

The Linux Kernel

7.3.0

Quick search

Go

Contents [Show]

This Page

Show Source

Linux Magic System Request Key Hacks

English

Documentation for sysrq.c

What is the magic SysRq key?

It is a ‘magical’ key combo you can hit which the kernel will respond to regardless of whatever else it is doing, unless it is completely locked up.

How do I enable the magic SysRq key?

You need to say “yes” to ‘Magic SysRq key (CONFIG_MAGIC_SYSRQ)’ when configuring the kernel. When running a kernel with SysRq compiled in, /proc/sys/kernel/sysrq controls the functions allowed to be invoked via the SysRq key. The default value in this file is set by the CONFIG_MAGIC_SYSRQ_DEFAULT_ENABLE config symbol, which itself defaults to 1. Here is the list of possible values in /proc/sys/kernel/sysrq:

- 0 - disable sysrq completely
- 1 - enable all functions of sysrq
- >1 - bitmask of allowed sysrq functions (see below for detailed function description):

2 = 0x2 - enable control of console logging level

4 = 0x4 - enable control of keyboard (SAK, unraw)

8 = 0x8 - enable debugging dumps of processes etc.

16 = 0x10 - enable sync command

32 = 0x20 - enable remount read-only

64 = 0x40 - enable signalling of processes (term, kill, oom-kill)

128 = 0x80 - allow reboot/poweroff

256 = 0x100 - allow nicing of all RT tasks

You can set the value in the file by the following command:

```
echo "number" >/proc/sys/kernel/sysrq
```

The number may be written here either as decimal or as hexadecimal with the 0x prefix. CONFIG_MAGIC_SYSRQ_DEFAULT_ENABLE must always be written in hexadecimal.

Note that the value of /proc/sys/kernel/sysrq influences only the invocation via a keyboard. Invocation of any operation via /proc/sysrq-trigger is always allowed (by a user with admin privileges).

How do I use the magic SysRq key?

On x86

You press the key combo *ALT-SysRq-<command key>*.

Note:

Some keyboards may not have a key labeled ‘SysRq’. The ‘SysRq’ key is also known as the ‘Print Screen’ key. Also some keyboards cannot handle so many keys being pressed at the same time, so you might have better luck with press *Alt*, press *SysRq*, release *SysRq*, press *<command key>*, release everything.

On SPARC

You press *ALT-STOP-<command key>*, I believe.

On the serial console (PC style standard serial ports only)

You send a **BREAK**, then within 5 seconds a command key. Sending **BREAK** twice is interpreted as a normal BREAK.

On PowerPC

Press *ALT - Print Screen* (or *F13*) - *<command key>*. *Print Screen* (or *F13*) - *<command key>* may suffice.

On other

If you know of the key combos for other architectures, please submit a patch to be included in this section.

On all

Write a single character to /proc/sysrq-trigger. Only the first character is processed, the rest of the string is ignored. However, it is not recommended to write any extra characters as the behavior is undefined and might change in the future versions. E.g.:

```
echo t > /proc/sysrq-trigger
```

Alternatively, write multiple characters prepended by underscore. This way, all characters will be processed. E.g.:

```
echo _reisub > /proc/sysrq-trigger
```

The *<command key>* is case sensitive.

What are the ‘command’ keys?

Command	Function
b	Will immediately reboot the system without syncing or unmounting your disks.
c	Will perform a system crash and a crashdump will be taken if configured.
d	Shows all locks that are held.
e	Send a SIGTERM to all processes, except for init.
f	Will call the oom killer to kill a memory hog process, but do not panic if nothing can be killed.
g	Used by kgdb (kernel debugger)
h	Will display help (actually any other key than those listed here will display help. but h is easy to remember :-)
i	Send a SIGKILL to all processes, except for init.
j	Forcibly “Just thaw it” - filesystems frozen by the FIFREEZE ioctl.
k	Secure Access Key (SAK) Kills all programs on the current virtual console. NOTE: See important comments below in SAK section.
l	Shows a stack backtrace for all active CPUs.
m	Will dump current memory info to your console.
n	Used to make RT tasks nice-able
o	Will shut your system off (if configured and supported).
p	Will dump the current registers and flags to your console.
q	Will dump per CPU lists of all armed hrtimers (but NOT regular timer_list timers) and detailed information about all clockevent devices.
r	Turns off keyboard raw mode and sets it to XLATE.
s	Will attempt to sync all mounted filesystems.
t	Will dump a list of current tasks and their information to your console.
u	Will attempt to remount all mounted filesystems read-only.
v	Forcefully restores framebuffer console
v	Causes ETM buffer dump [ARM-specific]
w	Dumps tasks that are in uninterruptible (blocked) state.
x	Used by xmon interface on ppc/powerpc platforms. Show global PMU Registers on sparc64. Dump all TLB entries on MIPS.
y	Show global CPU Registers [SPARC-64 specific]
z	Dump the ftrace buffer
0-9	Sets the console log level, controlling which kernel messages will be printed to your console. (0 , for example would make it so that only emergency messages like PANICs or OOPSes would make it to your console.)
R	Replay the kernel log messages on consoles.

Okay, so what can I use them for?

Well, unraw(r) is very handy when your X server or a svgalib program crashes.

sak(k) (Secure Access Key) is useful when you want to be sure there is no trojan program running at console which could grab your password when you would try to login. It will kill all programs on given console, thus letting you make sure that the login prompt you see is actually the one from init, not some trojan program.

Important:

In its true form it is not a true SAK like the one in a c2 compliant system, and it should not be mis-taken as such.

It seems others find it useful as (System Attention Key) which is useful when you want to exit a program that will not let you switch consoles. (For example, X or a svgalib program.)

reboot (b) is good when you’re unable to shut down, it is an equivalent of pressing the “reset” button.

crash (c) can be used to manually trigger a crashdump when the system is hung. Note that this just triggers a crash if there is no dump mechanism available.

sync (s) is handy before yanking removable medium or after using a rescue shell that provides no graceful shutdown -- it will ensure your data is safely written to the disk. Note that the sync hasn’t taken place until you see the “OK” and “Done” appear on the screen.

umount (u) can be used to mark filesystems as properly unmounted. From the running system’s point of view, they will be remounted read-only. The remount isn’t complete until you see the “OK” and “Done” message appear on the screen.

The loglevels **0-9** are useful when your console is being flooded with kernel messages you do not want to see. Selecting **0** will prevent all but the most urgent kernel messages from reaching your console. (They will still be logged if syslogd/klogd are alive, though.)

term (e) and **kill (i)** are useful if you have some sort of runaway process you are unable to kill any other way, especially if it’s spawning other processes.

“just thaw **it (j)**” is useful if your system becomes unresponsive due to a frozen (probably root) filesystem via the FIFREEZE ioctl.

Replay logs (R) is useful to view the kernel log messages when system is hung or you are not able to use dmesg command to view the messages in printk buffer. User may have to press the key combination multiple times if console system is busy. If it is completely locked up, then messages won’t be printed. Output messages depend on current console loglevel, which can be modified using sysrq[0-9] (see above).

Sometimes SysRq seems to get ‘stuck’ after using it, what can I do?

When this happens, try tapping shift, alt and control on both sides of the keyboard, and hitting an invalid sysrq sequence again. (i.e., something like *alt-sysrq-z*).

Switching to another virtual console (*ALT+Fn*) and then back again should also help.

I hit SysRq, but nothing seems to happen, what's wrong?

There are some keyboards that produce a different keycode for SysRq than the pre-defined value of 99 (see `KEY_SYSRQ` in `include/uapi/linux/input-event-codes.h`), or which don't have a SysRq key at all. In these cases, run `showkey -s` to find an appropriate scancode sequence, and use `setkeycodes <sequence> 99` to map this sequence to the usual SysRq code (e.g., `setkeycodes e05b 99`). It's probably best to put this command in a boot script. Oh, and by the way, you exit `showkey` by not typing anything for ten seconds.

I want to add SysRQ key events to a module, how does it work?

In order to register a basic function with the table, you must first include the header `include/linux/sysrq.h`, this will define everything else you need. Next, you must create a `sysrq_key_op` struct, and populate it with A) the key handler function you will use, B) a `help_msg` string, that will print when SysRQ prints help, and C) an `action_msg` string, that will print right before your handler is called. Your handler must conform to the prototype in `'sysrq.h'`.

After the `sysrq_key_op` is created, you can call the kernel function `register_sysrq_key(int key, const struct sysrq_key_op *op_p)`; this will register the operation pointed to by `op_p` at table key 'key', if that slot in the table is blank. At module unload time, you must call the function `unregister_sysrq_key(int key, const struct sysrq_key_op *op_p)`, which will remove the key op pointed to by 'op_p' from the key 'key', if and only if it is currently registered in that slot. This is in case the slot has been overwritten since you registered it.

The Magic SysRQ system works by registering key operations against a key op lookup table, which is defined in `'drivers/tty/sysrq.c'`. This key table has a number of operations registered into it at compile time, but is mutable, and 2 functions are exported for interface to it:

`register_sysrq_key` and `unregister_sysrq_key`.

Of course, never ever leave an invalid pointer in the table. I.e., when your module that called `register_sysrq_key()` exits, it must call `unregister_sysrq_key()` to clean up the sysrq key table entry that it used. Null pointers in the table are always safe. :)

If for some reason you feel the need to call the `handle_sysrq` function from within a function called by `handle_sysrq`, you must be aware that you are in a lock (you are also in an interrupt handler, which means don't sleep!), so you must call `__handle_sysrq_nolock` instead.

When I hit a SysRq key combination only the header appears on the console?

Sysrq output is subject to the same console loglevel control as all other console output. This means that if the kernel was booted 'quiet' as is common on distro kernels the output may not appear on the actual console, even though it will appear in the dmesg buffer, and be accessible via the `dmesg` command and to the consumers of `/proc/kmsg`. As a specific exception the header line from the sysrq command is passed to all console consumers as if the current loglevel was maximum. If only the header is emitted it is almost certain that the kernel loglevel is too low. Should you require the output on the console channel then you will need to temporarily up the console loglevel using `alt-sysrq-8` or:

```
echo 8 > /proc/sysrq-trigger
```

Remember to return the loglevel to normal after triggering the sysrq command you are interested in.

I have more questions, who can I ask?

Just ask them on the linux-kernel mailing list:

linux-kernel@vger.kernel.org

Credits

- Written by Mydraal <vulpyne@vulpyne.net>
- Updated by Adam Sulmicki <adam@cfar.umd.edu>
- Updated by Jeremy M. Dolan <jmd@turbogeek.org> 2001/01/28 10:15:59
- Added to by Crutcher Dunnivant <crutcher+kernel@datastacks.com>