NAME

btrfs-inspect-internal − query various internal information

SYNOPSIS

btrfs inspect−internal <subcommand> <args>

DESCRIPTION

This command group provides an interface to query internal information. The functionality ranges from a simple UI to an ioctl or a more complex query that assembles the result from several internal structures. The latter usually requires calls to privileged ioctls.

SUBCOMMAND

dump−super [options] <device> [device...]

Show btrfs superblock information stored on given devices in textual form. By default the first superblock is printed, more details about all copies or additional backup data can be printed.

Besides verification of the filesystem signature, there are no other sanity checks. The superblock checksum status is reported, the device item and filesystem UUIDs are checked and reported.

Note:

The meaning of option *−s* has changed in version 4.8 to be consistent with other tools to specify superblock copy rather the offset. The old way still works, but prints a warning. Please update your scripts to use *−−bytenr* instead. The option *−i* has been deprecated.

Options

−f|−−full

print full superblock information, including the system chunk array and backup roots

−a|−−all

print information about all present superblock copies (cannot be used together with *−s* option)

−i <*super*>

(deprecated since 4.8, same behaviour as *−−super*)

−−bytenr <*bytenr*>

specify offset to a superblock in a non−standard location at *bytenr*, useful for debugging (disables the *−f* option)

If there are multiple options specified, only the last one applies.

−F|−−force

attempt to print the superblock even if a valid BTRFS signature is not found; the result may be completely wrong if the data does not resemble a superblock

−s|−−super <bytenr>

(see compatibility note above)

specify which mirror to print, valid values are 0, 1 and 2 and the superblock must be present on the device with a valid signature, can be used together with *−−force*

dump−tree [options] <device> [device...]

Dump tree structures from a given device in textual form, expand keys to human readable equivalents where possible. This is useful for analyzing filesystem state or inconsistencies and has a positive educational effect on understanding the internal filesystem structure.

Note:

By default contains file names, consider that if you're asked to send the dump for analysis and use *−−hide−names* eventually. Does not contain file data.

Special characters in file names, xattr names and values are escaped, in the C style like **\n** and octal encoding **\NNN**.

Options

−e|−−extents
print only extent−related information: extent and device trees

−d|−−device
print only device−related information: tree root, chunk and device trees

−r|−−roots
print only short root node information, i.e. the root tree keys

−R|−−backups
same as *−−roots* plus print backup root info, i.e. the backup root keys and the respective tree root block offset

−u|−−uuid
print only the uuid tree information, empty output if the tree does not exist

−b <*block_num*>
print info of the specified block only, can be specified multiple times

−−follow
use with *−b*, print all children tree blocks of <*block_num*>

−−dfs (default up to 5.2)

use depth−first search to print trees, the nodes and leaves are intermixed in the output

−−bfs (default since 5.3)

use breadth−first search to print trees, the nodes are printed before all leaves

−−hide−names
print a placeholder *HIDDEN* instead of various names, useful for developers to inspect the dump while keeping potentially sensitive information hidden

This is:

- directory entries (files, directories, subvolumes)
- default subvolume
- extended attributes (name, value)
- hardlink names (if stored inside another item or as extended references in standalone items)

Note:
Lengths are not hidden because they can be calculated from the item size anyway.

−−csum−headers
print b−tree node checksums stored in headers (metadata)

−−csum−items
print checksums stored in checksum items (data)

−−noscan
do not automatically scan the system for other devices from the same filesystem, only use the devices provided as the arguments

−t <*tree_id*>
print only the tree with the specified ID, where the ID can be numerical or common name in a flexible human readable form

The tree id name recognition rules:

- case does not matter
- the C source definition, e.g. BTRFS_ROOT_TREE_OBJECTID
- short forms without BTRFS_ prefix, without _TREE and _OBJECTID suffix, e.g. ROOT_TREE, ROOT
- convenience aliases, e.g. DEVICE for the DEV tree, CHECKSUM for CSUM
- unrecognized ID is an error

inode−resolve [−v] <ino> <path>
(needs root privileges)

resolve paths to all files with given inode number *ino* in a given subvolume at *path*, i.e. all hardlinks

Options

−v (deprecated) alias for global *−v* option

logical−resolve [−Pvo] [−s <bufsize>] <logical> <path>
(needs root privileges)

resolve paths to all files at given *logical* address in the linear filesystem space

Options

−P skip the path resolving and print the inodes instead

−o ignore offsets, find all references to an extent instead of a single block. Requires kernel support for the V2 ioctl (added in 4.15). The results might need further processing to filter out unwanted extents by the offset that is supposed to be obtained by other means.

−s <*bufsize*>
set internal buffer for storing the file names to *bufsize*, default is 64KiB, maximum 16MiB. Buffer sizes over 64KiB require kernel support for the V2 ioctl (added in 4.15).

−v (deprecated) alias for global *−v* option

list−chunks [options] <path>
(needs root privileges)

Enumerate chunks on all devices. The chunks represent the physical range on devices (not to be confused with block groups that represent the logical ranges, but the terms are often used interchangeably).

Example output:

```
   Devid PNumber      Type/profile    PStart    Length      PEnd LNumber    LS
   ----- ------- ----------------- --------- --------- --------- ------- ------
       1       1       Data/single   1.00MiB  84.00MiB  85.00MiB      68 191.6
       1       2     System/DUP     85.00MiB  32.00MiB 117.00MiB      39 140.1
       1       3     System/DUP    117.00MiB  32.00MiB 149.00MiB      40 140.1
       1       4   Metadata/DUP    149.00MiB 192.00MiB 341.00MiB      59 188.4
       1       5   Metadata/DUP    341.00MiB 192.00MiB 533.00MiB      60 188.4
```

```
   1       6       Data/single 533.00MiB   1.00GiB   1.52GiB      49 169.9
   1       7       Data/single   1.52GiB  16.00MiB   1.54GiB      69 191.6
   1       8       Data/single   1.54GiB   1.00GiB   2.54GiB      17 100.9
   1       9       Data/single   2.54GiB   1.00GiB   3.54GiB      16  99.9
   1      10       Data/single   3.54GiB   1.00GiB   4.54GiB       1  71.4
   1      11       Data/single   4.54GiB   1.00GiB   5.54GiB      33 125.0
   1      12       Data/single   5.54GiB   1.00GiB   6.54GiB      50 170.9
   1      13       Data/single   6.54GiB 512.00MiB   7.04GiB      63 189.1
   1      14       Data/single   7.04GiB   1.00GiB   8.04GiB      51 171.9
```

- *Devid* −− the device id
- *PNumber* −− the number of the chunk on the device (in order)
- *Type/profile* −− the chunk type and profile
- *PStart* −− the chunk start on the device
- *Length* −− the chunk length (same for physical and logical address space)
- *PEnd* −− the chunk end, effectively *PStart* + *Length*
- *LNumber* −− the number of the chunk, in the logical address space of the whole filesystem
- *LStart* −− the chunk start in the logical address space of the whole filesystem, as it's a single space it's also called *offset*
- *Usage* −− chunk usage, percentage of used data/metadata of the chunk length

The chunks in the output can be sorted by one or more sorting criteria, evaluated as specified, in the ascending order (or descending when used like −*devid*,−*pstart*). By default the chunks are sorted by *devid* and *pstart*, this is most convenient for single device filesystems.

On multi−device filesystems it's up to the user what is preferred as the layout of chunks on e.g. striped profiles (RAID0 etc) cannot be easily represented. A logical view with corresponding underlying structure would be better, but sorting by *lstart*,*devid* at least groups devices of the given logical range. Can be also combined with *usage*.

This output can provide information for balance filters.

Options

−−**sort** *MODE*

sort by a column (ascending, prepend '−' for descending):

MODE is a comma separated list of:

devid − by device id (default, with pstart)

pstart − physical start (relative to the beginning of the device)

lstart − logical offset (in the logical address space)

usage − by chunk usage (percentage)

length − by chunk length

type − by chunk type (data, metadata, system)

profile − by chunk profile (single, DUP, RAID0, RAID1, ...)

−−raw raw numbers in bytes, without the *B* suffix

−−human−readable
print human friendly numbers, base 1024, this is the default

−−iec select the 1024 base for the following options, according to the IEC standard

−−si select the 1000 base for the following options, according to the SI standard

−−kbytes
show sizes in KiB, or kB with −−si

−−mbytes
show sizes in MiB, or MB with −−si

−−gbytes
show sizes in GiB, or GB with −−si

−−tbytes
show sizes in TiB, or TB with −−si

map−swapfile [options] <file>
(needs root privileges)

Find device−specific physical offset of *file* that can be used for hibernation. Also verify that the *file* is suitable as a swapfile. See also command**btrfs filesystem mkswapfile** and the Sw apfile feature <> description.

Note:
Do not use **filefrag** or *FIEMAP* ioctl values reported as physical, this is different due to internal filesystem mappings. The hibernation expects offset relative to the physical block device.

Options

−r|−−resume−offset
print only the value suitable as resume offset for file **/sys/power/resume_offset**

min−dev−size [options] <path>
(needs root privileges)

return the minimum size the device can be shrunk to, without performing any resize operation, this may be useful before executing the actual resize operation

Options

−−id <*id*>
specify the device *id* to query, default is 1 if this option is not used

rootid <path>
for a given file or directory, return the containing tree root id, but for a subvolume itself return its own tree id (i.e. subvol id)

Note:
The result is undefined for the so−called empty subvolumes (identified by inode number 2), but such a subvolume does not contain any files anyway

subvolid−resolve <subvolid> <path>
(needs root privileges)

resolve the absolute path of the subvolume id *subvolid*

tree−stats [options] <device>
(needs root privileges)

Print sizes and statistics of trees. This takes a device as an argument and not a mount point unlike other commands.

Note:
In case the the filesystem is still mounted it's possible to run the command but the results may be inaccurate or various errors may be printed in case there are ongoing writes to the filesystem. A warning is printed in such case.

Options

−b|−−raw
raw numbers in bytes, without the *B* suffix

−t <*tree_id*>
print only the tree with the specified ID, where the ID can be numerical or common name in a flexible human readable form

The tree id name recognition rules:

- case does not matter
- the C source definition, e.g. BTRFS_ROOT_TREE_OBJECTID
- short forms without BTRFS_ prefix, without _TREE and _OBJECTID suffix, e.g. ROOT_TREE, ROOT
- convenience aliases, e.g. DEVICE for the DEV tree, CHECKSUM for CSUM
- unrecognized ID is an error

−−human−readable
print human friendly numbers, base 1024, this is the default

−−iec select the 1024 base for the following options, according to the IEC standard

−−si select the 1000 base for the following options, according to the SI standard

−−kbytes
show sizes in KiB, or kB with −−si

−−mbytes
show sizes in MiB, or MB with −−si

−−gbytes
show sizes in GiB, or GB with −−si

−−tbytes
show sizes in TiB, or TB with −−si

EXIT STATUS

btrfs inspect−internal returns a zero exit status if it succeeds. Non zero is returned in case of failure.

AVAILABILITY

btrfs is part of btrfs−progs. Please refer to the documentation at <https://btrfs.readthedocs.io>.

SEE ALSO

mkfs.btrfs(8) <>