



CE Workgroup

Status of Embedded Linux

March 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 v3.3 – 18 Mar 2012 – 74 days
- Linux v3.4 – 20 May 2012 – 63 days
- 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-rc1 – 3 Mar 2013
 - Predict Linux v3.9 on April 30



CE Workgroup

Linux v3.3

- ARM large physical address extensions
 - See Catalin Marinas talk at ELC Europe 2011
- ALSA support for compressed audio
- New “charger manager” subsystem
 - Can partially resume to poll battery and re-suspend
- Android patches in staging
 - This is really cool



CE Workgroup

Linux v3.4

- Universal Flash Storage host controller drivers
 - See Documentation/scsi/ufs.txt
- Common clock framework
 - Unifies handling of subsystem clocks
 - See Documentation/clk.txt
- HSI (High-speed synchronous serial interface) framework
 - Used for communication between CPU and cellular modem engines



CE Workgroup

Linux v3.4 (continued)

- DMA buffer sharing API
- Remoteproc subsystem
 - Allows for control of other CPUs through shared memory
 - Rpmmsg is a new mechanism for communicating with other CPUs (running non-Linux)
 - See Documentation/remoteproc.txt and rpmmsg.txt



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
- Kerneldoc 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 (probable)

- Ftrace snapshots
 - Grab a snapshot of a running trace without stopping
- PowerPC support for transactional memory
- `CONFIG_EXPERIMENTAL=y`
 - And should be gone soon








CE Workgroup

Things to watch

- Android features
 - Volatile ranges
 - ARM FIQ -> KDB glue
- big.LITTLE
- SOC support for ARM (refactoring)





CE Workgroup

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff

Resources



CE Workgroup

Bootup Time

- Systemd in embedded
 - Systemd starts services and daemons on-demand
 - Angstrom uses systemd
 - <http://www.mattlmassey.com/2012/07/10/explorations-into-angstrom-syslog-and-systemd/>
 - People either love systemd or hate it, or both
 - Would be nice to get some boot time and size numbers to evaluate this



CE Workgroup

Graphics

- Ubuntu announces MIR
 - New display server that leverages Android GPU drivers
 - Replaces X, across multiple form factors (desktop, tablet, phone)
 - See <http://www.h-online.com/open/news/item/Canonical-reveals-plans-to-launch-Mir-display-server-Update-1815982.html>



CE Workgroup

Graphics (cont.)

- Lots of work around memory management between kernel, user-space and GPU
- Android has /dev/ion
 - A unified approach to buffer management and sharing between display, GPU, camera, codecs, etc, new in Ice Cream Sandwich
 - Replacement for pmem
- Mainline has Contiguous Memory Allocator (CMA) and dma-buf
 - <http://lwn.net/Articles/468044/> - CMA
 - <http://lwn.net/Articles/470339/> - dma-buf



CE Workgroup

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 about this later



CE Workgroup

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>



CE Workgroup

System Size

- Kernel size
 - Ezequiel Garcia's trace_analyze for kernel memory analysis (showed previously)
 - See talk at ELC 2013
 - Link-Time Optimization (LTO)
 - Hopefully showing up in mainline soon
- Cooperative memory relinquishment
 - Volatile Ranges
 - Lexmark work (membroker and ANR malloc)
 - See talk at ELC 2013 – "SystemWide Memory Management without Swap"



System Size (cont.)

- 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



Link Time Optimization

- See <http://lwn.net/Articles/512548/>
- Newer gcc (4.7) supports adding extra meta-data about routines (gimple) at compile time
- Linker can now do whole-program optimization at link time
- Andi Kleen has 74 patches that add support to the Linux kernel for LTO feature
 - Mark functions as 'visible' to avoid dead-code elimination
 - Adjust compilation flags to be consistent
 - Add dependencies to avoid conflicts for features which can't conform to LTO requirements (ftrace)



gcc Link-Time-Optimization of ARM Linux kernel

What is demonstrated

- Possibly the most ~~boring~~ ^{thrilling} demo ever
- Gcc has compile-time option to do link-time optimization
- Andi Kleen created patches to support this compiler option
 - He demonstrated on an Intel CPU
- This is first demonstration of LTO kernel running on ARM
 - **World's first, that I know of !!!**
- LTO supports whole-program optimization, at final link time
 - Slow link step, but good code optimizations

What was improved

System size and performance

6% reduction in image size (384K)

Kernel	non-lto	lto
Compile time	1m58s	3m22s
Image size	5.85M	5.46M
Meminfo Total	17804K	18188K
Meminfo Free	10908K	11260K
LTP time		

Hardware Information

TI panda board
mem=24M



pandaboard.org

Source code or detail technical information availability

<http://lwn.net/Articles/512548/>

[git://github.com/andikleen/linux-misc](https://github.com/andikleen/linux-misc)



CE Workgroup

LTO (cont.)

- **Cost:**
 - Longer kernel builds (extra 1.5 minutes)
 - More memory during build (up to 9G required for allyesconfig)
 - Subtle bugs from optimizations
 - E.g. duplicate code elimination caused a pointer comparison failure
- **Benefits:**
 - Size reduction – 380K (6%) on ARM
 - Performance: Unknown



CE Workgroup

LTO (cont. 2)

- Why am I so excited about this?
- Opens new possibilities for automatic kernel reduction techniques
 - It is not tractable to reduce kernel manually
 - Whole system optimization is a critical part of automatic reduction
- Note: This work obsoletes `-ffunction-sections`
- Takes Linux-tiny in a whole new direction



CE Workgroup

Volatile Ranges

- Work by John Stultz
 - Inspired by Android feature in ashmem
 - <http://lwn.net/Articles/468896/>
 - <http://lwn.net/Articles/500382/>
- Allows cooperation between the kernel and applications on "volatile" memory usage
- Overview:
 - Application notifies kernel about re-claimable memory areas
 - Not mainlined yet



CE Workgroup

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff



Resources



CE Workgroup

CEWG Contract Work 2012

- eMMC tuning
- Dynamic memory reduction
- Mainline FIQ debugger
- ConnMan support for WiFi direct
- Improve kexecboot
- Measure systemd and udev





CE Workgroup

eMMC tuning guide

- Description:
 - This project analysed EXT3, 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: work is almost complete
 - Document should be available very soon




CE Workgroup

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
 - See ELC 2013 talk




Drivers kmalloc





Linux kmalloc





Mainline FIQ debugger

- Description:
 - Add ARM FIQ glue code and integrate with existing kernel debugger
 - Allows use of ARM FIQ (non maskable interrupt) to activate a kernel debugger
- Status:
 - Developer worked on this independent of CE WG
 - Now called "NMI KGDB/KDB debugger"
 - Not just ARM FIQ glue code.
 - The generic driver is now `drivers/tty/serial/kgdb_nmi.c`
 - Further development (i.e. ARM-specific bits, and restricted mode) is in the following GIT tree:
 - [git://git.infradead.org/users/cbou/linux-nmi-kdb.git](http://git.infradead.org/users/cbou/linux-nmi-kdb.git)
<http://git.infradead.org/users/cbou/linux-nmi-kdb.git>



CE Workgroup

ConnMann WiFi direct

- Description:
 - Add support for WiFi direct to ConnMann wireless connection manager
- Contractor: contractor was acquired
- Status:
 - Considering project for 2013



CE Workgroup



Improve kexecboot

- Description:
 - Make improvements to kexecboot bootloader
 - Support load from network
 - UI improvements
- Contractor: Yuri Bushmelev
- Status:
 - Should be done in May



Measure systemd and udev

- Description:
 - Measure the overhead and performance of system and udev, as used in embedded systems
- Contractor: became unavailable
- Status: considering project for 2013




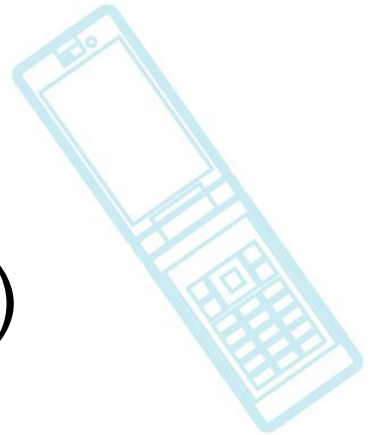




CE Workgroup

Other Projects

- Long Term Support Initiative (LTSI)
- Hardware fund







CE Workgroup

Long Term Support Kernel for Industry

- LTSI 3.4 is available now
- Many presentations available on status
 - See ELC 2013 presentation by Hisao Munakata
 - See ELC 2013 presentation by Tzugikazu Shibata
- Program for free hardware for LTSI kernel testing





CE Workgroup

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 – 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





CE Workgroup

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
 - "Kernel Testing Tools and Techniques" BOF by Matt Porter



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
- Android use in non-CE embedded
 - Headless android



CE Workgroup

Events

- ELC/Android Builders Summit – Feb 2012
- 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



CE Workgroup

Miscellaneous

- Has use of open source licenses peaked?
 - Interesting essay on moving to more free licenses (specifically, public domain)
 - Argument is that now individuals and companies will contribute even if license doesn't require it
 - Most developers understand benefits
 - <http://www.h-online.com/open/features/Why-it-s-time-to-stop-using-open-source-licences-1802140.html>



CE Workgroup

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_3.?.
- eLinux wiki - <http://elinux.org/>
 - Especially <http://elinux.org/Events> for slides
- Celinux-dev mailing list



CE Workgroup



Thanks!

