

OE Developers Meeting European Edition

Sunday October, 21, 2018
Edinburgh, Scotland, UK

Actions from OEDaM

- Metrics for layers to find out which layers need maintenance (Crofton) to improve user experience
 - **ACTION:** Generate survey
 - Jefro will work on this
 - **ACTION:** Tracker package
 - Christopher Clark will investigate
- Investigate Load Balancing openembedded.org
 - **ACTION:** Tom King + Sean will investigate
- Discuss YP compatible v2 & layer checking
 - **ACTION:** Armin will start a wiki about yocto-check-layer
 - **ACTION:** Khem will help Scott with docs review issues
- OSUOSL
 - **ACTION:** Board needs to put together an estimate to host
 - \$15K
 - **ACTION:** Jefro will help publicize the avenues available to donate to the project w/dreyna, behan
- Layer Quality (hatle)
 - **ACTION:** develop method for showing layer quality, e.g. ptests available, to be published in the layer index (Sean)
- Multiarch
 - **ACTION:** rp and koen together on multiarch with no sharp objects
 - Would like to fix multilib as well (arm/aarch64 fray)
 - If you set the lib path to /usr/koen, it won't build or run

Main Meeting:

- Board Member ballot
- OE Infrastructure
 - admins
- 2.7 Features
- Package Maintainers (reminder)
- Next meeting? Normally in US at ELC in spring. But in 2019 event is in August in San Diego. Skip? Arrange event in Spring? Where? (Maybe at SCaLE?)
- Tagged releases (Ruslan)
- Automated CVE (Common Vulnerabilities and Exposure) reporting tools (+ kernel) (Grygorii)
- Long-term OpenEmbedded support (Ruslan, Grygorii)
- Reproducible builds
- 'Patch accepted' notifications to contributors (Ruslan)
- Ptest vs oeqa runtime
- Ptests per layer

09:06 Welcome by crofton

- Bigger than last time in Prague
- An attendance sheet is being passed around
- You can join the organization
- He is looking for ideas why ppl show up and what they want from the meeting

09:22 everyone has said their names and why they are here

09:23 board member ballot

- Crofton requests more competitive elections next time

09:24 prior meeting minutes

- Jefro: Survey: no progress because of reasons
- Chris: tracker package (popularity contest)
 - The layer index now has more dependency information
 - "Pagerank" for packages, to find important packages
 - This data can be easily added to the REST API
 - May need to export additional data
 - Bill: clone count per repo would help
 - Koen: problematic with the git mirror
 - Crofton: can we sort layers by number of dependencies, find unmaintained but critical layers
 - Koen: instead of a large single unreviewed layer, we now have many layers, some of them with little review
 - Scott Murray (Konsulko) volunteered to implement automated review/checks
 - Richard: The layer now has many more data (bbappends, includes, ...) which could be used to find duplication (i.e. tune files)

- Find duplicate class/include names, find out why ppl are doing that.
- Ndec: will setup Google Analytics on layers.openembedded.org (and probably openembedded.org in general)
- Scott: we should run the existing checkers on known layers
- Load Balancing: relevant ppl are not present
- *Armin will start a wiki about yocto-check-layer*
 - No progress (yet)
 - Marco Cavallini: Are there rules to submitting layers to the index?
 - Some review, check that it's reusable for other projects (yocto compatible)
 - Crofton recommends to run yocto-check-layer (especially BSP layers)
 - Richard: the layer-index can find problems easily which are not detected by check-layer
 - Mark Hatle: Machine and distro checking works well
 - Richard: check-layer uses the sstate checksums to verify that a layer affects only things that it should touch (machine/distro level) when enabled
 - It has good signal to noise ratio
 - Mark: Adding a package config to a different package change the checksum even if it's not enabled. It's good at detecting changes that **may** affect results
 - Good tool to protect against unexpected changes brought in by a new layer
 - Mark: wants to be able to add many layers for the OS vendor use-case without causing problems (compared to the community side, where ppl. only add the layers that they actually need)
- Richard: he is trying to get layers to describe which releases they are compatible with (listed on the index)
- Mark: the index is usually very up-to-date with changes in the layer repo
- Mark: how can we handle "false positives", where the layer maintainer can demonstrate it's not a problem even if the hash changes?
 - Koen: we have the style guide and the commit guidelines(?)
 - Richard: Add that to the docs "how to make a safe layer"
 - Koen: we need to have examples in the docs. The mega manual is targeted at users, not at maintainers. Some of the documentation can cause problems for maintainers.
 - Richard: we can point Scott at these problems
- Mark: We need to have more tests including ptest, the test systems haven't ramped up as fast as we want to
 - Tim Orling: in meta-perl we have packages which haven't been tested for 5 years, we need ptest for that
 - Mark: we have the infrastructure, but not examples for simple tests. The learning curve is hard
 - Richard: we can have layers that only add tests at the image level, we have a lot of tools, but not much of that is document. He looked at the

ptests for the automated testing summit on thursday and counted test. It would be good to add these stats to the layer index

- Mark: Many of the ptests fail, so only check the delta against regressions
- Richard: will talk about testing later
- Crofton: Who would be interested in a day focused on testing? 10-20 ppl.
- Richard: the focus for OE is different than on the testing summing
- Tim Orling: Intel has moved the QA team every two years, making long term planning difficult
- Behan: We need more diversity regarding ppl. running the tests
- Ruslan (Cisco): Internal rule to have all used packages covered by ptest (~200)
 - They are compiling tests from source on the target (special case)
- Koen and rp (*multiarch*):
 - /usr/lib/<arch> instead of /usr/lib
 - There are *many* assumptions in recipes for the old layout. Lots of python code which does wrong checks
 - Debian is solving a different problem (less cross-compile)
 - Many hardcoded paths in init scripts "set LIBDIR to your first name"
 - Richard: for multilib, we solved a special case of that. For multiarch, we would need a multi-year development effort
 - Mark: A lot of time will go into upstreaming the changes into the tools
 - Richard: This year (after 15y), we may have a release with working multilib
 - Koen: The bigger distros are moving to multiarch ("universal docker container"), will make it easier for use to go there
 - Mark: redhat have ppl. paid to work on that, we can benefit
 - Mark: the kernel has support for binfmt for *many* years
 - many: OE is not just tiny devices, rather customized systems
 - Koen: Fixing multiarch bugs are beneficial even without multiarch

10:15 Infrastructure (Tom King)

- Equipment donated by HP is getting really old and needs replacement
 - Wiki and Patchwork has been moved off
 - Build infra is still there
 - core and meta-oe build test
 - ~6 machines (256 GB RAM, 12 Cores)
- Moved to Hurricane Electric (is now easier to access for him)
- 1'500 USD in the last 10?? years for infrastructure
- Would be good to close a testing loop via patchwork
- We need a campaign to get new build machines
- Mark: We need to document our requirements
- Tim: There should be a possibility to approach the LinuxFoundation because of how important OE is now

- Koen: Or approach bigger members with a concrete proposal
- Michael: For Yocto Project testing
 - 15 builders, 10Gbit bonded, 94TB NAS with SSD cache, 128/64 GB RAM
 - Full testing matrix with architectures
 - 40 different nightly targets
 - Some with virtual and real targets
 - Different package managers
 - Different kernel versions
 - Sysv vs systemd
 - Self-hosting tests
 - Building on different host distros
 - Full tests with cold sstate in 5 hours
 - Many tricks (with 2-3 hours with sstate)
 - Consolidated into 1 rack
 - Buildbot has been updated (web interface is now usable)
 - Queue and cancel
 - Funded by the Yocto Project
 - Marco Cavallini: We could also approach BSP/HW vendors
 - Tom: 20k per machine (7 years ago), this HW is owned by Tom

10:35 15 Minute Break

10:50 Richard: current status of OE

- Final stages of 2.6 (YP release)
- gcc 8, glibc 2.28, linux 4.18
- autobuilder changes: makes it easier to make config changes
 - Parallelizing the test runner?
 - <https://autobuilder.yoctoproject.org/typhoon/#/console>
 - Updates to toaster to use it as a configuration tool
- Python profile guided optimization (previously, our python was slower than from other distros)
 - Run under qemu (minor build time hit, on by default)
 - May be interesting for SDK native python
 - bitbake parsing is twice as fast
 - Not easy to backport to Sumo, many fixes were required for this in the python recipes
- Marco: can we have something simple like kbuild? Currently, one needs to know the package name, before it can be added.
 - Richard: Everyone's idea of simple and workflows are different.
- Mark: He has been working on making bitbake do some things the layer-index does offline (collect/extract information)
 - Would be a requirement for an advanced setup program
 - In the future, layer-index could use this code in bitbake via an API

- Behan: customers are building wrappers around bitbake to make it fit their environment
- Kernel dev source was reworked
- Locale compression to an archive is now possible
- We have reporting on sstate usage
 - The numbers are hard to understand
- Documentation and manual test cases have landed
 - Testopia in bugzilla is now defunct
 - We can now think about improving these test cases
- In 2.7, Intel wants to reduce their amount of testing
 - We need more automation (build appliance, toaster, eclipse plugins, CROPS)
 - Tim: The eclipse plugins need non-Intel devs to survive
 - WindRiver and IBM worked on this before, but are gone now
 - They are working on videos to show the existing features
 - He has work-in-progress repos on github for automated testing
 - David Reyna: Could test toaster with some help on the tools (selenium?)
 - We need people to lead the testing effort
 - Behan: eSDK fails in random ways
- We need machine readable test reports (which tests were run, results)
 - This is currently holding the 2.6 release (~ 1 week best case)
- In any batch of 30 patches run through the autobuilders, 1-2 fail in a corner case
 - musl builds often fail (everyone tests with glibc)
- The QA test the milestones, find 3-5 problems which were not found on the autobuilder (like error messages on the console)
 - ptest result collection is currently manual
- In 2.7 sstate hash equivalency stuff (leftover from 2.6, needs some debugging)
 - Can mark hashes as equivalent to avoid rebuilding all reverse depends
 - pseudo needs maintenance
 - Mark works on more integration between bitbake and the layer index
 - API to download and enable layers, automatically handle layer depends
 - Some parts of this are relevant to the autobuilder as well
 - Memory resident bitbake by default
 - Richard: documenting the problems would help
- Bill: Some Yocto ppl. had a meeting regarding testing and what would need to be done.
 - Running OE selftest against a set of fixed distros
- We currently have no testing of the package manager in the SDK
- Koen: He has seen that some ppl are doing tests under travis
 - Richard: Patrick Ohly has worked on this
- Mark: Someone is working on using Lava to run OE ptests on real hw
- Koen: public sstate is only for Poky, not no-distro
 - Richard: Were thinking of having some CDN for sstate
 - Denys: Continue talking about it on the OE architecture list
- Richard: security flags are enabled by default in Poky
- Richard: Releases going forward will need ppl stepping up to implement specific features

11:43 Runtime testing (ptest vs oeqa runtime)

- ptest tests a specific package, there is a wrapper which runs all installed ptests (ptest runner)
- Mark:
 - Testcase which was shipped with the source: ptest
 - Testcase for multiple components: system tests
 - DejaGNU like tests often 30 fails of 5000 passes, must be compared to previous run
- Jan Simon: IMAGE_FEATURES ptest vs installing individual packages
- Richard: Testing Hierarchy
 - Bitbake unit-tests -- bitbake selftest testing the API and datastore without metadata
 - bitbake-selftests
 - OE-selftests -- oe-selftest can execute bitbake multiple times, run devtool, recipetool, ...
 - meta/lib/oeqa
 - Runtime testing - things that run on the image
 - ptest is for tests run on the targets

11:56 'Patch accepted' notifications to contributors (Ruslan)

- Feedback for patches only when there are issues. When patches are accepted, it's not obvious.
- Richard: 50 extra emails for him to send doesn't scale. Original idea was to handle that via patchwork, but that's not really maintained.
- Ruslan: It's not easy to know which branches to look at to find patches.
- Koen: Sometimes the branch and the patch doesn't match
 - Peter: the list should be the authoritative one
- Richard: He can't just say no, because there needs to be explanation. This sometimes causes patches to get no reply if he has not time to explain.
- Richard: Thinking about moving the list to groups.io
- Patchwork is confused by the SPF/DKIM/DMARC workarounds in mailman

12:13 OEDAM 2019 location? SCaLE in Pasadena?

- Tom: Would be a Sunday, room is classroom-style, with decent networking
- Crofton: If we want to meet in spring, SCaLE would match
- Behan: The year after, LF will move OSS NA to March again
- Richard: It wouldn't need a full OE meeting at SCaLE, a set of core ppl would be enough
- Mark: The 6 month cadence is good, because it helps remind us of the relevant problems
- Tom: Will reserve the room tentatively
- Crofton: The list of 50 ppl. was full for about 1-2 months

12:30 Lunch

13:30 Richard shows some pages on the Yocto wiki

- <https://wiki.yoctoproject.org/wiki/Triage>
 - Weekly Call, run by Stephen Jolly, 30-60 minutes
 - currently 22 non-prioritized bugs
 - normally 10
 - previously > 50
 - Scope is OE-Core and bitbake
 - is a good way to keep the finger on the pulse of the project
 - it's not Yocto Project specific
- <https://autobuilder.yoctoproject.org/typhoon/#/console>
 - yocto-autobuilder2
 - <http://git.yoctoproject.org/cgi/yocto-autobuilder2>
 - yocto-autobuilder-helper
 - <http://git.yoctoproject.org/cgi/yocto-autobuilder-helper>
 - could be reused by jenkins
- <https://wiki.yoctoproject.org/wiki/BuildLog>
 - currently doesn't scale well
 - team: <https://wiki.yoctoproject.org/wiki/Swat>
 - Richard wants to make it easy to participate
 - <http://errors.yoctoproject.org/Errors/Latest/Autobuilder/>
- https://wiki.yoctoproject.org/wiki/Bug_Triage

14:02 Tagged releases (Ruslan)

- OE-Core has a 2018-04 tag
- Richard: OE didn't want to adopt the Yocto release scheme
- <https://wiki.yoctoproject.org/wiki/Releases>
- bitbake has version tags
- Bill: table in OE-Core README
- Ruslan: we only have tags for the releases, not for the stable releases
 - Richard: we should probably do that

14:10 Automated CVE (Common Vulnerabilities and Exposure) reporting tools (+ kernel) (Grygorii)

- The 'cvert' tool works better than the CVE-checker (it finds more)
 - <https://patchwork.openembedded.org/patch/153375/>
 - info from NVD database
- There were some problems with detecting the kernel CVEs correctly, but that was fixed
 - there is a kernel specific tool
 - Information available on https://github.com/nluedtke/linux_kernel_cves
- Richard: wants to avoid adding a second checker tool to the repo as it would confuse users. He is also not the right person to review the CVE patches.
 - Mark: Ross would likely be the correct technical contact

- Mark: The larger problem is that we don't actually have a process to handle CVEs. David has a talk on this on monday.
- Richard: He needs to know if there is agreement by the users to replace the existing tool with cvert.
- Tim: Ross will be able to continue working on this.
- Mark: Some companies are doing security response and submit patches upstream.
- Richard: Microsoft has been submitting CVE stuff to Sumo.
- Mark: There is a Yocto security list
 - <https://lists.yoctoproject.org/listinfo/yocto-security>
 - The current cve-checker can only handle ~10% of the CVEs in the database
- Richard: we have data in the recipes on which CVE should be fixed
 - Long standing TODO item to mine the data from code check ins

14:30 Long-term OpenEmbedded support (Ruslan, Grygorii)

- 5 year time scale, how to share this work
- Ruslan: Fixes are done only for recipes they are using, and testing only on their systems
- Richard: would like to see those branches, would like to see collaboration. They could use the same autobuilder scripts
- Crofton: who would be interested in long term support
 - ~35%
- Mark: for them to be Yocto blessed, they would need to have some level of testing and release engineering
- Richard: You'd need to run those releases on old server distro releases. May not be allowed to run on LF infrastructure
 - Crofton/Mark: needs some virtualization setup
- Josef: OE currently only provides mainline, if they need a LTS branch, they'd need to invest
- Mark: in telco, with virtualization/containers, they can change HW easier than before
- Richard: a large pain point is the server distros changing under us (glibc, rpc gen)
- Robert: need to build in a virtualized environment
- Richard: we'd need to setup a separate autobuild env for that

14:45 Reproducible builds (Ruslan)

- OE is mostly reproducible
- Richard: State is good, but hard to prove that it is reproducible
 - reproducibility for ARM was fixed
 - core-image-sato is 100% reproducible (contains gtk, x11)
 - Q: For what definition of reproducible? (darknighte)
 - A: Full, bit identical outputs
 - except the RPM-DB (because it's dynamically generated)
 - they are talking to the reproducible builds project
 - debug syms are now without paths
 - sstate is relocatable

- Marco Cavallini: how to make sure that a build is reproducible
 - Richard&Mark: build same input on the same system
 - the sstate checksums should be enough
 - the datestamp is fixed generated in a fixed way

14:50 Package Maintainers (reminder) (Armin)

- Richard: The push for having more maintainers has been successful, we need still more.
- Scott: Make a list of recipes that each maintainer would like to pass on?
 - Crofton: just ask
- Koen: Is there a document on which responsibilities and authorities a maintainer has?
 - Do we want to have more authority for the maintainer role?
 - Q: RP objects to a single maintainer being able to control a recipe/family of recipes? E.g. Khem controls gcc, etc
 - A: Everyone has a collective responsibility for *all* recipes. Just because you are a maintainer doesn't give you the ability to whatever you want.
- Richard: Ppl tend to develop areas of responsibility. Maintainers who work on a specific area can easily get patches in, but are also expected to fix the bugs they introduce.
- Sean(darknighte): RFC - how about using the OE TSC as an arbiter instead of a single person, e.g. Richard? Strong single maintainers could then be reviewed/over-ruled by TSC if their patches are causing problems?
- Denys: There were some qt5 patches were not on the mailing list, because they were merged by PR on github.
 - Richard: It should be documented on the README. Recommends against mixing PR and list.
- Richard: The layer maintainer should have the responsibility to fix critical problems (such as parsing errors)
- Richard: The OE TSC has remit over everything on git.oe.org
 - Current TSC: <https://www.openembedded.org/wiki/TSC>
 - Fray - Mark
 - Jama - Martin
 - RP - Richard
 - Khem
 - Paul
 - Richard: We didn't need a decision by the TSC for a long time. That's a sign the project is working well. It should not be an engineering organization.
 - Q: Do we need a separate engineering/arch function to help direct innovation/technical direction for the project for **OpenEmbedded**?
 - Richard: we have monthly tech calls and list, we just need more ppl involved

15:20 Coffee Break

15:40 Positioning Relative to Competitors

- Richard: We have good market share. He worries about changes in the developer community. We need more teaching of what is possible (evangelism).
- Rich Persaud: Use meta-virt to create a distributed test/build infrastructure. 50/50% internal builds/contributed builds.
 - LF is working on supply chain integrity, based on blockchain.
- Richard: OE/Yocto could help with creating containers. May have missed the window for that.
 - The sstate checksum algorithm could be changed to fit a supply chain integrity mechanism
 - sstate files can be gpg signs
- Crofton: we need more talks about OE at conferences (embedded, cloud)
- Koen&Mark: OE is the only way to have correct compliance for contains
- Richard: We need to talk about "How OE can fix license compliance for containers"
- Rich Persaud: Intel is forced to move to opensource FW, that firmware should be built with OE.

15:56 Crofton

- Why are we working on OE? Promotion? More popular?
- ????: We can collaborate on the common stuff and still customize as needed.
- ????: Customer started with buildroot, then needed features only available in OE (multiple images, BSPs in the layer)
- Behan: Layer design is the main feature, allows to work together across boundaries of responsibility
- ????: We need something like checkpatch for consistency
- Rich Persaud: Need best practices how to use layers, with explanation of why that is a best practice.
- Matt: Distro-Creep is a problem. Ppl dont know when to create a distro from an existing BSP layer.
- Tim: We still have too much tribal knowledge.
- Richard: We could change the quickstart from starting with a image build to creating their own layer.
- Robert: Has customers who are very happy with buildroot for simple projects.
 - Koen: ppl who use buildroot seriously have a lot of tooling around that
- Richard:
 - we check if the host has all required tools
 - we also enable all debug information, which slows down the build significantly
- Trevor: In the past, companies were focused on build everything from scratch. Now, they don't have enough time, so want to reuse. Board vendors should provide yocto layers.
- Mark: Sometimes board vendors explicitly build BSP layers to be incompatible with other vendors.
 - Often customers choose HW first then decide to run linux on it. Needs to be the other way round.
- Mark: SoC vendors are learning that their kernel integration is not their value add.

- Tom: License compliance is much easier
 - Mark: That mitigates the idiots.
- Mark: Would like to integrate support for SPDX information (signed info that something was reviewed by some org) into OE and store that with the output.
 - would need a way to pull external SPDX info
 - and a way to match sources to the SPDX information
 - Richard: the SPDX project has now moved to a format that can be included into a file header. Good tooling would make it more attractive to upstreams to move to these headers.
- Josef: If you want plagiarism checks use Blackduck, if you want compliance use fossology
- Richard: OE Classic git repo has been locked down
- Mark: Processing the SPDX headers takes longer than the build itself. So we need a cache.
 - There is currently no way to override that in an organization with authoritative information from the legal side.
 - Also could contain export information (regarding configuration)

16:45 Wishlist

- Koen: miniconference
 - Josef: Split between core devs and users, although there is overlap
 - Mark: Who would present at the user side? Users or devs?
 - Josef: He'd present. Layer Bootcamp is relevant to users. "show and tell" could be relevant for a user conf/track.
 - Tim: Kubernetes, Docker, gstreamer have user-focused tracks
 - Mark: Would like to hear from users.
 - Crofton: Try it at SCaLE? Instead of a meeting?
 - Tom: could work
 - Crofton: Would the user audience show up?
 - Behan: It's not a hardcore dev conference. There were 2 Yocto talks last year, there will be some this year as well. They have a 3 day Linux embedded conference/workshop before that sunday.
 - They are looking for people to teach these topics.
 - Could probably get 50 attendees
 - You'd get users to do presentations.
 - Topics ideas:
 - anti patterns
 - how to setup ptest
 - Behan: They have training material as well.