

Gnome Planet - Latest News

- [Release GCompris 26.0](#) (2026/02/03 23:00)

Today we are releasing GCompris version 26.0. We are also releasing the first official version of our companion tool for teachers: GCompris-teachers! You can find more information about it on the schools page. This new version contains 197 activities, including 2 new ones: "Drawing wheels" is an activity for drawing using a gear rotating in a cogwheel. "Multiple choice questions" is an MCQ activity. Note that this activity is hidden by default. It becomes visible after some datasets for it have been sent from GCompris-teachers. It also contains bug fixes and improvements on multiple activities. We ship translations for two more languages: Kannada and Tamil. It is fully translated in the following languages: Arabic Bulgarian Breton Catalan Catalan (Valencian) Greek Spanish Basque French Hebrew Croatian Italian Lithuanian Latvian Malayalam Dutch Polish Brazilian Portuguese Slovenian Albanian Swedish Turkish Ukrainian It is also partially translated in the following languages: Azerbaijani (87%) Belarusian (83%) Czech (98%) German (92%) UK English (96%) Esperanto (96%) Estonian (86%) Finnish (91%) Galician (97%) Hungarian (97%) Indonesian (98%) Georgian (88%) Kannada (85%) Macedonian (81%) Norwegian Nynorsk (89%) Portuguese (85%) Romanian (97%) Russian (97%) Sanskrit (97%) Slovak (78%) Swahili (88%) Tamil (84%) Chinese Traditional (85%) You can find packages of this new version for GNU/Linux, Windows, Android and Raspberry Pi on the download page. This update will also be available soon in the Android Play store, the F-Droid repository and the Windows store. Thank you all, Timothée & Johnny

- [GPG Update 2026](#) (2026/02/02 23:00)

A recent toot of mine got the response “friends don’t let friends use GPG” which, I suppose, is true enough. It certainly isn’t the attestation-friendly thing to use, and the opsec failures that are so easy with GPG-encrypted mail make it a hazard there. But for some things it’s all we’ve got, and I do like to sign Calamares releases and incidental FreeBSD things. And I am nominally the maintainer of the security/gnupg port on FreeBSD. So gpg.fail notwithstanding, here’s notes on my 2026 GPG key update. Previously in 2024 and 2025 I wrote down basically the same things: Things expire in about 13 months and I’ll have to remember then again, You can find my pubkey published on my personal and business sites, FreeBSD signature information is used rarely, but is available in the FreeBSD developers OpenPGP keys list, Codeberg will have signed commits in the Calamares repository with these keys. sec rsa4096/0x7FEA3DA6169C77D6 2016-06-11 [SC] [expires: 2027-02-03] Key fingerprint = 00AC D15E 25A7 9FEE 028B 0EE5 7FEA 3DA6 169C 77D6 uid [ultimate] Adriaan de Groot <groot@kde.org> uid [ultimate] Adriaan de Groot <adriaan@bionicbutton.org> uid [ultimate] Adriaan de Groot <adridg@freebsd.org> uid [ultimate] Adriaan de Groot <adriaan@commonscaretakers.com> ssb ed25519/0x55734316C0AE465B 2025-03-04 [S] [expires: 2026-08-26] ssb cv25519/0x064A54E8D698F287 2025-03-04 [E] [expires: 2026-08-26] ssb ed25519/0x14B6CC381BC256D6 2026-02-03 [S] [expires: 2027-02-28] ssb cv25519/0xD716006BBA771051 2026-02-03 [E] [expires: 2027-02-28]

- [Week 1 of Season of KDE 2026](#) (2026/02/02 17:36)

Hello everyone! Welcome to my first blog! I am Sayandeep Dutta, an undergraduate at SRM University. I learned about the awesome mentorship program, Season of KDE. Getting Started! started contributing to Mankala in December 2025. Got to know more about the project, interacted with mentors, and started with some small merge requests. I really like contributing to Mankala. Mankala has been a very interesting

game, and the guidance from the community is really good. Week 1: Development & Design In my first week, I had set up Mankala on my Ubuntu machine and started with the development. I created mockups for the proposed UI changes in MankalaNextGen. The mockups mainly included the main game page, a login page, a home page, and some other pages in the game, which were created and worked upon by me on Figma. Progress So Far I had created a merge request updating the MainMenu. Well, there are a couple of pages and components we need to work on, and the Main Menu is the most essential one to start with. What's next? In the upcoming week, I plan to: Complete the rest of the UI updates. Start implementing the new theme of MankalaNextGen. Thanks for reading. Stay tuned for more updates. ☺

- [Poor Man's Volumetric Light](#) (2026/02/02 10:42)

Building Volumetric Light Shafts with Real IES Profiles You know that feeling when the street lamps on a hazy evening cast these particular patterns of light? I've always been fascinated by those effects, so I decided to build a real-time volumetric lighting system using Qt 6.11's graphics capabilities, leveraging its QML shader integration and 3D rendering pipeline to recreate them with actual IES profiles. Streetlight patterns in the evening haze

- [Season of KDE 2026: Week 1 Progress for Automating Promo Data Collection](#) (2026/02/01 00:00)

Hi all! I'm CJ, and I'm participating in Season of KDE 2026 by automating portions of the data collection for the KDE promo team. This post is an update on the work I've done in the first week of SoK. My mentor gave me a light task to help me get set up and familiarize myself with the tools I'll be using for the rest of the project. The task was to automate the population of a spreadsheet that tracks follower and post counts for X (formerly known as Twitter), Mastodon, BlueSky, and Threads. The spreadsheet takes the follower and post counts of some of KDE's social media platforms and makes calculations based off that data. Important things to note: data from the sites is entered manually there are a lot of styles and formulas in the sheet Fetching Account Data Grabbing data was mostly no trouble. Mastodon and BlueSky were especially easy to work with. They have a public and well documented API that lets people collect all kinds of data in human-readable formats. One particular endpoint from both sources output account information, including follower and post counts, for a given account in neat JSON files (BlueSky, Mastodon). All it took were GET requests to these endpoints and it was smooth sailing. X and Threads proved a bit more finicky. Both of their APIs limit access to much of their functionality usually meaning webscraping methods are the most accessible for grabbing public account data. Threads shows users' follower counts out directly on an account's landing page, so processing a GET request to the URL of KDE's Threads account made it easy to grab. The problem is that there seems to be no direct way to grab the post count either through their API or with webscraping methods. For now, we've chosen to leave that be and circle back when I explore Threads more in the future. X presents a similar problem but there is an open-source frontend alternative named Nitter, instances of which lay all the stats information out in the open. The reliability of this method depends on public Nitter instances being available so it may be worth coming back around to this in a later part of the project, but for now it's a viable solution for getting follower and post counts. Inputting the Data Into the Spreadsheet With the data all fetched, all that was left is to add that data to the ODF spreadsheet. I had this down as the easy part of the task but in the end it wasn't so simple. The two major Python packages I found that can interpret and write ODF files: Pandas and pyexcel. Both of these have no problem reading data from the files, but when it comes to saving they don't preserve some elements of the spreadsheet. In the end we went the simple route which is to save the data to a separate ODF file using one of the Python-ODF interfaces and import that into the data sheet. This took a little finagling with formulas to get things working without popping errors into cells the sheet, but in the end we have an output ODF spreadsheet file containing the required data and the original spreadsheet with all the calculations pulling that data into its formulas, removing any requirement of a human interfacing with this

portion of data collection. Learned Lessons I feel like this week's task was a great first step into data collection automation. It was challenging without being too difficult to make progress on and forced me to explore different avenues for gathering data. On the confidence side, getting a (mostly) successful task out the gate helped me feel more comfortable with the tools and processes that will likely appear throughout the entirety of my SoK experience. Things will scale up from here on out though so I'm also keeping myself in check. From what I understand some of the most difficult parts of automated data collection come through having to interface with Javascript and not getting banned, both of which I've yet to come face-to-face with in any substantial capacity so far. Along with that I've face unexpected problems, such as the issue with modifying ODF files and that some websites don't play as well with certain browsers, which I don't have an easy way to test for yet. With these in mind I'm trying to tread lightly and be diligent with research and good practice as I continue on.

- [I'm Joining the Season of KDE 2026!](#) (2026/02/01 00:00)
- [This Week in Plasma: getting 6.6 ready for release](#) (2026/01/31 00:03)

Welcome to a new issue of This Week in Plasma! This week we reached that part of every Plasma release cycle where the bug fixes and polish for the upcoming release are still coming in hot and heavy, but people have also started to land their changes for the next release. So there's a bit of both here! Everyone's working really hard to make Plasma 6.6 a high-quality release. It's a great time to help out in one way or another, be it testing the 6.6 beta release and reporting bugs, or triaging those bugs as they come in, or fixing them! Notable New Features Plasma 6.7.0 The Emoji Selector window now lets you choose mixed skin tone groupings of emojis using a nice little user-friendly pop-up dialog. (Tobias Ozór, plasma-desktop MR #3426) After Before You're now able to set a global keyboard shortcut to clear the notification history. (Taras Oleksyn, KDE Bugzilla #408995) Notable UI Improvements Plasma 6.6.0 In the Application Dashboard widget, keyboard focus no longer gets stolen by selectable items that happen to be right under the pointer at the moment the widget is opened. (Christoph Wolk, KDE Bugzilla #510777) Breeze-themed checkboxes now always have an opaque background, which resolves an issue where their unchecked versions could be hard to see when overlaid on top of images. (David Redondo, KDE Bugzilla #511751) Plasma 6.7.0 System Settings' subcategory back button is now more of a traditional back button, eliminating a source of redundant page titles. There are still too many, but this reduces the count by one! (Nate Graham, systemsettings MR #390) After Before On System Settings' Notifications page, you can now always preview a notification sound even if sound is currently disabled for that notification. (Thomas Moerschell, plasma-workspace MR #6214) Plasma's network settings now expose additional L2TP VPN options that were previously unavailable. (Mickaël Thomas, plasma-nm MR #480) The old Air Plasma style (a lighter take on the original Oxygen style) is back, with fixes and improvements, too! (Filip Fila, oxygen MR #77) The Oxygen cursor theme received a small visual fix to improve the appearance of the busy cursor. (Filip Fila, oxygen MR #89) The cursor theme settings now show more accurate previews, which fixes issues like wobbling cursors, and makes the preview grid feel more stable. (Kai Uwe Broulik, plasma-workspace MR #6240) System Settings' various theme chooser pages are now consistent about whether you can delete the active theme (no), and also let you know why certain themes can't be deleted: because they were installed by the OS, not the "get new stuff" system. (Sam Crawford, plasma-workspace MR #6222) The Weather Report widget now shows a progress indicator while its popup is open but still loading the weather forecast from the server. (Bogdan Onofriichuk, kdeplasma-addons MR #993) Frameworks 6.23 Combobox pop-ups through QtQuick-based apps and System Settings pages now use the standard menu styling, rather than a custom style. (Nate Graham, qqc2-desktop-style MR #497) After Before Added styling for the new upstream SearchField component that was recently added to QtQuick itself. Now it looks like the KDE version that we created years ago before the Qt version existed. (Manuel Alcaraz Zambrano, qqc2-desktop-style MR #500) Notable Bug Fixes Plasma 6.5.6 Fixed an issue on operating

systems with asserts turned on (like KDE neon) that could sometimes crash Plasma when you launched apps. (David Edmundson, KDE Bugzilla #513312) Plasma 6.6.0 Fixed a rare issue that could leave KWin without control of the mouse and keyboard at login. (Vlad Zahorodnii, KDE Bugzilla #511611) Fixed a case where Plasma could crash when moving around a Weather Report widget on the desktop. (Bogdan Onofriichuk, KDE Bugzilla #514200) Fixed an issue in Discover when launched with Snap support that could prevent it being launched again after previously being closed. (Aleix Pol Gonzalez, KDE Bugzilla #507217) Fixed a case where the Plasma Bluetooth pairing wizard would fail to pair devices. (David Edmundson, KDE Bugzilla #495615) The logout screen no longer fails to take focus if you raised the focus stealing level to “Medium” or higher, and manually de-focusing it no longer breaks your ability to re-focus it. (Vlad Zahorodnii, KDE Bugzilla #514204 and Aleksey Rochev, KDE Bugzilla #511258) The custom size ruler for Plasma panels no longer sometimes appears on the wrong screen of a multi-screen setup. (Vlad Zahorodnii, plasma-workspace MR #6215) The “X notifications were received while Do Not Disturb was active” notification is no longer inappropriately saved to the notification history. (Kai Uwe Broulik, plasma-workspace MR #6223) Fixed two semi-related issues with widgets that made them not let you pick the same color after deleting it, or configure the same mouse action after deleting it. (Christoph Wolk, KDE Bugzilla #514983 and KDE Bugzilla #449389) The global menu on a newly-cloned panel on a different screen now realizes immediately that it’s on a new screen. (David Redondo, KDE Bugzilla #514907) Discover’s notification about ongoing updates no longer displays a nonsensically large number of updates under certain circumstances. (Harald Sitter, KDE Bugzilla #513676) Plasma Browser Integration no longer exports Microsoft Teams calls as controllable media sources; it always omitted them in the past, but Microsoft changed the URL again, so we had to adapt to that once more. (Kai Uwe Broulik, KDE Bugzilla #514870) You can now move focus from the Application Dashboard widget’s search field using arrow keys, and also type accented characters using dead keys or the compose key as the first character in a search while the search field isn’t explicitly focused. (Christoph Wolk, KDE Bugzilla #511146 and KDE Bugzilla #510871) Animated wallpapers explicitly set for the lock screen now play their animation as expected. (Taras Oleksyn, KDE Bugzilla #460910) Undoing the deletion of a panel widget no longer sometimes positions it far from where it was before. (Marco Martin, KDE Bugzilla #515107) Ending the renaming of an item on the desktop by clicking on another item no longer starts a drag-selection for no good reason. (Akseli Lahtinen, KDE Bugzilla #514954) The “Defaults” button on System Settings’ Accessibility page now works properly to reset non-default colorblindness modes. (Andrew Gigena, KDE Bugzilla #513489) Plasma 6.7.0 Fixed a very nasty KWin bug that could, under certain rare circumstances relating to intensive Alt+Tab usage, cause the screen to go black. This change may be backported to Plasma 6.6 if it’s deemed safe enough. (Xaver Hugl, KDE Bugzilla #514828) Frameworks 6.23 Fixed being unable to paste clipboard entries as text that were copied from a LibreOffice Calc spreadsheet cell and then re-arranged in the clipboard history list. (Alexey Rochev, KDE Bugzilla #513701) Notable in Performance & Technical Plasma 6.6.0 KRD (the KDE library for remote desktop support) no longer requires systemd. (David Edmundson, krd MR #141) Plasma 6.7.0 KWin now supports the ext-background-effect-v1 Wayland protocol. This adds support for standardized background effects like blur, opening the door for visual consistency across apps using these effects. (Xaver Hugl, KWin MR #4890) How You Can Help KDE has become important in the world, and your time and contributions have helped us get there. As we grow, we need your support to keep KDE sustainable. Would you like to help put together this weekly report? Introduce yourself in the Matrix room and join the team! Beyond that, you can help KDE by directly getting involved in any other projects. Donating time is actually more impactful than donating money. Each contributor makes a huge difference in KDE — you are not a number or a cog in a machine! You don’t have to be a programmer, either; many other opportunities exist. You can also help out by making a donation! This helps cover operational costs, salaries, travel expenses for contributors, and in general just keep KDE bringing Free Software to the world. To get a new Plasma feature or a bugfix

mentioned here Push a commit to the relevant merge request on invent.kde.org.

- [KWallet, SecretService, oo7: the story so far \(2026/01/30 16:15\)](#)

Continuing previous efforts to update the “secure passwords” story of the Plasma desktop, I’ve done some integration work between Plasma and oo7. Oo7 is a relatively recent SecretService provider written in Rust; it’s very nice, lightweight and cross-desktop. Very interesting for us supporting it as a first class citizen. We want in the end to replace completely the old KWallet system with something based on SecretService, have all our applications migrated transparently with user data as well, if possible with a cross-desktop backend implementation. We also want it to be as transparent as possible for the end user, not having any complicated first time setup, or dialogs that ask for a manual unlock when not needed. For the user, the whole system should be just invisible, except when looking up their passwords or when client applications add passwords to it themselves. As a first thing several months ago we did a transparent translation layer for the KWallet DBus interface, so now the “KWallet” service became just a translation layer between the kwallet api and the secretservice one. The old kwallleted service instead became “ksecretd” now providing only the SecretService interface, and that’s what we use by default now. However, it can be configured to use a different provider, so it can already be used with gnome-keyring or KeepassXC if the user wishes. In that case the user data will be migrated the first time. Also I’ve written a new application called KeepSecret. It is centered around SecretService, with a much more modern UI compared to KWalletManager and works also much better on mobile, so we can have also a modern password manager on Plasma Mobile. KeepSecret running with the oo7 backend, with a native password dialog. The last thing that we did so far is integration with a new SecretService provider on the block: oo7. If oo7 is just another SecretService provider, won’t it “just work” in Plasma? Not quite, but that’s actually a good thing. Oo7 wants to properly integrate with the environment it runs in, so we need to build that integration. When you use KWallet, KeepassXC, or gnome-keyring directly, they create their own password dialogs using their own toolkits and styles, therefore potentially looking a bit alien when run in a different environment than the one it was designed for. This is not ideal for evaluating options for a possible future replacement of KWallet (which piece by piece we are working towards), as we want the UI part completely integrated with the Plasma desktop. Instead, oo7 is built on the assumption that the environment it runs on will provide its own dialogs, a bit like the portal system for services like screen sharing and file picking. This requires active integration both on the Plasma part and the oo7 part, but on the other hand it will look native on Plasma, Gnome, and any other platform that implements the integration hooks. And that’s with the added benefit of sharing the password storage if someone every now and then goes desktop environment hopping. On the Plasma side we decided to employ a very small daemon which is launched only on demand with dbus activation. It will create the needed system dialogs, pass the user input back to oo7 and then quit when not needed anymore, so it won’t constantly use up system resources. This is already available in Plasma 6.6: the code for it is here. It’s designed first for oo7, but it can be used by any daemon which wishes to provide native password dialogs, so we might start to use it also from other components in the future, either other SecretService providers or completely different things which have similar needs like Polkit. On the oo7 part, Harald implemented the integration with its dbus interface, so with oo7 from main branch (or its next release) it can be used in place of KWallet (with all old applications still using the KWallet API using its backend as well), with all data migrated to it, and it will show native UI for its password dialogs... all the pieces of the puzzle slowly falling together. Needless to say this is still at an experimental stage, but it’s an important milestone in getting more independent from the aging KWallet infrastructure.

- [Web Review, Week 2026-05 \(2026/01/30 13:14\)](#)

Let’s go for my web review for the week 2026-05. Neocities Is Blocked by Bing Tags: tech, web, search, microsoft, vendor-lockin Huh? What’s

going on there? I don't see why they would exclude this domain completely, it makes no sense.

<https://blog.neocities.org/blog/2026/01/27/bing-block> The Enclosure feedback loop Tags: tech, ai, machine-learning, gpt, copilot, enclosure, vendor-lockin Interesting point. As we see the collapse of public forums due to the usage of AI chatbots, we're in fact witnessing a large enclosure movement. And it'll reinforce itself as the vendors are training on the chat sessions. What used to be in public will be hidden.

<https://michiel.buddingh.eu/enclosure-feedback-loop> No, Cloudflare's Matrix server isn't an earnest project Tags: tech, ai, machine-learning, copilot, matrix, cloudflare, security, failure Very in depth review of the mess of a Matrix home server vide coded at Cloudflare... all the way to the blog announcing it. Unsurprisingly this didn't go well and they had to cover their tracks several times. The response from the Matrix foundation is a bit underwhelming, it's one thing to be welcoming, it's another to turn a blind eye to such obvious failures. This doesn't reflect well on both Cloudflare and the Matrix Foundation I'm afraid. <https://nexy.blog/2026/01/28/cf-matrix-workers/> I Was Right About ATProto Key Management Tags: tech, social-media, bluesky, decentralized Indeed, it just can't be called decentralized... <https://notes.nora.codes/atproto-again/> Microsoft Gave FBI BitLocker Encryption Keys, Exposing Privacy Flaw Tags: tech, microsoft, security, privacy Are we surprised? Of course not... As soon as you backup the keys on someone else's server BitLocker can't do anything to ensure privacy.

<https://www.forbes.com/sites/thomasbrewster/2026/01/22/microsoft-gave-fbi-keys-to-unlock-bitlocker-encrypted-data/> ICE takes aim at data held by advertising and tech firms Tags: tech, advertisement, surveillance, politics, privacy What a surprise... No really who would have expected this could happen? I heard so many times "I have nothing to hide" over the years. When something like this happens you suddenly wish you were a bit more careful with your privacy and the privacy of the people around you.

https://www.theregister.com/2026/01/27/ice_data_advertising_tech_firms/ The Rise of Sanityware Tags: tech, privacy, surveillance, attention-economy Those are indeed getting more popular. In a way that's unfortunate, we shouldn't need them so much.

<https://thatshubham.com/blog/2026/Blogs-Are-Back> Tags: tech, rss, blog, tools Looks like a nice tool to help people to get into RSS.

<https://www.blogsareback.com/> Places to Telnet Tags: tech, networking, funny Telnet is not dead! We still have fun places to turn telnet clients to. <https://telnet.org/htm/places.htm> 10 Years of Wasm: A Retrospective Tags: tech, web, standard, webassembly, history Nice retelling of the story behind WebAssembly. <https://bytecodealliance.org/articles/ten-years-of-webassembly-a-retrospective> cppstat - C and C++ Compiler Support Status Tags: tech, c++, standard, tools Looks like an interesting tool to follow availability of C++ features in compilers.

<https://cppstat.dev/> Why I still teach OpenGL ES 3.0 Tags: tech, graphics, teaching, learning Good point, it is old but portable and carries the important concepts. This is a good teaching vehicle. Even though it's unlikely you'd use it in the wild much longer.

<https://eliasfarhan.ch/jekyll/update/2026/01/27/why-i-teach-opengles.html> SPAs Are a Performance Dead End Tags: tech, web, frontend, performance It's a solution for a problem long gone. SPAs should be the exception for highly interactive applications not the norm. Most web applications don't need to be a SPA and would be better off without being one. <https://www.yegor256.com/2026/01/25/spa-vs-performance.html>

Functional Core, Imperative Shell Tags: tech, functional, architecture, tests, tdd Clearly not a style which works for any and every applications. Still, it's definitely a good thing to aim towards such an architecture. It brings really nice properties in terms of testability and safety.

<https://www.destroyallsoftware.com/screencasts/catalog/functional-core-imperative-shell> How I estimate work as a staff software engineer Tags: tech, estimates, decision-making The approach is interesting. I wouldn't assume it's doable in every contexts though. What's sure is that you need to embrace the uncertainty and accept to go with the exercise. Estimates are needed to make decisions and help teams to sync.

<https://www.seangoedecke.com/how-i-estimate-work/> Is It Worth It? Tags: tech, failure, organisation, estimates Solving paper cuts pay off faster

than you'd think. <https://griffin.com/blog/is-it-worth-it> Things I've learned in my 10 years as an engineering manager Tags: tech, engineering, management, leadership Nice advice, there's a lot of variation on the role. And yet, some things seem to always be there. <https://www.jampa.dev/p/lessons-learned-after-10-years-as> Because coordination is expensive Tags: tech, team, organisation, communication, complexity The complexity and cost in organisations is indeed mostly about coordination. This is a difficult problem and largely unsolved in fact. <https://surfingcomplexity.blog/2026/01/24/because-coordination-is-expensive/> Douglas Adams on the English-American cultural divide over "heroes" Tags: culture The contrast is indeed very stark. I got my own bias and fondness for heroic failures. <https://shreevatsa.net/post/douglas-adams-cultural-divide/> Bye for now!

- [Call for Presentations - Qt World Summit 2026 in Berlin](#) (2026/01/30 11:59)

It's that time of year again; another Qt World Summit is approaching! We are looking for speakers, collaborators, and industry thought leaders to share their expertise and insights at the upcoming Qt World Summit 2026 on October 27-28, in Berlin, Germany. *Please note we are looking for live talks only.

- [December/January in KDE Itinerary](#) (2026/01/30 08:45)

In the past two months since the previous report KDE Itinerary got new vector-based map views and manual control over reservation cancellations, and there has been more work on reverse engineering proprietary train tickets, among many other things. New Features Vector maps The map views can now use a MapLibre-based vector map instead of the previous raster image map. This gives us smoother and continuous zoom, and should also allow localized labels eventually, something particularly important in regions you can't even read the local script. Vector-based map view in Itinerary. This is enabled in our nightly APK and Flatpak builds, other distribution channels might not have this enabled yet. Reservation cancellation It's now possible to manually mark reservations as canceled, rather than just importing that information from documents. This is useful for keeping unused reservations around but not have them considered for transfers, in the maps or for statistics, e.g. because you still need them for refund claims or travel cost reimbursement. Manually changeable cancellation status. Canceled reservations are now also clearly marked as such in the timeline view again, something that had gotten lost in the last timeline restyling. Infrastructure Work Ticket barcode reverse engineering Ticket barcodes can contain machine-readable information about the trip and can therefore be very useful for importing tickets into Itinerary. While ticket barcodes on international trips commonly use openly standardized formats for interoperability between different providers, this is unfortunately often not the case for domestic tickets. For those we have to resort to reverse engineering. Fortunately though, there's more people interested in what's in those barcodes, e.g. from a privacy or security perspective, and there's increasing collaboration. Specifically, the understanding of the Hungarian domestic ticket format has significantly increased recently, as the result of three different parties exchanging their observations and evaluating their respective findings against each other's ticket sample pools. A few magic numbers for product/tariff/discount codes are yet to be figured out, but there's now a pretty complete Kaitai spec able to decode versions 2 to 6 of the current format, e.g. using the Kaitai Web IDE. Many of the discoveries made there have resulted in improvements and fixes for MÁV and Volánbusz tickets in Itinerary. Documentation on ticket barcodes formats is also slowly being consolidated in the Train Ticket Wiki, where available also in appropriate machine-readable formats (Kaitai, Protobuf, ASN.1, etc). Reverse geocoding for location codes The work on automatic geocoding for train and bus reservations covered in the last report has been further expanded: Hotel and restaurant reservations are now also automatically geocoded. This makes automatic transfers work in more cases. Reverse geocoding on imported tickets only containing train station or airport codes is now also supported. This for example results in proper station names when scanning a train ticket

barcode that only contains station codes. Events There's two recent events that covered upcoming changes to Transitous, the public transport routing service used by Itinerary: 39C3 end of 2025. The Transitous Hack Weekend a few weeks ago. And more opportunities to meet the Itinerary and Transitous teams are coming up: FOSDEM on Jan 31-Feb 1 in Brussels, in particular with the Railways and Open Transport track. The next bi-annual OSM Hack Weekend in Karlsruhe on Feb 21-22. FOSSGIS-Konferenz on Mar 25-28 in Göttingen. I'll talk about our OSM indoor router there. The KDE Mega Sprint on Apr 6-11 in Graz right before Grazer Linux Tage will probably also include some Itinerary hacking. Fixes & Improvements Travel document extractor Added or improved travel document extractors for Booking.com, České dráhy, Cytric, Deutsche Bahn, gomus, KLM, MÁV and Volánbusz. Fixed a crash on referenced but non-existent PDF image masks (bug 513945). Handle RSP6 decoding failures properly, fixing a few fields containing "null" in that case. Added support for resolving Hungarian railway station codes. Added support for local SNCF pass barcodes in some regions of France. All of this has been made possible thanks to your travel document donations! Public transport data Added (partial) coverage in Japan and Thailand thanks to Transitous. Fixed access parameters to the VVS API for the Stuttgart area in Germany. All of this also directly benefits KTrip. Itinerary app Mark cancelled intermediate stops on the journey section map. Don't allow to open an empty place context menu. Show yearly sections for past trip groups in the trip group list. Fix showing entrance date picker for events without valid start date. Fix a timer overflow when monitoring for delays for trips in the slightly more distant future. Fixed several issues following from inconsistent ordering of timeline entries, such as two elements starting at exactly the same time and having no defined end time. This could result in duplicate or lost timeline entries under some circumstances. Correctly reset the trip group name field after closing the rename dialog. Fix overly aggressive input validation on manually entered flight numbers. Continuous Flatpak builds for testing the latest development version are now also available for ARM64. How you can help Feedback and travel document samples are very much welcome, as are all other forms of contributions. Feel free to join us in the KDE Itinerary Matrix channel.

- [This week in Bouncy Ball - new features land](#) (2026/01/30 08:08)

After some more nights spent with our beloved Bouncy Ball, I'm happy to announce that a new version is out on the KDE Store. Last week in Bouncy Ball | KDE Store page Prior to getting down and dirty with it, I again had a look at how the ball was implemented in KDE4..... Continue Reading →

- [Kaidan 0.15.0: Audio/Video Calls and Integrated Search Field](#) (2026/01/29 23:00)

Welcome Kaidan 0.15.0! This release adds experimental support for calls. In addition, it contains some very useful improvements and lots of fixes. Most of the work has been funded by NLnet via NGI Zero Entrust and NGI Zero Commons Fund with public money provided by the European Commission. Audio/Video Calls Kaidan has supported voice and video messages for a long time. Starting with this release, you can even have an audio or video call with a contact! An incoming call is indicated via a notification and you can either accept or reject it. Please note that there are still some features missing and some setups may not work properly. Especially, calls are only supposed to work on Linux at the moment. But we wanted to share the current achievements with you to get some feedback! Our goal is to extend the A/V calls functionality and make it available on other operating systems in the future. Notifications for Group Chat Replies Formerly, you got a notification if someone mentioned you in a group chat while the corresponding setting was enabled. But you could miss replies to your messages. Kaidan notifies you now on receiving replies as well. Message Input Field Focusing In contrast to the soft keyboard on a mobile device, which needs to be opened each time you want to enter something, your keyboard is always reachable on a desktop device. Why not make use of that circumstance? Kaidan ensures that the most relevant message input field stays focused to allow entering text without an additional click into the corresponding field. That way, you can

interact smoothly with the user interface and be more productive. Advanced Message Highlighting Kaidan 0.14 introduced highlighted messages if you opened their context menu. Messages are now also highlighted while they are being corrected or while you are choosing emojis to react to them. If another message was already highlighted before, that message is highlighted again once you sent the correction/reaction. Integrated Search Field With Kaidan, you can quickly search for chats and messages. But while searching for messages, the opened search bar reduced the space for messages. On mobile devices, the search bar even consumed unnecessary space within the chat list. Both problems are solved now! The search field is integrated into the main toolbar above the messages resp. the chat list. You can even focus each search field via an own keyboard shortcut to directly search without moving the cursor. Password Manager Fallback Since Kaidan's last version, passwords are stored in a password manager if the system provides one. But there was no fallback yet. It is now possible to use Kaidan even if no password manager is available. In that case, the passwords are stored in an unencrypted file. Once Kaidan detects a password manager on start, the unencrypted passwords are automatically migrated to the password manager. Changelog There are several other improvements. Have a look at the following changelog for more details. Features: Add support for audio/video calls (XEP-0166: Jingle, XEP-0167: Jingle RTP Sessions, XEP-0176: Jingle ICE-UDP Transport Method, XEP-0215: External Service Discovery, XEP-0320: Use of DTLS-SRTP in Jingle Sessions, XEP-0353: Jingle Message Initiation) (@melvo) Show busy indicator while saving captured image/video data (@melvo) Notify on receiving reply to own group chat message if 'On mention' notification setting is enabled (@melvo) Select file after opening in folder on Linux if supported (@melvo) Improve media capturing look/behavior (including preview after capturing image until image is saved) (@melvo) Restore focusing of last focused user interface elements (especially message input field) for various use cases (@melvo) Keep message bubble highlighted on reacting/correcting (@melvo) Allow to select message for correction via Ctrl+Up/Ctrl+Down (@melvo) Integrate search field into main toolbar increasing space for messages and, on mobile devices, even for chats in chat list (@melvo) Show message search field via Ctrl+Shift+F (@melvo) Display toolbar buttons on mobile devices exactly as on desktop devices (@melvo) Hide horizontal separator above top-most chat unless chat list is scrolled (@melvo) Store passwords in unencrypted file if no password manager is available or corresponding command-line option provided (@fazevedo) Migrate unencrypted passwords to password manager if available on start (@fazevedo) Bugfixes: Fix overlapping message bubble tail (@melvo) Fix medium preview hovering if hidden drop area info is hovered (@melvo) Fix updating OMEMO 2 keys for all use cases (@melvo) Fix deadlock on logout during upload of multiple files (@melvo) Fix creating additional database connection on wrong thread (@melvo) Fix sending/resetting whether message is being composed for various corner cases (switching chat, logging out, disabling corresponding setting) (@melvo) Fix updating last message on receiving initial message after setting up existing account in Kaidan for first time (@melvo) Fix resetting draft message after canceling message correction (@melvo) Fix resending failed message reaction (@melvo) Fix selecting previously selected message after changing reactions (@melvo) Fix restoring message highlighting and cancel ongoing correction/reply on removing corresponding message (@melvo) Fix displaying last message sender in chat list after draft message removal (@melvo) Notes: Kaidan requires Kirigami Addons 1.8 now Kaidan requires QXmpp 1.14 now Download Source code (.tar.xz) (sig signed with 04EFAD0F7A4D9724) Linux (Flatpak on Flathub) Or install Kaidan for your distribution:

- [Improving a part of KDE documentation generation as an SOK mentee \(2026/01/29 18:30\)](#)

The docs-kde-org project currently relies on an old, hard forked version of dblatex, which is bundled directly within the repository. This makes maintenance difficult. What's dblatex? From the man page: dblatex is a program that transforms your SGML/XML DocBook documents to DVI, PostScript or PDF by translating them into pure LaTeX as a first process. MathML 2.0 markups are supported, too. It started as a clone of

DB2LaTeX. My task is to swap the bundled dblatex with the one present in the upstream repos and make the PDF generation work reliably. Then switch the rendering backend from pdfLaTeX to XeLaTeX to support native unicode (to fix rendering issues for languages like Arabic, Chinese, Japanese, Korean, etc). If that doesn't work out, I'll extract the dblatex fork into its own repository to ease maintenance. Checkpoints to achieve Swap our current fork and use the latest upstream version. Also compare these two. Migrate from pdfLaTeX to XeLaTeX Verify our fixes and confirm if all the PDFs are being generated properly. I'll be sharing my progress and struggles (setting this website up was probably a struggle in itself) from time to time as I start this project. This work is part of Season of KDE, mentored by Johnny Jazeix.

- [Qt Creator 19 Beta released](#) (2026/01/29 10:29)

We are happy to announce the release of Qt Creator 19 Beta! You find a selection of improvements and fixes below. Please have a look at our change log for more detailed information.

- [Ooo, what does this disk do?](#) (2026/01/28 23:00)

I have a pile of hard drives. 3.5" Spinning rust. There's like a dozen of them, some labeled cryptically (EBN D2), some infuriatingly (1) and some not-at-all. Probably most of them work. But how to effectively figure out what is on them? FreeBSD to the rescue. Hotplug Just Works All the drives are SATA. I do have an IDE drive, I use it for opening beer bottles. And an ST-225 for old-time's sake. But SATA it is, and there's spare SATA data- and power-cables dangling out the side of my PC. This particular machine runs FreeBSD 14.3, and connecting a drive (data first, then power) yields some messages in the system log. Old-school, the command to read these is still dmesg, which prints: ada3 at ahcich0 bus 0 scbus0 target 0 lun 0 ada3: <WDC WD3200AAKS-00SBA0 12.01B01> ATA-7 SATA 2.x device ada3: Serial Number WD-WMAPZ0561055 ada3: 300.000MB/s transfers (SATA 2.x, UDMA6, PIO 8192bytes) ada3: Command Queueing enabled ada3: 305245MB (625142448 512 byte sectors) That matches information printed on the label of the disk (this drive is from 2007, has a 1 written on it in black marker – it may have been a drive in the English Breakfast Network server back then). It also tells me that the disk is registered with the system as ada3. FreeBSD's disk subsystem is a stack of "GEOM classes". The geom(4) manpage tells me that it is a modular disk I/O request transformation framework, but the important bit is geom(8), the command to query the disk subsystem. Running geom disk list ada3 tells me what is known about disks involved with ada3. This actually doesn't tell me much I don't already know: Geom name: ada3 Providers: 1. Name: ada3 Mediasize: 320072933376 (298G) Sectorsize: 512 Mode: r0w0e0 descr: WDC WD3200AAKS-00SBA0 lunid: 50014ee0aab8fe3 ident: WD-WMAPZ0561055 rotationrate: unknown fwsectors: 63 fwheads: 16 Well, actually this tells me that the dmesg and geom output are in mibi- and gibi-bytes, and that they're consistent. How about partitions on this disk, though? geom part list ada3 tells me (here I've removed several partitions, along with many other lines that are not very useful right now; the output is extensive): Geom name: ada3 scheme: GPT Providers: 1. Name: ada3p1 Mediasize: 819200 (800K) efimedia: HD(1,GPT,a43a2da3-bb5f-11e5-8d81-f5e4d894ddb1,0x22,0x640) label: (null) type: efi index: 1 2. Name: ada3p2 Mediasize: 315679277056 (294G) efimedia: HD(2,GPT,a43c9069-bb5f-11e5-8d81-f5e4d894ddb1,0x662,0x24bff9c0) label: (null) type: freebsd-ufs index: 2 So it is a GPT-partitioned disk, and at least one of the partitions is an "old-fashioned" UFS partition. That's FreeBSD before ZFS became the de-facto standard filesystem (maybe just for me). This disk is simple to deal with (further notes below)! After I'm done with the disk, I power it down first with camcontrol standby ada3. I can hear the disk stop spinning and then pull out the connectors (power first, then data). And move on to the next disk. After disconnecting, dmesg confirms that I unplugged the correct drive: ada3 at ahcich0 bus 0 scbus0 target 0 lun 0 ada3: <WDC WD3200AAKS-00SBA0 12.01B01> s/n WD-WMAPZ0561055 detached (ada3:ahcich0:0:0:0): Periph destroyed This way I can step through all of the drives in my pile and then jot down what disk does what (or, in rare cases, decide they can be zeroed out and re-used for something else).

Dealing with non-GPT disks The scheme reported by geom part list is GPT if you're sensible, but of course it is possible to bump into MBR and BSD disklabels as well. The geom subsystem abstracts all that away, and the only realy difference is the names of devices. Scheme BSD, also known as BSD disklabel, gives you ada3a and ada3d, rather than numbered partitions. It is possible to apply this to a whole disk. It is also possible to have a BSD disklabel inside an MBR partition, but that's a late-90s kind of setup. Scheme MBR gives you numbered partitions, but FreeBSD calls these slices instead of partitions and so you end up with ada3s1 instead of ada3p1. Dealing with UFS UFS partitions are pretty simple: mount them somewhere and things will be OK. I go for read-only, and I have a /mnt/tmp for all my arbitrary-disk-mounting activity. So mount -o ro /dev/ada3p2 /mnt/tmp it is (ada3p2 is the partition name found earlier). This particular disk turned out to be my main workstation drive around 2017, with a handful of still-interesting files on it. Ones I would not have missed if I hadn't looked, but it was nice to find a presentation PDF I gave to SIDN once-upon-a-time. After looking, umount /mnt/tmp to unmount and release the disk. After that, power-down and disconnect as described above. Dealing with ZFS Slightly newer disks might be ZFS. Here's the partition information for one: Providers: 1. Name: ada3p1 Mediasize: 250059309056 (233G) label: backup0 type: freebsd-zfs index: 1 It's a GPT partition with a label, and ZFS on it. That means it's part of a ZFS pool, and is probably referred to within the pool by its label. Possibly by its ID. In any case, ZFS needs to be imported, not mounted, because the filesystems are contained as part of the storage pool. The command to discover what is available is zpool import , which doesn't actually import anything. pool: zbackup config: zbackup ONLINE gpt/backup0 ONLINE That's encouraging: the pool consists of a single drive and everything is online. To make the filesystems in this pool available, I need to import it. I'll force it (-f, because it was probably untimely ripp'd from whatever machine it was in previously), without mounting any of the filesystems in it (-N), readonly (-o readonly=on) with a temporary name (-t zwhat), like so: zpool import -f -N -o readonly=on -t zbackup zwhat After importing the pool, zfs list tells me what filesystems are available: zwhat 36.2G 189G 25K /zwhat zwhat/zwhat/home 36.2G 189G 36.0G /tmp/foo/ I can mount a single ZFS filesystem just like a UFS filesystem. The ZFS filesystem knows where it would want to be mounted (/tmp/foo, I have no idea what I was doing back then), but we can treat it like a legacy filesystem: mount -t zfs zwhat/home /mnt/tmp After looking, umount /mnt/tmp to unmount and release the disk. After that, power-down and disconnect as described above. Dealing with Linux ext4 If the disk comes from a Linux machine, then it may have an ext4 filesystem on it. I still usually pick that when installing Linuxes. Here's the partition information for one: scheme: MBR Providers: 2. Name: ada3s2 Mediasize: 64428703744 (60G) rawtype: 131 type: linux-data There is ext4 support in the FreeBSD kernel, although it is named ext2 (and some ext4 filesystems use unsupported features, and then it won't mount). But for simple cases: mount -o ro -t ext2fs /dev/ada3s2 /mnt/tmp does the job. Dealing with linux-raid I found two disks, both WD Caviar Blue, labeled EBN D1 and EBN D2 with similar layouts. Those are linux-raid disks, and this is something I can't deal with in FreeBSD. Heck, I'm not confident I can deal with them under Linux anymore, either. Providers: 1. Name: ada3s1 Mediasize: 493928031744 (460G) efimedia: HD(1,MBR,0xa8a8a8a8,0x3f,0x398033d3) rawtype: 253 type: linux-raid Re-purposing Disks KDE has a lovely partitioning tool, but I wouldn't be me if I didn't go for the command-line approach. Make sure the disk isn't mounted anywhere, but is powered up. Zero out the first gigabyte or so of the disk: dd if=/dev/zero of=/dev/ada3 bs=1M count=1024 I guess this isn't strictly necessary, and geom warns during this operation that the GPT is corrupt and the backup GPT (at the other end of the disk) should probably be used. Ignore that. Destroy the partition table some more: geom part destroy -F ada3 , now it is really dead. Make a new GPT partition table on the disk: geom part create -s gpt ada3 Start adding partitions to the partition table. I (now) use labels with the last digits of the drive's serial-number. This drive gets a gigabyte of swap (just in case) and the rest is a ZFS partition which I can add to a pool later. geom part add -t freebsd-swap -s 1G -l swap-159666 ada3 geom part add -t freebsd-zfs -l zfs-159666 ada3 Why the labels-with-serial-numbers? Well, that's so that I can

subsequently create a ZFS pool from labeled partitions, and it remains obvious where the parts of the pool come from and also prevents name-collisions from naming everything backup0 and so.

- [Krita 5.2.15 bugfix release! \(2026/01/28 00:00\)](#)

Today we're releasing Krita 5.2.15. This is a bug fix release with a number of crash fixes and workarounds to improve use with the Xiaomi Pad. Changelog Fix crash when using a smudge brush (Bug 512243) Update GMic to 3.6.4.1 in krita/5.2 branch Fix crash when undoing liquefy (Bug 498696) Fix crash when an embedded profile is faulty (Bug 509875) Fix loading of TIFF files with JPEG compression Force SVG gradients to work in premultiplied alpha mode, counter to specification (Bug 502118, more info in commit) Fix Transform and Move shortcuts conflicting Timeline arrow key actions (Bug 513855) Don't compress touch events matched to shortcuts Treat 3+finger-touch cancels as ends (Bug 510993) Prevent rotation notice jitter around zero Allow combining zoom and rotation popup messages Disable context menu on overview widget (Bug 514238) Make touch scrolling not trigger menu (Bug 513413) Xiaomi stylus button and curve fixes (handle pageup and down) (more info in commit) Steeper tablet curve for Xiaomi devices (more info in commit) Don't flicker window on Xiaomi devices make xcodebuild an archival build to avoid unwanted debug entitlements. Download Windows If you're using the portable zip files, just open the zip file in Explorer and drag the folder somewhere convenient, then double-click on the Krita icon in the folder. This will not impact an installed version of Krita, though it will share your settings and custom resources with your regular installed version of Krita. For reporting crashes, also get the debug symbols folder. [!NOTE] We are no longer making 32-bit Windows builds. 64 bits Windows Installer: krita-5.2.15-setup.exe Portable 64 bits Windows: krita-5.2.15.zip Debug symbols. (Unpack in the Krita installation folder) Linux Note: starting with 5.2.11, the minimum supported version of Ubuntu is 22.04. [!WARNING] Starting with 5.2.11 has updated the AppImage runtime, which is known to be incompatible with the old versions of AppImageLauncher. Developers of the AppImage runtime suggest to remove or update AppImageLauncher. See this report: Issue 121 More AppImage troubleshooting info is available here: FUSE 64 bits Linux: krita-5.2.15-x86_64.AppImage MacOS Note: We're not supporting MacOS 10.13 anymore, 10.14 is the minimum supported version. MacOS disk image: krita-5.2.15.dmg Android We consider Krita on ChromeOS as ready for production. Krita on Android is still beta. Krita is not available for Android phones, only for tablets, because the user interface requires a large screen. 64 bits Intel CPU APK 64 bits Arm CPU APK 32 bits Arm CPU APK Source code krita-5.2.15.tar.gz krita-5.2.15.tar.xz md5sum For all downloads, visit <https://download.kde.org/stable/krita/5.2.15/> and click on "Details" to get the hashes. Key The Linux AppImage and the source .tar.gz and .tar.xz tarballs are signed. You can retrieve the public key here. The signatures are here (filenames ending in .sig).

- [GSOC: differences of expectations between students and organizations \(2026/01/28 00:00\)](#)

I've been administering KDE's participation in the Google Summer of Code program for the last few years (and mentoring on some). This post is just some personal thoughts on the differences between what the KDE organization expects and what usually the applicants want (I'm not in everybody heads, it's assumptions from my experience). I don't provide any solution (because I don't have any) and there is no judgment (both point of views are valid), just a personal point of view. Applicants point of view GSOC is a individual competitive program which can be a huge boost to start a career. There are usually two major reasons for the applicants to participate in GSOC: the money. the experience and the line in the CV for real life work. There is, for a lot of applicants, no genuine interest in the projects (at least at the beginning, the first interaction is often "I want to start contributing to KDE and prepare for GSOC"), and they mostly want a GSOC slot. Small digression, on one of my projects, two people already asked to contribute and if we were doing GSOC, but after telling them that contributions are welcomed and we will only propose a topic if we feel enough confidence that the contributor would stay after, they told us they preferred to choose another project. Contributors

consider GSoC as an end of their studies, not a project they want to contribute after: there is also a change of life to consider between university life and work life and a balance to achieve when starting the latter (which could be in a different town/country) which may be a reason for the drop of interest. Organization point of view GSoC is an opportunity to welcome new contributors to our community. The organization usually expects from GSoC: contributors willing to contribute regularly to their projects. implementation (partial or total) of the proposed ideas. having a positive return on investment (is the time spent mentoring worth the result?) GSoC is mostly a step to have contributors learning about the projects, and contribute in the long-term. Long-term is purposely vague; what we expected for the GCompris project was the contributor would contribute at least one year and mentor for the next GSoC (so a loop was present, and applicants would also gain an experience on mentoring, while giving some relief to the other experienced mentors) but it can vary depending on the project. For the return on investment, mentors do not consider it good having spent more than three months mentoring for projects they could have done in less time they spent to mentor the contributor. Helping someone grow is always personally rewarding but we still hope this person will also show us the results of the teaching by improving our projects in the future. Impacts of the differences There is a huge difference on the expectations. What can we do to reach a common ground where everyone is happy? Organizations cannot force anyone to contribute after the GSoC end. If money is the main motivation to stay, most organizations will not have money to hire applicants to work after (and do we want to hire someone only motivated by money? No, it's not OSS values, we love passion). Find ways to create a real interest for the organization. In KDE, we have the chance to have tons of projects with different topics (astronomy, education, digital painting, video editing, scientific plotting, games, desktop environment, system administration, internationalization, ...) where it is easy to participate to other projects from time to time and learn new things if we are curious enough. Be stricter on the entry point for organizations: explicitly say in the beginning that we expect a long-term relationship not just the equivalent of an internship. It should reduce the number of applicants and only keep the ones with a genuine interest (if applicants are honest of course). Organizations/Mentors could reduce their objectives of GSoC and consider that contributors are here to produce and spend less time on training/mentoring, expecting contributors already know the basics, but this would totally spoil the nature of the program. Another track is to find the projects where contributors stay / don't stay and understand the difference. Most of the big KDE applications don't manage to keep their contributors, while smaller ones do. Maybe as a contributor, it's because it's more difficult to take decisions, feel included/listened on large projects because they have a high maturity level/long-term contributors and a well-defined vision? Whereas on smaller projects, there are fewer constraints from the existing environment and it feels more rewarding and motivating to contribute there? To conclude, here are the statistics for KDE retention the last years: Year Started Completed Active after 1 year Active after 2 years Active after 3 years 2018 20 17 5 5 4 2019 24 22 9 8 8 2020 21 19 2 2 1 2021 16 15 8 6 3 2022 7 6 5 1 0 2023 9 7 3 3 2 2024 10 10 4 2 - 2025 15 12 6 - - That is, an average of 15 people sign up each year, of which an average of 90% finish 44% continued contributing 1 year later 28% continued contributing 2 years later 22% continued contributing 3 years later We would of course love having more contributors staying at least one year but it's almost half and a quarter that stay active for a few years!

- [KDE Plasma 6.6 Beta Release](#) (2026/01/27 00:00)

This is second beta of Plasma 6.6, which includes several bugfixes before final release on 17th February 2026. Here are the new modules available in the Plasma 6.6 beta: plasma-login-manager plasma-keyboard plasma-setup Some important features and changes included in 6.6 beta are highlighted on KDE community wiki page. View full changelog

- [Qt Creator 18.0.2 released](#) (2026/01/26 11:01)

We are happy to announce the release of Qt Creator 18.0.2! This release fixes a range of smaller issues like the persistence of the "Always save files before build" option and an issue that occurred when using some custom toolchains with vcpkg. It also updates the wizard templates for the Qt Safe Renderer.

- [Bouncy Ball will always bounce back](#) (2026/01/25 09:32)

Bouncy Ball. A proper KDE 4 classic. Many hymns of praise have been sung about this lovable desktop extension. Do You Remember the KDE Bouncy Ball Widget? Fear not, OMG! Ubuntu! You will bounce again! The KDE Ball Widget Bounces Back to Life It's a widget that lets you play with a bouncy ball on..... Continue Reading →

- [This Week in Plasma: fixing all the things](#) (2026/01/24 00:03)

Welcome to a new issue of This Week in Plasma! This week the Plasma team focused almost entirely on bug fixing. And let's let the results speak for themselves: we fixed 18 high and very high priority Plasma bugs, or 28% of all open ones! Lots of polishing for Plasma 6.6 to make it a great release. Notable New Features Plasma 6.7.0 Added a dedicated setup UI for configuring shared printers on Windows networks. (Mike Noe, KDE Bugzilla #406211) Notable UI Improvements Plasma 6.6.0 Desktop switching and Present Windows shortcuts now use the Meta key by default for more consistent system-wide behavior. (Antti Savolainen, kwin MR #8597) Plasma can now report printers' waste receptacle levels and notify users when they fill up. (Mike Noe, KDE Bugzilla #514525) KRunner's buttons have been reorganized to be consistent with other Plasma widgets, making the interface feel more familiar and coherent. (Taras Oleksyn, plasma-workspace MR #6203) Plasma 6.7.0 System Settings' Wi-Fi & Networking page now uses clearer Wi-Fi security labels, correctly showing WPA2 and WPA3 support for both Personal and Enterprise networks. (Lynne Megido, KDE Bugzilla #493238) There are now keyboard shortcuts for switching virtual desktops and opening the Present Windows effect that use the Meta key, to be consistent with other globally-scoped keyboard shortcuts. (Antti Savolainen, KDE Bugzilla #508187) Frameworks 6.23 Improved the visual fidelity of thumbnail images in open/save dialogs throughout Plasma and KDE apps. (Méven Car, KDE Bugzilla #489298) Notable Bug Fixes Plasma 6.5.6 Fixed an issue that could sometimes make KWin crash after periods of idleness. (Vlad Zahorodnii, KDE Bugzilla #513687) Fixed an issue that would make Plasma crash when you disabled widgets in the System Tray and clicked the dialog window's "OK" button rather than the "Apply" button. (David Edmundson, KDE Bugzilla #478625) Fixed an issue that could sometimes make KWin crash after repeatedly pressing the "Activate window demanding attention" shortcut (Meta+Ctrl+A by default) while multiple windows were demanding attention. (Vlad Zahorodnii, KDE Bugzilla #500748) Fixed a common case where Plasma could crash after certain games crashed first. (David Edmundson, KDE Bugzilla #506562) Fixed a common case where Plasma could crash when configured with a weather station from the Environment Canada source in its Weather Report widget. (Bohdan Onofriichuk, KDE Bugzilla #514553) Fixed a case where changing the visibility of the Media Player widget in the System Tray while music was playing could make Plasma crash. (David Edmundson, KDE Bugzilla #514823) Spectacle once again remembers the location where you last saved a screenshot the next time you save one. (Noah Davis, KDE Bugzilla #511649) Fixed an issue causing 24" 16:9 aspect ratio monitors to get the wrong default resolution. (Anton Golubev, kwin MR #8681) Plasma 6.6.0 Fixed a surprisingly common issue whereby KWin could sometimes crash when you frantically wiggled the pointer to try to stop a monitor from going to sleep. (Vlad Zahorodnii, KDE Bugzilla #487660) Fixed a case where KWin could crash when you deleted a virtual desktop that still had windows on it. (Vlad Zahorodnii, kwin MR #8680) Fixed another KWin crash, this one more random. (Xaver Hugl, kwin MR #8677) Fixed a long-standing issue whereby tooltips opened by buttons in Plasma widget popups could move onto the panel and get stuck there after you closed the widget popups. (Marco Martin, KDE Bugzilla #475646) Fixed an issue that made popups of panel widgets undesirably change their size when

you moved their panel to an adjacent screen edge. (Christoph Wolk, KDE Bugzilla #512273) Fixed an issue making certain sub-menus of Plasma widgets not have transparent backgrounds, which was especially visible with menu blurring turned on. (Marco Martin, KDE Bugzilla #513307) Fixed an issue in the HDR calibrator tool that made long pieces of text overflow from their boxes. (Nate Graham and David Edmundson, KDE Bugzilla #514687) Fixed an issue that made Plasma forget the IPSec certificate passwords of L2TP VPNs. (Mickaël Thomas, plasma-nm MR #460) Fixed an issue causing apps launched using D-Bus activation to be omitted from System Monitor's Applications table. (Arjen Hiemstra, KDE Bugzilla #510235) Fixed an issue in System Monitor that could make the Applications table's "Details" panel un-scrollable under certain circumstances. (Arjen Hiemstra, KDE Bugzilla #506150) Fixed an issue in Discover that could sometimes make Flatpak apps' languages packages fail to get grouped with the apps. (Harald Sitter, KDE Bugzilla #513111) Plasma 6.7.0 Fixed an issue that broke KRunner's Activities plugin from actually finding any activities. (Sam Morris, KDE Bugzilla #514000) Fixed an issue that caused long boot menu entries to be cut off in the Breeze GRUB Menu styling. (Sébastien Bouchard, KDE Bugzilla #513107) Frameworks 6.23 Fixed an issue that caused a large variety of crashes in Plasma and KDE apps related to devices appearing and disappearing. (Nicolas Fella, solid MR #232) Fixed an issue making KWallet crash on OpenSUSE-based operating systems. (Nicolas Fella, KDE Bugzilla #490788) Fixed an issue in Kirigami-based System Settings pages and apps that made it impossible to return to the first page when using a right-to-left language like Arabic or Hebrew. (Youssef Al-Bor3y, KDE Bugzilla #511295) Notable in Performance & Technical Plasma 6.6.0 Plasma's system monitoring infrastructure received further fixes to improve OpenBSD support. (Rafael Sadowski, libksysguard MR #454) PackageKit 1.3.4 Implemented support for DNF5 in PackageKit, which fixes a huge number of issues relevant to people using Discover on Fedora-based operating systems. (Neal Gompa, packagekit PR #931) How You Can Help "This Week in Plasma" still needs your help! Publishing these posts is time-consuming and needs community assistance to be sustainable. Right now there are two ways to help: Help put together the posts using the current mostly manual process Help automate the process Work can be coordinated in the relevant Matrix room. Beyond that, you can help KDE by directly getting involved in any other projects. Donating time is actually more impactful than donating money. Each contributor makes a huge difference in KDE — you are not a number or a cog in a machine! You don't have to be a programmer, either; many other opportunities exist. You can also help out by making a donation! This helps cover operational costs, salaries, travel expenses for contributors, and in general just keep KDE bringing Free Software to the world. To get a new Plasma feature or a bugfix mentioned here Push a commit to the relevant merge request on invent.kde.org.

- [In praise of memory-leak-detection](#) (2026/01/23 23:00)

Nicolas Fella wrote a bit about enabling memory leak detection in KDE CI, and I thought I'd add some comments from a different software engineering environment (e.g. my work-work). Guidelines for Memory Management At work-work, we have a slightly peculiar software environment: Bleeding edge C++ use - C++23 if possible, 26 features being experimented-with - and a strong reliance on C++ language and library features. Use C++ containers, `std::string` (with an `*implicit` "here be UTF-8-encoded strings"), `std::jthread`, business-logic in standard C++ wherever possible. Qt 5.6.3, from 2017, for UI code. That's pre-C++17. This leads to two guidelines about memory-management in our application: Use Qt's memory model and ownership only for QWidget's inheritance tree, .. well, ok, incidentally for things inheriting from QObject. Use C++ resource management for all other memory-management. This does mean that depending on where in the codebase you're working, you either have to be a "C++ purist" (as Nico puts it) or a "Qt pragmatist" (as I'll phrase it). Note that the guideline is not "use smart pointers", because that last guideline can be deconstructed a little (it wouldn't surprise me if this was inspired by Klaus Iglberger): If a class manages a resource, then that resource-management should be the only thing it does, If an object needs resources, then those resources should be

members (sub-objects) of a class as above. Often that means that memory-management is in the hands of `std::unique_ptr`, but `std::vector` does the job also. We have a handful of other classes for resource-management as well, like a wrapper for `FILE *` for those places where we need to interact with C's stdio library. Granted, `FILE *` is probably not a source of memory-leaks, but it is a managed resource, and forgetting one will lead to leaking file-descriptors in the long run. Leaky Libraries Dealing with memory leaks can be difficult particularly when they happen in libraries not under your control. As an example, I wrote this program (headers elided for brevity) which is a Qt5 program: `int main(int argc, char **argv) { QApplication a(argc, argv); QTimer::singleShot(5000, qApp, &QApplication::quit); return a.exec(); }` This is, like, the simplest Qt5 program. It sits there for five seconds and then quits. It also leaks 1662 bytes of memory, in 19 allocations, according to my ASAN output. They leak from allocations done in `libQt5Core`. Adding a `QLabel` and displaying some text adds leaks in `libfontconfig`, and the leaks keep accumulating as more libraries enter the mix. Granted, none of these leaks are large. I'm not even sure that they are impolite, but they're there and it makes tracking down leaks in my code harder, because there's some baseline of leakage that I need to ignore. Filtering out the noise becomes a serious undertaking. Takeaway Chasing all the memory-leaks is hard. Chasing any leaks in KDE Frameworks is something I can only applaud.

- [Detecting Memory Leaks in KDE CI](#) (2026/01/23 21:00)

Leaking memory is impolite. It's messy, it can suggest logic bugs, and thanks to AI grifters RAM is expensive. Unfortunately C++ makes it rather easy to leak memory. Fortunately we have tools to find such leaks. One such tool is Leak Sanitizer (LSAN) from the Address Sanitizer (ASAN) family. It's using compiler-based instrumentation for the code to reports any leaks after the program terminates. KDE's CI infrastructure has ASAN enabled for a while. However the leak sanitizer part was explicitly disabled, so no leaks were reported as part of the CI build and test run. This is because a number of projects have pre-existing memory leaks that would cause the C build to fail. Of course those should be fixed eventually, but in order to do that we need to know where they are. Also, for projects that currently don't have any leaks we want to enforce keeping it that way. A few lines of Python later KDE's CI system now allows to enable LSAN on a per-project basis. The `.kde-ci.yml` file that governs CI builds in KDE gained a new option `enable-lsan`. It is off by default for now for compatibility, but we may consider enabling it by default eventually. If enabled and any memory leak in a test is detected, that test and therefore the whole CI pipeline will fail. Now once we found a leak, what can we do about it? There's several options: Use good old `delete/free()`. Works, but is often rather error-prone. Use `QObject`'s parent mechanism. Works, but often frowned upon by C++ purists. Mostly makes sense for widget hierarchies. Smart pointers like `std::unique_ptr`/`std::shared_ptr`. These are great because they allow to express the ownership on an API level and make ownership transfers explicit. Often the best choice for business logic code. LSAN is now enabled for some Frameworks CI builds, but ideally it would be enabled for all KDE projects. And of course any leaks found along the way should be fixed. Happy leak-fixing!

- [Web Review, Week 2026-04](#) (2026/01/23 10:38)

Let's go for my web review for the week 2026-04. FOSS for digital sovereignty in the EU Tags: tech, europe, politics, foss Need inspiration for your answer to the European Commission call for evidence on open source? This is a good one.

<https://www.more-magic.net/posts/open-source-in-the-eu.html> Writing First, Tooling Second Tags: tech, web, self-hosting, blog, writing Get out and write indeed. You can fiddle with the tools later. <https://susam.net/writing-first-tooling-second.html> Attention Media ≠ Social Media Tags: tech, social-media, attention-economy, fediverse There was indeed another path for social media... Let's hope the Fediverse stay on its current course. <https://susam.net/attention-media-is-not-social-media.html> Why people believe misinformation even when they're told the facts Tags: tech, social-media, fake-news, attention-economy Or why the focus on fact checking is doomed to fail. You can't ignore our biases, the social

context, and above all the toxic architecture of the big social medias.

<https://theconversation.com/why-people-believe-misinformation-even-when-theyre-told-the-facts-271236> Deep reading can boost your critical thinking and help you resist misinformation – here's how to build the skill Tags: tech, reading, social-media Not all reading is born equal. The intent matters quite a lot. Build the skill, it'll last a life time.

<https://theconversation.com/deep-reading-can-boost-your-critical-thinking-and-help-you-resist-misinformation-heres-how-to-build-the-skill-268082> Giving University Exams in the Age of Chatbots Tags: tech, university, teaching, ai, machine-learning, gpt, ethics, foss Interesting ideas on how to approach teaching at the university. It gives a few clue on how to deal with chatbots during exams, can be improved but definitely a good start.

<https://ploum.net/2026-01-19-exam-with-chatbots.html> Why We've Tried to Replace Developers Every Decade Since 1969 Tags: tech, programming, business, leadership, complexity, history Good historical perspective about the attempts to get rid of developers. This never unfold as envisioned. This is mostly about the intellectual work to build artifacts handling the world complexity, and this doesn't go away.

<https://www.caimito.net/en/blog/2025/12/07/the-recurring-dream-of-replacing-developers.html> Malware Peddlers Are Now Hijacking Snap Publisher Domains Tags: tech, packaging, supply-chain, security New packaging ecosystems bring their new attack vectors. This is definitely a teething problem which will need to be addressed soon.

<https://blog.popey.com/2026/01/malware-purveyors-taking-over-published-snap-email-domains/> An adversarial coding test Tags: tech, security, ide Are you sure you want to trust that random project you got provided with? Really?

<https://runjak.codes/posts/2026-01-21-adversarial-coding-test/> Pocket TTS: A high quality TTS that gives your CPU a voice Tags: tech, ai, machine-learning, speech Sounds like a very interesting model (pun intended). It's really nice to pack that much performance in a smaller neural network.

<https://kyutai.org/blog/2026-01-13-pocket-tts> Personal infrastructure setup 2026 Tags: tech, self-hosting, infrastructure Nice ideas for setting up your own infrastructure at home.

<https://linderud.dev/blog/personal-infrastructure-setup-2026/> Unconventional PostgreSQL Optimizations Tags: tech, databases, postgresql, optimization, performance When everything obvious fails... there are still optimisation tricks available for your databases.

<https://hakibenita.com/postgresql-unconventional-optimizations> Why DuckDB is my first choice for data processing Tags: tech, data-science, databases, pandas, duckdb I definitely would like to have some time to fiddle with DuckDB more. It looks like a really neat alternative to something like pandas.

https://www.robinlinacre.com/recommend_duckdb/ Understanding C++ Ownership System Tags: tech, c++, memory, type-systems Nice introduction of the C++ ownership system. Nothing new under the sun obviously but since I still encounter developers struggling with this, such introductory material is nice to have handy for sharing.

<https://blog.aiono.dev/posts/understanding-c++-ownership-system.html> Rust's Culture of Semantic Precision Tags: tech, rust, culture This is indeed an important cultural trait in the Rust community. This can bring challenge when integrating Rust code into a context with more ambiguity.

<https://www.alilleybrinker.com/mini/rusts-culture-of-semantic-precision/> LikeC4 Tags: tech, architecture, c4, tools Still young and pretty much a one man show. This could turn into a nice tool to use C4 more productively.

<https://likec4.dev/#features> Pushing the smallest possible change to production Tags: tech, team, organisation This is good advice. Going for something extremely small first is a good way to on board in a new project.

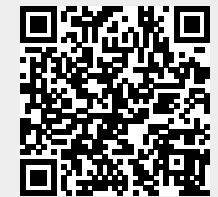
<https://ankursethi.com/blog/smallest-possible-change/> A year in the life of a Staff Engineer Tags: tech, tech-lead, leadership, management, product-management, engineering This is a good overview of what the Staff Engineer can be. There's of course a lot of variation depending on time, priorities and the culture of the organisation.

https://medium.com/@_davidanderson/a-year-in-the-life-of-a-staff-engineer-84ae9f6963c1 Work friends are the secret to great culture Tags: tech,

work, team, culture, leadership Feels a bit odd to go to such length to put it in numbers. And yet, it's clear that friendships in the workplace are a must. They should be fostered rather than stifled. [https://www.makeworkbetter.info/p/work-friends-are-the-secret-to-great The Most Powerful Pirate In History](https://www.makeworkbetter.info/p/work-friends-are-the-secret-to-great-The-Most-Powerful-Pirate-In-History) Tags: history Fascinating story. Some people shouldn't be forgotten. <https://www.youtube.com/watch?v=y2IkVIB96y4> Bye for now!

From:

<https://wiki.tromjaro.alexio.tf/> - **TROMjaro wiki**



Permanent link:

<https://wiki.tromjaro.alexio.tf/doku.php?id=news:planet:kde&rev=1583653345>

Last update: **2021/10/30 11:38**