



CE Workgroup

Status of Embedded Linux

March 2016

Tim Bird

Architecture Group Chair

LF CE Workgroup



CE Workgroup

Outline

Kernel Versions
Technology Areas
CE Workgroup Projects
Other Stuff
Resources



CE Workgroup

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff

Resources



CE Workgroup

Kernel Versions

- Linux v4.0 – 12 Apr 2015 – 63 days
- Linux v4.1 – 21 Jun 2015 – 70 days
- Linux v4.2 – 30 Aug 2015 – 70 days
- Linux v4.3 – 1 Nov 2015 – 63 days
- Linux v4.4 – 10 Jan 2016 – 70 days
- Linux v4.5-rc7 – currently at 60 days
 - Should see Linux v4.5 on March 13



CE Workgroup

Linux v4.0

- This version is not v3.20
- Android binder has security hooks
 - Can use SELinux security with it
- Non-volatile memory support patches
 - Can use filesystem in persistent memory
 - <http://lwn.net/Articles/610174/>
- UBIFS performance improvements



CE Workgroup

Linux v4.1

- New tracefs filesystem
- Kernel self-test 'install' target
- Ability to attach BPF programs to kernel probes
- I2C subsystem can function in slave mode
- Can configure kernel for single-user operation



CE Workgroup

Linux v4.2



- Linux security module stacking
 - See <https://lwn.net/Articles/635771/>
- F2FS supports per-file encryption
- Support for AMD GPUs
- Lots of pin control drivers:
 - Freescale, Mediatek, Allwinner, Qualcomm, Renesas
- Libnvdimm – non-volatile memory (NVM) management



CE Workgroup

Linux v4.3

- MOST (Media Oriented Systems Transport) support is in staging
 - MOST is a framework in automotive market for multimedia networking
- Ext3 removed
 - But ext4 code supports ext3 filesystems





CE Workgroup

Linux v4.4

- LightNVM feature
 - Take control of low-level SSD features
 - Will talk about this later
- Perf can build and load eBPF files
- Arm64 can have 16K pages
- Broadcom VC4 GPU (raspberry pi)
- Devfreq cooling – thermal management
- Various PWM drivers



CE Workgroup

Linux v4.5

- ARM multiplatform hits an important milestone
 - Major patch including lots of minor platforms
 - Many v6 and v7 platforms are now supported
- Not much else specific to embedded
 - Well, continued mainlining of drivers for SoC features



Things to watch (from past)

- Kdbus
 - Has hit some stumbling blocks getting merged
- Kernel tinification!
- RT-preempt
- Persistent memory
 - (NVM = Non-Volatile Memory)
- SoC mainlining progress



Things to watch (status)

- Kdbus (**stalled**)
 - Has hit some stumbling blocks getting merged
- Kernel tinification! (**stalled**)
- RT-preempt (**in progress**)
- Persistent memory (**in progress**)
 - Good talk on about issues:
 - “Making use of persistent memory”
 - <http://lwn.net/Articles/674752/>
- SoC mainlining progress (**in progress**)



CE Workgroup

Kernel process analysis

- Does kernel development scale
 - Does the kernel patch acceptance process scale?
- “How 4.4’s patches got to the mainline”
 - <http://lwn.net/Articles/670209/>
 - Has an interesting graph
- Tree is very flat
 - Only a few areas where patch flows through more than one maintainer tree
 - Networking is a good example



Kernel process analysis (cont.)

- Some conclusions:
 - Most maintainers push directly to Linus
 - Patch path is shorter than expected
 - Linus trusts his sub-maintainers
 - Linux pulls about 300 trees each release
 - Adding a little depth could scale the process even more, with no slow-down
 - Currently at about 1500 developers and 11,000 commits per release, but should be able to scale to many more
- Linus has pushed for “maintainer groups”
- There’s much more automated testing



CE Workgroup

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff

Resources



CE Workgroup

Bootup Time

- Mostly old news...
- XIP on x86
 - See <https://lwn.net/Articles/637532/>
- Asynchronous probing
 - Discussed at kernel summit
- Reduction in probe deferral
 - Explicit probe ordering can be used to get a specific subsystem (like display) up sooner
 - The “On-demand probing” patches were NAKed
 - Need to measure effect on overall boot time



Bootup Time (cont.)

- No talks at ELC this year
 - But boot time is NOT a solved problem
 - Boot time issues are unique per platform, and reductions tend not to be mainlinable
 - e.g. remove stuff not needed
- Some good previous talks:
 - ELCE 2014 - *12 Lessons Learnt in Boot Time Reduction* by Andrew Murray
 - ELC 2015 - *Fastboot Tools and Techniques* by John Mehaffey



Device Tree

- Device Tree Overlays
 - Seems to be working as intended
 - Will be session at ELC 2016 by Pantellis on making overlays independent of the base board
 - Should allow add-on boards to be used with different platforms
- New Maintainer – Frank Rowand
- Device Tree validation
 - Project by Matt Porter and others
 - Schema for binding language, validator for bindings and for device tree data
 - **Work is stalled**
- Updated Device Tree spec possibly in works
 - Want to update material and make it more available



CE Workgroup

Graphics

- Vulkan API from Khronos Group
 - Alternative to Direct3D or OpenGL
 - Reduce CPU overhead for CPU/GPU operations
 - AMD announced plans to open source the driver (but Intel and Valve already working it)
 - Version 1.0 is now available
 - Nvidia now supports it
- Qt license change
 - From LGPL 2.0 to LGPL 3.0
 - Companies scrambling to find alternative
 - GPL/LGPL 3.0 is undesirable for CE products



CE Workgroup

GPUs and OSS support

- Integrated GPUs
 - AMD, Intel, Nvidia, Qualcomm:Adreno
- GPU IP suppliers
 - ARM:Mali, Imagination:PowerVR, Vivante
- GPU support
 - Freedreno – Adreno (good progress)
 - ??? – for PowerVR (no progress)
 - Etnaviv – for Vivante (good progress)
 - Nouveau – for Nvidia (not sure of status)
 - Lima – for Mali (no progress)



CE Workgroup

Freedreno

- GPL driver for Adreno GPU on Qualcomm chips
 - 3xx supports OpenGL ES 3.0
 - 4xx supports OpenGL ES 3.1
- There are still some pieces that need work
 - Bug reports are appreciated
- Some interesting reverse-engineering tools developed for the project
 - <https://github.com/freedreno/freedreno/wiki/Reverse-engineering-tools>
- <http://lwn.net/Articles/638908/>



CE Workgroup

PowerVR

- In June 2015: Imagination Executive blogged:

Q: Is there plans to make/help/fund open PowerVR driver for Linux?

A: Yes, there is a plan and it is one of the things I've been working on for the past few months. Hopefully I'll have something more to share soon(-ish?).

Read more: <http://www.cnx-software.com/2015/06/18/open-source-linux-drivers-for-powervr-gpus-might-be-in-the-works/#ixzz3dSpJ9bhl>

- No word since then...



CE Workgroup

Etnaviv



- Etnaviv – for Vivante
 - Replaced 65K kernel driver with 6.5K driver
 - See ELCE 2015 talk: “Bringing up FOSS GPU Drivers on Freescale i.MX6 Systems” by Lucas Stach
 - Slides now available for this talk
 - Also see See <http://lwn.net/Articles/659391/>
- Stuff hit mainline in January:
 - <https://git.kernel.org/cgit/linux/kernel/git/torvalds/linux.git/commit/drivers/gpu/drm/etnaviv?id=a8c21a5451d831e67b7a6fb910f9ca8bc7b43554>
 - From “the etnaviv authors”



CE Workgroup

File Systems

- Proposals for UBIFS handling of MLC NAND
 - Lots of complexity due to MLC characteristics
 - See “NAND Support: (New?) Challenges for the MTD/NAND Subsystem” – Boris Brezillon (at ELC)
- EXT3 removed from kernel (4.3-rc1)





File Systems (cont.)

- ELC talks:
 - “Filesystem Considerations for Embedded Devices” – Tristan Lelong
 - Great talk with performance and robustness results for different file systems
 - Ext4, BTRFS, F2FS, XFS, NILFS2
 - Summary: F2FS is faster in many cases, EXT4 is mature
 - “Current Challenges in UBIFS” – ELCE 2015
 - Richard Weinberger



(new) LightNVM

- Framework for holding SSD parameters
- Allows kernel to manage flash translation layer
- SSDs have weird (black-box) FTL implementations
 - Are often optimized for FAT filesystems
 - Recent drives allow direct access to blocks
- See <http://lwn.net/Articles/641247/>
 - “The host primarily handles data placement, I/O scheduling, and garbage collection and leaves everything else to the SSD controller”



Networking

- Bluetooth:
 - Bluetooth 4.2 has better security, faster speeds
 - 6lowpan integration
 - Working on mesh networking
- New protocols for IOT
 - Thread – Nest’s low-power IP stack
 - Others (Sigfox, LoRaWan, etc.)
- Visible Light Communication (VLC)
 - Disney’s Linux Light Bulb
 - Low-bandwidth via LED-to-LED
 - Allows toy to have cheap transmitter/sensor



CE Workgroup

Power Management

- “Dynamic Audio Power Management”
 - ELCE 2015 talk by Lars Peter Clausen
 - New system that manages a graph of relationships between IP blocks on the system
 - Turns on/off power, or scales it, according to needs
 - Lots of graph traversal
 - Could be used for other systems (e.g. video)



Hardware power reduction

- **Passive WiFi**
 - Modulate reflected WiFi instead of broadcasting
 - 10,000x less power for mobile device
 - Research by University of Washington
 - Still in research, but promising
 - <http://www.networkworld.com/article/3037088/mobile-wireless/researchers-make-low-power-wi-fi-breakthrough.html>



Real Time – RT-preempt

- Linux Foundation Real-Time Linux Collaborative project
 - Thomas Gleixner is a Linux Foundation fellow
 - Should result in more stuff going upstream
 - One interesting note: press release says they'll meet regularly at ELC
 - Thomas will have keynote session at ELC 2016
- Latest RT-preempt is for 4.4 kernel
 - **Just released on this week!**
 - Tends to follow LTS releases
 - See <https://www.kernel.org/pub/linux/kernel/projects/rt/>



Real Time - other

- Xenomai 3.0 is out !! (actually, 3.0.1)
 - Uses Cobalt RT core
 - 3.0 supports both dual-kernel and single-kernel configurations (using RT-preempt)
 - See xenomai.org
- Some RT talks
 - ELCE 2015 – Practical Real-Time Linux – by Arnout Vandecappelle
 - Presentation on Xenomai at ELC 2016





Security

- “Making your own security modules” – Casey Schaufler
 - <http://lwn.net/Articles/674949/>
 - Promote experimentation by giving tips on how to write your own security modules
- Security module stacking
 - Added in kernel 4.2
 - See <https://lwn.net/Articles/635771/>
- New project for kernel security issues:
 - http://kernsec.org/wiki/index.php/Kernel_Self_Protection_Project



Security and IOT

- IOT raises lots of security issues
- See “Kernel security hacking for the Internet of Things” – Daniel Sangorrin (at LCJ)
 - Reduce attack surface
 - Can detect attacks by detecting variation from pre-determined behavior
 - Isolate critical software





CE Workgroup

System Size

- Kernel tinification project is **stalled**
 - **Tiny repository removed from linux-next**
 - **No activity in one year!**
- Single-user patches
 - Gets rid of users and groups
 - Saves about 25K
 - <http://lwn.net/Articles/631853/>
 - Mainlined in kernel v4.1
- Removal of kernel command-line parsing
 - Not mainlined



CE Workgroup

System Size (cont.)

- Intel X86 XIP patches
 - See <https://lwn.net/Articles/637532/>
- Nicolas Pitre has done work recently on supporting gcc --gc-sections
 - Lighter-weight option similar to LTO
- Linux Foundation announces IOT RTOS
 - Zephyr
 - Does this mean that we're giving up on Linux size reductions??
 - (maybe)



CE Workgroup

Testing

- Kselftest
- LTSI Test Project
- Kernelci.org
- Lots of automated testing talks at ELC 2016



CE Workgroup

kselftest

- Inside kernel source tree
 - Makefile target: ‘make kselftest’
- Ability to install tests mainlined in kernel v4.1
 - Cross-build now supported?
 - I didn’t have time to test this myself
 - <http://lwn.net/Articles/628625/>
- See “Linux Kernel Selftest Framework BoFs – Quality Control for New Releases” – Shuah Khan (at ELC)
- See <http://lwn.net/Articles/608959/>



CE Workgroup

LTSI test project

- Jenkins-based Test Automation (JTA)
 - I've asked to rename it Fuego
 - Waiting for my patch to be applied
- Available now
 - <https://bitbucket.org/cogentembedded/jta-public/>
- Hard to tell what adoption rate is
- Want to identify some specific verticals, and build tests for them
 - Otherwise, it's a solution in search of a problem



CE Workgroup

Kernelci.org

- Place to get free build/boot testing for your board
 - “ci” = continuous integration
 - Builds 126 trees continuously, then reports any errors
- <http://kernelci.org>
- ELC and ELCE 2015 - *Upstream Kernel Testing* – by Kevin Hilman
- Sony Mobile has a phone in this farm



CE Workgroup

Toolchains

- Khem Raj has added support to the Yocto Project for Clang (LLVM)
 - Builds all but about 45 packages
 - He has a mini-distro with kernel, musl, toybox, built with clang
 - Will have presentation at ELC 2016



Tracing





- eBPF to be used for dynamic tracing
 - Perf supports eBPF (in 4.4)
 - eBPF = extended Berkeley Packet Filter
- New tracefs filesystem
 - No longer part of debugfs
 - But all (psuedo) dirs and files the same
- Histograms (not mainlined yet)
- See “New (and Exciting!) Development in Linux Tracing – Elena Zannoni (at LCJ 2015)”



CE Workgroup

Miscellaneous

- J2
- Current LTS kernel version:
 - 4.4
- IOT news





CE Workgroup

J2

- Open hardware processor
- Formerly SH2, but patents have expired
- See <http://lwn.net/Articles/647636/>
“Resurrecting the SuperH architecture”
- Resurgence of nommu Linux?
- Someday might run Linux on 3-cent processors



CE Workgroup

Weird IOT news

- Microsoft released Windows 10 IoT kit for Raspberry PI
- “Linux” Foundation announces non-Linux RTOS for Internet of Things
- Google Brillo is available (by invitation)
 - Java-less, headless, Android



CE Workgroup

LF RTOS for IOT

- Zephyr – RTOS from Wind River
- Idea is to target devices that Linux will never support
 - E.g. sensors
- Attributes:
 - Apache 2 license
 - Minimal size – as small as 8K
 - Highly configurable
 - NoMMU
 - Networking: WiFi, Bluetooth, NFC
- Governed by committee



CE Workgroup

Short rant

- I'm not sure why this is called an RTOS
 - There's no RealTime requirement
- It should be called IoTOS

- (rant over)



CE Workgroup

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff





Resources



CE Workgroup

CEWG Projects

- Contract work
- Projects and initiatives









CE Workgroup

CEWG Contract Work

- Kernel string refactoring
- Device tree documentation
- LTSI test framework
- Shared distribution testing





CE Workgroup

Kernel string refactoring

- Description
 - Refactor kernel strings to reduce the space used for statically-defined strings
 - http://elinux.org/Refactor_kernel_strings
- Contractor: Wolfram Sang
- Based on results from last year's compressed printk investigation
 - Aiming for at least 50K of savings, depending on kernel config
- Project is in progress
 - Report expected at LinuxCon Japan



CE Workgroup

DT documentation





- Working on “guide” documentation
- Frank Rowand has been collecting data and giving talks
 - LinuxCon NA, ELCE, ELC and LCJ
- Will be put on elinux wiki at:
 - http://elinux.org/Linux_Drivers_Device_Tree_Guide



CE Workgroup

LTSI test framework

- (Discussed previously)







CE Workgroup

Shared Distribution Testing

- See “Shared Embedded Distribution” project (later in these slides)
- Project to test distribution on a few different hardware platforms
- Contractor: Tuan Hoang
- Status: Just starting









CE Workgroup

Projects and initiatives

- Civil Infrastructure
- Shared Embedded Distribution
- Device Mainlining
- LTSI
- eLinux wiki





CE Workgroup

Civil Infrastructure

- **Goals**
 - Solve problems with Linux for use in civil infrastructure systems
- **Status**
 - Recent Activity
 - BOFS at many recent events
 - Private meetings to discuss goals with interested companies
 - Have organized some companies to work on the project
- **Next steps:**
 - Activities are being planned
 - Presentation at ELC 2016



CE Workgroup

Shared Embedded Distribution

- **Goals**

- Create an industry-supported distribution of embedded Linux
 - Main goal is very long term support (15 years)

- **Status**

- Toshiba has created Yocto layer meta-Debian
- Presented at ELCE, ELC, and LCJ

- **Next steps**

- Get more companies collaborating on the project
- Presentation at ELC 2016



CE Workgroup

Device Mainlining

- http://elinux.org/CE_Workgroup_Device_Mainlining_Project
- Goal is to study obstacles to mainlining, and work to reduce obstacles
- Previous Activity
 - Developer survey in 2014
 - SIG/BOF meetings at ELCE, ELC, LCNA and Linaro Connect
 - Presentations about overcoming obstacles
 - See <http://lwn.net/Articles/647524/>
 - White paper (published at LCJ – June 2015)



Device Mainlining (cont.)



- Mobile phone source analysis
 - Phone kernels have between 1.1 and 3.1 million lines of code out-of-tree
- Published tools:
 - <https://github.com/tbird20d/upstream-analysis-tools>
- Ongoing Projects:
 - Presentation at Collab Summit 2016
 - Possibly create some training materials
 - Qualcomm report on mainline status at ELC 2016
 - Create tools for easier mainlining



CE Workgroup

Long Term Support Initiative

- LTSI 4.1 is latest kernel
- Many presentations available on status
- Latest project push is testing facility
 - See previous page on JTA test framework
- Kernel diversion measurement tool
 - Presentation at ELC 2016





CE Workgroup

eLinux wiki

- <http://elinux.org>
 - Web site dedicated to information for embedded Linux developers
 - The wikipedia of embedded linux!
- Hundreds of pages covering numerous topic areas: bootup time, realtime, security, power management, flash filesystem, toolchain, editors
- Lots of pages in last few years about low-cost development boards
- Please use and add to site



CE Workgroup

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff





Resources



CE Workgroup

Other Stuff

- Projects and Consortia
- Distros and Build Systems
- Events
- Hardware





Projects and Consortia

- Allseen Alliance – Peer-to-peer ad-hoc networking
 - AllJoyn is the name of the implementation
- Open Connectivity Foundation (OCF)
 - Adopted OIC/Iotivity technology, along with UPnP stuff
 - I think there's some kind of OCF/Allseen collaboration, but I'm not sure
- DroneCode – Open source UAV software
 - <http://www.dronecode.org/>
 - Have a good drone track at ELC 2016



Projects and Consortia

- Linaro
 - Linaro IoT and Embedded initiative (LITE)
 - Waiting to see their proposal
 - Run Linux on Cortex A and mbedOS on Cortex M
 - Unsure about licensing for Cortex M
- PRPL Foundation
 - Multi-company MIPS non-profit
 - Projects: PRPL OpenWRT, MIPS QEMU



CE Workgroup

Distros

- Android
 - Working on “N” version
 - New build system under development, using ‘go’ language and something called blueprints
 - Google switching to OpenJDK
 - Eliminates those troublesome Oracle Java libraries
- Tizen
 - Lots of security work
- CEWG Shared embedded distribution
 - (see previous slides)



Build Systems

- **OpenEmbedded/Yocto Project**
 - 2.0 (Jethro) released
 - 1.8 allowed builds and runs with Toaster (web interface)
 - HOB is gone
 - Presentation on Toaster at ELC 2016
- **Buildroot**
 - Configurable support for static linking
 - Improved support for package hashes
 - Better warnings about toolchain header safety issues
 - License reporting?



CE Workgroup

Events

- **Embedded Linux Conference Europe 2015**
 - October 5-7, 2015 - Dublin, Ireland
 - Lots of content - check for slides on elinux wiki
- **Embedded Linux Conference 2016**
 - April 4-6, 2016 - San Diego, USA
 - Please come if you can make it
 - Should be lots of fun
- **Embedded Linux Conference Europe 2016**
 - October 6-7, 2016 - Berlin, Germany



CE Workgroup

Hardware

- Samsung ARTIK 1 processor
 - 1MB RAM, 4MB Flash, BLE 4.0
 - 12x12mm
 - Running Nucleus (frowny face)
- Raspberry Pi Zero
 - \$5 computer
 - It came for free with the December issue of MagPI magazine
 - 1GHz, 512M RAM, faster than Pi 1
 - Good review at:
 - <https://www.youtube.com/watch?v=NFFQmdUc5Vg>



CE Workgroup

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff

Resources





CE Workgroup

Resources

- LWN.net
 - <http://lwn.net/>
 - If you are not subscribed, please do so
- Kernel Newbies
 - [http://kernelnewbies.org/Linux_\[34\].?](http://kernelnewbies.org/Linux_[34].?)
- eLinux wiki - <http://elinux.org/>
 - Especially <http://elinux.org/Events> for slides
- Celinux-dev mailing list



CE Workgroup



Thanks!

