

performance

PBS needs high IOPS performance. Benefit of ZFS would be that you can accelerate it using SSDs to store the metadata. But won't help that much with verify tasks (but still a bit as the HDDs are hit by less IO, because all the metadata that is read/written from the SSD doesn't have been read/written from the HDDs).

In general HDDs shouldn't be used with PBS, at least not if you store a lot of backups. And if you still do it, it's highly recommended to also use SSDs for storing the metadata.

If using L2ARC switch it to MFU also only to prevent writing a lot of data with every backup:
[tune_l2arc_for_backups](#)

How to calculate special VDEV size: [ZFS Metadata Special Device: Z](#)

So:

- RAIDZ on HDDs is very slow (due to IO OPS limit to single HDD). Throughput is $n \times \text{HDD}$, but IOOPS equals single HDD
- For HDDs use striped mirrors to multiply IOOPS (i.e. for 6 disk: 3 stripes per 2 hdd in mirror)
- Use DRAID - where groups of stripes are used by design.
- Use at least separate ZFS dataset for backup, then set block size to

```
zfs set recordsize=1M YourPoolName/DatasetUsedAsDatastore
```

- Disable atime

```
zfs set atime=off backup2
```

Don't use a raidz1/2/3 as PBS needs high IOPS performance and IOPS performance will only scale with the number of striped vdevs, not with the number of disks. So 20 disks in a raidz wouldn't be faster than a single HDD. And resilvering would also take forever.

From:

<https://niziak.spox.org/wiki/> - **niziak.spox.org**

Permanent link:

<https://niziak.spox.org/wiki/vm:proxmox:pbs:performance>

Last update: **2023/05/02 12:19**

