

# btrfs check

## checksum error at logical

```
BTRFS warning (device sdb4): checksum error at logical 271624814592 on dev /dev/sdb4, physical 270551072768, root 677114, inode 739, offset 14508666880, length 4096, links 1 (path: Windows_10/Snapshots/{8b68a2dd-457e-496f-96df-cf24806dfe2d}.vdi)
```

Reading of mentioned file returns IO error when hit 70% of progress.

Corruption detected by periodic scrub. File was not touched for long time, so probably bitrot occurs. This is single SSD data, so there is no mirror of data. There is a backup of data.

Possible solution:

- delete file to remove invalid entry from cksum tree and restore from backup
- try playing with mounting with `ro, rescue=ignoredatacsums` and copy bad file to another place and delete from source

## corrupt leaf invalid data ref offset

```
[4038494.678130] BTRFS critical (device sdb4): corrupt leaf: block=407659724800 slot=84 extent bytenr=406204694528 len=16384 invalid data ref offset, have 25442122467423744 expect aligned to 4096
[4038494.678133] BTRFS info (device sdb4): leaf 407659724800 gen 8637875 total ptrs 97 free space 5617 owner 2
...
[4038494.679249]           ref#9: extent data backref root 388257660928
objectid 99389691396278 offset 25442122467423744 count 11927552
[4038494.679251]           ref#10: (extent 407659724800 has INVALID ref type 192)
...
[4038494.679393] BTRFS error (device sdb4): block=407659724800 write time tree block corruption detected
[4038494.689126] BTRFS: error (device sdb4) in btrfs_commit_transaction:2389: errno=-5 I/O failure (Error while writing out transaction)
[4038494.689129] BTRFS info (device sdb4): forced readonly
[4038494.689131] BTRFS warning (device sdb4): Skipping commit of aborted transaction.
[4038494.689132] BTRFS: error (device sdb4) in cleanup_transaction:1944: errno=-5 I/O failure
```

```
btrfs check /dev/sdb4
```

```
ERROR: errors found in fs roots
found 338856771584 bytes used, error(s) found
```

```
btrfs rescue zero-log /dev/sdb4
btrfs check --repair /dev/mapper/pool2
```

## cache and super generation don't match

```
btrfs check --repair /dev/mapper/pool2
```

```
[3/7] checking free space cache
cache and super generation don't match, space cache will be invalidated
```

It is not real problem, but cache can be cleared:

```
btrfs check --repair --clear-space-cache v1 /dev/mapper/pool2
```

```
enabling repair mode
Opening filesystem to check...
Checking filesystem on /dev/mapper/pool2
UUID: 947b20e3-bc05-4f85-8989-d19178506209
Free space cache cleared
```

## bad file extent, some csum missing

```
btrfs check --repair /dev/mapper/pool2
```

```
[4/7] checking fs roots
root 649 inode 418 errors 1040, bad file extent, some csum missing
root 649 inode 419 errors 1040, bad file extent, some csum missing
ERROR: errors found in fs roots
```

Root number is subvolume id number, so mount subvolume to find files belonging to inodes

```
mount /dev/mapper/pool2 -o subvolid=649 /mountpoint
find /mountpoint -xdev -inum 418 -print
```

Find will report multiple files located on multiple subvolumes with the same inode. List subvolumes to get root id:

```
btrfs sub list /mountpoint
```

```
ID 257 gen 524652 top level 5 path @volume
ID 258 gen 524651 top level 257 path @volume/subvol1
```

```
ID 649 gen 524652 top level 257 path @volume/subvol12
<code>

===== Rescan hasn't been =====
<code>
[7/7] checking quota groups
Rescan hasn't been initialized, a difference in qgroup accounting is
expected
```

```
btrfs quota rescan /mountpoint
```

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



Permanent link:  
[https://niziak.spox.org/wiki/linux:fs:btrfs:btrfs\\_check](https://niziak.spox.org/wiki/linux:fs:btrfs:btrfs_check)

Last update: **2023/03/13 10:44**