



Linux Ubuntu 26.04 Manual Pages on command 'trace-cmd-split.1'

\$ man trace-cmd-split.1

TRACE-CMD-SPLIT(1)

libtraceefs Manual

TRACE-CMD-SPLIT(1)

NAME

trace-cmd-split - split a trace.dat file into smaller files

SYNOPSIS

trace-cmd split [OPTIONS] [start-time [end-time]]

DESCRIPTION

The trace-cmd(1) split is used to break up a trace.dat into small files. The start-time specifies where the new file will start at. Using trace-cmd-report(1) and copying the time stamp given at a particular event, can be used as input for either start-time or end-time.

The split will stop creating files when it reaches an event after end-time. If only the end-time is needed, use 0.0 as the start-time.

If start-time is left out, then the split will start at the beginning of the file. If

end-time is left out, then split will continue to the end unless it meets one of the requirements specified by the options.

OPTIONS

-i file

If this option is not specified, then the split command will look for the file named trace.dat. This options will allow the reading of another file other than trace.dat.

-o file

By default, the split command will use the input file name as a basis of where to write the split files. The output file will be the input file with an attached '.#\' to the end: trace.dat.1, trace.dat.2, etc.

This option will change the name of the base file used.

-o file will create file.1, file.2, etc.

-s seconds

This specifies how many seconds should be recorded before the new file should stop.

-m milliseconds

This specifies how many milliseconds should be recorded before the new file should stop.

-u microseconds

This specifies how many microseconds should be recorded before the new file should stop.

-e events

This specifies how many events should be recorded before the new file should stop.

-p pages

This specifies the number of pages that should be recorded before the new file should stop.

Note: only one of *-p*, *-e*, *-u*, *-m*, *-s* may be specified at a time.

If *-p* is specified, then *-c* is automatically set.

-r

This option causes the break up to repeat until end-time is reached (or end of the input if end-time is not specified).

```
trace-cmd split -r -e 10000
```

This will break up trace.dat into several smaller files, each with at most

10,000 events in it.

-c

This option causes the above break up to be per CPU.

```
trace-cmd split -c -p 10
```

This will create a file that has 10 pages per each CPU from the input.

-C cpu

This option will split for a single CPU. Only the cpu named will be extracted from the file.

```
trace-cmd split -C 1
```

This will split out all the events for cpu 1 in the file.

--top

This allows to keep the top buffer. The top buffer can be renamed using the -b option.

```
trace-cmd split --top
```

This will keep only the top buffer.

```
trace-cmd split --top -b old_top
```

This will keep only the top buffer and rename it 'old_top'.

-B buffer

This allows to keep the selected buffer. A buffer can be promoted to the top buffer using the -t option.

```
trace-cmd split -B timer -B sched
```

This will keep the 'timer' and 'sched' buffers.

```
trace-cmd split -B timer -t -B sched
```

This will keep the 'timer' and 'sched' buffers, with the events from the 'timer' buffer promoted to the top instance.

SEE ALSO

trace-cmd(1), trace-cmd-record(1), trace-cmd-report(1), trace-cmd-start(1),
trace-cmd-stop(1), trace-cmd-extract(1), trace-cmd-reset(1), trace-cmd-list(1),
trace-cmd-listen(1)

AUTHOR

Written by Steven Rostedt, <rostedt@goodmis.org[1]>

RESOURCES

<https://git.kernel.org/pub/scm/utils/trace-cmd/trace-cmd.git/>

COPYING

Copyright (C) 2010 Red Hat, Inc. Free use of this software is granted under the terms of the GNU Public License (GPL).

NOTES

1. rostedt@goodmis.org

<mailto:rostedt@goodmis.org>

libtracefs

01/22/2026

TRACE-CMD-SPLIT(1)