CE Workgroup

Embedded Linux Community Update

March 2020

Tim Bird

Sr. Staff Software Engineer, Sony Electronics

Architecture Group Chair, 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
- Not comprehensive!
  - Just stuff that I saw

Outline

Linux Kernel
Technology Areas
Conferences
Industry News
Resources

Outline

Linux Kernel

Technology Areas

Conferences

Industry News

Resources

Kernel Versions

- Linux v5.1 – 5 May 2019 – 63 days
- Linux v5.2 – 7 Jul 2019 – 63 days
- Linux v5.3 – 15 Sep 2019 – 70 days
- Linux v5.4 – 24 Nov 2019 – 70 days
- Linux v5.5 – 26 Jan 2020 – 63 days
- Current kernel = v5.6-rc4
  - Merge window is closed – no new features for 5.6
  - Expect 5.6 kernel on March 29 or April 5

Linux v5.1

- Deprecating support for a.out binaries
  - https://www.phoronix.com/scan.php?page=news_item&px=Linux-Dropping-A.Out
  - Use ELF from now on
- Lots of DRM changes
  - https://www.phoronix.com/scan.php?page=news_item&px=Linux-5.1-DRM-Changes
- More Y2038 work
  - More syscalls with 64-bit time values
    - See https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=b1b988a6a035 for a list of new syscalls (20 of them)

Linux v5.1 (cont.)

- Energy-Aware scheduling
  - New sysctl knob (/sys/kernel/sched_energy_aware)
  - Documentation/scheduler/sched-energy.txt
  - Documentation/power/energy-model.txt
- Improved idle behavior in tickless systems
  - Added timer-events oriented (TEO) CPU-idle governor
    - Uses timer interrupts timing instead of device interrupt timing for predicting next wake-up
  - See https://lwn.net/Articles/775618/

Linux v5.1 (cont1.)

- Modification to memfd for Android use case
  - Add F_SEAL_FUTURE_WRITE operation for memfd regions
    - Caller can continue to write to region, but others can't
  - Want to eliminate use of ashmem (legacy Android memory manager)
- F2FS has a new mode bit that disables copy-on-write behavior for a file (F2FS_NOCOW_FL)

Linux v5.2

- ext4 supports case-insensitive lookups
- New system calls for filesystem mounting
  - See https://lwn.net/Articles/759499/
- Support for ARM Mali GPUS
- New “mitigations=“ command-line option to control speculative execution features
- Improved support for gcc ‘-Wimplicit-fallthrough’
- Lots of BPF improvements
- Pressure stall monitors added

Pressure stall monitors

- Allow user-space to detect and respond quickly to memory pressure
- Monitor writes a stall notification specification to /proc/pressure/memory
  - Indicates to the kernel what frequency to check for stalls (which can be as little as .5 seconds)
- Monitor receives stall notification events (via poll())
- Android can use this to detect memory pressure and kill processes before the device becomes sluggish
  - See https://lwn.net/Articles/775971/

Linux v5.3

- New pidfd feature – to handle pid reuse
- Scheduler utilization clamping
  - (see next slide)
- 0.0.0/8 IPv4 address support
  - Allows 16 million new IPv4 addresses
- Added CONFIG_PREEMPT_RT
  - But not the final code yet
- init_on_alloc and init_on_free boot options
  - pre/post-initialize memory from heap allocations
- See https://kernelnewbies.org/Linux_5.3

Scheduler utilization clamping

- Extension to Energy Aware Scheduling
- Allows specifying minimum or maximum frequency for a process
- Can clamp user-visible (foreground) tasks to high minimum frequency
- Can clamp background tasks to low maximum frequency
- Helps conserve power while still keeping responsiveness
- See https://lwn.net/Articles/762043/

Linux v5.4

- EROFS graduated from staging
- exFAT added to staging
- fs-verity feature added
- boot-time entropy fix
  - Fix for commit that was reverted in late 5.3
  - Prevents get_random() from blocking on boot
  - Implementation based on clock jitter, by Linus himself
  - See https://lwn.net/Articles/802360/

Linux v5.5

- ARM64 has full support for ftrace
- MIPS supports kcov – coverage analysis
- KUnit testing framework added
- CPU scheduler's load-balancing algorithm was replaced
  - Follow-on to PELT (Per Entity Load Tracking) work
    - See https://lwn.net/Articles/732021/ for PELT info
- sysctl() system call was removed
  - Use /proc/sys/... instead

Linux v5.6

- WireGuard VPN feature added to kernel
- Work on 2038 issues for ALSA
  - New 64-bit structure for some operations
- Mechanism to disable SELinux at module load time is deprecated (system runtime)
  - Plan is to add a painful delay (increasing with each kernel release) in order to discourage future use
- Bootconfig tool to add super-long command-lines arguments to kernel
- F2FS gained support for compression

Interesting stats

- 285 new contributors in 5.5
  - Developers who have never contributed before
- The top 3 “reported-by” lines for bugfixes are for automated testing systems

Test system	Reported-bys	Percent
Hulk Robot	164	15.7%
Syzbot	125	12.0%
kbuild test robot	102	9.8%

- At least 14% of commits are fixes for bugs
- See https://lwn.net/Articles/798505/

- Table data from: https://lwn.net/Articles/810639/

More stats

- 90% of kernel developers are paid by their employer to work on the kernel
  - But there are areas that no one is paid to work on (dedicated)
    - E.g. There is no paid documentation person or team
  - There's still a lot of anxiety about unsupported or under-resourced areas of the kernel

Outline

Linux Kernel

Technology Areas

Conferences

Industry News

Resources

Technology Areas

- Audio
- BPF
- Documentation
- Filesystems
- Graphics
- Languages
- Memory
- Security
- Testing
- Tracing
- Toolchains
- Tools
- Build Systems

Audio

- Work in ALSA drivers to support new 64-bit structures
  - One benefit is cleanup of fields to avoid year-2038 problems
- New patch proposed for Qualcomm protection domain restart helpers
  - Feature specific to qualcomm SoCs
  - Allows AVS Audio to run in a separate address space
    - Can crash&recover without disrupting other domains

BPF

- Replacing kernel operations structures
  - Ability to replace a structure of function pointers
  - Can load new functions as BPF modules
  - Can create a structure in user-space to override the current in-kernel structure
  - Use a BPF program to replace the in-kernel structure
- Caveats:
  - Must be a structure pre-designated to support this kind of replacement
  - Is only used for TCP congestion-control algorithms (for now)
  - See https://lwn.net/Articles/811631/

BPF (cont.)

- Book Review: BPF Performance Tools
  - Looks like a nice book, and lots of neat tools
  - A sign that BPF should probably be taken seriously
  - See https://lwn.net/Articles/813114/
- Recent discussions about BPF integration with LSM for kernel runtime security instrumentation (KRSI) patch set
  - KRSI needs high performance (wants a special mechanism to replace crypto calls with static jumps)
  - LSM wants general mechanisms (no special cases for BPF)
  - See https://lwn.net/Articles/813261

Documentation

- Good article on how to contribute to kernel documentation:
  - https://lwn.net/Articles/810404/
- Specific Tasks:
  - Remove all warnings
    - Specific tips on how to address changes required in kerneldoc messages
  - Add unreferenced kerneldoc info
    - Use scripts/find-unused-docs.sh
  - Fix typos
    - This is a good place to start to learn process
    - Leave some typos for other beginners

Documentation Tasks (cont.)

- Specific Tasks (cont.):
  - Remove, fix or tag outdated documentation
    - Lots of old stuff
  - Organize content into better groups
  - Improve HTML look
    - Improve the style-sheet for HTML output
  - Make rst2pdf tool work with kernel docs
  - Write more documentation
    - Still lots of undocumented areas
- Would be nice to have automated testing to indicated “health status” of kernel docs
- See this video, from kernel recipes 2019:
  - https://www.youtube.com/watch?v=1LuAIUKqKDk

Side note on Japanese Docs

- Japanese translations of some docs are available
  - See Documentation/translations/ja_JP
- There was an effort in 2007 to do Japanese translation
- I'm not sure what happened to that effort
  - Wesbsite http://www.linux.or.jp/JF/ no longer works
- You might ask Tsugikazu Shibata about the status of the docs. He did some work in the Japanese docs in 2017.

Filesystems

- F2FS gets compression (already mentioned)
- New io_uring system for asynchronous I/O
  - Already have AIO system – this one is better
  - A ring buffer is shared between kernel and user-space
  - User-space can stuff opcodes (commands) into the buffer, and the kernel can execute them, without any syscalls
  - More complex operations are envisioned using BPF
    - That's under heavy discussion
  - See https://lwn.net/Articles/810414/

Graphics

- Panfrost driver for ARM Mali GPUs
  - Open Source driver
  - Some support by ARM for development work
  - Support for Mali T720, T820 and T860
  - Support for normal desktops (including Wayland)
  - Currently only support for OpenGL ES <= 2.0
  - Does not support Vulkan yet
- Source: LinuxConfAU talk by Robert Foss
  - https://linux.conf.au/schedule/presentation/68/

Graphics (cont.)

- “Everything Awesome about GPU Drivers”
  - LinuxConf AU talk by Daniel Vetter
  - Good overview of state of Linux kernel graphics stack
  - Source: https://linux.conf.au/schedule/presentation/86/

Languages

- Python

Python

- Python 2 is now unsupported
  - (since Jan 1, 2020)

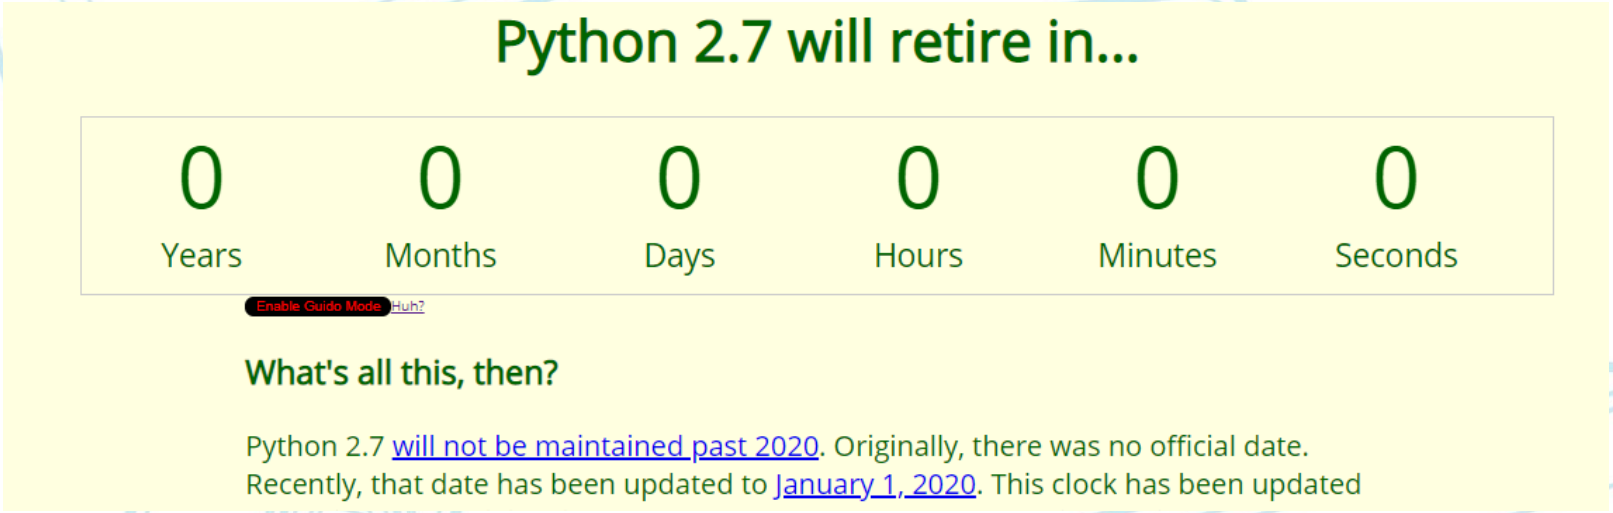

- There are lots of resource to help with conversion to Python 3
  - “2to3” tool, online guides, etc.

Python in embedded

- Snek by Keith Packard
  - New language - subset of Python
  - Made to run on Arduinos and other low-end systems
  - Can squeeze down to about 32kB without math functions
  - See https://lwn.net/Articles/810201/
- MicroPython = decent subset of python for embedded
  - 100s of kBs
  - There are variations like CircuitPython

Memory

- Facebook contributing OOMD to systemd
  - Facebook's OOMD is a replacement for Linux OOM handler
    - Faster and more configurable
  - Plan to contribute to systemd so it's automatically integrated into most Linux distributions
  - Probably take about 1 year to show up
  - See https://www.phoronix.com/scan.php?page=news_item&px=Systemd-Facebook-OOMD

Security

- Control-flow integrity for the kernel
  - New patch set for kernel hardening
    - Not mainlined yet – it's in development
  - Reduce ability to redirect code execution
  - Validate function pointer or return address on heap or stack
  - Compiler can collect function pointers into tables that are verified before being used
  - Arm has support for "shadow stacks"
    - Only stores return addresses
    - Shadow stack location is kept secret to prevent tampering
  - Requires Clang (LLVM 10) support for latest features
- See https://lwn.net/Articles/810077/

Security (cont.)

- Kernel Runtime Security Instrumentation (KRSI)
  - Allows to mitigate a security attack while it's in progress
  - Provides flexible hook for monitoring and mitigation
  - Implemented as LSM that can run eBPF programs
  - See https://lwn.net/Articles/798157/
  - Recent news:
    - Disagreements between LSM and BPF developers over ways to enhance performance
      - See BPF section

WireGuard VPN tunnel

- Faster and simpler than Ipsec and OpenVPN
- Aim is to be as easy to use as SSH
  - Simple generation of public/private key pairs
  - Similar mode of distribution for public keys
- Uses Linux ip commands to set up tunnel
- Allows roaming by both sides of tunnel
- Uses state-of-the-art cryptography
  - High-speed cryptography, suitable for embedded
- Amenable to security audits
  - Due to much simpler code base

Testing

- Kselftest
- Kunit test framework
- Linux Test Project
- Update on testing collaboration

Kselftest

- Sub-system test code inside kernel source tree
- Recent work:
  - Working on making sure tests cross-compile and install
    - Bpf is particularly difficult – it often requires the very latest (unreleased) LLVM compiler
  - Goal is to run with KernelCI
  - Also trying to reduce output differences
    - Some tests never adopted TAP format, like they were supposed to

Kunit test framework

- Set of patches for kernel unit testing
  - Accepted in v5.5
- Lots of changes recently
  - Allow tests to run on kernel boot or on module load
    - Originally tests were configured to run on a UML kernel (not on real hardware)
  - Changes to put log output into debugfs
  - Addition of KASAN (kernel address sanitizer) tests to Kunit

Linux Test Project (LTP)

- Adding syzkaller reproducers to LTP
  - Syskaller is a fuzzer that creates C programs to find kernel oopses
  - It saves all the ones what “work”
    - Called ‘reproducers’
    - There are thousands of them
  - LTP and syskaller developers working to add the reproducers to LTP
  - Result will be a nice regression test to make sure that the bug doesn’t come back in

Testing Collaboration

- "The magical fantasy land of Linux kernel testing"
  - Presentation by Russell Currey at LinuxConf AU
- Good overview of kernel testing landscape
  - Test matrix is very large
  - Lots of different stakeholders
    - Developers, maintainer, distros, end-users
  - People are starting to work together, but it's slow progress
  - First collaboration=unified results format
  - Source https://linux.conf.au/schedule/presentation/106/

Tracing

- New Bootconfig system
  - Extra boot configuration
  - Allows passing a large set of options to the kernel during boot
    - Was not a good fit for device tree
  - Passes a tree-structured key-value list
  - Data is loaded with initrd
  - Used primarily to pass kernel command line items for ftrace and early tracing
  - Mainlined in 5.6
  - See https://lwn.net/Articles/806002/

Tracing (cont.)

- Babeltrace 2.0 released
  - New version of the babeltrace trace manipulation toolkit
  - For viewing, converting, transforming, and analyzing traces
  - See https://lwn.net/Articles/810395/
  - And https://babeltrace.org

Kernel tracing overview

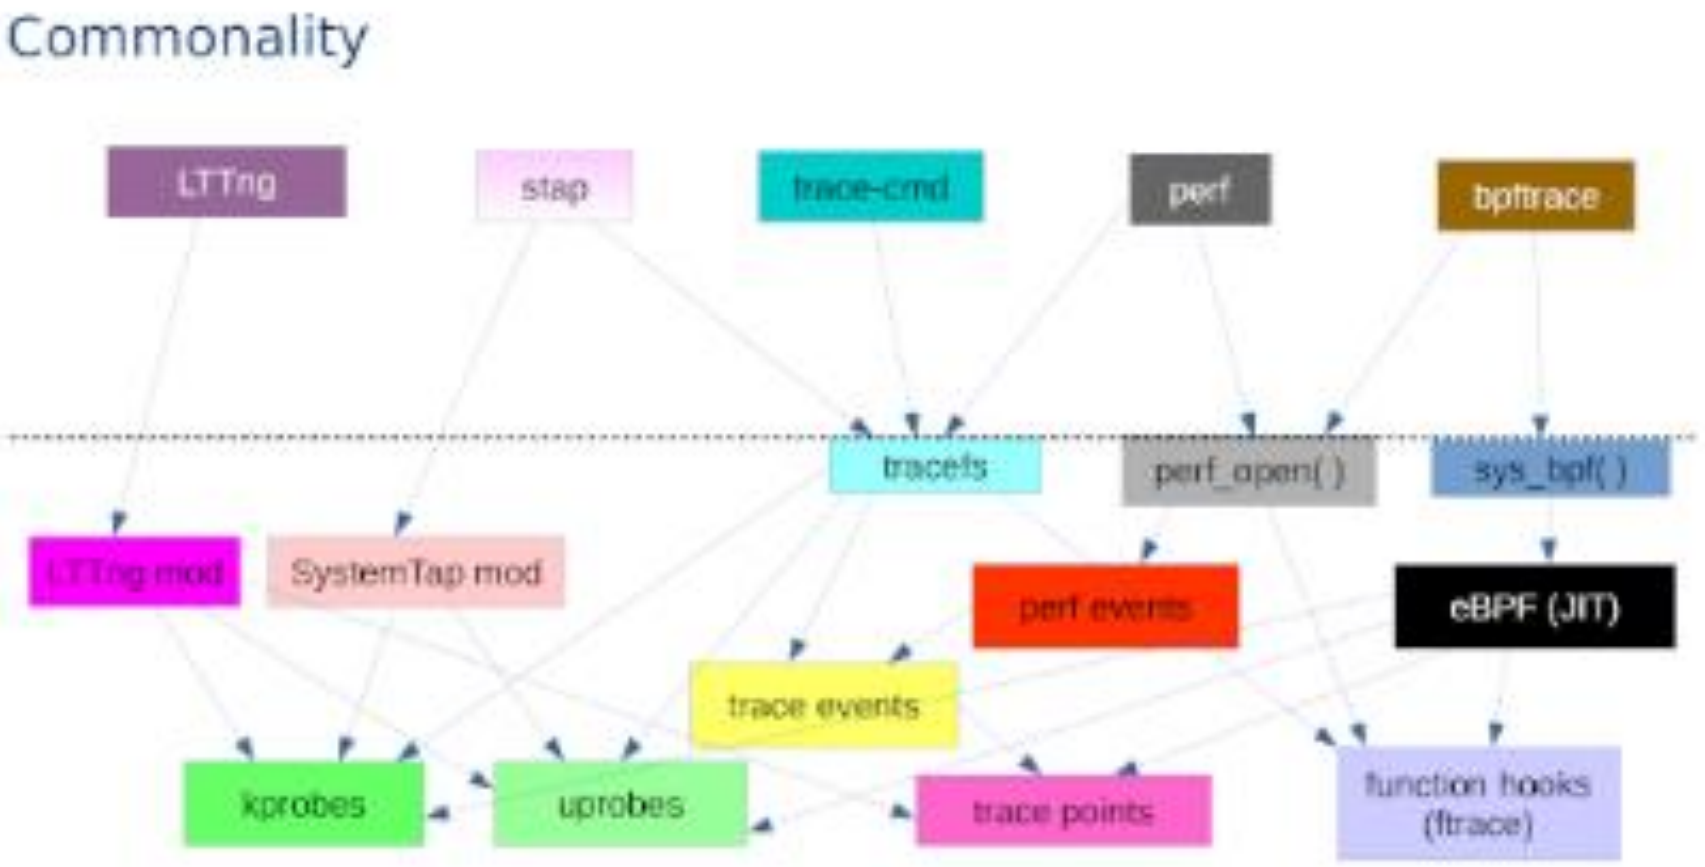

Toolchains

- Static analysis framework for GCC
  - Analyses intermediate representation of code
  - Uses a plugin architecture, for now
  - 2 facility categories for now:
    - memory allocation errors
    - file handling errors
  - Adds meta-data to the diagnostic message
    - Lots of information about the problem
    - e.g. Can indicate the Common Weakness Enumeration (CWE) entry for a problem
  - See https://lwn.net/Articles/806099/

Tools

- Continuing efforts to create tools for bridging gap between git and e-mail
  - New tool: get-lore-mbox
    - Can download a thread related to message id into a local .mbox format (mail archive)
    - Can download just the patch in the thread
    - Can automatically add tags (from different e-mails in the thread) to the patch
      - E.g. reported-by, acked-by, tested-by, etc.
    - Very handy for retrieving patches that were mangled by your email system
    - See https://lwn.net/Articles/811528/
- More work is going on to make additional tools
  - https://github.com/gitgitgadget/gitgitgadget
    - Converts from github pull request to e-mail patches

Tools (cont.)

- New patch attestation scheme
  - Allows users to sign patch messages send to email lists
  - New tool: attest-patches
  - Protects against patches coming from unconfirmed sources
  - Some kernel developers are not convinced it's needed
  - Under development – not deployed yet
- See https://lwn.net/Articles/813646/

Build systems

- Yocto Project

Yocto Project

- Yocto Project 3.0 release
  - Oct 23, 2019
  - Lots of work on build caching
  - Have new “build change equivalence” feature
    - Have hashes of source and binaries
    - Can detect already-built items, and retrieve them from cache
    - Speeds up build
    - Can share build artifacts using a server
  - Technology lends itself to reproducible builds
  - https://lwn.net/Articles/804640

Outline

Linux Kernel
Technology Areas
Conferences
Industry News
Resources

Conferences (past)

- Embedded Linux Conference 2019
  - August 21-23, San Diego, California, USA
- Linux Plumbers
  - September 9-11, Lisbon, Portugal
- ELC Europe 2019
  - October 28-30, Lyon, France
- Automated Testing Summit 2019
  - October 31, Lyon, France
- LinuxConf Australia 2020
  - January 13-17, Gold Coast

LinuxConf AU videos

- Schedule at: https://linux.conf.au/schedule/
  - Videos linked from session pages
    - On YouTube AND their own server
- I watched some:
  - Picolibc
    - Small C library for embedded by SiFIVE
  - OpenWRT for Energy monitoring
  - Everything Awesome about GPU drivers
  - Panfrost: Open Source meets Arm Mali GPUs
  - The magical fantasy land of Linux Kernel testing

LinuxConf AU impressions

- Some embedded topics
  - Whole track on RISC-V
    - Glibc, picolibc, software environments
  - Testing, energy monitoring, trusted boot, snek
- Not a ton of embedded content, but some worth viewing
- Lots of conferences provide videos now
  - This is a very welcome trend

Conferences - 2020

- Embedded Linux Conference 2020
  - June 22-24, Austin, Texas, USA
- Linux Plumbers
  - August 28, Halifax, Canada
- Open Source Summit Japan
  - September 15-16, Tokyo, Japan
- ELC Europe 2020
  - October 26-28, Dublin, Ireland

COVID-19 issues?

- Will conferences get canceled?
- No one knows what the status of COVID-19 will be this summer
- Everyone hopes that the virus will be under control in the next few weeks
- But it depends on many factors
  - China new infection rate is down this week
  - But South Korea, Italy, Iran are up
- ELC has not been canceled yet

COVID-19 infection rates

- Many of the China province graphs look like this: (this is for Hunan province)

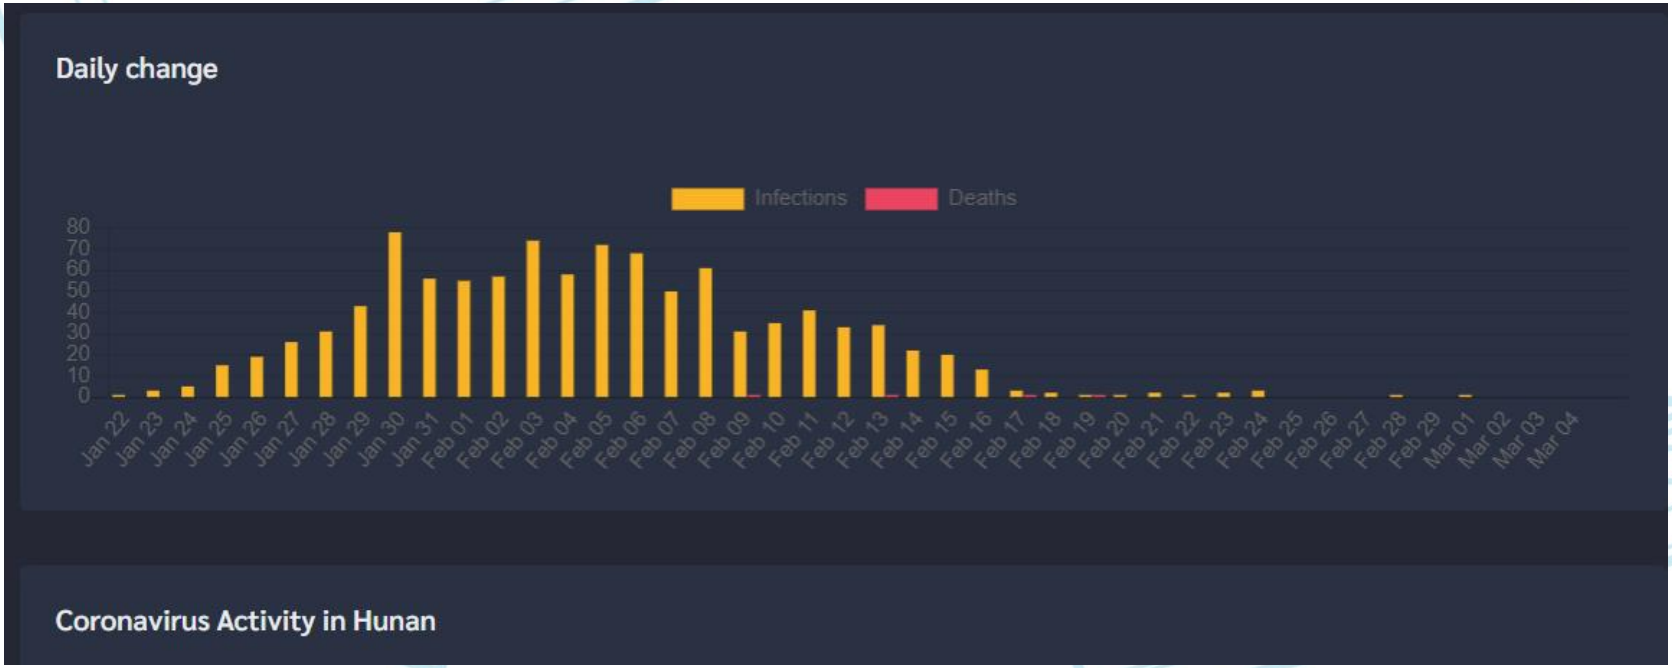

- See https://corona.help/ for good charts

ELC video salvaging

- ELC 2015 and ELCE 2015 videos were lost
  - Due to an error by Linux Foundation staff with LF's YouTube channel
  - ELC 2016 videos were recovered from original videographer
- I'm backing up all known videos
  - If you downloaded one of the missing videos, let me know
- Trevor Woerner also making a backup of ELC videos and slides
  - Also collecting videos from other events

Outline

Linux Kernel
Technology Areas
Conferences
Industry News
Resources

Industry News

- Trade associations
- Update on Gnome Foundation patent lawsuit
  - And patents (and OSS) in general
- FSF launching collaboration site
- New president of IBM
- Guido van Rossum retired
- Raspberry Pi STEM kit

Trade associations

- Linux Foundation
  - Subaru announces adoption of AGL software for infotainment in some cars
  - Uber and Microsoft announce OpenChain conformance
  - LF Core Infrastructure Initiative collaborates with Open Source Technology Improvement Fund (OSTIF)
    - OSTIF committing resources through LF's Community Bridge project
    - To enhance security audits of Open Source projects
      - LF already spent over $1M to complete 20 audits
  - See https://www.linuxfoundation.org/press-release/

Gnome Foundation sued by patent NPE (troll)

- Rothschild sued Gnome Foundation
  - Claims that “Shotwell” image management app infringes a patent
  - Patent is very generic
    - Uploading images over a wireless network using a filter criteria (e.g. subject identification)
- Gnome foundation is fighting back
  - Established a legal defense fund
- OIN has suggested that they will help
- See https://www.zdnet.com/article/leave-gnome-alone-this-patent-troll-is-asking-for-trouble/

Patent lawsuit update

- Gnome has filed a motion to dismiss
  - https://www.gnome.org/news/2019/10/gnome-files-defense-against-patent-troll/
- Gnome fundraiser raised more than needed:
  - $150,199 as of March 4
  - This is more than their goal of $125K
- Good overview of issues at:
  - https://www.eff.org/deeplinks/2019/12/how-patent-sorting-photos-got-used-sue-free-software-group

OIN update

- Microsoft joined OIN in October
  - Will now provide patents to pool to defend Linux from patent attacks
  - My how times change...
- OIN creates “United Patents Open Source Zone”
  - To defend OSS projects from patent trolls
  - see https://www.zdnet.com/article/open-invention-network-teams-up-with-ibm-linux-foundation-and-microsoft-to-protect-open-source-software-from-patent-trolls/

More on patents: DPL

- Law professors from Berkeley working on new license and pool system:
  - Defensive Patent License
- Not the same as traditional patent pools
  - Members of DPL contribute 100% of their patents to the pool
  - Patents only used for defense
  - A bunch of rules (e.g. for non-aggression between DPL members)
  - Encourage developers to use patent law like GPL uses copyright law
- See https://www.networkworld.com/article/2230668/the-defensive-patent-license-makes-patents-less-evil-for-open-source.html

New (freer) code collaboration site

- FSF is launching a new code collaboration site in 2020
  - Ostensibly to compete with github and gitlab
  - Want it to meet standards for ethics, freedom, etc.
    - No third party tracking
    - Proper license information
  - Plan to select an existing open source platform and enhance it to meet goals
  - In the planning stages now
  - See https://www.fsf.org/blogs/sysadmin/coming-soon-a-new-site-for-fully-free-collaboration

New president of IBM

- Jim Whitehurst, former CEO of RedHat will be the new president of IBM
  - Effective April 6, 2020
- Current president and CEO Virginia Rometty is retiring
- Indicates that knowledge of open source culture is extremely valuable
- See https://www.forbes.com/sites/jonobacon/2020/01/31/jim-whitehurst-becomes-president-of-ibm-why-he-gets-culture/#1a70802b6394

Guido van Rossum retired

- Guido created the Python language
- Retired from Dropbox in October
- Time marches on...
  - When will Linux people start retiring?
    - John “Maddog” Hall, early Linux pioneer is still active!

Raspberry Pi STEM kit

- New Sania Box embedded computer kit
  - Includes a Raspberry Pi 4 board and custom shield
    - Has multiple sensors, relay, pushbutton, LEDs, 7-segment display
  - Works with Raspberry Pi 4
  - Targeted at STEM education market
    - To teach basics of Python, electronics and IoT
  - Ships in April
- Was designed by 13-year old Sania Jain
  - I feel like I wasted my youth...
- See http://linuxgizmos.com/teen-launches-raspberry-pi-4-based-stem-kit/

Outline

Linux Kernel
Technology Areas
Conferences
Industry News
Resources

Resources

- LWN.net – https://lwn.net
  - If you are not subscribed, please do so
  - Some content is delayed by 2 weeks for non-subscribers (some links in this presentation)
- Linux Gizmos – https://linuxgizmos.com
- Phoronix - https://www.phoronix.com/
- Google

Thanks!