



Embedded Linux from scratch in 45 minutes


Michael Opdenacker

michael.opdenacker@bootlin.com

© Copyright 2004-2021, Bootlin.


Creative Commons BY-SA 3.0 license.

Corrections, suggestions, contributions and translations are welcome!





Welcome to the special edition of FOSDEM for Covid




<conspiracy>Note that this is the FOSDEM logo since 2014. Weird, isn't it?</conspiracy>

Image credits: https://commons.wikimedia.org/wiki/File:FOSDEM_logo.svg



- ▶ Founder and Embedded Linux engineer at Bootlin:
 - ▶ Embedded Linux **expertise**
 - ▶ **Development**, consulting and training
 - ▶ Focusing **only on Free and Open Source Software**
- ▶ Free Software contributor:
 - ▶ Current maintainer of the Elixir Cross Referencer, making it easier to study the sources of big C projects like the Linux kernel. See <https://elixir.bootlin.com>
 - ▶ Co-author of Bootlin's freely available embedded Linux and kernel training materials (<https://bootlin.com/docs/>)
 - ▶ Former maintainer of **GNU Typist**





Introduction



What I like in embedded Linux

- ▶ Linux is perfect for operating devices with a fixed set of features.
Unlike on the desktop, Linux is almost in every existing embedded system.
- ▶ Embedded Linux makes Linux easy to learn: just a few programs and libraries are sufficient. **You can understand the usefulness of each file in your filesystem.**
- ▶ The Linux kernel is standalone: no complex dependencies against external software. The code is in C!
- ▶ Linux works with just a few MB of RAM and storage
- ▶ There's a new version of Linux every 2-3 months.
- ▶ Relatively small development community. You end up meeting lots of familiar faces at technical conferences (like the Embedded Linux Conference).
- ▶ Lots of opportunities (and funding available) for becoming a contributor (Linux kernel, bootloader, build systems...).



Reviving an old presentation

- ▶ First shown in 2005 at the Libre Software Meeting in Dijon, France.
- ▶ Showing a 2.6 Linux kernel booting on a QEMU emulated ARM board.
- ▶ One of our most downloaded presentations at that time.

Embedded Linux From Scratch

Embedded Linux From Scratch

in 40 minutes!

Michael Opdenacker

Free Electrons

<http://free-electrons.com/>

nada + 40 min =



Created with OpenOffice.org 2.x



Embedded Linux From Scratch ... in 40 minutes!

© Copyright 2005-2008, Free Electrons

Creative Commons Attribution-ShareAlike 3.0 license

<http://free-electrons.com>

Sep 15, 2009



1



Things that changed since 2005

In the Linux kernel:

- ▶ Linux 2.6.x → 5.x
- ▶ *Bitkeeper* → *git*
- ▶ Linux is now everywhere, no need to convince customers to use it. It's even easier and easier to convince them to fund contributions to the official version.
- ▶ *devtmpfs*: automatically creates device files
- ▶ ARM and other architectures: devices described by the *Device Tree* instead of C code

And many more!

In the embedded environment:

- ▶ The Maker movement
- ▶ Cheap development boards
500+ EUR → 50-100 EUR
- ▶ The rise of Open Hardware
(Arduino, Beaglebone Black...)
- ▶ *RISC-V*: a new open-source hardware instruction set architecture



RISC-V: a new open-source ISA



- ▶ ISA: *Instruction Set Architecture*
- ▶ Created by the University of California Berkeley, in a world dominated by proprietary ISAs with heavy royalties (ARM, x86)
- ▶ Exists in 32, 64 and 128 bit variants, from microcontrollers to powerful server hardware.
- ▶ Anyone can use and extend it to create their own SoCs and CPUs.
- ▶ This reduces costs and promotes reuse and collaboration
- ▶ Implementations can be proprietary. Many hardware vendors are using RISC-V CPUs in their hardware (examples: Microchip, Western Digital, Nvidia)
- ▶ Free implementations are being created

See <https://en.wikipedia.org/wiki/RISC-V>




How to use RISC-V with Linux?

Hardware is now getting available

- ▶ Last minute news: BeagleV. The first affordable RISC-V board for the community, should be available at 150 USD in April:
<http://beagleboard.org/beaglev>.
- ▶ Icicle kit: with Microchip's PolarFire SoC and an FPGA with 254 K gates. Sold at 499 USD at CrowdSupply: <https://frama.link/dK1oanrd>
- ▶ Boards with the Kendryte K210 SoC. Sipeed MAix BiT only costs 13 USD at Seed Studio: <https://frama.link/QhBdPjsm>. Supported by Linux 5.8 but very limited, as its MMU is not supported by Linux.
- ▶ You can also synthesize RISC-V cores on FPGAs
- ▶ Before more hardware is available in 2021, you can get started with the QEMU emulator, which simulates a virtual board with *virtio* hardware

Already try it with JSLinux: <https://bellard.org/jslinux/>



BeagleV



Microchip PolarFire
SoC Icicle kit



Seed Studio Sipeed
MAix BiT



Goals

Show you the most important aspects of embedded Linux development work

- ▶ Building a cross-compiling toolchain
- ▶ Creating a disk image
- ▶ Booting a using a bootloader
- ▶ Loading and starting the Linux kernel
- ▶ Building a root filesystem populated with basic utilities
- ▶ Configuring the way the system starts
- ▶ Setting up networking and controlling the system via a web interface



Things to build today

- ▶ Cross-compiling toolchain: *Buildroot 2020.11.1*
- ▶ Firmware / first stage bootloader: *OpenSBI*
- ▶ Bootloader: *U-Boot 2021.01*
- ▶ Kernel: *Linux 5.11-rc3*
- ▶ Root filesystem and application: *BusyBox 1.33.0*

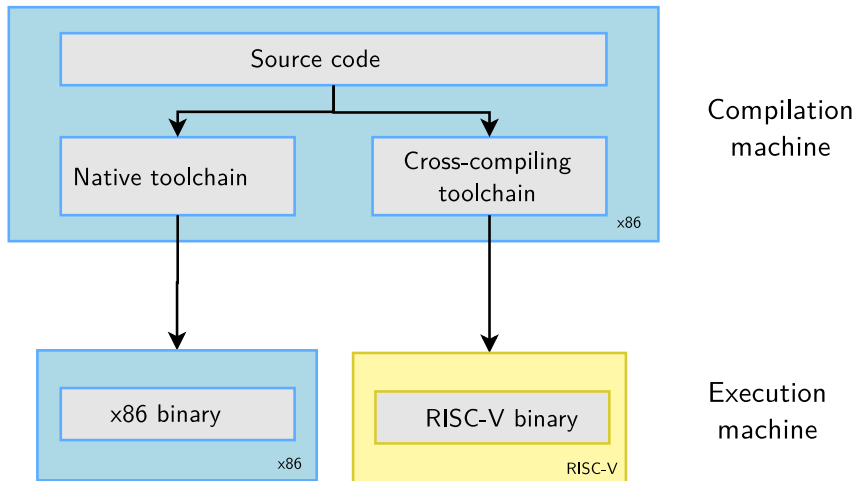
That's possible to compile and assemble in less than 45 minutes!



Cross-compiling toolchain



What's a cross-compiling toolchain?





Why generate your own cross-compiling toolchain?

Compared to ready-made toolchains:

- ▶ You can choose your compiler version
- ▶ You can choose your C library (glibc, uClibc, musl)
- ▶ You can tweak other features
- ▶ You gain reproducibility: if a bug is found, just apply a fix.
Don't need to get another toolchain (different bugs)




-
- The diagram illustrates the μClibc architecture. It features a large green circle labeled **μClibc** in the center. Inside this circle is a red circle labeled **System Call Interface**. Within the red circle, a flowchart shows the sequence of system calls: **Open the pattern** → **Calculate parameters** → **API calling** → **Network device** → **ICD interface** → **Open the pattern**. A red arrow at the bottom of the red circle indicates a cycle of **(about 280 system calls)**. To the right of the μClibc circle, there are two blue rectangular areas. The top one is labeled **Application (POSIX-compatible)** and includes examples like **BusyBox et al.**. The bottom one is labeled **Linux-specific Application**. Arrows indicate the flow of data: **data** flows from the Application to μClibc; **system calls** flow from μClibc to the Application; **system error** flows from μClibc to the Application; and **data** flows from the Application back to μClibc.

(<http://bit.ly/2zrGve2>)



Generating a RISC-V musl toolchain with Buildroot

- ▶ Download Buildroot 2020.11.1 from <https://buildroot.org>
- ▶ Extract the sources (`tar xf`)
- ▶ Run `make menuconfig`
- ▶ In Target options → Target Architecture, choose RISC-V
- ▶ In Toolchain → C library, choose musl.
- ▶ Save your configuration and run:
`make sdk`
- ▶ At the end, you have an toolchain archive in
`output/images/riscv64-buildroot-linux-musl_sdk-buildroot.tar.gz`
- ▶ Extract the archive in a suitable directory, and in the extracted directory, run: `./relocate-sdk.sh`



<https://asciinema.org/a/383836>



Testing the toolchain

- ▶ Create a new `riscv64-env.sh` file you can source to set environment variables for your project:

```
export PATH=$HOME/toolchain/riscv64-buildroot-linux-musl_sdk-buildroot/bin:$PATH
```

- ▶ Run `source riscv64-env.sh`, take a `hello.c` file and test your new compiler:

```
$ riscv64-linux-gcc -static -o hello hello.c
$ file hello
hello: ELF 64-bit LSB executable, UCB RISC-V, version 1 (SYSV), statically linked, not stripped
```

We are compiling statically so far to avoid having to deal with shared libraries.

- ▶ Test your executable with QEMU in user mode:

```
$ qemu-riscv64 hello
Hello world!
```



Hardware emulator



Finding which machines are emulated by QEMU

Tests made with QEMU 4.2.1 (Ubuntu 20.04)

```
sudo apt install qemu-system-misc
$ qemu-system-riscv64 -M ?
Supported machines are:
none                empty machine
sifive_e            RISC-V Board compatible with SiFive E SDK
sifive_u            RISC-V Board compatible with SiFive U SDK
spike               RISC-V Spike Board (default)
spike_v1.10         RISC-V Spike Board (Privileged ISA v1.10)
spike_v1.9.1        RISC-V Spike Board (Privileged ISA v1.9.1)
virt                RISC-V VirtIO board
```

We are going to use the `virt` one, emulating VirtIO peripherals (more efficient than emulating real hardware).



Booting process and privileges



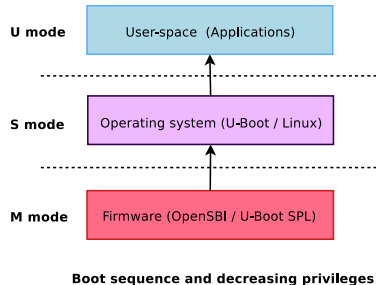
RISC-V privilege modes

RISC-V has three privilege modes:

- ▶ **U**ser (U-Mode): applications
- ▶ **S**upervisor (S-Mode): OS kernel
- ▶ **M**achine (M-Mode): bootloader and firmware

Here are typical combinations:

- ▶ **M**: simple embedded systems
- ▶ **M, U**: embedded systems with memory protection
- ▶ **M, S, U**: Unix-style operating systems with virtual memory





U-Boot bootloader



Environment for U-Boot cross-compiling

- ▶ Download U-Boot 2021.01 sources
- ▶ Let's add an environment variable to our `riscv64-env.sh` file for cross-compiling:

```
export CROSS_COMPILE=riscv64-linux-
```

- ▶ `CROSS_COMPILE` is the cross-compiler prefix, as our cross-compiler is `riscv64-linux-gcc`.



Cross-compiling U-Boot

- ▶ Find U-Boot ready-made configurations for RISC-V:

```
ls configs | grep riscv
```

- ▶ We will choose the configuration for QEMU and U-Boot running in S Mode:

```
make qemu-riscv64_smode_defconfig
```

- ▶ Now let's compile U-Boot (-j8: 8 jobs compile jobs in parallel)

```
make -j8
```




Firmware



OpenSBI: Open Supervisor Binary Interface

- ▶ Required to start an OS (S mode) from the Supervisor/Firmware (M mode)

```
git clone https://github.com/riscv/opensbi.git
cd opensbi
git checkout v0.8
make PLATFORM=generic FW_PAYLOAD_PATH=../u-boot-2021.01/u-boot.bin
```

- ▶ Run the above command every time you update U-Boot
- ▶ This generates the `build/platform/generic/firmware/fw_payload.elf` file which is a binary that QEMU can boot.



Starting U-Boot in QEMU

```
qemu-system-riscv64 -m 2G \  
-nographic \  
-machine virt \  
-smp 8 \  
-bios opensbi/build/platform/generic/firmware/fw_payload.elf \  

```

- ▶ `-m`: amount of RAM in the emulated machine
- ▶ `-smp`: number of CPUs in the emulated machine

Exit QEMU with `[Ctrl][a]` followed by `[x]`



Linux kernel



Environment for kernel cross-compiling

- ▶ Download Linux 5.11-rc3 sources from <https://kernel.org>
- ▶ Let's add two environment variables for kernel cross-compiling to our `riscv64-env.sh` file:

```
export CROSS_COMPILE=riscv64-linux-  
export ARCH=riscv
```

- ▶ `CROSS_COMPILE` is the cross-compiler prefix, as our cross-compiler is `riscv64-linux-gcc`.
- ▶ `ARCH` is the name of the subdirectory in [arch/](#) corresponding to the target architecture.



Kernel configuration

- ▶ Lets take the default Linux kernel configuration for RISCV:

```
$ make help | grep defconfig
defconfig          - New config with default from ARCH supplied defconfig
savedefconfig      - Save current config as ./defconfig (minimal config)
alldefconfig       - New config with all symbols set to default
olddefconfig       - Same as oldconfig but sets new symbols to their
nommu_k210_defconfig      - Build for nommu_k210
nommu_virt_defconfig      - Build for nommu_virt
rv32_defconfig       - Build for rv32
$ make defconfig
```

- ▶ We can now further customize the configuration:

```
make menuconfig
```



Compiling the kernel

```
make
```

To compile faster, run multiple **j**obs in parallel:

```
make -j 8
```

To **re**compile faster (7x according to some benchmarks), run multiple **j**obs in parallel:

```
make -j 8 CC="ccache riscv64-linux-gcc"
```

At the end, you have these files:

`vmlinux`: raw kernel in ELF format (not bootable, for debugging)

`arch/riscv/boot/Image`: uncompressed bootable kernel

`arch/riscv/boot/Image.gz`: compressed kernel



Booting the kernel



Booting the Linux kernel directly

We could boot the Linux kernel directly as follows

```
cd opensbi
make PLATFORM=generic FW_PAYLOAD_PATH=../linux-5.11-rc3/arch/riscv/boot/Image
cd ..

qemu-system-riscv64 -m 2G \
    -nographic \
    -machine virt \
    -smp 8 \
    -kernel opensbi/build/platform/generic/firmware/fw_payload.elf \
    -append "console=ttyS0" \
```

However, what we want to demonstrate is the normal booting process:

OpenSBI → U-Boot → Linux → Userspace



Booting the Linux kernel from U-Boot

- ▶ We want to show how to set the U-Boot environment to load the Linux kernel and to specify the Linux kernel command line
- ▶ For this purpose, we will need some storage space to store the U-Boot environment, load the kernel binary, and also to contain the filesystem that Linux will boot on.
- ▶ Therefore, let's create a disk image to give some storage space for QEMU



Disk image creation (1)

- ▶ Let's create a 128 MB disk image:

```
dd if=/dev/zero of=disk.img bs=1M count=128
```

- ▶ Let's create two partitions in this image

```
cfdisk disk.img
```

- ▶ A first 64 MB primary partition (type `W95 FAT32 (LBA)`), marked as bootable
- ▶ A second partition with remaining space (default type: `Linux`)



Disk image creation (2)

- ▶ Let's access the partitions in this disk image:

```
sudo losetup -f --show --partscan disk.img  
/dev/loop2
```

```
ls -la /dev/loop2*  
brw-rw---- 1 root disk 7, 2 Jan 14 10:50 /dev/loop2  
brw-rw---- 1 root disk 259, 11 Jan 14 10:50 /dev/loop2p1  
brw-rw---- 1 root disk 259, 12 Jan 14 10:50 /dev/loop2p2
```

- ▶ We can now format the partitions:

```
sudo mkfs.vfat -F 32 -n boot /dev/loop2p1  
sudo mkfs.ext4 -L rootfs /dev/loop2p2
```



Copying the Linux image to the FAT partition

- ▶ Let's create a mount point for the FAT partition:

```
mkdir /mnt/boot
```

- ▶ Let's mount it:

```
sudo mount /dev/loop2p1 /mnt/boot
```

- ▶ Let's copy the kernel image to it:

```
sudo cp linux-5.11-rc3/arch/riscv/boot/Image /mnt/boot
```

- ▶ And then unmount the filesystem to commit changes:

```
sudo umount /mnt/boot
```



Recompiling U-Boot for environment support

We want U-Boot be able to use an environment in a FAT partition on a virtio disk.

- ▶ So, let's reconfigure U-Boot with the following settings

```
make menuconfig
```

- ▶ `CONFIG_ENV_IS_IN_FAT=y`
- ▶ `CONFIG_ENV_FAT_INTERFACE="virtio"`
- ▶ `CONFIG_ENV_FAT_DEVICE_AND_PART="0:1"`

- ▶ Then recompile U-Boot

```
make -j8
```

- ▶ Then update the firmware loader:

```
cd ../opensbi  
make PLATFORM=generic FW_PAYLOAD_PATH=../u-boot-2021.01/u-boot.bin
```



Run U-Boot with an environment

- ▶ Add a disk to the emulated machine:

```
qemu-system-riscv64 -m 2G -nographic -machine virt -smp 8 \  
  -bios opensbi/build/platform/generic/firmware/fw_payload.elf \  
  -drive file=disk.img,format=raw,id=hd0 \  
  -device virtio-blk-device,drive=hd0 \
```

- ▶ In U-Boot, you should now be able to save an environment:

```
setenv foo bar  
saveenv
```



Booting Linux from U-Boot



Requirements for booting Linux

To boot the Linux kernel, U-Boot needs to load

- ▶ A Linux kernel image. In our case, let's load it from our virtio disk to RAM (find a suitable RAM address by using the `bdinfo` command in U-Boot):

```
fatload virtio 0:1 84000000 Image
```

- ▶ Possibly the image of an *Initramfs*, a filesystem in RAM that Linux can use.
- ▶ A *Device Tree Binary (DTB)*, letting the kernel know which SoC and devices we have. This allows the same kernel to support many different SoCs and boards.
 - ▶ *DTB* files are compiled from *DTS* files in [arch/riscv/boot/dts/](https://github.com/bootlin/arch-riscv/boot/dts/)
 - ▶ However, there is no such *DTS* file for the RISC-V QEMU virt board.
 - ▶ The *DTB* for our board is actually passed by QEMU to OpenSBI and then to U-Boot. See <https://tinyurl.com/y4ae5ptd>
 - ▶ In U-Boot, at least in our case, the *DTB* is available in RAM at address `${fdtcontroladdr}`



Linux kernel command line

- ▶ We need to set the Linux arguments (*kernel command line*)

```
setenv bootargs 'root=/dev/vda2 rootwait console=ttyS0 earlycon=sbi rw'
```

- ▶ `root=/dev/vda2`
Device for Linux to mount as root filesystem
- ▶ `rootwait`
Wait for the root device to be ready before trying to mount it
- ▶ `console=ttyS0`
Device (here first serial line) to send Linux booting messages to
- ▶ `earlycon=sbi`
Allows to have more messages before the console driver is initialized (*Early Console*).
- ▶ `rw`
Allows to mount the root filesystem in read-write mode.



Booting Linux

- ▶ Here's the command to boot the Linux `Image` file:

```
booti <Linux address> <Initramfs address> <DTB address>
```

- ▶ In our case:

```
booti 0x84000000 - ${fdtcontroladdr}
```

- ▶ So, let's define the default series of commands that U-Boot will automatically run after a configurable delay (`bootdelay` environment variable):

```
setenv bootcmd 'fatload virtio 0:1 84000000 Image; booti 0x84000000 - ${fdtcontroladdr}'
```

- ▶ Finally, we must save these new settings:

```
saveenv
```

- ▶ ... and boot our system (`boot` runs `bootcmd`):

```
boot
```



Building the root filesystem



BusyBox - About 10 years ago

Most commands in one binary!

```
[, [[, acpid, addgroup, adduser, adjtimex, ar, arp, arping, ash, awk, basename, beep, blkid, brctl, bunzip2, bzip2, cal, cat, catv, chat, chattr, chgrp, chmod, chown, chpasswd, chpst, chroot, chrt, chvt, cksum, clear, cmp, comm, cp, cpio, crond, crontab, cryptpw, cut, date, dc, dd, dealloct, delgroup, deluser, depmod, devmem, df, dhcprelay, diff, dirname, dmesg, dnsd, dnsdomainname, dos2unix, dpkg, du, dumpkmap, dumpleases, echo, ed, egrep, eject, env, envdir, envuidgid, expand, expr, fakeidentd, false, fbset, fbsplash, fdflush, fdformat, fdisk, fgrep, find, findfs, flash_lock, flash_unlock, fold, free, freeramdisk, fsck, fsck.minix, fsync, ftpd, ftpget, ftpput, fuser, getopt, getty, grep, gunzip, gzip, hd, hdparm, head, hexdump, hostid, hostname, httpd, hush, hwclock, id, ifconfig, ifdown, ifenslave, ifplugd, ifup, inetd, init, inotifyd, insmod, install, ionice, ip, ipaddr, ipcalc, ipcrm, ipcs, iplink, iproute, iprule, iptunnel, kbd_mode, kill, killall, killall5, klogd, last, length, less, linux32, linux64, linuxrc, ln, loadfont, loadkmap, logger, login, logname, logread, losetup, lpd, lpq, lpr, ls, lsattr, lsmod, lzmacat, lzop, lzopcat, makemime, man, md5sum, mdev, mesg, microcom, mkdir, mkdosfs, mkfifo, mkfs.minix, mkfs.vfat, mknod, mkpasswd, mkswap, mktmp, modprobe, more, mount, mountpoint, mt, mv, nameif, nc, netstat, nice, nmeter, nohup, nslookup, od, openvt, passwd, patch, pgrep, pidof, ping, ping6, pipe_progress, pivot_root, pkill, popmaildir, printenv, printf, ps, pscan, pwd, raidautorun, rdate, rdev, readlink, readprofile, realpath, reformime, renice, reset, resize, rm, rmdir, rmmmod, route, rpm, rpm2cpio, rtcwake, run-parts, runlevel, runsv, runsvdir, rx, script, scriptreplay, sed, sendmail, seq, setarch, setconsole, setfont, setkeycodes, setlogcons, setsid, setuidgid, sh, shasum, sha256sum, sha512sum, showkey, slattach, sleep, softlimit, sort, split, start-stop-daemon, stat, strings, stty, su, sulogin, sum, sv, svlogd, swapoff, swapon, switch_root, sync, sysctl, syslogd, tac, tail, tar, taskset, tcpsvd, tee, telnet, telnetd, test, tftp, time, timeout, top, touch, tr, traceroute, true, tty, ttysize, udhcpc, udhcpd, udpsvd, umount, uname, uncompress, unexpand, uniq, unix2dos, unlzma, unlzop, unzip, uptime, usleep, uudecode, uuencode, vconfig, vi, vlock, volname, watch, watchdog, wc, wget, which, who, whoami, xargs, yes, zcat, zcip
```

Source: `run /bin/busybox`



BusyBox - In 2019

[, [[, acpid, add-shell, addgroup, adduser, adjtimex, ar, arch, arp, arping, awk, base64, basename, bbconfig, bc, beep, blkdiscard, blkid, blockdev, bootchartd, brctl, bunzip2, busybox, bzip2, cal, cat, chat, chattr, chcon, chgrp, chmod, chown, chpasswd, chpst, chroot, chrt, chvt, cksum, clear, cmp, comm, conspy, cp, cpio, crond, crontab, cryptpw, ctyhack, cut, date, dc, dd, dealloct, delgroup, deluser, depmod, devmem, df, diff, dirname, dmesg, dnsd, dnsdomainname, dos2unix, dpkg, dpkg-deb, du, dumpkmap, dumpleases, echo, ed, egrep, eject, env, envdir, envuidgid, etherwake, expand, expr, factor, fakeidentd, fallocation, false, fatattr, fbset, fbsplash, fdflush, fdformat, fdisk, fgconsole, fgrep, find, findfs, flash_eraseall, flash_lock, flash_unlock, flashcp, flock, fold, free, freeramdisk, fsck, fsck.minix, fsfreeze, fstrim, fsync, ftpd, ftpget, ftpget, ftpput, fuser, getenforce, getopt, getsebool, getty, grep, groups, gunzip, gzip, halt, hd, hdparm, head, hexdump, hexedit, hostid, hostname, httpd, hush, hwclock, id, ifconfig, ifdown, ifenslave, ifplugd, ifup, inetd, init, insmod, install, ionice, iostat, ip, ipaddr, ipcalc, ipcrm, ipcs, iplink, ipneigh, iproute, iprule, iptunnel, kbd_mode, kill, killall, killall5, klogd, last, less, link, linux32, linux64, linuxrc, ln, load_policy, loadfont, loadkmap, logger, login, logname, logread, losetup, lpd, lpq, lpr, ls, lsattr, lsmod, lsof, lspci, lsscsi, lsusb, lzcat, lzma, lzop, lzopcat, makedevs, makemime, man, matchpathcon, md5sum, mdev, msg, microcom, minips, mkdir, mkdosfs, mke2fs, mkfifo, mkfs.ext2, mkfs.minix, mkfs.reiser, mkfs.vfat, mknod, mkpasswd, mkswap, mktmp, modinfo, modprobe, more, mount, mountpoint, mpstat, mt, mv, nameif, nanddump, nandwrite, nbd-client, nc, netcat, netstat, nice, nl, nmeter, nohup, nologin, nproc, nsenter, nslookup, ntpd, nuke, od, openvt, partprobe, passwd, paste, patch, pgrep, pidof, ping, ping6, pipe_progress, pivot_root, pkill, pmap, popmaildir, poweroff, printenv, printf, ps, pscan, pstree, pwd, pwdx, raidautorun, rdate, rdev, readahead, readlink, readprofile, realpath, reboot, reformime, remove-shell, renice, reset, resize, restorecon, resume, rev, rkill, rm, rmdir, rmmmod, route, rpm, rpm2cpio, rtcwake, run-init, run-parts, runcon, runlevel, runsv, runsvdir, rx, script, scriptreplay, sed, selinuxenabled, sendmail, seq, sestatus, setarch, setconsole, setenforce, setfattr, setfiles, setfont, setkeycodes, setlogcons, setpriv, setsebool, setserial, setuid, setuidgid, sh, sha1sum, sha256sum, sha3sum, sha512sum, showkey, shred, shuf, slattach, sleep, smemcap, softlimit, sort, split, ssl_client, start-stop-daemon, stat, strings, stty, su, sulogin, sum, sv, svc, svlogd, svok, swapoff, swapon, switch_root, sync, sysctl, syslogd, tac, tail, tar, taskset, tc, tcpsvd, tee, telnet, telnetd, test, tftpd, tftpd, time, timeout, top, touch, tr, traceroute, traceroute6, true, truncate, ts, tty, ttysize, tunc1, tune2fs, ubiattach, ubidetach, ubimkvol, ubirename, ubirmvol, ubirsvol, ubiupdatevol, udhcpc, udhcpd, udpsvd, uevent, umount, uname, uncompress, unexpand, uniq, unit, unix2dos, unlink, unlzma, unlzop, unxz, unzip, uptime, users, usleep, uudecode, uuencode, vconfig, vi, vlock, volname, w, wall, watch, watchdog, wc, wget, which, who, whoami, whois, xargs, xxd, xz, xzcat, yes, zcat, zcip




BusyBox - Dowloading

- ▶ Download BusyBox 1.33.0 sources from <https://busybox.net>
- ▶ Extract the archive with `tar xf`



BusyBox - Configuring

- ▶ Run `make allnoconfig`
Starts with no applet selected
- ▶ Run `make menuconfig`
 - ▶ In Settings → Build Options, enable
Build static binary (no shared libs)
 - ▶ In Settings → Build Options, set
Cross compiler prefix to `riscv64-linux-`
 - ▶ Then enable support for the following commands:
`ash`, `init`, `halt`, `mount`, `cat`, `mkdir`, `echo`, `ls`, `chmod`,
`uptime`, `vi`, `ifconfig`, `httpd`



<https://asciinema.org/a/384727>



BusyBox - Installing and compiling

- ▶ Compiling: `make` or `make -j 8` (faster)
Resulting size: only 276344 bytes!
Funny to see that we're using a 64 bit system to run such small programs!
- ▶ Installing in `_install/`: `make install`
- ▶ See the created directory structure and the symbolic links to `/bin/busybox`
- ▶ Installing to the root filesystem:

```
sudo mkdir /mnt/rootfs
sudo mount /dev/loop2p2 /mnt/rootfs
sudo rsync -aH _install/ /mnt/rootfs/
sudo umount /mnt/rootfs
```

```
.
├── bin
│   ├── ash -> busybox
│   ├── busybox
│   ├── cat -> busybox
│   ├── echo -> busybox
│   ├── ls -> busybox
│   ├── mkdir -> busybox
│   ├── mount -> busybox
│   ├── sh -> busybox
│   └── vi -> busybox
├── sbin
│   ├── halt -> ../bin/busybox
│   ├── ifconfig -> ../bin/busybox
│   └── init -> ../bin/busybox
└── usr
    ├── bin
    │   └── uptime -> ../../bin/busybox
    └── sbin
        └── httpd -> ../../bin/busybox
```



Completing and configuring the root filesystem (1)

- ▶ We need to create a `dev` directory.
The `devtmpfs` filesystem will automatically be mounted there
(as `CONFIG_DEVTMPFS_MOUNT=y`)
- ▶ Let's also create `/proc` and `/sys` so that we can also mount the `proc` and `sysfs` virtual filesystems on the target:

```
mount -t proc none /proc
mount -t sysfs none /sys
```

Without such virtual filesystems, commands such as `ps` or `top` can't work.



Completing and configuring the root filesystem (2)

Let's automate the mounting of `proc` and `sysfs`...

- ▶ Let's create an `/etc/inittab` file to configure Busybox Init:

```
# This is run first script:
::sysinit:/etc/init.d/rcS
# Start an "askfirst" shell on the console:
::askfirst:/bin/sh
```

- ▶ Let's create and fill `/etc/init.d/rcS` to automatically mount the virtual filesystems:

```
#!/bin/sh
mount -t proc nodev /proc
mount -t sysfs nodev /sys
```



Common mistakes

- ▶ Don't forget to make the `rcS` script executable. Linux won't allow to execute it otherwise.
- ▶ Do not forget `#!/bin/sh` at the beginning of shell scripts! Without the leading `#!` characters, the Linux kernel has no way to know it is a shell script and will try to execute it as a binary file!
- ▶ Don't forget to specify the execution of a shell in `/etc/inittab` or at the end of `/etc/init.d/rcS`. Otherwise, execution will just stop without letting you type new commands!



Add support for networking (1)

- ▶ Add a network interface to the emulated machine:

```
sudo qemu-system-riscv64 -m 2G -nographic -machine virt -smp 8 \  
-bios opensbi/build/platform/generic/firmware/fw_payload.elf \  
-drive file=disk.img,format=raw,id=hd0 \  
-device virtio-blk-device,drive=hd0 \  
-netdev tap,id=tapnet,ifname=tap2,script=no,downscript=no \  
-device virtio-net-device,netdev=tapnet \  

```

- ▶ Need to be `root` to bring up the `tap2` network interface



Add support for networking (2)

- ▶ On the target machine:

```
ifconfig -a  
ifconfig eth0 192.168.2.100
```

- ▶ On the host machine:

```
ifconfig -a  
sudo ifconfig tap2 192.168.2.1  
ping 192.168.2.100
```



Simple CGI script

```
#!/bin/sh
echo "Content-type: text/html"
echo
echo "<html>"
echo "<meta http-equiv=\"refresh\" content=\"1\">"
echo "<header></header><body>"
echo "<h1>Uptime information</h1>"
echo "Your embedded device has been running for:<pre><font color=Blue>"
echo `uptime`
echo "</font></pre>"
echo "</body></html>"
```

Store it in `/www/cgi-bin/uptime` and make it executable.



Start a web server

- ▶ On the target machine:

```
/usr/sbin/httpd -h /www
```

- ▶ On the host machine, open in your browser:
<http://192.168.2.100/cgi-bin/uptime>



What to remember

- ▶ Embedded Linux is just made out of simple components. It makes it easier to get started with Linux.
- ▶ You just need a toolchain, a bootloader, a kernel and a few executables.
- ▶ RISC-V is a new, open Instruction Set Architecture, use it and support it!
- ▶ In embedded Linux, things don't change that much over time. You just get more features.



Going further and thanks

- ▶ Drew Fustini's unmatched presentation about Linux on RISC-V:
<https://tinyurl.com/y6j8lfyz>
- ▶ Our "Embedded Linux system development" training materials (500+ pages, CC-BY-SA licence):
<https://bootlin.com/doc/training/embedded-linux/>
- ▶ All our training materials and conference presentations:
<https://bootlin.com/docs/>
- ▶ The Embedded Linux Wiki: presentations, howtos... contribute to it!
<https://elinux.org>
- ▶ Gratitude to Geert Uytterhoeven, my mentor in conference fashion, always wearing a smile and the most relevant and elegant conference T-shirts.

Questions?
Suggestions?
Comments?

Michael Opdenacker
michael.opdenacker@bootlin.com

Slides under CC-BY-SA 3.0
<https://bootlin.com/pub/conferences/2021/fosdem/>