



Business continuity for Proxmox VE workloads

NetApp virtualization solutions

NetApp
October 22, 2026

This PDF was generated from <https://docs.netapp.com/us-en/netapp-solutions-virtualization/proxmox/proxmox-business-continuity.html> on October 22, 2026. Always check docs.netapp.com for the latest.

Table of Contents

Business continuity for Proxmox VE workloads	1
Learn about Proxmox VE business continuity with SnapMirror active sync	1
Customer challenges	1
Solution overview	2
How the solution responds to failures	5
Choose between SnapMirror active sync and MetroCluster	7
Corosync considerations for a Proxmox VE stretch cluster	10
Integration with Storage Manager for Proxmox	11
NFS shared storage with SnapMirror active sync for NAS	12
Related information	16
Configure Proxmox VE business continuity with NAS storage	16
Step 1: Configure the mediator	17
Step 2: Protect the NFS SVM with SnapMirror active sync	18
Step 3: Configure a VIP LIF with BGP for layer 3 networks	19
Step 4: Create the Proxmox VE stretch cluster with a QDevice	21
Step 5: Provision protected NFS shared storage	22
Step 6: Configure HA and VM backups	23
Step 7: Test planned failover	24
Related information	25
Configure and use Storage Manager for Proxmox	25
Learn about Storage Manager for Proxmox	25
Install and set up Storage Manager for Proxmox	27
Configure Storage Manager for Proxmox	29
Provision storage	34
Protect virtual machines	45
Restore virtual machines or datastores	50
Clone virtual machines	54
Monitor operations	56
Summary	58

Business continuity for Proxmox VE workloads

Learn about Proxmox VE business continuity with SnapMirror active sync

Enterprises that run business-critical virtual machines (VMs) on Proxmox Virtual Environment (VE), and service providers that host customer VMs on Proxmox VE, need those VMs to keep running when a storage system or an entire data center goes offline. ONTAP SnapMirror active sync synchronously replicates the shared storage of a Proxmox VE cluster stretched across two sites and automatically fails over storage to the surviving site. You get a zero recovery point objective (RPO), automated recovery without scripts or runbooks, and the same shared storage at both sites, so Proxmox VE high availability (HA) can restart VMs without any storage reconfiguration. For NFS shared storage, the solution uses SnapMirror active sync for NAS.

Customer challenges

Building storage and site resilience for Proxmox VE with native tools alone is difficult:

- **Data loss after a site failure.** Proxmox VE storage replication and Proxmox Backup Server are asynchronous. Changes made since the last replication or backup are lost when the primary site fails.
- **Manual, slow recovery.** Recovering at a secondary site usually means promoting replicas, mounting storage under a different address, updating `/etc/pve/storage.cfg`, and restarting VMs. Each manual step adds downtime and the risk of error during an outage.
- **Stretched storage that competes with VMs for host resources.** Hyperconverged stretched storage, such as Ceph in stretch mode, runs on the Proxmox VE nodes and keeps two replicas of every object at each site. Replication, recovery, and rebalancing traffic consume CPU, memory, and network resources on the same nodes that run VMs.
- **Two quorum decisions that must agree.** Any two-site design needs a witness in a third failure domain, for the Proxmox VE cluster and for the storage. If the cluster and the storage choose different surviving sites after an inter-site link failure, VMs restart on nodes that can't reach the active storage.
- **A storage address that must work at both sites.** After a storage failover, Proxmox VE nodes must reach the shared storage at the same address. That address must be reachable at both sites, which usually means a layer 2 VLAN stretched between the sites.

Service provider challenges

Service providers that build cloud services on Proxmox VE face the same challenges, with tenant service level agreements (SLAs) on top:

- **The storage system is a shared point of failure.** Proxmox VE HA protects tenant VMs from a host failure, but not from a failure of the storage system that hosts their disks. One storage outage affects every tenant on that system.
- **No way to sell higher availability.** Without synchronous replication and automatic storage failover, a provider can't offer a premium availability tier with zero data loss and stronger uptime commitments.
- **Tenant isolation during recovery.** Recovery, maintenance, and disaster recovery tests must affect only the tenants involved, not the whole platform.

Enterprise challenges

Enterprises that move mission-critical workloads from other hypervisors to Proxmox VE expect the same site resilience they had before, such as zero data loss for databases and line-of-business applications and recovery without a manual runbook. Meeting that expectation with Proxmox VE native tools alone takes significant design, scripting, and operational effort.

Solution overview

The solution combines a Proxmox VE stretch cluster with SnapMirror active sync on NetApp AFF systems at two sites:


- **Zero RPO.** ONTAP replicates every write to the shared storage synchronously from the primary cluster at site A to the secondary cluster at site B.
- **Same shared storage at both sites.** After a failover, the secondary cluster presents the same storage to Proxmox VE, so the storage definitions in `/etc/pve/storage.cfg` stay valid.
- **Automatic failover.** A mediator acts as a witness outside the two sites. You can install ONTAP Mediator at a third site, or use ONTAP Cloud Mediator, a service hosted in the NetApp Console. When the primary storage or site fails, ONTAP fails over to site B without manual intervention.
- **Granular control.** You choose which workloads to protect, and you can fail over workloads individually, for example to move one tenant or application for site maintenance.
- **Use of the secondary copy.** You can clone the secondary copy for development, testing, or disaster recovery testing without affecting replication.

How the solution addresses the challenges

Challenge	How the solution helps	What you still plan for
Data loss and manual recovery	Synchronous replication gives zero RPO. ONTAP fails over automatically, and the secondary SVM presents the same NFS address, so <code>/etc/pve/storage.cfg</code> stays valid.	Proxmox VE HA still restarts VMs from a failed site.
Stretched storage on the hosts	Replication runs on the AFF controllers, not on the Proxmox VE nodes, so host CPU, memory, and network stay available for VMs. ONTAP keeps one copy at each site, and storage efficiency applies to both copies.	Like any synchronous replication, the solution needs a full second copy at site B, and replication adds load to both AFF clusters. Size both clusters for the replication workload.
Two quorum decisions	ONTAP Mediator, or ONTAP Cloud Mediator as a service in the NetApp Console, is a lightweight witness outside the data path. This solution shows how to align the corosync QDevice with the primary cluster so that both choose the same site.	A third location is still required for the QDevice. It's also required for ONTAP Mediator, but not for ONTAP Cloud Mediator. See Mediator options .

Challenge	How the solution helps	What you still plan for
A storage address at both sites	The secondary SVM takes over the same NFS LIF address, so no remount or storage reconfiguration is needed.	The NFS subnet must be reachable at both sites, typically with a stretched layer 2 VLAN. See Layer 3 networks with VIP LIFs and BGP for the qualification status of VIP LIFs.

The following diagram shows the solution architecture with NFS shared storage.



Site A and site B

Each site has an AFF two-node cluster and half of the Proxmox VE nodes. The primary cluster at site A serves the shared storage to all Proxmox VE nodes. Nodes at site B reach the primary cluster across the inter-site network. With NFS storage, the secondary SVM at site B stays dormant until a failover occurs.

Site C

A third failure domain hosts the corosync QDevice daemon (`corosync-qnetd`) and, if you choose option 1, ONTAP Mediator. The QDevice provides an extra vote so that the Proxmox VE cluster keeps quorum when one site fails.

NetApp Console

If you choose option 2, ONTAP Cloud Mediator runs as a service in the NetApp Console, which NetApp hosts and manages. Both AFF clusters connect to it over HTTPS, so you don't install or maintain a mediator

host.

Mediator

ONTAP Mediator or ONTAP Cloud Mediator helps ONTAP decide which cluster serves data. Use one of them, not both. For a comparison, see [Mediator options](#).

Proxmox VE stretch cluster

All Proxmox VE nodes belong to one cluster and share the same storage definitions. Proxmox VE HA restarts VMs from a failed site on the surviving nodes. Storage Manager for Proxmox (SMP) runs as an HA VM on the cluster to provision shared storage and to back up, restore, and clone VMs.

Mediator options

SnapMirror active sync uses a mediator as a passive witness. During a network partition or the failure of one cluster, ONTAP uses the mediator to decide which copy continues to serve I/O. The mediator isn't in the data path and doesn't start failovers itself. You can use one of the following options, but not both at the same time.

Consideration	Option 1: ONTAP Mediator	Option 2: ONTAP Cloud Mediator
Where it runs	A Linux host that you install and manage in a third failure domain, such as site C, with independent power and network infrastructure.	A cloud service in the NetApp Console. NetApp hosts and manages it, so you don't need a mediator host at a third site.
ONTAP version	ONTAP 9.19.1 or later for SnapMirror active sync for NAS.	ONTAP Cloud Mediator is available beginning with ONTAP 9.17.1. SnapMirror active sync for NAS needs ONTAP 9.19.1 or later.
Cluster connectivity	Both clusters connect to the mediator host. The default service port is 31784.	Both clusters connect outbound over HTTPS to the NetApp Console, directly or through an HTTP proxy. The recommended ping latency to the NetApp Console is under 200 ms.
Setup	Install the package on the Linux host, and then add it to the clusters with the host address and a mediator user.	Create a NetApp Console service account with the ONTAP Mediator Setup Role , and then add the mediator to the clusters with the organization ID, client ID, and client secret.
Operations	You patch and upgrade the host, and you keep it available.	No host to patch or upgrade. You depend on internet access from both sites and on the availability of the NetApp Console service.
Best fit	Environments that have a third site, or that don't allow the storage clusters to reach the internet.	Environments without a third site for a mediator, or that want to avoid operating one.

Keep the following in mind when you choose:

- **The QDevice still needs a third location.** ONTAP Cloud Mediator removes the mediator host, not the corosync QDevice host. The QDevice needs a small Linux host that all Proxmox VE nodes can reach on TCP port 5403, and it doesn't have the low-latency requirement of corosync.

- **Both clusters need access to the NetApp Console.** If the clusters use a proxy, install the proxy's root CA certificate on both clusters. NetApp recommends a different proxy server for each cluster.
- **Plan for the same failure behavior with either option.** When the mediator is unavailable, replication continues without data loss, but automatic and planned failovers aren't possible until the mediator is available again. With ONTAP Cloud Mediator, this includes a loss of connectivity from the clusters to the NetApp Console.



The mediator options in this section apply to SnapMirror active sync. For MetroCluster, this solution uses ONTAP Mediator or MetroCluster Tiebreaker. Before you deploy, verify that your ONTAP version, mediator option, and Proxmox VE version are listed together in the [NetApp Interoperability Matrix Tool \(IMT\)](#).

Benefits for service providers

- **VM availability beyond host HA.** Tenant VMs keep running when the storage system that hosts their disks fails. ONTAP fails over to the second AFF cluster, and the VMs resume I/O on the same shared storage without a restart.
- **A premium availability tier.** You can offer a zero RPO, storage-resilient VM service as an add-on, alongside standard tiers on unprotected storage in the same clusters.
- **Per-tenant protection and failover.** With separate protected storage for each tenant or service tier, you protect, fail over, and test disaster recovery for each tenant independently, without affecting other tenants.
- **Better use of both sites.** You can make some workloads primary at site A and others primary at site B, so both data centers serve active workloads.

Benefits for enterprises

- **Zero data loss for business-critical VMs** during a storage or site failure.
- **Automated recovery** without scripts, storage remounts, or changes to `/etc/pve/storage.cfg`.
- **Built into ONTAP**, with no extra software layer on the Proxmox VE nodes.
- **Day-to-day operations from one interface.** Storage Manager for Proxmox provisions the protected shared storage and backs up, restores, and clones VMs with ONTAP snapshots and FlexClone technology.

How the solution responds to failures


ONTAP and Proxmox VE make independent quorum decisions. The design aligns them so that both choose the same surviving site.

Event	ONTAP behavior	Proxmox VE behavior
Site A fails	The mediator confirms the failure, and ONTAP automatically fails over the shared storage to cluster B.	Site B nodes and the QDevice keep quorum. HA fences the site A nodes and restarts their VMs on site B nodes. VMs already on site B resume I/O after the storage failover.
AFF cluster A fails, Proxmox VE nodes at site A keep running	The mediator confirms that cluster A is unreachable, and ONTAP automatically fails over the shared storage to cluster B.	No quorum change and no VM restart. VMs on all nodes resume I/O on the same shared storage after the storage failover, with site A nodes reaching cluster B across the inter-site network.

Event	ONTAP behavior	Proxmox VE behavior
One AFF controller at site A fails	The partner controller takes over in a local HA takeover. No SnapMirror active sync failover occurs.	No impact. VMs pause I/O briefly during the takeover.
Site B fails	Cluster A continues to serve data. Replication resumes when site B returns.	Site A nodes and the QDevice keep quorum. HA restarts the site B VMs on site A nodes.
Inter-site link fails	Cluster A continues to serve data. No failover occurs.	The QDevice gives its vote to the site A partition when the tie breaker is aligned with the primary cluster. Site B nodes lose quorum and self-fence, and HA restarts their VMs on site A.
Site C fails, or the clusters lose access to the NetApp Console	Replication continues with zero data loss, but automatic and planned failovers aren't possible until the mediator is available. With ONTAP Cloud Mediator, a site C failure affects only the QDevice.	The cluster keeps quorum with the votes of the four nodes. No VMs are affected.
One Proxmox VE node fails	No impact.	HA restarts the VMs from that node on other nodes, preferring nodes at the same site.

Storage system failure

A storage system failure is the case that host-level HA alone can't handle, and the case where SnapMirror active sync adds the most value for service providers. The following diagram shows the state after AFF cluster A fails while every Proxmox VE node stays online, using NFS shared storage.



1. AFF cluster A becomes unavailable. The Proxmox VE nodes at site A keep running, so the cluster doesn't lose quorum and HA doesn't fence any node.
2. The mediator, ONTAP Mediator at site C or ONTAP Cloud Mediator in the NetApp Console, confirms that cluster A is unreachable, and ONTAP activates the secondary SVM on cluster B with the same NFS LIF address, volumes, and export policies.
3. All Proxmox VE nodes resume I/O on the same shared storage. Site B nodes reach cluster B locally, and site A nodes reach it across the inter-site network. VMs pause I/O during the failover and keep running without a restart.

When cluster A returns, replication resumes. Plan the failback to cluster A during a maintenance window so that site A VMs access local storage again.

Site failure

During a site failure, both the storage system and the Proxmox VE nodes at that site go offline. ONTAP fails over the shared storage to the surviving cluster, and Proxmox VE HA restarts the VMs from the failed site on the surviving nodes, as described in the table.

Choose between SnapMirror active sync and MetroCluster

NetApp MetroCluster configurations also deliver zero RPO and automatic recovery for a Proxmox VE stretch cluster. The two solutions differ in what they protect:

- **MetroCluster protects all workloads.** MetroCluster combines array-based clustering with synchronous replication. It mirrors the data in every mirrored aggregate, the nonvolatile memory (NVRAM), and the configuration of every SVM to the partner site, so all shared storage on the cluster is protected without per-workload setup.

- **SnapMirror active sync protects a subset of workloads.** You choose which SVMs or volumes to protect, and you can fail over each one independently. Unprotected workloads can share the same clusters.


Both solutions keep a full copy of the protected data at each site, need a witness in a third failure domain, and add replication load to the storage controllers. They differ in protection scope, failover granularity, and cluster design.



Before you deploy either solution, verify that your ONTAP version, MetroCluster or SnapMirror active sync configuration, and Proxmox VE version are listed together in the [NetApp Interoperability Matrix Tool \(IMT\)](#).

MetroCluster architecture with Proxmox VE

The following diagram shows a Proxmox VE stretch cluster with a four-node MetroCluster IP configuration and NFS shared storage.



MetroCluster IP clusters

Each site has an HA pair with mirrored aggregates. The MetroCluster IP back-end network mirrors data (SyncMirror) and NVRAM to the partner site synchronously.

SVMs at both sites

Each site can serve its own SVMs. For each SVM, MetroCluster keeps a stopped mirror SVM with the `-mc` suffix at the partner site, with the same configuration and LIF addresses.

NFS access from all nodes

Every Proxmox VE node mounts the shared storage of both active SVMs, SVM A1 and SVM B1, so that HA can restart any VM on any node. Nodes at site B reach SVM A1, and nodes at site A reach SVM B1, across

the inter-site network. This is the same access model as SnapMirror active sync.

Switchover

If a site fails, ONTAP Mediator or MetroCluster Tiebreaker at site C enables an automatic switchover. The mirror SVMs at the surviving site start and serve all shared storage with the same NFS addresses, and Proxmox VE HA restarts the VMs from the failed site. A single controller failure is handled by a local HA takeover, without a switchover. After you repair the failed site, you heal the aggregates and switch back.

How MetroCluster responds to failures

The following table shows how a MetroCluster IP configuration and the Proxmox VE stretch cluster respond to the same failures as in [How the solution responds to failures](#). In this example, SVM A1 is active at site A and SVM B1 is active at site B.

Event	ONTAP behavior	Proxmox VE behavior
Site A fails	ONTAP Mediator or MetroCluster Tiebreaker confirms the failure, and ONTAP switches over site A. SVM A1-mc starts at site B with the same NFS LIF address. SVM B1 keeps serving data.	Site B nodes and the QDevice keep quorum. HA fences the site A nodes and restarts their VMs on site B nodes. VMs that already ran on site B nodes keep running, and those on SVM A1 resume I/O after the switchover.
MetroCluster IP cluster A fails, Proxmox VE nodes at site A keep running	ONTAP Mediator or MetroCluster Tiebreaker confirms that cluster A is unreachable, and ONTAP switches over site A. SVM A1-mc starts at site B with the same NFS LIF address.	No quorum change and no VM restart. VMs on SVM A1 resume I/O after the switchover, with site A nodes reaching site B across the inter-site network.
One controller at site A fails	The HA partner at site A takes over in a local HA takeover. No switchover occurs.	No impact. VMs pause I/O briefly during the takeover.
Site B fails	ONTAP switches over site B. SVM B1-mc starts at site A with the same NFS LIF address. SVM A1 keeps serving data.	Site A nodes and the QDevice keep quorum. HA restarts the site B VMs on site A nodes, where they access SVM B1 at site A.
Inter-site link fails	No switchover occurs, because both clusters are still running. Each site keeps serving its own SVMs, and mirroring to the partner site is suspended until the link returns.	The QDevice gives its vote to the site A partition when the tie breaker is aligned with site A. Site B nodes lose quorum and self-fence. HA restarts VMs on SVM A1 on site A nodes. VMs on SVM B1 can't restart until the link returns, because site A nodes can't reach SVM B1.
Site C fails	Mirroring continues with zero data loss, but automatic switchover isn't possible until the witness is available. You can still switch over manually.	The cluster keeps quorum with the votes of the four nodes. No VMs are affected.
One Proxmox VE node fails	No impact.	HA restarts the VMs from that node on other nodes, preferring nodes at the same site.

The inter-site link failure is the main difference from a design with one primary site. Whenever both sites host

active SVMs, with MetroCluster or with two SnapMirror active sync SVMs in opposite directions, the corosync tie breaker can favor only one site. VMs whose active SVM is at the other site stay down until the link returns. To limit the impact, use redundant inter-site links on separate physical paths, and place the most critical VMs on the SVM at the site that the tie breaker favors.

Feature comparison

Consideration	SnapMirror active sync for NAS	MetroCluster IP
Protection scope	Selected SVMs and volumes	All mirrored aggregates and all SVMs in the configuration
Failover granularity	Each SVM independently	The whole site, by switchover
Unprotected workloads	Run on unprotected SVMs in the same clusters	Can use unmirrored aggregates, which go offline after a switchover until you bring them online
Cluster design	Two independent clusters, peered over an IP network	One configuration of two partner clusters, with a dedicated MetroCluster IP back-end network
Witness at site C	ONTAP Mediator or ONTAP Cloud Mediator	ONTAP Mediator or MetroCluster Tiebreaker
Copies of protected data	One at each site	One at each site, in mirrored aggregates
Replication path	SnapMirror synchronous over intercluster LIFs	SyncMirror and NVRAM mirroring over the MetroCluster IP back-end network
NFS address at both sites	Stretched layer 2 VLAN. VIP LIFs with BGP pending qualification	Stretched layer 2 VLAN, or VIP LIFs with BGP
Recovery after a site failure	Fail back each SVM	Heal the aggregates, and then switch back the site

Selection guidance

- Choose MetroCluster when all, or nearly all, workloads on the storage need zero RPO, for example when an enterprise runs only mission-critical VMs on the cluster, or a service provider offers the same availability tier to every tenant. You get one protection domain without per-workload configuration.
- Choose SnapMirror active sync when only a subset of workloads needs zero RPO, for example a premium tier for some tenants, or when you need to fail over or test disaster recovery for one tenant or application at a time.
- SnapMirror active sync for NAS isn't supported on a MetroCluster configuration. Use one solution for each pair of clusters.
- The corosync, QDevice, and HA guidance in this solution applies to both. With MetroCluster, align the QDevice tie breaker with the site that keeps serving data when the inter-site link fails, and use HA node affinity rules to keep VMs at the same site as their active SVM.
- MetroCluster IP also supports VIP LIFs with BGP for layer 3 networks between sites.

Corosync considerations for a Proxmox VE stretch cluster

Proxmox VE uses corosync for cluster membership and quorum. Corosync is more sensitive to the inter-site network than SnapMirror active sync, so plan the cluster network first.

Latency and network design

- SnapMirror active sync requires an inter-cluster round-trip time (RTT) under 10 ms. Proxmox recommends a corosync latency under 5 ms between all nodes, so the corosync requirement sets the distance limit for the solution.
- Corosync needs low, consistent latency rather than high bandwidth. Use a dedicated network for corosync, separate from storage and live migration traffic, so that storage bursts don't cause corosync timeouts.
- Configure at least two corosync links (for example, link 0 and link 1) on separate physical paths between the sites. Corosync uses Kronosnet to switch links automatically if one path fails.

Quorum with a QDevice at a third site

Split the Proxmox VE nodes evenly between site A and site B, and add a QDevice at site C. With four nodes and a QDevice, the cluster has five votes and needs three for quorum, so either site can fail without losing quorum. The QDevice connects to the cluster over TCP port 5403 and doesn't have the low-latency requirement of corosync, so it can share site C with ONTAP Mediator. If you use ONTAP Cloud Mediator, you still need a third location for the QDevice, but it can be a small Linux host that doesn't also host a mediator.

Tie breaker alignment with the primary cluster

When the inter-site link fails but both sites can still reach site C, ONTAP keeps the primary cluster at site A serving data. The QDevice must give its vote to the same site. By default, the QDevice chooses the partition that contains the lowest node ID.

BEST PRACTICE: Create the Proxmox VE cluster on a site A node and join the other site A nodes before you join site B nodes, so that site A holds the lowest node IDs. If the primary storage moves to site B for an extended time, update the QDevice `tie_breaker` setting to a site B node ID.

If the QDevice favored site B instead, site A nodes would self-fence, and HA would restart VMs on site B nodes that can't reach the active storage at site A.

HA behavior and VM placement

- Proxmox VE HA typically detects a failed node and recovers its VMs in about two minutes, including a 60-second watchdog fencing timeout. During a site failure, the SnapMirror active sync failover and the Proxmox VE HA recovery run in parallel.
- Use HA node affinity rules with node priorities to place VMs on the nodes at the same site as their primary storage, and to let them fail over to the other site.
- After a site failover, set the HA `failback` property to 0 for affected VMs until you fail the storage back. Otherwise, HA migrates VMs back to site A nodes as soon as they return, while the active storage is still at site B.
- Make sure the guest operating system disk timeouts are longer than the expected failover time, so that guests don't report I/O errors during a storage failover.

Integration with Storage Manager for Proxmox

Storage Manager for Proxmox (SMP) gives Proxmox VE administrators one interface to provision ONTAP shared storage and to back up, restore, and clone VMs by using ONTAP snapshots and FlexClone technology. SMP connects to the Proxmox VE cluster and to ONTAP through their APIs, so it doesn't need agents on the Proxmox VE nodes.

SMP and SnapMirror active sync protect against different kinds of failures. SnapMirror active sync keeps VMs

running through a storage or site failure, but it replicates every write synchronously, including an accidental deletion, a failed guest update, or data encrypted by ransomware. SMP snapshots give you point-in-time copies to recover from those logical failures. Together, they provide continuous availability and fast operational recovery from the same ONTAP storage.

Protection layer	Protects against	Recovery point	Managed with
SnapMirror active sync	Storage system failure and site failure	Zero data loss	ONTAP
SMP backups (ONTAP snapshots)	Accidental deletion, data corruption, failed updates, and ransomware	Last scheduled or on-demand backup	SMP
SnapMirror asynchronous fan-out or Proxmox Backup Server	Loss of both sites, and long-term retention	Last replication or backup	ONTAP or Proxmox Backup Server

Backup and recovery

SMP creates crash-consistent or file-consistent backups as ONTAP snapshots of the storage volumes, on demand or by policy, and restores individual VMs, storage, or resource groups:

- Use resource groups and backup policies to apply the same schedule and retention to the VMs of a tenant or an application.
- Use file-consistent backups for VMs that run the QEMU guest agent, so the guest file systems are quiesced before ONTAP creates the snapshot.
- SMP manages the storage that is active. After a SnapMirror active sync failover, run a discovery in SMP so that it maps the storage to cluster B before you run backup or restore operations.
- For a third copy outside the two sites, or for longer retention, add a SnapMirror asynchronous fan-out relationship from the protected volumes, or use Proxmox Backup Server.
- Tamperproof snapshots aren't supported with SnapMirror active sync, so don't enable snapshot locking in SMP for protected volumes.

VM cloning

SMP uses ONTAP FlexClone technology to create space-efficient VM clones for development, testing, and rapid provisioning. Clones that you create in protected storage consume replication bandwidth and capacity at both sites. Exclude clones and test storage that don't need business continuity from SnapMirror active sync protection.

SMP placement

- Run the SMP VM as an HA resource on a SnapMirror active sync protected shared storage, so that SMP stays available after a storage or site failure.
- Register the Proxmox VE cluster and both ONTAP clusters in SMP, so that SMP can reach the active storage at either site.

NFS shared storage with SnapMirror active sync for NAS

Beginning with ONTAP 9.19.1, SnapMirror active sync protects NAS workloads. The following considerations apply only to NFS shared storage. To deploy the solution with NFS storage, follow the [NAS configuration workflow](#).

SVM-level protection and network identity

- SnapMirror active sync for NAS protects an entire storage VM (SVM). You can protect all volumes in the SVM or exclude selected volumes. Only FlexVol volumes are supported, and the SVM can't contain LUNs.
- The secondary SVM stays dormant, with its LIFs down, until a failover. ONTAP replicates the SVM configuration, including LIF IP addresses, export policies, and NFS settings, so the secondary SVM presents the same NFS endpoint after a failover.
- Failover is non-disruptive for NFSv3 and later.
- SnapMirror active sync for NAS is active/passive, so VMs on site B nodes read and write across the inter-site network to the primary SVM. For active workloads at both sites, use two SVMs with opposite replication directions, one primary at each site, and align each VM's node affinity rule with the site of its primary SVM.
- SnapMirror active sync for NAS is supported on AFF A-Series and C-Series two-node clusters.

NFS client settings

Use NFSv3 or NFSv4.1 for the NFS storage. Keep the default hard mount behavior, and don't add the `soft` mount option, so that NFS I/O retries instead of failing during a failover.

Use multiple TCP connections to the storage for better performance and fault tolerance. Use either `nConnect` or session trunking, not both:

- **nConnect.** Opens multiple TCP connections to a single NFS LIF address, for example `nconnect=4`. `nConnect` works with NFSv3 and NFSv4.1 and with a single VIP LIF address, so it's the recommended option for SnapMirror active sync for NAS, including layer 3 networks.
- **Session trunking.** Groups multiple data LIFs into one NFSv4.1 session, for example `vers=4.1, trunkdiscovery, max_connect=16`. Session trunking requires NFSv4.1 or later. Use it only when the protected SVM has multiple data LIFs that clients can reach at both sites, such as in a stretched layer 2 network. For details, see [NFS trunking](#).

For the mount option commands and verification steps, see [Configure NFS storage for Proxmox VE](#).

NFS storage provisioning with SMP

SMP provisions an ONTAP volume, creates or selects an export policy, and mounts the volume as NFS shared storage on the Proxmox VE nodes that you select. For storage in a protected SVM:

- Provision the volume in the primary SVM, and select the NFS LIF, or the VIP LIF in a layer 3 network, as the preferred LIF. The secondary SVM uses the same address after a failover, so the storage definition stays valid.
- Mount the storage on all Proxmox VE nodes at both sites so that HA can restart VMs on any node.
- Make sure the export policy allows the Proxmox VE nodes at both sites. ONTAP replicates the export policy to the secondary SVM.
- After provisioning, verify that the volume is included in SnapMirror active sync protection.


Layer 3 networks with VIP LIFs and BGP

SnapMirror active sync for NAS requires the same IP address on the primary and secondary SVMs. When the NFS VLAN is stretched between sites, the LIF IP address moves with the SVM. When the sites are connected by a routed layer 3 network, the subnet of a site A LIF doesn't exist at site B.




At the time of publication, NetApp hasn't documented VIP LIFs with BGP as a qualified configuration for SnapMirror active sync for NAS. Before you plan a layer 3 design, confirm current support with NetApp. Otherwise, stretch the NFS VLAN between the sites. MetroCluster IP configurations support VIP LIFs with BGP.

ONTAP virtual IP (VIP) LIFs address this problem. A VIP LIF doesn't belong to any subnet. ONTAP advertises the VIP address to the peer routers as a /32 host route by using Border Gateway Protocol (BGP), so the address can move between sites without a stretched VLAN.



1. Each node at both sites has BGP LIFs in the system SVM that peer with the routers at its site. BGP LIFs are node specific and aren't replicated, so you configure them on both clusters.
2. The primary SVM hosts the NFS VIP LIF. Cluster A advertises the /32 host route, and the Proxmox VE nodes at both sites reach the storage through the routed network.
3. After a SnapMirror active sync failover, the VIP LIF comes up on the secondary SVM, cluster B advertises the host route, and the routers converge on site B.

When you plan BGP in ONTAP for the first time, it helps to know which SVM owns each LIF. BGP LIFs belong to the system SVM of the IPspace, not to the data SVMs. One set of BGP LIFs on each cluster serves the VIP LIFs of every data SVM in that IPspace, so you don't create BGP LIFs for each tenant. The following diagram shows the LIF and port layout on one cluster, with BGP LIFs on VLAN ports.



1. Each node has two ports for BGP, in this example the VLAN ports e0c-100 and e0d-200. Each port connects directly to a different router.
2. Each node has two BGP LIFs in the system SVM, one on each port, with a /31 subnet and a BGP peer group to one router. For the Default IPspace, the system SVM is the cluster SVM. For a custom IPspace, it's the SVM with the same name as the IPspace.
3. Data SVMs in the same IPspace host VIP LIFs and don't need BGP LIFs of their own. The BGP peer groups advertise the VIP routes for all SVMs in the IPspace.
4. ONTAP creates a VIP port, for example v0a, on each node when you create the first BGP peer group. VIP LIFs are hosted on the VIP port and are reachable through every port on the node that hosts a BGP LIF in the IPspace.

The home port of a BGP LIF can be a physical port, a VLAN port, an interface group, or a VLAN on an interface group. An interface group isn't required. NetApp Technical Report TR-4949 recommends two BGP LIFs on two separate ports for each IPspace, each directly connected to a different first-hop router.

Keep the following guidelines in mind:

- BGP LIFs don't fail over. Redundancy comes from two BGP LIFs per node, each peering with a different router. If a node loses all its BGP sessions for an SVM, ONTAP moves the VIP LIF to another node that has an active BGP session for that SVM.
- Use external BGP (eBGP) with a unique private autonomous system number (ASN) for each ONTAP node and each router.
- Use a /31 subnet (netmask 255.255.255.254) for each BGP LIF and its router. Assign the odd-numbered address to the BGP LIF and the even-numbered address to the router.
- ONTAP doesn't process routes from the routers. Configure the routers to accept /32 host routes from ONTAP and not to send route updates to the cluster.
- Proxmox VE nodes must not share a subnet with the VIP address. A redundant layer 3 router, for example with VRRP or HSRP, must be in the path between the nodes and the storage.
- Configure BGP peer groups with `-use-peer-as-next-hop true` so that each cluster learns its default route from its local routers. Routes that point to a site A gateway aren't valid at site B after a failover.

- Enable multipath routing so that outgoing VIP traffic uses all available BGP LIF ports.
- Lower the BGP hold time from the default of 180 seconds. TR-4949 recommends a hold time of 3 seconds. A router withdraws the route from a failed cluster only when the hold timer expires, and during a storage system failure the router at site A stays online.

Related information

- [ONTAP SnapMirror active sync architecture](#)
- [Prerequisites for ONTAP SnapMirror active sync](#)
- [Prepare to configure ONTAP Cloud Mediator](#)
- [Configure the ONTAP Cloud Mediator for SnapMirror active sync](#)
- [ONTAP MetroCluster considerations for virtual IP and BGP](#)
- [MetroCluster data protection and disaster recovery](#)
- [Configure and use Storage Manager for Proxmox](#)
- [TR-4949: BGP/VIP with ONTAP in the Datacenter](#)
- [Proxmox VE cluster manager](#)
- [Proxmox VE high availability](#)

Configure Proxmox VE business continuity with NAS storage

Protect Proxmox Virtual Environment (VE) VMs from a site failure by combining ONTAP SnapMirror active sync for NAS with a Proxmox VE stretch cluster. You configure a mediator, either ONTAP Mediator or ONTAP Cloud Mediator in the NetApp Console, protect the NFS storage VM (SVM), optionally make its NFS address portable across a routed network with a virtual IP (VIP) LIF and Border Gateway Protocol (BGP), build the stretch cluster with a QDevice, and then provision the shared storage and configure high availability (HA) and VM backups with Storage Manager for Proxmox. The result is zero data loss and automatic recovery of VMs at the surviving site.

About this workflow

Before you start, review the [solution architecture, failure behavior, and corosync design considerations](#). Replace the values in angle brackets, such as <primary_svm>, with the values for your environment.

Before you begin

- Two AFF A-Series or C-Series two-node clusters must run ONTAP 9.19.1 or later, one at site A and one at site B.
- SnapMirror and SnapMirror synchronous licenses, included in ONTAP One, must be installed on both clusters.
- The clusters must be peered in the default IPspace, with an inter-cluster round-trip time (RTT) under 10 ms.
- You must choose one mediator option. For a comparison, see [Mediator options](#).
 - For ONTAP Mediator, a Linux host must be installed at site C, or at another third failure domain, and reachable from the cluster and node management LIFs of both clusters.

- For ONTAP Cloud Mediator, both clusters must have outbound HTTPS access to the NetApp Console, directly or through a proxy, with a ping latency under 200 ms. You also need a NetApp Console organization and permission to create a service account.
- An SVM at site A must serve NFS to the Proxmox VE nodes. The SVM must contain only FlexVol volumes for NAS, with no LUNs, and its export policy must allow all Proxmox VE nodes.
- An even number of Proxmox VE 9.0 or later nodes must be installed, split evenly between site A and site B, with a corosync latency under 5 ms between all nodes.
- A Linux host at site C, or at another third location, must be available for the corosync QDevice daemon, reachable from all Proxmox VE nodes on TCP port 5403. You need this host with either mediator option.
- The NFS VLAN must be stretched between site A and site B. For a routed layer 3 network, a BGP-capable router must be available at each site, and you must confirm with NetApp that VIP LIFs with BGP are supported with SnapMirror active sync for NAS.
- Your ONTAP and Proxmox VE versions must be listed together in the [NetApp Interoperability Matrix Tool \(IMT\)](#).
- To provision shared storage and back up VMs, Storage Manager for Proxmox must be installed with network access to the Proxmox VE nodes and to both ONTAP clusters.

Step 1: Configure the mediator

Add a mediator to the cluster peer relationship before you initialize the SnapMirror active sync relationship. ONTAP needs the mediator to perform planned and automatic failovers. You can add the mediator from either cluster, and ONTAP adds it on the other cluster automatically. Use one mediator option, not both.

Choose one of the following options.

ONTAP Mediator

Use this option if you install the mediator on a Linux host at site C. For installation, see [Install or upgrade ONTAP Mediator](#).

Steps

1. On either cluster, add the mediator:

```
snapmirror mediator add -peer-cluster <peer_cluster> -type on-prem
-mediator-address <mediator_ip> -username <mediator_user>
```

Enter the mediator password when prompted.

2. Verify that the mediator is connected and that the quorum status is `true`:

```
snapmirror mediator show
```

ONTAP Cloud Mediator

Use this option to host the mediator as a service in the NetApp Console, so that you don't need a mediator host at a third site.

Before you begin

- The firewalls at both sites must allow HTTPS traffic from both clusters to the NetApp Console API endpoint. For the current endpoint hostname, see [Prepare to configure ONTAP Cloud Mediator](#).
- If the clusters use proxy servers, have the proxy IP address, port, and credentials available. Use a different proxy server for each cluster. The proxy's root CA certificate must be installed in ONTAP.
- ONTAP includes well-known root CA certificates, so you usually don't need to install the NetApp Console certificate. To check, run `security certificate show -common-name Sectigo*|USERTrust*|AAA*`.

Steps

1. In the NetApp Console, copy your organization ID.
2. In the NetApp Console, create a member service account:
 - a. Set the organization to the one where you use ONTAP Cloud Mediator.
 - b. Set the category to **Application**.
 - c. Set the role to **ONTAP Mediator Setup Role**.
 - d. Save the client ID and client secret when you create the account.
3. On either cluster, add the mediator:

```
snapmirror mediator add -peer-cluster <peer_cluster> -type cloud
-org-id <console_org_id> -service-account-client-id <client_id> -use
-http-proxy-local <true|false> -use-http-proxy-remote <true|false>
```

Enter the client secret when prompted.

You can also use System Manager. Select **Protection > Overview > Mediator > Add**, select **Cloud** as the mediator type, and enter the organization ID, client ID, and client secret.

4. Verify that the mediator type is `cloud`, the connection status is `connected`, and the quorum status is `true`:

```
snapmirror mediator show
```

Step 2: Protect the NFS SVM with SnapMirror active sync

Create a dormant destination SVM on the site B cluster and replicate the site A SVM to it. ONTAP replicates the data synchronously and replicates the SVM configuration, including the NFS LIF addresses and export policies. You can also complete this step in System Manager by selecting **Protect** for the SVM and choosing the `AutomatedFailOver` policy.

Steps

1. On the site B cluster, create the destination SVM:

```
vserver create -vserver <secondary_svm> -sub-type dp-destination
```

2. On the site B cluster, create an SVM peer relationship with the primary SVM:

```
vserver peer create -vserver <secondary_svm> -peer-vserver <primary_svm>
-peer-cluster <site_a_cluster> -applications snapmirror
```

3. On the site A cluster, accept the SVM peer request:

```
vserver peer accept -vserver <primary_svm> -peer-vserver <secondary_svm>
```

4. On the site B cluster, create the SnapMirror active sync relationship:

```
snapmirror create -source-path <primary_svm>: -destination-path
<secondary_svm>: -policy AutomatedFailOver
```

5. Initialize the relationship:

```
snapmirror initialize -destination-path <secondary_svm>:
```

6. Verify that the relationship is in sync:

```
snapmirror show -destination-path <secondary_svm>:
```

The Mirror State is Snapmirrored, and the Relationship Status is InSync.

7. Optionally, exclude volumes that don't need protection, such as scratch or test volumes:

```
volume modify -vserver <primary_svm> -volume <volume_name> -smas
-protection unprotected
```

Step 3: Configure a VIP LIF with BGP for layer 3 networks



At the time of publication, NetApp hasn't documented VIP LIFs with BGP as a qualified configuration for SnapMirror active sync for NAS. Confirm current support with NetApp before you complete this step. Otherwise, stretch the NFS VLAN between the sites and skip this step.

If the NFS VLAN is stretched between site A and site B, skip this step. If the sites are connected by a routed network, create BGP LIFs on every node of both clusters, and then host the NFS address on a VIP LIF in the primary SVM. The VIP address must not belong to the subnet of the Proxmox VE nodes.

Before you begin

- Each router must be configured to accept a BGP session from the BGP LIFs at its site, to accept /32 host routes, and not to send route updates to ONTAP.

- Each ONTAP node and each router must have a unique private ASN, and each BGP LIF and router pair must have a /31 subnet.

Steps

1. On each cluster, change to the advanced privilege level:

```
set -privilege advanced
```

2. On each node of both clusters, create a BGP configuration with a unique ONTAP ASN for the node, a hold time of 3 seconds, and the node management IP address as the router ID:

```
network bgp config create -node <node_name> -asn <node_asn> -holdtime 3  
-routerid <node_mgmt_ip>
```

3. Optionally, on each node of both clusters, create a VLAN port on each of the two ports that connect to the routers:

```
network port vlan create -node <node_name> -vlan-name <port_name>-  
<vlan_id>
```

You can also use physical ports, interface groups, or VLANs on interface groups as BGP LIF home ports. The ports must belong to a broadcast domain in the IPspace.

4. On each node of both clusters, create two BGP LIFs in the system SVM, one on each port, with a /31 subnet to the router on that port:

```
network interface create -vserver <system_svm> -lif <bgp_lif_name>  
-service-policy default-route-announce -home-node <node_name> -home-port  
<vlan_or_port> -address <bgp_lif_ip> -netmask-length 31
```

For the Default IPspace, <system_svm> is the cluster name. For a custom IPspace, it's the IPspace name. You create the BGP LIFs once for each node, and they serve the VIP LIFs of all data SVMs in the IPspace.

5. For each BGP LIF, create a BGP peer group with its router, and learn the default route from that router:

```
network bgp peer-group create -peer-group <peer_group_name> -ipspace  
<ipspace> -bgp-lif <bgp_lif_name> -peer-address <router_ip> -peer-asn  
<router_asn> -use-peer-as-next-hop true
```

6. Verify that each peer group is up:

```
network bgp peer-group show -ipspace <ipspace>
```

7. On each cluster, enable multipath routing:

```
network options multipath-routing modify -is-enabled true
```

8. On the site A cluster, create the NFS VIP LIF in the primary SVM:

```
network interface create -vserver <primary_svm> -lif <vip_lif_name>  
-role data -data-protocol nfs -home-node <node_name> -address  
<vip_address> -is-vip true -failover-policy broadcast-domain-wide
```

9. On the site A cluster, verify that the BGP status is `up` for the primary SVM on every node:

```
network bgp vservers-status show
```

10. On both routers, verify that the /32 route for the VIP address is learned from cluster A.

Step 4: Create the Proxmox VE stretch cluster with a QDevice

Create the cluster from a site A node so that site A nodes receive the lowest node IDs, and add a QDevice at site C. The QDevice gives the vote to site A when the inter-site link fails, which matches the SVM that ONTAP keeps active.

Steps

1. On the first site A node, create the cluster with two corosync links on separate networks:

```
pvecm create <cluster_name> --link0 <node_link0_ip> --link1  
<node_link1_ip>
```

2. On each remaining site A node, and then on each site B node, join the cluster:

```
pvecm add <first_node_ip> --link0 <node_link0_ip> --link1  
<node_link1_ip>
```

3. Verify that the site A nodes have the lowest node IDs:

```
pvecm nodes
```

4. On the QDevice host at site C or another third location, install the QDevice daemon:

```
apt install corosync-qnetd
```

5. On every Proxmox VE node, install the QDevice client:

```
apt install corosync-qdevice
```

6. On one Proxmox VE node, add the QDevice to the cluster:

```
pvecm qdevice setup <qdevice_ip>
```

7. Verify that the cluster is quorate with the QDevice vote:

```
pvecm status
```

For a four-node cluster, Expected votes is 5, Quorum is 3, and the Flags include Quorate Qdevice.

Step 5: Provision protected NFS shared storage

Provision the NFS storage in the primary SVM and mount it on all Proxmox VE nodes at both sites, using the NFS LIF or VIP address that both SVMs share. Storage Manager for Proxmox (SMP) creates the ONTAP volume, the export policy, and the Proxmox VE storage definition in one workflow. If you don't use SMP, create the volume in ONTAP and add it as NFS storage with the Proxmox VE CLI.

Choose one of the following methods to provision the storage.

Storage Manager for Proxmox

Use SMP to provision the storage. SMP must be installed, with the Proxmox VE cluster and both ONTAP clusters added as plans. For details, see [Configure and use Storage Manager for Proxmox](#).

1. In SMP, select **Provision > Add Volume**, and then select **NFS**.
2. In **ONTAP Details**, select the primary SVM, enter the volume name and size, and select the NFS LIF or VIP LIF as the preferred LIF.
3. Select an export policy that allows the Proxmox VE nodes at both sites, or create one.
4. In **Proxmox Details**, select the cluster and all nodes at both sites, select NFS v3 or NFS v4.1, and select the content types.
5. Make sure **Mount the datastore** is selected, and then select **Provision**.
6. On the site A cluster, verify that the volume is included in SnapMirror active sync protection:

```
volume show -vserver <primary_svm> -volume <volume_name> -fields  
smas-protection
```

CLI

Use the Proxmox VE CLI to add a volume that already exists in the primary SVM as NFS storage.

1. On one Proxmox VE node, add the NFS storage:

```
pvesm add nfs <storage_id> --server <nfs_lif_or_vip_address>
--export <export_path> --content images,rootdir --options
vers=<nfs_version>,nconnect=4
```

Use 3 or 4.1 for <nfs_version>. Don't add the `soft` mount option.

2. Verify that the storage is active on every node:

```
pvesm status --storage <storage_id>
```

Set the NFS mount options

Use `nConnect` to open multiple TCP connections to the NFS LIF or VIP address for better performance and fault tolerance. If the protected SVM has multiple data LIFs that clients can reach at both sites, you can use NFSv4.1 session trunking instead. Use either `nConnect` or session trunking, not both. For details, see [Configure NFS storage for Proxmox VE](#).

Steps

1. If you provisioned the storage with SMP, or without `nConnect`, set the mount options on one Proxmox VE node:

```
pvesm set <storage_id> --options vers=<nfs_version>,nconnect=4
```

To use session trunking instead, set:

```
pvesm set <storage_id> --options vers=4.1,trunkdiscovery,max_connect=16
```

Proxmox VE applies new mount options the next time the storage is mounted. Remount the storage on each node during a maintenance window, after you migrate the VMs off that node.

2. Verify the options in `/etc/pve/storage.cfg`.
3. On each Proxmox VE node, verify that there are multiple connections to the NFS server address:

```
ss -an | grep :2049
```

Step 6: Configure HA and VM backups

Add the VMs to HA with node affinity rules that prefer site A, where the primary SVM serves data. Then protect the VMs from logical failures, such as accidental deletion or ransomware, with SMP backup policies. SnapMirror active sync replicates those failures synchronously, so it doesn't replace point-in-time backups.

Steps

1. Add each protected VM to HA:

```
ha-manager add vm:<vm_id> --state started
```

2. Create a node affinity rule that prefers site A nodes and allows failover to site B nodes:

```
ha-manager rules add node-affinity <rule_name> --resources  
vm:<vm_id>,vm:<vm_id> --nodes  
"<site_a_node1>:2,<site_a_node2>:2,<site_b_node1>:1,<site_b_node2>:1"
```

3. Optionally, increase the disk timeout in each guest operating system so that it's longer than the expected storage failover time.
4. In SMP, select **Protect > Backup > Resource Group**, and create a resource group for the VMs or storage of each tenant or application.
5. Select **Protect > Backup > Policies > Add**, and create a policy with the backup frequency and retention for the resource group.

Select **VM Consistency** in **Advanced Options** for VMs that run the QEMU guest agent.

6. Run an on-demand backup, and then restore a test VM from it to verify the policy.

Step 7: Test planned failover

Validate the design before you put production VMs on the shared storage. A planned failover moves the active SVM to site B without data loss, so you can test it during a maintenance window and fail back afterward.

Steps

1. Live-migrate the test VMs to site B nodes.
2. Before the failover, prevent HA from moving the VMs back to site A nodes:

```
ha-manager set vm:<vm_id> --failback 0
```

3. On the site B cluster, start a planned failover:

```
snapmirror failover start -destination-path <secondary_svm>:
```

4. Monitor the failover until it completes:

```
snapmirror failover show
```

5. Verify that VM I/O continues and that the storage stays active on all Proxmox VE nodes.
6. If you use a VIP LIF, verify that both routers learn the /32 route from cluster B.
7. In SMP, run a discovery for the ONTAP plans, and then run an on-demand backup to verify that SMP protects the VMs from the SVM at site B.

8. Fail back to site A when the test is complete, and then set `failback` back to 1 for the test VMs.

Result

The Proxmox VE VMs keep running on the same NFS shared storage while the active SVM moves between sites, and the cluster recovers automatically from an unplanned site failure.

Related information

- [Protect with ONTAP SnapMirror active sync for NAS](#)
- [Configure ONTAP VIP LIFs](#)
- [Prepare to install ONTAP Mediator](#)
- [Prepare to configure ONTAP Cloud Mediator](#)
- [Configure the ONTAP Cloud Mediator for SnapMirror active sync](#)
- [snapmirror mediator add](#)
- [Configure NFS storage for Proxmox VE](#)
- [Proxmox VE corosync external vote support](#)

Configure and use Storage Manager for Proxmox

Storage Manager for Proxmox (SMP) provides comprehensive storage provisioning and lifecycle management for virtual machines in Proxmox environments using NetApp ONTAP storage systems. SMP enables administrators to manage storage directly, simplifying operations through automated provisioning, backup, restore, and cloning capabilities.

Learn about Storage Manager for Proxmox

Storage Manager for Proxmox (SMP) offers end-to-end storage management for virtual machines within Proxmox environments that utilize NetApp storage systems. By allowing administrators to directly manage storage, SMP simplifies both storage and data management tasks in Proxmox environments.

Key features

- **End-to-end storage management:** SMP provides complete lifecycle management for virtual machines, from provisioning storage to backup and recovery.
- **Direct storage management:** Proxmox administrators can manage storage directly, streamlining operations and reducing complexity.
- **Standalone installer:** SMP is delivered as a standalone installer that seamlessly connects to Proxmox nodes via APIs.

Capabilities

Each component of SMP is designed to enhance storage management efficiency:

- **Storage provisioning:** Easily allocate and manage storage resources for Proxmox hosts and associated virtual machines.
 - **NFS storage provisioning:** Set up and configure ONTAP NFS storage to be used with your Proxmox

environment.

- **iSCSI storage provisioning:** Implement thick LVM with multipath for robust and efficient iSCSI storage management.
- **Backup and recovery:** Perform reliable backups and swift recoveries to ensure data integrity and availability.
 - **Crash-consistent and file-consistent backups:** Ensure data integrity by setting up both crash-consistent and file-consistent backup mechanisms.
 - **VM restore operations:** Execute virtual machine restore processes to recover data and maintain system availability.
- **Storage efficiency:** ONTAP provides a comprehensive suite of storage efficiency technologies designed to minimize the storage required for data. These technologies not only reduce storage consumption but also enhance system performance, delivering:
 - **Rapid VM provisioning:** Streamline the provisioning of virtual machines for quicker deployment.
 - **Quicker Dataset Copies:** Speed up the duplication of datasets for efficient data management.

Prerequisites

Ensure you have completed the following before proceeding:

- Working Proxmox Virtual Environment (PVE) 9.x or later cluster.
- ONTAP All Flash FAS (AFF) storage with NFS and iSCSI protocol enabled at the SVM level. For instructions, see [Configure NFS datastores with ONTAP](#) and [Configure Proxmox VE 9.x for iSCSI with ONTAP storage](#).

For iSCSI LUN configuration, the following should be enabled:

Install multipath and iSCSI on all Proxmox hosts:

```
apt update
apt install open-iscsi multipath-tools -y
systemctl enable --now iscsid multipathd
```

Configure iSCSI initiator and connect via iSCSI (on Proxmox hosts):



Create or update /etc/multipath.conf:

```
touch /etc/multipath.conf
```

Or use:

```
cat << 'EOF' > /etc/multipath.conf
defaults {
    user_friendly_names yes
    find_multipaths yes
}
```

```
blacklist {  
    devnode "^sd[a]$"   
}  
EOF
```

- Proxmox hosts and ONTAP storage are network reachable from the SMP host.


Install and set up Storage Manager for Proxmox

To get started with SMP, use a Windows operating system on a designated virtual machine. SMP can be run on Windows Server 2019, 2022, and 2025.

Steps

1. Download the SMP package (~17MB) from the [NetApp Console Automation hub](#).

▼ View SMP download page




2. Unzip the package.
3. Run the executable file to install and start the service.



No extra package dependencies are required because SMP leverages the available APIs to orchestrate both provisioning and data protection workflows.

▼ View SMP installation wizard



Alternative installation using Shift toolkit


As an alternative to the standalone installer, you can install Proxmox VE Storage Manager using Shift toolkit:

- Open Shift Toolkit.
- Navigate to **Settings > Advanced > Storage Manager**.
- Enable Proxmox VE using the toggle switch.
- Wait for the system to automatically install the necessary components.
- After installation is complete, access Storage Manager by clicking **View** or by opening a web browser to https://<Shift_server_IP>:4001


For example: <https://10.61.187.117:4001>

4. Open a web browser and access the SMP UI by entering <http://<IP address specified during installation>:4001>

▼ View SMP login screen



▼ View SMP home screen



Configure Storage Manager for Proxmox

When the storage and connectivity to both the storage and Proxmox nodes have been configured properly, configure SMP to automate the provisioning and protection of virtual machines leveraging ONTAP Snapshot and FlexClone functionalities.

The first step is to discover and add the ONTAP storage endpoint and then the target Proxmox nodes to SMP.

Before you begin

Ensure the following prerequisites are met:

- SMP is installed
- SMP instance has connectivity to the Proxmox nodes and storage systems
- ONTAP cluster to provide NFS or iSCSI storage pools
- DNS resolution is in place when using DNS names. Otherwise, use IP addresses for discovery.
- The environment has supported versions of Proxmox Virtual Environment servers


About this task

After the connectivity is established between the hypervisor and storage, configuration takes a couple of clicks and about 2 minutes.

Steps

1. Log in to SMP.
2. From the SMP left navigation, select **Plan**.
3. Click **Add Plan**.

▼ View Add Plan




4. Click **Add ONTAP** to add the credentials of an ONTAP storage system.

You can add ONTAP using either the cluster endpoint or the SVM management LIF.

5. To add using the cluster endpoint, provide the following details:
 - a. Cluster IP address or FQDN
 - b. Username to access the ONTAP cluster
 - c. Password to access the ONTAP cluster
 - d. Select **Accept self-signed certificate** and click **Add Storage Cluster**

▼ View Add ONTAP: cluster endpoint



6. To add using the SVM management LIF, provide the following details:
 - a. SVM management LIF or FQDN
 - b. Username to access the SVM
 - c. Password to access the SVM
 - d. Select **Accept self-signed certificate** and click **Add Storage Cluster**

▼ **View Add ONTAP: SVM management LIF**

The screenshot shows the 'ONTAP SVM Details' configuration window in the NetApp Storage Manager for Proxmox. The 'Connection Type' is set to 'SVM'. The form contains the following fields: SVM Name (SMPProxSVM), SVM Management LIF IP (10.61.187.155), Admin Username (vasadmin), and Password (masked with dots). The 'Accept self-signed certificates' checkbox is checked. The 'Add Storage SVM' button is highlighted in red.

7. Click **Add Proxmox** to add the credentials of a Proxmox node to be registered with SMP.

To add Proxmox nodes, specify the IP address of one of the existing Proxmox nodes. The discovery job automatically queries the cluster to gather all essential configuration details.

8. Provide the following details:
 - a. Proxmox standalone IP address or FQDN
 - b. Username to access Proxmox
 - c. Password to access Proxmox host for performing inventory of the resources
 - d. Select **Accept self-signed certificate** and click **Continue**

▼ **View Add Proxmox: node credentials**

NetApp Storage Manager for Proxmox

Dashboard

Plan

Provision

Protect

Backup

Restore

Clone

Preserve

Process Logs

Proxmox Details

Node

Plan Name

SMPProxmox-Username

Proxmox Endpoint

10.193.172.225

Authentication Method

☒ Username / Password ☐ API Token

Proxmox Username

root

Proxmox Password

NetApp123

☒ Accept self-signed certificates

Cancel Add Site

Alternatively, use an API token to access the APIs by providing:

- API token ID
- API token secret



In this release, API token is only allowed for NFS operations.

▼ View Add Proxmox: API token

NetApp Storage Manager for Proxmox

Dashboard

Plan

Provision

Protect

Backup

Restore

Clone

Preserve

Process Logs

Proxmox Details

Node

Plan Name

SMPProxmox-APIToken

Proxmox Endpoint

10.193.172.221

Authentication Method

☐ Username / Password ☒ API Token

API Token ID

root@pam1My_Token_123

API Token Secret


☒ Accept self-signed certificates

Cancel Add Site

9. Verify that the discovery jobs are successful and the status for each plan is marked as **Success**.

Once added, SMP automatically performs a discovery scan and displays the VMs along with their associated metadata. Any existing storage volumes or mapped LUNs are also detected.


▼ View Plan discovery results



After you finish

- If any modifications are made to the Proxmox nodes or storage, run the discovery to fetch the latest information. Click the three dots against the plan name and click **Discover**.
- Ensure the ONTAP storage system is added first, so the discovery job can map existing resources accordingly. If Proxmox is added first, run a discovery job against ONTAP.
- If the credential or API token requires an update later, navigate to **Settings** and select the **Update Credentials** option.

▼ View Update Credentials



Provision storage

Provisioning a volume as a storage pool, commonly referred to as a datastore, creates a logical container for virtual machines and their virtual disks (QCOW2 or RAW). SMP enables provisioning a datastore and then attaching it either to a single host or to all hosts within a cluster.


Provision NFS storage

Use SMP to provision and mount an ONTAP volume as an NFS datastore.

Steps

1. In the SMP GUI, go to the **Provision** tab.
2. Navigate to **Add Volume**.

▼ View Provision tab: Add Volume



3. In the New Volume wizard, on the Quick Provision step, select **NFS**.



NFS is selected by default. Quick provisioning delivers a faster setup using default settings for standard deployments. Advanced provisioning offers expanded options for HA and BCDR capabilities, which will be introduced in future releases.

▼ View New Volume wizard: NFS

NetApp Storage Manager for Proxmox

Provision Storage

Quick Provision Advanced Provision

NFS iSCSI

ONTAP DETAILS

Storage Cluster: Select...

Volume Name: Size: GB

Storage Virtual Machine: Select... Aggregate Name: Select...

LIF: Select... Export Policy: Select...

PROXMOX DETAILS

Proxmox: Select... Nodes: Select...

NFS Version: 3

Content:

☐ Select All ☐ Disk image ☐ ISO image

☐ Container template ☐ Backup ☐ Container


☐ Snippets ☐ Import

☐ nconnect 4 ☒ Mount the datastore

Cancel Provision


4. In the **ONTAP Details** section, provide the following information:
 - a. Select the ONTAP cluster or the designated Storage Virtual Machine (SVM) from the **Storage Cluster** dropdown
 - b. Provide a unique value in the **Volume Name** field
 - c. Enter the required capacity for the datastore volume in the **Size** field
 - d. Choose ONTAP storage resources by selecting the appropriate options for:
 - Aggregate Name
 - Storage Virtual Machine (SVM)
 - Preferred LIF to be used for datastore access
 - e. Select an existing **Export Policy** or create a custom policy by clicking the (+) as needed

▼ **View ONTAP Details: NFS**




5. In the **Proxmox Details** section, provide the following information:
 - a. In the **Proxmox** field, choose the target cluster, then select all nodes that require access to the datastore
 - b. In the **NFS Version** field, select the preferred NFS protocol (NFS v3 or NFS v4.1)
 - c. In the **Content** section, select the content types to be enabled for this datastore:
 - Disk image
 - ISO image
 - Container template
 - Backup
 - Snippets
 - Import
 - Or select **All** content types
 - d. (Optional) Activate the **nconnect** option if multiple TCP connections are desired
 - e. Ensure **Mount the datastore** is checked so the datastore will be mounted across the selected Proxmox nodes

▼ View Proxmox Details: NFS



- Click **Provision** to create the ONTAP volume and register the datastore (storage pool) within the selected Proxmox cluster.

▼ View NFS provisioning result



NetApp Storage Manager for Proxmox Preview

Dashboard

Plan

Provision

Protect

Backup

Restore

Clone

Preserve

Process Logs

ONTAP Clusters: 1

Storage Virtual Machines: 21

Provisioned Volumes: 111

Storage: 2030

111 Provisioned Volumes

Search

Search... Add Volume

Storage Name	Protocol	Capacity	Mount Path	Proxmox Cluster	Mount Status	
sie_dev_vol_aakashgupta...	NFS	1075.20 GB	/mnt/pve/sie_dev_vol_aakash	SMPProxmox-Username	Mounted	...
sie_dev_vol_aakashgupta	NFS	1.00 GB	/mnt/pve/sie_dev_vol_aakash	SMPProxmox-ApiToken	Mounted	...
sie_dev_vol_aakashgupta	NFS	1075.20 GB	/mnt/pve/sie_vol_aka_tst_1	SMPProxmox-Username	Mounted	...
sie_dev_proxmox_varsinil_1	NFS	105.26 GB	/mnt/pve/varsinil_proxmox_vo	SMPProxmox-Username	Mounted	...
sie_dev_proxmox_varsinil	NFS	200.00 GB	/mnt/pve/varsinil_proxmox_vo	SMPProxmox-Username	Mounted	...
sie_dev_aka_test1	NFS	2.00 GB	/mnt/pve/sie_dev_aka_test1	SMPProxmox-ApiToken	Mounted	...
sie_dev_aka_demo_3	NFS	7.00 GB	/mnt/pve/sie_dev_aka_demo_3	SMPProxmox-Username	Mounted	...
sie_dev_aka_demo	NFS	5.00 GB	/mnt/pve/sie_dev_aka_demo	SMPProxmox-Username	Mounted	...
abcd9876	NFS	6.00 GB	/mnt/pve/abcd9876	SMPProxmox-ApiToken	Mounted	...
SMPDemoVol	NFS	2.00 GB	/mnt/pve/SMPDemoVol	SMPProxmox-ApiToken	Mounted	...

▼ View NFS datastore in Proxmox

PROXMOX Virtual Environment 9.1.1

Search

Documentation Create VM Create CT root@pam

Server View

Datacenter

Datacenter (nimravprox)

proxmox9-vm01

101 (nimrfs01)

102 (TVM-U18)

103 (ewaflead)

104 (testscsi)

107 (nim03)

112 (nimrfs06-88737)

113 (nimrfs01-1c20a)

Search

Summary

Notes

Cluster

Ceph

Options

Add Remove Edit

ID	Type	Content	Path/Target	Shared	Enabled	Bandwidth Limit
DNDdemo	NFS	Disk image	/mnt/pve/DNDdemo	Yes	Yes	
ISO_Proxmox	NFS	ISO image	/mnt/pve/ISO_Proxmox	Yes	Yes	
SMPDemoLVM	LVM	Disk image	/mnt/pve/SMPDemoLVM	Yes	Yes	
SMPDemoVol	NFS	Disk image	/mnt/pve/SMPDemoVol	Yes	Yes	
dnstst	NFS	Disk image	/mnt/pve/dnstst	Yes	Yes	

Edit a provisioned volume

Use SMP to modify volume settings after provisioning.

Steps

1. In the SMP GUI, locate the volume in the **Provision** tab.
2. Click the three vertical dots (:) next to the volume.
3. Select **Edit** from the context menu to open the volume settings.

▼ View Edit Volume dialog

4. Modify the volume settings as needed:

- **Resize volume:** Adjust the volume's storage capacity



Resizing is applicable to NFS volumes only.

- **Snapshot Locking:** Enable to prevent snapshots from being deleted or altered



This requires the SnapLock license to be installed on the cluster, and the cluster's compliance clock must be initialized.

- **Anti-Ransomware Protection:** Enable to proactively detect and alert on unusual activity that could signal a ransomware attack
- **Export policy:** Update or modify the export policy

5. Click **Save** to apply the changes.

Provision iSCSI storage

Use SMP to set up an iSCSI backend for virtual machines.

Steps

1. In the SMP GUI, go to the **Provision** tab.
2. Navigate to **Add Volume**.
3. In the New Volume wizard, on the Quick Provision step, select **iSCSI**.

▼ **View New Volume wizard: iSCSI**

NetApp Storage Manager for Proxmox

Provision Storage

Quick Provision | Advanced Provision

NFS | **iSCSI**

ONTAP DETAILS

Storage Cluster: Select...

Storage Virtual Machine: Select... | Aggregate Name: Select...

LUN Name: | Number of LUNs: 1

Capacity per LUN: | GB | Initiator Group: Select...

PROXMOX DETAILS

Proxmox: Select... | Nodes: Select...

LVM ID: | Storage VLAN: Select...

Content: ☐ Select All ☐ Disk image ☐ Container

☒ Mount the datastore ☐ Shared

Cancel | **Provision**

4. Under iSCSI protocol, in the **ONTAP Details** section, provide the following information:
 - a. Choose the appropriate ONTAP cluster or the designated Storage Virtual Machine (SVM) from the **Storage Cluster** dropdown
 - b. Select the options for:
 - Aggregate Name
 - Storage Virtual Machine (SVM)
 - c. Enter a unique value in the **LUN Name** field
 - d. Provide the required number of LUNs in the **Number of LUNs** field (default is 1; additional LUNs of the same size can be created in parallel)
 - e. Enter the desired size for each LUN in the **Capacity per LUN** field
 - f. Choose an **Initiator Group** (igroup) for iSCSI access control. An existing igroup can be selected or a new igroup can be created.

▼ **View ONTAP Details: iSCSI LUN configuration**

Provision Storage

Quick Provision | Advanced Provision

NFS | **ISCSI**

ONTAP DETAILS

Storage Cluster: SMPDemo212

Storage Virtual Machine: svm_sie_dev_func01

Aggregate Name: NSOL_NetApp_A400_T11U3_02_NV...

LUN Name: DemoISCSI

Number of LUNs: 1

Capacity per LUN: 2 GB

Initiator Group: **Custom**

PROXMOX DETAILS

Proxmox: Select...

Nodes: Select...

LVM ID: Select...

Storage VLAN: Select...

Content: ☐ Select All ☐ Disk image ☐ Container

☒ Mount the datastore ☐ Shared

Provision

▼ View iSCSI igroup selection

Provision Storage

Quick Provision | Advanced Provision

NFS | **ISCSI**

ONTAP DETAILS

Storage Cluster: SMPDemo212

Storage Virtual Machine: svm_sie_dev_func01

Aggregate Name: NSOL_NetApp_A400_T11U3_02_NV...

LUN Name: DemoISCSI

Number of LUNs: 1

Capacity per LUN: 2 GB

Initiator Group: **Custom**

PROXMOX DETAILS

Proxmox: Select...

Nodes: Select...

LVM ID: Select...

Storage VLAN: Select...

Content: ☐ Select All ☐ Disk image ☐ Container

☒ Mount the datastore ☐ Shared

Cancel | Provision

5. In the **Proxmox Details** section, provide the following information:
 - a. From the **Proxmox** dropdown, choose the target cluster, then select all nodes requiring access to the storage
 - b. Select the appropriate **LVM ID** that will represent the iSCSI-backed storage pool within Proxmox
 - c. In the **Nodes** field, choose all nodes that require access to the presented iSCSI targets
 - d. (Optional) If VLAN segmentation is required, choose the desired **Storage VLAN**
 - e. In the **Content** section, select the content types supported by this datastore:
 - Disk image
 - Container
 - Or select **All** content types


- f. Ensure **Mount the datastore** is selected so the storage is available across the chosen Proxmox nodes
- g. (Optional) Enable **Shared** if the storage pool will be accessed concurrently by multiple nodes

▼ View Proxmox Details: iSCSI


6. Click **Provision** to create the ONTAP LUN and register the datastore (storage pool) within the selected Proxmox cluster.

▼ View iSCSI LUN provisioning

▼ View iSCSI provisioning completed




▼ View iSCSI datastore in Proxmox



7. Validate that the NFS and iSCSI LUNs have been connected by accessing the Proxmox endpoint.

The attached storage is now available for provisioning virtual machines.

▼ View Proxmox storage view



Volume configuration options

When SMP provisions new ONTAP FlexVol volumes and LUNs, several default configuration values are automatically applied:

ONTAP Setting	Default Value
Snapshot reserve (-percent-snapshot-space)	0%
Fractional reserve (-fractional-reserve)	0%
Access time update (-atime-update)	Disabled
Minimum readahead (-min-readahead)	Disabled
Scheduled snapshots	None
Storage efficiency	Enabled
Volume guarantee	None (thin provisioned)
Volume autosize	grow_shrink
LUN space reservation	Disabled
LUN space allocation	Enabled

NConnect behavior

When using NFS storage with Proxmox VE, be aware of specific behavior in the Linux NFS client (as used by Debian, the base OS for Proxmox VE). The Linux kernel may reuse the mount parameters from the first time an NFS export is mounted on a node. If the same export is mounted again with different mount options, those new options may not be applied.

As a best practice, create and mount separate exports and use different NFS server IP addresses, such as distinct data LIFs. SMP automatically detects existing connections along with their associated nconnect values and applies the appropriate nconnect value to any new connections.

Protect virtual machines

When virtual machines run on ONTAP-backed NFS datastores, the most streamlined protection approach is to let Proxmox establish filesystem consistency, while ONTAP captures the actual backup point at the storage layer. ONTAP snapshots provide fast, space-efficient, point-in-time copies of the datastore volume and can be used later to recover entire datastores, individual files, or VM content.

SMP provides a simple orchestrator for crash-consistent and VM-consistent data protection and clone management. This software simplifies backup, restore, and clone lifecycles through the creation of backup policies that can be applied to multiple jobs. These policies can define schedule, retention, replication, and other capabilities. They allow optional selection of file-consistent snapshots; this approach uses QEMU's ability to quiesce I/O before taking an ONTAP snapshot.

SMP allows protecting an entire datastore, where one ONTAP NFS volume contains multiple VMs, groups of VMs that may span several datastores, or at an individual VM level.

Back up virtual machines

Once the ONTAP volumes have been provisioned, use the Backup option to list all the VMs running on the Proxmox cluster and perform crash-consistent or file-consistent backup of the virtual machines.

Steps

1. Access **Protect > Backup** from the left side navigation pane.


▼ View Protect Backup page

The screenshot shows the NetApp Storage Manager for Proxmox interface. The left sidebar contains navigation options: Dashboard, Plan, Provision, Protect, Backup (highlighted), Restore, Clone, Preserve, and Process Logs. The main content area displays a summary of VMs, Storage, Resource Groups, and Policies. Below this, a table lists 32 VMs. A red box highlights the table content.

VMID	Name	Backup	Backup Status	Last Backup Date	Proxmox Cluster	ONTAP Cluster
100	sie-dev-aka-tst-1-89a2e	🔄	-	-	SMPProxmox-Username	SMPProxClus
101	sie-dev-aka-tst-1-fc39f	🔄	-	-	SMPProxmox-Username	SMPProxClus
102	sie-dev-aka-tst-1-ce32c	🔄	-	-	SMPProxmox-Username	SMPProxClus
103	sie-dev-aka-tst-1	🔄	-	-	SMPProxmox-Username	SMPProxClus
104	sie-dev-aka-tst-1-305c7-47f84	🔄	-	-	SMPProxmox-Username	SMPProxClus
105	sie-dev-aka-tst-1-305c7-0754c	🔄	-	-	SMPProxmox-Username	SMPProxClus
106	sie-dev-aka-tst-1-305c7-15aa6	🔄	-	-	SMPProxmox-Username	SMPProxClus
107	sie-dev-aka-tst-1-fabfa-65967	🔄	-	-	SMPProxmox-Username	SMPProxClus

2. Select the virtual machine to back up and click **Backup**.

▼ View Backup dialog: VM selection



3. The backup workflow allows selection between crash-consistent and file-consistent protection.

File-consistent mode optionally uses QEMU fsfreeze to ensure in-guest filesystem quiescence before the snapshot is taken. If fsfreeze fails, the system falls back to using a Proxmox snapshot.



This fallback mechanism applies only to virtual machines stored on NFS volumes.

▼ View Backup dialog: consistency options

Start Time	End Time	Node	User name	Description	Status
Feb 25 02:55:28		proxmox9-...	root@pam	Shell	
Mar 10 08:51:01	Mar 10 08:51:01	proxmox9-...	root@pam	VM 116 - Delete Snapshot	OK
Mar 10 08:50:56	Mar 10 08:50:57	proxmox9-...	root@pam	VM 116 - Snapshot	OK

▼ View Backup progress

VMID	Name	Backup	Backup Status	Last Backup Date	Proxmox Cluster	ONTAP Clust
100	sie-dev-aka-tst-1-89a2e		Success	11/03/2026, 11:44:00	SMPProxmax-Username	SMPProxClus
101	sie-dev-aka-tst-1-fc39f		-	-	SMPProxmax-Username	SMPProxClus
102	sie-dev-aka-tst-1-ce32c		-	-	SMPProxmax-Username	SMPProxClus
103	sie-dev-aka-tst-1		-	-	SMPProxmax-Username	SMPProxClus
104	sie-dev-aka-tst-1-305c7-47f84		-	-	SMPProxmax-Username	SMPProxClus
105	sie-dev-aka-tst-1-305c7-0754c		-	-	SMPProxmax-Username	SMPProxClus
106	sie-dev-aka-tst-1-305c7-15aa6		-	-	SMPProxmax-Username	SMPProxClus
107	sie-dev-aka-tst-1-fabfa-65967		VM backup completed: 1 successful, 0 failed			



ONTAP snapshots are taken at the volume level, capturing all VM disks and configuration files within the NFS datastore in a single snapshot. When VM-consistent protection is selected, fsfreeze is issued only against the chosen VMs; all other VMs in the volume are included as crash-consistent within the same ONTAP snapshot.

Additional backup options

- To back up a single datastore, select the **Storage** tab and choose the **Backup** option.
- To back up multiple datastores as one unit or protect VM groups that span different datastores, use the **Resource Group** option. Resource groups allow datastores or VMs to be organized into logical protection sets.

▼ View Resource Group tab


The screenshot shows the NetApp Storage Manager for Proxmox interface. The left sidebar contains navigation options: Dashboard, Plan, Provision, Protect, Backup (selected), Restore, Clone, Preserve, and Process Logs. The main content area displays a summary of protection status for VMs, Storage, Resource Groups, and Policies. The 'Resource Groups' tab is selected, showing a table of VMs with their backup status and details.

VMID	Name	Backup	Backup Status	Last Backup Date	Proxmox Cluster	ONTAP Clust
100	sie-dev-aka-tst-1-89a2e	Success	Success	11/03/2026, 11:44:00	SMPProxmx-Username	SMPProxClus
101	sie-dev-aka-tst-1-fc39f	-	-	-	SMPProxmx-Username	SMPProxClus
102	sie-dev-aka-tst-1-ce32c	-	-	-	SMPProxmx-Username	SMPProxClus
103	sie-dev-aka-tst-1	-	-	-	SMPProxmx-Username	SMPProxClus
104	sie-dev-aka-tst-1-305c7-47f84	-	-	-	SMPProxmx-Username	SMPProxClus
105	sie-dev-aka-tst-1-305c7-0754c	-	-	-	SMPProxmx-Username	SMPProxClus
106	sie-dev-aka-tst-1-305c7-15aa6	-	-	-	SMPProxmx-Username	SMPProxClus
107	sie-dev-aka-tst-1-fabfa-65967	-	-	-	SMPProxmx-Username	SMPProxClus

Create a resource group

1. Open the **Resource Group** tab and select **Add**.
2. Enter a name for the new resource group.
3. Choose the resource type: **VM** or **Datastore**.
4. Select the appropriate **Protocol**.
5. Choose the specific VMs or datastores to include, based on the selected resource type.
6. Click **Create Resource Group** to complete the setup.

▼ View Create Resource Group dialog



Backups taken at the datastore (volume) level are always crash-consistent. A resource group can contain only one resource type: either VMs or datastores. For iSCSI-based storage, backups are always crash-consistent.

Backup workflow details

The high-level backup workflow includes:

1. **Define the protection scope:** Select whether the protection target is a datastore (ONTAP volume), a resource group (multiple volumes or VMs), or individual VM.
2. **Confirm the consistency method:** Choose how consistency is created before the ONTAP snapshot:
 - QEMU guest agent freeze/thaw for in-guest filesystem quiescing and Windows VSS integration.
3. **Dump the VM config files.**
4. **Take the ONTAP snapshot:** Capture an ONTAP snapshot of the datastore volume while the VMs are in a consistent state. This becomes the retained backup copy.
5. **Remove temporary Proxmox snapshots:** Delete the Proxmox snapshots after the ONTAP snapshot completes to prevent delta file growth.
6. **Apply retention and replication policies (using policies):** Apply the configured ONTAP retention policy and replicate the snapshot set if disaster recovery replication is enabled.

This approach minimizes storage space and network bandwidth by storing only changed information. Use Snapshot copies to make quick copies of the VM or datastore without affecting performance and then send them to a secondary system by using NetApp SnapMirror technology for longer-term, off-site data protection.


Automate VM backup tasks

Automate Proxmox VM backups by configuring a schedule. Creating regular backups is crucial for data protection and disaster recovery.

Steps

1. Go to **Protect > Backup > Policies** and click **Add** to add a new backup policy and schedule.

▼ View Backup Policies page



2. Configure general settings:
 - a. Enter a name for the policy
 - b. Select the **Policy Type** (Volume or Resource Group)
 - c. Choose the **Target** based on the selected type
 - d. Set the **Backup Frequency** (hourly, daily, weekly, or monthly)
 - e. Configure **Retention**, using either "Days to Keep" or "Number to Keep" (default retention value is 7)
 - f. Choose the **Replication Mode** (disabled in the current release):
 - Update SnapMirror after backup
 - Update SnapVault after backup



SnapMirror replication must be configured in ONTAP using the CLI or System Manager.

3. In **Advanced Options**, enable **VM Consistency** if Proxmox snapshots or fsfreeze should be used before triggering the ONTAP snapshot.
4. Click **Create Policy** to save the configuration.

▼ View Backup policy: schedule and retention

NetApp Storage Manager for Proxmox [Preview](#)

Dashboard
Plan
Provision
Protect
Backup
Restore
Clone
Preserve
Process Logs

Add Policy

Name
demoPolicy01

Policy Type
Volume

Select Volume
SMPDemoVol

Frequency
Hourly

Retention
Days to keep 7

Replication
☐ Update SnapMirror after backup
☐ Update SnapVault after backup

Advanced Options

▼ View Backup policy: replication and advanced options

NetApp Storage Manager for Proxmox [Preview](#)

Dashboard
Plan
Provision
Protect
Backup
Restore
Clone
Preserve
Process Logs

VMs

Total 32 Protected 1

Storage

Total 111 Protected 1

Resource Groups

Total 0 Protected 0

Policies

Total 1

1 Policies

VMs

Storage

Resource groups

Policies

Search...

Add

Policy Name	Type	Target	Schedule	Retention	Status	Last Backup
demoPolicy01	datastore	SMPDemoVol	hourly	7 days	Enabled	-

1 - 1 of 1 << < 1 > >>



A VM backup job can also be triggered immediately by clicking **Backup**, rather than waiting for a resource group's scheduled execution time.

Restore virtual machines or datastores

Use SMP to restore a Proxmox VM from backup.

Before you begin

Virtual machines must be powered off before a restore operation is performed.

Steps

1. Select the **Restore** option from the navigation pane on the left side of the web interface.
2. Select the virtual machine, storage, or resource group that needs to be restored.

▼ View Restore page

NetApp Storage Manager for Proxmox

Dashboard

Plan

Provision

Protect

Backup

Restore

Clone

Preserve

Process Logs

VMs: Total 32, Protected 1

Storage: Total 111, Protected 1

Resource Groups: Total 0, Protected 0

Policies: 1

26 VMs

VMs | Storage | Resource groups

Search...

VMID	Name	Proxmox Cluster	ONTAP Cluster	Protocol	Restore	Restore Status	Last I
100	sie-dev-aka-tst-1-89a2e	SMPProxmox-Username	SMPProxmox-Username	NFS		-	-
101	sie-dev-aka-tst-1-fc39f	SMPProxmox-Username	SMPProxmox-Username	NFS		-	-
102	sie-dev-aka-tst-1-ce32c	SMPProxmox-Username	SMPProxmox-Username	NFS		-	-
103	sie-dev-aka-tst-1	SMPProxmox-Username	SMPProxmox-Username	NFS		-	-
104	sie-dev-aka-tst-1-305c7-47f84	SMPProxmox-Username	SMPProxmox-Username	NFS		-	-
105	sie-dev-aka-tst-1-305c7-0754c	SMPProxmox-Username	SMPProxmox-Username	NFS		-	-
106	sie-dev-aka-tst-1-305c7-15aa6	SMPProxmox-Username	SMPProxmox-Username	NFS		-	-
107	sie-dev-aka-tst-1-fabfa-65967	SMPProxmox-Username	SMPProxmox-Username	NFS		-	-
108	sie-dev-aka-tst-1-1fbb1	SMPProxmox-Username	SMPProxmox-Username	NFS		-	-
111	sie-dev-aka-tst-1-350aa	SMPProxmox-Username	SMPProxmox-Username	NFS		-	-


3. Click the **Restore** icon.

Depending on the selection (whether VM, Storage, or Resource Group), the restore pop-up window will open.



Depending on the resource selected, the snapshots created for the respective components will be shown. For instance, snapshots created at the resource group level will only be shown when a resource group is selected. The same applies to storage. However, for virtual machines, all the snapshots created across the components will be displayed for selection.

▼ View Restore dialog: snapshot selection




4. Select the required options:


- Restore target:** Option to restore to same host or alternate host
- Host:** If alternate host is selected, the dropdown allows you to select any available host within the cluster
- Snapshot:** Choose the appropriate snapshot for point-in-time restore

5. Click **Restore**.

▼ View Restore in progress




▼ View Restore completed



When a storage-level restore is initiated, all virtual machines on that volume or LUN are powered off, and the system restores the entire storage entity as a whole.

▼ View Storage-level restore



Restore workflow details

The restore workflow varies depending on the underlying storage type and involves the following key steps:

1. **Define the restore scope:** Select whether the protection target is a datastore (ONTAP volume), a resource group (multiple volumes or VMs), or individual VM.
2. **Execute the restore operation:** The VM restore process differs based on storage protocol:

- **NFS-based restores:** When restoring a volume or datastore on NFS, SMP initiates a volume snap restore operation on the target volume. For VMs hosted on NFS volumes, SMP validates the restore parameters including VM ID, snapshot name, site IDs, and storage type. If the VM is running, SMP stops it by polling its status and issuing a stop command. The system then restores VM disks from the snapshot using a file-level restore operation.
- **iSCSI-based restores:** For iSCSI datastores or LUNs, SMP restores the LUN from the snapshot, verifies the restored device, performs an iSCSI rescan across Proxmox nodes, and ensures device accessibility. For VMs residing on iSCSI storage, SMP clones the snapshot LUNs, maps the cloned LUNs to the Proxmox host, and performs iSCSI discovery to import them. The restored storage is then added, and VM configurations are updated. VM disks are moved back to their original storage with progress tracked per VM. The process concludes with cleanup activities, including the removal of temporary storage and deletion of cloned LUNs.


Clone virtual machines

Use SMP to create clones of virtual machines for development, testing, or rapid provisioning.

Steps

1. Open the **Protect** section and select **Clone** from the left navigation menu.
2. From the VM list, locate the required VM and click **Clone** in the Actions column.
3. In the Clone VM dialog, choose the **Target Node** where the clone will be created.

▼ View Clone VM dialog



The screenshot shows the NetApp Storage Manager for Proxmox interface. On the left is a navigation menu with options: Dashboard, Plan, Provision, Protect (selected), Backup, Restore, Clone, Preserve, and Process Logs. The main area displays a summary of resources: 40 VMs, 113 Provisioned Volumes, 1 Resource Group, and 1 Policy. Below this is a table titled '27 VMs' with columns: VMID, Name, Proxmox Cluster, ONTAP Cluster, Protocol, and Actions. The table lists 11 VMs. In the 'Actions' column, the 'Clone' button for each VM is highlighted with a red box. The interface also includes a search bar and a 'Clone Lifecycle' tab.

VMID	Name	Proxmox Cluster	ONTAP Cluster	Protocol	Actions
100	sie-dev-aka-tst-1-89a2e	SMPProxmox-Username	SMPDemo212	NFS	Clone
101	sie-dev-aka-tst-1-fc39f	SMPProxmox-Username	SMPDemo212	NFS	Clone
102	sie-dev-aka-tst-1-ce32c	SMPProxmox-Username	SMPDemo212	NFS	Clone
103	sie-dev-aka-tst-1	SMPProxmox-Username	SMPDemo212	NFS	Clone
104	sie-dev-aka-tst-1-305c7-47f84	SMPProxmox-Username	SMPDemo212	NFS	Clone
105	sie-dev-aka-tst-1-305c7-0754c	SMPProxmox-Username	SMPDemo212	NFS	Clone
106	sie-dev-aka-tst-1-305c7-15aa6	SMPProxmox-Username	SMPDemo212	NFS	Clone
107	sie-dev-aka-tst-1-fabfa-65967	SMPProxmox-Username	SMPDemo212	NFS	Clone
108	sie-dev-aka-tst-1-1fb01	SMPProxmox-Username	SMPDemo212	NFS	Clone
111	sie-dev-aka-tst-1-350aa	SMPProxmox-Username	SMPDemo212	NFS	Clone

4. Select the **Snapshot** to use as the source for the clone.




In the current release, clone is performed against the active file system. In an upcoming release, the clone workflow will allow selection of a snapshot, in which case file clones are performed using a new volume created from the selected snapshot.

5. Specify the **Number of Copies**, **CPU cores**, **Sockets**, and **Memory** for the cloned VM.
6. (Optional) Enable **Power on after clone** if the cloned VM should start automatically after creation.

7. (Optional) Select **Create a new datastore** if the clone should be placed on a newly created datastore.
8. Under **Advanced Options**, enable **Network Configuration** if network settings should be applied automatically to the clone.
9. Click **Clone** to begin the cloning process.


The clone workflow provisions the new VM using ONTAP's rapid-cloning APIs and assigns a unique five-digit GUID to the cloned instance.

▼ View Clone operation



The cloned VM appears in both the Backup and Clone lists and can be protected using existing backup policies. To make the cloned VMs visible in these lists, run a discovery against the Proxmox and ONTAP plans.

▼ View Cloned VMs in lists



Clone workflow details

SMP automates rapid VM cloning by offloading the heavy lifting to ONTAP to perform storage-side clones instead of host-side file copies. Under the covers, SMP taps ONTAP FlexClone semantics (snapshot-based, writable, space-efficient clones) so clones are created quickly and consume little capacity initially; it avoids data reads/writes through the Proxmox host where possible.

This approach accelerates VM provisioning by pushing clone work into ONTAP, delivering fast, thin, storage-efficient clones suitable for Dev/Test, labs, and rapid refresh cycles.

By default, clone is performed against the active file system. The clone workflow also allows selection of a snapshot, in which case file clones are performed using a new volume created from the selected snapshot.

Monitor operations


Use the Process Logs and Dashboard views to monitor and track SMP operations.

Process logs

The Process Logs view provides a consolidated record of all operations performed through Storage Manager for Proxmox. Each entry displays the job status, job ID, action performed, start time, and duration. This makes it easy to track completed tasks, identify any failed or partial operations, and review the sequence of actions executed in the environment.

Filters and search options are available to quickly locate specific jobs or narrow the list by status.

▼ View Process Logs view




Dashboard

The Dashboard provides an at-a-glance overview of the entire Storage Manager for Proxmox environment. It summarizes key inventory such as the number of Proxmox clusters, hosts, datastores, ONTAP clusters, and protected virtual machines. Capacity usage, total storage volumes, protection policies, and resource groups are also displayed to give a clear picture of system utilization.

Cluster health is shown prominently, indicating the status of all connected nodes. The lower section of the dashboard highlights VM distribution (total, running, and stopped) and provides real-time Proxmox resource utilization for memory, storage, and CPU. This view helps quickly assess system health, identify resource pressures, and monitor overall operational state.

▼ View Dashboard



Summary

Storage Manager for Proxmox delivers a streamlined and unified way to manage protection, cloning, and recovery operations across Proxmox environments. By integrating directly with ONTAP, the solution brings the efficiency, reliability, and performance of enterprise-class storage into everyday VM workflows. ONTAP's space-efficient snapshots, rapid cloning capabilities, and built-in replication provide consistent and predictable data protection, while SMP simplifies these capabilities into an intuitive, task-focused interface.

Copyright information

Copyright © 2026 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.