



# *CE Linux Forum*

## **Status of Embedded Kernel Features**

**(July 2009)**





**Tim Bird - CELF AG Chair**



*CE Linux Forum*

# Outline


- Kernel Versions
- Technology Areas
- Distributions
- Resources





*CE Linux Forum*

# Interesting Stuff in Recent 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
  - Really nice to finally get into kernel
- BTRFS
  - Check-pointing log-structured file system
  - Mainlined in 2.6.29, BUT STILL EXPERIMENTAL
  - See <http://www.linux-mag.com/id/7308>
- NILFS
  - NILFS = New Implementation of a Log-Structured FS
  - Continuous checkpointing, ability to snapshot
  - Mainlined in 2.6.30
  - See <http://www.nilfs.org/>





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 BTRFS will fill this role



# System Size / Memory

- LZMA support
  - Support for LZMA kernel image compression
  - Still would like to see generic LZMA support in kernel (for e.g. SquashFS)
    - See <http://www.squashfs-lzma.org/>
- Compcache
  - Compressed SWAP in RAM
  - Reported at last Jamboree by Kosaki-san
- Mem\_notify patch
  - Also reported at last Jamboree
  - More on this later





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*

# 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/>





# Power Management


- Wakelocks submitted by Google in January
  - Kernel PM developers rejected the approach
  - But at least this resulted in a discussion of the requirements
    - See <http://lwn.net/Articles/318611>
- Re-work of runtime power management
  - Suspending individual devices
  - Will find out more tomorrow in talk by Rafael Wysocki



*CE Linux Forum*

# Security

- 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
- 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
  - Was just a distribution for LinkSys routers
  - But now supports other targets
- Chrome
  - Don't really know what's in it yet





*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*

# Questions and Answers

