

fedora<sup>f</sup>

10  
YEARS



# Introduction to Linux Kernel Development

*Your first patch*

**Levente Kurusa**

Linux kernel hacker & enthusiast

Fedora Project Hungary

<levex@linux.com>



# Topics

fedora<sup>f</sup>

10  
YEARS

# 1 General introduction

# 2 Building

# 3 Kernel API

# 4 Patching

# 5 Sending it off



fedora<sup>f</sup>

10  
YEARS

# General Introduction

fedora<sup>f</sup><sup>TM</sup>



# What is a kernel?

fedora<sup>f</sup>

10  
YEARS

The operating system's core

Responsibilities:

- # 1 Initialise the system for the user
- # 2 Protect the data
- # 3 Introduce abstractions to help the development of new software





# Kernel types

fedora<sup>f</sup>

10  
YEARS

- # 1 Microkernel
- # 2 Monolithic kernel
- # 3 Hybrid kernel

Some exotics:

- # 1 Exokernel
- # 2 Cache kernel



# Linux specifics

fedora<sup>f</sup>

10  
YEARS

# 1 Monolithic kernel

# 2 Modular

# 3 Error-happens-so-panic type





fedora<sup>f</sup>

10  
YEARS

# The source of panic()

" I remarked to Dennis that easily half the code I was writing in Multics was error recovery code. He said, "We left all that stuff out. If there's an error, we have this routine called panic, and when it is called, the machine crashes, and you holler down the hall, 'Hey, reboot it.'" "

-- Tom van Vleck to Dennis Ritchie

fedora<sup>f</sup>

10  
YEARS

# Building

fedora<sup>f</sup><sup>TM</sup>





# Kconfig

fedora<sup>f</sup>  
10  
YEARS

Allows for specific code to be excluded via optional preprocessor macros.

Why do we need this?

# 1 To allow different configurations!

# 2 Headers are lengthy!

*(and incomprehensible anyways)*



# Kconfig symbols

fedora<sup>f</sup>  
10  
YEARS

Specific symbols that will be converted to CONFIG\_\* preprocessor defines.

Types:

- # 1 bool (true or false) [y, n]
- # 2 tristate (built-in, module, false) [y, m, n]

They have default values





# Kconfig example

fedora<sup>f</sup>

10  
YEARS

```
config ATA_ACPI
```

```
bool "ATA ACPI Support"
```

```
depends on ACPI && PCI
```

```
default y
```

```
help
```

This option adds support for ATA-related ACPI objects. These ACPI objects add the ability to retrieve taskfiles from the ACPI BIOS and write them to the disk controller.

These objects may be related to performance, security, power management, or other areas.



# Makefiles are still in use!

fedora<sup>f</sup>

10  
YEARS

We still use makefiles to define the relations between the source files.

GNU Make gives us a pretty cool interface!

We don't reinvent the wheel!





# How to configure?

fedora<sup>f</sup>  
10  
YEARS

Configuration is easy thanks to GNU make!

- # 1 'make defconfig' → Create a default configuration for \$ARCH
- # 2 'make menuconfig' → Shows a cool ncurses menu, which allows pretty configuration
- # 3 'make randconfig' → Randomly configure the kernel
- # 4 'make xconfig' → For all the GUI fans out there!



# Building

fedora<sup>f</sup>  
10  
YEARS

Configuration completed, let's build.

`'make -j`nproc`'`

*(Looks way cooler than 'make -j4')*





# Post-build

fedora<sup>f</sup>  
10  
YEARS

To actually boot, you will need an initrd or initramfs.

The kernel file that needs to be booted will be at:  
`./$ARCH/boot/bzImage`

Or if you need the vmlinux file it will be in the root.

fedora<sup>f</sup>

10  
YEARS

# Introduction to Kernel API

fedora<sup>f</sup><sup>TM</sup>





fedora<sup>f</sup>

10  
YEARS

# No C library, no developer

Most developers flee when they see they can't use `printf()`

But wait, there is a replacement for that! `printfk()`

Actually, the relevant parts from the C library are there!

This means: `sscanf()`, `[vscn]printf()`, `kstrtoX()` etc...



fedora<sup>f</sup>

10  
YEARS

# Thread-safety concerns

Since the kernel supports Symmetric MultiProcessing, we need ways to achieve mutual exclusion.

This is what we have in store:

- # 1 mutex (with global/local IRQ disable, full flags save)
- # 2 semaphores!
- # 3 spinlock\_t





# A cool module system

Your device will most likely not be used by all users, so modules are a great way.

Modules are built-up like this:

# 1 `module_init(init function ptr);`

Shows which function will be called when the module is passed to `insmod`

# 2 `module_exit(exit function ptr);`

This function will be called when your module is `rmmod'd`.



# PCI drivers

fedora<sup>f</sup>  
10  
YEARS

PCI drivers are built up in the exact same way, but they also have a table consisting of `pci_device_id`'s.

To register your PCI driver, you would call:  
`pci_register_driver(ptr-to-pci_operations-struct);`

This together with a `probe()` function will suffice.





# kobjects

fedora<sup>f</sup>

10  
YEARS

Kobjects are primitives that are parts of a garbage collection system.

Kobjects are usually 'inlined' in a container struct which handles them.

They have a few helper functions i.e.  
`kobj_{put,get}(struct kobject *kobj);`

They have types! (`struct kobj_type`)

fedora<sup>f</sup>

10  
YEARS

# Patching

fedora<sup>f</sup><sup>TM</sup>





# Contributing to the Kernel

fedora<sup>f</sup>

10  
YEARS

- # 1 Reviewing patches for common errors
- # 2 Building randconfigs
- # 3 Adding support for new device
- # 4 Fixing a bug
- # 5 Doing staging work

This is what we will do!



# checkpatch.pl

fedora<sup>f</sup>  
10  
YEARS

Excellent utility to find code-style errors!

Sometimes it also finds bugs!

All patches have to be checkpatch-safe before being applied!





# Code style?

fedora<sup>f</sup>  
10  
YEARS

# 1 Code style contributes to the consistency of the kernel

# 2 Helps to spot bugs

# 3 Don't be afraid, this isn't that terribly bad code style!



# Developer Certificate of Origin

fedora<sup>f</sup>  
10  
YEARS

- # 1 I created this change and I have the right to post it; or
- # 2 The change is based on a previous change under the same license.
- # 3 The change was sent to me by someone who has certified any of these options.



fedora<sup>f</sup>

10  
YEARS

# Sending it off

fedora<sup>f</sup><sup>TM</sup>



fedora<sup>f</sup>  
10  
YEARS

# Thanks for your attention!

@ Levente Kurusa  
<levex@linux.com>

Developer Conference 2014, Brno, CZ.

<http://devconf.cz/f/105>