



## ***Linux Ubuntu 26.04 Manual Pages on command 'perf-ftrace.1'***

***\$ man perf-ftrace.1***

PERF-FTRACE(1)                      perf Manual                      PERF-FTRACE(1)

### NAME

perf-ftrace - simple wrapper for kernel's ftrace functionality

### SYNOPSIS

perf ftrace {trace|latency|profile} <command>

### DESCRIPTION

The perf ftrace command provides a collection of subcommands which use kernel's ftrace infrastructure.

'perf ftrace trace' is a simple wrapper of the ftrace. It only supports single thread tracing currently and just reads trace\_pipe in text and then write it to stdout.

'perf ftrace latency' calculates execution latency of a given function (optionally with BPF) and display it as a histogram.

'perf ftrace profile' show a execution profile for each function including total, average, max time and the number of calls.

The following options apply to perf ftrace.

### COMMON OPTIONS

-p, --pid=

Trace on existing process id (comma separated list).

--tid=

Trace on existing thread id (comma separated list).

-a, --all-cpus

Force system-wide collection. Scripts run without a <command> normally use -a by

default, while scripts run with a <command> normally don't - this option allows the latter to be run in system-wide mode.

-C, --cpu=

Only trace for the list of CPUs provided. Multiple CPUs can be provided as a comma separated list with no space like: 0,1. Ranges of CPUs are specified with -: 0-2.

Default is to trace on all online CPUs.

-v, --verbose

Increase the verbosity level.

## OPTIONS FOR PERF FTRACE TRACE

-t, --tracer=

Tracer to use when neither -G nor -F option is not specified: function\_graph or function.

-F, --funcs

List available functions to trace. It accepts a pattern to only list interested functions.

-D, --delay

Time (ms) to wait before starting tracing after program start.

-m, --buffer-size

Set the size of per-cpu tracing buffer, <size> is expected to be a number with appended unit character - B/K/M/G.

--inherit

Trace children processes spawned by our target.

-T, --trace-funcs=

Select function tracer and set function filter on the given function (or a glob pattern). Multiple functions can be given by using this option more than once. The function argument also can be a glob pattern. It will be passed to set\_ftrace\_filter in tracefs.

-N, --notrace-funcs=

Select function tracer and do not trace functions given by the argument. Like -T option, this can be used more than once to specify multiple functions (or glob patterns). It will be passed to set\_ftrace\_notrace in tracefs.

--func-opts

List of options allowed to set:

? call-graph - Display kernel stack trace for function tracer.

? irq-info - Display irq context info for function tracer.

-G, --graph-funcs=

Select function\_graph tracer and set graph filter on the given function (or a glob pattern). This is useful to trace for functions executed from the given function. This can be used more than once to specify multiple functions. It will be passed to set\_graph\_function in tracefs.

-g, --nograph-funcs=

Select function\_graph tracer and set graph notrace filter on the given function (or a glob pattern). Like -G option, this is useful for the function\_graph tracer only and disables tracing for function executed from the given function. This can be used more than once to specify multiple functions. It will be passed to set\_graph\_notrace in tracefs.

--graph-opts

List of options allowed to set:

? args - Show function arguments.

? retval - Show function return value.

? retval-hex - Show function return value in hexadecimal format.

? retaddr - Show function return address.

? nosleep-time - Measure on-CPU time only for function\_graph tracer.

? noirqs - Ignore functions that happen inside interrupt.

? verbose - Show process names, PIDs, timestamps, etc.

? thresh=<n> - Setup trace duration threshold in microseconds.

? depth=<n> - Set max depth for function graph tracer to follow.

? tail - Print function name at the end.

## OPTIONS FOR PERF FTRACE LATENCY

-T, --trace-funcs=

Set the function name to get the histogram. Unlike perf ftrace trace, it only allows single function to calculate the histogram.

-e, --events=

Set the pair of events to get the histogram. The histogram is calculated by the time difference between the two events from the same thread. This requires -b/--use-bpf option.

**-b, --use-bpf**

Use BPF to measure function latency instead of using the ftrace (it uses function\_graph tracer internally).

**-n, --use-nsec**

Use nano-second instead of micro-second as a base unit of the histogram.

**--bucket-range=**

Bucket range in ms or ns (according to -n/--use-nsec), default is log2() mode.

**--min-latency=**

Minimum latency for the start of the first bucket, in ms or ns (according to -n/--use-nsec).

**--max-latency=**

Maximum latency for the start of the last bucket, in ms or ns (according to -n/--use-nsec). The setting is ignored if the value results in more than 22 buckets.

## OPTIONS FOR PERF FTRACE PROFILE

**-T, --trace-funcs=**

Set function filter on the given function (or a glob pattern). Multiple functions can be given by using this option more than once. The function argument also can be a glob pattern. It will be passed to set\_ftrace\_filter in tracefs.

**-N, --notrace-funcs=**

Do not trace functions given by the argument. Like -T option, this can be used more than once to specify multiple functions (or glob patterns). It will be passed to set\_ftrace\_notrace in tracefs.

**-G, --graph-funcs=**

Set graph filter on the given function (or a glob pattern). This is useful to trace for functions executed from the given function. This can be used more than once to specify multiple functions. It will be passed to set\_graph\_function in tracefs.

**-g, --nograph-funcs=**

Set graph notrace filter on the given function (or a glob pattern). Like -G option, this is useful for the function\_graph tracer only and disables tracing for function executed from the given function. This can be used more than once to specify multiple functions. It will be passed to set\_graph\_notrace in tracefs.

**-m, --buffer-size**

Set the size of per-cpu tracing buffer, <size> is expected to be a number with appended

unit character - B/K/M/G.

-s, --sort=

Sort the result by the given field. Available values are: total, avg, max, count, name.

Default is total.

--graph-opts

List of options allowed to set:

? nosleep-time - Measure on-CPU time only for function\_graph tracer.

? noirqs - Ignore functions that happen inside interrupt.

? thresh=<n> - Setup trace duration threshold in microseconds.

? depth=<n> - Set max depth for function graph tracer to follow.

SEE ALSO

perf-record(1), perf-trace(1)

perf

05/25/2026

PERF-FTRACE(1)