



Rocky Enterprise Linux 9.2 Manual Pages on command 'sg_rbuf.8'

\$ man sg_rbuf.8

SG_RBUF(8) SG3_UTILS SG_RBUF(8)

NAME

sg_rbuf - reads data using SCSI READ BUFFER command

SYNOPSIS

```
sg_rbuf  [--buffer=EACH]  [--dio]  [--help]  [--mmap]  [--quick]
[--size=OVERALL] [--verbose] [--version] DEVICE
sg_rbuf [-b=EACH_KIB] [-d] [-m] [-q] [-s=OVERALL_MIB] [-t] [-v] [-V]
DEVICE
```

DESCRIPTION

This command reads data with the SCSI READ BUFFER command and then dis? cards it. Typically the data being read is from a disk's memory cache. It is assumed that the data is sourced quickly (although this is not guaranteed by the SCSI standards) so that it is faster than reading data from the media. This command is designed for timing transfer speeds across a SCSI transport.

To fetch the data with a SCSI READ BUFFER command and optionally decode it see the sg_read_buffer utility. There is also a sg_write_buffer utility useful for downloading firmware amongst other things.

This utility supports two command line syntaxes, the preferred one is shown first in the synopsis and explained in this section. A later section on the old command line syntax outlines the second group of options.

This is a Linux only utility and only works when DEVICE is an sg device (e.g. "/dev/sg1"). The sg_read_buffer utility has similar functionality and is ported to other OSes and within Linux can use bsg and normal block device names (e.g. "/dev/sdc").

OPTIONS

Arguments to long options are mandatory for short options as well.

-b, --buffer=EACH

where EACH is the number of bytes to be transferred by each READ BUFFER command. The default is the actual available buffer size returned by the READ BUFFER (descriptor) command. The maximum is the same as the default, hence this argument can only be used to reduce the size of each transfer to less than the device's actual available buffer size.

-d, --dio

use direct IO if available. This option is only available if the DEVICE is a sg driver device node (e.g. /dev/sg1). In this case the sg driver will attempt to configure the DMA from the SCSI adapter to transfer directly into user memory. This will eliminate the copy via kernel buffers. If not available then this will be reported and indirect IO will be done instead.

-h, --help

print usage message then exit.

-m, --mmap

use memory mapped IO if available. This option is only available if the DEVICE is a sg driver device node (e.g. /dev/sg1). In this case the sg driver will attempt to configure the DMA from the SCSI adapter to transfer directly into user memory. This will eliminate the copy via kernel buffers.

-O, --old

Switch to older style options. Please use as first option.

-q, --quick

only transfer the data into kernel buffers (typically by DMA from the SCSI adapter card) and do not move it into the user space. This option is only available if the DEVICE is a sg driver device node (e.g. /dev/sg1).

-s, --size=OVERALL

where OVERALL is the size of total transfer in bytes. The default is 200 MiB (200*1024*1024 bytes). The actual number of bytes transferred may be slightly less than requested since all transfers are the same size (and an integer division is involved rounding towards zero).

-t, --time

times the bulk data transfer component of this command. The elapsed time is printed out plus a MB/sec calculation. In this case "MB" is 1,000,000 bytes. The gettimeofday() system call is used internally for the time calculation.

-v, --verbose

increase level of verbosity. Can be used multiple times.

-V, --version

print out version string then exit.

NOTES

This command is typically used on modern SCSI disks which have a RAM cache in their drive electronics. If no IO to the magnetic media, or slower devices like flash RAM, is involved then the disk may be able to source data fast enough to saturate the bandwidth of the SCSI transport. The bottleneck may then be the DMA element in the HBA, the Linux drivers or the host machine's hardware (e.g. speed of RAM).

Various numeric arguments (e.g. OVERALL) may include multiplicative suffixes or be given in hexadecimal. See the "NUMERIC ARGUMENTS" section in the sg3_utils(8) man page.

EXAMPLES

On the test system /dev/sg0 corresponds to a fast disk on a U2W SCSI

bus (max 80 MB/sec). The disk specifications state that its cache is 4 MB.

```
$ time ./sg_rbuf /dev/sg0
```

READ BUFFER reports: buffer capacity=3434944,

offset boundary=6

Read 200 MiB (actual 199 MiB, 209531584 bytes),

buffer size=3354 KiB

real 0m5.072s, user 0m0.000s, sys 0m2.280s

So that is approximately 40 MB/sec at 40 % utilization. Now with the addition of the "-q" option this throughput improves and the utilization drops to 0%.

```
$ time ./sg_rbuf -q /dev/sg0
```

READ BUFFER reports: buffer capacity=3434944,

offset boundary=6

Read 200 MiB (actual 199 MiB, 209531584 bytes),

buffer size=3354 KiB

real 0m2.784s, user 0m0.000s, sys 0m0.000s

EXIT STATUS

The exit status of sg_rbuf is 0 when it is successful. Otherwise see the sg3_utils(8) man page.

OLDER COMMAND LINE OPTIONS

The options in this section were the only ones available prior to sg3_utils version 1.23 . Since then this utility defaults to the newer command line options which can be overridden by using --old (or -O) as the first option. See the ENVIRONMENT VARIABLES section for another way to force the use of these older command line options.

-b=EACH_KIB

where EACH_KIB is the number of Kilobytes (i.e. 1024 byte units) to be transferred by each READ BUFFER command. Similar to the --buffer=EACH option in the main description but the units are different.

-d use direct IO if available. Equivalent to the --dio option in the main description.

-m use memory mapped IO if available. Equivalent to the --mmap option in the main description.

-N, --new

Switch to the newer style options.

-q only transfer the data into kernel buffers (typically by DMA from the SCSI adapter card) and do not move it into the user space. Equivalent to the --quick option in the main description.

-s=OVERALL_MIB

where OVERALL_MIB is the size of total transfer in Megabytes (1048576 bytes). Similar to the --size=OVERALL option in the main description but the units are different.

-t times the bulk data transfer component of this command. Equivalent to the --time option in the main description.

-v increase level of verbosity. Can be used multiple times.

-V print out version string then exit.

ENVIRONMENT VARIABLES

Since sg3_utils version 1.23 the environment variable SG3_UTILS_OLD_OPTS can be given. When it is present this utility will expect the older command line options. So the presence of this environment variable is equivalent to using --old (or -O) as the first command line option.

AUTHOR

Written by Douglas Gilbert

REPORTING BUGS

Report bugs to <dgilbert at interlog dot com>.

COPYRIGHT

Copyright ? 2000-2017 Douglas Gilbert

This software is distributed under the GPL version 2. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

SEE ALSO

sg_read_buffer, sg_write_buffer, sg_test_rwbuf(all in sg3_utils)

