



CE Linux Forum

Status of Embedded Linux

September 2010

Tim Bird




CELF Architecture Group Chair



CE Linux Forum

Outline

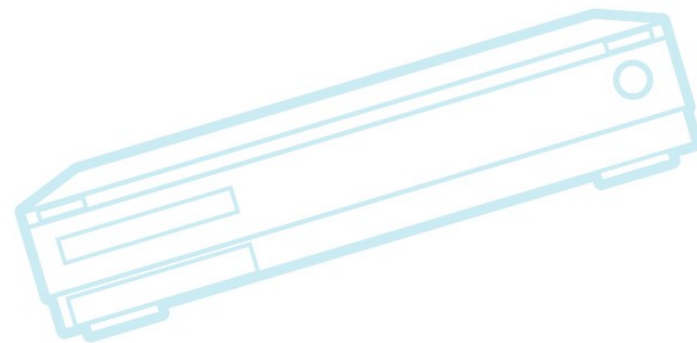

- Kernel Versions
- Technology Areas
- CELF Contract Work
- Embedded Distributions
- What to do?
- Resources





CE Linux Forum

Linux Kernel Versions





CE Linux Forum

Kernel Versions

- Linux v2.6.31 – 9 Sep 2009
- Linux v2.6.32 – 2 Dec 2009
- Linux v2.6.33 – 24 Feb 2010
- Linux v2.6.34 – 16 May 2010
- Linux v2.6.35 – 1 Aug 2010
 - 5 versions in 13 months
- Linux v2.6.35-rc5 – 20 Sep 2010
 - Probable 2.6.36 in a few weeks?



CE Linux Forum

Linux v2.6.31




- Ftrace features
 - Generic filters, function profiler, new tracepoints, new documentation
- SMACK security module logging
- Performance counters (“Perf”)
 - <http://lwn.net/Articles/311850/>
- kmemleak - detect kernel memory leaks
 - <http://lwn.net/Articles/187979/>



CE Linux Forum

Linux v2.6.32

- devtmpfs
 - Dynamic, fast population of /dev
 - <http://lwn.net/Articles/331818/>
- Timechart tool
 - New tool to create SVG chart of kernel events
 - See <http://blog.fenrus.org/?p=5>
- Runtime Power Management (core)
 - <http://lwn.net/Articles/347573/>
 - More on this later







CE Linux Forum

Linux v2.6.33

- LZO kernel compression
- Perf improvements
 - Probe, bench, diff, perl scripts, filters
 - ARM support
- Ramzswap/Compcache
- Android patches removed from -staging





CE Linux Forum

Linux v2.6.34

- LogFS
 - Log-structured flash filesystem
- Asynchronous suspend/resume
 - See <http://lwn.net/Articles/366915/>
- Cgroup memory threshold notifications
- Python scripting support for perf
 - <http://lwn.net/Articles/373842/>
 - OK – I like Python



CE Linux Forum

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
 - <http://lwn.net/Articles/387250/>
- Timer slack mechanism introduced
 - Allows for combining timers within a “slack” range, decreasing wakeups and saving power
 - <http://lwn.net/Articles/369549/>
- Ramoops driver
 - Record oops into RAM for later analysis



CE Linux Forum

Linux v2.6.36 (probably)



- 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



CE Linux Forum

Linux v2.6.36 (probably) cont.




- LIRC infrared controller driver suite
 - Support for infrared devices (e.g. TV remote controls)
 - Currently in staging – go use and test if you want it to move out of staging
 - 11 year-old project!





CE Linux Forum

Technology Areas







CE Linux Forum

Technology Areas

- Audio, Video, Graphics
- Bootup Time
- File Systems
- Power Management
- Real-time
- Security
- System Size
- Tracing/Tools

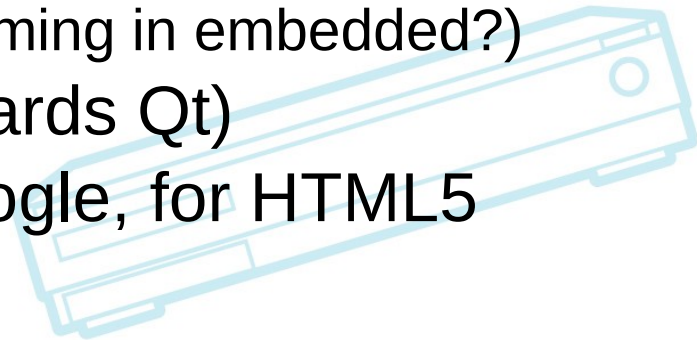
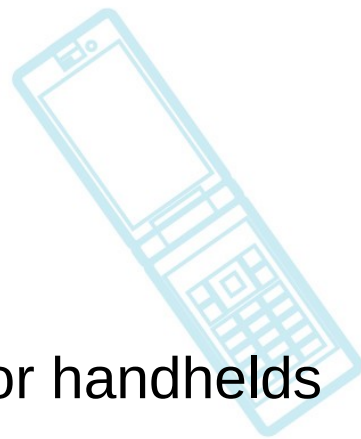





CE Linux Forum

Audio/Video/Graphics

- GStreamer
 - Is de-facto standard for video management for handhelds
- OpenCore/OMX
 - Android media layer
- DirectFB
 - Continued uptake in TV domain
- OpenGL ES
 - De-facto standard for 3D API (/Gaming in embedded?)
- Clutter, Qt (Meego moving towards Qt)
- WebM/VP8 – free codec by Google, for HTML5





CE Linux Forum

Bootup Time

- Lots of readahead work
 - Sreadahead
 - Ureadahead
- Asynchronous suspend/resume
- Kernel Execute-in-Place (XIP)
- Lots of work on unhibernate (resume from cold boot)
 - Example: Android boot in 1 second:
 - <http://www.linuxfordevices.com/c/a/News/Ubiquitous-QuickBoot/>
- LogFS – has very short mount times





Bootup patches to watch



- Speed up symbol resolution during module loading
 - Binary search in symbol table avoids linear lookup cost per symbol when loading a module
 - Developer, Alan Jenkins, reports saving 20% of time on coldboot
 - Depends on modules being loaded
 - Also, Carmelo Amorosso has patches for hash table lookup
 - Faster than binary search, but requires symbol table storage change
 - http://elinux.org/ELC_Europe_Presentations_2009



CE Linux Forum

File Systems - SquashFS

- Compressed, read-only FS
 - <http://www.linux-mag.com/cache/7357/1.html>
- Recent news:
 - Patches for LZMA support were submitted to kernel mailing list, but rejected
 - Out-of-tree patches are available
 - Support for LZO compression added for 2.6.36





CE Linux Forum

More File Systems

- LogFS
 - Log-structured file system which keeps meta-data on flash
 - Mounts really fast (no need to build in-memory tables to start accessing)
 - Mainlined (after much work) in 2.6.34
 - CELF supported this a few years ago
 - Introduction
 - <http://lwn.net/Articles/234441/>
 - Announcement of mainlining
 - <http://lwn.net/Articles/377741>



CE Linux Forum

Power Management

- Runtime Power Management
 - Adds ability to suspend and resume individual system components
 - See <http://lwn.net/Articles/347573/>
 - See Kevin Hillman's slides at:
 - http://elinux.org/ELC_2010_Presentations
- Asynchronous suspend and resume
 - Faster suspend and resume saves power
 - Usually associated with user-initiated suspend



FreeFoto.com



Power Management (cont.)

- Suspend blockers submitted by Google in April/May
 - A facility to enable “aggressive auto-suspend”
 - Successor to wake locks
 - There was much discussion about the technology
 - See <http://lwn.net/Articles/388131/>
 - See <http://lwn.net/Articles/390369/>
- Wakeup_counts
 - Community response to suspend blockers
 - Similar, but without some controversial parts
 - <http://lwn.net/Articles/393314/>
- Also, there are PM_QOS changes in progress



Real-time – RT-preempt

- Sleeping Spinlocks
 - Spinlock name cleanup merged in 2.6.33
 - See <http://lkml.org/2009/12/6/162>
 - Macros for doing switch from spinlocks to semaphores remains to be merged
- Lots of effort to eliminate use of “Big Kernel Lock”
 - CELF providing some support (hardware for developer)



CE Linux Forum

System Size / Memory

- smem
- Kernel image compression
- Ramzswap
- Cgroup mem notifications
- Function sections
- XIP
 - Not much new development (AXFS?), but still pretty important





CE Linux Forum

smem




- Tool for analyzing system memory
 - Better numbers than ps or top
 - Reports Proportional Set Size (PSS), which accounts better for shared pages, and doesn't lie, like RSS
 - Home page: <http://www.selenic.com/smem>
 - See <http://lwn.net/Articles/329458>
- 0.9 release was made in November
 - See <http://lwn.net/Articles/361497>
 - Release includes a man page and smemcap
 - smemcap = lightweight snapshotting tool for using smem with an embedded system
 - Smemcap applet added to busybox



CE Linux Forum

Kernel Image Compression

- LZO kernel image compression
 - See <http://lwn.net/Articles/350985>
 - Compression not so great, but faster than others
 - Merged in 2.6.33





CE Linux Forum

RamzSwap (compcache)




- Provides compressed swap, in memory
 - Uses LZO-compressed RAM based block device as swap disk
- Mainlined in 2.6.33
- See <http://lwn.net/Articles/334649>
- Home page: <http://compcache.googlecode.com/>
- <http://code.google.com/p/compcache/wiki/Performance>
 - Author writes: “Embedded devices: TODO”
 - Someone please test and post their results...



CE Linux Forum

Tracing

- Perf Counters
- Ftrace
- LTTng
- SystemTap





CE Linux Forum

Perf

- Started out as system for accessing hardware performance counters
- Has changed to support gathering tracepoint data
- Has scripting, to support capture and analysis of specific data sets
- Highlight: perf diff – shows difference from one trace to another
- Is increasingly popular with kernel developers



CE Linux Forum

FTrace



- System to provide kernel static and dynamic tracing
- Work recently has focused on trace-cmd and kernel shark
 - Trace-cmd = command line tool for easier trace generation
 - Kernel shark = trace visualization
- See <http://people.redhat.com/srostedt/ftrace-world.odp>



CE Linux Forum

Security

- Security Modules:
 - AppArmor
 - Finally mainlined in 2.6.36 (very probably)
 - Tomoyo
 - SMACK
 - DataState
 - New security module supporting file-content based access controls
 - Is experimental and out-of-tree right now
 - <http://lwn.net/Articles/399577/>





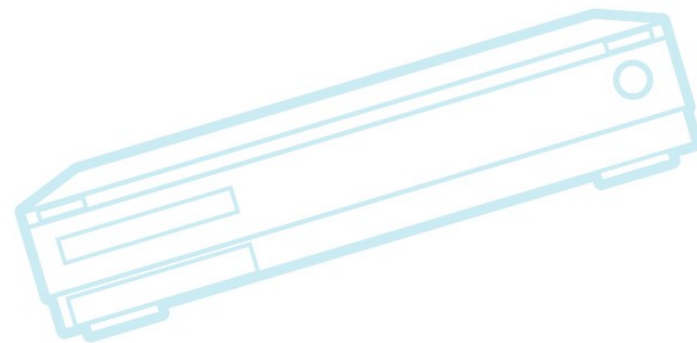



Miscellaneous Stuff

- Device Trees for ARM
 - Device Trees is a mechanism to pass info from bootloader to kernel
 - Supported in X86 and PPC
 - If all platforms supported it, it could allow more uniform device drivers
 - Also, allows for a single binary to run on multiple platforms
 - See Grant Likely's presentation here at LinuxCon
- Linaro
 - New consortium for enhancing Linux technologies for ARM processors



CE Linux Forum

Embedded Distributions









CE Linux Forum

Embedded Distributions

- Meego
- Android
- Ubuntu netbook
- WebOS
- Roll-your-own








CE Linux Forum

Meego!!

- Moblin transferred by Intel to Linux Foundation
- Moblin + Maemo = Meego
- Intel and Nokia joining forces to create single embedded distribution
- Released in March 2010
- Seeing uptake in tablets, in-vehicle systems
 - GENIVI





CE Linux Forum

Android



- Android 2.2 SDK (Froyo) released May 20, 2010
 - Added JIT for Dalvik
- Number of devices increasing
 - Non-phones starting to appear
 - Sony Internet TV (OK – shameless plug)
- Android patches removed from mainline kernel (2.6.33)
 - Patches lived in drivers/staging for a while, but were removed
 - Greg Kroah-Hartman talked about this at ELC
 - http://elinux.org/ELC_2010_Presentations



CE Linux Forum

Ubuntu netbook

- Only embedded distribution derived from desktop distribution
- Canonical and Linaro seem to be working together to refine for embedded use
- Canonical has been pushing Ubuntu for appliances








CE Linux Forum

WebOS

- HP purchased remains of PALM, and is expected to ship WebOS-based devices soon
- Don't know much about it







CE Linux Forum

Roll-your-own



- OpenEmbedded seems to have emerged as leading build system
 - Though there remain lots of options for niche markets
- Increasing trend to use one of the “major” embedded distributions, and customize from there





CE Linux Forum

Contract Work







CE Linux Forum

CELF Contract Work

- Busybox
 - Smemcap applet
 - bootchart applet
- Function Sections
- YAFFS2 mainline effort
- U-Boot ARM enhancements
- RT-preempt tutorial
- Trace format standard
- Kexecboot enhancements
- Flash filesystem testing





CE Linux Forum

What can you do to help?

- Work at top of tree
 - Version gap is still biggest problem with embedded people working in community
 - I know it's difficult
 - If your hardware isn't supported: bug your semiconductor vendor or use a different board for top-of-tree work
- Don't wait for someone else to test new features!!
 - Post results to eLinux wiki
 - Come to next conference and tell us what happened)



What to Test...

- New flash filesystems
 - LogFS, UBIFS, Yaffs2
- Boot-time:
 - Use devfs and measure improvement in boot times
- Size features:
 - Ramzswap on embedded platforms
 - Function sections
 - smem
- New tracing features:
 - trace-cmd, kernel shark, perf scripting
- Runtime power management



CE Linux Forum

Resources

- LWN.net
 - <http://lwn.net/>
 - If you are not a subscriber, please do so
 - This is an invaluable community resource that needs your support
- http://kernelnewbies.org/Linux_2_6_??
- eLinux wiki
 - <http://elinux.org/>
- Linux-embedded mailing list
 - <http://vger.kernel.org/vger-lists.html#linux-embedded>




CE Linux Forum

Embedded developer BOF


- Meeting info:
 - Tonight: 5:30 pm
 - Here in Roppongi Academy Hills 40
 - Quarante D
- Sponsored by CE Linux Forum
- Updates on more embedded topics
- Chance to socialize with embedded developers
- Refreshments and prizes



CE Linux Forum



Thanks!



Keep up the good work!

