CE Workgroup

Status of Embedded Linux

June 2017

Tim Bird

Architecture Group Chair

LF Core Embedded Linux Project

Nature of this talk…

- Quick overview of lots of embedded topics
- A springboard for further research
  - If you see something interesting, you have a link or something to search for

Outline

Kernel Versions
Technology Areas
CE Workgroup Projects
Other Stuff
Resources

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff

Resources

Kernel Versions

- Linux v4.7 – 24 July 2016 – 70 days
- Linux v4.8 – 2 Oct 2016 – 70 days
- Linux v4.9 – 11 Dec 2016 – 70 days
- Linux v4.10 – 19 Feb 2017 – 70 days
- Linux v4.11 – 30 Apr 2017 – 70 days
- Linux v4.12-rc5 -
  - v4.12 expected on 9 July, 2017

Linux 4.7

- Schedutil frequency governor
  - Use the load calculated by the scheduler instead of the average load over past little while
  - See http://lwn.net/Articles/682391/
- VFS layer can iterate through directories in parallel
- Ability to attach BPF programs to tracepoints
- Ftrace histogram triggers
  - Can tell tracer to accumulate events into buckets and give results, via the sysfs interface
- Android sync_file feature moved from staging
  - Support for explicit buffer fencing

Linux 4.8

- New kernel documentation system
- New pseudo-random number generator
  - See https://lwn.net/Articles/686033/
- ARM64 support for kexec and kprobes
- New timer wheel implementation
  - https://lwn.net/Articles/646950/
  - Better performance:
    - No more cascade operations
    - Quick determination of next timout
  - Automatically coalesces longer timeouts
  - Long timeouts have reduced resolution

Linux 4.9

- Virtually mapped kernel stacks
  - http://lwn.net/Articles/692953/
  - Allows to detect stack overruns
  - Cleans up kernel code, faster process creation
  - Only on x86, for now
- Greybus - https://lwn.net/Articles/715955/
- Timed samples for eBPF
- Modversions deprecated
  - See https://lwn.net/Articles/707520/

Linux 4.10

- Perf sched timehist
- Hybrid block polling
  - Supports polling for block I/O, but with a short delay (estimated) before the polling starts
    - Improves performance by queuing blocks as soon as device is ready (via polling)
    - Uses less CPU than full polling
- Support for ARM SoCs:
  - Huawei, Allwinner, Marvel, Renesas
- Posix timers are configurable
- Initramfs compression method is selectable
- New interface for system sleep state selection
  - /sys/power/mem_sleep
- UBIFS support for encryption

Linux 4.11

- New kernel refcount API
- TinyDRM subsystem added
- New statx() system call
  - https://lwn.net/Articles/707602/
  - 2038-safe time values
  - Mask of fields to obtain (for efficiency)
- Sched.h refactoring
  - Non-mainline code: watch out!

Linux 4.12 (expected)

- BFQ and Kyber block I/O schedulers
- Minitty prep work
  - Not full minitty implementation yet
- Proper support for USB type-C connectors
- AnalyzeBoot tool
  - Reads dmesg (and possibly ftrace log) and produces html graph of boot events
  - Part of Intel pm-graph tools project
    - https://github.com/01org/pm-graph
  - See tools/power/pm-graph/analyze_boot.py

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff

Resources

Bootup Time

- No new work in kernel, that I'm aware of
- Analyze_boot tool – new in in 4.12
- Some good previous talks:
  - ELCE 2014 - *12 Lessons Learnt in Boot Time Reduction* by Andrew Murray
  - ELC 2015 - *Fastboot Tools and Techniques* by John Mehaffey
- Android boot time ideas
  - ELC 2017 – *Improving the bootup speed of AOSP* – Bernhard Rosenkranzer

Bootup ideas from Bernhard

- Two approaches:
  - Improve cold boot
  - Enhance suspend/resume
- Areas analyzed for cold boot:
  - Package Manager scanning
  - Java class preloading
  - PM: force high CPU frequency during boot
  - IO: read-ahead, kernel compression, squashfs
  - Kernel modules – defer modules until later
  - Library and compiler optimizations

Device Tree

- Device Tree Overlays
  - Allow plugin-boards to be configured at runtime
  - Session at ELC 2016 by Pantellis
  - Not mainlined yet? – expected in 4.11?
- Device Tree validation
  - Schema for binding language, validator for bindings and for device tree data
  - This work stalled
- Updated Device Tree specification
  - Being discussed
  - Want to update material and make it more available
- See http://elinux.org/Device_tree_plumbers_2016_etherpad
  - And ELC 2017 Device Tree BOF – Frank Rowand

Graphics

- TinyDRM
  - Provides graphic support for small simple displays (eg displays over i2C or SPI)
  - Hope to replace framebuffer drivers over time
  - See https://www.phoronix.com/scan.php?page=news_item&px=TinyDRM-Patches-Posted
- GPU support:
  - ARM mali drivers status update
    - https://lwn.net/Articles/716600/
- Presentation
  - ELC 2017 *What Can Vulkan do for You?* - by Jason Ekstrand

File Systems

- UBIFS support for encryption (in 4.11)
- IO scheduling for solid state storage
- LightNVM
  - Software control of flash-translation layer
  - https://lwn.net/Articles/641247

Networking

- Bluetooth:
  - Bluetooth 5.0
    - Most features are on BLE codebase
      - Only 1 for “BL classic”
  - 800% data throughput increase
  - 4 times the range
  - Coexistence with wireless
    - Better error correction to handle noisy environments

Power Management

- New interface for system sleep state selection (in 4.10)
  - /sys/power/mem_sleep
- Operating-System-Directed Power-Management Summit
  - https://lwn.net/Articles/721573/
  - Energy-aware scheduling
  - A collection of scheduling talks that will make your head spin

Real Time

- SCHED_DEADLINE
  - ELC 2017 - *SCHED_DEADLINE: It's Alive* - by Juri Lelli
    - Energy Aware Scheduler support
    - Bandwidth reclaiming
      - Temporarily allow a task to exceed it's bandwidth, if no other process' deadline suffers
    - Support for Frequency scaling
    - Group scheduling
- Presentations:
  - ELC 2017 *Effectively Measure and Reduce Kernel Latencies for Real-time Constraints* – By Jim Huang
  - ELC 2017 *Real-Time Linux on Embedded Multicore Processors* – by Andres Ehmanns

Security

- Kernel hardening
  - http://kernsec.org/wiki/index.php/Kernel_Self_Protection_Project
    - Rare_write infrastructure
      - Keep some code and data read-only most of the time
      - https://lwn.net/Articles/724319/
  - GCC plugins for kernel security
    - Kernexec
      - Prevent kernel from executing user-space code
    - Structleak (mainlined in 4.11)
      - Zero out kernel structures passed to user space, under some conditions
    - See https://lwn.net/Articles/712161/
    - Randstruct
      - Randomize C structure layout
      - See https://lwn.net/Articles/722293/

Security Presentations

- ELC 2017 *Securing Embedded Linux Systems with TPM 2.0* – by Philip Tricca

System Size

- Initramfs compression method is selectable
- Nicolas Pitre work
- Configurable POSIX timers – in v4.10
  - https://lwn.net/Articles/701095/
- Mini TTY
  - Smaller implementation of TTY subsystem, for embedded
  - Saves about 38K
  - https://lwn.net/Articles/721074/
  - People wanted refactoring of full-size TTY instead of new small implementation, but Nicolas said that wasn't feasible

System Size (cont.)

- Shrinking the scheduler
  - Drops features and eliminates realtime and deadline scheduler classes
  - Saves about 20k
  - https://lwn.net/Articles/725376/
  - Lots of resistance to this
  - Code complexity increase is not worth saving 20k (according to Ingo Molnar)
  - Disagreement on whether Linux should support computers with sub-1MB memory

Size Presentations

- LinuxCon North America: *Running Linux on Tiny Peripherals* – by Marcel Holtmann
  - Got Linux to around 1MB for IOT sensor project
- ELC 2017 *Embedded Linux Size Reduction Techniques* – By Michael Opdenacker
  - Very good overview of existing reduction techniques and status
    - Formal Tinification project is stalled
    - Toybox and musl (smaller libc) are worth looking at

Testing

- Kselftest
- Fuego
- Kernelci.org
- LAVA V2

Kselftest

- Unit test system inside kernel source tree
- Recent work:
  - Lots more regression tests (preferred place for syscall compatibility/regression tests (over LTP)
  - Converting to TAP (Test Anything Protocol) for test output

Fuego

- New Test Framework for collaborating on tests and test infrastructure for Linux
- V1.1 features (April 2017)
  - Upgrade to latest Jenkins
  - Test script refactoring
  - Fuego container directory layout change
  - About 40 new tests
- V1.2 plans (RC in July)
  - Unified output format
    - Convert all test results to JSON
  - Support LAVA as a transport & board manager
  - Test dependency system
    - Board dynamic variables

Kernelci.org

- Place to get free build/boot testing for your board
  - Builds 126 trees continuously, then reports any errors
- http://kernelci.org
- Presentations:
  - ELC and ELCE 2016 – by Kevin Hilman
  - Linaro Connect:
    - Kernelci and lava update - See https://lwn.net/Articles/716600/
- The most successful public, distributed build and test system for Linux, in the world!

LAVA

- Linaro Automation and Validation Architecture
- V2
  - Job files now use Jinja2 templates
    - Was previously hand-written JSON
  - Jobs are run asynchronously, without polling,
  - ZeroMQ is used for communications.
  - ReactOBus is used to run jobs from messages.
  - Requires more explicit board configuration

Toolchains

- LLVM 4.0.0 is released
  - Some code size improvements from optimizations (GVNHoist)
  - Experimental support for LLVM coroutines
  - https://lwn.net/Articles/716979/
- Presentations:
  - ELC 2017 - *GCC/Clang Optimizations for Embedded Linux* – by Khem Raj

Tracing

- More perf tools (both in 4.10):
  - perf sched timehist
    - Analysis of scheduling events
  - perf c2c
    - Cacheline contention analysis
- Presentations:
  - ELC 2017 *Dynamic Tracing Tools on ARM/AArch64 Platform: Updates and Challenges* - by Hiroyuki Ishii
    - Great overview

Miscellaneous

- Year 2038 status:
  - 3 areas of work:
    - Converting all 32-bit timestamps to 64-bit in the kernel
      - e.g. New statx() system call
      - Many patches are in-progress (vfs layer, v4l, device-mapper, input subsystem)
    - C libraries
      - Lots of work in glibc to make everything backwards compatible
        - Even programs built with 32-bit timestamps should work
    - Distribution builds
  - See https://lwn.net/Articles/717076/

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff

Resources

Projects and initiatives

- Shared Embedded Distribution
- LTSI
- Fuego
- eLinux wiki

Shared Embedded Distribution

- Goals
  - Create an industry-supported distribution of embedded Linux
    - Main goal is very long term support (15 years)
- Status
  - Toshiba has created Yocto layer meta-Debian
  - Presented at ELCE, ELC, and LCJ
- Next steps
  - Improve coordination with Debian community

Long Term Support Initiative

- LTSI 4.9 is current LTSI kernel
  - Work is in progress on next release
- GregKH said
  - Expected delivery date: Sep 2017
  - Converting to upstream-first policy
- Presentation:
  - ELC 2017 *Using Linux as Long Term Working with the Community* – by Tsugikazu Shibata

Fuego - Linux Test Framework

- Working on lots of issues:
  - Command line tool
  - Test packaging
  - LAVA integration
  - Serial console transport
- Presentation:
  - ELC 2017 *BoF: Fuego Status and Roadmap* – by Tim Bird

eLinux wiki

- http://elinux.org
  - Web site dedicated to information for embedded Linux developers
    - The wikipedia of embedded linux!
- Hundreds of pages covering numerous topic areas: bootup time, realtime, security, power management, flash filesystem, toolchain, editors
- Lots of pages in last few years about low-cost development boards
- Please use and add to site

Outline

Kernel Versions
Technology Areas
CE Workgroup Projects
Other Stuff
Resources

Trade Associations

- Linaro still doing lots of great work
  - Lava v2 and kernelci
  - Now promoting Zephyr
  - Linaro Connect consistently has useful material
- Linux Foundation
  - Microsoft has joined the Linux Foundation as a platinum member
  - CE Workgroup officially changed its name to "Core Embedded Linux Project"

Conferences

- ELC 2017
  - Lots of great sessions
  - See: http://elinux.org/ELC_2017_Presentations
- Open Source Summit Japan
  - May 31-June 2, Tokyo
- Embedded Linux Conference Europe
  - October 23-25, Prague, Czech Republic
- Embedded Linux Conference
  - March 12-14, Portland, Oregon, USA
- Japan Jamborees
  - Continuing

ELC 2017 thoughts

- Linus and Dirk fireside chat
  - 4.10 release was calm
    - 4.9 was a bit bigger due to LTS pre-announcement
    - Linus thinks is healthier to not push things based on a deadline, but 4.9 wasn't too bad
  - Even after all these years, we see changes to core files
  - Linus said that Linux is general-purpose, so may not be appropriate for the lowest-footprint device
    - I feel vindicated

Outline

Kernel Versions
Technology Areas
CE Workgroup Projects
Other Stuff
Resources

Resources

- LWN.net
  - http://lwn.net/
  - If you are not subscribed, please do so
- Kernel Newbies
  - http://kernelnewbies.org/Linux_[34].?
- eLinux wiki - http://elinux.org/
  - Especially http://elinux.org/Events for slides
- Celinux-dev mailing list

Thanks!