Manage your storage

ASA r2

NetApp
February 11, 2026

Table of Contents

Manage your storage

Provision ONTAP SAN storage on the ASA r2 systems

When you provision storage, you enable your SAN hosts to read from and write data to ASA r2 storage systems. To provision storage, you use ONTAP System Manager to create storage units, add host initiators, and map the host to a storage unit. You also need to perform steps on the host to enable read/write operations.

Create storage units

On an ASA r2 system, a storage unit makes storage space available to your SAN hosts for data operations. A storage unit refers to a LUN for SCSI hosts or an NVMe namespace for NVMe hosts. If your cluster is configured to support SCSI hosts, you are prompted to create a LUN. If your cluster is configured to support NVMe hosts, you are prompted to create an NVMe namespace.

An ASA r2 storage unit has a maximum capacity of 128 TB. See the NetApp Hardware Universe for the most current storage limits for ASA r2 systems.

You add and map host initiators to the storage unit as part of the storage unit creation process. You can also add and map host initiators after you create the storage units.

Beginning with ONTAP 9.18.1, you can modify the snapshot reserve and enable automatic snapshot deletion when you create a storage unit. The snapshot reserve is the amount of space in the storage unit reserved specifically for snapshots. When snapshot reserve is set with automatic snapshot deletion, older snapshots are automatically deleted when space used by snapshots exceeds the snapshot reserve.

Learn more about snapshot reserve on ASA r2 systems.

Storage units are thin provisioned by default. Thin provisioning allows the storage unit to grow up to the size allocated but doesn't reserve the space in advance. Space is allocated dynamically from available free space as needed. This allows you to realize greater storage efficiency by *over provisioning* your available space. For example, suppose you have 1 TB of free space and you need to create four 1 TB storage units. Instead of immediately adding 3 TB of additional storage capacity to your system, you can create the storage units, monitor space utilization, and increase your storage capacity as the the storage units consume actual space. Learn more about thin provisioning.

Steps

1. In System Manager, select **Storage**; then select + Add .
2. Enter a name for the new storage unit.
3. Enter the number of units you want to create.

   If you create more than one storage unit, each unit is created with the same capacity, host operating system, and host mapping.

   To optimize workload balancing across the storage availability zone, create an even number of storage units.

4. Enter the storage unit capacity; then select the host operating system.

If you are creating more than one storage unit, each unit is created with the same capacity. Multiply the number of storage units you are creating by the desired capacity to ensure you have enough usable space. If you don't have enough free space and you chose to over provision, monitor utilization closely to avoid running out of space and losing data.

5. Accept the auto-selected **host mapping** or select a different host group for the storage unit to be mapped to.

   Host mapping refers to the host group that the new storage unit will be mapped to. If there is a pre-existing host group for the type of host you selected for your new storage unit, the pre-existing host group is auto-selected for your host mapping. You can accept the host group that is auto-selected or you can select a different host group.

   If there is no pre-existing host group for hosts running on the operating system you specified, ONTAP creates a new host group automatically .

6. If you want to do any of the following, select **More Options** and complete the required steps.

Option	Steps
Change the default Quality of Service (QoS) policy If the default QoS policy has not previously been set on the storage virtual machine (VM) on which the storage unit is being created, this option is not available.	a. Under **Storage and optimization**, next to **Quality of service (QoS)**, select . b. Select an existing QoS policy.
Create a new QoS policy	a. Under **Storage and optimization**, next to **Quality of service (QoS)**, select . b. Select **Define new policy**. c. Enter a name for the new QoS policy. d. Set a QoS limit, a QoS guarantee, or both. i. Optionally, under **Limit**, enter a maximum throughput limit, a maximum IOPS limit, or both. Setting a maximum throughput and IOPS for a storage unit restricts its impact on system resources so that it does not degrade the performance of critical workloads. ii. Optionally, under **Guarantee**, enter a minimum throughput, a minimum IOPS, or both. Setting a minimum throughput and IOPS for a storage unit guarantees that it meets minimum performance targets regardless of demand by competing workloads. e. Select **Add**.

Option	Steps
Change the default performance service level.	a. Under **Storage and optimization**, next to the **Performance service level**, select . b. Select **Performance**. ASA r2 systems offer two performance levels. The default performance level is **Extreme**, which is the highest available level. You can lower the level to **Performance**.
Modify the default snapshot reserve and enable automatic snapshot deletion.	a. Under **Snapshot reserve %**, enter the numeric value for the percentage of the storage unit space you want to allocate to snapshots. b. Select **Automatically delete older snapshots**.
Add a new SCSI host	a. Under **Host information**, select **SCSI** for the connection protocol. b. Select the host operating system. c. Under **Host Mapping**, select **New hosts**. d. Select **FC** or **iSCSI**. e. Select existing host initiators or select **Add initiator** to add a new host initiator. An example of a valid FC WWPN is "01:02:03:04:0a:0b:0c:0d". Examples of valid iSCSI initiator names are "iqn.1995-08.com.example:string" and "eui.0123456789abcdef".
Create a new SCSI host group	a. Under **Host information**, select **SCSI** for the connection protocol. b. Select the host operating system. c. Under **Host Mapping**, select **New host group**. d. Enter a name for the host group; then select the hosts to add to the group.

Option	Steps
Add a new NVMe subsystem	a. Under **Host information**, select **NVMe** for the connection protocol. b. Select the host operating system. c. Under **Host Mapping**, select **New NVMe subsystem**. d. Enter a name for the subsystem or accept the default name. e. Enter a name for the initiator. f. If you want to enable in-band authentication or Transport Layer Security (TLS), select ; then select your options. In-band authentication allows secure bidirectional and unidirectional authentication between your NVMe hosts and your ASA r2 system. TLS encrypts all data sent over the network between your NVMe/TCP hosts and your ASA r2 system. g. Select **Add initiator** to add more initiators. Format the host NQN as <nqn.yyyy-mm> followed by a fully qualified domain name. The year should be equal to or later than 1970. The total maximum length should be 223. An example of a valid NVMe initiator is nqn.2014-08.com.example:string

7. Select **Add**.

What's next?

Your storage units are created and mapped to your hosts. You can now create snapshots to protect the data on your ASA r2 system.

For more information

Learn more about how ASA r2 systems use storage virtual machines.

Add host initiators

You can add new host initiators to your ASA r2 system at any time. Initiators make the hosts eligible to access storage units and perform data operations.

Before you begin

If you want to replicate the host configuration to a destination cluster during the process of adding your host initiators, your cluster must be in a replication relationship. Optionally, you can create a replication relationship after your host is added.

Add host initiators for SCSI or NVMe hosts.

SCSI hosts

Steps

1. Select **Host**.
2. Select **SCSI**; then select + Add.
3. Enter the host name, select the host operating system and enter a host description.
4. If you want to replicate the host configuration to a destination cluster, select **Replicate host configuration**; then select the destination cluster.

   Your cluster must be in a replication relationship to replicate the host configuration.

5. Add new or existing hosts.

Add new hosts	Add existing hosts
a. Select **New hosts**. b. Select **FC** or **iSCSI**; then select the host initiators. c. Optionally, select **Configure host proximity**. Configuring host proximity enables ONTAP to identify the controller nearest to the host for data path optimization and latency reduction. This is only applicable if you have replicated data to a remote location. If you have not set up snapshot replication, you do not need to select this option. d. If you need to add new initiators, select **Add initiators**.	a. Select **Existing hosts**. b. Select the host you want to add. c. Select **Add**.

6. Select **Add**.

What's next?

Your SCSI hosts are added to your ASA r2 system and you are ready to map your hosts to your storage units.

NVMe hosts

Steps

1. Select **Host**.
2. Select **NVMe**; then select + Add.
3. Enter a name for the NVMe subsystem, select the host operating system and enter a description.
4. Select **Add initiator**.

What's next?

Your NVMe hosts are added to your ASA r2 system and you are ready to map your hosts to your storage units.

Map the storage unit to a host

After creating ASA r2 storage units and adding host initiators, map hosts to storage units to begin serving data. Storage units are mapped to hosts as part of the storage unit creation process. You can also map existing storage units to new or existing hosts at any time.

Steps

1. Select **Storage**.
2. Hover over the name of the storage unit you want to map.
3. Select ⋮; then select **Map to hosts**.
4. Select the hosts you want to map to the storage unit; then select **Map**.

What's next?

Your storage unit is mapped to your hosts and you are ready to complete the provisioning process on your hosts.

Complete host-side provisioning

After you have created your storage units, added your host initiators and mapped your storage units, there are steps you must perform on your hosts before they can read and write data on your ASA r2 system.

Steps

1. For FC and FC/NVMe, zone your FC switches by WWPN.

   Use one zone per initiator and include all target ports in each zone.

2. Discover the new storage unit.
3. Initialize the storage unit and a create file system.
4. Verify that your host can read and write data on the storage unit.

What's next?

You have completed the provisioning process and are ready to begin serving data. You can now create snapshots to protect the data on your ASA r2 system.

For more information

For more details about host-side configuration, see the ONTAP SAN host documentation for your specific host.

Clone data on ASA r2 storage systems

Data cloning creates copies of storage units and consistency groups on your ASA r2 system using ONTAP System Manager that can be used for application development, testing, backups, data migration or other administrative functions.

Clone storage units

When you clone a storage unit, you create a new storage unit on your ASA r2 system that is a point-in-time, writable copy of the storage unit you cloned.

Steps

1. In System Manager, select **Storage**.
2. Hover over the name of the storage unit you want to clone.
3. Select ⋮; then select **Clone**.
4. Accept the default name for the new storage unit that will be created as a clone or enter a new one.
5. Select the host operating system.

   A new snapshot is created for the clone by default.

6. If you want to use an existing snapshot, create a new host group, or add a new host, select **More Options**.

Option	Steps
Use an existing snapshot	a. Under **Snapshot to clone**, select **Use an existing snaphot**. b. Select the snapshot you want to use for the clone.
Create a new host group	a. Under **Host mapping**, select **New host group**. b. Enter a name for the new host group; then select the host initiators to include in the group.
Add a new host	a. Under **Host mapping**, select **New hosts**. b. Enter the a name for the new host; then select **FC** or **iSCSI**. c. Select the host initiators from the list of existing initiators or select **Add** to add new initiators for the host.

7. Select **Clone**.

What's next?

You have created a new storage unit that is identical to the storage unit you cloned. You are now ready to use the new storage unit as needed.

Clone consistency groups

When you clone a consistency group, you create a new consistency group that's identical in structure, storage units, and data to the consistency group you cloned. Use a consistency group clone to perform application testing or to migrate data. Suppose, for example, you need to migrate a production workload out of a consistency group. You can clone the consistency group to create a copy of your production workload to maintain as a backup until the migration is complete.

The clone is created from a snapshot of the consistency group being cloned. The snapshot used for the clone is taken at the point in time that the cloning process is initiated by default. You can modify the default behavior to use a pre-existing snapshot.

Storage unit mappings are copied as part of the cloning process. Snapshot polices are not copied as part of the cloning process.

You can create clones from consistency groups stored locally on your ASA r2 system or from consistency groups that have been replicated to remote locations.

Clone using local snapshot

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Hover over the consistency group you want to clone.
3. Select ⋮, then select **Clone**.
4. Enter a name for consistency group clone or accept the default name.
5. Select the host operating system.
6. If you want to dissociate the clone from the source consistency group and allocate disk space, select **Split clone**.
7. If you want to use an existing snapshot, create a new host group or add a new host for the clone, select **More Options**.

Option	Steps
Use an existing snapshot	a. Under **Snapshot to clone**, select **Use an existing snapshot**. b. Select the snapshot you want to use for the clone.
Create a new host group	a. Under **Host mapping**, select **New host group**. b. Enter a name for the new host group; then select the host initiators to include in the group.
Add a new host	a. Under **Host mapping**, select **New hosts**. b. Enter the name new host name; then select **FC** or **iSCSI**. c. Select the host initiators from the list of existing initiators or select **Add initiator** to add new initiators for the host.

8. Select **Clone**.

Clone using remote snapshot

Steps

1. In System Manager, select **Protection > Replication**.
2. Hover over the **Source** you want to clone.
3. Select ⋮, then select **Clone**.
4. Select the source cluster and storage VM; then enter a name for the new consistency group or accept the default name.
5. Select the snapshot to clone; then select **Clone**.

What's next?

You have cloned a consistency group from your remote location. The new consistency group is locally available on your ASA r2 system to use as needed.

What's next?

To protect your data, you should create snapshots of the cloned consistency group.

Split consistency group clone

When you split a consistency group clone, you dissociate the clone from the source consistency group and allocate disk space for the clone. The clone becomes a standalone consistency group that can be used independently of the source consistency group.

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Hover over the consistency group clone that you want you want to split.
3. Select **Split clone**.
4. Select **Split**.

Result

The clone is dissociated from the source consistency group and disk space is allocated for the clone.

Manage host groups

Create host groups on your ASA r2 system

On an ASA r2 system, a *host group* is the mechanism used to give hosts access to storage units. A host group refers to an igroup for SCSI hosts or to an NVMe subsystem for NVMe hosts. A host can only see the storage units that are mapped to the host groups to which it belongs. When a host group is mapped to a storage unit, the hosts that are members of the group, are then able to mount (create directories and file structures on) the storage unit.

Host groups are automatically or manually created when you create your storage units. You can optionally use the following steps to create host groups before or after storage unit creation.

Steps

1. From System Manager, select **Host**.
2. Select the hosts you want to add to the host group.

   After you select the first host, the option to add to a host group appears above the list of hosts.

3. Select **Add to host group**.
4. Search for and select the host group to which you want to add the host.

What's next?

You have created a host group and you can now map it to a storage unit.

Delete a host group on your ASA r2 system

On an ASA r2 system, a host group is the mechanism used to give hosts access to storage units. A host group refers to an igroup for SCSI hosts or to an NVMe subsystem for NVMe hosts. A host can only see the storage units that are mapped to the host groups to which it belongs. You might want to delete a host group if you no longer want the hosts in the group to have access to the storage units that are mapped to the group.

Steps

1. In System Manager, select **Storage**.
2. Under **Host mapping** select the host group you want to delete.
3. Select **Mapped storage**.
4. Select **More**; then select **Delete**.
5. Select to verify that you want to continue; then select **Delete**.

What's next?

The host group is deleted. The hosts that were in the group no longer have access to the storage units that were mapped to the host group.

Manage storage units

Modify storage units on ASA r2 storage systems

To optimize performance on your ASA r2 system, you might need to modify your storage units to increase their capacity, update QoS policies or to change the hosts that are mapped to the units. For example, if a new, critical application workload is added to an existing storage unit, you might need to change the Quality of Service (QoS) policy applied to the storage unit to support the performance level needed for the new application.

Increase capacity

Increase the size of a storage unit before it reaches full capacity to prevent a loss of data access that can occur if the storage unit runs out of writeable space. The capacity of a storage unit can be increased to 128 TB which is the maximum size allowed by ONTAP.

Modify host mappings

Modify the hosts that are mapped to a storage unit to assist in balancing workloads or reconfiguring system resources.

Modify QoS policy

Quality of service (QoS) policies guarantee that the performance of critical workloads is not degraded by competing workloads. You can use QoS policies to set a QoS throughput *limit* and a QoS throughput *guarantee*.

- QoS throughput limit

  The QoS throughput *limit* restricts the impact of a workload on system resources by limiting the throughput for the workload to a maximum number of IOPS or MBps, or IOPS and MBps.

- QoS throughput guarantee

  The QoS throughput *guarantee* ensures that critical workloads meet minimum throughput targets, regardless of demand by competing workloads, by guaranteeing that the throughput for the critical workload does not fall below a minimum number of IOPS or MBps, or IOPS and MBps.

Steps

1. In System Manager, select **Storage**.
2. Hover over the name of the storage unit you want to edit.
3. Select ⋮; then select **Edit**.
4. Update the storage unit parameters as needed to increase capacity, change the QoS policy, and update the host mapping.

What's next?

If you have increased the size of your storage unit, you must rescan the storage unit on the host for the host to recognize the change in size.

Move storage units on ASA r2 storage systems

If a storage availability zone is running low on space, you can move storage units to another storage availability zone to balance the storage utilization across the cluster.

You can move a storage unit while the storage unit is online and serving data. The move operation is non-disruptive.

Before you begin

- You must be running ONTAP 9.16.1 or later.
- Your cluster must consist of four or more nodes.

Steps

1. In System Manager, select **Storage**; then select the storage unit you want to move.
2. Select ⋮; then select **Move**.
3. Select the storage availability zone you want to move the storage unit to; then select **Move**.

Delete storage units on ASA r2 storage systems

Delete a storage unit if you no longer need to maintain the data contained in the unit. Deleting storage units that are no longer needed can help you free space needed for other host applications.

Before you begin

If the storage unit you want to delete is in a consistency group that is in replication relationship, you must remove the storage unit from the consistency group before you delete it.

Steps

1. In System Manager, select **Storage**.
2. Hover over the name of the storage unit you want to delete.

3. Select ⋮; then select **Delete**.
4. Acknowledge that the deletion cannot be undone.
5. Select **Delete**.

What's next?

You can use the space freed from the deleted storage unit to increase the size of storage units that need additional capacity.

Migrate storage VMs

Migrate a storage VM from an ASA cluster to an ASA r2 cluster

Beginning with ONTAP 9.18.1, you can non-disruptively migrate a storage virtual machine (VM) from any ASA cluster to any ASA r2 cluster. Migrating from an ASA cluster to an ASA r2 cluster allows you to adopt the simplified and streamlined architecture of ASA r2 systems for SAN-only environments.

Storage VM migration between ASA and ASA r2 storage systems is supported as follows:

From any of the following ASA systems:	To any of the following ASA r2 systems:
• ASA C800 • ASA C400 • ASA C250 • ASA A900 • ASA A800 • ASA A400 • ASA A250 • ASA A150 • ASA AFF A800 • ASA AFF A700 • ASA AFF A400 • ASA AFF A250 • ASA AFF A220	• ASA A1K • ASA C30 • ASA A90 • ASA A70 • ASA A50 • ASA A30 • ASA A20

For the most current list of ASA and ASA r2 systems, see NetApp Hardware Universe. ASA r2 systems are listed in NetApp Hardware Universe as "ASA A-Series/C-Series (New)".

You can migrate a storage VM to an ASA r2 cluster from an ASA cluster only. Migration from any other type of ONTAP system is not supported.

Before you begin

All nodes in the ASA r2 cluster and the ASA cluster must be running ONTAP 9.18.1 or later. The ONTAP 9.18.1 patch versions on the cluster nodes can vary.

Step 1: Verify the status of the ASA storage VM

Before you migrate a storage VM from an ASA system, there should be no NVMe namespaces or vVols present and each volume in the storage VM should contain only one LUN. Migration of NVMe namespaces and vVols is not supported. The architecture of ASA r2 systems requires that volumes contain a single LUN.

Steps

1. Verify that no NVMe namespaces are present in the storage VM:

   ```
   vserver nvme namespace show -vserver <storage_VM>
   ```

   If entries are displayed, the NVMe objects must be converted to LUNs or removed. See the `vserver nvme namespace delete` and the `vserver nvme subsystem delete` commands in the ONTAP command reference for more information.

2. Verify that there are no vVols present in the storage VM:

   ```
   lun show -verser <storage_VM> -class protocol-endpoint,vvol
   ```

   If any vVols are present, they should be copied to another storage VM and then deleted from the storage VM to be migrated. See the `lun copy` and `lun delete` commands in the ONTAP command reference for more information.

3. Verify that each volume in the storage VM contains a single LUN:

   ```
   lun show -verser <storage_VM>
   ```

   If a volume contains more than one LUN, use the `volume create` and `lun move` commands to create a 1:1 volume-to-LUN ratio. See the ONTAP command reference for more information.

What's next?

You are ready to create a cluster peer relationship between your ASA and ASA r2 clusters.

Step 2: Create a cluster peer relationship between your ASA and ASA r2 clusters

Before you can migrate a storage VM from an ASA cluster to an ASA r2 cluster, you need to create a peer relationship. A peer relationship defines network connections that enable ONTAP clusters and storage VMs to exchange data securely.

Before you begin

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

1. On the ASA r2 cluster, create a peer relationship with the ASA cluster and generate a passphrase:

```
cluster peer create -peer-addrs <ASA_cluster_LIF_IPs> -generate
-passphrase
```

The following example creates a cluster peer relationship between cluster 1 and cluster 2 and creates a system-generated passphrase:

```
cluster1::> cluster peer create -peer-addrs 10.98.191.193 -generate
-passphrase
Passphrase: UCa+6lRVICXeL/gq1WrK7ShR
            Peer Cluster Name: cluster2
            Initial Allowed Vserver Peers: -
            Expiration Time: 6/7/2017 09:16:10 +5:30
            Intercluster LIF IP: 10.140.106.185
Warning: make a note of the passphrase - it cannot be displayed again.
```

2. Copy the generated passphrase.
3. On the ASA cluster, create a peer relationship with the ASA r2 cluster:

```
cluster peer create -peer-addrs <ASA_r2_LIF_IPs>
```

4. Enter the passphrase generated on the ASA r2 cluster.
5. Verify that the cluster peer relationship is created:

```
cluster peer show
```

The following example displays the expected output for successfully peered clusters.

```
cluster1::> cluster peer show

Peer Cluster Name   Cluster Serial Number   Availability
Authentication
-----------------   ---------------------   --------------
--------------
cluster2             1-80-123456             Available      ok
```

Result

The ASA and ASA r2 clusters are peered and storage VM data can be securely transferred.

What's next?

You are ready to prepare your ASA storage VM for migration.

Step 3: Prepare for storage VM migration from an ASA to an ASA r2 cluster

Before you migrate a storage virtual machine (VM) from an ASA cluster to an ASA r2 cluster, you must run a migration pre-check and fix any required issues. You cannot perform the migration until the pre-check passes successfully.

Step

1. From your ASA r2 cluster, execute the migration pre-check:

   ```
   vserver migrate start -vserver <storage_VM> -source-cluster
   <asa_cluster> -check-only true
   ```

   If you need to fix any issues to prepare your ASA cluster for migration, the issue and the corrective action is displayed. Fix the issue and repeat the pre-check until it completes successfully.

What's next?

You are ready to migrate your storage VM from your ASA cluster to an ASA r2 cluster.

Step 4: Migrate an ASA storage VM to an ASA r2 cluster

After you have prepared your ASA cluster and created the necessary cluster peer relationship with the ASA r2 cluster, you can begin the storage VM migration.

When performing a storage VM migration, it is a best practice to leave 30% CPU headroom on both the ASA cluster and the ASA r2 cluster to enable the CPU workload to execute.

About this task

After the storage VM migration, clients are automatically cut over to the ASA r2 cluster and the storage VM on the ASA cluster is automatically removed. Automatic cutover and automatic storage VM removal are enabled by default. You can optionally disable them both and perform the cutover and the storage VM removal manually.

Before you begin

- The ASA r2 cluster must have enough free space to hold the migrated storage VM.
- If the ASA storage VM contains encrypted volumes, the onboard key manager or the external key manager on the ASA r2 system must be configured at the cluster level.
- The following operations cannot be running on the source ASA cluster:
  - Failover operations
  - WAFLIRON
  - Fingerprint
  - Volume move, rehost, clone, create, convert or analytics

Steps

1. From the ASA r2 cluster, start the storage VM migration:

```
vserver migrate start -vserver <storage_VM_name> -source-cluster
<ASA_cluster>
```

To disable automatic cutover, use the `-auto-cutover false` parameter. To disable the automatic removal of the ASA storage VM, use the `-auto-source-cleanup false` parameter.

2. Monitor the status of the migration

```
vserver migrate show -vserver <storage_VM_name>
```

When the migration is complete, the **status** displays as **migration-complete**.

> If you need to pause or cancel the migration before automatic cutover begins, use the `vserver migrate pause` and the `vserver migrate abort` commands. You must pause the migration before cancelling it. You cannot cancel the migration after cutover starts.

Result

The storage VM is migrated from the ASA cluster to the ASA r2 cluster. The storage VM's name and UUID, the data LIF name, IP address, and object names, such as the volume name, remain unchanged. The UUID of the migrated objects in the storage VM are updated.

What's next?

If you disabled automatic cutover and automatic storage VM removal, manually cut over your ASA clients to your ASA r2 cluster and remove the storage VM from the ASA cluster.

Cutover clients and clean up source storage VM after migration to an ASA r2 system

After a storage virtual machine (VM) is migrated from an ASA cluster to an ASA r2 cluster, by default, clients are automatically cut over to the ASA r2 cluster and the storage VM on the ASA cluster is automatically removed. If you chose to disable automatic cutover and removal of the ASA storage VM during the migration, you need to perform these steps manually after the migration is complete.

Manually cut over clients to an ASA r2 system after a storage VM migration

If you disable automatic client cutover during the migration of a storage VM from an ASA cluster to an ASA r2 cluster, after the migration is successfully complete, perform the cutover manually so the ASA r2 storage VM can serve data to clients.

Steps

1. On the ASA r2 cluster, manually execute client cutover:

```
vserver migrate cutover -vserver <storage_VM_name>
```

2. Verify that the cutover operation is complete:

```
vserver migrate show
```

Result

Data is being served to your clients from the storage VM on your ASA r2 cluster.

What's next?

You are now ready to remove the storage VM from the source ASA cluster.

Manually remove an ASA storage VM after migration to an ASA r2 cluster

If you disable automatic source cleanup during the migration of a storage VM from an ASA cluster to an ASA r2 cluster, after the migration is complete, remove the storage VM from the ASA cluster to free the storage space.

Before you begin

Your clients should be serving data from the ASA r2 cluster.

Steps

1. From the ASA cluster, verify that the status of the ASA storage VM is **Ready for source cleanup**:

```
vserver migrate show
```

2. Remove the ASA storage VM:

```
vserver migrate source-cleanup -vserver <storage_VM_name>
```

Result

The storage VM on your ASA cluster is removed.

ASA r2 storage limits

For optimal performance, configuration and support, you should be aware ofASA r2 storage limits.

For a complete list of the most current ASA r2 storage limits, see NetApp Hardware Universe.

ASA r2 systems support the following storage limits:

	Maximum per HA pair	Maximum per cluster
Consistency groups	256	256
Enterprise applications	100	350
Nodes	2	12

	Maximum per HA pair	Maximum per cluster
Replication groups	50	50
Storage availability zone size	2 PB	2 PB
Storage units	10,000	30,000
Storage unit size	128 TB	128 TB
Storage units per consistency group	256	256
Child consistency groups per parent consistency group	64	64
Storage virtual machines	• 256 (ONTAP 9.18.1 and later) • 32 (ONTAP 9.17.1 and earlier)	• 256 (ONTAP 9.18.1 and later) • 32 (ONTAP 9.17.1 and earlier)
Virtual machines	800	1200

Limits for SnapMirror asynchronous relationships

The following limits apply to storage units and consistency groups in a SnapMirror asynchronous replication relationship. For a complete list of the most current ASA r2 storage limits, NetApp Hardware Universe.

Limit maximum	Per HA pair	Per cluster
Consistency groups	250	750
Storage units	4,000	6,000

Limits for SnapMirror active sync relationship

The following limits apply to storage units and consistency groups in a SnapMirror active sync replication relationship. SnapMirror active sync is supported beginning with ONTAP 9.17.1 on two-node clusters only. Beginning with ONTAP 9.18.1, SnapMirror active sync is supported on four-node clusters.

For a complete list of the most current ASA r2 storage limits, NetApp Hardware Universe.

Limit maximum	Per HA pair
Consistency groups	50
Storage units	400

Copyright information

Trademark information