

NAS

Tout ce qui est relatif su stockage

- [Probleme de pool sur le nas truenas](#)
- [Move files and folders from a pool to another one \(more complex than I thought\)](#)

Probleme de pool sur le nas truenas

A la suite d'un remplacement d'un disque défectueux sur le nas truenas, j'ai du faire quelque chose de pas correct, parce qu'il m'indique toujours un disque du pool en problème alors qu'il n'existe plus...

ada0	1836E15408CD	111.79 GiB	freenas-boot
da0	Z1F2XZ07	2.73 TiB	vol1
da1	Z5051B2S	2.73 TiB	vol1
da2	Z505GKWH	2.73 TiB	vol1
da3	Z5051B5V	2.73 TiB	vol1
da4	Z501LWNS	2.73 TiB	vol1
da5	W1F4FEA3	2.73 TiB	vol1

/mnt/vol1			DEGRADED
RAIDZ2			DEGRADED
	da0		ONLINE
	da1		ONLINE
	da2		ONLINE
	da3		ONLINE
	da4		OFFLINE
	da5		ONLINE

Right click / Online

Resilvering Status

0.11%

Resilvering pool: vol1

Status: SCANNING

Now that this operation is ongoing, I am creating another little nas as a backup to store the data from my nas in case I have to recreate it.

The other nas consist of:

- A RPI4
- An external USB3 disk bay
- 5 sata disks (3 x 3 Tb and 2x 8 Tb)

For this one I have started to build a mergerfs/snapraid config. Even if I had great success with it (good speed) when it comes to transfert large amount of files it just fails because mergerfs is more a good fir with files that does not move too often and in my case, the transfert of all my files (videos, backups, etc) cause mergerfs to not catch up with the spread of infos among the other disks. I ended installing the LVM plugin, create a big PV/VG/LV with all the disks.

I have added an rsync service (but not used.. see below).

The transfert of the files is done using a rsync command in a screen session and.. wait.

```
|— Media
| |— Downloads << la ou le client torrent telecharge
| |— Films << torrent completed pour radarr
| |— Games
| |— img
| |— Iso
| |— media-nas
| |— media-user
| |— Music
| |— Musiques
| |— Photos
| |— Series << la ou medusa place les series
| |— tdarr
| |— TorrentsFiles
```


Move files and folders from a pool to another one (more complex than I thought)

Ok.

So I have this situation.

I use Freenas with 2 pools.

- One for my media and files
- Other one for downloading legitimate torrent files

I use a software I don't recall the name (call it thing*daar) to handle the search and download of the torrent files. It download them using a torrent client (Transmission) and copy them when finish to the pool where the files are.

But that make no real sense since it needs to copy the file so I have it twice and that I cannot use symlinks because NFS does not allow this in 2 separate pools.

As backup of those two pools, I use syncthing that sync those 2 "folders" into another NAS (based on openmediavault and a RPI4).

So I need to plan the move carefully

- Stop syncthing to sync for a while
- Move the files from pool A to pool B to have this folder structure:

Before:

```
\_POOL_A
  \_Download
\_POOL_B
  \_Media
```

After:

```
\_POOL_B
  \_Downloads_A
  \_Media
```

- Reconfigure my transmission docker config tthat use NFS to point this new location
Transnmission will be maybe a little confused about this so maybe we will need to re-

locate the download if it complains

- Do a file cleanup (so remove any duplicate and/or created symlinks from Download_A to Media)
- Reconfigure syncthing to only sync 1 folder/pool instead of 2

I guess that I need to list here the different components/vms and how they access the pools

- transmission vm
 - nas_ip:/mnt/vol1/pool_a >> /downloads
- rancher pod thing*daar
 - /media >> storage | nfs_media (nas_ip:/mnt/vol1/media)
- syncthing freenasyyhkk
- syncthing raspberrypie


2023-09-16:

So THIS weekend is THE weekend.

I have decided that this little project should go out of my mind for good :)

Let's start the journey by checking one little thing first: Do we have any space in the POOL_B to copy/move the files? In fact.. do we have enough free space at all??.

Let'S check this:



Name	Type	Used	Available	Compression	Compression Ratio	Readonly	Dedup	Comments
vol1	FILESYSTEM	8.89 TiB	1.89 TiB	lz4	1.01	false	OFF	
.bhye_containers	FILESYSTEM	70.52 MiB	1.89 TiB	Inherits (lz4)	1.01	false	OFF	
backup	FILESYSTEM	858.09 GiB	341.91 GiB	Inherits (lz4)	1.04	false	ON	backup
downloads	FILESYSTEM	1.53 TiB	171.33 GiB	Inherits (lz4)	1.00	false	OFF	downloads
locage	FILESYSTEM	37.52 GiB	1.89 TiB	lz4	1.81	false	OFF	
jails	FILESYSTEM	925.57 MiB	1.89 TiB	Inherits (lz4)	2.03	false	OFF	
longhorn	FILESYSTEM	3.88 GiB	16.12 GiB	Inherits (lz4)	1.05	false	OFF	loghorn
media	FILESYSTEM	3.16 TiB	349.91 GiB	lz4	1.01	false	OFF	Media
nextcloud	FILESYSTEM	9.19 GiB	1.89 TiB	Inherits (lz4)	1.17	false	OFF	Nextcloud data
pdfrbackup002-zgyer	VOLUME	40.63 GiB	1.93 TiB	Inherits (lz4)	1.59	false	OFF	
pdfrbackup01-2o3eer	VOLUME	58.12 GiB	1.95 TiB	Inherits (lz4)	1.55	false	OFF	
pdfrbackup02-e1qujm	VOLUME	45.81 GiB	1.92 TiB	Inherits (lz4)	1.92	false	OFF	
pdfrbackup02-h747fn	VOLUME	40.63 GiB	1.93 TiB	Inherits (lz4)	1.00	false	OFF	
pdrttest-e847ef	VOLUME	10.16 GiB	1.9 TiB	Inherits (lz4)	1.44	false	OFF	
pve-vms-backup	FILESYSTEM	2.92 TiB	1.89 TiB	Inherits (lz4)	1.01	false	OFF	
pvebackup+r7c6w	VOLUME	32.5 GiB	1.91 TiB	Inherits (lz4)	1.53	false	OFF	
rancher_vol	FILESYSTEM	26.3 GiB	23.7 GiB	Inherits (lz4)	1.21	false	OFF	Rancher
vms_backup	FILESYSTEM	120.92 GiB	379.08 GiB	Inherits (lz4)	1.00	false	OFF	Dataset pour backup de vms

Ouch... 1.8T free and the POOL_A is 1.5T.

...

Well, still I can make it!!! Anyway, as soon as the copy/move is done I can recover the space so that's good.

Now another question is: where I will do the copy command? Like in a screen session on the truenas box? That could make it....

Let's do it!

- First let stop the transmission client so it will not write while we play with this
- **Update:** We need to PAUSE the synching on freenas (setup as a jail) and on my RPI4 because the sync job on the source (by the way, the synching setup I use if on way only from the nas to the rpi4 / external usb disc) is configure to take what is in ../media with contain my files and will soon contain the download folder. So in order to prevent failed analyse job and unnecessary syncs, I will pause this for now.

That part is done in portainer since it got the pod configured there

- On the nas:

```
root@nas[~]# screen
root@nas[~]# mkdir /mnt/vol1/media/Media/downloads
root@nas[~]# rsync -avPh /mnt/vol1/downloads/* /mnt/vol1/media/Media/downloads/
```

- Update: I used this command instead because I have received from truenas a warning that if the storage capacity is below 20% it can cause some issues, so instead of copying the files, I have moved them using:

```
root@nas[~]# rsync --remove-source-files -ah --info=progress2 /mnt/vol1/downloads/*
/mnt/vol1/media/Media/downloads
```

- Obviously it's going to take a lot of time but that's fine since only the download part is affected and I do not plan to do any downloads. I will just check if I can exit and reenter the screen session to see if we are in budiness
- **Update:** Make sure the disc quota is set on the destination pool to something correct of the rsync will stop (requota execeeded)

2023-09-17

Cool! transfet is done and all seems to be ok. Well, I have first to modify the mount options on the truenas server to allow the network/ip from the transmission client to be able to connect to the server. in my case the client is natted so make sure you put in the server mount's option for the networks/ips allowed the address seen by the server and not the clien'ts ip (if you are natted).

I have modified the fstab entry in the client to point to the correct path (I have allowed the alldirs options on the truenas share side) and then restarted the box. Then the transmission client seems to be ok and all the sharing are green/working.

Nice! One big step done!

- I have made a test to create an hardlink (not a symlink like ln -s) and was able to access in the media/folder/<link> the file that is in download/torrent/<file> so it seems to be a

success

- thing*darr is configured to use the hardlink method as a preferred one so we should be good. I will test with a new torrent to see how it goes and if it is really working,

2023-09-23

Now there are some settings that need to be changed but on the Rancher side this time. I use Rancher as a K3S/pod management. One of these pods is the application that helps me search, organise and download the torrent files. Since it's a pod (more or less a docker pod) I need to configure the volume mapping.

On the Rancher side, I have configured a NFS Storage configuration that looks like this:

- Storage / PersistentVolumes

2023-10-21

Arf... I knew it in advance... I was lazy and have not finished this "project" yet...

Now what is the current situation

Well,

- The pool in the nas is full so I cannot add a single file to it (meaning that I can use any torrent for now)
- The previous pool for download is still there but cannot be deleted because truenas keeps saying that it's still used

Ok. So, I need first a little plan:

- Do some cleanup on the files and download folder because I know there are some duplicate files, then
- Finish the remapping of the pods of rancher that still use the old download pool and set a new one
 - One of the pods should use another download folder, like "download_2" for another transmission pod but right now, I cannot even create a folder into the main pool structure

```
_ \pool01
  _ \file_storage
    _ \files
      _ \download
        _ \download_2
```
- Make sure any syncing setup does not use the old pool anymore
- Try to delete the old pool to have some space

Now for the cleanup part, to discover any duplicated files, I use the ideas from that site:

<https://www.baeldung.com/linux/finding-duplicate-files>, and specially this piece of code:

```
awk -F '/' '{
    f = $NF
    a[f] = f in a? a[f] RS $0 : $0
    b[f]++ }
END{for(x in b)
    if(b[x]>1)
        printf "Duplicate Filename: %s\n%s\n",x,a[x] }' <(find . -type f -not -links 2)
```

Now I can safely delete some duplicate files (remember that I can use now hardlink so I will keep the copy from the download folder and hardlink to the files folder

I can now create my download_2 folder, reconfigure my rancher pod for the NFS part to point to this folder and then do the rancher cleanup.

Note: I have added "-not -links 2" because I do not want to list the files that are already harlinks in the output (meaning, the files that are set correctly between the download folder and the target folder)

2023-10-22

So let's do the easy part with syncthing. On the nas part (the source of the files)

Jails / Mount Points TrueNAS CORE © 2023 - iXsystems, Inc.

Mount Points of syncthing Filter Mount Points of sync... COLUMNS ACTIONS

Source	Destination	
/mnt/vol1/backup/burp	/mnt/vol1/iocage/jails/syncthing/root/backup/burp	⋮
/mnt/vol1/downloads	/mnt/vol1/iocage/jails/syncthing/root/backup/downloads	⋮
/mnt/vol1/media/Media	/mnt/vol1/iocage/jails/syncthing/root/backup/Media	⋮
/mnt/vol1/nextcloud	/mnt/vol1/iocage/jails/syncthing/root/backup/nextcloud	⋮

1 - 4 of 4

Delete

Delete Mount Point **/mnt/vol1/downloads?**



Confirm

CANCEL

DELETE MOUNT POINT

PersistentVolume: downloads Bound

Age: 681 days

Name *

downloads

Description

Any text you want that better describes this resource.

Persistent Volume Claim: default:downloads

Age: 681 days

Volume Plugin *

NFS Share

Capacity *

10

GiB

Plugin configuration

Customize

Path

/mnt/vol1/downloads

Server

192.168.2.50

Read Only

Yes

No

```
root@nas[~]# mkdir /mnt/vol1/media/Media/downloads_tvshows
root@nas[~]#
```

PersistentVolume: Create

Name *

download-medusa

Description

Any text you want that better describes this resource.

Volume Plugin *

NFS Share

Capacity *

10

GiB

Plugin configuration

Customize

Path

/mnt/vol1/media/Media/downloads_tvshows root@nas[~]#

Server

192.168.2.50

Read Only

Yes

No

vol-qgppy (Persistent Volume Claim)

Mount Point *

/downloads

Sub Path in Volume

Read Only

☐

Add Mount

Remove

DeploymentPodcontainer-0+ Add Container

Labels & AnnotationsNetworkingNode SchedulingPod SchedulingResourcesScaling and Upgrade PolicySecurity ContextStorage

Storage

Persistent Volume Claim

Volume Name *
vol-9h4zg

Persistent Volume Claim *
medusa-config

☐ Read Only

Persistent Volume Claim

Volume Name *
vol-augzn

Persistent Volume Claim *
media

☐ Read Only

Persistent Volume Claim

Volume Name *
vol-qgpy

Persistent Volume Claim *
downloads

☐ Read Only

Add Volume

PersistentVolumeClaim: Create

Namespace *
default

Name *
download-medusa

Description
Any text you want that better describes this resource

Volume ClaimCustomizeLabels & Annotations

Volume Claim

Source

☐ Use a Storage Class to provision a new Persistent Volume

☒ Use an existing Persistent Volume

Persistent Volume
download-medusa

Request Storage *
10 GiB

DeploymentPodcontainer-0+ Add Container

Labels & AnnotationsNetworkingNode SchedulingPod SchedulingResourcesScaling and Upgrade PolicySecurity ContextStorage

Storage

Persistent Volume Claim

Volume Name *
vol-9h4zg

Persistent Volume Claim *
medusa-config

☐ Read Only

Persistent Volume Claim

Volume Name *
vol-augzn

Persistent Volume Claim *
media

☐ Read Only

Persistent Volume Claim

Volume Name *
vol-z7gif

Persistent Volume Claim *
download-medusa

☐ Read Only

Add Volume

Starred

Cluster

Workload

Apps

Service Discovery

Storage

PersistentVolumes

StorageClasses

ConfigMaps

PersistentVolumeClaims

Secrets

Legacy

Longhorn

Rancher Backups

More Resources

PersistentVolumeClaims

Create

Download YAMLDel1 selected

Filter

State	Name	Status	Volume	Capacity	Access Modes	StorageClass	VolumeMode	Age
Namespace: default								
☐	Boundconfig	Bound	nfs-config	10Gi	RWO		Filesystem	930 days
☐	Boundconfig-radarr	Bound	pvc-6a4537c3-67fb-459e-8438-6633e074c42f	2Gi	RWO	longhorn	Filesystem	686 days
☐	Bounddata	Bound	nfs-postgres	10Gi	RWO		Filesystem	922 days
☒	Bounddownloads	Bound	downloads	10Gi	RWO		Filesystem	681 days
☐	Bounddownload-medusa	Bound	download-medusa	10Gi	RWO		Filesystem	3.3 mins
☐	Bounddownload-radaar	Bound	downloads-radaar	10Gi	RWO		Filesystem	34 days
☐	Boundmedia	Bound	nfs-media	10Gi	RWO		Filesystem	930 days
☐	Boundmedusa-config	Bound	pvc-04629553-b2e2-4338-9e67-ac065175f3cb	1Gi	RWO	longhorn	Filesystem	624 days

Starred

Cluster

Workload

Apps

Service Discovery

Storage

PersistentVolumes

StorageClasses

ConfigMaps

PersistentVolumeClaims

Secrets

Legacy

Longhorn

Rancher Backups

More Resources

PersistentVolumes

Create

Download YAMLDel1 selected

Filter

State	Name	Reclaim Policy	Persistent Volume Claim	Source	Reason	Age
☐	Availablebackup-node-1	Retain	n/a	HostPath		689 days
☐	Bounddownload-medusa	Retain	download-medusa	NFS Share		48 mins
☒	Releaseddownloads	Retain	n/a	NFS Share		681 days
☐	Bounddownloads-radaar	Retain	download-radaar	NFS Share		34 days
☐	Boundnfs-config	Retain	config	NFS Share		930 days
☐	Boundnfs-media	Retain	media	NFS Share		930 days
☐	Boundnfs-postgres	Retain	data	NFS Share		922 days
☐	Boundpvc-6a4537c3-67fb-459e-8438-6633e074c42f	Delete	config-radarr	Longhorn		686 days
☐	Boundpvc-9d9edbd5-082d-4706-8e16-5fc64b84bd3	Delete	rancher-backup-3	HostPath		491 days
☐	Boundpvc-04629553-b2e2-4338-9e67-ac065175f3cb	Delete	medusa-config	Longhorn		624 days

Delete Dataset *downloads*

The *downloads* dataset and all snapshots stored with it **will be permanently deleted.**

Enter **downloads** below to confirm.

downloads

☐ Confirm

CANCELDELETE DATASET

2023-11-12

So, what else I need to do?

Well, I still have the cleanup process to complete. Earlier in this pst, I put a script that can detect and list all duplicate filename. It is very handy to get that list because my "automation" trigger just started and I am thinking "since I have this list, I can use it in order to delete the destination file and instead to hard link!". Yep, that smart, I know. But first I need to validate that all the entries from the output are correct.

Luckily, the output, if valid, should always look like this:

```
Duplicate Filename: the-file.txt
./downloads/complete/the-file.txt
./destination_folder/the-file.txt
```

it should be like this everytime, everything else could be considered as not correct, So I just the script below to validate those entries in the file output (dup5.txt) from the other script.

```
#!/bin/bash
SAVEIFS=$IFS
IFS=$'\n'
filename="dup5.txt"

function check () {
    myarray=("$@")

    lin1="Duplicate Filename"
    lin2="./downloads"
    lin3a="./destination_folder"
    len_lin1=${#lin1}
    len_lin2=${#lin2}
    len_lin3a=${#lin3a}

    line1="${myarray[0]}"
    line2="${myarray[1]}"
    line3="${myarray[2]}"

    if [ ${line1:0:${len_lin1}} == "$lin1" ]; then echo "ligne 1 ok";else echo "ligne 1
pas ok at index $((idx -2));exit;fi
    if [ ${line2:0:${len_lin2}} == "$lin2" ]; then echo "ligne 2 ok";else echo "ligne 2
pas ok at index $((idx -1));exit;fi
    if [ ${line3:0:${len_lin3a}} == "$lin3a" ]; then echo "ligne 3 ok";else echo "ligne 3
pas ok at index $idx";exit;fi
```

```

}

readarray -t lines < "$filename"

let s=0
array=()
for ((idx=0; idx<${#lines[@]}; ++idx)); do
    echo "$idx" "${lines[idx]}"
    if ! [ "$s" -eq 3 ]; then
        array+=("${lines[idx]}")
        let s++
    fi
    if [ "$s" -eq 3 ];then
        check "${array[@]}" # ici go function to test the array
        let s=0
        array=()
    fi
done

IFS=$SAVEIFS

```

Now when I run it I can see the lines in the dup5.txt result file that contain some entries that are not correct. Once I did the cleanup on the files (so removing the files themselves if needed and/or removing the entries in the dup5.txt file in case the file name is the same but it is ok), then I can use the same script to delete the destination file and use a hardlink instead.

2023-11-22

So, I guess it's time to switch the 3T drives to the 8T ones. How to do this in Truenas?

First, let's take one offline (starting with the one with sector errors)

nas.digitalrat.lan/ui/storage/disks

TrueNAS CORE

systems

TrueNAS CORE © 2023 - iXsystems, Inc.

root
nas.digitalrat.lan

Dashboard

Accounts

System

Tasks

Network

Storage

Pools

Snapshots

VMware-Snapshots

Disks

Import Disk

Directory Services

Storage / Disks

Disks

Filter Disks

COLUMNS

	Name	Serial	Disk Size	Pool	
☐	ada0	1836E15408CD	111.79 GiB	freenas-boot	>
☐	da0	Z1F2XZ07	2.73 TiB	vol1	>
☐	da1	Z5051B2S	2.73 TiB	vol1	>
☐	da2	Z505GKWH	2.73 TiB	vol1	>
☐	da3	Z5051B5V	2.73 TiB	vol1	>
☐	da4	Z501LWNS	2.73 TiB	vol1	>

1 - 6 of 7

12

Nov 22 07:16:45 nas.digitalrat.lan 1 2023-11-22T07:16:45.045640-05:00 nas.digitalrat.lan smartd 1696 - - Device: /dev/da0 [SAT], 288 Offline uncorrectable sectors

Nov 22 07:46:45 nas.digitalrat.lan 1 2023-11-22T07:46:45.156894-05:00 nas.digitalrat.lan smartd 1696 - - Device: /dev/da0 [SAT], 288 Currently unreadable (pending) sectors

Nov 22 07:46:45 nas.digitalrat.lan 1 2023-11-22T07:46:45.156990-05:00 nas.digitalrat.lan smartd 1696 - - Device: /dev/da0 [SAT], 288 Offline uncorrectable sectors

nas.digitalrat.lan/ui/storage/pools

TrueNAS CORE

systems

TrueNAS CORE © 2023 - iXsystems, Inc.

root
nas.digitalrat.lan

Dashboard

Accounts

System

Tasks

Network

Storage

Pools

Snapshots

VMware-Snapshots

Disks

Import Disk

Directory Services

Storage / Pools

Pools

ADD

vol1 (System Dataset Pool) ONLINE 9.06 TiB (84%) Used | 1.71 TiB Free

Name	Type	Used	Available	Compression	Compression Ratio	Readonly	Dedup	Cache	Pool Actions
vol1	FILESYSTEM	9.06 TiB	1.71 TiB	lz4	1.01	false	OFF		Pool Options
.bhyve_containers	FILESYSTEM	70.52 MiB	1.71 TiB	Inherits (lz4)	1.01	false	OFF		Export/Disconnect
backup	FILESYSTEM	862.45 GiB	337.55 GiB	Inherits (lz4)	1.04	false	ON	be	Add Vdevs
iocage	FILESYSTEM	37.4 GiB	1.71 TiB	lz4	1.88	false	OFF		Scrub Pool
jails	FILESYSTEM	925.57 MiB	1.71 TiB	Inherits (lz4)	2.03	false	OFF		Status
longhorn	FILESYSTEM	4.34 GiB	15.66 GiB	Inherits (lz4)	1.05	false	OFF	lo	Expand Pool
media	FILESYSTEM	4.82 TiB	1.18 TiB	lz4	1.00	false	OFF		Media
nextcloud	FILESYSTEM	9.32 GiB	1.71 TiB	Inherits (lz4)	1.18	false	OFF		Nextcloud data
pdrbackup002-zgver	VOLUME	40.63 GiB	1.75 TiB	Inherits (lz4)	1.59	false	OFF		
pdrbackup01-2o3eer	VOLUME	58.12 GiB	1.77 TiB	Inherits (lz4)	1.55	false	OFF		
pdrbackup02-a1qujm	VOLUME	40.63 GiB	1.74 TiB	Inherits (lz4)	2.00	false	OFF		

Offline Disk

Offline disk da0?

☒ Confirm

CANCEL

OFFLINE

Pool Status

REFRESH

SCRUB

Status: FINISHED

Errors: 0

Date: 2023-11-19 00:00:13

Name	Read	Write	Checksum	Status	
▼ /mnt/vol1	0	0	0	DEGRADED	
▼ RAIDZ2	0	0	0	DEGRADED	
da5	0	0	0	ONLINE	⋮
da1	0	0	0	ONLINE	⋮
da3	0	0	0	ONLINE	⋮
da2	0	0	0	ONLINE	⋮
da0	0	0	0	OFFLINE	⋮
da4	0	0	0	ONLINE	⋮

Disks

Filter Disks

COLUMNS

	Name	Serial	Disk Size	Pool	
☐	ada0	1836E15408CD	111.79 GiB	freenas-boot	>
☐	da0	WD-CA2Y6LKK	7.28 TiB	N/A	>
☐	da1	Z5051B2S	2.73 TiB	vol1	>
☐	da2	Z505GKWH	2.73 TiB	vol1	>
☐	da3	Z5051B5V	2.73 TiB	vol1	>
☐	da4	Z501LWNS	2.73 TiB	vol1	>

Pool Status

REFRESH

SCRUB

Status: FINISHED

Errors: 0

Date: 2023-11-19 00:00:13

Name	Read	Write	Checksum	Status
/mnt/vol1	0	0	0	DEGRADED
RAIDZ2	0	0	0	DEGRADED
da5	0	0	0	ONLINE
da1	0	0	0	ONLINE
da3	0	0	0	ONLINE
da2	0	0	0	ONLINE
/dev/gptid/3da917f6-360a-11e8-a468-7085c24f6f21	0	0	0	OFFLINE
da4	0	0	0	ONLINE

/dev/gptid/3da917f6-360a-11e8-a468-7085c24f6f21	0	0	0	OFFLINE	
da4	0	0	0	ONLINE	Edit Online Replace

02 08:10:58 nas.digitalrat.lan da0: <ATA ST3000DM001-1CH1 CC29> s/n Z1F2XZ07 detached

02 08:10:58 nas.digitalrat.lan /da0:ms0:0:0:0: Beribh destroyed

Replacing disk /dev/gptid

/3da917f6-360a-11e8-

a468-7085c24f6f21

Member disk

da0

☐ Force

CANCEL

REPLACE DISK

Replacing Disk

30.00%

Replacing disk

Replacing Disk

Successfully replaced disk `/dev/gptid/3da917f6-360a-11e8-a468-7085c24f6f21`.

CLOSE

Pool Status

RESILVER

Status: SCANNING

Completed: 00%

Time Remaining: 293 days, 10 hours, 25 minutes, 33 seconds

Errors: 0

Date: 2023-11-22 08:22:01

Pool Status

RESILVER

Status: SCANNING

Completed: 00.04%

Time Remaining: 22 days, 1 hour, 29 minutes, 51 seconds

Errors: 0

Date: 2023-11-22 08:22:01

Pool Status

RESILVER

Status: SCANNING

Completed: 01.09%

Time Remaining: 1 day, 13 hours, 28 minutes, 9 seconds

Errors: 0

Date: 2023-11-22 08:22:01



root

nas.digitalrat.lan

Dashboard

Accounts

Storage / Pools

TrueNAS CORE® © 2023 - iXsystems

Pools

ADD

vol1 (System Dataset Pool)

ONLINE (Unhealthy)  | 9.11 TiB (85%) Used | 1.59 TiB Free

```
root@nas[~]# zpool status -xv
pool: vol1
```

```
state: ONLINE
status: One or more devices has experienced an unrecoverable error. An
       attempt was made to correct the error. Applications are unaffected.
action: Determine if the device needs to be replaced, and clear the errors
       using 'zpool clear' or replace the device with 'zpool replace'.
see: https://openzfs.github.io/openzfs-docs/msg/ZFS-8000-9P
scan: scrub in progress since Sat Nov 25 10:25:13 2023
      13.2T scanned at 162M/s, 12.8T issued at 157M/s, 13.3T total
      330M repaired, 96.53% done, 00:51:14 to go
config:
```

NAME	STATE	READ	WRITE	CKSUM
vol1	ONLINE	0	0	0
raidz2-0	ONLINE	0	0	0
gptid/65c1b0b9-4a60-11ed-863a-6cb3111b4436	ONLINE	0	0	0
gptid/7d7c80fe-8a65-11ee-bd6f-6cb3111b4436	ONLINE	0	0	14
(repairing)				
gptid/3c6894da-360a-11e8-a468-7085c24f6f21	ONLINE	0	0	0
gptid/51df798c-8ae1-11ee-bd6f-6cb3111b4436	ONLINE	0	0	10.5K
(repairing)				
gptid/1a44a6d9-893a-11ee-bd6f-6cb3111b4436	ONLINE	0	0	0
gptid/3e40c4e8-360a-11e8-a468-7085c24f6f21	ONLINE	0	0	0

```
errors: No known data errors
```

Resilvering Status

25.89%

Resilvering pool: vol1

Status: SCANNING

CLOSE

Disks

Filter Disks

COLUMNS

	Name	Serial	Disk Size	Pool	
☐	ada0	1836E15408CD	111.79 GiB	freenas-boot	>
☐	da0	WD-CA2ZDR3K	7.28 TiB	vol1	>
☐	da1	WD-CA2Y6LKK	7.28 TiB	vol1	>
☐	da2	WD-CA2ZAXNK	7.28 TiB	vol1	>
☐	da3	WD-CA2ZRPBK	7.28 TiB	vol1	>
☐	da4	WD-CA2Y4UHK	7.28 TiB	vol1	>
☐	da5	WD-CA2VD3EK	7.28 TiB	vol1	>

Pools

ADD

vol1 (System Dataset Pool)

ONLINE | 9.06 TiB (31%) Used | 19.87 TiB Free

Name	Type	Used	Available	Compression	Compression Ratio	Readonly	Dedup	Cache	Pool Actions
vol1	FILESYSTEM	9.06 TiB	19.87 TiB	lz4	1.01	false	OFF	Cache	Pool Options
.bhyve_containers	FILESYSTEM	70.52 MiB	19.87 TiB	Inherits (lz4)	1.01	false	OFF	Cache	Export/Disconnect
backup	FILESYSTEM	883.54 GiB	316.46 GiB	Inherits (lz4)	1.04	false	ON	Cache	Add Vdevs
> iocage	FILESYSTEM	37.42 GiB	19.87 TiB	lz4	1.88	false	OFF	Cache	Scrub Pool
> jails	FILESYSTEM	925.57 MiB	19.87 TiB	Inherits (lz4)	2.03	false	OFF	Cache	Status
longhorn	FILESYSTEM	4.52 GiB	15.48 GiB	Inherits (lz4)	1.05	false	OFF	Cache	Expand Pool
media	FILESYSTEM	4.82 TiB	1.17 TiB	lz4	1.00	false	OFF	Cache	More...

Expand pool vol1

Expand pool to fit all available disk space.

CANCEL

EXPAND POOL

Pools

ADD

vol1 (System Dataset Pool)

ONLINE  | 9.06 TiB (31%) Used | 19.87 TiB Free

