



# *CE Workgroup*

# Status of Embedded Linux

June 2011

Tim Bird

Architecture Group Chair




LF CE Workgroup



CE Workgroup

# Outline




- Kernel Versions
- Technology Areas
- CE Workgroup Projects
- Tools
- Embedded Distributions
- Industry Organizations
- Miscellaneous
- Discussion
- Resources





CE Workgroup

# Linux Kernel Versions





CE Workgroup

# Kernel Versions

- Linux v2.6.34 – 16 May 2010
- Linux v2.6.35 – 1 Aug 2010
- Linux v2.6.36 – 20 Oct 2010
- Linux v2.6.37 – 4 Jan 2011
- Linux v2.6.38 – 14 Mar 2011
- Linux v2.6.39 – 19 May 2011
  - 5 versions in 12 months
- Linux v3.0
  - Expected in August for 20<sup>th</sup> anniversary



CE Workgroup

# Linux v2.6.35

- User-space OOM notifier
- Cpuidle – idle pattern detection
  - Can detect when a periodic interrupt is causing a steady wakeup, and adjust next-wakeup accordingly
  - See <http://lwn.net/Articles/387250/>
- Timer slack mechanism introduced
  - Allows for combining timers within a “slack” range, decreasing wakeups and saving power
  - See <http://lwn.net/Articles/369549/>
- Ramoops driver
  - Record oops into RAM for later analysis



CE Workgroup

# Linux v2.6.36



- AppArmor – path-based security module
- Wakeup counts
  - Kernel-user interface to allow system to suspend aggressively without race conditions on wakup events
- New OOM killer
  - <http://lwn.net/Articles/391222/>
- More BKL removal
- LZO compression in SquashFS
- Runtime PM statistics



CE Workgroup

# Linux v2.6.37

- Jump labels
  - Eliminates (almost completely) the overhead when tracing calls are disabled
  - See <http://lwn.net/Articles/412072/>









CE Workgroup

# Linux v2.6.38

- Perf symbols abstraction
  - Added 'symfs' option for off-box analysis of perf.data
  - Should be good for embedded









CE Workgroup

# Linux v2.6.39

- Pstore
  - Store information from dying kernel into some persistent storage
  - Similar to mtdoops or ramoops
  - See <http://lwn.net/Articles/434821/>
- Device power domains for runtime PM
- ARM arch tree changes (just starting)





CE Workgroup

# Linux v3.0 (probable)




- Fast symbol resolution for module loading
  - Binary search instead of linear lookup for module linking
- POSIX alarm timers
  - Similar to Android Alarm Timers
- BKL function calls are now gone
- More ARM arch tree changes



CE Workgroup



# Things to watch

- ARM IRQ re-work
- ARM arch sub-tree refactoring
  - <http://lwn.net/Articles/443510/>
- Device trees
- More runtime PM improvements
  - Android effect on PM features







CE Workgroup



# Technology Areas





CE Workgroup

# Bootup Time

- **Readahead getting lots of attention**
  - Ureadahead in Ubuntu
  - See my presentation at ABS about readahead with Android
- **Snapshot boot**
  - Old topic, but still very popular
  - Requires work both inside and outside kernel
    - Not much mainlined
  - See ELC presentation by Kang Dongwook
- **Filesystem speedups**
  - CELF funding work in this area (more later)



CE Workgroup

# Bootup Time (cont.)

- XIP (Execute-In-Place)
  - Almost removed from kernel
    - Version in kernel was broken
    - Use of XIP on only out-of-tree platforms is a problem
- Bootloader improvements
  - Coreboot on x86
    - See “Really fast x86 boot” presentation at FOSDEM 2011
  - U-Boot ARM caching enhancements
- See presentation by Andrew Murray at ELC Europe 2010
  - Very good philosophy of boot time reduction



CE Workgroup

# Graphics



- 3D
  - OpenGL ES is de-facto standard everywhere
- 2D
  - Android had Skia, but is moving to...?
  - Meego used Clutter, Qt, and X
  - Framebuffer is going away, with acceleration required for larger screens
- Wayland
  - Intel moving towards Wayland
  - Replacement for X?
  - Support for multiple top-layer APIs
- Lots of work around memory management between kernel, user-space and GPU



CE Workgroup

# Graphics (cont.)

- Accelerated rendering is a big topic
  - Google introduced renderscript
    - Uses LLVM to do runtime retargeting of script to whatever capabilities device has
- Ability to support GPU in SOC is very important







CE Workgroup

# Graphics Drivers



- PowerVR graphics driver
  - PowerVR is being used lots of places
    - Intel adopted for Cedarview and it's already in Sodaville
    - Is in very many ARM SOCs
  - PowerVR driver is closed-source
  - Alan Cox submitted some driver pieces in February
    - Omitted anything relating to out-of-tree binary driver
    - See <http://thread.gmane.org/gmane.linux.kernel/1103793>



CE Workgroup

# Multimedia

- **Gstreamer**
  - Is still being used in TVs
    - Ex: Google TV uses it
- **Android media layer**
  - Stagefright – new media layer
    - Replaces OpenCore?
- **Codec wars**
  - WebM/VP8
    - Free codec by Google
    - Integrated into HTML5








CE Workgroup

# File Systems

- **UBIFS**
  - Replacing JFFS2 as default raw flash FS of choice
- **YAFFS2 is not in mainline yet**
  - Despite CELF funding
- **LogFS**
  - Appears to be abandoned





CE Workgroup

# File Systems (cont.)

- Google moving to Ext4 for future Android devices
  - Already using eMMC instead of raw flash
  - Developers said that main reason was SMP performance
- Want to optimize Linux filesystem layers for flash
  - See Arnd Bergmann's work on filesystem performance on cheap flash media (ELC 2011)



CE Workgroup

# Power Management

- Runtime Power Management
  - Relatively new ability to suspend and resume individual system components
  - See <http://lwn.net/Articles/347573/>
- See Magnus Damm's slides at: [http://elinux.org/ELC\\_2011\\_Presentations](http://elinux.org/ELC_2011_Presentations)
- Rafael Wysocki's presentation here at LCJ
- Device power domains



CE Workgroup




# System Size

- CE WG is reviving Linux-tiny project
- Bloatwatch still running – but who looks at it?
  - <http://www.selenic.com/bloatwatch>
  - Big increases in some kernel versions
- Xi Wang had a good talk at ELC 2010 about optimizing memory usage throughout the system
- User space is memory problem area now
  - OOM killer or OOM avoidance is big issue
    - Cgroup memory notifications
    - Android has it's own thing
      - Application lifecycle is key feature



CE Workgroup

# CE WG Contract Work






CE Workgroup

# CELF Contract Work 2010

- Bootchart and smemcap in busybox
- Function-sections
- YAFFS2 mainline effort
- SquashFS enhancements
- U-Boot ARM enhancements
- Trace format standard
- Kexecboot enhancements
- Flash filesystem testing









CE Workgroup

# Mainline YAFFS2 effort

- YAFFS2 is a popular NAND flash filesystem
  - Was used by Android in many devices
- 3 mainline attempts made, but hit some barriers
  - Currently stuck on some locking issues
- Outlook for mainline acceptance is uncertain





CE Workgroup

# Trace Format Standard




- Create a singled trace format standard for the embedded industry (CTF – Common Trace Format)
  - See <http://www.efficios.com/ctf>
  - Allows reuse of tools with data from different tracing systems
- BabelTrace trace conversion library
  - Converts trace formats into CTF (and back?)
  - Proof of concept conversion implementation
    - Can convert kernel messages with timestamps to CTF and back to text



CE Workgroup

# Contract Work 2011

- Mainline fast symbol resolution
- Mainline Device Firmware Upgrade (DFU) code in U-Boot
- Work on Linux tiny patches
- Improve UBIFS mount time
- Support read-only block filesystems on flash devices
- Flash filesystem testing







CE Workgroup

# Contract Work 2011 (cont.)

- Mainline the watchdog framework
- Extend bluetooth stack with Remote SIM Access protocol
- Kernel trace and debug documentation (on eLinux wiki)
- Mainline Android kernel features





CE Workgroup

# Contract Work Details

- Mainline fast symbol resolution
  - Change symbol lookup to use binary search instead of linear scan to speed up module loading
  - Already mainlined (Linux v3.0)
- Mainline DFU code in U-Boot
  - Device Firmware Upgrade (DFU) is an industry standard for upgrading and manipulating firmware in embedded devices
- Work on Linux tiny patches
  - Revive Linux-tiny patch set
  - Forward-port patches to latest kernel
  - Add more patches to improve kernel configurability



CE Workgroup

# Contract Work Details (2)

- Improve UBIFS mount time
  - Add logging or checkpointing to UBI to avoid bad-block scan of whole device on UBI attach
- Support read-only block filesystems on flash devices
  - Write block emulation layer to support read-only filesystems on top of MTD layer in kernel
  - Will allow Squashfs to be used on raw NAND flash media
- Flash filesystem testing
  - Publish performance results for each new kernel version



CE Workgroup

# Contract Work Details (3)



- Mainline the watchdog framework
  - Provides a generalized watchdog mechanism
    - Should provide easier method to add watchdogging to drivers and the kernel going forward
  - Original framework was written by Alan Cox and others
- Extend bluetooth stack with Remote SIM Access protocol
  - Allows for Linux bluetooth and telephony stack to utilize SIM in external device for operation
  - Primary use case is for Linux-based in-car system to utilize SIM in mobile device for calls, etc.



CE Workgroup

# Contract Work Details (4)




- Mainline Android kernel features
  - Goal is to incrementally reduce diff between Android and mainline kernels
  - Probably do pilot project to mainline Android logger code
    - If successful, will try other pieces









CE Workgroup



# Tools





CE Workgroup

# Tools




- QEMU
  - QEMU is being used everywhere, for device emulation (Android, Yocto)
  - Javascript QEMU implementation (!!)
- Eclipse
  - Is now de-facto “umbrella” tool for development
  - Need to pry seasoned developers away from command line
- Tracing
  - Common Trace Format standard exists



CE Workgroup

# Build Systems

- Yocto project
  - Umbrella project – has builder, eclipse tools, other things
  - OpenEmbedded and Yocto are getting integrated
  - Tons of talks at ELC 2011
- Still lots of custom build systems out there





CE Workgroup

# Embedded Distributions

- Meego
  - Version 1.2 released
  - Nokia switching to Windows Mobile
  - Still looking for products
- Android
- WebOS
  - HP may license OS to 3<sup>rd</sup> parties
- Legacy custom embedded
  - Still no “standard” embedded distribution



CE Workgroup

# Android

- Android 3.1 SDK (Honeycomb r2) released May, 2011
- Ice Cream Sandwich due Q4
  - Will unify mobile, tablet and TV platforms in one codebase
- Phone activations at 400,000 per day
- Dalvik ported to non-Android
  - Myriad Alien Dalvik for MeeGo
  - IcedRobot for native Linux
  - “Dalvik on Any Devices” – session later today



CE Workgroup

# Industry organizations

- Linux Foundation
  - Has lots of embedded-related projects
    - Yocto, Meego, CE Workgroup
    - Recently announced Meego TV workgroup
- CELF merger with LF
  - CELF is now the LF “CE Workgroup”
  - Now utilizing LF infrastructure
    - Should mean it’s easier for public to participate in CE WG initiatives
- Linaro
  - See David’s Rusling’s talk



CE Workgroup

# Miscellaneous

- **Unlockable bootloaders**
  - Announced by Motorola, Sony/Ericsson
  - Can unlock bootloader to install custom firmware
  - Wipes the phone to remove DRM-protected content
  - Motorola says you can re-lock by reinstalling vendor image
- **Increased use of Stack Overflow**
  - Great site for answering detailed development questions
  - See [www.youtube.com/watch?v=NWHfY\\_IvKIQ](http://www.youtube.com/watch?v=NWHfY_IvKIQ)
  - Google developers answer questions here



CE Workgroup

# Observations

- Rate of “general features for embedded” contributions to kernel seems low
  - We seem to have stalled on bootup time reduction, size reduction, realtime, security in embedded
    - Some problems and solutions shifted to user space
  - Hot areas in kernel:
    - Power management, ARM board support refactoring, GPU management (memory sharing, driver support)
- Still seeking ways to facilitate participation of embedded developers in community








CE Workgroup

# What are you working on?

- Good measure of what needs work is whatever developers spent a lot of time working on last year...
- What was that?







CE Workgroup

# Resources

- LWN.net
  - <http://lwn.net/>
  - If you are not subscribed, please do so
- Kernel Newbies
  - [http://kernelnewbies.org/Linux\\_2\\_6\\_??](http://kernelnewbies.org/Linux_2_6_??)
- eLinux wiki - <http://elinux.org/>
  - Especially <http://elinux.org/Events> for slides
- Linux-embedded mailing list
  - <http://vger.kernel.org/vger-lists.html#linux-embedded>



CE Workgroup



**Thanks!**

