



CE Workgroup

Status of Embedded Linux July 2014

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

Best of ...

Resources









CE Workgroup

Kernel Versions

- Pace of versions is consistent and good
- Kernel processes are working well





CE Workgroup

Kernel Versions

- Linux v3.11 – 2 Sep 2013 – 64 days
- Linux v3.12 – 3 Nov 2013 – 62 days
- Linux v3.13 – 19 Jan 2014 – 77 days
- Linux v3.14 – 30 Mar 2014 – 70 days
- Linux v3.15 – 8 Jun 2014 – 70 days
- Linux v3.16-rc6



CE Workgroup

Kernel Versions

- Linux v3.11 – 2 Sep 2013 – 64 days
- Linux v3.12 – 3 Nov 2013 – 62 days
- Linux v3.13 – 19 Jan 2014 – 77 days
- Linux v3.14 – 30 Mar 2014 – 70 days
- Linux v3.15 – 8 Jun 2014 – 70 days
- Linux v3.16-rc6
 - I predict 3.16 on Aug 16 – 69 days



CE Workgroup

Outline



Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff

Resources





CE Workgroup

Bootup Time

- Kernel can be quick (under 1 second)
 - But it takes a lot of work, per product
- Lots of resources available for tuning
 - See http://elinux.org/Boot_Time
 - Good recent presentation:
 - Michael Opdenacker – ELC 2014
 - Update on Boot Time Reduction Techniques with Figures
- More focus recently on user-space
 - Lots of people looking at systemd in embedded



Bootup Time

- Make sure UBI fastmap feature is on, if using UBIFS (available from 3.7 on)
 - Focus on firmware
 - Lots of U-boot optimizations in Michael's talk
 - Avoid copies, use best compression method, simplify scripting, trim down u-boot size, modularize u-boot, etc.
- Really looking forward to ELCE 2014
 - Have several good talks lined up for boot time reduction
- Long-term problem:
 - No automation for boot-time ye



CE Workgroup

File Systems





- Asynchronous startup helps a lot with resume on spinning media (SATA)
 - 10x reduction in resume speed
- Block layer on UBI flash translation layer
 - Read-only for now
 - Optimizations in IO scheduler for SSDs
 - Can now handle millions of IOPs



CE Workgroup

Memory Management

- ION memory allocator
- Now in staging







CE Workgroup

Power Management

- Evolution of power management in Linux
 - Suspend/resume, voltage and frequency scaling, longer sleep (tick reduction), runtime device power management, race-to-sleep (wakelocks/autosleep)
- New stuff starting to get crazy





CE Workgroup

Power Management

- Autosleep
- Power-aware scheduling
 - Big.LITTLE scheduling
 - Some products now shipping with this
 - Support is now mainlined for some processors
 - Baseline power-aware patches seem stalled
- Memory power management
- Full tickless



Power-aware scheduling:

- Small-task packing
 - Try to migrate tasks to allow more CPUs to go idle
- Task placement on mixed `cpu_power` systems
 - Move large tasks to faster CPUs
- Resources:
 - <http://lwn.net/Articles/546664> - overview
 - <http://lwn.net/Articles/552885> - some resistance
 - Ingo Molnar wants to consolidate this power stuff in the scheduler – rather than spread out into `power/cpufreq/cpuidle/scheduler` systems

This is big.LITTLE







CE Workgroup

System Size

- Renewed interest due to IOT
- Projects
 - MicroYocto
 - Linux for microcontrollers
- Kernel size
- Shared Library size





Renewed Interest due to IOT

- Is desirable to run Linux in very constrained devices
- Need to eliminate DRAM for power savings
 - Conversion to SRAM or STMRAM means no current draw when idle
- Can possibly work with on-die RAM
- Back to wanting 2M or less RAM



CE Workgroup

Projects




- Micro-Yocto project
 - by Tom Zanussi at Intel
 - 1.6MB SRAM, 8MB flash
 - Lots of kernel, user-space reductions
 - NET Diet patches, CONFIG_PROC_MIN
 - Link Time Optimization
 - LWIP in extreme case (lightweight IP – user-space IP stack)
 - Their own tracing system (microYocto hash triggers)
- Details in ELC 2014 presentation



CE Workgroup

Projects

- Linux for Microcontrollers
 - By Vitaly Wool at Softprise Consulting
 - 256K SRAM, 2MB NOR flash
 - 2.6.33 kernel, using XIP
- Flash usage:
 - 64K bootloader
 - 900K kernel text
 - 800K root filesystem (RO)
 - 196K config filesystem (RW)





Kernel Size

- Recent work on several different patches
 - NET Diet patches – Andi Kleen
 - Link Time Optimization – Andi Kleen
 - Min proc
 - Configurable syscalls and kernel features
 - `sys_sendfile`, `X86_IOPORT`,
`CONFIG_PTRACE`, `CONFIG_SIGNALS`
 - Resistance to NET Diet patches when they tried to mainline them
 - To be discussed at kernel summit in August
 - Maybe a new “linux-tiny” effort???



Shared Library reduction

- mklibs
 - <http://packages.debian.org/sid/mklibs>
- Cuts down shared libraries to match a specific set of executables
- Useful for large libraries like OpenGL and Qt
- Does not require source
- Warning: May miss dlopened libraries
- Available in Yocto but not Buildroot (as of 2013)



CE Workgroup

Security

- Kernel features
 - Kernel address space layout randomization
 - See <https://lwn.net/Articles/569635/>
- Heartbleed bug
 - Big problem caused by very old bug
 - <http://lwn.net/Articles/593809/>
- Linux Foundation creates fund for critical projects
 - See <http://lwn.net/Articles/595959/>



CE Workgroup

Tracing

- Ktap
 - Dynamic tracing, without the overhead of compiling into a module
 - Adds an interpreter to the kernel
 - Temporarily added to mainline in 3.13 (but subsequently removed)
 - Ingo Molnar objected – needs more work
 - See <https://lwn.net/Articles/572788/>



CE Workgroup

Device Tree





CE Workgroup

Device Tree

- continues its inexorable march
- All new SoC and most driver code needs to be DT-compatible
- New work on DT schemas and validators
 - By Tomasz Figa
 - See ELC 2014 talk “Trees need care: a solution to the device tree validation problem”
- More talks on device tree coming at ELCE 2014
 - Security, dynamic trees, tutorials, under-the-hood documentation



Device tree (cont.)

- Change in maintainership
 - Grant Likely transferred maintainership to others
 - Not enough review of bindings
- Discussion about having device tree be long-lived ABI to kernel
 - Should be usable by other operating systems
 - Maybe move out of kernel repository
- Lots of discussions planned at ARM mini-summit/Kernel Summit
 - Lots of presentations at ELC Europe this year
- See http://elinux.org/Device_Tree



CE Workgroup

Outline

Kernel Versions



Technology Areas

CE Workgroup Projects

Other Stuff

Best of ...

Resources





CE Workgroup

CEWG Contract Work 2014

- Compressed printk messages
- LTSI test framework
- Device-tree documentation
- ARM RO text and data overwrite detection
- Android boot time reduction
- CPU Shielding capability
- Config_numa support for ARM
- More robust UBIFS support



Contract Work Details

- Compressed printk messages
 - Size reduction while retaining messages
 - Contractor: Wolfrom Sang
 - Status: Will report preliminary results at LinuxCon North America in August
- LTSI test framework
 - Enhance Cogent test framework, in Jenkins, for use with LTSI kernel
 - Support board-independent deployment and control system, based on ttc
 - Contractor: Cogent
 - Status: not started



CE Workgroup

Contract Work (cont.)

- Device-tree documentation
 - Better docs needed for developers, sub-system maintainers, and binding reviewers
 - Contractor: Frank Rowand
 - Status: not started
- Overwrite detection for kernel text and read-only data
 - Something similar mainlined in 3.14
 - May not be needed
- Android boot time improvements
 - Work to bring up side stack (such as back-up camera) before rest of Android initializes
 - Contractor: Cogent
 - Status: not started



Contract Work (cont.)





- CPU Shielding capability
 - Ability to isolate a CPU at runtime for realtime work
- Add support for CONFIG_NUMA to ARM
 - To allow for handling some memory regions in a special way - even if memory appears uniform to kernel.
- More robust UBIFS support
 - Try to fix some robustness problems when power is lost or bits flip.



CE Workgroup

CEWG New Projects

- New projects for CE Workgroup, with focused areas
 - Based on Steering Committee meetings in May
- New project areas:
 - Internet of Things
 - Standard Distribution
 - Linux in Social Infrastructure
 - SoC Mainlining
 - Others...









CE Workgroup

CEWG New Projects

- New projects for CE Workgroup, with focused areas
 - Based on Steering Committee meetings in May
- New project areas:
 - Internet of Things
 - Standard Distribution
 - Linux in Social Infrastructure
 - SoC Mainlining
 - Others...









CE Workgroup

Project: Internet of Things

- Identify barriers to Linux use in IOT, and fix issues found
 - Focus on size, at the moment
 - May also focus on security
 - Supporting revival of Linux-tiny project may be part of this effort









CE Workgroup

Project: Standard Distribution

- Goal is to share maintenance burden for a standard embedded distribution
- Already discussed previously









CE Workgroup

Project: Linux in Social Infrastructure

- Identify issues with using Linux-based system for societal infrastructure projects
- Includes security, upgradability, long-term support








CE Workgroup

Project: SoC Mainlining

- Help companies collaborate on upstreaming support for their SoCs
- Identify barriers and overcome with training, best-practices documents, etc.





CE Workgroup

Outline

Kernel Versions



Technology Areas

CE Workgroup Projects

Other Stuff

Best of ...

Resources









CE Workgroup

Other Stuff

- Tools
- Testing Frameworks
- Build Systems
- Distributions
- Wiki









CE Workgroup

Tools

- LLVM
- Kernel is close to supporting LLVM without patches









CE Workgroup

Testing frameworks

- Kernel interest in automated testing increasing
 - Ktest.pl
 - Kevin Hillman automated test lab
- Linaro Lava
- Autotest






CE Workgroup

Build Systems



- Yocto project
 - Toaster project – big push to web-based interface
 - Micro-yocto
 - Tizen
- Buildroot
 - New: license compliance, rootfs overlays, eclipse plugin for toolchains and sdk
 - Defconfigs for lots of popular dev boards
 - Package dependency graphing
 - Build-time graphing









CE Workgroup

Distributions

- Poky
- Ångström
- MicroYocto
- Tizen
- Android
- Possible standard embedded distribution







CE Workgroup

▮ Poky

- ▮ Default distribution used by Yocto Project
- ▮ YP keeps trying to avoid shipping a distribution
 - ▮ Poky should be considered a sample distro, but not suitable for production (?)









CE Workgroup

Ångström

- Mature package-feed based embedded distro
 - Originally focused on handheld devices (e.g. Sharp Zaurus)
- Developed originally under OpenEmbedded
- Is shipped with a number of development boards









CE Workgroup

MicroYocto

- Focused on IOT requirements
 - Has special features for tiny systems
 - Kernel patches
 - Special tracing for dynamic memory usage monitoring
 - Currently focused on size and networking aspects of IOT









CE Workgroup

Tizen

- Descended from Meego, which descended from Maemo and Moblin
- Getting more widespread usage
 - Replacing Bada at Samsung
 - Shipping in TVs, phones and increasingly targeted at automotive
- Can be built with Yocto Project
- May be a serious competitor as general embedded distribution







CE Workgroup

Android

- Is increasingly used in non-mobile devices
 - See ABS 2014 talk by Gary Bisson
 - Google “cyborgstack” and “headless android”
- KitKat focused on size issues
 - Target devices with 512MB of space, to drop Android into low-end phones
- “L” release has lots of changes:
 - Project Volta – for power reduction
 - “Material Design” - 3d look with “floating” widgets
 - ART – Android Runtime
 - Ahead-of-time compiler
 - *Goodbye Dalvik*








CE Workgroup

Android

- Is increasingly used in non-mobile devices
 - See ABS 2014 talk by Gary Bisson
 - Google “cyborgstack” and “headless android”
- KitKat focused on size issues
 - Target devices with 512MB of space, to drop Android into low-end phones
- “L” release has lots of changes:
 - Project Volta – for power reduction
 - “Material Design” - 3d look with “floating” widgets
 - ART – Android Runtime
 - Ahead-of-time compiler
 - *Goodbye Dalvik*









CE Workgroup

Possible standard embedded distribution

- CEWG project to create and maintain a standard distribution for the embedded industry
- Just getting started with requirements definition
- May select an existing distribution
- If you are interested, please contact
 - Yoshitake Kobayashi





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
- Working on wiki projects:
 - Video transcription project



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_3.??
- eLinux wiki - <http://elinux.org/>
 - Especially <http://elinux.org/Events> for slides
- Celinix-dev mailing list



CE Workgroup



Thanks!









CE Workgroup

Extra Slides

- The following slides are just for reference, for embedded-related features introduced in recent kernel versions





CE Workgroup

Linux 3.12




- Full-system idle detection
 - Tricky rcu-based implementation to allow for fast indication of individual CPU idleness (using per-cpu variable), AND fast detection of global CPU idleness (single global variable)
- New cpu-idle driver that builds on multi-cluster power management
 - I.e. Getting closer to support for “big.LITTLE” CPU scheduling
- Lots of device drivers converting over to device tree
 - More on this later



CE Workgroup

Linux 3.13

- ARM big.LITTLE switcher support
- Scalable block layer for better SSD performance
- Power capping framework
- Ktap was put in, then taken out







CE Workgroup

Linux 3.14

- Kernel Address space layout randomization
- ION memory allocator in staging
- ARM text and RODATA protection
- 6lowpan emulation for BLE
- Zram compressed swap moved out of staging
 - Reports are that it is used in TVs and some Android devices
 - See <http://git.kernel.org/cgiit/linux/kernel/git/torvalds/linux.git/commit/?id=cd67e10ac6997c6d1e1504e3c111b693bfdbc148>









CE Workgroup

Linux 3.15

- UBI flash translation now has block device layer support
 - Read-only for now
- More support for LLVM compiler suite
 - Not finished yet, but getting closer








CE Workgroup

Linux 3.16

- Power-aware scheduling
- `decode_stacktrace.sh` – convert offsets in a stack trace to filenames and line numbers





CE Workgroup

Linux 3.13

