VMware Site Recovery Manager with ONTAP

Enterprise applications

NetApp
February 10, 2026

Table of Contents

VMware Site Recovery Manager with ONTAP

VMware Live Site Recovery with ONTAP

ONTAP has been a leading storage solution for VMware vSphere and, more recently, Cloud Foundation, since ESX was introduced into modern datacenters more than two decades ago. NetApp continues to introduce innovative systems, such as the latest generation of the ASA A-series, along with features like SnapMirror active sync. These advancements simplify management, enhance resiliency, and lower the total cost of ownership (TCO) for your IT infrastructure.

This document introduces the ONTAP solution for VMware Live Site Recovery (VLSR), formerly known as Site Recovery Manager (SRM), VMware's industry-leading disaster recovery (DR) software, including the latest product information and best practices to streamline deployment, reduce risk, and simplify ongoing management.

This documentation replaces the previously published technical report *TR-4900: VMware Site Recovery Manager with ONTAP*

Best practices supplement other documents such as guides and compatibility tools. They are developed based on lab testing and extensive field experience by NetApp engineers and customers. In some cases, recommended best practices might not be the right fit for your environment; however, they are generally the simplest solutions that meet the needs of the most customers.

This document is focused on capabilities in recent releases of ONTAP 9 when used in conjunction with ONTAP tools for VMware vSphere 10.4 (which includes the NetApp Storage Replication Adapter [SRA] and VASA Provider [VP]), as well as VMware Live Site Recovery 9.

Why use ONTAP with VLSR or SRM?

NetApp data management platforms powered by ONTAP are some of the most widely adopted storage solutions for VLSR. The reasons are plentiful: A secure, high-performance, unified protocol (NAS and SAN together) data management platform that provides industry-defining storage efficiency, multitenancy, quality of service controls, data protection with space-efficient snapshots, and replication with SnapMirror. All leveraging native hybrid multi-cloud integration for the protection of VMware workloads and a plethora of automation and orchestration tools at your fingertips.

When you use SnapMirror for array-based replication, you take advantage of one of ONTAP's most proven and mature technologies. SnapMirror gives you the advantage of secure and highly efficient data transfers, copying only changed file system blocks, not entire VMs or datastores. Even those blocks take advantage of space savings, such as deduplication, compression, and compaction. Modern ONTAP systems now use version-independent SnapMirror, allowing you flexibility in selecting your source and destination clusters. SnapMirror has truly become one of the most powerful tools available for disaster recovery.

Whether you are using traditional NFS, iSCSI, or Fibre Channel- attached datastores (now with support for vVols datastores), VLSR provides a robust first-party offering that leverages the best of ONTAP capabilities for disaster recovery or datacenter migration planning and orchestration.

How VLSR leverages ONTAP 9

VLSR leverages the advanced data management technologies of ONTAP systems by integrating with ONTAP tools for VMware vSphere, a virtual appliance that includes three primary components:

- The ONTAP tools vCenter plug-in, formerly known as Virtual Storage Console (VSC), simplifies storage management and efficiency features, enhances availability, and reduces storage costs and operational overhead, whether you are using SAN or NAS. It uses best practices for provisioning datastores and optimizes ESXi host settings for NFS and block storage environments. For all these benefits, NetApp recommends this plug-in when using vSphere with systems running ONTAP.
- The ONTAP tools VASA Provider supports the VMware vStorage APIs for Storage Awareness (VASA) framework. VASA Provider connects vCenter Server with ONTAP to aid in provisioning and monitoring VM storage. This enabled VMware Virtual Volumes (vVols) support and the management of VM storage policies and individual VM vVols performance. It also provides alarms for monitoring capacity and compliance with the profiles.
- The SRA is used together with VLSR to manage the replication of VM data between production and disaster recovery sites for traditional VMFS and NFS datastores and also for the nondisruptive testing of DR replicas. It helps automate the tasks of discovery, recovery, and reprotection. It includes both an SRA server appliance and SRA adapters for the Windows SRM server and the VLSR appliance.

After you have installed and configured the SRA adapters on the VLSR server for protecting non-vVols datastores, you can begin the task of configuring your vSphere environment for disaster recovery.

The SRA delivers a command-and-control interface for the VLSR server to manage the ONTAP FlexVol volumes that contain your VMware Virtual Machines (VMs), as well as the SnapMirror replication protecting them.

VLSR can test your DR plan nondisruptively using NetApp's proprietary FlexClone technology to make nearly instantaneous clones of your protected datastores at your DR site. VLSR creates a sandbox to safely test so that your organization and your customers are protected in the event of a true disaster, giving you confidence in your organization's ability to execute a failover during a disaster.

In the event of a true disaster or even a planned migration, VLSR allows you to send any last-minute changes to the dataset via a final SnapMirror update (if you choose to do so). It then breaks the mirror and mounts the datastore to your DR hosts. At that point, your VMs can be automatically powered up in any order according to your pre-planned strategy.

While ONTAP systems will allow you to pair SVMs in the same cluster for SnapMirror replication, that scenario is not tested and certified with VLSR. Therefore, it is recommended to only use SVMs from different clusters when using VLSR.

VLSR with ONTAP and other use cases: hybrid cloud and migration

Integrating your VLSR deployment with ONTAP advanced data management capabilities allows for vastly improved scale and performance when compared with local storage options. But more than that, it brings the flexibility of the hybrid cloud. The hybrid cloud enables you to save money by tiering unused data blocks from your high-performance array to your preferred hyperscaler using FabricPool, which could be an on-premises S3 store such as NetApp StorageGRID. You can also use SnapMirror for edge-based systems with software-defined ONTAP Select or cloud-based DR using NetApp Storage on Equinix Metal, or other hosted ONTAP services.

You could then perform test failover inside a cloud service provider's datacenter with near-zero storage footprint thanks to FlexClone. Protecting your organization can now cost less than ever before.

VLSR can also be used to execute planned migrations by leveraging SnapMirror to efficiently transfer your VMs from one datacenter to another or even within the same datacenter, whether your own, or via any number of NetApp partner service providers.

Deployment best practices

The following sections outline the deployment best practices with ONTAP and VMware SRM.

Use the latest version of ONTAP tools 10

ONTAP tools 10 provides significant improvements over previous versions, including the following:

- 8x faster test failover*
- 2x faster cleanup and reprotect*
- 32% faster failover*
- Greater scale
- Native support for shared site layouts

*These improvements are based on internal testing and may vary based on your environment.

SVM layout and segmentation for SMT

With ONTAP, the concept of the storage virtual machine (SVM) provides strict segmentation in secure multitenant environments. SVM users on one SVM cannot access or manage resources from another. In this way, you can leverage ONTAP technology by creating separate SVMs for different business units who manage their own SRM workflows on the same cluster for greater overall storage efficiency.

Consider managing ONTAP using SVM-scoped accounts and SVM management LIFs to not only improve security controls, but also improve performance. Performance is inherently greater when using SVM-scoped connections because the SRA is not required to process all the resources in an entire cluster, including physical resources. Instead, it only needs to understand the logical assets that are abstracted to the particular SVM.

Best practices for managing ONTAP 9 systems

As previously mentioned, you can manage ONTAP clusters using either cluster or SVM scoped credentials and management LIFs. For optimum performance, you may want to consider using SVM- scoped credentials whenever you aren't using vVols. However, in doing so, you should be aware of some requirements, and that you do lose some functionality.

- The default vsadmin SVM account does not have the required access level to perform ONTAP tools tasks. Therefore, you need to create a new SVM account. Configure ONTAP user roles and privileges using the included JSON file. This can be used for SVM or cluster scoped accounts.
- Because the vCenter UI plugin, VASA Provider, and SRA server are all fully integrated microservices, you must add storage to the SRA adapter in SRM the same way you add storage in the vCenter UI for ONTAP tools. Otherwise, the SRA server might not recognize the requests being sent from SRM via the SRA adapter.
- NFS path checking is not performed when using SVM-scoped credentials unless you first onboard clusters in ONTAP tools manager and associate them with vCenters. This is because the physical location is

logically abstracted from the SVM. This is not a cause for concern though, as modern ONTAP systems no longer suffer any noticeable performance decline when using indirect paths.

- Aggregate space savings due to storage efficiency might not be reported.
- Where supported, load-sharing mirrors cannot be updated.
- EMS logging might not be performed on ONTAP systems managed with SVM scoped credentials.

Operational best practices

The following sections outline the operational best practices for VMware SRM and ONTAP storage.

Datastores and protocols

- If possible, always use ONTAP tools to provision datastores and volumes. This makes sure that volumes, junction paths, LUNs, igroups, export policies, and other settings are configured in a compatible manner.
- SRM supports iSCSI, Fibre Channel, and NFS version 3 with ONTAP 9 when using array-based replication through SRA. SRM does not support array-based replication for NFS version 4.1 with either traditional or vVols datastores.
- To confirm connectivity, always verify that you can mount and unmount a new test datastore at the DR site from the destination ONTAP cluster. Test each protocol you intend to use for datastore connectivity. A best practice is to use ONTAP tools to create your test datastore, since it is doing all the datastore automation as directed by SRM.
- SAN protocols should be homogeneous for each site. You can mix NFS and SAN, but the SAN protocols should not be mixed within a site. For example, you can use FCP in site A, and iSCSI in site B. You should not use both FCP and iSCSI at site A.
- Previous guides advised creating LIF to data locality. That is to say, always mount a datastore using a LIF located on the node that physically owns the volume. While that is still the best practice, it is no longer a requirement in modern versions of ONTAP 9. Whenever possible, and if given cluster-scoped credentials, ONTAP tools will still choose to load balance across LIFs local to the data, but it is not a requirement for high availability or performance.
- ONTAP 9 can be configured to automatically remove snapshots to preserve uptime in the event of an out-of-space condition when autosize is not able to supply sufficient emergency capacity. The default setting for this capability does not automatically delete the snapshots that are created by SnapMirror. If SnapMirror snapshots are deleted, then the NetApp SRA cannot reverse and resynchronize replication for the affected volume. To prevent ONTAP from deleting SnapMirror snapshots, configure the snapshot autodelete capability to 'try'.

```
snap autodelete modify -volume -commitment try
```

- Volume autosize should be set to `grow` for volumes containing SAN datastores and `grow_shrink` for NFS datastores. Learn more about this topic at Configure volumes to automatically grow and shrink their size.
- SRM performs best when the number of datastores and thus protection groups is minimized in your recovery plans. Therefore you should consider optimizing for VM density in SRM-protected environments where RTO is of key importance.
- Use Distributed Resource Scheduler (DRS) to help balance the load on your protected and recovery ESXi clusters. Remember that if you plan to failback, when you run a reprotect the previously protected clusters

will become the new recovery clusters. DRS will help balance placement going in both directions.

- Where possible, avoid using IP customization with SRM as this can increase your RTO.

About array pairs

An array manager is created for each array pair. With SRM and ONTAP tools, each array pairing is done with the scope of an SVM, even if you are using cluster credentials. This allows you to segment DR workflows between tenants based on which SVMs they have been assigned to manage. You can create multiple array managers for a given cluster, and they can be asymmetric. You can fan out or fan in between different ONTAP 9 clusters. For example, you can have SVM-A and SVM-B on Cluster-1 replicating to SVM-C on Cluster-2, SVM-D on Cluster-3, or vice-versa.

When configuring array pairs in SRM, you should always add them in SRM the same way as you added them to ONTAP Tools, meaning, they must use the same username, password, and management LIF. This requirement ensures that SRA communicates properly with the array. The following screenshot illustrates how a cluster might appear in ONTAP Tools and how it might be added to an array manager.

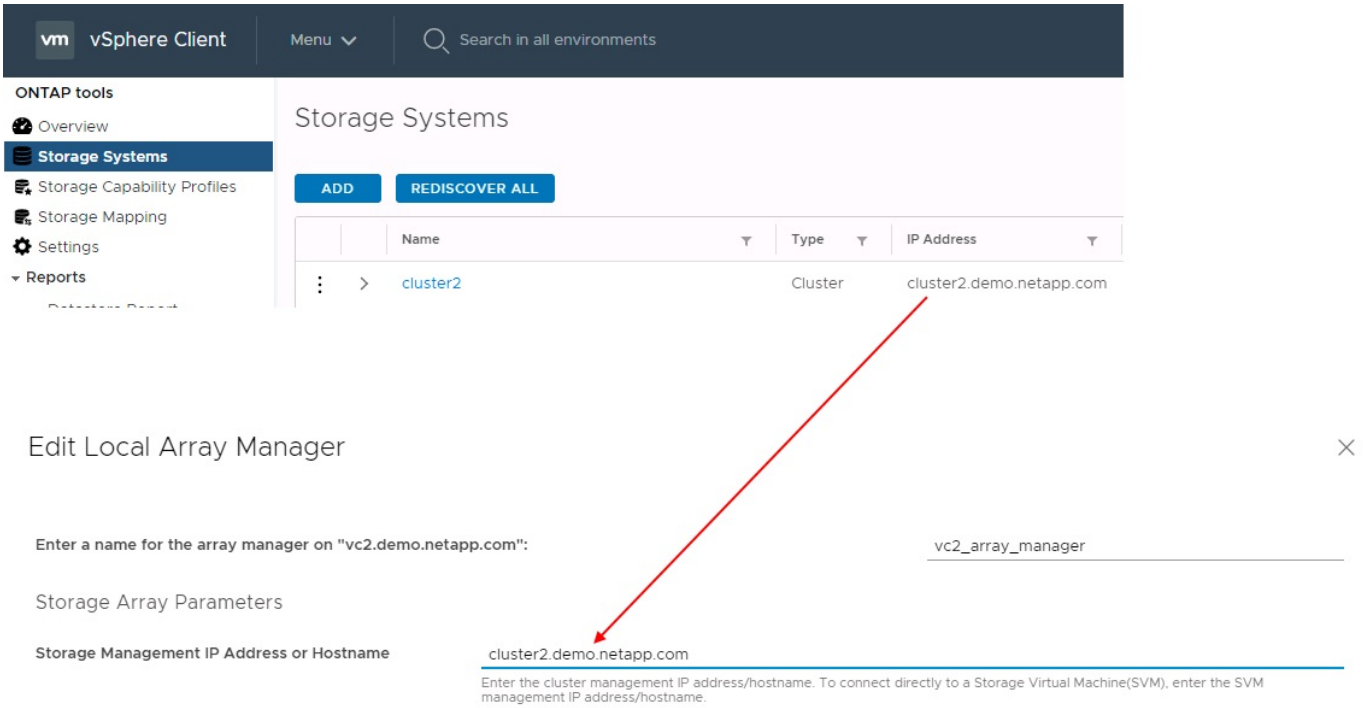

About replication groups

Replication groups contain logical collections of virtual machines that are recovered together. Because ONTAP SnapMirror replication occurs at the volume level, all VMs in a volume are in the same replication group.

There are several factors to consider with replication groups and how you distribute VMs across FlexVol volumes. Grouping similar VMs in the same volume can increase storage efficiency with older ONTAP systems that lack aggregate-level deduplication, but grouping increases the size of the volume and reduces volume I/O concurrency. The best balance of performance and storage efficiency can be achieved in modern ONTAP systems by distributing VMs across FlexVol volumes in the same aggregate, thereby leveraging aggregate-level deduplication and gaining greater I/O parallelization across multiple volumes. You can recover VMs in the volumes together because a protection group (discussed below) can contain multiple replication groups. The downside to this layout is that blocks might be transmitted over the wire multiple times because SnapMirror doesn't take aggregate deduplication into account.

One final consideration for replication groups is that each one is by its nature a logical consistency group (not to be confused with SRM consistency groups). This is because all VMs in the volume are transferred together using the same snapshot. So if you have VMs that must be consistent with each other, consider storing them in the same FlexVol.

About protection groups

Protection groups define VMs and datastores in groups that are recovered together from the protected site. The protected site is where the VMs that are configured in a protection group exist during normal steady-state operations. It is important to note that even though SRM might display multiple array managers for a protection group, a protection group cannot span multiple array managers. For this reason, you should not span VM files across datastores on different SVMs.

About recovery plans

Recovery plans define which protection groups are recovered in the same process. Multiple protection groups can be configured in the same recovery plan. Also, to enable more options for the execution of recovery plans, a single protection group can be included in multiple recovery plans.

Recovery plans allow SRM administrators to define recovery workflows by assigning VMs to a priority group from 1 (highest) to 5 (lowest), with 3 (medium) being the default. Within a priority group, VMs can be configured for dependencies.

For example, your company could have a tier-1 business-critical application that relies on a Microsoft SQL server for its database. So, you decide to place your VMs in priority group 1. Within priority group 1, you begin planning the order to bring up services. You probably want your Microsoft Windows domain controller to boot before your Microsoft SQL server, which would need to be online before your application server, and so on. You would add all these VMs to the priority group and then set the dependencies because dependencies only apply within a given priority group.

NetApp strongly recommends working with your application teams to understand the order of operations required in a failover scenario and to construct your recovery plans accordingly.

Test failover

As a best practice, always perform a test failover whenever a change is made to the configuration of protected VM storage. This ensures that, in the event of a disaster, you can trust that Site Recovery Manager can restore services within the expected RTO target.

NetApp also recommends confirming in-guest application functionality occasionally, especially after reconfiguring VM storage.

When a test recovery operation is performed, a private test bubble network is created on the ESXi host for the VMs. However, this network is not automatically connected to any physical network adapters and therefore does not provide connectivity between the ESXi hosts. To allow communication among VMs that are running on different ESXi hosts during DR testing, a physical private network is created between the ESXi hosts at the DR site. To verify that the test network is private, the test bubble network can be separated physically or by using VLANs or VLAN tagging. This network must be segregated from the production network because as the VMs are recovered, they cannot be placed on the production network with IP addresses that could conflict with actual production systems. When a recovery plan is created in SRM, the test network that was created can be selected as the private network to connect the VMs to during the test.

After the test has been validated and is no longer required, perform a cleanup operation. Running cleanup returns the protected VMs to their initial state and resets the recovery plan to the Ready state.

Failover considerations

There are several other considerations when it comes to failing over a site in addition to the order of operations mentioned in this guide.

One issue you might have to contend with is networking differences between sites. Some environments might be able to use the same network IP addresses at both the primary site and the DR site. This ability is referred to as a stretched virtual LAN (VLAN) or stretched network setup. Other environments might have a requirement to use different network IP addresses (for example, in different VLANs) at the primary site relative to the DR site.

VMware offers several ways to solve this problem. For one, network virtualization technologies like VMware NSX-T Data Center abstract the entire networking stack from layers 2 through 7 from the operating environment, allowing for more portable solutions. Learn more about NSX-T options with SRM.

SRM also gives you the ability to change the network configuration of a VM as it is recovered. This reconfiguration includes settings such as IP addresses, gateway addresses, and DNS server settings. Different network settings, which are applied to individual VMs as they are recovered, can be specified in the property's settings of a VM in the recovery plan.

To configure SRM to apply different network settings to multiple VMs without having to edit the properties of each one in the recovery plan, VMware provides a tool called the dr-ip-customizer. Learn how to use this utility, refer to VMware's documentation.

Reprotect

After a recovery, the recovery site becomes the new production site. Because the recovery operation broke the SnapMirror replication, the new production site is not protected from any future disaster. A best practice is to protect the new production site to another site immediately after a recovery. If the original production site is operational, the VMware administrator can use the original production site as a new recovery site to protect the new production site, effectively reversing the direction of protection. Reprotection is available only in non-catastrophic failures. Therefore, the original vCenter Servers, ESXi servers, SRM servers, and corresponding databases must be eventually recoverable. If they are not available, a new protection group and a new recovery plan must be created.

Failback

A failback operation is fundamentally a failover in a different direction than before. As a best practice, you verify that the original site is back to acceptable levels of functionality before attempting to failback, or, in other words, failover to the original site. If the original site is still compromised, you should delay failback until the failure is sufficiently remediated.

Another failback best practice is to always perform a test failover after completing reprotect and before doing your final failback. This verifies that the systems in place at the original site can complete the operation.

Reprotecting the original site

After failback, you should confirm with all stakeholders that their services have been returned to normal before running reprotect again,

Running reprotect after failback essentially puts the environment back in the state it was in at the beginning, with SnapMirror replication again running from the production site to the recovery site.

Replication topologies

In ONTAP 9, the physical components of a cluster are visible to cluster administrators, but they are not directly visible to the applications and hosts that use the cluster. The physical components provide a pool of shared resources from which the logical cluster resources are constructed. Applications and hosts access data only through SVMs that contain volumes and LIFs.

Each NetApp SVM is treated as a unique array in Site Recovery Manager. VLSR supports certain array-to-array (or SVM-to-SVM) replication layouts.

A single VM cannot own data—Virtual Machine Disk (VMDK) or RDM—on more than one VLSR array for the following reasons:

- VLSR sees only the SVM, not an individual physical controller.
- An SVM can control LUNs and volumes that span multiple nodes in a cluster.

Best Practice
To determine supportability, keep this rule in mind: to protect a VM by using VLSR and the NetApp SRA, all parts of the VM must exist on only one SVM. This rule applies at both the protected site and the recovery site.

Supported SnapMirror layouts

The following figures show the SnapMirror relationship layout scenarios that VLSR and SRA support. Each VM in the replicated volumes owns data on only one VLSR array (SVM) at each site.

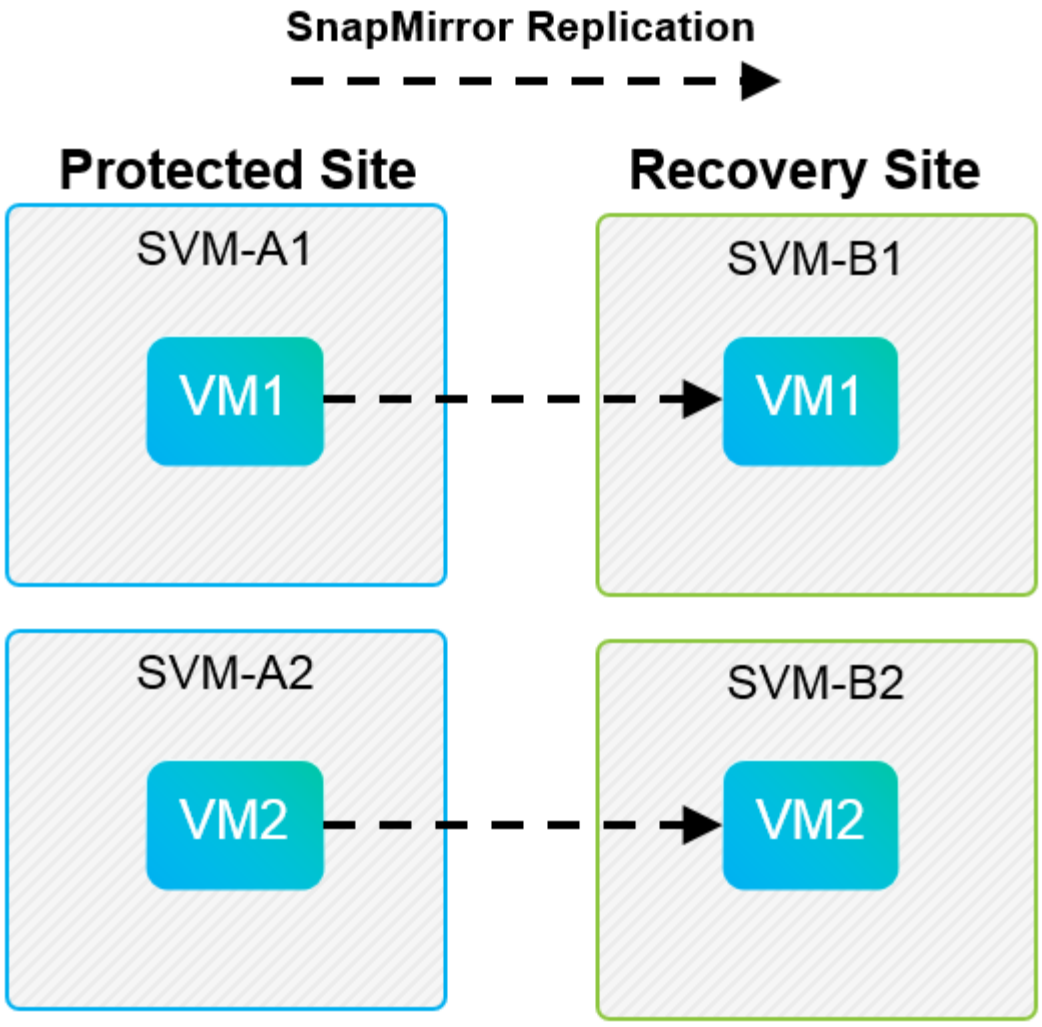
SnapMirror Replication
Protected Site
Recovery Site
SVM-A1
VM1
SVM-B1
VM1
SVM-A2
VM2
SVM-B2
VM2

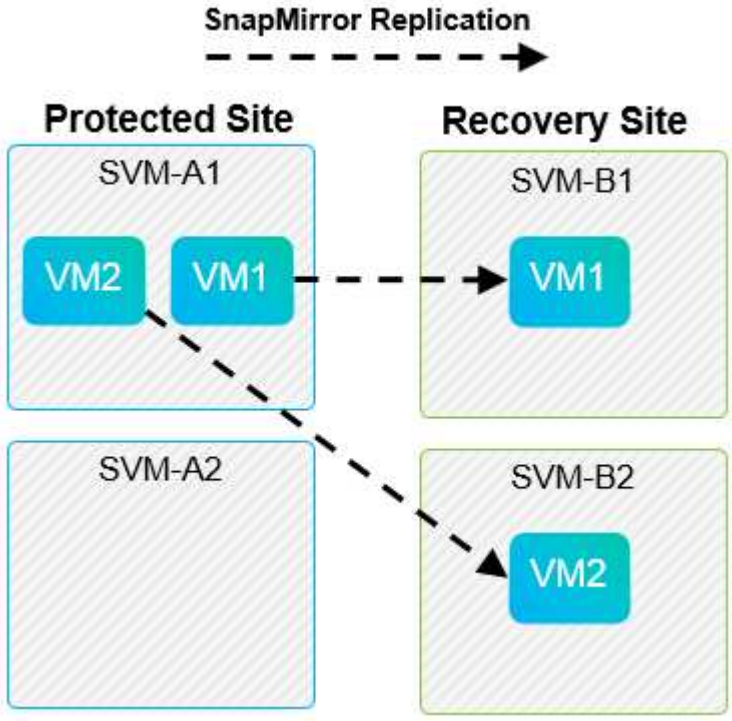
SnapMirror Replication
Protected Site
Recovery Site
SVM-A1
VM2
VM1
SVM-B1
VM1
SVM-A2
SVM-B2
VM2

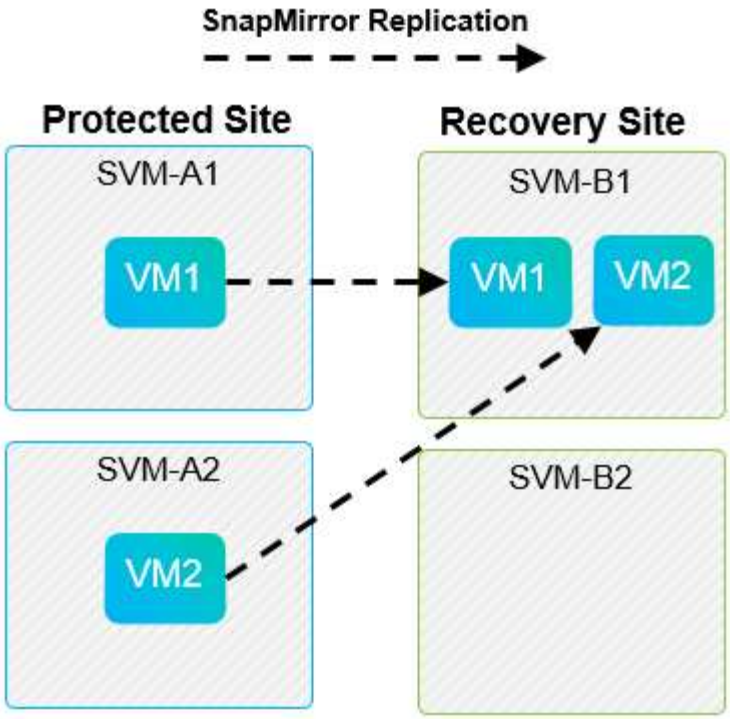

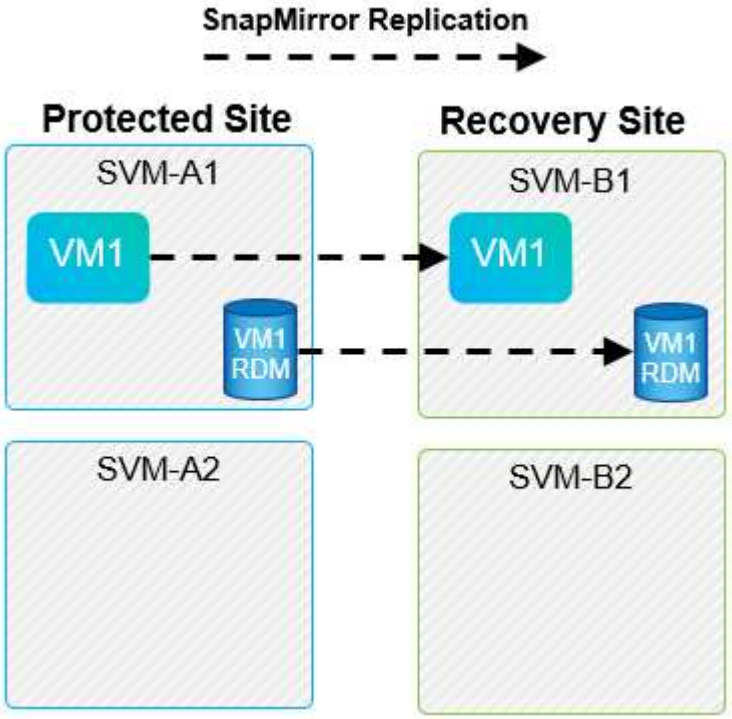

VMFS support with SnapMirror active sync

ONTAP tools 10.3 and later also support protecting your VMFS datastores with SnapMirror active sync(SMas). This enables transparent failover for business continuity between two datacenters (referred to as failure domains) that are relatively close together. Long-distance disaster recovery can then be orchestrated using SnapMirror asynchronous through the ONTAP tools SRA with VLSR.

Learn about ONTAP SnapMirror active sync

Datastores are collected together in a consistency group(CG), and the VMs across all datastores will all remain

write-order consistent as members of the same CG.

Some examples might be to have sites in Berlin and Hamburg protected by SMas, and a third site replica using SnapMirror asynchronous and protected by VLSR. Another example might be to protect sites in New York and New Jersey using SMas, with a third site in Chicago.

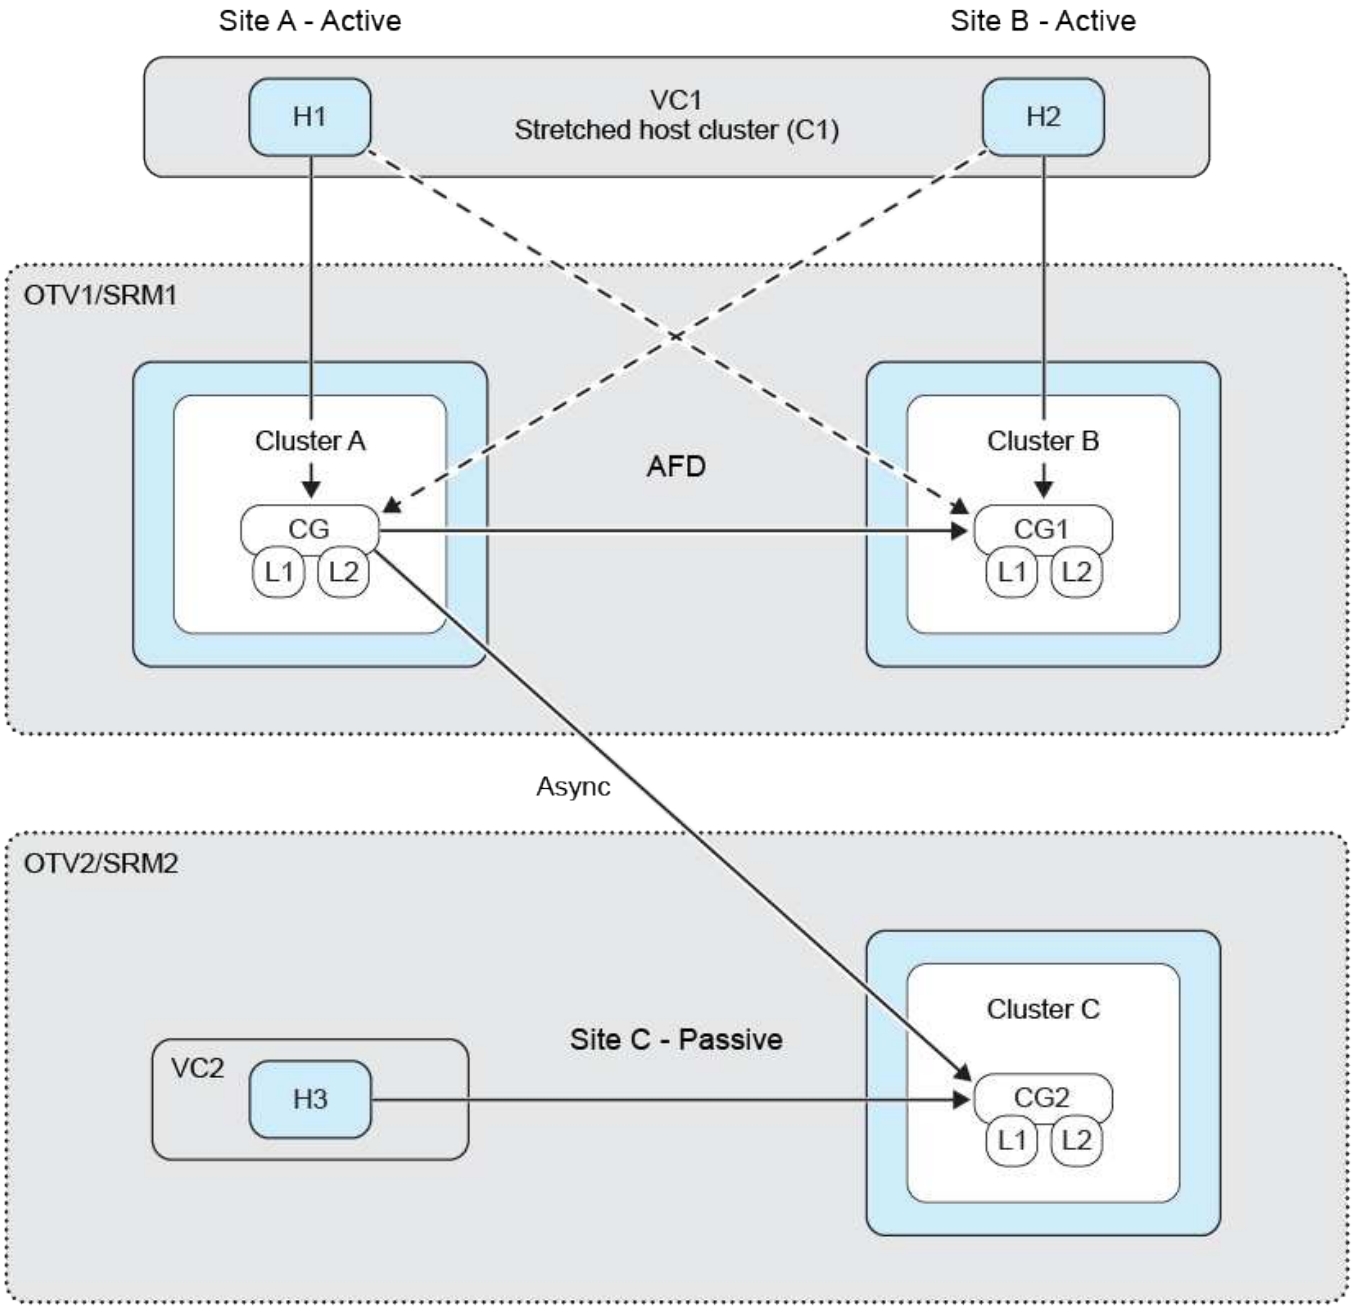

Supported Array Manager layouts

When you use array-based replication (ABR) in VLSR, protection groups are isolated to a single array pair, as shown in the following screenshot. In this scenario, `SVM1` and `SVM2` are peered with `SVM3` and `SVM4` at the recovery site. However, you can select only one of the two array pairs when you create a protection group.

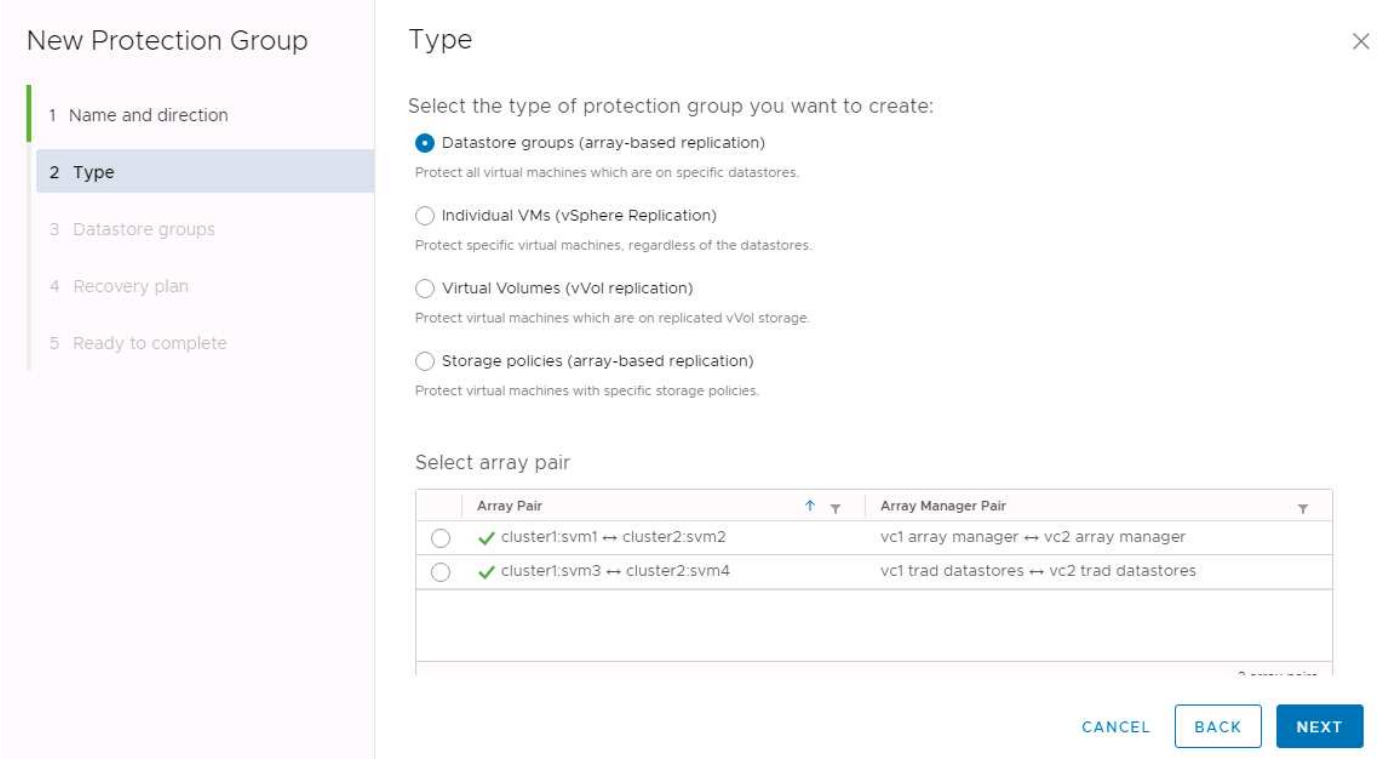

Unsupported layouts

Unsupported configurations have data (VMDK or RDM) on multiple SVMs that is owned by an individual VM. In the examples shown in the following figures, `VM1` cannot be configured for protection with VLSR because `VM1` has data on two SVMs.

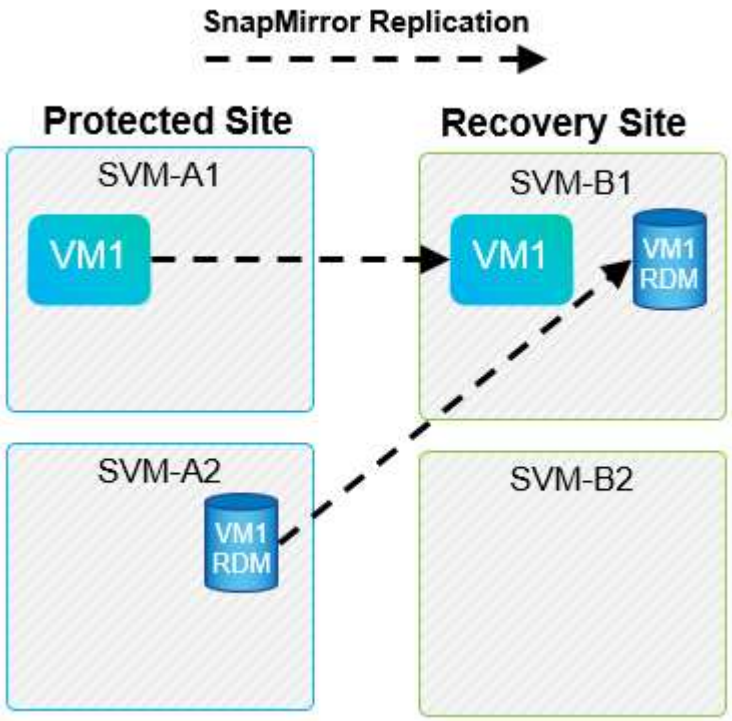

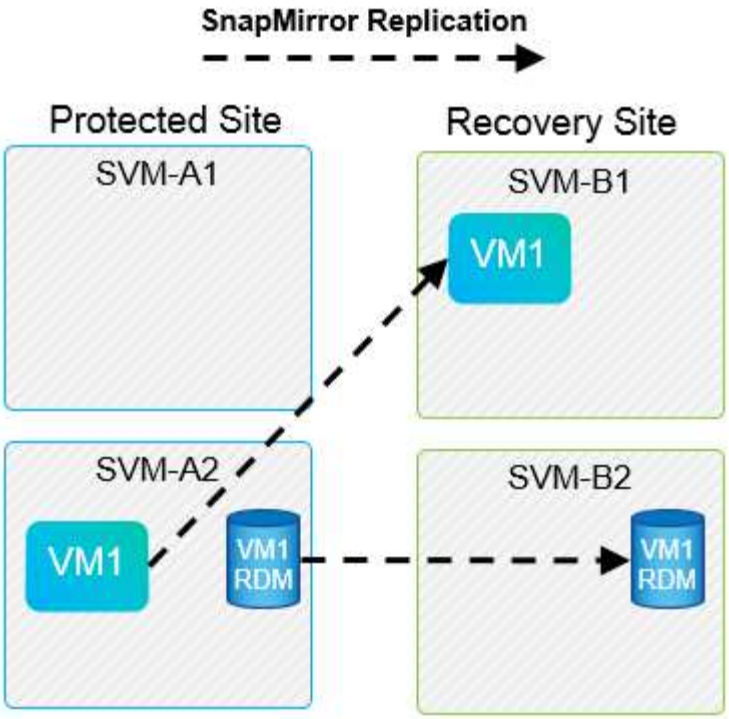

Any replication relationship in which an individual NetApp volume is replicated from one source SVM to multiple destinations in the same SVM or in different SVMs is referred to as SnapMirror fan-out. Fan-out is not supported with VLSR. In the example shown in the following figure, `VM1` cannot be configured for protection in VLSR because it is replicated with SnapMirror to two different locations.

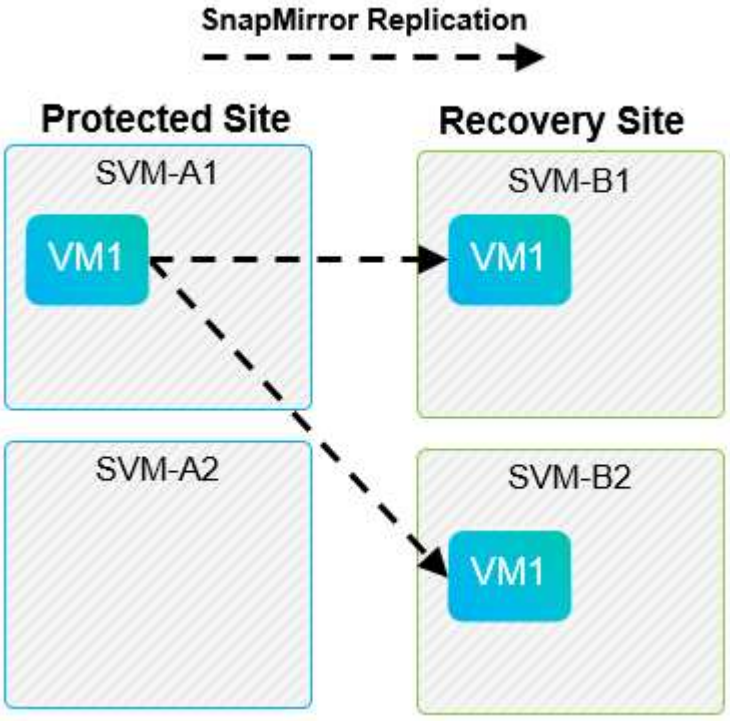

SnapMirror cascade

VLSR does not support cascading of SnapMirror relationships, in which a source volume is replicated to a destination volume and that destination volume is also replicated with SnapMirror to another destination

volume. In the scenario shown in the following figure, VLSR cannot be used for failover between any sites.

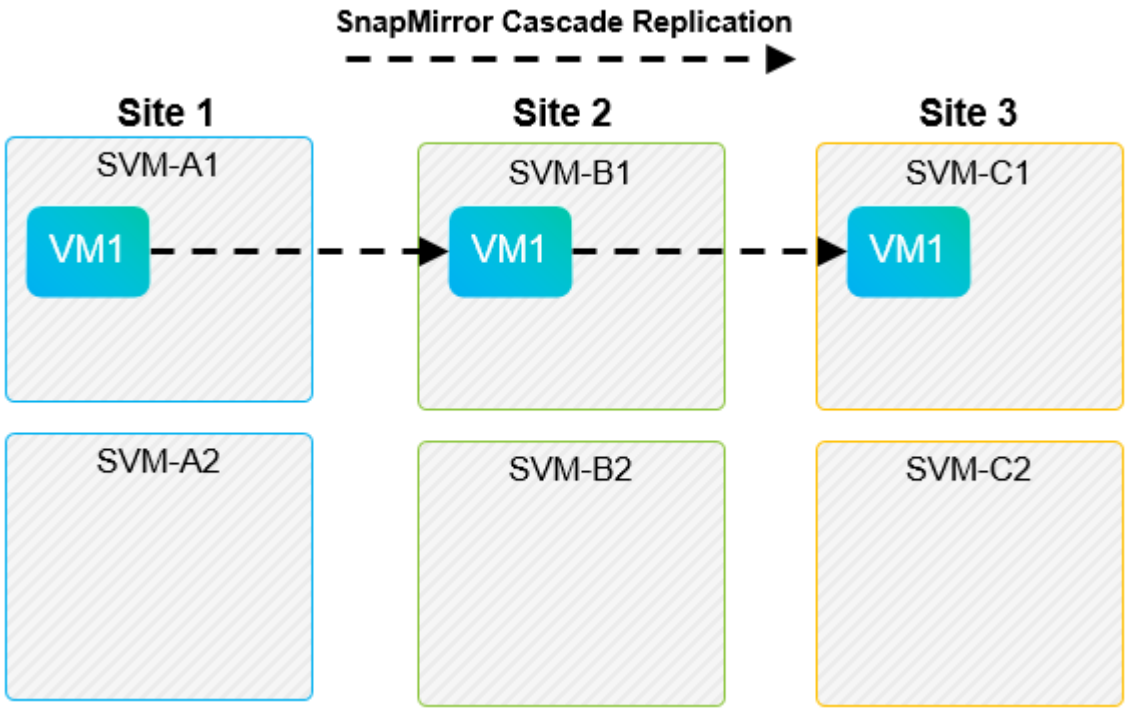

SnapMirror and SnapVault

NetApp SnapVault software enables disk-based backup of enterprise data between NetApp storage systems. SnapVault and SnapMirror can coexist in the same environment; however, VLSR supports the failover of only the SnapMirror relationships.

The NetApp SRA supports the `mirror-vault` policy type.

SnapVault was rebuilt from the ground up for ONTAP 8.2. Although former Data ONTAP 7-Mode users should find similarities, major enhancements have been made in this version of SnapVault. One major advance is the ability to preserve storage efficiencies on primary data during SnapVault transfers.

An important architectural change is that SnapVault in ONTAP 9 replicates at the volume level as opposed to at the qtree level, as is the case in 7-Mode SnapVault. This setup means that the source of a SnapVault relationship must be a volume, and that volume must replicate to its own volume on the SnapVault secondary system.

In an environment in which SnapVault is used, specifically named snapshots are created on the primary storage system. Depending on the configuration implemented, the named snapshots can be created on the primary system by a SnapVault schedule or by an application such as NetApp Active IQ Unified Manager. The named snapshots that are created on the primary system are then replicated to the SnapMirror destination, and from there they are vaulted to the SnapVault destination.

A source volume can be created in a cascade configuration in which a volume is replicated to a SnapMirror destination in the DR site, and from there it is vaulted to a SnapVault destination. A source volume can also be created in a fan-out relationship in which one destination is a SnapMirror destination and the other destination is a SnapVault destination. However, SRA does not automatically reconfigure the SnapVault relationship to use the SnapMirror destination volume as the source for the vault when VLSR failover or replication reversal occurs.

For the latest information about SnapMirror and SnapVault for ONTAP 9, see TR-4015 SnapMirror Configuration Best Practice Guide for ONTAP 9.

Best Practice
If SnapVault and VLSR are used in the same environment, NetApp recommends using a SnapMirror to SnapVault cascade configuration in which SnapVault backups are normally performed from the SnapMirror destination at the DR site. In the event of a disaster, this configuration makes the primary site inaccessible. Keeping the SnapVault destination at the recovery site allows SnapVault backups to be reconfigured after failover so that SnapVault backups can continue while operating at the recovery site.

In a VMware environment, each datastore has a universal unique identifier (UUID), and each VM has a unique managed object ID (MOID). These IDs are not maintained by VLSR during failover or failback. Because datastore UUIDs and VM MOIDs are not maintained during failover by VLSR, any applications that depend on these IDs must be reconfigured after VLSR failover. An example application is NetApp Active IQ Unified Manager, which coordinates SnapVault replication with the vSphere environment.

The following figure depicts a SnapMirror to SnapVault cascade configuration. If the SnapVault destination is at the DR site or at a tertiary site that is not affected by an outage at the primary site, the environment can be reconfigured to allow backups to continue after failover.

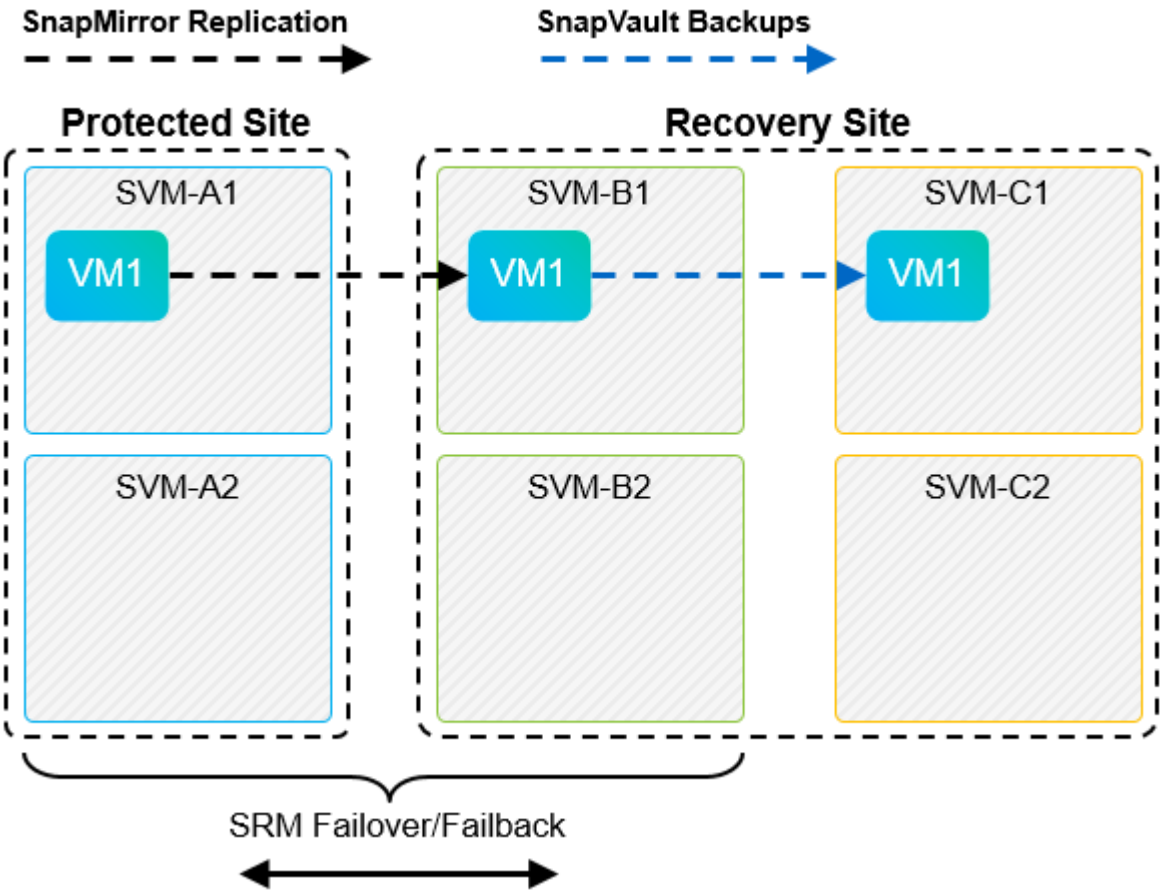

The following figure depicts the configuration after VLSR has been used to reverse SnapMirror replication back to the primary site. The environment has also been reconfigured such that SnapVault backups are occurring from what is now the SnapMirror source. This setup is a SnapMirror SnapVault fan-out configuration.

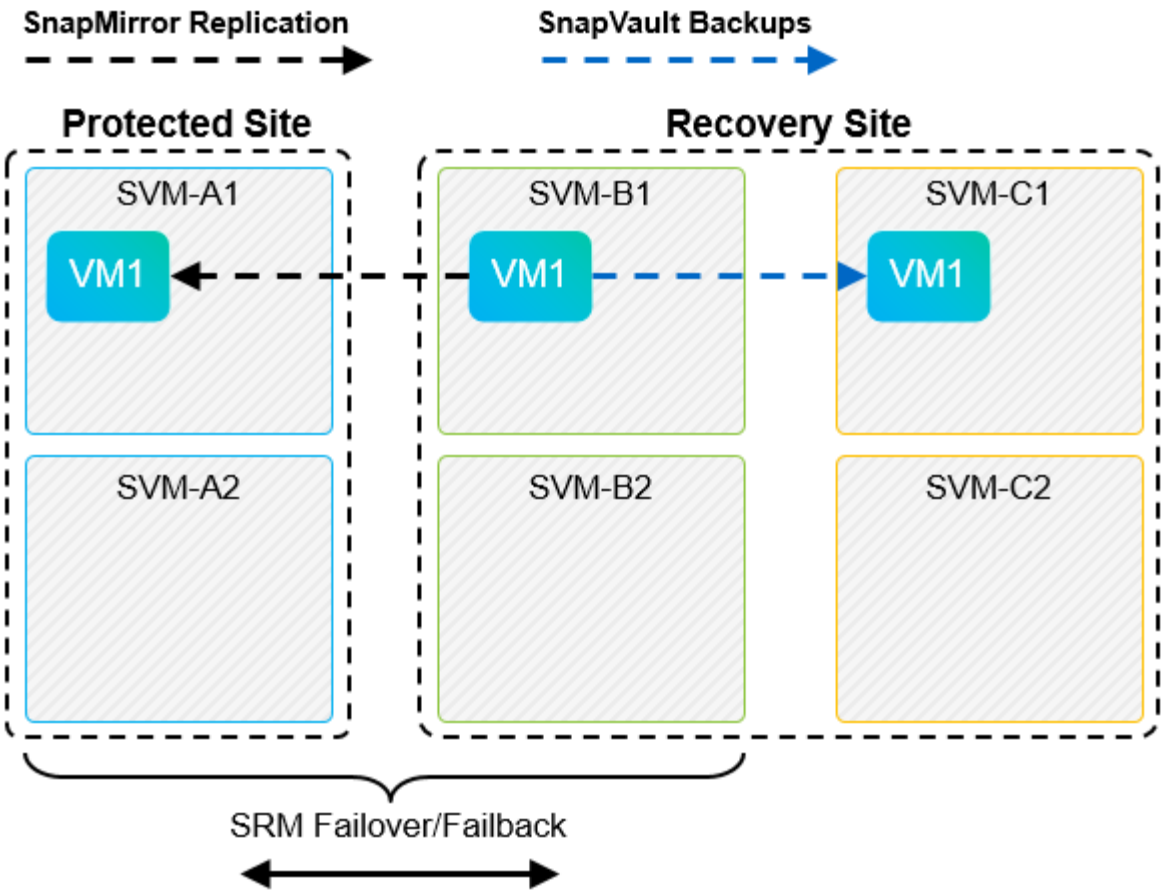

After vsrm performs failback and a second reversal of the SnapMirror relationships, the production data is back at the primary site. This data is now protected in the same way that it was before the failover to the DR site—through SnapMirror and SnapVault backups.

Use of Qtrees in Site Recovery Manager environments

Qtrees are special directories that allow the application of file system quotas for NAS. ONTAP 9 allows the creation of qtrees, and qtrees can exist in volumes that are replicated with SnapMirror. However, SnapMirror does not allow replication of individual qtrees or qtree-level replication. All SnapMirror replication is at the volume level only. For this reason, NetApp does not recommend the use of qtrees with VLSR.

Mixed FC and iSCSI environments

With the supported SAN protocols (FC, FCoE, and iSCSI), ONTAP 9 provides LUN services—that is, the ability to create and map LUNs to attached hosts. Because the cluster consists of multiple controllers, there are multiple logical paths that are managed by multipath I/O to any individual LUN. Asymmetric logical unit access (ALUA) is used on the hosts so that the optimized path to a LUN is selected and is made active for data transfer. If the optimized path to any LUN changes (for example, because the containing volume is moved), ONTAP 9 automatically recognizes and nondisruptively adjusts for this change. If the optimized path becomes unavailable, ONTAP can nondisruptively switch to any other available path.

VMware VLSR and NetApp SRA support the use of the FC protocol at one site and the iSCSI protocol at the other site. It does not support having a mix of FC-attached datastores and iSCSI-attached datastores in the same ESXi host or in different hosts in the same cluster, however. This configuration is not supported with VLSR because, during the VLSR failover or test failover, VLSR includes all FC and iSCSI initiators in the ESXi hosts in the request.

Best Practice
VLSR and SRA support mixed FC and iSCSI protocols between the protected and recovery sites. However, each site should be configured with only one protocol, either FC or iSCSI, not both protocols at the same site. If a requirement exists to have both FC and iSCSI protocols configured at the same site, NetApp recommends that some hosts use iSCSI and other hosts use FC. NetApp also recommends in this case that VLSR resource mappings be set up so that the VMs are configured to fail over into one group of hosts or the other.

Troubleshooting VLSRM/SRM when using vVols replication

When using ONTAP tools 9.13P2, the workflow within VLSR and SRM is significantly different when using vVols replication from what is used with SRA and traditional datastores. For example, there is no array manager concept. As such, `discoverarrays` and `discoverdevices` commands are never seen.

When troubleshooting, it is beneficial to understand the new workflows, which are listed below:

1. queryReplicationPeer: Discovers the replication agreements between two fault domains.
2. queryFaultDomain: Discovers fault domain hierarchy.
3. queryReplicationGroup: Discovers the replication groups present in the source or target domains.
4. syncReplicationGroup: Synchronizes the data between source and target.
5. queryPointInTimeReplica: Discovers the point in time replicas on a target.
6. testFailoverReplicationGroupStart: Begins test failover.
7. testFailoverReplicationGroupStop: Ends test failover.
8. promoteReplicationGroup: Promotes a group currently in test to production.
9. prepareFailoverReplicationGroup: Prepares for a disaster recovery.
10. failoverReplicationGroup: Executes disaster recovery.
11. reverseReplicateGroup: Initiates reverse replication.
12. queryMatchingContainer: Finds containers (along with Hosts or Replication Groups) that might satisfy a provisioning request with a given policy.
13. queryResourceMetadata: Discovers the metadata of all resources from the VASA provider, the resource utilization can be returned as an answer to the queryMatchingContainer function.

The most common error seen when configuring vVols replication is a failure to discover the SnapMirror relationships. This occurs because the volumes and SnapMirror relationships are created outside of the purview of ONTAP Tools. Therefore, it is a best practice to always make sure your SnapMirror relationship is fully initialized and that you have run a rediscovery in ONTAP Tools at both sites before attempting to create a replicated vVols datastore.

Additional Information

To learn more about the information that is described in this document, review the following documents and/or websites:

- ONTAP tools for VMware vSphere 10.x Resources
  https://mysupport.netapp.com/site/products/all/details/otv10/docs-tab

- ONTAP tools for VMware vSphere 9.x Resources
  https://mysupport.netapp.com/site/products/all/details/otv/docsandkb-tab
- TR-4597: VMware vSphere for ONTAP
  https://docs.netapp.com/us-en/ontap-apps-dbs/vmware/vmware-vsphere-overview.html
- TR-4400: VMware vSphere Virtual Volumes with ONTAP
  https://docs.netapp.com/us-en/ontap-apps-dbs/vmware/vmware-vvols-overview.html
- TR-4015 SnapMirror Configuration Best Practice Guide for ONTAP 9
  https://www.netapp.com/pdf.html?item=/media/17229-tr-4015-snapmirror-configuration-ontap.pdf
- VMware Live Site Recovery Documentation
  https://techdocs.broadcom.com/us/en/vmware-cis/live-recovery/live-site-recovery/9-0.html

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support Site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright information

Trademark information