



CE Workgroup

Status of Embedded Linux September 2016

Tim Bird

Architecture Group Chair

LF Core Embedded Linux Project



CE Workgroup

Outline

Kernel Versions
Technology Areas
CE Workgroup Projects
Other Stuff
Resources



CE Workgroup

Outline

Kernel Versions
Special: Areas
Kernel Summit Discussion
Review
Other Stuff
Resources



CE Workgroup

Ksummit-discuss issues

- Some of the most interesting discussions occur on ksummit-discuss mailing list
- A few big areas of discussion:
 - Stable workflow
 - Kernel testing
 - Git-series
 - GPL defense issues
 - Backport trees (and SoC mainlining issues)



CE Workgroup

KDSL: Stable workflow

- Some vendors don't trust the stable trees, and developers want to improve the quality of them
 - Some patches in stable caused regressions
 - Some users cherry-pick patches instead of basing product on whole tree
- Suggestions to improve quality:
 - Have more review of stable patches
 - But sub-system maintainers are already overworked
 - Have more testing of stable patches
 - This became a long thread about kernel testing
 - Identify commit that patch fixes
 - To allow for easier back-porting to multiple stable trees



KSDL: Kernel testing

- Would like to see new kernel code submitted with unit tests
 - Recommendation only – not required
- Kselftests is a good start, but needs more features
 - Examples, documentation, test interface, standard logging and reporting
- Need to declare test requirements
 - Test should return “not supported” instead of failure, if hardware is missing or requirements are not met
- Important to measure actual bugs caught



CE Workgroup

KSDL: Git-series

- Lots of discussions (in different threads) about how to use git to manage patches
- Josh Triplett introduced git-series, to maintain a set of patches in git
 - Similarities to quilt or stgit
 - Maintain patches as first-class objects, with meta-data and versioning for each patch



CE Workgroup

KSDL: GPL defense issues

- Big thread about when to enforce the GPL with lawsuit
 - See <http://lwn.net/Articles/698452/>
- Linus and Greg KH very reluctant to use lawsuits
 - It drives companies away instead of helping them join the community
- Busybox cited as example of legal enforcement that damaged the project
- Lots of disagreement on some issues



GPL defense issues (cont.)

- When to sue?
 - Can the threat make companies behave better?
 - If you never sue, is there any threat?
- What is objective?
 - Getting useful code (good for developers)
 - Bad actors don't have good code
 - Allowing end-users to get code for their devices (good for users)
 - End users don't have copyrights, and have no standing to sue
- Who can decide to sue?
 - Small group of copyright holders, or the SFC, or an individual?
 - Problem is if other stakeholders don't want to sue



CE Workgroup

KSDL: Backport trees (and SoC mainlining)



- Long thread about value of trees with backported features
 - LSTI and LSK (Linaro Stable Kernel)
- Some people don't like these trees
 - Reduces incentive to work with upstream
 - Loses shared testing with upstream
- Industry reality is that backporting is inevitable
 - Mobile vendors are not on upstream due to product cycle issues and legacy out-of-tree code
- Backport trees try to find balance in quality between back-porting to in-house tree and forward-porting patches to upstream



CE Workgroup

Ksummit-discuss issues

- And much more...
 - See <https://lists.linuxfoundation.org/pipermail/ksummit-discuss/>
- Many topics were put onto the kernel summit agenda
- Kernel summit planned for end of October, in Santa Fe, New Mexico





CE Workgroup

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff

Resources



CE Workgroup

Kernel Versions



- Linux v4.3 – 1 Nov 2015 – 63 days
- Linux v4.4 – 10 Jan 2016 – 70 days
- Linux v4.5 – 13 Mar 2016 – 63 days
- Linux v4.6 – 15 May 2016 – 63 days
- Linux v4.7 – 24 July 2016 – 70 days
 - By the way, my prediction was correct
- Linux v4.8-rc5
 - I predict 4.8 on 25 Sep 2016 (63 days)
- Greg KH already announced 4.9 as next LTS



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







CE Workgroup

Linux 4.6

- GPIO subsystem rework
- scripts/dtc/dtx_diff
 - Compare device trees in a number of formats
- Improved page-poisoning
 - Separate from debug, can set poison value to 0 (to clear pages after free for security reasons)





CE Workgroup

Linux 4.7

- Schedutil frequency governor
 - See <http://lwn.net/Articles/682391/>
- VFS layer can iterate through directories in parallel
- Ability to attach BPF programs to tracepoints
- Ftrace histogram triggers
 - Can tell tracer to accumulate events into buckets and give results, via the sysfs interface
- Android sync_file feature moved from staging



CE Workgroup

Sync_file

- Allows for explicit fencing for buffers by userspace
- How it works:
 - Producer driver sends the fence related to the buffer to userspace via a `sync_file`
 - An intermediary (e.g. a compositor) passes these fenced fds to DRM in an atomic commit
 - Consumer will not use the buffer for anything before the fence(s) signals
- This avoids a lot of waiting
- See [Documentation/sync_file.txt](#)



CE Workgroup

Linux 4.8 (predictions)

- New kernel documentation system
- New pseudo-random number generator
 - See <https://lwn.net/Articles/686033/>
- ARM64 support for kexec and kprobes
- New timer wheel implementation
 - <https://lwn.net/Articles/646950/>
 - Better performance:
 - No more cascade operations
 - Quick determination of next timeout
 - Long timeouts have reduced resolution
 - Automatically coalesces longer timeouts



CE Workgroup

Observations




- No embedded-specific features lately
 - Nothing for:
 - Boot-up time
 - System size
 - Embedded filesystems (are these done now?)
 - Embedded security
 - Realtime (well, maybe the timer wheel stuff)
 - Power management
- There is lots of processor support, and lots of device drivers for embedded
- I worry that people are using non-Linux
 - More on this later



CE Workgroup

Things to watch (from past)

- Kernel tinification!
- RT-preempt
- Persistent memory
 - (NVM = Non-Volatile Memory)
- SoC mainlining progress





CE Workgroup

Things to watch (status)

- Kernel tinification! (stalled)
- RT-preempt (only 10K lines left!)
- Persistent memory (in progress)
 - Good talk on about issues:
 - “Making use of persistent memory”
 - <http://lwn.net/Articles/674752/>
- SoC mainlining progress (stalled)



CE Workgroup

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff

Resources



Bootup Time

- Mostly old news...
- XIP on x86
 - See <https://lwn.net/Articles/637532/>
- Asynchronous probing
 - Discussed at last 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
- Device dependencies
 - May be discussed at next kernel summit



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
 - 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
- 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 specification is in progress
 - Want to update material and make it more available



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



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

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)





CE Workgroup




(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.)





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
 - Tends to follow LTS releases
 - See <https://www.kernel.org/pub/linux/kernel/projects/rt/>



Real Time - other

- Xenomai 3.0.1
 - Uses Cobalt RT core
 - 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



CE Workgroup

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
- New project for kernel security issues:
 - http://kernsec.org/wiki/index.php/Kernel_Self_Protection_Project



CE Workgroup

Security and IOT



- IOT raises lots of security issues
- See “Securing Embedded Linux” – Mike Anderson at ELC 2016
 - Who are attackers?
 - Secure boot techniques
 - Encryption
 - Physical security
 - Data Security
 - Network Security





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
- Fuego - 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

Fuego - LTSI test project

- Available now
 - <https://bitbucket.org/tbird20d/fuego/>
- Lots of work recently on wiki (documentation)
- Working on lots of issues in parallel
- Should be ready for demonstrations at ELCE (October)



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
 - 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)



CE Workgroup

Miscellaneous

- Next LTS kernel version:
 - 4.9
 - This is the earliest it's been announced!
 - We haven't even opened the 4.9 merge window
- Non-Linux announcements





Lots of non-Linux in IOT

- Zephyr – RTOS from Wind River
 - Apache 2 license
 - Minimal size – as small as 8K
 - Highly configurable
 - NoMMU
 - Networking: WiFi, Bluetooth, NFC
- Magenta – RTOS by Google
 - Fuchsia OS - Some attributes of Android
 - Based on LK
 - BSD license



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 provided at LinuxCon Japan



CE Workgroup

DT documentation

- Finished “guide” documentation
- Frank Rowand has been collecting data and giving talks
 - LinuxCon NA, ELCE, ELC and LCJ
- Is 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

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





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 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)



CE Workgroup

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:
 - **Create tools for easier mainlining**
 - **Patch submission tool**



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

- Distros and Build Systems
- Events





CE Workgroup

Distros

- Android
 - “Nougat” version released in August
 - 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)



CE Workgroup

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 2016
 - October 6-7, 2016 - Berlin, Germany
- Embedded Linux Conference
 - February 21-23, 2017 – Portland, Oregon





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!

