



## Dmesg and Linux kernel log



# Dmesg and Linux kernel log

Stable: 27.01.2020 - 13:03 / Revision: 27.01.2020 - 10:22

## Contents

1 Article purpose .....	2
2 Introduction .....	2
3 printk function .....	2
4 Linux kernel ring buffer .....	3
5 Loglevels .....	3
<b>5.1 loglevels values .....</b>	<b>4</b>
<b>5.2 Set loglevel filter value for console .....</b>	<b>5</b>
5.2.1 Default values .....	5
5.2.2 Using kernel command-line .....	5
5.2.3 Using sysfs in runtime .....	6
5.2.4 Using menuconfig before compilation .....	6
<b>5.3 Use loglevel in kernel source for log and trace .....</b>	<b>7</b>
5.3.1 Using printk .....	7
5.3.2 Using dedicated functions .....	7
6 earlyprintk .....	8
7 dmesg command .....	10
8 /var/log/messages file system entry .....	10
9 Dynamic debug message .....	11
10 References .....	11

## 1 Article purpose

The purpose of this article is to provide information about the Linux<sup>®</sup> kernel log including configuration, and to detail usage of dmesg command.

## 2 Introduction

Linux kernel is able to print log and trace messages, which are by default stored in a ring buffer.

The same messages can also be displayed, applying filter, on uart/console using serial port. This is defined in the kernel command-line, with the "console" parameter. See <sup>[1]</sup> for detail.

dmesg is a shell command on the kernel console, which also displays the content of the ring buffer, with filter or not (default).

## 3 printk function



The simplest way to get some debug information from the kernel code is by printing out various information with the kernel's equivalent of printf - the printk function and its derivatives.

```
printk("My Debugger is Printk\n");
```

See [elinux.org](http://elinux.org)<sup>[2]</sup> for reference. This information will be sent to the console, and also stored in a ring buffer.

You can also check to the [printk-format.txt](#)<sup>[3]</sup> document provided in the Linux kernel package to get detail about syntax and formatting.

## 4 Linux kernel ring buffer

The Linux kernel also manages a ring buffer to store log and trace messages.

The size of the buffer cannot be modified in runtime, and its default size value is  $2^{\text{CONFIG\_LOG\_BUF\_SHIFT}}$  bytes.

To change it, there are 3 possible ways:

- Modify CONFIG\_LOG\_BUF\_SHIFT value in defconfig file or use the config fragment file:

```
In example for 64K : CONFIG_LOG_BUF_SHIFT=16
```

- or use the Linux kernel [menuconfig](#) update

Location:

```
-> General setup  
-> Kernel log buffer size (16 => 64KB, 17 => 128KB)
```

- Or modify kernel arguments<sup>[4]</sup> in kernel command-line (via bootargs value in device tree, or directly in extlinux uboot config file)

```
bootargs = "root=/dev/mmcblk0p5 rootwait rw console=ttySTM0,115200 log_buf_len=65536";
```

This ring buffer can be displayed using *dmesg* command (see [dmesg](#)).

## 5 Loglevels

As reference, please see [elinux.org](http://elinux.org)<sup>[5]</sup>.

The log level is used by the kernel to determine the importance of a message and to decide whether it should be presented to the user immediately, by printing it to the current console.

For this, the kernel compares the log level of the message to the console\_loglevel (a kernel variable) and if the priority is higher (i.e. a lower value) than the console\_loglevel, the message will be printed to the current console. As example, if console\_loglevel=5, all messages with log level 0 to 4 will be displayed.



Please note that all messages with loglevel lower or equal to KERN\_INFO level are stored in the ring buffer.

## 5.1 loglevels values

Name	String	Meaning	alias functions	dev alias function
KERN_EMERG	"0"	<b>Emergency messages, system is about to crash or is unstable</b>	pr_emerg	dev_emerg
KERN_ALERT	"1"	Something bad happened and action must be taken immediately	pr_alert	dev_alert
KERN_CRIT	"2"	A critical condition occurred like a serious hardware /software failure	pr_crit	dev_crit
KERN_ERR	"3"	An error condition, often used by drivers to indicate difficulties with the hardware	pr_err	dev_err
KERN_WARNING	"4"	A warning, meaning nothing serious by itself but might indicate problems	pr_warning	dev_warn
KERN_NOTICE	"5"	Nothing serious, but notably nevertheless. Often used to report security events	pr_notice	dev_notice
KERN_INFO	"6"	Informational message e.g. startup information at driver initialization	pr_info	dev_info



Name	String	Meaning	alias functions	dev alias function
KERN_DEBUG	"7"	Debug messages	pr_debug, pr_devel if DEBUG is defined	dev_dbg

### "Loglevels table"

Important: please note that Higher priority message is loglevel 0

## 5.2 Set loglevel filter value for console

### 5.2.1 Default values

To determine your current console\_loglevel on the target you can verify with the following command:

```
Board $> cat /proc/sys/kernel/printk
7      4      1      7
current default_msg minimum default_console
```

The first integer shows you the current console\_loglevel; the second the default log level, see Use loglevel in the kernel source for log and trace.

This is defined at compilation:

- Current console loglevel via CONFIG\_CONSOLE\_LOGLEVEL\_DEFAULT=7 (defined in file lib/Kconfig.debug)
- Default message loglevel via CONFIG\_MESSAGE\_LOGLEVEL\_DEFAULT=4 (defined in file lib/Kconfig.debug)
- Minimum console loglevel via #define CONSOLE\_LOGLEVEL\_MIN 1 (defined in file include/linux/printk.h)
- Default console loglevel is equal to CONFIG\_CONSOLE\_LOGLEVEL\_DEFAULT

### 5.2.2 Using kernel command-line

The console loglevel can be also set via a kernel command-line parameter if you want to use a different value than one specify by CONFIG\_CONSOLE\_LOGLEVEL\_DEFAULT.

For example:

```
root=/dev/mmcblk0p5 rootwait rw console=ttySTM0,115200 loglevel=4
```

In that case only messages with a higher priority than KERN\_WARNING (means < 4, KERN\_EMERG to KERN\_ERR ) will be displayed on the console.

2 ways to add this command-line parameter which is set in the extlinux.conf file of boot partition:

- If using SD card, this is possible to edit the file on host PC:

```
Insert SD card on host PC
Check for mounting boot partition (i.e. /media/$USER/bootfs)
Check for your HW config (i.e. booting on mmc0 (SD Card) with ev1 board)
PC $> cd /media/$USER/bootfs/mmc0_stm32mp157c-ev1_extlinux/
PC $> gedit extlinux.conf
Add loglevel=8 at the end of APPEND line
Save and insert SD card on the board
```



- If using SD Card or eMMC, this is possible to edit the file directly on the board side:

```
When software is boot
Mount boot partition
Board $> mount /dev/mmcblk0p4 /boot (if not already done)
Update the kernel command line
Board $> cd /boot
Board $> cd mmc0_stm32mp157c-ev1_extlinux (case SD card on ev1 board)
Modify extlinux.conf to add loglevel=8 at the end of APPEND line by using a vi editor
Save and reboot the board
```

### 5.2.3 Using sysfs in runtime

To change your current console\_loglevel simply write to this file:

```
Board $> echo <loglevel> > /proc/sys/kernel/printk
```

or using `dmesg` command.

As example:

```
Board $> echo 8 > /proc/sys/kernel/printk           # Temporary increase loglevel to display
```

In that case, every kernel messages will appear on your console, as **all priority higher than 8 (lower loglevel values)** will be displayed.

Please note that after reboot, this configuration is reset.

### 5.2.4 Using menuconfig before compilation

As values are defined first at compilation step, this is also possible to set them ( `CONFIG_CONSOLE_LOGLEVEL_DEFAULT` and `CONFIG_MESSAGE_LOGLEVEL_DEFAULT`) using the Linux kernel Menuconfig tool (Menuconfig or how to configure kernel):

```
Symbol: CONSOLE_LOGLEVEL_DEFAULT [=7]
Location:
  Kernel hacking --->
    printk and dmesg options --->
      (7) Default console loglevel (1-15)

Symbol: MESSAGE_LOGLEVEL_DEFAULT [=4]
Location:
  Kernel hacking --->
    printk and dmesg options --->
      (4) Default message log level (1-7)
```



## 5.3 Use loglevel in kernel source for log and trace

### 5.3.1 Using printk

A loglevel information can be added in the `printk` function call, with the following syntax.

```
printk(KERN_ERR "something went wrong, return code: %d\n",ret);
```

When not present, default loglevel value is given by `CONFIG_MESSAGE_LOGLEVEL_DEFAULT` (usually "4" =`KERN_WARNING`)

### 5.3.2 Using dedicated functions

In the loglevels table above, there are some alias functions `pr_` and `dev_`.

These functions are defined to replace `printk + loglevel info inside`, in order to simplify syntax.

```
pr_err("something went wrong, return code: %d\n",ret);
```

`dev_` functions are taken one more parameter to provide more information about current device or driver where message is coming from.

- Example for `pr_info`

```
pr_info("%S%S%S at %s (irq = %d, base_baud = %d) is a %s\n",
        port->dev ? dev_name(port->dev) : "",
        port->dev ? ":" : "",
        port->name,
        address, port->irq, port->uartclk / 16, uart_type(port));
```

will display information below:

```
[ 0.919488] 40010000.serial: ttySTM0 at MMIO 0x40010000 (irq = 41, base_baud = 6046875) is
```

- Example for `dev_info`

```
dev_info(&pdev->dev, "interrupt mode used for rx (no dma)\n");
```

will display information below, including device reference automatically:

```
[ 1.046700] stm32-usart 40010000.serial: interrupt mode used for rx (no dma)
```



## 6 earlyprintk

earlyprintk is a Linux kernel debug feature useful to get traces for kernel issues which happen before the normal console is initialized.

- Linux kernel configuration

In order to enable earlyprintk feature, the Linux kernel configuration must activate **CONFIG\_DEBUG\_LL**, **CONFIG\_STM32MP1\_DEBUG\_UART** and **CONFIG\_EARLY\_PRINTK** using the Linux kernel Menuconfig tool (Menuconfig or how to configure kernel):

```
Symbol: DEBUG_LL
Location:
  Kernel hacking --->
    [*] Kernel low-level debugging functions

Symbol: STM32MP1_DEBUG_UART
Location:
  Kernel hacking --->
    [*] Kernel low-level debugging functions
      (*) Use STM32MP1 UART for low-level debug

Symbol: EARLY_PRINTK
Location:
  Kernel hacking --->
    [*] Early printk
```

- Serial port configuration

When enabling the Linux kernel configuration **CONFIG\_STM32MP1\_DEBUG\_UART**, it configures the addresses of the UART registers to be used.

By default, on STM32MP1 boards, UART4 is used for console for Linux kernel and by extension at all boot stages.

In case the UART port is different on a new board, you must apply the following changes:

- Update value for **CONFIG\_DEBUG\_UART\_PHYS**, to select the UART port for the debug console

```
Symbol: DEBUG_UART_PHYS [=0x40010000]
Location:
  Kernel hacking --->
    [*] Kernel low-level debugging port (Use STM32MP1 UART4 for low-level debug)
      (0x40010000) Physical base address of debug UART
```

- Update value for **CONFIG\_DEBUG\_UART\_VIRT**, to define the associated virtual address to be used





```
Symbol: DEBUG_UART_VIRT [=0xFE010000]
Location:
  Kernel hacking --->
  [*] Kernel low-level debugging port (Use STM32MP1 UART4 for low-level debug)
  (0xFE010000) Virtual base address of debug UART
```

**Following rules to be respected for defining the virtual address:**

- The 20 low weight bits (21 in case LPAE is enabled) must be kept in order to align region size of 1MB (2MB in LPAE is enabled).
- It must be mapped at the upper address of the vmlalloc area, in order to not be overwritten by kernel which is starting from lower addresses: i.e here we select 0xFE0xxxxx

```
CONFIG_DEBUG_UART_PHYS: 0x40010000 /* UART4 */
CONFIG_DEBUG_UART_VIRT: 0xFE010000
```

- Please find below table for USART/UART of STMP32MP1:

Name	Physical base address	Virtual base address
<b>USART1</b>	<b>5c000000</b>	<b>FE000000</b>
USART2	4000e000	FE00e000
USART3	4000f000	FE00f000
UART4	40010000	FE010000
UART5	40011000	FE011000
USART6	44003000	FE003000
UART7	40018000	FE018000
UART8	40019000	FE019000



**Note that the UART port used for console must be aligned for all components of the boot chain: FSBL(TF-A), SSBL(U-Boot) and Linux kernel**

**Especially because the Linux kernel do not configure all the setting registers for the UART port, as this is done by FSBL (U-Boot - How to debug)**

**See also TF-A - How to debug for FSBL changes)**

- Get trace

Earlyprintk traces are pushed automatically to the serial console defined as seen previously, and also added to the kernel ring log buffer.



## 7 dmesg command

As reference, please see man page<sup>[6]</sup>.

The Kernel ring buffer can be displayed using `dmesg` command. It will display on the console all the content of the ring buffer.

- It is possible to filter messages following the loglevels:

```
Board $> dmesg -n <loglevel>
```

In that case, only messages with a value lower (**not lower equal**) than the `console_loglevel` will be printed.

Here, `<loglevel>` can be a numeric value, but also a string:

```
Supported log levels (priorities):
emerg (0)
alert (1)
crit (2)
err (3)
warn (4)
notice (5)
info (6)
debug (7)
```

As example:

```
Board $> dmesg -n 8          # Temporary change loglevel to display messages up to debug level
OR
Board $> dmesg -n debug
```

In that case, every kernel messages will appear on your console, as **all priority higher than 8 (lower loglevel values)** will be displayed.

- It is possible to clear the dmesg buffer

```
Board $> dmesg -c          # Display the full content of dmesg ring buffer, and then clear it
Board $> dmesg -C          # Clear the dmesg ring buffer
```

## 8 /var/log/messages file system entry

An other way to display the content of the Linux kernel log is to look at the content of the file `/var/log/messages`.

It contains general system activity messages from the start-up. It also provides useful information about origin of the message, and log level.



---

## 9 Dynamic debug message

---

These messages are using the loglevel 7 (KERN\_DEBUG).

Please see [How to use the kernel dynamic debug article](#).

---

## 10 References

---

- <https://www.kernel.org/doc/html/latest/admin-guide/serial-console.html>
- [https://elinux.org/Debugging\\_by\\_printing](https://elinux.org/Debugging_by_printing)
- <https://www.kernel.org/doc/Documentation/printk-formats.txt>
- <https://www.kernel.org/doc/html/latest/admin-guide/kernel-parameters.html>
- [https://elinux.org/Debugging\\_by\\_printing#Log\\_Levels](https://elinux.org/Debugging_by_printing#Log_Levels)
- <http://man7.org/linux/man-pages/man1/dmesg.1.html>

SD memory card (<https://www.sdcard.org>) - NEW

former spelling for eMMC ('e' in italic)

System File System (See <https://en.wikipedia.org/wiki/Sysfs> for more details)

Low layer of STM32Cube

Universal Asynchronous Receiver/Transmitter

Universal Synchronous/Asynchronous Receiver/Transmitter

First Stage Boot Loader

Trusted Firmware for Arm Cortex-A

Second Stage Boot Loader

Das U-Boot -- the Universal Boot Loader (see [U-Boot\\_overview](#)) - NEW