



# CE Workgroup

# Status of Embedded Linux

June 2013

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 v3.5 – 21 July 2012 – 62 days
- Linux v3.6 – 30 Sep 2012 – 71 days
- Linux v3.7 – 10 Dec 2012 – 71 days
- Linux v3.8 – 18 Feb 2013 – 70 days
- Linux v3.9 – 28 Apr 2013 – 69 days
  - I predicted April 30 (only 2 days off)
- Linux v3.10-rc4 – 2 June 2013
  - I predict...



CE Workgroup

# Kernel Versions

- Linux v3.5 – 21 July 2012 – 62 days
- Linux v3.6 – 30 Sep 2012 – 71 days
- Linux v3.7 – 10 Dec 2012 – 71 days
- Linux v3.8 – 18 Feb 2013 – 70 days
- Linux v3.9 – 28 Apr 2013 – 69 days
  - I predicted April 30 (only 2 days off)
- Linux v3.10-rc4 – 2 June 2013
  - I predict... July 7, 2013



CE Workgroup

# Linux v3.5

- Kernel log rework
  - Structured printk (new format), with tags
  - <http://lwn.net/Articles/492125/>
- Support for writing NFC drivers
- Integration of ramoops and pstore
  - Part of work to support Android ram\_console
- Uprobes
  - User-space probes
  - <https://lwn.net/Articles/499190/>
- Autosleep



CE Workgroup

# Linux v3.6

- Android RAM console functionality integrated into pstore
- CANFD support for CAN protocol
  - CAN with flexible data rate
- LED oneshot mode
  - Sysfs interface for certain one-time LED/gpio manipulations
- "Suspend to Both"
  - Create resume image both in RAM and on disk
  - If power dies during suspend, disk image can be used to resume



CE Workgroup

# Linux v3.7

- ARM multi-platform support
  - See <http://lwn.net/Articles/496400/>
- ARM 64-bit support (Aarch64)
- Cryptographically signed kernel modules
  - See <https://lwn.net/Articles/470906/>
- Perf trace (alternative to strace)
  - Allows intermingling kernel trace events with `syscall` events
- Runtime power management for audio
- Kernaldoc system can output in HTML5 format





CE Workgroup

# Linux v3.8

- F2FS – flash-friendly file system
  - See <https://lwn.net/Articles/518988/>
- New thermal governor subsystem
- Memory control group support for accounting for kernel memory usage
  - Stack and slab accounting and limits
- Cpuidle support for big.LITTLE



CE Workgroup

# Linux v3.9



- Ftrace snapshots
  - Grab a snapshot of a running trace without stopping
- KVM virtualization for Cortex A15 processors
- PowerPC support for transactional memory
- CONFIG\_EXPERIMENTAL=y
  - And should be gone soon
- 'make menuconfig' now has "save" and "load" buttons



CE Workgroup

# Linux v3.9 (cont.)

- Descriptor-based GPIO
  - Access GPIOs by descriptor (e.g. by name in addition to by number)
  - Allows for grouping GPIOs - for “atomic” operations
    - Possibly useful for handling realtime issues
  - See <http://lwn.net/Articles/533632/>





CE Workgroup

# Linux v3.10

- Full tickless (more later)
- Single zImage for ARM
  - Lots more platforms support multi-platform kernels
  - Arnd Bergmann shooting for almost-complete coverage by v3.12
- Multi-cluster power management
  - Partial support for big.LITTLE PM
  - <https://lwn.net/Articles/539082/>



CE Workgroup

# Linux v3.10 (cont.)

- Multiple ftrace buffers
- Memory pressure control group support
  - Allows for notification if memory gets low
  - <http://lwn.net/Articles/531077/>



# Full tickless

- Full tickless (full dynamic tick)
  - Under some circumstance, some processors may run with no periodic ticks at all
  - Previous CONFIG\_NO\_HZ used dynamic tick, but only when CPU was idle
  - New option is tri-state: periodic, idle, full
  - Boot CPU cannot be 'full' tickless
  - CPU cannot be full tickless with more than one process
  - <https://lwn.net/Articles/549580/>



CE Workgroup

# Things to watch

- Android features
  - Volatile ranges
  - ARM FIQ -> KDB glue
- big.LITTLE MP scheduling
  - See <https://lwn.net/Articles/501501/>
  - See the In-Kernel-Switcher work
    - <https://lwn.net/Articles/549473/>
- Single zImage support on ARM
- Support for transactional memory instructions
  - Could be as big a deal (eventually) as locking primitives



# Things to watch (longer-term)

- Non-volatile mass memory
  - Interesting remarks by Linus in LinuxCon keynote question and answer
  - Skeptical it will happen “this year” – it’s always within a few years of happening.
  - Won’t change a lot of kernel algorithms
  - Will mostly change filesystems
    - Byte-addressable storage has big implications for long-term storage
  - Applications will still segregate data between persistent and non-persistent groups
  - Things take longer to change than people think





CE Workgroup

# Outline

Kernel Versions

**Technology Areas**

CE Workgroup Projects

Other Stuff





Resources



CE Workgroup

# Bootup Time

- Not much new stuff this time









CE Workgroup

# Graphics

- Not much new stuff this time





# File Systems

- F2FS – Samsung Flash-friendly filesystem
  - Mainlined in Linux version 3.8
  - Log-structured, with lots of tweaks
    - E.g. hot vs. cold data separation
  - See <http://elinux.org/F2FS>
- CE WG project to analyze filesystem performance on eMMC
  - More on next slides



CE Workgroup

# Flash Filesystem tuning

- CE Workgroup contracted with Cogent Embedded
- Goal: to test different block-based filesystems on flash media
  - Specifically, want to measure the effect of different kernel tuning options (IO scheduler, flash geometry vs. flash part attributes and workload characteristics)
- Result document is NOW available at:
  - [http://elinux.org/File\\_Systems#Comparison\\_of\\_flash\\_filesystems](http://elinux.org/File_Systems#Comparison_of_flash_filesystems)



# Power Management

- Autosleep
  - Wakelock-compatible solution by Rafael Wysocki
    - <http://lwn.net/Articles/479841/>
  - Rafael: *“This series tests the theory that the easiest way to sell a once rejected feature is to advertise it under a different name”*
  - Mainlined in v3.5
- Power-aware scheduling:
  - <http://lwn.net/Articles/512487>



# big.LITTLE scheduling





- Lots of work recently on big.LITTLE scheduling
  - Multi-cluster power scheduling
  - In-kernel-switcher work
  - See talk at LCJ by Nakagawa-san of Renesas
    - One User Space Approach to big.LITTLE MP System on Real Silicon
- Still waiting for real-product results ??



CE Workgroup

# System Size

- Kernel size
- Library size







# Kernel size

- Cooperative memory relinquishment
  - Volatile Ranges
  - Lexmark work (membroker and ANR malloc)
    - See talk at ELC 2013 – "SystemWide Memory Management without Swap"
- Tim Bird's presentation on advanced size optimization of the kernel
  - Notes and slides available at:  
[http://elinux.org/System\\_Size\\_Auto-Reduction](http://elinux.org/System_Size_Auto-Reduction)
  - (more later)



# Library reduction

- olibc – bionic libc
  - Has good features from Android, and is smaller and more configurable than glibc

glibc 2.11 : /lib/libc.so	→ 1,208,224 bytes
uClibc 0.9.30 : /lib/libuClibc.so	→ 424,235 bytes
bionic 2.1 : /system/lib/libc.so	→ 243,948 bytes

- See ELC 2013 talk by Jim Huang
- Kconfig for egllibc
  - Ability to configure parts of libc to use

libc-2.17.so reduced from	1.2M -> 830K
ld-2.17.so reduced from	128K -> 120K
libm-2.17.so reduced from	610K -> 580K

- See ELC 2013 talk by Khem Raj



CE Workgroup

# Advanced Size Optimization of the Linux Kernel





- “Auto-reduce” project
- Find automated ways to reduce the kernel
  - Link-time optimization
  - System call elimination
  - Kernel command-line argument elimination
  - Kernel constraint system
- Additional research
  - Link-time re-writing
  - Cold-code compression



CE Workgroup

# Security

- SMACK
- SE-Linux





CE Workgroup

# SMACK

- SMACK for Tizen
  - Simplified rule set (3 tiers, 40,000 rules)
- See <http://lwn.net/Articles/55278>
  - Published June 5, requires LWN.net subscription



CE Workgroup

# SE-Linux

- SE-Android
  - Implementation of SE-Linux for Android systems
- SE-Linux was previously too big for embedded
  - Early embedded SE-Linux required 2M
  - Desktop SE ruleset is 900,000 rules
- However, SE-Android only has 1658 rules and 263 types (71K policy size)
- <http://selinuxproject.org/page/SEAndroid>
  - Especially:  
[http://www.internetsociety.org/sites/default/files/Presentation02\\_4.pdf](http://www.internetsociety.org/sites/default/files/Presentation02_4.pdf)



CE Workgroup

# Tracing

- Ktap
  - Dynamic tracing, without the overhead of compiling into a module
  - Adds an interpreter to the kernel
  - Single module, that leverages ftrace, kprobes, etc.
  - Prints results in ASCII
  - Good session in LinuxCon Japan by Jovi Zhang



CE Workgroup

# Outline

Kernel Versions

Technology Areas

**CE Workgroup Projects**

Other Stuff

Resources





CE Workgroup

# CEWG Contract Work 2013

- Dynamic memory analysis (2012)
- eMMC tuning guide completed




# Dynamic memory analysis

- Description:
  - Instrument and collect data on kernel dynamic memory allocations
  - Make recommendations for areas where dynamic kernel memory usage could be reduced
- Contractor: Ezequiel Garcia
- Status:
  - Use existing kmem\_events (ftrace) infrastructure
  - Some patches already accepted upstream
  - New tool for visualization of kernel memory usage
  - See [http://elinux.org/Kernel\\_dynamic\\_memory\\_analysis](http://elinux.org/Kernel_dynamic_memory_analysis)
  - See ELC 2013 talk




# Drivers kmalloc





CE

# Linux kmalloc





CE Workgroup

# eMMC tuning guide





- Description:
  - This project analysed EXT4, BTRFS and F2FS on a variety of block-based flash parts on a few different development boards
  - Output is a document describing best practices for tuning Linux block-based filesystems for block-based flash filesystems
  - Also, methods and scripts for filesystem testing
- Contractor: Cogent Embedded
- Status: Just completed
  - Document at:  
[http://elinux.org/File\\_Systems#Comparison\\_of\\_flash\\_filesystems](http://elinux.org/File_Systems#Comparison_of_flash_filesystems)



CE Workgroup

# Other Projects

- Long Term Support Initiative (LTSI)







CE Workgroup

# Long Term Support Kernel for Industry

- LTSI 3.4 is available now
- Held a workshop at LinuxCon Japan
  - Discussed testing phase of project
  - Discussed promotion of project
- Program for free hardware for LTSI kernel testing





# Hardware

- Rise of cheap hardware
  - Lots of < \$200 boards
  - Raspberry Pi - \$35
  - New BeagleBone - ?? (<\$79)
- Lots of people have mobile phones or tablets
- No need for CE WG hardware program
- Anyone can learn embedded Linux
  - FYI – [code.org](http://code.org) – new site to teach programming





CE Workgroup

# Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

**Other Stuff**





Resources



CE Workgroup

# Other Stuff

- Tools
- Build Systems
- Events
- Miscellaneous





# Tools

- Cortex
  - Coredump filter
  - Generates sparse coredump
  - See ELC 2013 presentation by Tristan Lelong
    - "Debugging for production systems"
- Debugging techniques
  - Good overview by Kevin Dankwardt at ELC 2013
    - "Survey of Linux Kernel Debugging Techniques"



# Testing frameworks




- Autotest
  - Simple framework
  - Not cross-compiler aware?
- LAVA
  - Linaro test framework
- "Kernel Testing Tools and Techniques" BOF by Matt Porter at ELC 2013
- CE workgroup considering reviving test activity for 2013
  - Need input...



CE Workgroup

# Build Systems

- Yocto project
  - Lots of talks at ELC (and previous ELCs)
    - Sean Hudson – good introduction tutorial
    - Saul Wold – current status
  - Tutorials now online
- Buildroot still hanging in there





# Distributions

- Tizen – may be a serious competitor in embedded distros
  - Needs to open up a bit more (but it looks like it's happening)
  - Replacing Bada at Samsung
  - Shipping in phones??
- Android use in non-CE embedded
  - Headless android
- Yocto Project = the new in-house distro
- Angstrom
  - Very common on development boards



CE Workgroup

# Events

- LinuxCon Japan – May 29-31 2013
- Japan Jamborees
- LinuxCon US
  - September 2012 – New Orleans
- Embedded Linux Conference Europe 2013
  - October 21-23, 2013 – Edinburgh, Scotland
- Embedded Linux Conference 2014
  - April, 2013 – San Jose



CE Workgroup

# eLinux wiki



- <http://elinux.org>
  - Web site dedicated to information for embedded Linux developers
    - The wikipedia of embedded linux!
- Hundreds of page covering numerous topic areas: bootup time, realtime, security, power management, flash filesystem, toolchain, editors
- Working on new wiki projects:
  - Video transcription project
  - Tech Zones





# Miscellaneous

- Status of industry = Healthy
  - Over 1.4 billion devices shipped with embedded Linux
- Just had CELF 10<sup>th</sup> anniversary party
- Still going strong
- Personally, I'm excited about new job





CE Workgroup

Kernel Versions  
Technology Areas  
CE Workgroup Projects  
Other Stuff  
**Resources**





# Resources

- LWN.net
  - <http://lwn.net/>
  - If you are not subscribed, please do so
- Kernel Newbies
  - [http://kernelnewbies.org/Linux\\_3.?.](http://kernelnewbies.org/Linux_3.?.)
- eLinux wiki - <http://elinux.org/>
  - Especially <http://elinux.org/Events> for slides
- Celinux-dev mailing list
- LinuxCon Japan slides
  - <http://events.linuxfoundation.org/events/linuxcon-japan/program/presentations>



CE Workgroup



**Thanks!**

