



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

\$ man mcelog.8

MCELOG(8) Linux's Administrator's Manual MCELOG(8)

NAME

mcelog - Decode kernel machine check log on x86 machines

SYNOPSIS

mcelog [options] [device]

mcelog [options] --daemon

mcelog [options] --client

mcelog [options] --ascii

mcelog [options] --is-cpu-supported

mcelog --version

DESCRIPTION

X86 CPUs report errors detected by the CPU as machine check events

(MCEs). These can be data corruption detected in the CPU caches, in

main memory by an integrated memory controller, data transfer errors on

the front side bus or CPU interconnect or other internal errors. Pos?

sible causes can be cosmic radiation, instable power supplies, cooling

problems, broken hardware, running systems out of specification, or bad

luck.

Most errors can be corrected by the CPU by internal error correction

mechanisms. Uncorrected errors cause machine check exceptions which may

kill processes or panic the machine. A small number of corrected errors

is usually not a cause for worry, but a large number can indicate fu?

ture failure.

When a corrected or recovered error happens, the x86 kernel writes a record describing the MCE into a internal ring buffer available through the /dev/mcelog device. mcelog retrieves errors from /dev/mcelog, decodes them into a human readable format and prints them on the standard output or optionally into the system log.

Optionally it can also take more options like keeping statistics or triggering shell scripts on specific events. By default mcelog supports offlining memory pages with persistent corrected errors, offlining CPU cores if they developed cache problems, and otherwise logging specific events to the system log after they crossed a threshold.

The normal operating modes for mcelog are: running as a regular cron job (traditional way, deprecated), running as a trigger directly executed by the kernel, or running as a daemon with the --daemon option.

When an uncorrected machine check error happens that the kernel cannot recover from then it will usually panic the system. In this case when there was a warm reset after the panic mcelog should pick up the machine check errors after reboot. This is not possible after a cold reset.

In addition mcelog can be used on the command line to decode the kernel output for a fatal machine check panic in text format using the --ascii option. This is typically used to decode the panic console output of a fatal machine check, if the system was power cycled or mcelog didn't run immediately after reboot.

When the panic triggers a kdump kexec crash kernel the crash kernel boot up script should log the machine checks to disk, otherwise they might be lost.

Note that after mcelog retrieves an error the kernel doesn't store it anymore (different from dmesg(1)), so the output should be always saved somewhere and mcelog not run in uncontrolled ways.

When invoked with the --is-cpu-supported option mcelog exits with code 0 if the current CPU is supported, 1 otherwise.

OPTIONS

When the --syslog option is specified redirect output to system log.

The --syslog-error option causes the normal machine checks to be logged as LOG_ERR (implies --syslog). Normally only fatal errors or high level remarks are logged with error level. High level one line summaries of specific errors are also logged to the syslog by default unless mcelog operates in --ascii mode.

When the --logfile=file option is specified append log output to the specified file. With the --no-syslog option mcelog will never log anything to the syslog.

When the --cpu=cpotype option is specified set the to be decoded CPU to cpotype. See mcelog --help for a list of valid CPUs. Note that specifying an incorrect CPU can lead to incorrect decoding output. Default is either the CPU of the machine that reported the machine check (needs a newer kernel version) or the CPU of the machine mcelog is running on, so normally this option doesn't have to be used. Older versions of mcelog had separate options for different CPU types. These are still implemented, but deprecated and undocumented now.

With the --dmi option mcelog will look up the DIMMs reported in machine checks in the SMBIOS/DMI tables of the BIOS and map the DIMMs to board identifiers. This only works when the BIOS reports the identifiers correctly. Unfortunately often the information reported by the BIOS is either subtly or obviously wrong or useless. This option requires that mcelog has read access to /dev/mem (normally requires root) and runs on the same machine in the same hardware configuration as when the machine check event happened.

When --ignorenodev is specified then mcelog will exit silently when the device cannot be opened. This is useful in virtualized environment with limited devices.

When --filter is specified mcelog will filter out known broken machine check events (default on). When the --no-filter option is specified mcelog does not filter events.

When --raw is specified mcelog will not decode, but just dump the mcelog in a raw hex format. This can be useful for automatic post processing.

When a device is specified the machine check logs are read from device instead of the default /dev/mcelog.

With the --ascii option mcelog decodes a fatal machine check panic generated by the kernel ("CPU n: Machine Check Exception ...") in ASCII from standard input and exits afterwards. Note that when the panic comes from a different machine than where mcelog is running on you might need to specify the correct cputype on older kernels. On newer kernels which output the PROCESSOR field this is not needed anymore.

When the --file filename option is specified mcelog --ascii will read the ASCII machine check record from input file filename instead of standard input.

With the --config-file file option mcelog reads the specified config file. Default is /etc/mcelog/mcelog.conf See also CONFIG FILE below.

With the --daemon option mcelog will run in the background. This gives the fastest reaction time and is the recommended operating mode. If an output option isn't selected (--logfile or --syslog or --syslog-error), this option implies --logfile=/var/log/mcelog. Important messages will be logged as one-liner summaries to syslog unless --no-syslog is given. The option --foreground will prevent mcelog from giving up the terminal in daemon mode. This is intended for debugging.

With the --client option mcelog will query a running daemon for accumulated errors.

With the --cpumhz=mhz option assume the CPU has mhz frequency for decoding the time of the event using the CPU time stamp counter. This also forces decoding. Note this can be unreliable. on some systems with CPU frequency scaling or deep C states, where the CPU time stamp counter does not increase linearly. By default the frequency of the current CPU is used when mcelog determines it is safe to use. Newer kernels report the time directly in the event and don't need this anymore.

The --pidfile file option writes the process id of the daemon into file file. Only valid in daemon mode.

Mcelog will enable extended error reporting from the memory controller

on processors that support it unless you tell it not to with the --no-imc-log option. You might need this option when decoding old logs from a system where this mode was not enabled.
--version displays the version of mcelog and exits.

CONFIG FILE

mcelog supports a config file to set defaults. Command line options override the config file. By default the config file is read from /etc/mcelog/mcelog.conf unless overridden with the --config-file option.

The general format is optionname = value White space is not allowed in value currently, except at the end where it is dropped Comments start with #.

All command line options that are not commands can be specified in the config file. For example t to enable the --no-syslog option use no-syslog = yes (or no to disable). When the option has a argument use logfile = /tmp/logfile

For more information on the config file please see mcelog.conf(5).

NOTES

The kernel prefers old messages over new. If the log buffer overflows only old ones will be kept.

The exact output in the log file depends on the CPU, unless the --raw option is used.

mcelog will report serious errors to the syslog during decoding.

SIGNALS

When mcelog runs in daemon mode and receives a SIGUSR1 it will close and reopen the log files. This can be used to rotate logs without restarting the daemon.

FILES

/dev/mcelog (char 10, minor 227)

/etc/mcelog/mcelog.conf

/var/log/mcelog

/var/run/mcelog.pid

SEE ALSO

mcelog.conf(5), mcelog.triggers(5)

<http://www.mcelog.org>

AMD x86-64 architecture programmer's manual, Volume 2, System program?

ming

Intel 64 and IA32 Architectures Software Developer's manual, Volume 3,

System programming guide Chapter 15 and 16. <http://www.intel.com/sdm>

Datasheet of your CPU.

Mar 2015

MCELOG(8)