



CE Workgroup

Status of Embedded Linux

November 2011

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 v2.6.35 – 1 Aug 2010
 - Last longterm release for embedded = 2.6.35.14
- Linux v2.6.36 – 20 Oct 2010
- Linux v2.6.37 – 4 Jan 2011
- 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



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 wakeup 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



- 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 (probable)




- New pin control subsystem
 - Allows control of multiple pins as named groups, with multiplexing
 - See [Documentation/pinctrl.txt](#)
- 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

Things to watch

- ARM arch sub-tree refactoring
 - <http://lwn.net/Articles/443510/>
- Device trees
- More runtime PM improvements
- Android features
 - Especially after recent 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

- 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
 - Bootup time work = re-specialization of software



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.)

- /dev/ion -- a unified approach to buffer management and sharing between display, GPU, camera, codecs, etc, new in Ice Cream Sandwich
 - Looking forward to looking at code
- 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
 - 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



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
- 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

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
- Good talks recently:
 - Xi Wang at ELC 2011 about optimizing memory usage throughout the system
 - Darren Hart at ELCE 2011
- 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

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







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
- 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



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 is for Linux-based in-car system to utilize SIM in mobile device for telephony



CE Workgroup

Contract Work Details (4)

- Mainline Android kernel features
 - Goal is to incrementally reduce diff between Android and mainline kernels
 - Have CE Workgroup funding approval to do a pilot project to mainline Android logger code
 - If successful, will try other pieces
 - Almost finished creating broken-out patch set for android-common (diff against 3.0)
 - Linaro developer also has patches isolated into topic branches in stgit
 - Given recent input at kernel summit, I plan to broaden the scope of this project
 - If you are interested in this – email me!



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 at ELC Europe



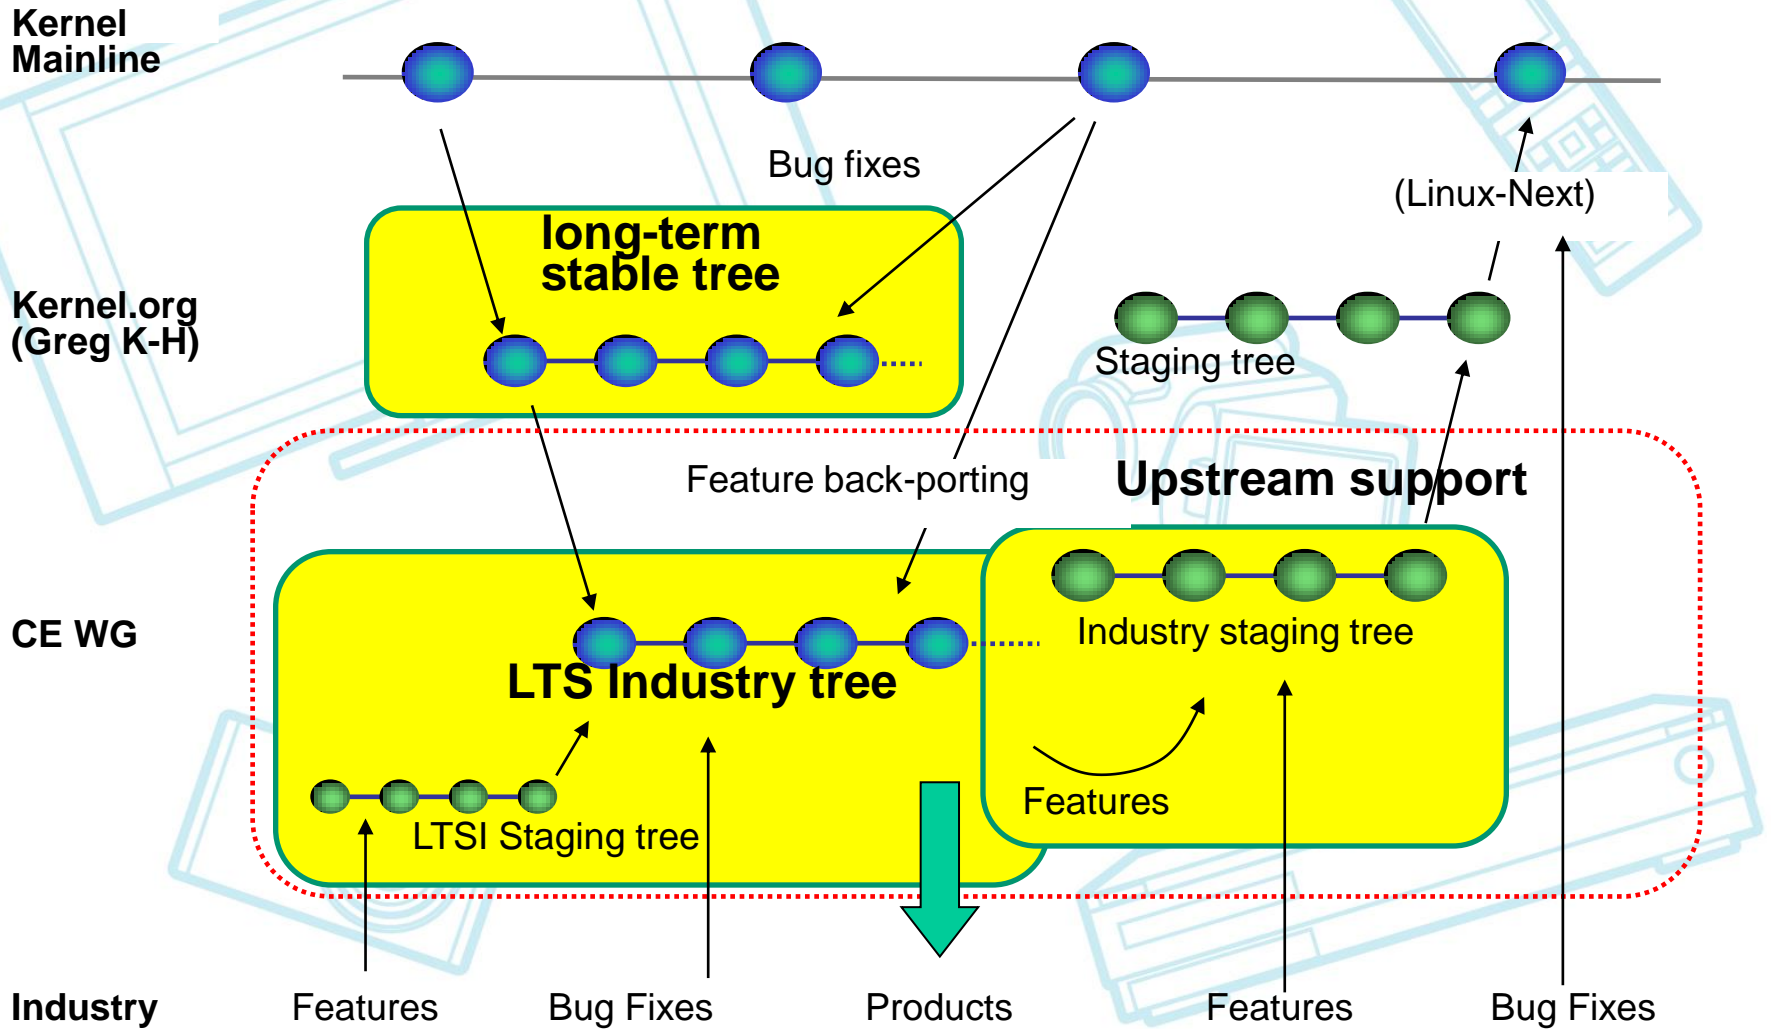
CE Workgroup

LSTI reasons

- Various effects contribute to low contribution rate from consumer electronics product teams
 - Version gap, product schedule impedance 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

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**
 - 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
 - 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 weeks ago
 - Nokia switching to Windows Mobile
 - Focus = HTML5 applications
 - <http://www.tizen.org/>
- WebOS
 - HP using it internally, but it won't be a platform (from HP) in the future
- 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 550,000 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
- 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
 - Except for when it’s out of commission
- Linaro
 - Doing lots of great stuff
 - See David Rusling’s ELC 2011 talk



CE Workgroup

Events

- **Android Builders Summit**
 - First one was last year
 - February 13-14
 - Redwood Shores, California
- **Embedded Linux Conference**
 - February 15-17
 - Redwood Shores, California
 - Call for papers going out soon
- **Embedded Linux Conference Europe 2012**
 - November 7-9
 - Barcelona, Spain





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
 - 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_??
- 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!

