

CE Workgroup

Status of Embedded Linux October 2015

Tim Bird

Architecture Group Chair

LF CE Workgroup




Outline

Kernel Versions
Technology Areas
CE Workgroup Projects
Other Stuff
Resources



Outline







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
- Linux v4.3-rc4 (as of yesterday)
 - Prediction for 4.3 release: 8 Nov 2015



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







- 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




- 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



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







Linux v4.0

- This version is not v3.20
 - Linus conducted a survey on Google+
 - 56% of respondents preferred 4.0
 - The name of this kernel is "hurr durr I'ma sheep"
- 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




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




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



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 Ext3 filesystems








Things to watch

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








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/



Outline





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
- XIP on x86
 - See https://lwn.net/Articles/637532/
- Deferred initcalls (patch still out-of-tree)
 - http://elinux.org/Deferred_Initcalls



Bootup Time (cont.)

- Kernel tinification project helps
 - Smaller size means shorter load times
- User-space speedups
 - Systemd in embedded
 - ELC 2015 Tuning systemd for Embedded by Alison Chaiken
- Some good 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 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/



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 working 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 or kernel build
- V2 of spec has been published still hashing out details



More devicetree stuff

- Frank Rowand is a new devicetree maintainer
 - Has been updating http://elinux.org/Device_Tree
 - Working on devicetree debugging
 - LCNA 2015 (and here) Solving Device Tree Issues by Frank Rowand
- Big DT session at plumbers this year
 - http://elinux.org/Device_Tree_presentations_pa pers_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 it)
- GPU support
 - Freedreno for Adreno
 - ??? for PowerVR
 - Etnaviv for Vivante
 - Nouveau for Nvidia
 - Lima for Mali



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/



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



Other OSS GPU drivers

- Etnaviv for Vivante
 - See http://www.x.org/wiki/Events/XDC2015/Program/ Stach_etnaviv.pdf
 - Replaced 65K kernel driver with 6.5K driver
- Nouveau for Nvidia
 - Nvidia published some GPU details to help open projects write driver (2013)
 - See 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)



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



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



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



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
- Latest release of RT-preempt is for 4.1 kernel
 - Tends to follow LTS releases



Real Time - other

- Xenomai 3.0 is at rc7 release
 - Uses Cobalt RT core
 - 3.0 supports both dual-kernel and single-kernel configurations (using RT-preempt)
 - See xenomai.org
- 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
 - Can detect attacks by detecting variation from pre-determined behavior
 - 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.)


- Nicolas Pitre has done work recently on supporting gcc --gc-sections
 - Lighter-weight option similar to LTO
- Some recent talks:
 - Optimize uClinux for ARM Cortex-M4 Jim Huang (at ELC)
 - Linux for Microcontrollers: From Marginal to Mainstream – Vitaly Wool (at ELC)
 - 840K .text, 132k .rodata, 86k .data (BT, no TCP/IP)
 - Pushing the limits of Linux on ARM Andreas Färber (at LCJ)




Testing

- Kselftest
- LTSI Test Project
- Kernelci.org









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/



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
- Please use JTA
 - Please send feedback to LTSI mailing list
 - https://lists.linuxfoundation.org/mailman/listinfo/ltsidev



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 (also here) Upstream Kernel Testing – by Kevin Hilman
- Sony Mobile has a phone in this farm



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 (non-GNU)
 - Call it LinuxNG?








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)




Miscellaneous

- Greybus
- J2
- Next LTS kernel version:
 - 4.1
- Weird IOT news









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






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




Weird IOT news

 Microsoft released Windows 10 IoT kit for Raspberry PI










Outline






CEWG Projects

- Contract work
- Projects and initiatives












CEWG Contract Work

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









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 just starting



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






LTSI test framework

(Discussed previously)











Projects and initiatives

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








Civil Infrastructure

- Goals
 - Solve problems with Linux for use in civil 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



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



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
 - Working to identify problem areas
- Published tools:
 - https://github.com/tbird20d/upstream-analysistools



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
 - Recently had conversations with Google and MediaTek



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)





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 lowcost development boards
- Please use and add to site




Outline







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 Interconnect Consortium
 - lotivity is the technology
- DroneCode Open source UAV software
 - http://www.dronecode.org/



Projects and Consortia

- Linaro
 - Just celebrated 5th anniversary
 - Linaro IoT and Embedded initiative (LITE)
 - Run Linux on Cortex A and mbedOS on Cortex M
 - Unsure about licensing for Cortex M
- PRPL Foundation (Multi-company MIPS non-profit)
 - Announced at ELCE 2014
 - Projects: PRPL OpenWRT, MIPS QEMU
 - OpenWRT summit tomorrow, here



Distros

- Android
 - Just released "M" version
 - New build system under development, using 'go' language and something called blueprints
- Tizen
 - Lots of security work
- AGL
 - Announced it will do it's own distro
- CEWG Shared embedded distribution
 - (see previous slides)



Build Systems

- OpenEmbedded/Yocto Project
 - 1.8 released
 - Can now do builds and runs with Toaster (web interface)
- Buildroot
 - Configurable support for static linking
 - Improved support for package hashes
 - Better warnings about toolchain header safety issues



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
- Embedded Linux Conference Europe 2016
 - October 6-7, 2016 Berlin, Germany



Hardware

- Intel and Micron 3D Xpoint memory
 - Non-volatile
 - Read/Write, Random access, Faster than NAND, Cheaper than flash
 - Not many details yet
- Is this the persistent memory we've been waiting for?



Outline


Kernel Versions
Technology Areas
CE Workgroup Projects
Other Stuff

Resources Final Impressions



Resources

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



Impressions – hardware price

- Steady decline in price of silicon
 - Cheapest Android phone = \$29
 - Lenovo A288t (Russian)
 - The Chip \$9 computer board
 - Estimate that cheapest Linux-capable SoC (with MMU) currently about \$3
- Still want to see Linux on cereal boxes
 - Less than \$1 for SoC, display, battery, input



Impressions – markets

- I worry that Linux will not be in the "things" part of IOT
 - Linux on IOT gateway is a no-brainer
 - Linux is too big for sensors
 - Rate of adoption of tinification patches is slow
 - Need a concerted, collaborative effort here
- In other areas Linux is already penetrating:
 - Drones, Industrial automation, Robotics
 - Automotive, Automated vehicles
 - Gateways, Civil infrastructure



Impressions

Embedded Linux is doing fine....

