NAS Adventure

—

Traveling the BSD route

FreedomBytes
2023

Who am I?

Desiderata

- **Home/SoHo NAS with**
  - Adequate fault-tolerance
  - Space usage optimization
  - Easy to maintain
  - Easy to access remotely

Network-Attached Storage

“NAS is a file-level (as opposed to block-level storage) computer data storage server connected to a computer network providing data access to a heterogeneous group of clients”

Why you cannot rely on proprietary solutions *ONLY*?

- **Vendor Lock-In**
- **Data Privacy and Security Concerns**
- **Cost and Scalability Constraints**
- **Lack of Customization and Control**

Why you should not rely on free / community ed. *ONLY*?

- **Vendor Lock-In**
- **Cost and Scalability Constraints**
- **Lack of Customization and Control**

Thus...

Bare Metal NAS

From GNU/Linux to FreeBSD

Brief summary of the GNU/Linux configuration

- Common NAS configuration running on GNU/Linux
- Why BSD?
- Quick comparison between ZFS vs BTRFS
- Benchmarking

Common NAS configuration running on GNU/Linux

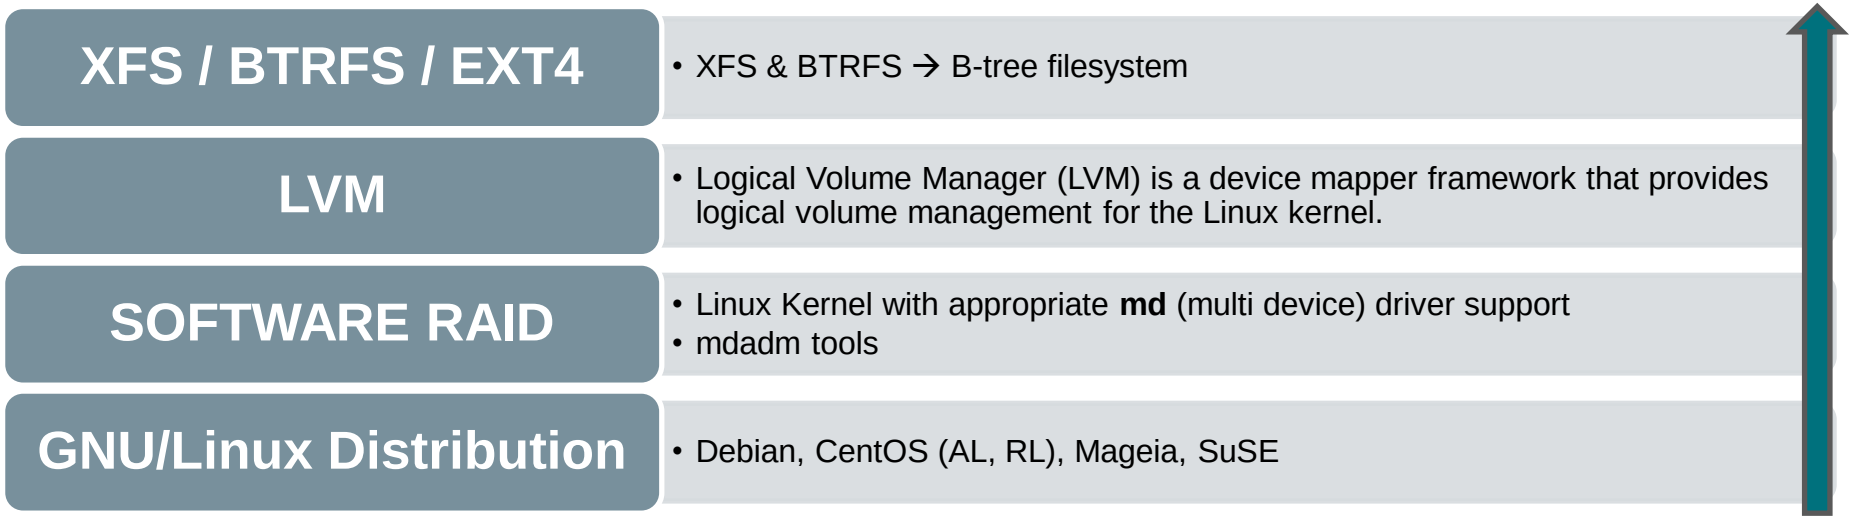

RAID: Redundant Array of *Inexpensive* Disks

Why BSD?

- ZFS File System
- Storage-Aware features
- Stability and Reliability
- DTrace

FreedomBytes
2023

ZFS vs BTRFS

ZFS

- Filesystem with volume management features
- Ported to various systems between 2006 and 2010
- Max Vol Size: 256 trillion yobibytes (2^{128} bytes)
- Max File Size: 16 exbibytes (2^{64} bytes)
- Max no. File:
  - Per directory: 2^{48}
  - Per file system: unlimited
- Max File Name Length: 255 ASCII (fewer for multibytes char standards)
- Features:
  - Forks
  - Attributes: POSIX, Extended attributes
  - FS Permissions: Unix, NFSv4 ACLs
  - Transparent compression
  - Transparent encryption
  - Deduplication: online
  - COW
- Supported OS (OpenZFS):
  - FreeBSD, Mac OS X Server > 10.5 (ro), NetBSD
  - GNU/Linux (requires third-party kernel module)

BTRFS

- Filesystem based on COW principle with a logical volume manager
- March 2009 (with Linux 2.6.29)
- Max Vol Size: 16 exbibyte (2^{64} bytes)
- Max File Size: 16 exbibytes (2^{64} bytes)
- Max no. File: 2^{64}
- Max File Name Length: 255 ASCII (fewer for multibytes char standards)
- Features:
  - ~~Forks~~
  - Attributes: POSIX, Extended attributes
  - FS Permissions: Unix, POSIX ACLs
  - Transparent compression
  - ~~Transparent encryption~~
  - Deduplication: offline
  - COW
- Supported OS:
  - Linux
  - ReactOS

Benchmarking

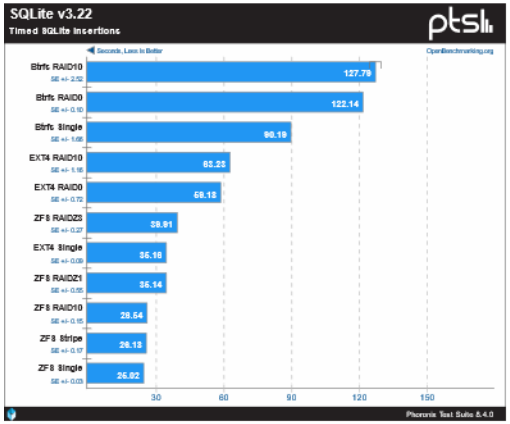

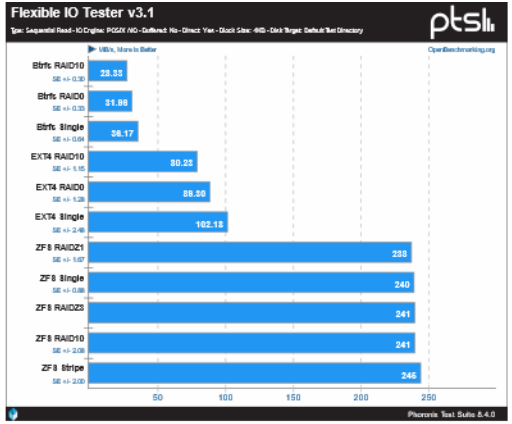

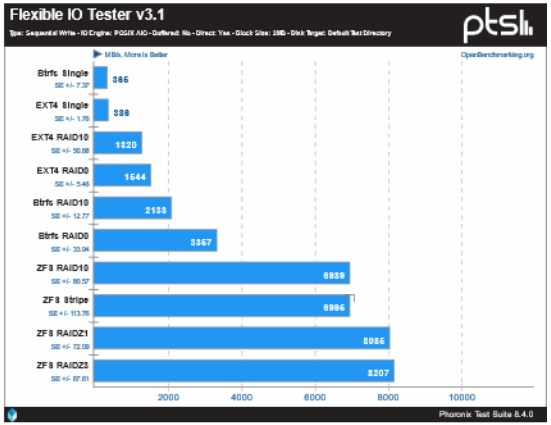

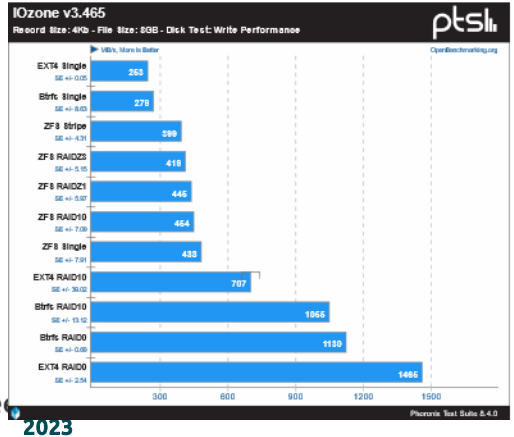

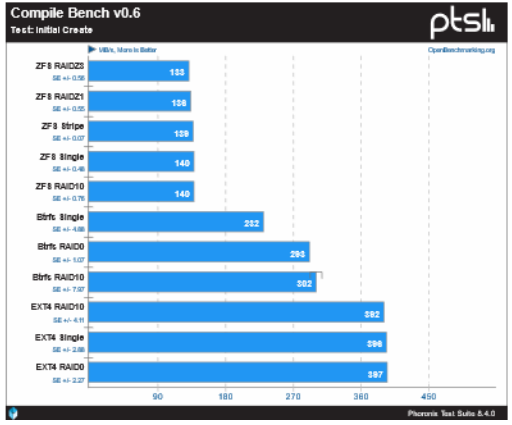

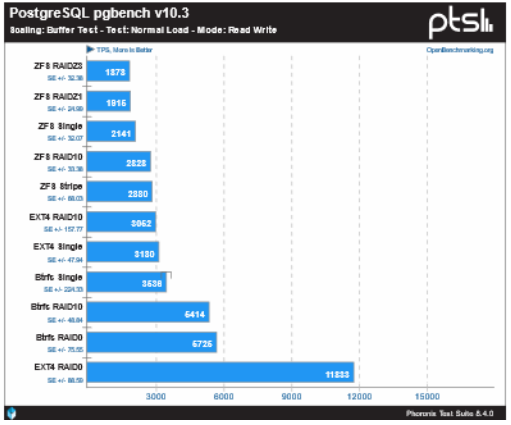

Brief summary of the BSD configuration

- FreeBSD initial setup
- Why ZFS and not UFS with gmirror?
- Configuration of the first mirrored pool
- Monitoring performances
- Checking data integrity!
- Configuring Samba server

INITIAL SETUP

boo+ufs+swap

/dev/ada8

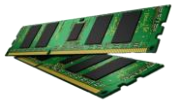

8GB+ [2]

freebsd-zfs

ada0

ada1

ada2

ada3

ada4

ada5

ada6

/etc/rc.conf:

```
zfs_enable="YES"
```

Then start the service:

```
# service zfs start
```

May 12 10:33:00 2022

FreedomBytes 2023

1 – using FreeBSD 13.2 RELEASE-p2
2 – currently using 16GB setup

Why ZFS and not UFS with gmirror?

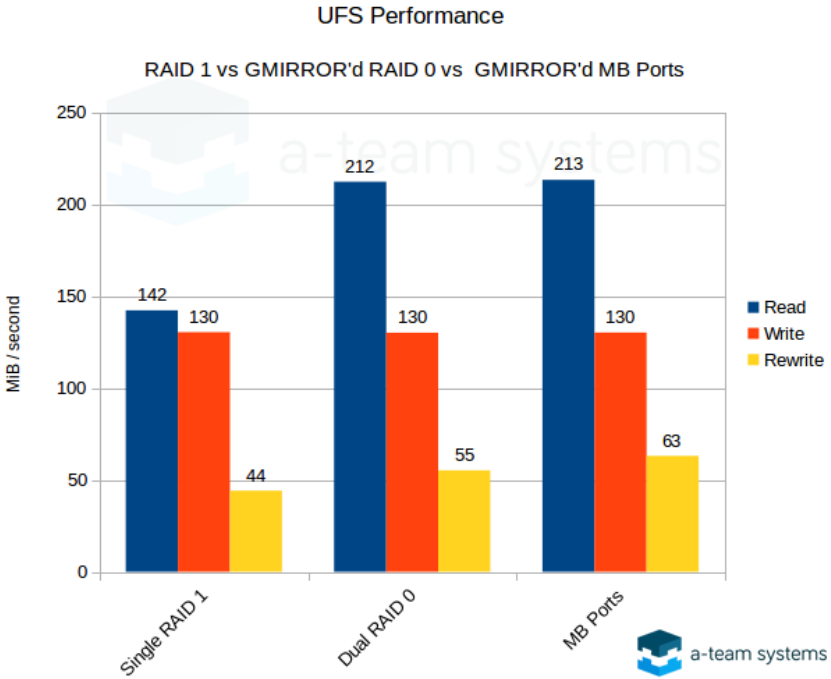

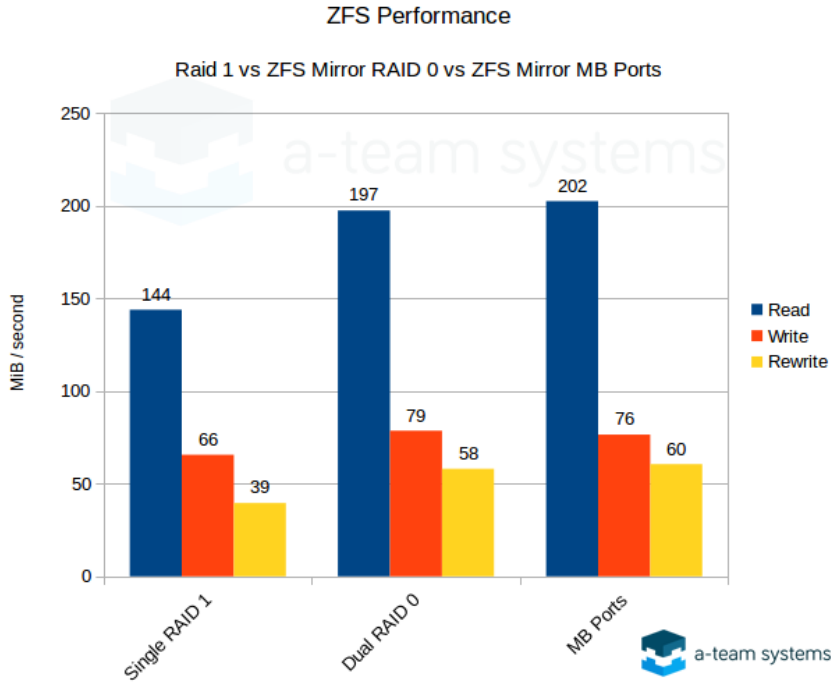

FreedomBytes
2023

Configuring the first mirrored pool

zpool create *storage_0-1* mirror ada0 ada1
zpool add *storage_0-1* mirror ada2 ada3
zpool add *storage_0-1* mirror ada4 ada6
zpool status

```
# zpool status
  pool: storage_0-1
 state: ONLINE
  scan: scrub repaired 0B in 03:47:11 with 0 errors on Wed Oct 18 03:35:00 2023
remove: Removal of vdev 2 copied 72K in 0h0m, completed on Sun Jan  8 18:12:43 2
023
        168 memory used for removed device mappings
config:

        NAME          STATE     READ WRITE CKSUM
        storage_0-1   ONLINE       0     0     0
          mirror-0    ONLINE       0     0     0
            ada0      ONLINE       0     0     0
            ada1      ONLINE       0     0     0
          mirror-3    ONLINE       0     0     0
            ada2      ONLINE       0     0     0
            ada3      ONLINE       0     0     0
          mirror-4    ONLINE       0     0     0
            ada4      ONLINE       0     0     0
            ada6      ONLINE       0     0     0

errors: No known data errors
```

Mountpoint

zfs mount

```
# zfs mount
storage_0-1                          /storage_0-1
```

Monitoring performances

zpool iostat -v

```
# zpool iostat -v
                capacity     operations     bandwidth
pool          alloc   free   read  write   read  write
------------  -----  -----  -----  -----  -----  -----
storage_0-1   1.88T  9.00T      0      0    243  13.5K
  mirror-0    1.78T  1.85T      0      0    129  3.70K
    ada0          -      -      0      0     65  1.85K
    ada1          -      -      0      0     63  1.85K
  indirect-1      -      -      0      0      0      0
  indirect-2      -      -      0      0      0      0
  mirror-3    88.9G  3.54T      0      0     94  4.64K
    ada2          -      -      0      0     46  2.32K
    ada3          -      -      0      0     47  2.32K
  mirror-4    14.7G  3.61T      0      0     19  5.16K
    ada4          -      -      0      0     10  2.58K
    ada6          -      -      0      0      8  2.58K
------------  -----  -----  -----  -----  -----  -----
```

Checking data integrity

zpool scrub *storage_0-1*

```
# zpool status
  pool: storage_0-1
 state: ONLINE
  scan: scrub repaired 0B in 03:47:11 with 0 errors on Wed Oct 18 03:35:00 2023
remove: Removal of vdev 2 copied 72K in 0h0m, completed on Sun Jan  8 18:12:43 2
023
        168 memory used for removed device mappings
config:

        NAME          STATE     READ WRITE CKSUM
        storage_0-1   ONLINE       0     0     0
          mirror-0    ONLINE       0     0     0
            ada0      ONLINE       0     0     0
            ada1      ONLINE       0     0     0
          mirror-3    ONLINE       0     0     0
            ada2      ONLINE       0     0     0
            ada3      ONLINE       0     0     0
          mirror-4    ONLINE       0     0     0
            ada4      ONLINE       0     0     0
            ada6      ONLINE       0     0     0

errors: No known data errors
```

Configure samba server

- **Thanks to ports!**
  - samba416-4.16.11 Free SMB/CIFS and AD/DC server and client for Unix
- **# pkg install samba416**
- **vi /usr/local/etc/smb4.conf**

```
[global]
workgroup = WORKGROUP
server string = Samba Server Version %v
netbios name = notyourbiz
wins support = No
security = user
passdb backend = tdbsam

[storage]
path = /storage_0-1/storage
valid users = criminal
writable  = yes
browsable = yes
read only = no
guest ok = no
public = no
create mask = 0666
directory mask = 0755
```

/etc/rc.conf:

```
samba_server_enable="YES"
```

System up and running

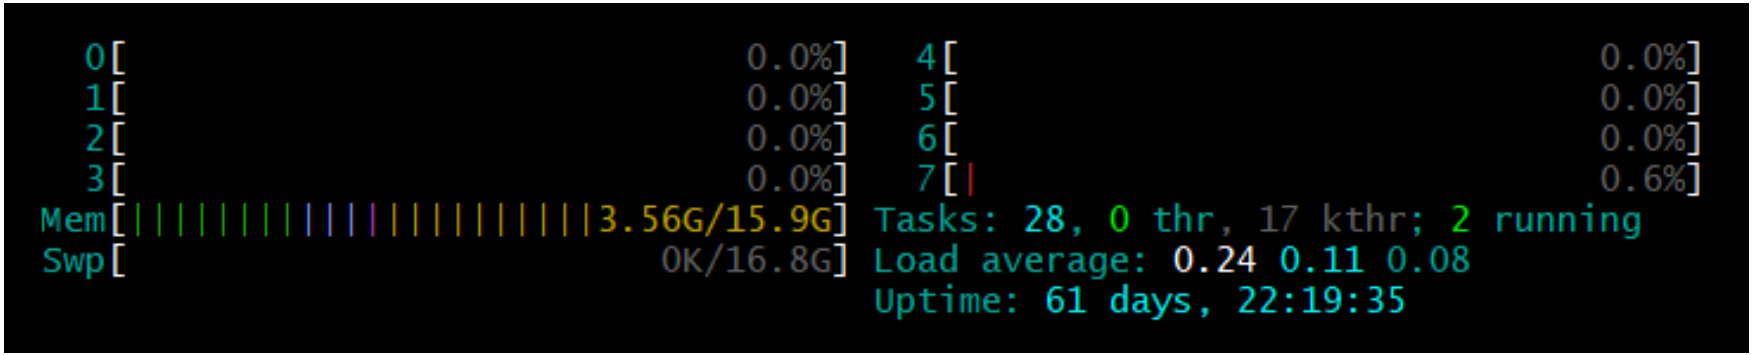

Mission Accomplished

Happy Hacking ;-)

FreedomBytes 2023

FreedomBytes

2023

backup

References

- https://docs.freebsd.org/en/books/handbook/zfs/
- https://openzfs.github.io/openzfs-docs/Project%20and%20Community/FAQ.html#do-i-have-to-use-ecc-memory-for-zfs
- https://www.klennet.com/notes/2019-07-04-raid5-vs-raidz.aspx
- https://www.klennet.com/notes/tags/zfs.aspx
- https://forums.freebsd.org/threads/openzfs-using-zpool-iostat-to-monitor-pool-perfomance-and-health.77521/
- https://docs.freebsd.org/en/books/handbook/firewalls/
- https://www.reddit.com/r/zfs/comments/mvho90/move_data_efficiently_between_datasets_in_the/
- https://www.phoronix.com/review/freebsd-12-zfs/2
- https://www.ateamsystems.com/tech-blog/freebsd-hardware-raid-vs-gmirror-vs-zfs/
- https://matteocroce.medium.com/linux-and-freebsd-networking-cbadcdb15ddd
- https://wiki.freebsd.org/NetworkPerformanceTuning

FreeBSD®

Some free / community solution

- TrueNAS (formerly FreeNAS, FreeBSD)
- XigmaNAS (formerly NAS4Free, FreeBSD)
- OMV (Open Media Vault, Debian)
- Rockstor (CentOS, Btrfs)
- NexentaStor (ex OpenSolaris, ex Illumos, ex Illumian…)

- Most of them provides community – free to use packages alongside enterprise-ready solution