



CE Workgroup

Status of Embedded Linux September 2015

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 v3.16 – 3 Aug 2014 – 57 days
- Linux v3.17 – 5 Oct 2014 – 63 days
- Linux v3.18 – 7 Dec 2014 – 63 days
- Linux v3.19 – 8 Feb 2015 – 63 days
- Linux v4.0 – 12 Apr 2015 – 63 days
- Linux v4.1 – 21 Jun 2015 – 70 days
- Linux v4.2 – 30 Aug 2015 – 70 days
- Currently at Linux v4.3-rc1
 - Prediction for v4.3: 8 Nov 2015



CE Workgroup

Linux v3.16

- Power-aware scheduling
- `decode_stacktrace.sh`
 - Converts offsets in a stack trace to filenames and line numbers
- F2FS large volume support



CE Workgroup

Linux v3.17

- Lots of ARM hardware support
 - Newly enabled ARM hardware
 - Rockchip RK3288 SoC
 - Allwinner A23 SoC
 - Allwinner A31 Hummingbird
 - Tegra30 Apalis board support
 - Gumstix Pepper AM335x
 - AM437x TI evaluation board
 - Other ARM boards with existing support also saw improvements with Linux 3.17
- Rework of "config-bisect" mode in ktest



CE Workgroup

Linux v3.18



- OverlayFS introduced
- Size reduction patch:
 - madvise and fadvise syscalls can be configured out
- More LLVM support
- New SOC support:
 - Hisilicon HiP04
 - Amlogic Meson6 (8726MX)
 - Renesas R-Car E2 (R8A77940)
 - Broadcom BCM63xx DSL
 - Atmel SAMA5D4



CE Workgroup

Linux v3.19

- F2FS now has a "fastboot" option
- Device tree overlay support
- Squashfs supports LZ4 compression
- Android "binder" code has been moved from the staging tree





CE Workgroup

Linux v4.0

- This version is not v3.20
 - Linus conducted a survey on Google+
 - 56% of respondents preferred 4.0
- 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 (preview)





- 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 that FS)



CE Workgroup

Things to watch

- Kdbus
 - Has hit some stumbling blocks getting merged
- Kernel tinification!
- RT-preempt (again)
- Persistent memory
- SoC mainlining progress





CE Workgroup

Kernel process improvements

- Kernel merge process is getting better.
- The percent of changes that are accepted after the merge window closes is trending down over time
 - In the 3.0 release, 19% of commits were after the merge window closed
 - In the 4.1 release, 10.5% of commits were after the merge window closed
 - See <https://lwn.net/Articles/650299/>



CE Workgroup

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff

Resources



CE Workgroup

Bootup Time

- F2FS filesystem has a new "fastboot" option
 - Skips some boot-time checks to reduce mount time
 - Sacrifices a little bit of normal performance
 - Due to more synching during normal filesystem operation
- Good talks:
 - ELCE 2014 - *12 Lessons Learnt in Boot Time Reduction* by Andrew Murray
 - ELC 2015 - *Fastboot Tools and Techniques* by John Mehaffey



CE Workgroup

Bootup Time

- Kernel speedups
 - Deferred initcalls (patch still out-of-tree)
 - http://elinux.org/Deferred_Initcalls
 - Kernel tinification project helps
 - smaller size means shorter load times
 - XIP on x86
 - See <https://lwn.net/Articles/637532/>
- User-space speedups
 - Systemd in embedded
 - ELC 2015 - *Tuning systemd for Embedded* by Alison Chaiken



CE Workgroup

Device Tree

- Device Tree is causing delays getting stuff upstream
 - DT maintainers are overloaded
 - Backwards compatibility is a problem
 - See “The Device Tree as a Stable ABI: A Fairy Tale?” – Thomas Petazzoni
- Device Tree Overlays
 - Useful for boards that have daughterboards (e.g. capes or shields) that need DTS changes at boot time.
 - “Transactional Device Tree & Overlays: Making Reconfigurable Hardware Work” - Pantelis Antoniou
 - Also see: <http://lwn.net/Articles/616859/>



CE Workgroup

Device Tree validation

- New work on validating device tree
 - Matt Porter is creating a formal binding document standard (schema for binding docs)
 - Frank Rowand implementing DTS parser (to be used with validator)
 - Tim Bird worked on a binding doc validator
- How it would work:
 - Binding docs are compared with binding schema
 - DTS entries are compared against binding doc and any errors are reported
 - Maybe add to checkpatch.pl
- To be discussed next week at Linaro connect



- Frank Rowand is a new devicetree maintainer
 - Has been updating http://elinux.org/Device_Tree
 - Working on devicetree debugging
 - LCNA 2015 - *Solving Device Tree Issues* by Frank Rowand (see presentations page for link)
- Big session at plumbers this year
 - http://elinux.org/Device_Tree_presentations_papers_articles



Graphics

- Vulkan API from Khronos Group
 - Alternative to Direct3D or OpenGL
 - Intent is to reduce CPU overhead for CPU/GPU operations
 - AMD announced plans to open source the driver (but Intel and Valve already working on a open source driver)
- GPU support
 - Freedreno – for Adreno
 - Nouveau – for Nvidia
 - Lima – for Mali
 - ??? – for PowerVR



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

- PowerVR SGX code leaked in November
- In June: 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>



CE Workgroup

Other OSS GPU drivers

- Nouveau – for Nvidia
 - Nvidia published some GPU details to help open projects write driver (2013)
 - See [https://en.wikipedia.org/wiki/Nouveau_\(software\)](https://en.wikipedia.org/wiki/Nouveau_(software))
 - See also <http://nouveau.freedesktop.org/wiki/>
- Lima – for Mali
 - Seems stalled – recent discussion of putting Mali DRM/KMS code into staging indicated that there needs to be an active user-space (but Lima appears to not be active)



File Systems

- SquashFS supports LZ4 compression
- OverlayFS
 - Support for read/write filesystem over the top of a read-only filesystem
 - Most common use-case is live CDs, but it can be useful for some embedded scenarios
- 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)



CE Workgroup

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



Power Management

- PM domains
 - See “Last One Out, Turn Off The Lights” - Geert Uytterhoeven (at ELC)
 - Good talk showing how to use this with device tree
- Idle and suspend to Idle
 - “The Art of Doing Nothing: Linux Low Power Idle” – Kristen Accardi (at LCJ)
 - “What is Suspend-to-Idle and How to Make It Work” – Rafael Wysocki (at LCJ)
- PowerTop/tuning
 - “Power Tuning Linux: A Case Study” – Alexandra Yates (at LCJ)
 - Was about tuning a laptop distro



CE Workgroup

Real Time

- RT-preempt patch set got a sponsor
 - That's good!
- Still have Xenomai (using Cobalt RT core)
- Good overview of existing RT solutions, and a new alternative
 - ELCE 2014 - *"rtmux: A thin multiplexer to provide hard realtime applications for Linux"*
- by Jim Huang
- Lots of people using PRUs (programmable real-time units)
 - See <http://lwn.net/Articles/639258/>



Security

- IOT raises lots of security issues
- See “Kernel security hacking for the Internet of Things” – Daniel Sangorrin (at LCJ)
 - Reduce attack surface
 - Use variation from pre-determined behavior to detect attacks
 - Isolate critical software
- Security module stacking
 - Added in kernel 4.2
 - See <https://lwn.net/Articles/635771/>



System Size

- Size project keeps nibbling away at items
- 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
 - Ability to make any command-line option static
 - Example for `initcall_debug` = saves 385 bytes
 - A lot of the savings are due to GCC constant folding
- Intel X86 XIP patches
 - See <https://lwn.net/Articles/637532/>



System Size (cont.)





- Some recent talks:
 - Optimize uClinux for ARM Cortex-M4 – Jim Huang (at ELC)
 - Target = STM32F4xx
 - Linux for Microcontrollers: From Marginal to Mainstream – Vitaly Wool (at ELC)
 - Target = STM32F2x
 - 840K .text, 132k .rodata, 86k .data (BT, no TCP/IP)
 - Pushing the limits of Linux on ARM – Andreas Färber (at LCJ)
 - Target = STM32F429



CE Workgroup

Testing

- Kselftest
- LTSI Test Project
- Kernelci.org





CE Workgroup

kselftest

- Inside kernel source tree
 - Makefile target: 'make kselftest'
- Ability to install tests mainlined in kernel v4.1
 - Cross-build should be worked on
 - <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)
- Available now
 - <https://bitbucket.org/cogentembedded/jta-public/>
- Several companies provided feedback at LTSI workshop meeting in Tokyo
 - CogentEmbedded will fix issues in next few months
- Please use JTA
 - Please send feedback to LTSI mailing list
 - <https://lists.linuxfoundation.org/mailman/listinfo/ltsi-dev>



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 2015 - *Upstream Kernel Testing* – by Kevin Hilman
- Sony Mobile has a phone in this farm



CE Workgroup

Tracing





- eBPF to be used for dynamic tracing
 - Ktap will not be merged (frowny-face)
- 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

- Greybus
- J2
- Next LTS kernel version:
 - 4.1







CE Workgroup

Greybus

- New fast bus for mobile device hotplugging
 - For project ARA (Google's modular phone)
 - Being worked on by Greg Kroah-Hartman
- <https://lwn.net/Articles/648400/>
- Work still needed in Android for support of dynamic hotplugging





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
- See the next presentation by Rob Landley



CE Workgroup

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff





Resources



CE Workgroup

CEWG Projects

- Contract work
- Projects and initiatives
- (conferences covered later)









CE Workgroup

CEWG Contract Work

- Kernel string refactoring
- Device tree documentation
- LTSI test framework





CE Workgroup

Kernel string refactoring

- Description
 - Refactor kernel strings to enable compiler optimizations which 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
 - Expect at least 20K of savings, depending on kernel config
- Project is just starting



CE Workgroup

DT documentation





- Are proceeding with “guide” documentation
- Frank Rowand has been collecting data and giving talks
 - LinuxCon NA, ELCE, ELC and LCJ
- Goal is to release by ELC Europe 2015
- Will be put on elinux wiki at:
 - http://elinux.org/Linux_Drivers_Device_Tree_Guide



CE Workgroup

LTSI test framework

- (Discussed previously)









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 social infrastructure systems

- **Status**

- **Recent Activity**

- BOFS at ELCE 2014 and ELC2015 and LCJ2015
- Private meetings to discuss goals with interested companies

- Working to define requirements in areas of functional safety and maintenance longevity

- **Next steps:**

- Hold additional meetings to define requirements



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



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 September 2014
 - SIG meetings at ELCE and ELC
 - Presentations about overcoming obstacles at ELCE 2014, ELC 2015, and LCJ 2015
 - 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
 - Working to identify problem areas
- Published tools:
 - <https://github.com/tbird20d/upstream-analysis-tools>





CE Workgroup

Big problem areas

| Area | Insertions range |
|----------|------------------|
| Mach-msm | 347K – 417K |
| Media | 120K – 360K |
| Video | 37K – 346K |
| Wireless | 80K – 250K |
| Sound | 74K – 240K |
| Input | 51K – 238K |
| Camera | 50K – 210K |
| GPU | 36K – 172K |
| Power | 44K – 94K |





- **Technical Projects:**
 - USB OTG charger integration
 - Broadcom wireless driver
 - http://elinux.org/Kernel_areas_of_focus_for_mainlining
- **Non-technical:**
 - Easy patch submission tool (no special mail settings required)
- **Engage with more companies and individuals**
 - Targets: Google and MediaTek
 - Meetings at Linaro Connect, ELC and kernel summit



CE Workgroup

Long Term Support Initiative

- LTSI 3.14 is latest kernel
- Many presentations available on status
- Latest project push is testing facility
 - See previous page on JTA test framework
- Considering multiple merge windows
- Will base next LTSI on 4.1 (LTS)





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

eLinux wiki

- I'm looking for people to take over specific sections (portals)
 - Boot time
 - Power management
 - Security
 - Etc.
- Main job is just to link to talks and articles on that topic
 - Should be 1 or 2 hours a week
- Contact me if you are interested



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





CE Workgroup

Projects and Consortia

- PRPL Foundation – Multi-company MIPS non-profit
 - Projects: PRPL OpenWRT, MIPS QEMU
- DroneCode – Open source UAV software group
 - <http://www.dronecode.org/>
 - First Linux Drone summit at ELC
 - Andrew Tridgell on ELC program committee



CE Workgroup

Distros

- Android
 - Getting ready for Marshmallow (6.0) release
 - App permissions
 - Fingerprint
 - Mobile payments
 - “Doze” mode to decrease power usage in standby
- Tizen
 - Lots of security work
- AGL
 - Announced it will do it’s own distro
- CEWG Shared embedded distribution
 - (see previous slides)






CE Workgroup

Build Systems

- OpenEmbedded/Yocto Project
- Buildroot

(nothing new to discuss this time)





CE Workgroup

Events

- LinuxCon Japan
 - June 3-5, 2015 – Tokyo, Japan
 - Slides at:
- Embedded Linux Conference Europe 2015
 - October 5-7, 2015 - Dublin, Ireland
- Embedded Linux Conference 2016
 - April 4-6, 2016 - San Diego



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_4.?.
- eLinux wiki - <http://elinux.org/>
 - Especially <http://elinux.org/Events> for slides
- Celinux-dev mailing list



CE Workgroup



Thanks!

