

Disks|ZFS|Pools|Virtual device

On this page you can prepare disks you like to use in one of your Zpools.

Pools	Datasets	Configuration
Virtual device	Management	Tools
	Information	I/O statistics
Name	Type	Description
hs0	spare	my Hot Spare
raidz0	raidz1	my RAIDZ

Following Virtual device types are possible:

Taken from the manpage: [zpool](#) - configures ZFS storage pools

Primary Type	Drive Quantity	Similar to	Description
Stripe	One or more	Raid-0	Drives act as pure storage, with no special safety/redundancy (similar to raid-0). With multiple similar-sized drives, this can increase performance, in that drives will be written to in parallel
Mirror	Two or more	Raid-1	Data is replicated in an identical fashion across all components of a mirror. A mirror with N disks of size X can hold X bytes and can withstand (N-1) devices failing before data integrity is compromised.
Raidz1	Three or more	Raid-5	Data and parity is striped across all disks within a raidz group. Addresses most of the weaknesses of Raid-5, allowing for better distribution of parity and elimination of the "RAID-5 write hole" (in which data and parity become inconsistent after a power loss). Listed in GUI as "Single Parity RAID-Z"
Raidz2	4 or more	Raid-6	Adds more redundancy/safety built in, such that as many as 2 drives can fail before losing data. Listed in GUI as "Double Parity RAID-Z"
Raidz3	5 or more	3disk redundancy	Yet more redundancy/safety. Up to 3 drives can fail before losing data. Listed in GUI as "Triple Parity RAID-Z"
Raidz	-	-	Synonym for Raidz1
Supporting Type	Description		
Hot Spare	A drive set aside and kept populated so that if any primary drive of a pool fails, the set-aside drive will be automatically used to replace the damaged/failed drive.		
Cache	One or more drives, usually SSD drives. Used to cache storage pool data. These devices provide an additional layer of caching between main memory and disk. For read-heavy workloads, where the working set size is much larger than what can be cached in main memory, using cache devices allow much more of this working set to be served from low-latency, low-seek-time media. Using cache devices provides the greatest performance improvement for random read-workloads of mostly static content, and filesystem metadata.		

Supporting Type	Description
Log	The ZFS Intent Log (ZIL) satisfies POSIX requirements for synchronous transactions. For instance, databases often require their transactions to be on stable storage devices when returning from a system call. NFS and other applications can also use fsync() to ensure data stability. By default, the intent log is allocated from blocks within the main pool. However, it might be possible to get better performance using separate intent log devices such as NVRAM or a dedicated disk.
Log (mirror)	Multiple log devices can also be specified, and they can be mirrored.

From: <https://www.xigmanas.com/wiki/> - **XigmaNAS**

Permanent link: https://www.xigmanas.com/wiki/doku.php?id=documentation:setup_and_user_guide:disks_zfs_pools_virtual_device&rev=1402309845

Last update: **2018/07/08 17:47**

