Full credit is given to the above companies including the OS that this PDF file was generated!

Red Hat Enterprise Linux Release 9.2 Manual Pages on 'blktrace.8' command

$ man blktrace.8

BLKTRACE(8) BLKTRACE(8)

NAME

blktrace - generate traces of the i/o traffic on block devices

SYNOPSIS

blktrace -d dev [-r debugfs_path] [-o output] [-w time] [-a ac?

tion] [-A action_mask] [-v]

DESCRIPTION

blktrace is a block layer IO tracing mechanism which provides detailed

information about request queue operations up to user space. There are

three major components: a kernel component, a utility to record the i/o

trace information for the kernel to user space, and utilities to ana?

lyse and view the trace information. This man page describes blktrace,

which records the i/o event trace information for a specific block de?

vice to a file.

The blktrace utility extracts event traces from the kernel (via the re?

laying through the debug file system). Some background details concern?

ing the run-time behaviour of blktrace will help to understand some of

the more arcane command line options:

- blktrace receives data from the kernel in buffers passed up through

the debug file system (relay). Each device being traced has a file

created in the mounted directory for the debugfs, which defaults to

/sys/kernel/debug -- this can be overridden with the -r command line

argument.

- blktrace defaults to collecting all events that can be traced. To limit the events being captured, you can specify one or more filter masks via the -a option.

  Alternatively, one may specify the entire mask utilising a hexadeci? mal value that is version-specific. (Requires understanding of the internal representation of the filter mask.)

- As noted above, the events are passed up via a series of buffers stored into debugfs files. The size and number of buffers can be specified via the -b and -n arguments respectively.

- blktrace stores the extracted data into files stored in the local di? rectory. The format of the file names is (by default) device.blk? trace.cpu, where device is the base device name (e.g, if we are trac? ing /dev/sda, the base device name would be sda); and cpu identifies a CPU for the event stream.

  The device portion of the event file name can be changed via the -o option.

- blktrace may also be run concurrently with blkparse to produce live output -- to do this specify -o - for blktrace.

- The default behaviour for blktrace is to run forever until explicitly killed by the user (via a control-C, or sending SIGINT signal to the process via invocation the kill (1) utility). Also you can specify a run-time duration for blktrace via the -w option -- then blktrace will run for the specified number of seconds, and then halt.

OPTIONS

-A hex-mask

--set-mask=hex-mask

Set filter mask to hex-mask (see below for masks)

-a mask

--act-mask=mask

Add mask to current filter (see below for masks)

-b size

--buffer-size=size

Specifies buffer size for event extraction (scaled by 1024). The

default buffer size is 512KiB.

-d dev

--dev=dev

Adds dev as a device to trace

-I file

--input-devs=file

Adds the devices found in file as devices to trace

-n num-sub

--num-sub-buffers=num-sub

Specifies number of buffers to use. blktrace defaults to 4 sub buffers.

-l

--listen

Run in network listen mode (blktrace server)

-h hostname

--host=hostname

Run in network client mode, connecting to the given host

-p number

--port=number

Network port to use (default 8462)

-s

--no-sendfile

Make the network client NOT use sendfile() to transfer data

-o basename

--output=basename

Specifies base name for input files. Default is device.blk? trace.cpu. Specifying -o - runs in live mode with blkparse (writing data to standard out).

-D dir

--output-dir=dir

Prepend file to output file name(s)

This only works when supplying a single device, or when piping the output via "-o -" with multiple devices.

-r rel-path

--relay=rel-path

Specifies debugfs mount point

-v

--version

Outputs version

-V

--version

Outputs version

-w seconds

--stopwatch=seconds

Sets run time to the number of seconds specified

FILTER MASKS

The following masks may be passed with the -a command line option, mul?

tiple filters may be combined via multiple -a command line options.

barrier: barrier attribute

complete: completed by driver

discard: discard / trim traces

fs: requests

issue: issued to driver

pc: packet command events

queue: queue operations

read: read traces

requeue: requeue operations

sync: synchronous attribute

write: write traces

notify: trace messages

drv_data: additional driver specific trace

REQUEST TYPES

blktrace distinguishes between two types of block layer requests, file

system and SCSI commands. The former are dubbed fs requests, the latter

pc requests. File system requests are normal read/write operations,

i.e. any type of read or write from a specific disk location at a

given size. These requests typically originate from a user process, but they may also be initiated by the vm flushing dirty data to disk or the file system syncing a super or journal block to disk. pc requests are SCSI commands. blktrace sends the command data block as a payload so that blkparse can decode it.

EXAMPLES

To trace the i/o on the device /dev/sda and parse the output to human readable form, use the following command:

```
% blktrace -d /dev/sda -o - | blkparse -i -
```

This same behaviour can be achieve with the convenience script btrace. The command

```
% btrace /dev/sda
```

has exactly the same effect as the previous command. See btrace (8) for more information.

To trace the i/o on a device and save the output for later processing with blkparse, use blktrace like this:

```
% blktrace /dev/sda /dev/sdb
```

This will trace i/o on the devices /dev/sda and /dev/sdb and save the recorded information in the files sda and sdb in the current directory, for the two different devices, respectively. This trace information can later be parsed by the blkparse utility:

```
% blkparse sda sdb
```

which will output the previously recorded tracing information in human readable form to stdout. See blkparse (1) for more information.

AUTHORS

blktrace was written by Jens Axboe, Alan D. Brunelle and Nathan Scott. This man page was created from the blktrace documentation by Bas Zoetekouw.

REPORTING BUGS

Report bugs to <linux-btrace@vger.kernel.org>

COPYRIGHT

terms of the GNU General Public License <http://www.gnu.org/li?censes/gpl.html>. There is NO WARRANTY, to the extent permitted by law.

This manual page was created for Debian by Bas Zoetekouw. It was de?rived from the documentation provided by the authors and it may be used, distributed and modified under the terms of the GNU General Pub?lic License, version 2.

On Debian systems, the text of the GNU General Public License can be found in /usr/share/common-licenses/GPL-2.

SEE ALSO

btrace (8), blkparse (1), verify_blkparse (1), blkrawverify (1), btt (1)