



CE Linux Forum

Status of Embedded Linux



Tim Bird - CELF AG Chair



CE Linux Forum

Outline


- Kernel Versions
- Technology Areas
- CELF Contract Work
- Embedded Distributions
- Community
- Conferences
- Resources





CE Linux Forum

Linux Kernel Versions







CE Linux Forum

Kernel Versions

- Linux v2.6.26 – 13 July 2008
- Linux v2.6.27 – 9 Oct 2008
- Linux v2.6.28 – 24 Dec 2008
- Linux v2.6.29 – 23 Mar 2009
- Linux v2.6.30 – 10 June 2009
- Linux v2.6.31-rc3 – 14 July 2009








CE Linux Forum

Linux v2.6.26

- **KGDB**
 - Finally, an in-kernel debugger gets mainlined









CE Linux Forum

Linux v2.6.27

- Ftrace core
- UBIFS
- Bits of Linux-tiny








CE Linux Forum

Linux v2.6.28

- Tracepoints
- Bits of Linux-tiny









CE Linux Forum

Linux v2.6.29

- SquashFS
- BTRFS
- Kernel mode setting
 - Ability to set graphics mode in kernel
- Asynchronous Function Calls









CE Linux Forum

Linux v2.6.30

- TOMOYO security module
- Integrity measurement
- Threaded interrupts
- NILFS









CE Linux Forum

Linux v2.6.31-rc3





- Ftrace filters
- kmemleak





CE Linux Forum

Technology Areas







CE Linux Forum

Technology Areas

- File Systems
- System Size
- Tracing
- Real-time
- Security
- Power Management
- Bootup Time








File Systems

- SquashFS
 - Compressed, read-only FS
 - Mainlined in 2.6.29
 - Was out-of-tree for years
 - CELF contracted author to mainline it
 - <http://www.linux-mag.com/cache/7357/1.html>
- UBIFS
 - General purpose flash filesystem, built on wear-leveling layer above flash
 - Mainlined in 2.6.27
 - See <http://lwn.net/Articles/276025>
- BTRFS
 - Check-pointing log-structured file system
 - Mainlined in 2.6.29, BUT STILL EXPERIMENTAL
 - See <http://www.linux-mag.com/id/7308>



More File Systems

- **NILFS**
 - NILFS = New Implementation of a Log-Structured FS
 - Continuous checkpointing, ability to snapshot
 - Mainlined in 2.6.30
 - See <http://www.nilfs.org/>
- **AXFS**
 - Advanced XIP File System
 - Still hanging around
 - Offered to Numonyx customers
 - Last mainline attempt was August 2008
- **LogFS**
 - Effort seems to have dropped off





CE Linux Forum

File Systems Issues

- Patches of interest:
 - VFAT patent workaround
 - 2 attempts by Andrew Tridgell to work around Microsoft VFAT long-name patent
 - First attempt was controversial, because functionality was lost
 - New approach preserves functionality
 - <http://lwn.net/Articles/339641>
 - VFS-based union mounts
 - See <http://lkml.org/lkml/2009/5/18/289>
- Some log-structured file system is needed for fast mounting
 - Possibly NILFS or BRTFS will fill this role



System Size / Memory



- Linux-tiny
 - CELF contracted with Free-Electrons
 - 8 patches mainlined in Fall by Thomas Petazzoni
- -ffunction_sections
 - Put each function in it's own linker section
 - Allows elimination of unused code
 - Patches sent last July, but I haven't seen any update since then
- Compression
 - LZMA support
 - Support for LZMA kernel image compression (up to 30% better than gzip)
 - Still would like to see generic LZMA support in kernel (for e.g. SquashFS)
 - See <http://www.squashfs-lzma.org/>
- XIP
 - Not much new development (AXFS?)



CE Linux Forum

Compressed swap


- Latest patch is “compcache”
- Uses LZO-compressed RAM based block device as swap disk
- NACKed because MM developer disliked the use of xvmalloc





Compcache results

RAM usage with/without compcache





CE Linux Forum

smem

- Description:
 - New tool to provide information about system memory usage
 - Reports *Proportional Set Size* (PSS), which accounts better for shared pages, and doesn't lie, like RSS
 - Supports visualization (charts)
- Author: Matt Mackall
- Status:
 - Presented at ELC 2009 – is in good shape
 - See <http://www.selenic.com/smem>



CE Linux Forum

Mem_notify patch

- Addition to cgroup system to support notification on low memory
- Patches not really noticed by community
- Some feedback was addressed, and patch was submitted again (last week)
- Work was funded by CELF



CE Linux Forum

Bootup Time



- Bootup time is now a hot topic (due to netbooks)
- Intel 5-second boot on Moblin
 - See <http://lwn.net/Articles/299483>
- Asynchronous function calls
 - Mainlined in 2.6.29
 - See <http://lwn.net/Articles/314808>
- *scripts/bootgraph.pl* for visualization of initcalls
- New Readahead techniques
- Aggressive application optimizations (especially for X)



CE Linux Forum

Bootup Time patches

- devtmpfs
 - A form of devfs again (oh no!)
 - 0.6 second faster boot-time
 - On desktops – who uses dynamic device nodes in embedded?









CE Linux Forum

Tracing

- Ftrace
- LTTng
- SystemTap
- Perf Counters





CE Linux Forum

FTrace



- Is a new system to provide kernel tracing
- Generic framework for adding tracing to the kernel
 - Provides multiple tracers, selectable at runtime
 - Infrastructure for tracepoint definition, data capture (ring buffer), tracer control and trace output
- Core mainlined in 2.6.27
 - More bits coming (generic filtering in 2.6.31)
- See <http://people.redhat.com/srostedt/ftrace-world.odp>



CE Linux Forum

Real-time

- RT-preempt
 - Interrupt threads
 - Mainlined 2.6.30
 - Sleeping Spinlocks
 - Thomas Gleixner and Ingo Molnar got busy with other stuff
 - Maybe a push to mainline this Fall.
- Xenomai
 - Dual-kernel Real-time kernel
 - Successor to Adeos/I-Pipe
 - See <http://www.xenomai.org/>







CE Linux Forum

Power Management

- Wakelocks submitted by Google
 - Kernel PM developers rejected the approach
 - But this opened a discussion of the requirements
 - See <http://lwn.net/Articles/318611>
- I'm not sure what else is going on lately







CE Linux Forum

Audio/Video/Graphics

- **GStreamer**
 - Is de-facto standard for video management for handhelds
- **DirectFB**
 - Continued uptake in TV domain
- **OpenGL ES**
 - De-facto standard for 3D API (/Gaming in embedded?)
- **Clutter (GTK UI Framework)**









CE Linux Forum

Security




- Tomoyo
- SMACK
 - CELF project to analyze SMACK
- Embedded SELinux
 - I haven't heard anything new





Security Modules

- Tomoyo Linux
 - Path-based security module
 - Mainlined in 2.6.30
 - Was a big deal to get a path-based LSM into the kernel
 - See <http://elinux.org/TomoyoLinux>
 - Presentations on using Tomoyo with Android at Japan Technical Jamborees 27 and 28
- Smack
 - Mainlined in 2.6.25 kernel
 - Seems good for embedded (low overhead)
 - CELF started project to evaluate TV use case
 - See <http://lwn.net/Articles/292291>
- Embedded SE Linux
 - Haven't heard much lately







CE Linux Forum

Security – other

- Integrity Measurement
 - Allows kernel to interact with TPM and ensure trusted operation
 - Mainlined in 2.6.30
 - See <http://lwn.net/Articles/137306>
 - Search: “Integrity Measurement Architecture IBM”









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
 - Is a new hot topic
- DLNA summit 2 – November in the Netherlands
 - CELF provided DLNA hardware and specs to OSS developers



CE Linux Forum



Embedded Distributions





Embedded Distributions





- Moblin
 - Intel transferred control to Linux Foundation
- Android
 - Number of devices seem to be exploding
 - Waiting to see consumer uptake
- Maemo
 - Maemo to use Qt (not a big surprise)
 - <http://lwn.net/Articles/341391> (subscriber only)
- OpenWRT
- Chrome
 - Don't really know what's in it yet





CE Linux Forum

Contract Work








CE Linux Forum

CELF Contract Work

- smem
- OOM notifications in cgroups
- SquashFS
- SMACK on TV analysis
- Flash filesystem testing and comparison
- DirectFB mainlining
- Matt Mackall – maintainer work





Contract Work Notes



- SMACK on TV analysis
 - White paper completed this week
 - Should show up on elinux wiki very soon
- Flash filesystem testing
 - Should be available on elinux wiki
- DirectFB mainlining
 - Adding support to Linux kernel for “fusion” exten



CE Linux Forum


Matt Mackall

- Is now in MAINTAINERS file as an embedded Linux maintainer
 - Joining David Woodhouse and Paul Gortmaker
- CELF pays for Matt to do miscellaneous development and advocacy for embedded Linux kernel work








CE Linux Forum



Community







CE Linux Forum

Community

- Need to build embedded linux community
 - Please use linux-embedded list
- MELD embedded linux developer portal
 - Supported by Monta Vista
- eLinux wiki







CE Linux Forum

eLinux Wiki

- Site for embedded Linux developers
- Very good for some technical areas
- Stagnant in some areas
- Promotion
 - Contests
 - Elinux.org working BOF at ELC
- Please use it and post stuff there







CE Linux Forum

Conferences





- Past
 - ELC Europe 2008
 - ELC 2009
 - Videos should be available soon
- Future
 - ELC Europe 2009
 - Japan Linux Symposium
 - ELC 2010
- See <http://www.embeddedlinuxconference.com/>





ELC Europe 2009

- Grenoble, France
- October 15, 16
- Speakers notified – program looks good!





CE Linux Forum

Japan Linux Symposium




- Tokyo, Japan
- October 21-23
- Linus Torvalds is speaking
- Kernel Summit scheduled for October 19-20
- Status:
 - Speakers notified – program shaping up



CE Linux Forum

ELC 2010

- **Current plan:**
 - April 12-14 in San Francisco, CA, USA
 - Co-Located with Linux Foundation Collaboration Summit
 - Call for presentations should go out shortly





CE Linux Forum

Resources



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



CE Linux Forum

Discussion

- How to take control of kernel away from Enterprise and Desktop developers?
 - Classic community answer – donate more patches than them
 - Any other ideas?
 - Distractions, bribes...





CE Linux Forum

Thanks!

- Keep up the good work!

