MR.PLANB FIELD GUIDE · PROXMOX VE

Proxmox Storage Decision Guide: ZFS, Ceph, LVM Thin, RAID

Which storage layer fits one server, a lab or a cluster, and the one pairing you should never build.

Updated **2026-10-29** · Full guide: www.mrplanb.com/blog/proxmox-storage-zfs-vs-ceph-vs-lvm-thin-vs-raid

> For most single-node Proxmox systems, choose ZFS for local redundancy and snapshots, or LVM Thin for simple local block storage. For a multi-node cluster that needs distributed shared storage, Ceph is the more natural fit. Hardware RAID still has a place, but under stacks designed for it, never under ZFS or Ceph.

4
storage options compared, each solving a different layer

10 Gbps
minimum dedicated network Proxmox recommends for hyperconverged Ceph

3+
production nodes before Ceph is worth evaluating

No
ZFS on a hardware RAID controller with its own cache

The four options side by side

They are not interchangeable products. RAID combines disks, ZFS is a filesystem and volume manager, LVM Thin is local thin provisioned block storage, and Ceph is a distributed storage system.

	ZFS	LVM Thin	Ceph	Hardware RAID
Layer	Filesystem and volume manager	Local thin block storage	Distributed storage	Disk combining in a controller
Shared across nodes	No, replication between nodes	No, local only	Yes	No
VM snapshots and clones	Yes, efficient	Yes, snapshots and linked clones	Yes, via RBD	Not by itself
Extras	Checksumming, compression	Simple and quick to understand	Survives host or disk failure	Battery backed write cache
Main risk	RAM use, pool design	Thin pool reaching 100 percent	Complexity, network demands	Controller hides disk health

Ask these questions first

Choose storage by failure model before anything else.

- [ ] Can one host fail?
- [ ] Can one disk fail?
- [ ] Do the VMs need live migration?
- [ ] How much recent data can be lost?
- [ ] How much operational complexity can the team support?
- [ ] Is HA a hard requirement? If so, decide the storage architecture before deploying VMs.

What I would choose for common setups

One small server, two SSDs
A ZFS mirror. Either disk can fail without losing the pool.

Server with a good RAID controller
Consider RAID plus LVM Thin if replacing the controller adds no real benefit.

Simple lab
LVM Thin for fast provisioning and snapshots, if you can restore from backup when the host dies.

Three or more production nodes
Evaluate Ceph if you need distributed shared storage and can support dedicated fast networking.

Rules not to break

- Do not put ZFS on a hardware RAID controller that manages its own cache. ZFS needs direct communication with the disks, so use an HBA or a controller mode that exposes them.
- Do not install Ceph on a single node just because it appears in the menu.
- Do not thin provision a pool without capacity alerts. A full thin pool turns a capacity problem into VM downtime.
- Do not treat RAID as a backup. Mirroring a deleted or encrypted file just gives you redundant copies of the same problem.
- If ZFS owns the disks, let ZFS handle redundancy. If a RAID controller owns them, give the logical volume to a layer that expects it.

Live migration and HA

Shared storage is the simplest model because all eligible nodes reach the same VM disks.

Storage	Shared?	Notes
Ceph, NFS, iSCSI, Fibre Channel	Yes	Different designs, same property
Local ZFS	No	Does not become shared by clustering; ZFS replication keeps copies on other nodes
Local LVM Thin	No	Proxmox can move local disks during migration

Sources and notes: Proxmox VE documentation: Storage, ZFS on Linux, LVM Thin, Hyper-Converged Ceph Cluster, Installation · Mr.PlanB blog: Proxmox Storage: ZFS vs Ceph vs LVM Thin vs RAID