



# *CE Workgroup*

# Status of Embedded Linux

February 2012

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 v2.6.38 – 14 Mar 2011 – 69 days
- Linux v2.6.39 – 19 May 2011 – 66 days
- Linux v3.0 – 21 July 2011 – 63 days
- Linux v3.1 – 24 Oct 2011 – 95 days
- Linux v3.2 – 4 Jan 2012 – 72 days
- Linux v3.3-rc3 - (as of 15 Feb)



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



- Fast symbol resolution for module loading
  - Binary search instead of linear lookup for module linking
- POSIX alarm timers
  - Similar to Android Alarm Timers
  - See <http://lwn.net/Articles/429935/>
- BKL function calls are now gone
- More ARM arch tree changes



CE Workgroup

# Linux v3.1

- Watchdog timer core
- New framework for handling power management domains was added
  - See `driver/base/power/domain.c`
- Multiple ARM SoCs now have device tree support







CE Workgroup

# Linux v3.2

- New pin control subsystem
  - Allows control of multiple pins as named groups, with multiplexing
  - See [Documentation/pinctrl.txt](#)
  - See ELC talk on Friday by Linus Walleij
- devfreq – DVFS for non-cpu devices
- PM QOS now supports per-device constraints
  - See [Documentation/power/pm\\_qos\\_interface.txt](#)
  - See <http://lwn.net/Articles/466230>



CE Workgroup

# Linux v3.3 (probable)

- ARM large physical address extensions
  - See Catalin Marinas talk at ELC Europe
- 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
  - Please don't use any interfaces from code in staging!!



CE Workgroup

# Things to watch

- ARM arch sub-tree refactoring
  - <http://lwn.net/Articles/443510/>
  - See Arnd Bergmann talk on Thursday
- Device trees
  - See Thomas Abraham talk today
- More runtime PM improvements
- Android features
  - Especially after October kernel summit
- Boot timing patches
  - See <http://lkml.org/lkml/2011/9/23/348>



CE Workgroup

# Outline



Kernel Versions

**Technology Areas**

CE Workgroup Projects

Other Stuff

Resources





CE Workgroup

# Bootup Time

- Not so much a kernel problem any more
  - Lots of talks and presentations
  - Good kernel techniques on eLinux wiki
    - [http://elinux.org/Boot\\_Time](http://elinux.org/Boot_Time)
- User-space is big problem area now
  - Features for overall performance
- See presentation by Andrew Murray at ELC Europe 2010
  - Very good philosophy of boot time reduction
    - Bootup time work = re-specialization of software





CE Workgroup

# Bootup Time technologies

- Bootloader improvements
  - Coreboot on x86
  - See “Really fast x86 boot” presentation at FOSDEM 2011
  - U-Boot ARM caching enhancements
- Snapshot boot
  - Old topic, but still very popular
  - Requires work both inside and outside kernel
    - Not much mainlined
  - See ELC 2011 presentation by Kang Dongwook



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
- Filesystem speedups
  - CELF funding work in this area (more later)
  - Readahead getting lots of attention
    - Ureadahead in Ubuntu
    - See Tim Bird presentation at ABS 2011 about readahead with Android



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





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

# Accelerated rendering

- 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 2011
    - 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
  - Still needs some boot time improvements
- **YAFFS2 is not in mainline yet**
  - Despite CELF funding
- **LogFS**
  - Appears to be abandoned
- **AXFS**
  - Advanced XIP File system – developed by Intel/Numonyx but never mainlined



CE Workgroup

# File Systems (cont.)

- Google moving to Ext4 for future Android devices
  - Already using eMMC instead of raw flash
  - Sad to see proprietary algorithms in black boxes responsible for storage performance
    - Lots of MMC optimized for serial workloads and FAT filesystems
- Want to optimize Linux block filesystem layers for flash
  - See Arnd Bergmann's talk at ELCE on filesystem performance on cheap flash media
  - See Ken Tough's ELC talk on Thursday





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)
- Device power domains
  - Set of devices sharing power resources (clocks, power planes, etc.)
  - See Rafael Wysocki's talks at LinuxCon Japan 2011 and ELC Europe 2011





CE Workgroup

# Power Management

- New attempt at wakelock-compatible solution by Rafael Wysocki
  - “Autosleep and wakelocks”
  - <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"*
  - Appears to be generating less heated discussion



CE Workgroup

# System Size

- CE WG has revived the Linux-tiny project
- Bloatwatch still running – but who looks at it?
  - <http://www.selenic.com/bloatwatch>
  - Big increases in some kernel versions
- Poky-tiny
- Good talks recently:
  - Xi Wang at ELC 2011 about optimizing memory usage throughout the system
  - Darren Hart at ELCE 2011 – poky-tiny
- User space is memory problem area now
  - OOM killer or OOM avoidance is big issue
    - Cgroup memory notifications
    - Android has it's own low memory killer
      - Application lifecycle is key feature



CE Workgroup

# Observations

- Rate of “general features for embedded” contributions to kernel seems low
  - Not a lot of progress recently on bootup time reduction, size reduction, 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

# Outline



Kernel Versions

Technology Areas

**CE Workgroup Projects**

Other Stuff

Resources







CE Workgroup

# CEWG 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

# CEWG 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
- Was a classic case of developer wanting to retain multi-platform support
  - This approach is rejected by community



CE Workgroup

# Trace Format Standard

- Create a singled trace format standard for the embedded industry (CTF – Common Trace Format)
  - See <http://www.fficios.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

# CEWG 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
- Flash filesystem testing
- Mainline the watchdog framework
- Extend bluetooth stack
- Kernel trace and debug documentation





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
  - Mostly mainlined in u-boot and kernel
- Work on Linux tiny patches
  - Revive Linux-tiny patch set
  - Forward-port patches to latest kernel
  - Add more patches to improve kernel configurability
  - Last work was with function-sections for kernel





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
  - Supposed to see patches this week...
- Flash filesystem testing
  - Publish performance results for each new kernel version
  - Lots of great data – charts and graphs!
  - Check out:  
[http://elinux.org/Flash\\_Filesystem\\_Benchmarks](http://elinux.org/Flash_Filesystem_Benchmarks)



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
  - Mainlined in 3.1
- 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 is for Linux-based in-car system to utilize SIM in mobile device for telephony
  - Mainlined in upstream bluez and kernel



CE Workgroup

# Projects

- Android mainline project
- Long Term Support Initiative (LTSI)
- Open Project Proposals for 2012





CE Workgroup

# Android Mainline Project



- Mainline Android kernel features
  - Goal is to incrementally reduce diff between Android and mainline kernels
- Interesting discussion at kernel summit
  - Would be nice to support Android with mainline kernel
  - Linus – we've taken sub-optimal stuff before
- Multi-party effort to mainline patches
  - CE WG, Linaro, and others
  - Greg KH put some files into drivers/staging
- Good discussion last week at Linaro Connect
- 3.3 kernel (with 12 lines of patches) boots AOSP



CE Workgroup

# Long Term Support Kernel for Industry

- CE Workgroup is initiating a new project for companies to collaborate on maintaining a kernel version for embedded products
  - Similar to long-term kernel maintained in enterprise space
  - Based on community long-term tree
- See presentation by Tsugikazu Shibata








CE Workgroup

# LSTI reasons

- Various effects contribute to low contribution rate from consumer electronics product teams
  - Version gap, product schedule impedance mismatch with mainline releases, focus on short-term rather than long-term solutions
- Want to create an area for collaboration between companies, as well as a staging ground for moving code to mainline

# LTSI project overview

- Project consists of three parts





CE Workgroup

# LSTI details

- The plan (subject to change):
  - 2-year overlapping releases
  - Bugfixes from community longterm tree and product trees
  - Backport of some features from mainline
  - Integration of some (a very small set) of out-of-mainline patches (e.g. LTTng, RT-preempt, Linux-tiny)
- Should have first release in early 2012



CE Workgroup

# Open Project Proposals

- Will be announced on celinux-dev mailing list
- Look for announcement, and proposal instructions on eLinux wiki very soon
  - [http://elinux.org/CELF\\_Open\\_Project\\_Proposal\\_2012](http://elinux.org/CELF_Open_Project_Proposal_2012)
- Please propose a project you think would benefit embedded Linux



CE Workgroup

# Outline



Kernel Versions

Technology Areas

CE Workgroup Projects

**Other Stuff**

Resources










CE Workgroup

# Other Stuff

- Tools
- Build Systems
- Distributions
- Android
- Industry Organizations
- Events
- Miscellaneous





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**
  - Perf, ftrace and LTTng 2.0
  - Common Trace Format standard



CE Workgroup

# Build Systems

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





CE Workgroup

# Embedded Distributions

- Tizen = MeeGo + Limo + (WAC technologies)
  - Was announced a few months ago
  - Nokia switching to Windows Mobile
  - Focus = HTML5 applications
  - <http://www.tizen.org/>
- WebOS
  - Open source announced
- Legacy custom embedded
  - Still no “standard” embedded distribution



CE Workgroup

# Android

- Android 4.0 SDK (Ice Cream Sandwich) released October 2011
  - Source released this week!
- Ice Cream Sandwich unifies mobile, tablet and TV platforms in one codebase
- Phone activations at 700,000 per day
- Dalvik ported to non-Android
  - Myriad Alien Dalvik for MeeGo
  - IcedRobot for native Linux
  - OpenMobile's ACL (Application Compatibility Layer)





CE Workgroup

# Distributions

1998 2000 2002 2004 2006 2008 2010

Embedix

Hardhat MontaVista

Denx ELDK

TimeSys Linux RT

LynuxWorks

Wind River Linux

Ångström

Poky

Yocto

Maemo

Moblin Meego Tizen


Limo

Embedded Debian

Ubuntu (embedded)

Android

Fire ?





CE Workgroup

# Industry organizations




- Linux Foundation
  - Has lots of embedded-related projects
    - Yocto, MeeGo, CE Workgroup
- CE Workgroup
  - Now utilizing LF infrastructure
    - Should mean it's easier for public to participate in CE WG initiatives
      - Was out of commission in fall
- Linaro
  - Doing lots of great stuff
  - See David Rusling's ELC 2011 talk



CE Workgroup

# Events

- Android Builders Summit
  - Just finished yesterday
  - Content will be online
- Embedded Linux Conference Europe 2012
  - November 7-9, 2012
  - Barcelona, Spain







CE Workgroup

# Miscellaneous

- 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

# 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
  - Some areas have lots of content – some need work





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\\_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!**

