MR.PLANB FIELD GUIDE · PROXMOX VE

Proxmox SSD Wearout: 7-Fix Checklist for Consumer Drives

What the Proxmox community suggests when consumer SSDs wear fast and enterprise drives are out of budget.

Updated **2026-10-29** · Full guide: www.mrplanb.com/blog/no-budget-enterprise-drives-proxmox-ssd-wearout

> You do not need enterprise SSDs to run Proxmox, but you do need a strategy. Find out what is writing to the disks, tune or rethink ZFS, stop mirroring mismatched drives, prefer larger drives, trim hidden writes, consider used enterprise SATA drives and keep real backups.

7	660MB/h	5s to 180s	8 mo vs 5 yr
fixes from the community thread	writes reported on a system with 3 idle containers	zfs_txg_timeout change discussed (default usually 5 s)	SSD life with and without a plan, per the verdict

Find the cause before you buy

The original setup was a Lenovo M720Q with mirrored consumer drives for local-zfs, barely six months old.

- [ ] Check whether ZFS, logging or temporary file churn is writing the most.
- [ ] Look at log verbosity: the OP had logging left on default.
- [ ] Check scratch disk activity, swap and unused system services.
- [ ] Run nvme-cli or read SMART stats in the GUI.
- [ ] Note whether every VM and container shares one mirrored pool.

The 7 fixes

1. Figure out what is eating the SSD: logs, scratch disks, temp files, swap.
2. Tune ZFS or rethink it: redirect logs to RAM with log2ram, or use ext4 with lvm-thin for basic needs.
3. Do not mirror an NVMe with a SATA drive. Split them by role instead.
4. Choose bigger drives: more cells to level over and a higher TBW.
5. Watch writes and trim the fat; consider raising zfs_txg_timeout only with UPS cover.
6. Look at used enterprise SATA drives such as the Intel DC S3620 or 4610.
7. Put backups ahead of mirrors: use Proxmox Backup Server or nightly snapshots.

ZFS tuning: what to know

Option	Benefit	Catch
log2ram	Moves log writes into RAM	Tune RAM disk size and write frequency; a full RAM disk can crash services
Higher zfs_txg_timeout	Fewer commits, fewer writes (default usually 5 s, 180 s or more discussed)	More data lost on power failure; only with UPS
ext4 with lvm-thin	Less unnecessary wear for basic needs	Gives up some ZFS features like snapshots and compression
ZFS on host and in VMs	None	Double or triple write amplification

Drive choices

Bigger consumer drive
A 2TB SSD instead of 512GB spreads writes across more cells and has a higher TBW.

Used enterprise SATA
Intel DC S3620 or 4610 are often cheap on eBay with plenty of life left.

Form factor limit
Affordable enterprise NVMe is mostly 22110, while the M720Q only fits 2280.

Workaround
U.2 to PCIe adapters were mentioned as one way around the form factor problem.

Backups beat mirrors

- Mirroring an NVMe with a slower SATA drive limits speed to the slowest link and doubles writes.
- Use the NVMe for VMs and the SATA drive for backups or light containers.
- Offload backups to another system, for example Proxmox Backup Server in a VM on an Unraid box.
- Mirrors do not protect against accidental deletion, OS-level failures or ransomware. Backups do.

Sources and notes: Proxmox community discussion on SSD wearout (r/Proxmox thread summarized on the page) · Proxmox Backup Server