ν	Specifies that a restore should fast-forward directly to the location on the tape where the file data resides instead of scanning the entire tape.
			For direct access recovery to work, the backup application must provide positioning information. If this variable is set to Y, the backup application specifies the file or directory names and the positioning information.
DMP_NAME	string	none	Specifies the name for a multiple subtree backup. This variable is mandatory for multiple subtree backups.
DUMP_DATE	return_value	none	You do not change this variable directly. It is created by the backup if the BASE_DATE variable is set to a value other than -1.
			The DUMP_DATE variable is derived by prepending the 32-bit level value to a 32-bit time value computed by the dump software. The level is incremented from the last level value passed into the BASE_DATE variable. The resulting value is used as the BASE_DATE value on a subsequent incremental backup.

ENHANCED_DAR_ENAB LEDY or NNSpecifies whether enhanced DAR functionality is enabled. Enhanced DAR directory DAR and DAR of files with NT Streams. It provides performance improvements.Specifies whether enhanced DAR functionality supports directory DAR and DAR of files with NT Streams. It provides performance improvements.	Environment variable	Valid values	Default	Description
 Pestore is possible only if the following conditions are met: ONTAP supports enhanced DAR. File history is enabled (HIST=Y) during the backup. The ndmpd.offset_map .enable option is set to on. ENHANCED_DAR_E NABLED variable is set to Y during restore. 	ENHANCED_DAR_ENAB LED	Y OR N	N	Specifies whether enhanced DAR functionality is enabled. Enhanced DAR functionality supports directory DAR and DAR of files with NT Streams. It provides performance improvements. Enhanced DAR during restore is possible only if the following conditions are met: • ONTAP supports enhanced DAR. • File history is enabled (HIST=Y) during the backup. • The ndmpd.offset_map .enable option is set to on. • ENHANCED_DAR_E NABLED variable is set to Y during restore.

Environment variable	Valid values	Default	Description
EXCLUDE	pattern_string	none	Specifies files or directories that are excluded when backing up data. The exclude list is a comma-separated list of
			file or directory names. If the name of a file or directory matches one of the names in the list, it is excluded from the backup.
			The following rules apply while specifying names in the exclude list:
			• The exact name of the file or directory must be used.
			• The asterisk (*), a wildcard character, must be either the first or the last character of the string.
			Each string can have up to two asterisks.
			 A comma in a file or directory name must be preceded with a backslash.
			 The exclude list can contain up to 32 names.
			Files or directories specified to be excluded for backup are not excluded if you set NON_QUO TA_TREE to Y simultaneo
			usiy.

Environment variable	Valid values	Default	Description
EXTRACT	Y, N, OF E	Ν	Specifies that subtrees of a backed-up data set are to be restored.
			The backup application specifies the names of the subtrees to be extracted. If a file specified matches a directory whose contents were backed up, the directory is recursively extracted. To rename a file, directory, or qtree during restore without using DAR, you must set the EXTRACT
EXTRACT_ACL	Y OR N	Y	Specifies that ACLs from the backed up file are restored on a restore operation. The default is to restore
			ACLs when restoring data, except for DARs (DIRECT=Y).

Environment variable	Valid values	Default	Description
FORCE	Y OR N	N	Determines if the restore operation must check for volume space and inode availability on the destination volume. Setting this variable to Y causes the restore operation to skip checks for volume space and inode availability on the destination path. If enough volume space or inodes are not available on the destination volume, the restore operation recovers as much data allowed by the destination volume space and inode availability. The restore operation stops when volume space or inodes are not available.
HIST	Y OR N	Ν	Specifies that file history information is sent to the backup application.Most commercial backup applications set the HIST variable to Y. If you want to increase the speed of a backup operation, or you want to troubleshoot a problem with the file history collection, you can set this variable to N.Image: Colspan="2">You should not set the HIST variable to N.Image: Colspan="2">You should not set the HIST variable to N.Image: Colspan="2">You should not set the HIST variable to N.

Environment variable	Valid values	Default	Descripti	on
IGNORE_CTIME	Y OR N	Ν	Specifies increment only its ct changed s previous i backup. Some app as virus s software, ctime valu the inode the file or have not of result, an backup m files that h changed. IGNORE_ should be increment taking an amount of because t was modi	that a file is not tally backed up if ime value has since the ncremental blications, such canning change the ue of a file within , even though its attributes changed. As a incremental ight back up nave not The CTIME variable e specified only if tal backups are unacceptable f time or space the ctime value fied.
			î	The NDMP dump command sets IGNORE_C TIME to false by default. Setting it to true can result in the following data loss: 1. If IGNOR E_CTI ME is set to true with a volume level increm ental ndmpc opy, it results in the 449

Environment variable	Valid values	Default	Description
IGNORE_QTREES	Y or N	Ν	Specifies that the restore operation does not restore qtree information from backed-up qtrees.
LEVEL	0-31	0	Specifies the backup level. Level 0 copies the entire data set. Incremental backup levels, specified by values above 0, copy all files (new or modified) since the last incremental backup. For example, a level 1 backs up new or modified files since the level 0 backup, a level 2 backs up new or modified files since the level 1 backup, and so on.
LIST	Y or N	Ν	Lists the backed-up file names and inode numbers without actually restoring the data.
LIST_QTREES	Y or N	Ν	Lists the backed-up qtrees without actually restoring the data.
MULTI_SUBTREE_ NAMES	string	none	Specifies that the backup is a multiple subtree backup. Multiple subtrees are specified in the string, which is a newline- separated, null-terminated list of subtree names. Subtrees are specified by path names relative to their common root directory, which must be specified as the last element of the list. If you use this variable, you must also use the DMP_NAME variable.

Environment variable	Valid values	Default	Description
NDMP_UNICODE_ FH	Y OR N	N	Specifies that a Unicode name is included in addition to the NFS name of the file in the file history information. This option is not used by most backup applications and should not be set unless the backup application is designed to receive these additional file names. The HIST variable must also be set.
NO_ACLS	Y or N	Ν	Specifies that ACLs must not be copied when backing up data.

Environment variable	Valid values	Default	Description
NON_QUOTA_TREE	Y OR N	N	Specifies that files and directories in qtrees must be ignored when backing up data. When set to Y, items in qtrees in the data set specified by the FILESYSTEM variable are not backed up. This variable has an effect only if the FILESYSTEM variable specifies an entire volume. The NON_QUOTA_TREE variable only works on a level 0 backup and does not work if the MULTI_SUBTREE_NAM ES variable is specified.
			(i) (i) (i) (i) (i) (i) (i) (i)
NOWRITE	Y or N	Ν	Specifies that the restore operation must not write data to the disk. This variable is used for debugging.

Environment variable	Valid values	Default	Description
RECURSIVE	Y OR N	У	Specifies that directory entries during a DAR restore be expanded. The DIRECT and ENHANCED_DAR_ENAB LED environment variables must be enabled (set to Y) as well. If the RECURSIVE variable is disabled (set to N), only the permissions and ACLs for all the directories in the original source path are restored from tape, not the contents of the directories. If the RECURSIVE variable is set to N or the RECOVER_FULL_PATHS variable is set to Y, the recovery path must end with the original path.
			(i) If the RECURSIV E variable is disabled and if there is more than one recovery path, all of the recovery paths must be contained within the longest of the recovery paths. Otherwise, an error message is displayed.
			For example, the following are valid recovery paths because all of the recovery paths are within foo/dir1/deepdir/my file:

Environment variable	Valid values	Default	Description
RECOVER_FULL_PATHS	Y OF N	N	Specifies that the full recovery path will have their permissions and ACLs restored after the DAR. DIRECT and ENHANCED_DAR_ENAB LED must be enabled (set to Y) as well. If RECOVER_FULL_PATHS is set to Y, the recovery path must end with the original path. If directories already exist on the destination volume, their permissions and ACLs will not be restored from tape.
UPDATE	Y or N	Y	Updates the metadata information to enable LEVEL based incremental backup.

Environment variables supported for SMTape

Environment variable	Valid values	Default	Description
BASE_DATE	DUMP_DATE	-1	Specifies the start date for incremental backups. BASE_DATE is a string representation of the reference snapshot identifiers. Using the BASE_DATE string, SMTape locates the reference snapshot. BASE_DATE is not required for baseline backups. For an incremental backup, the value of the DUMP_DATE variable from the previous baseline or incremental backup is assigned to the BASE_DATE variable. The backup application assigns the DUMP_DATE value from a previous SMTape baseline or incremental backup.
DUMP_DATE	return_value	none	At the end of an SMTape backup, DUMP_DATE contains a string identifier that identifies the snapshot used for that backup. This snapshot could be used as the reference snapshot for a subsequent incremental backup. The resulting value of DUMP_DATE is used as the BASE_DATE value for subsequent incremental backups.

Environment variable	Valid values	Default	Description
SMTAPE_BACKUP_SET _ID	string	none	Identifies the sequence of incremental backups associated with the baseline backup. Backup set ID is a 128-bit unique ID that is generated during a baseline backup. The backup application assigns this ID as the input to the SMTAPE_BACKUP_SET_I D variable during an incremental backup.
SMTAPE_SNAPSHOT_N AME	Any valid snapshot that is available in the volume	Invalid	When the SMTAPE_SNAPSHOT_N AME variable is set to a snapshot, that snapshot and its older snapshots are backed up to tape. For incremental backup, this variable specifies incremental snapshot. The BASE_DATE variable provides the baseline snapshot.
SMTAPE_DELETE_SNA PSHOT	Y OR N	N	For a snapshot created automatically by SMTape, when the SMTAPE_DELETE_SNA PSHOT variable is set to Y, then after the backup operation is complete, SMTape deletes this snapshot. However, a snapshot created by the backup application will not be deleted.
SMTAPE_BREAK_MIRR OR	Y OR N	N	When the SMTAPE_BREAK_MIRR OR variable is set to Y, the volume of type DP is changed to a RW volume after a successful restore.

Common NDMP tape backup topologies

NDMP supports a number of topologies and configurations between backup applications and storage systems or other NDMP servers providing data (file systems) and tape services.

Storage system-to-local-tape

In the simplest configuration, a backup application backs up data from a storage system to a tape subsystem attached to the storage system. The NDMP control connection exists across the network boundary. The NDMP data connection that exists within the storage system between the data and tape services is called an NDMP local configuration.

Storage system-to-tape attached to another storage system

A backup application can also back up data from a storage system to a tape library (a medium changer with one or more tape drives) attached to another storage system. In this case, the NDMP data connection between the data and tape services is provided by a TCP or TCP/IPv6 network connection. This is called an NDMP three-way storage system-to-storage system configuration.

Storage system-to-network-attached tape library

NDMP-enabled tape libraries provide a variation of the three-way configuration. In this case, the tape library attaches directly to the TCP/IP network and communicates with the backup application and the storage system through an internal NDMP server.

Storage system-to-data server-to-tape or data server-to-storage system-to-tape

NDMP also supports storage system-to-data-server and data-server-to-storage system three-way configurations, although these variants are less widely deployed. Storage system-to-server allows storage system data to be backed up to a tape library attached to the backup application host or to another data server system. The server-to-storage system configuration allows server data to be backed up to a storage system-attached tape library.

Supported NDMP authentication methods

You can specify an authentication method to allow NDMP connection requests. ONTAP supports two methods for authenticating NDMP access to a storage system: plaintext and challenge.

In node-scoped NDMP mode, both challenge and plaintext are enabled by default. However, you cannot disable challenge. You can enable and disable plaintext. In the plaintext authentication method, the login password is transmitted as clear text.

In the storage virtual machine (SVM)-scoped NDMP mode, by default the authentication method is challenge. Unlike the node-scoped NDMP mode, in this mode you can enable and disable both plaintext and challenge authentication methods.

Related information

User authentication in a node-scoped NDMP mode

User authentication in the SVM-scoped NDMP mode

NDMP extensions supported by ONTAP

NDMP v4 provides a mechanism for creating NDMP v4 protocol extensions without modifying the core NDMP v4 protocol. You should be aware of the NDMP v4 extensions that are supported by ONTAP.

The following NDMP v4 extensions are supported by ONTAP:

• Cluster Aware Backup (CAB)



This extension is supported only in the SVM-scoped NDMP mode.

- Connection Address Extension (CAE) for IPv6 support
- Extension class 0x2050

This extension supports restartable backup operations and Snapshot Management Extensions.

The NDMP_SNAP_RECOVER message, which is part of the Snapshot Management Extensions, is used to initiate a recovery operation and to transfer the recovered data from a local snapshot to a local file system location. In ONTAP, this message allows the recovery of volumes and regular files only.

The NDMP_SNAP_DIR_LIST message enables you to browse through the snapshots of a volume. If a nondisruptive operation takes place while a browsing operation is in progress, the backup application must reinitiate the browsing operation.

NDMP restartable backup extension for a dump supported by ONTAP

You can use the NDMP restartable backup extension (RBE) functionality to restart a backup from a known checkpoint in the data stream before the failure.

What enhanced DAR functionality is

You can use the enhanced direct access recovery (DAR) functionality for directory DAR and DAR of files and NT streams. By default, enhanced DAR functionality is enabled.

Enabling enhanced DAR functionality might impact the backup performance because an offset map has to be created and written onto tape. You can enable or disable enhanced DAR in both the node-scoped and storage virtual machine (SVM)-scoped NDMP modes.

Scalability limits for NDMP sessions in ONTAP

You must be aware of the maximum number of NDMP sessions that can be established simultaneously on storage systems of different system memory capacities. This maximum number depends on the system memory of a storage system.

The limits mentioned in the following table are for the NDMP server. The limits mentioned in the section "Scalability limits for dump backup and restore sessions" are for the dump and restore session.

Scalability limits for dump backup and restore sessions

System memory of a storage system	Maximum number of NDMP sessions
Less than 16 GB	8
Greater than or equal to 16 GB but less than 24 GB	20
Greater than or equal to 24 GB	36

You can obtain the system memory of your storage system by using the sysconfig -a command (available through the nodeshell).

Learn more about sysconfig -a in the ONTAP command reference.

About NDMP for FlexGroup volumes

Beginning with ONTAP 9.7, NDMP is supported on FlexGroup volumes.

Beginning with ONTAP 9.7, the ndmpcopy command is supported for data transfer between FlexVol and FlexGroup volumes.

If you revert from ONTAP 9.7 to an earlier version, the incremental transfer information of the previous transfers is not retained and therefore, you must perform a baseline copy after reverting.

Beginning with ONTAP 9.8, the following NDMP features are supported on FlexGroup volumes:

- The NDMP_SNAP_RECOVER message in the extension class 0x2050 can be used for recovering individual files in a FlexGroup volume.
- NDMP restartable backup extension (RBE) is supported for FlexGroup volumes.
- Environment variables EXCLUDE and MULTI_SUBTREE_NAMES are supported for FlexGroup volumes.

About NDMP with SnapLock volumes

Creating multiple copies of regulated data provides you with redundant recovery scenarios, and by using NDMP dump and restore, it's possible to preserve the write once, read many (WORM) characteristics of source files on a SnapLock volume.

WORM attributes on the files in a SnapLock volume are preserved when backing up, restoring and copying data; however, WORM attributes are enforced only when restoring to a SnapLock volume. If a backup from a SnapLock volume is restored to a volume other than a SnapLock volume, the WORM attributes are preserved but are ignored and are not enforced by ONTAP.

Manage node-scoped NDMP mode for FlexVol volumes

Manage node-scoped NDMP mode for FlexVol volumes overview in ONTAP

You can manage NDMP at the node level by using NDMP options and commands. You can modify the NDMP options by using the options command. You must use NDMP-specific credentials to access a storage system to perform tape backup and restore operations.

Learn more about options in the ONTAP command reference.

Related information

Commands for managing node-scoped NDMP mode

What node-scoped NDMP mode is

Commands for managing node-scoped NDMP mode in ONTAP

You can use the system services ndmp commands to manage NDMP at a node level. Some of these commands are deprecated and will be removed in a future major release.

You can use the following NDMP commands only at the advanced privilege level:

- system services ndmp service terminate
- system services ndmp service start
- system services ndmp service stop
- system services ndmp log start
- system services ndmp log stop

If you want to	Use this command
Enable NDMP service	system services ndmp on*
Disable NDMP service	system services ndmp off*
Display NDMP configuration	system services ndmp show*
Modify NDMP configuration	system services ndmp modify*
Display the default NDMP version	system services ndmp version*
Display NDMP service configuration	system services ndmp service show
Modify NDMP service configuration	system services ndmp service modify
Display all NDMP sessions	system services ndmp status
Display detailed information about all NDMP sessions	system services ndmp probe
Terminate the specified NDMP session	system services ndmp kill
Terminate all NDMP sessions	system services ndmp kill-all

If you want to	Use this command
Change the NDMP password	system services ndmp password*
Enable node-scoped NDMP mode	system services ndmp node-scope-mode on*
Disable node-scoped NDMP mode	system services ndmp node-scope-mode off*
Display the node-scoped NDMP mode status	system services ndmp node-scope-mode status*
Forcefully terminate all NDMP sessions	system services ndmp service terminate
Start the NDMP service daemon	system services ndmp service start
Stop the NDMP service daemon	system services ndmp service stop
Start logging for the specified NDMP session	system services ndmp log start*
Stop logging for the specified NDMP session	system services ndmp log stop*

• These commands are deprecated and will be removed in a future major release.

Learn more about system services ndmp in the ONTAP command reference.

User authentication in a node-scoped NDMP mode

In the node-scoped NDMP mode, you must use NDMP specific credentials to access a storage system in order to perform tape backup and restore operations.

The default user ID is "root". Before using NDMP on a node, you must ensure that you change the default NDMP password associated with the NDMP user. You can also change the default NDMP user ID.

Related information

Commands for managing node-scoped NDMP mode

Manage SVM-scoped NDMP mode for FlexVol volumes

Manage SVM-scoped NDMP mode for FlexVol volumes overview in ONTAP

You can manage NDMP on a per SVM basis by using the NDMP options and commands. You can modify the NDMP options by using the vserver services ndmp modify command. In the SVM-scoped NDMP mode, user authentication is integrated with the role-based access control mechanism.

You can add NDMP in the allowed or disallowed protocols list by using the vserver modify command. By

default, NDMP is in the allowed protocols list. If NDMP is added to the disallowed protocols list, NDMP sessions cannot be established.

You can control the LIF type on which an NDMP data connection is established by using the -preferred -interface-role option. During an NDMP data connection establishment, NDMP chooses an IP address that belongs to the LIF type as specified by this option. If the IP addresses do not belong to any of these LIF types, then the NDMP data connection cannot be established.

Learn more about vserver services ndmp modify in the ONTAP command reference.

Related information

Commands for managing SVM-scoped NDMP mode

What Cluster Aware Backup extension does

What SVM-scoped NDMP mode is

System administration

Commands for managing SVM-scoped NDMP mode in ONTAP

You can use the vserver services ndmp commands to manage NDMP on each storage virtual machine (SVM, formerly known as Vserver).

If you want to	Use this command		
Enable NDMP service	vserver services ndmp on		
	i	NDMP service must always be enabled on all nodes in a cluster. You can enable NDMP service on a node by using the system services ndmp on command. By default, NDMP service is always enabled on a node.	
Disable NDMP service	vserver	services ndmp off	
Display NDMP configuration	vserver	services ndmp show	
Modify NDMP configuration	vserver	services ndmp modify	
Display default NDMP version	vserver	services ndmp version	
Display all NDMP sessions	vserver	services ndmp status	
Display detailed information about all NDMP sessions	vserver	services ndmp probe	
Terminate a specified NDMP session	vserver	services ndmp kill	

If you want to	Use this command
Terminate all NDMP sessions	vserver services ndmp kill-all
Generate the NDMP password	vserver services ndmp generate-password
Display NDMP extension status	vserver services ndmp extensions show This command is available at the advanced privilege level.
Modify (enable or disable) NDMP extension status	vserver services ndmp extensions modify This command is available at the advanced privilege level.
Start logging for the specified NDMP session	vserver services ndmp log start This command is available at the advanced privilege level.
Stop logging for the specified NDMP session	vserver services ndmp log stop This command is available at the advanced privilege level.

Learn more about vserver services ndmp in the ONTAP command reference.

What Cluster Aware Backup extension does

CAB (Cluster Aware Backup) is an NDMP v4 protocol extension. This extension enables the NDMP server to establish a data connection on a node that owns a volume. This also enables the backup application to determine if volumes and tape devices are located on the same node in a cluster.

To enable the NDMP server to identify the node that owns a volume and to establish a data connection on such a node, the backup application must support the CAB extension. CAB extension requires the backup application to inform the NDMP server about the volume to be backed up or restored prior to establishing the data connection. This allows the NDMP server to determine the node that hosts the volume and appropriately establish the data connection.

With the CAB extension supported by the backup application, the NDMP server provides affinity information about volumes and tape devices. Using this affinity information, the backup application can perform a local backup instead of a three-way backup if a volume and tape device are located on the same node in a cluster.

Availability of volumes and tape devices for backup and restore on different LIF types

You can configure a backup application to establish an NDMP control connection on any of the LIF types in a cluster. In the storage virtual machine (SVM)-scoped NDMP mode,

you can determine the availability of volumes and tape devices for backup and restore operations depending upon these LIF types and the status of the CAB extension.

The following tables show the availability of volumes and tape devices for NDMP control connection LIF types and the status of the CAB extension:

NDMP control connection LIF type	Volumes available for backup or restore	Tape devices available for backup or restore
Node-management LIF	All volumes hosted by a node	Tape devices connected to the node hosting the node- management LIF
Data LIF	Only volumes that belong to the SVM hosted by a node that hosts the data LIF	None
Cluster-management LIF	All volumes hosted by a node that hosts the cluster-management LIF	None
Intercluster LIF	All volumes hosted by a node that hosts the intercluster LIF	Tape devices connected to the node hosting the intercluster LIF

Availability of volumes and tape devices when CAB extension is not supported by the backup application

Availability of volumes and tape devices when CAB extension is supported by the backup application

NDMP control connection LIF type	Volumes available for backup or restore	Tape devices available for backup or restore
Node-management LIF	All volumes hosted by a node	Tape devices connected to the node hosting the node- management LIF
Data LIF	All volumes that belong to the SVM that hosts the data LIF	None
Cluster-management LIF	All volumes in the cluster	All tape devices in the cluster
Intercluster LIF	All volumes in the cluster	All tape devices in the cluster

What affinity information is

With the backup application being CAB aware, the NDMP server provides unique location information about volumes and tape devices. Using this affinity information, the backup application can perform a local backup instead of a three-way backup if a volume and a tape device share the same affinity.

If the NDMP control connection is established on a node management LIF, cluster management LIF, or an
intercluster LIF, the backup application can use the affinity information to determine if a volume and tape device are located on the same node and then perform either a local or a three-way backup or restore operation. If the NDMP control connection is established on a data LIF, then the backup application always performs a three-way backup.

Local NDMP backup and Three-way NDMP backup



Using the affinity information about volumes and tape devices, the DMA (backup application) performs a local NDMP backup on the volume and tape device located on Node 1 in the cluster. If the volume moves from Node 1 to Node 2, affinity information about the volume and tape device changes. Hence, for a subsequent backup the DMA performs a three-way NDMP backup operation. This ensures continuity of the backup policy for the volume irrespective of the node to which the volume is moved to.

Related information

What Cluster Aware Backup extension does

NDMP server supports secure control connections in SVM-scoped mode

A secure control connection can be established between the Data Management Application (DMA) and NDMP server by using secure sockets (SSL/TLS) as the communication mechanism. This SSL communication is based on the server certificates. The NDMP server listens on port 30000 (assigned by IANA for "ndmps" service).

After establishing the connection from the client on this port, the standard SSL handshake ensues where the server presents the certificate to the client. When the client accepts the certificate, the SSL handshake is complete. After this process is complete, all of the communication between the client and the server is encrypted. The NDMP protocol workflow remains exactly as before. The secure NDMP connection requires server- side certificate authentication only. A DMA can choose to establish a connection either by connecting to the secure NDMP service or the standard NDMP service.

By default, secure NDMP service is disabled for a storage virtual machine (SVM). You can enable or disable the secure NDMP service on a given SVM by using the vserver services ndmp modify -vserver vserver -is-secure-control-connection-enabled [true|false] command.

NDMP data connection types

In the storage virtual machine (SVM)-scoped NDMP mode, the supported NDMP data connection types depend on the NDMP control connection LIF type and the status of the CAB extension. This NDMP data connection type indicates whether you can perform a

local or a three-way NDMP backup or restore operation.

You can perform a three-way NDMP backup or restore operation over a TCP or TCP/IPv6 network. The following tables show the NDMP data connection types based on the NDMP control connection LIF type and the status of the CAB extension.

NDMP data connection type when CAB extension is supported by the backup application

NDMP control connection LIF type	NDMP data connection type
Node-management LIF	LOCAL, TCP, TCP/IPv6
Data LIF	TCP, TCP/IPv6
Cluster-management LIF	LOCAL, TCP, TCP/IPv6
Intercluster LIF	LOCAL, TCP, TCP/IPv6

NDMP data connection type when CAB extension is not supported by the backup application

NDMP control connection LIF type	NDMP data connection type
Node-management LIF	LOCAL, TCP, TCP/IPv6
Data LIF	TCP, TCP/IPv6
Cluster-management LIF	TCP, TCP/IPv6
Intercluster LIF	LOCAL, TCP, TCP/IPv6

Related information

What Cluster Aware Backup extension does

Network management

User authentication in the SVM-scoped NDMP mode

In the storage virtual machine (SVM)-scoped NDMP mode, NDMP user authentication is integrated with role-based access control. In the SVM context, the NDMP user must have either the "vsadmin" or "vsadmin-backup" role. In a cluster context, the NDMP user must have either the "admin" or "backup" role.

Apart from these pre-defined roles, a user account associated with a custom role can also be used for NDMP authentication provided that the custom role has the "vserver services ndmp" folder in its command directory and the access level of the folder is not "none". In this mode, you must generate an NDMP password for a given user account, which is created through role-based access control. Cluster users in an admin or backup role can access a node-management LIF, a cluster-management LIF, or an intercluster LIF. Users in a vsadmin-backup or vsadmin role can access only the data LIF for that SVM. Therefore, depending on the role of a user, the availability of volumes and tape devices for backup and restore operations vary.

This mode also supports user authentication for NIS and LDAP users. Therefore, NIS and LDAP users can access multiple SVMs with a common user ID and password. However, NDMP authentication does not support Active Directory users.

In this mode, a user account must be associated with the SSH application and the "User password" authentication method.

Related information

Commands for managing SVM-scoped NDMP mode

System administration

Generate an NDMP-specific password for NDMP users

In the storage virtual machine (SVM)-scoped NDMP mode, you must generate a password for a specific user ID. The generated password is based on the actual login password for the NDMP user. If the actual login password changes, you must generate the NDMP-specific password again.

Steps

1. Use the vserver services ndmp generate-password command to generate an NDMP-specific password.

You can use this password in any current or future NDMP operation that requires password input.



From the storage virtual machine (SVM, formerly known as Vserver) context, you can generate NDMP passwords for users belonging only to that SVM.

The following example shows how to generate an NDMP-specific password for a user ID user1:

```
cluster1::vserver services ndmp> generate-password -vserver vs1 -user
user1
Vserver: vs1
User: user1
Password: jWZiNt57huPOoD8d
```

2. If you change the password to your regular storage system account, repeat this procedure to obtain your new NDMP-specific password.

How tape backup and restore operations are affected during disaster recovery in MetroCluster configuration

You can perform tape backup and restore operations simultaneously during disaster recovery in a MetroCluster configuration. You must understand how these operations are affected during disaster recovery.

If tape backup and restore operations are performed on a volume of anSVM in a disaster recovery relationship, then you can continue performing incremental tape backup and restore operations after a switchover and switchback.

About dump engine for FlexVol volumes

About dump engine for FlexVol volumes

Dump is a snapshot based backup and recovery solution from ONTAP that helps you to back up files and directories from a snapshot to a tape device and restore the backed up data to a storage system.

You can back up your file system data, such as directories, files, and their associated security settings, to a tape device by using the dump backup. You can back up an entire volume, an entire qtree, or a subtree that is neither an entire volume nor an entire qtree.

You can perform a dump backup or restore by using NDMP-compliant backup applications.

When you perform a dump backup, you can specify the snapshot to be used for a backup. If you do not specify a snapshot for the backup, the dump engine creates a snapshot for the backup. After the backup operation is completed, the dump engine deletes this snapshot.

You can perform level-0, incremental, or differential backups to tape by using the dump engine.



After reverting to a release earlier than Data ONTAP 8.3, you must perform a baseline backup operation before performing an incremental backup operation.

Related information

Upgrade, revert, or downgrade

How a dump backup works

A dump backup writes file system data from disk to tape using a predefined process. You can back up a volume, a qtree, or a subtree that is neither an entire volume nor an entire qtree.

The following table describes the process that ONTAP uses to back up the object indicated by the dump path:

Stage	Action
1	For less than full volume or full qtree backups, ONTAP traverses directories to identify the files to be backed up. If you are backing up an entire volume or qtree, ONTAP combines this stage with Stage 2.
2	For a full volume or full qtree backup, ONTAP identifies the directories in the volumes or qtrees to be backed up.
3	ONTAP writes the directories to tape.
4	ONTAP writes the files to tape.
5	ONTAP writes the ACL information (if applicable) to tape.

The dump backup uses a snapshot of your data for the backup. Therefore, you do not have to take the volume

offline before initiating the backup.

The dump backup names each snapshot it creates as snapshot_for_backup.n, where n is an integer starting at 0. Each time the dump backup creates a snapshot, it increments the integer by 1. The integer is reset to 0 after the storage system is rebooted. After the backup operation is completed, the dump engine deletes this snapshot.

When ONTAP performs multiple dump backups simultaneously, the dump engine creates multiple snapshots. For example, if ONTAP is running two dump backups simultaneously, you find the following snapshots in the volumes from which data is being backed up: snapshot for backup.0 and snapshot for backup.1.



When you are backing up from a snapshot, the dump engine does not create an additional snapshot.

Types of data that the dump engine backs up

The dump engine enables you to back up data to tape to guard against disasters or controller disruptions. In addition to backing up data objects such as a files, directories, qtrees, or entire volumes, the dump engine can back up many types of information about each file. Knowing the types of data that the dump engine can back up and the restrictions to take into consideration can help you plan your approach to disaster recovery.

In addition to backing up data in files, the dump engine can back up the following information about each file, as applicable:

- UNIX GID, owner UID, and file permissions
- · UNIX access, creation, and modification time
- · File type
- File size
- DOS name, DOS attributes, and creation time
- Access control lists (ACLs) with 1,024 access control entries (ACEs)
- Qtree information
- Junction paths

Junction paths are backed up as symbolic links.

• LUN and LUN clones

You can back up an entire LUN object; however, you cannot back up a single file within the LUN object. Similarly, you can restore an entire LUN object but not a single file within the LUN.



The dump engine backs up LUN clones as independent LUNs.

VM-aligned files

Backup of VM-aligned files is not supported in releases earlier than Data ONTAP 8.1.2.



When a snapshot-backed LUN clone is transitioned from Data ONTAP operating in 7-Mode to ONTAP, it becomes an inconsistent LUN. The dump engine does not back up inconsistent LUNs.

When you restore data to a volume, client I/O is restricted on the LUNs being restored. The LUN restriction is removed only when the dump restore operation is complete. Similarly, during a SnapMirror single file or LUN restore operation, client I/O is restricted on both files and LUNs being restored. This restriction is removed only when the single file or LUN restore operation is complete. If a dump backup is performed on a volume on which a dump restore or SnapMirror single file or LUN restore operation is being performed, then the files or LUNs that have client I/O restriction are not included in the backup. These files or LUNs are included in a subsequent backup operation if the client I/O restriction is removed.



A LUN running on Data ONTAP 8.3 that is backed up to tape can be restored only to 8.3 and later releases and not to an earlier release. If the LUN is restored to an earlier release, then the LUN is restored as a file.

When you back up a SnapVault secondary volume or a volume SnapMirror destination to tape, only the data on the volume is backed up. The associated metadata is not backed up. Therefore, when you try to restore the volume, only the data on that volume is restored. Information about the volume SnapMirror relationships is not available in the backup and therefore is not restored.

If you dump a file that has only Windows NT permissions and restore it to a UNIX-style qtree or volume, the file gets the default UNIX permissions for that qtree or volume.

If you dump a file that has only UNIX permissions and restore it to an NTFS-style qtree or volume, the file gets the default Windows permissions for that qtree or volume.

Other dumps and restores preserve permissions.

You can back up VM-aligned files and the vm-align-sector option. For more information about VM-aligned files, see Logical storage management.

What increment chains are

An increment chain is a series of incremental backups of the same path. Because you can specify any level of backup at any time, you must understand increment chains to be able to perform backups and restores effectively. You can perform 31 levels of incremental backup operations.

There are two types of increment chains:

- A consecutive increment chain, which is a sequence of incremental backups that starts with level 0 and is raised by 1 at each subsequent backup.
- A nonconsecutive increment chain, where incremental backups skip levels or have levels that are out of sequence, such as 0, 2, 3, 1, 4, or more commonly 0, 1, 1, 1 or 0, 1, 2, 1, 2.

Incremental backups are based on the most recent lower-level backup. For example, the sequence of backup levels 0, 2, 3, 1, 4 provides two increment chains: 0, 2, 3 and 0, 1, 4. The following table explains the bases of the incremental backups:

Backup order	Increment level	Increment chain	Base	Files backed up
1	0	Both	Files on the storage system	All files in the backup path
2	2	0, 2, 3	Level-0 backup	Files in the backup path created since the level-0 backup
3	3	0, 2, 3	Level-2 backup	Files in the backup path created since the level-2 backup
4	1	0, 1, 4	Level-0 backup, because this is the most recent level that is lower than the level-1 backup	Files in the backup path created since the level-0 backup, including files that are in the level-2 and level-3 backups
5	4	0, 1, 4	The level-1 backup, because it is a lower level and is more recent than the level-0, level-2, or level-3 backups	Files created since the level-1 backup

What the blocking factor is

A tape block is 1,024 bytes of data. During a tape backup or restore, you can specify the number of tape blocks that are transferred in each read/write operation. This number is called the *blocking factor*.

You can use a blocking factor from 4 to 256. If you plan to restore a backup to a system other than the system that did the backup, the restore system must support the blocking factor that you used for the backup. For example, if you use a blocking factor of 128, the system on which you restore that backup must support a blocking factor of 128.

During an NDMP backup, the MOVER_RECORD_SIZE determines the blocking factor. ONTAP allows a maximum value of 256 KB for MOVER_RECORD_SIZE.

When to restart a dump backup

A dump backup sometimes does not finish because of internal or external errors, such as tape write errors, power outages, accidental user interruptions, or internal inconsistency on the storage system. If your backup fails for one of these reasons, you can restart it.

You can choose to interrupt and restart a backup to avoid periods of heavy traffic on the storage system or to avoid competition for other limited resources on the storage system, such as a tape drive. You can interrupt a long backup and restart it later if a more urgent restore (or backup) requires the same tape drive. Restartable backups persist across reboots. You can restart an aborted backup to tape only if the following conditions are

true:

- The aborted backup is in phase IV.
- All of the associated snapshots that were locked by the dump command are available.
- The file history must be enabled.

When such a dump operation is aborted and left in a restartable state, the associated snapshots are locked. These snapshots are released after the backup context is deleted. You can view the list of backup contexts by using the vserver services ndmp restartable backup show command.

```
cluster::> vserver services ndmpd restartable-backup show
Vserver
           Context Identifier
                                              Is Cleanup Pending?
_____
vserver1 330e6739-0179-11e6-a299-005056bb4bc9 false
vserver1 481025c1-0179-11e6-a299-005056bb4bc9 false
vserver2 5cf10132-0179-11e6-a299-005056bb4bc9 false
3 entries were displayed.
cluster::> vserver services ndmpd restartable-backup show -vserver
vserver1 -context-id 330e6739-0179-11e6-a299-005056bb4bc9
                      Vserver: vserver1
           Context Identifier: 330e6739-0179-11e6-a299-005056bb4bc9
                  Volume Name: /vserver1/vol1
          Is Cleanup Pending?: false
           Backup Engine Type: dump
Is Snapshot Copy Auto-created?: true
                    Dump Path: /vol/vol1
   Incremental Backup Level ID: 0
                    Dump Name: /vserver1/vol1
    Context Last Updated Time: 1460624875
              Has Offset Map?: true
                Offset Verify: true
      Is Context Restartable?: true
             Is Context Busy?: false
                 Restart Pass: 4
             Status of Backup: 2
           Snapshot Copy Name: snapshot for backup.1
         State of the Context: 7
cluster::>"
```

How a dump restore works

A dump restore writes file system data from tape to disk using a predefined process.

The process in the following table shows how the dump restore works:

Stage	Action
1	ONTAP catalogs the files that need to be extracted from the tape.
2	ONTAP creates directories and empty files.
3	ONTAP reads a file from tape, writes it to disk, and sets the permissions (including ACLs) on it.
4	ONTAP repeats stages 2 and 3 until all the specified files are copied from the tape.

Types of data that the dump engine restores

When a disaster or controller disruption occurs, the dump engine provides multiple methods for you to recover all of the data that you backed up, from single files, to file attributes, to entire directories. Knowing the types of data that dump engine can restore and when to use which method of recovery can help minimize downtime.

You can restore data to an online mapped LUN. However, host applications cannot access this LUN until the restore operation is complete. After the restore operation is complete, the host cache of the LUN data should be flushed to provide coherency with the restored data.

The dump engine can recover the following data:

- · Contents of files and directories
- UNIX file permissions
- ACLs

If you restore a file that has only UNIX file permissions to an NTFS qtree or volume, the file has no Windows NT ACLs. The storage system uses only the UNIX file permissions on this file until you create a Windows NT ACL on it.



If you restore ACLs backed up from storage systems running Data ONTAP 8.2 to storage systems running Data ONTAP 8.1.x and earlier that have an ACE limit lower than 1,024, a default ACL is restored.

Qtree information

Qtree information is used only if a qtree is restored to the root of a volume. Qtree information is not used if a qtree is restored to a lower directory, such as /vs1/vol1/subdir/lowerdir, and it ceases to be a qtree.

- All other file and directory attributes
- Windows NT streams
- LUNs

• A LUN must be restored to a volume level or a qtree level for it to remain as a LUN.

If it is restored to a directory, it is restored as a file because it does not contain any valid metadata.

- $\,\circ\,$ A 7-Mode LUN is restored as a LUN on an ONTAP volume.
- A 7-Mode volume can be restored to an ONTAP volume.
- VM-aligned files restored to a destination volume inherit the VM-align properties of the destination volume.
- The destination volume for a restore operation might have files with mandatory or advisory locks.

While performing restore operation to such a destination volume, the dump engine ignores these locks.

Considerations before restoring data

You can restore backed-up data to its original path or to a different destination. If you are restoring backed-up data to a different destination, you must prepare the destination for the restore operation.

Before restoring data either to its original path or to a different destination, you must have the following information and meet the following requirements:

- The level of the restore
- The path to which you are restoring the data
- · The blocking factor used during the backup
- If you are doing an incremental restore, all tapes must be in the backup chain
- A tape drive that is available and compatible with the tape to be restored from

Before restoring data to a different destination, you must perform the following operations:

- If you are restoring a volume, you must create a new volume.
- If you are restoring a qtree or a directory, you must rename or move files that are likely to have the same names as files you are restoring.



In ONTAP 9, qtree names support the Unicode format. The earlier releases of ONTAP do not support this format. If a qtree with Unicode names in ONTAP 9 is copied to an earlier release of ONTAP using the ndmpcopy command or through restoration from a backup image in a tape, the qtree is restored as a regular directory and not as a qtree with Unicode format.



If a restored file has the same name as an existing file, the existing file is overwritten by the restored file. However, the directories are not overwritten.

To rename a file, directory, or qtree during restore without using DAR, you must set the EXTRACT environment variable to E.

Required space on the destination storage system

You require about 100 MB more space on the destination storage system than the amount of data to be restored.

The restore operation checks for volume space and inode availability on the destination volume when the restore operation starts. Setting the FORCE environment variable to Y causes the restore operation to skip the checks for volume space and inode availability on the destination path. If there is not enough volume space or inodes available on the destination volume, the restore operation recovers as much data allowed by the destination volume space and inode availability. The restore operation stops when there is no more volume space or inodes left.

Scalability limits for dump backup and restore sessions in ONTAP

You must be aware of the maximum number of dump backup and restore sessions that can be performed simultaneously on storage systems of different system memory capacities. This maximum number depends on the system memory of a storage system.

The limits mentioned in the following table are for the dump or restore engine. The limits mentioned in the scalability limits for NDMP sessions are for the NDMP server, which are higher than the engine limits.

System memory of a storage system	Total number of dump backup and restore sessions
Less than 16 GB	4
Greater than or equal to 16 GB but less than 24 GB	16
Greater than or equal to 24 GB	32



(;)

If you use ndmpcopy command to copy data within storage systems, two NDMP sessions are established, one for dump backup and the other for dump restore.

You can obtain the system memory of your storage system by using the sysconfig -a command (available through the nodeshell).

Learn more about sysconfig -a in the ONTAP command reference.

Related information

Scalability limits for NDMP sessions

Tape backup and restore support between Data ONTAP operating in 7-Mode and ONTAP

You can restore data backed up from a storage system operating in 7-Mode or running ONTAP to a storage system either operating in 7-Mode or running ONTAP.

The following tape backup and restore operations are supported between Data ONTAP operating in 7-Mode and ONTAP:

- Backing up a 7-Mode volume to a tape drive connected to a storage system running ONTAP
- Backing up an ONTAP volume to a tape drive connected to a 7-Mode system
- Restoring backed-up data of a 7-Mode volume from a tape drive connected to a storage system running ONTAP
- Restoring backed-up data of an ONTAP volume from a tape drive connected to a 7-Mode system

• Restoring a 7-Mode volume to an ONTAP volume



A 7-Mode LUN is restored as a LUN on an ONTAP volume.
You should retain the ONTAP LUN identifiers when restoring a 7-Mode LUN to an existing ONTAP LUN.

• Restoring an ONTAP volume to a 7-Mode volume



An ONTAP LUN is restored as a regular file on a 7-Mode volume.

Delete restartable contexts

If you want to start a backup instead of restarting a context, you can delete the context.

About this task

You can delete a restartable context using the vserver services ndmp restartable-backup delete command by providing the SVM name and the context ID.

Steps

1. Delete a restartable context:

vserver services ndmp restartable-backup delete -vserver vserver-name -context -id context_identifier.

```
cluster::> vserver services ndmpd restartable-backup show
Vserver
         Context Identifier
                                         Is Cleanup Pending?
vserver1 330e6739-0179-11e6-a299-005056bb4bc9 false
vserver1
         481025c1-0179-11e6-a299-005056bb4bc9 false
vserver2 5cf10132-0179-11e6-a299-005056bb4bc9 false
3 entries were displayed.
cluster::>
cluster::> vserver services ndmp restartable-backup delete -vserver
vserver1 -context-id 481025c1-0179-11e6-a299-005056bb4bc9
cluster::> vserver services ndmpd restartable-backup show
Vserver
         Context Identifier
                                         Is Cleanup Pending?
______ ____
         330e6739-0179-11e6-a299-005056bb4bc9 false
vserver1
         5cf10132-0179-11e6-a299-005056bb4bc9 false
vserver2
3 entries were displayed.
cluster::>"
```

How dump works on a SnapVault secondary volume

You can perform tape backup operations on data that is mirrored on the SnapVault secondary volume. You can back up only the data that is mirrored on the SnapVault secondary volume to tape, and not the SnapVault relationship metadata.

When you break the data protection mirror relationship (snapmirror break) or when a SnapMirror resynchronization occurs, you must always perform a baseline backup.

Related information

• snapmirror break

How dump works with storage failover and ARL operations

Before you perform dump backup or restore operations, you should understand how these operations work with storage failover (takeover and giveback) or aggregate relocation (ARL) operations. The -override-vetoes option determines the behavior of dump engine during a storage failover or ARL operation.

When a dump backup or restore operation is running and the -override-vetoes option is set to false, a user-initiated storage failover or ARL operation is stopped. However, if the -override-vetoes option is set to true, then the storage failover or ARL operation is continued and the dump backup or restore operation is aborted. When a storage failover or ARL operation is automatically initiated by the storage system, an active dump backup or restore operation is always aborted. You cannot restart dump backup and restore operations even after storage failover or ARL operations complete.

Dump operations when CAB extension is supported

If the backup application supports CAB extension, you can continue performing incremental dump backup and restore operations without reconfiguring backup policies after a storage failover or ARL operation.

Dump operations when CAB extension is not supported

If the backup application does not support CAB extension, you can continue performing incremental dump backup and restore operations if you migrate the LIF configured in the backup policy to the node that hosts the destination aggregate. Otherwise, after the storage failover and ARL operation, you must perform a baseline backup prior to performing the incremental backup operation.



For storage failover operations, the LIF configured in the backup policy must be migrated to the partner node.

Related information

High Availability

How dump works with volume move

Tape backup and restore operations and volume move can run in parallel until the final cutover phase is attempted by the storage system. After this phase, new tape backup and restore operations are not allowed on the volume that is being moved. However, the current operations continue to run until completion.

The following table describes the behavior of tape backup and restore operations after the volume move operation:

If you are performing tape backup and restore operations in the	Then
storage virtual machine (SVM) scoped NDMP mode when CAB extension is supported by the backup application	You can continue performing incremental tape backup and restore operations on read/write and read-only volumes without reconfiguring backup policies.
SVM-scoped NDMP mode when CAB extension is not supported by the backup application	You can continue performing incremental tape backup and restore operations on read/write and read-only volumes if you migrate the LIF configured in the backup policy to the node that hosts the destination aggregate. Otherwise, after the volume move, you must perform a baseline backup before performing the incremental backup operation.



When a volume move occurs, if the volume belonging to a different SVM on the destination node has the same name as that of the moved volume, then you cannot perform incremental backup operations of the moved volume.

How dump works when a FlexVol volume is full

Before performing an incremental dump backup operation, you must ensure that there is sufficient free space in the FlexVol volume.

If the operation fails, you must increase the free space in the Flex Vol volume either by increasing its size or by deleting the snapshots. Then perform the incremental backup operation again.

How dump works when volume access type changes

When a SnapMirror destination volume or a SnapVault secondary volume changes state from read/write to read-only or from read-only to read/write, you must perform a baseline tape backup or restore operation.

SnapMirror destination and SnapVault secondary volumes are read-only volumes. If you perform tape backup and restore operations on such volumes, you must perform a baseline backup or restore operation whenever the volume changes state from read-only to read/write or from read/write to read-only.

How dump works with SnapMirror single file or LUN restore

Before you perform dump backup or restore operations on a volume to which a single file or LUN is restored by using SnapMirror technology, you must understand how dump operations work with a single file or LUN restore operation.

During a SnapMirror single file or LUN restore operation, client I/O is restricted on the file or LUN being restored. When the single file or LUN restore operation finishes, the I/O restriction on the file or LUN is removed. If a dump backup is performed on a volume to which a single file or LUN is restored, then the file or LUN that has client I/O restriction is not included in the dump backup. In a subsequent backup operation, this file or LUN is backed up to tape after the I/O restriction is removed.

You cannot perform a dump restore and a SnapMirror single file or LUN restore operation simultaneously on the same volume.

How dump backup and restore operations are affected in MetroCluster configurations

Before you perform dump backup and restore operations in a MetroCluster configuration, you must understand how dump operations are affected when a switchover or switchback operation occurs.

Dump backup or restore operation followed by switchover

Consider two clusters: cluster 1 and cluster 2. During a dump backup or restore operation on cluster 1, if a switchover is initiated from cluster 1 to cluster 2, then the following occurs:

- If the value of the override-vetoes option is false, then the switchover is aborted and the backup or restore operation continues.
- If the value of the option is true, then the dump backup or restore operation is aborted and the switchover continues.

Dump backup or restore operation followed by switchback

A switchover is performed from cluster 1 to cluster 2 and a dump backup or restore operation is initiated on cluster 2. The dump operation backs up or restores a volume that is located on cluster 2. At this point, if a switchback is initiated from cluster 2 to cluster 1, then the following occurs:

- If the value of the override-vetoes option is false, then the switchback is cancelled and the backup or restore operation continues.
- If the value of the option is true, then the backup or restore operation is aborted and the switchback continues.

Dump backup or restore operation initiated during a switchover or switchback

During a switchover from cluster 1 to cluster 2, if a dump backup or restore operation is initiated on cluster 1, then the backup or restore operation fails and the switchover continues.

During a switchback from cluster 2 to cluster 1, if a dump backup or restore operation is initiated from cluster 2, then the backup or restore operation fails and the switchback continues.

About SMTape engine for FlexVol volumes

About SMTape engine for FlexVol volumes

SMTape is a disaster recovery solution from ONTAP that backs up blocks of data to tape. You can use SMTape to perform volume backups to tapes. However, you cannot perform a backup at the qtree or subtree level. SMTape supports baseline, differential, and incremental backups. SMTape does not require a license.

You can perform an SMTape backup and restore operation by using an NDMP-compliant backup application. You can choose SMTape to perform backup and restore operations only in the storage virtual machine (SVM) scoped NDMP mode.



Reversion process is not supported when an SMTape backup or restore session is in progress. You must wait until the session finishes or you must abort the NDMP session.

Using SMTape, you can back up 255 snapshots. For subsequent baseline, incremental, or differential backups, you must delete older backed-up snapshots.

Before performing a baseline restore, the volume to which data is being restored must be of type DP and this volume must be in the restricted state. After a successful restore, this volume is automatically online. You can perform subsequent incremental or differential restores on this volume in the order in which the backups were performed.

Use snapshots during SMTape backup

You should understand how snapshots are used during an SMTape baseline backup and an incremental backup. There are also considerations to keep in mind while performing a backup using SMTape.

Baseline backup

While performing a baseline backup, you can specify the name of the snapshot to be backed up to tape. If no snapshot is specified, then depending on the access type of the volume (read/write or read-only), either a snapshot is created automatically or existing snapshots are used. When you specify a snapshot for the backup, all the snapshots older than the specified snapshot are also backed up to tape.

If you do not specify a snapshot for the backup, the following occurs:

· For a read/write volume, a snapshot is created automatically.

The newly created snapshot and all the older snapshots are backed up to tape.

• For a read-only volume, all the snapshots, including the latest snapshot, are backed up to tape.

Any new snapshots created after the backup is started are not backed up.

Incremental backup

For SMTape incremental or differential backup operations, the NDMP-compliant backup applications create and manage the snapshots.

You must always specify a snapshot while performing an incremental backup operation. For a successful incremental backup operation, the snapshot backed up during the previous backup operation (baseline or incremental) must be on the volume from which the backup is performed. To ensure that you use this backed-up snapshot, you must consider the snapshot policy assigned on this volume while configuring the backup policy.

Considerations on SMTape backups on SnapMirror destinations

• A data protection mirror relationship creates temporary snapshots on the destination volume for replication.

You should not use these snapshots for SMTape backup.

• If a SnapMirror update occurs on a destination volume in a data protection mirror relationship during an SMTape backup operation on the same volume, then the snapshot that is backed up by SMTape must not

be deleted on the source volume.

During the backup operation, SMTape locks the snapshot on the destination volume and if the corresponding snapshot is deleted on the source volume, then the subsequent SnapMirror update operation fails.

• You should not use these snapshots during incremental backup.

SMTape capabilities

SMTape capabilities such as backup of snapshots, incremental and differential backups, preservation of deduplication and compression features on restored volumes, and tape seeding help you optimize your tape backup and restore operations.

SMTape provides the following capabilities:

- · Provides a disaster recovery solution
- · Enables incremental and differential backups
- Backs up snapshots
- Enables backup and restore of deduplicated volumes and preserves deduplication on the restored volumes
- · Backs up compressed volumes and preserves compression on the restored volumes
- · Enables tape seeding

SMTape supports the blocking factor in multiples of 4 KB, in the range of 4 KB through 256 KB.



You can restore data to volumes created across up to two major consecutive ONTAP releases only.

Features not supported in SMTape

SMTape does not support restartable backups and verification of backed-up files.

Scalability limits for SMTape backup and restore sessions in ONTAP

While performing SMTape backup and restore operations through NDMP or CLI (tape seeding), you must be aware of the maximum number of SMTape backup and restore sessions that can be performed simultaneously on storage systems with different system memory capacities. This maximum number depends on the system memory of a storage system.



SMTape backup and restore sessions scalability limits are different from NDMP session limits and dump session limits.

System memory of the storage system	Total number of SMTape backup and restore sessions
Less than 16 GB	6

System memory of the storage system	Total number of SMTape backup and restore sessions
Greater than or equal to 16 GB but less than 24 GB	16
Greater than or equal to 24 GB	32

You can obtain the system memory of your storage system by using the sysconfig -a command (available through the nodeshell).

Learn more about sysconfig -a in the ONTAP command reference.

Related information

- Scalability limits for NDMP sessions
- Scalability limits for dump backup and restore sessions

What tape seeding is

Tape seeding is an SMTape functionality that helps you initialize a destination FlexVol volume in a data protection mirror relationship.

Tape seeding enables you to establish a data protection mirror relationship between a source system and a destination system over a low-bandwidth connection.

Incremental mirroring of snapshots from the source to the destination is feasible over a low bandwidth connection. However, an initial mirroring of the base snapshot takes a long time over a low-bandwidth connection. In such cases, you can perform an SMTape backup of the source volume to a tape and use the tape to transfer the initial base snapshot to the destination. You can then set up incremental SnapMirror updates to the destination system using the low-bandwidth connection.

How SMTape works with storage failover and ARL operations

Before you perform SMTape backup or restore operations, you should understand how these operations work with storage failover (takeover and giveback) or aggregate relocation (ARL) operation. The -override-vetoes option determines the behavior of SMTape engine during a storage failover or ARL operation.

When an SMTape backup or restore operation is running and the -override-vetoes option is set to false, a user-initiated storage failover or ARL operation is stopped and the backup or restore operation complete. If the backup application supports CAB extension, then you can continue performing incremental SMTape backup and restore operations without reconfiguring backup policies. However, if the -override-vetoes option is set to true, then the storage failover or ARL operation is continued and the SMTape backup or restore operation is aborted.

Related information

Network management

High Availability

How SMTape works with volume move

SMTape backup operations and volume move operations can run in parallel until the storage system attempts the final cutover phase. After this phase, new SMTape backup operations cannot run on the volume that is being moved. However, the current operations continue to run until completion.

Before the cutover phase for a volume is started, the volume move operation checks for active SMTape backup operations on the same volume. If there are active SMTape backup operations, then the volume move operation moves into a cutover deferred state and allows the SMTape backup operations to complete. After these backup operations are completed, you must manually restart the volume move operation.

If the backup application supports CAB extension, you can continue performing incremental tape backup and restore operations on read/write and read-only volumes without reconfiguring backup policies.

Baseline restore and volume move operations cannot be performed simultaneously; however, incremental restore can run in parallel with volume move operations, with the behavior similar to that of SMTape backup operations during volume move operations.

How SMTape works with volume rehost operations

SMTape operations cannot commence when a volume rehost operation is in progress on a volume. When a volume is involved in a volume rehost operation, SMTape sessions should not be started on that volume.

If any volume rehost operation is in progress, then SMTape backup or restore fails. If an SMTape backup or restore is in progress, then volume rehost operations fail with an appropriate error message. This condition applies to both NDMP-based and CLI-based backup or restore operations.

How NDMP backup policy are affected during ADB

When the automatic data balancer (ADB) is enabled, the balancer analyzes the usage statistics of aggregates to identify the aggregate that has exceeded the configured high-threshold usage percentage.

After identifying the aggregate that has exceeded the threshold, the balancer identifies a volume that can be moved to aggregates residing in another node in the cluster and attempts to move that volume. This situation affects the backup policy configured for this volume because if the data management application (DMA) is not CAB aware, then the user has to reconfigure the backup policy and run the baseline backup operation.



If the DMA is CAB aware and the backup policy has been configured using specific interface, then the ADB is not affected.

How SMTape backup and restore operations are affected in MetroCluster configurations

Before you perform SMTape backup and restore operations in a MetroCluster configuration, you must understand how SMTape operations are affected when a switchover or switchback operation occurs.

SMTape backup or restore operation followed by switchover

Consider two clusters: cluster 1 and cluster 2. During an SMTape backup or restore operation on cluster 1, if a switchover is initiated from cluster 1 to cluster 2, then the following occurs:

- If the value of the -override-vetoes option is false, then the switchover process is aborted and the backup or restore operation continues.
- If the value of the option is true, then the SMTape backup or restore operation is aborted and the switchover process continues.

SMTape backup or restore operation followed by switchback

A switchover is performed from cluster 1 to cluster 2 and an SMTape backup or restore operation is initiated on cluster 2. The SMTape operation backs up or restores a volume that is located on cluster 2. At this point, if a switchback is initiated from cluster 2 to cluster 1, then the following occurs:

- If the value of the -override-vetoes option is false, then the switchback process is aborted and the backup or restore operation continues.
- If the value of the option is true, then the backup or restore operation is aborted and the switchback process continues.

SMTape backup or restore operation initiated during a switchover or switchback

During a switchover process from cluster 1 to cluster 2, if an SMTape backup or restore operation is initiated on cluster 1, then the backup or restore operation fails and the switchover continues.

During a switchback process from cluster 2 to cluster 1, if an SMTape backup or restore operation is initiated from cluster 2, then the backup or restore operation fails and the switchback continues.

Monitor tape backup and restore operations for FlexVol volumes

Monitor tape backup and restore operations for FlexVol volumes overview

You can view the event log files to monitor the tape backup and restore operations. ONTAP automatically logs significant backup and restore events and the time at which they occur in a log file named backup in the controller's /etc/log/ directory. By default, event logging is set to on.

You might want to view event log files for the following reasons:

- Checking whether a nightly backup was successful
- · Gathering statistics on backup operations
- For using the information in past event log files to help diagnose problems with backup and restore operations

Once every week, the event log files are rotated. The /etc/log/backup file is renamed to /etc/log/backup.0, the /etc/log/backup.0 file is renamed to /etc/log/backup.1, and so on. The system saves the log files for up to six weeks; therefore, you can have up to seven message files (/etc/log/backup.[0-5] and the current /etc/log/backup file).

Access the event log files

You can access the event log files for tape backup and restore operations in the /etc/log/ directory by using the rdfile command at the nodeshell. You can view these event log files to monitor tape backup and restore operations.

About this task

With additional configurations, such as an access-control role with access to the spi web service or a user account set up with the http access method, you can also use a web browser to access these log files.

Steps

1. To access the nodeshell, enter the following command:

node run -node node_name

node name is the name of the node.

2. To access the event log files for tape backup and restore operations, enter the following command:

rdfile /etc/log/backup

Related information

System administration

What the dump and restore event log message format is

Dump and restore event log message format overview

For each dump and restore event, a message is written to the backup log file.

The format of the dump and restore event log message is as follows:

type timestamp identifier event (event info)

The following list describes the fields in the event log message format:

• Each log message begins with one of the type indicators described in the following table:

Туре	Description
log	Logging event
dmp	Dump event
rst	Restore event

- timestamp shows the date and time of the event.
- The identifier field for a dump event includes the dump path and the unique ID for the dump. The identifier field for a restore event uses only the restore destination path name as a unique identifier. Logging-related event messages do not include an identifier field.

What logging events are

The event field of a message that begins with a log specifies the beginning of a logging or the end of a logging.

It contains one of the events shown in the following table:

Event	Description
Start_Logging	Indicates the beginning of logging or that logging has been turned back on after being disabled.
Stop_Logging	Indicates that logging has been turned off.

What dump events are

The event field for a dump event contains an event type followed by event-specific information within parentheses.

The following table describes the events, their descriptions, and the related event information that might be recorded for a dump operation:

Event	Description	Event information
Start	NDMP dump is started	Dump level and the type of dump
End	Dumps completed successfully	Amount of data processed
Abort	The operation is cancelled	Amount of data processed
Options	Specified options are listed	All options and their associated values, including NDMP options
Tape_open	The tape is open for read/write	The new tape device name
Tape_close	The tape is closed for read/write	The tape device name
Phase-change	A dump is entering a new processing phase	The new phase name
Error	A dump has encountered an unexpected event	Error message
Snapshot	A snapshot is created or located	The name and time of the snapshot
Base_dump	A base dump entry in the internal metafile has been located	The level and time of the base dump (for incremental dumps only)

What restore events are

The event field for a restore event contains an event type followed by event-specific information in parentheses.

The following table provides information about the events, their descriptions, and the related event information that can be recorded for a restore operation:

Event	Description	Event information
Start	NDMP restore is started	Restore level and the type of restore
End	Restores completed successfully	Number of files and amount of data processed
Abort	The operation is cancelled	Number of files and amount of data processed
Options	Specified options are listed	All options and their associated values, including NDMP options
Tape_open	The tape is open for read/write	The new tape device name
Tape_close	The tape is closed for read/write	The tape device name
Phase-change	Restore is entering a new processing phase	The new phase name
Error	Restore encounters an unexpected event	Error message

Enabling or disabling event logging

You can turn the event logging on or off.

Steps

1. To enable or disable event logging, enter the following command at the clustershell:

options -option_name backup.log.enable -option-value {on | off}

on turns event logging on.

off turns event logging off.



Event logging is turned on by default.

Error messages for tape backup and restore of FlexVol volumes

Backup and restore error messages

Resource limitation: no available thread

• Message

Resource limitation: no available thread

Cause

The maximum number of active local tape I/O threads is currently in use. You can have a maximum of 16 active local tape drives.

Corrective action

Wait for some tape jobs to finish before starting a new backup or restore job.

Tape reservation preempted

• Message

Tape reservation preempted

Cause

The tape drive is in use by another operation or the tape has been closed prematurely.

Corrective action

Ensure that the tape drive is not in use by another operation and that the DMA application has not aborted the job and then retry.

Could not initialize media

• Message

Could not initialize media

Cause

You might get this error for one of the following reasons:

- The tape drive used for the backup is corrupt or damaged.
- $\circ\,$ The tape does not contain the complete backup or is corrupt.
- $\circ\,$ The maximum number of active local tape I/O threads is currently in use.

You can have a maximum of 16 active local tape drives.

Corrective action

• If the tape drive is corrupt or damaged, retry the operation with a valid tape drive.

- If the tape does not contain the complete backup or is corrupt, you cannot perform the restore operation.
- If tape resources are not available, wait for some of the backup or restore jobs to finish and then retry the operation.

Maximum number of allowed dumps or restores (maximum session limit) in progress

• Message

Maximum number of allowed dumps or restores (maximum session limit) in progress

Cause

The maximum number of backup or restore jobs is already running.

Corrective action

Retry the operation after some of the currently running jobs have finished.

Media error on tape write

• Message

Media error on tape write

Cause

The tape used for the backup is corrupted.

Corrective action

Replace the tape and retry the backup job.

Tape write failed

• Message

Tape write failed

Cause

The tape used for the backup is corrupted.

Corrective action

Replace the tape and retry the backup job.

Tape write failed - new tape encountered media error

• Message

Tape write failed - new tape encountered media error

Cause

The tape used for the backup is corrupted.

Corrective action

Replace the tape and retry the backup.

Tape write failed - new tape is broken or write protected

• Message

Tape write failed - new tape is broken or write protected

Cause

The tape used for the backup is corrupted or write-protected.

Corrective action

Replace the tape and retry the backup.

Tape write failed - new tape is already at the end of media

• Message

Tape write failed - new tape is already at the end of media

Cause

There is not enough space on the tape to complete the backup.

Corrective action

Replace the tape and retry the backup.

Tape write error

• Message

Tape write error - The previous tape had less than the required minimum capacity, size MB, for this tape operation, The operation should be restarted from the beginning $% \left(\frac{1}{2} \right) = 0$

Cause

The tape capacity is insufficient to contain the backup data.

Corrective action

Use tapes with larger capacity and retry the backup job.

Media error on tape read

• Message

Media error on tape read

Cause

The tape from which data is being restored is corrupted and might not contain the complete backup data.

Corrective action

If you are sure that the tape has the complete backup, retry the restore operation. If the tape does not contain the complete backup, you cannot perform the restore operation.

Tape read error

• Message

Tape read error

Cause

The tape drive is damaged or the tape does not contain the complete backup.

Corrective action

If the tape drive is damaged, use another tape drive. If the tape does not contain the complete backup, you cannot restore the data.

Already at the end of tape

• Message

Already at the end of tape

Cause

The tape does not contain any data or must be rewound.

Corrective action

If the tape does not contain data, use the tape that contains the backup and retry the restore job. Otherwise, rewind the tape and retry the restore job.

Tape record size is too small. Try a larger size.

• Message

Tape record size is too small. Try a larger size.

Cause

The blocking factor specified for the restore operation is smaller than the blocking factor that was used

during the backup.

Corrective action

Use the same blocking factor that was specified during the backup.

Tape record size should be block_size1 and not block_size2

• Message

Tape record size should be block size1 and not block size2

Cause

The blocking factor specified for the local restore is incorrect.

Corrective action

Retry the restore job with block size1 as the blocking factor.

Tape record size must be in the range between 4KB and 256KB

• Message

Tape record size must be in the range between 4KB and 256KB

Cause

The blocking factor specified for the backup or restore operation is not within the permitted range.

Corrective action

Specify a blocking factor in the range of 4 KB to 256 KB.

NDMP error messages

Network communication error

Message

Network communication error

Cause

Communication to a remote tape in an NDMP three-way connection has failed.

Corrective action

Check the network connection to the remote mover.

Message from Read Socket: error_string

Message

Message from Read Socket: error string

Cause

Restore communication from the remote tape in NDMP 3-way connection has errors.

Corrective action

Check the network connection to the remote mover.

Message from Write Dirnet: error_string

• Message

Message from Write Dirnet: error string

Cause

Backup communication to a remote tape in an NDMP three-way connection has an error.

Corrective action

Check the network connection to the remote mover.

Read Socket received EOF

• Message

Read Socket received EOF

Cause

Attempt to communicate with a remote tape in an NDMP three-way connection has reached the End Of File mark. You might be attempting a three-way restore from a backup image with a larger block size.

Corrective action

Specify the correct block size and retry the restore operation.

ndmpd invalid version number: version_number ``

• Message

ndmpd invalid version number: version number

Cause

The NDMP version specified is not supported by the storage system.

Corrective action

Specify NDMP version 4.

ndmpd session session_ID not active

• Message

ndmpd session session ID not active

Cause

The NDMP session might not exist.

Corrective action

Use the ndmpd status command to view the active NDMP sessions.

Could not obtain vol ref for Volume volume_name

• Message

Could not obtain vol ref for Volume vol name

Cause

The volume reference could not be obtained because the volume might be in use by other operations.

Corrective action

Retry the operation later.

Data connection type ["NDMP4_ADDR_TCP"|"NDMP4_ADDR_TCP_IPv6"] not supported for ["IPv6"|"IPv4"] control connections

Message

```
Data connection type ["NDMP4_ADDR_TCP"|"NDMP4_ADDR_TCP_IPv6"] not supported for ["IPv6"|"IPv4"] control connections
```

Cause

In node-scoped NDMP mode, the NDMP data connection established must be of the same network address type (IPv4 or IPv6) as the NDMP control connection.

Corrective action

Contact your backup application vendor.

DATA LISTEN: CAB data connection prepare precondition error

• Message

DATA LISTEN: CAB data connection prepare precondition error

Cause

NDMP data listen fails when the backup application has negotiated the CAB extension with the NDMP

server and there is a mismatch in the specified NDMP data connection address type between the NDMP_CAB_DATA_CONN_PREPARE and the NDMP_DATA_LISTEN messages.

Corrective action

Contact your backup application vendor.

DATA CONNECT: CAB data connection prepare precondition error

• Message

DATA CONNECT: CAB data connection prepare precondition error

Cause

NDMP data connect fails when the backup application has negotiated the CAB extension with the NDMP server and there is a mismatch in the specified NDMP data connection address type between the NDMP_CAB_DATA_CONN_PREPARE and the NDMP_DATA_CONNECT messages.

Corrective action

Contact your backup application vendor.

Error:show failed: Cannot get password for user '<username>'

Message

Error: show failed: Cannot get password for user '<username>'

Cause

Incomplete user account configuration for NDMP

Corrective action

Ensure that the user account is associated with the SSH access method and the authentication method is user password.

Dump error messages

Destination volume is read-only

• Message

Destination volume is read-only

Cause

The path to which the restore operation is attempted to is read-only.

Corrective action

Try restoring the data to a different location.

Destination qtree is read-only

• Message

Destination qtree is read-only

Cause

The qtree to which the restore is attempted to is read-only.

Corrective action

Try restoring the data to a different location.

Dumps temporarily disabled on volume, try again

• Message

Dumps temporarily disabled on volume, try again

Cause

NDMP dump backup is attempted on a SnapMirror destination volume that is part of either a snapmirror break or a snapmirror resync operation.

Corrective action

Wait for the snapmirror break or snapmirror resync operation to finish and then perform the dump operation.



Whenever the state of a SnapMirror destination volume changes from read/write to readonly or from read-only to read/write, you must perform a baseline backup.

Related information

- snapmirror break
- snapmirror resync

NFS labels not recognized

• Message

Error: Aborting: dump encountered NFS security labels in the file system

Cause

NFS security labels are supported Beginning with ONTAP 9.9.1 when NFSv4.2 is enabled. However, NFS security labels are not currently recognized by the dump engine. If it encounters any NFS security labels on the files, directories, or any special files in any format of dump, the dump fails.

Corrective action

Verify that no files or directories have NFS security labels.

No files were created

• Message

No files were created

Cause

A directory DAR was attempted without enabling the enhanced DAR functionality.

Corrective action

Enable the enhanced DAR functionality and retry the DAR.

Restore of the file <file name> failed

• Message

Restore of the file file name failed

Cause

When a DAR (Direct Access Recovery) of a file whose file name is the same as that of a LUN on the destination volume is performed, then the DAR fails.

Corrective action

Retry DAR of the file.

Truncation failed for src inode <inode number>...

• Message

```
Truncation failed for src inode <inode number>. Error <error number>. Skipping inode.
```

Cause

Inode of a file is deleted when the file is being restored.

Corrective action

Wait for the restore operation on a volume to complete before using that volume.

Unable to lock a snapshot needed by dump

• Message

Unable to lock a snapshot needed by dump

Cause

The snapshot specified for the backup is not available.

Corrective action

Retry the backup with a different snapshot.

Use the snap list command to see the list of available snapshots.

Learn more about snap list in the ONTAP command reference.

Unable to locate bitmap files

• Message

Unable to locate bitmap files

Cause

The bitmap files required for the backup operation might have been deleted. In this case, the backup cannot be restarted.

Corrective action

Perform the backup again.

Volume is temporarily in a transitional state

• Message

Volume is temporarily in a transitional state

Cause

The volume being backed up is temporarily in an unmounted state.

Corrective action

Wait for some time and perform the backup again.

SMTape error messages

Chunks out of order

• Message

Chunks out of order

Cause

The backup tapes are not being restored in the correct sequence.

Corrective action

Retry the restore operation and load the tapes in the correct sequence.

Chunk format not supported

• Message

Chunk format not supported

Cause

The backup image is not of SMTape.

Corrective action

If the backup image is not of SMTape, retry the operation with a tape that has the SMTape backup.

Failed to allocate memory

• Message

Failed to allocate memory

Cause

The system has run out of memory.

Corrective action

Retry the job later when the system is not too busy.

Failed to get data buffer

• Message

Failed to get data buffer

Cause

The storage system ran out of buffers.

Corrective action

Wait for some storage system operations to finish and then retry the job.

Failed to find snapshot

• Message

Failed to find snapshot

Cause

The snapshot specified for the backup is unavailable.

Corrective action

Check if the specified snapshot is available. If not, retry with the correct snapshot.

Failed to create snapshot

• Message

Failed to create snapshot

Cause

The volume already contains the maximum number of snapshots.

Corrective action

Delete some snapshots and then retry the backup operation.

Failed to lock snapshot

• Message

Failed to lock snapshot

Cause

The snapshot is either in use or has been deleted.

Corrective action

If the snapshot is in use by another operation, wait for that operation to finish and then retry the backup. If the snapshot has been deleted, you cannot perform the backup.

Failed to delete snapshot

• Message

Failed to delete snapshot

Cause

The auto snapshot could not be deleted because it is in use by other operations.

Corrective action

Use the snap command to determine the status of the snapshot. If the snapshot is not required, delete it manually.

Failed to get latest snapshot

• Message

Failed to get latest snapshot

Cause
The latest snapshot might not exist because the volume is being initialized by SnapMirror.

Corrective action

Retry after initialization is complete.

Failed to load new tape

• Message

Failed to load new tape

Cause

Error in tape drive or media.

Corrective action

Replace the tape and retry the operation.

Failed to initialize tape

• Message

Failed to initialize tape

Cause

You might get this error message for one of the following reasons:

- The backup image is not of SMTape.
- The tape blocking factor specified is incorrect.
- The tape is corrupt or damaged.
- The wrong tape is loaded for restore.

Corrective action

- If the backup image is not of SMTape, retry the operation with a tape that has SMTape backup.
- If the blocking factor is incorrect, specify the correct blocking factor and retry the operation.
- If the tape is corrupt, you cannot perform the restore operation.
- If the wrong tape is loaded, retry the operation with the correct tape.

Failed to initialize restore stream

• Message

Failed to initialize restore stream

Cause

You might get this error message for one of the following reasons:

• The backup image is not of SMTape.

- The tape blocking factor specified is incorrect.
- The tape is corrupt or damaged.
- The wrong tape is loaded for restore.

Corrective action

- If the backup image is not of SMTape, retry the operation with a tape that has the SMTape backup.
- If the blocking factor is incorrect, specify the correct blocking factor and retry the operation.
- $\circ\,$ If the tape is corrupt, you cannot perform the restore operation.
- $\circ\,$ If the wrong tape is loaded, retry the operation with the correct tape.

Failed to read backup image

• Message

Failed to read backup image

Cause

The tape is corrupt.

Corrective action

If the tape is corrupt, you cannot perform the restore operation.

Image header missing or corrupted

• Message

Image header missing or corrupted

Cause

The tape does not contain a valid SMTape backup.

Corrective action

Retry with a tape containing a valid backup.

Internal assertion

• Message

Internal assertion

Cause

There is an internal SMTape error.

Corrective action

Report the error and send the etc/log/backup file to technical support.

Invalid backup image magic number

• Message

Invalid backup image magic number

Cause

The backup image is not of SMTape.

Corrective action

If the backup image is not of SMTape, retry the operation with a tape that has the SMTape backup.

Invalid backup image checksum

• Message

Invalid backup image checksum

Cause

The tape is corrupt.

Corrective action

If the tape is corrupt, you cannot perform the restore operation.

Invalid input tape

• Message

Invalid input tape

Cause

The signature of the backup image is not valid in the tape header. The tape has corrupted data or does not contain a valid backup image.

Corrective action

Retry the restore job with a valid backup image.

Invalid volume path

• Message

Invalid volume path

Cause

The specified volume for the backup or restore operation is not found.

Corrective action

Retry the job with a valid volume path and volume name.

Mismatch in backup set ID

• Message

Mismatch in backup set ID

Cause

The tape loaded during a tape change is not a part of the backup set.

Corrective action

Load the correct tape and retry the job.

Mismatch in backup time stamp

• Message

Mismatch in backup time stamp

Cause

The tape loaded during a tape change is not a part of the backup set.

Corrective action

Use the smtape restore -h command to verify the header information of a tape.

Job aborted due to shutdown

• Message

Job aborted due to shutdown

Cause

The storage system is being rebooted.

Corrective action

Retry the job after the storage system reboots.

Job aborted due to snapshot autodelete

• Message

Job aborted due to snapshot autodelete

Cause

The volume does not have enough space and has triggered the automatic deletion of snapshots.

Corrective action

Free up space in the volume and retry the job.

Tape is currently in use by other operations

• Message

Tape is currently in use by other operations

Cause

The tape drive is in use by another job.

Corrective action

Retry the backup after the currently active job is finished.

Tapes out of order

• Message

Tapes out of order

Cause

The first tape of the tape sequence for the restore operation does not have the image header.

Corrective action

Load the tape with the image header and retry the job.

Transfer failed (Aborted due to MetroCluster operation)

• Message

Transfer failed (Aborted due to MetroCluster operation)

Cause

The SMTape operation is aborted because of a switchover or switchback operation.

Corrective action

Perform the SMTape operation after the switchover or switchback operation finishes.

Transfer failed (ARL initiated abort)

• Message

```
Transfer failed (ARL initiated abort)
```

Cause

While an SMTape operation is in progress if an aggregate relocation is initiated, then the SMTape operation is aborted.

Corrective action

Perform the SMTape operation after the aggregate relocation operation finishes.

Transfer failed (CFO initiated abort)

• Message

Transfer failed (CFO initiated abort)

Cause

The SMTape operation is aborted because of a storage failover (takeover and giveback) operation of a CFO aggregate.

Corrective action

Perform the SMTape operation after the storage failover of the CFO aggregate finishes.

Transfer failed (SFO initiated abort)

• Message

Transfer failed (SFO initiated abort)

Cause

The SMTape operation is aborted because of a storage failover (takeover and giveback) operation.

Corrective action

Perform the SMTape operation after the storage failover (takeover and giveback) operation finishes.

Underlying aggregate under migration

• Message

Underlying aggregate under migration

Cause

If an SMTape operation is initiated on an aggregate that is under migration (storage failover or aggregate relocation), then the SMTape operation fails.

Corrective action

Perform the SMTape operation after the aggregate migration finishes.

Volume is currently under migration

• Message

Volume is currently under migration

Cause

Volume migration and SMTape backup cannot run simultaneously.

Corrective action

Retry the backup job after the volume migration is complete.

Volume offline

• Message

Volume offline

Cause

The volume being backed up is offline.

Corrective action

Bring the volume online and retry the backup.

Volume not restricted

• Message

Volume not restricted

Cause

The destination volume to which data is being restored is not restricted.

Corrective action

Restrict the volume and retry the restore operation.

NDMP configuration

Learn about ONTAP NDMP configuration

You can quickly configure an ONTAP 9 cluster to use the Network Data Management Protocol (NDMP) to back up data directly to tape using a third-party backup application.

If the backup application supports Cluster Aware Backup (CAB), you can configure NDMP as *SVM-scoped* or *node-scoped*:

• SVM-scoped at the cluster (admin SVM) level enables you to back up all volumes hosted across different

nodes of the cluster. SVM-scoped NDMP is recommended where possible.

• Node-scoped NDMP enables you to back up all the volumes hosted on that node.

If the backup application does not support CAB, you must use node-scoped NDMP.

SVM-scoped and node-scoped NDMP are mutually exclusive; they cannot be configured on the same cluster.



Node-scoped NDMP is deprecated in ONTAP 9.

Learn more about Cluster Aware Backup (CAB).

Before configuring NDMP, verify the following:

- You have a third-party backup application (also called a Data Management Application or DMA).
- You are a cluster administrator.
- Tape devices and an optional media server are installed.
- Tape devices are connected to the cluster through a Fibre Channel (FC) switch or locally attached.
- At least one tape device has a logical unit number (LUN) of 0.

Learn about ONTAP NDMP configuration workflow

Setting up tape backup over NDMP involves preparing for NDMP configuration, verifying the tape device connections, enabling tape reservations, configuring NDMP at the SVM or node level, enabling NDMP on the cluster, configuring a backup user, configuring LIFs, and configuring the backup application.



Prepare ONTAP NDMP configurations

Before you configure tape backup access over Network Data Management Protocol (NDMP), you must verify that the planned configuration is supported, verify that your tape drives are listed as qualified drives on each node, verify that all nodes have intercluster LIFs, and identify whether the backup application supports the Cluster Aware Backup (CAB) extension.

Steps

1. Refer to your backup application provider's compatibility matrix for ONTAP support (NetApp does not qualify third-party backup applications with ONTAP or NDMP).

You should verify that the following NetApp components are compatible:

- The version of ONTAP 9 that is running on the cluster.
- The backup application vendor and version: for example, Veritas NetBackup 8.2 or CommVault.

- The tape devices details, such as the manufacturer, model, and interface of the tape drives: for example, IBM Ultrium 8 or HPe StoreEver Ultrium 30750 LTO-8.
- The platforms of the nodes in the cluster: for example, FAS8700 or A400.



You can find legacy ONTAP compatibility support matrices for backup applications in the NetApp Interoperability Matrix Tool.

- 2. Verify that your tape drives are listed as qualified drives in each node's built-in tape configuration file:
 - a. On the command line-interface, view the built-in tape configuration file by using the storage tape show-supported-status command.

<pre>cluster1::> storage tape show-supported-status</pre>							
Node: cluster1-1							
	Is						
Tape Drives	Supported	Support Status					
Certance Ultrium 2	true	Dynamically Qualified					
Certance Ultrium 2 Certance Ultrium 3	true true	Dynamically Qualified Dynamically Qualified					

b. Compare your tape drives to the list of qualified drives in the output.



The names of the tape devices in the output might vary slightly from the names on the device label or in the Interoperability Matrix. For example, Digital DLT2000 can also be known as DLT2k. You can ignore these minor naming differences.

c. If a device is not listed as qualified in the output even though the device is qualified according to the Interoperability Matrix, download and install an updated configuration file for the device using the instructions on the NetApp Support Site.

NetApp Downloads: Tape Device Configuration Files

A qualified device might not be listed in the built-in tape configuration file if the tape device was qualified after the node was shipped.

- 3. Verify that every node in the cluster has an intercluster LIF:
 - a. View the intercluster LIFs on the nodes by using the network interface show -role intercluster command.

Learn more about network interface show in the ONTAP command reference.

b. If an intercluster LIF does not exist on any node, create an intercluster LIF by using the network interface create command.

```
cluster1::> network interface create -vserver cluster1 -lif IC2 -role
intercluster
-home-node cluster1-2 -home-port e0b -address 192.0.2.68 -netmask
255.255.255.0
-status-admin up -failover-policy local-only -firewall-policy
intercluster
cluster1::> network interface show -role intercluster
         Logical Status Network
                                          Current
Current Is
Vserver Interface Admin/Oper Address/Mask
                                          Node
Port
     Home
_____ ____
_____ ___
cluster1 IC1
                  up/up 192.0.2.65/24 cluster1-1
e0a true
cluster1 IC2
                  up/up 192.0.2.68/24 cluster1-2
e0b true
```

Learn more about network interface create in the ONTAP command reference.

Network management

4. Identify whether the backup application supports Cluster Aware Backup (CAB) by using the documentation provided with the backup application.

CAB support is a key factor in determining the type of backup you can perform.

Verify ONTAP NDMP tape device connections

You must ensure that all drives and media changers are visible in ONTAP as devices.

Steps

1. View information about all drives and media changers by using the storage tape show command.

```
cluster1::> storage tape show
Node: cluster1-01
Device ID
                 Device Type Description
Status
_____ ____
_____
sw4:10.11
               tape drive HP LTO-3
normal
            media changer HP MSL G3 Series
0b.125L1
normal
0d.4
                 tape drive IBM LTO 5 ULT3580
normal
                  media changer IBM 3573-TL
0d.4L1
normal
. . .
```

- 2. If a tape drive is not displayed, troubleshoot the problem.
- 3. If a media changer is not displayed, view information about media changers by using the storage tape show-media-changer command, and then troubleshoot the problem.

```
cluster1::> storage tape show-media-changer
Media Changer: sw4:10.11L1
 Description: PX70-TL
       WWNN: 2:00a:000e11:10b919
       WWPN: 2:00b:000e11:10b919
Serial Number: 00FRU7800000 LL1
      Errors: -
Paths:
                      Initiator Alias Device State
Node
Status
____
                                        _____
_____
cluster1-01
                     2b mc0 in-use
normal
. . .
```

Enable tape reservations for ONTAP NDMP backup operations

You must ensure that tape drives are reserved for use by backup applications for NDMP backup operations.

About this task

The reservation settings vary in different backup applications, and these settings must match the backup application and the nodes or servers using the same drives. See the vendor documentation of the backup application for the correct reservation settings.

Steps

1. Enable reservations by using the options -option-name tape.reservations -option-value persistent command.

The following command enables reservations with the persistent value:

```
cluster1::> options -option-name tape.reservations -option-value
persistent
2 entries were modified.
```

2. Verify that reservations are enabled on all nodes by using the options tape.reservations command, and then review the output.

```
cluster1::> options tape.reservations
cluster1-1
  tape.reservations persistent
cluster1-2
  tape.reservations persistent
2 entries were displayed.
```

Configure SVM-scoped NDMP

Enable SVM-scoped NDMP on the ONTAP cluster

If the DMA supports the Cluster Aware Backup (CAB) extension, you can back up all the volumes hosted across different nodes in a cluster by enabling SVM-scoped NDMP, enabling NDMP service on the cluster (admin SVM), and configuring LIFs for data and control connection.

Before you begin

The CAB extension must be supported by the DMA.

About this task

Turning off node-scoped NDMP mode enables SVM-scoped NDMP mode on the cluster.

Steps

1. Enable SVM-scoped NDMP mode:

cluster1::> system services ndmp node-scope-mode off

SVM-scoped NDMP mode is enabled.

2. Enable NDMP service on the admin SVM:

cluster1::> vserver services ndmp on -vserver cluster1

The authentication type is set to challenge by default and plaintext authentication is disabled.



For secure communication, you should keep plaintext authentication disabled.

3. Verify that NDMP service is enabled:

cluster1::> vserver services ndmp show

```
VserverEnabledAuthentication type------------------cluster1truechallengevs1falsechallenge
```

Enable backup users for ONTAP NDMP authentication

To authenticate SVM-scoped NDMP from the backup application, there must be an administrative user with sufficient privileges and an NDMP password.

About this task

You must generate an NDMP password for backup admin users. You can enable backup admin users at the cluster or SVM level, and if necessary, you can create a new user. By default, the users with the following roles can authenticate for NDMP backup:

- Cluster-wide: admin or backup
- Individual SVMs: vsadmin or vsadmin-backup

If you are using an NIS or LDAP user, the user must exist on the respective server. You cannot use an Active Directory user.

Steps

1. Display the current admin users and permissions:

security login show

Learn more about security login show in the ONTAP command reference.

2. If needed, create a new NDMP backup user with the security login create command and the appropriate role for cluster-wide or individual SVM privileges.

You can specify a local backup user name or an NIS or LDAP user name for the -user-or-group-name parameter.

The following command creates the backup user <code>backup_admin1</code> with the <code>backup</code> role for the entire cluster:

cluster1::> security login create -user-or-group-name backup_admin1
-application ssh -authmethod password -role backup

The following command creates the backup user <code>vsbackup_admin1</code> with the <code>vsadmin-backup</code> role for an individual SVM:

```
cluster1::> security login create -user-or-group-name vsbackup_admin1
-application ssh -authmethod password -role vsadmin-backup
```

Enter a password for the new user and confirm.

Learn more about security login create in the ONTAP command reference.

3. Generate a password for the admin SVM by using the vserver services ndmp generate password command.

The generated password must be used to authenticate the NDMP connection by the backup application.

```
cluster1::> vserver services ndmp generate-password -vserver cluster1
-user backup_admin1
Vserver: cluster1
   User: backup_admin1
Password: qG5CqQHYxw7tE57g
```

Configure ONTAP LIFs for SVM-scoped NDMP

You must identify the LIFs that will be used for establishing a data connection between the data and tape resources, and for control connection between the admin SVM and the backup application. After identifying the LIFs, you must verify the service and failover policies are set.



Beginning with ONTAP 9.10.1, firewall policies are deprecated and wholly replaced with LIF service policies. For more information, see Manage supported traffic.

ONTAP 9.10.1 or later

Steps

1. Identify the intercluster LIF hosted on the nodes by using the network interface show command with the -service-policy parameter.

network interface show -service-policy default-intercluster

Learn more about network interface show in the ONTAP command reference.

2. Identify the management LIF hosted on the nodes by using the network interface show command with the -service-policy parameter.

network interface show -service-policy default-management

3. Ensure that the intercluster LIF includes the backup-ndmp-control service:

network interface service-policy show

Learn more about network interface service-policy show in the ONTAP command reference.

- 4. Ensure that the failover policy is set appropriately for all the LIFs:
 - a. Verify that the failover policy for the cluster-management LIF is set to broadcast-domainwide, and the policy for the intercluster and node-management LIFs is set to local-only by using the network interface show -failover command.

The following command displays the failover policy for the cluster-management, intercluster, and node-management LIFs:

cluster1::> network interface show -failover Logical Home Failover Failover Vserver Interface Node:Port Policy Group _____ _____ _____ cluster cluster1 clus1 cluster1-1:e0a local-only cluster Failover Targets: cluster1 cluster mgmt cluster1-1:e0m broadcast- Default domain-wide Failover Targets: IC1 cluster1-1:e0a local-only Default Failover Targets: IC2 cluster1-1:e0b local-only Default Failover Targets: cluster1-1 c1-1 mgmt1 cluster1-1:e0m local-only Default Failover Targets: cluster1-2 c1-2 mgmt1 cluster1-2:e0m local-only Default Failover Targets:

b. If the failover policies are not set appropriately, modify the failover policy by using the network interface modify command with the -failover-policy parameter.

cluster1::> network interface modify -vserver cluster1 -lif IC1
-failover-policy local-only

Learn more about network interface modify in the ONTAP command reference.

5. Specify the LIFs that are required for data connection by using the vserver services ndmp modify command with the preferred-interface-role parameter.

cluster1::> vserver services ndmp modify -vserver cluster1
-preferred-interface-role intercluster, cluster-mgmt, node-mgmt

6. Verify that the preferred interface role is set for the cluster by using the vserver services ndmp show command.

```
cluster1::> vserver services ndmp show -vserver cluster1
            Vserver: cluster1
            NDMP Version: 4
            .....
Preferred Interface Role: intercluster, cluster-mgmt, node-mgmt
```

ONTAP 9.9 or earlier

Steps

1. Identify the intercluster, cluster-management, and node-management LIFs by using the network interface show command with the -role parameter.

The following command displays the intercluster LIFs:

The following command displays the cluster-management LIF:

The following command displays the node-management LIFs: cluster1::> network interface show -role node-mgmt Logical Status Network Current Current Is Vserver Interface Admin/Oper Address/Mask Node Port Home _____ _____ ----- ----cluster1 cluster1-1 mgmt1 up/up 192.0.2.69/24 cluster1-1 eOM true cluster1-2 mgmt1 up/up 192.0.2.70/24 cluster1-2 eOM true

Learn more about network interface show in the ONTAP command reference.

- 2. Ensure that the firewall policy is enabled for NDMP on the intercluster, cluster-management (cluster-mgmt), and node-management (node-mgmt) LIFs:
 - a. Verify that the firewall policy is enabled for NDMP by using the system services firewall policy show command.

The following command displays the firewall policy for the cluster-management LIF:

cluster1::> system services firewall policy show -policy cluster Vserver Policy Service Allowed _____ cluster cluster 0.0.0.0/0 dns 0.0.0.0/0 http https 0.0.0.0/0 ndmp 0.0.0.0/0 ndmps 0.0.0.0/0 0.0.0.0/0 ntp rsh 0.0.0.0/0 0.0.0.0/0 snmp ssh 0.0.0.0/0 telnet 0.0.0.0/0 10 entries were displayed.

The following command displays the firewall policy for the intercluster LIF:

```
cluster1::> system services firewall policy show -policy
intercluster
Vserver Policy Service Allowed
_____
           cluster1 intercluster dns
                  http
                          _
                  https
                          _
                  ndmp 0.0.0/0, ::/0
                  ndmps
                           _
                  ntp
                  rsh
                  ssh
                          _
                  telnet
                         _
9 entries were displayed.
```

The following command displays the firewall policy for the node-management LIF:

```
cluster1::> system services firewall policy show -policy mgmt
Vserver Policy Service Allowed
              ---- ---- ----
_____
                                   _____
cluster1-1 mgmt
                    dns
                             0.0.0.0/0, ::/0
                     http 0.0.0/0, ::/0
                     https
                             0.0.0.0/0, ::/0
                     ndmp
                              0.0.0/0, ::/0
                     ndmps
                             0.0.0.0/0, ::/0
                     ntp
                              0.0.0.0/0, ::/0
                     rsh
                             0.0.0.0/0, ::/0
                     snmp
                     ssh
                             0.0.0.0/0, ::/0
                     telnet
10 entries were displayed.
```

b. If the firewall policy is not enabled, enable the firewall policy by using the system services firewall policy modify command with the -service parameter.

The following command enables firewall policy for the intercluster LIF:

cluster1::> system services firewall policy modify -vserver cluster1 -policy intercluster -service ndmp 0.0.0.0/0

3. Ensure that the failover policy is set appropriately for all the LIFs:

a. Verify that the failover policy for the cluster-management LIF is set to broadcast-domainwide, and the policy for the intercluster and node-management LIFs is set to local-only by using the network interface show -failover command.

The following command displays the failover policy for the cluster-management, intercluster, and node-management LIFs:

cluster1::> network interface show -failover Logical Home Failover Failover Vserver Interface Node:Port Policy Group ----- ---------- ----cluster cluster1_clus1 cluster1-1:e0a local-only cluster Failover Targets: cluster1 cluster mgmt cluster1-1:e0m broadcast-domainwide Default Failover Targets: IC1 cluster1-1:e0a local-only Default Failover Targets: IC2 cluster1-1:e0b local-only Default Failover Targets: cluster1-1 cluster1-1_mgmt1 cluster1-1:e0m local-only Default Failover Targets: cluster1-2 cluster1-2 mgmt1 cluster1-2:e0m local-only Default Failover Targets:

b. If the failover policies are not set appropriately, modify the failover policy by using the network interface modify command with the -failover-policy parameter.

```
cluster1::> network interface modify -vserver cluster1 -lif IC1
-failover-policy local-only
```

Learn more about network interface modify in the ONTAP command reference.

4. Specify the LIFs that are required for data connection by using the vserver services ndmp modify command with the preferred-interface-role parameter.

```
cluster1::> vserver services ndmp modify -vserver cluster1
-preferred-interface-role intercluster,cluster-mgmt,node-mgmt
```

5. Verify that the preferred interface role is set for the cluster by using the vserver services ndmp show command.

Configure node-scoped NDMP

Enable node-scoped NDMP on the ONTAP cluster

You can back up volumes hosted on a single node by enabling node-scoped NDMP, enabling the NDMP service, and configuring a LIF for data and control connection. This can be done for all nodes of the cluster.



Node-scoped NDMP is deprecated in ONTAP 9.

About this task

When using NDMP in node-scope mode, authentication must be configured on a per-node basis. For more information, see the Knowledge Base article "How to configure NDMP authentication in the 'node-scope' mode".

Steps

1. Enable node-scoped NDMP mode:

cluster1::> system services ndmp node-scope-mode on

NDMP node-scope-mode is enabled.

2. Enable NDMP service on all nodes in the cluster:

Using the wildcard "*" enables NDMP service on all nodes at the same time.

You must specify a password for authentication of the NDMP connection by the backup application.

```
cluster1::> system services ndmp on -node *
```

```
Please enter password:
Confirm password:
2 entries were modified.
```

3. Disable the -clear-text option for secure communication of the NDMP password:

Using the wildcard "*" disables the -clear-text option on all nodes at the same time.

cluster1::> system services ndmp modify -node * -clear-text false

4. Verify that NDMP service is enabled and the -clear-text option is disabled:

```
cluster1::> system services ndmp show
```

Node	Enabled	Clear text	User Id		
cluster1-1	true	false	root		
cluster1-2	true	false	root		
2 entries were displayed.					

Configure ONTAP LIFs for node-scoped NDMP

You must identify a LIF that will be used for establishing a data connection and control connection between the node and the backup application. After identifying the LIF, you must verify that firewall and failover policies are set for the LIF.



Beginning with ONTAP 9.10.1, firewall policies are deprecated and wholly replaced with LIF service policies. For more information, see Manage supported traffic.

ONTAP 9.10.1 or later

Steps

1. Identify the intercluster LIF hosted on the nodes by using the network interface show command with the -service-policy parameter.

network interface show -service-policy default-intercluster

2. Ensure that the intercluster LIF includes the backup-ndmp-control service:

network interface service-policy show

- 3. Ensure that the failover policy is set appropriately for the intercluster LIFs:
 - a. Verify that the failover policy for the intercluster LIFs is set to local-only by using the network interface show -failover command.

cluster1::> network interface show -failover Failover Logical Home Failover Vserver Interface Node:Port Policy Group _____ _____ _____ _ _____ cluster1 IC1 cluster1-1:e0a local-only Default Failover Targets: cluster1-2:e0b local-only IC2 Default Failover Targets: cluster1-1 cluster1-1 mgmt1 cluster1-1:e0m local-only Default Failover Targets:

b. If the failover policy is not set appropriately, modify the failover policy by using the network interface modify command with the -failover-policy parameter.

cluster1::> network interface modify -vserver cluster1 -lif IC1
-failover-policy local-only

Learn more about network interface show, network interface service-policy show,

and network interface modify in the ONTAP command reference.

ONTAP 9.9 or earlier

Steps

1. Identify the intercluster LIF hosted on the nodes by using the network interface show command with the -role parameter.

- 2. Ensure that the firewall policy is enabled for NDMP on the intercluster LIFs:
 - a. Verify that the firewall policy is enabled for NDMP by using the system services firewall policy show command.

The following command displays the firewall policy for the intercluster LIF:

```
cluster1::> system services firewall policy show -policy
intercluster
Vserver Policy Service Allowed
_____
        -----
cluster1 intercluster dns
                   http
                   https
                   ndmp
                          0.0.0.0/0, ::/0
                   ndmps
                   ntp
                           _
                   rsh
                           _
                   ssh
                           _
                   telnet
                           _
9 entries were displayed.
```

b. If the firewall policy is not enabled, enable the firewall policy by using the system services firewall policy modify command with the -service parameter.

The following command enables firewall policy for the intercluster LIF:

```
cluster1::> system services firewall policy modify -vserver
cluster1 -policy intercluster -service ndmp 0.0.0.0/0
```

- 3. Ensure that the failover policy is set appropriately for the intercluster LIFs:
 - a. Verify that the failover policy for the intercluster LIFs is set to local-only by using the network interface show -failover command.

cluster1::>	ter1::> network interface show -failover						
	Logical	Home	Failover				
Failover							
Vserver	Interface	Node:Port	Policy	Group			
cluster1	IC1	cluster1-1:e0a	local-only				
Default							
			Failove	r			
Targets:							
	T Q Q	1 1 0 01	•••••				
Defeult	102	Cluster1-2:eUb	local-only				
Delault			Failoro	~			
Targets.			rallove	L			
largets.							
cluster1-1	cluster1-1 momt1	cluster1-1.eOm	local-only				
Default			rocar onry				
Deradre			Failove	r			
Targets:			0				
2							

b. If the failover policy is not set appropriately, modify the failover policy by using the network interface modify command with the -failover-policy parameter.

```
cluster1::> network interface modify -vserver cluster1 -lif IC1
-failover-policy local-only
```

Learn more about network interface show and network interface modify in the ONTAP command reference.

Configure backup applications for ONTAP NDMP configuration

After the cluster is configured for NDMP access, you must gather information from the cluster configuration and then configure the rest of the backup process in the backup application.

Steps

- 1. Gather the following information that you configured earlier in ONTAP:
 - The user name and password that the backup application requires to create the NDMP connection
 - The IP addresses of the intercluster LIFs that the backup application requires to connect to the cluster
- 2. In ONTAP, display the aliases that ONTAP assigned to each device by using the storage tape alias show command.

The aliases are often useful in configuring the backup application.

```
cluster1::> storage tape show -alias

Device ID: 2a.0

Device Type: tape drive

Description: Hewlett-Packard LTO-5

Node Alias Mapping

stsw-3220-4a-4b-02 st2 SN[HU19497WVR]

...
```

3. In the backup application, configure the rest of the backup process by using the backup application's documentation.

After you finish

If a data mobility event occurs, such as a volume move or LIF migration, you must be prepared to reinitialize any interrupted backup operations.

Replication between NetApp Element software and ONTAP overview

You can ensure business continuity on an Element system by using SnapMirror to replicate snapshots of an Element volume to an ONTAP destination. In the event of a disaster at the Element site, you can serve data to clients from the ONTAP system, and then reactivate the Element system when service is restored.

Beginning with ONTAP 9.4, you can replicate snapshots of a LUN created on an ONTAP node back to an Element system. You might have created a LUN during an outage at the Element site, or you might be using a LUN to migrate data from ONTAP to Element software.

Configure replication of NetApp Element software and ONTAP.

Copyright information

Copyright © 2025 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.