

# Linux Kernel Hardening with the Yocto Project

## Embedded World 2026

Michael Opdenacker

Root Commit

Mar. 10, 2026



© 2024-2026 Root Commit. Licensed under CC BY-SA 4.0.

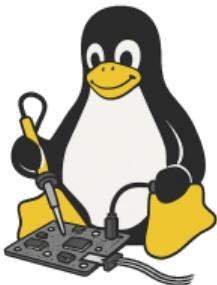
Embedded Linux consultant and trainer

- <https://rootcommit.com/about/michael-opdenacker/>
- Former founder of [Bootlin](#)
- New founder of [Root Commit](#)
- Free Software enthusiast and advocate (member of [April.org](#))



## Consulting and engineering work

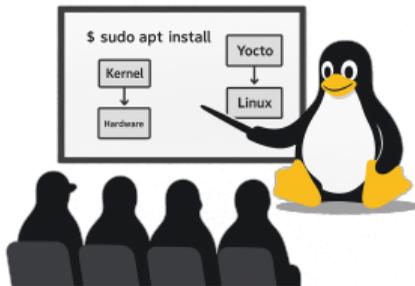
- Yocto Project  
System implementation, porting to new versions, new features
- Linux Kernel  
Driver development, board bring-up, debugging
- Embedded Linux  
Boot time, bug fixing, security and other optimizations



## Training — <https://rootcommit.com/training>

- Yocto Project and OpenEmbedded — Free Materials!
- Linux kernel, board support, driver development  
Free Materials after next course!
- Embedded Linux
- Linux Boot Time Reduction

What's special: focus on practical activities, interactivity and learning techniques. At the heart of the community!

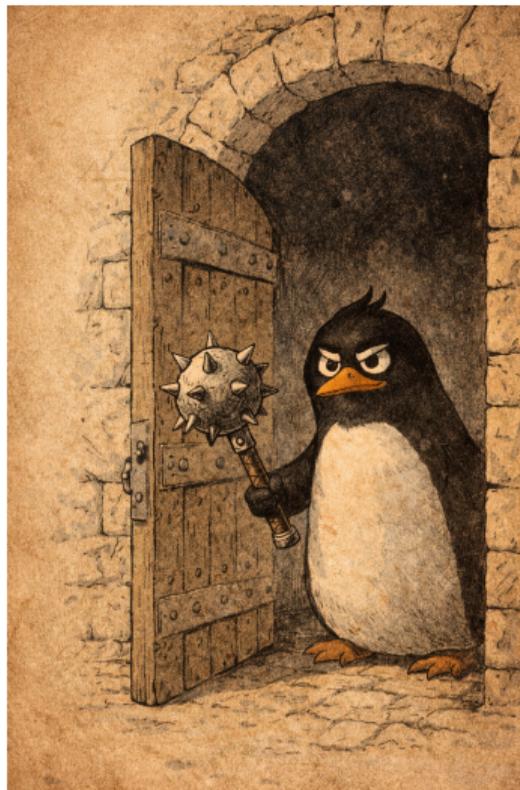


# Kernel Hardening

- The kernel is the cornerstone and stronghold of a Linux based system.
- It provides pretty good security by default, but there are developer friendly settings that are very attacker friendly too!
- If the kernel is compromised, there is almost no limit to what an attacker can do.

There are two main types of techniques to make your kernel harder to compromise:

- Reducing the attack surface
  - Only keep the features and drivers needed in your system
  - Good fit for dedicated embedded systems anyway
  - Also helps to boot faster 😊
- Using kernel hardening features
  - Features making it harder to exploit unknown or future vulnerabilities
  - Developed through the [Kernel Self-Protection](#) project.



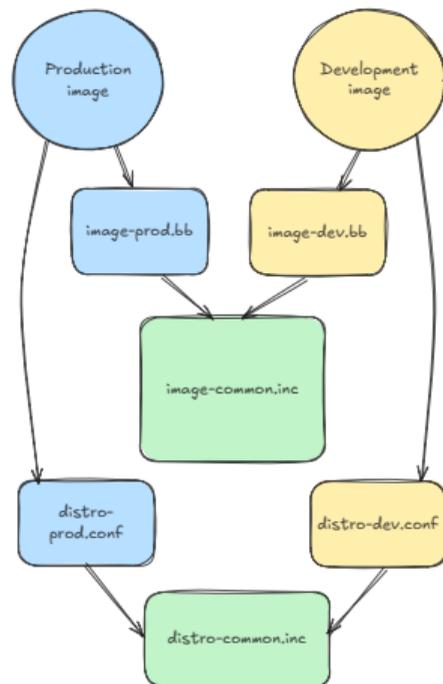
# Kernel Hardening

## Project Setup

Prepare for production early during development

- Maintain a production image alongside the development one
- Development settings isolated from the start from the production one
- Less risk of forgetting debugging features in the production image

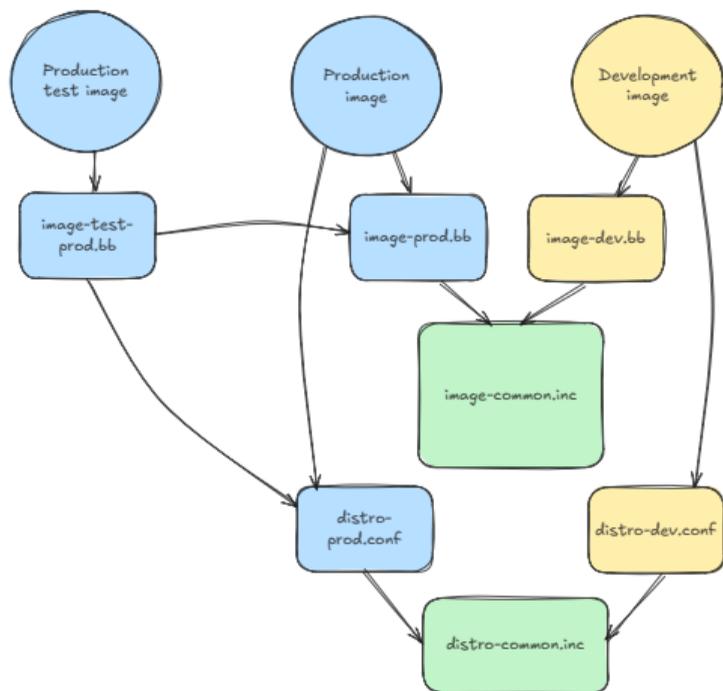
Production and development images



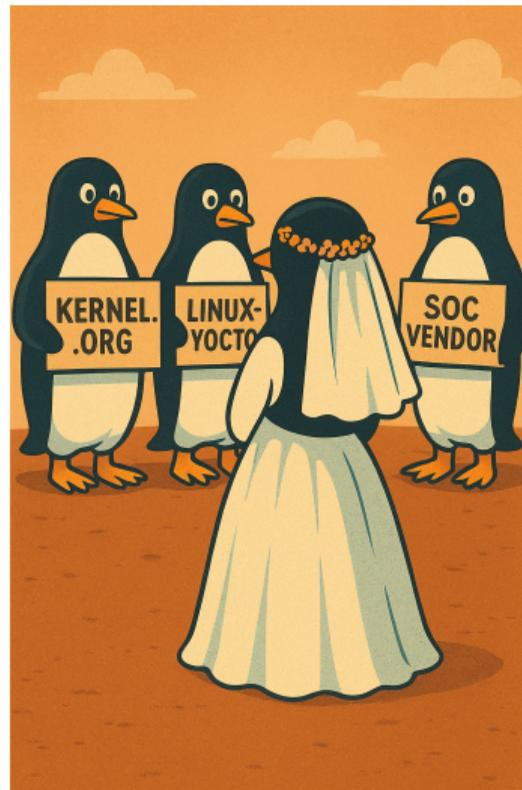
At this stage, we need some tools to test the production image

- Don't want to put them in the development image, will get false positives with development tools and features.
- Don't want to leave them in the production image, could be used to reveal weak spots. They may also pull unwanted dependencies.
- Solution: create a third "Production test" image.

Production test, production and development images



- An **LTS kernel** from [kernel.org](https://kernel.org) is the best solution to get kernel updates quickly
- **linux-yocto** is pretty well supported, if you can wait a few weeks for updates.
- **Vendor kernels** support your hardware well, but they are not meant to include vulnerability fixes in a timely fashion. They can also be very outdated.



- Add this to `conf/local.conf`:  
`INHERIT += "cve-check"`
- This will add a CVE task to the recipes you're building
- You may also want to ignore CVEs that are irrelevant to Poky and OE-core:  
`include conf/distro/include/cve-extra-exclusions.inc`
- To speed up NVD database downloads, request a unique key (`NVDCVE_API_KEY`)  
`NVDCVE_API_KEY = "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx"`
- Then, `bitbake` your regular image,  
and the checks will be run (without running the other tasks if not necessary)
- You can also run checks on specific recipes:  
`$ bitbake -c cve_check linux-yocto`

# Kernel Hardening

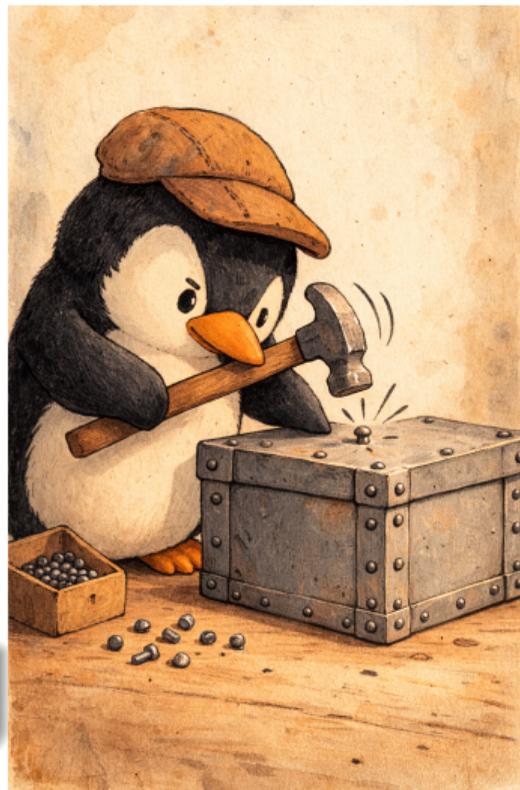
## Detect Hardening Opportunities

kernel-hardening-checker from Alexander Popov

- <https://github.com/a13xp0p0v/kernel-hardening-checker>
- Introduced only in Walnascar (meta-oe)
- Reports all the kernel configuration and command line settings which could be changed to harden the running kernel.
- My contributions:
  - Updated meta-oe to support version 0.6.10.2 (many recent changes), supporting a `-a` option and the RISC-V architecture ([link](#))
  - Backported to Scarthgap too.

```
recipes-core/images/image-prod-test.bb
```

```
IMAGE_INSTALL += "kernel-hardening-checker"
```





Try it on your PC!

- Clone the repository:

```
$ git clone https://github.com/a13xp0p0v/kernel-hardening-checker.git
```

- Run it directly from the repository:

```
$ /path/to/kernel-hardening/checker/bin/kernel-hardening-checker -a
```

Run it on the embedded target:

```
$ kernel-hardening-checker -a
```



- Kernel configuration: remove unnecessary features
  - This reduces the attack surface, and boot time too!
  - Yocto stores kernel modules in separate packages. This helps to identify unnecessary ones!
- See [security/self-protection](#) (kernel documentation) for details about kernel protection techniques.

## The Linux

### Kernel

6.17.0-rc2

#### Quick search

06

#### Contents

- Development process
- Submitting patches
- Code of conduct
- Maintainer handbook
- All development process docs
- Care API
- Driver APIs
- Subsystems
  - Core subsystems
    - Human interfaces
    - Networking interfaces
    - Storage interfaces
  - Other subsystems
    - Accounting
    - ARMv8 CPU frequency and voltage scaling code in the Linux/ARM kernel
    - ELFv2 Subsystem
    - FSMA
    - IO/Mem Subsystem
    - Intel(R) IO
    - IOATDMA
    - Intel Frequency Governor (FGT)
    - 3D Graphics Subsystem
    - Workload Support
    - Virtualization support
    - Hardware Monitoring
  - Security/Initialization
    - Crash API
    - PIE Documentation
    - LSM support
    - IO Subsystem
    - ASoC Subsystem
    - ASoC Multimedia Device
    - Documentation
    - PCI Subsystem
    - VMX Subsystem
    - TDX Subsystem
- Locking
- Licensing rules
- Writing documentation
- Development tools
- Testing guide
- Harding guide
- Tracing
- Toolbox Inspector
- Linking/arch
- Build
- Administration
  - Build system
  - Reporting Issues
  - Userspace tools
  - Userspace API
  - Firmware
  - Firmware and DeviceTree
- FSU architectures
- Unsorted documentation
- Translations

## Kernel Self-Protection

English

Kernel self-protection is the design and implementation of systems and structures within the Linux kernel to protect against security flaws in the kernel itself. This covers a wide range of issues, including removing entire classes of bugs, blocking security flaw exploitation methods, and actively detecting attack attempts. Not all topics are explored in this document, but it should serve as a reasonable starting point and answer any frequently-asked questions (frequently welcome, of course!).

In the worst-case scenario, an attacker (an unprivileged local attacker has arbitrary read and write access to the kernel's memory. In many cases, bugs being exploited will not provide this level of access, but with systems in place that defend against the worst case we'll cover the more limited cases as well.) A higher bar, and one that should still be kept in mind, is protecting the kernel against a "privileged" local attacker, since the root user has access to a vastly increased attack surface. (Obviously when they have the ability to load arbitrary kernel modules.)

The goals for successful self-protection systems would be that they are effective, on by default, require no opt-in by developers, have no performance impact, do not impede kernel debugging, and have tests. It is understood that all these goals can be met, but it is worth explicitly mentioning them, since these aspects need to be explored, dealt with, and/or accepted.

### Attack Surface Reduction

The most fundamental defense against security exploits is to reduce the areas of the kernel that can be used to redirect execution. This ranges from limiting the exposed APIs available to userspace, making in-kernel APIs hard-to-use (accessing, maintaining the areas of writable kernel memory, etc.).

### Strict kernel memory permissions

When all of kernel memory is writable, it becomes trivial for attacks to redirect execution flow. To reduce the availability of these targets the kernel needs to protect its memory with a tight set of permissions.

### Executable code and read-only data must not be writable

Any areas of the kernel with executable memory must not be writable. While this obviously includes the kernel text itself, we must consider all additional places too: kernel modules, JIT memory, etc. (There are temporary exceptions to this rule to support things like instruction alternatives, breakpoints, kprobes, etc. If these must exist in a kernel, they are implemented in a way where the memory is temporarily made writable during the update, and then returned to the original permissions.)

In support of this are `CONFIG_STRICT_KERNEL_SORT` and `CONFIG_STRICT_MODULE_SORT`, which seek to make sure that code is not writable, data is not executable, and read-only data is neither writable nor executable.

Most architectures have these options on by default and not user selectable. For some architectures like ARM that wish to have these be selectable, the architecture's booting can select `ARCH_OPTIONAL_STRICT_KERNEL_SORT` to enable a Kconfig prompt. `CONFIG_ARM_OPTIONAL_STRICT_KERNEL_SORT_DEFAULT` determines the default setting when `ARCH_OPTIONAL_STRICT_KERNEL_SORT` is enabled.

### Function pointers and sensitive variables must not be writable

vast areas of kernel memory contain function pointers that are looked up by the kernel and used to continue execution (e.g. descriptor/vector tables, file/buffer/write operation structures, etc). The number of these variables must be reduced to an absolute minimum.

Many such variables can be made read-only by setting their `const` and/or they live in the `.rodata` section instead of the data section of the kernel, gaining the protection of the kernel's strict memory permissions as described above.

For variables that are initialized once at `_init` time, these can be marked with the `__ro_after_init` attribute.

What remains are variables that are updated rarely (e.g. GDT). These will need another infrastructure similar to the temporary exceptions made to kernel code mentioned above that allow them to spend the rest of their lifetime read-only (for example, when being updated, only the CPU thread performing the update would be given unacceptably write access to the memory.)

### Segregation of kernel memory from userspace memory

The kernel must never execute userspace memory. The kernel must also never access userspace memory without explicit expectation to do so. These rules can be enforced either by support of hardware-based

# Kernel Hardening

## Implement Hardening

## Using config fragments

- Possible for recipes inheriting the `linux-yocto` class
- Convenient to manage settings separately
- But adds unpredictability in case of settings

```
my-linux.bb
```

```
SRC_URI += " \  
    file://kallsyms.cfg \  
    file://randstruct-full.cfg \  
    file://io-strict-devmem.cfg \  
"
```

```
io-strict-devmem.cfg
```

```
CONFIG_DEVMEM=y  
CONFIG_IO_STRICT_DEVMEM=y
```

## Using one complete configuration file

- Guaranteed to work as intended  
No issues because of conflicting parameters
- Can be updated easily:  
`bitbake -c menuconfig linux-yocto`  
`bitbake -c savedefconfig linux-yocto`

```
my-linux.bb
```

```
SRC_URI += "file://defconfig"
```

- First apply configuration changes
  - First remove unnecessary kernel features (time consuming!)
  - Then follow applicable recommendations of `kernel-hardening-checker`
  - Do it little by little, test at each iteration
  - Keep track of all intermediate configurations (if rewinding is needed)
- Then tune the remaining kernel command line and `sysctl` settings  
Some are no longer relevant after kernel configuration changes

# Selected Kernel Hardening Settings (1)

```
CONFIG_FORTIFY_SOURCE:
    Harden common str/mem functions against buffer overflows

Detect overflows of buffers in common string and memory functions
where the compiler can determine and validate the buffer sizes.

Symbol: FORTIFY_SOURCE [=y]
Type : bool
Defined at security/Kconfig.hardening:216
Prompt: Harden common str/mem functions against buffer overflows
Depends on: ARCH_HAS_FORTIFY_SOURCE [=y] && (!X86_32 || !CC_IS_CLANG [=n] || CLANG_VERSION [=0]>=160000)
Location:
-> Security options
-> Kernel hardening options
-> Bounds checking
-> Harden common str/mem functions against buffer overflows (FORTIFY_SOURCE [=y])

(100%)
```

```
CONFIG_KFENCE:
    KFENCE: low-overhead sampling-based memory safety error detector

KFENCE is a low-overhead sampling-based detector of heap out-of-bounds
access, use-after-free, and invalid-free errors. KFENCE is designed
to have negligible cost to permit enabling it in production
environments.

See <file:Documentation/dev-tools/kfence.rst> for more details.

Note that, KFENCE is not a substitute for explicit testing with tools
such as KASAN. KFENCE can detect a subset of bugs that KASAN can
detect, albeit at very different performance profiles. If you can
afford to use KASAN, continue using KASAN, for example in test
environments. If your kernel targets production use, and cannot
enable KASAN due to its cost, consider using KFENCE.

Symbol: KFENCE [=n]
Type : bool
Defined at lib/Kconfig.kfence:6
Prompt: KFENCE: low-overhead sampling-based memory safety error detector
Depends on: HAVE_ARCH_KFENCE [=y]
Location:
-> Kernel hacking
-> Memory Debugging
-> KFENCE: low-overhead sampling-based memory safety error detector (KFENCE [=n])

(95%)
```

```
CONFIG_ARM_DEBUG_MX:
    Warn on W+X mappings at boot

Generate a warning if any W+X mappings are found at boot.

This is useful for discovering cases where the kernel is leaving
W+X mappings after applying NX, as such mappings are a security risk.

Look for a message in dmesg output like this:

    arm/mm: Checked W+X mappings: passed, no W+X pages found.

or like this, if the check failed:

    arm/mm: Checked W+X mappings: FAILED, <N> W+X pages found.

Note that even if the check fails, your kernel is possibly
still fine, as W+X mappings are not a security hole in
themselves, what they do is that they make the exploitation
of other unfixed kernel bugs easier.

There is no runtime or memory usage effect of this option
once the kernel has booted up - it's a one time check.

If in doubt, say "Y".

(73%)
```

```
CONFIG_HARDENED_USERCOPY:
    Harden memory copies between kernel and userspace

This option checks for obviously wrong memory regions when
copying memory to/from the kernel (via copy_to_user() and
copy_from_user() functions) by rejecting memory ranges that
are larger than the specified heap object, span multiple
separately allocated pages, are not on the process stack,
or are part of the kernel text. This prevents entire classes
of heap overflow exploits and similar kernel memory exposures.

Symbol: HARDENED_USERCOPY [=y]
Type : bool
Defined at security/Kconfig.hardening:225
Prompt: Harden memory copies between kernel and userspace
Location:
-> Security options
-> Kernel hardening options
-> Bounds checking
-> Harden memory copies between kernel and userspace (HARDENED_USERCOPY [=y])
Implies: STRICT_DEVMEM [=n]

(100%)
```

# Selected Kernel Hardening Settings (2)



```
CONFIG_FTRACE:                                     Tracers
Enable the kernel tracing infrastructure.

Symbol: FTRACE [=y]
Type : bool
Defined at kernel/trace/Kconfig:200
Prompt: Tracers
Depends on: TRACING_SUPPORT [=y]
Location:
-> Kernel hacking
-> Tracers (FTRACE [=y])

< > [x] >
```

```
CONFIG_KPROBES:                                    Kprobes
Kprobes allows you to trap at almost any kernel address and
execute a callback function. register kprobe() establishes
a probe point and specifies the callback. Kprobes is useful
for kernel debugging, non-intrusive instrumentation and testing.
If in doubt, say "N".

Symbol: KPROBES [=y]
Type : bool
Defined at arch/Kconfig:117
Prompt: Kprobes
Depends on: HAVE_KPROBES [=y]
Location:
-> General architecture-dependent options
-> Kprobes (KPROBES [=y])
Selects: KALLSYMS [=y] && EXECMEM [=y] && NEED_TASKS_RCU [=y]
Selected by [n]:
- KGDB_HONOUR_BLOCKLIST [=n] && KGDB [=n] && HAVE_KPROBES [=y] && MODULES [=y]

< > [x] >
```

```
CONFIG_KALLSYMS:                                    Load all symbols for debugging/ksymbols
Say Y here to let the kernel print out symbolic crash information and
symbolic stack backtraces. This increases the size of the kernel
somewhat, as all symbols have to be loaded into the kernel image.

Symbol: KALLSYMS [=y]
Type : bool
Defined at init/Kconfig:1986
Prompt: Load all symbols for debugging/ksymbols
Visible if: EXPERT [=y]
Location:
-> General setup
-> Configure standard kernel features (expert users) (EXPERT [=y])
-> Load all symbols for debugging/ksymbols (KALLSYMS [=y])
Selected by [y]:
- KPROBES [=y] && HAVE_KPROBES [=y]
Selected by [n]:
- LATENCYTOP [=n] && DEBUG_KERNEL [=y] && STACKTRACE_SUPPORT [=y] && PROC_FS [=y] && FRAME_POINTER [=n] || \
MIPS || PPC || S390 || MICROBLAZE || ARM [=y] || ARC || X86
- DEBUG_KMEMLEAK [=n] && DEBUG_KERNEL [=y] && HAVE_DEBUG_KMEMLEAK [=y]
- CODE_TAGGING [=n]
- LOCKDEP [=n] && DEBUG_KERNEL [=y] && LOCK_DEBUGGING_SUPPORT [=y]
- FUNCTION_TRACER [=n] && FTRACE [=y] && HAVE_FUNCTION_TRACER [=y]
- STACK_TRACER [=n] && FTRACE [=y] && HAVE_FUNCTION_TRACER [=y]

< > [x] >
```

```
CONFIG_KEXEC:                                       Enable kexec system call
kexec is a system call that implements the ability to shutdown your
current kernel, and to start another kernel. It is like a reboot
but it is independent of the system firmware. And like a reboot
you can start any kernel with it, not just Linux.

The name comes from the similarity to the exec system call.

It is an ongoing process to be certain the hardware in a machine
is properly shutdown, so do not be surprised if this code does not
initially work for you. As of this writing the exact hardware
interface is strongly in flux, so no good recommendation can be
made.

Symbol: KEXEC [=y]
Type : bool
Defined at kernel/Kconfig.kexec:20
Prompt: Enable kexec system call
Depends on: ARCH_SUPPORTS_KEXEC [=y]
Location:
-> General setup
-> Kexec and crash features
-> Enable kexec system call (KEXEC [=y])
Selects: KEXEC_CORE [=y]

< > [x] >
```



# Selected Kernel Hardening Settings (3)

**Poison kernel stack before returning from syscalls**

CONFIG\_KSTACK\_ERASE:

This option makes the kernel erase the kernel stack before returning from system calls. This has the effect of leaving the stack initialized to the poison value, which both reduces the lifetime of any sensitive stack contents and reduces potential for uninitialized stack variable exploits or information exposures (it does not cover functions reaching the same stack depth as prior functions during the same syscall). This blocks most uninitialized stack variable attacks, with the performance impact being driven by the depth of the stack usage, rather than the function calling complexity.

The performance impact on a single CPU system kernel compilation sees a 1% slowdown, other systems and workloads may vary and you are advised to test this feature on your expected workload before deploying it.

Symbol: KSTACK\_ERASE [=n]  
Type : bool  
Defined at security/Kconfig.hardening:88  
Prompt: Poison kernel stack before returning from syscalls  
Depends on: HAVE\_ARCH\_KSTACK\_ERASE [=y] && (GCC\_PLUGINS [=y] || CC\_HAS\_SANCOV\_STACK\_DEPTH\_CALLBACK [=n])  
Location:  
-> Security options

( 87%)



**Disable heap randomization**

CONFIG\_COMPAT\_BRK:

Randonizing heap placement makes heap exploits harder, but it also breaks ancient binaries (including anything libc5 based). This option changes the bootup default to heap randomization disabled, and can be overridden at runtime by setting /proc/sys/kernel/randomize\_va\_space to 2.

On non-ancient distros (post-2000 ones) N is usually a safe choice.

Symbol: COMPAT\_BRK [=y]  
Type : bool  
Defined at mm/Kconfig:303  
Prompt: Disable heap randomization  
Location:  
-> Memory Management options  
-> Disable heap randomization (COMPAT\_BRK [=y])

(100%)



**Enable core dump support**

CONFIG\_COREDUMP:

This option enables support for performing core dumps. You almost certainly want to say Y here. Not necessary on systems that never need debugging or only ever run flawless code.

Symbol: COREDUMP [-y]  
Type : bool  
Defined at fs/Kconfig.binfmt:171  
Prompt: Enable core dump support  
Visible if: EXPERT [=y]  
Location:  
-> Executable file formats  
-> Enable core dump support (COREDUMP [=y])

(100%)



**Allow writing to mounted block devices**

CONFIG\_BLK\_DEV\_WRITE\_MOUNTED:

When a block device is mounted, writing to its buffer cache is very likely going to cause filesystem corruption. It is also rather easy to crash the kernel in this way since the filesystem has no practical way of detecting these writes to buffer cache and verifying its metadata integrity. However there are some setups that need this capability like running fsck on read-only mounted root device, modifying some features on mounted ext4 filesystem, and similar. If you say N, the kernel will prevent processes from writing to block devices that are mounted by filesystems which provides some more protection from runaway privileged processes and generally makes it much harder to crash filesystem drivers. Note however that this does not prevent underlying device(s) from being modified by other means, e.g. by directly submitting SCSI commands or through access to lower layers of storage stack. If in doubt, say Y. The configuration can be overridden with the bdev\_allow\_write\_mounted boot option.

Symbol: BLK\_DEV\_WRITE\_MOUNTED [=y]  
Type : bool  
Defined at block/Kconfig:77  
Prompt: Allow writing to mounted block devices  
Depends on: BLOCK [=y]  
Location:  
-> Enable the block layer (BLOCK [=y])

( 94%)



# Selected Kernel Hardening Settings (4)

```
CONFIG_MODULE_SIG_ALL: Automatically sign all modules

Sign all modules during make modules_install. Without this option,
modules must be signed manually, using the scripts/sign-file tool.

Symbol: MODULE_SIG_ALL [=y]
Type : bool
Defined at kernel/module/Kconfig:280
Prompt: Automatically sign all modules
Depends on: MODULES [=y] && (MODULE_SIG [=y] || IMA_APPRAISE_MODSIG [=n])
Location:
-> Enable loadable module support (MODULES [=y])
-> Module signature verification (MODULE_SIG [=y])
-> Automatically sign all modules (MODULE_SIG_ALL [=y])

< [xit] > (100%)
```

```
CONFIG_MODULE_SIG_FORCE: Require modules to be validly signed

Reject unsigned modules or signed modules for which we don't have a
key. Without this, such modules will simply taint the kernel.

Symbol: MODULE_SIG_FORCE [=y]
Type : bool
Defined at kernel/module/Kconfig:273
Prompt: Require modules to be validly signed
Depends on: MODULES [=y] && MODULE_SIG [=y]
Location:
-> Enable loadable module support (MODULES [=y])
-> Module signature verification (MODULE_SIG [=y])
-> Require modules to be validly signed (MODULE_SIG_FORCE [=y])

< [xit] > (100%)
```

```
CONFIG_STRICT_DEVMEM: Filter access to /dev/mem

If this option is disabled, you allow userspace (root) access to all
of memory, including kernel and userspace memory. Accidental
access to this is obviously disastrous, but specific access can
be used by people debugging the kernel. Note that with PAT support
enabled, even in this case there are restrictions on /dev/mem
use due to the cache aliasing requirements.

If this option is switched on, and IO_STRICT_DEVMEM=n, the /dev/mem
file only allows userspace access to PCI space and the BIOS code and
data regions. This is sufficient for dosemu and X and all common
users of /dev/mem.

If in doubt, say Y.

Symbol: STRICT_DEVMEM [=y]
Type : bool
Defined at lib/Kconfig.debug:1981
Prompt: Filter access to /dev/mem
Depends on: MMU [=y] && DEVMEM [=y] && (ARCH_HAS_DEVMEM_IS_ALLOWED [=n] || GENERIC_LIB_DEVMEM_IS_ALLOWED [=y])
Location:
-> Kernel hacking
-> Filter access to /dev/mem (STRICT_DEVMEM [=y])
Implied by [y]:

< [xit] > ( 97%)
```

```
CONFIG_IO_STRICT_DEVMEM: Filter I/O access to /dev/mem

If this option is disabled, you allow userspace (root) access to all
io-memory regardless of whether a driver is actively using that
range. Accidental access to this is obviously disastrous, but
specific access can be used by people debugging kernel drivers.

If this option is switched on, the /dev/mem file only allows
userspace access to 'idle' io-memory ranges (see /proc/iomem) This
may break traditional users of /dev/mem (dosemu, legacy X, etc...)
if the driver using a given range cannot be disabled.

If in doubt, say Y.

Symbol: IO_STRICT_DEVMEM [=y]
Type : bool
Defined at lib/Kconfig.debug:1921
Prompt: Filter I/O access to /dev/mem
Depends on: STRICT_DEVMEM [=y]
Location:
-> Kernel hacking
-> Filter access to /dev/mem (STRICT_DEVMEM [=y])
-> Filter I/O access to /dev/mem (IO_STRICT_DEVMEM [=y])

< [xit] > (100%)
```

More such settings: <https://gitlab.com/rootcommit/yocto-kernel-hardening/-/tree/main/graphics/more-kconfig>

- Module signing is automatic with `CONFIG_MODULE_SIG_ALL`.  
Also turn on `CONFIG_MODULE_SIG_FORCE`.
- The signing key is generated on the fly and the public key is built into the kernel.
- Not necessary to keep the private key.
- Still possible to use your own keys. Yocto can help you:  
<https://ejaaskel.dev/yocto-hardening-kernel-module-signing/>

See kernel documentation for details: [admin-guide/module-signing](#)

Not completely trivial!

- No universal way to change the default command line (check the bootloader recipes).
- However, you can always add parameters to the command line through the kernel *bootconfig* mechanism.

See kernel documentation:  
[admin-guide/bootconfig](#)

```
linux-stable_%.bbappend
```

```
FILESEXTRAPATHS:prepend := "${THISDIR}/${PN}:"  
KBUILD_DEFCONFIG = ""  
  
do_configure:prepend() {  
    install -m 0644 ${UNPACKDIR}/cmdline.bootconfig ${S}/  
}  
  
SRC_URI += " \  
    file://defconfig \  
    file://cmdline.bootconfig \  
    "
```

```
defconfig
```

```
...  
CONFIG_BOOT_CONFIG=y  
CONFIG_BOOT_CONFIG_EMBED=y  
CONFIG_BOOT_CONFIG_EMBED_FILE="cmdline.bootconfig"
```

```
cmdline.bootconfig
```

```
kernel {  
    page_alloc.shuffle=1  
    hash_pointers=always  
    nosmt  
}
```

```
recipes-core/sysctl/sysctl-config.bb
```

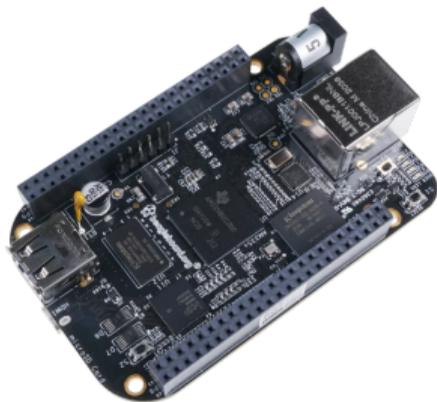
```
SUMMARY = "Custom sysctl configuration for kernel hardening"  
LICENSE = "MIT"  
LIC_FILES_CHKSUM = "file://${COMMON_LICENSE_DIR}/MIT;md5=0835ade698e0bcf8506ecda..."  
  
SRC_URI = "file://99-kernel-hardening.conf"  
  
do_install() {  
    install -Dm 0644 ${UNPACKDIR}/99-kernel-hardening.conf \  
        ${D}${sysconfdir}/sysctl.d/99-kernel-hardening.conf  
}  
  
FILES:${PN} += "${sysconfdir}/sysctl.d/99-kernel-hardening.conf"
```

```
files/99-kernel-hardening.conf
```

```
# Kernel hardening settings  
kernel.oops_limit = 100  
kernel.warn_limit = 100  
kernel.perf_event_paranoid = 3  
kernel.kptr_restrict = 2  
user.max_user_namespaces = 0  
net.ipv4.conf.default.accept_redirects = 0  
net.ipv6.conf.all.accept_redirects = 0  
net.ipv6.conf.default.accept_redirects = 0  
net.ipv6.conf.all.accept_ra = 0  
net.ipv6.conf.default.accept_ra = 0  
fs.protected_fifos = 2  
fs.protected_regular = 2  
kernel.yama.ptrace_scope = 3
```

- Many default kernel options are developer friendly, not security friendly.
- You don't want many such options in a production embedded system.
- Some hardening features are meant to find issues in kernel code, but have a significant performance cost.  
You may remove them in production systems after running robustness tests.
- Not all recommended security features are relevant:
  - Some of them may not be available on your hardware
  - You may not need them in your specific system.

- The only hardware board officially supported by Yocto (meta-yocto-bsp layer)
- Before hardening:
  - Compressed kernel size: 5.5 MB
  - Number of OK: 159
  - Number of FAIL: 125
- After hardening:
  - Compressed kernel size: 4.7 MB  
Note: minimal effort to reduce kernel size
  - Number of OK: 278
  - Number of FAIL: 6 🎉



- The only hardware board officially supported by Yocto (meta-yocto-bsp layer)
- Before hardening:
  - Compressed kernel size: 5.5 MB
  - Number of OK: 159
  - Number of FAIL: 125
- After hardening:
  - Compressed kernel size: 4.7 MB  
Note: minimal effort to reduce kernel size
  - Number of OK: 278
  - Number of FAIL: 6 🎉  
which are actually OK for me 😊



Can't select these on my platform

option_name	type	reason	decision	desired_val	check_result
CONFIG_STACKPROTECTOR_PER_TASK	kconfig	self_protection	defconfig	y	FAIL: is not found
CONFIG_UBSAN_TRAP	kconfig	self_protection	kspp	y	FAIL: CONFIG_UBSAN_ENUM is not "is not set"
CONFIG_ARM_SMMU	kconfig	self_protection	a13xp0p0v	y	FAIL: is not found
CONFIG_ARM_SMMU DISABLE BYPASS BY DEFAULT	kconfig	self_protection	a13xp0p0v	y	FAIL: is not found
CONFIG_MODULES	kconfig	cut_attack_surface	kspp	is not set	FAIL: "y"
kernel.modules_disabled	sysctl	cut_attack_surface	kspp	1	FAIL: "0"

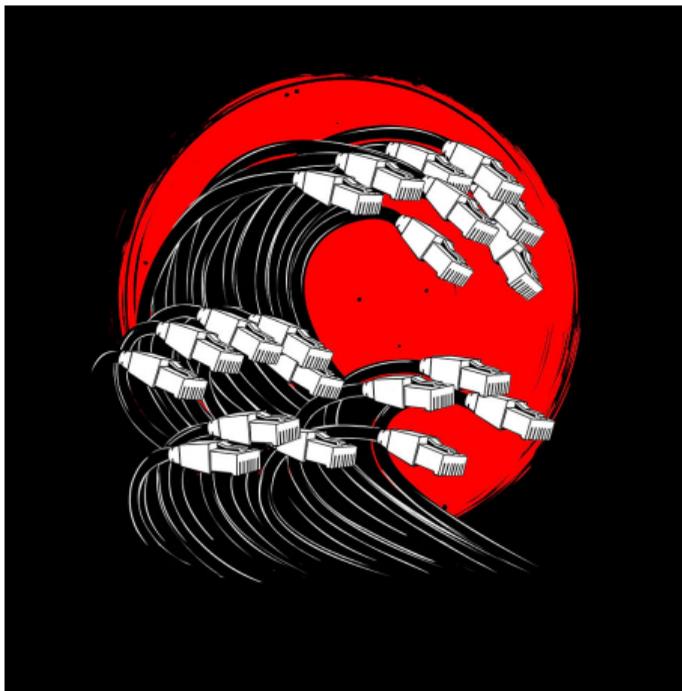
[+] Config check is finished: 'OK' - 278 (suppressed in output) / 'FAIL' - 6

Want to keep modules. Safer with module signing.

# Wrapping Up

- `kernel-hardening-checker` from Alex Popov is your best friend.  
Find it in the `meta-oe` layer (meta-openembedded repository)
- Default kernel configurations are developer friendly, not secure enough
- There is lot of functionality you don't need in an embedded system  
This helps to reduce the attack surface
- Pretty easy to increase security without sacrificing functionality and performance (not substantially)
- Don't apply recommendations blindly
- This protects against future vulnerabilities.  
But don't forget to fix known ones!

- Michael Opendacker — Securing Yocto Built Systems  
<https://rootcommit.com/pub/conferences/2025/elce/yocto-security/yocto-security.pdf>
- Root Commit Blog — Yocto Security: Kernel Hardening  
<https://rootcommit.com/2025/yocto-security-kernel-hardening/>
- Root Commit's Yocto Project training course  
<https://rootcommit.com/training/yocto>
- Esa Jääskelä — Yocto Hardening  
Series of blog posts about many aspects of the topic  
Fun and exhaustive!  
<https://ejaaskel.dev/yocto-hardening/>



Darknet Diaries podcast shop:  
<https://shop.darknetdiaries.com>



## Booth 4-648

- Have Yocto goodies: coffee ☕, cookies 🍪
- Our best swag: our smiles 😊
- Come and see how welcoming the Yocto community is.

## Questions? Comments?

- [mo@rootcommit.com](mailto:mo@rootcommit.com)
-  <https://fosstodon.org/@MichaelOpdenacker>
- XMPP: [omichael@conversations.im](mailto:omichael@conversations.im)
- Signal: [rootcommit.01](https://rootcommit.01)
- Slides available under the CC-BY-SA 4.0 license  
<https://rootcommit.com/pub/conferences/2026/embedded-world/yocto-kernel-hardening/>
- Sources (L<sup>A</sup>T<sub>E</sub>X):  
<https://gitlab.com/rootcommit/yocto-kernel-hardening/>
- Code used in this demo:  
<https://gitlab.com/rootcommit/meta-kernel-hardening-demo>

