

Choosing the Right Storage Backend for Kubernetes PVCs

Julien Laurenceau | Published 2025-12-15 | Updated 2026-09-30

<https://pepitedata.com/kubernetes-pvc-storage-backend/>

PepiteData | LinkedIn: Julien Laurenceau | X: @Julien_Lau



Choosing a PVC backend is an architecture decision I review; the [Ceph storage consulting](#) page describes the cluster side.

Most discussions about PVC focus on **on-prem deployments**, but many of these technologies are equally relevant **in the cloud**, especially when using managed block storage like **AWS EBS**, which caps at around **64K IOPS**.

By contrast, **OpenEBS LocalPV (ZFS or LVM)** can attach directly to **NVMe-backed local disks**, unlocking **true NVMe performance (>1M IOPS)**, up to 15× faster than EBS.

-> This makes **ZFS LocalPV** a must-know option for anyone running **Kubernetes on AWS, GCP, or Azure**, particularly for workloads like **Kafka, Cassandra, or MongoDB**, which already handle replication at the application level. In such cases, **you don't need a replicated storage backend (Ceph, EBS, etc.)** and removing that extra layer can **dramatically reduce latency and cost**.

[building block] Ceph RBD

- **Type:** Distributed block storage (RADOS Block Device)

- **Pros:**
- **Cons:**
- **Best for:** large multi-node clusters needing strong HA and scalability

[component] OpenEBS

✂ ZFS LocalPV

- **Type:** Node-local storage
- **Pros:**
- **Cons:**
- **Best for:** ultra-fast RWO workloads, both **on-prem** and **in-cloud NVMe-backed nodes**

[launch] Mayastor (SPDK + NVMe-oF)

OpenEBS is gradually **deprecating cStor** in favor of **Mayastor**, its next-gen replicated engine.

- **Type:** Distributed storage built on **SPDK** (user-space) and **NVMe-over-Fabrics**
- **Pros:**
- **Cons:**
- **Best for:** modern replicated setups needing high availability with lower overhead

⚙ DRBD / LINSTOR (+ OCFS2 for RWX)

- **Type:** Block-level synchronous replication between nodes
- **Pros:**
- **OCFS2:**
- **Cons:**
- **Best for:** compact HA clusters, critical databases

⚖ Summary

Article content (see online article)

[tip] Key takeaways

- Need *scale-out resilience* → **Ceph RBD**
- Need *pure performance (on-prem or cloud)* → **OpenEBS ZFS LocalPV**
- Need *modern lightweight HA* → **OpenEBS Mayastor (SPDK/NVMe-oF)**
- Need *simple, robust HA* → **DRBD + OCFS2**

[launch] **Pro tip:** If you're running **Kafka, Cassandra, MongoDB, or Redis** on **Kubernetes in AWS**, you likely **don't need EBS or Ceph**, use **OpenEBS ZFS LocalPV** instead and get near-bare-metal NVMe performance.

[question] What's your go-to storage setup for **RWO PVCs** in production? #Kubernetes #Storage #Ceph #OpenEBS #ZFS #Mayastor #SPDK #DRBD #NVMe #AWS #EBS #DevOps #SRE

<https://www.linkedin.com/>

Update 2026-09-29. This post is now part of [Ceph in Production: RBD, RGW and CephFS](#), the field guide to Ceph in production.

Related posts

- [Ceph in Production: RBD, RGW and CephFS](#)
- [Ceph RBD mirroring for Kubernetes storage in production](#)
- [Why file count matters as much as file size on object storage](#)
- [GPUs changed everything. Storage is the bottleneck again.](#)

PDF snapshot of the public article. The online version is authoritative; links remain clickable.